

FREE DVD fedora 24  &  Linux Mint 18
"Cinnamon" 32-bit

ALTERNATIVE
FILE MANAGERS

LINUX
PRO



LINUX

MAGAZINE

SEPTEMBER 2016

ALTERNATIVE FILE MANAGERS

Find the perfect tool for your work habits

Webmin

Why this classic configurator is still an admin favorite

Chat Pack

We check out some alternatives to HipChat and Slack

Qt4-fsarchiver

Back up files and create disk images with this easy archive tool



Magic URL

Connect to your server from a cell phone – without a fixed IP address

Open365

LibreOffice and KDE Kontact in the Cloud

Mofo

This secretive Linux was designed to defeat state censorship

XMP Toolkit

Manipulate image metadata

Issue 190 US\$ 12.99
September 2016 CAN\$ 13.99



WWW.LINUXPROMAGAZINE.COM

HETZNER
ONLINE

We're always on guard

Security made in Germany

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, Intel Xeon and Xeon Inside are brands of the Intel Corporation in the USA or other countries.



Dedicated Root Server PX61-NVMe

Intel® Xeon® E3-1275 v5
Quad-Core Skylake Processor
64 GB DDR4 ECC RAM
2 x 512 GB NVMe Gen3 x4 SSD
Guaranteed 1 Gbit/s bandwidth
100 GB Backup Space
30 TB traffic inclusive*
No minimum contract
Setup Fee \$132

monthly \$ **65**.50

Data Storage in Germany

Keeping your data safe and secure in Hetzner Online's own German data centers in accordance with the high standard of data protection in Germany.

Located in the heart of Nuremberg and in Falkenstein/Vogtland, Germany, Hetzner Online's two state-of-the-art data center parks provide an excellent and safety infrastructure for your Internet business.

* There are no charges for overage. We will permanently restrict the connection speed if more than 30 TB/month are used. Optionally, the limit can be permanently cancelled by committing to pay \$1.30 per additional TB used.

www.hetzner.de/us

OLD AND NEW

Dear Linux Pro Reader,

I have put in my time as the news writer for this magazine at various moments over the last 11 years, and one way or another, the European Commission has always been part of the big stories. I don't even live in Europe, but I've penned many stories on the European Commission rulings on privacy, patent reform, and antitrust issues. The European Commission is very interested in high-tech policy; the official website lists topics for "Cybersecurity and Digital Privacy," "Online Trust," and "Content and Media" in the highest level of its directory.

People who complain about consumer security in the European Union (EU) shouldn't even bother moving to the US. The US has almost no protections from data mining, cookie tracking, mailing list rental, and other excesses of the digital age. The US also has fewer restrictions on software patents and more protections for digital copyright holders. Although most of the Free Software crowd still has some issues with privacy, patents, and corporate overreach in the EU, most observers believe the US and EU have very different viewpoints on digital matters.

Now that the UK has voted to exit the EU, it is interesting to consider where they might emerge on this digital spectrum once they become masters of their own national destiny. Will they continue with policies similar to the EU, or will they drift somewhere closer to the US? Many people believe the UK has always been a little more like the US than the rest of Europe. Will it come to occupy some kind of middle ground on these high-tech, hot-button issues? The UK definitely takes the business of business very seriously, but at the same time, they aren't as much under the influence of the colossal high-tech companies that throw their weight around in the US.

It is hard for me to speculate on how things will turn out for the UK tech scene, but it is interesting to consider that, whatever the people of the UK do, they will be designing a solution for the world we know *right now*. The European Union dates back to the Maastricht treaty in 1993, which is well in the "recent events" space for most endeavors but is ancient history in tech. Issues like cyberprivacy, software

patents, and bandwidth throttling weren't on anyone's mind 23 years ago, and many of the problems we face stem from the fact that our legal institutions weren't designed to cope with these kinds of issues.

Most observers say the Brexit transition will be gradual, with any major changes still a few years away, but by the time any changes actually happen, the framework for those changes will already be in place. The discussions are probably already beginning for what the new system will look like. And to the extent that these hot-button, high-tech political issues boil down to business interests versus the interests of everyday people, you can be confident the business interests are already working on them. For instance, the EU has been moving toward a two-tiered, unitary patent system for years, which the UK will presumably be released from, unless they negotiate some form of further participation as part of an exit strategy. Any company in the UK that would benefit from the lax, pro-business, US-style software patents, for instance, has probably already held some strategy sessions on how they are going to make this happen. And that means, of course, that anyone who opposes such measures had better get busy *now* and not wait for three years to join a protest march.

Joe

Joe Casad,
Editor in Chief





LINUX PRO MAGAZINE

news

- 8 SUSE + MICROSOFT**
 - SUSE joins the Microsoft Test Drive program
 - SiFive launches a line of open source chips
- 9 FEDORA 24**
 - Fedora 24 is released
 - Canonical announces cross-distro support of Snaps
 - More online
- 10 OWNCLOUD**
 - ownCloud forked as Nextcloud
 - HP Enterprise unveils a futuristic memory-driven computer
 - More online

SERVICE

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

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.

ALTERNATIVE FILE MANAGERS

Beyond usual suspects like Dolphin, Nautilus, and Thunar are a garden of alternative file managers.

12 Alternative File Managers

We look at four Linux file managers that might be off your radar: XFE, Sunflower, Vifm, and Double Commander

Community Notebook

- 86 Doghouse: RISC vs. CISC**
maddog revisits the RISC vs. CISC controversy.
- 88 Free Fonts**
How the SIL Font License enabled the rise of the free font movement.

- 92 Free Laptop Project**
The earth-friendly EOMA68 laptop promises to be easy to repair, easy to upgrade, and easy to secure ... but will it catch on?



HIGHLIGHTS

24 WEBMIN

Meet a web-based graphical management tool for every distribution.

54 XMP TOOLKIT

Embed metadata for digital media with the Extensible Metadata Platform.

66 OPEN365

An alternative to Microsoft Office 365 and Google productivity software.

80 MOFO LINUX

This distro aims to protect your data, anonymity, and Internet activity.

FEATURES

30 Magic URL

With a highly secure Linux server, you don't need a fixed IP address to connect over a cellphone network.



44 Qt4-fsarchiver

Back up files, complete partitions, and create disk images with a single mouse click.

24 Webmin

A distribution-independent graphical management tool.

38 Chat Alternatives

InspIRCd, Rocket.Chat, and Mattermost: alternatives to HipChat and Slack.

50 Ask Klaus!

New fdisk behavior, booting Knoppix from USB, and using PHP mail() in Apache.

54 XMP Toolkit

Develop applications that store metadata in files.

60 Charly – testssl.sh

Use a shell script that talks plain English to integrate and test SSL certificates.


62 Perl – Secretary Problem

When is it statistically favorable to stop searching or gambling? Solid algorithms point the way.

LINUXUSER

66 Open365

Office suite, email client, and video chat run as software-as-a-service on the network.



70 Command Line – Package Managers

Seeking the universal Linux package manager.

74 Workspace – Pixeluvo

A capable tool for editing images and photos on Linux.

80 Mofo Linux

Achieve secure digital communications even in places where it is politically or ideologically discouraged.

REVIEW

20 Tool Tips

Tool tests on the fast track.



fedora 24 f
&
Linux Mint 18
"Cinnamon" 32-bit

SEE PAGE 6 FOR DETAILS

On the DVD

fedora 24 
Workstation 64-bit

ISSUE 190 SEPT 2016

LINUX

 Linux Mint 18
"Cinnamon" 32-bit

ISSUE 190 SEPT 2016

LINUX
MAGAZINE

Although the Linux Magazine DVD has been better and is to the best of our knowledge free of malicious software and defects, Linux Magazine cannot be held responsible, and is not liable for any damages, loss, or damage to data and computer systems caused by the use of the DVD.

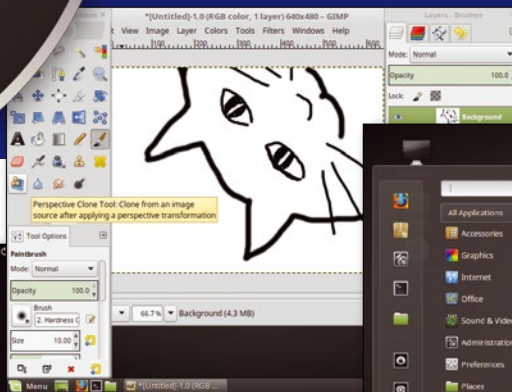
**TWO TERRIFIC
DISTROS**
**DOUBLE-SIDED
DVD!**

Fedora 24 Workstation (64-bit)

Fedora is a Red Hat sponsored Linux with a global community of users and developers. The default desktop is GNOME, which has been updated to v3.20. Desktop users will notice better integration of Qt applications and upgrades to Firefox 45, LibreOffice v5.1, and NetworkManager 1.2. Thunderbird now includes the Lightning calendar and scheduling extension by default. Developers will be interested in upgrades to Python 3.5, Mono 4.2, Ruby 2.3, glibc 2.23, and GCC 6. Fedora 24 has two new applications for the BBC micro:bit: Mu, a Python editor for writing code, and uFlash (micro flash), a utility that flashes programs onto the device.

Linux Mint 18 Cinnamon (32-bit)

This long-term support release (until 2021) with Cinnamon v3.0 is named "Sarah." The main window and Preferences window now use stack widgets and subtle animations, with better support for alternative themes. Update Manager also received a number of improvements, both visual and under the hood, and update policy concepts of regressions, stability, and security are explained. The X-Apps project properly integrates a number of GNOME applications for the GTK desktop environment, providing the same set of core applications across desktops.



ADDITIONAL RESOURCES

- [1] Fedora project: <https://getfedora.org>
- [2] Fedora 24 release notes: https://docs.fedoraproject.org/en-US/Fedora/24/html/Release_Notes/
- [3] Fedora wiki: https://fedoraproject.org/wiki/Fedora_Project_Wiki
- [4] Linux Mint: <https://linuxmint.com>
- [5] Mint 18 release notes: https://linuxmint.com/release_sarah_cinnamon_whatsnew.php
- [6] Mint user guides: <https://linuxmint.com>

Defective discs will be replaced. Please send an email to cs@linuxpromagazine.com.



What?!
Archives
come with my
digital subscription?



Archives + Current!

Sign up for a digital subscription to get the latest issues of Linux Magazine, PLUS access to archive articles.

shop.linuxnewmedia.com/digisub

That's a lot of articles!

NEWS

Updates on technologies, trends, and tools

THIS MONTH'S NEWS

08 SUSE Joins Microsoft Test Drive Program

- SUSE Joins the Microsoft Test Drive Program
- SiFive Launches A Line of Open Source Chips

09 Fedora 24 Released

- Fedora 24 Released
- Canonical Announces Cross-Distro Support of Snaps
- More Online

10 ownCloud Forked

- ownCloud Forked
- HP Enterprise Unveils a Futuristic Computer
- More Online

SUSE Joins the Microsoft Test Drive Program

Long-time Microsoft partner SUSE has joined the Microsoft Test Drive program to offer free test drives of SLES for HPC (High Performance Computing) through the Azure Marketplace. The Linux vendor has also joined the Microsoft Enterprise Cloud Alliance.

According to SUSE, SLES seamlessly integrates with Azure cloud services to deliver an easily manageable cloud environment. More than 11,400 applications are certified and supported on SUSE Linux Enterprise Server on the growing network of global data centers with Azure, including Microsoft Government Cloud.



© lassedesignen, Fotolia.com

Linux has been steadily increasing its presence in Microsoft's Azure Cloud. According to Mark Russinovich, CTO of Microsoft Azure, one out of three Azure cloud machines runs Linux.

"SUSE's growth on Azure means more options for customers," said Naji Almahmoud, Vice President of cloud and software alliances for SUSE. "It is yet another indicator of SUSE's commitment to helping enterprise customers be successful."

The Microsoft SUSE partnership dates back to 2006, when SUSE was owned by Novell.

SiFive Launches a Line of Open Source Chips

A San Francisco-based company known as SiFive is trying to bring the open source development model to the chip industry. The company has announced its first Freedom family of system-on-a-chip (SoC) products, including the Freedom U500 and Freedom E300 platforms. SiFive is a fabless semiconductor company, similar to AMD. The company doesn't fabricate the chip but outsources it to manufacturers.

The SiFive founders developed a free and open RISC-V instruction set architecture at the University of California, Berkeley. This new architecture is the foundation of the two new chip platforms.



According to SiFive, the Freedom family of chips represents a fundamentally new approach to designing and producing SoCs that redefines traditional silicon business models and reverses the industry trend of rising licensing, design, and implementation costs.

“The semiconductor industry is at an important crossroads. Moore’s Law has ended, and the traditional economic model of chip building no longer works,” said Yunsup Lee, co-founder of SiFive and one of the original creators of RISC-V. “Unless you have tens — if not hundreds — of millions of dollars, it is simply impossible for smaller system designers to get a modern, high-performance chip, much less one customized to their unique requirements.”

Fedora 24 Released

The Red-Hat-sponsored Fedora project released version 24 of the popular Linux distribution. This release comes in three versions: workstation, cloud, and server. The main highlight of the release is the arrival of version 2.3 of glibc (GNU C Library). The base packages of Fedora are rebuilt using the new library to offer better code optimization and stability.

Fedora Cloud users will be able to use OpenShift Origin, a Kubernetes distribution, to orchestrate containers. Matthew Miller of Fedora wrote in Fedora Magazine, “Fedora is on its way to being the best platform for containerized applications, from base Fedora container images to a full-featured platform as a service to run and manage them.”

Fedora is known as a Gnome distribution and users will get the latest Gnome 3.20 with this release. Fedora 24 also packs Flatpak, a new approach to app packaging for Linux distribution; users can install Flatpak from official repositories. Users will also be able to run full system upgrades from the Software app. Users can also test Wayland, the successor of the XOrg display server.

Fedora 24 and its spins are available for download from: <https://getfedora.org/>.



Canonical Announces Cross-Distro Support for Snaps

Canonical, the parent company of Ubuntu, announced cross-distro support for Snaps, a new package management system developed by Canonical. Although both Red Hat and SUSE have said they did not collaborate on the project, Canonical says it worked with developers from Fedora, Arch Linux, and Gentoo to take Snaps beyond Ubuntu.

Linux is in need of a different approach to software packaging and deployment. Even Linus Torvalds, the creator of the Linux kernel, is not happy with the current situation. He didn’t offer binaries of his SubSurface diving log software for Linux desktop, whereas binaries exist for Windows and macOS.

Dirk Hohndel, the maintainer of SubSurface, explained, “The current situation with dozens of distributions, each with different rules, each with different versions of different libraries, some with certain libraries missing, each with different packaging tools and packaging formats... that basically tells app developers ‘go away, focus on platforms that care about applications’.”

Red Hat and SUSE have expressed misgivings about Snaps, citing the Contributor License Agreement (CLA) requirements by Canonical as a potential deal breaker.

Red Hat is reportedly backing Flatpack, another Snaps-like approach to app distribution on Linux. This is not the first time Canonical has been at odds with the larger Linux community. Similar conflicts have erupted previously, such as Upstart vs systemd and Mir vs Wayland.

MORE ONLINE

Linux Pro Magazine

www.linuxpromagazine.com

Off the Beat • Bruce Byfield

How Graphical Installers Introduced the User
Last weekend, I was exploring GuixSD, the distribution that introduces a universal package manager developed by The GNU Project.

Is the Cost of Diversity a Lack of Innovation?

Journalism is the first draft of history. That is why I was glad to see Christine Hall’s article on how the reactions to GNOME 3 created more choices on the free desktop; too often, the community forgets its own history.

Why Debian Policy Is Important to Package Quality

Unless you are a Debian maintainer, you probably haven’t read the Debian Policy Manual. However, when Ubuntu started promoting Snappy packages as a more secure solution to package management, the claim was challenged, not by reference to the technical structure of Debian packages, but to the Debian Policy Manual.

The Nostalgia of Windows is Everyday Linux

A few days ago, I read a mailing list discussion about the advantages of running a computer in the 1980s. A few, like the lack of Digital Rights Management (DRM), were points well-taken.

Productivity Sauce • Dmitri Popov

Turn Gists into Articles with GistLog
So, you wrote an article in Markdown and you want to publish it with a minimum of fuss.

Block Ads on Android with Block This

While there are several adblocking apps available for Android, they are not without their shortcomings.

Create Presentations on Android with Slide

Creating slides on an Android device can be as much fun as poking yourself in the eye with a sharp object. Unless you use something like Slide.

Frictionless File Upload with curldrop

Need to move a file from one Linux machine to another? curldrop is just the tool for the job. Written in Python, this simple application is basically a tiny server that accepts file uploads.

MORE ONLINE

ADMIN HPC

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

Optimizing Your NFS Filesystem • Jeff Layton
NFS is the most widely used HPC filesystem. It is very easy to set up and performs reasonably well for small to medium clusters as primary storage.

Fast Help • Laurent Coquerel and Praveen Mosur

Quickassist technology offloads computationally intensive compression and encryption tasks to provide a performance boost for Intel processors.

ADMIN Online

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

Virtual networks with Hyper-V in Windows Server 2016 • Thomas Joos

Microsoft expands Hyper-V's functionality in Windows Server 2016. One interesting new feature is that Hyper-V containers are included in Technical Preview 4 (TP4), which let you virtualize Docker in Hyper-V.

Managing Linux Filesystems

Georg Schönberger and Werner Fischer

Imagine a filesystem as a library that stores data efficiently and in a structured way.

Live Migration of Virtual Machines with KVM
Oliver Frommel

The KVM hypervisor has been a powerful alternative to Xen and VMware in the Linux world for several years.



ing a seven-digit investment in the new company. Both Mache and Karlitschek will serve as the managing directors. Nextcloud is setting up a foundation to manage the project and trademark.

Within 12 hours of the Nextcloud announcement, ownCloud Inc, closed their US offices. ownCloud said in a blog post that the main lenders in the US cancelled their credit, which forced them to close the US office and terminate the contracts of eight employees.

A majority of former ownCloud employees, who quit the company after Karlitschek's departure, have already joined Nextcloud. Karlitschek said Nextcloud welcomes all ex-employees of ownCloud Inc.

Nextcloud said it will honor the existing contracts of those ownCloud customers who want to switch to Nextcloud. Nextcloud will be available in two versions: enterprise and community. The first release of Nextcloud is expected in July. Contributors won't be required to sign a contributor license agreement to contribute to Nextcloud.



HP Enterprise Unveils a Futuristic Computer

Hewlett Packard Enterprise (HPE) is inviting the open source community to collaborate on their ambitious project known as The Machine. Back in 2014, HP introduced the concept of The Machine, which places memory instead of the processor at the core of the computer.

According to HP, "Imagine a computer with hundreds of petabytes of fast memory that remembers everything about your history, helps inform real time situational decisions, and enables you to predict, prevent, and respond to whatever the future brings."

The Machine is based on a new concept that HP calls "Memory-Driven Computing." Memory-driven computing collapses the memory and storage into one vast pool of memory called "universal memory."

Bdale Garbee, veteran Linux developer and HP Fellow, said in an interview that Linux is the primary operating system HP is targeting with The Machine. Teams at HPE are working very closely with the Linux kernel community.

Garbee said that this is the first time a major company has taken a project at such an early stage to the open source community. HPE has released four tools to help the open source community in getting involved with The Machine: 1) Fast optimistic engine for data unification services, 2) fault-tolerant programming model for non-volatile memory, 3) Fabric Attached Memory Emulation, and 4) performance emulation for non-volatile memory bandwidth.



HPE says these tools enable existing communities to capitalize on memory-driven computing in ways that could lead to breakthroughs in machine learning, graph analytics, and event processing.

HPE plans to release their internal research work on The Machine to the open source community.

More Craft. Less Cruft.

Dec. 4 – 9, 2016
BOSTON

Now in its 30th year, LISA is the premier IT operations conference where systems engineers, operations professionals, and academic researchers share real-world knowledge about designing, building, and maintaining the critical systems of our interconnected world.

LISA16 will feature talks and training from:

Jane Adams, Data Research and Acceleration, Two Sigma
Chastity Blackwell, System Administrator
Caskey L. Dickson, Azure SRE, Microsoft
Tyler Fitch, Chef
Alice Goldfuss, Site Reliability Engineer, New Relic
Brendan Gregg, Senior Performance Architect, Netflix
Mitchell Hashimoto, Founder and CEO, HashiCorp
Radia Perlman, EMC Fellow
Theodore Ts'o, Linux Kernel Engineer and Storage Systems Architect, Google

Register by November 10 and save!

Sponsored by the USENIX Association



LISA16

usenix.org/lisa16



Alternative Linux file managers

Managing the Jungle

Working with files and folders is a typical daily task on any computer. We tested four lesser known file managers that significantly simplify handling content on mass storage media. *By Erik Bärwaldt*

Every Linux desktop environment has its own file manager. Outside the mainstream, beyond the kings of the hill (e.g., Dolphin, Nautilus, Thunar, PCManFM, and Nemo), many smaller and usually less popular file management tools access the content of mass storage devices. By targeting special audiences, these tools often have special functions that are missing in the major league tools on Linux.

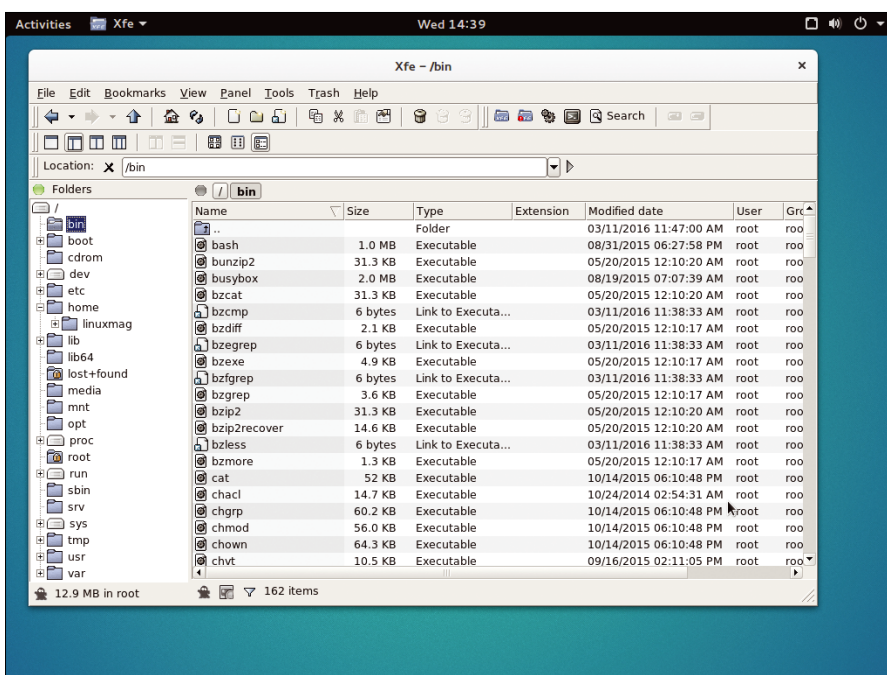


Figure 1: XFE is based on the classic window-splitting of older file managers under Windows with a directory tree displayed on the left.



XFE: For Oldies

The first contender is X File Explorer, or XFE for short, which was specifically designed for aging computers whose performance is no longer sufficient for state-of-the-art desktops. This oldie has been around since 2002, and unlike many other lean file managers, XFE comes with an appealing graphical interface. Most popular distributions provide the software in their repositories, from which you can easily set up the tool via Synaptic, YaST, or Apper.

If you use a distribution whose repositories still don't have the latest version 1.41, you can pick up the source code (as well as pre-built packages for some distros) from the project website [1]. XFE is based on the FOX Toolkit; it runs very quickly with frugal resource requirements, has no other dependencies, and thus remains compatible with all common interfaces.

Visually, XFE is strongly reminiscent of older file managers on Microsoft operating systems. On the left, you will find the tree view with the folder hierarchy of the mass storage device, and in the larger area on the right are the files of the active folder. The display areas, an address bar, a toolbar, and a menubar are arranged horizontally (Figure 1).

To begin, you should configure the software to suit your needs in the *Edit | Preferences* menu and take a look at the list of external programs for viewing and modifying files. Under the *Programs* tab, you will find applications that not every distribution has on board out the box (Figure 2).

As its own viewers, XFE includes two display programs for text and image files, XFW and XFI: XFW allows supports editing. In addition the Archive Manager XFP, which shows

the contents of packed files. For all other tasks, XFE reverts to the default programs from the Linux treasure trove.

The file manager also lets you customize the user interface and set the display options. For example, if you check the *Hidden files* box in the Panel menu, you will see normally invisible files and directories. If the interface is too crowded, turn off individual columns by unchecking the respective options in the View menu.

The software performs all changes in real time, removing the need to relaunch.

You can also change the list views, allowing, for example, a Midnight Commander-style two-pane view. Depending on the selected view, the program automatically modifies the menu line and adapts all functions in a context-sensitive way.

Keyboard Controlled

XFE also supports drag and drop, both within the file manager and in the respective work environment. The lean file manager thus functionally blends in well with various interfaces.

Experienced users in particular prefer to use time-saving keyboard shortcuts when working with files and folders. XFE can also be controlled with keyboard shortcuts that you can customize to suit your requirements. The *Edit | Preferences | Key Bindings* menu has a dialog with numerous predefined shortcuts. Clicking the *Modify key bindings* button opens a list of predefined shortcuts, which you can then modify. XFE also takes into account its own menu items (Figure 3).

Archive

XFE reads and unzips archives, relying on external software to do so. However, Xarchiver, the default archiving tool, is usually

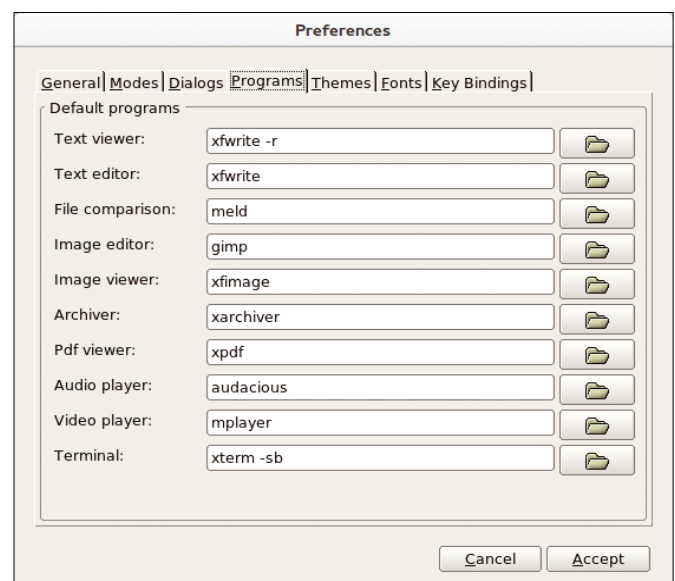


Figure 2: In the Programs tab, you can define the default tools for certain file types.

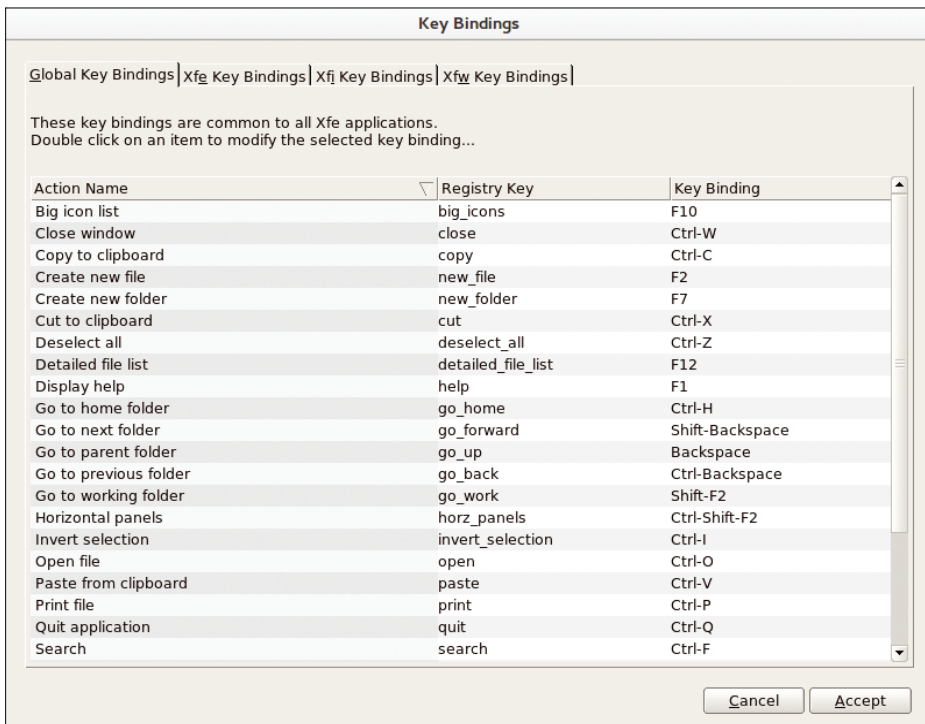


Figure 3: Numerous shortcuts allow almost complete keyboard-based control of XFE.

missing on major desktops like KDE, which have their own desktop-specific tools.

This explains why XFE pops up a routine in a small window asking which tool to use when you click on an archive. You can then enter the command for the desired archiving program in the appropriate command line. This sets up the tool in question as the tool to call for working with archives of the format currently being accessed. Additionally, you can enter this software in the Preferences dialog as the general-purpose archiving tool.

Sunflower: The Clone

Much like XFE, Sunflower's lean design also makes it useful for older hardware, but without appearing jaded. The graphical interface is oriented on the well-known Norton Commander of the DOS era, both functionally and in terms of the window structure. Only Fedora and OpenMandriva have Sunflower in their repositories. Because the Sunflower project has binary packages of the latest version of the file manager for an exceptionally large number of Linux distributions, you should get the latest version 0.3 of the software from the project homepage [2].

After a fast-paced start, a two-pane interface appears, supplemented with a menubar at the top and a command line at the bottom of the screen. Unlike XFE, Sunflower has

no toolbar, and an overview of the predefined function keys, as in Norton Commander or Midnight Commander, is also missing.

To compensate for this, the tab structure of the software lets you open multiple Windows simultaneously and even move them within the hierarchy. For example, you can click on an icon in either the left or right pane to enable a terminal that appears in its own tab. Pressing the Tab key lets you toggle quickly between the two display areas (Figure 4).

Sunflower is somewhat less consistent in terms of controls than XFE; you need to unlock some functions in a dialog. Selecting *Edit | Preferences* takes you to the options, which are clearly divided into groups. You will find the most important options in the *Item List (Hidden files | Show hidden files)* group and the *Plugins* group (*Archive Support, Basic rename options,*

and *Basic find file options* entries).

These options directly influence the capabilities of the individual files and folders. Like XFE, Sunflower relies on third-party programs for many editing functions. Therefore, you will want to choose a powerful editor for text files in the drop-down box of the Edit section under the *View & Edit* group of the Preferences dialog (Figure 5). You can even choose LibreOffice as your default editor, which is particularly useful if you often view and edit files in ODF format.

Sunflower supplements the file-specific context menu, which can be opened by right-clicking the respective file, to reflect the




Figure 4: Simply click on the terminal prompt icon to open a terminal in the current directory.



enabled options. However, the tool only modifies the display of folder and file hierarchies in the main window after you press *Save* in the Preferences dialog.

Vifm: The Spartan

As the unusual name suggests, Vifm [3] is a program for users who often work with the venerable editor Vi or its slightly more convenient counterpart, Vim. Vifm, with an ncurses interface, is accordingly the most spartan tool. Its keyboard shortcuts are similar to those used in the editor and therefore allow for very quick navigation in large sets of data.

Vifm is in the software repositories of almost every major distribution. Also, the project provides the file manager on its website as a source package and in the form of prebuilt binary packages for various distributions, all licensed under GPLv2. The setup process does not create a launcher in the application menu, so you launch it from a terminal window.

Vifm also uses a two-pane view mode and can be controlled with the keyboard alone (Figure 6). For begin-

ners who are not so knowledgeable of the operation and configuration of the Vi editor, Vifm provides fairly detailed documentation [4], along with a guide to all the keyboard shortcuts.

The file manager offers the usual functions, an option for renaming multiple files in one step, and freely definable key-

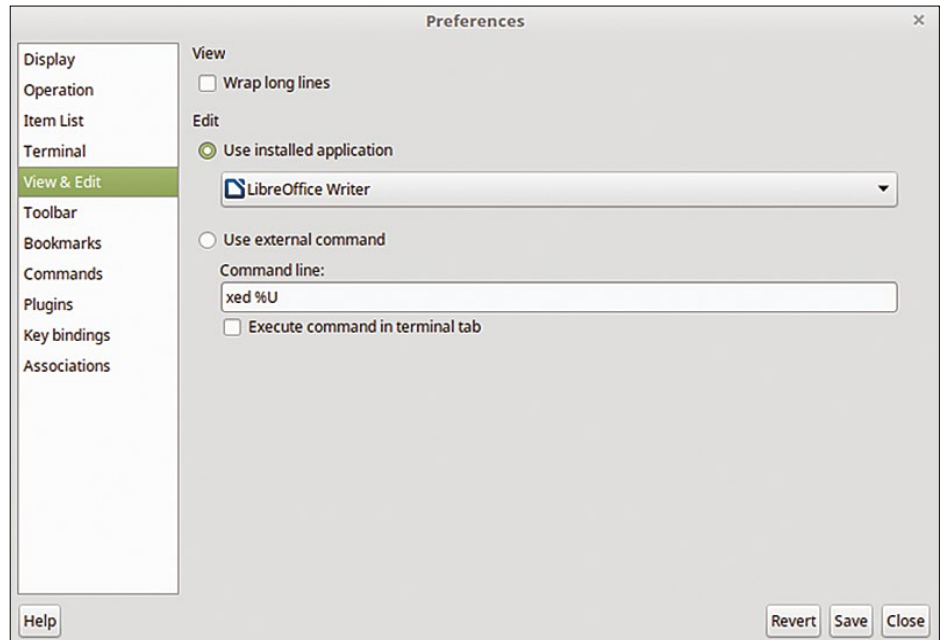


Figure 5: Sunflower also offers many options for flexibly customizing the file manager to your own needs.

BACKUP AND BARE METAL RECOVERY FOR LINUX SERVERS

Without good bare metal recovery, what good are backups? In 1999 Microlite delivered the first backup and bare metal recovery (including P2V) software for Linux, and we've been improving it ever since.

Try a free 60 day demo today.

Learn why thousands of sites around the world protect their Linux servers with *BackupEDGE*.

New Partners Welcome

<http://www.microlite.com/welcome.html>

Supporting the industry's largest variety of storage resources, including...

Amazon S3
Google Cloud Storage
CIFS / SMB / NFS Shares
FTP / FTPS Servers
Blu-ray Disk BD-RE
USB Disk / Flash
Traditional Tape & More

MICROLITE
CORPORATION

Developers of BackupEDGE

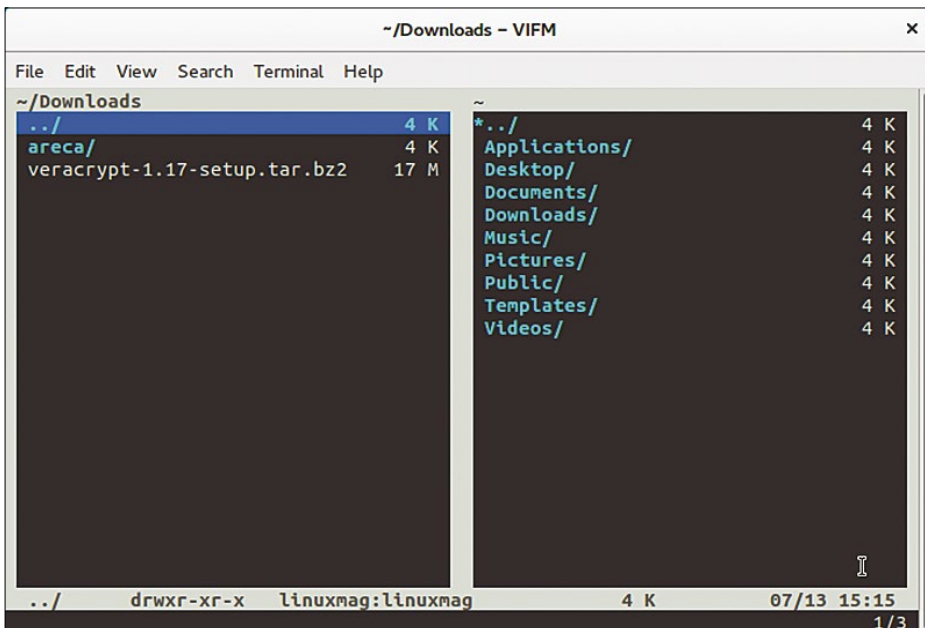


Figure 6: The console-based Vifm is visually austere, but you can navigate efficiently with just the keyboard.

board shortcuts. An undo/redo function lets you roll back commands you accidentally executed or run them again if needed. Using external programs, the file manager also lets you look at files of various formats, without having to change to the respective application.

Much like graphical file managers, Vifm easily adjusts to your needs: You can use the program in a single pane, adjust the display columns to suit your own needs, and define the colors within the framework of the ncurses interface. Extensive

launch on various BSD derivatives.

The program, which has been under development since 2007, can be found in the repositories of all major distributions. However, if you prefer to have the latest version, you can pick up a binary for various distributions from the project web site. You will also find portable binaries for distributions that do not support RPM or DEB.

In some cases the installation routine does not put a launcher in the menu; however, you call `./doublecmd` in the terminal after you install

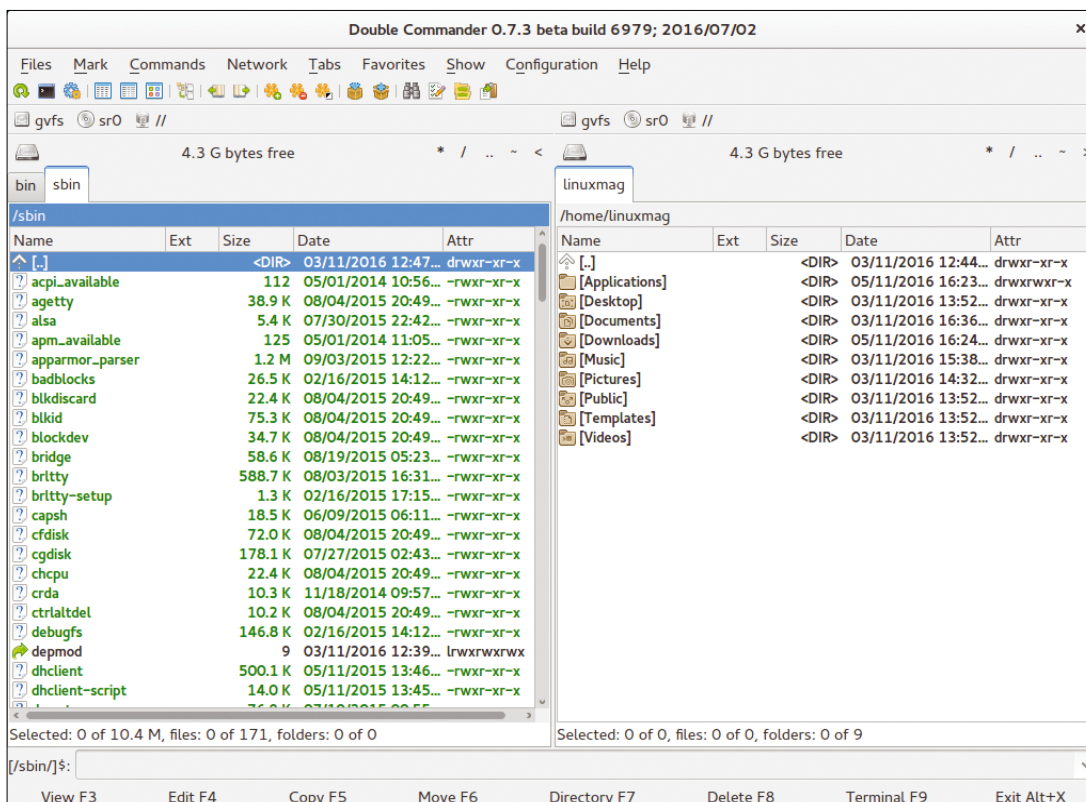


Figure 7: The Russia-born Double Commander offers a very extensive user interface.

search and filter options round out the offerings.

For efficient navigation of a very large number of files, Vifm offers a bookmark feature that lets you quickly jump to selected files in nested folder hierarchies. Finally, it offers a command-line option that lets users execute shell commands from within the file manager, which makes changing to a second terminal tab unnecessary in many cases.

Double Commander

Double Commander [5], which comes out of Russia and is licensed under GPLv2, offers a feature set without equal. The program, written in Object Pascal, comes in 32- and 64-bit flavors as well as in various packages for Gtk2- and Qt4-based interfaces. Additionally, the cross-platform application will

launches into a two-pane view, like most programs for managing files, but it is anything but ordinary with its huge number of controls (Figure 7).

In addition to the obligatory menu bar and a button bar for fast access to important functions, you will also see a bar displaying the drives on the system, which lets you click to change from one drive to another. Double Commander can also handle network drives mounted via SMB/



CIFS and FTP, so they also appear in the view. Another line above the list view shows the active drive.

On the right side of the panels are some buttons that change the view. At the bottom of the program window a bar with function key assignments give you access to menu items at the push of a button or with a keyboard command.

Using the *Show | Tree View Panel* menu, you can switch to a tree view that Double Commander displays on the left in the program window while simultaneously reducing the size of the two main panels. This view significantly speeds up navigation in large folder hierarchies.

In the list view of mass storage content, the program displays the permissions for each file with different color highlighting; the two panes can thus become very colorful at times. Along with the main editing functions, Double Commander also supports tabs and right-click context menus for files.

Viewers and Players

The software uses third-party programs, such as Kate, Gwenview, and Okular on KDE for displaying file content. Double Commander also has its own viewers and editors, which you enable by pressing the function keys. For example, F4 starts a very simple integrated text editor and F3 opens a viewer that also displays hex and binary files.

Double Commander supports various archive formats, including ZIP, TAR, GZ, LZMA, and RAR archives, as well as DEB and RPM. The software even offers a directory synchronization tool; you will not need an external program for simpler data retrieval tasks. A batch option rounds off the functionality and enables fast renaming of multiple files in a single run. For quick navigation in large databases, the software lets you create bookmarks that support rapid content discovery.

Configuration

Because of its huge feature set, careful configuration of Double Commander is recommended. You can reach the configuration window via the *Configuration | Options* menu item.

In the left pane, a configuration tree divides options into groups, and the corresponding context-sensitive settings are displayed in the right window. Here, you specify display functions and external programs, define or modify keyboard shortcuts, change the appearance of the software, or add new function buttons to the button bar, among other things (Figure 8).

On Windows, the proprietary, commercial Total Commander has gained a large fan base as an all-round file manager. Its numerous extensions can be used in several formats for Double Commander and are available for download from the corre-

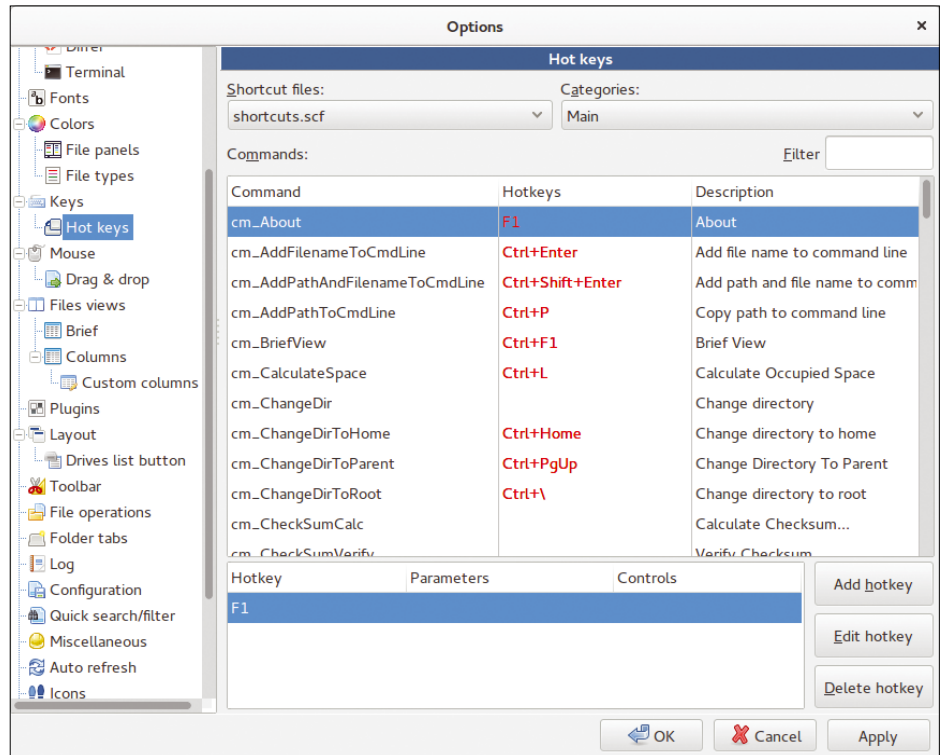


Figure 8: Double Commander does not lack configuration options.

sponding plugin page [6]. Plugins set up the file manager to perform tasks beyond plain vanilla file management.

Conclusions

You would be hard pressed to find another operating system that offers as many programs for file management as Linux – the choice is accordingly difficult. Resource-saving contenders like XFE and Sunflower are excellent for use on older hardware, regardless of the desktop environment.

Vifm, in contrast, is only likely to find favor with users who are already familiar with the Vi(m) text editor: If you do not want to do without mouse control and a graphical interface, you will need to look elsewhere. If you are looking for a multitasking file manager that can even handle synchronization tasks and include network drives, your best option is Double Commander. ■■■

INFO

- [1] XFE: <http://roland65.free.fr/xfe/index.php?page=home>
- [2] Sunflower: <http://sunflower-fm.org>
- [3] Vifm: <https://vifm.info>
- [4] Vifm documentation: <https://vifm.info/manual.shtml>
- [5] Double Commander: <http://doublecmd.sourceforge.net>
- [6] Total Commander plugins: <http://www.ghisler.com/plugins.htm>



Find us on
Facebook

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

SUBSCRIBE NOW



6
issues
per year!

The New IT

New tools, new threats, new technologies...Looking for a guide to the changing world of system administration?

shop.linuxnewmedia.com

AND SAVE 30%

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, Solaris, and popular varieties of Unix.

**REAL-WORLD
PROBLEMS
SOLVED!**

ADMIN

Network & Security

shop.linuxnewmedia.com

Tool tests on the fast track *By Uwe Vollbracht*

TOOL TIPS

htop 2.0.1

Function: Clear-cut process manager

Source: <http://hisham.hm/htop>

License: GPLv2

Alternatives: Top, ps

```

Terminal - vollbracht@LMLab: ~
CPU: [|||||] ] Tasks: 115, 146 thr; 1 running
Mem: [|||||] ] Load average: 0.42 0.27 0.14
Swp: [|||] ] Uptime: 00:50:58

PID USER PRI NI VIRT RES SHR S CPU% MEM% TIME+ Command
1 root 20 0 6584 2884 2036 S 0.0 0.6 0:02.50 /sbin/init splash
1722 20 0 56584 6744 4656 S 0.0 1.3 0:00.11 /usr/lib/udisk
1729 20 0 56584 6744 4656 S 0.0 1.3 0:00.00 /usr/lib/udisk
1727 20 0 56584 6744 4656 S 0.0 1.3 0:00.00 /usr/lib/udisk
1725 20 0 56584 6744 4656 S 0.0 1.3 0:00.00 /usr/lib/udisk
1723 20 0 56584 6744 4656 S 0.0 1.3 0:00.00 /usr/lib/udisk
1665 20 0 65484 5052 4140 S 0.0 1.0 0:00.02 /usr/lib/upower
1674 20 0 65484 5052 4140 S 0.0 1.0 0:00.00 /usr/lib/upower
1673 20 0 65484 5052 4140 S 0.0 1.0 0:00.00 /usr/lib/upower
1429 vollbracht 20 0 52624 6112 5728 S 0.0 1.2 0:00.06 /usr/bin/gnome
1548 vollbracht 20 0 52624 6112 5728 S 0.0 1.2 0:00.00 /usr/bin/gnome
1536 vollbracht 20 0 52624 6112 5728 S 0.0 1.2 0:00.00 /usr/bin/gnome
1524 vollbracht 20 0 52624 6112 5728 S 0.0 1.2 0:00.00 /usr/bin/gnome
1430 vollbracht 20 0 52624 6112 5728 S 0.0 1.2 0:00.00 /usr/bin/gnome
1423 vollbracht 20 0 6320 2884 2172 S 0.0 0.6 0:00.06 /lib/systemd/sd-pam
1424 vollbracht 20 0 7548 956 0 S 0.0 0.2 0:00.00 /usr/sbin/cups
1296 20 0 14720 3592 2788 S 0.0 0.7 0:00.02 /usr/sbin/cups
  
```

Many admins turn to Top to monitor the usage of their computers. The classic tool presents a dynamic view of the processes running on your system and the resources needed. Htop, which made a giant leap from version 1.0.3 to 2.0.0 in early February and has since published the first bugfix release 2.0.1, proves to be a good alternative. The ncurses interface not only allows horizontal and vertical scrolling but also lets users manipulate processes by pressing function keys.

The developers added a few more interesting new features to the tool. Htop is now platform independent and thus is available not only for Linux but also for FreeBSD, OpenBSD, and OS X. Also, the tool relies on ncurses 6, which supports scrolling with the mouse wheel. You can install from the sources on systems that use the older ncurses version 5, but mouse wheel support does not work in this case.

★★★★★ The new release, finally available after a wait of more than four years, is a huge success for the htop developers. Not only will it satisfy system administrators on Linux machines, but finally also BSD and OS X users, too. ■■■

Metastore 1.1.0

Function: Metadata store files

Source: <https://github.com/przemoc/metastore>

License: GPLv2

Alternatives: None

```

Terminal - vollbracht@LMLab: ~
./bilder/CIMG0184.JPG: no difference
./bilder/CIMG0176.JPG: no difference
./bilder/CIMG0185.JPG: no difference
./bilder/CIMG0177.JPG: no difference
./bilder/CIMG0186.JPG: no difference
./bilder/CIMG0170.JPG: no difference
./bilder/CIMG0178.JPG: no difference
./bilder/CIMG0187.JPG: no difference
./bilder/CIMG0171.JPG: no difference
./bilder/CIMG0179.JPG: no difference
./bilder/CIMG0180.JPG: no difference
./bilder/CIMG0188.JPG: no difference
./bilder/CIMG0172.JPG: changing metadata
./bilder/CIMG0172.JPG: changing group from ww-data to vollbracht
./bilder/CIMG0172.JPG: changing mode from 0600 to 0644
./bilder/CIMG0181.JPG: no difference
./bilder/CIMG0189.JPG: no difference
./bilder/CIMG0190.JPG: changing metadata
./bilder/CIMG0190.JPG: changing group from ww-data to vollbracht
./bilder/CIMG0173.JPG: no difference
./bilder/CIMG0182.JPG: no difference
./bilder/CIMG0191.JPG: changing metadata
./bilder/CIMG0191.JPG: changing mode from 0600 to 0644
vollbracht@LMLab: ~$
  
```

Because the git versioning tool does not cover all metadata for files, directories, and links, resourceful developers wrote a tool to handle these tasks: metastore stores information in separate files. Users call the tool with the `-s` option and pass in the path to the source. The results end up in the `.metadata` file in the current folder.

By default, Metastore is restricted to details of the owner, group, access privileges, and extended attributes (`xattr`). If you want to cover the last modification time, you need to add `-m` to the command line. The `-c` parameter tells Metastore to compare the stored information with the current version of the data and output any changes to the standard output. The `-a` option lets users reset the metadata. To cover removed or empty directories, you would add `-e`.

★★★★★☆☆ Metastore does a great job, not only in Git repositories but also in providing rudimentary protection against intruders, as well as in archives in which the tool ensures that all metadata match after unpacking on the target system. ■■■

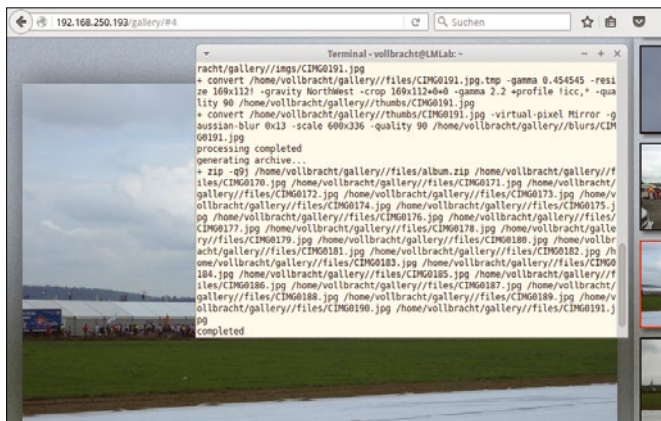
fgallery 1.8.1

Function: Create simple web galleries

Source: <https://www.thregr.org/~wavexx/software/fgallery>

License: GPLv2

Alternatives: booh, AGG



If you want to show your holiday snaps outside of Facebook, Flickr, and other services on the web and have your own web space, your best bet is a photo album generator. Many solutions rely on dynamic content and therefore require a database back end. For a more resource-friendly approach, try `fgallery`: All it needs to open an exhibition is `ImageMagick`, the `JSON:PP` Perl module, and `exiftran`.

The Perl script does not just ensure the correct orientation of the images, it also sorts them in chronological order and by timestamp. `fgallery` also scales the thumbnails. Their image quality (default: 90 percent) is defined by the `--quality` parameter. If the collection contains many portraits, `-f` can optimize thumbnail cropping. Optionally, you can use `-j` to parallelize the process of creating snapshots. The result is a static gallery that contains some JavaScript elements for navigation.

★★★★☆☆ `fgallery` is fast, its use is intuitive, and the results are okay, but users who want a more sophisticated design, a more complex preview, or embedded videos might want to look for an alternative. ■■■

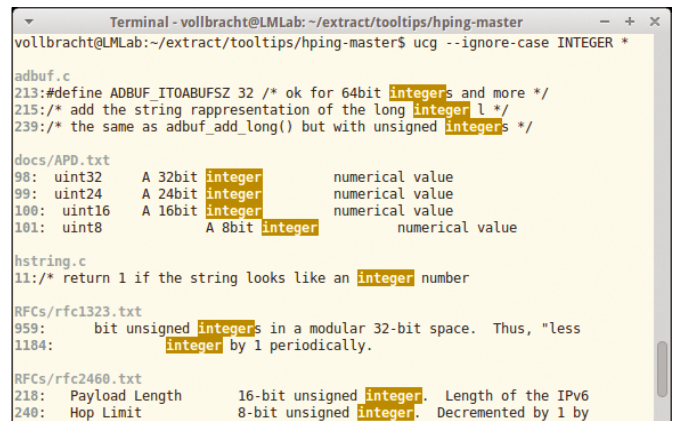
UniversalCodeGrep 0.2.1

Function: Fast grep for long source code

Source: <https://gvansickle.github.io/ucg>

License GPLv3

Alternatives: grep, ack



Working through large amounts of source code by keywords or regular expressions can take some time. The `UniversalCodeGrep` (`ucg`) tool shortens the wait. The developers implemented the program in the C++11 programming language and leveraged its support for parallel operations. `UniversalCodeGrep` also relies on the PCRE libraries and their just-in-time compilation for regular expressions, which also contributes performance gains.

For a simple search, you just feed `ucg` a search pattern and the target directory or individual file. By default, the tool takes upper- and lowercase into account, but you can prevent this with `--ignore-case`. If you want to avoid interpreting numbers numerically in a search pattern, enabling `--literal` will help. The `--word-regexp` parameter ensures that a regular expression only accepts whole words as hits, which reduces the result set.

By default, `ucg` also looks at files in subdirectories, which you can stop by stipulating `--no-recurse`. Individual folders can be skipped with `--ignore-dir`. To leverage the benefits of parallelization, mentioned previously, set `-j` with the number of processes you want `UniversalCodeGrep` to launch for the search.

★★★★☆☆ If you believe the measurements published on the project website, the speed benefits are minor – on our lab system, `grep` even had its nose in front. Nevertheless, `ucg` is worth looking at because of its interesting options. ■■■

hping 3.0.0

Function: Analyze network traffic

Source: <http://wiki.hping.org>

License: GPLv2

Alternatives: Nmap

```
Terminal - root@LMLab:/home/vollbracht
using enp0s17, addr: 192.168.250.193, MTU: 1500
Scanning 192.168.250.96 (192.168.250.96), port 22,25,110,145,80,443,6777,53
8 ports to scan, use -V to see all the replies
-----+-----+
|port| serv name | flags |ttl| id | win | len |
-----+-----+
 25 smtp      : ..R.A... 64 45288  0 46
 53 domain   : ..R.A... 64 45544  0 46
110 pop3     : ..R.A... 64 45800  0 46
145          : ..R.A... 64 46056  0 46
6777        : ..R.A... 64 46312  0 46
All replies received. Done.
Not responding ports: (22 ssh) (80 http) (443 https)
root@LMLab:/home/vollbracht# hping -V -1 192.168.250.96
using enp0s17, addr: 192.168.250.193, MTU: 1500
HPING 192.168.250.96 (enp0s17 192.168.250.96): icmp mode set, 28 headers + 0 data bytes
len=46 ip=192.168.250.96 ttl=64 id=59574 tos=0 iplen=28
icmp_seq=0 rtt=3.1 ms
len=46 ip=192.168.250.96 ttl=64 id=59575 tos=0 iplen=28
icmp_seq=1 rtt=1.8 ms
^C
--- 192.168.250.96 hping statistic ---
2 packets transmitted, 2 packets received, 0% packet loss
```

The name suggests close kinship with ping – but hping goes well beyond the functionality of the legacy tool (i.e., sending ICMP packets). In addition to TCP and UDP, hping can send RAWIP. It also has a traceroute mode and is therefore useful to search for bugs in networks and analyze or discover security vulnerabilities. The tool works on Linux, various BSD variants, Solaris, Windows, and OS X.

Users control hping in the shell. Parameters nest between the hping3 command and the hostname of the target system. If you want to stress the target, you can increase the number of packets per second to 10 with --fast; or even to 100 packets per second with --faster. If this is not sufficient, fire --flood at the target to rev up to maximum speed. For each protocol, you can adapt the packets sent. In addition to flags such as SYN, ACK, or RST, users can set the TCP sequence number or adjust the size of the TCP window. The --scan switch expects one or more ports as parameters; --traceroute switches to traceroute mode.

The new version 3 is compatible with the previous hping2. New features include the built-in TCL interpreter. Users call hping3 with no other information and are taken to an interactive shell. Simple application examples can be found in the wiki.

★★★★★ Hping is a powerful analysis tool that should be part of any admin toolbox. The current version 3 extends the performance to include practical scripting functions. ■■■

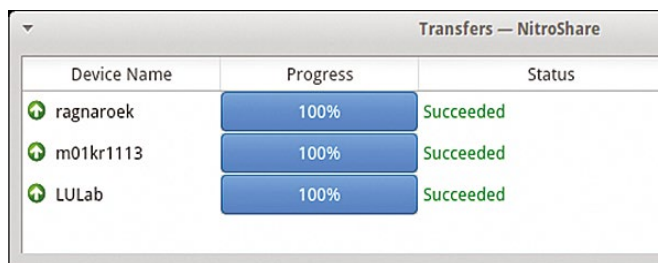
NitroShare 0.3.1

Function: Exchanging files on a LAN

Source: <https://nitroshare.net>

License: MIT

Alternatives: Transfer on LAN, Dukto



If you are worried about setting up Samba or NFS on a home network, you will probably resort to USB sticks or cloud services to exchange data between several computers. This also works across platforms and does not require any lengthy server configuration.

NitroShare provides another option for this purpose. The P2P solution for exchanging data on a LAN is implemented in C++; it relies on Qt and works on Linux, Windows, and OS X. On all systems, the installation is manual. After launching, a NitroShare icon nests into the toolbar of the graphical interface.

Pressing the right mouse button pops up the context menu for users to initiate the data transfer. You can then select individual files or entire directories; the program finds other active NitroShare clients on the local network. A dialog box displays the names and operating system. After selecting the desired target system, the data migrates across the network. A status window informs you of the progress. By default, the incoming data ends up in the user's download directory.

Mutual authentication takes place between the NitroShare clients; communication on the network is unencrypted. NitroShare relies on TCP. The broadcast search for clients uses port 40816, and the transfer takes place on 40818.

★★★★☆ NitroShare should only be used on trusted networks, but if this sounds like your home network, then the tool offers a convenient way of moving data from point A to point B – regardless of the operating system. ■■■



"It looks like the drone industry has chosen their go-to event!" —Robert Rodriguez, President of the Society of Aerial Cinematography

InterDrone

The International Drone Conference and Exposition

September 7-9, 2016

Paris Hotel, Las Vegas

www.InterDrone.com



The Largest Commercial Drone Show in the World!

InterDrone is Three Awesome Conferences:

Drone TECHCON

For Drone Builders

Content will focus on advanced flying dynamics, chips and boards, airframe considerations, hardware/software integration, sensors, power issues, and software development.

Drone ENTERPRISE

For Flyers, Buyers and Drone Service Businesses

Classes focus on enterprise applications such as precision agriculture, surveying, mapping, infrastructure inspection, law enforcement, package delivery, and search and rescue.

Drone CINEMA

For Flyers Engaged in Aerial Photography and Videography

Class content includes drone use for real estate and resort marketing, action sports and movie filming, newsgathering — any professional activity where the quality of the image is paramount.



125+ Exhibitors • 110+ Classes and Panels • InterDrone Film Festival • Women In Drones Luncheon



Demos • Keynotes • 100+ Industry Expert Speakers

Registration Open!

A BZ Media Event



Keep watch over your server system with Webmin

STEADY HAND

Every leading Linux distribution offers some kind of built-in GUI management tool, but the world still has room for a flexible, distribution-independent tool like Webmin. *By Marco Fioretti*

Webmin [1] is a web-based interface for system administration of Unix-like systems. The first part of this tutorial explains what Webmin is and when and why you might want to use it. The second part describes how to set up Webmin and use it for some basic administration tasks. Finally, I'll show you some advanced Webmin features.

If you want to configure your Linux computers in the best possible way or be

able to rescue them when something goes wrong, there are no shortcuts: You need to know the basics of the command line and understand how to edit shell scripts and configuration files manually. However, just because you use the command line doesn't mean you like it, and this is where interfaces like Webmin can help.

Of course, if Webmin were nothing more than a mouse-first version of some system administration utilities, it would be much less useful now than

when it first appeared. These days, every desktop-oriented Linux distribution provides a suite of graphical management tools. However, Webmin still provides some benefits for today's networks.

First, Webmin is "network transparent." It doesn't matter whether you're actually sitting in front of the computer you need to manage or half a planet away, if you configure your firewall properly, you will be able to manage

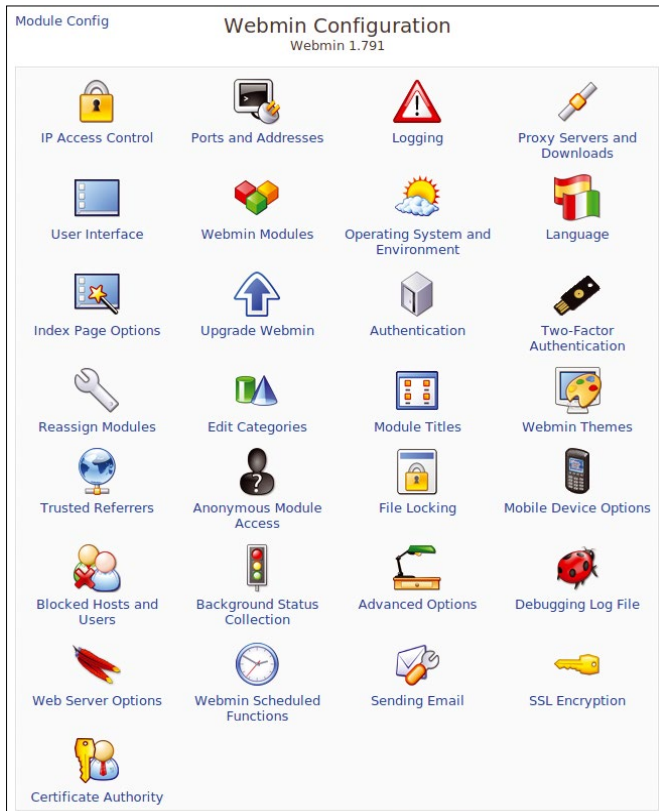


Figure 1: The Webmin Configuration page gives an idea of how complex and flexible this tool can be, but you can do lots of useful stuff with it.

your computers in the same way. Plus, you can operate Webmin through any browser, even if you cannot connect via SSH. Even in cases where Webmin is little more than a web-based text editor for configuration files, it is still worth using.

Second, running a personal server is becoming a more common task. If you care about your blogging, photographs, and email, you really should locate these services on a server that is actually yours. Many Linux distributions that are optimized for such scenarios don't run any desktop environment – nor should they. With Webmin, you can manage many parts of your servers through the same interface you use on your own desktop.

Third, Webmin lets you delegate. If the basic management tasks are possible from any browser with a

Linux that Webmin supports, plus Solaris, FreeBSD, and HP/UX.

Finally, for all the reasons above, Webmin can also be a great (self-)training environment. You can configure programs and tasks with Webmin and then

few clicks, you can ask colleagues, friends, or relatives, who might not even know Linux, to look after your computer when you're sick or traveling.

Fourth, Webmin hides differences among distributions. Things have improved since the 1990s, but in many cases you still cannot reuse exactly the same script or command across distributions. Webmin does its best to hide all that. The Webmin form to create a user, for example, will always be basically the same for all the versions of

study how it modified the configuration files of your system, or you can set up Webmin in a computer class to let your students do the same thing.

Gotchas and Limitations

If you carefully read the previous paragraphs, you have probably noticed some limitations in what Webmin can do. The limitations arise from the fact that Webmin, no matter how convenient it is, is ultimately just a configuration and management assistant.

To use Webmin, you must first install Linux and possibly other programs (not to mention Webmin itself) the traditional way, but you cannot install Webmin on systems where you do not have root access.

Overview and Initial Configuration

Technically speaking, Webmin is a set of Perl scripts that can be extended with many modules, and the configuration includes a small web server. At a low level, all Webmin does is translate (or reapply) your commands as changes in the standard, textual configuration files used to configure a Linux system.

The good thing is that, although Webmin directly updates the usual configuration files, it doesn't overwrite all their contents. Inside those files, Webmin only changes the parts that you tell it to change and leaves the rest alone. In

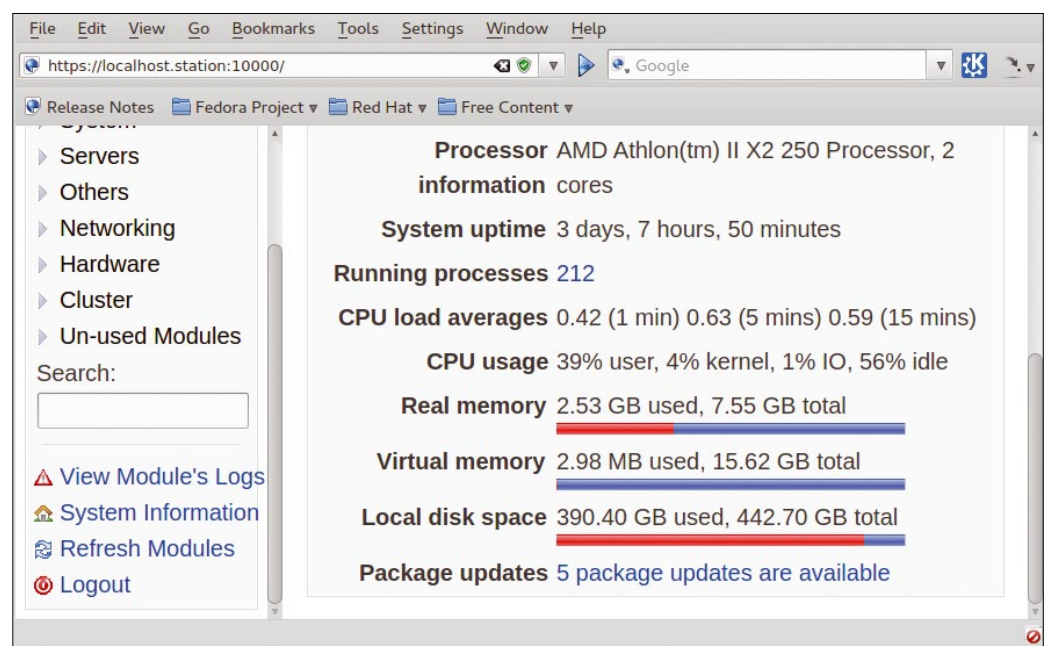


Figure 2: Part of the initial Webmin page, with the default theme. The main panel shows system performance and status, whereas the left sidebar groups all the available functions in several areas.

other words, almost all the careful changes and customizations you did manually before you started using Webmin should be unchanged. In general, you should be able to mix Webmin and good old command-line typing as you wish.

Installing and Configuring Webmin

Webmin is not generally included in the package repositories of most Linux distributions. To install it, you must download the files from the website in the format most compatible with your distribution, then follow the instructions. If you use a Linux distribution based on the .rpm or .deb packaging formats, just download the corresponding binary package and install it with your usual software manager. You might have to force your package manager to ignore dependencies or install the dependencies manually yourself, but usually this does not create any problem. On some distributions, you have to start Webmin manually after the installation by issuing a command like the following as root:

```
service webmin start
```

At that point, you will be able to log in to Webmin by pointing your browser to the address `https://localhost:10000`. If

you are accessing Webmin from another computer, replace `localhost` with the hostname or IP address of the system you want to access. If you want to manage your computer remotely, you will also have to tell the firewall to forward incoming traffic to TCP port number 10000.

Because Webmin generates its own SSL certificate when you install it – to use encrypted (HTTPS) connections – you’ll also get the standard warning from your browser that the certificate cannot be trusted; in this case, you may ignore the warning.

Figure 1 shows what Webmin looks like when you start using it. You can change the look and feel by choosing another of the (few) themes included, but be warned: None of the themes will win a web design award anytime soon (Figure 2). However, they all do the job, and for the most part, Webmin themes also work in text browsers that support frames.

Webmin is well documented, and you should spend some time browsing the resources listed in the *System | System Documentation* section. The search box at the bottom of the sidebar will help you to find the functions you need and assist you with installing the corresponding packages if they aren’t already in your computer. Most Webmin windows include links to relevant instructions or

chapters of the online documentation [2] to facilitate their usage.

Modules

You can augment Webmin functionality by adding modules, many of which are included in the basic package. To see the modules, click on *Un-used Modules* in the sidebar. To manage your modules, or install new ones, click on *Webmin Modules* in the configuration panel shown in Figure 1. This opens a relatively intuitive form that lets you download and install modules, control who can use them, and generate clones. A Webmin module clone is a separate instance of the module that runs with different settings. To learn more about modules, visit the corresponding documentation page [3].

Basic Configuration

The first things you should learn to do with Webmin are, of course, the first things you would do on any Linux computer.

Go to *Hardware | System Time*, for example, to set system and hardware time or to synchronize your computer to a time server.

Click on *System | Bootup and Shutdown* to get a list of which system services are available and to stop or automatically start the services at every boot. The current status of system services is also visible under *Others | System and Server Status*.

Installation and updates of software, including new versions of Webmin, happen in the *Software Packages* and *Software Package Updates* panels in the *System* section of the sidebar (Figure 3). These panels are more primitive than you are likely accustomed to if you work with Synaptic or similar tools, but they are not difficult to use. Again, they work pretty much in the same way, no matter what Linux distribution is underneath.

Webmin also includes a basic file manager at *Others | File Manager*, and you will even find backup utilities. The File manager looks quite bare, but it has all the main functions for a program of this kind, including a means for creating zipped archives and the possibility for getting files directly from any Internet server.

As far as backups go, Webmin includes two interfaces with different pur-

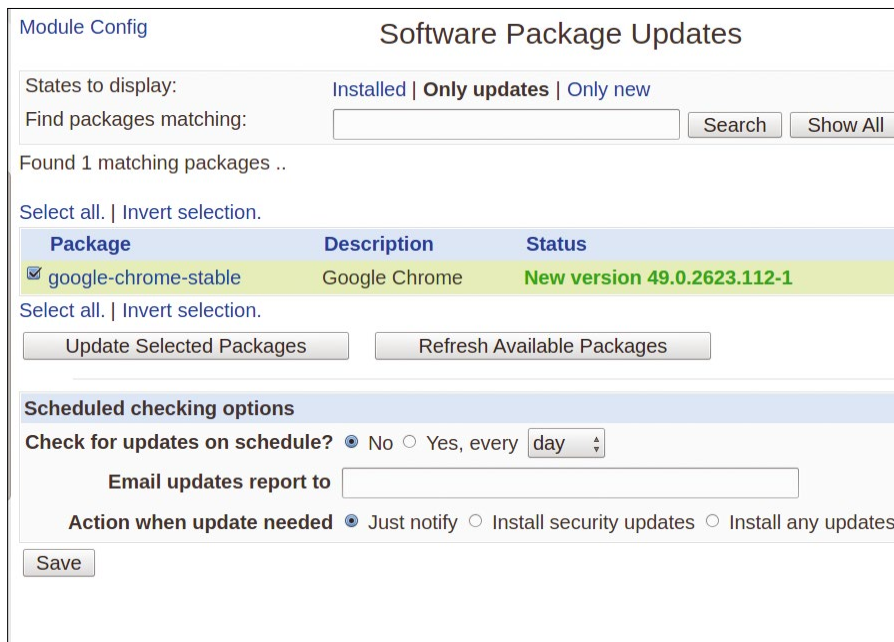


Figure 3: The Software Package Updates panel is a good example of what working in Webmin looks like. The appearance might be dull, but all pages pack a lot of options in a clear way, and the interface even displays the full output of many commands.

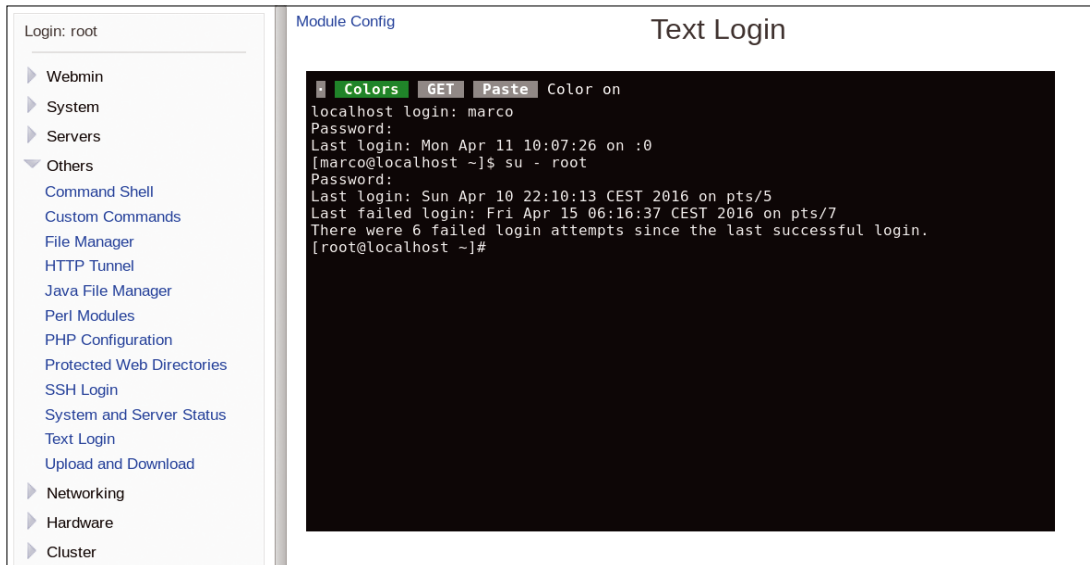


Figure 4: A built-in Ajax terminal gives you a command prompt right inside your browser.

poses. The one in *Webmin | Backup Configuration Files* is aimed mainly at saving all or parts of the configuration files you edit via Webmin, plus its own. In that panel, you can manually add other files to back up, restore from previous backups, or schedule regular backups.

The window under *System | Filesystem Backup* is reserved for backups of full, generic directories in your filesystem. In both interfaces, you can tell Webmin to schedule the backups at regular intervals and save the resulting archive to an external server via FTP or SSH.

Linux or Webmin Users? Both, of Course

If you go to *System | Users and Groups*, you will see a list of all the standard Linux/Unix accounts on your systems, which should not be confused with Webmin accounts. This part of Webmin lets you do almost everything you need with system users and groups: You can create or delete them, move users among groups, change user privileges, and so on. If needed, you may even click on *Run batch file* to manage many users at once from an uploaded or local text file. Click on *Instructions and batch format* to see the exact (simple) syntax you should use in the text file. The one thing about system users that you cannot manage in this panel is the disk space quotas for the `/home` partition. That function has its own window under *System | Disk Quotas*.

Webmin doesn't just help you manage the users of a Unix/Linux system. You

can also create users for Webmin itself – or at least its own, separate rules about what system users may do if they log into Webmin. This functionality is handled in the *Webmin | Webmin Users* page. A whole section of the online documentation [4] explains how to create Webmin users and how to specify which users will

having to use the root account.

The Webmin authentication code that makes this possible needs the Perl module *IO:Pty*, which might not be installed by default. Most distributions have it in their standard repositories, however, so if Webmin complains, just install it and try again.

have access to which functions or modules.

You might not need bother about Webmin users at all, if you are the only user of your system, except for one thing: click on *Unix User Authentication*, select *Allow users who can run all commands via sudo to login as root*, then save and log off. At that point, you will be able to log in and work in Webmin using the account of any user with sudo privileges, instead of

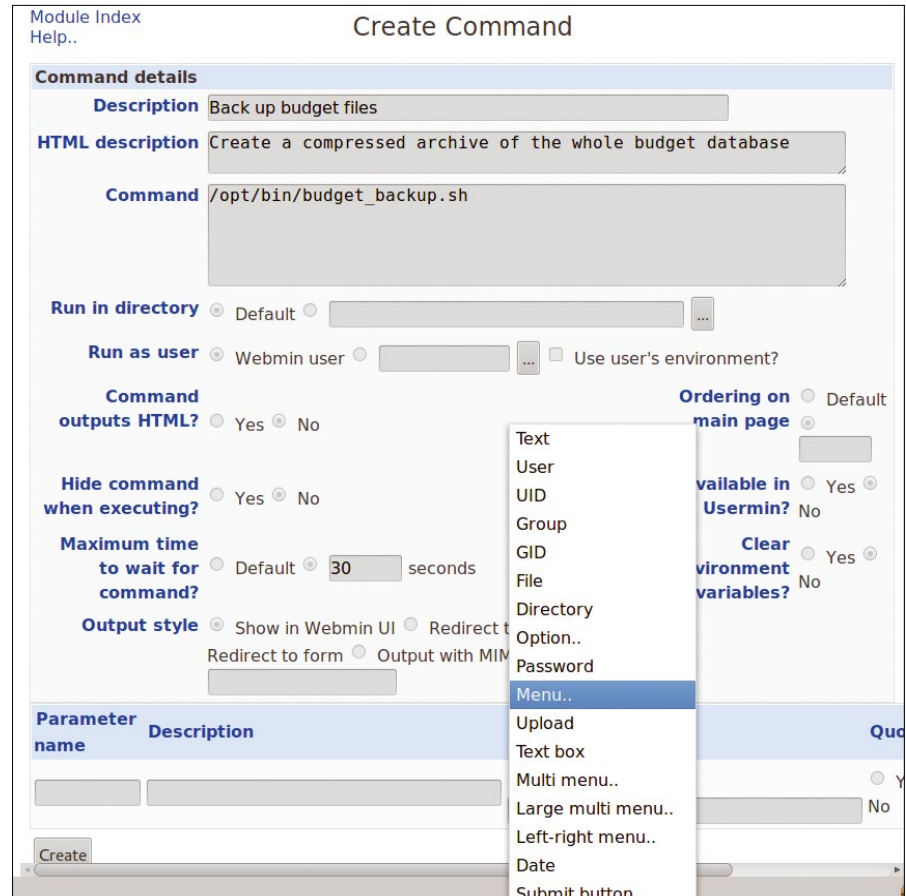


Figure 5: You can add buttons as you please, attaching them to your custom scripts, controlling which parameters they have, and defining how they display their output.

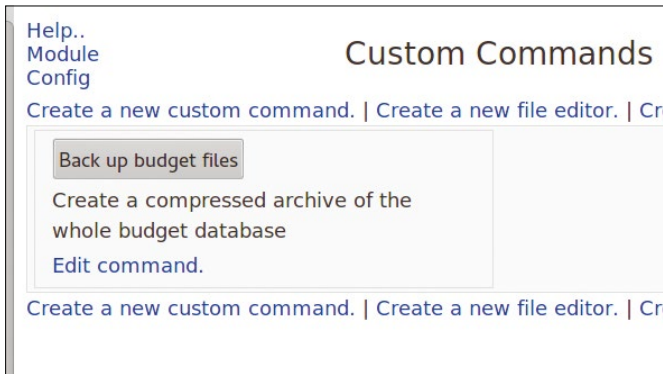


Figure 6: The *Back up budget files* button created with the form in Figure 5 will stay, by default, in the Custom Commands window.

Log Files, System Monitoring, and Automatic Jobs

Even with logfiles, Webmin clearly distinguishes between its own stuff and the operating system in which it runs. The *Webmin | Webmin Actions Log* window is a search interface for Webmin's internal log. Use it whenever you need to figure out exactly what you did in all of Webmin or just in some of its modules.

Rotation of system logfiles – and definition of post-rotation commands for each of them – has its own panel in the *System* section. Monitoring system performance is split between *Networking | Bandwidth Monitoring* and *Others | System and Server Status*. In the latter window, you can activate tens of system monitors, and you can set how often they should run and to which email address they should send updates and warnings.

Speaking of automated actions, the most useful part of Webmin might be its interface to set up and control cron jobs – that is, automatic execution of specific commands or scripts at regular intervals. From the visual point of view, the form in *System | Scheduled Cron Jobs* might be outdated, but it makes it very, very easy to create cron jobs that run *only* when you want. You can even specify a non-recurring range of dates in which each job should be executed and define which users can have their own cron jobs.

Firewall and Servers

The most critical part of the *Networking* section is probably the configuration of your firewall. By default, Webmin includes two modules for firewall configuration, called *FirewallD* and *Linux Firewall*. Which one is more appropriate de-

pends on which distribution you are running. Luckily, Webmin is smart enough to advise you, in most cases, which you should use.

Out of the box, Webmin knows how to configure the servers corresponding to the essential services of any Linux sys-

tem: remote access via SSH; websites with Apache; databases with MySQL; and email with Postfix, Procmail, and SpamAssassin. Many more servers, from DNS to DHCP, have separate modules that you can activate with a few clicks.

On one hand, this part of Webmin is the least useful: This is not Webmin's fault, it's just the nature of the beast. Regardless of which flavor of Linux you use, any configuration interface for services like running websites or email can only do so much for its users. Webmin has many practical forms that help a lot – and

save lots of time – to set each variable of those servers; however, it's still up to you to know what those variables mean.

Beyond the GUI: Advanced Webmin Usage

No software interface is really useful if you cannot customize at least part of it to your specific needs or occasionally do something that is not specifically listed in its menus. Webmin provides several functions that make these things possible even without hacking its source code.

To execute a command that you would otherwise have to enter at a shell prompt, click on *Others | Command Shell* and type the command. To do serious terminal work, choose *Others | Text Login* (Figure 4). This option opens a terminal inside your browser in which you can run all the commands and scripts you want inside a secure SSH session. If there's a command that you want to assign to a clickable button for frequent use, add it in the *Others | Custom Commands* page shown in Figures 5 and 6.

These features alone make Webmin flexible enough for most users, but there is more: If you want to save even more

LISTING 1: Using Webmin from a Perl Script

```
01 chop($url = `cat url.txt`);
02 eval {
03     $server = Frontier::Client->new('url' => $url);
04 };
05 $@ && die "Failed to create server : $@";
06
07 $jobs = $server->call("cron::list_cron_jobs");
08 print "Found ", scalar(@$jobs), " cron jobs\n";
09 ($already) = grep { $_->{'user'} eq 'root' &&
10                 $_->{'command'} eq 'echo foo' } @$jobs;
11
12 if ($already) {
13     print "Deleting cron job for $already->{'user'}\n";
14     $server->call("cron::delete_cron_job", $already);
15     print "Done deletion\n";
16 }
17 else {
18     print "Adding cron job for root\n";
19     $job = { 'user' => 'root',
20            'active' => 1,
21            'command' => 'echo foo',
22            'special' => 'weekly' };
23     $server->call("cron::create_cron_job", $job);
24     print "Done\n";
25 }
```

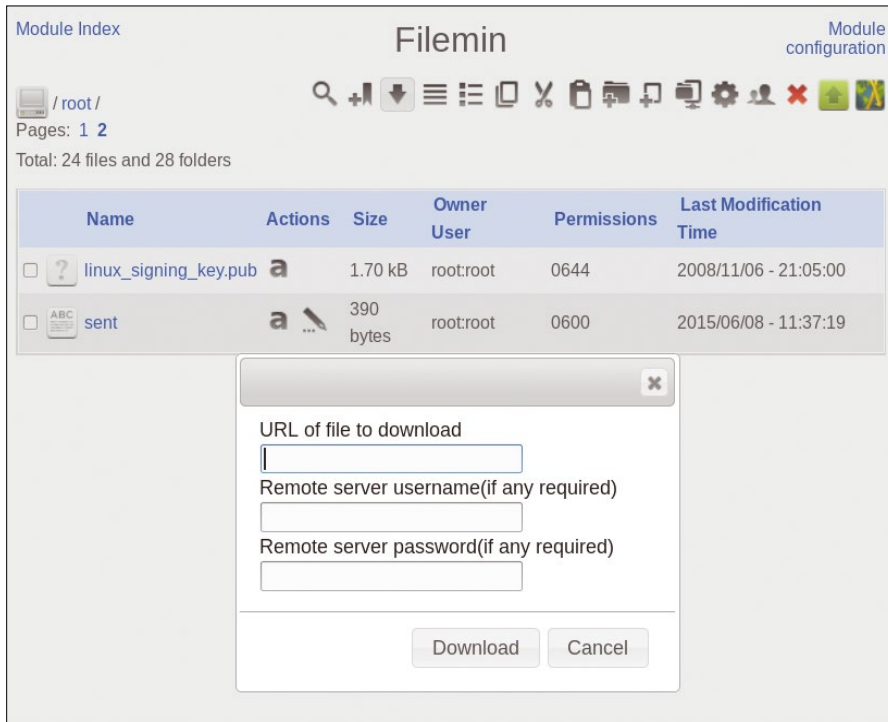


Figure 7: The Webmin file manager is no match for software like Dolphin or Nautilus, but it has all you need to move, rename, upload, and download files, even remotely.

time, you can write Perl scripts that will use Webmin on your behalf, using the XML Remote Procedure Call protocol. To understand how this trick works, look at Listing 1, which shows the essential parts of one of the sample scripts provided by the Webmin developers [5]. This script connects to your Webmin with the credentials in the single line of a file called `ur1.txt` (line 1), which has the following format:

```
https://USERNAME:PASSWORD@
yourserver.com:10000/xmlrpc.cgi
```

Lines 2 to 6 create an object that connects to the specified Webmin instance.

Using the `cron::list_cron_jobs` method, lines 7 to 10 check whether a cron job that executes the command `echo foo` as user `root` already exists. If yes, it is deleted (remember, this is just an example). Otherwise (lines 17 to 25), the server object is told to connect to Webmin again to create the cron job with the specified parameters. Of course, in real life, the user and command terms specified in the script would be parameters passed to the script.

Conclusion

Webmin might be visually outdated,



but even in 2016, it is a really useful tool for both expert and aspiring system administrators (Figure 7). Experienced users will save lots of time thanks to Webmin, and beginners will discover that Webmin is a useful tool for learning about Linux configuration. You'll find that Webmin has many more options than I can describe in this article. If you want further information, look at the Webmin extensions or spinoffs, called Usermin, Virtualmin, and Cloudmin.

Usermin creates a web-based interface in which ordinary users can read their email in a browser, change their password, and perform other simple tasks autonomously. Virtualmin gives Webmin the ability to create and manage multiple, semi-permanent virtual hosts on the same computer: Each host will also have its own copy of Webmin with a separate administrator account and no access at all to any other part of the system.

Cloudmin is similar to Virtualmin, but, as you can guess from its name, it is optimized for on-demand cloud computing. A Cloudmin admin can quickly create, resize, reconfigure, and restart on demand many independent clouds, each with its own parameters. ■■■

INFO

- [1] Webmin: <http://www.webmin.com>
- [2] Documentation: <http://doxfer.webmin.com>
- [3] Module documentation: http://doxfer.webmin.com/Webmin/Webmin_Modules
- [4] User documentation: http://doxfer.webmin.com/Webmin/Webmin_Users
- [5] Sample scripts: <http://www.webmin.com/xmlrpc.zip>

AUTHOR

Marco Fioretti is a freelance author, trainer, and researcher based in Rome, Italy, who has been working with Free Software since 1995 and on open digital standards since 2005. Marco also is a Board Member of the Free Knowledge Institute (<http://freenknowledge.eu>) and a team member of the Digital DIY project (<http://www.didiy.eu>).



Access your servers securely with a Magic URL

Magic Act

With a highly secure Linux server, you don't need a fixed IP address to connect over a cellphone network from anywhere on the planet if you have a Magic URL. *By Chris Binnie*

Much of my time I work with enterprise technologies, but when I get the chance to view Linux as a hobby again, I find myself solving problems by using only a simple script and a limited number of tools. My last foray into customizing Linux was securing a CCTV (closed circuit television) server that helps protect my home. In this article, I'll demonstrate how I solved a problem by using what I'm going to call a "Magic URL."

By Magic URL I mean that despite having what some would consider to be a highly secure Linux server (a firewall, two-step verification, fail2ban, TCP Wrappers, and htaccess), I don't need a fixed IP address to connect over a cellphone network from anywhere on the planet, just a spare 30 seconds, an SSH client, and access to a valid two-step verification code. I'll assume you have an Apache web server installed, preferably with HTTPS for a little more htaccess password secrecy. An htaccess configuration file is used to add simple password controls and is loaded and

executed by the Apache web server software [1].

It's worth saying that I'm sure you can solve similar problems in other ways, and some will suit your needs to a greater extent; however, this solution works for me, and you might find you can apply it or modify it to work for you in a future scenario. I've no doubt that after tinkering with this solution, I'll end up using something I've learned here on another project.

Escapology

How I arrived at having a Linux server hosted in the cloud and locked down tightly to host my CCTV footage has a significant background, but I don't want to lose track of how to set up a Magic URL, so I'll briefly just explain the tools I'm using before beginning.

None of these tools are mandatory, and once you see how they integrate with each other, you should be able to choose which tools you can use with relative ease. For those who prefer a quick solution, I'll also offer details of

the bare minimum required to achieve functionality.

For clarity, all the CCTV server does is receive uploaded footage (images and videos) of any motion-triggered events at my home. For obvious reasons, I won't go into too many details, but in essence, it's an image repository and a simple viewing platform to watch for unusual events at home day or night. Generally it will pull about 20GB of data off a single camera in a month. Incidentally, it costs pennies to run on a monthly basis, and camera hardware (even running glorious HD) is very affordable.

Fail2Prepare/Prepare2Fail

I'm a big fan of a super-slick tool called fail2ban [2]. It cleverly provides a mechanism to add firewall rules if it spots repeated failed logins in system logfiles. Netfilter's iptables is used in most cases on Linux, because it's so powerful and highly accessible.

In my case, on the CCTV server, I watch fastidiously and block anyone (automated or otherwise) from trigger-

```
Chain fail2ban-ssh (1 references)
pkts bytes target      prot opt in      out     source      destination
 34 1656 RETURN          all  --  *      *      0.0.0.0/0  0.0.0.0/0
```

Figure 1: Fail2ban stopping repeated incorrect password attempts over SSH.

```
Chain fail2ban-apache-chris (1 references)
pkts bytes target      prot opt in      out     source      destination
 10   520 REJECT      all  --  *      *      37.187.114.171 0.0.0.0/0
 1     60 REJECT      all  --  *      *      120.132.95.96  0.0.0.0/0
 0     0 REJECT      all  --  *      *      169.229.3.90   0.0.0.0/0
                                reject-with icmp-port-unreachable
                                reject-with icmp-port-unreachable
                                reject-with icmp-port-unreachable
```

Figure 2: I'm being really strict with 404 and other errors by banning every offender for a long time.

ing an HTTP 404. Here's the slightly unusual bit: I ban an IP address for a very long period of time for even triggering a single 404. This action might seem a bit draconian, but it's for good reason: I'm the only person accessing the CCTV server, and I know where to view the latest events with impunity. If you're completely new to fail2ban, it's well worth a look at an article I wrote in *Linux Magazine* [3] about how to keep pesky robots from filling your logs with noise when they generate HTTP 404s and, indeed, anything above a friendly HTTP 200.

On the CCTV server, I've possibly gone a little overboard with my fail2ban configuration, because sometimes troubleshooting a problem can take a little longer than it should; however, I have three iptables-driven fail2ban configs running. First, I use a config to deal with a popular remit for fail2ban: protecting against repeated SSH login attempts. In Figure 1 you can see how iptables reports it, the label is *fail2ban-ssh*.

Second, as you can see in Figure 2, I have a very strict set of rules for a non-existent page that I call *fail2ban-apache-chris*. The output in Figure 2 is abbreviated because of all the offenders (ranging from search engines to automated attacks). Admittedly, the rule doesn't have the best of names, but I'm the only sys admin on the server, so I know what it does.

You might want to compare the config examples in the fail2ban article mentioned earlier [3] with the fail2ban-apache-chris rules I run here (Figure 3). Notice the mention of *MAGIC-URL-PATH* within the *ignoreregex* field, because I talk about that later.

The final fail2ban ruleset was originally running within fail2ban-apache-chris to look for HTTP 401 login errors (i.e., users who don't know the password to .htaccess pages), but I found it too cumbersome to include alongside the Magic URL functionality, so I separated it out. In Figure 4 you can see the filter *fail2ban-apache-denied*, which

deals solely with HTTP 401 errors. Anyone who attempts to access my Magic URL subdirectories with the wrong login and password (also known as *HTTP 401 Unauthorized* errors) is

promptly struck down by lightning. It took a little while to get this ruleset just as I wanted, but the rules essen-

tially catch login failures over htaccess (Figure 5).

Thaumaturgy

As mentioned, fail2ban, like the other security tools used in this setup, isn't absolutely necessary to get the Magic URL functionality working, but it certainly adds to my peace of mind. It's also highly configurable, and I would suggest getting the other aspects of the Magic URL working first (so you're sure they're running for you as expected) and then adding fail2ban and two-step authentication with Google Authenticator afterward.

The CCTV images and video footage are uploaded by FTP with *vsftpd* into a user's home directory. The CCTV cameras are locked down by IP address thanks to TCP Wrappers, a host-based networking ACL system for *nix systems used to filter network access to Internet protocol servers. If you're new to TCP Wrappers and don't want to use iptables to lock down the likes of SSH or other network-accessible services, you might want to read an article I wrote for *ADMIN* magazine on that topic [4]. I use TCP Wrappers on this server to limit access to *sshd*, *vsftpd*, and the Magic URL script to a few static IP addresses. How it integrates with that functionality will become clear soon.

Within a subdirectory off the web root lives my main viewing page for the latest image (captured during the most recent motion-detected event). If I want to dig deeper, I can download an MPEG video

```
[Definition]
failregex = <HOST>.*[A-Z]* /(cms|admin|archiver|bbs|editor|common).*"
failregex = <HOST>.*\" (502|500|417|416|415|414|413|412|404|405|403|400)
ignoreregex = .*\"GET \/(MAGIC-URL-PATH|something).*
```

Figure 3: Fail2ban is catching the IP addresses of those requesting any non-existent pages on the web server and banning them with the fail2ban-apache-chris filter.

```
Chain fail2ban-apache-denied (1 references)
pkts bytes target      prot opt in      out     source      destination
1825 149K RETURN          all  --  *      *      0.0.0.0/0  0.0.0.0/0
```

Figure 4: To make scripting easier, fail2ban has a separate filter just for HTTP 401 errors.

```
[INCLUDES]
before = apache-common.conf

[Definition]
failregex = ^%(_apache_error_client)s (AH01617: )?user .*? authentication failure for "\S*": Password Mismatch(, referer: \S+)?$
            ^%(_apache_error_client)s (AH01618: )?user .*? not found(, referer: \S+)?\s*$
ignoreregex =
```

Figure 5: The fail2ban-apache-denied filter catches HTTP login errors.

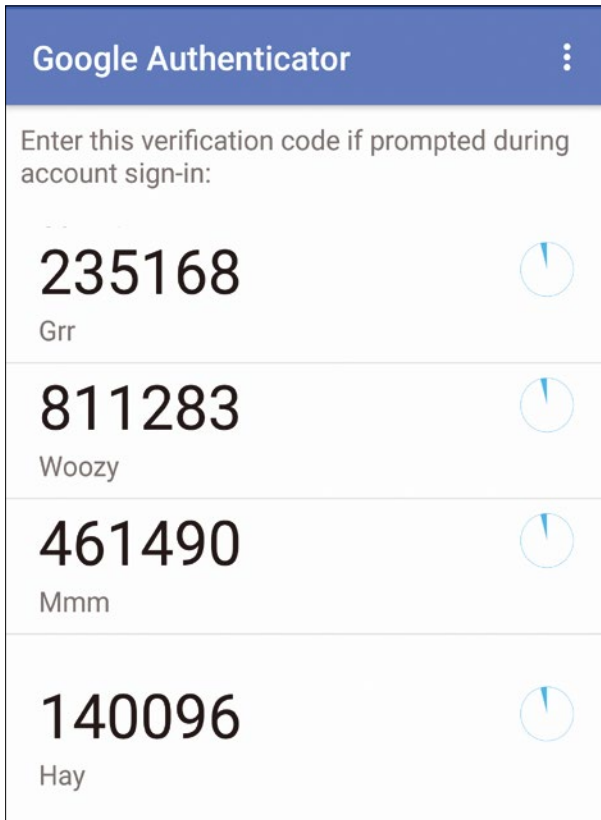


Figure 6: Ephemeral Google Authenticator codes on my smartphone.

or JPG images onto my smartphone to look at an event in more detail, having been presented with a prettified directory listing of all of the footage residing under capture dates and timestamp sub-directories.

Those who don't approve of security through obscurity might debate this next approach. I thought for a while about making this aspect much more sophisticated, and I might improve it in the future, time permitting. My issue with the core of this simple design is that, essentially, my Magic URL is statically named. Don't get me wrong: In addition to just being impossible to guess, the URL is secured by fail2ban, which would make short shrift of attackers' attempts to access it. (They'd be banned for more than a week if they got the password wrong once.)

To my mind, however, the best security comprises of a number of carefully considered layers. In the future, I might change my unguessable static URL every hour. I could generate a value, along the same vein as two-step authentication does, or use some other apparently random value known only to me, or I might opt for simplicity and just append the

current hour of the day to the end of the URL.

Say, for example, that under the web root of my Linux server I wanted to connect to my server at two o'clock in the afternoon by calling a subdirectory for my Magic URL (e.g., `chris-magic-url/`). I would have a cron job make sure my directory was renamed `chris-magic-url-14/` at 1400hr.

For the rest of this article I'll assume that `chris-magic-url/` is the name of my subdirectory. If you're diving straight into using fail2ban from the start, you should look again at the `ignoreregex` line in Figure 3, which I will use to replace `MAGIC-URL-PATH` with `chris-magic-url`.

I'll try not get too bogged down with the many security options available, but I do want to mention two-step authentication (also known as *two-step auth* or *two-step verification*) at this stage. It's a fantastic way of bolstering SSH security immediately after you've entered your password. You have to enter a time-limited code (as with a dongle or RSA token).

Figure 6 shows how to view Google Authenticator codes on a smartphone once it's set up. In case you're wondering, I don't ever name my authentication services by their real name in case my phone is lost or stolen. There's nothing wrong with some security through obscurity in my book; it all helps to slow an attacker down or make things more difficult to figure out.

It's a Kind of Magic

I've explained that two-step and fail2ban are optional components, but I should really have one mechanism locking `sshd` to specific IP addresses at the very least in case I make a mistake with my Magic URL config. I could easily use iptables on it's own, for example, but I need to be careful not to confuse that with my

fail2ban rules. The Magic URL functionality will allow otherwise untrusted dynamic IP addresses to connect briefly over SSH to the server, so for clarity, to do so, I'll follow these steps below:

1. Set up my `.htaccess` files, which means that even if my Magic URL is guessed, a username and password is needed to get access to it via a Browser (and fail2ban won't take too kindly to any failed passwords at all if it's configured correctly!).
2. Write a script (and break it up into segments to keep it simple) to check to see if I need to act after my Magic URL has been accessed and tidy up previous SSH access entries, as well.
3. Create a simple cron job to set the Magic URL, so it runs every minute.

Sleight of Hand

Even novice sys admins have probably heard of `.htaccess` files, which provide some web servers (Apache being the most popular) with some clever functionality, like cleaning up URLs or password protection of web pages. I'll create two simple files, each with just a single line, to add password protection to my `chris-magic-url/` directory.

To keep my password secure I'll save the password file a level above web root (so it's not viewable by web visitors); in this case, that's `/var/www`. In that directory the `.htpasswd` file looks like this:

```
hack-and-defend:$apr1$jbx7Soz$Z
c1t5T9p6KzHMXj8wqmWHR
```

As you can see, my username is `hack-and-defend` and my password is an MD5 hash. If you're new to how this works, I must tell you about it, because it's really

LISTING 1: .htaccess Config

```
AuthName "These aren't the droids
you're looking for. Go about your
business."
AuthUserFile /var/www/.htpasswd
AuthType Basic
<Limit GET POST>
Require valid-user
Order Deny,Allow
Deny from all
Allow from 10.10.10.10
Allow from 12.12.12.0/24
Satisfy Any
</Limit>
```


LISTING 2: A Simple .htaccess File

```
AuthName "You're definitely not welcome here. Move along."
AuthUserFile "/var/www/.htpasswd"
AuthType Basic
Require user hack-and-defend
```

very clever. The password is just obfuscated characters using MD5 and written as a hash. The actual password is never saved to disk. The login process asks for a password, creates a hash of it, and simply compares it to the one saved on the disk. If it matches, I'm allowed in. You have to love intelligent security that uses such well-considered simplicity.

The next step is easy, too. At the top of the web root, I create two directories. (With some fine tuning, the two directories could become one, but I'll keep using two directories for simplicity.) The first directory is called `chris-magic-url/` and the second is `secret-service/`. In my case, `secret-service/` is the CCTV footage viewing page I mentioned. You might have an application or something precious sitting in your `secret-service/` directory.

For clarity, my Magic URL will open up SSH access to any IP address, and the `secret-service/` directory isn't strictly related. I thought it was worth mentioning, though, because in addition to the Magic URL functionality, you can also do what I think are pretty clever things with IP addresses and password protection with `htaccess` files. Here's a quick look to give you some food for thought.

You might want to improve the Magic URL functionality to let only specific IP addresses even attempt that first login via

.HTACCESS FILES

In case you aren't aware, the prepended dot for these `htaccess` files not only makes them hidden files on the filesystem, but modern web servers also know not to serve these files directly to a browser (e.g., so your MD5 hash is never exposed). For example, Apache has this configuration in the `apache2.conf` file in my Debian Jessie build:

```
# The following lines prevent .
# htaccess and .htpasswd files from
# being
# viewed by web clients.
#
<FilesMatch "^\.ht">
    Require all denied
</FilesMatch>
```

`htaccess`. Think about something like this example in Listing 1. You can see that not only do I ask for a password for ac-

cess to my web server files, but I also have configuration in place that doesn't ask for a password if either the single IP address 10.10.10.10 connects or the whole /24 network 12.12.12.0 connects. That's very convenient.

The clever line is `Satisfy Any`. Be warned that you should test this carefully before trusting it implicitly. I really like this functionality, and the nice thing about `htaccess` is that any changes you make to it via a script (e.g., adding a trusted IP address via Magic URL functionality) are live immediately, and no web server restarts are required.

Back to the simpler `htaccess` approach and the second `htaccess` file: I'm going to insist that everyone is asked for a password, and in this case, `fail2ban` will keep me safe from repetitive attacks. You'll be glad to know that the second file is straightforward too and looks like Listing 2. This time the file resides in the directory `/var/www/html/chris-magic-url/` and I call the file `.htaccess` (see the ".htaccess Files" box).

I'll also put a copy of my second file (`.htaccess`) in my `secret-service/` directory to protect my secret application (e.g., the viewing page in my case). Again, you don't need to do this to get the Magic URL functionality working, and it's optional, it's simply to give you some context for where your prized assets might live in relation to your Magic URL directory.

The only other thing I'll need inside my virtual host files for `htaccess` to function correctly (as it applies to my Debian Jessie build and Apache version 2.4.10-

10 + `deb8u4`) is an entry in the `/etc/apache2/sites-available/` virtual host entries. Within the `000-default.conf` file, and the `default-ssl.conf` file if you're using SSL/TLS, I add the following before closing the `</VirtualHost>` tag:

```
<Directory "/var/www/html">
    AllowOverride AuthConfig
</Directory>
```

Clearly, I need to alter the directory on which I am enabling password protection by editing the line:

```
Directory "/var/www/html"
```

Now I just need to tell Apache that I've made those changes, if they weren't there before, with a quick config reload as follows for Debian derivatives on systemd machines (first line) or the SysV equivalent for older versions (second line):

```
# systemctl reload apache2
# /etc/init.d/apache2 reload
```

On Red Hat derivatives, your package might be called `httpd` and not `apache2`, so use

```
# systemctl reload httpd
# /etc/init.d/httpd reload
```

for `systemd` and `SysV` systems, respectively.

After using the Bash script [5] for the first time, which I'll look at shortly, and triggering a Magic URL, a quick peek inside the `/var/www/html/chris-magic-url/` directory reveals the previously unseen files, and the Magic URL directory contains:

```
.htaccess
index2.html
index.html
```

LISTING 3: Receive Email on Attack

```
ALL: ALL: SPAWN ( \
    echo "\n\
    TCP Wrappers\: Connection refused\n\
    By\: $(uname -n)\n\
    Process\: %d (pid %p)\n\
    User\: %u\n\
    Host\: %c\n\
    Date\: $(date)\n\
    " | /usr/bin/mail -s "Connection to %d blocked" chris@binnie.tld) &
```

LISTING 4: Magic URL Bash Script

```

01 #!/bin/bash
02
03 # Create Magic URL dir if not present
04 mkdir -p /var/www/html/chris-magic-url/
05
06 # Check if I ran last time and tidy up if so
07 CHECK=`tail -n2 /etc/hosts.allow | head -n1`
08
09 if [ "${CHECK}" = "# START MAGIC URL IP ADDRESS #" ] ; then
10 sed -i '/# START MAGIC URL IP ADDRESS #/d' /etc/hosts.allow # delete marker in file
11 sed -i '$d' /etc/hosts.allow # delete last line
12 fi
13
14 # Create date variable
15 DATE=$(date +%H:%M:%S %m-%d-%Y)
16
17 #####
18 # Create the first index file and avoid pesky browser caching (this is not infallible)
19 echo -e "<html><head><meta http-equiv=\"cache-control\" content=\"no-cache\"><meta http-equiv=\"expires\"
  content=\"0\"><meta http-equiv=\"expires\" content=\"Tue, 01 Jan 1980 1:00:00 GMT\"><meta http-equiv=\"pragma\"
  content=\"no-cache\"><meta http-equiv=refresh content=\"5;URL='index.html'\">" > /var/www/html/chris-magic-url/index.html
20
21 echo ${DATE} >> /var/www/html/chris-magic-url/index.html
22 #####
23
24 #####
25 # Empty second index file but make sure you know it exists and the last time it was written to
26 echo ${DATE} > /var/www/html/chris-magic-url/index2.html
27 #####
28
29 if [[ ! -z ${GSM} ]]
30 then
31
32 echo "# START MAGIC URL IP ADDRESS #" >> /etc/hosts.allow
33 echo "sshd: ${GSM}" >> /etc/hosts.allow
34
35 echo -e "<html><head><meta http-equiv=\"cache-control\" content=\"no-cache\"><meta http-equiv=\"expires\"
  content=\"0\"><meta http-equiv=\"expires\" content=\"Tue, 01 Jan 1980 1:00:00 GMT\"><meta http-equiv=\"pragma\"
  content=\"no-cache\"><meta http-equiv=refresh content=\"5;URL='index2.html'\">" > /var/www/html/chris-magic-url/index.html
36
37 echo -e "<html><head><meta http-equiv=\"cache-control\" content=\"max-age=0\"><meta http-equiv=\"cache-control\"
  content=\"no-cache\"><meta http-equiv=\"expires\" content=\"0\"><meta http-equiv=\"expires\" content=\"Tue, 01 Jan 1980
  1:00:00 GMT\"><meta http-equiv=\"pragma\" content=\"no-cache\">" > /var/www/html/chris-magic-url/index2.html
38
39 echo ${DATE} >> /var/www/html/chris-magic-url//index2.html
40 echo -e "<br>${GSM}" >> /var/www/html/chris-magic-url//index2.html
41 echo -e "<br>Aphorism: Time is short, carpe diem." >> /var/www/html/chris-magic-url//index2.html
42
43 iptables -D Fail2ban-ssh -s ${GSM} -j REJECT
44 iptables -D Fail2ban-apache-chris -s ${GSM} -j REJECT
45 iptables -D Fail2ban-apache-denied -s ${GSM} -j REJECT
46
47 mail -s "Magic triggered ${GSM}" chris@binnie.tld < /dev/null
48 sed -i '/chris-magic-url/d' /var/log/apache2/access.log
49 /etc/init.d/apache2 reload
50 else
51 echo No match
52 fi

```

Remember, to get the bare minimum working, you only need `htaccess` files in place, TCP Wrappers, the Bash script, and a cron job.

Hat Trick

It's important to note that single lines in the script could easily use `iptables` as opposed to TCP Wrappers. I prefer the simplicity of dealing with the `/etc/hosts.allow` file from TCP Wrappers for a task like this. Hopefully, it also makes it easier to follow and avoids confusion because I'm looking at `fail2ban` `iptables` entries too.

To set up TCP Wrappers for SSH, I first put a global deny in the `/etc/hosts.deny` file (Listing 3). This example includes a nice little trick to receive email whenever someone or something is rejected by a TCP Wrapper rule. One caveat: If this makes your server prone to denial of service attacks because someone can generate a lot of email, then disable it.

The `ALL:ALL` is the only entry you need to deny *all* IP addresses for *all* services in that file if you don't want to receive email. Try it and see if you like the func-

tionality. I love this feature of TCP Wrappers, because it helps me keep an eye on misconfigurations and attacks via the occasional email.

In `/etc/hosts.allow` you should add (comma delimited) IP addresses to which you want to allow SSH access:

```
sshd: 22.22.22.22, 11.11.11.11
```

I won't edit `/etc/hosts.allow` directly any further because the Magic URL script will take care of writing to that from now on. Be warned not to disconnect your existing SSH session, however, or you might lose remote access. Test the Magic URL set up using a virtual machine with a remote console (or one which has out-of-band access) if you don't want to risk being accidentally locked out of your server.

Rabbits in Hats

I mentioned the other security measures running on my CCTV server earlier. In a moment, I will share snippets of their configuration, which I use to integrate with the Magic URL function-

ality – I hope without confusing the matter in hand.

First, though, I'll look at the heart of the Magic URL mechanism: a short Bash script. This main script, which will be run by a cron job, is called `/usr/local/etc/magic-url-script`. I won't go into every detail, but I'm sure you'll get the idea. You should download a version of the script [5] as opposed to entering it by hand to avoid typos.

The first section (Listing 4, lines 1-27) begins by tidying up any previously added Magic URL IP addresses in the `/etc/hosts.allow` file, creating the pertinent directory, clearing index files of old data, and creating a timestamp variable called `#{DATE}`.

As you can see, although it's noisy, it's just a case of overwriting the `index.html` and `index2.html` files in the `/var/www/html/chris-magic-url` directory. I then simply write a timestamp to them so I can see easily in a browser when they were last updated.

In the second section (lines 29-38), I'm calling the IP address I allow to connect to the GSM server (Global Sys-

A Webzine for High-Performance Computing Specialists

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

ADMIN
Network & Security



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

tem for Mobile Communications, which is the older name for all things related to cellphone data and now called 4G, 3G, etc.). To check inside the web server access logs to see if anyone has triggered the Magic URL and had their IP addresses grabbed, define the GSM variable with:

```
GSM=`cat /var/log/apache2/access.log |
grep --line-buffered chris-magic-url |
awk '{print $1}'|tail -n1`
```

If you've messed around with your Apache logging, you might want to change the `print $1` value to `print $3`, for example. Just copy that line and test it on the command line to get it right.

From here onward, the important stuff happens. To start off, I'm appending the `/etc/hosts.allow` file with a line mentioning the ``${GSM}`` IP address to allow SSH access, with a comment above it that acts as a marker. In case you're wondering, TCP Wrappers doesn't need to be restarted, because the config is picked up immediately and is "live."

The two next busy-looking blocks of code simply create non-caching HTML files to ease the pain of your browser. If you're using a smartphone, that could be a welcome addition because they try to cache as much as possible to reduce data tariff costs.

That opening line for the `if` statement (line 29) simply runs through the commands that follow if the ``${GSM}`` variable isn't empty after the Magic URL has been spotted in the web server logs.

The rest of the script (lines 40-53) dutifully creates the index files and then redirects the browser to `index2.html` when the Magic URL has been successfully triggered.

Now look at lines 44-46 starting with `iptables`, which are there for `fail2ban`, if used. They remove entries for the newly allowed dynamic IP address – in my case, one for each filter (i.e., repeated SSH login attempts, non-existent pages, and HTTP 401 login errors). Next, the script then sends email telling me the Magic URL was triggered (line 48).

Finally, I delete all mentions of the Magic URL from the web server's access logs using `sed` and do a relatively painless Apache reload (lines 49-50) so the logs keep updating correctly.

Note that because of the way some web server logs are created, one "hit" on the website might create two log entries; moreover, as `meta` refreshes take place (redirecting the browser to `index2.html`), a few instances of `chris-magic-url` could be generated in the web logs.

As a result, the SSH login window can sometimes be extended to two or three minutes; a third run of the script should clean up `/etc/hosts.allow` access. You could create another timestamp file to catch only one instance of the Magic URL being triggered rather than rely on the web server logs directly. For my purposes, this oddity doesn't cause concern, but you could certainly finely tune this aspect.

Although you might want a slicker way of solving this problem on a very busy high-traffic production server, you have many other ways to do so, such as by using other temporary files and not reloading Apache config at all. I'll leave you to experiment to your heart's content.

In case you're wondering, the `No match` section (line 52) can be used if you manually run the script for testing. If no Magic URL entry exists in your web server logs, you will see `No match` echoed to your screen.

Magicians

The last stage is to run the Bash script every minute. In my `/etc/crontab` file I have the following reference to run the `/usr/local/etc/magic-url-script` file:

```
* * * * *      root
/usr/local/etc/magic-url-script >
/dev/null 2>&1
```

For cron to pick up the new run-every-minute cronjob in the crontab, you might need to run the following command on a current Linux version,

```
# systemctl restart cron
```

or this command on the SysV equivalent:

```
# /etc/init.d/cron restart
```

Note that I intentionally left the root user running this cron job so that everything works without permission issues, but ideally, you should run it as a non-privileged user. Change the parts that

would otherwise need to be run by root and prepend `sudo` to your privileged commands (e.g., when you append to the `/etc/hosts.allow` file).

The Great Reveal

The proof is in the pudding, as they say. To gain access to the server, I first visit the Magic URL in a browser on the device on which I want to log in over SSH. I include the Magic URL's `htaccess` login and password and build them into the bookmark of the URL for ease:

```
https://hack-and-defend:chrisbinnie@cctv.server.tld/chris-magic-url/
```

The `htaccess` user `hack-and-defend` has the password `chrisbinnie` in this case. The server hostname is `cctv.server.tld`. Because the Magic URL lives off the web root, I can simply append it to the hostname (the `chris-magic-url/` section).

By design, the next thing seen in the browser wouldn't thrill an attacker and immediately pique their interest. If your password is securely passed over the (optional) HTTPS connection and accepted by the web server, you'd be presented with an unglamorous page that simply offers a (mobile-friendly) blank, white page with nothing more than a timestamp displayed. There's nothing of interest to see for anyone that gets this far. In my case it looks like this:

```
11:11:01 11-11-2016
```

You can see from the timestamp that the page was dutifully updated by the cron job within the last minute, and because I can see the page, it means I've been allowed access via `htaccess`.

If you cast your mind back to the HTML echoed by the Bash script, you won't be too surprised to see the page auto-refresh within a few seconds and jostle you onward to the second URL, which appears in the browser's address bar:

```
https://cctv.server.tld/
chris-magic-url/index2.html
```

On that page, I see an equally perplexing message:

```
11:12:01 11-11-2016
154.138.22.6
Aphorism: Time is short, carpe diem.
```

The first line confirms when that page was last updated (a minute after the first index page). The second line is the IP address picked up in the web server logs, and the third line is a shameless geek joke.

Why did I put the joke there? It's to remind me that I only have a minute (or less depending on when I see that page) to spawn an SSH session with my CCTV server before being locked out again. In this way, I prevent lingering, insecure connections. Even on my slow GPRS my smartphone connects within a few seconds without breaking a sweat (I use the Blowfish cipher on SSH, which is a little lighter over slow mobile connections), so this time interval presents no problems.

In my case, I fire up my favorite SSH client on my smartphone, JuiceSSH, and, et voila, my connection is accepted – even using an otherwise unwelcome dynamic IP address – and I'm presented with a login prompt. As I've said, I prefer the belt-and-braces approach to accessing servers. So, at this point, I am also presented with a two-step verification prompt for a code generated by the Google Authenticator app on my smartphone, which is yet another security layer with which to contend.

Although I can only spawn a session within a minute, I'm still allowed SSH access for as long as the session is open, which suits me just fine. I might prune some footage from the storage drive on the server or check out something else of interest.

A reminder that while you're testing this (assuming you're not going for the Fort Knox approach straightaway with iptables, fail2ban, and two-step auth initially) the only file that is really being changing at this stage is for access by sshd via the `/etc/hosts.allow` file.

To test whether you can log in correctly, just run:

```
# cat /etc/hosts.allow
```

If all is well, you should see that your cellphone IP address has now been added to the end to allow SSH access; for example:

```
# START MAGIC URL IP ADDRESS #
sshd: 123.123.123.123
```

If you're not sure what your dynamically allocated IP address is, then just ask your friendly, neighborhood Google by entering

```
https://www.google.com/#q=what%20is%20my%20ip
```

in your browser's address bar.

Dance the Two-Step

Earlier I promised to show you how I integrated other security layers with my Magic URL. You briefly saw how I tidied up any fail2ban iptables rules for my dynamic IP address within the main script. Here, I will take a quick look at some two-step authentication configuration.

I will leave you to research the finer details of setting up your two-step verification with the excellent Google Authenticator. However, should you want to add functionality to your Magic URL, wherein IP addresses you know and trust don't have to run through the two-step code-entering process but instead are just asked for a password, then you can use code like this:

```
echo -e "+ : ALL : 51.15.2.14/24
+ : ALL : ${GSM}
- : ALL : ALL" > /etc/security/access-local.conf
```

As you can see, `${GSM}` is the IP address extracted from the Magic URL functionality. If asked, I would say use this with caution and probably only allow fixed-line IP addresses to connect over SSH without running through the two-step process.

The End

Although this approach to opening up access is rudimentary in places, I would like to think it encourages sys admins to develop their lateral thinking. The fact that you can get htaccess to deal with unique IP addresses, rename the Magic URL every minute with a new timestamp derivative (or via another security token), and incorporate other failsafes into a security design raises many questions about how far it's possible to go in securing a server using relatively basic tools.

If you know that you won't be able to carry your SSH key with you for a period

of time, you could take two-step authentication scratch codes with you on a piece of paper (or just switch off two-step completely) and still get SSH access relatively securely from anywhere on the planet using this approach.

Something else I might try in the future is spawning another type of task or application with a Magic URL. I didn't want a big clunky *Submit* button embedded within a web page to allow Magic URL functionality, but you could easily customize this approach to run any service or script you like.

Having read about the Magic URL I use for my CCTV server, I hope you will attempt some of these methods in a creative way to suit your own security needs. ■■■

AUTHOR

Chris Binnie's new book, *Linux Server Security: Hack and Defend* [6], shows how hackers launch sophisticated attacks to compromise servers, steal data, and crack complex passwords, so you can learn how to defend against these attacks. In the book, he also talks you through making your servers invisible, performing penetration testing, and mitigating unwelcome attacks. You can find out more about Linux servers and security at his website (<http://www.linux-server-security.com>).

INFO

- [1] htaccess: <http://www.htaccess-guide.com>
- [2] fail2ban: http://www.fail2ban.org/wiki/index.php/Main_Page
- [3] "Intrusion Detection with fail2ban" by Chris Binnie: <http://www.linux-magazine.com/Online/Features/Intrusion-Detection-with-Fail2ban>
- [4] "Secure Your Server with TCP Wrappers" by Chris Binnie: <http://www.admin-magazine.com/Articles/Secure-Your-Server-with-TCP-Wrappers>
- [5] Code for this article: <ftp://ftp.linux-magazine.com/pub/listings/magazine/190>
- [6] Binnie, Chris., *Linux Server Security: Hack and Defend*. Wiley, 2016, <http://www.amazon.com/Linux-Server-Security-Hack-Defend/dp/1119277655/>

HipChat and Slack alternatives

Chitchat



HipChat and Slack are in when it comes to communicating with customers or colleagues, but data security is out. InspIRCd, Rocket.Chat, or Mattermost are better alternatives.

By Martin Loschwitz

When you want to communicate with colleagues over a distance – whether between offices or between cities – email is always an option, but not the best if you just need a quick answer to a simple question or need some face time. Most companies have established real-time communication through the use of some kind of chat service, and employees often use these same systems for chatting with customers.

Two services have become particularly popular in recent years: HipChat [1] and Slack [2]. Although they are easy to use, your data resides on the hosts' servers in the US, which does not have a good record of digital privacy and security. In this article, I present three alternatives to Slack and HipChat: InspIRCd [3], Mattermost [4], and Rocket.Chat [5]. A test team set out to discover whether you need your own client software with these three alternatives and to determine how well each candidate accommodates business use.

Basics

HipChat and Slack are successful for many reasons. First, they are easy to use within a web browser, which means it

isn't necessary to install a client (although HipChat does have a desktop client if you want it). Second, the software provider normally hosts both services, so you don't have to run your own infrastructure, and there's no need for troubleshooting. Finally, both systems also provide useful additional features, such as formatting for code, automatic link previews, or the integration of photos in ongoing chat sessions.

However, hosting is what worries most companies. Anyone who communicates with colleagues over a chat service can't be sure that internal or confidential information doesn't also use the same channels. With HipChat, the data ends up directly on Amazon servers. It's pretty similar with Slack. Companies who want to protect sensitive data could feel uncomfortable about placing their data on servers in a country where data privacy has a rather tarnished image, and you have to address the inevitable question of how to prevent company data from being leaked.

InspIRCd

Much has changed in the technology of real-time chat over the past 15 years.

Some chat software builds on the IRC protocol, which does without encryption, whereas others are web based and provide TLS out of the box. InspIRCd's name reveals that it is an IRC server, although that's no proof of added value: Almost all IRC servers are related, and rarely are they easy for administrators to operate.

The IRC protocol is one of the oldest communication protocols of the Internet. Although it is documented in detail in an RFC [6], no established IRC servers are actually based on this RFC. Instead, over the years, many admins thought it was a good idea to expand the protocol with more and newer functions; thus, a zoo of inflated IRC servers has developed.

InspIRCd claims to be different. It isn't a fork of the UnrealIRCd [7]; rather, it is a completely new development programmed in C++. The developers promise that InspIRCd therefore doesn't suffer from the bugs and limitations that make life difficult for administrators with other IRC daemons.

As early as the installation process you might become dubious of InspIRCd, because you won't find it in any server distributions. Neither SUSE, Red Hat,

Debian, nor Ubuntu have packages. The developers only provide an installer for Windows. Anyone wanting to use the service on a Linux distribution will have to fire up the compiler. Although it's feasible, it's not exactly convenient.

Bundles of Features

A look at InspIRCd's list of features gives more reason for doubting its ease of use. On the program's wiki, the developers compare their work with other popular IRC servers and determine quite succinctly that InspIRCd can do pretty much everything that has ever been invented in terms of functionality with IRC servers. Basic configuration is pretty easy in the beginning, but you need to do plenty of reading of the relevant documentation to use the advanced features (Figure 1).

InspIRCd incorporates all its functionality in very little code, so compared with UnrealIRCd, it manages with fewer than half of the lines of code. However, it still supports SSL encryption if the client provides the necessary prerequisites. InspIRCd can also be combined with additional services. For example, if you are in need of a bot that manages certain channels, you can link it with InspIRCd.

A Question of the Client

Unlike HipChat or Slack, InspIRCd is just a server. The developers don't offer a matching client, and why should they? All platforms have the choice of a number of IRC clients: M-IRC for Windows and the cross-platform command-line program `irssi` are just two examples.

This plethora of clients is both a curse and a blessing. With a wealth of alternatives, not only is the choice of client imposed on the user, but various functions will be unavailable. IRC is a plain text product with no automatic previewing of links, so although clients can integrate features on their own, achieving a consistent behavior can be problematic. For example, client A might show text prefixed with the tag `<code>` in fixed-pitch font without wrapping, whereas client B might display such a tag as normal text.

Another problem you'll encounter is with file sharing: IRC doesn't offer a way to store files for download. The exchange of files in IRC occurs via DCC (Direct Client-to-Client). The name itself suggests that a direct connection must be established between the two clients

```
server name="penguin.omega.org.za"
description="Waddle World"
network="Omega">

<admin name="Johnny English"
nick="MIS"
email="MIS@the.best.secret.agent">

<bind address="" port="6000" type="clients">
<bind address="" port="6660-6669" type="clients" ssl="gnutls">
<bind address="" port="7000,7001" type="servers">
<bind address="1.2.3.4" port="7005" type="servers" ssl="gnutls">

<power diepass="diepass" restartpass="restartpass" pause="2">

<include file="file.conf">

<connect allow="196.12.*" password="secret">
<connect allow="*" timeout="60" flood="20" threshold="1" pingfreq="120" sendq="262144"
rcvq="8192" localmax="3" globalmax="3">
<connect allow="*" port="6660">
<connect deny="69.254.*">
<connect deny="3ffe::0/32">

<class name="Shutdown" commands="DIE RESTART REHASH LOADMODULE UNLOADMODULE RELOAD">
<class name="ServerLink" commands="CONNECT SQUIT RCONNECT MKPASSWD MKSHA256">
<class name="BanControl" commands="KILL GLINE KLINE QLINE ELINE">
<class name="OperChat" commands="WALLOPS GLOBOPS SETIDLE SPYLIST SPYNAMES">
<class name="HostCloak" commands="SETHOST SETIDENT SETNAME CHGHOST CHGIDENT">
```

Figure 1: Configuring InspIRCd is a test of patience for administrators and requires reading a lot of documentation.

for the exchange to work. However, virtually all popular firewalls prevent such connections – and reconfiguring the firewall isn't the only challenge, because port forwarding also needs to work properly if one of the two clients goes online via NAT.

Data exchange via DCC has almost completely disappeared from the everyday life of larger IRC networks for this very reason. Very few administrators want to poke holes in their firewall.

Web-based IRC clients such as Kiwi IRC [8] (Figure 2), which operate on the same host as InspIRCd, provide a possible solution. You can even enforce their use at the InspIRCd level, meaning a completely homogeneous client landscape can

exist. However, such a setup again requires installation and configuration work. The result is an easier way for customers or suppliers to get in direct contact with individual employees. All it takes is a suitably prepared link to the web IRC client.

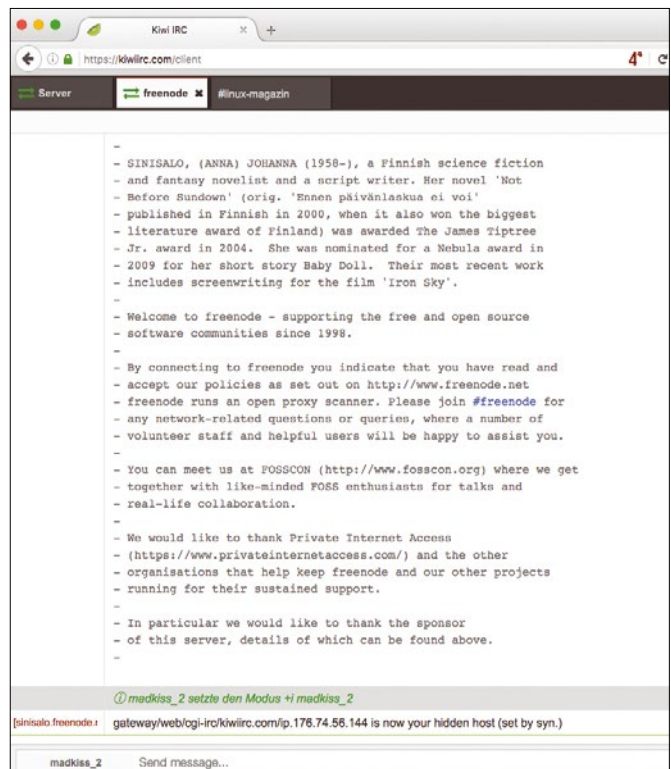


Figure 2: Web-based IRC clients such as Kiwi IRC help achieve a certain consistency on the client side.

Mattermost

The data protection problems experienced by HipChat and Slack were noticed by many community members several years ago, giving rise to IRC protocol alternatives. Mattermost, written in Golang and React, is a typical example: Its authors promote it as a “modern Slack alternative.” In fact, Mattermost, Slack, and HipChat have unmistakable similarities: All three services offer a native web interface that can be used without an additional client, and each has native clients that can be run on desktops.

The appearance is always the same: Typical convenient functions such as previewing web links are standard, as are stylesheets that show source text with correct formatting. Unlike HipChat and Slack, Mattermost can be run locally, which means that use of the ser-

vice doesn't automatically mean data is outsourced to external servers.

High Installation Overhead

Mattermost doesn't make it easy. The instructions for common Linux distributions on the program website extend over many screens, starting with the installation of a database – the Mattermost developers recommend PostgreSQL.

Mattermost also needs its own disk space for large files. For example, if you load a large file during an ongoing chat, the file automatically ends up in a separate data directory. It can then be downloaded by other chat users.

Once the database is up and running, you still have plenty of setup work left. Installing Mattermost itself is the next item on the list. As with InspIRCd, you can search in vain for pre-built packages

for popular server distributions. Operating the service will inevitably lead to downloading a tarball, which then needs to be unpacked on the target host. The Mattermost configuration file, written in JSON, is where you can store the data for logging into the PostgreSQL database.

The third step involves installing Nginx: The developers recommend running Mattermost behind a firewall. In the setup, Nginx serves as a proxy server that accepts incoming connections and forwards them to the Mattermost server. Because Nginx also works as an SSL terminator, Mattermost can be operated as an SSL service with upstream Nginx. When the PostgreSQL, Mattermost, and Nginx trio are ready, the next step is the final configuration: Administrators can do this work in the Mattermost web-based System Console.

A Cornucopia of Functions

Once Mattermost is up and running, the fun can begin. Thanks to a rich set of features and the aforementioned formatting features, Mattermost provides a system-wide full-text search. Themes let you customize the look to your personal preferences or company requirements (Figure 3).

Mattermost is clearly aimed at Slack converts. The Slack apps contributed by third parties, for example, extend Slack, adding many functions that are missing in the plain vanilla version. Mattermost has an interface for webhooks that makes it easy to integrate Slack apps into Mattermost chats. Moreover, logs and accounts can be imported from Slack directly into Mattermost, making migration significantly easier. Anyone using a Slack theme can import it into Mattermost and will even find more configuration possibilities there than in the original.

Many clever community functions, such as GitHub or GitLab external systems, can be integrated directly into the service using webhooks (Figure 4). The developers see a webhook as a URL that calls the respective version control system if a certain action occurs. Slack can be replaced easily with Mattermost, because Mattermost uses the same syntax for webhook calls as Slack.

A separate plugin for GitLab connects the service even more closely with Mattermost: Pull requests, comments, and

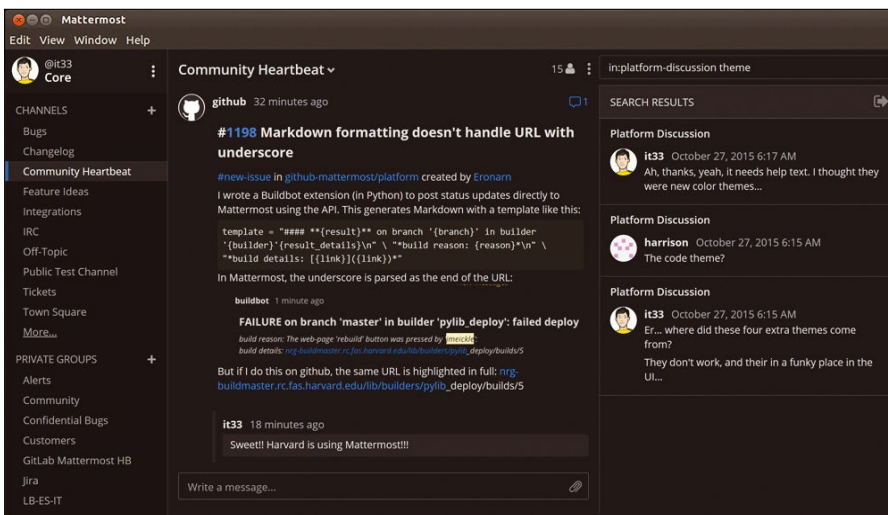


Figure 3: Themes are part of Mattermost's functionality: A client in company colors is therefore conceivable.

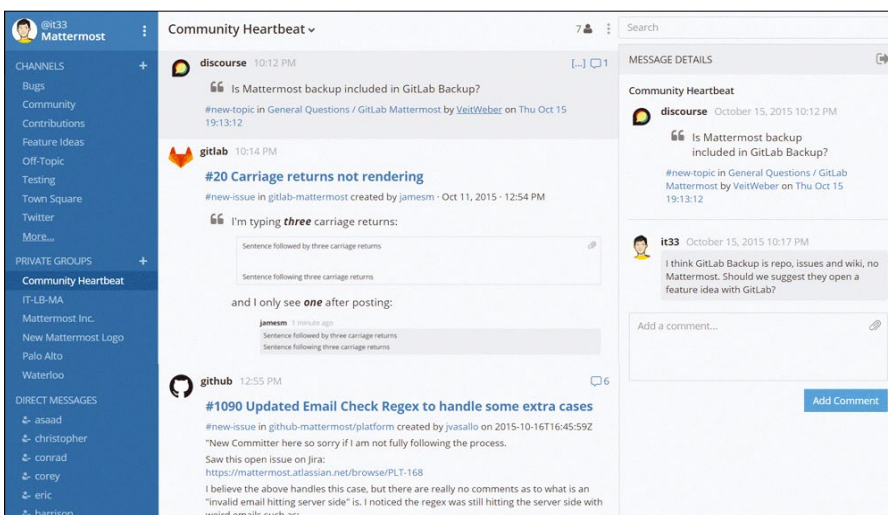


Figure 4: Mattermost has various features for DevOps, such as the integration of GitHub using webhooks.

issues can thus be presented natively in the chat software.

If you are running GitLab with an LDAP connection, you can use GitLab as a single sign-on source to ensure that all employees log on to Mattermost directly with their LDAP credentials. Separate user management provides the option to map a company's hierarchy directly in the communication system in Mattermost. This even works if LDAP is used as a source for user data. An IRC gateway translates between Mattermost and any IRC server.

A Successful Product

All in all, working with Mattermost is great fun, once you have battled your way through the installation. For testing purposes, Mattermost images are available both in Amazon's cloud service and in the form of Docker containers, but to use the service in production requires the manual setup. Matterfront desktop clients are available for OS X, Windows, and Linux, and Mattermost apps for Apple iOS and Android. Mattermost has a corresponding theme for mobile devices.

If the developers were to provide complete packages for various Linux systems and make it possibly to install the software sensibly and automatically, Matter-

most might well be considered the Slack and HipChat killer.

Rocket.Chat

The third test candidate, Rocket.Chat, is positioned in the same segment as Mattermost and aims to provide an alternative to HipChat and Slack.

Rocket.Chat also serves as an all-inclusive package: The service provides a web interface, although access via a native client is also possible. The biggest difference is that Rocket.Chat offers both normal text chat and the exchange of voice messages or video chat with WebRTC. Rocket.Chat itself is implemented in Node.js, meaning the operation requires a functioning Node.js installation.

Before thinking about Node.js, you first need to prepare the target system for Rocket.Chat, which insists on having access to a database. Like Mattermost, Rocket.Chat stores persistent data such as chat content in its database so that it can be accessed later. The Rocket.Chat developers decided on MongoDB as the database of choice, so the MongoDB configuration needs to be completed before Rocket.Chat can be installed.

Cumbersome Installation

Once MongoDB is installed, Rocket.Chat follows: As with InSPIRCd and Matter-

most, the Rocket.Chat developers don't provide packages for SUSE, Red Hat, or Debian/Ubuntu. Again, the installation is via tarball, which you need to unpack on the target system in a directory of your choice. MongoDB is launched by simply calling `node`: It has no init script or systemd unit file. Therefore, you need to create your own if you want to start Rocket.Chat automatically.

In the installation guide, the developers point out that manual installation isn't, in their opinion, the best and fastest way to get a Rocket.Chat instance up and running. Instead, they recommend the use of hosting platforms like Sandstorm [9] or Heroku [10]. However, that would wipe out the advantage of a self-hosted system and control over your data. Although Docker containers are available as a middle ground, many administrators dislike these black boxes because they must rely completely on the benevolence of the container's authors and their willingness to provide updates for the containers as needed (e.g., for security reasons).

Decent UI, Handy Features

A look at the feature list reveals that Rocket.Chat is clearly guided by Slack. Whether the connection is made from a web browser or the native desktop app,

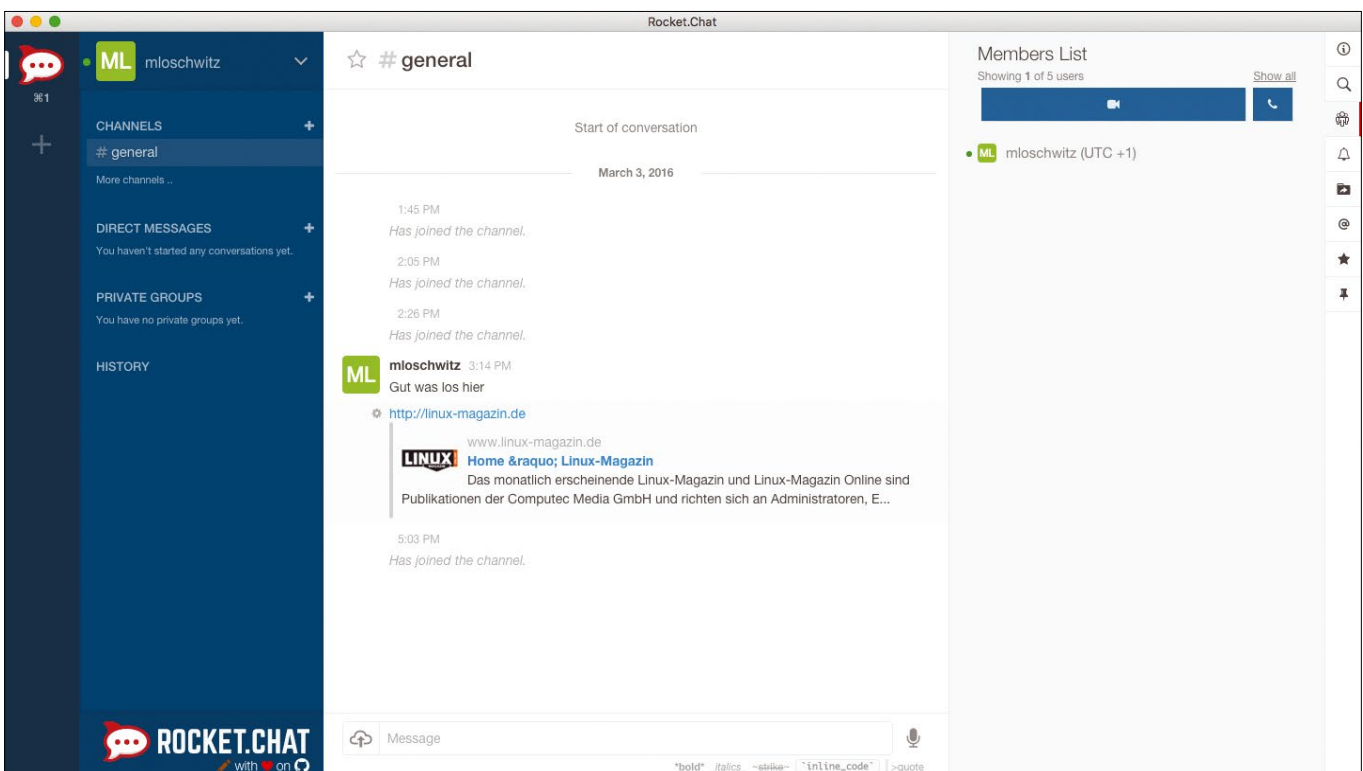


Figure 5: Rocket.Chat is another alternative to Slack and has an automatic URL preview.

Rocket.Chat's graphical interface is pleasant and unobtrusive. Standard features such as the automatic preview display for posted links work perfectly (Figure 5). Rocket.Chat stores files uploaded for sharing with other chat participants permanently on the server. As with Mattermost, you should have enough disk space available.

A Separate API

Rocket.Chat provides its own API that follows the REST principle. For example, according to the public Rocket.Chat documentation, it is possible to log on to the API with user access and to query all unread messages from a specific channel. However, the API has another purpose: Rocket.Chat implements the same functions with its API as Mattermost does with webhooks.

GitHub, GitLab, Jira, and Confluence can therefore be used to access the Rocket.Chat API by URL for certain events. The commit message for a change in GitLab thus ends up directly in the chat channel of the developers involved, if so desired. If you don't want to write the integration yourself, you can use Hubot [11], which makes sure that GitLab and GitHub directories are correctly integrated into Rocket.Chat.

In addition to the API, Rocket.Chat provides a plugin interface that can be used to expand the service at will. For example, the ScreenShare extension can export your display to another user directly from Rocket.Chat. Rocket.Chat provides appropriate plugins for

Chrome and Firefox in the respective extension stores.

If you want to operate more than one Rocket.Chat installation, you will appreciate the federation feature, which can be used to link several servers to form a network. Federation also offers gateway features for IRC, Slack, Skype, and Asterisk. Installing the OAuth extension for Rocket.Chat keeps the user databases of all servers in sync.

Awkward Issue: Native Apps

Only shortly before this article was written did the Rocket.Chat developers released a new version 1.1.0 of their native desktop application. (It was at version 1.3.1 at time of publication.)

Many users don't like web-based chats because they can't be integrated into a desktop as well as an actual chat program. The OS X program icon, which can display a symbol for new messages in the dock, is cited as an example. Also, receiving the chat client's notifications in the system tray is useful because users usually keep an eye on that location.

Rocket.Chat wasn't always up to speed in the past regarding desktop applications: The previous version of Rocket.Chat.Electron – the name of the Rocket.Chat desktop client – was pretty unreliable on OS X, for example. The developers have now solved the problem for v1.2.0.

The Electron project currently offers packages for OS X, Windows, and Linux. The developers also thought about mobile users: Rocket.Chat clients are available both in Google Play and the Apple App Store. (Figure 6).

Conclusions

Hip.Chat and Slack face stiff competition in the form of InspIRCd, Rocket.Chat, and Mattermost. These programs sometimes even outdo their role models in the feature list. However, hosted services still offer one major advantage: The chat system is ready

immediately after logging on. Any further configuration at most extends to the appearance or behavior of the client, not its basic functionality.

If you want to use one of the tools presented here, you will not be able to avoid a certain amount of installation and configuration work. Operating the chat solution is also more complex; for example, security updates for individual components need to be installed on a regular basis to avoid security holes in the chat environment.

In return, however, a company can be sure that their chat content is stored on a separate server, which is an important factor in the EU because companies are bound by laws relating to data protection.

All three approaches presented here have specific advantages and disadvantages. Anyone looking for a communication solution for teams should certainly consider hosted solutions rather than giving Slack or Hip.Chat their money and data. Mattermost appears to be the strongest candidate at the moment, even if the lead is only marginal. If you want a proven protocol and versatility in terms of clients, combined with a few frills, InspIRCd is a good choice. ■■■

INFO

- [1] HipChat: <https://www.hipchat.com>
- [2] Slack: <https://slack.com>
- [3] InspIRCd: <http://www.inspircd.org>
- [4] Mattermost: <http://www.mattermost.org>
- [5] Rocket.Chat: <https://rocket.chat>
- [6] RFC 2812: <https://tools.ietf.org/html/rfc2812>
- [7] UnrealIRCd: <https://www.unrealircd.org>
- [8] Kiwi IRC: <https://kiwiirc.com>
- [9] Sandstorm: <https://sandstorm.io>
- [10] Heroku: <https://www.heroku.com>
- [11] Hubot: <https://github.com/github/hubot>

AUTHOR

Martin Gerhard Loschwitz works as a cloud architect at SysEleven in Berlin. He is also an official member of the Debian project and has been a Debian developer for more than 12 years.




Figure 6: Rocket.Chat has clients for iOS and Android.

PRESENTED BY BARNES & NOBLE

WORLD




NEW YORK
HALL OF SCIENCE

7TH ANNUAL

Maker Faire®

ny sci

GREATEST
Show
& TELL
ON
Earth



OCT 1+2 7TH ANNUAL WORLD MAKER FAIRE
NEW YORK HALL OF SCIENCE



PRESENTING

BARNES & NOBLE

GOLDSMITH SPONSORS



SILVERSMITH SPONSORS



Makeblock



GET YOUR
TICKETS
TODAY!



makerfaire.com

BROUGHT
TO YOU BY

Make:



Data backups, images, and more with Qt4-fsarchiver

Fast Copier

Qt4-fsarchiver lets you back up files, complete partitions, and create disk images with a single mouse click. *By Erik Bärwaldt*

Linux supports a variety of filesystems, includes security software for encryption, and can handle a variety of compression methods and formats. The Linux environment thus offers perfect conditions for handling system and data backups in heterogeneous environments. The graphical Qt4-fsarchiver tool removes virtually all limitations.

Qt4-fsarchiver is a graphical version of the very useful fsarchiver tool, which the project developers describe as a "... tool that allows you to save the contents of a file-system to a compressed archive file. The file-system can be restored on a partition which has a different size and it can be restored on a different file-system."

Qt4-fsarchiver is included in the repositories of Mandriva derivatives Mageia, OpenMandriva, ALT Linux, and Rosa Linux. On Ubuntu, you need to add the appropriate PPA [1] to the package sources before easily installing the software using the package manager. For Linux Mint, Debian "wheezy," and Debian "jessie" [2], as well as openSUSE and Fedora [3], you'll find precompiled packages on SourceForge. If you are looking for a special treat, try the Qt4-fsarchiver Live DVD, which weighs in at approximately 960MB. The Live DVD is

based on Ubuntu 14.04.2 LTS and includes UEFI support.

After completing the installation you will find an icon for Qt4-fsarchiver in the main menu below Tools, System Tools, or Utilities depending on the distribution. The program requires root privileges, the system prompts you for the administrator password each time you launch it. Please keep in mind that the program relies on the fsarchiver command-line application, which must also be installed. When installing via the package manager, fsarchiver often automatically ends up on your hard disk. However, if you retrieve the software from a third-party source, you typically need to install the command-line tool manually on your disk.

Program Interface

Qt4-fsarchiver comes with a not entirely intuitive interface that just looks complex to use. It is thus important to be familiar with the most important functions of the program, although detailed knowledge of mass media partitioning can prove to be helpful. Qt4-fsarchiver first pops up a message window at startup pointing you to the `Readme` file. However, it actually specifies an incorrect directory; the file is in fact located

below `/usr/share/doc/qt4-fsarchiver/doc/`. Take a look at the content; it offers both some installation instructions and tips on using the program.

After the dialog closes, the application window opens. Besides a menubar at the top and a button bar on the left below it, you will find an area with a variety of settings, and two list windows on the right show you all the identified partitions and file structures in a tree view. The partition list also shows you the filesystems on the partitions. However, Qt4-fsarchiver only fully supports Btrfs, NTFS, and ext2/3/4; the outdated FAT(32) and VFAT are only partially supported. USB flash drives are therefore only useful as a source or target if you format them with one of the supported filesystems. ISO9660 and ZFS remain completely sidelined. Qt4-fsarchiver supports the established XFS filesystem, however, so the program can also be useful in the server environment.

At the bottom in the program window is a tray with time entries showing how many files and directories you added to the backup and how many are already backed up. A progress bar completes the visual display of the backup operation (Figure 1).

Lead Image © Volodymyr Vasylykiv, 123RF.com

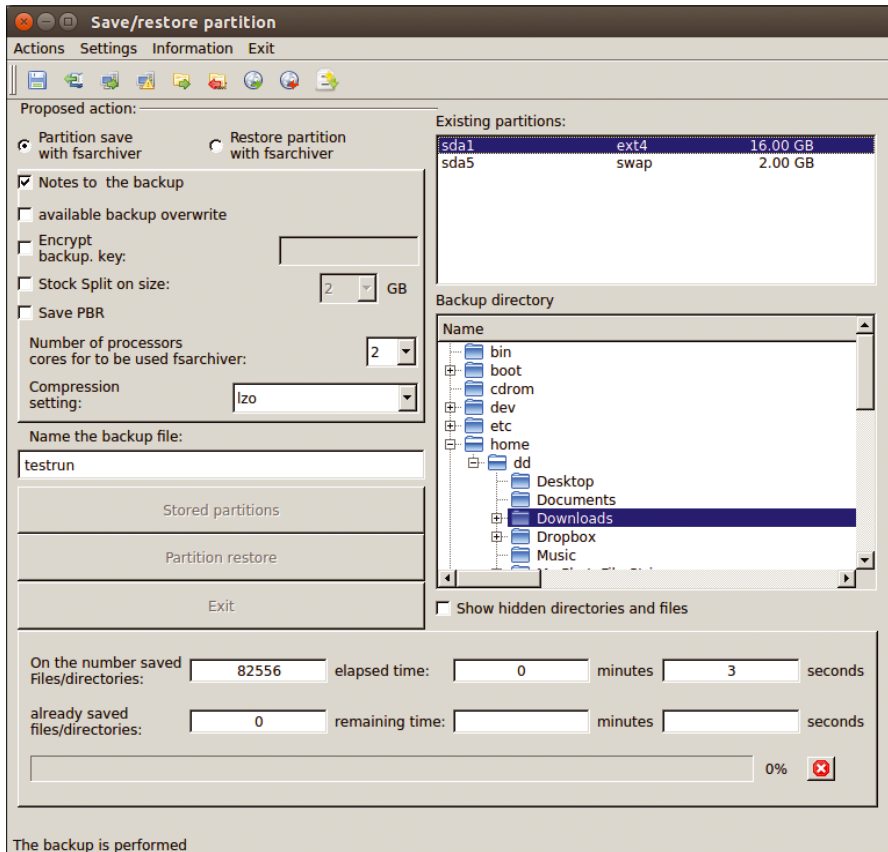


Figure 1: The Qt4-fsarchiver application window looks somewhat confusing at first glance.

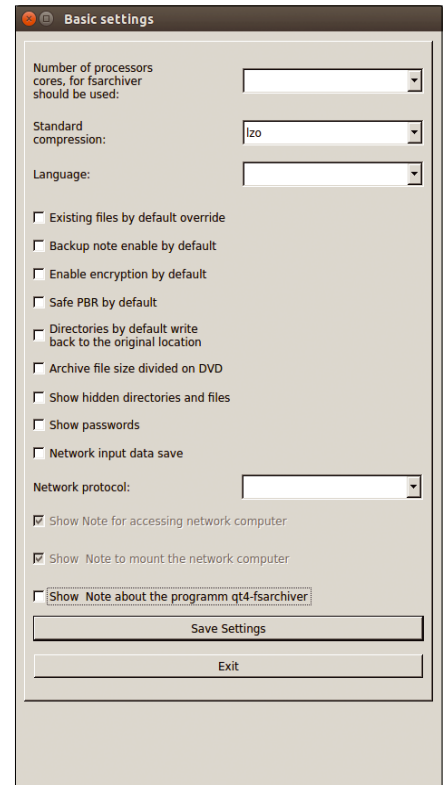


Figure 2: The basic settings are used if you do not define your own options for individual backups.

Qt4-fsarchiver offers many settings and options for the backup run in the program window. Additionally, you can go to *Settings | Basic settings* to complete a basic configuration if you do not change any of the default options for a backup run. In a neatly arranged dialog, you can define how the software backup behaves for backups and restores and specify different options for backing up on computers in the network. After completing the changes, save them by pressing the *Save Settings* button (Figure 2).

Backup Run

The graphical interface implements almost the entire instruction set of the fsarchiver command-line program; this means that you initially need to make some settings on the left in the program window. When backing up a partition, you can first define whether you want to encrypt the backup. Check the *Encrypt backup key* option to do so and enter the desired key with a length of 6 to 64 characters to the right in the input box.

During the backup, large partitions with extensive databases can be split into several parts, for example, to match

the size of individual sections to the capacity of a USB stick or a CD/DVD. To back up the first block of a bootable partition, also check the option *Save PBR*.

On multicore systems, you can significantly speed up the backup run by telling the program to use multiple processor cores. To the right of the option, type the appropriate *Number of processor cores to be used for fsarchiver*. Then, define the Qt4-fsarchiver backup compression format you want to use. The corresponding drop-down menu lists a variety of compression methods – with different compression rates – that can make the backup significantly smaller depending on the data types.

Next, you can define a name for the backup in the box at the bottom.

Following this, you still need to select the drive to be backed up in the list of *Existing partitions* on the right. Finally, choose the target directory to which Qt4-fsarchiver will write the backup in the directory tree below.

To start the backup run, click *Save partition* center left in the application window. Qt4-fsarchiver now shows you a small window with a summary of the

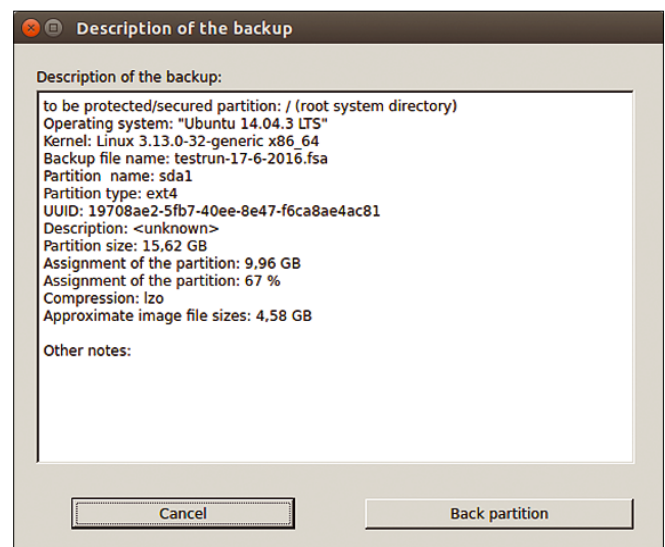


Figure 3: Qt4-fsarchiver summarizes all settings prior to the backup in a separate window.

task (Figure 3) and after clicking the *Back partition* button, stores the compressed backup at the specified location. During the backup, additional information is displayed at the bottom of the main program window. After the backup is complete, Qt4-fsarchiver presents statistics on the number of backed up directories, files, and links in another message window. The software also stores a small text file with this data in the target directory.

No operating system will boot without the master boot record or its successor for large disks with a capacity of more than 2TB, the GUID partition table (GPT). Thus, it's advisable to create a separate backup for the MBR or GPT. The appropriate backup and restore functions can be found in the *Actions* menu below *Save MBR/GPT* and *Restore MBR/GPT*. During the backup, you specify the backup directory in a dialog and – if the computer has multiple storage media – the hard drive or SSD whose MBR or GPT you want to back up. With a single click on *Save MBR/GPT*, you can store a copy of the relevant sectors at the desired location. Similarly, you can restore a previously saved MBR or a GPT in the same dialog.

Qt4-fsarchiver also offers a quick backup and restore function for individual directories, which can also recursively backup or restore. The corresponding settings dialog is opened by se-

lecting *Actions | Save directory*; you then only need to specify the source and target. The program compresses the backup by default; you can also choose from several compression formats and optionally encrypt the backup. Click on *Directory save* to start the backup run (Figure 4).

Backup via LAN

On networks, Qt4-fsarchiver also lets you create complete backups on other computers on the LAN. To do so, you need to install a Samba client on the system to be backed up, because either a Windows PC or a Linux system with a directory shared via CIFS VFS is used as the backup target. In any case, you must share a suitable directory for storing the backup on the target.

Next, select *Actions | Save partition over a network* to call the routine for backing up to a remote host. The software now automatically checks the shares on the network and presents a short dialog for authenticating against the target machine. Then Qt4-fsarchiver mounts the share and presents the same dialog as when creating local partition backups.

Images

Qt4-fsarchiver can clone entire hard drives or SSDs on request. In the *Actions* menu select the *Clone Harddrive, create/restore disk image* entry. In the

source drive and the file path where you want Qt4-fsarchiver to store the image. Then, click on *Create harddrive image* at bottom right. In the same window, you can also restore an existing image on a mass storage device by selecting the *Image of a hard disk write back* option. The corresponding dialogs do not offer any further selection options, such as the number of cores or encryption mechanisms. The program does compress the images, however.

Data Recovery

To restore a backup, got to the *Actions* menu and select the *Restore partition* entry. The program now shows same dialog as when creating a backup, but grays out irrelevant boxes and automatically enables the option *Restore partition with fsarchiver*.

If you encrypted the backup, you must enter the associated key in the *Decrypt backup key* box. Again, you then define how many processor cores Qt4-fsarchiver should use to decrypt. In the last step, specify the target partition and enter the relevant file name in the *Name the backup file* box. For the restore, the target partition must be unmounted; otherwise, the software cannot restore any data.

Apart from this, Qt4-fsarchiver is very flexible: For example, the target partition can be formatted with a different filesystem than the original source. The program easily restores the data to the new partition, restoring the original formatting in the process. Also, restoring the data to a smaller partition than the original source drive works. You obviously need to make sure that the uncompressed data will fit in the smaller target partition. The fact that Qt4-fsarchiver stores a small text file containing statistical data with each backup is a very useful feature: This information helps you determine the original footprint of the uncompressed data.

You restore individual directories or restore a backup from a remote computer in the same way; the dialogs are similar to those for backing up. You can launch the routines via *Actions | Restore partition restore over a network* or *Actions | Restore directory*. The program shows you the restored backup files continuously in a list window. If Qt4-fsar-

chiver opens (Figure 5), the software shows the mass storage devices. Select the hard disk or SSD to clone, specify the destination drive, and then click the *Clone Harddrive* button at the bottom right in the window.

If you want to generate an image file of a mass storage device, select the *Image of a hard disk to create* radio button in the same window and choose the

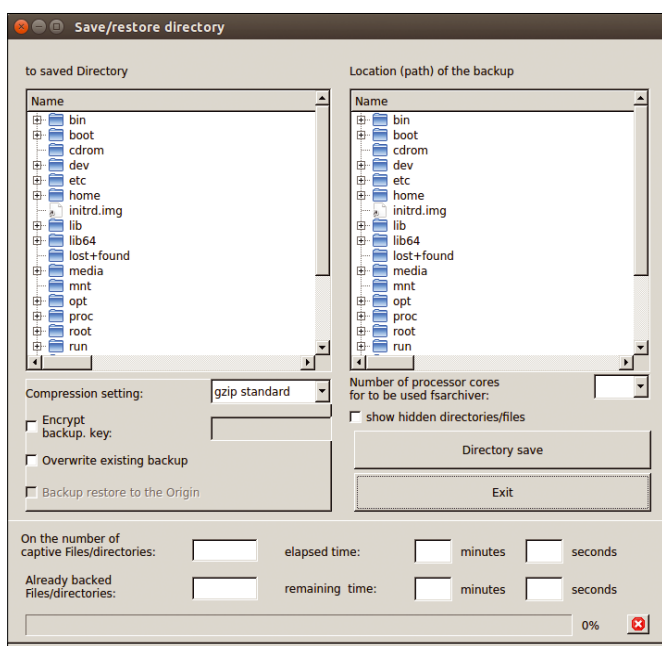


Figure 4: Using this dialog, you can back up individual directories in Qt4-fsarchiver.

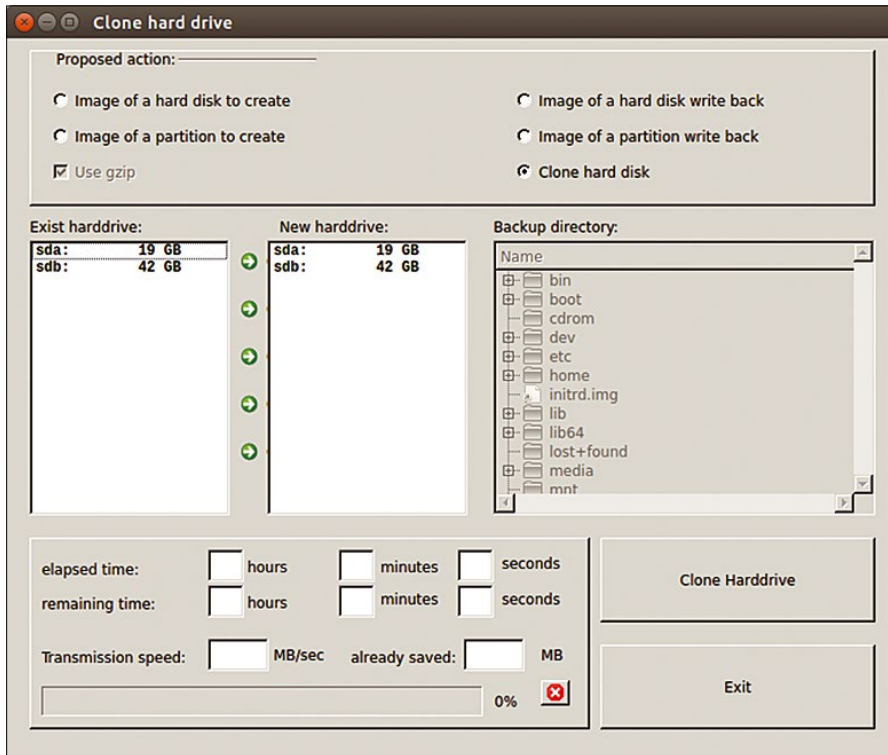


Figure 5: With a few clicks of the mouse, users can clone a mass storage device or create an image.

chiver is unable to restore one or more files, it ignores them and continues restoring the next file.

Live Variant

The Live Qt4-fsarchiver CD is primarily used to back up Windows machines and restore them after a failure or loss of data. You can also deploy it for cases in which you want to restore a backed up partition on a drive that normally would be mounted in the filesystem in stationary operation.

In addition to the backup software, the Live medium provides two useful tools for detecting rootkits under Linux in the form of the RootKit Hunter (rkhunter) [4] and chkrootkit [5] programs. Fortunately, the developers of the Live version eliminated anything that did not focus on the specific purpose of the system. Therefore, the Unity user interface is amazingly clean and can be used immediately, even by novices (Figure 6).

The classic program menu in the top panel is missing. The start menu is arranged vertically on the left edge of the screen and mainly offers launchers for file and disk management software: You will thus also find GParted and a System Profiler in the menu next to Qt4-fsarchiver. You can also run the Terminal to

access the command-line tools without going through a menu.

The Live version also does not support all filesystems. For example, the Live medium will not run on systems manufactured by Apple that use the HFS and HFS+ filesystems. The Live version is not suitable for BSD derivatives with the ZFS filesystem. However, no functional

restrictions are expected in heterogeneous environments with current Microsoft operating systems or Linux-only networks.

Conclusions

Qt4-fsarchiver proves to be a very useful piece of software when you need a simple option for backing up and restoring files and partitions. However, it requires a short learning curve and, for backups over the network, a remote computer with a proprietary communication protocol or operating system. Qt4-fsarchiver shows its flexibility on standalone systems under Linux, where it can replace several backup programs that are only capable of backing up parts of a computer system. In the test, Qt4-fsarchiver was extremely stable, reliable, and fast, and its use of resources on less powerful systems is impressive. ■■■

INFO

- [1] Qt4-fsarchiver (PPA): <https://launchpad.net/~dieterbaum/+archive/ubuntu/qt4-fsarchiver>
- [2] Qt4-fsarchiver (DEB): <http://sourceforge.net/projects/qt4-fsarchiver/files/deb-Pakete/>
- [3] Qt4-fsarchiver (RPM): <http://sourceforge.net/projects/qt4-fsarchiver/files/rpm-Pakete/>
- [4] RKHunter: <https://en.wikipedia.org/wiki/Rkhunter>
- [5] Chkrootkit: <http://www.chkrootkit.org/>

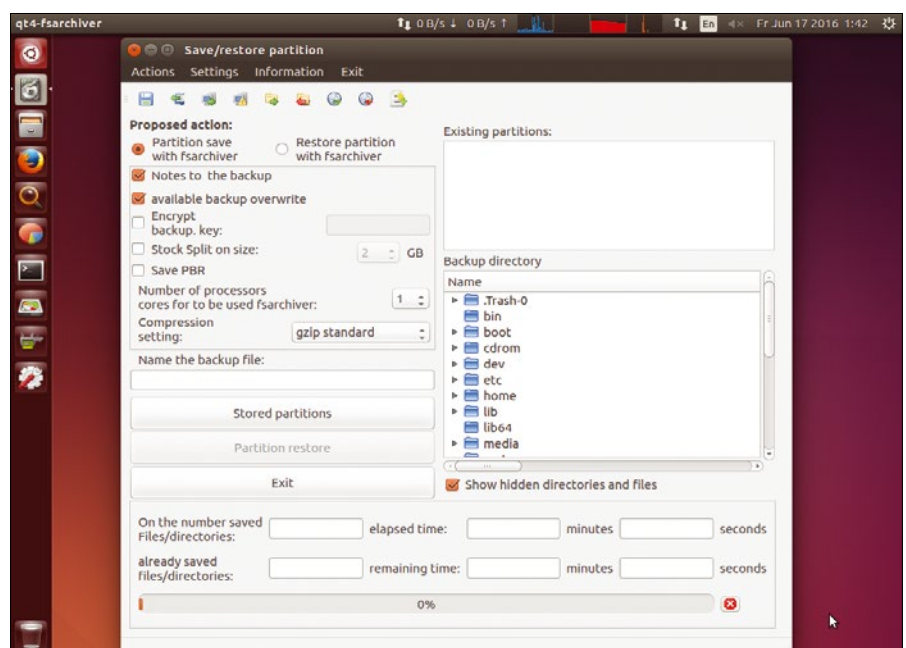


Figure 6: The Live version is limited to meaningful software and very clear-cut compared with Ubuntu.

Your Roadmap to the Open Hardware Revolution ...

An exciting world of projects, tips, and skill-building tutorials awaits you in every issue of Raspberry Pi Geek.

Order your subscription today and tune in to the revolution!

shop.linuxnewmedia.com

Print Sub

Carry our easy-to-read print edition in your briefcase or backpack - or keep it around the lab as a permanent reference!

Digital Sub

Our PDF edition is a convenient option for mobile readers.

6 print issues with 6 DVDs or
6 digital issues for only

\$59.95 £37.50 €44.90



We're also in
Google Play
and the
iTunes Store!



iTunes Store




Google US



Google UK



Want to find out what's in the next issue?



Sign up for our newsletter
www.raspberry-pi-geek.com/mc/subscribe

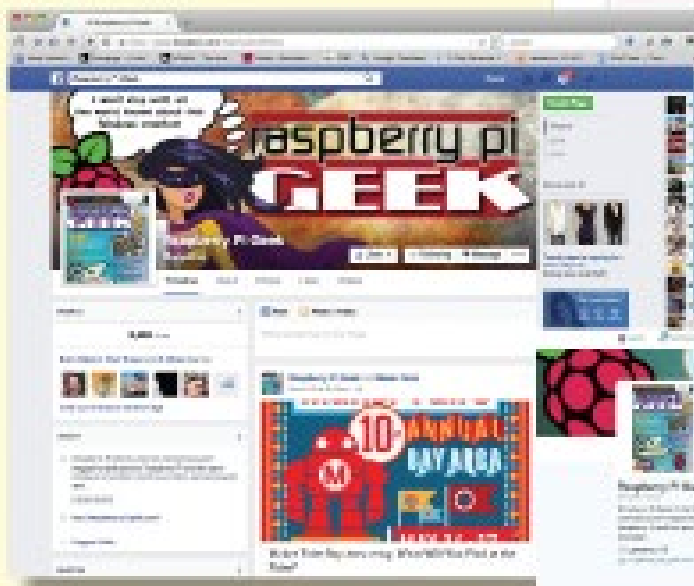
Visit us online
www.raspberry-pi-geek.com



Like us on Facebook
www.facebook.com/RasPiGeek



Follow us on Twitter
[@RasPi_Geek](https://twitter.com/RasPi_Geek)



Klaus Knopper answers your Linux questions

Ask Klaus!

By Klaus Knopper

New Linux fdisk Partitioning Tool

? I've read your article about repartitioning the hard disk using the famous `fdisk` command-line tool. I tend to partition a disk in a way that starts and ends partitions at cylinders, in order to adapt to the mechanical specs of the disk. However, the newer versions of `fdisk` that come with Knoppix 7.7 no longer show the disk geometry in cylinders and heads; all settings are based on sectors of 512 or 4096 bytes, depending on the disk's physical sector size.

Do I have to install an old version of `fdisk` in order to restore the old, cylinder-based behavior?

💡 To illustrate what happens, start an `fdisk` session. Unless you quit `fdisk` with the `w` command for "write," nothing will be changed or saved on disk, so make sure you quit with `q` when leaving the program (unless you

really want to change the disk's partitioning).

```
$ sudo fdisk /dev/sda
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only
until you decide to write them.
Be careful before using the write
command.

Command (m for help):
```

Next, list the current partitions (Listing 1). As you can see, there is one "dos" (or UEFI) FAT32 partition, two Linux data partitions, and one swap partition. All Starts and Ends are given in 512-byte "sectors," whereas the data range is stated in blocks of 1KB. (This may be confusing at first glance, but the partition size is actually easier to think about in decimals than fractions of a kilobyte.)

So, how to get to the physical disk layout from there? The `m` command of `fdisk`



KLAUS KNOPPER

Klaus Knopper is an engineer, creator of Knoppix, and co-founder of LinuxTag expo. He works as a regular professor at the University of Applied Sciences, Kaiserslautern, Germany. If you have a configuration problem, or if you just want to learn more about how Linux works, send your questions to: klaus@linux-magazine.com

LISTING 1: Listing Partitions in fdisk

```
Command (m for help): p

Disk /dev/sda: 30.0 GB, 30016659456 bytes, 58626288 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000c12bc

   Device Boot      Start         End      Blocks   Id  System
/dev/sda1  *           2048        411647        204800   c   W95 FAT32 (LBA)
/dev/sda2                411648       25577471       12582912   83   Linux
/dev/sda3           25577472       50743295       12582912   83   Linux
/dev/sda4           50743296       58626287        3941496   82   Linux swap / Solaris
```

LISTING 2: fdisk Commands

```
Command (m for help): m
Command action
  a  toggle a bootable flag
  b  edit bsd disklabel
  c  toggle the dos compatibility flag
  d  delete a partition
  g  create a new empty GPT partition table
  G  create an IRIX (SGI) partition table
  l  list known partition types
  m  print this menu
  n  add a new partition
  o  create a new empty DOS partition table
  p  print the partition table
  q  quit without saving changes
  s  create a new empty Sun disklabel
  t  change a partition's system id
u  change display/entry units
  v  verify the partition table
  w  write table to disk and exit
  x  extra functionality (experts only)
```

reveals the option (Listing 2).

As you can guess from the **g** option, this version of **fdisk** also handles the GPT/GUID partitioning scheme, which is often used for disks larger than 2TB. Option **u** will switch display and input of partition limits to the traditional heads/cylinders/sectors scheme (Listing 3).

Now the start/end values are given in cylinders, while partition size is still displayed in kilobyte blocks. Note the “deprecated” warning message: Modern disks don’t necessarily report physical “cylinders.”

Instead of changing units inside the program, **fdisk** can be started with “cylinder” units as the default unit, restoring the old behavior, by adding **-u=cylinders** after the **fdisk** command.

Listing 4 shows a quick look into the “expert mode” partition display, splitting the partitions addresses again into cylinders, heads, and sectors:

NAND flash-based solid-state disks (SSDs) in particular do not even have physically magnetic media organized in cylinder units; however, the operating system might require this information report in cylinders, heads, and sector size of tracks to calculate the correct size of the disk the traditional way.

Therefore, cylinder-size-based limits might not align to the physical medium or memory chips as desired when partitioning along “cylinder” limits; instead, the preferred (and probably more honest) method is now working with sector units, which also allows a more fine-grained space adjustment of partitions. You can still use the “cylinder” unit setting for partitioning, but it most likely won’t give you any advantage over partitioning in sector units.

Don’t forget to leave **fdisk** with **q**, to leave with no changes written to disk.

LISTING 3: Changing fdisk Display Units to Cylinders

```
Command (m for help): u
Changing display/entry units to cylinders (DEPRECATED!).

Command (m for help): p

Disk /dev/sda: 30.0 GB, 30016659456 bytes, 58626288 sectors
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000c12bc
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sda1	*	1	26	204800	c	W95 FAT32 (LBA)
/dev/sda2		26	1593	12582912	83	Linux
/dev/sda3		1593	3159	12582912	83	Linux
/dev/sda4		3159	3650	3941496	82	Linux swap / Solaris

LISTING 4: fdisk Expert Mode


```
Command (m for help): x

Expert command (m for help): p


Disk /dev/sda: 255 heads, 63 sectors, 3649 cylinders
```

Nr	AF	Hd	Sec	Cyl	Hd	Sec	Cyl	Start	Size	ID
1	80	32	33	0	159	6	25	2048	409600	0c
2	00	159	7	25	31	39	568	411648	25165824	83
3	00	31	40	568	159	9	86	25577472	25165824	83
4	00	159	10	86	80	63	577	50743296	7882992	82

Booting Knoppix from USB Flash Disk

 I installed Knoppix on a small computer with a tiny internal hard disk by flash-knoppix, so it boots with the same method it does from DVD or USB flash disk, saving space by the base system staying compressed. The overlay data partition is encrypted. Now I would like to extend partition sizes to fill up the entire disk and readjust the filesystem by booting from another Knoppix-installed USB flash disk so that none of the hard disk partitions are mounted. But, apparently, booting from the USB flash disk still results in both hard disk partitions being mounted, and I’m stuck at the boot screen asking me for my encryption password.

How can I boot from USB when Knoppix is already installed on the internal hard disk; is it possible?

 The reason for Knoppix seeming to ignore the USB drive is that the operating system can’t know from which media its kernel was loaded and started.

After the CPU switches to “protected mode,” the computer “forgets” which medium was selected in the BIOS boot menu. Knoppix uses a standard search order that tries DVD drives first, then hard disks, and flash drives in the order detected by the kernel and initial ram-disk scripts. Because of this, it might find your KNOPPIX directory on the hard disk first, before the DVD or USB flash disk are even ready to use.

You can tell Knoppix to search for its files and setup for overlay images by a boot option, thereby overriding the standard search path:

```
knoppix fromhd=/dev/sdb1
```

Assuming that your flash drive is `/dev/sdb`, the first partition containing the KNOPPIX data would be `/dev/sdb1`. The `fromhd` option just means to search for a partition on a specific medium giving its Linux device name; it is not specific to hard disks.

Furthermore, you can skip mounting the encrypted image or partition by adding boot option `noimage`. In the case of an encrypted overlay partition, you can also just enter an arbitrarily wrong encryption password three times (two retries), so mounting the encrypted partition will fail, and the system will continue to boot without the encrypted overlay.

Using the PHP mail() Function in Apache

? I’m learning to write dynamic web pages using Apache with PHP on Linux. I stumbled across the `mail()` function of PHP, which is supposed to

send the text enclosed in the function parameters, so I can send data input from the web page by email to my mail account instead of writing to a database or file. But, it does not seem to work, Mail is never sent out. What am I missing?

💡 In PHP, `mail()` requires a working setup of a mail delivery agent on the computer running Apache/PHP; that is, it calls `sendmail` to send mail to an email address.

If `sendmail` is not installed, or if it isn’t configured in a way that allows sending mail to the desired addresses, the `mail()` function of PHP will fail silently.

To check whether `sendmail` is configured correctly, you can invoke the command manually and try to send email. As `<mailaddress>`, use `<your.username>@<your.providers.mail.domain>` for sending email over the Internet, or `<username>@localhost` for sending mail to a local user present on the computer running Apache/PHP.

```
$ sendmail <mailaddress>
From: Klaus <mailaddress>
To: Klaus <mailaddress>
Subject: Test
Hello, World!
.
```

The single dot in the last input line tells `sendmail` to finish the message text body and send it off.

Now, if you get an error message (or Delivery Status Report – DSN) telling you the mail could not be sent, or the `sendmail` command is not even installed,




Figure 2: Postfix setup.

you need to install and configure local email. Most distros use Postfix, which is a very complete mail server and transport agent that fits the majority of mail server purposes, including the possibility to add spam filters. It also has a choice of a few straightforward and easy basic settings when installed first (Figure 1).

On Debian-based systems, such as Knoppix and Ubuntu, installation can be done with

```
sudo apt-get install postfix
```

which will run an interactive basic setup during installation. If Postfix is already installed, you get back to the initial setup (Figure 2) with the command:

```
sudo dpkg-reconfigure postfix
```

Local only is the safest option, but it only allows sending mail to users on the local system running Apache + PHP. *Internet Site* will attempt to send mail directly to the mail server of the designated recipient address, and allows receipt of mail via the SMTP protocol on the machine running Postfix. Thus, it’s a potential target for spammers.

With Postfix installed and configured, the `sendmail` command should now be available for the PHP `mail()` function to work correctly.

The following lines,

```
<?php
$message = "Hello, World";
if( mail('user@mailserver', '
    'My subject line', '
    $message) )
    echo "Mail sent out successfully.";
else
    echo "Mail could not be sent.";
```

are an excerpt of an HTML page containing PHP code for sending mail. ■■■

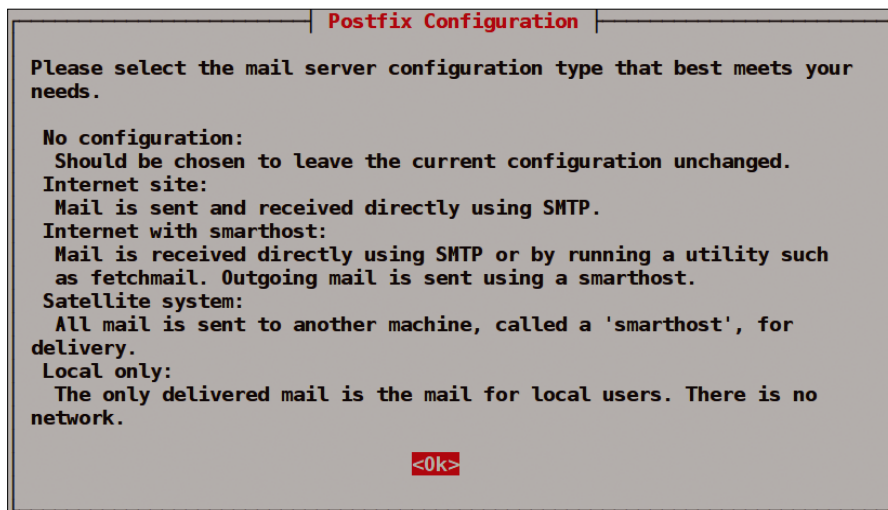


Figure 1: Initial Postfix configuration screen.

RISK-FREE TRIAL!



Practical.
Technical.
Elegant.

SUBSCRIBE NOW


**3 ISSUES
+ 3 DVDs**

\$3

shop.linuxnewmedia.com

£3 / €3 / \$3 / \$9 rest of world
Depending on the region you live in.
Terms and conditions: <http://goo.gl/SSSQer>





Storing metadata in files

EMBEDDED

With the Python XMP Toolkit, programmers can develop applications that store metadata in files. *By Daniel Stender*

Metadata is information about the contents of objects like graphic files, digital videos, or documents and can include the name of a photograph, the author, and the project, as well as details about the license used to publish the item. Subject catalogs, thumbnails, or coordinates of image areas containing faces or other items are also part of the larger realm of metadata.

A range of standards control the exchange of datasets across digital objects. The prevailing technique is to retain the metadata in the file itself and copy that data along with the file. For example, a widespread technical standard for metadata in digital photography is the Exchangeable Image File (Exif) format [1] for JPEG, JFIF, and TIFF images.

With Exif, digital cameras and smartphones can usually write a full range of

technical information for inclusion in image data, such as date and time, camera model, focal length, and shutter speed. The standard even includes thumbnails and GPS coordinates.

XMP

Adobe introduced the XMP (Extensible Metadata Platform) standard in 2001. The current version appeared as a three-part specification between 2012 and 2014 [2], with ISO 16684-1-2012 as the core component [3]. XMP provides a comprehensive system for describing the metadata for digital media, consisting of a data model, a serialization or implementation model, and directions for embedding metadata in various file formats.

XMP's data model first stipulates that the metadata within files should all reside in a so-called "packet". That data comes in three variants, as follows: (1) Simple data fields assign an individual content item (e.g., the date of origin) to an element. (2) Structured fields bind an element with several pieces of content (e.g., the dimensions of a document). (3) Lists present data in different ways: as sequential and ordered lists, unordered lists (e.g., keywords),

and alternative lists (e.g., a title in different languages).

Embedded

XMP-compatible software embeds the packet directly in the objects described. In doing so, it supports a large range of file formats. XMP packets can be written in image files (JPEG, PNG, TIFF), multimedia formats (AVI, MPEG, MOV, MP3), and digital documents (PDF, PS), as well as in many other formats. The software sets up the metadata in a way that also allows other programs unable to handle the packets to process the files' content without problem.

In TIFF images, for example, XMP uses a block in the image file directory (IFD) with the tag number `0x02BC`; the data slots into PDF files, formatted as metadata objects of the *XML* subtype. For unrecognized formats, such as DjVu, the XMP software creates packets in the form of additional sidecar files with the `.xmp` file extension.

You can set the scale of an XMP packet within a file flexibly. The key feature of XMP, however, as its name makes clear, is its expandability. The XMP specification describes many data fields for differ-

AUTHOR

Daniel Stender <http://www.danielstender.com> has been keeping busy e-documenting since finishing a humanities degree. He is a Debian developer and, among other things, the maintainer of the Python XMP Toolkit.

Lead Image © Nasir Khan, 123RF.com

ent purposes; XMP users can augment a file's metadata with their own elements, offering practically unlimited opportunities to expand XMP as desired. For instance, in PDF files it can store identification data used and designed internally for scanned invoices and letters in paperless offices.

Metadata of Exif and other standards may appear side by side in the same graphic file with XMP. Exif tags can also be written in XMP, allowing you to place Exif data in file formats that do not support it by design (e.g., PNG and GIF).

In practice, however, writing several instances of metadata to a single file frequently becomes chaotic. If you want to use individual metadata standards uniformly in image files, you are best advised to follow the guidelines set for this purpose by the Metadata Working Group [4], a consortium of well-known producers of hardware and software.

RDF/XML

XMP packets are programmed in World Wide Web Consortium RDF/XML [5], an XML syntax for serializing Resource Description Framework (RDF) graphs. For the most part, two points could keep the practitioner busy. First, to deal with XMP, you at least need basic XML knowledge. Second, the syntax always disassembles the elements to include individual data (e.g., *title* and *date of origin*) in groups or classes sorted by given XML namespaces.

The elements defined by the Dublin Core Metadata Initiative are among the primary components of XMP. Dublin Core [6] is a rudimentary set comprising 15 core elements belonging to the `dc` namespace, which also uniformly describes objects of different types, such as photos, videos, and books.

It is therefore mostly elements like `dc:title`, `dc:creator`, `dc:description`, and so on that are present in XMP packets of various file formats. The description elements in the `xmp` namespace function as the further core components of XMP: for instance, `xmp:CreateDate` and `xmp:CreatorTool`.

The `xmlns` element creates namespaces and generally binds them with a uniform resource identifier (URI) – usually an HTTP address. If you know XML, you know that such HTTP addresses are only for registration. They do not need to con-

tain real websites on the Internet. The creators can define namespace identifiers on publication of their sets. The `dc` namespace, for instance, belongs to the URI <http://purl.org/dc/elements/1.1/>, and the `xmp` namespace belongs to <http://ns.adobe.com/xap/1.0/>.

Reading Lamp

Users can very easily sort the content of an XMP packet with the `exempi` command-line tool by using the `-x` option. This tool comes from the library of the same name found on Debian in the `exempi` package and in `exempi-tools` on openSUSE.

Listing 1 shows a shortened version of a typical XMP packet. The `x:xmpmeta` root element first wraps the data in the XMP packet. After this, the `rdf:RDF` tag creates the `rdf` namespace with the URI <http://www.w3.org/1999/02/22-rdf-syntax-ns#>.

The XMP elements always reside in `rdf:Description` blocks. For these, the element `rdf:about` proves to be obligatory, although it always remains empty in XMP. XMP data is written in different styles. Commonly, the elements of a description class are collected in a `Description` environment, and then a respective namespace is created.

LISTING 1: PDF File XMP Packet

```

01 <x:xmpmeta xmlns:x="adobe:ns:meta/" x:xmp:tk="Exempi + XMP Core 5.1.2">
02 <rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#">
03 <rdf:Description rdf:about=""
04   xmlns:dc="http://purl.org/dc/elements/1.1/">
05 <dc:format>application/pdf</dc:format>
06 <dc:title>
07 <rdf:Alt>
08 <rdf:li xml:lang="x-default">XMP Specification Part 1: Data Model,
   Serialization, and Core Properties</rdf:li>
09 </rdf:Alt>
10 </dc:title>
11 <dc:creator>
12 <rdf:Seq>
13 <rdf:li>Adobe Systems Incorporated</rdf:li>
14 </rdf:Seq>
15 </dc:creator>
16 </rdf:Description>
17 <rdf:Description rdf:about=""
18   xmlns:xmp="http://ns.adobe.com/xap/1.0/">
19 <xmp:CreateDate>2012-03-21T08:55:04Z</xmp:CreateDate>
20 <xmp:CreatorTool>FrameMaker 8.0</xmp:CreatorTool>
21 <xmp:ModifyDate>2012-03-21T08:57:17-07:00</xmp:ModifyDate>
22 <xmp:MetadataDate>2012-03-21T08:57:17-07:00</xmp:MetadataDate>
23 </rdf:Description>
24 <rdf:Description rdf:about=""
25   xmlns:pdf="http://ns.adobe.com/pdf/1.3/">
26 <pdf:Producer>Acrobat Distiller 8.1.0 (Windows)</pdf:Producer>
27 <pdf:Copyright>Copyright 2012, Adobe Systems Incorporated. All rights
   reserved.</pdf:Copyright>
28 </rdf:Description>
29 <rdf:Description rdf:about=""
30   xmlns:xmpMM="http://ns.adobe.com/xap/1.0/mm/">
31 <xmpMM:DocumentID>uuid:3806c052-1c0d-4599-b3ff-38c0441bbb3a</xmpMM:DocumentID>
32 <xmpMM:InstanceID>uuid:abc82103-c360-4b74-8812-ff4c81eacea</xmpMM:InstanceID>
33 </rdf:Description>
34 </rdf:RDF>
35 </x:xmpmeta>

```

Three Dublin Core elements – `dc:format`, `dc:title`, and `dc:creator` – are found in the XMP packet. The title of the document is in an alternative list (`rdf:Alt`) in several language versions (`xml:lang`), and the publisher is in a sequential list (`rdf:Seq`).

Core XMP elements such as `xmp:CreateDate` and `xmp:ModifyDate`, which bear date stamps, are found in a further `Description` environment. It states here that Framemaker 8.0 has produced this document (`xmp:CreatorTool`, line 20). Specific description elements for PDF documents (`xmlns:pdf`) and data fields for media management (`xmlns:xmpMM`) follow in other blocks. Table 1 delivers an overview of some XMP elements and classes.

Python XMP Toolkit

XMP applications can be programmed without great effort with the help of a few Python libraries (see the “XMP and Exif with Python” box). The Python XMP Toolkit [7] was developed by the European Space Agency (ESA), among others, to manage images from the Hubble Telescope (Figure 1). The current version is 2.0.1.

If you install the XMP Toolkit’s Linux package, you also get the necessary `Exempi` library on your computer. Until now, the Toolkit has only been available

XMP AND EXIF WITH PYTHON

Free Python libraries for programming Exif applications also can handle XMP, they reside alongside the XMP Toolkit (although the Exif libraries do not support the same file formats). None of these tools is implemented in pure Python; instead, they are all bindings to available C- or C++ libraries.

Under the hood, the XMP Toolkit is a wrapper written with ctypes around `Exempi` [8], an offshoot of Adobe’s official XMP Software Development Kit (current version 2.3.0, which is based on Adobe XMP SDK 4.1.1).

`Pyexiv2` [9] is a binding to the `Exiv2` C++ library [10], implemented with `Boost.Python`, which developers can use to pro-

gram applications for Exif, IPTC, and XMP metadata (current version 0.3.2). Because no one is developing `Pyexiv2`, a switch to `GExiv2` is recommended.

`GExiv2` [11] is a wrapper around the `Exiv2` library for the `GObject` programming environment (current version 0.10.3). The software supports `GObject` introspection, which Python programmers can access via `PyGObject` [12]. To do this, you need the `gir1.2-gexiv2`, `python-gi`, and `python3-gi` packages (e.g., on a Debian system). Then, use the following command:

```
from gi.repository import GExiv2
```

to import the library.

in a few distributions, such as in Debian and its offshoots, where it is within the `python-libxmp` and `python3-libxmp` packages. Alternatively, you can install it from the Python Packet Index [13].

The online documentation currently misses some of its parts [14]; Debian users are better off installing the `python-libxmp-doc` documentation package [15]. Alternatively, programmers can collect the documentation from the GitHub repo or scour the source code directly for the docstrings.

The `libxmp.files.XMPFiles` class controls the handling of files in the Python

XMP Toolkit, and the `libxmp.core.XMPMeta` class offers a range of methods (functions) for manipulating XMP packets in memory. For contact with XMP, the Toolkit defines its own complex data object, with which Python’s usual XML tools cannot cope (although this is not necessary). The example in Listing 2 demonstrates a few simple operations with the Toolkit in an IPython session.

Action Mode

The listing script imports both main classes and opens the requested file, loading the XMP packet via `get_xmp()` onto the `myxmp` memory object. If the file called still does not contain an XMP packet, the Toolkit creates an empty template with `x:xmpmeta` and the basic RDF framework in memory.

The `consts` module offers a range of substitutes for the common namespaces, meaning `consts.XMP_NS_DC`, for instance,




Figure 1: The Python XMP Toolkit helped sort the Hubble Telescope’s numerous images.

TABLE 1: A Selection of XMP Elements and Namespaces

Description	Content	Format
Dublin Core (http://purl.org/dc/elements/1.1/)		
<code>dc:title</code>	Title of document or item	Alternative list with <code>xml:lang</code>
<code>dc:creator</code>	Producer (person or organization)	Ordered list
<code>dc:publisher</code>	Name of the publishing entity	Unordered list
<code>dc:subject</code>	Collection of keywords	Unordered list
<code>dc:language</code>	Language of the item	Unordered list with RFC 306 tags
<code>dc:format</code>	File format of the object	MIME type
<code>dc:identifier</code>	ISBN/ISSN, URN, DOI, and others	Text
XMP Core Elements (http://ns.adobe.com/xap/1.0/)		
<code>xmp:CreateDate</code>	Object’s date of production	Date stamp
<code>xmp:CreatorTool</code>	Tool of production	Text
<code>xmp:MetadataDate</code>	Modification date of the metadata	Date stamp
<code>xmp:ModifyDate</code>	Modification date of the object	Date stamp
<code>xmp:Rating</code>	Rating of the tool	Score from -1 to 5
XMP Rights Management (http://ns.adobe.com/xap/1.0/rights/)		
<code>xmpRights:Marked</code>	Copyright marking	True/false
<code>xmpRights:Owner</code>	Rights holder	Unordered list
<code>xmpRights:UsageTerms</code>	License/use terms	Alternative list with <code>xml:lang</code>
XMP Media Management (http://ns.adobe.com/xap/1.0/mm/)		
<code>xmpMM:DocumentID</code>	Identifier of the object	GUID stamp
<code>xmpMM:InstanceID</code>	Identifier of an object instance	GUID stamp

represents the Dublin Core URI, and `consts.XMP_NS_XMP` represents that for the core XMP elements.

The `get_localized_text()` method returns certain datasets localized with `xml:lang` in alternative lists (e.g., those with `x-default` from `dc:title`). The XMP object can be manipulated in a targeted way with `set_property()`, for instance, by converting the `x-default` language setting to `en`. The `set_localized_text()` changes localized data in alternative lists, and in `dc:title`, it would expand a short German title with the localization `xml:lang=de`.

Next, `get_property()` again requests the dataset of the `xmp:MetadataDate` element this time. This is a date stamp in ISO 8601 format. Developers can create a new date stamp (`now`) with the Python library's `datetime` module [16] and overwrite `xmp:MetadataDate` in the XMP packet with `set_property()`.

Alternatively, the `set_property_datetime()` method deals with date stamps. The `can_put_xmp()` method checks whether the opened file is write protected. If this is not the case, `put_xmp()` writes the file and `close_file()` closes it.

Creative Commons

The non-profit Creative Commons organization distributes a variety of easily applicable free licenses for digital works [17] that artists and authors worldwide use to protect their copyrights. Creative Commons recommends XMP as a metadata format. At the same time, the organization posts guidelines on how artists can best utilize the information of the license they use [18] and embed this in their published files. This creates a win-win situation: The license strengthens the position of the author, and the users establish the license conditions in an uncomplicated and reliable way from the objects themselves.

Creative Commons has defined the XML namespace `http://creativecommons.org/ns#` for XMP, as well as for other purposes; this includes the following four elements `cc:license`, `cc:morePermissions`, `cc:attributionURL`, and `cc:attributionName`. The XMP Toolkit makes this namespace available under the `consts.XMP_NS_CC` key. Fur-

LISTING 2: Python XMP Toolkit Demo

```
In [1]: from libxmp import XMPFiles, XMPMeta, consts

In [2]: xmpfile = XMPFiles(file_path='XMPSpecificationPart1.pdf',
openforupdate=True)

In [3]: myxmp = xmpfile.get_xmp()

In [4]: print myxmp.get_localized_text(consts.XMP_NS_DC, 'title', '', 'x-default')
XMP Specification Part 1: Data Model, Serialization, and Core Properties

In [5]: myxmp.set_property(consts.XMP_NS_DC, 'title[1]?xml:lang', 'en')

In [6]: myxmp.set_localized_text(consts.XMP_NS_DC, 'title', '', 'de',
'XMP-Spezifikation, Teil 1')

In [7]: print myxmp.get_property(consts.XMP_NS_XMP, 'MetadataDate')
2012-03-21T08:57:17-07:00

In [8]: import datetime
now = datetime.datetime.now().strftime("%Y-%m-%dT%H:%M:%S+01:00")

In [9]: myxmp.set_property(consts.XMP_NS_XMP, 'MetadataDate', now)

In [10]: print myxmp.get_property(consts.XMP_NS_XMP, 'MetadataDate')
2016-02-21T19:33:26+01:00

In [11]: xmpfile.can_put_xmp(myxmp)
True
```


HACKATHON

**REGISTER
NOW!**

WORKSHOPS
**LATEST KNOWLEDGE
NEWEST DEVELOPMENTS
PIONEERING PERSPECTIVES**
netways.de/osmc

presented by


NETWAYS®

thermore, the organization recommends the use of the two elements `xmpRights:Marked` and `xmpRights:UsageTerms` from the XMP rights management schema (Table 1).

A copy of the license notice should also (according to their advice) end up in `dc:rights`. Listing 3 shows a Python script that writes the metadata recommended by Creative Commons into a

JPG file (not yet containing an XMP packet) with XMP Toolkit for the Attribution-ShareAlike 4.0 license. The result of these efforts is presented in Listing 4.

The `xmpRights:WebStatement` text element transports a callable URL, by which details regarding the authorship and conditions of use of a file can be verified. The Creative Commons creators

suggest installing a reference under the address given here that also contains the hash value of the published file [19]. In this case, it can be manually or automatically compared to raise the alarm in case of forgeries.

Conclusion

XMP involves a sophisticated and professional metadata system that compiles

LISTING 3: Generating Creative Commons-Compliant Metadata

```
01 from libxmp import XMPFiles, XMPMeta, consts
02 image = XMPFiles(file_path='1001.jpg', open_forupdate=True)
03 xmp = image.get_xmp()
04
05 xmp.register_namespace(consts.XMP_NS_CC, 'cc')
06 xmp.set_property(consts.XMP_NS_CC, 'license', 'http://creativecommons.org/licenses/by-sa/4.0/')
07 xmp.set_property(consts.XMP_NS_CC, 'attributionName', 'Max Mustermann')
08 xmp.set_property(consts.XMP_NS_CC, 'attributionURL', 'http://example.com/bilder/1001.jpg')
09
10 xmp.register_namespace(consts.XMP_NS_XMP_Rights, 'xmpRights')
11 xmp.set_property_bool(consts.XMP_NS_XMP_Rights, 'Marked', True)
12 usage = 'This work is licensed to the public under the Creative Commons Attribution-ShareAlike 4.0 International License'
13 xmp.set_property(consts.XMP_NS_XMP_Rights, 'UsageTerms', usage)
14
15 xmp.register_namespace(consts.XMP_NS_DC, 'dc')
16 xmp.set_localized_text(consts.XMP_NS_DC, 'rights', '', 'x-default', usage)
17
18 image.put_xmp(xmp)
19 image.close_file()
```

LISTING 4: Created XMP Package

```
<x:xmpmeta xmlns:x="adobe:ns:meta/" x:xmptk="Exempi + XMP Core 5.1.2">
  <rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
    <rdf:Description rdf:about="" xmlns:cc="http://creativecommons.org/ns#">
      <cc:license>http://creativecommons.org/licenses/by-sa/4.0/</cc:license>
      <cc:attributionName>Max Mustermann</cc:attributionName>
      <cc:attributionURL>http://example.com/bilder/1001.jpg</cc:attributionURL>
    </rdf:Description>
    <rdf:Description rdf:about="" xmlns:xmpRights="http://ns.adobe.com/xap/1.0/rights/">
      <xmpRights:Marked>True</xmpRights:Marked>
      <xmpRights:UsageTerms>This work is licensed to the public under the Creative Commons
Attribution-ShareAlike 4.0 International License</xmpRights:UsageTerms>
    </rdf:Description>
    <rdf:Description rdf:about="" xmlns:dc="http://purl.org/dc/elements/1.1/">
      <dc:rights>
        <rdf:Alt>
          <rdf:li xml:lang="x-default">This work is licensed to the public under the Creative Commons
Attribution-ShareAlike 4.0 International License</rdf:li>
        </rdf:Alt>
      </dc:rights>
    </rdf:Description>
  </rdf:RDF>
</x:xmpmeta>
```

schemas for different objects and applications. The comprehensive set of descriptive components can even be expanded with elements not provided by the specification.

XMP supports a wide range of file formats in which to embed metadata packets; it is only on exceptional occasions that no solution exists. Nothing else is comparable, especially when it comes to PDF documents, although this is not surprising considering the standard's author (Adobe).

End users can request XMP metadata with ease from the command line, and programmers can develop special XMP applications with a range of C/C++ libraries. Python bindings are available for the existing XMP and Exif libraries, which can also handle XMP.

Exiftool (in the *libimage-exiftool-perl* Debian package or *perl-image-exiftool* package for openSUSE and others) and the command-line tool from the Exiv2 library (in the Debian *exiv2* package) are also useful for editing XMP data at the command line. ■■■

INFO

- [1] Exif specification: http://www.cipa.jp/std/documents/e/DC-008-2012_E.pdf
- [2] XMP specification: <http://www.adobe.com/devnet/xmp.html>
- [3] XMP ISO standard press release: http://www.iso.org/iso/home/news_index/news_archive/news.htm?refid=Ref1525
- [4] MWG guidelines: <http://www.metadataworkinggroup.com/specs>
- [5] RDF 1.1 XML syntax: <https://www.w3.org/TR/rdf-syntax-grammar>
- [6] DCMI metadata terms: <http://dublincore.org/documents/dcmi-terms/>
- [7] Python XMP Toolkit in GitHub: <https://github.com/python-xmp-toolkit/python-xmp-toolkit/>
- [8] Exempi: <https://libopenraw.freedesktop.org/wiki/Exempi/>
- [9] Pyexiv2: <http://tilloy.net/dev/pyexiv2/>
- [10] Exiv2: <http://exiv2.org>
- [11] GExiv2: <https://wiki.gnome.org/Projects/gexiv2>
- [12] PyGObject: <https://wiki.gnome.org/Projects/PyGObject>
- [13] Python XMP Toolkit in PyPI: <https://pypi.python.org/pypi/python-xmp-toolkit/>
- [14] XMP Toolkit documentation: <http://python-xmp-toolkit.readthedocs.org/en/latest/>
- [15] Debian packages: <https://packages.qa.debian.org/p/python-xmp-toolkit.html>
- [16] Python datetime library: <https://docs.python.org/2/library/datetime.html>
- [17] Creative Commons licensing types: <https://creativecommons.org/share-your-work/licensing-types-examples/>
- [18] Creative Commons XMP properties: <https://wiki.creativecommons.org/wiki/XMP>
- [19] WebStatement: https://wiki.creativecommons.org/wiki/Web_Statement

Shop the Shop

shop.linuxnewmedia.com

Missed an issue?

You're in luck.

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

shop.linuxnewmedia.com

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



The sys admin's daily grind: testssl.sh

A Key Role

Back in the day, integrating and function testing an SSL certificate was an easy thing to do. Now, you can use a shell script that talks plain English, despite the Babylonian confusion of key protocols and ciphers. *By Charly Kühnast*

Web servers with SSL certificates will likely be the rule. The administrator's responsibility for keeping track of which SSL and TLS settings are up to date on which servers tends to grow proportionally. I have found a trusty helper for this task, `testssl.sh` [1], a shell script with many capabilities. For example, typing

LISTING 1: `testssl -U https://<server>`

```
01 Service detected:          HTTP
02 Testing vulnerabilities
03
04 Heartbleed (CVE-2014-0160)      not vulnerable (OK) (timed out)
05 CCS (CVE-2014-0224)           not vulnerable (OK)
06 Secure Renegotiation (CVE-2009-3555) not vulnerable (OK)
07 Secure Client-Initiated Renegotiation not vulnerable (OK)
08 CRIME, TLS (CVE-2012-4929)    not vulnerable (OK)
09 [...]
```

LISTING 2: `testssl -P <address>`

```
01 Testing server preferences
02
03 Has server cipher order?    yes (OK)
04 Negotiated protocol        TLSv1.2
05 Negotiated cipher          ECDHE-RSA-AES256-GCM-SHA384, 256 bit ECDH
06 Cipher order
07 [A long list of ciphers offered]
```

```
testssl -U https://<server>
```

runs scan tests for popular exploits like Heartbleed, Poodle, and Crime (see output in Listing 1), and I can trigger each test option individually. To discover whether a server is vulnerable for Logjam, I would just type:

```
testssl.sh --logjam <address>
```

To test whether the server offers its ciphers in a defined order (from strong to weak), I type:

```
testssl -P <address>
```


Listing 2 shows the strongest encryption method the server was able to negotiate. The script additionally helps me discover whether the server supports HTTP/2 with:

```
testssl.sh -Y <address>
```

It might then respond with *Service*

CHARLY KÜHNAST

Charly Kühnast is a Unix operating system administrator at the Data Center in Moers, Germany. His tasks include firewall and DMZ security and availability. He divides his leisure time into hot, wet, and eastern sectors, where he enjoys cooking, freshwater aquariums, and learning Japanese, respectively.



detected: HTTP HTTP2/ALPN http/1.1 (offered).

The tests are not restricted to HTTPS. I can test a mail server like this:

```
testssl.sh --starttls smtp <address>
```

which gives me an equally good quality of analysis as for the web protocol – and it is very exhaustive in this case. If you don't specify any parameters testssl.sh runs through the whole arsenal of test routines (Figure 1).

Aha

Finally, a tip from the developer of testssl.sh: If you additionally install the aha [2] tool, you can easily convert the console output to HTML:

```
testssl.sh <address> | aha >Z  
/<Path>/servertest.html
```

After wrapping this up in a small shell script, I can then use cron to check the TLS and SSL status of my servers on a weekly basis. ■■■

INFO

[1] Testssl.sh: <https://github.com/drwetter/testssl.sh>

[2] Aha: <https://github.com/theZiz/aha>

```

Testing protocols (via sockets except TLS 1.2, SPDY+HTTP2)
SSLv2      not offered (OK)
SSLv3      not offered (OK)
TLS 1      offered
TLS 1.1    offered
TLS 1.2    offered (OK)
SPDY/NPN   not offered
HTTP2/ALPN http/1.1 (offered)

Testing ~standard cipher lists
Null Ciphers      not offered (OK)
Anonymous NULL Ciphers not offered (OK)
Anonymous DH Ciphers not offered (OK)
40 Bit encryption not offered (OK)
56 Bit encryption not offered (OK)
Export Ciphers (general) not offered (OK)
Low (<=64 Bit)    not offered (OK)
DES Ciphers       not offered (OK)
Medium grade encryption not offered (OK)
Triple DES Ciphers offered
High grade encryption offered (OK)

Testing robust (perfect) forward secrecy, (P)FS -- omitting Null Authentication/Encry
PFS is offered (OK)  ECDHE-RSA-AES256-GCM-SHA384 DHE-RSA-AES256-GCM-SHA384 DHE-RSA-AES

Testing server preferences
Has server cipher order?  yes (OK)
Negotiated protocol      TLSv1.2
Negotiated cipher        ECDHE-RSA-AES256-GCM-SHA384, 256 bit ECDH
Cipher order
TLSv1:  ECDHE-RSA-AES256-SHA ECDHE-RSA-AES128-SHA DHE-RSA-AES256-SHA DHE-RSA-AES
TLSv1.1: ECDHE-RSA-AES256-SHA ECDHE-RSA-AES128-SHA DHE-RSA-AES256-SHA DHE-RSA-AES
TLSv1.2: ECDHE-RSA-AES256-GCM-SHA384 ECDHE-RSA-AES128-GCM-SHA256 ECDHE-RSA-AES256

```

Figure 1: Testssl.sh even checks mail server certificates.



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

ADMIN
Network & Security

www.admin-magazine.com/newsletter

Solving optimal stopping problems with statistical algorithms

Enough Is Enough

When is it statistically favorable to conclude, say, the search for a suitable employee with a chosen one? Solid algorithms point the way to success. *By Mike Schilli*



The saying goes that you should quit when you're ahead. Making a decision on when to conclude a selection process is often not that easy in real life. In mathematical statistics, a puzzle known as the "secretary problem" [1] sums this up nicely. A number of candidates are applying for a secretary position at a company. In the selection process, the employer must decide after each interview, whether to accept the candidate or reject them and hope for a more suitable applicant later on. The decision is final, the employer is not allowed to invite back any rejected applicants.

An employer who is thinking logically should carefully check the first few candidates and not simply take the first one.

Later, when the line

starts to come to an end, the employer will possibly take on a more-or-less suitable candidate for fear of being left with only unsuitable candidates, and – not having any other options – grudgingly just bite the bullet, and grab whoever is left.

Ideal: 37 Percent

How should the employer proceed to mathematically ensure the best possible odds of finding an above-average candidate? How long should the try-before-buy phase last, in which the employer explores the unknown capabilities of candidates before immediately snapping up the one that is better than all the previous ones and hopefully scores favorably against future candidates further down the line? Mathematicians have been tackling stop problems [2] for thousands of years.

As the recently published book *Algorithms to Live By* [3] entertainingly explains, the answer is 37 percent, or more precisely $1/e \approx 37\%$. Given 100 candidates, an employer is well advised to check out the first 37 without any intention to employ anyone, just to gain some idea of the talents available on the labor market. As of candidate number 38, smart bosses then immediately sign a

MIKE SCHILLI

Mike Schilli works as a software engineer in the San Francisco Bay Area. He can be contacted at mschilli@perlmeister.com. Mike's homepage can be found at <http://perlmeister.com>.




Figure 1: The Gaussian bell curve of normally distributed events (from Wikipedia; public domain).

contract if one candidate is obviously better than all the previous ones.

Love-Crazed Stargazers

This task is also known as the wedding problem. The story of the famous stargazer Johannes Kepler, who was desperately looking for a new wife after the death of his first wife in 1611, is legendary. With meticulous mathematical skills, he launched himself into a lengthy process of research, looking at no less than 11 candidates. Completely exhausted and remorseful, he then returned to candidate number five, proposed to her, and to his great joy luckily received a glowing “Yes!” as an answer to his question.

Simulating Reality

To work your way through the secretary problem with a Perl program, the script needs to make some assumptions about the distribution of the supply in the labor market. The controversial “20-70-10” rule of American tycoon Jack Welch [4] dictates that the 20 percent top performers in a company return exceed expectations, 70 percent work quite well, and 10 percent are unproductive. For the experiment, we let the labor market also follow this scheme; the available

Because the candidates arrive in random order, the script needs a random generator that outputs a normally distributed pool of talents. The `random_normal` function of the `Math::Random` CPAN module already has such a generator; in Listing 1 [5], it uses the parameters `100_000`, `10`, `2` to return 100,000 random values that fluctuate around the mean value of 10 with a standard deviation of 2.

To check whether the distribution also roughly meets the requirements visually, Listing 1 uses the `get_histogram()` function from the `Statistics::Histogram` CPAN module to print a histogram of distributed values as ASCII art on the command line. The bell shape is indeed visible (Figure 2), and although the maximum of the random data is at around 10, the frequency drops off towards both

LISTING 1: histo-normal

```
1 #!/usr/local/bin/perl -w
2 use strict;
3 use Math::Random qw( random_normal );
4 use Statistics::Histogram;
5
6 my @data = random_normal( 100_000, 10, 2 );
7
8 print get_histogram( \@data, 20 );
```

talents thus roughly lie on the bell curve of a Gaussian normal distribution (Figure 1).

sides, and there are hardly any values less than 4.5 or greater than 15.5.

The script for simulating the secretary interviews (Listing 2) takes on the randomly arriving candidates one by one. Following the 37 percent rule, it breaks down the `@data` array into two parts. The first part, stored in `@look`, contains the first 37 percent of the random numbers, which are merely for analysis of the labor market without any intention to employ; the second part (`@act`) contains the other 63 percent; the employer will immediately make a decision here if one of the candidates is better than all previous encountered ones. Line 41 stops the script when this occurs.

If no better candidate shows up by the end of the array, the boss has lost out and has to employ the last candidate as a final option, which is what line 46 does. If the script is called with the `--verbose` option, line 9 of Listing 2 switches `Log4perl`'s logging function to debug mode to produce diagnostic messages on how the hiring process is proceeding.




Figure 2: Histogram of events following a bell curve.

LISTING 2: secretary

```

01 #!/usr/local/bin/perl -w
02 use strict;
03 use Math::Random qw( random_normal );
04 use Log::Log4perl qw(:easy);
05 use Getopt::Long;
06
07 GetOptions( \my %opts, "verbose" );
08
09 Log::Log4perl->easy_init(
10     $opts{ verbose } ? $DEBUG :
11     $INFO );
12
13 my $nof_candidates = 10;
14 my @data = random_normal(
15     $nof_candidates, 10, 2 );
16
17 my $ratio = 37.0;
18 my $cutoff = $nof_candidates *
19     $ratio / 100;
20
21 my @look = splice @data, 0, $cutoff;
22 my @act = @data;
23
24 my $best = 0;
25
26 for my $talent ( @look ) {
27     DEBUG "Testing: $talent";
28
29     if( $talent > $best ) {
30         $best = $talent if $talent > $best;
31         DEBUG "New best: $best";
32     }
33 }
34
35 for my $talent ( @act ) {
36     DEBUG "Acting on: $talent";
37
38     if( $talent > $best ) {
39         DEBUG "Taking: $talent";
40         print "$talent\n";
41         exit 0;
42     }
43 }
44
45 DEBUG "Forced to take $act[ -1 ]";
46 print "$act[ -1 ]\n";

```

\$ secretary --verbose

```

2016/06/12 12:03:45 Testing: 11.0533785516305
2016/06/12 12:03:45 New best: 11.0533785516305
2016/06/12 12:03:45 Testing: 8.87076704287439
2016/06/12 12:03:45 Testing: 11.5266307422267
2016/06/12 12:03:45 New best: 11.5266307422267
2016/06/12 12:03:45 Acting on: 9.65402858507352
2016/06/12 12:03:45 Acting on: 9.60358693687708
2016/06/12 12:03:45 Acting on: 10.3845934764761
2016/06/12 12:03:45 Acting on: 8.27611196394389
2016/06/12 12:03:45 Acting on: 11.6820667491877
2016/06/12 12:03:45 Taking: 11.6820667491877
11.6820667491877

```

Figure 3: The selection process for secretaries, first run.

\$ secretary --verbose

```

2016/06/12 12:03:43 Testing: 12.8466752939257
2016/06/12 12:03:43 New best: 12.8466752939257
2016/06/12 12:03:43 Testing: 8.85942045102728
2016/06/12 12:03:43 Testing: 8.76202779689708
2016/06/12 12:03:43 Acting on: 11.9118154283905
2016/06/12 12:03:43 Acting on: 9.15624316640158
2016/06/12 12:03:43 Acting on: 9.01033623524895
2016/06/12 12:03:43 Acting on: 7.09819801235453
2016/06/12 12:03:43 Acting on: 10.1725661008826
2016/06/12 12:03:43 Acting on: 10.2328990079837
2016/06/12 12:03:43 Acting on: 9.36355947693828
2016/06/12 12:03:43 Forced to take 9.36355947693828
9.36355947693828

```

Figure 4: This procedure does not always turn out well for the boss.

The script then outputs a number of messages to standard output at run time. Figure 3 shows that the bar for the candidates during the screening stage was at 11.05, and then at 11.52. In employment mode, a number of candidates with suitability values below the previous record of 11.52 then follow, before one with a value of 11.68 gets the job.

Although the procedure has been proven to be optimal, it still often goes wrong. Figure 4 shows a run in which a candidate with an extremely high suitability value (12.84) interviews during the test phase, a value that puts all the others later in the employment round to shame. Bad luck, the boss has to take the last candidate with a rating of 9.36.

How often do these embarrassing and significant suboptimal outcomes occur in the algorithm? To measure this, Listing 3 calls the secretary script 100 times and measures how often a candidate is given preference although they are below average. The Stats::Histogram CPAN module shows the distribution.

Figure 5 shows that the procedure finds a useful candidate in about 80 percent of the cases, including some superstars with ratings of 12 and more. However, it failed in 20 percent of cases where the boss had to take on the last candidate, whether they were a star or a clown. All told, the

LISTING 3: secretarystats

```
01 #!/usr/local/bin/perl -w
02 use strict;
03 use Log::Log4perl qw(:easy);
04 use Statistics::Histogram;
05 use Getopt::Long;
06
07 GetOptions( \my %opts, "verbose" );
08
09 Log::Log4perl->easy_init(
10     $opts{ verbose } ? $DEBUG :
11     $INFO );
12
13 my @data = ();
14
15 my $clowns = 0;
16 my $stars = 0;
17
18 for ( 1..100 ) {
19     my $talent = `./secretary`;
20     chomp $talent;
21
22     DEBUG "Result: $talent";
23
24     if( $talent >= 10 ) {
25         $stars++;
26     } else {
27         $clowns++;
28     }
29
30     push @data, $talent;
31 }
32
33 INFO get_histogram( \@data, 20 );
34
35 INFO "Stars: $stars Clowns: $clowns";
```

```
$ ./secretary-stats
2016/06/12 11:35:31 Count: 100
Range: 5.313 - 15.188; Mean: 11.456; Median: 11.824; Stddev: 1.919
Percentiles: 90th: 13.574; 95th: 14.281; 99th: 14.685
 5.313 - 5.617: 1 ###
 5.617 - 6.621: 1 ###
 6.621 - 7.374: 1 ###
 7.374 - 8.201: 2 #####
 8.201 - 8.644: 4 #####
 8.644 - 9.109: 3 #####
 9.109 - 9.597: 5 #####
 9.597 - 10.107: 8 #####
10.107 - 10.643: 7 #####
10.643 - 11.204: 8 #####
11.204 - 11.793: 9 #####
11.793 - 12.409: 13 #####
12.409 - 13.056: 21 #####
13.056 - 13.733: 7 #####
13.733 - 14.444: 5 #####
14.444 - 15.188: 5 #####
2016/06/12 11:35:31 Stars: 77 Clowns: 23
```

Figure 5: In 20 percent of cases, the boss has to take the last candidate.

strategy is optimal; check out [1] for the mathematical proof.

Las Vegas and Everyday Life

Just as there are correct and incorrect ways of winning at certain games in Las Vegas – and although I am amazed at how few people actually take them into consideration to at least statistically give the casino as little chance as possible to win – it is important in everyday life to keep the upper hand in randomly driven processes and always consider the most likely possible outcome of such a situation.

However, what is demonstrably the optimal strategy, if you do the math, is sometimes of little practical value for stop problems. In a casino that pays back triple your bet after tossing a coin with a winning percentage of 50 percent, the theoretically correct approach would be to keep on playing and never quit. But, as you can imagine, the players would go home bankrupt every time if they followed this strategy. ■■■

INFO

- [1] Secretary problem: https://en.wikipedia.org/wiki/Secretary_problem
- [2] Optimal stopping: https://en.wikipedia.org/wiki/Optimal_stopping
- [3] Christian, Brian, and Tom Griffiths. *Algorithms to Live By: The Computer Science of Human Decisions*. Henry Holt and Co. April 2016: <http://algorithmstoliveby.com>
- [4] Vitality curve: https://en.wikipedia.org/wiki/Vitality_curve
- [5] Listings for this article: <ftp://ftp.linux-magazine.com/pub/listings/magazine/190>

GOT CLUSTER?

ADMIN HPC NEW DASH

- HPC Up Close**
 - Monitoring HPC Systems
 - Watson in the Cloud
 - SC14 Call for Workshop Proposals
 - Data Visualization Tools
- Further Reading**
 - Fixing Disks with Parted Magic
 - Exploring the Xeon Phi
 - Tuning I/O Patterns in Python
 - Blue Waters Open for Business
- Partner Link**
Intelligent High Availability DAS for Windows
- Raspberry Pi Geek Offer**
Subscribe to Raspberry Pi Geek and receive either Issue 1 or Raspberry Pi Geek for FREE!
Use code HOLSFPG2013 at checkout. Click here for full details. (Offer ends 2014.)
- Most Read**
 - Tuning I/O Patterns in Fortran 90**
In the second article of this three-part series, we look at simple write examples in Fortran 90 and track the output with strace to see how...
 - Build a Network Attached Storage System with FreeNAS**
We provide an overview of FreeNAS and ZFS, its main filesystem, then show you how to set up and maintain a FreeNAS installation.
 - XSEDE Releases Toolset**
The new tools will improve compatibility of local campus clusters with the XSEDE supercomputer network.
- Highlights**
 - Monitoring HPC Systems: What Should You Monitor?**
In rapidly growing HPC installations, you need to understand what is happening within the system to make improvements or simply to justify the...
 - Watson in the Cloud**
IBM launches three new services for its well-trained, Jeopardy-busting behemoth.
 - SC14 Calls for Workshop Proposals**
Entries should provide "interaction and in-depth discussions of stimulating topics of interest to the HPC community."
 - Data Visualization Tools**
Logfiles contain extremely useful traces of everyday events – if only they weren't so boring. Fortunately, special tools help you generate and...

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

hpc.admin-magazine.com/Newsletter



Open365 puts LibreOffice, Kontact, and Jitsi in the cloud

Open All Year

Office suites, email clients, and video chats often run as a “software as a service” on the network. Open365 is pitting web-based LibreOffice together with Kontact and Jitsi against the top dogs Google Drive and Office 365.

By Christoph Langner

Many people working on screens switch between multiple devices rather than using a single computer – typically PCs, laptops, smartphones, and tablets. For users, this means that they don’t just need to synchronize the data between the computers, they also need compatible software for all these devices. Not least because “Software as a Service” (SaaS) is becoming increasingly important: Instead of installing software locally, you can use cloud-based services that can be accessed from all devices no matter what platform is used.

As you know the big software corporations have a foot in the door here already: Microsoft has offered its own web-based Office365 since mid-2011 in the form of a fee-based subscription model. Google also clawed its way into the market 10 years ago with Google Docs, Sheets, and Slides. In typical Google-style, the use of these services remains free. So, the question is: Google or Microsoft; if you use SaaS, you’ll be giving up control of your data.

A Whole Load of Cloud

Another alternative is now entering the scene in the form of Open365 [1]: The service combines the well-known open source tools LibreOffice as an office package, Seafile [2] as a synchronization tool, KDE Kontact [3] for managing addresses and emails, and Jitsi [4] as an instant messenger on an interface.

The service is currently (at the time of writing) in the beta phase, but registration is open to any interested parties. In addition to the web interface, Open365 provides a desktop client and an Android app [5] that, similarly to Dropbox, is used for synchronizing data. The developers are currently still working on the iOS version of the mobile app.

The Seafile web interface (Figure 1), which slightly expands Open365, serves as the central component of Open365. In this interface, you can organize your data in libraries (*Libraries*) and encrypt them as required. Altogether 20GB of disk space is available free of charge. (See the “Installation” box for more information.)

Similarly to Dropbox and unlike Google Drive, where Linux support is still missing, Open365 also provides a PC program for Linux, Mac OS X, and Windows that synchronizes data between PC and online storage in the background (Figure 2).

LibreOffice on the Web

The Seafile front end of Open365 provides the option for uploading data and synchronizing these data between different computers via the client. Data or even whole libraries can also be shared with other people – this includes the option to add these files to a library. There is also a personal wiki, an editor for text files along with Markdown formatting, an image viewer with a gallery view, and

an HTML5-based player for various video formats in the front end. However, the most important component is probably the web-based LibreOffice.

Open365 includes the word processor Writer (Figure 3), the spreadsheet Calc (Figure 4), and the presentation tool Impress from LibreOffice 5.1.0.3. The programs provide the same functionality as a locally installed LibreOffice. Technically, the implementation of Linux programs is based on the web desktop from EyeOS [6], whose first versions were published under an open source license [7].

Open365 provided a set of sample files in the *My Library* folder for direct testing, including current Microsoft Office documents. You can save your own documents in any library and then, for example, back up the files using the Open365 client on your PC, where you can continue editing the file without any problems. The client then synchronizes local changes with the online store and the data therefore always remains up to date.

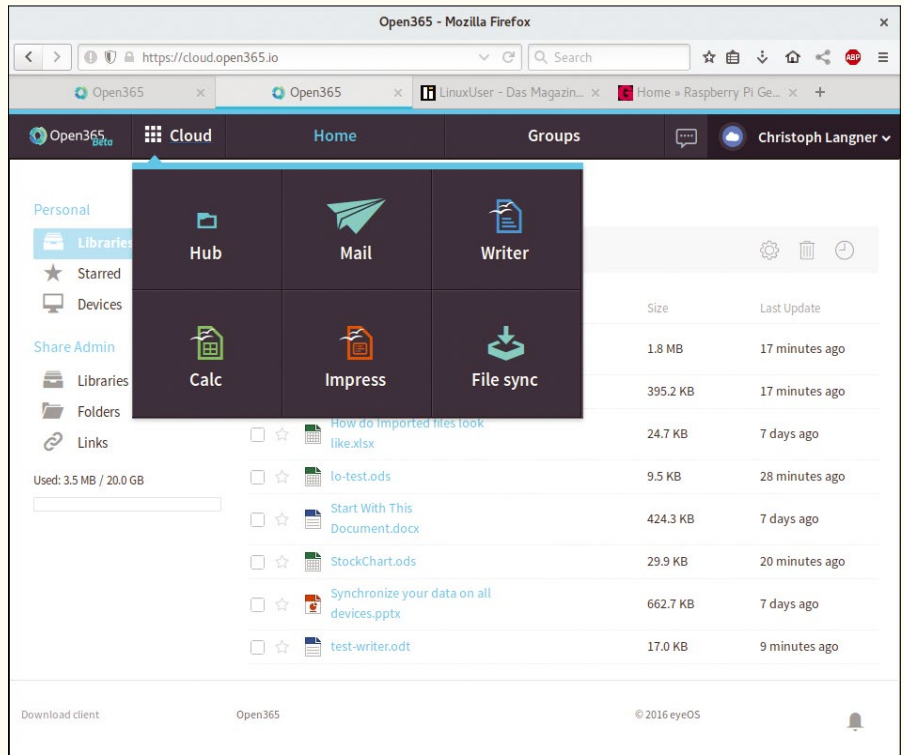


Figure 1: The Open365 web interface comes from Seafile and has been expanded to include LibreOffice and a voice and video chat function.

Voice and Video Chats

As well as the LibreOffice suite with Kontact 4.13.0.15 from KDE 4.13.3, Open365 also includes a comprehensive personal information manager (PIM) with email, address book, and calendar functions. Your account name in the style of “fredbloggs@open365.io” is automatically used as the email address. For this to appear properly for the recipient later, click *Mail* in the *Cloud* drop-down, open the email function in the sidebar under *Mail*, and then correct the name from the menu under *Settings | Configure KMail*.

Within Open365, the service also provides the option to send short messages to other users. You can access this function from the hub via the chat icon next to your name. You can search through the complete set of users in the sidebar that opens *Contacts* – Open365 lists the name here that you specified when creating your account. There are no options for protecting privacy so far – such as hiding your own name in the search.

You can start a call in the chat window by clicking *Start Video Conference*. Internally, the service uses the open source video chat solution Jitsi and the protocol WebRTC [8]. Thus, you don't need special drivers or apps – all modern web browsers should work.

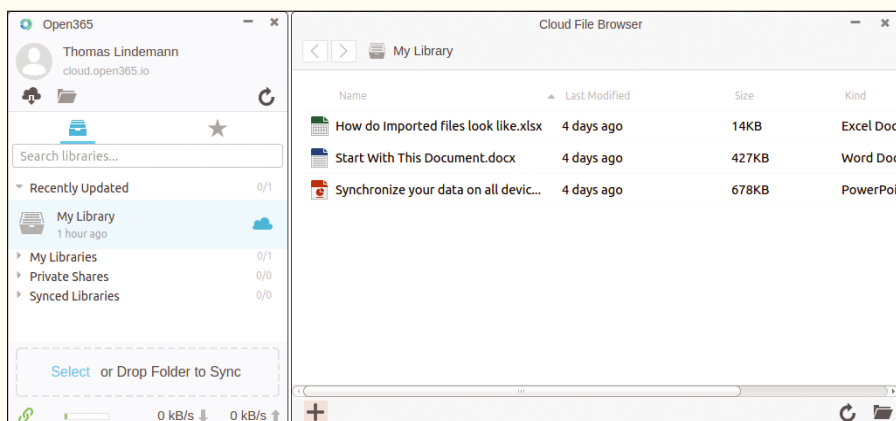


Figure 2: The client provided by Open365 for desktop PCs for synchronizing data also originates from the Seafile project.

INSTALLATION

Like the service's web interface, the Open365 client is based on Seafile. You can install the program via the package manager in the form of a DEB file. Open365 professes to support Ubuntu from version 14.04, but you can also install the package on Debian 8 "jessie." Both Ubuntu 16.04 and Debian Unstable are still missing the Qt4 libraries requested by the package as a dependency. This means you have to wait for an update by the Open365 developers. Various virus scanners kick in when running the set up in the client's Windows version – the reason for this wasn't clear by time of publication [10].

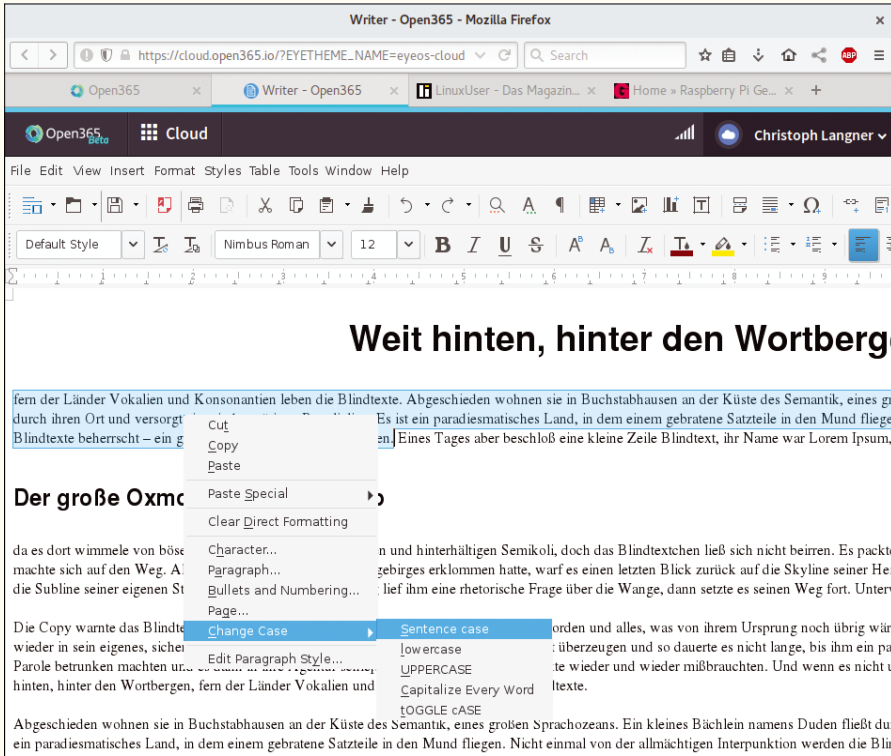


Figure 3: Open365 includes Writer, Calc, and Impress.

However, Chrome proved to be a bit awkward in the test compared to Firefox. The Google browser asked for permission to access the computer's webcam when starting a chat but then didn't show either the camera image or the operating elements of the video chat.

Conclusions

On the homepage, Open365 states: "You can download it and deploy to your own servers or use it online for free." This statement, however, is not yet true. Although Open365 bears the word "Open" in its name and uses open source foundations such as LibreOffice, Seafile, and Jitsi, the service's source code is still missing. Without that, it's pretty tricky to host an instance of the service on your own hardware. So, you'll be making your data available with Open365 and are therefore reliant on the provider's goodwill. However, an admin in the forum writes that the provider is working intensively on the release of the source code [9].

Open365 isn't bad in practice. Regardless of the beta status, all the service's components are up and running. Anyone who likes working with LibreOffice and often uses a computer where the Office package is missing will appreciate a readily available Office suite and 20GB of free cloud storage on top in the form of Open365. However, the missing privacy settings are uncomfortably conspicuous: As a user, you should be able to decide whether you want to publish your presence and the plain names of all users of the service. ■■■

INFO

- [1] Open365: <https://open365.io>
- [2] Seafile: <https://www.seafile.com>
- [3] Kontakt: <https://userbase.kde.org/Kontakt/en>
- [4] Jitsi: <https://jitsi.org>
- [5] Open365 for Android: <https://play.google.com/store/apps/details?id=io.open365.cloud>
- [6] EyeOS: <http://www.eyeos.com>
- [7] Open source version of EyeOS: <https://sourceforge.net/projects/eyeos>
- [8] WebRTC: <https://webrtc.org>
- [9] Open365 on premises: <https://support.open365.io/index.php/en/forum/englishsuport/13-open365-on-premises#26>
- [10] Virus scanners warn against Windows version: <https://support.open365.io/index.php/es/forum/englishsuport/41-virus>

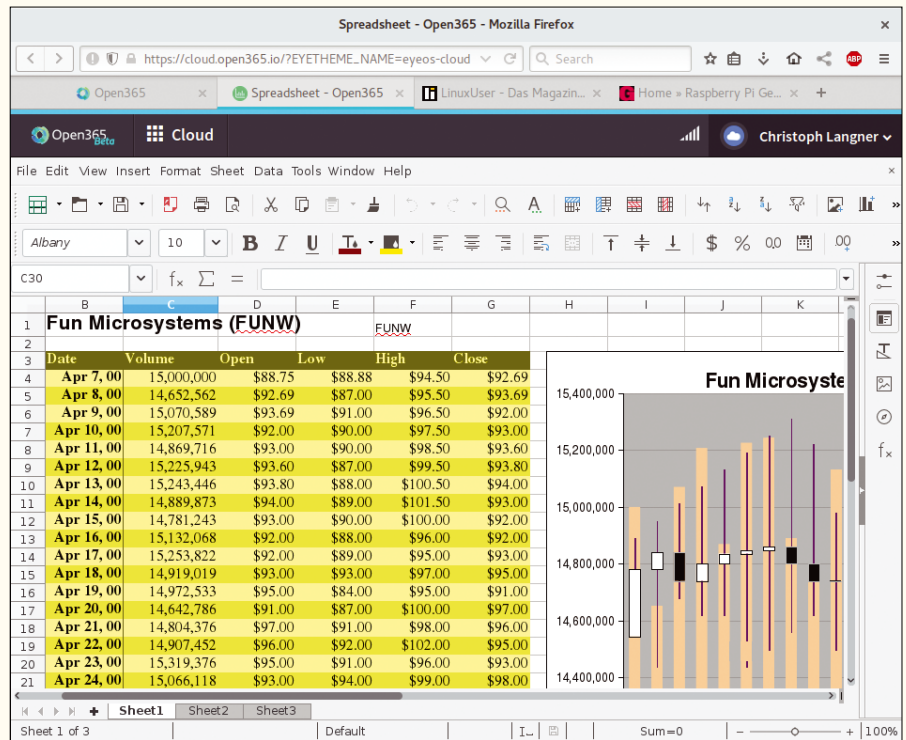



Figure 4: Using Open365, you can edit LibreOffice documents and those in Microsoft's Office Open XML format.

Raspberry Pi

HANDBOOK



In case you missed it last time...

**You ordered your Raspberry Pi...
You got it to boot...what now?**

The Raspberry Pi Handbook takes you through an inspiring collection of projects. Put your Pi to work as a:

- media center
- web server
- IR remote
- hardware controller
- and much more!

Discover Raspberry Pi's special tools for teaching kids about programming and electronics, explore Wolfram Mathematica, and find out how to integrate your Rasp Pi system with LEGO Mindstorms.

**THE ONLY RASPBERRY PI REFERENCE
YOU'LL EVER NEED!**



ORDER ONLINE:

shop.linuxnewmedia.com/rpi



The race for a universal package manager

Package Deal

We look at four contenders that want to become the Linux universal package manager, although the reality is still on the horizon. *By Bruce Byfield*

BRUCE BYFIELD

Bruce Byfield is a computer journalist and a freelance writer and editor specializing in free and open source software. Bruce's most recent book, *Designing with LibreOffice*, was released under a Creative Commons License in March 2016. You can buy or download his book at <http://designingwithlibreoffice.com/download-buy/>. 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>.

So far, 2016 is likely to be remembered as the year in which universal package managers were debated. With the announcement of Ubuntu's Snappy packages, followed quickly by Fedora's and Red Hat's release of its Flatpak format, the race is on for a method of software installation that will work on any distribution and make .deb, .rpm, and the rest footnotes in Linux history.

Advocates of a universal package manager claim that it will make software installation more efficient by eliminating dependency problems. Additionally, many candidates include extra security and extra features, such as package installation by ordinary users for their own use. However, whether any improvements will result is questionable. Debian developers in particular insist the way to

eliminate software installation problems lies in enforcing a strict package structure, as described in the Debian Policy Manual [1].

In fact, Josh Triplett, a long-time Debian contributor, went so far as to suggest that Debian's reputation is a consequence of the Debian Policy, and not the .deb format. According to Triplett, "Debian without the .deb format would still be Debian; Debian without [the] Debian Policy would just be Sourceforge or Rpmfind," and nothing more than a random collection of software [2].

This claim is borne out by the specialized installers available for some programming languages, which configure aspects of the language, but not external dependencies. A new format would not eliminate such problems, but more detailed package guidelines might.

Besides, as many observers point out, the attempt to establish a universal standard in free software tends simply to add to the number of choices. Perhaps Flatpak and Snappy will simply replace DEBs and RPMs as the two main rivals in package management.

Still, whether a universal package manager is needed or not, the fact that two candidates are backed by corporate interests means that attempts to establish one as the Linux standard are likely to continue. Here are the four candidates most often mentioned:

AppImage

Under various names, including klik and PortableLinuxApps, AppImage has existed since 2004 [3]. No application is needed to manage AppImage packages, and each packaged application and its dependencies are stored in a single file. Running an AppImage package requires only changing permissions to make it executable, then deciding whether to add a desktop launch the first time you click the package (Figure 1).

Historically minded people mention AppImage as a possible option for a universal package manager, but the possibility seems remote, even though a number of applications, including Audacity, Blender, and Krita offer AppImages [4].

For one thing, the single static file takes far more space than a standard .deb or .rpm file. For example, the sample on the project's home page for SubSur-

face, a diving diary favored by Linus Torvalds, requires seven and a half times the memory of the same `.deb` package. This difference might not matter on a modern system but probably makes AppImage unsuitable for an older machine.

More importantly, AppImage is not designed with security in mind. Even the project's homepage suggests that AppImage is best suited for installation from upstream sources – which means, among other things, a trusted source. Users can use permissions to limit who can run an AppImage package, but security goes no further. Other security modifications would be necessary to make AppImage a suitable candidate.

Flatpak

Flatpak [5] development is led by Red Hat developer Alexander Larsson. It uses the same basic command for building and managing packages, as well as enabling repositories, differing only in the options available for each task. Perhaps because Flatpak has the examples of other package managers to draw upon, instead of being developed from scratch, it has some of the clearest and most consistent names for options and commands ever to find their way into a man page. In particular, unlike DEBs, Flatpak provides separate commands for installing and updating, instead of giving them the same name.

Packages for Flatpak are sandboxed by having separate run-time environments and by using standard Linux tools such as `cgroups` [6] and `namespaces` [7] to assign resources to each package and isolate them from each other. Because these tools are standard in recent kernels, they do not significantly add to the size of packages, especially because Flatpak can use already installed dependencies as well as new ones.

Additionally, each version of each Flatpak is a separate branch in the Git-like repository OSTree [8]. This arrangement makes it easy for different releases of the same application to co-exist on the same system, and for packages to be rolled back and deleted cleanly. The use of OSTree also gives Flatpak additional levels of security, with checksums for each individual file and for checking out branches that are in use.

Moreover, using OSTree also has the advantage of making packages editable with the `ostree` command. Some overlap with Flatpak commands results, but `ostree` also has a number of advanced tools, such as administrative cleanup and built-in diff. Users can ignore `ostree` and most likely get along fine much of the time, but learning the command might come in handy at times.

The disadvantage of using OSTree is that only one branch (i.e., software version) can be active at one time. That means that system upgrades do not go into effect until after a reboot.

Flatpak runs on many major distributions, but, as I write, few packages or repositories are available. Probably the best tool for exploring Flatpak is the LibreOffice pack-

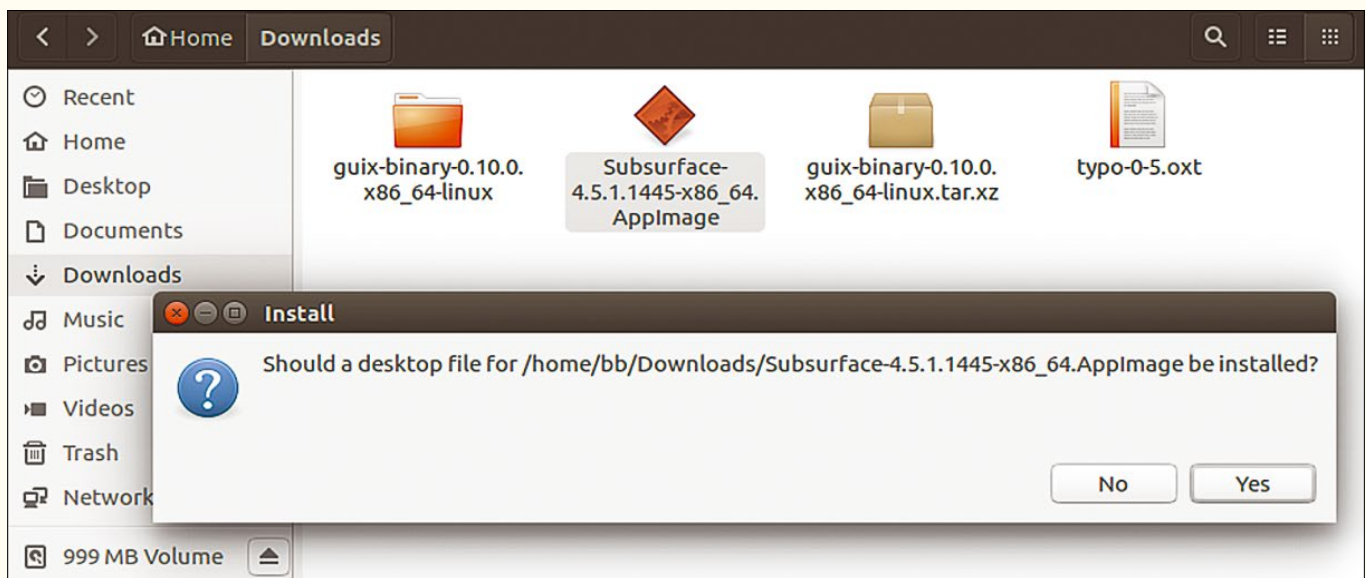


Figure 1: AppImage offers to create a launcher the first time one of its packages is run.

age, whose download page also details clearly how to set up a repository [9]. Should your system segfault (i.e., trigger a memory access violation) in the process, try again with the most recent kernel you can find.

Guix

Guix [10] is a member of the GNU Project, the software collective associated with the Free Software Foundation (FSF). In fact, early in 2016, Guix became the first package manager to be ported to GNU/Hurd, the long-awaited alternative to the Linux kernel [11].

Guix is also the package manager for the GuixSD distribution [12], although it can run on any distribution, even alongside another project manager. It is not available as a package itself, but its installation, while not graphical, is still relatively easy to use with the instructions.

Guix (Figure 2) is based on an older package manager called Nix [13]. The two parted company several years ago, and today Guix is written in GNU Guile, while Nix uses a variety of programming language. What both have in common is the metaphor of package management as a mathematical function that is affected by scripts and other inputs but cannot itself alter the system environment or edit files outside of its own directories beneath /gnu/store.

This arrangement has numerous advantages. Unlike with a .deb package, a Guix package cannot be half-installed; either it is completely installed, a feature that makes a rollback to an early version of a piece of software much easier, or it is not installed at all. The isolation of packages means that ordinary users can install software for their own use that does not challenge security, because each package is isolated from the others. Most importantly, package isolation creates a level of security comparable to that of a container without the extra overhead.

From a technical perspective, these features make Guix a serious competitor to Flatpak and Snappy. If it lacks a corporate backer, the FSF may be a suitable substitute, especially because its main interest is to keep the core parts of an operating system available as a free license.

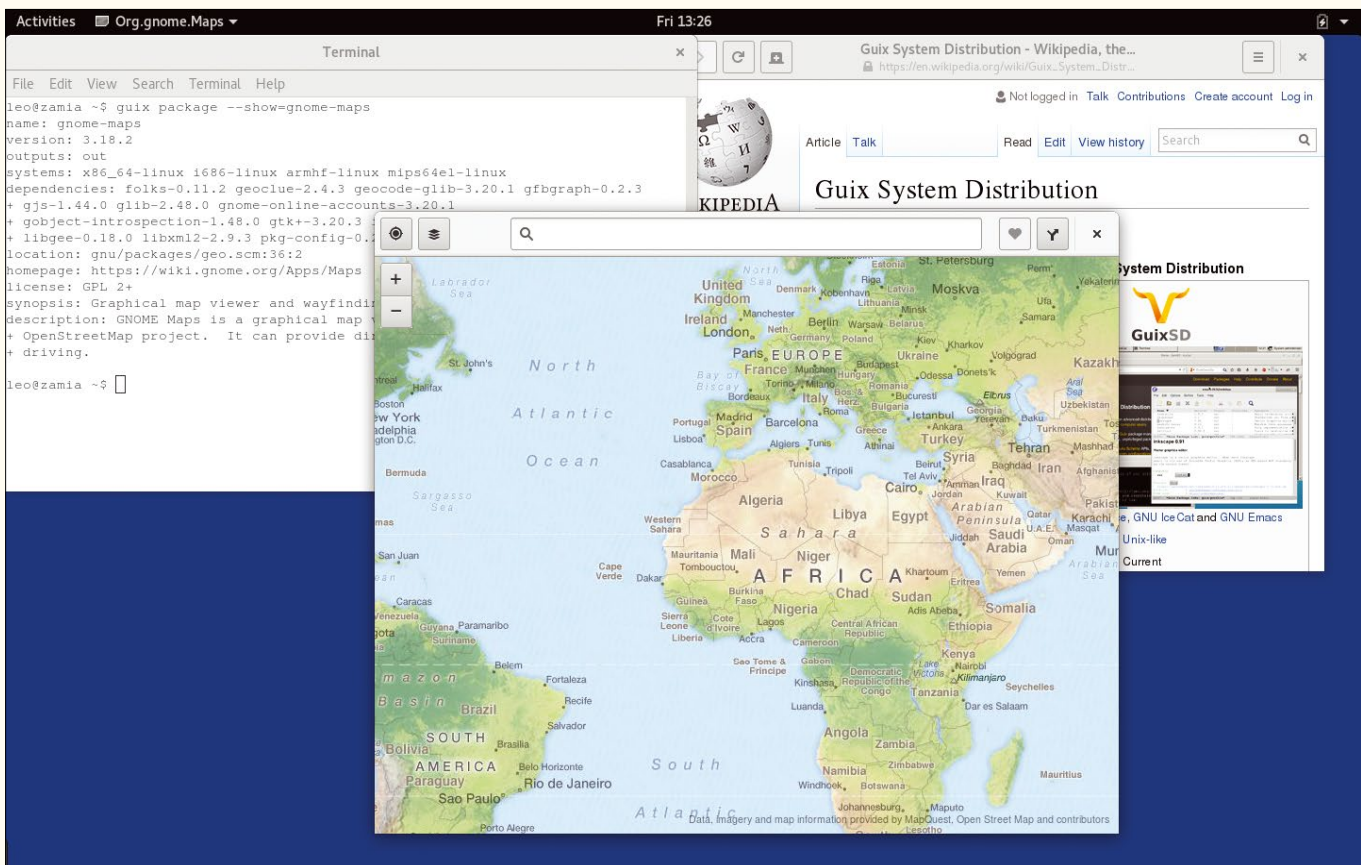


Figure 2: Output of Guix while installing a package.



```
bb@bb-VirtualBox: ~$ sudo snap install freecad
221.57 MB / 221.57 MB [=====] 100.00 % 1.57 MB/s

Name      Version  Rev  Developer  Notes
freecad   0.17     4    vejmarie   -
```

Figure 3: Using the snap command to install FreeCAD.

Snappy

Snappy packages (Figure 3) got off to a rocky start when Canonical Software, which controls Ubuntu, announced them as though they had already become the universal package manager. A typical headline in the media was “Snap Packages Become Universal Binary Format for All GNU/Linux Distributions” [14], when the truth was less dramatic.

Originally developed for embedded systems, Snappy’s packages, like AppImage’s, are reminiscent of the static tarballs used in the 1990s, combining all libraries, programs, and scripts in a single file. The main difference from static tarballs is that .snap packages add containers and draw on recent Linux features for an added level of security.

Shortly after Snappy was announced, the rumor went around that its packages were as bloated as AppImage’s. The source of the rumor was a LibreOffice snappy package whose debugging tools made it appear 500% larger than the equivalent .deb or .rpm when the truth was closer to 20%. Still, even a 20% difference in size can be significant on busy systems, or older ones.

Even more seriously, while Snappy, Flatpak, and Guix all offer similar features – everything from extra security to rollbacks and the co-existence of multiple versions of the same application – Flatpak and Guix do so with less overhead. True, Snappy packages can be built without containers in developer builds by using the --devmode option, yet that only creates a new danger. Although Snappy documentation insists that packages built with --devmode will never be included in official repositories, it is easy to imagine users preferring unofficial repositories to save memory – and, in doing so, removing a major part of Snappy security. Such a practice is more unlikely in Flatpak or Guix, whose security is architectural and probably more difficult to bypass.

Some users, too, might balk at a universal package manager being developed largely inside Canonical and under a contributors’ agreement that gives Canonical almost complete ownership of any software that is produced [15].

However, Canonical remains one of the most influential Linux-based companies. Already, there seems to be more software available in .snap packages than in any of the other formats, so Canonical might win the competition through publicity and availability alone.

Waiting for the Endgame

Universal package managers still have a ways to go. As I write, for example, graphical versions of Flatpak, Guix, and Snappy have yet to appear, although they are a high priority. In places, too, although basic functionality is available, there is still plenty of room for enhancements. As Fedora’s replacement of yum with dnf shows, adjustments to package management can be more difficult than many expect [16].

In the end, which candidate becomes universal – if any – is unlikely to be decided for several years. Possibly, while the competition fuels the commercial rivalry between Canonical and Red Hat, few distributions will see the need for a universal package manager, all the more so because the conversion to a new format would be a major effort.

Moreover, although the .deb and .rpm rivalry can sometimes be a nuisance, it is at least a known nuisance. And whereas DEBs had an advantage around the turn of the millennium, RPMs have been roughly functionally equivalent for more than a dozen years.

Under these circumstances, we can expect a lot of discussion and little resolution. Unless some statement or new feature highlights the issues involved, I expect the situation to remain much the same. ■■■

INFO

- [1] Debian Policy: <https://www.debian.org/doc/debian-policy/>
- [2] Strengths of DEBs: <https://lists.debian.org/debian-devel/2016/06/msg00287.html>
- [3] AppImage: [https://en.wikipedia.org/wiki/AppImage_\(packaging_method\)](https://en.wikipedia.org/wiki/AppImage_(packaging_method))
- [4] AppImage apps: <https://bintray.com/probono/AppImages>
- [5] Flatpak: <http://flatpak.org/>
- [6] cgroups: <https://en.wikipedia.org/wiki/Cgroups>
- [7] Namespaces: https://en.wikipedia.org/wiki/Linux_namespaces
- [8] OSTree: <https://ostree.readthedocs.io/en/latest/>
- [9] LibreOffice Flatpak: <https://www.libreoffice.org/download/flatpak/>
- [10] Guix: <https://www.gnu.org/software/guix/>
- [11] Adding GNU/Hurd support to GNU Guix: <https://fosdem.org/2016/schedule/event/guixhurd/>
- [12] GuixSD: http://dustycloud.org/misc/talks/guix/chicagolug_2015/guix_talk.html
- [13] Nix: <https://nixos.org/nix/>
- [14] LibreOffice SNAP package: <https://skyfromme.wordpress.com/2016/06/16/a-third-of-a-libreoffice-snap/>
- [15] Canonical contributor agreement: <http://www.ubuntu.com/legal/contributors/submit>
- [16] “Will DNF Replace Yum?” by Bruce Byfield: <http://www.linux-magazine.com/Online/Features/Will-DNF-Replace-Yum>





Editing images with Pixeluvo

Pixel Love

Nimble, slick, and user-friendly, Pixeluvo is a capable tool for editing images and photos on Linux. This article can help you to master Pixeluvo's basics. *By Dmitri Popov*

DMITRI POPOV

Dmitri Popov has been writing exclusively about Linux and open source software for many years, and his articles have appeared in Danish, British, US, German, Spanish, and Russian magazines and websites. Dmitri is an amateur photographer, and he writes about open source photography tools on his Scribbles and Snaps blog at scribblesand-snaps.wordpress.com.

When it comes to photo retouching and image editing, Gimp is probably the most popular choice among Linux users. After all, it's a mature and powerful application available in the software repositories of many Linux distributions. You can find plenty of Gimp tutorials on the web, and due to the application's popularity, finding answers to your Gimp-related questions is relatively easy. Better still, Gimp is free in every sense of the word.

So, why then would you opt for something like Pixeluvo [1], a closed-source

image and photo editing tool that costs money? For a number of reasons, actually. Gimp's interface is an acquired taste with its fair share of shortcomings (e.g., take a look at how the save and export features work). The application is loaded with features, and as such, it's not the easiest tool to master. And in most cases, you won't even need all the functionality Gimp has to offer. In short, if you are looking for an image editing application that offers just the right amount of functionality wrapped into a polished and user-friendly interface, then Pixeluvo is worth the money.

Getting Started with Pixeluvo

The Linux version of Pixeluvo is available as DEB and RPM packages for the 32-bit and 64-bit platforms. Download the appropriate package, and install it using the default package manager. To do this on Ubuntu and its derivatives, run the `sudo dpkg -i pixeluvo_x.x.x._xxx.deb` command.

To start working on an existing image, choose *File | Open* and select the desired file. You can also open an image or a photo by dragging it onto Pixeluvo's working area. Additionally, you can push images from other applications to Pixeluvo via the *Open With* context menu command.

Pixeluvo puts essential tools at your fingertips and neatly tucks the rest under various menus. As a result, the application features an unobtrusive interface that makes it easy to perform basic actions (Figure 1). The main toolbar on the left side of the main window gives you access to standard tools such as Crop, Straighten, Red Eye Removal, Dodge/Burn, and others. The palette in the upper-right corner displays a thumbnail of the currently opened image, while the slider underneath it can be used to zoom in and out (you can also do this using the mouse scroll wheel while holding down the Ctrl key). When you zoom in on the image, drag a rectangle area in the palette to quickly move around it. Hover the mouse over the palette, and you can use the buttons in the sliding drawer to switch between the thumbnail, histogram, and info views.

The toolbar gives you access to several advanced yet easy-to-use tools. The Spot

Heal tool, for example, lets you effortlessly remove specks and small blemishes on the image. Select the tool, point the mouse to the spot you want to remove, draw over it while holding down the right mouse button, and watch the spot magically disappear. Although the Spot Heal feature can be useful for fixing small blemishes, the Clone tool is suitable for removing large unwanted objects by painting them out using other parts of the image. Select an area of the image to clone by holding down the Alt key and clicking at the desired position. Then use the mouse to paint out the specific part of the image.

Magic Wand can come in rather useful when you need to quickly select regions of similar color. Click and hold down the right mouse button to start the selection, and then drag the mouse to specify the color matching tolerance (i.e., the size of the selection). Click on the appropriate icon in the options bar to extend the current selection or subtract from it (or use the Shift and Alt keys while making a selection).

In addition to an assortment of color adjustment tools, Pixeluvo features a handful of ready-made filters that let you quickly apply a specific effect to the image. Choose *Color | Quick Color* (or press Ctrl + R), and click on the desired filter. Use the *Blend* and *Vignette* slider to adjust the effect's strength and vignetting, and then press *OK* to apply the filter. The options under the *Effects* menu give you access to various effects: from Tilt-Shift and Local Contrast to Film Grain and Light Leaks.

Although some effects are more useful than others, there are a couple of entries in the list that can come in handy in many situations. The Soften Skin effect, for example, offers an easy way to make portraits more pleasing, whereas the Reduce Noise, Sharpen, Boost Detail, and Enhance Clarity effects can improve the overall image quality with a minimum of effort.

Need to convert a photo to black and white or duotone? Pixeluvo has several ready-made filters for that. Choose *Color | Monochrome*, pick the filter you like in the *Monochrome* palette, adjust the *Blend* and *Vignette* parameters if needed, and hit *OK* (Figure 2).

If you find the default monochrome filter selection limited, the Threshold tool gives you more control over black-and-white conversion. This tool converts an image to black and white based on the brightness of each pixel. The threshold value controls which pixels are converted to black and which are converted to white. Pixels that are brighter than the threshold value become white, while all other pixels become black.

Several other creature comforts can help you to get good results with a minimum of effort. As the name suggests, the Replace Color tool can be used to easily replace a color in the image with another color with different hue, saturation, and lightness properties (Figure 3). Although the cross-process effect has become a cliché, when used sparingly it can liven up otherwise dull images. And, Pixeluvo's Cross Process tool lets you apply the effect with consummate ease. Photo Filter is another handy tool that simulates various color filters. It can be used to remove color casts or make the image warmer or cooler. Used creatively, this tool can also help to produce interesting effects.




Figure 1: Pixeluvo features an unobtrusive interface.

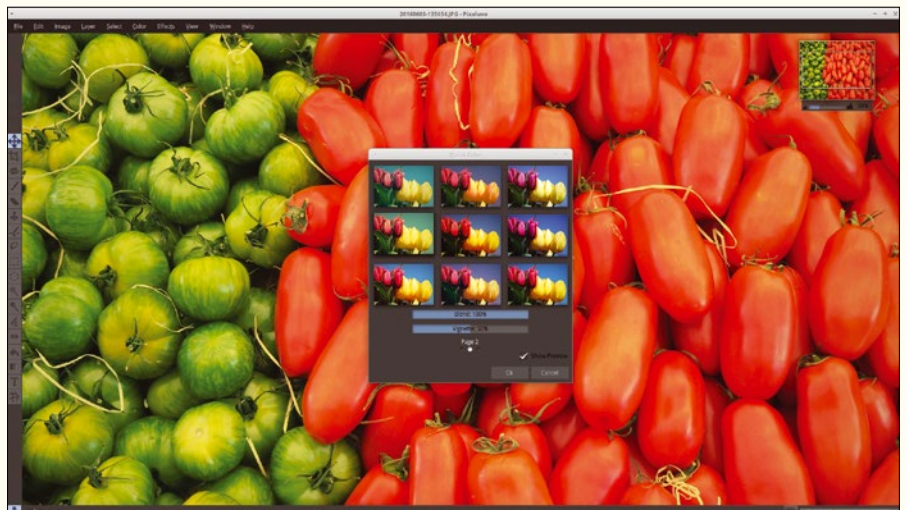


Figure 2: The Quick Color feature offers a collection of ready-made filters.

Pixeluvo also offers a comprehensive set of color adjustment tools, including Curves and Levels as well as tools for adjusting saturation, brightness, contrast, etc. (Figure 4). All these tools follow the established paradigm, so if you understand their basic operation principles, you won't have trouble using these tools in Pixeluvo.

To make the image editing process more efficient, Pixeluvo provides extensive support for keyboard shortcuts, and practically all tools and features can be activated via dedicated keys. Better still, the tools in the main toolbar have single-key bindings (e.g., Q for Crop, H for Spot Heal, etc.), so they are easy to remember and use.

Adjustment Layers and Masks

Pixeluvo supports nondestructive editing via so-called adjustment layers (Figure 5). Instead of applying color adjustments directly to the image, you can create a separate layer on top of it and perform edits there. This approach has several obvious advantages. For starters, the original image remains untouched, and you can hide and undo edits by simply disabling or removing the layer. You can also separate different adjustments by creating individual layers. More importantly, if you apply adjustments directly to the image, you can't change them later. Using adjustment layers solves this problem, too.

Pixeluvo's layer functionality follows the established norms, so if you have ever worked with layers in Gimp (or other image editors), you'll quickly find your way around. And, if you have no previous experience working with layers, Pixeluvo is an ideal platform for mastering this functionality, thanks to the application's focus on simplicity and user-friendliness.

Creating an adjustment layer is a matter of choosing the appropriate entry from the *Layer | New Adjustment Layer* menu. You can then perform adjustments in a regular manner. Changes you make to a layer apply to all other layers below it. This means that you can control the scope of specific edits by rearranging layers in the current layer stack. The *Layers* panel in the lower-right corner of the working area can be used to add and manage layers.

The panel contains the current layer stack, and you can rearrange layers by dragging them up and down with the mouse. It's also possible to modify adjustments of a specific layer at any time by double-clicking on the *Settings* icon



Figure 3: Creature comforts like the Replace Color tool simplify many image editing actions.

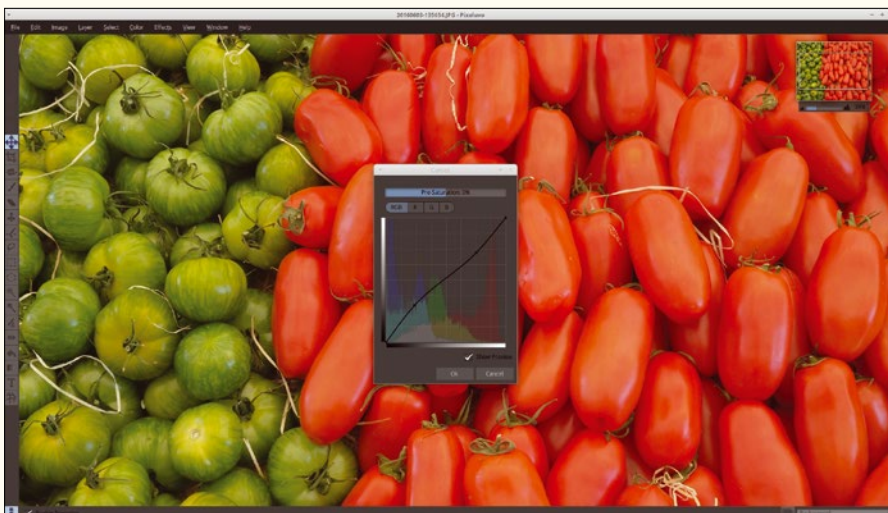


Figure 4: The application includes standard tools for curve and level adjustments.




Figure 5: Pixeluvo supports nondestructive editing via adjustment layers.

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

LINUX UPDATE
EXPLORING THE WORLD OF LINUX

- SuperCollider
- Epoch 1.0 Released
- Wayland 1.4 Challenges X11
- Reading the Tea Leaves with Canonical's Financial Statement
- Qubes OS

FEATURED ARTICLES

SuperCollider
SuperCollider is a versatile and powerful sound generator that can be used to create, filter, and modulate sound. In this article, we tackle the basics.

Epoch 1.0 Released
New alternative for init comes with a small footprint and minimal configuration. (more)

Wayland 1.4 Challenges X11
X marks the target for the next-generation windowing system. (more)

Reading the Tea Leaves with Canonical's Financial Statement
Strictly speaking, the recently published annual report for the Canonical Group, the developers of Ubuntu, covers only its United Kingdom subsidiary for the year ending March 31, 2013. However, in practice, the report provides glimpses into both Canonical's global activities and the challenges that the company faces before it can become profitable. (more)

Qubes OS
Most operating systems claim to be secure; however, most fail terribly. Newcomer Qubes OS tries a different approach, relying on a microkernel and pervasive virtualization. (more)

FURTHER READING

- New Worm Attacks Linux Devices
- Will DNF Replace Yum?
- Sneaky New Linux Attack Discovered
- Fixing Disks with Parted Magic
- YouTube to MP3

Partner Link
Save Money on Storage with a Cluster in a Box

DrupalCon Austin
Get "Drupal 8 Ready" at DrupalCon Austin
DrupalCon Austin will be a can't-miss event for anyone who wants to be up to speed on Drupal 8 and understand its full potential.
With a full week of sessions, training courses, keynotes, vendor exhibits, and FUN, DrupalCon Austin is the place to be!
Reserve your spot today for the event, June 2-6, 2014.

GIMP Handbook Out Now!

MOST READ
Android Admin Tools
Smartphones and tablets have lost their gadget status and become part of the system administrator's tool kit. We look at the... (more)



www.linuxpromagazine.com/mc/subscribe




Figure 6: Pixeluvo makes it easy to work with masks.

next to the desired layer and tweaking the available parameters.

In addition to layers, Pixeluvo supports masks that control which part of the layer should be visible (Figure 6). This functionality can be put to a wide variety of creative uses. For example, if you want to desaturate only a specific area of the image, select the desired area of the image using one of the available selection tools and add a new saturation adjustment layer. Pixeluvo is intelligent enough to automatically convert the selection into the mask in the added adjustment layer, so any adjustments you make will apply only to the selected region.

You can add a mask to an existing layer, too. Click on the *Add Layer Mask* icon next to the existing layer, and Pixeluvo turns the current selection into the mask. Mastering layers and masks requires some experimenting, but once you've figured out how to use them, you can achieve rather impressive results.

Pixeluvo and RAW Files

Although Pixeluvo isn't billed as a RAW processor, it's perfectly capable of handling RAW files from most manufacturers (Figure 7). The application gives you full control over exposure, color, and tone curves. You can also adjust brightness and contrast as well as apply the cross-processed filter. Instead of adjusting individual parameters manually, you can let Pixeluvo do the donkey job by pressing the *Auto-Exposure* button. This usually produces decent results. Pixeluvo places the RAW file in its own layer, and you can use adjustment layers to perform non-destructive editing.

Final Word

It's hard not to like Pixeluvo. The right mix of features, attention to detail, and a polished user-friendly interface makes Pixeluvo a worthy alternative to existing image editing applications on Linux. Pixeluvo costs \$34, a small price to pay for such a great and genuinely useful tool. It's also a good way to support the developer and his efforts to maintain the Linux version of the application. If you want to give Pixeluvo a try before parting with your money, you can download a 30-day trial from the application's website. ■■■




Figure 7: The application can handle RAW files, too.

INFO

[1] Pixeluvo: www.pixeluvo.com

LOST YOUR BOOKSTORE?

LET US BE YOUR BOOKSTORE



Browse our shop for single issues of *ADMIN*, *Linux Pro*, *Linux Magazine*, *Ubuntu User*, *Drupal Watchdog*, and *Raspberry Pi Geek* – delivered right to your door.

■ shop.linuxnewmedia.com/single

Better yet, subscribe, and you won't need a bookstore.

■ shop.linuxnewmedia.com/subs



shop.linuxnewmedia.com

DIGITAL AND PRINT EDITIONS AVAILABLE!

5 Tor Alternatives
options for peer-to-peer networking

Free DVD!
NOOBS 1.9.0

DISCOVER DRUPAL'S
PASSIVE SECURITY
MODULES





Mofo Linux tears down virtual walls

Crossing Frontiers

Mofo Linux enables secure digital communications even in places where it is politically or ideologically unwelcome.

By Ferdinand Thommes

Mofo Linux claims to help facilitate global freedom of information. It comes with the tools needed to work around politically motivated firewalls in countries such as China, Iran, Turkey, Thailand, Syria, the United Arab Emirates, and other countries with repressive regimes. The toolset includes various VPN flavors and other tools such as I2P, Cjdns, and Lantern along with Tor, Tor Browser, and Tor Messenger. The system also offers Arab and Chinese localization in addition to English.

Mofo Linux [1] was forked from Porteus, the portable distribution based on Slax and thus ultimately on Slackware, in 2011. Meanwhile, the system has matured to a fully fledged distribution and changed its underpinnings to Ubuntu 15.10. The

live image with installer is available for downloading from SourceForge [2] and weighs in at about 1.6GB. You can burn the image onto an optical disk or transfer it to a USB stick using Universal Netboot Installer (Unetbootin) and the dd console tool. Alternatively, you can try it out in live mode with virtual machines like VirtualBox. The project took over the installer virtually unchanged from Ubuntu. It offers to set up Logical Volume Manager (LVM) as well as to encrypt the entire system or just the home directory.

Tor or I2P

Mofo Linux uses the well-known Unity 7 Desktop interface (Figure 1); however, there are some fairly obscure icons on the left side of the screen. This is where the project locates applications that serve the specific purpose of the distribution. At the top, below the mandatory search box, is the Tor Browser [3]. It directs outgoing requests by the browser across multiple servers of the Tor anonymization network to conceal the identity of the user.

The anonymous overlay network Invisible Internet Project (I2P) [4] is also pre-installed. Although Tor mainly offers anonymous proxy servers for various Internet services including IRC, email, and HTTP, everything primarily occurs within the network in the case of I2P. This also removes the need for a node database of directory servers as required for Tor, and which runs the risk of de-anonymization under certain circumstances. Another benefit: The service uses end-to-end-data encryption. Tor only offers this if you use the “hidden services.”

However, the Tor Browser failed to launch in the present image, delivering instead a fairly meaningless error message. Actually, the Tor Browser Launcher should start when you first start Tor to download and set up the web browser. In testing, we needed to visit the *Ubuntu Software Center* (Figure 2) and install the Tor Browser Launcher at the command line with the `torbrowser-launcher` command to get Tor working and create a working Tor Browser (Figure 3).

Sharing Securely

Below I2P in the bar is the button that launches OnionShare [5]. The tool is used for a secure exchange of files and requires a Tor Browser running in the background to use its anonymization services. For this

AUTHOR

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




Figure 1: Ubuntu's default desktop Unity 7 provides the user interface for Mofo.

purpose, a small web server is launched locally, which Tor makes available online as a *hidden service*. The files are assigned a random URL via which they can be downloaded. The size of the file does not matter. A document on GitHub reveals more technical details [6].

The next item in the bar is Bitmask VPN [7], one of several VPN clients available in Mofo. Bitmask (Figure 4) provides Internet connections and email encryption. You can choose between predetermined VPN providers, such as Riseup, or use your own.

Restrictive Corners

The next tool is aimed primarily at users who travel to the countries where use of the Internet is subject to restrictions or censorship. The Lantern proxy service [8] tries to work around blocked websites with up-to-date information on the situation in the country (Figure 5). The service relies on peer-to-peer techniques where a network of users share bandwidth in countries without restrictions with users from

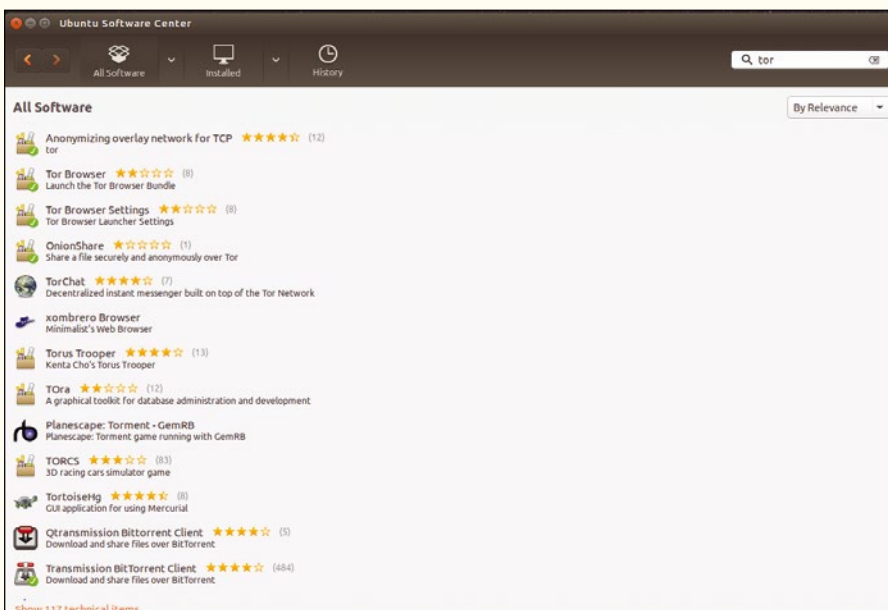


Figure 2: Before you can use the allegedly installed Tor Browser, you first need to download and install it.

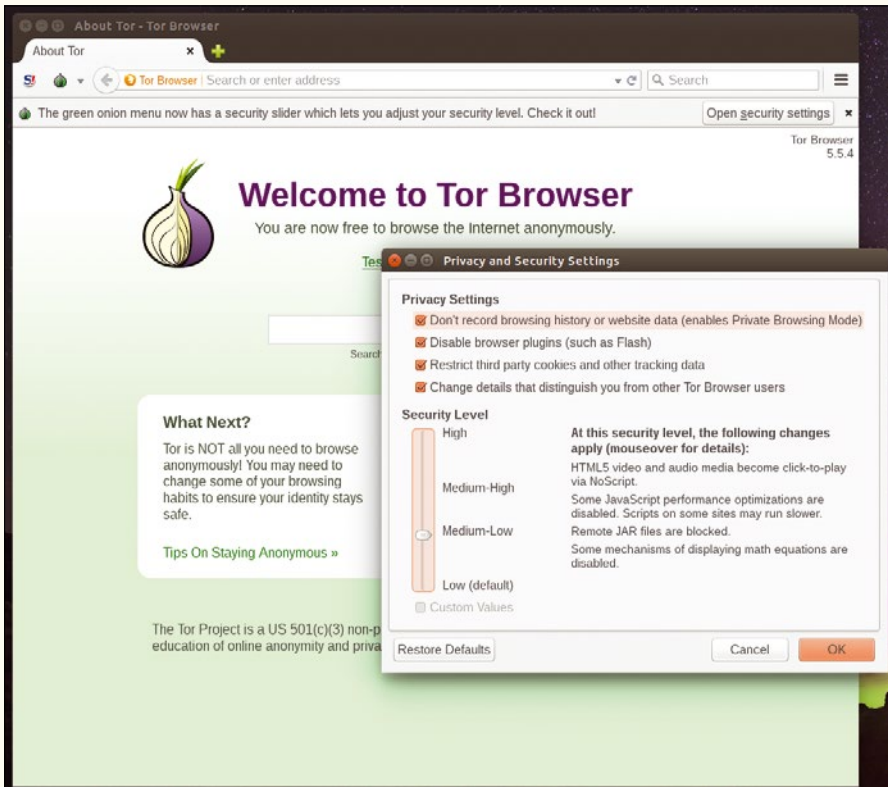


Figure 3: The security level of the Tor Browser can be easily adapted to your own needs using the slider.

countries where access is subject to restrictions.

The U.S. State Department funded the service with \$ 2.2 million. An interesting detail is that Lantern developer Adam Fisk operated the LimeWire file-sharing client up until 2010 when it was prohibited by a court ruling in the United States. In doing so, he developed the techniques that make Lantern eligible for funding today. So far, China's repeated attempts to block the service have been unsuccessful.

You need to be very careful with Popcorn Time [9] (Figure 6). It's a BitTorrent client with an integrated media player, that is, more or less Netflix with a built-in copyright infringement. Users of this convenient service are fiercely pursued by the copyright holders of the works provided and receive mass warnings. Why the developer integrated the service into a distribution trimmed for preserving privacy remains unclear.

Container Tricks

VeraCrypt [10] is responsible for encrypting the entire system, individual partitions, or "containers." Containers in this

case consist of files of a fixed size, which the system treats like virtual drives after decryption. The software offers the option of hiding one container in another, making it impossible to prove its existence. VeraCrypt (Figure 7) was founded in 2012 as a fork of TrueCrypt, which was abruptly discontinued by its developers in 2015 under what are still somewhat mysterious circumstances. VeraCrypt has been considered its successor since then, because the code is largely compatible.

The hidden containers in particular make encryption interesting for the purposes of

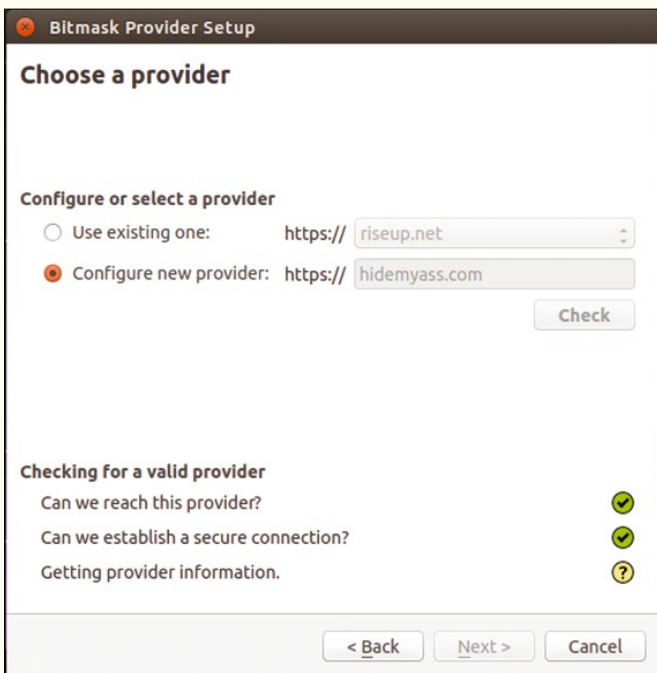


Figure 4: Bitmask VPN supports easy connecting of the computer with predefined or freely selectable VPN services.

Mofo. Encryption can be set up so that less important data is subject to normal encryption, while the hidden container contains the genuinely sensitive data. If the authorities ask the owner of a computer to hand over the password for encryption, the user can seem to cooperate without revealing compromising data. The distribution also includes zuluCrypt and eCryptfs as an alternative to VeraCrypt.

Another tool to help secure connections is offered by the CjdnS [11] network protocol, which also forms the basis of the Hyperboria network [12]. The principle here is that of mesh networking, as used in Germany by the Freifunk network. Remotely coupled mesh networks, whose single node is no longer connected

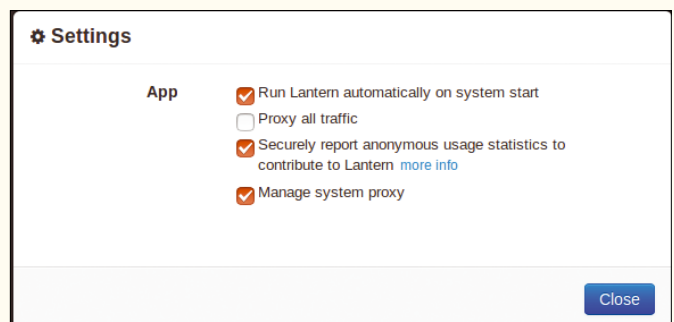



Figure 5: The Lantern proxy service only needs a few settings to operate.

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

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.



FREE DVD INSIDE!

Linux Mint 17.3
"Cinnamon" 32-bit

ORDER ONLINE:

shop.linuxnewmedia.com/specials

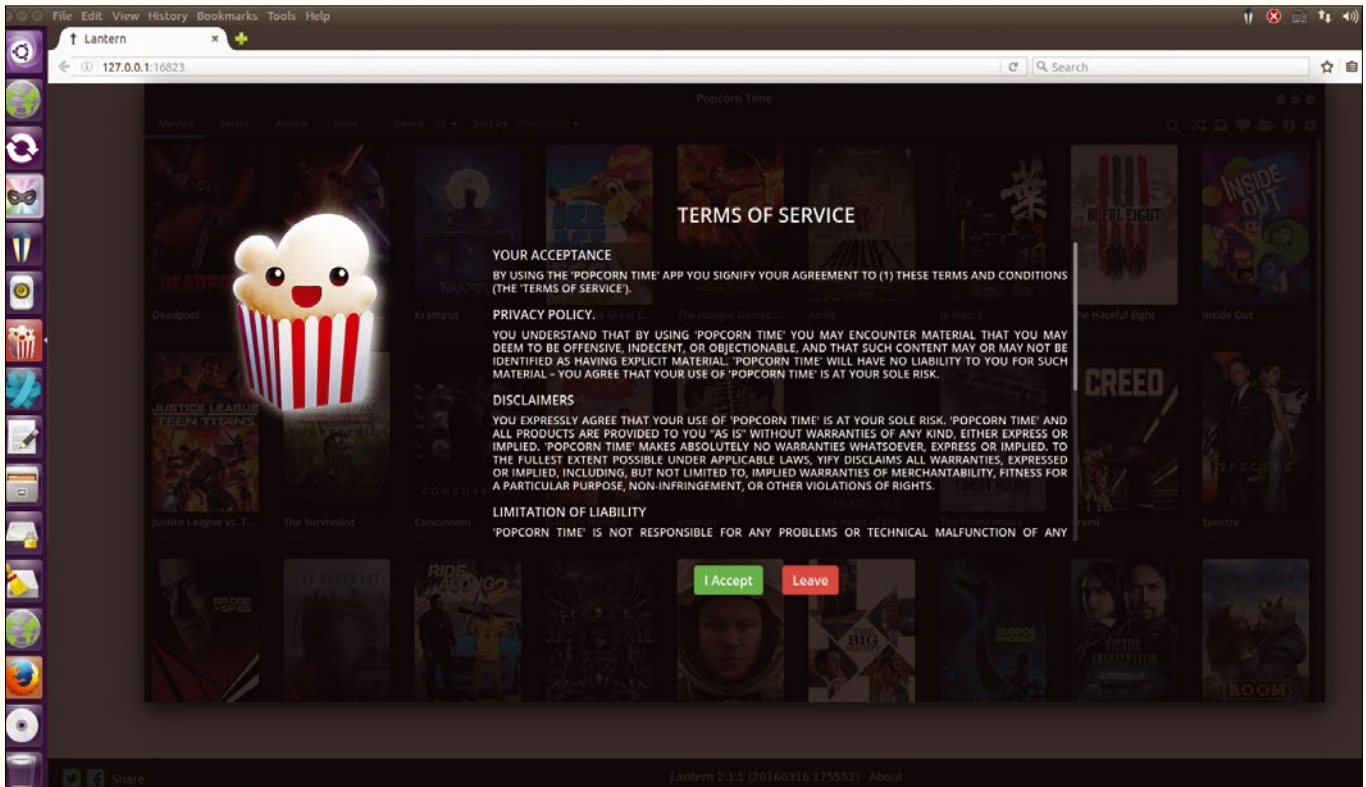


Figure 6: Why exactly the developers chose to integrate Popcorn Time – a dangerous BitTorrent client – remains unclear.

to a central access provider, more or less form their own network on the Internet if the network manages to grow to a reasonable size. Cjdns provides functions not only for managing peer-to-peer connections but also for encrypting the transport routes.

Conclusions

Mofo Linux is a quickly installed and easy-to-use distribution based on Ubuntu. The operating system is completely secondary, leaving the stage to the tools that fulfill the actual role of the distribution. It is becoming increasingly important to protect information to avoid jeopardizing the owner or other persons. Mofo Linux takes care of this and offers a range of well-integrated tools to do so. Only the initial failure of the Tor Browser to launch, despite being a key feature of the distribution, could scare off less savvy users. The developers need to improve here with a view to ensuring ease of use, which Mofo otherwise consistently offers. ■■■

INFO

- [1] Mofo Linux: <http://mofolinux.com>
- [2] Mofo Linux download: <https://sourceforge.net/projects/mofolinux/>
- [3] Tor Browser: <https://www.torproject.org/projects/torbrowser.html>
- [4] I2P: <https://en.wikipedia.org/wiki/I2P>
- [5] Onionshare: <https://onionshare.org/>
- [6] Security of Onionshare: <https://github.com/micahflee/onionshare/blob/master/SECURITY.md>
- [7] Bitmask: <https://bitmask.net>
- [8] Lantern: <https://getlantern.org>
- [9] Popcorn Time: <http://popcorn-time.to>
- [10] VeraCrypt: <https://veracrypt.codeplex.com>
- [11] Cjdns: <https://en.wikipedia.org/wiki/Cjdns>
- [12] Hyperboria: <https://hyperboria.net>

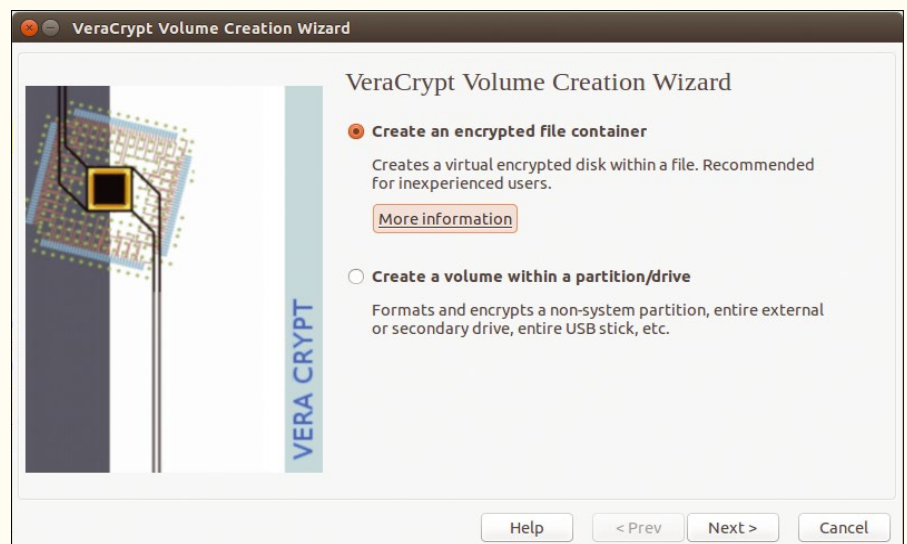


Figure 7: VeraCrypt encrypts systems, partitions, and USB flash drives.

IT Highlights at a Glance



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

- ADMIN HPC
- Linux Update
- ADMIN Update
- Raspberry Pi

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

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

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

Raspberry Pi: www.raspberry-pi-geek.com/mc/subscribe

The continuing drama of RISC vs. CISC

Full Speed Ahead

Crank the compilers and turn up the system clock; maddog explains the basic differences between RISC and CISC architectures. *By Jon "maddog" Hall*

Recently, it was announced that Hewlett Packard had won a long-outstanding lawsuit against Oracle, citing that Oracle had violated a business agreement requiring support for the Itanium processor. I remember when the Itanium was announced. All I could do was slap my head and say, "Again?" How many times did we have to invent Ultra Wide Instruction Set Computing (UWISC) in computer science?

At that time Digital was putting all of its thrust behind a Reduced Instruction Set Computing (RISC) processor named Alpha, after years of making Complicated Instruction Set Computing (CISC) hardware, such as the PDP-11 and VAX.

The basic difference between CISC and RISC architectures is that for every instruction the programmer codes, tens, hundreds, or even thousands of tiny little instructions (called "microcode") in a CISC computer tell the hardware exactly what to do to solve a potentially complex set of tasks, whereas in a RISC computer, each machine language instruction basically does one tiny little thing. You might need tens, hundreds, or thousands of RISC instructions to do the same work as a CISC instruction – but then again, you might not. It depends on what the compiler generates.

CISC instructions are fabricated to reflect what the user may do. For example, you might have a single CISC instruction that compresses a string of bytes in memory or creates a JPEG. Naturally, these instructions tend to lag the industry, because the standard for what is being done has to be created first. Unfortunately, as the industry shifts, the CISC instruction may become less useful, yet it still has to remain in the instruction set until it can be retired – a long and painful process.

In a CISC computer, a significant amount of space on the chip holds the microcode. This space takes away from the space available for more registers, more cache, and more threads than would be allocated on a RISC processor chip. This is especially true of a 64-bit architecture, which tends to take up more space for registers and cache than a 32-bit system. To be fair, the extra machine-code-level instructions that a RISC processor needs to do the work of a single CISC instruction also slows down the task and takes up cache inside the CPU.

Because it can be difficult for a compiler of a high-level language to match up exactly what it is trying to do with the instruction set of a particular processor, CISC machines may be slower in some processes than RISC machines. RISC machines tend to put a lot of emphasis on highly optimized compilers and high-speed clocks for the CPU.

The argument of CISC vs. RISC has been going on since the days of Maurice Wilkes and Alan Turing, with Dr. Wilkes taking the side of CISC and Alan Turing taking the side of RISC. I am firmly in the RISC camp.

In a UWISC, such as the Itanium, multiple instructions are processed in parallel to make the machine execute "faster." The compiler loads up the CPU with instructions and data and then the CPU executes all of the instructions at one time, making the flow "parallel." This is great in theory if your application allows the parallelism that the CPU is capable of doing to match up with what is needed. Unfortunately, most of the time it does not.

UWISC has other issues. In the VAX line, a program called AXE tested each instruction (and there were a lot of them), each addressing mode (and there were a lot of them), and each data type (and there were a lot of them). Running at a million instructions a second, the AXE program took a month to run through all the combinations ONE TIME.

On a RISC machine, because of the small number of instructions and the almost non-existence of addressing modes, the same AXE program might take a day, shortening the design cycle for the processor.


How long would it take to run AXE for a UWISC system, with all the combinations for a CISC computer and multiple instructions in parallel? AXE might take a year to complete one run.

I believe that Intel created the Itanium (which many people called the Itanic) to stop AMD from duplicating their instruction set. But AMD fooled them, extended the 32-bit instruction set to make a 64-bit system, and blew past Intel in the 64-bit space. The Itanic was taking on water and starting to sink.

I am sure that this article will generate lots of comments. Some people will disagree, but I will stick with Alan Turing. Crank the compilers, turn up the system clock, add more cores and more cache, full steam ahead! ■■■

THE AUTHOR

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®.



MORE UBUNTU!



Can't get enough Ubuntu? We've got a whole lot more!

Ubuntu User is your roadmap to the Ubuntu community. In the pages of **Ubuntu User**, you'll learn about the latest tools, best tricks, and newest developments in the Ubuntu story.

Ubuntu User helps you explore the treasures of open source software within Ubuntu's expansive repositories. We'll bring you exclusive interviews with Ubuntu leaders, keep you current on the exciting Ubuntu community, and answer your most perplexing Ubuntu questions. Learn how to choose a video editor, find the perfect tool to customize your desktop, and configure and manage Ubuntu systems using the best admin tools.

DON'T MISS ANOTHER ISSUE!



HUGE SAVINGS OFF THE NEWSSTAND PRICE!

SUBSCRIBE NOW: SHOP.LINUXNEWMEDIA.COM

How Free Software Got Its Fonts

Sharing Alike

With the advent of free software for non-programmers, users ran into a licensing dilemma in a world of proprietary fonts.

By Bruce Byfield

Most Linux users soon hear of the influence of the GNU General Public License (GPL) in the development of free software. However, fewer have heard of the influence of the SIL Open Font License, although it is as important for design as the GPL has been for software. Just as the GPL is responsible for the development of free software, so the SIL Open Font License enabled the rise of the free font movement, making Linux a practical choice for designers and artists.

Today, it is the most popular free license for fonts, although few know its story.

Today, many distributions install with a variety of fonts, some for enabling multiple languages, and others for design. Yet, until the early years of the millennium, fonts were an afterthought

in free software. Free licensed fonts were packaged in most distributions, but they were intended for on-screen display, particularly in terminals. The emphasis was on monospaced fonts – fonts whose letters all took up exactly the same space – the kind familiar at the command line. If the quality was generally poor, no one minded, because the fonts were good enough for on-screen viewing, and most people using them were developers, not designers.

Around the turn of millennium, the needs of users started to change. For one thing, Gimp had developed to the point that it could be used for professional work. Just as importantly, in July 2000, Sun Microsystems released the StarDivision [1] code that became OpenOffice.org, and later Apache OpenOffice and LibreOffice, giving free software its first

full office system. According to rumor, the StarDivision developers were told that they would have to use the Writer word processor to document their work, which meant that free software also had its first intermediate desktop publisher.

The problem was that fonts at the time were either proprietary or public domain. Obviously, proprietary fonts would be incompatible with free software, but what happened when a font was embedded in a GPL document? At the same time, font designers were concerned that, if they released their work under a free license, it might be included in collections of low-quality fonts, and their reputation might suffer as a result.

This body of issues was not covered by the GPL, then in its second version. In fact, the growing feeling was that the GPL was not suited to artistic works [2], which place an emphasis on authorship that is missing in most software projects. A similar perspective led to the GNU Free Documentation License [3] in the same era.

Some of the Creative Commons licenses were more suitable for fonts, such as the Attribution-ShareAlike, but while they were an improvement over the GPL for fonts, neither they nor any other free license was ideal, either. Yet, if these concerns could be settled, free software could potentially attract a whole new body of users who were not developers.

The GPL Font Exception

One attempt to address these issues was the GPL Font Exception [4]. Written by David Turner in 2005, the Font Exception is an optional clause at the end of the license:

“As a special exception, if you create a document which uses this font, and embed this font or unaltered portions of this font into the document, this font does not by itself cause the resulting document to be covered by the GNU General Public License. This exception does not however invalidate any other



reasons why the document might be covered by the GNU General Public License. If you modify this font, you may extend this exception to your version of the font, but you are not obligated to do so. If you do not wish to do so, delete this exception statement from your version.”

In other words, the document and embedded fonts are licensed separately, and the license of one does not affect the license of the other.

The Font Exception satisfied the Free Software Foundation’s concerns, and was used by Red Hat in 2007 when it first released its Liberation fonts [5], a collection of fonts that are metrically equivalent to Times New Roman, Arial, and Courier, occupying the same space as these three popular fonts, but different from them in design. However, the Font Exception did not address the concerns of font designers, so perhaps it is no accident that Google later relicensed the Liberation Fonts and released them under the Apache License 2.0, under the name CrosCore fonts [6]. The Font Exception is still used today, but it has become a minority preference.

The Rise of SIL

Another solution was the license under which the Bitstream font family was licensed by Gnome in 2003. The Bitstream Vera License [7] granted users “the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software” so long as the copyright accompanied any use of the fonts. In addition, no derivative fonts could use the words “Bitstream” or “Vera” in their names and could not be sold alone.

Around the same time, SIL International [8], a Christian academic organization, was developing fonts to promote literacy in developing countries and to encourage the use of minority and previously unwritten languages. In 2006, Victor Gaultney, one of SIL’s designers, released Gentium [9], a font he had designed for his Master of Arts in Typeface Design at the University of Reading under a license inspired by the Bitstream Vera License.

Gentium is important in the history of free fonts for two reasons. First, instead of being a derivative of a proprietary font, it was an original design (Figure 1).

Shop the Shop

shop.linuxnewmedia.com



RASPBERRY PI ADVENTURES



COOL PROJECTS FOR GEEKS OF ALL AGES

RASPBERRY PI ADVENTURES

is a one-volume special edition magazine for curious Raspberry Pi beginners. This easy, hands-on guide starts with an introduction to computers and offers a series of special hands-on projects illustrating many of the most popular uses for the Raspberry Pi.

ORDER YOUR VERY OWN ISSUE!



ORDER ONLINE: shop.linuxnewmedia.com/se23

abcdefghijklmnopqrstuvxyz
ABCDEFGHIJKLMNPOQRSTUVWXYZ
1234567890- = !@#\$%*()_ +
The quick brown fox jumped over the lazy
dog.

Figure 1: Gentium, the font that started it all.

INFO

- [1] StarDivision: <https://en.wikipedia.org/wiki/StarOffice>
- [2] Licensing artistic work: <http://programmers.stackexchange.com/questions/297102/licensing-for-artistic-work-used-inside-gplv3-licensed-software>
- [3] Why not GPL for manuals?: <https://www.gnu.org/licenses/gpl-faq.en.html#WhyNotGPLForManuals>
- [4] GPL font exception: <https://www.gnu.org/licenses/gpl-faq.en.html#FontException>
- [5] Liberation fonts: <https://fedorahosted.org/liberation-fonts/>
- [6] CrosCore fonts: <https://packages.debian.org/testing/fonts-croscore>
- [7] Bitstream Vera License: <https://www.fontsquirrel.com/license/Bitstream-Vera-Sans-Mono>
- [8] SIL International: <http://www.sil.org/>
- [9] Gentium: <https://www.linux.com/news/gentium-award-winning-font-joins-free-software-world>
- [10] SIL Open Font License: http://scripts.sil.org/cms/scripts/page.php?site_id=nrsi&id=OFL
- [11] Google Fonts: <https://www.google.com/fonts>
- [12] Open Font Library: <https://fontlibrary.org/>
- [13] Charis SIL: <http://software.sil.org/charis/>
- [14] Doulos SIL: <http://software.sil.org/doulos/>
- [15] Graphite: http://scripts.sil.org/cms/scripts/page.php?site_id=projects&item_id=graphite_home
- [16] Graphite fonts: http://scripts.sil.org/cms/scripts/page.php?site_id=projects&item_id=graphite_fonts

Based on calligraphic designs, it is a design of undeniable quality and beauty.

Second, its license was the first version of the SIL Font License [10], Revised in 2007, the SIL Font License went on to become the most popular license for free fonts. Although the requirement that a font using the license cannot be sold alone is technically a violation of the GPL, the requirement is so easy to fulfill that the SIL Font License is generally accepted as a free license. If you go to Google Fonts [11] or the Open Font Library [12], you will find that the SIL Font License is used by at least two-thirds of the available fonts.

The Legacy of SIL International

SIL International was also among the first to release fonts intended to have a complete set of Unicode characters; in

fact, “Gentium” is Latin for “belonging to all nations.” One of Gentium’s notable features is long ascenders – upper strokes on letters like “b” or “h” – that give plenty of room for diacritical marks. Other fonts like Charis SIL [13] and Doulos SIL [14] might lack the beauty of Gentium (Figure 2), but their versatility set a standard for inclusion that other free font designers have strained to follow.

Additionally, SIL International developed Graphite [15], a small font system that automatically inserts small capitals, old-style figures, and other advanced typographical features. Graphite is supported by applications such as LibreOffice and Mozilla, but only a couple of dozen fonts [16] are designed to use it.


Still, in as little as 10 years, SIL International and its license have transformed free software. Where once high-quality free fonts could be number on two hands, now there are hundreds. Where once Linux was not a platform for designers, now professionals can use tools like LibreOffice, Inkscape, Krita, and Scribus. Where once many doubted that font designers would care to participate in free software, today there are dozens working alone and in teams. Without SIL International and its license, undoubtedly free software would be much less developed today. ■■■

ABCDEFGHIJKLMNPOQRSTUVWXYZ
1234567890- = !@#\$%*()_ +
The quick brown fox jumped over the
lazy dog.

ABCDEFGHIJKLMNPOQRSTUVWXYZ
1234567890- = !@#\$%*()_ +
The quick brown fox jumped over the lazy
dog.

Figure 2: Charis (top) and Doulos (bottom) were among the first fonts to attempt to support the full range of Unicode characters.

An entire year of
**RASPBERRY PI PROJECTS
AND OPERATING TIPS**
for one low price!
Two bundles to choose from:



Hurry while supply lasts!

ORDER NOW!



<http://shop.linuxnewmedia.com/us/magazines/raspberry-pi-geek/catchup.html>

A campaign to redefine free software and hardware

Libre Laptop

The earth-friendly EOMA68 is an open hardware laptop that promises to be easy to repair, easy to upgrade, and easy to secure. *By Bruce Byfield*

The Four Freedoms [1] have been central to free software for years. Now, Luke Leighton is offering additional concerns in a crowdfunding campaign to build the EOMA68 [2], a free laptop that is environmentally responsible and security-conscious and, if successful, could help establish open hardware as a presence in manufacturing.

Leighton is a long-time contributor to free software (or, as he prefers to call it, *Software Libre*), having contributed his first patches for Samba functions in 1992. As a result of his experience, he realized that one of the key disadvantages of free software has always been that those who build it often benefit the least. He compares his own experience of having to work on construction sites at one point in his career to the discovery in 2014 of security vulnerabilities in OpenSSH and other major applications because of a lack of funding for sufficient developers.

At the same time, he became aware that companies like Google have regularly taken advantage of free software and relicensing its work under BSD-style licenses without contributing their work back to the Linux kernel. His conclusion is that “even companies that tell us that they’re doing us a favor by working with open source are not really interested in acting in an ethical or responsible fashion.”

Realizing that complaining about these problems accomplished nothing, Leighton resolved to develop ethical, privacy-respecting hardware. “From there,” he says, “I would be able to set up an ethical business fund that gave something back to the actual people who had made the business successful – and ultimately do things like fund the creation of a useful and powerful privacy-respecting processor.”

The result was the EOMA68 Standard (Figure 1), which first made the news as a result of Leighton’s collaboration in 2013-2014 with



Figure 1: An early version of the laptop showing 3D-printed parts.



Figure 2: The EOMA68 laptop. Notice the bamboo frame around the screen and the home-made housing for the computer card.

Aaron Seigo to build Vivaldi, an attempt to build a free tablet, and the Improv development board that was intended to help jumpstart free-licensed hardware. Neither succeeded, apparently as a result of inexperience, miscommunication, and the setbacks [3] that trouble most first efforts to product hardware.

From this experience, Leighton says he has learned that keeping control of product design, specifications, and production is essential. He is also determined not to deal with venture capital and is convinced that running an ethical business by dealing with unethical ones “is not going to work.” At the same time, he retains some of the manufacturing contacts that will be needed to bring a new product to market

A New Breed of Laptop

Looking back on Vivaldi and Improv, Leighton concludes that their failure may have been just as well, because, two years ago, the EOMA68 Standard was not ready for large-scale production.

Today, building on the lessons learned from past experience, Leighton is planning to build a 15.6-inch laptop with a casing made of bamboo, wood, and parts created with a 3D printer whose specifications are available (Figure 2). The case includes peripherals such as the screen, keyboard, track-

pad, and WiFi and is completed by a computer card housed in recycled PCMCIA cases that slot into the case (Figure 3). This arrangement makes the laptop environmentally friendly and easy to repair or upgrade.

In fact, Leighton expects users will do many of the repairs themselves using parts from doll houses, radio-control models, and lumberyards, rather than depending on vendors for parts. This expectation began at a time when Leighton needed a simple replacement on a prototype’s case and realized that plywood was more economical than “wast[ing] my time and perfectly good plastic to print out a damn rectangle. It wasn’t until afterwards that I retrospectively added the whole eco thing on top, but it’s pretty ironic that this was driven by my laziness and cost-cutting.”

To help make the case module self-repairable, the 3D-printed parts are broken down into 35 separate pieces so that damage to one piece does not require the reprinting of the entire case, as it does



Figure 3: Computer cards slide into the rest of the laptop, allowing quick changes of operating systems. They are housed in recycled PCMCIA cards.

with the pi-top [4], a Raspberry Pi-based laptop that also allows the printing of replacement parts.

Initially, the computer cards will be loaded with a free-licensed bootloader and provide either Parabola GNU/Linux-libre or Debian GNU/Linux as the operating system. Both will use a GNU Linux-libre [5] kernel, a version of the Linux kernel with all the proprietary blobs removed, an Allwinner A20 dual-core processor, 2GB of RAM, and 8GB of NAND flash.

Parabola is used in the hopes of the laptop receiving Respects Your Freedom certification [6] from the Free Software Foundation, since it is one of the few completely free distributions for ARM processors. Debian is likely to be more popular, but it is technically non-free because it allows easy access to non-free software. To make the two easily distinguishable, the Parabola version is called the Libre Tea Computer Card, and the Debian version is the Practically Perfect.

If produced, the EOMA68 laptop will be among the first laptops with a free license for everything from the bootloader to the operating system. It will be modular, allowing for easy upgrades and repairs; economical, because parts can be replaced rather than the entire computer; and ecologically responsible, using some recycled parts and giving much of the device a lifetime greater than the three to five years of comparable laptops.

Challenges and Definitions

The EOMA68 standard has the potential to redefine the definition of free software, linking it to environmentalism and privacy concerns in a way that the Free Software Foundation has attempted for more than a decade. Moreover, with even limited success, it could help ease the way for other free hardware devices.

The obstacles in its way, however, remain formidable. For one thing, its specifications are mediocre at best, which may be why, after two weeks of the fundraising campaign, it has reached less than 20 percent of its goal of \$150,000.

Another concern is whether Leighton's past efforts will handicap this current effort. The Vivaldi-Improv efforts failed to reach the required levels for manufacturing, which may have destroyed the proj-

ect's reputation. Moreover, these previous efforts were dealing with quantities of fewer than 2,500 units and could have typecast the project as a small-time effort that is unlikely to scale to anything like the massive orders made by major computer vendors.

Granted, Leighton is aware of these potential problems, noting that working with his past contacts again "will be hard." However, getting manufacturers to understand the requirements of free software or to take a second chance remain serious difficulties. Even if the EOMA68 laptop is brought to market, that is no guarantee of success, because winning over vendors and customers requires fighting many of the battles with manufacturers all over again.

Still, with luck, Leighton may be able to overcome these difficulties and move on to other products, such as a free tablet or – his own ambition – a Digital SLR camera. His success may depend on how many people can be convinced that the project is "at its heart an ethical initiative to give us our freedom of choice back at the hardware level" and not judge it by specifications alone.

The task is difficult. Yet, as Leighton himself observes, sitting back and complaining will only guarantee that nothing will change. If nothing else, efforts like his may help to expand the free software definition beyond its current limitations. ■■■

INFO

- [1] Free software definition: <https://www.gnu.org/philosophy/free-sw.html#translations>
- [2] EOMA68: <https://www.crowdsupply.com/eoma68/micro-desktop#products-top>
- [3] Open hardware setbacks: <http://www.linux-magazine.com/Online/Features/MakePlayLive-Releases-New-Single-Board-System>
- [4] pi-top: <https://lwn.net/Articles/667040/>
- [5] GNU Linux-libre: <http://www.fsfla.org/ikiwiki/selibre/linux-libre/index.en.html>
- [6] Respects Your Freedom certification: <https://www.fsf.org/resources/hw/endorsement/respects-your-freedom>

Shop the Shop

shop.linuxnewmedia.com

Want to subscribe?

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

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

shop.linuxnewmedia.com

DIGITAL & PRINT SUBSCRIPTIONS



SPECIAL EDITIONS



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.

GUADEC 2016

Date: August 12–14, 2016

Location: Karlsruhe, Germany

Website: <https://2016.guadec.org>

Gnome users, developers, foundation leaders, individuals, and others come together to meet collaborators from chat rooms and mailing lists, to network, to visit old friends and make new ones, and to have fun at events, workshops, BoFs, and hackfests.

LinuxCon Europe

Date: October 4–6, 2016

Location: Berlin, Germany

Website: <http://events.linuxfoundation.org/events/linuxcon-europe>

Learn from the leading maintainers, developers, and project leads in the Linux community from around the world. Choose from among 100+ sessions, ranging from tutorials to deep technical dives and everything in between.

OpenStack Summit

Date: October 25–28

Location: Barcelona, Spain

Website: <https://www.openstack.org/summit/>

Welcome to users, developers, and administrators of OpenStack cloud software. This five-day conference includes keynotes, breakout sessions, and working sessions. Learn about new products in the ecosystem and participate in hands-on workshops.

EVENTS

GUADEC 2016	August 12–14	Karlsruhe, Germany	https://2016.guadec.org/
IFA	September 2–7	Berlin, Germany	http://www.ifa-berlin.de/
ownCloud Contributor Conference 2016	September 9–15	Berlin, Germany	https://conference.owncloud.org/conference/oCC2016/
Node Interactive Europe	September 15–18	Amsterdam, Netherlands	http://events.linuxfoundation.org/events/node-interactive-europe
Kieler Open Source and Linux-Tage	September 16–17	Kiel, Germany	http://www.kilux.de/
LPI Partner Day 2016	September 21–22	Kassel, Germany	https://www.lpice.eu/cfp
systemd.conf2016	September 28–October 1	Berlin, Germany	https://conf.systemd.io/
LinuxCon Europe	October 4–6	Berlin, Germany	http://events.linuxfoundation.org/events/linuxcon-europe
ContainerCon Europe	October 4–6	Berlin, Germany	http://events.linuxfoundation.org/events/containercon-europe
Icinga Camp San Diego 2016	October 18	San Diego, California	https://www.icinga.org/community/events/icinga-camp-san-diego/
OpenStack Summit 2016	October 25–28	Barcelona, Spain	https://www.openstack.org/summit/
Node.js Interactive North America	November 29–December 2	Austin, Texas	http://events.linuxfoundation.org/events/node-interactive
Open Source Monitoring Conference	November 29–December 2	Nürnberg, Germany	https://www.netways.de/events_schulungen/osmc/uebersicht/
IT-Tag 2016	December 12–15	Frankfurt, Germany	http://www.ittag.informatik-aktuell.de/

CALL FOR PAPERS

We are always looking for good articles on Linux and the tools of the Linux environment. Although we will consider any topic, the following themes are of special interest:

- System administration
- Useful tips and tools
- Security, both news and techniques
- Product reviews, especially from real-world experience
- Community news and projects

If you have an idea, send a proposal with an outline, an estimate of the length, a description of your background, and contact information to edit@linux-magazine.com.



The technical level of the article should be consistent with what you normally read in *Linux Magazine*. Remember that *Linux Magazine* is read in many countries, and your article may be translated into one of our sister publications. Therefore, it is best to avoid using slang and idioms that might not be understood by all readers.

Be careful when referring to dates or events in the future. Many weeks could pass between your manuscript submission and the final copy reaching the reader's hands. When submitting proposals or manuscripts, please use a subject line in your email message that helps us identify your message as an article proposal. Screenshots and other supporting materials are always welcome.

Additional information is available at:

http://www.linux-magazine.com/contact/write_for_us.

AUTHORS

Erik Bärwaldt	12, 44
Swapnil Bhartiya	8
Chris Binnie	30
Bruce Byfield	70, 88, 92
Joe Casad	3
Marco Fioretti	24
Jon "maddog" Hall	86
Klaus Knopper	50
Charly Kühnast	60
Christoph Langner	66
Martin Loschwitz	38
Dmitri Popov	74
Mike Schilli	62
Daniel Stender	54
Ferdinand Thommes	80
Uwe Vollbracht	20

CONTACT INFO

Editor in Chief

Joe Casad, jcasad@linux-magazine.com

Managing Editor

Rita L Sooby, rsooby@linux-magazine.com

Localization & Translation

Ian Travis

News Editor

Swapnil Bhartiya

Copy Editors

Amber Ankerholz, Amy Pettle

Layout

Dena Friesen, Lori White

Cover Design

Lori White

Cover Image

© nomadsoul1 and Maxim Kazmin, 123RF.com

Advertising – North America

Ann Jesse, ajesse@linuxnewmedia.com
phone +1 785 841 8834

Advertising – Europe

Penny Wilby, pwilby@sparkhausmedia.com
phone +44 1807 211100

Publisher

Brian Osborn, bosborn@linuxnewmedia.com

Marketing Communications

Gwen Clark, clark@linuxnewmedia.com
Linux New Media USA, LLC
616 Kentucky St.
Lawrence, KS 66044 USA

Customer Service / Subscription

For USA and Canada:
Email: cs@linuxpromagazine.com
Phone: 1-866-247-2802
(Toll Free from the US and Canada)
Fax: 1-785-856-3084

For all other countries:

Email: subs@linux-magazine.com

Phone: +49 89 99 34 1167

Fax: +49 89 99 34 1198

www.linuxpromagazine.com – North America

www.linux-magazine.com – Worldwide

While every care has been taken in the content of the magazine, the publishers cannot be held responsible for the accuracy of the information contained within it or any consequences arising from the use of it. The use of the disc provided with the magazine or any material provided on it is at your own risk.

Copyright and Trademarks © 2016 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 Germany

Distributed by COMAG Specialist, Tavistock Road, West Drayton, Middlesex, UB7 7QE, United Kingdom

Published in Europe by: Sparkhaus Media GmbH, Putzbrunner Str. 71, 81749 Munich, Germany.

Issue 191 / October 2016

Samba

Approximate
UK / Europe Sep 05
USA / Canada Sep 30
Australia Oct 31
On Sale Date

The Samba file service is a popular option for networks supporting Windows, Linux, and Mac clients. Next month, we look inside Samba, study some Samba benchmarks, and examine Samba in clustered environments.

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

DON'T MISS A SINGLE ISSUE!

Ubuntu User is the first print magazine created specifically for Ubuntu users. Ease into Ubuntu with the helpful Discovery Guide included in each issue, or advance your skills with in-depth technical articles, HOW-TOs, reviews, tutorials, and community reports.

SUBSCRIBE NOW!

4 issues per year for only
£ 24.90 / EUR 29.90 / US\$ 39.95

- ✓ Don't miss a single issue!
- ✓ Huge savings – Save more than 35% off the cover price!
- ✓ Free DVD – Each issue includes a Free DVD!



shop.linuxnewmedia.com



Lightning Fast

All Flash NVMe

S E R V E R S O L U T I O N S



1U 10 NVMe
SYS-1028U-TN10RT+



2U 48 NVMe
SSG-2028R-NR48N



2U 24 NVMe
SYS-2028U-TN24R4T+

Hot-Swap NVMe

New Generation SSD All-Flash Drive Technology

- Intel® Xeon® processor E5-2600 v4 Product Family Supported
- Highest Performance up to 3GB/s Read and 750K IOPS Per drive
- Satisfies the Performance Demands of VDI, Cloud, Storage, and Server Caching
- 2.5" U.2 (SFF-8639) Form Factor serviceability advantage vs. standard Add-On Cards
- Hot-swap cartridge design protects against unintentional removal



Intel Inside®. Powerful Productivity Outside.



Learn More at www.supermicro.com

©Super Micro Computer, Inc. Specifications subject to change without notice. Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its US-based other countries.