

FREE  
DVD

**Knoppix 7.6 & netrunner OS**  
The World's Greatest Live Distro

**SECURING LINUX**  
Expert tips for a safer system

**LINUX**



# LINUX **PRO**

## MAGAZINE

FEBRUARY 2016

# SECURING LINUX

**Expert tips for locking down your Linux server system**

### Host-based IDS

Investigate past break-ins with a back-to-basics detection tool

### What's New

With the Ardour 4.4 digital audio workstation

### LibreOffice Tricks

Integrate emoticons with your spreadsheet



### GPS Tracking

**Build your own tracking app in Python**

### Admidio


Manage member accounts for your website

### Mhdds

Group hard disks with the multi-drive filesystem

### The New SQL

Much more than just a relational database



Issue 183 US\$ 12.99  
February 2016 CAN\$ 13.99



# Time to remodel!

## The latest vServer generation

- *Top performance*
- *Triple-redundant storage*
- *Super speedy SSDs*
- *guaranteed RAM*
- *Snapshot function*
- *Free Setup*



### e.g. vServer CX10

- Complete root access
- 1 vCore
- 1 GB RAM
- 25 GB SSD
- 2 TB Traffic\*
- 1 Snapshot

monthly \$ **4.50**

Enjoy the optimal combination of performance-enhancing guaranteed RAM, generous CPU resources, and speedy SSDs, with the maximum amount of security for your data thanks to triple-redundant storage!

\* There are no charges for overage. We will permanently restrict the connection speed if the traffic limit is exceeded. Optionally, the limit can be permanently cancelled for \$ 1.50 per additional TB used.

All prices exclude Sales Tax and are subject to the terms and conditions of Hetzner Online GmbH. Prices are subject to change. All rights reserved by the respective manufacturers.

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

# A TALE OF TWO TOOLS

Dear Linux Pro Reader,

One interesting item in the chat rooms these days is Mozilla's announcement that it plans to end its involvement with the famous Mozilla Thunderbird mail client [1]. Firefox and Thunderbird used to be twins – two hot, new built-for-the-future desktop tools rising from the ashes of the old Mozilla browser community, which itself had risen from the ashes of a pioneering Internet company known as Netscape. The Firefox web browser and Thunderbird mail client saw their first major releases within a month of each other back in 2004, and both embodied the vision of a sleek, sensible, single-purpose tool to replace the dinosaurs of an earlier era.

Firefox rode that vision to glory, challenging Microsoft's Internet Explorer for browser dominance in what was sweet revenge for the Netscape faithful. The Firefox browser lost a bit of its fire with the rise of the Google's Chrome browser, but it remains a popular option for Linux, Mac OS, and Windows users around the world. Thunderbird, however, flew to a different mountain. The Thunderbird mail client is still a top choice in many open source circles, but the rest of the world seems to have moved on.

The rise of mobile devices and cloud computing has left conventional desktop mail clients like Thunderbird to drift. According to the Email Client Marketshare website [2], 85% of the marketshare for mail clients belongs to mega-companies Apple, Google, and Microsoft. Another leading option, Yahoo! Mail, is a cloud-based webmail option. Thunderbird, weighing in at 10th place in market share, is actually the highest ranking third-party, conventional desktop mail reader on the list, but its market share is down to around 1%.

The appeal of reading web-based mail on any desktop or any mobile device leads to overwhelming popularity for the Apple, Google, and Microsoft products, and even in the business environment, more sophisticated groupware tools now integrate mail with other forms of collaboration.

So desktop email clients aren't as important as they used to be. Does that really mean Thunderbird (the 10th most popular email client) has no place in the portfolio? According to Mozilla Foundation Executive Chairwoman Mitchell Baker, "Many inside Mozilla, including an overwhelming majority of our leadership, feel the need to be laser-focused on activities like Firefox that can have an industry-wide impact. With all due respect to Thunderbird and the Thunderbird community, we

have been clear for years that we do not view Thunderbird as having this sort of potential."

Baker, and others in the Mozilla community, refer to the cost of Thunderbird as a "tax" upon Mozilla resources. It seems fair to ask, what exactly is getting taxed in order to pay that tax to Thunderbird? Look a little closer, and you'll see that Firefox is big, big money. Mozilla used to get millions of dollars per year from Google for making Google the default search engine, and they signed a big new contract with Yahoo! for US search a year ago. So the twins have gotten a bit asymmetric, with Firefox operating as a multi-million dollar business and Thunderbird as a steady, but not-so-lucrative basic open source development project.

What corporation wouldn't want to get rid of the non-remunerative stuff to focus on the ultra-remunerative stuff? But then, Mozilla Foundation isn't supposed to be a corporation, or at least, part of it isn't. Actually, the for-profit Mozilla Corporation is a wholly owned subsidiary of the non-profit Mozilla Foundation, which means the organization has a kind of double identity.

I'm willing to admit that desktop email clients are probably not the future of the Internet, but I wonder about the argument that an open source organization as big and well heeled as Mozilla Foundation should abandon all efforts that aren't slated to have an "industry-wide impact." The FOSS world supports thousands of open source projects that have no potential to overturn the existing order but exist simply because they play a vital role within the open source ecosystem.

The good news is that Mozilla seems committed to finding a good home for the Thunderbird project. Baker writes that she and other Mozilla execs "... want to make sure that Thunderbird has the right kind of legal and financial home, one that will help the community thrive." A smooth transition would be best for all parties. One possible candidate is the Apache Foundation, which has taken in other lost projects in the past. An Apache scenario might work well for everyone, but if it happens, I probably won't be the only one to wonder, how come Apache can afford to support Thunderbird but Mozilla can't? ■■■

Joe

Joe Casad,  
Editor in Chief



## INFO

[1] Mitchell Baker on Thunderbird Future: <https://groups.google.com/forum/#!topic/mozilla.governance/kAyVlhEcXg>

[2] Email Client Market Share: <https://emailclientmarketshare.com/>



# LINUX PRO MAGAZINE

## news

### 8 THIS MONTH'S NEWS

- Adobe Flash renamed.
- Rasp Pi generates weak SSH keys.

### 9 NEWS

- VMware Photon Controller goes open source.
- New point-and-click exploit kit appears in the wild.
- Online News.

### 10 RANSOMWARE!

- Red Hat releases RHEL 7.2.
- Dell systems ship with faulty root certificates.
- Ransomware reaches Linux.

# Securing Linux

Secure your system from the first boot, and know how to find the evidence if you're ever attacked.

**12 Host-Based IDS** A host-based intrusion detection system is a simple but powerful tool for finding traces of an attacker's footprint.

**16 Securing Linux** When you set up a new Linux server, take a few extra steps to ensure the system is truly secure.

## 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.

## Community Notebook

- 78 Doghouse: Raspberry Pi Zero Just chill. You'll get one.
- 80 Event Report: SUSECon 2015 Highlights from Amsterdam.
- 82 Open Arts Open source hardware and software in the arts.

- 88 openSUSE Tumbleweed Richard Brown on the rolling release process.
- 92 Kernel News Reports on kernel communities, conundrums, and controversies.



## HIGHLIGHTS

**12 HIDS**  
Detect evidence of break-ins with a host-based intrusion detection system.


**52 ADMIDIO**  
Web-based membership management for associations.

**74 MHDDFS**  
A union filesystem that accommodates the growth of data across media.

**82 OPEN ARTS**  
Open source hardware and software encourages creativity in the arts.

## FEATURES

**26 Modern SQL**  
Discover the state of the art in modern SQL databases.



**40 Ardour**  
Version 4.4 of this free digital audio workstation impresses with optimum stability and well-conceived, functional extensions.

**32 DocBook XML Publishing**  
Three documentation tools: DAPS, Publican, and Yelp.

**48 Perl: Producing Labels**  
Bring order to chaos with Dymo LabelWriter and Perl.

**37 Charly: Gatling**  
A lightweight web server that can accommodate old PCs and nanocomputers.

**52 Admidio**  
Manage membership data, events, appointments, and participants online.

**38 Ask Klaus!**  
Strategies for getting around flash drive limitations and updating OSs on flash drives.

## REVIEW

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

## LINUXUSER

**56 Calc Conditional Formatting**  
Integrate graphical information alongside the data it represents.

**60 Command Line – Users and Groups**  
Create and maintain multiple accounts and manage your system more effectively.

**64 Workspace – Geofix**  
Record and track locations and build a simple Python-based app in the process.

ID	Date/Time	Lat/Long	Longitude	Altitude	Snapshot	Map
132	20151117-144004	56.34286762	10.18639207	geo:56.34286762,10.18639207		
137	20150816-142348	39.57032293	2.48144079	geo:39.57032293,2.48144079		
137	20150807-142321	51.51509365	-8.21916498	geo:51.51509365,-8.21916498		
78	20150409-212927	56.12547853	10.21242532	geo:56.12547853,10.21242532		
65	20150531-151652	35.42674676	139.77223074	geo:35.42674676,139.77223074		
65	20150531-111347	35.38899468	139.75304374	geo:35.38899468,139.75304374		

**70 Video Conferencing**  
Jitsi video and audio conferences in web browsers.

**74 mhddfs**  
Group directories or hard disks on a union filesystem to create a single, large, virtual filesystem.



## netrunner OS & Knoppix 7.6

The World's Greatest Live Distro

- Easy desktop experience
- Plasma 5.4
- New look and feel
- Elegant design

- \* 32- and 64-bit versions
- \* Runs directly from DVD
- \* Includes dozens of powerful management and troubleshooting tools

SEE PAGE 6 FOR DETAILS

# On the DVD





## Knoppix 7.6

The popular Knoppix is the pioneering Live Linux system loved by Linux power users and admins around the world. Knoppix comes with a powerful toolkit of expert utilities for system maintenance and troubleshooting, as well as an impressive collection of everyday desktop applications. Version 7.6 comes with Linux kernel 4.2.2, the Wine 1.7.50 Windows integration tool, and experimental support for UEFI boot and Secure Boot.

## Netrunner 17

Netrunner is a Kubuntu-based Linux alternative with multimedia codecs and other useful extras for a pleasing desktop experience. Version 17 includes the Plasma 5.4 desktop, LibreOffice 5, VirtualOffice 5, and an elegant collection of new wallpaper and graphic elements.



## ADDITIONAL RESOURCES

- [1] Knoppix: <http://www.knopper.net/knoppix/index-en.html>
- [2] Knoppix.net forums: <http://www.knoppix.net/forum/>
- [3] Knoppix 7.6 release notes: <http://knopper.net/knoppix/knoppix760-en.html>
- [4] Netrunner 17: <http://www.netrunner.com/netrunner-17-released/>
- [5] Netrunner forums: <http://forums.netrunner.com/>

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

# MOBILE USERS

search for us today at your digital newsstand!

## Only a swipe away!

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



## GOOGLE PLAY MAGAZINES

ADMIN Magazine

Linux Pro Magazine

Ubuntu User

Raspberry Pi Geek



## APPLE NEWSSTAND

ADMIN Magazine

Linux Pro Magazine

Ubuntu User

Raspberry Pi Geek



Visit our apps page for more information: [www.medialinx-shop.com/apps](http://www.medialinx-shop.com/apps)

# NEWS

Updates on technologies, trends, and tools

## THIS MONTH'S NEWS

### 08 Adobe Flash Renamed

- Gone in a Flash
- Rasp Pi Generates Weak SSH Keys

### 09 VMware Open-Sources Photon Controller

- Photon Controller Goes Open Source
- New Point and Click Exploit Kit Appears in the Wild
- Online News

### 10 Ransomware Reaches Linux

- Red Hat Releases RHEL 7.2
- Dell Systems Ship with Faulty Root Certificate
- Ransomware Reaches Linux
- Online News

## ■ Gone in a Flash

Adobe Flash, which has been the subject of many security alerts during the past two years, is getting a new start, with rewritten code and a new name. Adobe announced that it is renaming the Flash Professional tool included in its Creative Cloud suite Adobe Animate CC.

Adobe's announcement does not mention the numerous security problems associated with Flash; however, it does point out the need to incorporate native HTML5 support.

The Canvas feature of the new HTML standard offers a means for developers to create animations without the need for tools like Flash. The new Adobe Animate edition will leverage HTML5 features while still providing "first-class" support for Flash SWF and AIR formats.



© Misha, Fotolia.com

## ■ Rasp Pi Generates Weak SSH Keys

The Debian-based Raspbian Linux system, which runs on the tiny and popular Raspberry Pi single-board computer systems, appears to have a problem generating potentially weak SSH keys. Because the Rasp Pi doesn't come with a monitor, many Rasp Pi owners use SSH as a primary means of communicating with the system. And, although many consumer-end Pis are sitting behind firewalls on small home networks (uh, how safe are those home firewalls?), Internet-connected Raspberry Pis have started to appear as web servers, weather stations, remote photography experiments, and security cameras.

*The Register* quotes a Rasp Pi message board note, "Many Linux distributions stockpile random seed data during installation, and then use that to prime the pool during first boot-up, but Raspbian doesn't work that way – it starts up ready to go straight from the SD card, and thus suffers from low entropy."

On current Raspbian systems, hardware random number generation isn't enabled by default. The system uses the random data in the `/dev/urandom` pool to generate a host key, but the pool doesn't have enough entropy at the early stage where the keys are created. The Raspbian developers say they will fix the issue in the next release. In the meantime, users who are concerned about SSH security should use the Pi's on-board hardware random number generator to see the `/dev/urandom` file and regenerate SSH host keys.



## MORE ONLINE

## Linux Pro Magazine

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

## Off the Beat • Bruce Byfield

Review: Digital Painting with Krita 2.9

The painting application Krita has been limping along for years with scattered and partial documentation. Consequently, Scott Petrovic's *Digital Painting with Krita 2.9* is a welcome first step toward the kind of thorough documentation that it deserves.

## Pi-Top and the Do-It-Yourself Feeling

I spent several days last week assembling a Pi-Top, the do-it-yourself laptop that runs on a Raspberry Pi. I have limited experience with hardware, so if a mistake could be made, I probably made it. In the end, I concluded that I have a defective battery, although I have been able to explore the assembled laptop thanks to the power cord. Yet, despite the ups and down and the anxiety, I found myself enjoying the experience in a way that I haven't enjoyed computing since I first discovered Linux.

## Paw Prints • Jon "maddog" Hall

Raspberry Pie Zero: Just Chill a Bit

My friend Álvaro Justen recently showed me the announcement about the Raspberry Pi Zero (hereafter just known as "Zero"), which is selling for \$5, and is actually included for free in the *MagPi* magazine that sells in the UK. I take off my hat (again) to the Raspberry Pi people for creating such a computer at such a low price point, and although they have already manufactured "tens of thousands" of the units, in most cases the Zero was sold out "overnight" (or in some cases, even faster).

## Productivity Sauce • Dmitri Popov

Putting Context Menus in Nautilus to Practical Uses

Dolphin is not the only file browser that makes it possible to add custom commands to the context menu. If you happen to use Nautilus as your preferred file browser, you can add commands by placing executable scripts in the `~/local/share/nautilus/scripts/` directory. In addition to popular languages like Python and Perl, you can use good old Bash to write scripts.

## Tame ThinkPad's Fan with thinkfan

Most Linux distributions run perfectly well on ThinkPad laptops right out of the box, save for a few minor annoyances. Fan control is one of them.

## Photon Controller Goes Open Source

Virtualization giant VMware continues its effort to embrace the container business with the announcement that its Photon Controller tool has been officially released as open source. According to VMware, Photon Controller, which was officially announced in August of this year, is "... a new infrastructure stack optimized for containers and cloud-native apps."

The container revolution has added complications to VMware's market position as a leading virtualization alternative. Photon Controller will make it easier to integrate container support with VMware's virtualization management infrastructure, and it should put the company in a better position to compete with container management technologies from Docker, Google, and other vendors.

Open sourcing the Photon Controller should stir up activity among developers and integrators, which VMware hopes will lead to adoption of its complete virtualization and cloud management stack.

## New Point-and-Click Exploit Kit Appears in the Wild

A new chapter in the history of malware opens with the arrival of 3ROS, an exploit kit with a user-friendly GUI designed to make it easier for beginners to infect and exploit Windows, Mac, Linux, and Android victims. The tool, as described in the *Virus Bulletin*, resembles many admin GUIs that incorporate various tools into a simple point-and-click interface. The user can browse for an exploit or a malware payload through an easy series of menus.

According to a report in *The Register*, 3ROS will "... make pwning as easy as ordering pizza." The system is believed to originate from Russia or Ukraine and appears to be the basis for the Hunter exploit kit, which is also gaining popularity as a user-friendly attack tool.

Cyberattack used to be considered a black art practiced by shady and sophisticated computer experts, or at least, hackers with mid-level command-line skills. Even the so-called "script kiddies" seemed to aspire to the expert role and loved to navigate other people's directories from a remote terminal window. Although GUI-based attack tools have existed in the past, they typically weren't particularly polished or friendly. 3ROS and Hunter might indicate the further development of the intrusion industry. Advanced attackers can now focus on development and reconnaissance and leave the nuts-and-bolts exploit and delivery to beginners.

The actual malware included in 3ROS doesn't appear to be especially current or state of the art. The disadvantage of an elaborate GUI is that it takes longer to update and adapt the system to incorporate new exploits and stay ahead of detection services. Only time will tell if 3ROS and Hunter become a model for other exploit kits, or if the industry opts for more conventional but flexible approaches.



## MORE ONLINE

## ADMIN Online

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

#### Application Virtualization with Docker • Thorsten Scherf

Half-lives of products are becoming shorter and shorter on today's virtualization market. This year, the buzz is all about Docker.

#### What's New in PostgreSQL 9.4 • Bernd Helmle

The PostgreSQL Global Development Group recently introduced the new major version 9.4 of the popular free database, which includes innovative functions as well as a whole range of changes regarding speed and functionality.

#### Innovations in PowerShell 5 • Thomas Wiefel

Windows 10 brings an updated, fifth release of PowerShell that vastly simplifies the task of managing modules and software packages. The scripting language now also handles various formatted output from commands and selection lists.

## Red Hat Releases RHEL 7.2

Red Hat has announced the release of Red Hat Enterprise Linux (RHEL) 7.2, the latest version of the company's enterprise subscription-based Linux distribution. According to Red Hat, RHEL 7.2 includes security updates and some significant network performance improvements. Red Hat also continues its drive for better container and cloud



redhat.

support, with updates to Kubernetes, Cockpit, and the Docker engine, as well as enhancements to the Atomic host platform.


Perhaps the most interesting new development in RHEL

7.2 is its support for Red Hat's add-on Insight service, which the company describes as "an add-on operational analytics offering designed to increase IT efficiency and reduce downtime through the proactive identification of known risks and technical issues."

## Dell Systems Ship with Faulty Root Certificate

Several news sources report that users have discovered a preinstalled digital certificate shipped with Dell Inspiron 5000 series notebooks and the XPS 15 model laptop that could allow an attacker to successfully impersonate a trusted website. The eDellRoot certificate creates a gaping security hole that is similar to the problems created by the Superfish security issue experienced by Lenovo earlier this year.

Some users report that the certificate reinstalls itself when users attempt to remove it; however, Dell has posted instructions that it says will permanently remove eDellRoot. According to the Dell Knowledge Base, "The certificate is not malware or adware. It was intended to provide the system service tag to Dell online support, allowing us to quickly identify the computer model, making it easier and faster to service our customers. This certificate is not being used to collect personal customer information. It's also important to note that the certificate will not reinstall itself once it is properly removed using the recommended Dell process."



## Ransomware Reaches Linux

The security firm Dr. Web warns of a new strain of ransomware attacks that is targeting Linux server systems. The attack, which is known as Linux.Encoder.1, is currently preying on web servers. Recent versions gain entry through a flaw in the Magento CMS, but the attack could be adapted to other vectors.

Linux has largely been free of the ransomware craze, in which the attacker encrypts the victim's data and demands a ransom to restore the files. With all the Linux servers on the web, however, experts consider it a matter of time before intruders find ways to exploit Linux users. According to a report in the Register, this attack encrypts the public\_html, www, webapp, backup, .git, and .svn directory and asks for a payment of 1 Bitcoin to release the data.

The advice is the same as always: back up your data, and install security updates to minimize the possibility of intrusion.



© Jens Hilberger, 123RF.com



# DebConf 15

Heidelberg, Germany

<http://debconf15.debconf.org>

This year we hosted the largest annual Debian Conference ever which would not have been possible without all of our sponsors.

## Thank you.



Open Source  
for Business



sipgate



example42

ubuntu<sup>®</sup>  
Supported by Canonical



hudson river trading



FUJITSU

ARM



TWO SIGMA



DebConf16 will be held in Cape Town, South Africa



For more information please visit: <http://debconf16.debconf.org>



### Detect evidence of break-in attempts with host-based intrusion detection systems

# Stop, Thief!

A host-based intrusion detection system is a simple but powerful tool for finding traces of an attacker's footprint. *By Tobias Eggendorfer*

If your existence depends on extreme security, you could put your computers in a bunker and avoid external network connections, but unless you happen to be dealing with nuclear missiles, the realities of doing business might require you to configure some form of Internet access, and the price for joining the information superhighway is a need to watch for attackers attempting to penetrate your network.

IT professionals deploy a dizzying variety of tools for protecting the modern network, from firewalls to sophisticated intrusion prevention frameworks. (See the box titled "NIDS, IPS, HIDS.") One useful class of tools that receives less attention than it used to is the humble host-based intrusion detection system (HIDS). HIDS systems perform checks on a single host computer to look for indications of a recent attack.

HIDS works on the principle that every attacker leaves tracks. In forensics, this assumption is known as Locard's exchange principle: Interaction always causes change. The trick is to identify this change and draw the right conclusions [1].

The HIDS was once a primary line of defense, although today's enterprise networks often opt for bigger and more comprehensive network-based systems. In the spirit of Unix, however, a simple tool with a specific purpose is often a healthy addition to the admin toolkit. A typical HIDS system takes one of the following approaches for detecting the footprint of a past attack:

- track changes to the filesystem
- analyze logfiles
- search for malware

In this article, I describe some popular HIDS tools that highlight each of these approaches. You'll also learn about some all-in-one tools that combine these techniques with a single tool.

### AUTHOR

Tobias Eggendorfer is a professor of IT security in Weingarten (Baden-Wuerttemberg, Germany). He also holds foundation lectures on theoretical computer science and is a freelance IT consultant and external data protection officer (<http://www.eggendorfer.info>).



### NIDS, IPS, HIDS

When the alarm bells ring, the occupants of a house become aware that a burglar is present. In the world of computers, a network-based intrusion detection system (NIDS) silently raises an alert and helps the administrator track or intercept the attacker.

A related tool called an intrusion prevention system (IPS) automatically raises the drawbridge before the assumed attacker can even cross the moat; for example, by adjusting the firewall rules – either by blocking entry or by redirecting to a honeypot, where the attacker can rummage to their heart's content. Administrators can learn from observing the attacker's approach and thus improve the security of their systems.

Inspecting ransacked cabinets and broken window panes is an admittedly late way of detecting intrusions, but this approach is precisely what any host-based intrusion detection system (HIDS) does. Because latecomers typically miss out on everything in life, HIDS is not a very popular approach on today's networks. But this attitude fails to account for the fact that computer crackers often do not leave any obvious traces. Without an HIDS, the victim might not ever discover that they have had uninvited guests. A simple tool for finding information about what the attacker changed also tells the administrator where to look for malware, which is a big help during forensic analysis.

## COVER STORIES

**HOST-BASED IDS** 12  
Track the tell-tale signs of an intruder.

**SECURING LINUX** 16  
How can you secure your Linux system? Let us count the ways.



## Changes to the Filesystem

An attacker gets extreme pleasure from breaking into a third-party system through a buffer overflow. But once the intruder is inside, the work becomes much more monotonous. The attacker starts by installing a rootkit to provide convenient access. To set up a rootkit, the intruder adds software to the system and executes this software – which is something a system administrator might notice when running `ps`. To use a backdoor to control the target system, the attacker needs a network listener, resulting in an open port that would normally be visible to `netstat`. Attackers need to prevent their programs from appearing in the output of typical admin tools. Traditionally, intruders replaced the executables for common admin utilities like `ps` and `netstat` with fake versions that conceal their presence. More modern approaches, such as kernel rootkits, install a separate kernel module that blocks the visibility of rootkit processes in the kernel itself.

Administrators can identify Linux tools replaced by intruders on the basis of the different modification and creation times, the file sizes, or the fact that binaries use different inodes. In practical terms, administrators who are not using some form of automated security system will never notice the change because manual checks are far too time-consuming. Classical HIDS tools like Tripwire [2], AIDE [3], or the newer AFICK [4] provide an automated means for finding filesystem changes.

The syntax these free tools use in their configuration files is very similar. AFICK, which is written in Perl, actually advertises its Tripwire compatibility. All three tools create a list of critical files and directories, along with their modification and creation times, sizes, inode numbers, and permissions. Comparing these parameters with the results of later checks helps identify attacks.

Of course, capable attackers are likely to build their malware so that the file sizes are identical; they will set identical access times using `touch`, keep the file permissions, and bundle the modified files into the same inodes (`cat newps > ps`). The HIDS,

however, does not just keep records of the visible data but also of the files' hash values.

AIDE even offers the ability to apply multiple hash processes in parallel to make collisions even less likely: For MD5, the probability is less than  $2^{64}$ , and less than  $2^{128}$  for SHA-256.

The probability of a collision in both procedures is thus less than  $2^{192}$ , that is about  $1.5 \times 10^{58}$ . The comparison: If you play the lottery, your chances of winning a jackpot are  $7 \times 10^9$ .

AIDE is a good example of an HIDS tool that looks for file changes. On Debian-based systems, install AIDE with

```
apt-get install aide
```

You can modify the configuration to suit your needs. To change the configuration, edit the `/etc/aide/aide.conf` file, and for special cases, also `/etc/aide/aide.conf.d`. The objective is to make sure AIDE keeps an eye on all the relevant programs and configuration files.

The `aide.conf` file is also where your own macro definitions will reside. Studying the macro definitions can save you a huge amount of work when it comes to writing rules. AIDE detects any new files created – for example, as a consequence of a core dump, which is sometimes associated with a buffer overflow. (Core dumping is disabled by default on most current flavors of Linux [5].)

The default configuration offers meaningful rulesets for log files and for determining the system status. The configuration file even contains a good explanation for how to derive your own ruleset.

AIDE uses `update-aide.conf` to update its configuration; you can type `aide -D` to perform a syntax check. The command

```
aideinit -y -f
```

initializes the database. AIDE now computes the hashes – a process that took some 10 minutes on the virtual machine in my lab. Because a cron job is automatically installed when you install AIDE, the tool checks the changes and mails them to the root user. AIDE can also output the results to any URL (`aide.conf report_url` entry) so you can collect reports in a central location.

AIDE emulates its role models so closely – you can easily migrate the setup I just described to Tripwire.

## Logs, Logs, Logs

Any admin's daily chore list should include checking the log files. Did a process unexpectedly terminate? Was the process restarted – either by a watchdog or by an attacker attempting to conceal their tracks? Do any other clues indicate attempted intrusions? Regularly investigating the logs helps you identify problems. Fortunately, some tools help administrators combat any tendency toward laziness – these tools analyze logfiles, filter out potentially critical messages, and mail the results to the admin.

Logsentry [6] (formerly called Logcheck) is a useful analysis tool: This simple but complete HIDS consists entirely of shell scripts and configuration files. Logsentry saves adminis-



trators the trouble of thinking about the right `grep` commands to detect attacks on the basis of logfile analysis. Download Logsentry, then install it with the following commands:

```
tar -xzf logcheck-1.3.17.tar.gz
cd logcheck-1.3.17
make linux
```

It is a good idea to set up a cron job to check your logs regularly. Still working as root, type

```
crontab -e
```

to open the crontab for editing. The line

```
0 * * * * /usr/local/etc/logcheck.sh
```

makes sure the `logcheck.sh` test script runs every hour. If you want to check more than the three logfiles configured by default, modify `logcheck.sh` from line 172 using the preceding lines as an example. Logcheck defines its rules in multiple configuration files below `/etc/logcheck`.

Everything in `logcheck.hacking` appears right at the top of the mail message to root. The `logcheck.violations` follow slightly lower down. Everything that is not excluded by `ignore` ends up in the message – including *Unusual system events*. The only important thing is to make sure you don't leave

`logcheck.ignore.empty` – this would be equivalent to a wildcard that ignores everything, thus rendering the tool useless.

Logsentry and its siblings save administrators some work, but this is actually *all* they do. Anyone who needs to inspect their logfiles regularly is likely to have programmed a simple search script of their own that filters out non-critical messages. Logsentry does exactly the opposite: When you set up the tool, you need to know exactly what critical messages look like. The tool only filters for known suspects, not for assumed non-suspects.

Also, the practice of checking logfiles relies on the attacker lacking ability. Any capable attacker will clean the logfiles on the guest system before leaving, and Logsentry will not find anything suspicious. The attacker will probably have enough time to clean up, depending on how often cron runs Logsentry. In general, a tool that looks for attackers by checking logfiles is only useful if it runs in combination with other systems.

## Searching for Malware

If log analysis and file checks reveal the aftermath of an infection, it makes sense to look for the malware code. This task is handled by malware scanners, which also exist for Linux. ClamAV [7] is one example. The demarcation between malware scanners and host-based intrusion detection systems is not clearly defined; most antivirus scanners use both signatures and heuristics to detect attacks.

Programs such as Chkrootkit [8] and RKHunter [9] (a tool written in Perl) use a more specific approach; both of these tools were last updated in 2014, and you can install both on Ubuntu as root with the following commands:

```
apt-get install chkrootkit
apt-get install rkhunter
```

Chkrootkit is a shell script that attempts to identify rootkits based on various parameters. You can run Chkrootkit without specifying any parameters. Among other things, the Chkrootkit script uses strings to detect suspicious strings, checks the last log to see if anyone has removed entries, checks to see if `ps` is not showing certain process IDs, and looks for missing information in `/proc/` (Figure 1).

RKHunter has a more sophisticated approach. The hunter's configuration file is named `/etc/rkhunter.conf`. (I had some issues setting up RKHunter on my Ubuntu system – see the box titled "Ubuntu Fix.")

You can initialize RKHunter and check the system for rootkits and conspicuous files:

### UBUNTU FIX

I had to do some troubleshooting when installing RKHunter on my Ubuntu system. The program had a need for `lwp-request`, although it was not installed. I changed this as follows:

```
cpan Mock::LWP::Request
```

However, Perl dumped the tool into `/usr/local/bin/lwp-request`, much to RKHunter's surprise – this means that you need to edit the `/usr/bin/` path manually in the configuration file.

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

**ADMIN HPC**  

- Monitoring HPC Systems
- Watson in the Cloud
- SC14 Call for Workshop Proposals
- Data Visualization Tools

**Highlights**  
**Monitoring HPC Systems: What Should You Monitor?**  
In rapidly growing HPC installations, you need to understand what is happening within the system to make improvements or simply to justify the...  
**Watson in the Cloud**  
IBM launches three new services for its well-trained, Jeopardy-busting behemoth.  
**SC14 Calls for Workshop Proposals**  
Attendees should provide "interaction and in-depth discussions of stimulating topics of interest to the HPC community."  
**Data Visualization Tools**  
Logfiles contain extremely useful traces of everyday events – if only they weren't so boring. Fortunately, special tools help you generate and...  
**Most Read**  
**Tuning I/O Patterns in Fortran 90**  
In the second article of this three-part series, we look at simple while examples in Fortran 90 and track the output with trace to see how...  
**Build a Network Attached Storage System with FreeNAS**  
We provide an overview of FreeNAS and ZFS, its main filesystem, then show you how to set up and maintain a FreeNAS installation.  
**XSEDE Releases Toolset**  
The new tools will improve compatibility of local campus clusters with the XSEDE supercomputer network.

**Further Reading**

- Fixing Disks with Parlet Magic
- Exploring the Xeon Phi
- Tuning I/O Patterns in Python
- Blue Waters Open for Business

**Partner Link**  
Intelligent High Availability DAS for Windows and Linux

**Raspberry Pi Geek Offer**  
  
 Subscribe to Raspberry Pi Geek starting with Issue 2 and receive either Issue 1 or Raspberry Pi Handbook for FREE!  
 Use code HCLSRPG013 at checkout.  
 Click here for full details. (Offer expires January 31, 2014.)

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



```
rkhunter --propupd
rkhunter --check
```

Figure 2 gives an impression of what RKHunter considers to be a potential attack. You can run both Chkrootkit and RKHunter on a daily basis using cron, and they will thus become part of your HIDS.

## All in One

Instead of three tools that complement one another, some solutions combine all three functions into a single tool. Whether or not this approach is superior is very much a question of taste. Notable all-in-one solutions include Samhain [10] and OSSEC [11]. Both of these all-in-one tools offer an agent mode that runs sensors on remote systems, which then transfer their data to a central computer for evaluation.

```
ROOTDIR is '/'
Checking `amd'... not found
Checking `basename'... not infected
Checking `biff'... not found
Checking `chfn'... not infected
Checking `chsh'... not infected
Checking `cron'... not infected
Checking `crontab'... not infected
Checking `date'... not infected
Checking `du'... not infected
Checking `dirname'... not infected
Checking `echo'... not infected
Checking `egrep'... not infected
Checking `env'... not infected
Checking `find'... not infected
Checking `fingerd'... not found
Checking `gpm'... not found
Checking `grep'... not infected
Checking `hdparm'... not infected
Checking `su'... not infected
Checking `ifconfig'... not infected
Checking `inetd'... not infected
Checking `inetdconf'... not found
Checking `identd'... not found
Checking `init'... not infected
Checking `killall'... not infected
Checking `ldsopreload'... not infected
Checking `login'... not infected
Checking `ls'... not infected
Checking `lsof'... not infected
Checking `mail'... not infected
Checking `mingetty'... not found
Checking `netstat'... not infected
Checking `named'... not found
Checking `passwd'... not infected
```

Figure 1: Chkrootkit making various attempts to detect signs of typical rootkit activity.

```
Checking the network...
Performing checks on the network ports
Checking for backdoor ports [ None found ]
Checking for hidden ports [ Skipped ]
Performing checks on the network interfaces
Checking for promiscuous interfaces [ None found ]
Checking the local host...
Performing system boot checks
Checking for local host name [ Found ]
Checking for system startup files [ Found ]
Checking system startup files for malware [ None found ]
Performing group and account checks
Checking for passwd file [ Found ]
Checking for root equivalent (UID 0) accounts [ None found ]
Checking for passwordless accounts [ None found ]
Checking for passwd file changes [ None found ]
Checking for group file changes [ None found ]
Checking root account shell history files [ OK ]
Performing system configuration file checks
Checking for an SSH configuration file [ Found ]
Checking if SSH root access is allowed [ Warning ]
Checking if SSH protocol v1 is allowed [ Not allowed ]
Checking for a running system logging daemon [ Found ]
Checking for a system logging configuration file [ Found ]
Checking if syslog remote logging is allowed [ Not allowed ]
Performing filesystem checks
Checking /dev for suspicious file types [ None found ]
Checking for hidden files and directories [ Warning ]
```

Figure 2: One part of RKHunter's work is checking for illegal material.

If you are paranoid and have sufficient free time to read all the email, you will be happy to hear that multiple HIDS tools running in parallel do not get in each other's way. In other words, you can run a combination of Tripwire, AIDE, AFICK, Samhain, and OSSEC on a single computer and let the systems monitor one another. But some caution is advised: Updates that arrive at the same time can easily send the system into an infinite loop.

Look online for detailed HowTo information on Samhain [12] and OSSEC [13]. Using these tools is not exactly rocket science if you know how the components work.

Having said this, each additional line of code adds the risk of security-critical programming errors. For example, OSSEC 2.7 to 2.8.1 had a privilege escalation problem in June 2015: CVE-2015-3222 [14]. Version 2.8.2 was the first to provide a working patch. Whatever tool and version you use for your network, the task of managing your HIDS should include watching for security updates: You will not want your security tools to cause new vulnerabilities.

## Tracking Down the Thief

Host-based intrusion detection systems come in many varieties, with features that often complement one another, and pre-built packages are available that combine all of the required functions. HIDS logs give admins who have fallen victim to a successful attack important information that lets them understand the attacker's actions, plug the gaps, and possibly even bring the evildoer to justice.

Many tools reach their performance limit when the attacker has root privileges and sufficient time to disable security mechanisms and clean up their tracks. Because an HIDS by design only detects the results of an attempted attack, it can only be one building block in the overall security architecture. A network intrusion detection system, firewalls, restrictive user rights, regular security audits of your systems, and efforts to raise your users' security awareness complete the puzzle. ■■■

## INFO

- [1] Locard's principle: [https://en.wikipedia.org/wiki/Locard%27s\\_exchange\\_principle](https://en.wikipedia.org/wiki/Locard%27s_exchange_principle)
- [2] Tripwire: <http://sourceforge.net/projects/tripwire/>
- [3] AIDE: <http://aide.sourceforge.net>
- [4] AFICK: <http://afick.sourceforge.net>
- [5] HowTo Enable Core Dumps: [http://en.linuxreviews.org/HOWTO\\_enable\\_core-dumps](http://en.linuxreviews.org/HOWTO_enable_core-dumps)
- [6] Logsentry: <http://sourceforge.net/projects/sentrytools/files/logcheck%201.x/logcheck-1.1.1/>
- [7] ClamAV: <http://www.clamav.net/index.html>
- [8] Chkrootkit: <http://www.chkrootkit.org>
- [9] RKHunter: <http://rkhunter.sourceforge.net>
- [10] Samhain: <http://www.la-samhna.de/samhain/>
- [11] OSSEC: <http://www.ossec.net>
- [12] Samhain manual: [http://www.la-samhna.de/samhain/MANUAL-2\\_4.pdf](http://www.la-samhna.de/samhain/MANUAL-2_4.pdf)
- [13] OSSEC manual: <http://ossec-docs.readthedocs.org/en/latest/manual/installation/index.html>
- [14] CVE-2015-3222: <http://www.openwall.com/lists/oss-security/2015/06/11/1>



### Setting up a secure Linux server

# Lock It

When you set up a new Linux server, take a few extra steps to ensure the system is truly secure. *By Chris Binnie*

I recently signed up for a cloud-based server system to replace my last co-located physical server, which had given up the ghost. The sign-up for the new system was painless, aside from the need to relinquish my precious credit card details, and within no more than a minute, I had a brand new Debian 8.1 “Jessie” server with a static IPv4 address.

It’s relatively unusual for me to manually build boxes from scratch. However, because this Debian system is destined for service as a development server, I thought, for once, I’d take note of the steps I took to secure it.


The eclectic mixture of the following steps comprise a unique, personal-preference approach, but perhaps this summary will give you food for thought about how to secure your own servers. Start with these steps and add your own. Some of the settings you might be thinking about will be missing from this list (because the spectrum of possible security steps is simply too large to cover succinctly), but you might also find some ideas you hadn’t considered.

This article shows what I did to set up a Debian Jessie server from scratch. The plan, as always, is to keep the number of packages to a minimum to help reduce security holes and reduce the need for ongoing package updates and maintenance.

This article focuses on hardening a single Debian server system. Keep in mind that maintaining the overall security of your network requires many other steps, including maintaining security for routers, firewalls, intrusion prevention services, and much more.

### Services

Once I had changed my DNS to point at the static IP address proffered by the cloud host, I logged into the system over SSH and immediately changed the root password to one of my liking (and which hadn’t arrived over email!). A simple suggestion to the cloud provider would be, procedurally speaking, only offering the initial root password from their web interface (which uses HTTPS). Email is horribly insecure, and we’re talking about a fresh server installation that is supposed to be trusted for a long time to come. To their credit, the cloud provider does encourage users to create an SSH key immediately which is the safest way of accessing SSH as opposed to passwords.







```

jcasad@ThinkPad:~$ sudo free -m
[sudo] password for jcasad:
              total        used          free      shared  buffers   cached
Mem:           1863         1623           240         168          20         436
-/+ buffers/cache: 1165        698
Swap:          5491         1001         4490
jcasad@ThinkPad:~$

```

Figure 1: Checking RAM memory with `free -m`.

Once I logged in as “root” user over SSH, I got to work. The superuser should *never* be able to log in remotely over SSH without a user of less privilege to let them in through the front door first, so I shuddered at the sight of the prompt initially. I’ll deal with that soon.

The first command I ran was to check that my new ISP had provisioned the size of machine correctly. I checked the RAM specifically:

```
# free -m
```

See Figure 1 for an example of the `free -m` command. In my case, the command showed me under the top line of the column “used” that I had only used a remarkable 93MB of RAM to boot up my Debian instance without many running services (compare that to many, many more megabytes required by a fresh Windows installation).

Additionally, it showed me 401MB of free RAM under the “free” column, which meant I had been specified a tiny 512MB instance as expected. That’s no problem for my needs when Linux is so fantastically efficient.

Next, I rapidly moved onto securing the box a little more to my liking. This next step was to shut down any services that exposed network ports to the Internet; initially, all I wanted open was my SSH server. To achieve this, I ran the following command:

```
# lsof -i | grep LISTEN
```

On the left-hand side, I saw service names (usually correlating to start/stop filenames in the `/etc/init.d` directory) and on the right-hand side the network ports which they had opened. Those ports are usually listed as “service names” translated by the file `/etc/services` from numbers to names.

I immediately spotted two services that weren’t welcome (and potentially a security risk) and sought to shut them down. I figured out that `/etc/init.d/rpcbind` was responsible for the first service so, remembering that `init` was a thing of the past and this was a Systemd build, I promptly typed:

```
# systemctl stop rpcbind
```

followed by this command so `rpcbind` wouldn’t start up again after a reboot:

```
# systemctl disable rpcbind
```

Then, having seen that another port was open, I eventually figured out that the Systemd startup service called `nfs-common` was responsible and disabled it also, simply replacing `rpcbind` with that service name in the preceding commands.

If you get stuck (e.g., because you’re not using Debian Jessie’s version of the `lsof` command but another distribution’s) try tweaking the command below. It works for me on Jessie. The command searches for any listening services and outputs the name of the service along the user account responsible for spawning the service.

```
# lsof -i | grep LISTEN | awk '{print $1,$3}'
```

```
rpcbind rpc
master root
sshd root
httpd root
httpd apache
```

In addition to shutting off services, you might also want to use this command to experiment with changing services to more secure user accounts. With some trial and error, you can successfully switch some of these services to non-root user accounts, which can only be a good thing for increasing your security.

## Ssshhhh!

My next step was to move the SSH daemon `sshd` off the predictable port 22. Edit the file `/etc/ssh/sshd_config` carefully, uncommenting the top line shown in Figure 2 and changing 22 to a port number of your choice like 2222. This precaution will stop automated port 22 attacks from harassing your server.

You’ll also want to prevent the root user from logging in as we mentioned. Set the `PermitRootLogin` parameter to `no`:

```
PermitRootLogin no
```

Next, I’ll quickly create that less-privileged user that I mentioned, but first, I’ll make sure only the users I want to have permission to log in remotely. I add a line to the SSH config file (at the top of the file to remind me that I’ve made these changes the next time I view the file):

```
AllowUsers lionel, luis, neymar
```

Only three users can SSH into the box now. I need to make I create these accounts before logging out or restarting the OpenSSH server or I will be locked out. With a few pertinent questions, I type:

```
# adduser lionel
```

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
```

Hardening the SSH config file further is for another day (but you should definitely do it and

Figure 2: In the `sshd_config` file, alter the top line to say port 2222 or similarly obscure port number after uncommenting the line by removing the hash or pound symbol.



understand what you have changed). Restart `sshd` to make the changes live:

```
# systemctl restart ssh
```

To put my mind at rest, I'll do a quick check to ensure SSH is still allowing me to log in and peek at Systemd's `syslog` (by running `journalctl -a` and looking at the last few entries) for errors. I still need to become fully familiar with the freshly released Systemd errors, but I don't see anything in the file relating to the changes I've made, so I'm happy for now.

The next step is to open up `/etc/hosts.allow` and fiddle with TCP wrappers briefly by removing anything else that isn't commented and adding the IP addresses I connect from:

```
sshd: 11.22.33.44, 123.123.123.123, 9.8.7.6
```

Now only three IP addresses can log in over SSH, and the user must be one of the triumvirate's user names configured previously, such as `lionel`. I then add the following line to `hosts.deny`:

```
sshd: ALL
```

And, as if by magic, I'm all set. Remember not to shut your initial login window until you've tested a fresh SSH login.

## Command History

I always struggle with `Ctrl-R` to reverse search through the command history, so I prefer to enable `Page-Up` and `Page-Down` key functionality. This change is not security related, but it helps me zip through the command line much more efficiently.

Figure 3 shows the two lines in the file `/etc/inputrc` file I need to uncomment. I'll just wait for the next reboot to set it live; it's not urgent unless you've got loads of typing ahead.

## Home Sweet Home

Not all distributions tighten up your home directories by default. Assume my user name is `lionel`, and I don't want `Luis` or `Neymar` to see the files in my home directory. It's easy to keep other users out of my home directory by running the following command as root:

```
# chmod -R 750 /home/lionel
```

If you're keen to stop all home directories being accessed by others, simply replace `lionel` with an asterisk.

## Hey, Mr. Postman

The next step is to set up the `logwatch` tool [1] and configure it to send you a daily digest of any important events found in your system logs. Bear in mind that, when you're finished with these simple steps, you'll also have a fully functional

```
"\e[5~": history-search-backward
"\e[6~": history-search-forward
```

**Figure 3:** Boost your typing speeds on the command line with `Page-Up` and `Page-Down` search capability.

SMTP server. Clearly, I'm zooming through these commands for brevity; I en-

courage you to take a closer look if you're not sure what you're doing.

Straight out of the box, the Postfix mail server will be fully functioning. Choose *Internet Site* and add your full hostname (including FQDN) during the configuration process. With a quick shuffle, you can then change the old-school, Sendmail-style `/etc/aliases` file so that your personal address receives your root user email.

Install Postfix as follows:

```
# apt-get install postfix
```

Edit the `aliases` file and add your external address so you receive the root user's mail:

```
# pico -w /etc/aliases
```

Make sure the line starting `root:` looks similar to the following:

```
root: chris@binnie.tld
```

Update the aliases DB file by executing the Postfix `newaliases` command:

```
# newaliases
```

Finally, install the `Logwatch` package; even as a fresh install it's fabulously functional:

```
# apt-get install logwatch
```

The easiest way to execute `Logwatch` is to set up a cron job that runs at the default `cron.daily` predefined time, which is 06:25hr on my Jessie 8.1 build. Adjust the time in the `/etc/crontab` file or move the `/etc/cron.daily/00logwatch` file somewhere else and execute it specifically at another time from within the `/etc/crontab` file.

Next up run `Logwatch` and then check your inbox:

```
# /etc/cron.daily/00logwatch
```

The results are mind-blowingly useful. For more on monitoring your logs with `logwatch`, see the article at the *ADMIN Magazine* website [2].

`Logwatch` provides invaluable information about your system (including disk space usage, logins over SSH, new users/groups, service restarts, and unusual log entries). Additionally, you receive a list of the packages that you've installed following the completion of your minimal install.

The packages are also conveniently separated into `Installed`, `Upgraded`, and `Removed` sections. I find these lists are an excellent point for future reference, archived in my email with the subject line "Logwatch for <hostname> (Linux)."

If you ever encounter discrepancies between two servers due to package version or functionality, it is a two-second lookup to scan over how they were built. If you don't want root to directly receive mail via your `aliases` file, you can add `MailTo =` followed by an email address in the config file `/etc/logwatch/conf/logwatch.conf`. You might need to create this file if it doesn't already exist.



I would be remiss not to mention the task of securing of your shiny new SMTP server further. Even if relaying is switched off by default, you still need to take some steps to secure your mail server. The Ask Ubuntu site has a useful post on securing a Postfix mail server [3].

## Passwords

If I am expecting a few different (non-sys admin users) to be logging into this box, I should also consider password security. If you have ever stumbled across a list of common passwords, you'll know why admins are very concerned about letting everyday users have access to server systems [4].

The quality of the typical user password is not good news if you haven't locked down access to your server by IP address or some other alternative method.

Without these precautions you are definitely vulnerable to automated dictionary attacks. Thankfully, you can strengthen your password rules on Linux with a sophisticated package called `libpam-cracklib`.

The following configuration works on my Debian "Wheezy" system. Install `libpam-cracklib` with:

```
# apt-get install libpam-cracklib
```

Then, edit the following file:

```
# pico -w /etc/pam.d/common-password
```

Adjust the following command with the necessary settings (comment out the old config line rather than deleting it so you can revert if necessary):

```
password requisite 2
pam_cracklib.so retry=2 2
minlen=10 difok=6 dcredit=2 ocredit=2
```

This hardened configuration means a user is only permitted two password error retries and demands at least 10 characters for a password.

My favorite setting (to annoy colleagues) is `difok=6`, which means only six characters of the last password can be reused and the last two entries must be a minimum of two numbers and two symbols (or "special characters") present within the password. To activate these settings, on Wheezy at least, you can run the following command:

```
# pam-auth-update
```

Try a few tests – carefully, so you can still log back in. For troubleshooting, you might want to check whether `UsePAM yes` is present in the file `/etc/ssh/sshd_config` file. Restart the

**TABLE 1: Cracklib Password Enforcement Options**

Option	Description
<code>retry=N</code>	Prompt user at most N times before returning with error. The default is 1.
<code>difok=N</code>	This argument will change the default of 5 for the number of character changes in the new password that differentiate it from the old password.
<code>minlen=N</code>	The minimum acceptable size for the new password (+1 if credits are not disabled as per the default). More detail for this option is available in the manual.
<code>dcredit=N</code>	(N >= 0) The maximum credit for having digits in the new password. If you have less than or N digits, each digit will count +1 towards meeting the current <code>minlen</code> value. The default for <code>dcredit</code> is 1, which is the recommended value for <code>minlen</code> less than 10. (N < 0) This is the minimum number of digits that must be met for a new password.
<code>ucredit=N</code>	(N >= 0) The maximum credit for having upper case letters in the new password. If you have less than or N uppercase letters, each letter will count +1 towards meeting the current <code>minlen</code> value. The default for <code>ucredit</code> is 1, which is the recommended value for <code>minlen</code> less than 10. (N < 0) This is the minimum number of uppercase letters that must be met for a new password.
<code>lcredit=N</code>	(N >= 0) The maximum credit for having lowercase letters in the new password. If you have less than or N lowercase letters, each letter will count +1 towards meeting the current <code>minlen</code> value. The default for <code>lcredit</code> is 1, which is the recommended value for <code>minlen</code> less than 10. (N < 0) This is the minimum number of lowercase letters that must be met for a new password.
<code>ocredit=N</code>	(N >= 0) The maximum credit for having other characters in the new password. If you have less than or N other characters, each character will count +1 towards meeting the current <code>minlen</code> value. The default for <code>ocredit</code> is 1, which is the recommended value for <code>minlen</code> less than 10. (N < 0) This is the minimum number of other characters that must be met for a new password.
<code>minclass=N</code>	The minimum number of required classes of characters for the new password. The default number is zero. The four classes are digits, upper and lower letters and other characters. The difference to the credit check is that a specific class if of characters is not required. Instead N out of four of the classes are required.
<code>maxrepeat=N</code>	Reject passwords that contain more than N same consecutive characters. The default is 0, which means that this check is disabled.
<code>maxsequence=N</code>	Reject passwords that contain monotonic character sequences longer than N. The default is 0, which means that this check is disabled. Examples of such sequence are "12345" or "fedcb". Note that most such passwords will not pass the simplicity check unless the sequence is only a minor part of the password.
<code>maxclassrepeat=N</code>	Reject passwords that contain more than N consecutive characters of the same class. The default is 0, which means that this check is disabled.
<code>reject_username</code>	Check whether the name of the user in straight or reversed form is contained in the new password. If so, the new password is rejected.
<code>gecoscheck</code>	Check whether the words from the GECOS field (usually full name of the user) longer than 3 characters in straight or reversed form are contained in the new password. If any such word is found, the new password is rejected.
<code>enforce_for_root</code>	An "error" or "failed check" message prevents the root user from changing the password. This option is off by default, which means that the message about the failed check is printed, but root can change the password anyway.
<code>use_authok</code>	This argument is used to force the module to not prompt the user for a new password but use the one provided by the previously stacked password module.



OpenSSH server with `systemctl restart ssh` if you change the configuration.

Table 1 shows some of the excellent options for enforcing stricter passwords from the `pam_cracklib` manual, which covers `libpam-cracklib`. Experiment with these settings to your heart's content.

If you want to flex your Cracklib muscles, you could also go one step further and introduce your own dictionaries for Cracklib to check against. Look online for an old but nicely-written piece on using Cracklib with Django [5].

### Other Ideas

My next thoughts would be how to keep a close check on the filesystem. You might want to create MD5 checksums of your key files to compare them later if you suspect a compromise. Alternatively, you could install a filesystem integrity check, such as AIDE (Advanced Intrusion Detection Environment), to keep an eye on filesystem changes.

You might also want to configure iptables to catch any services that fire up unexpectedly through either conspiracy or misconfiguration.

Another interesting option for securing your system is port knocking [6]. Port knocking lets you close off your SSH server port to the outside world – switch off pings and make your server completely invisible. Then, the only way to wake up the SSH Server is to connect to a predefined sequence of ports.

If you are offering your users remote access through SFTP, you might want to lock down their access by forcing them to work within their home directories and not to run other system commands. Simply configure an SFTP chroot of sorts [7].

You might also consider employing a tool that automatically examines and locks down your system, reporting on security issues. Tools such as Bastille Linux or Lynis provide this kind of securing auditing service.

The upside of tools like Bastille Linux and Lynis is they speed up the process of security hardening and make it so you don't need to know all of the ins and outs of what you are securing in great detail. The downside is that sys admins become less interested in understanding security at a deeper level and ultimately risk losing the benefits of learning how to secure systems themselves. It's very easy to type Y and accidentally render a system useless if you don't understand exactly what the application is asking you.

### It Never Ends

Clearly security isn't just about the server configuration. I could have gone into much further detail about configuring your server and looked at other parts of the network as well. What about checking for dangerous files with setuid permissions? You would only want to allow a select few files to have permissions that mean the user running the file inherits the permissions of the file.

Try tinkering with this command-line example, if I've caused you concern in this regard:

```
# find / -xdev \( -perm -4000 \) -type f -print0 | xargs -0 ls -hal
```

With this command we can check which files, on our entire system, are configured with permission 4000 or the SUID bit

(which will show up as `s` in the file's permissions after a directory listing). You can remove any nasty surprises that show the `s` bit as being present with:

```
# chmod -s /home/chrisbinnie/nasty_file
```

To protect my filesystems further, I might consider SELinux at this stage of the build too. I might have also go secret-agent-style and have a peek at the US National Security Agency's recommendations [10].

When putting Apache on this box, I would immediately hide the `cgi-bin` directory and change the `ServerTokens` settings to `Prod`. I also would encourage a quick refresher on hardening Apache by reading the excellent article at the Geekflare site [11].

The never-ending list of security checks simply keeps on going. If you want to read further, look at OpenSCAP from the National Institute of Standards and Technology (NIST) [12]. OpenSCAP works on "automatically verifying the presence of patches, checking system security configuration settings, and examining systems for signs of compromise."

I hope this article has piqued your interest enough to keep on exploring the fascinating side of system security. It is a good idea to build and configure your server systems with a close eye on security if you want to return at a later date and find them working in the same way that you left them. ■■■


### INFO

- [1] Logwatch: <http://www.logwatch.org/>
- [2] "Lean on Logwatch" by Chris Binnie, *ADMIN*, issue 25, 2015: <http://www.admin-magazine.com/Archive/2015/25/Learn-on-Logwatch/%28language%29/eng-US>
- [3] How to Secure Postfix on Ubuntu Server: <http://askubuntu.com/questions/418340/how-to-secure-postfix-on-ubuntu-server>
- [4] "Revealed: The Top 25 Most Common Passwords": <http://news.sky.com/story/1412124/revealed-the-top-25-most-common-passwords>
- [5] "Using Cracklib to Require Stronger Passwords": <http://thegarywilson.com/blog/2006/using-cracklib-to-require-stronger-passwords/>
- [6] "Protect Your Network with Port Knocking" by Chris Binnie, *ADMIN*, issue 23, 2014: <http://www.admin-magazine.com/Archive/2014/23/Port-Knocking>
- [7] "Limiting Access with SFTP Jails on Debian and Ubuntu": <https://www.linode.com/docs/tools-reference/tools/limiting-access-with-sftp-jails-on-debian-and-ubuntu>
- [8] Bastille Linux: <http://bastille-linux.sourceforge.net/>
- [9] Lynis: <https://cisofy.com/lynis/>
- [10] NSA Guide for Secure OS Configuration: [https://www.nsa.gov/ia/mitigation\\_guidance/security\\_configuration\\_guides/operating\\_systems.shtml](https://www.nsa.gov/ia/mitigation_guidance/security_configuration_guides/operating_systems.shtml)
- [11] "Apache Web Server Hardening and Security Guide": <http://geekflare.com/apache-web-server-hardening-security/>
- [12] OpenSCAP Portal: [http://www.open-scap.org/page/Main\\_Page](http://www.open-scap.org/page/Main_Page)

Shop the Shop

shop.linuxnewmedia.com

# RASPBERRY PI ADVENTURES



**COOL  
PROJECTS  
FOR GEEKS  
OF ALL  
AGES**



## RASPBERRY PI ADVENTURES

is a one-volume special edition magazine for curious Raspberry Pi beginners. This easy, hands-on guide starts with an introduction to computers and offers a series of special hands-on projects illustrating many of the most popular uses for the Raspberry Pi.

**Great Gift Idea!**

**ORDER YOUR VERY OWN ISSUE!**



ORDER ONLINE:

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

Tool tests on the fast track *By Uwe Vollbracht*

# TOOL TIPS

## Mtools 2.3

Testing multicast connections

Source: <https://github.com/troglobit/mtools>

License: Public Domain

Alternatives: Mz, Netsniff-ng

```
Terminal - vollbracht@LMLab-1504_b:~/extract/tooltips/mtools-master
vollbracht@LMLab-1504_b:~/extract/tooltips/mtools-master$ ./msend
Now sending to multicast group: 224.1.1.1
Sending msg 1, TTL 1, to 224.1.1.1:4444:
Sending msg 2, TTL 1, to 224.1.1.1:4444:
Sending msg 3, TTL 1, to 224.1.1.1:4444:
Sending msg 4, TTL 1, to 224.1.1.1:4444:
Sending msg 5, TTL 1, to 224.1.1.1:4444:
Sending msg 6, TTL 1, to 224.1.1.1:4444:
Sending msg 7, TTL 1, to 224.1.1.1:4444:
Sending msg 8, TTL 1, to 224.1.1.1:4444:
Sending msg 9, TTL 1, to 224.1.1.1:4444:
Sending msg 10, TTL 1, to 224.1.1.1:4444:
vollbracht@LMLab-1504_b:~/extract/tooltips/mtools-master$ ./mreceive
Now receiving from multicast group: 224.1.1.1
Receive msg 1 from 192.168.250.193:4444:
Receive msg 2 from 192.168.250.193:4444:
Receive msg 3 from 192.168.250.193:4444:
Receive msg 4 from 192.168.250.193:4444:
Receive msg 5 from 192.168.250.193:4444:
Receive msg 6 from 192.168.250.193:4444:
Receive msg 7 from 192.168.250.193:4444:
Receive msg 8 from 192.168.250.193:4444:
Receive msg 9 from 192.168.250.193:4444:
vollbracht@LMLab-1504_b:~/extract/tooltips/mtools-master$
```

Multicast connections are used to send data from a single sender to a group of receivers at the same time. This technology is mainly used in multimedia applications – for example, distributing live streams. Two tools, `msend` and `mreceive`, allow users to perform simple tests.

To check the sender, call `msend`. The tool then uses the default network interface in the IPv4 224.1.1.1 address range to send UDP packets on port 4444 at intervals of one second. At the receiving end, you need to call `mreceive` – again without any options – to listen to packets in the same IPv4 range on port 4444. To send or receive the data on a specific interface, you need to define this at the command line using `-i`; `mreceive` will also take multiple interfaces if needed.

By default, `msend` sets the TTL of the packets to 1, thus causing the first router to drop them and keeping them on the local network. Admins who need to test multiple routers on large networks can increase this value using `-t`.

★★★★☆ Mtools provides two practical helpers for testing multicast connections. It loses points for a lack of IPv6 support. ■■■

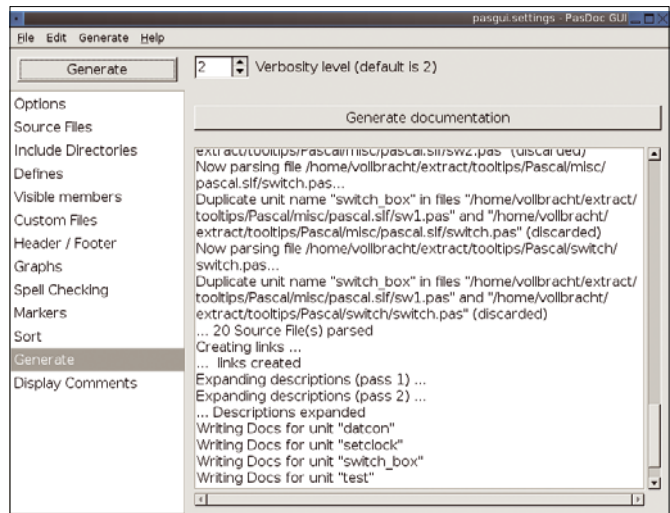
## PasDoc 0.14.0

Pascal source code documentation

Source: <http://pasdoc.sipsolutions.net>

License: GPLv2

Alternatives: None



Documenting their own source code is a pesky task for many developers. PasDoc relieves part of this burden for Pascal and Delphi programmers. It works its way through the units and uses the comments, functions, and object classes to generate clear-cut API documentation.

The project website has versions of the tool, which is written in Pascal, for Windows, OS X, and Linux. The tarball with the Linux version includes both a command-line tool and a graphical variant. You can use the partitions on the left-hand side to decide which directories and files PasDoc should include in the documentation. Optionally, the tool adds headers and footers to the documentation and even checks the spelling.

PasDoc generates documents in LaTeX and HTML format. To modify the layout of the HTML files, the program uses CSS files, which the user can define. The roadmap includes XML support for future versions.

★★★★★ The command-line variant is excellent for larger projects, giving programmers the ability to update program manuals automatically. ■■■

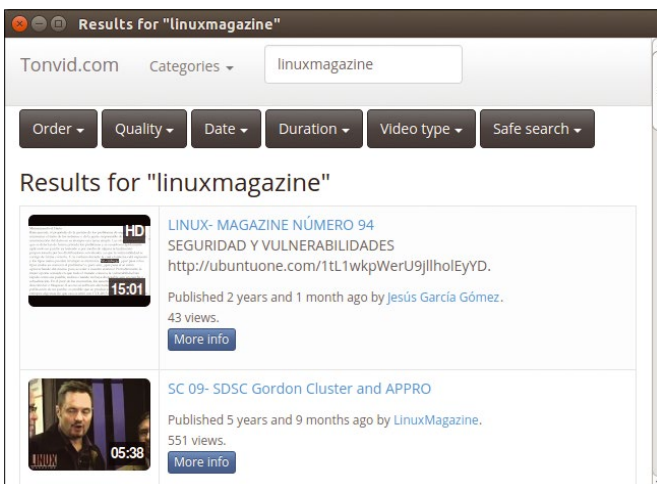
## SMTube 15.8.0

YouTube browser and client

Source: <http://www.smtube.org>

License: GPLv2

Alternatives: Minitube, Gtk YouTube Viewer



The SMTube video browser and player is part of the SMPlayer front end. It comes with a clear-cut user interface for finding YouTube movies, filtering the list of results, and playing movies.

When first launched, SMTube lists movies that are popular with other users; later, the tool will show you the category you most recently selected. The client supports 13 additional categories (e.g., animation, news, sports, science, and technology). To play movies, SMTube relies on SMPlayer by default. You can use the configuration dialog to set up an alternative player, such as MPlayer or VLC.

In contrast to earlier SMTube versions, you no longer download movies directly in the client, but via a detour. You need to specify a tool such as uGet or Youtube-dl in the configuration. The reason for these changes is that the video platform itself changes frequently.

★★★★☆ The YouTube client for the Linux desktop leaves nothing to be desired. Because YouTube regularly changes its interfaces, readers should choose the latest SMTube version to ensure full viewing pleasure. ■■■

## Jq 1.5

JSON stream editor

Source: <https://stedolan.github.io/jq>

License: MIT

Alternatives: Sed, Awk



Many programs rely on JSON when it comes to configuring, exchanging, and storing data. Users who want to process JSON files in the shell will be interested in taking a look at Jq. The tool, which is implemented in C, reads files from the standard import, and filters and outputs them in a user-defined format. It does not have any dependencies and is immediately ready for action after downloading and setting the x flag.

When calling the program, you need to specify a data stream. Command-line parameters influence the way that jq filters the input. For example, `-s` tells the tool not to process each object as it arrives, but first to read everything. In contrast, the `-R` parameter tells jq to expect the data in raw format and pass the data to the filter as simple strings, line by line. The tool interprets spaces as separators. The manual recommends quoting parameters with single quotes to avoid the shell interpreting Jq control elements.

The filtered results are sent to standard output, where again blanks act as separators. Various options format the output, color-highlight various messages, or use tabs as separators. The comprehensive documentation describes all the options in detail and uses examples to explain how users can create simple or advanced filters with regular expressions.

★★★★★ Jq reliably filters JSON files in the shell. Users are advised to consult the excellent manual before starting to use the tool. ■■■

### Pen 0.30.0

Load balancer

Source: <http://sourceforge.net/projects/penloadbalancer>

License: GPLv2

Alternatives: Balance, LVS

```
Terminal - vollbracht@LMLab-1504.b: ~
vollbracht@LMLab-1504.b:~$ pen -f -C 192.168.250.193:8081 -d 8080 192.168.250.5:80:1 192.168.250.62:80:1 128.127.66.142:80:1 37.221.192.98:80:1
2015-08-30 10:12:51: read_cfg(null)
2015-08-30 10:12:51: Before: conns = (nil), connections_max = 0, clients = (nil), clients_max = 0
2015-08-30 10:12:51: expand_conntable(500)
2015-08-30 10:12:51: After: conns = 0x9004c48, connections_max = 500, clients = 0xb7275008, clients_max = 2048
2015-08-30 10:12:51: pen 0.30.0 starting
2015-08-30 10:12:51: servers:
2015-08-30 10:12:51: 0 192.168.250.5:80:1:0:0:0
2015-08-30 10:12:51: 1 192.168.250.62:80:1:0:0:0
2015-08-30 10:12:51: 2 128.127.66.142:80:1:0:0:0
2015-08-30 10:12:51: 3 37.221.192.98:80:1:0:0:0

vollbracht@LMLab-1504.b:~$ penctl 192.168.250.193:8081 connection
Connection 0:
state = 0
downfd = -1, upfd = -1
client = 0, server = 1
pend = 0
vollbracht@LMLab-1504.b:~$ penctl 192.168.250.193:8081 servers
0 addr 192.168.250.5 port 80 conn 0 max 1 hard 0 weight 0 prio 0 sx 0 rx 0
1 addr 192.168.250.62 port 80 conn 0 max 1 hard 0 weight 0 prio 0 sx 419 rx 205
2 addr 128.127.66.142 port 80 conn 0 max 1 hard 0 weight 0 prio 0 sx 0 rx 0
3 addr 37.221.192.98 port 80 conn 0 max 1 hard 0 weight 0 prio 0 sx 0 rx 0
vollbracht@LMLab-1504.b:~$
```

On Linux, Windows, OS X, FreeBSD, HP-UX, and Solaris, the Pen load balancer fields requests for TCP and UDP protocols and forwards them to the hosts defined on launching the tool; it supports both IPv4 and IPv6.

To set up simple load distribution, users specify an address and port, on which Pen expects requests after the pen command. This is then followed by a list of target servers, for which you need to state the IP address, port, and maximum permissible number of connections. All of these values are separated by colons, for example, 128.127.66.142:80:1.

Pen manages the incoming requests in a table that can handle up to 2,048 clients by default. The -c option lets you modify this value. Connections that have not been used for a long time are automatically removed to create space for new ones. If a computer fails, Pen contacts the remaining systems and ensures that requests are directed to the servers with the lowest load at the time.

The archive also includes the penctl tool, which lets admins manage load balancers at run time. It reports the current load status, modifies the parameters of the target systems, or excludes clients from access in a targeted way.

★★★★☆ You should test Pen in your own environment. Because the man page only comes with a couple of application examples, this can be a question of trial and error. ■■■


### PingChecker 15.8.16

Ping for multiple computers

Source: <http://sourceforge.net/projects/pingchecker>

License: GPLv3

Alternatives: Nagios, Monitorix



To test whether a computer is reachable on the network, most users resort to ping. However, if you want to test multiple systems in one fell swoop, you can use PingChecker instead. The tool accepts computer names or IP addresses as input, checks whether they are reachable, and summarizes the results in a useful overview.

The current version of the Python tool contains a GUI version in addition to the command-line version. Users can type a list of computers to test manually in either version, or they can pass in a list as a text file. Manual entry in the graphical user interface is restricted to 10 addresses. PingChecker tests all the stated systems, transmitting four pings in each case. At the command line, users can follow the tool's actions live and in color – computers that are reachable are displayed in green, and those that are unreachable are shown in red.

The graphical interface shows the overall results in a new dialog. Additionally, PingChecker creates two text files, ping\_results.txt and ping\_results.csv, which makes it easy to process the results downstream in a database or spreadsheet.

★★★★☆ PingChecker gives you a quick and uncomplicated approach to checking the availability of multiple systems. Most appreciated is the fact that PingChecker automatically writes its results to two text files, which are easily processed downstream. ■■■



# 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)

Surprises in SQL: State-of-the-art options in the standard query language

# New Voice



In recent times, many developers have come to view SQL as inflexible and limited, but this classical database language has some tricks and special features that many users don't know about. *By Markus Winand*

**S**tate-of-the-art SQL can do much more than you might think. Despite its popular image as a fairly limited database tool, SQL is no longer restricted to the relational data model but can also handle nested objects and structured documents, features more commonly associated with later technologies like NoSQL. Of course, it all depends on what you call SQL. Not all vendors implement all the features of the various SQL standards that have appeared through the years. In this article, I take you on a tour of some of interesting tricks available through standards-based SQL.

## 92 and 99

The SQL:92 standard is the starting point for the complete, classical SQL

database system we think of today. SQL:92 was already the second major version of the SQL standard, and it achieved a certain level of completeness as an embodiment of the classic relational model. However, developers knew even back in 1992 that the relational model is not ideal for all data.

The third major version of the SQL standard in 1999 brought an end to the plain vanilla relational SQL. All signs pointed to object-oriented programming. The standard featured the concept of the object-relational database, but a couple of years too late as it turned out. Object-relational mappers (ORMs) had already

### LISTING 1: PHP SQL Pseudocode

```
$product_categories = array(SELECT category FROM product_categories);
foreach ($product_categories as $category) {
    SELECT *
    FROM products
    WHERE category = $category
    ORDER BY popularity
    LIMIT 3;
}
```

Lead Image © Olly, fotolia.com

### LISTING 2: Two Attempted Joins

```
SELECT *
FROM product_categories k
JOIN products p ON (p.category = k.category)
ORDER BY popularity
LIMIT 3

SELECT *
FROM products
WHERE category IN (SELECT category FROM product_categories)
ORDER BY popularity
LIMIT 3
```

begun to build a bridge between object-oriented programming and the relational data model.

SQL:1999 also introduced some other new features, such as loops. Even though SQL:1999 broke with many traditions, it still remained a declarative language. In other words, it is impossible by definition to tell a database how to execute a query. As long as the results are okay, the database has full freedom. This declarative nature made it difficult for the standard to address loops, because a loop defines the solution's approach. The trick SQL used in SQL:1999 was to define constructs that can only be executed as loops. Programmers can then use these constructs as loops, which should not really exist in a declarative language.

As an example, consider the problem of converting the PHP and SQL code

### LISTING 3: Invalid Subquery

```
SELECT *
FROM product_categories k
CROSS JOIN (SELECT *
FROM products p
WHERE p.category = k.category
ORDER BY popularity
LIMIT 3
) top_products
```

### LISTING 4: With LATERAL

```
SELECT *
FROM product_categories k
CROSS JOIN LATERAL (SELECT *
FROM products p
WHERE p.category = k.category
ORDER BY popularity
LIMIT 3
) top_products
```

from Listing 1 into pure SQL. The example starts by loading a list of categories into a PHP array and then runs a further query for each category; the query returns the three most popular products from each category.

Performance-conscious users could object at this point that placing database queries in loops of an imperative language – PHP in this case – could cause trouble. A join is typically a better option to make sure that the database returns all the required data at the same time. In this case, there is a problem: the example cannot be handled with a simple join. The examples from Listing 2 show that, no matter how you twist and turn, the result is always the three most popular products all told – and not the three most popular per category.

The problem is that LIMIT does not act in a category-specific way. To achieve the desired results, you need to run LIMIT in a subquery that is restricted to a specific category (Listing 3). However, this is no longer valid SQL; subqueries in the FROM clause cannot access data external to the subquery.

The WHERE clause that uses k.category to access a table external to the subquery is thus invalid. Or at least, that was the case in SQL:92. SQL:1999 supports this kind of access if the user precedes the subquery with the new LATERAL keyword (see Listing 4).

This query is equivalent to the PHP code in Listing 1, except that the database executes the loop itself and the latencies between the application and the database are thus avoided. Figure 1 shows the similarities between the Foreach loop and LATERAL.

Another advantage of using SQL for everything is that the user can process the results with SQL downstream. For example, you could sort the overall results differently using an ORDER BY clause or write the results directly to a caching table with INSERT INTO ... SELECT .... The latter case avoids transporting data from the database to the application and then back to the database.

The benefit of being able to process the results with SQL farther down the line applies to any SQL query, of course. As a general rule, users should not assume that SQL databases are just storage bins. Data processing is often easier with SQL than with other programming languages. The results are typically more correct and the performance better. This approach only fails when faced with massively parallel access, such as experienced by Google and Facebook. If you are that big, however, you will definitely have the means to build a proprietary solution, and until you get there, the flexibility that SQL offers is often the better approach.

SQL:1999 introduced a second construct that can also be used like a loop:

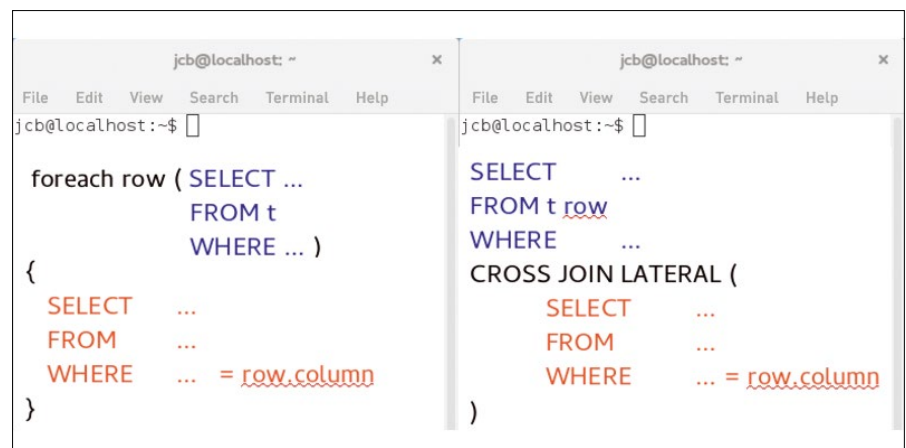


Figure 1: PHP pseudocode and SQL LATERAL loops compared.

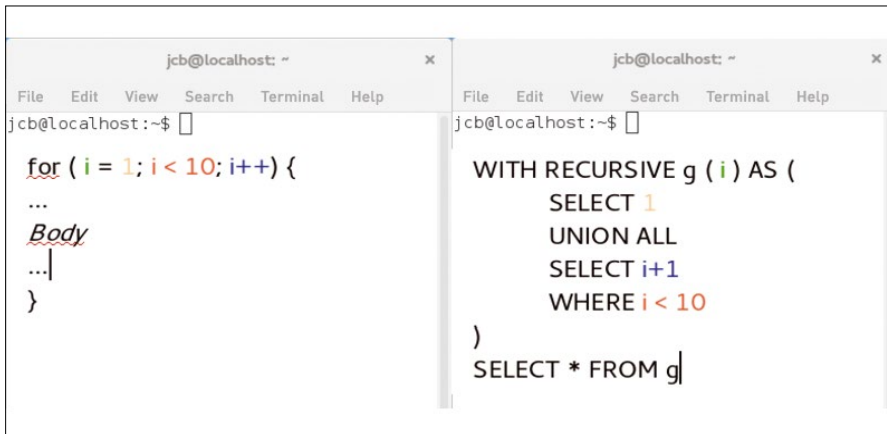


Figure 2: A WITH RECURSIVE loop in SQL.

WITH RECURSIVE. The details of how this works are fairly complex and well beyond the scope of this article, but Figure 2 shows the basics.

The WITH RECURSIVE variant has three benefits:

- It is better supported by today's databases than LATERAL – for example, it is also supported by SQLite (see Figure 3).
- Data can pass from one iteration to the next.
- It is possible to formulate a dynamic termination condition.

The disadvantage is that the loop body cannot be transferred one-to-one to SQL: It must be merged with the part after UNION ALL.

WITH RECURSIVE supports some important use cases. The query shown in Figure 2 is a row generator; it simply returns 10 numbered rows – very practical for generating test data. A more important use case is traversing graphs, such as finding the shortest connection between two persons on a social network.

### Rapid Steps – SQL:2003

After ditching relational-only thought patterns in SQL:1999, in only took four years for the next major revision of the SQL standard. The focus of SQL:2003 was on two points: XML and analytical functions.

XML support is interesting because SQL databases became document stores, as one would call them today. In SQL:2003, users can store XML both as text and as validated documents that can be processed using SQL and XQuery.

Although some databases support the XML extension today, XML was un-

able to assert itself in web development. Its competitor JSON was simply too attractive. All popular SQL databases have introduced features for handling JSON documents in recent years. These extensions are purely proprietary – each database offers a different feature set.

But the second SQL:2003 focus – analytic functions – has asserted itself. In particular, the window functions are supported by many databases today, and they vastly simplify data preparation.

A window function lets the database programmers use aggregate functions, such as SUM or COUNT, without GROUP BY. Of course, you still need to define the rows to use for the aggregate, but if you don't use GROUP BY, because rows are grouped, you can now use the new OVER clause directly after the aggregate function.

The following query clarifies the effect: it adds an additional column to a query (SELECT \*) that has the number of rows in the overall result when done:

```
SELECT *
, COUNT(*) OVER ()
FROM [...]
```

It makes no difference how you proceed: JOIN, WHERE, HAVING, GROUP BY, ORDER BY – everything is possible. The COUNT(\*) OVER() window does not have any side effects on the rest of the query. The empty bracket in the OVER clause means the COUNT function runs against all the rows of the result.

The window function essentially says: Count all rows! The result of this query is thus an additional column containing the number of rows in the result. To clarify things once again, the row count comes with each row and is thus returned multiple times; however, this does not mean it is determined multiple times!

Of course, there is no need to return the same results umpteen times with the window function. You can use the OVER clause to define the rows against which to run the function. The most important tool for defining the rows is PARTITION BY:

```
SELECT *
, COUNT(*) OVER (PARTITION BY category)
FROM [...]
```

You can read the preceding expression as: Count the rows with the same value in the category column!. In other words, PARTITION BY delimits the rows just like GROUP BY, but it does not group the rows like GROUP BY does; instead, it simply specifies the rows to which the window function is applied.

If you have multiple rows of the same category, you will receive the same result in each case of these rows. Remember that a window function does not have any side effects on the remaining results of the query. In particular, this means that multiple window functions can be used in a single query without them influencing one another. As an example, if you need both the number of

### LISTING 5: ORDER BY with OVER()

```
SELECT *
, SUM(value) OVER
(OORDER BY time ROWS BETWEEN
UNBOUNDED PRECEDING AND CURRENT
ROW)
, AVG(value) OVER
(OORDER BY time ROWS BETWEEN 3
PRECEDING AND 3 FOLLOWING)
FROM [...]
```

### LISTING 6: With Value Ranges

```
SELECT *
, AVG(value) OVER (ORDER BY date
RANGE BETWEEN INTERVAL '3' DAYS PRECEDING
AND INTERVAL '3' DAYS FOLLOWING)
FROM [...]
```

rows in the overall results, and the number of rows per category, you can use the relevant above examples together in a query.

The rows that are visible for a window function can be further delimited if the rows are sorted. That allows `OVER` clauses to mean something like, *all rows before or three rows before to rows after*. “Before” and “after” refer to a sort order, which is freely definable using `ORDER BY` in the `OVER` clause. This means you can, say, compute a subtotal (*Sum of all rows before*) or a moving average (*Average of three rows before to three rows after*) Listing 5); `value` and `time` are the column names in this example.

The formulation of the `BETWEEN` range follows similar rules to `BETWEEN` in the `WHERE` clause: The start must be somewhere before the finish, and the specified values themselves are part of the range. The `AVG` function in the example thus generates the average value of up to seven rows: the three rows before, the current row itself, and the next three rows. If you don’t have enough rows to fill the window, then it is smaller – this would be the case in the first row, for example, which cannot be preceded by three rows.

You can define the window size by row, but also by value range. For example, you can define windows that cover *three days before and three days after* – no matter how many rows these days cover (Listing 6) – just define the `BETWEEN` range with `RANGE` instead of `ROWS`.

In addition to the well-known aggregate functions, SQL:2003 also introduced ranking functions that strictly require an `OVER` clause with `ORDER BY`.

## Not Just Aggregate Functions

These ranking functions in particular are `ROW_NUMBER`, `RANK`, and `DENSE_RANK`. As

the name would suggest, `ROW_NUMBER` lets you enumerate rows. `RANK` and `DENSE_RANK` are less intuitive. Both return a rank as per the `ORDER BY` clause. `RANK` and `DENSE_RANK` differ from `ROW_NUMBER` in that *ex aequo* placements (ties) with `RANK` and `DENSE_RANK` take the same rank, as is typical in sports. In other words, you could have two first-place contenders.

`RANK` and `DENSE_RANK` differ in the question of how to rank the next contender. `RANK` omits a placement in this case – thus returning 3; `DENSE_RANK` uses placements without gaps – that is, 2 in this example.

SQL:2011 introduced more window functions for access to individual rows of a sorted data window. `LAG` and `LEAD` let the user access the previous and next row. `FIRST_VALUE`, `LAST_VALUE`, and `NTH_VALUE` provide access to the first, last, or *n*th row of the window. The example in Listing 7 illustrates this function by means of a competition.

The first column determines the placement using `RANK`. This step is followed by the name and the number of points achieved. The last two columns show the gap to the competitor just ahead and to the winner. The query uses `LAG` to access the score of the competitor in front and `FIRST_VALUE` to see how many points the winner has. The number of points achieved by the current candidate is deducted (-points) to give you the required gaps.

A word of warning on this example: this `OVER` clause uses `ORDER BY` without using `BETWEEN` to delimit the window. In this case, a default clause takes effect: `RANGE BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW`. This default is surprising for two reasons. First, the subsequent rows are excluded from the window. Second, this exclusion is based on `RANGE` not on `ROWS`. In other words, you still

have competitors with the same score in the window.

## And the Performance?

The main benefit of window functions is the variety of options you have with simple syntax. Window functions are also typically much faster than their purely relational equivalents. Despite the huge range of options that window functions offer, you can easily estimate the additional overhead they cause in a database. You just need to remember that most window functions are performed in two simple steps: preparation and application.

In the preparation phase, you need to make sure the data is sorted as per the `PARTITION BY` and `ORDER BY` clauses. In other words, the columns of the `PARTITION BY` clause – if they exist – are placed in front of the `ORDER BY` clause, and the data is sorted accordingly. You can omit this sorting if the data is already sorted in this order. The data could be in order if it was sorted previously – for example, for a different window function – or, as in the previous example, by a global `ORDER BY` clause. Sorting is also unnecessary if the data comes from an index with appropriate sorting.

The additional overhead of the preparatory phase is thus mainly equivalent to the overhead for an `ORDER BY` clause for each different `OVER` clause. Identical `OVER` clauses do not cause this overhead multiple times. The previous example requires no preparatory overhead, because the global `ORDER BY` clause has already sorted the data as required in advance.

After this preparation, the individual window functions are then executed. This overhead can be roughly compared with normal use of aggregate functions with `GROUP BY`, with the exception that the function is executed for each row in the worst case. Of course, each SQL vendor does its best to avoid this worst case.

## Time Travel with SQL:2011

The final trailblazing SQL feature I’ll examine in this article is temporal and bi-temporal tables. Although the term *temporal* itself does not give you many clues, this feature encompasses functions that everyone has already used.

### LISTING 7: LAG and FIRST

```
SELECT RANK() OVER(ORDER BY points DESC) "Placement"
, name
, points
, LAG(points) OVER(ORDER BY points DESC)
  -- points "gap to the previous competitor"
, FIRST_VALUE(points) OVER(ORDER BY points DESC)
  -- points "gap to the first placed competitor"
FROM [...]
ORDER BY points DESC
```

Generally speaking, you can use temporal and bi-temporal tables to make data changes traceable. Consider a customer database: how does the user handle the fact that a customer has changed their name? Is it sufficient to overwrite the name with a simple UPDATE? Or do you need to be able to see that the name has changed? Where traceability was required, users previously had to take care of this themselves. That is, mark the old row in the customer database as deprecated and create a new row with the new name.

SQL:2011 now offers time-versioned tables. An UPDATE of this kind of table would look something like the following:

```
UPDATE customers
FOR PORTION OF validity
FROM DATE'2015-01-01'
TO DATE'9999-12-31'
SET name = 'New name'
WHERE id = 123
```

The user defines in the UPDATE – immediately after the table name – the period of time for which the new name is valid. To allow this to happen, the table has to be prepared appropriately by creating two columns for the start and end times and grouping the data to create an application time period (validity in this example). SQL does all the rest. Whatever you need to achieve, the desired state is done – for example, an UPDATE of the last valid row and an INSERT for the new name.

SQL:2011 naturally offers everything else you need to handle these tables, including constraints that understand the time concept (WITHOUT OVERLAPS), but also queries with a specific time:

```
SELECT *
FROM customers
WHERE id = 123
AND validity CONTAINS DATE '2015-12-01'
```

This query returns the new name, because SQL periods always include the start time and exclude the end time.

Application time versioning is designed to reflect reality. For instance, the valid start for a new name is the day of the wedding – not the time when the new name is entered on the system. The application can choose validity arbitrarily – which explains the “application time” label.

Validity periods are different for system-versioned tables. Once a table is system versioned by means of appropriate ALTER TABLE instructions, the database handles versioning in a completely transparent way. All SELECT, UPDATE, ... instructions still continue to work unchanged; you don't need to modify the application. However, the application cannot choose the validity periods itself. Instead, the database always uses the current system time.

All INSERT, UPDATE, and DELETE instructions always take immediate effect. System versioning thus does not reflect reality but always references the time at

which the change was made in the database.

Both types of versioning – application time or system versioning – are optional and can be used independently of one another. It is also possible to apply both kinds of versioning to a single table at the same time; however, this approach creates bi-temporal tables. Application time versioning would use the wedding date as the date for the change, but system versioning uses the time at which the change was entered on the system, which could be weeks later. (Probably very few newly-weds inform their banks of a name change just one day after the ceremony.)

### Which Databases?

If you try out one or more of the functions described in this article, you might discover they don't work! Not all SQL databases are state-of-the-art. To be more precise: the SQL standard only mandates a couple of functions. Most of the functions in the standard – including all the functions described in this article – are optional. Additionally, many manufacturers still regard SQL:92 as the definitive statement on SQL. Just imagine this headline for comparison: “Leading NoSQL database now also supports Windows 3.1!”

These modern SQL functions are useful, even though just a few databases (including some open source databases) support them. If you work with SQL on a regular basis, it is worth your while to learn these state-of-the-art SQL features so you can select a database that supports the functionality when you need it. Figure 3 shows some SQL database systems that support the functionality described in this article. ■■■

### AUTHOR

As a freelance trainer, **Markus Winand** specializes in teaching developers efficient techniques for SQL. He wrote *SQL Performance Explained* and is currently working on his next book, which is being published on an on-going basis at <http://modern-sql.com>.

 Find us on Facebook

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

	DB2 LUW	MySQL	Oracle	PostgreSQL	SQL Server	SQLite
LATERAL	✓	✗	✓	✓	✓	✗
WITH RECURSIVE	✓	✗	✓	✓	✓	✓
Window Function	✓ <sub>0</sub>	✗	✓	✓	✓ <sub>2</sub>	✗
Temporal Tables	✓	✗	✓ <sub>1</sub>	✗	✗ <sub>3</sub>	✗

<sup>0</sup>No NTH\_VALUE (in version 10.5)  
<sup>1</sup>Similar function available as “Flashback”  
<sup>2</sup>No NTH\_VALUE (in version 2014)  
<sup>3</sup>Announce for SQL server 2016

Figure 3: Availability of selected SQL functions.

Shop the Shop

shop.linuxnewmedia.com

Want to subscribe?


Need training?

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


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

shop.linuxnewmedia.com

DIGITAL & PRINT SUBSCRIPTIONS



SPECIAL EDITIONS



TRAINING



- LPIC-1 LPI 101 - CompTIA Linux+ LX0-101
- LPIC-1 LPI 102 - CompTIA Linux+ LX0-102
- LPIC-1 - CompTIA Linux+ 101 + 102



Publishing systems for DocBook XML

# From the Source

Good software requires detailed documentation. The DocBook Authoring and Publishing Suite, Publican, and Yelp all let you generate a variety of documents – from man pages to eBooks. *By Mathias Huber*

**D**ocumentation for standard software tools can be written using the simplest resources. For example, the ReadMe files included in most source code packages can be written with the use of any text editor. In recent years, developers have also used the Markdown format [1], which offers bulleted lists, listings, and web links and is formatted neatly by code hosting provider GitHub.

However, users of powerful software suites, particularly commercial software, need and expect genuine manuals. The DocBook [2] format, which exists in SGML and XML flavors, satisfies this need, prevailing over more lightweight competitors, such as reStructuredText [3].

DocBook provides everything a developer needs for a professional manual –

## AUTHOR

**Mathias Huber** is an avid Linux user and works in the documentation department of thin client specialist IGEL Technology.

from tables of content and tables of figures to cross references and indexes. DocBook has also established itself as a kind of lingua franca for technical documentation. Free software can handle this language just as well as proprietary applications, such as Adobe FrameMaker.

If you use DocBook XML, you do not need to worry about whether it will still be possible to read and maintain your documents in a few years' time. Of course, this assumes that you are not fazed by working with XML source code, although this work can be made easier by choosing a suitable editor.

## On Board

The DocBook schemas and a collection of XSL style sheets for converting into HTML and other formats are parts of the package scope for virtually all Linux distributions. Using them, it is possible to implement single-source publishing workflows (i.e., create man pages and info, HTML, PDF, and EPUB documents from the same source code).

A practical publishing system, however, needs to offer more – especially if multiple users need to collaborate. It needs to specify where the numerous XML, image, and other files from the documentation project end up; automate the process of building output files; and make sure the cross references actually work. Additionally, a good system needs to create the framework for a new documentation project in a single step.

## Test Site

Those who don't fancy programming all that themselves can draw on the work of others who have published the source code for their software. The SUSE Linux documentation team has developed the DocBook Authoring and Publishing Suite (DAPS) [4] for in-house projects and openSUSE. Red Hat, which runs projects like Fedora or oVirt, meets similar requirements with Publican [5]. The third candidate is the group of programs by the Gnome project centered on the Yelp [6] help browser.



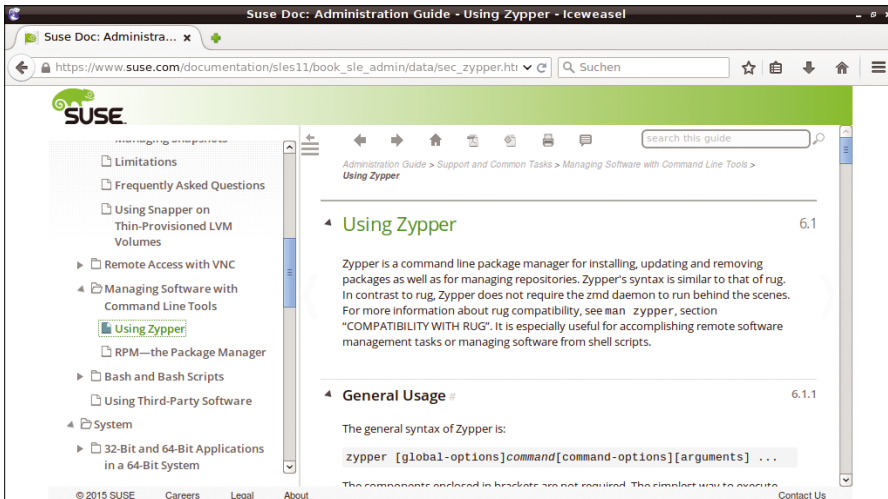


Figure 1: The official SUSE documentation in WebHelp format is created using the DAPs.

### DAPS

The openSUSE and SUSE Linux Enterprise manuals are created using DAPS (Figure 1), so you can simply install the suite on these distributions from the package repository. The installation guide [7] lists the required packages for other flavors of Linux. Thanks to Git, the code for the suite ends up on the hard disk and can be run directly in the cloned directory with just a few tweaks.

DAPS follows the usual DocBook recipe and converts the DocBook XML into other formats such as (X)HTML, HTML Help, man page, EPUB, or WebHelp using `xsltproc` and the DocBook stylesheets.

WebHelp is included in the latest DocBook stylesheets, offering an expandable navigation bar and presenting SLES documentation clearly on the web. Another ability is to convert to XSL Formatting Objects (XSL-FO), which serve as raw material for producing PDFs using the Apache Formatting Objects Processor (FOP) [8]. Bash scripts and Makefiles glue the components together.

The central command is `daps`; it has zillions of options and parameters, from checks, through transformations, to packaging options. The `daps-init` command creates a new document; Figure 2 shows the results.

The DC file contains the key settings for the publishing project; the DocBook XML file is in the XML directory. It can also be used as a main file for unitized XML and, for example, integrate chapters from separate files. The `images` directory shows the variety of formats that DAPS processes: the software also auto-

matically generates output versions for Dia, Xfig and SVG when publishing.

While typing, authors can make use of suitable extensions, as required, for the Emacs, Vim, and jEdit editors. DAPs also supports authors with various options that validate the DocBook XML, check cross references and web links, and find required or orphaned files. The spell-checker is not fazed by the XML tags thanks to Aspell.

The suite also fulfills advanced documentation requirements by creating several versions of a document using conditional text and profiling to replace product names or exclude sections and features that only apply to a specific product version. It is also possible to add comments and draft watermarks for proofreaders and translators.

As well as the output formats mentioned, DAPS also creates tarballs, which contain either the XML source code or HTML Help. On request, it also packages the documentation for RPMs or the Gnome and KDE help systems.

### Lost in Translation

One of the documentation team's important tasks is to collaborate with

translators. This involves localizing XML source code, replacing images, and managing updates. The SUSE team uses the `lclodrop` DAPS option for exporting files and then imports the completed translation using `unpack-lclodrop`. The SUSE developers have also created an open source tool in the form of DocManager [9], which cooperates with DAPs and manages the workflow of documentation projects with the help of metadata.

A comprehensive user manual [10] and a quick start guide – made with DAPS, of course – provide users with a detailed explanation on how to use the complex suite and leave virtually no questions unanswered.

### Publican

Publican [5] is available in the Red Hat, Fedora, Debian, and Ubuntu package archives. If you have openSUSE, note that it can only be installed on version 12.1 or later. Just as with DAPS, developers can also check out the publishing system from Git [11] to use the most recent version.

Red Hat's Publican system also uses DocBook XSL stylesheets and Apache FOP to generate various output formats. Perl serves as an integration language. Important information, such as the DocBook version and the document class, can be found in the `~/.publican.cfg` and `<project_name>/publican.cfg` files.

The `publican` command supports a large number of actions and has a very long man page: `publican create gener-`

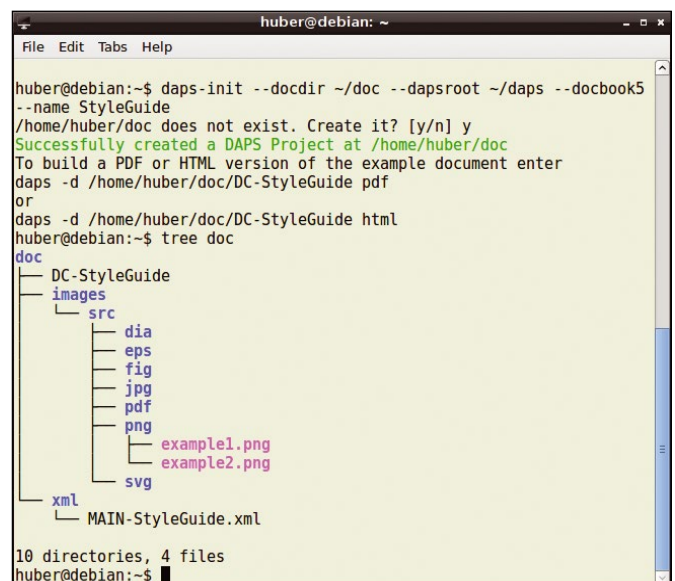


Figure 2: The `daps-init` command creates the directory structure for a new documentation project.

ates a new documentation project, `publican build` produces output from HTML to PDF and help plugins for Eclipse, and `publican clean` deletes temporary files.

The build action also supports the `test` format, which doesn't build anything but checks whether a particular language version of a document can be built without errors:

```
publican build --formats=test
--langs=de-DE
```

The Publican developers came up with a viable workflow for internationalization and localization and have added it to the Publican command. A document project

has a separate directory for each language version; the directory is named per the ISO language code, such as `pt-BR` for Brazilian Portuguese. The

```
publican update_po --langs=hi-IN,ru-RU
```

command tells Publican to extract the character strings to be translated for Hindi and Russian from the DocBook XML and package them in appropriate directories.

The PO format from GNU gettext [12] is used in this case, and widespread tool support is available. The format is also suitable for translation portals from open source projects, for which volunteers

choose work packages to translate themselves. Publican currently supports more than 40 languages; standard content, such as formatting conventions, are already localized.

Publican provides functions for identifying unused XML or image files, for displaying used or prohibited DocBook tags, or for packaging as aids for a documentation project workflow. Packaging is not just for RPMs, but also for importing into CMS Drupal or as a website with a search engine-optimized sitemap.

## Brands

As a publishing system for RHEL, Fedora, Gimp, and the like, Publican uses brands to adapt documents to the corporate design. Debian has its own Publican brand in its package archive, and a generic-looking common package is in the public domain.

Designers can also create their own brands. The

```
publican create_brand --name=LM
--lang=en-US
```

command creates a style directory for *Linux Pro Magazine* (Figure 3).

Its files contain the word `SETUP` in the places that need customization. They are easy to find using the `grep` command (Figure 4):

```
grep -rHn --color=auto 'SETUP' *
```

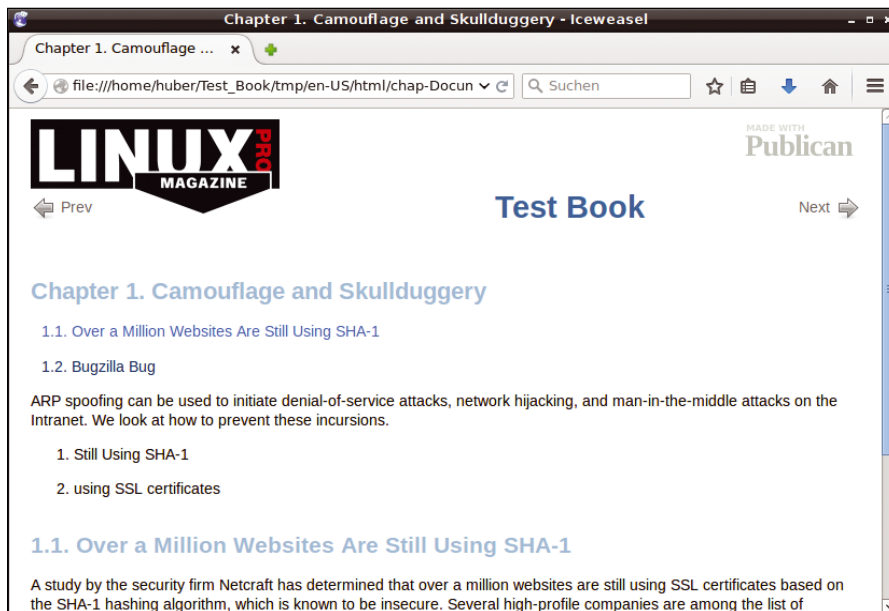


Figure 3: A simple brand can be designed for *Linux Magazine* within minutes with Publican and applied to a document.

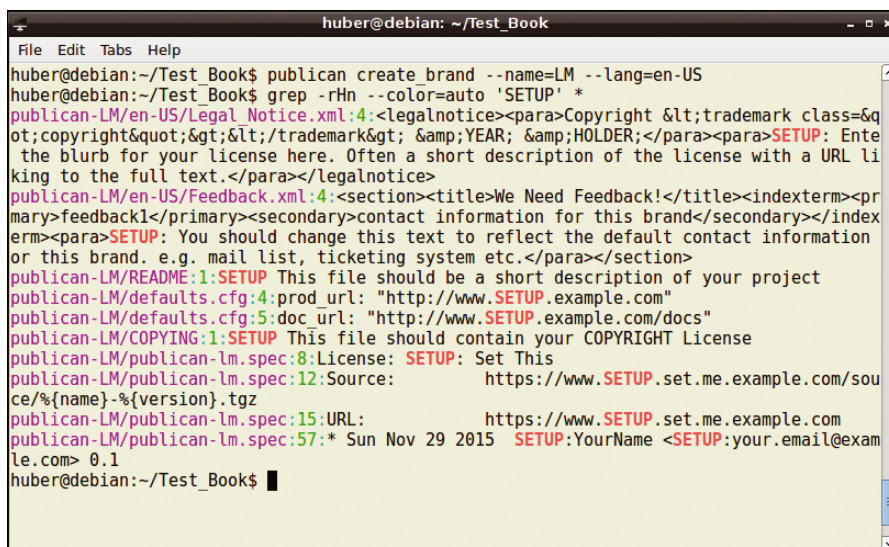


Figure 4: Publican can generate a unique look for documents. Developers need to customize passages containing the `SETUP` keyword.

Furthermore, standard image files, such as the logo, need to be replaced and the `overrides.css` stylesheets need to be edited. Publican can package brands as RPMs or as tarballs. If you don't want to create a package, you can use the `--brand_dir` option to specify which brand directory you want to use. Red Hat's team rounds off the Publican tool with a detailed manual [13].

## Yelp

The help browser for the Gnome desktop environment goes by the name of Yelp [6]. It can display DocBook documents, among other things, without the need for conversion up front (Figure 5). It comes with the `yelp-tools` package, which according to the Yelp homepage contains small programs for creating, editing, managing, and publishing documents. The DocBook and Mallard [14]

input formats are supported. Mallard is similar to DocBook but uses a simpler repertoire of XML elements and uses a separate file for each text section.

If you need information about the `yelp-build` tool, for example, you will notice that the documentation available on the Yelp website is pretty sparse. Although the tool supports DocBook versions 4 and 5, it seems to be a second class citizen:

The attractive EPUB target format is the reserve of Mallard, as is the option to use a cache for converting (Figure 6).

Reading material on `yelp-new` for new documents is sparse as well, but `yelp-check` is worthy of note: It validates DocBook XML, checks cross references and external links, ensures the integrity of image files, and finds extraneous files.

Ironically, what is most missed in Yelp tools is a decent manual to provide users with an overall picture of how to create a documentation project and how the individual tools fit into the workflow, from initial stages to output formats. The Gnome wiki [15] is not much help and cites outdated or even missing pages.

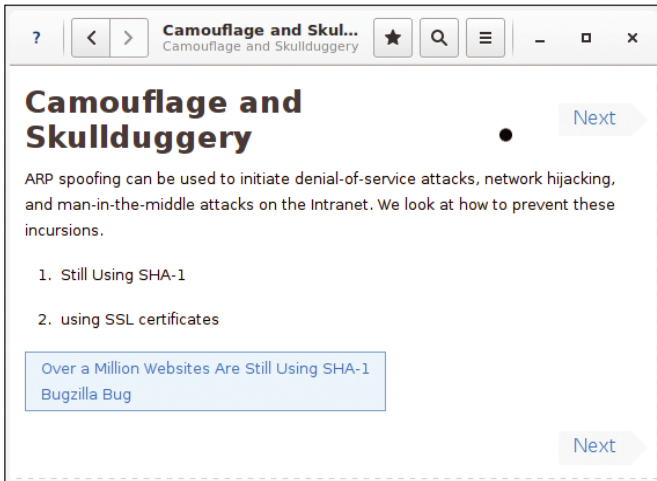


Figure 5: Gnome's help browser Yelp displays DocBook documents, but the documentation is less complete.

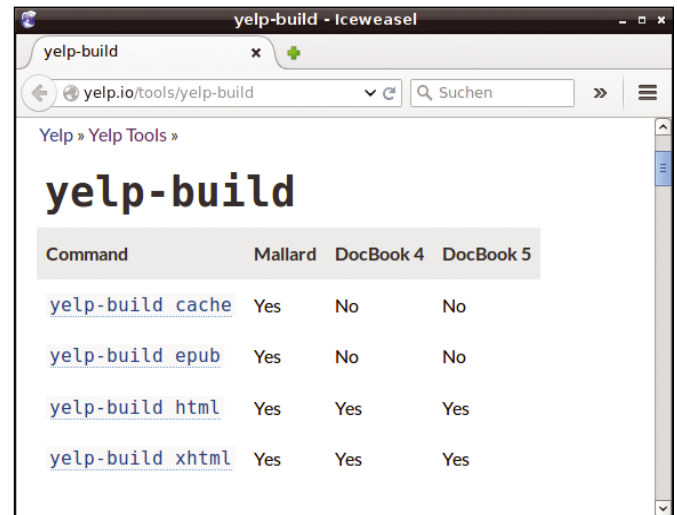


Figure 6: The Yelp documentation is sparse and shows a strong preference for the Mallard format.

# 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

### Summary

DAPS and Publican are full-fledged systems that DocBook professionals – and newcomers, too – can use to manage their documentation projects from start to finish. They both go above and beyond individual manuals because they support sets (i.e., collections of books) with cross references that are guaranteed to work.

Both tools reign in their particular domains: DAPS supports WebHelp output and many graphics formats; the manual also reveals how developers can migrate existing DocBook projects to the suite. Publican, on the other hand, makes it easy to create unique brands. The SUSE group is quite active in terms of layout and design, too, but its focus is on customizing DocBook XSL style sheets. Team member Thomas Schraitle devotes a whole book to the topic as an open source project [16].

Another advantage of the Red Hat tool is that Publican has integrated localization, which will appeal to anyone wanting to implement widespread language support. Publican also packages RPMs for Fedora and RHEL and can handle Red Hat’s Brew build system.

Gnome’s Yelp draws the short straw, mainly because of the poor documentation for the documentation tool. It is also the reason for the gaps in the general overview in Table 1. However, the lack

of DocBook capability only matters if you insist on using this format; after all, Gnome help can be written pretty easily with Mallard [17].

### Outlook

Topic-oriented publishing manages sections with similar content as standalone components, compiles them to create larger units, and rehashes the results. Mallard shows a tendency toward topic orientation, but the buzzword among

documentation authors is Darwin Information Typing Architecture (DITA) [18]. Like DocBook, the XML architecture specification is maintained by an Organization for the Advancement of Structured Information Standards (OASIS) working group.

A collection of useful tools is already available under the Apache license in the form of the DITA Open Toolkit [19]. The Yelp developers are experimenting with DITA, as well. ■■■

### INFO

- [1] Markdown: <http://daringfireball.net/projects/markdown/syntax>
- [2] DocBook Project: <http://docbook.sourceforge.net>
- [3] reStructuredText: <http://docutils.sourceforge.net/rst.html>
- [4] DAPS: <http://opensuse.github.io/daps/>
- [5] Publican: <https://fedorahosted.org/publican/>
- [6] Yelp: <http://yelp.io>
- [7] DAPS installation: <https://github.com/openSUSE/daps/blob/master/INSTALL.adoc>
- [8] Apache FOP: <http://xmlgraphics.apache.org/fop/>
- [9] DocManager: <https://github.com/openSUSE/docmanager>
- [10] DAPS documentation: <https://opensuse.github.io/daps/doc/index.html>
- [11] Publican Git repository: <https://git.fedorahosted.org/git/publican.git>
- [12] gettext: <https://www.gnu.org/software/gettext/>
- [13] Publican documentation: [https://jfearn.fedorapeople.org/en-US/Publican/4.0/html/Users\\_Guide/](https://jfearn.fedorapeople.org/en-US/Publican/4.0/html/Users_Guide/)
- [14] Mallard: <http://projectmallard.org>
- [15] Gnome Yelp wiki: <https://wiki.gnome.org/Apps/Yelp/>
- [16] DoCookBook: <http://doccookbook.sourceforge.net>
- [17] “Mallard” by Mario Blättermann, *Linux Pro Magazine*, issue 132, November 2011, pg. 54
- [18] Dita community: <http://dita.xml.org>
- [19] DITA Open Toolkit: <http://www.dita-ot.org>

**TABLE 1: Feature Overview**

Name	DAPS	Publican	Yelp
License	GPLv2 or GPLv3	GPLv2+ and Artistic License 1.0	GPLv2+
Version	2.1	4.3	3.16
XSL formatter	DocBook XSL	DocBook XSL	Yelp XSL
PDF formatter	Apache FOP	Apache FOP, Wkhtmltopdf	Not specified
Glue code	Bash, Make	Perl	C
Own themes	Yes	Yes	Not specified
Source text input formats	DocBook XML 4 and 5	DocBook XML 4.5 and 5	Mallard, Docbook XML 4 and 5
XML spellcheck	Yes	Rudimentary	Not specified
Link check	Yes	Yes	Yes
Draft watermarks	Yes	Yes	Not specified
Annotations	Yes	Yes	Yes
Custom layouts	Yes	No	Not specified
Own themes	No	Yes	Not specified
Output formats	HTML, HTML single, WebHelp, TXT, MAN, PDF, PDF grayscale, Mobi, EPUB	HTML, HTML single, HTML desktop, TXT, MAN, PDF, EPUB, Eclipse	HTML, XHTML
Localization	Own tool DocManager	Installed, uses GNU gettext	Via xml2po
Extras	RPM, Emacs macros, XML Bigfile, partial builds	RPM, Brew integration, website, Drupal export, accessibility (Section 508)	DocBook display
Conditional text	Yes	Yes	Not specified

## The sys admin's daily grind: Gatling

# Apache Under Fire

Western aficionados and sys admin Charly are about to set up a Gatling in a field that is normally home to Apache. Read on to discover why blogger Fefe is to blame. *By Charly Kühnast*

I admit to watching (too) many westerns as a child; in other words, I'm fully aware of what a Gatling gun is: an automatic weapon with multiple rotating barrels. In old movies, soldiers used its infernal din to send the horses of attacking Apaches galloping mad. And, this always worked until an Apache warrior crept up from behind and torched the ammunition chest.

As a rapid fire weapon, the Gatling lends its name to the web server I will be looking at today. The server's binary file is a lightweight compared with its name giver, weighing in at just 100KB. Its RAM requirements are minimal, too, because Gatling [1] does not fork. On an old PC, Gatling is just as fast as on my Raspberry Pi. The daemon is even included with the plain vanilla Raspbian distribution.

The server is designed for delivering static web pages as quickly as possible. To allow this to happen, its programmer,

Felix "Fefe" von Leitner, whose widely read blog [2] naturally also runs on Gatling, gave the server just as many features as the number of shells you could slot into a revolver (IPv6, TLS/SSL, simply equipped virtual hosts, login via the `.htaccess` mechanism).

If you want this tool, you need to build it from the source code, but before doing so, remember to install the IO-API library [3], which was also implemented by Fefe. Launched without any parameters, Gatling serves up the data from the current directory, but you can home in on a target like this:

```
gatling -u nobody -c /var/www
```

After Gatling has shot its way out of port 80, it drops its root privileges and continues running with the "nobody" account – which is also the name of a character played by actor Terence Hill 42 years ago [4]. The chroot parameter

`-c /var/www` helps the sheriff bundle off the data into the county jail. And, `-R` number fends off mass attacks by telling Gatling to restrict the number of HTTP requests per minute and client.

## This Is the End of the Line, Django!

The log format is exotic (Figure 1), but it does help you understand the timing. An HTTP connection times out after a nerdy 23 seconds (last line). I tend to change this timeout value to 10 or fewer seconds using the `-T seconds` paragraph, to release the sockets as quickly as possible. The reason for this change is that I tend to run Gatling on the front line, where access attempts are like a heavy barrage. The web server responds with static HTML, for example, when a server farm enters a maintenance window ("Please try again later") or in the case of blogs that are churned out by static site generators.

Gatling even offers a simple FTP server (which you can disable with the `-F` switch). And, if you allow general access to the FTP directory, users can upload files. Conversely, you can only download files that you are allowed to read.

Backed up by a Gatling, "Nobody" can relax and coolly ride off into the sunset as the closing credits roll by. ■■■

```

charly: gatling - Konsole
File Edit View Bookmarks Settings Help
root@funghi:~# gatling -u nobody -c /var/www/default/
starting_up 0 :: 80
start_ftp 0 :: 21
start_smb 0 :: 445
accept 8 127.0.0.1 41666 1 http
accept 9 127.0.0.1 41667 2 http
accept 10 127.0.0.1 41668 3 http
HEAD/304 8 / 0 Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/45.0.2454.85 Safari/537.36 [no_referrer] localhost
request done 8 459 223
close/read0 9
close/read0 10
timeout 8
close/timeout 8
  
```

Figure 1: Gatling's log format is anything but standard, but it's pretty accurate about timing.

## 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.

## INFO

- [1] Gatling: <http://www.fefe.de/gatling/>
- [2] Fefe's blog: <http://blog.fefe.de> (in German)
- [3] Libowfat: <http://www.fefe.de/libowfat/>
- [4] *My Name is Nobody*: [https://en.wikipedia.org/wiki/My\\_Name\\_Is\\_Nobody](https://en.wikipedia.org/wiki/My_Name_Is_Nobody)

Klaus Knopper answers your Linux questions

# Ask Klaus!

By Klaus Knopper

**?** Hi Klaus, my Linux PC is “off the road.” Is it possible to offer follow-up disks after installing your Knoppix OS, for people that do not have Internet access capable of downloading files off the OS server?

Regards,  
Cy O’Hara




Figure 1: The flash-knoppix installer.

**💡** Like most distros, there are independent providers of DVDs or USB flash disks for Knoppix, too. Most of their addresses are listed on the “Knoppix Vendors” website [1]. Because the main feature of Knoppix is NOT having to get installed in the first place, but running Live from DVD or USB, new versions are not meant to be used as an update to a hard disk installation. However, the *flash-knoppix* installer allows you to update flash disks already containing a Knoppix installation (Figure 1). This process lets you overwrite files on the FAT32-formatted system partition but keep your personal data and settings on the second partition intact.

**?** Hi Klaus, I noticed that you replied to my email about compiling and linking 64-bit code in issue 166 of *Linux Magazine* last year [2]. Thank you so much! I also noticed that you installed *gcc-multilib* in Knoppix now. Thanks again! Would you consider including *g++-multilib* in Knoppix in the future as well?

If I may be so bold, I have another question for you. Knoppix is the best distro I have found by far for USB flash drives. If I install another distro, such as

Ubuntu, directly to a USB flash drive like I would to an HDD, it can pause and be unresponsive often. How come USB flash drives are more restrictive than a normal SSD drive? What steps does Knoppix take to get around the USB flash drive limitations? I have tried moving, for example, `/tmp` and `/var/log` to `tmpfs` but that’s not enough. Does Knoppix move `/home` directly to the ramdisk and start up and sync it on shutdown?

Kind Regards, Chad

**💡** Unfortunately, multiarch can be somewhat heavy concerning dependencies, so I’d rather like to keep run-time and development files for multiple architectures at a minimal level.

About USB speed: Reducing write cycles and finding a well-balanced scheme for parallel writes and reads is a challenge for hardware and software and depends greatly on the USB flash drive hardware design. SSDs are more optimized for asynchronous writes than most USB flash pens, so no latencies should occur when writing a large amount of data to disk. However, some flash internal controllers block read access during writes, so you might experience noticeable delays or “freezes” during the drive’s internal data reorganization, which can be quite annoying.



### 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)

You'll have to figure out which brand supports quicker parallel read/writes. Knoppix delays writes somewhat in order to write larger chunks of data at once, which is more efficient and enhances the lifetime of most flash pens. However, on slow drives, this can cause longer freezes during a write cycle, so these optimizations are not always efficient. Avoiding frequent writes by putting /tmp, /var/tmp, and other locations for often-written volatile data on a ramdisk can help as a work-around for slow writes.

Also, the filesystem RAM cache can be helpful. I have experimented with mov-

ing /home to ram and back to disk, but it has a huge disadvantage: Your RAM could get filled up with huge amounts of temporary data and thus crash the system. And, if you don't shut down the computer properly or are impatient, data may not get written back in time. Thus, I prefer to use buffered filesystem writes on an AuFS overlay with a second partition, so that reading from the read-only system overlay and writing on the read-write flash partition overlay can go in parallel. Apparently, this works pretty well in most cases.

The fact that the main system data is stored in compressed form (files KNOPPIX and KNOPPIX\*) is also helpful for performance, because it speeds up the physical read speed by the compression factor – usually between one and five times the uncompressed read speed. Decompress-

ing data in memory is much quicker than waiting for read response of magnetic, optical, or flash physical media.

Check to see what's mounted where in Knoppix with the df command after booting from a USB flash disk (Figure 2).

The /KNOPPIX mountpoint(s) are constructed from read-only compressed data, whereas /KNOPPIX-DATA is the read/write partition on the flash disk installation. Optimally, the controller managing NAND flash should handle parallelized read and write operations to different segments of the flash memory, but most often with cheap drives, it's implementation is busy waiting for any pending operation to complete in a sequential order, which makes data throughput much slower than necessary. ■ ■ ■

## INFO

- [1] Knoppix vendors: <http://knopper.net/knoppix-vendors/>
- [2] "Ask Klaus!" by Klaus Knopper. *Linux Magazine*, issue 166, September 2014, pg. 50: <http://www.linux-magazine.com/Issues/2014/166>

```
knoppix@Microknoppix:~$ df
```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
rootfs	998556	48	998508	1%	/
/dev/sda1	6322100	4269676	2052424	68%	/mnt-system
tmpfs	2097152	0	2097152	0%	/ramdisk read-only
/dev/loop	10068394	10068394	0	100%	/KNOPPIX
/dev/loop1	360086	360086	0	100%	/KNOPPIX1
/dev/sda2	9392860	35472	9357388	1%	/KNOPPIX-DATA
unionfs	9392860	35472	9357388	1%	/usr read-write
tmpfs	20480	2800	17680	14%	/run
tmpfs	10240	4	10236	1%	/UNIONFS/var/lock
tmpfs	102400	32	102368	1%	/UNIONFS/var/log
tmpfs	2097152	8	2097144	1%	/tmp
udev	20480	0	20480	0%	/dev
tmpfs	2097152	0	2097152	0%	/dev/shm
tmpfs	200184	4	200180	1%	/run/user/1000

Figure 2: Seeing what is mounted where.

# DON'T MISS A SINGLE ISSUE!


The first print magazine created specifically for Ubuntu users! Ease into Ubuntu with the helpful Discovery Guide, or advance your skills with in-depth technical articles, HOW-TOs, reviews, tutorials, and much, much more.

**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



Professional music production with Ardour

# MODERN CLASSIC

Version 4.4 of the free Ardour digital audio workstation impresses with optimum stability and well-conceived, functional extensions. *By Hartmut Noack*

The Ardour project introduced version 4 of its music production suite in April 2015. As is often the case with version jumps, the software experienced several inconsistencies and errors, as well as many new capabilities. Version 4.4 [1], which was released in mid-October 2015, has been tidied up; it features new developments (Figure 1) and is even leaner and more stable than its predecessor.

You can now complete many tasks more easily than before without having to forego tried and tested complex methods. For example, Ardour 4.4 is immediately ready for use on a system without a JACK audio server – it just uses ALSA as an interface. Of course, the performance of this combination is greatly inferior to a setup with JACK.

If you then install the JACK Audio Connection Kit and fire it up before starting Ardour, you'll get not only the full performance you're used to with Ardour 3, you'll get even more. (See also

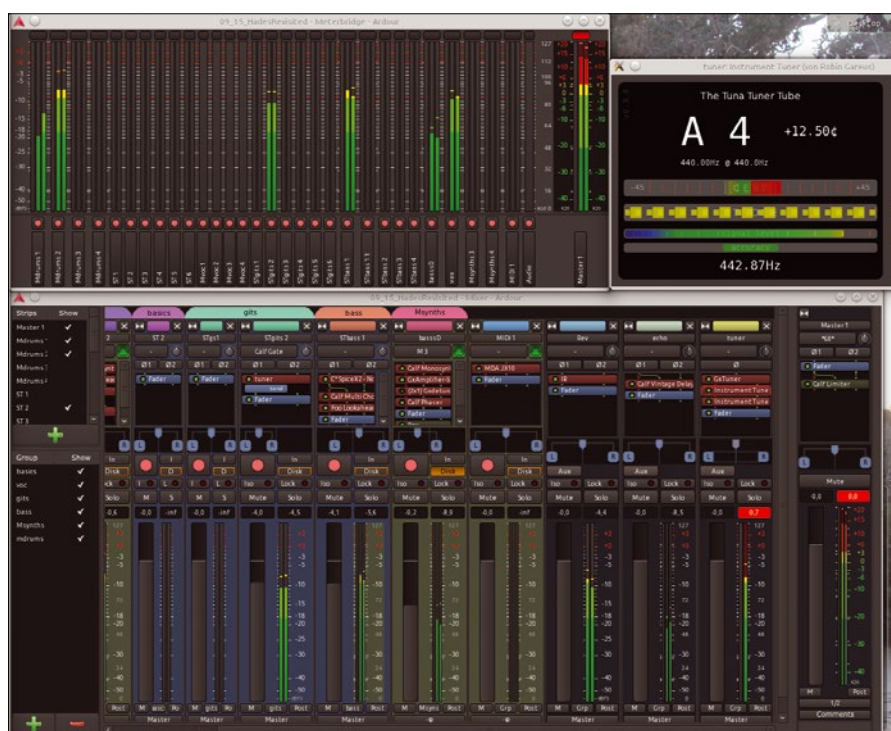


Figure 1: Ardour 4.4 comes with the clear-cut Meterbridge (top left), a fully equipped mixer, and the instrument tuner plugin created by Robin Gareus.

Lead Image © Karel Miragaya, 123RF.com



the “Ardour Anywhere – Even Without JACK” box.)

Ardour’s user interface has been almost completely rewritten, and multiple optimizations result in lower memory requirements and quicker response for almost all functions. The 80 percent lower memory requirements announced at the release of Ardour 4 prove to be accurate in practice. Any popular PC from the past five years should be happy with the software’s minimal requirements. See “Installing Ardour 4.4” to get started.

## Less Is More

After launching Ardour 4.4 with an empty project, free -m displays an additional memory consumption of around 200MB because Ardour preloads several functions in the background. Consumption only increases by another 50MB when you’re setting up a new MIDI

track with three large plugins. The system load increases by about 50 percent if you start the required modules Calf Monosynth and Guitarix independently from Ardour. This behavior by the application indicates that Ardour’s performance requirements are strongly dependent on what exactly you want to do in the application. A single MIDI track together with its plugins can be used normally on a low-powered system as long as you don’t have a “fat” desktop” and a load of system services working in the background.

However, if you’re aiming for a typical music production with 30 or more tracks and dozens of plugins, you’ll need a correspondingly powerful computer; in fact, it might well need to be a superfast machine. In the test, Ardour 4 ran reliably and smoothly with 40 tracks and 120 plugins on a computer with Intel i5-2500 CPU (3.3GHz) and

## ARDOUR ANYWHERE – EVEN WITHOUT JACK

Ardour was developed from the outset as a reference application for the JACK audio server. French developer Stephane Letz ported JACK to Apple’s operating system to make it possible to run Ardour on Mac OS X. Soon enough, people wanted an Ardour port for Microsoft systems, and a JACK for Windows was created. However, these ported JACK servers didn’t prove so popular in third-party operating system environments and didn’t provide the same quality as JACK for Linux in terms of performance and stability.

Ardour 4 ended this era because it also works directly with the sound system of

the underlying operating system. This also applies to Linux: For the first time, Ardour users can now select and configure ALSA as a sound system. Ardour’s internal capabilities don’t suffer, and the complex signal flow in the mixer works as usual.

As long as the plugins you use are also available on all platforms, it is possible to exchange Ardour projects arbitrarily with Ardour variants from band colleagues who use other platforms. However, JACK-specific functions, such as the use of external JACK software as a recording source or synchronizing with the drive from another JACK program, will not work with ALSA.

## INSTALLING ARDOUR 4.4

Ardour is and will remain free software; you can clone the source code anonymously and for free from the Git version management system by entering:

```
$ git clone git://git.ardour.org/ardour/ardour.git 4.0
```

In principal, it is possible for reasonably experienced users to do this themselves using waf [2], but it’s also pretty time-consuming work. The git statements can also cause problems on occasion, because Ardour has seen some fairly major development.

Therefore, you might want to download a prepared installation package from Ardour.org after making a donation of at

least \$1 using PayPal. If you register on the site and pay \$50, you get all the upgrades for the current and next generation at no extra cost. The same applies to subscribers who pay \$10 each month. However, you need a regular credit card registered with PayPal; users of prepaid cards have reported that they couldn’t get the subscription.

If you either don’t want to use PayPal or don’t like the idea of doing it yourself using the source code and its uncertainties, the only option left is to wait until the latest version appears in the repositories of the distribution you use. Experience shows that this only takes a few weeks on multimedia distros like KXStudio.

## OPTIMIZING ARDOUR PROJECTS

Ardour leaves it to you to decide what system resources a project consumes. Each new track takes as much power as required by the functions in the track, above and beyond the basic need of around 200MB of RAM after startup. Ardour is less focused on economy and more on smooth, comfortable operation. The cache for large pulse code modulation (PCM) data in the audio tracks results in reduced disk access. For example, even big jumps in very long recordings happen surprisingly fast with relatively few tracks. As you increase the number of tracks, you will quickly reach the capacity of your RAM.

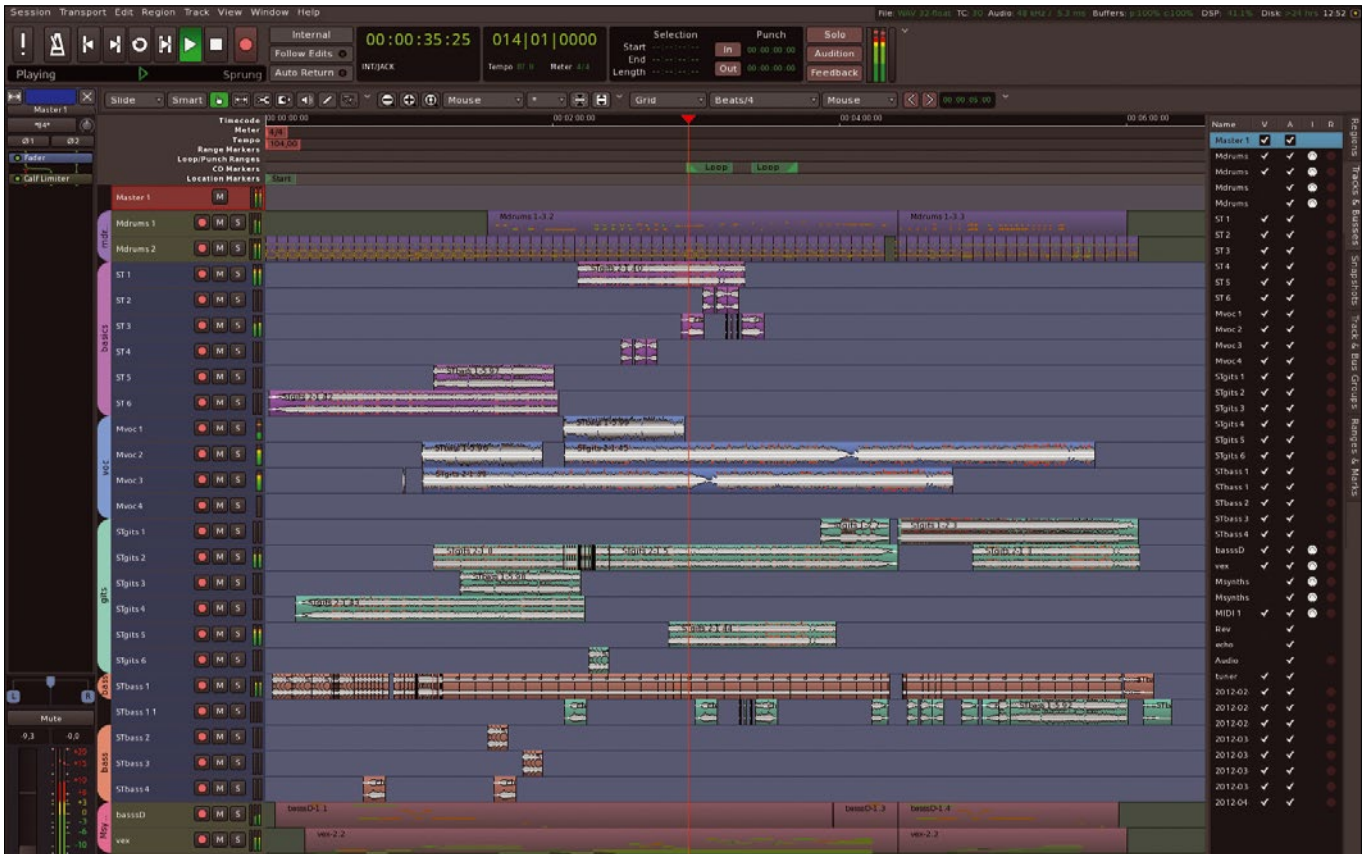
Various strategies help to utilize available system resources optimally. First, you can specifically adapt the settings for the cache (Figure 3). Large buffers on a few tracks allow big jumps without a noticeable delay, even in long recordings. However, if you load a shorter project with many identical buffer settings, Ardour could use up all the system memory, because it tries to preload a set period of time on the PCM data near the playback cursor. This involves significantly more data with shorter pieces and many tracks than with a long piece with just a few tracks.

Add-ons also devour system capacity. Because Ardour computes all operations in real time, the plugins in all tracks also need the corresponding CPU and RAM resources at the same time. For example, a really economical three-band equalizer, if it is active in every track by default, can mutate into a hungry monster given 40 tracks or more.

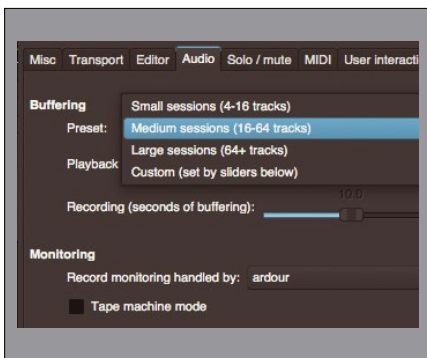
Like all modern professional digital audio workstations (DAWs), Ardour also provides different options for implementing a more efficient usage strategy and still not have to do without any functions. Regions on tracks can be calculated individually or as sections drawn with the range tool (*Bounce*). For this, Ardour applies all plugins to the data in the track and exchanges raw data for the result. If you then turn off all the plugins, you will get exactly the same sound from the track.

Ardour does this automatically if you explicitly request the whole track to be frozen. Ardour also provides the *Freeze* and *Unfreeze* options in the context menu of each track. *Thaw* restores the track to the state before freezing, activates all plugins again and disperses the audio data general view generated when freezing back into the original regions.

Additionally, you can calculate individual regions or sections in the track using the *Bounce (with processing)* function (Figure 4), that is, to apply a single effect to certain parts of a track and then disable it. Unlike frozen tracks, the track as a whole can still be operated on as normal afterward.



**Figure 2:** More than 40 audio/MIDI tracks, a lot of plugins, and still only 50 percent CPU load with 2.7ms input latency: Ardour 4 requires significantly less power than its predecessor.



**Figure 3:** You can adjust audio material buffering for default project types with a single click or define a custom setting with sliders.



**Figure 4:** The two Bounce functions in the context menu for regions work like the *New Region* function for regions. The calculated region replaces the corresponding source region.

8GB of RAM with a latency of only 5.3ms (Figure 2). However, I was using a very respectable sound card in the form of an MAudio 1014 PCI. The same project, but on a similar computer with a simple USB 2 sound card, only ran without interruptions with JACK settings for 10ms latency.

Ardour supports virtually all common techniques for music production, including synchronizing videos. Sound can be recorded on audio tracks, which may have up to 12 discrete channels. Using these channels, you can implement various surround-sound techniques in addition to stereo, such as 7.1 sound for game soundtracks. Ardour 4 can also control systems for wave-field synthesis, which is another step into the future of surround sound.

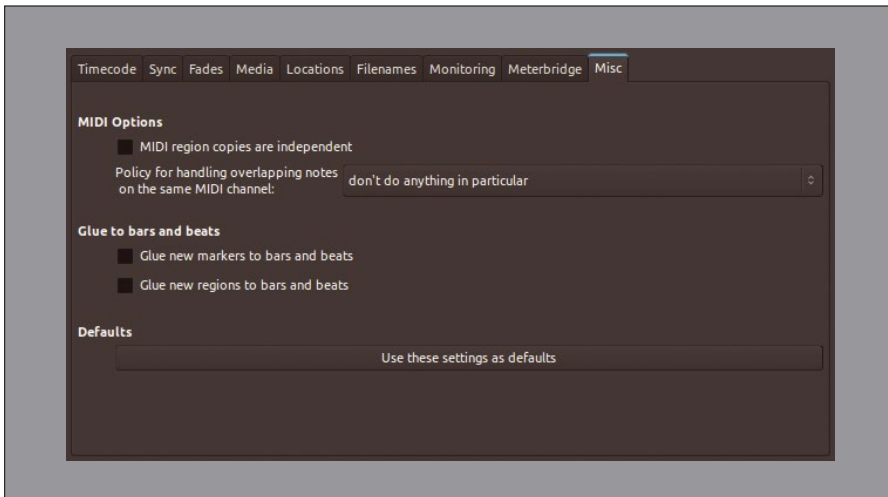
MIDI tracks record notes from connected instruments and events from controllers, which can also be drawn and edited using pen tools. Ardour does not provide its own editor for this; as in Traktion, the editing tools work directly in the MIDI region in the track of the Editor window. Robin Gareus has completely redesigned the way MIDI

signals are passed through to Ardour 4 tracks so that the plugins that output MIDI signals themselves can also be connected.

The new version can handle large projects really well. In our lab, I was able to edit a three-hour recording without any noticeable delays and to extend it by an additional eight stereo tracks with other music recordings. Only a seven-minute project with 52 channels led to a slight decline in the response speed and jump-started the fan. See the "Optimizing Ardour Projects" box for ways to improve the efficiency of your work.

### Tending Old Flames

At first glance, only the advanced wizard for starting or creating projects really stands out in Ardour 4. The interface also looks pretty much the same as in Ardour 3, although it is much easier to read. Whereas the background, graphics, and controls were still monochrome in Ardour 3, subtle gradients now grace Ardour 4. The display of small elements and fonts is sharper and better harmonized.



**Figure 5:** You can now control the basic settings for MIDI and manage the involvement of regions for specific projects.

The switch to the Cairo 2D library, which also significantly lowered the system load for the basic application, is responsible for the visual improvements. Graphics cards with 2D hardware acceleration make the Ardour 4 graphical user interface quicker than the Gtk elements used by older versions. Thanks to Cairo, from Ardour 4.2 onward, it is possible to scale labels and control elements continuously like buttons and regulators.

If you take a closer look, you'll see new tools in familiar places. Some tools

supplement what was already there and others replace old functions. However, Ardour 4 doesn't focus on reducing (sometimes lamented) functional diversity. No functions from Ardour 3 were dropped in Ardour 4. The greatly developed tool for project settings opens up customization options (Figure 5); you can use it to adjust many global settings specifically for each project.

The editor provides a few new modes for cutting and transporting regions. When cutting in *Ripple* mode, all the re-


gions to the right slide to the new end of the regions to the left. If you move a region, all those to the right of it move to compensate. *Play loop range* is now a regular playback mode.

Ardour has always provided comprehensive support for controller hardware, such as the popular drive controllers and mixer from Mackie and Behringer. All mixer and plugin functions can also be controlled using signals from MIDI keyboards and controllers. However, classical MIDI signals only have a resolution of 7 bits, which only allows 127 gradations for the velocity of played notes and controller settings. Ardour 4 now supports advanced MIDI devices that have 14-bit signals to provide 16,000 gradations. Ardour has supported the even more accurate and even more flexible Open Sound Control (OSC) for many years. OSC is a protocol for transmitting high-res signals with high sampling rates via network. The receiving software interprets these signals as parameter changes.

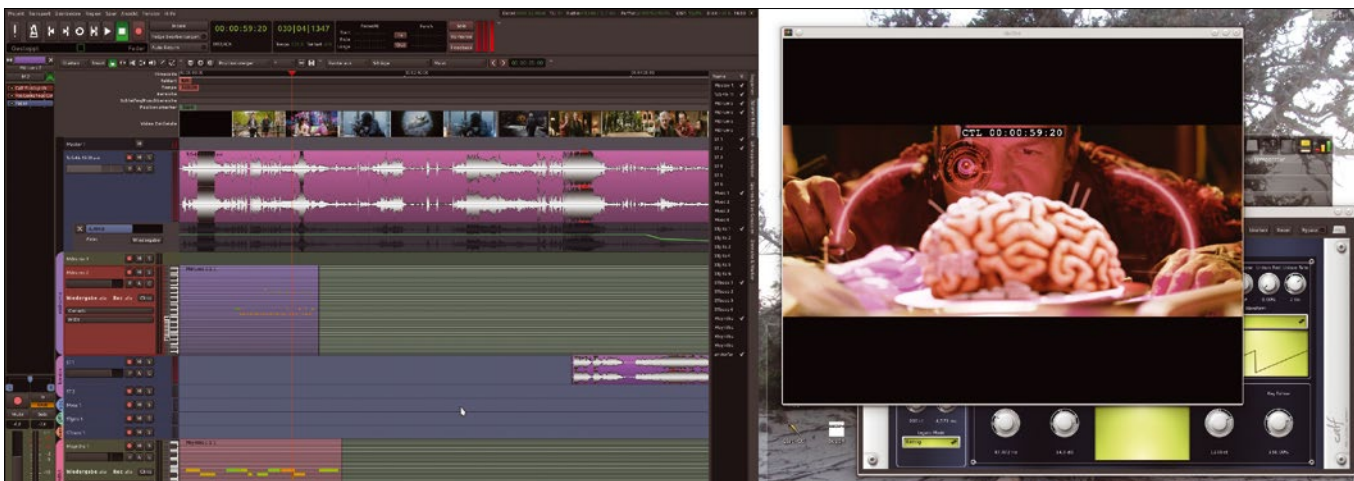
Various minor improvements are evident across the whole interface, such as a new scissors tool; splitting at the operating point by pressing *S* also is still functional. Mixer inputs now have sliders for the input level, and the *Save as* function now creates a complete copy of

the project in a new folder. The old snapshot function still exists here, too.

Since Ardour 4.2, it has been possible to roll the regions during playback while the playback pointer re-



**Figure 6:** The tool bar in Ardour 4.4 offers new modes and a pair of scissors, with various new functions awaiting in the menus, as well.



**Figure 7:** Video post-processing works best with two monitors: Ardour's video window behaves with abandon in the foreground, while the plugin interface remains in the background.

mains fixed in the middle. All tools are familiar with a new Smart mode, in which they adapt to the respective usage

scenario. The *Smart* switch for the default pointer tool can be used like the range tool if the mouse pointer is at the

top of the track. The MIDI editing tool on the far right can also draw notes using Smart mode (Figure 6).

### ARDOUR AS AN AUDIO EDITOR

Ardour operates very differently as a real-time DAW than typical sound editing programs such as Audacity, but it still works excellently as an editor. The Bounce and Render functions in particular are the key to this. Ardour uses them in different ways to write audio material on tracks into new audio files.

The arrangement of loops with recurring musical patterns is a typical application for a sound editor. For example, musicians like to cut individual sounds from an improvised bass line and arrange them so that an optimal loop is created for the respective piece. Using the *Split/Separate* function, you can extract an appropriate section of a region that you can then edit easily and precisely as a new region. The note grid built into the background of the editor helps place the notes in meter positions that match the speed of the song.

You can set the *Snap Mode* in the top middle drop-down list to *Grid* or the less rigorous *Magnetic*, or you can disable it completely by choosing *No Grid*. If disabled, regions and sections can be cut and moved with pixel-level accuracy. This also allows you to build de-

liberately small variations in a loop. The resolution of the grid can be set to all common values between a bar and a 128th note (*Beats/128*) for the mode to the right of the drop-down list.

The result often includes dozens of sometimes subtly positioned individual regions that can complicate the arrangement. The Range tool can help: You can use it to drag a contiguous block across all the regions in the loop. In doing so you are best to set it to a rough value like a bar or half-note. Right-clicking the new area then unveils the various Range functions. If you've only cut and arranged, using the pre-mixer type is a good idea. It just recomposes the original material and otherwise leaves the PCM data unchanged.

The alternative post-mixer (Figure 8) also renders all settings and plugins for the mixer channel into the new file created in this way. This is useful, for example, if you also want to add adjustments or a special effect to the cut material. If you write this effect in the new region, you can then delete or deactivate it in the track to save performance.



Figure 8: A new area with signal processing also renders plugins in the new region, such as the sophisticated Calf Multi Chorus and the very useful SpiceX2. You can then remove them both from the project.

### Prime Time

Working with videos is pretty easy (Figure 7). No weaknesses or unusual system loads were apparent in the lab. The thumbnails for the video timeline let you navigate intuitively in the film. On request, Ardour can automatically import an existing original soundtrack from the movie, and I experienced no inconsistencies between sound and vision in testing.

Ardour's ability to configure and balance tracks and masters with more than two stereo channels makes it possible to mix modern multichannel film sound. Ardour might not be able to export proprietary formats such as DTS or Dolby Digital directly, but the multichannel WAV files that it outputs can be processed in a corresponding encoder. A free tool in the form of Aften [4] is also available for Dolby Digital, and the commercial encoder software can also process Ardour's outputs.

The MIDI tracks appear mature and solid in Ardour 4. Everything works perfectly with normal operations, such as recordings from MIDI drum kits and keyboards, but I also didn't have any issues

with more complex actions like editing recording during playback. Ardour 4 also provides a few new functions for MIDI composers.

If you to select a few notes and right-click, the context menu will *Quantize* or *Legatize* these notes only, for example. In the Transform dialog, you can make various idiosyncratic transformations as well as a steady crescendo. The tempo for the whole piece and for tempo change markers can be changed with the *Tap tempo* button.

### Conclusions

The code name for Ardour 4.4 is "Discreet Music," which is probably a reference to the record by Frank Zappa. It's also possible that Ardour's chief developer Paul Davis and his team wanted to make clear the kind of music production for which they optimized the software: classical recordings with live musicians at a professional level.

The suite absolutely meets this requirement. Particularly in terms of reliability, Ardour 4.4 has no weaknesses – in around 30 hours of testing there

wasn't a single serious problem: not even crashes caused by random editing during operation. (However, please see "Better Safe than Sorry.")

The beautiful, new interface and the well-conceived, functional extensions are another good reason to download Ardour 4.4 from the Ardour website as soon as possible. The donation of at least \$1 is certainly a good investment. ■■■

### INFO

- [1] Paul Davis on Ardour 4.4: <https://community.ardour.org/node/10682>
- [2] Building Ardour yourself: [http://ardour.org/building\\_linux.html](http://ardour.org/building_linux.html)
- [3] Ardour error messages: <http://tracker.ardour.org>
- [4] Free Dolby digital encoder Aften: <http://aften.sourceforge.net>

### AUTHOR

**Hartmut Noack** (<http://lapoc.de>) works as a lecturer, writer and musician in Celle and Hanover. He has always thought that free software and homemade music fit together brilliantly. When he is not sitting in front of his Linux audio workstation, you will find him hanging about on web servers. You can download some CC-licensed musical samples of his work with free music software from Hartmut's own server.

### BETTER SAFE THAN SORRY

Ardour lets you control very complex operations with just a few milliseconds of delay. However, it is a pretty ambitious thing to attempt and doesn't always work perfectly. The program therefore creates backups in the background so that no more than a few minutes work is lost if the program crashes. You also need to take a few common precautions.

The principle of "Save often, save early" applies as always. So, always press Ctrl+S after making a big change to your project and before performing potentially tricky operations, or go to the menu and select *Session | Save*. In hundreds of hours of work using Ardour, I haven't had a single issue with reloading a saved session.

The delicate procedures in Ardour 4.4 also include the time-stretch function and inserting plugins or tracks during operation. Both are much safer when operation is stopped. The call for plugin interfaces and their settings was also sometimes problematic in older versions of Ardour; I had no problems in the test with the current version. Rule of thumb: If a plugin works smoothly after installation, the same will apply to when in operation.

In some situations, it is essential that absolutely nothing goes wrong, such as in

recording sessions with musicians or when using Ardour as a mixer at concerts. It is worth taking the following precautions in such cases:

- Set up all the tracks and plugins before using them and test them thoroughly. Only add more elements when not in operation during breaks at concerts.
- Don't make any changes to plugins while a recording is running.
- Where possible, always use native mixer elements for adjusting effect signals in the mix (faders from broadcasting channels).
- Control effects that are difficult or impossible to implement with the mixer elements by using the automation functions or a MIDI/OSC remote control.

Ardour's mixer provides extensive options for the last two points. You can, for example, place special effects on a broadcasting bus and mix them using the fader without having to touch the interfaces of the plugins involved. From Ardour 4 onward, it is possible to automate discrete switches in the mixer so that you can also switch plugins on and off without using the mouse. Additionally, you can control each action in the mixer using MIDI and OSC signals (Figure 9).



**Figure 9:** When holding down the Ctrl key, clicking any controller in the mixer with the middle mouse button connects the next incoming controller signal to the current controller. Individual controllers from plugins such as Calf Monosynth can be placed directly in the channel.

# SUBSCRIBE NOW



## 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)

Producing stickers and QR codes on a label printer with Perl

# Label Maker

Labels bring order to the mess of wires hiding behind the Perlmeister's home routers and organize the treasures hoarded in a multitude of boxes. With just some tweaking, the Dymo LabelWriter even prints on Linux. *By Mike Schilli*



**F**or decades, I have organized my network cables with permanent labels (Figure 1) printed on a portable device by Brother. However, it bugs me that every time I print a label the device wastes raw material (Figure 2), which I have to buy in the form of fairly expensive cartridges.

Brother's engineers deliberately seem to have built the machine to use twice as much label ribbon as I actually need, boosting the ribbon cartridge turnover as a side effect! If there is a hell for committing such wanton waste, I expect that the product managers responsible for this feature will be there some time soon. Apart from this, typing strings like `192.168.0.1` takes ages on the unorthodox

keyboard; using a desktop computer would be far quicker.

### Faster than Manual

Recently, I found a label printer on eBay that I was able to connect with my Ubuntu desktop via a USB port. The LabelWriter 450 Turbo by Dymo (Figure 3) cost me around \$40 secondhand, and I got it working in no time. What you need is the printer's CUPS system driver, which is available as source code [1].

After installing a couple of additional packages, such as `libcups2-dev` and `libcupsimage2-dev`, with `apt-get`, I was able to complete the build using `./configure; make` without any trouble. Then, `sudo make install` installed the CUPS files. It wasn't until later that I discov-



**Figure 1:** Fast diagnostics thanks to systematic organization: the configured IP address and WiFi SSID on the router.

### 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>.



ered that my Ubuntu distribution has an easy-to-install package and that

```
sudo apt-get install printer-driver-dymo
```


would do the job in one fell swoop.

After installing the CUPS drivers, Ubuntu detected the new label printer without much ado. A call to

```
lpstat -p -d
```



**Figure 2:** The portable Brother label printer wastes expensive raw material for labels.



**Figure 3:** The LabelWriter 450 Turbo also prints on Linux.

shows the available printers and their states (e.g., *printer LabelWriter-450-Turbo is idle*). If you are interested in the label sizes the driver supports, you can type

```
lpoptions -p LabelWriter-450-Turbo -l
```

to query them. Below *Printers* in the system settings, you should see a dialog like that in Figure 4 with a *Printer Options* submenu, in which you can con-

figure the dimensions of the labels you are using.

I bought a roll of labels numbered “30330” featuring 500 return address labels. I wasn’t really worried about what professional shippers use these labels for; it was the handy 19x51mm format that interested me.

The labels are shot vertically out of the thermo-printer, so I needed to adjust the lettering to match; in other words, rotated through 90 degrees in landscape. Now, how does the computer actually send the text to be printed to the device?

The CUPS printing system [2] uses the `lpr` command for this:

```
lpr -P LabelWriter-450-Turbo -o PageSize=w54h144.1 label.pdf
```

As you can see, the option `-P LabelWriter-450-Turbo` tells CUPS to select the label printer if some other device is configured as the default printer. The label dimensions are set using the `PageSize` option, which expects the length and width – not in millimeters or inches, but in points.

Some brief research on Google showed that you need to multiply millimeters by around 1.8 to convert to points (Figure 5). To match the 19x51mm labels, I thus needed rectangles in a format of 54x144.1 points; the parameter I used – `PageSize= w54h144.1` – defined precisely that size.

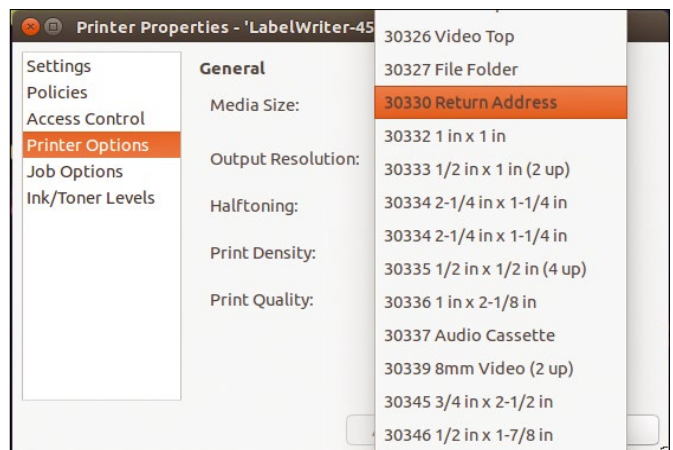
The final argument the `lpr` command expects is the name of the PDF file with the text to be printed.

### Formatting as PDF

Listing 1 [3] generates the required print layout with the help of the CPAN PDF::Create module. The script expects a string at the command line, for example,

```
label-format "Huzzah!"
```

and dumps it into the middle of a PDF document in landscape format measuring 144.1x54 points. Note that the text needs to be quoted if it contains blanks. Some shell dialects also expect you to quote the exclamation marks, which the shell would otherwise interpret as a call to its history function.



**Figure 4:** Ubuntu detects the label printer and serves up the label size configuration options.

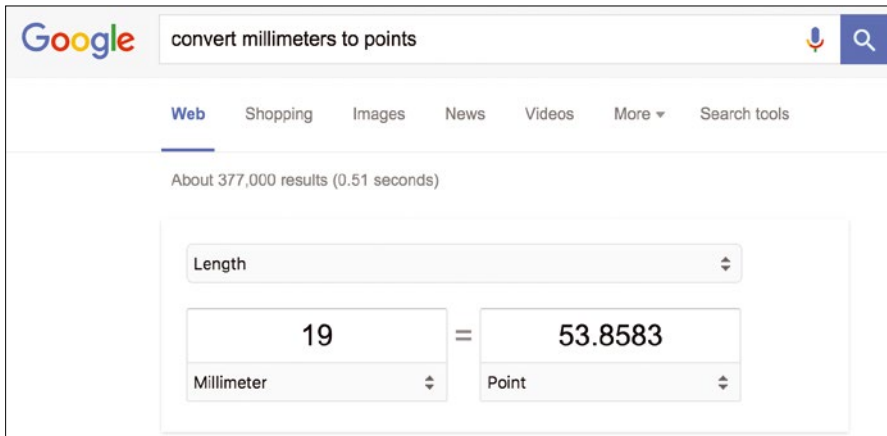


Figure 5: According to the conversion table I used, 19 millimeters are about 54 points.

Looking at the code, the reasoning behind the expression

```
20 / length( $string ) * 11;
```

in line 16 of Listing 1 is that my experiments revealed that a string of eleven 20-point characters would precisely fill out the label lengthwise. Longer strings necessitate a linear font size reduction to fit on the label. This simple formula works amazingly well even though the Helvetica font I chose does not use fixed-width spacing but assigns a proportional amount of space to each letter in the document.

The `stringc()` method called in line 28 prints the string passed in as its fifth argument in the middle of the PDF document. The font size is the second

parameter, followed by the  $x$  and  $y$  coordinates of the center of the label. The  $x$  value runs from left to right in the document, and the  $y$  value from the bottom to the top edge. For some strange reason, the text string was always too low in the document in my experiments; this prompted me to introduce the `$adjust` variable in line 17 to move the string upward by one seventh of the font size.

Thanks to these simple tricks, the label printer positions both short and slightly longer lines of text with up to 25 characters perfectly at the center of the document; the call to `cmd-print` fires up the printer, and the label comes flying out.

The process for handling non-standard characters, like umlauts, is some-

thing special. Because the string passed in at the command line is already UTF-8 encoded, but Perl interprets it as ASCII by default, you would normally see the printer output cryptic characters rather than what you wanted if left to its own devices. The `_utf8_on` function from the `Encode` module tells the interpreter that the string you passed in is already UTF-8 encoded and that there is thus no need to recode.

The CPAN module for creating PDFs can do more than just position strings on a white background; it can also position images. Of course, a thermal printer that can only print black without grayscale limits your artistic options somewhat, but what about printing sticky labels with barcode?

You can scan QR codes with a smartphone, opening up the option of assigning inventory numbers to devices printed as QR codes on labels that you could then scan and, if needed, store next to an asset number in a database or a filing system like Evernote, say, for a scanned owner's manual.

Listing 2 processes the text passed in at the command line,

```
dymo-qr "Linux Pro Magazine!"
```

and uses the CPAN `Image::PNG::QRCode` module to generate the QR code shown in Figure 6. The free QR Reader iPhone scanner app detected the QR code on the

### LISTING 1: label-format

```
01 #!/usr/local/bin/perl -w
02 use strict;
03 use PDF::Create;
04 use Encode qw( _utf8_on );
05
06 my( $string ) = @ARGV;
07 _utf8_on( $string );
08
09 die "usage: $0 string" if !defined $string;
10
11 my $pdf = PDF::Create->new(
12     filename => 'label.pdf' );
13
14 my $width      = 144;
15 my $height     = 54;
16 my $font_size = 20 / length( $string ) * 11;
17 my $adjust     = $font_size/7;
18
19 my $page = $pdf->new_page(
20     'MediaBox' => [ 0, 0, $width, $height ] );
21
22 my $f1 = $pdf->font(
23     Subtype => 'Type1',
24     Encoding => 'WinAnsiEncoding',
25     BaseFont => 'Helvetica'
26 );
27
28 $page->stringc( $f1,
29     $font_size,
30     $width/2 + 10,
31     ($height - $font_size)/2 + $adjust,
32     $string );
33
34 $pdf->close;
35
36 system( "acroread label.pdf" );
```

label still in the printer, even without me tearing off the label. Figure 7 shows the plain text version in the app.

Because the CPAN module I used to generate the QR code can only output PNG-formatted images, Listing 2 uses a

clone of the CPAN `Image::Magick` module as an easy option for converting image formats. Installing the module involves some manual attention because it needs a number of developer packages from various libraries. Fortunately, a

kind person from Ubuntu went through the trouble of generating a `perlmagick` package that provides the feature set with the `Graphics::Magick` module.

The module converts the temporary PNG file managed using the `File::Temp` module into a JPG file, which is then picked up by `PDF::Create` and dumped into a PDF document. The `label.pdf` file that this process creates is then forwarded to the label printer with the previously shown `lpr` command.

The only disadvantage with my label printer is that you can't buy rugged plastic labels for it. It only does black thermal printing on white paper or transparent film, both of which are self-adhesive. That said, however, the more expensive printer model called `LabelWriter Duo`, seems to have more options. I might just go bargain hunting again in the near future. ■■■



Figure 6: The scanning tool looks for the QR code in the photo, ...

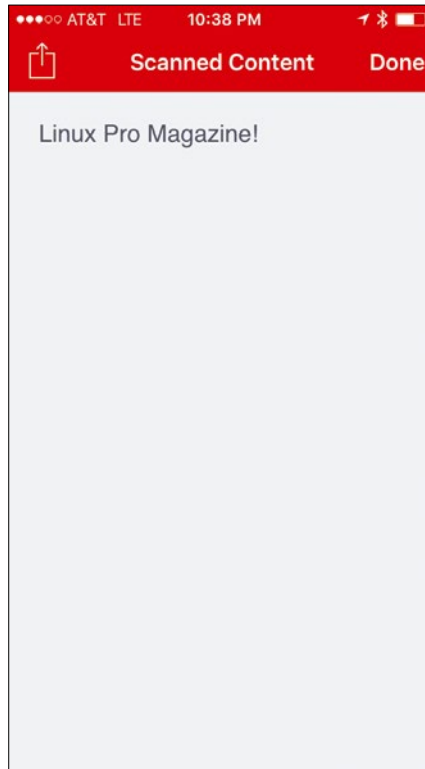


Figure 7: ... identifies the code, and outputs the encoded message.

### INFO

- [1] Linux SDK for the label printer:  
<http://var.dymo.com/US/resources/sdk/linux/>
- [2] Sweet, Michael. *CUPS: Common UNIX Printing System*. Sams Publishing, 2001.
- [3] Code for this article:  
<ftp://ftp.linux-magazine.com/pub/listings/magazine/183/>

### LISTING 2: `dymo-qr`

```
01 #!/usr/local/bin/perl -w
02 use strict;
03 use Image::PNG::QRCode 'qrpng';
04 use Graphics::Magick;
05 use File::Temp qw( tempfile );
06 use PDF::Create;
07
08 my( $text ) = @ARGV;
09 die "usage: $0 text" if !defined $text;
10
11 my $width      = 144;
12 my $height     = 54;
13 my $pdffile    = "label.pdf";
14
15 my( $fn, $pngfile ) =
16     tempfile( UNLINK => 1,
17             SUFFIX => ".png" );
18
19 ( my $jpgfile = $pngfile ) =~
20     s/\.png$/\.jpg/;
21
22 qrpng ( text => $text,
23         out  => $pngfile );
24
25 my $img = Graphics::Magick->new;
26 $img->Read( $pngfile );
27 $img->Write( $jpgfile );
28
29 my $pdf = PDF::Create->new(
30     filename => $pdffile );
31
32 my $jpg = $pdf->image( $jpgfile );
33
34 my $page = $pdf->new_page(
35     'MediaBox' => [ 0, 0, $width, $height ] );
36
37 $page->image( 'image' => $jpg,
38     'xscale' => 0.5, 'yscale' => 0.5,
39     'xpos'   => 50, 'ypos'   => 0 );
40
41 $pdf->close;
```

Manage membership data through your website with Admidio

# Members Only



Admidio manages membership data online no matter what the platform and also manages event dates, times, and participants elegantly. *By Andreas Reitmaier*

Large associations often use commercial software to manage membership data. Smaller associations with more limited budgets sometimes face the difficult choice of spending money on an expensive commercial alternative or spending time and energy on rolling their own membership management solution. An open source tool called Admidio offers a better option. Admidio is a free user management tool for websites and organizations, and it even comes with modules for supporting member lists, managing events, building a guestbook or photo album, and creating a user-friendly download area.

## Setting Up

Admidio is a classic web application that requires at least PHP 5.3.7 and MySQL 5.0 or PostgreSQL 9.0 for the server. The installation is similar to other Linux tools: First, download the latest version from the developer's homepage [1] and unzip the archive on your computer. Re-

move the components related to the software version from the folder name. Then, move the folder to a web server using an FTP program.

You should assign the 777 rights for the `adm_my_files` directory on the server, because the server needs to write some files there. The software creates the con-



Figure 1: Admidio requests access information for the MySQL database in the setup dialog.

figuration file during installation. You can perform the initial configuration via the `http://<Server>/admidio/index.php` page in the browser. This page lets you specify data about the server on which the software is running (Figure 1), along with details on the association and the administrator.

If everything is fine, the installation is performed. Under certain circumstances, the setup script might not succeed in saving the `config.php` file in the `adm_my_files` directory. In this case, you can download the file and then manually copy it into the `admin` directory using FTP.

## Tour

Admidio comes up with a clutter-free interface. The larger, left pane shows current information (first for the manufacturer then for the user). On the right, you'll find a login pane and an overview of the available modules.

After logging in, the program only displays the modules that the respective user is allowed to use in the main pane. The first user created is automatically made the administrator and therefore has access to the entire management area. The left-hand side displays an explanation about the functions of the individual modules; on the right in the sidebar are links to the individual sections (Figure 2).

The panes with the messages, the current notifications, and the appointments department – where there is the option to register for activities – are the most important panes for normal users (i.e., association members). It is also the administrator's responsibility to manage the membership data. Among other things, this includes activating members who have registered directly via the website.

## New Members

The Admidio homepage offers new members the opportunity to register themselves, either in short form or in detail. Alternatively, you can disable this function and provide a classic application form as a PDF file. When new members register, they receive a message saying that the administrator first needs to activate their login. The administrator also receives a notification via email so that no new registration is missed.

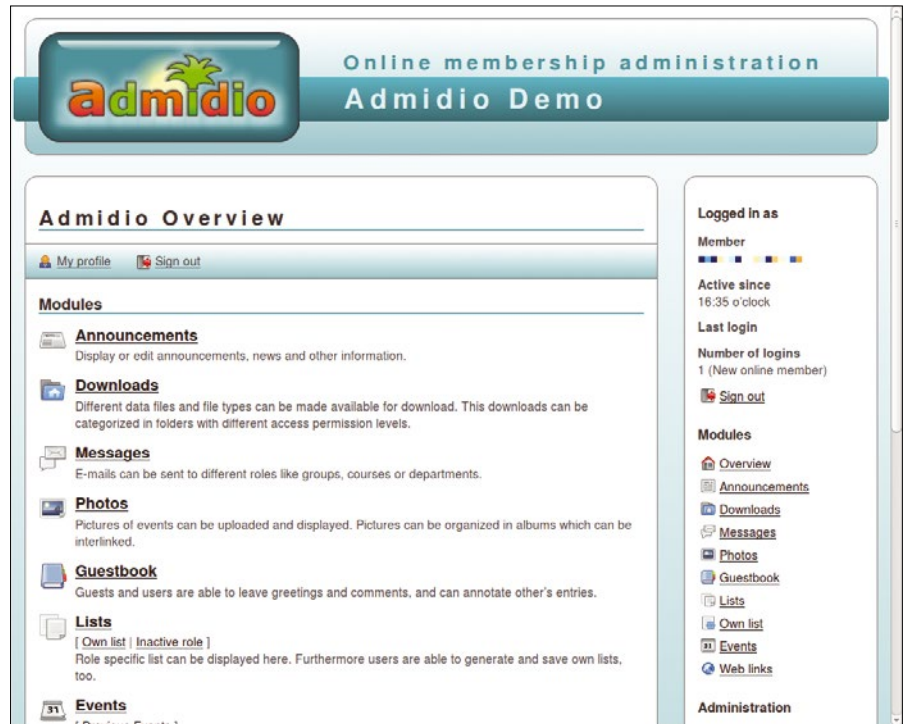


Figure 2: The plain Admidio interface is well integrated on the existing pages and can easily be adapted to other designs.

The list of new registrations can be found under *New registrations*. You can open the data via the *Assigning registration* (plus) button and add additional information to it as required (Figure 3). You can then assign the new member a role – the *Members* role by default – and save the data. The member is thus created and can log in with the entered data.

All members appear in the *User management* menu. You can access the data records here and make changes to the roles and the individual data. To

change the role, access the data record for the person in question by clicking on the name. Within the view, you will find the *Additional Functions* menu at the top right. You can then adapt the role as desired via *Change role memberships*.

## In Practice

However, Admidio can do much more than just manage membership data: It is also possible to set up online registration for courses or other events with the help of the software. The easiest way to

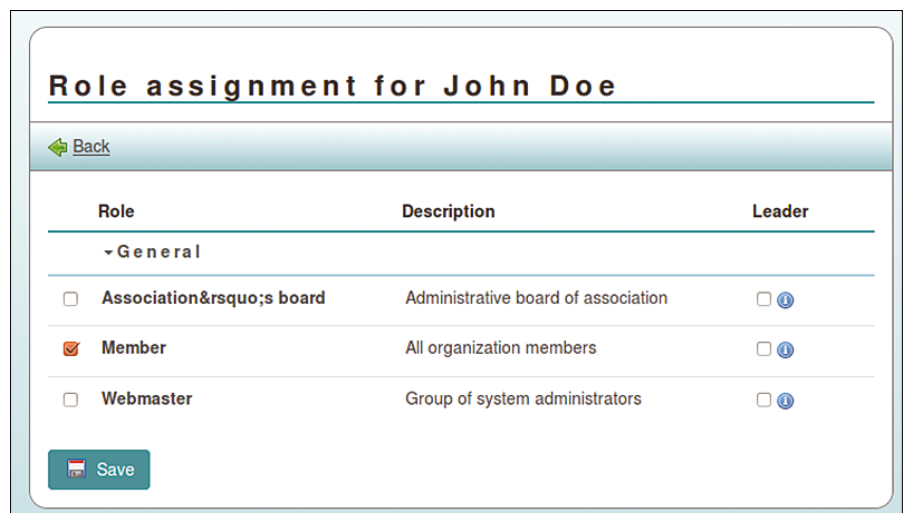


Figure 3: Admidio is only familiar with three roles by default: Member, Association board, and Webmaster. You can add other roles as required.

**Figure 4:** Participants and visitors can register directly for an appointment in Admidio.

do this is to use Admidio's appointment management. For each appointment, you can determine whether members are allowed to register for them. This provides an easy way to create a registration system, because the appointments do not have to be one-off events, they can include courses and other regular activities.

You need to login as an administrator to create an appointment. In *Events*, you will find the *Create Events* button, which you use to access the input mask

list stand out. This can be used for an annual general meeting or compulsory courses, for example. Under that, you can activate the *Enable registration to this appointment* option as required. The software then shows the *Participate* button in addition to the information in the list. Members can then register for the event using this button. Admidio initially shows the registration dialogue when *Participate* is clicked for interested parties who are not yet registered in the online system.

for a new entry. Here you first need to enter the usual data (e.g., date, location and description). The interesting *Visibility & Registration* is in the third pane down. Here you first need to specify which group is allowed to see the event. The group *All (also Visitors)* is the default selection, which you can define using the existing roles within the system.

Below this, you'll find the *Highlight date* option, which makes the entry in the

The *Show participants* button allows administrators to take a look at all currently available applications at any time (Figure 4).

## Integration

Admidio can be accessed as a separate page or integrated into an existing website. You have a couple of options for the latter: If you want to use all of Admidio's functions, it makes sense to integrate it inconspicuously via a link and thus to use an inline frame (`iframe`). Alternatively, you can integrate specific areas into an existing website. Admidio provides some direct links for this purpose, which you can integrate very easily into existing HTML or PHP pages.

You need to understand the structure of Admidio when deciding between these two methods. The software is based on modules and plugins. Modules are all components that provide functionality in the default installation. Plugins, as in other systems, provide additional functions, which you can activate in just a few steps. You need to take a couple of additional steps to integrate such plugins into other websites.

For example, you would use the list of events for registering in a course. This is a module because it represents a basic function. You can integrate this within an external page in the form of an `iframe`. An example call for the *Events* module is shown in the first line of Listing 1. You can show the complete list of appointments if the space provided is big enough, and you can customize the display by attaching a corresponding parameter to the HTML call [2].

Use an event's ID as the parameter to integrate a single event in a website. The easiest way to find this is to access the events as an administrator and then open the desired event in the detail view. In the URL, you will see the `dat_id` parameter, which indicates the corresponding ID. Add this value to the call (Listing 1, line 2). However, the *Events* headline displayed here doesn't really fit that well with a single appointment. So, if necessary, you

### LISTING 1: Example Call

```
http://www.example.org/admidio-file/adm_program/modules/dates/dates.php
http://example.biz/admidio/adm_program/modules/dates/dates.php?id=2
http://example.biz/admidio/adm_program/modules/dates/dates.php?id=2&headline=Registration
```

### LISTING 2: Integrate common.php File

```
<?php
include_once($_SERVER['DOCUMENT_ROOT']. '/<Path/to/Admidio>/adm_program/system/common.php');
?>
```

### LISTING 3: Integrate Access to Plugins

```
<?php
include($_SERVER['DOCUMENT_ROOT']. '/<Path/to/Admidio>/adm_plugins/login_form/
login_form.php');
?>
```

can specify another headline using the headline parameter (line 3).

Integrating modules doesn't really create major problems either with a static website or with content management systems (CMS). The cooperation with WordPress is particularly simple – at least when it comes to integrating the plain appointment calendar in one of the sidebars, for example. A special plugin for this is available in the official WordPress plugin directory [3]. It displays a list of the upcoming events with just a few clicks. The advantage is that the display appears in the exact design of the theme used. A detailed manual within Admidio's online documentation describes the integration into Joomla.

## Plugins

The procedure is a bit more complex with plugins. They are a bit different with Admidio than you might be used to with WordPress or Joomla. There are small function definitions, which you first need to copy manually into the adm\_plugins directory. Then, you place a ref-

erence to it in the desired place on the Admidio website.

First, you integrate the common.php file at the start of a PHP site on an external website (Listing 2). You can also rename a simple HTML page from <Page>.html to <Page>.php. Make sure that the path to the Admidio folder corresponds exactly to the environment on the server concerned.

You integrate access to the actual plugins (Listing 3) at the exact point in the source code where you want to see the plugin output – usually a special area referred to as widgets. You should thoroughly test the construct to make sure incorrect output or formatting doesn't destroy the layout.

Integrating additional modules or plugins into an existing website is accomplished in the same way. However, this approach only makes sense for individual cases in which the user needs to input something directly, such as when making appointments. However, if you need several components, you should consider integrating the whole Admidio interface in the website.

## Conclusions

Admidio proves to be clear and relatively easy-to-adapt software for membership management in associations. It is relatively easy to integrate plugins, modules, or the complete interface into an existing website. For standard modules, this should even be possible for users who are less familiar with working with HTML files. Integrating the plugins, however, requires a bit of background knowledge. The documentation provides numerous notes and aids, and there is also an active forum. ■■■

## INFO

- [1] Download Admidio: <http://www.admidio.org/index.php?page=download> (in German)
- [2] List of modules and associated parameters: [http://www.admidio.org/dokuwiki/doku.php?id=de:2.0:individueller\\_aufruf\\_der\\_module](http://www.admidio.org/dokuwiki/doku.php?id=de:2.0:individueller_aufruf_der_module) (in German)
- [3] WordPress plugin for Admidio events: <https://wordpress.org/plugins/admidio-events/>

# IT Highlights at a Glance



Too busy to wade through press releases and chatty tech news sites? Let us deliver the most relevant news, technical articles, and tool tips – straight to your inbox.

Linux Update • ADMIN Update • ADMIN HPC • Raspberry Pi

Keep your finger on the pulse of the IT industry.

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

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

Raspberry Pi: [www.raspberry-pi-geek.com/mc/subscribe](http://www.raspberry-pi-geek.com/mc/subscribe)



## Conditional formatting in LibreOffice spreadsheets

# Sparklines

Integrate graphical information alongside the data it represents with conditional formatting and sparklines.

By Bruce Byfield

### AUTHOR

Bruce Byfield is a computer journalist and a freelance writer and editor specializing in free and open source software. In addition to his writing projects, he also teaches live and e-learning courses. In his spare time, Bruce writes about Northwest coast art. You can read more of his work at <http://brucebyfield.wordpress.com>

Spreadsheets are all about detailed data and its manipulation. Sometimes, though, you need a summary for gathering general impressions and discerning trends. That is where conditional formatting becomes useful in LibreOffice Calc – as a graphical summary for interpreting general trends in data at a glance, rather than studying the figures closely.

Contrary to first impressions, Conditional Formatting in Calc shares only part of its name with Conditional Fields

or Conditional Styles in Writer. All the three features have in common is that their appearance changes depending on context. In the case of conditional formatting, for instance, the format changes when the content of the involved cells changes.

Conditional formatting is an extension of the idea of sparklines, a concept named and popularized by data visualization theorist Edward Tufte, who devoted 20 pages to the concept in his book, *Beautiful Evidence* [1]. Sparklines are small line graphs that fit on a line of text or in a single spreadsheet cell. They contain no figures, but from their shape, you can see the general trends in the data they map.

According to Tufte, the advantages of sparklines [2] are that they:

*... vastly increase the amount of data within our eyespan and intensify statistical graphics up to the everyday routine capabilities of the human eye-brain system for reasoning about visual evidence, seeing distinctions, and making comparisons. And data graphics are no longer a special occasion in a separate place with a frame on some slide with a label "Fig. 17-B" .... Providing a straightforward and contextual look at intense evidence, sparkline graphics give us some chance to be approximately right rather than exactly wrong.*

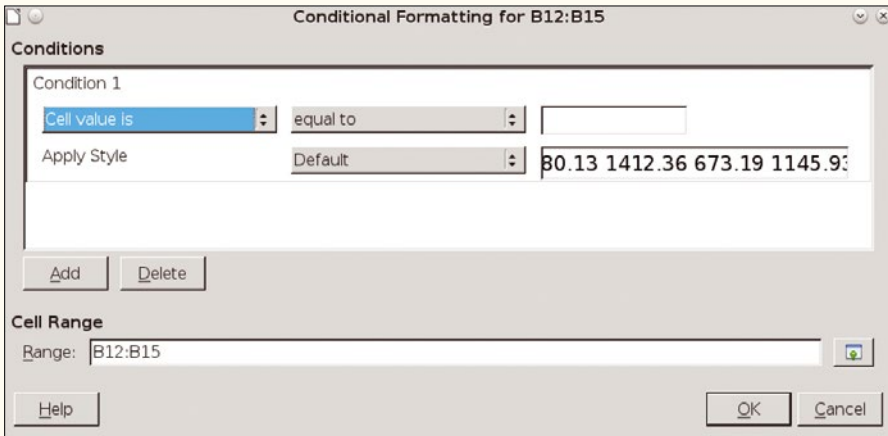
Excel supports sparklines, but Calc does not, although some proof-of-concept macros were circulating a few years ago. However, conditional formatting offers several alternatives to sparklines that provide the same visual advantages while adding some complexity to the basic idea.

### How Conditional Formatting Works

Conditional formatting in Calc automatically changes the appearance of selected cells. The automatic format can be the application of a cell style or of an indicator that visually presents data in much the same way as a graph or a chart.

To use Calc's conditional formatting, select the cells to work with, then click *Format | Conditional Formatting | Condition*, and select an item from the submenu. All types of conditional formatting can be configured from *Condition*. The *Color Scale* and *Data Bar* menu





**Figure 1:** The Conditional Formatting window sets the kind of information to be indicated by the formatting.

items are shortcuts to options available under *Condition*. If necessary, you can edit the range of cells affected at the bottom of the Conditions window (Figure 1).

The Conditions pane has three fields. From left to right, they are:

- The general condition.
  - The filter to refine the condition.
  - The numerical value that must be present to activate the conditional formatting.
- Below these three fields, you have the option to select a cell style.


By changing the general condition, you can choose from the following types of conditional formatting:

- *Cell value is*: Applies the selected cell style when the numerical value is met. This type of formatting is useful for emphasizing target values in a range of cells. It cannot be used for cells formatted for text.
- *Formula is*: Applies the selected cell style to cells in which the designated formula is used. The formula is typically one that viewers of the sheet want to find easily.
- *Date is*: Applies the selected cell style to cells in which the designated filter is used, from *Today to Last Week*. This type is especially useful for locating recent information.
- *All Cells | Color Scale (2 entries)*: Creates a gradient of two colors. The fields refer not to formulas, but to target values. The color scale is especially useful for showing high and low values in a range of cells at a glance (Figure 2).
- *All Cells | Color Scale (3 entries)*: Like a color scale for two entries, except that a third target value is added, often a midpoint using the value *Percent*.
- *All Cells | Data Bar*: A gradient that creates a graph-like representation, typically showing how far a cell value is above or below a designated norm (Figure 3).
- *All Cells | Icon Set*: Adds a set of icons to summarize the contents of cells (Figure 4). Available icons include arrows, flags, checkmarks, bar graphs, emoticons, and quartered circles. Each icon set has three or four icons, with each depicting a different state. For example, traffic light icons or emoticons might designate if the results were above, below, or equal to projections.


Once you have set conditional formatting, you can select *Manage* from the submenu to see a summary of the instances of Conditional Formats and *Edit* the selections (Figure 5). You might also want to add a caption to the formatted cells to explain what is being displayed.

### Adding a Sparkline


If you want to use a sparkline, you can create one manually with the following procedure:



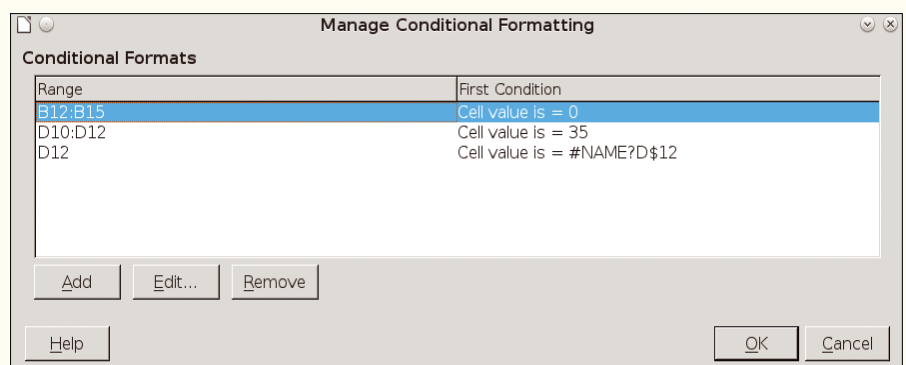
**Figure 2:** A color scale, with yellow indicating the lowest value and cyan the highest. Mid-range values are different shades of green.



**Figure 3:** A data bar in which high values are indicated by more blue in a cell.



**Figure 4:** An icon set annotates results. Here, emoticons indicate that two results were neutral, and two were better than neutral.



**Figure 5:** Click *Manage* in the Conditional submenu to Add, Edit, or Remove conditional formatting.


	A	B	C	D	E
1	Q1	Q2	Q3	Q4	
2	\$145.78	\$75.31	\$97.40	\$149.98	

Figure 6: You can edit a chart, then shrink it down to create a sparkline.

1. Highlight the range of cells for the sparkline.
2. Select *Insert | Object | Chart*.
3. Create a *Line* graph. Use the *Lines Only* subtype or, if you want to emphasize the points on the graph, the *Points and Lines* subtype.
4. Remove each element of the chart except the line by deleting it or setting it to the same color as the background (usually white).
5. Select the chart and drag it while pressing the Shift key, so it changes size without distortion (Figure 6).
6. Position the line graph by the range of cells it represents.

The sparkline will update if you click it or save and reopen the spreadsheet after changing one or more figures in the range of cells.

### New Dimensions to Spreadsheets

Both conditional formatting and sparklines are relatively new concepts. So far as anyone can tell, they appear to be less than 20 years old. As the workaround for sparklines in Calc indicates, they are extensions of charts and graphs, except at a higher level and with less detail. Also, unlike charts and graphs, they are integrated with the data they represent.

Think of them as the graphical equivalent of captions – an annotation that helps you make sense of your data. Even if you do not use sparklines or conditional formatting when you first design a spreadsheet, you might decide to add them later, especially in a spreadsheet that has a long, active lifetime. ■■■

### INFO

- [1] Tufte, Edward. *Beautiful Evidence*. Graphics Pr, 2006, [http://www.edwardtufte.com/tufte/books\\_be](http://www.edwardtufte.com/tufte/books_be)
- [2] Sparklines: [http://www.edwardtufte.com/bboard/q-and-a-fetch-msg?msg\\_id=0001OR](http://www.edwardtufte.com/bboard/q-and-a-fetch-msg?msg_id=0001OR)

## 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)

Shop the Shop

shop.linuxnewmedia.com

# GIMP HANDBOOK



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

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

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

Order online:

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



**FOR WINDOWS, MAC OS, AND LINUX USERS!**



## Adding and managing users and groups

# Beyond the Basics

The command line provides all the necessary tools for creating and maintaining multiple accounts. By understanding these commands, you can more effectively manage your system. *By Bruce Byfield*

### BRUCE BYFIELD

Bruce Byfield is a computer journalist and a freelance writer and editor specializing in free and open source software. In addition to his writing projects, he also teaches live and e-learning courses. In his spare time, Bruce writes about Northwest Coast art. You can read more of his work at <http://brucebyfield.wordpress.com>

If you have ever done any system administration, chances are you have added an account or a group.

However, both tasks offer possibilities beyond the plain command. Not only do numerous options exist to modify the basic command, but a variety of other commands are available for manipulating users and groups and viewing their activities.

Users and groups are means of controlling access to a Linux system. A user account gives normal access to a system, whereas user membership in a group gives access to different hardware, subsystems, and files. Typically, each user account is matched with a group of the

same name, so that users can access the files in their home directory. In some distributions, only the ordinary user created during installation has full non-root access, and you might have to add new users to each group to which the first user belongs.

### Adding and Deleting

In Debian-based distributions, the easiest way to create a new user is with `adduser`, a script that leads you through the process. After you enter the command `adduser NAME`, you are prompted for a name, a password, and optional contact information that in effect turns a list of users into a contact list. Other

information, such as the user ID (UID), is created for you, starting with 1000 (Figure 1). Non-Debian distributions also include `adduser`, but in most cases it is an alias for `useradd`.

In all distributions, the basic command for creating users is `useradd`. However, unlike `adduser`, entering the command and a username is not enough. To start, you need to create a home directory for the account with `--create-home (-m)`. If you do not want the home directory to be a subfolder of `/home`, you need to specify the base directory with `--base-dir DIRECTORY (-b)`. When making a home directory, you probably want to include the option `--skeleton DIRECTORY (-k)` to add default files to it.

Additionally, you'll probably want to specify the password with `--password PASSWORD (-p)` and groups beside the account's own group with `--group GROUPS (-G)`. Other characteristics of the account will be those listed by entering `useradd -D`, an option that can also be used for editing the defaults with `useradd -D OPTION`. You can also specify the account's shell with `--shell SHELL (-s)` and its UID with `--uid UID (-u)` (Figure 2).

The `useradd` command can also set a couple of options for user's passwords, although they seem to be little used on smaller systems. For example, with `--expire DATE (-e)`, you can set the date on which an account's password expires. Usually, you will want to accompany `--expire` with `--inactive DAYS (-f)`, to set the number of days after expiration to disable an account. When the account is disabled, its files are preserved, but the user cannot log in.

Except for the options for password expiration, groups have a similar set of commands. Both the Debian `addgroup` command and the more generally used `groupadd` have options similar to `adduser`. In both, you can specify the group ID (GID) and a password. The main difference from the basic user commands is that in both group commands you can use the `--system` option to create a group that helps to run the system, instead of one to which users can be assigned.

To remove users, Debian-based systems have `groupdel`, with the convenient option `--backup-to DIRECTORY`, which automatically removes all groups the user is in. The exception is the user's private group, which cannot be deleted until the user is removed. The `userdel` command has the option to `--remove` the home directory or, in case the user has files elsewhere in the system, to `--remove-all-files`. By contrast, `userdel` has only the options to `--force (-f)` deletion or `--remove` to delete the user's home directory.

### Usermod and Groupmod

The `usermod` and `groupmod` commands are for editing users and groups after they are created. Many of the `usermod` options mirror those of `useradd`, including `--shell SHELL (-s)`, `--uid UID (-u)`, `--expiredate DATE`, and `--inactive DAYS (-f)`.

To this set of commands, `usermod` adds `--login NAME (-l)`, which can only be changed when the user is not logged in, and `--password PASSWORD (-p)`.

Other options change the groups to which the account belongs. With `--gid GROUP (-g)`, the root user can change the initial group name or GID for an account. Group membership is modified by `GROUPS (-G)`, in preference to editing `/etc/group` in

```
root@nanday:/etc# adduser plw
Adding user `plw' ...
Adding new group `plw' (1006) ...
Adding new user `plw' (1005) with group `plw' ...
Creating home directory `/home/plw' ...
Copying files from `/etc/skel' ...
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
Changing the user information for plw
Enter the new value, or press ENTER for the default
  Full Name []:
  Room Number []:
  Work Phone []:
  Home Phone []:
  Other []:
Is the information correct? [Y/n]
```

Figure 1: The `adduser` tool (not to be confused with `useradd`) provides a wizard-like script for creating users.

```
root@nanday:/etc# useradd -mkU --password aratinga -U macaw
root@nanday:/etc# finger macaw
Login: macaw                               Name:
Directory: /home/macaw                     Shell: /bin/sh
Never logged in.
No mail.
No Plan.
```

Figure 2: The `useradd` command requires careful planning of options to set up a new user.

a text editor, which does not update `/etc/gshadow`. The groups are specified in a comma-separated list with no whitespace.

An especially useful option for `usermod` is `--lock (-L)`, which prevents anyone from using the account to log in. The lock is represented by adding an exclamation mark (!) at the start of the password. The lock is applied with no warning or confirmation message and can be removed with `--unlock (-U)`. As you might expect, neither can be used together or with `--password (-p)`, although changing the password would be just as effective in preventing the account from being used.

The `groupmod` command has far fewer options than `usermod`, possibly because it can potentially have farther-reaching effects – in fact, on systems that use `sudo` rather than the root account, careless use could leave you unable to do any administration. At any rate, `groupmod` uses only three options: `--gid` to change the GID, `-n` to change the group name, and `--password (-p)` to add or change the group password – an option that may be useful on a large system but often not the average home setup.

### Other Administration Commands

Several other commands for both users and groups also exist. For example, `groups USER` lists the different groups to which the specified user belong. The same information can be obtained from `id USER`. At one time, the `users` command could be used to see user activity listed in the logs. However, because the logs are binary now that Systemd is used in most major distributions, you need to use `journalctl` to read them or obtain some of the same information with a combination of `finger` and `who` instead.

On systems that use shadow files to help conceal passwords, you can use a set of four commands to set up and edit the system, making sure that key files are in sync.

- `--pwconv` creates `/etc/shadow` from `/etc/passwd`.
- `--pwunconv` creates `/etc/passwd` from `/etc/passwd` and `/etc/shadow`, then removes `shadow`.
- `--grpconv` creates `/etc/gshadow` from `/etc/group`.
- `--grpunconv` creates `/etc/group` from `etc/group` and `etc/gshadow`, then removes `gshadow`.

On larger systems, where user accounts are created and deleted regularly, these commands can help avoid possible problems.

A somewhat safer alternative for syncing files is `grpck`. This command edits `/etc/group` and `/etc/gshadow`, the file that helps hide group information on some systems (Figure 3). To be specific, `grpck` checks the validity and uniqueness of each group's name, GID, and members, as well as looking for matching entries in `/etc/gshadow` and removing duplications or obsolete or corrupted information.

With the addition of the `--read-only (-r)` option, `grpck` lists entries that need correcting without making any changes – an option that should be run first to avoid any problems. The command can also use `--sort (-s)` to arrange information alphabetically, instead of adding newer entries at the bottom of the list.

### Tools for a Multiuser System

As these commands show, the topics of users and groups is far more complex than most people typically believe. For those on Debian-based systems, the convenience of `adduser` and `addgroup` scripts is paid for by a certain laziness (as well as an inevitable confusion between the commands). For others, especially on home systems, some of the available options are a reminder of the built-in tools for encouraging the regular changing of passwords.

Still, the assortment of commands for administering users and groups should come as no surprise. Most Unix-like systems are designed for multiple accounts, so naturally the command line provides the tools for creating and maintaining them as well. By learning some of these commands, you can not only help your system run better, but more efficiently as well. ■■■

```
root@nanday:~# grpck --read-only
'bb' is a member of the 'cdrom' group in /etc/group but not in /etc/gshadow
'bb' is a member of the 'floppy' group in /etc/group but not in /etc/gshadow
```

**Figure 3:** Using `grpck --read-only` lists differences between `/etc/group` and `/etc/gshadow` but does not make any changes.


Shop the Shop

shop.linuxnewmedia.com



## Subscribe Today!

Each issue of Raspberry Pi Geek is a new adventure, with ingenious applications and cool projects for Raspberry Pi and the open hardware revolution.



**For a limited time only!**

Order issues #4, #5, and #6 of Raspberry Pi Geek for only US\$ 30\*!

\* £ 14.99 / € 23.99 / \$ 33.50 rest of world



Image © Daniel Villeneuve, 123RF.com



➤ shop.linuxnewmedia.com



Build your own tracking app

# Snap and Track

Build a simple and handy solution for recording and tracking locations; in the process, learn how to automate Android and build a simple Python-based app. *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).

**A** dedicated GPS tracker or a GPS tracking app can come in handy in many situations. Recording and saving your movements can help you geocorrelate photos, keep track of places you've visited, maintain a simple travelogue, and more. Although several excellent apps can be used to transform an Android device into a capable GPS tracker, most of them are one-trick ponies designed to record your location continuously and save the geographical data in the GPX format.

What if you need a tool that does more than that? For example, suppose you want to use your Android device to snap a photo and then save the relevant data – such as date and time, latitude

and longitude – in a database or a text file. Wouldn't it be great to have a web-based app that runs on a Linux server and lets you view your snapshots and data using a regular browser? If you have a working knowledge of Python scripting, you can build a tool like this yourself, and this article shows you how to do that.

## General Notes and Preparatory Work


The solution you are about to build (I'll call it Geofix) consists of two parts: (1) a simple Python script running on the Android device that automates the process of taking snapshots, collecting the relevant data, and saving it in an SQLite database and (2) a Python-based web app that acts as a front end for viewing the snapshots and the data in the database.

Android doesn't natively support Python scripting, so to run Python scripts on your Android device, you need to install the Scripting Layer for Android (SL4A) and Python for Android packages on it. To do this, grab the latest APK packages of SL4A [1] and Python for Android [2] for the appropriate architecture and install them on the Android device.

Most mainstream Linux distributions come with Python installed by default. That means any machine running a popular Linux distribution like Debian or Ubuntu can be used to run the Geofix web app. Instead of writing a web app from the ground up, you can save a lot of time and effort by using a dedicated web framework. You can choose from several excellent Python web frameworks, and Bottle [3] is probably the simplest and most lightweight of these – which makes it a perfect tool for building the Geofix web app. To install Bottle on a Debian or Ubuntu-powered machine that will serve the web app, run the following commands as root:

```
apt-get install python-pip
pip install bottle
```

Although Geofix is a relatively simple solution, listing all of its code here is not very practical, so before you proceed, grab the latest version of Geofix [4] from the GitHub repository or clone it using:

```
git clone 
https://github.com/dmpop/geofix.git
```



## The Android Python Script

With all the pieces in place, you are ready to dissect and study the `geofix.py` Python script that runs on the Android device. This relatively short and simple script performs several actions. To keep things tidy and organized, the script uses dedicated directories for storing the database and snapshots.

The script starts by checking whether the directory specified as a value of the `geofix_dir` variable exists. If the directory doesn't exist, the script creates it, along with the `snapshots` sub-directory for storing photos:

```
if not os.path.exists(geofix_dir):
    os.makedirs(geofix_dir)
    os.makedirs(geofix_dir + 'snapshots/')
```

Thanks to so-called facades (friendly wrappers for Android APIs), the SL4A and Python for Android can access and make use of various Android functions. The script uses three facades to toggle location features and obtain geographical data:

```
droid = android.Android()
droid.startLocating()
droid.eventWaitFor('location', int(wait))
location = droid.readLocation().result
droid.stopLocating()
```

The `startLocating()` facade enables the locating functionality, the `eventWaitFor()` facade pauses the script for the period of time specified in the `wait` variable to obtain geographical data, and the `stopLocating()` facade disables locating. The period of time required to obtain geographical data depends on several factors, and you might need to adjust the default `wait` value to give the script enough time to receive the data.

Instead of modifying the wait value in the script, you can add a dialog that lets you specify the wait time on the fly. To do this, replace the `wait = 9000` line with the following code:

```
droid = android.Android()
droid.dialogCreateAlert("Wait time")
wait = droid.dialogGetInput("Specify wait time in milliseconds:").result
```

The locating functionality in Android uses two sources to obtain geographical data: the built-in GPS module and mobile network. This ensures that the device is able to obtain the data even if one of the sources is not available. The code block shown in Listing 1 takes care of extracting latitude and longitude values from the received geographical data.

The script first attempts to obtain latitude and longitude values from the network source (lines 1-5). If this fails, the script tries to extract coordinates from the GPS source (lines 7-11). Next, the script constructs a string containing the coordinates in the `digiKam`-compatible format (see the “Geotagging Photos in `digiKam`” box) and a URL that shows the exact location on the `OpenStreet-Map` map:

```
digikam = 'geo:' + lat + ',' + lon
osm = ↵
'http://www.openstreetmap.org/↵
index.html?mlat=' + lat + '&mlon=' ↵
+ lon + '&zoom=18'
```

### LISTING 1: Obtaining Latitude and Longitude Values

```
01 try:
02     coords = location['network']
03     lat = str(coords['latitude'])
04     lon = str(coords['longitude'])
05     droid.makeToast('Network coordinates: ' + lat + ' ' + lon)
06 except (KeyError):
07     try:
08         coords = location['gps']
09         lat = str(coords['latitude'])
10         lon = str(coords['longitude'])
11         droid.makeToast('GPS coordinates: ' + lat + ' ' + lon)
12     except (KeyError):
13         droid.makeToast('Geofix failed to obtain coordinates.')
14     sys.exit()
```

### LISTING 2: Saving Coordinates

```
01 f_path = geofix_dir + 'geofix.csv'
02 f = open(f_path, 'a')
03 f.write(str(dt) + ', ' + str(lat) + ', ' + str(lon) + ', ' + digikam + ', ' + osm + '\n')
04 f.close()
```

### LISTING 3: Writing Data to geofix.sqlite

```
01 sql_query = "INSERT INTO geofix (dt, lat, lon, digikam, osm_url)
              VALUES ('%s', '%s', '%s', '%s', '%s')" % (dt, lat, lon, digikam, osm)
02 conn = sqlite3.connect(geofix_dir + 'geofix.sqlite')
03 conn.execute(sql_query)
04 conn.commit()
05 conn.close()
```

### LISTING 4: Creating the Database

```
01 conn = sqlite3.connect(geofix_dir + 'geofix.sqlite')
02 conn.execute("CREATE TABLE geofix (id INTEGER PRIMARY KEY, dt VARCHAR, lat VARCHAR,
              lon VARCHAR, digikam VARCHAR, osm_url VARCHAR)")
```

```
droid.cameraInteractiveCapturePicture(geofix_dir + 'snapshots/' + dt + '.jpg')
```

To install `geofix.py` on your Android device, put it in the `s14a/scripts` directory. Although you can launch the script through the SL4A app, you can also create a shortcut to the script on the home screen. This way, you can run the script with a single tap.

## Scripting with Termux

The SL4A and Python for Android combination is not your only option for automating Android. The Termux project [5] equips Android with a terminal emulator and Linux environment. Among other things, this means you can automate Android using Bash shell scripts that make use of popular Linux tools. Better still, the Termux:API add-on enables access to Android's APIs, so you can integrate Android functionality into shell scripts.

Both Termux and Termux:API are available on Google Play Store and F-Droid, so you can install the apps from your preferred source. To replicate the functionality of the `geofix.py` Python script, you need to install a couple of Linux tools, including `termux-api`, which enables Termux:API functionality, `jq` for manipulating data in the JSON format, and `sqlite` for working with SQLite databases. To do this, launch Termux and run the following commands:

```
#apt update
#apt upgrade
#apt install termux-api jq sqlite
```

## GEOTAGGING PHOTOS IN DIGIKAM

If you happen to use digiKam as your preferred photo management application, you can use it together with Geofix to geotag photos. Both `geofix.py` and `geofix.sh` scripts save the geographical coordinates in the digiKam-compatible format. So, to geotag one or several photos in digiKam, just copy the desired coordinates in the digiKam format (e.g., `*geo:56.1831455,10.1182492*`) from Geofix. Then, switch to digiKam, select the photos you want to geotag, and choose *Image / Geo-location*. Select the photos, right-click on the selection, and choose *Paste coordinates*.

The current date and time, latitude, longitude, generated digiKam string, and OpenStreetMap URL are then saved in the `geofix.csv` comma-separated file (Listing 2).

In addition to the comma-separated file, the script also writes the data into the `geofix.sqlite` SQLite database, which the web app uses as its back end (Listing 3). If the database doesn't exist, the script creates it before writing the data into it (Listing 4).

Finally, the script takes a photo using the `cameraInteractiveCapturePicture` facade. The photo is saved in the dedicated `snapshots` directory using the date and time stamp as the photo's file name:

The `geofix.sh` script that automates the required tasks is not particularly complicated, but a couple of commands need some explanation. The `termux-location` command outputs location data in the JSON format, and to extract latitude and longitude, the script needs to pipe the output to the `jq` tool:

```
lat=$(termux-location | jq '.latitude')
lon=$(termux-location | jq '.longitude')
```

The `termux-toast` command can only display text from `stdin` (the standard input connection), so to show a pop-up message, you need to pipe it to `termux-toast` using the `echo` command:

```
echo "Coordinates: $lat, $lon" | termux-toast
```

The `termux-camera-photo` command requires a camera ID value assigned to the `-c` or `--camera` parameter. Usually, the back camera has ID 0, and the front camera has ID 1. To check the IDs of the cameras on your Android device, run the `termux-camera-info` command.

Termux lets you automate tasks on Android using the environment and tools familiar to most Linux users, but it does have a couple of drawbacks compared with SL4A and Python for Android. To run the script (or any Bash shell script for that matter), you need to start the Termux app first and then execute the appropriate command. This can quickly become a serious nuisance if you use the script often. Also, the `termux-camera-photo` command doesn't provide any visual feedback, and you have no control of shooting settings.

## The Web App

The Geofix web app is based on the Python Bottle framework that was covered in a previous issue [6], so to get a grasp of the framework's basics, it's a good idea to peruse that article first.

Geofix is not exactly a complex web app, and it consists of three key components: the `main.py` Python script, which is the core of the app, several `.tpl` template files, and a CSS stylesheet that controls the overall appearance of the app. The `main.py` script has four functions responsible for retrieving data from the `gefix.sqlite` database, displaying the obtained data, together with the accompanying photos as a web page, showing an individual photo, and deleting photos.

# LOST YOUR BOOKSTORE?

## LET US BE YOUR BOOKSTORE

Browse our shop for single issues of *ADMIN*, *Linux Pro*, *Linux Magazine*, *Raspberry Pi Geek*, and *Ubuntu User* – delivered right to your door.

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

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

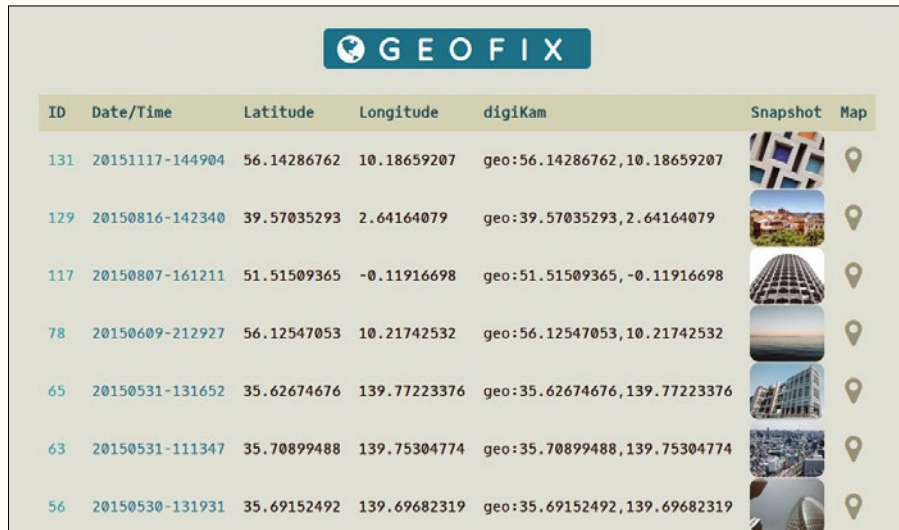
■ [shop.linuxnewmedia.com/subs](http://shop.linuxnewmedia.com/subs)



# shop.linuxnewmedia.com

DIGITAL AND PRINT EDITIONS AVAILABLE!



















ID	Date/Time	Latitude	Longitude	digiKam	Snapshot	Map
131	20151117-144904	56.14286762	10.18659207	geo:56.14286762,10.18659207		
129	20150816-142340	39.57035293	2.64164079	geo:39.57035293,2.64164079		
117	20150807-161211	51.51509365	-0.11916698	geo:51.51509365,-0.11916698		
78	20150609-212927	56.12547053	10.21742532	geo:56.12547053,10.21742532		
65	20150531-131652	35.62674676	139.77223376	geo:35.62674676,139.77223376		
63	20150531-111347	35.70899488	139.75304774	geo:35.70899488,139.75304774		
56	20150530-131931	35.69152492	139.69682319	geo:35.69152492,139.69682319		

Figure 1: Geofix web app in action.

Again, if you read the Bottle article first, you should have no problems figuring out how the Geofix web app works. The app's main page is not overloaded with features and information. The only thing you need to know is that the ID number of each entry links to the delete page, whereas the map marker points to the OpenStreetMap map showing the exact geographical position of the selected entry.

## Practicalities

The Geofix app doesn't resize photos; instead, it shrinks them to thumbnail size on the fly. Adding a lot of photos taken at the maximum resolution supported by your Android device's camera to the Geofix web app will dramatically increase load time and generate a lot of unnecessary traffic. You can solve this problem in several ways; for example, you can configure the default camera app to take photos at a lower resolution (e.g., 640x480).

Alternatively, you can use Termux to resize the photos with the ImageMagick `mogrify` tool (to install it in Termux, use the `apt install imagemagick` command). The following command resizes all JPG images in the current directory to 600px on the longest side:

```
mogrify -resize "800x800>" *.jpg
```

If you use the `geofix.sh` script, you can add the command below to resize the photo automatically after it has been taken:

```
mogrify -resize "800x800>" $geofix_dir"/snapshots/"$dt.jpg
```

To make use of the Geofix web app, you need to copy the `geofix.sqlite` database and all accompanying photos from the `Geofix` directory on your Android device to the static directory inside the web app. You have several options here, including a dedicated synchronization tool like `Syncthing` [7]. Alternatively, you can opt for running `rsync` in Termux.

## Conclusion

Despite its simplicity, a tool like Geofix can handle a variety of useful tasks: from tracking your trips and recording interesting places to helping you geotag photos. With a bit of tweaking, you can put Geofix to other practical uses, too. For example, you can add a description field and modify the web app to transform it into a no-frills travel diary. If you integrate the geocoder Python library [8] into the web app, you can use Geofix for reverse geocoding.

In short, with a bit of creative thinking, you can make Geofix do all kinds of useful and clever tricks. ■■■

## INFO



- [1] SL4A on GitHub: [github.com/kuri65536/sl4a](https://github.com/kuri65536/sl4a)
- [2] Python for Android on GitHub: [github.com/kuri65536/python-for-android](https://github.com/kuri65536/python-for-android)
- [3] Bottle: [bottlepy.org/docs/dev/index.html](http://bottlepy.org/docs/dev/index.html)
- [4] Geofix on GitHub: <https://github.com/dmpop/geofix>
- [5] Termux: [termux.com](http://termux.com)
- [6] "Using the Bottle Framework to Build Python Apps" by Dmitri Popov. *Linux Magazine*, issue 174, May 2015, pg. 80, <http://www.linux-magazine.com/Issues/2015/174/Workspace-Bottle>
- [7] Syncthing: [syncthing.net](http://syncthing.net)
- [8] Geocoder Python library: [pypi.python.org/pypi/geocoder](http://pypi.python.org/pypi/geocoder)

# LINUX UPDATE

## Need more Linux?

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

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



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



### Jitsi Videobridge and Jitsi Meet

# Meeting Place

Jitsi Videobridge and its front end Jitsi Meet offer video and audio conferences in web browsers with real-time chat, Prezi presentations, screen sharing, and Etherpad document editing. *By Mela Eckenfels and Heike Jurzik*

**M**any video conferences these days use a Google or Microsoft solution. The two providers stand out primarily because they are user friendly and have very good video and audio quality; however, the downside is obvious: The data moves through a foreign server. The Jitsi project [1] provides an open source, encrypted, and sustainable alternative. In addition to the messenger available for Linux, Windows, and OS X desktops, the developers also provide a web application called Jitsi Meet [2].

Two components – Jitsi Videobridge and Jitsi Meet – work together to help images and sound reach the users’ computers. Videobridge is a web real-time communication (WebRTC)-compatible Selective Forwarding Unit (SFU) that expands the Jitsi Messenger desktop client to add videoconferencing capabilities. Jitsi Meet is the front end to Videobridge and is implemented in JavaScript. The Nginx web server [3] and the Prosody XMPP server [4] are also required on the server side. (See the “Test Environment” box for a description of the server on which Jitsi Meet was run and the computer and Internet connections of the participants.)

Users can take part in meetings with WebRTC-enabled browsers. The current versions of Opera, Google Chrome, and Mozilla Firefox already support the standard; additional functions such as desktop sharing require browser add-ons. Microsoft’s Internet Explorer and Apple’s Safari browser do not currently include this technology.

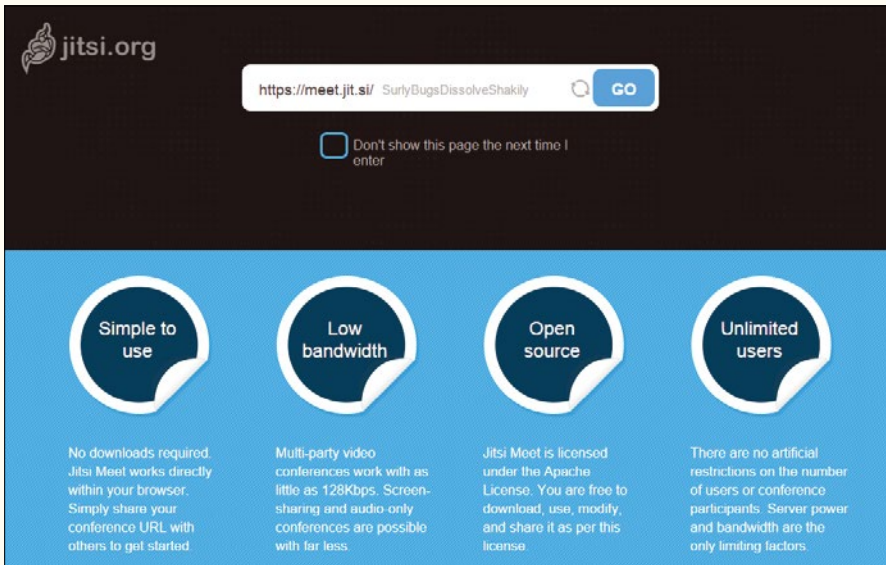
Debian server administrators have an easy task. As described in the Quick Install Guide [5], they just need to add the nightly build repository to the software sources and install the `jitsi-meet` package. The package manager automatically installs additionally required components Nginx, Prosody, OpenJDK 7, Jitsi Meet, and some graphics, video, and audio libraries.

Administrators need to intervene twice: to enter the domain name or IP address in a dialog box for the computer via which Jitsi conferences should be accessible and to designate an SSL certificate on the server. Alternatively, you can press Enter and Jitsi Meet creates a self-signed certificate.

Anyone using a different distribution just needs to follow the server installation instructions [7], which explain how to set up the Prosody XMPP server and the Nginx web server; look after the Java Runtime Environment; download, unzip, and configure the Jitsi Videobridge sources; and set up the Jitsi Conference Focus and Jitsi Meet components.

### Conference Room

Once the installation has been completed, users can call up the Jitsi Meet instance address in their browser. They need to enter either a unique name for



**Figure 1:** Jitsi Meet generates random names on request for the virtual conference rooms. Alternatively, users can define something of their own in the box at the top.

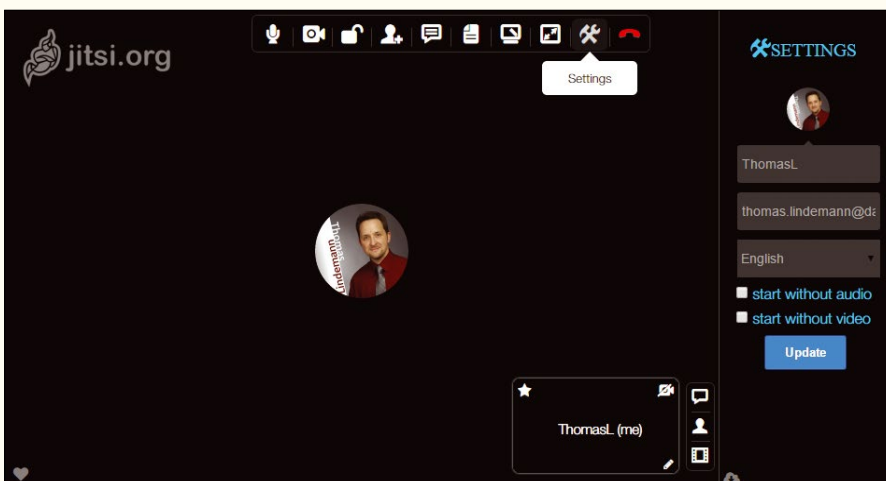
the virtual conference room in the box at the top or confirm the random identifier by pressing *Go* (Figure 1). In some tests at the end of September, the software refused to allocate unique names, and Jitsi Meet conferences fell victim to error 404. This was the case on both their own servers and in the Jitsi project. The bug now seems to be fixed.

The first participant automatically receives moderator rights as the initiator of the conference and can mute individual or all participants, protect the room, or throw people out. It is not possible to designate further visitors as moderators. Upon entering the conference room, all participants first confirm that Jitsi Meet is allowed to access the camera and the microphone. It is only possible to disable the transmission hardware when visiting a second time. To do so, users need to move the mouse to the top to show the Jitsi Meet icons, open the pretty lean *Settings* (Figure 2) by clicking the tool icon, and remove the checkmarks.

Which icons are displayed at the top depends on the extensions installed. Here, users can find buttons to lock the room to further visitors, access the invitation link, open the chat bar on the right edge, switch to full-screen mode and leave the conference.

The project server [6] provides two to three extra buttons – depending on the type of day – for access to Etherpad and Prezi files and your own desktop. However, instructions for installing these modules are missing, and the testers didn't succeed in getting the modules going on their own server.

Jitsi Meet displays the participants at the bottom and the moderator is marked with a star. If you move the mouse over a thumbnail and the small bars, information



**Figure 2:** The camera and the microphone can only be switched off on a second visit.

## TEST ENVIRONMENT

In our lab, Jitsi Videobridge and Jitsi Meet were installed on a LAMP server with six CPU cores, 8GB of memory, 300GB of disk space, and a 64-bit version of Debian Jessie 8.2. The Jitsi components came from the nightly build repository described in the Quick Install Guide [5]. Jitsi Meet and Jitsi Videobridge were stored on the disk on September 24, 2015, and the Jigasi component on October 7, 2015.

The testers also met in a conference room that ran on a server from the Jitsi project [6]. They accessed Jitsi Meet instances from Berlin via a VDSL 50 connection (50Mbps), from Karlsruhe, Germany via cable (also 50Mbps), and from Altena, Germany with T-DSL 1000 (1Mbps). In doing so, they used various operating systems and browsers: Linux Mint 17.2 with Firefox 41.0, OS X Yosemite 10.10.4 with Chrome 45.0.2454.99 and Firefox 40.0.3, and Ubuntu 15.04 with Chromium 45.0.2454.85 and Firefox 41.0. Two tablets were also used: a Samsung Galaxy Tab 3 and Tolino Shine 8.9 (both with Android 4.2.2).

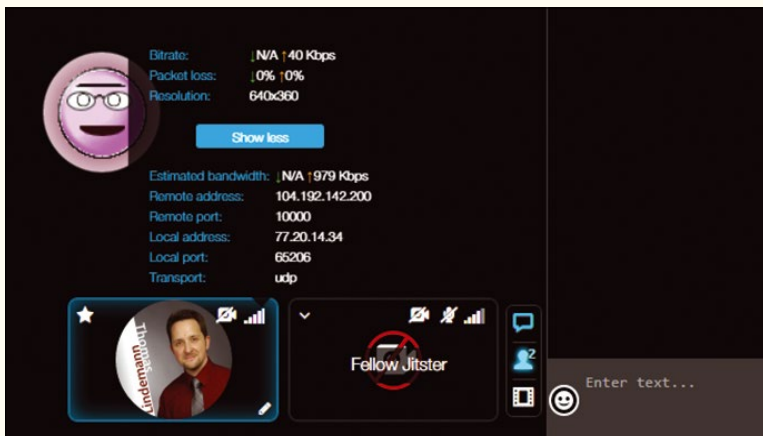


Figure 3: Jitsi Meet provides information about the connection quality.



Figure 4: Prezi presentations appear as independent conference participants in Jitsi Meet.

about the connection quality is displayed. Only your own thumbnail provides information about the IP address, upstream and downstream, and package loss via *Additional Information* (Figure 3). Clicking the icons in the bottom right displays the conference participants in the right sidebar, the text chat, or a contact list.

### Generous

Screen sharing isn't completely problem free, as Jitsi founder and project manager Emil Ivov reports in the video cast FLOSS Weekly [8]. The ability to view another user's desktop directly from within the browser opens up vulnerabilities that websites could exploit with malicious code.

The feature therefore needs an additional browser plugin on the client side. Chrome users will find the Jitsi Desktop Streamer extension in the Chrome Web Store [9]. There is currently an experimental add-on called Jidesha [10] for Firefox.

Desktop sharing didn't work with either Firefox or Chrome on the test server. The testers had a bit more success in a conference room on the Jitsi server – but not with Firefox. The Chrome user only saw incomplete windows and black surfaces instead of an xterm.

The shared Prezi presentation, which only ran on the Jitsi server and not on the test server, gave a slightly better impression. Conference participants had to enter the address for a Prezi into the field provided; the presentation then appeared as a new participant with a preview window at the bottom. In this way, visitors could choose at any time whether they wanted to look at a different user or the Prezi, through which they could move independently (Figure 4). However, control remains with the participant who loaded it in the conference – all the others follow this participant's movements through the document.

Loading the Prezi takes several minutes, and not everything went smoothly. Even with participants connected via VDSL, everything appeared with a few minutes' delay. The Jitsi server also showed all of the controlling owner's actions to the other users with a second time delay. One of the participant's Chrome browser also had trouble with alignment – the Prezi appeared back-to-front.

Thanks to the SIP gateway Jigasi [11], conference participants can add users with SIP phones to the meetings in the browser. It doesn't matter whether it's a hard phone or a soft phone. Administrators can install the Jigasi extension on the Jitsi Meet server. Debian users will find a corresponding package in the Jitsi repository. Everyone else should check out the GitHub repository and build the tool from the resources. A SIP account is also required because Jigasi itself acts as an SIP client, which Jitsi Meet uses to call other participants from the conference.

### Visit from the Outside

The Debian package configuration routine requests the necessary settings for the SIP account (server, username, and password), writes them in the Jigasi setup file `/etc/jitsi/jigasi/sip-communicator.properties`, and stores an entry in the XMPP server configuration (`/etc/prosody/conf.avail`). Admins who install Jigasi manually take care of the necessary settings themselves. They then need to replace the `<<JIGASI_SIPUSER>>`, `<<JIGASI_SIPSERVER>>`, and `<<JIGASI_SIPPWD>>` tags in the `/etc/jitsi/jigasi/sip-commu-`



nicator.properties file with the SIP account data. The password must be base64 coded.

Next, the admin turns to the XMPP server to enter the callcontrol subdomain for the gateway component and the associated secret key into the configuration file in the /etc/prosody/conf.avail directory. System administrators who have installed the Jigasi Debian package should also take a look at the Prosody configuration file. The testers installed and removed the Jigasi component several times, but a new entry ended up in the file for each installation.

Systemd appeared to be responsible for the Jigasi start on the Debian test machine. Administrators need to restart the service after changing the configuration using:

```
systemctl restart jitsi-videobridge
systemctl restart jigasi
```

Anyone who has compiled the sources themselves needs to start the Jigasi service using the

```
/usr/share/jigasi/jigasi.sh
```

script, as described in the instructions.

A small phone icon appears in conference rooms in the web interface; clicking on it opens a box in which the inviting participant needs to enter the SIP ID (Figure 5). If the person being called picks up, that individual can take part in the conversation, but not with video support.

## Reporting

All three testers agreed: The sound quality on the project's Jitsi Meet server was significantly better than on the self-hosted instance, even though their own server had enough power under the hood. Firefox 41.0 was a bit awkward in Ubuntu on the client side, whereas it was a bit more cooperative in Linux Mint. Participants with the Mozilla browser were often left with nothing. For example, the third conference participant disappeared several times as soon as the second visitor left the meeting. The moderator couldn't see or hear the third participant – only the second guest could see and hear everyone else.

The video quality often left a lot to be desired. Only the tester connected via VDSL was generally satisfied. The first participant (cable) and the second (DSL) at times had very pixelated images of each other. The first visitor's image froze several times for the others despite a high-speed Internet connection.

All the computers found it difficult to work with the supposedly lightweight tool. The OS X computer was under a lot of strain, even though it was equipped with 16GB of RAM. Another computer's Cinnamon interface left the service several times.

Jitsi will only be serious competition to Google Hangout, Skype, and the like when the developers eliminate these faults. The software is not currently suitable for use in professional environments but is fine for private individuals or tech-savvy workgroups.

Even though it's possible to install and set up Jitsi Meet quickly on your own server, the conference tool leaves behind a very moderate overall impression, which one of the testers described as the "VW Effect": The fact that Jitsi Meet shone in a few public presentations might be because of highly optimized test environments. ■■■

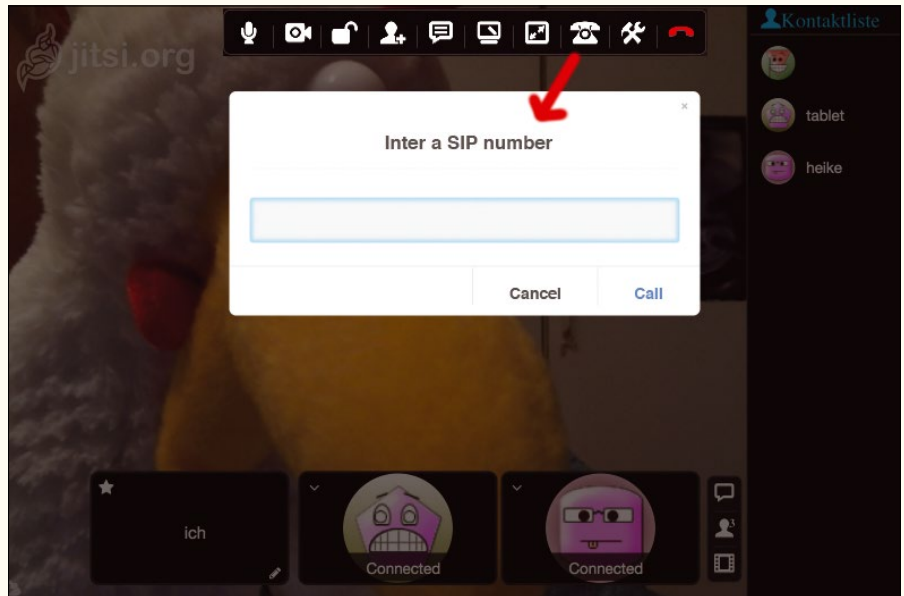



Figure 5: Users with SIP phones can take part in Jitsi Meet conferences thanks to Jigasi.

## INFO

- [1] Jitsi: <http://jitsi.org>
- [2] Jitsi Meet: <https://jitsi.org/Projects/JitsiMeet>
- [3] Nginx: <http://nginx.org>
- [4] Prosody IM: <https://prosody.im>
- [5] Quick install on Debian systems: <https://github.com/turint/jitsi-meet/blob/master/doc/quick-install.md>
- [6] Jitsi Meet on Project Servers: <https://meet.jit.si>
- [7] Manual installation of Jitsi Meet: <https://github.com/jitsi/jitsi-meet/blob/master/doc/manual-install.md>
- [8] Floss Weekly, episode 293: Jitsi Meet: <https://www.youtube.com/watch?v=w--WDNbBE54>
- [9] Chrome Web Store: <https://chrome.google.com/webstore/category/apps>
- [10] Jidesha add-on: <https://addons.mozilla.org/en-US/firefox/addon/jidesha1/>
- [11] Jigasi: <https://github.com/jitsi/jigasi>



## Use mhdds to group hard disks and directories

# United

The multi-hard drive disk filesystem (mhdds) combines directories or hard disks on a union filesystem to create a single, large, virtual filesystem that you can use both locally and via Samba or NFS. *By Ferdinand Thommes*

### AUTHOR

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

**E**stablishing a reliable system and keeping track of a continually growing collection of movies and audio files can be very time consuming. What makes matters worse is that multimedia data typically resides on various disks.

This is where mhdds enters the game: Using a union filesystem, it groups files from different locations to create a virtual directory. The tool not only combines existing data, it also provides details about free storage space on the individual filesystems. (See the box “What Is a Union Filesystem?”)

Consequently, it is no longer a problem to use small disks to store a music collection that extends over three disks. You could just as easily store rock music on one disk, classical tracks on another, and e-books on the third. What happens, however, if your rock music disk is full, but your e-book disk still has room to spare? Things start to become untidy again.

An alternative would be to create a RAID [1] array, but you would always have to compromise between keeping your data safe and using storage space; it does not appear to be a viable solution for the example in this article. The use of LVM [2] only makes sense with RAID for reasons of data safety, and again this does not help solve the problem presented here.

Fortunately, mhdds offers precisely the functionality that most users need in this case: If you run out of space on one of the grouped disks, the data can be migrated in the background to a different disk with free space without the user even noticing. By default, mhdds reserves 4GB on each disk for emergencies: If needed, you can use

```
mlimit=<Limit>
```

to reduce this value down to as little as 100MB at the outset.

### Transparent Write Access

For the virtual array to work, mhdds – in contrast to UnionFS; AuFS, as commonly used by Live media; or OverlayFS, which was recently added to the kernel – not only makes read access transparent, but also data writes. Whereas legacy union filesystems rely on the copy-on-write (COW) [3], here mhdds not only writes to the top level of the filesystem, but to all underlying levels, too.

Mhdds stores files that you add to the virtual array on the first hard disk, as long as it has sufficient space (i.e., as long as the `mlimit` is still upheld). After this, it checks the remaining disks in sequence to see if they have sufficient space. If none of the `mlimits` on the disks meet the requirements, mhdds uses the disk with the most space.

Mhdds always stores files atomically, avoiding the kind of file splitting that you see with LVM. This works on all popular Linux filesystems, including

## LISTING 1: Example Disks

```
$ df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda1       80G   50G   30G   63% /mnt/Movies
/dev/sdb1       40G   35G    5G   88% /mnt/Photos
/dev/sdc1       60G   10G   50G   17% /mnt/Music
```

Samba and NFS, because both return correct information about occupied and free space on the respective filesystems. SSHFS does not meet this criterion and the mhdds developers thus warn against integrating it.

If mhdds notices during a write that the disk in question does not have enough space, it moves the data it has already written to another disk with more space as a background operation and continues the write action on that disk. The writing program does not notice this. In other words, you can work with the virtual filesystem as if you were working on a single large disk.

No matter where data resides or how much space is available on individual disks, you only see the complete remaining free space. If you later buy a disk with enough capacity and decide to stop using the smaller disks in the mhdds array, or if you want to use the smaller disks elsewhere, you can simply copy the content of the virtual filesystem to the new disk and unmount the smaller disks.

## Flexible Storage

Mhdds is available from the repositories of most distributions; you can thus use your distribution's package manager for the install. If you prefer to build mhdds yourself, you can pick up the source code from the mhdds Subversion repository [4]. Using the tool is very easy in practice. In the following example, I use three hard disks: sda1, sdb1, and sdc1; Listing 1 shows the situation at the start.

You can now create a new mountpoint for the array you will be creating and assign the permissions by typing:

```
mkdir /mnt/media
chmod 775 /mnt/media
```

From now on, the FUSE filesystem, which as installed to fulfill one of mhdds's dependencies, comes into its own with its ability to migrate kernel space functions to userspace.

You do not need to be root to use mhdds; a normal user account is fine. The account simply needs to belong to the fuse group. You can ensure this by typing:

```
addgroup <User> fuse
```

Now create the new array (Listing 2, line 1); the `-o allow_other` option allows other users to create files.

Additionally, you can specify the `mlimit` parameter, as mentioned before, but options really belong in `/etc/fstab`. Assuming that the mount works, you will see output as shown in Listing 2. All three disks are mounted; all logged-in users have access, and the limit is 4GB. The results, viewed using `df -h`, should look something like Listing 3.

## LISTING 2: Mount Disks to Filesystem

```
$ mhdds /mnt/Movies,/mnt/Photos,/mnt/Music /mnt/media -o allow_other
option: allow_other (1)
mhdds: directory '/mnt/Movies' added to list
mhdds: directory '/mnt/Photos' added to list
mhdds: directory '/mnt/Music' added to list
mhdds: move size limit 4294967296 bytes
mhdds: mount point '/mnt/media'
```

## WHAT IS A UNION FILESYSTEM?

Union filesystems, such as UnionFS, AuFS, OverlayFS, and mhdds, have been growing in popularity in the Linux universe ever since Klaus Knopper released the first version of the Knoppix Live system. These filesystems allow users to mount multiple filesystems on top of one another concurrently. For the user, the whole enchilada looks like a single unit, and on a Live distribution, it lets users install packages without changing the system. OverlayFS has recently seen use with containers and atomic updates.

## LISTING 3: New Filesystem

```
$ df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda1       80G   50G   30G   63% /mnt/Movies
/dev/sdb1       40G   35G    5G   88% /mnt/Photos
/dev/sdc1       60G   10G   50G   17% /mnt/Music
mhdds           180G   95G   85G   53% /mnt/media
```

### LAUNCHING MHDDS WITH SYSTEMD

Mhdds does not come with a service file for Systemd. For this reason, you need to create the file, named `/etc/systemd/system/mnt-media.mount`, then copy the script from Listing 6 to the file. The command

```
systemctl daemon-reload
```

then reloads the file so the service can be started with the `systemctl enable mnt-virtual.mount` command at boot time. You can then type

```
systemctl start mnt-virtual.mount
```

for an automatic start.

As you can see, the software has created the new filesystem; the total capacity is that of the sum total of the individual disks, and the same is true of the free space. The next task is to provide this setup automatically at boot time. To do so, add a new line to your `/etc/fstab` file (Listing 4).

If you do experience problems, it is a good idea to use another option to define where the software creates a logfile and to define the verbosity level for mhdds's output (Listing 5). For more details, refer

to the mhdds man page [5].

If needed, you can add more disks to the array at any time. To do so, unmount the array, restart the software, and add the disks. Then add the mountpoint to your `/etc/fstab` to mount the array automatically.

If you want to stop using the program, remove the line from the `/etc/fstab` file and delete the mountpoint for the array. If you have a distribution that uses Systemd, you can launch mhdds via the init system; the "Launching mhdds with Systemd" box describes this option.

### On the Safe Side

The driver, which is what mhdds is at the end of the day, focuses on a single task in the classic Unix style, and it does its job well. However, it does not offer any kind of

backup in the case of failure. A disk failure in the array will therefore cause loss of data. One drawback in practice is that you do not know where the software will store a new file and thus do not know what data you stand to lose if a disk dies on you.

The only remedy is to back up the data involved. Mhdds is often used in combination with SnapRAID [6] to add a modicum of safety. Beyond this, you can also mirror the array one-to-one. To do so, create a second mhdds instance on your backup disk and synchronize the two instances using Rsync or a similar tool [7].

### Conclusions

The only drawback in using mhdds is the five percent performance hit in data transfer speeds, attributable to FUSE, which you can probably live with. Otherwise, everything is in favor of using mhdds, especially because you have the option of mounting the array via Samba and NFS – this opens up many more options. In this way, you can automate network-attached storage (NAS), video disk recorder (VDR) TV software, or a media center of your choice. Mhdds is not only suitable for multimedia files, it is particularly useful in use cases such as this, in which the size of the collection grows in the course of time. ■■■

#### LISTING 4: Creating the Filesystem at Bootup

```
mhdds#/mnt/Movies,/mnt/Photos,/mnt/Music /mnt/media fuse defaults,
allow_other,mlimit=500M 0 0
```

#### LISTING 5: Mhdds Options

```
mhdds#/mnt/Movies,/mnt/Photos,/mnt/Music /mnt/media fuse defaults,
allow_other,mlimit=500M,logfile=/var/log/mhdds.log,loglevel=1 0 0
```

#### LISTING 6: `/etc/systemd/system/mnt-media.mount`

```
[Unit]
Description = Mount mhdds virtual volume
[Mount]
What = mhdds#/mnt/Movies,/mnt/Photos,/mnt/Music
Where = /mnt/media
Type = fuse
Options = allow_other,mlimit=2G 0 0
[Install]
WantedBy = multi-user.target
```

### INFO

- [1] RAID: <https://en.wikipedia.org/wiki/RAID>
- [2] LVM: [https://en.wikipedia.org/wiki/Logical\\_Volume\\_Manager\\_\(Linux\)](https://en.wikipedia.org/wiki/Logical_Volume_Manager_(Linux))
- [3] Copy-on-write: <https://en.wikipedia.org/wiki/Copy-on-write>
- [4] Source code: <http://svn.uvw.ru/mhdds/trunk/>
- [5] mhdds man page: <http://manpages.ubuntu.com/manpages/karmic/man1/mhdds.1.html>
- [6] SnapRAID: <http://www.snapraid.it/>
- [7] Backup: <http://zornsoftware.codenature.info/blog/why-i-ditched-raid-and-greyhole-for-mhdds.html>

# 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**

The run on Raspberry Pi Zeros is temporary

# Don't Panic

The Raspberry Pi Zero sold out soon after it was released. maddog puts the situation into perspective. *By Jon "maddog" Hall*

Chill! I advised people who were upset because they could not get a Raspberry Pi Zero two days after it was announced. Despite the fact that tens of thousands of them had been produced, the supply (including all the *MagPi* magazines that included them in the issue) had been sold out.

Despair ran rampant on the Internet because (mostly) people who already had a Raspberry Pi A/B/B+ or Raspberry Pi 2 Model B could not get the very latest model.

Chill! I said and advised folks to remember the early days of the first Raspberry Pi when they spent hours trying to find some outlet that would have them in stock. After what seemed like forever, the manufacturing of the Raspberry Pis caught up with the demand, and people could get them whenever they wanted.

In the meantime, some people are selling (or trying to sell) those *MagPi* magazines on eBay for five to ten times what they paid, because they bought up copies knowing that someone would pay more money for the same magazine just to have the Raspberry Pi Zero NOW!

Chill! Because this is not the situation of the original Raspberry Pi, when the designers thought there was only a market for 1000... oh...10,000... oh... 100,000 units, and a couple of years later five million Raspberry Pis had been sold. This time, the Raspberry Pi Foundation must have known that they would sell hundreds of thousands of boards, and they will be prepared to manufacture enough to meet demand. It is just a matter of time.

Other than putting a hook on each of your Raspberry Pis and hanging them on your Christmas tree, what are you going to do with the RPi Zero? A friend of mine said he was interested in

building a Beowulf high-performance cluster with them. "Think how small and cheap that would be," he said. "I can buy eight of them for only \$40!"

True, but then, you'd have to buy some type of networking to allow them to communicate: either a wireless 802.11 device or a USB Ethernet device. These use a fair amount of power, and with only one OTG microUSB connector on the RPi Zero, you probably would have to get a powered hub (and power supply

for the hub). Each RPi Zero in the cluster would also need a Class 10 SD card of (at least) 4GB, which would add

a few dollars to each single-core node of your supercomputer. By the time you have purchased your eight nodes, hubs (and power supplies), SD cards, and networking cards, you are probably looking at \$30 per node.

You might argue that the eight RPi Zeros give you a lot of memory, but that memory comes in half-gigabyte "clumps," and each clump has to load a duplicate of the operating system, networking software, and application code, leaving very little room for data. And, the overhead of message passing uses a lot of CPU, so the RPi Zero cluster will probably not be very efficient.

Now you have a fairly inefficient "supercomputer" that with its hubs, multiple power supplies, cables holding Ethernet controllers (did I mention the 16-port Ethernet switch you need to connect them?) probably costs around a couple hundred dollars to build.

On the other hand, you can purchase an eight-core, 64-bit ARM system with two GB of RAM and wireless 802.11n built-in, running at 1.2 GHz for \$99 [1]. The efficiency of OpenMP or OpenCL to control your parallel threads on a shared-memory model for your application instead of OpenMPI across networking interfaces will be breathtaking (this is not a slight on OpenMPI; every tool has its place).

"I need a Zero!" I can hear the cries coming now. Of course you do, and you will get it.

Almost 50 years ago, the first portable electronic calculators were produced. Displaying through Nixie tubes, they could add, subtract, multiply, and divide, and they cost about \$400. I never bought one of these calculators because I had more powerful computers to do my calculations at school and work. These computers were not portable, but they were much more powerful than the calculators of the day. Eventually, the more portable calculators became so inexpensive that they were given away as prizes for opening accounts at banks and, later still, given away at trade shows. I never bought a "calculator" in my life.

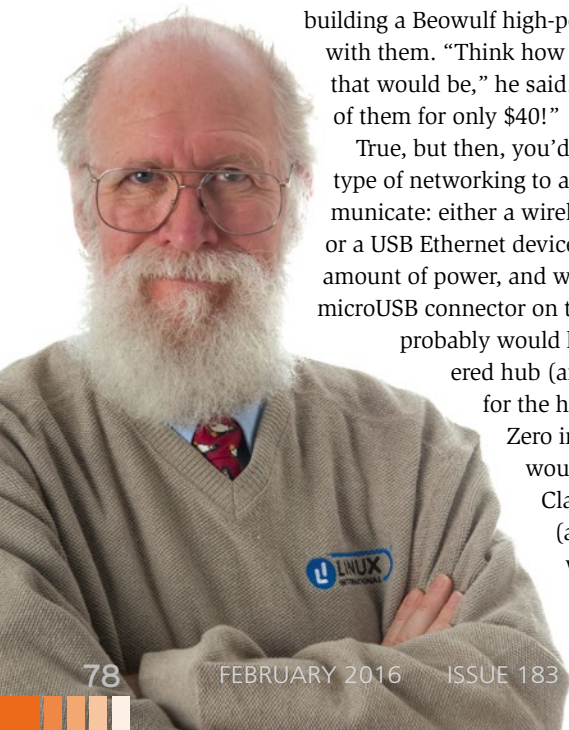
The miracle of the Raspberry Pi Zero is not that you need a Raspberry Pi Zero, but that you will soon be able buy as many Raspberry Pi Zeros as you want for \$5 a piece. ■■■

## INFO

- [1] HiKey (LeMaker version): <http://www.lenovator.com/product/86.html>

## 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®.



Shop the Shop

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

# EXPERT TOUCH




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

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

You'll learn how to:

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

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



**FREE DVD INSIDE!**

The world's greatest Live distro!

**KNOPPIX 7.4.2**

ORDER ONLINE:

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

### Conference coverage

# SUSECon 2015

We provide some highlights from SUSECon 2015.

By Swapnil Bhartiya

**T**his year's SUSECon was held November 2-6 in Amsterdam, the capital of the Netherlands. This was the first SUSECon I've attended, and I was impressed with the magnitude of the event.

The venue was Beurs van Berlage [1], which was designed as a commodity exchange by architect Hendrik Petrus Berlage and constructed between 1896 and 1903. SUSE's creative teams did a magnificent job of SUSE-fying the inside of the building.

### The Event

My perception is that the conference had three focus areas: 1)

ISVs and partners; 2) developers and DevOps; and 3) new talent. SUSE brought many of their big and small partners under the same roof: IBM, Intel, HP, SAP, Fujitsu, etc. IBM came with their LinuxONE mainframe, which was a center of attraction, with many people trying to get their picture taken with this monstrous machine.

The conference also featured hundreds of hands-on workshops and talks, which ranged from topics like OpenStack and Docker to the very basics of Linux. The target audience for these sessions was developers and DevOps teams. Full days were dedicated to certification exams and other training. The conference was a complete package: You had vendors who bring business, developers and admins

who handle the infrastructure of these businesses, and fresh blood.

On the first day, Nils Brauckmann, President & General Manager of SUSE gave a media briefing and talked about the company's growth, engagement with partners, and future plans. I later had a one-on-one interview with Brauckmann, where he explained that the core focus of the company is on its strength – which is the server business – because so many enterprise customers run their own massive data centers.

At the same time, SUSE is also investing resources in “new” technologies, such as software-defined everything and cloud computing. A major share of this investment is going into people. “Thirty-six percent of net additional dollars we can spend, to a large degree, goes into people. No surprise there, because what we do is very people intensive. We actually opened new positions so that, at the end of fiscal year 2016, our staffing will be 26 percent larger than it has been before. We are bringing a lot of new people into the business,” said Brauckmann. SUSE is on a hiring spree, which is good news for Linux experts.

He talked about the innovations SUSE has made in OpenStack. He claimed that SUSE was the first company to bring non-disruptive upgrade capabilities to OpenStack. And, although the company is very strong in OpenStack, they are also spreading their wings. During the conference, SUSE announced that they were joining the Cloud Foundry [2]. SUSE will be primarily focusing on collaborating with SAP on the BOSH OpenStack Cloud Provider Interface (CPI) project to help Cloud Foundry users



### THE AUTHOR

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



more easily run the applications of their choice on public or private clouds.

I spoke with many SUSE executives during the event, and I gathered that SUSE focuses a lot on working with partners, instead of building their own stack, which can create a kind of vendor lock. Brauckmann said there might be other players in the market – both closed source as well as open source – whose strategy is to create their own stack from top to bottom. Then, they go to their customers and say “you have to use our stack.” That creates vendor lock-in for their own stack.

Brauckmann said, “That is not SUSE’s strategy. We do what we do best – our core competencies – and then we engage in partnerships to customize, optimize, to make a solution that is best of breed for our customers to consume. That’s why we are so strong with other partners like SAP and other ISVs.”

That evening, press and others went on a nice canal cruise for dinner, where we dined on a six-course meal.

### Keynotes

The second day, Michael Miller, Vice President of Global Alliances & Marketing, took the stage for the keynote and showed a parody of the Maroon 5 “Sugar” video [3], which set the mood for the rest of the conference. He then invited Brauckmann on stage, and the two discussed the company’s growth in the past year.

Brauckmann said SUSE had another record-breaking year in 2015, as billings grew by more than 25 percent. He stated that really large transactions, more than \$1 million, grew more than 47 percent. He said that this growth further proves SUSE’s ability to support large mission-critical deployments with high quality and high performance and meet enterprise customer requirements.

Ross Mauri, General Manager of z Systems, IBM Corporation, joined Miller on stage and talked about the long relationship between IBM and SUSE around the mainframe. During the keynote, they announced elastic pricing for LinuxONE servers, which allows companies of different sizes to use this powerful machine.

Chiseki Sagawa, Corporate Vice President, Head of Global Software Center at Fujitsu, then took the stage and talked about the decades-long collabora-

tion between the two companies. Sagawa shared the fact that the name Fujitsu has little to do with Mount Fuji. The company was the outcome of a collaboration between Furukawa Electric Company and the German conglomerate Siemens. When Miller asked about his experience at the recent OpenStack event that was organized in Tokyo, Sagawa said it was weird to see so many strange people from outside of Japan wearing jeans and t-shirts in Tokyo. There was a round of applause from the audience, then things got more serious. Sagawa said that they had been in the Linux world for a very long time but had never released their internally developed software as open source. He said that although SUSE is a champion of open source, Fujitsu is a very proprietary company. That’s changing, however. For the first time, Fujitsu is releasing software as open source. This product is the Open Service Catalog Manager, which is used for managing the cloud marketplace.

Sagawa added that the company is making a major shift from proprietary to open source because the speed of innovation is so fast that no single company can keep up with this pace. They have to have open source collaboration to keep up. At the same time, their customers’ dependence on IT infrastructure is increasing, so they have to avoid vendor lock-in, and open source is the answer. Miller asked whether, since the Open Service Catalog Manager is an open source project, could all those people working on it now wear t-shirts and jeans? Sagawa rode the joke and said yes.

Later, Miller was joined by Jürgen Kreuziger, SVP, Cloud & Lifecycle Management at SAP AG to talk about the engagement between the two companies.

After that, SUSE announced the recipients of their first “Always Open” Customer and Partner of the Year awards: S2 Inc., Sicoob, Westinghouse, and Kongsberg Spacotec.



### Conclusion

The end of the keynotes led to the beginning of workshops and talks. The next three days were full of talks, hands-on workshops, and tutorials. I used some of my time to interview key SUSE executives and openSUSE members. I also participated in the launch of openSUSE Leap 42.1.

The closing keynote was dedicated to technical demonstrations and focused on containers and Docker. SUSE engineers gave a demo of JEOS (Just Enough Operating System), a bare minimum system used to deploy containers.

Then, Miller and Alan Clark (Chairman of the Board of the OpenStack Foundation) took the stage and discussed SUSE’s presence in the OpenStack space. Clark said Intel organizes an OpenStack contest called “Rule The Stack,” and SUSE has won all three times. At the end of the conference, they gave out several prizes for the contests that ran throughout the event. And then the beers started coming to the stage, and the conference hall converted into a party hall. ■■■

### INFO

- [1] Beurs van Berlage: [https://en.wikipedia.org/wiki/Beurs\\_van\\_Berlage](https://en.wikipedia.org/wiki/Beurs_van_Berlage)
- [2] SUSE Joins Cloud Foundry: <https://www.suse.com/company/press/2015/suse-joins-cloud-foundry-foundation.html>
- [3] SUSE, yes please: [https://www.youtube.com/watch?v=M9bq\\_alk-sw](https://www.youtube.com/watch?v=M9bq_alk-sw)

Free arts software: Satlisca and Malisca

## Artistic Perspective

Austrian programmer and media artist Michael Aschauer goes on river cruises of a special kind, creating unusual impressions, thanks to free hardware and software. *By Erik Bärwaldt*

River cruises in Europe – particularly on the Rhine and Danube – are growing in popularity. However, unlike tourists who are just looking to relax and let the landscape drift by, Viennese photo and video artist Michael Aschauer succeeds in capturing the cultural landscapes and habitats along major rivers from various perspectives and thus provides distinctive views of these rivers with a high recognition factor.

In doing so, Aschauer – armed with free image editing software, which he developed himself for this purpose – also transforms a rather monotonous, sedate-looking panorama of these rivers into something more exciting by concisely portraying the activity on the river with its different facets, viewed by ship and satellite and sometimes over distances of more than several thousand kilometers.

### The Artist

Michael Aschauer [1], was born in Austria's Styria region in 1977 and has been regarded as a photo and video artist of international repute for more than 15 years. After studying digital art and visual media design in Vienna, he has demonstrated his prowess with numerous works and publications.

Exhibitions in various countries – including Brazil, Great Britain, China, the USA, Bulgaria, Spain, France, Switzerland, and Germany – have established him as part of the international art scene. Several internationally renowned art prizes and various grants from galleries also demonstrate that his artistic achievements are accordingly valued (Figure 1).

However, Aschauer is not just an artist in the original sense of the word but also a virtuoso in the world of bits and bytes. As a Linux professional, he has devel-



Figure 1: One of Michael Aschauer's works at an exhibition in Schattendorf, Austria, in 2010.

oped applications based on free software for several educational and workshop projects over 15 years. In terms of his artistic work, Aschauer is not fazed by the task of writing his own programs if he doesn't find an application to suit his needs. And, of course, his laptop – which accompanies him on all of his trips – is equipped with a Debian derivative and exclusively free software with which he manages his digital universe.

## The Idea

One of the focuses of Michael Aschauer's work is portraying rivers. In doing so, the artist approaches his subject in two different ways: From an external point of view, his works track the flow from a bird's eye view by using satellite images. In contrast, Aschauer also photographs the banks of the river from the vessel navigating it for a view from the inside looking out. He groups these photos in multiple-line panoramic views to create an overall perspective. These projects – “River Studies,” the “Danube Panorama Project,” or the “Nile Studies” – were realized on the world's major rivers over a period of many years.

The artist himself pursues two goals with this experiment: On one hand, he seeks to portray the habitats and cultural space of the river in all its aspects, while on the other, he faces the technical challenge of creating a realistic-looking panoramic view from a series of digital photographs from different perspectives without breaks in the photographic media.

## The Technology

Michael Aschauer welcomed me to his studio in Vienna. On the table in front of him, there was what looked like a fairly clumsy black camera – viewed superficially. It reminded me of a commercially available CCTV surveillance camera from years gone by. In stark contrast, it had a silver, fixed-focus lens from the 1970s affixed – as typically used with analog movie cameras of the period (Figure 2).

Although existing satellite images acted as the source material for creating the river panoramas from a bird's eye view, Aschauer used his special camera by Elphel to record the river from on board a boat. He explained: “These cameras by Elphel are not only characterized by their consistent use of open hardware components, they also use embedded

Linux as their operating system. This gives you various options for adapting the camera to suit your individual needs, because it contains programmable logic components.”

The Elphel 353 [2] used by Aschauer with its five-megapixel sensor – in combination with a fixed-focus lens (with a focal length of 8 or 16mm) – generates customized video images in an extremely wide format of 2592x48 pixels. The frame rate of the software-based line scan on the computer can be up to 400 frames per second. Aschauer has to adjust this frame rate manually for his recordings, taking various factors into consideration in his computations, such as the speed at which the vessel is traveling or the distance to the riverbank.

Alternatively, it would also be possible to allow the camera to use its internal line scan method – in this case, the camera sends a continuous data stream of frames captured line by line to the computer, and the computer generates an image from them. The Aptina sensor in the Elphel 353 achieves a rate of up to 2,300 lines per second. However, the camera's internal line scan mode turned out to be not particularly stable, as the artist explained. Aschauer thus opted for the slower but more precise method of using software computations on the computer.

To ensure a uniform image, it is essential to adjust the camera's frame rate permanently. Because this is a manual process, Aschauer also needs a powerful preview function: “To be able to read the RTSP [real-time streaming protocol] data stream from the camera, display a small thumbnail of this, and then save the stream as a video file, I programmed a special software tool that relies on GStreamer as its back end. In terms of metadata, the software also stores the GPS data in the file.” However, the artist wasn't satisfied with this. He said, “To add more convenience, I also created a small controller from an Arduino minicom-



Figure 2: The camera equipment, including a fixed-focus lens.



Figure 3: The hardware (from left to right): GPS receiver, bag with batteries, laptop, and camera.

puter and a MIDI controller that I had taken apart – this made controlling the camera recording far easier.” Figure 3 shows what this looks like.

### The Software

Aschauer put a huge amount of time and work into developing a program capable of computing the panoramic views; at the time, there simply wasn't an applica-

tion on the market suitable for his purposes. In particular, evaluating the satellite images caused problems because legacy programs were unable to create a panoramic view in the form of a uniform river course from them.

This prompted Aschauer to develop two programs, both of which he released under the GPL and published on the GitHub platform: Satlisca [3] and Malisca [4]. A historic form of representation provided the impetus for the development of Satlisca, which is designed for working with satellite images. In historic literature on cartography dating back to the first half of the 19th century, the artist found several comparative representations of the courses of rivers that visualized the river in a symbolic way in strips (Figure 4).

Based on this model, and using publicly accessible images from the NASA Landsat 7 program and OpenStreetMap, Aschauer developed the Satlisca application, which uses various parameters and line-scanning methods to create an image with a geographically accurately rotated image of the course of the river in question as a straight line. The artist used this approach to create satellite panoramas of rivers, including the Nile, Amazon, Yangtze, Ob, Yellow, Mekong, Ganges, Brahmaputra, Euphrates, Rhine, and Danube. Aschauer then plotted hard copies of these unusual views of the major rivers, which he exhibited in a 120x250cm format in an exhibition titled “What If You Would Pull Rivers To A Straight Line?” (sic) in 2010 and 2011.

Malisca, which is written in the C programming language, also acts as a line scanner, but it's designed for use with a movie camera. Malisca relies on the GStreamer sources, OpenGL, and various GPS libraries for this. Thanks to a small Python front end called Camcontrol, the artist was able to manage the Elphel 353 camera precisely and thus create a uniform image of the river bank independent of the actual geographic conditions. Armed with this equipment, Aschauer created panoramas of the Danube, Nile, Ganges, Brahmaputra, and Amazon rivers (Figure 5).

### Animation

Aschauer, however, was not satisfied with the photographic renderings of the river courses to propagate his own exper-



Figure 4: A length comparison of long rivers dating back to 1817. Published with the kind approval of David Rumsey, Cartography Associates, San Francisco [5].

imental style of cartography. In a further step, he decided to transcend the boundaries between static documentation and realistic-looking animation. To do so, the video artist generated a movie from the existing static material using DCP Cinemascope; it presents the five rivers with their different characteristics in short summaries of 10 minutes each. The 51-minute movie created this way was first aired in Vienna in May 2013 [6].

Beyond this project, the programming artist also developed a web player based on the GeoDjango framework and the OpenLayers library to let viewers experience the river panorama in a realistic way in a web browser. Playing back the panoramas gives viewers the impression of traveling on a boat along the river themselves. The River View Player [7], of which version 2.0 beta now exists, also gives the viewer the ability to change the boat's speed (i.e., the speed of the animation) and set the zoom factor (Figure 6). Additionally, a topographical map based on OpenStreetMap is available to help viewers trace the route of the "river cruise" (Figure 7).

## Conclusions

Michael Aschauer's river studies impressively demonstrate the options offered to creative minds by free and open source software, along with hardware that offers transparent documentation. Aschauer has succeeded in transcending the borders between legacy and experimental geography, while also creating a totally new visualization approach that combines static and dynamic animated elements.

Thanks to his profound knowledge as a software developer, Aschauer was able to write his own software tools to suit the special requirements of his own style of photographic cartography. Of course, he released these programs under the GPL, thus making them freely available to anyone interested in using them. It is, not least, this willingness to contribute to the free software universe that explains its amazing development dynamic that has already outpaced proprietary solutions. ■■■

### INFO

- [1] The artist's website:  
<http://m.ash.to/en/Main/News?from=Main.Index>
- [2] Info on the camera:  
<http://www3.elphel.com/nc353>
- [3] Satlisca:  
<https://github.com/backface/satlisca>
- [4] Malisca:  
<https://github.com/backface/malisca>
- [5] Historic map material:  
<http://www.davidrumsey.com/>
- [6] "River Studies" (film):  
[http://m.ash.to/de/Projects/RiverStudies\\_Cinemascope](http://m.ash.to/de/Projects/RiverStudies_Cinemascope)
- [7] River View Player:  
<http://play.riverstudies.org/>




Figure 5: The Amazon offers a uniform panorama over long distances

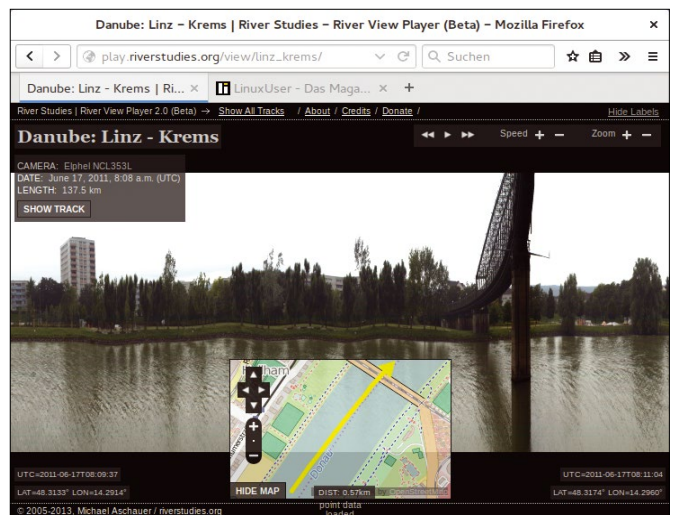


Figure 6: The River View Player by artist Michael Aschauer takes viewers on a boat trip.

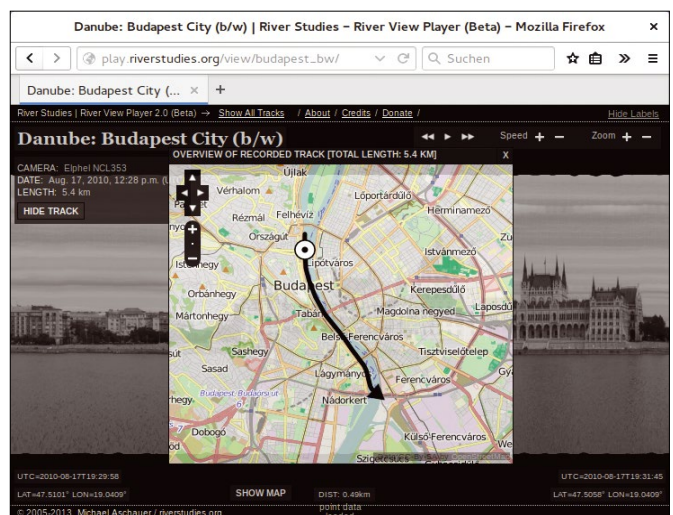


Figure 7: Thanks to OpenStreetMap, the River View Player shows the viewer the current location.

# 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.

4

RASPBERRY PI GEEK

ISSUE 10



# 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!



A stable and well-tested rolling release

# Rollin' rollin' rollin'

We talked with Richard Brown at SUSECon about openSUSE Tumbleweed and the rolling release process. *By Swapnil Bhartiya*

There are two types of users: 1) Those who want a very stable, rock solid environment where they can run their applications without worrying about breaking things; and 2) Those who want the latest technologies to play with. openSUSE is now targeting both use cases with rock solid Leap and rolling release Tumbleweed. At SUSECon, I sat down with Richard Brown, who has played an instrumental role in openSUSE Tumbleweed. Brown chairs the openSUSE Board and recently joined SUSE (Figure 1).

## The Brief History of Tumbleweed

The name Tumbleweed has been around for a very long time; kernel developer Greg Kroah-Hartman started

the project back in 2010. Kroah-Hartman called it “a repo that is a rolling updated version of openSUSE, containing the latest ‘stable’ versions of packages for people to use.” It wasn’t a rolling release distribution in the true sense; in essence, it was actually rolling updates on top of the stable releases of openSUSE.

“Since the Tumbleweed project lacked manpower, Greg ended up doing all the work. I tried Tumbleweed a couple of times and things would break after a while, and it would take Greg to go and fix them,” Brown said.

Tumbleweed enjoyed some popularity initially, but because the entire stack was not moving forward, the user interest faded and the user-base declined.

Brown had been championing the idea of a rolling release, though. He used to talk to Kroah-Hartman about it but struggled with Kroah-Hartman’s vision of rolling updates. In the end, most conversations ended in typical open source fashion: If you want this change, go ahead and do it.

And, that’s what happened in the end. The openSUSE community went ahead and created a rolling release. But, it was not Kroah-Hartman’s Tumbleweed; the new Tumbleweed was created from the Factory code branch (Figure 2).

## The Mutation

“To talk about what is Tumbleweed today, we have to talk about Factory actually,” said Brown. Factory was the

### THE AUTHOR

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



Figure 1: Richard Brown at SUSECon.





Figure 2: openSUSE Tumbleweed.

development code base for openSUSE for over a decade. It was similar to Debian unstable or Fedora Rawhide: Everything went there without any testing. Some brave users would use it; it would work for a while, and then occasionally it would spontaneously combust.

A while ago, openSUSE developers started playing with the Factory development branch, trying to find ways in which they could develop it in a more reliable way. Their work led to the creation of OpenQA, a tool that can automatically run an operating system through rigorous testing.

When you are testing a Linux-based operating system, you are looking at hundreds of different scenarios: how and where a user may install it, what desktop environments will be used, how different components or the entire stack will be updated. Testing all those scenarios is a herculean task if done manually, especially when dealing with a continuously moving target like a rolling release. OpenQA was developed to automate the testing process. OpenSUSE developers soon realized that, with OpenQA, they had inadvertently created a true rolling release distribution. So, for a brief period, openSUSE had two rolling release distributions: Factory and Tumbleweed.

Almost one year ago, the openSUSE community decided that it made no sense to have two rolling release distributions. It was causing user confusion as well as duplication for developers. They decided to keep only one rolling release, which was based on Factory. The idea of openSUSE rolling release came from

Kroah-Hartman, though, so they wanted to give him credit. They asked his permission to use the name Tumbleweed for this rolling release.

Brown recalled that Kroah-Hartman was really happy about the idea because this was what he had envisioned; this is what he would have done himself, if maintaining his own version had not been too much work for one person. Thus Factory became the place where all crude development work happens, and Tumbleweed became the rolling distribution associated with the Factory development process. There is no longer a distribution called Factory.

## How Stuff Works

Open Build Service (OBS) and OpenQA are the key components of Tumbleweed. The developers have a very simple workflow; everything from Gnome, KDE, Apache, etc. goes to OBS. Package maintainers/developers self-organize into groups looking after a particular stack or group of packages. OpenSUSE developers create a project in the build service called a “devel project,” and it’s those devel project submissions that are accepted into staging and then the Factory code base.

“Factory can’t be used by any human being, but we have a process called staging where we get these incoming pull requests, and then we use the power of OBS to make what in essence is a distribution,” said Brown.

OBS basically compiles an openSUSE distribution from these packages for the sole purpose of testing. According to Brown, “it’s a special spin of openSUSE Tumbleweed, where we have

everything that previously worked in Tumbleweed plus these new experimental changes. So, we create this special distribution where changes have been staged, for the sole purpose of testing. Then, we throw this distribution at our OpenQA, which will actually check it quite extensively, but in a very quick manner. It will check if we could accept the new kernel, new GCC change, or whatever without breaking the distribution. If anything breaks, then those submissions can't get into Tumbleweed."

The moment something breaks, OpenQA sends a message to the developer informing him why the package can't be accepted. The developer can fix it and submit it again.

Once everything is staged, checked, and validated, it goes into Factory proper, where everything goes through extensive testing. Factory proper is the actual project called Factory, where everything is integrated and tested together. Stagings only include a subset of changes, grouped together normally based on which devel project they come from.

For example, suppose both KDE and Gnome want to update at the same time. In that case, KDE's proposed changes will be put into a staging project, let's call it Staging:K, and Gnome's proposed changes will be put into a different one: Staging:G. In each of those staging projects, only the KDE or the Gnome changes exist, so openQA can confirm that, on their own, the changes don't break anything. Thus, these changes can be accepted into Factory.

Factory is where everything mixes together, however, so it's entirely possi-

ble that new Gnome might break new KDE, or vice versa, so it's in Factory where the more in-depth testing takes place (Figure 3).

OpenQA plays an important role, because changes made to one package might not break something, but when other changes are made to different packages and put together, something else might break. That's what trips up other distributions at the last moment. Other distributions either test things manually "or just cross their fingers that it won't break. Take, for example, the Arch approach, where the philosophy is very much that 'you should be doing this to learn how to fix it yourself'," said Brown.

He added that the openSUSE project considers it extremely important to create a stable and reliable operating system. "We try to make a rolling release that is easier to live with. We have a responsibility as a distribution house. We are still producing a distribution, and we want it to work for us as well because most of us are running Tumbleweed."

Seeing the capabilities of OpenQA, even Fedora and SUSE have started using it for internal testing.

### Time to Go to Press

Not everything is automated. "There is some human oversight, as OpenQA may show errors that we know are acceptable. We don't support upgrades from openSUSE 13.2 to Tumbleweed, for example, but run the check through OpenQA to see how badly is it broken. So, OpenQA will fail that test but we know it's not a deal breaker," said Brown. In cases like these, the developers will go ahead with the distribution. However, even then critical issues can appear, and the release won't happen unless a human has looked at it and fixed it. "Once everything is ready and green, we see the power of OBS, which publishes the images magically," said Brown.

### Conclusion

I have been using the new Tumbleweed for a while, and now I know what makes it so stable. Brown actually joked that the terms *stable* and *rolling release* are contradictory, but openSUSE Tumbleweed has succeeded in combining the concepts. ■■■

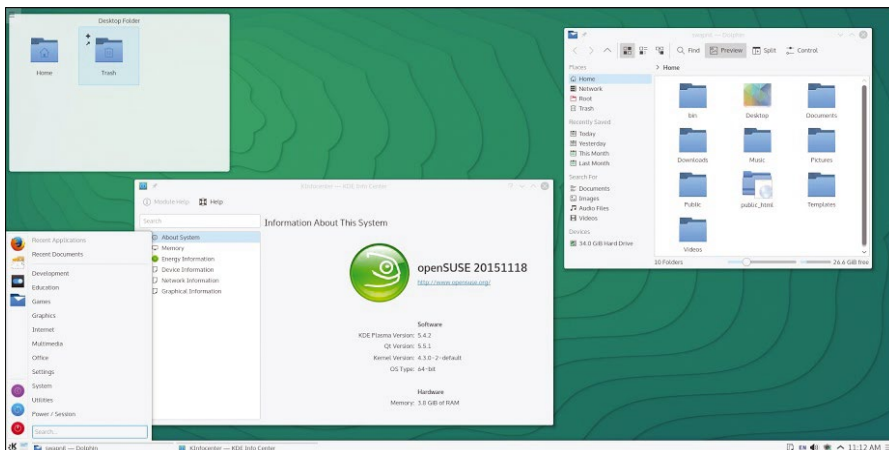


Figure 3: Tumbleweed system info.

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!



# Zack's Kernel News

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

By Zack Brown

## Incremental Updates to the OOM Killer

Michal Hocko gathered up a few ideas from Mel Gorman and Oleg Nesterov about how to improve the OOM (out-of-memory) killer. This is a bizarre little corner of the kernel whose job is to decide which processes to kill when RAM has become so overused that it threatens to lock up the whole system. The idea is that if it chooses correctly, the OOM killer can restore the system to usability. Of course, if it chooses incorrectly, nobody's happy.

As Michal put it, "The OOM killer currently allows to kill only a single task in a good hope that the task will terminate in a reasonable time and frees up its memory. Such a task (oom victim) will get an access to memory reserves via `mark_oom_victim` to allow a forward progress should there be a need for additional memory during exit path."

However, a variety of folks, such as Tetsuo Handa, had shown that under some workloads the oom victim could get trapped in the D state and never exit, thus holding onto the extra memory reserves indefinitely.

Michal posted a patch to create a new kernel thread called `oom_reaper` that would reclaim that memory if possible. The problem with this, Michal noted, was that if the `oom_victim` was in the process of core dumping and didn't have enough memory to complete the job, it would terminate without producing a proper debuggable dump. He justified this however by saying that "the overall system health is more important than debuggability of a particular application."

Typically, the kernel developers would try to avoid creating a whole new kernel thread because the system can only tolerate so many threads at a time, and the more kernel threads there are, the fewer user threads there can be. Michal felt that it was necessary in this case because an out-of-memory condition naturally makes normal execution paths less reliable – that's the whole point of having an OOM killer. An independent kernel thread would have the best chance of actually getting to run.

Johannes Weiner felt that Michal put too much emphasis on making the most common cases run smoothly. He agreed that this would be necessary, but he felt that the big problem was how to resolve memory deadlocks. There had to be some way for the OOM killer to switch from trying to kill one process to trying

to kill another – or perhaps even all – running processes in order to keep the kernel itself up. Before solving that problem, any other consideration seemed cosmetic to him.

However Michal said he wasn't trying to solve that issue; he just wanted to make one troubling class of deadlocks and corner cases go away. He said there was never going to be a truly correct OOM killer implementation because it was always going to be a mishmash of heuristics and best guesses. Instead of trying to solve the problem in the code, Michal thought it would be better to allow the system administrator to choose their own policy via configuration files. That way, some amount of customization could be based on the use case.

Tetsuo agreed with Johannes, though, saying that "the last resort solution has higher priority than smoothening the common case." He added, "Please do consider offering the last resort solution first. That will help reducing unexplained hangup/reboot troubles."

But, Michal really wanted to stay on topic. He said that he didn't have a last resort solution to offer and that all the last resort proposals he'd seen had their own difficult problems to solve. He was offering an incremental improvement to the current OOM killer, and, he said, "I really do not want to make this thread yet another mess of unrelated topics."

There followed a technical discussion of Michal's patch with Tetsuo and others, with only minor deviations into larger problems that he didn't want to address. He reined those in, and at one point, Mel Gorman gave his thumbs up for the patch, offering a couple of minor suggestions. The technical discussion continued for a bit, with Michal posting fixes and updates in response to comments.

The OOM killer is really one of the thorny roses of the Linux kernel. It's a lovely, beautiful thing to have until you actually try to touch it, and then you notice the blood and stinging pain. Michal's effort to clamp down on any broader discussion of overall OOM killer behavior shows how difficult it can be to make even minor changes to that part of the kernel. Everyone wants the big magical fix: the perfect rose that has no thorns.

## Locks Are Hard; Watch Out!

Linus Torvalds had some advice for anyone writing kernel code: Be careful with locks!

## ZACK BROWN

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

Some folks were doing some work with kthreads, and they needed to use some locking, so they posted this as part of their patch:

```
while (!trylock(worker)) {
    if (work->canceling)
        return;
    cpu_relax();
}
queue;
unlock(worker);
```

Linus immediately said:

*People, you need to learn that code like the above is **\*not\*** acceptable. It's busy-looping on a spinlock, and constantly trying to **\*write\*** to the spinlock.*

*It will literally crater performance on a multi-socket SMP system if it ever triggers. We're talking 10x slowdowns, and absolutely unacceptable cache coherency traffic.*

*These kinds of loops absolutely **\*have\*** to have the read-only part. The 'cpu\_relax()' above needs to be a loop that just tests the lock state by **\*reading\*** it, so the cpu\_relax() needs to be replaced with something like*

```
while (spin_is_locked(lock)) cpu_relax();
```

*instead (possibly just "spin\_unlock\_wait()") – but the explicit loop might be worth it if you then want to check the "canceling" flag independently of the lock state too).*

*In general, it's very dangerous to try to cook up your own locking rules. People **\*always\*** get it wrong.*

Linus went on to say:

*... people need to realize that locking is harder than they think, and not cook up their own lock primitives using things like trylock without really thinking about it a **\*lot\***.*

*Basically, 'trylock()' on its own should never be used in a loop. The main use for trylock should be one of:*

- [some]thing that you can just not do at all if you can't get the lock.
- avoiding ABBA deadlocks: if you have an A->B locking order, but you already hold B, instead of 'drop B, then take A and B in the right order', you may decide to first 'trylock(A)' – and if that fails you then fall back on the 'drop and relock in the right order'.

*But if what you want to create is a 'get lock using trylock', you need to be very aware of the cache coherency traffic issue at least.*

*It is possible that we should think about trying to introduce a new primitive for that 'loop\_try\_lock()' thing. But it's probably not common enough to be worth it – we've had this issue before, but I think it's a 'once every couple of years' kind of thing rather than anything that we need to worry about.*

*The 'locking is hard' issue is very real, though. We've traditionally had a **\*lot\*** of code that tried to do its own locking, and not getting the memory ordering right etc. Things that happen to work on x86 but don't on other architectures etc.*

## POSIX Compliance

Linus Torvalds and Al Viro had an interesting exchange recently about POSIX and practicality. Al had noticed that the behavior of the readlink() system call seemed wrong in certain places, and he wanted to fix them.

The readlink() system call returns the name of the thing pointed to by a symbolic link. One problem Al noticed was that calls to readlink() would update the



`atime` value – the time of the most recent access – even in the event of an input/output error or an out-of-memory error. Typically, depending on the kind of access, Linux might update the `atime` of an underlying filesystem object (the file, directory, or whatnot), or the symlink itself. But, regardless, according to POSIX, `readlink()` should not update the `atime` after either of those errors.

Linus replied, “I really don’t think anybody cares, but I also don’t think anybody cares about the current behavior, so we can certainly fix it to match POSIX wording.”

Another problem Al noticed was that `readlink()` was supposed to fail when used on anything other than a symlink. And, according to POSIX, it should fail with an `EINVAL` error (invalid argument). But, Al found a weird and unlikely case where `readlink()` could succeed when run on a directory. The Andrews File System (AFS) is a distributed filesystem that automounts directories as the user attempts to go into them. Those directories look like directories to the `stat()` system call. But, instead of using the `open()` system call, AFS uses `readlink()` to find out what it should open.

Al thought that the whole AFS approach was a kludge that should be fixed, but Linus came back with his position on POSIX compliance in general:

*I don’t think POSIX is necessarily relevant here.*

*We have had magic file behavior outside the scope of POSIX before, and we will have it in the future. It makes perfect sense to use `readlink()` for management tools for automounting, even if the normal operation is to treat the thing as a directory.*

*Not everything is within the domain of POSIX.*

Al pointed out that since the AFS directories could only be opened via `readlink()`, while appearing to `stat()` to be a directory, there was no way for user code to know that it should open those directories with `readlink()` instead of the usual `open()`.

Linus replied that only the AFS management tools needed to know about it, and they already did because it was their own filesystem. Linus added:

*Not everything has to be “generic”. Sometimes it’s good enough to just have the ability to get the work done. Now, if it turns out that others also want to do this,*

*maybe somebody decides “let’s add flag `-V` to ‘ls’, which forces a ‘`readlink()`’ on all the targets, whether links or not, and shows the information”.*

*I could imagine other special files having “a single line of information about the file” that they’d expose with `readlink()`. Who knows?*

*So there is \*potential\* for just making it generic, but that doesn’t mean that it necessarily has to act that way.*

Linus went on to say:

*... it’s not necessarily just `readlink()` either. I still think it might be a perfectly fine idea to allow non-directories to act as directories in some case (by exposing “`readdir`” and “`lookup`”).*

*But `readdir()` really doesn’t sound horrible either. How about unix domain sockets (or named pipes) giving their link information when you do `readdir()` on them?*

*Quite frankly, I think allowing those kinds of unified interfaces is better than the current situation where you have to use a “`getpeername()`” system call etc. If it’s a filesystem object, why not allow filesystem operations to work on it?*

*We expose some things in `/proc` as symlinks, things that actually would work better as non-symlinks, exactly \*because\* we want to expose not just the end result of what they point to, but also a \*description\* of what they point to. So we have those odd “pseudo-symlinks” in `/proc` that don’t actually really do a pathname walk on the symlink content they expose, but still \*look\* like symlinks just because `readdir()` is such a useful thing to have.*

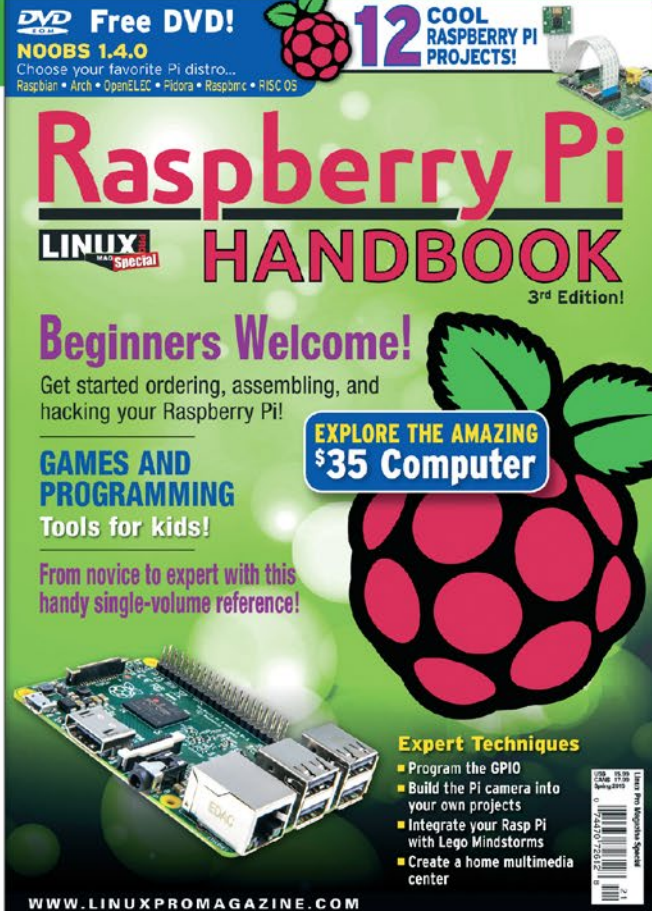
Al cautioned against exposing `lookup()` for non-directories, which he said would cause a major nightmare with locking code. He added, “The situation is convoluted enough as it is; playing with parallel lookups is going to be interesting in itself and I’d rather not mix it with attempts to accommodate for hybrid objects”

At that point, the discussion veered back to a technical consideration of AFS behaviors, but it’s interesting to see the two of them go at it. Al basically knows everything, while Linus has a clearer grasp of big-picture issues. Having said that, though, Al has made it clear in the past that if certain big-picture issues go the wrong way, he’d fork off his own “VirAl” kernel and take it in his own direction. He probably has the clout to bring a lot of developers with him, but so far it’s never come to that. ■■■

Shop the Shop

shop.linuxnewmedia.com

# Raspberry Pi HANDBOOK



**DVD** Free DVD!  
**NOOBS 1.4.0**  
Choose your favorite Pi distro...  
Raspbian • Arch • OpenELEC • Pidora • Raspbmc • RISC OS

**12 COOL RASPBERRY PI PROJECTS!**

## Raspberry Pi HANDBOOK

LINUX Special  
3rd Edition!

### Beginners Welcome!

Get started ordering, assembling, and hacking your Raspberry Pi!

**GAMES AND PROGRAMMING**  
Tools for kids!

From novice to expert with this handy single-volume reference!

**EXPLORE THE AMAZING \$35 Computer**

**Expert Techniques**

- Program the GPIO
- Build the Pi camera into your own projects
- Integrate your Rasp Pi with Lego Mindstorms
- Create a home multimedia center

WWW.LINUXPROMAGAZINE.COM

*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

# 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).



### Embedded Linux Conference

**Date:** April 4–6, 2016

**Location:** San Diego, California

**Website:** <http://events.linuxfoundation.org/events/embedded-linux-conference/>

ELC, the technical conference for companies and developers using Linux in embedded products, presents the 12th year of sessions dedicated exclusively to embedded Linux and embedded Linux developers.

### Apache: Big Data

**Date:** May 9–11, 2016

**Location:** Vancouver, British Columbia

**Website:** <http://events.linuxfoundation.org/events/apache-big-data-north-america>

Apache projects are the foundation of many Big Data platforms. Join other professionals working in Big Data, ubiquitous computing, and data engineering and science to accelerate the state of the art.

### ApacheCon North America

**Date:** May 11–13, 2016

**Location:** Vancouver, British Columbia

**Website:** <http://events.linuxfoundation.org/events/apachecon-core-north-america>

Join the open source community to learn about and collaborate on the technologies and projects driving the future of open source, web technologies, and cloud computing.

## EVENTS

<b>FAST '16</b>	February 22–25	Santa Clara, California	<a href="https://www.usenix.org/conference/fast16">https://www.usenix.org/conference/fast16</a>
<b>Mobile World Congress</b>	February 22–25	Barcelona, Spain	<a href="http://www.mobileworldcongress.com/">http://www.mobileworldcongress.com/</a>
<b>Icinga Camp</b>	March 1	Berlin, Germany	<a href="https://www.icinga.org/community/events/icinga-camp-berlin/">https://www.icinga.org/community/events/icinga-camp-berlin/</a>
<b>Open Networking Summit Cebit</b>	March 14–17	Hannover, Germany	<a href="http://opennetsummit.org/">http://opennetsummit.org/</a>
<b>NSDI '16</b>	March 16–18	Santa Clara, California	<a href="https://www.usenix.org/conference/nsdi16">https://www.usenix.org/conference/nsdi16</a>
<b>Chemnitzer Linux-Tage</b>	March 19–20	Lake Tahoe, California	<a href="https://chemnitzer.linux-tage.de/2016/en">https://chemnitzer.linux-tage.de/2016/en</a>
<b>Collaboration Summit</b>	March 29–31	Lake Tahoe, California	<a href="http://events.linuxfoundation.org/events/collaboration-summit">http://events.linuxfoundation.org/events/collaboration-summit</a>
<b>Embedded Linux Conference</b>	April 4–6	San Diego, California	<a href="http://events.linuxfoundation.org/events/embedded-linux-conference">http://events.linuxfoundation.org/events/embedded-linux-conference</a>
<b>Cloud Expo Europe</b>	April 12–13	London, England	<a href="http://www.cloudexpo-europe.com/">http://www.cloudexpo-europe.com/</a>
<b>Linux Storage Filesystem and MM Summit</b>	April 18–19	Raleigh, North Carolina	<a href="http://events.linuxfoundation.org/events/linux-storage-filesystem-and-mm-summit">http://events.linuxfoundation.org/events/linux-storage-filesystem-and-mm-summit</a>
<b>Vault Linux Storage and Filesystems Conference</b>	April 20–21	Raleigh, North Carolina	<a href="http://events.linuxfoundation.org/events/vault">http://events.linuxfoundation.org/events/vault</a>
<b>Open Source Data Center Conference</b>	April 26–28	Berlin, Germany	<a href="https://www.netways.de/en/events_trainings/osdc/overview/">https://www.netways.de/en/events_trainings/osdc/overview/</a>
<b>Grazer Linxstage 2016</b>	April 29–30	Graz, Austria	<a href="https://www.linuxstage.at/">https://www.linuxstage.at/</a>
<b>Re:publica</b>	May 2–4	Berlin, Germany	<a href="https://re-publica.de/">https://re-publica.de/</a>
<b>Open Tech Summit</b>	May 5	Berlin, Germany	<a href="http://opentechsummit.net/">http://opentechsummit.net/</a>
<b>Apache Big Data North America</b>	May 9–11	Vancouver, BC, Canada	<a href="http://events.linuxfoundation.org/events/apache-big-data-north-america">http://events.linuxfoundation.org/events/apache-big-data-north-america</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	82
Swapnil Bhartiya	80, 88
Chris Binnie	16
Zack Brown	92
Bruce Byfield	56, 60
Joe Casad	3, 8
Mela Eckenfels	70
Tobias Eggendorfer	12
Jon "maddog" Hall	78
Mathias Huber	32
Heike Jurzik	70
Klaus Knopper	38
Charly Kühnast	37
Hartmut Noack	40
Dmitri Popov	64
Andreas Reitmaier	52
Mike Schilli	48
Ferdinand Thommes	74
Uwe Vollbracht	22
Markus Winand	26

## 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

Joe Casad

### Copy Editor

Amber Ankerholz

### Layout

Dena Friesen, Lori White

### Cover Design

Lori White

### Cover Image

© stylephotographs & Maxim Borovkov, 123RF.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

Darrah Buren, [dburen@linuxnewmedia.com](mailto:dburen@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.

Microsoft, Windows, and the Windows logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

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 184 / March 2016

# Systemd

Approximate	
UK / Europe	Feb 06
USA / Canada	Mar 04
Australia	Apr 04

On Sale Date	
--------------	--

Most Linux users know Systemd is gradually replacing the venerable sysVinit service as the leading Linux startup tool, but the daemon is in the details. Next month, we explore the practical side of working with Systemd on real-world Linux systems.

## 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 © Sakonboon Sansri, 123RF.com

# START

## Enigma 2016

January 25–27, 2016, San Francisco, CA, USA  
[enigma.usenix.org](http://enigma.usenix.org)

## FAST '16: 14th USENIX Conference on File and Storage Technologies

February 22–25, 2016, Santa Clara, CA, USA  
[www.usenix.org/fast16](http://www.usenix.org/fast16)

## NSDI '16: 13th USENIX Symposium on Networked Systems Design and Implementation

March 16–18, 2016, Santa Clara, CA, USA  
[www.usenix.org/nsdi16](http://www.usenix.org/nsdi16)

Co-located with NSDI '16

CoolDC '16: USENIX Workshop on Cool Topics on Sustainable Data Centers  
March 19, 2016  
[www.usenix.org/cooldc16](http://www.usenix.org/cooldc16)

## SREcon16

April 7–8, 2016, Santa Clara, CA, USA

## USENIX ATC '16: 2016 USENIX Annual Technical Conference

June 22–24, 2016, Denver, CO, USA  
Submissions due February 1, 2016  
[www.usenix.org/atc16](http://www.usenix.org/atc16)

Co-located with USENIX ATC '16:

HotCloud '16: 8th USENIX Workshop on Hot Topics in Cloud Computing  
June 20–21, 2016  
Submissions due March 8, 2016  
[www.usenix.org/hotcloud16](http://www.usenix.org/hotcloud16)

HotStorage '16: 8th USENIX Workshop on Hot Topics in Storage and File Systems  
June 20–21, 2016  
Submissions due March 10, 2016  
[www.usenix.org/hotstorage16](http://www.usenix.org/hotstorage16)

SOUPS 2016: Twelfth Symposium on Usable Privacy and Security  
June 22–24, 2016  
Paper registration due March 1, 2016  
[www.usenix.org/soups2016](http://www.usenix.org/soups2016)

## SREcon16 Europe

July 11–13, 2016, Dublin, Ireland

## USENIX Security '16: 25th USENIX Security Symposium

August 10–12, 2016, Austin, TX, USA  
Submissions due February 18, 2016  
[www.usenix.org/sec16](http://www.usenix.org/sec16)

Co-located with USENIX Security '16

WOOT '16: 10th USENIX Workshop on Offensive Technologies  
August 8–9, 2016

CSET '16: 9th Workshop on Cyber Security Experimentation and Test  
August 8, 2016

ASE '16: 2016 USENIX Workshop on Advances in Security Education  
August 9, 2016

HotSec '16: 2016 USENIX Summit on Hot Topics in Security  
August 9, 2016  
[www.usenix.org/hotsec16](http://www.usenix.org/hotsec16)

## OSDI '16: 12th USENIX Symposium on Operating Systems Design and Implementation

November 2–4, 2016, Savannah, GA, USA  
Abstracts due May 3, 2016  
[www.usenix.org/osdi16](http://www.usenix.org/osdi16)

## LISA16

December 4–9, 2016, Boston, MA, USA  
[www.usenix.org/lisa16](http://www.usenix.org/lisa16)

### USENIX Open Access Policy

USENIX is the first computing association to offer free and open access to all of our conferences proceedings and videos. We stand by our mission to foster excellence and innovation while supporting research with a practical bias. Your membership fees play a major role in making this endeavor successful. Please help us support open access. Renew your USENIX membership and ask your colleagues to join or renew today!

[www.usenix.org/membership](http://www.usenix.org/membership)





# FatTwin™

## Evolutionary 4U Twin Architecture

4/8-Node w/ Front I/O  
Highest Density & Efficiency

4-Node w/ Front I/O  
Highest Power Efficiency

4/8-Node w/ Rear I/O  
Highest Storage Density & Efficiency  
8 hot-swap 3.5" drive bays in 1U

GPU/Intel® Xeon Phi™  
w/ Front I/O



### Eight, Four or Two hot-pluggable Servers (Nodes) in a 4U form factor. Each Node supports up to:

- Up to 36 Cores per node and 145W TDP dual Intel® Xeon® Processor E5-2600 v3 product family
- 1TB DDR4-2133MHz memory in 16 DIMM slots
- 1 PCI-E 3.0 x16, and 1 MicroLP PCI-E 3.0 x8 slots (Rear I/O models)
- 8 SAS 3.0 (12Gbps) ports with LSI® 3108/ 3008 controller
- 10 SATA 3.0 (6Gbps) ports with Intel® C612 controller, RAID 0,1,5,10
- 2 NVMe ports
- Dual 10GBase-T or Dual Gigabit Ethernet LAN options
- Redundant Titanium (96%+)/Platinum (95%+) Level Digital power supplies
- Integrated IPMI 2.0 plus KVM with dedicated LAN



Learn more at [www.supernmicro.com/FatTwin](http://www.supernmicro.com/FatTwin)



© Super Micro Computer, Inc. Specifications subject to change without notice. Intel, the Intel logo, the Intel Inside logo, Xeon, and Intel Xeon Phi are trademarks of Intel Corporation in the U.S. and/or other countries. All other brands and names are the property of their respective owners.