FREE DVD

ubuntu 16.04 Desktop (64-bit) & ubuntu MATE 16.04 Desktop (32-bit)

CC BY-SA 4.0

LOG CHECK

Monitor your systems with a central log server

ubuntu MATE 16.04 Desktop (32-bit)
ubuntu 16.04 Desktop (64-bit)
LINUX
DVD
JUST RELEASED!

LINUX PRO MAGAZINE

JULY 2016

LOG CHECK

Monitoring your systems with a central log server

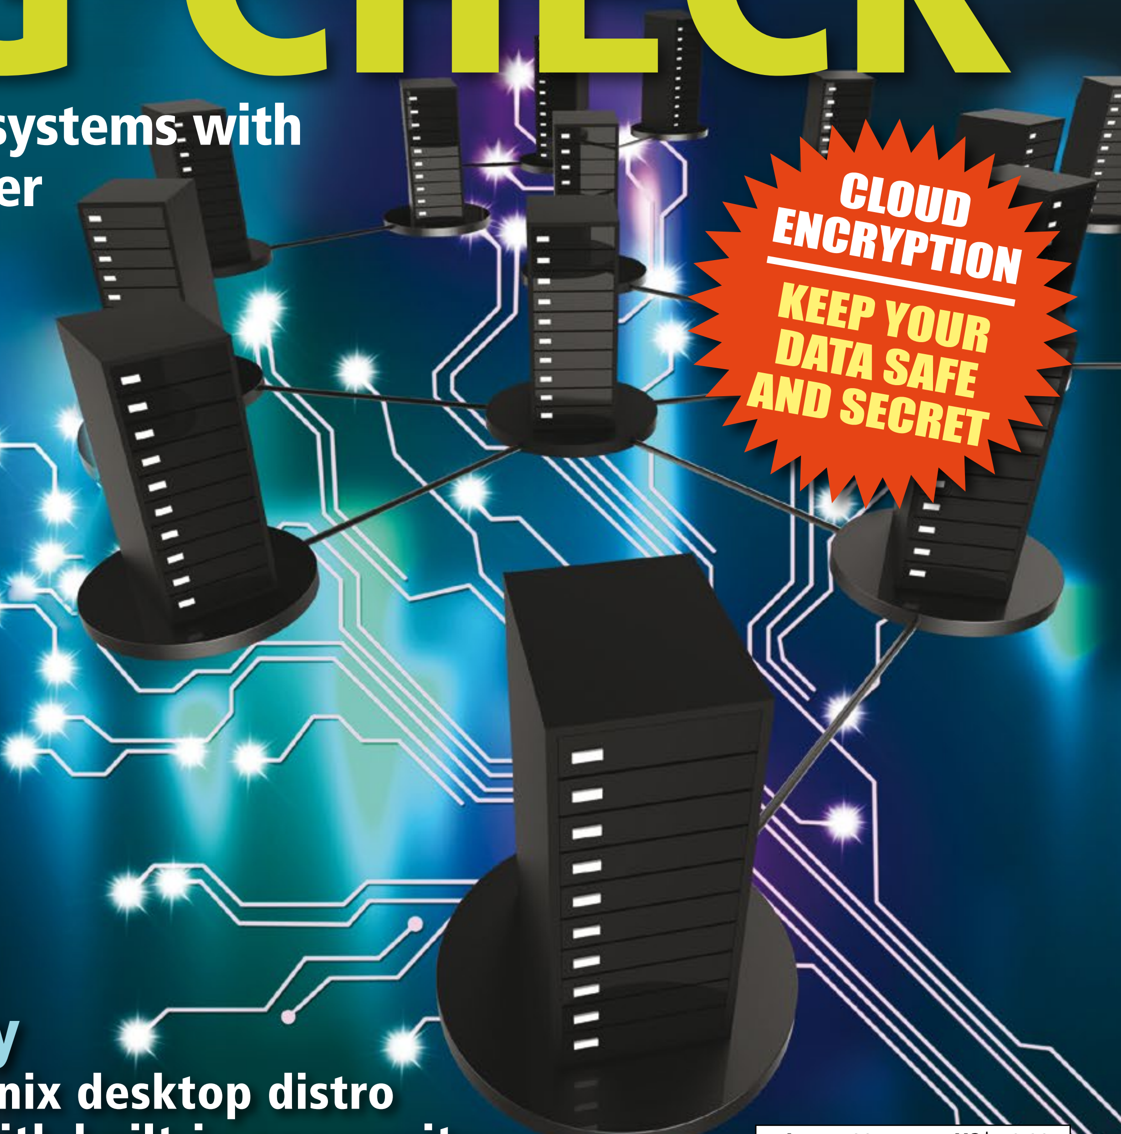

CLOUD ENCRYPTION

KEEP YOUR DATA SAFE AND SECRET

Automation Tricks

Use a Perl script to watch your eBay account

Choosing a Spam Filter

Persistent Firewall

Reload iptables rules automatically

Privacy

The Whonix desktop distro comes with built-in anonymity

Issue 188
July 2016
US$ 12.99
CAN$ 13.99

AnyDesk

Fast and secure remote desktop

File Sync

Osync and Freehold offer smooth online backup

Lightworks

Free video editor with lots of extras

WWW.LINUXPROMAGAZINE.COM

HETZNER ONLINE

Dedicated Servers for all users.

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

Dedicated Root Server PX61-SSD

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

monthly $ **68**

Dedicated Root Server PX61-NVMe

Intel® Xeon® E3-1275 v5
Quad-Core Skylake Processor
64 GB DDR4 ECC RAM
2 x 512 GB NVMe Gen3 x4 SSD
Guaranteed 1 Gbit/s bandwidth
100 GB Backup Space
30 TB traffic inclusive*
No minimum contract
Setup Fee $136

monthly $ **68**

* There are no charges for overage. We will permanently restrict the connection speed if more than 30 TB/month are used. Optionally, the limit can be permanently cancelled by committing to pay $1.35 per additional TB used.

www.hetzner.de/us

THE EXTRA MILE

Dear Linux Pro Reader,

The illusive interface between business and the open source community is difficult to define or even describe to the uninitiated. No one who was writing the management textbooks back when I was in college would believe the mysterious melding of community and commerce that is an everyday part of the FOSS ecosystem today. Conventional wisdom would claim the community isn't really a community or that a business helping to build this community could only be acting as some kind of PR-inspired charity. The whole idea that investing in community is a healthy and enlightened way to pursue vital business interests is lost on many in the business world even now, which is why it is important to celebrate when *real* businesses do *real* good things for the open source community.

Tag1 is a consulting company that does a lot of web development. Much of their work is with the Drupal open source content management system. Drupal has an active and vibrant user community around the world that gets together online and in person – at the popular DrupalCon conference series. But the developers, designers, and webmasters at Tag1 realized the Drupal community was missing something important: a magazine. Although they had no previous experience in publishing, Tag1 launched *Drupal Watchdog*, a print magazine for the whole Drupal community. *Drupal Watchdog* covered the whole spectrum of the Drupal experience, from design, to development, to daily tasks, to the management and marketing that surrounds the web content industry – and the emphasis was on the *whole* community. *Drupal Watchdog* had authors, and even advertisers, who worked for other consulting companies, but the Tag1 team wanted to reach everyone. *Drupal Watchdog* quickly became a favorite at the DrupalCon Conference, and readers subscribed and signed up for back issues through the *Drupal Watchdog* website.

After 10 semi-annual issues, Tag1 realized *Drupal Watchdog* was ready for a change. A company that was only interested in narrow, short-term goals could have stopped publishing right then. They were, after all, a company of consultants, not editors and publishers. They also could have sold off all the pieces, as the management textbooks would probably have advised them to do: maximize the return by liquidating the mailing list, the newsletter names, and the website. But Tag1 wanted *Drupal Watchdog* to keep growing and flourishing. So they went to extra effort and expense to find a professional publishing company that understood the open source community and help *Drupal Watchdog* reach a new audience of readers around the world.

And that company is … our company: Linux New Media, publishers of *Linux Magazine*, *Linux Pro Magazine*, *Ubuntu User*, and other publications that serve the FOSS community. We're proud to welcome *Drupal Watchdog* to the Linux New Media family and excited to consider all the synergies we'll stir up with the Drupal community as we make *Drupal Watchdog* into a regular quarterly and send it out through our global distribution network.

If we succeed with our efforts to take *Drupal Watchdog* to the next level, it will help the whole Drupal community, which will, in turn, help Tag1 and all other Drupal consultants. You and I know it, and Tag1 knows it: Community works! We'll just have to be patient and wait for those old-school authors of the management textbooks to catch up.

Welcome *Drupal Watchdog*, and here's to Tag1 for going the extra mile.

Joe

Joe Casad,
Editor in Chief

FREE DVD
ubuntu & ubuntuMATE
JUST RELEASED!
LINUX PRO MAGAZINE
LOG CHECK
Monitor your systems with a central log server
LOG CHECK
Monitoring your systems with a central log server
Automation Tricks
Use a Perl script to watch your eBay account
Choosing a Spam Filter
Persistent Firewall
Reload iptables rules automatically
Privacy
The Whonix desktop distro comes with built-in anonymity
CLOUD ENCRYPTION
KEEP YOUR DATA SAFE AND SECRET
AnyDesk
Fast and secure remote desktop
File Sync
Dsync and Freehold offer smooth online backup
Lightworks
Free video editor with lots of extras
WWW.LINUXPROMAGAZINE.COM

news

Log Check

System logs offer clues for tracking intruders and troubleshooting problems.

SERVICE

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

HIGHLIGHTS

FEATURES

LINUXUSER

REVIEW

ubuntu MATE
16.04 Desktop (32-bit)
CC BY-SA 4.0

- First Mate LTS version
- Reduced CPU requirements
- One-click installs
- Extended systemd support

ubuntu
16.04 Desktop (64-bit)

- Updated GNU toolchain
- Apt privilege separation
- systemd init
- Linux kernel 4.4

SEE PAGE 6 FOR DETAILS

On the DVD

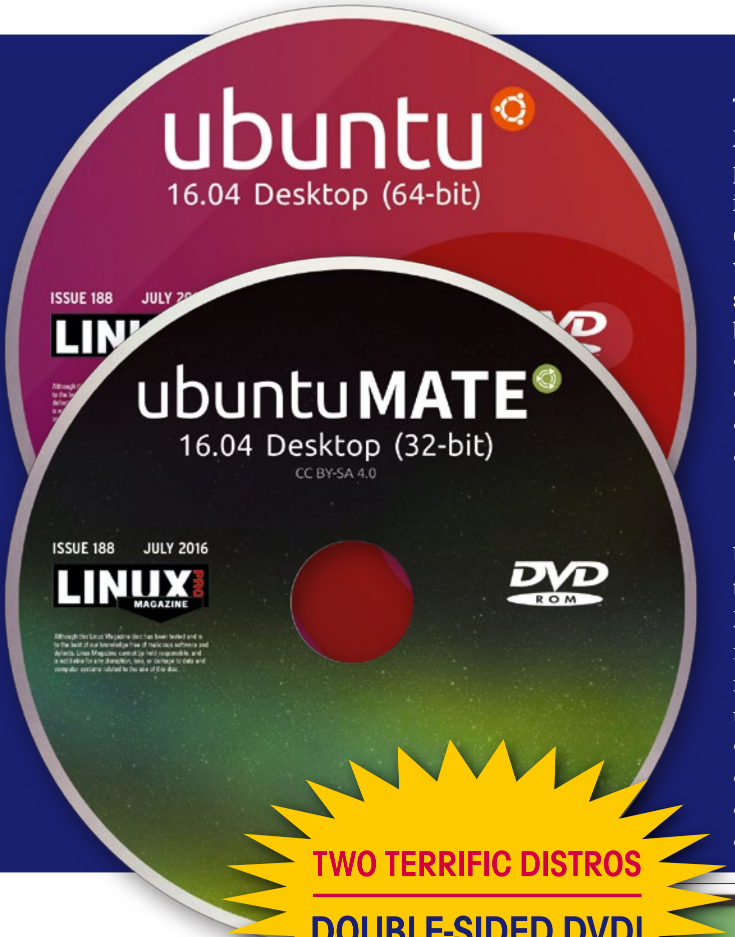

TWO TERRIFIC DISTROS

DOUBLE-SIDED DVD!

Ubuntu 16.04 (64-bit)

The newest long-term support (LTS) version of Ubuntu (Xenial Xerus) now arrives with search lens off by default for increased privacy. From the command line, the new Snap container packaging system resides alongside the legacy DEB system (do `sudo apt dist-upgrade` first). The Software Center is gone, and Ubuntu Software Store (from Gnome Software) now handles GUI software search and installation chores. Xenial Xerus also sports numerous bug fixes and improvements.

- systemd init
- Linux kernel 4.4
- Updated GNU toolchain
- Apt privilege separation

Ubuntu 16.04 Mate (32-bit)

Ubuntu Mate is designed for "those who want the most out of their computers and prefer a traditional desktop metaphor." The Mate desktop, which is based on Gnome 2, is full-featured enough for the modern user but still accommodates older computers with modest hardware profiles. In this version of Mate, you'll find many updates and improvements.

- Reduced CPU requirements
- First Mate LTS version
- One-click installs
- Extended systemd support

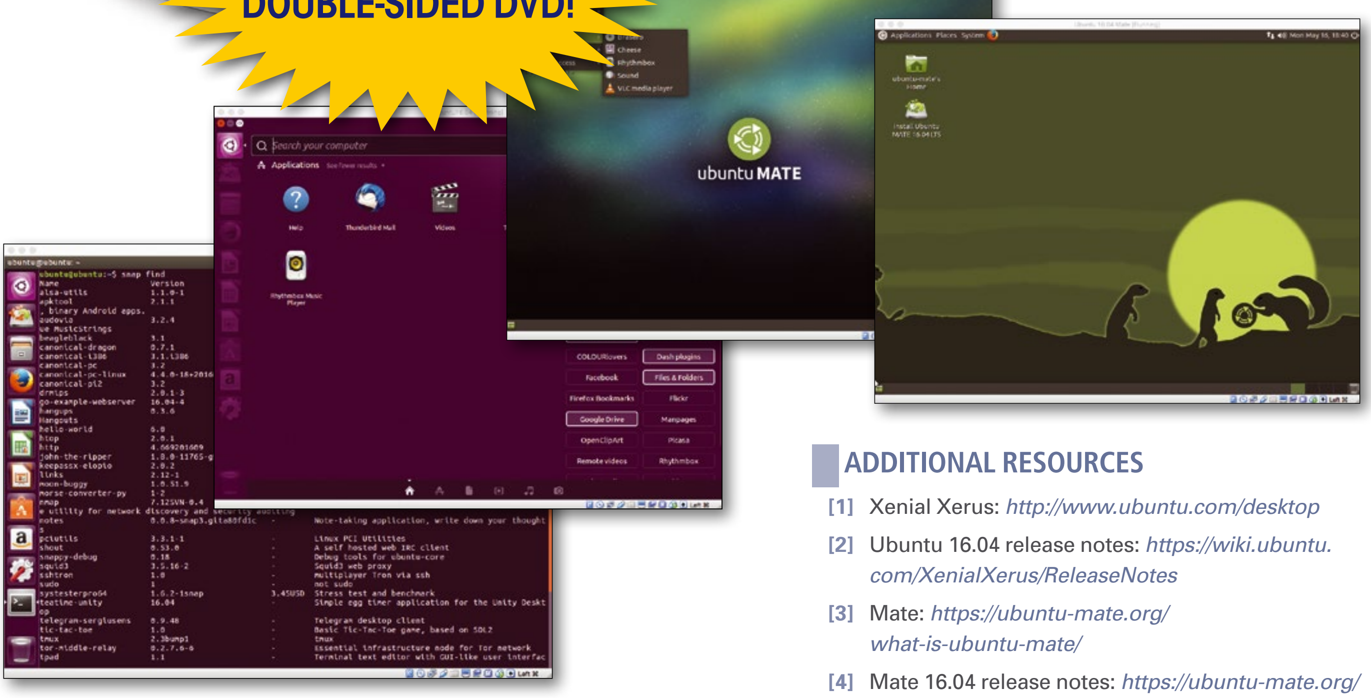

```
ubuntu@ubuntu:~$ snap find
Name                  Version
alsa-utils            1.1.0-1
apktool               2.1.1
, binary Android apps.
audovia               3.2.4
ue MusicStrings
beagleblack           3.1
canonical-dragon      0.7.1
canonical-i386        3.1.i386
canonical-pc          3.2
canonical-pc-linux    4.4.0-18+2016
canonical-pi2         3.2
drmips                2.0.1-3
go-example-webserver  16.04-4
hangups               0.3.6
Hangouts
hello-world           6.0
htop                  2.0.1
http                  4.669201609
john-the-ripper       1.8.0-11765-g
keepassx-elopio       2.0.2
links                 2.12-1
moon-buggy            1.0.51.9
morse-converter-py    1-2
nmap                  7.12SVN-0.4
e utility for network discovery and security auditing
notes                 0.0.8-snap3.gita80fd1c      -        Note-taking application, write down your thought
s
pciutils              3.3.1-1                     -        Linux PCI Utilities
shout                 0.53.0                      -        A self hosted web IRC client
snappy-debug          0.18                        -        Debug tools for ubuntu-core
squid3                3.5.16-2                    -        Squid3 web proxy
sshtron               1.0                         -        multiplayer Tron via ssh
sudo                  1                           -        not sudo
systesterpro64        1.6.2-1snap                 3.45USD  Stress test and benchmark
teatime-unity         16.04                       -        Simple egg timer application for the Unity Deskt
op
telegram-sergiusens   0.9.48                      -        Telegram desktop client
tic-tac-toe           1.0                         -        Basic Tic-Tac-Toe game, based on SDL2
tmux                  2.3bump1                    -        tmux
tor-middle-relay      0.2.7.6-6                   -        Essential infrastructure node for Tor network
tpad                  1.1                         -        Terminal text editor with GUI-like user interfac
```

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

ADDITIONAL RESOURCES

[1] Xenial Xerus: *http://www.ubuntu.com/desktop*

[2] Ubuntu 16.04 release notes: *https://wiki.ubuntu.com/XenialXerus/ReleaseNotes*

[3] Mate: *https://ubuntu-mate.org/what-is-ubuntu-mate/*

[4] Mate 16.04 release notes: *https://ubuntu-mate.org/blog/ubuntu-mate-xenial-final-release/*

NEWS

Updates on technologies, trends, and tools

THIS MONTH'S NEWS

OwnCloud Founder Resigns

Frank Karlitschek, the founder and CTO of ownCloud has resigned from the company that he founded some four years ago. He resigned due to conflicts between the interests of the company and the interests of the ownCloud community.

According to Karlitschek, the company failed to recognize the contributions and achievements of the ownCloud community. He said that the company has "a tendency to control the work too closely and discuss things internally."

Announcing his resignation, Karlitschek wrote in a blog post, "I thought a lot about this situation. Without sharing too much, there are some moral questions popping up for me. Who owns the community? Who owns ownCloud itself? And what matters more, short-term money or long-term responsibility and growth? Is ownCloud just another company, or do we also have to answer to the hundreds of volunteers who contribute and make it what it is today? These questions brought me to the very tough decisions: I have decided to leave my own company today. Yes, I handed in my resignation and will no longer work for ownCloud, Inc."

Karlitschek will continue to lead the open source ownCloud project, as long as the community allows him to do so. The conflict between the community and the company creates the possibility for an ownCloud fork.

Ubuntu 16.04 LTS Released

Canonical has announced the arrival of Ubuntu 16.04. The new release is a long-term support (LTS) release that's suitable for servers and enterprise customers. LTS releases are supported for five years, whereas regular releases are supported for nine months.

Ubuntu 16.04 comes with many new technologies and features aimed at enterprise users. This release introduces snaps, a new package format for containerized Ubuntu apps. Snaps will continue to co-exist with .deb-based packages for a foreseeable time. Canonical has added new redundancy features to the ZFS filesystem. The company has also added LXD, a pure container hypervisor with support for the recently released OpenStack Mitaka. With this release, Ubuntu is now available on IBM's z Systems and LinuxONE mainframe machines.

On the desktop, Ubuntu 16.04 comes with Unity 7.x, which disables the Dash online search by default. Users now have the ability to customize the appearance of menu items and can change the location of the Unity Launcher through the Unity Tweak tool. With this release, Ubuntu has dropped the homegrown Ubuntu Software Center and moved to Gnome Software.

Ubuntu 16.04 and its official flavors are available for download immediately.

Microsoft Visual Studio Code v1.0 Is Available for Linux

Microsoft has released the first stable version of Visual Studio Code (VS Code) for Linux and Mac OS X. The company claims that more than 500,000 developers are actively using VS Code each month. Visual Studio Code is Microsoft's code editor, which the company recently open sourced. Visual Studio Code is based on the Atom Shell (now known as Electron) framework, which is developed by GitHub and used in GitHub's text editor Atom.

Microsoft is offering the stable version of VS Code in .deb and .rpm binaries. The company is also offering Insider builds to provide developers with early access to new features and extensions currently in development.

In a blog post, the Visual Studio Code team writes, "VS Code was initially built for developers creating web apps using JavaScript and TypeScript. But in less than 6 months since we made the product extensible, the community has built over 1000 extensions that now provide support for almost any language or runtime in VS Code."

Open Container Initiative Announces Image Format Project

Open Container Initiative (OCI), a Linux Foundation Collaborative Project, has announced a new initiative called the OCI Image Format project. The primary goal of the new project is to create a software container image format spec with security and federated naming as key components. The OCI Image Format project is hosted on GitHub.

According to the announcement, "This represents an expansion of the OCI's first project, OCI Runtime Spec, that focuses on how to run containers. Industry leaders are collaborating to enable users to package and sign their application, then run it in any container runtime environment of their choice – such as Docker or rkt. With the development of the new OCI Image Specification for container images, both vendors and users can benefit from a common standard that is widely deployable across any supporting environment of the user's choice."

Jonathan Boulle of CoreOS, a company that develops a lightweight Linux distribution for container deployment, praised the formation of the OCI Image Format project and wrote in a blog post, "An open, and openly implementable container image format specification underpins all the portability goals of containers, allowing users to build and package a container once, sign it, and run it in a variety of vendor implementations and platforms, in the cloud and on-premises."

Qualcomm Bug Threatens Millions of Android Devices

FireEye, a cybersecurity firm, has found a flaw in Android devices running Qualcomm chips. The vulnerability has existed in Android devices for the last five years, and it affects devices with Qualcomm processors running Android 4.3 and older Android systems. Devices running newer versions of Android take advantage of SEAndroid, but FireEye says they are still affected to some extent.

According to a FireEye blog post, "This vulnerability allows a seemingly benign application to access sensitive user data, including SMS and call history, and the ability to

MORE ONLINE

Linux Pro Magazine

www.linuxpromagazine.com

Off the Beat • Bruce Byfield

Why the Ubuntu Tablet Matters

I am not generally a technophile. I don't obsess over hardware stats, and I judge hardware on how well it does its job. Yet recently, I found myself anticipating the release of three pieces of hardware. One is the pi-topCEED, the Raspberry Pi-based computer for education.

Hating Microsoft

Hating Microsoft has been part of open source from the start. Infamous for its executives calling Linux unAmerican, the equivalent of communism and a cancer, Microsoft has been the arch-enemy, working behind the scenes in reality almost as much as in the minds of conspiracy theorists, the proprietary and corporate antithesis of everything that open source is about.

The Satisfactions of a Free License

Offer free software developers money, and they are practical enough to accept it. However, what keeps many of them at work are the intrinsic rewards, not the external ones.

Ordering a custom computer

"You can't customize your computer," a sysadmin once admonished me. "The parts need to be carefully matched with one another, and that's not a job for an amateur." Despite this advice, I have ordered customized workstations for over two decades.

Paw Prints • Jon "maddog" Hall

Brazil: Free and Open Source Culture Is Economics, Not Politics

Over the years people have accused Free and Open Source Culture (FOSC) of being a "religion." Other people have used FOSC as a political tool, assigning the advocacy of FOSC to one political party; usually the "left," "liberal" or (as some people call them) "progressive" party. FOSC is none of these.

Productivity Sauce • Dmitri Popov

Ansiweather: Weather in the Terminal

Sometimes the simplest tool can also prove to be an indispensable one. Take Ansiweather, for example: This one-trick pony displays the current weather conditions and forecast right in the terminal, and that's all it does.

perform potentially sensitive actions, such as changing system settings or disabling the lock screen."

FireEye informed Qualcomm of the bug in January, and Qualcomm released a fix by April, making it available to all vendors. Google pushed the fix to Nexus devices in May. Although Google secured its own Nexus devices, the company has no control over the rest of the Android ecosystem. Carriers and Android hardware vendors control software updates on their own Android devices, and users of these devices will remain vulnerable unless these companies update the software.

Windows 10 Pro Loses Critical Features

Business customers running Windows 10 Pro will no longer be able to use the Group Policy feature to restrict employees from accessing the Windows Store. Microsoft made this change last month with the upgrade to version 1511 of Windows 10. After this upgrade, users can't disable Windows Store access through Group Policy. According to Microsoft's support page, "This behavior is by design. In Windows 10 version 1511, these policies are applicable to users of the Enterprise and Education editions only."

A Microsoft spokesperson told ZDNet "Windows 10 Pro offers a subset of those capabilities and is recommended for small and mid-size businesses looking for some management controls, but not the full suite necessary for IT pros at larger enterprises."

Businesses need tighter control over their systems, and Microsoft is encouraging enterprise customers to use the Windows 10 Enterprise edition, which lets customers restrict access to Windows Store through AppLocker or Group Policy.

JBoss Vulnerability Could Lead to SamSam Ransomware

Researchers at Cisco Talos found a vulnerability in JBoss that can be exploited by SamSam ransomware. Cisco Talos said in a blog post, "As part of this investigation, we scanned for machines that were already compromised and potentially waiting for a ransomware payload. We found just over 2,100 backdoors installed across nearly 1600 IP addresses." The research firm says they estimate over 3.2 million machines are at risk.

SamSam is distributed through compromised servers and then holds victim systems for ransom. Attackers are using the JexBoss open source tool to test and then exploit JBoss application servers. Once they gain access to the network, they start encrypting Windows systems using SamSam.

Cisco Talos suggests that if your server is vulnerable, the first piece of advice is to remove external access to the server. "Ideally, you would also re-image the system and install updated versions of the software," the firm said in the blog post.

New Exploit Bypasses Windows AppLocker

A new Windows vulnerability allows attackers to install any application on Windows systems, bypassing AppLocker. AppLocker is a feature of Windows 7 and Windows Server 2008 R2 that allows admins to manage application access to users. This serious flaw targets business users and not just home users, and it affects the latest Windows 10 systems, as well as earlier versions of Windows going all the way back to Windows 7.

The vulnerability was accidentally discovered by Casey Smith, who realized that the Windows command-line utility Regsvr32 can be exploited to bypass AppLocker by registering and unregistering DLLs. Because this method doesn't touch the system registry, system admins won't find any trace of changes to the system.

Microsoft has not yet released a fix for the vulnerability; however, users can mitigate it by blocking Regsvr from the Windows Firewall.

MORE ONLINE

Prune Your Photo Library with fdupes
If your photo library contains thousands of photos, chances are it has duplicate files lurking in its corners. But finding and removing these unwelcome guests can be tricky, unless you use the fdupes tool for the job.

Instant Pronounceable Passwords with passwds.ninja
The passwds.ninja web app can come in rather handy when you need an easy-to-remember password.

ADMIN HPC

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

Interview with the Developer of Singularity • Jeff Layton
Sometimes we see the names of people working on the Linux kernel or other high-profile projects, but we don't hear much about these people working behind the scenes that are contributing their knowledge for the greater good.

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

ADMIN Online

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

Backups Using rdiff-backup and rsnapshot • Georg Schönberger
The first step in ensuring comprehensive backups is to consider where the backups should be stored; therefore, a separate backup server is often used that connects to other computers and initiates the backups.

Integrating FreeIPA with Active Directory • Thorsten Scherf
A directory service usually provides a wealth of information on top of the classic user and group accounts, including machine and service accounts, security rules, and possibly DNS information and other data that administrators would like to store centrally to deliver to clients in the domain.

Freeing your Data from Ransomware • Thomas Gronenwald and Stefan Becker
One of the latest trends among cyber criminals is to encrypt files on a network, forcing the user to pay a ransom for the decryption of their data.

Centralized log management with Graylog

Watching the Logs

System logs offer clues for tracking intruders and troubleshooting problems. If you're in charge of a whole network, wouldn't you rather monitor all your logs from a single central point? Graylog and its companion components let you manage all your logs from a single interface. *By Thomas Gronenwald, Giuseppe Dispinzeri, and Michael Classes*

Logfiles chronicle the state of the system, and experienced admins know to check the logs for messages when a problem arises. If you only administer one computer and it is sitting on your desk, the task is easy. But if you're taking care of several systems on a diverse network, keeping up with all the logfiles can be a major chore.

Several commercial tools fill the role of managing and monitoring log messages across the network, but you don't have to spend big to get big-time log monitoring capabilities. This article describes how to configure network monitoring using a configuration centered around the Graylog log server.

Logging Server Architecture

Graylog is an open source log management tool, providing central storage, processing, and analysis of log messages from servers, clients, or network devices. The Graylog log server is based on Java and offers a means for combining several server nodes in a cluster for high availability and scalability.

The architecture of a typical Graylog implementation is shown in Figure 1. Central elements for the Graylog central logging systems include:

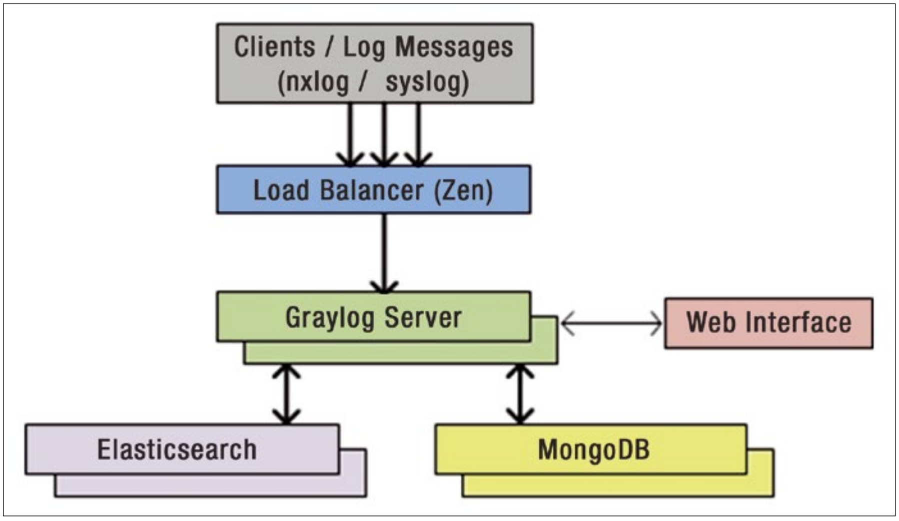

Figure 1: Simplified schematic diagram of a Graylog architecture.

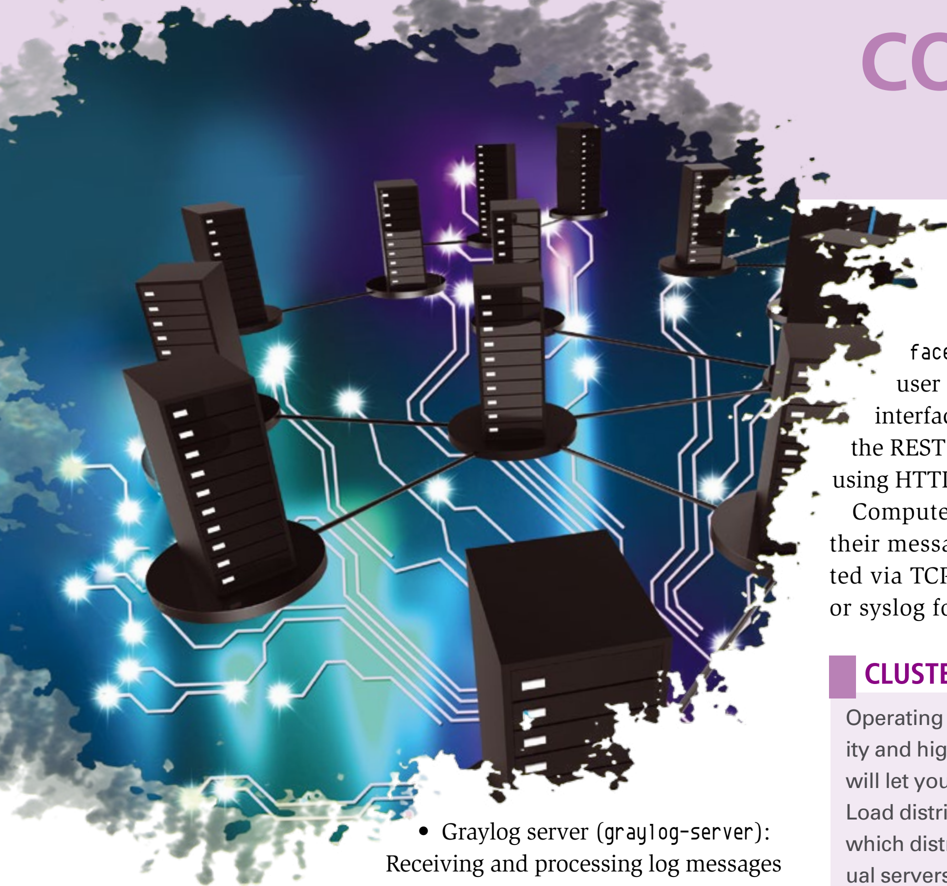

- Graylog server (`graylog-server`): Receiving and processing log messages and alarms.
- Web interface (`graylog-web-interface`): Web-based access to the log management software in a browser for administration, configuration, and monitoring of Graylog.
- MongoDB: Storage of configuration and metadata of the Graylog server.
- Elasticsearch: an index and search server that offers a convenient interface for accessing log messages.

As you will learn later in this article, many users prefer to implement Graylog in a cluster configuration with a load balancing tool to distribute the log messages across the Graylog servers (as shown in Figure 1). See the box titled "Clusters" for more on configuring Graylog in a cluster configuration.

The Graylog web interface (`graylog-web-interface`) lets you authenticate the user with a separate user account or LDAP. Communication between the web interface and the Graylog server (`graylog-server`) relies on the REST protocol (HTTP-based), which you can protect using HTTPS.

Computers on the network act as clients, transmitting their messages to the log server. Log messages are transmitted via TCP or UDP in GELF (Graylog extended log format) or syslog format, and you can use TLS to encrypt the com-

CLUSTERS

Operating the Graylog server as a cluster will help with scalability and high availability (HA). Adding more Graylog server nodes will let you handle a higher number of log messages per minute. Load distribution is carried out by an upstream load balancer, which distributes the log messages for processing to the individual servers. The extra nodes also increase resilience, because the logging server continues to work if one Graylog server fails.

The Elasticsearch server and MongoDB database can also operate with multiple instances (Figure 2). The configuration described in this article supports a MongoDB database as a replica set, with the first instance of the database (MongoDB1) running in "Master" mode and the second instance (MongoDB2) in "Slave" mode. The master database is configured so that it is automatically replicated to the second instance. The replica set of the database boosts data availability, and you can easily extend it to include more database instances.

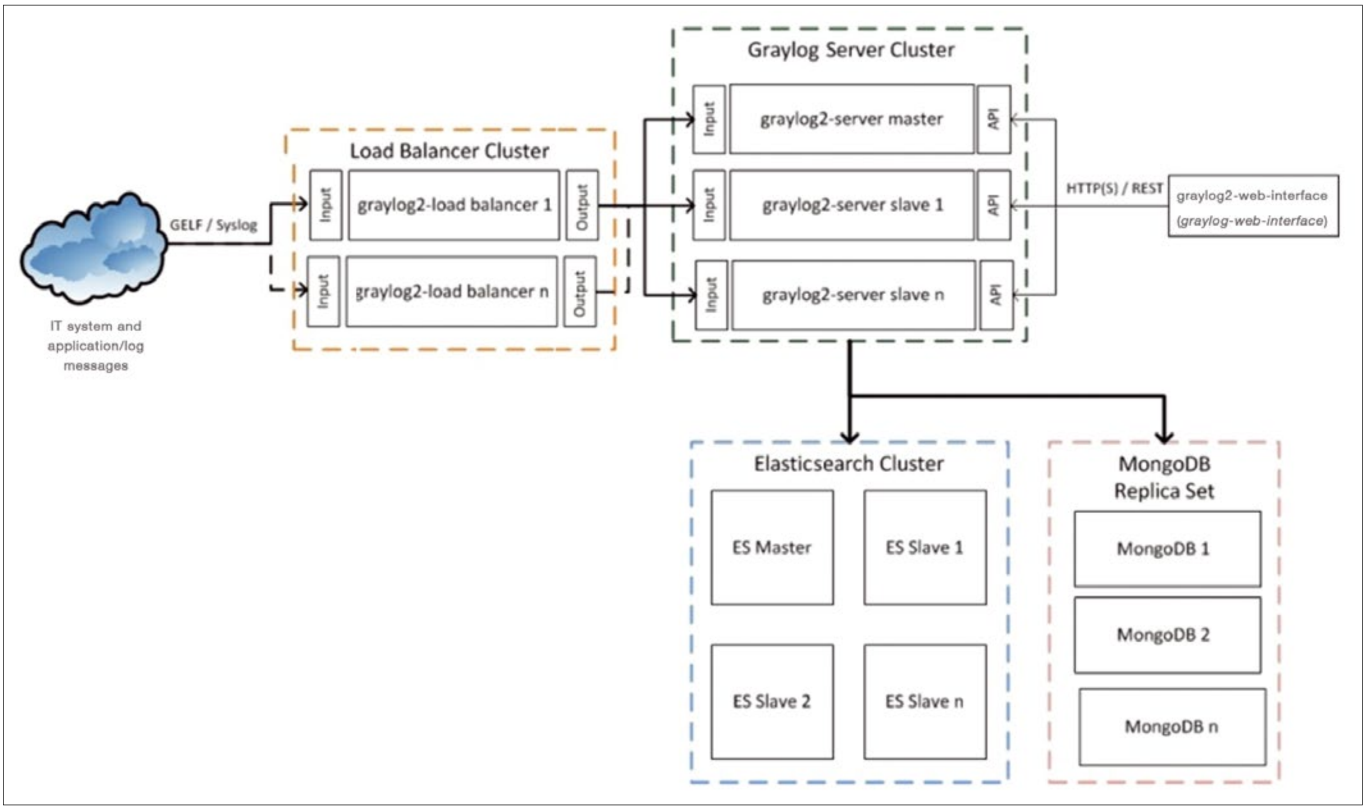

Figure 2: **The Graylog components operate in clusters for scalability and high availability.**

munication between the Graylog server and its back end, a MongoDB database.

You'll need to run the Graylog server on a Linux system with Java 7 or later.

Installing Graylog

Setting up Graylog starts with retrieving the component packages. I will configure the various components as virtual machines running a Debian 7 (wheezy) operating system as a basis.

Follow the installation steps in Listing 1 for `graylog-server` on `graylog-ms` (master) and `graylog-node1` (slave) instances; then, set up the MongoDB database. To set up MongoDB, run the commands in Listing 2 for the VMs `graylog-ms` and `graylog-node1`.

After you install Graylog, the next step is the configuration. For the `graylog-ms` master VM, perform the following changes to the `/etc/graylog/server/server.conf` configuration file:

- `is_master = true`: The virtual machine `graylog-ms` is the master. If you use several `graylog-server` systems, only one system can be the master.
- `rest_listen_uri = http://graylog-ms:12900/`: URI of the interface REST API, which must be accessible when using a cluster. Replace the name of the `graylog-ms` system with the corresponding IP address of the system or define it in your `/etc/hosts` configuration file.
- `password_secret = complex password`: Password required for encrypting other passwords and for generating random strings (salts). The value used for `password_secret` must be identical on all `graylog-server` instances. You can create a complex password, for example, with the

```
pwgen n 1 s 96
```

  command.
- `root_password_sha2 = SHA2 hash value`: Hash value of the password for logging via the web interface with the user name `admin`. Create the hash value using

```
echo -n MyComplexesPassword | shasum -a 256
```

- `elasticsearch_max_docs_per_index = 20000000`: Number of log messages to keep per index. This value is the default and relates to the Elasticsearch component.
- `elasticsearch_max_number_of_indices = 20`: Total number of indexes. The value of `20` is the default and relates to the Elasticsearch component. The total number of possible log messages that can be stored is calculated by multiplying the configuration parameter `elastic-search_max_docs_per_index` by the `elasticsearch_max_number_of_indices`.
- `elasticsearch_shards = 2`: The total number of shards (i.e., allocation of indexes to systems with the Elasticsearch component). The value depends on the number of Elasticsearch components.
- elasticsearch_replicas = 1: The total number of original copies of the indexes. The value of this parameter depends on the Elasticsearch component. A value of 1 creates a copy of all log messages on the master VM, as well as on `es-node1`.
- `mongodb_*`: Change the following configuration parameters for integrating the MongoDB database:

```
mongodb_host = *IP_address of "graylog-ms"*
mongodb_database = graylog2
mongodb_port = 27017
```

  You can leave the remaining MongoDB configuration parameters as the defaults.

For the `graylog-node1` VM, modify the `/etc/graylog/server/server.conf` configuration file with all the same settings you configured for `graylog-ms`, except for the following parameters:

- `is_master = false`
- `rest_listen_uri = http://graylog-node1:12900/` (URI of the interface REST API, which must be accessible when using a cluster.) Replace the name of the `graylog-node1` system with the IP address of the system or define it in your `/etc/hosts` configuration file.
- `mongodb_*`: Change the following configuration parameters for integrating the MongoDB database:

```
mongodb_host = *IP_address of "graylog-node1"*
mongodb_database = graylog2
mongodb_port = 27017
```

For a detailed description of each MongoDB configuration parameter, see the Graylog documentation [1].

Graylog Web Interface

For the `graylog-web-interface` VM, you need to make some changes to the `/etc/graylog/web/web.conf` Graylog configuration file:

```
graylog2-server.uris=↩
  "http://graylog-ms:12900/,↩
   http://graylog-node1:12900/"
```

Replace the names `graylog-ms` and `graylog-node1` with the IP addresses of the systems or define them in the `/etc/hosts` configuration file.

LISTING 1: Installing Graylog

```
$ wget https://packages.graylog2.org/repo/packages/graylog-1.0-repository-debian7_
  latest.deb
$ sudo dpkg -i graylog-1.0-repository-debian7_latest.deb
$ sudo apt-get install apt-transport-https
$ sudo apt-get update
$ sudo apt-get install graylogserver
```

LISTING 2: Setting up MongoDB

```
$ sudo apt-key adv --keyserver keyserver.ubuntu.com --recv 7F0CEB10
$ sudo apt-get install lsb-release
$ echo "deb http://repo.mongodb.org/apt/debian "$(lsb_release-sc)"/mongodb-org/3.0
  main" | sudo tee /etc/apt/sources.list.d/mongodb-org-3.0.list
$ sudo apt-get update
$ sudo apt-get install -y mongodb-org
```

Use the following command

```
Application.secret=<complex password>
```

to enter the password required to encrypt other passwords and generate random strings (salts).

Configuring Elasticsearch

The Elasticsearch tool lets you index, organize, and search on the log messages in the Graylog message database. To configure the Elasticsearch component, you need to set up on the `es-master` and `es-node1` VMs with the following commands:

```
$ wget https://download.elastic.co/elasticsearch/elasticsearch/⮐
  elasticsearch-1.5.2.deb
$ sudo dpkg -i elasticsearch-1.5.2.deb
```

To access the logging server, you will need the web interface. Set this up on the `graylog-web-interface` VM:

```
$ wget https://packages.graylog2.org/⮐
  repo/packages/graylog-1.0-repository-⮐
  debian7_latest.deb $ sudo dpkg -i ⮐
  graylog-1.0-repository-debian7_latest.deb
$ sudo apt-get install apt-transport-https
$ sudo apt-get update
$ sudo apt-get install graylog-web
```

For the master VM, make the following changes to the configuration file `/etc/elastic-search/elasticsearch.yml`:

- `cluster.name: graylog-production` – Unique identifier of the cluster for the Elasticsearch component.
- `node.name: es-master` – Unique name of the node in the cluster of the Elasticsearch component.
- `node.master: true` –The node acts as a master in the Elasticsearch cluster.
- `node.data: true` – The node (Elasticsearch component) stores data.
- `index.number_of_shards: 2` – See `elasticsearch_shards` in the Graylog server configuration.
- `index.number_of_replicas: 1` – See `elasticsearch_replicas` in the Graylog server configuration.
- `discovery.zen.ping.multicast.enabled: false` – Disable the multicast discovery function to avoid sending multicast requests to determine the nodes in the cluster.
- `discovery.zen.ping.unicast.hosts: ["es-master:9300", "es-node1:9300"]` – A list of nodes that make up the Graylog production cluster. Replace the names in `/etc/hosts` or use the IP addresses for your systems accordingly.

You will find the complete and detailed description of each configuration parameter in the documentation [1].

Make all the same changes to the `/etc/elasticsearch/elasticsearch.yml` configuration file for the `es-node1` VM except for the following:

- `node.name: es-node1` – Unique name of the node in the cluster of the Elasticsearch component.
- `node.master: false` – The node does not act as master in the cluster.

Setting Up the Load Balancer

The Zen load balancer will distribute the log message traffic among the Graylog servers. We used the Zen load balancer community edition [2]. The current stable version at the time of this article was version 3.05.

You can use the Zen administration panel web interface to configure the load balancer. The web interface is reachable at *https://IP_address_of_load_balanceer:444*. The username and the password are *admin* by default.

To configure the load balancer so log messages are split across two Graylog servers, you need to create a *farm* in the web interface. A farm is a profile that contains the configuration for a specific network protocol (such as TCP, UDP, or HTTP) and an algorithm for load balancing. After you have created a new farm, adjust the additional configuration parameters by adding the Graylog servers that receive the log messages and entering their IP addresses and ports.

This example assumes the load balancer is configured to balance the load between two Graylog servers that receive log messages from clients via UDP with `nxlog` and `syslog`. The IP addresses and associated ports of the systems are as follows:

```
IP address "graylog-lb": 192.168.15.86
IP address "graylog-ms": 192.168.15.86
IP address "graylog-node1": 192.168.15.86
UDP port "nxlog": 12201
UPD port "syslog": 1514
```

Each UDP port has a farm. The names for the farms are *GraylogL4xNAT-UDP-12201* and *GraylogL4xNAT-UDP-1514*. The configuration parameters for the *GraylogL4xNAT-UDP-12201* farm are

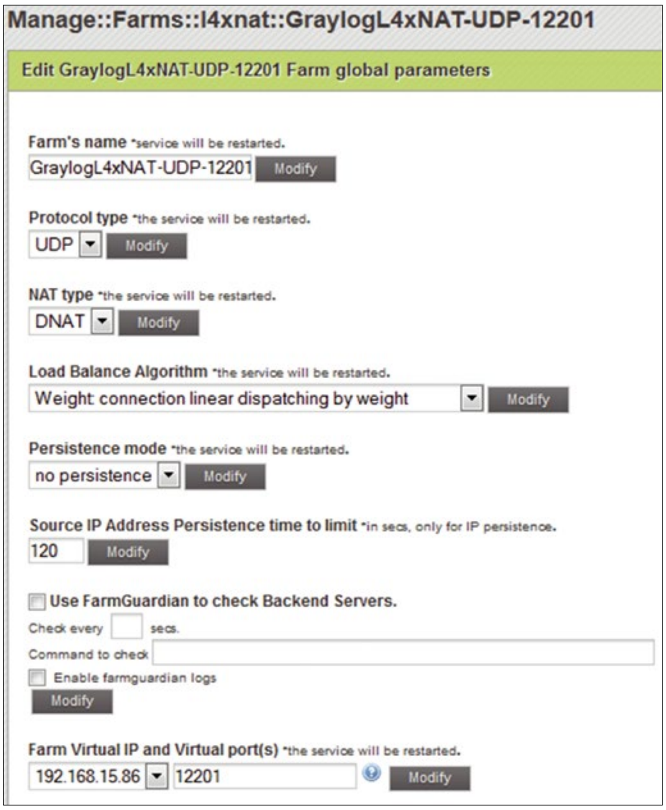

Manage::Farms::l4xnat::GraylogL4xNAT-UDP-12201

Edit GraylogL4xNAT-UDP-12201 Farm global parameters

Farm's name *service will be restarted.
GraylogL4xNAT-UDP-12201

Protocol type *the service will be restarted.
UDP

NAT type *the service will be restarted.
DNAT

Load Balance Algorithm *the service will be restarted.
Weight: connection linear dispatching by weight

Persistence mode *the service will be restarted.
no persistence

Source IP Address Persistence time to limit *in secs, only for IP persistence.
120

□ Use FarmGuardian to check Backend Servers.
Check every secs.
Command to check
□ Enable farmguardian logs

Farm Virtual IP and Virtual port(s) *the service will be restarted.
192.168.15.86 12201

Figure 3: Configuration parameters for a farm – the UDP transport protocol is an important setting.

shown in Figure 3. Note that you select the *UDP* network protocol as the *Protocol type*. For load balancing, I have set the algorithm to *Weight connection linear dispatching by weight* as an example. Load distribution depends on the weighting; you set up the weighting in the next step for the two Graylog servers. Then enter the IP addresses and the corresponding ports of the two Graylog servers and set the weighting or priority according to the load distribution. (See the detailed description of each configuration parameter in the Zen documentation [3].)

Transferring Log Messages

Once you get the Graylog server up and running, you'll need a way for the other systems to forward their log messages to Graylog. Syslog (via TCP or UDP) is a useful choice as a client tool because it is available on most Linux systems and is typically supported by managed network devices such as routers, switches, and firewalls.

For systems that do not use syslog by default (e.g., Windows), you'll need the NXLog client software. NXLog Community Edition [4] supports multithreaded log management and various log message formats (syslog, CSV, GELF, JSON, XML, Windows EventLog). In addition to several Windows platforms, NXLog runs on several versions of Linux, as well as BSD and Android. NXLog, an open source program available free of charge [5], is a good option for mixed networks with a both Window and Linux clients. Linux – Debian 7 (wheezy) here – offers an up-to-date version available as a DEB package" (`nxlog-ce-x.x.x_debian-wheezy.deb`). Use the following command to install:

```
$ sudo dpkg -i nxlog-ce-x.x.x_debian-wheezy.deb
```

The NXLog configuration syntax is identical on Windows and Linux. On Windows platforms, the configuration file usually is located under `C:\Program Files(x86)\nxlog\conf\nxlog.conf`. The default installation configuration file is below `/etc/nxlog/nxlog.conf` on Debian 7. To transfer all the log messages stored in the event log of a Windows 7 client to the log server, you need an `nxlog.conf` configuration file that looks like Listing 3.

GELF [6] offers a number of advantages over syslog. See the entry for the IP address (192.168.15.86) and UDP port (12201) of the `graylog-lb` load balancer in Listing 3 (lines 20 and 21).

To transfer all the log messages stored in the Debian Linux logfile `/var/log/messages` to the Graylog server, add the entries in Listing 4 to the `nxlog.conf` configuration file.

For details on the individual configuration parameters, see the documentation for NXLog [4].

LISTING 3: nxlog.conf

```
01 define ROOT C:\Program Files (x86)\nxlog
02 Moduledir %ROOT%\modules
03 CacheDir %ROOT%\data
04 Pidfile %ROOT%\data\nxlog.pid
05 SpoolDir %ROOT%\data
06 LogFile %ROOT%\data\nxlog.log
07
08 <Extension gelf>
09      Module xm_gelf
10 </Extension>
11
12 <Input in>
13        Module im_msvistalog
14 # For windows 2003 and earlier use the following:
15 # Module im_mseventlog
16 </Input>
17
18 <Output out>
19        Module om_udp
20        Host 192.168.15.86
21        Port 12201
22        OutputType GELF
23 </Output>
24
25 <Route 1>
26        Path in => out
27 </Route>
```

Setting Up the Graylog Inputs

For the Graylog server to receive `graylog-ms` and `graylog-node1` log messages from the clients, you need to create and configure *inputs*. To select an input, choose *System* in the Graylog web interface. Under the Inputs heading, select an input from the drop-down menu.

As you can see in Figure 4, the input setting specifies the network protocol used and the format of the log messages – you create a new input by clicking *Launch new input*. You then specify the configuration parameters, such as the port to use to finish

LISTING 4: nxlog.conf Additions

```
User nxlog
Group nxlog

LogFile /var/log/nxlog/nxlog.log
LogLevel INFO

<Extension _syslog>
 Module xm_gelf
</Extension>

<Input in>
 Module im_file
 File "/var/log/messages"
</Input>

<Output fileout1>
       Module om_udp
       Host 192.168.15.86
       Port 12201
       OutputType GELF
</Output>

<Route 1>
       Path in => fileout1
</Route>
```

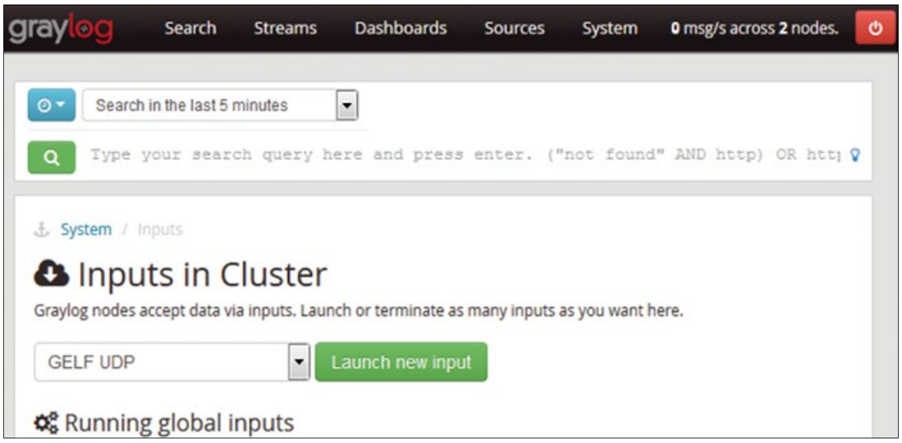

Type your search query here and press enter. ("not found" AND http) OR htt

Inputs in Cluster

Graylog nodes accept data via inputs. Launch or terminate as many inputs as you want here.

Running global inputs

Figure 4: Selecting an input in the Graylog web interface.

creating the input. Figure 5 shows finished inputs for the UPD ports 12201 (GELF format) and 1514 (syslog format), which are available on two Graylog servers.

Collecting Data with Streams

Now that the infrastructure is completed, it is time to start collecting data. To keep track on specific events in log messages, such as failed login attempts in SSH or Windows Remote Desktop from Graylog, you need *streams*. A stream is a function in Graylog that analyzes log messages in real time based on defined criteria to classify them into predefined categories. For example, a rule might tell Graylog to collect all messages that include the string "root."

Streams are also the basis for defining Graylog alerts. In a stream, you can define conditions, such as, when to notify via email. For example: "If more than eight failed attempts to log in to the SSH service are made within one minute, generate and send an alert."

Streams are set up in the Graylog web interface in the Streams menu, by selecting *Create stream*. Then, choose a title for the new stream, and optionally a description, and click on *Create stream and continue*. Now define the rules (Stream rules) for detecting certain log messages, such as failed logon attempts. You can see an example of a stream for failed logon attempts detected on the Windows Remote Desktop (RDP) service in Figure 6.

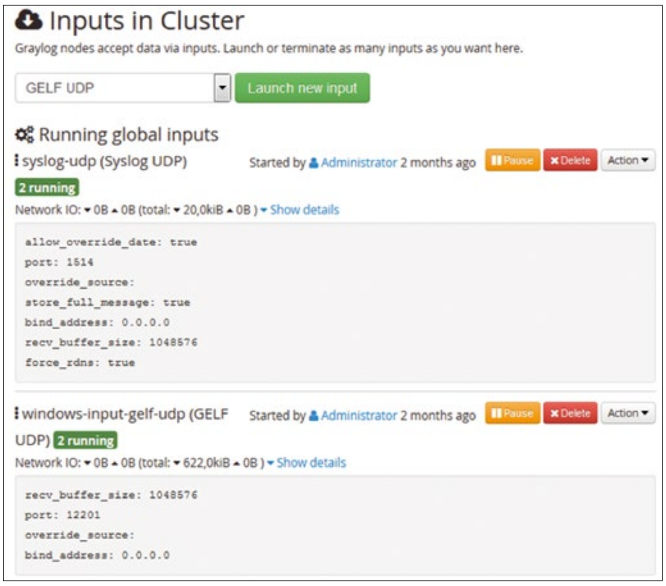

Inputs in Cluster

Graylog nodes accept data via inputs. Launch or terminate as many inputs as you want here.

Running global inputs

syslog-udp (Syslog UDP) Started by Administrator 2 months ago

Network IO: ▾ 0B ▴ 0B (total: ▾ 20,0kiB ▴ 0B)

```
allow_override_date: true
port: 1514
override_source:
store_full_message: true
bind_address: 0.0.0.0
recv_buffer_size: 1048576
force_rdns: true
```

windows-input-gelf-udp (GELF UDP) Started by Administrator 2 months ago

Network IO: ▾ 0B ▴ 0B (total: ▾ 622,0kiB ▴ 0B)

```
recv_buffer_size: 1048576
port: 12201
override_source:
bind_address: 0.0.0.0
```

Figure 5: Summary of the inputs created in Graylog.

To add a rule, select *Add stream rule*. Then set values for the fields *Field*, *Type*, and *Value*. The values for the first rule in the example are: `Field: Channel`, `Type: match exactly`, and `Value: Security`.

You will want to decide which events to classify as critical in advance of implementation and then define them as stream rules.

Conclusion

With a suitable stream and ruleset, you can analyze critical log messages in real time and define appropriate escalation measures in the advent of trouble. In combination with an audit policy, you can handle a large volume of information.

Graylog and its flotilla of related applications provide a powerful, elegant, and inexpensive solution for centralizing log management on a diverse network. ■■■

INFO

[1] Graylog: *https://www.graylog.org/*

[2] Zen Load Balancer: *https://www.zenloadbalancer.com/*

[3] Zen Load Balancer Administration Guide: *https://www.zenloadbalancer.com/zlb-administration-guide-v305/*

[4] NXLog Community Edition Reference Manual: *http://nxlog.org/docs/nxlog-ce/nxlog-reference-manual.html*

[5] NXLog: *http://nxlog-ce.sourceforge.net/download*

[6] Graylog Extended Log Format: *http://docs.graylog.org/en/latest/pages/gelf.html*

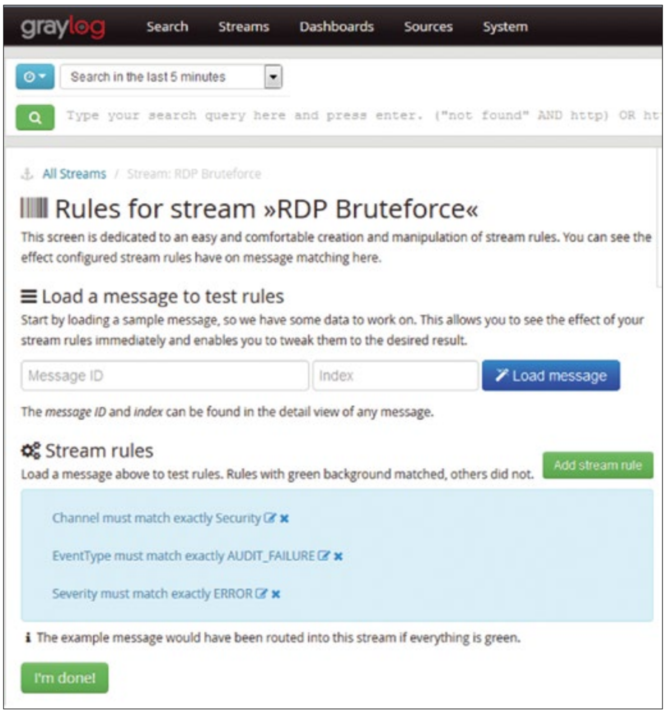

Type your search query here and press enter. ("not found" AND http) OR ht

Rules for stream »RDP Bruteforce«

This screen is dedicated to an easy and comfortable creation and manipulation of stream rules. You can see the effect configured stream rules have on message matching here.

Load a message to test rules

Start by loading a sample message, so we have some data to work on. This allows you to see the effect of your stream rules immediately and enables you to tweak them to the desired result.

The *message ID* and *index* can be found in the detail view of any message.

Stream rules

Load a message above to test rules. Rules with green background matched, others did not.

Channel must match exactly Security

EventType must match exactly AUDIT_FAILURE

Severity must match exactly ERROR

The example message would have been routed into this stream if everything is green.

Figure 6: Tracking failed RDP logon attempts.

Tool tests on the fast track *By Uwe Vollbracht*

TOOL TIPS

Tiny Applications 20130215

Toolkit for admins
Source: *http://sourceforge.net/projects/tinyapps*
License: GPLv3 and AFLv3
Alternatives: None

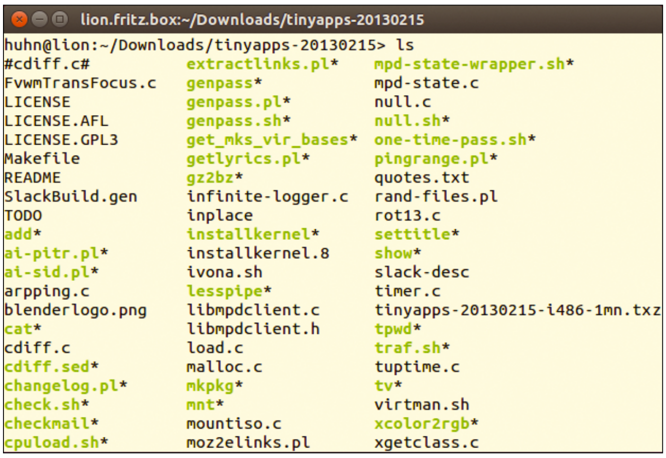

```
huhn@lion:~/Downloads/tinyapps-20130215> ls
#cdiff.c#          extractlinks.pl*    mpd-state-wrapper.sh*
FvwmTransFocus.c   genpass*            mpd-state.c
LICENSE            genpass.pl*         null.c
LICENSE.AFL        genpass.sh*         null.sh*
LICENSE.GPL3       get_mks_vir_bases*  one-time-pass.sh*
Makefile           getlyrics.pl*       pingrange.pl*
README             gz2bz*              quotes.txt
SlackBuild.gen     infinite-logger.c   rand-files.pl
TODO               inplace             rot13.c
add*               installkernel*      settitle*
ai-pitr.pl*        installkernel.8     show*
ai-sid.pl*         ivona.sh            slack-desc
arpping.c          lesspipe*           timer.c
blenderlogo.png    libmpdclient.c      tinyapps-20130215-i486-1mn.txz
cat*               libmpdclient.h      tpwd*
cdiff.c            load.c              traf.sh*
cdiff.sed*         malloc.c            tuptime.c
changelog.pl*      mkpkg*              tv*
check.sh*          mnt*                virtman.sh
checkmail*         mountiso.c          xcolor2rgb*
cpuload.sh*        moz2elinks.pl       xgetclass.c
```

The Tiny Applications collection contains more than 60 short scripts and programs that aim to make working life easier for administrators. Tips for compiling or starting the tools are mostly located in the comments; only a few of the tools come with man pages. The readme file delivers a short description of all the tools.

The collection includes helpful applications from a variety of areas. For example, `arpping` contacts IP addresses using the ARP protocol, `cdiff` colorizes `diff` output, and `cpuload.sh` shows the CPU load. Similarly useful utilities include `genpass` for creating complex passwords and `moz2elinks.pl` for converting Mozilla bookmarks for use in the ELinks text-based browser. A few tools are well-known as independent applications.

Admins who want to contribute a script to the set can contact the developer; the email address is available in the readme file.

★★☆☆☆ The Tiny Applications collection is well arranged and contains some useful aids for system operators. Points must be deducted for the documentation, which is not up to date, and the docs cover scripts that are no longer included in the most recent collection. ■■■

CertMgr 0.2.49

Manage SSL certificates
Source: *http://sourceforge.net/projects/certmgr*
License: GPLv3
Alternatives: Xca, Kleopatra

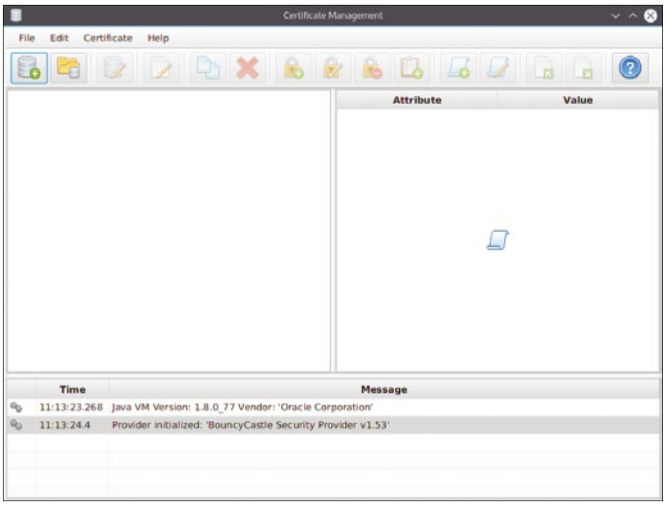

Time	Message
11:13:23.268	Java VM Version: 1.8.0_77 Vendor: 'Oracle Corporation'
11:13:24.4	Provider initialized: 'BouncyCastle Security Provider v1.53'

Many admins manage their SSL certificates in the classic way from the shell using `openssl`. If you are searching for a graphical solution, you should give the Java program CertMgr a chance. CertMgr requires at least Java 8.

The interface shows a summary of the available certificates, the data from the certificate selected, and a list of the tasks executed, along with timestamps. CertMgr manages the certificates in what it calls *stores*. Users define a validity, the algorithm, and a key length. Each new certificate that lands in this store then receives these standard values. Admins have the option to adjust the values for each certificate individually.

CertMgr offers functions for importing and exporting and can re-sign or disable existing certificates. We were unable to create new certificates with the current version in our lab tests. The Java program stubbornly displayed the progress bar and had not completed even after 45 minutes.

★★☆☆☆ Maintaining existing SSL certificates is an extremely smooth process. For creating new certificates, however, admins are better off sticking to the proven command-line tools. ■■■

Lead Image © Kheng Ho Toh, 123RF.com

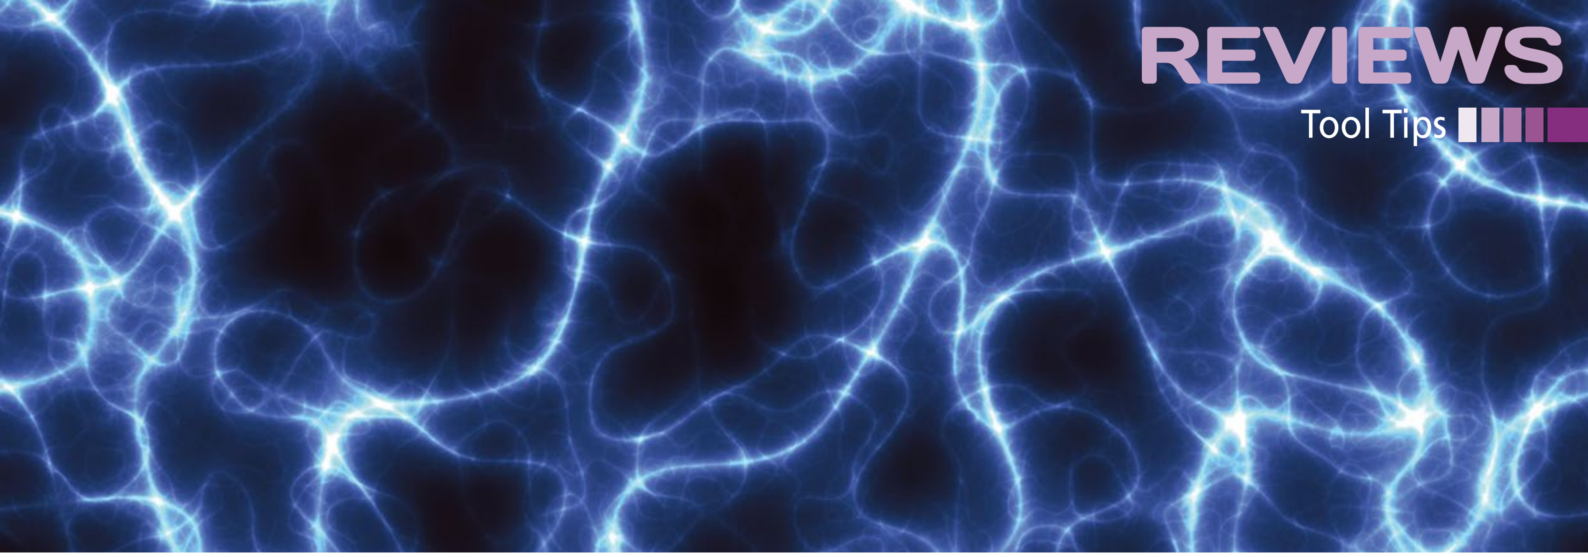

Difftree 0.5.8

Compare directories with one another
Source: *http://www.uberadmin.com/Projects/difftree*
License:GPLv3
Alternatives: Diff, Rsync

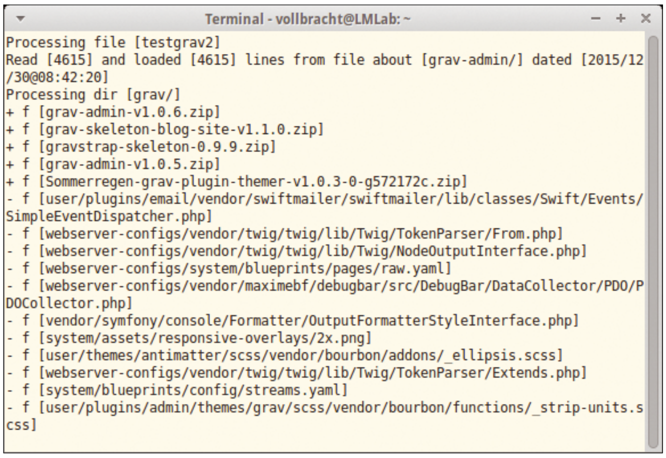

```
Processing file [testgrav2]
Read [4615] and loaded [4615] lines from file about [grav-admin/] dated [2015/12
/30@08:42:20]
Processing dir [grav/]
+ f [grav-admin-v1.0.6.zip]
+ f [grav-skeleton-blog-site-v1.1.0.zip]
+ f [gravstrap-skeleton-0.9.9.zip]
+ f [grav-admin-v1.0.5.zip]
+ f [Sommerregen-grav-plugin-themer-v1.0.3-0-g572172c.zip]
- f [user/plugins/email/vendor/swiftmailer/swiftmailer/lib/classes/Swift/Events/
SimpleEventDispatcher.php]
- f [webserver-configs/vendor/twig/twig/lib/Twig/TokenParser/From.php]
- f [webserver-configs/vendor/twig/twig/lib/Twig/NodeOutputInterface.php]
- f [webserver-configs/system/blueprints/pages/raw.yaml]
- f [webserver-configs/vendor/maximebf/debugbar/src/DebugBar/DataCollector/PDO/P
DOCollector.php]
- f [vendor/symfony/console/Formatter/OutputFormatterStyleInterface.php]
- f [system/assets/responsive-overlays/2x.png]
- f [user/themes/antimatter/scss/vendor/bourbon/addons/_ellipsis.scss]
- f [webserver-configs/vendor/twig/twig/lib/Twig/TokenParser/Extends.php]
- f [system/blueprints/config/streams.yaml]
- f [user/plugins/admin/themes/grav/scss/vendor/bourbon/functions/_strip-units.s
css]
```

Difftree compares the contents of two directories. According to statements from its developer, this command-line tool primarily wins points for its speed. Difftree is not just fast, however; it also includes a few extras. The tool can contrast several folders on request. To contrast folders, users run `dt` and then enter the directories to be compared, separating them with a space.

In its output, as well as the folder name, Difftree displays the owner, the access privileges, the size, and all timestamps. The parameter `-q` can optionally accelerate all of this, in which case Difftree only detects differences in the size and new and removed files.

The tool can calculate MD5 or SHA256 checksums on request, and the results can then be written to a text file with the `-w` command. Users can retrieve these files for comparison in future uses of the program. In combination with Cron, you can thus create your own monitoring routines in a very elegant way.

★★★★★ Difftree is convincing due to its speed, flexibility, and simple operation. Numerous examples on the project site and man page help out with the first steps. ■■■

Scriptform 1.0

Generate web forms
Source: *https://github.com/fboender/scriptform*
License: GPLv3
Alternatives: None

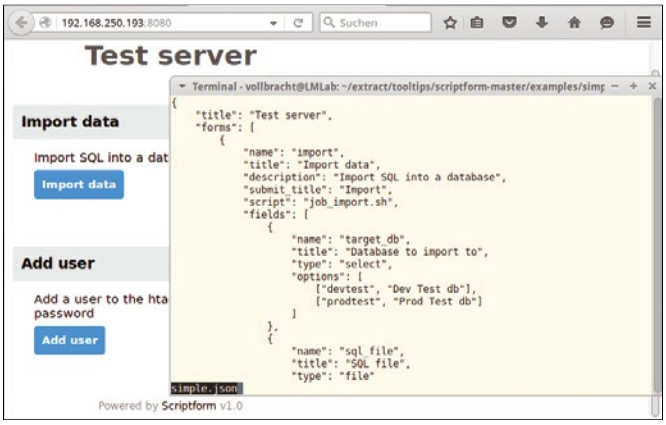

```
{
    "title": "Test server",
    "forms": [
        {
            "name": "import",
            "title": "Import data",
            "description": "Import SQL into a database",
            "submit_title": "Import",
            "script": "job_import.sh",
            "fields": [
                {
                    "name": "target_db",
                    "title": "Database to import to",
                    "type": "select",
                    "options": [
                        ["devtest", "Dev Test db"],
                        ["prodtest", "Prod Test db"]
                    ]
                },
                {
                    "name": "sql_file",
                    "title": "SQL file",
                    "type": "file"
```

If you need to create interactive websites with form fields, along with a web server, you require a scripting language like PHP or Perl. Scriptform steps up to ease these tasks for programmers. The Python tool comes with its own web server, so learning a scripting language is no longer necessary. Instead, users render the page structure in JSON format and pass in this file upon running Scriptform.

As long as users do not define any other port when opening the program, the server is accessible on `localhost:80`, which requires root privileges. Alternatively, the server runs with simple user privileges on a port higher than 1024.

In addition to simple entry fields, Scriptform also allows more complex forms with access restrictions via the use of HT Auth. Users have the option to assign preset entries to form fields. The archive contains meaningful examples, which can serve as models. Even starting programs or scripts is possible with JSON forms.

The programs run on the underlying operating system with Scriptform's privileges. The Python tool automatically generates protocols. Users can find the log files in the directory from which they started the tool.

★★★★☆ Scriptform is a useful little helper for creating simple or complex web forms. If you do want to work with the Python tool, however, you will need to become accustomed to using JSON format. ■■■

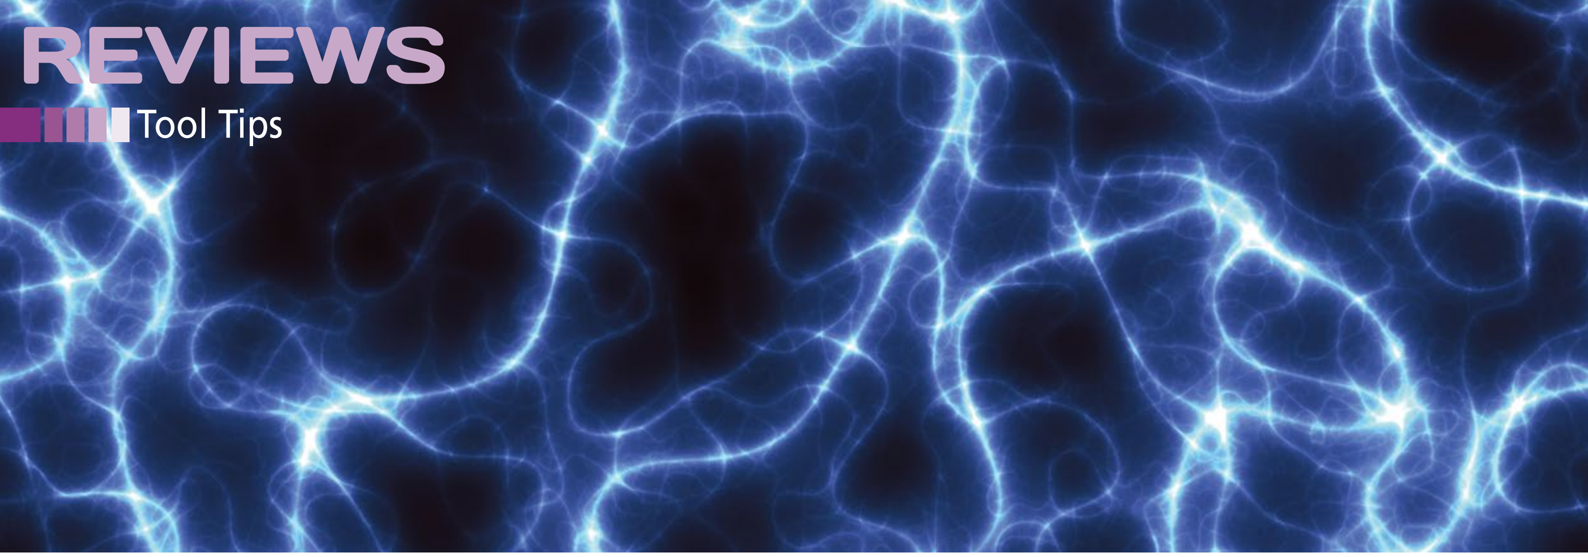

Duply 1.11.1

Console wrapper for Duplicity
Source: *http://duply.net*
License: GPLv2
Alternatives: Duplicity

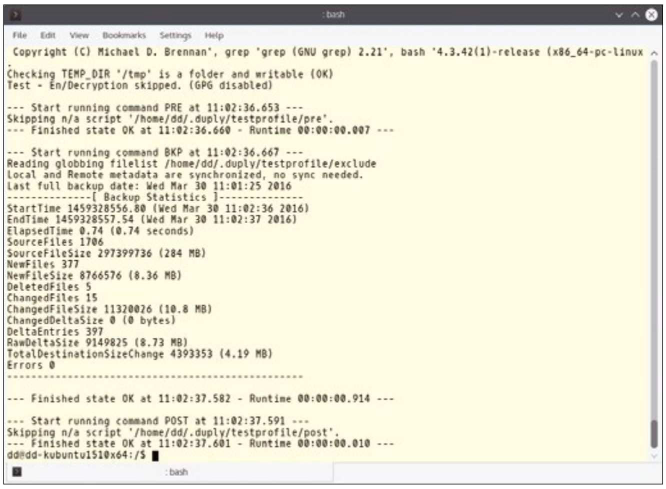

```
 Copyright (C) Michael D. Brennan', grep 'grep (GNU grep) 2.21', bash '4.3.42(1)-release (x86_64-pc-linux
.
Checking TEMP_DIR '/tmp' is a folder and writable (OK)
Test - En/Decryption skipped. (GPG disabled)

--- Start running command PRE at 11:02:36.653 ---
Skipping n/a script '/home/dd/.duply/testprofile/pre'.
--- Finished state OK at 11:02:36.660 - Runtime 00:00:00.007 ---

--- Start running command BKP at 11:02:36.667 ---
Reading globbing filelist /home/dd/.duply/testprofile/exclude
Local and Remote metadata are synchronized, no sync needed.
Last full backup date: Wed Mar 30 11:01:25 2016
--------------[ Backup Statistics ]--------------
StartTime 1459328556.80 (Wed Mar 30 11:02:36 2016)
EndTime 1459328557.54 (Wed Mar 30 11:02:37 2016)
ElapsedTime 0.74 (0.74 seconds)
SourceFiles 1706
SourceFileSize 297399736 (284 MB)
NewFiles 377
NewFileSize 8766576 (8.36 MB)
DeletedFiles 5
ChangedFiles 15
ChangedFileSize 11320026 (10.8 MB)
ChangedDeltaSize 0 (0 bytes)
DeltaEntries 397
RawDeltaSize 9149825 (8.73 MB)
TotalDestinationSizeChange 4393353 (4.19 MB)
Errors 0
-------------------------------------------------

--- Finished state OK at 11:02:37.582 - Runtime 00:00:00.914 ---

--- Start running command POST at 11:02:37.591 ---
Skipping n/a script '/home/dd/.duply/testprofile/post'.
--- Finished state OK at 11:02:37.601 - Runtime 00:00:00.010 ---
dd@dd-kubuntu1510x64:/$ 
```

The Duplicity backup tool creates encrypted backups on remote systems and is well suited for saving data in potentially insecure environments. The Duply shell script aims to simplify working with Duplicity. Duply saves recurring settings in profiles, automates the process of importing and exporting GPG keys, and lets you run scripts before or after running Duplicity.

Before the first backup, you can create a new profile with the `duply <Name> create` command. Duply generates a subdirectory with the profile name and stores in it a rudimentary configuration file, which you can later edit in the text editor. You need to enter the source and target directory and can optionally create entries for the GPG key.

If you prefer to do without encryption, you can set the variable `GPG_key` to `disabled`. The setup file also stores the maximum size of the archive and the number of full backups. `duply <Name> backup` then creates the first backup copy.

The script also supports additional parameters. For instance, `status` provides information about the available backups, and `purge` removes obsolete ones. You will find use cases and a manual on the project website, although a man page is lacking.

★★★★☆ Duply makes it significantly easier to work with the proven backup tool Duplicity and supports users who are creating profiles or managing their backup copies. ■■■

Xplico 1.1.1

Forensic network analysis
Source: *http://www.xplico.org*
License: GPLv2
Alternatives: Wireshark, Sysdig

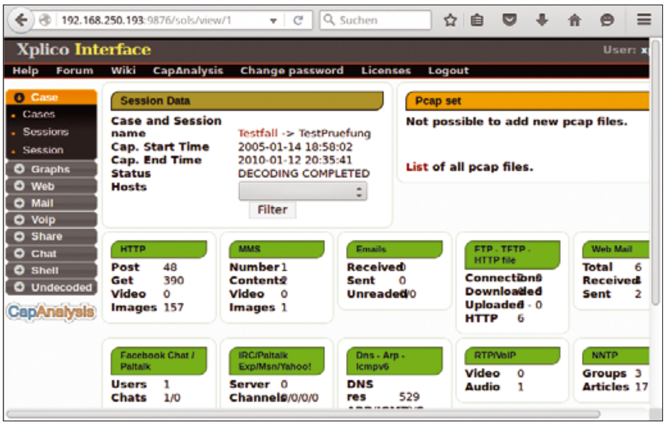

Session Data
Case and Session name Testfall -> TestPruefung
Cap. Start Time 2005-01-14 18:58:02
Cap. End Time 2010-01-12 20:35:41
Status DECODING COMPLETED
Hosts

Pcap set
Not possible to add new pcap files.
List of all pcap files.

HTTP
Post 48
Get 390
Video 0
Images 157

MMS
Number 1
Contents 2
Video 0
Images 1

Emails
Received 0
Sent 0
Unreaded 0

FTP - TFTP - HTTP file
Connections 0
Downloaded 0
Uploaded - 0
HTTP 6

Web Mail
Total 6
Received 4
Sent 2

Facebook Chat / Paltalk
Users 1
Chats 1/0

IRC/Paltalk Exp/Msn/Yahoo!
Server 0
Channels 0/0/0/0

Dns - Arp - Icmpv6
DNS
res 529

RTP/VoIP
Video 0
Audio 1

NNTP
Groups 3
Articles 17

On Linux, several programs can record data traffic, including `tcpdump`, `nmap`, Wireshark, or Snort. The applications rely on the free programming interface Pcap to capture packets directly on the network interface.

Xplico assists users during the subsequent filtering of the records. The tool extracts data from TCP and UDP packets and can operate at the command line or in a web interface. Xplico supports more than 100 different application layer protocols, including SMTP, POP, IMAP, and HTTP, as well as various messenger and VoIP protocols.

From the shell, users run `xplico` and then define an input type using `-m`. The tool accepts individual Pcap files and whole directories. Additionally, real-time analysis of a network interface is possible; the wiki explains all options in detail. Xplico stores its analyses in the directory `xdecode` and organizes them in subdirectories below that path by IP address and protocol.

In the web interface, users first create a new case and open a session that loads the selected Pcap file. Xplico shows several categories for the individual protocols. An Apache sample configuration for the Xplico interface is provided by the documentation.

★★★★★ Xplico is convincing across the board – in the shell and in the browser. The tool processes the data clearly, making it easier for users to analyze recorded network traffic. ■■■

New Magazine for Raspberry Pi Users!
PACKED with terrific tips, techniques, and projects
Together we'll conquer the world...
Adventures in the Raspberry Pi Wonderland!
Your roadmap to OPEN HARDWARE treasure
Got to get home! I'm missing my Raspberry Pi!
Sure I trust my Raspberry Pi...
raspberry pi GEEK
Photo Tricks
Time-lapse video and motion detection with the Rasp Pi camera
BrickPi
Raspberry Pi meets Lego Mindstorms
Talk to your Pi
with a TV remote
Emergency PC
Outfit your Rasp Pi as a backup desktop system
WIND POWER
Power your Pi off-grid with an inexpensive wind turbine
COOL PROJECTS:
HOME AUTOMATION
Lights, cameras, appliances ... control and monitor up to 50 home devices
DISCOVER THE NEW

Adventures in the Raspberry Pi Wonderland!
Your roadmap to OPEN HARDWARE treasure
Please tell me more about this most astonishing device
And now, Dearie, how about a little Pi?
I got a new Pi and I'm feeling the music!
Don't you touch my Raspberry Pi!
raspberry pi GEEK
3D PRINTING
Awesome cat toy launcher gets a 3D-printed frame
Odroid
NoIR
Turn your Rasp Pi into an infrared camera
Pydio
Create a handy home file server
Turn your Pi into an remote control device
20 Cool Rasp Pi extensions
Steampunk Name
PiFace Control and Display
Put an LCD screen, IR receiver, and buttons right on top of your Pi
More Recon on
BPi-R1
A Banana Pi-based
Set up a Minecraft server on your Pi

Camouflaged operating system: Whonix

Anonymous Traveler

The Whonix desktop operating system lets you use the web without revealing your identity. *By Thomas Joos*

Many Internet users want to protect their privacy on the Internet, without disclosing personal information unnecessarily. The special Linux distribution Whonix [1], which incorporates The Onion Router (Tor) network, lets you do so for free.

If you want to try out Whonix, your best bet is to install it on a virtual machine (VM). Although physical hardware would work just as well – and you don't even need particularly new or powerful hardware – you would need two machines, because Whonix consistently separates the Internet physically from the computer on which you work, either with the use of two VMs or two separate physical systems. It is easy to set up and use Whonix: You only need to import two VMs, and a wizard then connects them to the Tor network.

The Architecture

Two VMs or two computers form the basis of the Whonix Linux distribution. One machine used as the connection gateway to the Tor network [2] is known as the Whonix-Gateway on the Whonix network. The other machine accommodates the applications with which you work. To begin, you set up the gateway, and it then sets up the connection to the Internet instead of connecting directly to the Internet; the wizard can also connect the gateway via a proxy server.

Because the workstation is on a separate network, Whonix keeps it from being contaminated by viruses or other malware and keeps your IP address from becoming public. The Whonix-Workstation can only access the Inter-

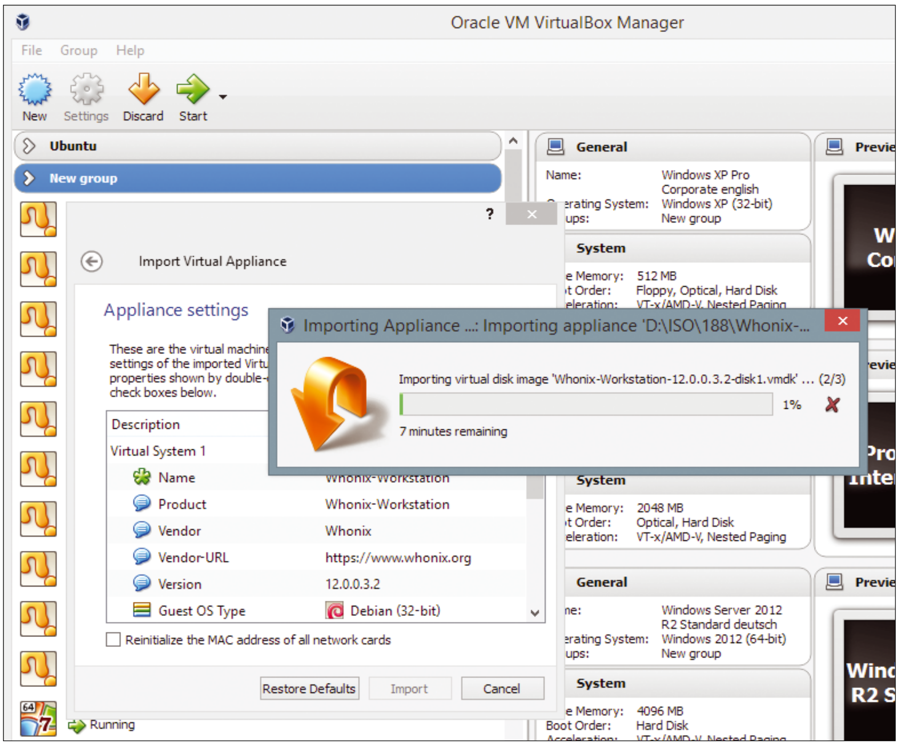

Oracle VM VirtualBox Manager

Import Virtual Appliance

Appliance settings

These are the virtual machine settings of the imported Virtu… properties shown by double-… check boxes below.

Description

Virtual System 1

Name	Whonix-Workstation
Product	Whonix-Workstation
Vendor	Whonix
Vendor-URL	https://www.whonix.org
Version	12.0.0.3.2
Guest OS Type	Debian (32-bit)

□ Reinitialize the MAC address of all network cards

Importing Appliance ...: Importing appliance 'D:\ISO\188\Whonix-...

Importing virtual disk image 'Whonix-Workstation-12.0.0.3.2-disk1.vmdk' ... (2/3)

1%

7 minutes remaining

Name: Windows XP Pro Corporate english
…erating System: Windows XP (32-bit)
…ups: New group

…e Memory: 512 MB
…t Order: Floppy, Optical, Hard Disk
…eleration: VT-x/AMD-V, Nested Paging

…e Memory: 2048 MB
…t Order: Optical, Hard Disk
…eleration: VT-x/AMD-V, Nested Paging

…ne: Windows Server 2012 R2 Standard deutsch
…erating System: Windows 2012 (64-bit)
…ups: New group

…e Memory: 4096 MB
Boot Order: Hard Disk

Figure 1: Importing Whonix VMs into VirtualBox.

Lead Image © Mohamad Razi Bin Husin, 123RF.com

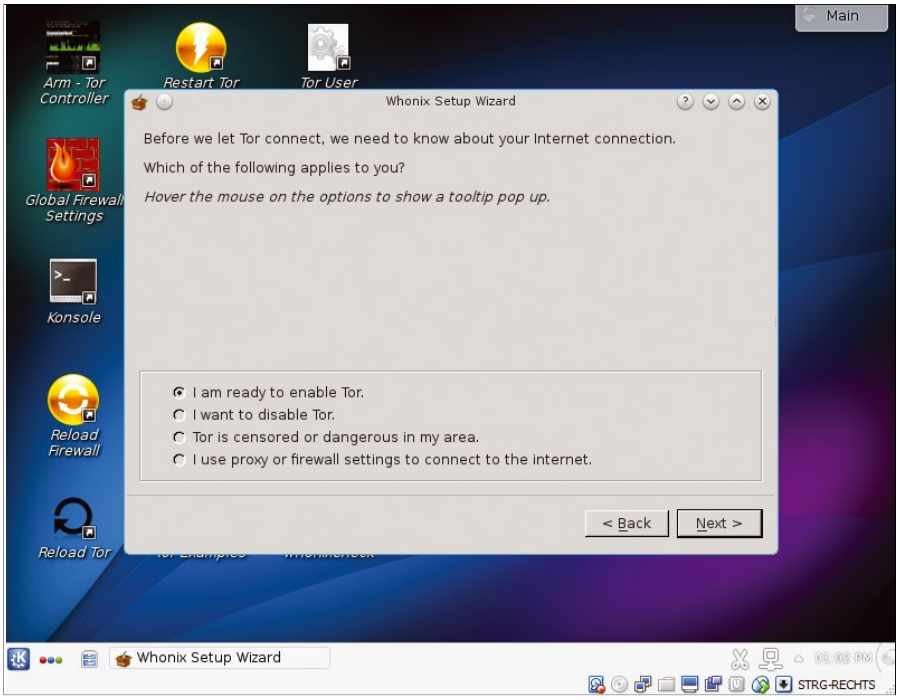

Whonix Setup Wizard

Before we let Tor connect, we need to know about your Internet connection.

Which of the following applies to you?

Hover the mouse on the options to show a tooltip pop up.

● I am ready to enable Tor.
○ I want to disable Tor.
○ Tor is censored or dangerous in my area.
○ I use proxy or firewall settings to connect to the internet.

< Back | Next >

Figure 2: During installation, Whonix establishes a connection to the Tor network to guarantee anonymous surfing on the Internet.

net via the Tor router installed on the Whonix-Gateway.

Installation and Setup

Qubes, KVM, and VirtualBox can virtualize the environment; unfortunately, VMware vSphere and Qemu cannot. The easiest way to install the two VMs, both available as OVA files, is in VirtualBox. To do so, you only need to import an appliance (Figure 1) by setting up the gateway in the first step and the workstation in the second step.

After the installing the environment, a setup wizard helps adapt the two machines to your requirements, where you can change such settings as the number of processors for the VM or the size of available memory. When first set up, Whonix launches a setup wizard that creates the connection to the Tor network (Figure 2). Also, you can define here whether Whonix should update automatically in the future.

In the course of the setup, you can also decide which repository to use. If you will be deploying Whonix in a production environment, the best choice is the *Whonix Stable Repository*. Alternatively, you can choose the *Whonix Testers Repository* or the *Whonix Developers Repository*.

After all the options are set up, the connection to the Tor network is opened automatically. If necessary, Whonix also downloads updates in the background. To access the latest versions, it is advisable to update the repositories first. On Whonix, you can do this by typing:

```
apt-get update
apt-get upgrade
```

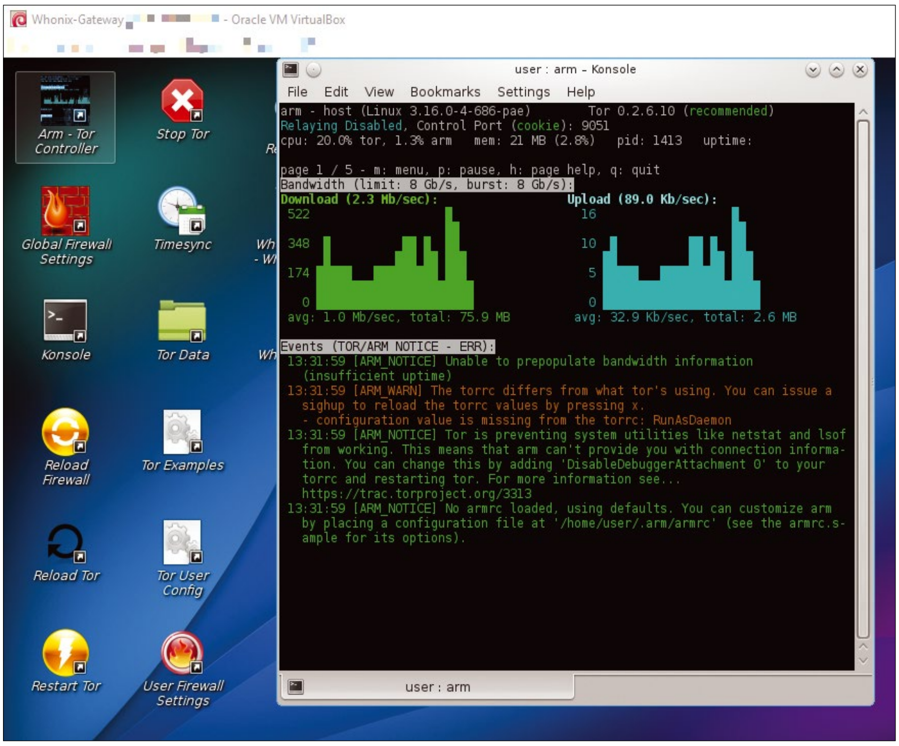

```
arm - host (Linux 3.16.0-4-686-pae)        Tor 0.2.6.10 (recommended)
Relaying Disabled, Control Port (cookie): 9051
cpu: 20.0% tor, 1.3% arm   mem: 21 MB (2.8%)   pid: 1413   uptime:

page 1 / 5 - m: menu, p: pause, h: page help, q: quit
Bandwidth (limit: 8 Gb/s, burst: 8 Gb/s):
Download (2.3 Mb/sec):                    Upload (89.0 Kb/sec):
 522                                       16
 348                                       10
 174                                        5
   0                                        0
 avg: 1.0 Mb/sec, total: 75.9 MB           avg: 32.9 Kb/sec, total: 2.6 MB

Events (TOR/ARM NOTICE - ERR):
 13:31:59 [ARM_NOTICE] Unable to prepopulate bandwidth information
   (insufficient uptime)
 13:31:59 [ARM_WARN] The torrc differs from what tor's using. You can issue a
   sighup to reload the torrc values by pressing x.
   - configuration value is missing from the torrc: RunAsDaemon
 13:31:59 [ARM_NOTICE] Tor is preventing system utilities like netstat and lsof
   from working. This means that arm can't provide you with connection informa-
   tion. You can change this by adding 'DisableDebuggerAttachment 0' to your
   torrc and restarting tor. For more information see...
   https://trac.torproject.org/3313
 13:31:59 [ARM_NOTICE] No armrc loaded, using defaults. You can customize arm
   by placing a configuration file at '/home/user/.arm/armrc' (see the armrc.s-
   ample for its options).
```

Figure 3: Via the Whonix-Gateway, you can monitor the use of the Tor Internet connection from the Whonix-Workstation.

The gateway needs to be running for you to use Whonix; you can iconize the window without worry because there's nothing to configure.

Clicking the *WhonixCheck* icon makes sure everything is working and that the gateway is up to date and connected to the Tor network. If several workstations are connected to the Whonix-Gateway, the traffic can be monitored with the *Arm-Tor Controller* desktop shortcut. When launched, the tool shows statistics about current uploads and downloads (Figure 3).

Whonix integrates a firewall that can be set up with the *Global Firewall Settings* desktop shortcut. The settings are password protected – the default password is *changeme* – and configuration changes are by finalized by clicking on the *Reload Firewall* desktop shortcut.

With the *Whonix Setup* icon, you can launch the wizard for connecting to the Tor network, which is necessary, for example, if you want to use a different Internet gateway for the connection. It is also possible to connect the gateway to a proxy server through the wizard.

Working with Whonix

Once the gateway is running, everything else happens on the Whonix-Workstation, which is also imported

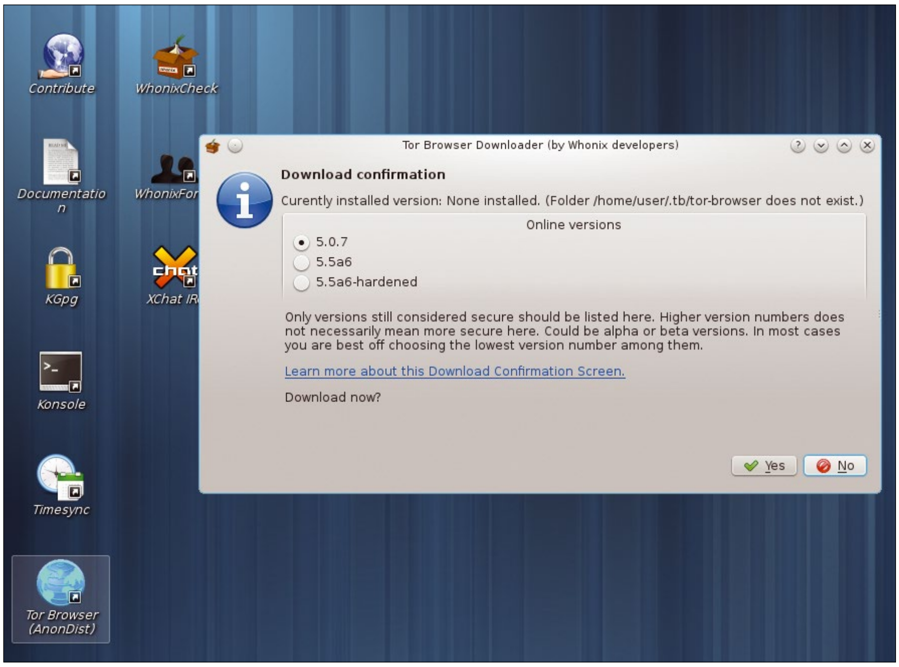

Tor Browser Downloader (by Whonix developers)

Download confirmation

Curently installed version: None installed. (Folder /home/user/.tb/tor-browser does not exist.)

Online versions

● 5.0.7
○ 5.5a6
○ 5.5a6-hardened

Only versions still considered secure should be listed here. Higher version numbers does not necessarily mean more secure here. Could be alpha or beta versions. In most cases you are best off choosing the lowest version number among them.

Learn more about this Download Confirmation Screen.

Download now?

Yes No

Figure 4: When you first start the Tor browser, the application downloads from the Internet and then installs automatically.

into VirtualBox as a VM, just like the gateway. To work without interruption, you will want to assign the workstation more virtual CPUs and more memory. The default username is *user* and the password, again, is *changeme*. The Tor browser downloads automatically when you first start the workstation and proceeds to install itself (Figure 4).

After launching the browser, you can see the successful connection to Tor at top right. Also, you can see that the "No Script" extension is installed, which prevents scripts running on Internet pages without permission.

In addition to your own workstation opening connections to the Internet via the Whonix-Gateway, any computer or virtual machine can use this gateway for the same purpose. For this to happen, the gateway has two network adapters. One of the adapters communicates with the public Internet, and the other adapter is for private communication with the connected workstations. Through this network interface, multiple VMs or multiple physical computers can connect to the Internet via the Whonix-Gateway without problem.

Tails Alternative

One alternative to Whonix is Tails [3]. With this Debian-based distro, you can access the Internet anonymously via a Live DVD or USB stick, but without the same level of security because Tails lacks the physical separation of the components that Whonix offers.

Because Tails, in contrast to Whonix, runs as a Live system, it is up and running even faster and is already protected from attacks at startup because it runs from a DVD (Figure 5). The download size is approximately 910MB, so the system is not intended to replace your current operating system; rather, it serves as a mobile surfing environment or secure environment for sensitive missions.

You can activate Windows Stealth Mode in Tails, during which time the environment behaves just like a Windows 8 system on the Internet. The connection to the Internet is set up via the Tor network. Like Whonix, Tails does not allow direct Internet connections. Programs are accessed via *Applications* in the Start menu.

In addition to the Tor browser, the email tool Claws Mail and several instant messaging programs are available on the Tails desktop. Tails can also be used as a secure online banking solution. The distribution also offers a screen keyboard.

By pressing the Tor icon and selecting *Preferences*, you can set up a proxy server in Tails. With the help of the Live DVD, users can install the system on a USB stick or an SD card. To do this, use the *Applications* | *Tails* | *Tails Installer* menu item.

Rating

Users or administrators who need a secure and anonymous Internet connection and want to use it not just once, but permanently, can take advantage of the options that Whonix offers. The security distribution can be set up very efficiently with VirtualBox. Of course, the ability to install the gateway on a physical computer is also useful on professional networks. Multiple workstations can thus be connected securely to the Internet.

For those who do not have comprehensive knowledge of Linux or want to use a secure Linux system on Windows workstations, Whonix is perfectly suited to the task. After the initial setup, through which an easy-to-use wizard guides you, you can securely and anonymously surf the Internet. Additionally, the gateway offers a connection option for other Linux distributions, whether virtualized or physical, and you do not need a license. ■■■

INFO

[1] Whonix: *https://www.whonix.org*

[2] Tor: *https://www.torproject.org*

[3] Tails: *https://tails.boum.org/index.en.html*

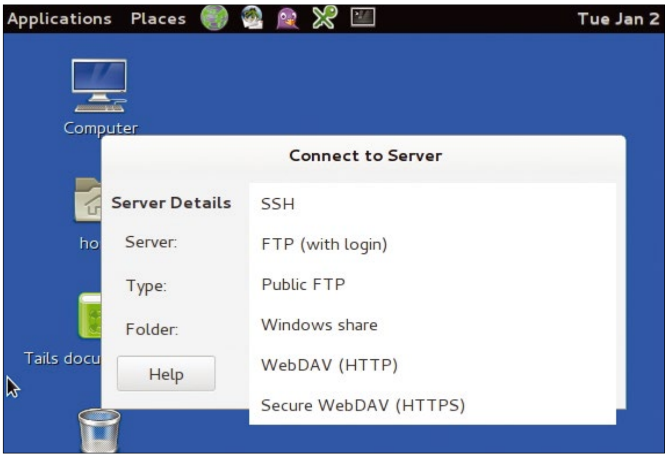

Figure 5: Tails, an alternative to Whonix, is ready faster, but it does without the two-component safety principle.

YUNEEC
ELECTRIC AVIATION

InterDrone

Encryption with VeraCrypt

Hidden

The VeraCrypt encryption software comes with a handy graphical interface, and the ability to hide a container in an encrypted volume adds a unique professional feature: plausibly deniable encryption. *By Peter Kreußel*

When the TrueCrypt developers dissuaded people from further use of its software with an ominous security warning [1], many users were confused and concerned about their privacy, especially in the Windows camp, where TrueCrypt was a popular open source encryption solution (see the "TrueCrypt" box).

In the meantime, TrueCrypt fork VeraCrypt [2], which dates back to 2013, has inherited its predecessor's followers and introduced Linux support in 2014. Given that the Linux kernel already ciphers directories or entire partitions, why would Linux users want to embrace a program with a black spot in its history? VeraCrypt provides some solid reasons for doing so.

Plausible Reasons

One strong motive for the use of VeraCrypt is its guaranteed "plausibly deniable encryption": The encrypted container can embed a hidden inner container (Figure 1). Should you ever be forced to

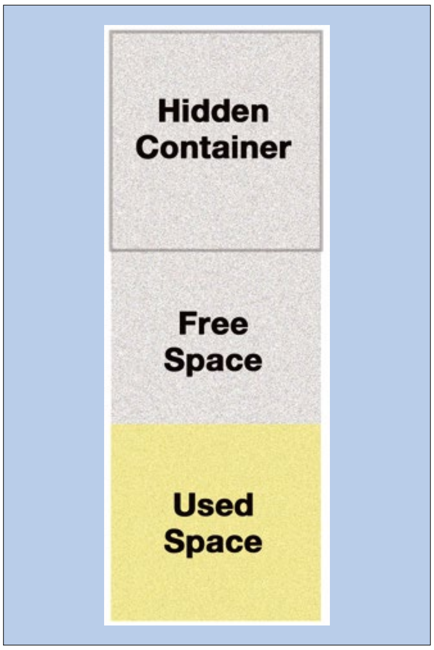

Figure 1: In free space in a VeraCrypt container pre-formatted with white noise, you can hide another container. Without knowledge of a separate password, users cannot view metadata about its extent or the encrypted container itself.

Lead Image © bowie15, 123RF.com

TRUECRYPT

By the spring of 2015, the open source and free encryption software TrueCrypt stood alone. Some users, however, were disturbed because the developers were never identified, leading to speculation. At the end of May 2015, the developers terminated the project and advised users to switch to non-open-source Windows onboard encryption with the words, "Using TrueCrypt is not secure as it may contain unfixed security issues."

Clarity about the actual security of the software was achieved by an independent security audit [3]. However, except for some problems with Windows drivers, the examiners only objected to the low number of hash iterations required to derive the key, which was too small for the computing power of its day. This failed to slow down attackers attempting to brute force passwords; containers with weak passwords were therefore easier to crack. VeraCrypt improved this point promptly, but it also made mounting encrypted objects take considerably more time.

Google employees finally found two critical vulnerabilities that were not directly related to encryption, allowing attackers on Windows [4] – given certain conditions [5] – to gain administrative privileges. The Windows version of VeraCrypt ironed out these weaknesses in the meantime.

PLAUSIBLE DENIABILITY

Some countries (e.g., the UK) by law compel computer owners to disclose their passwords on demand for encrypted data [6]. With the standard Linux encryption tools dm-crypt/LUKS [7], you could be in trouble. A partition encrypted in this way can be identified readily, and the user would not be able to deny its existence (Figure 4) and thus the presence of encrypted data.

The same is true for normal VeraCrypt volumes: Good encryption does not allow any conclusions as to the encrypted data; the content of a container thus looks from the outside like a random numeric sequence. By contrast, unencrypted data (text, video, images) always exhibits certain regularities. The difference can be demonstrated statistically, thus revealing encrypted files.

Precisely the quality that reveals the existence of encrypted filesystems gives VeraCrypt the ability to create a secure hiding place in an inner container. The inner container looks like a random bit sequence and transitions seamlessly and undetectably past statistical analysis into the outer container.

In practice, when creating the outer container, VeraCrypt first overwrites the intended disk space with a random number sequence. A second step embeds a hidden container with its own password. When opening a VeraCrypt volume, you then decide with the choice of a password whether to unlock the outer or inner container.

In the outer container, you will want to store a sufficient number of alibi files as camouflage. The inner container hides in the free space, remaining invisible, unless you know the corresponding password. This is also true of VeraCrypt itself: The content of the outer container will overwrite the hidden volume without warning if it becomes too big. To prevent this, you enter a kind of mixed mode in which you enter the passwords of both containers: Only then will the software detect the position of the inner container and prevent overwriting.

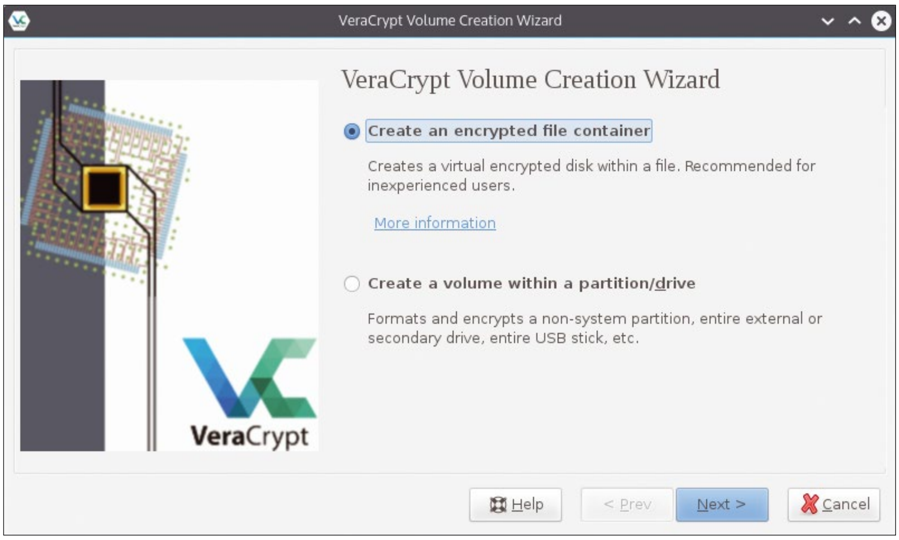

VeraCrypt Volume Creation Wizard

● Create an encrypted file container

Creates a virtual encrypted disk within a file. Recommended for inexperienced users.

More information

○ Create a volume within a partition/drive

Formats and encrypts a non-system partition, entire external or secondary drive, entire USB stick, etc.

Figure 2: The intuitively designed VeraCrypt dialog lets you create encrypted containers, even without prior knowledge or studying the manual.

reveal your encryption password, you could do so for the outer container only (see the box "Plausible Deniability").

Without the second password, you cannot even prove the existence of an inner container. After unlocking the outer container, it appears to be a blank space. Information relating to its extent is encrypted with the second password in a special reserved memory space. The metadata, like the entire inner container, looks like random values before you unlock them separately.

Although standard Linux tools dm-crypt and eCryptfs [8] are well suited for integration with the operating system (e.g., to encrypt the entire system or the home partition), in contrast, the VeraCrypt GUI lends itself to opening containers for particularly security-critical files as needed. To do this, you create a file-based container with a few mouse clicks (Figure 2); the container can be used not only on Linux, but also on Mac OS X and Windows.

The simple user interface (Figure 3) also handles the task of mounting encrypted volumes, which the program mounts transparently in the filesystem below `/mnt` or `/media`. Alternatively, VeraCrypt encrypts entire partitions. The command-line option `--text` eliminates the need to start the graphical user interface;

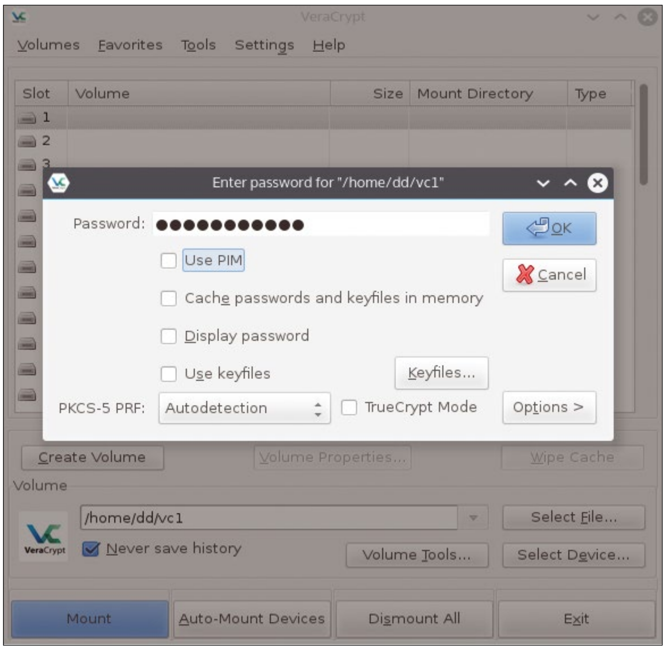

Slot	Volume	Size	Mount Directory	Type
1				
2				
3				

Enter password for "/home/dd/vc1"

Password: ●●●●●●●●●●●

□ Use PIM

□ Cache passwords and keyfiles in memory

□ Display password

□ Use keyfiles

PKCS-5 PRF: Autodetection

□ TrueCrypt Mode

Volume

/home/dd/vc1

☑ Never save history

Figure 3: The VeraCrypt GUI mounts the encrypted volume and provides access to all other functions.

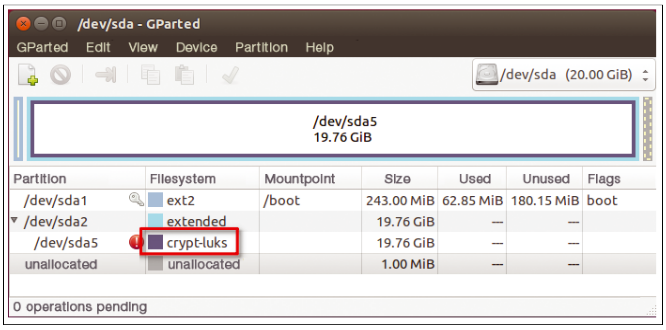

Partition	Filesystem	Mountpoint	Size	Used	Unused	Flags
/dev/sda1	ext2	/boot	243.00 MiB	62.85 MiB	180.15 MiB	boot
/dev/sda2	extended		19.76 GiB	—	—	
/dev/sda5	crypt-luks		19.76 GiB	—	—	
unallocated	unallocated		1.00 MiB	—	—	

0 operations pending

Figure 4: Partitions encrypted with the Linux on-board solutions dm-crypt/LUKS appear in the partition table as such.

you can control all the functions from the command line or with a script.

Secure?

Features like plausibly deniable encryption or a practical GUI are of little use if the underlying encryption method proves to be insecure. As always with security issues, you can only follow circumstantial evidence with known factors; potentially unknown vulnerabilities remain undetected.

To the best of my knowledge and belief, the security of VeraCrypt looks good. The software has a long history in open source: It is based on TrueCrypt, which in turn was based on Encryption for the Masses (E4M), launched in 1997 [9]. The TrueCrypt heritage might initially cause some concern, but the VeraCrypt developers understandably explain how they ironed out its known vulnerabilities [10]; in any case, they only affected the Linux version in part. The developers also subjected the code to two static analyses, which revealed some critical programming errors. An expert audit of VeraCrypt itself is still pending.

The software is available from SourceForge [11] in the form of an installer, which only installs a binary and some additional files. As always with security-related software, it pays to verify the integrity of the installation files with `sha512sum`. Compiling the software turns out to be difficult at present: The current openSUSE and Ubuntu releases include a compiler that uses the new C++ ABI by default, but not all of the utilities you need are available in this format.

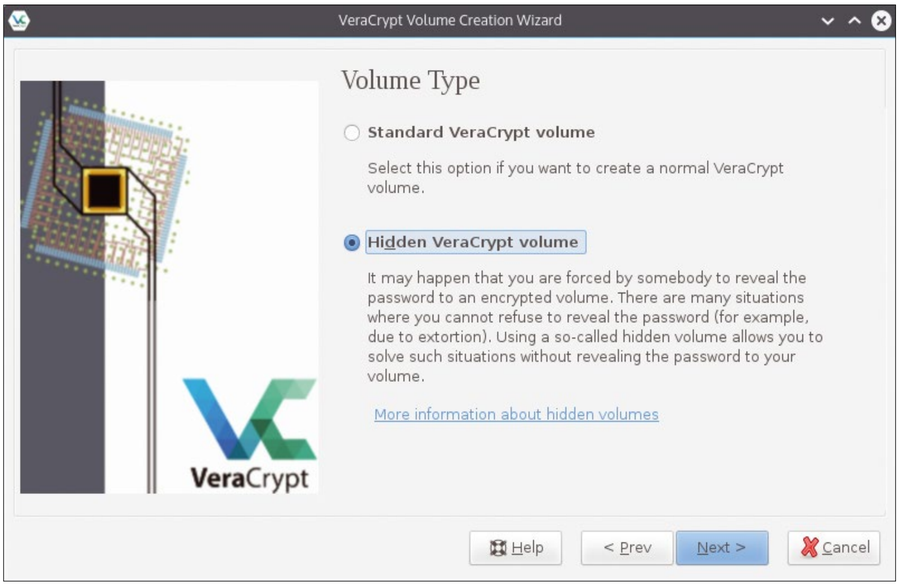
VeraCrypt Volume Creation Wizard

Volume Type

○ Standard VeraCrypt volume

Select this option if you want to create a normal VeraCrypt volume.

● Hidden VeraCrypt volume

It may happen that you are forced by somebody to reveal the password to an encrypted volume. There are many situations where you cannot refuse to reveal the password (for example, due to extortion). Using a so-called hidden volume allows you to solve such situations without revealing the password to your volume.

More information about hidden volumes

Figure 5: After choosing the *Hidden VeraCrypt volume* option as the Volume Type, the wizard generates both the outer and an embedded hidden container in one pass.

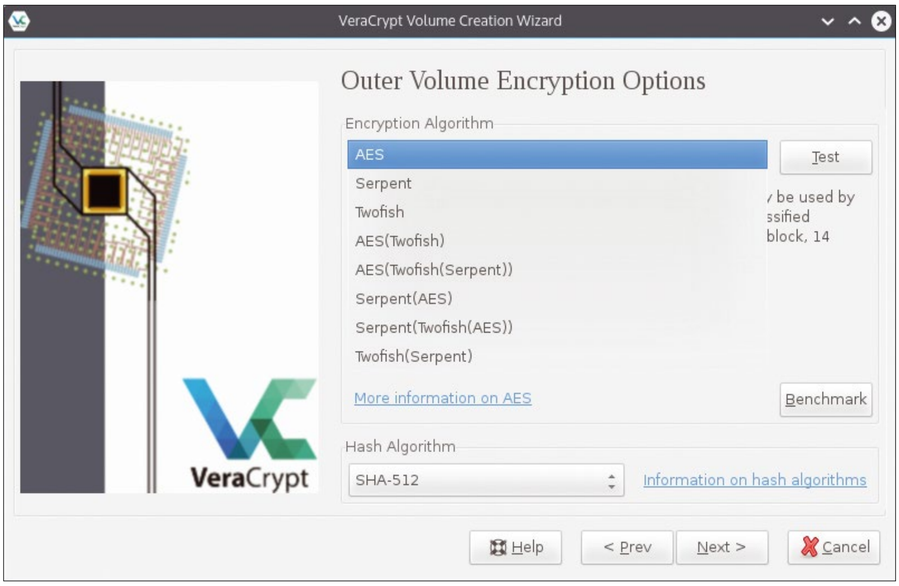
VeraCrypt Volume Creation Wizard

Outer Volume Encryption Options

Encryption Algorithm

AES
Serpent
Twofish
AES(Twofish)
AES(Twofish(Serpent))
Serpent(AES)
Serpent(Twofish(AES))
Twofish(Serpent)

/ be used by
ssified
block, 14

More information on AES

Hash Algorithm

SHA-512

Information on hash algorithms

Figure 6: VeraCrypt lets you choose between various proven encryption algorithms.

Handy

The current documentation [12] for VeraCrypt leaves no questions unanswered. The basic functions of the software can be used without reading the manual anyway, thanks to the intuitively designed graphical interface. The *Create Volume* button starts the Volume Creation Wizard. You first need to decide whether you want to create a container or encrypt a hard disk partition. Then the wizard asks whether you want to create a standard volume or a container with an embedded hidden partition for plausibly deniable encryption (Figure 5).

You always need to create a standard outer container. To do so, stipulate a file path in which the software will create the container or the device file of a disk partition (e.g., `/dev/sda3`). In the Encryption Options dialog, the Encryption Algorithm default is *AES* and the Hash Algorithm default is *SHA-512*, which offer good runtime performance and impeccable security features from today's perspective.

Alternative encryption algorithms (Figure 6) are available in line with the common practice in cryptography of keeping all sensitive components interchangeable. Should future attack vectors compromise the current secure process, you can then change the algorithm but continue using the familiar software.

After entering the desired volume size, type your password twice or select one or more keyfiles, which may consist of any number of files. For the filesystem, VeraCrypt uses the system global default, *FAT*; more sophisticated filesystems, such as NTFS and ext2/3/4, are also available for use. Of course, selecting an ext filesystem will impair compatibility with Windows. In the final dialog box, click on *Format* to start the process of generating the container.

Key Function

To mount a volume, just access the container or device file of a disk partition in the main dialog by clicking *Select File* or *Select Device*. In the list box in the upper half of the window, you then select a

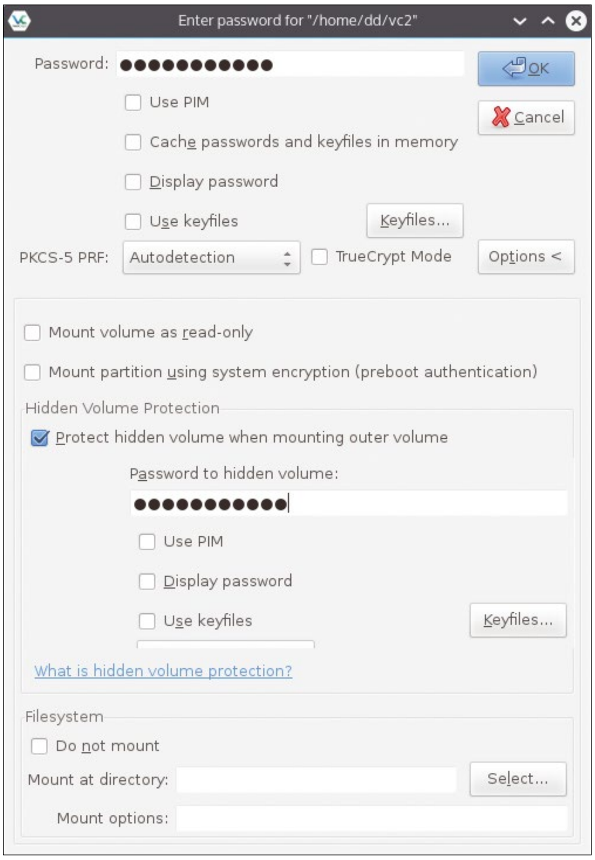

Enter password for "/home/dd/vc2"

Password: ●●●●●●●●●●● OK

□ Use PIM Cancel

□ Cache passwords and keyfiles in memory

□ Display password

□ Use keyfiles Keyfiles...

PKCS-5 PRF: Autodetection □ TrueCrypt Mode Options <

□ Mount volume as read-only

□ Mount partition using system encryption (preboot authentication)

Hidden Volume Protection

☑ Protect hidden volume when mounting outer volume

Password to hidden volume:

●●●●●●●●●●●●

□ Use PIM

□ Display password

□ Use keyfiles Keyfiles...

What is hidden volume protection?

Filesystem

□ Do not mount

Mount at directory: Select...

Mount options:

Figure 7: You need to know the passwords of both the outer and the hidden containers to enable the protected mode that prevents the inner container from being overwritten.

volume under Slot. On Linux, you define a slot number (`<n>`) in the mountpoint `/media/veracrypt<n>`, and on Windows, a drive letter. If you want to mount the object in a different directory, click *Options* in the Mount dialog, and in *Mount at directory,* enter the target path.

Then specify the password or keyfile. Whether you mount the outer container or – if available – the inner container, is decided exclusively by the password and keyfile. As explained in the “Plausible Deniability” box, you will always use the inner volume in daily operation, but if you mount an outer container with an embedded inner container, in *Options* you need to check the *Protect hidden volume when mounting outer volume* box and enter the password for the inner container, too (Figure 7).

Clicking *Dismount* unmounts the device, and VeraCrypt explicitly clears the password from memory. Besides this, VeraCrypt offers a *Dismount All* button, which closes all open containers as soon as possible. Because the software needs the password constantly during operation, you cannot prevent it remaining in memory while containers are active. This also applies to both of the Linux on-board solutions.

If the system crashes with open containers, you face a residual risk: In case of memory shortage, the operating system writes sections of the main memory to a swap file on the hard disk. Under certain circumstances, the password for active VeraCrypt volumes could thus survive a system power off – if you suspend to disk, this happens in any case. However, this usually only proves to be risky if it falls into the hands of a forensics professional with appropriate knowledge of the system.

You can avoid this risk with full system encryption (e.g., as set up by the Linux installer on Ubuntu) because this also encrypts the swap partition.

Systematically Hidden

On Windows, VeraCrypt’s capabilities go much further: In the hidden inner container on Windows you can install a second hidden operating system whose existence is not demonstrable (Figure 8). The special bootloader used for this does not work with Linux, and a posting from the VeraCrypt forum suggests that this situation is not likely to change any time in the near future [13].

Of course, the free operating systems offer many well-known solutions: `cryptsetup`, a tool that offers full system encryption and comes with many Linux installers, has been able to unlock VeraCrypt volumes since version 1.6.7 from spring 2015. However, both Ubuntu 15.10 and openSUSE 42.1 still use older versions; only Arch Linux already uses the current Cryptsetup release. It relies on a shell script to unlock the root filesystem embedded in the initial ramdisk, which is a file archive the kernel mounts provisionally as root at bootup. This archive contains kernel modules for the filesystems and a shell. Scripts create the conditions for mounting the final root filesystem.

To cooperate with VeraCrypt volumes, you need to extend the standard version of this script; for users with shell skills, this is not too difficult. The initial ramdisks of other distributions [14] work in a similar way to those for Arch, which is why the process can be transferred in principle to this script [15].

Listing 1 shows a section of the `/usr/lib/initcpio/encrypt` shell script, which

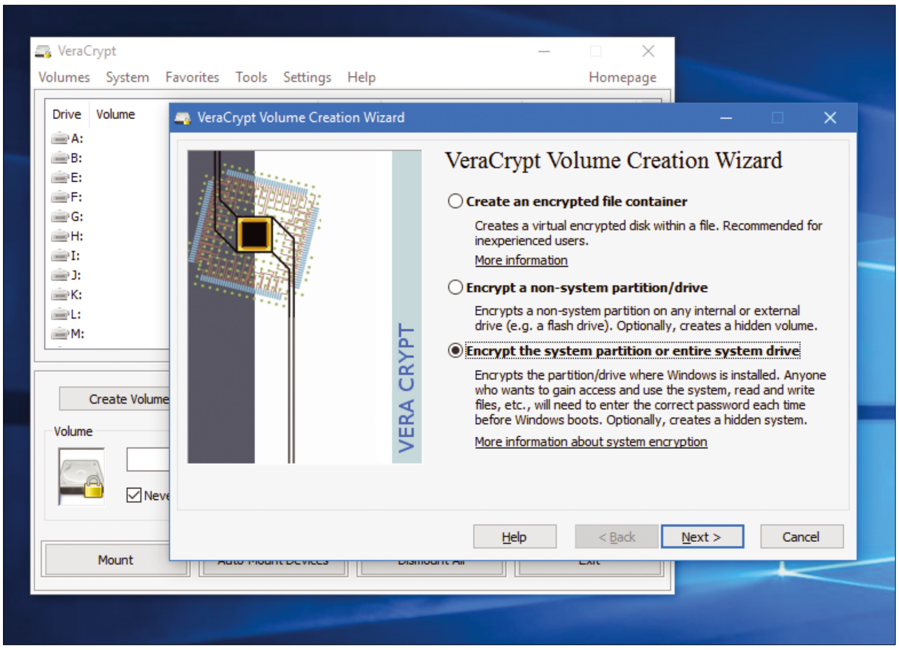

VeraCrypt
Volumes System Favorites Tools Settings Help Homepage
Drive Volume
A: B: E: F: G: H: I: J: K: L: M:
Create Volume
Volume
Never
Mount Auto-Mount Devices Dismount All Exit
VERA CRYPT

VeraCrypt Volume Creation Wizard

VeraCrypt Volume Creation Wizard

○ Create an encrypted file container

Creates a virtual encrypted disk within a file. Recommended for inexperienced users.

More information

○ Encrypt a non-system partition/drive

Encrypts a non-system partition on any internal or external drive (e.g. a flash drive). Optionally, creates a hidden volume.

● Encrypt the system partition or entire system drive

Encrypts the partition/drive where Windows is installed. Anyone who wants to gain access and use the system, read and write files, etc., will need to enter the correct password each time before Windows boots. Optionally, creates a hidden system.

More information about system encryption

Help < Back Next > Cancel

Figure 8: On Windows, VeraCrypt supports encrypting a complete operating system in a hidden partition. The Linux version does not have such a function.

LISTING 1: Unlocking LUKS

```
01 if [ ${dopassphrase} -gt 0 ]; then
02   echo ""
03   echo "A password is required to access the ${cryptname} volume:"
04   #loop until we get a real password
05   while ! eval cryptsetup open --type luks ${resolved} ${cryptname}
     ${cryptargs} ${CSQUIET}; do
06     sleep 2;
07   done
08 fi
```

asks for the password that will unlock the root filesystem. Line 5 shows the call to dm-crypt via `cryptsetup` with the parameter `--type luks`, which is the most common encryption format for Linux partitions.

To use VeraCrypt instead of LUKS, you just need to replace the parameter `--type luks` with `--type tcrypt --veracrypt` for a TrueCrypt container with the VeraCrypt sub-format. To ensure that encryption with LUKS still works, all you need is a simple `if` clause to check for the existence of the `vera=1` boot parameter and store the container type in a variable (Listing 2).

LISTING 2: Using VeraCrypt

```
if [ -n "${vera}" ]; then
  type="--type tcrypt --veracrypt"
else
  type="--luks"
fi
```

In the Cryptsetup call, simply replace `--type luks` with the `${type}` variable. Dm-crypt does not automatically search for a hidden container; instead, you need to stipulate the additional `--tcrypt-hidden` option. The script thus handles the task of checking for both versions. In the tuned version, it tries the first password entered for the outer container. If it fails to unlock the container, it again prompts for a password, but now tries this on the inner container, then again on the outer container, and so on.

This process is not particularly convenient, especially because it takes time for VeraCrypt to reject an incorrect password. However, at least this process does not reveal any evidence of the hidden container. You need to replace the code in Listing 1 with the code shown in Listing 3 for this. You can download the complete `vencrypt` script file for Arch Linux online [16].

LISTING 3: Unlocking VeraCrypt

```
#loop until we get a real password
hidden=""
while ! eval cryptsetup open ${type} ${hidden} -T1 ${resolved} ${cryptname}
  ${cryptargs} ${CSQUIET}; do
  if [ -n ${vera} ] && [ "${hidden}" == '' ]; then
    hidden="--tcrypt-hidden"
  else
    hidden=""
  fi
 sleep 2;
done
```

LISTING 4: Using a Second Disk

```
$ rsync -aAXv --exclude={"/dev/*","/proc/*","/sys/*","/tmp/*","/run/*",
  "/mnt/*","/media/*","/lost+found","/boot/"} / /mnt/veracrypt1
```

LISTING 5: /etc/fstab

```
/dev/sdb5             /boot  ext4  rw,relatime,data=ordered  0 1
/dev/mapper/veracrypt1  /      ext4  rw,relatime,data=ordered  0 1
```

Now the question arises as to how to install Arch Linux instead of Windows in a hidden VeraCrypt partition. The easiest way is by preparing an Arch Linux machine with the VeraCrypt GUI in a partition on a second hard drive and mounting the hidden container. It must be at least the same size as the current system. You need to consider whether you want to set up an alibi system in the outer container when defining the size of the container.

Installing on a second disk is safer because you do not need to touch the bootloader of the old system. You decide in the BIOS/EFI which system boots. An Rsync call [17] transfers the current system to the mounted VeraCrypt partition – in the example in Listing 4, to `/mnt/veracrypt1`.

In addition to the root partition, an encrypted system needs an unencrypted boot partition, for which you need to create another 100MB partition. Format it with:

```
mkfs.ext4 /dev/sd<xN>
```

You replace `<xN>` with the identifier for the matching device file and copy to it the contents of the `/boot` directory from the current system.

Startup Aid

What now follows are some steps already familiar to Arch Linux users from the initial installation of their system. On the mirrored system, first change the fifth line in `<Mountpoint>/etc/default/grub` to:

```
GRUB_CMDLINE_LINUX=⮐
  "cryptdevice=/dev/sd<xN>:veracrypt1 ⮐
   vera=1"
```

The `cryptdevice` keyword points to the partition encrypted with VeraCrypt. The colon is followed by the name of the mapper in `/dev/mapper/` (e.g., `veracrypt1` here), which VeraCrypt uses to access the currently mounted partition.

Copy the modified `vencrypt` script available online [16] to the `<Mountpoint>/usr/lib/initcpio/hooks/` directory. To be able to install in the initial ramdisk,

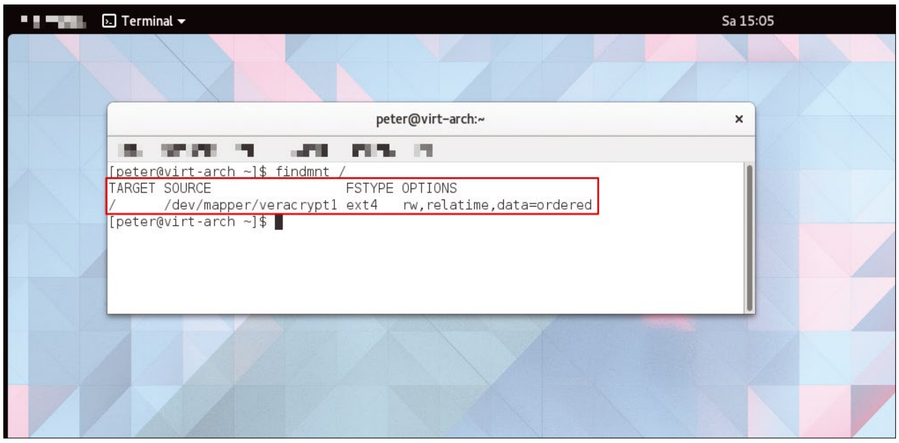

```
[peter@virt-arch ~]$ findmnt /
TARGET SOURCE                 FSTYPE OPTIONS
/      /dev/mapper/veracrypt1 ext4   rw,relatime,data=ordered
[peter@virt-arch ~]$
```

Figure 9: **You've done it! Arch Linux is running in a hidden VeraCrypt container mounted under** `/dev/mapper/veracrypt1`.

you need to duplicate `<Mountpoint>/usr/lib/initcpio/install/encrpyt` as `vencrypt`.

Now register the modified `vencrypt` script under `<Mountpoint>/etc/mkinitcpio.conf`; optionally, replace the existing original Arch Linux `encrpyt` version. The *HOOKS* line responsible for this [18] might then read as follows:

```
HOOKS="base udev autodetect modconf ⮐
       block filesystems keyboard ⮐
       keymap vencrypt fsck"
```

The order is important: `vencrypt` needs to follow `filesystems` and `keyboard` but must occur before `fsck`. The `keymap` lets you change the keyboard mapping if needed; otherwise, leave any existing hooks in your system unchanged.

Now it is time to use `chroot` (change root) to change to the new system. The arch-chroot script from the Arch Linux installation medium [19] handles this task; you can call it with `./arch-chroot <Mountpoint>`. The Ubuntu guide for fixing GRUB [20] also starts a chroot environment. Now you should change `/etc/fstab` as shown in Listing 5.

Now mount `/boot` and install GRUB on the new disk. On BIOS systems, use the `grub-install /dev/sd<X>` command; for EFI systems, check out the Arch wiki [21]. Next, generate a grub configuration with the parameters inserted into `/etc/default/grub` with:

```
$ grub-mkconfig -o /boot/grub/grub.cfg
```

The only thing missing is the initial ramdisk, which you can created with `mkinitcpio -p linux`. Now the encrypted system should boot from the second hard drive. Acknowledge the first password prompt after the GRUB menu by pressing Enter; the password is for the outer volume and will fail accordingly. After entering the password for the hidden volume at the second prompt, the system derived from the unencrypted original system boots (Figure 9).

For maximum deniability, use Rsync to mirror a Linux system in the outer container as well, and give it the same `/etc/fstab` file. The entry in the bootloader is fine for both systems and therefore does not require any further modification.

Conclusions

VeraCrypt impresses in three scenarios: (1) Access to VeraCrypt-encrypted objects is possible across platforms with Linux, Mac OS X, and Windows; (2) the GUI is ideal for volumes unlocked only when needed, whereas the Linux on-board tools play to their strengths with system-integrated, permanently mounted filesystems; (3) hidden VeraCrypt containers cannot be demonstrated to exist "by design," which adds security that you might need depending on the political situation in your country.

VeraCrypt comes with a bootloader that starts Windows systems in hidden containers. However, with an up-to-date Cryptsetup binary and some modifications to the initial ramdisk, this function can be emulated under Linux, too. Incidentally, VeraCrypt on Linux uses the kernel's dm-crypt mechanism for encryption on the fly, as do the Linux on-board methods, thus removing the need for a separate kernel module that could compromise system stability. ■■■

INFO

[1] TrueCrypt: *http://truecrypt.sourceforge.net*

[2] VeraCrypt: *https://veracrypt.codeplex.com*

[3] TrueCrypt audit: *http://blog.cryptographyengineering.com/2015/04/truecrypt-report.html*

[4] Rights escalation: *https://code.google.com/p/google-security-research/issues/detail?id=538*

[5] Rights escalation: *https://code.google.com/p/google-security-research/issues/detail?id=537*

[6] Key disclosure laws: *https://en.wikipedia.org/wiki/Key_disclosure_law*

[7] dm-crypt/LUKS: *https://wiki.archlinux.org/index.php/Dm-crypt*

[8] eCryptfs: *http://ecryptfs.org*

[9] E4M: *https://en.wikipedia.org/wiki/E4M*

[10] Security fixes: *https://veracrypt.codeplex.com/discussions/569777*

[11] Installation: *http://sourceforge.net/projects/veracrypt/files/*

[12] Documentation: *https://veracrypt.codeplex.com/documentation/*

[13] Windows bootloader: *http://sourceforge.net/p/veracrypt/discussion/technical/thread/a010f9bc/*

[14] Ubuntu initramfs: *https://wiki.ubuntu.com/Initramfs*

[15] openSUSE dracut: *https://www.kernel.org/pub/linux/utils/boot/dracut/dracut.html*

[16] Code for this article: *ftp://ftp.linux-magazine.com/pub/listings/magazine/188*

[17] Full-system backup with Rsync: *https://wiki.archlinux.org/index.php/Full_system_backup_with_rsync*

[18] Initcpio hooks: *https://wiki.archlinux.org/index.php/mkinitcpio#HOOKS*

[19] Chroot helper script: *https://projects.archlinux.org/arch-install-scripts.git/tree/arch-chroot.in*

[20] Fixing GRUB in Ubuntu via chroot: *https://help.ubuntu.com/community/Grub2/Installing#via_ChRoot*

[21] EFI in Arch wiki: *https://wiki.archlinux.org/index.php/Unified_Extensible_Firmware_Interface*

FREE DVD
clearOS 7.2 64-bit
openSUSE LEAP 42.1
Measuring Performance
ADMIN Network & Security
FREE DOUBLE-SIDED DVD!
ADMIN
Network & Security
JUN/JUL 2016
MEASURING PERFORMANCE
FIND BOTTLENECKS WITH I/O BENCHMARKING
DIGITAL BLACKMAIL
Surviving the scourge of ransomware
Tracking Memory Errors
Convenient Wedding
Can you really manage a Linux server with PowerShell?
Veil-Evasion
Slip your pen test past antivirus scanners
Windows Server 2016
Privileged Access Management and other new security features
AWS CloudFormation
Work smarter with DevOps in the cloud
Puppet
in the VMware universe
ADMIN Jun/Jul 2016
US$ 15.99
CAN$ 17.99
Windows • Linux • Unix • Solaris
WWW.ADMIN-MAGAZINE.COM
6 issues per year!

Automatically restore firewall filter rules

Reloaded

The Linux iptables packet filter lacks an easy way to load rules automatically after restarting a system, but you can automate this process several ways. *By Frank Hofmann*

Lead Image © stylephotographs, 123RF.com

Every administrator has to determine how to protect a network reliably against unauthorized access and ensure that the (sub) network fulfills its task as expected and is not misused as a starting point for malicious activities. At the network level, you have various ways and means at your disposal, including, for example, managing credentials for authenticating users (e.g., via PAM and LDAP), the appropriate selection of correctly configured services, and correct network device configurations.

Routers and firewalls are often used to isolate individual network segments, and the Linux kernel manages the firewall rules in the system's RAM. I demonstrate different methods for permanently storing your firewall settings by using iptables, thus removing the need to enter them again whenever you reboot.

Firewall Protection

From a network perspective, a firewall's primary objective is selecting and limiting the network packets that reach a network via a monitored interface on another network. In practice, firewalls are deployed not only as software directly

LISTING 1: Typical iptables

```
# iptables -F
# iptables -P INPUT  DROP
# iptables -P OUTPUT DROP
# iptables -P FORWARD DROP
# iptables -A INPUT -p tcp --dport 22 -s 192.168.45.0/24 -j ACCEPT
# iptables -A OUTPUT --sport 22 -d 192.168.45.0/24 --state ESTABLISHED -j ACCEPT
```

LISTING 2: iptables Manual Method

```
# iptables-save > /etc/rules
# cat /etc/rules
# Generated by iptables-save v1.4.21 on Tue Dec 8 23:03:26 2015
*filter
:INPUT DROP [3:604]
:FORWARD DROP [0:0]
:OUTPUT DROP [0:0]
-A INPUT -s 192.168.45.0/24 -p tcp -m tcp --dport 22 -j ACCEPT
-A OUTPUT -d 192.168.45.0/24 -p tcp -m tcp --sport 22 -j ACCEPT
COMMIT
# Completed on Tue Dec 8 23:03:26 2015
[...]
# iptables-restore < /etc/rules
```

LISTING 3: Customized

```
allow-hotplug eth0
iface eth0 inet dhcp
  pre-up /usr/local/sbin/firewall-on.sh
  post-down /usr/local/sbin/firewall-off.sh
```

on the system to be protected, but also in the form of separate hardware appliances that often combine various services under one roof.

In short, a firewall is a packet filter. Filtering is performed, for example, on the basis of (1) the IP address, (2) the connection state, (3) the MAC address, or (4) the payload of a packet.

The iptables package by the Netfilter project [1] includes the iptables command, which covers cases 1 and 2; for case 3, ebtables [2] steps in. In the future, nftables will combine the iptables, ip6tables, ebtables, and arptables projects under one roof. For case 4, evaluating and limiting user traffic, you need an application-level firewall, such as the Squid [3] proxy. In this article, I only look at iptables.

Operations

The iptables command works in chains with rules (filters) and actions (targets). On the basis of these criteria, the software fields packages (ACCEPT) and discards (DROP or REJECT) or forwards and rewrites them (MASQUERADE).

Listing 1 shows a typical call sequence. In this example, only connections via port 22 with a source IP address between 192.168.45.0 and 192.168.45.255 are allowed to pass. The entry in line 1 clears all the existing rules; the commands in lines 2 to 4 set the default behavior.

Line 5 adds a filtering rule to the chain for received packets. This limits the accepted packages (-j ACCEPT) to TCP packets (-p tcp) for the target port 22 (--dport 22) from the subnet 192.168.45.x (-s 192.168.45.0/24). Line 6 allows also the corresponding response packets (-j ACCEPT) – original from port 22 here (--sport 22) to the target network 192.168.45.x (-d 192.168.45.0/24) and only for existing connections (--state ESTABLISHED).

To apply these commands to IPv6, you need to use the new ip6tables, which uses the same syntax.

Validity

As mentioned previously, the iptables rules apply only as long as the system is running. The software does not store the rules persistently; it only keeps them in main memory. If you turn off the computer or reboot it, the current firewall settings will be lost. To prevent having to enter them every time you reboot, you need to store the rules permanently and enable them automatically when the computer starts up.

The variants discussed here are based on my own experience and the firewall entries from the Debian wiki [4]. The intent was to use only built-in tools. The variations also show a selection of solutions – you can decide which best fit your purposes.

Manual

With the manual version, you use the iptables commands and tools. The ipta-

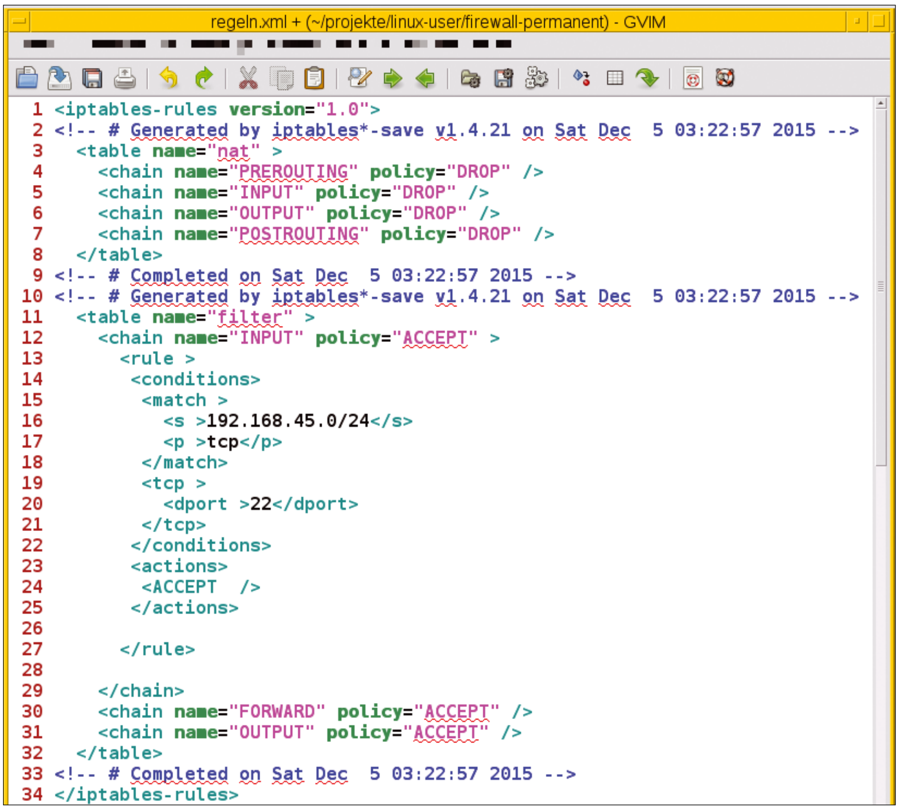
regeln.xml + (~/projekte/linux-user/firewall-permanent) - GVIM

```
<iptables-rules version="1.0">
<!-- # Generated by iptables*-save v1.4.21 on Sat Dec  5 03:22:57 2015 -->
  <table name="nat" >
    <chain name="PREROUTING" policy="DROP" />
    <chain name="INPUT" policy="DROP" />
    <chain name="OUTPUT" policy="DROP" />
    <chain name="POSTROUTING" policy="DROP" />
  </table>
<!-- # Completed on Sat Dec  5 03:22:57 2015 -->
<!-- # Generated by iptables*-save v1.4.21 on Sat Dec  5 03:22:57 2015 -->
  <table name="filter" >
    <chain name="INPUT" policy="ACCEPT" >
      <rule >
       <conditions>
        <match >
          <s >192.168.45.0/24</s>
          <p >tcp</p>
        </match>
        <tcp >
          <dport >22</dport>
        </tcp>
       </conditions>
       <actions>
        <ACCEPT  />
       </actions>

      </rule>

    </chain>
    <chain name="FORWARD" policy="ACCEPT" />
    <chain name="OUTPUT" policy="ACCEPT" />
  </table>
<!-- # Completed on Sat Dec  5 03:22:57 2015 -->
</iptables-rules>
```

Figure 1: The iptables-xml tool converts iptables rules to an XML structure that you can reuse in other applications.

`bles-save` command reads the current firewall rules. Because the program sends them to standard output, you need to redirect the output to a file by using the (>) redirection operator (Listing 2). This file has a specific, compact format.

You can restore these rules later with the command `iptables-restore` (Listing 2, last line). This approach is simple and clear-cut, but still not automated. However, you can manage that task with the variations that follow.

Automated

Runlevels contain shell scripts that run with a system runlevel change. To load the previously stored firewall rules, you need to create an appropriate shell script. In it, you can either call `iptables-restore` or create a separate call to iptables in the desired order for each individual rule.

The variant presented here is far more customized. It involves expanding the entries for the network interfaces in the `/etc/network/interfaces` file. The entries allow you to specify a script or command that runs when the system activates or disables a network interface (Listing 3).

The first command in line 3 is specified after the keyword `pre-up` and executes the command before activating the interface. The `post-down` statement in line 4 refers to the time after disabling the interface. Therefore, the firewall only processes certain rules when a particular interface is switched on.

Resourceful Debian developers also identified the problem and devised a solution, which they dubbed `iptables-persistent` [5]. It has been included since Debian 5 "squeeze" and combines the variants featured in Listing 3.

For this to happen, `iptables-persistent` creates two files – `/etc/iptables/rules.v4` for IPv4 and `/etc/iptables/rules.v6` for IPv6 – and stores the current firewall rules during package configuration. Additionally, it creates a suitable init script named `/etc/init.d/netfilter-persistent`, which you can call via `service netfilter` with the usual switches: `start`, `stop`, `reload`, and `restart`.

Future Useful Information

In the context of iptables, several small programs facilitate everyday admin life. They include `iptables-apply` and `iptables-xml`, both of which are part of iptables, as well as `iptables-converter`.

The `iptables-apply` shell script helps you test firewall rules remotely. It allows you to roll back rule changes if they would take down the current connection.

To do this, the script asks whether the changes to the rules are okay. If you do not respond within a certain period of time, it rolls back the changes. This reduces the risk that you lock yourself out when changing the firewall rules, although you still have the option of communication via the serial interface, but only if the remote computer is set up for it, because the firewall rules do not take it into consideration under normal circumstances.

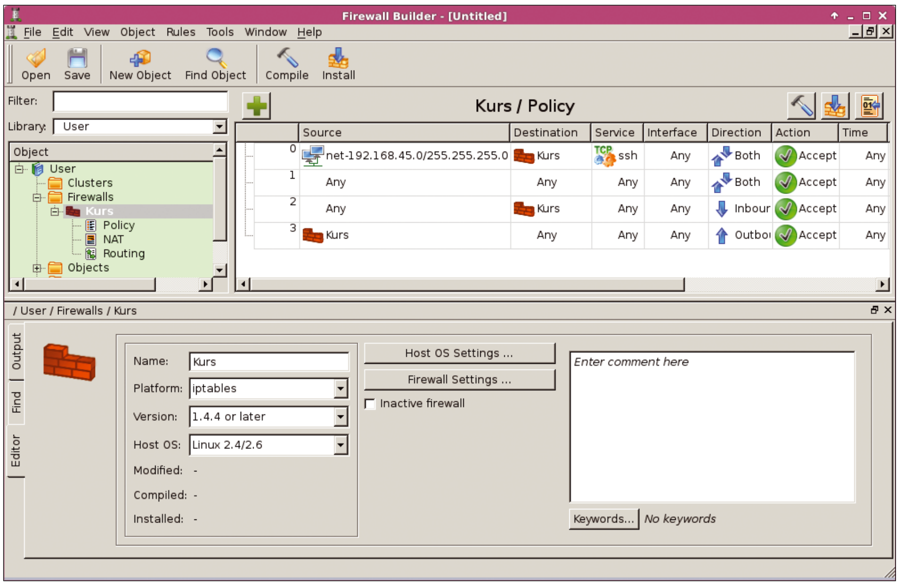

Kurs / Policy

	Source	Destination	Service	Interface	Direction	Action	Time
0	net-192.168.45.0/255.255.255.0	Kurs	ssh	Any	Both	Accept	Any
1	Any	Any	Any	Any	Both	Accept	Any
2	Any	Kurs	Any	Any	Inbour	Accept	Any
3	Kurs	Any	Any	Any	Outboı	Accept	Any

/ User / Firewalls / Kurs

Name: Kurs
Platform: iptables
Version: 1.4.4 or later
Host OS: Linux 2.4/2.6
Modified: -
Compiled: -
Installed: -
□ Inactive firewall

Figure 2: Graphically editing rules in Firewall Builder. You can see the imported rules created by `iptables-converter`.

Access would also be possible via a KVM switch, which acts as a toggle switch that connects a set of devices (keyboard, video, mouse) to more than one computer. In this way, you can control multiple computers with one keyboard, mouse, and monitor.

The `iptables-xml` script contains firewall rules in XML format. Several graphical tools for firewalls understand and process this format. Figure 1 shows an example of an XML file generated by this tool.

The `iptables-converter` program is useful, as well. It converts `iptables` statements to the format that `iptables-save` and `iptables-restore` use. Thus, you can convert existing shell scripts with sequences of `iptables` statements directly to the required format. The graphical `fwbuilder` (Firewall Builder) tool [6] also reads the files and generates customized package checks (Figure 2). ■■■

INFO

[1] Netfilter: *http://www.netfilter.org*

[2] ebtables: *http://ebtables.netfilter.org*

[3] Squid: *http://www.squid-cache.org*

[4] Debian firewall wiki: *https://wiki.debian.org/DebianFirewall*

[5] Saving iptables firewall rules permanently: *https://www.thomas-krenn.com/en/wiki/Saving_Iptables_Firewall_Rules_Permanently*

[6] Firewall Builder: *http://www.fwbuilder.org*

AUTHOR

Frank Hofmann (*http://www.efho.de*) works in Berlin, Germany for Büro 2.0, an open source expert network, as a service provider specializing in printing and layout. Since 2008 he has also coordinated the regional meeting of the Berlin-Brandenburg region Linux usergroup. He is the co-author of the Debian package management book (*http://www.dpmb.org*).

THANKS

The author thanks Werner Heuser and Axel Beckert for their criticism and suggestions in the run-up to this article.

Klaus Knopper answers your Linux questions

Ask Klaus!

By Klaus Knopper

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*

UEFI Boot Problems

I just bought a new Lenovo Notebook, and I'm trying to start Knoppix from USB flash disk. Knoppix seems to start to a certain point, but then I get no graphics, just a black screen. None of the cheatcodes seem to help. What happens, how to fix?

With the recent versions, Knoppix supports booting on UEFI-enabled computers. However, the graphics card in UEFI boot mode is set to a VESA graphics mode by the firmware, not to "text mode" as known for non-UEFI systems. A visible indication of this is the missing boot logo and the fact that after the kernel has loaded, the usual text messages that indicate progress during the Knoppix boot are missing until the mode setting kicks in and resets the screen resolution. Because the graphics chipset had been in a VESA mode previously, the X server can fail to set the native resolution; hence, the "black screen" or system freeze effect that you observed.

The only workaround I found so far is disabling UEFI and reverting to *Legacy* or CSM (compatibility support module) boot mode, which re-enables text mode for the bootloader and puts the graphics card in a default mode that lets Xorg start normally. Some notebooks, especially MacBooks with DVD drives, do not boot in UEFI mode automatically when selecting to boot from CD or DVD. For this purpose, the Knoppix "boot-only" CD is useful. An ISO for the CD resides in the `KNOPPIX` folder. The boot-only CD starts the bootloader and kernel from the CD, then searches for other drives that contain the `KNOPPIX` folder (e.g., a USB flash disk), and continues booting from there.

Joining PDFs

Is there a simple way to concatenate several hundred PDF files into a single, continuous PDF under Linux?

The easiest way that requires no special program is using the PostScript/PDF interpreter Ghostscript directly. The original command line is somewhat lengthy, so I usually put this in a script (e.g., `pdfcat.sh`):

```
#!/bin/bash
gs -q -sPAPERSIZE=a4 ↩
      -dNOPAUSE ↩
      -dBATCH ↩
      -sDEVICE=pdfwrite ↩
      -sOutputFile=- "$@"
```

If you need a different paper size than A4 (e.g., letter), change the `PAPERSIZE` option accordingly.

The syntax for concatenating all files in a directory is:

```
pdfcat.sh directory/*.pdf > all.pdf
```

ADMIN Network & Security
MEASURING PERFORMANCE
LINUX PRO MAGAZINE
TUNING SSDs
raspberry pi GEEK
UBUNTU user
ALTERNATIVES
Drupal Watchdog
COMMUNITY

GOOGLE SPECIAL
Google
PERL
POWERHOUSE PERL
Raspberry Pi HANDBOOK
FREE FROM XP
GIMP HANDBOOK
LINUX SHELL HANDBOOK
RASPBERRY PI ADVENTURES

Local data encryption for cloud storage

Cloud Master

Synchronizing your data in the cloud is practical, but it's risky if you don't encrypt your data. Desktop encryption utilities offer various levels of security and ease. *By Jens Kubieziel*

Backing up data in the cloud sounds easy and useful; little wonder then that many individuals and companies take advantage of this opportunity. Companies that offer these services take your locally stored data and sync it with a storage service on the Internet – often automatically. In this way, you keep your directories synchronized without having to worry about backups. Of course, such services offer not only benefits but also risks. The burning issue to consider is: Who has access to you data?

Sometimes cloud providers automatically scan the uploaded files to check them for unacceptable content (e.g., child pornography or copyright infringement) [1]. In the first case, suspicious data is sent to investigating authorities, and in the second, the algorithm locks the sharing feature. If you opt for a free version of a service, you might start receiving advertisements based on the content of your evaluated files.

BOXCRYPTOR

One of the more popular encryption programs is Boxcryptor [2]. Although the classic version [3] still supports Linux, the current versions do not. As the manufacturer states on its website: "It will not be supported on upcoming versions of these operating systems. Therefore, we can't guarantee that the Classic version will work on them." In this article, I only mention the software for the sake of completeness, because you have many other alternatives.

As the Snowden documents reveal, the NSA is also interested in data of any kind. From the beginning, administrators of cloud services also have had access to user data. Additionally, many companies in Europe are increasingly unsure about what happens when they offload their personal data onto servers elsewhere in the world.

Key Services

One solution is end-to-end encryption, wherein the user encrypts all the data on the local machine (with exclusive possession of the key) and then uploads to the cloud server. Some applications promise to handle the encrypted data in an easy and user-friendly way, but Linux also has on-board resources to help you achieve your goal. Here, I present a range of programs and discuss their advantages and disadvantages (see also the "Boxcryptor" box).

PanBox

One candidate goes by the name of PanBox [4]. The software was released in early 2015 and was funded by Germany's Federal Department of Justice and Consumer Protection. The developers of the program, the Fraunhofer Institute for Secure Information Technology and Sirrix AG, emphasize that privacy by design was an important precondition. In other words, the software model is not privacy by policy, which would have to rely on the existence of benevolent operators or laws. Instead, its strength is its verifiability and design. The sources are available from a public archive on GitHub [5], so the code is open to examination for bugs and verification of safeguards with regard to the security of the software.

PanBox is available under GPLv3. In addition to the open source variant, an Enterprise version of PanBox specially targets public authorities and companies [6], offering access to a directory service (LDAP) and a public key infrastructure (PKI).

A ZIP file for Linux and other operating systems, which you just download

Lead Image © Luciano De Polo, 123RF.com

INSTALLING PANBOX

Because PanBox 1.1.0 is a Java application, you need a current version of Java. The developers recommend the Java version by Oracle. Because the software uses strong encryption, you also need the Java Cryptography Extension (JCE), which is available in the download section of Java SE [8].

You still have to satisfy other dependencies. In the ZIP file, a README lists the libraries required in Ubuntu, Arch Linux, Fedora, and Gentoo. I used Ubuntu 15.10 for the article and thus installed the following additional software:

```
sudo apt-get install dbus-java-bin python-appindicator python-nautilus
   libbluetooth-dev python-notify python-gtk2 python-dbus
```

To launch the software, you run the `start.sh` script.

and unpack into a subdirectory, is available for download from the Sirrix AG site [7]. After installing and launching for the first time (see the box "Installing PanBox"), the software checks to see whether an identity already exists on the computer. If not, it wants to know the first name, last name, and email address of the user. In the next dialog, PanBox asks you to enter a device name. This name lets you map to devices easily. The hostname is preset.

Finally, you need to enter a password. According to the manual [9], the password must be at least eight characters in length and "include a random combination of upper/lower case letters, digits, and special characters." In the test, eight characters was the only discernible limit; passwords comprising only numbers or letters were accepted by PanBox without comment. After entering the password, the software displays the previously entered data and generates the identity in the next step. Finally, it opens the application window (Figure 1).

At first launch, PanBox will detect a Dropbox installation, at which time it looks for your existing Dropbox folder and launches a second setup wizard that proceeds to integrate your Dropbox share. The wizard requires an access token, which is a kind of password that tells the Dropbox service that PanBox is allowed to access your Dropbox account. Clicking on *Refresh* launches a browser, and a Dropbox website asks whether PanBox is allowed to access the service. If you allow this, the page displays the access token which you need to copy and enter in PanBox at the end of the configuration.

Although the press releases for PanBox promise a simple and user-friendly approach, I had some irreconcilable differences with the application. I began with a fresh installation of Ubuntu 15.10 as the basis. After starting the application, a white window without any menus and the tray icon appeared. The application window did not respond to my attempts to close it; even clicking on *Exit* in the tray icon resulted in no discernible response. Working on Ubuntu 15.10 was therefore impossible.

Under Debian the software launched and opened a functional application window. During use, it crashed repeatedly when I tried to change the language (Figure 2). Apparently, the interface still needs some work. Under the vastly simplified Directory share, the application shares a folder of your choice set up in

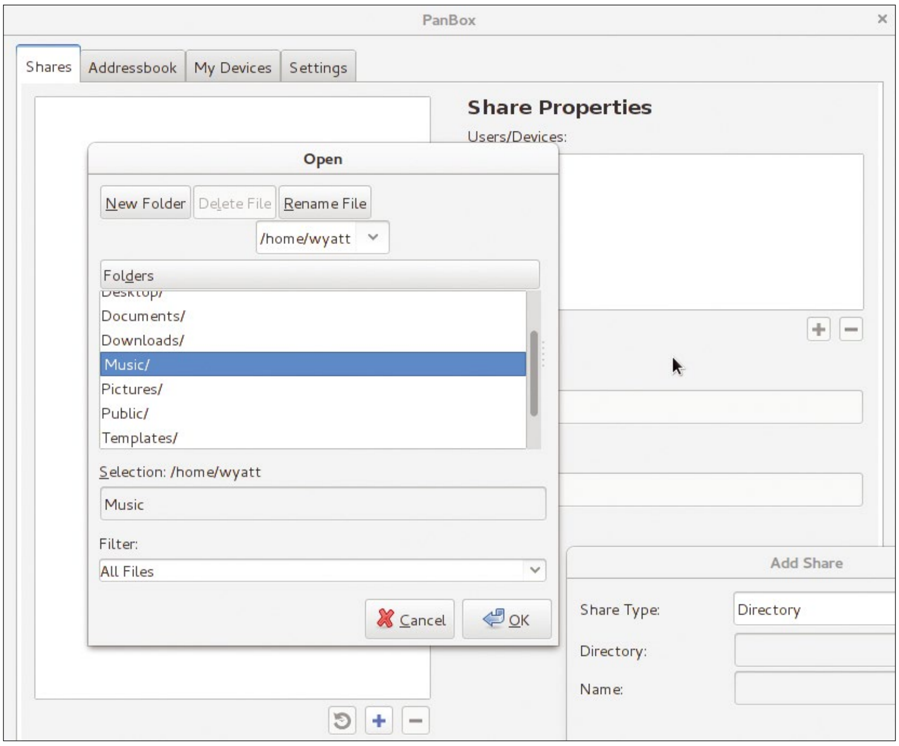

/home/wyatt

Selection: /home/wyatt

Music

Filter: All Files

Share Type: Directory

Directory:

Name:

Figure 1: **PanBox encrypts files on the desktop and then passes them on to Dropbox. The software is available in open source and Enterprise versions.**

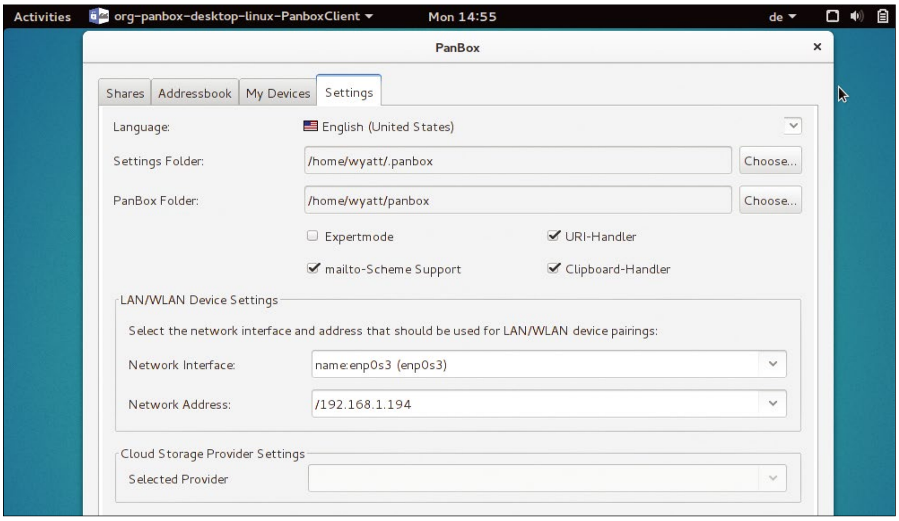

Language: English (United States)

Settings Folder: /home/wyatt/.panbox

PanBox Folder: /home/wyatt/panbox

☐ Expertmode ☑ URI-Handler

☑ mailto-Scheme Support ☑ Clipboard-Handler

LAN/WLAN Device Settings

Select the network interface and address that should be used for LAN/WLAN device pairings:

Network Interface: name:enp0s3 (enp0s3)

Network Address: /192.168.1.194

Cloud Storage Provider Settings

Selected Provider

Figure 2: **More than once PanBox froze in the middle of its activity for no particular reason.**

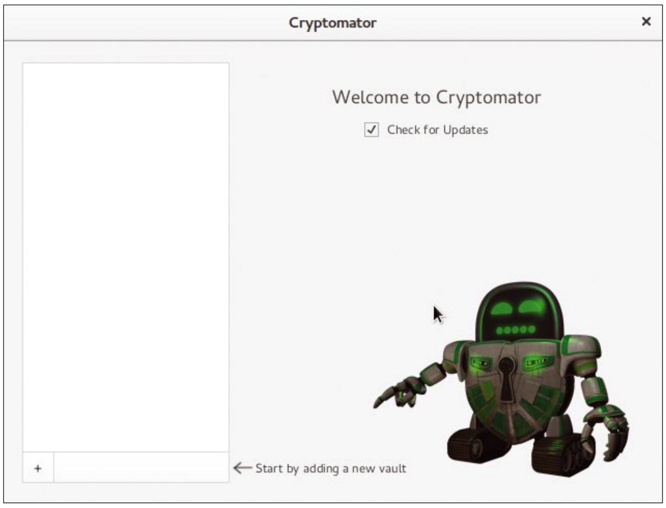

Figure 3: The Cryptomator mascot appears on launch; an update function is also offered.

the *Shares*) tab. The program distinguishes between a Dropbox share, which works only with Dropbox integration, and a Directory share, which is for users with any other cloud provider. Unfortunately, I was unable to test the implementation because of the problems already described.

Cryptomator

Cryptomator [10] offers another solution to encrypting data for the cloud. Developer Sebastian Stenzel mainly works on the software, which is available under the MIT license. The source code is on GitHub [11]. If you look at the project homepage, you will note that the authors go to great lengths to explain the cryptographic contexts. That and the fact that it is free software ensure confidence and allow third parties to develop programs that interact with Cryptomator.

Apparently the developers have put much thought into how they want to protect the software. To begin, Cryptomator generates a key from your password and a random value (salt). For this, it uses the `scrypt` key derivation algorithm [12], making it difficult to brute-force the key. Usually, a hash function (MD5, SHA-1, SHA-256) picks up the password together with the salt. The output from the hash function is the key.

With the use of special hardware and algorithms, attackers can try several million values per second and very quickly guess a password under certain circumstances. The Scrypt algorithm therefore performs the hash function multiple times and also uses a large amount of RAM. As a result, its slows down attempts to guess the password.

If the operator selects the correct parameters for Scrypt, even special hardware will not achieve more than a few hundred attempted guesses per second. The system later uses the key encryption key (KEK) calculated in this way to decrypt the master key.

For other cryptographic actions, Cryptomator uses the AES algorithm relying on cypher block chaining (CBC) or counter (CTR) mode. The SHA-256 hash algorithm is used as the basis for further operations and is a good choice, because it rules out many attack vectors.

A window with the mascot of the software (Figure 3) appears after installation (see the box "On Disk"). Adding a vault begins in the lower left corner. A vault is a folder in which the encrypted content is stored as well as the key. Although the key cannot be used easily for decryption because of the key generation process already mentioned, it would appear to be safer to store it outside of the directory. The vault then lies within a folder hosted by the cloud provider.

To create the vault, you need to specify a directory. In the second step, you create a password. The software has no restrictions in terms of the length or the characters used. Minimum requirements would be desirable: After all, despite a good derivation function such as Scrypt, simple passwords can be guessed pretty quickly.

Before Cryptomator opens the vault, it prompts you again for the password you just composed. At the same time, it acti-

ON DISK

Two possibilities exist for installing of Cryptomator on Linux: The developers offer a ready-made package for Debian-based systems that can be installed with:

```
dpkg -i <filename>
```

The package stores the software in the `/opt/Cryptomator` directory. For other systems, you can launch the software's `jar` file by calling:

```
java -jar <filename>
```

Alternatively, you can run the file `/opt/Cryptomator/Cryptomator`.

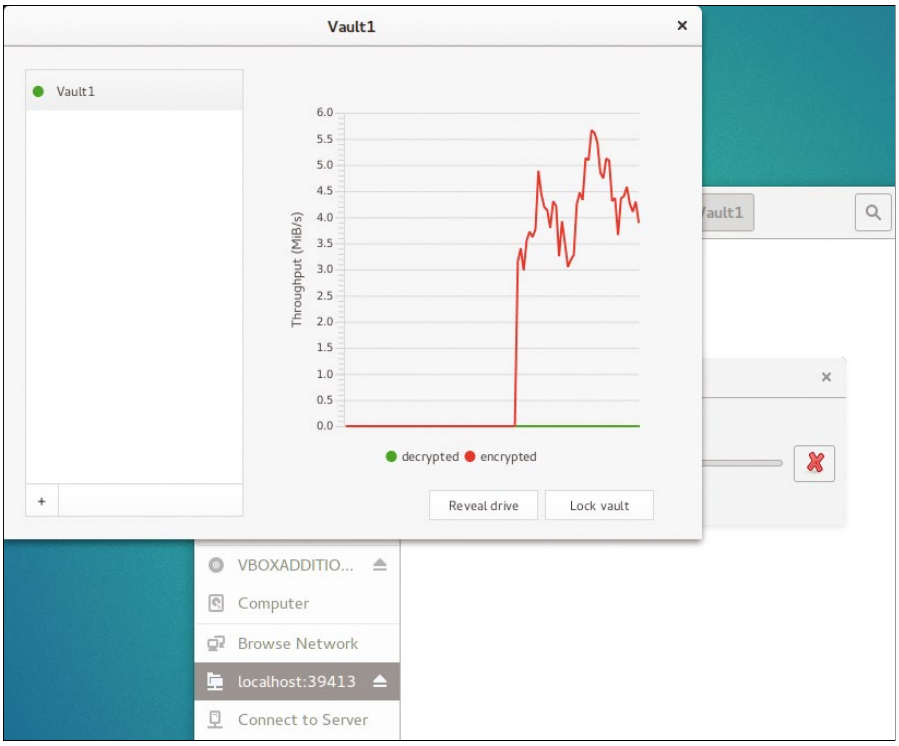

Figure 4: Cryptomator graphically displays the encryption activity when receiving a file. Below you can see the WebDAV server integrated in Gnome Files.

vates a local WebDAV server on a random port in the background. Nautilus and Gnome Files (file manager) support WebDAV and automatically display the new directory.

For the command line, you need a program like Cadaver [13] to your manage files. When you move files from your hard disk to the WebDAV directory, Cryptomator encrypts them automatically and places them in the vault (Figure 4). After a short learning curve, Cryptomator was easy to use, with no appreciable defects, such as crashes, in the lab test. Another positive aspect is that the developers are still actively working on the software.

Old Companions

In addition to the programs designed for encrypting data in the cloud, Linux offers other cryptotools, including GnuPG [14] and the obsolete TrueCrypt [15], although it is now available as a fork named VeraCrypt [16].

GnuPG is typically used to encrypt email (e.g., for use with Enigmail [17]). Thanks to the tool, any Linux user can encrypt files locally and drop them into a cloud folder to synchronize the data. However, this solution requires some manual work. For example, you need to make sure not to send unencrypted content to the cloud folder or decrypt files in the wrong folder and then accidentally sync them to the cloud. If you still want to use GnuPG, you should think about using scripts that handle the important work steps, thus preventing errors.

Using TrueCrypt or its successor VeraCrypt (Figure 5), you can create encrypted containers that a folder intended for the cloud can then synchronize. Only users who know the password and have the container mounted will see the content. The disadvantages are that small changes to the files in the container affect the entire container, which means that unless the provider supports block-level synchronization, the service syncs the complete container, even if just one character changes in a file.

eCryptfs [18], which Ubuntu uses [19], is also worthy of mention. The kernel-based encrypted filesystem is mounted on an existing filesystem (e.g., ext4). eCryptfs then creates two directories (`~/Private` and `~/.Private`) by default and writes the encryption information to the headers of the files to be encrypted. eCryptfs automatically stores files that the user saves to `~/Private` as encrypted files in `~/.Private`. From there, you can upload to the cloud. However, you can customize the software to suit your own needs and locate the encrypted folder in Dropbox. eCryptfs is not foolproof, however, because an attacker that gets hold of the password can read the encrypted user directories.

Conclusions

Programs that have been designed especially for encrypting cloud data seem to offer the better approach all told. Cryptomator is currently only available as a beta version; nonetheless, it still demonstrated the strongest performance in the lab. It also appears to be the easiest to use of the tested programs.

PanBox offers more features and therefore also requires more time for training. Assuming that the bugs encountered during testing are resolved in the future, PanBox also seems to be a strong candidate for handling data for the cloud. ∎∎∎

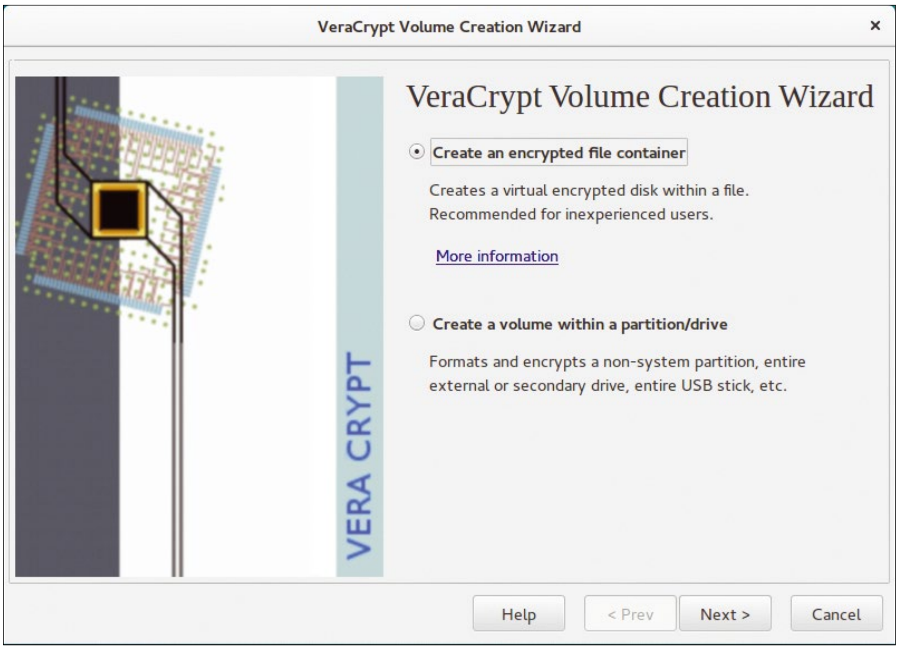
VeraCrypt Volume Creation Wizard

VeraCrypt Volume Creation Wizard

● Create an encrypted file container

Creates a virtual encrypted disk within a file. Recommended for inexperienced users.

More information

○ Create a volume within a partition/drive

Formats and encrypts a non-system partition, entire external or secondary drive, entire USB stick, etc.

Figure 5: After the developers gave up on TrueCrypt, VeraCrypt took over.

INFO

[1] "Is your cloud drive really private? Not according to fine print" by Rosa Golijan, NBC News, March 15, 2013: *http://www.nbcnews.com/technology/your-cloud-drive-really-private-not-according-fine-print-1C8881731*

[2] Boxcryptor: *https://www.boxcryptor.com*

[3] Boxcryptor with Linux support: *https://www.boxcryptor.com/en/classic*

[4] PanBox: *https://www.sirrix.com/content/pages/66722.htm*

[5] PanBox code: *https://github.com/Sirrix-AG/PanBox*

[6] PanBox Enterprise: *https://www.sirrix.com/content/pages/67191.htm*

[7] PanBox download: *https://www.sirrix.de/content/pages/66538.htm*

[8] Oracle Java: *https://www.oracle.com/java/index.html*

[9] PanBox manual: *https://cybersecurity.rohde-schwarz.com/sites/default/files/download/panbox_benutzerhandbuch_-_en.pdf*

[10] Cryptomator: *https://cryptomator.org*

[11] Cryptomator source code: *https://github.com/cryptomator/cryptomator*

[12] Scrypt: *https://www.tarsnap.com/scrypt.html*

[13] Cadaver: *http://www.webdav.org/cadaver/*

[14] GnuPG: *https://www.gnupg.org*

[15] TrueCrypt: *http://truecrypt.sf.net*

[16] VeraCrypt: *https://veracrypt.codeplex.com*

[17] Enigmail: *https://www.enigmail.net*

[18] eCryptfs: *http://ecryptfs.org*

[19] eCryptfs in Ubuntu: *https://help.ubuntu.com/community/EncryptedPrivateDirectory*

Spam Filter Mechanics

Wise Choice

Spam filters have different modes of operation. Understanding how they work can help you choose which to use. *By Bruce Byfield*

These days, the choice of spam filters comes down to Bogofilter [1] and SpamAssassin [2]. Other choices, like DSPAM [3], are no longer in development. A few other choices (e.g., SpamBayes [4]) are available, but when an email reader offers a plugin, it is almost always for either Bogofilter or SpamAssassin.

What is less often discussed is which filter is the best to use in which circumstances. Instead, most users simply nod solemnly when they read that both involve "Bayesian filtering." Most of us – including many who use the phrase – have no idea what Bayesian filtering is, but it sounds scientific and reassures us that either choice is acceptable.

In fact, learning that Bogofilter and SpamAssassin are "Bayesian" is useless for choosing between them. To call them Bayesian means nothing more than their structure is based on the 18th-century work of Thomas Bayes [5] in statistics and probability. More specifically, both apply Bayes' work by collecting words and assigning a probability that each word indicates spam. The more suspect words contained in an email, the greater the chance it is spam. However, to make an informed choice between spam filters requires considerably more detail.

Bogofilter

Bogofilter has its roots in "A Plan for Spam" [6], a 2002 essay by English developer Paul Graham. After trying to develop filters based on the identifying characteristics of spam, Graham concluded that beyond a certain point, the more rules he added, the more false positives he obtained – that is, the more email messages that were incorrectly identified as spam.

Graham's solution was to parse his samples of spam and non-spam into tokens, or individual words, and use Bayesian tools to assign each token the possibility that it indicates spam, biasing them slightly in favor of not being spam to minimize false positives. By examining the top 15 tokens in the header and body of each new email message, he calculated the possibility that it was spam. If the probability was greater than 0.9, the message was considered spam.

According to Graham, the advantage of this statistical approach is that it refers to something real – the probability of being spam – and worked with both neutral and spam-indicating words.

However, he also recognized that the more personalized the filter was, the more accurate it would be. For this reason, he also included the possibility of using white lists to indicate non-spam, or "ham," and black lists to indicate spam.

After reading Graham's essay, Eric S. Raymond founded the Bogofilter project. Today, Bogofilter is maintained by other

Lead Image © Jakub Jirsak, 123RF.com

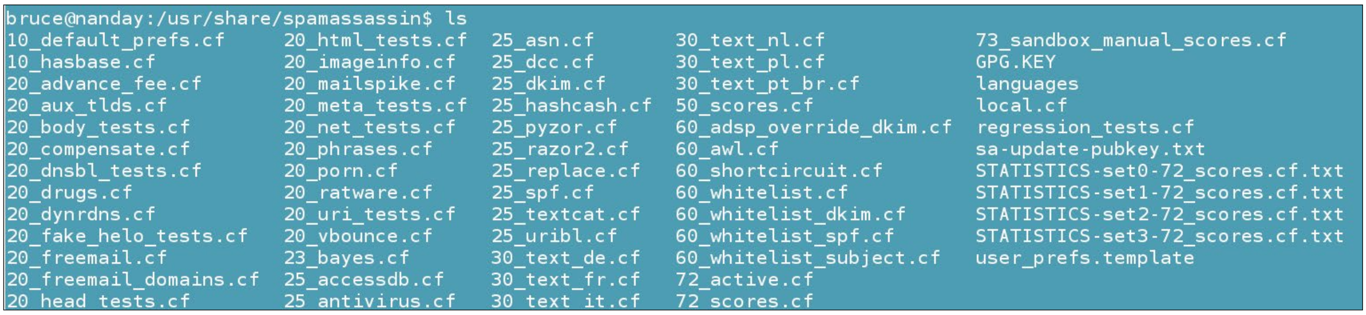

```
bruce@nanday:/usr/share/spamassassin$ ls
10_default_prefs.cf     20_html_tests.cf  25_asn.cf       30_text_nl.cf            73_sandbox_manual_scores.cf
10_hasbase.cf           20_imageinfo.cf   25_dcc.cf       30_text_pl.cf            GPG.KEY
20_advance_fee.cf       20_mailspike.cf   25_dkim.cf      30_text_pt_br.cf         languages
20_aux_tlds.cf          20_meta_tests.cf  25_hashcash.cf  50_scores.cf             local.cf
20_body_tests.cf        20_net_tests.cf   25_pyzor.cf     60_adsp_override_dkim.cf regression_tests.cf
20_compensate.cf        20_phrases.cf     25_razor2.cf    60_awl.cf                sa-update-pubkey.txt
20_dnsbl_tests.cf       20_porn.cf        25_replace.cf   60_shortcircuit.cf       STATISTICS-set0-72_scores.cf.txt
20_drugs.cf             20_ratware.cf     25_spf.cf       60_whitelist.cf          STATISTICS-set1-72_scores.cf.txt
20_dynrdns.cf           20_uri_tests.cf   25_textcat.cf   60_whitelist_dkim.cf     STATISTICS-set2-72_scores.cf.txt
20_fake_helo_tests.cf   20_vbounce.cf     25_uribl.cf     60_whitelist_spf.cf      STATISTICS-set3-72_scores.cf.txt
20_freemail.cf          23_bayes.cf       30_text_de.cf   60_whitelist_subject.cf  user_prefs.template
20_freemail_domains.cf  25_accessdb.cf    30_text_fr.cf   72_active.cf
20_head_tests.cf        25_antivirus.cf   30_text_it.cf   72_scores.cf
```

Figure 1: SpamAssassin includes more than 50 tests to detect spam.

developers,and has refined Graham's calculations based on Gary Robinson's suggestions [7]. The modern refinements include recognizing MIME types, treating each hostname and IP address as a separate token (rather than dividing them up into separate words), and ignoring dates and Message-IDs as irrelevant. However, the basic approach remains that advocated by Graham.

The mathematically inclined can learn more about how Bogofilter assigns the probability of an email being spam by following the links and reading the man page for the filter. However, the most important point for the average user is that Bogofilter relies on statistical probability, supplemented by each user's list of spam and ham. Advocates of this approach emphasize its simplicity, as well as its lower number of false positives once it is trained – that is, once the white and black lists are produced. These lists are contained in the `.bogofilter` folder in your home directory.

SpamAssassin

SpamAssassin takes a different approach from Bogofilter. SpamAssassin's main approach is to identify the characteristics of spam and then run tests to locate them. Many tests, although not all, rely heavily on regular expressions to catch variations of words and phrases.

You can view the Perl scripts used by SpamAssassin in `/usr/share/spamassassin` (Figure 1). More than 50 are listed in my current installation of Debian Stable. From their number alone, you can tell they are a varied lot, but they include tests for the common indicators of spam in headings, in the bodies of email, and in HTML code, as well as tests for recognizing offers for anti-viruses, drugs, and pornography. In the English version, some basic tests for French, German, and Italian are also included. The scripts also include a Bayesian probability test similar

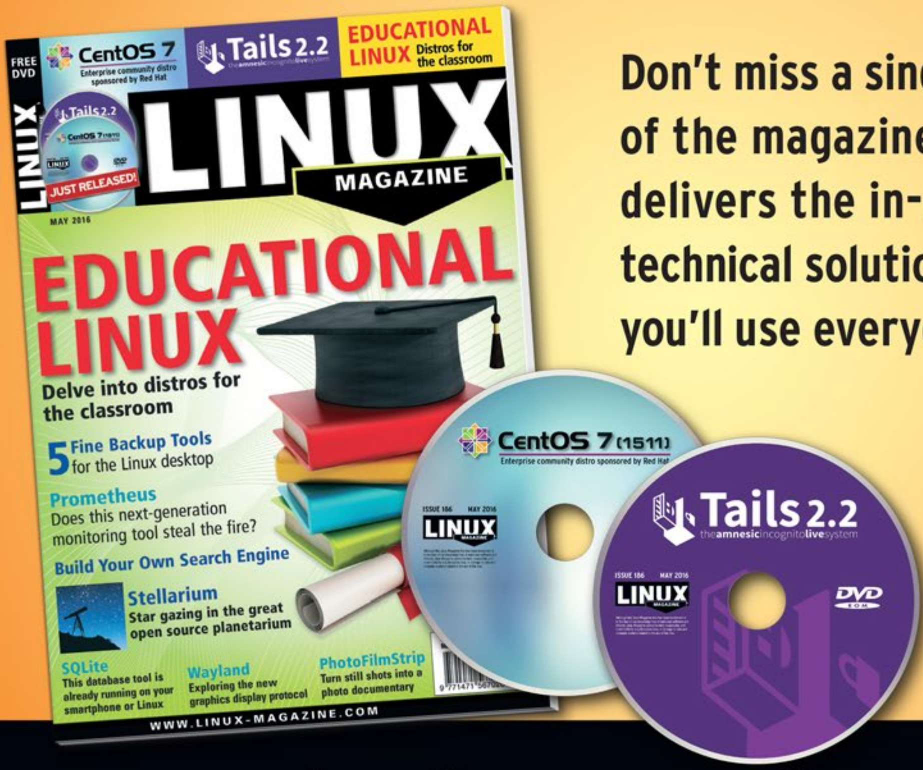

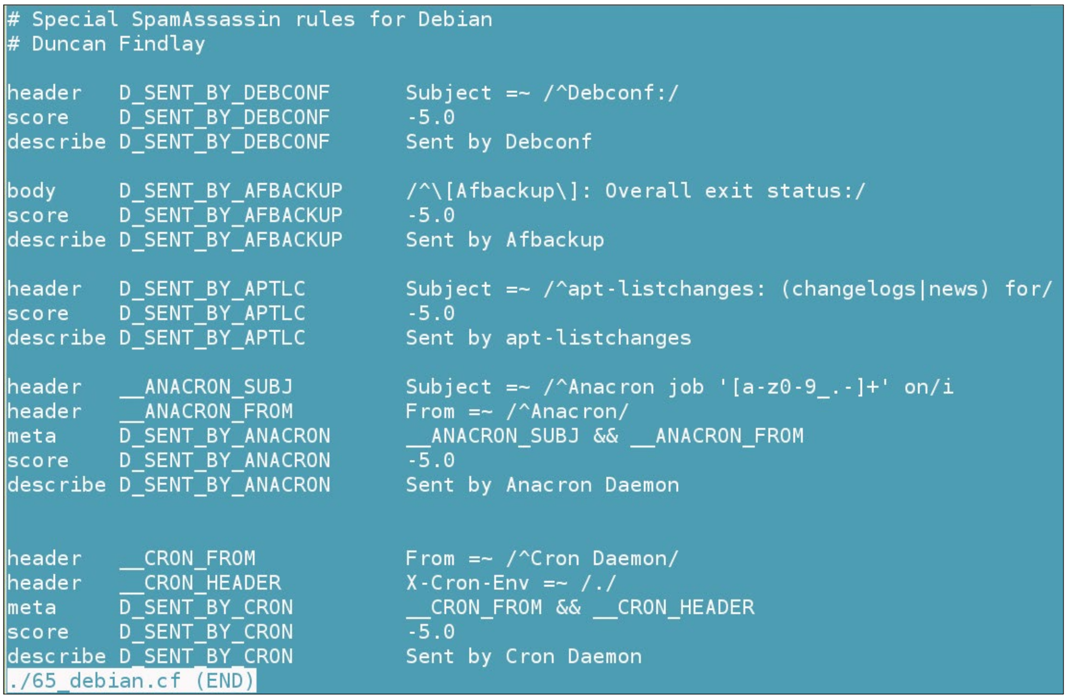

```
# Special SpamAssassin rules for Debian
# Duncan Findlay

header   D_SENT_BY_DEBCONF      Subject =~ /^Debconf:/
score    D_SENT_BY_DEBCONF      -5.0
describe D_SENT_BY_DEBCONF      Sent by Debconf

body     D_SENT_BY_AFBACKUP     /^\[Afbackup\]: Overall exit status:/
score    D_SENT_BY_AFBACKUP     -5.0
describe D_SENT_BY_AFBACKUP     Sent by Afbackup

header   D_SENT_BY_APTLC        Subject =~ /^apt-listchanges: (changelogs|news) for/
score    D_SENT_BY_APTLC        -5.0
describe D_SENT_BY_APTLC        Sent by apt-listchanges

header   __ANACRON_SUBJ         Subject =~ /^Anacron job '[a-z0-9_.-]+' on/i
header   __ANACRON_FROM         From =~ /^Anacron/
meta     D_SENT_BY_ANACRON      __ANACRON_SUBJ && __ANACRON_FROM
score    D_SENT_BY_ANACRON      -5.0
describe D_SENT_BY_ANACRON      Sent by Anacron Daemon


header   __CRON_FROM            From =~ /^Cron Daemon/
header   __CRON_HEADER          X-Cron-Env =~ /./
meta     D_SENT_BY_CRON         __CRON_FROM && __CRON_HEADER
score    D_SENT_BY_CRON         -5.0
describe D_SENT_BY_CRON         Sent by Cron Daemon
./65_debian.cf (END)
```

Figure 2: **Debian adds its own SpamAssassin tests to the already comprehensive list.**

to Bogofilter's, as well as white and black lists for individual customization.

Additionally, /etc/spamassassin includes a test developed for Debian that looks for spam involving anacron, cron, and debconf (Figure 2), as well as plugins installed with each recent version. Each test assigns an email message a positive or negative value, which is added to the results of other tests to determine whether the email is ham or spam. Unlike Bogofilter, exactly what these values represent is uncertain, although considering many users probably have no understanding of Bayesian analysis, much the same could also be said for Bogofilter, of course.

With all these tests, SpamAssassin exemplifies the basic security principle of "defense in depth." Unlike Bogofilter, it does not rely on one or two approaches, but on a wide variety of defenses. A piece of spam might slip by a single SpamAssassin test, but the odds of it slipping by all of them is unlikely.

Context Is Everything

Both Bogofilter and SpamAssassin are available as plugins for major email readers and generally require little customization. Both also have high success rates. However, because black and white lists greatly improve each filter's accuracy, be wary of the various comparisons online. Your own results are likely to be very different from those posted, especially before you have trained the filter to suit your personal email.

In fact, the filters are so different in their approaches and so dependent on how they are trained that deciding in any objective sense which one is most effective is almost impossible. To some extent, your decision as to which filter to use may depend on whether you prefer Bogofilter's single, all-encompassing approach or SpamAssassin's defense in depth.

Even more importantly, your choice will depend on context. To start, if the speed of filtering matters, Bogofilter is much faster than SpamAssassin for the simple reason that it runs fewer tests. If you ordinarily receive several hundred email messages in the first download of the day, SpamAssassin runs so many tests that you might be unable to access your email for five minutes – a delay that you might consider worse than manually deleting spam.

By contrast, in my experience, Bogofilter requires several days of training before it reaches full effectiveness. On one hand, stopping to train Bogofilter in the middle of other tasks can be a nuisance, especially because it seems to require several examples before it recognizes posts on a mailing list as ham. On the other hand, SpamAssassin is so comprehensive that it generally identifies spam more accurately without training. If you prefer to minimize training, SpamAssassin is probably the filter you want.

Another consideration is how many false positives you have once your filter of choice has been trained. My experience is that, once trained, Bogofilter has fewer false positives. Just as Graham observed,adding more rules, the way SpamAssassin does, beyond a certain point seems to increase false positives.

Still another consideration is that SpamAssassin is reactive. It adds tests in response to the latest tactics used by spammers but appears to be slower to discard tests that are no longer needed – if it does so at all. Similarly, if new spamming tactics appear, you might temporarily have less effective filtering until a new software release is made. However, because Bogofilter relies on probability rather than on spam characteristics, it might not have the same problems – at least not to the same extent.

As you can see, the decision of which filter to use has no absolute answer. However, once you understand how both filters work, you can at least make a more informed choice to accommodate your preferences and your needs. If nothing else, you can choose the lesser of two evils. ■■■

INFO

[1] Bogofilter: *http://bogofilter.sourceforge.net/*

[2] SpamAssassin: *http://spamassassin.apache.org/*

[3] DSPAM: *https://en.wikipedia.org/wiki/DSPAM*

[4] SpamBayes: *http://spambayes.sourceforge.net/*

[5] Thomas Bayes: *https://en.wikipedia.org/wiki/Bayesian*

[6] "A Plan for Spam" by Paul Graham: *http://www.paulgraham.com/spam.html*

[7] Robinson's refinements: *http://www.linuxjournal.com/article/6467*

UBUNTU user

MORE UBUNTU!

Linux Mint 17.3 + Kubuntu 15.10

The two most popular Ubuntu derivatives on one double-sided DVD!

Mint + Kubuntu THE 2 BEST UBUNTU DERIVATIVES

UBUNTU user

EXPLORING THE WORLD OF UBUNTU

DISCOVER THE WORLDS BEYOND UBUNTU

ALTERNATIVES

- 4 Mini distros you have to try
- Build your own distro with PCLinux OS
- Bodhi gets a new desktop: Moksha
- Power Squared: MX Linux
- BunsenLab's "Hydrogen" - So light it floats

SyncThing Sync all the things between all your devices

PLUS

Back up and keep all your stuff safe with Qt4-fsarchiver

Come to grips with the Scribus desktop publishing program

Design, build & 3D print with OpenSCAD

Computer running slow? Troubleshoot with Dstat

US$ 15.99 CAN$ 17.99 Spring 2016

DISCOVERY GUIDE

New to Ubuntu? Check out our special section for first-time users! p. 83

- How to install Ubuntu 15.10
- Get all your multimedia working
- Go online with NetworkManager
- Package management

SPRING 2016 WWW.UBUNTU-USER.COM

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 sys admin's daily grind: Prettyping and Asciiflow

Block Heroes

Columnist Charly is delighted that people still program useful tools for the terminal. Here, he looks at one tool that transforms boring ping data into colorful statistics and another that publishes a construction set for ASCII graphics on the network. *By Charly Kühnast*

One weapon for command-line warriors is Prettyping [1], a shell script that wraps around the ping command. It reads its tasks, keeps a record of run times and packet losses, and shows at the command line in block graphics the average values since starting the tool and for the past 60 seconds (Figure 1).

The script runs on any system with Bash and Awk (i.e., also on OS X and probably also in the new Linux environment on Windows 10). Prettyping detects whether it is running in a terminal and how wide the terminal is, then scales the output accordingly. If you think the output is a little too clownish, you can switch to a more staid monochrome display using `--nocolor`. Prettyping passes on to ping any parameters that it isn't familiar with.

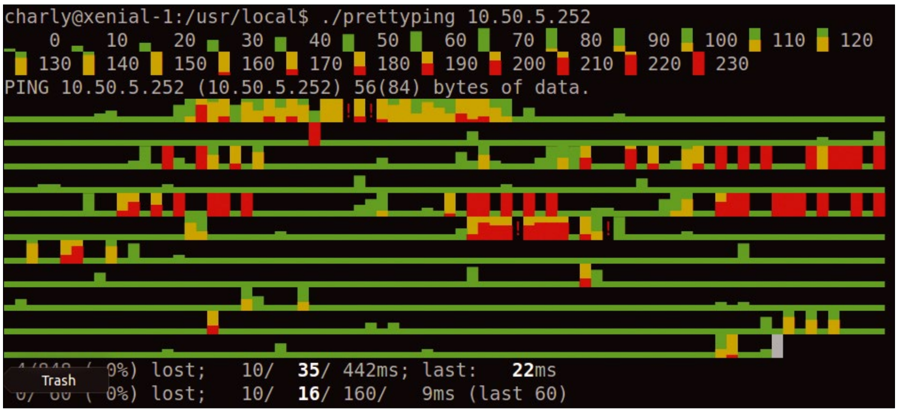

```
charly@xenial-1:/usr/local$ ./prettyping 10.50.5.252
    0 ▁ 10 ▂ 20 ▃ 30 ▄ 40 ▅ 50 ▆ 60 ▇ 70 █ 80 ▁ 90 ▂ 100 ▃ 110 ▄ 120
 ▅ 130 ▆ 140 ▇ 150 █ 160 ▁ 170 ▂ 180 ▃ 190 ▄ 200 ▅ 210 ▆ 220 ▇ 230
PING 10.50.5.252 (10.50.5.252) 56(84) bytes of data.
          0%) lost;   10/  35/ 442ms; last:   22ms
 0/ 60 (  0%) lost;   10/  16/ 160/    9ms (last 60)
```

Figure 1: It's a colorful world folks: Prettyping visualizing ping statistics in the terminal.

Everything <--|__ASCII__|

If you read RFCs, you will occasionally see small ASCII graphics that show connections more compactly than is possible with sentences. Authors typically painstakingly create such charts with boxes and arrows with ASCII symbols, such as plus and minus signs, (back-)slashes, and pipes. Naturally, at some point, various ASCII graphic victims have written tools – but none are as easy and intuitive to use as Asciiflow [2].

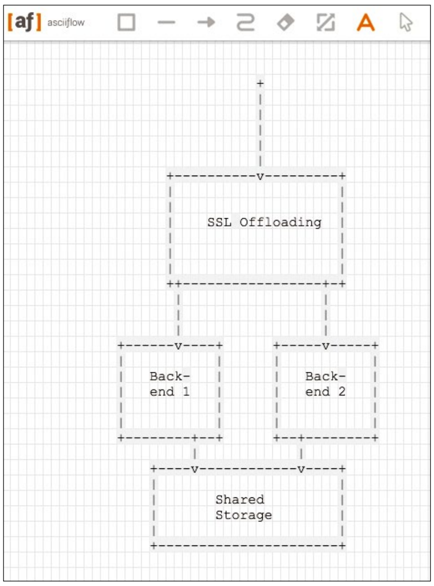

```
                  +
                  |
                  |
                  |
                  |
        +---------v---------+
        |                   |
        |   SSL Offloading  |
        |                   |
        |                   |
        ++----------------+-+
         |                |
         |                |
     +---v----+      +----v---+
     |  Back- |      |  Back- |
     |  end 1 |      |  end 2 |
     |        |      |        |
     +------+-+      +-+------+
            |          |
       +----v----------v----+
       |     Shared         |
       |     Storage        |
       |                    |
       +--------------------+
```

Figure 2: A few boxes drawn in Asciiflow often say more than a thousand words.

Asciiflow is a website that at first looks like a blank sheet of graph paper. In a toolbar at the top, you can select boxes, lines, arrows, text, and so on, and then simply draw on the blank sheet using the mouse (Figure 2). Once you are happy with your work, you just press the export symbol and – hey, presto – the finished ASCII graphic appears in your clipboard. Asciiflow also has an import function. My verdict on it: \o/. ■■■

INFO

[1] Prettyping: *http://denilson.sa.nom.br/prettyping/*

[2] Asciiflow: *http://asciiflow.com*

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.

For Linux, Windows, and Mac OS X
Ubuntu Studio
GIMP SPECIAL 2016
LINUX MAGAZINE
Look in the gimp/ folder for Windows and Mac OS installation files.

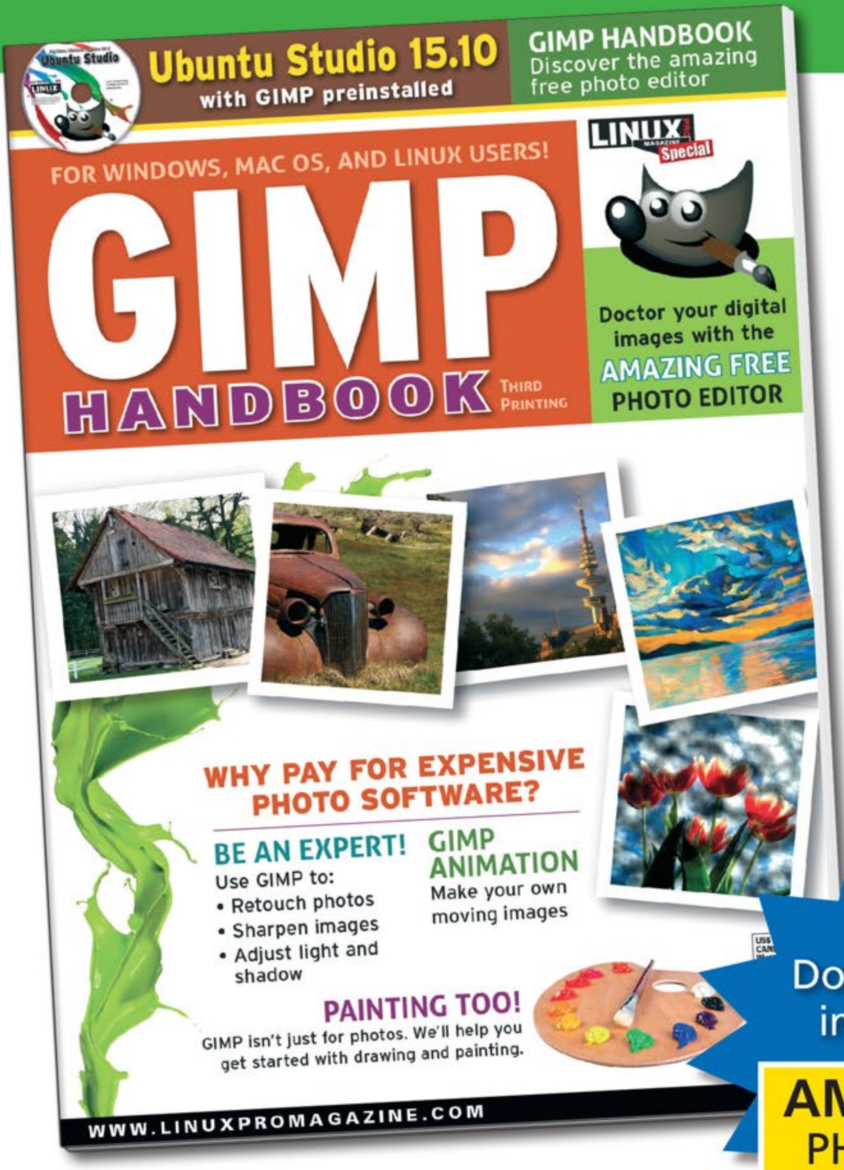

Ubuntu Studio 15.10 with GIMP preinstalled

GIMP HANDBOOK Discover the amazing free photo editor

LINUX MAGAZINE Special

FOR WINDOWS, MAC OS, AND LINUX USERS!

GIMP HANDBOOK THIRD PRINTING

Doctor your digital images with the AMAZING FREE PHOTO EDITOR

WHY PAY FOR EXPENSIVE PHOTO SOFTWARE?

BE AN EXPERT! Use GIMP to:

- Retouch photos
- Sharpen images
- Adjust light and shadow

GIMP ANIMATION Make your own moving images

PAINTING TOO! GIMP isn't just for photos. We'll help you get started with drawing and painting.

WWW.LINUXPROMAGAZINE.COM

Static code analyzers for JavaScript, PHP, Python, and the Linux shell

Script Doctor

Admins daily use scripts to automate tasks, generate web content, collect and parse data, and perform many other tasks. A few sophisticated tools can tell admins where script problems lurk. *By Tim Schürmann*

Administrators are likely to throw together a shell script quickly during the stress of everyday admin life, but sometimes it takes only a few days before the script self-destructs. Meticulous, time-consuming, and typically hectic troubleshooting commences. Thank goodness for static code analyzers.

Code analyzers examine source code for errors and typical problems, such as typos that a human is likely to overlook – from uninitialized variables to incorrectly used semicolons. Predefined test rules decide whether an error exists, and programmers can specify their own test criteria in some cases.

These code analyzers are referred to as "static" because they only see the source code; they cannot make any guesses about future performance. Some tools that perform data flow analysis and track variables through the program code can find unused or unnecessary variables.

For dynamically typed languages, such as Python, the variable type is only defined at run time, so the analysis tools need to guess what data could eventually reside in what variables. The result is that they sometimes mark as an error what is actually correct (false positive).

Linty

Many static code analyzers evaluate bad programming practices and thus act as style checkers as well. The analysis is based on commonly accepted coding style guides (e.g., the PEP 8 [1] standard, in the case of Python). Some tools generate statistics and make suggestions for improvements. For example, if the analyzer detects many identical rows, it can advise the developer to outsource them into a separate function.

Most static code analyzers are command-line programs, so the developer can include them in (shell) scripts or their own toolchain. Only a few tools include a front end, which usually just

TABLE 1: JavaScript Static Code Analyzers

Name	License	URL
Closure Tools	Apache 2.0	*https://developers.google.com/closure/*
Flow	BSD	*http://flowtype.org*
ESLint	MIT	*http://eslint.org*
JSHint	MIT Expat and JSON	*http://jshint.com*
JSLint	Modified MIT (open source, non-free)	*http://jslint.com*
JSPrime	MIT	*https://github.com/dpnishant/jsprime*
plato	MIT	*https://github.com/es-analysis/plato*
TAJS	Apache 2.0	*https://github.com/cs-au-dk/TAJS*

Lead Image © Ewa Walicka, Fotolia.com

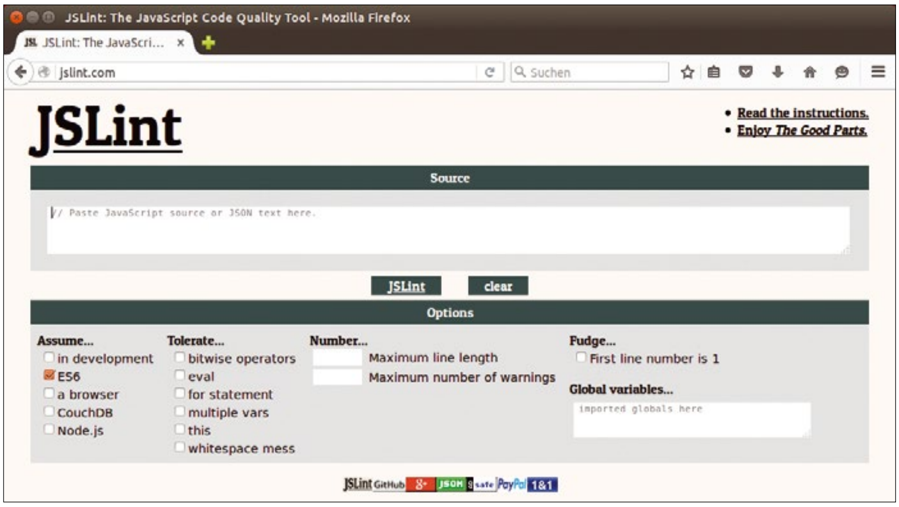

JSLint: The JavaScript Code Quality Tool - Mozilla Firefox

jslint.com

JSLint

- Read the instructions.
- Enjoy *The Good Parts.*

Source

```
// Paste JavaScript source or JSON text here.
```

JSLint | clear

Options

Assume...
☐ in development
☑ ES6
☐ a browser
☐ CouchDB
☐ Node.js

Tolerate...
☐ bitwise operators
☐ eval
☐ for statement
☐ multiple vars
☐ this
☐ whitespace mess

Number...
Maximum line length
Maximum number of warnings

Fudge...
☐ First line number is 1

Global variables...
imported globals here

Figure 1: Enter the JavaScript code in the large text box on the JSLint website and press the *JSLint* button to run a check.

shows the output of the command-line version. Ideally, the tools integrate text editors and IDEs, so the programmer sees coding errors when entering the code.

Lint, which was one of the first static code analyzers, took a close look at C programs, and its name prompted the similar naming of other tools. For example, Pylint takes cares of Python programs. Apart from the names and goals, these "linters" have nothing in common.

JavaScript

JavaScript has several linters (Table 1), such as JSLint by Douglas Crockford (Figure 1). The tool is written in JavaScript, and you can try it out directly on the home page; the source code is on GitHub [2]. JSLint checks the style of programming in addition to syntax and reveals some structural problems.

The tool takes a more restrictive approach, in that it enforces version 6 of the ECMAScript standard, which means a semicolon at the end of each statement and no allowance for the == operator. JSLint is released under a modified MIT license. Users can use the tool only for good, but not for "evil purposes." The Free Software Foundation classifies the license as non-free.

The alternative JSHint, like JSLint, checks syntax, complains about bad programming style, and reveals typical problems, such an implicit type conversion. JSHint in the current version supports ECMAScript 3, 5.1, and 6. You can try JSHint directly at the project homepage (Figure 2). Developers can integrate JSHint into their own web pages as a JavaScript module.

A command-line version of JSHint requires Node.js. JSHint feels at home in many text editors, such as Vim, Emacs, Sublime, Atom, and Brackets. The tool is released under the JSON license, which has provisions similar to the license used with JSLint.

ESLint is released under a plain vanilla MIT license. In contrast to JSLint and JSHint, the ESLint tool packs each check rule into a separate plugin that the programmer can enable and disable as required – even at run time. Users can contribute their own plugins, as well. The supplied rules cover syntax errors, check for best practices (e.g., the use of `eval()`), complain about overly complex constructs, and point out bad programming style.

The tool supports ECMAScript 6; however, users need to enable support for the standard explicitly. You can test ESLint directly on the website [3]; a command-line version is also available that requires Node.js.

Google provides various tools for JavaScript programmers as closure tools, including the Closure Compiler (Figure 3). Closure Compiler converts JavaScript statements into "compact, high performance" code and also performs syntax checks and variable type checks, warning you about common problems.

Closure Compiler is available online on Appspot [4], and it runs locally as a Java program in a command-line version or via the REST API. It supports ECMAScript 3, 5, and parts of version 6. As a supplement, Google provides the Closure Linter, which investigates programming style [5] according to the Google JavaScript style guide [6] and calls foul in case of missing semicolons, among other problems.

Flow, developed by Facebook, primarily focuses on data flow and examines variable content and type. It complains, for example, if a function tries to compute something the programmer has assigned to a string. That said, Flow only checks JavaScript files that the developer has identified with an appropriate comment. What's more, you need to add annotations to the variables to state their desired types.

Using the appropriate plugins, you can integrate the tool with Vim, Emacs, and

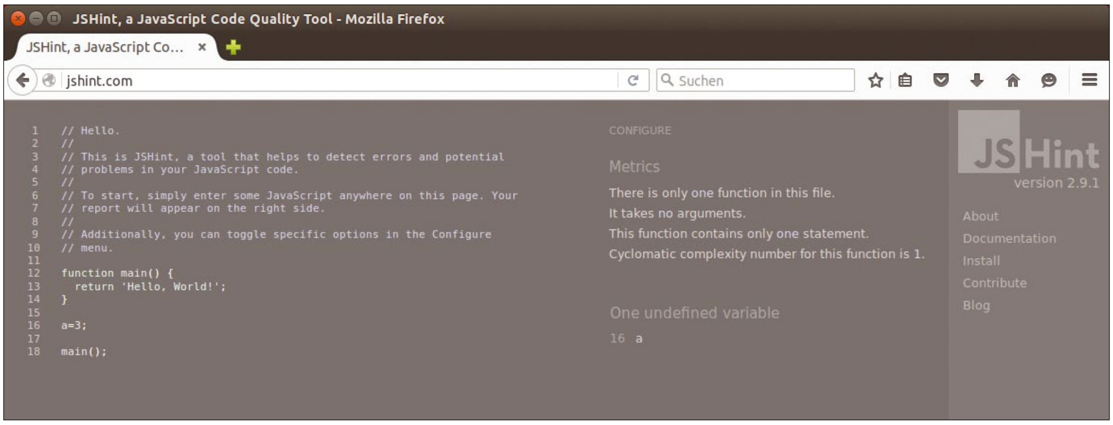

JSHint, a JavaScript Code Quality Tool - Mozilla Firefox

jshint.com

```
// Hello.
//
// This is JSHint, a tool that helps to detect errors and potential
// problems in your JavaScript code.
//
// To start, simply enter some JavaScript anywhere on this page. Your
// report will appear on the right side.
//
// Additionally, you can toggle specific options in the Configure
// menu.

function main() {
  return 'Hello, World!';
}

a=3;

main();
```

CONFIGURE

Metrics

There is only one function in this file.
It takes no arguments.
This function contains only one statement.
Cyclomatic complexity number for this function is 1.

One undefined variable

16 a

Figure 2: JSHint is an alternative to JSLint, which also runs in a browser. The software detects bad programming style and implicit type conversions.

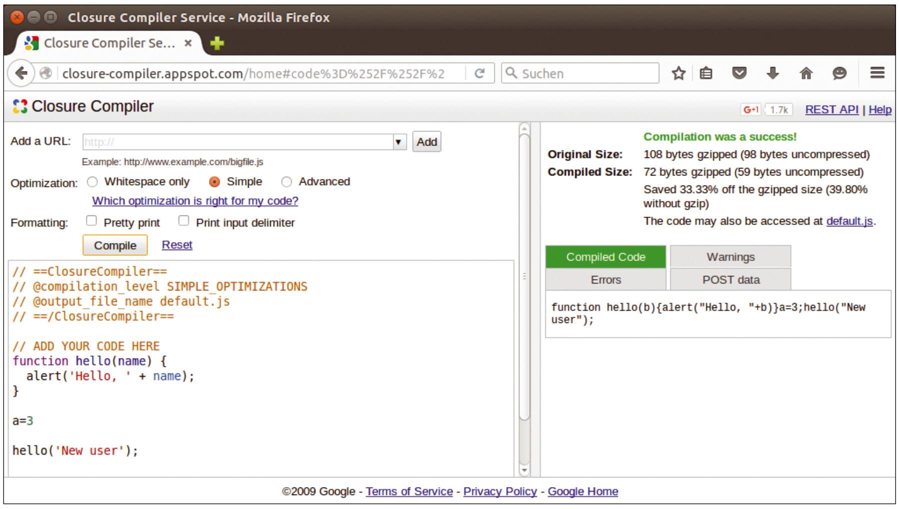
Closure Compiler Service - Mozilla Firefox
Closure Compiler Se...
closure-compiler.appspot.com/home#code%3D%252F%252F%2
Suchen
Closure Compiler
G+1 1.7k
REST API | Help

Add a URL: http:// Add
Example: http://www.example.com/bigfile.js
Optimization: ○ Whitespace only ● Simple ○ Advanced
Which optimization is right for my code?
Formatting: □ Pretty print □ Print input delimiter
Compile Reset

```
// ==ClosureCompiler==
// @compilation_level SIMPLE_OPTIMIZATIONS
// @output_file_name default.js
// ==/ClosureCompiler==

// ADD YOUR CODE HERE
function hello(name) {
  alert('Hello, ' + name);
}

a=3

hello('New user');
```

Compilation was a success!

Original Size: 108 bytes gzipped (98 bytes uncompressed)
Compiled Size: 72 bytes gzipped (59 bytes uncompressed)
Saved 33.33% off the gzipped size (39.80% without gzip)
The code may also be accessed at default.js.

```
function hello(b){alert("Hello, "+b)}a=3;hello("New
user");
```

©2009 Google - Terms of Service - Privacy Policy - Google Home

Figure 3: The Closure Compiler removes all unnecessary characters from source code and corrects it at the same time: Here, it has removed and added semicolons.

Nuclide. Moreover, you can run a Flow server in the background and incrementally re-investigate JavaScript files as they are changed. The tool itself is written in OCaml.

The development of JSPrime has already been on ice for two years; however, the tool explicitly checks JavaScript code for some security vulnerabilities. Additionally, JSPrime can handle minified Java code. You control JSPrime via a web client.

Many other tools available on the Internet simply leverage other applications. For example, TAJS relies on the Closure Compiler, whereas Plato puts the JSHint and ESLint team to work.

PHP

PHP developers can also rely on various helpers (Table 2). PHP checks the source code itself for syntax errors by calling a `.php`file with the `-l` parameter set.

PHPLint [7] goes beyond this capability, handling both PHP5 and PHP7 applications (Figure 4). In addition to searching for syntax errors, PHPLint also starts data flow analysis, ensures that exceptions are correctly handled, verifies function signatures, and ensures consistent types.

Developers can help PHPLint by adding comments to the code that explicitly state the variable types. As a bonus, the tool will generate documentation on request directly from the source code. PHPLint is written in PHP and requires the PHP CLI command-line interpreter; some distributions put the command-line interpreter in a separate package.

PHP mess detector (PHPMD; Figure 5) positions itself as an alternative to PHPLint. The mess detector is based on the JMD static code analyzer for Java [8]. PHPMD searches the source code for errors, suboptimal code, unnecessarily complicated expressions, and unused parameters or methods. PHPMD outputs the results either in text, HTML, or XML format. Although the tool has been around since 2009, the developers still consider it a young project and apologize for the sparse number of rules supplied.

Phan, another contender programmed entirely in PHP7, really is quite young. Among other things, Phan tests whether source code will run in PHP7. In its analysis, Phan also considers PHPDOC comments, such as `@depricated` or `@param`, as well as generics, namespaces, traits, and variadic functions.

RIPS was popular for many years. Unlike the other tools, RIPS is a web application written in PHP. The PHP programmer uses a form to pass in the script you wish to test, and RIPS then outputs a variety of facts with a colorful layout (Figure 6). In particular, it looks for vulnera-

TABLE 2: Static Code Analyzers for PHP

Name	License	URL
Phan	MIT	*https://github.com/etsy/phan*
PHP_CodeSniffer	BSD	*https://github.com/squizlabs/PHP_CodeSniffer*
PHPLint	BSD	*http://www.icosaedro.it/phplint/index.html*
PHPMD	BSD	*https://phpmd.org*
RIPS	GNU GPLv3	*http://rips-scanner.sourceforge.net*

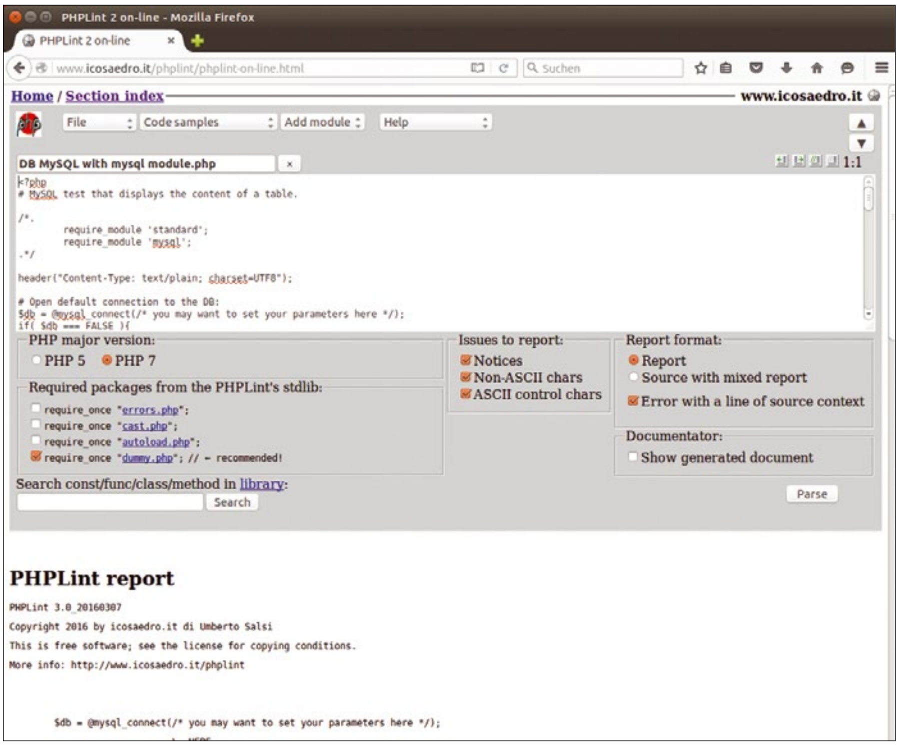
PHPLint 2 on-line - Mozilla Firefox
PHPLint 2 on-line
www.icosaedro.it/phplint/phplint-on-line.html
Suchen
Home / Section index
www.icosaedro.it
File
Code samples
Add module
Help
DB MySQL with mysql module.php
1:1

```
<?php
# MySQL test that displays the content of a table.

/*.
        require_module 'standard';
        require_module 'mysql';
.*/

header("Content-Type: text/plain; charset=UTF8");

# Open default connection to the DB:
$db = @mysql_connect(/* you may want to set your parameters here */);
if( $db === FALSE ){
```

PHP major version: ○ PHP 5 ● PHP 7

Required packages from the PHPLint's stdlib:
□ require_once "errors.php";
□ require_once "cast.php";
□ require_once "autoload.php";
☑ require_once "dummy.php"; // - recommended!

Issues to report:
☑ Notices
☑ Non-ASCII chars
☑ ASCII control chars

Report format:
● Report
○ Source with mixed report
☑ Error with a line of source context

Documentator:
□ Show generated document

Search const/func/class/method in library: Search

PHPLint report

```
PHPLint 3.0_20160307
Copyright 2016 by icosaedro.it di Umberto Salsi
This is free software; see the license for copying conditions.
More info: http://www.icosaedro.it/phplint


        $db = @mysql_connect(/* you may want to set your parameters here */);
```

Figure 4: You can try out PHPLint directly in the browser: Drop PHP code into the input field and then run the parser by pressing the button bottom at right.

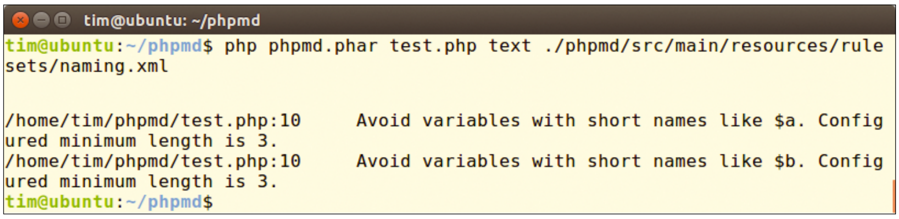

```
tim@ubuntu: ~/phpmd
tim@ubuntu:~/phpmd$ php phpmd.phar test.php text ./phpmd/src/main/resources/rule
sets/naming.xml

/home/tim/phpmd/test.php:10     Avoid variables with short names like $a. Config
ured minimum length is 3.
/home/tim/phpmd/test.php:10     Avoid variables with short names like $b. Config
ured minimum length is 3.
tim@ubuntu:~/phpmd$
```

Figure 5: Developers need to tell PHPMD the file to be examined, the output format, and the desired rule set. Here PHPMD checks variable names.

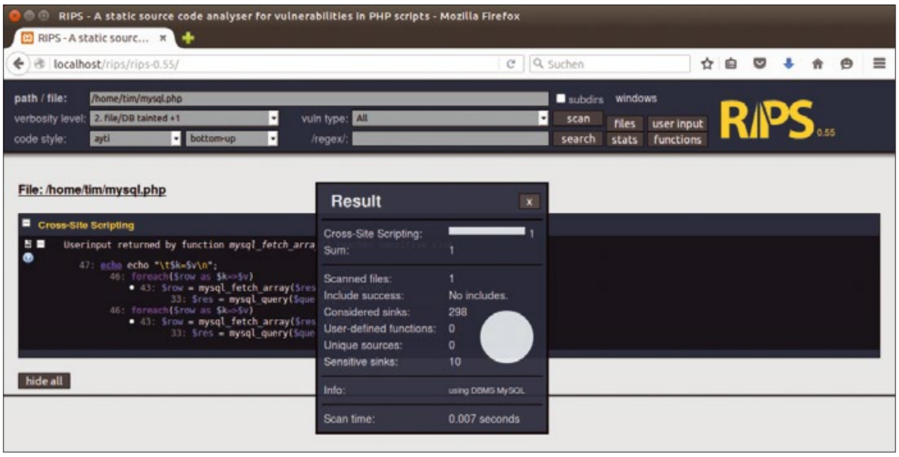

RIPS - A static source code analyser for vulnerabilities in PHP scripts - Mozilla Firefox

localhost/rips/rips-0.55/

path / file: /home/tim/mysql.php
verbosity level: 2. file/DB tainted +1
code style: ayti bottom-up
vuln type: All
/regex/:

File: /home/tim/mysql.php

Cross-Site Scripting

```
Userinput returned by function mysql_fetch_arra
    47: echo echo "\t$k=$v\n";
        46: foreach($row as $k=>$v)
            • 43: $row = mysql_fetch_array($res
                33: $res = mysql_query($que
        46: foreach($row as $k=>$v)
            • 43: $row = mysql_fetch_array($res
                33: $res = mysql_query($que
```

Result

Cross-Site Scripting: 1
Sum: 1

Scanned files: 1
Include success: No includes.
Considered sinks: 298
User-defined functions: 0
Unique sources: 0
Sensitive sinks: 10

Info: using DBMS MySQL

Scan time: 0.007 seconds

Figure 6: RIPS is a web application written in PHP; above all, it identifies vulnerabilities.

TABLE 3: Python Static Code Analyzers

Name	License	URL
Bandit	Apache 2	*https://github.com/openstack/bandit*
Flake8	MIT	*https://pypi.python.org/pypi/flake8*
Prospector	GNU GPLv2	*https://github.com/landscapeio/prospector*
Pyflakes	MIT	*https://github.com/pyflakes/pyflakes*
Pylint	GNU GPLv2	*https://www.pylint.org*

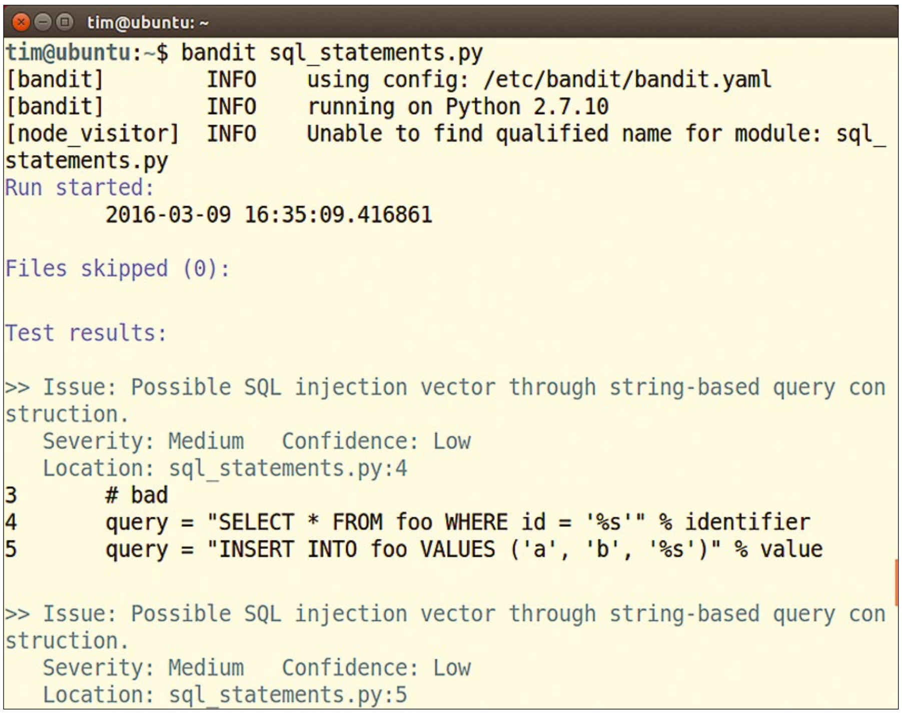

```
tim@ubuntu: ~
tim@ubuntu:~$ bandit sql_statements.py
[bandit]        INFO    using config: /etc/bandit/bandit.yaml
[bandit]        INFO    running on Python 2.7.10
[node_visitor]  INFO    Unable to find qualified name for module: sql_
statements.py
Run started:
        2016-03-09 16:35:09.416861

Files skipped (0):

Test results:

>> Issue: Possible SQL injection vector through string-based query con
struction.
   Severity: Medium   Confidence: Low
   Location: sql_statements.py:4
3       # bad
4       query = "SELECT * FROM foo WHERE id = '%s'" % identifier
5       query = "INSERT INTO foo VALUES ('a', 'b', '%s')" % value

>> Issue: Possible SQL injection vector through string-based query con
struction.
   Severity: Medium   Confidence: Low
   Location: sql_statements.py:5
```

Figure 7: Bandit investigates Python code for security vulnerabilities. The tool comes from the OpenStack universe, where Python plays a supporting role.

bilities, including SQL injections, cross-site scripting, and code execution.

The founder of RIPS stopped developing the tool in 2013, and only a few bug fixes have been incorporated in the past few years. Originally, a completely new version 1.0 was supposed to be developed, but there are no signs of it materializing. Because RIPS is not familiar with current attack techniques and cannot handle modern versions of PHP, you can only use it for an initial assessment.

Strictly speaking, PHP_CodeSniffer does not belong to the group of static code analyzers. The tool checks to see whether the source follows one or multiple coding standards. Among others, the PEAR standards, PSR-1, and PSR-2 are available. PHP_CodeSniffer also tests for a few typical sources of error, such as functions not properly declared. The tool returns the results either as a simple text message or in another format, such as XML, CSV, or JSON. Users can even define test settings in `php.ini`. For example, a call to

```
phpcs -d include_path=⮐
  .:/php/includes test.php
```

would add the `/php/includes` directory to the include path when checking the `test.php` file. PHP_CodeSniffer can process PHP7 applications and, as a bonus, can even parse JavaScript code, with some restrictions.

Python

Among the various Python linters (Table 3), Pylint is undoubtedly one of the classic static code analyzers. On request, the tool checks code for compliance with the PEP 8 Style Guide for Python Code. Pylint also helps with refactoring by tracking double code, among other things. On request, Pylint generates appropriate UML diagrams from the Python code. An optional `type` parameter even checks whether all of the parameters accepted by the Python script are consistent and properly documented for later users. Pylint also produces a variety of statistics that list the number of duplicate rows, for example.

Programmers can use plugins to extend the functionality of the tool. On request Pylint uses multiple CPU cores at the same time, speeding up the process, especially for large-scale source code.

TABLE 4: Shell Script Static Code Analyzers

Name	License	URL
bashate	Apache	*https://pypi.python.org/pypi/bashate/*
checkbashisms	GNU GPLv2	*http://ftp.halifax.rwth-aachen.de/debian/pool/main/d/devscripts/*
ShellCheck	GNU GPLv3	*https://github.com/koalaman/shellcheck*

You can also integrate Pylint with various IDEs and text editors, such as Emacs, Vim, and Eclipse, and it can be used with continuous integration tools such as Apycot, Hudson, or Jenkins. Its development is significantly driven by Logilab, a company that offers commercial services related to Pylint.

One alternative to Pylint is Pyflakes. Although it works faster than its competitors, Pyflakes does not check programming style. Also, the tool individually inspects each script file, so it does not see the bigger picture, meaning that it discovers fewer errors.

Its competitor Bandit (Figure 7) examines Python code for typical vulnerabilities; thus, it is recommended on top of Pylint and Pyflakes. The tool, which comes from the OpenStack universe, particularly examines XML processing, network code (e.g., FTP, Telnet, HTTP, and SSL connections), problematic SQL queries, and encryption. Users can enable and disable the completed tests individually or add their own tests.

Several tools for Python also harness other competitors for their analysis. For example, Flake8 tests Python scripts fed to it for errors with the help of Pyflakes, relying on PEP 8 for style checking and Ned Batchelder's McCabe script for generating statistics. Prospector deploys up to eight additional tools on top of Pylint and Pyflakes to check the source code and prepares output from these tools.

Shell

That just leaves the shell, for which a few helpers also exist (Table 4). Most shell interpreters support the `-n` command-line parameter, which is used to test the script passed in for syntax errors. The `-u` parameter tells the shell to stop working as soon as it discovers an undefined variable.

ShellCheck (Figure 8) reveals many more script errors, including incorrectly used quotes, problematic test queries, typical newbie mistakes, and even frequently used but sub-optimal commands. ShellCheck also checks the programming style to discover whether as many shell interpreters as possible can handle the script. On request, ShellCheck can be integrated with Vim, Emacs, Sublime, and Atom. The tool presents the results either as plain text, XML data, or in JSON format.

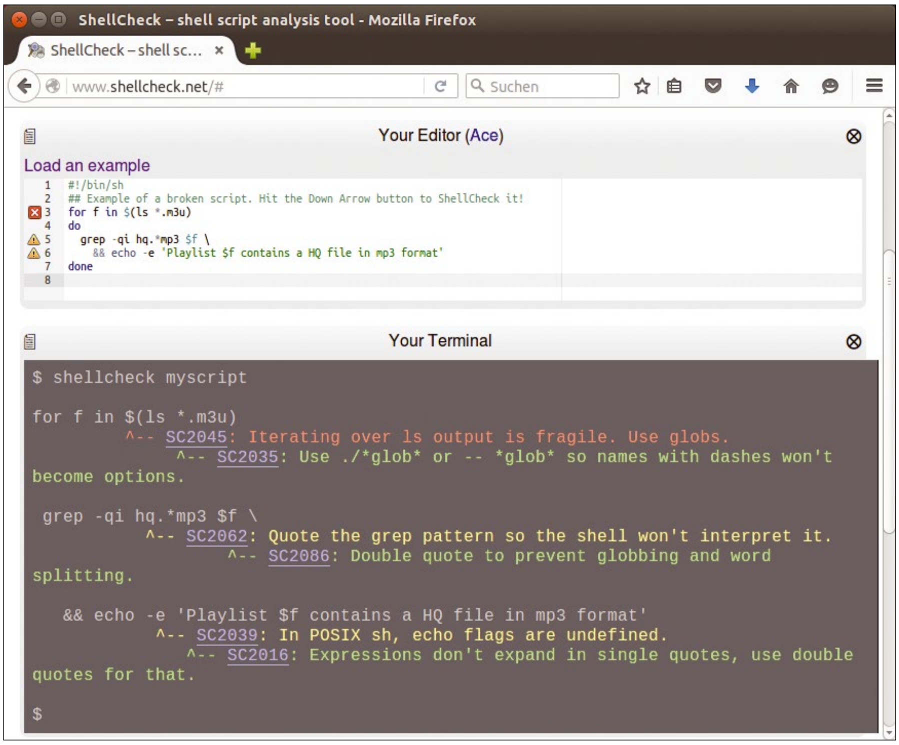

```
#!/bin/sh
## Example of a broken script. Hit the Down Arrow button to ShellCheck it!
for f in $(ls *.m3u)
do
  grep -qi hq.*mp3 $f \
    && echo -e 'Playlist $f contains a HQ file in mp3 format'
done
```

```
$ shellcheck myscript

for f in $(ls *.m3u)
         ^-- SC2045: Iterating over ls output is fragile. Use globs.
              ^-- SC2035: Use ./*glob* or -- *glob* so names with dashes won't
become options.

 grep -qi hq.*mp3 $f \
           ^-- SC2062: Quote the grep pattern so the shell won't interpret it.
                   ^-- SC2086: Double quote to prevent globbing and word
splitting.

   && echo -e 'Playlist $f contains a HQ file in mp3 format'
            ^-- SC2039: In POSIX sh, echo flags are undefined.
               ^-- SC2016: Expressions don't expand in single quotes, use double
quotes for that.

$
```

Figure 8: On the ShellCheck website, developers can even test their program code directly. The evaluation is carried out while typing.

If you only want to check whether your shell script will run on as many interpreters as possible, you can use checkbashisms as an alternative. The tool is part of the *devscripts* package in some distributions and as a separate package in others. Checkbashisms complains about all constructs that are not POSIX compliant.

Finally, bashate from the OpenStack developers performs a few style checks and checks for a few dangerous commands and discouraged commands.

Conclusions

Static code analyzers quickly identify errors that script programmers can easily overlook in the heat of battle. However, the tools are only as smart as their rules, and they usually only cover a small part of all possible problematic constructs.

The tools presented in this article do not examine program logic. In other words, they will not understand whether an online shop needs to multiply the prices in the shopping cart or just add them. Developers should thus not consider static code analyzers as a panacea, but merely as a useful building block in a comprehensive test plan. ■■■

INFO

[1] PEP 8 – Style Guide for Python Code: *https://www.python.org/dev/peps/pep-0008/*

[2] Source code for JSLint: *https://github.com/douglascrockford/JSLint*

[3] ESLint online: *http://eslint.org/demo/*

[4] Closure Compiler online: *http://closure-compiler.appspot.com/home*

[5] Closure Linter: *https://developers.google.com/closure/utilities/*

[6] Google JavaScript Style Guide: *https://google.github.io/styleguide/javascriptguide.xml*

[7] PHPLint online: *http://www.icosaedro.it/phplint/phplint-on-line.html*

[8] JMD: *https://pmd.github.io*

Managing multiple systems in parallel with SaltStack

Pulling the Strings

Professionals often turn to SaltStack to manage server farms in parallel. When used properly, the same technology also saves work in small networks. *By Valentin Höbel*

Even in smaller IT environments, managing systems as consistently as possible pays dividends. The effort required to learn and use massive tools for configuration management (e.g., Ansible, Chef, Puppet) rarely pays off, because you frequently need to complete simple, one-time tasks; when these present themselves, writing recipes (Chef) or manifests (Puppet) wastes too much time.

Alternatively, many admins use self-made scripts, which they execute to hosts in rotation from a list that is laborious to maintain. Although these methods work, you must take the particulars of the respective platforms into account and cope with possible timeouts while opening connections. This effort can be reduced by turning to a proven solution. The right tool should be able not only to abstract the differences of distributions and operating systems but also to provide its own communication channel and offer modules with pre-built commands and macros. Normally, projects from the orchestration and remote execution environment offer something similar. Thanks to its simple operation, speed, and scalability, SaltStack [1] stands out from the competition in these cases.

AUTHOR

Valentin Höbel works as a cloud architect for the VoIP specialists NFON AG in Munich, Germany. If he is not playing table soccer in his free time, you are likely to find him investigating the latest open source technologies.

SaltStack

The systems and software architect Thomas S. Hatch faced the challenge of centrally maintaining an infrastructure that was partly inconsistent. To master it, he wrote his own software in Python; SaltStack (Salt, for short) emerged and was eventually published as an open source project in March 2011.

SaltStack is based on a central master, whose commands are executed on the target systems by Salt “minions.” The communication between the master and the minions does not take place via SSH; rather, it relies on the well-known ZeroMQ [2] message bus library using an asynchronous approach. By design, various modules offer a full palette of commands and command sequences with which you can execute specific tasks. Additionally, another module facilitates running Linux commands. As soon as you initiate tasks with the master, the tool monitors execution on the minions and subsequently displays the outcome.

It does not matter whether you only use Salt on Linux, run different distributions on the master host and the target systems, or even go with FreeBSD. Salt supports a wide range of operating systems, although some of them (e.g., Windows) only as minions [3]. A list of all the operating systems in question can be found online [4]. You can deliver standardized commands with the modules supplied, and the modules on the minions then translate these into platform-specific commands.

Installation

As the master, you use a physical or virtual machine; even for large installations, the requests are so small that the master does not overtask the minions. Usually, you can install both the master and the minions from the official repository. However, where differences are too large, the available Salt versions can run into compatibility problems because of different module packages. You are better off running all minions compatibly with the same version of Salt. For the sake of ease in my tests, I prepared a set

Lead Image © perseomedusa, 123RF.com

of virtual machines (VMs) that can see one another. An Ubuntu system hosts the master, and SUSE, Debian, and Fedora VMs serve as minions. You can see a more detailed overview along with OS versions and IPs in Table 1.

On Ubuntu 14.04, you can install the master with the

```
sudo apt-get install salt-master
```

command. To install the correct minion on Debian 7, modify the package sources by adding the following line to the `/etc/apt/sources.list` file

```
deb http://debian.saltstack.com/debian ↩
    wheezy-saltstack main
```

and updating with `sudo apt-get update`. The minion is then installed by typing one of the following commands,

```
# On Debian
sudo apt-get install salt-minion
# On openSUSE 13.2
sudo zypper install salt-minion
# On Fedora 23
sudo dnf install salt-minion
```

depending on your Linux installation.

Configuration

On startup, the minions try to connect to the master by default. If they have not configured a target of this kind, the software automatically searches the DNS domain for a *salt* entry. You can therefore either create a corresponding A record that points to the master (in my test, the target was the IP address 192.168.178.39), or you can configure the Salt minions manually.

On any system running a minion, you can edit the `/etc/salt/minion` configuration file and search for the line that begins `#master:`. You then remove the comment character and modify the line to

```
master: 192.168.178.39
```

(Figure 1). Replace the IP address from the example with any address that shows on the master in your network. After the changes, you need to restart the minion.

If the local system or network has a firewall that potentially filters the communication between master and minions, open TCP ports 4505 and 4506 on all systems involved. In case of questions or problems, your best option is to check out the appropriate section in the official documentation [5].

The connection between the Salt master and its minions is encrypted, and both sides authenticate. Before the minions can accept commands from a master, they first need to store their keys on the master, and the master has to accept the keys. When a minion connects to a new master for the first time, this exchange occurs automatically. On the master, you can view the list of minions (Figure 2) that have submitted your key using:

TABLE 1: Lab VM Overview

Hostname	OS	IP	Role
ub1404	Ubuntu 14.04	192.168.178.39	Master
deb7	Debian 7	192.168.178.40	Minion
linux-x3b4	openSUSE 13.2	192.168.178.41	Minion
fedora23	Fedora 23	192.168.178.44	Minion

```
sudo salt-key -L
```

If you trust these minions straightaway, you can accept all keys in one fell swoop with the

```
sudo salt-key -A
```

command and confirm by typing *Y*. If you want to check the identity of the individual minions, on the other hand, help with `salt-key` can be found on the man page. Running `sudo salt-key -L` again verifies acceptance (Figure 2).

All minions listed in green beneath *Accepted Keys* will accept commands from the master. For test purposes, you should check that the master sees the minions online, using the

```
sudo salt-run manage.status
```

command. Figure 3 shows the result.

Choice of Target

Someone with an army of subordinates at their command will either direct their orders to all the group's members, or just a selection. The same holds for a Salt master: here, too, it decides which minions it communicates with.

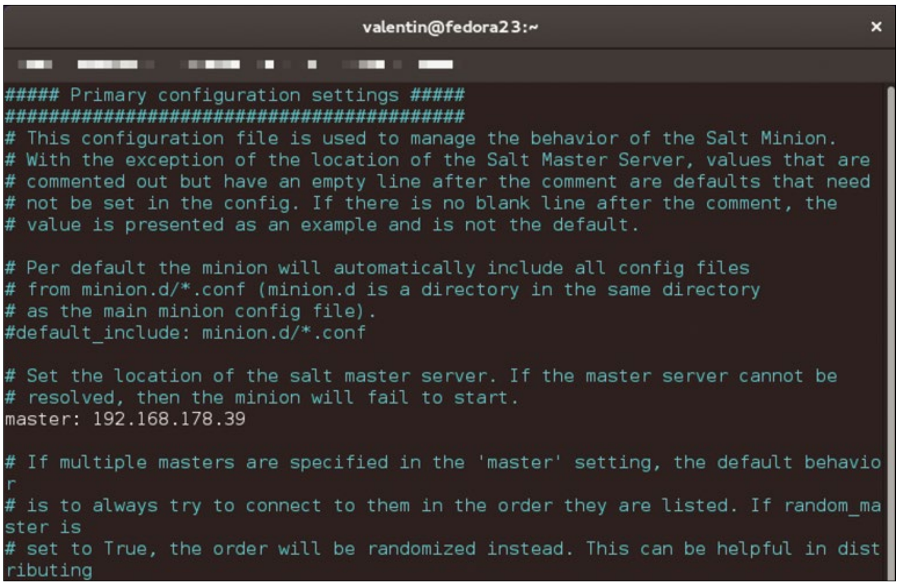

```
##### Primary configuration settings #####
##########################################
# This configuration file is used to manage the behavior of the Salt Minion.
# With the exception of the location of the Salt Master Server, values that are
# commented out but have an empty line after the comment are defaults that need
# not be set in the config. If there is no blank line after the comment, the
# value is presented as an example and is not the default.

# Per default the minion will automatically include all config files
# from minion.d/*.conf (minion.d is a directory in the same directory
# as the main minion config file).
#default_include: minion.d/*.conf

# Set the location of the salt master server. If the master server cannot be
# resolved, then the minion will fail to start.
master: 192.168.178.39

# If multiple masters are specified in the 'master' setting, the default behavio
r
# is to always try to connect to them in the order they are listed. If random_ma
ster is
# set to True, the order will be randomized instead. This can be helpful in dist
ributing
```

Figure 1: If automatic configuration fails, set the Salt master's options manually.

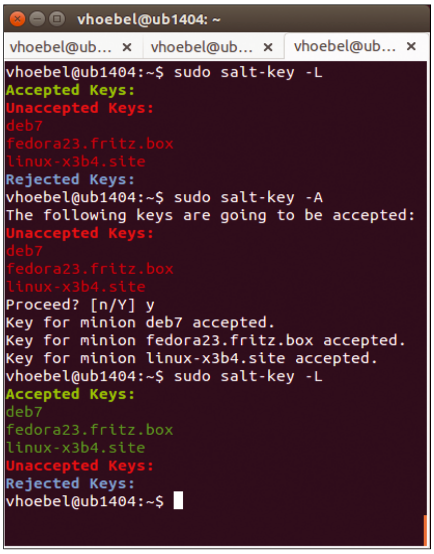

```
vhoebel@ub1404:~$ sudo salt-key -L
Accepted Keys:
Unaccepted Keys:
deb7
fedora23.fritz.box
linux-x3b4.site
Rejected Keys:
vhoebel@ub1404:~$ sudo salt-key -A
The following keys are going to be accepted:
Unaccepted Keys:
deb7
fedora23.fritz.box
linux-x3b4.site
Proceed? [n/Y] y
Key for minion deb7 accepted.
Key for minion fedora23.fritz.box accepted.
Key for minion linux-x3b4.site accepted.
vhoebel@ub1404:~$ sudo salt-key -L
Accepted Keys:
deb7
fedora23.fritz.box
linux-x3b4.site
Unaccepted Keys:
Rejected Keys:
vhoebel@ub1404:~$
```

Figure 2: Verification procedure: Salt makes sure the components know one another.

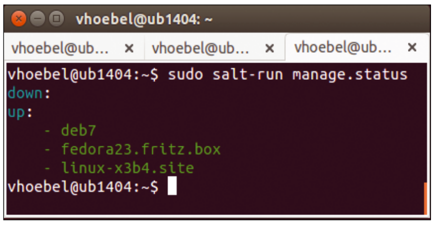

```
vhoebel@ub1404:~$ sudo salt-run manage.status
down:
up:
    - deb7
    - fedora23.fritz.box
    - linux-x3b4.site
vhoebel@ub1404:~$ 
```

Figure 3: If you have registered the minions properly on the master, nothing stands in the way of communication between the computers.

The primary command you will use on SaltStack is built to that effect. After entering the command (`salt`), you name the targets (choose all with `"*"`) and specify a `<module>.<function>` pair, which in this case checks the connections with the minions:

```
$ sudo salt "*" test.ping
linux-x3b4.site:
  True
fedora23.fritz.box:
  True
deb7:
  True
```

Targeting the minions (i.e., choosing the targets) is one of two central elements in working with the command. Salt provides a wide range of options just to reach a particular host:

```
$ sudo salt "deb7" test.version
deb7:
  2015.5.3
```

With Salt, you can also select minions using regular expressions and patterns. For example, entering `d*` lets you incorporate all systems whose names begin with the letter `d`.

Alternatively, `file[1-5]` can be applied to the minions with the `file1` to `file5` identifier.

To address the three minions in the test configuration shown here, for instance, you could use the starting letters of their names, after which any character string could follow:

```
$ sudo salt "[d,l,f]*" test.version
linux-x3b4.site:
  2014.1.11
deb7:
  2015.5.3
fedora23.fritz.box:
  2015.5.5
```

You can link together almost any number of criteria to select a system on what SaltStack calls compound matchers. Activate these keywords with `-C` before choosing the targets. Inside the quotes, you are free to enter several pieces of information, linking them with `and` or `and not`. This simple example shows how to select all minions then exclude the Debian host:

```
$ sudo salt -C "* and not deb*" test.ping
linux-x3b4.site:
  True
fedora23.fritz.box:
  True
```

You can select minions on the basis of attributes like the operating system, the IP address, and many other factors. More information can be found in the “Attributes” section.

Incidentally, if you get the impression that a minion name is not optimal, you can set the name in the `/etc/salt/minion_id` file. Generally, you can assign any name, although for day-to-day use, a fully qualified domain name (FQDN) or a combination of role and number have proven successful.

The `test` module used here mostly does not see day-to-day use. Instead, it helps with debugging and fast connectivity checking. Further functions of `test` can be found in the online documentation [6]. Likewise, if you want to take a comprehensive look at the choice of target systems on Salt, there are relevant documents on the web [7].

Attributes

A Salt minion is aware of certain information about the system on which it is running and stores this in so-called grains. Among other things, grains provide information about the operating system, the network configuration, the hardware, or the BIOS.

In the following example, the first command delivers a list of all grain names on all minions, whereas the second command provides their values, as seen in Figure 4:

```
sudo salt "*" grains.ls
sudo salt "*" grains.items
```

Sometimes you only need certain information, so you only need to retrieve single grains, such as the operating system:

```
$ sudo salt “*” grains.item os
deb7:
  os: Debian
linux-x3b4.site:
```

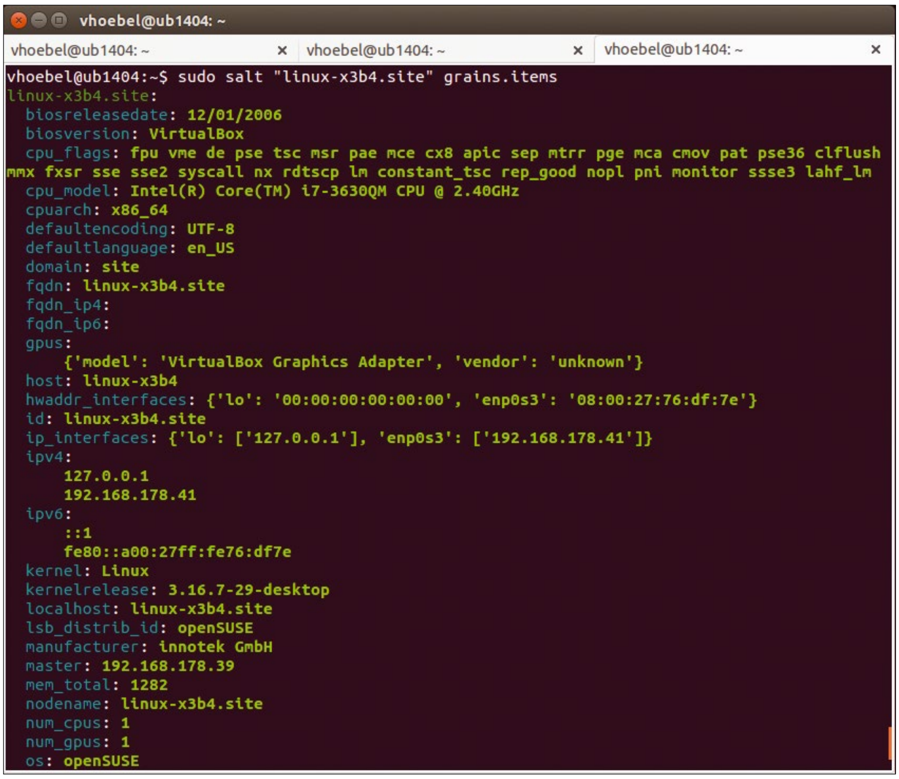

```
vhoebel@ub1404:~$ sudo salt "linux-x3b4.site" grains.items
linux-x3b4.site:
  biosreleasedate: 12/01/2006
  biosversion: VirtualBox
  cpu_flags: fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat pse36 clflush
mmx fxsr sse sse2 syscall nx rdtscp lm constant_tsc rep_good nopl pni monitor ssse3 lahf_lm
  cpu_model: Intel(R) Core(TM) i7-3630QM CPU @ 2.40GHz
  cpuarch: x86_64
  defaultencoding: UTF-8
  defaultlanguage: en_US
  domain: site
  fqdn: linux-x3b4.site
  fqdn_ip4:
  fqdn_ip6:
  gpus:
      {'model': 'VirtualBox Graphics Adapter', 'vendor': 'unknown'}
  host: linux-x3b4
  hwaddr_interfaces: {'lo': '00:00:00:00:00:00', 'enp0s3': '08:00:27:76:df:7e'}
  id: linux-x3b4.site
  ip_interfaces: {'lo': ['127.0.0.1'], 'enp0s3': ['192.168.178.41']}
  ipv4:
      127.0.0.1
      192.168.178.41
  ipv6:
      ::1
      fe80::a00:27ff:fe76:df7e
  kernel: Linux
  kernelrelease: 3.16.7-29-desktop
  localhost: linux-x3b4.site
  lsb_distrib_id: openSUSE
  manufacturer: innotek GmbH
  master: 192.168.178.39
  mem_total: 1282
  nodename: linux-x3b4.site
  num_cpus: 1
  num_gpus: 1
  os: openSUSE
```

Figure 4: By design, Salt recognizes numerous important attributes of a system on which a compatible minion is running.

```
  os: openSUSE
fedora23.fritz.box:
  os: Fedora
```

Furthermore, minions can be addressed on the basis of grains, which makes it possible to target minions across several relevant conditions, such as those with a specific CPU architecture or those with 984MB of working memory that do not have the same name as the openSUSE test system:

```
$ sudo salt -G "cpuarch:x86_64" test.ping
$ sudo salt -C "mem_total:984 ↩
  and not linux-x3b4.site" test.ping
```

By default, Salt has been designed to define a few important grains, but in many cases, it will make sense to add your own definitions (e.g., with information about the computers' locations or roles). SaltStack lends support with its comprehensive documentation [8].

Variety of Modules

Although SaltStack lets you send shell commands directly to minions, in practice most administrators will tend to rely on pre-made modules that cover a multitude of scenarios, from managing users or packages to interacting with databases. All of the modules offer commands tailored for the various platforms and options for processing the output if needed. For this reason, you will want to give modules preference over shell commands, as long as they offer the desired functionality.

A complete list of all available Salt modules can be found online [9]. Please note that different numbers of modules are available depending on the Salt version you use. Only the latest version supports all the modules. If you are looking for more of a challenge, you will find tips in the documentation to write and integrate your own modules [10]. Salt also supports virtual modules. When you deploy these modules, you will not be able to tell the difference, but under the hood, these extensions run other modules that then complete the required tasks.

One such virtual module, for instance, is the `pkg` package module: A look at the official documentation [11] shows that it forwards commands to standard modules, which Salt selects according to the platform on the target system. As soon as you look into using a module, it is a very sensible idea to discover all the extension's functions by reading the official documentation.

The `pkg` module contains numerous important calls, such as letting you install or remove repositories across multiple platforms, manage repositories, and handle pending updates. A full overview of the commands can be found online for systems with the Apt package manager [12].

The calls from this sub-module are, in fact, broadly comparable under the hood to the modules that work with Yum or Zypper. To discover across multiple platforms whether a certain package is installed on a system, simply request the version:

```
$ sudo salt "*" pkg.version grep
deb7:
  2.12-2
linux-x3b4.site:
  2.20-2.4.1
fedora23.fritz.box:
  2.21-7
```

To reinstall a package, run `pkg.install` (Listing 1). You can tell from the output whether the package was already installed, updated, or reinstalled, as the case may be. If a system provides no response at all, the package was already available in its up-to-date version. If you want to update the package lists explicitly (`apt-get update`), simply add the option `refresh=True` before the name of the package.

Processing files on Salt is simple, just like managing packages. Using the `file` module [13], you can not only request metadata on files (e.g., the current Linux privileges, the owner, the file size, and a few other facts), but also delete, append, or manipulate them using regular expressions. One module.function pair that proves to be very useful is `file.append`, with which you can append one or more lines to the file. Using `file.comment` and `file.comment_line`, you can comment out the lines of a file, and `file.rename` lets you can change its name. In the online documentation [13], you will find many examples demonstrating how to manipulate files simultaneously on multiple systems.

Listing 2 shows how to add one entry to the `/etc/motd` file. If you want to build up your knowledge with Salt, one experiment to follow this would be to reverse the change by first commenting out the newly added lines with a file module, then deleting them with another call.

Managing Users

If you manage several systems but do not use a directory service, you still might want to manage users centrally. Unsurprisingly, Salt offers a module, `pw_user`, that lets you add new users, delete user accounts, and modify individual attributes, such as the home directory [14]; you can even rename existing users. Listing 3 shows how you can create a new user with the `user`, rather than the `pw_user`, keyword. Before deleting a specific user, you could first list the users on all the systems.

LISTING 1: Reinstalling a Package

```
$ sudo salt "*" pkg.install refresh=True unzip
linux-x3b4.site:
        ----------
fedora23.fritz.box:
        ----------
deb7:
        ----------
        unzip:
            ----------
            new:
                6.0-8+deb7u5
            old:
                6.0-8+deb7u4
```

LISTING 2: Adding a Line to a File

```
$ sudo salt "*" file.append /etc/motd
  "This system is managed by Salt."
linux-x3b4.site:
  Wrote 1 lines to "/etc/motd"
deb7:
  Wrote 1 lines to "/etc/motd"
fedora23.fritz.box:
  Wrote 1 lines to "/etc/motd"
```

LISTING 3: Creating a New User

```
$ sudo salt "*" user.add intern shell=/bin/bash
deb7:
  True
linux-x3b4.site:
  True
fedora23.fritz.box:
  True
```

LISTING 4: Other SaltStack Commands

```
$ salt-cp "*" "/home/vhoebel/motd" "/etc/motd"
$ sudo salt-run jobs.list_job
$ salt-run jobs.lookup_jid "20140408130655508732" | less
```

Note that Salt offers a module (ssh) to manage the SSH service, so when creating users you can combine modules if you need to generate an SSH key for a new user (ssh-keygen -t rsa) and roll it out to all the existing systems. The module's documentation [15] demonstrates how that works. If you want to do more than just monitor your systems (e.g., collect status information), take a look at the status module [16]. Running status.all_status returns all status data (a lot) for the minion specified. To access the hard disk data in a targeted way, use status.diskusage, and to check the load status, use status.loadavg.

If Salt does not offer a module for a functionality of interest, you can simply run shell commands on the systems by using the cmdmod module [17], which takes the form cmd.<function>. The cmd.run pair takes a command you pass in between double or single quotes (Figure 5) and returns the output. Make sure you run the functions from this module with cmd and not with cmdmod.

More than Modules

In addition to modules, Salt offers a few commands that cover other functionality. The salt-cp command, for example, copies data from the master to the minions (Listing 4, first line), allowing you to maintain /etc/hosts, /etc/motd, or other configuration files centrally. However, this command was only intended as an emergency solution. Generally, you will be able to manipulate existing files in the most targeted way with the file module or use the management module for configuration files.

The salt-run command interacts with Salt; for example, it can be used to create a list of unreachable minions:

```
sudo salt-run manage.down
```

On Salt, every use of a module is dubbed a job, and salt-run lets you generate a list of the last jobs (Listing 4, second line). If you would like to see the output of a certain Salt job again, specify the job ID (Listing 4, last line). For the individual commands, salt-run makes use of its own modules, called "runners" [18].

Conclusions

SaltStack offers almost unlimited system management possibilities because of its large number of integrated modules. Even without the use of in-house configuration management, your computers will be very much under your command, regardless of whether you need to manage users, keep your systems up to date, or just get an overview of particular servers.

Even in heterogeneous networks with Windows machines, you can use Salt. Thanks to its built-in software administration, you can install standard programs like Firefox on Windows clients, and other functions from Linux are also available. Salt proves to be a versatile talent: It makes the use of many tools and DIY scripts redundant. ■■■

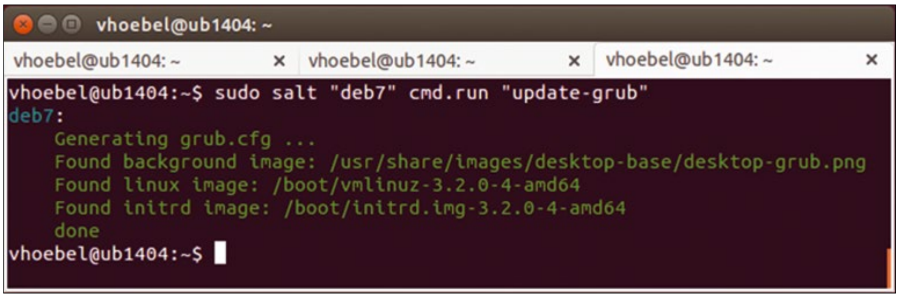

```
vhoebel@ub1404:~$ sudo salt "deb7" cmd.run "update-grub"
deb7:
    Generating grub.cfg ...
    Found background image: /usr/share/images/desktop-base/desktop-grub.png
    Found linux image: /boot/vmlinuz-3.2.0-4-amd64
    Found initrd image: /boot/initrd.img-3.2.0-4-amd64
    done
vhoebel@ub1404:~$
```

Figure 5: In many cases, it makes sense to run shell commands directly.

INFO

[1] SaltStack project: *http://SaltStack.com/community/*

[2] ZeroMQ: *http://zeromq.org/*

[3] Windows minions: *https://docs.SaltStack.com/en/latest/topics/installation/windows.html*

[4] Supported operating systems: *https://docs.SaltStack.com/en/latest/topics/tutorials/salt_bootstrap.html#supported-operating-systems*

[5] Opening ports for Salt: *https://docs.SaltStack.com/en/latest/topics/tutorials/firewall.html*

[6] test module: *https://docs.SaltStack.com/en/latest/ref/modules/all/salt.modules.test.html*

[7] Targeting minions: *https://docs.SaltStack.com/en/latest/topics/targeting/*

[8] Defining grains: *https://docs.SaltStack.com/en/latest/topics/targeting/grains.html*

[9] Built-in Salt modules: *https://docs.SaltStack.com/en/develop/ref/modules/all/index.html*

[10] Writing your own modules: *https://docs.SaltStack.com/en/latest/ref/modules/*

[11] pkg module: *https://docs.SaltStack.com/en/latest/ref/modules/all/salt.modules.pkg.html*

[12] Package module for Apt-based systems: *https://docs.SaltStack.com/en/latest/ref/modules/all/salt.modules.aptpkg.html#module-salt.modules.aptpkg*

[13] file module: *https://docs.SaltStack.com/en/latest/ref/modules/all/salt.modules.file.html*

[14] User management with Salt: *https://docs.SaltStack.com/en/latest/ref/modules/all/salt.modules.pw_user.html*

[15] SSH commands: *https://docs.SaltStack.com/en/latest/ref/modules/all/salt.modules.ssh.html*

[16] status module: *https://docs.SaltStack.com/en/latest/ref/modules/all/salt.modules.status.html*

[17] Shell commands with Salt: *https://docs.SaltStack.com/en/develop/ref/modules/all/salt.modules.cmdmod.html*

[18] Integrated runner modules: *https://docs.SaltStack.com/en/latest/ref/runners/all/index.html*

Your Roadmap to the Open Hardware Revolution …

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

Order your subscription today and tune in to the revolution!

shop.linuxnewmedia.com

Adventures in the Raspberry Pi Wonderland! Your roadmap to OPEN HARDWARE treasure
Don't you touch my Raspberry Pi!
raspberry pi GEEK
Set up a Minecraft server on your Pi
Sun Display Use your Raspberry Pi to view NASA space images!
Overclock your Rasp Pi CPU and improve media center playback
Mux Expand your Arduino's analog capacity with inexpensive multiplexers
NOOBS 1.4.1 Access 7 great Rasp Pi distros from one easy menu!
Special Bonus! Boot your Pi to the Cozy personal cloud!
FREE DVD!
WWW.RASPBERRY-PI-GEEK.COM
Discover the tiny PI ZERO the newest addition to the Raspberry Pi family
This formula will shrink my Raspberry Pi!
raspberry pi GEEK
RASPBERRY PI ZERO Seriously, a computer for $5?
Noise Monitor How loud is your party?
Arduino Drivers Write drivers for unsupported hardware
RasPlex Stream multimedia to multiple TVs
Missile Launcher Control a toy with a wireless keyboard
Stop motion animation with Pi Camera and Scratch
Add a Touchscreen to your RASPBERRY PI
This magnificent experiment depends on my Raspberry Pi
raspberry pi GEEK
BlinkStick LEDs without the GPIO
Sunrise Sunset Pi Zero speeds up time on a model train set
littleBits Pet Feeder Feeding Fifi when you're away
Pi 3 is Here! with Bluetooth and built-in wireless networking
Track Your Trek Logging bike data with a Pi and sensors
Free DVD! NOOBS 1.7.0
Easy menu lets you boot to:
WWW.RASPBERRY-PI-GEEK.COM

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

Digital Sub
Our PDF edition is a convenient option for mobile readers.

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

$59.95 £37.50 €44.90

Download on the App Store

GET IT ON Google play

iTunes Store

Google US

Google UK

Apple, the Apple logo, iPhone, iPad, and iPod touch are trademarks of Apple Inc., registered in the U.S. and other countries. App Store is a service mark of Apple Inc.

Perl script monitors payments and feedback in eBay sales

Watchful Seller

We show how to use a screen-scraper and an application for the official eBay API to trigger an alarm on incoming eBay customer feedback and detect errors in the monthly billing statement. *By Mike Schilli*

If I ever get around to writing my memoirs, I am going to include a lengthy chapter about my life's motto, which is: "Anything you don't constantly monitor is guaranteed to go belly up when you least expect it."

True to this motto, I wrote a script this month to receive an email immediately when one of my eBay customers has left feedback about a transaction.

Simply Scraping

As a quick and dirty solution, and to first avoid having to register as a developer with eBay, I wrote a screen-scraper that extracts the current feedback score from the slew of HTML on the feedback page, before saving the counter and raising the alarm during subsequent runs if the value has increased.

When I click on my username on eBay.com, the browser lands on the feedback page in Figure 1. Selecting `view-source` in the browser shows the HTML code from Figure 2, and a text search for the string `362` (the current

MIKE SCHILLI

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

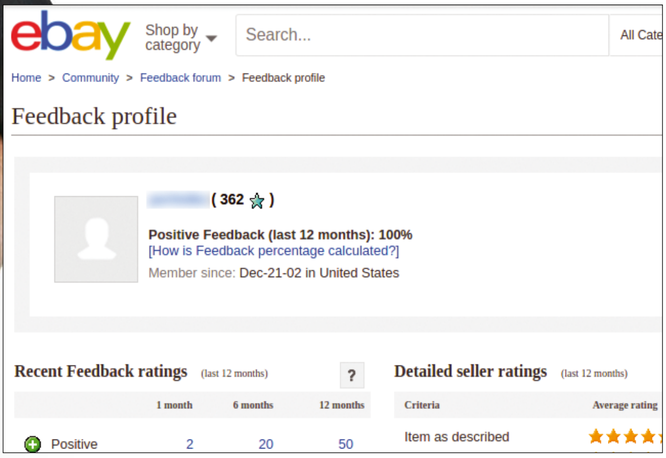

Home > Community > Feedback forum > Feedback profile

Feedback profile

[redacted] (362 ☆)

Positive Feedback (last 12 months): 100%
[How is Feedback percentage calculated?]
Member since: Dec-21-02 in United States

Recent Feedback ratings (last 12 months)

	1 month	6 months	12 months
Positive	2	20	50

Detailed seller ratings (last 12 months)

Criteria	Average rating
Item as described	★★★★

Figure 1: The score for previously received customer comments is on eBay's feedback page – in this case, it is 362.

feedback score) shows that this number is found in a markup tag of the `mbg-l` class:

```
<span class="mbg-l">( 362<img src="...
```

An XPath processor such as HTML::TreeBuilder::XPath can easily retrieve the content of this tag. The query

```
/html/body//span[@class="mbg-l"]
```

locates all `span` elements in the HTML body that have a `class` attribute with a value of `mbg-l`. The double slash in the expression indicates that the requested elements can exist at an arbitrary nesting depth beneath the HTML body tag.

The XPath query then spits out a string such as `( 362) (...`. From there, with a regular expression in Perl, extracting the score is child's play. Listing 1 does exactly this in the `feedback_fetch` function in line 70, and then saves the score that was found in a cache file. Then, on the next run, it compares the value with the one obtained at that point, and fires off an email if the number has increased [1].

Email at Your Command

Writing and reading the cache data, which saves the feedback score confirmed during the last run as its only value, is handled by the `slurp` and `blurt` functions. These are from the CPAN Sysadm::Install module, which exports them in line 3. To retrieve the eBay page, the script uses the slimmed-down LWP::Simple module, whose `get` method simply runs an HTTP request for the specified URL and, if successful, returns the content of the page that was found.

On Linux, it is often very easy to send mail with the `/usr/bin/mail` utility; the CPAN Mail::DWIM module is recommended for those wanting more flexibility. Because I already replaced `/usr/bin/mail` on my Linux box at home with a Perl script that can cope with my ISP's requirements, I left things as is. Figure 3 shows the email that landed in my Gmail inbox.

Before firing it up, the parameters in the script header still need to be tuned to match the local conditions. In line 8, the `$nick` variable receives the name assigned to the eBay user, whose feedback counter will be monitored by `ebay-feedback`. The URL is valid for US eBay accounts, whereas UK accounts run on *ebay.co.uk* instead. Finally, in the `$mail_to` variable, line 23 requires the email address of the user to whom the alerts should be sent.

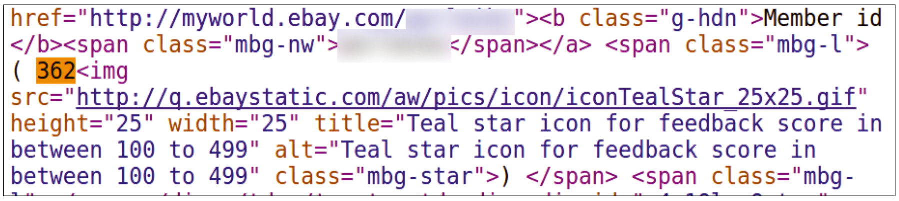

```
href="http://myworld.ebay.com/[redacted]"><b class="g-hdn">Member id
</b><span class="mbg-nw">[redacted]</span></a> <span class="mbg-l">
( 362<img
src="http://q.ebaystatic.com/aw/pics/icon/iconTealStar_25x25.gif"
height="25" width="25" title="Teal star icon for feedback score in
between 100 to 499" alt="Teal star icon for feedback score in
between 100 to 499" class="mbg-star">) </span> <span class="mbg-
```

Figure 2: The feedback score is found in an HTML element of the `mbg-l` class: here, `362`.

LISTING 1: ebay-feedback

```
01 #!/usr/local/bin/perl -w
02 use strict;
03 use Sysadm::Install qw(:all);
04 use LWP::Simple ;
05 use HTML::TreeBuilder::XPath;
06 use Log::Log4perl qw(:easy);
07
08 my $nick       = "my-ebay-name";
09 my $ebay_url   =
10   "http://feedback.ebay.com" .
11   "/ws/eBayISAPI.dll?ViewFeedback2" .
12   "&userid=$nick";
13
14 my( $home )    = glob "~";
15 my $data_dir   = "$home/logs";
16 my $cache      =
17   "$data_dir/ebay-feedback.cache";
18 my $log_file   =
19   "$data_dir/ebay-feedback.log";
20
21   # mail prefs
22 my $mailer     = "/usr/bin/mail";
23 my $mail_to    = 'my@email.com';
24
25 mkd $data_dir if !-d $data_dir;
26
27 Log::Log4perl->easy_init( {
28   level => $DEBUG,
29   file => ">>$log_file" } );
30
31 my $last_feedback;
32
33 if( -f $cache ) {
34   $last_feedback = slurp $cache;
35 }
36
37 my $feedback = feedback_fetch();
38
39 if( !defined $last_feedback or
40   $last_feedback != $feedback ) {
41
42   $last_feedback ||= 0;
43
44   INFO "New feedback: $feedback";
45   INFO "Sending mail to $mail_to";
46
47   open PIPE,
48     "| $mailer -s 'New Ebay Feedback: " .
49     "$feedback' $mail_to";
50
51   print PIPE <<EOT;
52 Ebay feedback changed: It's $feedback now
53 and was $last_feedback yesterday:
54
55     $ebay_url
56
```

Monitoring the Monitor

How can you monitor whether the script is still functioning, or whether eBay has changed its page layout? Listing 1 records all the operations for that purpose via Log::Log4perl in the log file `~/logs/ebay-feedback.log`. Additionally, for instance, a daily Nagios script could trawl through the data that show up there based on a pattern such as "OK" under the current date, and sound the alarm if it is not found."

eBay Can't Add

You might not realize that some people out there actually add up the individual line items of a computer-generated invoice and check whether the bottom line is right. But, I must confess, sometimes I cannot resist – especially for the monthly eBay invoice when just four items are listed, as in Figure 4.

Every second grader with a calculator could verify in a breeze that 4.40 + 0.45 + 1.00 + 0.39 adds up to the total sum of $6.24. Yet, eBay's Java jockeys are apparently not aware of the article on floating-point arithmetic that every programmer should read [2]. Otherwise, they'd realized that CPUs save floating-point numbers inaccurately; for that reason, astonishing rounding errors can arise during addition.

That explains why eBay brazenly charged me $6.25, although the total sum for the line items should, without a doubt, have come to $6.24. I had already called into the customer service center a while ago on a different issue and spent a half hour talking to a representative from overseas to report a discrepancy of a few cents. I will save this one-cent complaint for a very special day.

Auditor

The script in Listing 2 helps to automatically conduct this audit every month. It relies on the official eBay API, identifying itself with a token for the eBay web service and retrieving the previous month's invoice as an XML document. The line items are found there under the `NetDetailAmount` tag, as is the total amount under `InvoiceBalance`. The pipelined `invoice-check` in Listing 3 is

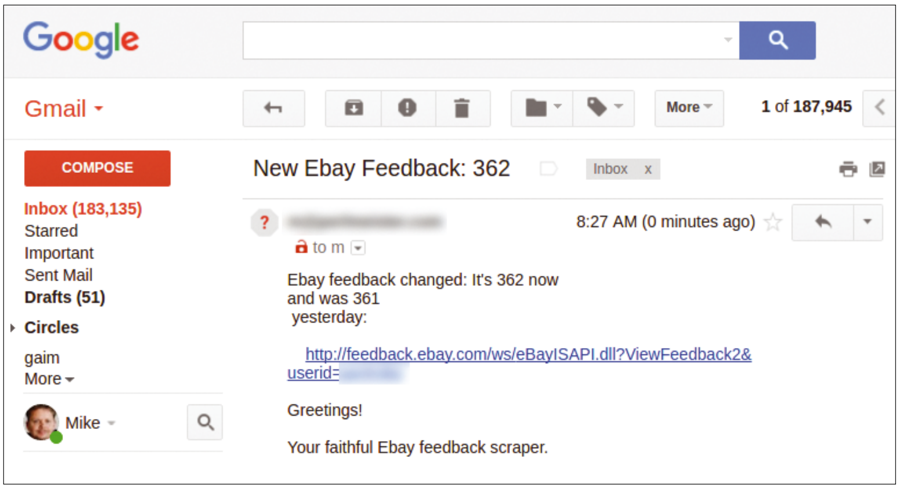

New Ebay Feedback: 362

[redacted]

8:27 AM (0 minutes ago)

to m

Ebay feedback changed: It's 362 now
and was 361
yesterday:

http://feedback.ebay.com/ws/eBayISAPI.dll?ViewFeedback2&
userid=[redacted]

Greetings!

Your faithful Ebay feedback scraper.

Figure 3: A message informs the user about the newly arrived feedback from a buyer.

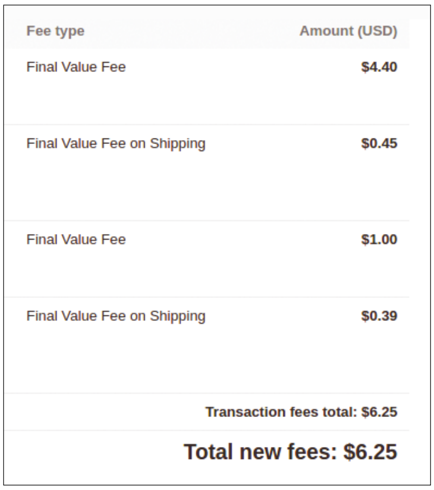

Fee type	Amount (USD)
Final Value Fee	$4.40
Final Value Fee on Shipping	$0.45
Final Value Fee	$1.00
Final Value Fee on Shipping	$0.39
	Transaction fees total: $6.25
	Total new fees: $6.25

Figure 4: On eBay, 4.40 + 0.45 + 1.00 + 0.39 does not come to $6.24, but $6.25.

LISTING 1: ebay-feedback (continued)

```
Greetings!

Your faithful Ebay feedback scraper.
EOT
  close PIPE;

  blurt $feedback, $cache;

} else {
  INFO "Feedback unchanged ($feedback).";
}

###########################################
sub feedback_fetch {
###########################################
  INFO "Fetching $ebay_url";

  my $content = get $ebay_url;

  if( !defined $content ) {
    ERROR "Fetching $ebay_url failed";
    return undef;
  }

  my $tree= HTML::TreeBuilder::XPath->new;
  $tree->parse( $content );

  my( $text ) = $tree->findvalue(
      '/html/body//span[@class="mbg-l"]');

  if( $text =~ /\s*(\d+)/ ) {
    return $1;
  }

  ERROR "Pattern in page not found";
  return undef;
}
```

then able to check whether eBay calculated the sum of the line items correctly (Figure 5).

Key to the Kingdom of Data

An application that reads or writes eBay user data must identify itself to the platform with a token. To accomplish this, the developer must register with the eBay Developers Program and enter an application name. In the normal flow for a smartphone app, for example, the user would be guided through the login process on eBay, and, upon confirmation, receive a token that the app can use going forward. If developers only need a token for test purposes and for their own account, they can instead click the `Sign in to Production` button on the eBay Developer Program sign-up page and receive a token this way (Figure 6) [3].

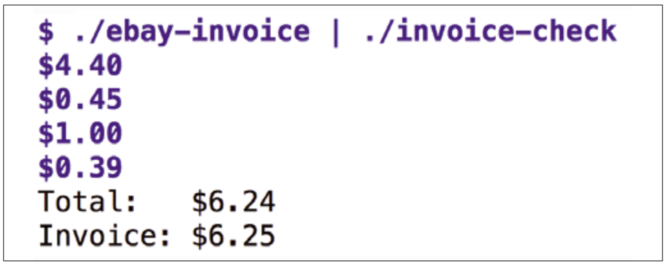

```
$ ./ebay-invoice | ./invoice-check
$4.40
$0.45
$1.00
$0.39
Total:   $6.24
Invoice: $6.25
```

Figure 5: Two scripts fetch the last eBay invoice and check its veracity.

LISTING 2: ebay-invoice

```
#!/usr/local/bin/perl -w
use strict;
use LWP::UserAgent;
use DateTime;
use Path::Tiny;

my $dt_today = DateTime->today;
my $dt = DateTime->new(
  year  => $dt_today->year,
  month => $dt_today->month,
  day   => 1,
);
my $invoice_date =
  $dt->subtract( days => 1 )->ymd;

my $ua = LWP::UserAgent->new;

$ua->default_header(
  "X-EBAY-API-CALL-NAME", "GetAccount" );
$ua->default_header(
  "X-EBAY-API-COMPATIBILITY-LEVEL", 863 );
$ua->default_header(
  "Content-Type", "text/xml" );
$ua->default_header(
  "X-EBAY-API-SITEID", "0" );

my $token = path( "token" )->slurp;
chomp $token;

my $body  = <<EOT;
<?xml version="1.0" encoding="utf-8"?>
<GetAccountRequest
 xmlns="urn:ebay:apis:eBLBaseComponents">
<RequesterCredentials>
    <eBayAuthToken>$token</eBayAuthToken>
</RequesterCredentials>
<AccountHistorySelection>SpecifiedInvoice
</AccountHistorySelection>
<InvoiceDate>$invoice_date</InvoiceDate>
</GetAccountRequest>
EOT

my $resp = $ua->post(
  "https://api.ebay.com/ws/api.dll",
  Content => $body );

if( $resp->is_error ) {
    die $ua->message;
}

print $resp->decoded_content;
```

LISTING 3: invoice-check

```
#!/usr/local/bin/perl -w
use strict;
use XML::Simple;
use Math::Currency;

my $xml   = join "", <>;
my $ref   = XMLin( \$xml );
my $total = 0;
my @items = ();

for my $entry ( @{ $ref->{
 "AccountEntries" }->{ AccountEntry } } ) {

    # ignore payments
  next if $entry->{ ItemID } == 0;

  my $amount = $entry->{
    "NetDetailAmount" }->{ content };

    # ignore free items
  next if $amount == 0;

  my $mc = Math::Currency->new( $amount );
  print "$mc\n";
  $total += $mc;
}

my $invoice_amount =
  $ref->{ AccountSummary }->{
    "InvoiceBalance" }->{ content };

print "Total:   $total\n";
print "Invoice: \$$invoice_amount\n";
```

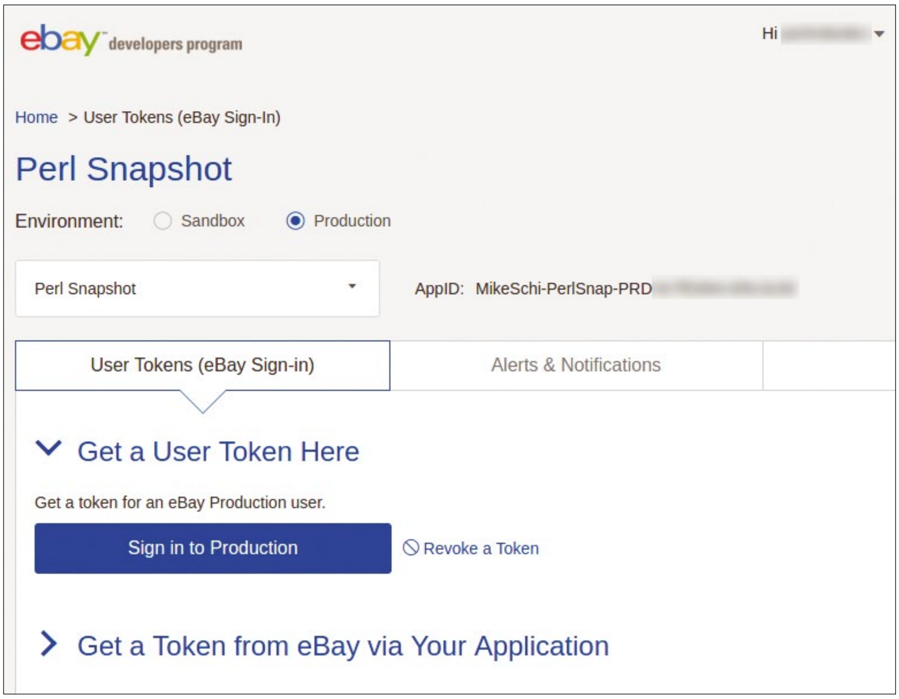

ebay developers program

Hi [redacted] ▾

Home > User Tokens (eBay Sign-In)

Perl Snapshot

Environment: ○ Sandbox ● Production

Perl Snapshot ▾

AppID: MikeSchi-PerlSnap-PRD[redacted]

User Tokens (eBay Sign-in) | Alerts & Notifications

Get a User Token Here

Get a token for an eBay Production user.

Get a Token from eBay via Your Application

Figure 6: Developers can get a token for their own eBay account for test purposes.

This token is then valid for the account actually in live production, as opposed to the sandbox in which developers can test out their apps until they are sure that they are ready for production (Figure 7).

Requests for personal data then require the token to be included, embedded in the `RequesterCredentials` tag in the request's XML (see Listing 2, line 35). It consists of an 872-character hex string that must be placed between the `eBayAuthToken` tags without line breaks.

Saint Bureaucratius

The documentation on the eBay Developers Program pages [3] is organized quite chaotically; the individual APIs overlap for historical reasons. The API Reference page for the `GetAccount` request [4] shows which parameters it needs to retrieve an invoice created on a specific date.

As Listing 2 shows, the post request sent over the eBay API will then still require a few HTTP header values alongside the XML body. The `X-EBAY-API-SITEID` header value determines which eBay department and country the request is sent to, and the `0` value is valid for eBay.com in the United States, whereas applications can reach eBay.co.uk with the number 3.

The `X-EBAY-API-COMPATIBILITY-LEVEL` header specifies the lowest version of the web API with which the client is able to work. Currently, version 959 is the most up to date; the value entered in the script (863) is for academic purposes only. To get the output in XML, line 22 sets the `Content-Type` header to `text/xml`. The fact that the application runs the `GetAccount` server method is already stated in the XML body of the post request, but Saint Bureaucratius demands that it appear again in the `X-EBAY-API-CALL-NAME` header in line 19.

eBay always generates bills at the end of the month. To jump from the current day to the last day of the past month (i.e., the last billing date), Listing 2 sets the value for the current day to 1 (i.e., the first of the current month); then it subtracts one day from there. It therefore lands on the last day of the past month.

The `DateTime` function `ymd` converts the date for the March 2016 invoice, for example, into the `20160331` format, which the eBay API recognizes if the `SpecifiedInvoice` value is entered in the request under `AccountHistorySelection`. The date of this invoice must then be added with `InvoiceDate`.

The script looks for the sensitive auth token in a file named `token`, from which it reads it in line 27, removing a potential newline character at the end of the file in line 28 and then integrating the key into the XML blob via the `$token` variable as of line 35. The date of the invoice is found in the `$invoice_date` variable and is also added to the XML; line 38 interpolates the value.

If an error occurs during the retrieval of data, line 47 aborts the script and prints the error report from eBay. Typical problems include the portal reacting badly to unwanted spaces and line breaks in the XML data, although it does fortunately describe them in detail in its error reports. If all goes smoothly, line 50 prints the returning XML to the standard output, where the next stage of the processing pipeline accesses the data.

Adding Up (Correctly)

The script in Listing 3 grabs the XML blob from Listing 2 and searches inside it for the line items (Figure 8) and the total, relying on the XML::Simple CPAN module to do so. It then verifies whether the addition of the line items produces the total sum listed. The beauty of XML::Simple is that it converts the entire XML into a huge Perl data structure in which a script can labor away at its convenience. Tags are converted into hashes, and it creates Perl arrays from lists of single entries. For text that occurs between two XML tags as their content, XML::Simple creates a hash entry named `content`. Using this name, for instance, in line 30 in Listing 3, the script

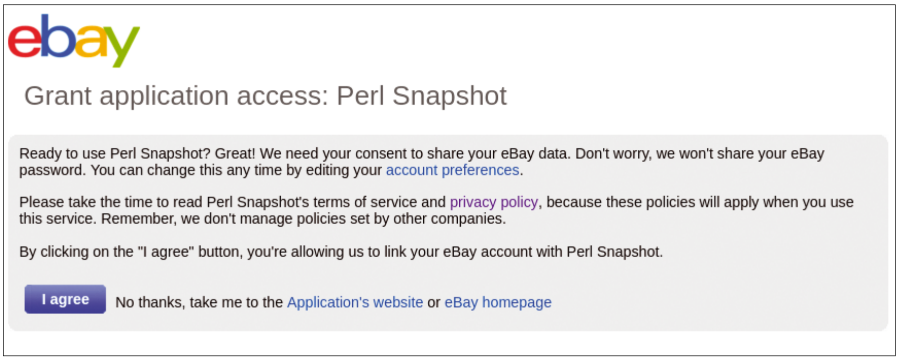

ebay

Grant application access: Perl Snapshot

Ready to use Perl Snapshot? Great! We need your consent to share your eBay data. Don't worry, we won't share your eBay password. You can change this any time by editing your account preferences.

Please take the time to read Perl Snapshot's terms of service and privacy policy, because these policies will apply when you use this service. Remember, we don't manage policies set by other companies.

By clicking on the "I agree" button, you're allowing us to link your eBay account with Perl Snapshot.

No thanks, take me to the Application's website or eBay homepage

Figure 7: After login, eBay asks whether the developer really trusts their app.

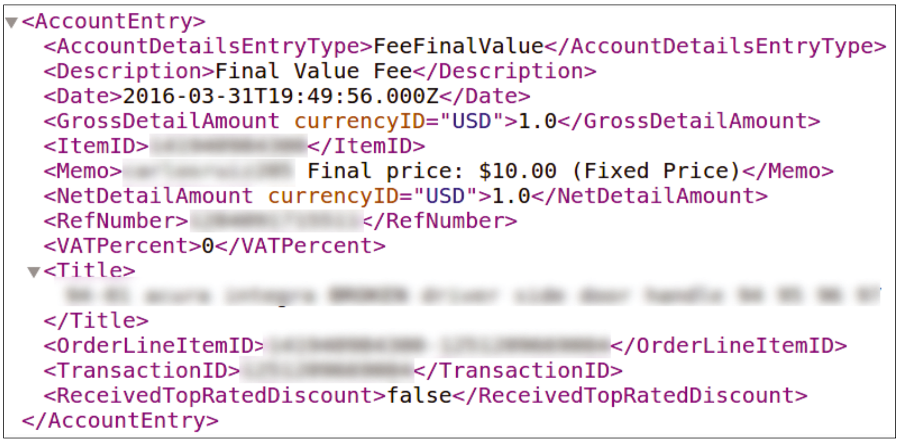

```
▼<AccountEntry>
   <AccountDetailsEntryType>FeeFinalValue</AccountDetailsEntryType>
   <Description>Final Value Fee</Description>
   <Date>2016-03-31T19:49:56.000Z</Date>
   <GrossDetailAmount currencyID="USD">1.0</GrossDetailAmount>
   <ItemID>[redacted]</ItemID>
   <Memo>[redacted] Final price: $10.00 (Fixed Price)</Memo>
   <NetDetailAmount currencyID="USD">1.0</NetDetailAmount>
   <RefNumber>[redacted]</RefNumber>
   <VATPercent>0</VATPercent>
  ▼<Title>
     [redacted]
   </Title>
   <OrderLineItemID>[redacted]</OrderLineItemID>
   <TransactionID>[redacted]</TransactionID>
   <ReceivedTopRatedDiscount>false</ReceivedTopRatedDiscount>
 </AccountEntry>
```

Figure 8: The invoice in XML format lets you add single entries to verify whether the total is correct.

extracts the invoice's total amount from the `InvoiceBalance` tag.

As I already said, you are strongly advised not to simply add up dollar (or Euro) and cent amounts as floating-point values [2]. Instead, careful programmers should calculate everything in cents and process these as integers to subsequently separate the last two numbers from the result and account for them in cents.

The Math::Currency CPAN module does this correctly and formats the result attractively, even if the input data isn't provided correctly, as in Ebay's one-digit decimal-point format (e.g., `$1.6`). However, the course of time has taken its toll on the CPAN module; its last release is six years old, and the test suite it includes rattles through the test with creaks and groans and fails in the end. The module does work, however, so I simply forced the CPAN shell to install it by typing:

```
cpanm --force Math::Currency
```

The XML invoice contains not only the line items and the running total, but also the amounts already transferred to eBay by the customer from previous bills. Line 15 acknowledges by looking for `ItemIDs` with a value of `0`, screening out all fees unrelated to individual item sales.

Monitoring the Checkout Assistant

The error in the eBay invoice brought me around to the idea of also monitoring the line items on my receipt when I next visit the supermarket. Who would notice if the total shown there was a few cents off from the sum total of the line items? No one, except me, of course! I've already made a plan to monitor them with a script that gobbles up the OCR data from the receipt and runs the test. Potentially something you can look forward to reading in a future column. ■■■

INFO

[1] Listings for this article: *ftp://ftp.linux-magazine.com/pub/listings/magazine/188/Perl*

[2] "What Every Computer Scientist Should Know About Floating-Point Arithmetic": *http://docs.oracle.com/cd/E19957-01/806-3568/ncg_goldberg.html*

[3] eBay Developers Program application keys: *https://developer.ebay.com/my/keys*

[4] API documentation for `GetAccount` to retrieve an account's monthly eBay invoice: *http://developer.ebay.com/Devzone/xml/docs/Reference/ebay/GetAccount.html*

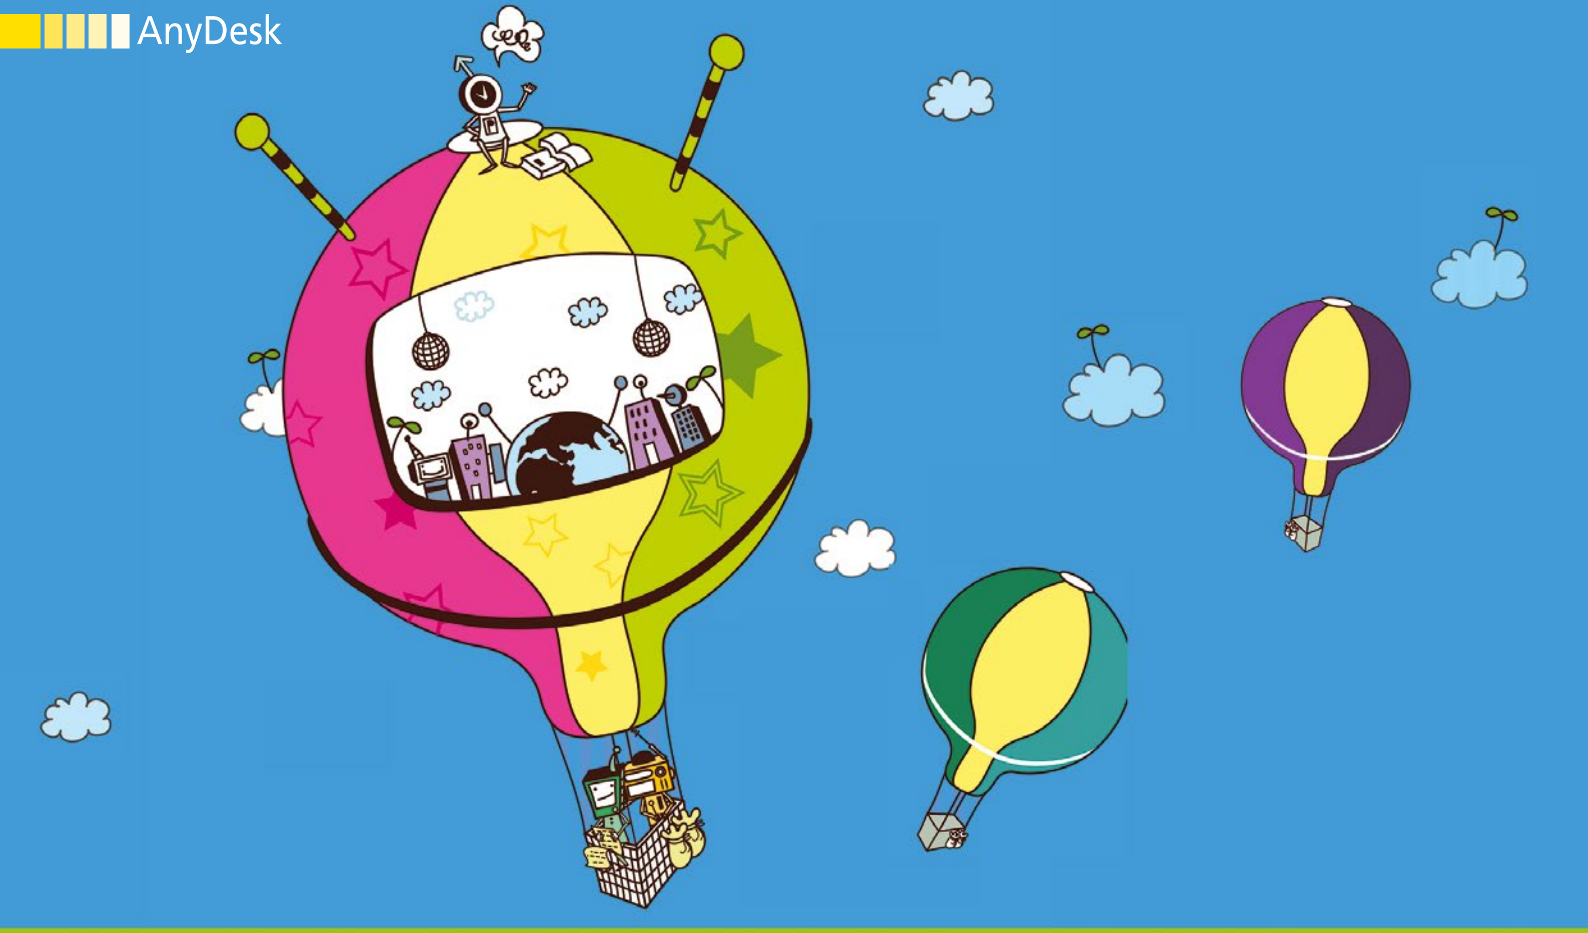

AnyDesk in competition with TeamViewer

Working Remotely

Remote control software is frequently used to work on remote computers and for group work. AnyDesk wants to gain a foothold in this niche with an innovative technical concept. *By Ferdinand Thommes*

Many experienced PC users help family and friends manage and maintain their computers. However, if their family and friends live far away, they usually use remote access programs (also known as remote desktop programs). The software displays the remote computer's desktop on the local screen (in simplified form) so that the helper can work on the remote computer as if sitting in front of it. Typical functions of this type of software include screen sharing, remote maintenance, and file transfer; they also often allow audio and video chats.

Competition

The proprietary software TeamViewer [1] is the top dog among remote desktop applications. The company has developed a comfortable position in the market in recent years and provides its software free of charge for private use. Around two years ago, competition arose from within the company's own ranks: Three former employees put the proprietary tool AnyDesk [2] on the market. The developers wholeheartedly claim that AnyDesk is the fastest remote desktop application in the world and that it provides new dimensions for working on remote computers. The company behind the software, philandro, sees the software as the first of a series of products that aim to protect privacy and put cloud services back into private hands.

After the end of the one-year beta phase for the Windows version in the summer of 2015, a beta version was re-

AUTHOR

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

leased in November for Linux and BSD derivatives. Despite being a beta version, the Linux version now has the same version number 2.1.1 as the Windows release. AnyDesk provides its software as a tarball with sources, as well as in the form of packages for Debian and its offshoots, various Fedora versions, Red Hat Enterprise Linux (RHEL), Mageia, openSUSE, and SUSE Linux Enterprise Server (SLES). Versions for other platforms, including Mac OS X, iOS, and Android, are already under development, although the developers have not yet named a release date.

Comparative Values

I tested the latest version for Debian and Windows in conjunction with various Linux guests and Windows 7 and 10. I didn't bother with benchmarks between the candidates: They aren't actually necessary for determining whether the candidate works faster and more stutter-free than its competitors or displays a clearer picture.

The manufacturer provides its own benchmark results in a PDF document that is available for download [3], which you should, however, treat with caution for obvious reasons. TeamViewer, RDP, Google Remote Desktop, Screenhero, Splashtop, and AnyDesk all compete with each other in this space.

Optimized Image Processing

The higher performance of AnyDesk in several areas is mainly apparent in two key areas. AnyDesk with DeskRT [4] uses a video codec specifically developed for transmitting graphical desktop interfaces for all platforms. It transmits screen content in a special compression process with up to 60 images per second, although the software is limited to the modified screen captures. AnyDesk also keeps up to 100 screen images in a buffer, which can be used when needed rather than being retransmitted.

DeskRT has been optimized further to compress large areas of color, high contrasts, sharp edges, repetitions of pixels, and moving image content, as are often found in graphical user interfaces. Additionally, the deep integration in the respective operating system, the optimization of multiprocess architecture, and AnyDesk's zero-copy design aim to ensure that the image data is transferred to the screen in as few processing steps as possible.

The server infrastructure, which allows the networking of AnyDesk participants outside the LAN, has also been optimized for using failsafe telecommunications applications. Based on the Erlang programming language, it provides a run-time system and a substantial set of libraries (OTP) for constructing high-availability systems.

A second point only concerns the Linux version. The AnyDesk developers created a native Linux application based on GTK, which is available for most distributions as both a 32-bit and 64-bit version. TeamViewer, on the other hand, relies on Wine as an intermediate layer – providing only half-hearted support for Linux and resulting in some disadvantages. However, perhaps spurred by the competition, the manufacturer has promised an increase in efficiency of up to 30 percent and file transfers that are 15 times faster in the new TeamViewer 11.

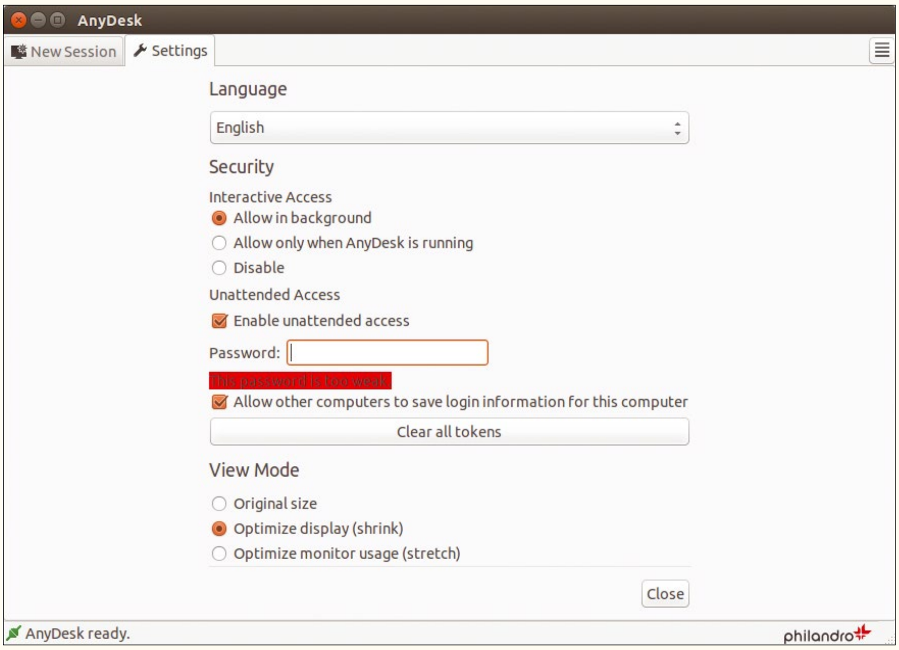

Language

English

Security

Interactive Access

● Allow in background

○ Allow only when AnyDesk is running

○ Disable

Unattended Access

☑ Enable unattended access

Password:

This password is too weak

☑ Allow other computers to save login information for this computer

View Mode

○ Original size

● Optimize display (shrink)

○ Optimize monitor usage (stretch)

AnyDesk ready.

Figure 1: The settings allow silent remote access to the computer – in theory. This often fails in practice, because the software doesn't always note the access password.

The Right Settings

Working with AnyDesk is pretty easy. In Windows, all you need to do is run the 1.7MB EXE file. At the end of the session, the software will ask whether you want to install the program. If you copy it onto a USB flash drive, the software

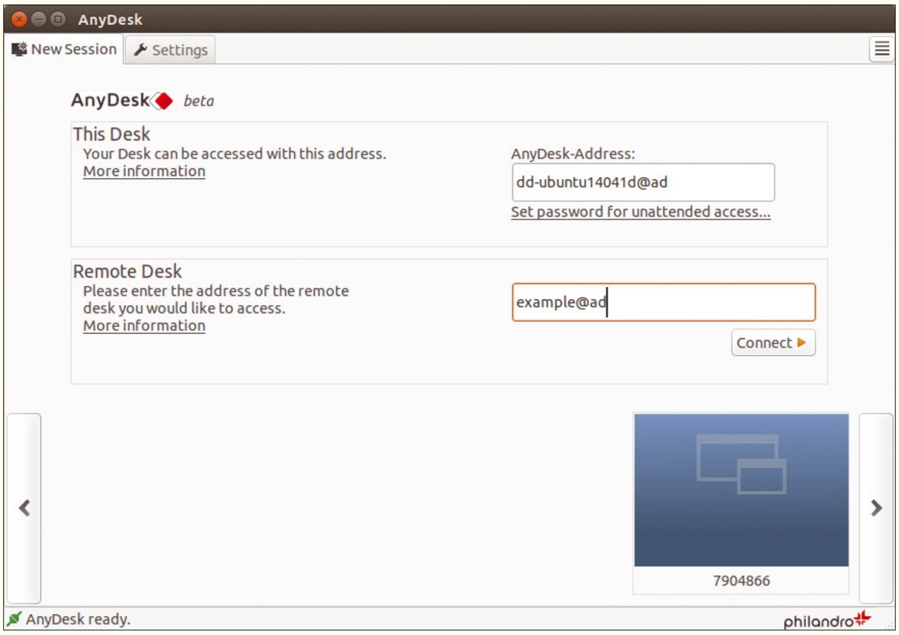

This Desk
Your Desk can be accessed with this address.
More information

AnyDesk-Address: dd-ubuntu14041d@ad
Set password for unattended access...

Remote Desk
Please enter the address of the remote desk you would like to access.
More information

example@ad

7904866

AnyDesk ready.

Figure 2: AnyDesk's main window is essentially limited to input on the computer to which you want to connect.

can be used anywhere. However, according to the developers, such a stand-alone version will not be available for Linux in the foreseeable future.

After starting up the interface, open the *Settings* (Figure 1), which are hidden in the top-right corner behind the horizontal lines. Here you can activate the password-protected silent mode, which allows you to access an unattended PC remotely – the only requirement is that AnyDesk is running on both machines.

The software also lets you log in to the computer without entering a password. This is an option when first starting in silent mode when entering the password.

The Interface

AnyDesk's main window appears self-explanatory at first glance and is divided into two areas: *This Desk* and *Remote Desk* (Figure 2). The former is usually the host name of the computer at which you are sitting.

To establish a connection with another computer on the LAN, you need to enter its hostname or IP address in the second input field. This creates a direct session via the 7070/TCP port without the need to go through an external AnyDesk server. On the other hand, AnyNet (AnyDesk's network) handles sessions outside the local network. To do so, the software on the other side generates a unique key when started, and you can enter this key in the input field next to *Remote Desk*.

Once the connection is established, a window opens on the computer to be administered that indicates the connection request. The other side can then also specify the permissions (Figure 3). The functions that are always available are controlling the remote computer, using the clipboard, and file sharing. A small symbol next to the settings icon also makes it possible to exchange text messages (Figure 4).

Mainly Bright

AnyDesk left a mixed impression during testing with the beta version for Linux. The good news: In terms of security, the tool uses TSL 1.2, better known as SSL, in the latest version. The program cryptographically verifies all connection participants. Both the speed and the quality of the display were impressive with limited strain on bandwidth. Basic access to computers worked reliably in the test, whether on remote PCs or on the local network.

However, several errors currently affect working with AnyDesk. For example, I recorded several disconnections in the Linux network, after which I had to terminate the AnyDesk process using the `kill` command before I could restore the connection.

It wasn't possible to use silent mode on the LAN, either – apparently because AnyDesk hadn't saved the stored password. These errors occurred more frequently in

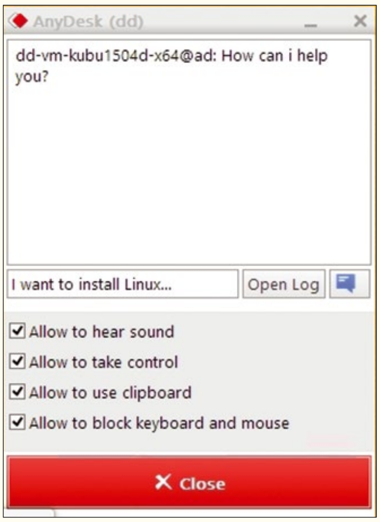

dd-vm-kubu1504d-x64@ad: How can i help you?

I want to install Linux...

☑ Allow to hear sound
☑ Allow to take control
☑ Allow to use clipboard
☑ Allow to block keyboard and mouse

Figure 3: AnyDesk requires explicit approval of the connection setup on the remote computer. This is where the permissions granted are specified.

environments based on the Qt framework; GTK environments had significantly fewer errors.

Conclusions

AnyDesk is certainly currently suitable for brief repair sessions. Unlike TeamViewer, noticeable input delays rarely occur. One particular highlight is the option to operate AnyDesk via terminal [5]. However, the errors mentioned above currently make group work or work over longer periods difficult on another computer in the local network – at least in the Linux version.

AnyDesk therefore has a technically good foundation. However, the developers still need to make some improvements in the Linux version to make it serious competition for TeamViewer. These improvements might take some time because of the complexity of the application, but I'll definitely take another look at AnyDesk later in the year. ■■■

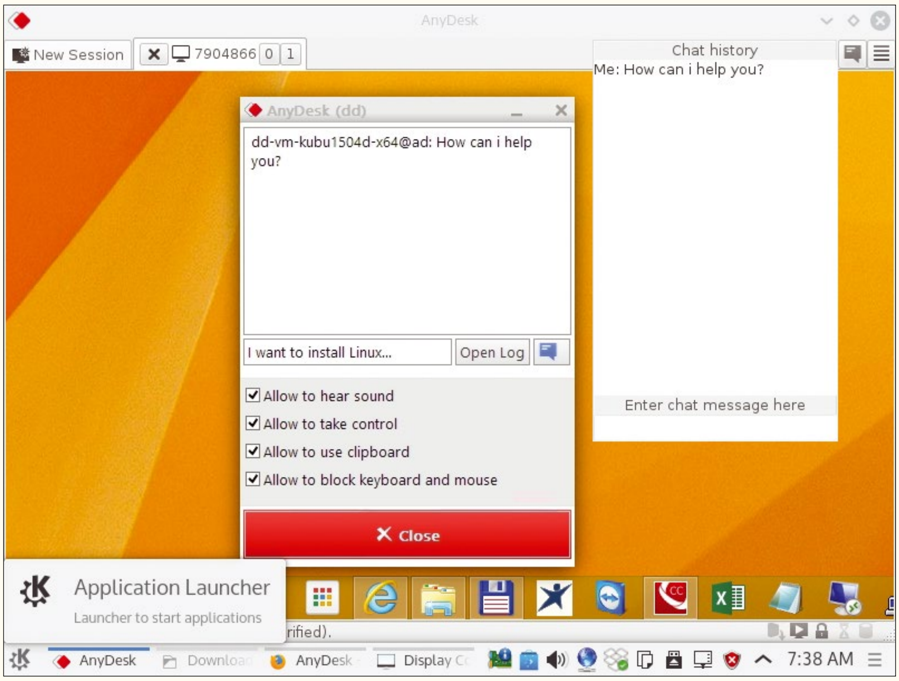

Me: How can i help you?

dd-vm-kubu1504d-x64@ad: How can i help you?

I want to install Linux...

☑ Allow to hear sound
☑ Allow to take control
☑ Allow to use clipboard
☑ Allow to block keyboard and mouse

Figure 4: The integrated chat function makes it possible to exchange messages directly.

INFO

[1] TeamViewer: *https://www.teamviewer.com*

[2] AnyDesk: *http://anydesk.com/remote-desktop*

[3] AnyDesk benchmark: *http://anydesk.com/benchmark/anydesk-benchmark.pdf*

[4] DeskRT video codec: *http://anydesk.com/technology*

[5] Starting Anydesk via terminal: *http://support.anydesk.com/knowledgebase/articles/441867-command-line-interface*

DVD
Linux Mint 17.3 + Kubuntu 15.10
Mint + Kubuntu
UBUNTU user
EXPLORING THE WORLD OF UBUNTU
DISCOVER THE WORLDS BEYOND UBUNTU
ALTERNATIVES
4 Mini distros you have to try
Build your own distro with PCLinux OS
Bodhi gets a new desktop: Moksha
Power Squared: MX Linux
SyncThing
Sync all the things between all your devices

File synchronization with Osync and Freehold

Sync and Host

Sync files and host them on the web with a minimum of effort using Osync and Freehold. By Dmitri Popov

File synchronization and hosting applications come in all shapes and sizes, but if simplicity and speed are your primary requirements, then Osync and Freehold are right up your alley. Both tools are light on resources, easy to deploy, and straightforward in use.

File Syncing from the Command Line with Osync

Rsync is probably the most popular command-line backup tool on the Linux platform. Even though mastering rsync basics doesn't require a lot of effort, things can quickly become somewhat tricky when you move beyond simple backup commands and scripts. Enter Osync [1], a Bash shell script that acts as a user-friendly wrapper for rsync and lets you make the most of rsync's functionality.

Osync has no other dependencies besides rsync, and the latter is available in the software repositories of all mainstream Linux distributions. To install rsync on Ubuntu and its derivatives, run the `sudo apt-get install rsync` command. Note that Osync is designed for Bash, and the script might not work with other shells like Zsh or Fish. Osync comes with an installer script that simplifies the process of deploying the tool on a Linux system. Clone the project's GitHub repository and then run the supplied installer script (the last command must be run as root):

```
git clone https://github.com/deajan/osync
cd osync/
./install.sh
```

The installer copies a handful of scripts, including `osync.sh`, to the `/usr/local/bin` directory and creates the `osync-srv` script for starting Osync on boot. To check whether Osync has been installed properly, run the `osync.sh` command, which should return the current version number and usage info (Figure 1).

Unlike rsync, which supports only one-way synchronization, Osync can sync the contents of two directories in both directions. This means that all changes (i.e., modifications, additions, and deletions) in Directory A are mirrored in Directory B, and vice versa. Better still, Osync handles file deletions in a rather intelligent manner: When a file is deleted in one directory, Osync moves its copy in the other directory to the special `.osync_workdir/deleted` folder instead of simply deleting it. This ensures that you can always retrieve the deleted files, if needed. To purge the deleted files after a specified period of time, you can use a dedicated Osync option (more about it later).

To perform synchronization, Osync requires two parameters: `--initiator` (path to the source directory) and `--target` (path to the target directory):

```
osync.sh --initiator=/path/to/source ⮐
  --target=/path/to/target
```

Like any other command-line tool, Osync supports several useful options.

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 *scribblesandsnaps.wordpress.com*.

The `--dry` option, for example, can be used to test synchronization commands before actually running them, and `--verbose` generates a detailed output during synchronization. The `--on-change` option performs synchronization automatically when the contents of the source directory are changed. And, if you set up a cron job to run Osync at a specified schedule, you should use the `--silent` option to suppress all output.

In addition to the command-line options, you can also specify Osync-specific parameters. For example, Osync supports the `CREATE_DIRS` parameter. When enabled, it instructs Osync to create the target directory if it doesn't exist. Unlike the regular options, Osync-specific parameters must be specified before the Osync command:

```
CREATE_DIRS=yes osync.sh --initiator=/path/to/source --target=/path/to/target
```

Because Osync is powered by rsync, the tool supports remote synchronization via SSH. To make use of this feature, you need to specify the SSH URL as part of the target directory path:

```
osync.sh --initiator=/path/to/source --target=ssh://user@remotehost.com:port//↩
  path/to/target
```

Specifying options and parameters on the fly is fine for occasional synchronization, but if you plan to use Osync on a regular basis, it makes sense to create a configuration file and point Osync to it. You also can use the supplied `sync.conf` file as a template. At the very minimum, you need to specify the correct paths for the `INITIATOR_SYNC_DIR` and `TARGET_SYNC_DIR` options.

If the target directory is on a remote host, you also need to add the path to your private SSH key to the `SSH_RSA_PRIVATE_KEY` option. You can tweak several other useful options to your liking. For example, if you want to exclude certain files from being synced, you can create a file containing exclusion rules and add its path to the `RSYNC_EXCLUDE_FROM` option. Osync comes with an `exclude.list.example` example file you can use as a starting point. Keep in mind that the exclusion file must be in the same directory as the configuration file.

Osync also supports several options that control backup and deletion. Although there is probably no need to tweak them, understanding what they do can give you a better idea how Osync handles synchronization conflicts and deletions. When the `CONFLICT_BACKUP` option is enabled, Osync saves backup copies of files that have been modified in the source directory in the special `.osync_workdir/backups` folder inside the target directory.

Enabling the `CONFLICT_BACKUP_MULTIPLE` option effectively adds no-frills versioning capabilities to Osync; whenever you edit a file in one directory and perform synchronization, Osync creates a backup version of the file in the other directory. With this feature activated, you can quickly end up with a large number of backup files. Fortunately, Osync has the dedicated `CONFLICT_BACKUP_DAYS` option that purges the backup copies after a specified number of days.

Osync can also send email alerts, which can come in rather handy when you run Osync unattended. To enable this feature, configure the required options (email addresses and SMTP connection info) in the `ALERT OPTIONS` section of the configuration file.

Once you've configured all the options, save the configuration file under the `osync.conf` name in the `/etc/osync`

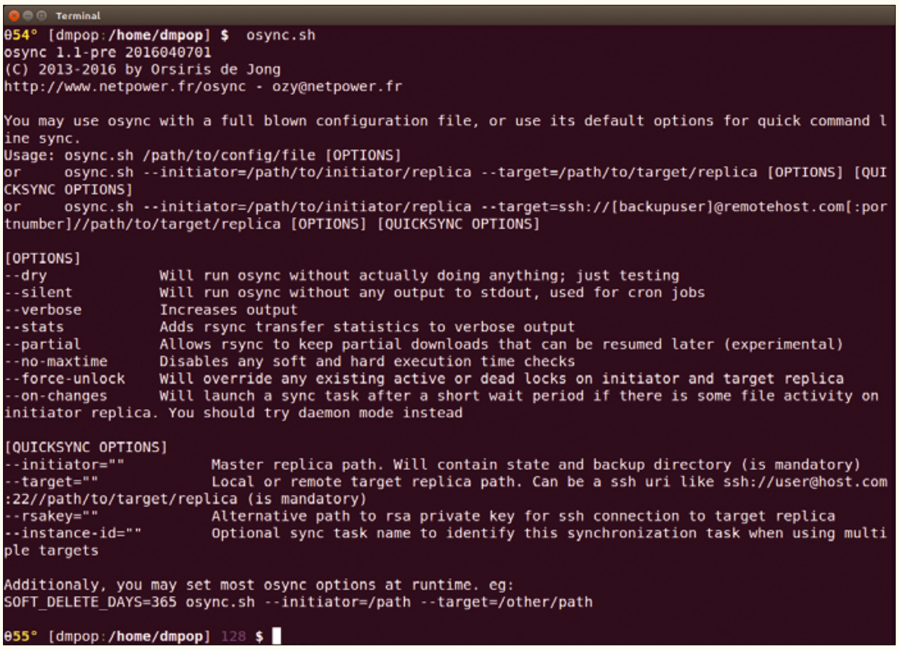

```
054° [dmpop:/home/dmpop] $  osync.sh
osync 1.1-pre 2016040701
(C) 2013-2016 by Orsiris de Jong
http://www.netpower.fr/osync - ozy@netpower.fr

You may use osync with a full blown configuration file, or use its default options for quick command l
ine sync.
Usage: osync.sh /path/to/config/file [OPTIONS]
or     osync.sh --initiator=/path/to/initiator/replica --target=/path/to/target/replica [OPTIONS] [QUI
CKSYNC OPTIONS]
or     osync.sh --initiator=/path/to/initiator/replica --target=ssh://[backupuser]@remotehost.com[:por
tnumber]//path/to/target/replica [OPTIONS] [QUICKSYNC OPTIONS]

[OPTIONS]
--dry              Will run osync without actually doing anything; just testing
--silent           Will run osync without any output to stdout, used for cron jobs
--verbose          Increases output
--stats            Adds rsync transfer statistics to verbose output
--partial          Allows rsync to keep partial downloads that can be resumed later (experimental)
--no-maxtime       Disables any soft and hard execution time checks
--force-unlock     Will override any existing active or dead locks on initiator and target replica
--on-changes       Will launch a sync task after a short wait period if there is some file activity on
initiator replica. You should try daemon mode instead

[QUICKSYNC OPTIONS]
--initiator=""          Master replica path. Will contain state and backup directory (is mandatory)
--target=""             Local or remote target replica path. Can be a ssh uri like ssh://user@host.com
:22//path/to/target/replica (is mandatory)
--rsakey=""             Alternative path to rsa private key for ssh connection to target replica
--instance-id=""        Optional sync task name to identify this synchronization task when using multi
ple targets

Additionaly, you may set most osync options at runtime. eg:
SOFT_DELETE_DAYS=365 osync.sh --initiator=/path --target=/other/path

055° [dmpop:/home/dmpop] 128 $ 
```

Figure 1: Run the osync.sh command to view brief usage info.

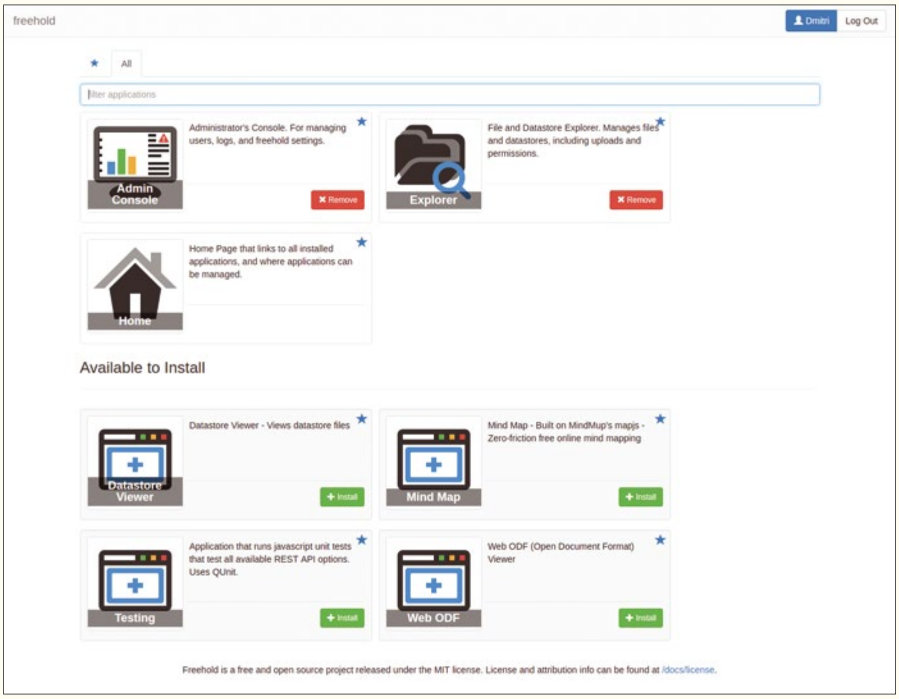

Administrator's Console. For managing users, logs, and freehold settings.

File and Datastore Explorer. Manages files and datastores, including uploads and permissions.

Home Page that links to all installed applications, and where applications can be managed.

Datastore Viewer - Views datastore files

Mind Map - Built on MindMup's mapjs - Zero-friction free online mind mapping

Application that runs javascript unit tests that test all available REST API options. Uses QUnit.

Web ODF (Open Document Format) Viewer

Freehold is a free and open source project released under the MIT license. License and attribution info can be found at /docs/license.

Figure 2: Freehold comes with several apps including Admin Console and Explorer.

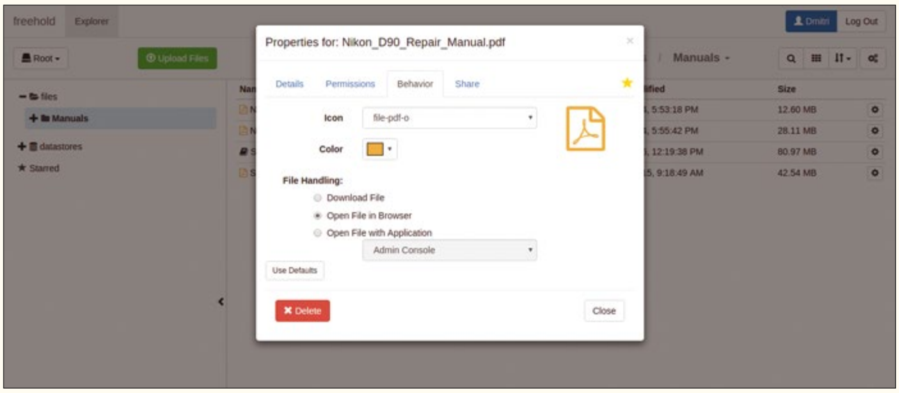

Properties for: Nikon_D90_Repair_Manual.pdf

Icon: file-pdf-o

Color

File Handling:

○ Download File
● Open File in Browser
○ Open File with Application: Admin Console

Figure 3: Editing file properties.

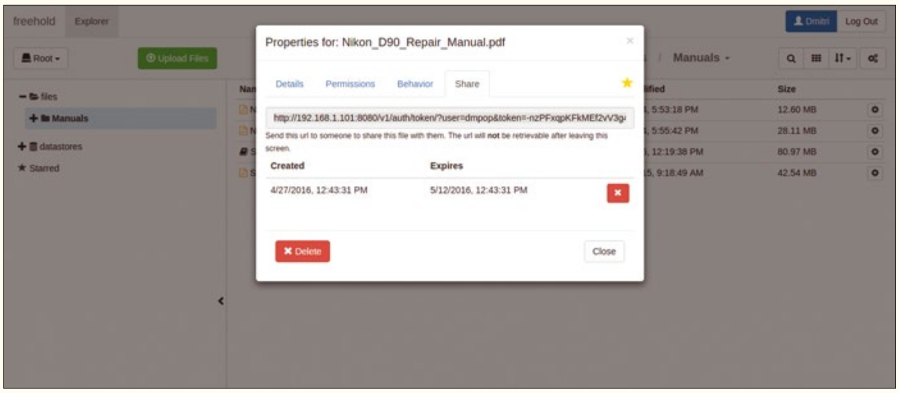

Properties for: Nikon_D90_Repair_Manual.pdf

http://192.168.1.101:8080/v1/auth/token/?user=dmpop&token=-nzPFxqpKFkMEf2vV3g

Send this url to someone to share this file with them. The url will **not** be retrievable after leaving this screen.

Created	Expires
4/27/2016, 12:43:31 PM	5/12/2016, 12:43:31 PM

Figure 4: Files in Freehold can be shared via automatically generated links.

directory and move the exclusion file there, too. Then, run Osync using the `osync.sh /etc/osync/osync.conf` command.

This is all fine and dandy, but what if you need to sync multiple source and target directories? The `osync-batch.sh` script can help you with that. It reads all configuration files in the specified directory and performs synchronization for each profile. So, to keep multiple source/target directories sets in sync, you need to create a `.conf` profile for each set as described above and point the `osync-batch.sh` script to the directory where these configuration files are stored:

```
osync-batch.sh --path=/etc/osync
```

Besides running Osync manually and via a cron job, you can configure the tool to run on boot using a systemd service. To do this, make sure that the `osync.conf` file is properly configured and stored in the `etc/osync` directory, and then run the `systemctl enable osync-srv@osync.conf` command as root.

Freehold: Synchronization and Hosting Combo

Written in Go, Freehold [2] has two things going for it: simplicity and speed. This application is very light on resources and runs blazingly fast even on modest hardware like Raspberry Pi. Better still, Freehold is distributed as a ready-to-go binary, so you can deploy it in a matter of minutes. Grab the appropriate tarball from the project's website, extract the archive, switch to the resulting directory, and run the `./freehold` command. By default, Freehold runs on port 80, but you can easily change that by editing the `port` option in the `~/.config/freehold/settings.json` file. Point your browser to your Freehold installation, specify the desired user name and password, log in using the created credentials, and you should see the application's dashboard.

Before you can start using Freehold, you need to add the Admin Console and Explorer apps to it (Figure 2). The former allows you to manage users, backups, and application settings, and the latter is used to manage files. Installing both apps is a matter of hitting the *Install* button next to each item in the *Available to Install* section.

Populating Freehold with files couldn't be easier. Launch the Explorer app, switch to the *files* section in the navigation sidebar, and drop the files from your machine onto the main area to upload them. Alternatively, you can use the dedicated *Upload Files* button.

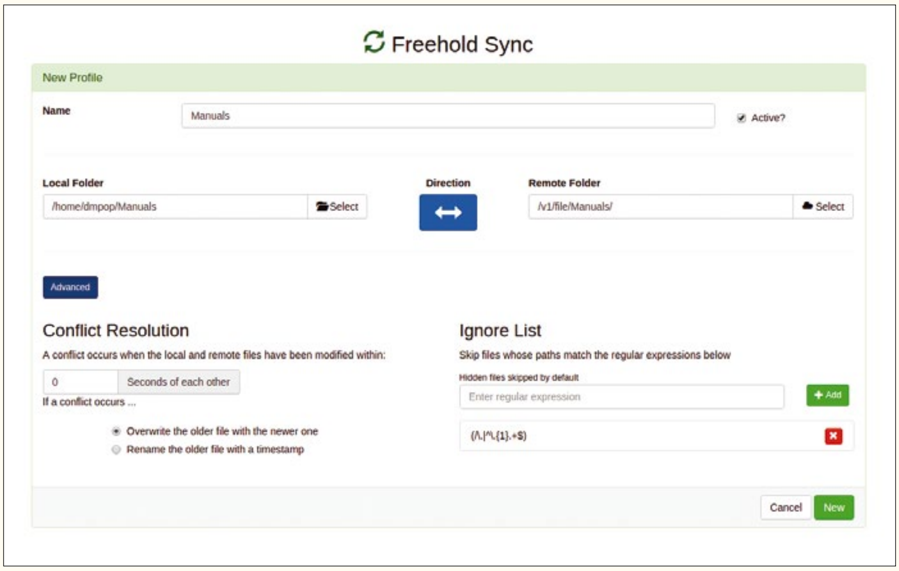

Freehold Sync

New Profile

Name: Manuals ☑ Active?

Local Folder: /home/dmpop/Manuals Select

Direction: ↔

Remote Folder: /v1/file/Manuals/ Select

Advanced

Conflict Resolution

A conflict occurs when the local and remote files have been modified within:

0 Seconds of each other

If a conflict occurs ...

● Overwrite the older file with the newer one
○ Rename the older file with a timestamp

Ignore List

Skip files whose paths match the regular expressions below

Hidden files skipped by default

Enter regular expression + Add

(^.|^\.{1}.+$) ✖

Cancel New

Figure 5: Freehold-Sync client turns Freehold into a synchronization platform.

Click on the *Properties* icon next to the file item, and you can tweak the file's properties in the *Properties* window (Figure 3). Here, you can star the file, modify its permissions and owner, and change the default color and handling behavior. The latter two properties are global (i.e., they apply to all files of the same type). The *Sharing* section of the *Properties* window lets you generate a sharable link to the current file with a single press on the *New Link* button (Figure 4). The system automatically sets an expiration date for the link, but you can remove it, if necessary.

Files in Freehold can be organized into folders and subfolders, and you can switch between the icon and list views. You can also sort items by name, owner, modification date, and size. A search feature is available for finding specific files and documents.

The Admin Console app provides access to administrative functions, such as adding and managing user accounts, viewing logs, tweaking Freehold's settings, and creating backups. All features in the administration dashboard are straightforward in use. Creating a user account, for example, is a matter of pressing the *New User* button under the *Users* tab and filling out the required fields. Because Freehold works perfectly well with the default configuration, there is no need to tweak the settings. However, if you decide to do that, brief descriptions provide clues to what each configuration option does.

In addition to hosting and sharing files, Freehold can act as a synchronization platform courtesy of the dedicated Freehold-Sync tool (Figure 5). Like Freehold, the sync client is written in Go, and it requires no installation. Grab the Freehold-Sync client for your platform, unpack the downloaded archive, switch to the resulting directory, and run the `./freehold-sync` command.

Point your browser to `127.0.0.1:6080`, and you should see Freehold-Sync's dashboard. Hit the *New Profile* button to create and configure a synchronization profile. Add paths to the local directory and a remote Freehold folder you want to keep in sync, specify a syncing direction, and give the profile a name. Press the *Advanced* button if you want to modify the default conflict resolution behavior and add an ignore rule. Hit *New* to save the profile and the client will keep the specified directories in sync.

Final Word

Both Osync and Freehold are really nifty tools. Osync transforms rsync into a capable bidirectional synchronization tool that boasts a range of genuinely useful features and is supremely easy to use. Freehold, in turn, is a no-nonsense platform for hosting and sharing files that can be deployed in a matter of minutes even on modest hardware like Raspberry Pi. Better still, the accompanying Freehold-Sync client provides an easy way to keep your data in sync across multiple machines. ■■■

INFO

[1] Osync: *github.com/deajan/osync*

[2] Freehold: *tshannon.bitbucket.org/freehold*

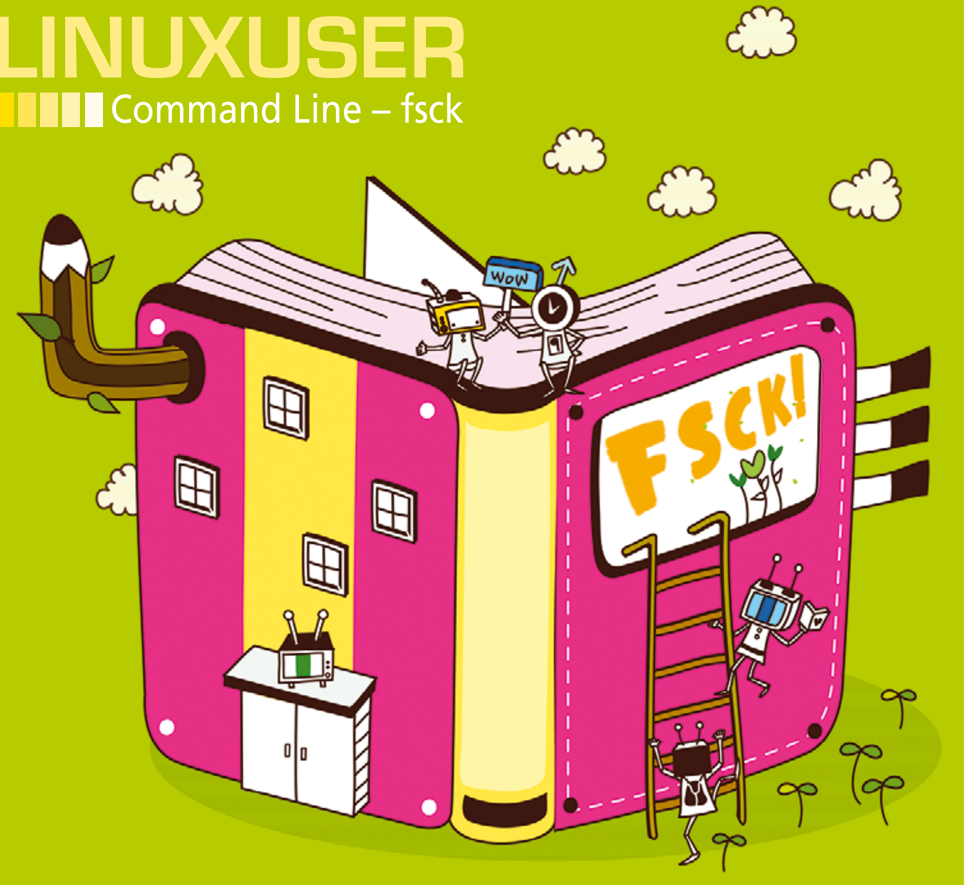

Get to know the fsck command

Situation Normal, All Fscked Up

Learn how to use fsck's capabilities to solve filesystem problems. By Bruce Byfield

The fsck command [1] is often used as a euphemism for a well-known swear word – and not just because it is a four-letter word that starts with *f* and ends with *k*. Typically, the command becomes relevant only when a filesystem needs repairing before the computer will finish booting or starting a graphical interface, and the repair work requires human intervention. In such circumstances, users may feel like swearing as they try to cope with their limited knowledge of the command. However, with more knowledge of fsck's background and options, such moments become much less alarming.

An abbreviation of FileSystem Consistency Check, fsck is installed by default on Linux systems as part of the *util-linux* package. Properly speaking, it is not a separate command at all, but a front end for filesystem-checking commands such as e2fsck, dosfsck, and fsckvfat, all of which work in very similar ways, usually in close association with /etc/fstab [2], which lists the available filesystems. These commands can still be run separately, but, thanks to fsck, users no longer need to remember each of them. Instead, in most cases fsck calls the appropriate command as needed, and the individual commands are only needed in advanced circumstances that few users are likely to encounter.

For convenience, either run fdisk -l (Figure 1) or else run less /etc/fstab before using fsck (Figure 2). Either command gives you a list of partitions on the systems, which can help you ensure that you make all necessary repairs. Make sure you run fsck on the correct drive. Finding the correct drive is essential, because it is possible to bypass fsck's defaults and run on a mounted drive, which can permanently corrupt the drive. Instead, use umount [3] to unmount the filesystem first. Should the filesystem not be unmountable – as often happens with the root partition at boot time – start a Live DVD such as GParted to run fsck.

If you forget to unmount a drive before checking it, fsck will start to run, but it will warn you before proceeding about what you are doing. Usually, a drive that is used mainly for storage can be checked without being unmounted, as long as you close other applications that are running. However, checking active filesystems – like /root or /usr – can crash your system and even permanently corrupt it. On the whole, you should err on the side of caution and pay attention to fsck's warning about mounted drives, no matter how they are used (Figure 3).

Fsck has a standard command structure of:

```
[COMMAND] [OPTIONS] [FILESYSTEM]
```

Filesystems can be added to the command in a comma-separated list, and each can be specified as a device: such as /dev/sda1, a label, or a mountpoint. Add the -A option to check all drives without naming any filesystem in the command, or type the command without any mount point to run fsck one filesystem at a time, in the order they are listed in /etc/fstab.

You can continue running the command, but you do risk making a change

BRUCE BYFIELD

Bruce Byfield is a computer journalist and a freelance writer and editor specializing in free and open source software. Bruce's most recent book, *Designing with LibreOffice*, was released under a Creative Commons License in March 2016. You can buy or download his book at *http://designingwithlibreoffice.com/download-buy/*. In addition to his writing projects, he also teaches live and e-learning courses. In his spare time, Bruce writes about Northwest Coast art. You can read more of his work at *http://brucebyfield.wordpress.com*.

that may crash your system. As you work, you might notice that fsck does not include any GNU options, with specific names and two hyphens before them. This lack does not affect functionality but presumably reflects the fact that fsck was written long after the basic GNU/Linux system. Also, the command is intended not for ordinary users. Instead, it is intended for administrators, who in theory do not require built-in guides of GNU options to functionality in their options.

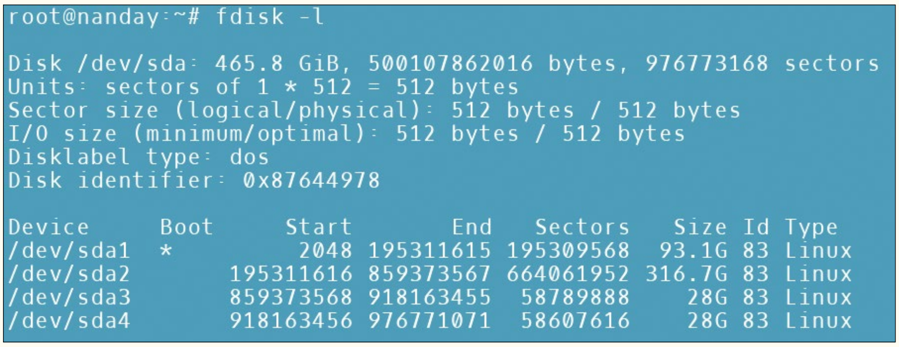

```
root@nanday:~# fdisk -l

Disk /dev/sda: 465.8 GiB, 500107862016 bytes, 976773168 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x87644978

Device     Boot     Start       End   Sectors   Size Id Type
/dev/sda1  *         2048 195311615 195309568  93.1G 83 Linux
/dev/sda2       195311616 859373567 664061952 316.7G 83 Linux
/dev/sda3       859373568 918163455  58789888    28G 83 Linux
/dev/sda4       918163456 976771071  58607616    28G 83 Linux
```

Figure 1: Before running fsck, run fsck -l to see a list of available partitions.

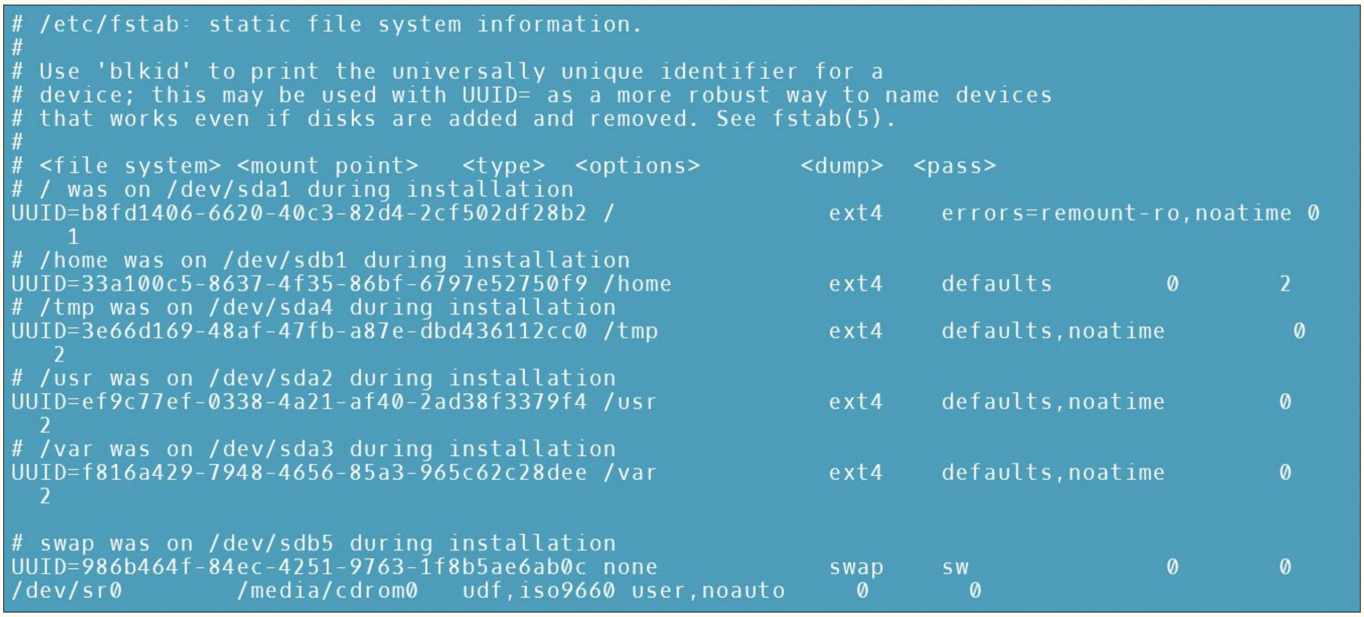

```
# /etc/fstab: static file system information.
#
# Use 'blkid' to print the universally unique identifier for a
# device; this may be used with UUID= as a more robust way to name devices
# that works even if disks are added and removed. See fstab(5).
#
# <file system> <mount point>   <type>  <options>       <dump>  <pass>
# / was on /dev/sda1 during installation
UUID=b8fd1406-6620-40c3-82d4-2cf502df28b2 /               ext4    errors=remount-ro,noatime 0
    1
# /home was on /dev/sdb1 during installation
UUID=33a100c5-8637-4f35-86bf-6797e52750f9 /home           ext4    defaults        0       2
# /tmp was on /dev/sda4 during installation
UUID=3e66d169-48af-47fb-a87e-dbd436112cc0 /tmp            ext4    defaults,noatime         0
   2
# /usr was on /dev/sda2 during installation
UUID=ef9c77ef-0338-4a21-af40-2ad38f3379f4 /usr            ext4    defaults,noatime        0
  2
# /var was on /dev/sda3 during installation
UUID=f816a429-7948-4656-85a3-965c62c28dee /var            ext4    defaults,noatime        0
  2

# swap was on /dev/sdb5 during installation
UUID=986b464f-84ec-4251-9763-1f8b5ae6ab0c none            swap    sw              0       0
/dev/sr0        /media/cdrom0   udf,iso9660 user,noauto     0       0
```

Figure 2: Opening /etc/fstab provides useful information, including the filesystems, their UUID, and, in the sixth column for each filesystem, the order in which to check them.

Filesystem Selection Options

About half of fsck's options set how the command selects the filesystems to run upon. The most important of these options is -A, which tries to save time by checking all filesystems listed in /etc/fstab in a single pass instead of by running the command multiple times. By itself, -A begins with the root filesystem, than moves on to others according to fstab's sixth column, pass_no. Filesystems with a pass_no of 1 follow the root system, then those with a pass_no of 2. Filesystems with the same priority are checked in parallel, if possible. By contrast, filesystems with a pass_no of 0 are not checked at all.

One limitation of -A is that it cannot run in parallel on RAID systems or encrypted disks. An even more important one is that it does not detect filesystems not listed in /etc/fstab, which means that it might fail to detect a filesystem that has a problem.

Two other options modify -A. When -P is included, the root filesystem is checked in parallel with the other filesystems. The man page for fsck warns that if -P is used and the root filesystem requires repairs, other commands, such as e2fsck may be corrupted, although normally the practice might save time. By contrast, -R is much safer, because it signals to skip the root filesystem. This is useful if you know that it has already been mounted – as generally happens if a problem emerges while the system is booting.

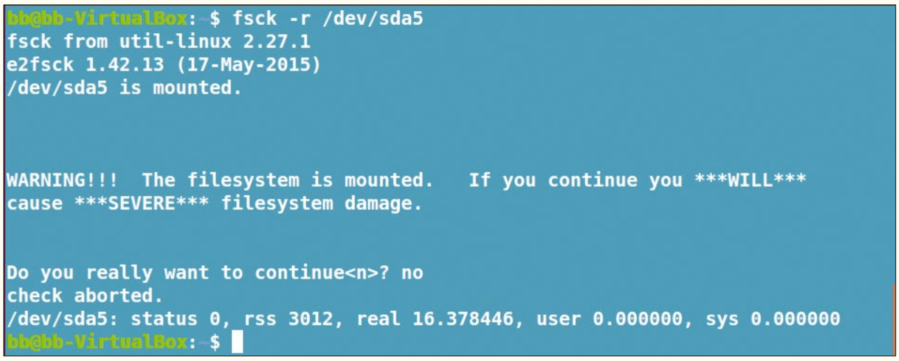

```
bb@bb-VirtualBox:~$ fsck -r /dev/sda5
fsck from util-linux 2.27.1
e2fsck 1.42.13 (17-May-2015)
/dev/sda5 is mounted.



WARNING!!!  The filesystem is mounted.   If you continue you ***WILL***
cause ***SEVERE*** filesystem damage.


Do you really want to continue<n>? no
check aborted.
/dev/sda5: status 0, rss 3012, real 16.378446, user 0.000000, sys 0.000000
bb@bb-VirtualBox:~$ 
```

Figure 3: The fsck command warns when you are about to check a mounted drive.

Additionally, multiple filesystems can be set to run one after another with the -s option. This option is especially useful when fsck runs in interactive mode (see below), asking for your input before performing any repairs.

Repair Options from e2fsck

For repairs, fsck borrows options from the commands for which it fronts. That is especially true in the cases of e2fsck, which repairs the popular ext2/3/4 filesystems.

By default, fsck runs interactively, asking for confirmation before attempting any repairs. However, users can change this behavior to suit themselves. For example, when -p is added to the command, the command repairs (or preens, in the jargon of fsck) problems that can be fixed without the user making any decision. If the option encounters a problem it cannot handle, then it prints a description of the problem and exits. The -a option acts the same way, but it exists only for backward compatibility with older releases. The man page advises that -p be used whenever possible rather than -a.

Neither -p nor -a can be used with -n or -y. With -n, all requests for decisions by the user are answered automatically with "No", whereas with -y, all requests are answered automatically with "Yes". Running -n causes few problems, because any problems are simply printed out as the command runs, and users can simply try running fsck again with other options. However, -y could aggravate the problems, so it should be used carefully.

Options to Change fsck's Behavior

Other options for fsck affect how the command runs rather than where. As in most commands, verbosity (-V) assures users that the command ran successfully or else helps pinpoint problems. Advanced users also have the -r option, which lists such statistics as the exit status, the time fsck took to run, and the user and system CPU time in a comma-separated list.

Usually, the filesystem format can be read from /etc/fstab. However, if you want fsck to run on a filesystem not listed in /etc/fstab or have any other dificulty, you can hard-code its format with -t FILESYSTEM.

All in all, fsck can be a tricky command to run, and adding options can sometimes make the situation worse. Under these circumstances, using -M to skip mounted filesystems can sidestep the most common problems. Similarly, once you have decided on which options to use, you should use -N to simulate the command without actually doing anything, so that you can avoid unexpected, or even fatal results (Figure 4).

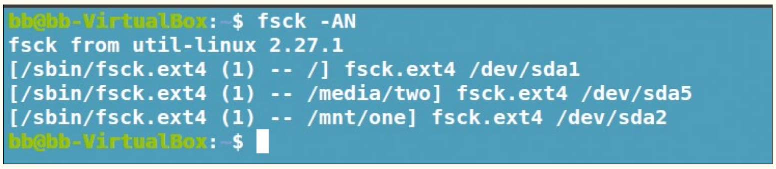

```
bb@bb-VirtualBox:~$ fsck -AN
fsck from util-linux 2.27.1
[/sbin/fsck.ext4 (1) -- /] fsck.ext4 /dev/sda1
[/sbin/fsck.ext4 (1) -- /media/two] fsck.ext4 /dev/sda5
[/sbin/fsck.ext4 (1) -- /mnt/one] fsck.ext4 /dev/sda2
bb@bb-VirtualBox:~$
```

Figure 4: The -N option simulates running fsck without actually running it. Here, fsck simulates running the -A command. Notice the use of square brackets.

Environmental Variables

Environmental Variables can have as much effect on fsck's behavior as the available options. The command's man page includes these variables:

- FSCK_FORCE_ALL_PARALLEL: Sets fsck to run in parallel. This setting helps fsck to run on RAID systems or on filesystems on the same device.
- FSCK_MAX_INST: Sets the maximum number of filesystem checks that can run at the same time so that system resources are overloaded. 0 means that no limit is set and is currently the default, although the man page hints that might change in later releases.
- PATH: Sets the directories where fsck should look for filesystem checkers. Fsck searches /sbin, /sbin/fs.d, /sbin/fs, /etc/fs, and /etc, then the directories in the path.
- FSTAB_FILE: Sets an alternate location for /etc/fstab, allowing users to test fsck with less worry over consequences.

Taking Control

Faced with the need to check a filesystem, many users – including me – often simply enter the filesystem. Sometimes, that is all that is needed, but it leaves you helpless if the problem remains unsolved.

If you learn more about fsck's capabilities, however, you can consider the options systematically instead of panicking. Just remember that, when all repairs are made, you should remount each filesystem or else reboot. ■■■

INFO

[1] fsck: *https://en.wikipedia.org/wiki/Fsck*

[2] fstab: *https://en.wikipedia.org/wiki/Fstab*

[3] mount (Unix): *https://en.wikipedia.org/wiki/Mount_Unix*

Shop the Shop shop.linuxnewmedia.com

RASPBERRY PI ADVENTURES

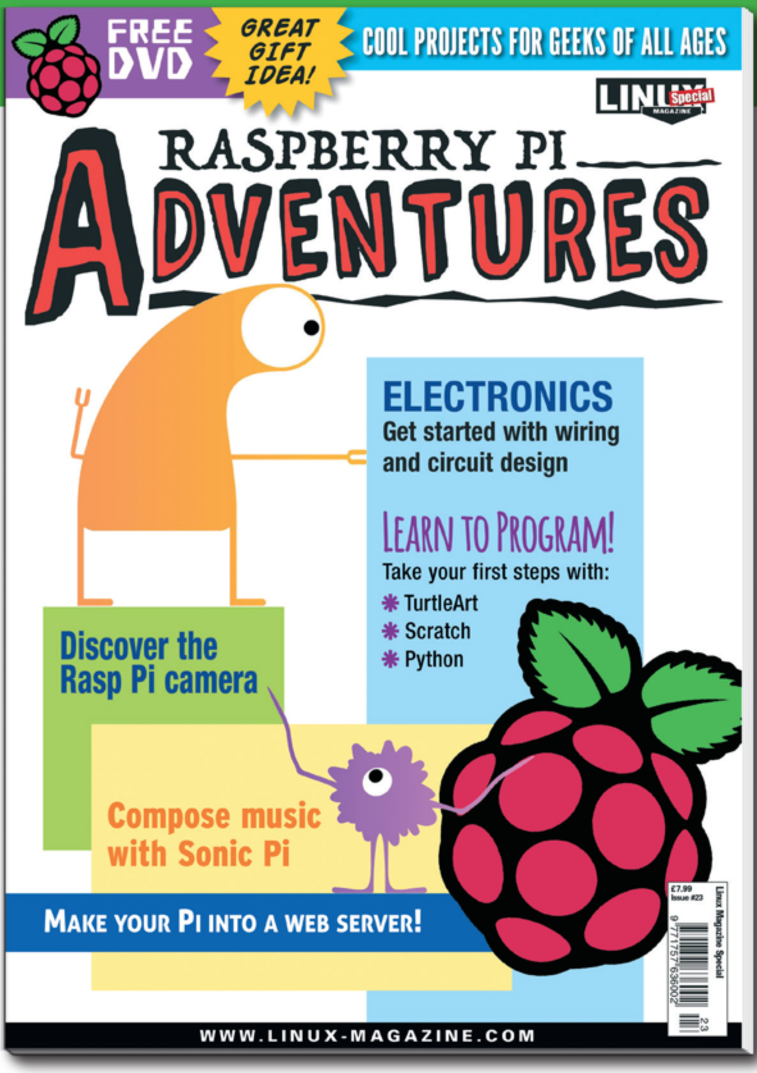
FREE DVD
GREAT GIFT IDEA!
COOL PROJECTS FOR GEEKS OF ALL AGES
LINUX Special MAGAZINE
RASPBERRY PI ADVENTURES
ELECTRONICS
Get started with wiring and circuit design
LEARN TO PROGRAM!
Take your first steps with:
TurtleArt
Scratch
Python
Discover the Rasp Pi camera
Compose music with Sonic Pi
MAKE YOUR PI INTO A WEB SERVER!
WWW.LINUX-MAGAZINE.COM

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

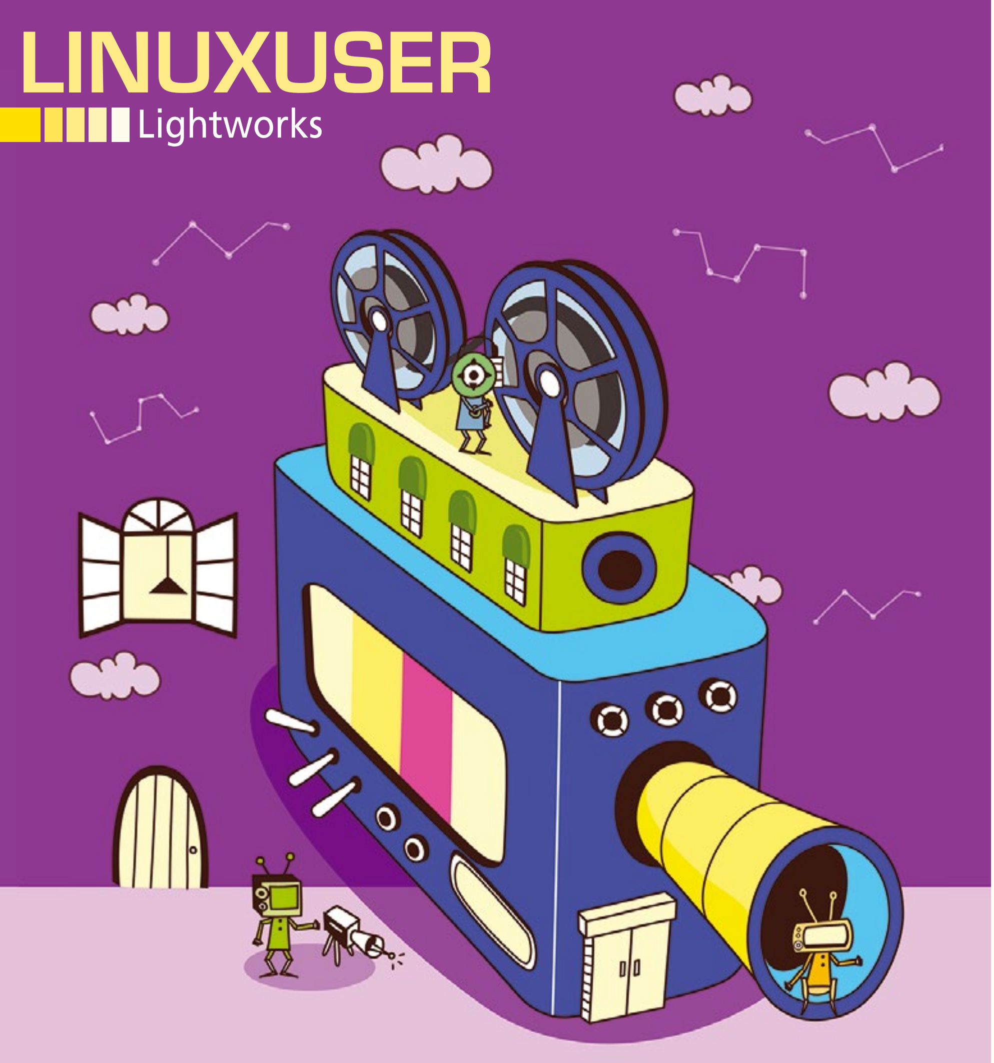

Cut videos and add effects using Lightworks

Quick Step

The free editing program Lightworks Free makes small video editing projects easy. If you need more, use the Pro version. *By Andreas Reitmaier*

The video editing software Lightworks [1] organizes files into projects. Creating a corresponding directory at the filesystem level to access video, audio, and image files will avoid problems later. If you copy any material to this new directory that you will want to use later, you will be able to access it quickly and easily via the Import dialog.

You don't need to give too much thought to the file formats, because Lightworks supports a wide range from the professional and consumer sectors [2]. However, the output format depends on the software version you are using (see the "Versions" box for more information).

Importing Media

Lightworks is organized into various levels – with the top level called the project (Figure 1). On startup, a dialog appears either displaying existing projects or prompting you to create a new one. You just need to enter a name in the text box, possibly supplemented with a particular frame rate – the default here is always *Auto*.

The interface initially appears to be clearly organized. You will see just one small field with a few icons on the left-hand edge (Figure 2). The second symbol from the top is for the *Import* function. Here, you will the find the *Places* entry, which lists some default directories, as well as local and removable media.

At the bottom of the window, you will see the *Create Link* entry. Here, you can set whether the software copies the selected files to a local directory or leaves them in their storage location.

Now, you can select a file and import it. The program then displays an overview of the media in the Content Manager, which lists all the media and media sections you are using in the project. Double-clicking an entry opens a preview, which provides the option to define the desired section.

Creating a Scene

Next, you can use the toolbar on the edge of the screen to open a first edit. This is the area where you can arrange and cut media, effects, and titles. Now drag and drop the first file from the Content Manager into the editor.

The individual tracks will appear there – usually one video and two audio tracks (audio tracks in waveform view). You will see the video and some tools in a slightly smaller preview.

You can perform simple cuts fairly quickly here: If you just want to cut a bit from the start and end of a clip, first move the cursor to the start. Then, place a start or *in* mark in the preview or editor. Next, move the cursor to the position you want to cut the clip. Finally, you can define an *out* mark (Figure 3).

You'll see two delete functions in the editor window: You can either remove the selected area and replace it with "Black" or delete it and leave it to the program to close the gap. The latter proves to be the easiest option in many

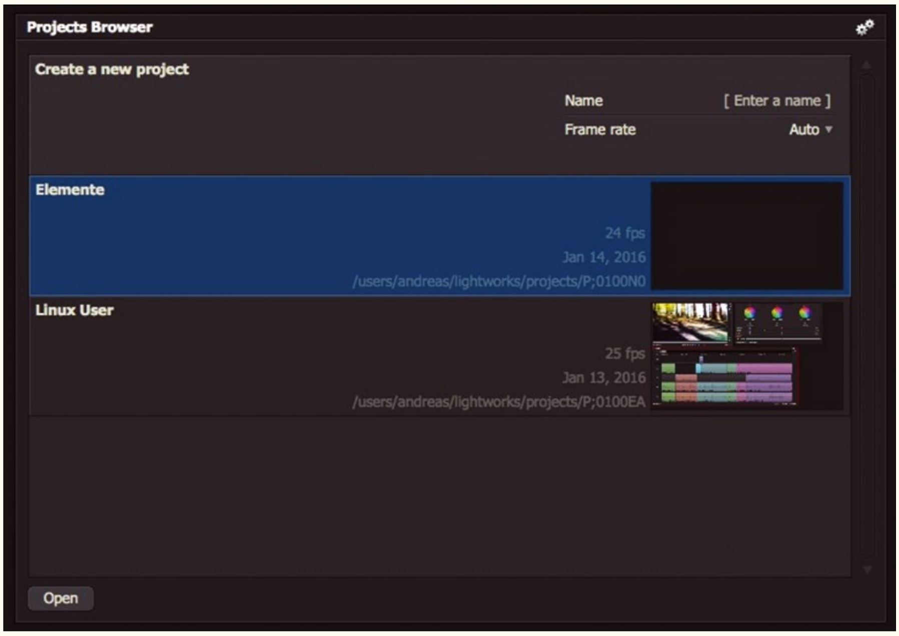

Projects Browser

Create a new project

Name [Enter a name]

Frame rate Auto

Elemente
24 fps
Jan 14, 2016
/users/andreas/lightworks/projects/P;0100N0

Linux User
25 fps
Jan 13, 2016
/users/andreas/lightworks/projects/P;0100EA

Open

Figure 1: Lightworks organizes the work into projects – and you can work on several projects at once. The application keeps track of all media, settings, and the window arrangement.

cases. Conversely, you can add new clips this way, provided you move the other components.

Preparing the Opening Credits

You usually only deal with components like the opening and closing credits once the project is mostly complete. In this test, however, I dealt with them first. In larger projects, these steps should be considered the final tasks or be performed in separate edits.

In this test, a cut sequence served as a background video for the opening credits. You might prefer to use a still image or a photo, because they are calmer and distract less from the title. If you use a video sequence, optimizing the clip is worthwhile. You can use a simple filter to make the background blurry, so it doesn't distract.

An *Effects* button can be found at the bottom right of the timeline window. The software provides a pretty wide selection of filters. The *Blur* effect used in this example is in the *Video* | *Stylize* section. Drag the filter onto the clip to which you want to apply

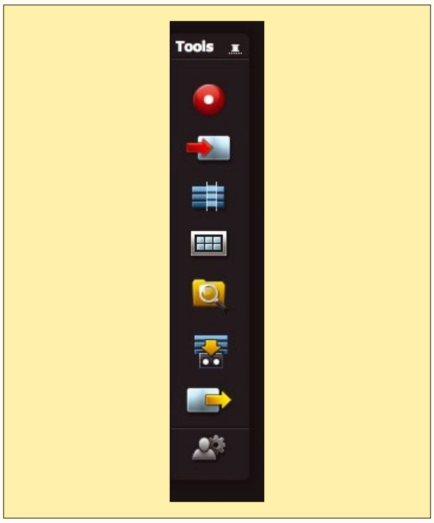

Figure 2: Lightworks initially requires only a minimal set of commands. You can access all other options using keyboard shortcuts, the settings in each window, or context menus.

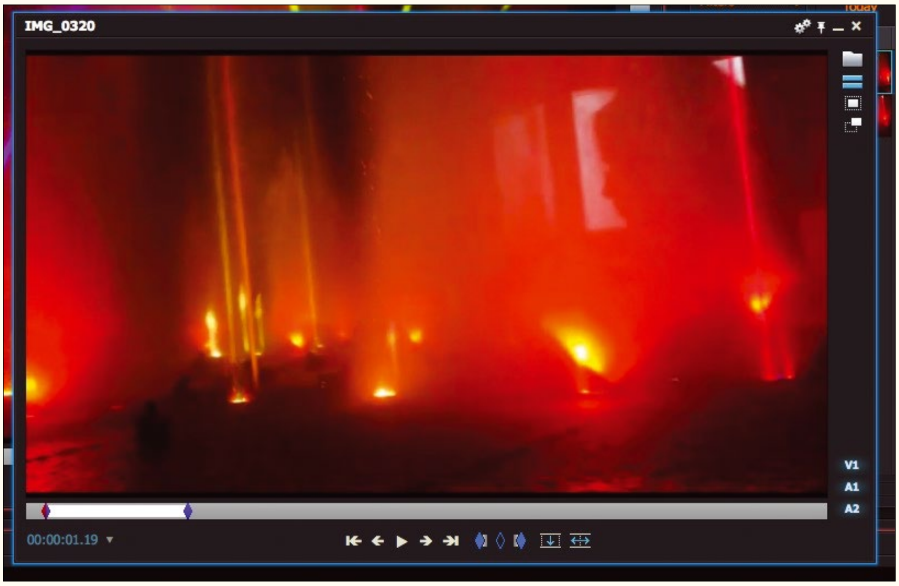

Figure 3: Select the section of a clip you want to use in the editor with *in* and *out* marks. You can choose multiple sections of a file.

VERSIONS

Two versions of Lightworks are available: Free and Pro. You can choose to buy a subscription or purchase the software outright for the Pro version. The price for the subscription is $24.99 per month or $174.99 per year. Purchasing a license outright costs $437.99.

The main difference between the Free and Pro versions is the export options. The Free version only creates MP4/H.264 videos with a resolution of 720p. The Pro version supports more video formats, diverse formats for exchange with editing systems and cameras, and direct Blu-ray and DVD export. Lightworks Pro also provides some special options for editing.

the effect. The program then opens the Effects editor, which you can use to change the strength of the blur filter, as required. Other filters also provide more options.

If you want to customize or remove the effect later, right-click the corresponding section of the video. *Effects*, the top entry in the context menu, now appears with various options (Figure 4).

Creating the Opening Credits

Next up are the opening credits or title. For these, you need to access the effects selection. You'll find text objects in the *Video* section under *Titles* and *Video* | *Lower Thirds*. The software provides simple titles, which can, however, move through the image to some extent. *Lower Thirds* contains typical information that shows messages within a video.

As with the blur effect used previously, you can drag the selected title effect (e.g., the *Roll* effect) onto the video clip over which you want to display the text. As before, you can edit the content of the effect in the corresponding editor window.

The text slides onto the image from the bottom to the top, so you can enter multiple lines of text. You can also determine the font type, select *Bold* or *Italic*, and specify the *Size* and *Opacity* of the font. You can even animate both the size and opacity of the font. This process is relatively simple: Just put the playback cursor where you want to insert a keyframe – that is, a point at which an animation is activated.

Now find the controller whose values you want to animate. This value is Opacity in the example. You'll find a small button to activate between the name and the control-

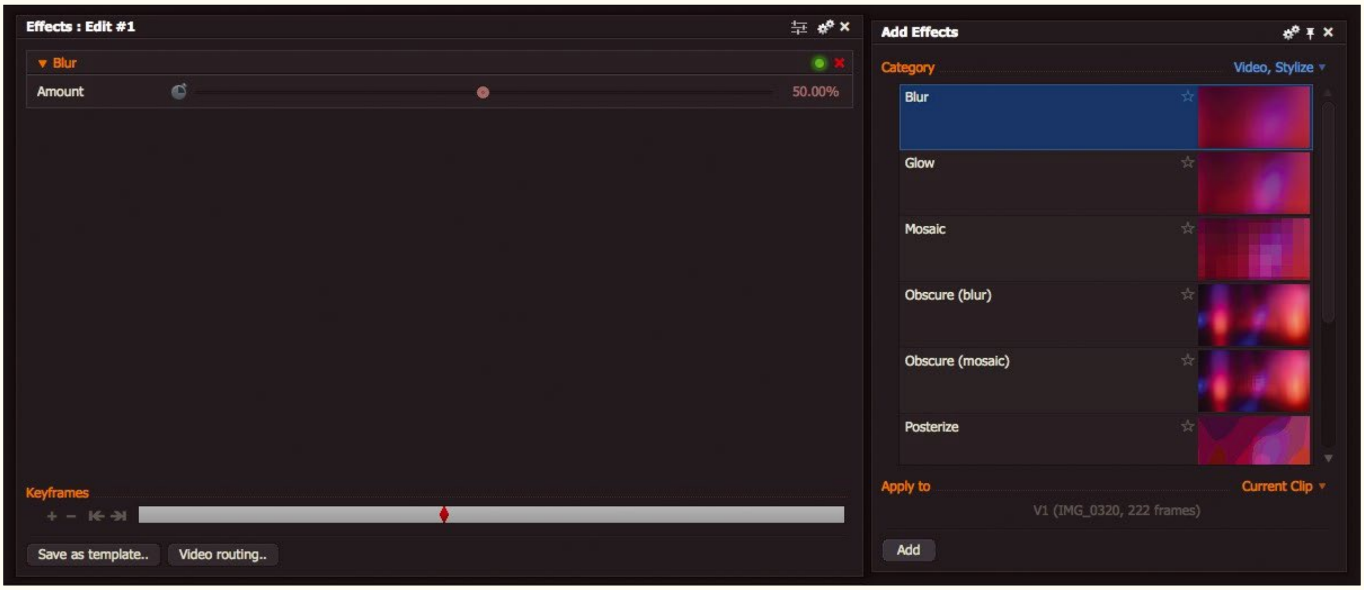
Effects : Edit #1

Blur

Amount 50.00%

Add Effects

Category: Video, Stylize

Apply to: Current Clip

V1 (IMG_0320, 222 frames)

Figure 4: In the selection of effects, you will see a rough preview applied to the clip in the current cursor position.

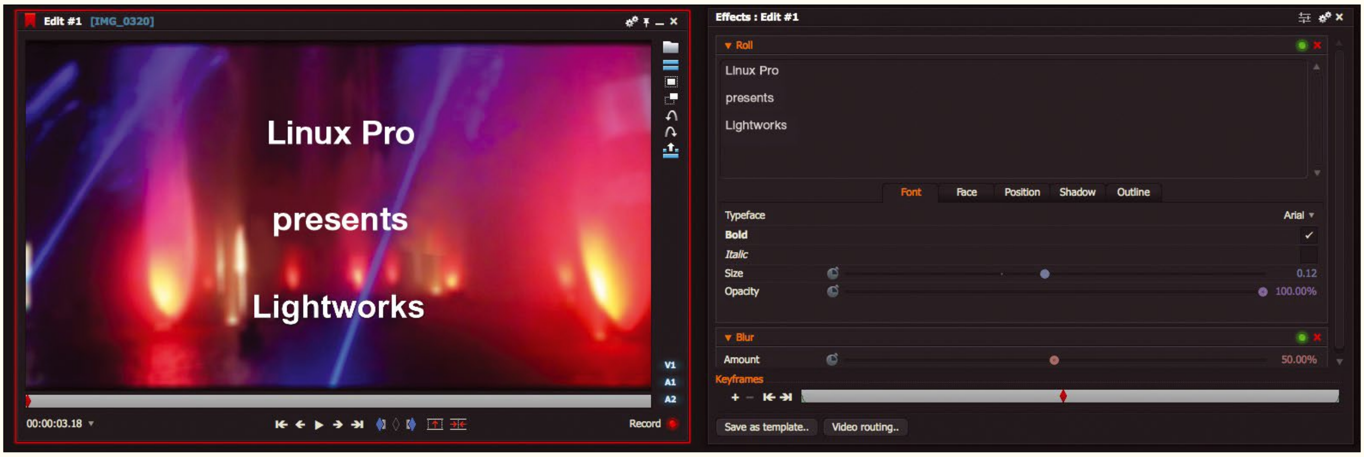
Edit #1 [IMG_0320]

Linux Pro

presents

Lightworks

00:00:03.18

Effects : Edit #1

Roll

Linux Pro
presents
Lightworks

Typeface: Arial
Bold: ✓
Italic
Size: 0.12
Opacity: 100.00%

Blur

Amount: 50.00%

Figure 5: The program lets you animate many effects. An editor opens as soon as you activate the effect and lets you control elements such as position, opacity, and color.

ler. As soon as you move the slider, the software will place a keyframe with the specified value at the point the cursor is resting.

These markings are at the bottom in the Effects editor. There, you have the option to edit them or even save them as a template for future videos. This can be quite handy for the opening credits if you plan to produce similar projects.

The text effect has other parameters that also animate the software to a large extent. For example, you'll find color settings under *Face*, which you can also control using keyframes. *Position* provides more precise settings for positions and intervals. *Shadow* and *Outline* determine the appearance of each letter (Figure 5).

The color effects are not perhaps crucial for practical use. However, the progress controller, which you'll find under *Position*, saves you from having to split the opening credits into several suitable pieces.

A title doesn't usually start at exactly the same time as the film. In Lightworks, you can use the *Position* | *Progress* function for a delayed start. Just place the cursor in the position at which the software should incorporate the title.

Then, you can activate the button for the keyframes and drag the progress controller to 100 percent. You should also place a keyframe with 100 percent in the position where the title should end. Then, add a keyframe with 0 percent at both the beginning and the end of the clip.

Cutting Scenes

Next up is the process of inserting additional clips into the project. You have two options for dealing with film material: If the material already consists of lots of little clips, and you want to use all of them, just drag them into the editor window. For a more complex arrangement, first create the number of required video and audio tracks. Right-clicking will take you to the context menu with the command *Add Tracks*.

However, the simple method for stringing clips together is usually the one to use when dealing with the software for the first time. To begin, open the next video in the clip editor by double-clicking it from the Content Manager. Under the video, you'll find various controls and buttons for placing *in* and *out* markers. You can use them to define the start and end of the section you want to add to the film.

You now have two options to integrate the selected part: The button with the downward arrow inserts the selection at the point where the cursor is and overwrites the existing material until the clip comes to an end. The button with the outward arrow inserts the selection in the same place, but this moves the existing material backwards.

Dissolving

You can create a transition in the same way as the effects you made use of for the opening credits. Just open up the Effects palette (which is the button at the bottom

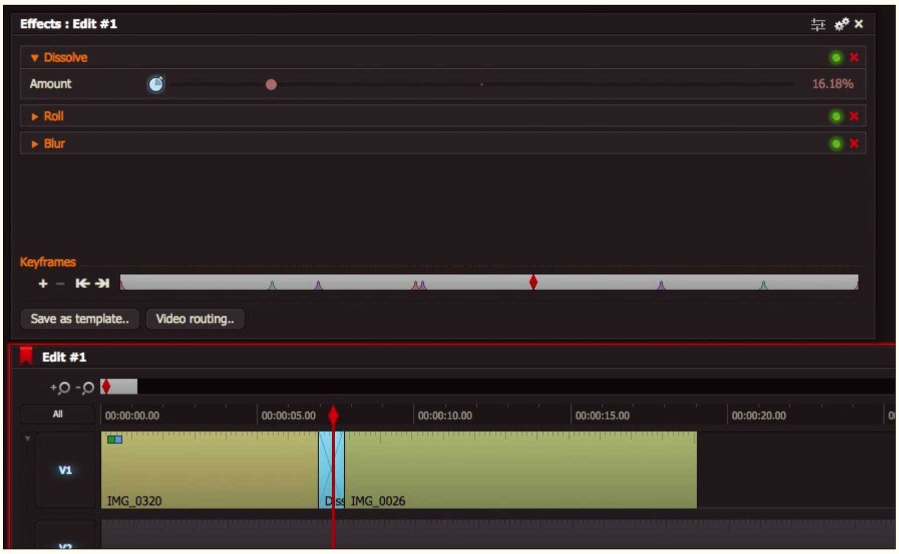

Amount 16.18%

Figure 6: To create a transition between two clips, drag the selected effect from the palette onto the interface between the clips.

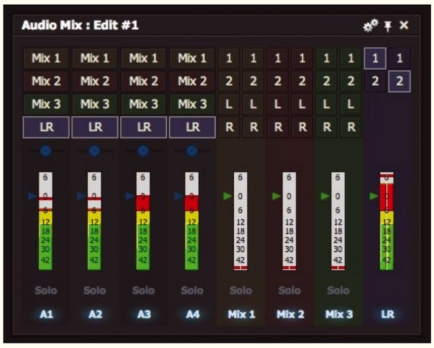

Figure 7: The audio mixer provides a quick mix down, but you can create more complex situations pretty easily.

right of the timeline window). You'll find a handful of transition effects in the *Video | Mixes* section. Lightworks tries to show a preview of the transition in the effect selection, although this isn't always particularly useful.

Now drag the desired effect into the interface between two clips (Figure 6). As before, you have the option to configure the effect in the editor – sometimes with multiple parameters, depending on the type. *Transitions* make it possible to animate most options.

Working with Sound

The software handles audio files much like their video counterparts: An audio clip contains audio tracks with a black background whose video track is essentially disabled. You can also just use the audio part of videos by disabling the video track.

To do so, open a clip in the Content Manager by double-clicking it. You'll find the track display in the bottom-right corner of the editor with the markers *V1*, *A1*, and *A2*. You can select the *in* and *out* points of the audio clip for the video track. Then, you can replace this selection with an existing audio track or insert the selection as described for the videos (Figure 7).

A small, integrated mixer for audio files can be used to adjust the film's sound as required. Just open up the Settings menu – either via the button with the gear icon in one of the windows or via the context menu. Then, choose the *Audio Mixer Panel* entry.

Using this button, you can set the input channels to mix channels, which you can in turn assign to the left or right output and thus control the corresponding volume everywhere. This way, you can coordinate a video's background music and original sound, among other things.

Export

The biggest limitation with the free version of Lightworks is its export function. You can open the Export dialog via the toolbar (Figure 8), and two options are available under *Format*: the *Lightworks archive* is for saving projects for processing later, and *YouTube* is Lightworks' name for MP4/H.264 files. These are again limited to 720p format. Only the Pro version allows you to export files to other, more professional formats.

Choose the *YouTube* export format. Then, set the *Frame rate* and resolution. Under Destination, you can determine a location on the hard disk and give the file a name. You can then start the export by clicking *Start*, or you can upload the video directly to the video platform.

You just need to enter your username and password. In testing, I experienced some issues with the direct upload option, which occasionally terminated. This isn't Lightworks' fault, but rather the fault of the website's interface. Exporting onto a hard disk before uploading via the browser is thus an easier option.

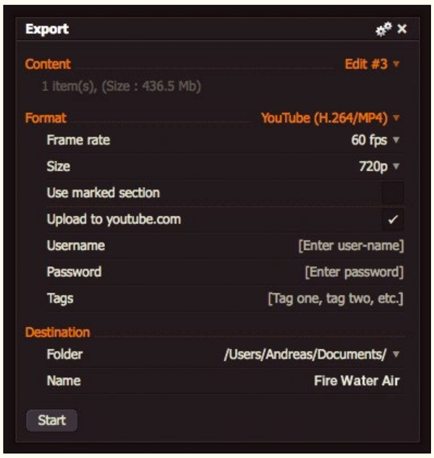
Export
Content — Edit #3
1 item(s), (Size : 436.5 Mb)
Format — YouTube (H.264/MP4)
Frame rate — 60 fps
Size — 720p
Use marked section
Upload to youtube.com — ✓
Username — [Enter user-name]
Password — [Enter password]
Tags — [Tag one, tag two, etc.]
Destination
Folder — /Users/Andreas/Documents/
Name — Fire Water Air
Start

Figure 8: The free version of Lightworks provides only a single file format for video export. If you need professional formats, you need to pay for the Pro version.

Conclusions

Lightworks is generally aimed at professional filmmakers. However, the attractive pricing means anyone can cut videos professionally for free as long as they don't need a full HD output format. People swapping from other editing programs might need a bit of time to get used to it, because not every mouse click appears logical at first. Those using the program frequently, however, will quickly achieve their goal using Lightworks, and they can refer to numerous tutorials [3] and a detailed manual [4] for additional help. ■■■

INFO

[1] Lightworks: *https://www.lwks.com*

[2] Features: *https://www.lwks.com/index.php?option=com_content&view=article&id=102&Itemid=213*

[3] Tutorials: *http://www.lwks.com/index.php?option=com_content&view=article&id=162&Itemid=246&start=v12_5*

[4] Manual: *http://www.lwks.com/index.php?option=com_docman&task=doc_download&gid=198*

IT Highlights at a Glance

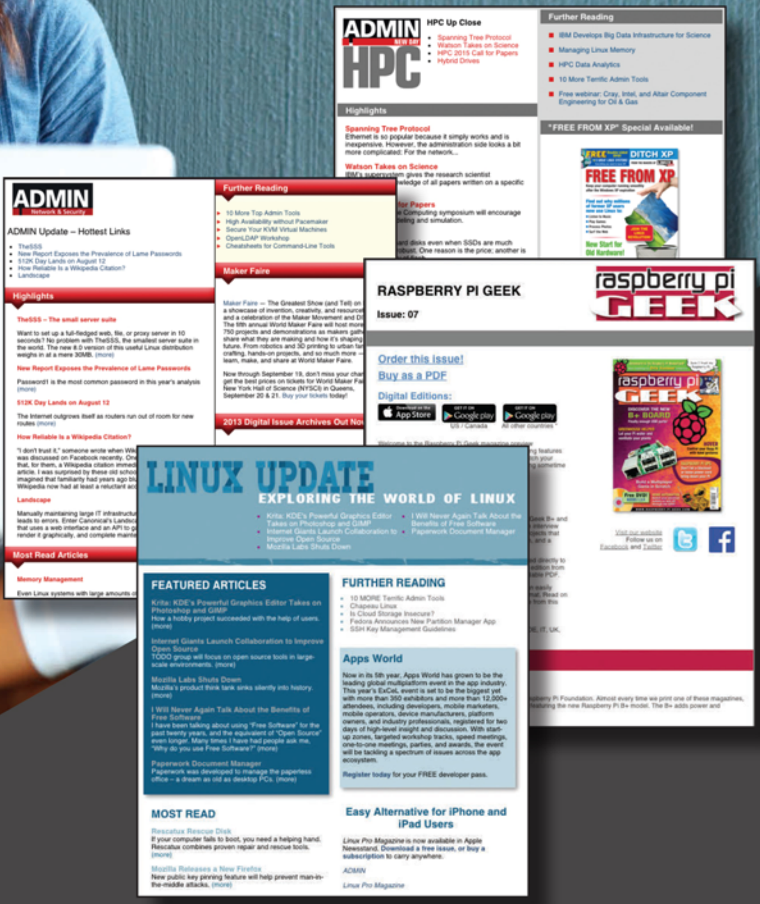

ADMIN NEW DAY HPC

HPC Up Close

- Spanning Tree Protocol
- Watson Takes on Science
- HPC 2015 Call for Papers
- Hybrid Drives

Highlights

Spanning Tree Protocol
Ethernet is so popular because it simply works and is inexpensive. However, the administration side looks a bit more complicated: For the network...

Watson Takes on Science
IBM's supersystem gives the research scientist
wledge of all papers written on a specific

for Papers
e Computing symposium will encourage
eling and simulation.

ard disks even when SSDs are much
obust. One reason is the price; another is

Further Reading

- IBM Develops Big Data Infrastructure for Science
- Managing Linux Memory
- HPC Data Analytics
- 10 More Terrific Admin Tools
- Free webinar: Cray, Intel, and Altair Component Engineering for Oil & Gas

"FREE FROM XP" Special Available!

ADMIN Network & Security

ADMIN Update – Hottest Links

- TheSSS
- New Report Exposes the Prevalence of Lame Passwords
- 512K Day Lands on August 12
- How Reliable Is a Wikipedia Citation?
- Landscape

Highlights

TheSSS – The small server suite

Want to set up a full-fledged web, file, or proxy server in 10 seconds? No problem with TheSSS, the smallest server suite in the world. The new 8.0 version of this useful Linux distribution weighs in at a mere 30MB. (more)

New Report Exposes the Prevalence of Lame Passwords

Password1 is the most common password in this year's analysis (more)

512K Day Lands on August 12

The Internet outgrows itself as routers run out of room for new routes (more)

How Reliable Is a Wikipedia Citation?

"I don't trust it," someone wrote when Wik
was discussed on Facebook recently. One
that, for them, a Wikipedia citation immed
article. I was surprised by these old schoo
imagined that familiarity had years ago bl
Wikipedia now had at least a reluctant acc

Landscape

Manually maintaining large IT infrastructu
leads to errors. Enter Canonical's Landsca
that uses a web interface and an API to g
render it graphically, and complete mainte

Most Read Articles

Memory Management

Even Linux systems with large amounts o

Further Reading

- 10 More Top Admin Tools
- High Availability without Pacemaker
- Secure Your KVM Virtual Machines
- OpenLDAP Workshop
- Cheatsheets for Command-Line Tools

Maker Faire

Maker Faire — The Greatest Show (and Tell) on
a showcase of invention, creativity, and resource
and a celebration of the Maker Movement and DI
The fifth annual World Maker Faire will host more
750 projects and demonstrations as makers gath
share what they are making and how it's shaping
future. From robotics and 3D printing to urban fa
crafting, hands-on projects, and so much more —
learn, make, and share at World Maker Faire.

Now through September 19, don't miss your char
get the best prices on tickets for World Maker Fa
New York Hall of Science (NYSCI) in Queens,
September 20 & 21. Buy your tickets today!

2013 Digital Issue Archives Out No

RASPBERRY PI GEEK

Issue: 07

raspberry pi GEEK

Order this issue!

Buy as a PDF

Digital Editions:

App Store | Google play US / Canada | Google play All other countries *

Visit our website
Follow us on
Facebook and Twitter

LINUX UPDATE

EXPLORING THE WORLD OF LINUX

- Krita: KDE's Powerful Graphics Editor Takes on Photoshop and GIMP
- Internet Giants Launch Collaboration to Improve Open Source
- Mozilla Labs Shuts Down
- I Will Never Again Talk About the Benefits of Free Software
- Paperwork Document Manager

FEATURED ARTICLES

Krita: KDE's Powerful Graphics Editor Takes on Photoshop and GIMP
How a hobby project succeeded with the help of users. (more)

Internet Giants Launch Collaboration to Improve Open Source
TODO group will focus on open source tools in large-scale environments. (more)

Mozilla Labs Shuts Down
Mozilla's product think tank sinks silently into history. (more)

I Will Never Again Talk About the Benefits of Free Software
I have been talking about using "Free Software" for the past twenty years, and the equivalent of "Open Source" even longer. Many times I have had people ask me, "Why do you use Free Software?" (more)

Paperwork Document Manager
Paperwork was developed to manage the paperless office – a dream as old as desktop PCs. (more)

FURTHER READING

- 10 MORE Terrific Admin Tools
- Chapeau Linux
- Is Cloud Storage Insecure?
- Fedora Announces New Partition Manager App
- SSH Key Management Guidelines

Apps World

Now in its 5th year, Apps World has grown to be the leading global multiplatform event in the app industry. This year's ExCeL event is set to be the biggest yet with more than 350 exhibitors and more than 12,000+ attendees, including developers, mobile marketers, mobile operators, device manufacturers, platform owners, and industry professionals, registered for two days of high-level insight and discussion. With start-up zones, targeted workshop tracks, speed meetings, one-to-one meetings, parties, and awards, the event will be tackling a spectrum of issues across the app ecosystem.

Register today for your FREE developer pass.

MOST READ

Rescatux Rescue Disk
If your computer fails to boot, you need a helping hand. Rescatux combines proven repair and rescue tools. (more)

Mozilla Releases a New Firefox
New public key pinning feature will help prevent man-in-the-middle attacks. (more)

Easy Alternative for iPhone and iPad Users

Linux Pro Magazine is now available in Apple Newsstand. Download a free issue, or buy a subscription to carry anywhere.

ADMIN

Linux Pro Magazine

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

- ADMIN HPC
- ADMIN Update
- Linux Update
- Raspberry Pi

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

Admin and HPC:	www.admin-magazine.com/newsletter
Linux Update:	www.linuxpromagazine.com/mc/subscribe
Raspberry Pi:	www.raspberry-pi-geek.com/mc/subscribe

©Yuri Arcurs, 123RF.com

A look at Microsoft's love for Linux

New Love

Open Source means more than just giving out a few pieces of code to an Open Source project; maddog explains. *By Jon "maddog" Hall*

For months, the technical news media has been telling us how much Microsoft is "Open." They point to announcements from Microsoft about how they are cooperating with The Linux Foundation on putting code into the Linux kernel that will allow Linux hypervisors to support Microsoft virtual environments better. They point to the number of patches they have contributed to the Linux kernel and how they work with Canonical to put "Linux" functionality on top of Windows 10.

Companies that were almost crushed by Microsoft in the early days are now partnering with the Redmond giant. People whom I have known for years and are otherwise well respected in the "Open Source" community take these crumbs of code and crow that Microsoft has seen the light.

Nothing could be further from the truth. Microsoft only has one partner: Microsoft. And nothing demonstrates this more than Microsoft's attempts to restrict the browser for Windows 10 to Edge, because Edge will work better with Windows 10 and be integrated better with Cortana, and Edge will integrate better with the rest of Windows 10 applications, or so they say. This stinks of the browser wars of years ago, spawned by Internet Explorer.

Microsoft's "love for Open Source" did not start with the rise to power of their latest CEO. It started years ago when Microsoft started attending and sponsoring Free and Open Source Software (FOSS) events such as OSCON. Tim O'Reilly, the publisher of many books that Open Source people know and love, would invite Microsoft to his events. Microsoft never reciprocated by inviting FOSS people to their conferences to talk about the value of Free and Open Source, however. That might have caused too many Microsoft customers to question why their favorite vendor was not following Open Source best practices, but rather forcing them to go from one disastrous operating system upgrade to another.

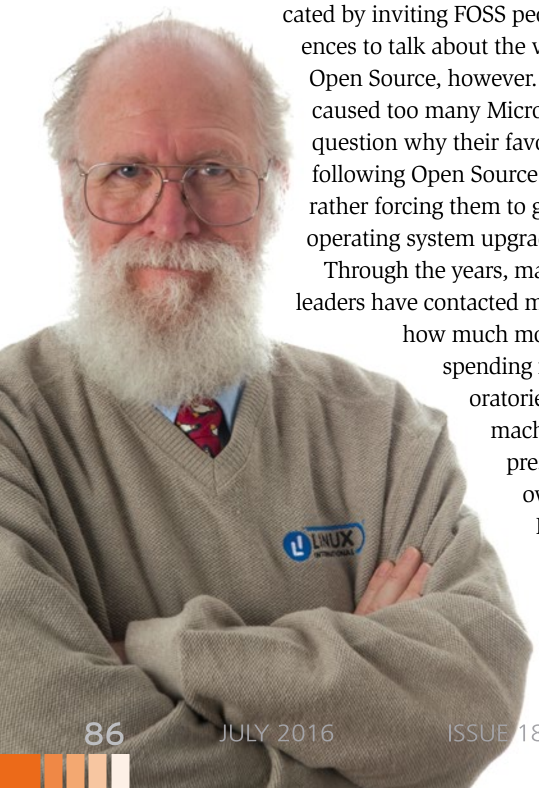

Through the years, many "Open Source" leaders have contacted me and tried to tell me how much money Microsoft was spending in "Open Source" laboratories with people and machines. This did not impress me, because my own company (Digital Equipment Corporation) spent over a billion dollars a year with equipment and engineers to turn out a fine Unix operating system, only to show that their main goal in life was to sell OpenVMS and Windows NT. When you have a lot of money, you can afford to have some dalliances. And, you might even learn a thing or two for your own products, but it does not show an understanding of Open Source.

Openness also has to do with your business tactics, and that is another place where Microsoft really fails. For years, they have been quietly approaching licensees of Android and demanding royalties on patents that they claim are their own inventions. Part of the settlement for these "patent infringements" is that the victim cannot make public which pieces of code that Microsoft claims are infringing. Friends of this practice (and of Microsoft) state that Microsoft has a right to be compensated for their work through the patents. With this point I have no argument. However, one of the traditional balances of patent law is that if some other non-patented way can be determined to do the same thing, the victim of the patent suit should be able to choose that different path. Because Microsoft has never approached the Linux community as a whole, the developers of the Linux kernel (including Android) have no knowledge of which pieces of code would have to be replaced.

Other pundits of Open Source (such as my friend Simon Phipps) have pointed out that Microsoft should join the Open Invention Network (OIN), which allows members to use patents for self-defense or to get patent royalties from closed source companies who do not share the values of Open Source. OIN simply says (and I paraphrase here) "if the code that is violating the patent is part of an open source project, then you cannot sue them or the end user for patent infringement, and they (in turn) cannot sue you for patent infringement on your open source code."

Even if the CEO of Microsoft loves Open Source and secretly uses GNU/Linux in his home office, what really spells love for Open Source is embracing it for your customer's needs and uses. My observation of Microsoft sales and management teams in various countries shows that there is much to be developed in good Open Business Practices.

"Open Source" involves more than just giving out a few pieces of code to an Open Source project or deciding that your closed source product (Microsoft SQL) will sell more copies on top of GNU/Linux than it will just running on your own platforms. ■■■

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®.

A Webzine for High-Performance Computing Specialists

Most Visited
Getting Started
Latest Headlines
Home – ADMIN | The resource for
Home – ADMIN | The resource for all system administrators
http://www.admin-magazine.com/
Google
ADMIN
NEW DAY
HPC
Explore the new world of system administration!
ADMIN
Pen Testing
Home
Articles
News
ADMIN
Shop
About Us
SEARCH
Sign up for our Newsletter
Email Address
Subscribe

Data Recovery
4000 nodes can turn out lots of data. Sure can you can store it, but can you provide the level of fault tolerance admins expect with more conventional data processing environments? We'll show you how to ensure your HPC data will survive a crash -- and how to get it back.

Recent Highlights
- AMD's new power architecture
- HPC in the clouds
- How raw is your raw data?
- Replacing a node without shutting down
- Tune up your cluster

Whitepapers
- Overclocking and HPC
- 10 tips for longer CPU life
- Building a more secure cluster
- Increasing uptime
- Saving money with heartbeat

HPC Resources
- Video: SGI's unveils a green
- Introduction to high perform
- NSA high performance
- Video: AMD and the

Shop
Subscribe to ADMIN Magazine

What's New
- Keeping your clusters on budget
- Building a better compute node
- Red Hat launches a $20 million HPC center
- High performance coding techniques
- Investing in InfiniBand
- HPC overload: Avoiding staff burnout

BLOGS
- Douglas Eadline
- HPC Destinations

VENDORS
- AMD
- CRAY
- DELL
- SGI

VIDEOS
- SGI's Goes Green
- AMD announces
- Bulldozer

ARTICLES
WHITEPAPERS
WRITE FOR US
ABOUT US

If you use Linux... ...you need Linux Magazine!

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

http://hpc.admin-magazine.com

ADMIN
Network & Security

An Interview with the author of *Designing with LibreOffice*

Writing with Style

By Rita L. Sooby

Bruce Byfield has been the writer of the Command Line column for *Linux Pro Magazine* (LPM) since mid 2009. In addition to his "Off the Beat" blog he writes for us at the LPM website [1], he keeps a personal "Off the Wall" blog [2], where he addresses art, specifically Pacific Northwest art and First Nations artists, writing, feminism, and life in general.

Bruce recently published *Designing with LibreOffice*, a book in which he guides the reader through the style, template, typography, and design tools available in LibreOffice to create precise and visually comprehensible documents. He says, "By taking advantage of styles and templates, you can concentrate on self-expression, rather than format" [3].

LibreOffice [4] is a popular free office suite that includes word processing, spreadsheet, multimedia, drawing, database, and math editing applications. This set of programs is so feature rich, that most users take advantage of a very small subset of its capabilities.

Writing a book that helps writers grasp the software's vast functionality is not a venture many would be willing to attempt. Bruce describes how he came to the project:

> "I was at OSCON in 2000 at which Sun Microsystems released the code that would become OpenOffice.org, and later LibreOffice. I knew right then that I would like to write about the code, but I didn't know how to manage or structure such a large project. I made four tries with one editor and finally admitted that I didn't know what I was doing. I was left with 1,400 unpublishable pages, which I cannibalized for short articles for years.
>
> "But I always wanted to write that book. Finally, talking about styles and designs gave me the unifying topic that would allow me to shape the book and contain it to a decent length, so it's a very heady feeling to see it in print after three years of work."

To find out more about his book, I asked Bruce a few questions.

LPM: *Designing with LibreOffice* came out in March 2016. Were you commissioned to write this book, or was it an independent project?

Bruce Byfield: A bit of both. Jean Hollis Weber, who until March 2016 was managing the LibreOffice documentation volunteers, asked me to write a book about styles and templates. I went a little crazy and went far beyond that original topic. Luckily for me, she thought what I was doing was worth encouraging.

LPM: As a professional writer, how challenging was this project? Did you run into any problems, or was it smooth sailing?

Lead Image © Author, 123RF.com

BB: I soon learned that writing a 90,000-word book was vastly different than writing the 1,200-word articles I usually write. I had to do much more planning, and finding a suitable structure for talking about both LibreOffice and typography took an unexpected amount of revision.

However, the biggest problem is eliminating typos and other errors. I copyedited it, and so did Jean. We even got Lee Schlesinger, my former editor at Linux.com, to read the final manuscript so we could get a fresh perspective. Yet all three of us still missed a lot. I swear the typos breed in the night after a file is saved.

LPM: You released your book under a Creative Commons Attribution Sharealike License [5], sometimes also known as "copyleft." Share with us what this means and why you decided to publish your work in this way.

BB: The license means that anyone can copy, share, and revise the contents as they please, so long as they give me credit and release any work that changes or borrows the content under the same license.

I have been writing about free software and free licenses for almost 15 years, so I would be a hypocrite if I published any other way.

Moreover, free software has shaped a good part of my life. I have sometimes criticized its shortcomings, especially the lack of diversity, but the truth is that free software gave me a sense of direction when I badly needed one and has allowed me over the years to hobnob with brilliant and talented people. I thought I was overdue to give something back, and, having criticized for so long as a writer, to make myself a target for once.

Also, I was curious. Would having free downloads affect sales of the hard copy book? What would the ratio of download to sales be? Since Jean gave me an advance, I am able to explore such questions.

Mostly, however, offering a free book is its own reward. Like volunteer coders, much of my immediate pay for the book is in credit, and since the book was released, I've received an embarrassing amount of that. The enthusiasm from others for what I wrote is unexpected, but very welcome.

LPM: The title alerts the reader that this is not a how-to book for LibreOffice. Rather, you address design or, more specifically, typography. What clues did you see that indicated a need for this kind of book?

BB: Thanks mainly to Jean, LibreOffice has some of the best documentation in free software. However, the limitation of any manual is that, while it explains everything you can do, it doesn't explain when or why you should choose different types of formatting. In a large application like LibreOffice, that means that people have a hard time taking full advantage of the features.

Also, people often talk about LibreOffice as if it were just another office suite. However, Writer is not just a word processor, but an intermediate desktop publisher as well. Users can do far more in Writer than they ever could imagine in Microsoft Word. When I was a technical writer, I did several manuals of well over 500 pages in OpenOffice.org, its predecessor. I'm not sure that would be even possible in Word, but if it is, the experience would be nightmarish. I wanted more people to know what the capabilities were.

Another clue was that people often think of typography as design that calls attention to itself, with pages full of garish colors and decorative fonts formatted into illegibility. That was understandable when the personal computer first came in, and everyone wanted to experiment with all the new tools that were available, but almost forty years later, we need to get beyond that type of excess. I wanted to stress that typography is about design that works quietly in the background.

In addition, I wanted to explain that part of effective design is features that allow easy editing and revision over a long period of time. That's one reason why styles are so important. Similarly, the humble *Hide* feature in Paragraph and Character styles allows you to maintain two different versions of a document in one file, which makes keeping the versions in sync over several versions much easier.

LPM: In the Introduction you mention that, for the most part, typography should be "hidden." What do you mean?

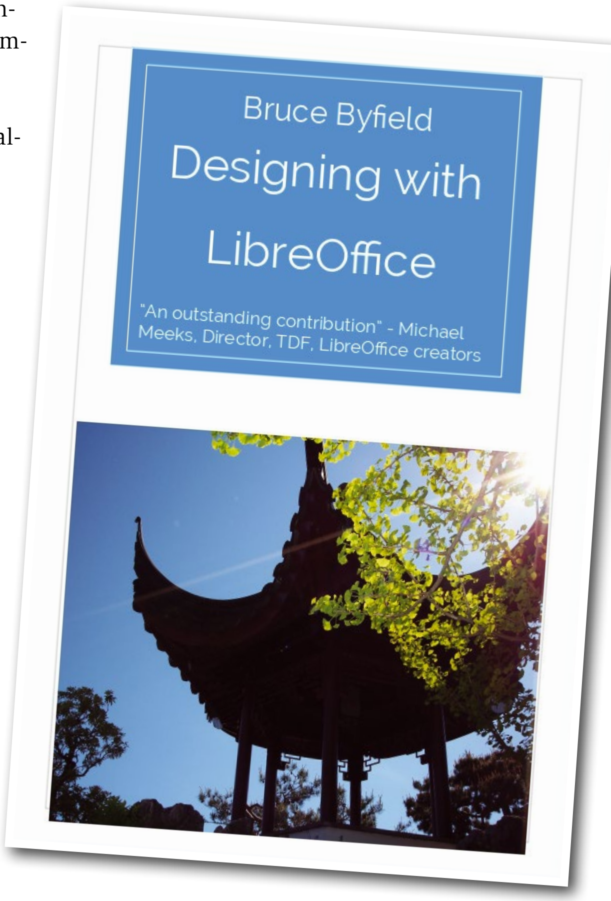
Bruce Byfield
Designing with LibreOffice
"An outstanding contribution" - Michael Meeks, Director, TDF, LibreOffice creators

BB: People sometimes ask why I used the Sun Yat-Sen Garden in Vancouver, Canada, as a motif. If I'm feeling flippant, I reply that if O'Reilly can use covers unassociated with the subject matter of its books, then so can I.

However, the real reason is the analogy between typography and the philosophy of feng shui on which Chinese Classical gardens are based. Everything in a Classical garden is meticulously planned – the positioning of rocks and tiles, the angle of corridors, the contrast between plants and trees, and absolutely everything else. In the Imperial Gardens that the English burned, even the exact position of ornaments on shelves was determined. Yet the end result is supposed to look absolutely natural, so that anyone strolling through the garden doesn't notice how carefully planned everything is. I thought that a good analogy for what typography is supposed to do, so much so that I personally paid for permission to use photos of the Garden.

LPM: You wrote an article in 2004 called "Replacing FrameMaker with OOo Writer," which implied that you could use LibreOffice's predecessor OpenOffice.org for desktop publishing (DTP). What elements did OpenOffice have then that made it suitable for DTP, and what additions have been made to LibreOffice since then that improve on this capability?

BB: Right from the days when the code was part of StarDivision, it has always had formatting features that are far ahead of other word processors. Writer allows control and precision that is more characteristic of DTP than of word processors.

The story I heard from a Sun Microsystems employee years ago is that the original programmers for StarDivision, the proprietary office suite that became OpenOffice.org then later LibreOffice, were told that they would have to use what they wrote for their own documentation. I'm not sure whether the story is true, but it would explain why Writer is so sophisticated compared to most word processors.

At any rate, very few changes have been needed. However, LibreOffice has cleaned up the code and made the features available for different types of styles more consistent. A major headache was removed by including an option to embed fonts, which means you can share documents without worrying whether the recipient has the fonts installed that you use.

I should also mention the Typography toolbar extension, which makes advanced features easier to apply. In effect, it changes Writer from an intermediate desktop publisher to an advanced one.

LPM: Creating esthetically pleasing documents is often considered an art practiced by specially trained graphic artists. With the help of your book, how challenging do you think it will be for the average reader to understand and apply the concepts you present?

BB: Design takes practice, of course. What I've tried to do is give the background so that users can practice on their own.

I suspect that many users will probably dip into the sections they want. However, if they read the book from the beginning, it shows how, by beginning by choosing fonts and their line spacing, you can have a framework for many of your other design choices. With this approach, users do not have to rely on vague impressions or random choices. Instead, design becomes more of a science, and less of an art. I hope that this approach will demystify design and give users more confidence in setting up their own designs.

LPM: Share any other thoughts you might have, and tell us where we can get your book.

BB: The book's reception has exceeded all expectations. I thought it might receive a few thousand downloads in the first year. Instead, in six weeks, it had over 13,500 downloads, with another 8,500 to be distributed on a DVD that comes with a magazine, and offers for French, German, Spanish, and Chinese translations. I feel like I've fallen down a rabbit hole, and any minute now I'm going to see Alice and a white rabbit hurrying by.

The next step is to release parts of the book in smaller volumes, for those who only want part of its information. This step is also a good way of revising the book, which I otherwise would feel too overwhelmed to attempt.

You can download the book at the book's website [3]. The *Download/Buy* tab also has a link to Lulu.com, where you can order a hard copy of the book. In whichever medium readers see the book, I hope that readers find it educational and useful. ■■■

INFO

[1] Off the Beat: *http://www.linuxpromagazine.com/Online/Blogs/Off-the-Beat-Bruce-Byfield-s-Blog*

[2] Off the Wall: *https://brucebyfield.com*

[3] Designing with LibreOffice: *http://designingwithlibreoffice.com*

[4] LibreOffice: *http://www.libreoffice.org*

[5] CC BY-SA 3.0: *http://creativecommons.org/licenses/by-sa/3.0/*

Zack's Kernel News

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

By Zack Brown

Random Number Generation on Modern Systems

Stephan Müller recently pointed out that `/dev/random` has been showing signs of age relative to modern environments like embedded systems, solid-state drives, massively parallel systems, and virtualized systems. The problem is how to identify good sources of entropy on all systems, so that `/dev/random` really does produce random numbers that are equally random across all environments.

Stephan's approach, LRNG (Linux Random Number Generator), seeks to solve that problem and especially to provide proper entropy sources during boot time. He also wanted LRNG to have a lower performance effect on parallel systems and allow accelerated cryptographic primitives. Crypto primitives are simple, reliable tools that are used as building blocks of larger scale security systems. Massively parallel systems have to implement security protocols on all nodes, and having good cryptographic speed can benefit that.

Stephan gave a link to a scholarly article he'd written that described his approach [1]. Beyond the technical details, Stephan chose to release his design under a dual license – either the GPL (version number unspecified) or a more BSD-ish license that allowed closed-source binary distribution.

In terms of implementation, Stephan explained, "The patches do not replace or even alter the legacy /dev/random implementation but allows the user to enable the LRNG at compile time. If it is enabled, the legacy /dev/random implementation is not compiled. On the other hand, if the LRNG support is disabled, the legacy /dev/random code is compiled unchanged. With this approach you see that the LRNG is API and ABI compatible with the legacy implementation."

Nikos Mavrogiannopoulos read the PDF and noticed that in both the traditional `/dev/random` implementation and Stephan's LRNG implementation, the random number generator would be "minimally" seeded 112^6 bits of entropy. Nikos said, "Unfortunately one of the issues of the /dev/urandom interface is the fact that it may start providing random numbers even before the seeding is complete." And based on the text, Nikos concluded that LRNG suffered from this problem as well. He said, "That's a serious limitation […], since most/all newly deployed systems from 'cloud' images generate keys using /dev/urandom (for sshd for example) on boot, and it is unknown to these applications whether they operate with uninitialized seed."

Nikos liked the rest of Stephan's implementation, but he felt that if `/dev/random` was going to be replaced, any new implementation should ensure "that the kernel seed buffer is fully seeded prior to switching to userspace."

Stephan reassured Nikos that the `getrandom()` system call would block until the appropriate amount of seed data had been obtained; after which, he said `getrandom()` would behave like `/dev/urandom`. Alternatively, he said, "you may use the /proc/sys/kernel/random/drbg_minimally_seeded or drbg_fully_seeded booleans. If you poll on those, you will obtain the indication whether the secondary DRBG feeding /dev/random is seeded with 112 bits (drbg_minimally_seeded) or 256 bits (drbg_fully_seeded). Those two booleans are exported for exactly that purpose: allow user space to know about initial seeding status of the LRNG."

Nikos pointed out that user code would need to have some way to tell whether `getrandom()` existed on a given system. "Today," he said, "due to libc not having the call, we can only use /dev/urandom and applications would most likely continue to do so long time after getrandom() is introduced to libc."

Stephan explained:

> *Implement the syscall yourself with syscall(). If you get ENOSYS back, revert to your old logic of seeding from /dev/urandom.*
>
> *If you know you are on kernels > = 3.14, you could use the following steps in your library:*
>
> *1) poll /proc/sys/kernel/random/entropy_avail in spaces of, say, one second and block your seeding process until that value becomes non-zero 2) if you unblock, seed from /dev/urandom and you have the guarantee of having a /dev/urandom seeded with 128 bits.*

Nikos didn't like that explanation at all. He replied, "That's far from a solution and I wouldn't recommend to anyone doing that. We cannot expect each and every program to do glibc's job. The purpose of a system call like getrandom is to simplify the complex use of /dev/urandom and eliminate it, not to make code handling randomness in applications even worse."

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.**

Theodore Ts'o replied, "Yes, but if glibc is falling down on the job and refusing to export the system call (I think for political reasons; it's a Linux-only interface, so Hurd wouldn't have it), then the only solution is to either use syscall directly (it's not hard for getrandom, since we're not using 64-bit arguments which gets tricky for some architectures), or as Peter Anvin has suggested, maybe kernel developers will have to start releasing the liblinux library, and then teaching application authors to add -linux to their linker lines."

But Nikos felt that the "political" issue was significant. If the system call wasn't available on a given system, "they have an almost impossible task to simulate getrandom() on kernels which do not support it. One may agree with their concerns, but the end result is that we have not available that system call at all, several years after it is there."

Ted rejoined, "The whole *point* of creating the getrandom(2) system call is that it can't be simulated/emulated in userspace. If it can be, then there's no reason why the system call should exist." He suggested a range of technical implementation possibilities. Or, he said, "you can let the application author specify some kind of 'I want to run in insecure mode', via some magic glibc setting. You could probably default this to 'true' without a huge net reduction of security, because most application authors weren't getting this right anyway."

Elsewhere, Ted had his own objections to Stephan's code. He pointed out that some of the entropy sources might not contain true entropy and would therefore lead to insecure random number generation. Stephan replied that any individual source of entropy, such as the "jitter" source Ted mentioned, could be removed to satisfy Ted's concerns.

Sandy Harris spoke out in favor of the jitter source, saying, "Jitter, havege and my maxwell(8) all claim to get entropy from variations in timing of simple calculations, and the docs for all three give arguments that there really is some entropy there." He gave a link to a PDF discussing the issue [2].

Pavel Machek also had some objections. He noticed that Stephan's code seemed to be dependent on the hardware having a high-resolution clock on board. He asked, "What goes on if high resolution timer is not available?" Stephan replied, "If there is no high-resolution timer, the LRNG will not produce good entropic random numbers." He listed the 14 architectures for which the Linux kernel did not implement a high-resolution timer and pointed out that none of those were large-scale architectures. He said, "Please note that also the legacy /dev/random will have hard time to obtain entropy for these environments. The majority of the entropy comes from high-resolution time stamps. If you do not have them and you rely on Jiffies, an attacker has the ability to predict the events mixed into the pools with a high accuracy. Please remember the outcry when MIPS was identified to have no get_cycles about two or three years back."

But Stephan added, "the patch I offer leaves the legacy /dev/random in peace for those architectures to not touch the status quo." Pavel replied, "… that's the major problem – right? Makes it tricky to tell what changed, and we had two RNGs to maintain." Stephan said:

I would rather think that even the legacy /dev/random should not return any values in those environments. The random numbers that are returned on these systems are bogus, considering that the only noise

source that could deliver some entropy excluding timestamps (if you trust the user) are the HID event values. And for those listed systems, I doubt very much that they are used in a desktop environment where you have a console.

If everybody agrees, I can surely add some logic to make the LRNG working on those systems. But those additions cannot be subjected to a thorough entropy analysis. Yet I feel that this is wrong.

My goal with the LRNG is to provide a new design using proven techniques that is forward looking. I am aware that the design does not work in circumstances where the high-res timer is not present. But do we have to settle on the least common denominator knowing that this one will not really work to begin with?

By the end of the discussion, most of the objections to Stephan's code seemed on track to finding decent solutions or workarounds. This is one of those situations where an older implementation of a kernel feature just isn't cutting it anymore because the industry has moved in directions that hadn't been predicted (massively parallel systems, etc.), and so the code needs to be updated to do the best it can to support what exists in the world. Because of that, even if Stephan's code ends up having missing pieces and other remaining problems, it's likely to still go into the kernel in one form or another, just as an improvement over what was there before. After that, future patches would continue to address the remaining problems where possible.

Randomizing Memory Locations to Secure Against Attack

Thomas Garnier implemented ASLR (Address Space Layout Randomization) for kernel memory on x86-64 systems. ASLR is used to prevent attackers from writing security exploits based on a known location of code in memory. A weak form of ASLR has existed in the Linux kernel since 2005 and has been supplemented by various patch sets for use in security-oriented Linux distributions ever since. Thomas wanted to bring proper ASLR to the main tree itself. Thomas explained, "This security feature mitigates exploits relying on predictable kernel addresses. These addresses can be used to disclose the kernel modules' base addresses or corrupt specific structures to elevate privileges." He went on, "Knowing the base address and physical memory size, an attacker can deduce the PDE virtual address for the vDSO memory page. This attack was demonstrated at CanSecWest 2016, in the 'Getting Physical Extreme Abuse of Intel Based Paged Systems' [3] (see second part of the presentation). Similar research was done at Google leading to this patch proposal. Variants exists to overwrite /proc or /sys objects' ACLs leading to elevation of privileges."

To implement his solution, he explained, "Entropy is generated using the KASLR early boot functions now shared in the lib directory (originally written by Kees Cook). Randomization is done on PGD & PUD page table levels to increase possible addresses. The physical memory mapping code was adapted to support PUD level virtual addresses. An additional low memory page is used to ensure each CPU can start with a PGD aligned virtual address (for realmode)."

There was no significant debate on the mailing list. H. Peter Anvin and others had some minor technical issues and bugs to report against Thomas's patch, but no one expressed any doubts about adding the feature itself.

Security has always been a central element of Linux development, but it has never received the amount of testing it's gotten in recent years. In the old days, the biggest threats were from spammers wanting to set up their own botnets or individual hackers looking for thrills, and for many years the more tempting target of such attacks would be Windows machines. Nowadays the United States, China, Russia, and many other countries devote significant resources to cyber warfare, and Linux presents a very tempting target because it is essentially the back end for every significant service on the Internet. The Linux developers are having to shore up security features that were not necessarily tended very carefully for many years.

Eventually, the tremendous focus on world-wide cyber warfare will result in a much stronger and more secure Linux kernel in all respects. For now, the developers are having to play catch-up. Ultimately, the pace of kernel development will always leave it susceptible to new vectors of attack, but hopefully within a few years most existing attack vectors will be nailed down.

Tracking Removable Devices

Wade Mealing posted some patches to implement a new "audit subsystem," that would log when devices were added to or removed from a running system. Along with the subsystem, he included a set of user tools to sift through the audit logs and track specific devices or events. For his initial implementation, he included support for USB devices only, although he hoped to extend that to other subsystems as well.

Oliver Neukum felt that the project might not be worth it and should at least be publicly debated before anything serious was implemented. In terms of specific implementation, he suggested that Wade stick to generic functions rather than being quite so USB specific.

Bjørn Mork agreed that the project needed a public debate. Specifically, he pointed out that there had already been earlier discussions, with different conclusions from Wade's proposal. He said:

Greg has already asked the obvious questions and made the obvious 'do this in userspace using the existing uevents' proposal. I did not see any followup to his last message, so I assumed this audit thing would return to the drawing board with a userspace implementation [4].

It was quite surprising to instead see a USB specific kernel implementation duplicating existing device add/remove functionality. Why? The provided reason makes absolutely no sense at all. Userspace tools are as intelligent as you make them. And 'decoded, filtered or ignored' implies policy, which IMHO has no place in the kernel in any case.

Bjørn concluded, "I think the generic layer implementation is already there. The proposed USB specific solution adds nothing, as pointed out by Greg the last time this was discussed."

Greg Kroah-Hartman joined in the chorus of implementing Wade's patches in userspace and to catch all device types rather than just USB.

Steve Grubb, however, spoke out partially in favor of a kernel-based implementation, saying, "The audit system has to do everything possible to make sure that an event is captured and logged. Does the uevent netlink protocol ever drop events because the user space queue is full? If the uevent interface drops events, then it's not audit quality in terms of doing everything possible to prevent the loss of a record. If this were to happen, how would userspace find out when a uevent gets dropped? I may have to panic the machine if that happens depending on the configured policy. So, we need to know when it happens. If on the other hand it doesn't ever drop events, then it might be usable."

Paul Moore supported Steve's statements, saying:

Audit has some odd requirements placed on it by some of its users. I think most notable in this particular case is the need to take specific actions, including panicking the system, when audit records can't be sent to userspace and are 'lost'. Granted, it's an odd requirement, definitely not the norm/default configuration, but supporting weird stuff like this has allowed Linux to be used on some pretty interesting systems that wouldn't have been possible otherwise. Looking quickly at some of the kobject/uevent code, it doesn't appear that the uevent/netlink channel has this capability.

It also just noticed that it looks like userspace can send fake uevent messages; I haven't looked at it closely enough yet, but that may be a concern for users which restrict/subdivide root using a LSM … although it is possible that the LSM policy could help here. I'm thinking aloud a bit right now, but for SELinux the netlink controls aren't very granular and sysfs can be tricky so I can't say for certain about blocking fake events from user space using LSMs/SELinux.

Greg said he'd never seen uevent drop an event in 10 years of watching. He asked (several separate times, as it turned out) what the use case for Wade's code really was. Wade replied:

The goal of these message is to let a system administrator see in the audit logs, that a device has been plugged in and the basic details about this. Having this only in user space means that (and Greg alludes to this) that this will be for human eyes only and not be machine usable in the kernels. Without it being in kernel, it can't be extended for manipulation by auditctl at some point in the future.

Specifically I am trying to create a well formed audit trail when devices are added or removed from the system by the user space audit tools. The implementation at the moment does not do any filtering, but rather creates the raw audit events.

In some ways this is similar to a decorated class in say java. In this case the class is unaware it is being decorated yet we can monitor what is happening in that class without polluting the class code with messy log or trace information.

I don't see either kernel or user-space applications create add or remove events in the audit subsystem. I understand that some events are placed into uevents (To be intercepted by udevd), while this also exports the same information it is not in the audit subsystem in kernel.

Burn Alting also offered his own list of abilities that he hoped would be provided by Wade's code:

– when was a (possible) removable media device plugged into a system and what were the device details – perhaps my corporation has a policy on what devices are 'official' and hence one looks for alternatives, and/or,

– was it there at boot? (in case someone adds and removes such devices when powered off), and eventually

– has an open for write (or other system calls) occurred on designated removable media? (i.e. what may have been written to removable media – cooked or raw) – Yes, this infers a baseline of what's connected or an efficient means of working out if a device is 'removable' at system call time.

In essence, I need to know if and how removable media is being used on my systems. The definition of 'removable' is challenging, but my idea would be for one to be able to define it via the auditd interface.

Clearly, Burn's purpose would be to implement security features. In a later post he acknowledged that a determined hacker could get past these audits, but that it was important to thwart the less skilled hackers.

Greg didn't respond directly to these feature desires, but later he seemed to back off from his opposition to them, saying, "It's not an easy problem, good luck all!"

That seems to be far from the last word, however. The feature seems shrouded in controversy, not least of which will inevitably be: If the kernel can't stop a determined hacker, then don't such features just amount to code clutter? For now, it seems that the immediate objections have been withdrawn, and a kernel-based audit trail is still in the works. ■■■

INFO

[1] "Linux Random Number Generator – A new approach" by Stephan Müller: *http://www.chronox.de/lrng/doc/lrng.pdf*

[2] "The maxwell(8) random number generator" by Sandy Harris: *ftp://ftp.cs.sjtu.edu.cn:990/sandy/maxwell/Maxwell.pdf*

[3] "Getting physical: Extreme abuse of Intel based paging systems" by Nicolas A. Economou and Enrique E. Nissim, *https://goo.gl/ANpWdV*

[4] Extending usb to do device auditing: *http://www.spinics.net/lists/linux-usb/msg137671.html*

FEATURED EVENTS

Users, developers, and vendors meet at Linux events around the world. We at *Linux Magazine* are proud to sponsor the Featured Events shown here.

For other events near you, check our extensive events calendar online at *http://linux-magazine.com/events.*

If you know of another Linux event you would like us to add to our calendar, please send a message with all the details to *events@linux-magazine.com*.

USENIX ATC '16

Date: June 22-24, 2016

Location: Denver, CO

Website: *https://www.usenix.org/conference/atc16*

Leading systems researchers gather to gain insight into virtualization, system and network management and troubleshooting, cloud computing, security, and more with refereed papers, industry talks, a poster session, and BoF sessions.

GUADEC 2016

Date: August 12–14, 2016

Location: Karlsruhe, Germany

Website: *https://2016.guadec.org*

Gnome users, developers, foundation leaders, individuals, and others come together to meet collaborators from chat rooms and mailing lists, to network, to visit old friends and make new ones, and to have fun at events, workshops, BoFs, and hackfests.

ContainerCon North America '16

Date: August 22–24, 2016

Location: Toronto, Ontario, Canada

Website: *http://events.linuxfoundation.org/events/containercon*

Learn how to automate, deploy, and scale workloads using container technologies: From hardware virtualization to storage and software-defined networking, containers are driving "cloud native."

EVENTS

EclipseCon France	June 7-9	Toulouse, France	http://www.linuxpromagazine.com/Resources/Event-Calendar#event_68229
Pi and More 9	June 11	Trier, Germany	http://piandmore.de/en
Tübix	June 11	Tübingen, Germany	http://www.tuebix.org/
SLAC 2016	June15-17	Berlin, Germany	https://www.heinlein-support.de/secure-linux-administration-conference
ISC High Performance	June 19–23	Fankfurt, Germany	http://www.isc-hpc.com/
Deutsche OpenStack Tage	June 21-22	Cologne, Germany	https://openstack-tage.de/
2016 USENIX Technical Conf.	June 22-24	Denver, Colorado	https://www.usenix.org/conference/atc16
Maker Faire Kansas City	June 25-26	Kansas City, Missouri	http://www.makerfairekc.com/
Debconf	July 3-9	Cape Town, South Africa	https://wiki.debconf.org/wiki/DebConf16
Texas Linux Fest	July 8-9	Austin, Texas	http://2016.texaslinuxfest.org/
Tech Open Air 2016	July 13-15	Berlin, Germany	http://toa.berlin/
GUADEC 2016	August 12-14	Karlsruhe, Germany	https://2016.guadec.org/
LinuxCon North America	August 22-24	Toronto, ON, Canada	http://events.linuxfoundation.org/events/linuxcon-north-america
ContainerCon North America '16	August 22-24	Toronto, ON, Canada	http://events.linuxfoundation.org/events/containercon
IFA	Sept. 2-7	Berlin, Germany	http://www.ifa-berlin.de/
systemd.conf2016	Sept. 28 - Oct. 1	Berlin, Germany	https://conf.systemd.io/
LinuxCon Europe	October 4-6	Berlin, Germany	http://events.linuxfoundation.org/events/linuxcon-europe
ContainerCon Europe	October 4-6	Berlin, Germany	http://events.linuxfoundation.org/events/containercon-europe

Images © Alex White, 123RF.com

CALL FOR PAPERS

We are always looking for good articles on Linux and the tools of the Linux environment. Although we will consider any topic, the following themes are of special interest:

- System administration
- Useful tips and tools
- Security, both news and techniques
- Product reviews, especially from real-world experience
- Community news and projects

If you have an idea, send a proposal with an outline, an estimate of the length, a description of your background, and contact information to *edit@linux-magazine.com*.

The technical level of the article should be consistent with what you normally read in *Linux Magazine*. Remember that *Linux Magazine* is read in many countries, and your article may be translated into one of our sister publications. Therefore, it is best to avoid using slang and idioms that might not be understood by all readers.

Be careful when referring to dates or events in the future. Many weeks could pass between your manuscript submission and the final copy reaching the reader's hands. When submitting proposals or manuscripts, please use a subject line in your email message that helps us identify your message as an article proposal. Screenshots and other supporting materials are always welcome.

Additional information is available at:
http://www.linux-magazine.com/contact/write_for_us.

AUTHORS

Swapnil Bhartiya	8
Zack Brown	92
Bruce Byfield	44, 76
Joe Casad	3, 8
Giuseppe Dispinzeri	12
Thomas Gronenwald	12
Jon "maddog" Hall	86
Valentin Höbel	56
Frank Hofmann	34
Thomas Joos	22
Michael Klassen	12
Klaus Knopper	38
Peter Kreußel	26
Jens Kubieziel	40
Charly Kühnast	48
Dmitri Popov	72
Andreas Reitmaier	80
Mike Schilli	62
Tim Schürmann	50
Rita L Sooby	88
Ferdinand Thommes	68
Uwe Vollbracht	18

CONTACT INFO

Editor in Chief
Joe Casad, jcasad@linux-magazine.com

Managing Editor
Rita L Sooby, rsooby@linux-magazine.com

Localization & Translation
Ian Travis

News Editor
Swapnil Bhartiya

Copy Editor
Amber Ankerholz

Layout
Dena Friesen, Lori White

Cover Design
Dena Friesen, Lori White

Cover Image
© Stuart Miles, 123RF.com

Advertising – North America
Ann Jesse, ajesse@linuxnewmedia.com
phone +1 785 841 8834

Advertising – Europe
Penny Wilby, pwilby@sparkhausmedia.com
phone +44 1807 211100

Publisher
Brian Osborn, bosborn@linuxnewmedia.com

Marketing Communications
Gwen Clark, glark@linuxnewmedia.com
Linux New Media USA, LLC
616 Kentucky St.
Lawrence, KS 66044 USA

Customer Service / Subscription
For USA and Canada:
Email: cs@linuxpromagazine.com
Phone: 1-866-247-2802
(Toll Free from the US and Canada)
Fax: 1-785-856-3084

For all other countries:
Email: subs@linux-magazine.com
Phone: +49 89 99 34 1167
Fax: +49 89 99 34 1198

www.linuxpromagazine.com – North America

www.linux-magazine.com – Worldwide

While every care has been taken in the content of the magazine, the publishers cannot be held responsible for the accuracy of the information contained within it or any consequences arising from the use of it. The use of the disc provided with the magazine or any material provided on it is at your own risk.

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.

Approximate

UK / Europe	Jul 02
USA / Canada	Jul 29
Australia	Aug 29

On Sale Date

Issue 189/ August 2016

Filesystems

In Linux (like Unix) everything is a file, so the filesystem has special importance. Next month we look at some popular filesystems for Linux, including ZFS, Btrfs, and ext 2/3/4.

Lead Image © Womue, fotolia.com

25th USENIX Security Symposium

AUGUST 10–12, 2016 • AUSTIN, TX

The USENIX Security Symposium brings together researchers, practitioners, system administrators, system programmers, and others interested in the latest advances in the security and privacy of computer systems and networks. The Symposium will span three days, with a technical program including refereed papers, invited talks, panel discussions, posters, a Work-in-Progress session, Doctoral Colloquium, and Birds-of-a-Feather sessions (BoFs).

The following co-located events will occur before the Symposium:

WOOT '16: 10th USENIX Workshop on Offensive Technologies, August 8–10

CSET '16: 9th Workshop on Cyber Security Experimentation and Test, August 8

FOCI '16: 6th USENIX Workshop on Free and Open Communications on the Internet, August 8

ASE '16: 2016 USENIX Workshop on Advances in Security Education, August 9

HotSec '16: 2016 USENIX Summit on Hot Topics in Security, August 9

Register by July 18 and save.

www.usenix.org/sec16

usenix
THE ADVANCED
COMPUTING SYSTEMS
ASSOCIATION

www.usenix.org/facebook www.usenix.org/gplus www.usenix.org/linkedin

twitter.com/usenix www.usenix.org/youtube

SUPERMICRO®

MicroBlade

Intel® Xeon® Processor E5-2600 v4 Product Family

High Density • High Performance • High Efficiency • Cost-Effective

Enterprise, Data Center, Web Applications, HPC and Cloud Computing Solutions

Industry's Densest Intel® Xeon® DP & UP Blade Solutions

3U MicroBlade

Up to 28 UP/ 14 DP Intel® Xeon® Nodes

6U MicroBlade

Up to 56 UP/ 28 DP Intel® Xeon® or 112 Intel® Atom™ Nodes

MBI-6128R-T2

MBI-6128R-T2X

Intel® Xeon® E5-2600 v4/v3

MBI-6218G-T41X

MBI-6118G-T41X

Intel® Xeon® D-1541

MBI-6219G-T

MBI-6118D-T2/T2H

MBI-6118D-T4/T4H

Intel® Xeon® E3-1200 v5/v4/v3

MBI-6418A-T7H/T5H

Intel® Atom™ C2750/ C2550

MicroBlade

Industry's Densest DP/UP Servers, supports up to:

- 112 Intel® Atom™ processor C2000 nodes per 6U
- 56/28/14 Intel® Xeon® D-1541 UP nodes per 6U/3U
- 56/28/14 Intel® Xeon® E3-1200 v5/v4/v3 UP nodes per 6U/3U
- 28/14 Intel® Xeon® E5-2600 v4/v3 DP nodes per 6U/3U
- High-efficiency Supermicro system design (as low as 10W per node)
- Redundant (N+1 or N+N) Titanium level (96%+) digital power supplies
- Redundant 10/2.5/1GbE SDN switches with 40Gbps/10Gbps uplinks
- Redundant chassis management modules with unified remote management software
- Front-loading nodes for easy access and servicing

intel XEON inside™

TITANIUM We Power the Cloud X10 SERVER PLATFORMS

Learn more at www.supermicro.com