

FREE  
DVD

Linux Mint 17.3  
"Rosa" Mate Edition 32-bit

& Clonezilla  
64-bit

**FILESYSTEMS**

Btrfs, ZFS, and the story of ext

LINUX  
PRO  
MAGAZINE



# LINUX

MAGAZINE

AUGUST 2016

# FILESYSTEMS

Choose the best Linux filesystem for your data

**The Long Strange History of Ext**

**Relax-and-Recover**  
Create a Linux rescue disk

**Matrix**  
Videoconferencing  
the FOSS way

**7 Remote desktop applications**



**Maru OS**  
Debian desktop  
on a smartphone

**Yabs**  
Bill clients and keep  
track of your finances

**Firejail**  
Is this the world's  
easiest sandbox?

**Snapper**  
Manage your  
Btrfs snapshots



Issue 189	US\$ 12.99
August 2016	CAN\$ 13.99
0 74470 58049 2 08	

**HETZNER**  
ONLINE

**Dedicated Servers**  
**for all users.**

★ ★ ★  
ATTRACTING  
LOYAL CUSTOMERS  
OF ALL TYPES

Since 1997

**GAMER  
ON  
BOARD**

**I ♥  
SOFTWARE  
DEVELOPERS**

**I ♥  
PIXELS**



**MANAGER**



**NERD**

**WORLD'S BEST  
ADMIN**

**I ♥ My  
Ceo**



## Dedicated Root Server PX61-SSD

Intel® Xeon® E3-1275 v5  
Quad-Core Skylake Processor  
64 GB DDR4 ECC RAM  
2 x 480 GB SATA SSD  
Guaranteed 1 Gbit/s bandwidth  
100 GB Backup Space  
30 TB traffic inclusive\*  
No minimum contract  
Setup Fee \$136

monthly \$ **68**

## 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 \$136

monthly \$ **68**

\* 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 \$1.35 per additional TB used.

[www.hetzner.de/us](http://www.hetzner.de/us)

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.



# THE GREAT DEBATE

Dear Linux Pro Reader,

The history of high tech in politics writes another chapter this season, with Twitter taking on new importance. Barack Obama pioneered the use of Twitter as a campaign vehicle in his 2008 and 2012 campaigns, carefully coordinating his tweets to complement a comprehensive strategy of voter outreach.

This time, candidate Donald Trump appears to be using Twitter even more than Obama did. In fact, one could argue that Twitter is Trump's primary mode of communication – every day a few more Trump tweets appear in the news headlines. With all those tweets, you might be thinking that Trump spends even more time and energy studying and coordinating what he says, but judging from the content of some of his messages, he doesn't appear to be passing the contents by any kind of professional politician or handler. He really seems to be tweeting whatever happens into his head. Is this a missed opportunity for more careful coordination? Maybe not. Twitter thrives on high velocity. Trump just keeps the tweets coming, and therefore, keeps himself constantly in the news. Even if he tweets something that no experienced political manager would ever let him say, it doesn't matter, because, before anyone can even start a controversy over it, he has already moved on and tweeted something else.

What is truly more interesting than the rise of Twitter is the apparent demise of email. Barack Obama used to carry his BlackBerry around wherever he went to send instructions and receive reports in a steady stream of email messages. If one report is correct, Trump doesn't really email at all but tends to scrawl handwritten notes on paper and pass them on to assistants to email. In a recent interview, fellow billionaire and reality star Mark Cuban recounts a past dust-up with the Donald: "He doesn't email, right, so I got an email from his assistant that is scanned from a piece of paper with his comments on it. All it said was 'what happened?'" [1].

Hillary Clinton certainly has her own email issues, having been the subject of a lengthy investigation over which email server she used to send which message when she was Secretary of State.

Email is altogether too old fashioned for this year's news – too easy to spoof, too easy to hack. Too many private email messages have ended up on too many blog sites. If I were a presidential candidate, I would be very careful about putting too much strategic information in an email message.

Email is good for communicating details: specific instructions, bulleted lists, detailed commentary, counter-commentary on previous commentary. What is Twitter good for communicating?

- Emotions
- Platitudes
- Links to other stuff on the Internet

The tiny character limit does not let you actually develop any kind of argument or background evidence for whatever it is you are claiming. Better to tag names and things with emotionally charged adjectives, so as to encapsulate a system of thought, or a person's life, in a single word. Thus Trump does not refer to his opponent as "Hillary Clinton" but as "crooked Hillary Clinton" and to his senate nemesis as "goofy Elizabeth Warren."

Senator Warren is quite ready to play the Twitter game herself, taking to the medium more than once to call Trump a "bully" and use terms like "stupid" and "reckless."

It has occurred to me that the 140-character limit for a Twitter message is pretty similar to about how much text someone could shout at you when they drive past you in a car traveling at 30 miles per hour.

American political discourse has thus evolved from the 1858 Lincoln-Douglas debates [2], in which each candidate could speak for 90 minutes in an unmoderated forum to elucidate their positions on the issues facing the nation, to our new paradigm, which is the equivalent of people driving around in cars shouting two-sentence insults about each other and circling back several times a day with new insults or new ways to formulate the old insults.

Interestingly, the candidates didn't have cars back in the days of the Lincoln-Douglas debates, or even microphones with public address systems to amplify their oratory. Not only do we have cars and microphones, we also have the Internet, as well as cell phones, social networks, hairless cats, and sonic toothbrushes. We are sooo much more advanced.

Joe

Joe Casad,  
Editor in Chief



## INFO

[1] Mark Cuban interview: <http://www.msnbc.com/all-in/watch/mark-cuban-on-his-relationship-with-trump-703460419620>

[2] Lincoln-Douglas debates: [https://en.wikipedia.org/wiki/Lincoln%E2%80%93Douglas\\_debates](https://en.wikipedia.org/wiki/Lincoln%E2%80%93Douglas_debates)



# LINUX PRO MAGAZINE

## news

### 8 GOOGLE LAWSUIT

- CoreOS announces distributed storage system
- Google beats Oracle in the Android lawsuit

### 9 LINUX IS LARGEST EVER SOFTWARE DEV PROJECT

- Linux is the largest development project in history
- Linksys says Yes to open source firmware
- TeamViewer hit with DoS attack
- More online

### 10 INDUSTRIAL MALWARE

- New Stuxnet-like malware found in the wild
- Russian FindFace app is a privacy nightmare
- 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.

# Filesystems

In Linux, everything is a file, but different filesystems handle metadata, optimization, and other functions differently. We look at three popular filesystems.

**12 Btrfs** Does this powerful new filesystem represent the future of Linux?

**16 ZFS** ZFS is a first-class filesystem for big iron, but for various reasons, it is still awaiting widespread Linux adoption.

**22 Ext2/3/4** A brief tour of ext filesystem history will give you some insight into how this 25-year-old Linux filesystem works – and how it has evolved to meet users' needs.

## REVIEWS

**28 Tool Tips**  
Tool tests on the fast track.

**32 Matrix**  
An open standard for decentralized communication.

**36 Remote Administration**  
Remote desktop applications are useful assistants in the personal and professional spheres. We look at a few candidates.

## Community Notebook

**90 Krita 3.0**  
The free sketching and painting program has new features and improvements in version 3.0.

**94 Doghouse – Open Source Endurance**  
A tale of long-lived Free and Open Source projects.



## HIGHLIGHTS

**12 BTRFS**  
A filesystem with advanced features for data centers and desktops.



**32 MATRIX**  
A decentralized communication tool with a choice of clients and servers.

**48 MARU OS**  
Carry Debian on your smartphone, so you can leave your laptop at home.

**66 YABS**  
This accounting software for small businesses simplifies administrative tasks.

## FEATURES

**42 Snapper**  
Restore Btrfs snapshots to return a system to a previous state.



**52 Relax-and-Recover**  
A simple Bash script generates rescue media from running systems, takes care of backups, and migrates computers to new hardware.

**48 Maru OS**  
Dump the laptop, because Maru OS provides a desktop Debian on a smartphone.

**58 Charly – Etckeeper**  
Restore configuration files to a state from better days.

**56 Ask Klaus!**  
Klaus Knopper answers questions about auto-login, Gnome Commander, and Live discs.

**60 Perl – IFTTT Home Automation**  
Switch power outlets on and off with home automation tools.


## LINUXUSER

**66 Yabs**  
Keep track of your SMB finances easily.



**70 Command Line – Firejail**  
Sandboxing is as easy as typing eight letters before a command.

**76 Workspace – Document Management**  
Transform piles of paper into an organized and searchable digital library with Scan Tailor and Paperwork.



**84 ICgrep**  
A modern, parallel, Unicode-enabled alternative to grep.



**Linux Mint 17.3**  
"Rosa" Mate Edition 32-bit  
&  
**Clonezilla**  
64-bit

SEE PAGE 6 FOR DETAILS

# On the DVD

 Linux Mint 17.3  
"Rosa" Mate Edition 32-bit

ISSUE 189

LINUX

clonezilla  
64-bit

ISSUE 189 AUG 2016

LINUX  
MAGAZINE

Although this Linux Magazine disc has been tested and  
in the best of our knowledge free of viruses and other  
malicious code, Magazine cannot be held responsible for  
any damage to your system or data that may occur  
as a result of using this disc.

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

## Linux Mint 17.3 "Rosa" Mate (32-bit)

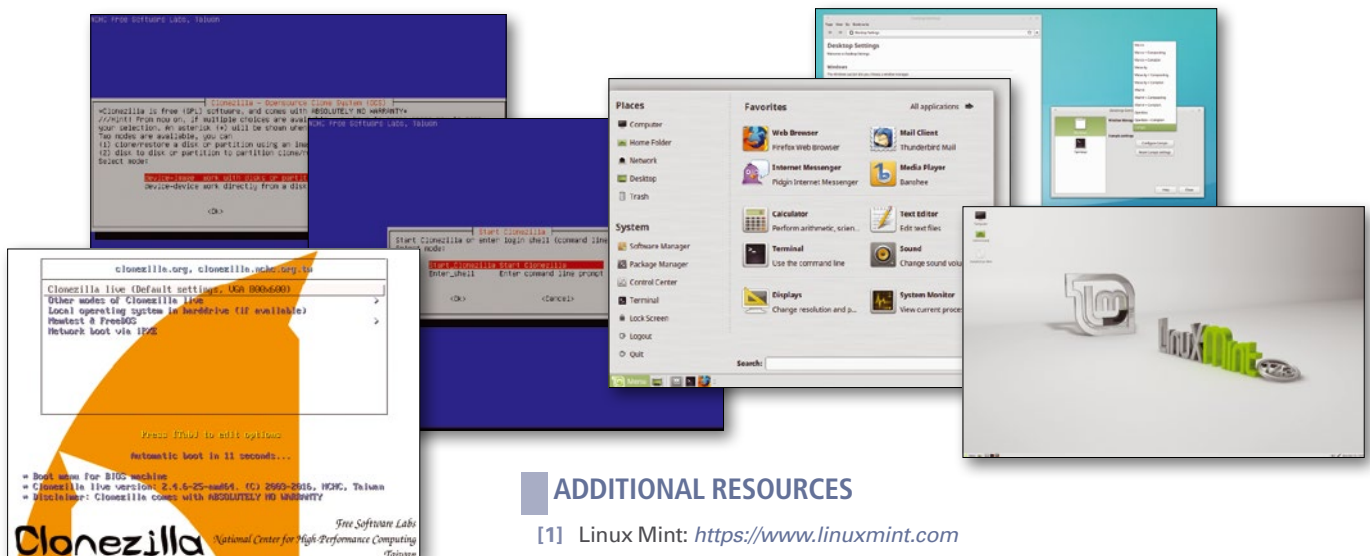
This long-term support version of Linux Mint, built on Debian and Ubuntu, features the Mate v1.12 desktop. The "Rosa" Desktop Settings offer multiple window managers, including Marco, Metacity, Xfwm4, Openbox, and Compiz. The Compton compositing manager is configured to prevent screen tearing. Two new commands show information about which window/compositing managers are running (`wm-detect`) and allow you to recover the default window manager (`wm-recovery`).

- Improved touchpad support
- More checks in Update Manager
- Improved multimonitor support
- Updated Xorg, Mesa, and kernel
- Improved HiDPI support

## Clonezilla Live 2.4.6-25 (64-bit)

The Clonezilla partition and disk imaging/cloning Live distro is suitable for system deployment, bare metal backup, and recovery of systems. It supports a multitude of 32- and 64-bit filesystems (e.g., ext, ReiserFS, XFS, JFS, Btrfs, FAT, NTFS, HFS+, UFS, MINIX, VMFS), facilitating Linux, Windows, Mac OS, BSD, and VMware cloning, with only used blocks saved and restored. Sector-to-sector copying supports cloning of other filesystems.

- LVM2 support
- MBR and GPT partitions
- BIOS or UEFI boot
- Unattended mode with numerous options
- One to many restore



## ADDITIONAL RESOURCES

- [1] Linux Mint: <https://www.linuxmint.com>
- [2] New features in Mate: [https://www.linuxmint.com/rel\\_rosa\\_mate\\_whatsnew.php](https://www.linuxmint.com/rel_rosa_mate_whatsnew.php)
- [3] Clonezilla: <http://www.clonezilla.org>
- [4] Clonezilla enhancements and changes: <https://sourceforge.net/p/clonezilla/news/>

Defective discs will be replaced. Please send an email to [cs@linuxpromagazine.com](mailto:cs@linuxpromagazine.com).





# RISE HIGHER

EACH ISSUE OF DRUPAL WATCHDOG OFFERS TOOLS, TIPS, AND BEST PRACTICES FOR BETTER DRUPAL WEBSITES.



Renew or subscribe now!

SUBSCRIPTIONS NOW AVAILABLE WORLDWIDE!

Visit <http://drupalwatchdog.com/subscribe>

# NEWS

Updates on technologies, trends, and tools

## THIS MONTH'S NEWS

### 08 CoreOS Announces Torus

- CoreOS Announces Distributed Storage System
- Google Beats Oracle in the Android Lawsuit

### 09 Linux Is Largest Development Project

- Linux Is The Largest Development Project in History
- Linksys Says Yes to Open Source Firmware
- TeamViewer Hit with DoS Attack
- More Online

### 10 Stuxnet-Like Malware Found in the Wild

- New Stuxnet-Like Malware Found in the Wild
- Russian FindFace App Is a Privacy Nightmare
- More Online

## CoreOS Announces Distributed Storage System

CoreOS has announced a new distributed storage tool called Torus. Torus is a scalable distributed storage system for application containers and cluster orchestration platforms like Kubernetes.

According to CoreOS, many of the available open source storage solutions are not designed for the GIFEE (Google Infrastructure for Everyone Else) approach, which creates massive clusters of inexpensive and small machines. Commercial solutions offer GIFEE functionality, but they can be very expensive and difficult to integrate with other systems.

The goal of the Torus project is to build GIFEE into an open source storage solution. Torus uses the etcd distributed key-value database to store and retrieve file or object metadata. etcd provides a solid, battle-tested base for core distributed systems operations that must execute rapidly and reliably. Torus acts as a building block and enables various types of storage, including distributed block devices or large-object storage.

Torus is written in Go and available on GitHub.

## Google Beats Oracle in the Android Lawsuit

A jury unanimously found that Google did not violate Oracle's copyright in its Android Operating system. The jury found that, while Google did use Oracle's Java APIs, they were used under fair use doctrine, which means Google didn't need Oracle's permission. As a result, Oracle is not entitled to any damages. The decision brings another end to the \$9 billion lawsuit Oracle filed against Google. Oracle is likely to appeal, but once a jury decides on the facts, an appeals court has limited options for overturning the decision.

The courtroom battle between Google and Oracle has gone on for several years. Google achieved an earlier victory when Judge William Alsup decided the API was

not copyrightable, but Oracle went to the Federal Circuit court, which decided that APIs are subject to copyright. The case then went back to a lower court to decide whether Google's use of the APIs was fair.





Although Google has won this particular case, it's not a victory for developers. EFF's Parker Higgins wrote in a blog post, "While developers of interoperable software can take some comfort in the fact that reimplementations may be fair use, a simpler and fairer solution would simply have been to recognize API labels as a system or method of operation not restricted by copyright."

## Linux is the Largest Software Development Project in History

Greg Kroah-Hartman, a leading Linux kernel developer, delivered a keynote at CoreOS Fest on May 10-11 in Berlin, Germany. Kroah-Hartman shared some impressive numbers about Linux kernel development. He noted that in 2015, more than 440 companies were involved with the kernel and more than 4,000 people contributed to it. The latest release of Linux (4.5) has more than 21 million lines of code. Every single day, more than 10,800 lines of code are added, 5,300 lines of code are removed, and 1,875 lines of code are modified. The current pace of development boils down to more than eight changes per second.

Kroah-Hartman said, "It's the largest software development project ever, in the history of computing – by the number of people using it, developing it, and now using it, and the number of companies involved. It's a huge number of people." But all of these changes are for a reason: They bring security fixes and introduce new features. Kroah-Hartman urged companies and communities using Linux to keep up with the changes. He praised the updating mechanisms used by projects such as ChromeOS and CoreOS that keep the system up-to-date without user intervention.

Kroah-Hartman said that vendors should provide a system for regular updates and encourage users to take advantage of it. He sent a very strong message to computer vendors saying: "... your machine is insecure unless you're running my kernel, or based on my kernel, or based on another [recent] one. If you're not taking these fixes, then it is insecure."



## Linksys Says Yes to Open Source Firmware

The networking hardware vendor Linksys says it will allow users to install open source firmware on some of its devices. Beginning in June 2, 2016, the FCC (U.S. Federal Communication Commission) will require that device manufacturers limit the user's access to wireless router configuration settings. The primary reason behind this requirement is to ensure that users do not operate these routers beyond their licensed radio frequencies, which interferes with other services, such as FAA Doppler weather radar systems. In response to the FCC's requirements, most vendors decided the easiest path was just to ban open source firmware on their devices.

However, Linksys chose to comply with the FCC while still giving users the ability to install custom firmware. Linksys is working with chip maker Marvell and OpenWrt developers to find a midway point between the FCC requirements and the user's rights. The company plans to separate the RF wireless data from the firmware, so even if a user installs OpenWrt, it will not allow the use of non-licensed radio frequencies. Linksys will allow custom firmware only on WRT routers; the rest of its routers will block custom firmware.

## TeamViewer Hit with DoS Attack

The website for the TeamViewer remote control/remote access service came under a denial of service (DoS) attack on June 1 that caused a massive service outage. The at-

## MORE ONLINE

### Linux Pro Magazine

[www.linuxpromagazine.com](http://www.linuxpromagazine.com)

#### Off the Beat • Bruce Byfield

##### Why Licensing Limits Ebooks

One year in which traditional books outsells ebooks, and suddenly headlines are announcing, "Books are back," citing a small increase in book sales and a small decline in ebook sales.

#### Exploring Desktop mode on the Aquaris M10 Ubuntu Edition

I am writing this article on an Aquaris M10 Ubuntu Edition tablet, using LibreOffice with a mouse and keyboard.

#### Waiting for LibreOffice Table Styles

Like most word processors, LibreOffice Writer includes character and paragraph styles. However, Writer draws much of its power as a mid-level desktop publisher from the addition of frame, list, and page styles.

#### Paw Prints • Jon "maddog" Hall

##### A Nice Place, A Warm Summer's Day, and Lua

Last year I was in Rio de Janeiro, and a couple of friends and I were just finishing a delightful day in the botanical gardens...

#### I Give My Heart to You

Hello, my friends, please forgive me for using this method of telling you what has happened to me this past week, but since I have been writing about Free and Open Source Software and Culture in this blog for many years, I hope to reach most of those people who care about my writing and actions through this medium.

#### Productivity Sauce • Dmitri Popov

##### Burn Image Files with Style Using Etcher

Many Linux distributions nowadays are distributed as ISO or IMG images, and you need a specialized tool to burn them onto SD cards or USB drives.

#### Block Ads and Malware Sites with a Unified Host File

Instead of blocking ads and malware sites using a specialized browser extension, you can enable the blocking at system level.

#### Get Weather Forecast from the Command Line with wtr.in

Of all the weather forecast applications and tools I have tried and reviewed over the years, wtr.in is probably the most ingenious and useful one.

## MORE ONLINE

## ADMIN HPC

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

## Ceph Jewel • Martin Loschwitz

When Red Hat took over Inktank Storage in 2014, thus gobbling up the Ceph object store, a murmur was heard throughout the community: on the one hand, because Red Hat already had a direct competitor to Ceph in its portfolio in the form of GlusterFS, and on the other because Inktank was a such a young company – for which Red Hat had laid out a large sum of cash.

## A Container for HPC • Jeff Layton

Containers have become an important part of the IT industry because (1) they are more efficient than full (hardware-level) virtualization and (2) the container workflow readily supports DevOps.

## ADMIN Online

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

## Ease Your Network Inventory Pain with Spiceworks • Chris Binnie

Network device inventory is one of those common pain points that affects every system administrator who has ever tried to script or invent a custom solution.

## What's New in Ansible 2.0

## Konstantin Agouros

According to Ansible's changelog, the name "Ansible" comes from Ursula K. Le Guin's 1966 novel *Rocannon's World*. Her communicator of the same name operates at greater than the speed of light, without delay, with any place in the universe.

## Manage Resources with AWS CloudFormation • Oliver Arafat

The cloud has changed where and how companies invest their IT budgets. Instead of spending money on data centers and servers, without really knowing how they will be used, the customer only pays for the computing resources needed.

## Security First with the Hiawatha Web Server Hans-Cees Speel

Security on the Internet is vital. The Hiawatha web server, created by Hugo Lensink, is a small (and free) web server that subscribes to the principle of "security by default."

tack came just two weeks after many users complained that their accounts were compromised, alleging that TeamViewer had been hacked.

TeamViewer has denied any such hack and has blamed users for carelessness, stating, "TeamViewer is appalled by any criminal activity; however, the source of the problem, according to our research, is careless use, not a potential security breach on TeamViewer's side." The company said that they were not hacked nor was there any security hole. TeamViewer suggested that users should take some precautions to protect themselves.

Responding to the reports of the DoS attack, TeamViewer did admit an attack on its infrastructure. But they refused any connection with the alleged reports of hacking. The company said in a statement, "Some online media outlets falsely linked the incident with past claims by users that their accounts have been hacked and theories about would-be security breaches at TeamViewer. We have no evidence that these issues are related."

If you are a TeamViewer user, you should follow these best practices to keep your account safe.

## New Stuxnet-Like Malware Found in the Wild

Researchers at the security firm FireEye have found mysterious malware, named Irongate, which is designed to target industrial processes, specifically, ICS/SCADA equipment that is manufactured by Siemens.

Irongate masks its malicious activities by replacing a Dynamic Link Library (DLL) with a malicious DLL. The DLL works as a broker between a programmable logic controller (PLC) and the monitoring software. It records five seconds of normal traffic and then replays it – the way Keanu Reeves loops tape in the movie *Speed*. This strategy allows attackers to hide from process operators.

The malware has the same attack traits that were seen in the infamous Stuxnet, which was allegedly created by Israel and the US to sabotage Iran's nuclear program.

FireEye said in a blog post, "Our analysis finds that IRONGATE invokes ICS attack concepts first seen in Stuxnet, but in a simulation environment. Because the body of industrial control systems (ICS) and supervisory control and data acquisition (SCADA) malware is limited, we are sharing details with the broader community."



©Natalia Lukiyanova, 123RF

## Russian FindFace App is a Privacy Nightmare

The machine-learning capability that social networks are gaining is becoming a privacy nightmare. Anyone can take your photo then find everything about you via social networks. That's exactly what the Russian site FindFace is doing.

FindFace is bringing a very powerful facial recognition technology to VK, formerly known as VKontakte, the "Facebook of Russia."


What worries privacy advocates is that the site has a high accuracy rate when it comes to identifying a total stranger simply by snapping their picture and uploading to the FindFace site.

Antivirus firm Kaspersky wrote in a blog post, "If you upload ideal photos, that were taken when your target was posing, everything works just great. The program has successfully found 9 of 10 test 'victims' in the office."

The algorithm powering the facial recognition is developed by Russia-based N-Tech. Lab, which beat Google's face recognition software at the MegaFace challenge.



# 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

**6th Edition!**

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



**FREE DVD INSIDE!**

The world's greatest Live distro!

**KNOPPIX 7.4.2**

ORDER ONLINE:

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



### Btrfs and the future of the filesystem

# New Butter

The Btrfs filesystem offers advanced features such as RAID, subvolumes, snapshots, checksums, and transparent compression, but do desktop users really need all that power? *By Ferdinand Thommes*

**A** filesystem functions below the operating system, ensuring that abstract data is converted into physical address attributes such as tracks and sectors. Some filesystems go beyond this basic functionality. One powerful and popular filesystem for Linux is Btrfs. The Btrfs filesystem [1], which is affectionately pronounced “ButterFS,” is sometimes called the *next generation* filesystem. Btrfs is a copy-on-write filesystem [2] originally developed by Oracle Corporation and masterminded by Chris Mason. In some ways, Btrfs is best understood as an implementation of the Solaris 10 transactional filesystem ZFS [3] for the Linux platform. Oracle acquired ZFS in 2010 when it acquired Sun Microsystems. Btrfs is free software under the GPL and was adopted into Linux kernel 2.6.29 early in 2009.

Btrfs was declared suitable for production use in April 2013. Btrfs developer Chris Mason moved from SUSE, where he worked on ReiserFS, to Oracle, and he

## COVER STORIES

**BTRFS** 12  
Is this innovative Linux filesystem ready for the real world?

**ZFS** 16  
Sun’s souped-up filesystem offers many advantages for Linux systems, if you can work around the licensing issues.

**EXT2/3/4** 22  
The ext filesystem is still very popular, but can it hold its own against the powerful newcomers?

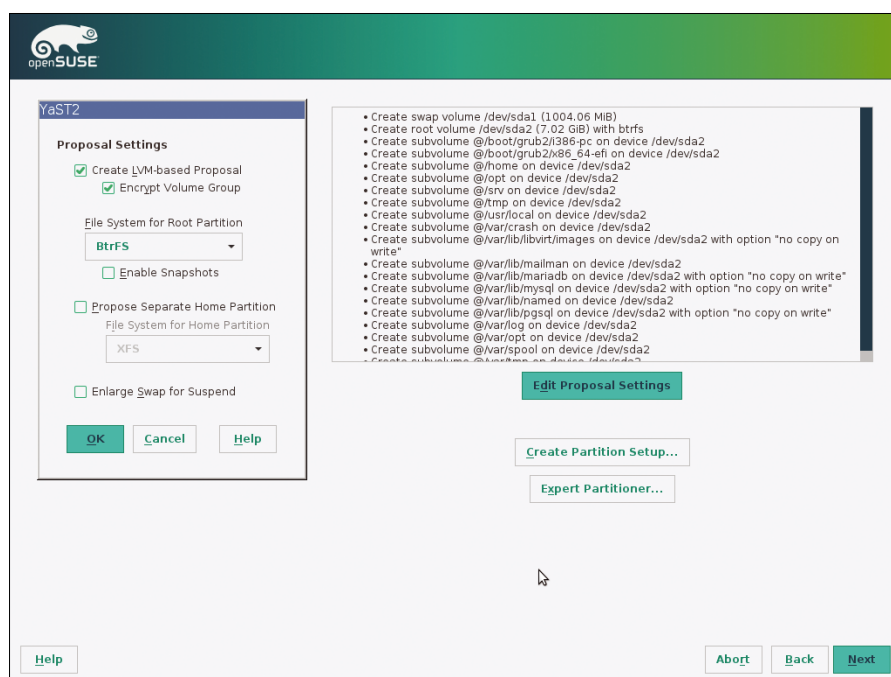


Figure 1: Btrfs as the default subvolume when installing openSUSE.





plains the origin of the name Btrfs, which expands to “B-tree FS.” Another feature in the filesystem structure is pseudo-subdirectories, or subvolumes, that act like separate drives but only store changes from the main volume and thus do not occupy a great deal of space. Additionally, Btrfs can never run out of inodes [6], which is possible with the ext family, even if free space is available on the partition.

## Differences and USPs

In contrast to the currently unofficial default Linux filesystem ext4, Btrfs offers some features that are generally not attributed to the functionality of a filesystem but are hugely popular, especially in professional environments such as data centers. Useful features include built-in Logical Volume Management functionality and built-in RAID levels 0, 1, 5, 6, and 10. Compared with traditional hardware or software RAID, the advantage to Btrfs is that it can distinguish between used and unused data blocks and thus save much time on recovery. Moreover, Btrfs has mastered the art of converting RAID levels on the fly.

Btrfs offers optional transparent data compression, deduplication [7], checksums for data and metadata, filesystem checks and defragmentation at run time, and optimizations for solid-state drives (SSDs). Additionally, `btrfs-convert` can reversibly convert existing ext3 and ext4 filesystems to Btrfs, and reverting to ext4 is possible as long as the snapshot created when converting the metadata of the ext4 system is not deleted. The Btrfs filesystem encryption feature, although planned from the outset, has not yet been implemented.

## Snapshots Undo Mistakes

One outstanding feature for desktop users is the ability to create snapshots on a running system (Figure 2). Snapshots are images of a subvolume frozen in time. Particularly in connection with

has worked at Facebook for several years, where Btrfs is widely used in the back end. Btrfs is

now no longer limited to Linux; the WinBtrfs [4] project offers what are still experimental drivers for Windows. In the Linux world, Oracle started using Btrfs in its *Unbreakable Linux* release, version 2, four years ago, and SUSE in SUSE Linux Enterprise Server 12 (SLES 12) and openSUSE 13.2 (Figure 1). OpeSUSE Leap uses Btrfs as the default for the root partition; most other distributions include Btrfs in their archives and offer it as an alternative in the installer. Fedora plans to make Btrfs the default with Fedora 24.

## Trees, Nodes, and Leaves

Btrfs is not just well equipped for future desktops, it is also well suited for data centers. A maximum filesystem size of 16 exbibytes (EiB) corresponds to 16 million terabytes, or just over a trillion bytes. (Compare this with ext4, which can only handle 1EiB.) The filesystem can consist of one file of that size, or you can store 18 trillion files, and filenames can comprise 255 bytes. Whereas the directories in ext4 are organized in a table as an HTree, Btrfs uses a B-tree [5], with metadata in the nodes and the data in the leaves.

The idea is to reduce access times for file elements by reducing the height of the tree structure and expanding horizontally. The use of a B-tree also ex-

```

dd@linux-k8oj:~$ df -h
Filesystem                Size      Used Avail Use% Mounted on
devtmpfs                   994M         0  994M   0% /dev
tmpfs                      1001M    232K 1001M   1% /dev/shm
tmpfs                       1001M     2.0M  999M   1% /run
tmpfs                       1001M         0 1001M   0% /sys/fs/cgroup
/dev/mapper/system-root    5.7G     5.1G  292M  95% /
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/opt
/dev/mapper/system-root    5.7G     5.1G  292M  95% /tmp
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/log
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/tmp
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/lib/pgsql
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/lib/mysql
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/lib/named
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/lib/mariadb
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/spool
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/lib/mailman
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/lib/libvirt/images
/dev/mapper/system-root    5.7G     5.1G  292M  95% /var/crash
/dev/mapper/system-root    5.7G     5.1G  292M  95% /usr/local
/dev/mapper/system-root    5.7G     5.1G  292M  95% /srv
/dev/mapper/system-root    5.7G     5.1G  292M  95% /opt
/dev/mapper/system-root    5.7G     5.1G  292M  95% /home
/dev/sda1                   379M     35M  321M  10% /boot
dd@linux-k8oj:~$

```

Figure 2: Activating a snapshot during installation.

the Snapper [8] tool developed by openSUSE (Figure 3), snapshots allow rolling back a system with problems to its previous state, with relatively little effort. Once Snapper is configured, the ability to create a snapshot before a system upgrade removes the worries about upgrades for many users. In SUSE, snapshots are handled graphically via a YaST module for Snapper; external GUIs exist for other operating systems. Btrfs snapshots are less limited than with LVM, because you do not need to allocate space first and run the risk of an overflow. Even a large number of snapshots will not slow the system down.

## Internals

Under the hood, Btrfs emulates its role model ZFS and relies on a technique called copy-on-write (CoW). The idea of CoW is to avoid unnecessary copies and copying actions. Therefore, changed blocks are only written physically if they are changed by one of the entities. For filesystems, CoW means that changed blocks are first written to free space without overwriting the original data; only the metadata is modified accordingly. CoW thus supports the creation of snapshots, because the original blocks are retained and form the snapshot. The metadata of the snapshot is stored with the first snapshot. In other snapshots, only changes to data blocks from the previous snapshot are incrementally stored.

Another advantage of the transaction-based CoW method is that the filesystem is always consistent, because changes to data do not overwrite the existing data but cause a block rebuild. This approach offers wear benefits on conventional hard drives with moving heads, because the mechanical system is not overtaxed; the write heads do not have to look for old data to overwrite but can simply write changes to a free part of the disk. This feature increases the write speed with less complexity than journaling-based filesystems like ReiserFS, ext3, and ext4.

## Cards on the Table

The opinions on whether a filesystem needs the enormous functionality of Btrfs are as diverse as the views of its current production capability. However, many different application scenarios are implemented on Linux, and the kernel supports a number of different filesystems.

Whether or not Btrfs will inherit ext4's crown is currently unclear, because Btrfs's extensive feature set is widely regarded as overkill for normal desktop systems, and because the very useful snapshot function is also available on ext4, Btrfs currently serves as more of an alternative for specific use cases. Theodore T'so, the developer of the ext family, sees Btrfs as a suc-

cessor to ext4 that is pushing its development limits. In addition to enthusiastic desktop users, the target group for Btrfs includes data centers, cloud computing environments, and large-scale storage, such as the massive storage required by developer Mason's current employer, Facebook.

## Yes, No, But ...

Btrfs requires some training time, which is not the case with conventional filesystems. Users also need to be confident about working at the terminal, because Btrfs does not integrate well with graphical file managers.

The undeniable advantage that Btrfs offers is currently bought at the price of additional coordination and configuration time for ordinary desktop users. One possible exception is openSUSE, which uses Btrfs as the default and thus offers good support for the system. ■■■

## AUTHOR

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

## INFO

- [1] Btrfs: [https://btrfs.wiki.kernel.org/index.php/Main\\_Page](https://btrfs.wiki.kernel.org/index.php/Main_Page)
- [2] CoW: <https://en.wikipedia.org/wiki/Copy-On-Write>
- [3] OpenZFS: [http://open-zfs.org/wiki/Main\\_Page](http://open-zfs.org/wiki/Main_Page)
- [4] WinBtrfs: <https://github.com/maharmstone/btrfs>
- [5] B+ tree: [https://en.wikipedia.org/wiki/B%2B\\_tree](https://en.wikipedia.org/wiki/B%2B_tree)
- [6] inode: <https://en.wikipedia.org/wiki/Inode>
- [7] Deduplication: [https://en.wikipedia.org/wiki/Data\\_deduplication](https://en.wikipedia.org/wiki/Data_deduplication)
- [8] Snapper GUI: <https://github.com/ricardomv/snapper-gui>

 Find us on Facebook  
<http://www.facebook.com/linuxpromagazine>

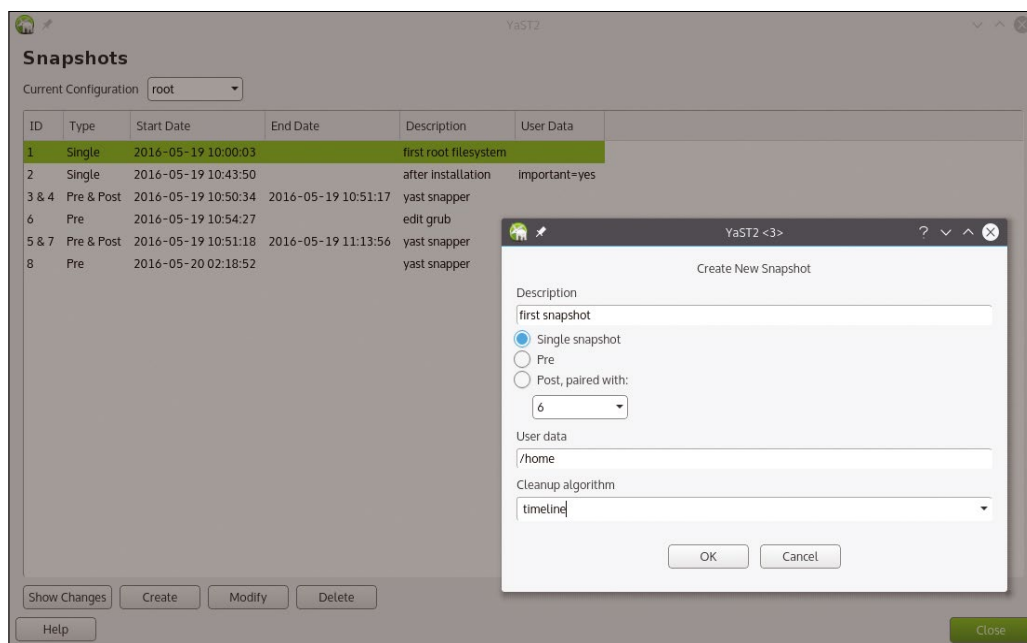


Figure 3: Creating a snapshot with Snapper.









### The professional filesystem ZFS

# Revolutionary

ZFS is a first-class filesystem for big iron, but for various reasons, it is still waiting for widespread Linux adoption. *By Erik Bärwaldt*

**W**ith over 15 years of development, ZFS [1] is one of the oldest of the current Unix-like filesystems. The ZFS filesystem was originally developed by Sun Microsystems for the Solaris operating system and was published for the first time in 2005. ZFS was originally intended as a closed-source, proprietary filesystem for high-end Solaris storage environments.

When Sun open-sourced Solaris in 2005 with the OpenSolaris project, ZFS went with it. Oracle acquired Sun in 2009 and, in 2010, Oracle declared it was returning ZFS to a closed-source development model. Thanks to the beauty of open source licensing, the ZFS community continued to develop and maintain the open source version of ZFS. The umbrella project for ZFS development is now known as OpenZFS.

The open version of ZFS is licensed under the Common Development and Distribution License (CDDL) [2]. The CDDL is recognized as a Free and Open Source license by the Free Software Foundation and the Open Source Initiative, but the CDDL has limited copyleft protection and is thus considered incompatible with the Linux kernel's GPL. Because of this incompatibility, it isn't easy for developers to integrate ZFS directly with the Linux kernel. Although various workarounds are possible, the license incompatibility has slowed ZFS adoption in Linux. As you will learn in this article, ZFS is typically implemented as a separate add-on module or as a Filesystem in Userspace (FUSE) [3] in Linux environments.

### Why ZFS

ZFS was the first of a new generation of filesystems built for our era of large disk drives and inexpensive storage. Until the end of the 1980s, the extremely high cost of storage space meant that most popular filesystems were designed for economy. Storage capacities for IDE and SCSI hard drives continued to grow through the beginning of the 1990s, and conventional filesystems were increasingly pushed to their limits: The call for more data security came in larger companies with their own IT infrastructure and data centers with large storage clusters, and this additional protection was primarily achieved using data mirroring.

At the same time, concepts of volume management were developed that, with the help of a logical volume manager (LVM), condensed several physical disks into a logical network and thus overcame the capacity limits of individual mass storage systems. However, the LVM was frequently combined with the old filesystems, meaning that their restrictions also determined the overall network performance.

As the complexity of the overall mass storage subsystem grew, so too did the administrative overhead and the risk of a lack of data integrity due to transmission and memory errors. The concepts of data mirroring also required outrageously expensive additional hardware in the early years that was usually offered in the form of plug-in cards for high-end systems.

These printed circuit boards normally supported both data mirroring and the distribution of data stocks via various physical mass storage techniques, meaning it was



possible to achieve a significant speed gain when retrieving data through parallel access mechanisms.

The ZFS developers took into account all these vulnerabilities and integrated functions in the filesystem that had previously only been possible using external solutions. ZFS is therefore not a filesystem in the original sense of the word, but rather a combined solution: The ZFS filesystem integrates an internal LVM and creates a storage pool that automatically manages the mass storage.

The individual pools automatically adjust their size as soon as the total capacity changes, such as when a new physical mass storage device is added to the system. ZFS performs the modification of the pools transparently, meaning you won't need to do any manual administrative work [4]. ZFS also automatically creates redundancies in order to improve data security.

Using snapshots, it is possible to duplicate a defined state of the system so that you can reconstruct the system based on the snapshot in the event of a failure. This technology, referred to as copy-on-write, works transparently in the background while the filesystem is active. The filesystem can also ensure economical use of the existing space thanks to the integrated data compression, which can lead to significant resource savings, depending on the type of data.

The integrity check is another innovation of ZFS: ZFS ensures data integrity between mass and working storage using checksums for all blocks. The filesystem even provides mechanisms for self-healing. If, based on various checksums, a variant turns out to be corrupt among the redundantly stored data, ZFS repairs the error using the stored copy. As a user, you can also trigger

a manual test run so that the data integrity is guaranteed at all times. These mechanisms eliminate the need for costly and time-consuming manual filesystem checks, which can take days in large data networks.

Last but not least, the developers have trimmed ZFS for high operating speeds: The filesystem provides significant speed increases during data transfer thanks to various cache levels, both in the working memory (adaptive replacement cache, ARC) and on device-based caches (cache-vdev, L2ARC). You won't risk data loss if the system fails, because the cached content is duplicated for both cache levels.

The Btrfs filesystem implements some of the advanced features pioneered by ZFS with a licensing model that is more compatible with Linux. (See the article on Btrfs elsewhere in this issue.) But Btrfs is relatively new to the mainstream and was only declared stable in 2013. ZFS, on the other hand, has been in active use for years.

## Requirements

ZFS was designed for servers and data centers, and it needs plenty of resources, including a 64-bit processor and sufficiently large working memory. You should provide 1GB of RAM for every terabyte of capacity. You might be able to use ZFS with fewer hardware resources, but you wouldn't be able to make full use of its benefits.


ZFS works with 128-bit-wide pointers and manages to deal with enormous storage capacities: The maximum capacity of a ZFS filesystem is 16 exabytes, where the maximum size of a file is identical. The hardware industry is unlikely to exhaust such capacities in the near future. Despite its origins as a filesystem for huge storage systems, ZFS is sometimes a competitive option for home servers, especially if you're using a recent computer with multiple hard disks or SSDs.

## In Linux

So far, ZFS is widely available on Unix-like operating systems, especially in various BSD derivatives. It is, for example, integrated in FreeBSD from version 8.0 as a stable filesystem, and it is the default filesystem in PC-BSD. In Linux, the copyright and licensing problems interfere with the direct integration of ZFS with the kernel.

ZFS on FUSE [5] is a popular option that avoids the licensing problems associated with Linux kernel integration. Like other FUSE systems, ZFS on FUSE operates ZFS in userspace. You can even use ZFS on FUSE to run ZFS on a 32-bit system. However, ZFS on FUSE is only available in an obsolete ZFS version; the project was terminated some time ago, but you can still find binary packages in the repositories of many distributions. Of course, putting ZFS in userspace causes some speed and stability issues.

The ZFS on Linux project, which integrates a self-compiled kernel module, proved much more sustainable. ZFS on Linux uses a newer version of ZFS that



ensures compatibility with Solaris 10, FreeBSD, and OpenSolaris. The software has been regarded as stable since version 0.6.3 and can therefore be used in productive environments. Pre-compiled packages exist for various Linux distributions [6].

Canonical made an attempt to increase the dissemination of ZFS under Linux in the winter of 2016. In an announcement from mid-February the company behind Ubuntu said that, as of the Ubuntu 16.04 LTS “Xenial Xerus” release, the ZFS filesystem will be implemented in the kernel. Canonical planned to use a kernel module by OpenZFS that has been under development since 2013.

But because this kernel module is also under the CDD license, the announcement provoked massive opposition from the community, as well as from the Linux developer group. The Free Software Foundation (FSF) also expressed worries about this step. Canonical, however, believes the integration of ZFS in Linux is long overdue: ZFS’s security and enormous capacity make it an ideal filesystem for cloud environments and clusters.

## In Practice

Because of the licensing issues, the integration of ZFS in the Linux universe is currently proceeding very hesitantly. Existing solutions are only suitable for use in larger environments, and most desktop distros still don’t support ZFS.

Canonical stirred up much dust with its announcement to integrate ZFS in Ubuntu 16.04 LTS. On closer investigation, this “integration” turns out to be far from finished; even power users are likely to fail when faced with the practical implementation.

The Ubuntu Installer Ubiquity lacks an option for using ZFS when setting up the mass storage. Users, instead, need to switch to the command line and set up the entire system manually. Completing the configuration requires good knowledge of the nomenclature and workings of ZFS. Users of Ubuntu and its official derivatives will find a guide to installing with virtualized containers on the Internet [7].

The Arch derivative Antergos has distinguished itself as a pioneer of ZFS integration on the Linux desktop. Without much media fanfare, the Antergos developers have integrated ZFS on the basis of OpenZFS and also modified the Cnchi installer.

Although Cnchi does not allow users to harness the full capabilities

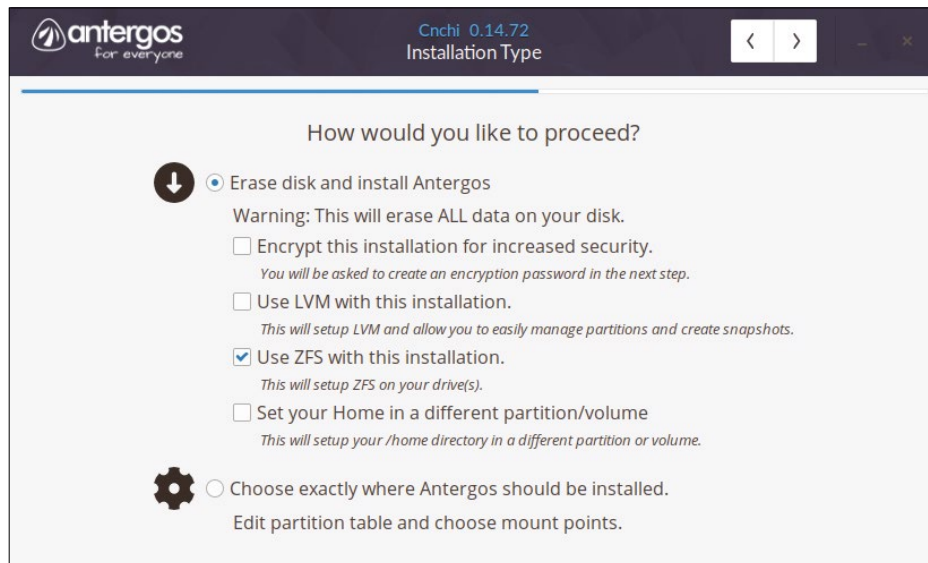


Figure 1: ZFS is already fully integrated with the new Cnchi Installer on Antergos.

of ZFS, it does provide an easy option for using the operating system with ZFS out of the box. You can then define further configuration options, if required, at the command line. During the installation of Antergos from within the Live system, you first define whether you want to use ZFS with your own home volume (Figure 1).

In a further step, the installer then prompts you for the basic filesystem parameters (Figure 2). You can specify the block size and define the ZFS operating mode, from a selection list in *Pool type*. Various RAID levels are available, so you can also enable mirroring or striping in systems with multiple physical storage media.

Please note that Cnchi is context sensitive: Faulty or incompatible settings prevent users from moving to the next step, which is done, in each case, by clicking on the right arrow at top right in the install window. It is thus impossible to define a RAID level as the pool type for computers with only one physical storage medium: In this case, the corresponding option must stay set to *None*.

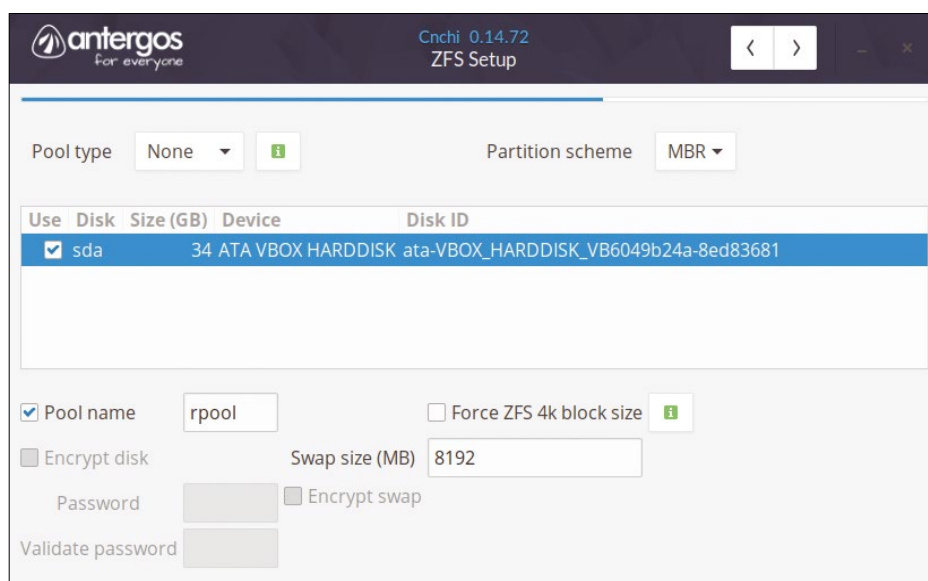


Figure 2: Users set up basic options, such as the RAID level or block size, in the Antergos Installer.





Antergos also sets up a swap partition without further intervention. The installer automatically configures the GRUB boot manager on legacy BIOS and on EFI systems, thus removing the need for manual adjustments. The routine also creates a 512MB boot partition formatted with ext4.

During the first reboot after installation, you need to avoid calling the alternative kernel variants, which Antergos offers in the GRUB boot manager in addition to the default kernel: These options do not include ZFS support and therefore cause an immediate crash.

## Details

After the initial setup, you then configure the ZFS parameters at the command line if necessary. For example, you can specify features such as data compression and quota definitions.

```
[erik@2540p-i ~]$ sudo -s
[sudo] Passwort für erik:
[root@2540p-i erik]# zfs get compressratio rpool
NAME PROPERTY VALUE SOURCE
rpool compressratio 1.01x -
[root@2540p-i erik]# zpool get dedupratio rpool
NAME PROPERTY VALUE SOURCE
rpool dedupratio 1.00x -
[root@2540p-i erik]# zpool list
NAME SIZE ALLOC FREE EXPANDSZ FRAG CAP DEDUP HEALTH ALROOT
rpool 232G 5,30G 227G - 1% 2% 1.00x ONLINE -
[root@2540p-i erik]#
```

Figure 3: ZFS supports compression and deduplication.

Just two commands – `zfs` and `zpool` – let you configure and manage your ZFS filesystem. The `zfs` command configures the filesystem and `zpool` manages the disks. To enable data compression for a storage pool, type the following command at the root prompt:

```
# zfs set compression=on pool_name
```

If you created multiple storage pools, you can find the one you need, including its current status data, with the `zpool list` command. After a system reboot, query the compression rate by typing the command:

```
# zfs get compressratio pool_name
```

The command shows the data in clear text in a neatly arranged list. Compression rates that can be achieved for previously compressed data, such as images or video and audio files, are low. The compression feature works much better with plain text files (Figure 3).

## Deduplication

ZFS also supports efficient mechanisms for data deduplication. This technique is used to avoid the redundant storage of

# 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



data, which also saves space. You can enable data deduplication using the command:

```
# zfs set dedup=on pool_name
```

To determine the duplication rate in the further course of working with the filesystem, use the following command:

```
# zpool get dedupratio pool_name
```

Clearly noticeable deduplication rates are usually only obtained with memory sizes of more than 4GB and several pools. Larger installations, in which multiple virtualized operating systems work in containers, are better suited for deduplication than desktop PCs with just one pool.

### Snapshots

Another common technique for enhancing the data security of a filesystem is creating snapshots. ZFS provides very easy and efficient snapshot capability thanks to its copy-on-write technology; older data is always saved, because the filesystem stores new datasets in free memory blocks instead of overwriting the existing data. The following command creates snapshots of a volume:

```
# zfs snapshot pool_name/volume@snapshot_name
```

The system creates the snapshot in the current directory in a hidden folder with the extension `.zfs`.

Although this folder is not displayed with the `ls -lisa` command, you can change to it with a simple `cd .zfs`. This `.zfs` directory has a `snapshot/` subfolder, where you will find your snapshot(s).

The individual snapshots are accessible as in conventional file hierarchies:

You can retrieve or copy files from them, and you can create backups. The snapshot technology built into ZFS is far more flexible than you will find in other filesystems (Figure 4).

### Conclusions

ZFS is one of the most innovative and modern filesystems available for Unix-flavored operating systems today. Restrictions of obsolete filesystems fall away completely with ZFS, and ZFS comes with a range of features, such as RAID, snapshots, and data compression. ZFS also cuts a good figure on the desktop. The less sophisticated Btrfs provides some of the enhanced features of ZFS, but Btrfs has its own issues. It is time to resolve the restrictive license disputes around ZFS and make the ZFS filesystem fully suitable for use on Linux. ■■■

### INFO


- [1] ZFS: <http://docs.oracle.com/cd/E19253-01/819-5461/zfsover-2/>
- [2] Common Development and Distribution License: <https://opensource.org/licenses/CDDL-1.0>
- [3] Filesystem in User Space (FUSE): [https://en.wikipedia.org/wiki/Filesystem\\_in\\_Userspace](https://en.wikipedia.org/wiki/Filesystem_in_Userspace)
- [4] Storage pools: <http://docs.oracle.com/cd/E19253-01/819-5461/gaypk/index.html>
- [5] ZFS on FUSE: <http://zfs-on-fuse.blogspot.co.uk/>
- [6] ZFS on Linux: <http://zfsonlinux.org/>
- [7] ZFS with Ubuntu: <https://insights.ubuntu.com/2016/02/16/zfs-is-the-fs-for-containers-in-ubuntu-16-04/>

```
[root@2540p-i eriks-daten]# zfs snapshot rpool/eriks-daten@Test-Snapshot
[root@2540p-i eriks-daten]# cd .zfs
[root@2540p-i .zfs]# ls
shares  snapshot
[root@2540p-i .zfs]# cd snapshot/
[root@2540p-i snapshot]# ls
Test-Snapshot
[root@2540p-i snapshot]#
```

Figure 4: Create and access ZFS snapshots from a terminal window.

Shop the Shop

shop.linuxnewmedia.com



# RASPBERRY PI ADVENTURES



COOL PROJECTS FOR GEEKS OF ALL AGES

ORDER YOUR VERY OWN ISSUE!



ORDER ONLINE: [shop.linuxnewmedia.com/se23](http://shop.linuxnewmedia.com/se23)



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

Better yet, subscribe, and you won't need a bookstore.

■ [shop.linuxnewmedia.com/subs](http://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







### Life and times of the classic ext Linux filesystem

# The Filer

The ext filesystem celebrates its 25th birthday next year. A brief tour of ext history will give you some insights into how this classic Linux filesystem works – and how it has evolved to meet users' needs. *By Valentin Höbel*

**W**hen Linus Torvalds developed the first early versions of Linux in 1991, he used the Minix filesystem by Andrew S. Tanenbaum. The Minix filesystem was part of the legacy of the Minix Unix clone [1], which Tanenbaum created for teaching purposes. Some restrictions, such as a maximum file name length of 14 characters or the file size limit of 64MB, gave rise to the need for a filesystem specially developed for Linux: The birth of the Extended Filesystem, or ext, followed shortly afterward.

French software developer Rémy Card released the first version of ext in 1992, making it possible to save files of up to 2GB on Linux. The permissible length for file names grew to 255 characters. Although ext got many things right in the first version, it was hardly up to professional requirements. For some applications, ext suffered from severe fragmentation, and it was impossible to store different time-stamps for access and modifications of the inode and file.

## Successor Ext2

When designing ext2, the developers adopted many best practices and principles of the then-widespread Unix Berkeley Fast File System [2]: Accordingly, an ext2 filesystem divides the storage medium into blocks from a logical perspective, then strings the blocks together. The default block size is 4KB.

The blocks should generally be at least as large as the hard disk sectors. At that time, the typical size was 512 bytes, which, from a logical perspective, is a block containing eight sectors. For several years, however, many hard disks have come with 4KB sectors, so a block directly maps to a sector.

Ext2 also adds blocks together to create groups; given a block size of 4KB, 32,768 blocks typically form a group (of about 128MB). A modern hard disk with a capacity of 2TB holds many thousands of blocks and around 16,000 block groups.

The division into blocks and block groups helps to organize the storage space logically and optimize read and write access. As a rule, the system writes files within the same block group in order to minimize fragmentation and access times to the storage medium. If the system writes a file that is greater than the configured block size, the file accordingly occupies a number of blocks.

The division into fixed block sizes offers a decisive disadvantage: If a file does not occupy the block size of 4096 bytes (4KB), it is wasting space. A 96-byte file fills a complete 4KB block, and a 5092-byte file fills two. In either case, 4000 bytes are lost.





```

boot: xxd
+ x boot: xxd
valentin@carbon:/boot$ sudo debugfs /dev/sda1
debugfs 1.42.9 (4-Feb-2014)
debugfs: dump grub grub_inode_dump
debugfs: q
valentin@carbon:/boot$ xxd <grub_inode_dump |head -n 12
0000000: 5985 0000 0c00 0102 2e00 0000 0200 0000  Y.....
0000010: 0c00 0202 2e2e 0000 5a85 0000 1800 1001  .....Z.....
0000020: 6766 7862 6c61 636b 6c69 7374 2e74 7874  gfxblacklist.txt
0000030: 5b85 0000 1000 0702 6933 3836 2d70 6300  [...i386-pc.
0000040: 5c85 0000 1000 0602 6c6f 6361 6c65 0000  \.....locale..
0000050: 9786 0000 1000 0502 666f 6e74 7300 0000  .....fonts...
0000060: 7885 0000 1400 0b01 756e 6963 6f64 652e  x.....unicode.
0000070: 7066 3277 9986 0000 1000 0701 6772 7562  pf2w.....grub
0000080: 656e 7600 7985 0000 7c03 0801 6772 7562  env.y...|...grub
0000090: 2e63 6667 7985 0000 6c03 0c01 6772 7562  .cfgy...l...grub
00000a0: 2e63 6667 2e6e 6577 0000 0000 0000 0000  .cfg.new.....
00000b0: 0000 0000 0000 0000 0000 0000 0000 0000  .....
valentin@carbon:/boot$

```

Figure 1: A look at the user-readable content of a directory inode using `debugfs` and the `xxd` hex viewer.

### Inodes

An ext filesystem stores the files without metadata on the hard disk, so it needs a way to manage the size of the files, their ownerships, and the access permissions. Additionally, the filesystem needs space for the exact location on the hard drive so that the system can find the file quickly.

The developers of ext2 used inodes for this purpose: Every file and directory is represented by an inode, which contains the ownership and access information. The name *inode* stands for “index node,” which is why the abbreviation *i-node* often appeared in the early years of ext2. An inode in ext2 has a default size of 128 bytes.

### Directories

From the perspective of ext, folders are nothing more than special files that are home to a list of files. Each entry associates a file name with an inode number, the length, and content from the actual file name. When accessing a file, the system just needs to take a look at the inode, which represents the directory holding the file (Figure 1).

File name and inode number mappings do not need to be singular: A completely different file name that points to an already-referenced inode number is known as a hard link. A hard link typically is not readily visible to applications and users, and it can only point to objects in its own filesystem. Incidentally, the same principle applies to subdirectories. A subdirectory is also a special file, which the inode

of the overlying directory references with its inode number as a file.

If you type `ls -a` to list the contents of a folder, you will notice the `.` and `..` entries. These entries are special properties of directories: When it creates a new folder, the system automatically generates them, and they cannot be deleted. Ext2 lists the entries `.` and `..` in the directory with the inode number of the current or parent directory. The root directory always resides in inode number 2, so that the system can find it faster.

### Superpowers

For an ext filesystem to be able to organize itself, it keeps a record of some important configuration parameters and current states. This information lives in the superblock, which always follows the first 1024 bytes of a storage medium (Figure 2).

This fixed position is important, for example, for the mount operation or for checking the filesystem. Among other things, the superblock contains the following:

```

dd@dd-ubu1204lts-d: ~
dd@dd-ubu1204lts-d:~$ sudo tune2fs -l /dev/sdb1
tune2fs 1.42 (29-Nov-2011)
Filesystem volume name:   HDD2
Last mounted on:         <not available>
Filesystem UUID:         180038a4-1a17-4cc6-a81b-e948ef009968
Filesystem magic number: 0xEF53
Filesystem revision #:   1 (dynamic)
Filesystem features:     has_journal ext_attr resize_inode dir_index filetype extent
flex_bg sparse_super large_file huge_file uninit_bg dir_nlink extra_isize
Filesystem flags:        signed_directory_hash
Default mount options:   user_xattr acl
Filesystem state:        clean
Errors behavior:         Continue
Filesystem OS type:      Linux
Inode count:             524288
Block count:             2096474
Reserved block count:    104823
Free blocks:             2026187
Free inodes:             524277
First block:             0
Block size:              4096
Fragment size:          4096
Reserved GDT blocks:     511
Blocks per group:        32768
Fragments per group:    32768
Inodes per group:        8192
Inode blocks per group:  512
Flex block group size:   16
Filesystem created:      Fri May 27 10:16:56 2016
Last mount time:         Fri May 27 10:17:34 2016
Last write time:         Fri May 27 10:22:04 2016
Mount count:             3
Maximum mount count:    -1
Last checked:            Fri May 27 10:16:56 2016
Check interval:          0 (<none>)
Lifetime writes:         258 MB
Reserved blocks uid:     0 (user root)
Reserved blocks gid:     0 (group root)
First inode:             11
Inode size:              256
Required extra isize:    28
Desired extra isize:     28
Journal inode:           8
Default directory hash:  half_md4
Directory Hash Seed:    ba43ef2c-f990-47e8-9a75-2eeb89a28e9b
Journal backup:          inode blocks
dd@dd-ubu1204lts-d:~$

```

Figure 2: The `tune2fs` tool reads the superblock of an ext2 filesystem.

```

dd@dd-ubu12041ts-d: /boot/grub
dd@dd-ubu12041ts-d: /boot/grub$ stat grub.cfg
  File: 'grub.cfg'
  Size: 4591          Blocks: 16          IO Block: 4096   regular file
Device: 805h/2053d   Inode: 793291      Links: 1
Access: (0444/-r--r--r--)  Uid: ( 0/   root)   Gid: ( 0/   root)
Access: 2016-05-27 08:47:00.000000000 +0200
Modify: 2016-05-27 08:47:00.000000000 +0200
Change: 2016-05-27 08:47:00.000000000 +0200
 Birth: -
dd@dd-ubu12041ts-d: /boot/grub$ █

```

Figure 3: The `stat` command offers a convenient view of the three timestamps.

- the exact block size
- the number of blocks and block groups
- the number of free and used inodes
- the date of the last successful mount
- the last change date
- the name of the operating system that created the filesystem
- the current state of the filesystem.

The superblock is more or less the heart of the ext filesystem. To avoid damage to the superblock that could destroy the filesystem, ext2 creates several copies of the superblock distributed across block groups. For repairs, tools rely on these backup copies.

## JOURNALING WITH THE EXT FILESYSTEM

A change in the ext filesystem takes effect in many places: If a new file is added, the filesystem reserves blocks and an inode position. Additionally, ext creates an inode, writes the data, changes the last access time in the inode of the corresponding directory, and updates the statistics in the superblock. Also, the kernel does not write the data directly to the hard drive. It first keeps the data in memory but tells the software that the operation was completed. It is only after a specified interval that the kernel writes several changes to disk all at once using an optimized approach.

If an unforeseen event like a crash occurs before the data is written to the disk, the crash is likely to make the filesystem inconsistent. For example, it is not clear what data the kernel has actually written to the hard disk and what is only in temporary memory. At boot time or during mounting, the system therefore checks the filesystem. Depending on the size of the filesystem and speed of the medium, this test takes as long as one to two episodes of your favorite TV show. At worst, you then still need to perform a manual repair.

A journal makes tests of this kind superfluous: The filesystem spreads out a kind of safety net. Under ext3, the changes are not written directly; instead, the kernel first writes the information to the journal (a kind of log book). Depending on the configuration, the journal might only contain the metadata (e.g., the number of reserved blocks and all other non-user data) or metadata and user data. By default, many distributions only save the metadata for performance reasons.

After writing all the relevant changes to the journal, the operation (transaction) is deemed complete; a commit winds it up. At the

Each block group has a group descriptor that holds some information about the block group. For security reasons, all block groups store all group descriptors of the ext filesystem.

Additionally, each block group contains inode and block bitmaps. These bitmaps act as registers that list all the inodes and blocks of each block group. A bit with a value of 0 marks an inode or block as free; a value of 1 indicates an occupied inode or block. A bitmap must always fit in one block, which also limits the size of a block group.

## Storage in Action

If you create a new file or folder on ext2, the filesystem first makes a decision on the exact storage location. For an empty filesystem, the data could theoretically end up anywhere, but for performance reasons, it makes sense to store

contiguous data as close together as possible. Ext2 thus tries to put subfolders and files in the block group in which the directory that holds them resides. The filesystem assumes that this data is related and the system will retrieve it, in quick succession if necessary. If the block group fills up, ext2 is forced to divert to another location. To determine which blocks and inodes are usable for storage, ext2 looks at the inode and block bitmaps.

## Good References

Ext2 also introduced support for symbolic links, or symlinks for short. A symlink is a special type of file that only contains a ref-

erence to another file. On the next opportunity, the kernel now writes the changes to disk, which ensures a new, consistent state. The actual data is already on the hard drive in the default setting, but the system only references the data in the filesystem after the journal commit.

If a system crash occurs between writing the journal and committing the changes, the system only needs to transfer the transactions from the journal to the filesystem when checking the filesystem during the next boot.

However, for complete safety, you need to make sure that the system keeps the changed content in the journal. You define how the journal works when mounting the filesystem with `mount -o data=<mode_name>` (e.g., `mount -o data=journal`). However, the journal mode affects the performance of the filesystem, which is why it is not enabled by default, although it provides the best safety.

The most used journal mode is ordered (often the default), which initially writes the contents to the filesystem and only then updates the metadata in the journal (Figure 4). The third mode, `writback`, only writes the metadata to the journal and leaves it to the kernel to decide when it writes the content to the HDD.

Depending on the configuration, a filesystem journal offers quite good insurance against system crashes. The journal also can be swapped out to another storage medium if necessary, so it does not impair the performance of the ext filesystem. However, journaling also causes far more write accesses to the underlying storage medium. For storage media with limited numbers of write cycles, such as flash storage, journaling therefore can drastically shorten the service life.





erence to another file or directory instead of data. In contrast to hard links, Symlinks point to targets on all filesystems, not just their own.

Unlike hard links, you can recognize a symlink. Applications that access the files or folders hiding behind symlinks treat them as ordinary files. But when it matters, such as when creating backups, a program needs to handle a symlink separately.

Symlinks exist independently of the target. If you delete the file that is referenced by a symlink, the link will remain – after all, it is a separate file with its own inode. The potential downside: If you change the name or location of a file, you also need to modify the symlink accordingly.

The filesystem stores all symlinks fewer than 60 bytes in size directly in the associated inode. Storing the symlink in the inode avoids the association of a separate block in the ext2 filesystem. (Actually, this same procedure is also used for normal files: If the content of a file takes up fewer than 60 bytes, the filesystem writes the data directly to the associated inode as “inline data.”)

## Extended Attributes

Version 2 of the ext filesystem was the first to offer the possibility to use extended attributes. One the best known of these

attributes is the immutable bit: You can use the command `chattr +i` as an administrator to protect a file against modification or deletion.

Moreover, the filesystem manages extended access privileges (Access Control Lists, or ACLs for short). ACLs make it possible to grant different users and groups access to files and directories – or to revoke this access. As you might expect, the extended attributes reside directly in the inode.

```
dd@dd-ubu12041ts-d: ~
dd@dd-ubu12041ts-d:~$ sudo dumpe2fs /dev/sdb1 |grep Journal
dumpe2fs 1.42 (29-Nov-2011)
Journal inode:          8
Journal backup:        inode blocks
Journal features:      (none)
Journal size:          128M
Journal length:        32768
Journal sequence:      0x00000006
Journal start:         0
dd@dd-ubu12041ts-d:~$
```

**Figure 4:** The `dumpe2fs` command displays information about the journal. The mode illustrated in this figure, inode blocks, corresponds to ordered mode.

# Shop the Shop

# shop.linuxnewmedia.com

## Missed an issue?

You're in luck.

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

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

**GET IT NOW!**

SAVE TIME ON DELIVERY WITH OUR ALTERNATIVE PDF EDITIONS



Another innovation from ext2 is multiple timestamps: The filesystem manages three different timestamps for files and folders (Figure 3). These timestamps allow tools such as `ls` to see when the system last accessed an object or something last changed at the level of the inode or data.

## Compressed

A few years after the release of ext2, the `e2compr` patch appeared on the web; this patch extends the ext2 driver in the Linux kernel to include support for compression [3]. `e2compr` only acts on files, while directories, inodes, superblocks, and other data related to the system remain unaffected.

After integrating the patch, you have the option of using

```
$ chattr +c <filename>
```

to set an extended attribute for a file. If you select a directory for compression, the system will automatically compress all newly added files.

## Not Quite Perfect

Despite many advantages over the first ext version, ext2 struggled with various limitations. For example, because of an internal limit, a directory must not exceed 31,998 directories, although this number probably rarely plays a role in practice. Because of the way ext2 manages files, the filesystem gets into trouble if you have more than 10,000 files in a folder.

Ext2 is at least better than ext in terms of the maximum file size and filesystem size: With kernel versions 2.4 or later, an ext2 filesystem can weigh in at up to 16TB and manage files of sizes up to 2TB (assuming the block size is 4KB).

With a Linux kernel 2.4 or earlier, the size of block devices was still limited to 2TB, which thus also limited an ext2 filesystem to this size.

## Jump to Ext3

By the turn of the millennium, ext2 had developed into a stable and widespread filesystem. Over time, however, ext2 had difficulty dealing with the growing volumes of data on ever-growing hard disks. Some developers, including Scotsman Dr. Stephen Tweedie, attempted some performance improvements and useful changes in the code.

However, the architecture of ext2 limited its evolution, which meant that it eventually made sense to move to a new version of the filesystem. Tweedie and other developers took the plunge and launched ext3, which was downwardly compatible to ext2.

One major innovation in ext3 was a feature that Tweedie had already developed in 2000 as an extension of ext2: the journal. This feature safeguards a filesystem against inconsistencies in case of a sudden crash (see box titled “Journaling with the Ext Filesystem”). Ext3 launched as a slightly improved ext2 with a journal.

As mentioned previously, ext2 does not scale well if it is home to many thousands of files in subdirectories. When the system is asked to search for a given file, it needs to conduct a painstaking examination of all the entries in the directory inode.

To be better prepared for the future, the developers have already given the filesystem an HTree index structure. This tree-

like structure allows ext3 to organize the contents of directories more effectively, making the search for files much faster.

The feature was originally developed for ext2, but it did not make it into the official source code at the time. Although it is enabled by default in ext4, ext3 still expects you to enable directory indexing manually. Also, ext3 lets you resize the filesystem while it is still mounted.

Additionally, you can define how the kernel behaves when it does not understand some of the filesystem metadata – for example, in case of damage. Depending on the configuration, the kernel can still mount the filesystem with restricted write access (writes are not recommended for reasons of data safety).

Similar to Windows filesystems, ext3 is prone to fragmentation when in constant use. Depending on how the free blocks are distributed across the filesystem, the system probably cannot always write contiguous data consecutively. The remedy for this problem is block preallocation, which ext3 uses to reserve the blocks before it actually needs them, which allows the filesystem to store the data as closely together as possible.

Incidentally, an online tool for defragmentation does not exist for ext3; however, several developers have released tools that at least optimize the distribution of the data somewhat. By adding a journal, you can convert an ext2 filesystem to an ext3 filesystem and vice versa.

The proximity to its predecessor is the biggest disadvantage of ext3: Because many structures resemble one another, ext3 initially lacked some features that competing filesystems already had. An ext3 filesystem with a block size of 4KB can accommodate a maximum file size of 2TB, and the filesystem can only grow to 16TB.

A few years after the publication of ext3, several companies and developers were still working on extensions and improvements (especially in terms of performance and stability). However, a lively debate developed on the kernel mailing list about whether further changes would really offer improvements to the fundamental problems or whether they would burden the existing users with even more disadvantages.

Finally, the developers agreed in 2006 to stop most of the work on ext3. Instead, they decided to copy the code of the filesystem to a new branch named *ext4* and only add new features and make substantial changes there. More than two years later, ext4 made it into kernel 2.6.28, gradually becoming the standard for many distributions.

## Ext4

Ext4 supports 64-bit processors, which makes it possible to create files with a size of 16TB for the first time, based on a block size of 4KB. The filesystem itself can grow to 1 exabyte (EB; approximately 1 million TB); this value is of a purely theoretical nature for most systems. In most cases, it is advisable to limit the filesystem size to around 16TB – and possibly to create several filesystems side by side.

Ext2 and ext3 store data in block bitmaps that represent the physical storage space. Under ext4, “extents” take over this feature in order to group contiguous blocks on the storage medium. A single extent on ext4 includes up to 128MB of contiguous space, and one inode contains up to four extents. If a file covers more than four extents, the filesystem indexes the re-





maining extents in a tree structure. Thanks to its use of extents, the ext filesystem can offer vastly improved performance for large files, as well as better defragmentation.

One feature painfully missed in ext3 was added by ext4: A new kernel function tells ext4 to reserve space for a file up front (“pre-allocation”). Ext4 fills the area with zeros, then tries to guarantee that the storage space really is available and that it is as contiguous as possible.

Through further modifications, ext4 now supports more than 32,000 subdirectories per folder – theoretically almost an infinite number, because the HTree mechanism now always intervenes. If you need more than the standard 64,000 permissible subdirectories, select the `dir_nlink` feature.

The journal added in ext3 has seen a meaningful improvement: Checksums now prevent the potential risk of defective metadata, which could destroy the filesystem. The infamous filesystem check also runs significantly faster under ext4 because unused storage areas in the filesystem are marked as such. The software just skips unused areas, which speeds up the check significantly. Fans of very precise timestamps also appreciate the fact that ext4 supports time-stamping with nanosecond accuracy. Ext4 now also has options for improving support for SSD storage media.

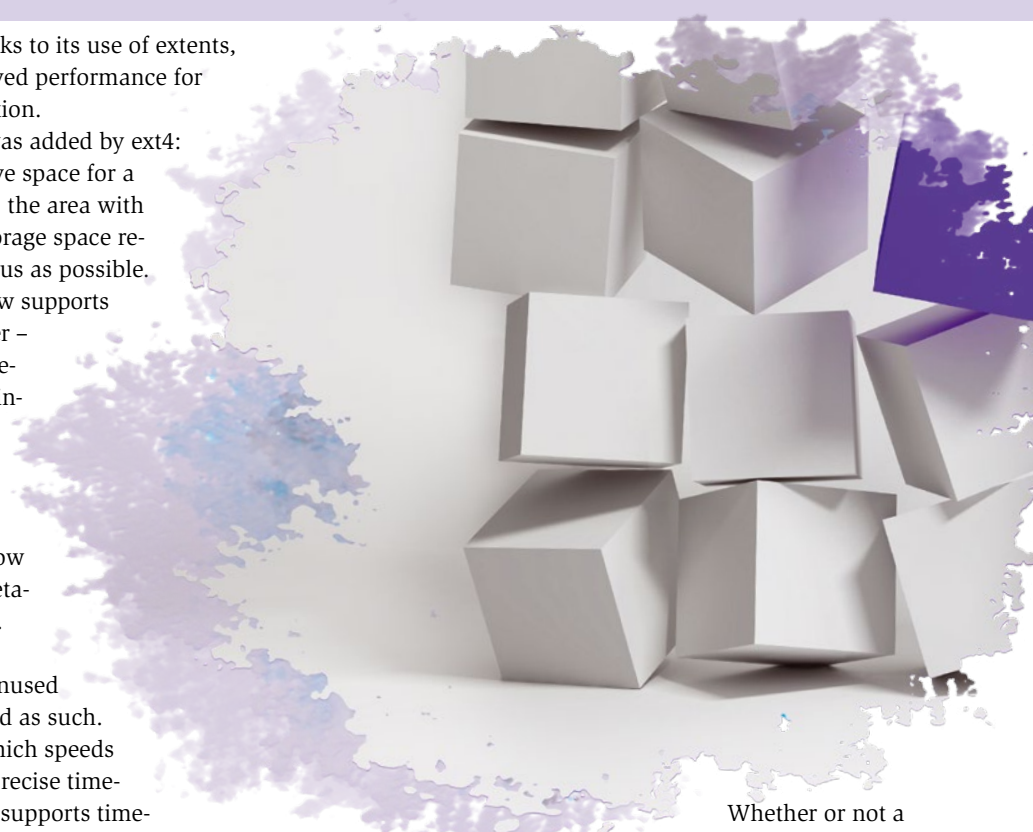
With the 4.1 kernel in mid-2015, experimental transparent encryption found its way into the ext4 filesystem [4]. As of kernel 4.4 this feature fulfills all the preconditions for future use. Incidentally, Google is the driving force behind ext4 encryption. Presumably, Google wants to offer better security features for Android and Chrome OS. If you are currently experimenting with a state-of-the-art kernel and newer userspace tools (e.g., `tune2fs`), check out some of the discussions online [5].

Ext4 provides backward compatibility with ext3 and ext2, which makes it possible to mount the predecessors as ext4 filesystems. However, you cannot mount an ext4 filesystem directly as an ext3 filesystem. To do so, you would first have to disable some features when creating the filesystem by adjusting the limits to reflect ext3’s capabilities.

Ext4 introduced quotas in kernel 3.6 that allow admins to limit the storage space available to users and user groups. Project quotas were also introduced with the recently release of kernel version 4.5. Project quotas let you limit storage space either for directory hierarchies or for files scattered across multiple folders. The patches were already submitted at the end of 2014 and described in a post on the ext4 mailing list [6].

## Future

Ext4 has been the default filesystem in many distributions over the past few years; after all, it is well suited for most purposes and is considered stable. Meanwhile, some distributors are moving to XFS (Red Hat) or Btrfs (openSUSE). Btrfs, in particular, is viewed as the leading contender to replace ext4 in the mid-term. (See the article on Btrfs elsewhere in this issue.) Btrfs includes many features that are unlikely to be added to ext4 in the foreseeable future, including filesystem snapshots, support for true online defragmentation, and copy-on-write operations.



Whether or not a fifth version of the ext filesystem will happen is currently uncertain.

In the past, the next version of the filesystem was always launched as soon as enough changes accumulated to justify a next higher version, while the developers sought to keep the existing ext version as stable as possible.

A patch description was posted on the ext4 mailing list by Oracle Developer, Darrick J. Wong back in May 2014; his idea was at least to introduce a number of flags under the ext5 banner [7]. During the discussion, other developers plainly stated, however, that the suggested changes were not sufficient to found a new filesystem. More likely, the developers will continue actively working on ext4 for now. But in the mid-term, many users will probably change to Btrfs. ■■■

## AUTHOR

**Valentin Höbel** works as a Cloud Architect for VoIP specialists NFON AG in Munich, Germany. In his spare time, he works on his Raspberry Pi or participates in open source projects.

## INFO

- [1] Minix: <http://www.minix3.org>
- [2] Unix filesystem: [https://en.wikipedia.org/wiki/Unix\\_File\\_System](https://en.wikipedia.org/wiki/Unix_File_System)
- [3] E2compr: <http://e2compr.sourceforge.net>
- [4] Ext4 encryption: <https://lwn.net/Articles/639427/>
- [5] Blog post on ext4 encryption: <http://blog.quarkslab.com/a-glimpse-of-ext4-filesystem-level-encryption.html>
- [6] Ext4 Quota patches: <https://lwn.net/Articles/623835/>
- [7] Ext5 patch: <https://www.marc.info/?l=linux-ext4&m=139898619610519&w=1>



Tool tests on the fast track *By Uwe Vollbracht*

# TOOL TIPS

## DK Tools 4.2.2

Function: Comprehensive tool collection  
 Source: <http://dktools.sourceforge.net>  
 License: BSD  
 Alternatives: GNU Tools

```
Terminal - vollbracht@LMLab:~/extract/tooltips/dktools-4.2.2
vollbracht@LMLab:~/extract/tooltips/dktools-4.2.2$ dk-pwgen
Wwqt3 gh4i
vollbracht@LMLab:~/extract/tooltips/dktools-4.2.2$ dk-pw
dk-pwd dk-pwgen
vollbracht@LMLab:~/extract/tooltips/dktools-4.2.2$ dk-pwd -l
dktools-4.2.2
vollbracht@LMLab:~/extract/tooltips/dktools-4.2.2$ dk-pwd
/home/vollbracht/extract/tooltips/dktools-4.2.2
vollbracht@LMLab:~/extract/tooltips/dktools-4.2.2$ dk-pwgen
7@0yyy0b7
vollbracht@LMLab:~/extract/tooltips/dktools-4.2.2$ dk4info
Used directories:
-----
--sysconfdir=/usr/local/etc
--bindir=/usr/local/bin
--libdir=/usr/local/lib
--datarootdir=/usr/local/share
--localstatedir=/usr/local/var

Configured at built time:
-----
--prefix=/usr/local
--sysconfdir=/usr/local/etc
--localstatedir=/usr/local/var
```

Developer Dirk Krause is collecting useful aides for the Linux shell under the name “DK Tools and Libraries.” The latest DK Tools 4.2.2 package contains more than 30 tools, including programs from the areas of graphics conversion, software development, file, user and directory management, word processing, database, and printer management.

The `dk-t2I` or `dk-t2h` tools provide support for converting text files into LaTeX or HTML format. `dk-pwgen` generates complex passwords, and `dk-eradi$` produces random content for overwriting a disk. The collection also includes alternatives to known shell programs with new features. The `dk-cat` tool normalizes the output’s line length, and `dk-pwd` just shows the name of the current directory without the path.

An extensive man page is available for every tool. Users can find more information and an overview of all the aides included in the package on the project website. Along with a thorough description of each tool, you will also find several examples of use.

★★★★★ The DK Tools package includes many practical tools for shell workers. The developer maintains the collection in an exemplary manner and regularly publishes new versions. ■■■

## Midnight Commander 4.8.15

Function: File manager for the shell  
 Source: <https://www.midnight-commander.org>  
 License: GPLv3  
 Alternatives: Vifm, Last File Manager

Midnight Commander (MC) is the bedrock of the shell file manager. The Ncurses-based tool was written by Miguel de Icaza, who developed the free clone of the DOS program Norton Commander in 1994.

MC provides a management interface in a two-column panel layout and tools for viewing and editing files. MCedit supports comprehensive syntax highlighting and integrates external tools such as the spell checker Ispell. MCview doesn’t just display plain text files, it also copes very well with man pages and HTML files. MC has supported CIFS and FTP for some time to provide access to remote directories.

Each new version provides numerous improvements. The current version fixes a segfault when using SFTP protocol. Other new features include the upgrade to version 2.26 of glibc and the display of up to nine columns – MC previously had been limited to two. Lua developers can look forward to improved MCedit highlighting.

★★★★★ The powerful file manager for the shell doesn’t need to hide behind graphics. MC is cutting edge. ■■■

### Datamash 1.1.0

Function: Analyzing and processing data

Source: <http://www.gnu.org/software/datamash>

License: GPLv3

Alternatives: Sed, Awk

```
Terminal - vollbracht@LMlab: ~/extract/tooltips/datamash-1.1.0/examples
vollbracht@LMlab:~/extract/tooltips/datamash-1.1.0/examples$ datamash -H -s -g 1
3 count 2 collapse 2 < genes_h.txt | head -5
GroupBy(name2) count(name) collapse(name)
ABCC1 1 NM_004996
ABCC10 2 NM_001198934,NM_033450
ABCC11 3 NM_033151,NM_145186,NM_032583
ABCC12 1 NM_033226
vollbracht@LMlab:~/extract/tooltips/datamash-1.1.0/examples$ datamash -s -g 13 c
ount 2 collapse 2 < genes.txt | head -5
ABCC1 1 NM_004996
ABCC10 2 NM_001198934,NM_033450
ABCC11 3 NM_033151,NM_145186,NM_032583
ABCC12 1 NM_033226
ABCC13 2 NR_003088,NR_003087
vollbracht@LMlab:~/extract/tooltips/datamash-1.1.0/examples$ datamash -H -g 2 co
unt 3 mean 3 pstdev 3 < scores_h.txt
GroupBy(Major) count(Score) mean(Score) pstdev(Score)
Arts 19 68,947368421053 10,143567752966
Business 11 87,363636363636 4,9409739060732
Health-Medicine 13 90,615384615385 8,8625266543026
Social-Sciences 15 60,266666666667 16,643183456165
Life-Sciences 12 55,333333333333 19,728715675944
Engineering 13 66,538461538462 19,10141108205
vollbracht@LMlab:~/extract/tooltips/datamash-1.1.0/examples$
```

Many users write their own scripts or put together constructions from `sed`, `awk`, and `cut` to evaluate and process structured program output on the shell. This work can be done more conveniently and in a one-liner with `Datamash`.

The tool imports the data via standard input. Parameters sort the input before processing, rearrange the headers, and define an alternate delimiter (tab is the default). Even lines with a variable number of fields don't put `Datamash` off its stride.

Orders such as `groupby` or `crostab` mean that `Datamash` only examines certain fields or columns. The tool is also familiar with more than 10 numerical and 20 statistical functions. The extensive man page and the well-documented examples on the project site help to get you started.

★★★★☆ `Datamash` is a useful tool for analyzing and evaluating structured data outputs. Users can use most of the functions with clever combinations of `sed`, `awk` and `cut` calls. `Datamash` makes achieving goals more convenient and faster. ■■■

### F3 6.0

Function: Testing the size of flash media

Source: <http://oss.digirati.com.br/f3>

License: GPLv3

Alternatives: none

```
huhn@wily: ~
root@wily:~# f3probe --destructive --time-ops /dev/sdc
F3 probe 6.0
Copyright (C) 2010 Digirati Internet LTDA.
This is free software; see the source for copying conditions.

WARNING: Probing normally takes from a few seconds to 15 minutes,
but it can take longer. Please be patient.

Bad news: The device '/dev/sdc' is a counterfeit of type limbo
You can "fix" this device using the following command:
f3fix --last-sec=32033311 /dev/sdc

Device geometry:
    *Usable* size: 15.27 GB (32033312 blocks)
    Announced size: 976.56 GB (204800000 blocks)
    Module: 1.00 TB (2^40 Bytes)
    Approximate cache size: 511.00 MB (1046528 blocks), need-reset=no
    Physical block size: 512.00 Byte (2^9 Bytes)

Probe time: 49'25"
Operation: total time / count = avg time
    Read: 1'31" / 2097717 = 43us
    Write: 47'47" / 8373468 = 342us
    Reset: 336.7ms / 2 = 168.3ms
```

Plenty of users are familiar with this phenomenon: The cheap 128GB USB stick is already full after using 64GB. This is because of a hardware manufacturer gimmick that makes the operating system think there is more space than is actually available. Windows users can test new storage media using the proven program `H2testw`. Linux users test flash media using the `F3` test suite.

`F3` contains the command-line tools `f3read`, `f3write`, `f3fix`, `f3brew`, and `f3probe`. The `f3write` tool writes random data onto a flash device and `f3read` reads it again. Users can access the two tools together with a USB memory mount point. Depending on the size of the data medium, you may want to allow for some waiting time. If the data volumes don't match in the end, this might be an indication of a "fake flash."

Anyone who doesn't want to start the read and write cycles manually can use `f3brew`. The tool writes 1GB blocks on the drive, performs a reset, and reads the data again. The `f3probe` tool checks the blocks of a medium and requires root privileges; parameters control thoroughness and storage needs. `f3fix` is a tool for repairing the partition table of flash media. It requires the `Parted` libraries.

★★★★☆ `F3` determines the actual size of flash memories and includes fake information. The tools `f3probe`, `f3brew`, and `f3fix` are still marked as experimental, but didn't cause problems when tested. ■■■



### Sauvegarde 0.0.7

Function: Client-server backup solution  
 Source: <https://github.com/dupgit/sauvegarde>  
 License: GPLv3  
 Alternatives: Urbackup, Bareos, Bacula

```
Terminal - vollbracht@LMLab: ~/extract/tooltips/sauvegarde-master
1453507574, mtime: 1453507574, size: 7580, filename: /home/vollbracht/extract/to
oltips/dktools-4.2.2/dk3maus.o, owner: vollbracht, group: vollbracht, uid: 1000,
gid: 1000, link:
file backend: --> type 1, inode: 828859, mode: 33204, atime: 1453507581, ctime:
1453507575, mtime: 1453507575, size: 15988, filename: /home/vollbracht/extract/t
ooltips/dktools-4.2.2/dk3maul.o, owner: vollbracht, group: vollbracht, uid: 1000,
gid: 1000, link:
file backend: --> type 1, inode: 828860, mode: 33204, atime: 1453507581, ctime:
1453507575, mtime: 1453507575, size: 15564, filename: /home/vollbracht/extract/t
ooltips/dktools-4.2.2/dk3mad.o, owner: vollbracht, group: vollbracht, uid: 1000,
gid: 1000, link:
file backend: --> type 1, inode: 828861, mode: 33204, atime: 1453507581, ctime:
1453507575, mtime: 1453507575, size: 9332, filename: /home/vollbracht/extract/t
ooltips/dktools-4.2.2/dk3maid.o, owner: vollbracht, group: vollbracht, uid: 1000,
gid: 1000, link:
file backend: --> type 1, inode: 828866, mode: 33204, atime: 1453507581, ctime:
1453507576, mtime: 1453507576, size: 17916, filename: /home/vollbracht/extract/t
ooltips/dktools-4.2.2/dk3maid8.o, owner: vollbracht, group: vollbracht, uid: 100
0, gid: 1000, link:
```

Sauvegarde is a new project. The backup solution includes three tools: `cdpfglserver`, `cdpfglclient`, and `cdpfglrestore`. The former is a central server component that `cdpfglclient` sends its data to and `cdpfglrestore` helps with restoring. The server listens on all interfaces on port 5468 by default. The incoming data ends up in a two-stage directory structure under `/var/tmp/cdpfgl/server`. Sauvegarde supports deduplication.

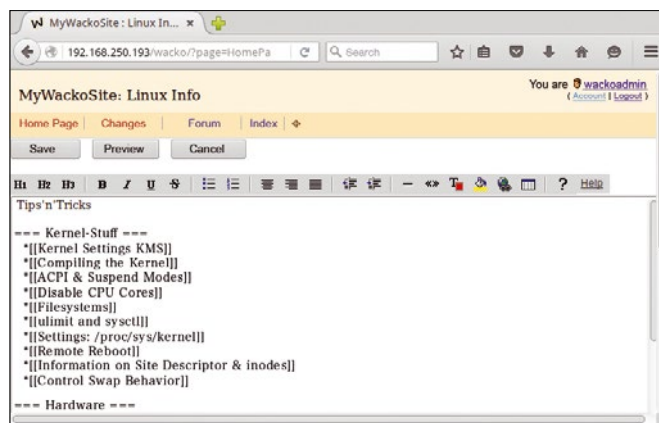
You can start the `cdpfglclient` tool on the systems you want to back up. In the configuration file, you can specify the directories that the client should keep an eye on, and you can exclude certain files or file patterns from the backup. The client backs up the data as soon as it detects a change. During the transfer to the server, it breaks the data down into 16KB blocks and computes checksums for checking. Information about the backups ends up under `/var/tmp/cdpfgl` in a SQLite database.

There currently is no proper documentation – just an installation manual and some configuration files in the GitHub repository. Users need to search for all other information themselves, although compilation from the sources can be done quickly. Users should also note the list of dependencies on GitHub.

★★★★☆ Sauvegarde is already a reliable client-server backup solution. However, because of the meager documentation, it may take a while before users manage to get a meaningful configuration. ■■■

### WackoWiki 5.4.3

Function: Resource-saving wiki  
 Source: <http://wackowiki.sourceforge.net>  
 License: BSD  
 Alternatives: DokuWiki, Moin Moin



Several wiki solutions are vying for the attention of Linux users. WackoWiki is located somewhere between top dogs like DokuWiki or MediaWiki and the minor alternatives such as Wiki on a Stick and LionWiki. The rather speedy engine provides a wiytiwyg editor, a multilingual interface, access management, data uploads, mail alerts, and more. WackoWiki requires PHP and MySQL.

Once you've unpacked the archive in the web server's document root directory, you can set up WackoWiki either using a text editor or via the web interface. The latter only makes sense if the server is isolated in a local environment or supports HTTPS. The SMTP server setup via the web front end wasn't successful in the test. Mail delivery is required for automatically registering new users. You can also set up the feature manually in the `config.php` file.

Working with the editor is convenient – a preview function helps users with composing their posts. The wiki administrator gives users different rights and restricts access to individual pages, as required. WackoWiki pages keeps an eye on sites and sends notifications via email if changes are detected. Registered users can determine the appearance of the wiki and the language in their own settings.

★★★★☆ WackoWiki is a compact, resource-saving wiki engine. Thanks to the user management system, it is suitable for semi-professional use, although the web-based configuration does have weaknesses. ■■■



Now Appearing on

# APPLE NEWSSTAND

New age convenience...

Our inspired IT insights  
are only a tap away.

Look for us on  
Apple Newsstand  
and the iTunes store.

Download  
a FREE issue of  
each publication  
now!





The Matrix decentralized communication tool tested

# Talk Talk



One tool to rule all online communication: one tool to find them, one tool to bring them all in, and the Matrix to bind them. An open standard for decentralized communication enters the scene. *By Mela Eckenfels and Heike Jurzik*

In the last 20 years, users have seen a number of communication tools come and go on the Internet. Text-based classics like IRC and ICQ made way for audio and video chatting with Skype, Google Hangouts, and others. If you were in Facebook groups five years ago, like-minded people now meet on WhatsApp, Snapchat, or another mobile messaging application. It's difficult remembering your way around all the apps and programs and remembering who you meet on what channel.

Matrix [1] promises to reorganize online communication. The web-based RTC and HTTP protocol is released under the Apache license and aims to bring all the services together under a single roof: Everyone chats in the style they like but can still reach everyone else. End-to-end encryption and a free choice of clients are just a few of the features the developers are planning. In addition to text, video, and audio communication, the service is looking to provide interfaces for the Internet of Things.

## Distributed but Shared

Matrix does not route communication via a central server, nor does it connect the participants' computers directly. Instead, home servers collaborate. They also act as account servers. Each user can be reached via a unique ID on the matrix network. The ID comprises the nickname and the server address, as in `@Mela:matrix.org` or `@hej:matrix.org`.

This means that Matrix not only solves the problem of often very ugly nicknames you see on IRC or Skype with system-wide unique usernames, but it also avoids central identity management failures, which are targets primed for denial-of-service attacks.

Distributed home servers also prevent netsplits. The major IRC networks in particular suffer from this temporary breakdown of servers into two or more subnets. The Matrix concept ensures that home servers only receive the traffic of a channel on which users are logged and are actively using the channel. Another advantage of this approach is that it avoids overloading the entire network.

The Matrix ecosystem [2] is already the playground of a number of applications. Although many clients were at the early or late beta stage when this issue went to press, the server situation is not looking as good: Only two of the six projects had left the alpha stage.

However, the Matrix developers are aiming for more than just real-time communication. Their goal is to counteract the fragmentation of the messenger market. Bridges will expand home servers, adding the ability to communicate with other networks and protocols. A text message gateway is in late beta, and bridges for IRC, Slack, and XMPP are still tagged early beta and alpha. No components are currently tagged "stable."

In addition to the Android and iOS clients, our lab team also installed the text-based IRC client WeeChat and integrated a Matrix plug in. We also chatted in the browser-based Vector.im [3]. The functionality and appearance are strongly reminiscent of the good old IRC. Matrix channels have several participants marked with a hashtag, much like IRC channels, and additionally

show the name of the home server (e.g., #matrix:matrix.org and #lm:matrix.org).

The clients display a list of available chat rooms. If you want to start a private conversation with another participant, you can open another room on the respective client and invite that participant to join you. Once they have consented, the conversation can begin.

Commands like /join, /nick, /me, /op, /deop, /kick, and /ban also work exactly as on IRC, but Matrix has additional functions for uploading and downloading files, audio and video telephony (some clients), a search function, and a cross-client cache. Each home server not only stores the account information, but also the communication in a history, and synchronizes this data with other home servers and clients. Switching between desktop and mobile devices is thus a seamless process.

## Just Talk

Much like Slack [4], the current star of the team communication world, the founder or the user with admin status can make the history of a channel accessible and searchable from the beginning. This means that participants added later can read past discussions. In the settings of the Matrix room, admins define which messages remain in the history (Figure 1). Individual messages about sensitive information can also be deleted selectively with a single mouse click.

We missed tools for improving the way communication via Matrix is organized. Mechanisms for adding other users to a personal contact list, as in ICQ, Skype, and others, would be nice. Also bookmarks for individual postings and a function for liking or pinning critical messages, as in competitor Slack, are not available yet.

Matrix uses Web RTC for audio and video conferences. Some clients display icons for initiating a phone call in public chat rooms and private conversations. If you click on one of the symbols in a Matrix channel with many visitors, you do not need to worry that you will be calling all the participants at the same time on all devices – a banner message simply indicates that a conference is currently going on in the chat room (Figure 2), and the user can take part in it with a single mouse click. Unlike a private chat, an audio or video call started here specifically addresses the chat partner, who then accepts or rejects at the push of a button.

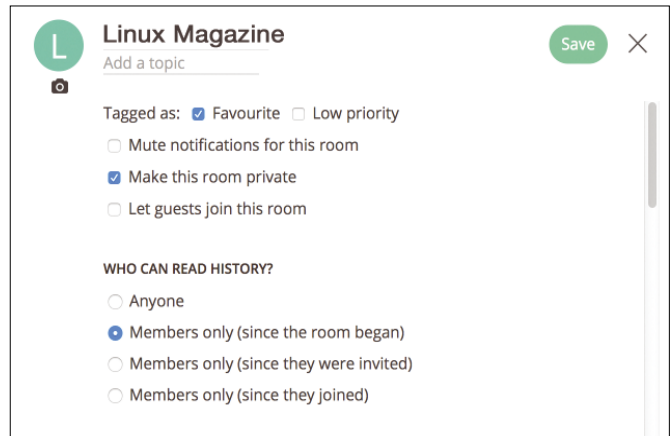


Figure 1: Administrators define who can access what part of the channel history in the room settings.

Phone calls to other Matrix users in our lab on the Matrix.org home servers were more or less reliable between web and Android clients and two browsers using Vector.im (Figure 3). Users with plain text clients are left out in the cold, and the feature has also not yet made its way into the I-OS app. Future versions for iPhone and iPad will be playing catching up here.


Another thing we really liked was that users can configure alerts both on the client side and for every channel and every conversation. Everyone can thus decide freely whether their smartphone, tablet, or browser should beep or write messages to the lock screen.

The Matrix concept envisages an open federation of home servers,

# LINUX UPDATE

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

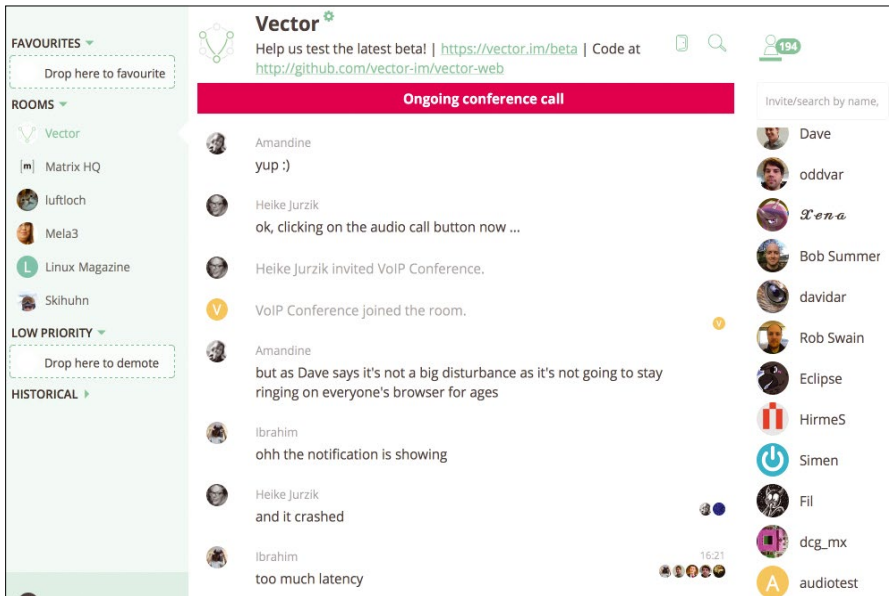
- Original articles on real-world Linux
- Linux news
- Tips on Bash scripting and other advanced techniques
- Discounts and special offers available only to newsletter subscribers



Ft Photography, Fotolia

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





**Figure 2:** On a Matrix channel, active phone calls are shown to new participants as a banner at the top.

which means that each user can add their own server to the existing Matrix network. Our lab team followed the instructions at the Matrix GitHub page [5]. We installed the software on Debian 8.3 (Jessie), which worked without problems. Setting up the SSL certificates is somewhat tricky. If they are not valid or are self-signed, some clients refuse authentication. The approach described on the Matrix blog [6], with a certificate from Let's Encrypt, an Nginx web server, and a redirect to port 443, should help.

Although the Matrix team attaches great importance to security, some parts of the system are currently ex-

cluded. Communication between home servers is TLS-encrypted, and the chat history is not readily visible. The promised end-to-end encryption for clients will be available shortly according to the FAQs. But the metadata ends up on the home servers without encryption, allowing admins or potential intruders to discover at any time who was talking to whom.

### Building Bridges

The open approach, the ability to connect to the Matrix network with any client, and running your own home server on the network – all of this makes the concept very interesting. Whether Ma-

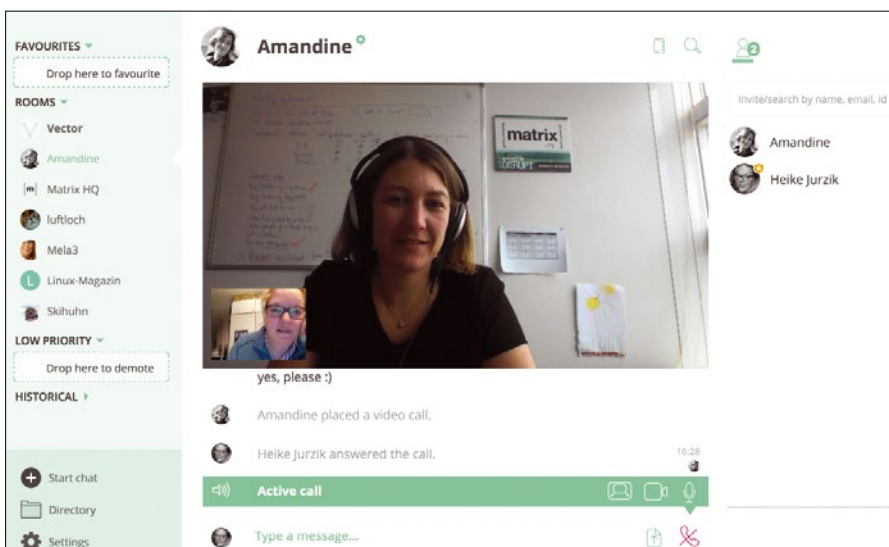
trix can revolutionize online communication not only depends on how user friendly the clients are, but also on how easy it is to set up your own home server. Our lab team is a little concerned that a kind of “diaspora effect” could occur here. The decentralized diaspora social network [7], highly acclaimed as an alternative to Facebook, with a guarantee of absolute control over your own data, never has been easy to install and manage.

Even with Matrix, the majority of users will probably chat via home servers provided by a third party, and Matrix's competitors see the major problems in precisely this. The developers of the Psyc [8] communication tool view the open federation of home servers and their identity management as a problem. It is naive, say the Psyc developers [9], to trust an arbitrary third-party server and its identity information.

It remains to be seen whether other vendors jump on the bandwagon and build bridges to Matrix. Other open standards, such as IRC or XMPP, are not the problem. Google, Facebook, Skype, and others, however, are likely to be keen to keep their users in their own ecosystems. So far, it is the smaller projects that are giving Matrix a chance, including the Freifunk Kiel [10] project, which links a Vector.im interface to the IRC channel and is especially looking forward to the history of past postings, according to the operator. ■■■

### INFO

- [1] Matrix: <http://matrix.org>
- [2] List of Matrix projects: <http://matrix.org/docs/projects/try-matrix-now.html>
- [3] Vector.im: <https://vector.im/beta>
- [4] Slack: <https://slack.com>
- [5] Matrix home server with Synapse: <https://github.com/matrix-org/synapse>
- [6] Home server with Let's Encrypt/ Nginx: <https://matrix.org/blog/2016/02/10/advanced-synapse-setup-with-lets-encrypt>
- [7] diaspora: <https://diasporafoundation.org>
- [8] Psyc: <http://www.psyc.eu>
- [9] Psyc developers on Matrix: <http://about.psyc.eu/Matrix>
- [10] Freifunk Kiel: <http://freifunk.in-kiel.de>



**Figure 3:** Video calls with other Matrix users already work perfectly. (Here, *Linux Magazine* writer Heike Jurzik talks to Matrix co-founder Amandine Le Pape.)


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



## Remote desktop applications

# Long-Distance Relationship



Remote desktop applications allow remote access to machines, including desktop sharing, which makes them useful assistants in both the personal and professional spheres. We look at a few candidates. *By Ferdinand Thommes*

If you need to administer computers over a long distance – whether for work or to help out friends or relations – the time for remote desktop software has arrived. It allows admins to repair the desktops of far-flung relatives in their free time. The software also comes in useful for businesses looking to provide support for their customers from afar. Desktop sharing is a core function, which means if a developer wants to present a new website to their colleagues, they can do so with ease.

In the end, roughly two types of remote desktop applications can be distinguished, with fairly fluid definitions. On remote desktop applications using virtual network computing (VNC), one machine operates as the server. It sends the image of its desktop as a compressed bitmap to the client. The programs run on

the server; the user manages them on the client by sending mouse movements and keystrokes to the server.

During desktop sharing, applications that use remote desktop protocol (RDP) or proprietary protocols send screen content as terminal servers. This is sent over the network to one or more clients (multi-channel) in the form of primitives, a mixture of images and information. They can even transmit the desktop sound if needed. The users then work on the server's desktop, alone or in parallel with other users, as if sitting directly in front of it.

Both software categories overlap when it comes to their functionality. Newer software, for example, makes it possible to simultaneously chat, send files, make calls over the Internet, and transfer videos. This article looks at the way long-distance relationships work

with RealVNC [1], TightVNC [2], TeamViewer [3], AnyDesk [4], NX [5], and its free offshoot X2Go [6]. We will be comparing functions and performance and considering the technology behind them.

## VNC: Framebuffer Ahead

RealVNC and TightVNC are part of a group of remote desktop applications that rely on VNC [7]. The basic technology behind VNC is the platform-independent remote framebuffer protocol (RFB) [8]. It works at framebuffer level, so it functions for window-based systems like Windows, OS X, or X11, but also makes cross-platform connections possible. RFB transfers the screen content as bitmaps, where the server only responds to a client-side `FramebufferUpdateRequest` with the changes since the last request (`FramebufferUpdate`).

In contrast with the members of the VNC camp, the remote support tools TeamViewer and AnyDesk use proprietary protocols. NoMachine's NX and the free implementation X2Go fall under the category of terminal servers, though they fulfill the purpose of a remote desktop program. Their advantage over VNC lies in their complete encryption of data traffic. In contrast with some VNC applications, their users do not have to manually set up port forwarding if the remote machine is not part of the local network.

### Versatile Use

Use of VNC is versatile. Sales reps can connect with company branches over the protocol. Because not all VNC variants encrypt, the admins often secure these sessions with the help of secure shell (SSH) tunnels. Together, they beat the problem of port forwarding at the same time. Companies with a number of field staff can save money thanks to free VNC clients, because tools like TeamViewer or AnyDesk charge fees for commercial use. At the same time, the data exchange remains under their own control.

The admin also controls servers that require a graphical interface via VNC. The hypervisor VirtualBox offers VNC to operate a virtual machine without display (headless) [9].

### RealVNC

The remote desktop application RealVNC (Figure 1) uses the RFB protocol. It operates across platforms and can be set up on Unix-flavored systems like Linux, OS X, Solaris, HP-UX, and AIX, and also on Windows and the Raspberry Pi. The current Linux distributions offer

their own packages for RealVNC. In Debian, these are `vnc4server` and `xvnc4viewer`; for RPM-based distributions, they are `realvnc-vnc-server` and `realvnc-vnc-viewer`. Meanwhile, apps are also available for the viewer for Android and iOS.

The software is written in Java and licensed under the GPL along with proprietary licenses. The developers offer four editions of RealVNC. The Open Edition costs nothing, although you do have to register and activate the software. Since version 5, there has also been a Free Edition as well as the Open Edition; this demands a free license key.

The Personal Edition and Enterprise Edition are both commercial releases, and they essentially differ in their target group. The Personal Edition is aimed at home users and small businesses (as well as Windows users with Vista or later installed), whereas the Enterprise Edition is for larger firms. Unlike the free Open Edition, both of these encrypt and allow users to exchange files, chat, and print on remote devices.

As the admin, you install VNC (in one of its variants) on the machine you want to control. You then run a VNC viewer on the client, which attempts to access the remote device. Viewer and server then establish a point-to-point connection, which allows secure communication when encrypted.

RealVNC is mainly suited to remotely accessing your own machine, but it is also a remote helpdesk application. The pricing model for the commercial variants depends on the number of desktops you want to monitor, as well as virtual ones. The vendor charges \$30 per desk-

top for the Personal Edition, while the Enterprise Edition comes in at \$44.

### TightVNC

Another remote desktop application is TightVNC (Figure 2), which likewise works with VNC and thus with RFB as a protocol. The software is licensed under the GPL, allowing users to work with it freely both in personal and commercial environments. TightVNC is written in C, C++, and Java, and it is available for the Linux and Windows platforms. Mixed operation of both systems is also possible; the vendor also offers a Java viewer, as well as an Android viewer with Remote Ripple [10].

Since the Java viewer has been around, there is no longer any separate Linux version of the client. This existed only up to TightVNC 1.3.10. The software can also be installed via the package manager of most distributions.

TightVNC is compatible with the client or server components of other VNC implementations. The tool focuses on compression, specifically using JPEG and zlib. This is also how it copes with low bandwidths. Users can watch movies and play DirectX games, even with a reduced image refresh rate on broadband connections. TightVNC encrypts passwords individually; complete encryption is on the agenda. To heighten security, the developers recommend tunneling connections over SSH.

With TightVNC, users control the desktops of remote machines. In version 2.0, the program introduced autoscaling, adjusting the size of the display in the viewer to that of the remote machine, regardless of the client's resolution. The

```

devil@siductionbox:~$ xvnc4viewer

VNC Viewer Free Edition 4.1.1 for X - built Aug 17 2015 11:14:27
Copyright (C) 2002-2005 RealVNC Ltd.
See http://www.realvnc.com for information on VNC.
Server: 85.214.207.215 :1

Fri Apr 8 19:25:36 2016
CConn:    connected to host 85.214.207.215 port 5901
CConnection: Server supports RFB protocol version 3.8
CConnection: Using RFB protocol version 3.8
Password:
Fri Apr 8 19:25:46 2016
TXImage:  Using default colormap and visual, TrueColor, depth 24.
CConn:    Using pixel format depth 6 (8bpp) rgb222
CConn:    Using ZRLE encoding

Fri Apr 8 19:26:18 2016
CConn:    Throughput 7777 kbit/s - changing to hextile encoding
CConn:    Throughput 7777 kbit/s - changing to full colour
CConn:    Using pixel format depth 24 (32bpp) little-endian rgb888
CConn:    Using hextile encoding
    
```

Figure 1: The admin establishes the connection with RealVNC via the command line.

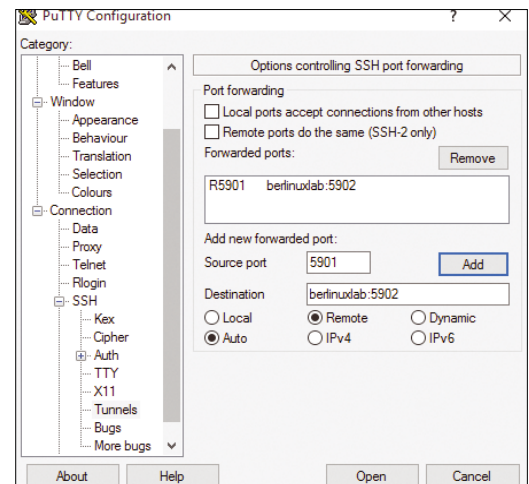
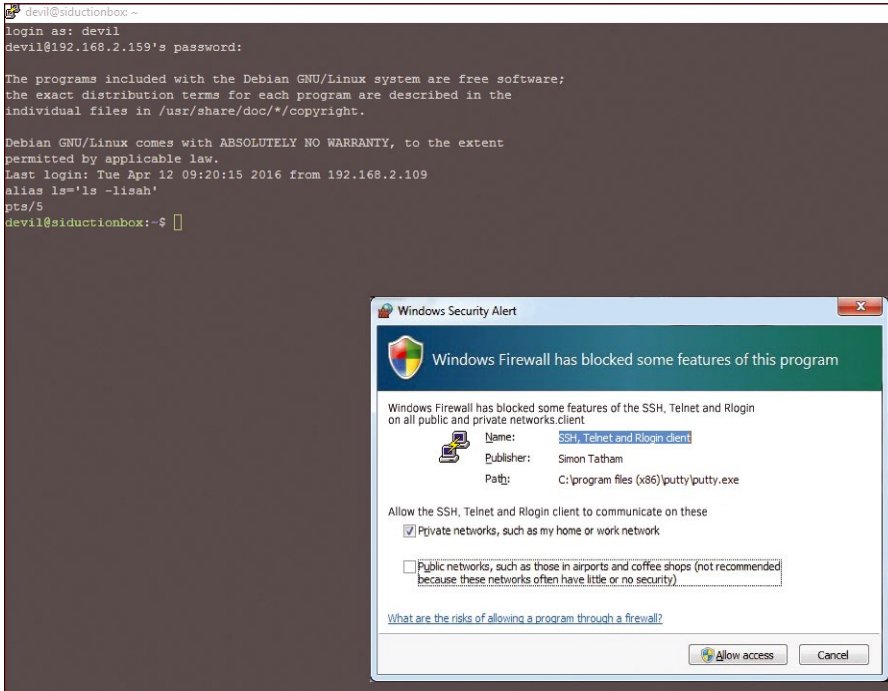



Figure 2: TightVNC over a SSH tunnel with PuTTY.





**Figure 3:** On Windows, access to the VNC server often requires modifications to the firewall. When using TightVNC, the admin should also install additional software to free up resources.



**Figure 4:** TeamViewer offers a range of options on the graphical interface's navigation bar.

admin can reduce the system resources required for the Windows platform (Figure 3) by installing the DF Mirage Driver Hooks [11] on the server side.

Additionally, TightVNC has spawned various offshoots, such as RemoteVNC, TightVNC Portable, TurboVNC, and TigerVNC [12]. The latter of these departed from its motherhood in 2009 and is Fedora's standard VNC application.

Unlike TightVNC, TigerVNC is equipped with extensions that let users authenticate and encrypt with transport layer security (TLS). Its software focuses mainly on 3D display and video applications.

### Lookalikes

TeamViewer and AnyDesk have much in common but also exhibit clear differences. Both applications are commercially designed, use proprietary protocols, and offer their own free versions for personal use. Both offer plenty of comfort – thanks to user-friendly GUIs (Figure 4) – penetrate firewalls without manual intervention from the admin, and require hardly any configuration.

The admin manages the access privileges via black- and whitelists.

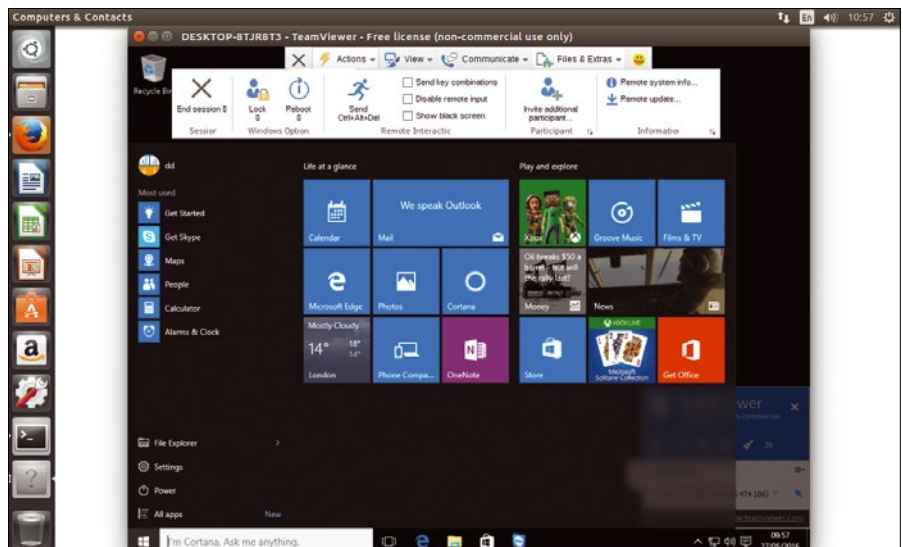
### TeamViewer

TeamViewer (Figure 5) is a contender both for personal and business use. The remote maintenance software offers screen sharing, video conferencing, file

transfers, chats, and VPN. It is available for Windows, OS X, and Chrome OS, as well as for RPM-based and Debian-based Linux distributions. Unlike AnyDesk, it does not run natively on Linux, but requires a Wine environment. Only a 32-bit version is available for Linux, meaning a multi-architecture system is necessary. TeamViewer supports Android, iOS, and Blackberry with apps in both directions, and there is a viewer for Windows Phone.

TeamViewer connects machines over secured data channels and initiates the encryption via a 2048-bit RSA public private-key exchange. After this, the software uses symmetrical 256-bit AES end-to-end encryption. This process is considered secure by current standards. Additionally, Verisign code-signs the software, with the intent of guaranteeing that only legitimate variants are in circulation. TeamViewer also offers commercial customers two-factor authentication.

For their payment model, the company relies on the sale of licenses. The customer pays a one-time fee and can then use the software without restriction. The company serves up licenses in four different flavors: Business, Premium, Corporate, and Enterprise. A Business license costs \$550, and the customer can use it on up to three devices. It includes a session channel with unrestricted endpoints, and the basic functions of a management console. For \$1,150 (Premium), or \$2,300 (Corporate), more devices can be integrated, and the management console offers more functions in line with this. The Enterprise plan is aimed at requirements



**Figure 5:** With TeamViewer, a Linux device can even comfortably access a Windows one.

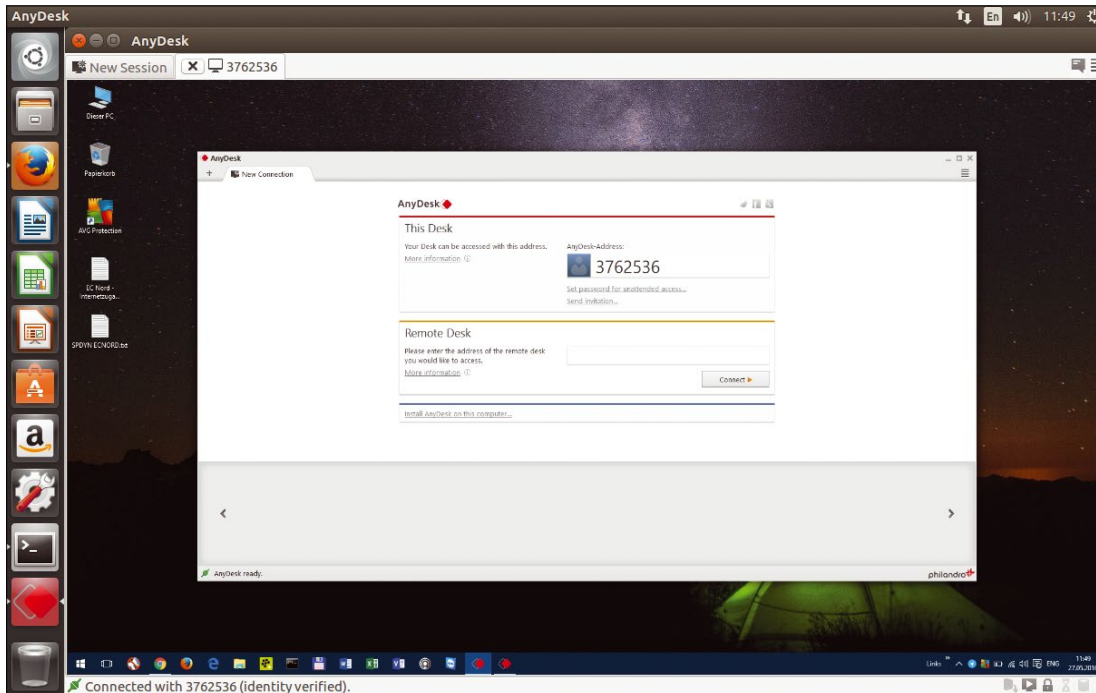


Figure 6: Windows session with AnyDesk on Ubuntu. The proprietary protocol transfers data in a flash.

of large corporations and for service providers, following consultation.

### AnyDesk

AnyDesk Software GmbH is still quite a young company, having been founded in 2014 by three former TeamViewer employees. In the same year, AnyDesk emerged (Figure 6) as a beta version for Windows XP, 7, 8.x, and 10. The Windows version can also be started with the simple click of a mouse on the .exe file (which is around 1MB in size), making it portable.

Meanwhile, AnyDesk is also available in stable versions for Linux and BSD. Native versions (for either 32- or 64-bit systems) for Debian, Ubuntu, Fedora, RHEL, Mageia, openSUSE, and SLES, as well as a generic Linux version are on hand for Linux admins. A Mac version and apps for Apple and Android tablets are planned.

Although user-friendliness is definitely a subjective question, the AnyDesk folks are well placed when it comes to the speed and quality of the transferred desktop. They are assisted by a proprietary protocol that transmits the data with DeskRT.

The video codec for graphical user interfaces – specially developed by AnyDesk – compresses the data. The protocol transfers the screen content thanks to a special compression process with up to

60 images per second. At the same time, AnyDesk caches up to 100 screen content captures in a buffer; thus it does not need to resend when need arises.

The developers have integrated AnyDesk into the respective operating system at a low level and optimized it for multi-process architectures. This means the image data reaches the screen after just a few processing steps. The company's servers, shared worldwide, use Erlang/OTP [13], a middleware for building shared, highly available systems, based on the Erlang programming language developed by Ericsson. TLS 1.2

### NoMachine NX

The terminal server software NX (Figure 7) from Italian software smithy NoMachine is the equivalent of Microsoft's RDP [14] for the Linux world and is particularly suited to connections that only use low bandwidth. NX increases the efficiency of the X11 protocol by compressing the data in the network traffic and creating a cache for data that have already been transferred. Furthermore, the in-house NX protocol reduces the round trips between X client and X server used extensively by X11; these increase the latency of the connection (Figure 8).

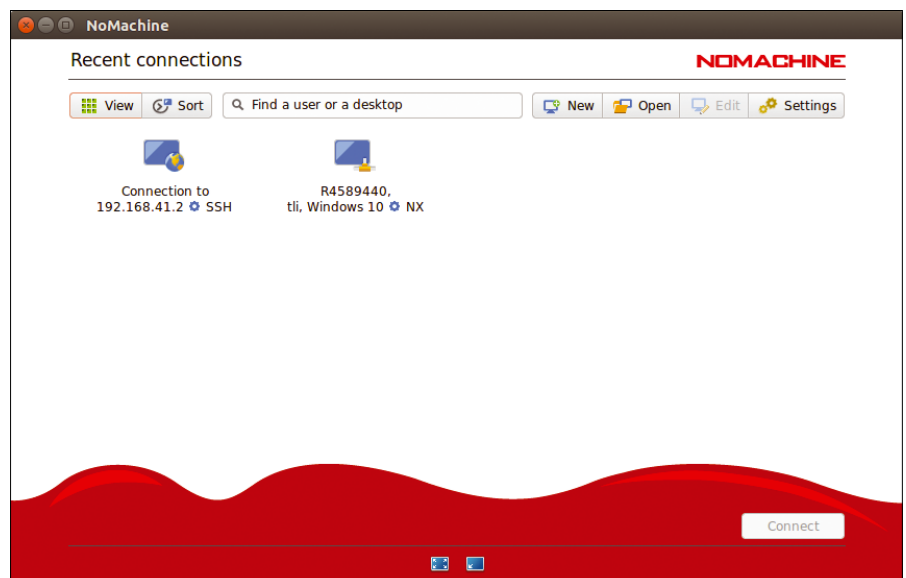


Figure 7: NoMachine's NX saves the last connections, facilitating quick access.



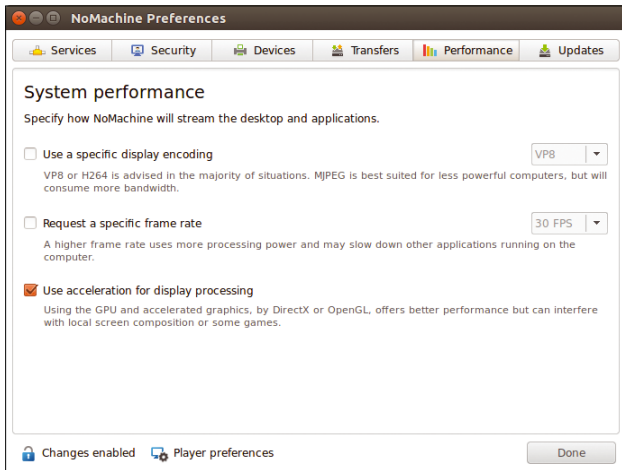


Figure 8: If you tinker with the individual options, you can squeeze out even more bandwidth.

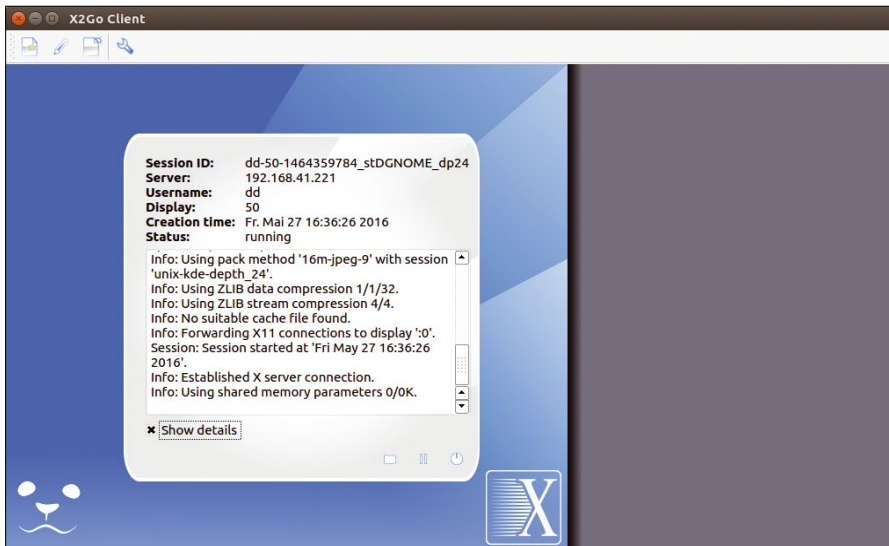


Figure 9: Configuring a session on X2Go, which uses the free NX protocol version 3.0.

NX transfers the data through a SSH tunnel. This relies on the client-server model by which the NX server can forward sessions to VNC servers or via RDP to Windows terminal servers, compressing the data traffic again. The NX client runs on Linux, Windows, OS X, and on other stationary and mobile platforms.

The server and client components of

NX have been proprietary since version 4.0. The developers offer a wide spectrum of Enterprise products that admins can test free of charge for 30 days [15]. NoMachine links its licensing model to the number of CPU cores used, along with the respective platform. The chosen distribution also plays a role on Linux. Prices start at around \$700 and are uncapped.

### X2Go

A popular, free alternative to NX is the terminal server X2Go (Figure 9), which similarly operates under the client-server principle. It is available for Linux, Maemo, and Windows. The client software runs on thin clients, PCs, web browsers, and mobile devices. You can access the sessions over a web client with a Firefox add-on. Like Free NX (no longer in active development), X2Go makes use of the libraries of NoMachine NX version 3, which are still subject to the GPL.

Based on this library, the X2Go project has built up a number of extensions, including an alternative graphical interface in the form of PyHoca-GUI. Because X2Go is popular with businesses, too, there is comprehensive documentation and a support forum. The client and alternative GUI are part of the standard repertoire in most distributions. For the server, web plugin, and many other extensions, the admin usually has

TABLE 1: Comparing Remote Desktop Applications

	RealVNC	TightVNC	TeamViewer	AnyDesk	NoMachine NX	X2Go
Platforms	Linux, Windows, OS X, Solaris, HP-UX, AIX	Linux, Windows, Android (Viewer)	Linux, Windows, Chrome OS, OS X, Android, iOS, Windows Phone (Viewer)	Windows, Linux, BSD	Windows, OS X, Linux, Solaris, Sharp Zaurus, Sony PlayStation 2, HP Compaq iPaq, Android, iOS	Linux, Maemo, Windows
License	GPL/proprietary	GPL	Closed source	Closed source	Closed source (since version 4.0)	GPL, AGPL
Protocol	VNC	VNC	Proprietary	Proprietary	NX protocol	NX protocol v3.0
Desktop sharing	Yes	Yes	Yes	Yes	Yes	Yes
Data exchange	Commercial only	No	Yes	Yes	Yes	Yes
Chat	No	No	Yes	Yes	No	No
Videoconferencing	No	No	Yes	Yes	No	No
Free version	Yes	Yes	Yes	Yes	No	Yes
Encryption	Commercial only	No	Yes	Yes	Yes (SSH tunnel)	Yes
Automatic port forwarding	No	No	Yes	Yes	Yes	Yes
On the market since	2002	2001	2006	2014	2004	2007

to integrate the X2Go repository. The project is subject to the GPLv2 and AGPLv3.

### Handling

The options for using the remote desktop applications mentioned in this article are varied and differ in their methodology, protocols used, and the licenses they are subject to. In fact, all the vendors have free variants on offer, although only TightVNC and its derivatives such as TigerVNC and the terminal server X2Go offer completely free software under the GPL.

As a common denominator in the area of functionality, all candidates in the test access remote machines. The VNC variants' field of activity is in the LAN, as well as in SSH-secured access to their own infrastructure from a distance. A variety of functions, as TeamViewer or AnyDesk offer their users, is usually not needed by them. Access to the desktop of the remote machine and data exchange is usually enough.

TeamViewer and AnyDesk are proprietary but offer free versions for personal use. They can be put into operation by less technically experienced users more easily than the VNC variants. Likewise, NoMachine NX is simple to handle thanks to its graphical user guidance, whereas the free variant X2Go does not create any great problems for moderately experienced Linux users.

One alternative, that requires very little effort and is ready for use in a minute, comes in the form of Google's Chrome Remote Desktop [16], which you can install as a browser extension.

The tested applications' user-friendliness varies; all candidates do well in typical usage on questions of the speed and quality of the display. Here, TightVNC is especially suited for slow connections due to its compression method. An easy-to-operate, versatile, and free piece of software comes in the form of X2Go.

If the license does not play a role, the proprietary newcomer AnyDesk is an option. It sets new standards in terms of compression, speed, and display quality with the specially developed DeskRT video codec. Here, there is still a limitation on Linux: If the controlling part uses a Qt-based graphical interface, it occasionally crashes. According to a statement from the developer, a solution is imminent.


The diversified portfolio of remote desktop applications covers many areas of use. Besides functionality, selection of license and commercial costs are among the deciding factors. In private application, experience shows that user-friendliness outweighs license considerations. As usual with Linux, the headache of decision-making falls on the user and depends on their own intended use (Table 1). ■■■

### INFO

- [1] RealVNC: <https://www.realvnc.com>
- [2] TightVNC: <http://www.tightvnc.com>
- [3] TeamViewer: <https://www.teamviewer.com/en/>
- [4] AnyDesk: <http://anydesk.com>
- [5] NX: <https://www.nomachine.com>
- [6] X2Go: <http://wiki.X2Go.org/doku.php>
- [7] VNC: [https://en.wikipedia.org/wiki/Virtual\\_Network\\_Computing](https://en.wikipedia.org/wiki/Virtual_Network_Computing)
- [8] RFB: [https://en.wikipedia.org/wiki/Remote\\_Framebuffer\\_Protocol](https://en.wikipedia.org/wiki/Remote_Framebuffer_Protocol)
- [9] VirtualBox Headless: <https://grenville.wordpress.com/2012/03/14/using-vnc-with-headless-VirtualBox-vms/>
- [10] Remote Ripple: <https://play.google.com/store/apps/details?id=com.glavsoft.rviewer&hl=en>
- [11] DF Mirage: <https://sourceforge.net/projects/mirage-driver/>
- [12] TigerVNC: <http://tigervnc.org>
- [13] Erlang/OTP: <http://learnyousomeerlang.com/what-is-otp>
- [14] RDP: [https://en.wikipedia.org/wiki/Remote\\_Desktop\\_Protocol](https://en.wikipedia.org/wiki/Remote_Desktop_Protocol)
- [15] NX Enterprise: <https://www.nomachine.com/enterprise>
- [16] Chrome Remote Desktop: <https://chrome.google.com/webstore/detail/chrome-remote-desktop/gbchcmhahfdphkkmpfmihenigjpp?hl=en>

Shop the Shop

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



# GIMP HANDBOOK

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



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



FOR WINDOWS, MAC OS, AND LINUX USERS!





## Restore Btrfs files at the push of a button with Snapper

# Oh, Snap!

Thanks to Btrfs snapshots, Snapper lets you restore a system to a previous state in case of failure.

By Ferdinand Thommes

**T**he Snapper tool uses Btrfs or Logic Volume Management (LVM) to manage system states, giving users an easy way to restore working files in case of problems. OpenSUSE offers a matching module for YaST. For other operating systems and filesystems, you will find a neutral user interface on GitHub.

### WARNING

Different distributions and experts assess the security of the fledgling Btrfs filesystem differently, depending to a great extent on the application and the features you use. Btrfs lead developer Chris Mason, who deploys the filesystem on the back end belonging to his employer Facebook, says it is ready for production use. The MeeGo mobile operating system also relies on Btrfs.

I was unable to find any reports of data loss in connection with Btrfs and Snapper in the SUSE support channels. To try out Snapper, it makes sense to set up a test system first, or at least create backups so that no valuable data is lost.

OpenSUSE developers trust the Btrfs filesystem [1]. Since openSUSE 13.2, or version 12 of SUSE Linux Enterprise Server (SLES), it has been the filesystem of choice for the root partition. On the other hand, the CoreOS distribution, which focuses strongly on container technology, recently dropped Btrfs, because the filesystem was too buggy in production use (see the “Warning” box).

Work on the filesystem has been progressing for about eight years now. Like the ZFS [2] OS developed by Sun for Solaris 10, Btrfs offers features that are not usually part of the filesystem’s repertoire. In addition to an extended storage area, these features include support for RAID0, 1, 5, 6, and 10; snapshots and sub-volumes; and the ability to defragment partitions on the fly. It is also possible to convert ext3/4 partitions to Btrfs, after creating a backup, with the `btrfs-convert` utility from the Btrfs toolbox.

### Snapping openSUSE

To test Snapper with Btrfs in the test lab, I used the latest release of openSUSE

Leap 42.1 and a current Ubuntu 16.04 beta, whereas the experiments with ext4 used Siduction (aka Debian “sid”) underpinnings. Because a current kernel is one of the most important prerequisites for working with Btrfs and snapshots, I updated all three systems to the latest versions before starting the experiments. In SUSE, that meant installing kernel 4.1.15 as well as `btrfsprogs 4.1.2-10.1` and `snapper 0.2.9.1`. Additionally, I installed `pam_snapper` to launch Snapper when logging in to the system. Siduction and Ubuntu were updated in a similar style.

One factor that prompted the choice of openSUSE as a test platform was that SUSE developer Arvin Schnell was the main developer behind Snapper and that the distro has a graphical tool that facilitates the task of managing the snapshots. However, Snapper can also be found in the archives of other distributions, such as Arch [3], Debian, and Ubuntu. With Debian, I could also test whether Snapper works on ext4, because the statements on the Internet are somewhat contradictory. Additionally, I wanted to convert a Debian installation with ext4 to Btrfs and then test Snapper.

### Configuration

After installing openSUSE Leap 42.1, Snapper is already set up for the root

Lead image © citalliance, 123RF.com

```

dd@dd-ubu1604d-test:~
To run a command as administrator (user "root"), use "sudo <command>".
See "man sudo_root" for details.

dd@dd-ubu1604d-test:~$ df -h
Filesystem      Size  Used Avail Use% Mounted on
udev            982M   0  982M   0% /dev
tmpfs           201M   3.7M  197M   2% /run
/dev/sda1       31G   3.6G   25G  13% /
tmpfs           1001M 252K 1001M   1% /dev/shm
tmpfs           5,0M   4.0K   5,0M   1% /run/lock
tmpfs           1001M   0 1001M   0% /sys/fs/cgroup
/dev/sda1       31G   3.6G   25G  13% /home
tmpfs           201M   4.0K  201M   1% /run/user/108
tmpfs           201M  116K  200M   1% /run/user/1000

dd@dd-ubu1604d-test:~$
dd@dd-ubu1604d-test:~$ sudo btrfs filesystem show
[sudo] password for dd:
Label: none  uuid: 9374c4aa-38f0-42e6-a5d0-7bb5f08fd905
Total devices 1 FS bytes used 3.32GiB
devid    1 size 30.10GiB used 6.02GiB path /dev/sda1

dd@dd-ubu1604d-test:~$

```

Figure 1: Disk space usage on Btrfs: The output from the commands `df -h` and `btrfs filesystem show` is significantly different.

partition. In this state, it creates a snapshot before and after each action with YaST or Zypper. It is also possible to set up hourly snapshots via the timeline; alternatively, you can create snapshots manually. The software removes most snapshots according to a defined schedule, but it does keep some of them by default.

If you leave the default settings from the `/etc/snapper/configs/root` file, then the snapshots will quickly occupy a large amount of space. You need to consider this when partitioning. A root partition that is normally fine with 30GB of disk space, will need between 100 and 300GB of space for snapshots – depending on the intended Snapper configuration.

Unfortunately, commands such as `df -h` or `du`, when used to display used and free space for the various snapshot partitions residing on subvolumes, do not work as expected, as evidenced by the many occurrences of the *no space left on device* error message among the complaints associated with Btrfs. In general, it is therefore advisable to use the `btrfs filesystem show` command instead of the classic tools (Figure 1).

If you also want to use Snapper for the home partition, you need to convert the filesystem for this partition from XFS to Btrfs during the installation. If Home lives on the root partition, Snapper does not cover it.

After the kernel update,

```

$ sudo zypper ar http://download.opensuse.org/
-f http://download.opensuse.org/
  repositories/Kernel:/stable/
  standard/ kernel
$ zypper ref
$ zypper dup -r kernel

```

### LISTING 1: Snapshot Configuration

```

# subvolume to snapshot
SUBVOLUME="/"
# filesystem type
FSTYPE="btrfs"
[...]
# run daily number cleanup
NUMBER_CLEANUP="yes"
# limit for number cleanup
NUMBER_MIN_AGE="1800"
NUMBER_LIMIT="10"
#create hourly snapshots
TIMELINE_CREATE="yes"
[...]
#cleanup hourly snapshots after some time
TIMELINE_CLEANUP="yes"
# limits for timeline cleanup
TIMELINE_MIN_AGE="1800"
TIMELINE_LIMIT_HOURLY="3"
TIMELINE_LIMIT_DAILY="2"
TIMELINE_LIMIT_MONTHLY="0"
TIMELINE_LIMIT_YEARLY="0"

```

# GOT CLUSTER?

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

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



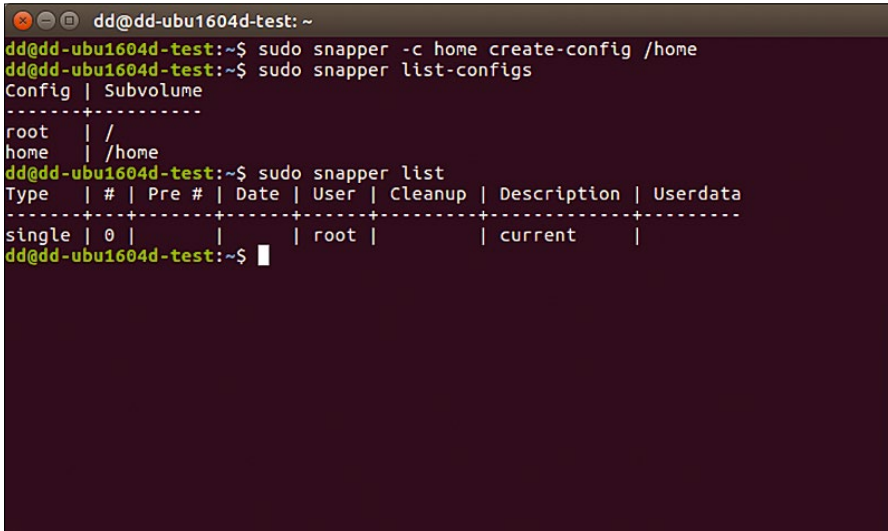


Figure 2: You can add configurations to create snapshots of other areas of the filesystem.

but before updating the system via YaST or Zypper, you need to edit the `/etc/snapper/configs/root` file to specify the intervals for creating and deleting snapshots. For a desktop system, it makes sense to lower the default values in some places. Listing 1 shows the most critical parameters as examples. The man page for `snapper-configs` [4] provides further insight. For a detailed list of parameters, see the SUSE documentation [5].

If you want to create a snapshot of the home partition, first create the appropriate configuration with the

```
sudo snapper -c home create-config/home
```

command and then edit it as shown for root. To check that the system has integrated the two subvolumes correctly, type:

```
sudo snapper list-configs
```

In addition to root, you should have an entry for home (Figure 2).

The `etc/snapper/filters` file lists the files the software never restores, so you can enter the files here that you want to exclude from the restore. Now that the configuration is complete, you can manage Snapper in the graphical user interface (Figure 3). The corresponding man page explains how to manage Snapper at the command line [6].

### GUI Control

In the YaST Control Center below *Software* is a *Miscellaneous* section with a

*Snapper* module that lets you create, view, compare (Figure 4), restore, or delete snapshots in a GUI. In addition to automatic snapshots, whether cron-controlled or generated by updates, pressing the *Create* button takes a new snapshot manually (Figure 5), which makes sense after a major change has been made to files.

If you encounter any problems later, you can use *Restore Selected* to revert to a state

before the change [7] (Figure 6). If the system is in a state that no longer allows you to roll back within the existing session, you can also select a snapshot in GRUB after a restart and boot to it.

### Ubuntu and Debian with ext4 and Btrfs

On Debian sid/Siduction and Ubuntu, I tried to get Snapper and the Snapper GUI [8] running using ext4. Several attempts delivered mixed results that do not suggest a predictable experience; therefore, I would advise against the use of ext4 with Snapper right now. Because I had already installed all components and configurations for Snapper and the Snapper GUI, I decided to convert the existing ext4 filesystem to Btrfs.

The Debian instructions on how to perform such a conversion [9] were useful. After a subsequent reboot,

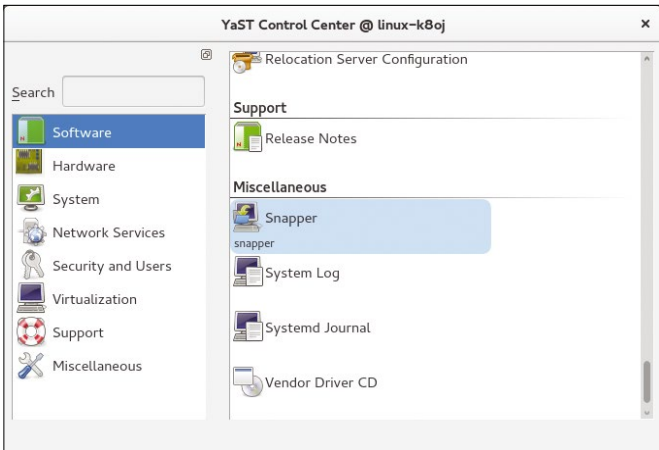


Figure 3: Thanks to a GUI, openSUSE makes it easy to manage, compare, and restore snapshots as needed.

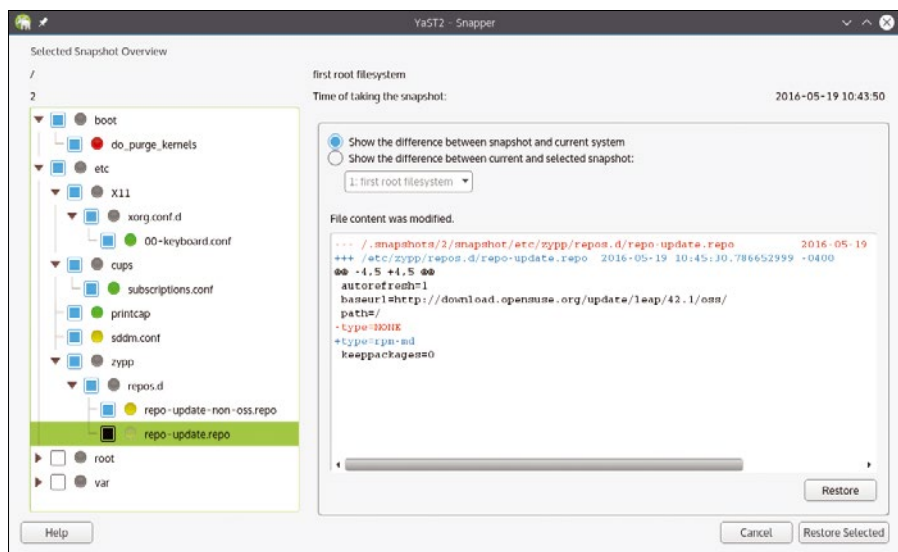


Figure 4: The graphical interface provided by openSUSE as a YaST module for Snapper lets you compare two snapshots.

Snapper and the Snapper GUI worked well together. Note, however, that it is faster to set up a new system with Btrfs than to convert from ext4. The process described here was more of an exercise to see whether the conversion could work at all.

While re-installing Ubuntu with Btrfs as the filesystem on a 30GB partition, I took advantage of putting the system under the control of LVM [10]. If you want to do this, select the *Use LVM* checkbox during installation, then right-click the partitions below *Others* at the bottom, which lets you choose Btrfs as the filesystem (Figure 7). The rest of the installation proceeds as usual. After re-starting the installed system, you should first upgrade it, because each new kernel of late has introduced optimizations for Btrfs.

Next, install the *snapper* package; *libpam-snapper*, if needed, which launches Snapper when you log in to the system; and *apt-btrfs-snapshot*, which ensures that a snapshot is created after every upgrade with Apt. Then install *snapper-gui*. As a precondition, you only need to install the *git* and *python3-setuptools* packages on a recent Ubuntu. Then, type

```
git clone --depth=1 https://github.com/ricardomv/snapper-gui.git
cd snapper-gui/
sudo python3 setup.py install
sudo snapper create-config /
sudo snapper-gui
```

to download the code for and install the GUI, create a basic configuration for snapshots of the entire system, and launch the GUI, which is relatively self-explanatory with its sparse controls (Figure 8).

You can create configurations and open, remove, and view the differences between snapshots. You use the left Ctrl key to select the snapshots, and the changes are then displayed when you click on *Changes* (Figure 9). All told, the Snapper GUI offers roughly the same features as the YaST module in open-SUSE. In the end, however, with a little practice, managing Snapper in a terminal window is not rocket science.

Snapper on Ubuntu, Debian, and probably other operating systems uses the same configuration file for snapshots




Figure 5: Creating a snapshot manually.

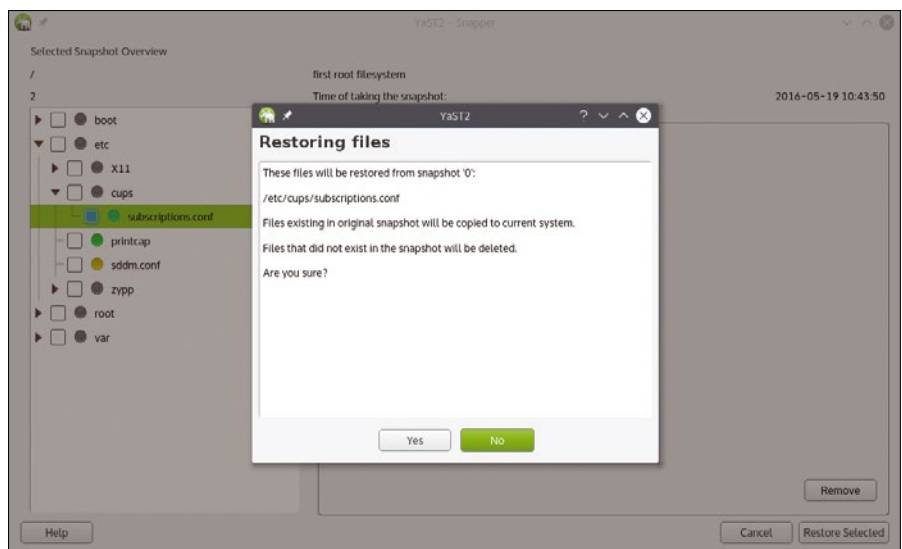


Figure 6: If you made a configuration error, you can restore a working version of a file from a snapshot.

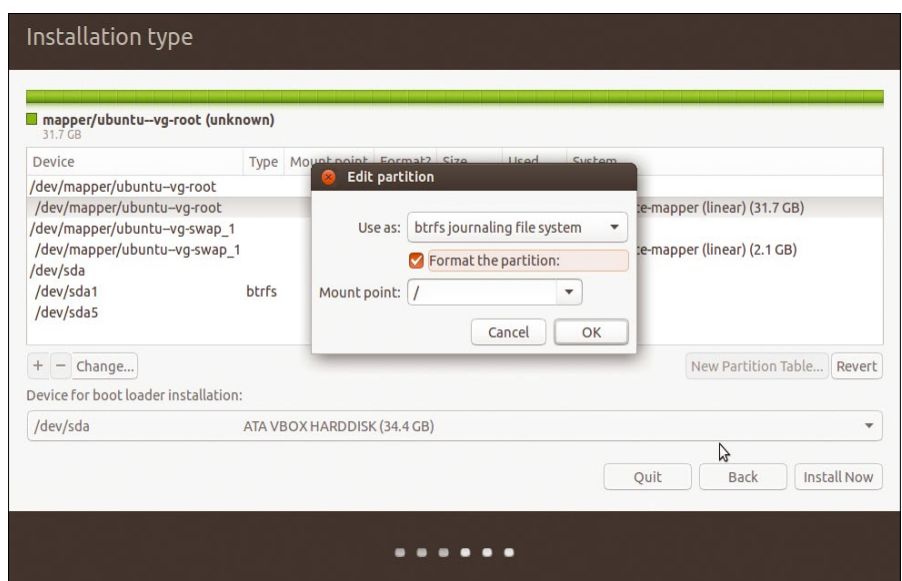


Figure 7: Selecting Btrfs in the installer.



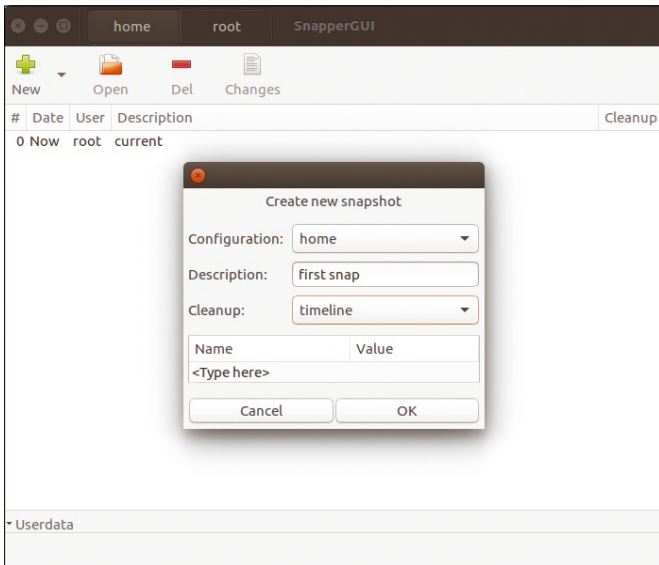


Figure 8: Creating the first snapshot.

as on openSUSE. Before you start production operation, you will want to check all the configurations you created in `/etc/snapper/configs` and modify them to match Listing 1 or to suit your own needs. All further management steps are as already described; the commands described for openSUSE also work on Debian and Ubuntu.

Although I chose to use LVM, it is not mandatory because Btrfs already has LVM functionality built in. Setting it up in this case can be viewed as more of a test as to whether the duplicate functions would interfere

with each other, which could not be proved in the lab. An error with the `apt-btrfs-snapshot` package, which is supposed to ensure that a snapshot is automatically created after every upgrade with Apt, currently prevents it from working; therefore, before a system upgrade, you should use the command line or the GUI to create a snapshot.

distributions. The software provides a valuable service by letting you roll back to a prior state before upgrading.

If you operate a stable system and do not install a lot of software daily, it is worthwhile calculating Snapper's overhead. Under certain circumstances it might be advisable to disable the snapshots launched every hour by cron and to only use automated snapshots when installing or upgrading.

The stability of Btrfs is an important factor. During the tests on three operating systems, more than 100 snapshots were created automatically and manually; they were compared, rolled back, and used to launch snapshot systems in GRUB. Under Btrfs, Snapper and the Snapper GUI on Ubuntu were stable; no errors endangered the data in any way. All the same, you should remember that Snapper is not a replacement for complete backups: In fact, it is more like the well-known Windows System Restore feature. ■■■

## Conclusions

Snapper optimizes the Btrfs snapshot feature, making it a valuable helper on systems subject to continuous change, such as developers and users of rolling release

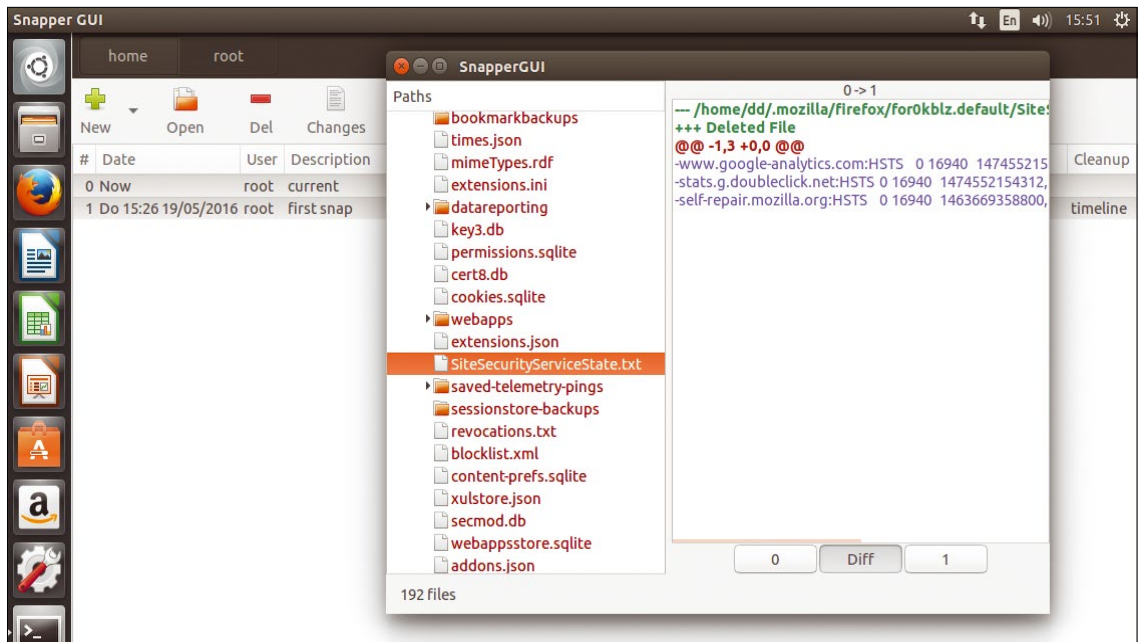



Figure 9: Differences between two snapshots.

## INFO

- [1] Btrfs: <http://en.wikipedia.org/wiki/Btrfs>
- [2] ZFS: <https://en.wikipedia.org/wiki/ZFS>
- [3] Snapper in Arch: [https://wiki.archlinux.org/index.php/Snapper#Create\\_a\\_new\\_configuration](https://wiki.archlinux.org/index.php/Snapper#Create_a_new_configuration)
- [4] Snapper configuration: <http://snapper.io/manpages/snapper-configs.html>
- [5] SUSE documentation for Snapper: [https://www.suse.com/documentation/sles-12/book\\_sle\\_admin/data/sec\\_snapper\\_config.html](https://www.suse.com/documentation/sles-12/book_sle_admin/data/sec_snapper_config.html)
- [6] Snapper man page: <http://snapper.io/manpages/snapper.html>
- [7] Rolling back files: [https://www.suse.com/documentation/sles-12/book\\_sle\\_admin/data/sec\\_snapper\\_snapshot-boot.html](https://www.suse.com/documentation/sles-12/book_sle_admin/data/sec_snapper_snapshot-boot.html)
- [8] snapper-gui: <https://github.com/ricardomv/snapper-gui>
- [9] Debian how-to: [https://www.debian-administration.org/article/675/How\\_to\\_convert\\_a\\_Wheezy\\_or\\_newer\\_system\\_to\\_btrfs](https://www.debian-administration.org/article/675/How_to_convert_a_Wheezy_or_newer_system_to_btrfs)
- [10] LVM: [https://en.wikipedia.org/wiki/Logical\\_volume\\_management](https://en.wikipedia.org/wiki/Logical_volume_management)

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



Debian 8 on the Nexus 5

# Pocket Linux

Maru OS provides a desktop Debian on a smartphone. *By Marcel Hilzinger*

Today's Android devices have plenty of RAM and viable processors: How great would it be to present the software directly on smartphones on an actual Debian system? Maru OS [1] makes this a possibility.

The system created by American developer Preetam D'Souza combines two operating systems in one firmware. Its custom ROM (v0.2.3 at press) for the Nexus 5 (it only works on this device so far) is available for download under a free license, providing both Android 5.1

in the Android Open Source Project version and Debian 8 "jessie." The Debian system only comes into play if the user connects the Nexus 5 to an HDMI-capable monitor via an adapter.

This solution immediately offers several advantages over the classic Android system: It is much easier to use Debian on a larger monitor than Android, which is particularly optimized for small displays and touchscreens. Additionally, basically all applications provided in the Debian project for the ARM platform start on the Debian system. Along with LibreOffice (v4.3), this also includes tools such as Gimp or the popular Apache web server. Even some rather exotic programs like Blender and LaTeX run on the Android smartphone.

Not surprising, however, is that the focus is more on the use of lightweight applications. The Nexus 5 might have plenty of memory with 2GB of RAM, but the built-in Snapdragon 800 CPU proves

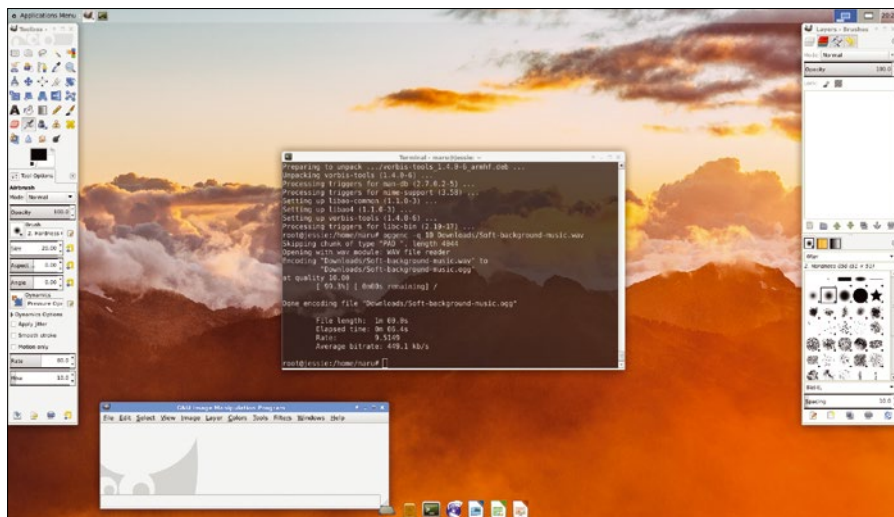


Figure 1: The encoder oggenc converts a 60-second WAV file on Nexus 5 in Debian in 6.4 seconds.

Lead image © watchara rojjanasain, 123RF.com

a bit weak when it comes to processor-intensive tasks.

For comparison, an Intel i5-5200U processor takes about 1.2 seconds to convert a 60-second WAV file (16-bit, 44.1kHz) into Ogg Vorbis format with the encoder call `oggenc -q 10`. The Snapdragon 800, on the other hand, needs more than six seconds to work on this task (Figure 1).

### Download and Installation

Maru was still in a (closed) beta phase, when tested for this article, but it was already easy to use the system on the Nexus 5. Anyone wanting to use Maru on a daily basis still has a few hurdles to overcome. For example, the Android part of Maru OS (Figure 2) can currently only be updated by reinstalling the system. The developer cannot provide Android's usual over-the-air updates. Users need to be flexible because the system is evolving rapidly.

Another drawback is that Maru uses Android 5.1, but Android 6 "Marshmallow" has been available for the Nexus 5 since October 2015. Additionally, the audio routing doesn't seem to work on the Debian system. When tested, the Nexus 5 didn't provide an audio signal

via the HDMI output. The audio output works as usual on Android.

The source code was only available for beta testing when this article went to press, but you can now get a ZIP archive of around 700MB on the website's Downloads page [1]. This includes all the necessary system images and the tools required for the installation such as the Android Debug Bridge (`adb`) and `fastboot`.

The best approach for users is to activate the developer options on the Nexus 5, unzip the ZIP file in a folder, and access the `install.sh` script. If necessary, this will unlock the smartphone's bootloader. Caution: This will delete all the data on the Nexus 5.

The ZIP package contains the installer and an `uninstall.sh` script that, on request, re-equips the Nexus 5 with the original firmware from Google [2]. An unzipped version of the original firmware needs to exist on the computer for this to work. The script still expects the absolute path to the Nexus 5 firmware to then start the installation there. Running the `uninstall.sh` script will erase all data on the smartphone.

### How Does It Work?

Although previous attempts by Motorola, Ubuntu, and other manufacturers used virtualization or emulation and failed because of too little RAM in most cases [3], Maru uses the LXC lean container system from [Linuxcontainers.org](http://Linuxcontainers.org) [4]. The Debian system thus directly accesses Android's kernel v3.4, rather than Jessie's standard kernel (3.16). This kernel is then solely responsible for the complete virtualization. It acts toward the Debian system as if it is a PC based on the ARM architecture.

Separation from the Android system takes place through special kernel control groups (cgroups) that hierarchically assign all LXC container processors to a Linux process. Put more simply: LXC is a cross between a simple chroot and a complete virtualization such as Xen or KVM. It uses fewer resources but can't, for example, simulate any other operating system, because it is still just a Linux kernel.

### Practicality

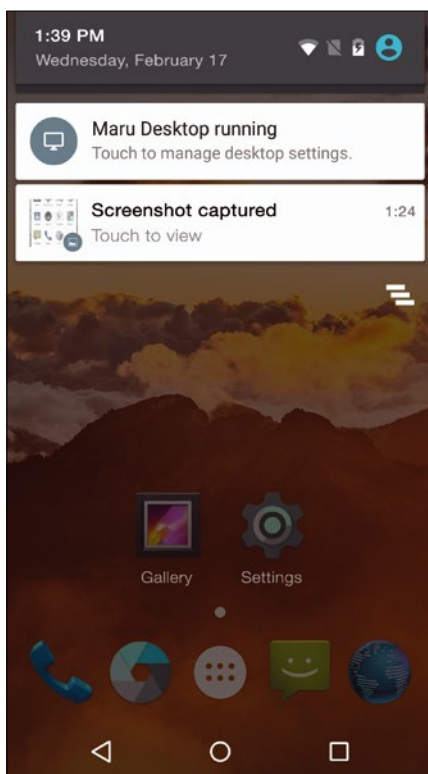
Initially, the Android system is just pure Android (Figure 3). Anyone who also

needs typical Google apps, such as the Play Store or Google Maps, should also download and load onto the Nexus 5 one of the images for the ARM platform and Android 5.1 from Open GApps [5]. The images provided there differ from many Google Apps. F-Droid [6], a free app store for Android apps, is an alternative software source. You can, at the very least, install a useful email client and a handful of practical tools.

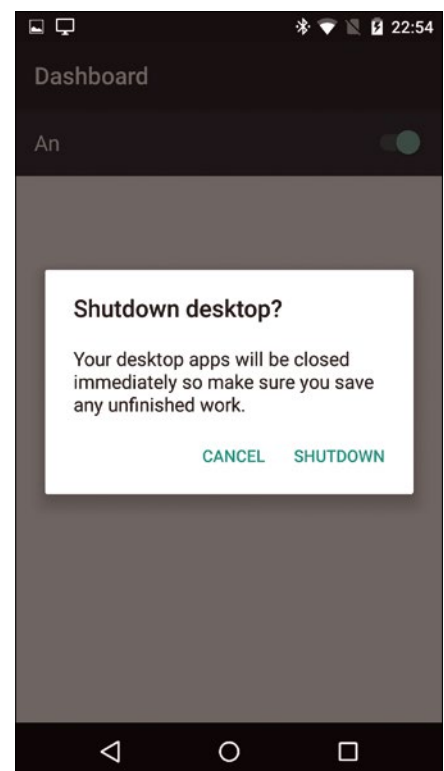
The Debian system is pure Debian with Xfce as a desktop. Unlike a desktop installation, many tools are missing – although users skilled in using Debian can retrofit them using `apt-get`. The lean system has proven to be advantageous. The only restrictions occur when no Debian package is available for ARMv7. Because Nexus 5's Snapdragon 800 isn't an ordinary x86 or x86\_64 processor, it doesn't accept any i386 or x64 command sets.

When tested, no significant problems occurred when using LibreOffice, Gimp, ImageMagick, or several other tools. It was also possible to use Blender, as mentioned in the introduction (Figure 4); however, a powerful PC should probably be used for rendering.

Maru is only suitable on the go when used together with selected accessories. Users need a Bluetooth keyboard or a



**Figure 2:** Maru's Android desktop. The Debian system starts automatically when connected to a monitor.



**Figure 3:** You can shut down Debian via a separate menu item on the Android system.



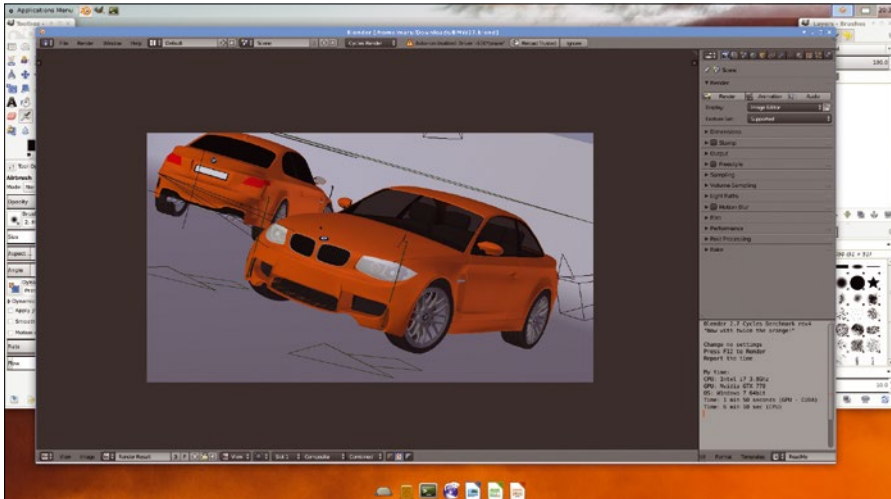


Figure 4: Even resource-intensive applications such as Blender run on the Nexus 5, although much more slowly.

Bluetooth mouse, or both, and a special adapter for converting the Nexus 5's micro-USB port into an HDMI port. Carrying these devices around all the time would be almost as annoying – or even more annoying – than taking a small netbook. However, Maru is a very useful tool for anyone who has a Bluetooth keyboard with an integrated touchpad and the SlimPort adapter.

The question of how to use the application still remains. Monitors with HDMI input abound – almost every modern hotel room has one. Meeting rooms also

often have a projector with HDMI input. Therefore, you could, for example, start a presentation created in LibreOffice natively (provided no sound is needed). Another scenario involves various work in the shell or programming. Whether you use Emacs or Vim, it is easier to work in Debian using a keyboard than on the Android system (Figure 5).

However, Maru users shouldn't simply trust their system with sensitive data, because there are no safeguards to protect the Debian container against unauthorized access. If the Nexus 5 is con-

nected to an HDMI-capable monitor, the Linux system starts without asking for a password. However, you could create a TrueCrypt container in Debian or send encrypted mail using Mutt and GnuPG.

## Conclusions

Maru OS is an exciting project. Unlike many predecessors, the performance of Debian on the Nexus 5 is pleasing. With the additional Linux system, users can complete many tasks that would be impossible with an Android.

Maru is currently still missing multimedia capabilities and comprehensive localization. However, now that the project has released its source code, this is likely to change quickly. ■■■

## INFO

- [1] Maru OS: <http://maruos.com>
- [2] Nexus firmware files: <https://developers.google.com/android/nexus/images>
- [3] Test-driving the Motorola Atrix's Lapdock: <http://www.infoworld.com/article/2623673/android/test-driving-the-motorola-atrrix-s-lapdock.html>
- [4] LXC: <https://linuxcontainers.org/lxc/introduction/>
- [5] Google Apps image: <http://opengapps.org>
- [6] F-Droid: <https://f-droid.org>

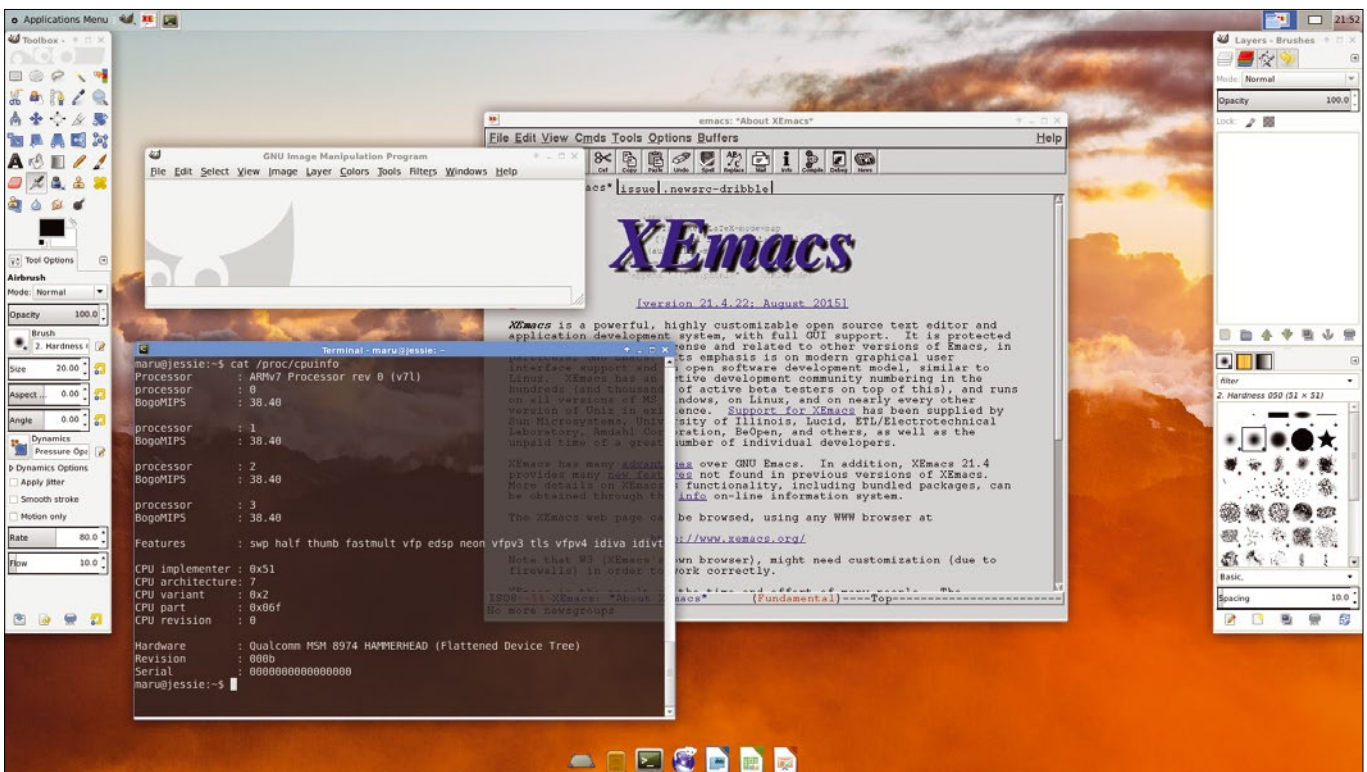


Figure 5: Many other tools are available for Linux-savvy programmers besides Emacs.



# 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**



### Disaster recovery framework

# Unwind

A simple Bash shell script can be a very powerful tool: Relax-and-Recover generates rescue media from running systems, takes care of backups, and helps when migrating computers to new hardware or converting to virtual machines. *By Oliver Hoffmann*

If you are looking to migrate the system and data on your sad, old laptop to a new device or you have a server that's been standing in the corner for years consuming electricity that would be better off as a virtual machine, the simple Relax-and-Recover [1] shell script can help. This disaster recovery framework for Linux licensed under GPLv2 creates rescue disks (USB sticks, CDs, PXE boot images) and saves the data – optionally in cooperation with an external backup solution.

During subsequent recovery of the data, the tool proves to be a flexible partner that also takes hardware changes into account. Relax-and-Recover is thus not just an emergency tool, it is also a useful migration tool. Physical computers thus become virtual machines through Physical-to-Virtual (P2V) migration; the way back to physical hardware with Virtual-to-Physical (V2P) migration is also possible, and you can even migrate from one virtualization solution to another through Virtual-to-Virtual (V2V).

Our test team worked on Debian 8 ("jessie") and Ubuntu 16.04 (Xenial Xerus) with Relax-and-Recover current stable version 1.17.2 and the latest version 1.18 from the GitHub repository (March 2016) [2]. We produced a rescue medium and a backup of a desktop computer. We also backed up a legacy server, combining the script with an external backup solution to do so. After this, the rescue systems and backups formed the basis for converting the physical hardware to virtual machines.

### Well Prepared

Although Relax-and-Recover comes with a number of Linux distributions, the latest version is not necessarily on tap. The developers offer complete packages of the current versions on the project website [1]. You need to take

care to resolve the dependencies on your system yourself. Usually, the package manager points out the missing components.

Alternatively, you can check out the program from the GitHub repository [2]. For this purpose, you need to install the *git*, *syslinux*, *extlinux*, and *binutils* packages and then execute the following commands:

```
git clone https://github.com/rear/rear.git
cd rear
make deb
```

On RPM-based systems, the package-building command is `make rpm`. After installing, *rear* resides in the `/usr/sbin` directory; you need to be root to start the script.

### LISTING 1: Launching Relax-and-Recover

```
$ sudo rear format /dev/sdb
USB device /dev/sdb must be formatted with ext2/3/4 or btrfs file system
Please type Yes to format /dev/sdb in ext3 format:
```

Lead Image © Lucato, 123RF.com

```

root@wily:~
root@wily:~# rear format /dev/sdb
USB device /dev/sdb must be formatted with ext2/3/4 or btrfs file system
Please type Yes to format /dev/sdb in ext3 format: Yes
root@wily:~# blkid
/dev/sda5: UUID="3f5ad96a-5144-4bee-871e-97446a22cdb8" TYPE="swap" PARTUUID="36bb84a4-05"
/dev/sda1: UUID="859eadd7-50f2-491c-b4b0-c1a6d6a5042b" TYPE="ext4" PARTUUID="36bb84a4-01"
/dev/sdb1: LABEL="REAR-000" UUID="24a8047b-e54d-45e9-89f4-21b7606070a7" SEC_TYPE="ext2" TYPE="ext3"
PARTUUID="ad2dec1-01"
root@wily:~# █
    
```

Figure 1: Relax-and-Recover formats the USB medium and writes the label *REAR-000*.

The next task is to format the backup device, which overwrites all the data. Relax-and-Recover comes with its own command to do so; to launch the process, type the command followed by device name of the removable medium, which must not be mounted (Listing 1). It is important to type *Yes* exactly as shown, otherwise Relax-and-Recover will quit.

Formatting is not entirely without problems. On one of the test systems, a USB stick needed a fresh partition table before the medium was accepted. If the software quits with an error, it immediately refers you to the logfile, which contains detailed information about the problem. If all goes well, the tool formats the stick with the filesystem speci-

fied in the output and labels it *REAR-000* (Figure 1).

### Rescue the Desktop!

A backup of your own work computer would be a good first trial of the software. After formatting your USB device, go to the configuration file `/etc/rear/local.conf`, which only contains a few comments by the developers.

Using a text editor, you can define, for example, the following requirements:

```

OUTPUT=USB
BACKUP=NETFS
BACKUP_URL=usb:///dev/disk/by-label/REAR-000
    
```

The rescue `initramfs` that contains the required startup files should end up on the USB medium. Alternative options for `OUTPUT` are `ISO`, `PXE`, `RAMDISK`, or `OBDR` (see the manual [3] for details). Additionally, the type of backup is defined: `NETFS` is the internal Relax-and-Recover method, which creates a

simple `tar.gz` archive of all data on the system. In this case Relax-and-Recover stores the rescue image and the backup itself on the device labeled *REAR-000*.

The next step depends on whether you only need a bootable rescue system or a system with a full backup on top. Listing 2 shows how to create the first alternative. The `-v` option enables verbose message output. In this case, Relax-and-Recover only generates a bootable medium, which you can use to recover the basic environment of your system (e.g., disk layout, drivers, network card configuration), without programs and data, which are then backed up in a second step.

It is much more convenient to let Relax-and-Recover perform the backup task, too. If the USB removable drive is large enough, the command `rear mkbackup` backs up everything to a single medium. If you don't have not enough space there, external hard drives or network drives can be used. The location for the backup is defined by the `BACKUP_URL` variable in the `/etc/rear/local.conf` configuration file. It expects values that begin with `file://`, `nfs://`, `tape://`, or `cifs://`. You can enter authentication information for Samba shares in the `/etc/rear/.cifs` file (again, take a look at the manual [3] for a sample configuration).

After starting the rescue system, a selection menu (Figure 2) appears listing available rescue images and, in a second step, existing backups. You can also decide whether Relax-and-Recover should do everything automatically or if you want to intervene, in which case you log in as `root` without a password. The recovery process is launched by `rear recover`. Again, the `-v` switch ensures detailed output.

In the course of the restore, the program creates the partitions and reformats them with the original filesystem. The swap space ends up at the same location, and the script mounts the device and then restores the data. Finally, Relax-and-Recover installs the GRUB bootloader. Before rebooting the computer, you can see the results of your work in `/mnt/local`, as the script reveals (Figure 3).

### Teamwork

After the reboot, the system will be in the same state it was before calling `rear`

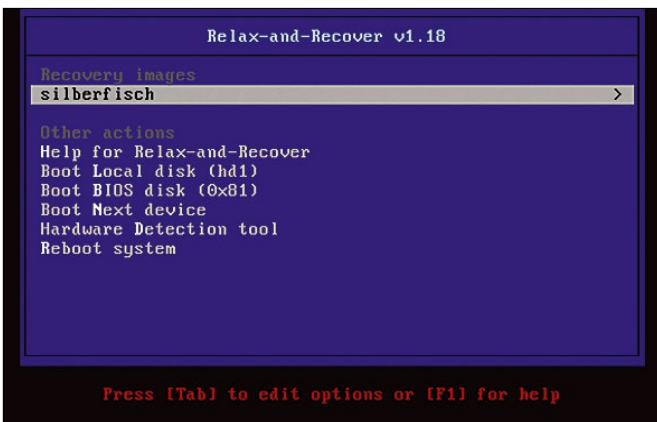


Figure 2: The Relax-and-Recover boot menu showing the available images (e.g., *silberfisch* here ) and alternative boot methods.

#### LISTING 2: Rescue Disk

```

$ sudo rear -v mkrescue
Relax-and-Recover 1.17.2 / Git
Using log file: /var/log/rear/rear-craptop.log
Creating disk layout
Creating root filesystem layout
TIP: To login as root via ssh you need to set up/root/.ssh/authorized_keys or
SSH_ROOT_PASSWORD in your configuration file
Copying files and directories
Copying binaries and libraries
Copying kernel modules
Creating initramfs
Writing MBR to /dev/sdb
Copying resulting files to usb location
    
```



```

4) Go to Relax-and-Recover shell
5) Continue recovery
6) Abort Relax-and-Recover
#? 5
Start system layout restoration.
Creating partitions for disk /dev/sda (msdos)
Creating filesystem of type ext4 with mount point / on /dev/sda1.
Mounting filesystem /
Creating swap on /dev/sda5
Disk layout created.
Decrypting disabled
Restoring from '/tmp/rear.TTuM2IuNo2tKdFk/outputfs/rear/silberfisch/20160319.164
B/backup.tar.gz'
Restored 5438 MiB [avg 43172 KiB/sec]OK
Restored 5438 MiB in 130 seconds [avg 42840 KiB/sec]
Restore the Mountpoints (with permissions) from /var/lib/rear/recovery/mountpoint
permissions
Patching '/proc/1817/mounts' instead of 'etc/mtab'
Updated initramfs with new drivers for this system.
Installing GRUB2 boot loader
Updating udev configuration (70-persistent-net.rules)

Finished recovering your system. You can explore it under '/mnt/local'.

RESCUE silberfisch:~ # _
    
```

**Figure 3:** If you want to check the results of your recover operation before rebooting, look in /mnt/local.

mkrescue or rear mkbackup. Interaction is possible and necessary if the disk size has changed; you need to edit the /var/lib/rear/layout/diskrestore.sh script generated by the restore in a text editor. Thanks to detailed logfiles and the ability to change scripts on the fly during the recover process, you can even move systems to larger or smaller hard disks, although you can run into problems if you move to an enhanced IDE (EIDE) disk.

Relax-and-Recover is also at home in data centers, where it can back up or migrate servers with a significantly larger amount of data storage, compared with a desktop PC. You can even integrate an existing backup solution. Relax-and-Recover works with, among others, Bacula, Bareos, HP Data Protector, CommVault Galaxy (versions 5 to 7, and possibly later versions), Symantec NetBackup, EMC (formerly Legato) NetWorker, and IBM Tivoli Storage Manager.

Teamwork with an external backup program is not just an elegant way to combine the recovery and restore process, it also offers other benefits. The NETFS backup method, with its simple .tar.gz archive, does not take into account running databases and lock files. If you do not want to create a database dump manually and stop processes before running Relax-and-Recover, you will probably prefer a complete solution that takes care of exactly such problems.

The test team combined Bareos [4], a Bacula fork, with Relax-and-Recover and added the following information to the /etc/rear/local.conf configuration file:

```

OUTPUT=ISO
OUTPUT_URL=nfs://<IP>/<Share>/<Hostname>
BACKUP=BAREOS
BAREOS_CLIENT=<Bareosclient-fd>
    
```

The ISO image as a rescue system resides on a computer accessible via NFS. After a server crash, or when migrating to a virtual machine, the ISO boots via

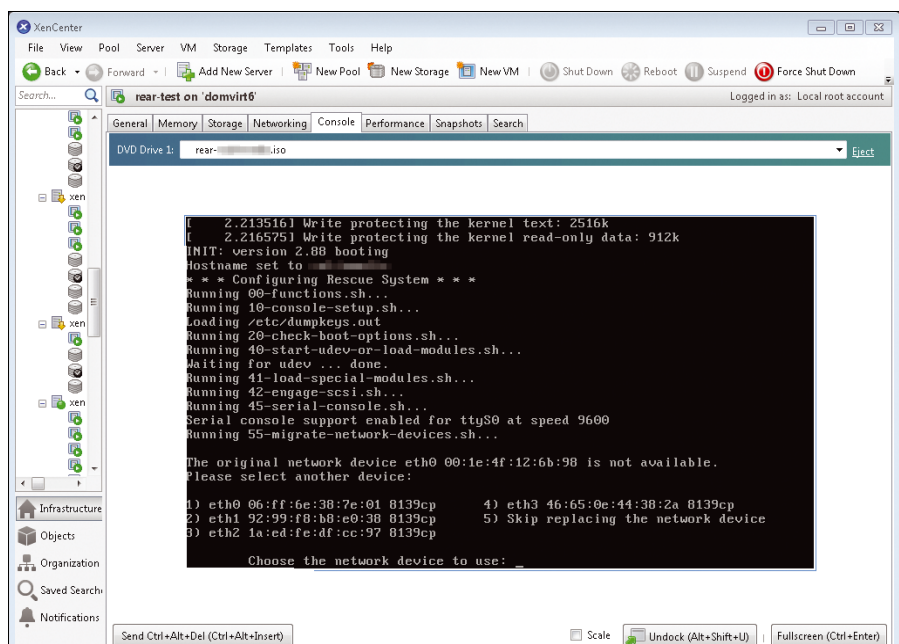
PXE – given an appropriately configured TFTP server. BAREOS\_CLIENT defines the name of the Bareos file daemon, which is responsible for backing up the client data.

The file daemon itself and its configuration are stored in the ISO image; it then autonomously takes care of restoring the data later on. A mandatory requirement, however, is that it access the backup server via bconsole, which is the tool that communicates with the Bareos Director.

One other consideration in this scenario is that everything beyond a minimum rescue system is best stored in the backup, including directories such as /etc, /usr, and /var, but not /tmp, /proc, and the like. This will ensure that the system administrator finds everything again as expected.

## Migration Aid

Relax-and-Recover also helps administrators with the process of moving a system to a different environment – whether physical or virtualized hardware. Transferring a hardware-based server to a virtual machine (P2V) is quite simple, but, as always, the devil is in the details – of the virtualization software, in this case. As a first experiment, the test team converted an application server (Ubuntu 12.04 LTS) with a MySQL database from an old Dell PC with a faulty RAID drive to a virtual machine on a XenServer [5].



**Figure 4:** The physical NIC eth0 is no longer available. You select virtual hardware provided by XenServer and listed here.

Creating the rescue system and the backup did not differ from the steps discussed previously. To boot the image, we created an ISO storage repository (NFS ISO storage) in the XenServer configuration. As the VM template, we chose *Other*; the other templates refused to cooperate.

During the recovery process, Relax-and-Recover asked a few questions about the network card (Figure 4) and the hard disk, because of, first, the modified virtual hardware and, second, the different drive size. We had to explicitly confirm the device file change from `/dev/sda` to `/dev/xvda` and had no problems downsizing the disk.

When migrating an old laptop (Debian 8) P2V, the test team had to trick both VMware [6] and VirtualBox [7]. The rescue system again landed on a USB stick. To boot from this medium, USB 2.0 support had to be activated on the virtual machine. Moreover, we used a special boot manager named Plop [8], because VMware and VirtualBox cannot boot natively

from USB. The subsequent rear recover command then proceeded without any glitches.

### Pure Relaxation

Relax-and-Recover showed hardly a weakness during our tests. Only cloning a laptop with a solid-state disk and the subsequent migration to an older computer with an EIDE drive failed. Even customizing the scripts did not help in this case; a bug report has been submitted. The program is constantly learning new tricks, which makes us optimistic. For example, the upcoming version 1.18 includes support for NVM Express, an interface to connect SSDs via PCI Express, for the first time.

In all of our other experiments, the disaster recovery tool was convincing. With a few simple commands, you can generate rescue systems and then provide them on CD, USB media, or via a PXE boot image. Relax-and-Recover also proves extremely cooperative when it comes to backups. Not only does it have its own minimal backup

routines, it works with known and flexible backup software from other vendors. ■■■

### AUTHOR

**Oliver Hoffmann** works as a system administrator at Digital Online Media (DOM). Thanks to Relax-and-Recover, he not only has full mastery over his own laptop, he can rely on this software for customer systems.

### INFO

- [1] Relax-and-Recover: <http://relax-and-recover.org>
- [2] Relax-and-Recover on GitHub: <https://github.com/rear/rear>
- [3] Handbook: <https://github.com/rear/rear/blob/master/doc/user-guide>
- [4] Bareos: <https://www.bareos.org>
- [5] XenServer: <http://xenserver.org>
- [6] VMware: <http://www.vmware.com>
- [7] VirtualBox: <https://www.virtualbox.org>
- [8] Plop boot manager: <https://www.plop.at/en/bootmanagers.html>

# Linux Magazine on the go

Subscribe to our PDF Edition

- Quick & easy access
- Read PDFs offline
- Keyword Search



**GET YOUR  
FIRST  
3 ISSUES  
FOR ONLY  
CAN & US \$3  
UK £3  
EUR €3**

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

Ft Photography, Fotolia



## Klaus Knopper answers your Linux questions

# Ask Klaus!

By Klaus Knopper

### Gpm and Auto-Login

**?** Dear Klaus:  
It seems like it should be possible using `gpm` or perhaps `mev` to have a mouse click either send a `startx` or `Kodi` command to the shell. For example, the `man` [page for the] `gpm` command indicates that something like

```
gpm -S "kodi:startx:sudo shutdown -h now"
```

could be used to set up the mouse so that the left mouse button starts Kodi, the middle button starts X windows, and the right button shuts down the system, but it must be more complicated than this because this does not work.

The second issue is how to auto-login at the command line without a password to begin with.

Kind regards, Charles

**💡** Indeed, the `-S` special command in `gpm` is really special. Before getting to the point, I'll tell you some more about `Gpm`.

`Gpm` is a "mouse driver" for the Linux text console that allows you to select, copy and paste text, and navigate in ncurses-enabled programs such as `dialog`. For example, you can use the text browser elinks on the text console and click on links or menus with the "text mouse cursor" that's enabled with `Gpm`.

Because no `Xorg` as a graphical interface is running on the text console, `Gpm` has to provide its own drivers. In the case of an old PS/2 mouse, this is done with

```
gpm -m /dev/psaux -t ps2
```

(which I use in `Qemu/KVM` for testing mouse emulation), or

```
gpm -m /dev/input/mice -t evdev
```

for some USB mice. Each mouse type has its own special features and buttons; for example

```
gpm -m /dev/input/mice -t help
```

shows a list of available drivers. Note that `Gpm` launches itself as a background process with (almost) no visible status or error messages. It logs its output to `syslog`, where you can read it using `systemd's journalctl` command, or

```
tail -f /var/log/messages
```

for `SysVinit` Linux distros. So, if `Gpm` fails, you will have to check the log to determine why. Often the error message simply tells to use the correct combination of `-m /dev/mouse-device` and `-t driver`.

Once you have `Gpm` running with the correct driver, you can select text with the left button and paste it with the middle mouse button or use a combination of left and right button as the middle mouse button replacement. Until you get the buttons and mouse movements to work correctly with the `-t` and `-m` arguments, it is useless trying to add the `-S` option.

To make sure no background `Gpm` process is still running, you can terminate `Gpm` by issuing `gpm -k`.

Now the `-S` option comes into play, allowing you to associate commands with special gestures. The short version of the online help may be a little misleading: The commands set with

```
gpm -m /dev/input/mice -t driver Z
-S 'left_button_command:Z
```



### 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](mailto:klaus@linux-magazine.com)

```
middle_button_command=⌘
right_button_command'
```

are NOT executed when simply pressing a button. You first have to activate “command mode” with a special gesture the man page calls “triple clicking the left and right mouse button,” which you can accomplish by holding the left button down and tapping the right button three times. Doing so should produce a message on the console telling you to release all mouse buttons and start the associated actions by clicking on the desired button.

For example, enter:

```
gpm -m /dev/psaux -t ps2 ⌘
-S "echo left:echo middle: echo right"
```

Now press and hold the left mouse button and then triple-click on the right button. Next, release all buttons and click the left, middle, or right mouse button just once. Depending on which button you clicked, the text *left*, *middle*, or *right* should appear. Similarly, the commands you sent in your question would be executed. Only commands that normally work on the text console can be executed successfully by Gpm button gestures. The `startx` command in your example will start the graphical desktop as root (i.e., the user running Gpm).

About your second question: Depending on your Linux distribution and init system, there are `getty` processes running on text consoles `/dev/tty1` through `/dev/tty10` that use the `login` command for authentication. A typical example for SysVinit is given in the init configuration file `/etc/inittab`:


```
1:2345:respawn:/sbin/getty 38400 tty1
2:2345:respawn:/sbin/getty 38400 tty2
```

Change this to:

```
1:2345:respawn:/bin/bash -login >⌘
/dev/tty1 </dev/tty1 2>&1
2:2345:respawn:/bin/su -1 ⌘
-c /bin/bash knopper >⌘
/dev/tty2 </dev/tty2 2>&1
```

In this example, the first line gives you a root shell on console 1, and the second line starts a shell for user *knopper* on console 2.

## Gnome Commander Problem

 Hi Klaus:

Thanks for your earlier help with mounting partitions.

I have been using Gnome Commander 1.2.8.17 in Linux Mint 17.2 Cinnamon 64 bit for some time, being spoiled with dual-screen filesystems for a long time on proprietary systems.

Double-clicking normally opens many file types, but recently my GCMD has stopped opening HTML and related file types. I may have also clobbered the same for some other file types. I have not located any folder or file that I can try to edit to repair my losses to file types, if any exist. Am I likely to guess correctly what to do if I find an editing point?

I have not been able to find any hints how to repair this. A right-click on the lost file type offers an *Open* or *Open with* entry to a fix, but there is no information on how to proceed to repair this feature. As I have many GCMD bookmarks set up, I hesitate to uninstall and reinstall to see if that restores all I had before I managed to kill my “open HTML file type” feature.

Further, what accidental mousing or other actions must I have done to wipe the open (HTML, etc.) file types? Also, will just upgrading to a later version repair my lost file-type opening features? What backups do I need to save my bookmarks and any other GCMD settings I may want to keep?

Also, what would be nice to have are links to where more GCMD guides or manuals may be found if any are available. Searching within forums seems a never-ending process of chasing unrelated tangents.

Thanks in advance for even considering to help.

Dee Shoolingin



The `gnome-commander` program apparently uses the standard mechanism defined for freedesktop.org-compatible desktop systems for associating file types (MIME types) to commands for “opening.” You should find a configuration file in your home directory called `.local/share/applications/defaults.list`, which looks like this:

```
[Default Applications]
application/pdf=evince.desktop
```

```
application/zip=xarchiver.desktop
application/vnd.ms-excel=⌘
openoffice.org-calc.desktop
application/vnd.ms-powerpoint=⌘
openoffice.org-writer.desktop
```


The listed applications are referenced by their corresponding `.desktop` files, not by the actual command names. Those `.desktop` files are usually created during program installation by the package manager in the directory `/usr/share/applications`.

It may be that “opening HTML files” was associated with `iceweasel.desktop`, which has now become `firefox.desktop`, and the old entry is still present in `defaults.list`. You can edit the file with any text editor. A restart of Gnome Commander may be required for change notification.

See also the main documentation for Gnome Commander [1].

If you cannot find the `.local` directory in your file manager, make sure that “hidden” files and directories (those starting with a dot) are set to be visible.

## Live Discs

 Hello Klaus, Could you please tell me how to run Knoppix 7.7 without installing? I could not find instructions in the magazine or on the DVD. I expected a “Read Me” file or something. The two pages in the magazine did not mean much to me (only a beginner to Linux).

I have installed Mint from a burnt ISO disc without too much trouble.

Thanks, Henry Young



Your computer should be set to start from DVD, as for your “Mint” installation disk. Once started from DVD, Knoppix runs Live, including the desktop, and does not require a hard disk installation. However, you can install Knoppix to flash disk, which makes it run about five times faster.

If the Knoppix DVD is not bootable on your computer, please check the surface for visible scratches or defects. You may get a free replacement if the DVD included in the magazine was damaged for some reason. ■■■

## INFO

[1] Gnome Commander docs: <https://gcmd.github.io/doc.html>



### The Sys Admin's Daily Grind: Etckeeper

# Fountain of Youth

Configuration files change over the life of almost every Linux system, but not always for the better. Sometimes individual /etc files literally degenerate. Don't despair: You can restore a state from better days. *By Charly Kühnast*

**E**tckeeper is part of many distributions and is otherwise available from GitHub [1]. Of course, Git also needs to be installed on the computer. Anyone who hasn't used Git on their system should configure a few fundamental settings:

```
git config --global user.name "Charly"
git config --global user.email "charly@example.com"
git config --global core.editor "vim"
```

Because I execute the etckeeper and git commands with sudo, these settings prevent the root user from being entered as a committer.

## Preserving Youth

To begin, I need to initialize the new repository, then I save all /etc files in it for the first time:

```
cd /etc
sudo etckeeper init
sudo etckeeper commit "Initial etc commit"
```

Some distributions, such as Ubuntu, execute this step automatically when installing Etckeeper. In this case, I see output like:

```
> On branch master
> nothing to commit, working directory clean
```

Sometimes I don't want all the data in /etc/ in the repository, so I tell Git which directories Etckeeper should ignore in the .gitignore file right after the # end section [...] etckeeper line:

```
# end section managed by etckeeper
ghostscript/*
```

Now, however, all the data from /etc/ since the first commit is already in the repository, so I have to remove the undesired files manually:

```
sudo git rm -r --cached ghostscript/*
```

To demonstrate, I change a little something in /etc/postfix/main.cf; in fact, sudo git status shows that Git notices

the difference when comparing the data with the repository (Figure 1). I can now check in the new version:

```
sudo git commit -a -m "main.cf changed"
```

An overview of all actions is provided with sudo git log command:

```
commit 9695e06a8175bd5cf485316f20d8fb2
6d6fcc1e49
Author: Charly <charly@example.com>
Date: Wed May 11 14:18:50 2016 +0200

    main.cf changed
```

It's important to realize that this process implements versioning – there's no backup! Etckeeper and Git help you comprehend changes in configuration files and roll them back if necessary, but that's a lot for anyone who's configured something incorrectly. ■■■

## INFO

[1] Etckeeper: <https://github.com/joeyh/etckeeper>

```
On branch master
Changes to be committed:
  (use "git reset HEAD <file>..." to unstage)

    modified:   .etckeeper

Changes not staged for commit:
  (use "git add <file>..." to update what will be committed)
  (use "git checkout -- <file>..." to discard changes in working directory)

    modified:   postfix/main.cf
```

Figure 1: If someone has configured something in /etc, Git notices the difference when comparing with the repository.

## 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.



# RISK-FREE TRIAL!



Practical.  
Technical.  
Elegant.

**SUBSCRIBE NOW**

**3 ISSUES  
+ 3 DVDs**

**\$3**

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

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





### Switch power outlets on and off with home automation tools

# Click Clack

The “If This Then That” web service lets you simply click together home automation components, instead of laboriously integrating them via their APIs. However, you still have to open the hood for the nitty gritty where applicable. *By Michael Schilli*

**W**hat’s that? There’s a gathering storm, and suddenly, as if by magic, the outside shutters on the house are pulled in? Oh, this must be home automation at work! Unfortunately, it’s not that easy to integrate individual compo-

nents in a smart home into a coherent overall system; many manufacturers like to do their own thing and deliberately don’t play well with gadgets manufactured by the competition.

#### WeMo

I recently bought the WeMo Switch [1] on a whim (Figure 1). It can be used to switch electrical devices on and off, both mechanically and by using the mobile phone app either via local WiFi or over the Internet.

The remote switch generates its own WiFi signal with the identifier *WeMo-xxx* after plugging it into the wall socket. After downloading the iOS or Android mobile app (Figure 2) onto your smartphone and adjusting your WiFi setting to point to the temporary WeMo wireless signal, the app connects to the switch and prompts you to enter the SSID and password of your home WiFi.

The application sends the data to the minicomputer in the switch, which then accesses the Internet via the home WiFi and automatically logs into the Belkin WeMo service. Neither an email address nor registration is required – ev-

erything works anonymously (well, except that WeMo might know your home WiFi credentials).

After a successful setup, the app communicates with the WeMo switch either via WiFi or, if the mobile phone is outside its range, via the service provided by WeMo’s parent company, Belkin, on the open Internet (Figure 3).

Tapping a green button on the app to deploy the switch gets old pretty quick, but you also have the option to flip it on or off programmatically in a scripting language, because who in their right mind can afford to take their hands off the keyboard just to turn on the light? In all seriousness, only by connecting the switch with sensors in the house and using a bit of programming logic can you create a “smart home” that, for example, turns on the outside garage light when it’s both dark and your smartphone’s Bluetooth signal is approaching.

#### On/Off Command

The WeMo switch runs a small web server connected to the home WiFi, so a simple Perl script (e.g., Listing 1) [2] can connect to and operate the mechanical switch with HTTP requests. The program uses the `Power::Outlet::WeMo` CPAN module, whose constructor receives the IP address of the switch and then communicates with the device’s server behind the scenes as an HTTP client using the `on()` and `off()` methods.

#### MIKE SCHILLI

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





Figure 1: The WeMo Switch powers electrical devices on and off using the WeMo app.

The module also supports the methods `switch()` (for switching to the opposite state) and `query()` (for requesting the current state). The whole thing could be done pretty easily using a simple web client, but the module abstracts the underlying URLs, thus making programming convenient and tidy.

### WeMo Vobiscum?

The wireless router dynamically assigns the switch an IP address. So, to make sure the script will find the device, even after a reboot, it is advisable to configure a static lease in the router by adding the device's MAC address. However, WeMo implements the UPnP protocol, so the script in Listing 2 can also learn its IP address with the CPAN module `Net::UPnP::ControlPoint`.

### LISTING 1: wemo-onoff

```
01 #!/usr/bin/perl -w
02 use strict;
03 use Power::Outlet::WeMo;
04
05 my $lamp = Power::Outlet::WeMo->new(
06     host => "192.168.1.139" );
07
08 $lamp->on;
09 sleep 1;
10 $lamp->off;
```

status. At the time of testing, I had only a single UPnP device on my network; the output in Figure 4 shows, among other things, that the UPnP query found the switch's web server at IP 192.168.1.139.

Because the switch can be actuated via the Internet as well as on the local LAN, the device can get around the local router's NAT firewall by speaking the STUN protocol [3] with a Belkin server on the Internet.

Of course, this isn't exactly safe, because a small error in the implementation could suddenly produce a huge dark net for remote-controlled electrical devices on the open Internet [4]. The operation of critical components such as ground-to-air defense systems, sauna heaters, or dentist drills must therefore be strongly discouraged.

### If This Then That

A serious disadvantage of the device is obviously that it defines its own proprie-

As a search target, the script sets the parameter `st` to the value `upnp:rootdevice`, specifying that it is interested in all UPnP root devices in the network. Using the parameter `mx` and the value 3, it sets the maximum wait time to 3 seconds, until it aborts waiting for a device to report its

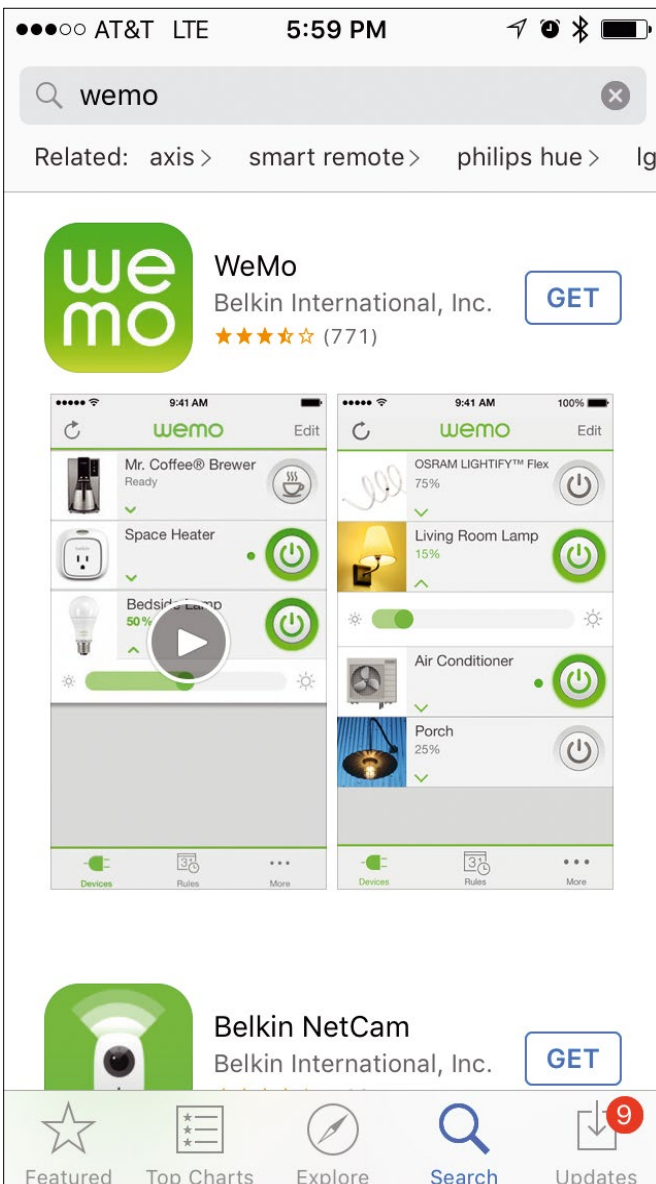


Figure 2: The app controls the WeMo Switch over WiFi or the web.

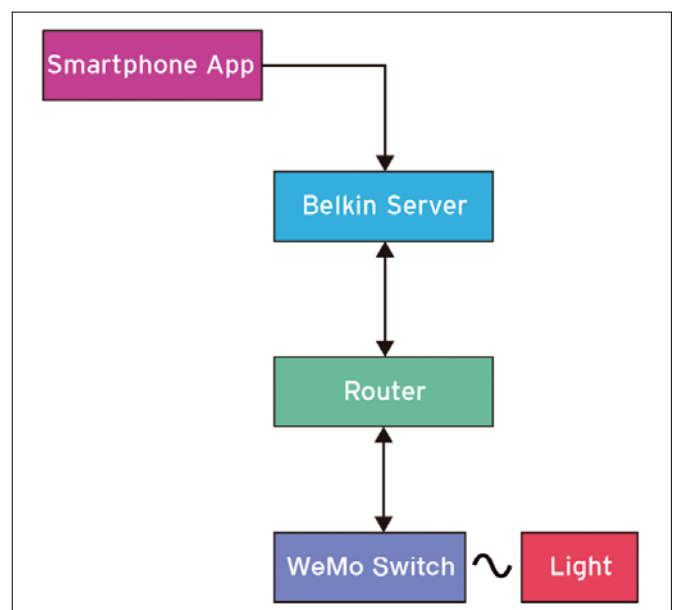


Figure 3: The smartphone app communicates with the WeMo Switch located behind a firewall via a server on the open Internet.



### LISTING 2: wemo-search

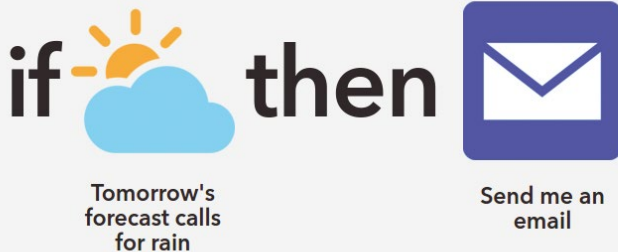
```
01 #!/usr/local/bin/perl -w
02 use strict;
03 use Net::UPnP::ControlPoint;
04
05 my $upnp = Net::UPnP::ControlPoint->new();
06
07 my @devices = $upnp->search(
08     st => 'upnp:rootdevice',
09     mx => 3 );
10
11 foreach my $device (@devices) {
12     print $device->getdevicetype(), "\n";
13     print $device->getfriendlyname(), "\n";
14     print $device->getssdp(), "\n";
15 }
```

```
$ ./wemo-search
urn:Belkin:device:controllee:1
WeMo Switch
HTTP/1.1 200 OK
CACHE-CONTROL: max-age=86400
DATE: Sun, 08 May 2016 17:42:11 GMT
EXT:
LOCATION: http://192.168.1.139:49153/setup.xml
OPT: "http://schemas.upnp.org/upnp/1/0/"; ns=01
01-NLS: 90eb38c6-1de1-11a2-9797-95c8d0e715ca
SERVER: Unspecified, UPnP/1.0, Unspecified
X-User-Agent: redsonic
ST: upnp:rootdevice
USN: uuid:Socket-1_0-221338K010162C::upnp:rootdevice
```

Figure 4: The script launches a UPnP query on the WiFi and finds the WeMo switch at IP 192.168.1.139.

### IFTTT

Make powerful connections with a simple phrase:



Continue

Figure 5: The IFTTT service triggers actions when predefined events occur.

tary communication protocol. Unfortunately, this is the norm among home automation products. As a result, they only work with the individual manufacturer's not particularly original apps, and the open source community has either reverse-engineered the protocols used, had to dive into the de-

tails of another quirky API, or been left out in the cold entirely.

However, fortunately, the trendy Internet company IFTTT [5] has set out to provide a convenient abstraction layer around this protocol confusion. Using the "If This Then That" paradigm, it lets you create logical connections between events. Integrating these events using an API would require a lot of brainpower and developer hours in the real world. Instead, the service generates rules from conditions (*This*), which in turn trigger actions (*That*). With just a few mouse clicks, you can permanently establish these rules.

The *This* clause defines an input channel, such as "At this time" or "When it rains" or "When this YouTube video goes viral." *That* refers to initiated actions such as "Send me an email" or "Post this tweet" or "Turn on this device." The example in Figure 5 illustrates how IFTTT sends email if it's going to rain tomorrow.

IFTTT has all sorts of ready-made channels that trigger an event: Weather Channel events like sunrise, rain, or sunshine are triggered as those natural phenomena occur at the user's location; various security components also trigger an alarm if a sensor in the house reports a break-in. Falling or rising stock prices can also trigger events.

People wanting to use a channel on access-controlled servers like GitHub or Gmail don't need to store a password on IFTTT.com; instead, you are guided properly through the OAuth web flow of these applications. The applications then prompt you to confirm that IFTTT.com may now perform selected actions under your user ID on the specific service. As usual with OAuth, you can withdraw permission at any time. The password never leaves the original service, and IFTTT.com works with tokens of limited validity.

Because programmers are often recruited from the open source scene to code the nitty gritty of the interface between IFTTT and a service (which the end user subsequently just has to select), the channels are consistently very good and implemented with the latest security guidelines in mind.

The arsenal of output actions triggered by events includes email and text messages, push notifications, on/off

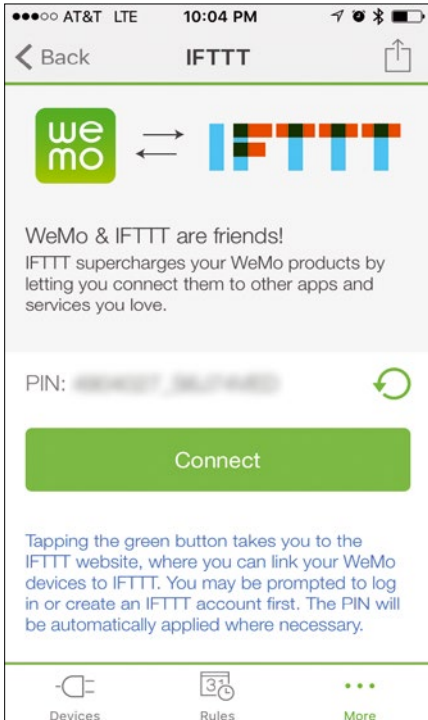


Figure 6: The WeMo app generates a PIN ...

switches like WeMo, and more. The possible combinations of inputs and outputs, called “recipes,” therefore appear to be almost unlimited, and the contributing developers imagination assumes astounding levels.

### Setting Rules with a Mouse Click

Communication between the smartphone app and the Belkin server can only be achieved by use of an SDK, but the WeMo app already has built-in integration with IFTTT.com. If you want to switch the WeMo on and off using IFTTT, simply subscribe to the WeMo Switch channel and provide a PIN to

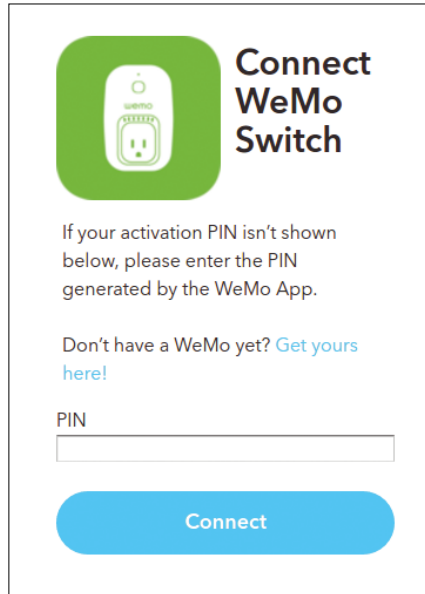


Figure 7: ... which IFTTT accepts and can use to control the WeMo Switch remotely.

IFTTT.com, which the WeMo app generates on request in the Settings menu (Figures 6 and 7).

Now, IFTTT has the keys to the kingdom of remotely controlling the WeMo switch. In conjunction with a triggering input channel, it can now switch on local electrical devices when predefined events occur, as well as with uncommon events, such as rising share prices, or the arrival of tweets.

### Alarm for GitHub Issue

If you have a project on GitHub, you can now simply create a new IFTTT recipe to switch on a light or siren each time a new GitHub issue is created in your project. To do so, you need to subscribe to the GitHub channel and the WeMo Switch channel and then link the

*New Issue* section of the GitHub channel as an input to the WeMo Switch channel as an output action (Figure 8).

The IFTTT server needs to continuously check many such recipes, so once the input is initiated, it might take a while for the output to be acti-

vated. In my tests, the delay between creating a GitHub issue and the switching action was about two minutes. Nevertheless, it is fascinating to see how easy this kind of link can be made and how many possibilities for home automation are opened this way.

### SmartThings Integrated

A startup company called SmartThings [6] in Silicon Valley has stepped forward to unite the more or less proprietary home automation protocols such as Z-Wave [7], ZigBee, and WeMo in the “SmartThings Hub” and to expand into new dimensions with open protocols and an actively integrated developer community. This startup company now belongs to Samsung and the platform is still rather unstable, but it’s a move in the right direction.

I have been using the SmartThings Hub at home for a while now and can confirm that it still has its quirks and share of buggy software. It has been available in the US for some time, and rumor has it that it might be offered in the UK soon. I’ll keep experimenting with it, and I will be writing about it here eventually. ■■■

### INFO

- [1] WeMo Switch at Amazon.com: [http://www.amazon.com/WeMo-Switch-Enabled-Control-Appliances/dp/B00BB2MMNE?ie=UTF8&keywords=Belkin%20Wemo%20Home%20Automation%20Switch&qid=1465308400&ref=sr\\_1\\_1&sr=8-1](http://www.amazon.com/WeMo-Switch-Enabled-Control-Appliances/dp/B00BB2MMNE?ie=UTF8&keywords=Belkin%20Wemo%20Home%20Automation%20Switch&qid=1465308400&ref=sr_1_1&sr=8-1)
- [2] Listings for this article: <ftp://ftp.linux-magazine.com/pub/listings/magazine/189>
- [3] STUN (Session Traversal Utilities for NAT): <https://en.wikipedia.org/wiki/STUN>
- [4] WeMo flaws: <https://threatpost.com/researchers-find-serious-flaws-in-wemo-home-automation-devices/104300/>
- [5] IFTTT: <https://ifttt.com/recipes>
- [6] Samsung SmartThings Hub: <https://www.smarthings.com>
- [7] “Turn On” by Mike Schilli, *Linux Pro Magazine*, issue 184, March 2016, <http://www.linux-magazine.com/Issues/2016/184/Perl-Z-Wave>

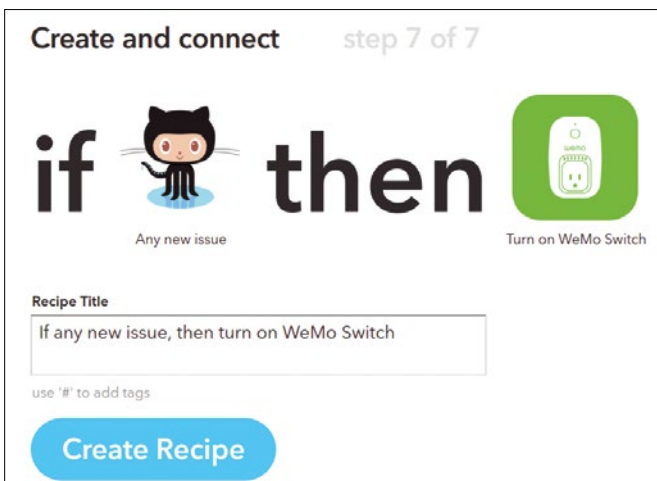


Figure 8: Each new GitHub issue switches on the WeMo Switch.



# Subscribe today and join the revolution!

Each issue of Raspberry Pi Geek is an adventure, with ingenious applications and cool projects for Raspberry Pi, Arduino, and other maker-board systems!

Discover the secrets that will empower you to envision and build your own Raspberry Pi inventions.

## News


The latest developments from the front lines of the Raspberry Pi community, including information on new products and upcoming events.

## Features

Special articles on new technologies and topics of interest to the Rasp Pi community.

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

**\$59.95 £37.50 €44.90** [shop.linuxnewmedia.com](http://shop.linuxnewmedia.com)



RASPBERRY PI GEEK

# raspberry pi GEEK



**24 Rasp Pi Extensions:** The resourcefulness of private and commercial users has led to the creation of many practical and novel extensions for the Raspberry Pi.

**32 Solar Power Management:** A successful solar power project requires data analysis and the ability to modify the system to take advantage of prevailing weather conditions.



## News

- 7 Raspberry Pi 2  
Raspberry Pi turns 3  
Flotilla electronics platform
- 8 Rasp Pi desktop GUI  
PiJuice
- 9 Official Pi Case, WiFi dongle  
RasPiO Duino  
MagPi gets official



## Service

- 3 Welcome
- 6 DVD
- 98 Masthead

## Features

- 10 IR Remote Control  
Turn a Raspberry Pi into an IR remote control for a DSLR, TV, or any other device with an IR port.
- 16 Monitoring Temperature  
Retrofitting a Raspberry Pi with a cooling element does not cost much; however, does the Rasp Pi really need one?
- 20 BitScope Micro  
The BitScope Micro oscilloscope and logic analyzer is a capable tool for electronics enthusiasts.
- 24 Rasp Pi Extensions  
We provide an overview of 20 useful supplementary circuit boards.





# RASPBERRY PI GEEK



## Highlights

**10** IR Remote Control: Turn a Raspberry Pi into an infrared remote control.

**42** Arduino Name Badge: Build a programmable name tag with an Arduino brain.

**54** Monitoring a Nest Box: A nest box complete with camera keeps track of hatchlings.

**73** Automobile Climate Control: An old iPad and an Arduino Mega team up.

## Projects

**32** Solar Power Management  
Tracking the sun and monitoring data to optimize a solar power system.

**42** Arduino Name Badge  
A programmable name badge with a Steampunk theme.

**46** Modding the Robosapien, II  
Hacking a Robosapien robot, Part 2: Control many motors by adding a port expander.

**54** Monitoring a Nest Box  
A camera in a nest box lets you know how the newest hatchlings are faring.

**62** littleBits: callBit  
A call button for a bed-ridden housemate.

**68** LED Shows  
Two LED projects.

**73** Automobile Climate Control  
Create an in-dash climate control app.

## Skills

**80** SwitchDoc Labs: Logic Analyzer  
Use a logic analyzer to isolate problems in your hardware setup.


## Kid Stop

**88** Lego Projects  
Lego Mindstorms NXT plus Rasp Pi equals cool projects.

**92** Scratch: PicoBoard Sensors  
A PicoBoard connects Scratch projects to the physical world.

## Community

**96** Rasp Pi on the Space Station  
We talk with Tim Peake, British ESA Astronaut, about the Raspberry Pis pegged for the International Space Station.



## NOOBS 1.4.0

Kernel 3.18.5  
Updated for Rasp Pi 2

See p6 for details!

RASPBERRY PI GEEK

ISSUE 10

## Skills

Timely tutorials to help you build your skills with Linux and other underlying technologies.

## Kid Stop

Special projects for kids, including a new Scratch programming exercise in every issue.



## Projects

Do-it-yourself, real-world projects that let you learn by doing.

## DVD

Every print issue comes with a free DVD!

## Community

Interviews and reports on Raspberry Pi meet-ups around the world.

Get your Pi on!







## Billing for small businesses and freelancers

# Bookkeeper

Finances and accounting are among the more unpopular tasks in small businesses. With Linux and the free program Yabs, however, you can easily keep track of your finances.

By Erik Bärwaldt

**T**raditional business software is usually primarily aimed at larger organizations and therefore includes many modules that small businesses don't need. The administrative overhead often increases with the complexity of the software, and small businesses in particular are thus reluctant to use such solutions. However, Yabs (Yet another Business Software) provides a lean and ballast-free solution without compromising on operating convenience.

### Installation

Yabs is available for download as an about 37MB zip archive at GitHub [1]. Yabs is a Java application, so you'll need a Java runtime environment to use the program. The business software cooper-

ates with both the proprietary Java environment from Oracle and with the free OpenJDK version. Because OpenJDK is now found in virtually all package management systems of major Linux distributions, it is worth installing it from the software repository of your Linux derivative. If you aren't sure whether a Java environment has already been installed on your system, you can check the status by entering the `java -version` command in the terminal.

Then, unzip the Yabs zip archive into a designated directory using a tool such as Ark or PeaZip. You can then move it to a program directory of your choice later. Third-party programs are typically stored in the `/opt/` directory in Linux. After the subsequent change to the Yabs program directory, you can enable the software by entering the `java -jar yabs.jar` command. No entry for Yabs is created in your desktop's menu tree, so you need to create a corresponding starter – if you want to start the software from the menu via mouse click in the future.

Yabs opens a wizard when the program is first launched. The wizard steps can be completed without changes if you use the software on a single system. Unlike many larger business programs, Yabs doesn't need a separately installed database back end to operate; instead, it works with Apache Derby, a Java-based lightweight database. At the same time, other databases can also be used if necessary. The database structure is automatically created by the installation routine.

Once you have run through the wizard, the actual program window opens. The window is set up very intuitively: In a horizontal menubar at the top of the screen, you will find a button bar that provides quick access to some important features. In the two-part main area of the window, there are vertically arranged buttons on the left that are assigned to the groups *Contacts*, *Accounting*, *Products*, and *Extras* depending on the context. If you click on one of the buttons, a large input and overview area opens on the right-hand side of the window (Figure 1).

### Basic Configuration

To perform the initial configuration with your data, select the *Tools | Settings* menu. All the options are now shown

### AUTHOR

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

as colored icons in the *Settings* tab (Figure 2).

The first step is to enter your company data in the *Company Information* menu. After completing this work, click the *Save* button on the bottom right and then click the green button to get back to the main view of the settings menu. No other configurations need to be performed at this point, meaning you can start with creating your databases.

## Master Data

The next step involves recording all the master data required for the most extensive automated accounting possible. To create records for your customers, click the *Customer* button on the left in the program window in the *Contacts* group and then, after *Add*: at the top in the middle, the *Customer* button on the right in the contact list.

Next, enter your customer's basic data in the gray-shaded entry dialog that opens in the *Addresses* tab. Click the *Contact* tab to enter contact information such as phone and fax numbers. You can enter more data in the *Account Data* and *Info* tabs (free text can be entered in the *Info* section), and you can either print it or just use it internally.

We spotted an inconsistency in the *Account Data* segment: As usual, the account number and bank sort code is asked for here in all input dialogs in the *Contacts* groups. Because this data has now been replaced by IBAN and BIC, program developer Andreas Weber told us that the BIC number needs to be entered in the *Bank Sort Code* field and the IBAN in the *Account Number* field. Weber informed us that the field names would be changed to match the new guidelines in future versions.

You can also create your own keys in the *Settings* tab (e.g., payment conditions), which only apply to current customers, however. Once you've completed your entries, save the record by clicking the disk icon in the button bar in the top-left corner. If necessary, you can add different contacts or additional addresses under the primary address by clicking the *Add Address* button and then filling out the dialog that opens. The new address is then stored on the far right in the program window next to the address field in its own vertically displayed tab.

You can record your product range in a similar way: To do this, click the *Products* button on the bottom left and then the *New Article* button. An input dialog now opens on the right-hand side of the program window where you can record your desired products. You can also enter a large amount of free text in the input dialog and add a product image at the bottom left. The software also allows automatic calculations using the corresponding input fields and tabs (Figure 3).

Yabs also lets you enter services as a product. To do so, use the *New Service* button on the left in the program window. This opens a similarly structured entry dialog where you can enter all the necessary data. You can again save the data by clicking the disk icon.

As part of the conventional product management system, the program also

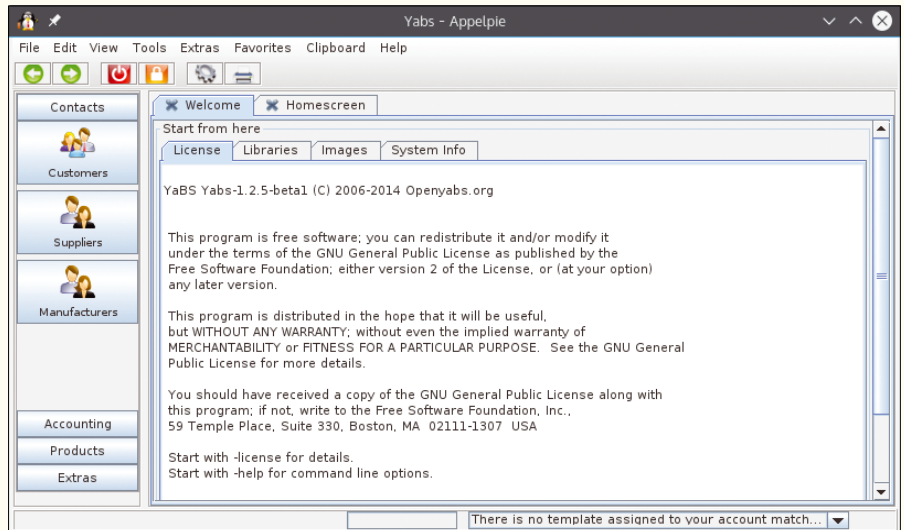


Figure 1: The Yabs program window is clearly structured.

and then, after *Add*: at the top in the middle, the *Customer* button on the right in the contact list.

Next, enter your customer's basic data in the gray-shaded entry dialog that opens in the *Addresses* tab. Click the *Contact* tab to enter contact information such as phone and fax numbers. You can enter more data in the *Account Data* and *Info* tabs (free text can be entered in the *Info* section), and you can either print it or just use it internally.

We spotted an inconsistency in the *Account Data* segment: As usual, the account number and bank sort code is asked for here in all input dialogs in the *Contacts* groups. Because this data has now been replaced by IBAN and BIC, program developer Andreas Weber told us that the BIC number needs to be entered in the *Bank Sort Code* field and the IBAN in the *Account Number* field. Weber informed us that the field names would be changed to match the new guidelines in future versions.

You can also create your own keys in the *Settings* tab (e.g., payment conditions), which only apply to current customers, however. Once you've completed your entries, save the record by clicking the disk icon in the button bar in the top-left corner. If necessary, you can add different contacts or additional addresses under the primary address by clicking the *Add Address* button and then filling out the dialog that opens. The new address is then stored on the far right in the program window next to the address field in its own vertically displayed tab.

You can record your product range in a similar way: To do this, click the *Products* button on the bottom left and then the *New Article* button. An input dialog now opens on the right-hand side of the program window where you can record your desired products. You can also enter a large amount of free text in the input dialog and add a product image at the bottom left. The software also allows automatic calculations using the corresponding input fields and tabs (Figure 3).

Yabs also lets you enter services as a product. To do so, use the *New Service* button on the left in the program window. This opens a similarly structured entry dialog where you can enter all the necessary data. You can again save the data by clicking the disk icon.

As part of the conventional product management system, the program also

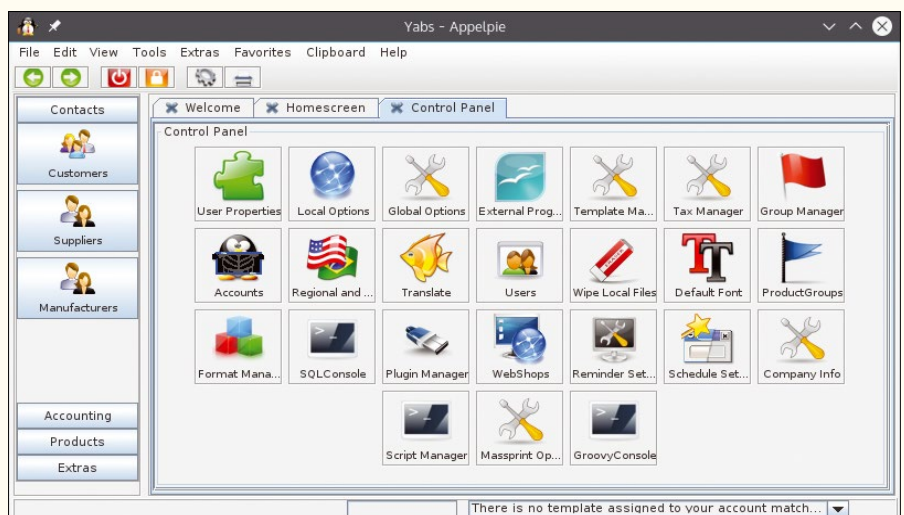


Figure 2: You can access the settings dialog via the colored icons.



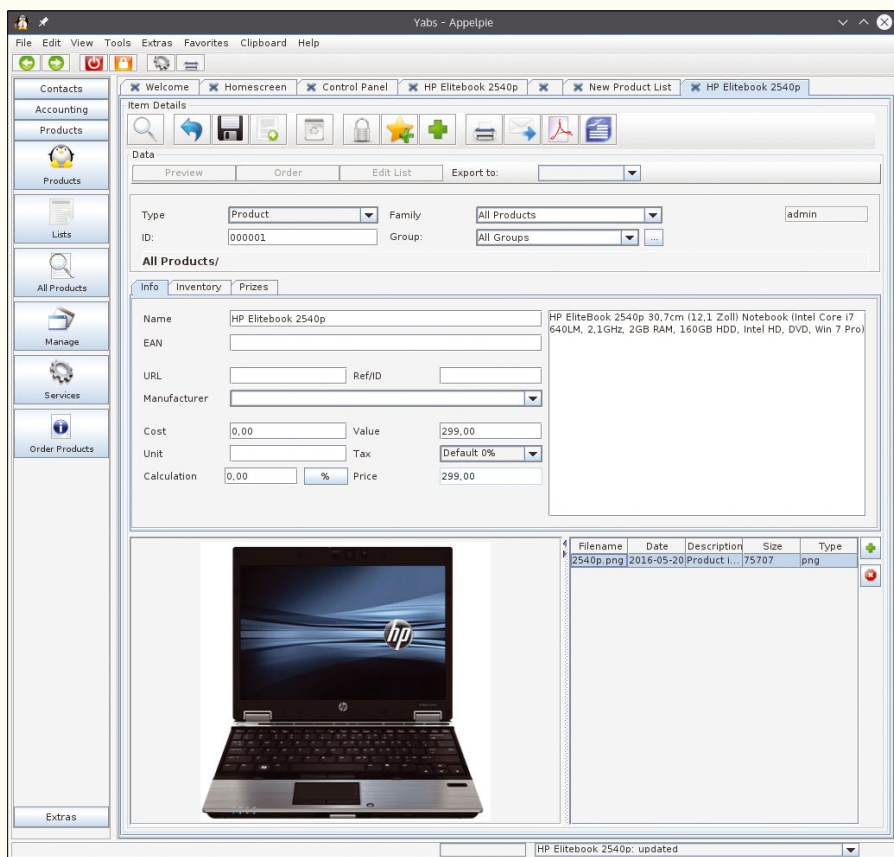


Figure 3: Extensive descriptions can be entered in the product screen.

The easiest way to create operations is to access the corresponding entry dialog using the *New Offer*, *New Job*, or *New Invoice* buttons. Yabs then opens a new entry page in a separate tab where you can enter the corresponding data. Free text input is supported here; however, you can also just use the relevant data from the article master by entering an article number and then clicking on the *Search* button. Additional information, such as free text or images, can be integrated in the input screen using the *Additional info* tab on the input screen. Once you have fully filled out the form, save it by clicking the disk icon (Figure 4).

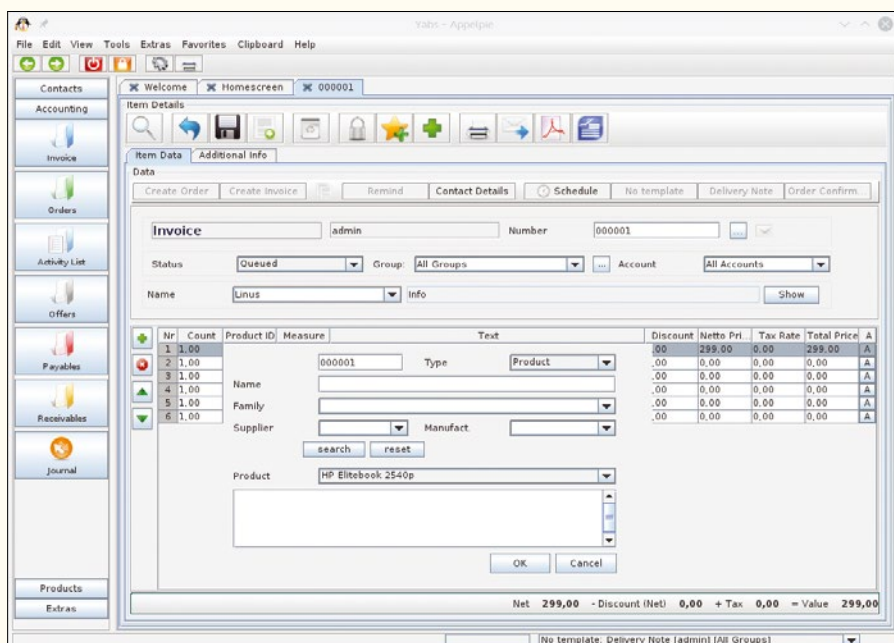


Figure 4: You can easily adopt articles from the article master when creating the operations.

makes storage easier. Yabs can, for example, warn you if inventory falls below minimum values for certain articles. To set this up, open the *Inventory* tab in the window for the respective product and, in the *Inventory Management* section, enter the current inventory and the warning threshold. At this point, the software should inform you whether there is not enough inventory for the particular item. You can also define a supplier for the article in this status window, although the supplier can also be taken from the stored data in the *Contacts* group via the selection box. The software can also handle multiple suppliers for an article.

## In Writing

The software bundles the accumulated transactions in the *Operations* group. Here you will find all the relevant forms from quotation to invoicing that are required in business transactions. A journal function is also implemented, which provides a quick overview of business performance during a period that you can select. An output table in the *Operations* group can be used to record outputs ordered by accounts and to manage them.

## Printing Securely

In order for you to print the individual stored procedures, they need to be connected to appropriate forms. Additionally, you need to define the desired printer in the separate settings dialog.

The program developer provides several Yabs templates for different types of documents that you can download from the project site. These templates are available in ODT and PDF format, and all of them are in German. However, you can customize the templates using Libre/OpenOffice. After downloading the templates, extract the zip archive with the templates and then load the desired template into Writer (Figure 5).

Once you have customized your template, you can link it to the appropriate form. To do so, select the *Manage templates* entry in the *Tools | Settings*

menu and click the *Import file* button in the middle at the bottom to integrate the corresponding template. After selecting the desired template in the file manager that opens, you can define which form the template should be linked to in the *Type* section of the settings windows.

The printer to be used must be defined in the *Printer name* field. If you click the *Search* button to the right of the name field for the printer, a list with all the printers connected to the computer or that can be accessed in the network appears. You just need to select the desired printer using the mouse and then close the dialog by pressing *OK*. Then, save the modified template by clicking *Save* at the bottom right.

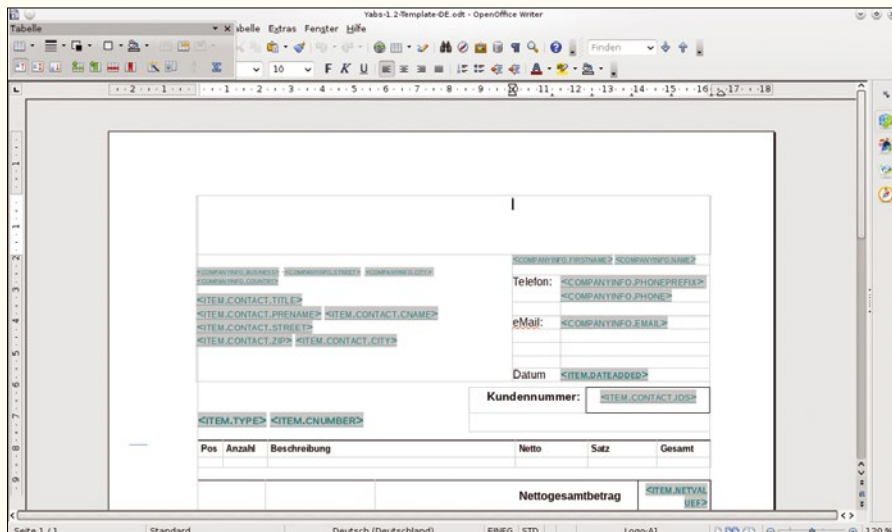


Figure 5: Manage templates in LibreOffice or OpenOffice.

If you want to check the respective form before printing, you can view the document by clicking the *Preview* tab at the top right. However, for this purpose, the first line in the respective list view must be marked. You can also export the new document as either an ODT file or a PDF file by clicking the corresponding buttons at the top right of the main window. Yabs then opens an export dialog box, where you specify the storage path. The document is then saved in the specified location; the respective form is, of course, also stored in the internal database of Yabs.

Once all the preparatory work has been completed and the document meets your expectations, you can trigger the printing process by clicking the printer icon either at the top right in the program window's button bar or in the preview window. The software then opens a printer dialog, where you can, for example, select certain pages to be printed in multipage documents.

## Alarm!

Unfortunately, not all customers are well known for paying their bills on time. This is why inventory management and accounting programs should also always include open item management for automatically monitoring incoming payments. To this end, Yabs has integrated deadline monitoring, which automatically opens when the program starts. Here, in the *Home screen* tab, you can view open items, open orders, and open offers, and you can generate recurring bills for a certain period at the bottom.

Yabs has integrated a small calendar management tool where you can enter both one-off and recurring events. You can add new entries in a separate dialog by clicking the *New event* button at the bottom right under the event list. The calendar and the recorded events are listed on the right in the application's main window. Thus, you always have an overview of upcoming activities as well as those from the past.

## Conclusions


Yabs did particularly well in the brief test thanks to its practicality for small businesses and freelancers who just want to map daily transactions and don't need sophisticated accounting modules. No complicated installation and external databases are required for Yabs. The user interface might take a bit of getting used to, but once you've come to terms with the logic of the program, you'll quickly produce results with Yabs.

Despite the Java substructure, the software works swiftly and stably. The existing documentation needs revision, and there is currently no documentation in English. However, some aspects of the software, such as form customization using the templates, are not good enough. Thus, you should try a few exercises with the software and adapt the forms before starting to work with them. ■■■

## INFO

[1] Download Page: <https://github.com/anti43/openYabs/releases/tag/release> (in German)





## Sandboxing with Firejail

# Secure Play

Firejail makes sandboxing as easy as typing eight letters in front of a command. *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>.

**S**andboxing [1], or isolating processes and applications in their own environment, is a long-established practice in Linux. Unfortunately, although it is efficient, it can be difficult to configure and use. Even containers and virtual machines have not improved the process much, because they are only as secure as their configuration. What makes Firejail [2] so different is that it makes sandboxing easy and can do far more if you are willing to learn how to configure it. This simplicity has made Firejail the center of attention in less than a year.

Firejail is a structural security solution; it is configured to prevent intrusions rather than react to them the way an anti-virus program does. Instead of adding daemons and other applications, it works by creating a restricted environment with its own set of solutions, running within user space and using features that are already a part of the Linux kernel, such as seccomp-bpf [3]. The result is sandboxing that requires far fewer system resources than traditional solutions, such as creating a chroot jail [4], and is easy to customize.

In fact, Firejail installs with 64 security profiles for popular applications, ranging

```
bb@nanday:/etc/firejail$ ls
audacious.profile      fbreader.profile      mathematica.profile    spotify.profile
bitlbee.profile        filezilla.profile     midori.profile        steam.profile
chromium-browser.profile  firefox.profile      mupen64plus.profile  telegram.profile
chromium.profile       generic.profile       nlocal.net            thunderbird.profile
clementine.profile    gnome-mplayer.profile  opera-beta.profile    totem.profile
conkeror.profile       google-chrome-beta.profile  opera.profile         transmission-gtk.profile
deadbeef.profile       google-chrome.profile  parole.profile        transmission-qt.profile
deluge.profile         google-chrome-stable.profile  pidgin.profile       uget-gtk.profile
disable-common.inc     google-chrome-unstable.profile  qbittorrent.profile  unbound.profile
disable-devel.inc     hexchat.profile       quassel.profile       vlc.profile
disable-mgmt.inc      icecat.profile        rhythmbox.profile     webserver.net
disable-secret.inc     icedove.profile       rtorrent.profile     weechat-curses.profile
dnscrypt-proxy.profile  iceweasel.profile     seamonkey-bin.profile  weechat.profile
dropbox.profile        kmail.profile         seamonkey.profile     whitelist-common.inc
empathy.profile        login.users           server.profile        wine.profile
evince.profile         Mathematica.profile    skype.profile         xchat.profile
```

Figure 1: Firejail ships with security profiles for common Linux applications, plus a default profile for anything else.

from Firefox and KMail to XChat and Wine, as well as a generic profile used automatically for applications and processes that lack a custom profile (Figure 1). All profiles use a basic syntax with one item per line that can easily be learned by studying the profiles that are installed along with the command (Figure 2).

At the most basic level, running Firejail could not be simpler. Running the command by itself opens a sandbox of the default shell. To run an application in a sandbox, simply preface the command, using the following standard structure:

```
firejail COMMAND COMMAND-OPTIONS
```

However, you can add options to Firejail to configure the command’s environment more concisely. If you constantly use the same options, you might prefer to customize the command’s security profiles or to create your own profiles. For some purposes, you might even want to maintain multiple profiles of the same command, specifying the one to use with `--profile=FILENAME`.

## Configuring Sandboxes

The default options for running Firejail go far beyond those commonly used by other commands, such as setting the location of the sandbox’s logfile with `--output=LOGFILE`. Many of Firejail’s options show their relation to one another by beginning with the same prefix, a custom that simplifies learning about them.

To start, the protection afforded by a security file can be enhanced by using `--blacklist=` and `--white-list=` to define directories or files that cannot or can be accessed from within the sandbox. Alternatively, blacklisted or whitelisted directories or files can be listed in the profile, each item given on a separate line starting with `blacklist` or `whitelist`.

Other standard behavior can be modified with options. For instance, `-c` closes the sandbox as soon as the command is run once, making it an approximate equivalent of `sudo`. Similarly, `--chroot=DIRECTORY` creates a traditional chroot jail that runs from its own root directory, and `--ignore=COMMAND` cancels the use of one command listed in an existing profile and `-- COMMAND` (with no option letter) skips a command listed in the profile. Although a Bash shell is assumed, you can also specify another shell with `--csh`, `--zsh`, or `--shell=SHELL`. Sandboxed environments can similarly be defined with `--env=NAME=VALUE`.

Still other options can define how a sandbox is used. For example, `--noroot` runs the sandbox with no root user, limiting what you or any intruder can use it for. The `--read-only=DIRECTORY OR FILE` option also restricts what files can be accessed from the sandbox. By contrast, `--write-etc` and `--write-var` give access to key administration files within the sandbox. Another option is to define who can use the sandbox by specifying `--user=USER`.

```
# kmail profile
noblacklist ${HOME}/.gnupg
include /etc/firejail/disable-mgmt.inc
include /etc/firejail/disable-secret.inc
include /etc/firejail/disable-common.inc
include /etc/firejail/disable-devel.inc
blacklist ${HOME}/.pki/nssdb
blacklist ${HOME}/.lastpass
blacklist ${HOME}/.keepassx
blacklist ${HOME}/.password-store
blacklist ${HOME}/.wine
caps.drop all
seccomp
protocol unix,inet,inet6,netlink
netfilter
noroot
tracelog
```

Figure 2: Instead of using command-line options, you can create custom profiles, with one option per line. You can study the predefined profiles to learn how to create your own.



```
bb@nanday:/etc/firejail$ firejail --list
2195:bb:/usr/bin/firejail iceweasel
3718:bb:firejail chromium
bb@nanday:/etc/firejail$ firejail --tree
2195:bb:/usr/bin/firejail iceweasel
  2196:bb:/usr/bin/firejail iceweasel
    2197:bb:iceweasel
3718:bb:firejail chromium
  3719:bb:firejail chromium
    3720:bb:/usr/lib/chromium/chromium
      3725:bb:/usr/lib/chromium/chrome-sandbox /usr/lib/chromium/chromium --type=zygote
      3726:bb:/usr/lib/chromium/chromium --type=zygote
      3728:bb:/usr/lib/chromium/chromium --type=zygote
      3761:bb:/usr/lib/chromium/chromium --type=renderer --enable-features=use-new-media-cache<use-new-media
```

Figure 3: Viewing processes in list and tree views.

Because a sandbox is isolated from the main system, its networking capabilities can be defined separately. Firejail includes a host of options, such as `--ip=ADDRESS`, `--hostname=NAME`, `--dns=ADDRESS`, and `--net=ETHERNET`, all of which can be configured separately from the main system. Other related options can be identified by their prefix, so that IP addresses, for example, can also be defined by `--ip=NONE` and `--ipv6=NAME`. These options are too numerous to detail, so consult Firejail’s lengthy man pages for more information.

A sandbox can also be limited by the maximum file size that a process can create using `--rlimit-fsize=NUMBER` or by the maximum number of files that can be opened by a process with `--rlimit-nofile=NUMBER`. Processes can also be limited by the maximum number of processes a user can open with `--rlimit-nproc=NUMBER`.

Ordinarily, by definition, a sandbox has no communication with the system’s main directory structure. After all, that is what makes a sandbox secure. However, you view directories and files with `--ls=PID OR NAME` or `--get=PID OR NAME`. Additionally, you can read a list of top processes within sandboxes with `--top`, or you can view all sandboxed processes using `--list` or a hierarchical view using `--tree` (Figure 3). Another option is to overlay the present working directory with the sandbox so that the two can be read as a single directory, using `--overlay` or `--overlay-tmpfs`.

If you are using Firejail for more than a quick security feature, you might also want to use `--name=NAME` to make the sandbox you create easier to identify or for keeping track of multiple sandboxes. If nothing else, a named sandbox can be easily closed by adding the option `--shutdown=NAME`.

### Improving Security Even More

For many users, these options are more than enough to make use of Firejail. However, you can improve the security of Firejail with some advanced options.

One option is to manipulate Linux Capabilities, a kernel feature that divides root privileges into as many as three dozen functions. For example, `CAP_SYS_MODULE` lets you load modules from anywhere, instead of restricting the search path to `/lib/modules`,

whereas `CAP_SYS_BOOT` enables rebooting, and `CAP_SYS_NICE` lets you change the priority of processes and threads. Using Linux Capabilities, you can increase the security of a sandbox by limiting what the root user can do in the sandbox. With `--caps.keep=CAPABILITY, CAPABILITY`, you can specify the capabilities that can be used within the sandbox; conversely, with `caps.drop=CAPABILITY, CAPABILITY`, you can prevent the listed capabilities from being used in the sandbox. You could even ban all capabilities altogether with `--caps.drop=all`, although `--noroot` would work just as well.

Another set of options that begins with `private` separates the sandbox even further from the main system by creating a temporary filesystem that lasts only until the sandbox shuts down. Each of these options creates a particular directory:

- `--private=DIRECTORY` creates a user’s home directory.
- `--private-bin=FILE, FILE` creates the `/bin` directory, adding applications in a comma-separated list.

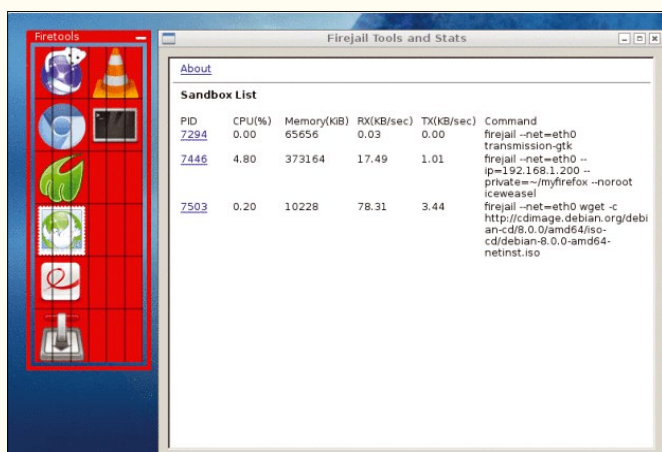


Figure 4: Firejail runs on the desktop as Firetools, but the advantage over the command line is minimal.

- `--private-dev` creates a limited `/dev` directory.
- `--private-etc=FILES, DIRECTORIES` creates the sandbox's `/etc` directory, adding applications in a comma-separated list.
- `--private-tmp` mounts an empty temporary filesystem on top of `/tmp`, so that the two directories display together.

As often as not, you would want to use these options together, so you might prefer to create a private profile rather than adding all these options separately.

A third set of options manipulates `seccomp` [5], the kernel function that controls application sandboxing. Firejail's man page includes almost 50 system calls, including `mount`, `swapon` and `swapoff`, and `reboot`. Using `--seccomp=SYSTEMCALL, SYSTEMCALL` enables the sandbox system calls, whereas `--seccomp.drop=SYSTEMCALL, SYSTEMCALL` specifically disables system calls.

## Running Firejail

Firejail is also available as a desktop application called Firetools (Figure 4). However, Firetools has no advantage over the carefully named options of Firejail, except possibly to divide them up into small chunks of information. I admit that I personally find it confusing compared with the Firejail man page.

Still, no matter how you access these options, Firejail remains a powerful new weapon in the security toolbox. Thanks to the ready-made profiles, it is simple enough to begin using without reading the documentation, except their example. Yet, thanks to the broad range of options and the lower overhead, it is as useful for the sys admins of corporate networks as for a lone user on a workstation.

If you wanted, you could use Firejail for every application. However, your web browser, email reader, and chat applications are probably mostly what you need to sandbox. A convenient way to set them up is to use launchers on the panel or desktop, because the launchers can store a complicated command structure. ■■■

## INFO

- [1] Sandbox: [https://en.wikipedia.org/wiki/Sandbox\\_computer\\_security](https://en.wikipedia.org/wiki/Sandbox_computer_security)
- [2] Firejail: <https://firejail.wordpress.com/>
- [3] "Yet another new approach to seccomp" by Jonathan Corbet: <http://lwn.net/Articles/475043/>
- [4] chroot: <https://en.wikipedia.org/wiki/Chroot>
- [5] seccomp: <https://en.wikipedia.org/wiki/Seccomp>

# Drupal Watchdog

Subscribe to Drupal Watchdog  
Now published 4 times per year!

Drupal Watchdog is THE technical journal for the Drupal community. Each issue brings you the latest news and information written by key contributors from the Drupal community.

Each issue covers

- ★ Coding
- ★ Design
- ★ Workflow
- ★ And much more!

US \$44.95, UK £24.90, EU €34.90

Digisub (PDF) available online.

(Prices for 1-year print subscription. Other regions please see our website.)





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





### Scan Tailor and Paperwork

# Papers, Please

Transform piles of paper into a neatly organized and searchable digital library with Scan Tailor and Paperwork.

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](http://scribblesand-snaps.wordpress.com).

**M**agazine articles, important documents, receipts, and whatnot – paper is still the most commonly used storage and distribution medium, and it’s not going anywhere in the foreseeable future. The problem with paper is that it tends to pile up and take up all available space at an astonishing speed. Worse yet, every time you discard documents you think you no longer need, there is a risk of throwing away something important.

Fortunately, tools like Scan Tailor and Paperwork provide a solution to this conundrum. Using these applications, you can set up an efficient system for turning paper documents into

a searchable library of scanned and cleaned up files.

### Processing Scanned Pages with Scan Tailor

Most mainstream Linux distributions come with a scanning utility pre-installed. Ubuntu, for example, ships with Simple Scan, a no-frills tool that’s more than adequate for all but the most complex scanning tasks. But scanning documents is only half of the battle. In most cases, you might want to clean up and tweak scanned pages, and this is where Scan Tailor [1] comes into the picture (Figure 1). This application is designed for post-processing scanned pages, and it allows you to split and deskew pages, add and remove borders, as well as generate cleaned up files. Keep in mind, though, that Scan Tailor is not a scanning application, so you need to scan pages before you start using the application.

Scan Tailor is available in the software repositories of many popular Linux distributions, so you can install it using the default package manager. On Ubuntu, installing Scan Tailor is a matter of running the `sudo apt install scantailor` command.

After you launch Scan Tailor, you need to set up a new project (a directory containing scanned pages ready for processing and various project-related files). To do this, press the *New Project* button and specify the folder containing scanned pages as the input directory. You should then see a list of all scanned files in the *Files in Project* pane of the *Project Files* dialog window.

At this point, you can exclude certain pages from the project, if necessary. Then, press the *OK* button. If the project files have an unspecified or incorrect dots per inch (DPI) resolution, Scan Tailor prompts you to fix it. In the *Fix DPI* dialog window, select the *All Pages* entry under the *Need Fixing* tab, specify the correct DPI resolution (300x300 is a good starting point), and press *Apply*. Press *OK* to create and open the project.

Although Scan Tailor’s tools are easy to master, you need to understand the overall workflow to use them optimally. Scan Tailor is a batch processor that has several stages: *Fix Rotation*, *Split Pages*, *Deskew*, *Select Content*,



Figure 1: Scan Tailor interface in all its spartan beauty.

Margins, and Output (Figure 2). For each stage, you need to adjust the available settings for each page in the project. When you press the *Run* button, the configured action is applied to the pages.

Here is how this works in practice. Assume you have a project consisting of three scanned pages, and you start post-processing from the *Split Pages* stage. This action can split double-page scans into two separate pages. Select the first page in the thumbnail sidebar to open the page in the working area. Select the *Split Pages* stage, and Scan Tailor should automatically detect the appropriate layout.

If the application fails to do this, press the *Change* button, choose the manual mode, and select the appropriate layout in the *Page Layout* section. To apply this configuration to other pages, press *Change* and enable the *All* option in the *Scope* section. Alternatively, you can specify settings for each page individually by selecting the next page in the thumbnail sidebar and adjusting the available options. When you're done, press the *Run* button next to the *Split Pages* stage to run the action on the pages.

Although the *Fix Orientation* and *Deskew* stages in Scan Tailor can be skipped, other stages must be completed before you can move to the *Output* stage. This last stage is where the application generates processed pages in the TIFF format (Figure 3). Besides the output resolution, you can select



Figure 2: Pages in Scan Tailor are batch-processed in stages.

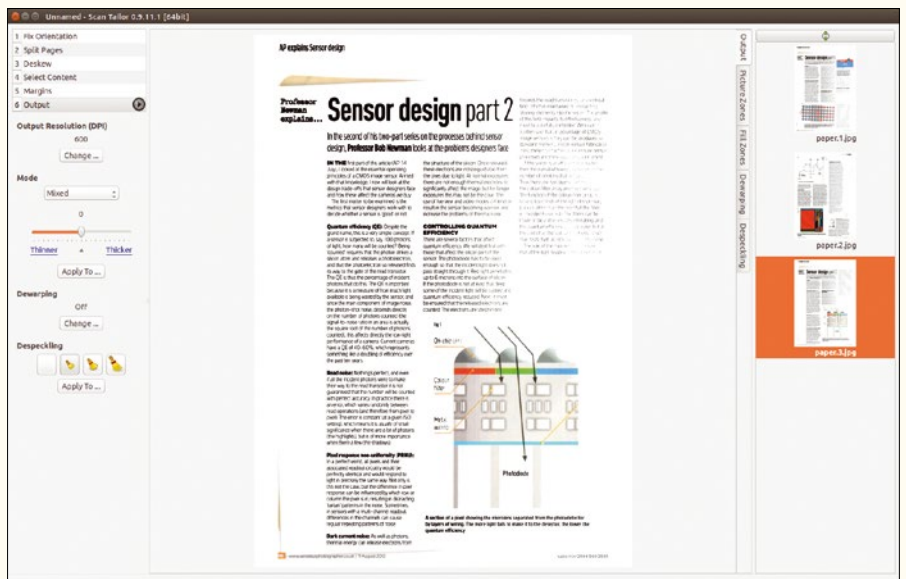


Figure 3: Generating output result.



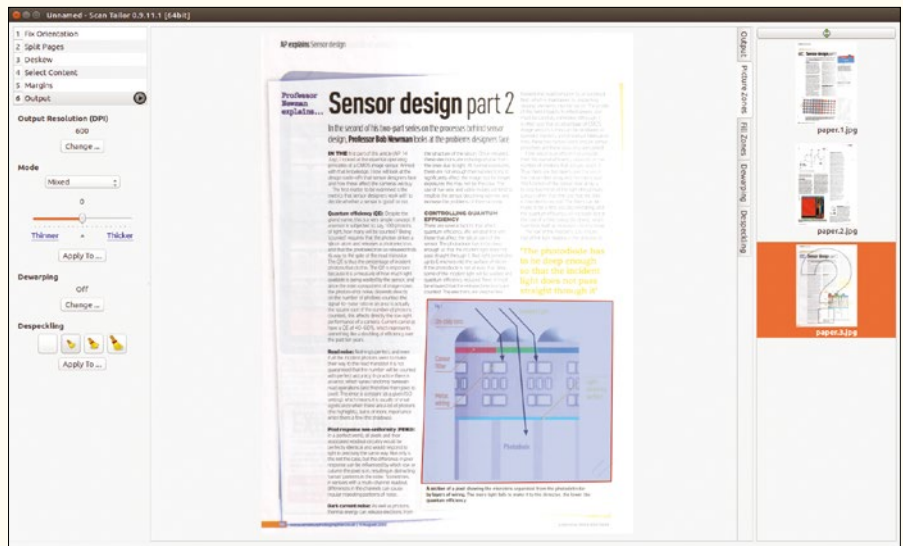


Figure 4: Scan Tailor lets you specify picture zones on a page.

one of three modes: Black and White (generates black-and-white pages), Color/Grayscale (produces color or grayscale pages), and Mixed. The latter can be used to generate pages where text is treated as black-and-white areas, while images are handled as color zones.

In many cases, Scan Tailor does a good job of identifying images, but the application also allows you to specify so-called picture zones manually (Figure 4). To do this, switch to the *Picture Zones* tab on the right side of the main working area. Use the mouse to draw boxes around images on the page. To clean up the page even further and reduce the size of the final file, you can mark large empty white areas on the page under the *Fill Zones* tab. Scan Tailor features two more useful tools for fixing pages.

The Dewrapping tool allows you to manually straighten the page, which can come in rather handy when working with scans of wrinkled or otherwise distorted pages (Figure 5). This tool can be used for perspective adjustment, too, so it's perfectly suited for processing images taken with a mobile device. Finally, the Despeckling tool can clean up the page by removing small straight artifacts. Once you've configured the output settings, press *Run* next to the *Output* stage to generate processed TIFF images in the specified output folder.

At this point, Scan Tailor's job is done. You need to assemble multiple pages into a single document, run them through an optical character recognition software, and perform other operations. However, those are beyond the scope of Scan Tailor's

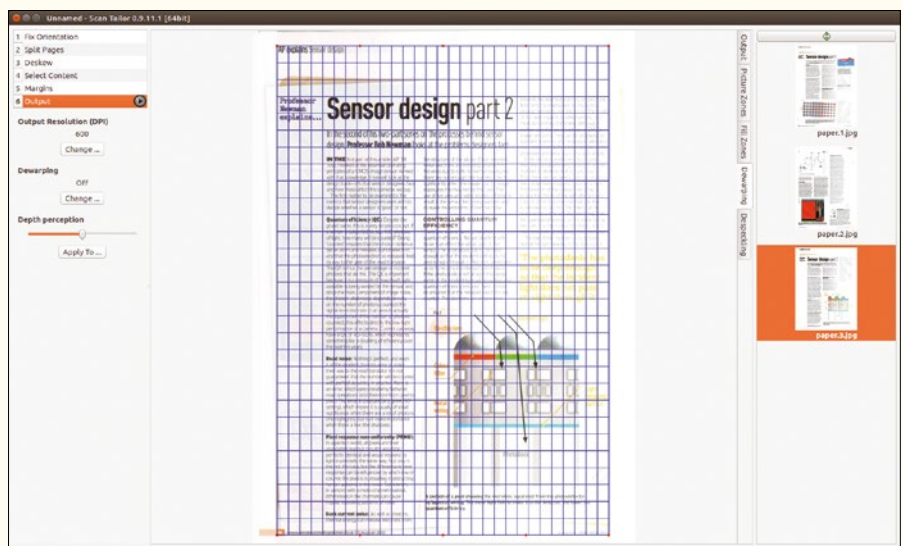



Figure 5: The Dewrapping tool is useful for straightening warped pages.

# Subscribe now!

Don't miss a single issue of the magazine that delivers the in-depth technical solutions you'll use everyday!



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



UK £ 49.90 | Europe € 79.90 | USA / Canada US\$ 99.95 Rest of World (by Airmail) US\$ 109.90  
DVD included with print edition • Terms and conditions: <http://goo.gl/SSSQer>

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



functionality. So, how do you actually turn a collection of single pages into a well-organized and searchable library? Enter Paperwork [2].

### Building an Archive with Paperwork

Paperwork offers powerful yet easy-to-use tools for archiving, organizing, and searching scanned documents (Figure 6). Written in Python, it is easy to install using a few simple commands. On Ubuntu, you can start by installing the required packages:

```
sudo apt install python-pip python-setuptools python-dev python-pil libenchant-dev
```

Next, install Paperwork by running the `sudo pip install paperwork` command. Finally, you need to check for remaining dependencies and install them using the `paperwork-chkdeps` command. Once you've done that, you can launch Paperwork by issuing the `paperwork` command in the terminal.

The first step is to populate the application with scanned pages, and Paperwork provides two ways to do that. If you already have scanned pages, you can import them into Paperwork. This option is ideal for importing pages processed with Scan Tailor. However, Paperwork doesn't support the TIFF format, so you need to convert `.tif` files from Scan Tailor into the JPEG format first.

The easiest way to do this is to use the *mogrify* tool that is part of the ImageMagick suite. In the terminal, switch to the directory containing the `.tif` file and run the `mogrify -format jpg *.tif` command. Then, switch to Paperwork, select the *New Document* item in the left sidebar, and choose *Import file(s)* from the *Scan* menu. Add all the pages you need, and Paperwork will import them into a single document – no need to collate them manually.

You can also scan pages directly into Paperwork using the *Scan* button. When you press the *Scan* button for the first time, Paperwork prompts you to configure some basic settings (Figure 7). To change the default directory for the Paperwork library, specify the desired location in the *Work directory* field. If Paperwork has successfully detected the connected scanner, you should see its name in the *Device* field. Then select the appropriate scan type from the *Source* drop-down list. Interestingly, Paperwork can scan not only regular pages (the Normal source) but also slides (Transparency) and negatives (Negative).

Obviously, you wouldn't want to use Paperwork as your preferred tool for scanning negatives, but it can be useful for maintaining records for physical negatives and transparencies. Select the desired scanning resolution (300 is good, 600 is even better) and make sure that the OCR (Optical Character Recognition) option is enabled.

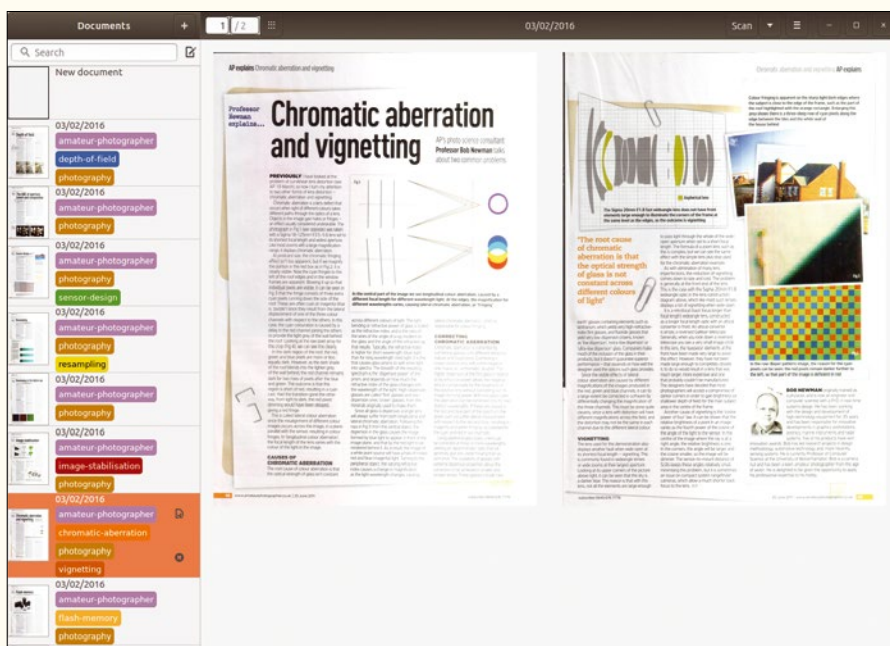


Figure 6: Paperwork is an ideal tool for building a searchable library of scanned pages.

Press the *Scan* button to perform a test scan and make sure that everything works properly. Close the *Settings* dialog, choose the *New document* item in the left sidebar, and press *Scan* to start scanning pages to a new document.

Although Paperwork is not built for processing scanned pages, it has basic cropping and rotating tools. To access them, click on the page you want to edit, and press the *Edit* button.

When you scan pages in Paperwork, it automatically performs optical character recognition, and the application allows you to run full-text search queries on all scanned documents as if they were regular text files. Start typing a search term in the *Search* field, and Paperwork returns a list of pages containing the search term. Better still, the application highlights all occurrences of the search term in each returned page. This functionality alone makes Paperwork an indispensable tool for managing scanned pages, but the application has yet another clever trick up its sleeve.

You can add multiple labels to each document in the library and assign a unique color to each label. This allows you to visually identify specific documents in the library, as well as search documents by their labels. In addition to that, you can assign multiple keywords to each document and run keyword-based search queries.

To add labels and keywords to a document, press the *Properties* button next to it (Figure 8). In the *Properties* sidebar, press **+**, give the label a name, and assign a color to it. Enter the desired keywords in the *Additional keywords* field.

To search documents by labels and keywords, press the *Advanced search* button next to the *Search* field. In the *Search* dialog window, you can configure search criteria that include multiple rules. For example, the advanced search query in Figure 9 finds all documents added to the library between December 31, 2015, and May 25, 2016, and containing the *amateur-photographer* and *sensor-design* labels.

## DIGITIZE DOCUMENTS WITH OPEN NOTE SCANNER

A regular scanner is not the only option you have for scanning paper. If your Android device has a decent camera, you can use it to scan paper documents using a specialized app like Open Note Scanner [3]. This is not the most advanced scanning app for Android, but it does the job, and it's released under the GPLv3 license.

Open Note Scanner captures and processes pages automatically, and it works best with pages that have dark (or preferably black) borders around them. To get the best possible result, you might want to place a loose page on a black surface. When activated, Open Note Scanner detects the page's boundaries, captures the documents, corrects the perspective of the captured image, and converts it to black and white. If you plan to do post-processing in Scan Tailor, you can turn off the image-processing functionality in Open Note Scanner.

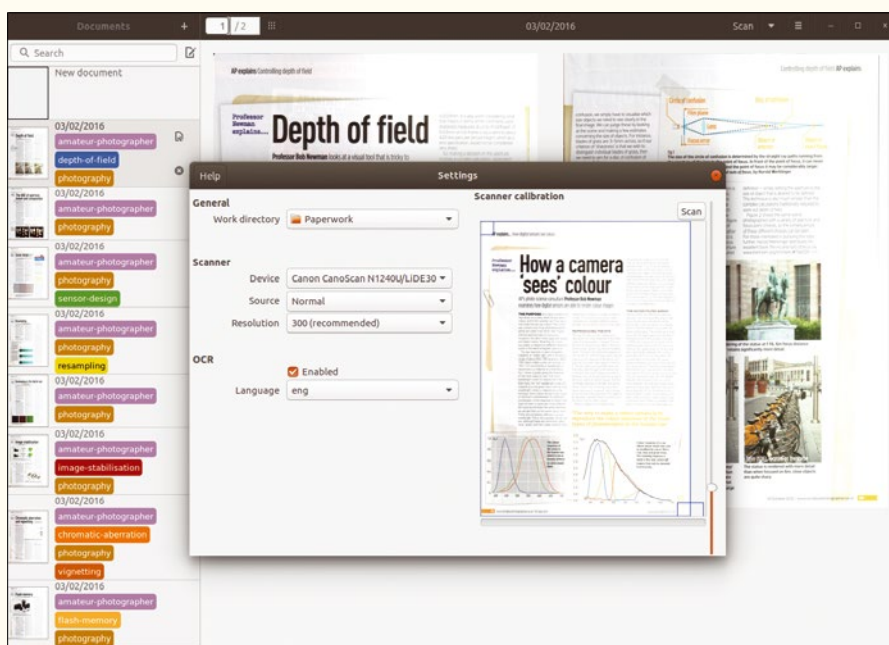


Figure 7: Configuring Paperwork's settings.



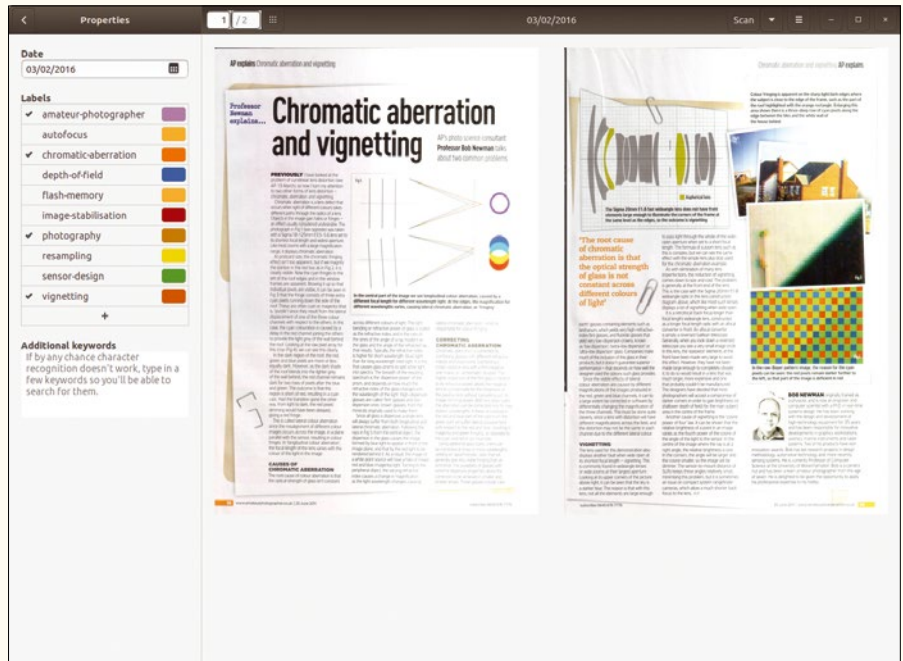


Figure 8: Assigning labels and keywords to a document.

Paperwork stores scanned pages as image files in a dedicated directory, but the application also allows you to export individual pages and entire documents in the library as PDF files (see the “Digitize Documents with Open Note Scanner” box for additional information). This feature gives you an easy way to produce collated documents that can be read on practically any platform. To export a page or a document, choose *Export | Page* or *Export | Document* from the *Options* (a.k.a. Hamburger) menu. Select the desired paper size and quality, specify a path and name for the output file, and select *Export*.

### Final Word

Scan Tailor and Paperwork make a powerful combination for processing and organizing scanned documents. So, if you need to bring order to your paper chaos, these tools will handle the task with aplomb. ■■■

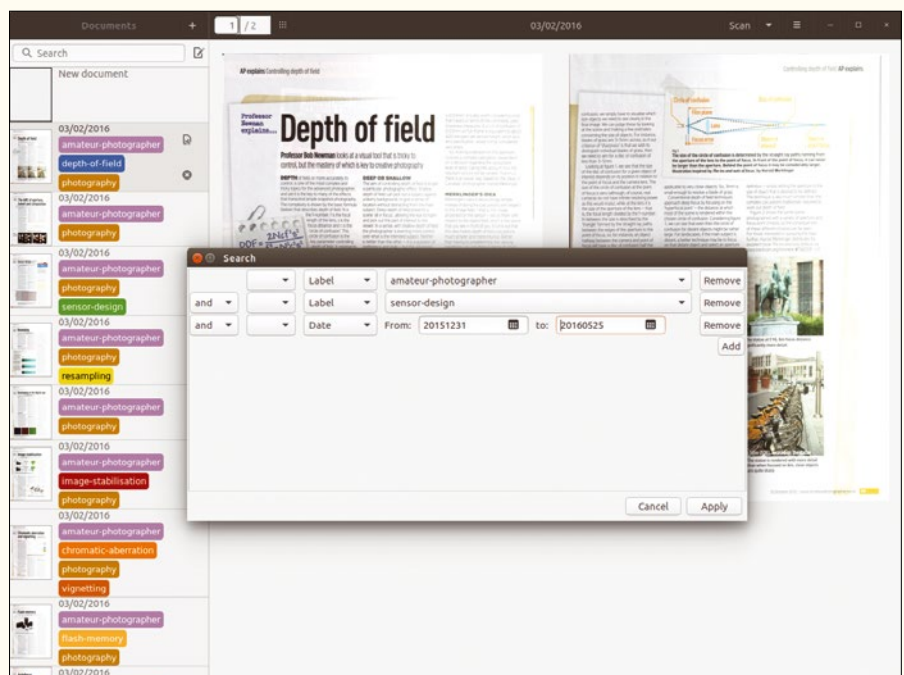


Figure 9: Paperwork lets you create advanced queries.

### INFO

- [1] Scan Tailor: [scantailor.org](http://scantailor.org)
- [2] Paperwork: [github.com/jflesch/paperwork](https://github.com/jflesch/paperwork)
- [3] Open Note Scanner: [github.com/ctodobom/OpenNoteScanner](https://github.com/ctodobom/OpenNoteScanner)

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





## Faster search filter: ICgrep

# Well Filtered

One of the most common tasks when working on computers involves browsing texts for search patterns. Here, ICgrep offers a modern, parallel, and Unicode-enabled alternative to the classic `grep`. *By Karsten Günther*

**G**rep is arguably the most important text-browsing tool in Unix. Although its origins are not entirely clear, likely, the name for this tool developed from the command `g/re/p` (global, regular expression, print) from the Unix standard editor `Ed`.

In any case, `grep` searches entries line by line for certain formulated patterns, as regular expressions. Depending on the options used, the tool can display the matches, their location, number, and so on.

With today's texts, this concept is stretched to its limits in several respects. For example, modern systems no longer use the antiquated ASCII character set with its 128 characters or its larger, 256-character siblings such as Latin 1 or

increasingly Unicode [1] (see the "Using Unicode Characters" box).

Another problem is the definition of input lines. A plethora of variants already existed for identifying or inducing a line break with the previous, non-Unicode-enabled systems [2]. The default in Unix is `NL` (New Line), and there are other options with Unicode, in particular `NEL` (New Line), but also `LS` (Line Separator) or `PS` (Paragraph Separator).

The relatively low operating speed of `grep` – with large input volumes and complex regular expressions – makes matters worse. Even in the best case scenario, regexes [3] are pretty complicated to enter. When editing complex expressions, the computer requires a corresponding amount of computing power. On top of that, there are also conditional, non-greedy, and greedy regexes – which attempt to generate matches for samples to include as much text as possible – and many other advanced variants. All `grep` varieties process regular expressions at run time so that the complexity is mapped directly in the calculation time.

An online study [4] shows just how quickly ICgrep [5] works in comparison to the standard `grep`. It describes a scenario in which the standard `grep` needs more than a minute to search for a pattern, whereas ICgrep handles the same tasks in just over a second. According to the ICgrep developers, there might be differences in speed with a factor of up to 100.

## Installation

Installing ICgrep proved to be a little tricky and quite time consuming in the test; its prerequisites are `Subversion`, `Cmake`, and `Clang` (Listing 1, line 1).

To begin, load the latest ICgrep version onto your computer using `Subversion`. The message `Checked out, revisionNumber` marks the end of this action (line 5). Now switch to the `icgrep1.0/icgrep-devel/` directory. There is an installation guide, among other things, in the `README-icgrep-1.00a.txt` file. It essentially consists of two parts: the interpretation and the installation of `Low Level Virtual Machine (LLVM)` and the subsequent construction of ICgrep itself.

To install `LLVM`, it is important to generate the required `Makefiles`. To do this, switch to the `llvm-build/` directory (line 6) and let `Cmake` perform its role there (line 7). After a short time, you

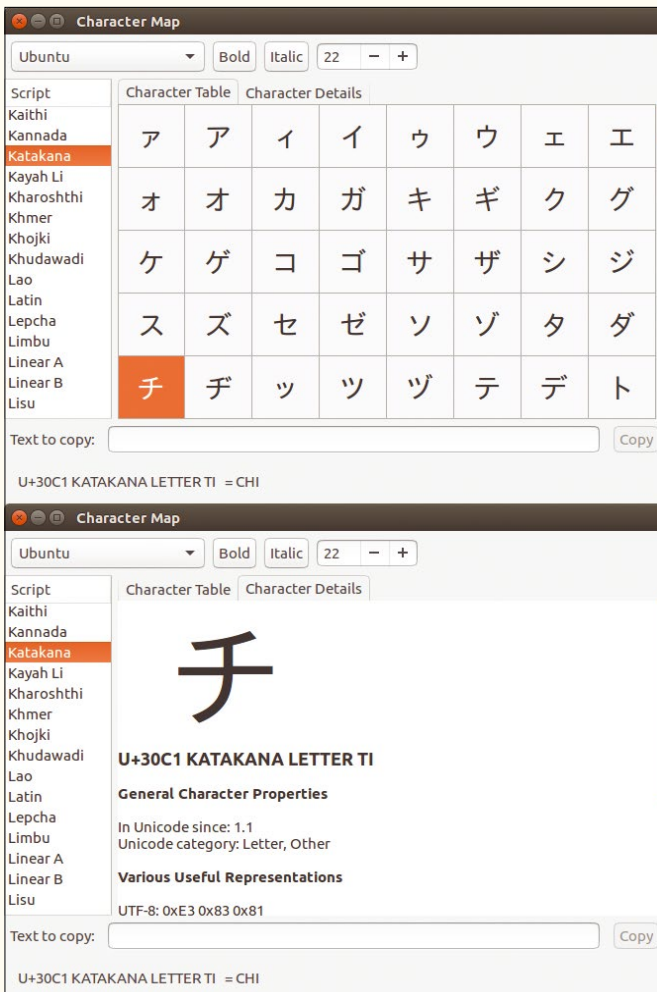
can initiate the build of the LLVM using `make && make install` (line 10). Depending on the computing power, the interpretation process takes about an hour, and the subsequent installation only takes a few seconds to complete.

The second step involves creating ICgrep (Listing 2). To do so, switch to the `icgrep1.0/icgrep-devel/icgrep-build/` directory and create the Makefiles there using the command from the second line. The compiler automatically detects which CPU extensions the program can use to speed up processing. Alternatively, you can present one of the options `SSE2`, `SSE3`, `SSE4_1`, `SSE4_2`, `AVX1`, or `AVX2` using the `-DSIMD_SUPPORT` button.

The next step is to generate the executable file (Listing 2, line 15) without options using `make` – this again takes a little longer. The generated program `icgrep` can be made smaller using `strip`. Last but not least, you can install the program manually in the `/usr/local/bin/` directory.

## In Practice

You can largely control ICgrep – and its cousins `grep` or `agrep` – using options. As this tool is a piece of research software, it contains two or three sets of options with different relevance for typical users. As usual, `-help` shows the most important ones; `-help-list` goes a step further: It generates a short, compact list of important buttons that, however, do require quite a bit of knowledge of regular expressions. The `-help-list-hidden` option ultimately gives you the full list of options; however, many of these are only important in special cases.



**Figure 1:** The Gucharmap tool systematically lists the characters of many languages.

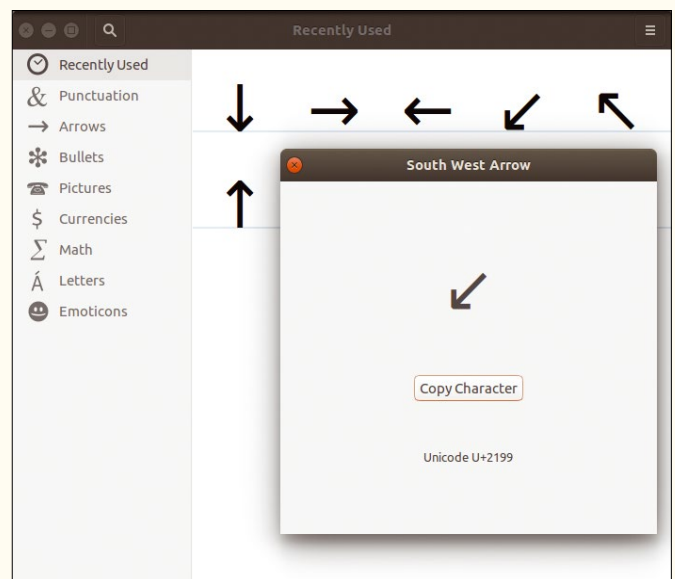
## USING UNICODE CHARACTERS

Unicode aims to assign a digital code to each existing meaningful character or text element of all known writing cultures and character systems. The character range is equally extensive – encoding a single character requires several bytes.

Due to the large scope of the Unicode character range, it often isn't possible to enter characters on the keyboard directly because there simply aren't any keys for them. However, in Linux, Qt and GTK allow you to enter characters using key codes. To do so, you just need to precede the codes with the keyboard shortcut `Ctrl+U` or `Ctrl+Shift+U` – but not all programs support this.

It is often easier just to search for these – especially if you don't know the code for a particular character. There are two utilities for GTK-based desktops that are both used regularly: `Gucharmap` (Figure 1) and `Gnome Characters` (Figure 2). The latter is particularly useful for finding special characters and copying their codes.

You can define patterns, that is, regular expressions, using the `-e` button (“expression”). If there is a risk that the shell will detect and evaluate or replace metacharacters in it, you should include the pattern between the single



**Figure 2:** The Gnome Characters tool sorts special characters into different, often immediately insightful classes.



### LISTING 1: Creating ICgrep

```
01 $ sudo apt-get install cmake clang subversion
02 [...]
03 $ svn co http://parabix.costar.sfu.ca/svn/tags/icgrep1.0
04 [...]
05 Checked out, revision 5012.
06 $ cd icgrep1.0/icgrep-devel/llvm-build/
07 $ cmake -DCMAKE_INSTALL_PREFIX=../libllvm -DLLVM_TARGETS_TO_BUILD=X86
    -DLLVM_BUILD_TOOLS=OFF -DLLVM_BUILD_EXAMPLES=OFF -DCMAKE_BUILD_TYPE=Release
    -DCMAKE_CXX_COMPILER:FILEPATH=/usr/bin/clang++
    -DCMAKE_C_COMPILER:FILEPATH=/usr/bin/clang ../llvm-3.5.0.src
08 [...]
09 -- Build files have been written to: /home/jluther/icgrep1.0/icgrep-devel/llvm-build
10 $ make && sudo make install
11 Scanning dependencies of target LLVMSupport
12 [ 0%] Building CXX object lib/Support/CMakeFiles/LLVMSupport.dir/APFloat.cpp.o
13 [...]
14 [100%] Built target LTO
15 [sudo] password for jluther:
16 -- Installing: [...]
```

or double quotes. An equals sign can (but does not have to) follow the call, which may occur more than once:

```
$ icgrep -e Patt -e=ern
```

These types of expressions then stand for a logical OR: Condition 1 OR Condition 2 OR both must be met. If the number of patterns is too large and the commands get too confusing, it is advisable to outsource the pattern to an external file, one per line. You can specify this file using `icgrep -f <SampleFile>`. Table 1 summarizes the most important buttons.

One of ICgrep's significant shortcomings is that there's no option for recursively editing, like GNU `grep` does with the `-r`

button. Instead, you have to resort to the combination of Find, Xargs, and ICgrep:

```
$ find <Path> -type f | xargs icgrep <Sample>
```

Alternatively, today many shells, including Bash, support recursive directory searching using `**`. However, this functionality must be activated via `shopt -s globstar`.

The `shopt globstar` command then shows `globstar` on. Caution: This process might lead to overly long argument lists and error messages like `-bash: /usr/local/bin/icgrep: The argument list is too long`.

The other important options currently missing in ICgrep include `-v` ("invert match"), `-C`, `-A`, `-B` (displays context lines), `-m` (defines maximum number of matches), `-a`, `-b` (interprets inputs as text or binary), and `--color` (highlights matches in color). The developers are planning to implement some of them soon [6].

### Testing Regexes

With large data streams, it is advisable to test regular expressions first before applying them to all the data. Tools that visualize regular expressions make life easier for less experienced users. Two such tools are available in Ubuntu in the repositories: Kodos and Kiki [7].

Although Kodos can only be installed on older distributions with difficulty, Kiki is less demanding

### LISTING 2: Generate the Executable File

```
01 $ cd ../icgrep-build
02 $ cmake -DCMAKE_BUILD_TYPE=Release -DCMAKE_CXX_COMPILER:FILEPATH=/usr/bin/clang++
    -DCMAKE_C_COMPILER:FILEPATH=/usr/bin/clang ../icgrep-1.00a
03 [... Tests ...]
04 -- Performing Test AVX2
05 -- Performing Test AVX2 - Failed
06 -- Performing Test AVX1
07 -- Performing Test AVX1 - Failed
08 -- Performing Test SSE4_2
09 -- Performing Test SSE4_2 - Failed
10 -- Performing Test SSE4_1
11 -- Performing Test SSE4_1 - Success
12 -- Configuring done
13 -- Generating done
14 -- Build files have been written to: /home/jluther/icgrep1.0/icgrep-devel/icgrep-build
15 $ make
16 Scanning dependencies of target RegExpADT
17 [ 2%] Building CXX object CMakeFiles/RegExpADT.dir/re/re_re.cpp.o
18 [... Compilation ...]
19 [100%] Building CXX object CMakeFiles/icgrep.dir/compiler.cpp.o
20 Linking CXX executable icgrep
21 [100%] Built target icgrep
```

### LISTING 3: Test File

```
$ java -Xmx1200M -Dfile.encoding=windows-1251 -jar bin/regexper.jar -f <TestFile>
```

**TABLE 1:** Important Options in ICgrep

Option	Function
-e=<Sample>	Defines a sample, can be specified multiple times
-f=<File>	Sample file
-c	Only counts matches, doesn't show them
-n	Displays the line numbers of matches
-H	Displays file name
-i	Ignores upper/lowercase
-disable-Unicode-linebreak	Only NL as line break
-normalize-line-breaks	Normalizes line breaks
-dump-generated-IR	Shows the regular expressions used
-print-RES	Shows named expressions

(Figure 3). It uses regular expressions defined and generated by Python, which it then applies to a previously entered text.

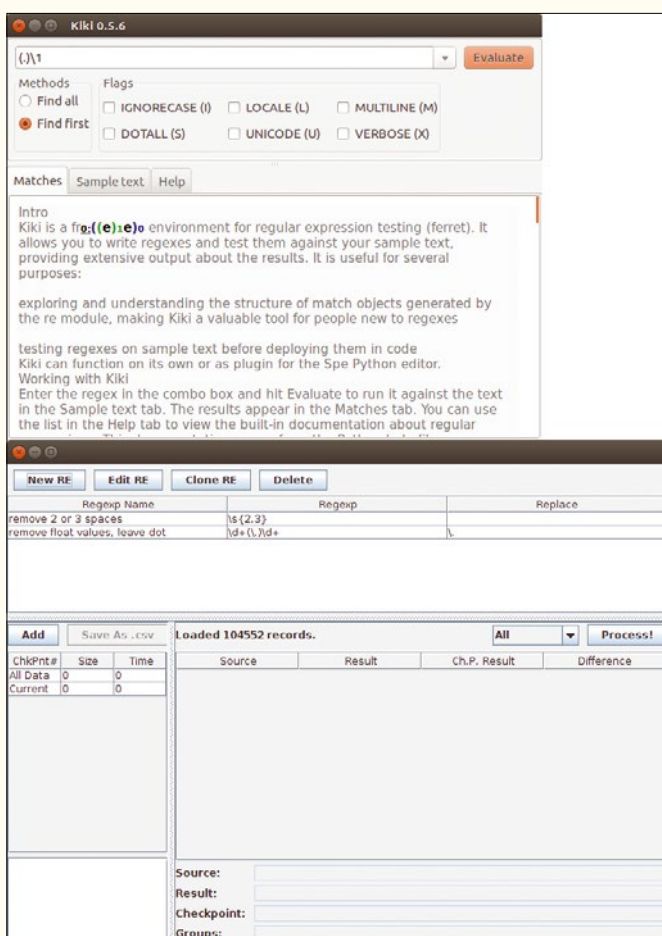
The regex testing tool `regexper` [8] works somewhat differently. The Java program naturally uses Java-based regular expressions and applies them to a previously loaded test file (Listing 3). The test data may be relatively large, which sometimes proves advantageous.

All of these tools, however, have weaknesses when it comes to processing Unicode characters. Therefore, when constructing regular expressions, they are only suitable for initial tests.

## Conclusions

Unless you have a need for Unicode-specific features, the conventional `grep` variants are enough for browsing small data streams or files – not least because these variants

includes special offshoots for different applications, such as `agrep` [9]. However, browsing through large data streams for complex patterns is an entirely different matter: Here, `ICgrep` is a remarkable new development, and it can be used free of charge. ■■■



**Figure 3:** Kiki (above) and `regexper` (below) test regular expressions before using them on large amounts of data.

## INFO

- [1] Unicode: <https://en.wikipedia.org/wiki/Unicode>
- [2] Unicode newline: <https://en.wikipedia.org/wiki/Newline>
- [3] Regular expressions: [https://en.wikipedia.org/wiki/Regular\\_expression](https://en.wikipedia.org/wiki/Regular_expression)
- [4] ICgrep Demonstration Rev 2: [http://international-characters.com/download/icGREP\\_Demonstration.pdf](http://international-characters.com/download/icGREP_Demonstration.pdf)
- [5] ICgrep: <http://www.icgrep.com/>
- [6] ICgrep roadmap: <http://parabix.costar.sfu.ca/wiki/ParabixRegexRoadMap>
- [7] Kiki: <https://code.google.com/p/kiki-re/>
- [8] Regexp testing tool: <http://sourceforge.net/projects/regexper?source=recommended>
- [9] Agrep: <https://www.tgries.de/agrep/>



# 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](http://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](http://www.raspberry-pi-geek.com/mc/subscribe)


Visit us online  
[www.raspberry-pi-geek.com](http://www.raspberry-pi-geek.com)



Like us on Facebook  
[www.facebook.com/RasPiGeek](http://www.facebook.com/RasPiGeek)



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







What's new in Krita 3.0

# Drawing on Success

Krita, a free sketching and painting program, has new features and improvements in version 3.0. *By Bruce Byfield*

**K**rita [1] is well known for having evolved from a toy to a tool for professional artists [2]. In late May, the Krita Foundation announced Krita 3.0, the Animation Release, after the conclusion of a successful crowdfunding campaign [3] to support ambitious last-minute features. The campaign raised more than £42,000 (~\$47,500), with a third of its £30,000 goal pledged in less than two days. I talked with Krita maintainer Boudewijn Rempt about the introduction of animation, improved text tools, and revision of vector graphics – three of the major features that debuted in version 3.0 [4].

“Feature requests are a tough thing to get right,” Rempt observes. Sometimes, new users ask for something that is obscure or fits poorly with the other features. Other times, a developer might want a feature for themselves. However, usually, that’s not how the project develops its roadmaps.

“Usually it’s pretty clear from discussions that artists have on the chat channel or the forum, or even in our bug tracker, where the real problems are. And, of course, we have many artists in the development team. They join us at sprints, and they are always around on the #krita IRC channel. So, gradually, a sort of consensus about the priority of different feature requests arises, and we start asking people to do mock-ups and user-experience designs.”

Another way features become priorities is what Rempt calls the “organic” method. Over time, the general dissatisfaction or demand for a feature increases, until a consensus is reached that it needs to be improved or rewritten.

“Right now,” Rempt says, “we’re getting people who ask whether it isn’t time to start improving the brush engines again, [so] next year, it might be time to spend some really focused time on them again.”

## Introducing Animation

One of the major goals that developed organically is animation. According to Rempt, four different people tried five different times to implement animation, “which shows it’s something people really want.” In the original planning, basic animation features were scheduled for v3.1, but they were ready for v3.0 after extensive work was done in 2015.



Artwork by Wolthera van Hóvell tot Westerfliet (<http://wolthera.info>) CC BY-SA 4.0 [11]

Animation features for v3.0 include importing and exporting an image sequence, dragging and dropping frames, and onion skinning [5], or working on multiple frames at the same time. Still to come in v3.1 are integration with sound and export to movie formats.

The v3.0 release, Rempt says, is “usable for short clips of hand-drawn animation. But we’re not yet in a stage where professional animators actually know about Krita. The professional world is dominated by closed-source software like Toon Boom [6] or TVPaint [7], but with the recent opening up of OpenToonz [8], I hope that free and open software will become more accepted.”

## Text Tools

The crowdfunding page gives a long list of improvements to Krita’s text tools, ranging from improving the basic interface and working with text on a path to using ligatures and composition overlays – essentially, frames within frames that are formatted separately. Many of the improvements will help work with multiple languages, including changing the direction for text flow and making translation easier. At the same time, as Krita becomes more independent of Calligra Suite, “we’re going to drop the weird office stuff like the bibliography tool, the paragraph sections, and the semantic markup.”

The page also states that changes to Krita’s text tools will focus on “creating poster layouts, comics, or game cards.” Asked whether such a goal means that Krita is moving toward desktop publishing, Rempt replies, “Well, ... not really desktop publishing.

We’re now firmly focusing on the need to add occasional, but beautiful text to an image for comics or webcomics, or playing cards or designs – not tables, or headers or footers, or text that’s connected through different shapes.”

Typographical layout as done in Scribus [9] is apparently not under consideration for now.

## Rendering Objects Using SVG

Krita has supported saving in scalable vector graphics (SVG) [10] for years. However, the new release also sees Krita rendering objects using SVG instead of Open Document Graphics (ODG), the format that is also used by LibreOffice.

The use of ODG originated because, when Calligra Suite was KOffice, Krita was more closely associated with the rest of Calligra Suite. Given that Calligra Suite was pioneering the use of ODG, at the time it made sense to have graphics displayed as ODG rather than SVG.

When KOffice was ported to the Qt4 toolkit, each application was organized around basic objects, such as a page or a spreadsheet, and documents were built from the addition of such objects.

“It worked, but never really well,” Rempt recalls. “When we ported KOffice, renamed to Calligra Suite by then, to Qt5, we decided that it was time to cut the cords and make sure that the office applications would be focused on office needs and the art application, Krita, on art needs.” Because ODF has “some odd limitations,” Krita would transition to using vector graphics for its basic objects. ODF, Rempt says, “just isn’t the right format” for a graphics application, whereas using SVG would improve interoperability with Inkscape and web browsers, “and even with LibreOffice, which




“Merida” by Ramon Miranda (<http://www.ramonmiranda.com>)  
CC BY-SA 4.0 [11]



Artwork by David Revoy (<http://davidrevoy.com/>) CC BY-SA 4.0 [11]





Artwork by Tyson Tan CC BY-SA 4.0 [11]

has also been emphasizing the importance of SVG recently.”

Some work on this transition was done in 2012 but never finished. Now, “we’ll have to do a lot of design first,” but Krita developers hope that Krita will be stronger for not having to support word processors or presentation applications.

### Not in the Release: Improved Brush Tools

Although Rempt is excited about the improvements in v3.0, he also speaks with some regret about features that will not be ready for the release. “I really had wanted to finish the stacked brush engine, where you could make a stack of brush presets and set a blending mode, random offset, and masking, and then paint with all of them.” Rempt managed to design an interface but not to implement the engine.

Another issue that has been delayed is improving the brush engine. “It’s too big and can be both confusing and intimidating,” Rempt says, and resists being redesigned. “Then there’s the need to figure out a way to access favorite brush presets with shortcuts, which is coupled with the problem that newcomers to Krita have issues with the various ways of actuating an eraser mode. All those things need careful thought and then a final decision on the approach. We haven’t found any consensus here yet.”

### Roadmaps and Reasons

All in all, Krita 3.0 is a combination of careful planning and opportunity.

“Making 2016 the Year of Vector and Text has been on the roadmap for two, three years,” Rempt says. “Animation came out of the blue, though.”

Like most maintainers, Rempt wishes for more developers. However, having managed several successful crowdfunding campaigns, he hopes that in the future, more developers can be hired full-time to write major features. “We’re right now in a kind of awkward situation, sort of suffering from growing pains – the user base is getting big enough that it becomes quite hard to keep up with the feedback, whether in the form of bugs or wishes.”

Still, although developing Krita can be challenging, Rempt sounds like he has no doubt about the value of it:

“In this day and age, you have to do digital art, or you’re not in the same place as the other artists. But the big, important, recognized tools are closed source and out of reach for many people. They’re also [proprietary], so artists cannot experiment with changing their tools.

“And that’s, in the end, a big reason why I’m working on Krita. I almost haven’t got time to paint or sculpt anymore, but that doesn’t matter that much: What matters is that we’re creating something that gives people with the urge to create art a tool that’s available to all of them, that allows them to compete, and that they can change and experiment with, and that can never, ever be taken away from them.” ■■■

### INFO

- [1] Krita: <https://krita.org/>
- [2] Krita: Listening to the Pros: <http://www.linux-magazine.com/Online/Features/Krita-KDE-s-Powerful-Graphics-Editor-Takes-on-Photoshop-and-GIMP>
- [3] Kickstarter campaign: <https://www.kickstarter.com/projects/krita/krita-2016-lets-make-text-and-vector-art-awesome>
- [4] Release notes: <https://krita.org/krita-3-0-release-notes/>
- [5] Onion skinning: [https://en.wikipedia.org/wiki/Onion\\_skinning](https://en.wikipedia.org/wiki/Onion_skinning)
- [6] Toon Boom: <http://www.toonboom.com/>
- [7] TVPaint: <http://www.tvpaint.com/v2/content/article/home/>
- [8] OpenToonz: <https://opentoonz.github.io/index.html>
- [9] Scribus: <https://www.scribus.net>
- [10] SVG: [https://en.wikipedia.org/wiki/Scalable\\_Vector\\_Graphics](https://en.wikipedia.org/wiki/Scalable_Vector_Graphics)
- [11] Attribution-ShareAlike 4.0 International: <https://creativecommons.org/licenses/by-sa/4.0/>





### A tale of long-lived Free and Open Source projects

# Longevity

By Jon “maddog” Hall

I was reading an Atlanta Linux Enthusiasts (ALE) mailing list recently that has been around a long time. Many of the people on that list have been reading and writing to the list for 20 years. One of them had just installed a version of the Common Desktop Environment (CDE), which was developed when I was working at Digital Equipment Corporation (DEC) more than 20 years ago.

The person writing about it talked about “Motif,” a set of widgets that create the button boxes, scroll bars, and menus for the X Window System, and some of the various applications that came with CDE, such as the Motif Window Manager (MWM), xsnow (which made tiny snowflakes fall down the screen and build up on top of the windows and menubars), xroach (which had cockroaches scurry around and behind your windows), xeyes (eyes that would follow your cursor), xearth (a spinning “almost 3D” picture of the earth), and xfish tank (obviously, a fish tank for your system). It also came with some simple text editors, a “desktop,” and various text-based startup files to tailor the environment.

During the exchange of messages, some commenters indicated that everyone thought CDE was dead in 2002; however, apparently the people that owned the code decided to open source it. Since that time, CDE has actually flourished despite all of the competing desktops out there.

This made me think of several other stories of Open Source software longevity.

When I was introduced to Unix in Bell Laboratories, two text processing formatters, `nroff` and `troff`, used control statements embedded in the text to tell the formatting program when and how much to indent, when to break a line, whether hyphenation was to occur, and so on. It was NOT What-You-See-Is-What-You-Get (WYSIWYG) by any stretch of the imagination. `Nroff` was for character cell and paper terminals, and `troff` was for phototypesetters – specifically, the very expensive Graphic Systems CAT phototypesetter located at Bell Labs, New Jersey, in 1971.

At first attempts were made to create a filter that would take the computer-assisted typesetter (CAT) commands and convert them into commands for other typesetters and (later) laser printers, but this effort was not com-

pletely successful. Later, a different version of `troff` came out called “device-independent `troff`” (`ditroff`), which produced device-independent code that was easy for post-processors to use.

`Troff` (and to a certain extent `nroff`) also could use pre-processors to expand macros in the stream of the input text. These macros could help set up pictures (`pic`), mathematical expressions (`eqn`), tables (`tbl`), and manual pages, affectionately known as “`man`” pages. The `man` command would call `nroff` (or `troff`) to format manual pages to the needed size for devices on the fly. `Ditroff` could invoke these macro packages as options on its command line.

I had used `nroff` and `troff` quite a bit while I was at Bell Labs, and when I went to DEC in 1983, I took my documents and macros that I had written with me. The `*roff` packages were now 12 years old.

I continued to write documents in `*roff`, because WYSIWIG packages of the day were not very good for documents of more than one or two pages and were decidedly NOT useful for documents that needed indices and tables of content.

Eventually, the rights to a lot of the `ditroff` software were purchased by a company called SoftQuad. They put out a product that sold for several thousand dollars (in 1994), and DEC bought a copy for the software support and became a reseller.


About the same time, the GNU project started working on a replacement called `groff`, which produced very good PDF output, bypassing a lot of the device support issues of `ditroff` and its filters. I remember putting up a printout on my cubicle showing the output of the SoftQuad product (which really did not look that good and could not even do parallel lines correctly) compared with what the PDF file from `groff` produced, which was “perfect.” I pointed out that SoftQuad’s product cost several thousand dollars and `groff` was “free.” This was not appreciated by Digital’s SoftQuad product manager.

Another, shorter, example of software longevity is the RAND Mail Handling System, known as MH. A command-line mail system, it was made graphical by wrapping a Tk-based GUI around it and calling it `xmh`, and then adding MIME extensions and calling it `exmh`. When Motif was added to it, the name changed again. All this time I just kept using it.

Therefore, it does not surprise me that CDE lives on as Free and Open Source Software. Once code is “Opened,” projects can last a very, very long time. ■■■

## 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®.



# An entire year of **RASPBERRY PI PROJECTS AND OPERATING TIPS**

## for one low price!

### Two bundles to choose from:



# 2014



# 2015

## Hurry while supply lasts!

**ORDER NOW!**



<http://shop.linuxnewmedia.com/us/magazines/raspberry-pi-geek/catchup.html>



# FEATURED EVENTS



Users, developers, and vendors meet at Linux events around the world. We at *Linux Magazine* are proud to sponsor the Featured Events shown here. For other events near you, check our extensive events calendar online at <http://linux-magazine.com/events>.

If you know of another Linux event you would like us to add to our calendar, please send a message with all the details to [events@linux-magazine.com](mailto:events@linux-magazine.com).

## 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 North America

**Date:** August 22–24, 2016

**Location:** Toronto, ON, Canada

**Website:** <http://events.linuxfoundation.org/events/linuxcon-north-america>

Join developers, sys admins, architects, and technical talent for education, collaboration, and problem-solving to further the Linux platform. Speakers include Linus Torvalds and Cory Doctorow.

## 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.

## EVENTS

Tech Open Air 2016	July 13–15	Berlin, Germany	<a href="http://toa.berlin/">http://toa.berlin/</a>
GUADEC 2016	August 12–14	Karlsruhe, Germany	<a href="https://2016.guadec.org/">https://2016.guadec.org/</a>
LinuxCon North America	August 22–24	Toronto, ON, Canada	<a href="http://events.linuxfoundation.org/events/linuxcon-north-america">http://events.linuxfoundation.org/events/linuxcon-north-america</a>
ContainerCon North America '16	August 22–24	Toronto, ON, Canada	<a href="http://events.linuxfoundation.org/events/containercon">http://events.linuxfoundation.org/events/containercon</a>
IFA	September 2–7	Berlin, Germany	<a href="http://www.ifa-berlin.de/">http://www.ifa-berlin.de/</a>
systemd.conf2016	Sep. 28 – Oct. 1	Berlin, Germany	<a href="https://conf.systemd.io/">https://conf.systemd.io/</a>
LinuxCon Europe	October 4–6	Berlin, Germany	<a href="http://events.linuxfoundation.org/events/linuxcon-europe">http://events.linuxfoundation.org/events/linuxcon-europe</a>
ContainerCon Europe	October 4–6	Berlin, Germany	<a href="http://events.linuxfoundation.org/events/containercon-europe">http://events.linuxfoundation.org/events/containercon-europe</a>
Open Source Monitoring Conference	November 29– December 2	Nürnberg, Germany	<a href="https://www.netways.de/events_schulungen/osmc/uebersicht/">https://www.netways.de/events_schulungen/osmc/uebersicht/</a>
IT-Tage 2016	December 12–15	Frankfurt, Germany	<a href="http://www.ittage.informatik-aktuell.de/">http://www.ittage.informatik-aktuell.de/</a>

# CALL FOR PAPERS

We are always looking for good articles on Linux and the tools of the Linux environment. Although we will consider any topic, the following themes are of special interest:

- System administration
- Useful tips and tools
- Security, both news and techniques
- Product reviews, especially from real-world experience
- Community news and projects

If you have an idea, send a proposal with an outline, an estimate of the length, a description of your background, and contact information to [edit@linux-magazine.com](mailto:edit@linux-magazine.com).



The technical level of the article should be consistent with what you normally read in *Linux Magazine*. Remember that *Linux Magazine* is read in many countries, and your article may be translated into one of our sister publications. Therefore, it is best to avoid using slang and idioms that might not be understood by all readers.

Be careful when referring to dates or events in the future. Many weeks could pass between your manuscript submission and the final copy reaching the reader's hands. When submitting proposals or manuscripts, please use a subject line in your email message that helps us identify your message as an article proposal. Screenshots and other supporting materials are always welcome.

Additional information is available at:

[http://www.linux-magazine.com/contact/write\\_for\\_us](http://www.linux-magazine.com/contact/write_for_us).

## AUTHORS

Erik Bärwaldt	16, 66
Swapnil Bhartiya	8
Bruce Byfield	70, 90
Joe Casad	3
Mela Eckenfels	32
Karsten Günther	84
Jon "maddog" Hall	94
Marcel Hilzinger	48
Valentin Höbel	22
Oliver Hoffmann	52
Heike Jurzik	32
Klaus Knopper	56
Charly Kühnast	58
Dmitri Popov	76
Mike Schilli	60
Ferdinand Thommes	12, 36, 42
Uwe Vollbracht	28

## CONTACT INFO

### Editor in Chief

Joe Casad, [jcasad@linux-magazine.com](mailto:jcasad@linux-magazine.com)

### Managing Editor

Rita L Sooby, [rsooby@linux-magazine.com](mailto:rsooby@linux-magazine.com)

### Localization & Translation

Ian Travis

### News Editor

Swapnil Bhartiya

### Copy Editors

Amber Ankerholz, Amy Pettle

### Layout

Dena Friesen, Lori White

### Cover Design

Dena Friesen, Lori White

### Cover Images

© maksym yemelyanov, 123RF.com  
© Christos Georghiou, fotolia.com

### Advertising – North America

Ann Jesse, [ajesse@linuxnewmedia.com](mailto:ajesse@linuxnewmedia.com)  
phone +1 785 841 8834

### Advertising – Europe

Penny Wilby, [pwilby@sparkhausmedia.com](mailto:pwilby@sparkhausmedia.com)  
phone +44 1807 211100

### Publisher

Brian Osborn, [bosborn@linuxnewmedia.com](mailto:bosborn@linuxnewmedia.com)

### Marketing Communications

Gwen Clark, [clark@linuxnewmedia.com](mailto: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](mailto:cs@linuxpromagazine.com)  
Phone: 1-866-247-2802  
(Toll Free from the US and Canada)  
Fax: 1-785-856-3084

For all other countries:

Email: [subs@linux-magazine.com](mailto:subs@linux-magazine.com)

Phone: +49 89 99 34 1167

Fax: +49 89 99 34 1198

[www.linuxpromagazine.com](http://www.linuxpromagazine.com) – North America

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

While every care has been taken in the content of the magazine, the publishers cannot be held responsible for the accuracy of the information contained within it or any consequences arising from the use of it. The use of the disc provided with the magazine or any material provided on it is at your own risk.

Copyright and Trademarks © 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 190 / September 2016

# Chroot

**Approximate**  
UK / Europe Aug 06  
USA / Canada Sep 02  
Australia Oct 03  
**On Sale Date**

After a break-in, the wily intruder tries to expand the invasion, unless you take steps to limit the damage. A chroot jail keeps the rest of your system out of reach.



## Preview Newsletter

The Linux Magazine Preview is a monthly email newsletter that gives you a sneak peek at the next issue, including links to articles posted online.

Sign up at: [www.linux-magazine.com/newsletter](http://www.linux-magazine.com/newsletter)

Lead Image © victor-zastol'skiy / 23RF.com



# 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.

**KEYNOTE SPEAKERS:**

**Jane Adams**, Data Research and Acceleration, Two Sigma

**Mitchell Hashimoto**, Founder, HashiCorp

**John Rouse**, CTO, EMC

**Register by November 10 and save!**

Sponsored by the USENIX Association

**LISA**16



[usenix.org/lisa16](http://usenix.org/lisa16)





# 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](http://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.