

FREE DVD



**BROWSER SHOOTOUT**  
We compare Firefox, Chrome, Opera, and Vivaldi

**IS FIREFOX FINALLY BACK?**

**LINUX**  
MAGAZINE



Double-Sided DVD  
**INSIDE!**

# LINUX **PRO** MAGAZINE

APRIL 2018


## IS FIREFOX BACK?

We compare Mozilla's reborn browser with recent versions of Chrome

### Font Tricks

From the comfort of your office suite

Write a script that changes your YouTube metadata



### Kernel Self-Protection

Tune in to this practical project for tips on safer coding

### Lakka

Turn your Rasp Pi into a gaming console



### NoMachine

Remote desktop solution

### Tkinter Library

Control your maker creations from a smartphone

## LINUXVOICE

### FOSS Picks

- KIT Scenarist
- Oil Bash Alternative
- Etcher ISO to USP

- maddog: When chips go bad
- LFTP: Better, faster, more versatile FTP
- Bringing order to ideas with Compendium



### Tutorial

Over-the-Top Video Transitions

Issue 209  
Apr 2018  
US\$ 15.99  
CAN\$ 17.99



**HETZNER**  
ONLINE

**LOTS OF CLOUD**

**FOR LITTLE MONEY!**

**\$3.10**



## Cloud Server CX11

Intel® Xeon® Skylake  
1 vCPU  
2 GB RAM  
20 GB NVMe SSD  
20 TB traffic\*  
Intuitive Cloud Console  
Free DDoS protection  
Located in Germany



**\$3.10** per month

## The time has come for truly thrifty cloud hosting

Our clever Cloud Console interface makes it easy to use our Hetzner Cloud servers for all your IT needs. With it, you can create server instances almost instantly, usually in under 10 seconds. Hetzner Cloud servers naturally have flexible billing. Each server has a monthly price. But if you choose to use the server for less than a month, you will be billed at the hourly rate.

**cloud.hetzner.com**

\* With 20 TB of included traffic, you'll have lots of bandwidth for your projects, regardless of which Hetzner Cloud package you choose. But if you need, you may add more for an extra \$ 1.25 a month per TB.

All prices exclude VAT and are subject to the terms and conditions of Hetzner Online GmbH. Prices are subject to change. All rights reserved by the respective manufacturers. Intel, Intel Logo, Xeon and Xeon Inside are brands of the Intel Corporation in the USA or other countries.

# BACK TO THE FUTURE

Dear Reader,

Facebook has become a huge colossus in the advertising industry, sucking the air out of the ad market and seizing a gigantic market share with their low prices and fine-tuned demographic targeting. This tremendous surge has led to large profits and millions of paying ad customers. One customer that might not be paying for long is Unilever, a vast consumer conglomerate with properties like Lipton and Ben & Jerry's. At the Interactive Advertising Bureau's annual leadership meeting in Palm Desert, California, Unilever spokesman Keith Weed denounced social media companies like Facebook for what he called an "erosion of trust" [1]. Although he had several comments about transparency and truthfulness, some of his strongest words were about what he called, "...an environment where our consumers can't trust what they see online." In what sounds like an ultimatum, he added, "Unilever will not invest in platforms that do not protect our children or which create division in our society."

The problem, of course, is the role Facebook and other big social media companies play in spreading false information. Many experts believe the fake news epidemic may have influenced the 2016 US election. Where did all this bad information come from? From shadowy, fake accounts, and it took months after the election to even sort out what the ads were and where they started.

For a little context on why this new phenomenon is so weird, it might be best to start with how ad buying used to work (and still does work) at conventional media companies. In the old model, ad sales occurred within an active relationship between an ad buyer rep, working for the company that is advertising, and an ad seller rep, working for the media company. After the deal is made, a creative team working for the ad buyer creates an ad and sends it to the media company. Someone within the production team of the media company receives the ad image and logs it, taking whatever steps are necessary to schedule it, or, in the case of a print publication, place it in an issue.

Along the way, the ad crosses the hands and eyes of several humans, and at any point, someone can raise a question if the ad is offensive or misleading. The media company doesn't want an offensive ad on their website or in their magazine because they have a reputation that is associated with the contents of the publication. The company buying the ad also has

a reputation to protect. The ad buyer pays attention to where the ad goes and makes sure it appears in a place where it won't be associated with nefarious and creepy content.

In the new social media advertising model, however, ads are like water. You turn on the tap and they flow, more or less automatically, from the point they are envisioned, seeping out into the Internet according to the nuances of a predetermined demographic profile.

The popular conception is that Facebook and other social media companies have brought innovation to ad targeting through their command of personal profiling data, which is certainly true. But the other part of the story that receives far less information is that they have put themselves in black ink by automating the ad purchase and upload process, which allows them to operate at a massive scale without the need for human intervention. In other words, part of their "innovation" has nothing to do with technical brilliance and is simply that they operate without the traditional protections built into the system – like the manager of a swimming pool who "innovates" by laying off all the lifeguards.

In November 2017, Mark Zuckerberg told shareholders the company is serious about combating the perception that Facebook is a platform for false information [2]. He announced a new initiative that he warned could "...significantly impact our profitability going forward." The initiative will increase the total expenses for the company between 45% and 60% over 2017 levels, and much of the money will reportedly go to hiring thousands of new employees who will take some personal ownership of ads going on the site.

If Facebook follows through on their promise to wrangle some human judgment back into the ad process, they should be commended for taking the problem seriously. If they succeed, they will see their profits decline, and they will gradually start to look more like an everyday online company and less like a supernatural disrupter. They might even have to raise their ad prices, which might bring some competition back and make the other online advertising venues seem not-so-obsolete after all. But hey, that's progress.

Joe

Joe Casad,  
Editor in Chief



## Info

- [1] Facebook Gets Weed-Whacked: Unilever Exec May Axe Ads over Social Network's Toxic Posts: [https://www.theregister.co.uk/2018/02/12/unilever\\_facebook\\_toxic\\_content/](https://www.theregister.co.uk/2018/02/12/unilever_facebook_toxic_content/)
- [2] Facebook: Foiling Rogue Posts and Ads Will Significantly Hit Profits: <https://www.fastcompany.com/40490540/facebook-defeating-fake-news-and-rogue-ads-will-significantly-hit-profits>

# LINUX MAGAZINE

## WHAT'S INSIDE

The release of Firefox 57 last November offered new hope that Mozilla's free browser was back at last and ready to compete with Chrome. This month we take up the challenge with a head-on performance comparison of Firefox 57 with Chrome, Vivaldi, and Opera. We also offer some tips on how to clean up the user data stored with your browser and how to configure custom privacy settings.

Elsewhere inside:

- **Kernel Protection** – the Kernel Self-Protection initiative promotes safer coding practices for the Linux kernel community.
- **YouTube Metadata** – write a script that changes your YouTube metadata automatically, allowing you to roll back to previous profile settings.

Check out MakerSpace for a look at how to control a maker project remotely with the Tkinter Python library, and join us at LinuxVoice for LFTP, the Compendium mind mapper, and more great FOSS Picks from Graham.

## SERVICE

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

## NEWS

### 08 News

- LibreOffice 6.0 Released
- Red Hat Acquires CoreOS
- Red Hat Enterprise Linux 7.5 Beta Out
- Torvalds Is Not Happy with Intel's Patch, Calls It Garbage
- More Than 2,000 WordPress Sites Infected by Malware

### 12 Kernel News

- Preventing the Kernel from Tainting
- Encrypting printk() Output
- New Kernel Bug Reporting Bot

### 16 Interview

Dell Technologies VP Josh Bernstein talks about {code} and the value of open source.

## REVIEWS

### 36 Font Management

Users of contemporary office applications often want to do more than just compose letters or design simple flyers. But the well-known office suites on Linux sometimes make font management a test of patience, with the errors literally hidden in the details.



## COVER STORIES

### 18 Web Performance

Will the Firefox 57 Quantum release help Mozilla regain its former glory? We compare Firefox 57 with the Chrome, Opera, and Vivaldi browsers.

### 24 Browser Clean-Up

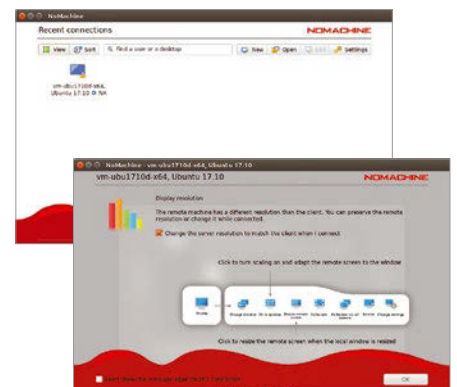
Web browsers collect a large amount of data about the user's browsing habits. If you care about privacy, you might want to clean up your browser and configure some custom privacy settings.

### 30 Firefox in Transition

What's new in Firefox 57 Quantum – and why does it matter?

### 42 NoMachine

NoMachine offers remote access to a PC and much more.



IN-DEPTH

46 Kernel Protection

The kernel hacker initiative, Kernel Self-Protection, promotes safe programming techniques to keep attackers off the network, and, if they do slip through the net, mitigate the consequences.



50 Charly's Column: Sysdig

Charly invites Sysdig, the jack-of-all-trades among system diagnostic tools, into his office for a quick checkup.

52 Serverless

Serverless computing takes cloud abstraction to a whole new level. Find out why more developers are turning to this exciting new programming paradigm.

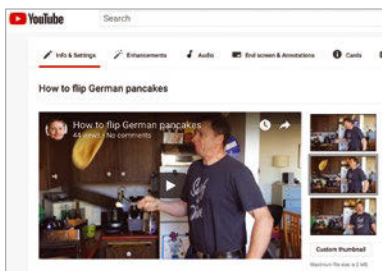


54 Command Line – Tilix

A terminal emulator with innovative features makes working from the command line easier and more efficient.

58 Programming Snapshot – YouTube Metadata

Instead of manually editing the metadata of YouTube movies, video craftsman Mike Schilli dips into YouTube's API spell book and lets a script do the work.



MAKERSPACE

62 Open Hardware – Ebook Publishing


LibreOffice, Calibre, and Sigil help would-be authors with do-it-yourself ebook publishing.

66 Lakka

With the right software, you can turn the Raspberry Pi into a versatile console for retro games.

70 Tkinter

Use Tkinter to control your Rasp Pi projects from a smartphone or tablet.



LINUXVOICE

75 Welcome

This month in Linux Voice

77 Doghouse – Chip Replacement

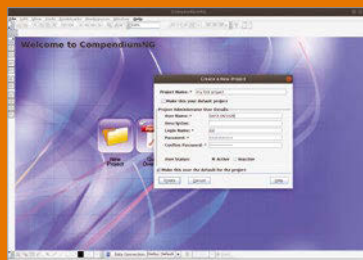
Chip replacement isn't always the answer for hardware problems.

78 LFTP

LFTP is an alternative to the FTP command set that supports many protocols and offers countless parameters.

82 Compendium

Compendium helps bring order to ideas and thoughts and link them to form a complex map from which you can derive arguments and decisions.

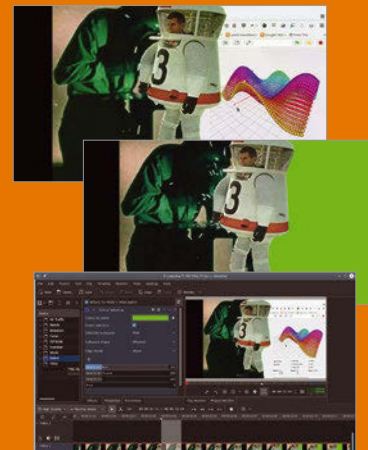


86 FOSSPicks

Graham Morrison looks at KIT Scenarist, Inboxer 1.0, Tungsten, rtl\_433, Oil 0.3.0, Etcher 1.3.0, and more!

92 Tutorial – Video Editing

Most video editors supply you with a generic catalog of transitions, usually in the shape of tired wipes and fades. But what if you wanted something a little more special?



# On the DVD



ubuntu Budgie<sup>®</sup>  
17.10 (64-bit)

## Ubuntu 17.10 "Budgie"

This up-and-coming Ubuntu spin features the Budgie desktop environment. The simple yet practical Budgie is designed to provide "...a familiar, modern, and functional experience whilst getting out of the user's way." Because this release is all Ubuntu under the hood, you get the bug fixes, security updates, and enhancements that come with Ubuntu 17.10 "Artful Aardvark," including improved swap performance and driverless printing support.



TrueOS<sup>®</sup> 17.12

## TrueOS Desktop 17.12


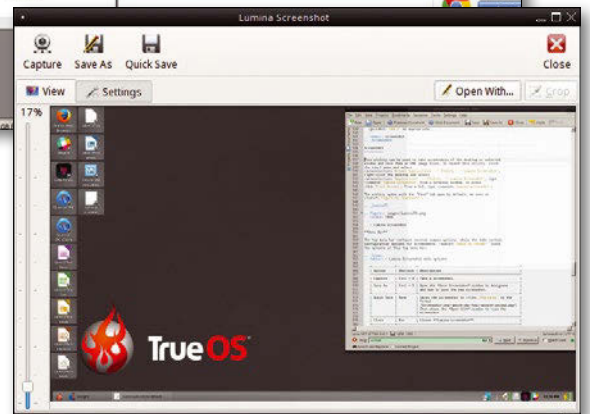
TrueOS is a desktop OS based on FreeBSD. The FreeBSD project is a Unix-like cousin of Linux known for stability, security, and a tight-knit development community. The goal of the TrueOS project is to deliver the benefits of BSD while lowering the learning curve with an easier end-user experience. The latest version features graphical remote management and enhanced system update features.

ISSUE 209  
APR 2018  
LINUX  
MAGAZINE

DVD  
ROM

TWO TERRIFIC  
DISTROS

DOUBLE-SIDED  
DVD!




## Additional Resources

- [1] Ubuntu Budgie: <https://ubuntubudgie.org/>
- [2] Budgie Desktop Environment: <https://budgie-desktop.org>
- [3] TrueOS: <https://www.trueos.org/>
- [4] TrueOS Blog: <https://www.trueos.org/blog/>
- [5] TrueOS Handbook: <https://www.trueos.org/handbook/trueos.html>

Defective discs will be replaced. Please send an email to [subs@linux-magazine.com](mailto:subs@linux-magazine.com).

Prague, Czech Republic



# Conference

Open Source Software


- ◆ Embedded Systems
- ◆ Cloud ◆ Containers ◆ More

May 25 - 27, 2018

[events.opensuse.org](https://events.opensuse.org)

# NEWS

Updates on tech



## THIS MONTH'S NEWS

### 08 Open Source Office Suite

- LibreOffice 6.0 Released

### 09 Red Hat Grows Kubernetes Share

- Red Hat Acquires CoreOS
- Red Hat Enterprise Linux 7.5 Beta Out
- More Online

### 10 Faulty Meltdown/Spectre Patches

- Torvalds Calls Intel's Patch Garbage
- More than 2,000 WordPress Sites Infected by Malware

## LibreOffice 6.0 Released

The Document Foundation announced the release of LibreOffice 6.0, the biggest release of the fully open source office suite since LibreOffice 5.x was released in 2016.

One of the biggest improvements in LibreOffice 6.0 is the introduction of Notebook Bar, a ribbon interface that makes it easier to perform many tasks without having to dig through menus. You can enable the feature from the Advanced settings. The new release offers many different modes that users can choose for their own workflows, including a sidebar and a minimalistic single toolbar.

LibreOffice 6.0 claims to offer better file compatibility with Microsoft Office documents. The release also offers the ability to export documents as ePub, an ebook format. "OOXML interoperability has been improved in several areas: import of SmartArt and import/export of ActiveX controls, support of embedded text documents and spreadsheets, export of embedded videos to PPTX, export of cross-references to DOCX, export of MailMerge fields to DOCX, and improvements to the PPTX filter to prevent the creation of broken files," said Italo Vignoli, one of the cofounders of the Document Foundation.

There is still no release of LibreOffice for mobile platforms that competes with Microsoft Office and Google Docs. LibreOffice Viewer has largely been an app to view documents on mobile devices. But The Document Foundation promises that the upcoming release of LibreOffice Viewer for Android will be able to create new documents. "It will offer a tab-based toolbar with formatting options and will let users add pictures either from the camera or from a file stored locally or in the cloud," said Vignoli.

LibreOffice is available for the cloud through Collabora. LibreOffice 6.x will bring more capabilities to LibreOffice Online. "New features introduced with LibreOffice 6.0 aim to align the functionality of the desktop and cloud versions, especially in areas where users expect similar behavior," said Vignoli.

Collabora offers Collabora Online Developer Edition (CODE), a free of cost solution that's based on the latest version of LibreOffice. Users can run it on their own servers with other open source solutions, like Nextcloud for storage. Collabora also

offers a paid version of LibreOffice Online as part of its Collabora Cloudsuite.

LibreOffice 6.0 is available for Linux, Mac OS, and Windows (<https://www.libreoffice.org/download/download/>).



**LibreOffice**  
The Document Foundation



## Red Hat Acquires CoreOS

Red Hat has acquired CoreOS, one of the most innovative players in the Kubernetes and container space, for \$250 million. The acquisition was completed in January 2018, making it the first acquisition of Red Hat in 2018.

Red Hat and CoreOS are among the top contributors of the Kubernetes project. The developers of the two companies have already been working together, and now they will work under the same umbrella.

The flagship product of CoreOS is Tectonic, a fully upstream distribution of Kubernetes that automates operational tasks, enables platform portability, and supports multicluster management. Red Hat already has a container application platform called OpenShift that brings Docker and Kubernetes to enterprise customers for orchestration and management of containers.

Tectonic, along with other CoreOS open source projects, will strengthen Red Hat's container technologies. More importantly, Red Hat now employs some of the brightest people in the Kubernetes world. CoreOS has been a very disruptive company in the distributed systems space. CoreOS introduced Container Linux, a Linux distribution designed for containerized workloads. Red Hat responded with Red Hat Enterprise Linux Atomic Host (RHEL Atomic Host), a variant of Red Hat Enterprise Linux Server (RHEL Server) designed and optimized to run Linux Containers. Now Red Hat has become the owner of Container Linux, just the way they own CentOS, an open source project that once competed with Red Hat Enterprise Linux in the server space.

CoreOS has been instrumental in the creation of container image and run-time specification. They came out with rkt, a competitor of Docker Engine to solve security-related concerns around Docker Engine. Docker later addressed these concerns by creating containerd and donating it to CNCF as an open source project. CoreOS also created etcd, a distributed key-value store for a distributed system's most critical data.

Another widely used product from CoreOS is Clair, an open source project for the static analysis of vulnerabilities in application containers (currently including appc and Docker). All of these are fully open source projects that are used widely by the companies in the container space.

CoreOS joins the growing list of companies that Red Hat has acquired in the last couple of years to bring cloud native technologies to the company. In 2017, Red Hat acquired Codenvy, a provider of cloud-native development tools, and Permabit Technology Corporation, a provider of software for data deduplication, compression, and thin provisioning.



## Red Hat Enterprise Linux 7.5 Beta Out

Red Hat has announced the beta of RHEL 7.5, which supports alternative architectures, with variants available for IBM Power, IBM z System, and ARM deployments, as well as x86.

Security is certainly the key highlight of this release. Red Hat said in a press release that RHEL 7.5 beta comes with security improvements and usability enhancements for cloud and remotely hosted systems that can more securely unlock Network Bound Disk Encrypted devices at boot time, designed to eliminate the need for manual intervention in an often inconveniently timed boot process.

This release also integrates Red Hat Ansible Automation with OpenSCAP, which enhances the ease of automating the remediation of compliance issues and enables administrators to scale policies across their environment more efficiently.

RHEL 7.5 beta also improves compliance for accurate timestamping and synchronization needs with the addition of failover with bonding interfaces for Precision Time Protocol (PTP) and Network Time Protocol (NTP).

RHEL 7.5 Beta enhances usability for Linux administrators, Windows administrators new to the platform, and developers seeking self-service capabilities along with an easier-to-use cockpit administrator console. The console is designed to simplify the interface for managing storage, networking, containers, services, and more for individual systems.

## MORE ONLINE

### Linux Magazine

[www.linux-magazine.com](http://www.linux-magazine.com)

### ADMIN HPC

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

#### Log Everything • Jeff Layton

To be a good HPC system administrator for today's environment, you need to be a lumberjack.

#### 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/>

#### Proxmox Virtualization Manager

Martin Loschwitz

Proxmox has been working for years on their virtualization manager that does not require a cloud. We look at the improved reliability and new features in version 5.0.

#### Better Compression of Web Pages

Stefan Wintermeyer

Google develops a software tool that is a genuine alternative to Gzip, with improved website compression rates that save bandwidth for server operators.

#### 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.

### ADMIN DevOps Focus

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

#### Auditing Docker Containers in a DevOps Environment • Chris Binnie

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

Sys admins will love the automated creation of a “known-good” bootable snapshot to help speed recovery and rollback after patching, helping IT teams feel more confident that their systems are in working order.

Users can download the beta for testing (<https://access.redhat.com/downloads/>).

## Torvalds Calls Intel’s Patch Garbage

Intel’s woes are not going away. After releasing the patches for Spectre/Meltdown, the company is asking users to stop installing these patches until a better version is out.

“We recommend that OEMs, cloud service providers, system manufacturers, software vendors, and end users stop deployment of current versions on specific platforms,” Navin Shenoy, executive vice president of Intel wrote in an announcement, “as they may introduce higher than expected reboots and other unpredictable system behavior.”

Red Hat has already reverted the patches that the companies earlier released for the RHEL family of products, after reports of rebooting problems.

Linus Torvalds, the creator of Linux, reserves the harshest words for Intel. “... I really don’t want to see these garbage patches just mindlessly sent around,” wrote Torvalds on the LKML mailing list.

Though not everyone on the mailing list thought it was such a bad thing. One maintainer said, “Certainly it’s a nasty hack, but hey – the world was on fire and in the end we didn’t have to just turn the data centres off and go back to goat farming, so it’s not all bad.”

Another maintainer chimed in and said, “As a hack for existing CPUs, it’s just about tolerable — as long as it can die entirely by the next generation.”

Torvalds didn’t buy either argument. “That’s part of the big problem here. The speculation control cpuid stuff shows that Intel actually seems to plan on doing the right thing for Meltdown (the main question being *when*). Which is not a huge surprise, since it should be easy to fix, and it’s a really honking big hole to drive through. Not doing the right thing for Meltdown would be completely unacceptable,” said Torvalds. “So the IBRS garbage implies that Intel is *not* planning on doing the right thing for the indirect branch speculation. Honestly, that’s completely unacceptable too.”

## More than 2,000 WordPress Sites Infected by Malware

If you are a WordPress admin, you need to check if your site is infected by the infamous cloudflare.solutions malware.

A few month ago, researchers at Sucuri, a web security company, discovered two infections related to cloudflare.solutions. The company reports that the malware and attack are back.

The malware is a bigger threat to WordPress-powered sites that offer e-commerce services because it is designed to steal payment details. “If hackers manage to steal the admin credentials, they can just log into the site without relying on a flaw to break into the site,” wrote *The Hacker News*.

Although the new attack is not as widespread as the original, the return of the malware does show that website admins didn’t protect their sites after the first attack. It’s very likely that most WordPress admins may not even be aware of the problem.

According to *The Hacker News*, “More than 2,000 WordPress websites have once again been found infected with a piece of crypto-mining malware that not only steals the resources of visitors’ computers to mine digital currencies but also logs visitors’ every keystroke.”

“To clean up a website that has been compromised with this infection, you’ll need to remove the malicious code from the theme’s `functions.php`, scan the `wp_posts` table for possible injections, change all WordPress passwords(!), and update all server software, including third-party themes and plugins,” wrote Sucuri in a blog post.

If you are a WordPress admin, you might want to try the Sucuri plugin to check how to clean your website of any infected code.



© wamsler, 123RF.com

# ISC HIGH PERFORMANCE 2018

FRANKFURT, GERMANY

- 5 DAYS
- 450 SPEAKERS
- 150 EXHIBITORS
- 3500 ATTENDEES



SAVE THE DATE!



Platinum Sponsors



Beyond Moore's Law | Exascale Systems | Climate Change | HPC and Electric Power Grid Control |  
Big Data Analytics | Cosmology and HPC | Human Brain Modeling and Related Big Data Challenges |  
Future Applications for Quantum Computers | Robotics | What's New with Cloud Computing for HPC? |  
Future Challenges for Programming Models and Languages | The Rise of Containerized HPC |  
Artificial Intelligence on HPC Platforms

[isc-hpc.com](http://isc-hpc.com)

# Zack's Kernel News



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

*By Zack Brown*

## 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**.

## Preventing the Kernel from Tainting

Matthew Garrett recently posted a patch to allow users to select at compile time whether the kernel would be “tainted” when loading unsigned modules. In the Linux world, tainting refers to whether the kernel is entirely open source or not. For example, if a user loads a binary-only module from a third-party vendor, it would taint the kernel because the module is not open source. The reason this is significant is that the Linux developers don’t want to respond to bug reports that could be related to closed-source blobs of code to which they don’t have access. Typically if a user sends in a bug report using a tainted kernel, the Linux developers will ask that user to reproduce the bug using an untainted kernel.

It wasn’t immediately clear why Matthew wanted this code added to the kernel. He only said, “Distributions may wish to provide kernels that permit loading of unsigned modules based on certain policy decisions.”

Rusty Russell said Matthew’s explanation was too vague and asked for clarification.

Matthew pointed out that the kernel also included a run-time license checker that reads a variable included in each module, specifically identifying the license under which the module was released. Anything other than an open source license, he said, would taint the kernel. In which case, he reasoned, there was no need to taint unsigned modules, since the license checker would do it anyway.

This didn’t make sense to Rusty, since it didn’t answer the question of why one would want to add Matthew’s patch in the first place. So Matthew offered this explanation: “A distribution may ship a kernel with signed modules. In some configurations, the signatures are irrelevant – there’s no mechanism to verify that the correct kernel was loaded in the first place, so for all you know the signature validation code has already been removed at run time. In that scenario you’re fine with users loading unsigned kernel modules, and there’s

no benefit in tainting the kernel. But the same kernel may be booted under circumstances where it *is* possible to validate the kernel, and in those circumstances, you want to enforce module signatures, and so `sig_enforce` is set.”

Matthew went on to say that a main motivation for this patch was “in order to avoid dealing with user questions about why loading locally built modules now taints the kernel.”

Jessica Yu came into the discussion at this point, saying that she still didn’t see the point of going to all this trouble to disable a useful feature. She said, “I find it a bit contradictory to have `CONFIG_MODULE_SIG` enabled and at the same time expect the kernel to behave as if the option wasn’t enabled.”

But Matthew was still unable to give a clear explanation of why the patch would be useful. At one point, Ben Hutchings speculated, “The current state of affairs is that Debian doesn’t have the mechanism in place to sign modules with a trusted key. If we were to allow third parties to add signatures in some way (I think that’s what Matthew’s interested in doing), we would have to enable `CONFIG_MODULE_SIG`, but that would cause modules to be tainted by default.”

Matthew didn’t clarify, and the thread petered out.

This kind of thing happens occasionally. A company will want to construct some sort of loophole in Linux policy and will try to justify it either vaguely or with made-up scenarios that aren’t what the company really wants to do.

## Encrypting `printk()` Output

Dan Aloni recently posted a patch to encrypt all `printk()` output (i.e., all kernel messages to the console). In this way, Dan hoped to eliminate possible attack vectors, by preventing hostile users from seeing what the kernel was up to. He said, “any information leak has the potential to make exploitation easier,” and pointed out that he didn’t use any special encryption scheme in his code, only the normal cryptographic keys that the

kernel already used to sign firmware blobs and kernel modules.

Steven Rostedt replied that this did not seem like a very useful feature. Apparently, the only real users might be device manufacturers, hiding kernel output from regular users as well as from attackers. Steven said he wasn't opposed to a feature like this in principle, but only if folks like Google and other phone producers said it would truly be useful.

While not working for Google, Daniel Micay, nonetheless an android security engineer, pointed out that Android already used SELinux to prevent regular users from seeing whatever `printk()` sent to `dmesg`. He added, "Unprivileged processes including apps can't access `dmesg`, `debugfs`, `sysfs` (other than a tiny set of exceptions), `procfs` outside of `/proc/net` and `/proc/PID` (but it uses `hidepid=2`), and a few whitelisted files, etc. That's mostly done with SELinux along with using it for `ioct1` command whitelisting to have per-device, per-domain whitelists of commands instead of using their global `seccomp-bpf` policy for it."

To Steven, this seemed to indicate that Android would not enable Dan's feature, since `printk()` output was already restricted via SELinux. Absent any further information from device engineers, he said, Dan's patch didn't seem worth it.

The thread died there. It's interesting, because it leaves open the question of why the patch was developed in the first place. Was it to stop attackers from breaking through security? Was it to stop regular users from rooting their device? It's unclear.

## New Kernel Bug Reporting Bot

Dmitry Vyukov announced the new syzbot open source automated bug identification system tailored for the Linux kernel. The idea behind this project was to automate the entire process, from bug identification, to sending the reports into the Linux developers, tracking the progress of each bug as the developers submit patches to fix them, and testing those proposed patches.

The problem was in the bug tracking portion of syzbot. There was currently no way to associate a given bug report with a future patch that might be intended to fix it. And so far, Dmitry had not discovered any way to purely automate a solution –

he felt he would need collaboration from the developer community.

One possibility he had considered would be to use the ancient and still functioning Kernel.org Bugzilla system to track everything. Another would be for developers to tag Git commits with an ID provided by syzbot.

Greg Kroah-Hartman was opposed to using Bugzilla, but wanted to stick to using the official Linux Kernel Mailing List if possible.

Andrey Ryabinin suggested adding the ID number to the *From:* line of the syzbot emails. For example, *From: syzbot-{hash} <syzkaller@googlegroups.com>*. This way the tag would be likely to be kept via the normal kernel development sign-off process.

Linus Torvalds liked Andrey's idea, but offered a slight modification:

*Put the hash in the original email instead, and use "+" instead of ",", because that's the standard email ("latter part doesn't matter"). Skip the proper name part entirely, since it doesn't really add anything.*

*And then make it show up in the Cc: part or Reported-by:, so it would just look something like Reported-by: syzkaller + 1234abc@googlegroups.com and you're all done.*

Dmitry liked this idea and implemented it in syzbot.

Meanwhile, Eric Biggers suggested that syzbot should do more to expose its internal data to developers. He wanted a dashboard of some kind, so that developers could see the status of each bug, when it was last seen, how frequently, and so on. He also felt it would be useful for syzbot to automatically send out reminders about old bugs that hadn't been fixed yet.

And in the case of bugs that were still technically unsolved, but that hadn't been reproduced in awhile, Eric felt they should be automatically invalidated, to prevent bug reports from piling up without bound, yet also without usefulness.

Dmitry replied, saying he'd like to expose the UI, as soon as he'd done some more work to spruce it up a bit. But he added that all the data was accumulating, so it would all be visible as soon as the dashboard came online.

Dmitry was reluctant to automate email reminders though, saying he was

afraid they wouldn't be warmly received by the developers.

Regarding a public UI, Ozgur Karatas remarked that a graphical UI probably wouldn't be necessary and wouldn't be as valuable as a command-line interface.

But Dmitry said the web-based UI was already partially implemented so it would probably appear ahead of a command-line tool. Also, syzbot was implemented in "the cloud," so it didn't lend itself naturally to the command line. However, the whole thing was open source, so Dmitry invited Ozgur, and anyone else who pleased, to contribute patches to the UI or any other part of the code.

Meanwhile, Eric W. Biederman voiced some issues with syzbot's overall value. He said, "It tells us things are wrong but not how to reproduce the problem. Apparently syzbot will test fixes, but that doesn't help when more information is needed to track down the problem. The long of the short of it is that I don't care about bug reports that no one can reproduce and no human cares about." Eric also pointed out that "syzbot is written in a language (Go) that switches which kernel thread things run in at arbitrary times. That is absolutely not productive to understanding what is happening when things break. I have heard too many complaints from container run times that they can't make what should be a couple of lines change but is completely non-trivial because someone chose Go for their implementation language. Whatever bene-

fits Go has, it is not a programming language I would choose for fine and reproducible control of kernel interfaces."

Dmitry replied that syzbot did try to include information about how to reproduce bugs. But he said, "Unfortunately, localizing kernel bugs is hard and is not possible in all cases. The root cause of this is actually in the kernel itself, not in syzbot. Things would be much simpler if we would work on a single-threaded, deterministic user-space library. Then we would get precise reproducers in 100% of cases. But kernel is a concurrent, parallel, non-deterministic system that constantly accumulates state."

Dmitry also said that in a lot of cases, the bug was trivial to understand simply by looking at the crash report – missing input checks, off-by-one errors, and so on.

As far as the programming language issue, Dmitry replied, "this is not true. The part of syzkaller that actually executes syscalls is written in C++ from day one." He continued, "It does explicit, manual thread scheduling; compiled as static binary to avoid any variance due to dynamic loading; does not use C++ run-time support nor `malloc` to avoid unexpected `mmap` calls. This is mostly for the reasons you outlined."

The discussion petered out around there. But the essential issues seemed to have been resolved – bug tracking. Beyond that, the rest will probably evolve over time to suit the developers' needs. ■■■



## THE ULTIMATE raspberry pi GEEK ARCHIVE

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



**Order Now!**

[shop.linuxnewmedia.com](http://shop.linuxnewmedia.com)

# OPEN SOURCE DATA CENTER CONFERENCE

JUNE 12 – 13 | BERLIN | GERMANY



**SIMPLIFYING COMPLEX IT  
INFRASTRUCTURES WITH OPEN SOURCE**



**REGISTER NOW!**

**FIRST SPEAKERS ONLINE!**

**MITCHELL HASHIMOTO**

HASHICORP

**MIKE PLACE**

SALTSTACK

**GIANLUCA ARBEZZANO**

INFLUXDATA

**OSDC.DE**

Close to the {code} with Dell’s Josh Bernstein

# COMMUNITY MATTERS

Dell’s expansive {code} project is a cornerstone of the company’s open source strategy. Dell Technologies VP Josh Bernstein talks about {code} and the value of open source.

By Swapnil Bhartiya



**A**s a vice president of technology at Dell Technologies, Josh Bernstein has many responsibilities, but the project closest to his heart is {code} [1], which Dell inherited by acquiring EMC. {code} is a collection of open source projects maintained or supported by Dell. According to Bernstein, “A lot of times older legacy organizations don’t really get open source. They don’t understand how to market it. {code} was created to develop a brand for the community, to make Dell relevant to a community that did not interact with them, didn’t trust large enterprise brands. The {code} team is a collection of open

source engineers and advocates inside Dell Technologies that are working to better support the community and engage Dell Technologies with the community.”

Dell is not new to open source. They worked on Dynamic Kernel Module Support (DKMS) and Redfish, and they wrote open source drivers for hardware on the Linux kernel. Some of the big open source projects that came directly from Dell include EdgeX Foundry, which is now a Linux Foundation collaborative project. OpenSwitch is another big project to which Dell contributed. Most Linux users may be familiar with the project Sputnik, which offers high-end Dell systems with fully supported Linux distributions.

“Dell was doing all this open source work, but nobody knew about it,” said Bernstein. “EMC was doing a lot of work specifically in the storage data center infrastructure. When the two companies were combined, there was actually a lot of synergy between the two. We naturally found each other and just increased everybody’s voice in the whole community.”

Open source is as much about technology as it is about people and culture. It’s also about legal frameworks. A small mistake may release a company’s intellectual property as open source or make them a violator of the GNU GPL. Companies need to create an internal environment that nurtures the open source mentality while offering legal frameworks to developers.

One of the projects within the {code} umbrella is called DevHigh5 [2]. The DevHigh5 project allows any developer in the company to bring their own project to the open source community, offering the internal support and legal guidance needed to establish a successful project.

“We help them with legal guidelines, licensing, infrastructure and make sure that whatever project is released its done with the best interest of the open source community. It has to be transparent and it has to be consumable by the community,” said Bernstein.

The {code} team contributes to more than 100 projects in the open source



community. The project website provides links to an impressive list of sites hosted by {code}, as well as other important open source projects to which the team is contributing. The {code} team maintains a GitHub page that has 67 repositories [3].

## Container Storage Interface

One of the most ambitious projects the {code} team is working on is the Container Storage Interface (CSI) [4]. Although Kubernetes is enjoying an explosive adoption, Bernstein believes that it’s not getting adopted in production as heavily as it should.

“There is a lot of dev-tests; there is a lot of internal PoC (proof of concept), but when you want to run it in production, you want to run meaningful and valuable applications that add business value,” said Bernstein. “Many of those applications require storage. In an effort to fill a gap in the ecosystem and make it a production ready platform, storage is very important.”

CSI is designed to provide storage orchestration between container run times and back-end storage infrastructure so persistent applications can run in containers. Many container storage interfaces exist, but currently there is no standardization.

CSI aims to fix the existing fragmentation. The initial specification for CSI was led by Mesosphere. Later on the Kubernetes community joined the efforts.

“Our contributions have been in implementing CSI inside of Kubernetes and also providing two drivers to the open source community to use CSI with our products,” said Bernstein.

At KubeCon Austin, the team released CSI drivers for Dell EMC ScaleIO (csi-scaleio) and vSphere (csi-vsphere). The {code} team also released a generic NFS driver, a generic VFS driver, and a generic block driver.

“We are working very hard so that CSI becomes a Cloud Native Computing Foundation (CNCF) project,” said Bernstein, “We really think that storage has been left out of the container ecosystem for so long.”

## Business and Open Source

Selling open source doesn’t always come through direct sales or subscriptions; it can also come through becoming a core component of the open source ecosystem. You don’t make money *from* open source; you make money *around* open source. {code} is a great model for other companies to follow, to build an open source culture within the company and also create actual projects that are useful in real use cases. ■■■

### Info

- [1] Code community: <https://thecodeteam.com/>
- [2] DevHigh5: <https://thecodeteam.com/devhigh5/>
- [3] Code team GitHub page: <https://github.com/thecodeteam>
- [4] CSI: <https://github.com/container-storage-interface>

 Find us on  
**Facebook**

<http://www.facebook.com/linuxpromagazine>

**The intersection of  
DEVELOPMENT and OPERATIONS**  
**Check out our new ADMIN DevOps corner!**  
[www.admin-magazine.com/DevOps](http://www.admin-magazine.com/DevOps)

SPONSORED BY



**Linux  
Professional  
Institute**



We compare four recent web browsers

# Faster and Higher

Will the Firefox 57 “Quantum” release help Mozilla regain its former glory?

We compare the latest Firefox with the Chrome, Opera, and Vivaldi

browsers. *By Ferdinand Thommes*

The release of Firefox 57 “Quantum” in mid-November caused a stir in the IT community. The Firefox browser was once a state-of-the-art browser, with a huge market share and millions of loyal followers, but over the years, newer and more advanced alternatives have taken away some of the limelight. After years of development and 12 months of re-engineering, the Mozilla developers think they have what they need to get back in the game. The new Firefox is described as twice as fast as the version released a year ago, with a 30 percent savings in memory usage. In addition, Firefox 57 is supposed to be as fast as Google Chrome.

The word on the street is that Firefox is ready for a head-to-head performance comparison. We decided to test Firefox versions 51 and 57 against Chrome 62. We also tested the Opera 49 and Vivaldi 1.12 browsers. Opera, Vivaldi, and Chrome all use the Blink rendering engine – a WebKit fork – so we did not expect any big differences between Blink-based contestants.


We ran the tests on a mid-range Lenovo ThinkPad X220 laptop with 4GB of main memory. The browsers were installed in their default configuration on a newly installed Debian “Sid.” We put all test candidates through a series of special browser benchmarks and ran the benchmarks five times on each browser to ensure that no outlier distorted the mean.

The tests used were Speedometer [1], JetStream [2], ARES-6 [3], Kraken [4], WebXPRT [5], MotionMark [6], and HTML5test [7]. These benchmarks are designed to determine the speed at which the browsers process HTML and JavaScript. The WebXPRT benchmark goes even further and also includes image processing and DNA sequencing in the

## Browser Benchmarks

The study described in this article used the following benchmarks. Descriptions are from the project websites.

- Speedometer [1] – a browser benchmark that measures the responsiveness of web applications. Speedometer uses demo web applications to simulate user actions such as adding to-do items.
- JetStream [2] – a JavaScript benchmark suite focused on the most advanced web applications.
- ARES-6 [3] – measures the execution time of JavaScript’s newest features.
- Kraken [4] – JavaScript benchmark.
- WebXPRT [5] – uses scenarios created to mirror the tasks you do every day.
- MotionMark [6] – a graphics benchmark that measures a browser’s capability to animate complex scenes at a target frame rate.
- HTML5test [7] – measures the level of support for HTML5.



### Mixed Results

We put the browsers through a series of freely available benchmarks. See the box entitled “Browser Benchmarks” for a description of the tests used in this article. In the Speedometer (Figure 1) and ARES-6 (Figure 2) tests, Chrome and its two relatives are clearly in the lead. With WebXPRT (Figure 3) and Kraken (Figure 4), Firefox 57 is ahead, but not as clearly. With its strong focus on graphics performance, the MotionMark test (Figure 5) sees Opera and Vivaldi as winners and Firefox 57 slightly ahead of Chrome. JetStream (Figure 6) ranks Chrome as the winner with Opera in second place.

The HTML5test measures the degree of browser compliance with the HTML5 standard. Of a possible 555 points, the browsers with the Blink engine reach around

test run.

ARES-6 focuses on the speed at which new JavaScript routines are implemented. Except for the ARES-6 and Kraken benchmarks, higher values mean better results.

The HTML5test determines the degree of a browser’s compliance with the HTML5 standard. In addition, we looked at the candidates’ memory consumption when opening 10 tabs at browser launch time.

### Head to Head

At first glance, the figures reveal that Mozilla Firefox 51 was far off the performance of the other browsers – at least in terms of speed. If you have ever run Firefox alongside Chrome, you don’t need statistics to confirm that Firefox releases of the recent past have had trouble keeping up.

Firefox 57, however, tells a different story. In some tests, the new Firefox came out ahead of Chrome; in others, it was still behind, but the gap has narrowed. The browsers that use the Blink engine exhibit very similar behavior; sometimes Chrome is slightly in front, sometimes Opera or Vivaldi, sometimes both. For the most part, the variations lie within a range of nuances that play no role in real use.

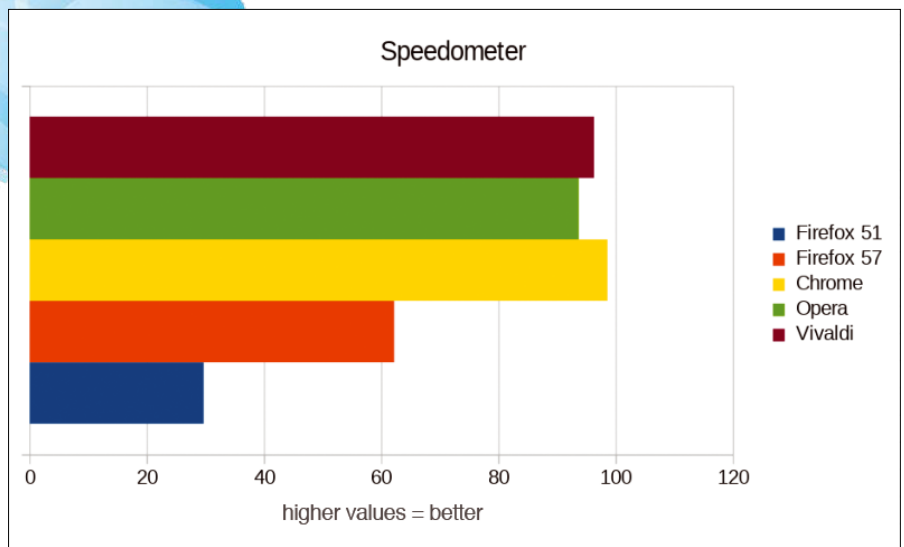


Figure 1: Speedometer sees the browsers with the Blink engine clearly in front. Firefox 57 closes the gap, doubling the score of Firefox 51.

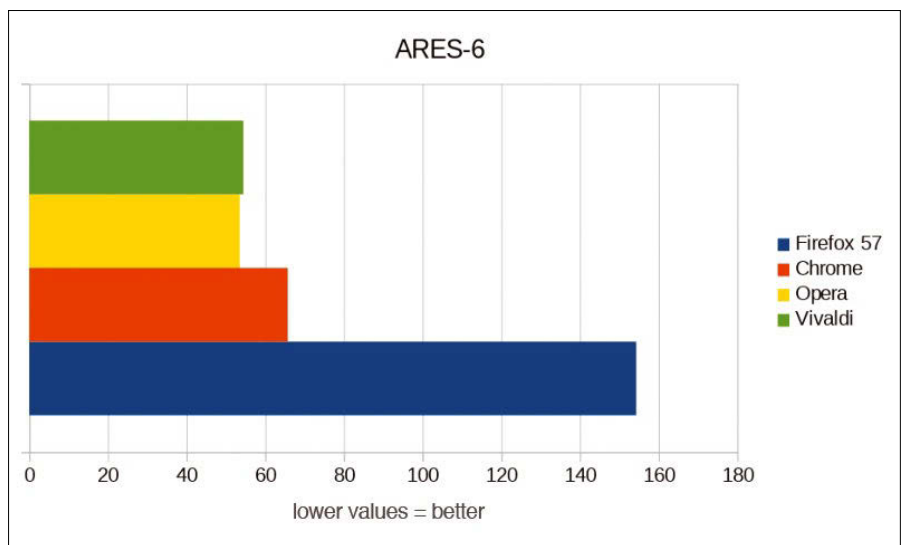
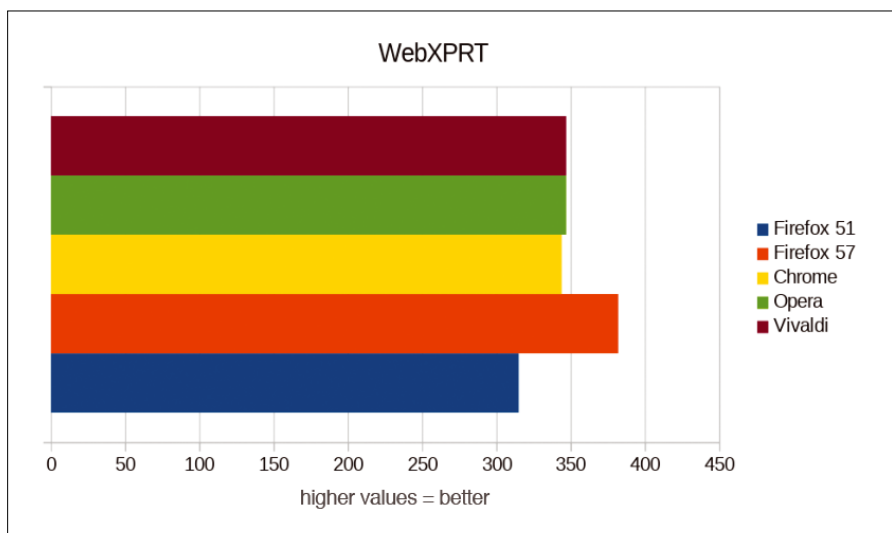
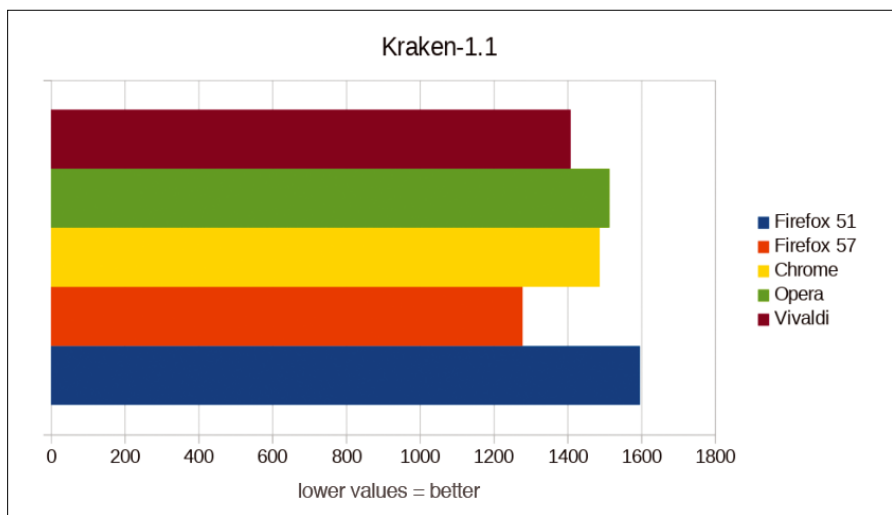


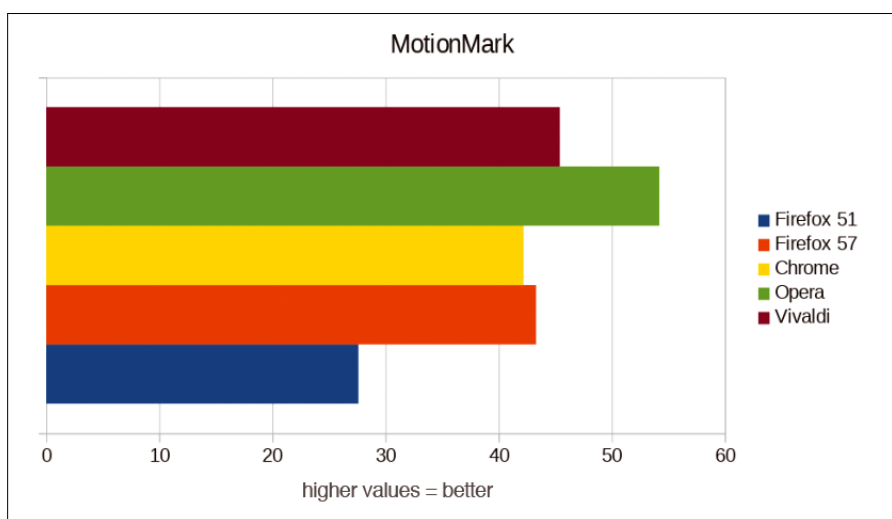
Figure 2: ARES-6 indicates an advantage for Chrome, Opera, and Vivaldi. With Firefox 51, the test stopped.



**Figure 3:** In WebXPRT, the most comprehensive test of our course, Firefox 57 is the best. The rest of the field is approximately level, with Firefox 51 trailing behind.



**Figure 4:** The Kraken test shows a level field, with Firefox 57 in front and Firefox 51 trailing behind.



**Figure 5:** The results of MotionMark, which tests graphics and animation capabilities, put Opera and Vivaldi on top, followed by Firefox 57 and Chrome.

520 points, and the two Mozilla browsers lag slightly behind with 465 and 481 points (Figure 7).

On fairly up-to-date hardware, the detectable speed differences between the browsers are confined to narrow limits. Firefox 57 feels as though it's back in line with Chrome and others, even though it's still slightly behind in some of the measurements.

### Mozilla Saves Memory

The situation is completely different when it comes to memory usage. We opened 10 tabs with the same web pages in each of the browsers and then determined memory usage. For Mozilla browsers, you can type `about:memory` in the address line to check the memory. Click on *Measure* in the top-left corner; Firefox will calculate the RAM usage and break it down to individual processes. The uppermost displayed value (`--explicit`) indicates the overall memory usage.

For browsers that use Blink as a web engine, use the Task Manager to determine memory usage. You can open the Task Manager via the More Tools menu or the key combination `Shift + Esc`. The Task Manager lists the open web pages separately with their memory consumption. The usage of the browser itself and the sandbox with GPU processing are shown separately at the top of the list.

The result of the memory test confirms the widely held belief that Chrome and its descendants are real storage guzzlers. Whereas Chrome, Opera, and Vivaldi use between 337MB and 376MB of memory when opening the 10 web pages, Firefox 51 is content with 249MB, and Firefox 57 occupies just 170MB.

If you work with many open tabs and don't have much RAM, Firefox is definitely better. Computers with 4GB of main memory, which are still quite common, quickly reach their limits with Chrome, Opera, and Vivaldi if you add many open tabs to normal system consumption (Figure 8).

### Quantum Flow Remains Interesting

Mozilla still has a few arrows in its quiver. The Quantum project, to which Firefox 57 owes its nickname, is still a work in the progress. In the near future,



Quantum rendering will become more significant. The WebRender [8] project, which aims to speed up the presentation of websites with the help of the GPU, will also bring further optimization. The further expansion of the Quantum project will be known as *Quantum Flow*, and it could bring further benefits to Firefox after the current release.

### Conclusions

If you prefer open source software and have remained loyal to Mozilla and Firefox over the past few years, Firefox 57 will reward your loyalty. With its new Quantum engine, Firefox takes a big step to closing the speed gap with the Blink-based web browsers, and it comes out ahead on memory usage.

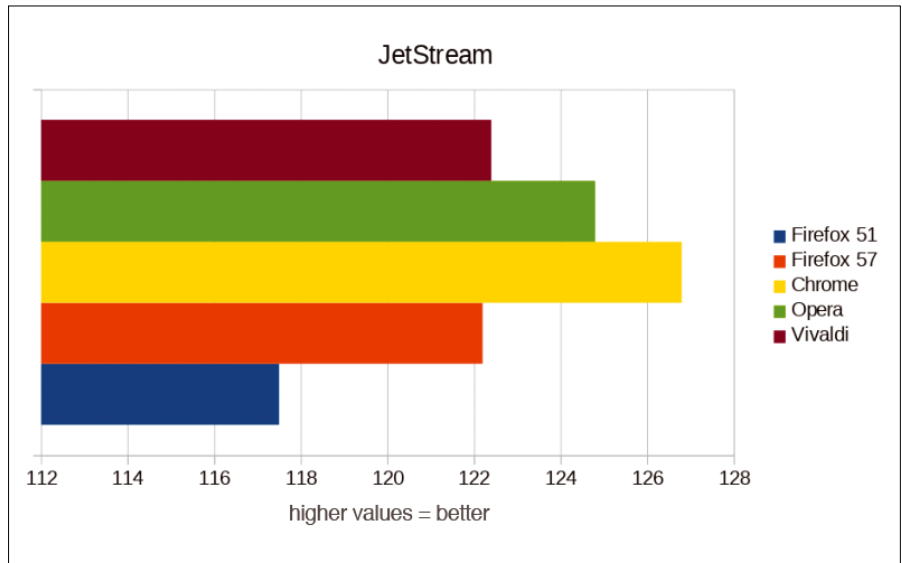
Note that the tests described in this article did not address third-party extensions. In the past, Firefox add-ons often had a negative effect on the surfing experience, due to slow execution or even crashes. With the new plug-in architecture for web extensions, these problems should be a thing of the past, but we will wait until more add-ons are running under the new Firefox before continuing the comparison. ■■■

### Author

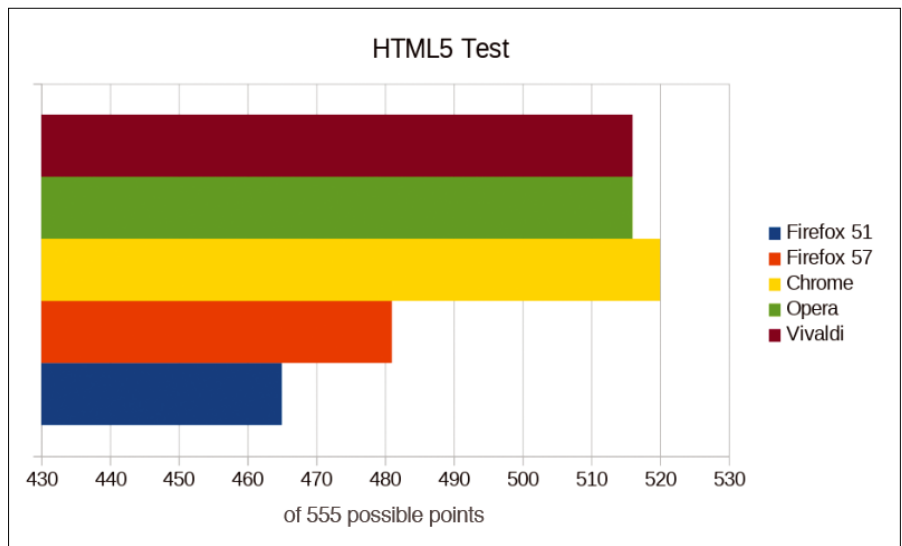
**Ferdinand Thommes** lives and works as a Linux developer, freelance writer, and tour guide in Berlin.

### Info


- [1] Speedometer: <http://browserbench.org/Speedometer/>
- [2] JetStream: <http://browserbench.org/JetStream>
- [3] ARES-6: <http://browserbench.org/ARES-6>
- [4] Kraken: <https://krakenbenchmark.mozilla.org/>
- [5] WebXPRT: <http://www.principledtechnologies.com/benchmarkxprt/webxprt/>
- [6] MotionMark: <http://browserbench.org/MotionMark/>
- [7] HTML5test: <https://html5test.com/>
- [8] WebRender: <https://hacks.mozilla.org/2017/10/the-whole-web-at-maximum-fps-how-webrender-gets-rid-of-jank/>



**Figure 6:** JetStream shows a clear advantage for Chrome and Opera. Vivaldi and Firefox 57 are approximately equal and outpace Firefox 51.



**Figure 7:** Mozilla needs to do some work on W3C HTML5 compliance.



**Figure 8:** A clear winner in memory usage: Firefox 57 is way ahead, followed by its predecessor in version 51.

# ADMIN

Network & Security

COMPLETE ARCHIVE DVD  
ISSUES 00-39

ISSUE 40/2017

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



OVER  
**3,000**  
PAGES OF UP-CLOSE  
SYSTEM  
ADMINISTRATION

## Smart and Practical Information for Network Professionals



---

# 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!

Clear off your bookshelf and complete your ADMIN library with this powerful DVD!

**ORDER NOW!**  
[shop.linuxnewmedia.com](http://shop.linuxnewmedia.com)



Customizing the way your browser stores and organizes data

# Sweeping Up

Web browsers collect a large amount of data about the user's browsing habits. If you care about privacy, you might want to clean up your browser and configure some custom privacy settings. *By Frank Hofmann and Mandy Neumeier*

**T**he web browser is one of the most popular desktop applications. Regardless of whether you prefer Firefox, Chrome, Opera, w3m, Internet Explorer, Safari, or another browser, you probably spend a lot of time on the Internet. To make life easier, modern web browsers (and web servers) collect information about us and our behavior. Thanks to this system of data storage and collection, we benefit from amenities such as automatic completion of URLs and access credentials, as well as predictive content buffering. But keeping all this browser data around also has some disadvantages – especially if you're concerned about privacy. This article offers some tips for cleaning up browser user data. The examples in this article use Firefox 57, but other browsers have similar features.

## Basic Settings

To tweak Firefox settings, choose *Edit | Settings*. Alternatively, you could just type *about:preferences* in the address bar or click on the button with the three horizontal dashes in the upper right corner (Figure 1) and then select the gearwheel icon.

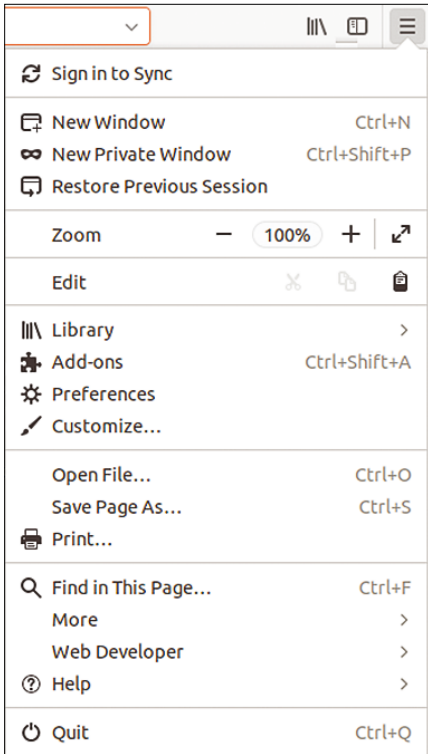
The Settings dialog window contains a selection menu with higher-level terms on the left and detailed information on the right. Many details are hidden in sub-menus and additional windows, making orientation difficult. For now, we will focus on the settings related to data protection (Figure 2). Choose *Privacy & Security* from the menu on the left side of the Preferences page, or enter *about:preferences#privacy* in the address bar.

## History and Activities

The settings in the *Tracking Protection* section let you block websites known to keep track of your activities. With the default settings, Firefox uses the basic version of the Disconnect block list [1]. Click the *Change Block List* button to adjust the filter rules. You can switch between the two modes: *Basic protection* and *Higher protection*. Mozilla recommends that you not block all trackers to allow websites to work as expected. You







**Figure 1:** In Firefox 57, you can open the settings via the menu.

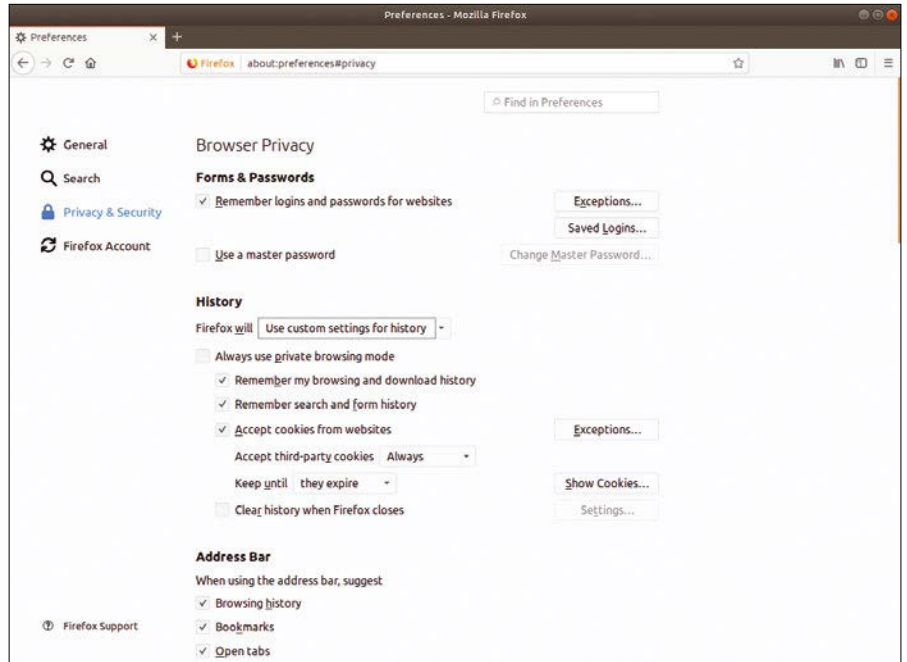
The History section lets you specify how Firefox will log the progress of the websites you visit. By default, Firefox automatically remembers your browsing history. For more privacy, change the setting to *Never remember history*, or choose *Use custom settings for history* for user control of the privacy settings (Figure 3).

To see the history, press Ctrl + H or select *Library | History* (Figure 4).

Press Delete to delete a recorded entry. Use the input field above to start a search for the specified term and quickly jump to the desired entry. To remove the entire history, choose *Library | History | Clear Recent History*.

### Removing Cookies

When you meet someone for coffee, the host will often offer you cookies. If you visit a website with your browser, it too will offer you a few cookies, in this



**Figure 2:** Browser settings for privacy and security.

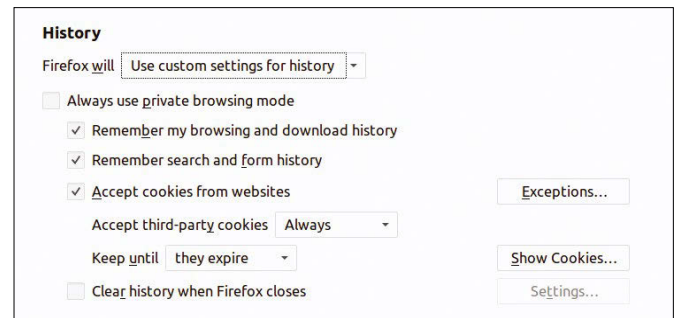
case in the form of small text files saved locally by the browser.

The files you receive help the web page to identify you on your next visit by means of a stored session ID. For example, you don't have to register again, such as with a forum or a webmail. Other information, such as whether and how often you visit the site and which settings you use for surfing, also ends up in the browser's cookie jar.

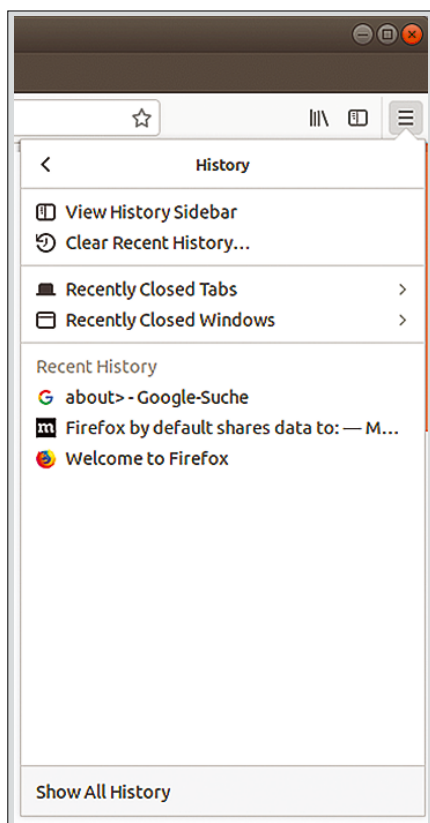
Not all cookies taste good. For example, cookies can be used to inject malware into the browser, to sniff critical data, or to spy on your surfing behavior.

When you call up a website, you often receive what are known as third-party cookies in addition to cookies from the domain you actually accessed. For example, if you navigate to a travel agent's website, you will often find one of the typical Facebook *like* buttons next to the individual offers. This element, which is retrieved directly from the Facebook servers, now sets a Facebook cookie when the travel provider's website is loaded; the social media service can then read this cookie in order to present you with targeted advertising.

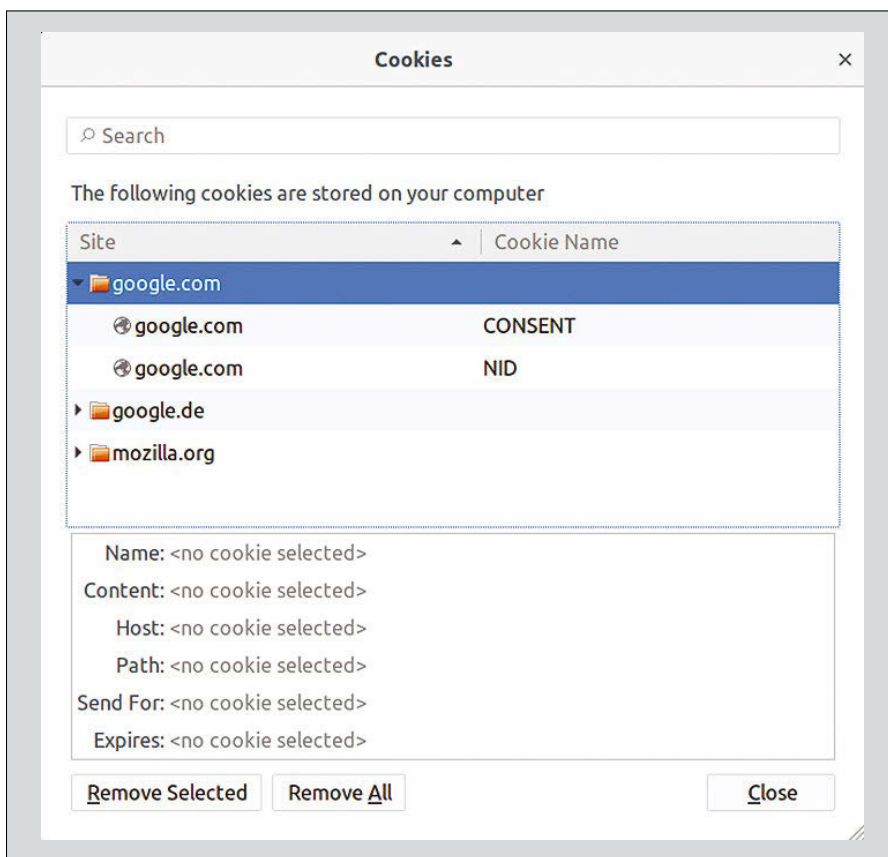
You can view the stored cookies by clicking the *Show Cookies* button (Figure 5) and removing them if necessary. Keep in mind that this restricts the functionality of many websites if



**Figure 3:** In the History settings, you can also set the validity of cookies.



**Figure 4:** Firefox displays the history of web pages opened in the browser in the sidebar or via the menu.



**Figure 5:** There is hardly any current website that does not leave cookies in the user's browser.


you completely block cookies: For example, a typical online store or forum throws you out of the system every time you visit a page, so you would have to log in again and again.

### Emptying the Browser Cache

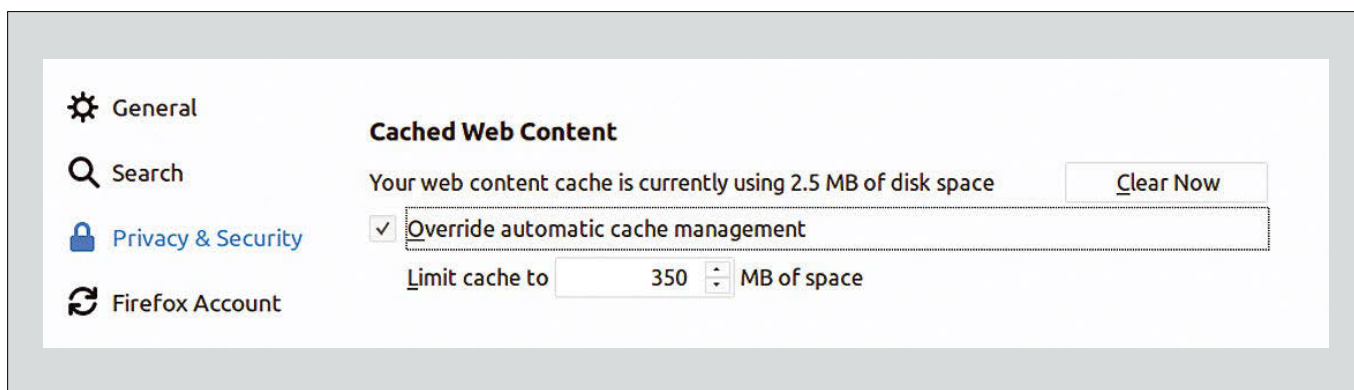
The browser cache serves as a buffer for the Internet pages you visit. The browser stores resources such as page content and images in case you open the website later. For example, large-format images or recurring page elements such as stylesheets do not have to be reloaded from the web when a new page is created. This approach reduces the data volume and speeds up the page build-up.

However, the cache only fully leverages these advantages if it is static, and if the website has not been updated on the server in the meantime. Dynamic content changes every time it is called, which means cached data is always obsolete.

By default, Firefox automatically organizes the cache; the default size is 350MB (Figure 6). If necessary, you can disable automatic cache management and limit the cache size. Clicking on *Clear Now* causes Firefox to remove all traces from the hard disk. Alternatively, you can just delete the data from the past few hours in the menu using *Library |History| Clear recent history*.



**Figure 7:** Deleting form data and other user-specific data in the History menu.



**Figure 6:** The cache's size can be limited in the Firefox settings.



## Deleting Form Data

Firefox remembers everything you type in input fields on the web, such as your name and email address when you subscribe to a newsletter or your address when you order from an online shop. If you are concerned about data privacy, you can delete this data in the menu under *Library* | *History* | *Clear recent history*. In the window that then appears (Figure 7), select the desired period of time, click *Details*, and check *Form & Search History*.

## Blocking and Filtering

In order to reduce the risk of attacks via the web browser, it helps to suppress additional JavaScript code. Popular Firefox add-ons for suppressing scripts include NoScript [2] and uMatrix [3]. Ad blockers such as uBlock Origin [4] ensure that the browser only loads what you want it to load. Firefox also lets you surf in a private window (*New private window*), in which Firefox does not remember cookies or form data.

Another option is to lock the web browser into a secure environment, such as a virtual machine. At first glance, this approach might seem complicated, but it definitely slows down attackers, and it typically doesn't have a noticeable effect on performance given the capabilities of today's hardware. In Firefox 57, the browser now outsources many tasks to a sandbox that runs separately from the rest of the system [5], which reduces the importance of running the browser in a virtual machine.

## Conclusions

Tidying up is often tedious, but it does get rid of the ballast. It is a good idea to clear out the browser cache once in a while; cookies and form data also belong in the virtual trash can. The same applies to the history and bookmarks – discard all entries that you no longer need. ■■■

## Info

- [1] Disconnect: <https://disconnect.me>
- [2] NoScript: <https://addons.mozilla.org/en-US/firefox/addon/noscript/>
- [3] uMatrix: <https://addons.mozilla.org/en-US/firefox/addon/umatrix/>
- [4] uBlock Origin: <https://addons.mozilla.org/en-US/firefox/addon/ublock-origin/>
- [5] "Firefox 57 Brings Better Sandboxing on Linux": <https://www.bleepingcomputer.com/news/security/firefox-57-brings-better-sandboxing-on-linux/>

## Author

**Frank Hofmann** works on the road, preferably in Berlin, Geneva, and Cape Town, as a developer, trainer, and author. He is coauthor of the *Debian Package Management Book* (<http://www.dpmb.org/index.en.html>).

**Mandy Neumeyer** works in tourism and has been living in South Africa for nine years. She is currently earning additional income as a digital nomad.

# 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](http://shop.linuxnewmedia.com)

**GET IT NOW!**

SAVE TIME ON DELIVERY WITH OUR ALTERNATIVE PDF EDITIONS



# 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](http://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](http://shop.linuxnewmedia.com)



## Mozilla reboots with Firefox 57

# Breakthrough

What's new in Firefox 57 Quantum – and why does it matter so much? *By Ferdinand Thommes*

**F**irefox 57 Quantum is probably the most important release in the Mozilla browser's 15-year history. Firefox 57 (FF57) is the biggest upgrade cycle yet for Firefox and the stakes are high: Public acceptance of FF57 is important for Mozilla's future health.

Market observers say the overall Firefox marketshare is somewhere between 8 and 12 percent. By comparison, Google Chrome weighs in closer to 60 percent. Firefox fares better within the Linux community, though, where many users prefer it for ideological reasons while others are attracted to its impressive range of extensions. But even among the Linux crowd, Chrome has been closing the gap in recent years.

## Setting the Stage

One mistake Mozilla made was the failed excursion into Firefox OS. Mozilla spent a great deal of money on the ill-fated mobile OS, and the Firefox OS effort used up valuable developer resources, thus slowing down work on Mozilla's flagship browser. In addition, integrating the online bookmarking service Pocket and the WebRTC service Hello added more bloat. Mozilla also annoyed some customers by displaying ads on the New Tab page and by collecting too much data.

However, the biggest reason for Firefox's decline in popularity is its 15-year-old technical substructure, which locked the entire browser into a single process. Chrome, which benefits from a more modern architecture, has gained a reputation for offering better performance.

The Firefox developers have known for years that they must take action if they want to keep the browser relevant for a new generation of users.

Over the last several years, Firefox has gradually undergone some extensive modifications. Firefox





developers have been slowly completing revisions that have been in the experimental phase for years, including the Servo browser engine written in Rust [1].

The browser has also undergone two facelifts. The first new design – and the start of an unprecedented overhaul – was known as Australis [2]. Originally planned for Firefox 24, Australis only appeared with Firefox 28 in April 2014. The overhaul still divides the user base today, but it succeeded in delivering a more usable and feature-rich interface.

Work on the Servo browser engine began in 2012. The developers leveraged the parallelization capabilities of Rust, which is syntactically similar to C, to enable a multiprocess architecture with multiple threads on several cores, as well as GPU-supported browsing.

With the Servo engine, Firefox moved away from the single-process architecture that prevented it from keeping up with the speed and stability of other modern browsers. Rust is faster than C++; it also makes applications less vulnerable by preventing memory access errors and buffer overflows. At the same time, Rust is designed to rule out the most frequent causes of crashes and avoid race conditions.

Servo is not fully implemented in the official Firefox release version; however, Mozilla began using parts of Servo in its previous web engine, Gecko, some time ago. The first component to replace the previous C++ code in Firefox 48 (FF48) was the MP4 metadata parser. Replacing the metadata parser let the developers leverage Rust's security benefits in an area that is constantly under attack.

## Quantum Brings Everything Together

Over a year ago, Mozilla released the Quantum [3] project, a new web engine bolstered with Servo technologies. Quantum's goal is to provide a foundation for multiprocessing that will improve performance, as well as prevent a frozen tab or a crashed web page from paralyzing the entire browser.

A subset of Quantum called Electrolysis has actually been around since 2009 (E10S) [4]. Development was interrupted several times and then picked up speed again in 2014 with a strengthened team. In FF48, a very limited number of users were allowed to test the new architecture, but only without add-ons. At first, only two processes were used, one each for the interface and content. E10S was successively distributed to other users in later releases. The goal was sandboxing for tabs and isolation of extensions in separate processes. In FF57, the developers have further strengthened the sandbox to avoid endangering the entire system in the event of a security threat. It is thus no longer possible for the sandbox to access your home directory.

## Add-Ons and Other Changes

Mozilla has also been working to make Firefox extensions more compatible

with other browsers, which will make it easier to develop add-ons. The name of the project behind this change is WebExtensions API [5]. Many developers have complained that Mozilla mainly provides high-level APIs, while many add-ons require low-level APIs. The WebExtensions API, which is required for FF57 onwards, forces you to convert all add-ons to the new system. Most of the widely used extensions are available now in WebExtensions format, but the change has caused some programmers to walk away from Firefox add-on development. Others have tried to rewrite their extensions, but have failed because of missing or incomplete interfaces.

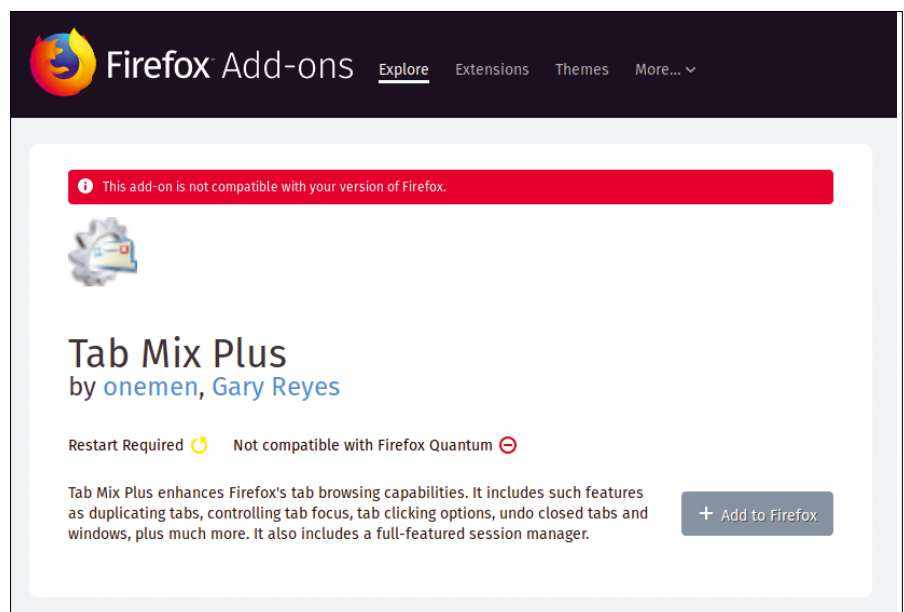
Many users have been worried about preserving their favorite add-ons. Add-on support is especially important for Mozilla because its popular extension library is one of the main reasons for using Firefox.

As of FF57, popular add-ons, such as DownThemAll, Beyond Australis, OmniSidebar, Puzzle Bars, Classic Toolbar Buttons, and Forecast Plus, are no longer available, at least for the time being. Other popular add-ons, such as Tab Mix Plus (Figure 1), are in transition to the new format.

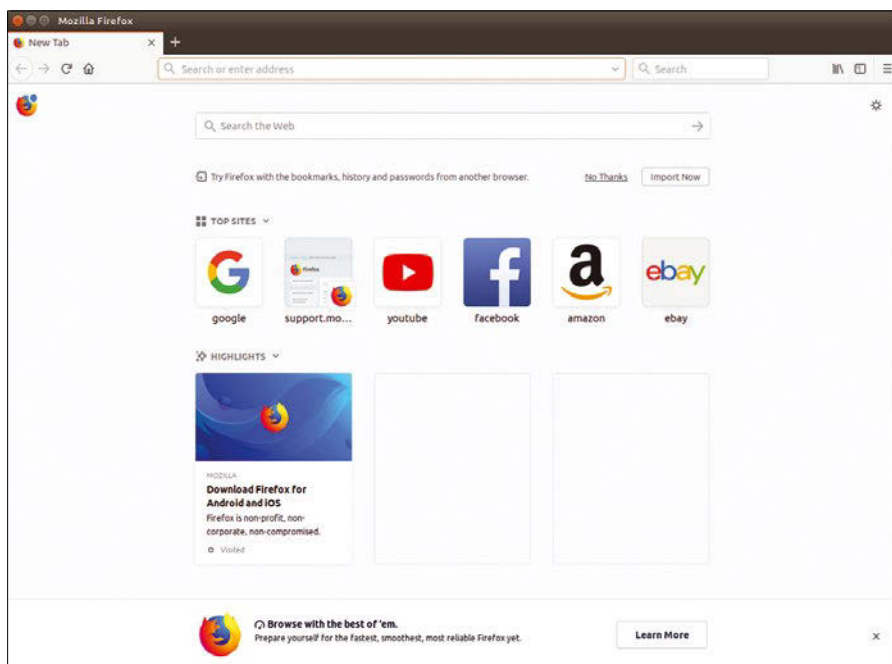
However, in the long run, the switch to the new add-on architecture appears to be an essential step. The previously used technologies, such as the XML User Interface Language (XUL), were outdated and incompatible with other extension standards, often causing browser crashes.

With the Photon UI [6] project, developers have modified the user interface, menus, and settings. Among other things, Photon UI is intended to work better with high-resolution displays. In addition, the technology detects whether the input comes from a mouse or a touch screen and adjusts the menu sizes accordingly. Here, too, the developers initially faced rejection.

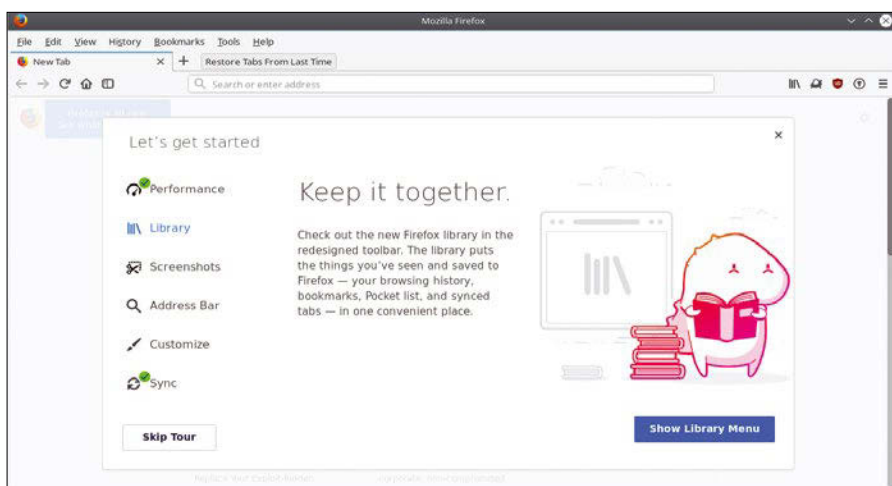
A new CSS engine written in Rust runs in parallel on several CPU cores, both on the desktop and on mobile devices. In recent months, the developers have also eliminated almost 500 bugs that affected browser performance.



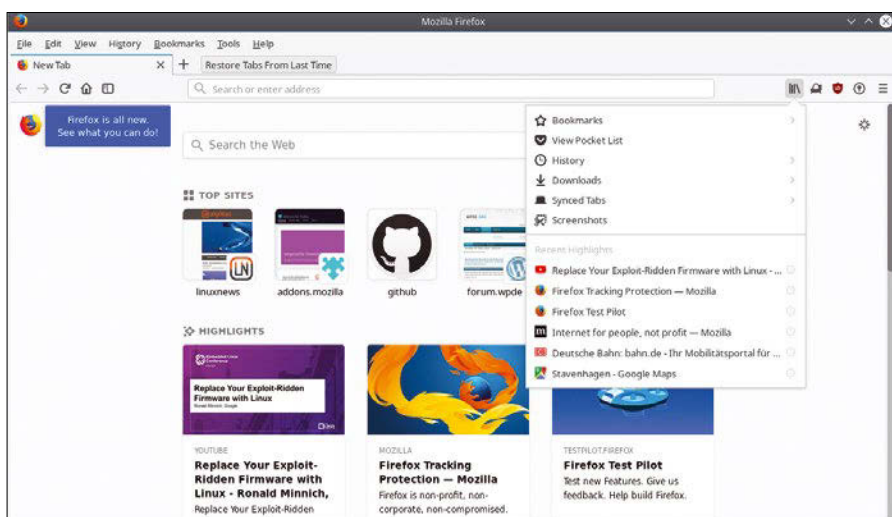
**Figure 1:** The Tab Mix Plus add-on is gradually being rewritten as a WebExtension.



**Figure 2:** Via Activity Stream, Firefox now integrates Pocket reading suggestions into the new tab page.



**Figure 3:** With a tour, Mozilla introduces Firefox Quantum's new features.



**Figure 4:** The Library menu combines frequently used functions such as history, downloads, and bookmarks.

### Further Innovations

Activity Stream introduces an innovation that users only could access previously through the experimental Test Pilot [7]. In early 2017, Mozilla acquired Read It Later, the developers of Pocket. Pocket is now integrated deeper into Firefox via Activity Stream. If you open a new tab, you will see both searches and frequently visited websites, as well as recommended reading selected by Pocket (Figure 2).

In a new tab with the Firefox logo, the browser offers a tour that introduces the interface's most important new features (Figure 3). Also included is a new *Library* menu to the right of the address bar in the form of a book stack icon. The menu combines bookmarks, history, synchronized tabs, downloads, and the Screenshots tool (Figure 4).

An available feature in Chrome and other browsers for a long time, Mozilla now merges the address and search fields in FF57. The main menu is now easier to navigate, with small icons in a space-saving list (Figure 5).

### Useful Improvements

In addition to the further improved and more prominently placed Screenshots tool (Figure 6), which you can now access via the three dots in the right of the address field, Firefox Quantum offers two more useful functions. The first is the *Send Tabs* feature, which lets you to send open tabs to other devices. In the context menu of a tab, the new *Send Tab to Device* menu item supports cross-platform opening of tabs (Figure 7). In the Android version, the new feature is found in the menu below the *Share* icon and then at the top of the other options.

FF57 also has a function for automatically filling out address forms. The current version will only work in the USA, but the address feature will gradually roll out in other countries.

At the last minute, the Firefox developers also released a tracking protection feature that was originally intended for Firefox 58. You can choose whether the protection should only apply to private windows or to all windows. You can also choose between two tracking lists. However, keep in mind that strict tracking protections can sometimes impair the presentation of websites.





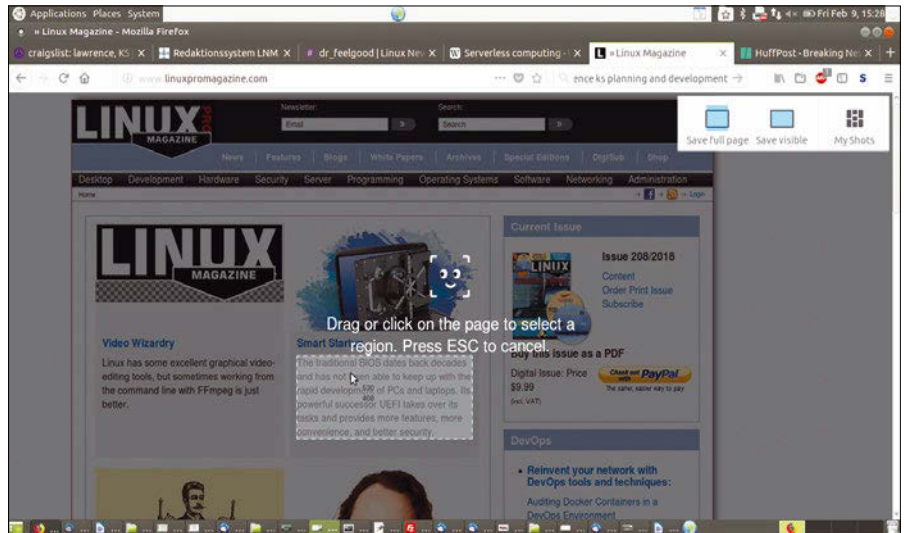
### Summary and Outlook

The experts differ on whether Mozilla will be able to stop the downward trend in marketshare. Mozilla's former technical director, Andreas Gal, who left the company two years ago, sees Firefox as doomed [8]. He also predicts a browser death race that only Chrome and Safari can survive. However, Gal and other critics forget that the Mozilla Foundation does not need to make a profit. Mozilla's head of marketing, Jascha Kaykas-Wolff, says that Firefox must be more than just a niche product, but not necessarily as widespread as Google, Facebook, or Amazon. Mozilla is primarily concerned with preserving Internet freedom and driving innovation, not maximizing downloads.

Mozilla has done much for the open Internet and is already planning for interactions with the network of networks in the future. Ongoing developments address artificial intelligence, augmented and virtual reality, and the open source Common Voice [9] project, which aims to improve speech recognition. A voice interface would completely change the browser of the future.

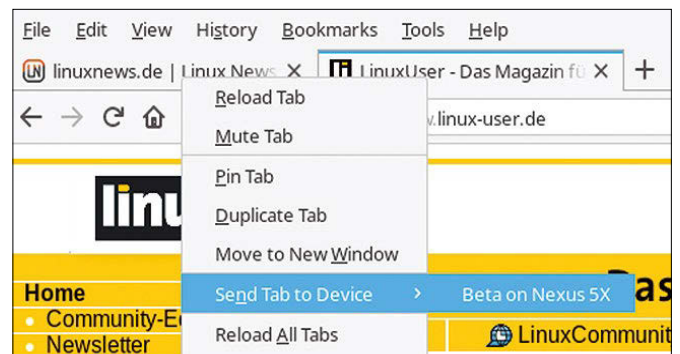
### Conclusions

FF57 puts Mozilla on a comeback path. All told, the team turned 75 percent of the code inside out – equivalent to seven million lines of code – and the software shows this at every turn. In terms of technology, Mozilla advertises that the new browser is twice as fast as Firefox 52 but requires 30 percent less main memory. However, the Mozilla browser has only just made up the gap with its competitors from Mountain View and

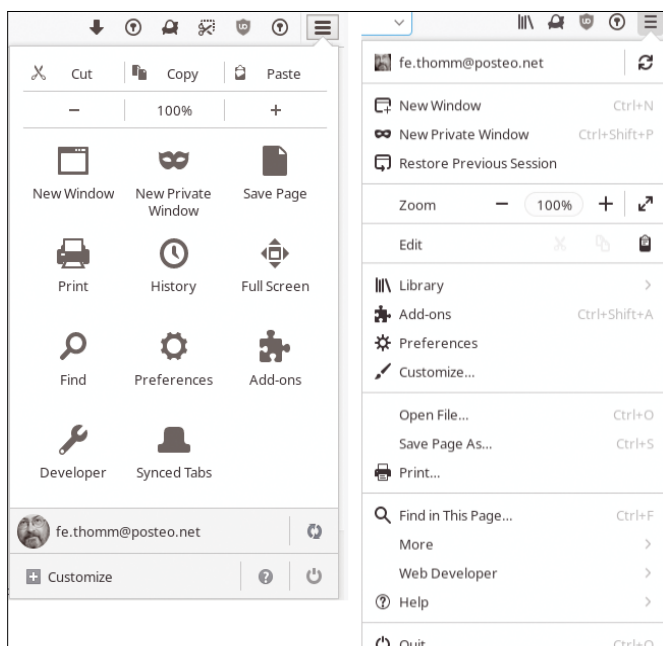


**Figure 6:** Firefox 57 makes it possible to grab screenshots directly in the browser, including saving them to the cloud. However, Opera is already far ahead of the game here.

Redmond. FF57 levels the playing field and joins the ranks of modern browser architectures. However, it remains to be seen whether the changes will be sufficient to win back old users and attract new ones. ■■■




**Figure 7:** A new feature lets you send tabs to other devices linked to your Firefox account.



**Figure 5:** The main menu is easier to navigate: Instead of large icons, there are now small icons and text.

### Info

- [1] Rust: <https://github.com/rust-lang/rust/>
- [2] Australis: <https://blog.mozilla.org/ux/2013/11/australis-is-landing-in-firefox-nightly/>
- [3] Quantum: <https://medium.com/mozilla-tech/a-quantum-leap-for-the-web-a3b7174b3c12>
- [4] E10S: <https://asadotzler.com/2016/06/06/firefox-48-beta-release-and-e10s/>
- [5] WebExtensions: <https://developer.mozilla.org/en-US/Add-ons/WebExtensions>
- [6] Photon UI: <https://www.askvg.com/review-new-photon-user-interface-ui-of-mozilla-firefox/>
- [7] Test Pilot: <https://testpilot.firefox.com/>
- [8] "Chrome Won": <https://andreasgal.com/2017/05/25/chrome-won/>
- [9] Common Voice: <https://blog.mozilla.org/internetcitizen/2017/06/19/commonvoice/>



---

# 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](http://shop.linuxnewmedia.com)





## Font management with current office suites

# WRITTEN EXAMINATION

Users of contemporary office applications often want to do more than just compose letters or design simple flyers. But the well-known office suites on Linux sometimes make font management a test of patience, with the errors literally hidden in the details. *By Erik Bärwaldt*

**T**hanks to their WYSIWYG (“What You See Is What You Get”) representation and good printer support under Linux, modern office programs have long been suitable for typologically and graphically more demanding tasks.

However, depending on the character set and preferred fonts used, the home print studio can also be a pain in the neck: For example, if you need special characters or want to integrate new fonts into the office suite, you need to delve deeper into font management under Linux. In this article, I explain where problems lie in some office packages’ font management.

### Basic Information

Developers have adapted font management to new needs in the course of computer history; things have become more professional in terms of fonts, which were originally implemented in the hardware, especially in the last three decades.

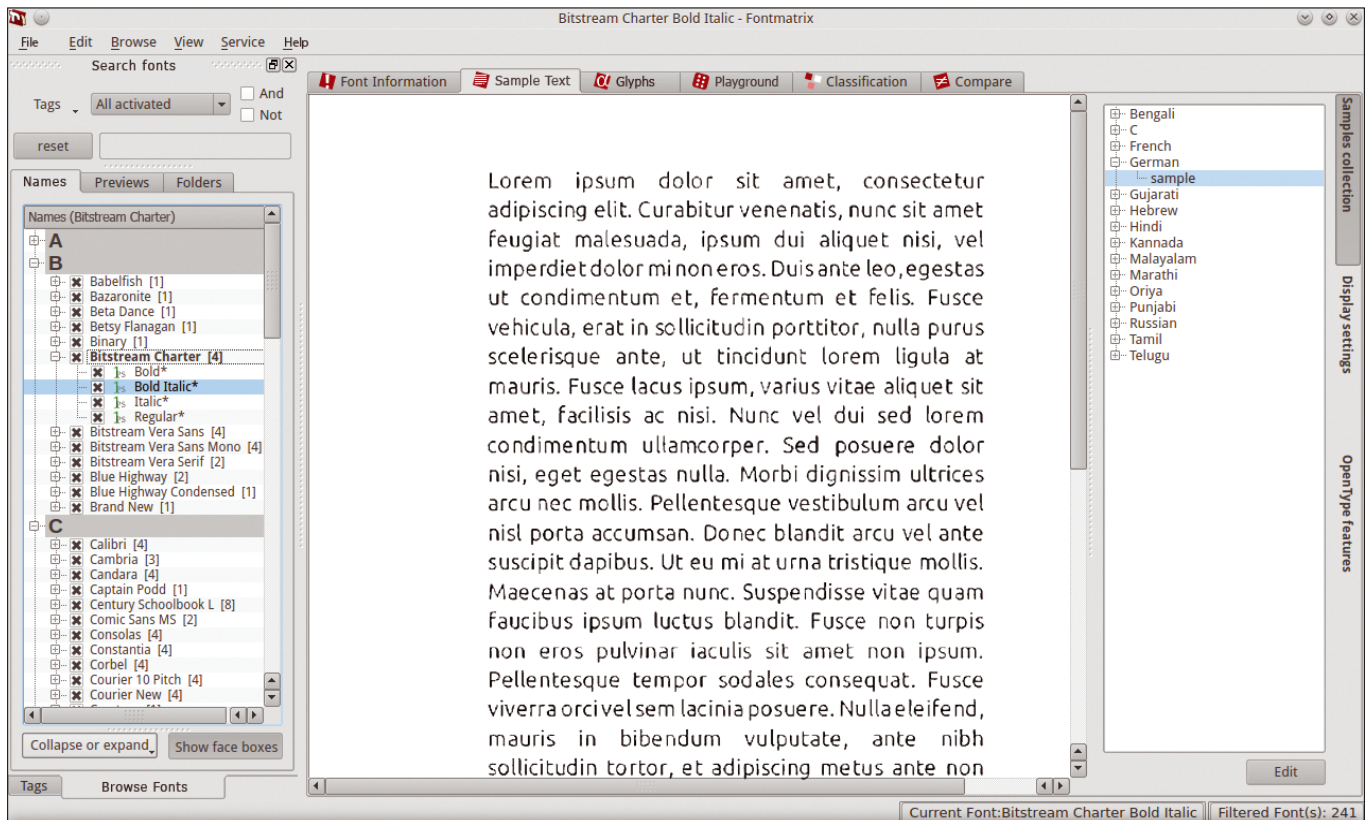
On Linux, as with all current operating systems, preinstalled system fonts are placed on your mass storage device when you install the operating system. Since 2004 in Linux, this has been the task of the *fontconfig* library [1], which is an integral part of every Linux derivative. These system fonts not only take care of screen content, but they are also

used for printer control. Linux is particularly flexible: It not only supports bitmap fonts but also vector fonts, which scale without any loss of quality. They have therefore established themselves on the market.

Adobe’s PostScript [2] page description language, which was developed in the 1980s and supplemented by the vector fonts mentioned above, is still regarded as the standard for high-quality document design. PostScript is also used by professional print studios.

Under Linux, the office suite selected by the distribution usually sends its print jobs to the print server in the form of PostScript files. This formats the data for

Lead Image © PhotoCreate, Fotolia.com



**Figure 1:** Fontmatrix provides very effective functions for font management.

the active printer using the Ghostscript interpreter [3]. Linux systems not only print documents in the usual office formats, but also PDF files, which are closely related to PostScript.

The globally installed font families are usually located in the `/usr/share/fonts/` subdirectory under Linux and are usually managed in a graphical font management system such as Fontmatrix (Figure 1) [4].

## Special Case: Asian Script Systems

Asian typeface systems pose a particular challenge to the standard office suites. These include in particular the Chinese, Japanese, and Korean writing systems, some of which even differ within their native countries.

If you want to use these font systems and ensure that the text displays correctly, you first need to install the required character sets manually, and you have to set up the word processors in such a way that they can display text in Asian fonts onscreen and output them correctly on the printer. Additionally, you need to activate input methods such as the Smart Common Input Method (SCIM) [5] or Intelligent Input Bus

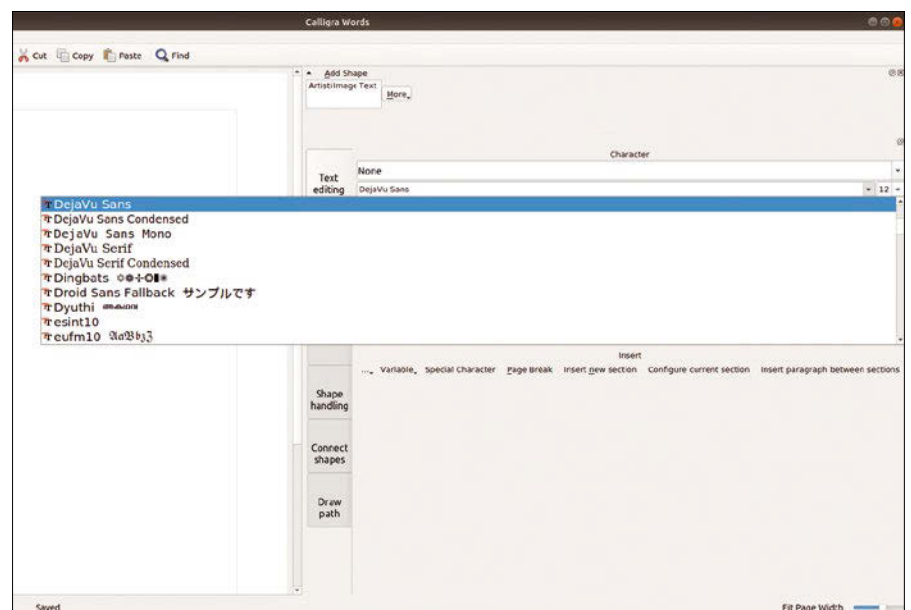
(IBus) [6] in the operating system to ensure trouble-free text input from the keyboard.

## Calligra

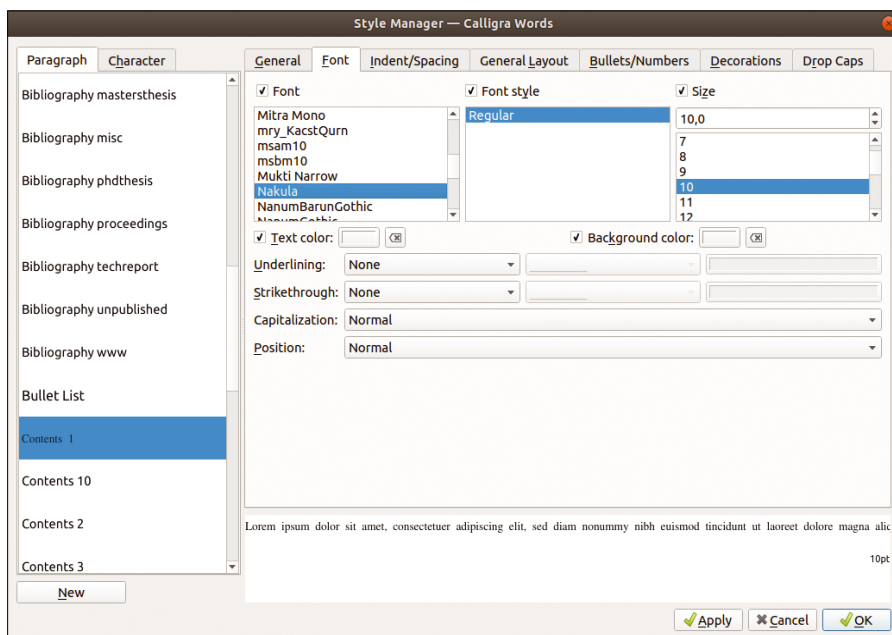
Originally developed for the KDE desktop and comprising several individual applications, Calligra Suite [7] is now regarded as a useful office package for

everyday use. Like font management in most office suites, Calligra relies on the fonts installed in the operating system, which you can conveniently choose in a selection field – including the font size.

Remarkably, Calligra can fill the horizontal usable area of widescreen monitors better than most other office suites,



**Figure 2:** The numerous imported special effects fonts are seldom or never used in everyday life.



**Figure 3:** With style sheets, you can set up fixed fonts in Calligra.

which usually place their buttonbars horizontally at the top of the window. This placement reduces the vertical area available for displaying documents. Calligra moves the buttonbars to a vertical window segment at the right edge of the screen.

Calligra adopts the system fonts without any further selection options, which has a noticeable disadvantage: Because many Linux distributions pre-install numerous special effects fonts of questionable quality, Calligra's selection list contains fonts that are rarely used and are therefore superfluous, making font selection rather confusing (Figure 2).

### Stylistic

Similar to format management in other word processing systems, style management brings some order to the typeface jungle in Calligra. For this purpose, open the *Styles | Style Manager* dialog and optionally define a *Font*, *Font style*, and *Size* in the *Font* tab. In a small area at the bottom right, you are given a text preview while selecting a font and the various attributes (Figure 3).

On the downside, the Calligra suite's support for non-standard writing systems has been only rudimentary up to now. Although it lets you write text from right to left, as is customary in Arabic, and insert Asian fonts, it does not render them completely correctly. Errors also often occur with copy-and-

paste inserts, because Calligra shows single characters only as empty rectangles, which makes it impossible to insert or recreate a text completely in an Asian font.

### LibreOffice and OpenOffice

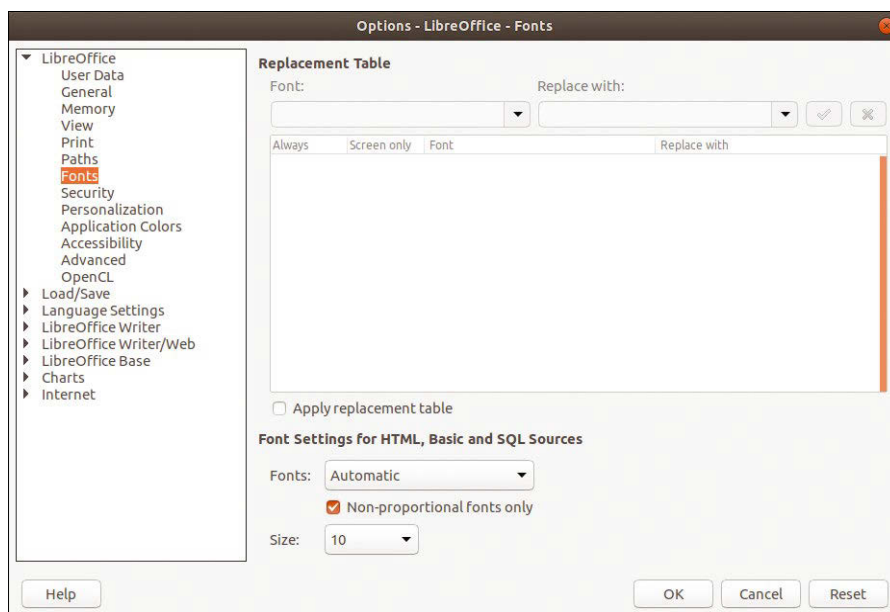
The most widely used office package under Linux, LibreOffice [8], offers very extensive font management in the buttonbar. As with all large office applications, you choose the desired font and font size using the appropriate selection boxes and symbols and set font styles and attributes.

LibreOffice and OpenOffice [9] divide fonts internally into screen and printer fonts, which causes some pitfalls: The X server fonts are for screen display only, whereas the PostScript fonts are intended for both screen display and printing. Print fonts are handled by the FreeType library [10] in LibreOffice and OpenOffice; FreeType is available for both display forms. X server fonts therefore do not appear in the font selection box in LibreOffice.

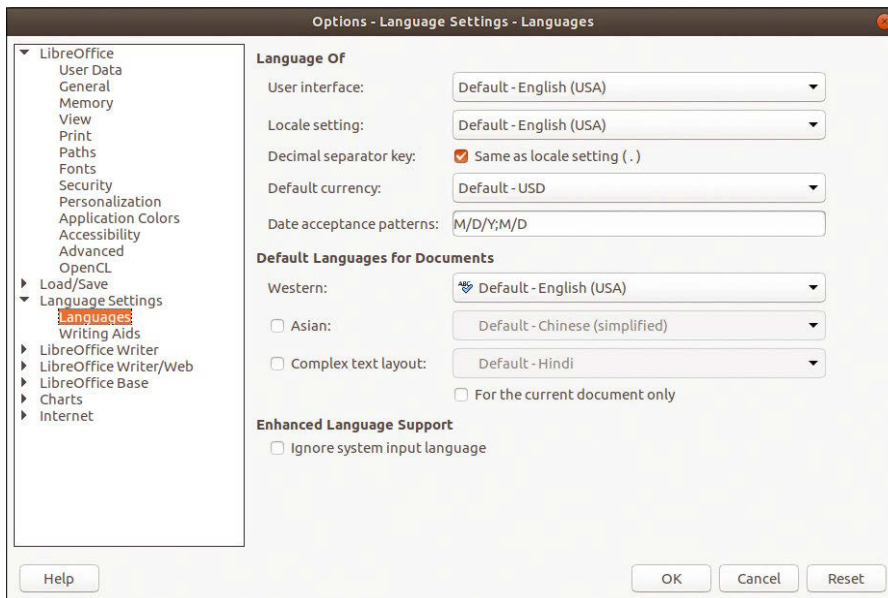
If you want to use an X server font under LibreOffice or OpenOffice, you need to install it in PostScript or TrueType format. For this purpose, use the font management system integrated into all major desktop environments or a font manager such as Fontmatrix.

Alternatively, you can select a similar font in the replacement table to replace the missing character set. Both LibreOffice and OpenOffice offer a clear-cut dialog, which you will find in the *Tools | Options* menu. In the *LibreOffice* or *OpenOffice.org* area, select the *Fonts* option. In the list view that appears, enable the replacement table by checking the box to the left of the *Apply replacement table* entry. In the *Font* and *Replace with* selection fields, choose the desired fonts to transfer them to the table by clicking on the button with the checkmark to the right.

Before enabling settings, it is usually worth taking a look at the replacement table documentation that comes with LibreOffice [11]. The office suites allow



**Figure 4:** A replacement table lets you define replacements for fonts that do not exist on the system.



**Figure 5:** With just a few clicks, you can get OpenOffice and LibreOffice ready for Asian fonts.

various settings combinations that influence whether the replacement for the originally selected font appears only on the screen or in the print output, depending on the installation (Figure 4). Clicking on the *OK* button at the bottom of the dialog finally applies the configured replacement table.

Asian typeface systems have not posed insurmountable problems for LibreOffice and OpenOffice for some versions. However, you do need to configure them. In the first step, change the input method in the configuration dialog. In the *Tools | Options | Language Settings | Languages* menu, select the *Asian* (to enable Asian languages) and *Complex text layout (CTL)* entries.

These ensure that the office suites display Asian fonts correctly and mirror the layout (Figure 5).

You then need to copy the appropriate character sets into the designated directory to be able to select them in the office suite. To enter text with the corresponding font, modify the default language in the *Tools | Language* submenu. This can either be set for the active selection, the paragraph, or the entire text.

## OnlyOffice

The OnlyOffice Desktop Editors office suite [12], which is strongly focused on collaborative work, also relies on the operating system's font management for

a document's output design. Unlike other candidates, it does not offer a WYSIWYG preview for font selection; rather, it just lists the fonts available on the system. Because the software is strongly oriented toward cloud-based services, you will not find any options for font substitution.

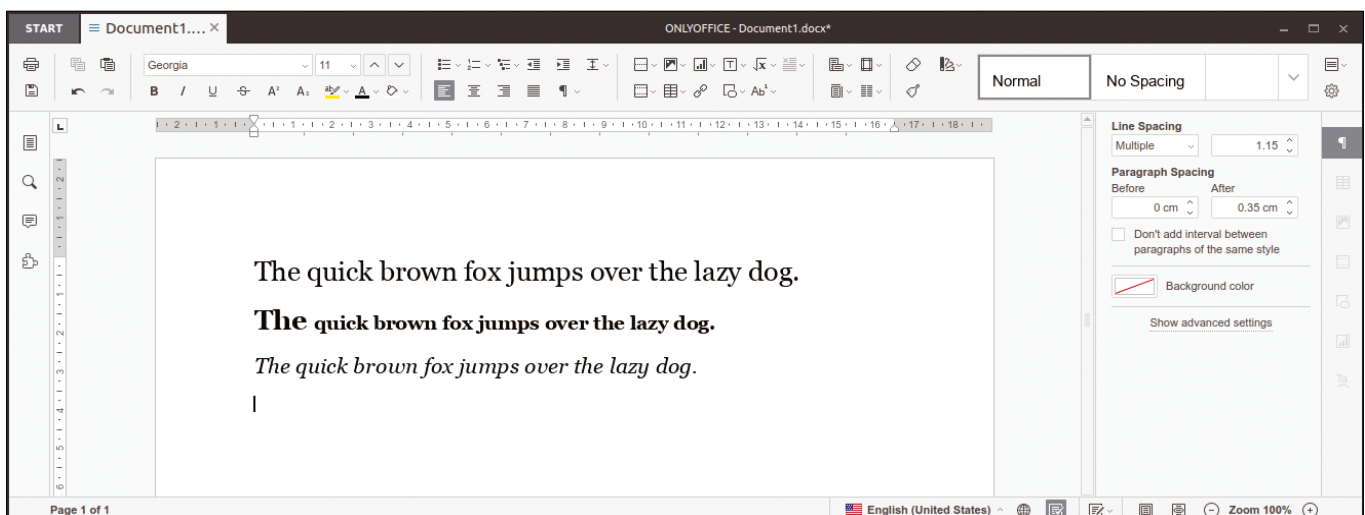
As with the other office solutions, selection boxes for font and font size appear in the respective program window. You can also change attributes using buttons. However, further options are not available.

## Automatic

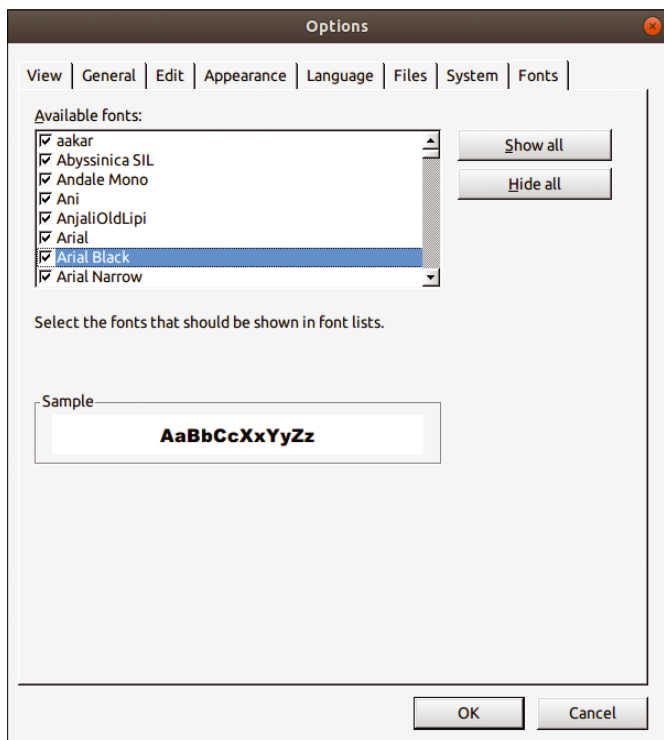
If you receive a document with a font that is missing on your local system, OnlyOffice attempts to replace the font automatically, including any attributes and sizes, with a font installed in the existing operating system. However, it incorrectly displays the original font in the selection field of the menubar, without making it available. If you then add text, it appears in the replacement font. In the case of longer text with many additions between individual paragraphs, the font changes permanently in the selection field.

Depending on the substitute font automatically selected by the program, the entire typeface can suffer: Different widths of individual fonts can cause completely different line breaks, which sabotages the entire layout of a document, and different font heights cause a very uneven typeface.

Additionally, our lab revealed that after exporting such a modified docu-



**Figure 6:** With OnlyOffice, font management leaves a great deal to be desired, especially when working collaboratively on different platforms.



**Figure 7: SoftMaker Office is the only application in the test that brings order to font confusion.**

ment, unexpected results occurred when opening it in other word processors. For the additions, OnlyOffice initially used one of the system fonts, which looked rather unattractive on the screen because of different character widths and font heights. After reimporting to OpenOffice, the document appeared in its original font, and the substitute fonts from OnlyOffice had completely disappeared, which again affected the layout.

For collaborative, cross-platform work with OnlyOffice, it is therefore a good idea to ensure that the computers are equipped with identical system fonts. Only then do the documents change hands without time-consuming postprocessing on the individual systems (Figure 6).

Even if you install fonts at a later date, the software sometimes incorrectly displays imported text with empty rectangles instead of the desired characters. OnlyOffice is unable to cope with Asian type systems.

### SoftMaker Office

SoftMaker Office [13] has a long and continuous development history and is available across platforms. The software also uses the fonts installed in the operating system but has its own font manager on board, as well, which you can

limited extent: Nevertheless, a good half-dozen randomly selected standard fonts remained in the selection field. If you only want to use your own fonts, you can't do so.

SoftMaker Office offers useful support for Asian type systems. You can first install appropriate character sets in the font directory then access them after specifying an individual search path in the *Tools | Options | System* menu. Note that you have to enter the search path as a complete path and set the corresponding language under *Format | Character | Font* in the *Language* selection field. Once you have successfully adapted the input method for your

access via *Tools | Options*. In this settings window, select the *Fonts* tab to activate or deactivate existing fonts by ticking the checkbox (Figure 7).

However, this function displayed weaknesses in the practical test. Theoretically, for example, you remove all fonts in the font dialog of the Options window by clicking *Hide all* or by unchecking individual boxes to the left of the respective fonts. That only worked to a

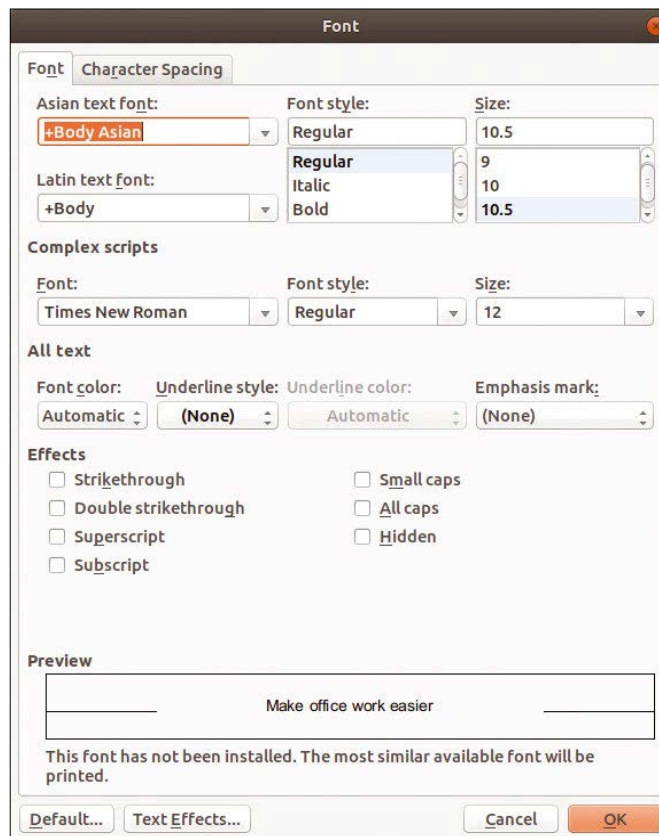
desktop environment, you can then input Asian text.

TextMaker, SoftMaker Office's word processor, copies and pastes text blocks from nonstandard writing systems into a document without complaint: Only WPS Office and TextMaker were able to master font substitution without incorrect characters. TextMaker uses the preset SimSun font for Chinese typesetting systems.

### WPS Office

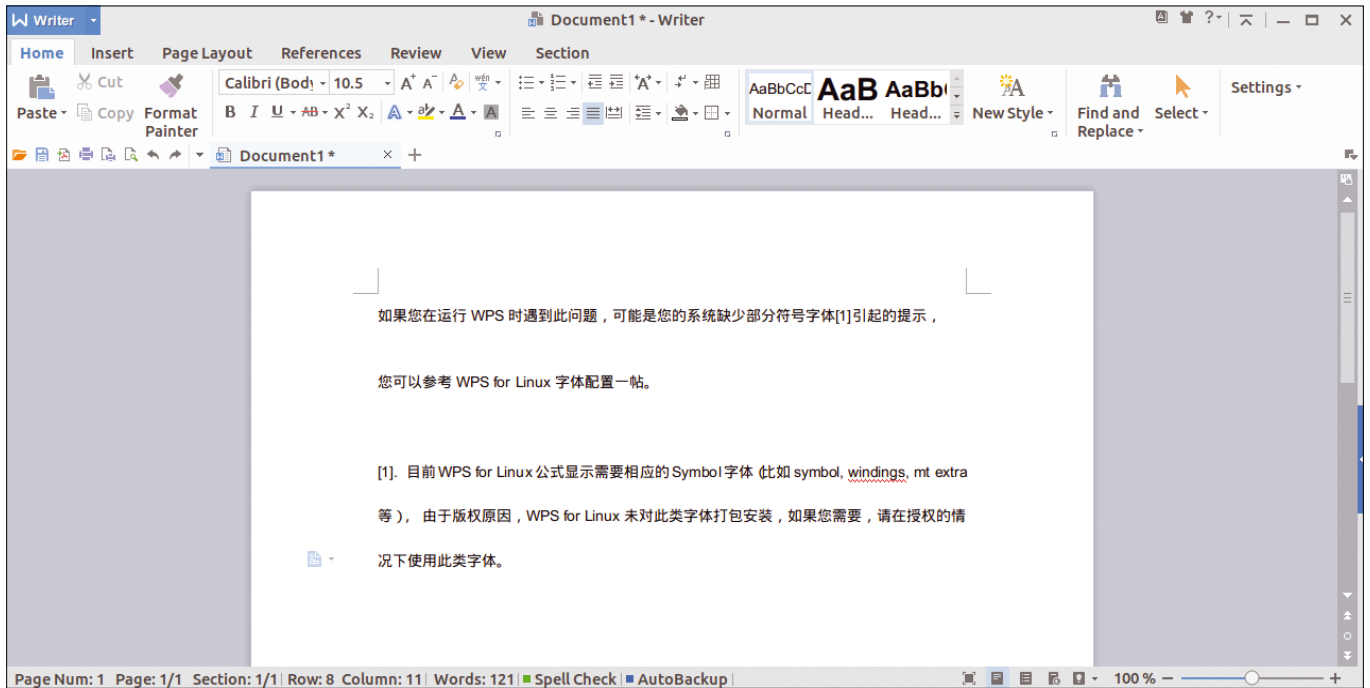
WPS Office [14], developed by Kingsoft in China for more than 20 years, has its own font management system, in contrast to many other office packages. However, at first glance, the individual applications' interfaces hardly differ from those of the other office packages: The selection list for fonts and the usual buttons for font attributes are available.

The selection field for the existing fonts uses the fonts installed in the operating system and displays them accordingly in the preview. However, in the *Format | Font* menu, you will also find a very detailed settings dialog in which



**Figure 8: Highs and lows: Although font management (top) is exemplary, WPS Office slips up with the preview (bottom).**





**Figure 9:** No surprise: The Chinese-based WPS Office can easily manage Chinese typefaces.

you can select a standard font from the list for a document and add various attributes and font styles.

Because of its origin, WPS Office distinguishes between Asian and Western typesetting systems and allows the configuration of different character sets for the respective systems. The lower part of the configuration dialog has a small preview of the chosen font.

However, WPS Office is a bit of a nuisance: In the test for *Linux Magazine*, some fonts showed up in the preview with incorrect spacing. If fonts are not installed, WPS Office also displays a message in this dialog box and uses a similar replacement font (Figure 8).

WPS generates replacement fonts with the use of an internal table. If you copy text from Asian font systems into WPS Writer, the word processor automatically switches to the SimSun font. You do not need to install any fonts manually, and WPS displays the text

correctly on the screen and during printing (Figure 9).

## Conclusions

Under Linux, almost all office suites have a sophisticated font management system that can handle even the most demanding creative tasks (see Table 1). Weaknesses are only revealed in the details: For example, not all office suites can cope with non-Roman font systems. The installation of additional fonts can be problematic if you do not want to make them available globally on the system.

Deactivating unnecessary fonts that automatically end up on the mass storage device during the operating system's installation does not affect any of the candidates except SoftMaker Office. Therefore, it is recommended that you also use a program for professional font management, such as Fontmatrix. This is the only way to exploit the possibilities of creative design fully. ■■■

## Info

- [1] fontconfig: <https://www.freedesktop.org/wiki/Software/fontconfig/>
- [2] PostScript: <https://en.wikipedia.org/wiki/PostScript>
- [3] Ghostscript: <https://www.ghostscript.com>
- [4] Fontmatrix: <https://github.com/fontmatrix/fontmatrix>
- [5] SCIM: [https://en.wikipedia.org/wiki/Smart\\_Common\\_Input\\_Method](https://en.wikipedia.org/wiki/Smart_Common_Input_Method)
- [6] IBus: <https://github.com/ibus/ibus/wiki>
- [7] Calligra Office: <https://www.calligra.org>
- [8] LibreOffice: <https://www.libreoffice.org>
- [9] OpenOffice: <https://www.openoffice.org/?redirect=soft/>
- [10] FreeType: <https://freetype.org>
- [11] Replacement table documentation: <https://help.libreoffice.org/5.4/en-US/text/shared/optionen/01010700.html>
- [12] OnlyOffice Desktop Editors: <https://www.onlyoffice.com/apps.aspx>
- [13] SoftMaker Office: <http://www.softmaker.com/en/>
- [14] WPS Office for Linux: <https://www.wps.com/linux>

**Table 1:** Font Management for Office Suites

	Calligra Office	LibreOffice/OpenOffice	OnlyOffice	SoftMaker Office	WPS Office
Font selection possible	Yes	Yes	Yes	Yes	Yes
Preview function	Yes	Yes	Limited	Yes	Yes
Integrated font management	No	No	No	Yes	No
Activate/deactivate fonts	No	No	No	Yes	No
External writing systems	No	Yes	No	Yes	Yes
Replacement tables	No	Yes	No	Yes	Yes
Manual selection of font directory	No	No	No	Yes	No

NoMachine remote desktop solution on Linux

# Radical

Need remote access to a PC? NoMachine offers this and much more. *By Erik Bärwaldt*

**T**oday, two or more PCs are often used in private households, in addition to tablets and smartphones. Often, one of the desktop computers is used as a server that the other systems access – for example, to transfer data, but above all to stream multimedia content. Remote desktop programs let you work on a remote computer without the systems exchanging large amounts of data.

Software for remote access to computers on the intranet also simplifies the maintenance of the other computers from your own workstation. In enterprises, remote desktop software performs the same task, but also allows access to servers via thin clients, for example. If such lightweight PCs are used instead of conventional desktop computers, this saves on costs.

Linux offers many options to take advantage of remote desktop applications. In pure Linux environments, programs based on Virtual Network Computing (VNC) are often used, such as TightVNC [1] or the KDE duo Krdc/Krfb. However, you are much more independent with NoMachine [2], cross-platform remote desktop software that offers some additional highlights.

VNC [3] and its derivatives are based on the Remote Framebuffer Protocol,

which in older versions causes a fairly substantial flow of data along with some technical disadvantages, such as rudimentary scaling of the screen resolution in real time. In contrast, NoMachine uses the NX protocol, which minimizes data overhead and is therefore suitable for use over low-bandwidth networks.

The NX protocol compresses data and caches previously transferred files, so the volume drops considerably during the transfer. No additional software is required on other platforms, so the software always works with similar transfer speeds in heterogeneous environments. Additionally, NX uses an SSH tunnel to transmit the data in the professional versions.

## Installation

NoMachine is available free of charge in several versions as DEB and RPM packages on NoMachine's website [4]. A tarball archive also exists for each hardware architecture, which you can install independent of the existing package manager and distribution.

After downloading the 45MB package, install it using the appropriate package management tools. In most cases, the routine integrates the application into the existing menu structure.

The setup also ensures that the program is automatically activated at system startup, so you will find the red NoMachine icon in your desktop system tray when you restart your computer. However, for some distributions with the MATE desktop, automatic integration into the menus does not work, so you have to add the corresponding starters manually.

The freely accessible software version is available as a client-server application under a proprietary license, and the essential components of the NX protocol and the associated libraries are licensed under the GPL. For professional users, NoMachine offers a commercial licensing model that includes various support options and additional features: Although the free version does not support the SSH protocol and only one connection to a remote system is established at a time, these restrictions are removed in the paid Enterprise versions.

## First Use

After installing on all computers involved, you can launch the software by clicking on the corresponding launcher. In the first dialog, the wizard shows some options that allow you to connect to remote computers.



**Figure 1:** NoMachine remote desktop software relies on a mix of proprietary and free components and a tiered licensing model.

After clicking *Continue*, you reach an area where the software browses the local network for computers on which the program is also running. It then displays these computers along with their IP addresses or hostnames. Make sure you launch the software on the participating systems before the first connection is established (Figure 1).

First, change the application's user interface to your desired language. To do this, on the administration machine used to display the remote desktops, click the *Settings* button and choose the *Appearance* tab of the new dialog. Next, select the desired language option in the *Text* drop-down box.

Further optional settings can be found in the *Transfers*, *Folders*, and *Privacy* tabs. In *Transfers*, decide whether it should be possible to exchange files between the participating machines and, if so, up to what maximum size. In the *Folders* tab, you can define where the program stores the transferred data, among other things. In the *Privacy* tab, you can set different options for deleting logfiles. By default, the program automatically removes all logs at the end of a session.

You can set some basic settings for the appropriate service by clicking on the *Connection preferences* button at bottom left: Here you can configure, among other things, whether it loads automatically every time the PC is started and whether the system is visible on the in-

tranet. You can also change the port if necessary; by default, the tool listens on port 4000. If you change this value on the server, you will also need to enter the new port number on all systems that you control remotely from here; otherwise, no contact will be established.

When the configuration is complete, click *Done* and relaunch the computer system to apply the modified options.

## Remote Control

After relaunching, select the desired remote system from the displayed list in the corresponding dialog and establish the connection by clicking on the *Con-*

*nect* button. The program will now guide you through a multilevel dialog box with instructions on how to use the software, in the course of which you can authenticate against the remote system.

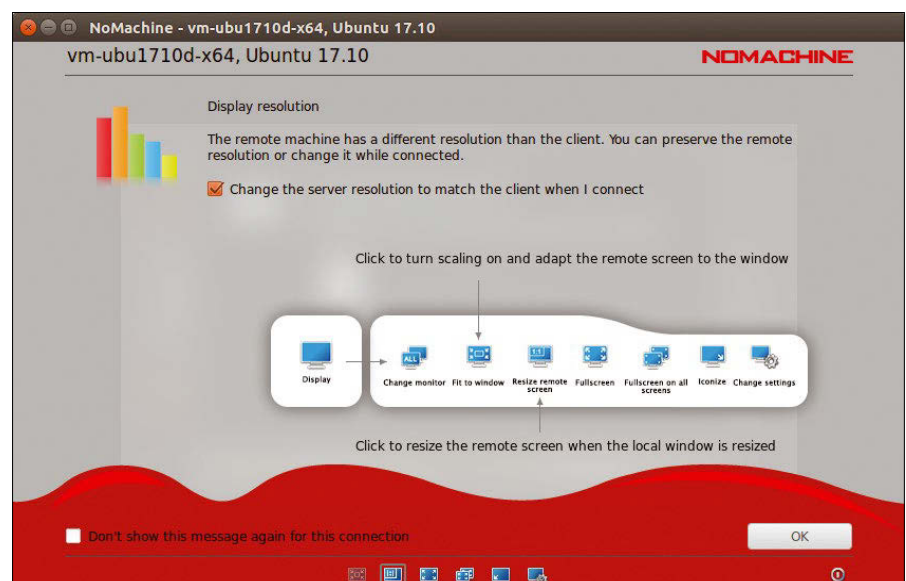
Unlike similar programs such as Any-Desk and TeamViewer, NoMachine does not use separate accounts for this purpose, but uses a user account and the associated password on the remote computer. Therefore, when you log on to the remote system with root credentials, you have administrative privileges.

When using the program for the first time, it is strongly recommended that you read the notes on the meaning of the symbols that appear in the individual dialogs, because they provide important information for controlling the screen and scaling (Figure 2).

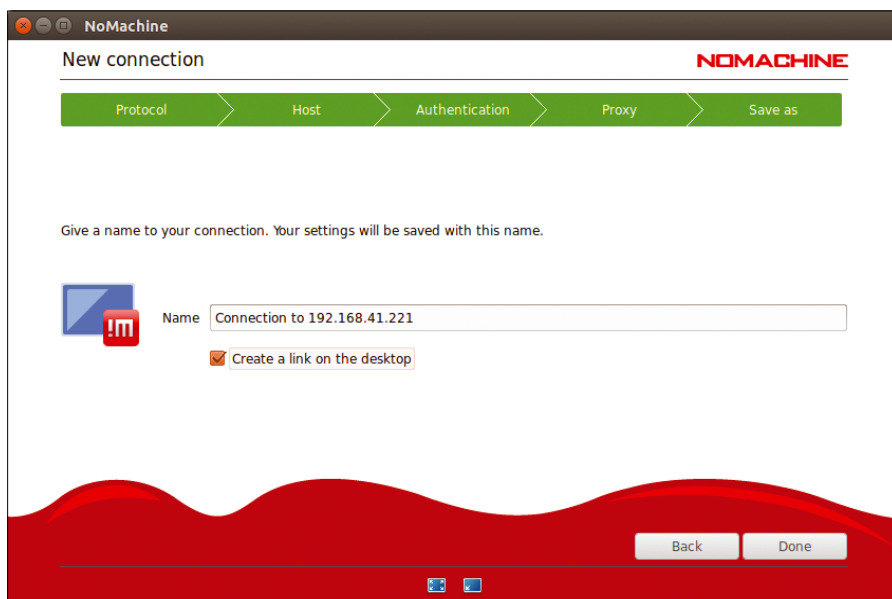
## Manual

Additionally, the software lets you manually connect to other computers. To do this, press the *New* button and open the appropriate wizard in the connection display. In the first step, define the protocol; only the NX protocol is enabled in the free variant. Next, enter the IP address of the remote host and the port you are using in the second dialog.

In the third step, define the authentication method. There are several options to choose from, although in the private sector and small companies with few workstations, it is usually fine to log in with the system login. In larger environments, authentication using a private



**Figure 2:** The program window provides you with information on how to operate the software.



**Figure 3:** You can manually set up new connections in just a few steps.

key, chip card, or Kerberos ticket is recommended.

Now, tell the software whether you want to establish the connection directly or use a proxy server. Finally, assign a name for the configuration, and exit the wizard by clicking *Done*, thus adding the new connection to the list (Figure 3).

If you are connected to a remote NoMachine server and you see its screen contents on your own screen, it makes sense to tweak some more details. If you mouse over the upper right corner of the duplicated screen, a tab opens; then, left-click to see the dialog with the settings. This is where you can adjust the remote system's screen resolution and window size to suit your local client and also modify other parameters using the various groups displayed.

There are options for enabling the microphone on the remote NX server, changing the volume, or recording the entire session. The option to integrate drives and printers in the *Devices* special menu is particularly relevant for terminal operation. NoMachine supports bidirectional use. This means, for example, that USB drives connected locally to the client are available to the user on the remote machine. Additionally, the software detects USB or other drives attached to the remote computer and lists them in the settings (Figure 4).

By clicking on the respective drive, you can integrate it into the client system. The program creates a folder in

which you can access the corresponding content as well as local data. A green dot next to the drive name in the settings indicates successful integration. You can close the settings window and release the remote system screen by clicking *Done*.

### Signs of Disintegration

NoMachine does not automatically scale the remote system's contents on the local client. As a result, errors occur relatively frequently, especially with the different screen resolutions of the systems involved. To make changes to the resolution at any time, the software lets you open the dialog with the settings by clicking *Ctrl + Alt + 0*. The most common

options – full-screen operation of the remote system on the local screen and use in a local window – are available at the push of a button.

After opening the dialog, you will find the options to customize the screen at the bottom of the window (Figure 5). The screen symbol on the right-hand side of the screen lets you set the general options. If you want to adjust them during a session, you can do this on the fly with the other icons. You may have to change an input device's settings to avoid seeing double when you look at the mouse cursor.

While the on-the-fly settings take effect immediately, options modified for all sessions are generally only enabled after restarting the software.

### Status Displays

The software provides information about a connection during a session. To do this, call the dashboard during operation with *Ctrl + Alt + 0* and select the *Connection* option. The most important data now appears in a new dialog box (Figure 6).

Alternatively, you can view status information with the use of an applet in the desktop's system tray. During the program installation, the routine also creates a *Show the Server Status* launcher, which also reveals important connection details when the application is active. In addition to the actual server status, you can see the connected users and transmission processes in the resulting dialog.



**Figure 4:** NoMachine allows you to use different peripheral devices.

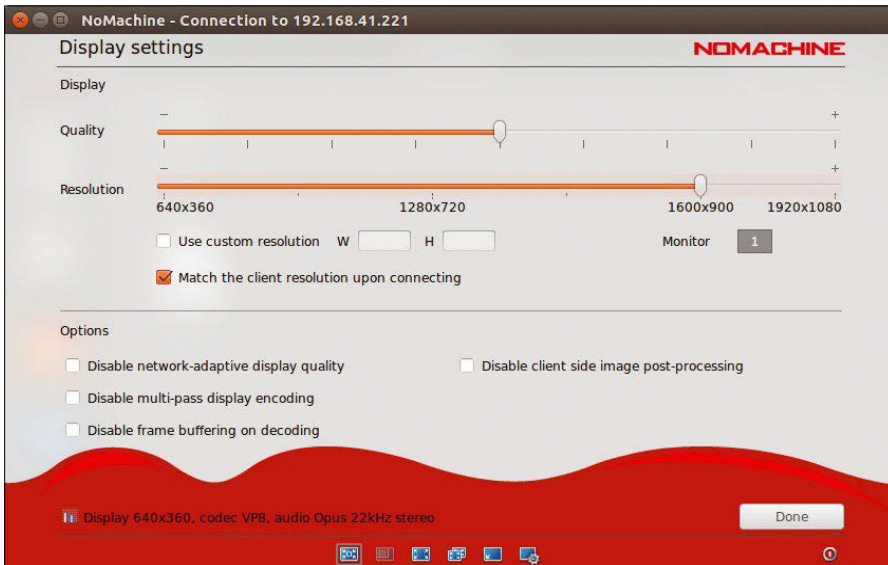


Figure 5: In a separate dialog box, you can adjust the screen resolution.

## Conclusions

NoMachine was not entirely successful in the field test: The free version is stable and offers a well-thought-out operating solution, but the transmission rates left much to be desired, even in a wireless infrastructure based on current standards. Competitors such as AnyDesk and TeamViewer build up content much faster – no matter which graphics cards and which native resolution you use.

The software plays to its strengths in remote maintenance scenarios rather than environments where thin clients help to save costs in larger IT infrastructures. In this case, the program runs in full-screen mode on the client, giving the user the impression that they are sitting

in front of the remote machine. Bidirectional data transmission facilitates work in such environments.

The program's lean design also proves that it does not necessarily have to be difficult to configure a client-server solution. NoMachine is therefore recommended for all users who want to work seamlessly on remote computers in an intranet from the local system. ■■■

## Info

- [1] TightVNC: <http://www.tightvnc.com>
- [2] NoMachine: <https://www.nomachine.com>
- [3] Info about VNC: [https://en.wikipedia.org/wiki/Virtual\\_Network\\_Computing](https://en.wikipedia.org/wiki/Virtual_Network_Computing)
- [4] Download: <https://www.nomachine.com/download/linux&id=1>



Figure 6: The software lists important status characteristics, if needed.

# GOT CLUSTER?

**HPC Up Close**

- Let the Editor Wars Begin!
- Los Alamos Double-Checks Wiring, Saves \$2 Million
- XSEDE Starts a New Era
- TKperf

**Further Reading**

- Useful NFS Options for Tuning and Management
- New Service Will Adapt HPC Code for Next-Generation Hardware
- Singularity – A Container for HPC
- Performance Management Tools
- Desired State Configuration for Linux

**Highlights**

**Let the Editor Wars Begin!**  
Editors, particularly command-line editors, are an important tool for high-performance computing systems administrators. We point out some...

**XSEDE Starts a New Era**  
NSF announces details for new XSEDE 2.0 phase

**Los Alamos Double-Checks Wiring; Saves \$2 Million**  
A review team for the Trinity supercomputer finds savings with aluminum cable and other innovations

**TKperf!**  
Linux offers several low-level tools for investigating solid state drives (SSDs). Admins are well advised to deploy these tools immediately...

**Most Read**

**SSD Tuning**  
Solid-State drives (SSDs) are quite different from old-fashioned hard disks, and you'll have to learn some new techniques if you want to tune up...

**Boosting Performance with Intel's QuickAssist Technology**  
QuickAssist technology offloads computationally intensive compression and encryption tasks to provide a performance boost for Intel processors.

**Workshop on Workflows Announced**  
Free training session for the HPC community will take place at 12 US locations

**Altair Case Study**

Watch the **Wayne State University** video on using PBS Pro to support their environment for research

**Get to Know ADMIN Magazine**



Get to know ADMIN with a risk-free ADMIN Trial Subscription.

Try ADMIN magazine for two issues delivered to your inbox over four months with Trial Subscription.

You'll be sent the next two issues of ADMIN at no charge.

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

[hpc.admin-magazine.com/Newsletter](http://hpc.admin-magazine.com/Newsletter)



The Kernel Self-Protection project aims to make Linux more secure

# KERNEL KEEPER

Security vulnerabilities in the kernel often remain undetected. The kernel hacker initiative, Kernel Self-Protection, promotes safe programming techniques to keep attackers off the network, and, if they do slip through the net, mitigate the consequences. *By Tobias Eggendorfer*

**A**ny Black Hat who finds a previously unknown vulnerability in the Linux kernel has hit the jackpot. Potentially millions of servers and embedded devices are suddenly open to attack, and the attacker can usually gain root privileges. Users clearly don't want this to happen, and kernel makers try to prevent such events.

Based on pure theory, strict coding standards and a sophisticated software quality management system ensure that loopholes are found immediately and eliminated before the release. A shining and rare example is OpenBSD, which in 20 years' time has only had two significant security breaches [1]. Although I am certainly a supporter of the vigilant approach followed by OpenBSD, we have to be realistic: The Linux kernel contains mountains of code that no one can review with the required depth; dependencies vary, and thus so do the possible attack vectors. (See the box entitled "Harmless Start" for a complex example.)

The complexity of the Linux kernel means that it is likely to carry legacy ballast and bugs for an indefinite period of time. At the end of 2010 [4], Jonathan Corbet checked how long the safety-relevant bugs eliminated in that year had existed until discovered: 22 of the 80 loopholes examined had been in the code for more than five years!

Practical experience leads to an approach that simultaneously makes attacks more difficult and reduces the consequences of exploitable code weaknesses. This two-pronged approach is the goal of the Kernel Self-Protection [5] project.

## Break-In Technology for Everyone

Viewed through the looking glass with sufficient hindsight, most attacks on programs work in a similar way. An attacker tries to add new program code to a running process, which the hijacked process then executes with its privileges. The added code can be SQL or shell commands, or typically, binary

code in kernel attacks. In order to inject this code, attackers exploit programming errors that allow them to determine memory contents and manipulate the program counter.

The program counter is a CPU register that points to the next instruction to be executed. For Intel processors, the 16-bit programs had the Instruction Pointer (IP), the 32-bit world had the Extended IP (EIP), and the 64-bit world has a somewhat-morbid Relative Instruction Pointer (RIP). Almost all attacks aim to change the contents of this register in such a way that it points to one of the attacker's commands instead of the next command intended by the programmer.

Direct write access to this register is virtually impossible, which makes a small detour necessary. When a program calls a function, the program copies the return address (i.e., the point at which the program will continue running after calling the function) to the stack. Since the stack also contains local function variables that the attacker may be able

Lead Image © stylephotographs, 123RF.com

## Harmless Start

Some security breaches are well hidden. CVE-2015-7547 started with an innocuous-sounding bug report on Glibc 2.20 [2]: A programming error that had probably existed since Glibc 2.9 caused a program crash. Six months later, it turned out that a skillful combination of access turned the error into an attack.

The Glibc programmers had provided a buffer of 2048 bytes on the stack for a DNS response, which they wanted to put back on heap if the response was larger. But this sometimes went wrong, because the function queried an IPv4 and an IPv6 address and tried to use the rest of the buffer in the second answer. Only if that didn't work was the bigger buffer requested on heap. However, the variable that held the pointer for this buffer was not updated, so that access to the stack continued. A stack overflow was made possible by the fact that the size check no longer fit the situation.

Taking advantage of this (now remedied) error situation involves some complications for the attacker: For example, they

would have to send DNS packets larger than 2KB. According to the vulnerability investigators, the attacker therefore either needs control over a DNS server or at least the ability to spoof DNS packets as the man-in-the-middle [3]. Alternatively, an attack could also be triggered by clever timing. To do this, the second response has to arrive after the time out and the first response needs to be larger than 2KB. For the second response, Glibc then incorrectly sets the buffer size to "large," while it only allocates a small buffer to the stack. This allows a stack overflow.

Because the attack requires an unusual combination of parameters and complex dependencies, it is difficult to detect. One measure that could help detect such an attack is a canary (described later in this article). If a check reveals the canary value has changed, a buffer overflow has occurred. Other measures include address space layout randomization. A non-executable stack would have made the attack more difficult and thus reduced the risk.

to manipulate through input, this return address is a popular destination.

## A Leap Back to Ruin

The attacker has several options for changing the address. A buffer overflow is a bit rough: The attacker simply overfills a variable with excessive amounts of data and floods the return address (Figure 1). Slightly more subtle attacks rely on format string vulnerabilities, which make it possible to both read and manipulate the stack.

Integer overflows are another possibility; for example, an attacker might exploit the fact that signed and unsigned integers have different value ranges. The

unsigned number 128 could become a -128. If the program only checks whether a certain upper limit is complied with, for example, because there is a maximum of 100 values on the stack, -128 is technically OK. Instead of occupying memory locations 0 to 100 as the programmer intended, the range now extends from -128 to 100, which can trigger an overflow situation.

## Run for Cover

An attacker can attempt to manipulate the stack or heap. Both are storage areas in which data is normally stored. Therefore, a popular avoidance strategy, which Kernel Self-Protection also advocates, is

to mark memory areas as either executable or non-executable (in other words, to declare some areas as code and others as data). Newer CPUs have a bit for this in the page descriptor. AMD refers to this as the No Execute (NX) and Intel as the Execute Disable (XD) bit. Linux kernel supports write protection in this context.


Although these access rights do not provide 100 percent protection, they at least make the attack more difficult. For example, advanced attackers could still gain control of the system with a *Return to Libc* attack, which writes the entry address of `execve()` to the stack along with suitable parameters [6] [7].

Successful attackers first only manage to influence the memory of a hijacked process. It's a good thing that modern processors contain functions that separate the kernel memory from application memory. In this case, a kernel function must not execute any commands that have been injected into the user space. This protection also makes attacks more difficult.

## Canary vs. Overflow

A stack overflow overwrites parts of the stack. Protection can be provided by a canary, a bit sequence prior to the return address that the attacker cannot predict. If the patterns no longer match, that means somebody overwrote the stack. This is a warning, in the same way that a canary warned miners of mine gas.

A skillful attacker can work around this canary protection. For example, a technique called trampolining bends existing pointers in the program so that they point to the return address. However, this approach requires considerably more effort and the use of vulnerable pointers. Similarly, an attacker might try to read out the



**Figure 1:** If an attack succeeds in storing more data in a buffer on the stack than the programmer had intended, the attacker can overwrite the valid return address in the buffer overflow.

canary value and adjust the buffer overflow appropriately, but the programmer can counteract this step by randomizing the canary.

Shadow stacks provide a similar line of defense: They keep a copy of the return address, which the attacker is unlikely to access. The process checks whether the copy matches the original before it is returned [8].

The heap, as well as the stack, is sometimes vulnerable to overflows. Attackers often take advantage of use-after-free errors: The programmer releases memory with `free()`, but then uses the pointer again later on. Such errors can be corrected by more stringent checks when accessing memory areas in the kernel.

### Invitation to Gamble

Attackers rely on well-known facts relating to the operating system's memory layout to inject their code (Figure 2). If the kernel randomizes the design of the addresses, attacking is like gambling. Although there are attacks like heap spraying, the overhead is growing. If both the kernel and the kernel modules are loaded at random addresses, and possibly in a random order, attacks on the kernel modules are more difficult. Similarly, random stack base addresses increase the effort required for a successful attack.

The more chance comes into play, the more important it is for the attacker to find out something about the memory structure or the value of canaries through information exposure. Attackers take advantage of the fact that released memory is usually not overwritten. Uninitialized variables or format string vulnerabilities also help.

To avoid this kind of information gathering, it is important to overwrite the memory immediately after releasing it. In addition, functions can be accessed via IDs and a table, analogous to the interrupt vector table, instead of directly

using their addresses. This precaution means that the addresses outside the kernel are not known and are more difficult for attackers to predict.

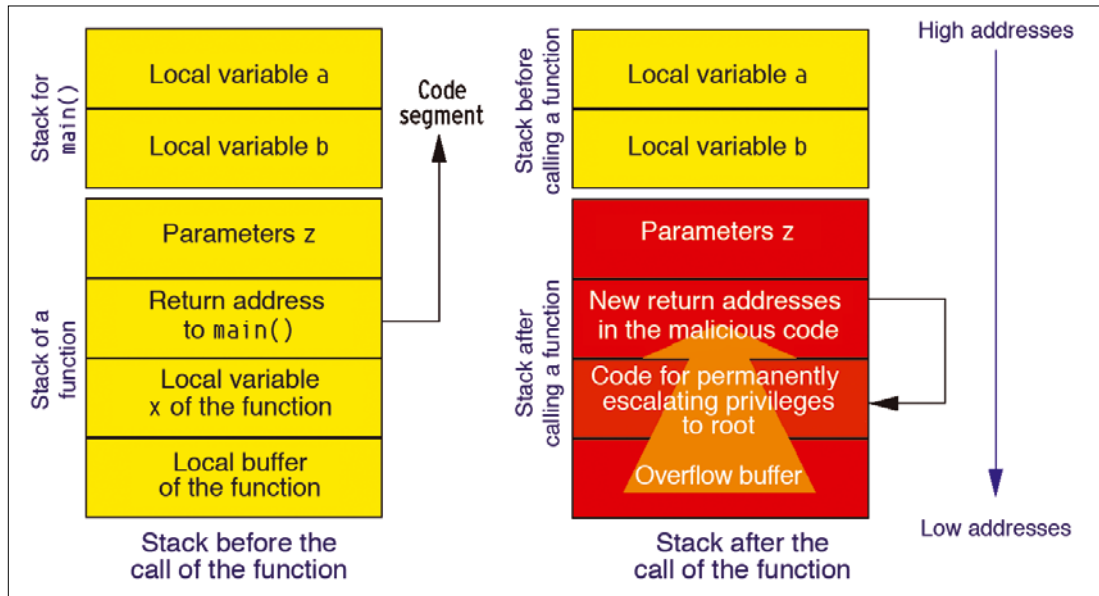
### Modular Is Cool

Attackers who have managed to inject malicious code in spite of the resistance usually look to install a rootkit, which they can do simply by loading their own kernel module into the system. To prevent this, the admin could compile the kernel statically – without module support – but a static kernel is often inconvenient, unless it is used with an embedded device. Alternatively, the Self-Protection project suggests that only a few locally logged-in users should be allowed to load modules.

### Conclusions

You can bet the Linux kernel has many more security problems that aren't yet listed in the CVE databases. The goal of the Kernel Self-Protection is to establish self-defense functions such as address space layout randomization to make the attacker's task more difficult and limit the damage of a successful attack.

When it comes to safe programming practices, doing one thing doesn't mean giving up on another. Targeted code reviews and intensive quality management should be an essential part of any programming effort. ■■■



**Figure 2:** Normally, the distribution of the address space is defined on a 32-bit system of the I-386 architecture. The heap and stack are moving towards each other from two sides as their size grows dynamically.

### Author

**Tobias Eggendorfer** is a professor of IT security in Ravensburg-Weingarten and a freelance IT consultant (<http://www.eggendorfer.info>).

### Info

- [1] OpenBSD security page: <https://www.openbsd.org/security.html>
- [2] "In send\_dg, the recvfrom function is NOT always using the buffer size of a newly created buffer," CVE-2015-7547: [https://sourceware.org/bugzilla/show\\_bug.cgi?id=18665](https://sourceware.org/bugzilla/show_bug.cgi?id=18665)
- [3] Patch for CVE-2015-7547: <https://www.sourceware.org/ml/libc-alpha/2016-02/msg00416.html>
- [4] Corbet, Jonathan. "Kernel vulnerabilities: old or new?": <https://lwn.net/Articles/410606/>
- [5] Kernel Self-Protection: <https://www.kernel.org/doc/html/latest/security/self-protection.html>
- [6] The GNU C Library Reference Manual, "Executing a File": [https://www.gnu.org/software/libc/manual/html\\_node/Executing-a-File.html](https://www.gnu.org/software/libc/manual/html_node/Executing-a-File.html)
- [7] C0ntex. "Bypassing non-executable-stack during exploitation using return-to-libc": [http://infosecwriters.com/text\\_resources/pdf/return-to-libc.pdf](http://infosecwriters.com/text_resources/pdf/return-to-libc.pdf)
- [8] Dang, Maniatis, and Wagner. "The Performance Cost of Shadow Stacks and Stack Canaries": <https://people.eecs.berkeley.edu/~daw/papers/shadow-asiaccs15.pdf>





**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](http://shop.linuxnewmedia.com/subs)

**GET IT NOW!**  
FAST DELIVERY  
WITH OUR PDF EDITION

## The sys admin's daily grind: Sysdig

## Alpha Beast

In this issue, sys admin columnist and tool veterinarian Charly Kühnast invites Sysdig, the jack-of-all-trades among system diagnostic tools, into his surgery for a quick checkup. The project promises to unite the functionality of *lsof*, *iftop*, *netstat*, *tcpdump*, and others.

By Charly Kühnast

Where an alpha beast claims to replace an entire herd, the bar is naturally fairly high. Of course, the Wire-shark authors, who are also the people behind the Sysdig [1] project, are no beginners. The software only performs well if you have root privileges; otherwise, it can't access all the required system areas. If you launch the tool without parameters, a steady stream of system messages scrolls by: It meticulously logs every single syscall. To thin out the thicket, Sysdig uses what it calls chisels. You can find out which chisels exist with the `sysdig -cl` command.

The chisels are sorted into categories (Net, IO, application, logs, and so on). For example, the Performance category has a chisel named `netlower`. I decided to pass in a time value of 10 milliseconds as a parameter:

```
sysdig -c netlower 10
```

Now Sysdig keeps listing processes whose network IO is slower than

10 milliseconds – on my home network, this means the SmokePing probes to the garden Raspberry Pis and some Munin connections.

You can output a list of the processes with the most frequent mass storage accesses by typing:

```
sysdig -c topprocs_file
```

The following reveals the entity causing the most network traffic:

```
sysdig -c topconns
```

A replacement for `top` can be found in:

```
sysdig -c topprocs_cpu
```

The built-in automatic analysis of bottlenecks is particularly informative. Typing

```
sysdig -c bottlenecks
```

generates a list of processes whose syscalls take a suspiciously long time. This is a great approach to searching for bottlenecks.

## Depth on the Interface

If you like a more interactive approach, try `csysdig`. The tool displays the information provided by Sysdig in a continuously updated ncurses interface. Called without parameters, the start screen reminds one of `htop`, but

pressing F2 takes you to a list of *Views* that correspond to the categories to which Sysdig assigns its chisels, and you can access them quickly and easily.

For example, if you choose the *Spectrogram-File* view, you are treated to a graphic like that shown in Figure 1: It shows the file access latency distribution, in which each line represents one second. At the time of grabbing the screenshot, an `apt dist-upgrade` was running, hence the high read and write load highlighted in red.

The *Views* overview showcases one of the specialities of Sysdig and *Csysdig*: You can restrict analyses to applications that run in containerized systems such as Docker or Kubernetes. Thus, admins can quickly and easily identify any performance fluctuations in containerized software.

My conclusions: Used only as a replacement for `top` and `netstat`, Sysdig is like taking a sledgehammer to crack a nut, but the many easily parameterized analyses of file and network latencies are a real help. If I have to dig down into individual syscalls, I can save a trace file and filter it until I find what I want. Here, at last, you can finally see the signature of the Wireshark makers. ■■■

### Info

[1] Sysdig: <https://www.sysdig.org>

### Author

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

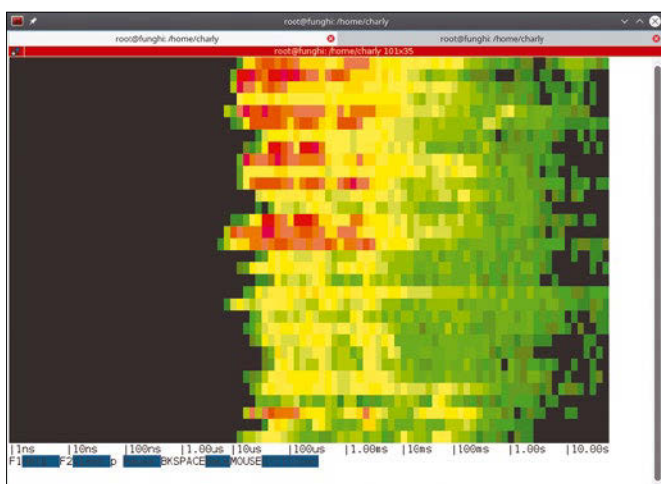


Figure 1: Abstract art? Nope: Concrete latencies during file access as a spectrogram, applied over time.

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!





Introducing the new serverless paradigm

# Look No Hands

Serverless computing takes cloud abstraction to a whole new level. Find out why more developers are turning to this exciting new programming paradigm. *By Swapnil Bhartiya*

**S**erverless is the next stage in the evolution of IT infrastructure. Serverless computing is a deployment model where functions or applications are uploaded to a platform and then executed, scaled, and billed in response to the demand needed for the moment. Users run their applications without spending resources on the complexity that comes with servers – provisioning, management, scaling, and more.

“The main drive behind this continuous evolution is to make it easier for a developer to focus on adding business value with their applications, without having to worry about the underlying infrastructure,” says Peter Chadwick, di-

rector of product management for cloud and systems management at SUSE.

Serverless computing provides two different kinds of functionalities: Backend as a Service (BaaS) and Functions as a Service (FaaS). BaaS replaces the subset of custom server-side functions of an application with API-based, third-party services. A good example of BaaS is Google’s Firebase. Firebase offers functionalities like analytics, databases, messaging, and crash reporting that can be plugged into an app through an API. Google runs and manages these services. A mobile developer doesn’t have to worry about maintaining a “server.” Google manages all of it, and you just pay when the application runs.

FaaS is a new way of deploying server-side software functions that can be triggered by events or HTTP API requests. According to an upcoming white paper by Cloud Native Computing Foundation (CNCF), developers deploy small units of code to FaaS, which are executed as needed as discrete actions, scaling without the need to manage the underlying infrastructure. This model removes the need for a virtual machine (VM) or a container. You just need an FaaS platform that runs these functions, which could consist of very basic code. FaaS is currently more popular than BaaS and is driving the adoption of serverless computing as an alternative for the IT industry.

Lead Image © bolina, 123RF.com

## Isn't That What PaaS Does?

If you think FaaS is trying to do what Platform as a Service (PaaS) projects like Cloud Foundry do – abstract the underlying infrastructure – you are not wrong. According to Chadwick, some customers taking their existing applications and put them in Cloud Foundry without rewriting everything. In theory, Cloud Foundry could be considered a FaaS platform.

From a developer's perspective, however, serverless offers a higher level of abstraction. PaaS abstracts the infrastructure and enables developers to create applications without worrying about the infrastructure. FaaS abstracts everything around the functions of that application. All you see is the function that you need to run. "It's a new architectural model for developers, where they decompose their applications into sets of functions that are triggered by events or HTTP API requests," says Docker's Patrick Chanezon.

Unlike traditional applications, these functions don't sit idle waiting to run. FaaS waits for the event to occur and then starts a container to run these functions; then it kills the container once it finishes its job. The developer has nothing to do with running the container – it is all done by the FaaS platform.

You can access a fully managed serverless platform through a cloud provider like Google [1], Amazon [2], or Microsoft [3], or you can build your own serverless solution using a locally hosted open source container platform. In either case, the technologies around FaaS are relatively new, and there is no standardization, so interoperability could be a challenge.

## Event Horizon

FaaS allows the user to create functions and respond to events. Imagine a user has uploaded an image or video file to the server through a mobile app. That is an event that could trigger a function that runs image recognition code, creates a thumbnail, or sends the video file to another service. Another good use case for serverless is an Internet of Things (IoT) device; for example, a temperature sensor could create an event that triggers a function to instruct a thermostat to adjust the temperature.

As you might have noticed, this "serverless" platform isn't really serverless. A server is still back there, with all the complexity required to manage VMs or spinning containers, but FaaS hides all that complexity.

## Back to the Future

To help users bring their traditional applications to the containerized world, players like Docker, IBM, and Microsoft have been working together to create solutions that help customers leverage technologies like serverless. Because serverless is a relatively new technology, the scope of use cases is still growing. Serverless is not a silver bullet for every application, but it is a great fit for a workload that sees sporadic demand, is concurrent, can be parallelized, and is asynchronous.

Serverless caters to developers who are more focused on the business logic and who want to scale out that business logic for the components that make up some of the application's back-end processes.

Batch processing is a good use case of serverless. "When you have a lot of data, such as images, you can send them in small batches that can then be processed one by one by serverless functions. IoT is another use case, where many events are happening, sending a great deal of data from devices. These events can trigger functions that can start a workflow," says Chanezon.

You may have noticed that Amazon Prime can provide you with information about the actors in a scene you are watching in a movie or TV show. This use case is well suited to serverless: The application can identify actors in each frame and pull relevant information from the database to display on the screen. Serverless can help a banking customer deposit a check using a mobile device that triggers a function when the device takes a picture of the check.

Machine learning will become a core part of serverless strategy, as it helps in automating all those triggers and actions. The possibilities are immense. "FaaS comes into play with back-end jobs that are very manual and task oriented," says CNCF's Ken Owens. "Companies can take those manual jobs and automate to create a smaller set of functions that can be independently scaled, deployed, and triggered."

## Buyer Beware

As exciting as it sounds, serverless is a very new concept that is continuously evolving, so tread carefully. The industry has not yet agreed on standards or best practices, and documentation is hard to find. In some cases, you might need to watch out for vendor lock-in, especially with BaaS. ■■■

## Info

- [1] Google Serverless: <https://cloud.google.com/functions/>
- [2] Amazon Lambda: <https://aws.amazon.com/lambda/>
- [3] Microsoft Azure Serverless: <https://azure.microsoft.com/en-us/services/functions/>

## Author

**Swapnil Bhartiya** is a writer and journalist covering Linux and open source for more than 10 years. He is also a science fiction writer whose stories have been broadcast on Indian radio and published in leading Indian magazines. He founded an open source web magazine while living in Europe. Swapnil currently resides in Washington, DC.

## Too Swamped to Surf?

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

Subscribe today to our free ADMIN Update newsletter and receive:

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



[www.admin-magazine.com/newsletter](http://www.admin-magazine.com/newsletter)



Getting started with the Tilix terminal emulator

# The EMULATOR

A terminal emulator with innovative features makes working from the command line easier and more efficient. *By Bruce Byfield*

**W**hen people talk about the command line these days, more often than not they mean a terminal emulator – a terminal running from a desktop environment. The most common exceptions are system administrators or old Unix hackers who are running without a desktop. Yet, despite the widespread use of terminal emulators, they have evolved little, if any, in the last two decades. As far as average users are concerned, the functional difference between, for example, Gnome Terminal and Konsole is minimal. What is unusual about Tilix (formerly Terminix) [1] is that it offers a wealth of innovative features that make working from a terminal emulator more efficient, from the basic features on the titlebar to the options available by drilling down into the preferences and profile settings.

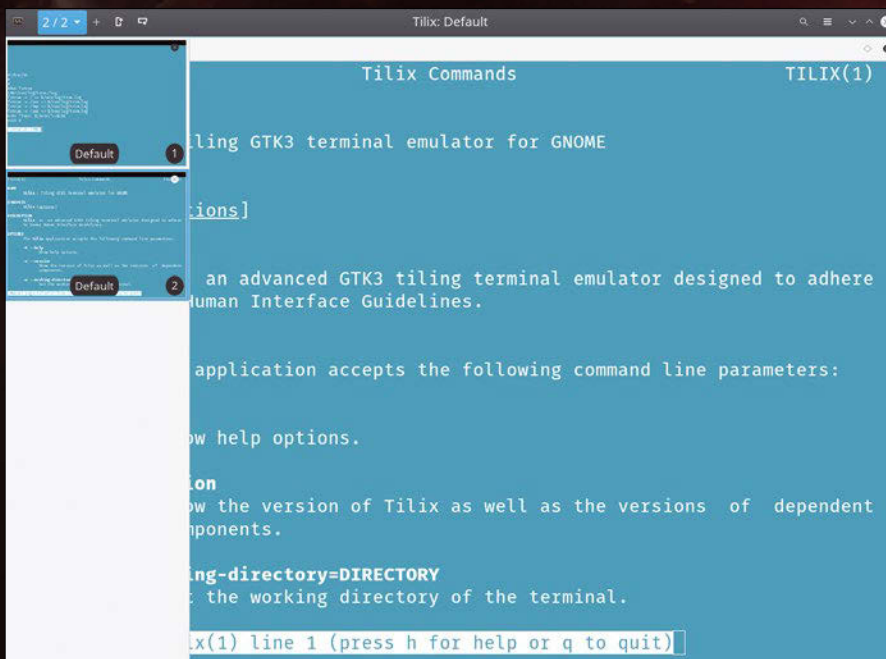
**Author**

**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>

In order to use Tilix’s full set of features, you may need to make some adjustments. The problem is that Tilix uses the Virtual Terminal Emulator (VTE) Gtk+ 3 widget. VTE’s configuration includes `/etc/profile.d/vte.sh`, which overrides the `PROMPT_COMMAND` in order to set which current directory to use when it is running. However, while

this arrangement works in some distributions, such as Fedora, other distributions, such as Arch and Ubuntu, execute scripts in `/etc/profile.d` only for login shells. Since by default VTE does not run in login shells, the result is that some of Tilix’s features, such as notifications, will not work immediately after installation.

Lead Image © Luká Gojda, 123rf.com



**Figure 1:** Tilix offers the option of displaying open terminals in a sidebar.

```

TILIX(1)          Tilix C
ommands          TILIX(1)  #!/bin/sh
#
#
NAME             #Run fstrim
                LOG=/var/log/trim./log
                fstrim -v / >> /var/log/trim.log
                fstrim -v /usr >> /var/log/trim.log
                fstrim -v /tmp >> /var/log/trim.log
                fstrim -v /var >> /var/log/trim.log
                echo "Time: $(date)">>$LOG
                :
SYNOPSIS         tilix [options]
DESCRIPTION     tilix is an advanced GTK3 tiling
                terminal emulator designed to adhere
                to Gnome Human Interface Guidelines.
OPTIONS         The tilix application accepts the following
                command line parameters:
                -h --help          Show help options.
                -v --version      Show the version of Tilix
line 1 (press h for help or q to quit)

#!/usr/bin/fontforge
i=1
while ( i<$argc )
  Open($argv[i])
  Generate($argv[i]:r + ".otf")
  i = i+1
endloop
convert.pe (END)

```

**Figure 2:** Terminals can be split to display multiple terminals in the same window.

What this situation means is that, in order to make use of Tilix’s full set of features, you may need to modify the `.bashrc` or `.zshrc` file in your home directory to execute `vte.sh` directly or set each profile on its Command Page to run as a login shell, as described on the Tilix website [2]. Alternatively, you may decide to ignore the issue, since even without the adjustment, Tilix retains its basic

functionality and still has plenty of features that are unaffected by VTE’s configuration.

## Default Features in the Main Window

Tilix’s titlebar includes several key features besides the configuration icon on the far left and the usual titlebar buttons for maximizing, minimizing, and

closing the window on the far right. Second from the left is a toggle for a sidebar that displays all the open terminal windows in a pane on the left of the window, like the one for showing presentation slides in LibreOffice’s Impress.

If you prefer, you can make the sidebar open on the right by clicking the configuration icon and selecting *Preferences | Appearance | Place the sidebar on the right*. Starting with Tilix 1.7, you also have the option of using tabs instead of the sidebar. However, one advantage of the sidebar is that, if your close-up vision is up to the task, you can see exactly what each terminal contains – and, even if you can’t, you may be able to distinguish a terminal by the general appearance of its contents. And changing terminals in the sidebar is certainly more efficient than a separate desktop window for each terminal. You can add terminals by clicking the plus sign third from the left in the titlebar (Figure 1).

However, Tilix does not have to create a new window or tab to give you access to another prompt – although you can create either by using the `tilix` command from within a running instance, as explained in the man page. Instead, you can use the next two icons from the left to split the current screen vertically or horizontally (Figure 2). The tools on the far right before the standard titlebar icons are designed to make working with split screens easier. You can search for text in the active prompt, name the active prompt to make it easier to find, or even synchronize all the prompts, so that what you type in one is entered automatically in the others. For instance, if each prompt has a different current working directory, you could enter the `ls` command just once to see the content of each directory and then enter `clear` to remove the output at each prompt before turning synchronization off (Figure 3). In addition, each prompt can be dragged and dropped to a new position in the Tilix window in order to keep your work organized.

Within each prompt, you also have other options. From the titlebar’s context menu (Figure 4), you can set the current prompt to read-only or password-protected. The context menu also gives you

```

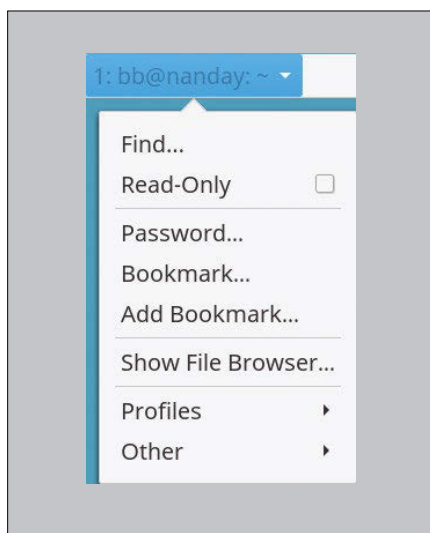
bb@nanday:~/Downloads$ ls
pg2850-images.epub
pureos-8.0-live-amd64.hybrid.iso
sfu-continuing-studies-jan-jun-2018-viewbook.pdf
S005302823-abbiegardner7-320.zip
S005407495-jp Pettit-320.zip
Some aspects of wage payments and coinage in ancient Rome (1st-3rd century AD), in J. Lucassen (ed.), Wages and Currency: Global and Historical Comparisons, Peter Lang, Bern-Berlin, 2007, p. 77-96. | Johan van Heesch - Academia.edu.html
TNS_BruceByfield_FreelanceWriterAgreement.pdf
Tracks - Jocelyn Pettit.zip
Using_OT_in_LibreOffice_1-2.pdf
VancityVisa(1).pdf
VancityVisa.pdf
view-from-bed.JPG
what-things.pdf
Wishes_on_a_Neon_Sign.high.zip
writer2epub
writer2epub.oxt
writer2xhtml.oxt
bb@nanday:~/Downloads$

bb@nanday:~/scripts$ ls
bluetooth-autostart.sh
bluetooth-autostart.sh-convert.pe
convert.pe
spectre-meltdown-checker.sh
trim.sh
bb@nanday:~/scripts$

bb@nanday:~/Pictures/Art/sheldon-dennis$ ls
100_1621.JPG
100_1640.JPG
howdesignpainting1.jpg
how design painting.jpg
how designpainting2.jpg
firstyearmask2.jpg
firstyearmask.jpg
night-waver.jpg
shilton-bear-9ugs.png
bb@nanday:~/Pictures/Art/sheldon-dennis$

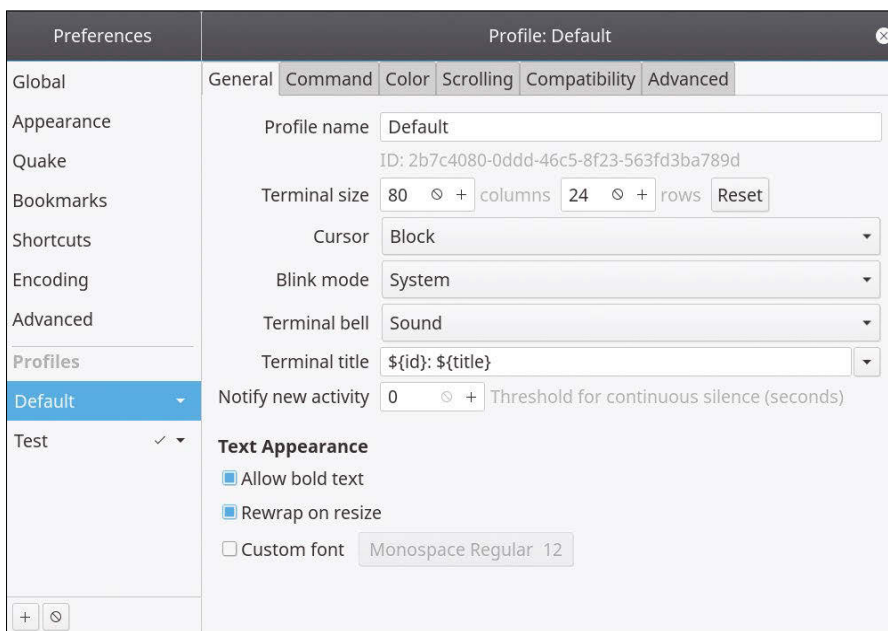
```

**Figure 3:** Synchronizing prompts can reduce typing. Here, entering the `ls` command at one prompt displays the contents of the working directories at all three prompts.



**Figure 4:** Each terminal has its own menu of options.

access to bookmarks to increase navigation speed and opens your desktop’s default file manager. From the Other sub-menu, the context menu also lets you save output to file, a more efficient alternative to creating a file than copying and pasting to it.



**Figure 5:** Each profile has six tabs full of cosmetic and functional customizations.

### Cosmetics and Functionality

The controls in the main window would be enough by themselves to justify Tilix. However, they are just the beginning. To start with, like most terminal emulators, Tilix profiles support a wide variety of customizations. From *Preferences | Appearances*, you can set general options such as the default theme. These general options can be overwritten by specific profiles (Figure 5). For instance, the General page for each profile has options for setting the default terminal size, the cursor, notifications, and font size, weight, and family. Other tabs set background and foreground colors and scrolling options. All the customizations in a profile override the general options in *Preferences | Appearances* (Figure 6), which include the default theme.

However, what makes Tilix stand out is not cosmetic features but functionality. If you want to be notified when a running process not currently displayed is finished or to use triggers to run certain commands when defined conditions are met, you will need to configure VTE as defined above. But even without VTE being fully functional, Tilix still has numerous features to help streamline functionality. For example:

- From *General | Terminal* in any profile, you can change the contents of each prompt’s title – a feature that can be useful in helping to locate it when multiple prompts are open.
- Right-clicking at a prompt allows you to paste. From *Preferences | Global*, you can choose, among other options, to use the advanced paste dialog, which gives options to convert spaces to tabs, adjust carriage returns, or copy by selecting characters with the mouse.
- Save and restore the current state of a terminal from *Preferences | Global*.
- Click hyperlinks in the terminal to open them in your desktop’s default web browser.
- Set a Quake-style hotkey for quick opening of a terminal that is designed to work in Wayland. *Preferences | Quake* sets the size and position of the terminal that opens with this feature, as well as the option to have it open on all workspaces, and to add it to the active monitor when multiple ones are used.
- Customize automatic profile switching in a profile’s Advanced page.

The total number of features is actually impossible to list and detail without making this article at least twice as long. Many have the virtue of simplicity, yet if you are anything like me, after using them for a while, you may wonder how such an obvious feature has escaped being standard in other terminal emulators.

## MOBILE USERS

search for us today at your digital newsstand!



Only a swipe away!

Download our convenient digital editions for your iPad, iPhone, or Android device.

ADMIN Magazine

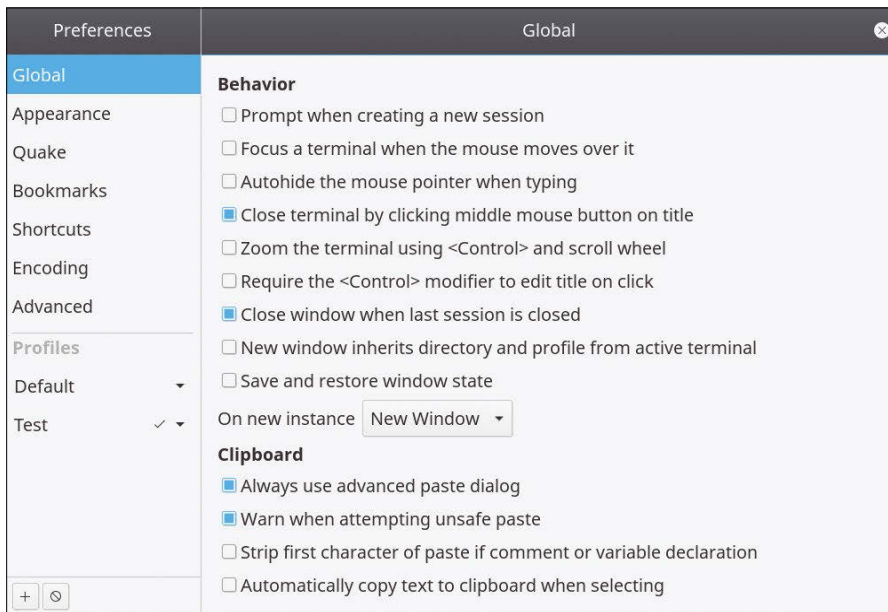


Linux Magazine



Visit our apps page for more information  
[shop.linuxnewmedia.com/us/apps](http://shop.linuxnewmedia.com/us/apps)





**Figure 6:** Tmux's Preferences set default options that can be overridden by a profile.

### Some Drawbacks

Tmux is in rapid development, which can sometimes cause problems. The problems may be minor. For example, Custom Links, which will open in a web browser when clicked, are still

listed on *Preferences | Advanced*, when experimentation shows that links no longer have to be entered for this feature. Nor has the documentation been kept up to date with the features in every case.

More seriously, in my experience, settings may work inconsistently, either requiring a restart of Tmux to apply or else not working at all. Very occasionally, Tmux freezes or crashes.

Still, these drawbacks are not surprising. According to the Debian package pages, the main Tmux package is almost six times larger than the Bash package – 35,000 KB to Bash's 6,000. Tmux is still small enough that it can be used with limited memory, but the difference in size reflects how many more features are packed into Tmux than a standard shell.

What I like about Tmux is that its developers are rethinking a part of Linux that everyone has taken for granted for years. Even if you don't use all of its features, you are still likely to find half a dozen that make working from the command line easier and more efficient. Tmux, I submit, is free software at its best. ■■■

### Info

- [1] Tmux: <https://gnunn1.github.io/tmux-web/>
- [2] VTE configuration: <https://gnunn1.github.io/tmux-web/manual/vteconfig/>

Shop the Shop → [shop.linuxnewmedia.com](http://shop.linuxnewmedia.com)

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

Want to subscribe?



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

➤ [shop.linuxnewmedia.com](http://shop.linuxnewmedia.com)

DIGITAL & PRINT SUBSCRIPTIONS



SPECIAL EDITIONS



## Programmatically change YouTube metadata

# Film Maker

Instead of manually editing the metadata of YouTube movies, video craftsman Mike Schilli dips into YouTube's API spell book and lets a script automatically do the work. *By Mike Schilli*

If one day someone should invite me onto a TV talk show, in which I could brag about my best productivity tricks, I would say this: The most effective way forward is to be able to make quick changes at any time, which you can revert later on – just as quickly and almost without loss – if they turn out to be duds.

Software development with Git is a good example of this: Git makes it possible to boldly add new functions to programs in huge leaps or to rearrange the code on a large scale. If the whole thing turns out to be a crazy idea after a while, the developer simply reverts everything in just a second without anyone noticing what pipe dreams they were chasing in the meantime.

That's why I religiously put everything I produce under version control. My blog posts, my articles, my programs: The version control system lets me revert to yesterday's status at any time should the need arise, and for any reason.

But up to now, I have not been able to get the procedure to work for my YouTube videos. For example, if I change the title or even just one tag of one of my internationally-respected, quality movies (Figure 1), I can't simply tell YouTube

**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](mailto:mschilli@perlmeister.com) he will gladly answer any questions.



later on: Today, all I've done is produce nonsense; please undo everything and make it look like yesterday.

**Metadata with Version Control**

Therefore, I want to entrust the metadata of my YouTube existence to a local YAML file (Listing 1) and version it with Git. A script then regularly checks the entries in the file, reads the ID of each movie from the YAML data, and checks whether everything is set up according to the file (i.e., whether the title in the file matches the

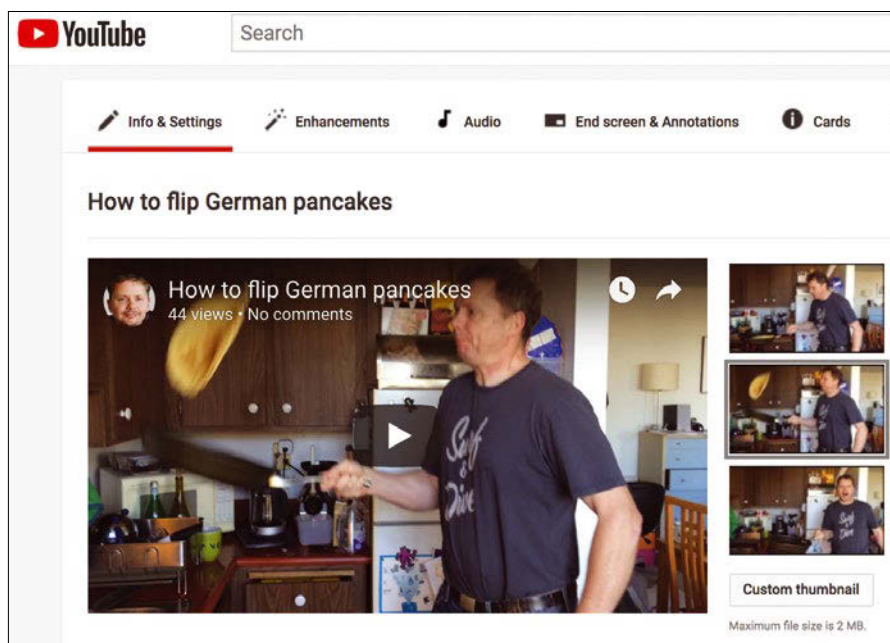
video page's metadata). In the event of differences, the script adapts the on-line YouTube data to the YAML data. Listing 1 [1] defines a long list of videos with their IDs and titles. It could easily be extended with tags, a thumbnail image, a release date, or other metadata.

**At the User's Service**

But before the script is allowed to access or change the user data, Google

**Listing 1: videos.yaml**

```
1 videos:
2   - id: _3i5yVoTvCs
3     title: "How to flip German pancakes"
4   - id: brPfe66FC24
5     title: "Tivo Stream Cooling Fan Replacement"
6   - id: 2qxXhW7RxsY
7     title: "Rio Portable Beach Shelter Assembly Instructions"
```



**Figure 1:** Mike Schilli as a pancake tossing chef on YouTube.

## Listing 2: youtube-sync

```

01 #!/usr/bin/python
02 import httplib2
03 import os
04 import sys
05 import yaml
06
07 from apiclient.discovery import build
08 from apiclient.errors import HttpError
09 from oauth2client.client import \
10     flow_from_clientsecrets
11 from oauth2client.file import Storage
12 from oauth2client.tools import \
13     argparser, run_flow
14
15 CLIENT_SECRETS_FILE = "client-secrets.json"
16 YOUTUBE_READ_WRITE_SCOPE = \
17     "https://www.googleapis.com/auth/youtube"
18 YOUTUBE_API_SERVICE_NAME = "youtube"
19 YOUTUBE_API_VERSION = "v3"
20
21 def get_authenticated_service(args):
22     flow = flow_from_clientsecrets(
23         CLIENT_SECRETS_FILE,
24         scope=YOUTUBE_READ_WRITE_SCOPE)
25
26     storage = Storage("oauth2.json");
27     credentials = storage.get()
28
29     if credentials is None or \
30         credentials.invalid:
31         credentials = \
32             run_flow(flow, storage, args)
33
34     return build(YOUTUBE_API_SERVICE_NAME,
35                 YOUTUBE_API_VERSION,
36                 http=credentials.authorize(
37                     httplib2.Http()))
38
39 def video_update(youtube, id, title):
40     response = youtube.videos().list(
41         id=id, part='snippet').execute()
42
43     if not response["items"]:
44         print("Video '%s' was not found." % id)
45         sys.exit(1)
46
47     snippet = response["items"][0]["snippet"]
48
49     if snippet['title'] == title:
50         print("%s: Unchanged" % id)
51         return
52
53     snippet['title'] = title
54
55     try:
56         youtube.videos().update(
57             part='snippet',
58             body=dict(
59                 snippet=snippet, id=id)).execute()
60     except HttpError, e:
61         print("HTTP error %d: %s" % \
62             (e.resp.status, e.content))
63     else:
64         print("Updated OK")
65
66 if __name__ == "__main__":
67     args = argparser.parse_args()
68     youtube = get_authenticated_service(args)
69
70     stream = open("videos.yaml", "r")
71     all = yaml.load(stream)
72     for video in all['videos']:
73         video_update(youtube, video['id'],
74                     video['title'])

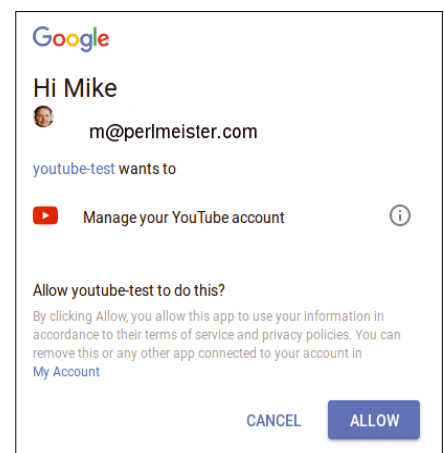
```

has to grant the new application access to the account holder's data, because Joe Public is not allowed to mess around with my videos. Google obtains the user's consent in Listing 2 using the Google API [2] to guide the user's browser to a Google page on which the user logs in and then confirms that the application actually should have access (Figure 2).

To do this, API jockeys need to create a new project on the Google Cloud Platform Console [2] (Figure 3). Then navigate to *Create credentials* (Figure 4) and select *OAuth client ID* [3] (not *API key*, that's for project management only). Since it is a desktop program

and not a web application, you have to select *Other* in the selection menu for the application type. The strings then produced by Google for the *Client ID* and *Client Secret* (Figure 5) need to be entered into a JSON file as shown in Listing 3.

After the script's first run, when the user has successfully granted permission in the browser, the browser branches to a page with a *The authentication flow has completed* message. The script drops an OAuth 2 access token into the `oauth2.json` file, thus granting access to the user's data for future calls without the user having to agree again.



**Figure 2:** The first time the script is called, it opens a browser that asks for the user's consent.

Listing 3: client-secrets.json

```

1 {
2   "installed": {
3     "client_id": "XXX",
4     "client_secret": "YYY",
5     "redirect_uris": ["http://localhost",
6       "urn:ietf:wg:oauth:2.0:oob"],
7     "auth_uri": "https://accounts.google.com/o/oauth2/auth",
8     "token_uri": "https://accounts.google.com/o/oauth2/
9     token"
10  }
11 }

```

This works until the access token expires. The corresponding expiration date is also noted in the JSON file. The file also contains a refresh token that the script can use to request a new token after the current access token expires. This practically works infinitely – unless the user goes to the Google Console (Figure 5) and revokes access for the client, in which case Google pulls the plug.

The `get_authenticated_service()` function in line 21 of Listing 2 defines `YOUTUBE_READ_WRITE_SCOPE` as the access scope, thus requesting read and write access. The interaction with the browser and the underlying OAuth 2 token dance is nicely abstracted by the Google API SDK; the script only calls the `flow_from_clientsecrets()` and `run_flow()` functions from the two `oauth2client.client` and `oauth2client.tools` packages. This works equally well in the browser on both Linux and Mac OS.

Line 70 reads the YAML file with the locally stored video metadata; line 71 iterates over all movies located there. For each of them, it calls the `video_update()` function defined in line 39, which first retrieves the online movie's metadata by its ID using `youtube.videos().list()`. Data retrieved is limited to the `snippet` area, which YouTube uses to designate the title, tags, description, and a few other fields.

Line 49 retrieves the movie title from this metadata and compares it with the local version. If the two titles do not match, line 56 triggers the `update()` method in a `try` block. The function expects the originally retrieved metadata along with the new, modified title as parameters. If an error occurs during transfer, line 61 prints the HTTP error message; otherwise line 64 reports `Updated OK` – and the metadata on YouTube now exactly matches the metadata stored locally in the version control system.

## Installation and Outlook

The SDK needed to run the script is available as a Python package from the standard repository and can be installed via `pip`:

```

pip install --user ?
--upgrade google-api-python-client

```

After the script's first run – with the browser window popping up at the beginning to obtain the user's consent – it produces the following output with the YAML file in Listing 1:

```

<$> ./youtube-sync
  _3i5yVoTvCs: Unchanged

```

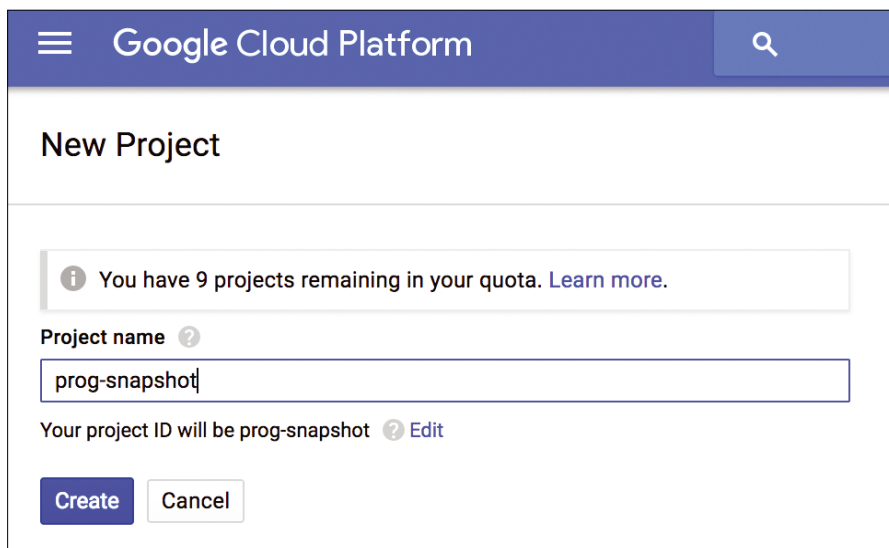


Figure 3: Creating a new project on the Google Cloud Platform Console.

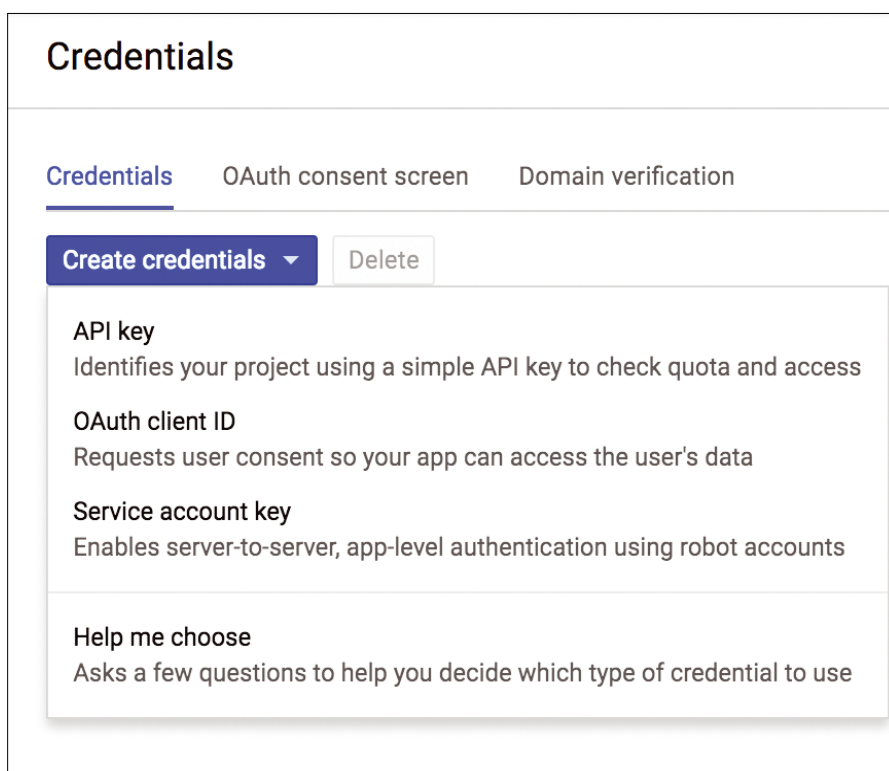


Figure 4: Google produces the client ID and key.

```
brPfe66FC24: Unchanged
```

```
2qxXhW7RxsY: Unchanged
```

Since all title strings in the retrieved YouTube metadata match the YAML data, the script does not make any changes. However, if a title text changes in the YAML file, the script fires it off to YouTube, which refreshes the movie's metadata. The script confirms such an action with Updated OK.

The script can be extended as needed, so that it can also upload videos from the hard disk to YouTube if required and automatically synchronize the local status of your movie collection with the publicly accessible movies on YouTube in each run. If you version your local collection and its

metadata with a version control system, you can jump back and forth in time, mak-

ing changes, and reverting if an idea turns out to be less than brilliant. ■■■

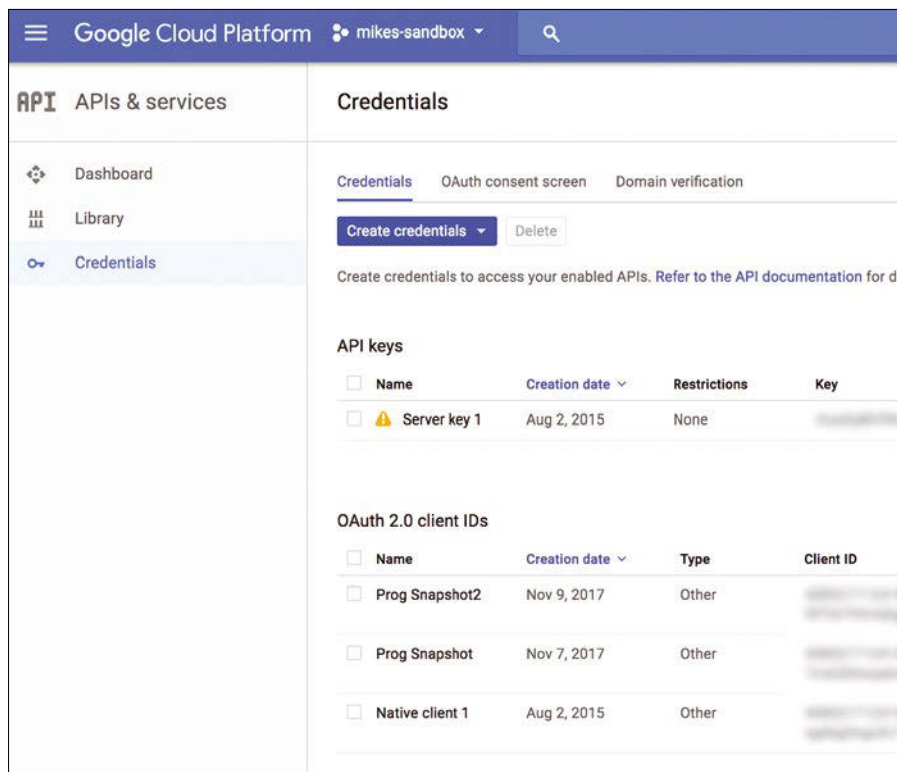


Figure 5: Active API keys on Google Console.

## Info

- [1] Listings for this article: <ftp://ftp.linux-magazine.com/pub/listings/linux-magazine.com/209/>
- [2] Google Cloud Platform Console: <https://console.cloud.google.com/apis>
- [3] "OAuth 2.0 for Mobile & Desktop Apps": <https://developers.google.com/identity/protocols/OAuth2InstalledApp>


# A Webzine for High-Performance Computing Specialists



**ADMIN**  
Network & Security

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

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



# MakerSpace

Desktop tools for self publishing

## DIY Ebook Publishing

LibreOffice, Calibre, and Sigil help would-be authors with do-it-yourself ebook publishing. *By Bruce Byfield*

**S**elf-publishing is not usually considered part of the Maker movement. Yet the two have in common a do-it-yourself approach and increasing popularity. Although the word processor was a major advance over the typewriter, until recently, the gap between home publications and professional ones loomed large. In the last few years, however, thanks to advances in free software such as the HarfBuzz font renderer [1] and the widespread availability of export filters, the differences between amateur and professional publications has all but disappeared.

The change is particularly noticeable in ebook publication. In 2016, almost 150,000 ebooks were published in the United States alone [2], and that number has almost certainly continued to rise in the last two years. Today, instead of working with exorbitantly priced online services, would-be authors can prepare their works for Internet download using such desktop tools as LibreOffice, Calibre, and Sigil.

The major ebook formats are PDF, EPUB [3], and MOBI [4]. Each is available in several different versions. PDF is a legacy format, having been in use long before the rise of ebooks. Being a close relative to the PostScript printing language, PDF is ideal for exporting complex formatting, but it has the disadvantage of having fixed font sizes that cannot be changed in many e-readers and software browsers.

By contrast, EPUB and MOBI sometimes have trouble with complex formatting, but they are both widely used. For many purposes, their output is indistinguishable. EPUB is notable mainly for being free-licensed. MOBI is notable for being, in the slightly modified AZW form (aka Kindle format), the format required for an ebook posted on Amazon. Because of its license, EPUB remains the most popular filter for software applications, although MOBI and AZW support is also sometimes available.

### EPUBS and LibreOffice

LibreOffice has had extensions for exporting to ebook formats – mostly, to EPUB – for years. All these extensions can be added by clicking *Tools | Extension Manager*, navigating to the downloaded extension, and then restarting LibreOffice.

Until now, extensions for working with ebook formats were severely limited [5]. For example, eLAI is intended primarily for the ILIAS e-learning platform, making its EPUB export too specialized for many uses. Writer2ePUB is no better; although MOBI export is listed among its preferences, it is only possible with the installation of KindleGen, which is apparently an extension that is no longer available. In addition, most of the extensions choke on all except the simplest formats, and several are seriously outdated. None of these legacy extensions should be used, despite being still available.

The results are much better with the EPUB export filter integrated directly into the menus of LibreOffice 6.0. This filter is available from *File | Export As*. To have LibreOffice do all the work for you, click *Export Directly As EPUB* in the submenu. However, if you want to choose the EPUB version or if the file is split by heading or page break, click instead *Export As EPUB*. Both produce an acceptable EPUB file not only for manuscripts, but also for many documents with complex formatting, making them a vast improvement over the obsolete extensions (Figure 1).

However, unlike other ebook export tools, LibreOffice's tools do require you to correctly place title pages rather than positioning them for you. Also, if you require any ebook format other than EPUB, you will need another tool to do the conversion.

## Calibre

Despite an eccentric interface, Calibre is one of the leading applications for all matters to do with ebooks. It is probably most used for managing ebooks and reading devices. However, it also has extensive tools for fine-tuning ebooks and supports a wide variety of formats. In fact, with Calibre, you can not only convert files to ebook formats, but also convert ebook files to desktop formats such as DOCX, PDF, RTF, and TXT.

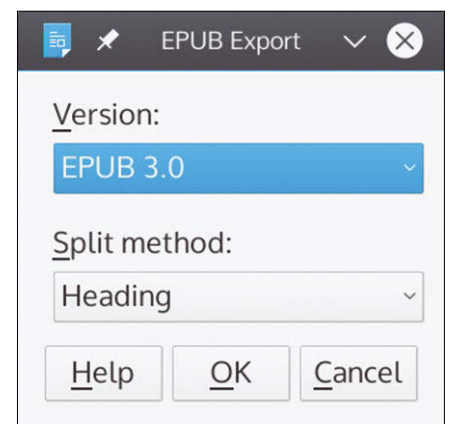
All books managed by Calibre must be registered by clicking the *Add books* icon. When you select a book being managed by Calibre by clicking the *Convert books* icon, you are taken to a second window. There, you can select the *Input* format on the window's top left and the *Output* format on the top right. Once you have done that, preparing an export is simply a matter of starting at the top of the left-hand pane and working through the subscreens one at a time (Figure 2).

Calibre's conversion window starts with the *Metadata* – defining information such as the Title, Author, and Publisher. You can also select a file for the cover. From there, you work downwards, selecting the default font and line spacing in *Look & Feel*, and formatting in *Table of Contents*. Contrary to what you might expect from experience with a word processor, *Page Setup* is not so much about features like headers and footers – although you can adjust margins – as choosing an output profile typical of a particular type of e-reader.

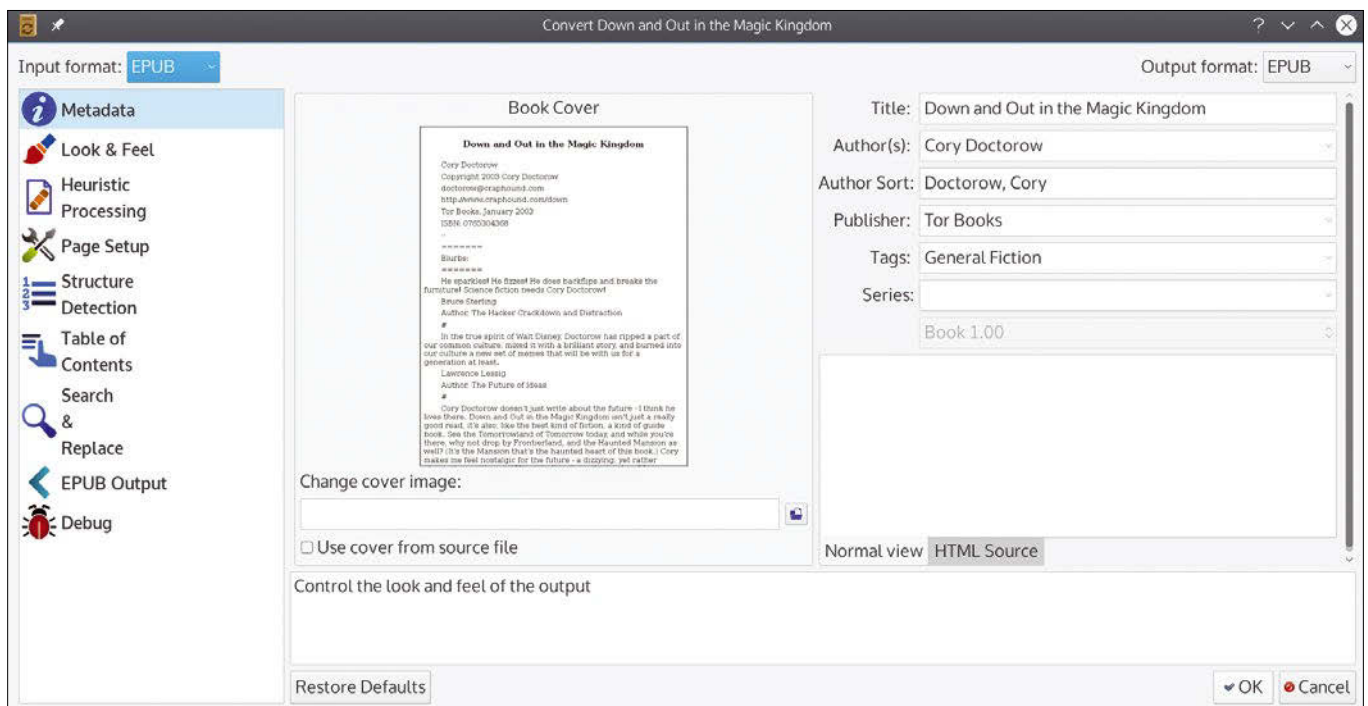
Mixed in with these features are tools for manual editing, such as *Heuristic Processing*, *Structure Detection*, and *Search & Replace*. If you are having problems, you can also use *Debug*. With these tools, you can manually edit the output file to correct problems before you create another exported file. For all but the sim-

plest layout, a file is likely to require at least some tweaking before the output is satisfactory. Once the output is tweaked, you can select standard *Output* options before generating the final output.

What is peculiar about Calibre's options is that in many cases, you are not manipulating the output directly. Instead, you are selecting options, and leaving Calibre to make the changes for you. The manual options are just in case, and you may be lucky and not need to use them at all. In the same way, the defaults of the fields on each subscreen can be left unchanged, so there is no need for anxiety options.



**Figure 1:** LibreOffice 6.0 offers limited options for controlling output to EPUB, but often handles complex formatting without tweaks from users.



**Figure 2:** Calibre offers numerous options for exporting to ebook formats, but in an oddly hands-off interface.

Except for metadata, many options do not specifically need to be changed. This mixture of detailed options and intelligent defaults makes Calibre the ebook producer of choice for many. As a matter of fact, I suspect that Calibre is the tool being used by third-party conversion services.

## Sigil

Those who want a more hands-on approach than Calibre allows may prefer Sigil. Hovering a minor version away from its first general release, Sigil is an editor that resembles a traditional word processor. It differs from a word processor in that it is designed to edit the text and CCS style sheets and to manage the document's elements, including images, fonts, audio, video, and the table of contents. Only EPUB is supported (Figure 3).

The best way to learn Sigil's capabilities is to open an existing file and study how it is structured. If nothing else, editing an existing file for your own purposes can save time, especially when managing the document's elements.

Sigil is especially useful for generating a file's contents. Organizing the text by sections, it offers a basic but adequate editor. If problems cannot be fixed in the text editor, you can drop down to the

Misc directory and view and edit the content.opf file directly if you are familiar with markup languages and CSS. Similarly, using the *Generate Table of Contents* icon in the toolbar, you can choose which headings to include, either selecting them individually or by style.

As elements are added via the editor, they are cataloged automatically in the other sections of Sigil, where they provide a different perspective on the document. For example, instead of searching for an image, you can edit it from the Images directory or see what fonts are used without checking the markup.

Sigil has its limitations. Its interface does not, for instance, handle tables, although you could add one by editing the markup directly. In addition, its default display for paragraph styles is extremely small for unknown reasons – although, fortunately, the display includes a zoom. All the same, Sigil can help correct problems when LibreOffice and Calibre simply lack the tools required.

## Conclusion

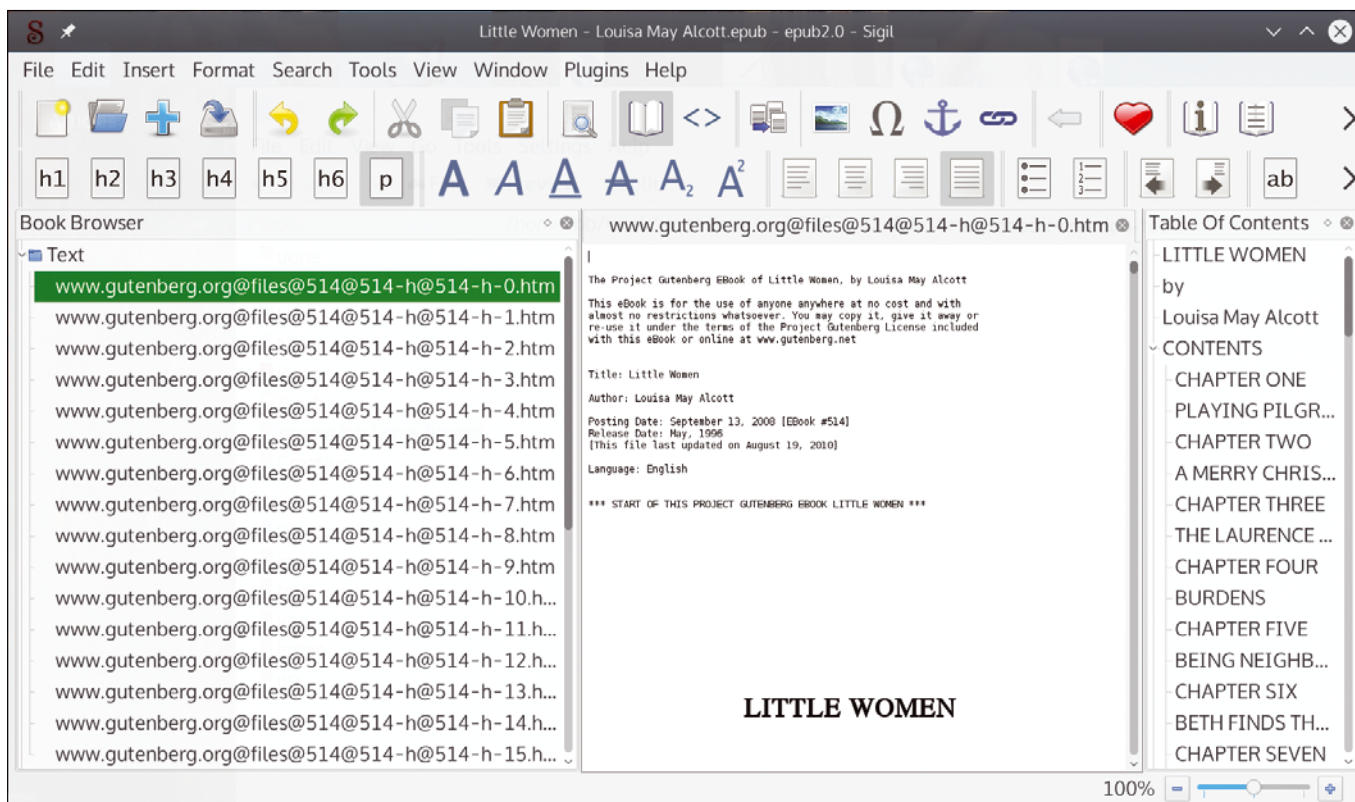
Numerous other scripts exist for creating ebooks. Debian especially offers a thorough set of scripts for creating ebooks from AbiWord, Apple Pages, DocBook XML, Microsoft Works, WordPerfect, and

others [6]. Before publishing, you may also want to verify your output, either using an online verifier or a free tool such as EpubCheck [7] from your distributor's package repositories.

For general purposes, LibreOffice 6.0 or Calibre should produce acceptable output. However, keep Sigil for hands-on troubleshooting. Between these three tools, you should have all you need for do-it-yourself ebook publishing. ■■■

## Info


- [1] HarfBuzz: <https://en.wikipedia.org/wiki/HarfBuzz>
- [2] Ebook self-publishing stats: <https://www.statista.com/statistics/249036/number-of-self-published-books-in-the-us-by-format/>
- [3] EPUB: <https://en.wikipedia.org/wiki/EPUB>
- [4] MOBI: <https://en.wikipedia.org/wiki/Mobipocket>
- [5] LibreOffice ebook conversion extensions: <https://extensions.libreoffice.org/@@search?SearchableText=epub>
- [6] Debian scripts for ebook conversion: <https://packages.debian.org/search?keywords=epub&searchon=names&suite=stable&section=all>
- [7] EpubCheck: <https://packages.debian.org/stretch/epubcheck>



**Figure 3:** In contrast to LibreOffice or Calibre, Sigil approaches ebook creation much like a word processor.



# 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



# MakerSpace

## Retro games on the Raspberry Pi Gamer's Paradise

With the right software, you can turn the Raspberry Pi into a versatile console for retro games. *By Erik Bärwaldt*

**I**f you are a gaming console enthusiast, the Raspberry Pi is a powerful accessory for your hobby. In addition to recalboxOS and RetroPie, Lakka is now climbing into the ring on this small-board computer and bringing old console games to life.

If you are of a certain age, you might remember the entertaining games from the 80s with simple block graphics, an equally simple plot, and background noise that usually consisted of a few notes. These technically and visually simple games have recently become popular again, and you don't have to invest in consoles or online games: A cheaper and more flexible Raspberry Pi is more fun, and you can play on your TV set, just like the old days.

The retro gaming platform Lakka [1], a LibreELEC derivative with the RetroArch desktop, a Rasp Pi, and a TV or monitor can mutate into several gaming consoles thanks to a wide variety of emulation options.

You don't even have to own a Rasp Pi 3 (RPi3). An older model works, too, with restrictions. A hardware guide [2] shows which game emulators are compatible with which edition of the Rasp Pi. Lakka's programmers have done their homework: The LibreELEC derivative does not require any laborious manual configuration or connection to a modern TFT display: Even an old tube TV is enough.

### Setting Up Lakka


After downloading the approximately 300MB Lakka image from the project page [3] and unpacking it, transfer the image file to the microSD card with the command:

```
$ dd if=Lakka-RPi2.arm-<Version>.img \
of=/dev/sd<X> bs=1M
```

Replace <Version> with the version number of your downloaded image and the <X> in the device name of the target with the correct device name. Some distributions also address the card as /dev/mmcblk0; if so, then specify this block device as the destination. Partition names like sdb1 do not work here.

After restarting the Rasp Pi, the dashboard appears with RetroArch's main menu; navigate through it with the arrow keys. Press Esc to return to the main menu from submenus. Early on, you should check Internet access. If you are using an older Rasp Pi model, you can connect via Ethernet. With the RPi3, Lakka detects the wireless module and offers the corresponding configuration in the Settings menu (Figure 1).

The Settings menu can be accessed by activating the gearwheel symbol in the dashboard and then selecting the *WiFi* option. In the *Information | Network Information* submenu, the current IP address appears next to the local host IP address if a network connection is available.



**Figure 1:** The wireless RPi3 module instantly recognizes Lakka and offers the corresponding options.

The developers are constantly improving Lakka, so the next step is to update the system by going to *Online Updater* | *Update Lakka* and following the on-screen instructions.

## Customizing

The customization options for the system are in the Settings menu. To change the menu hierarchy to a different language, choose the *User* sub-menu. Use the arrow keys to change to the *Language* point and select the desired language. Clicking *Esc* takes you back to the Settings menu and saves your changes.

To enable the mouse, choose *Settings* | *Menu*. First slide the switch for the *Show Advanced Settings* option to *On*, and then move the slider beside the *Mouse Support* option to *On*.

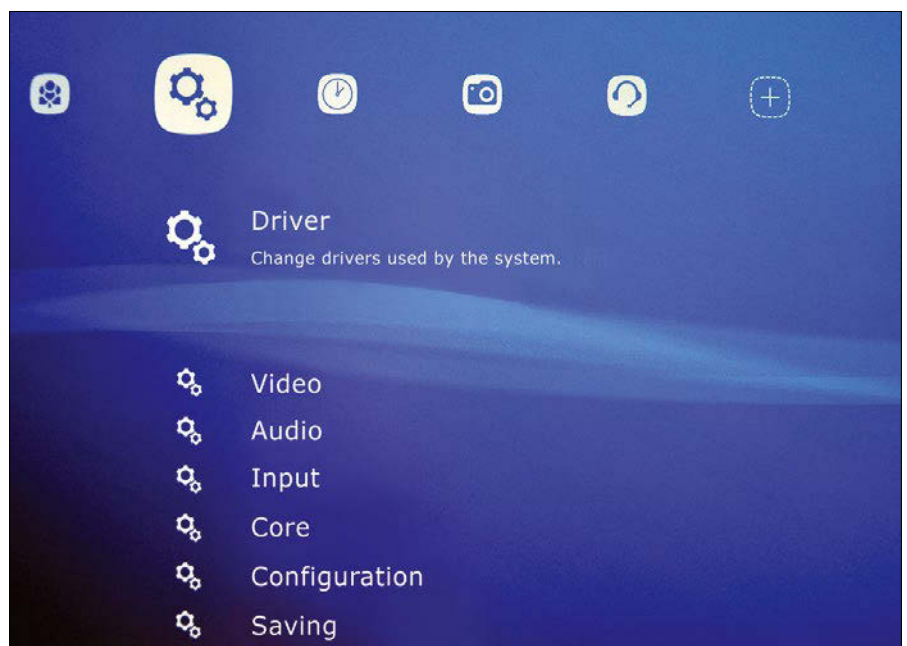
The somewhat confusing menu structure in Lakka often requires a look at the documentation [4]. The software is still under development, so some settings seem to be provisional – in particular, settings for graphics options, such as shadow and transparency. In the test, for example, changing the menu design led to switches for individual menu groups, some of which were not operational – making the system virtually unusable.

Although you might need to adjust the video and audio settings in some games, the *Driver* submenu lets you modify the loaded kernel modules,

which might be necessary if emulators do not emit sound, for example. Alternative Linux modules are available in the system. Use the arrow keys to adjust the settings; the system accepts the changes immediately (Figure 2).

If you prefer to work on the command line, you can also customize Lakka by editing the `/storage/.config/retroarch/retroarch.cfg` file. When editing manually, note that you need to deactivate RetroArch with the

```
systemctl stop retroarch
```



**Figure 2:** The console emulator settings are made in a separate intuitive menu.

command before starting to work. After completing the configuration, restart the service with the

```
systemctl start retroarch
```

command.

## Control

In the 1980s, game consoles were primarily controlled by joysticks and, later, multifunctional controllers. To recreate this retro gaming experience, the software supports various controllers; according to the website [5], they can be partially addressed over Bluetooth.

In these tests, I used a Sony DualShock 4 controller that usually can be operated over Bluetooth or wired over the USB bus. However, it was not possible to connect the controller to the RPi3 via Bluetooth. Even a test with the model 2 and an external USB Bluetooth dongle by Asus, which is otherwise easily supported by Linux, failed to solve the problem. Therefore the controller could only be used over USB. Lakka automatically detected and initialized the device when connecting the cable.

Controllers have different control elements, so you should check which buttons trigger which functions. Also, some old DOS games developed for traditional PCs did not support joysticks or controllers and were controlled by the keyboard.

## Play with Me!

After all the preparatory work has been completed, you can integrate the desired games, known as ROMs, into the system by transferring them to the `roms/` directory of the SD card. (Please see the “Gray Area” box.) If you emulate several different gaming consoles, you should create a separate subdirectory for each console to store the ZIP archives or ISO images.

To activate the desired console, select the *Load Core* option from the main menu. In addition to hardware emulators for pure consoles, you will find DOSBox for emulating a DOS system. Next, load the desired game by choosing *Load Content | Start Directory | Load Archive*.

The menu disappears and the game screen appears. Use the controller to navigate through the game. Lakka automatically saves scores and configurations, allowing you to resume interrupted games from the History menu at your old score (Figure 3).

Lakka also makes it possible to retrieve ROMs from the network by clicking *Netplay* in the main menu. In a clear-cut dialog, you can define whether to use the Rasp Pi as a host or a client.

## Conclusions

Lakka on the Raspberry Pi offers hours of gaming fun without tying the user to a specific console model. The game pace was positive in the test on the RPi3.

## Gray Area

From a legal point of view, the use of emulators and ROMs on computer systems proves to be problematic. If the emulator contains code from the original console, it may violate copyright laws, so downloading and using it may be illegal. However, many manufacturers release their software for old consoles that are no longer manufactured. Therefore, it is recommended only to buy emulators from reputable websites.

Game ROMs are also subject to copyright. For example, some manufacturers, such as Nintendo, explicitly point out on their websites that free downloading of ROMs is prohibited [6]. However, many games have not been available on the normal gaming market for many years, and often the rights owners no longer exist or have been taken over by other companies.

However, some DOS games ran so fast they could not be played in DOS emulation. In other games, the player immediately lost track of the game in quickly changing complex scenarios. You shouldn't use very old games on high-resolution monitors, because scaling the content creates ugly block effects.

Lakka offers many configuration options, decent hardware support, and easy-to-use menus. The few weaknesses still to be found under the hood are probably a result of the rapid on-going development of the software and are likely to disappear in future releases. ■■■

## Info

- [1] Lakka: <http://www.lakka.tv>
- [2] Hardware guide: <http://www.lakka.tv/doc/Hardware-support/>
- [3] Download Lakka: <http://www.lakka.tv/get/linux/rpi2/>
- [4] Documentation: <http://www.lakka.tv/doc/Home/>
- [5] Input settings: <http://www.lakka.tv/doc/Input-settings/>
- [6] Nintendo rights granted: [https://www.nintendo.com/corp/legal.jsp#download\\_rom](https://www.nintendo.com/corp/legal.jsp#download_rom)



**Figure 3:** The History menu allows you to reload scores after a break.

**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](http://shop.linuxnewmedia.com)



# MakerSpace

Control your Rasp Pi projects with Python's Tkinter library

## The Tinkerer

Use Tkinter to control your Rasp Pi projects from a smartphone or tablet. *By Pete Metcalfe*

**I**f you are looking for an easy way to control your Raspberry Pi projects, consider Python's Tkinter library. In about 20 lines of Python code, you can create a graphic program to control a Rasp Pi rover project (Figure 1). This article shows how to create a Python


Tkinter application to control the Rasp Pi's General Purpose Input Output (GPIO) pins. I will also show you how to set up a desktop shortcut that you can use to easily launch the application, and I will configure a Virtual Network Computing (VNC) server so that a smartphone, tablet, or PC can control the Rasp Pi remotely.

### Python Tkinter Library

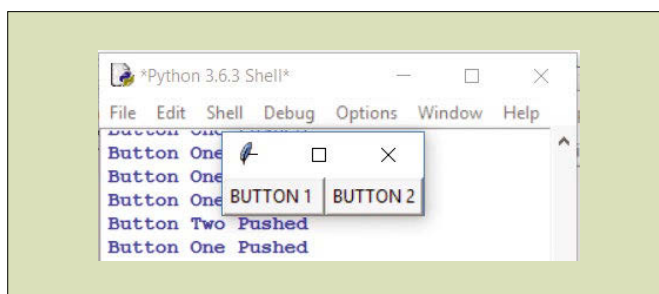
Tkinter [1] is bundled with the standard Rasp Pi Python installation. The base Tkinter library contains 17 different widgets that you can use for graphic applications. These widgets can be defined with customizable colors, sizes, text, and commands. There are three geometry managers (Pack, Grid, and Place) that can be used to arrange the widgets. For Rasp Pi hardware projects, I found the Grid geometry manager was a good fit. Listing 1 is a simple test program that presents two buttons in a grid and prints feedback when the buttons are pressed. (Note: Python 2.7 references the library as Tkinter, whereas Python 3 uses tkinter, lowercase). Figure 2 shows this test program with some button feedback.

### Tkinter Rasp Pi GPIO Motor Example

The Rasp Pi communicates with external hardware through its GPIO pins.



**Figure 1:** Tkinter-controlled rover.



**Figure 2:** Simple Tkinter example with feedback.

**Listing 1: tk\_test.py**

```

01 import Tkinter # Python 2.7 library, for 3.0 use tkinter
02
03 top = Tkinter.Tk()
04
05 def push_one():
06     print ("Button One Pushed")
07
08 def push_two():
09     print ("Button Two Pushed")
10
11 B_ONE = Tkinter.Button(top, text ="BUTTON 1 ", command = push_one).
12     grid(row=1,column=1)
13 B_TWO = Tkinter.Button(top, text ="BUTTON 2", command = push_two).
14     grid(row=1,column=2)
15
16 top.mainloop()

```

These GPIO pins only support 5V connections, so relays or custom *tops* are required to connect to higher-powered devices. Very low-powered motors can be directly connected to Rasp Pi GPIO pins; however, it is highly recommended that you use a Rasp Pi motor top or some other protection hardware.

Different Rasp Pi tops will vary on the GPIO pins used for their motor connections. For our hardware setup, GPIO pins 19 and 21 were connected to the left and right motors. The GPIO pins are accessed in Python using the *RPi.GPIO* library. The *RPi.GPIO* library is preinstalled in the Raspbian image. Listing 2 shows the full code required to drive two motors. For this example, pins 19 and 21 are set as outputs using the `GPIO.setup` call. A `motor` function is created to turn each

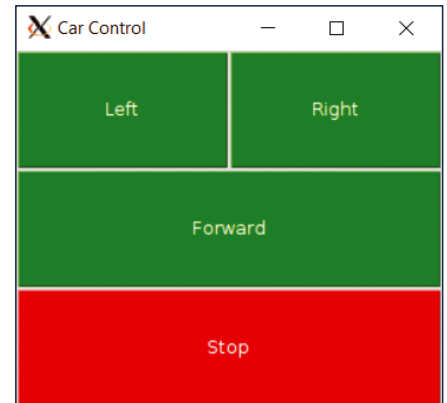
motor on or off. The Tkinter code was enhanced to make the buttons larger and add color. Also, the Python `lambda` function was used for passing the

**Listing 2: tk\_gpio.py**

```

01 import Tkinter # Python 2.7 library, for 3.0 use tkinter
02 import RPi.GPIO as GPIO
03
04 # Setup GPIO motor pins 24 (left) 29 (right)
05 lpin = 19
06 rpin = 21
07 GPIO.setmode(GPIO.BCM)
08 GPIO.setwarnings(False)
09 GPIO.setup(lpin, GPIO.OUT)
10 GPIO.setup(rpin, GPIO.OUT)
11
12 top = Tkinter.Tk()
13 top.title("Car Control")
14
15 #Set pin outs
16 def motor(lstate,rstate):
17     GPIO.output(lpin,lstate)
18     GPIO.output(rpin,rstate)
19     print (lpin,lstate, rpin,rstate)
20
21 B_Left = Tkinter.Button(top, text ="Left", bg = "green", fg = "white", width= 15, height= 5,
22     command = lambda: motor (True,False)).grid(row=1,column=1)
23 B_Right = Tkinter.Button(top, text ="Right", bg = "green", fg = "white", width= 15, height= 5,
24     command = lambda: motor (False,True)).grid(row=1,column=2)
25 B_Forward= Tkinter.Button(top, text ="Forward", bg = "green", fg = "white", width= 33, height= 5,
26     command = lambda: motor (True,True)).grid(row=2,column=1,columnspan=2)
27 B_Stop = Tkinter.Button(top, text ="Stop", bg = "red", fg = "white", width= 33, height= 5,
28     command = lambda: motor (False,False)).grid(row=3,column=1,columnspan=2)
29
30 top.mainloop()

```

**Figure 3: Rover controls.**

motor states. Figure 3 shows the Tkinter GPIO control program.

**Rasp Pi Shortcut**

To make the project easier to access, you can place a quick launch icon or shortcut on the Rasp Pi desktop (Figure 4). To configure a shortcut, add a file to the user's `$HOME/desktop` directory:

```
nano $HOME/desktop/tk_gpio.desktop
```

The name, icon, and execution command are defined inside this file (Listing 3).

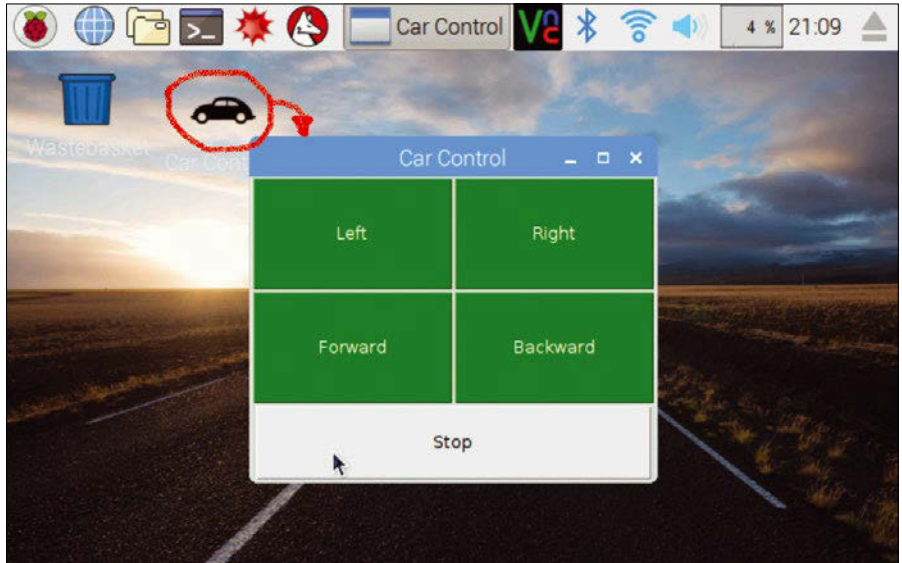
Depending on how you plan to run your program, a shortcut might be enough; however, if you would like to control the project remotely, VNC is a good option.

**VNC**

VNC [2] is a platform-independent desktop sharing system. When the VNC server is enabled on the Rasp Pi, a remote VNC client can view and access the Rasp Pi desktop as if it were connected locally. You can install the VNC client software on Android, Windows, Mac OS, iOS, and Linux devices.

**Listing 3: tk\_gpio.desktop**

```
01 [Desktop Entry]
02 Name=Car Controls
03 Comment=Python Tkinter Car Control Panel
04 Icon=/home/pi/car1.png
05 Exec=python /home/pi/tk_gpio.py
06 Type=Application
07 Terminal=false
08 Categories=None
```

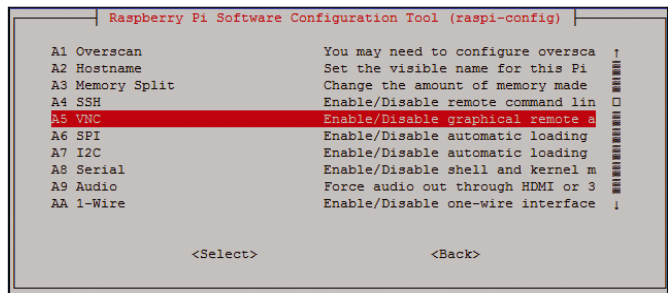


**Figure 4: Placing a shortcut on the Rasp Pi.**

The VNC server is preinstalled on the Raspbian image. To enable the VNC server, enter:

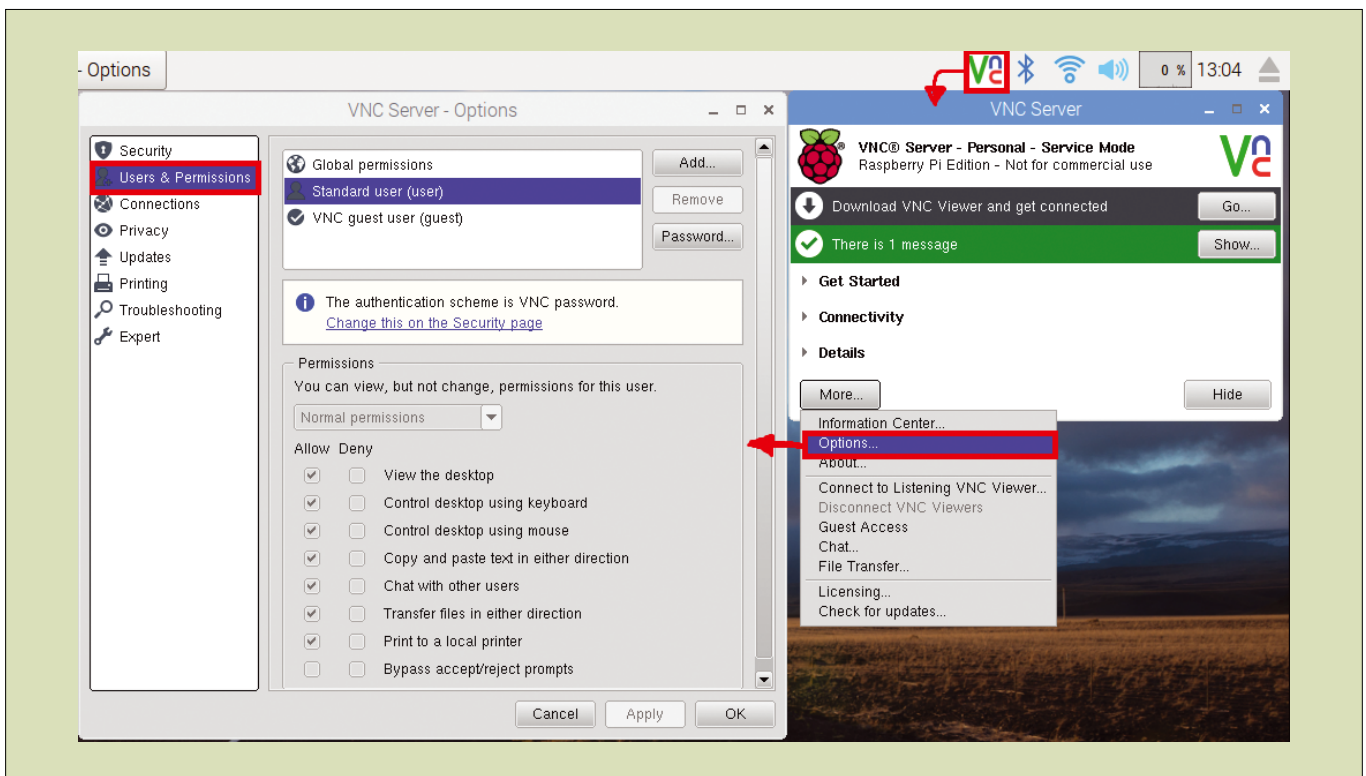
```
sudo raspi-config
```

In raspi-config, select *interfacing option*, and then choose *VNC* and enable (Figure 5).



**Figure 5: Configuring the Rasp Pi for a VNC interface.**

After the VNC server is enabled, you'll need to set up the VNC user security.



**Figure 6: Configuring a VNC user.**





**Figure 7:** Tkinter-controlled rocket launcher.

Click on the VNC tray icon and select the *More* button, and then choose *Options*. In the *Options* dialog, select *Users & Permissions* (Figure 6).

To install the VNC client, go to the RealVNC site [3] and select the required operating system.

When your Rasp Pi boots without an HDMI monitor connected to the desktop, the resolution will be at a low setting (typically 800x600). You can adjust this resolution either in the client setup or on the VNC server in a `~/.vnc/config` file.

**Summary**

Tkinter is a simple way to build controls for your Rasp Pi projects without a lot of code. To enhance your applications, take a look at Tkinter's widgets and add-on libraries, such as Tix [4]. Figure 7 shows an enhancement to the rover project with a missile launcher and graphic buttons for the controls. ■■■

**Info**

- [1] Python Tkinter documentation <https://docs.python.org/3/library/tkinter.html>
- [2] RealVNC Viewer download: <https://www.realvnc.com/en/connect/download/viewer/>
- [3] RealVNC Connect download: <https://www.realvnc.com/en/connect/download/vnc/>
- [4] Tix – Extension widgets for Tkinter: <https://docs.python.org/3/library/tkinter.tix.html>

**Author**

You can investigate more neat projects by Pete Metcalfe and his daughters at <https://funprojects.blog>.

Shop the Shop

[shop.linuxnewmedia.com](http://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](http://shop.linuxnewmedia.com)

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

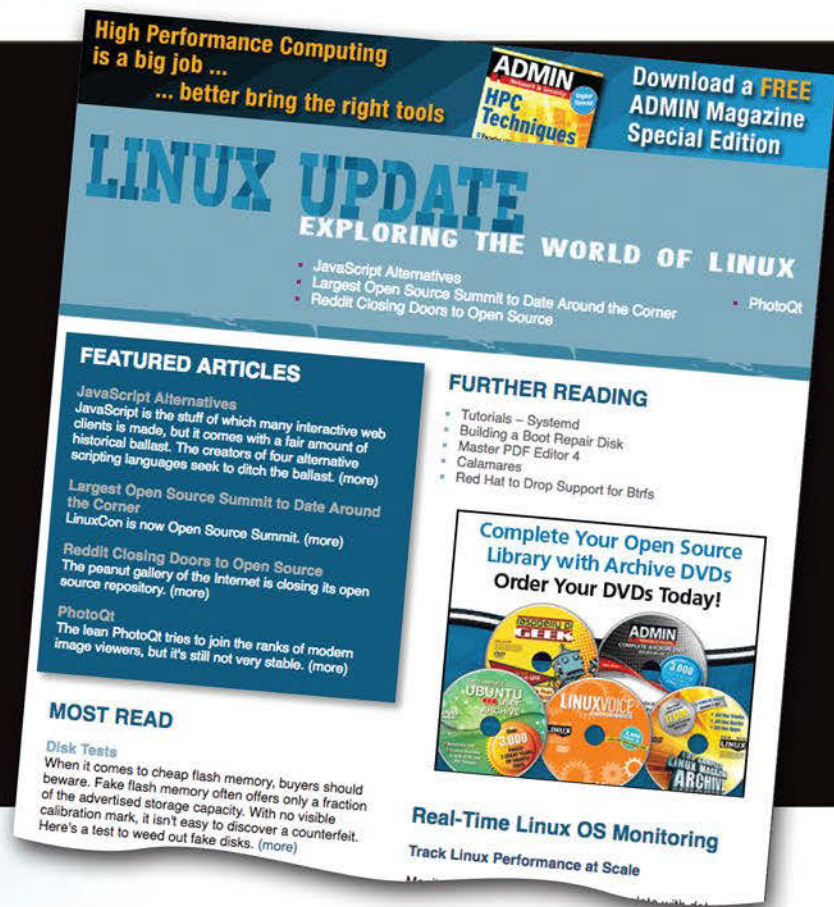


# LINUX UPDATE

## 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](http://www.linuxpromagazine.com/mc/subscribe)

**Linux users are always looking for better tools** to go farther and make life easier, and the good news is, with so many thousands of free tools out there right now, your Linux experience is a never-ending voyage of discovery. LFTP, for instance, is a file transfer alternative that has been around for years, but many Linux users still don't know about it. This month, we take a look at LFTP and show you why many users consider it more versatile and powerful than FTP. We also describe how to organize ideas with the Compendium mind mapping tool, and we show you how to add quirky, over-the-top transitions to your home video creations.



Image © Alexandr Moroz, 123RF.com

# LINUXVOICE ▶

**Doghouse – Chip Replacement** 77

*Jon "maddog" Hall*  
Chip replacement isn't always the answer for hardware problems.

**LFTP** 78

*Ferdinand Thommes*  
LFTP is an alternative to the FTP command set that supports many protocols and offers countless parameters.

**Compendium** 82


*Daniel Tibi*  
Compendium helps bring order to ideas and thoughts and link them to form a complex map from which you can derive arguments and decisions.

**FOSSPicks** 86

*Graham Morrison*  
Graham Morrison looks at KIT Scenarist, Inboxer 1.0, Tungsten, rtl\_433, Oil 0.3.0, Etcher 1.3.0, and more!


**Tutorial – Video Editing** 92

*Paul Brown*  
Most video editors supply you with a generic catalog of transitions, usually in the shape of tired wipes and fades. But what if you wanted something a little more special?



# 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](http://shop.linuxnewmedia.com/specials)

# MADDOG'S DOGHOUSE



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

Chip replacement isn't always the answer. BY JON "MADDOG" HALL

## When chips go bad

**L**ast month, I wrote about the Meltdown and Specter issues. The vendors and the community are working hard to evaluate the whole problem and see what types of solutions exist. Many people are demanding that the chips be replaced. While this would be wonderful, in reality it is unlikely to ever happen, for many reasons.

This problem has been going on a long time, over many iterations of chips. The fabrication plants for creating the older chips have started making more modern chips. Going back in time to replace these older chips would be difficult, and even if the chip makers could do that, many of the chips are soldered to the motherboard. Removing the chip from the motherboard and resoldering it would be very expensive. Even if replacing the chips could be done, unless the chip is an exact replacement, the motherboard circuitry could be different, and the instructions to boot and run the operating system would have to change.

Perhaps it is a more modern chip, manufactured only a short time ago. The customer would have to identify the chip, find out if it is indeed affected, and submit that chip to the vendor to replace. Many of the same issues would exist, although the operating system might still be adaptable to the newer chip.

### Another Time and Place

VAXstation 3100 computer systems were desktop computers designed by Digital Equipment Corporation (DEC) around 1986, one of DEC's first forays into "high volume" manufacture, which then measured in hundreds of thousands, not hundreds of millions.

DEC had a whole team of engineers on this product. Hardware circuitry, case design, manufacturing engineers, and of course both VMS and Ultrix (DEC's Unix product) engineers. We met every week in a project meeting.

This product utilized daughter cards for the RAM memory, way before the concept of industry standard SIMMs of memory. DEC had used memory chips from two different Japanese memory manufacturers. Not wanting to be too "cutting edge," DEC used memory chips that were shipping in the millions of units all over the world to many manufacturers, including DEC. The VAXstation 3100 was also the first system that DEC created that had only parity error detection, not ECC correctable memory.

The systems were headed toward field test. Because this was such an important product, there were a *huge* number of units

made, not the typical 20-30 field test units that a larger system might have generated.

As the units were being readied to field test, it was noticed that Ultrix kept crashing while VMS was rock solid. The hardware engineers made comments about how terrible the Ultrix code was.

We eventually proved that the memory chips from one of the Japanese memory vendors was "forgetting," even though it was being properly clocked and fed the right amount of electricity – something that was *not* supposed to happen. If you did not read or write to that company's memory chip for over 45 seconds, it "forgot." Not all the time ... only about half the time.

Wait! Why only Ultrix and not VMS? Because VMS always read into memory from disk when it was swapping in a swapped-out process. This was "refreshed" RAM that had not been written to for a long period. Unix (and Ultrix) realized that if the RAM had not been changed, what was the sense of "refreshing" it, since it still contained the information it had before? Right?

We brought in the Japanese firm and proved to them what was happening. That was when I learned the Japanese words that could not be printed here.

Wait a minute! Why had this not shown up before in those millions and millions of other systems with that vendor's chips?

Most servers of the day had ECC correctable memory. If one chip forgot, the others recreated the missing bit, called a "soft hit," something only noticeable if you looked for it in the system log.

PCs (which only had parity memory detection, not ECC correction) usually ran with small, active memories, so the chips were always being accessed. And if the system did crash, it was probably due to that crappy Microsoft operating system, right?

The memory company begged us not to announce this. DEC decided not to announce it, since it literally would bankrupt the memory company, which would have solved nothing. DEC did insist they buy back every chip they had ever sold us for the price we paid, and when the fix was done, replace the chips at current market price. The difference in price allowed DEC to replace every chip in every board owned by DEC's customers (including manufacture and field service on-site replacement) and still make twelve million dollars of profit on the replacement.

DEC never told which company had the bad chips, and I never will either. ■■■

# Enhanced

LFTP is an alternative to the FTP command set, which supports many protocols and offers countless parameters. **BY FERDINAND THOMMES**

**A**lthough pretty much outdated, the File Transfer Protocol (FTP) still plays a significant role. For 20 years, LFTP has offered a greatly expanded command set for the command line that handles secure transmissions, without being excessively difficult to handle.

FTP dates back to 1985. Designed for transferring files between two computers, FTP is a relic from the infancy of modern IT. Today it has lost much of its former importance, not least because of some serious security problems: It neither encrypts the FTP username and password, nor protects the data against manipulation during transfer.

FTP via SSL (FTPS) or SSH (SFTP) counteracts these weaknesses. The two methods are considered to be equally secure, and they encrypt both data and metadata during transmission. In most cases, SFTP proves to be more flexible and much simpler in practice.

## Alternatives

FTP alternatives have been around for a long time – on Linux, the venerable secure copy (`scp`) has been used to transfer data to other machines from the command line. In the meantime, WebDAV (Web Distributed Authoring and Versioning), a kind of extension of HTTP with new commands, has also developed into a useful replacement.

Nevertheless, FTP is a marathon runner for various reasons, most often because many web hosts only offer FTP and leave WebDAV or SSH out of the equation. FTP is thus likely to be around for a while.

LFTP is a very powerful extension of the original FTP command set that has been under development since 1997. It is developed by Alexander Lukyanov, is licensed under the GNU Public License (GPL), and was last updated to version 4.8.0 in July 2017 [1]. LFTP supports a wide range of protocols, including plain vanilla FTP, along with FTPS, SFTP, HTTP, HTTPS, HFTP, FXP (see the “FXP” box), Fish, and BitTorrent.

## Special Features

LFTP stands out from most other command-line FTP clients with advanced features such as recursive mirroring and the ability to update entire directory trees and manage bookmarks. Several simultaneous sessions to different servers are implemented by assigning commands to *slots* in a separate shell. Additionally, scheduled transfers can be queued and scheduled, and bandwidth use can be managed. If a transmission breaks down, LFTP resumes it later on at the same location. If you exit LFTP, it continues to process the active tasks in the background.

## Client-Server Model

LFTP can be found in the repositories of most distributions; thus, you can usually install it with the package manager.

To test LFTP without an active server, it is best to install an FTP server on a second local machine, to which LFTP then establishes a connection. For example, the very secure FTP daemon (`vsftpd`) [2] is a good way to go, because it does not require any configuration for simple tests. Before you can start, all you have to do is remove the `#` comment sign in front of the `#write_enable=YES` line in `/etc/vsftpd.conf`.

Last but not least, create a directory on the client computer for the test and fill it with some data.

## Which LFTP?

If you want to use LFTP with FTPS, you might need to compile the package with SSL support

### FXP

FXP, which stands for File Exchange Protocol, is a method defined in FTP, but not named there, for controlling file transfers between two servers on a client. The files move directly from server to server, saving time. Many current FTP clients support FXP.

yourself. Use the command line to check whether this is necessary by entering the

```
which lftp
```

command to determine where LFTP has been installed. After calling

```
ldd /<path>/<to>/lftp
```

check its output to see whether one of the lines starts with *libssl* (Figure 1); for example:

```
$ ldd /usr/local/bin/lftp | grep ssl
```

If this is not the case, and you want to use FTPS instead of SFTP, follow the instructions in the “DIY LFTP” box.

At system startup, LFTP will search the */etc/lftp.conf*, *~/.lftprc*, and *~/.lftp/rc* files (in that order) for configuration. Directives that you create in the home directory take precedence over the global configuration in the */etc* directory. The instructions always start with *set* and follow the

```
set <Option>:<Switch>[on|off]
```

scheme.

## Practice

For test purposes, it is best to use two local computers that serve as the server and client. LFTP will behave in the same way as if you were using remote servers, but it will be easier to locate sources of problems and errors. Unless you specify otherwise, LFTP uses the SFTP protocol. Ports other than standard port 21 are entered at the end of the connection command in the form *:<port>*.

For more feedback, add the *-d* (debug information) option to the *lftp* call. The *-h* (*--help*) parameter offers little help when it comes to the application's options. An intelligible version of the LFTP man page can be found on the Lifewire website [3].

## DIY LFTP

On Debian, you need to install the following dependencies:

- *libgnutls30*
- *libgcc1*
- *libreadline-dev*
- *libssl-dev*
- *libtinfo5*
- *libunwind8*
- *libghc-zlib-dev*

Depending on the version of the system you are using, you might need other versions of these packages. If further software is missing, entering *./configure* will usually let you know. Build LFTP as shown below with the usual three-card trick – *./configure*, *make*, and *make install* – and use *ldd* to test your home-grown LFTP version for SSL support.

```
$ wget http://lftp.yar.ru/ftp/
lftp-4.8.0.tar.gz

$ tar xzf lftp-4.8.0.tar.gz

$ cd lftp-4.8.0

$ ./configure --with-openssl=/usr/lib

$ make

$ sudo make install

$ ldd /usr/local/bin/lftp | grep ssl

libssl.so.1.0.0 => /lib/x86_64-linux-gnu/
libssl.so.1.0.0 (0x00007f8e89dbd000)
```

```
test@test-VirtualBox:~$ ldd /usr/bin/lftp
linux-vdso.so.1 => (0x00007ffef79c6000)
libgnutls.so.30 => /usr/lib/x86_64-linux-gnu/libgnutls.so.30 (0x00007fd2b2aec000)
libz.so.1 => /lib/x86_64-linux-gnu/libz.so.1 (0x00007fd2b28d0000)
libreadline.so.7 => /lib/x86_64-linux-gnu/libreadline.so.7 (0x00007fd2b2686000)
libutil.so.1 => /lib/x86_64-linux-gnu/libutil.so.1 (0x00007fd2b2483000)
libtinfo.so.5 => /lib/x86_64-linux-gnu/libtinfo.so.5 (0x00007fd2b225a000)
libdl.so.2 => /lib/x86_64-linux-gnu/libdl.so.2 (0x00007fd2b2054000)
libidn.so.11 => /lib/x86_64-linux-gnu/libidn.so.11 (0x00007fd2b1e21000)
libstdc++.so.6 => /usr/lib/x86_64-linux-gnu/libstdc++.so.6 (0x00007fd2b1a99000)
libc.so.6 => /lib/x86_64-linux-gnu/libc.so.6 (0x00007fd2b16d2000)
libp11-kit.so.0 => /usr/lib/x86_64-linux-gnu/libp11-kit.so.0 (0x00007fd2b146d000)
libtasn1.so.6 => /usr/lib/x86_64-linux-gnu/libtasn1.so.6 (0x00007fd2b125a000)
libnettle.so.6 => /usr/lib/x86_64-linux-gnu/libnettle.so.6 (0x00007fd2b1022000)
libgmp.so.10 => /usr/lib/x86_64-linux-gnu/libgmp.so.10 (0x00007fd2b0da2000)
libhogweed.so.4 => /usr/lib/x86_64-linux-gnu/libhogweed.so.4 (0x00007fd2b0b6f000)
/lib64/ld-linux-x86-64.so.2 (0x00005566f1ffc000)
libm.so.6 => /lib/x86_64-linux-gnu/libm.so.6 (0x00007fd2b0866000)
libgcc_s.so.1 => /lib/x86_64-linux-gnu/libgcc_s.so.1 (0x00007fd2b064f000)
libffi.so.6 => /usr/lib/x86_64-linux-gnu/libffi.so.6 (0x00007fd2b0447000)
libpthread.so.0 => /lib/x86_64-linux-gnu/libpthread.so.0 (0x00007fd2b0227000)
test@test-VirtualBox:~$ ldd /usr/local/bin/lftp | grep ssl
libssl.so.1.0.0 => /lib/x86_64-linux-gnu/libssl.so.1.0.0 (0x00007f8dd2dba000)
test@test-VirtualBox:~$
```

Figure 1: Using FTPS will not work without *libssl* support.

To get used to the LFTP shell, first download or upload files from the server. To do this, log on to the server with the `lftp <user>@[<FQDN> | <IP>]` command; for example:

```
lftp ft@192.168.178.123
```

After the connection is established, you are in the LFTP shell.

The `ls` command lists the current directory content on the server. With a preceding exclamation mark (!), you can redirect the command to the client; `!ls` shows the contents of the client's current local working directory (Figure 2). Some commands can be executed with a preceding small `l` (ell, for "local") on the client, and other commands can be preceded by either switch. Usage is not completely consistent. To simplify things, you can define command aliases to suit your own requirements.

### Up and Down

To upload a file from the client to the server, use the `put` command, to which you add the name of the file in question (`put <file>`). The `-d` parameter used for logging in shows you what is happening in the background in detail. If you want to put the file in a certain directory, go there with the `cd` command and then run `put`.

Following the same pattern, `get` lets you download a file. However, with `pget`, LFTP offers an alternative that distributes the data transfer over several parallel connections, thus speeding up the download on servers with slow connections. In practice, you would first change to the directory in which you want to store the downloaded file `!cd` or `lcd` followed by `get <file>` or `pget <file>`.

### Mirroring and Synchronizing

You can use the `mirror` command to move a folder or an entire directory tree to a server. Use `lpwd` to make sure that you are in the correct local directory. If necessary, change this with `lcd`. You can do the same with `pwd` and, if necessary, `cd` on the server side. The command to upload the directory is simply `mirror -R`, the reverse path from the server to the client works with `mirror`.

Alternatively, you can specify the directories with the

```
mirror -R /<Server folder>/ ↵
      /<Client path>/
```

command. After the comparison, if you change a file or the file repository in the source directory and then execute the `mirror` command again, LFTP simply syncs the changes with the other side. The `mirror` command also supports several parallel connections, which you can initiate with `mirror -P` or define more precisely with:

```
mirror -parallel=<n>
```

LFTP plays to one of its strengths in the case of large volume transmissions, in particular by displaying a progress bar and the transmission speed, which you can reduce if required. If a transmission fails once, it can be resumed with the `-c` parameter, even if you terminated LFTP in the meantime.

Likewise, LFTP can handle several transmissions at the same time, and they can be sent off into the background as desired and then moved into the foreground again. To view details of the actions performed, take a look at the `~/lftprc/` directory, where logs of all sessions that document the actions in detail reside.

When all work is complete, exit LFTP with `exit` or `bye` and return to the shell command line.

When all work is complete, exit LFTP with `exit` or `bye` and return to the shell command line.

### Advanced Functions

In addition to transferring files and directories between client and server, LFTP has many additional features and hundreds of control parameters. The most important commands are briefly summarized in Table 1.

For example, you can use the `at <time>` time control transmission parameter; `queue` lets you

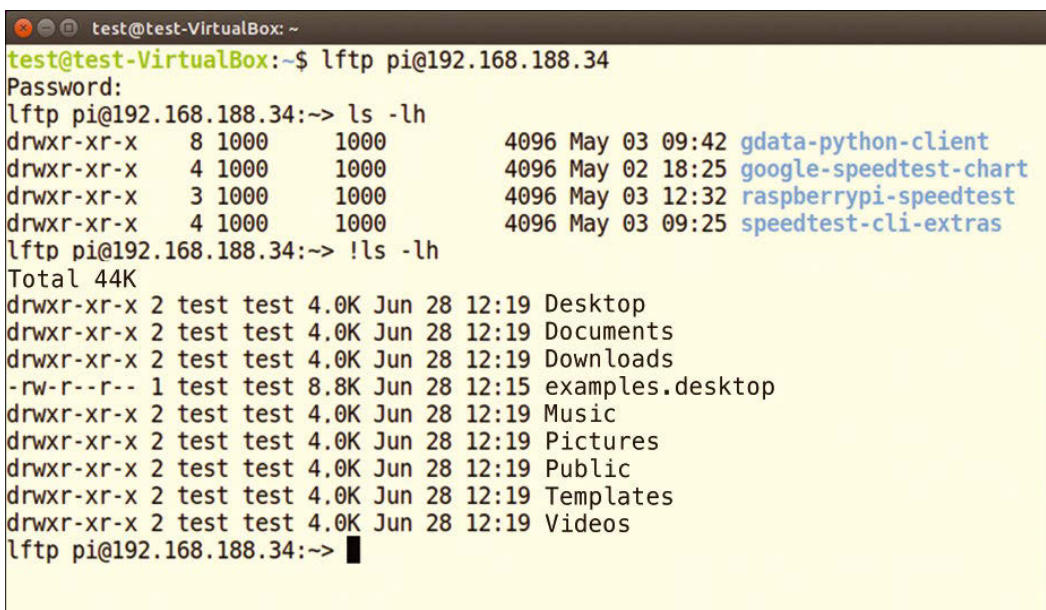


Figure 2: Bash commands with preceding exclamation marks act on the client.



**Table 1: Command Reference**

Command	Function
<b>File Transfers</b>	
<code>put &lt;File&gt;</code>	Transfer file from client to server
<code>get &lt;File&gt;</code>	Transfer file from server to client
<b>Sync Entire Directories</b>	
<code>mirror -R</code>	Sync directory from client to server
<code>mirror</code>	Sync directory from server to client
<b>Bookmarks</b>	
<code>bookmark add &lt;Name&gt;</code>	Create bookmarks
<code>bookmark list</code>	List bookmarks
<code>open &lt;Name&gt;</code>	Open bookmarks
<b>Time Control and Batch Processing</b>	
<code>at &lt;hh:mm&gt; &lt;Command&gt;</code>	Specify time
<code>queue mirror &lt;Folder&gt;</code>	Create queue (directory)
<code>queue mirror &lt;File&gt;</code>	Create queue (just one file)
<code>queue start</code>	Start queue

create queues; and `-n <number>` divides the file transmission into several parts. If you want to repeat an action later, create a bookmark with `bookmark add <name>`, and `bookmark list` lets you view all the bookmarks you have created.

Generally, you can initiate parallel transfers with `-P (--parallel)=<n>`. LFTP supports FXP transfers from server to server that excludes the client. You can download web pages with HTTP – even the Torrent protocol is supported. LFTP can also be controlled with scripts.

### Conclusions

The very powerful and versatile LFTP is comparatively easy to master. Many commands are the same or similar to those of the shell or the conventional FTP command set. The application is

still very much active, even after 20 years of development, as can be seen by the release of the fourth update in 2017.


LFTP, on the other hand, lacks up-to-date documentation. Many of the instructions found online are no longer correct. Nevertheless, if you appreciate the command line and want to or need to use FTP, I strongly recommend LFTP as a more powerful alternative. ■■■

### Info

- [1] LFTP: <https://lftp.yar.ru>
- [2] vsftpd: <https://security.appspot.com/vsftpd.html>
- [3] LFTP command reference: <https://www.lifewire.com/lftp-linux-command-4093434>


Shop the Shop

[shop.linuxnewmedia.com](http://shop.linuxnewmedia.com)



# GIMP HANDBOOK

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



Order online: [shop.linuxnewmedia.com/specials](http://shop.linuxnewmedia.com/specials)



FOR WINDOWS, MAC OS, AND LINUX USERS!

# Mind Map

Compendium helps bring order to ideas and thoughts and link them to form a complex map from which you can derive arguments and decisions.

BY DANIEL TIBI

Collecting ideas, thoughts, and information is important in a wide variety of fields. Visualizing this process enables you to identify networks, which makes it easier to evaluate and use the data. The English author Tony Buzan developed the principle of mind mapping for this purpose.

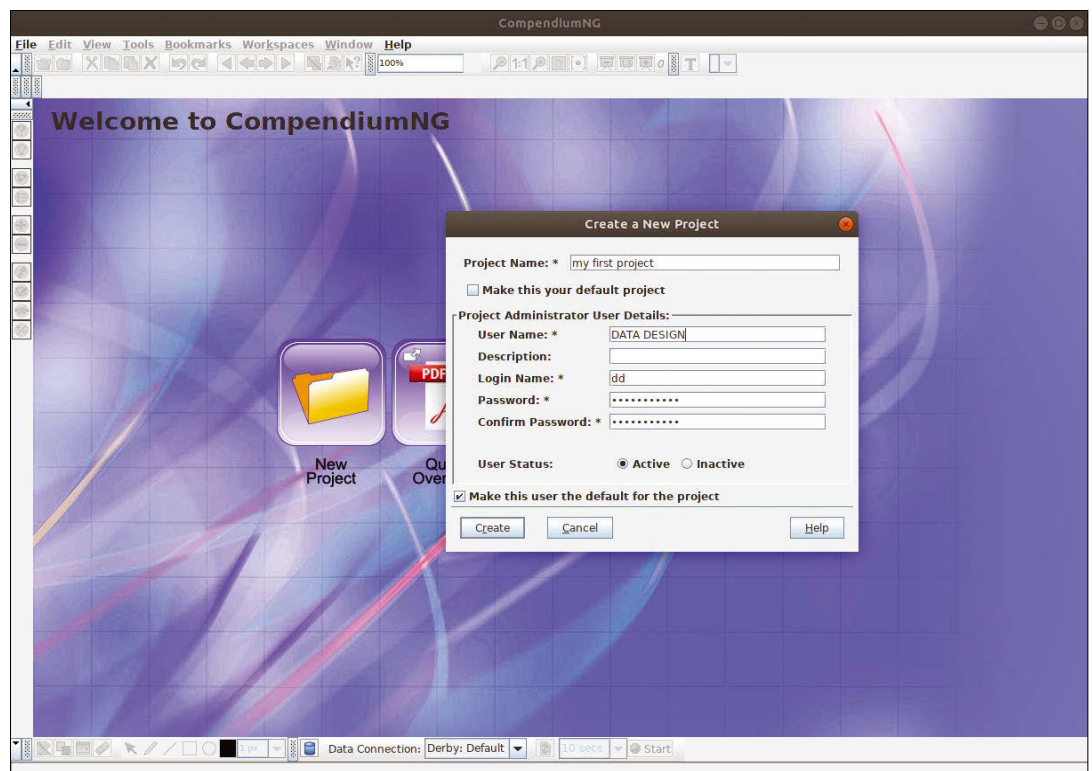
Compendium [1] is a powerful program for creating mind maps on the computer. The program doesn't just replace a notepad: At seminars in combination with a projector, you can create collaborative mind maps that you can then email directly to all participants.

Compendium is not included in the package sources of common distributions, so you need to download the current version from GitHub [2]. The program, written in Java, runs on any

platform but requires a current Java Runtime Environment [3] on the system. To install, run the file you downloaded. The program is licensed under the GNU Lesser General Public License (LGPL), which you accept in the first screen. Next, select a target directory for the installation, ideally `/opt/compendium/`.

Now, you need to decide whether the installer should create an additional launcher (i.e., a corresponding script for Linux and Mac OS X or a batch file for Windows). In this way, if you install the software on a USB stick for use on different computers with different operating systems, you will always have the correct launcher at hand.

Before the installation begins, you can decide whether you want to create a shortcut in the Start menu. Just in case you want to remove



**Figure 1:** If necessary, use the documentation provided, which comprises text files and several videos to get you started with Compendium. Alternatively, you can launch directly into a new project.

the program later, the software also creates the `Uninstaller/` subdirectory in the selected target directory, where you will find the appropriate program for the uninstall. Complete the installation by clicking on the *Done* button. The software is now ready for use.

### Creating a Project

After the installation, you can either start Compendium by clicking on the shortcut or by using the

```
sh /<path to>/CompendiumNG.sh
```

command from the terminal. The first time you start the program, a small window appears in which you can call up instructions in a PDF file or watch movies with an introduction. Click on the *New Project* button to get started. In the dialog box that follows, enter a name for your project and specify a username and password (Figure 1).

The next time the program is launched, Compendium displays a list of all projects on the computer. Select the one you want to continue working on and log in with its username and password. To manage users, open a dialog with *Tools | User Manager*. Use the *Users* tab to add users to the project.

If you work alone, use the options in the *Single User System Options* tab instead, and activate *Set current user as default for this project*. If the current project is the only one, also enable *Set current database as default project for Compendium*.

### Nodes

Compendium stores information in nodes. Each new project already contains two such nodes: an Inbox and a Trash Bin. Users working together on a project can send each other messages within the program that then appear in the recipient's Inbox. Even if you're working alone on a project and don't use the Inbox, you cannot remove it from the desktop.

Node points and the information stored in them are first sent to the Trash when they are deleted. Double-clicking the icon opens the Trash Bin and displays all the information it contains. You then have the option of either restoring the information or deleting it permanently.

The software supports 10 types of nodes: Question, Answer, Map, List, Pro, Con, Reference, Note, Decision, and Argument. You can create a new node by right-clicking where you want to insert it. In the context menu, choose *Create Node* for the appropriate type and assign an intuitive name.

This basic principle applies to all nodes. The corresponding symbol with the name you assigned appears in the workspace. Double-clicking the icon opens a new dialog box in which you

can save information. Depending on the type of node, the software displays and processes the data differently.

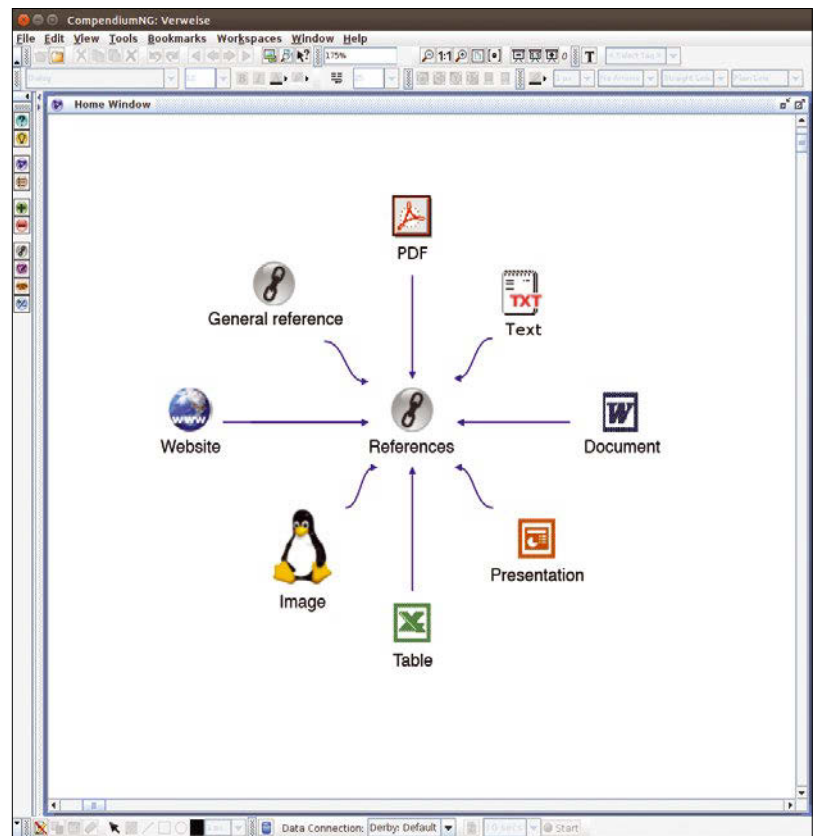
The node types Question, Answer, Pro, Con, Note, Argument, and Decision contain text. They only differ in appearance by their different icons, but they fulfill different purposes within the mind map.

The type of node can change during the development process of a map: For example, an Answer can turn into an Argument and later into a Decision. An asterisk at the top left of the node indicates that the node contains text. If you mouse over the asterisk, the program displays this text; double-clicking the icon lets you edit the text.

The Reference node type is used to include external information. After double-clicking, a dialog opens containing another text field for an external link, in addition to the usual text field. Enter a shortcut to a file on your computer or an address on the Internet.

For a file, you can decide whether the shortcut refers to the location of the file in the filesystem or whether Compendium stores a copy of the file in the project folder, which is useful if you use different computers to work on a project.

If you have a web page or a file format that Compendium recognizes, the link is automatically



**Figure 2:** In Compendium, you can include links to external content. If the program recognizes the linked file format, it displays a suitable icon and opens the external content in the correct program.

assigned an appropriate icon and double-clicking it opens the page or the file in the associated program (Figure 2). By clicking *Tools | Linked Files Browser*, you can see a list of all linked files and the locations in the filesystem.

The Map and List nodes represent two different ways of displaying information. A Map is a mind map within a mind map, in which you arrange information on an area using nodes and connections as desired. However, the software displays the information as subitems in a list and offers you the option of sorting them according to certain criteria.

You can move and arrange nodes in the workspace as required. By right-clicking on a node, you can change its type with the *Change Type To* menu option (e.g., change it from Answer to Argument and then to Decision).

### Networking

Compendium does much more than simply gather information in nodes: It lets you network information, providing an overview of interrelationships and facilitating evaluation. Networking can be created by means of graphical connections such as lines and arrows. Alternatively, you can use keywords to create context.

Graphical connections in the form of lines and arrows visualize relationships between nodes. Thus, an arrow from a study Reference to an Argument indicates that the paper sup-

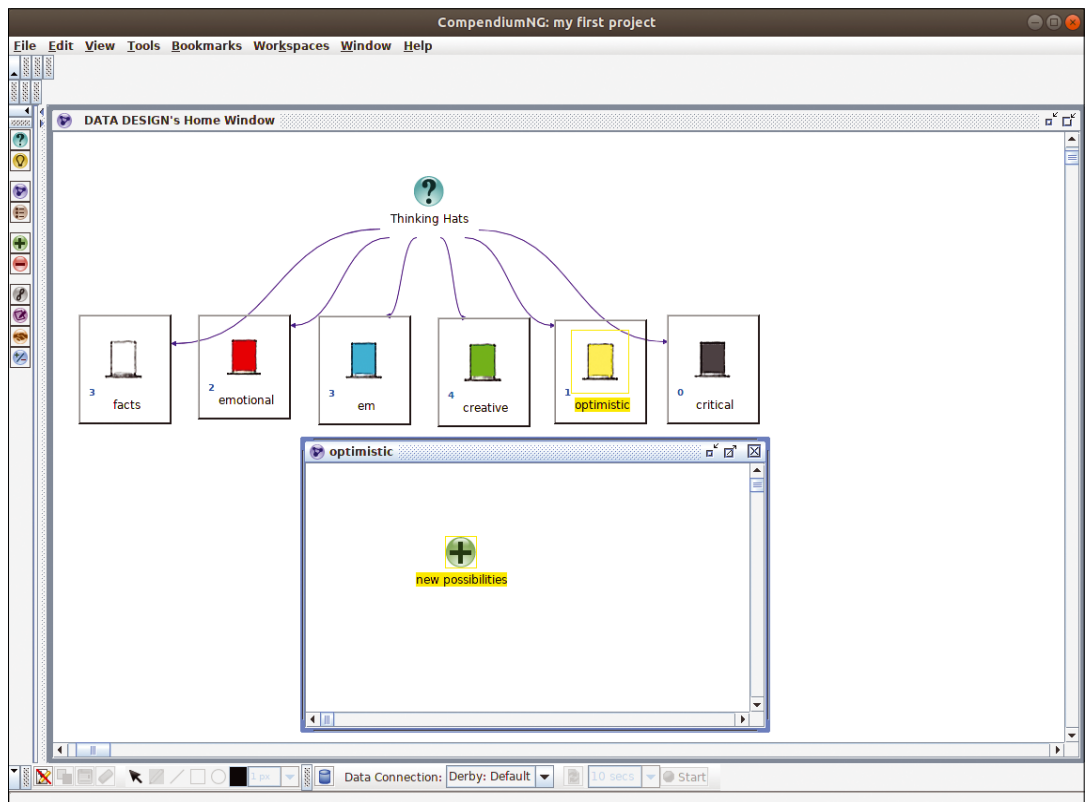
### Six Thinking Hats

The Six Thinking Hats (6TH) technique was developed by Edward de Bono. Using this method, you assign information and arguments to six hats, each a different color and each of which offers a different approach to the problem. In this way, you ensure that the process of decision making takes into account different approaches to a problem by allowing individuals or members of a discussion group to put on different hats and present and evaluate arguments to match the hat.

ports the Argument. From the Argument, an arrow points to a Decision indicating that the Argument had an effect on it.

When you mouse over an existing node, you can create a graphical connection. A small arrow appears on each of the four sides of the node. A left-click on one of these arrows opens a selection menu listing the types of nodes. From this list, you select the kind of node you want to create and connect to the current node by dragging a line. The new node appears on the desktop, connected to the starting point by a line.

The line first appears as an unlabeled arrow from the new to the old node. You can adjust the



**Figure 3:** In Compendium, you can combine the 6TH method with a mind map. Each of the six thinking hats has its own mind map. The number on the left side of each node indicates the number of nodes in the thinking hat's mind map.

appearance of the line by right-clicking on it and selecting the *Contents* menu item. A dialog box appears with a text field in which you can enter a description for the line, if desired. In the same dialog, you can select a color for the line from a list, where a certain color implies a certain meaning (e.g., *Argument for*, *Argument against*, *Reply to*, *Report against*, *Further explanation*, etc.). This association of colors with meanings simplifies the overview.

Later, you can change the text of a line by simply clicking on it and overwriting, and you can change the meaning of a line and its color assignment by right-clicking on the line. To assign a new meaning and color, select *Change Type To* and make the change. To adjust the direction of the arrows, right-click and select the *Change Arrows To* item.

### Keywords

Keywords are used to link nodes that have something in common. By clicking on the icon with the letter *T* in the upper toolbar, you can open the keywords view in the left margin. There, the software lists all entries, and you can create new entries or delete existing ones.

To add a keyword to a node, click on it and select the keywords you want to assign from the list by checking the boxes. The letter *T* now appears at top left in the corresponding node, indicating that keywords exist for this node.

### Searching and Evaluating

The *Edit | Search* menu option lets you call the program's search function and then proceeds to search for node names, content, and keywords, among other things. The Map and List node types let you evaluate such a search.

To do so, you first create a new node of the Map or List type and keep the window open with the empty contents of this new node before running the desired search. As a result, the software provides a list of all nodes that meet the search criteria.

From this list, select all the nodes that you want to insert into the new Map or List and save them using *Insert into View*. In this way, the same node can appear in different views and thus in different contexts, which in turn makes it possible to highlight it from different angles.

For example, a node with a new idea can appear together with a positive evaluation in a Map with new options, but also in a Map with risks, where it is negatively evaluated. Changes you make to a node in a view always affect all other views as well.

### Thinking Hats

Compendium was developed to create mind maps, but it is also suitable for the Six Thinking

Hats technique (see the "Six Thinking Hats" box). As an example, create a new project, in which the focus will be on a specific question about an appropriate node. Starting from this one node, now create six new nodes, one for each thinking hat. Next, you convert each of these six nodes into a thinking hat by right-clicking on a node and selecting from the context menu *Change Type To | Stencils | De Bono 6 Hats Thinking*.

Thinking hats act as mind maps within a mind map (Figure 3). Double-clicking on a thinking hat node opens a new workspace in which you can develop a mind map based on the approach of the respective thinking hat.

### Sharing Mind Maps

Sharing a jointly created mind map with all participants is easy by email. The receivers do not need to have Compendium installed on their computers: If required, you can export mind maps into various formats, including JPEG and HTML, with the usual *File | Export*.

If you choose the JPEG format, the software simply generates a graphic with the appearance of the main window. It only maps the individual nodes and their connections to each other, but not the contents of the nodes. Exporting to HTML format does more: The hypertext page contains the contents of the nodes, which you can display by clicking on the corresponding anchor.

### Conclusions

Mind maps are an excellent way to connect and interlink information to give you an overview and simplify your evaluation process. Compendium is a good tool for the job. It collects and combines information structured as a classic mind map, adds keywords, searches the information on the basis of various criteria, and focuses the information more sharply. The ability to export mind maps to HTML makes them easy to share. Because Compendium is platform-independent, it opens further possibilities by letting you use the program on different computers and work in groups without any compatibility problems. ■■■

### Info

- [1] Compendium:  
<https://github.com/CompendiumNG/CompendiumNG/wiki>
- [2] Installation file:  
<https://github.com/CompendiumNG/CompendiumNG/wiki/CompendiumNG-installer>
- [3] Java Runtime Environment:  
<http://www.oracle.com/technetwork/java/javase/downloads/jre8-downloads-2133155.html>

# 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**

Screenwriter Studio

## KIT Scenarist

While there are plenty of applications, and plenty of additions to Vim and Emacs, that help you write a novel, not many can help you write a screenplay for theater, video, television, or films. This may be due to screenplays being more specific, more niche, and perhaps a little lower down on the spectrum of dream careers for people who can follow one word with another. But they're also (theoretically) easier to write than a novel and offer another potential avenue for people wanting to find a path into professional writing. With

the odds stacked against authors getting their work published, additional avenues are always worth pursuing.

KIT Scenarist is a Russian application written for screenwriters. It gets off to a great start by offering to run with a preference chosen from 12 different languages, and either a light or dark theme. You then get to choose which modules to enable by default, such as the research module for collecting your background information, the cards module to help you build a script from a series of ideas written on

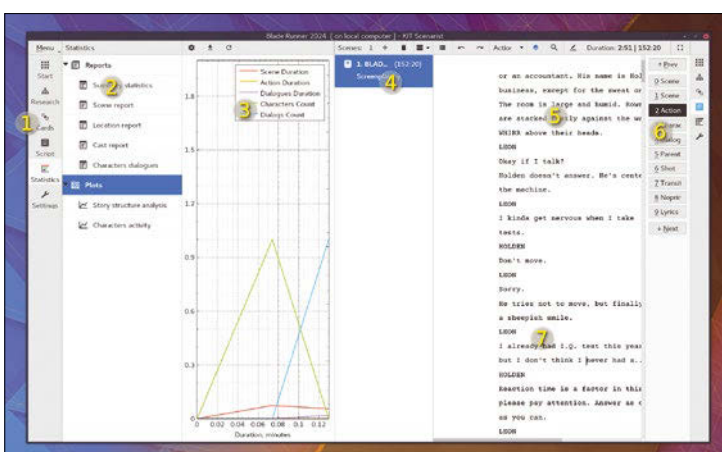
help you charge per brilliant word, the script module itself with support for scenes and standard scriptwriter tools, and, finally, a starter template, such as the "Final Draft Screenplay" template. However, you can just as easily create your own templates or apply a template to an already worked-on script to change its layout and structure.

The application feels more like an IDE for scriptwriters than an augmented text editor. When you start a project, for example, you're guided through creating characters and adding background details for them, creating scenes, adding dialogue, and pulling the whole thing together into a script that can easily be exported as FDX, PDF, Fountain, and DOCX formats. For writers, simple things such as dynamic word count are seldom implemented, but Scenarist gets this right, as it does the automatic text formatting, which gets out of the way and lets you write your dialogue or your descriptions without having to worry about syntax. You can easily see the plot status, organize the threads of your story, and change things around without having to resort to a pair of scissors and randomly sequenced bits of paper.

There's even a review mode where editors can leave comments, and your research can include external directories and images, all selectable from a module tab that can be opened alongside the editor if you opt for the two-panel mode. Scripts have a very specific layout and format, and KIT Scenarist can generate this output from your content without you needing to know the specifics of how to generate the exact formatting. This also includes essential details such as the timing for each line of dialogue, reflected in the chart for the statistics overview for the entire script. Scenarist is remarkably mature and capable. The Qt user interface makes it feel like a native Linux application, and it's easy to see why there are more than 1,000 script writers using the application already.

**Project Website**

<https://kitscenarist.ru/en/index.html>



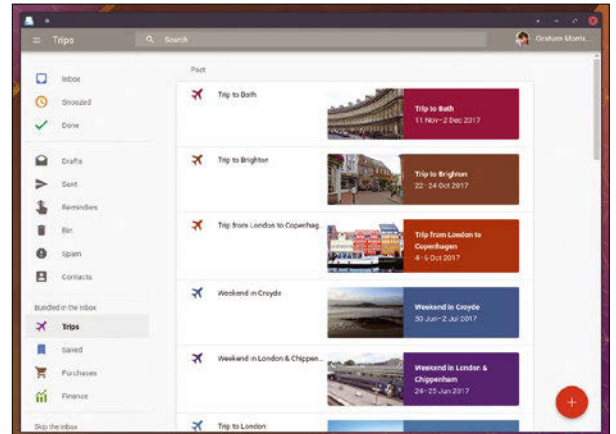
1. Edit modes: Scenarist covers each step of the script-writing process, from research to printing. 2. Reports: Characters, scenes, and locations can all be described separately. 3. Graphs: See how you hit timing targets for production planning. 4. Screenplay: Scenarist is also a serious editing tool, splitting your script into scenes and sections. 5. Editor: Add timing information, and import character and scene details from your overviews. 6. Toolbar: All the elements that make up a screenplay can be dragged into and written easily. 7. Review mode: Generate a PDF script or collaborate with other editors.

## Gmail client

# Inboxer 1.0

**T**hose of us who have foolishly allowed Google to pervade most aspects of our online lives often feel we're missing out when it comes to desktop or standalone applications. This is why you see third-party Drive utilities, for example, or communication clients including Hangout support. This is a little ironic considering Google likely wouldn't have had the same success without its foundation of Linux and open source, but Google still does plenty of proactive things for Linux and open source in other areas, such as the wonderful Summer of Code, so it perhaps equals out. Although developers don't have official Google apps, they have plenty of motivation to create their own versions.

Inboxer is fundamentally a web wrapper around Google's very useful Inbox, which is itself an advanced version of Gmail that attempts to do clever things to the endless stream of email you typically get. It will prioritize messages it thinks are important, automatically bundle groups of related messages together, allow you to set reminders with ease, or delay your attention on an incoming email until later. When you first start to use Google Inbox, it can feel like magic, as bundles are automatically created to list your trips, finance messages, or purchases. It can make a real difference to chaotic inboxes (and you can still access your old version via Gmail). If you're happy giving up manual obsessive control over



**Google's Inbox is great, but Inboxer is better; it turns a tab in your browser into a desktop application.**

your inbox, Inbox actually works very well, and so does Inboxer. Inboxer uses Electron to pull the web application out of the web to better integrate Inbox with your desktop. It looks more like your typical Linux application, it runs as a separate process, and it embeds desktop notifications, unread badges, and a system tray. If you find yourself using Inbox more and more, you should definitely check out Inboxer.

**Project Website**

<https://github.com/denysdovhan/inboxer>


## 3D renderer

# Tungsten

**T**he point of entry for this discovery is actually the excellent GitHub repository of its author, Benedikt Bitterli, which contains some wonderfully small projects that simulate various aspects of two-dimensional and three-dimensional fluid dynamics. These tools use your GPU to generate, mostly in real time, simulations of how fluid moves, from turbulent storms to explosive mushrooms of rising heat. You can watch the eddies of particles move for hours; even when rendered in grayscale, the output is fascinating. The C++ and GPU shader code that accompanies these projects, some of which are more than five years old, provides a wonderful set of examples and starting frameworks for your own experiments. If you don't understand the code,

clone the repository; building each project is usually no more difficult than typing `make` after grabbing a few OpenGL dependencies (e.g., `libgl-mesa-dev`).

The real gem in this collection, though, is the Tungsten Renderer. This more ambitious project is written in C++ 11 to exploit Intel's Embree ray-tracing library and multithreading. It's a 3D renderer that supports many different geometric shapes and material models (e.g., hair, plastic, and coat surfaces), photon mapping, light tracing, path tracing, and atmospheric effects. The output is hyper-realistic, and there are some excellent JSON-formatted example files to play with. Its offline rendering and beautiful light models are reminiscent of earlier Amiga versions of an application called Real3D.



**The Tungsten renderer includes an editor, command-line rendering tools, and even a server to send and monitor jobs.**

But Tungsten is more academic than this (indeed Blender, the graphical editor, is much more primitive by comparison), but the whole project is also easier to understand and will help anyone interested in the computer science of light and material composition within a three-dimensional scene, especially if they've some experience with OpenGL and graphical primitives.

**Project Website**

<https://github.com/tunabrain>

## Data receiver

## rtl\_433

**H**undreds of small embedded devices receive small chunks of data transmitted at 433.92MHz from a sensor somewhere. Weather sensors commonly send data to a receiver located somewhere dry and warm, whereas different home automation systems can use the data from a variety of different temperature sensors to do things like stop pipes from freezing, warm a house before you get home, or open the gates as your car approaches. Home automation is one of the biggest consumers of this kind of technology, but every vendor seems to implement its own set of protocols for its own set of hardware, making it very difficult for those of us trying to implement an

open source system that may integrate some of their kit. Thanks to much of this hardware being built around the Realtek RTL2832 receiver chip, a chip that can also be found in many common USB DVB digital television receiver dongles, all that was needed to turn those proprietary dongles into open source data sources was an open source implementation of their driver protocols. This is exactly what rtl\_433 does. It supports nearly 100 separate devices, from car key fobs to oil tank level monitors, and grabs the data those devices are broadcasting. With a simple rtl\_433 command, that data can be output to your Linux command line and, from there, transformed into almost any-

Like a pen in the hands of an author, the humble rtl\_433 utility can take simple readings and transform them into smart home automation controllers.

thing. With the Domoticz home automation system, for instance, you can very easily execute a read command at set intervals and use this to update a virtual sensor within a home automation system, triggering an alarm, perhaps, or sending a new heating oil order. These small tools are the epitome of what's made Unix/Linux successful, and it's the best thing about open hardware in general.

## Project Website

[https://github.com/merbanan/rtl\\_433](https://github.com/merbanan/rtl_433)

## Bash alternative

## Oil 0.3.0

**I**t often surprises users of proprietary operating systems, and Microsoft Windows users in particular, that not only does Linux provide choice when it comes to desktop environments, it also offers a plethora of choice between different kinds of shells or terminals. Many of these shells have been around for years, even decades, but not Oil. Oil is a new shell that aims to be broadly compatible with Bash and the way most of us are used to working, while at the same time upgrading Bash's scripting abilities with more serious programming credentials. Most of us have written Bash scripts to tie certain convenience functions together, but these scripts are seldom used for more ambitious programming projects where Python might be used, for

example. And that's exactly the niche Oil is targeting, both in wanting to be a programming platform itself, and – almost ironically – by being written in Python.

From a shell perspective, Oil upgrades the capabilities of the script syntax to offer the same advantages you'd expect from a modern programming language – they become reproducible on other systems or run in containers, run within data science and scientific computing environments, and even run to build distributions, for example. From a programming perspective, it implements proper arrays, script modes, richer constructs for concurrency, and parallelism. There's even the author's secret plan to use Oil as the basis for a distributed operating system. Oil

Upgrade Bash scripting by creating an entire program from a string literal with Oil.

even manages to do all this without losing humble compatibility with Bash, including a totally new and “very compatible” Bash parser. It's even possible to convert Bash scripts to the language used by Oil, so although you don't get the convenience features you may have worked into your own Bash environment, you can switch and test Oil without any lengthy relearning, which definitely makes Oil worth a look.

## Project Website

<https://github.com/oilshell/oil>

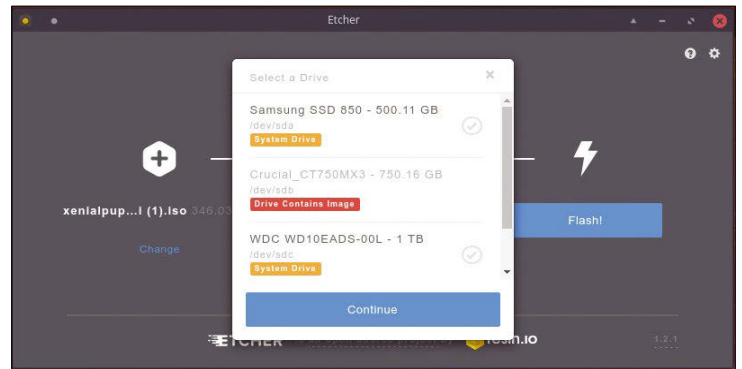


ISO to USB

# Etcher 1.3.0

If you've ever written instructions on how to write an ISO image to a USB drive, you'll know this seemingly simple task is threaded with danger. Not only is the concept mystifying – a newbie will wonder why an ISO needs to be written in a special way, for instance – it's difficult to explain without going into too much detail. You want to point the reader at a simple three-step procedure without having to explain why dropping the ISO on the mounted USB drive won't work. There are lots of tools for writing an ISO file to an external USB drive, from the Russian roulette of `dd` to the clunky ubiquity of UNetbootin, but no tools combine cross-platform support, open source, and ultra-sleek visuals with a procedure simple

enough to explain with a few hundred words – none apart from Etcher, that is. Etcher is one of those rare applications that combines both form and function without any noticeable trade-off, unless you count the 53MB download or its use of the bloated Electron platform. It splits the process into three simple steps: Choose the ISO, select the destination, and click *Flash!*. Etcher will take care of everything else. It also makes clear which destination devices are system drives, which are external, and which already contain an image, all of which is vital if you're preparing to install your first Linux system. It's also



Etcher looks better than any USB ISO burning tool deserves to look.

constantly being updated, and this release adds useful write delays, a native application menu, improved error handling, and fixes when scanning a drive. If you find yourself needing to help someone through their first Linux installation, regardless of their operating system, Etcher should be close to the top of your Linux tools for them to try.

**Project Website**

<https://github.com/resin-io/etcher>

Lottery predictor

# fortune-mod

If you started using Linux or even Unix some considerable time ago, especially if it was before you started using a mouse with X windows, you could always rely on a few things. The first was that you could typically send the output from the `banner` command directly to the default line printer connected to the network; the second was that the `finger` command was actually useful for seeing who else was using a machine; and the third was that you could rely on a snippet of wisdom being displayed whenever you logged in to your console. That snippet of wisdom was courtesy of a simple command called `fortune`, which was typically set to run automatically either when you connected or logged out.

Remarkably, `fortune-mod` is a continuation of that original project and is polishing up for the milestone release of version 2.0. Sadly, this doesn't include spinning cube effects or surround sound, but it does update the codebase to use the much more modern `cmake`, fixes lots of typos in the data files that still exist, and builds support for Travis testing. This update basically makes `fortune` fighting fit for the 21st century, and if you've not used it for a couple of decades, it's definitely worth a revisit. It may seem jokey – and some of the quips are – but it can also offer a moment of genuine thought now that we typically spend so much time in front of screens. And what's great about building the new



Happiness is having a scratch for every itch. – Ogden Nash

version for yourself is that you can see the hard work that's gone into the code, as well as the 41 data files that contain the wisdom. This project is far more than a Perl script with a few embedded quotes.

**Project Website**

<https://github.com/shlomif/fortune-mod>

## Terminal IRC client

## Irssi 1.1.0

These pages have often included updates to the WeeChat IRC client. What makes WeeChat so good is that it runs from a terminal, it's mature, and it's crammed full of features, scripts, and potential. But there has always been one big rival to WeeChat, and one that many users even consider superior, and that's Irssi. Irssi has been around longer than WeeChat, almost two decades, and like WeeChat, its age makes it incredibly stable and secure. Although it is older, it shares many of the same features of WeeChat, including windows, plugins, and horizontal splits, but you'll have to wait for 1.2.0 for the long, long-awaited vertical split. The main difference is configurability; whereas WeeChat may have more features

built in, Irssi may need more plugins – apart from its excellent theming – but Irssi does include options WeeChat has to emulate via plugins, such as the excellent set of tools that send message notifications across your desktop and even to your Android phone. (Thanks, IrssiNotifier.)

Terminal IRC clients like WeeChat and Irssi make sense because with IRC you're always dealing with text, which means there's no extra overhead in issuing internal commands for the client from the input prompt, and this is how Irssi works, too. Type `/connect` followed by the server name, for instance; it's the same for the huge number of internal commands both clients support. It starts off intimidating but quickly works just like any other

```

Irssi v1.1.0 - http://www.irssi.org
20:23 Irssi commands:
20:23 accept die knock notice reconnect unnotify
20:23 action disconnect knockout notify scrollbar unquery
20:23 admin echo lastlog op server unsilence
20:23 alias eval layout oper servlist upgrade
20:23 away exec links part set uptime
20:23 ban flushbuffer list ping sethost userhost
20:23 beep foresch load query silence ver
20:23 bind format log quit squery version
20:23 cat hash users quote squit voice
20:23 cd help msg rawlog stats wait
20:23 channel highlight me recode statusbar wall
20:23 clear ignore mirccdc reconnect time wallops
20:23 completion info mode redraw toggle who
20:23 connect invite noid rebash topic whois
20:23 ctcp ircnet msg reload trace whowas
20:23 cycle json names resize ts window
20:23 dcc join nickltp restart unalias
20:23 dehighlight kick netsplit reconnections unban
20:23 doop kicoban network rnrjoins unignore
20:23 devoice kill nick save unload
20:23 ! Irssi client: irssi 1.1.0 (20180108 1443)
20:23 ! Irssi: Not connected to server.

```

**Terminal IRC clients make a lot of sense because everything remains in the text domain.**

part of IRC. Irssi 1.1.0 is a major update and includes a year's worth of work, which changes age-old commands, such as the `/server` command replaced with the `/connect` command, and adds new commands, such as an accessible global history, where elements can now be deleted. If you've found WeeChat to be just a little too foreboding, try Irssi!

**Project Website**  
<https://irssi.org>


## Image viewer

## qimgv

Image viewers might not seem the most exciting application category, but when you think how many images you typically view each day, choosing one other than the default option for your desktop could save you time as well as a few CPU cycles in the process. Qimgv is a new, cross-platform image editor that's worthy of your consideration, mainly because it's quick and has a super-click user interface that doesn't get in the way of actually viewing images. Click or select an image, and qimgv displays it. If you mouse to the top of the window, preview thumbnails for images found within the same folder appear listed horizontally. This is a great feature because you're seldom looking for a single image in iso-

lation, and very few of us name our picture files anything meaningful, leaving us to scan quickly through folders of similarly named images.

Qimgv also has a great set of keyboard shortcuts for moving between images, zooming in and out, and entering full-screen mode, and it even provides a basic set of editing tools. This is the application's best feature, because you often need to crop or resize an image without wanting to resort to something like Gimp, and that's exactly what these shortcuts enable you to do, as well as copy and paste between directories. The dark theme is also worth mentioning, as is the native support for simple WebM videos. This fledgling, royalty-free format needs greater exposure,



**A super-fast Qt-based image viewer with support for WebM video playback.**

and while it's a little strange to find video playback in an image viewer, it means you need one less application – especially if you find yourself looking for such a tool on Windows!

**Project Website**  
<https://github.com/easymodo/qimgv>

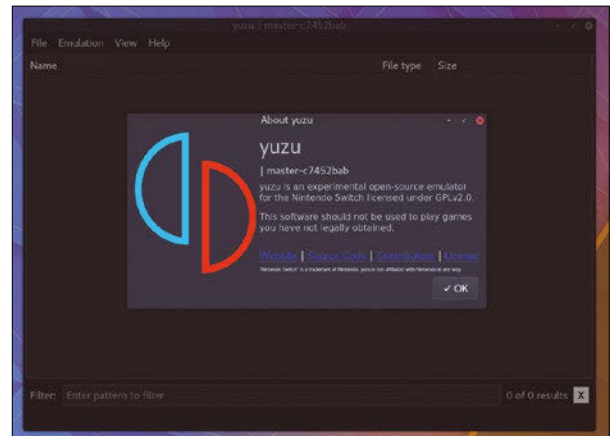
## Switch emulator

## yuzu

It's important to point out that although I'm dedicating this page to two cutting edge game console emulators, this in no way is an endorsement of software piracy. Firstly, neither of these emulators are particularly better at playing real games than the original (certainly not yuzu, yet). Even the 3DS emulator Citra lacks the real thing's inputs and 3D screen, spoiling and wasting the design of and immersion into the games that run on the real thing. Secondly, the vast majority of those who are fascinated by emulators are fascinated because of what they're doing on a technical level, and not because of the potential end result. What they're doing is virtualizing real hardware with clever bits of code, nailing the timing and the

chip interactions to such an extent that the original games and home-brew software you can sometimes run on these emulators thinks they are running on the real thing.

With that in mind, yuzu is the beginning of an exceptional emulator of Nintendo's current-generation console, Switch. Nintendo Switch uses 3DS-related hardware and some elements of Android, along with some FreeBSD code, so it's perhaps no real surprise. But in many ways, Switch is still nascent, and despite its current success, the future of Nintendo relies on its continued profitability. Vivaly, games don't yet run on yuzu. You can't really do anything with it – yet – but the fact that it already exists has many people worried. This is a le-



Although the yuzu emulator doesn't yet play games, it's likely to start adding support for 2D titles first.

gitimate concern, because gaming will likely never recover from losing a publisher as innovative and brave as Nintendo. However, as an open source project that reverse engineers a piece of hardware, it's an incredible piece of software, and one that follows the long tradition of emulation.

## Project Website

<https://github.com/yuzu-emu/yuzu>

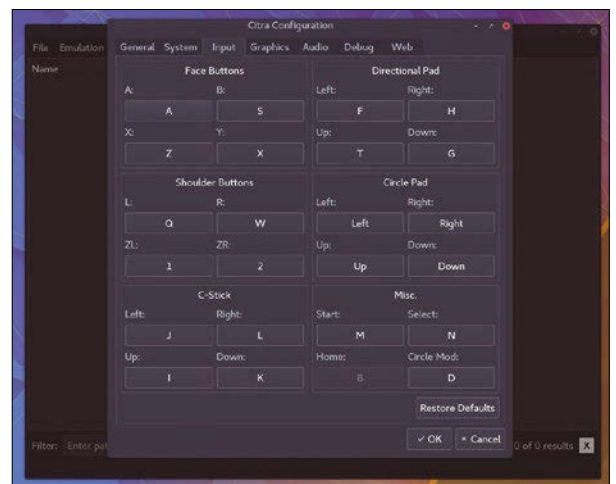
## 3DS emulator

## Citra

Citra is a Nintendo 3DS emulator from the same team behind the trailblazing yuzu Switch emulator (see above). The big difference is that 3DS has a far more modest hardware specification, because it's considerably older. It also helps that this emulator has been in development for a considerable time. As a result, Citra is relatively mature. This can be seen with the setup wizard that helps you choose between versions and automatically download the binaries of whatever you choose. It also means you can play real games if you have the hardware facility and permission to copy them from your cartridges. The main Qt-based UI window is almost identical to yuzu, and it's obvious that both emulators and hardware have a lot in common. But the big difference is that, with

Citra, you can actually play games. Like the Wine compatibility list for Windows titles, the wiki contains a list of games and a rating of how well they work: 36 titles are rated as having "Perfect" emulation, and a whopping 62 reportedly work "Great."

The games themselves will need to be sourced before you can play anything. Using your own cartridges is likely the most legal option. Games are also encrypted against your original hardware, which can create more legal and moral hurdles, but Citra is capable of excellent performance if you get that far. Like many emulators, modules are used to perform the CPU, system, input graphics, and audio hardware. Each of these can be configured separately and changed to adjust performance for your specific system. In partic-



Map any key or controller input to the game you want to play.

ular, you can change the rendering's native resolution so that the output supersedes the original viewport's size. The original 3DS selling point, of course, was its 3D screen. Unfortunately, this isn't yet supported, even if you have the hardware, but touch is emulated, as is network gameplay, so you get almost the whole experience.

## Project Website

<https://citra-emu.org/>

# Gobbling Up

## Creating custom transitions with FFmpeg

Most video editors supply you with a generic catalog of transitions, usually in the shape of tired wipes and fades. But what if you wanted something a little more special? FFmpeg to the rescue.

BY PAUL BROWN

A transition is a way of moving a film's narrative from one scene to the next. It differs from a cut in that a transition is gradual: The new scene wipes in from the left or fades in while the old scene fades out, etc. A cut, on the other hand, just jumps to the new scene.

Transitions in videos are like transitions in PowerPoint presentations: Unless used ironically or in *Star Wars* movies, they are generally considered tacky and old-fashioned. The stock transitions

that come with most graphic video editors are the worst in that respect. However, custom – let's call them "vanity" – transitions can be fun, especially if they are over the top, which is what I was aiming for when I started this project.

In preparing a promotional video for an issue of *Raspberry Pi Geek* [1], instead of using the stock transitions, I thought it would be fun to have an animation cross the screen, uncovering the next scene in its wake. You can see the effect at the 30-second mark in Figure 1.

### Vanity Transitions 101

The basic idea is you create and animate your image moving across some transparent frames. You leave the space to the left of your character transparent (each frame is a PNG image), but you fill the region to the right of the animation with a flat, primary color (usually green). Figure 2 shows what a single frame halfway through the animation would look like.

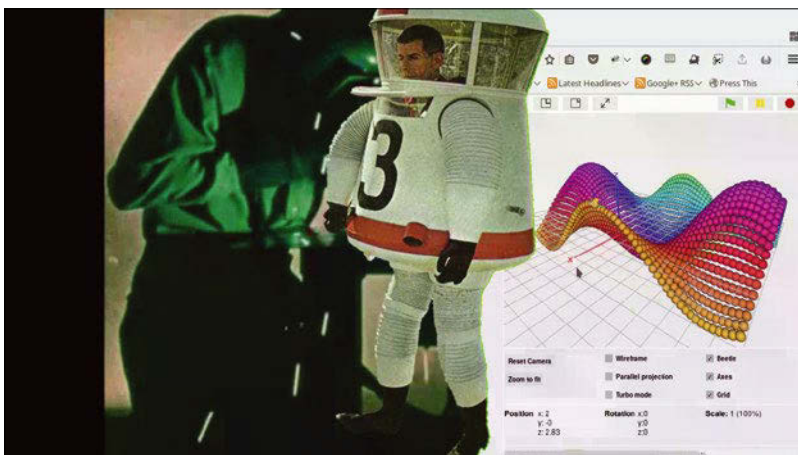
Then, in your video editor, you place your animation over the first scene. As the right side of your animation is transparent, the first scene will show through on the right (Figure 3). You merge them together and then use the colored area with what is called a chroma key filter to cut the green bit and make the area transparent so your second clip shows through on the right, as shown in Figure 1.

With Kdenlive [2], you cannot do this in one go: You have to create a project with the animation on top of the top video clip, render it, and then use the new rendered clip to merge it with the bottom clip using Kdenlive's chroma key filter (Figure 4).

That is because Kdenlive treats clips like a stack of cards: You cannot create a new card by fusing two cards together and then merging that with something else.

With Natron [3], things are different. Each clip is a node, and you can extract streams (shown as white arrows in Figure 5) from each clip node

**Figure 1:** Mr. Spaceman walks across the screen, uncovering a new scene in his wake.



**Figure 2:** A frame from the animation: transparent to the left and filled with a color you will use in your chroma key effect to the right.

and merge them using different filters, which are also nodes. You can extract streams from the merged clips and merge them further. This allows you to merge the animation with the top clip (the blue node in Figure 5) and then apply a chroma key effect to the result (the green node in Figure 5), merging the composite stream with the bottom clip. The advantage is you build your effect up node by node and only have to render once to produce the final result.

That got me thinking... Remember what other video-manipulation software works with streams [4]? FFmpeg [5].

## FFmpeg Transitions

Suppose I want to create a vanity transition in which an animated Pacman gobbles up one scene to reveal another one below. In the spirit of this section, you can do everything using command-line tools, including the Pacman animation – if you want to try, I used ImageMagick [6] for everything, from the Pacman drawings to the generation of the 48 frames used in the transition.

For this exercise, you will need the 720p version of *Sintel* [7] and the 720p version of *Tears of Steel* [8]. Both movies were produced by the Blender Foundation and are distributed under a permissive license.

Note that the two films are not really 720p; that is, they are not 1280 pixels wide by 720 pixels high. Formatted for widescreens, they are actually 1280x534p. You can confirm this using `ffprobe` with the file.

The `ffprobe` tool gives you more useful information, like the Storage Aspect Ratio (SAR). The SAR tells you the shape of the pixels that make up each frame. A SAR of 1:1 is a perfect square. A different SAR would define an oblong shape. Before you proceed, both of your films must have the same SAR, preferably 1:1.

If necessary, you can change your movie's SAR like this:

```
ffmpeg -i <original_video> -filter_complex 2
"setsar=sar=1" <corrected_video>
```

You may also want to cut out a one-minute clip to work with:

```
ffmpeg -i Sintel.2010.720p.mkv -ss 00:05:00 2
-t 00:01:00 Sintel_cut.mp4
```

Otherwise, every time you want to render something, you will have to render two complete films, and that takes a long time.

The `-ss` parameter tells FFmpeg the starting time from which to cut. `-t` tells FFmpeg the clip's length. The preceding command will give you a clip that starts at the five-minute mark in the *Sintel* film and that lasts for one minute.




Figure 3: The first clip overlaid with the animation and its chroma key layer.

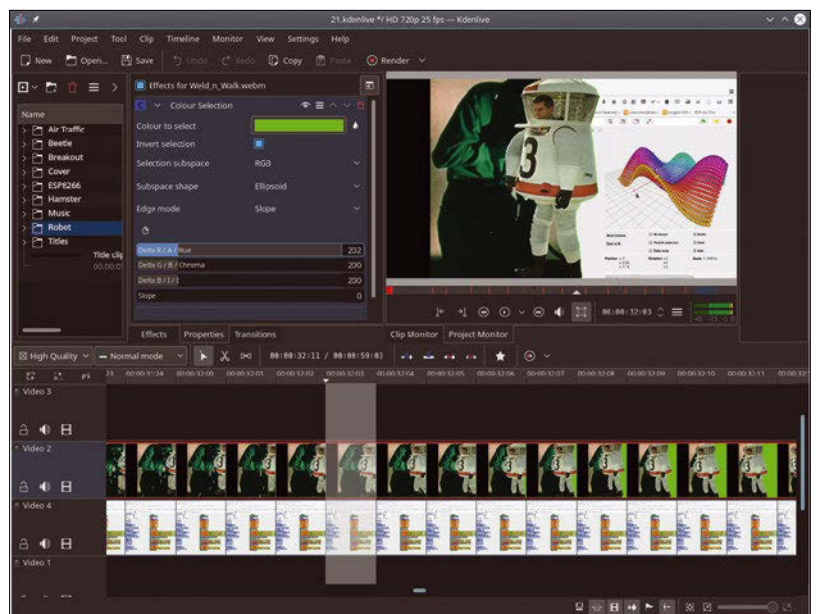


Figure 4: In Kdenlive, you have to merge and render the animation with the top layer before you can merge it with the bottom layer.

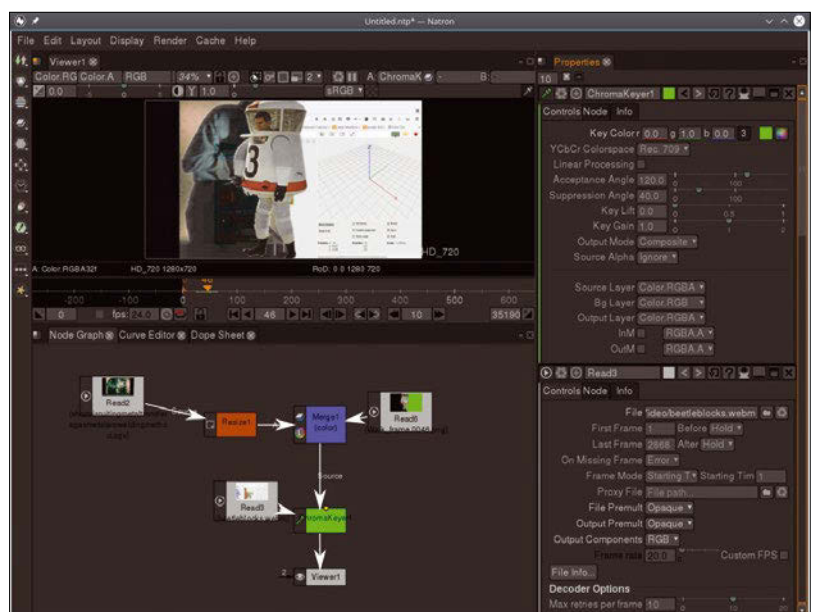


Figure 5: Merging the top clip with the animation and then using a chroma key filter to merge with the bottom clip in Natron.




Figure 6: A PNG image with a green area overlaid onto the *Sintel* clip.

Do the same with the *Tears of Steel* film. Now that you have your two clips, you can start experimenting with the chroma key filter.

First try using a static PNG image with a green section. Create a frame-sized (1280x534) image called `mask.png`, which is half transparent and half green. Then overlay your image on the *Sintel* clip like this:

```
ffmpeg -i Sintel_cut.mp4 -i mask.png \
-filter_complex "overlay" Sintel_cut_ck.mp4
```

The resulting `Sintel_cut_ck.mp4` clip should look like Figure 6.

Now you can merge that clip with the *Tears of Steel* clip using the chroma key filter:

```
ffmpeg -i tears_of_steel_cut.mp4 \
-i Sintel_cut_ck.mp4 \
-filter_complex "[1:v] format=rgba, \
chromakey=0x00FF00:0.3 [ckout]; \
[0:v] [ckout] overlay; \
amix=inputs=2" Sintel_ToS_ck.mp4
```

As discussed in my previous article on FFmpeg [4], you can label input streams and then the derivative streams. This makes passing the streams from one filter to the next, and finally to the output, easier and non-ambiguous.

In this case:

- `[0:v]` is the first input stream (`tears_of_steel_cut.mp4`).

- `[1:v]` is the second input stream (`Sintel_cut_ck.mp4`). Remember that `Sintel_cut_ck.mp4` is the result of overlaying the chroma key image onto the *Sintel* clip.

- `[ckout]` is the stream that is the result of applying the `chromakey` filter to `[1:v]`.

In this example, after converting the second stream (`[1:v]`) to the RGBA color space, you pass two parameters to the `chromakey` filter:

1. `0x00FF00` uses the classical RRGGBB format to tell FFmpeg the color you want to use as the chroma key.
2. `0.3` is the degree of similarity with your chosen color. `0.01` will only affect pixels exactly the same color as your chosen color (so only pure green pixels), while `1.0` will affect all pixels, regardless of their color. `0.3` is just enough to avoid the “green ghosting” effect often seen around the contours of objects placed in front of a green screen.

Also notice that you use the `amix` filter to mix together the sound streams of both clips. If you don't use `amix`, the final clip will have the soundtrack from one clip or the other, but not from both.

Figure 7 shows how all this works out in the finished video.

That is how Kdenlive would do things. But, as mentioned above, FFmpeg is more similar to Natron. Since it can process input streams live and turn them into other streams, it can do everything in a single pass:

```
ffmpeg -i tears_of_steel_cut.mp4 \
-i Sintel_cut.mp4 -i mask.png \
-filter_complex "[1:v] [2:v] overlay [cmask]; \
[cmask] format=rgba, \
chromakey=0x00FF00:0.3 [ckout]; \
[0:v] [ckout] overlay; \
amix=inputs=2" Sintel_ToS_ck.mp4
```

In the preceding instruction, you create the `[cmask]` stream by overlaying stream `[2:v]` (`mask.png`) onto stream `[1:v]` (`Sintel_cut.mp4`). The `chromakey` is then applied to the new `[cmask]` stream. The resulting stream, `[ckout]`, is overlaid onto stream `[0:v]` (`tears_of_steel_cut.mp4`).

And now the grand finale: Using an animation's sequence of frames (as opposed to a static image) for the chroma key layer isn't substantially more complicated:

```
ffmpeg -i tears_of_steel_cut.mp4 \
-i Sintel_cut.mp4 -framerate 24 \
-i frames/pmframe%03d.png \
-filter_complex "[1:v] [2:v] overlay [cmask]; \
[cmask] format=rgba, \
chromakey=0x0000FF:0.3 [ckout]; \
[0:v] [ckout] overlay; \
amix=inputs=2" Sintel_ToS_anim.mp4
```



Figure 7: *Sintel* and *Tears of Steel* merged using the chroma key filter.

There are only two new things that are noteworthy in the preceding command:

1. The input for the third input stream does not point to an image, but to a directory called `frames/` containing a sequence of images (Figure 8) with names like `pmframe000.png`, `pmframe001.png`, `pmframe002.png`, etc. (`pmframe%03d.png` means `pmframe` followed by three digits.)


2. `-framerate 24` tells FFmpeg how many frames per second it must use when creating the video stream from the sequence of images.

When you play back your video, you'll see your animated Pacman crossing the frame, uncovering the next scene in your video (Figure 9).

### Things to Try

You can try correcting the sound fading out of the *Sintel* video and into the *Tears of Steel* video. FFmpeg has nearly as many audio filters as it does video filters, so finding one that fades from one audio stream to another shouldn't be hard. In fact, the `acrossfade` filter [9] looks promising.

The second thing worth tinkering with is the timing of the transition. As it stands, the transition starts right at the beginning of the clip. You could cut out the moment you want the transition to start within *Sintel* and *Tears of Steel*, create the clip with the transition, and then glue all the bits together using separate instructions. FFmpeg supplies a way to edit tasks like this on the fly. You may want to check out the `trim` [10]




**Figure 8:** A few frames from the Pacman animation used as a transition. Notice the use of blue (`0x0000FF`) as the chroma key color, because green clashes with some of the colors in the *Sintel* video layer.

and `concat` [11] filters for that little bit of extra video-editing wizardry. ■■■

### Info

- [1] *Raspberry Pi Geek #21* promo video: <https://youtu.be/nlmU7NlXrMA>
- [2] Kdenlive: <http://kdenlive.org/>
- [3] Natron: <https://natron.fr/>
- [4] *Video Wizardry* by Paul Brown, *Linux Magazine*, issue 206, January 2018, <http://www.linux-magazine.com/Issues/2018/206/Tutorials-FFmpeg>
- [5] FFmpeg: <https://ffmpeg.org/>
- [6] ImageMagick: <http://www.imagemagick.org/>
- [7] *Sintel*: <https://durian.blender.org/download/>
- [8] *Tears of Steel*: <https://mango.blender.org/download/>
- [9] `acrossfade`: <https://ffmpeg.org/ffmpeg-filters.html#acrossfade>
- [10] `trim`: <https://ffmpeg.org/ffmpeg-filters.html#trim>
- [11] `concat`: <https://ffmpeg.org/ffmpeg-filters.html#concat>

**Figure 9:** Pacman chews up his scenes.



# 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](mailto:events@linux-magazine.com).



## Linux Storage Filesystem and MM Summit

**Date:** April 23–25, 2018

**Location:** Park City, Utah

**Website:** <https://events.linuxfoundation.org/events/lsfmm-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.

## LinuxFest Northwest 2018

**Date:** April 28–29, 2018

**Location:** Bellingham Technical College, Washington

**Website:** <https://linuxfestnorthwest.org/conferences/lnfw18>

LFNW features presentations and exhibits on free and open source topics, as well as Linux distributions & applications, InfoSec, and privacy; something for everyone from the novice to the professional!

## openSUSE Conference 2018

**Date:** May 25–27, 2018

**Location:** Prague, Czech Republic

**Website:** <https://events.opensuse.org/conference/oSC18>

The annual openSUSE community event that brings people from around the world together to meet and collaborate. The organized talks, workshops, and BoF sessions provide a framework around more casual meet ups and hack sessions.

### Events

CloudFest 2018	March 10–16, 2018	Rust, Germany	<a href="https://www.cloudfest.com/">https://www.cloudfest.com/</a>
Embedded Linux Conference	March 12–14, 2018	Portland, Oregon	<a href="http://events.linuxfoundation.org/events/embedded-linux-conference">http://events.linuxfoundation.org/events/embedded-linux-conference</a>
OpenIoT Summit	March 12–14, 2018	Portland, Oregon	<a href="http://events.linuxfoundation.org/events/openiot-summit">http://events.linuxfoundation.org/events/openiot-summit</a>
Open Networking Summit	March 26–29, 2018	Los Angeles, California	<a href="https://events.linuxfoundation.org/events/open-networking-summit-europe-2018/">https://events.linuxfoundation.org/events/open-networking-summit-europe-2018/</a>
NSDI '18	April 9–11, 2018	Renton, Washington	<a href="https://www.usenix.org/conference/nsdi18">https://www.usenix.org/conference/nsdi18</a>
heise Security Tour 2018	April 10, 12, 18, 24, and 26, 2018	Numerous European cities	<a href="https://www.heise-events.de/securitytour">https://www.heise-events.de/securitytour</a>
2018 HPC for Wall Street	April 16, 2018	New York, New York	<a href="http://www.flaggmgmt.com/linux/">http://www.flaggmgmt.com/linux/</a>
Cloud Foundry Summit North America	April 18–20, 2018	Boston, Massachusetts	<a href="https://www.cloudfoundry.org/event/nasummit2018/">https://www.cloudfoundry.org/event/nasummit2018/</a>
Linux Presentation Day 2018.1	April 21, 2018	Europe-wide in many cities	<a href="http://www.linux-presentation-day.org/">http://www.linux-presentation-day.org/</a>
Linux Storage Filesystem and Memory Management Summit	April 23–25, 2018	Park City, Utah	<a href="http://events.linuxfoundation.org/events/lsfmm-2018/">http://events.linuxfoundation.org/events/lsfmm-2018/</a>
LinuxFest Northwest	April 28–29, 2018	Bellingham, Washington	<a href="https://linuxfestnorthwest.org/conferences/lnfw18">https://linuxfestnorthwest.org/conferences/lnfw18</a>
KubeCon + CloudNativeCon	May 2–4, 2018	Copenhagen, Denmark	<a href="https://events.linuxfoundation.org/events/kubecon-cloudnativecon-europe-2018/">https://events.linuxfoundation.org/events/kubecon-cloudnativecon-europe-2018/</a>
openSUSE Conference 2018	May 25–27, 2018	Prague, Czech Republic	<a href="https://events.opensuse.org/conference/oSC18">https://events.opensuse.org/conference/oSC18</a>
Texas Linux Fest	June 8–9, 2018	Austin, Texas	<a href="https://2018.texaslinuxfest.org/">https://2018.texaslinuxfest.org/</a>
CeBIT 2018	June 11–15, 2018	Hannover, Germany	<a href="http://www.cebit.de/en/">http://www.cebit.de/en/</a>
Open Source Data Center Conf.	June 12–13, 2018	Berlin, Germany	<a href="https://osdc.de/">https://osdc.de/</a>



# 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](mailto: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](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	36, 42, 66
Swapnil Bhartiya	8, 16, 52
Paul Brown	92
Zack Brown	12
Bruce Byfield	54, 62
Joe Casad	3, 75
Mark Crutch	75
Tobias Eggendorfer	46
Jon "maddog" Hall	77
Frank Hofmann	24
Charly Kühnast	50
Vincent Mealing	75
Pete Metcalfe	70
Graham Morrison	86
Mandy Neumeyer	24
Mike Schilli	58
Ferdinand Thommes	18, 30, 78
Daniel Tibi	82

## Contact Info

### Editor in Chief

Joe Casad, [jcasad@linux-magazine.com](mailto:jcasad@linux-magazine.com)

### Managing Editor

Rita L Sooby, [rsooby@linux-magazine.com](mailto:rsooby@linux-magazine.com)

### Localization & Translation

Ian Travis

### News Editor

Swapnil Bhartiya

### Copy Editor

Amy Pettle

### Layout

Dena Friesen, Lori White

### Cover Design

Dena Friesen

### Cover Image

© DesignMachine and Kadriya Gatina, 123RF.com

### Advertising

Brian Osborn, [bosborn@linuxnewmedia.com](mailto:bosborn@linuxnewmedia.com)  
phone +49 89 3090 5128

### Marketing Communications

Gwen Clark, [gclark@linuxnewmedia.com](mailto:gclark@linuxnewmedia.com)  
Linux New Media USA, LLC  
616 Kentucky St.  
Lawrence, KS 66044 USA

### Publisher

Brian Osborn

### Customer Service / Subscription

For USA and Canada:  
Email: [cs@linuxpromagazine.com](mailto: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](mailto:subs@linux-magazine.com)

[www.linuxpromagazine.com](http://www.linuxpromagazine.com) – North America

[www.linux-magazine.com](http://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.

**Approximate**

UK / Europe

Apr 07

USA / Canada

May 04

Australia

Jun 04

**On Sale Date**

Issue 210 / May 2018

# Libreboot

Many Free Software devotees are weary of dealing with the complications of opaque and inscrutable proprietary BIOS systems. Libreboot is a free BIOS or UEFI replacement that will initialize hardware and boot the operating system with no backdoors and no mystery.

## 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](http://www.linux-magazine.com/newsletter)

Image © Viacheslav Krisanov, 123RF.com

# COMPLETE YOUR LIBRARY

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



**ORDER YOURS TODAY!**  
[shop.linuxnewmedia.com](http://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!

# Application-Optimized Server Solutions

Exactly the Best Performance & Density at Enterprise Scale



## The Most Efficient Data Center in the World

A Fortune 100 Enterprise deploys over 50,000 Supermicro® MicroBlade™ servers at its Silicon Valley data center. With a PUE of 1.06, and Intel® Xeon® Scalable Processors it achieves an 88% improvement in energy efficiency and over \$13M in annual savings. For Exactly the Best Enterprise Computing Solution for your Application, Choose Supermicro.



*Intel Inside®. Powerful Productivity Outside.*



Application-Optimized Server Solutions

*Learn more at [supermicro.com](http://supermicro.com)*