

FREE DVD

NetBSD 7.1.1

fedora 27 Workstation 64-bit

Double-Sided DVD INSIDE!

FEDORA 27 We test drive the newest Fedora

TERMINAL TUNING

LINUX MAGAZINE

LINUX MAGAZINE PRO

MARCH 2018

TERMINAL TUNING

Tips and tools for a better Bash shell

What's New in Fedora 27

5 To-Do Managers Stay productive with a free task-management tool

YunoHost Personal server for a private cloud

Eelo This new phone puts your privacy first

Rasp Pi Tricks Spice up the light around your TV with Lightpack

BackBox Specialty Linux for security audits



LINUXVOICE

- Apt vs apt-get - Do you know about the new Debian package tool?
Get better colors on your monitor with DisplayCAL
maddog - Meltdown and Spectre



FOSS Picks

- Jotting notes with Joplin
Pioneer space adventure

Tutorial

- USB over the network with USB/IP

Issue 208 Mar 2018 US\$ 15.99 CAN\$ 17.99



REMEMBERING THE RULE

Dear Reader,

The recent action by the US Federal Communications Commission to reverse previous policy and rescind net neutrality received lots of attention. Of course the geeks didn't like it, but a sizable majority of ordinary citizens didn't like it much either. In fact, many wondered who *did* like it, aside from broadband service providers and those who believe the way to help America succeed is to yield to the wishes of broadband service providers. If you follow the high tech news, you know that much has already been written on this decision – including in this magazine, so I won't dwell on it for long.

What I really want to mention is a report from researchers at the Berkman Klein Center for Internet and Society, at Harvard University. The authors studied 40 community-owned ISPs and compared them to similar plans offered by commercial providers. The result was that municipal broadband providers typically charged less and offered more transparent pricing, with fewer gimmicks.

This might seem obvious – why wouldn't governments, who just want to break even, be able to provide Internet service for less than businesses, who are trying to make a profit? Plausible as this may seem, the USA runs on a deeply held conviction that businesses do things more efficiently than governments.

Unfortunately, as illuminating as this study on municipal broadband might seem, it does not totally resolve the business vs. government (capitalism vs. socialism) ideological divide. The capitalists would readily argue that business is better because it promotes competition. In some cases, they are probably correct, based on time-honored rules learned in business school and the tendency of governments to move lumberingly. But the fact is, there isn't much real competition in the broadband ISP market, so broadband isn't a very good example of a free market.

The gulf between these sweeping ideological viewpoints has led to the present state of dysfunction. One party gets in power and imposes net neutrality. The other party gets in and eliminates it. The next time the power structure in Washington shifts, the net neutrality pendulum will swing back again – unless Congress does something about it first, and (the truth be told) Congress really doesn't *do* very much.

Info

- [1] *City-Owned Internet Services Offer Cheaper and More Transparent Pricing*: <https://arstechnica.com/tech-policy/2018/01/city-owned-internet-services-offer-cheaper-and-more-transparent-pricing/>

So instead of working to solve the problem at a national or international level, with broad-reaching policy decisions by government-appointed czars and kingpins, the other approach is to just act locally, within your own community, and support municipal broadband.

Most municipal broadband systems are funded through bonds and do not receive direct government subsidies. They function very much like companies, but the difference is: The customers and the shareholders are the same people. They don't face the weird dilemma that broadband companies must face, where the customers and the shareholders are different people, and the best way to return value for the shareholders is to yank around the customers, which is basically what is happening with the end of net neutrality.

Acting locally is often a more reliable way of affecting public debate. You might not have a chance against an army of lobbyists in your national capital, but at your local city hall, the field really is a little closer to level. And if enough small towns and cities take up the challenge of providing their own Internet service so they can get it their own way, the commercial Internet providers might just wake up one day and remember another rule they probably learned in business school. *"Oh yeah...the customer is always right."*

Joe

Joe Casad,
Editor in Chief



LINUX MAGAZINE

WHAT'S INSIDE

This month we present some interesting utilities for enhancing your command-line experience. Other highlights include:

- **YunoHost** – A cool Linux alternative that easily installs an array of important server applications (page 22).
- **Lightpack** – Check out this weird technology that lights up the wall behind your TV (page 56).

Over at LinuxVoice, you'll learn about the new apt package tool and how it differs from apt-get. We'll also show you how to calibrate your monitor with DisplayCAL (page 78).

SERVICE

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

NEWS

08 News

- Dell Kickstarts 2018 with a Brand New Linux Laptop
- Linus Torvalds Rips Intel for Meltdown and Spectre Flaws
- LibreOffice-Based CODE 3.0 Released
- Google Announces Kubeflow to Bring Kubernetes to Machine Learning
- Critical Flaw in phpMyAdmin

12 Kernel News

- Implementing Digital Rights Management In-Kernel
- Improving Lighting Controls
- Updating printk()

REVIEWS

22 YunoHost

YunoHost offers a wide range of services on a proven Debian platform that you can host yourself.



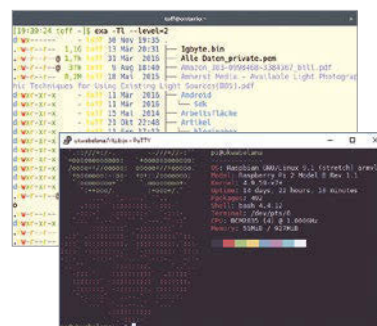
28 Fedora 27

Fedora Linux is a favorite of developers, sys admins, and tech-savvy users. We check out the latest release.

COVER STORY

16 Terminal Tuning

Tired of the same old Bash? We explore some helpful tools for extending and expanding your shell experience.



32 To-Do Managers

Busy people have busy calendars full of appointments and tasks. Linux to-do-managers help manage the clutter in a controlled way.



IN-DEPTH


- 42 **Command Line – unoconv**
A hidden utility in the LibreOffice toolbox, unoconv offers a wide array of import and export filter options.
- 46 **Programming Snapshot – Protectli**
What is making the lights on the router flicker so excitedly? An intruder? We investigate with pfSense on a Protectli micro appliance and a screen scraper to email the information.
- 50 **BackBox Linux**
Specializing in security and forensics, BackBox Linux is ideal for vulnerability assessment, but it also works well as an everyday desktop distribution.



- 54 **Charly's Column – Tsung**
Charly uses the Tsung load generator instead of human users as beta testers.

MAKERSPACE

- 56 **Lightpack**
Ambilight lights up the wall behind Philips TVs with LEDs mounted on the device to enhance the onscreen visual impression. All you need is Lightpack and a Rasp Pi.
- 60 **KiCad**
KiCad helps you design printed circuit boards with up to 32 layers and supplies schematics and assembly diagrams in the popular Gerber format.



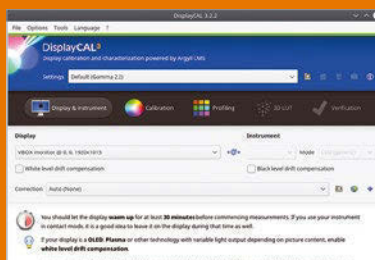
- 66 **Open Hardware – eelo**
Mandrake creator, Gaël Duval has big plans for a free-licensed phone.



LINUXVOICE

- 71 **Welcome**
This month in Linux Voice.
- 72 **Doghouse – Security**
A serious security problem requires your attention.
- 74 **apt vs. apt-get**
The apt command-line utility is a successor to the well-known apt-get, offering simpler installation and maintenance for the DEB packages used with Debian, Ubuntu, Knoppix, and many other Linux distros.

- 78 **DisplayCAL**
If you're tired of inaccurate colors on your monitor, try calibrating and profiling using a colorimeter and DisplayCAL.



- 84 **FOSSPicks**
Ocenaudio 3.3.6, Otter Browser, Joplin, WeeChat 2.0, Mailspring, Siril 0.9.7, SuperTuxKart 0.9.3, and more!



- 92 **Tutorial – USB/IP**
USB/IP lets you use USB devices connected to other machines as if they were plugged directly into your computer.

On the DVD

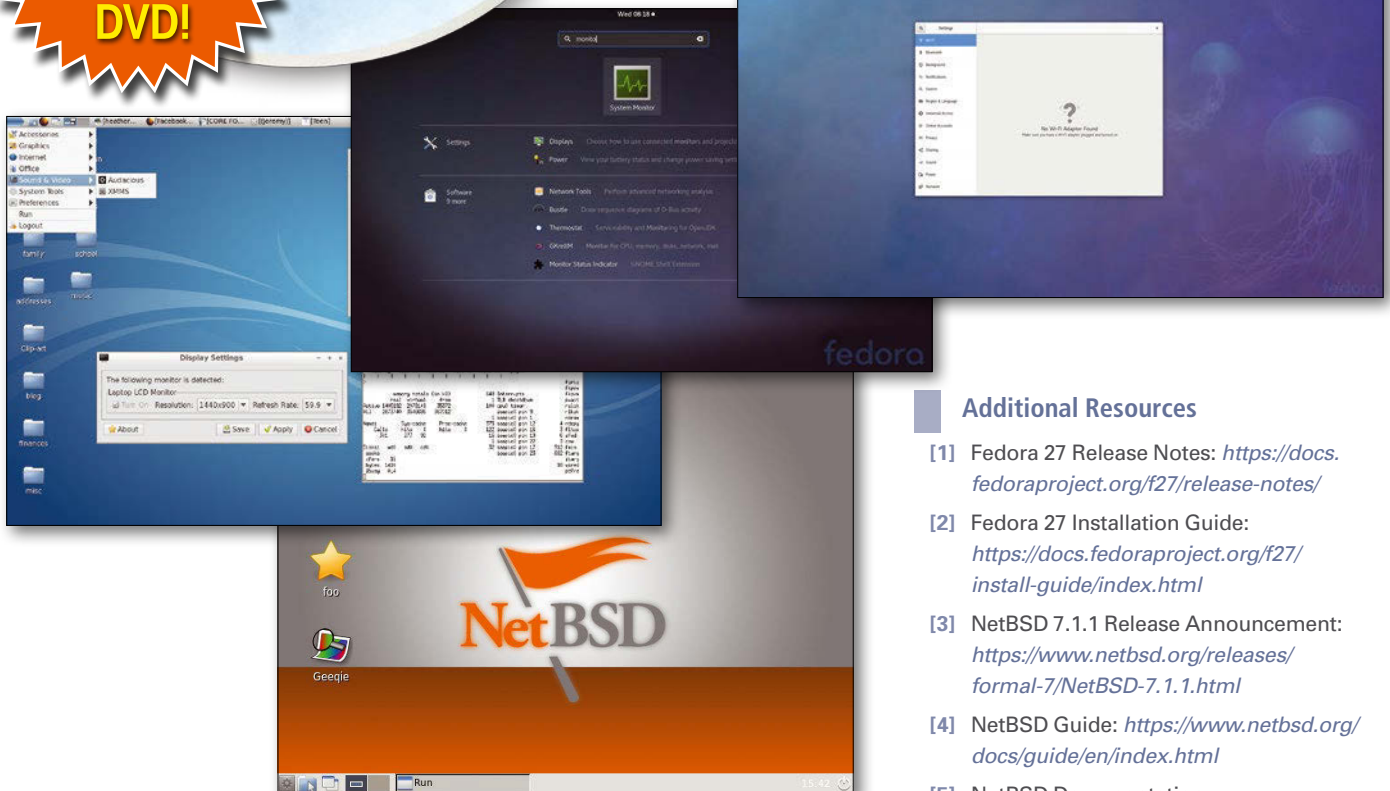


Fedora 27

Fedora is a Red-Hat-sponsored, all-Free Linux alternative with a reputation for the cutting edge. Many of the enterprise innovations that will one day make their way into Red Hat Enterprise Linux get their start in Fedora. Fedora 27 features the Gnome 3.26 desktop, with improved Display and Network configuration panels and a better system search feature. The latest Fedora also comes with support for color emojis, as well as many improvements with LibreOffice 5.4 and Fedora Media Writer.

NetBSD 7.1.1

The NetBSD "Unix-like" operating system is part of the famous BSD family of OS distributions that predated the development of Linux. The BSDs are known for stability, security, and a loyal community of users. NetBSD 7.1.1 is a security/critical update of NetBSD 7.1, which introduced Linux compatibility improvements, a new driver for the Google Computer Engine, better support for NVidia graphics cards, and other enhancements.



Additional Resources


- [1] Fedora 27 Release Notes: <https://docs.fedoraproject.org/f27/release-notes/>
- [2] Fedora 27 Installation Guide: <https://docs.fedoraproject.org/f27/install-guide/index.html>
- [3] NetBSD 7.1.1 Release Announcement: <https://www.netbsd.org/releases/formal-7/NetBSD-7.1.1.html>
- [4] NetBSD Guide: <https://www.netbsd.org/docs/guide/en/index.html>
- [5] NetBSD Documentation: <https://www.netbsd.org/docs/>

Defective discs will be replaced. Please send an email to subs@linux-magazine.com.

SUPERMICRO®

MicroBlade®

Highest-Density • High-Performance • High-Efficiency • Cost-Effective
Enterprise, Data Center, Web Applications, HPC and Cloud Computing Solutions



6U MicroBlade



3U MicroBlade

- Supports Intel® Xeon® Scalable Processors
- Best density and power efficiency with up to 56 UP or 28 DP server nodes
- M.2 NVMe SSD and HW SAS3 RAID support
- Ethernet switches supporting 1G, 2.5G and 10G downlinks with 96% fewer cables
- Redundant 2200W Titanium Level (96% efficiency) digital power supplies
- Battery Backup Power (BBP) modules with integrated power supply and cooling




Intel Inside®. Powerful Productivity Outside.

Learn more at [supermicro.com/X11](https://www.supermicro.com/X11)

NEWS

Updates on tech



THIS MONTH'S NEWS

08 Dell XPS 13 Developer Edition

- Dell Kickstarts 2018 with a Brand New Linux Laptop
- Linus Torvalds Rips Intel for Meltdown and Spectre Flaws

09 A Community Version of LibreOffice Online

- LibreOffice-Based CODE 3.0 Released

10 A Multicloud Solution for Machine Learning

- Google Announces Kubeflow to Bring Kubernetes to Machine Learning
- Critical Flaw in phpMyAdmin

Dell Kickstarts 2018 with a Brand New Linux Laptop

Dell is one of the only major PC vendors that sells Linux preloaded systems. The company has announced a brand new laptop from the XPS 13 family that runs Ubuntu Linux. The 7th generation XPS 13 Developer Edition (9370) is powered by Intel's 8th generation Quad Core processor and features a brand new chassis and a new display.

These machines are a result of Project Sputnik, which started back in 2012 and celebrated its 5th anniversary last year.

Dell XPS 13 Developer Edition features 8th Generation Intel Quad Core, i5 (US and Canada only), and i7 versions. It comes with three different configurations for memory – 4GB, 8GB, or 16GB Dual Channel SDRAM. Users can get up to 1TB of storage. The system also comes with a choice of the UltraSharp 4K Ultra HD (3840x2160) InfinityEdge touch display or the FHD (1920 x 1080) InfinityEdge display. It comes with a USB Type C port for transfer and two Thunderbolt 3 ports with PowerShare, DC-In, and DisplayPort. It also has a built in SD card reader.

Dell's engineering teams work with partners to develop drivers for supported Linux distributions. These machines comes with the LTS release of Ubuntu and offer a complete out-of-the-box experience with full support for touch screen. Although Ubuntu is the officially supported OS, users can wipe the hard drive and install any OS of their choice without worrying about the hardware warranty.

High demand for the system means that Dell is also making it available in some European countries, including the UK, Ireland, Germany, Austria, France, Italy, Spain, Switzerland (French and German), Belgium, Netherlands, Sweden, Norway and Denmark.



Photo by NordWood Themes on Unsplash

Linus Torvalds Rips Intel for Meltdown and Spectre Flaws

Linus Torvalds, the creator of the Linux kernel, is not known for mincing words when it comes to core technology. The world is still recovering from the shock waves of Meltdown and Spectre vulnerabilities, which affect almost every platform, including Intel, AMD, and ARM64. The news just broke that SPARC is also affected.

Out of all these companies, Intel gets the most criticism. Intel is the dominant player, so their chips dominate the market, which means more affected users. The company

MORE ONLINE

knew about the vulnerabilities for months. More than one team of researchers found the flaw and informed Intel about it. What's the possibility that it was also known to spy agencies and state-sponsored hackers?

Intel is facing the public's wrath, and its CEO sold all of his stock in the company (which is allowed by the bylaws) before the vulnerability went public.

However, Torvalds is concerned about only one thing: technology. Torvalds wrote in on the Linux kernel mailing list:

"I think somebody inside of Intel needs to really take a long hard look at their CPUs and actually admit that they have issues instead of writing PR blurbs that say that everything works as designed.

.. and that really means that all these mitigation patches should be written with "not all CPUs are crap" in mind.

Or is Intel basically saying "we are committed to selling you shit forever and ever, and never fixing anything"?

Because if that's the case, maybe we should start looking towards the ARM64 people more.

Please talk to management. Because I really see exactly two possibilities:

– Intel never intends to fix anything

OR

– These workarounds should have a way to disable them.

Which of the two is it?"

Many kernel developers have refrained from sharing their views on the subject because they either work for Intel or an Intel partner, so it's great to see Torvalds openly talking about it.



LibreOffice-Based CODE 3.0 Released

Collabora Productivity, a UK-based company that offers a cloud-based LibreOffice solution, has announced the release of CODE 3.0.

CODE is the community version of LibreOffice Online, which is available free to anyone who wants to run LibreOffice in their own cloud. In a press release, Collabora Productivity stated, "CODE is the LibreOffice Online solution with the latest developments, perfect for home users that want to integrate their own online Office Suite with their preferred File Share and Sync solution. It allows editing of richly formatted documents directly from a web browser, with excellent support for all popular office file formats, including text documents (docx, doc, odt, ...), spreadsheets (xlsx, xls, ods, ...), and presentations (pptx, ppt, odp, ...)."

Michael Meeks, general manager of Collabora Productivity, told us that 3.0 is an interesting release in which they have started to bring parts of the rich LibreOffice functionality to the browser. Combined with collaboration, it's easy to deploy and powerful to use. "In the Office world, people have a choice of any two of feature-depth, collaboration, or web deployment. We're starting to provide all three," said Meeks.

CODE 3.0 comes with many new features, including full-feature editing dialog, as seen in the desktop version of LibreOffice. The main purpose of CODE is to provide users early access to the very latest feature additions and updates to LibreOffice Online, to enable them to develop, test to make it better, and contribute back to LibreOffice.

Collabora sells a CODE-based commercial version called Collabora Online.

Linux Magazine

www.linux-magazine.com

ADMIN HPC

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

HPC Compilers • Carla Guillen and Reinhold Bader

If you compile software on an expensive supercomputer, it's a good idea to select the languages and compilers with particular care. We report on tried-and-proved tools used on SuperMUC, a supercomputer at the Leibniz Supercomputing Center in Germany.

ADMIN Online

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

Project management with Phabricator
Oliver Frommel

Bug trackers are the least thing that can be deployed to combat chaos in daily IT work. The free Phabricator software provides this service and many others.

Automated compliance testing with InSpec
Martin Loschwitz

Don't equate compliance through certification with security, because compliance and security are not the same. We look at automated compliance testing with InSpec for the secure operation of enterprise IT.

Correctly integrating containers
Thomas Fricke

If you run microservices in containers, they are forced to communicate with each other – and with the outside world. We explain how to network pods and nodes in Kubernetes.

ADMIN DevOps Focus

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

Docker Forensics • Chris Binnie


The handy auditd package can help track down weaknesses in your system before, during, or after an attack.

Google Announces Kubeflow to Bring Kubernetes to Machine Learning

After Kubernetes and TensorFlow, Google has now released Kubeflow, a new open source project that makes it easy to consume machine learning (ML) stacks with Kubernetes.

Kubernetes is being touted as the cloud Linux, and an increasing number of people are employing it in different use cases. Machine learning is one of the fastest growing use cases for Kubernetes, but it's quite a challenge to get the entire machine learning stack up and running.

"Building any production-ready machine learning system involves various components, often mixing vendors and hand-rolled solutions. Connecting and managing these services for even moderately sophisticated setups introduces huge barriers of complexity in adopting machine learning," said David Aronchick and Jeremy Lewi, project manager and engineer, respectively, on the Kubeflow project. "Infrastructure engineers will



often spend a significant amount of time manually tweaking deployments and hand rolling solutions before a single model can be tested."

Kubeflow solves this problem because it makes using ML stacks on Kubernetes fast and extensible. It's hosted on GitHub, and the repository contains three components: JupyterHub, to create and manage interactive Jupyter notebooks; a TensorFlow (TF) Custom Resource Definition (CRD) that can be configured to use CPUs or GPUs and adjusted to the size of a cluster with a single setting; and a TF Serving container.

Kubeflow is a multicloud solution. If you can run Kubernetes in your environment, you can run Kubeflow.

Critical Flaw in phpMyAdmin

A security researcher has found a critical flaw in phpMyAdmin that allows an attacker to damage databases. According to Hacker News, "The vulnerability is a cross-site request forgery (CSRF) attack and affects phpMyAdmin versions 4.7.x (prior to 4.7.7)."

The vulnerability was discovered by researcher Ashutosh Barot. Barot wrote in a blog post, "In this case (phpMyAdmin), a database admin/Developer can be tricked into performing database operations like DROP TABLE using CSRF. It can cause devastating incidents! The vulnerability allows an attacker to send a crafted URL to the victim and if she (authenticated user) clicks it, the victim may perform a DROP TABLE query on her database."

On its advisory page, phpMyAdmin wrote that "by deceiving a user to click on a crafted URL, it is possible to perform harmful database operations such as deleting records, dropping/truncating tables, etc." The phpMyAdmin project has already released a patch and suggests users either apply the patch to the existing installs or upgrade to phpMyAdmin 4.7.7 or newer.


phpMyAdmin is an open source tool for managing MySQL over the web. It supports a wide range of functions, including management of database, tables, columns, relations, indexes, users, permissions, etc. via the user interface, instead of using a command-line interface. This ease of use has made phpMyAdmin a very popular tool for hosting providers.



Image © Yevgeniya Ponomareva, Fotolia.com

Help your friends and colleagues 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

Zack's Kernel News



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

By Zack Brown

Implementing Digital Rights Management In-Kernel

Content providers are always interested in ways to stream audio and video in such a way that the data cannot be copied by the recipient. Sean Paul recently posted a patch that the Chrome OS team has been using for three years to control content on Exynos, MediaTek, and Rockchip hardware. The patch can be used to turn content protection entirely off, it can request that content protection be enabled by the hardware driver, and it can actually stream protected data.

The patch was received with suspicion by kernel developers.

Pavel Machek specifically said that he couldn't see any case where a user would set the feature to anything other than "off." He also asked, "If kernel implements this, will it mean hardware vendors will have to prevent user[s] from updating the kernel on machines they own?" And wondered, "If this is merged, does it open kernel developers to DMCA threats if they try to change it?"

Daniel Vetter made the case that this particular patch would only encrypt data over a cable, using High-bandwidth Digital Content Protection (HDCP). It wouldn't implement any other aspects of content protection and was thus a generic data security feature. He added, "If you want to actually lock down a machine to implement content

Author

The Linux kernel mailing list comprises the core of Linux development activities. Traffic volumes are immense, often reaching 10,000 messages in a week, and keeping up to date with the entire scope of development is a virtually impossible task for one person. One of the few brave souls to take on this task is **Zack Brown**.

protection, then you need secure boot without an unlockable bootloader and a pile more bits in userspace. If you do all that, only then do you have full content protection. And yes, then you don't really own the machine fully."

Pavel replied, "This patch makes it more likely to see machines with locked down kernels, preventing developers from working with systems on their own, running hardware. That is evil and [a] direct threat to the Free software movement."

He added, "Users compiling their own kernels get no benefit from it. Actually it looks like this only benefits Intel and Disney. We don't want that." And he concluded, "it does not belong in kernel."

Alex Deucher suggested that the patch might be useful for "sensitive video streams in government offices where you want to avoid a spy potentially tapping the cable to see the video stream." He added that it was not just Intel and Disney who would benefit, but also "just about every SoC manufacturer and Google and Amazon and a ton of other companies and organizations."

Alex suggested that if the kernel folks didn't want a patch like this, then they should also remove support for encrypted filesystems and encrypted virtual machines. To which Pavel replied, "Encrypted filesystems benefit users. Encrypted video is designed to work against users. In particular, users don't have encryption keys for video they generate. I'd have nothing against [a] feature that would let users encrypt video with keys they control."

Meanwhile, Sean pointed out that his patch only enabled features that were already present in the machine's hardware. He wasn't implementing anything new, just giving the kernel the ability to control hardware features

that were already present. He said, “those registers exist and `_can_` be used for HDCP; it’s just that now you know about it. Having all of the code in the open allows users to see what is happening with their hardware; how is this a bad thing?”

And along the same lines, Daniel also said to Pavel, “you can’t claim to speak for the entire kernel and FLOSS community of users and developers. The feature is optional: It does not enforce additional constraints on users but exposes additional functionality already present in hardware, for those who wish to opt in to it. Those who wish to avoid it can do so, by simply not making active use of it.”

At this point, Alan Cox came in, saying that he was speaking for himself at this time, and not for his employer, Intel. He said:

“The upstream policy has always been that we don’t merge things which don’t have an open usable user space. Is the HDCP encryption feature useful on its own? What do users get from it?”

If this is just an enabler for a lump of binary stuff in Chrome OS, then I don’t think it belongs; if it is useful standalone, then it seems it does belong?”

He also added in response to Alex’s scenario regarding secure government communications, “Last time I checked HDCP did not meet government security requirements – which is hardly surprising since you can buy \$10 boxes from China to de-HDCP video streams.”

Daniel pointed out that everything going into Chrome OS was open source and said that Chrome OS had very strict requirements about what could go into the userspace side of things as well.

That was the end of the discussion. It looks as if some very heavy hitters are opposed to this going into the kernel. Although in a case like this, it’s impossible to predict what Linus Torvalds will ultimately decide. Like Alan says, if there’s a legitimate use for the code, Linus would be more likely to include it; while if the only use is to lock people out of their own systems, Linus would refuse. It’s also possible that someone might notice a way that this hardware

could be used to benefit regular users, while completely failing to satisfy the large content providers. In that case, Linus would be likely to simply repurpose the hardware to do the user-beneficial thing, in complete disregard of the original intent of the hardware.

I would guess that, as it stands, the code is DOA, and the Chrome OS people will have to keep maintaining it on their own codebase for the foreseeable future.

Improving Lighting Controls

Enric Balletbo i Serra posted a patch to adjust the way Linux controls back-lit screens. As he pointed out, the human eye perceives changes to light levels differently in low-light situations. He felt that the CIE 1931 lightness formula represented the proper way to calculate Linux screen behavior during user configuration. He said, “This patch adds support to compute the brightness levels based on a static table filled with the numbers provided by the CIE 1931 algorithm.”

There were a few technical comments. Daniel Thompson felt the data table could be made smaller, and Pavel Machek felt the table could probably just be generated on the fly. But Enric replied, “This was discussed a bit in previous RFC which had the code to generate the table on the fly [...]. The use of a fixed table or an on-the-fly table is something that I’ll let the maintainers decide. I’ve no strong opinion on the use of the fly table.”

Other technical comments revolved around the relationship between the actual amount of light and the way a human would perceive it, and how best to adjust those numbers.

But there were no outcries against the patch. Everyone seemed in agreement that human perceptions should be given preference in Linux over those of other animals, even in spite of New Zealand’s 2015 legal recognition of animals as sentient beings. Hopefully a dolphin-centric version of Enric’s code will be coming soon.

Updating `printk()`

Sergey Senozhatsky posted a patch to give the `printk()` system call its own thread of execution on the running system. This has been an ongoing effort,

with many versions of the patch coming down the pike. The basic issue is `printk()` is not safe to call everywhere in the kernel. The reason is that the call to output its log messages, `console_unlock()`, may loop forever, which means that if `printk()` is called in an atomic (uninterruptible) context, it could lock the system. This isn’t really a danger, since kernel code knows to avoid that situation. But it does mean that `printk()` may not be called to log messages that really should be logged.

Sergey’s patch offloads the entire question to another thread, whose primary purpose is to recognize and break out of loops, returning control to the system. There are also certain emergency circumstances, like during a system panic, where normal interrupts no longer take place, and `printk()` must continue to log messages even without its dedicated thread.

As with everything having to do with bootup and shutdown, the code is insane. Hence, the many versions submitted for review.

Petr Mladek preferred a separate approach from Steven Rostedt, in spite of that code being complex and prone to false bug reports from people who couldn’t figure out what it was doing. On top of that, it didn’t fully solve the problem of locking the system. However, Petr preferred it because, as messy as it was, it represented a more modular approach that was still less complicated and insane than Sergey’s code.

Tejun Heo said he didn’t care which approach got in; he just wanted something that would work (i.e., not lock the system). But he felt both patches were way too complicated.

Meanwhile, although Steven’s version theoretically could still lock the system, no one has been able to reliably reproduce such a case. Steven came in at this point to remark, “I still don’t believe there is one. And it’s all hand waving until there’s an actual report that we can lock up the system with my approach.”

Sergey objected to Steven’s approach, because he said it required calling `printk()` from a CPU that was not in atomic mode. But he said, “what happens if there is NO non-atomic CPU or that non-atomic simply misses the `console_owner != false` point?”

The bottom line, Sergey said, was that the code had to do the exact right thing at the exact right time on the exact right CPU, and the exact right time was a vanishingly tiny window.

But this, in fact, was the case that Steven felt had not yet been proved. He asked if anyone could confirm that there was indeed a way to lock the system like that. Before going with Sergey's more complex approach, Steven wanted real evidence that his own not-quite-as-complex approach was truly insufficient.

Sergey replied that at his company the engineers encountered this problem fairly frequently; he posted some info about how it came about.

But Petr was still skeptical. He said, the "console_lock() owner is able to sleep. Therefore there is no risk of a soft lockup. Sure, many messages will get stacked in the meantime, and the console owner my get then passed to another owner in atomic context. But do you really see this in the real life?"

As he put it, "My current view is that Steven's patch could not make things worse. I was afraid of possible deadlock, but it seems that I was wrong. Other than that, the patch should make things just better because it allows you to pass the work from time to time a safe way."

But Sergey replied, "we are not looking for a solution that does not make things worse. We are looking for a solution that improves the thing."

And Petr replied that they should just push the code into the kernel and see what shook loose. If there were bug reports, then the kernel developers could act on them. But he didn't want to preemptively try to patch a bug that no one would ever encounter, at the cost of increasing code complexity. To which Steven agreed wholeheartedly.

Sergey began in earnest to try to design a sequence of events to trigger the lockup that he believed was there. But he also remarked, "I don't even understand what our plan is. I don't see how are we going to verify the patch. Over the last 3 years, how many emails do you have from Facebook or Samsung or Linaro, or any other company reporting the printk-related lockups? I don't have that many in my Gmail inbox, to be honest. and this is not because there were no printk lockups ob-

served. This is because people usually don't report those issues to the upstream community. Especially vendors that use outdated kernels, that are approx 1-2 year(s) behind the mainline. And I don't see why it should be different this time. It will take years before vendors pick the next LTS kernel, which will have that patch in it. But the really big problem here is that we already know that the patch has some problems. Are we going to conclude that 'no emails = = = no problems'? With all my respect, it does seem like, in the grand scheme of things, we are going to do the same thing, yet expect a different result."

He said in conclusion that Steven's patch did not meet the real world needs he was encountering at his own company, and that Steven's code resulted in worse behavior than the current version of printk(), given that it would sleep longer, cause userspace applications to time out, and still allowed the system lockup that Sergey wanted to fix.

Finally, Sergey did post some process traces showing a lockup. But Steven didn't find these convincing. He said, "The traces I've seen from you were from non-realistic scenarios. But I have hit issues with printk()s happening that cause one CPU to do all the work, where my patch would fix that. Those are the scenarios I'm talking about."

Steven and Sergey went back and forth for a bit, each growing more and more frustrated with the other.

Other people started jumping in at this point, and eventually Tejun returned, preferring Sergey's approach because he, too, was seeing the lockups and needed them to be addressed. He said to Steven, "I tried your v4 patch and ran the test module and could easily reproduce RCU stall and other issues stemming from a CPU getting pegged down by printk flushing." He continued, "this isn't a theoretical problem. We see these stalls a lot. Pre-emption isn't enabled to begin with. Memory pressure is high, and OOM triggers and printk starts printing out OOM warnings; then, a network packet comes in, which triggers allocations in the network layer, which fails due to memory pressure, which then generates memory allocation failure messages, which then generates netconsole

packets, which then tries to allocate more memory and so on. It's just that there's no one else to give that flushing duty too, so the ping-ponging that your patch implements can't really help anything."

Tejun concluded, "You argue that it isn't made worse by your patch, which may be true, but your patch doesn't solve actual problems and is most likely an unnecessary complication, which gets in the way for the actual solution."

Steven looked at Tejun's scenario, but said, "WTF! You are printing 10,000 printk messages from an interrupt context??? And to top it off, I ran this on my box, switching printk() to trace_printk() (which is extremely low overhead). And it is triggered on the same CPU that did the printk() itself on. Yeah, there is no hand off, because you are doing a shitload of printks on one CPU and nothing on any of the other CPUs. This isn't the problem that my patch was set out to solve, nor is it a very realistic problem."

Well, the discussion is ongoing. We have multiple people accusing each other of not reading what they're writing, curses flying back and forth, and people generally talking past each other. There's no end in sight.

And yet, even in the heat of frustration and disagreement, both sides are still taking each other seriously and trying to address each other's concerns. Steven has begun to think that Tejun's scenario is a bug in another part of the kernel code and has begun trying to diagnose that. At which point, possibly, both Sergey and Tejun would agree that Steven's code addresses the only real problems that remain. Meanwhile, both Sergey and Tejun have been trying to post real-world scenarios, process traces, and code that reveal the bug that Sergey's code attempts to fix, but that Steven's does not.



And of course, this debate has been ongoing for quite a while already, with multiple patches, and multiple debates, occurring over several years. There's no way to know where this will lead or how it will pan out. The code is already insanely complex. The patches to fix it are more complex still. And workloads that reproduce the problems may or may not be related to what the developers are trying to fix. ■■■

THE COMPLETE

LINUXVOICE

ARCHIVE DVD: ALL 32 ISSUES!

3,400 PAGES OF



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 EPUB, PDF, and HTML format!

Order now!
shop.linuxnewmedia.com

Useful helpers for the shell

Terminal Tuning

Some users have been working at the command line for years and wouldn't change a thing about it, but others ride a wave of constant innovation. We explore some helpful tools for extending and expanding your shell experience.

By Christoph Langner

The Bash shell is one of the central building blocks of any Linux system. Beginners are often a little lost, but they know they'll never become experts without some basic knowledge of the mystical command line. Advanced users love the convenience of navigating their Linux life using simple and practical text commands. For many users, the Bash shell is like everything else in Linux: Sure it's great, but it could always get better.

With a few tricks and some useful helper tools, you can add new capabilities and features to the venerable Bash shell. This article explores some favorite terminal tuning tricks.

Better History

Bash stores each command in the `~/.bash_history` file. You can display the command history by typing `history`. Bash searches the history if you press `Ctrl + R` in the terminal and then enter a search term. If you repeat the `Ctrl + R` key combination, the shell jumps to the next hit; `Ctrl + Shift + R` takes it in the other direction.

Searching helps you dig up old commands so you don't have to type them again. The search function always sorts the hits from new to old. The last entered commands always appear first when searching with `Ctrl + R`.

Artful use of the Bash search feature can lead to some time-saving tricks. For instance, if you enter a long command that you don't want to have to type again, add a short comment string to the end of it using a hash tag. For example, instead of `long command`, enter `long command #quick`. The shell interprets the text after the hash tag as a comment and does not execute it. However, if you want to access the command later, just type `Ctrl + R` and `quick`, and Bash will jump to the desired command.

Hstr

Perhaps you think it would be more convenient if the most frequently used commands were to appear first instead of last. Exactly this and much more is possible with Hstr [1], also known as the Bash and Zsh Shell History Suggest Box, a special tool designed to let you "easily view, navigate, search, and manage your command history."

Not all distributions include Hstr directly in the package archives; for example, Ubuntu requires you to integrate a PPA. Installation instructions for several Linux distros are available on the Hstr GitHub project page [2].

After you set up Hstr, transfer its default configuration to the configuration file for the Bash shell and reload it so that the changes takes effect without a new logon (Listing 1). Hstr uses the `hh` command.

Now when you launch a command search in the shell using `Ctrl + R`, Hstr will start automatically: Instead of just one command, you will see an Ncurses interface with the most frequently entered commands. If you type in a search term, Hstr filters out

Listing 1: Hstr Configuration

```
$ hh --show-configuration >> ~/.bashrc
$ source ~/.bashrc
```


marked as favorites with Ctrl + F and brought back to life by changing the Hstr sort order with Ctrl + Shift + 7.

Better Lists

The `ls` command, which lists the content of a directory, is one of the most common Bash commands. If you're looking for more information, try extending the simple `ls` command with `ls -alh`. The `-alh` options turn the simple list into a complete overview, including file permissions, ownership, file size, and more. The command also outputs hidden files and directories.

The `ls` command also supports colors, but usually only color highlighting for directories and links. The alternative Exa [3], which is written in Rust, offers significantly more convenience. Only a few distributions offer Exa in their package sources, although you will find it in the Arch User Repository (AUR). If your distro doesn't provide a package for Exa, check out the project homepage,



the corresponding hits and sorts them according to their frequency (Figure 1). Use the arrow keys to scroll through the list; pressing the Enter key executes the currently selected command. Use a tabulator to transfer the command to the shell so that you can still make changes.

If a command with security-critical data slips into your history (e.g., the password when transferring data with ftp), you can remove this command from the Bash history by pressing Del. Frequently used commands can be

```
toff@ontario:~/bin
toff@ontario$ htm
Type to filter, UP/DOWN move, DEL remove, TAB select, C-f add favorite, C-g can
-HISTORY - view:ranking (C-7) - match:exact (C-e) - case:insensitive (C-t) - 3-
make-artikel --html
make-html
joe index.html
epiphany --private-instance index.html
epiphany --new-window index.html
epiphany -i index.html
sf -t -i -s -E --user artikel.txt index.html
7z e erimake-artikel --html
sf -t -i -s -U artikel.txt index.html
make-artikel --html#
sf -t -i -s -E artikel.txt index.html
sf -t -i -s -E artikel.txt test.html
sf -t -i -s artikel.txt test.html
epiphany -a index.html
sf -t -i artikel.txt test.html
which make-html
epiphany index.html
epiphany index.html &
sf -t artikel.txt test.html
sf artikel.txt test.html
```

Figure 1: Hstr sorts the Bash history by frequency of execution, not chronologically.

Table 1: Exa Important Options

Switch	Long Form	Comment
Display Options		
-l	--oneline	Shows each file and folder in a new line
-l	--long	Displays metadata, such as file permissions and size
-R	--recurse	Lists the contents of subdirectories recursively (depth option: --level=<level_number>)
-T	--tree	Displays a tree structure (depth option: --level=<level_number>)
	--color <when>	Color output is never, always, or automatic
	--color-scale	Highlights the display of the file size in color for large files
Filtering and Sorting Options		
-a	--all	Shows hidden files and folders
-d	--list-dirs	Treats directories like files
-r	--reverse	Reverses the sort order
-s <option>	--sort=<option>	Sorts by name, size, created, newest, or oldest, among other options
	--group-directories-first	Lists directories before files

```

toff@ontario:~/bin
[11:15:36 toff bin]$ ll
Permissions Links Size Blocks User Group Date Modified Name
-rwxr-xr-x 1 2,1Ki 8 toff users 18 Dez 2015 android-scrout
-rwxrwxr-x 1 164 8 toff users 24 Jan 2013 andyshooter
-rwxr-xr-x 1 81 8 toff users 4 Jan 2014 ausgrossmachklein.sh
-rw-r--r-- 1 458 8 toff users 16 Okt 2015 backup-lui.sh~
-rwxr-xr-x@ 1 742 16 toff toff 6 Sep 18:42 dd++
-rwxr-xr-x 1 4,2Ki 16 toff users 2 Mai 2010 duext
-rwxr-xr-x 1 667 8 toff users 12 Feb 2013 echo_lc.py
-rwxrwxr-x 1 11Ki 24 toff users 21 Dez 2011 f.sh
-rwxrwxr-x 1 4,5Ki 16 toff users 20 Jan 2012 ipwebcam.sh
-rwxrwxr-x 1 514 8 toff users 7 Mai 2012 lnmount
-rwxrwxr-x 1 32 8 toff users 15 Dez 2011 lnmount
-rwxr-xr-x 1 74 8 toff users 25 Feb 2016 lui
-rwxr-xr-x 1 1,2Ki 8 toff toff 21 Aug 16:00 make-artikel
-rwxr-xr-x 1 378 8 toff users 29 Jan 2013 maxdome-desktop
-rwxr-xr-x@ 1 4,1Ki 24 toff users 25 Feb 14:23 pishrink.sh
-rwxr-xr-x 1 82 8 toff users 15 Mai 2014 png2jpg
-rwxr-xr-x 1 80 8 toff users 15 Mai 2014 png2jpg-400
-rwxr-xr-x 1 80 8 toff users 15 Mai 2014 png2jpg-768
-rwxr-xr-x 1 71 8 toff users 15 Mai 2014 png2jpg~
-rwxr-xr-x 1 741 8 toff users 1 Mär 2013 qrcodegenerator
-rwxrwxr-x 1 520 8 toff users 20 Apr 2012 rainbow
-rwxrwxrwx 1 284Ki 576 toff users 8 Feb 11:32 sf

```

Figure 2: As an alternative to `ls`, Exa provides more information and presents the output in a clearly arranged format with colors.

where the developer provides a zipped binary for 64-bit systems, which you can simply copy as `exa` to `/usr/local/bin/` or `~/bin/`.

The output of the `exa` command corresponds to the output of `ls`, except that Exa adds more colors to the overview (Figure 2). For example, it highlights archives in red, images in light purple, videos in dark purple, backup copies in gray, and directories in light blue. Exa gets even more interesting when you enable further options (see Table 1). For example, with `exa -bghHHS`, Exa displays the file permissions, including file sizes, affiliations, and the last modification data, in a neatly sorted list.

Other colors help you keep track of things such as file permissions. If you decide on a version of the command you like best, create an alias for it with

```

toff@ontario:~
[19:39:24 toff ~]$ exa -Tl --level=2
drwx----- - toff 30 Nov 19:35 .
-rw-r--r-- 1,1G toff 13 Mär 20:31 | 1gbyte.bin
-rw-r--r--@ 1,7k toff 31 Mär 2016 | Alle Daten_private.pem
-rw-r--r--@ 37k toff 9 Aug 18:40 | Amazon_303-0998468-3384367_bill.pdf
-rw-r--r-- 8,2M toff 18 Mai 2015 | Amherst Media - Available Light Photograph
hic Techniques for Using Existing Light Sources(BBS).pdf
drwxr-xr-x - toff 11 Mär 2016 | Android
drwxr-xr-x - toff 11 Mär 2016 | Sdk
drwxr-xr-x - toff 15 Mai 2014 | Arbeitsfläche
drwxr-xr-x - toff 21 Okt 22:48 | Artikel
drwxr-xr-x - toff 11 Sep 17:12 | bloginbox
drwxr-xr-x - toff 21 Okt 22:49 | cirrus7
drwxr-xr-x - toff 12 Sep 15:47 | countdown
drwxr-xr-x - toff 8 Sep 14:59 | icydock
drwxr-xr-x - toff 8 Sep 12:30 | lightpack
drwxr-xr-x - toff 14 Sep 16:12 | pixtend
drwxr-xr-x - toff 17 Feb 2016 | Aufnahmen
-rw-r--r--@ 499k toff 10 Mai 12:34 | auto backup 20170510 0954-201705100954.pr
o
-rw-r--r-- 11k toff 7 Jan 2:36 | bewerbung.tex
drwxr-xr-x - toff 30 Nov 13:36 | Bilder
-rw-r--r-- 7,8M toff 1 Jul 2014 | erika-hering.jpg
-rw-r--r-- 280k toff 29 Dez 2015 | erika-puenktchen.jpg

```

Figure 3: The Exa tree view allows you to keep track of even the most branched directory structures.

```
alias ll='exa -bghHHS'
```

and place the alias in the `~/bashrc` file. The preceding alias would let you call the command by simply typing `ll`.

Exa's ability to display entire folder structures recursively in a tree view is also practical (Figure 3). To display a tree, supplement the call with `exa --tree` or `exa -T`. You can limit the tree depth with the `--level=<n>` option, where `<n>` is the number of levels you want to display. Alternatively, switch off the tree view and let Exa list just the subfolders with the `-R` or `--recurse` switch. Again, you can restrict the depth with the `--level` option.

The shortcuts for sorting the output also make everyday life in the shell easier. For example, to find the largest files in a folder quickly, call Exa with:

```
exa -lr --sort=size
```

The `-r` switch reverses the order of the output so that the largest files are at the beginning of the list.

Better Unpacking

Linux users often have to deal with file archives in ZIP, RAR, or tarball (`tar.gz`, `tar.bz`) format. The commands for unpacking these archives are as different as the formats: `unzip`, `unrar`, `tar xzf`.

The `tar` command, in particular, regularly challenges users with its cryptic syntax. Particularly annoying is an archive full of files that are not stored in a separate subdirectory when unpacked but instead clutter the current folder. Such shortcomings are addressed by the `Dtrx` [4] unpacker – the abbreviation stands for “do the right extraction.”

All the popular distributions have `Dtrx` in their package sources; under Ubuntu, the installation includes a number of tools for unpacking. On the other hand, with Arch, you have to install the necessary tools yourself. As a wrapper, `Dtrx` is not able to unpack archives itself, but it composes the required syntax for you. You only have to remember `dtrx <file_name>` to unpack an archive.

Figure 4 shows `Dtrx` in action: The badly built `bad-tarball.tar.xz` tarball does not contain a base subdirectory, so all the data will end up in the current folder when unpacking with `tar xf`. Additionally, `tar` adopts the file permissions without changes – in this case, you are not allowed to edit the files without first modifying the permissions using `chmod` or a file manager. `Dtrx` corrects all of this in one command.



Better Terminal

Perhaps you are looking for an alternative to the terminal emulator itself? If so, it's worth taking a look at Tilix [5]. The Tilix terminal emulator outperforms most other terminal applications and offers some interesting enhanced features. For example, Tilix allows you to divide the window into several sub-terminals. You can display the current workload in an Htop window zone, output system information in the second window zone, and continue working as usual in the main section (Figure 5). You can arrange the individual sub-windows by dragging and dropping or by pulling a sub-window out of the application window as a new terminal.

Palette Pick

Some users like to deploy different colors in the terminal to highlight text and organize the information. Many terminal programs provide color profiles, which you can enable in the settings with just a few clicks of the mouse. If you want colors but you're weary of the drab, elementary color palette used with

```

toff@ontario:/tmp/test
[15:37:23 toff test]$ tar xf bad-tarball.tar.xz
[15:37:32 toff test]$ ll -T --level=2
Permissions Links  Size Blocks User Group Date Modified Name
drwxr-xr-x    4    -    - toff toff  1 Dez 15:37
dr-xr-xr-x    3    -    - toff toff 14 Sep  2:15
drwxr-xr-x    2    -    - toff toff 14 Sep  2:15
-rw-r--r--    1 63Ki  128 toff toff  1 Dez 15:34
-r--r--r--    1 3,9Ki   8 toff toff 14 Sep  2:15
dr-xr-xr-x    2    -    - toff toff 14 Sep  2:15
-rw-r--r--    1 14Ki   32 toff toff 14 Sep  2:15
-rw-r--r--    1 1,1Ki   8 toff toff 14 Sep  2:15
-r--r--r--    1 1,1Ki   8 toff toff 14 Sep  2:15
-r--r--r--    1 1014    8 toff toff 14 Sep  2:15
-r-xr-xr-x    1 146Ki  296 toff toff 14 Sep  2:15
-r--r--r--    1 9,7Ki  24 toff toff 14 Sep  2:15
-r--r--r--    1 3,5Ki   8 toff toff 14 Sep  2:15
[15:37:37 toff test]$

toff@ontario:/tmp/test
[15:38:04 toff test]$ dtrx bad-tarball.tar.xz
[15:38:12 toff test]$ ll -T --level=2
Permissions Links  Size Blocks User Group Date Modified Name
drwxr-xr-x    3    -    - toff toff  1 Dez 15:38
drwx-----  4    -    - toff toff  1 Dez 15:38
drwxr-xr-x    3    -    - toff toff 14 Sep  2:15
-rw-r--r--    1 3,9Ki   8 toff toff 14 Sep  2:15
drwxr-xr-x    2    -    - toff toff 14 Sep  2:15
-rw-r--r--    1 1,1Ki   8 toff toff 14 Sep  2:15
-rw-r--r--    1 1014    8 toff toff 14 Sep  2:15
-rwxr-xr-x    1 146Ki  296 toff toff 14 Sep  2:15
-rw-r--r--    1 9,7Ki  24 toff toff 14 Sep  2:15
-rw-r--r--    1 3,5Ki   8 toff toff 14 Sep  2:15
-rw-r--r--    1 63Ki  128 toff toff  1 Dez 15:34
[15:38:16 toff test]$
    
```

Figure 4: Dtrx automatically corrects read and write permissions during unpacking and ensures order by creating a subdirectory.

Shop the Shop

shop.linuxnewmedia.com

Missed an issue?

You're in luck.

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

shop.linuxnewmedia.com

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



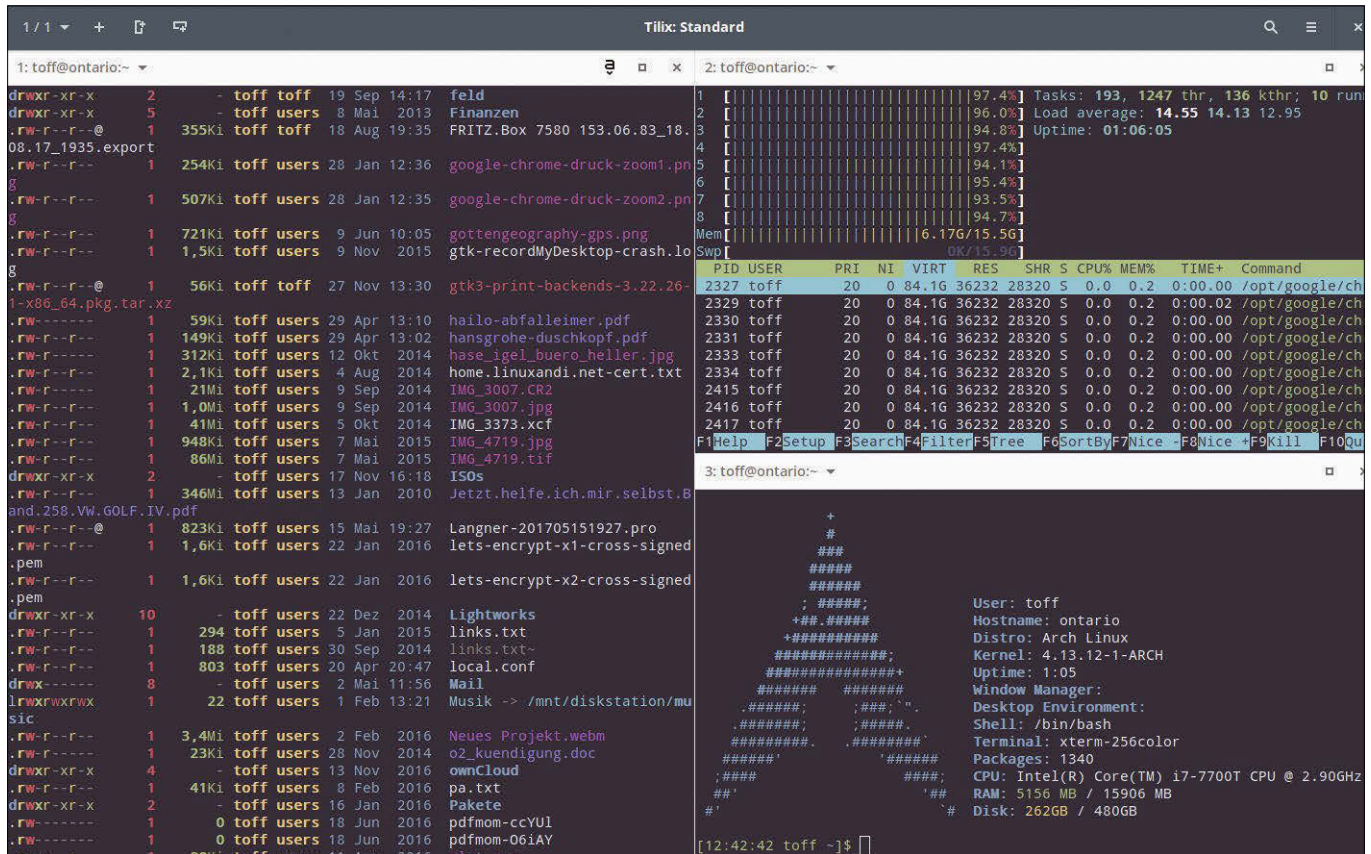


Figure 5: The modern terminal emulator Tilix offers many useful functions. For example, you can divide the window into several sub-windows.

most terminal emulators, you might want to try Nord [6], which the project website describes as “an arctic, north-bluish color palette.” Nord is specially designed “to achieve optimal focus and readability for code syntax highlighting and UI.” The Nord palette “consists of a total of sixteen, carefully selected, dimmed pastel colors for an eye-comfortable, ... yet colorful ambiance.”

The Nord pastel color profile is available for numerous programs, from classic terminal applications and Putty (also for Windows) to corresponding profiles for editors or development environments (Figure 6). On its GitHub page, the project links to the different applications.

Conclusions

If you’re looking for something more than the usual shell, the tools and extensions described in this article will bring some extra power and flash to your command-line experience. When you’re ready for more shell enhancements, you might want to check out Awesome Shell [7], which calls itself “a curated list of awesome command-line frameworks, toolkits, guides, and gizmos.” For almost all your applications and needs, there is a tool in the open source ocean to help you – you just have to find it. ■■■

Info

- [1] Hstr: <https://github.com/dvorka/hstr>
- [2] Hstr Installation: <https://github.com/dvorka/hstr/blob/master/INSTALLATION.md>
- [3] Exa: <https://the.exa.website>
- [4] Dtrx: <https://brettcsmith.org/2007/dtrx/>
- [5] Tilix: <https://github.com/gnunn1/tilix>
- [6] Nord: <https://github.com/arcticicestudio/nord>
- [7] Awesome Shell: <https://github.com/alebcay/awesome-shell>

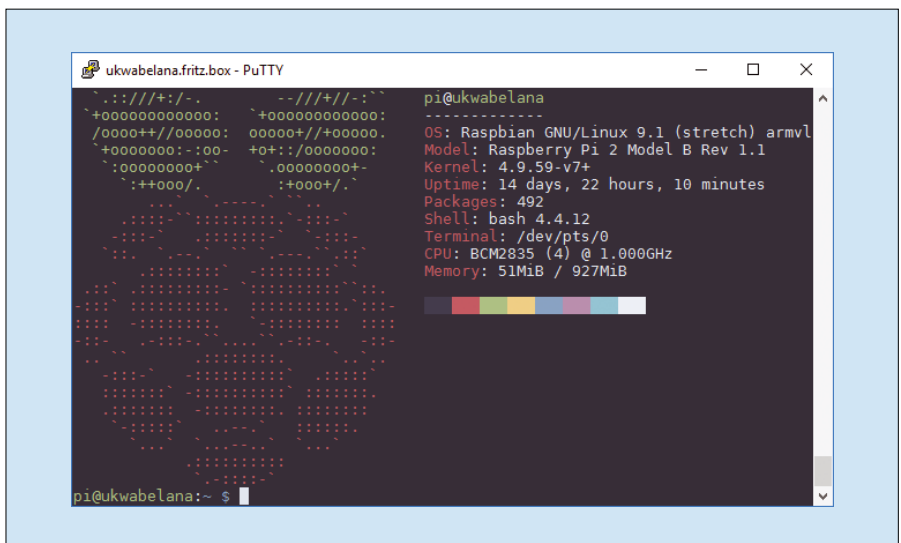



Figure 6: If you prefer an “arctic, north-bluish” color palette in the terminal, try out Nord.

EXPERT TOUCH



Linux professionals stay productive at the Bash command line – and you can too!

The Linux Shell special edition provides hands-on, how-to discussions of more than 300 command-line utilities for networking, troubleshooting, configuring, and managing Linux systems. Let this comprehensive reference be your guide for building a deeper understanding of the Linux shell environment.

You'll learn how to:

- Filter and isolate text
- Install software from the command line
- Monitor and manage processes
- Configure devices, disks, filesystems, and user accounts
- Troubleshoot network connections
- Schedule recurring tasks
- Create simple Bash scripts to save time and extend your environment

8th Edition!

The best way to stay in touch with your system is through the fast, versatile, and powerful Bash shell. Keep this handy command reference close to your desk, and learn to work like the experts.

ORDER ONLINE:

shop.linuxnewmedia.com/specials

Self-hosting solution YunoHost for the private cloud

With Simple Means

YunoHost offers a wide range of services on a proven Debian platform that you can host yourself. *By Ferdinand Thommes*



YunoHost [1] is a Linux distribution that offers a basic server platform with single sign-on for hosting web, mail, XMMS, and several dozen other services. According to the project website, YunoHost's goal is to provide "...a server operating system aimed at making self-hosting accessible to everyone."

YunoHost comes with the nginx web server, MariaDB, mail transfer agent Postfix, and the IMAP server Dovecot, including the Rspamd spam filter. The Metronome IM XMPP server, OpenLDAP, Dnsmasq, and the SSOwat transparent authentication system for single sign-on are also available, as is Let's Encrypt for creating SSL certificates.

YunoHost also provides a community repository of validated helper scripts to install other services and applications. The repository, which is hosted on GitHub, provides customized versions of the applications preconfigured to integrate into YunoHost's single sign-on environment.

YunoHost is similar to other personal server systems, such as FreedomBox [2],

Sandstorm [3], UBOS [4], or the recently discontinued arkOS. All of these projects have the goal of letting the user operate a server and install web services in the easiest possible way.

YunoHost is based on Debian and is available for download for the i386, AMD64, PowerPC, and ARM architectures. A demo site [5] lets you test the user interface in advance.

Well Equipped

The range of official apps for various web services includes 21 packages [6], which you can install with just one click. This list encompasses well-known services such as WordPress, Nextcloud, Roundcube, and DokuWiki, as well as lesser known tools such as the Baikal [7] CardDAV server, the RainLoop [8] web-mail application, and the ZeroBin [9] encrypted paste service.

The list of unofficial apps [10] developed and maintained by the community is much longer: It includes around 100 programs, such as a sync server for Firefox [11], the Gogs [12] Git service,

the Jenkins [13] CI server, Mattermost [14] as an alternative to Slack, the Piwigo [15] photo gallery, and many more.

But that's not the end: The *Apps in Progress* section offers several apps you can test, but without any guarantees.

Which Platform?

YunoHost will run on the Raspberry Pi 3, and even on earlier versions of the Rasp Pi. On the small scale of a home network, the Rasp Pi might be satisfactory, but if you are concerned about performance and throughput, the limits of the Rasp Pi are quickly evident. You can also install YunoHost on a local computer (see the "Installing on Your PC" box) or on a VServer with a professional hosting service provider. If Debian is already running on the computer on the base system, you can use a script to set up YunoHost [16].

You can also use a VirtualBox, VMWare, or KVM virtual machine as the basis for your tests. An image for the virtualization-focused Vagrant development environment is also available [17].

Installing on Your PC

To install under VirtualBox or on a 32- or 64-bit PC, download the appropriate ISO image from the project website. Leave the default disk-partitioning values. The install, including an update, takes just a few minutes to complete; the system then automatically reboots to a command line. Now log in as *root* with a password of *yuno-host* and start the post-install routine by calling `yunohost tools postinstall`. The routine requires the main domain, as well as the future administrative password, and then configures all services accordingly.

Installing on Rasp Pi

To install YunoHost on a Rasp Pi, first download the image for the ARM platform and store it on an SD card, preferably a Class 10 card with at least an 8GB capacity. Be sure to choose the target device carefully; otherwise you will lose the data on a partition that your computer needs. The `dd` command is fine for writing the data, as are graphical tools like Etcher under Linux or Rufus under Windows. The corresponding ARM version of the Rasp Pi is based on Raspbian 8 “Jessie.”

After copying the image to the SD card, insert the card into the slot on the Rasp Pi and connect the Ethernet cable, the power supply, and, optionally, a display and a keyboard. The boot process takes about 90 seconds.

The next step is to determine the Rasp Pi’s IP address. If you have connected a display, the computer will show you its IP address. If no display is connected, the `hostname -i` command on the computer’s console helps. Alternatively, you can find the address via the router’s web interface. Under Windows, you can use the Advanced IP Scanner [18] program to display IP addresses logged on to the local network. You can do the same under Linux with on-board tools:

```
$ sudo arp-scan --localnet | \
grep Raspberry
```

You now have to finish the installation. Without a display, continue via SSH or call the `http://<RasPi-IP>` address in a browser to continue graphically. Since YunoHost uses a self-signed certificate

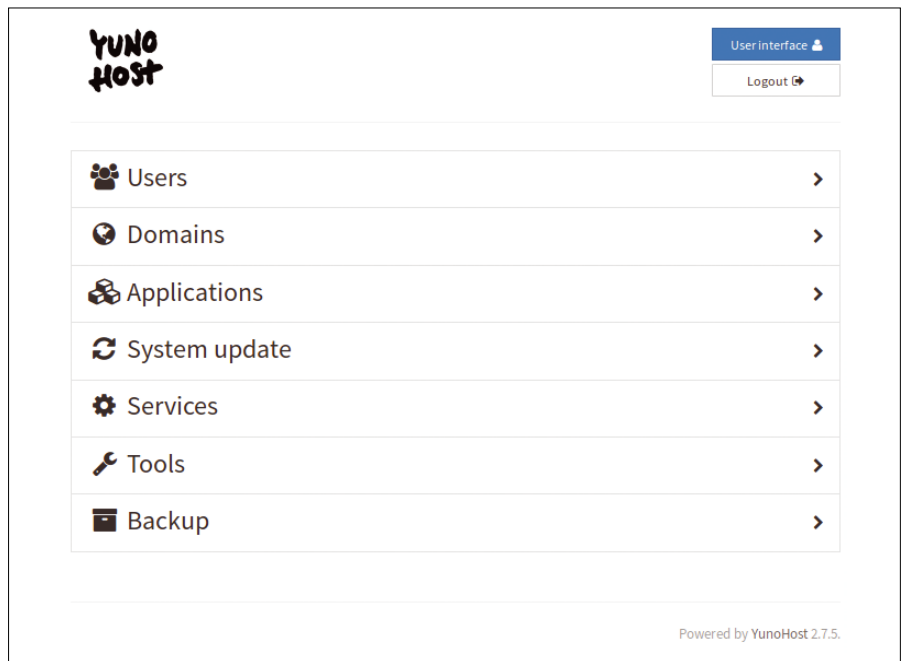


Figure 1: A first look at the back end shows all categories, including the area where you configure the services on the host.

by default, the browser displays the message that the connection is not secure. You can ignore this message in a private environment and allow an exception to reach the interface (Figure 1).

If you want to grant other users external access to services, you should use your own domain. Use the `yunohost domain cert-install` command to create an SSL certificate with Let’s Encrypt and install it so that it secures the connection via HTTPS in the future. However, this approach presup-

poses that the domain registrar first adapts the DNS A record so that it points to the domain’s IP address. It usually takes 24 hours or longer for the change to propagate.

The first step in the interface is to set up a domain under which you can access the server (Figure 2). Guided configuration in the browser is easier for most users than working with SSH in a terminal. The first question is whether you already have a domain that you want to use. The example assumes that

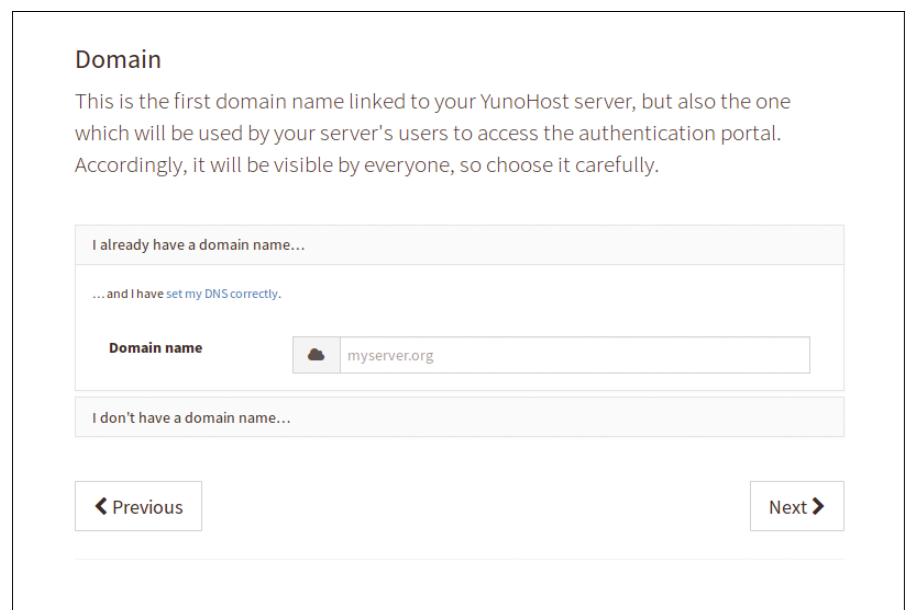


Figure 2: If you want to allow users to access the services on the host, configure a domain.

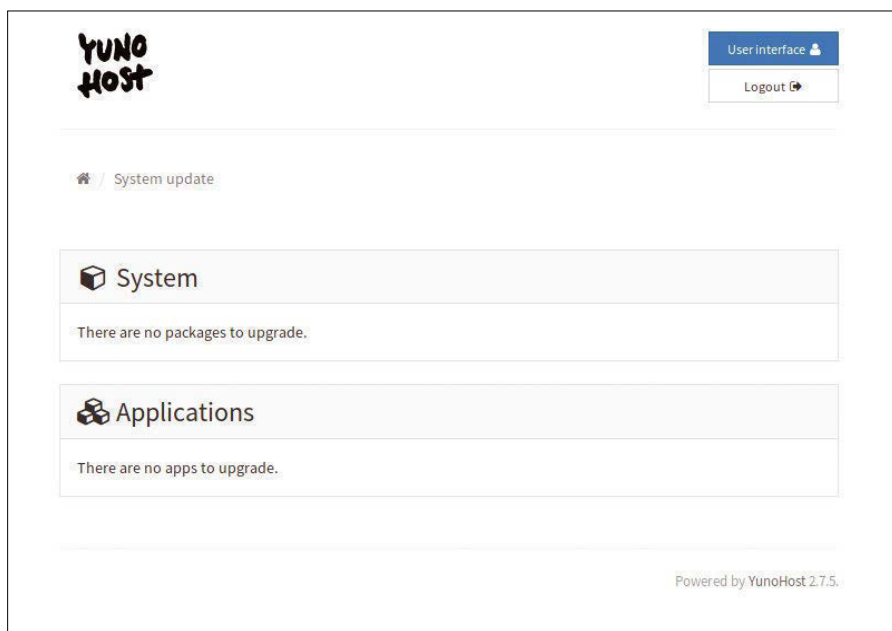


Figure 3: I recommend periodically updating the system to patch security vulnerabilities in the installed packages.

you don't have a domain, so the second option is the right choice.

However, if you want to use a self-defined local domain, such as yunohost.local, enter it along with the IP address in the /etc/hosts file. Under Windows, you will find the corresponding file under %SystemRoot%\system32\drivers\etc\.

DNS Knowledge

To provide services beyond the host machine, adjust the router's DNS settings. A look at the documentation [19] is helpful. However, beginners can follow the example and use one of the two YunoHost domains for a start. Multiple domains are possible, and you are allowed to mix your own domains and subdomains with YunoHost domains.

In the example without your own domain, enter the desired subdomain in the dialog and select one of the two domains offered to you. When assigning names, keep in mind that everyone who logs into the server uses this default domain.

In the example, ft is used as a subdomain. The default domain is ft.nohost.me. A first attempt with yunohost as a subdomain failed, because the system was not able to complete the configuration – but for no apparent reason.

Time for a Cup of Tea

After creating a server administration password, you'll need to set up the net-

work, which may take a few minutes. You will then find yourself in the administration interface.

The first step is to update the system. Import new packages from both the Debian server and YunoHost (Figure 3) under the *System update* category. Be sure to update the system at regular intervals, because only an up-to-date server system offers sufficient security.

For all actions, you will see a status display at the top right-hand corner of the screen that displays the actions currently running. A list of actions in progress is displayed when you mouse over the items.

Next, you should create at least one user to log in to the user area. Afterwards, start installing applications. For this test, the first choices were Baikal, Wallabag, and Nextcloud. The installation and basic configuration required only one mouse click (Figure 4).

To test the installed applications, log in as the previously created user and access the services (Figure 5). The initial configuration of the apps I installed proved to be very good. For example, Nextcloud already had the required LDAP information.

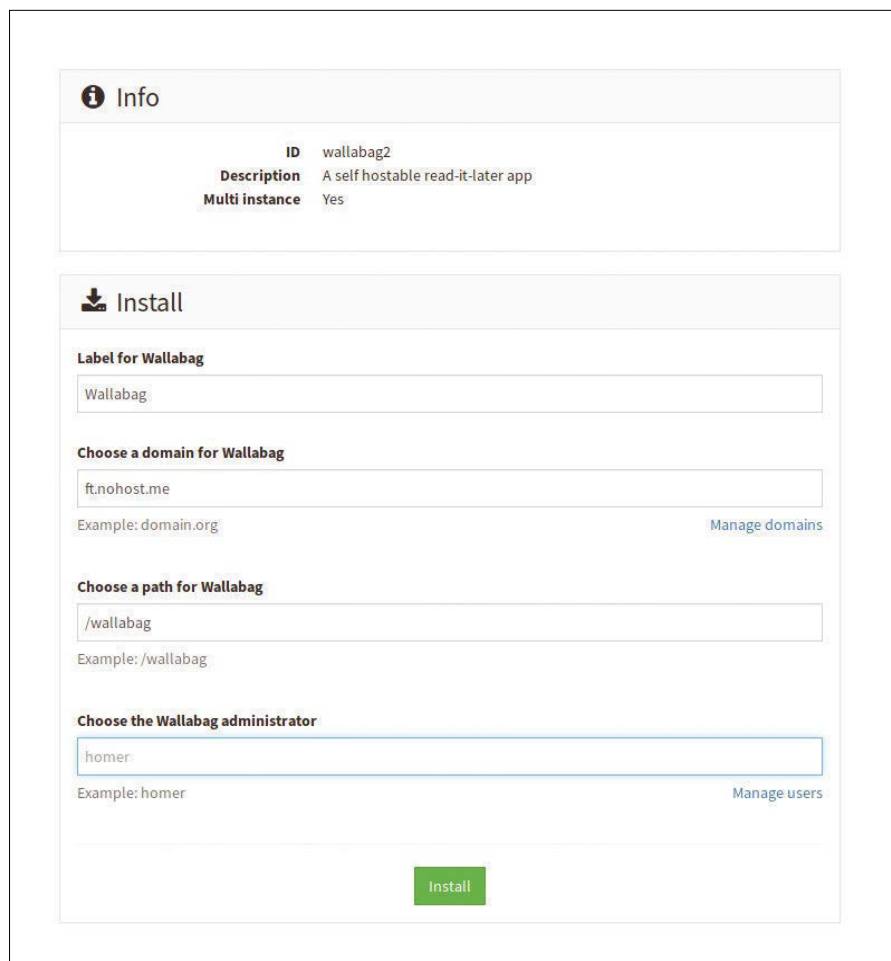


Figure 4: Wallabag was chosen for the test; the package comes from the company's official repository.

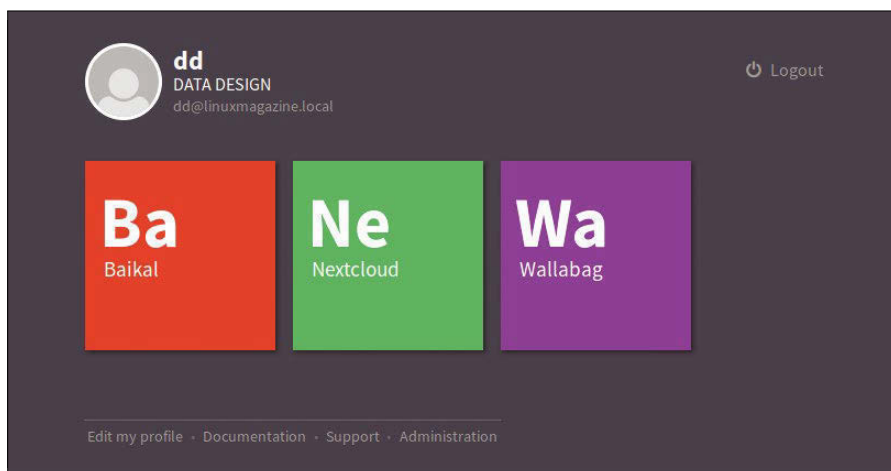


Figure 5: You can log in to all apps using single sign-on.

Full of Apps

The easy installation steps only apply to the company's official apps. For the far larger number of custom apps, which come from and are managed by the community, you'll find an *Install user-defined app* submenu at the end of the official apps list.

Enter the application's GitHub URL. You can find the GitHub URL with the information provided for the app. For example, I installed the Piwigo photo gallery, which had the URL https://github.com/YunoHost-Apps/piwigo_ynh (Figure 6).

In the *Domains* category, it is possible to define additional domains, change the standard, and add certificates. In the test, I created a subdomain of my own domain as the second domain, which worked fine after DNS was set to the IP address defined in `/etc/hosts` (or alternatively on the router).

Tools

The *Services* category allows you to start and stop system services. Behind the *Tools* tab, you will find the settings for the firewall, as well as tools for diagnosis and system messages. You can also change the root password.

The *Backup* category lets you back up the system, user data, and applications (Figure 7). This function is still considered experimental, but it worked perfectly during the test, including restoring the data (Figure 8).

Access via SSH

You can check the results for yourself using SSH. The complete configuration of the server is under `/home/yunohost.conf/`, and the backup is located under `/home/`

`yunohost.backup`. The system offers many more commands, options, and parameters via the command line than through the web interface. You can view these command-line options with `yunohost -h`.

In the initial configuration, administration via SSH is only possible as root, since the system manages the users via LDAP. A script and some other tricks help users reach the system via SSH [20].

In any case, I recommend that you protect the security of SSH access with a key instead of the root password and then generally prohibit the use of passwords for SSH. As soon as user access is working, you should also prohibit logging in as root in the `/etc/sshd.conf` file.

Conclusions

Opinions differ on self-hosting systems such as YunoHost: Some people think that such systems are created by lazy admins who are not familiar with server administration and thus endanger themselves and the Internet. On the other hand, many users need this kind of system to manage data on the network without external help.

The present system makes it much easier to set up a well-provisioned server, but responsible use still requires some knowledge, including securing SSH access, configuring DNS settings, and configuring a certificate with Let's Encrypt.

On a Rasp Pi, the system performs excellently. Well-considered details, such

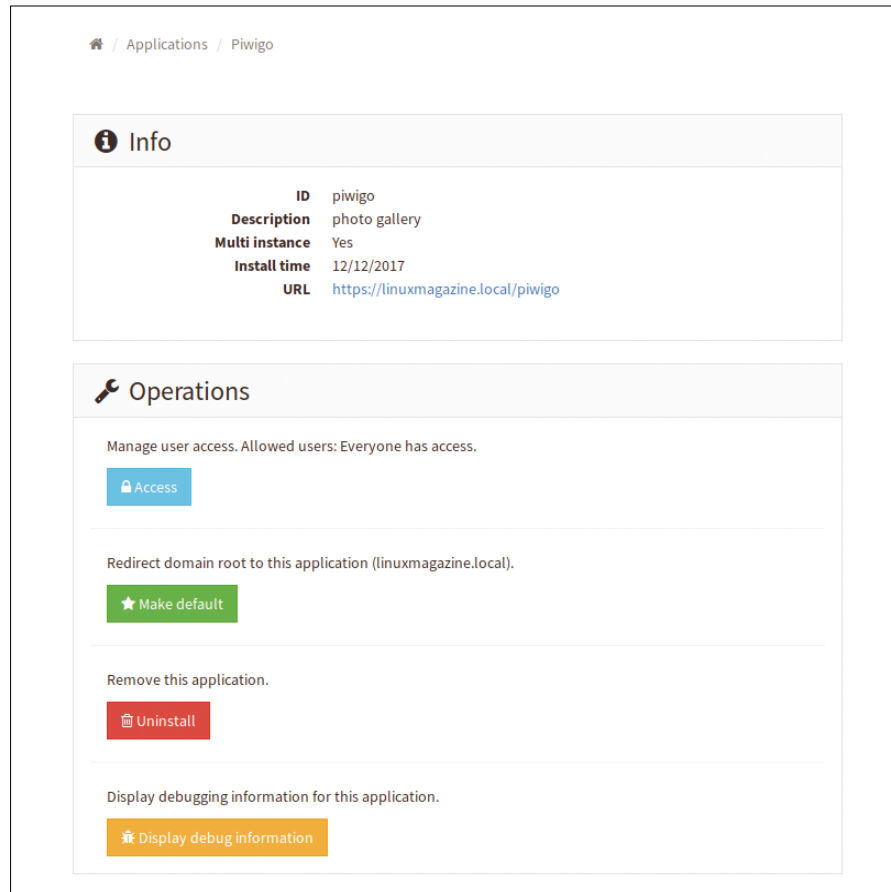


Figure 6: If the packages from the official repository do not meet all your needs, you can turn to community software.

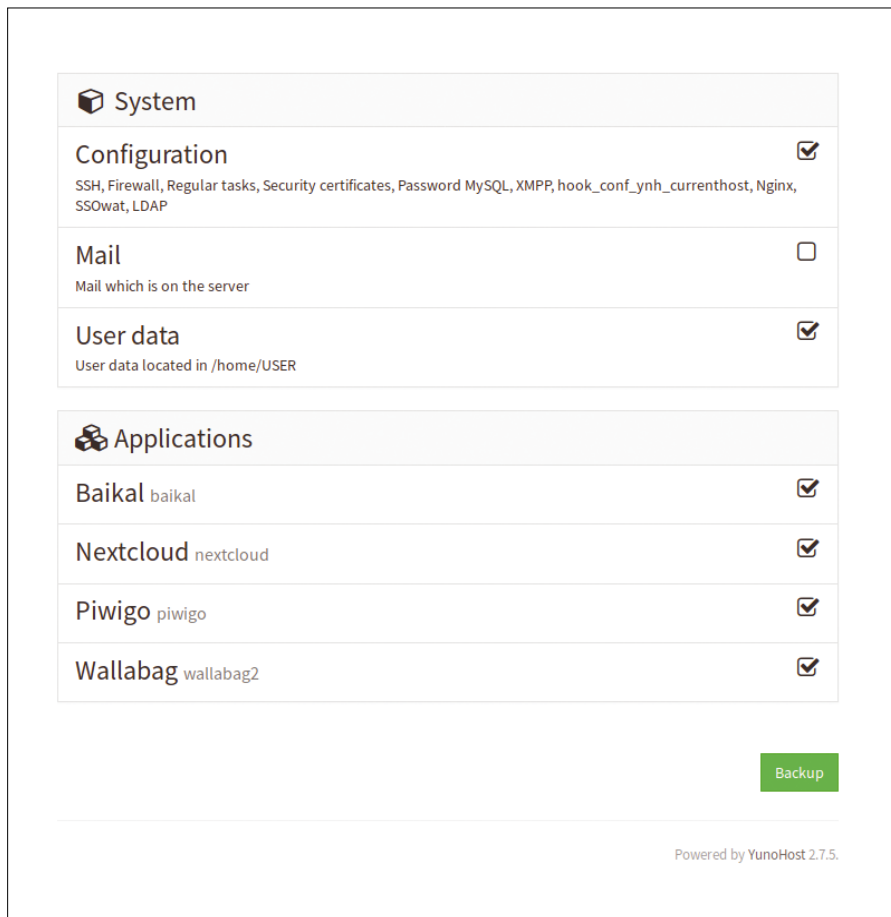


Figure 7: Specify which data you want to include in the backup.

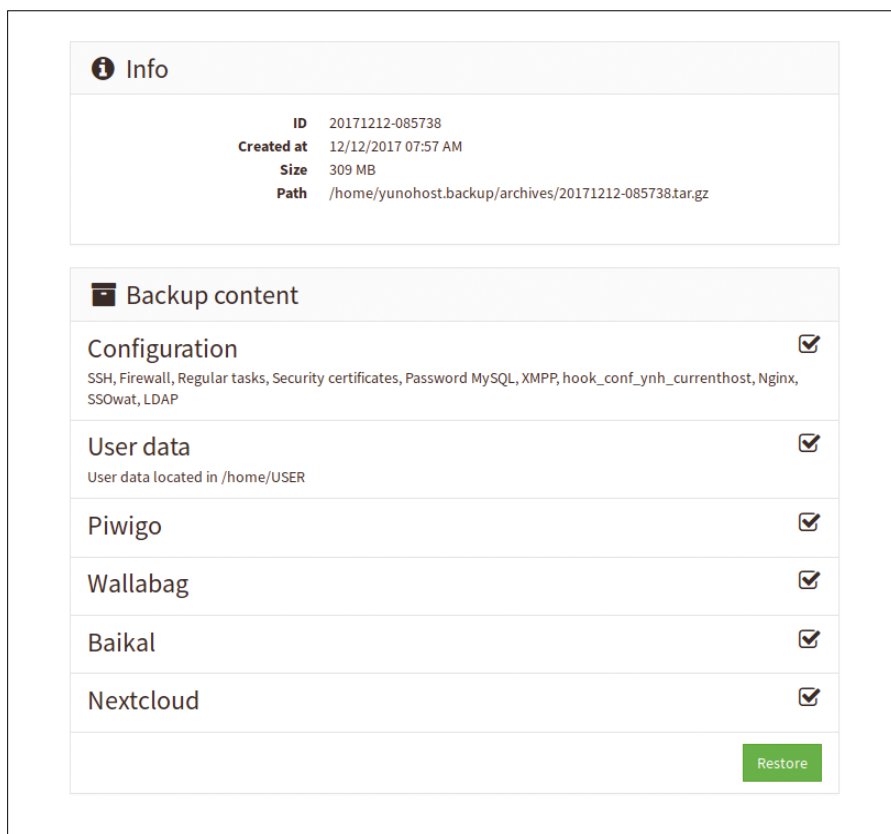


Figure 8: The test succeeded in importing the data from the backup back into the system.

as a preconfigured Fail2Ban [21] to protect against intrusion or user administration via LDAP, are surprising positives. The large selection of apps covers all areas of web services.

YunoHost is suitable for home offices and small businesses. The distribution, which is very solidly set up with Debian as a base, might also work for larger assignments. The fact that YunoHost is free software and under active development means that it could have a very bright future. ■■■

Info

- [1] Download: <https://yunohost.org/>
- [2] FreedomBox: <https://en.wikipedia.org/wiki/FreedomBox>
- [3] Sandstorm: <https://sandstorm.io>
- [4] UBOS: <https://ubos.net/>
- [5] Demo: <https://yunohost.org/#/try>
- [6] Standard apps: <https://yunohost.org/#/apps>
- [7] Baikal: <http://sabre.io/baikal/>
- [8] RainLoop: <https://www.rainloop.net/>
- [9] ZeroBin: <http://www.linux-magazine.com/Online/Blogs/Productivity-Sauce/Host-Your-Own-Secure-Pastebin-with-ZeroBin>
- [10] Community apps: <https://yunohost.org/#/apps>
- [11] Firefox sync server: https://github.com/YunoHost-Apps/ffsync_ynh
- [12] Gogs: https://github.com/YunoHost-Apps/gogs_ynh
- [13] Jenkins: https://github.com/YunoHost-Apps/jenkins_ynh
- [14] Mattermost: https://github.com/YunoHost-Apps/mattermost_ynh
- [15] Piwigo: https://github.com/YunoHost-Apps/piwigo_ynh
- [16] Manual installation: https://yunohost.org/#/install_manually
- [17] Vagrant box: <https://github.com/YunoHost/Vagrantfile>
- [18] Advanced IP Scanner: www.advanced-ip-scanner.com/
- [19] Documentation: <https://yunohost.org/#/docs>
- [20] SSH as user: <https://forum.yunohost.org/t/ssh-disconnects-after-successful-login/256>
- [21] Fail2Ban: <https://en.wikipedia.org/wiki/Fail2ban>

COMPLETE YOUR LIBRARY

Order a digital archive bundle and save at least **50% off** the digisub rate!



ORDER YOURS TODAY!
shop.linuxnewmedia.com



You get an **entire year of your favorite magazines** in PDF format that you can access at any time from any device!



What's new in Fedora 27

New Hat

Fedora Linux is a favorite of developers, sys admins, and other tech-savvy users. The new release sports a new Gnome and other useful enhancements. *By Swapnil Bhartiya*

The Red-Hat-sponsored Fedora project recently released Fedora 27, and I couldn't wait to take a look. Fedora Linux holds a very special place in my heart. In fact, Fedora was the distro that led me to the desktop Linux world. And I'm not the only one who holds Fedora in high regard: Linux creator Linus Torvalds uses it as his primary distro. Linus once told me in an interview, "One of the reasons I like Fedora is they tend to be fairly good about new kernels." He also said he appreciated the efforts the Fedora community, especially Red Hat, put towards kernel development. "They do have lot of kernel engineers. So one of the reasons I ended up going with Fedora is just that they do a good job on the side I cared about."

Fedora Linux comes in three editions: Workstation, Server, and Cloud. This article will focus on the Workstation edition, but it is important to know that, in the case of Fedora, "workstation" does not mean "end user." Fedora's target audience is people like Linus Torvalds – developers and power users who are looking for early versions of the latest Linux updates. Or, as Fedora Project Manager Matthew Miller once explained to me,

"Fedora Workstation appeals to our traditional default user base in many ways, but will have a greater focus on software developers and content creators – and, as the name implies, on people who don't want to do all of their computing from their phone." According to Miller, the Fedora project has always thought of its target user as someone who could also be a potential contributor.

I tested Fedora 27 on multiple machines – Dell XPS 13 Developer Edition,

Dell Precision 5720 AIO, a custom built PC with Intel i7 Core 4780k and GTX 1070 Ti, VirtualBox, VMWare Fusion, and Parallels Desktop. All of my systems have a minimum of 16GB and maximum of 64GB of RAM, so I have really not tested it on an underpowered system. Installation was flawless on each of the bare metal systems and virtual machines.

I do have some mixed feelings about Anaconda, the Fedora installer. It's not

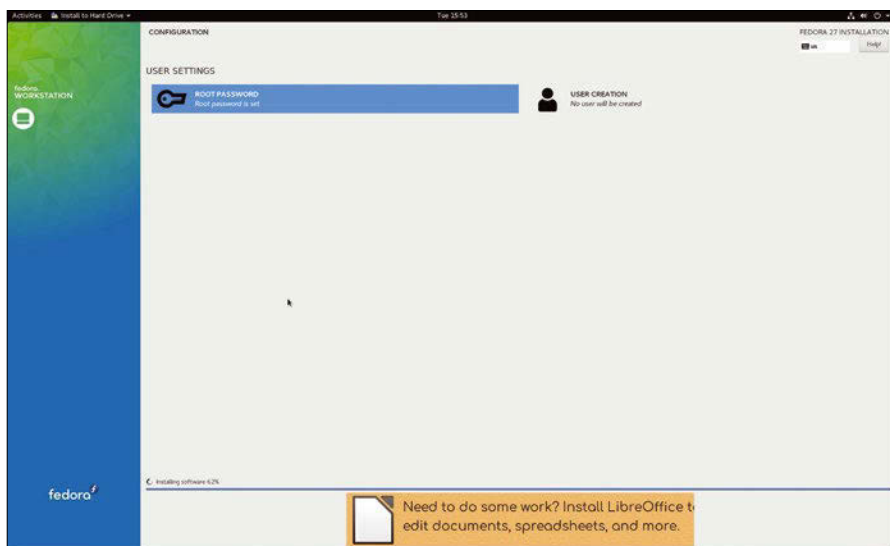


Figure 1: The installer offers no clear instructions about creating a user.

Lead Image: Craig Whitehead, Unsplash.com

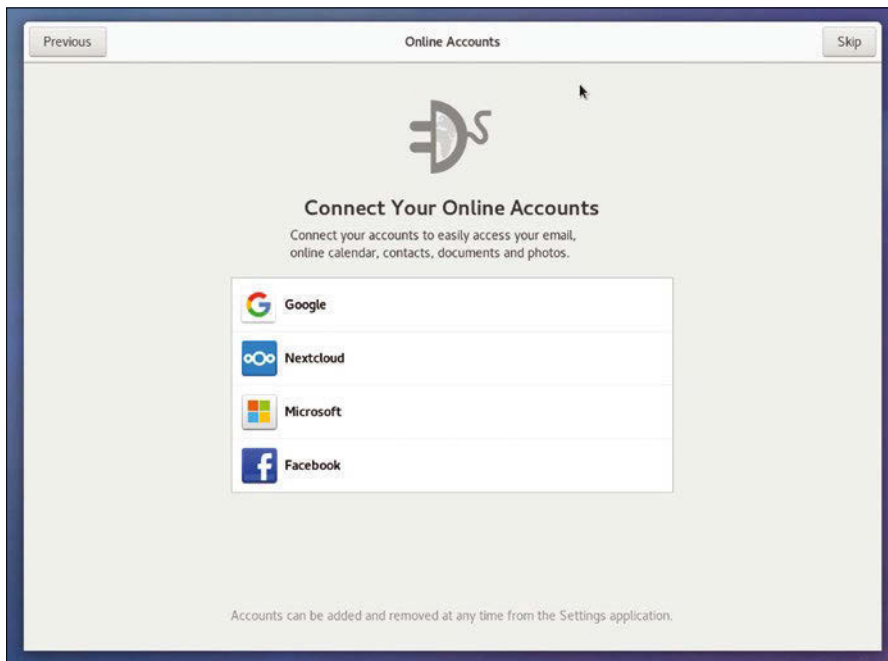


Figure 2: You can easily configure online services on Fedora 27.

very intuitive and could be confusing to a new user. For instance, the option for creating a system user is easy to miss (Figure 1). The inconsistency in the button placement, and the instructions and warning display in an almost unreadable font color, fall short of expectations for a cutting edge distro.

First Impressions

Once you manage to install Fedora successfully, you will see a welcome screen that offers to configure keyboard layout and fine-tune privacy settings. You can also disable location tracking, disable automatic problem reporting, and enable access to non-free software.

The wizard offers Gnome integration with online services, including Google, Nextcloud, Microsoft, and Facebook (Figure 2). If you forget to create a system user during the installation process, you will also have to create a user.

Fedora is known as a Gnome distribution, although alternative Fedora “spins” versions are available for other desktops. Red Hat is the leading contributor to the Gnome project, and Gnome is the default desktop environment for Red Hat Enterprise Linux, so the investment in Gnome is quite evident. The close connection between Red Hat and Gnome means that Fedora users get early access to many of Gnome’s new features and functionalities. For example, in previous years, Fedora users got a early look at

the Gnome Software package tool, as well as Google Drive integration for Gnome’s Files file manager.

Fedora 27 includes the latest stable release of Gnome, version 3.26, which brings some changes to the UI and many improvements under the hood. As someone who uses 4K or HiDPI monitors, I like the options available through Fedora/Gnome. You can choose tiny fonts and icons with massive screen real estate (akin to four 1080p monitors on the same screen) or scale the full screen down to a boring 1080p.

Gnome has done some impressive work for those who use multiple displays. Fedora now supports fine-grained configuration for all connected monitors using the redesigned display settings. A preview of the content showing on the connected monitors lets you set resolution or scaling for best experience. Some of the advanced display features require the Wayland display server protocol, which is now available as an option on the login screen along with the aging X11. The only downside of using Wayland is that it may give you a hard time if you are using any dedicated GPU cards, but Fedora is not known for out-of-the-box GPU support. (My Fedora systems never survive updates if I use non-free drivers for NVidia cards, so I resort to using integrated graphics for my Fedora machines.)

The system settings utility, which is simply called Settings (Figure 3), has a brand new user interface that replaces the good old grid style. I like the refreshing look, which cuts down on clicks, but I struggled to find the monitor settings, which were buried in the *devices* settings.

That’s when I noticed an improved search in Gnome (Figure 4). Not only does it show more items as a list, it can also search system functions like shutdown, logout, etc. My only gripe with Gnome search is the lack of content search, where I can search for an item like SSN and it shows a preview of the

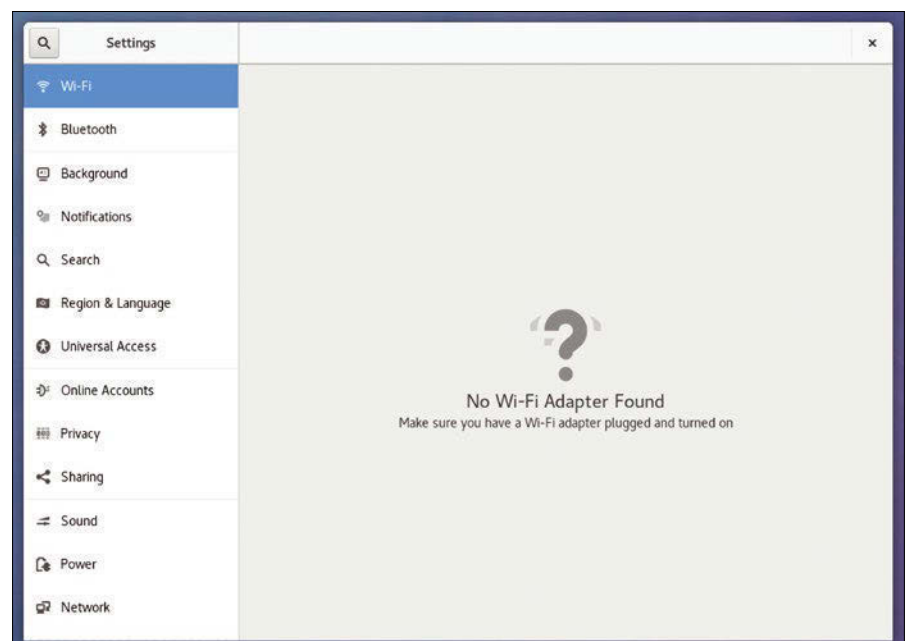


Figure 3: The Settings utility has a new user interface.

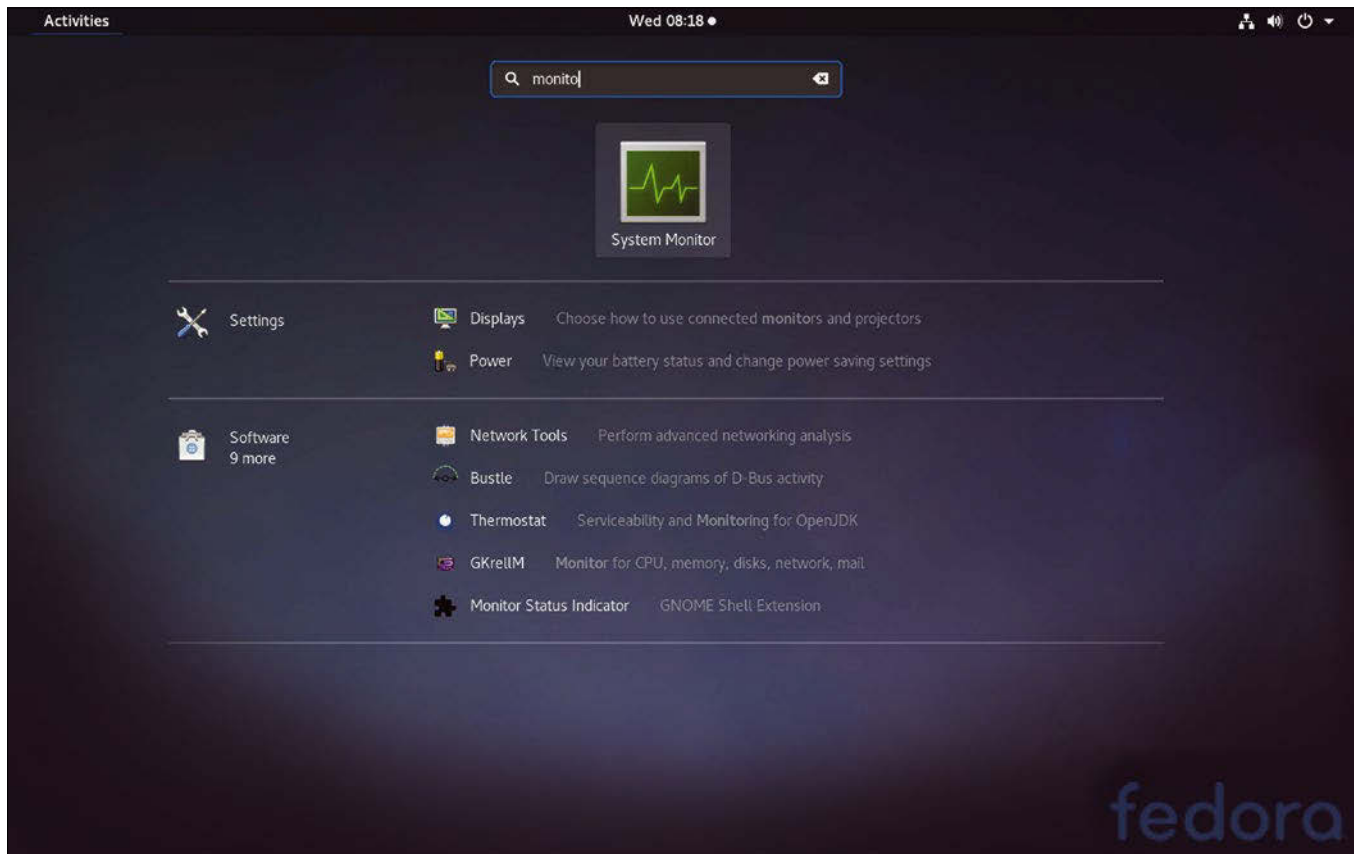


Figure 4: The search feature is much improved.

document that has my SSN. (I like the way Spotlight does this on Mac OS.)

Speaking of Mac OS, Fedora 27 brings emoji support to desktop Linux (Figure 5). Users can now embed color emojis in their documents, but, unlike Mac OS, Gnome does not provide dedicated keyboard shortcuts. You have to open the new Characters application and then drag and drop any emoji into your document. The good news is that developers are working on an emoji picker that can be integrated with other applications.

The default Gnome web browser of Gnome, called Web, now lets you sync your passwords, bookmarks, and browsing history using Firefox Sync, which is integrated with the desktop. No worries about forgetting your passwords. I wish there was a system-wide sync service so that once I log into my user account on Fedora, my emails, contacts, calendars, passwords, etc. are automatically synced – that way I don't have to reconfigure them every time I upgrade to a new Fedora release. On the positive side, Gnome's calendar, mail, and contacts have partial support for offline work; the changes get synced when you get connected.

The Gnome Tweak Tool, which is simply called Tweaks, has an interesting change. Tweaks is designed to let the user customize the desktop with a few advanced configuration options. Tweaks now has an option that lets you move the windows buttons to the left. I wonder if they are doing this for Ubuntu users who are switching back to Gnome

after Ubuntu abandoned the Unity desktop, which placed the buttons on the left side. You can also choose to display the battery percentage for laptops and disable the touchpad when typing. This last feature is quite handy, as I get annoyed on my Dell XPS 13 when my cursor hops from one place to another while I am typing.

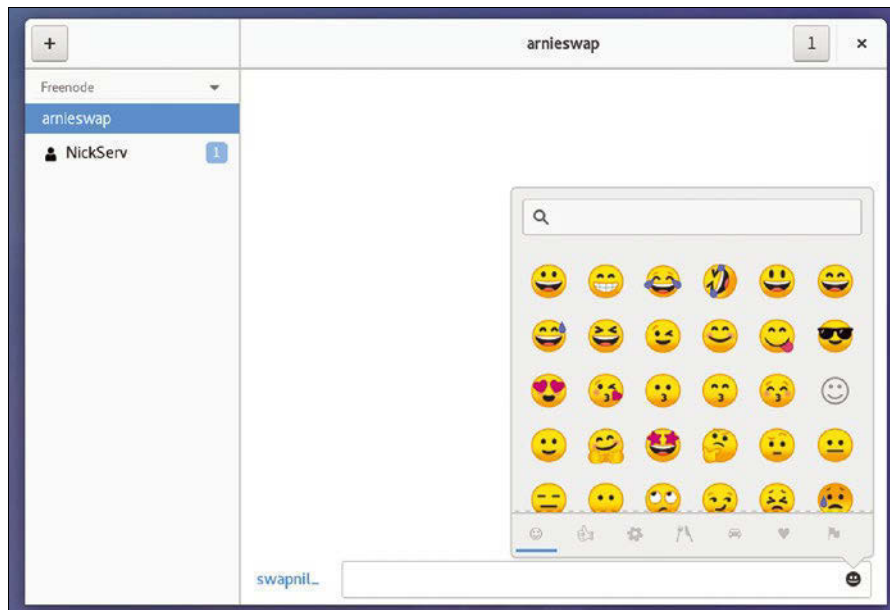


Figure 5: Choosing emojis in Fedora.

External Hardware Support

Fedora automatically detected my networked Brother laser printer, and it was ready for printing without the need for additional drivers. Gone are the days of desktop Linux users struggling with printer drivers. My wireless Apple keyboard and touchpad also worked without any issues; the only minor issue was function key mapping on the keyboard, but I am used to that by now. I had to install `exfat` drivers to access my Panasonic GH5 camera storage, which required more extra work than expected. I had to manually add `rpmfusion` repositories in order to get `exfat`. Unfortunately, Gnome software was not able to find `fuse-exfat`, so I had to resort to the good old CLI:

```
$ sudo dnf install fuse-exfat
```

I mounted the camera and opened the `.RW2` raw images that even Windows 10 can't open. Gnome has come a long way with its support for high resolution previews of raw images.

I also was able to play `.mov` files from the camera. And the most refreshing improvement was out of the box support for MP3, since the patents on MP3 have expired. Red Hat legal gave Fedora permission to ship MP3 encoding in Fedora, so users can now enjoy MP3-encoded songs without any issues.

I can easily mount my iPad Pro, iPhone X, Pixel 2, and Galaxy Note 8 to transfer files.

A Bit about Developers and Sys Admins

Fedora 27 also comes with some significant improvements for developers. The `debuginfo` metadata packages have been split up into smaller sub-packages, making it possible to install just the `debuginfo` for one specific sub-package or library. The new Fedora also comes with version 2.26 of `glibc`, which offers many new features, including a per-thread `malloc` cache that significantly improves the `malloc` API family of functions on certain multithreaded workloads. `Unicode 10.0.0` includes improved support for character encodings, character type info, and transliteration tables. Fedora 27 also ships with `Node.js 8`.

Sys admins may be aware of Fedora's effort to achieve *modularity*. Modularity makes the life cycle of applications independent of the platform and other applications and, as a result, makes the system less likely to break with an application or system upgrade.

Fedora 27 also brings improved support for SSDs. If you create new partitions using `LUKS` encryption, they will automatically be configured to use `TRIM`, which increases the lifetime of SSDs without affecting encrypted data security.

The new Fedora also offers improved Kerberos support, including the new Kerberos Cache Manager (KCM), which, according to the developers, is "better suited for a containerized environment." `OpenSSH Server` now adheres to

system-wide crypto policies, making life easier for sys admins. The developers have removed support for insecure `SSH-1` protocol and upgraded `Samba` to version 4.7. Fedora 27 is the first Fedora version to include `Samba AD` domain controller support.


For developers or sys admins who work on ARM systems, the good news is that Fedora 27 includes disk images for 64-bit ARM devices like `Pine64` and `Raspberry Pi 3`.

Conclusion

Fedora has taken some flack through the years from users who do not understand its primary purpose. Fedora Linux is a cutting edge distribution that's designed for developers and sys admins who want to keep an eye on new technologies coming into the enterprise space. It's not meant to play `Steam` games with the latest `Nvidia` cards; you can do that with Fedora, but that's not its target market. Most of Fedora's new features are geared towards developers and sys admins, because that is Fedora's target audience.

If you're strictly a desktop user, however, you might still have an interest in Fedora because of its tight integration with `Gnome`; Fedora remains the distribution that introduces the latest and greatest features from the `Gnome` world.

Fedora has come a long way and is a distro that is well suited for a wide range of use cases. I use it for my work – and so does Linus. ■■■



THE ULTIMATE
raspberry pi
GEEK
ARCHIVE

GET 22 ISSUES OF RASPBERRY PI GEEK
ON ONE FULLY-SEARCHABLE DVD!



Order Now!
shop.linuxnewmedia.com

To-do managers regulate appointments and tasks

Organizational Talent



Busy people often keep busy calendars full of appointments and tasks. In order to keep an eye on things, Linux to-do-managers help manage the clutter in a controlled way.

By Erik Bärwaldt

Yellow sticky notes on your desk are becoming a thing of the past. Word has spread that computers manage appointments and projects far more flexibly and reliably. Instead of relying on a jumble of little notes, you can refer to a carefully managed digital to-do list. However, time-management applications vary considerably, so we decide to take a closer look at BasKet [1], Getting Things Gnome (GTG) [2], Makagiga [3], RedNotebook [4], and Task Coach [5].

Basic Information

Large project management systems are often based on a client-server architecture. Small Getting Things Done (GTD) managers usually only run on the desktop. The option to map out larger projects and integrate external resources is usually missing. However, to-do managers also need to manage different task groups, some of which consist of individual tasks.

Time management applications also need to make it possible to integrate external sources, such as documents that you need to complete a task. An easy-to-understand interface and the option to create backups are important features for any viable task management solution. I have picked five candidates out of the very extensive pool of GTD software (see the “Not Considered” box) and tested them for practicality.

BasKet

The BasKet [1] task management system, optimized for the KDE desktop, is included in the software repositories of most major Linux distributions. The application uses baskets to manage tasks and appointments; the program primarily acts as a digital notepad or index card box. When first launched, the main window appears somewhat cluttered due to the short descriptions in the large display segment on the window’s right side.

On the window’s left side, there are already several baskets with sample data,

Not Considered

Under Linux, there are additional task managers available. I have left out some of these – even quite well-known ones – due to lack of updates or free licenses. Tasque [6] has long been one of Gnome’s standard applications. However, apart from a small improvement in Romanian localization, there has been no further software development for years. Tracks [7], a web-based GTD manager with a built-in web server, was last updated more than two years ago.

ThinkingRock [8], a very extensive Java program, doesn’t offer a free license for the current 3.X versions, and there is only one commercial version.

The iKog [9] command-line task manager consists of a Python script. The last files offered for download are from 2011. My guess is that this software is no longer maintained.

Lead Image © Javarman, 123RF.com

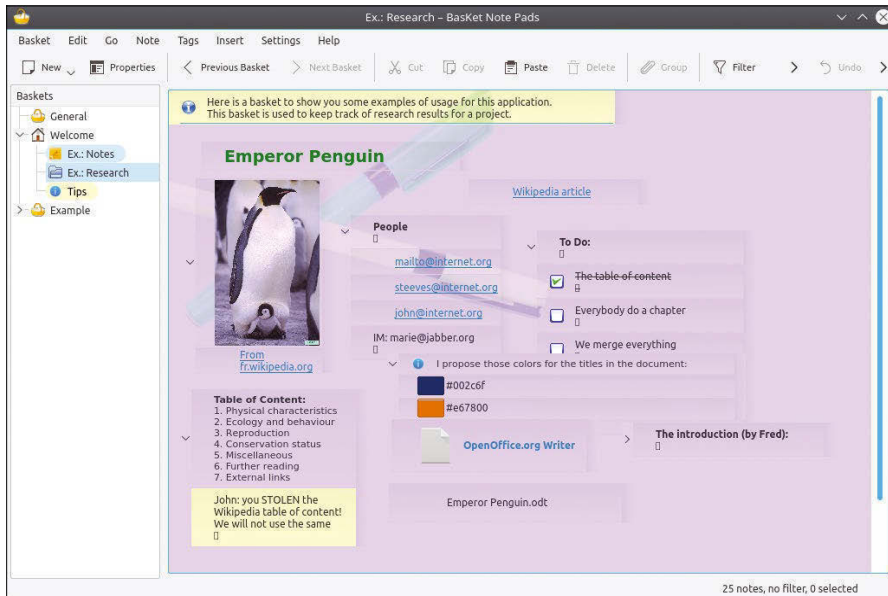


Figure 1: Basket can be very flexible with different source documents.

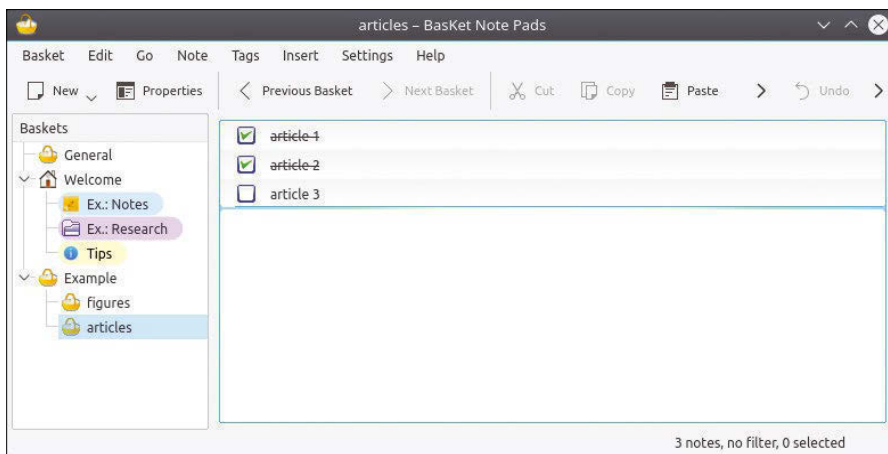


Figure 2: Completed tasks can be checked off by clicking; they are then shown scored through.

which show that Basket integrates unstructured data very flexibly into the baskets: images as well as links (Figure 1). Users can edit images with third-party programs. The individual elements form a frame in the display area and can be grouped freely.

For a to-do list, define a new basket in the left window area by clicking on *New* (far left). Then in the dialog box, assign a name and define the layout. A single-column display in the right-hand pane is suitable for to-do lists. If the job consists of several tasks, you can click on the basket and define subtasks via *New | New Sub-Basket*. In the right pane, enter the individual data by clicking in the free area and selecting the *Text* option from the context menu. For example, Basket records the appointment including the corresponding activities.

Once you have created an appointment, click on *Tags* in the menubar and choose *To do!!!*. A small box with a blue border appears to the left of the appointment entry; you can click on this as soon as you have completed the task (Figure 2).

Basket does not offer automatic reminders or automatic deletion of completed tasks. You therefore have to keep an eye on the baskets. In order to track the progress of task elements, click on the small triangle to the left of the task description and on the *Progress* option in the context menu. An empty progress bar appears in the line, which is filled in 25 percent increments by clicking on it. This shows at a glance how much of the task has been completed. You can assign priorities to tasks using the same approach.

If several users work on one computer system, you can password protect baskets to prevent other users from accessing them. To assign a password to a basket, use the dialog in the *Basket | Password* menu, which also provides password protection for subordinate task lists.

Basket also allows you to back up and restore your personal data via the *Basket | Backup & Restore* menu option. Select the appropriate directories in a clearly arranged window according to your wishes and then create a basket (Figure 3).

In order to use the existing data for other purposes or to import into Basket from third-party applications, the program offers some format converters. The data import can be called up via the *Basket | Import* menu option; the converter handles Tomboy [10], KNotes [11], Tux-Cards [12], and Sticky Notes [13] formats. The export possibilities are less extensive: The software only exports personal datasets to simple HTML or the program's own archive format.

GTG

The GTG [2] to-do application was originally developed for the Gnome desktop, but it also runs in other work environments. It is included in most software repositories and can therefore be easily installed using the package manager. The software is partly written in Python and partly in JavaScript. GTG also has an add-on for Firefox and Thunderbird, which enables additional scheduling in a web browser. Unfortunately, there are no ports for other platforms.

GTG starts with an inconspicuous list view that presents tutorials as tasks. A click on one of the entries opens the actual task window, which provides an initial overview of the GTG functions (Figure 4).

The actual program window consists of a menubar with four buttons and an input field, which allows for a quick se-



Figure 3: Basket also offers to backup your personal data.

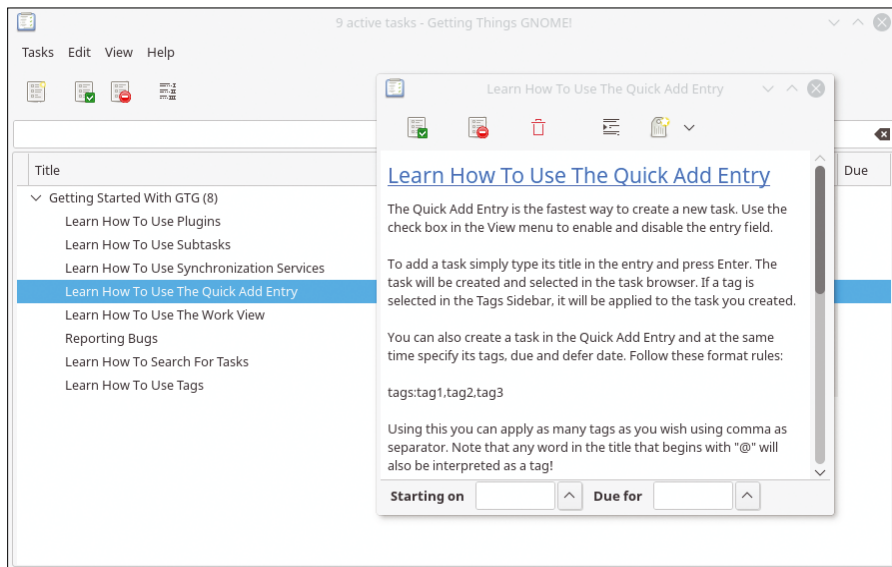


Figure 4: The GTG user interface is very simple.

lection of functions and provides an input line.

To create a new task, select *New Task* in the Tasks menu or click on the *New Task* icon top left. In the subsequent pop-up window, enter all relevant data as free text and select the task's start date by clicking on the drop-down menu to the right of the *Starting on* input field to access a calendar. In the *Due for* field, determine the task's completion date.

After you click on the drop-down menu to the right of the date field, a calendar is available. The entry automatically lands in the main window's list area, it appears with a *Start date* and *Due on date*.

You can display a task with several steps via *Tasks | New Subtask*. The data you enter is indented below the main task. If several subtasks are defined for a primary task, the entire list is opened by clicking on the small minus symbol to the left of the main task and can be removed at any time using the plus symbol. This gives you a better overview of extensive task lists.

Remove tasks from the list by clicking on the *Select as completed* button at the top of the program window in the buttonbar. You can view these tasks again later by selecting *View | Closed Tasks Pane* (Figure 5).

Right-clicking on an entry and selecting the *Delete* option from the context

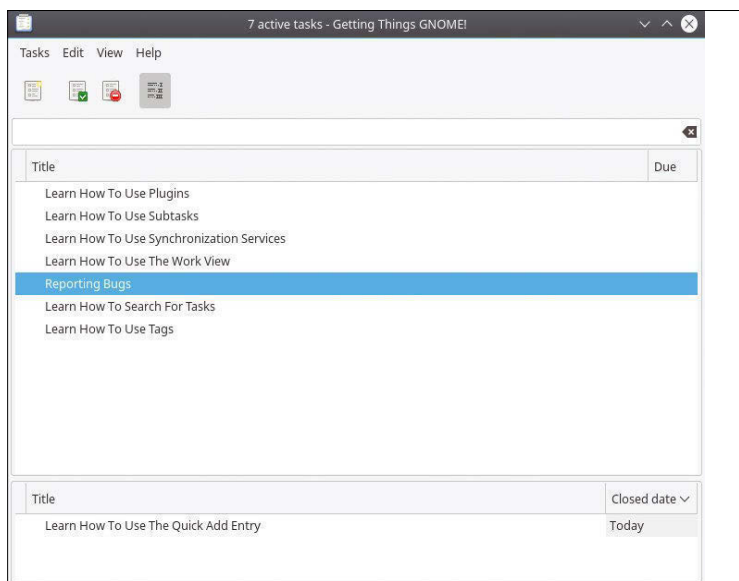


Figure 5: GTG displays completed and pending tasks in the same window.

menu deletes primary and secondary tasks. The second icon from the right side of the buttonbar marks a task as “not to be completed.” It then disappears from the task list and appears at the bottom of the checked tasks list.

GTG synchronizes its data with Gnote [14] or Tomboy [10]. synchronize, select the data in Tomboy or Gnote and then transfer to GTG. You can access the sync function in GTG via *Edit | Synchronization Services*.

GTG has a modular structure and can integrate plugins. The *Edit | Plugins* menu option calls plugins that are already integrated. A separate pop-up window lists existing extensions; you can enable them by checking the boxes to the left of the plugins. After restarting the software, you will find additional entries in the menubar and the buttonbar to match the plugin functions (Figure 6).

Because GTG is a very compact desktop application, the program is not capable of exchanging data with other schedulers or project managers. A freely configurable reminder function with acoustic or optical signals is also missing.

Makagiga

Makagiga [3], a to-do manager for the desktop written in Java, is available under the free Apache 2.0 license. Makagiga comes in three versions for Linux (32- and 64-bit systems, as well as a portable version) and requires a Java Runtime Environment (JRE) on the installation system. It does not

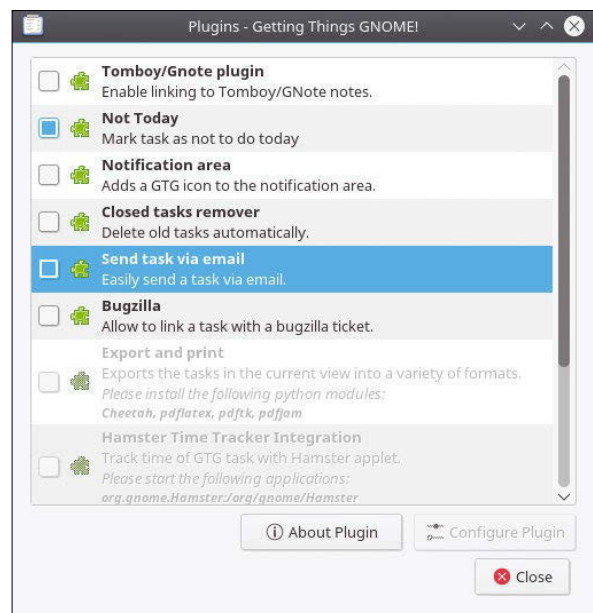


Figure 6: GTG is expandable with plugins.

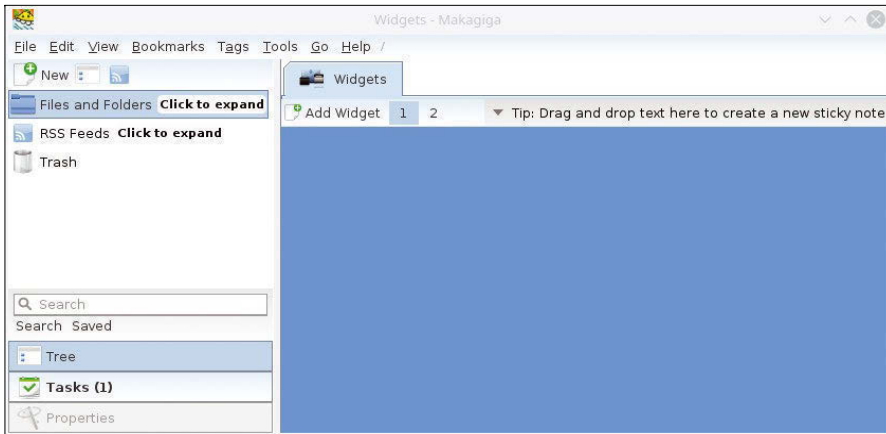


Figure 7: The heavyweight among personal GTD managers: Makagiga.

matter whether you use OpenJDK or Oracle's JRE.

Makagiga's functionality goes far beyond that of a plain vanilla to-do manager: In addition to a pinboard and an appointment calendar with corresponding to-do lists, the software also provides an RSS reader and can be extended with various plugins. The range of plugins extends from simple text-viewer modules for certain formats to a wiki tool for viewing and editing wiki texts, a graphics program, and an electronic topographic map based on OpenStreetMap [15]. You can download plugins from the repository and integrate them into the task folders in Makagiga. Automatic integration with the uniform interface means that there is no need for time-consuming installation work.

The software is available as a 33MB 7z archive, which you unpack and drop into a directory. Then call Makagiga via `java -jar makagiga.jar`. Since the program directory also comes with icons and a small script called `makagiga.sh`, which creates a starter, it integrates into a menu structure without any problems. After the call, Makagiga starts surprisingly quickly, showing a program window with an empty pinboard (Figure 7).

When you create a new folder or tab via *New* (top left), it appears on the right in the pinboard, where Makagiga has already preconfigured a to-do list and widgets. You can stick your notepads (widgets) to the pinboard, and add a second page to the bulletin board by clicking on 2 below the tabs. The content of the tabs on the right varies, depending on the folders on the left in the tree view.

When you enter content in your virtual desktop, you do not have to save it separately. When you quit, Makagiga au-

tomatically stores all the content before closing and shows you the same tabs and folders when you restart so that you can continue working seamlessly.

In order to provide topic- or task-specific news quickly and clearly, many websites offer their own RSS feeds. Makagiga provides a built-in RSS reader.

Setting up a new feed is very easy: Copy the feed address to the clipboard and paste it into the Makagiga setup dialog. You can do this via *New | Add RSS Feed*. After inserting the link, click on the *Preview* button at the window's bottom right. Makagiga now loads the latest feed message and displays it in the *Preview* area. If you use the *Create* button bottom right, the link appears on the left side of the program window in the *RSS Feeds* section (Figure 8).

At startup, the software updates the feeds and tells you how many messages it received for each feed since the last

startup. Clicking on the feed shows the titles of the latest messages in a list view.

Clicking on a message opens it on the right in a new tab; you can navigate between the messages with the help of the arrow keys (Figure 9).

Makagiga lets you extend the virtual desktop with plugins [16] (Figure 10). For an overview of the available plugins, click on *Plugins | Tools* in the *Tools* menu. A list of the available plugins now opens on the virtual desktop, some of which are already installed (*Installed*) or permanently implemented (*Built-In*).

Clicking on *Get More Plugins* above the list opens an additional dialog that lists further plugins hosted on SourceForge. To add one of the plugins, click on it and then on the *Install* button in the dialog's lower left corner. A pop-up window shows brief information about the extension; relicking on *Install* downloads the plugin and integrates it into Makagiga.

At the same time, the routine displays a short note on how to call the plugin. You can also remove an existing extension from the plugin list by clicking on the red *X* icon to the right of the plugin and confirming the uninstall in another dialog.

Makagiga implements time management in the form of to-do lists and, a calendar. You can activate both via the *Tools* menu; they then appear as separate tabs on the right of the virtual desktop. The *Add Task*: input line lets you enter tasks in to-do lists and expects a short description of the task. Clicking on

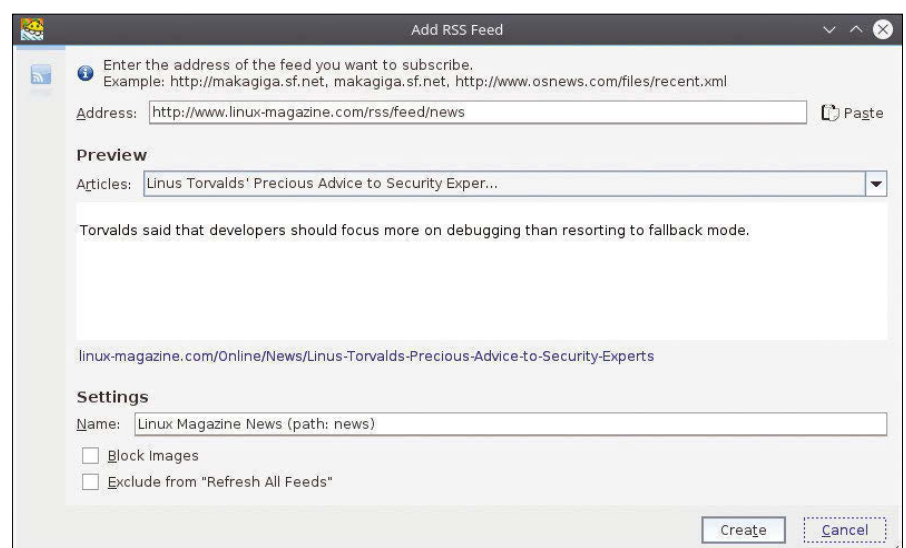


Figure 8: Makagiga's functionality goes beyond the simple to-do list. Among other things, the software integrates RSS feeds.

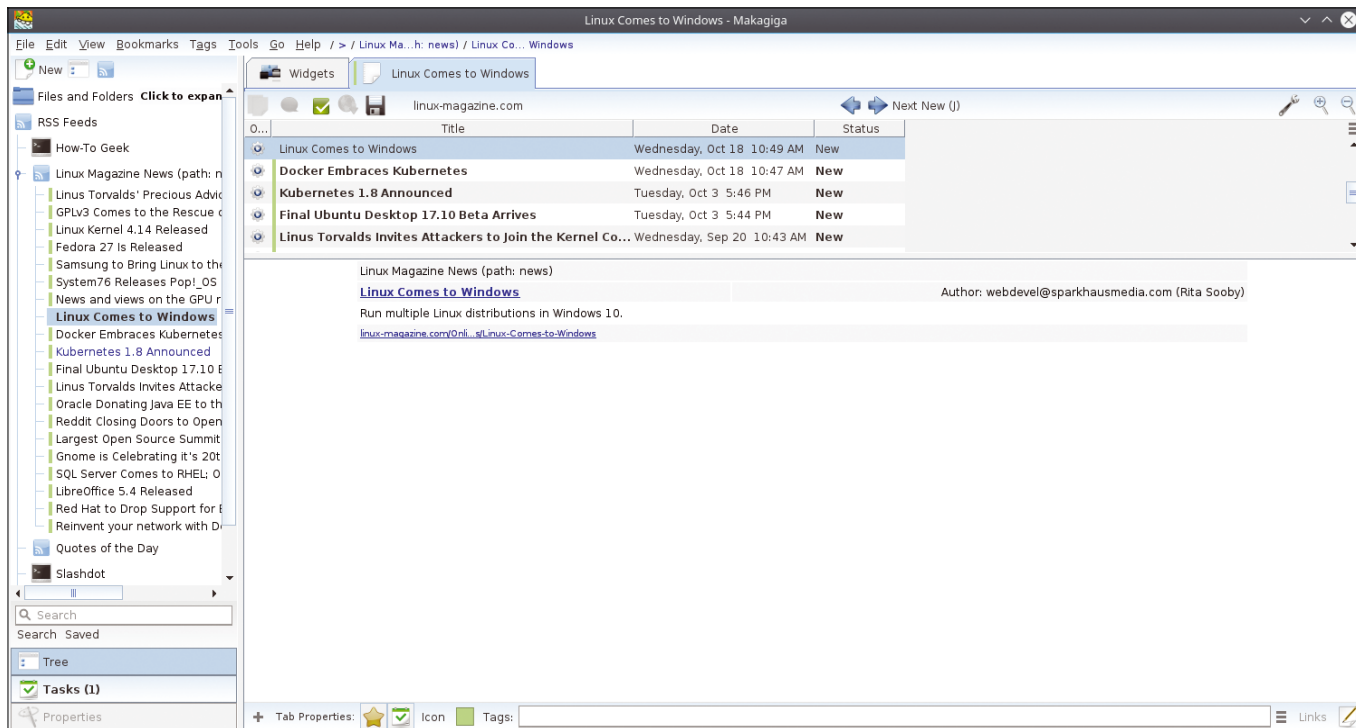


Figure 9: RSS feeds keep you up to date on specific topics.

the + symbol to the right of the line adds the new task to the overview.

You can open the calendar via the *Set Date/Time* link right in the overview line. Select the desired date and time for completing the task. After you click on *OK*, the data is placed in the to-do list. In the *Priority* column, specify the priority with which Makagiga should handle a task. In the *Tools* menu, also define the

alarm functions for the individual tasks in *Settings*. Color gradients are available in the to-do list or messages in the system tray when a time limit is reached.

RedNotebook

Developers have been working continuously on the GPLv2 software RedNotebook [4] for several years now. RedNotebook is available in the software repositories of

most major Linux distributions. The software comes with a clear-cut user interface that requires practically no training. RedNotebook is designed as a journal application and therefore has an extended range of functions compared to plain vanilla task planners: The program also includes a calendar with a small word processor that provides basic formatting options and various export options for the finished texts.

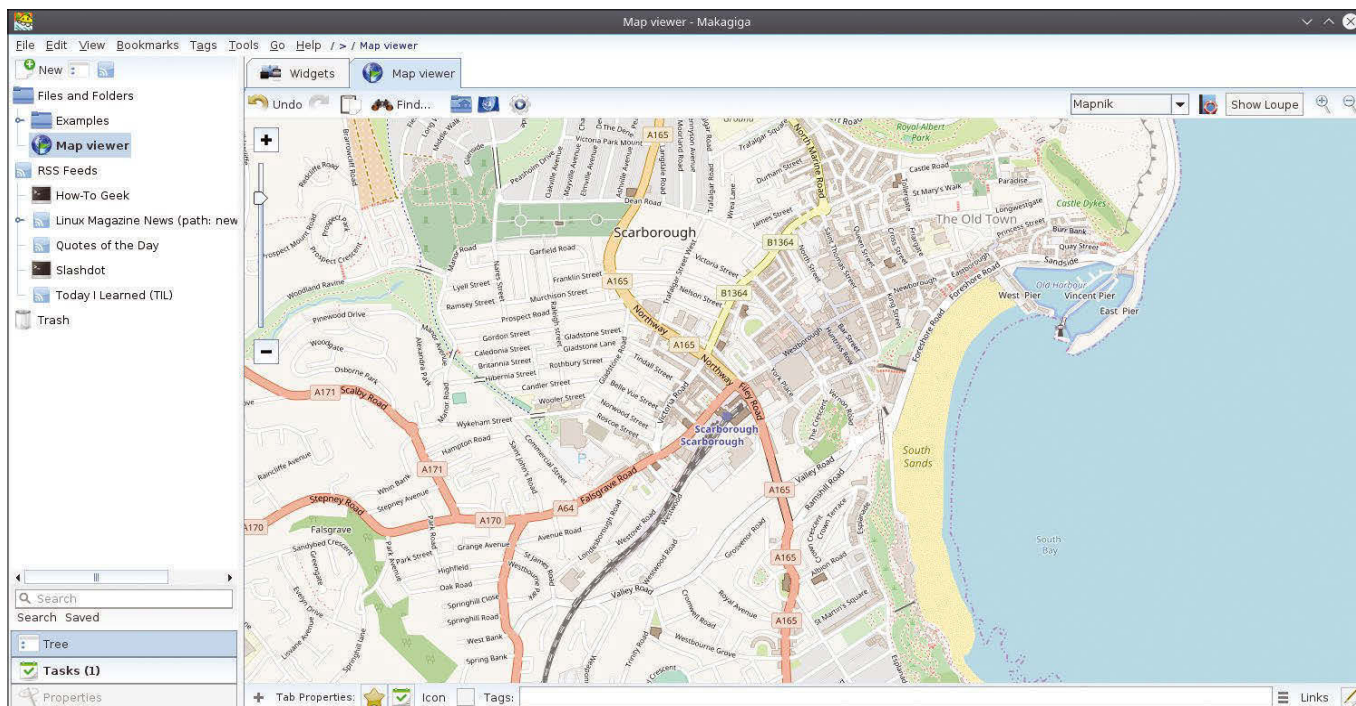


Figure 10: Makagiga even replaces maps with an OpenStreetMap plugin.

After you start the application, you will find yourself in a window divided into four parts. A horizontal menubar appears at the top and a buttonbar is below it. In the program window's left-hand area, there is a vertically arranged calendar and below it a word cloud (Figure 11). To the right is a large section for text input and editing.

The buttonbar controls the calendar and the display area separately. A special feature of RedNotebook is the word cloud: Certain words appear from the text that you have keyworded with the help of hashtags. The words appear in different sizes, depending on how frequently they occur.

If you click on one of the words, the software searches for it in the text body to the right of the picture and then displays it with a yellow background.

Since the software has good online documentation, you can easily familiarize yourself with the two operating modes, *Edit* and *Preview*, which you can select on the buttonbar.

After startup, I recommend that you first delete the sample text in the editing area. To do so, enable the editor in the buttonbar by clicking on the *Edit* button. On first use, changing from the *Edit* to the *Preview* button is somewhat confusing, because the software is actually in edit mode and not in display mode. The previously grayed out *Template*, *Insert*, and *Format* buttons are now enabled, together with options.

When entering text, select important words that you want to index for the word cloud by prepending a hashtag. Your selections appear immediately in red letters in the text to attract attention. When you save the journal and exit the program, the word cloud updates. The keywords appear in the word cloud the next time you start the software.

By the way, you do not necessarily have to use the *Save* dialog via the *Journal | Save* or *Journal | Save as* menu, because RedNotebook automatically saves the newly generated content. The more often keywords appear in the text, the larger they are in the word cloud. The same applies to the selections: RedNotebook displays frequently selected words in a larger typeface (Figure 12).

In order to quickly find entries with certain keywords in a well-filled calendar, use the highlighting function. For

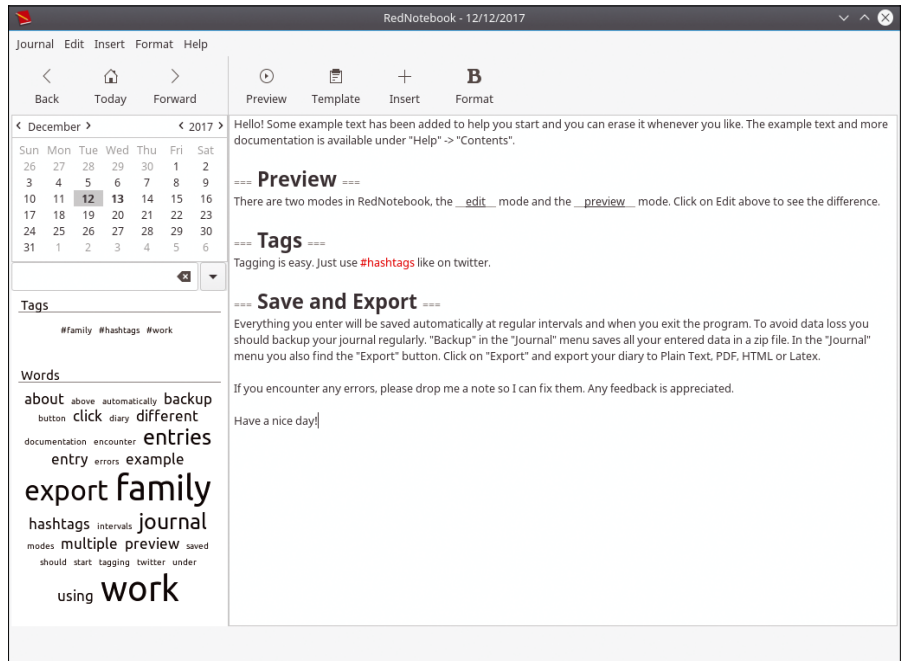


Figure 11: RedNotebook organizes tasks using a word cloud (left).

this purpose, there is an input field on the left of the program window between the calendar and the keyword and selection lists. Type the search term into this input field; the display disappears while typing and a small list view with the results of the search term appears in its place. The list has only two columns with the date of the entry and the location where the search term appears in the text.

Clicking on one of the entries immediately changes the calendar view and shows the month and year of the entry.

RedNotebook also shows the entry on the window's right side. Thus, you do not have to flip through the calendar view with the arrow keys to jump to an entry with the search word.

RedNotebook uses format texts in edit mode; only a few formatting options are appropriate for the program's purpose. To select a formatting option, click on the *Format* button in the program window's upper right corner and select the desired format from the context menu.

You cannot enter fonts or special formatting options, such as tables, indents,

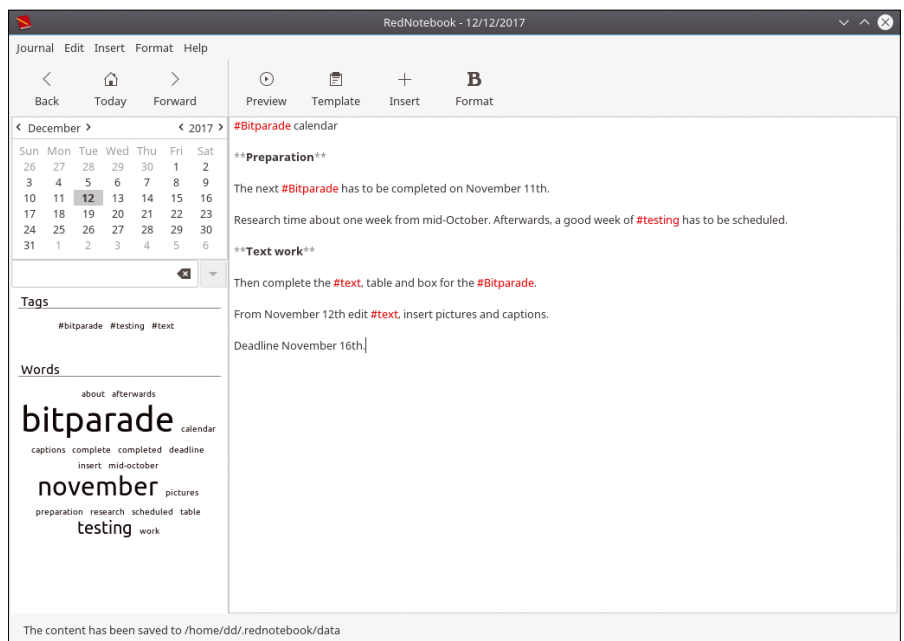


Figure 12: Important words are indexed in RedNotebook.

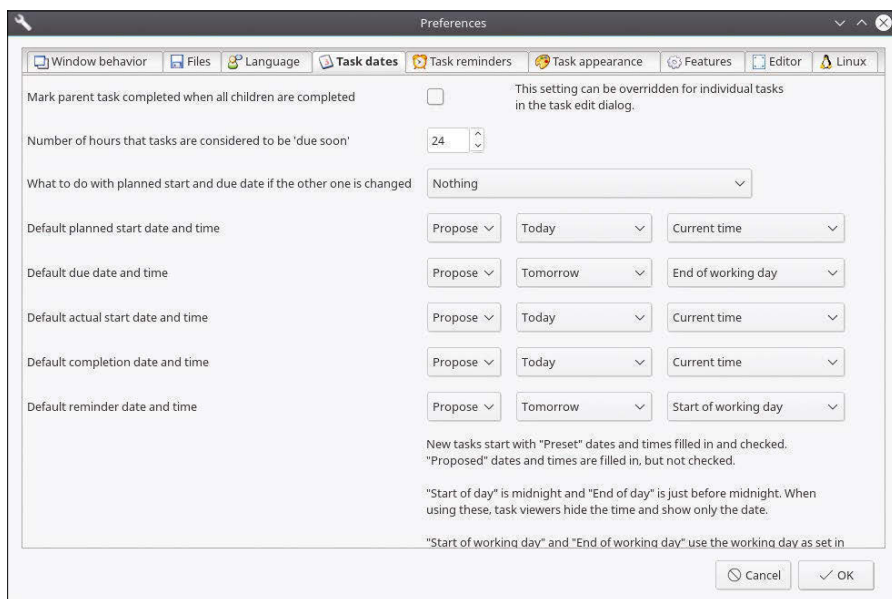


Figure 13: The Task Coach settings dialog is very extensive.

or footnotes, but you can draw lines, insert manual line breaks, or optionally insert images and links via the *Insert* button. The list function lets you create lists with partially indented contents. You can also transfer files to the editor page by first selecting them from a file manager via *Insert | File* and then dropping them into the editor window as direct links.

Finally, in display mode, you can access files using distribution-specific third-party programs, such as image viewers. The *Template* button takes you to some examples that allow you to quickly create and format text. This is where you will find templates for a phone memo, a report, or a travel report. You also can add your own templates.

If you select *Create New Template* from the *Template* menu, a help page listing the template's various formatting op-

tions will open after you enter a name. Once you have designed your template, click on the *Save* button to add it to the list. After restarting the software, your new template will appear in the list.

RedNotebook can export your journal to various file formats and is thus cross-application and even cross-platform capable. In addition to LaTeX and simple text, HTML and PDF are also available as export formats. In a separate export dialog – accessible via *Journal | Export* – the software exports entries from a freely selectable period to the desired target format in five steps.

You can specify filter criteria to exclude certain keywords from the export; you can also decide whether to use plain text or keep the markup in the target format. Finally, choose a target path for the backup. The software creates the target

file after displaying an overview, which includes the formatting, but not the additional headers or footers with page numbers.

RedNotebook also lets you back up the data via *Journal | Backup*. After specifying a target path, the program creates a zip archive where the filename consists of the current date. The zip archive contains plain text documents with special tags. If you want to edit the documents later, you cannot use a conventional text editor.

Task Coach

Task Coach [5] is one of the standard task managers under Linux, thanks to years of development and maintenance. The software appears in the repositories of almost all the major distributions. The latest version and the source code are also available from the project site [5]. The program, released under the GNU GPLv3, offers many functions of a professional task management tool, and learning to use it does not take long.

Before using Task Coach for the first time, it is a good idea to complete the basic configuration in the *Settings* dialog. Select *Edit | Settings*. In the configuration window, the *Task Appointments* and *Task Reminders* tabs are especially relevant. You can define various options for the reminder function and time intervals (Figure 13).

To use Task Coach effectively, I recommend that you first create individual categories. Then assign the tasks to be created at a later date to each category. The dialog for creating a category lets you enter free text as a description and

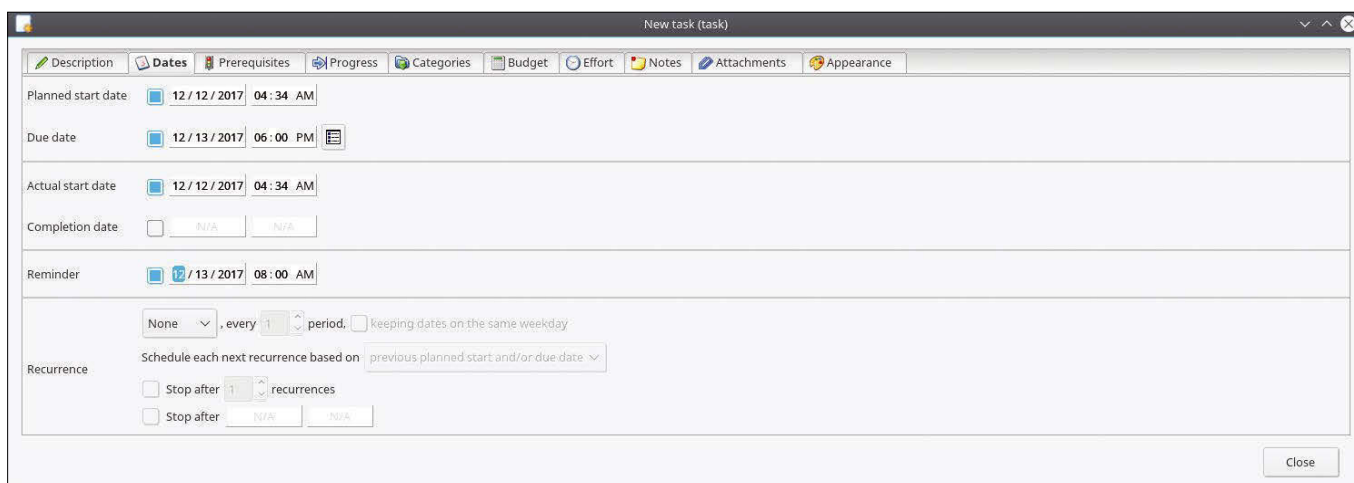


Figure 14: Tasks can be scheduled in great detail in Task Coach.

also link attachments. Once you have created a category, assign the corresponding tasks to it in the next step.

The task dialog is far more complex than the category dialog. Press the *New task* button directly above the task pane in the program window. The interface is similar to the category dialog: In the *Description* tab, assign a name for the task and, if necessary, a description as free text. You can also assign priorities.

In the *Dates* tab, define the start date, the planned due date, and the actual closing date. You can also configure a reminder time.

In the *Prerequisites* tab, you define subtasks that Task Coach lets you enter in a similarly designed dialog. Subtasks need to be processed first as a prerequisite for the main task. The *Progress* tab lets you enable a progress bar, and the *Categories* tab assigns categories to tasks. Task Coach displays the categories that have already been entered.

In the *Budget* tab, enter a working time quota, as well as an hourly rate or a fixed budget. The time-tracking features make Task Coach suitable for freelancers and service providers who implement customer projects and therefore require precise scheduling and invoicing. The *Effort* tab collects

the associated work data. The *Notes* and *Attachments* tabs let you add user-defined texts and integrate files. Clicking on the *Close* button terminates task entry and transfers the results to the main window (Figure 14).

When you launch the program later on, the software automatically checks the individual categories and tasks for deadlines. When it identifies due dates, Task Coach displays a small window that reminds the user of the appointment. Task Coach offers the option of marking this task as completed, to avoid the software reminding you again later. Click on the *Mark task completed* button (Figure 15). In response, Task Coach places a checkmark to the left of the task in the main window and highlights the task in green. Task Coach marks overdue appointments in red if they have not been marked as completed.

The buttonbar with its quick functions, located directly above the two list areas for tasks and categories, also lets you mark a task as inactive. In such cases, the reminder function is automatically deactivated. It is also possible to hide inactive and completed tasks (i.e., remove them from the lists).

If you want to assign new dates to several similar tasks, you can also define an existing task as a template.

Task Coach also lets you import and export data. The software supports CSV or plain text format for importing to-do lists. Task Coach is more flexi-

ble when it comes to exporting: You can choose between HTML, CSV, text, or iCalendar data formats. This ensures good compatibility with other appointment management systems, as numerous applications support the iCalendar format [17].

Conclusions

The to-do managers discussed in this article cover a wide range of tasks (see Table 1 for an overview). If you need a plain old task manager without an automatic reminder function, Basket, RedNotebook, or GTG. The first two stand out because they can also handle unstructured data.

Task Coach, on the other hand, focuses on professional time management and offers a reminder function and warnings to a large number of employees in the case of overlapping appointments. Makagiga serves users who need a complete virtual desktop for highly complex tasks and want to use different third-party sources. ■■■

Info

- [1] Basket: <https://github.com/kelvie/basket>
- [2] GTG: <https://github.com/getting-things-gnome>
- [3] Makagiga: <http://makagiga.sourceforge.net>
- [4] RedNotebook: <http://rednotebook.sourceforge.net>
- [5] Task Coach: <http://taskcoach.org>
- [6] Tasque: <https://wiki.gnome.org/Apps/Tasque>
- [7] Tracks: <http://www.getontracks.org>
- [8] ThinkingRock: <https://www.trgtd.com.au>
- [9] iKog: <http://www.henspace.co.uk/ikog/>
- [10] Tomboy: <https://wiki.gnome.org/Apps/Tomboy>
- [11] KNotes: <https://www.kde.org/applications/utilities/knotes/>
- [12] TuxCards: <http://www.tuxcards.de>
- [13] Sticky Notes: <https://store.kde.org/p/1127372>
- [14] Gnote: <https://wiki.gnome.org/Apps/Gnote>
- [15] OpenStreetMap: <https://www.openstreetmap.org>
- [16] Makagiga plugins: <http://makagiga.sf.net/plugins.php>
- [17] iCalendar: <https://de.wikipedia.org/wiki/ICalendar>

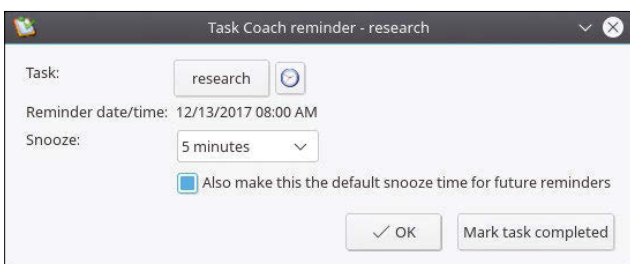



Figure 15: Task Coach also has a reminder function.

Table 1: Overview of To-Do Managers

	Basket	GTG	Makagiga	RedNotebook	Task Coach
Cross-platform	No	Yes	Yes	Yes	Yes
Calendar	No	Yes	Yes	Yes	Yes
Subtasks	Yes	Yes	Yes	Yes	Yes
Memory function	No	No	Yes	No	Yes
Tagging	No	Yes	No	Yes	No
Automatic backup	Yes	No	Yes	Yes	Yes
Data import	Yes	Yes	Yes	No	Yes
Data export	Yes	Yes	Yes	Yes	Yes
Backup	Yes	Yes	Yes	Yes	Yes
Links	Yes	Yes	No	Yes	Yes
Attachments	Yes	Yes	Yes	Yes	Yes
Templates	No	No	Yes	Yes	Yes





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





Customizing file formats with unoconv

FLEXIBLE IMPORT/EXPORT

A hidden utility in the LibreOffice toolbox, unoconv offers a wide array of import and export filter options for use at the command line. *By Bruce Byfield*

LibreOffice is designed to save, import, or export one file at a time, using standard filter settings. The File menu allows you to choose PDF export options, but for most other types of files, you must use the default filter settings. If you want to save multiple files, or adjust the filter settings, you need to shift to the command line and run unoconv [1], a little known Python script that gives you greater control, both with a wide array of import and export filter options.

Unoconv is short for Universal Network Objects (UNO) conversion, a reference to the UNO API used by both LibreOffice and OpenOffice [2]. Bindings for UNO are available for most C++, Java, and Python

compilers, and the API is used to create extensions, as well as to provide support for formats not visible in the LibreOffice desktop window, such as the obsolete LibreOffice 1.0 file formats.

Unsurprisingly, unoconv requires access to LibreOffice's resources. The easiest way to provide this access is to install unoconv on a system that already has LibreOffice installed. However, as detailed in the man page, you can also use the `--connect (-c)` option followed by a comma-separated list to define and connect to the location of a remote LibreOffice instance or `--listener (-l)` to have unoconv detect one.

Unoconv's basic command structure (Figure 1) is:

```
unoconv [FILE].EXTENSION
```

Other files can be added, either in a space-separated list or by using regular

expressions. The command structure assumes that you are exporting the file(s) to PDF format, which is probably the most widely used operation for the command. The extension is the quickest way to specify the type of file, although alternatively you can use the option `--doctype (-d) [TYPE]`, specifying document, graphics, presentation, or spreadsheet. Formulas, databases, or charts are not supported by unoconv – no doubt due to lack of demand, since these types of documents have existed in LibreOffice and its predecessor OpenOffice.org for over a decade. If you prefer to see confirmation that the command has been successfully carried out, you can also add up to three `--verbose (-v)` options – without at least one, unoconv only displays error messages, and the only sign of a completed conversion is the return to the command prompt.

```
Input file: ./nuisance-complaint.odt
Selected output format: Portable Document Format [.pdf]
Output file: file:///home/bb/nuisance-complaint.pdf
```

Figure 1: Unoconv's basic command structure. The `--verbose` option has been added to show the script's operation.

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. You can read more of his work at <http://brucebyfield.wordpress.com>

The following list of document formats are currently available:

```

bib      - BibTeX [.bib]
doc      - Microsoft Word 97/2000/XP [.doc]
doc6    - Microsoft Word 6.0 [.doc]
doc95   - Microsoft Word 95 [.doc]
docbook - DocBook [.xml]
docx    - Microsoft Office Open XML [.docx]
docx7   - Microsoft Office Open XML [.docx]
fodt    - OpenDocument Text (Flat XML) [.fodt]
html    - HTML Document (OpenOffice.org Writer) [.html]
latex   - LaTeX 2e [.ltx]
mediawiki - MediaWiki [.txt]
odt     - ODF Text Document [.odt]
ooxml   - Microsoft Office Open XML [.xml]
ott     - Open Document Text [.ott]
pdb     - AportisDoc (Palm) [.pdb]

```

Figure 2: A few of unoconv's supported formats. Unoconv supports dozens of formats, some of which are not listed in the desktop interface.

If you want to change the export format, add the option `--format= (-f)`. The supported formats for both exports and imports are displayed by running `unoconv --show`. Supported formats include text, CSV, dBase, HTML, PDF, several versions of Microsoft Office formats, StarOffice formats (LibreOffice's original ancestor), common graphic formats, and, of course, current LibreOffice formats (Figure 2).

In addition, unoconv also includes several different housekeeping op-

tions. If a file's attributes matter, you can add `--preserve` so that the output file has the same attributes as the original file. For batch conversions, you might want to use `--output (-o)` to place all the output files in a separate directory, rather than have them mixed together with the original files. The output file can also be password protected by adding:

```
--password= [PASSWORD]
```

Still another interesting option is to set the output file to the same format as the original, then add `--template (-t) [FILE]` to add styles from another file to the output – a command-line version of the *Load Styles* feature in the *Styles and Formatting* window on LibreOffice's desktop interface.


Import and Export Filter Settings

For many users, the default filter settings are all that is needed. However, you can adjust both import and filter settings to your own preferences, using `--export (-e) [SETTING]` or `--import (-i) [SETTING]`. Among other purposes, this ability can be used as an easy method for adjusting the character encoding or date formats in the original file.


Filter settings are added directly after `--export (-e)` or `--import (-i)`, with a separate option for each setting. For text and CSV files, these settings are introduced by `FilterOptions=` and completed by a comma-separated list unique to the format. In the list, settings can be left blank (,) or at the end of the list omit-

GOT CLUSTER?

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



hpc.admin-magazine.com/Newsletter



ted altogether, forcing the use of the default settings.

By contrast, PDF and graphics exports and imports are added after `--export (-e)` or `--import (-i)`, with a separate option for each setting. In other words, to set a password and set the highest image resolution to 300dpi in a PDF file, the command would include:

```
--export PermissionPassword=abcdef
--export MaxImageResolution=300
```

A complete list of standard import and export settings is available online [3], but it is far too long to mention here. However, different types of files have their own set of filter options.

Text Export and Import

For text import, the most common setting to customize is the encoding. A single value can be entered, such as

```
--import FilterOptions=76
```

which would set the encoding to UTF-8. However, for exporting text from a spreadsheet, the `FilterOptions` fields are encoding, field-separator, text-delimiter,

Table 1: Common PDF Export Settings

```
EmbedStandardFonts
EnableCopyingOfContent
EnableTextAccessForAccessibilityTools
EncryptFile
ExportBookmarks
ExportNotes
ExportNotesPages
FirstPageOnLeft
InitialView
MaxImageResolution
OpenInFullScreenMode
PageLayout
PageRange
OpenInFullScreenMode
PageLayout
PageRange
RestrictPermissionPassword
Zoom
```

Table 2: Unique Settings for Graphic Formats

.jpg	ColorDepth, Quality
.png	Compression, InterlacedMode
.gif	Transparency, InterlacedMode
.eps	ColorFormat, Compression, Preview, Version

`quote-all-text-cells`, and `save-cell-content-as-shown`.

CSV Text File Import

CSV files have four basic settings. In order, they are the field separator, the text delimiter, the encoding, and the first line in the file to convert to or from a spreadsheet. For example,

```
--export FilterOptions=44,34,76,2
```

will set commas as the field separator, a double quotation mark as the text delimiter, UTF-8 as the encoding, and the first line in the file to the second. In theory, at the end of the settings, you could add the date format for each column, so that:

```
--export FilterOptions=
44,34,76,2,1/5,2/5,3/5
```

would specify that the date formats for the first three columns would be YY/MM/DD. Any other columns would use the date format already specified for them.

PDF Export Settings

Unoconv does not import PDF files. However, its settings for PDF export are numerous. The man page gives only a list of settings without a full explanation of any of them. However, a full explanation, as well as valid options for each setting, is available online [4].

As a quick reference, you can also open LibreOffice and click *File | Export to PDF* to study the options in the PDF dialog window. For instance, on the *General* tab, the drop-down list of choices for *Reduce image resolution* shows the valid setting for `MaxImageResolution=`. Other options, such as `EnableCopyingOfContent=`, can be completed with *True* or *False*, or, like *Changes*, with *0* for off and *1* for on. With research, you should be able to use many of the available settings. Some, like `EmbedStandardFonts` only need to be listed. Each setting should be entered with a separate `--export` option, if only to make revision of the command easier.

Some of the most commonly used PDF export settings are listed in Table 1. Their names should make most of their purposes reasonably clear.

If researching all the settings is more work than you want to undertake – especially for a single operation you are unlikely to repeat – the default filter for PDF export should be acceptable for most purposes.

Graphics Export Settings

Unoconv’s graphic support for exporting includes `.jpg`, `.png`, `.gif`, and `.eps` file formats, entered in lowercase letters. All these formats support `Height`, `Width`, and `Resolution` settings, making unoconv a convenient way to standardize all the illustrations for an article or a book. In addition, each format supports other settings (Table 2). As with PDF export settings, the man page lists only the settings for graphic formats, leaving you to research valid settings on your own.

The Conversion Toolbox

Unoconv does have a few gaps in its functionality. Its man page lists several specialized scripts that you might need, including `asciidoc-odf` for converting plain text to the Open Document Format used in LibreOffice and other free office applications [5] and `docbook2odf` for converting DocBook XML to Open Document Format [6].

However, in many cases, unoconv offers exact imports and exports for common free office formats. It combines the advantages of sensible and easy-to-use defaults with thorough flexibility – if only you take the time to research the possibilities. Especially if you need to do batch file conversions, unoconv remains one of the best hidden utilities in the LibreOffice/OpenOffice toolbox. ■■■

Info

- [1] unoconv: <https://github.com/dagwieers/unoconv>
- [2] UNO API: <https://api.libreoffice.org/>
- [3] Standard import and export settings: https://wiki.openoffice.org/wiki/Documentation/DevGuide/Spreadsheets/Filter_Options
- [4] PDF export: http://wiki.services.openoffice.org/wiki/API/Tutorials/PDF_export
- [5] `asciidoc-odf`: <http://github.com/dagwieers/asciidoc-odf>
- [6] `docbook2odf`: <http://open.comsultia.com/docbook2odf/>

Now Appearing on

APPLE NEWSSTAND

New age convenience...

Our inspired IT insights
are only a tap away.

Look for us on
Apple Newsstand
and the iTunes store.

Download
a FREE issue of
each publication
now!



Home network monitoring with pfSense, Protectli, and a screen scraper

Servile Guardian

What is making the lights on the router flicker so excitedly? An intruder? We investigate with pfSense on a Protectli micro appliance and a screen scraper to email the information.

By Mike Schilli

It's a shame that no routers simply display the network packet addresses that pass through them on an LED display. Because I'm curious about what's going on in my home network, on the advice of a work colleague, I bought a micro appliance from the Chinese company Protectli (Figure 1) [1], which runs the FreeBSD-based open source firewall pfSense. The

box is about four by four inches in size and passively cooled, so there is absolutely no fan or other noise.

The installation is a piece of cake – simply load the distribution from the pfSense Community Edition website [2] onto a bootable USB stick, insert an mSATA disk and RAM into the Protectli's small case, and after booting, say yes to the installation prompts. Badda-bing badda-boom, pfSense's web GUI is up and running (Figure 2).

Guardian at the Gate

The Protectli appliance is directly connected to the Internet-facing interface (in my case, a DSL modem to the ISP). On the LAN side, it provides access to the Internet for all devices connected to my home network (in my case, a series of routed subnets).

Equipped with a four-core Celeron, it's powerful enough to look at every

single packet, create statistics or even intervene when needed, and block certain communication attempts according to predefined firewall rules. If I want to know why router lights are flickering, I only need to call the pfSense GUI to see who is streaming Spotify, watching Netflix, or ordering on Amazon (Figure 2).

In addition to traditional terminal-based tools such as `pftop`, the firewall GUI also offers very elegant add-on

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 go to mschilli@perlmeister.com he will gladly answer any questions.



Figure 1: The Protectli micro appliance used as a router.



	Application	L4 Proto	Client	Server	Duration	Breakdown	Actual Thpt	Total Bytes	Inf
Info	SSL	TCP	Wohni.localdomain:62853	5 :https	2 min, 19 sec	Server	0 bps	110.7 MB	5
Info	ICMPV6	IPv6-ICMP	fe80::1:1		16 days, 3 h, 11 min, 39 sec	Client	0 bps	19.44 MB	Ro
Info	OpenVPN	UDP	Wohni.localdomain:52928	:openvpn	11 h, 59 min, 59 sec	Client Server	298.2 bps	6.61 MB	
Info	Unknown	UDP	client	:dhcpv6-	5 days, 22 h, 10 min, 43 sec	Client	0 bps	4.87 MB	
Info	SSL	TCP	Wohni.localdomain:51144	:https	2 min, 32 sec	Client	0 bps	4.75 MB	wo
Info	SSL	TCP	Wohni.localdomain:56559	:https	3 min, 41 sec	Client	157.14 Kbit	5.88 MB	wo
Info	OpenVPN	TCP	Wohni.localdomain:58982	:openvpn	2 days, 18 h, 33 min, 28 sec	Client Server	0 bps	4.41 MB	

Figure 2: Finally, a device to monitor all of your home’s network traffic.

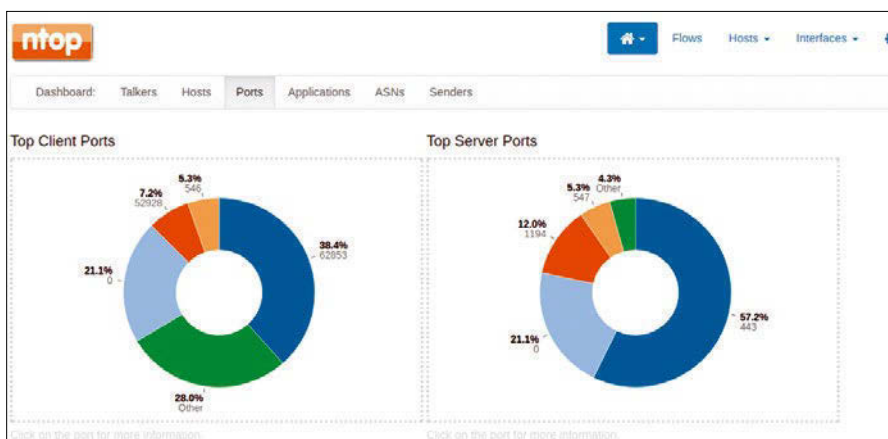


Figure 3: Pie charts for server and client ports.

packages such as ntopng, so you can browse through pie charts and HTML tables to find out who uses the most bandwidth or contacts computers in dubious countries (Figure 3).

Unfortunately, there is no official API for the GUI, only a FauxAPI [3], which runs as an add-on package on the pfSense distribution and provides limited access to the firewall’s internals.

Keys for Protectli

To check at regular intervals what is happening on the Protectli box, I thought it would be easy to write a screen scraper [4] that logs in periodically and automatically to the box’s login page (Figure 4) and scans and mails the data displayed on the dashboard.

However, the first hurdle between a command-line client and the juicy network data is the login page so definitly presented by pfSense. A look at the HTML code (Figure 5) reveals that the two fields for accepting the username and password are dubbed `username-fld` and `password-fld`, and the submit button goes by the name `login`. The Python scraper quickly thrown together in

Listing 1 [5], which uses the `selenium` module installed using `pip3` to simulate a browser, searches for and finds these elements using the `find_element_by_name()` function.

For the `webdriver.Firefox()` call to work with the system’s Firefox browser, the Linux distribution needs the `geckodriver` program, which is available as a TAR file [6]. You need to unpack this and dump the binary that falls out of the archive into a path that can be found somewhere in `$PATH`. The script opens the browser, takes you to the login page, autofills the form fields, and then clicks on the `Login` button. The `selenium` module is often used for testing WebGUIs and makes it really

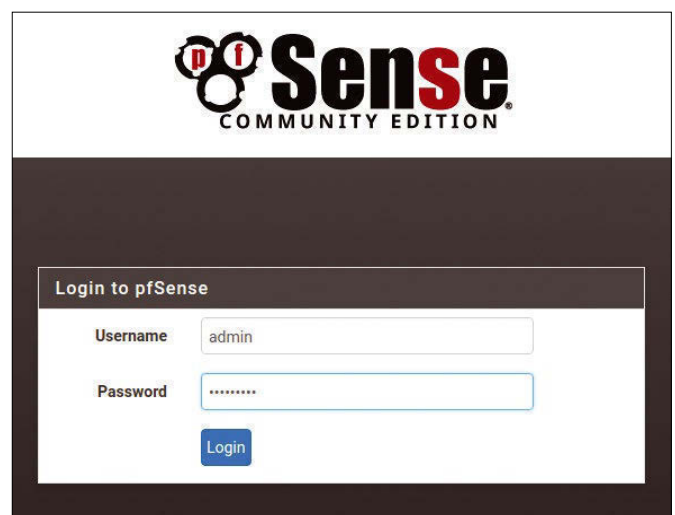


Figure 4: First hurdle: Login.

```

<form method="post" action="/index.php" class="form-horizontal"><input type='hidden' name='__csrf_magic'
" />
  <div class="form-group">
    <label for="usernamefld" class="col-sm-3 control-label">Username</label>
    <div class="col-sm-9 col-md-7">
      <input type="text" class="form-control" name="usernamefld" id="usernamefld" placeholder="Enter your username" autocorrect="off"
hellcheck="false">
    </div>
  </div>
  <div class="form-group">
    <label for="passwordfld" class="col-sm-3 control-label">Password</label>
    <div class="col-sm-9 col-md-7">
      <input type="password" class="form-control" name="passwordfld" id="passwordfld" placeholder="Enter your password">
    </div>
  </div>
  <div class="form-group">
    <div class="col-sm-offset-3 col-sm-9 col-md-7">
      <button type="submit" class="btn btn-primary" name="login">Login</button>
    </div>
  </div>
</form>
...

```

Figure 5: The login page's source code contains fields for the username and password.

easy to simulate a user sitting in front of a web user interface.

After the pfSense login screen, the pfSense dashboard page (Figure 6) with the firewall overview data is saved in the saved.png file by the new screen scraper by calling save_screenshot() as a trigger. The values to be filled in are read by the script from the creds.yaml file, which is read from disk; the data is then stored as a username and password accessible in the creds dictionary (Figure 7).

One Man Went to Mow

Listing 2 is used to periodically send the collected data to an email address. It bundles the PNG file created by List-

Listing 1: dash-scraper.py

```

01 #!/usr/bin/python3
02
03 from selenium import webdriver
04 import yaml
05
06 driver = webdriver.Firefox()
07
08 creds = yaml.safe_load(open('creds.yaml', 'r'))
09
10 driver.get('https://192.168.241.1')
11
12 user_field =
13     driver.find_element_by_name("usernamefld")
14
15 pass_field =
16     driver.find_element_by_name("passwordfld")
17
18 user_field.send_keys(creds['pfsense_user'])
19 pass_field.send_keys(creds['pfsense_password'])
20
21 driver.find_element_by_name('login').click()
22
23 driver.save_screenshot('saved.png')
24
25 driver.close()

```

Listing 2: mail.py

```

01 #!/usr/bin/python3
02
03 import smtplib
04 import yaml
05 from email.mime.multipart import MIMEMultipart
06 from email.mime.text import MIMEText
07 from email.mime.image import MIMEImage
08 from email import encoders
09
10 creds = yaml.safe_load(open('creds.yaml', 'r'))
11
12 attachment = 'saved.png'
13 body = "The latest pfSense Dashboard.â
14
15 msg = MIMEMultipart()
16 msg['From']=creds['from_email']
17 msg['To']=creds['to_email']
18 msg['Subject']='pfSense Status'
19
20 msgText = MIMEText(
21     '<b>%s</b><br><br>'
22     % (body, attachment), 'html')
23 msg.attach(msgText)
24
25 fp = open(attachment, 'rb')
26 img = MIMEImage(fp.read())
27 fp.close()
28 img.add_header('Content-ID', "<{}>".format(attachment))
29 msg.attach(img)
30
31 server = smtplib.SMTP(creds['smtp_server'], 587)
32 server.starttls()
33 server.login(creds['smtp_user'], creds['smtp_password'])
34 server.sendmail(creds['from_email'],
35     creds['to_email'], msg.as_string())
36 server.quit()

```

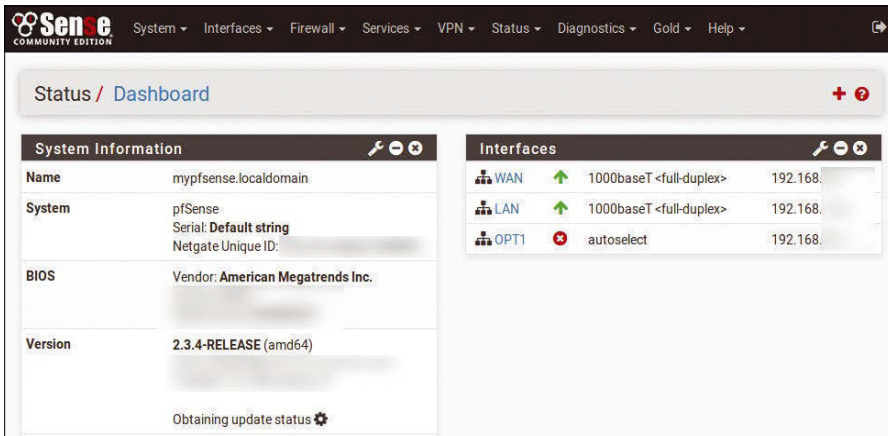



Figure 6: pfSense firewall status information.

ing 1 into an email as an attachment and sends it via an SMTP server. To read the security-relevant username and password variables for the SMTP server, line 10 retrieves the same YAML file as before and stores its contents in the creds dictionary.

The script then builds an HTML body – with introductory text and an IMG link to the attached image, so that a webmail client can display it graphically. As of line 31, Listing 2 establishes a connection to the SMTP server, whose address is also retrieved from the creds.yaml file as the smtp_server: mail.provider.net entry. The script uses port 587 and transmits the data in TLS-encrypted form. It

```
$ cat creds.yaml
pfsense_user: admin
pfsense_password: ***
from_email: m@perlmeister.com
to_email: m@perlmeister.com
smtp_server: mail.provider.net
smtp_user: ***
smtp_password: ***
```

Figure 7: Sensitive data in the creds.yaml file.

appends the screenshot in MIME format and adds an email content ID header with the name of an imaginary file in square brackets.

Figure 8 shows how the mail arrives in Gmail. Called as a cron job once a day, this keeps the home owner up to date with what’s happening on the local network. ■■■

Info

- [1] Protectli micro appliance: <https://www.amazon.com/gp/product/B01GIVQI3M>
- [2] pfSense Community Edition download: <https://www.pfsense.org/download/>
- [3] pfSense FauxAPI: https://github.com/ndejong/pfsense_fauxapi
- [4] Jarmul, Katharine, and Richard Lawson. *Python Web Scraping*, 2nd ed. Packt Publishing, 2017
- [5] Listings for this article: <ftp://ftp.linux-magazine.com/pub/listings/linux-magazine.com/208/>
- [6] geckodriver: <https://github.com/mozilla/geckodriver/releases/tag/v0.19.0>

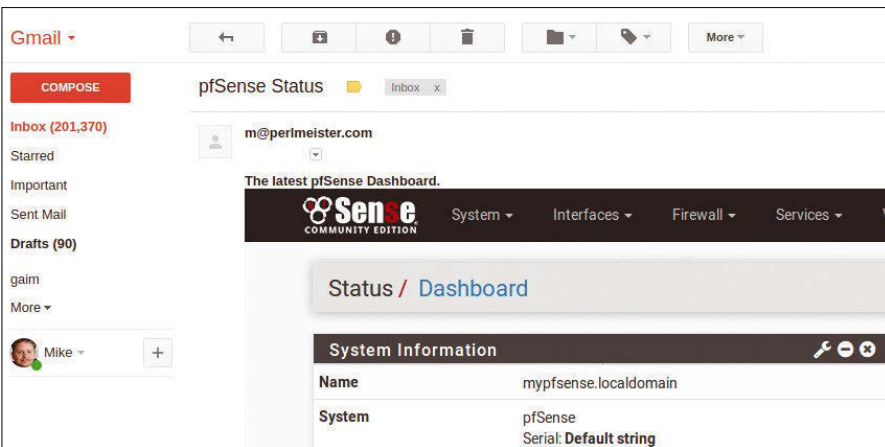


Figure 8: The screen-scraped dashboard arrives by email.

The intersection of DEVELOPMENT and OPERATIONS


Check out our new ADMIN DevOps corner!

www.admin-magazine.com/DevOps

SPONSORED BY



Linux Professional Institute



Ubuntu derivative BackBox Linux for security analysis

(In)secure?

Specializing in security and forensics, BackBox Linux is not only good for a vulnerability assessment, but thanks to its lean substructure, it is also suitable as a desktop distribution.

By Erik Bärwaldt

Many security-related Linux derivatives focus on a specific area of IT security and only take other problems into account marginally, or not at all. In contrast, the Italian Ubuntu derivative BackBox Linux, which has been in continuously development for several years, addresses most of the security-relevant issues faced by administrators of small and medium-sized networks.

You can pick up the approximately 2.5GB ISO image from the BackBox project page [1] for both 32- and 64-bit architectures. Here, you can choose whether you want to download the image directly or via BitTorrent. You can either specify a donation amount or enter 0 in the corresponding field.

As the minimum system requirements, the developers specify a computer with 1GB of RAM and 10GB of free space on mass storage. The screen resolution

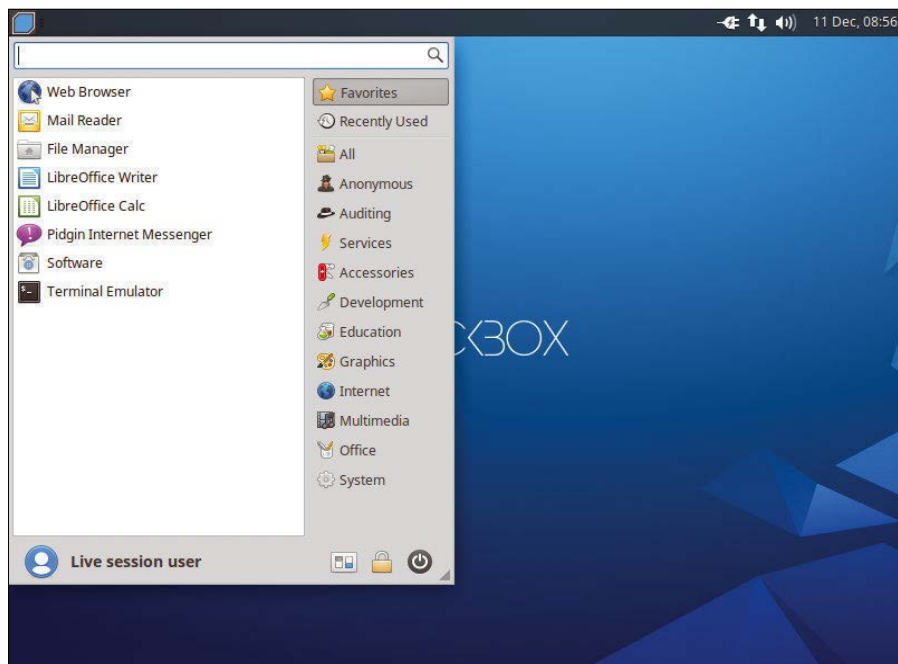


Figure 1: The Xfce desktop does has no surprises; only a couple of entries in the start menu indicate the distribution's peculiarities.

Lead Image © Andrea De Martin, 123RF.com

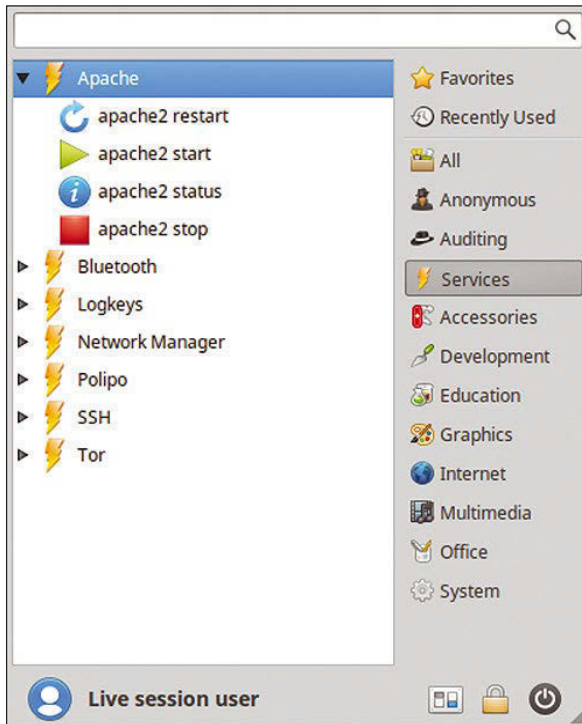


Figure 2: Preconfigured services can be launched by clicking in BackBox's start menu.

should be at least 800x600 pixels; both a USB memory stick and a DVD can be used as boot drives.

Much like Ubuntu, the system starts up automatically without intervention on booting. If you press a key in the first five seconds, a selection menu appears in which you can specify the localization, if needed. After a short time, an Xfce v4.11 desktop appears in a conventional design with a horizontal panel at the top of the screen (Figure 1).

At first glance, three unusual categories stand out in the start menu: the Anonymous, Auditing, and Services submenus indicate the distribution's universal applicability for security-related tasks. On the desktop itself, in addition to the obligatory recycle bin and some icons for existing drives and folders, there is also a launcher to install the system on the mass storage device.

Clicking on the launcher activates Ubiquity, Ubuntu's default installation wizard, which appears in brighter colors on BackBox than on the original Ubuntu.

Software

BackBox Linux is based on Ubuntu 14.04, with kernel 4.4 underpinnings and systemd. The distribution comes

with many popular pre-installed programs, including LibreOffice 4.2.8, Firefox 50.0.2, Gimp 2.8.10, and Thunderbird 45.5.1. However, multimedia programs such as VLC and Audacity are missing, and you won't find a Games submenu.

For this purpose, in the Services folder, BackBox offers the option to start or stop various services and processes with a mouse click. These include the Apache web server, the Tor anonymization network, the PostgreSQL database, and the SSH daemon. The status of these services can be easily determined with a mouse click (Figure 2).

You can use the BleachBit delete tool to remove unnecessary and obsolete data, which mainly includes rotating logfiles, but also temporary storage or histories that show a user's behavior on the computer. BackBox offers two starters: one for users and one for administrators with extended permissions.

In the Anonymous menu, you can start, stop, or determine the current status of the *anonymous* service by clicking the like-named scripts (e.g., *anonymous start*). The anonymous script lets you generate an incorrect MAC address for one of the network interfaces integrated into the system. The service also lets you modify the hostname and finally starts the Tor anonymizing service.

The entire script runs interactively in the terminal. To open an Internet connection secured by a Tor network, click on *Vidalia* in the Internet menu.

The graphical Vidalia control panel, which provides a great deal of interesting information about Tor in a visual format, automatically establishes the connection to the service (Figure 3).

In the Anonymous menu you will also find a launcher that calls a script to erase the memory before switching off. Switch this function on or off in the terminal after entering your authentication data.

Unlike other IT security distributions, BackBox does not provide the Firefox web browser with add-ons by default. You will therefore want to reinstall and configure uBlock Origin, Ghostery, or HTTPS Everywhere to stop unmanageable tracking by advertising agencies.

Pidgin, the preinstalled chat client, optionally supports encrypted communication using the OTR plugin, but you have to install this extension manually. In the Thunderbird email client, I also recommend setting up add-ons such as Enigmail, which allow messages to be simply encrypted with OpenPGP or GnuPG.

Core

The core of BackBox is the Auditing menu, with 16 groups of security-related applications. In addition to tools for penetration tests, you will also find many programs for network and web app analysis, as well as forensic tasks. It also contains well-known sniffers and vulnerability analysis programs.

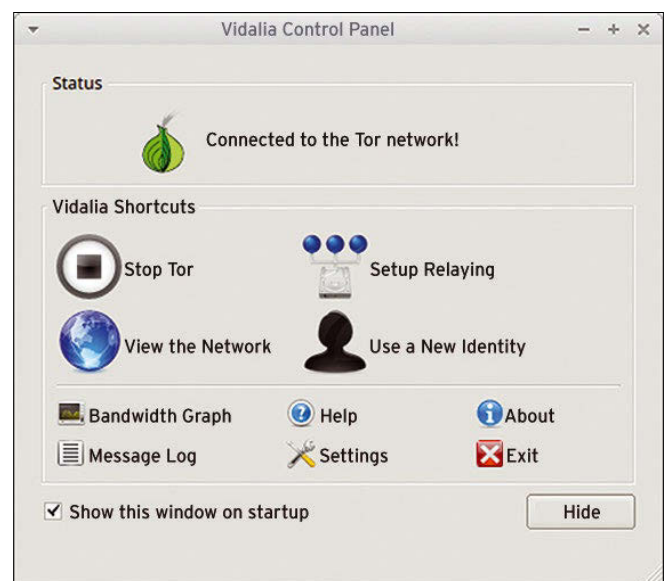


Figure 3: The Tor anonymization service can be conveniently managed with Vidalia.

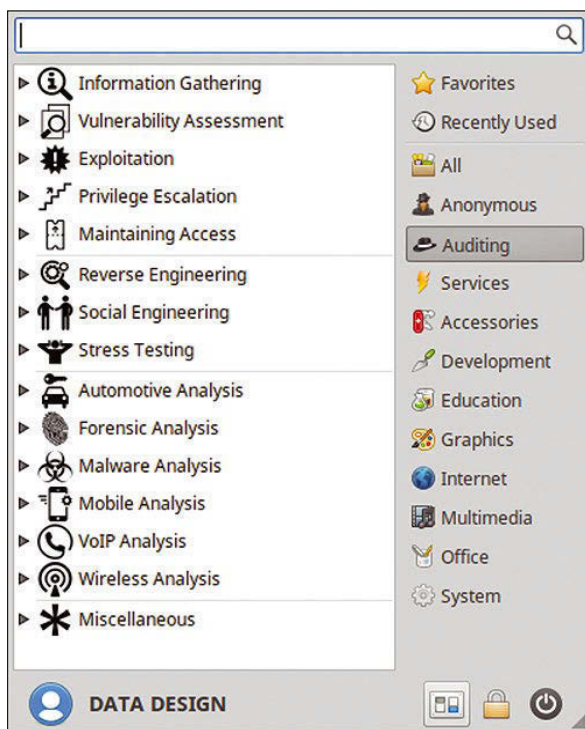


Figure 4: BackBox has the right software for many problems.

The developers also incorporated some unusual tools into the system: For example, tools from the *Automotive Analysis* submenu can be used to read the CAN bus of motor vehicles, whereas *Mobile Analysis* deals with smartphones and tablet PCs. The Android and iPhone submenus group the respective applications by system type.

The *Wireless Analysis* submenu, which distinguishes between Bluetooth and WiFi applications, is similarly laid out. Under *WiFi | Cracking*, you will find programs for checking WiFi encryption, such as *Aircrack-ng*. The *Scanning* submenu combines applications for detecting WiFi networks, including the Kismet program.

Various applications for LAN analysis offer significant benefits: Password crackers, some of which can be used for specific file types, make it possible to trace password strength. With network sniffers such as Ettercap and Wireshark, you can also quickly identify weak points in the intranet by recording data transfer (Figure 4).

In addition to software that exclusively deals with computer and data security, you will find tools for reconstructing data in various subgroups in the Forensic Analysis menu. This includes – especially in the *Data Recovery* submenu –

applications such as PhotoRec, TestDisk, Scalpel, and Foremost. The same menu also has applications for file and partition analysis.

All-Rounder

Thanks to the lean Xfce desktop, BackBox is already very efficient with resources by nature. For example, with its graphical desktop and all commonly started services, the system only occupies about 500MB of RAM, which means that it runs smoothly, even on poorly equipped computers. Because most of the supplied tools start in the terminal, they do not require any additional memory.

Thanks to the low hardware requirements, BackBox is also suitable as a solid all-rounder for daily desktop use. For this reason, the developers implemented the graphical Synaptic front end for package management, which also makes it easy to install updates.

Package management already specifies all package sources, including those from third-party providers and the Tor Project, so you can keep your security-

relevant software up to date at all times. Therefore, Synaptic lists more than 48,000 packages for the BackBox system (Figure 5).

Conclusions

With BackBox, you get a solid operating system. On the one hand, it is suitable as a specialized Linux derivative for security audits and anonymous Internet surfing with Tor. On the other hand, it can also be used as a general Linux desktop. The system is based on Ubuntu 14.04 LTS and therefore comes with the complete Ubuntu software inventory. Additionally, the developers provide their own repository, which is primarily used to keep BackBox's integrated safety-relevant tools up to date.

Because of its slim Xfce desktop and low system requirements, the Italian Ubuntu derivative explicitly addresses users of older hardware who work with limited resources. Therefore, the distribution should be of interest to all users who are looking for an everyday system that also takes many security aspects into account. ■■■

Info

[1] BackBox download: <https://backbox.org>

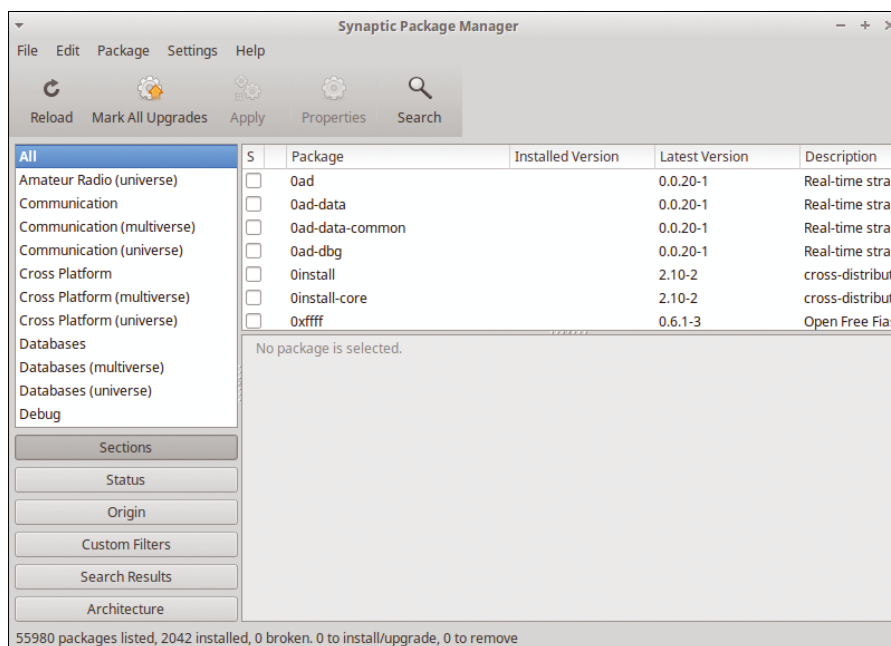


Figure 5: Thanks to Synaptic package management, BackBox can be easily updated and extended with additional software.

Shop the Shop

shop.linuxnewmedia.com


Want to subscribe?

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



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

▶▶ shop.linuxnewmedia.com ◀◀

DIGITAL & PRINT SUBSCRIPTIONS



SPECIAL EDITIONS



The sys admin's daily grind: Tsung

More Lust for Load

How many users can the database take? When does a CMS throw in the towel? In order to explore performance limits, Charly Kühnast uses the Tsung load generator instead of human users as beta testers. *By Charly Kühnast*

If I want to test how much load a (perhaps even distributed) system can take, I launch a load generator. Some time ago, I praised Siege [1] in my column, which I still consider to be a good barrage tool. However, most load generators fire unrealistically from all barrels and do not simulate the behavior of a real user. Tsung [2] can do this better.

Tsung evolved in several evolutionary steps from a tool that ran load tests against Jabber/XMPP servers. Under the fear-inspiring name of idx-Tsunami, it was given multiprotocol capabilities. Since 2014, the development of idx-Tsunami has petered out. Tsung has simply taken the basis and continued developing Tsunami's codebase.

XMPP is still one of the services that Tsung can deploy to cause unrest on its test servers. On top of this, Tsung supports HTTP with and without TLS, WebDAV, SOAP, PostgreSQL, MySQL, AMQP, MQTT, and LDAP. All protocols are integrated via a plugin engine, so further protocols can follow at any time.

Planning the Attack Using XML

Using XML configuration files, the Tsung user designs their load test scenarios in detail. For example, you can stipulate that the requests should not only originate from one machine, but that several load generators (or clients) should play a key role. I can assign more or less work

to clients with different performance characteristics by using weighting. I can also configure several back-end servers. IPv4 and IPv6 are allowed for the connections, also in mixed mode.

The details of the requests that Tsung uses to stress the servers can be configured within a wide range. In order to simulate realistic user behavior, the software does not torment the servers with constant fire on request, but instead makes well-planned pauses, just as a human user would if he or she were looking at the content of a website and then clicking on it.

Reception Center

If you want to make it even more realistic, use the supplied recorder: After starting, it records the behavior of one or more users, and Tsung replicates this session later. For example, variables can be brought into play when simulated users enter data in a search mask.

I can bundle a group of requests into one transaction. Tsung understands this term as a logically related request, for example: A user calls the website, authenticates themselves (say, using OAuth), then accesses the sub-page using the search function, and submits a search query.

Statistics Reveal All to the Administrator

In addition to the existing evaluations of the load behavior for the back-end servers, Tsung also generates reports on the performance of such transactions (Figure 1). These statistics are, as expected, more useful for the behavior of the systems in production than synthetic flak tests – and that's exactly what I like about Tsung. ■■■

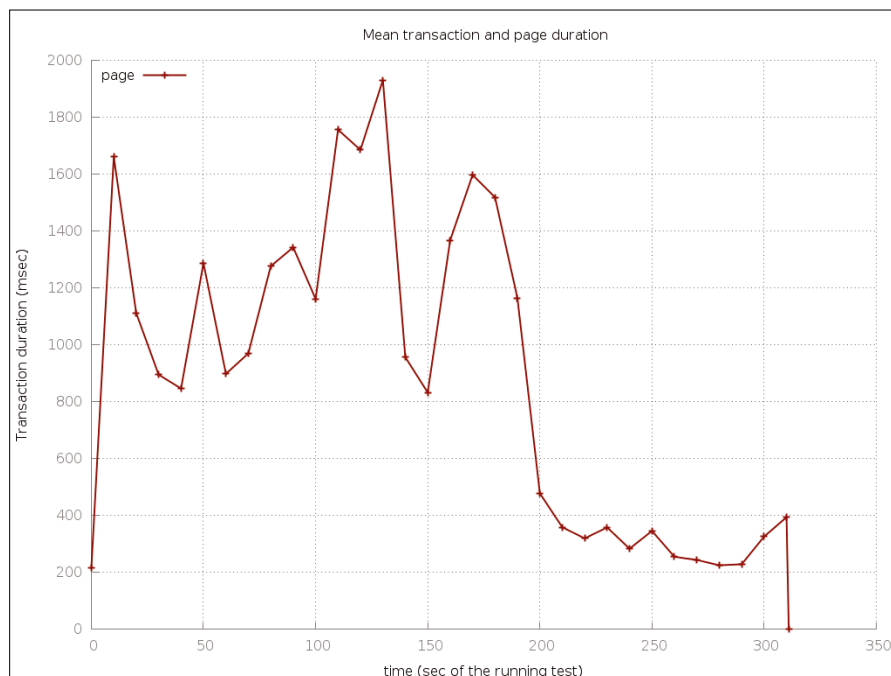


Figure 1: In this report, Tsung has the transaction time per call on the y axis and the sequence of the benchmark on the x axis.

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.

Info

[1] Siege: <https://www.joedog.org/siege-home/>


[2] Tsung: <https://github.com/processone/tsung/>

LINUX UPDATE

Need more Linux?

Our free Linux Update newsletter delivers insightful articles and tech tips to your mailbox twice a month. You'll discover:

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



www.linuxpromagazine.com/mc/subscribe



MakerSpace

Raspberry Pi media center
with DIY ambient light

Let There Be Light

Ambilight lights up the wall behind Philips TVs with LEDs mounted on the device to enhance the onscreen visual impression. With Lightpack and a Rasp Pi media center, every TV can be upgraded with ambient light. *By Christoph Dyllick-Brenzinger*

The Ambilight technology, originally developed by Philips, lights up the wall behind the television set by monitoring the colors in the current image and generating the appropriate color to extend the image content beyond the monitor surface. The “extended” image is easier on the eyes, because the field of vision is enlarged for the viewer.

An Ambilight system basically comprises three components: at least three edges (left, right, and top) on the back of the TV for strips of RGB LEDs, a control unit that ensures the individual LEDs light up accordingly, and a signal to the control unit provided by software that continuously analyzes the video image.

If you don’t want to buy a Philips Ambilight TV, but you already watch movies and TV series on the Kodi multimedia platform, you can upgrade to ambient lighting in a DIY project with either a Raspberry Pi with Raspbian, a Kodi distribution like Open/LibreELEC, or a Linux PC. Various preexisting solutions such as Lightpack [1], Lightberry [2], or AmbiLED [3] remove the need for further programming

For the test, I chose the Lightpack PC [4] package by Woodenshark, which made the best overall impression. It originated from a Kickstarter campaign and

costs \$89, for which you receive 10 pre-assembled LED strips, an unobtrusive control unit, an AC power supply with adaptors, cables, and mounting accessories. All components can be attached with double-sided adhesive tape without causing damage to the TV.

On the software side, Boblight [5] and Hyperion [6] are two popular solutions. Thanks to Hyperion’s less complicated configuration, you can achieve the desired results faster. Therefore, the ambient light system setup in this example comprises Lightpack PC, Hyperion, and a Raspberry Pi.

Setup

When attaching the LED strips to the back of the screen, be careful not to block the connections and any attachment points that can be used on the TV. In the Lightpack packaging, you will find adhesive pads and cable ties that facilitate cable management (Figure 1).

Lightpack’s manufacturer recommends that the strips should be applied in the “Andromeda” pattern, which promises the most uniform illumination possible (Figure 2). Gradually mount the LEDs in a clockwise direction, starting at the bottom left-hand side of the TV housing; then, connect the strips to the control unit in the same order.



Figure 1: The complete Lightpack kit mounted behind a TV set. A classic PC delivers the image data.



Figure 2: In the "Andromeda" arrangement, you attach the LED strips clockwise from the bottom left.

To make the assembly of individual modules as easy as possible, lay the screen display-side down on a soft surface so that the relevant points remain easily accessible without the fear of scratching the screen. After connecting the power supply unit and USB cable, you need to teach the

Raspbian system how to install the ambient lighting system.

Hyperion

To set up Hyperion, download an installation script offered on the project page, make it executable, and then call it with root privileges:

```
$ cd /tmp
$ wget https://raw.githubusercontent.com/hyperion-project/hyperion/master/bin/install_hyperion.sh
$ chmod +x install_hyperion.sh
$ sudo ./install_hyperion.sh
```

The program then prompts you to reboot the system.

Hyperion imports the installation routine together with a unit for systemd, but the service cannot start successfully yet. Because the system does not automatically detect how many LED strips it is supposed to control and where they are located, you have to carry out the rest of the configuration manually.

For the LED setup, download the program HyperCon [7] (suitable for all current operating systems) onto your computer. Hyperion and HyperCon do not necessarily have to run on the same system: The configuration file written by HyperCon can be copied to the media center computer at a later stage.

The Java application does not require an installation procedure, but it does require the Java Runtime Environment version 1.7 or later, which comes with the Raspbian "Stretch" distro with the Pixel desktop. Listing 1 shows how to view the Java version and then launch the application. You need to be in the same directory as the JAR file for this step.

In the HyperCon application window, first change *Type* to *Lightpack*. The *Serial #* input field (Figure 3) then appears

below it. This serial number does not correspond to the sequence of digits printed on the back of the control; instead, you need to look it up in the Raspberry Pi's kernel log (Listing 2). However, this only works if you have previously connected

Listing 1: Launching HyperCon

```
$ java -version
java version "1.8.0_65" Java(TM) SE Runtime Environment (build 1.8.0_65-b17)
Java HotSpot(TM) Client VM (build 25.65-b01, mixed mode)
$ java -jar HyperCon.jar
```

Listing 2: Finding the Lightpack Serial Number

```
$ dmesg
[...]
May 19 13:41:05 homie kernel: [ 1.001954] usb 1-3: New USB device strings: Mfr=1, Product=2,
SerialNumber=220
May 19 13:41:05 homie kernel: [ 1.001955] usb 1-3: Product: Lightpack
May 19 13:41:05 homie kernel: [ 1.001956] usb 1-3: Manufacturer: lightpack.googlecode.com
May 19 13:41:05 homie kernel: [ 1.001957] usb 1-3: SerialNumber: 12345678901234
[...]
```

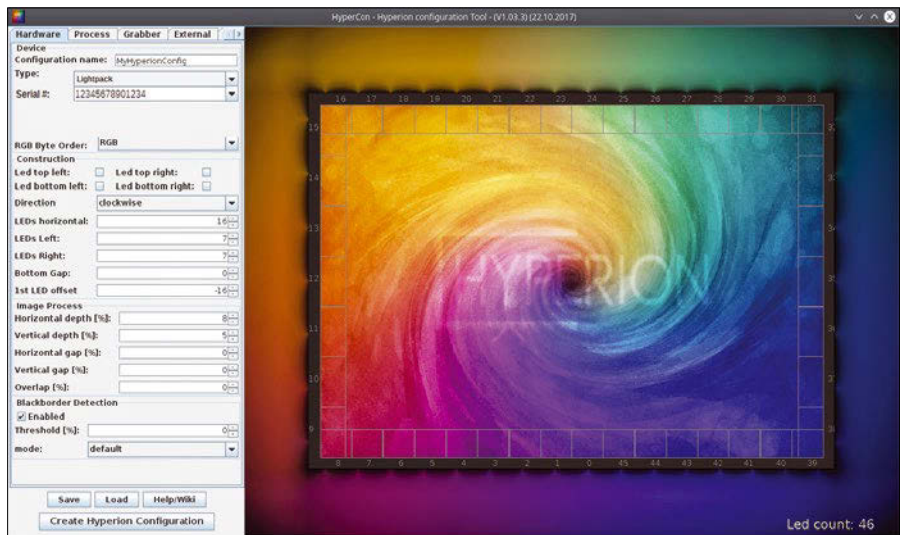


Figure 3: HyperCon makes creating a configuration file for the Hyperion service easy.

the Lightpack control unit to the Rasp Pi via USB.

After entering the serial number, reduce the number of LED strips to 10. To do this, first enter 4 in the *LEDs horizontal* field under Construction; then, enter 2 for *LEDs left*, *LEDs right*, and *Bottom gap*; accept the remaining values and click *Save*.

Now create the configuration file by clicking *Create Hyperion Configuration* and save it as `hyperion.config.json` in

the home directory. Next, copy it with root privileges to `/etc/hyperion/` and start the Hyperion service:

```
$ sudo cp ~/hyperion.config.json
/etc/hyperion
$ sudo systemctl restart hyperion
```

All LEDs should now light up and flash in several colors. The flashing colors show you that Hyperion is communicating successfully with Lightpack. To ensure that the service will load automatically at system startup, enable it in `systemd`:

```
$ sudo systemctl enable hyperion
```

If the service cannot be started immediately, you might have an incorrect setting in the configuration file. In `hyperion.config.json`, check that HyperCon also searches for the effects in `/usr/`

`share/hyperion/effects/` and not only under `/storage/hyperion/effects` (Listing 3).

Video Grabbing

If the colors of the LEDs do not yet match the screen content, the software has to evaluate the video signal. Add-ons for a Rasp Pi media center with a Kodi installation like *Boblight* [8] or *AmbiBox* [9] analyze the video signal supplied by Kodi and transmit the control signals to the Hyperion service.

Instead of picking up the video signals in the application software, Hyperion can also evaluate the X server image directly. *Ambilight* thus supports all screen content. Hyperion comes with the *hyperion-x11* module, which you launch with the X server by creating the `hyperion-x11.desktop` file in the `~/.config/autostart/` folder and adding the contents of Listing 4.

After restarting the system, the LEDs respond to changes in the screen content. To try this out, simply move a window that is not too large and with the brightest possible background (e.g., a file manager) back and forth on your desktop.

Optimization

The most difficult part of the configuration is defining the order and position of the LEDs in the "leds" section of the `/etc/hyperion/hyperion.config.json` configuration file. The entry for a single LED (or, in the case of Lightpack, for an LED strip) is shown in Listing 5.

Each entry can be assigned uniquely by the index value. A pair of coordinates defines the position, or rather the area, in which the system should

react to the contents displayed on the screen. The values run from 0.0000 to 1.0000. The information in Listing 6 shows the setting for the upper left corner of the screen.

However, in the case of the devices I used in the lab test, the order of the connections

Listing 3: Hyperion Effects Paths

```
// EFFECT PATH
"effects" :
{
  "paths" :
  [
    "/storage/hyperion/effects",
    "/usr/share/hyperion/effects" ]
},
```

Listing 4: Desktop File

```
[Desktop Entry]
Version=1.0
Encoding=UTF-8
Name=Script
Type=Application
Exec=hyperion-x11
Icon=
Terminal=false
StartupNotify=false
Hidden=false
GenericName=
GenericName[en_US]=
```

Listing 5: Single LED Entry

```
[...]
{
  "index" : 4,
  "hscan" : { "minimum" : 0.7000, "maximum" : 0.9000 },
  "vscan" : { "minimum" : 0.3000, "maximum" : 0.5000 }
},
[...]
```

Listing 6: Upper Left Corner

```
"hscan" : { "minimum" : 0.0000, "maximum" : 0.1000 },
"vscan" : { "minimum" : 0.0000, "maximum" : 0.1000 }
```

on the control unit and the index number did not match; thus, I had to reposition either the connections or the LEDs accordingly: I opted for the latter to keep the cable routing and arrangement on the control unit as clean as possible.

The Lightpack kit contains 10 LED strips, and the control unit

controls each of these by reference to a unique index. It makes sense to isolate a certain area in the configuration file and assign it to the various lighting elements one after another.

With the LED configuration from Listing 7, the active area for the first LED strip is in the upper left corner of the screen. First place all other LED strips on the bottom right area, and then apply the changes by restarting the Hyperion service:

```
$ sudo systemctl enable hyperion
```

Once you have assigned all elements, the alignment can be corrected by means of coordinates. Figure 4 shows a simple configuration. The values are then transferred to the appropriate LEDs.

Color Intensity

In practice, the system should now light the area behind the TV to match

Listing 7: LED Configuration

```
[...]
{
  "index" : 0,
  "hscan" : { "minimum" : 0.0000, "maximum" : 0.1000 },
  "vscan" : { "minimum" : 0.1000, "maximum" : 0.1000 }
},
{
  "index" : 1,
  "hscan" : { "minimum" : 0.9000, "maximum" : 1.0000 },
  "vscan" : { "minimum" : 0.9000, "maximum" : 1.0000 }
},
{
  "index" : 2,
  "hscan" : { "minimum" : 0.9000, "maximum" : 1.0000 },
  "vscan" : { "minimum" : 0.9000, "maximum" : 1.0000 }
},
[...]
```

hscan 0.1 - 0.3 vscan 0.1 - 0.3	hscan 0.3 - 0.5 vscan 0.1 - 0.3	hscan 0.5 - 0.7 vscan 0.1 - 0.3	hscan 0.7 - 0.9 vscan 0.1 - 0.3
hscan 0.1 - 0.3 vscan 0.3 - 0.5			hscan 0.7 - 0.9 vscan 0.3 - 0.5
hscan 0.1 - 0.3 vscan 0.5 - 0.7			hscan 0.7 - 0.9 vscan 0.5 - 0.7
hscan 0.1 - 0.3 vscan 0.7 - 0.9	hscan 0.3 - 0.5 vscan 0.7 - 0.9	hscan 0.5 - 0.7 vscan 0.7 - 0.9	hscan 0.7 - 0.9 vscan 0.7 - 0.9

Figure 4: Hyperion divides screen content into areas and uses them to determine the color signal sent to the LEDs.

the screen content. However, the brightness of the LEDs, and thus the intensity of the effects, may leave much to be desired.

If necessary, you can also control luminance with the `hyperion.config.json` configuration file. The "transform" section (Listing 8) is responsible for the intensity of the color effects. An optimal setting requires a successful balance between saturation ("saturationGain") and brightness ("luminanceGain").

Tests have shown that brightness values above 1.5000 have a counterproductive effect: The colors appear rather pale, if not white. The best way to find out the optimum configuration for your requirements is to conduct your own experiments.

Conclusions

With a Rasp Pi, the relatively inexpensive Lightpack, Hyperion software, and a little configuration work, you can save yourself the cost of an Ambilight TV. However, Lightpack only "lights up" content played back by the Media Center, not conventional television programs.

The biggest task is reworking the configuration: The file created by the setup tool requires some adjustments. If you

like to experiment with media center distributions and avoid modern smart TVs with all the extras, this task should not deter you. ■■■

Listing 8: JSON transform Section

```
"transform" :
[
  {
    [...]
    {"hs1" :
    {
      "saturationGain" : 1.0000,
      "luminanceGain" : 1.0000,
      "luminanceMinimum" : 0.0000
    }
  },
  [...]
]
```

Author

Christoph Dyllick-Brenzinger is technical director of datamate Service GmbH, which develops preconfigured desktop PCs and server solutions in solid wood housings for private and commercial customers.

Info

- [1] Lightpack: <http://www.lightpack.tv>
- [2] Lightberry: <http://lightberry.eu>
- [3] AmbiLED HD: <http://www.ambiledhd.com>
- [4] Lightpack PC: <https://store.lightpack.tv/products/lightpack>
- [5] Boblight: <https://code.google.com/archive/p/boblight>
- [6] Hyperion: <http://hyperion-project.org>
- [7] HyperCon: <https://github.com/hyperion-project/hypercon>
- [8] Kodi Boblight add-on: http://kodi.wiki/view/Add-on:XBMC_Boblight
- [9] Kodi AmbiBox add-on: <http://kodi.wiki/view/Add-on:AmbiBox>



MakerSpace

Design printed circuit boards with KiCad

Boardwalk

KiCad helps you design printed circuit boards with up to 32 layers, checks for optimal placement, and supplies schematics and assembly diagrams in the popular Gerber format for submission to PCB manufacturers. *By Harald Zisler*

KiCad was originally launched as simple schematics software; today, it supplies blueprints for components used at the CERN nuclear research center in Geneva, Switzerland. Accordingly, CERN's scientists are actively involved in the development of KiCad, which is available for Linux, FreeBSD, Mac OS, and Windows [1].

Schematics created with this program are suitable for importing into other pro-

grams, depending on the format. KiCad comes with numerous features and, unlike the free versions of commercial counterparts such as Eagle Light Edition or Target 3001!, there are no restrictions on its use and scope. See the "Basics" boxout if you are a newcomer to the field.

The program is available for many distributions. You can download the latest version from the project page, where you also will find the source code if you want to compile the software yourself.

BASICS

In KiCad, you create the schematic first. After checking it carefully, you then design the printed circuit board (PCB), which brings the components and their "footprints" together – footprints being the connection geometries of the various components. Components can be wired (conventional, through-hole devices) or surface-mount devices (SMDs): Wired components are located on the upper side of the printed circuit board, with the soldered side and the conducting paths mirrored on the underside. SMD components are soldered directly on the PCB surface, revealing the circuit diagram in the printed conductors. Information about components can be found in libraries [2] maintained by volunteers [3].

BACKPORT

As root, add the package source for KiCad to the `/etc/apt/sources.list` file; then, update the package data and install the software and the associated help files:

```
# echo "deb http://httpredir.debian.org/debian jessie-backports main contrib
non-free" >> /etc/apt/sources.list
# apt-get update
# apt-get -t jessie-backports install kicad
```

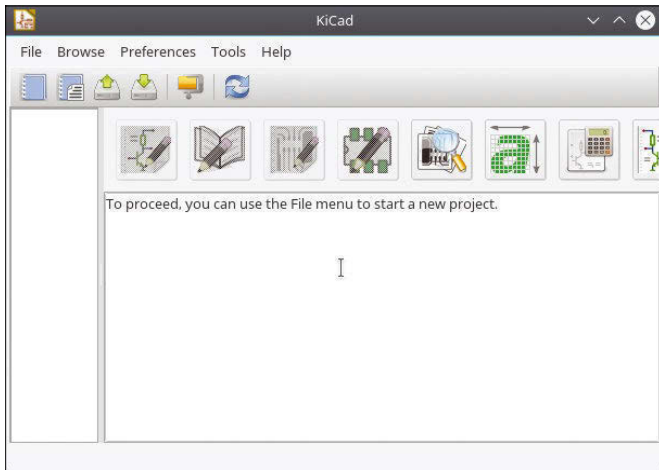


Figure 1: After the first call, the KiCad program window looks quite empty.

Users of Debian-based systems who use an older version of the distribution can access backports (see the “Backport” box). The Unstable and Testing branches already contain the latest version of the software.

Creating a Project

When you run the program (Figure 1), you are prompted to create a PCB design

project. To do this, either click on the left or take the classic *File | New Project* route.

For the sake of clarity, create a new folder (in Figure 2, this is *first*), assign a name to the project (e.g., *first* here), and press *Save*. The left column of the program window contains a tree view of the project and its components with

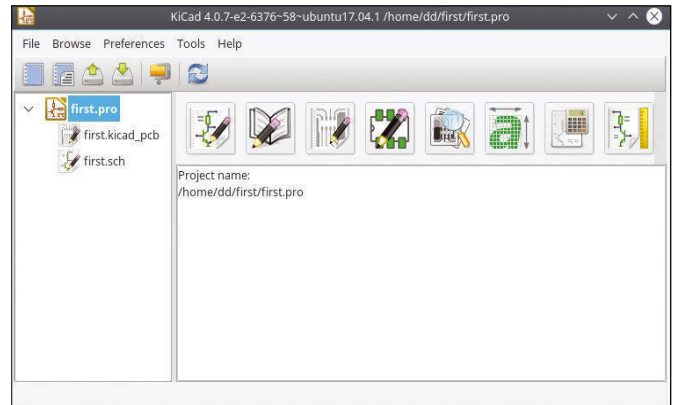


Figure 2: After creating a project, a multitude of functions become available in the main window.

the first set of files (see Table 1), and the individual components are no longer grayed out.

In the File menu, you can open existing projects or create new ones. You can also create or read archives. Under *Browse*, you have the option of editing files with a text editor of your choice, which is determined under *Preferences*, where you can also configure the language for the user interface and the PDF viewer. Additionally, you can define paths and set up keyboard shortcuts (called hotkeys). Table 2 shows the shortcuts for the main window. Each tool also has its set of hotkeys.

KiCad manages projects and calls up a number of tools, which you can access either from the Tools menu or directly in the toolbar (Table 3). The Help menu offers a manual and a tutorial for beginners.

Small Circuit

A design for a small 5V power supply is a good example of how to trace all the important work steps. Click on the









Table 1: Project Files

Extension	Content
.pro	Entries for project management
Shared	
.kicad_uks	Adjustments to the worksheet
.net	Netlist
.cmp	Component assignments in the footprint schematic (create with the Pcbnew board layout editor; import into Eeschema)
Eeschema (Schematic Editor)	
*_cache.lib	Copies of the parts used in the schematic
.dcm	Documentation of the component library (information about components)
.lib	Component library with information on design and connections
.sch	Schematic file without components
Pcbnew (Printed Circuit Board Editor)	
fp-lib-table	List of footprint libraries
.kicad_mod	Footprint file with description
.kicad_pcb	Printed circuit board with board layout
.pretty	Directory of the footprint library folder
Output Files for Printed Circuit Board Production	
.drl	Drilling data in Excellon format
.gbr	Gerber file for producing the printed circuit board
.pos	Positioning instructions for insertion machines
Output Files for Documentation	
.dxf	Plot in DXF format
.pdf	Plot in PDF format
.plt	Plot in HPGL format
.ps	Plot in PostScript format
.rpt	Report file (text)
.svg	Plot in SVG format

Table 2: Hotkeys

Shortcut	Function
Ctrl+O	Load project
Ctrl+S	Save project
Ctrl+N	New project
Ctrl+T	New project from a template
Ctrl+E	Eeschema – schematic editor
Ctrl+L	Schematic library editor
Ctrl+P	Pcbnew – printed circuit board editor
Ctrl+F	PCB footprint editor
Ctrl+G	GerbView utility
Ctrl+B	Bitmap2Component utility
Ctrl+C	PCB calculator
Ctrl+Y	PI editor

Table 3: Components

Symbol	Task
	Eeschema schematic editor
	Schematic library editor
	Pcbnew printed circuit board editor
	PCB footprint editor
	GerbView
	Bitmap2Component utility
	PCB calculator
	PI editor

new project you just created and start *Eeschema*. If you see any error messages, they only appear on the first call, so you can simply ignore them.

Click on the first icon, *Page settings*, and set up your project (Figure 3). To fill

the desktop with components, you can either click *Place | Component* or choose the third symbol from the top in the right vertical toolbar.

Click in the schematic sheet to position the component. You will now see the Choose Component selection dialog (Figure 4). Click *OK* to position the component. Each component used appears in the *History* of the Choose Component window, making it easy to reuse components. To find a component, enter a search term in the *Filter* field.

As long as you hold the component with the mouse cursor while placing, you can rotate it through 90 degrees by pressing *R*. If you insert an incorrect component, you can remove it by pressing *Ctrl + Z* or right-click and choose *Delete Component*.

You can edit components you have already positioned by selecting the element with the left mouse button and then right-clicking to open a context menu (Figure 5). A prompt often appears (*Clarify Selection*) asking what exactly you want to edit. In this dialog, select the component.

Adding the required data for the components you place uses the context menu, as well. Select *Edit Component*,

and enter the reference and the value (Figure 6). For example, for a resistor, you might choose *Resistor: R1, 100 Ohm*. These entries appear later in various lists and in the layout for assembly.

Once all the components are on the desktop, connect them either by selecting *Place | Wire* or by clicking the fifth icon from the top in the right-hand toolbar. KiCad shows some weaknesses here: You always need to drag from an unconnected pin on one component to an already otherwise connected component.

Where three lines meet, the program creates a node. If two connections cross and you want to connect them, you click *Place | Junction*. If the component labeling is not to your liking, you can right-click, select *Move Value*, and place the text in a better place.

Make sure all components in the schematic are completely referenced and provided with values before using *Tools | Annotate Schematic* and *Tools | Generate Netlist File* to generate the result. If necessary, print or plot the schematic from the File menu.

In the next step, use *Tools | Assign Component Footprint* to merge the components with their footprints. These geometric properties include the assignment of the connections and housing dimensions. The program returns error messages the first time it is called, but you can continue working after restarting this step.

On the left-hand side of the window are the component classes and in the middle are the components from the schematic; on the right-hand side you

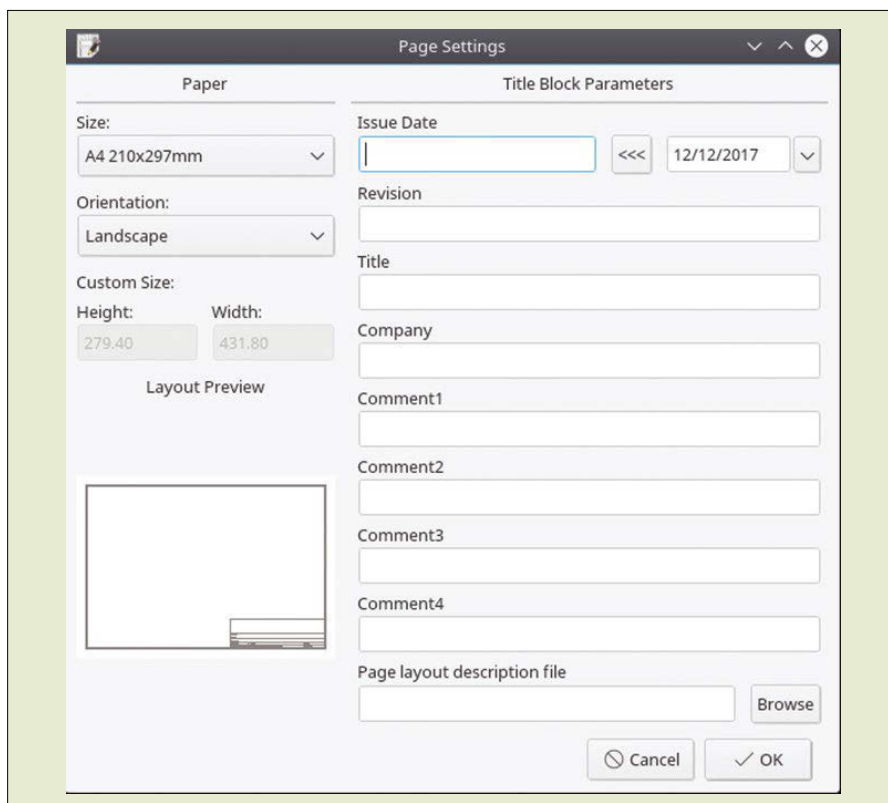


Figure 3: You will thank yourself later if you document your work well in the beginning.

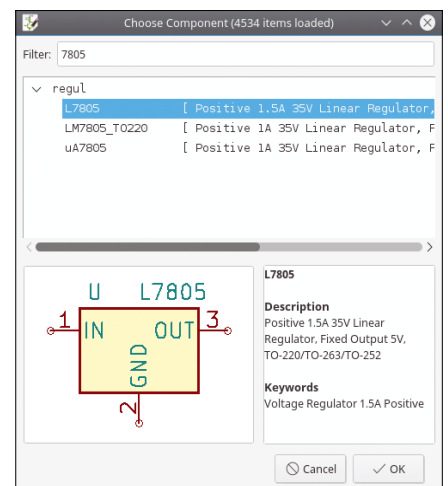


Figure 4: Each project starts by selecting the required components.

can select the matching footprint element (Figure 7). If you want to view the dimensions and contacts of a component

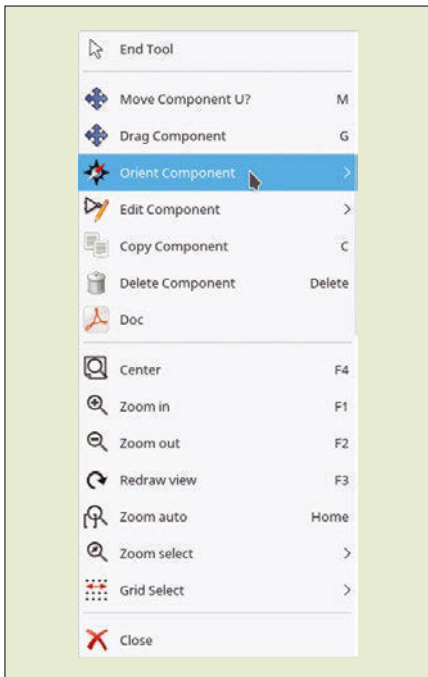


Figure 5: You can edit a component retroactively after placing it.

when selected, click on the third symbol from the left side of the window (*View selected footprint*). Figure 8 shows a diode used in the example.

Make sure you process all components before saving your work and closing the

assignment window. To be on the safe side, click on *Tools | Generate Netlist File* again and then click on *Tools | Layout Printed Circuit Board*.

Now set the grid, the units of measurement, and the grid width. To do

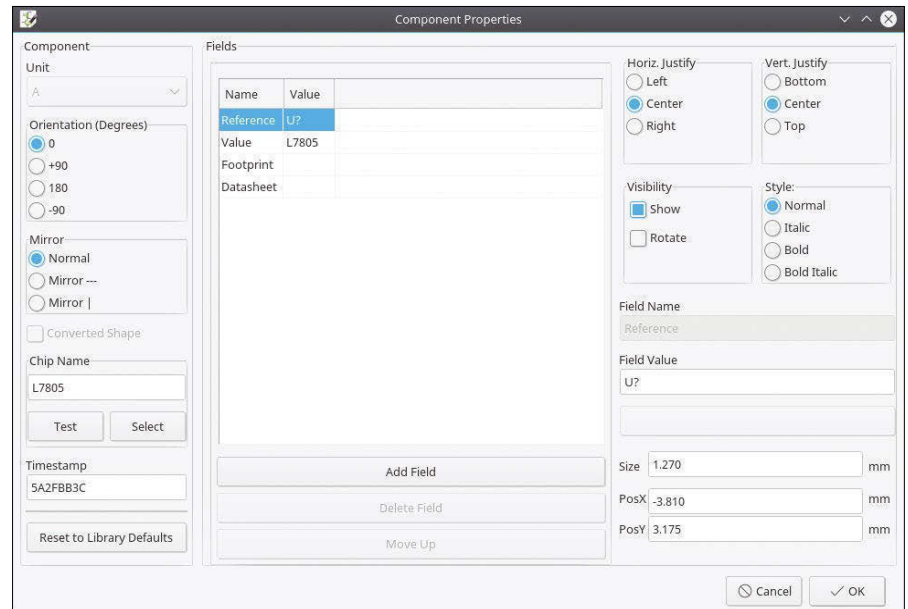


Figure 6: Gradually add suitable references and values to the individual components.

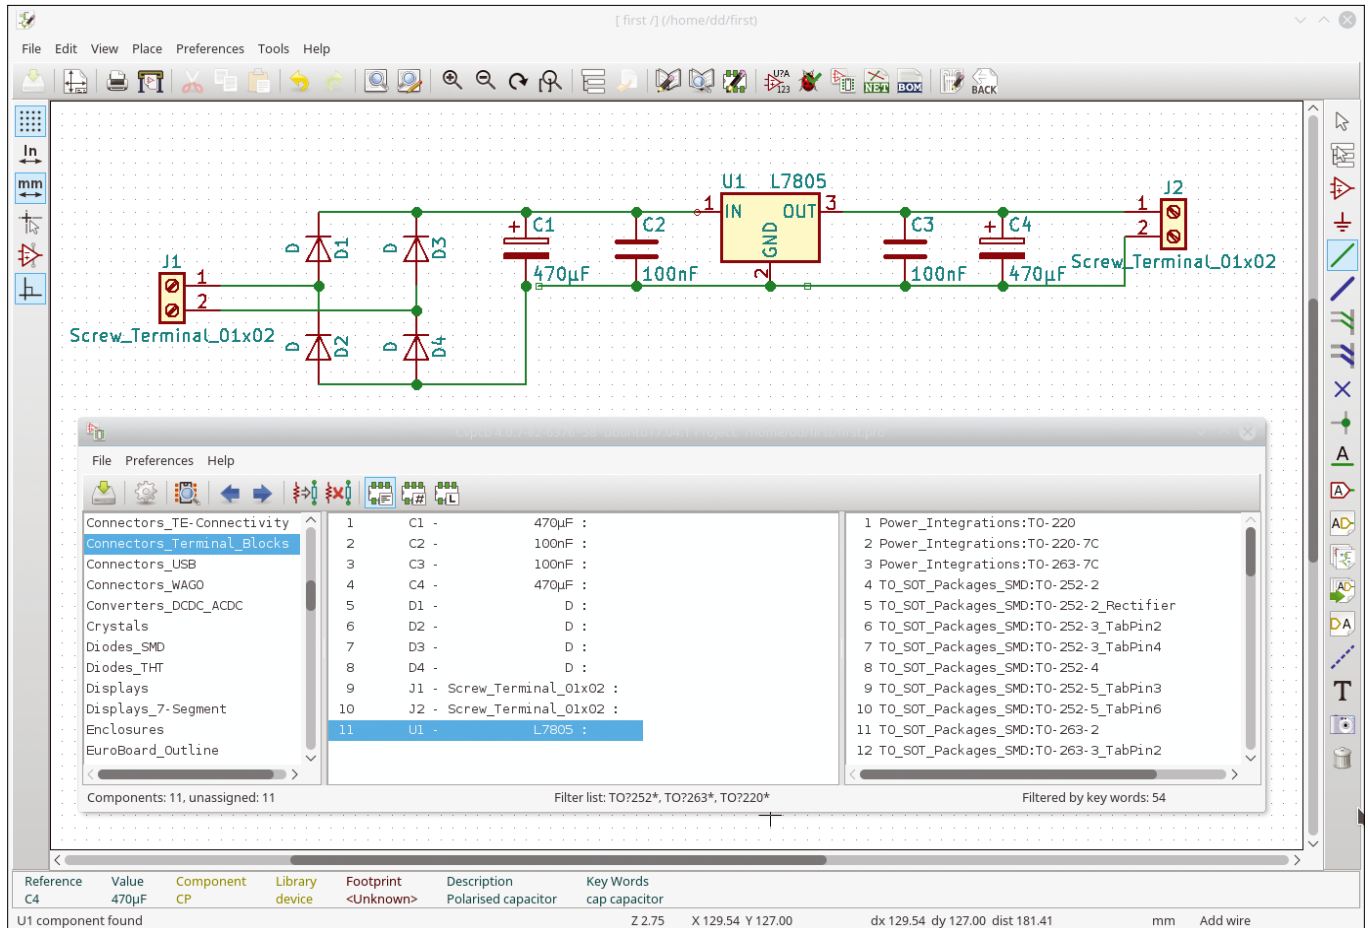


Figure 7: Once you have designed the entire circuit, connect the circuit symbols to the matching footprints.

this, use the *Show grid/Hide grid* icon (second down on the left) and the *Units in millimeters* icon for metric dimensioning (fifth from the top). The grid width can be changed under *Dimensions | Grid* or with the *Grid* drop-down along the top of the schematic window.

The next step is to determine the size of the board. To do this, start the layer manager (left toolbar, second icon from bottom). Set *Edge.Cuts* as the current layer by positioning the arrow to the left

of the entry; the corresponding checkbox must be filled (Figure 9).

Now draw the outline of the board. You can find this function under *Place | Line or Polygon* or by clicking on the eighth symbol from the top right of the layer manager.

To read the schematics data, use the *Tools | Netlist | Read Current Netlist* or the *Net* symbol at top. The components can be found as a bundle on the desktop. Now switch to Footprint mode (top toolbar, fourth button from the right).

Right-click in the free workspace to open the context menu and select *Global Spread and Place | Automatically Place All Footprints*. The program now

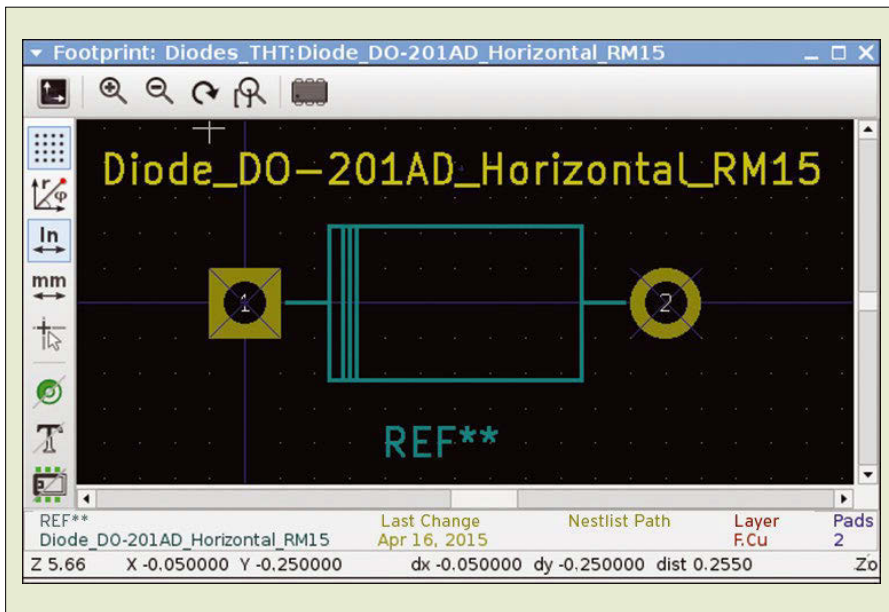


Figure 8: Upon request, KiCad can display the footprint data in detail.

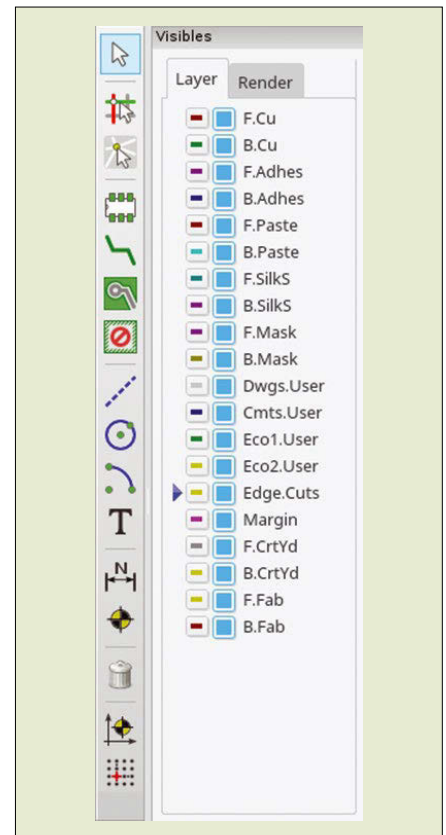


Figure 9: Use the layer manager to define the final size of the planned board.

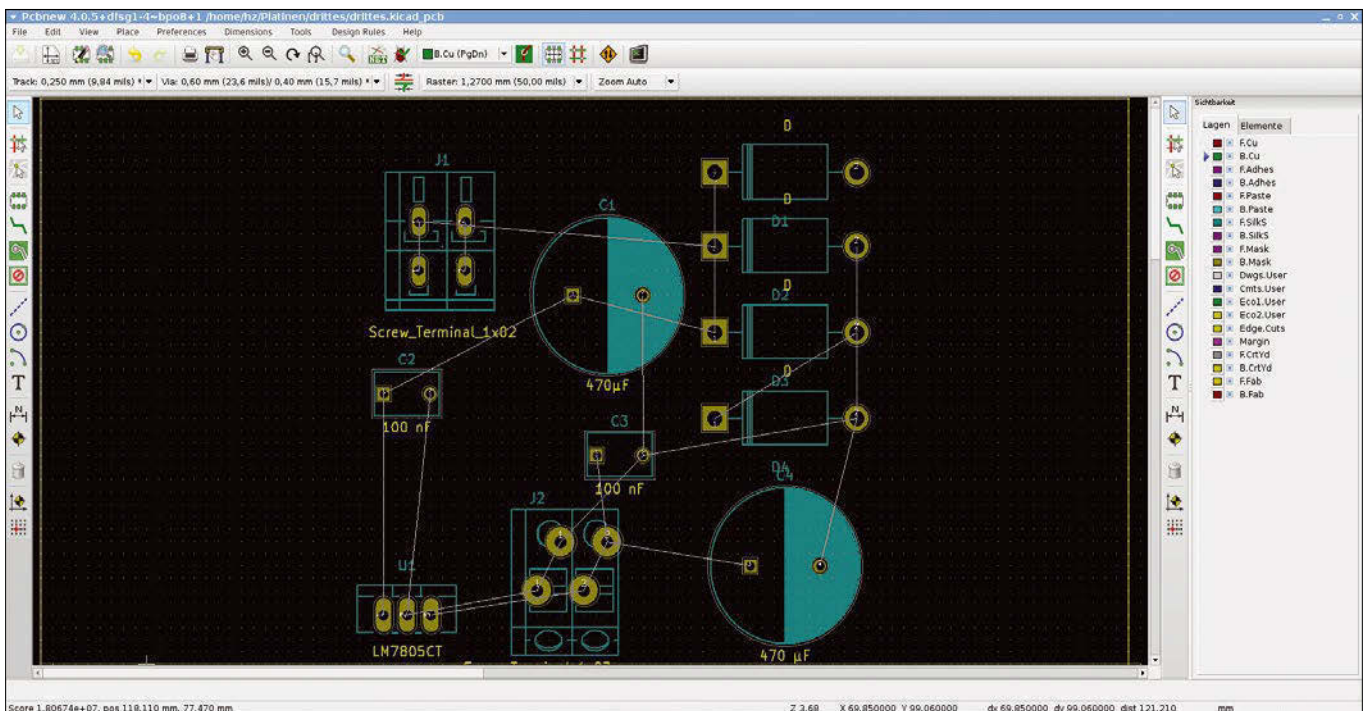


Figure 10: Automatic arrangement of the components does not always give you optimal results. If necessary, you can correct them manually.

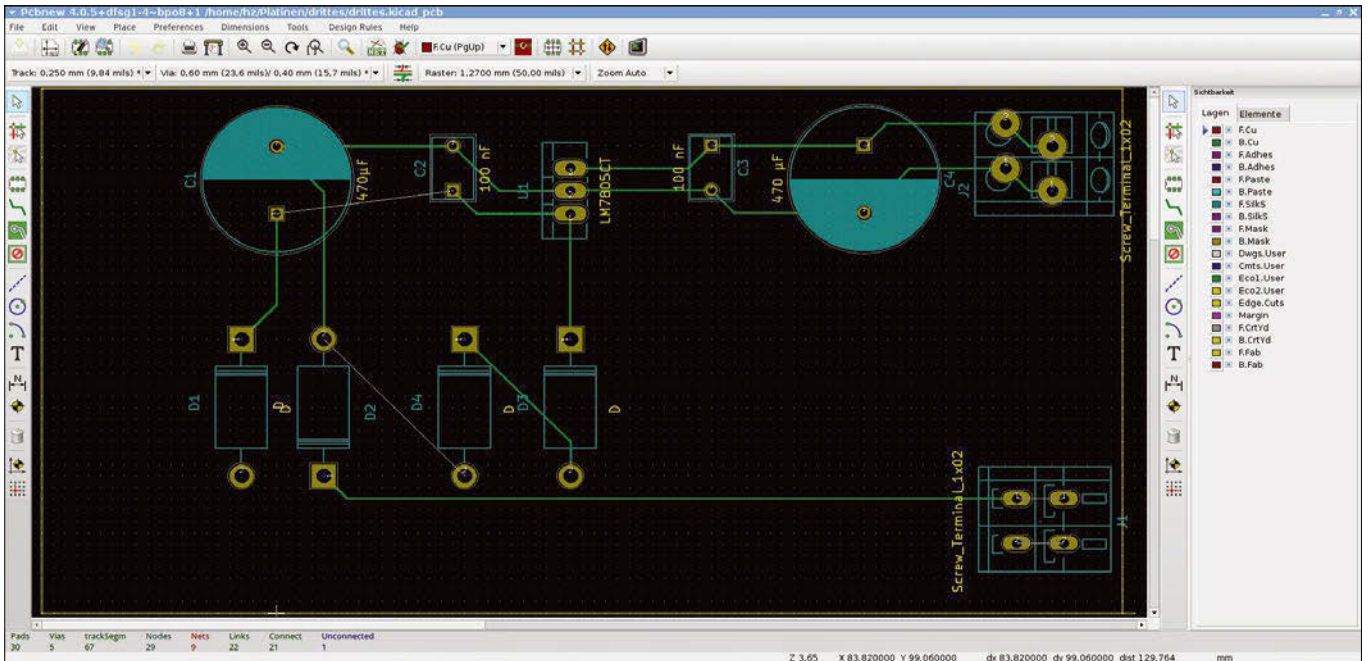


Figure 11: Once you have put all the elements in place, connect them using conducting paths whose properties you can edit later, if necessary.

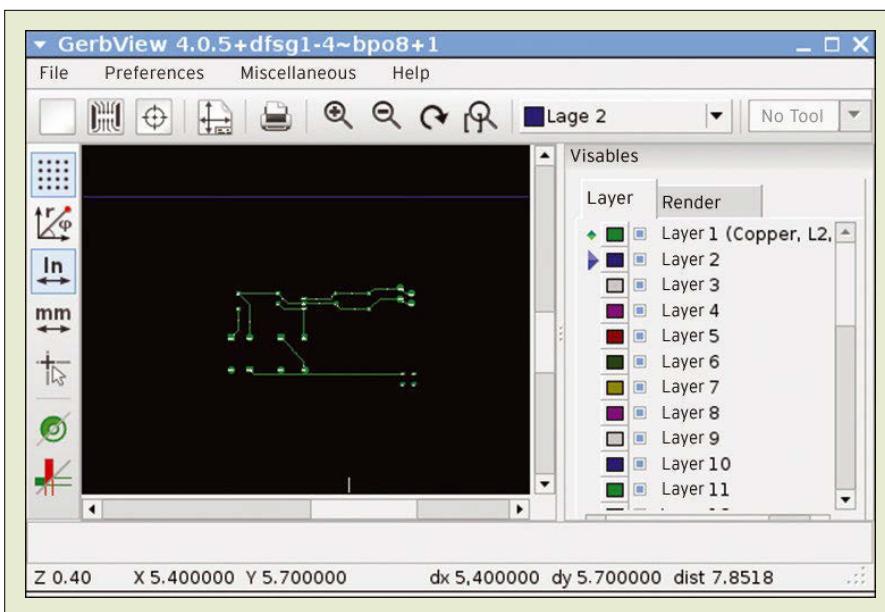


Figure 12: In GerbView, you can see whether all elements of the end product are in the right place on the PCB.

arranges the components more or less sensibly within the frame. To make your corrections, you might need space for heat sinks or wires.

To move a component, again choose Footprint mode (Figure 10) and right-click the component in question. A menu opens in which the top entry always starts with *Footprint*. Expand this menu item to move the component or drag, rotate, flip, and so on.

Always be sure to use the correct layer when using wired devices – use

the bottom layer (*B.Cu*) when dragging the conducting paths. You can change the thickness of the conducting paths or insert ground areas later (Figure 11).

Results

To print or plot the results, use the *File* menu. With the GerbView tool, you can create a true-to-scale proof (Figure 12). Place the components on top of it and check that they match the conducting paths, drill marks, and other elements. After this quality control,

send the Gerber files to a PCB service provider. To check the correct position of critical parts in advance, create a Gerber plot (*File | Plot*) in Pcbnew.

Conclusions


KiCad lets you design the layout for circuit boards from the first diode to the finished product. This complex process requires appropriate knowledge, but the program does not pose any major hurdles for an experienced user after a learning phase. Unlike commercial competitors, KiCad often offers a great deal of flexibility because of its Python interface. In terms of functions, it compares well to its free counterparts. ■■■

Info

- [1] KiCad: <http://kicad-pcb.org>
- [2] Component library: <http://kicad-pcb.org/libraries/download/>
- [3] Docs: <http://kicad-pcb.org/help/documentation/>

Author

Harald Zisler has worked with Linux and FreeBSD for many years. He writes magazine articles and books on the topics of technology and computing. The fourth edition of *Computer Networks* was recently published at Rheinwerk Publishing.



MakerSpace

Free-licensed and secure phones

You Say Goodbye, and I Say eelo

With the eelo project, Gaël Duval has big plans for a free-licensed phone, with an accompanying app store and online services. *By Bruce Byfield*

Suddenly, efforts to produce secure, free-licensed phones are everywhere. One of the most likely to succeed is eelo [1], a project that aims not only for a free phone, but the app store and online services to accompany it. The project's Kickstarter campaign [2] reached its original goal of EUR25,000 (~\$30,000) in less than a week, and, as I write with nine days left in the campaign, it has collected over twice that amount with EUR69,000 and has a strong chance of reaching EUR100,000 (~\$119,000).

One reason why eelo has a high chance of success is that it is being headed by Gaël Duval [3] (Figure 1). Long-time free software users may remember Duval as the founder of Mandrake (later Mandriva, and the ancestor of Mageia, PCLinuxOS, and OpenMandriva), the distribution that, around the turn of the millennium, led in Linux usability. Duval left Mandrake Software in 2006 and has since been a serial entrepreneur, involved in such startups as Ulteo [4], as well, he says, as investing “in a dozen startups last year.” Although some of Duval’s ven-

tures have had only modest success, they should give him a familiarity with business that is rare among those attempting to develop open hardware.

For Duval, eelo is a highly personal project. As he explains on his blog [5], in 2017, with over a decade as a dedicated iPhone and Mac OS user, “I realized that I had become lazy and that my data privacy had vanished. Not only wasn’t I using Linux anymore as my main operating system, but I was also using a proprietary OS on my smartphone. And I was using Google more and more. Search of course, but also Google Mail, Google Drive, and Google Docs. And Google Maps.” Without meaning to, he had entered “voluntary servitude,” and lost his privacy, with information leaking from all sorts of sources (Figure 2).

Duval decided that the way “to reconquer my privacy” was to create a nonprofit that would start to build an alternative. He named the project “eelo” as a reference to moray eels, which he describes as “fish that can hide in the sea” [6]. He also began work with Hathibelagal Ashraff, an Indian developer for mobile devices, and Rhandros Dembicky, a Brazilian interface designer, to produce a proof of concept study and to develop a roadmap.



Figure 1: Gaël Duval, the founder of the Mandrake distribution, is now creating the ambitious eelo project.

Eventually, Duval intends to create a nonprofit foundation that, over the next three years, will develop both a free phone as well as an ecosystem of services to support it. This emphasis on services is in marked contrast to the efforts of Purism’s Librem 5 phone [7], which is apparently being developed in the hopes that the combination of a Linux distribution’s package software repository and high-end security will be attractive enough that users will not miss the selection of apps offered by Google Play or iOS’ App Store. As Duval says, “I’m not sure that we are exactly in the same market as Purism.”

Instead, Duval has said several times that “I want something with more privacy for Mum and Dad.” By implication, that means something “attractively designed” and easy to use, so that any level of user can have security. Moreover, according to the crowdfunding page, at least in the beginning, eelo’s offering will be – as much as possible – open source, meaning that initially eelo may not meet Purism’s high standards for avoiding proprietary tools. However, Duval adds that “this will improve over time, step by step. Releasing the perfect thing from the beginning would be a utopia.”

Building an Infrastructure

Eelo’s first step is to build an operating system based on LineageOS [8], the fork of custom CyanogenMod ROM [9], one of the most comprehensive efforts to provide a free version of Android. However, while Duval describes the core of LineageOS as “usable and performing well,” he suggests that “the design is not very attractive and there are tons of micro-details that can be show-stoppers for a regular user. [...] Unless you are a geek, LineageOS is not realistically usable if you don’t want Google Inside” [10]. Launchers, icons, and notifications on the desktop and a control center are all areas that Duval intends to improve or add.

Before the campaign began, Duval’s team had already managed to get a proof-of-concept version of the eelo operating system running on one phone model, the LeEco Le2, and, as I write, was awaiting the arrival of another model for testing.

The campaign page also includes a link to Duval’s Twitter feed, so that would-be users can request that the project work to support other phones [11]. However, in the interest of time, this support means dealing with the proprietary drivers found in the average smart-

phone, which can be difficult to reverse engineer. In the short term, eelo will likely have to continue using proprietary firmware, although Duval does talk about creating some means of auditing this firmware to ensure that it is not leaking data secretly. Another alternative may be to partner with a project such as Fairphone [12], which is also attempting to produce a free phone and faces a similar limitation.

Other issues concern the support ecosystem for eelo’s operating system. For example, the project is considering replacing Google Play with app stores such as F-Droid [13] or APKPure [14], two app stores that specialize in free-licensed apps – and, in the long run, perhaps, with its own app store. Another alternative is to use Yalp Store [15], which acts as an anonymous proxy for Google Play users. Similarly, Google services may be replaced by microG [16], and Google’s SafetyNet Attestation API, which checks whether a device complies with Google’s environment with Magisk [17], as part of the infrastructure that will help eelo to run apps for unrooted phones. Search engines like DuckDuckGo or Qwant may also be used as replacements for Google, while multimedia services like Facebook will

be left for users to decide to install. In other cases, such as Gmail, eelo may have to build its own services.

The first round of such decisions and developments is scheduled for the project’s first year, with the first results due in the second quarter of 2018. In the second year, eelo plans to add its own web services where necessary and an eelo operating system for computers, as well as setting up the project’s nonprofit foundation. According to the roadmap, the third year will see the continued develop-




Figure 2: Eelo’s illustration of data leaks on a modern smartphone.



Figure 3: What the eelo phone may look like.

ment of these efforts, as well as an eelo smartphone (Figure 3) and digital assistant.

Duval blogs, “I want eelo to be a non-profit project ‘in the public interest’. I think operating systems and web services should be a common resource: as I explained a few year ago [...] this is infrastructure, like phone networks, rail tracks, roads. [...] Nonprofit doesn’t mean nothing will be for sale. Probably some eelo smartphone will be for sale, and some premium services will be available for corporates. But profit won’t be the first focus of eelo. Eelo will be for users first, for everyone who cares about their data privacy.”

Creating a New Market

Of all the attempts to build a free phone to date, eelo is by far the most

ambitious and comprehensive. While Duval says that he does not expect that an eelo smartphone will be “something people will find at Walmart,” the project is clearly being planned with the possibility of being more than a niche service in mind. “I think

that we can partner with some phone makers, because some of them will be interested that we do not have Google preinstalled,” says Duval. “And of course, there is a huge market in corporations who want to control their data better.”


How far eelo will go towards its goals remains to be seen. However, so far, eelo is receiving considerable attention, both from the media and potential supporters. “I’m very surprised, actually,” says Duval, “because most of the time I hear, ‘People don’t care.’ My point of view is that people do care when they know.”

At the very least, eelo has helped people to start talking. And who knows? When secure and free-licensed phones finally arrive, eelo just might be a major leader in a field it has helped to create. ■■■

Info

- [1] eelo: <https://www.eelo.io/#>
- [2] eelo’s Kickstarter campaign: <https://www.kickstarter.com/projects/290746744/eelo-a-mobile-os-and-web-services-in-the-public-in>
- [3] Gaël Duval: <https://www.indidea.org/gael/blog/>
- [4] Ulteo: <https://www.ulteo.com/> (french)
- [5] Duval’s blog: <https://www.indidea.org/gael/blog/leaving-apple-google-eelo-odyssey-introduction/>
- [6] How eelo got its name: <http://www.digitaltechinsider.com/mobile-technology/eelo-in-conversation-with-gael-duval-creator-of-the-google-less-android-project/>
- [7] Librem 5: <https://puri.sm/shop/librem-5/>
- [8] LineageOS: <https://lineageos.org>
- [9] CyanogenMod ROM: <http://www.cyanogenmodroms.com/>
- [10] The problem with LineageOS: <https://www.indidea.org/gael/blog/leaving-apple-google-eelo-odyssey-part1-mobile-os/#more-1031>
- [11] Duval’s Twitter page: https://twitter.com/gael_duval
- [12] Fairphone: <https://www.fairphone.com/en/>
- [13] F-Droid: <https://f-droid.org/>
- [14] APKPure: <https://apkpure.com/app>
- [15] Yalp store: <http://www.droidviews.com/yalp-store-google-play-store-alternative-to-install-apps-on-android/>
- [16] microG: <https://microg.org/>
- [17] Magisk: <https://magiskmanager.com/>

7 Years of ADMIN on One DVD



This searchable DVD gives you 40 issues of ADMIN, the #1 source for:

- systems administration
- security
- monitoring
- databases
- and more!



ORDER NOW!
shop.linuxnewmedia.com



Shop the Shop

shop.linuxnewmedia.com

GIMP HANDBOOK



**SURE YOU
KNOW LINUX...**
but do you know **Gimp?**

- Fix your digital photos
- Create animations
- Build posters, signs, and logos

Order now and become an expert in one of the most important and practical open source tools!

Gimp
Doctor your digital images with the
AMAZING FREE PHOTO EDITOR!

Order online:
shop.linuxnewmedia.com/specials



FOR WINDOWS, MAC OS, AND LINUX USERS!



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

Expand your Linux skills with:

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

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

Subscribe now!
shop.linuxnewmedia.com/subs

GET IT NOW!
FAST DELIVERY WITH OUR PDF EDITION


The **apt-get** command is a well known and much-loved tool for many users of Debian-based Linux systems. **apt-get** is the go-to utility for installing and managing applications at the command line – or at least it used to be. A new utility called **apt** solves some of the problems and inconsistencies that **apt-get** users have learned to ignore. This month we introduce you to **apt** and show you some of the differences between **apt** and **apt-get**. We also show you how to profile a monitor for better color reproduction, and we delve into a little known utility called USB/IP that lets you access USB-connected devices attached to other computers on your network.



Image © Alexandr Moroz, 123RF.com

LINUX VOICE

Doghouse – Security <i>Jon “maddog” Hall</i> A serious security problem requires your attention.	72
apt vs. apt-get <i>Ferdinand Thommes</i> The apt command-line utility is a successor to the well-known apt-get, offering simpler installation and maintenance for the DEB packages used with Debian, Ubuntu, Knoppix, and many other Linux distros.	74
DisplayCAL <i>Karsten Günther</i> If you’re tired of inaccurate colors on your monitor, try calibrating and profiling using a colorimeter and DisplayCAL.	78
FOSSPicks <i>Graham Morrison</i> Ocenaudio 3.3.6, Otter Browser, Joplin, WeeChat 2.0, Mailspring, Siril 0.9.7, SuperTuxKart 0.9.3, and more!	84
Tutorial – USB/IP <i>Paul Brown</i> USB/IP lets you use USB devices connected to other machines as if they were plugged directly into your computer.	92



MADDOG'S DOGHOUSE

A serious security problem requires your attention.

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®.

Meltdown and Spectre

I was lying in my bed in the early morning of January 3rd, 2018 when a tweet woke me from my sleep; not a tweet from a bird (nor from the president of the United States), but a tweet from a security researcher discussing two rather bad hardware issues with a large number of CPU chips.

These issues, now known as "Meltdown" and "Spectre," have been rocking the computing world for the past couple of days. Unlike many other security exploits, these are not really "fixable" by a simple software patch, are not operating-system specific, and cannot be avoided by telling your mother not to post her password on her computer screen.

Both have to do with modern hardware architecture and an issue called "out-of-order instruction execution," used to speed up the processor. Sometimes this feature is used to fetch instructions on both sides of a branch (both the "true" and the "false" side), so as soon as the condition is known, the instructions are ready to execute. Access to this "pre-fetched" data could allow a carefully crafted user-level program to access kernel memory, and once that happens, any data on the machine is vulnerable to be read, including passwords, security certificates, and so forth.

At first this was painted as a problem mostly or even solely with Intel processors, but as people investigated it was projected that at least some AMD and some ARM processors are affected. AMD has now stated that none of its processors have this problem. It also depends on when the processor was designed (starting in 1965) and what class of processor it is (Intel 64-bit processors seem to be prime culprits), as not all processors perform out-of-order execution.

The bad news is that the more powerful and expensive the processor is (think server systems, high-end desktops, gaming systems, and superior grade notebooks), the more likely it is to have this feature. Cloud server systems are particularly vulnerable, since they typically run lots of applications at one time.

It is conceivable that even applets and web-based applications could trigger this type of exploit, although the Apache Software Foundation (as an example) has taken steps to make sure that Apache based applications cannot exploit this by restricting access to the high-precision timers that can be used for the exploit. Google has also applied these types of fixes to their systems.

As with many security exploits, this was "embargoed" by the security research groups as people scrambled, first to understand the problem and then to find a solution. Eventually, as solutions were found, the problem was "leaked," which resulted in my early morning tweet. Before I wrote anything or contacted anyone, I checked sources for the "leak," and, feeling assured that the information was both true and correct, I sent out messages to other people I knew.

Understand that these are reasonably hard exploits to utilize, and the application trying to exploit them needs to be on the local machine. But in the day of shared cloud servers, the heavy use of containers, and web-based applications, this is not a rare occurrence.

There are now tens of thousands of people in the world who have the expertise to exploit this, and perhaps hundreds who would have the expertise and the desire. With today's Internet, you really only need a few such people.

The Linux kernel has released a patch that stops most of the exploits, and Microsoft and Apple will do the same. By the time you read this, your favorite distribution should have a replacement kernel ready for you.

Lack of detailed information about which processor is in your server, desktop, tablet, or device, and whether that processor is affected by this exploit, will make it harder to determine if you need the patches. This is a problem on two fronts.


Current patches in the Linux kernel (and I assume in other affected operating systems) cause a slowdown. The slowdown occurs because the kernel no longer shares address space with user programs, and this causes overhead for the hardware as the system shifts back and forth between kernel memory and user memory during interrupts and system calls. The percentage of slowdown depends on the programs being executed, and typically heavy I/O programs (such as database engines or network-heavy programs) take the brunt of the slowdown.

You can partially disable these patches by booting with `nopt i` as a kernel boot option if your environment is secure or you are using a processor that is unaffected.

This exploit has roots back to 1995; there are thousands of systems still being used that will not get their kernel updated. Desktops running Microsoft Windows XP (yes, they are still out there), as well as older Apple systems, are just two examples.

Caveat emptor. ■■■

A Webzine for High-Performance Computing Specialists



If you work with high-performance clusters, or if you're ready to expand your skill set with how-to articles, news, and technical reports on HPC technology.

ADMIN
Network & Security

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

apt vs. apt-get

The `apt` command-line utility is a successor to the well known `apt-get`, offering simpler installation and maintenance for the DEB packages used with Debian, Ubuntu, Knoppix, and many other Linux distros. **BY FERDINAND THOMMES**

The Debian package management infrastructure is a simple yet powerful system that has been a distinguishing feature of Debian since the project's early years. At the core of the Debian package management is `dpkg` [1], a low-level tool used for installing and removing Debian `.deb` packages. The bigger and more comprehensive Advanced Package Tool (APT) [2] is a higher-level tool that includes `dpkg` on the back end. APT performs most of the other functions associated with a package management system, including integrating external repositories.

The APT package system supports a number of front-end applications, both command-line and GUI-based, that act as a user-friendly interface for managing Debian packages. In the past, most users who wanted to work at the command line depended on the classic `apt-get` tool as an interface to the Debian package system. Through the years, `apt-get` and a constellation of other supporting utilities, such as `apt-cache`, have provided

nearly effortless package management for Debian users, as well as users of Debian-based distros like Ubuntu and Knoppix, who take the time to master the commands.

However, some users were not happy about the complex, and often confusing, system of `apt-get` commands and options. Since the release of Debian 8 "Jessie" and Ubuntu 16.04 "Xenial Xerus," all users of Debian-based systems can enjoy a new command-line interface to the APT system: `apt`. The `apt` utility consolidates features that were once spread among `apt-get`, `apt-cache`, and other commands, and it simplifies many of the command options for an easier and friendlier user experience.

The `apt` utility has been used in the Debian "Unstable" edition since 2014, and it became official with Debian 8. Because `apt-get` still works in recent releases of Debian and other Debian-based distros, many users have not taken the time to learn about `apt`. This article introduces the `apt`

Table 1: Differences

Function	<code>apt-get</code>	<code>apt</code>
Install package	<code>apt-get install <package></code>	<code>apt install <package></code>
Remove package	<code>apt-get remove <package></code>	<code>apt remove <package></code>
Remove package including configuration	<code>apt-get purge <package></code>	<code>apt purge <package></code>
Update sources	<code>apt-get update</code>	<code>apt update</code>
Update packages (without removing or reinstalling)	<code>apt-get upgrade</code>	<code>apt upgrade</code> ¹
Update packages (with removing and reinstalling)	<code>apt-get dist-upgrade</code>	<code>apt full-upgrade</code>
Remove unnecessary dependencies	<code>apt-get autoremove</code>	<code>apt autoremove</code>
Search package	<code>apt-get search <package></code>	<code>apt search <package></code>
Display package information	<code>apt-cache show <package></code>	<code>apt show <package></code>
Display active package sources in detail	<code>apt-cache policy</code>	<code>apt policy</code>
Display available and installed package versions	<code>apt-cache policy <package></code>	<code>apt policy <package></code>
New Commands		
Edit packages sources	–	<code>apt edit-sources</code>
List packages by criteria	<code>dpkg --get-selections > list.txt</code>	<code>apt list</code>
Set/change package status	<code>echo <package> hold dpkg --set-selections</code>	<code>apt-mark <package></code>
¹ Corresponds to <code>apt-get upgrade --install new-pkgs</code>		

package management tool and explores some of the differences between `apt` and `apt-get`.

Note: The `apt` tool discussed in this article is not the same as the `apt` Python wrapper created by Linux Mint developers a few years ago.

apt vs apt-get

`apt` does not guarantee downwards-compatibility with `apt-get`, but many (though not all) command options are interchangeable. For a full comparison, type `apt --help` and `apt-get --help` and compare the results.

Table 1 shows some of the important commands side by side. As you can see, many commands are the same if you just exchange `apt-get` for `apt`.

One exception is the `apt upgrade` command. The old command: `apt-get upgrade` installed the latest versions of all packages currently installed on the system from the repository, but it didn't remove existing packages or retroactively add dependencies. The new `apt upgrade` installs new packages when they are added as

```
root@dd-ubu1704d-x64-vm: /home/dd
root@dd-ubu1704d-x64-vm:/home/dd# apt-get update
Hit:1 http://security.ubuntu.com/ubuntu zesty-security InRelease
Hit:2 http://us.archive.ubuntu.com/ubuntu zesty InRelease
Hit:3 http://us.archive.ubuntu.com/ubuntu zesty-updates InRelease
Hit:4 http://us.archive.ubuntu.com/ubuntu zesty-backports InRelease
Reading package lists... Done
root@dd-ubu1704d-x64-vm:/home/dd#
```

Figure 1: Use `apt-get update` to update the lists using the packages provided by the repository, but do not expect to see any additional information.

```
root@dd-ubu1704d-x64-vm: /home/dd
root@dd-ubu1704d-x64-vm:/home/dd# apt-get update
Hit:1 http://security.ubuntu.com/ubuntu zesty-security InRelease
Hit:2 http://us.archive.ubuntu.com/ubuntu zesty InRelease
Hit:3 http://us.archive.ubuntu.com/ubuntu zesty-updates InRelease
Hit:4 http://us.archive.ubuntu.com/ubuntu zesty-backports InRelease
Reading package lists... Done
root@dd-ubu1704d-x64-vm:/home/dd#
```

Figure 2: The `apt update` command not only updates the package lists, but also shows you if and how many new versions of the packages exist in the repository.

```
root@dd-ubu1704d-x64-vm: /home/dd
root@dd-ubu1704d-x64-vm:/home/dd#
root@dd-ubu1704d-x64-vm:/home/dd# apt list --upgradable
Listing... Done
appport/zesty-updates,zesty-updates 2.20.4-0ubuntu4.9 all [upgradable from: 2.20.4-0ubuntu4]
appport-gtk/zesty-updates,zesty-updates 2.20.4-0ubuntu4.9 all [upgradable from: 2.20.4-0ubuntu4]
apt-transport-https/zesty-updates 1.4.6-17.04.1 amd64 [upgradable from: 1.4]
bind9-host/zesty-updates 1:9.10.3.dfsg.P4-10.1ubuntu5.3 amd64 [upgradable from: 1:9.10.3.dfsg.P4-10.1ubuntu4]
bluemoon-cups/zesty-updates,zesty-security 5.43-0ubuntu1.1 amd64 [upgradable from: 5.43-0ubuntu1]
bluemoon-obexd/zesty-updates,zesty-security 5.43-0ubuntu1.1 amd64 [upgradable from: 5.43-0ubuntu1]
ca-certificates/zesty-updates,zesty-updates,zesty-security,zesty-security 20170717~17.04.1 all [upgradable from: 20161130]
cracklib-runtime/zesty-updates 2.9.2-3ubuntu1 amd64 [upgradable from: 2.9.2-3]
curl/zesty-updates,zesty-security 7.52.1-4ubuntu1.4 amd64 [upgradable from: 7.52.1-4ubuntu1]
distro-info-data/zesty-updates,zesty-updates 0.33ubuntu0.3 all [upgradable from: 0.33]
dnstools/zesty-updates 1:9.10.3.dfsg.P4-10.1ubuntu5.3 amd64 [upgradable from: 1:9.10.3.dfsg.P4-10.1ubuntu4]
evince/zesty-updates,zesty-security 3.24.0-0ubuntu1.3 amd64 [upgradable from: 3.24.0-0ubuntu1]
evince-common/zesty-updates,zesty-updates,zesty-security,zesty-security 3.24.0-0ubuntu1.3 all [upgradable from: 3.24.0-0ubuntu1]
file/zesty-updates,zesty-security 1:5.29-3ubuntu0.1 amd64 [upgradable from: 1:5.29-3]
firefox/zesty-updates,zesty-security 57.0.1+build2-0ubuntu0.17.04.1 amd64 [upgradable from: 52.0.1+build2-0ubuntu1]
fwupd/zesty-updates 0.8.1-3ubuntu0.1 amd64 [upgradable from: 0.8.1-3]
gdb/zesty-updates,zesty-security 7.12.50.20170314-0ubuntu1.1 amd64 [upgradable from: 7.12.50.20170314-0ubuntu1]
gdbserver/zesty-updates,zesty-security 7.12.50.20170314-0ubuntu1.1 amd64 [upgradable from: 7.12.50.20170314-0ubuntu1]
ghostscript/zesty-updates,zesty-security 9.19-0ubuntu7.6 amd64 [upgradable from: 9.19-0ubuntu7.1]
ghostscript-x/zesty-updates,zesty-security 9.19-0ubuntu7.6 amd64 [upgradable from: 9.19-0ubuntu7.1]
gir1.2-gdkpixbuf-2.0/zesty-updates,zesty-security 2.36.5-3ubuntu0.2 amd64 [upgradable from: 2.36.5-3]
gir1.2-javascriptcoregtk-4.0/zesty-updates,zesty-security 2.18.3-0ubuntu0.17.04.1 amd64 [upgradable from: 2.16.1-1]
gir1.2-soup-2.4/zesty-updates,zesty-security 2.56.0-2ubuntu0.1 amd64 [upgradable from: 2.56.0-2]
gir1.2-webkit2-4.0/zesty-updates,zesty-security 2.18.3-0ubuntu0.17.04.1 amd64 [upgradable from: 2.16.1-1]
gnome-calendar/zesty-updates 3.24.2-0ubuntu0.1 amd64 [upgradable from: 3.24.0-0ubuntu1]
gnome-desktop3-data/zesty-updates,zesty-updates 3.24.2-0ubuntu0.1 all [upgradable from: 3.24.0-0ubuntu1]
gnome-settings-daemon-schemas/zesty-updates,zesty-updates 3.24.2-0ubuntu0.1 all [upgradable from: 3.24.0-0ubuntu2]
gnome-software/zesty-updates 3.22.7-0ubuntu3.17.04.7 amd64 [upgradable from: 3.22.7-0ubuntu3]
gnome-software-common/zesty-updates,zesty-updates 3.22.7-0ubuntu3.17.04.7 all [upgradable from: 3.22.7-0ubuntu3]
gnome-software-plugin-snap/zesty-updates 3.22.7-0ubuntu3.17.04.7 amd64 [upgradable from: 3.22.7-0ubuntu3]
grub-common/zesty-updates 2.02-beta3-4ubuntu2.2 amd64 [upgradable from: 2.02-beta3-4ubuntu2]
grub-pc/zesty-updates 2.02-beta3-4ubuntu2.2 amd64 [upgradable from: 2.02-beta3-4ubuntu2]
grub-pc-bin/zesty-updates 2.02-beta3-4ubuntu2.2 amd64 [upgradable from: 2.02-beta3-4ubuntu2]
grub2-common/zesty-updates 2.02-beta3-4ubuntu2.2 amd64 [upgradable from: 2.02-beta3-4ubuntu2]
imagemagick/zesty-updates,zesty-security 8:6.9.7.4+dfsg-3ubuntu1.2 amd64 [upgradable from: 8:6.9.7.4+dfsg-
```

Figure 3: Use `apt list --upgradable` to provide some color for a better overview of the packages for which the repository provides a more recent version.

```

root@dd-ubu1704d-x64-vm: /home/dd
Preparing to unpack .../apt_1.4.6-17.04.1_amd64.deb ...
Unpacking apt (1.4.6-17.04.1) over (1.4) ...
Setting up apt (1.4.6-17.04.1) ...
Created symlink /etc/systemd/system/timers.target/wants/apt-daily-upgrade.timer
-> /lib/systemd/system/apt-daily-upgrade.timer.
(Reading database ... 168839 files and directories currently installed.)
Preparing to unpack .../apt-utils_1.4.6-17.04.1_amd64.deb ...
Unpacking apt-utils (1.4.6-17.04.1) over (1.4) ...
Preparing to unpack .../xserver-common_2%3a1.19.3-1ubuntu1.3_all.deb ...
Unpacking xserver-common (2:1.19.3-1ubuntu1.3) over (2:1.19.3-1ubuntu1) ...
Preparing to unpack .../libsystemd0_232-21ubuntu7.1_amd64.deb ...
Unpacking libsystemd0:amd64 (232-21ubuntu7.1) over (232-21ubuntu2) ...
Setting up libsystemd0:amd64 (232-21ubuntu7.1) ...
(Reading database ... 168839 files and directories currently installed.)
Preparing to unpack .../0-xserver-xorg-core_2%3a1.19.3-1ubuntu1.3_amd64.deb ...
Unpacking xserver-xorg-core (2:1.19.3-1ubuntu1.3) over (2:1.19.3-1ubuntu1) ...
Preparing to unpack .../1-libnss-resolve_232-21ubuntu7.1_amd64.deb ...
Unpacking libnss-resolve:amd64 (232-21ubuntu7.1) over (232-21ubuntu2) ...
Preparing to unpack .../2-libpam-systemd_232-21ubuntu7.1_amd64.deb ...
Unpacking libpam-systemd:amd64 (232-21ubuntu7.1) over (232-21ubuntu2) ...
Preparing to unpack .../3-systemd_232-21ubuntu7.1_amd64.deb ...
Unpacking systemd (232-21ubuntu7.1) over (232-21ubuntu2) ...
Progress: [ 7%] [#####.....]
    
```

Figure 4: apt full-upgrade uses a new element in the form of a progress bar to inform users about the upgrade's progress.

dependencies of packages to be updated, although it still doesn't remove previously installed packages.

The new command structure offers subtle improvements that save time and extra steps. For instance, apt-get lets you move a package to /var/cache/apt/archives/ and then install the package using apt-get install package_name, and the software automatically manages the dependencies. apt eliminates the need to push things around; all you need is apt install package_name, specifying the full path if necessary.

Design Error Corrected

apt fixes some errors of the original apt-get implementation. For example, in addition to apt-get for editing Debian packages, the legacy system used the apt-cache command for outputting information about the packages. The new apt command thoroughly cleans up, combining functions from both old commands and structuring them in a better way. apt provides features of the most commonly used apt-get and apt-cache commands, adding useful information that once required additional options.

For example, apt-get update only reports execution at the end (Figure 1), but apt update shows

the number of updatable packages and offers an option to display a structured list of possible updates (Figure 2). This list even provides a better overview, adding color to the otherwise drab display (Figure 3).

The actual package update command has been given a new visual element: Both apt upgrade and apt full-upgrade or apt dist-upgrade now show a progress bar that informs users of the update progress (Figure 4).

The update bar also appears when removing packages with apt remove or apt purge. apt remove removes the package itself, and apt purge removes the entire associated configuration – if it is located outside the home directory: apt never tampers with configurations in the user directory.

Two New Commands

The apt show command sorts the output alphabetically and suppresses some less important information that you might see with apt-cache show. The dist-upgrade parameter is assigned the more significant name of full-upgrade, although the two commands behave in the same way.

The apt list and apt edit-sources commands are new: apt list in connection with the --installed or --upgradeable options outputs lists of the installed or upgradeable packages; apt edit-sources <List> opens /etc/etc/apt/sources.list.d/ with the specified list in the editor for editing (Figure 5).

Conclusions

It was high time to detox and dust off Debian's command-line package management tool. The new apt utility has been available since 2014, but habits change slowly, and apt has still not completely replaced apt-get in the real world. The differences between apt-get and apt still cause confusion in Internet forums, especially on Ubuntu and Linux Mint. Switching to apt is definitely worthwhile: The new command is faster, more logically structured, and easier to use. ■■■

Info

- [1] dpkg: https://en.wikipedia.org/wiki/Debian_Package_Manager
- [2] APT: https://en.wikipedia.org/wiki/Advanced_Packaging_Tool

```

GNU nano 2.8.7                               Datei: /etc/apt/s
# debian loadbalancer
deb http://ftp.debian.org/debian/ unstable main contrib non-free
# deb-src http://httpredir.debian.org/debian/ unstable main contrib non-free
# deb http://ftp.debian-ports.org/debian unstable main
deb http://incoming.debian.org/debian-build builddd-unstable main contrib non-free
deb http://ftp.debian.org/debian/ testing main contrib non-free
# deb http://ftp.debian.org/debian/ experimental main contrib non-free
    
```

Figure 5: The apt tool opens the specified file for manual changes in the standard editor with the command apt edit-sources debian. list.

Celebrating 25 Years of Linux!

**ORDER NOW
and SAVE 25%**
on 7 years
of *Ubuntu User!*



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

ENTER CODE: SAVE25



shop.linuxnewmedia.com

Color Coordinated with DisplayCAL

If you work with digital images, you know the colors on your monitor don't always match what you see in real life. If you're looking for a more accurate representation of color, try calibrating and profiling your monitor using a colorimeter and DisplayCAL. **BY KARSTEN GÜNTHER**

Anyone who has ever tried to reconcile an electronic image with nature will be familiar with the problem: The colors in the image almost always differ from what you see with your eyes. This problem has several causes. Colors are lost on the way from the camera to the image. The human eye adapts much better than a camera to different lighting conditions and automatically supplements missing information. Another problem is that color values often shift as the image makes its way through the chain of devices, from the camera, to the monitor, and finally the printer.

Each device can only absorb and process colors to a limited extent. Image-processing experts use the term gamut [1] to describe the possible colors a device can produce by internal mixing. Colors outside the gamut appear in replacement colors, which leads to distortions.

Different colors have different effects on the human eye. The eye is particularly sensitive to tones in the green range and can detect far more nuances of green in nature than a monitor shows. For mixed colors such as yellow – composed of green and red on the monitor – the printer may sometimes be better than the monitor assuming that this color is available as a process color.

What the brain perceives as “white” is defined by a whole series of factors, not just the perceived wavelengths. Cultural conditions influence what our brain recognizes as “pure white.” For example, Asians prefer red shades of white, and Central Europeans prefer blue.

The differences in the way colors are perceived, combined with natural limitations of electronic devices and the imprecise replacement of colors outside the gamut, mean the colors on your monitor will never exactly match the colors you see with your eye. However, it is possible to bring the truest possible color reproduction to your electronic device using a tool called a colorimeter.

A colorimeter is a specialized hardware tool that is used to profile the color properties of a monitor or other display device. The colorimeter attaches to the computer, typically through a USB port, and hangs in front of the monitor. Colorimetry software running on the computer then proceeds through a series of checks, using the colorimeter to measure the colors produced by the monitor. The software running on the computer then produces a color profile in ICC format (see the box entitled “ICC Profile”). The profile tells the system how to adjust the color settings to provide truer color representation. After you load the generated color profile to the graphics card, the colors will display more accurately on the monitor.

Calibrate and Profile

In order to achieve the truest possible color reproduction, the monitor must first be calibrated and then profiled. Calibration is a preliminary step that sets certain neutral properties for the monitor.

Until a few years ago, calibration and profiling in Linux were performed in the terminal window using rather cryptic tools of the ArgyllCMS [2] color management system. Today, you can use DisplayCAL [3], which acts as a front end for the ArgyllCMS system.

DisplayCAL, which is available in the package repositories for many popular Linux distributions, works with several leading colorimeters (see the box entitled “Supporting Colorimeters.”) Colorimeters are available in different price categories; be sure to use some common sense when choosing a tool: You don't need a EUR300 (~\$366) measuring device to profile a EUR100 (~\$122) monitor. If you

ICC Profile

The International Color Consortium (ICC) maintains a specification for defining the color attributes for a device and mapping the transition from a source to a target color space.

prefer open hardware and open source software, it is definitely worth taking a look at ColorHug (see the "All Free" box).

Keep in mind that, in older colorimeters, the filters in the cameras age, which has an effect on quality. Always store the colorimeter in a dark box to reduce this effect.

DisplayCAL in Action

By default, DisplayCAL profiles quite extensively and more thoroughly than the standard software

supplied with most colorimeters. The process takes considerably longer, but it ultimately achieves better results.

Install the colorimeter on the USB port (not via a passive hub), and then start the software. After the logo has been displayed, the user interface opens (Figure 1).

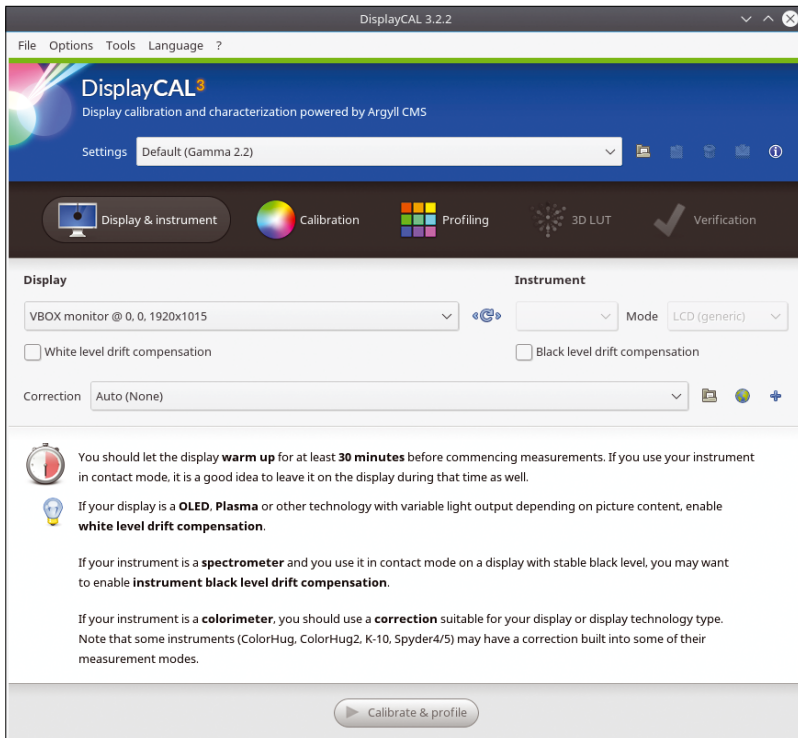


Figure 1: DisplayCAL starts with a few dialogs that provide you with important information on how to use the software.

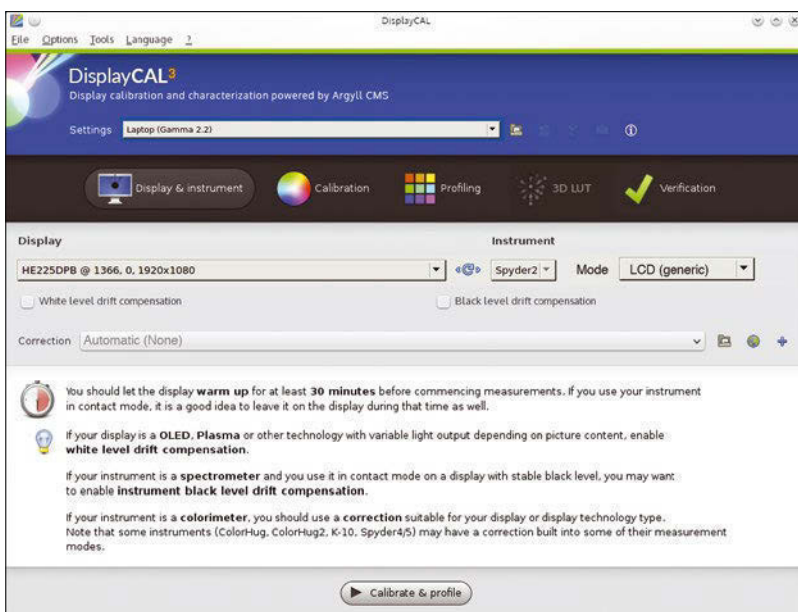


Figure 2: The first settings for calibration are made using the *Display* and *Instrument* drop-down menus.

All Free

The problems with proprietary software in commercial products led to the ColorHug project in 2011. ColorHug is a completely free colorimeter that works under Linux [4]. The same manufacturer has produced a professional version of the device known as ColorHug+, but it has a similarly professional price tag.

Supported Colorimeters

According to the project documentation, DisplayCAL supports the following colorimeters:

- CalMAN X2 (treated as i1 Display 2)
- Datacolor/ColorVision Spyder 2
- Datacolor Spyder 3 (since ArgyllCMS 1.1.0)
- Datacolor Spyder 4 (since ArgyllCMS 1.3.6)
- Datacolor Spyder 5 (since ArgyllCMS 1.7.0)
- Hughski ColorHug (Linux support since ArgyllCMS 1.3.6, Windows support with newest ColorHug firmware since ArgyllCMS 1.5.0; fully functional Mac OS X support since ArgyllCMS 1.6.2)
- Hughski ColorHug2 (since ArgyllCMS 1.7.0)
- Image Engineering EX1 (since ArgyllCMS 1.8.0)
- Klein K10-A (since ArgyllCMS 1.7.0. The K-1, K-8, and K-10 are also reported to work)
- Lacie Blue Eye (treated as i1 Display 2)
- Sencore ColorPro III, IV, & V (treated as i1 Display 1)
- Sequel Imaging MonacoOPTIX/Chroma 4 (treated as i1 Display 1)
- X-Rite Chroma 5 (treated as i1 Display 1)
- X-Rite ColorMunki Create (treated as i1 Display 2)
- X-Rite ColorMunki Smile (since ArgyllCMS 1.5.0)
- X-Rite DTP92
- X-Rite DTP94
- X-Rite/GretagMacbeth/Pantone Huey
- X-Rite/GretagMacbeth i1 Display 1
- X-Rite/GretagMacbeth i1 Display 2/LT (the HP DreamColor/Advanced Profiling Solution versions of the instrument are also reported to work)
- X-Rite i1 Display Pro, ColorMunki Display (since ArgyllCMS 1.3.4. The HP DreamColor, NEC SpectraSensor Pro, and SpectraCal C6 versions of the instrument are also reported to work).

The first step is to configure any hardware-dependent settings. Select the option labeled *Colorimeter corrections from other display profiling software*, and specify the group of colorimeters to which the connected device belongs. The software usually selects the correct entry. A firmware installation step often follows.

Either use the software provided by the manufacturer, which is usually no problem even with outdated versions, or you can load the manufacturer's

software from the website and enter *Select file*.

The program reports an error on some Spyder colorimeter models, but the device appears to work in spite of the error. Once the firmware is loaded, the device will appear as activated in the *Instrument* area (Figure 2).

You can select the type of screen you want to calibrate using the *Mode* drop-down menu. With laptop displays, the results are usually significantly poorer than with separate LCD monitors. If there are multiple screens available for the system – for example, because you are running an external monitor on the HDMI port of a laptop – select the desired device in the first, unnamed field.

Before clicking on the *Calibrate & profile* button at the bottom of the screen, you should configure some additional settings. Click the *Calibration* button to view the *Calibration settings* (Figure 3).

Normally the values for the *Whitepoint* and the *White level* will be fine if you choose *As measured*. The *Tone curve* is usually set correctly at *Gamma 2.2*. If you only work with sRGB profiles, sRGB might be the better choice. Many modern colorimeters allow you to calibrate at high speed instead of using the preset speed – if so, you can also choose the high-speed option.

The *Profile quality* slider (Figure 4) in *Profiling Settings* has a special meaning. Setting this slider to *High* yields far better results. The other settings are best left at the default values.

The *Calibrate & profile* button takes you to a test field where you can place the colorimeter. Then press the *Start measurements* button to start the

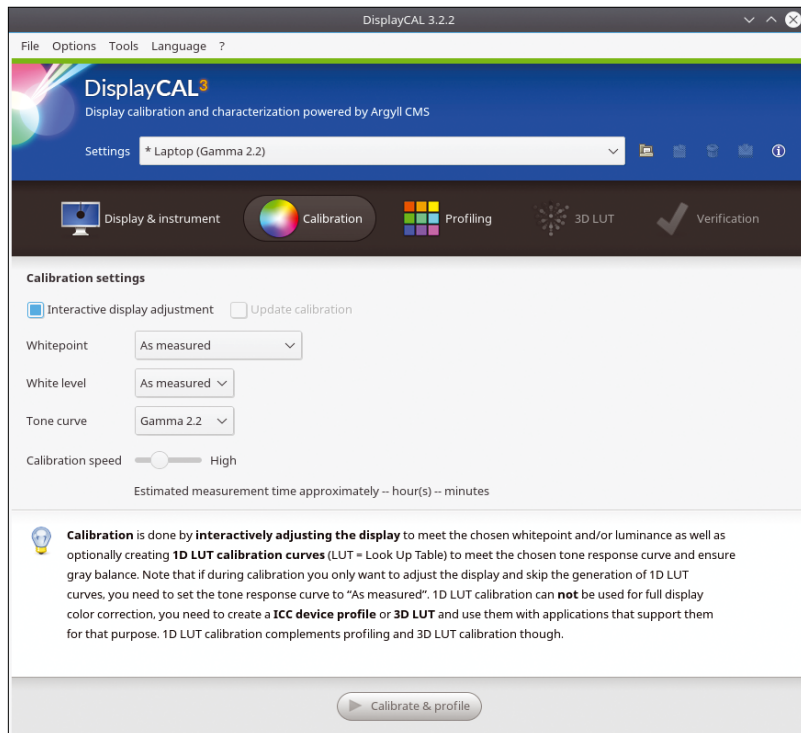


Figure 3: You can create defaults for calibration.

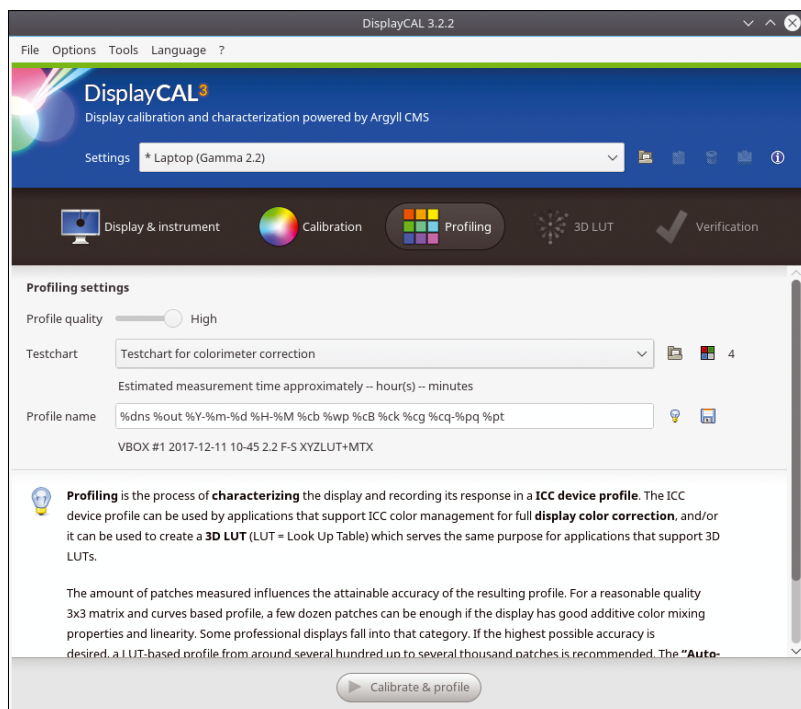


Figure 4: In *Profiling settings*, define basic settings for the color profile to be created.

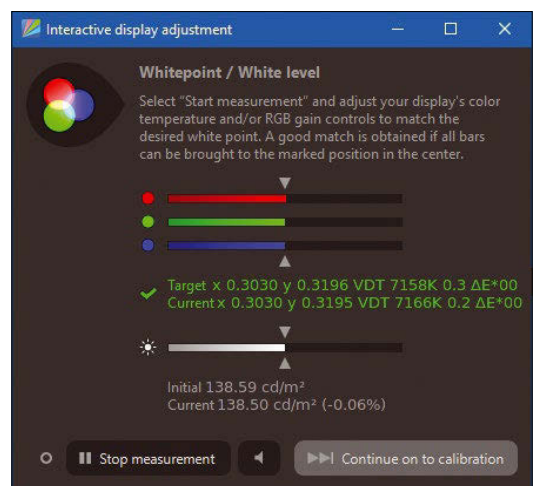


Figure 5: Calibration, including setting the basic monitor properties, is performed via a dialog.

TIP
Some manufacturers provide the driver data required from the firmware in the form of separate files with the *.ccmx* or *.ccss* file extension. Load these files manually via the *Correction* field (Figure 2).

Interactive Display Adjustment dialog (Figure 5). With laptop displays, you usually have no way to make adjustments. In the case of an external monitor, activate *neutral playback* in the Monitor menu and then start profiling at this point via *Start measurement*.

With a few short tests for the basic colors, the colorimeter determines the current settings. Now adjust the three RGB bars and the brightness on the monitor to equal (medium) values. DisplayCAL stays in this dialog until you end the measurement with *Stop measurement*.

Start the profiling, which you can perform using a long row of color fields, via the *Continue on to calibration* button. In some cases, you will see the differences with the naked eye; in others, you can hardly see the image change. Some dialogs provide information about the progress of the process (Figure 6). Depending on the colorimeter and the profile quality, the measurements will take some time.

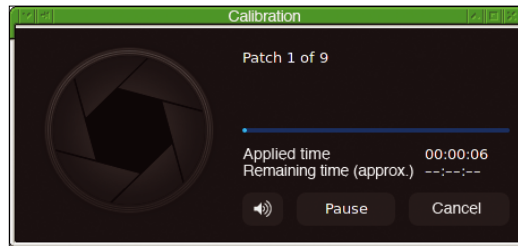


Figure 6: DisplayCAL comments on the actual measurements with progress dialogs.

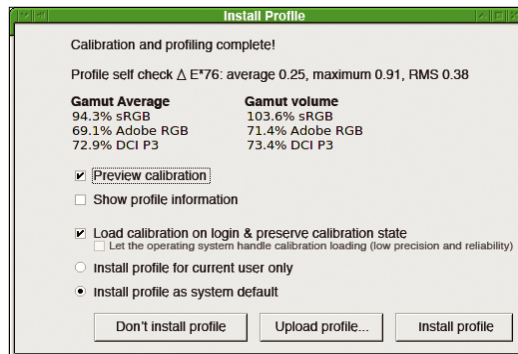


Figure 7: You can view the results of the measurement as graphical output via *Show profile information*.

Results

It is best to take the measurements in the evening in a dark room where no external light falls on the monitor. Disable the screensaver. The button with the loudspeaker symbol causes the software to give acoustic feedback as long as measurements are being taken.

On completion, the dialog shown in Figure 7 appears. The program shows which part of the standard RGB color space (sRGB) the monitor covers. *Gamut Volume* lets you evaluate the scope of the monitor's color space. The Gamut settings reveal that laptop screens often prove inferior to stand-

alone monitors, which restricts the use of a mobile computer for image processing.

The *Show profile information* option (Figure 7) offers an easier-to-interpret and more informative variant of the result representation (Figure 8). Of particular interest are the edge areas of the measured profiles, displayed in a standard color space. The edge areas appear in the graphic with a colored border. For comparison, see one of the predefined profiles, which then

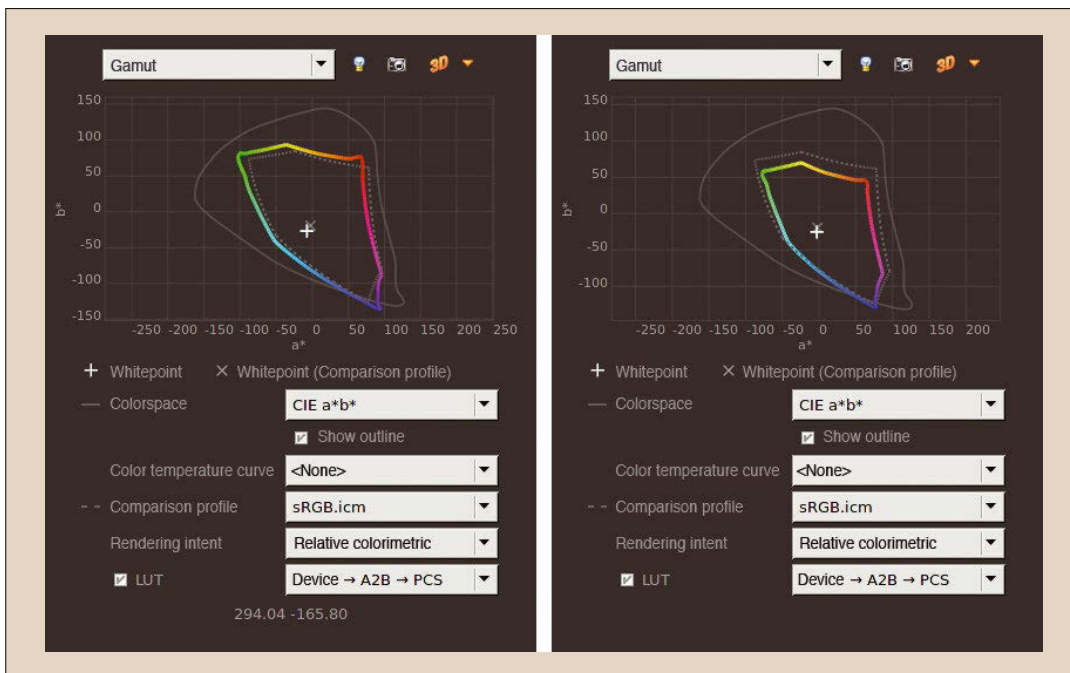


Figure 8: The diagram for the external monitor (left) shows that it largely covers sRGB and sometimes even shows more colors. On the other hand, the laptop screen (right) has coverage gaps on almost all edges of the color space.

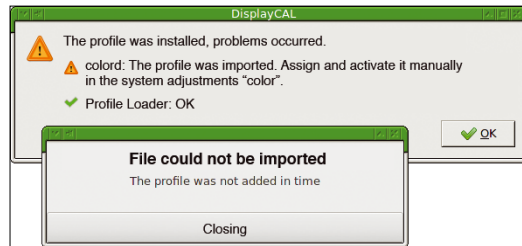


Figure 9: Under XFCE, the color profiles could not be installed automatically, but I was able to integrate them manually.

appears with a dashed gray line, under *Comparison profile*.

The relevant profile, in this case, is sRGB, which all screens should completely support. The profile initially uses *Absolute Colorimetric* as *Color Transfer* (Rendering Intent) – not a good choice. It is better to switch to *Relative colorimetric*.

To avoid losing track if you have several ICC profiles on the computer, the software offers the possibility of naming the profiles according to an adjustable scheme. The first component to appear is the (internal) display name, followed by the date of calibration. The same information, but with more detail, can be found in the profile information displayed, which is where you will find the manufacturer and sometimes the name of the model under *Device*.

The profiles generated by the application, as well as a number of additional files, are stored by the program under `.local/share/displayCAL/storage/` in your home directory. Not everything that is there is of direct benefit to you. If you have any problems, you should take a look at the logfile.

As a final step, the program now offers to load the profile you just created. Depending on the desktop environment, this step either works without any problems or not at all (Figure 9).

Manually Loading Profiles

Many Linux systems use the Color Daemon (`colord`) for managing color profiles. Some desktop environments offer GUI tools that act as an interface to `colord`, such as the Gnome Color Manager tool. The other option is to load the profile using a command-line tool such as `xcalib`:

```
xcalib profile_name.icc
```

`xcalib` is a “tiny monitor calibration tool” for X Window and MS Windows environments. The author describes the tool as “postcardware.” (It has a

Listing 1: Loading a Profile

```
$ xcalib -s 1 B156XW02\ 2013-09-12\ 2.2\ MQ-HQ\ 3xCurve+MTX.icc
$ xcalib -red 2 50 100 -a
```

GPL license, but the author requests you send him a postcard if you use it.)

See your local package manager for more on obtaining `xcalib`. With Arch Linux, the `xcalib` utility is contained in a separate package of the same name. You can also find the source code online [5]. The complete, not very meaningful documentation for the program is in a README file. `xcalib` requires profiles that contain the VCGT tag, which is always the case for the data generated with DisplayCAL.

If you have multiple screens, use the `-s` option to specify the number of the screen that will receive the profile. (Screen 0 is the default if you don't specify the `-s` option.)

The command from the first line of Listing 1 loads the specified profile for the display 1, an additionally connected monitor.

With `xcalib`, you can use the `-a` option to change loaded profiles during operation. The argument for one of the basic colors (`-red`, `-blue`, or `-green`) is three values: the gamma from 1 to 5, the brightness in percent, and the contrast in percent (Listing 1, second line). The adjustments remain active until they are switched off, unless you explicitly delete them using the `-c` option.

To automatically launch `xcalib`, use the bash configuration script `.bashrc` or call a script using the desktop's autostart mechanism.

If you have problems with `xcalib`, you can use the `xicclu` utility from the Argyll package [6]. `xicclu` is similar to `xcalib` but provides some different options. One disadvantage of `xcalib` is that it does not automatically detach itself from the terminal.

Conclusions

For users who are serious about editing images, it is definitely worth the time to create a profile of the monitor at least once. The result is always surprisingly positive, and the higher the monitor's quality, the more satisfying it is. Photos from the expensive SLR cameras finally look good, and shade and light have structure. Whether you want to go to the effort of checking the monitor regularly and possibly investing in a better colorimeter depends on your requirements. Normally, an annual recalibration run for a previously-profiled monitor is sufficient. ■■■

Info

- [1] Gamut: <https://en.wikipedia.org/wiki/Gamut>
- [2] Argyll: <http://argyllcms.com>
- [3] DisplayCAL: <https://displaycal.net>
- [4] ColorHug: <http://www.hughski.com>
- [5] xcalib: <http://xcalib.sourceforge.net>
- [6] xicclu: <http://www.argyllcms.com/doc/xicclu.html>

IT Highlights at a Glance



ADMIN HPC Up Close

- Spanning Free Protocol
- Watch Takes on Science
- HPC 2013 Call for Papers Hybrid Drive!

Further Reading

- IBM Develops Big Data Infrastructure for science
- Managing Linux Memory
- HPC Data Analytics
- 10 More Terabit Admin Tools
- Free webinar: CPU, I/O, and AIO Component Engineering for Cx & Ops

FREE FROM XP! Special Available!

ADMIN Update - Hottest Links

- Theft53
- New Report Explores the Prevalence of Loose Passwords
- 152K Day Lands on August 12
- How Reliable is a Wikipedia Citation?
- Leakage

Highlights

Further Reading

- IBM's TrustAdmin Dash
- High Availability without Penetration
- Secure Your VMs Virtual Machines
- OpenCL Workshop
- Checklists for Command-Line Tools

Maker Fairs

Block Level

Cloud Initiative

LINUX UPDATE EXPLORING THE WORLD OF LINUX

- Kyle Koenig's Powerful Graphics Editor Takes on Photoshop and GIMP
- Internet Giants Launch Collaboration to Improve Open Source
- 1000 group will focus on open source tools in large-scale environments (more)
- Mozilla Labs: Shuts Down
- Mozilla's product think tank sinks silently into history (more)
- I Still Haven't Again Talk About the Benefits of Free Software
- Register for OpenStack Summit

FEATURED ARTICLES

FURTHER READING

- 10 More Terabit Admin Tools
- Cloud Storage Security
- Patrols Anonymous' True Path: Manager App
- SQL Key Management Guidelines

Apps World

Easy Alternative for iPhone and iPad Users

MOST READ

ADMIN

Linux Pro Magazine

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

ADMIN HPC • ADMIN Update • Linux Update

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

Admin and HPC: www.admin-magazine.com/newsletter

Linux Update: www.linuxpromagazine.com/mc/subscribe

FOSSPicks

Sparkling gems and new releases from the world of Free and Open Source Software



Graham tears himself away from updating Arch Linux to search for the best new free software. **BY GRAHAM MORRISON**

Audio editor

Ocenaudio 3.3.6

The Audacity audio editor is excellent and powerful, but its user interface (UI) reflects the lack of any major updates for years. Like Gimp, there has been very little to compete with its capabilities, making Audacity the go-to application for any audio processing job, big or small. Ocenaudio, however, is a strong contender for Audacity's crown, with many small updates over the last 12 months and a modern Qt-based UI. It has had dark and light color schemes for a while, as well as high-DPI support from version 3.1.0. The UI scrolls seamlessly into and out

of a waveform and lets you scale the view from the small overview envelope so that it's easy to navigate through large files. Most importantly, Ocenaudio also works well with PulseAudio; Audacity still seems to struggle.

Monitoring is a real strength in both application, but even here Ocenaudio wins. It has a large, high-detail volume unit (VU) meter, the ability to see a spectrogram alongside the amplitude waveform, and a floating FFT analytics window for specific slices of audio time. It's a pity the internal effects don't run in real time in the background or that their

parameter windows don't offer more feedback (e.g., amplitude reduction or activity of the compressor and gate), but neither do Audacity's. What Ocenaudio does do, though, is allow you to see the audio waveform within the plugin window and use this to select a part to play back, during which you can change its parameters and apply the effect when you've dialed something suitable. This is a good half-way compromise between small preview clips and fully fledged real-time integration. The effects are good, too, especially the noise reduction. The Noise Reduction window lets you take a fingerprint profile of the background noise and process the audio in combination with a gate to remove the noise. It works brilliantly with the kind of noise you get from a low signal or cheap microphone.

Ocenaudio also offers a quick, powerful, and precision-based editor that can easily replace Audacity for most tasks. Place the cursor somewhere within the audio file and press record, for example, and Ocenaudio automatically records in-place. This is perfect for overdubs or replacing swear words in podcasts without going through the arduous process of recording something else separately and pasting it in. If you do need to work with several files at once, Ocenaudio has a handy side panel that lets you switch between them instantly or quickly paste to a new buffer that can easily be edited and pasted back into another of your recordings. Finally, to round off the feature set, its default file format is WAV, as used by almost every other audio application outside of Mac OS, except for Audacity, which requires you to export your audio and defaults when saving to a project file. The real winner in all of this is us, of course, as we now have a choice. Both projects are wonderful in their own way, and competition may just push the envelope with any new versions.

Project Website
<http://www.ocenaudio.com/>



1 File management: Open more than one file, tab between them, and paste to new ones. **2 Zoom and regions:** Quickly scale the view and create regions for looping and overdubs. **3 VU meters:** High-precision metering to ensure no clips or glitches. **4 Spectrogram:** Alongside the amplitude waveform, the spectrum shows the frequency range of audio over time. **5 FFT analysis:** A spectrum view of a single slice of time. **6 Effects:** Real-time previews from inside the plugin window and excellent noise reduction. **7 Routing:** Control all inputs and outputs to PulseAudio. **8 Audio properties:** Make notes, add artwork, and see region markers.

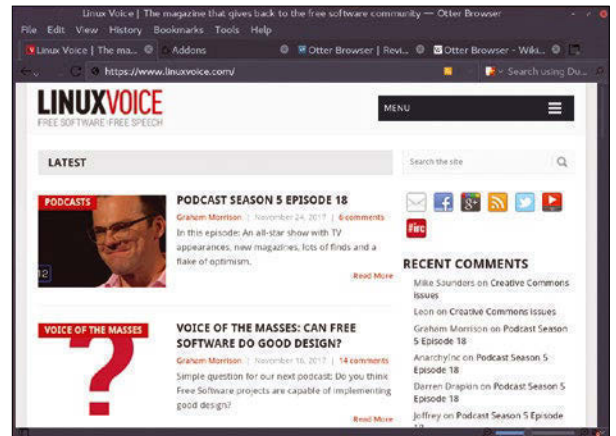
Web browser

Otter Browser

Many alternative web browsers were developed by simplifying or re-engineering the open source engines behind Chrome and Firefox. Firefox itself was born from a similar motivation to split the browser component from the Mozilla Application Suite, but there has been neither enough reason nor the mechanism to recreate the past glory of the Opera browser. However, that's exactly what Otter Browser does. If you don't use Opera, you won't be aware that a schism occurred in its user base between the release of 12.x and the following 15.x, when Opera switched from its Presto engine to WebKit and, subsequently, to the Blink engine used by nearly everything else. Many users felt some of Opera's best features were lost

in the transition, such as Opera's wonderful tab management and bookmark systems.

Otter is built using Qt and Qt's WebEngine for rendering web pages. Thanks to this, Otter fits perfectly into most Linux desktops, regardless of whether they're based on the Qt toolkit or not. Unlike most browsers, Otter will assimilate your font, color, and theme settings. The user interface is a little old school but useful; you'll also notice the downloading metrics in the status bar, such as the total size of a page, the speed, and the time it takes to load, as well as the original speed-dial shortcuts for quickly accessing your favorite sites. An integrated zoom slider is also a great way of making sites easier to read without need-



Session management is one of the best things about Otter Browser; sessions can be saved, managed, and loaded just like sets of open bookmarks.

ing to remember a keyboard shortcut, and everything is quick, stable, and system efficient; plus, of course, it has the features that made Opera 12 so popular. There has also been a constant stream of updates since its original release in 2013, making this an ideal browser for those of us still hankering for Opera's glory days.

Project Website

<https://otter-browser.org/>

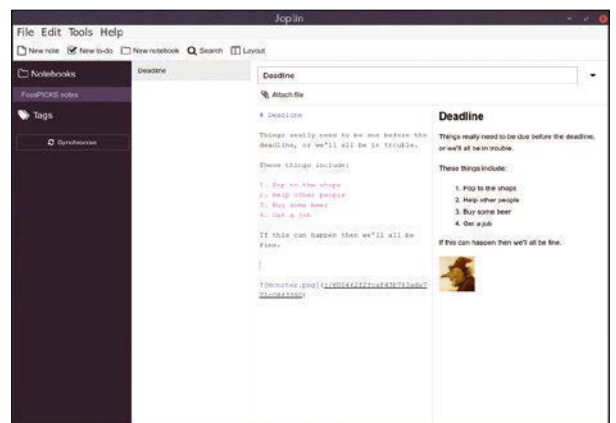
Note taking

Joplin

Note taking applications attempt to solve a difficult problem. They allow you to take notes and organize things that you don't think you'll remember and then hopefully allow you to retrieve those notes without knowing exactly what you made a note about. Among the many different solutions, most modern versions use the ubiquity of the services to solve the retrieval problem. Evernote and Google Keep, for example, store notes online and can use apps to set reminders that impose themselves into your day-to-day existence. Many of us are looking for self-hosted solutions that provide the same ubiquity; this is something Joplin gets very close to offering. For a start, it's available on almost every platform you can think of, including

iOS, Android, Linux, Mac OS, and Windows. There's even both a terminal version and a GUI version for the ultimate in accessibility. Hopefully, someone is working on the Amiga version.

You write notes using Markdown, which means if you already know its simple syntax, you can quickly create and format your notes, from headings and sub-headings to numbered lists, blocks of code, and direct quotes. It's the perfect format for this application, because it looks almost as good as raw text as it does rendered. Joplin will render the Markdown output in real time as you type and colorize your source material. You can also add attachments, and images are shown inline. Finally, you associate an alarm with a note, and `notify-osd` will be used on



For an application with such a simple premise, Joplin features excellent documentation that can really help you get the most out of your note taking.

Linux to inform you of when that deadline is hit. The equivalent notification system will be used on your phone or other operating system, and notes can be synchronized across different versions of the application using Microsoft's OneDrive. Hopefully, Nextcloud support is coming in the next update, creating a completely open source solution.

Project Website

<https://github.com/laurent22/joplin/>

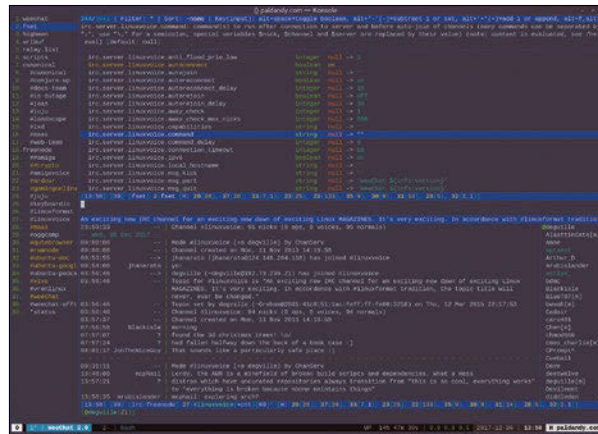
Terminal IRC

WeeChat 2.0

WeChat, the terminal-based IRC client, is one of the best reasons for using the command line. It's perhaps more powerful than almost any other client, and yet, the user interface is minimal, intuitive, and fast. It's harder to use than GUI-equivalent applications, but unlike the majority of those, WeeChat has never stopped innovating. Need Perl regular expressions to create an alert when something is mentioned? Want a push notification sent to your phone when your nicklist is used? Want to set up a WeeChat relay to use a single connection with multiple clients (e.g., your phone)? WeeChat does all of this and much more with an incredibly quick release cadence, as this 2.0 release proves.

The main new feature is the integration of `fset`, previously a plugin. This is another portal into the wonderful world of WeeChat settings configuration, alongside the interactive `iset` plugin that most of us use to change settings during a session. Without installing anything further, you can now issue the `/fset` command to open a new buffer from which you can browse and change anything from the complete set of WeeChat and running plugin parameters. It's quick and relatively easy to use and is a great step forward for WeeChat general configuration.

Scripts can now be written in PHP, alongside the many other languages supported by WeeChat. Many small additions help plugin makers and configuration tweekers, including recursive evaluation



Forget Slack, with clients like WeeChat, the supremely flexible and far more open IRC is the best way to communicate.

of variables within the buffer list and new signals. It's also possible to have two completely different bar items, each with its own configuration options, which is great for having different parts of your layout showing different servers or sets of channels.

Project Website
<https://weechat.org/>

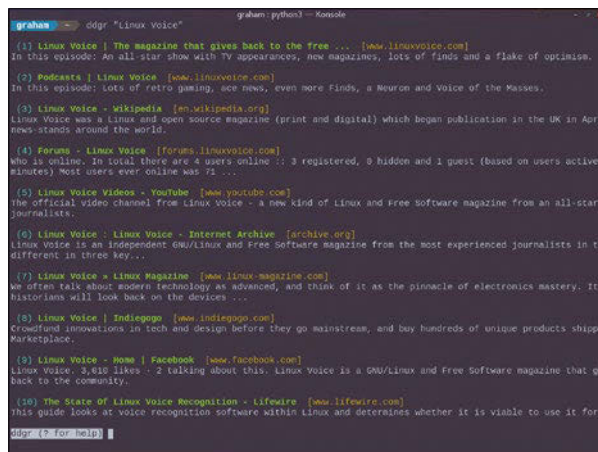
Terminal search

ddgr

This may seem rather esoteric, but because I've covered similar commands that deal with rival search engines, it seems more than fair to do the same for Duck Duck Go, a search engine that respects your privacy. It won't share your details, and it doesn't store anything in relation to your searches, unless you ask it to save your color and search preferences in a local cookie. If you've not used Duck Duck Go, or you haven't used it in a while, it's really improved over the last 12 months and now delivers results that are mostly comparable to Google. And when it's worse than Google, simply adding a `!g` to your search delivers results directly from Google through Duck Duck Go. This applies to other search en-

gines and web portals too, such as Wikipedia or even eBay, making the web interface to Duck Duck Go a powerful portal to the open web.

All of which leads to `ddgr`, a command-line interface to Duck Duck Go searches. After installation, you just type `ddgr` followed by your search phrase, and the results will be delivered as numbered output. Select a number and the result is opened in your default web browser. There's even `bang` support for external searches, such as `!g` for Google, and you can limit the number of



Search Duck Duck Go, and every other search engine supported by Duck Duck Go, directly from the command line.

results, fire off new searches from the prompt, and specify MIME file types for searches. If you need to search while in something like Vim, it's the perfect solution, and one that could easily be worked into your own scripts and sites, for instance, if you want to add dynamic linking or cache bookmarks.

Project Website
<https://github.com/jarun/ddgr>

... it's really improved over the last 12 months and now delivers results that are mostly comparable to Google.

Text editor

Suplemon

The Sublime Text editor has rightly become incredibly popular. Its minimal and beautiful user interface can adapt to any writing and programming style, and its exceptional plugin system can transform the editor into a fully fledged IDE, whether you're coding for Arduino or the cloud. One of its most innovative features is the ability to make multiple selections, rather than a single selection block used by nearly every other editor, which makes it easy to perform search and replace operations in a single pass or quickly change multiple instances of a variable name or the tags in a web document. This feature has been transposed into Suplemon, a console-based text editor that puts this

"multicursor" text editing at the top of its features list, alongside TextMate theme support, auto-complete, and a tabbed view with multiple files open at once.

To illustrate how these features work, load up a new file and create more than one cursor using the mouse. When you now start typing, the same text will appear at each cursor position. This is particularly useful when you're starting a new file because you can quickly populate it with similar arguments for methods, as well as the structure used by each of the methods you're going to write, before switching back to single-cursor mode. As with the Nano text editor, the command-line interface isn't intimidating because most of the important keyboard shortcuts

```

Suplemon Editor v0.1.63 - 0:19:06 38.12 [mplyosc.cpp]
#define MIDI_DEVICE "/dev/FLUID_Synth (/qsynth).128.0.0"

mPlayOSC: mPlayOSC()
{
}

osc_address = "127.0.0.1";
osc_list_port = 8000;
osc_host_port = 8000;

thread = 0;

// thread group to play note off durations
for (int i=0; i<NOTE_THREADS; i++){
  note_thread_grp.emplace_back(mPlayOSC: threadOSCnoteOff, i);
}

mPlayOSC: mPlayOSC(char *osc_address, int _list_port, int _hostport)
{
  osc_address = _address;
  osc_list_port = _list_port;
  osc_host_port = _hostport;

  thread = 0;

  // thread group to play note off durations
  for (int i=0; i<NOTE_THREADS; i++){
    note_thread_grp.emplace_back(mPlayOSC: threadOSCnoteOff, i);
  }
}

```

If you want to try some of Sublime Text's best features without paying for the editor, check out Suplemon.

are shown at the bottom of the text file, making it easy to get started, easy to know how to save, and easy to open up extra help. It's also a more powerful choice than Nano, which makes this an ideal editor upgrade if you find yourself spending more and more time in the terminal.

Project Website

<https://github.com/richrd/suplemon>

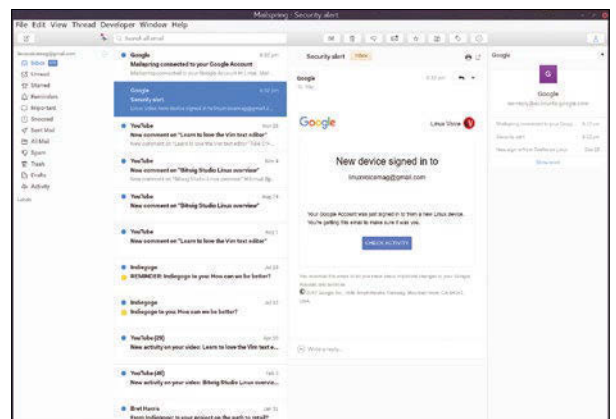
Desktop email

Mailspring

Nothing can beat a well-designed desktop email client. They are more efficient and faster than their web-based counterparts, they make it easier to host your own mail server or use a server other than Google's, and they integrate better with your desktop. Best of all, you don't get distracted by your email when you're being distracted by YouTube. But they've suffered a little at the hands of the web revolution; especially on Linux, their design and UIs haven't kept up with the times. Gears gets close with its conversational view, and KDE's Kube looks promising (if it can ever make it to release), but otherwise, we're usually left with Thunderbird or apps from five

years ago. Mailspring, however, is a breath of fresh air. It may be because it's part of a commercial product that seems to take a lot of its inspiration from Mac mail, but it's mostly because it's a fork of Nylas Mail, a beautiful client that attempted to bridge the cloud-desktop divide before closing in 2017.

First, there's a major negative with this application: You need to create a Mailspring account to be able to use it. While most of the application is open source, the company behind Mailspring wants to monetize its software by offering a Pro version. If this is acceptable, you'll be able to step through the login and add your own mail account details to the application. You can then enter your cloud provider details or



You need to register online to use Mailspring, but the app itself may be well worth the loss of complete control in certain circumstances.

provide the details for your own mail server. With your data sacrificed to the cloud, you get to experience the UI, and it's definitely one of the best on Linux. The main view to enable is three panels, with your folders on the left, your subject inbox, and the messages on the right. But whether it's worth the account creation is debatable.

Project Website

<https://getmailspring.com/>

Astronomy imaging

Siril 0.9.7

Astronomy and Linux often go hand in hand. It can be found controlling the largest and the smallest telescopes and is just as happy number crunching vague signals into planets orbiting a distant sun as it is guiding your personal telescope to the Horsehead Nebula in your backyard. There's plenty of open source astronomy software to choose from too, taking you from the photorealism of the Stellarium planetarium to the fully four-dimensional space exploration of Celestia. Although Linux has a constellation of brilliant applications for processing photos, there has never been a decent applica-

tion for processing a sequence of space photographs to create a single image ... drum roll ... until now. And Siril is awesome. Many problems befall amateur astronomy photography, but the two principle problems can't be solved: The night sky is dark, and the Earth is rotating on its axis. Because the night sky is dark,

photographs of space need to have long exposures. But because the Earth is rotating, unless you move the camera perfectly to compensate, the astronomical objects you're imaging will move to smear across your photograph. The only solution is to take shorter exposures, reducing the light hitting your lens, making images darker and the background noise easier to see. If you take several of these images, known as a stack, and then align and process them with some clever algorithms designed to tell the difference between the stardust and the CMOS pollution, you can create a single strong



Taken on a phone, this image is the result of Siril processing 10 three-second exposures of the sky (ISO 3200).

your phone sitting atop your car. See the above image for proof. Set a high ISO and RAW image output if you can and import these as a sequence into Siril. These are first converted to FITS or SER files, useful for high bit depths or importing videos, before being manipulated into a sequence by Siril. First, a little pre-processing can be used to darken and remove as much thermal noise as possible before Siril's incredibly smart alignment algorithm detects and removes the effects of the Earth's rotation from your sequence of images. Depending on the size and number of your images, as well as the algorithm you select, this can take either a few seconds or a full kettle's worth of minutes. Finally, the images are stacked and processed with a selection of noise-reducing and star or planetary detail-enhancing filters. These will further extract the background noise and attempt to detect color outside of local atmospheric conditions; the process is similar to creating an HDR image in something like Hugin. After tweaking, a final output can be generated with some patient parameter tweaking, some of which has real-time feedback, and many difficult control points and values to change; however, patience is always rewarded with stellar photography.

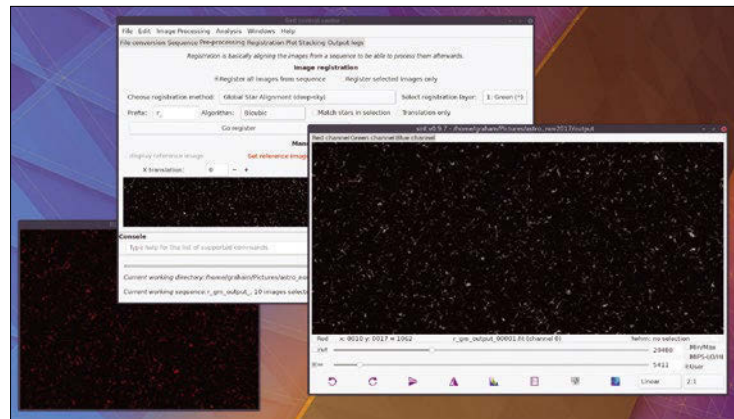
Project Website
<https://free-astro.org/index.php/Siril>

If you take several of these images ... you can create a single strong image that is greater than the sum of its parts.

tion for processing a sequence of space photographs to create a single image ... drum roll ... until now. And Siril is awesome.

image that is greater than the sum of its parts. This is what Siril does.

Siril imitates a left-right wizard stepping you through the entire process, although you can skip steps at any time. First, you need the images. These can be from a camera connected to a telescope, but they could equally be from



Siril is complex, but you can get by with adjusting values and checking the output. Fortunately, you also can access excellent online and PDF manuals.

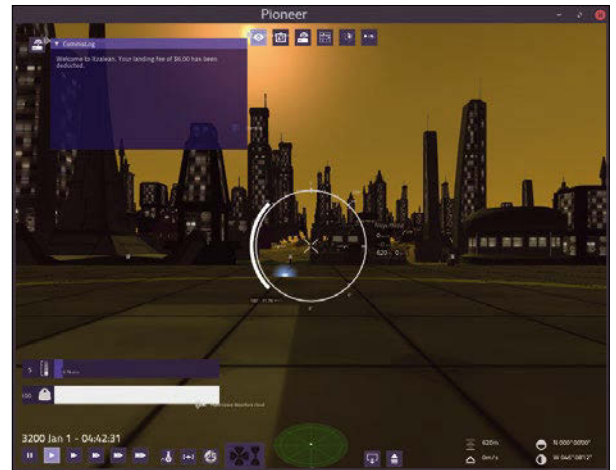
Space adventure

Pioneer

Like the Voyager 2 space probe, the game Pioneer has been around for a long time. It hasn't quite been around since 1977 when the probe was launched, but 2006 is an epoch in open source gaming. Also like Voyager, Pioneer is now interstellar and based on a technology from a previous generation, starting life as an attempt to bring something of the then-neglected Frontier – Elite 2, a 16-bit-era game – to Linux and open source. David Braben's Frontier was remarkable, taking the infinitely large open universe of Elite and populating it with curves, textures, missions, more craft, and a much richer universe. Braben's great genius has always been combining procedurally generated space with enough humanity to warm

up the place; he did it with the star and planet names in Elite, with the human contacts in the Frontier games, and from terrain generation to the AI interaction in his modern creation, Elite Dangerous. Pioneer does this too.

Pioneer captures the spirit of those original games alongside the cold hard vacuum of space. Classical music fills the communal space stations and the game sticks rigidly to the Newtonian flight model of Frontier 2. Although there's no multiplayer, the universe is a combination of procedurally generated systems and real systems. One of the start locations is Barnard's Star, for example. However, Pioneer isn't a copy. It's simply inspired by the same values as its models and offers a very similar game style in a completely new



Unlike Elite Dangerous, but just like Frontier II, you can land within cities on terraformed planets in Pioneer.

universe. Fly, taxi, trade, fight, upgrade, and explore the beautifully rendered 3D universe. If you were a fan of the originals, you'll feel right at home. And if you've never flown in 1980s space, climb aboard, Commander Jameson!

Project Website
<https://pioneerwiki.com>

Racing game

SuperTuxKart 0.9.3

SuperTuxKart has been on the nitro (i.e., for power) for some time now. Starting life as a quaint MarioKart homage, it's now the master of its own domain, and that domain just keeps getting richer, more playable, and more populated. It's now even available on Android, for free, and it's still open source! The basic play still features a story mode, where you drive around unlocking various challenges. This is much like Diddy Kong Racing, for example, where different lands reveal a set of different races against the AI-driven competition. There are time trial modes, grand prix cups, and multiplayer arenas, just as you'd expect, and like similar games, you'll be playing against cute avatars loosely based on those you see

around Free Software circles: Gnu, Wilber, Puffy, Beastie, Konqi, Kiki, Pidgin, and Amanda will all look familiar, and the story mode has you face the evil NoloK.

This new version has taken over a year to develop and features two new tracks. The first, Candela City, is a Paris-like city at night time. The second new level is Cornfield Crossing, and this looks like a classic Nintendo level set around a summer farm, complete with jumps through the barn, a roller coaster section into the desert, and banana skins. Keeping up with the times, the game also embeds a screen recorder for recording a video of your fastest laps. It must be working, because plenty of players have already shared theirs, adding a new di-




The awesome racing game SuperTuxKart is still open source and is now even available on Android!

mension to an excellent racer. As ever, if you donate more than \$3.99 to help fund development, you also get access to the gift package tracks. This is definitely worth it considering how many things this project gets right, and it's still at its heart an incredibly playable racer.

Project Website
<https://supertuxkart.net/>

SUBSCRIBE NOW



GET IT FAST
with a digital subscription!

The New IT

New tools, new threats, new technologies...Looking for a guide to the changing world of system administration?

shop.linuxnewmedia.com

AND SAVE 30%

6 issues per year!

Explore the new world of system administration

It isn't all Windows anymore - and it isn't all Linux. A router is more than a router.

A storage device is more than a disk. And the potential intruder who is looking for a way around your security system might have some tricks that even you don't know. Keep your network tuned and ready for the challenges with the one magazine that is all for admins.

Each issue delivers technical solutions to the real-world problems you face every day. Learn the latest techniques for better:

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

on Windows, Linux, and popular varieties of Unix.

**REAL-WORLD
PROBLEMS
SOLVED!**

ADMIN
Network & Security

shop.linuxnewmedia.com

Finders, Keepers

The Linux kernel has many interesting but unknown services. USB/IP, in particular, is one that you'll probably wonder why you have never encountered. USB/IP lets you use USB devices connected to other machines on your network as if they were plugged directly into your computer.

BY PAUL BROWN

This happened to me recently: I have an oldish Brother multifunction printer/scanner/fax machine. It is an okay thing to have and quite useful, but it is as dumb as brick and has no network capabilities at all. This means you must have a computer plugged into it directly if you need to use it. If you want to share it over a network, which of course you do, you have to make the computer a server and have it manage the network side of things.

Brother supplies some drivers for Linux, but they are closed source. This means that, if something doesn't work, you face a familiar conundrum: Firstly, Brother will not assign an engineer to sort out a problem for a system used by a minority of desktop users and a machine that is at least seven years old. Secondly, you can't solve things yourself because the drivers are closed source.

Unfortunately for me, something doesn't work.

Brother's scanner driver refuses to play ball with SANE's network printing protocol. That means, if I want to scan without lifting my laptop off the table, walking into the living room, unplugging it from the server, and attaching my computer directly to the lunk, I have to proceed as follows:

1. Open an SSH session from my laptop to the server.
2. Scan from the command line with `scanimage`.
3. Copy the file that `scanimage` creates over to my laptop using `scp`, and open it in an image visualizing program.
4. Make sure that the scan is the correct in size, color, contrast, and alignment.
5. If not, get up and adjust the piece of paper and the `scanimage` flags
6. Start all over again.

This is less than ideal, especially when you have graphical utilities, like, say, XSane [1], a comprehensive graphical application for all your scanning needs.

Is there a way to make the process less cumbersome? Thankfully, there is. It's called USB/IP (read: USB over IP); it is a set of modules and tools that comes with all modern kernels from 3.17 onwards. USB/IP allows you to set up a USB server on one machine (e.g., on your printer/scanner server) and share it with the rest of your network as if it were connected directly to the client machine.

USB/IP in Practice

This may sound like a godsend, but USB/IP, precisely because not very many people use it, is a bit finicky. All distributions provide the modules, as they are a standard part of modern kernels. The suitability of the tools you need for setting up the server and sharing to the client, however, is a different matter. The combination that worked for me was to use an up-to-date Debian as a server, and machines loaded with Ubuntu and KDE neon (which is currently based on Ubuntu 16.04) as clients. If you try anything else, the steps you follow should be very similar, but your mileage may vary.

Although the USB/IP modules (`usbip-core`, `usbip-host`, and `vhci-hcd`) are standard, you will have to install the tools that allow you to set up a server and start a client. In Debian, the package you need is called simply `usbip`, but Ubuntu's `usbip` package contains out-dated and flawed versions of the software. I have no idea why this exists – very confusing. If you have installed the `usbip` package before reading this, purge it now. What you are looking for is a package called `linux-tools-generic`. Install that.

Once you have everything installed on your server and clients, connect to the machine that will act as the USB/IP server and load two modules into the kernel:

```
su -c "modprobe usbip-core; modprobe usbip-host"
```

Then to start the server, still as root, do:

```
su -c "usbipd"
```

By starting the server like this, you will be able to see what it is doing. Your computer will display the machines that try to connect and show whether they are successful. This will help you troubleshoot during the first few runs. Once you are confident everything works as it should, you may want to run the server with the `-D` flag:

```
su -c "usbipd -D"
```

This will push `usbipd` to the background and have it run as a daemon.

You can now check what USB devices you can share. To do this, use the instruction shown in Listing 1, line 1. The `list` command is, as its name implies, for listing devices. The `-l` flag tells `usbip` to list the devices connected “locally.” As you will see later, you use a similar instruction to find out which devices are available on the server from the client.

As Listing 1 shows, there seems to be three devices you can share, but really there are only two: The last one (`busid 1-4`) is the machine’s internal USB hub doohickey (pretty sure that’s its technical name). The two devices you can share are `busid 1-1`, a Sanyo webcam, and `busid 1-2`, the multi-function printer/scanner/fax machine I mentioned earlier.

Notice that, in the case of the webcam, the server has no idea what the device is. But, then again, it doesn’t have to: The task of figuring out what a device is and how to use it falls to the client.

The next step is binding a device to the server, so it can be shared. To bind a device, you have to pass the device’s bus ID to the server:

```
su -c "usbip bind -b 1-1"
```

The `-b` stands for “bus ID,” and, as per Listing 1, the instruction above shares the Sanyo video camera.

To share the printer/scanner, you should do this:

```
su -c "usbip bind -b 1-2"
```

To stop sharing a device, you should use the `unbind` command. The command:

```
su -c "usbip unbind -b 1-2"
```

disconnects the printer/scanner from the server, for example.

Client Side

To get the Ubuntu/neon client set up (see the “Ubuntu Paths” box), load the modules you need into the kernel:

Ubuntu Paths

Ubuntu saves the USB/IP tools in a weird place that is not in your `$PATH`. This makes calling `usbip` cumbersome, since you have to prepend the whole path (`/usr/lib/linux-tools-<version number>/`) each time you want to use it.

This gets boring quickly, so you may want to make a soft link to the tool somewhere in your path. As it is a command you should issue only with root privileges, you can do:

```
sudo ln -s /usr/lib/flfl  
linux-tools-<version number>/usbip /usr/sbin/
```

For the rest of the article, I’ll assume you have done that.

```
sudo modprobe usbip-core  
sudo modprobe vhci-hcd
```

To query the server and see what devices it has on offer, use another variant of the `list` instruction:

```
sudo usbip list -r <server>
```

where `<server>` is the name or IP address of the USB/IP server. Listing 2 shows what querying the server looks like.

The listing gives you the two devices, the printer/scanner and the webcam, that the server is sharing.

Listing 1: Shareable Devices

```
$ su -c "usbip list -l"
- busid 1-1 (0474:025f)
  Sanyo Electric Co., Ltd : unknown product (0474:025f)

- busid 1-2 (04f9:01eb)
  Brother Industries, Ltd : MFC-7320 (04f9:01eb)

- busid 1-4 (8087:0a2a)
  Intel Corp. : unknown product (8087:0a2a)
```

Listing 2: Querying the Server

```
$ sudo usbip list -r 192.168.1.24
Exportable USB devices
=====
- 192.168.1.24
  1-2: Brother Industries, Ltd : MFC-7320 (04f9:01eb)
    : /sys/devices/pci0000:00/0000:00:14.0/usb1/1-2
    : (Defined at Interface level) (00/00/00)

  1-1: Sanyo Electric Co., Ltd : unknown product (0474:025f)
    : /sys/devices/pci0000:00/0000:00:14.0/usb1/1-1
    : Miscellaneous Device / ? / Interface Association (ef/02/01)
```

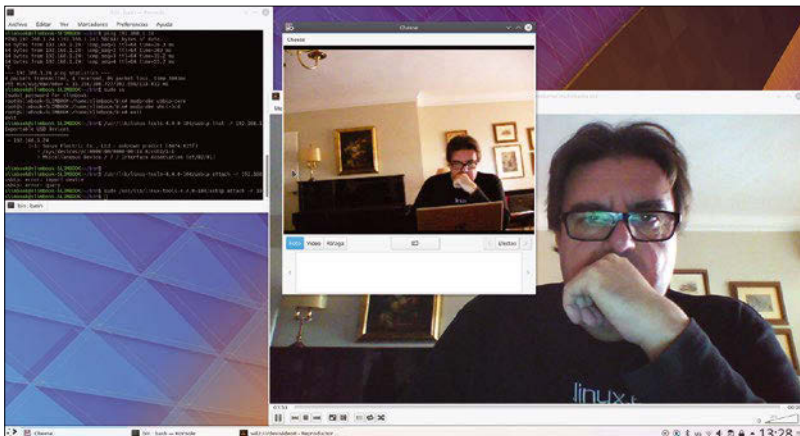


Figure 1: A question of perspective: video feed from the laptop camera (bottom right) and from a webcam connected to the USB/IP server several feet away (center).

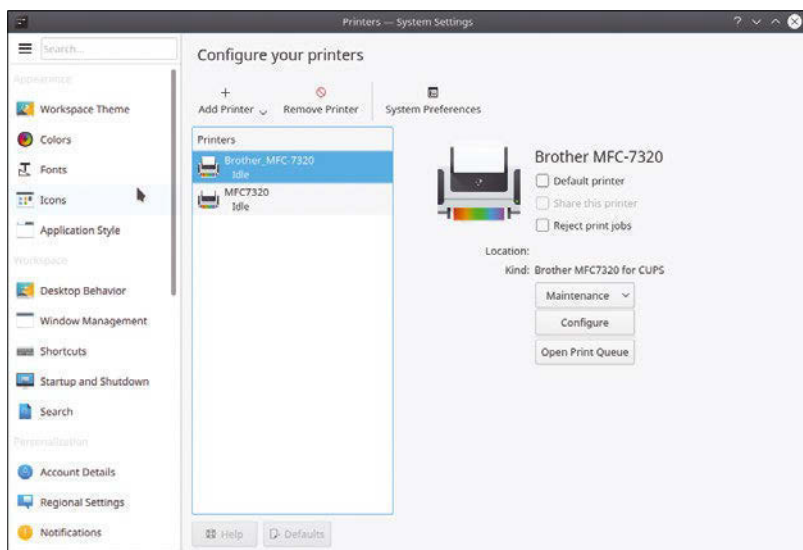


Figure 2: Two printers in one. The top printer is the virtual USB/IP printer, and the bottom one is detected via CUPS, although they are both the same physical machine.

To start using a device, you have to first attach it :

```
sudo usbip attach -r <server> -b <bus ID>
```

As before, <server> is your server machine’s name or IP (-r stands for “remote”) and <bus ID> is the bus ID of the device, as shown using the `lsusb` command in Listing 2.

To attach the Sanyo webcam shown in Listing 2, for example, you should do this:

```
sudo usbip attach -r 192.168.1.24 -b 1-1
```

Listing 3: Client USB Devices

```
$ lsusb
<Bus 003 Device 004: ID 0474:025f Sanyo Electric Co., Ltd>
Bus 003 Device 001: ID 1d6b:0002 Linux Foundation 2.0 root hub
Bus 002 Device 001: ID 1d6b:0003 Linux Foundation 3.0 root hub
Bus 001 Device 003: ID 058f:d102 Alcor Micro Corp.
Bus 001 Device 002: ID 8087:0a2a Intel Corp.
Bus 001 Device 001: ID 1d6b:0002 Linux Foundation 2.0 root hub
```

Note that some versions of `usbip` use `-h` (for “host”) instead of `-r`.

If everything goes well, the device should now show up on your system as if it were plugged directly into one of your USB ports. Listing 3 shows what the list of USB devices looks like on the client. Figure 1 shows what it looks like to see yourself from your laptop’s webcam and from a webcam connected to the USB server several feet away.

Attaching the printer/scanner had the interesting effect of providing the client with two printers: one network printer via CUPS and one virtual USB-connected printer (Figure 2), even though both are the same physical machine.

Just to show how the client sees the printers differently, take a look at Figure 3. On the left is what you get if you tell the CUPS network printer to print a test page – it is the server that sends the test page to the printer, hence the Debian/CUPS logos. By contrast, on the right is the test page from the virtual USB/IP printer. As the machine “thinks” the printer is connected locally to one of its USB ports, it sends its own, internal test page.

With a bit of fiddling (and by “a bit” I mean “hours of hair-wrenching frustration”) you can get even the most stubborn scanner to work as it should (Figure 4).

Detaching

Devices that are attached to a client via USB/IP are blocked for the rest of the network, so at some point, you will want to detach them and let others use them.

To detach a device, first you need to know to which virtual port each device is attached. Listing 4 shows an example from a client with the printer/scanner and webcam attached. If you are not sure which device is which, look at the vendor/product line and compare the numbers in brackets with what you get from `lsusb`.

Armed with the port number, you can detach a device with:

```
sudo usbip detach -p <port number>
```

So, to detach the printer/scanner (in Listing 4), do:

```
sudo usbip detach -p 01
```

Caveats

Although USB/IP is pretty awesome, there are some disadvantages you may want to consider before deploying it for everyday use. For starters, as it is underused, many distributions have scatty sup-

port for the USB/IP tools `usbip` and `usbipd`. Many distros include old versions, and different versions have different parameters. Some versions don't work at all.

Secondly, you have to attach and detach devices as root, which makes it inconvenient and a bit dangerous. Also, while a device is attached, it is blocked for other clients, and you have to remember to detach it when you have finished using it.

If you have a large number of client machines, by using USB/IP, you will have to install drivers onto each and every one of them. If you want to share a printer, CUPS may be the way to go, as you only have to install and maintain drivers on the server. If your scanner plays well with SANE's network system, this would also be the best way of cutting back on installing, configuring, and maintaining drivers.

Make It Better

That said, you can make USB/IP much friendlier if, for example, you load modules at boot time both on the server and client and turn `usbip` into a service on the server.

USB/IP also allows you to use a Raspberry Pi as a server for devices – even if the Pi doesn't understand them. You see, many vendors with proprietary drivers only provide drivers for Intel-based machines and do not supply any for the Pi, since it is an ARM-based computer. This is the case of the printer/scanner that kicked all of this off. With USB/IP, however, you can still use the Pi as a printer/scanner server because it doesn't actually have to load any drivers: The Pi just has to forward the whole USB connection to the client machine, where the driver is installed.

Is that cool or what? ■■■

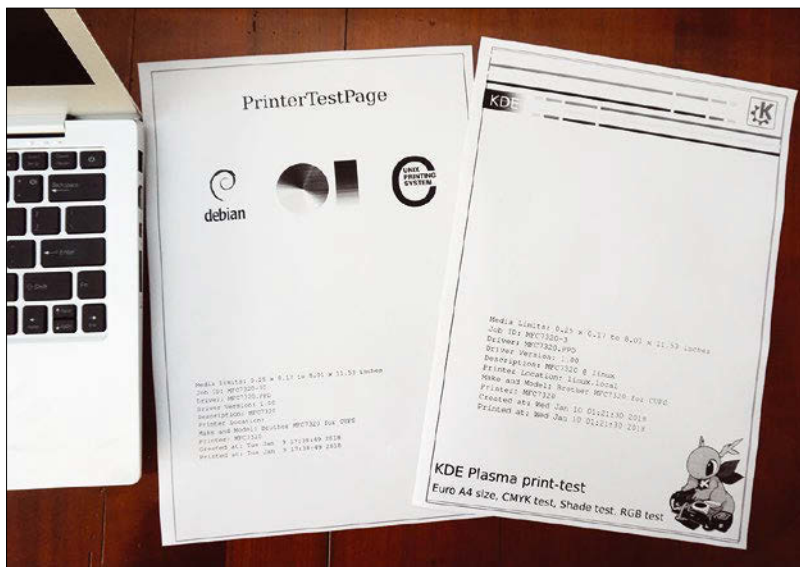


Figure 3: Same printer, but different test pages: on the left, the test page from the networked CUPS printer; on the right, the test page from the USB/IP printer.

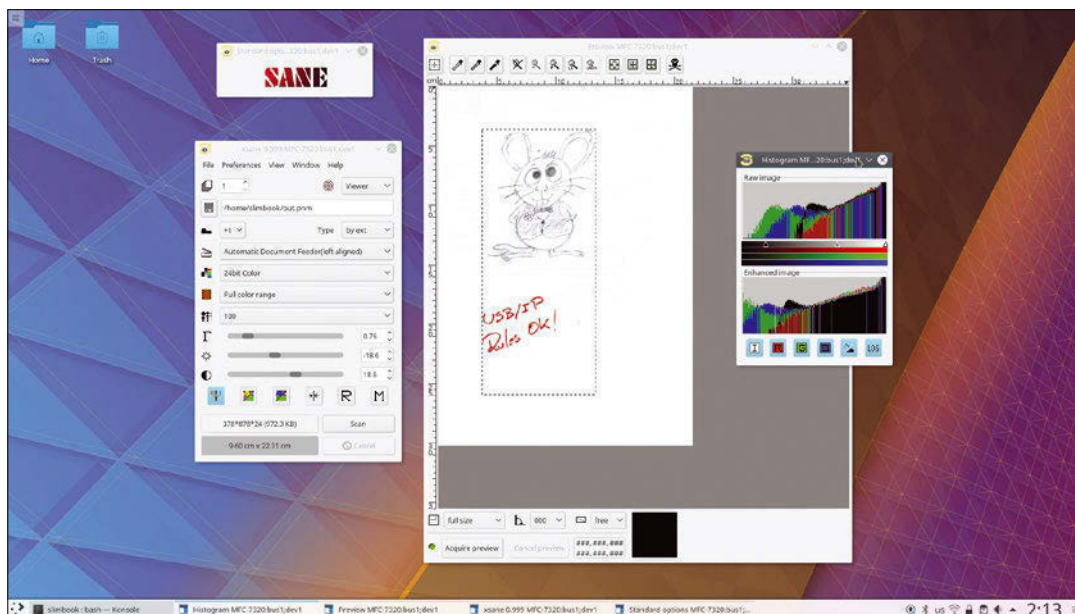
Listing 4: Connected Ports

```
$ sudo usbip port
Imported USB devices
=====
Port 00: <Port in Use> at High Speed(480Mbps)
  unknown vendor : unknown product (0474:025f)
  8-1 -> usbip://192.168.1.24:3240/1-1
        -> remote bus/dev 001/002
Port 01: <Port in Use> at Full Speed(12Mbps)
  unknown vendor : unknown product (04f9:01eb)
  8-2 -> usbip://192.168.1.24:3240/1-2
        -> remote bus/dev 001/004
```

Info

[1] XSane: <http://xsane.org/>

Figure 4: Beat your scanner into submission and make it allow scanning from the network thanks to USB/IP.



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.



Embedded Linux Conference

Date: March 12–14, 2018

Location: Portland, Oregon

Website: <http://events.linuxfoundation.org/events/embedded-linux-conference>

The Embedded Linux Conference is a vendor-neutral technical conference for companies and developers using Linux in embedded products. ELC has a large collection of sessions dedicated exclusively to embedded Linux and embedded Linux developers.

Open Networking Summit

Date: March 26–29, 2018

Location: Los Angeles, California

Website: <http://events.linuxfoundation.org/events/open-networking-summit-north-america>

Join business and technical leaders across enterprise, cloud, and service providers to share information, highlight innovation, and discuss the future of open networking and orchestration.

Linux Storage Filesystem and MM Summit

Date: April 23–25, 2018

Location: Park City, Utah

Website: <https://events.linuxfoundation.org/events/lfsfmm-2018/>

Experts map out and implement improvements to the Linux filesystem, storage, and memory management subsystems that will find their way into the mainline kernel and Linux distributions in the next 24 to 48 months.

Events

Open Source Leadership Summit	March 6–8, 2018	Sonoma Valley, California	http://events.linuxfoundation.org/events/open-source-leadership-summit
SCALE 16x	March 8–11, 2018	Pasadena, California	http://www.socallinuxexpo.org/scale/16x
Chemnitz Linux Days 2018	March 10–11, 2018	Chemnitz, Germany	https://chemnitzer.linux-tage.de/2018/en/
Maker Faire Ruhr	March 10–11, 2018	Ruhr, Germany	https://www.makerfaire-ruhr.com/
CloudFest 2018	March 10–16, 2018	Rust, Germany	https://www.cloudfest.com/
Embedded Linux Conference	March 12–14, 2018	Portland, Oregon	http://events.linuxfoundation.org/events/embedded-linux-conference
OpenIoT Summit	March 12–14, 2018	Portland, Oregon	http://events.linuxfoundation.org/events/openiot-summit
Open Networking Summit	March 26–29, 2018	Los Angeles, California	https://events.linuxfoundation.org/events/open-networking-summit-europe-2018/
NSDI '18	April 9–11, 2018	Renton, Washington	https://www.usenix.org/conference/nsdi18
heise Security Tour 2018	April 10, 12, 18, 24, and 26, 2018	Numerous European cities	https://www.heise-events.de/securitytour
2018 HPC for Wall Street	April 16, 2018	New York, New York	http://www.flagmgmt.com/linux/
Cloud Foundry Summit North America	April 18–20, 2018	Boston, Massachusetts	https://www.cloudfoundry.org/event/nasummit2018/
Linux Presentation Day 2018.1	April 21, 2018	Europe-wide in many cities	http://www.linux-presentation-day.org/
Linux Storage Filesystem and Memory Management Summit	April 23–25, 2018	Park City, Utah	http://events.linuxfoundation.org/events/lfsfmm-2018/
KubeCon + CloudNativeCon	May 2–4, 2018	Copenhagen, Denmark	https://events.linuxfoundation.org/events/kubecon-cloudnativecon-europe-2018/
Texas Linux Fest	June 8–9, 2018	Austin, Texas	https://2018.texaslinuxfest.org/

CALL FOR PAPERS

We are always looking for good articles on Linux and the tools of the Linux environment. Although we will consider any topic, the following themes are of special interest:

- System administration
- Useful tips and tools
- Security, both news and techniques
- Product reviews, especially from real-world experience
- Community news and projects

If you have an idea, send a proposal with an outline, an estimate of the length, a description of your background, and contact information to edit@linux-magazine.com.



The technical level of the article should be consistent with what you normally read in *Linux Magazine*. Remember that *Linux Magazine* is read in many countries, and your article may be translated into one of our sister publications. Therefore, it is best to avoid using slang and idioms that might not be understood by all readers.

Be careful when referring to dates or events in the future. Many weeks could pass between your manuscript submission and the final copy reaching the reader's hands. When submitting proposals or manuscripts, please use a subject line in your email message that helps us identify your message as an article proposal. Screenshots and other supporting materials are always welcome.

Additional information is available at:

http://www.linux-magazine.com/contact/write_for_us.

NOW PRINTED ON recycled paper from 100% post-consumer waste; no chlorine bleach is used in the production process.

Authors

Erik Bärwaldt	32, 50
Swapnil Bhartiya	8, 28
Paul Brown	92
Zack Brown	12
Bruce Byfield	42, 66
Joe Casad	3
Mark Crutch	71
Christoph Dyllick-Brenzinger	56
Karsten Günther	78
Jon "maddog" Hall	72
Charly Kühnast	54
Christoph Langner	16
Vincent Mealing	71
Graham Morrison	84
Mike Schilli	46
Ferdinand Thommes	22, 74
Harald Zisler	60

Contact Info

Editor in Chief

Joe Casad, jcasad@linux-magazine.com

Managing Editor

Rita L Sooby, rsooby@linux-magazine.com

Localization & Translation

Ian Travis

News Editor

Swapnil Bhartiya

Copy Editor

Amy Pettle

Layout

Dena Friesen, Lori White

Cover Design

Lori White

Cover Image

© ndul, 123RF.com

Advertising

Brian Osborn, bosborn@linuxnewmedia.com
phone +49 89 99 34 11 48

Publisher

Brian Osborn, bosborn@linuxnewmedia.com

Marketing Communications

Gwen Clark, gclark@linuxnewmedia.com
Linux New Media USA, LLC
616 Kentucky St.
Lawrence, KS 66044 USA

Customer Service / Subscription

For USA and Canada:
Email: cs@linuxpromagazine.com
Phone: 1-866-247-2802
(Toll Free from the US and Canada)
Fax: 1-785-856-3084

For all other countries:

Email: subs@linux-magazine.com
Phone: +49 89 99 34 11 67
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 © 2018 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, 616 Kentucky St., Lawrence, KS, 66044, USA. Periodicals Postage paid at Lawrence, KS and additional mailing offices. Ride-Along Enclosed. POSTMASTER: Please send address changes to Linux Pro Magazine, 616 Kentucky St., Lawrence, KS 66044, USA.

Published monthly in Europe as Linux Magazine (ISSN 1471-5678) by: Sparkhaus Media GmbH, Ziebländstr. 1, 80799 Munich, Germany.

Issue 209 / April 2018

Browser Shootout

Approximate
UK / Europe Mar 05
USA / Canada Mar 30
Australia Apr 30
On Sale Date

After all these years and all this new technology, the web browser is still the most popular application for many desktop users. Which browser is best? It all depends on what you do and what you need. Next month, we compare some Linux browser alternatives, including popular applications like Firefox and Chrome, as well as some command-line browsers and other lesser-known options.



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

Lead Image © Dejan Bozic, 123RF.com

Happy 25th Anniversary to Linux!

Help us celebrate
25 years of Linux
with the All-Time
Archive DVD of
Linux Magazine!

Order today and
get 191 issues of
Linux Magazine on
one handy DVD!




You get 17,000 pages of practical know-how on one searchable disc - that's 191 issues including 40+ issues that were not included in the previous release.

Order Now! Shop.linuxnewmedia.com

HETZNER
ONLINE

**THE FUTURE
STARTS NOW!**

**WITH THE NEW
HETZNER CLOUD**



cloud.hetzner.com