

FREE  **RED HAT TRAINING VIDEOS**
20+ HOURS OF GUIDES | PLUS 20 SOFTWARE PACKAGES & TUTORIAL ASSETS

www.linuxuser.co.uk

Linux User

THE ESSENTIAL MAGAZINE FOR THE GNU GENERATION



& Developer™

POWER UP YOUR PC

TRIPLE BOOT

Run three distros on one computer



WIN
10 CASES FOR YOUR PI 2

50

Pages of guides

- » Network tools » BTRFS setup
- » Penetration testing & more...

Pi car computer



Explore the world using an open source, portable satnav

Secure your Pi



Top tricks and techniques for protecting your pocket PC



MinnowBoard Max

Intel's upgraded development board brings a 64-bit processor



OpenPi

Build your own circuit board now



DigitalEdition
GreatDigitalMags.com
ISSUE 150



FROM THE DIRECTOR OF THE DARK KNIGHT TRILOGY AND INCEPTION

INTERSTELLAR

OWN IT ON DIGITAL HD FROM MARCH 17
OWN IT ON BLU-RAY™ AND DVD FROM MARCH 30



Synology®

Store and Stream movies with Synology NAS
For details on where to buy, please visit www.synology.com



DS214play



DS415play



To celebrate the release of the latest Christopher Nolan epic, Interstellar (on Digital HD, Blu-ray™ and DVD), Synology is giving you the chance to win great prizes so please join us at The Gadget Show Live (8-12 April) or visit our Facebook page



7-12 APRIL 2015 • NEC BIRMINGHAM

Welcome

to issue 150 of Linux User & Developer

Get
Linux User
cheaper
every issue


Page 50

Your team of Linux experts...

Rob Zwetsloot studied aerospace engineering, using Python to model complex simulations. Rapiro returned this month and, emboldened, demanded an upgrade. Rob obliged (p.68) then turned back to his latest project: a secure (p.42), avant-garde (p.38), multibooting (p.16) system for the splendid **LU&D** Flying Fortress.




Jon Masters is a Linux kernel hacker who has been working on Linux for some 19 years, since he first attended university at the age of 13. Jon lives in Cambridge, Massachusetts, and works for a large enterprise Linux vendor. You can find his brilliant Kernel Column on pages 14-15 this month.



Joey Bernard is a true renaissance man, and he splits his time between building furniture, helping researchers with scientific computing problems and writing Android apps. This month in his brilliant Python column (p.72) he takes us through the first part of a three-part guide to programming a Raspberry Pi digital assistant.



Mihalis Tsoukalos is a UNIX system administrator with expertise in programming, databases and maths. He has been using Linux since 1993. Mihalis shows us how to discover essential network information by using `lsof` in this issue (p.26), as well as how to create smart publications in LaTeX (p.30).



Christian Cawley is a former IT and software support engineer, and since 2010 he has written for computer and smartphone users both online and in magazines. In another special four-page guide, this month Christian shows us the best ways to physically and digitally protect a Raspberry Pi from harm (p.64).



Gareth Halfacree is our resident news reporter and brings us the latest developments from all over the open source world, starting on page 8. Also this month, Gareth tests out Intel's new Minnowboard Max to find out how it stacks up against its predecessor, and its ARM-based rival boards.



This issue

- » Triple your PC's potential
- » Build an open source satnav
- » Start using the btrfs file system
- » Troubleshoot your home network



Welcome to the latest issue of **Linux User & Developer**, the UK and America's favourite Linux and open source magazine.

We've packed a lot into these pages for you this month, covering a real range of useful software and techniques. Kicking off with our Triple Boot feature, first we show you how to install three different

distros onto your computer – including Windows, since it's still necessary for many of us from time to time (but we soften the blow with a sneak peak at the upcoming new version) – and then run through some guides to partitioning, maintenance and GRUB styling. Turn to page 16 to get started.

Elsewhere this month, we're penetration testing with Pentoo's impressive suite of tools, learning how to use LaTeX, migrating to the btrfs file system, building a private YouTube and more. We also return to our Qt series, and round up the very best music players for Linux in this month's group test.

On the Raspberry Pi side of things we're taking you through our most ambitious project yet – with the guidance of Pi pro Liam Fraser, you'll be setting up SainSmart's brilliant 9-inch touchscreen and Adafruit's Ultimate GPS kit in order to make an open source satnav, complete with a music player and local weather functionality as well as Navit maps. Have fun!

Gavin Thomas, Deputy Editor

Get in touch with the team:
linuxuser@imagine-publishing.co.uk

Facebook:
Linux User & Developer

Twitter:
@linuxusermag

Buy online
imagineshop.co.uk

Visit us online for more news, opinion, tutorials and reviews:

www.linuxuser.co.uk

Contents



Subscribe & save!
50 Save up to 50% on the shop price. US customers can subscribe via page 76



16 Triple Boot masterclass
Run three different distros on your computer

OpenSource

- 08 News**
The biggest stories from the open source world
- 12 Opinion Columns**
Expert views on open source and free software
- 14 Kernel Column**
The latest on the Linux Kernel with Jon Masters
- 94 Letters**
Your views on the magazine and the open source scene

Features

- 16 Triple Boot**
From partitioning to styling up custom GRUB menus
- 52 Raspberry Pi Car Computer**
Make an open source satnav
- 42 Penetration testing with Pentoo**
Find vulnerabilities in your home network
- 88 Q & A**
Your hardware and software questions answered

Tutorials

- 26 Get network information with lsof**
Troubleshoot your network by extracting the required information
- 30 Write a book in LaTeX**
Typeset, beautify and customise a whole book using LaTeX code
- 34 Host your own media gallery with MediaGoblin**
Share your media with friends and family on a private, YouTube-like platform
- 38 Switch to the btrfs file system**
Get ahead of the curve by setting up this next-generation storage technology
- 46 Program a client-server application using Qt**
Find out how to start programming sockets by building your own client and server

FileSilo.co.uk




96 Free downloads
Find out what we've uploaded to our digital content hub FileSilo for you this month

Reviews


77 Music players
What's the best Linux software for listening to music online and offline?



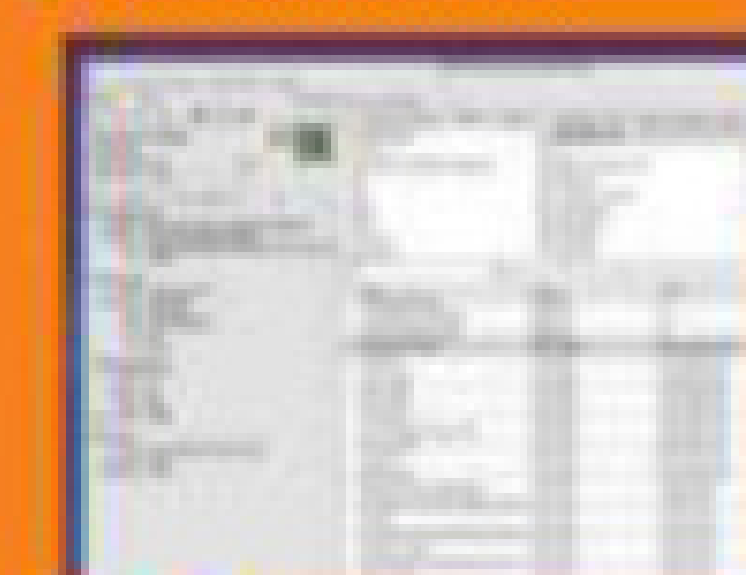
Nightingale



Tomahawk



CMus



GMusicBrowser



82 MinnowBoard Max
Has Intel listened to its customer feedback?



84 OpenPi
Can you design your own PCB using this open hardware?



52 Practical Raspberry Pi
Take your Pi out on the road by turning it into a car computer, paint electric circuits, secure your Raspberry Pi, add controls to your robot and create a digital assistant.

Join us online for more Linux news, opinion and reviews www.linuxuser.co.uk

PCI Express Ltd

TBS UK DISTRIBUTOR

EUMETSAT COMPATIBLE PROFESSIONAL PCIe TUNERS

EUMETSAT is the European operational satellite agency for monitoring weather, climate and the environment

The service provided by EUMETSAT helps to enhance and safeguard the daily lives of European citizens.

They aid meteorologists in identifying and monitoring the development of potentially dangerous weather situations and in issuing timely forecasts and warnings to emergency services and local authorities, helping to mitigate the effects of severe weather and protecting human life and property.

EUMETSAT delivers satellite data in real-time to users worldwide



The TBS6983 and TBS6908 are Professional level digital satellite TV Tuner cards with PCI Express interface. They support not only normal DVB-S2/DVB-S QPSK, 8PSK, but also CCM, ACM, VCM, Multi Input Stream, 16APSK, 32APSK, Generic Stream Mode which most satellite receiving devices can't support.

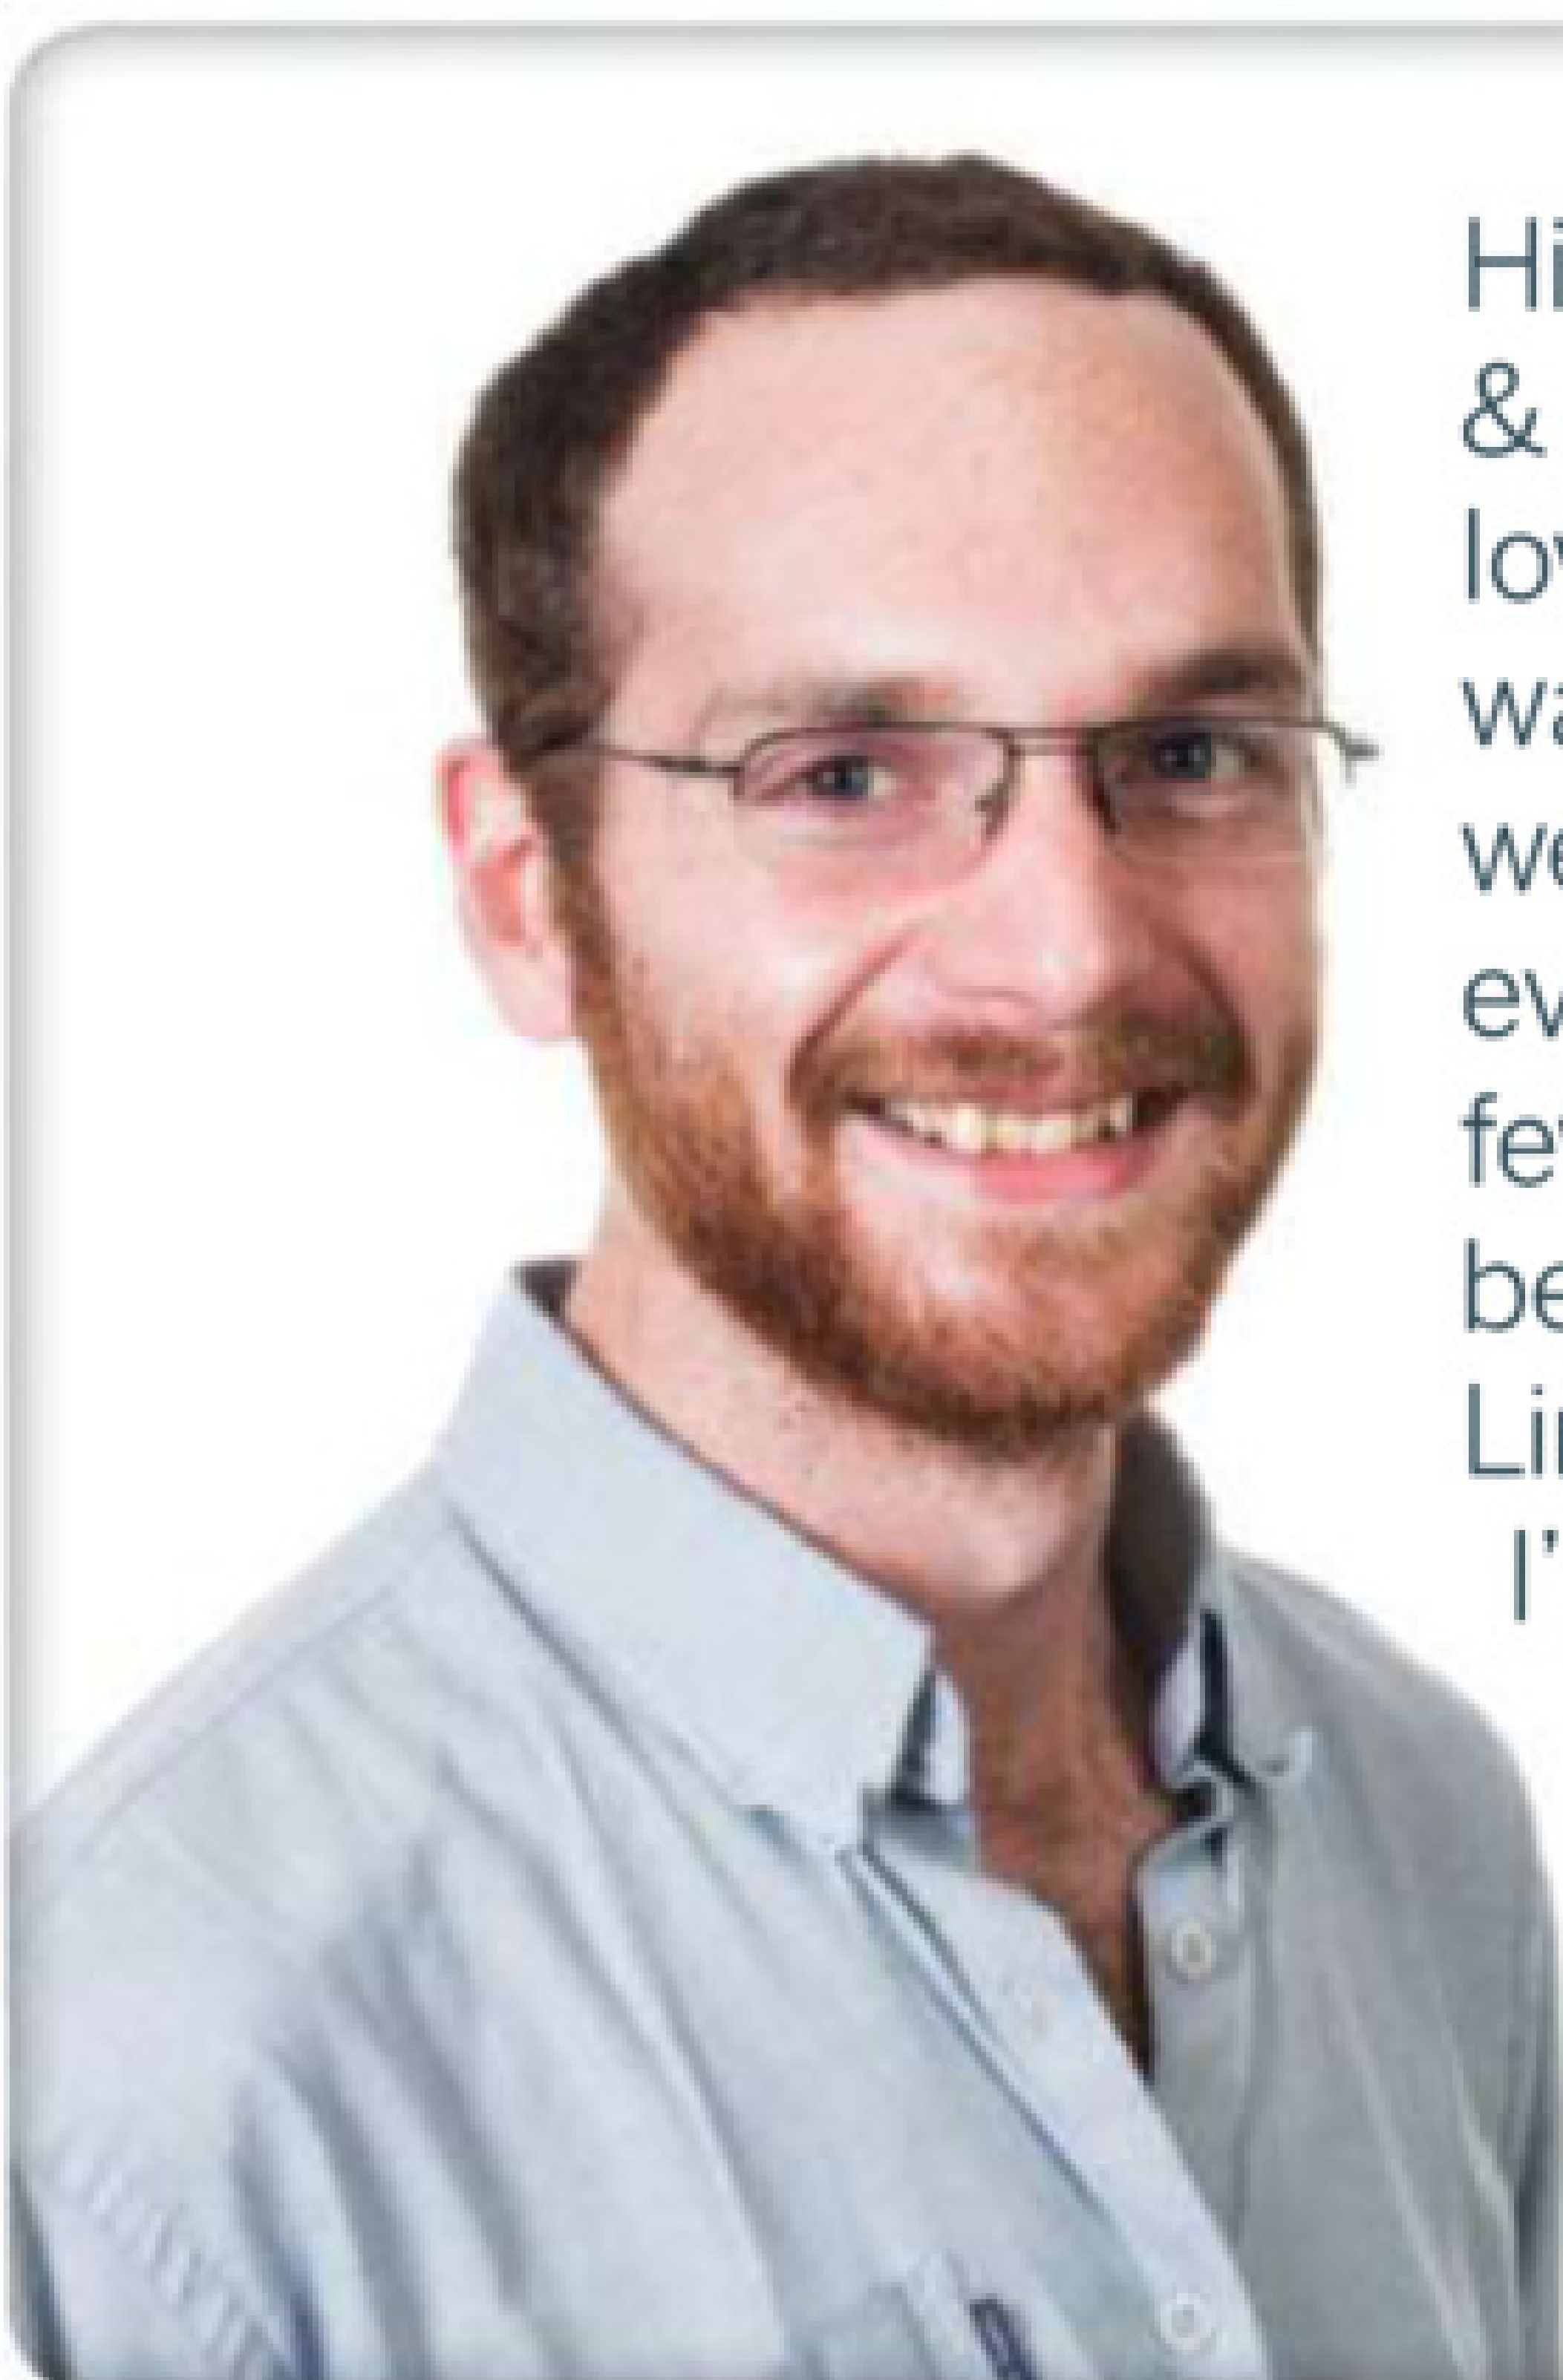
With use of dedicated TBS tools, those special streams can be captured. Windows BDA driver and Linux driver up to the latest kernel 3.X are ready.

WWW.TBSCARDS.CO.UK
Tel: 01892 833 007

Available online: www.valueav.co.uk | www.radtv.co.uk

Join our Reader Panel today!



Take our 3-minute survey at linuxuser.co.uk/survey and win a place on our panel



Hi. We love making Linux User & Developer and we hope you love reading it too. This year we want to make it even better, and we want you to be involved in its evolution. By answering just a few quick questions, you could be selected to join our first ever Linux User & Developer panel.

I'm excited to hear your views on the magazine and how we can improve it for you.

Gavin Thomas
Deputy Editor




Join our panel and help us make the magazine even better!



Take the survey on any device


Complete our survey to join our **Linux User & Developer** panel and:

- See new ideas and changes before anyone else
- Get invited to attend special events in your area
- Have a chance to come and meet the Linux User & Developer team
- Test third-party products for review
- Shape the future issues of the magazine
- Enter into exclusive competitions



Take the chance to get closer to the team than ever before

Only takes 3 mins!



Complete our survey and win your place today
linuxuser.co.uk/survey

OpenSource

08 News | 12 Opinion | 94 Letters

EMBEDDED

Snappy Ubuntu Core offered as IoT solution

Hoping for involvement from the Raspberry Pi community

Canonical founder Mark Shuttleworth is pushing his company heavily as the go-to operating system provider for Internet of Things (IoT) projects, via the new Ubuntu Core.

"I'm inspired by the inventors and innovators who are creating incredible machines – from robots that might clean to drones that follow us at play, to smarter homes which use energy more efficiently or more insightful security systems," Shuttleworth wrote in a blog post regarding his company's recent release. "Proving the power of open source to unleash innovation, most of this stuff runs on Linux – but it's a hugely fragmented and insecure kind of Linux. Every device has custom 'firmware' that lumps together the OS and drivers and device-specific software, and that firmware is almost never updated."

Shuttleworth's suggestion to resolve this issue is for IoT vendors to standardise on Ubuntu Core and its 'snappy' app system. "The snappy system and Ubuntu Core are perfect for distributed, connected devices that need security updates for the OS and applications but also need to be reliable and self-healing," he claimed. "Snappy is much better than package dependencies for robust, distributed devices."

Snappy apps and the Ubuntu Core distribution vary from traditional packaging

solutions by offering atomic upgrades and the ability to roll back at any point. They're positioned by Canonical as the solution to systems that require predictability and reliability and the company is working to convince developers of those benefits, beginning at the hobbyist level.

The recent launch of the Raspberry Pi 2, which brings an upgrade with the ARMv7-based quad-core Broadcom BCM2836 system-on-chip processor, came with the release of an Ubuntu Core image. This is joined by images for a variety of other hobbyist-targeted development boards, including the open source BeagleBone Black and Odroid-C1. These, Shuttleworth has promised, will be kept updated thanks to a partnership with the companies behind the hardware. "In this release of Ubuntu Core we've

added a hardware abstraction layer where platform-specific kernels live. We're working commercially with the major silicon providers to guarantee free updates to every device built on their chips and boards.

"If you're an inventor or a developer of apps that might run on devices, then Ubuntu Core is for you," Shuttleworth boasted. "With an app store (well, a 'snapp' store) built in and access to the work of thousands of communities collaborating on Github and other forums, with code for robotics and autopilots and a million other things instantly accessible, I can't wait to see what people build.


"I for one welcome the ability to install AI on my next camera-toting drone, and am glad to be able to do it in a way that will get patched automatically with fixes for future Heartbleeds!"

“ Snappy is better than package dependencies ”



Above Founder Mark Shuttleworth claims Canonical's Ubuntu Core will centralise IoT

Right You can get Ubuntu Core for the Pi from raspberrypi.org/downloads



DISTRO

CentOS releases regular rolling builds

Aims for monthly media releases

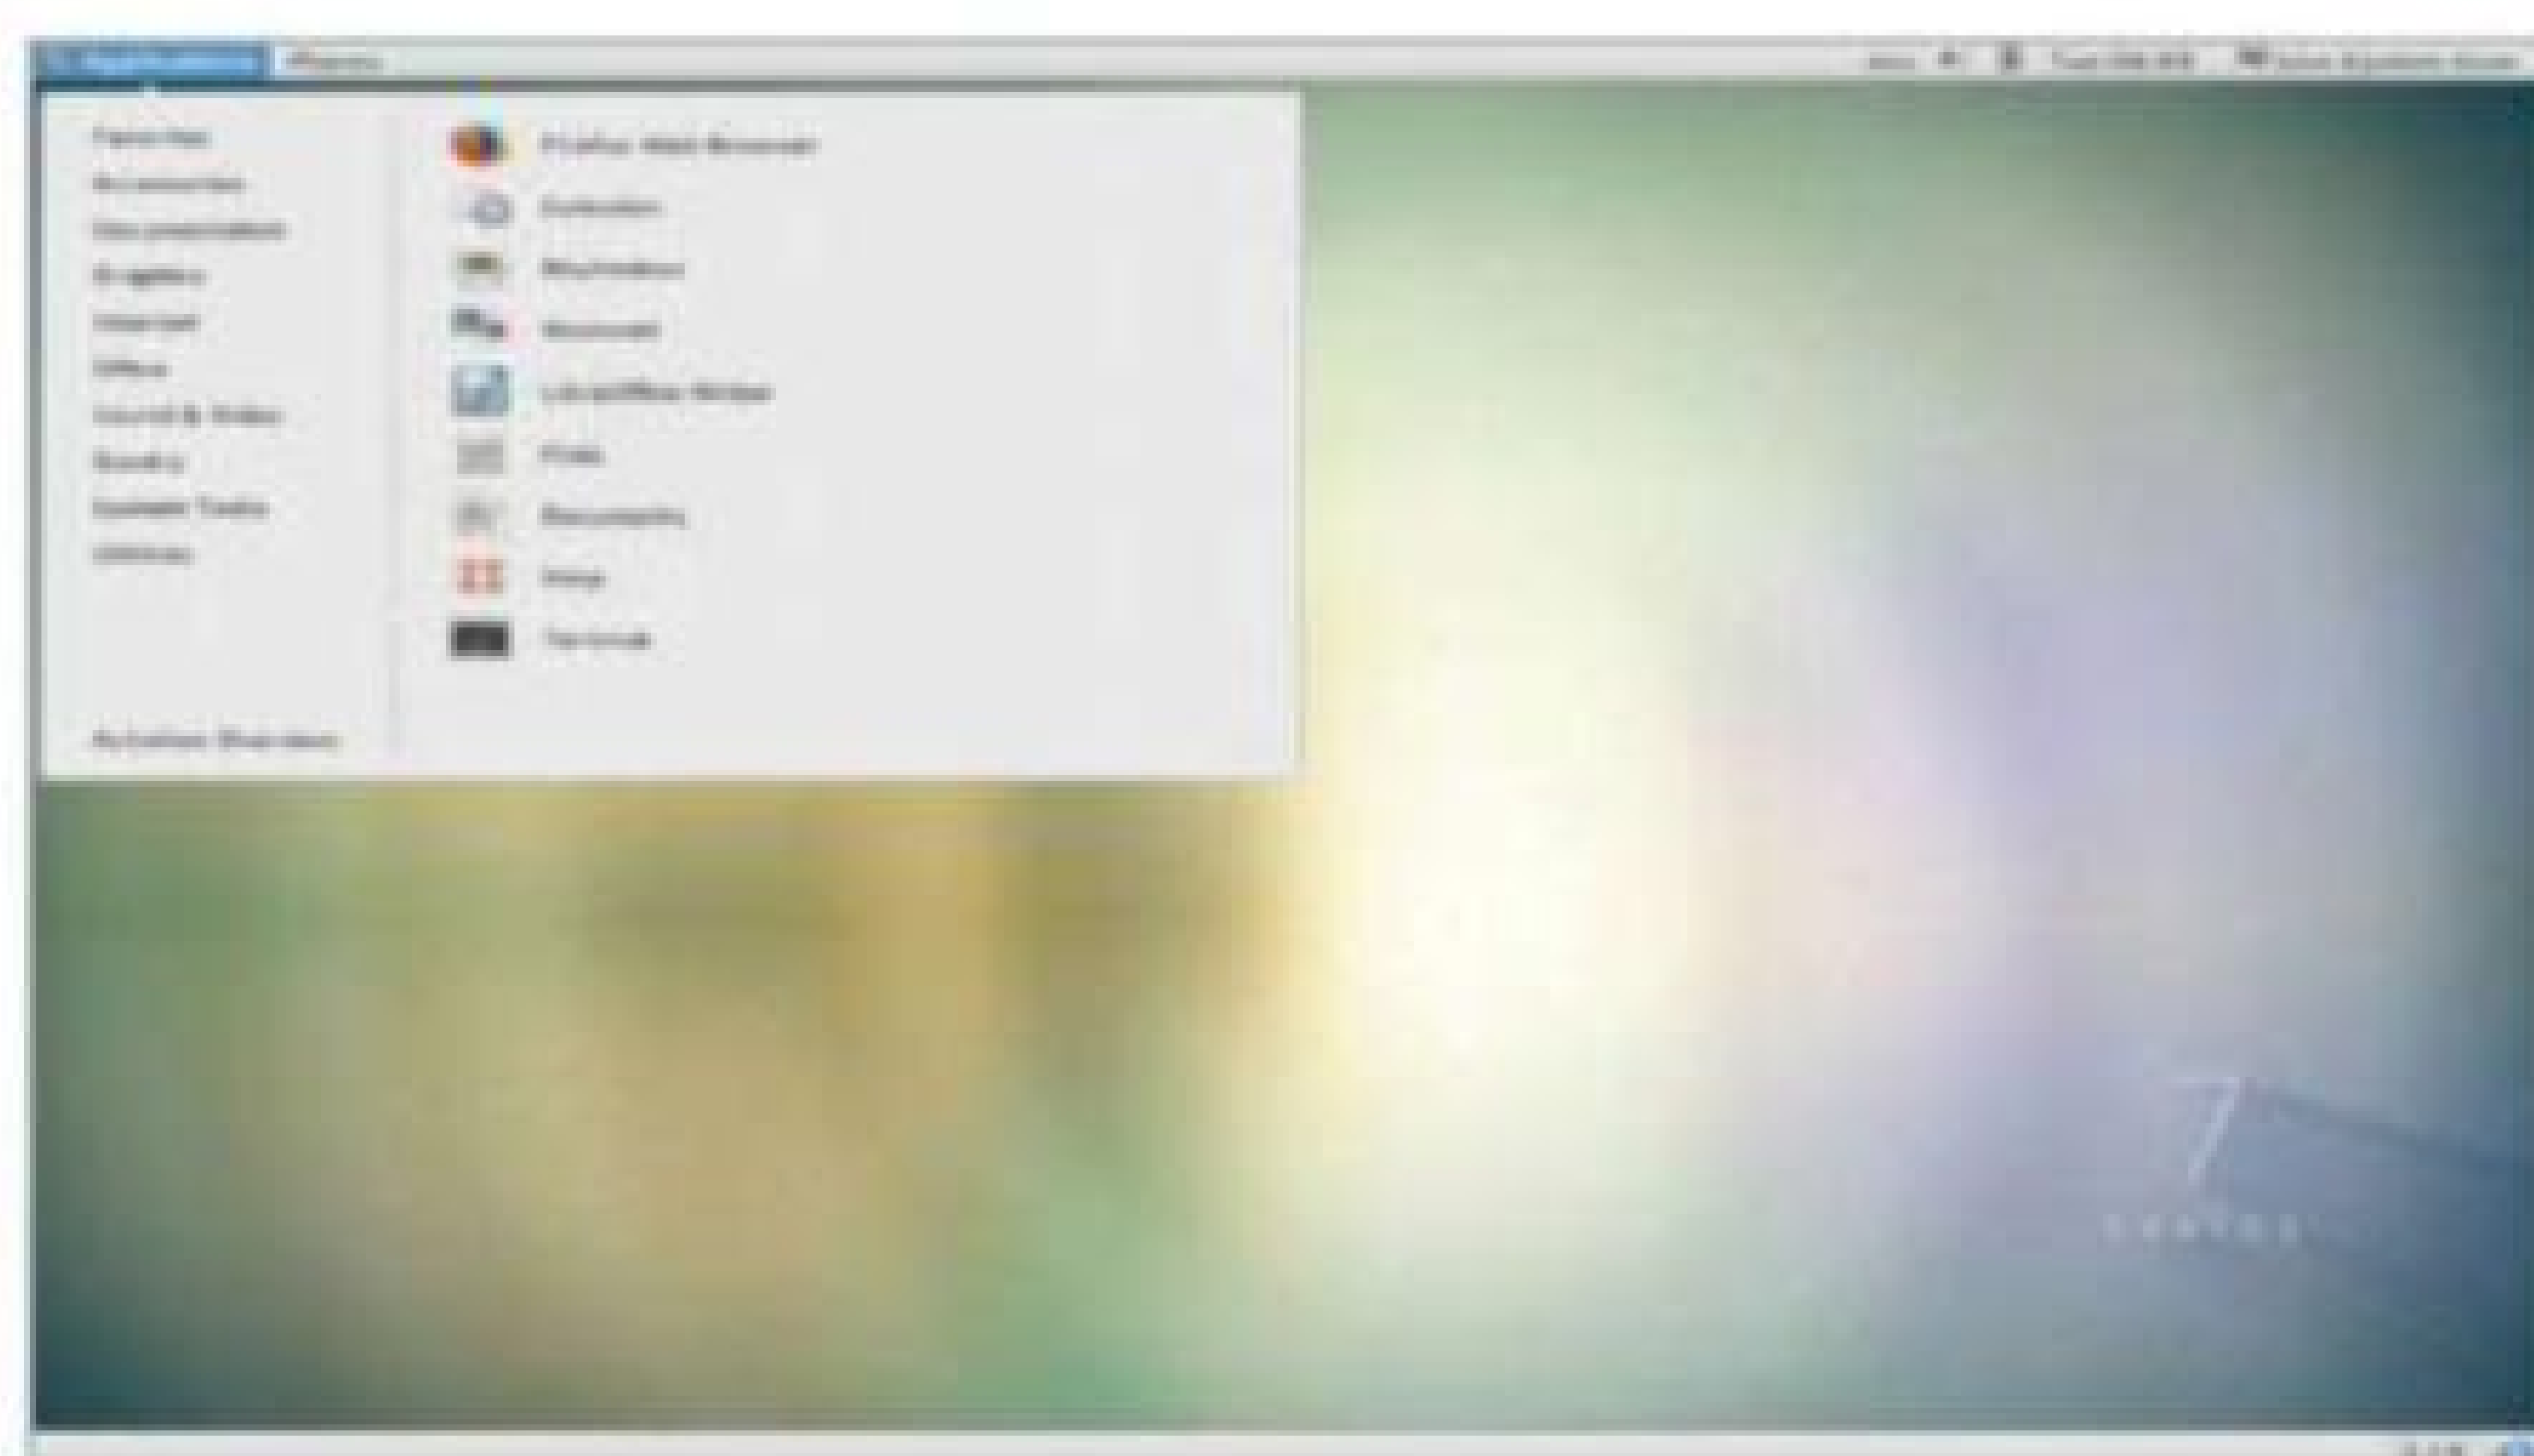
The CentOS Project has announced the availability of 'rolling builds' of its distribution, designed to offer a rolling-release update process to users who want it.

"CentOS Linux rolling builds are point-in-time snapshot media, rebuilt from original release time to include all updates pushed to mirror.centos.org's repositories," explained project lead of the launch Karanbir Singh. "Machines installed from this media will have all of these updates

pre-included and will look no different when compared with machines installed with older media that have been yum-updated to the same point in time."

The project aims to release new versions of the rolling builds at the end of each month, with a couple of days lag behind build time and release for testing. Singh has also indicated that special releases will be made following the discovery and patching of critical security vulnerabilities, such as the recent Heartbleed and Shellshock exploits.

CentOS' Karsten Wade has distanced the builds from other distributions' typical nightly builds, thanks to their additional testing. "As the release cycles progress, we'll be pulling in more images, such as CentOS Linux 7 live media, and probably future releases coming from project special interest groups (SIGs)," he added in the official launch announcement.



Above CentOS has 'rolling builds' that it aims to release monthly

DEVELOPMENT

Facebook opens machine learning modules

The research arm of social networking giant Facebook has released a series of modules designed to improve the ability of the Torch development environment to handle deep-learning tasks.

"These modules are significantly faster than the default ones in Torch and have accelerated our research projects by allowing us to train larger neural nets in less time," claimed Facebook's Soumith Chintala of the release. "Progress in science and technology

accelerates when scientists share not just their results, but also their tools and methods. This is one of the reasons why Facebook AI Research (FAIR) is committed to open science and to open-sourcing its tools."

Included are modules for GPU-optimised calculations of large convolutional nets and networks commonly used in natural language processing applications. The majority of the modules are based on Nvidia's CUDA GPGPU offload language.

Linux calendar

13th – 16th April

ApacheCon North America

- » Hyatt Austin, Texas
- » USA
- » events.linuxfoundation.org

Attracting developers across numerous Apache projects including Hadoop, Cassandra, Spark, Mesos, BigTop, CloudStack, Lucene and Solr, ApacheCon includes lectures, labs and community-building sessions.

17th – 19th April

Open Source Days 2015

- » Akademia Techniczno-Humanistyczna, Bielsko-Biala
- » Poland
- » 2015.opensource-days.mikstura.it

One of the largest free open source events in Europe, Open Source Days has been running since 2007 and brings together community members and developers from a wide range of different projects.

21st – 23rd April

Open Source Data Centre Conference

- » Hotel "MOA", Berlin
- » Germany
- » osdc.de

This three-day event includes workshops on Docker, Vagrant, NoSQL and Logstash, with presentations by Nigel Kersten, Luca Gibelli, Matthias Klein, Dr. Udo Seidel and more.

30th April

Technology.UG Leeds

- » DoubleTree by Hilton, Leeds
- » UK
- » technologyug.co.uk

A free event focused on networking, organised by a user group aimed at IT professionals. This season's event aims to look at virtualisation and hyper-convergence in the data centre.



The new Zumwalt Class Destroyer is a state-of-the-art stealth ship

KERNEL

US Navy picks Linux for new Destroyer

Linux base for Shipboard Computing Environment

The US Navy has released details of its latest project, the DDG-1000 Zumwalt Class Destroyer and a Total Shipboard Computing Environment (TSCE) which is based wholly on the Linux kernel.

While the US Navy had previously used Solaris to power the control stations on board its ship, a report issued in 2014 indicated that it would gradually begin transitioning over to Linux as this “overcomes hardware obsolescence issues with the Solaris-based control stations and provides lower-cost software updates.”

One of the first ships to benefit from this change has been confirmed as the DDG-1000 Zumwalt Class Destroyer, of which the Navy plans to commission three. While previous announcements made by the Navy have concentrated on the ship itself, including the use of an innovative ‘tumblehome’ hull designed to pierce waves and reduce its radar signature to one below that of a ship half its size, new releases showcase the benefits of Linux as part of its new multi-level Ship’s Mission Centre.

Featuring 16 million lines of code, the Navy’s TSCE has been built to its specifications by

defence contractor Raytheon and runs from 16 IBM blade servers distributed around the ship. An internal network connects Common Display Stations (CDS) to these, with any station being able to be reconfigured on-the-fly to take over any task and featuring touchscreens and standard USB interfaces for peripherals.

The US military’s move to Linux as a platform is the opposite of that employed by the Royal Navy in the early 2000s, when defence contractor BAE Systems standardised all its future development on the proprietary Windows 2000 platform from Microsoft.

OPEN SOURCE

LibreOffice 4.4 overhauled

“The most beautiful LibreOffice ever,” boasts TDF

The Document Foundation has released LibreOffice 4.4 with considerable attention having been paid to its complete user experience (UX) design.

“LibreOffice 4.4 has got a lot of UX and design love, and in my opinion is the most beautiful ever,” claimed Jan ‘Kendy’ Holesovsky, a member of the The Document Foundation’s Membership Committee and the leader of the LibreOffice design team, at the launch. “We have completed the dialog conversion, redesigned menu bars, context menus, toolbars, status bars and rulers to make them much more useful. The Sifr monochrome icon theme is extended and now the default on OS X. We also developed a new Color Selector,

improved the Sidebar to integrate more smoothly with menus, and reworked many user interface details to follow today’s UX trends.”

The new release also brings digital signing support during PDF export, new default templates, new fonts to better support files loaded from a Microsoft environment, improved change tracking and considerably cleaner code, among other improvements.

“LibreOffice 4.4 is the next in a series of releases that incrementally improve not only the features but also the foundation of the Free Software office suite,” claimed TDF’s Michael Meeks, while admitting that “it is, of course, not perfect yet.”

Gnome releases public app sandbox demos

Having previously worked on application sandboxing internally, the Gnome Project has released the first publicly-available implementation of its concepts.

“We have been working on putting together a working implementation of these [sandbox] ideas,” explained Matthias Clasen of his project’s release. “Alexander Larsson has made steady progress, and we are now almost at the point where it is useful for other people to start playing with it.”

Clasen claims that sandboxing will make it easier for third parties to create and give out applications for multiple distributions, plus run them more securely and make it easier to develop for Linux.

OPEN SOURCE

MediaFire releases Linux toolkit

Built out of VP's "personal need"

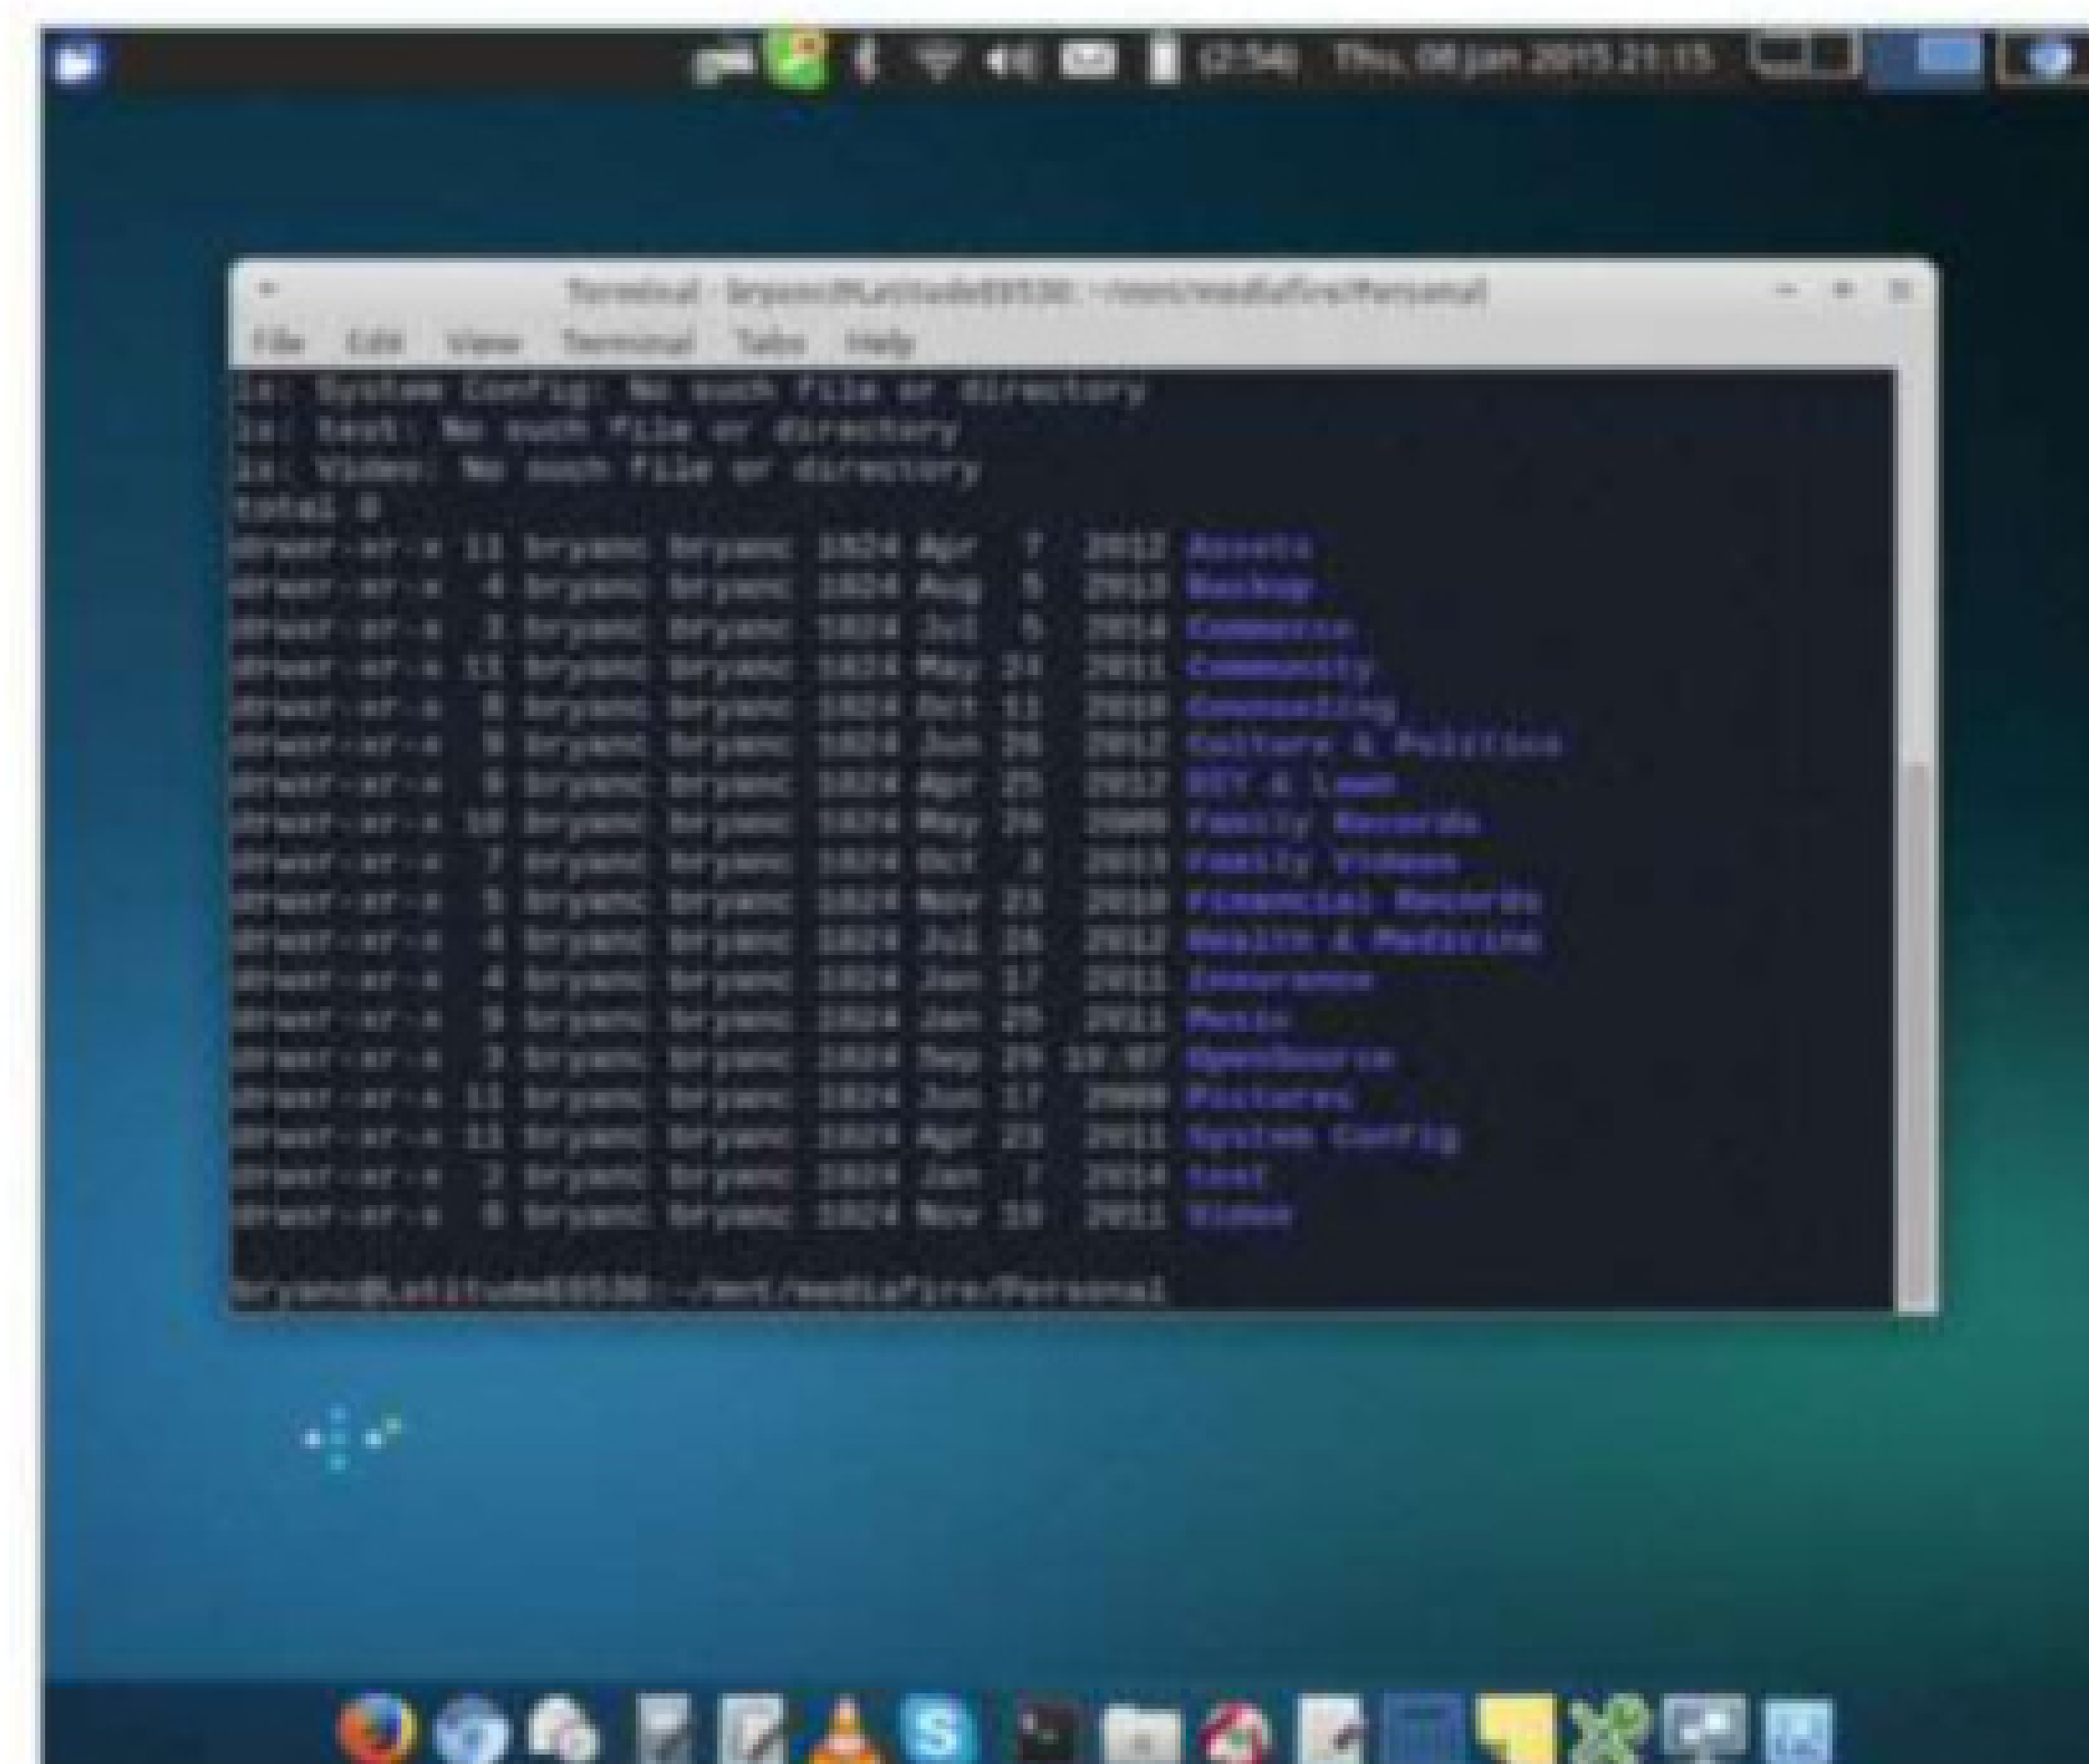
Cloud storage specialist MediaFire has announced the release of its first toolkit for Linux users made available under the GNU General Public Licence v2.

Describing the company as "huge fans of the open source community," MediaFire's vice president of software engineering Bryan Christ says the release marks the first time that the company has offered open source code for Linux users.

Launched under the name mediafire-fuse, the project includes a FUSE module for userspace mounting of cloud storage and a shell-like interface designed to resemble FTP for file upload, download and management tasks.

"As an avid Linux user and open source developer, the tools first took form a few years ago when a personal need arose to access my account without a browser," Christ explained. "After sitting untouched for quite some time, it seemed like a matter of good stewardship and practicality to breath new life into it. I find it personally rewarding when we can contribute to back to that ecosphere."

The mediafire-fuse toolkit is available to download now from the project's GitHub page, which includes compilation, runtime configuration and fstab mounting instructions for users: github.com/MediaFire/mediafire-fuse.



Above MediaFire's Linux toolkit includes an FTP-like interface and FUSE module

SECURITY

Glibc hit by 'Ghost' vulnerability

2013 patch not marked as critical

A security vulnerability in the GNU C library glibc, stretching back 14 years, has been discovered with numerous systems proving at risk despite a patch being released for the flaw in 2013.

Dubbed 'Ghost', the vulnerability was publicised by security firm Qualys and relates to a buffer overflow issue in the library's gethostbyname functionality. The company showed that glibc could be exploited remotely using a specially-formed email to gain a remote shell on an affected server.

While the flaw in the software stretches back across almost all releases since November 2000, the issue was discovered and patched in May 2013 but was not marked a security patch. As a result, numerous distros – in particular, long-term support distros – had not yet migrated to the fixed branch of glibc. Even when Google rediscovered the vulnerability in Chrome OS in early 2014, it wasn't brought to general attention until Qualys released its report in 2015. While the majority of supported distros have released backported patches to close the hole, the ubiquity of glibc means that many platforms, especially embedded systems, could be vulnerable.

Ghost joins similar long-standing issues recently discovered in widely-distributed packages, including GnuTLS and Bash, indicating that open source projects could benefit from improved analysis rather than reliance on the 'many eyes' concept and its hope that someone is looking for issues.

ACHIEVE STABILITY WITH IMPROVED KNOWLEDGE

Our excellent courses provide unrivalled training.

100% of our clients rate our PostgreSQL training programme as excellent. Book your place and gain access to expert knowledge of the core code.

+44 (0)870 766 7756
2ndquadrant.com/knowledge

2ndQuadrant Professional PostgreSQL

THE FREE SOFTWARE COLUMN

Freedom is slavery

Freedom of expression is a precondition for an open society, so why are encrypted communications under threat from government?



Richard Hillesley writes about art, music, digital rights, Linux and free software for a variety of publications

In the wake of the Charlie Hebdo killings an estimated two million people marched through Paris in the cause of freedom of speech, and Voltaire was misquoted in a thousand newspaper headlines. We are agreed that freedom of expression is not negotiable and we will defend it to the death, but we are often less sanguine about those whose ideas we oppose, as was illustrated by the incongruity of the French comedian, Dieudonné, being charged days later as an 'apologist for terrorism' for expressing empathy with one of the gunmen, writing "I feel like Charlie Coulibaly" on his Facebook page.

Many noted a similar incongruity in the presence at the head of the march of leaders and representatives of a number of countries where free speech is restricted and the lines between truth and freedom of expression are blurred. The Egyptians imprison Al Jazeera journalists for reporting the crackdown on the Muslim Brotherhood. The Turks imprison more journalists than any other country in the world. The Israelis have killed journalists operating in Gaza. The Saudis sentenced a liberal blogger, Raif Badawi, to 1000 lashes and 10 years in prison for the crime of apostasy, which isn't a crime in most countries.

"Encryption is mathematics, not technology"

As Neil Gaiman expressed it in a different context (bit.ly/1f76ktS), if you are going to stand up for the "freedom to write, freedom to read, freedom to own material that you believe is worth defending, it means you're going to have to stand up for stuff you don't believe is worth defending, even stuff you find actively distasteful, because laws are big blunt instruments that do not differentiate between what you like and what you don't, because prosecutors are humans and bear grudges and fight for re-election, because one person's obscenity is another person's art. Because if you don't stand up for the stuff you don't like, when they come for the stuff you do like, you've already lost."

Our own government's response to the threat to free speech was to suggest a clamp-down on internet security measures, a further restriction on freedom of speech. David Cameron posed the question: "In our country, do we want to allow a means of communication between people which we [the Government] cannot read?" His reply to his own question was "no, we must not." In other words, the government must have access to all of our private communications – an action worthy of any of the governments of which we claim to disapprove, although we are allied to many of the countries with governments of which we claim to disapprove. As Matthew Bloch noted "If the Prime Minister is proposing banning or hobbling our encryption standards by statute, he is proposing to limit freedom of expression. That's the freedom he claimed 'we would never give up' less than a week ago."

During the second world war, in a time of total war, our government read all the communications of our enemies. This could be justified by the exigencies of an existential threat to our democratic freedoms. We are not in a similar

situation. There are real threats, many of which are a consequence of our involvement in a wide range of recent conflicts across the middle east during recent decades, but these threats are no greater than those posed by the Provisional IRA three or four decades ago.

The Charlie Hebdo killers were already known to the French police and the killer of Lee Rigby was already known to the British police. The shortcomings were not in accumulation of data, but in attention to detail. Access to everything we hear or say isn't consistent with the democratic society we say we are trying to protect, and is unlikely to change the reality it claims to alter.

Cameron's declaration that our transactions should be readable by GCHQ amounts to a ban on the right to privacy. Even if this were not the case, the proposal to ban or restrict the use of encryption is stupid. The absence of encryption would drive any IT company out of Britain. Internet transactions would be impossible and online interaction with foreigners would run a considerable risk of being illegal.

Encryption is mathematics, not technology. Are we also to ban mathematical algorithms, and would this be policed by the algorithm police? Would GCHQ arrest anyone who did their sums wrong or disguised their communications to hide an affair?

In theory, the British government could demand backdoors into the encryption tools used by corporate entities, but a backdoor for one is a backdoor for all. What's the point in locking the front door if the back door is open, and those who live outside the law will always find a way in. Will the relinquishment of our right to privacy really give an advantage to GCHQ that it doesn't already have, and would companies like Google even choose to comply, given that their services would

THE OPEN START-UPS COLUMN

Learning freedom

The education market was impenetrable to FOSS companies, until schools' defences were broken

be compromised? Free software, of course, would have to be banned.

The government has also promised that, if re-elected, we will be treated to a new and revised Snooper's Charter (bit.ly/1qRlkkZ), which will allow GCHQ to gather data on a massive scale to analyse and profile our online habits for use by the police and other unspecified public bodies. "Powers that are currently being challenged in the courts, but are in practice available to GCHQ under programmes like TEMPORA, would become an everyday policing tool."

At the same time, the government has proposed to rewrite the Human Rights Act and replace it with a Bill of Rights. The text of this rehashed bill, though promised some time ago, has yet to appear, presumably because some of the rights will have to be denied to us, in the name of 'free speech' or some other pretext. Either that or it will be a rewrite of the current act, written in the spirit of the fictional Pierre Menard, who rewrote *Don Quixote* word for word and wound up with a masterpiece that was identical in every particular to the original *Quixote*. Menard's *Quixote*, we are told, was superior to the original, however, because it was "much richer in allusion" – it had to be "considered in the light of world events since 1602."

As Edward Snowden revealed, our privacy and human rights have been low in the priority scale for all our political leaders in their rush to prosecute proselytising wars, placate the owners of the presses or save us from ourselves. The Prime Minister may have won some tabloid headlines and proved to someone he is tough on terrorism, but if he is serious in his aim of banning encryption then we have to hope, like James Ball in the Guardian (bit.ly/1BWIXIC), that he finds someone else to do his job. "He could begin with a concussed kitten on a ketamine trip," writes Ball, "and work up from there."



Richard Smedley is a Unix and networking jack-of-all-trades, specialising in free and open source software

After a good start in the 1980s, with the UK government providing schools with open source code to use on the BBC Model B, free software languished for years, eclipsed by large and independent providers wedded to the Microsoft ecosystem (with the US also supporting companies writing proprietary code for the Mac).

Following the NHS Spine debacle (~£20 billion spent on a cancelled system) and austerity budgets, free software use has been encouraged for government contracts. But education changes slowly and its biggest disruption has come from the Raspberry Pi – a "risk-free environment" enabling teachers to be creative, according to Raspberry Jam pioneer Alan O'Donohoe. Raspberry Pi Foundation education director Clive Beale emphasises that the Pi is a "tool to make computing normal", solving problems in anything from music to metalwork.


We go to press shortly after BETT, the UK's education technology fair, where FOSS has been under-represented for years, mirroring its position in the education market. Beale told us how pleasantly surprised he'd been to see the Pi cropping up on many stalls, reflecting the constant spread of the Welsh-made board.

Helped by interest in the Internet of Things, *Minecraft* and physical computing, as well as curriculum changes, the Pi has won a place in UK schools where other free software businesses failed. But Raspberry Pi Trading Ltd, employing the engineers to develop the boards, funnels all of its profits into the Raspberry Pi Foundation, an educational charity that supports schools in their use of the Pi. From a behaviourist point of view, this is a successful social enterprise with a successful (outsourced) marketing department. That marketing was out in force at BETT, with not just the usual Pi gadgets but also the Astro Pi competition, where winning schools send their code to the ISS with British astronaut Tim Peake.


The Foundation is not the only company with a disruptive business model: TLM, The Learning Machine, provides school-age qualifications in IT that remove the administrative burden from teachers, keeping costs insignificant, and are neutral to software platforms. Add certificates to this, which demand pupils' time in giving back to the IT community, and you have a successful model for sneaking ideas about FOSS into schools.

We asked founder Ian Lynch how this works as a business: "TLM has taken the long-term view that freemium models are the sustainable future, where a service provides the revenue stream and the content and technology supporting the service is free," he told us. "We provide quality assurance for learning through certificates that provide performance points for schools in the league tables and funding for courses in Further Education. This is a disruptive model, so keeping costs down and customer convenience up is a key part of competitive strategy."

All the supporting content is CC-licensed, "minimising cost to customers while maximising convenience," says Lynch. "At scale we can still be more than 50% profitable. Open source strategies can take time – it needs patience but it makes for long term competitive advantage."



storemags



JON MASTERS

The kernel column

Jon Masters updates us as Linus contemplates Linux 4.0 and debate continues over how to handle the year 2038 ‘Y2K’ problem



Jon Masters is a Linux kernel hacker who has been working on Linux for some 19 years, since he first attended university at the age of 13. Jon lives in Cambridge, Massachusetts, and works for a large enterprise Linux vendor, where he is driving the creation of standards for energy efficient ARM-powered servers

Linus Torvalds released the final 3.19 kernel roughly on cue, noting that “nothing all that exciting happened [since the 3.19-rc7 release candidate], and while I was tempted a couple of times to do an rc8, there really wasn’t any reason for it.” As mentioned in last month’s issue, the new kernel includes a number of exciting new features: support for Intel’s MPX Memory Protection Extensions (which we covered in detail previously), a new HSA driver for AMD GPU devices, enhanced RAID 5 and 6 support in Btrfs, and the final promotion of Android’s Binder IPC mechanism out of the kernel’s staging tree. As usual, KernelNewbies have an excellent summary of the various patches with links to commits: kernelnewbies.org/Linux_3.19.

One of the more miscellaneous features merged into Linux 3.19 is that of DeviceTree Overlays. These are relevant to embedded devices, such as the RaspberryPi, which uses a DeviceTree to describe the many assorted platform (non-discoverable) devices that form the System-on-Chip (SoC) upon which it is built. Modern SoCs have dozens or even hundreds of devices, including a plethora of IO interfaces (such as USB and networking). In most cases, these are directly attached to the CPU cores on the same chip with no intervening PCI-like enumerable buses in between. This means that determining the physical topology and structure of the system at boot time necessitates the provision of a DeviceTree data structure. Overlays extend DeviceTrees by allowing for the provision of run-time extensions that describe the devices contained within Ras Pi add-ons or BeagleBone capes (for example). Thus, Linux 3.19 will make life much easier for those using custom Ras Pi boards.

Linux 4.0?

With the release of Linux 3.19 came the near-immediate opening of the merge window for the subsequent release. The merge window is a period of up to several weeks (the actual duration varying depending upon how Linus is feeling), during which intrusive and disruptive changes are allowed to the kernel. Traditionally these could be quite impactful, and while it is still true that significant churn happens in the merge window, these days most of the 10,000+ changesets – collections of patches – applied during this small window of time have been through an entire previous kernel 7-8 week cycle in Stephen Rothwell’s linux-next test kernels, or the equivalent maintainer development tree for certain non-core parts of the kernel.

In much the same way that Linux development has changed to be less eventful in terms of churn, so has the significance of kernel version

numbers. There was a time when kernels adopted odd/even numbering schemes in which an unstable 2.5 kernel series preceded the stable 2.6 series. But that was done away with during 2.6 development and there never was a 2.7 series. This was in large part due to the transition to the modern git-driven source development process and the 8 week kernel development cycle. As a consequence, Linux 3.0 came about not because of some new Earth-shattering development, but simply because Linus felt that the version numbers had grown too high. This same logic has been applied again, with Linus warning, “We are getting to release numbers where I have to take off my socks to count that high again”. Thus it is highly likely that the next kernel will be Linux 4.0, and even more likely that the change in version number will be entirely meaningless, except in product marketing literature which will likely use it to full effect.

The merge window for what will become the next kernel has already brought with it some goodies. For example, Linux has gained support for lazytime, a kernel file system mount time option that improves file system performance by intentionally delaying the update of certain file access times, so that simply reading files won’t result in many writes updating the associated metadata. This concept is of course not new. Kernels have long supported mount options such as relatime that are in widespread use (and the default for a number of distributions). These existing options also change the kernel behaviour – by only updating file access times under certain conditions – but they break strict POSIX compliance in the process. The new lazytime option instead performs the update of file access times but it stores this data in memory. It will not be written back to disk until the kernel has a reason to otherwise write file metadata, or a certain amount of time has passed (24 hours in the current implementation).

Storemags.com

Binding Android APIs



Linux 3.19 promotes Android's Binder out of the Linux staging tree and into the real part of the kernel. Binder is an Inter-Process Communication (IPC) mechanism, used to allow two different processes (applications, or 'apps') on an Android phone or device to communicate with one another by having one process call a runtime method provided by the other. It contrasts with traditional Unix IPC mechanisms such as SystemV IPC (SysV IPC). For the longest time, Binder lived as a driver module contained within Android patchsets that needed to be applied to an upstream Linux kernel before it could be used on an Android device. Several years ago, work to move Binder into the Linux staging tree took place and it was moved into the special drivers/staging subdirectory of the kernel, a place intended for unstable test driver code that isn't quite ready for production, to be made available to those early adopters who nonetheless have a need or a desire to use it.

Binder was never really a conventional driver in the sense that it supported some specific hardware add-on devices, but its presence in staging served two purposes. The first was political: Linux is still seen as a Unix-like operating system to many developers. Those people did not like the notion of Google inventing a new IPC mechanism and thrusting it upon the masses without a fight. Thus, moving the code into staging enabled those with concerns to be placated, while beginning to carry the code in the same location as the remainder of the Linux kernel source. Secondly, keeping Binder out-of-tree meant that it had to be maintained separately and then patched into kernels before they could be used on Android systems. After a sufficient period of time there was a certain acceptance that millions of phones and other devices are running Android and all of these are using Binder. As Greg Kroah-Hartman said when promoting it out of staging: "No matter what comes in the future, we are going to have to support this API".

Ongoing development

There was much debate over the past month concerning the continuing problem of the year 2038. This is the point in time at which the 32-bit time_t used in 32-bit Linux, Unix and Unix-like systems overflows such that time is seen to retrogress back to the year 1901. The cause of the 'bug' is simple: 32-bit integers can only encode 32 bits of information. Unix systems view time as the passage of seconds since a magic epoch of January 1 1970 (about the birth of Unix), and they encode time as a signed quantity of seconds added or taken from this value. Modern 64-bit Linux, Unix and Unix-like systems have updated APIs that allow for an unfathomable amount of time (292 billion years) to be represented in a full 64 bits, but there are many legacy 32-bit systems running 32-bit code, and there will be more 64-bit systems running legacy 32-bit code using special APIs, such as x32 and ILP32.

“Linux 3.19 promotes Android's Binder out of the Linux staging tree”



OUR 24/7 SUPPORT IS THE ULTIMATE SAFETY NET

Who else can offer a response within 15 minutes?

We offer unrivalled around-the-clock backup to PostgreSQL users. Platinum Support provides complete peace of mind.

+44 (0)870 766 7756
2ndQuadrant.com/safety_net

2ndQuadrant Professional PostgreSQL

TRIPLE BOOT

Two operating systems are so last year – here's how to start using three of them to maximise your productivity



Dual booting is a staple of being a Linux user these days.

Classically, a lot of people think of this as Linux and Windows coexisting together on one machine. There are people who just have two Linux distros though and there are many reasons why this could be the case – testing on two systems, one for leisure and one for work, is just one example of why two can be useful.

We can easily take a step beyond that though. It's doable (storage space permitting) to have three operating systems residing side-by-side. Whether you're having two Linux distros and a Windows install or simply three Linux distros, the concept is quite similar to dual booting and a natural extension of the practice.

Over the course of this feature, we will teach you how to perfectly partition your system, from a fresh hard drive to a pre-existing install, as well as a few tips on the best methods of installing the systems to get them to work together.

Why go for triple boot?

As we noted above, there are many reasons why you would want to both dual boot and triple boot, and they depend entirely on how you use your computer and how often you need to use different environments for different tasks.

One reason is often Windows – however we feel about it, many of us need it in our day-to-day lives. It could be something as simple as enjoying playing new games, which aren't always supported on Linux, or it could be the case that you are a designer who needs to use the industry standard Photoshop or InDesign. You can even install OS X for a Hackintosh build if that's more to your taste.

A key reason to further extend a dual boot setup is to preserve your main distro – the one containing the bulk, if not all, of your personal data and media. There are innumerable reasons as to why you may want or need to use different distros on a regular basis, and sometimes live-booting or virtualising just doesn't cut it – in such cases, it is incredibly convenient to have a third partition onto which you can install the distro you temporarily need to use. Non-Linux OSs aside, it could be something like wanting to have, for example, a Pentoo partition for testing alongside your main Debian distro, with a third slot for distro-hopping. It really is down to you!

As with any tutorial of this nature...
Make sure you back up everything you need first!

Installation order

Windows first, primary last

Installing Windows first can actually make the installation process a lot smoother – this is also good news if you've already got an existing or new Windows computer. The major benefits of installing Windows first is that you don't have to mess around with recovering and rebuilding GRUB at the end of the installation process, and it won't try and overwrite your currently installed Linux distros during its own installation process. The GRUB benefits also apply to installing your primary distro last, as you'll then be able to easily modify and update it later.

Linux first, Windows second

Performing installation in this order has its advantages by more easily tracking what you're installing and where. If you're setting up a disc from scratch or already have a Linux distro installed, you can use GParted straight away to get the disc formatted to your specifications. This means that while you're still in the live disc you can do the first installation. This can save a lot of time if you have limited resources for creating live discs or live USBs – you're already in Linux to edit the partitions, so why not install it?

Partitioning

Set up your hard drive so that space is usefully split between your different distros

In the first partition we have our primary default distro – this could be your work or leisure distro. We recommend about 20 GB ext4 partition for a Linux distro

The same as the first partition – a 20 GB ext4 partition for your other version of Linux. If you're triple booting Linux, you can use a third one of these quite easily

This partition is larger as we will install Windows here. Make it NTFS so Windows can see it during installation. It's best to give Windows 100 GB of space to be safe

Partition	File System	Label	Size	Used
/dev/sdc1	unknown	Microsoft reserved partition	128.00 MB	---
/dev/sdc3	ext4		19.53 GiB	486.07 MiB
/dev/sdc4	ext4		19.53 GiB	486.07 MiB
/dev/sdc5	ntfs		97.66 GiB	67.50 MiB
/dev/sdc6	ntfs		324.92 GiB	74.60 MiB
/dev/sdc2	linux-swap		4.00 GiB	0.00 B

We'll create a shared storage partition for all operating systems. This can be made up of all the remaining space and it's best to keep it to NTFS so everything can use it

This is the swap partition, used with the RAM when Linux is running. Similar to Windows' page file system but that resides on the main Windows installation partition

If you plan to set up the hard drive and install the operating systems from scratch, it's best to use GParted – found on most live distros or any maintenance distribution

Above we have listed what we believe to be a useful setup for triple booting your system, but this is only a guideline. The 20 GB sizes for the Linux distros take into account just purely installing packages – in many ways it's a very liberal estimation of how much space you're going to use, however this depends on your development habits and what kind of software you are planning on using. The order is also fairly arbitrary – it won't make any difference to disc speed but it may make sense to you personally.

While we do recommend a shared storage partition, the file structure of the home folder in Linux and Windows is quite different, which can easily complicate things.

Windows and Linux both allow you to mount specific parts of the partition to specific locations in their hierarchy though, which can make it a lot easier and quicker to organise. However, another option that you can think about is splitting up the storage partition between the two.

Installing Linux

Installing Linux alongside another distribution can be easy with the right distro

Installing Linux has been reduced down to the bare minimum interactions these days for a lot of distros. Ubuntu, Fedora, openSUSE and all your major, modern distros have their own graphical wizard either shared between them or created for the distro. Usually it's a case of just hitting install and overwriting the disk, but if the hard drive partitions are set up as we suggested then you won't want to do that.

In these distros though, you will be able to do one of two things: install alongside or specifically set up a place to install the root file system. The root file system is the core of the distro, where all the files to run it are stored, and it is represented by a '/'. Installing some distros from scratch will make it create a separate partition for the boot files or home folder and you can certainly make your custom partition setup do that as well.

On less advanced distros – perhaps those designed for older systems – the installation process can be a lot more involved. It will require you to know what partitions you've got and where they live on your system. While they may have their own partition software it is likely to be a lot more manual, so in this case we'd still recommend using GParted with another live distro first to get the partitions sorted beforehand. During this process make a note

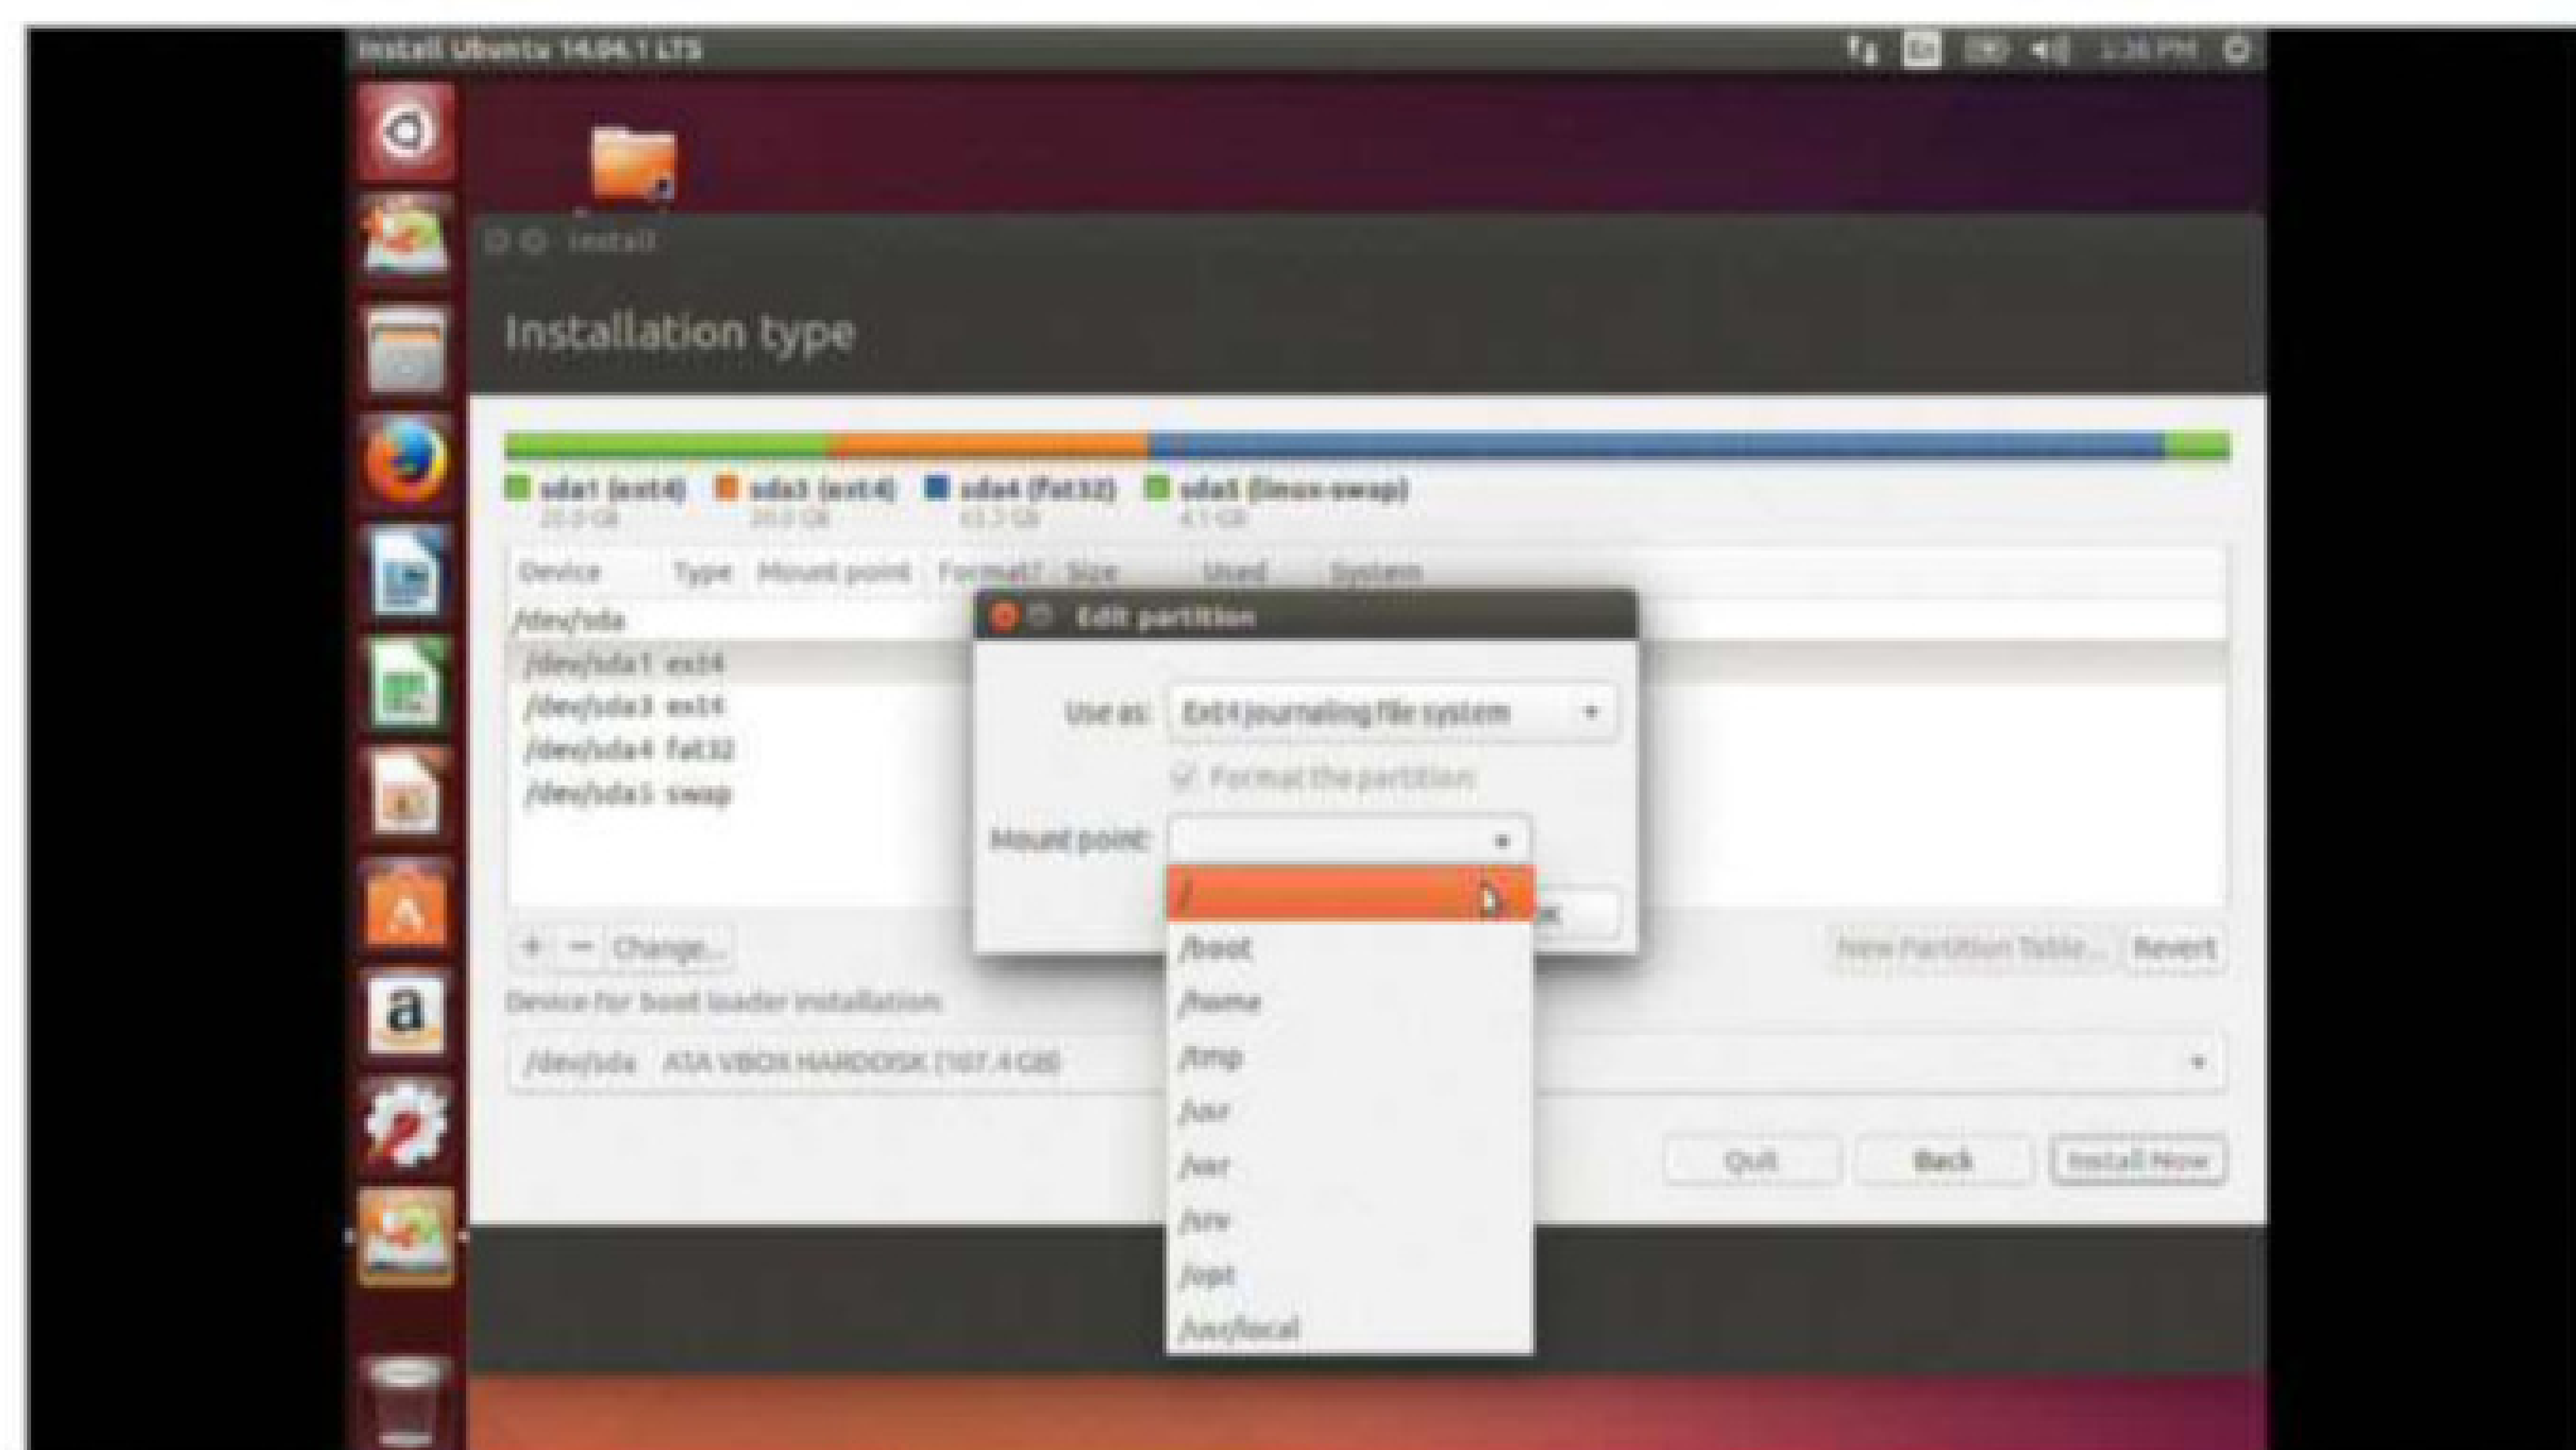


Above Always read the disk options during installation – it's easy to erase the disk entirely

of what the hard drives are called during the partitioning process – this will be something like /dev/sda for the hard drive and sda1, sda2, sda3, etc for the individual partitions. These numbers won't always be in the order you expect though, so it's best not to guess it.

All distros, live and installed, will automatically use any swap partition on the system. You only need one of these for your system as you are just running one distro at a time, and you don't need to set it as the correct swap when installing either.

For further installation advice always make sure to read the available options and if all else fails, seek out original documentation on the distro's website.



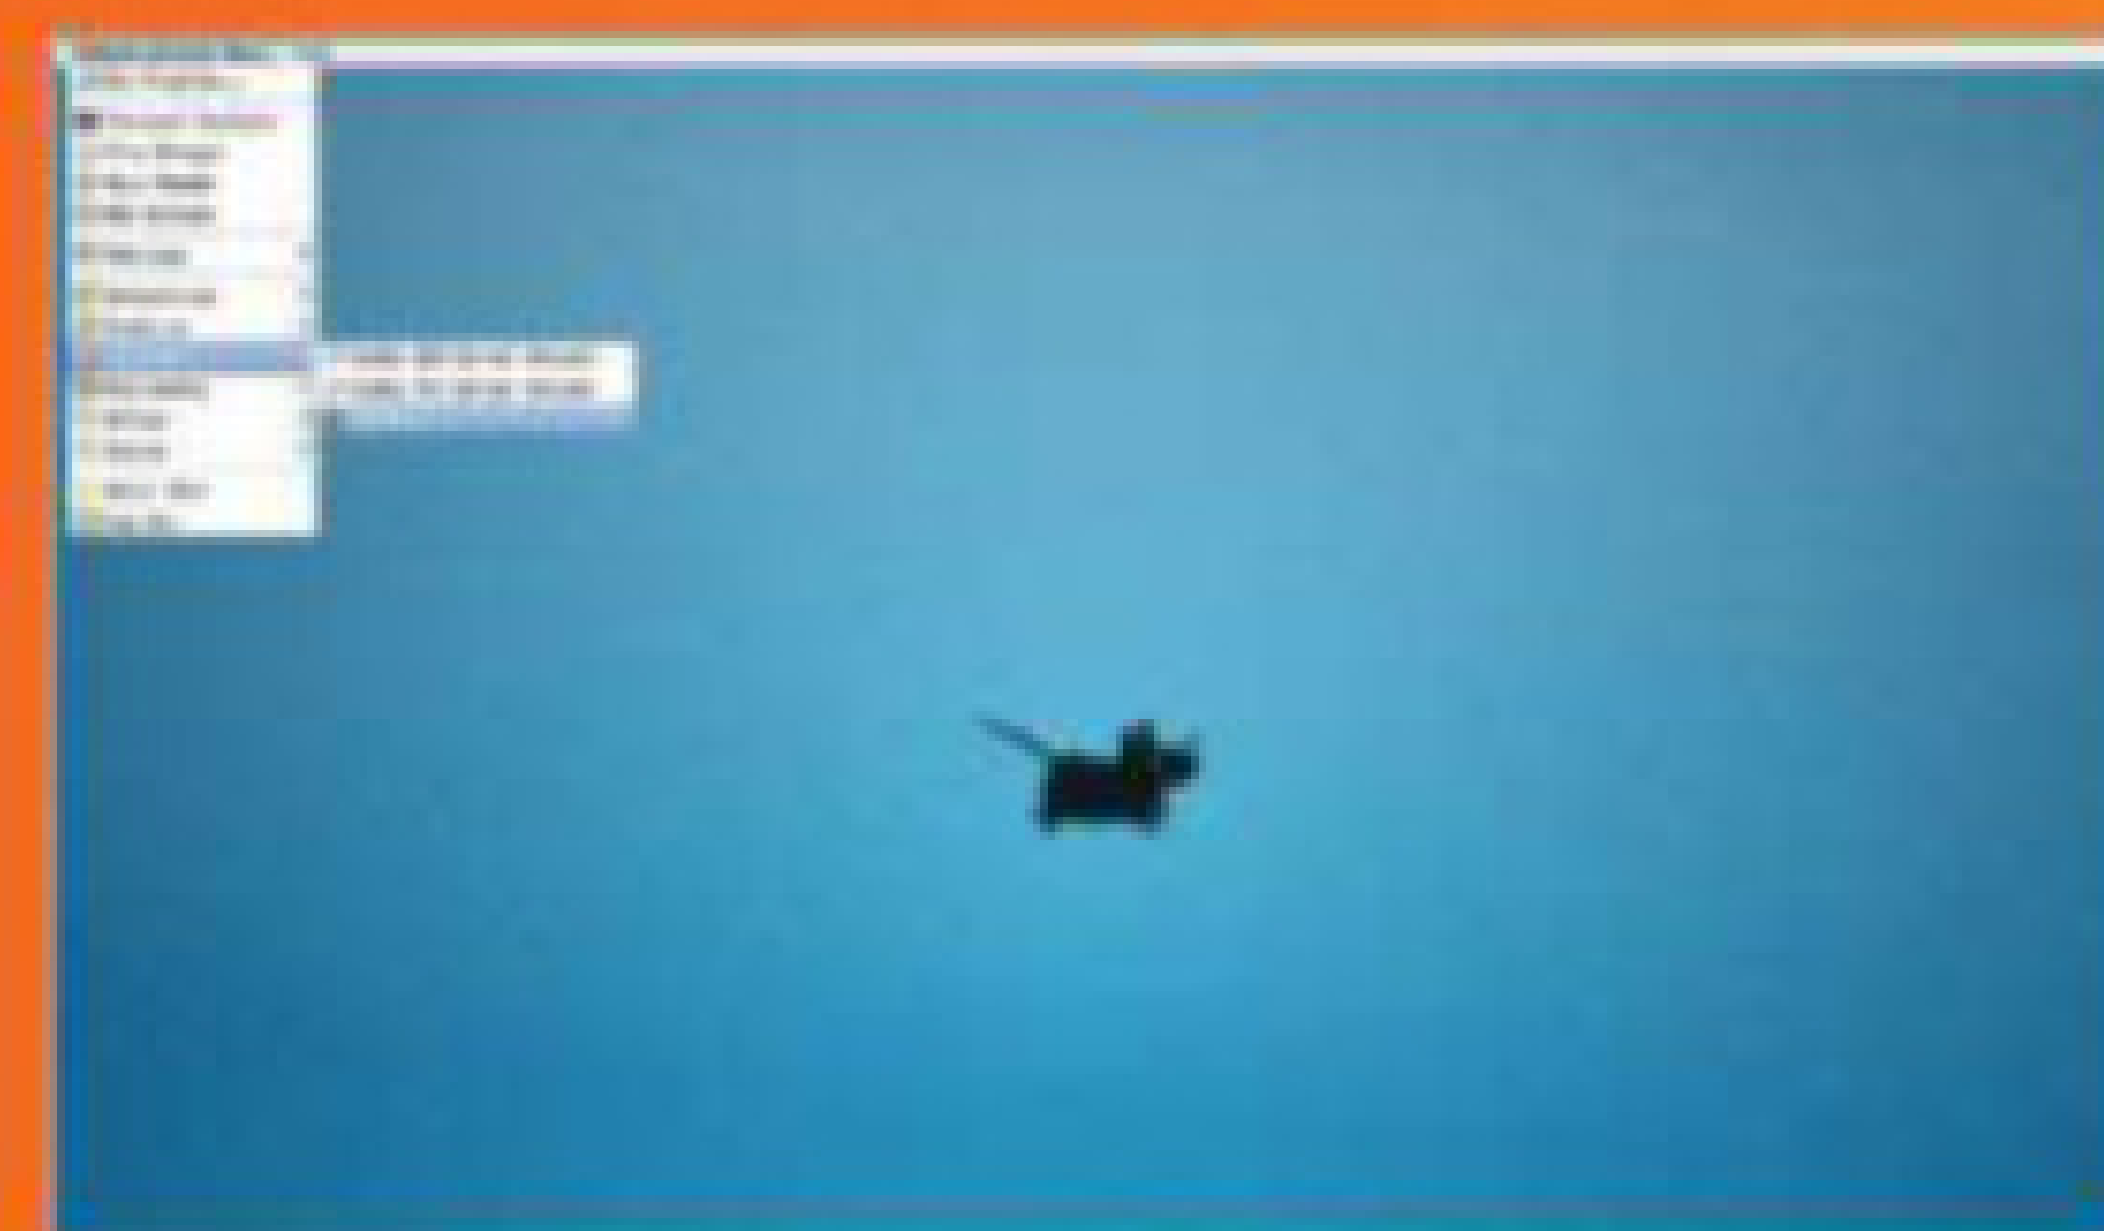
Above In Ubuntu you can select custom partitions for different areas of the distro's files

Types of distros



Everyday elementary OS

A beautiful looking distro that is easy to use and yet still offers everything you would want Linux for.



Development Arch Linux

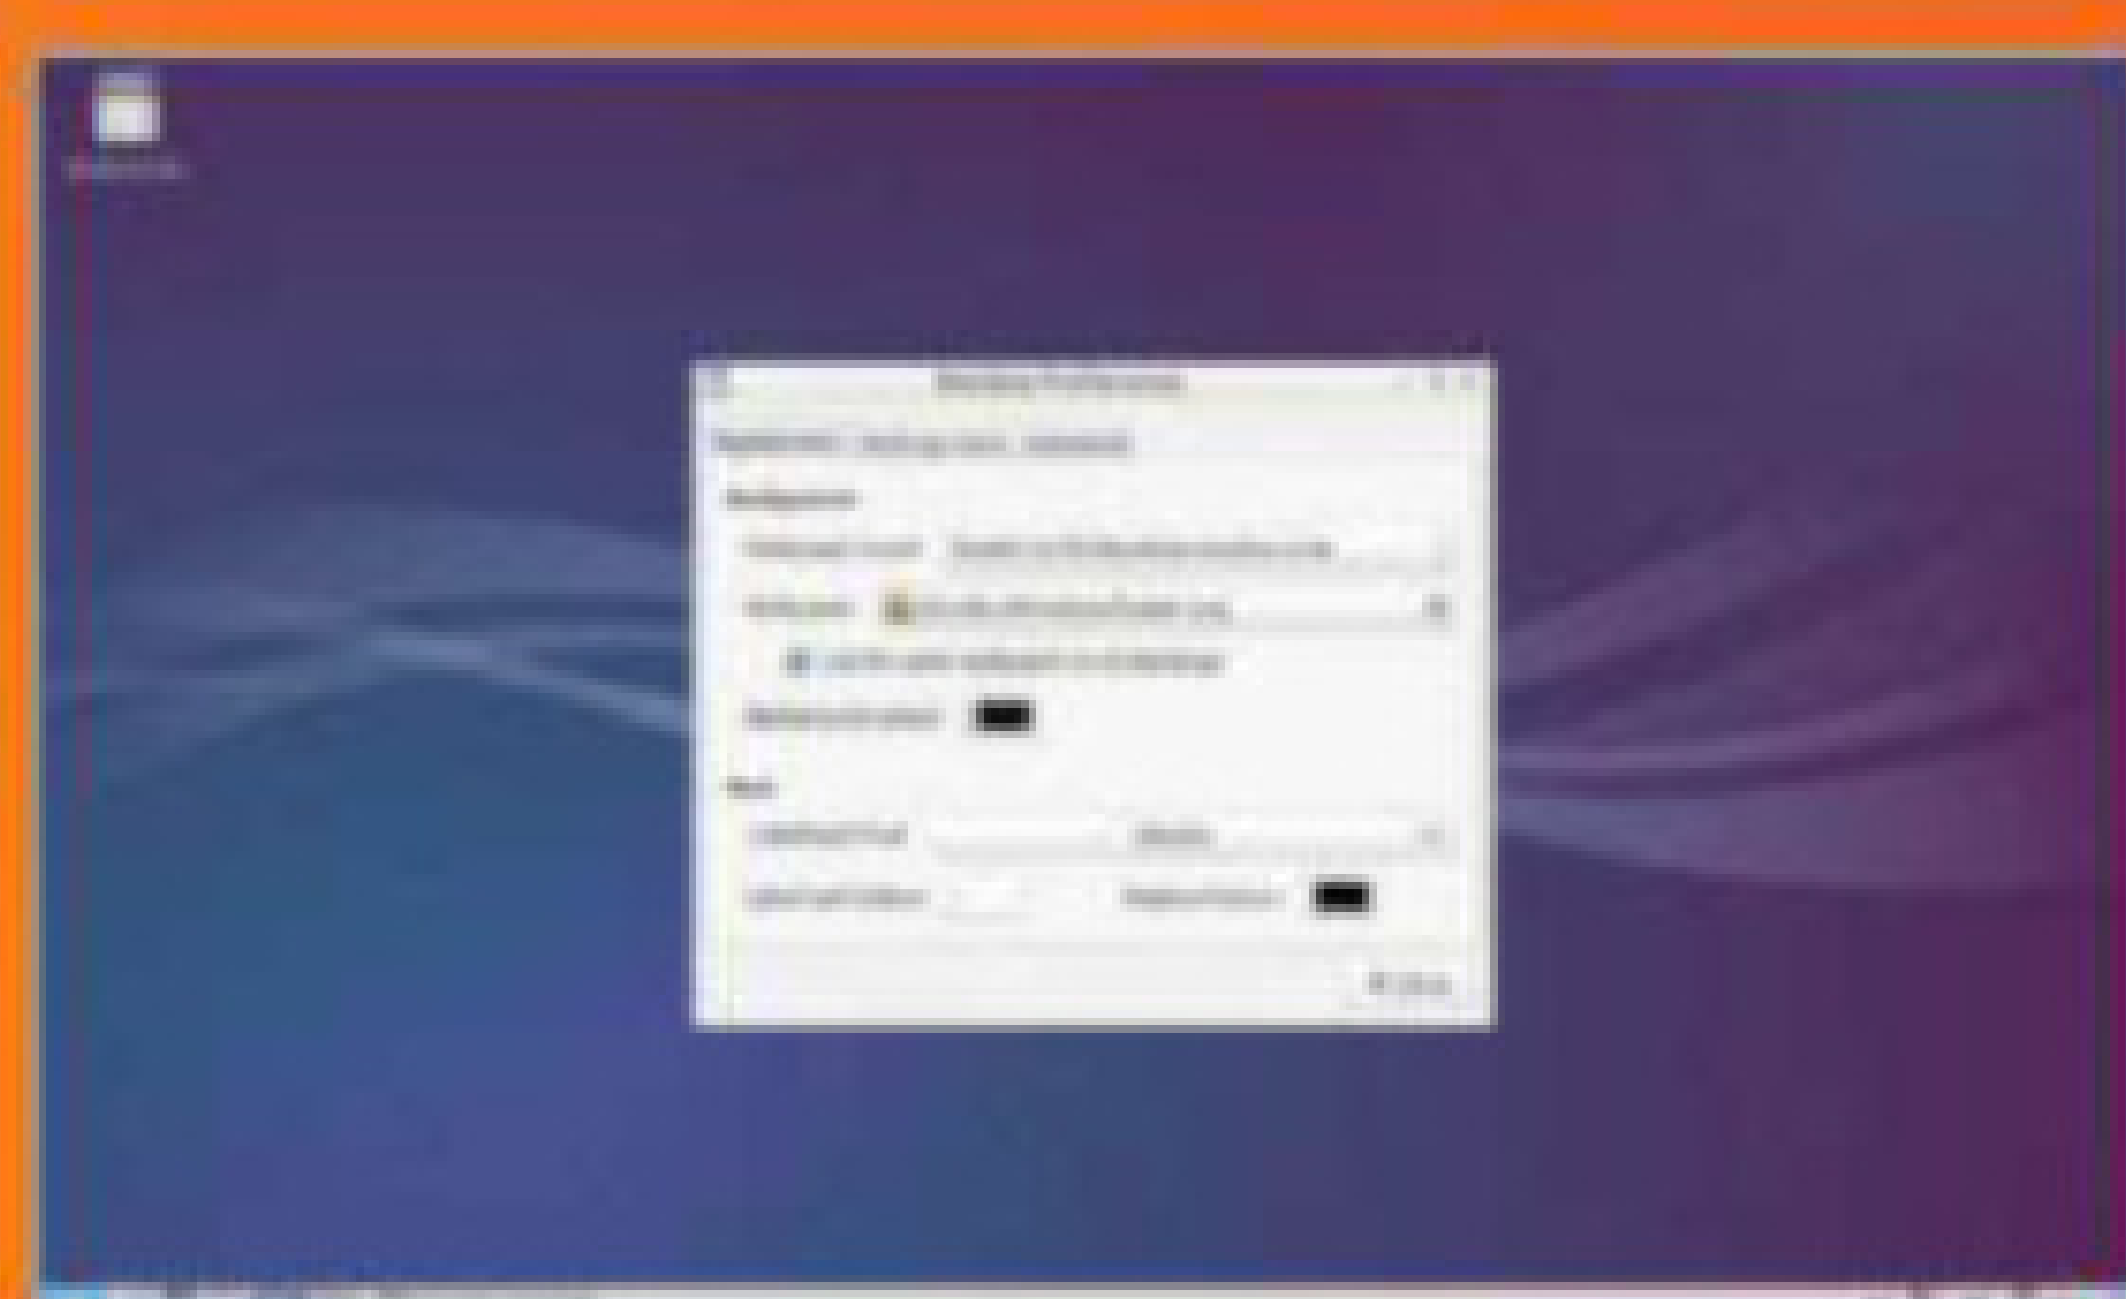
A serious distro for serious Linux users, Arch has everything that you need to develop your skills.



Entertainment OpenELEC

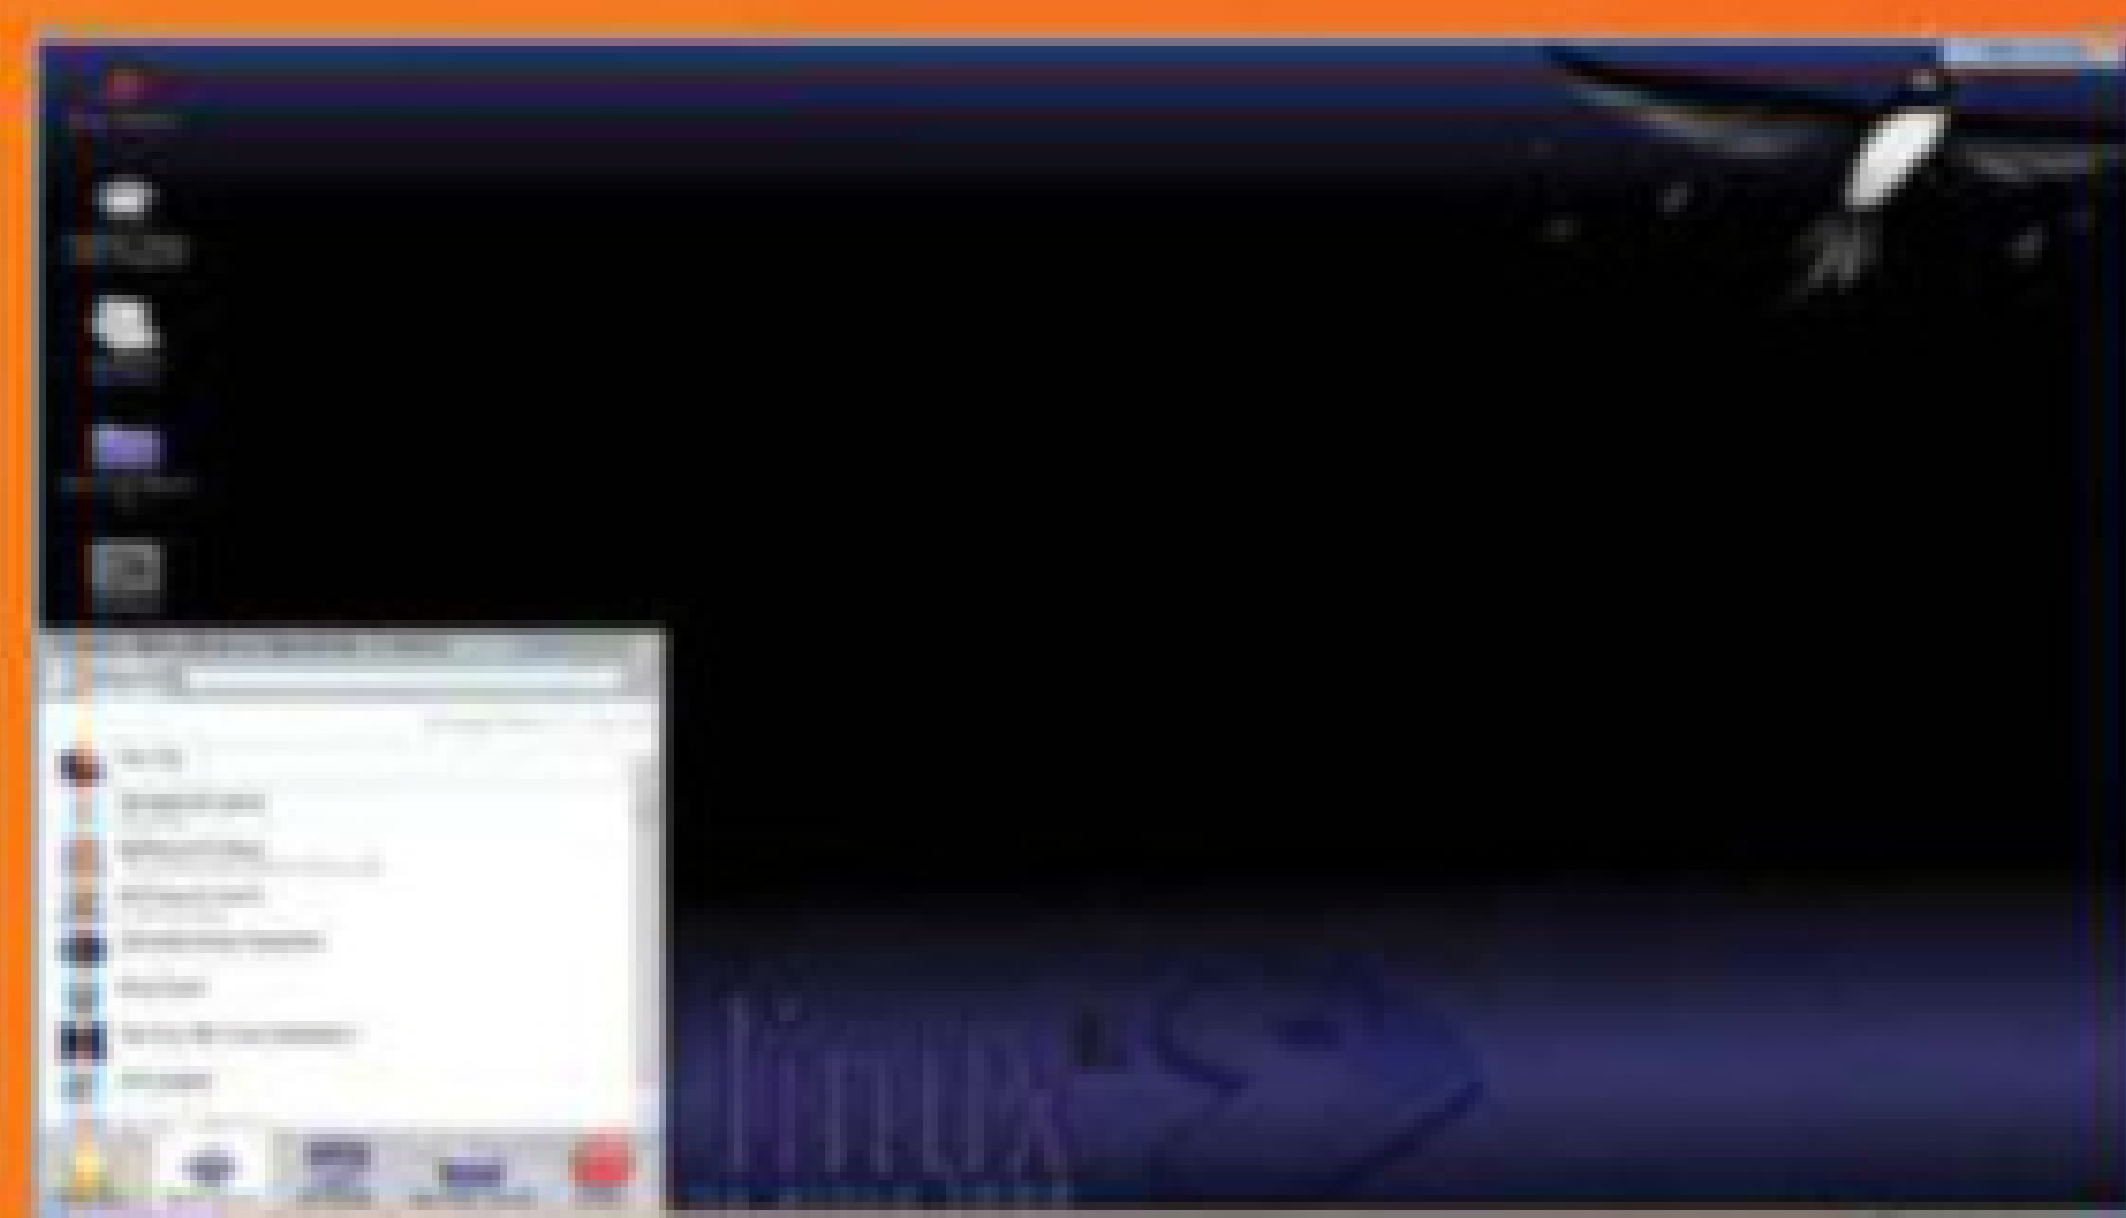
Fully explore and experiment with Kodi using OpenELEC. It is the best way to view media on Linux.

Choosing distros to use is easier if you think in terms of their categories



Lightweight Lubuntu

Being light without sacrificing anything is one of Lubuntu's major pros. There are many more advantages though.



Rolling release Gentoo

You can get the very latest packages and updates in one of the most customisable setups around.



Working desktop openSUSE

It's great for enterprise and it's also great if you want to just do some work without any distractions.

Recovering GRUB

Get the boot menu back if things go wrong or manage it with a different distro

One of the issues you may find with installing one or two distros next to each other is that you might mess up GRUB, the boot manager used by Linux to actually boot into the distros and other operating systems. You may also want to manage it with your 'default' OS. Both of these can be fixed using our recovering GRUB guide below but unfortunately this won't help you recover an operating system you've written over.



01 Live boot Any of the distros we've been using will work for this – you can even technically do it from another installed Linux distro if you're already inside it. When you boot into the live Linux, you may need to install the grub package. Ensure that you do it from the terminal before continuing.

03 Restore GRUB Assuming your primary partition is sda1, your installation hard drive is sda and you mounted it to mnt, you can now restore GRUB using the data from the partition that you just mounted with:

```
$ grub-install --root-directory=/mnt /dev/sda
```

02 Mount the hard drive Some distros like Ubuntu, for example, will let you click and just auto-mount the internal hard drive from the live environment. However, you can also do this in the terminal. Mount the primary install partition to a logical spot using something like:

04 Reboot After a reboot, GRUB should be back to normal and at the very least you can boot into your main distro. From there you may need to update GRUB further – to do this, open the terminal and perform the following two commands:

```
$ mount /dev/sda1 /mnt/
```

```
$ update-grub  
$ sudo grub-install
```



Above Your GRUB components all live inside the etc folder of your Linux distro

Windows installation

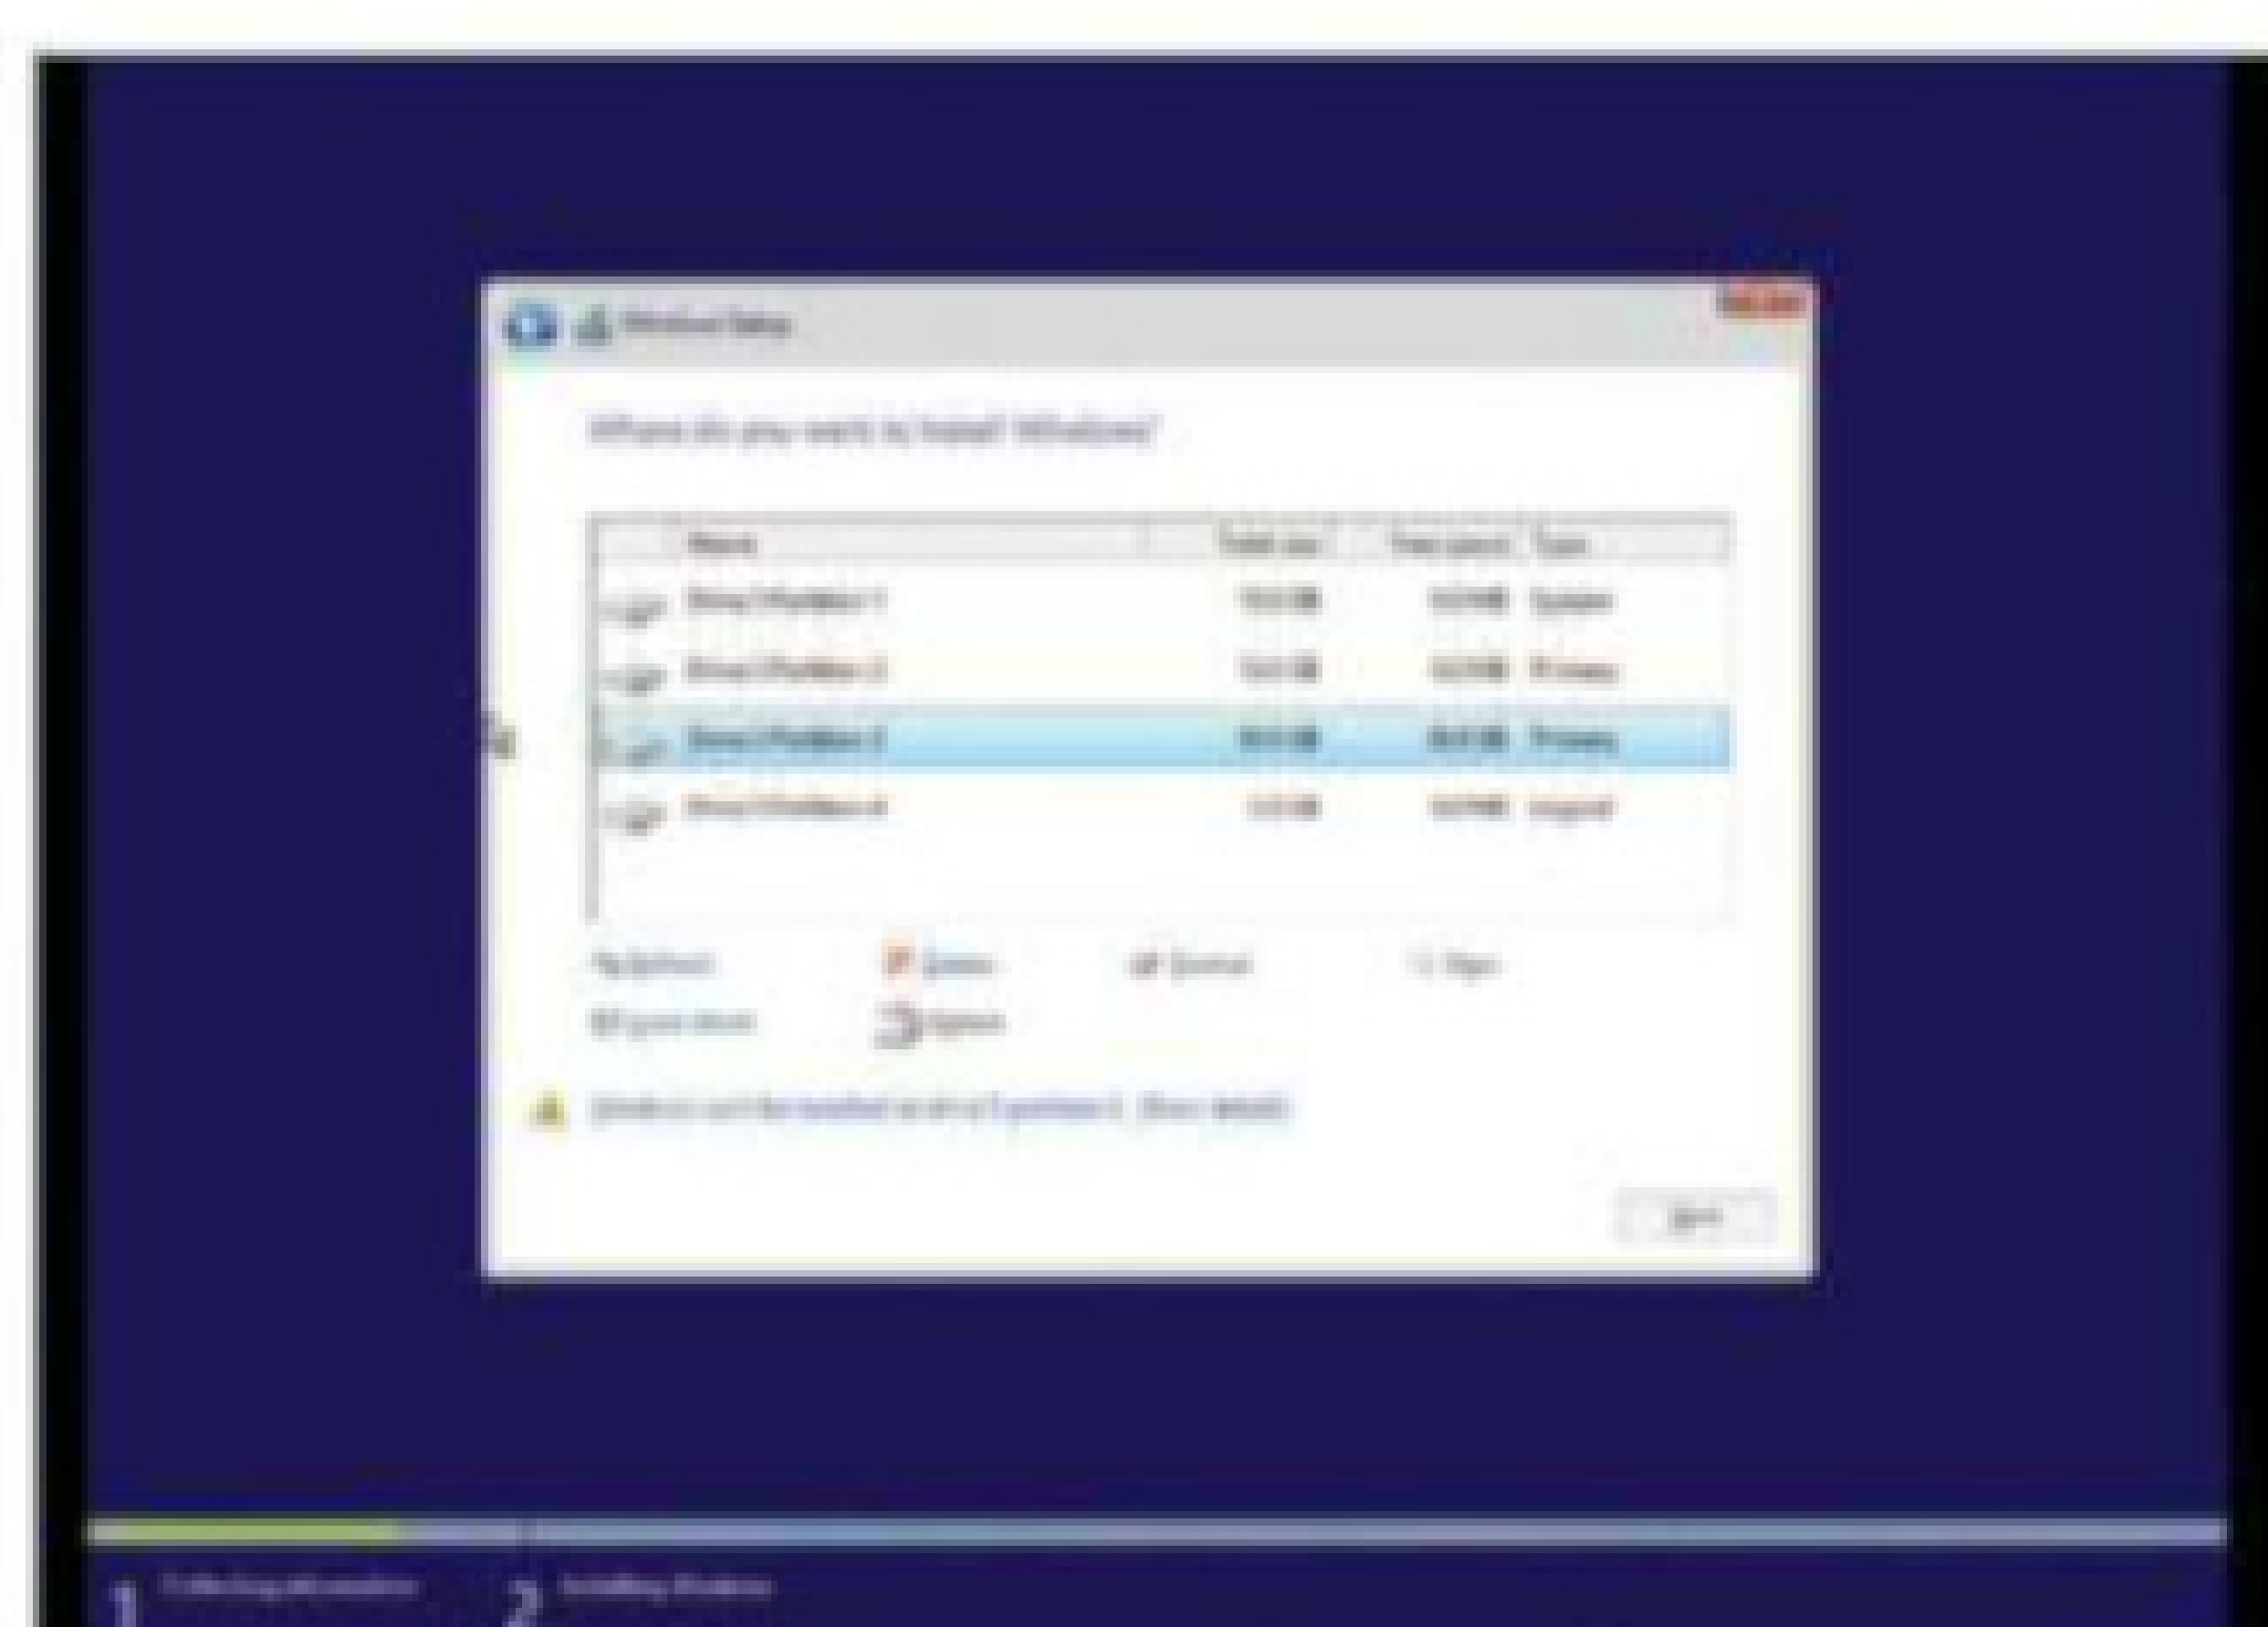
Microsoft's operating system is a necessary evil to some, but here's how to summon it safely from its dark pit

If you've already got Windows installed, we suggest looking over the page to figure out the best way to prepare for installing other operating systems alongside it. Otherwise, installing Windows when Linux is already there is not quite as easy as its Linux counterparts. Windows would prefer to completely wipe the disc and set itself up as the ruler of your computer, but with some persuasion you can get it to play nice.



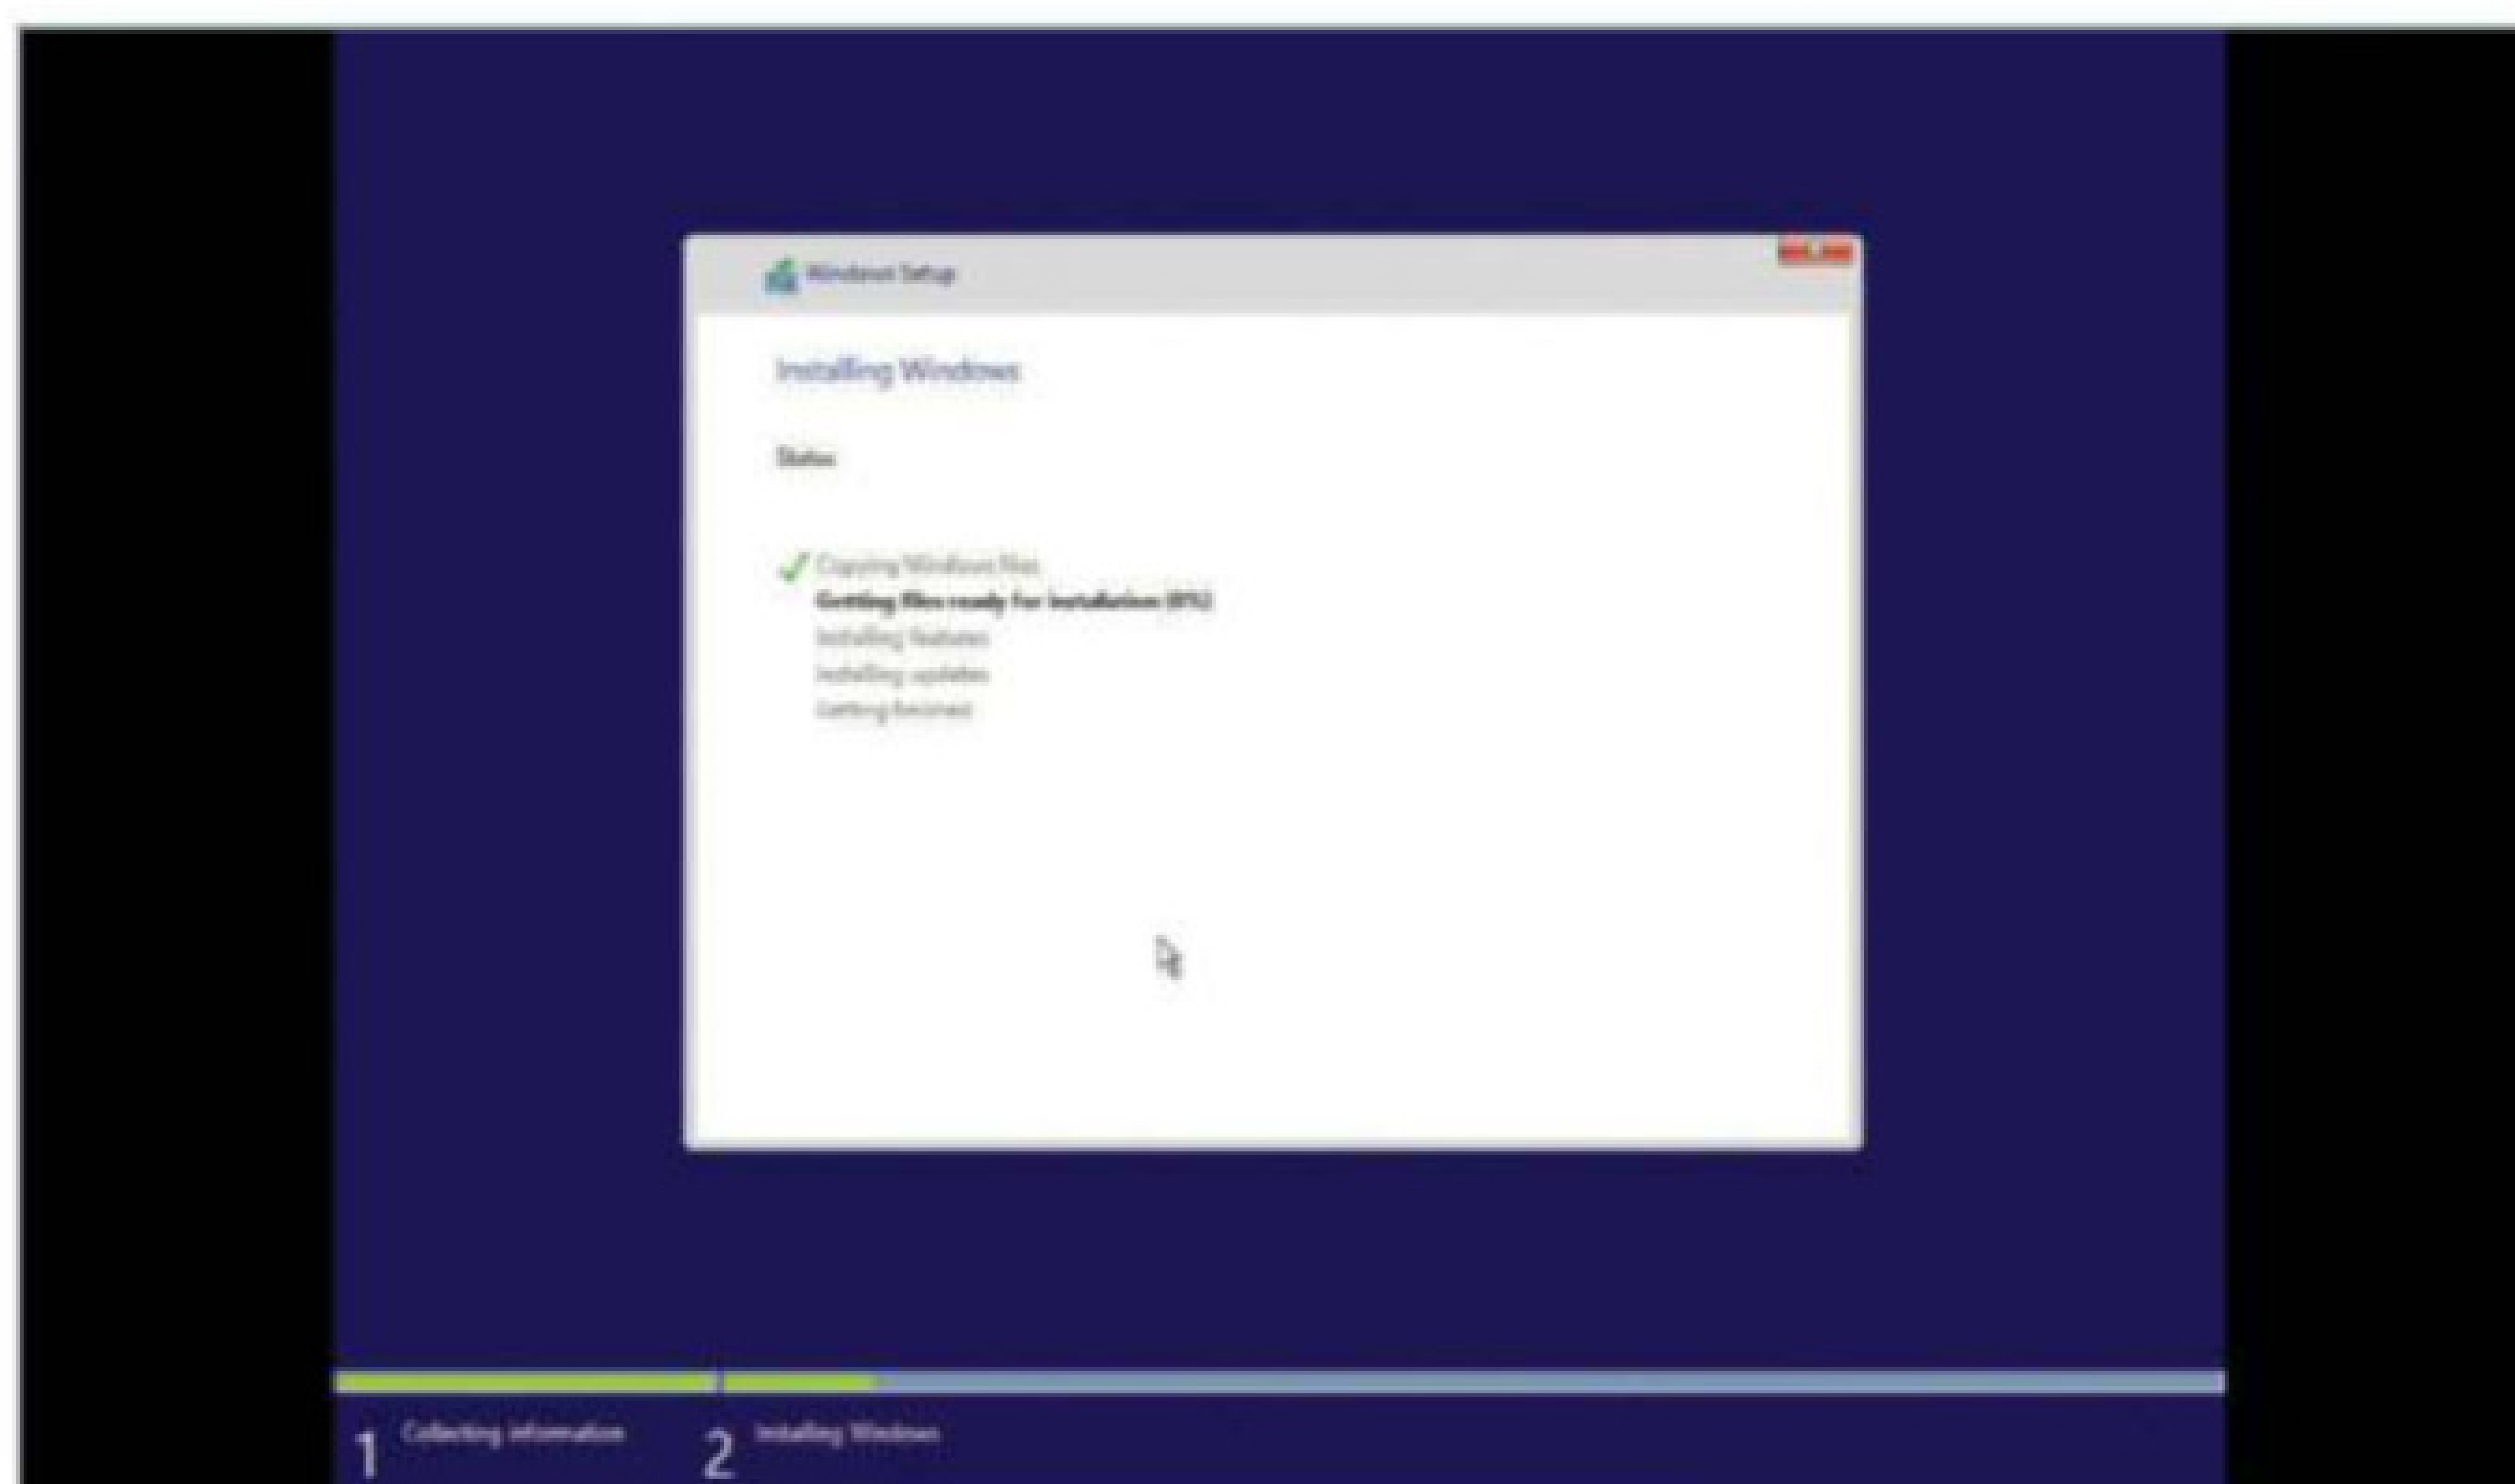
01 Prepare to install

Put the disc in and boot your PC. The first step in installing is to set the language as you normally would. Click install and then agree to the licence. After this, you'll be asked how you want to install; click on the Custom install option to install from scratch.



02 Storage

Select the empty partition we created before in GParted as the place to install Windows to. It will recognise it as NTFS, with the Linux partition as unknown – it will also completely reformat this partition once again, so make sure there's nothing on there.



Above Feeling adventurous? Grab and burn the Windows 10 Technical Preview from bit.ly/1y8MoE2

03 Wait a while

Windows can take a while to install and will go through several phases, including rebooting once or twice during the process. Leave it alone and it will do its thing without any interruptions. Now is a great time to go and make yourself a cup of tea.




04 Personalise your install

On Windows 10 (or Windows 8.1, if you're playing it safe), you'll now need to make some basic settings for your data and the way that your system should work. Go through the wizard with your own choice of settings before entering in your account details.

05 Access a microsoft account

You'll need to log into or set up a Microsoft-based account in order to use the latest versions of Windows. If you already have one then you can enter it here and log in. Otherwise, you'll need to link an email address to a new account.



06 Final setup

Wait a while and the wizard will grab your account settings and any other data you may have associated with a Windows 10 installation under your Microsoft account. After this, it will bring up the desktop and allow you to start using it.

Shrink partitions

The NTFS file system isn't suited to editing its space use but there are ways around it

“You need to be careful when shrinking the partition”

Defrag Linux?

While ext3 and ext4 are much more efficient, they can still exhibit fragmentation. Shrinking of the partition might still corrupt files if parts of them are located towards the end of the disc, although it's much more unlikely. If you're shrinking right down to the line, you're asking for trouble. Linux defragmenters are few and far between though (however defragfs is worth a try: sourceforge.net/projects/defragfs), so one of the best methods is to temporarily move all of your personal files onto portable storage, while setting up the partitions and such, and be conservative. Again it is worth mentioning to always back up important files on these kinds of tutorials anyway.



Above Defragfs gives you quite detailed and useful output in the terminal


If you're starting with Windows already installed onto your system, then you'll need to start making space so that you can install your other distros. This means shrinking the install partition on Windows, which almost always by default includes the partition that carries all of your files and documents and such – the contents of which used to be known as My Documents.

Therefore, you may need to start moving files off your computer before shrinking partitions and re-organising your hard drive. As we've shown on the partitioning page, you should allow for more space on Windows anyway and you can even use its main partition as your general storage space rather than keeping it completely separate.

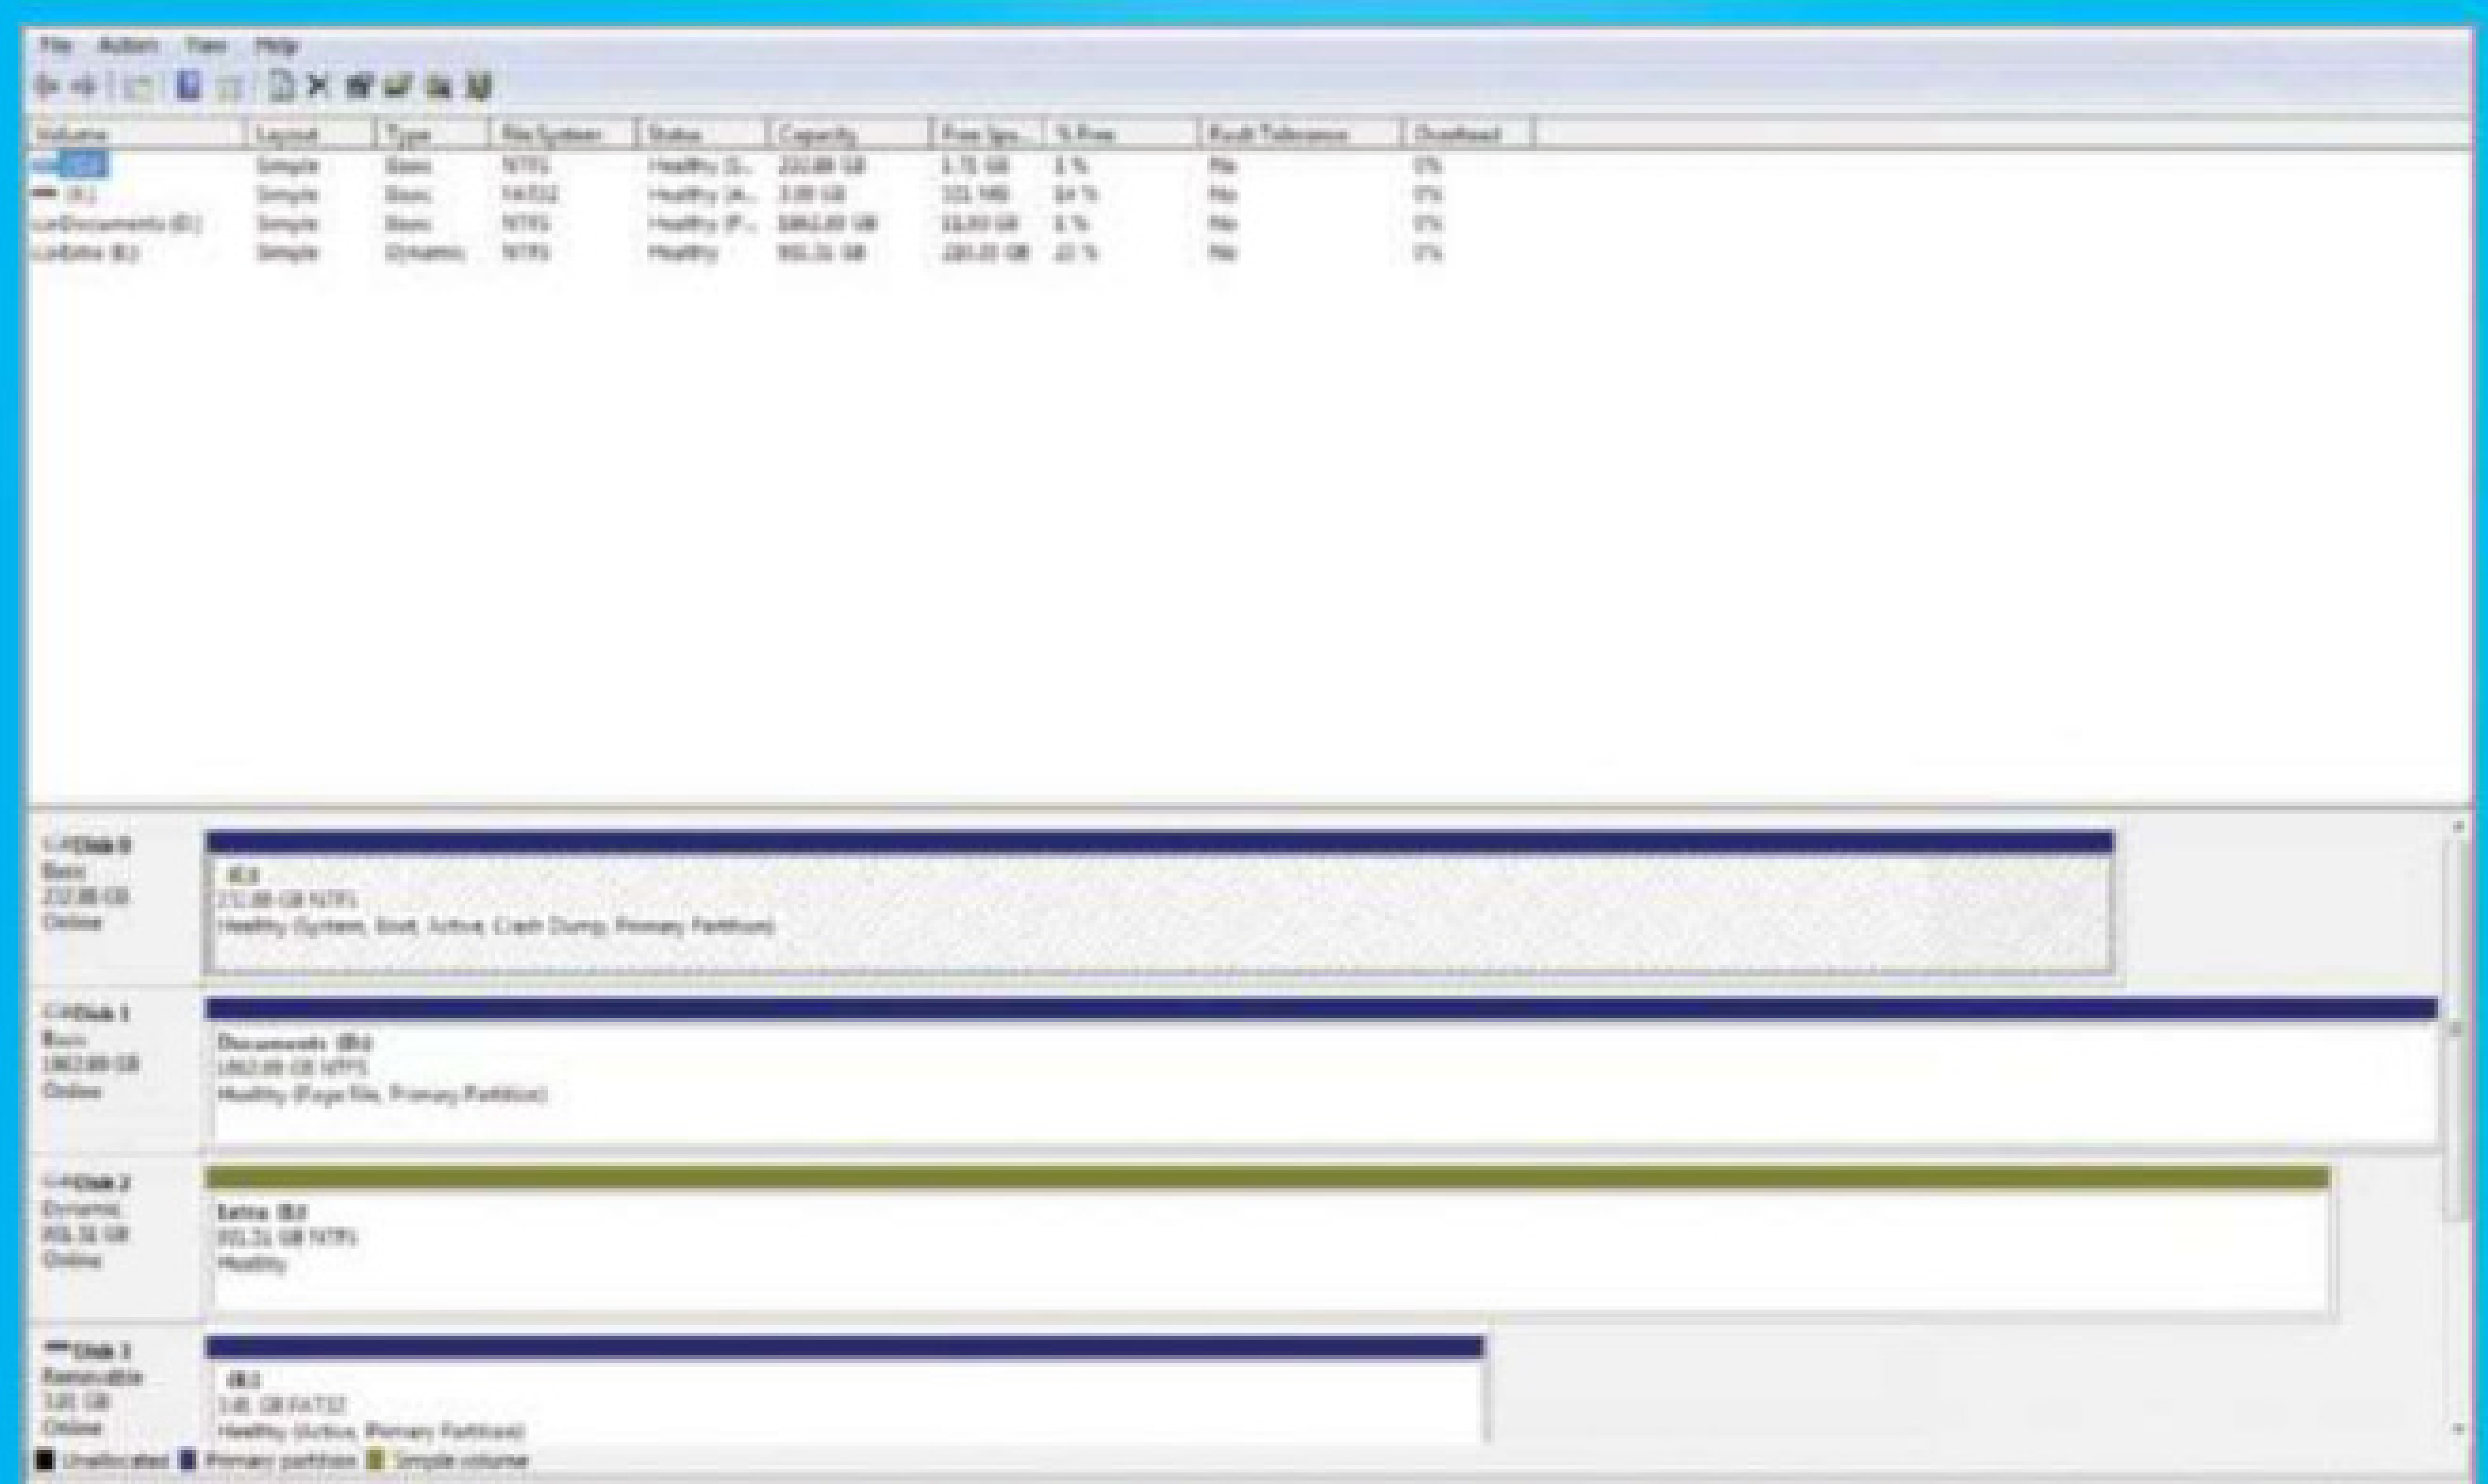
Either way, you need to be careful when shrinking the partition – files usually become fragmented and split themselves up across the disc over the course of many writes and rewrites and moving around. In the days of XP this could become quite ridiculous, however since then Microsoft have implemented auto-defragmentation systems into Windows that generally keep it more tidy on a disc level.

It's not perfect though, and you may feel the need to perform a defrag before shrinking the partition. For this we recommend UltraDefrag (bit.ly/168iFAR), which is an open source defragging application. You can make it perform an active defrag that will remove files from the end of the partition – where we will be reclaiming space – and even perform a boot-time defrag so that system files can be moved as well.

Once that is done, you can use GParted to start setting your disc up. Do not move the Windows partition though as it will completely ruin your installation. Instead, keep it where it is at the start and install everything around it.



Above We recommend using UltraDefrag if you want to have a pass at defragmenting



Above Windows has tools for this in Windows Vista and up – you can use those over GParted

Editing GRUB

The boot menu is automatically created by your distros, so it may need some tidying

Modern GRUB, GRUB 2, is very smart. Merely installing it as part of most distros has it look to see what else is on your system and add it to the boot menu. If you've used the GRUB recovery pages in this guide, you'll know that these can be updated at any time to include new distros.

There is a lot more you can do with GRUB though, from simply changing the default selection on the boot menu to customising the naming and placement of operating systems on it. After each update you can save the changes with `sudo update-grub`.

01 Default selection and time-out

On the boot screen the default selection will be the first distro in the list. This will automatically be selected when the timer hits zero. As a quick way of changing the default to better suit your needs, in a terminal use `nano /etc/grub.d/00_header` and then search for the following two lines:

```
GRUB_DEFAULT=0
GRUB_TIMEOUT=5
```

02 Manual order change

After every `update-grub`, the `grub.cfg` file is updated, usually located in `/boot/grub/grub.cfg`. Open it up with `nano` and scroll down

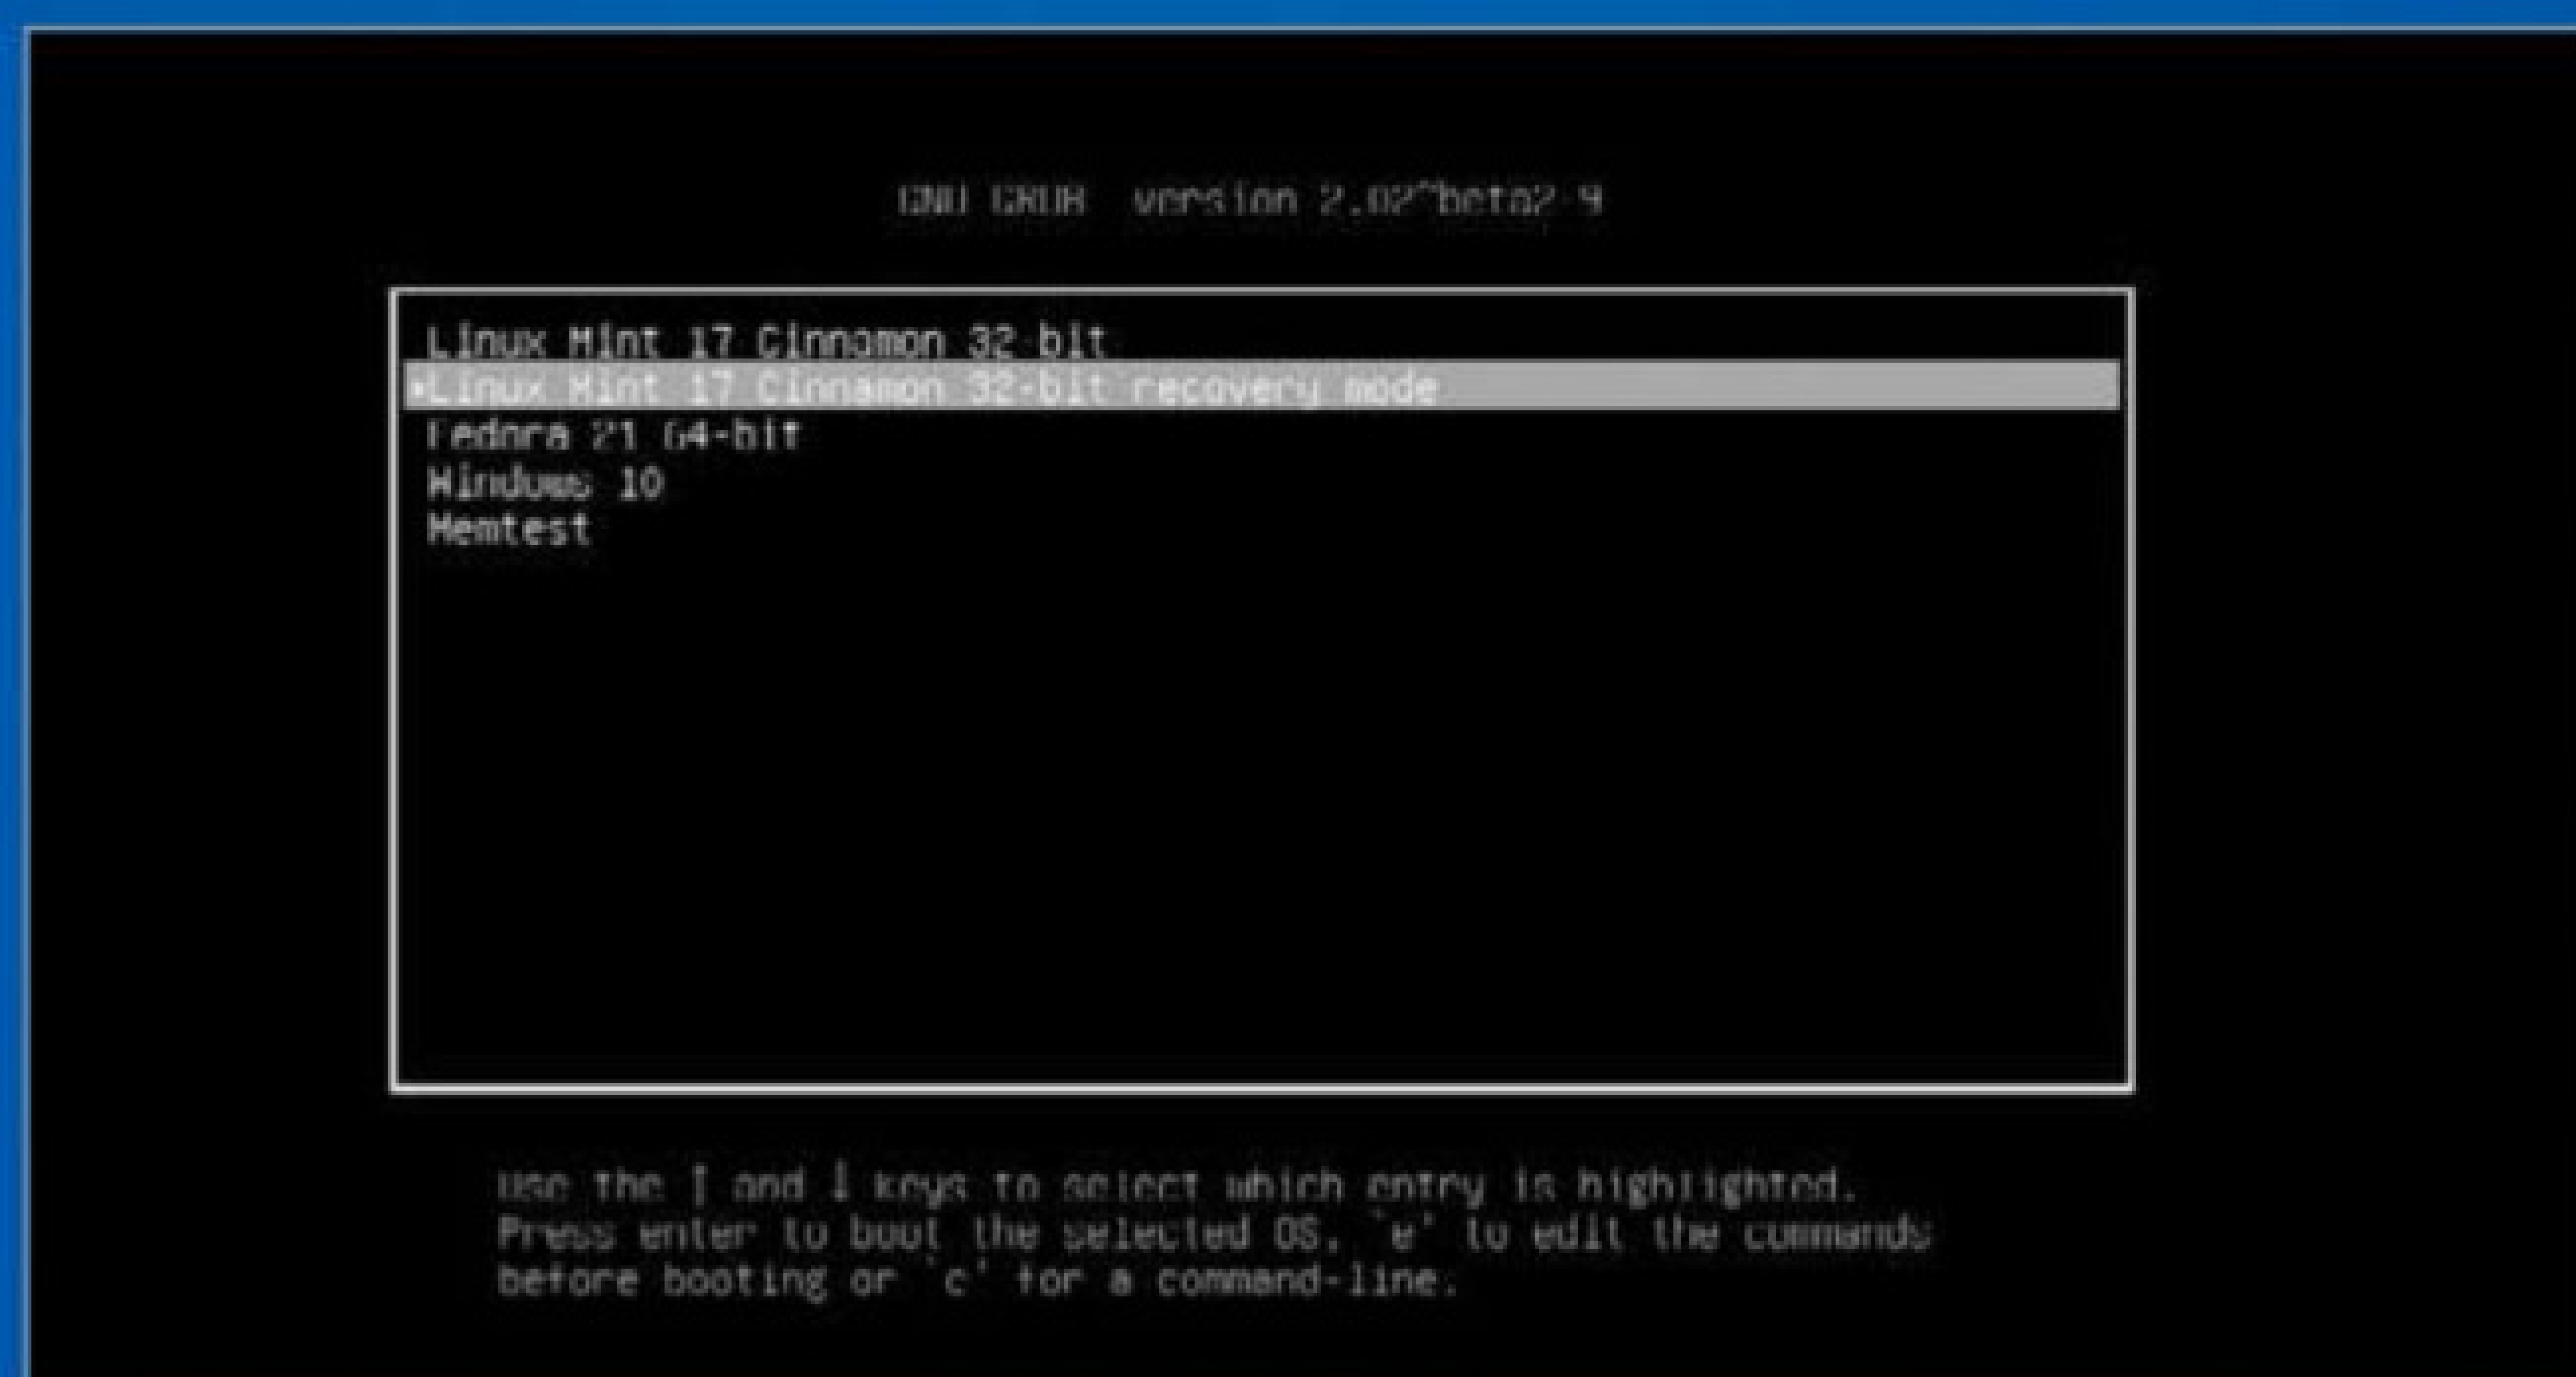
to see the default and time-out changes we made, as well as the individual boot scripts. You can manually move these around in the `cfg` and save it, but it will be overwritten every time you do an `update-grub`.

03 Quick order change

The order of boot menu placement depends on the number of the files located in the `/etc/grub.d` folder. The Linux you're using will have a custom script titled `10_linux`, whereas everything else will be discovered using the `30_os-prober` list. Changing the number on the two will move them around in the list, for example `09_os-prober` will occur before `10_linux`.

04 Best order change

The most effective and permanent method of changing the order is to create custom scripts for your three distros and order them properly in the `grub.d` folder. Other Linux distros will use a very similar script to the `10_linux` file and you can find a template in `40_custom`. Windows is done slightly differently to boot into its chainloader. Once you've got these setup, you may need to do some maintenance on the scripts every few months, but it should keep your GRUB menus in perfect order.



UEFI

When Windows 8 launched there was a lot of furore over UEFI and secure boot. For good reason though, as secure boot would not allow you to install other operating systems alongside Windows. As most PCs and laptops come with Windows as standard, this meant that it would cause major problems for Linux users. Luckily, distros began to adapt and implement software so that even though motherboards still had UEFI and secure boot, they would be able to boot without too much of a problem.

These days there's not a huge problem in installing Linux alongside Windows as most distros have a solution in place, and you can easily deactivate secure boot to get the installation underway. If you do come across any issues though, Google should help you out right away.



“ Windows is done slightly differently to boot into its chainloader ”

Styling GRUB

If you want to get really fancy, you can always theme GRUB

GRUB isn't, by default, the prettiest thing in the world, so if you're going to be running with your triple boot setup for a while then you might want to give it a little polish.

There are some great guides out there in the wild that walk you through the entire process, including how to set the option titles, the splash image, plus the colours of each element as well as the fonts. Explained in a little more detail just to the right, one great site to check out for this is dedoimedo.com.



Above Adding custom backgrounds to GRUB2 is much simpler than you'd think



Above You can also customise the menu options themselves with fonts and colours

Resources

If you need to know more about booting, GRUB and the different distros, try these resources

DistroWatch distrowatch.com

Like the idea of triple booting your system but can't quite figure out exactly what distros to use? DistroWatch compiles one of the most complete list of Linux distributions on the Internet. It keeps tabs on the updates and release cycles of all the major distros, and also has archives of all the update news for each of the listed distros. Every one has its own categories and a little explanation so you can figure out if the distro is what you're looking for before trying out a live disc.

There is a ranking table of distros that seem to be popular on the site, which may help you discover new and excellent operating systems, and an upcoming release schedule so you can plan what distros to get ready to install in advance.

UNetbootin unetbootin.sourceforge.net

Who has DVDs lying around to burn images to these days? Well... we do! We have a load left that we might need in the future (you never know), but in general we still prefer to use a bootable USB stick to create our live media. Especially when not every computer has a disc drive any more.

UNetbootin is the perfect piece of software to do this and it works across all platforms. It quickly writes ISOs to a bootable USB stick and also has the ability to download a preset selection of distros if you haven't hunted down an ISO yet. It can also add a little bit of re-usable storage to Ubuntu-based distributions. Otherwise, you can simply navigate to the location of an ISO on your system and write the USB from there.



Dedoimedo GRUB 2 tutorials bit.ly/1yoXSk1

Featured in the distant past of issue 117 of this magazine, a full GRUB 2 reference site can be found at dedoimedo.com. It includes a much more in-depth discussion of how GRUB 2 works, how the config file is built and used, the different directories and more. There is also a lot more info on creating your own custom boot scripts for different distributions, a few more recovery tips and ways to customise the look of the boot menu.

It is kept up to date with the latest GRUB 2 changes, so even if you have a problem in the future and need some help, it should remain an excellent resource to go to.

StoreMags

NEXT GENERATION 1&1 CLOUD SERVER

Easy to use – ready to go!

The new 1&1 Cloud Server offers all the advantages of dedicated hardware performance combined with the flexibility of the cloud!

FLEXIBLE & AFFORDABLE

Customised configuration

- SSD, RAM and CPU are all independently flexible, and can be adjusted exactly to your requirements



Transparent costs

- **NEW:** Billing by the minute
- **NEW:** The clearly structured cost overview enables efficient planning and management

EASY & SECURE

1&1 Cloud Panel

- **NEW:** The innovative, user-friendly interface with smart administration, simplifies the management of your server

Security

- The high-performance 1&1 Data Centres are among the safest in Europe
- Backups and snapshots prevent accidental data loss
- The integrated firewall protects your server from online threats

ALL-INCLUSIVE

Top performance


- **NEW:** Set up your cloud server in under 1 minute
- **NEW:** Premium SSD with the highest possible performance
- **NEW:** Private networks, professional API, load balancers, firewalls and many more easy-to-configure server features
- **NEW:** Virtualisation with the leading technology from VMware®
- **NEW:** Ready-to-use applications included: WordPress, Drupal™ and Magento®
- Parallels® Plesk 12
- Unlimited traffic



Storemags DOMAINS | MAIL | HOSTING | eSHOPS | SERVERS



Enter your e-mail address to get started



TRY FREE FOR 1 MONTH!*

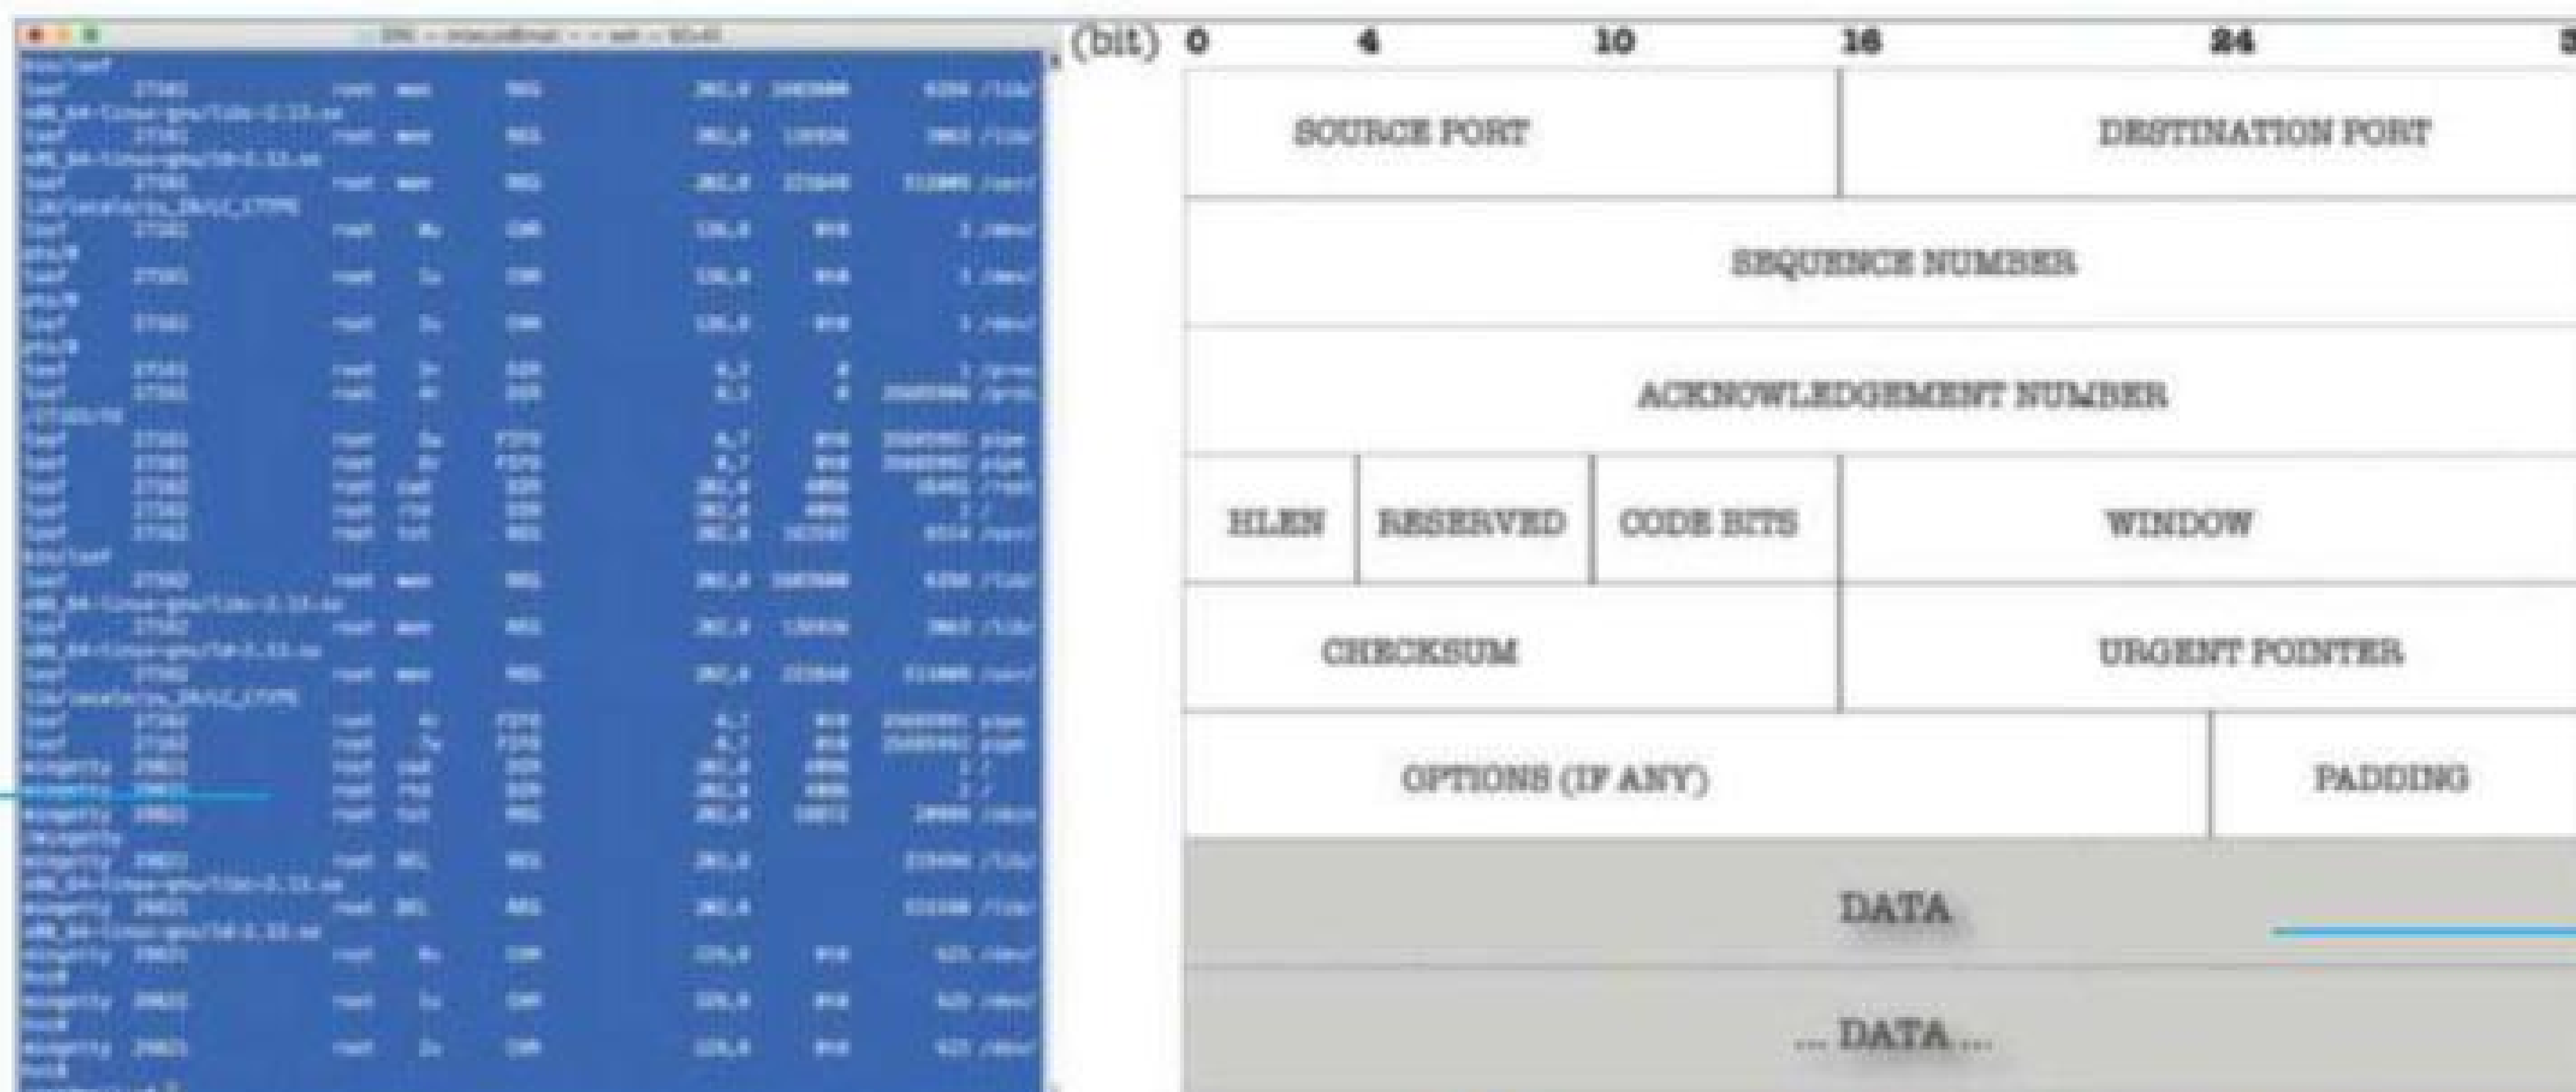
☎ 0333 336 5509

*1&1 Cloud Server 1 month free trial with no payment details required. Visit www.1and1.co.uk for full offer details, terms and conditions. Windows is a registered trademark of Microsoft® Corporation. Linux is a registered trademark of Linus Torvalds.



1and1.co.uk

Running **lsof** without command line parameters produces a lot of output, so it can be hard to find information



This is the format of a TCP packet, which is very useful to know when troubleshooting network problems

Here you will find the general format of an IP network packet



Running **lsof -i** lets you to list all network to get the information you want

Get your network information using lsof

Advisor



Mihalis Tsoukalos is a UNIX administrator, a programmer (UNIX and iOS), a DBA and a mathematician. He has been using Linux since 1993

Learn how to utilise **lsof** to acquire important network related information about your Linux machine

Resources

Lsof
Root privileges

Lsof (list open files) is a handy Linux utility created by **Victor A. Abell**. However, don't think that it is only about files because, as you may already know, UNIX handles every device as a file. Since all UNIX devices are files, the border between regular files, directories and network interfaces is very blurred.

Networks often fail because software and hardware stop working – the most difficult part is finding out exactly where the problem resides. This tutorial will present various handy **lsof** usages with the confidence that they will help you during network troubleshooting.

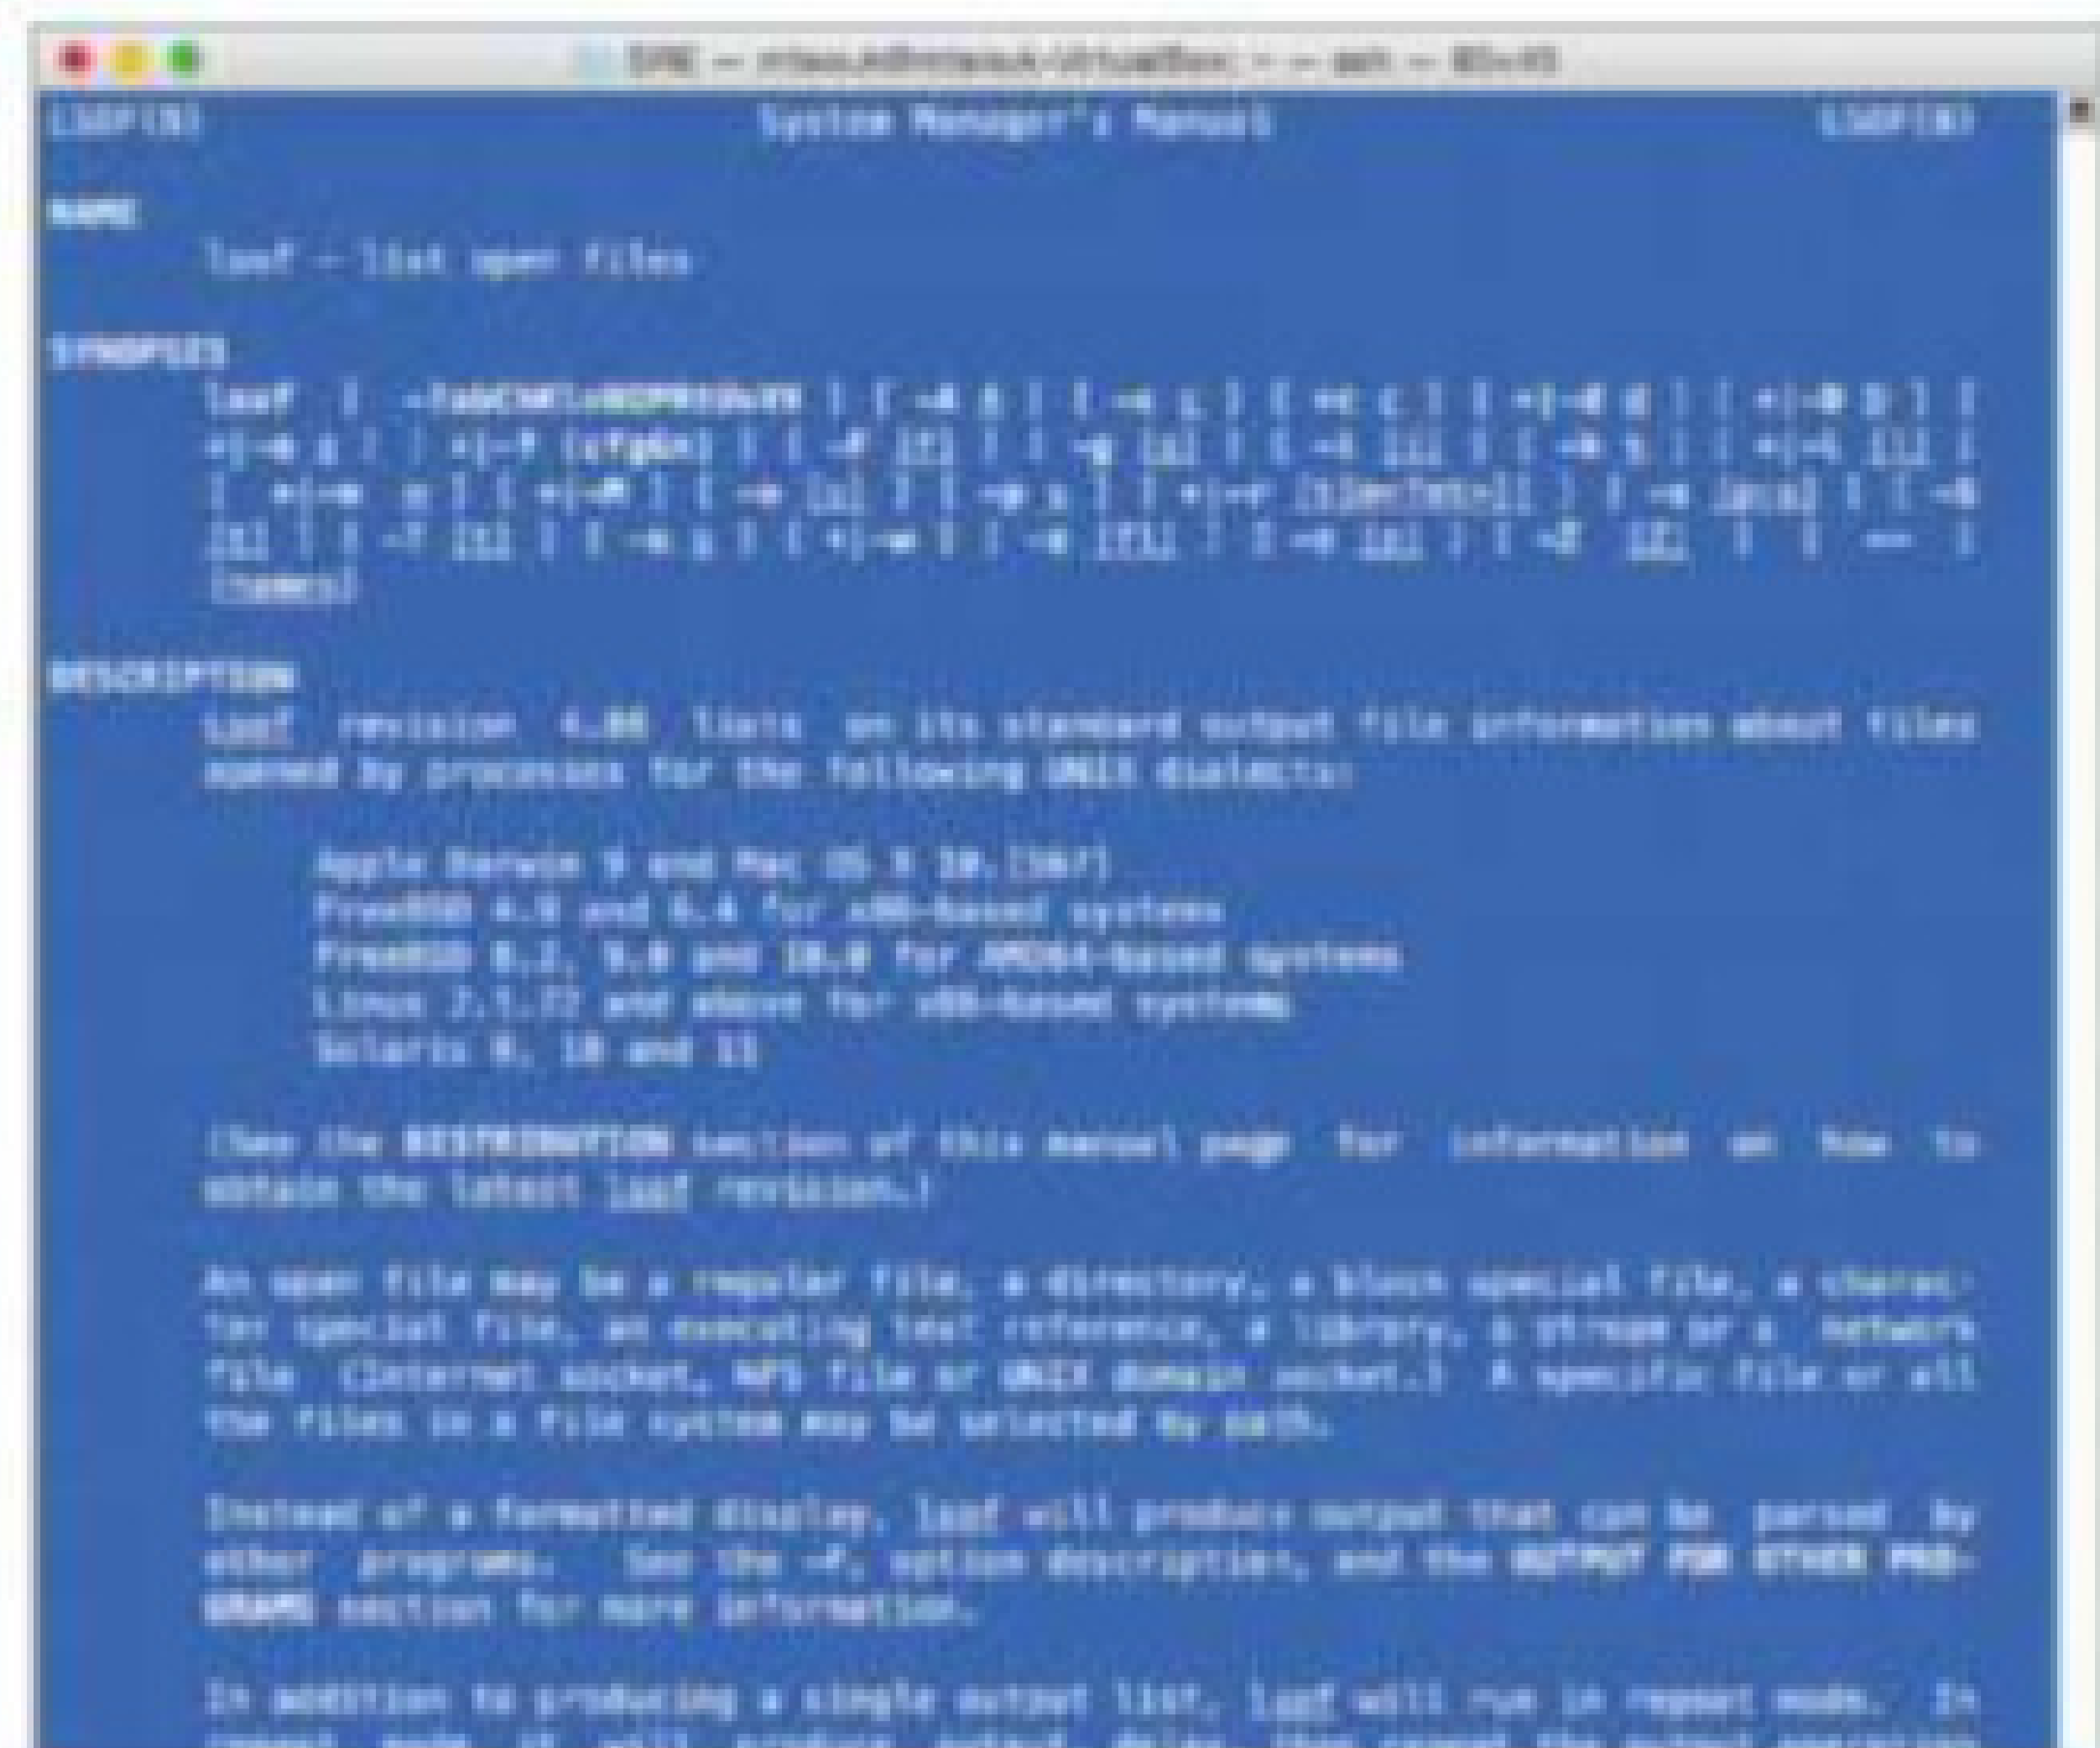
The two main drawbacks of **lsof** are that it can only display information about the local machine (aka localhost) and that it needs administrative privileges in order to run. Usually, you don't execute **lsof** without any command line parameters because you will find its output busy and difficult to read:

```
# lsof | wc
6759 64564 846201
```

It is useful to pay attention to the command line options of **lsof** and use them as needed.

Get network information using lsof

Troubleshoot your network by identifying its problems



01 About lsof

When you're in doubt about which command line option to use, the first place to look for help should be the **man** page of lsof. If you want to have the **man** page as a separate file, you can convert it into Postscript format by executing the **man** command using the following options:

```
$ man -t ps lsof > lsof.ps
```

02 What to ask when you have a network problem

When you have networking problems, many questions may come up such as: is there a problem with the local network or the server? Is the problem DNS related? Is there anything wrong with routing? Is it a hardware or software issue? Is the Internet connection working?

Utilities such as lsof can help you answer some of these questions and find and correct the root of the problem.

03 Which program listens to a TCP port?

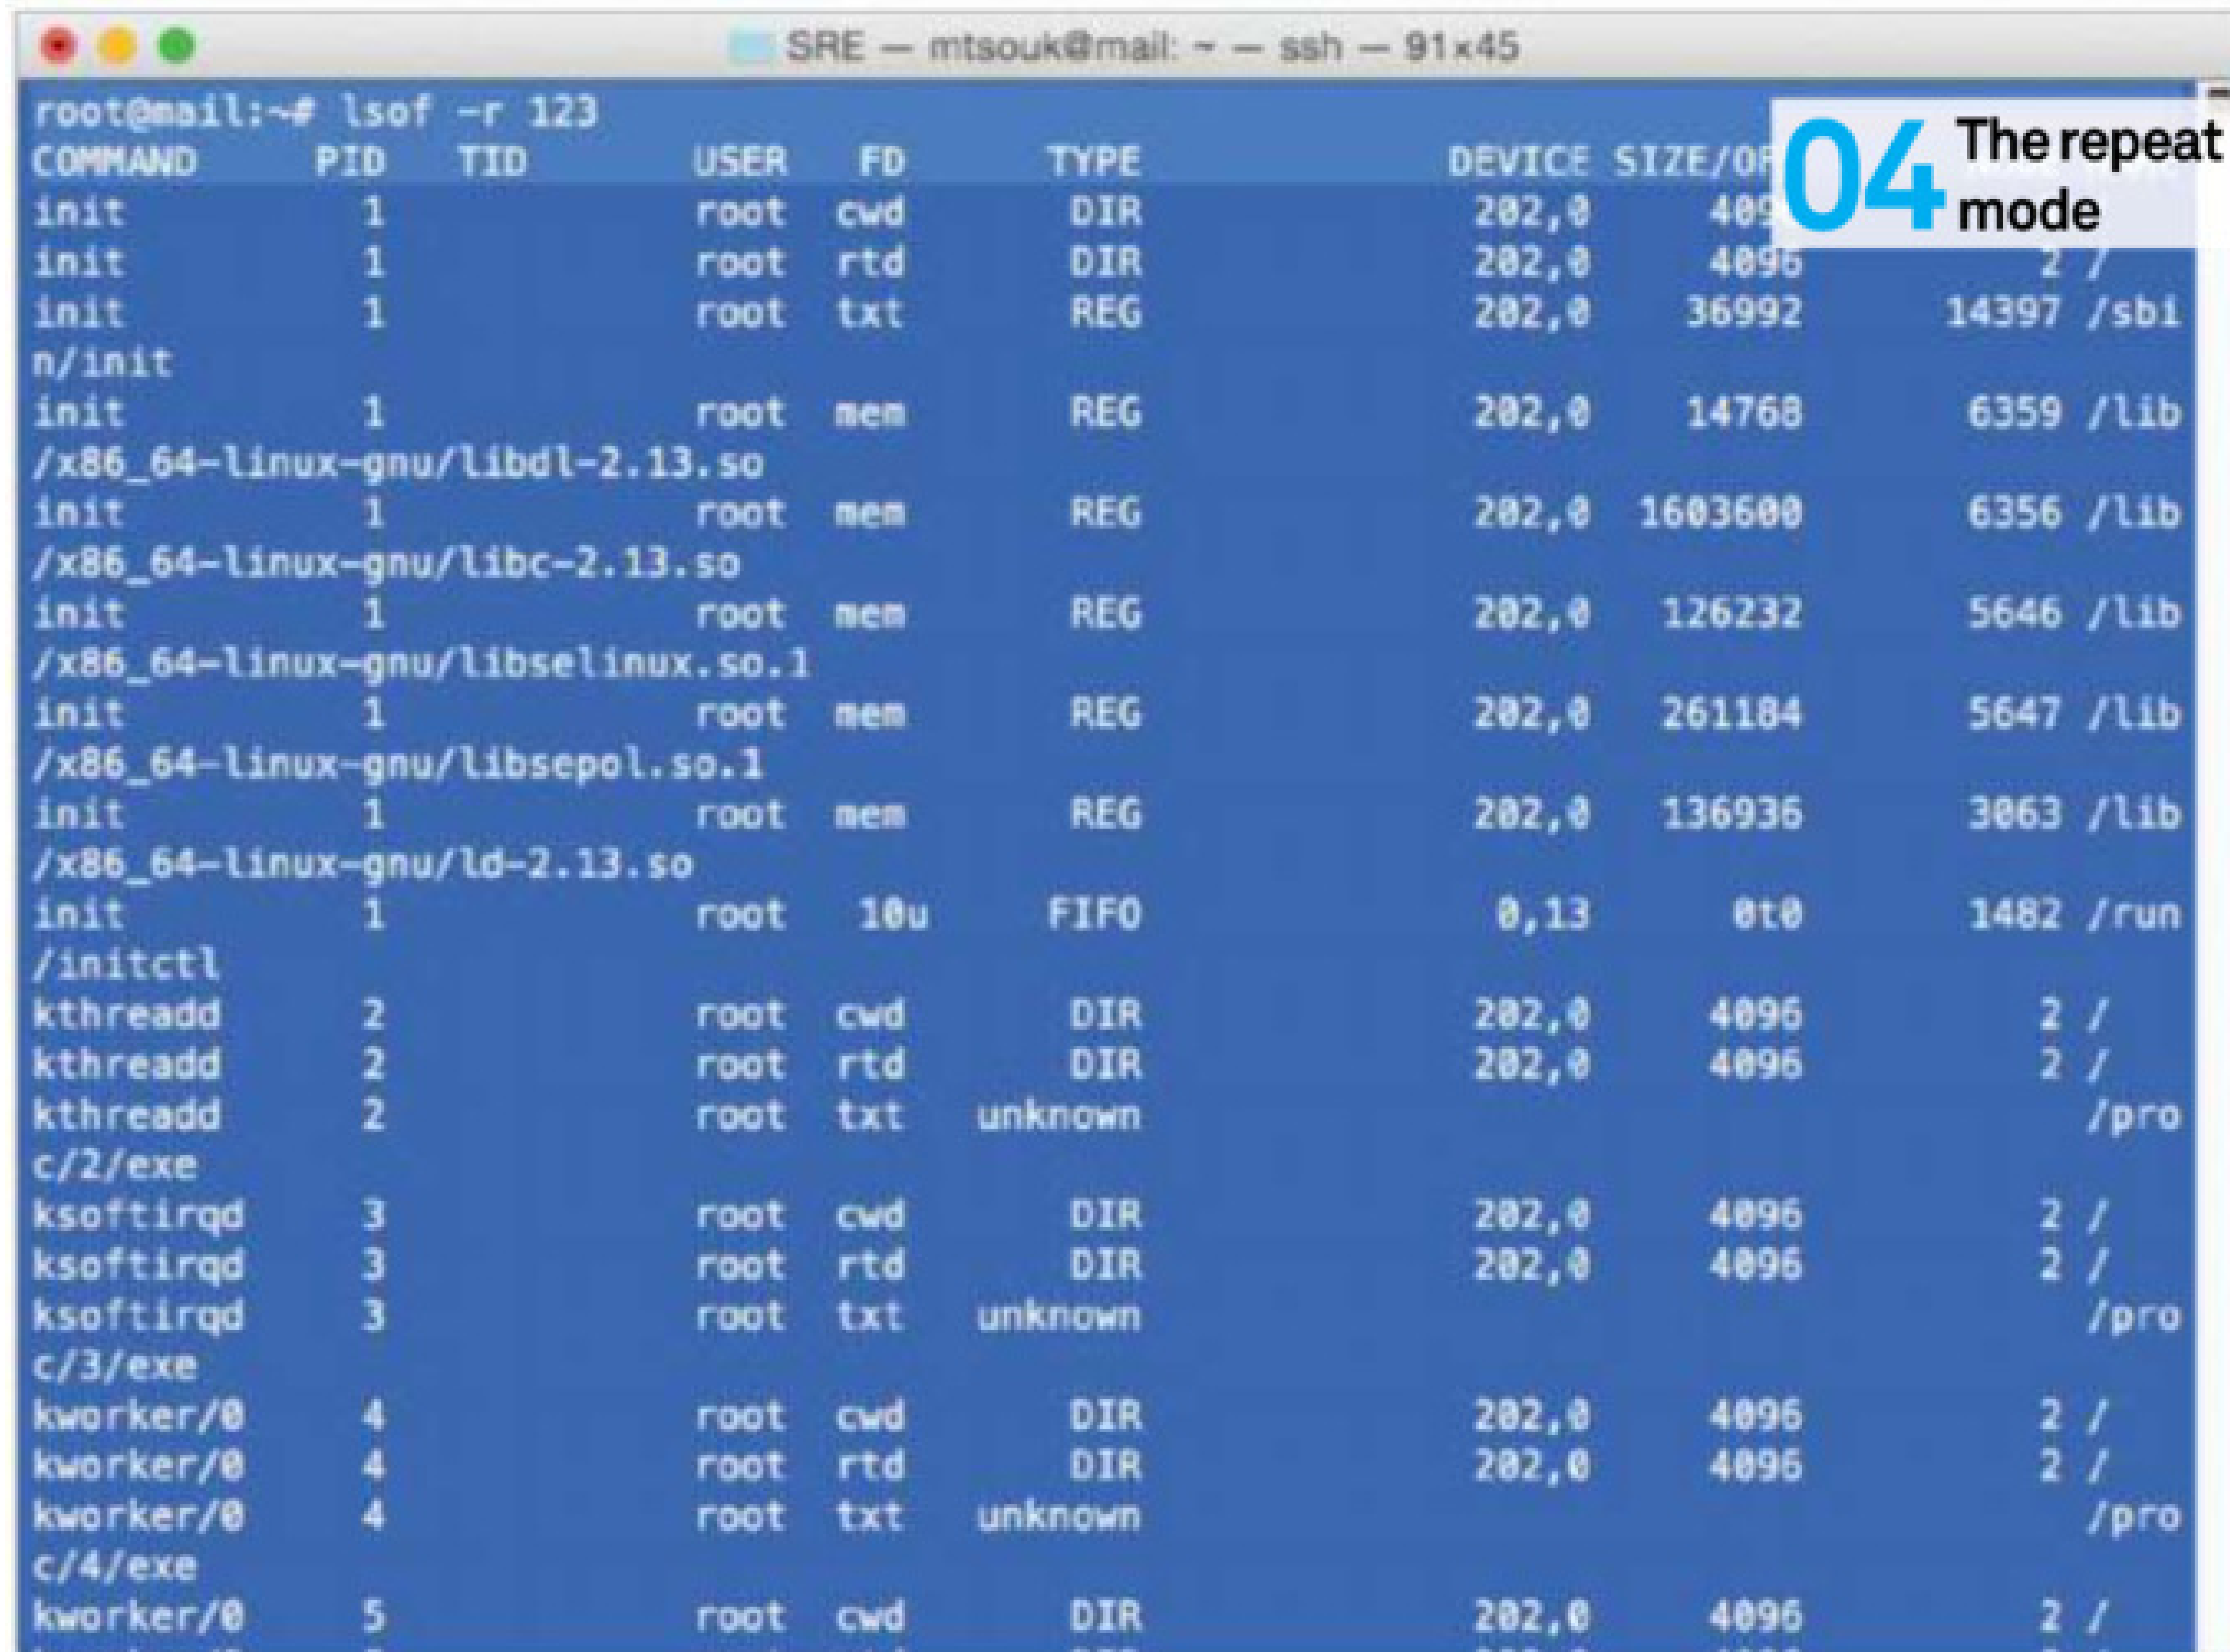
The **-i** option tells lsof to only display network related information. Therefore, the following command shows all Internet network files:

```
# lsof -i
```

The first question you will usually want answered is which server program occupies a TCP port on your Linux machine. To find out which program listens to a TCP specific port, run the following command:

```
$ sudo lsof -nP -iTCP -sTCP:LISTEN
```

Other possible states of a TCP connection are CLOSED, SYN-SENT, SYN-RECEIVED, ESTABLISHED, CLOSE-WAIT, LAST-ACK, FIN-WAIT-1, FIN-WAIT-2, CLOSING and TIME-WAIT.

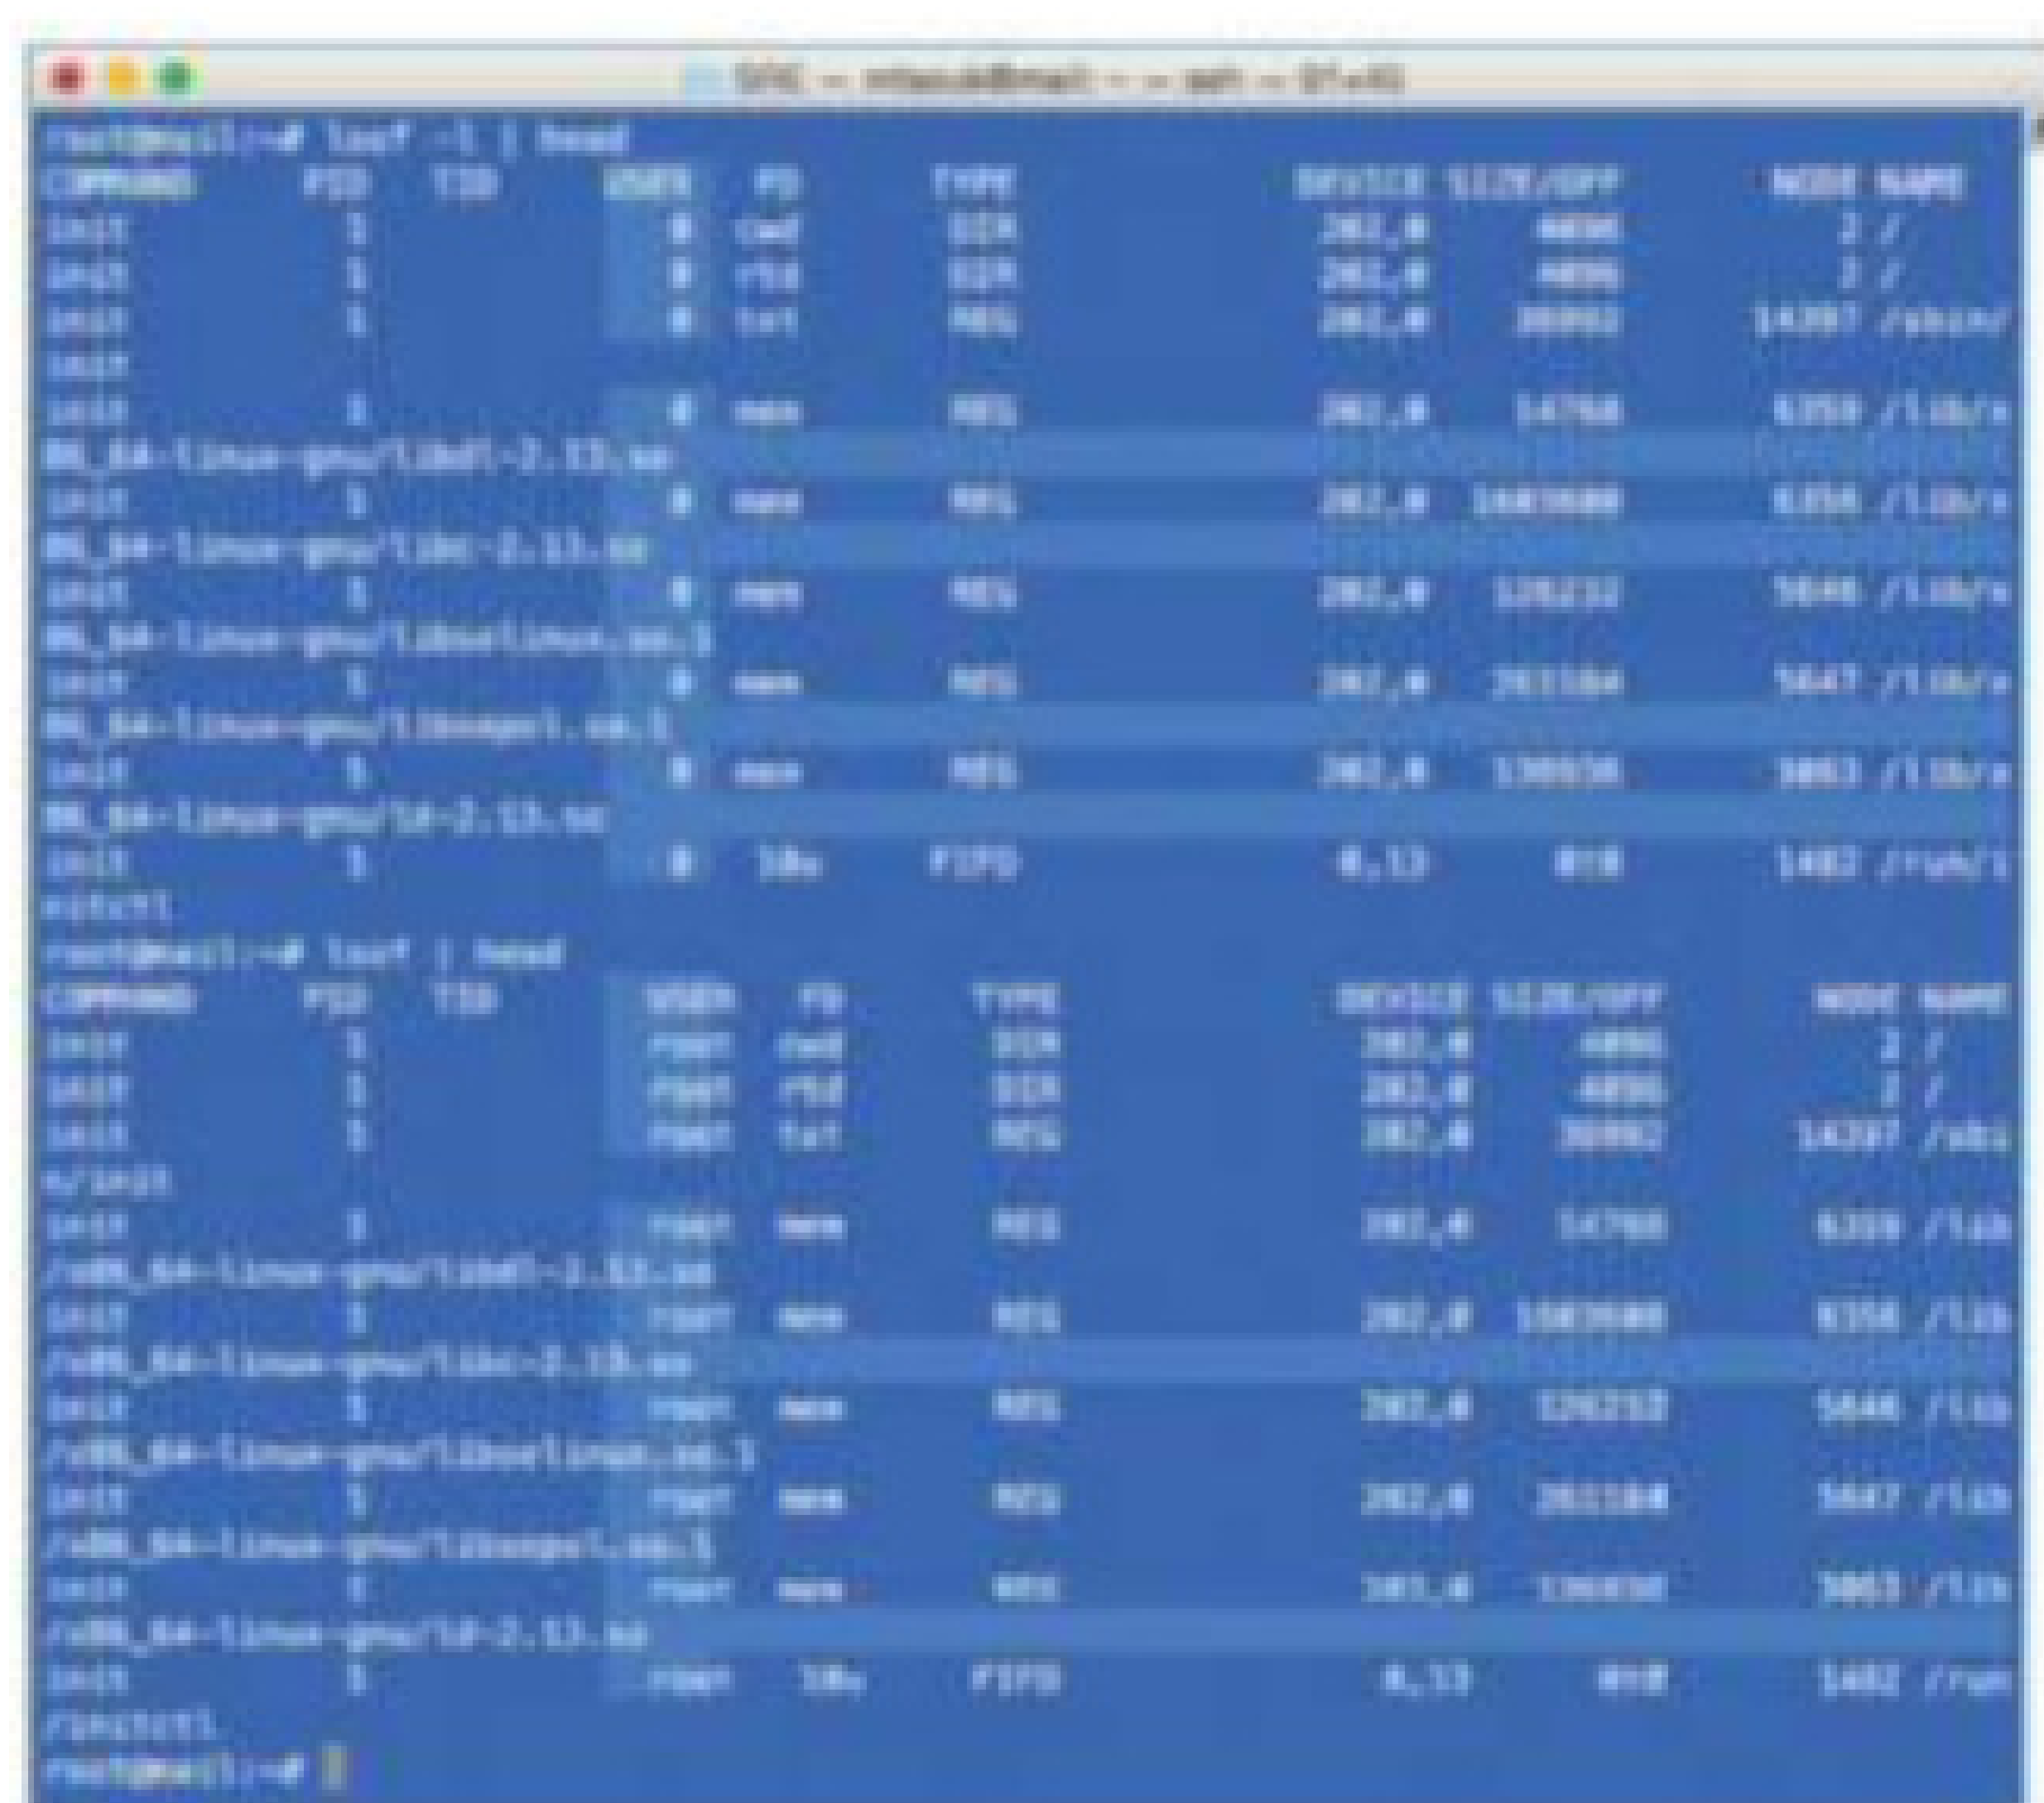


04 The repeat mode

04 The repeat mode

Running lsof with the **-r** option puts lsof into a repeat mode which runs forever, so you'll need to terminate lsof manually. The **+r** option will also put lsof in repeat mode – the difference is that it will automatically terminate lsof when a circle has no new output to print.

When lsof is in repeat mode, it prints new output every 't' seconds (a circle) – the default value of t is 15 seconds. You can change the default value by typing an integer after **-r** or **+r**.



05 User ID numbers to login names

If converting user ID to login name is working improperly, you can disable it using the **-l** parameter. This option can be combined with a range of other options for providing additional speed to the output.

06 Find all established SSH connections

The following command finds all established SSH connections to the local machine:

```
# lsof | grep sshd | grep EST
```

The following slightly different version will present exactly the same information:

```
# lsof -i | grep ssh | grep EST
```

The second version is quicker because the **-i** option makes lsof print less information.

07 Which processes are accessing a file?

You can find the processes that are accessing **/var/run/mysqld/mysqld.sock** by running the following command:

```
# lsof /var/run/mysqld/mysqld.sock
```

If you are only interested in the process ID of the processes accessing a file, use the **-t** option:

```
# lsof -t /var/run/mysqld/mysqld.sock
```

The process ID can be used for killing a process.

Get network information using lsof

Troubleshoot your network by identifying its problems

```
root@mail:~# lsof -c /^.$/
COMMAND  PID USER  FD   TYPE    DEVICE  SIZE/OFF      NODE NAME
md        222 root   cwd   DIR     202,0    4096        2 /
md        222 root   rtd   DIR     202,0    4096        2 /
md        222 root   txt   unknown /proc/222/exe
su        25096 root   cwd   DIR     202,0    4096      50899 /home/mtsouk
su        25096 root   rtd   DIR     202,0    4096        2 /
su        25096 root   txt   REG     202,0   36816     16480 /bin/su
su        25096 root   mem   REG     202,0   18672     5599 /lib/x86_64-linux-gn
u/libattr.so.1.1.0
su        25096 root   mem   REG     202,0   17112     5604 /lib/x86_64-linux-gn
u/libcap.so.2.22
su        25096 root   mem   REG     202,0    8160    1220656 /lib/x86_64-linux-gn
u/security/pam_cap.so
su        25096 root   mem   REG     202,0    6000     5947 /lib/x86_64-linux-gn
u/security/pam_permit.so
su        25096 root   mem   REG     202,0    5912     5927 /lib/x86_64-linux-gn
u/security/pam_deny.so
su        25096 root   mem   REG     202,0   35104     6358 /lib/x86_64-linux-gn
u/libcrypt-2.13.so
su        25096 root   mem   REG     202,0   56056     5962 /lib/x86_64-linux-gn
u/security/pam_unix.so
su        25096 root   mem   REG     202,0   10208     5942 /lib/x86_64-linux-gn
u/security/pam_mail.so
su        25096 root   mem   REG     202,0   14360     5929 /lib/x86_64-linux-gn
u/security/pam_env.so
su        25096 root   mem   REG     202,0  126232     5646 /lib/x86_64-linux-gn
u/libselinux.so.1
su        25096 root   mem   REG     202,0    6016     5950 /lib/x86_64-linux-gn
u/security/pam_rootok.so
su        25096 root   mem   REG     202,0   47616     6462 /lib/x86_64-linux-gn
u/libnss_files-2.13.so
su        25096 root   mem   REG     202,0   43560     6478 /lib/x86_64-linux-gn
u/libnss_nis-2.13.so
su        25096 root   mem   REG     202,0   89056     6410 /lib/x86_64-linux-gn
u/libnsl-2.13.so
su        25096 root   mem   REG     202,0   31584     6413 /lib/x86_64-linux-gn
u/libnss_compat-2.13.so
su        25096 root   mem   REG     202,0   14768     6359 /lib/x86_64-linux-gn
u/libdl-2.13.so
su        25096 root   mem   REG     202,0  1603600    6356 /lib/x86_64-linux-gn
u/libc-2.13.so
```

16 Use regular expressions

16 Use regular expressions

Lsof supports regular expressions. In order to find all commands with exactly two characters, you can execute this command:

```
# lsof -c /^.$/
```

Regular expressions start and end with a slash. The ^ character denotes the beginning of a line whereas \$ denotes the end of one. Each dot represents a single character.

17 Choose between IPv4 and IPv6

Lsof lists both IPv4 and IPv6 connections by default but you can choose the kind of connections you want to display. The next command displays IPv4 connections:

```
# lsof -i4
```

The following command displays IPv6 connections only:

```
# lsof -i6
```

```
root@mail:~# lsof -i4 | head
COMMAND  PID USER  FD   TYPE    DEVICE  SIZE/OFF      NODE NAME
tcpdump  1197 root   4u   IPv4  25466450  000  TCP *:*:80 (LISTEN)
vshd     2784 root   3u   IPv4  2343283  000  TCP *:*:22 (LISTEN)
vshd     2784 root   4u   IPv6  2343283  000  TCP *:*:22 (LISTEN)
dnlClient 2980 root   5u   IPv4  4817  000  UDP *:*:53
dnlClient 2980 root   99u  IPv4  4804  000  UDP *:*:53
dnlClient 2980 root   23u  IPv4  4805  000  UDP *:*:53
nfsd     4872 nfsd   14u  IPv4  31363731 000  UDP *:*:111
nfsd     4872 nfsd   17u  IPv4  31363732 000  UDP *:*:111
nfsd     4872 nfsd   18u  IPv4  31363730 000  UDP localHost:111
COMMAND  PID USER  FD   TYPE    DEVICE  SIZE/OFF      NODE NAME
tcpdump  1197 root   4u   IPv4  25466450  000  TCP *:*:80 (LISTEN)
vshd     2784 root   3u   IPv4  2343283  000  TCP *:*:22 (LISTEN)
vshd     2784 root   4u   IPv6  2343283  000  TCP *:*:22 (LISTEN)
dnlClient 2980 root   5u   IPv4  4817  000  UDP *:*:53
dnlClient 2980 root   99u  IPv4  4804  000  UDP *:*:53
dnlClient 2980 root   23u  IPv4  4805  000  UDP *:*:53
nfsd     4872 nfsd   14u  IPv4  31363731 000  UDP *:*:111
nfsd     4872 nfsd   17u  IPv4  31363732 000  UDP *:*:111
nfsd     4872 nfsd   18u  IPv4  31363730 000  UDP localHost:111
```

18 Disable DNS and resolve port number

Lsof uses the data found in the /etc/services file to map a port number to a service. You can disable this functionality by using the -P switch as follows:

```
# lsof -P
```

Lsof uses DNS to display the name of a host instead of its IP address. You can turn off this functionality as follows:

```
# lsof -n
```

19 Process lsof output

You can process lsof output and create useful summaries. It's easy to find the total number of TCP and UDP connections with the help of **awk**. The following command easily does the trick:

```
# lsof -i | awk '{print $8}' | sort | uniq -c | grep -v NODE
```

The **lsof -i** command lists all Internet connections, **awk** extracts the eighth field and **sort** produces the output. Then the **uniq -c** command counts how many times a line exists. Finally, the **grep -v** command deletes the line with the 'NODE' string in it.

You can get details about all active Internet connections using the following command:

```
# lsof -i -n -P | grep ESTABLISHED | awk '{print $1, $9}' | sort -u
```

20 Other useful networking tools to consider

As it happens with every tool, lsof is not a panacea because it can't solve all kinds of network problems. Other tools that you may want to consider for solving network related problems are netstat, netcat, DTrace, tcpdump, Wireshark, tshark or traceroute.

Despite the existence of the other tools, lsof is a worthy tool that is useful to learn.

21 More about the lsof output

The output of lsof has various columns. The COMMAND column contains the first nine characters of the name of the UNIX command associated with the process. The PID column shows the process ID of the command, whereas the USER column displays the name of the user that owns the process. The TID column shows the task ID - a blank TID indicates a process and the FD column stands for file descriptor. Its values can be cwd, txt, mem and mmap. The TYPE column displays the type of the file: regular file, directory or socket. The DEVICE column contains the device numbers, separated by commas. The value of the SIZE/OFF column is the size of the file or the file offset in bytes. The value of the NODE column is the node number of a local file. Finally, the NAME column shows the name of the mount point and file system the file is located or the Internet address.

Tutorial

This is the main website of the LaTeX project where you can find further useful information

The screenshot shows the LaTeX website interface. At the top, it says "LaTeX – A document preparation system". Below this, there are several sections: "LaTeX the product", "The LaTeX3 project", "Documentation", "Getting LaTeX", and "Getting help". On the left side, there is a navigation menu with links like "Introduction", "LaTeX news", "Documentation", "Books", etc. Below the main content, there is a "Contents" section with a table of contents listing chapters and sections. To the right, a terminal window shows the execution of the `latexmk` command, displaying the process of generating PDF files from LaTeX source files.

Create your own table of contents, list of figures or list of tables easily

Execute the `latexmk` Perl script in silent mode; `Latexmk` is a handy automatic LaTeX document generation routine that saves you time

Write a book in LaTeX

Advisor



Mihalys Tsoukalos is a UNIX administrator, a programmer (UNIX and iOS), a DBA and a mathematician. He has been using Linux since 1993

Learn how to typeset, beautify and customise a whole book using LaTeX and its extensive packages

Resources

LaTeX project

www.latex-project.org

Comprehensive TeX Archive Network

www.ctan.org

Latexmk

<http://users.phys.psu.edu/~collins/software/latexmk-jcc>

The LaTeX Companion, 2nd Edition, Addison Wesley, 2004

Donald Knuth created TeX in the 1970s as a computer language for use in typesetting, after wanting to write books on computer algorithms and finding that no previous typesetting system was acceptable for him. Its origins resonate with the story of the C programming language that was created in order to write the UNIX operating system – TeX had many low-level commands and it was difficult to use, so Leslie Lamport decided to make it a little easier to digest. He programmed many higher-level TeX commands and that became LaTeX, which is widely used today.

Many people want to typeset a large body of work or a book of around 350 pages, for example, but the biggest concern can be the choice between using LaTeX or Adobe InDesign. LaTeX often proves itself to be the best option, so this tutorial will go through the process and some of the modifications and additions that can be made to a book LaTeX-style.

Although this article assumes that you are already familiar with LaTeX, if you have never heard of it and are wondering why it's so popular, you should know that it helps many book and dissertation authors construct their work easily.

01 Install LaTeX

First you need to install LaTeX. There is a plethora of packages out in the wild that contain the name LaTeX, so running `apt-cache search latex` will return a lot of output. You can get a more helpful output by simply executing the following command:

```
$ apt-cache search latex | grep ^latex
```

Your journey to using LaTeX will start by running the following command:

```
$ sudo apt-get install latex209-base latex209-bin
```

```
root@mail:~# sudo apt-get install latex209-base latex209-bin
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following extra packages will be installed:
 fonts-fonts-t1cmexcs tex-common texlive-base
 texlive-latex-base texlive-latex-base-doc texlive-latex
 texlive-latex-extra texlive-latex-extra-doc
Suggested packages:
 perl-tk
The following NEW packages will be installed:
 fonts-fonts-t1cmexcs tex-common texlive-base texlive-latex
 texlive-latex-base texlive-latex-base-doc texlive-latex
 texlive-latex-extra texlive-latex-extra-doc
0 upgraded, 12 newly installed, 0 to remove and 0 not upgraded.
Need to get 88.7 MB of archives.
After this operation, 182 MB of additional disk space will be used.
Do you want to continue? [Y/n] Y
Get:1 http://ppa.archive.ubuntu.com/ubuntu trusty/main tex-common all 4.84-0ubuntu1 [4,719
 kB]
Get:2 http://ppa.archive.ubuntu.com/ubuntu trusty/main fonts-fonts-t1cmexcs all 2.894.4-3 [38,308
 kB]
Get:3 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-base all 2.894.4-3 [12.4 MB]
Get:4 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]
Get:5 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-base-doc all 2.894.4-3 [14,719
 kB]
Get:6 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]
Get:7 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]
Get:8 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]
Get:9 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]
Get:10 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]
Get:11 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]
Get:12 http://ppa.archive.ubuntu.com/ubuntu trusty/main texlive-latex-base-doc all 2.894.4-3 [14,719
 kB]

```

02 Use LaTeX

After installing the basic LaTeX packages you should try to compile the consequent LaTeX code to make sure that everything works exactly as you expected:

```
\documentclass{article}
\title{The LaTeX version of the Hello
World program}
\author{Mihalis Tsoukalos}
\date{January 2015}
\begin{document}
\maketitle
Hello world!
\end{document}
```

You can compile the document using the following command:

```
$ latexmk helloWorld.tex
```

The previous command produces a DVI file that you can convert into PDF format as follows:

```
$ dvi2pdf helloWorld.dvi
```

If you view the generated file you can easily understand that the default LaTeX styles produce high quality output. The problem is that if everyone is using the default styles, all books produced using LaTeX will look exactly the same!

```
root@mail:~# sudo apt-get install texlive-full
sudo: unable to resolve host mail
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following package was automatically installed and is no longer required:
 libnss3-ld
Use 'apt-get autoremove' to remove it.
The following extra packages will be installed:
 aglfn chktex clisp cm-super cm-super-minimal context context-modules dvidvi dvipng
 feynmf fonts-comfortaa fonts-freefont-otf fonts-gfs-artemisla fonts-gfs-baskerville
 fonts-gfs-bodoni-classic fonts-gfs-complutum fonts-gfs-didot fonts-gfs-didot-classic
 fonts-gfs-gazis fonts-gfs-neohellenic fonts-gfs-olga fonts-gfs-porson
 fonts-gfs-solomos fonts-gfs-theokritos fonts-hosny-amiri fonts-inconsolata
 fonts-ipaefont-gothic fonts-ipaefont-mincho fonts-ipafont-gothic
 fonts-ipafont-mincho fonts-junicode fonts-linuxlibertine fonts-oflb-asana-math
 fonts-sil-gentium fonts-sil-gentium-basic fonts-stix fragmaster ko.tex-extra-hlfont
 lacheck latex-cjk-all latex-cjk-chinese latex-cjk-chinese-arphic-bkai00mp
 latex-cjk-chinese-arphic-bsmi00lp latex-cjk-chinese-arphic-gbsn00lp
 latex-cjk-chinese-arphic-gkai00mp latex-cjk-common latex-cjk-japanese
 latex-cjk-japanese-wadalab latex-cjk-korean latex-cjk-thai latex-fonts-thai-tlwg
 latex-sanskrit latexdiff lcdf-typetools libffcall1 libfile-which-perl libmagick++5
 libplot2c2 libpstoedit0c2a libreheadline5 m-tx musixtex otf-freefont pfb2t1c2pfb pmx
 pstoedit purifyeps swath t1utils tex4ht tex4ht-common texlive-bibtex-extra
 texlive-doc-ar texlive-doc-bg texlive-doc-cs+sk texlive-doc-de texlive-doc-en
 texlive-doc-es texlive-doc-fi texlive-doc-fr texlive-doc-it texlive-doc-ja
 texlive-doc-ko texlive-doc-mn texlive-doc-nl texlive-doc-pl texlive-doc-pt
 texlive-doc-rs texlive-doc-ru texlive-doc-si texlive-doc-th texlive-doc-tr
 texlive-doc-uk texlive-doc-vi texlive-doc-zh texlive-fonts-extra
 texlive-fonts-extra-doc texlive-formats-extra texlive-games texlive-humanities
 texlive-humanities-doc texlive-lang-african texlive-lang-arabic texlive-lang-armenian
 texlive-lang-cjk texlive-lang-croatian texlive-lang-cyrillic texlive-lang-czechslovak
 texlive-lang-danish texlive-lang-dutch texlive-lang-english texlive-lang-finnish
```

04 Install the LaTeX perl script

03 Pros and cons

LaTeX has its advantages. It comes with many default styles and if they are sufficient for you, then start writing immediately! It's also fast to type and edit text, so you can use traditional UNIX tools to make changes to your LaTeX code. As a consequence, you don't need a very fast computer to run LaTeX. Additionally, LaTeX files are easily interchangeable between different platforms. One of its biggest advantages is that the output looks highly professional because LaTeX is a pro tool.

However, you do have to compile the LaTeX code to see its output, and LaTeX needs time and experimentation to set up and change the default document styles, which can be annoying.

04 Install the latexmk Perl script

This command installs a highly recommended Perl script to improve the whole book-building process:

```
$ sudo apt-get install latexmk
```

From now on you can translate your LaTeX code into PDF format using the following command:

```
$ latexmk -pdf book.tex
```

Most of the packages you will need are installed by default. Installing a new LaTeX package

can be tricky – the best thing to do is read the documentation and follow the instructions before trying to install it. For this tutorial it's also worth installing the `texlive-full`, `texlive-math-extra`, `texlive-extra-utils`, `texlive-generic-extra` and `texlive-latex-extra` packages.

05 Organise the files

The name of the core LaTeX file will be `book.tex`. Instead of using a big and unmanageable LaTeX file, you can put the book chapters and the appendices in separate files using the `include` LaTeX command:

```
\include{ch1}
\include{apA}
```

Here, LaTeX will search for files named `ch1.tex` and `apA.tex` in the current working directory. This is the most important task of the whole process because it allows you to edit smaller and more manageable files.

We will use four chapters and two appendices to make our example book, but you can use more or less pages on future projects if you wish. After creating the files for all the book chapters and appendices, and referencing them inside `book.tex`, you are ready to start adding content.

06 Create a bibliography

LaTeX offers BibTeX for dealing with bibliographic references. All BibTeX entries will be located in a file called booklist.bib. Each BibTeX entry has the following format:

```
@book{Kernighan84,
  author = {Brian W. Kernighan and Robert
           Pike},
  title = {The UNIX Programming
           Environment},
  year = {1984},
  isbn = {013937681X},
  publisher = {Prentice Hall Professional
              Technical Reference},
}
```

The following LaTeX commands are used for beautifying and inserting the bibliography into this project:

```
\bibliographystyle{alpha}
\addcontentsline{toc}{chapter}
{Bibliography}
\bibliography{booklist}
```

You can add a reference to the Kernighan84 book in your LaTeX code as follows:

```
\cite{Kernighan84}
```

07 Sort content

LaTeX generates pages with table of contents, list of figures and list of tables. Create a table of contents (TOC) using this command:

```
\tableofcontents
```

If you only want to include entries for chapters and sections, include the next command before the previous command in Step 6:

```
\setcounter{tocdepth}{1}
```

In order to include the list of figures you should insert the following command:

```
\listoffigures
```

If you have tables in your book, you can generate the list of tables as follows:

```
\listoftables
```

08 Mini table of contents

If your book chapters are on the large side you may need to add a mini table of contents at the beginning of each chapter. There is a LaTeX package called minitoc that does the job. The next LaTeX code adds support for the mini TOC functionality:

```
\usepackage{minitoc}
\setcounter{minitocdepth}{1}
```

The previous commands should be combined with the `\dominitoc` command that must be used once to turn on the mini TOC functionality. Finally, type `\minitoc` at the beginning of each chapter to actually add the mini table of contents.

09 Format code

The book is going to contain multiple blocks of code to showcase certain techniques, therefore the code must be well formatted. Using the following LaTeX code you can define a new style called MyCode for formatting code:

```
\usepackage{reysize, color, fancyvrb}
\CustomVerbatimEnvironment{MyCode}
{Verbatim}
  {numbers=left,frame=lines,framerule=
  1pt,rulecolor=\color{blue},framesep=4mm}
```

Creating a style saves you time. Additionally, if you want to make a global change to the output format of code in your book, all you have to do is go to its definition, make the changes and recompile the book. Also, if you want to create a new style similar to an existing one, you can copy and paste the existing style definition code, choose a different name and make the modifications you want.

You can now use the MyCode style inside LaTeX using the following command:

```
\begin{MyCode}[fontsize=\relsize{-1}]
#include <stdio.h>
int main (int argc, const char * argv[])
{
  printf("Hello, World!\n");
  return 0;
}
\end{MyCode}
```

If you change the `\relsize` value from `-1` to `+1`, the code size is going to get bigger.

10 Include a header and footer

The following LaTeX code changes the header and the footer of each book page. Right pages (RO) have a different rule to left pages (LE). It's also worth noting that the first page of each chapter (CO) has another rule that puts the page number at the footer of it without adding any header information.

```
\renewcommand\headrule
{
  \color[gray]{0.5} \hrule height 1pt
  width\headwidth
  \vspace{-4pt}}
\fancyhead[RO]{%
  {\bfseries\S \thesection\ \textbar\ \
  LARGE\thepage\normalsize\ \textbar}}
}
\fancyhead[LE]{%
  {\bfseries\textbar\ \LARGE\thepage\
  normalsize\ \textbar\ \S \
  thesection\ \textbar\ \leftmark}}
}
\fancypagestyle{plain}{%
  \fancyhf{}
  \fancyfoot[CO]{%
    {\bfseries\LARGE\thepage}}
  \renewcommand{\headrulewidth}{0pt}%
}
```


12 Generate an index

Index

```
-apple-dashboard-region CSS command, 111  
.lproj, 182  
.wdgt, 62  
  
Aho, 151  
Ajax, 6  
alert, 145  
alert() function, 282  
AllowFileAccessOutsideOfWidget, 68  
AllowFullAccess, 68  
  
C++, 6, 78, 122  
calculator, 4  
Calendar Widget, 16  
Carbon, 122, 123  
Carnegie-Mellon, 1  
CFBundleDisplayName, 68  
CFBundleIdentifier, 68, 182  
CFBundleName, 68, 182  
CFBundleVersion, 68  
CGL, 96
```

You should first insert the following code for the modifications to work:

```
\usepackage{fancyhdr}  
\pagestyle{fancy}  
\fancyhf{}  
\renewcommand{\chaptermark}[1]{\markboth{#1}{}}
```

The last command is for resetting the existing definition of the header and the footer.

11 Format figures

You can easily alter the existing format of a picture caption by using the following LaTeX commands:

```
\usepackage{caption}  
\DeclareCaptionLabelSeparator{par}{\par}  
\DeclareCaptionFormat{dashedlabel}{\textbf{---} #1 \textbf{---}#2#3}  
\captionsetup{format=dashedlabel,margin=1cm,singlelinecheck=true,font=small,labelfont={sc,bf},textfont=it,justification=centerlast,labelsep=par}
```

Now you will be able to add a specific figure in your book as follows:

```
\begin{figure}[tb]  
  \centering  
  \fbox{\scalebox{0.45}{\includegraphics{figures/ch1/f0101.jpg}}}  
  \caption{Adding a figure}  
  \label{ch10:fig1}  
\end{figure}
```

12 Generate an index

The following command helps you create a book index:

```
\usepackage{makeidx}
```

The next LaTeX command puts the index information into the generated file:

```
\printindex
```

This code adds the word Knuth to the index:

```
\index{Knuth}
```

13 Add tables

This code produces an elegant table to blend in with the rest of the book:

```
\begin{table}[t]  
  \centering  
  \setlength{\extrarowheight}{2.5pt}  
  \setlength{\extrarowheight}{2.5pt}  
  \begin{tabular}{cccc}  
    \hline \addlinespace  
    \textbf{Name} & \textbf{Size in Kbytes} & \textbf{\# of files} & \textbf{\# of keys in Info.plist} \\ \hline  
    amazonsearch.wdgt & 144 & 8 & 13 \\ \hline  
    Tags - HTML.wdgt & 88 & 17 & 11 \\ \hline  
    Temperature Monitor.wdgt & 416 & 50 & 10 \\ \hline  
    Wikipedia.wdgt & 696 & 125 & 12 \\ \hline  
  \end{tabular}  
  \caption{Statistical Information.}  
  \label{ch3:stats}  
\end{table}
```

1

The Dashboard

1.1 First Section of Chapter 1

The main disadvantages of L^AT_EX are that you will have to compile the L^AT_EX code to see its output and that it takes more time and experimentation to set up and change the default document styles.

14 Some other useful info

Producing a more fancy and professional chapter heading can be done easily using the `fncychap` package:

```
\usepackage[Bjornstrup]{fncychap}
```

15 Work with the NOTE style

Before talking about the NOTE style you should first download the `picins.sty` file and put it where your LaTeX files are:

```
$ wget http://tug.ctan.org/tex-archive/macros/latex209/contrib/picins/picins.sty
```

The NOTE style was defined as follows:

```
\usepackage{picins}  
\usepackage{amsthm}  
\newcommand\NOTE{\linethickness{1mm}\parpic(14mm,5mm)[s1]{NOTE}\noindent}
```

You can now add a note as follows:

```
\NOTE A small note doesn't look very elegant
```

The only drawback of the style is that it needs to have at least two lines of text to look pleasant on the page.

16 How it worked

Our example book was built step-by-step: a minimal `book.tex` file was generated first, then chapter and appendix files without any data were added to the project. At this point it was time to write the book by adding text, images, tables, bibliography and index entries, while experimenting with the look of the existing LaTeX styles, making changes to them, adding new ones and correcting errors. This is the safest procedure to follow when writing a book using LaTeX – LaTeX was able to process all the files, as well as the index and the bibliography, without any difficulties.

Here, you can add new files or create a collection of media files. Files can be assigned to collections at any point of time

Update account settings here and track media files using the Media processing panel, which shows details about uploads

All creative commons licences are available here to assign under an available media file

This section lets you upload a media file and update the relevant information regarding that file

Host your own media gallery with MediaGoblin

MediaGoblin provides a way to share videos, photos or audio recordings with your loved ones, without worrying about privacy

Advisor

Nitish Tiwari is a software developer by profession and an open source enthusiast by heart. As well as writing for leading open source magazines, he helps firms set up and use open source software for their business needs



YouTube is not only a website anymore, it's become a phenomenon. Millions of hours are spent – or wasted – daily in watching videos of cats, dogs and humans doing strange things. With the predictive playlist appearing just after you finish a video, it is sometimes very difficult to close the window and you go on and on, watching one video after the other. But YouTube is a dangerous place for personal videos and other media that you don't want strangers to access. Though it has an option to make your videos private, you don't really know how private it is. So we need to find a solution that has the perfect match of convenience and privacy.

This is where MediaGoblin comes in. This open source tool can help you organise, host and stream media from your own PC without having to worry about privacy. If you are a power user, you can also have it run on a web server and let other people add their media. There are a range of other useful features available too, like tags and collections to name just two. In this tutorial, we will begin with the steps to install MediaGoblin on Ubuntu and then proceed to get it working and see it in action. We have used Ubuntu 14.04 as the host system and MediaGoblin's clone from their Gitorious repo.

Resources

MediaGoblin home page
mediagoblin.org

Host your own media gallery

Share media without worrying about privacy

02 Set up the database

```
nitish@nitish-ubuntu: /srv/mediagoblin.example.org/mediagoblin
en us
Removing obsolete dictionary files:
* No PostgreSQL clusters exist; see "man pg_createcluster"
Setting up postgresql-client (9.3+154) ...
Setting up python-egenix-mxtools (3.2.7-1build1) ...
Setting up python-egenix-mxdatetime (3.2.7-1build1) ...
Setting up python-psycopg2 (2.4.5-1build5) ...
Processing triggers for ureadahead (0.100.0-16) ...
Setting up postgresql-9.3 (9.3.5-0ubuntu0.14.04.1) ...
Creating new cluster 9.3/main ...
  config /etc/postgresql/9.3/main
  data   /var/lib/postgresql/9.3/main
  locale en_IN
  port   5432
update-alternatives: using /usr/share/postgresql/9.3/man/man1/postmaster.1.gz to
provide /usr/share/man/man1/postmaster.1.gz (postmaster.1.gz) in auto mode
* Starting PostgreSQL 9.3 database server [ OK ]
Setting up postgresql (9.3+154) ...
Processing triggers for libc-bin (2.19-0ubuntu6) ...
nitish@nitish-ubuntu: /srv/mediagoblin.example.org/mediagoblin$ sudo -u postgres
createuser -A -D mediagoblin
nitish@nitish-ubuntu: /srv/mediagoblin.example.org/mediagoblin$ sudo -u postgres
createdb -E UNICODE -O mediagoblin mediagoblin
nitish@nitish-ubuntu: /srv/mediagoblin.example.org/mediagoblin$
```

01 Sort out dependencies

MediaGoblin is a full-fledged media-streaming platform and therefore it has a few software dependencies that you will need to take care of before installing MediaGoblin. Let's take a look at these dependencies and how to install them.

Python 2.6 or 2.7 This interprets the MediaGoblin source code.

Python-lxml Binds certain C libraries to Python.

Git For downloading and updating the repository.

SQLite/PostgreSQL This is where everything is stored. SQLite is the default option and works fine for small set-ups, but you need to use PostgreSQL if you expect more users.

Python Imaging Library This adds image-processing capabilities to Python interpreter.

virtualenv This is used to create isolated Python environments.

You can install all these on a Debian based system, using the apt-get command. It can be done with a single command:

```
sudo apt-get install git-core python
python-dev python-lxml python-imaging \
python-virtualenv
```

02 Set up the database

As we said before, the default SQLite database doesn't perform well for deployments involving more than two or three users. So, if you are planning to have more than three users, it's recommended to use the PostgreSQL database. To set it up for MediaGoblin, first download and install the packages using apt-get:

```
sudo apt-get install postgresql
postgresql-client python-psycopg2
```

Note that it has other required packages too. The installation process creates a user (postgres) with sufficient privileges to handle the database, but keeping security in mind we will create a separate user for MediaGoblin. To create the new user, type:

```
sudo -u postgres createuser -A -D
mediagoblin
```

Once the user is created, create the database:

```
sudo -u postgres createdb -E UNICODE -O
mediagoblin mediagoblin
```

Here the first 'mediagoblin' is the user name and the second is the name of database.

03 MediaGoblin user

If you followed the previous step properly, you'd have noticed that we didn't add a password for the user named mediagoblin. So how does the system authenticate the user? This is done via the local Unix authentication. Local Unix authentication allows a system user to connect to any PostgreSQL database on the system without a password. To enable this, you need to create a system user with same name as the PostgreSQL database user. So we now need to create an unprivileged system user named mediagoblin. Note that the user can be underprivileged because MediaGoblin doesn't need any special privileges to run. This also makes the system more secure. Run this command to create the user:

```
adduser --system mediagoblin
```

Since it doesn't have password, you can't login to this account but a switch is possible using:

```
sudo -u mediagoblin /bin/bash
```

Once created, you can then use this user account for all further steps.

04 Install MediaGoblin

To start the installation, you need to create a working directory for MediaGoblin. This is where the git repository will be downloaded. Create the directory using:

```
sudo mkdir -p /srv/mediagoblin.example.org && sudo chown -hR mediagoblin /srv/mediagoblin.example.org
```

As you can see, we create the directory with elevated privileges (root) and then change the owner to our underprivileged mediagoblin user. Let's now clone the MediaGoblin repo to this folder. First switch to the mediagoblin user and then change the directory to the working directory we just created:

```
cd /srv/mediagoblin.example.org
```

Now start cloning:

```
git clone git://gitorious.org/mediagoblin/mediagoblin.git
```

Then move to the 'mediagoblin' folder created – `cd mediagoblin`. Initialise the repo and then fetch the data:

```
git submodule init && git submodule update
```

You'll notice that we didn't take code from the stable revision but instead the master branch of the git repository. MediaGoblin is under rapid development so it makes sense to use the master, at least until a consistent release.

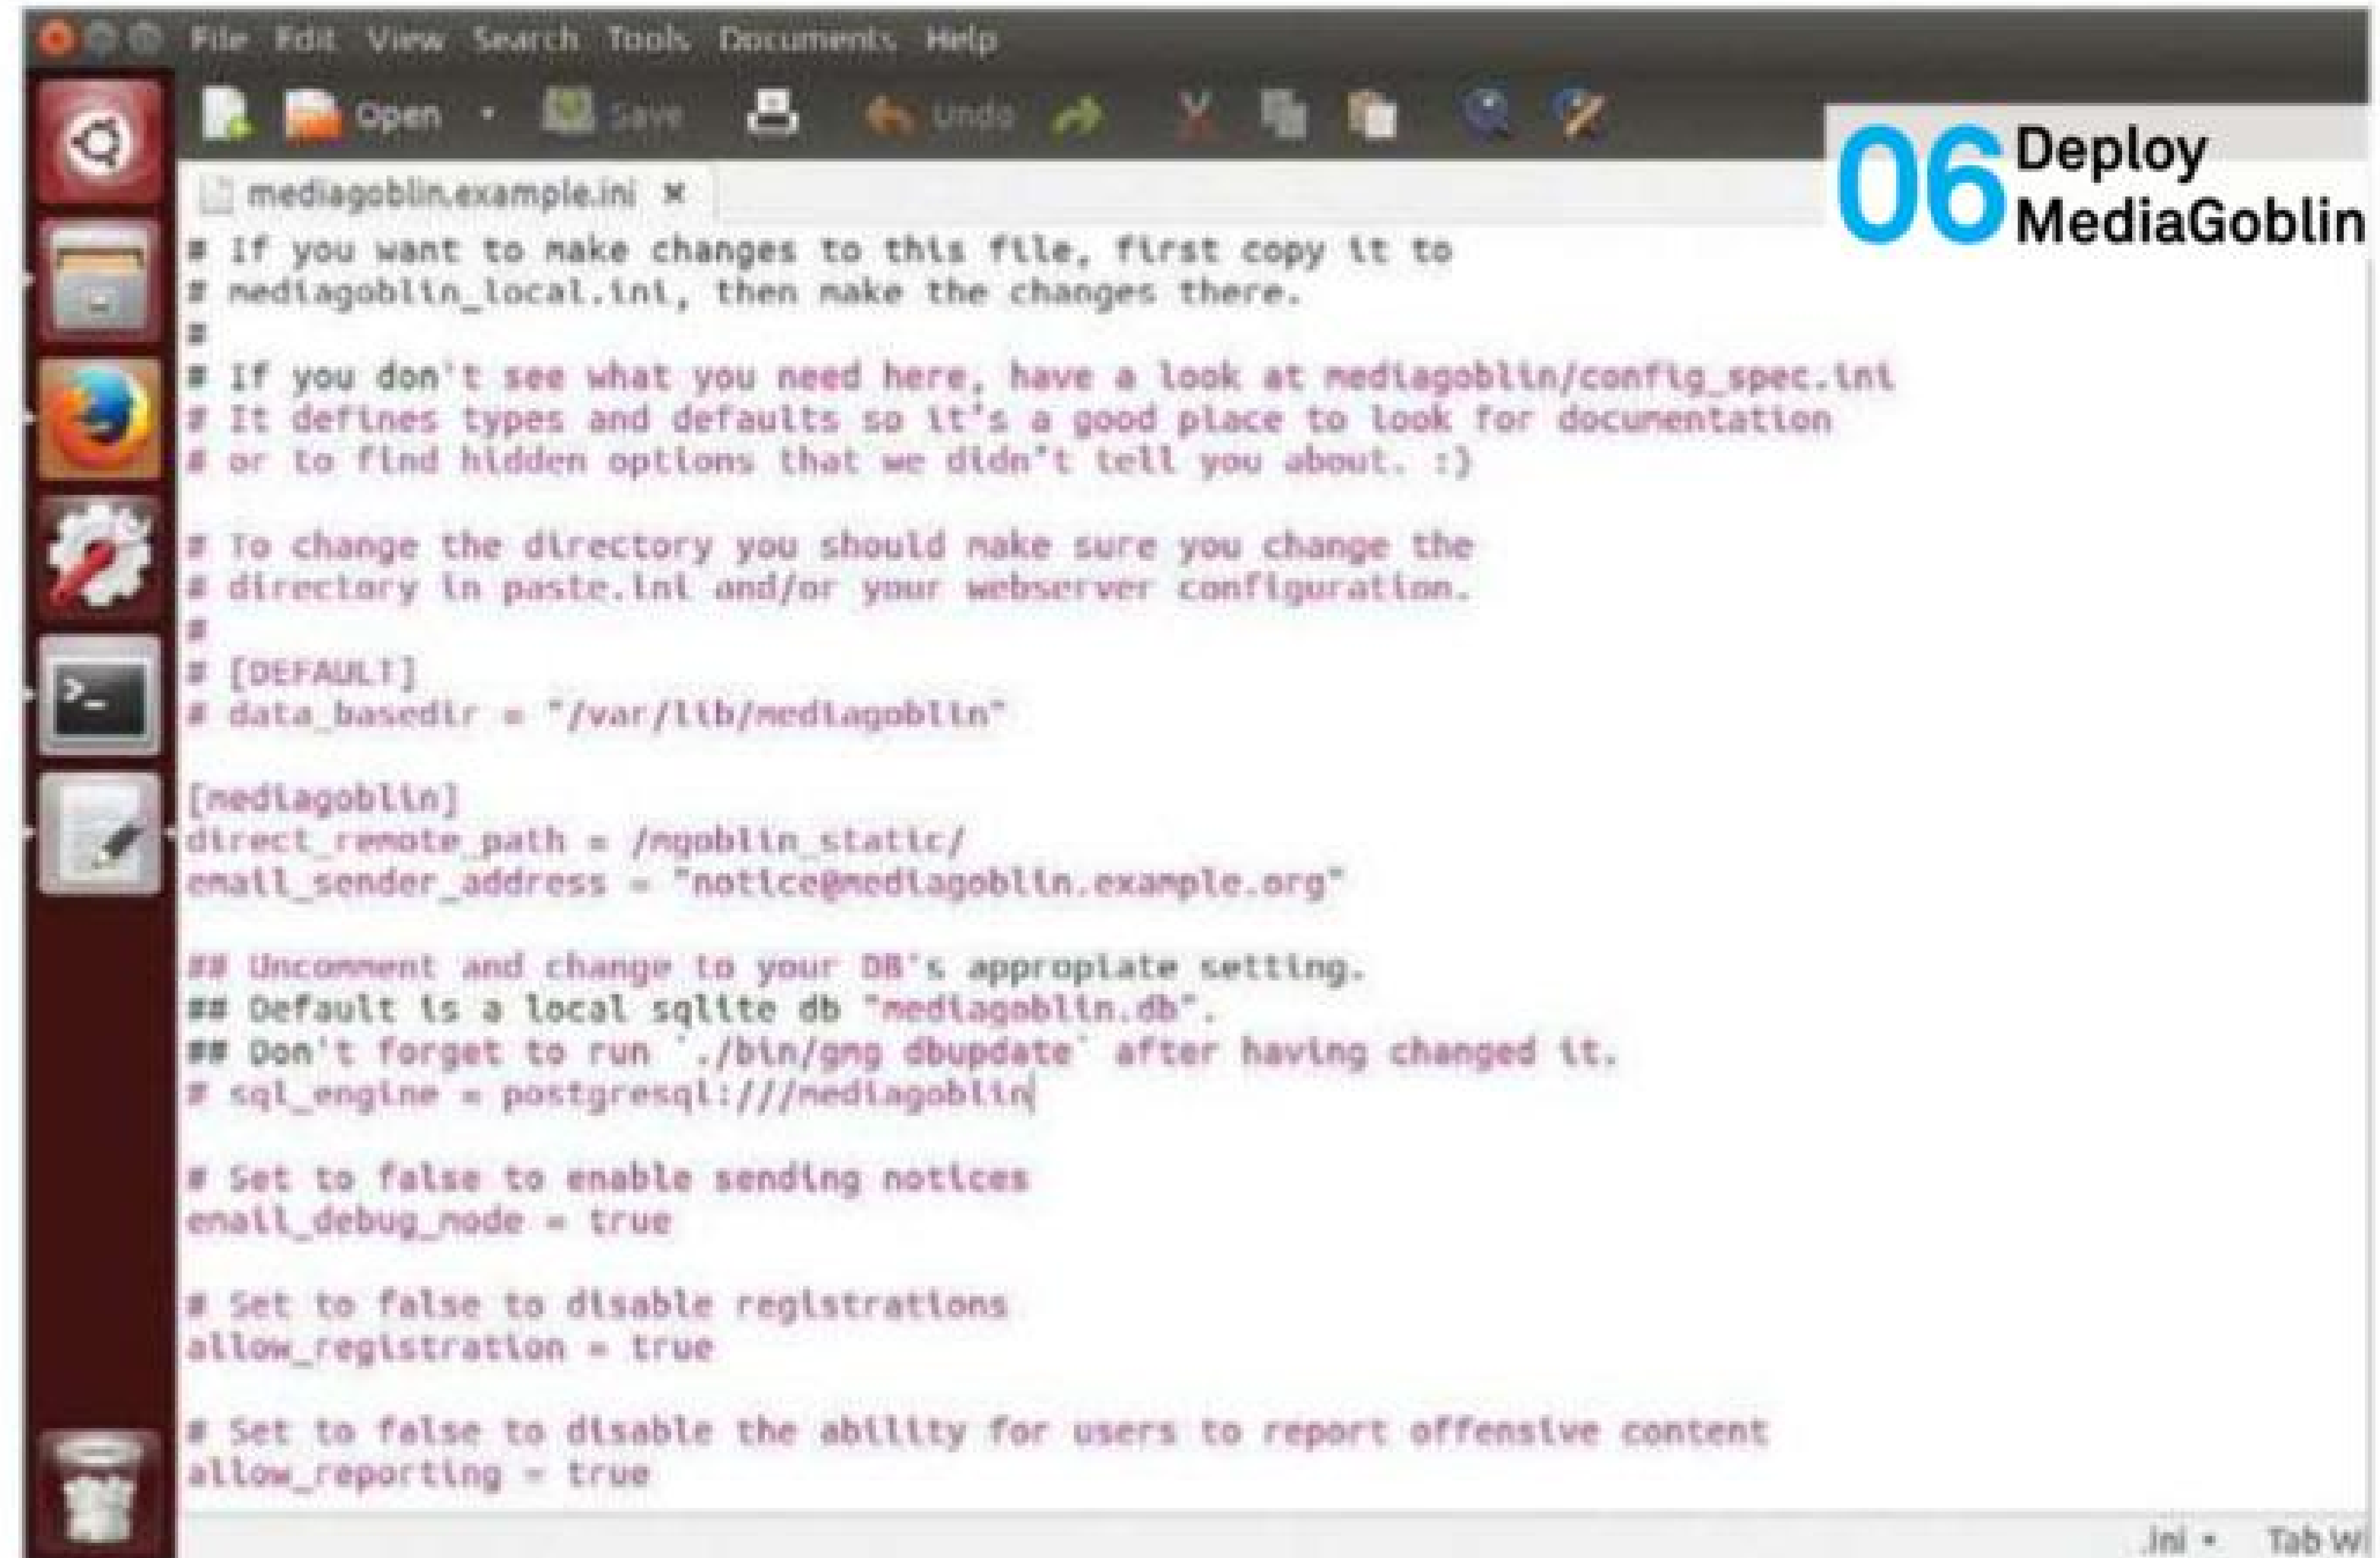
05 Install Virtualenv & others

MediaGoblin uses virtualenv – a tool to help manage the dependencies by creating isolated Python environments. It's already available in the package, so set it up by using:

```
(virtualenv --python=python2 --system-site-packages . || virtualenv --python=python2 .) && ./bin/python setup.py develop
```

If you are feeling adventurous, you can also try the experimental deploy system (shell script) instead of the earlier command:

```
./experimental-bootstrap.sh && ./configure && make
```



MediaGoblin doesn't need any special privileges to run

This script sets up virtualenv and also helps you keep it updated by running `make update`, but as per the developers of MediaGoblin, this is still under development and may break. To update the codebase at a later point of time simply run:

```
git submodule update && ./bin/python setup.py develop --upgrade && ./bin/gmg dbupdate
```

You also need to install Flup before the setup concludes. Install it using:

```
./bin/easy_install https://pypi.python.org/pypi/flup/1.0.3.dev-20110405
```

We will see more on Flup and FastCGI in the following sections.

06 Deploy MediaGoblin

Now that dependencies are set up and MediaGoblin is installed, we will edit the MediaGoblin configuration files – specifically the `mediagoblin.ini` file located inside `/srv/mediagoblin.example.org/mediagoblin`. Here are the changes required:

- Set `email_sender_address` as the ID you want to use for sending system mails.
- Uncomment the line `sql_engine = postgresql:///mediagoblin` if you are using PostgreSQL.
- Edit `direct_remote_path`, `base_dir` and `base_url` as per the root of virtual host.

Now update the database using `./bin/gmg dbupdate`. This populates the database with MediaGoblin data structures. Finally, test the MediaGoblin server using:

```
./lazyserver.sh --server-name=broadcast
```

You should now be able to connect on your browser port 6543.

07 Flup and FastCGI

MediaGoblin uses FastCGI for deployment and FastCGI needs a server. So we need Flup – a FastCGI server. We already installed Flup in step five, so let's now see why FastCGI is so important. Later you will learn a FastCGI setup with an Nginx server to serve MediaGoblin pages.

Host your own media gallery

Share media without worrying about privacy

FastCGI is a protocol to interface interactive programs with a web server – it's an improvement over CGI (common gateway interface). CGI, while easy to implement, had problems in scaling since separate processes were created for each web request – a huge overhead for the host OS. FastCGI solves this by using persistent processes to serve series of web requests; moreover, these processes are owned by FastCGI server (Flup in our case) and not the web server. This de-couples webserver and FastCGI server, allowing effective scaling. Now any server that supports FastCGI can be used for MediaGoblin. Nginx is a good option because of easy configuration and setup.

08 Nginx setup

Nginx has been slowly rising in the ranks of the web server of choice and is currently one of the most used web servers. An acronym for Engine X, it is a high-performance HTTP server. It does support a lot of other protocols too but those are out of scope for us here. Let's go straight to the server set up. Create a configuration file at `/srv/mediagoblin.example.org/nginx.conf` and create a symbolic link into a directory that will be included in your nginx configuration with one of the following commands (as the root user):

```
ln -s /srv/mediagoblin.example.org/nginx.conf /etc/nginx/sites-enabled/
```

This way, a change in one file automatically reflects in the other. You need to then add the contents to the configuration file, as shown in the screenshot below. Remember to change the fields as per your local paths. Once done, restart nginx using `sudo /etc/rc.d/nginx restart`. If everything goes well, start MediaGoblin using:

```
cd /srv/mediagoblin.example.org/mediagoblin/ ./lazyserver.sh --server-name=fcgi fcgi_host=127.0.0.1 fcgi_port=26543
```

Visit mediagoblin.com to see an example MediaGoblin gallery in action.

09 MediaGoblin home

The setup process is a little lengthy, and for the novice user it may seem a complex task, but the steps are simple and you just need to follow them one at a time. Once you have successfully completed the process, you can enjoy uninterrupted media streaming for you and your loved ones.

```
server {
#####
# Stock useful config options, but ignore them :)
#####
include /etc/nginx/mime.types;

autoindex off;
default_type application/octet-stream;
sendfile on;

# Gzip
gzip on;
gzip_min_length 1024;
gzip_buffers 4 32k;
gzip_types text/plain text/html application/x-javascript text/javascript text/xml text/css;

#####
# Mounting MediaGoblin stuff
# This is the section you should read
#####

# Change this to update the upload size limit for your users
client_max_body_size 8m;

# prevent attacks (someone uploading a .txt file that the browser
# interprets as an HTML file, etc.)
add_header X-Content-Type-Options nosniff;

server_name mediagoblin.example.org www.mediagoblin.example.org;
access_log /var/log/nginx/mediagoblin.example.access.log;
error_log /var/log/nginx/mediagoblin.example.error.log;

# MediaGoblin's stock static files: CSS, JS, etc.
location /mgoblin_static/ {
    alias /srv/mediagoblin.example.org/mediagoblin/mediagoblin/static/;
}

# Instance specific media:
location /mgoblin_media/ {
    alias /srv/mediagoblin.example.org/mediagoblin/user_dev/media/public/;
}

# Theme static files (usually symlinked in)
location /theme_static/ {
    alias /srv/mediagoblin.example.org/mediagoblin/user_dev/theme_static/;
}

# Plugin static files (usually symlinked in)
location /plugin_static/ {
    alias /srv/mediagoblin.example.org/mediagoblin/user_dev/plugin_static/;
}

# Mounting MediaGoblin itself via FastCGI.
location / {
    fastcgi_pass 127.0.0.1:26543;
    include /etc/nginx/fastcgi_params;

    # our understanding vs nginx's handling of script_name vs
    # path_info don't match :)
    fastcgi_param PATH_INFO $fastcgi_script_name;
}
}
```


08 Nginx setup

The first step after you're ready with your own MediaGoblin instance is to create an account. This is because you can browse the collections anonymously but you need an account to upload media. To create an account click on the 'Log in' button on the top-right corner. In the log in page that appears next, click on 'Create one here' to open the user registration page. Fill the details in the registration page and you are good to go! Just log in with the credentials and then click on your user name to be redirected to your profile page. Here you have the options to upload media and manage your profile.

10 Add media

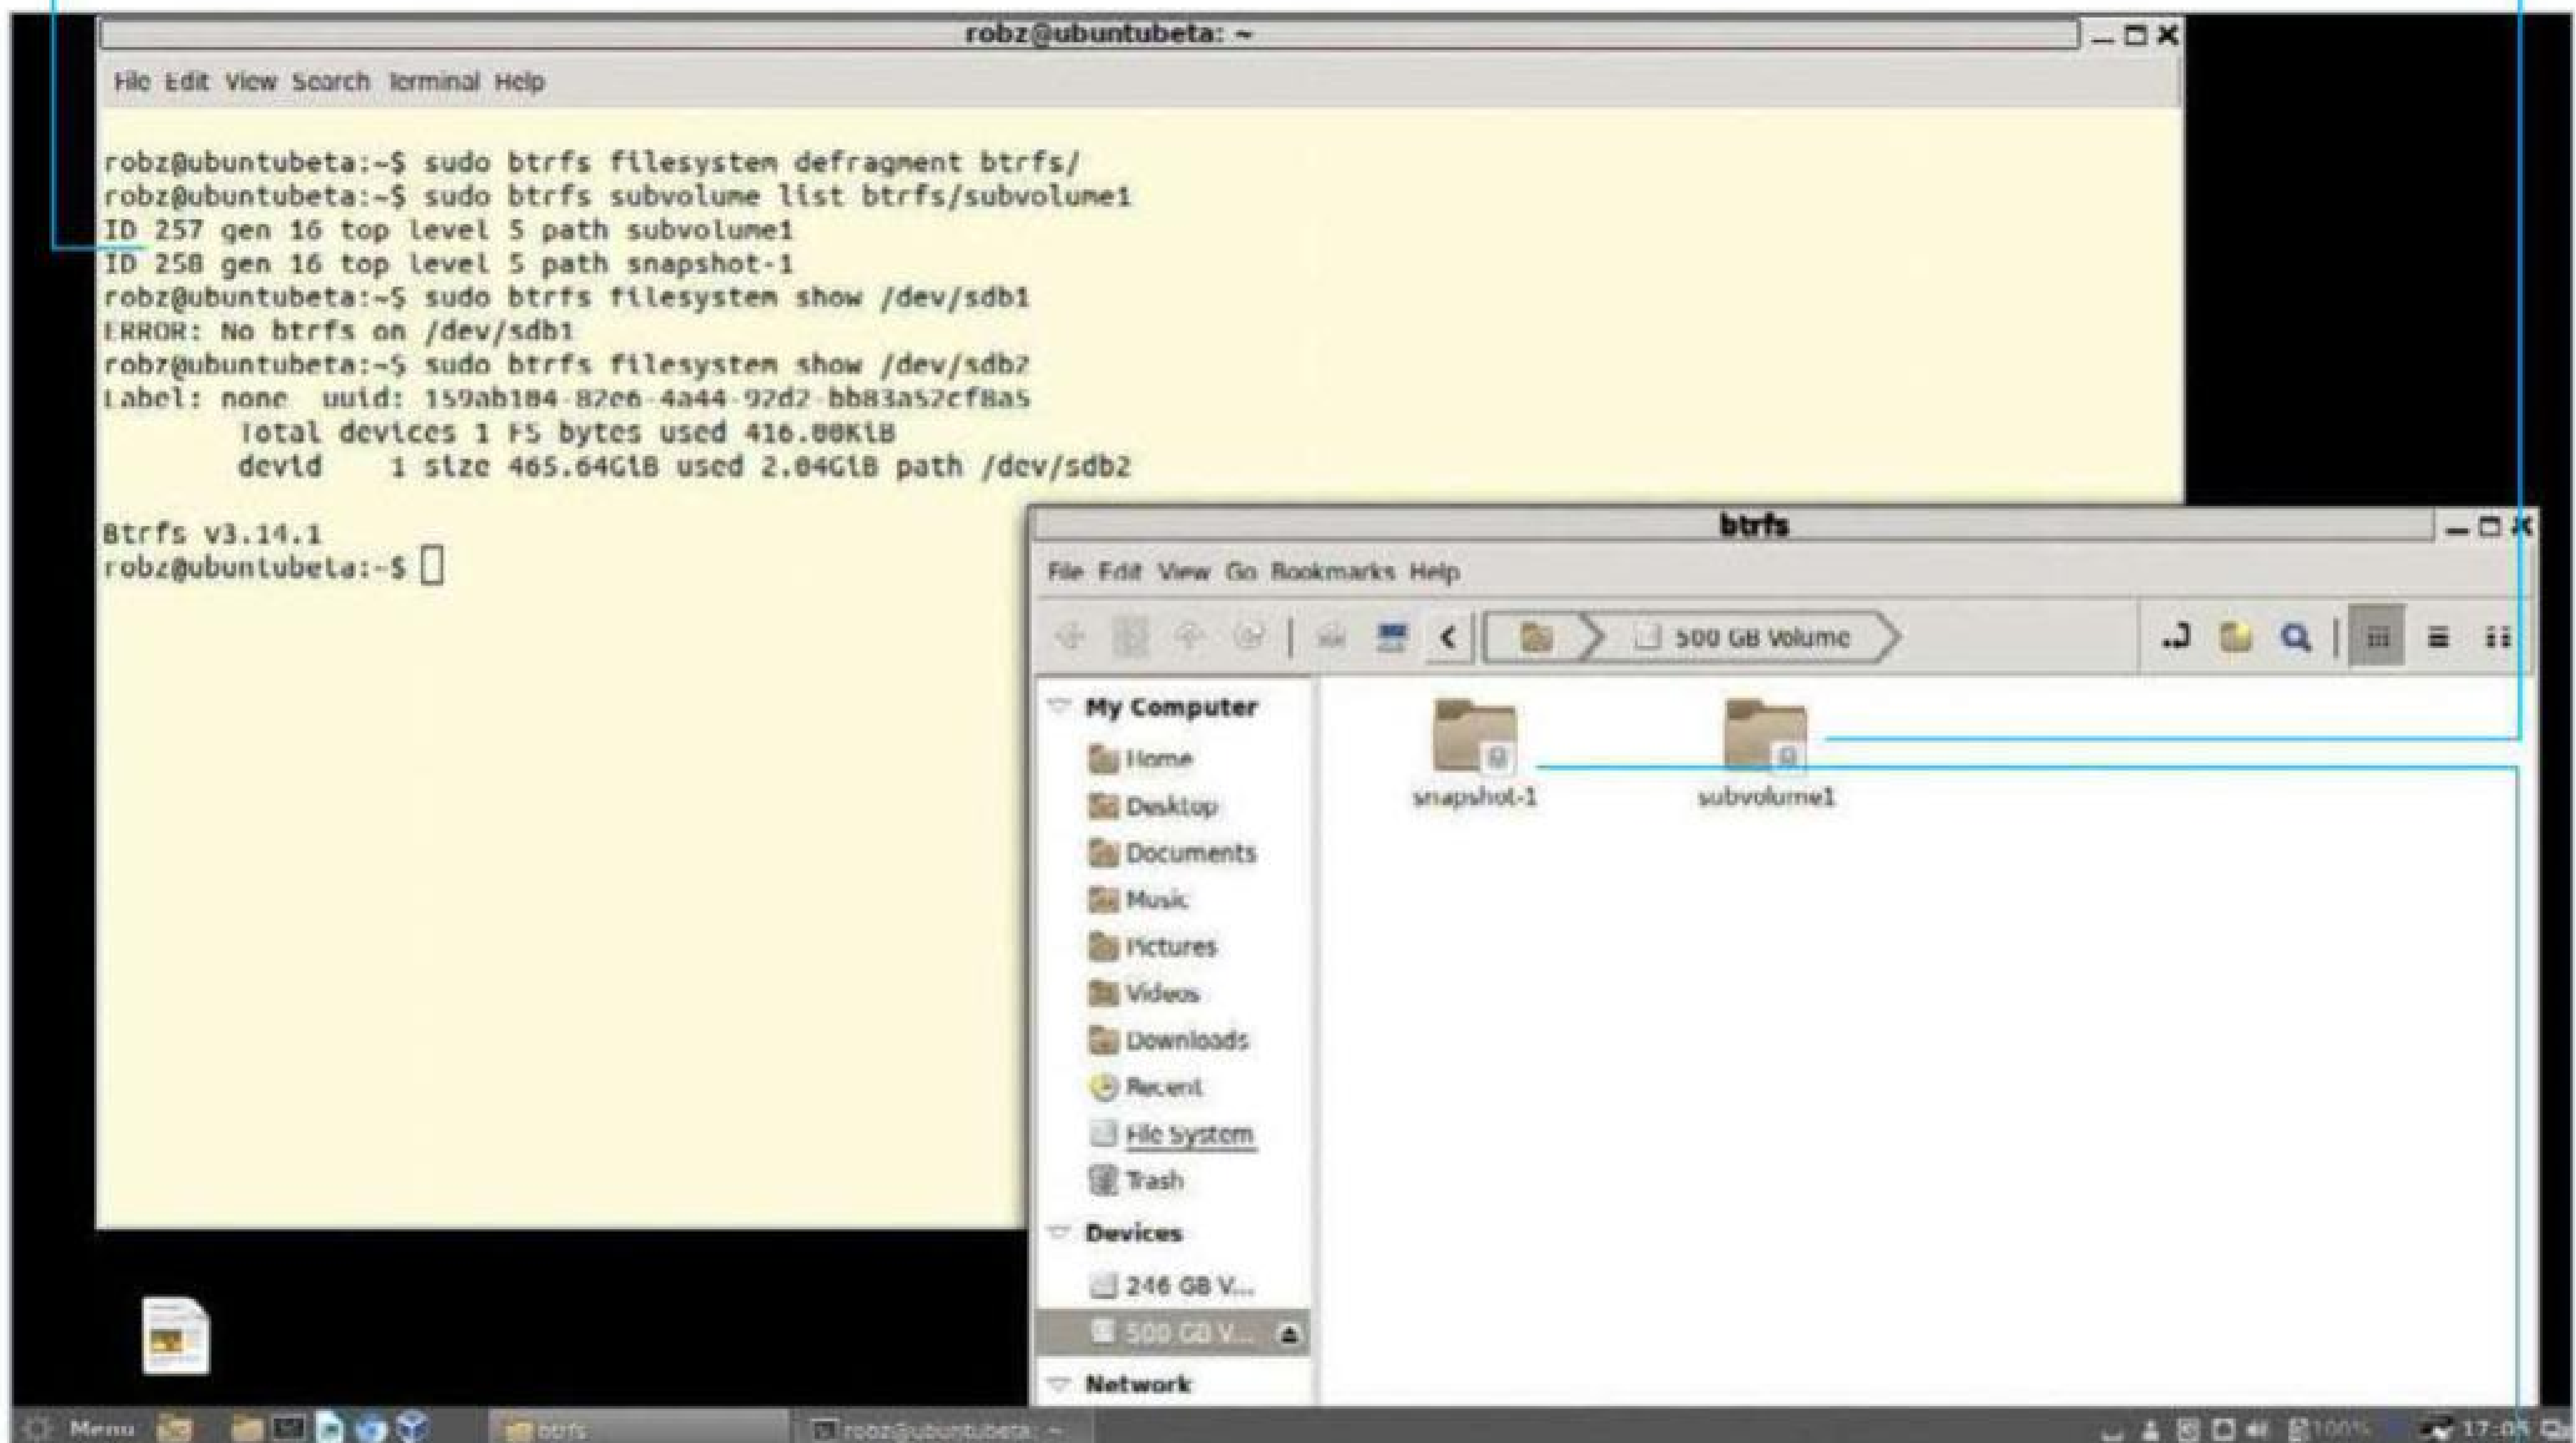
Adding a media file is a breeze, just click on the 'Add media' button on the right side of your profile page. In the next page, upload the file, set the title of the media, add a

description, add tags and set the license you'd like to assign to the media. Finally click on the 'Add' button to upload the file. On the left side of the page, you may notice the 'Browse Collections' link. This option lets you browse collections created by other users (if you are on a multi-user environment). A collection is a group of media files logically bundled together, generally to represent an event or other such scenarios. Note that media files can be added to collections at any point in time and not just during the upload. To create a collection yourself, click on the top-right icon to reveal the account related options and then click on 'Create new collection'. You can then add the title and description to add your own collection.



Create multiple types of btrfs file systems, from single hard drives to RAID 10 arrays

Create subvolumes within the original file system with different mount options



Back up volumes with snapshots that you can roll back to with ease

Switch to the btrfs file system

Advisor



Rob Zwetsloot models complex systems and is a web developer proficient in Python, Django and PHP. He loves to experiment with computing

Discover how to set up and use all the great features of the next generation of file system, btrfs

Resources

Gparted <http://gparted.org>
btrfs-tools

The ext3 and ext4 file systems have been a mainstay of Linux for a long time – very standard, partition-based file systems that have some neat defragmentation policies. They're used all the time by desktop distros and servers and such, even within LVMs that some distributions like to create.

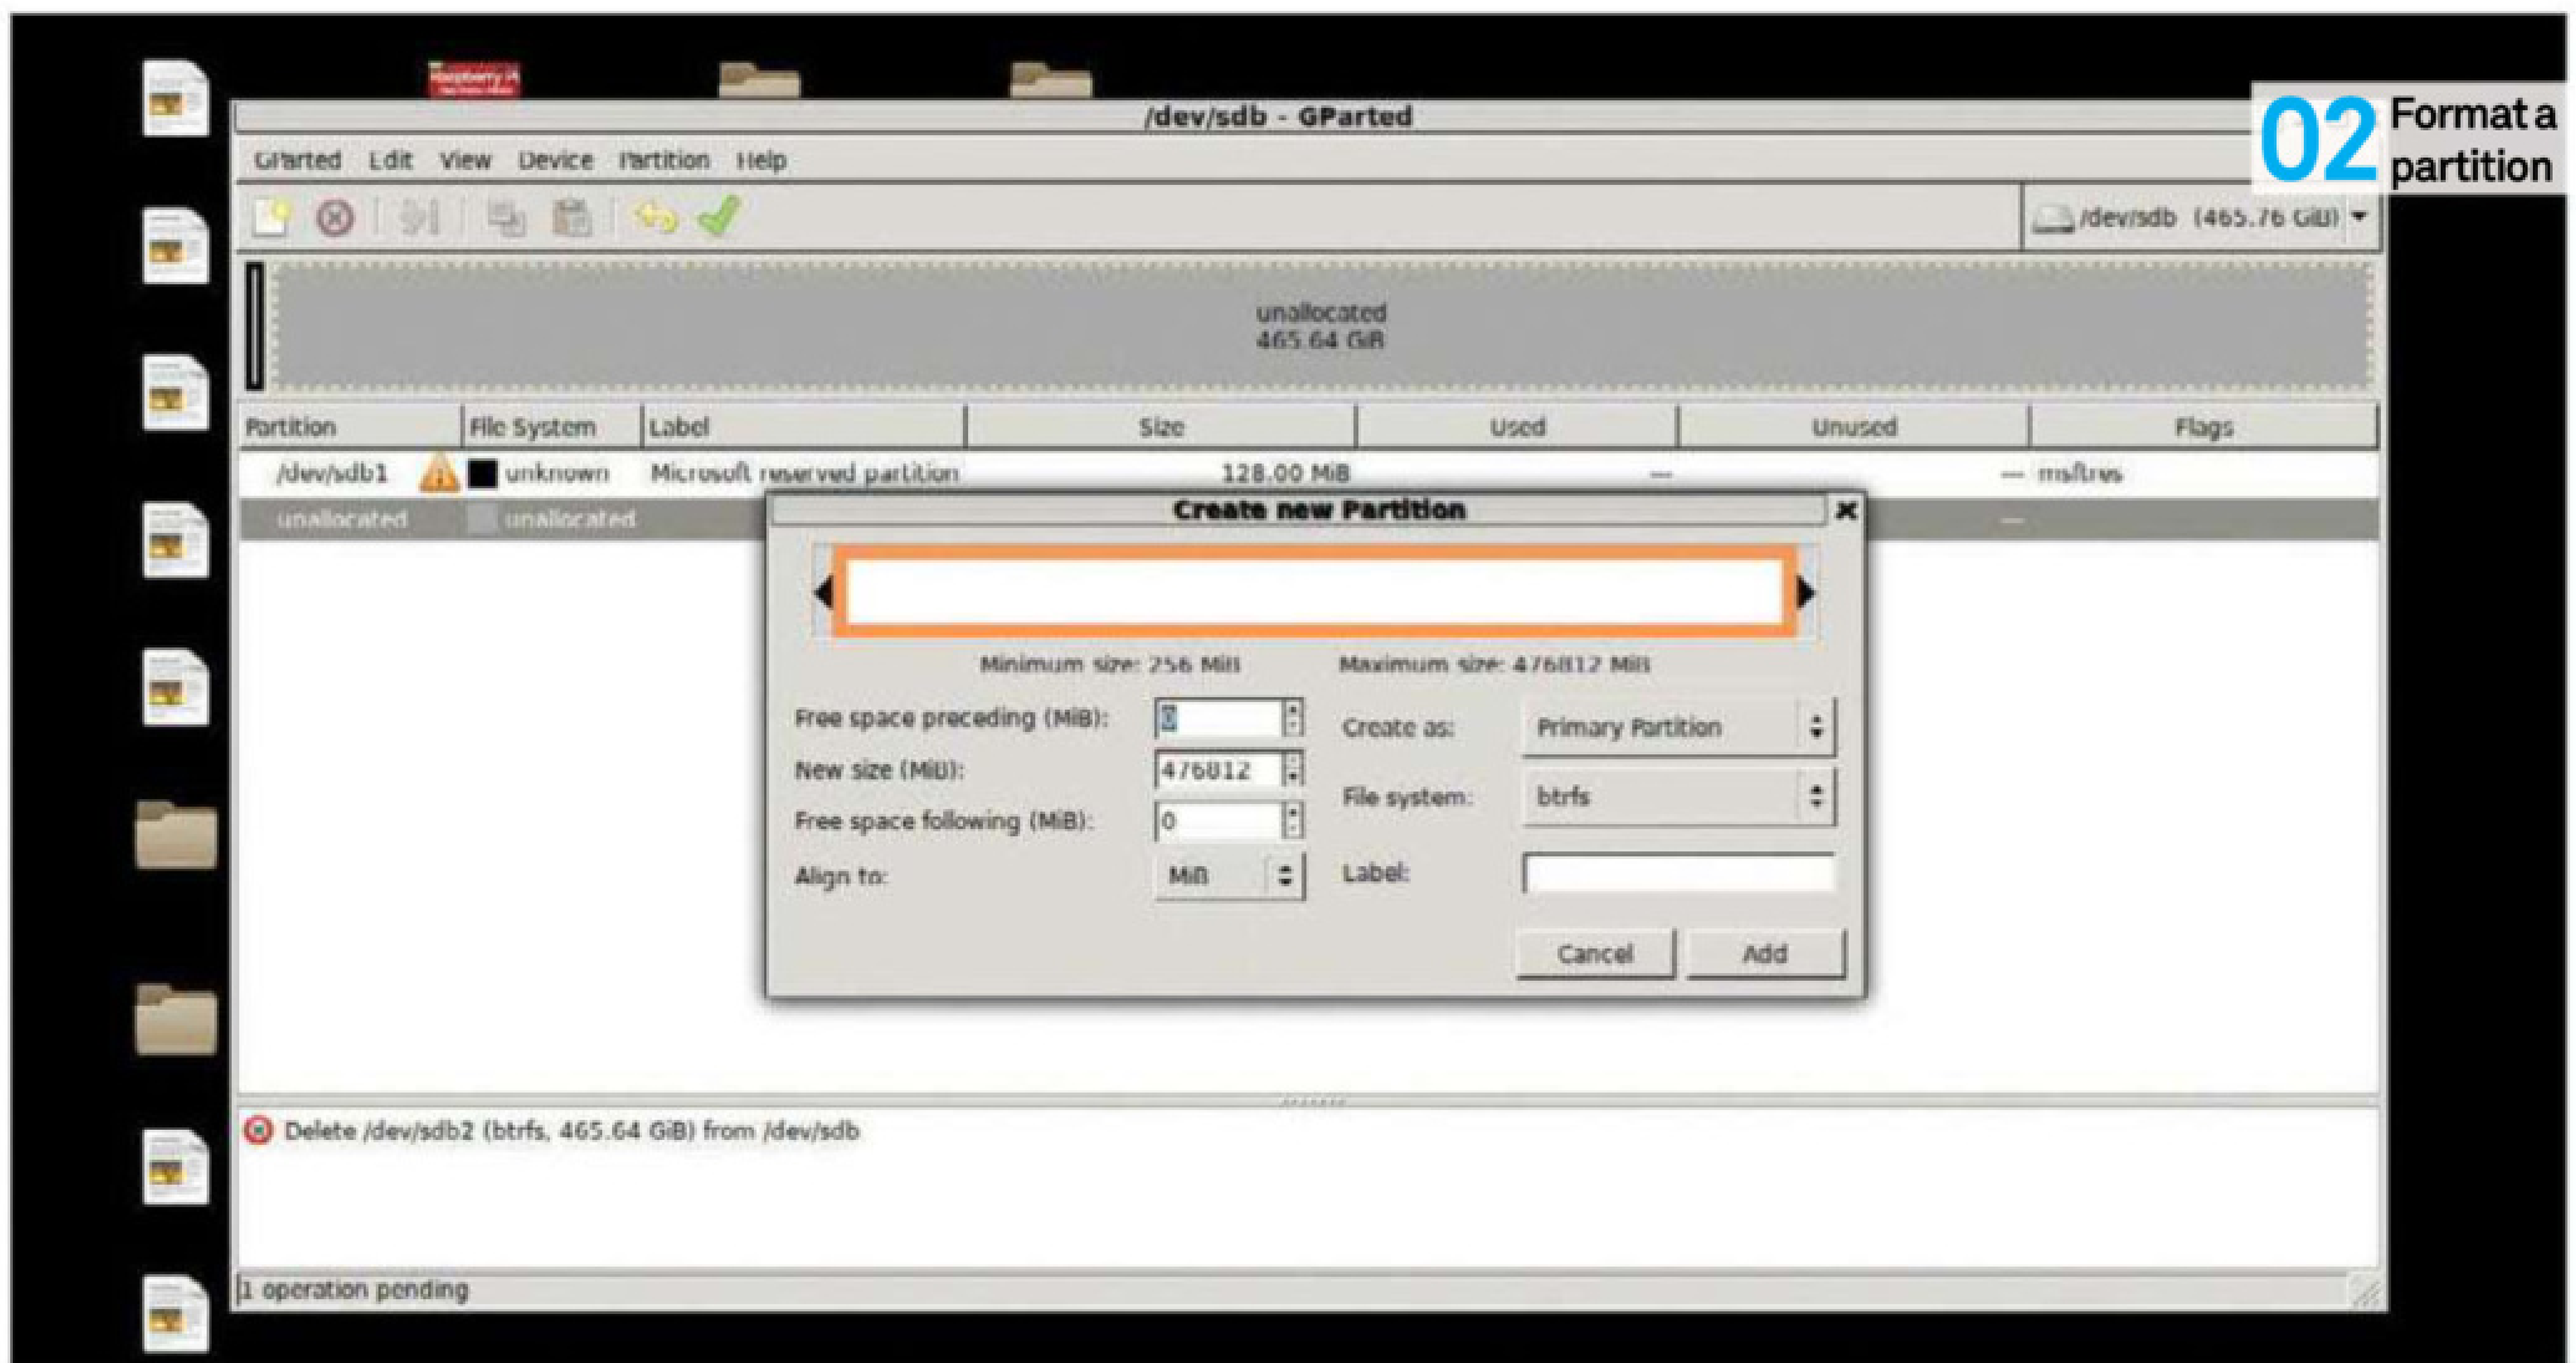
There's always a new challenger though, and for the past few years Fedora has been desperately trying to get btrfs in as the default file system of the distro. It's considered the

next step in Linux file systems and has some excellent features that should make it a perfect choice for the evolution of Linux.

With the ability to create sub-partitions within itself that are easier to manage than LVM, a degree of roll-back, advanced compression and defragmentation tools tailored to itself and more, btrfs will revolutionise the way we use storage. We'll show you how you can get ahead of the curve in this tutorial.

Switch to the btrfs file system

Get ahead of the curve with this next-gen storage technology



02 Format a partition



01 Find the right tools
The very first step is to install the appropriate setup tools. While it's supported in the kernel, not every distro holds the right packages to manage it by default. In Debian and Ubuntu distros, install it with:

```
$ sudo apt-get install btrfs-tools
```

02 Format a partition
Now the tools are installed, you can set up your btrfs partition on your system using gparted if you want. Launch it and navigate to the device that you want to format. You can either reformat existing partitions or make a smaller btrfs partition instead.

03 Go into the terminal
You can also do this in a terminal now the tools are installed. Open it up and use `fdisk -l` if you need to figure out the names of partitions and hard drives on your system. Once that's done, create your btrfs partition with:

```
$ mkfs.btrfs /dev/sda
```

“ You can always check the details of your btrfs partitions ”

04 Create a RAID
One of the best uses for btrfs, due to the way it stores data, is using a similar technique to put a series of hard drives into RAID 0. For now, do it in the terminal with:

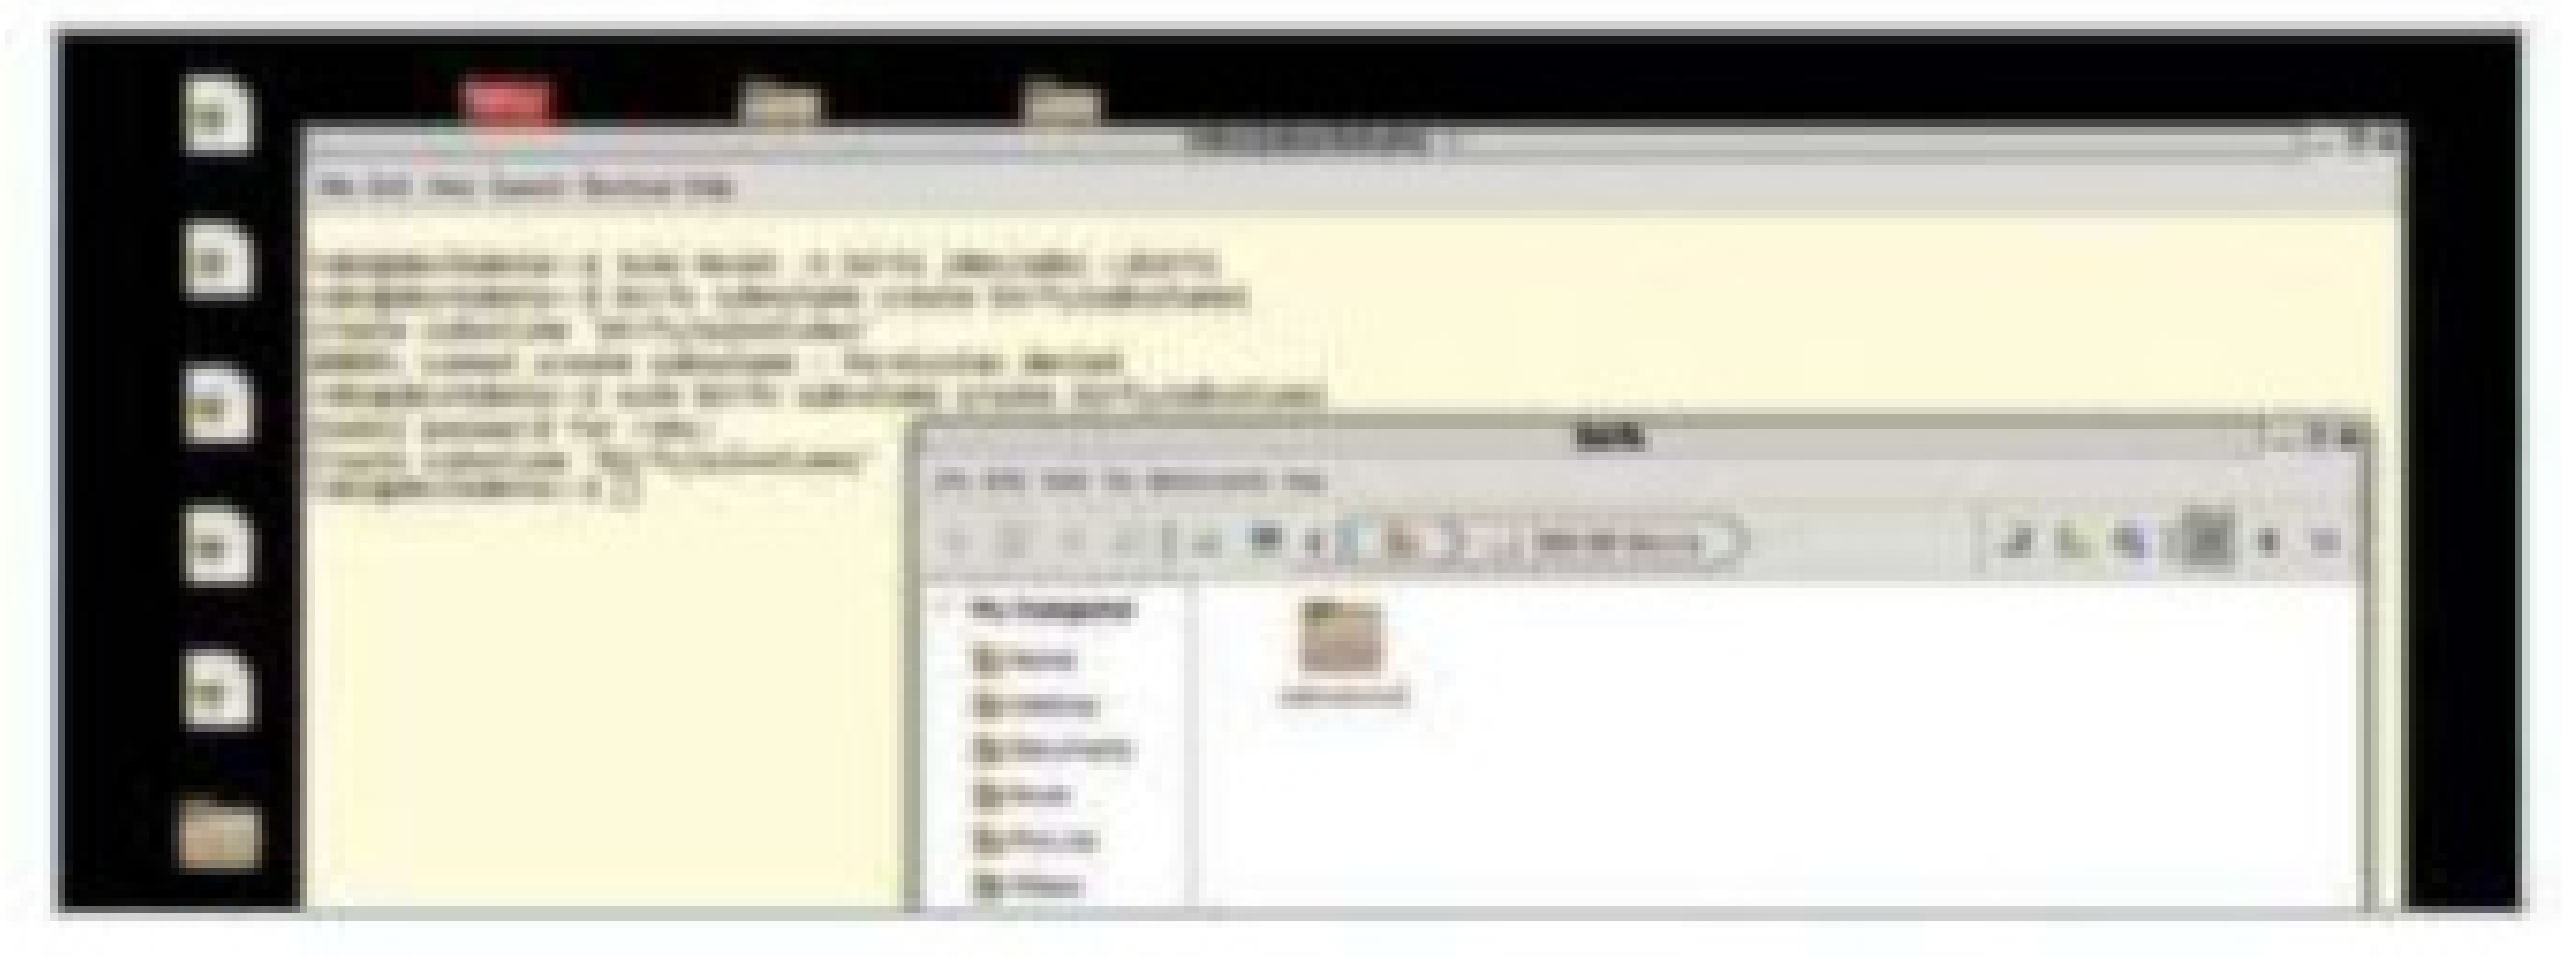
```
$ mkfs.btrfs /dev/sdb /dev/sdc /
```

05 Grab your file system details
Just in case you've forgotten, you can always check the details of your btrfs partitions and hard drives. This also works with individual hard drives in an array, which is why we mentioned RAID in the previous step. In a terminal, use:

```
$ btrfs file system show /dev/sda
```

06 Mount the file system
With most modern distros you can easily mount the btrfs system with a click. However, to do it manually and place it within a specific section of your file system you can do it simply in the terminal with something like:

```
$ mount -t btrfs /dev/sdb ~/btrfs
```

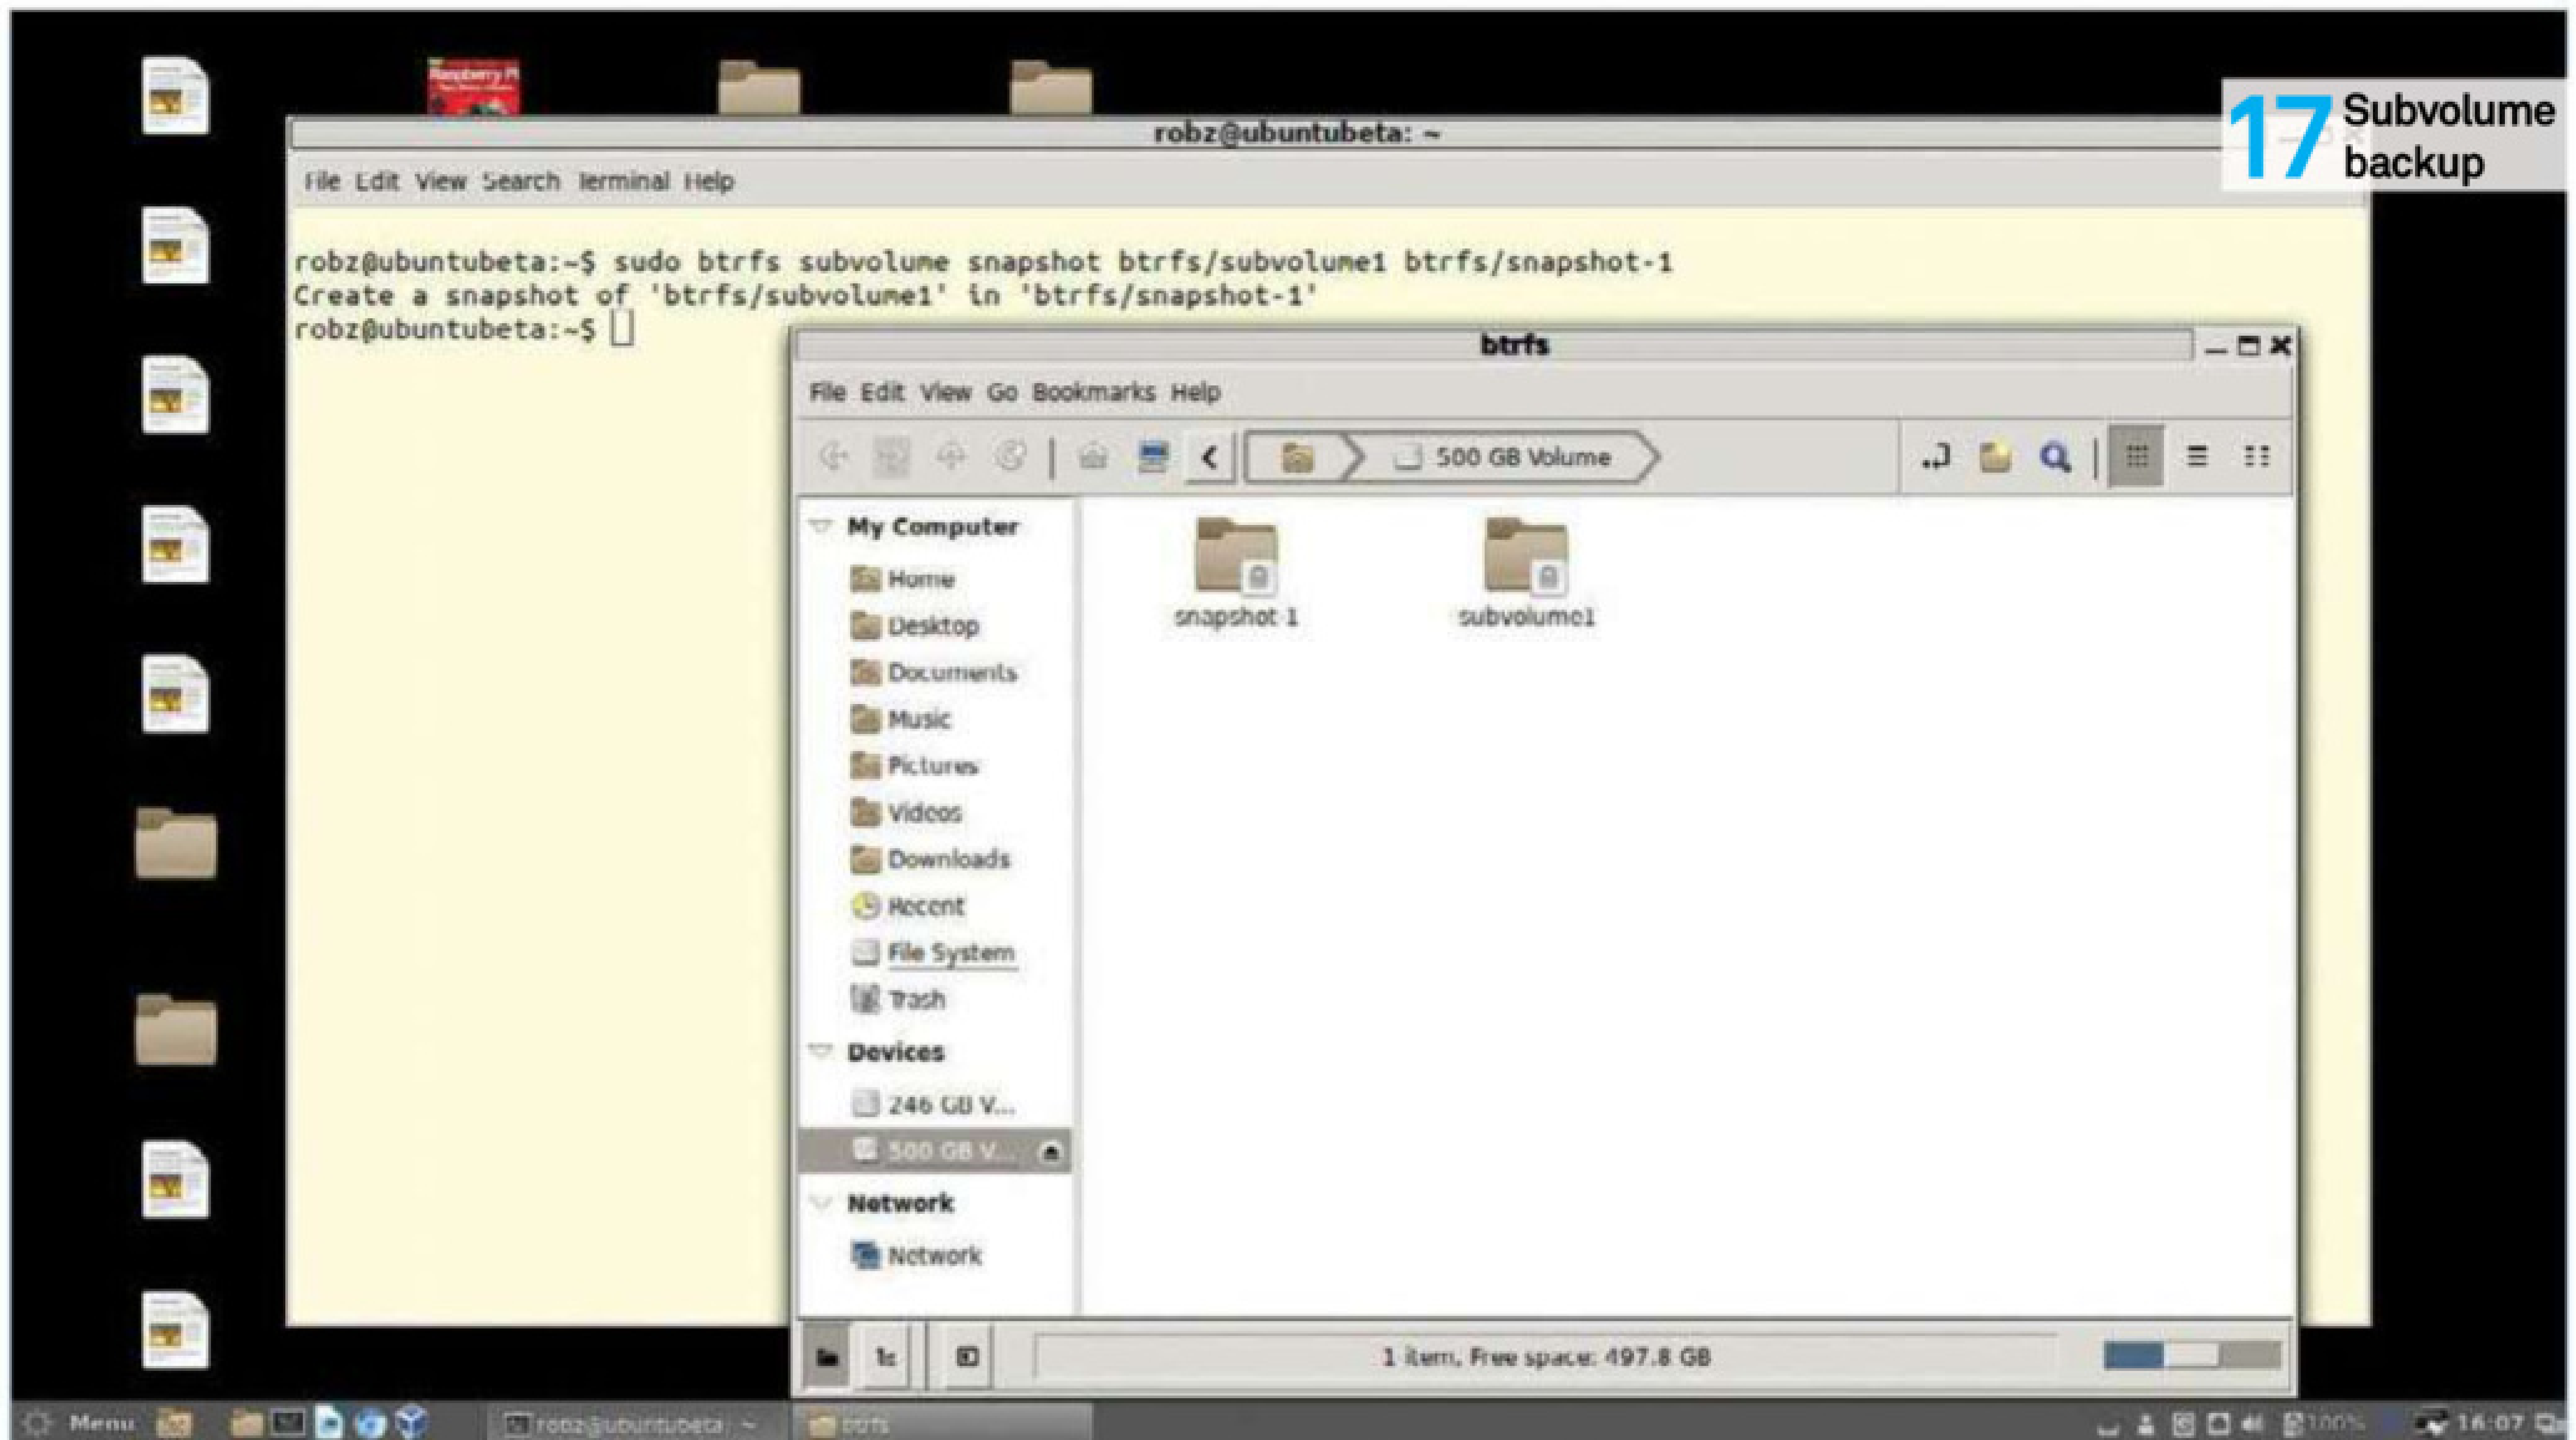


07 Make a subvolume
You can create subvolumes within btrfs that can have different mount options and mount points, and only uses space when they need it. Create one in the terminal using:

```
$ btrfs subvolume create btrfs/subvolume1
```

08 Create a partition
We've created hard drives, small RAIDs and subvolumes, but you can also create btrfs partitions using a similar method as before. Convert a partition into btrfs by going into the terminal and using something like:

```
$ mkfs.btrfs single /dev/sdb2
```

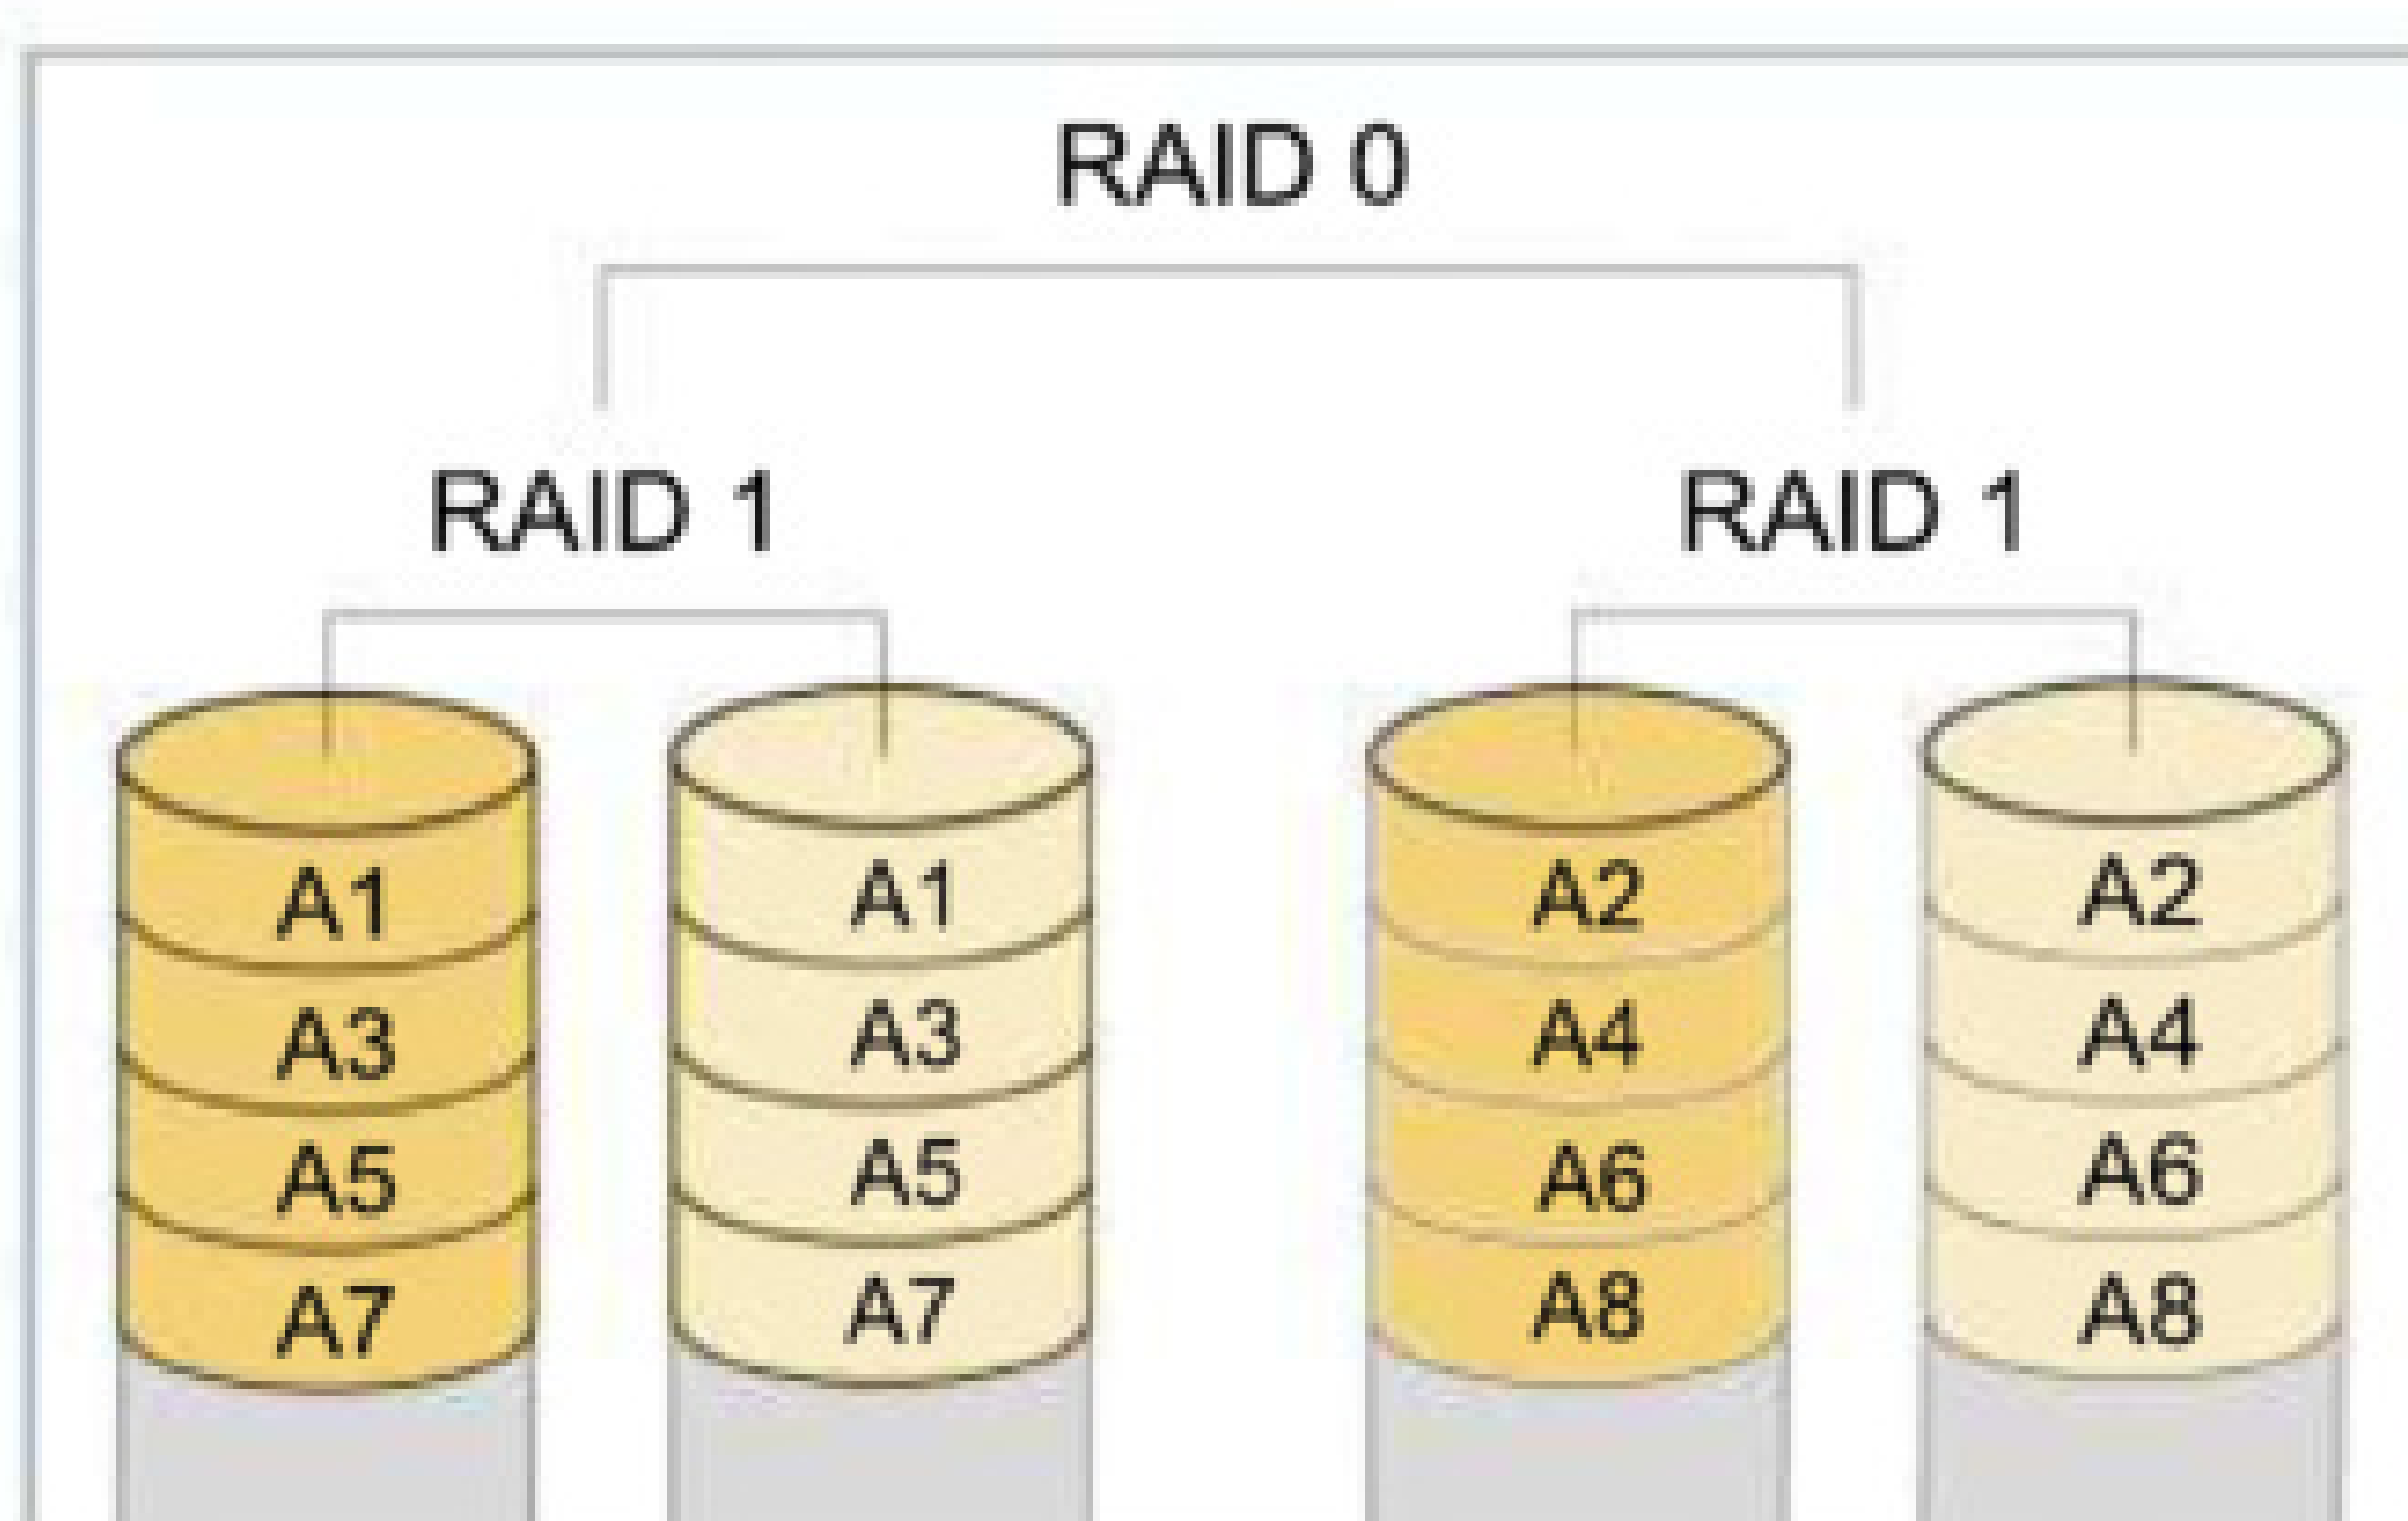



17 Subvolume backup

09 Give permission

The way the drives, partitions and subvolumes have been created means you need to access them with a root account to use them properly. We can give them normal permissions with the following:

```
$ sudo chown 777 btrfs/
```



10 The RAID

You can use btrfs to create different types of RAID setups – by default it will do RAID 0, but it also supports RAID 1, RAID 10 (1+0), RAID 5 and RAID 6. Choosing which RAID to use comes down to personal preference and how much effort you put in to maintain your setup.

“ If you’re doing a hardware RAID, having the metadata be redundant may be useful ”

11 RAID 1

RAID 1, where data is duplicated across drives for redundancy, can be created on btrfs using a modified version of the command used to create btrfs over a selection of drives. The `-d` option refers to how the data will be used.

```
$ mkfs.btrfs -d raid1 /dev/sdb /dev/sdc
```

12 Have non-redundant metadata

If you’re doing a hardware RAID, having the metadata be redundant may be useful – otherwise you can lose your data if the metadata gets corrupted. This requires the `-m` option and can be used on a single hard drive with ‘single’ or RAID 0 with:

```
$ mkfs.btrfs -m raid0 /dev/sdb /dev/sdc /
```

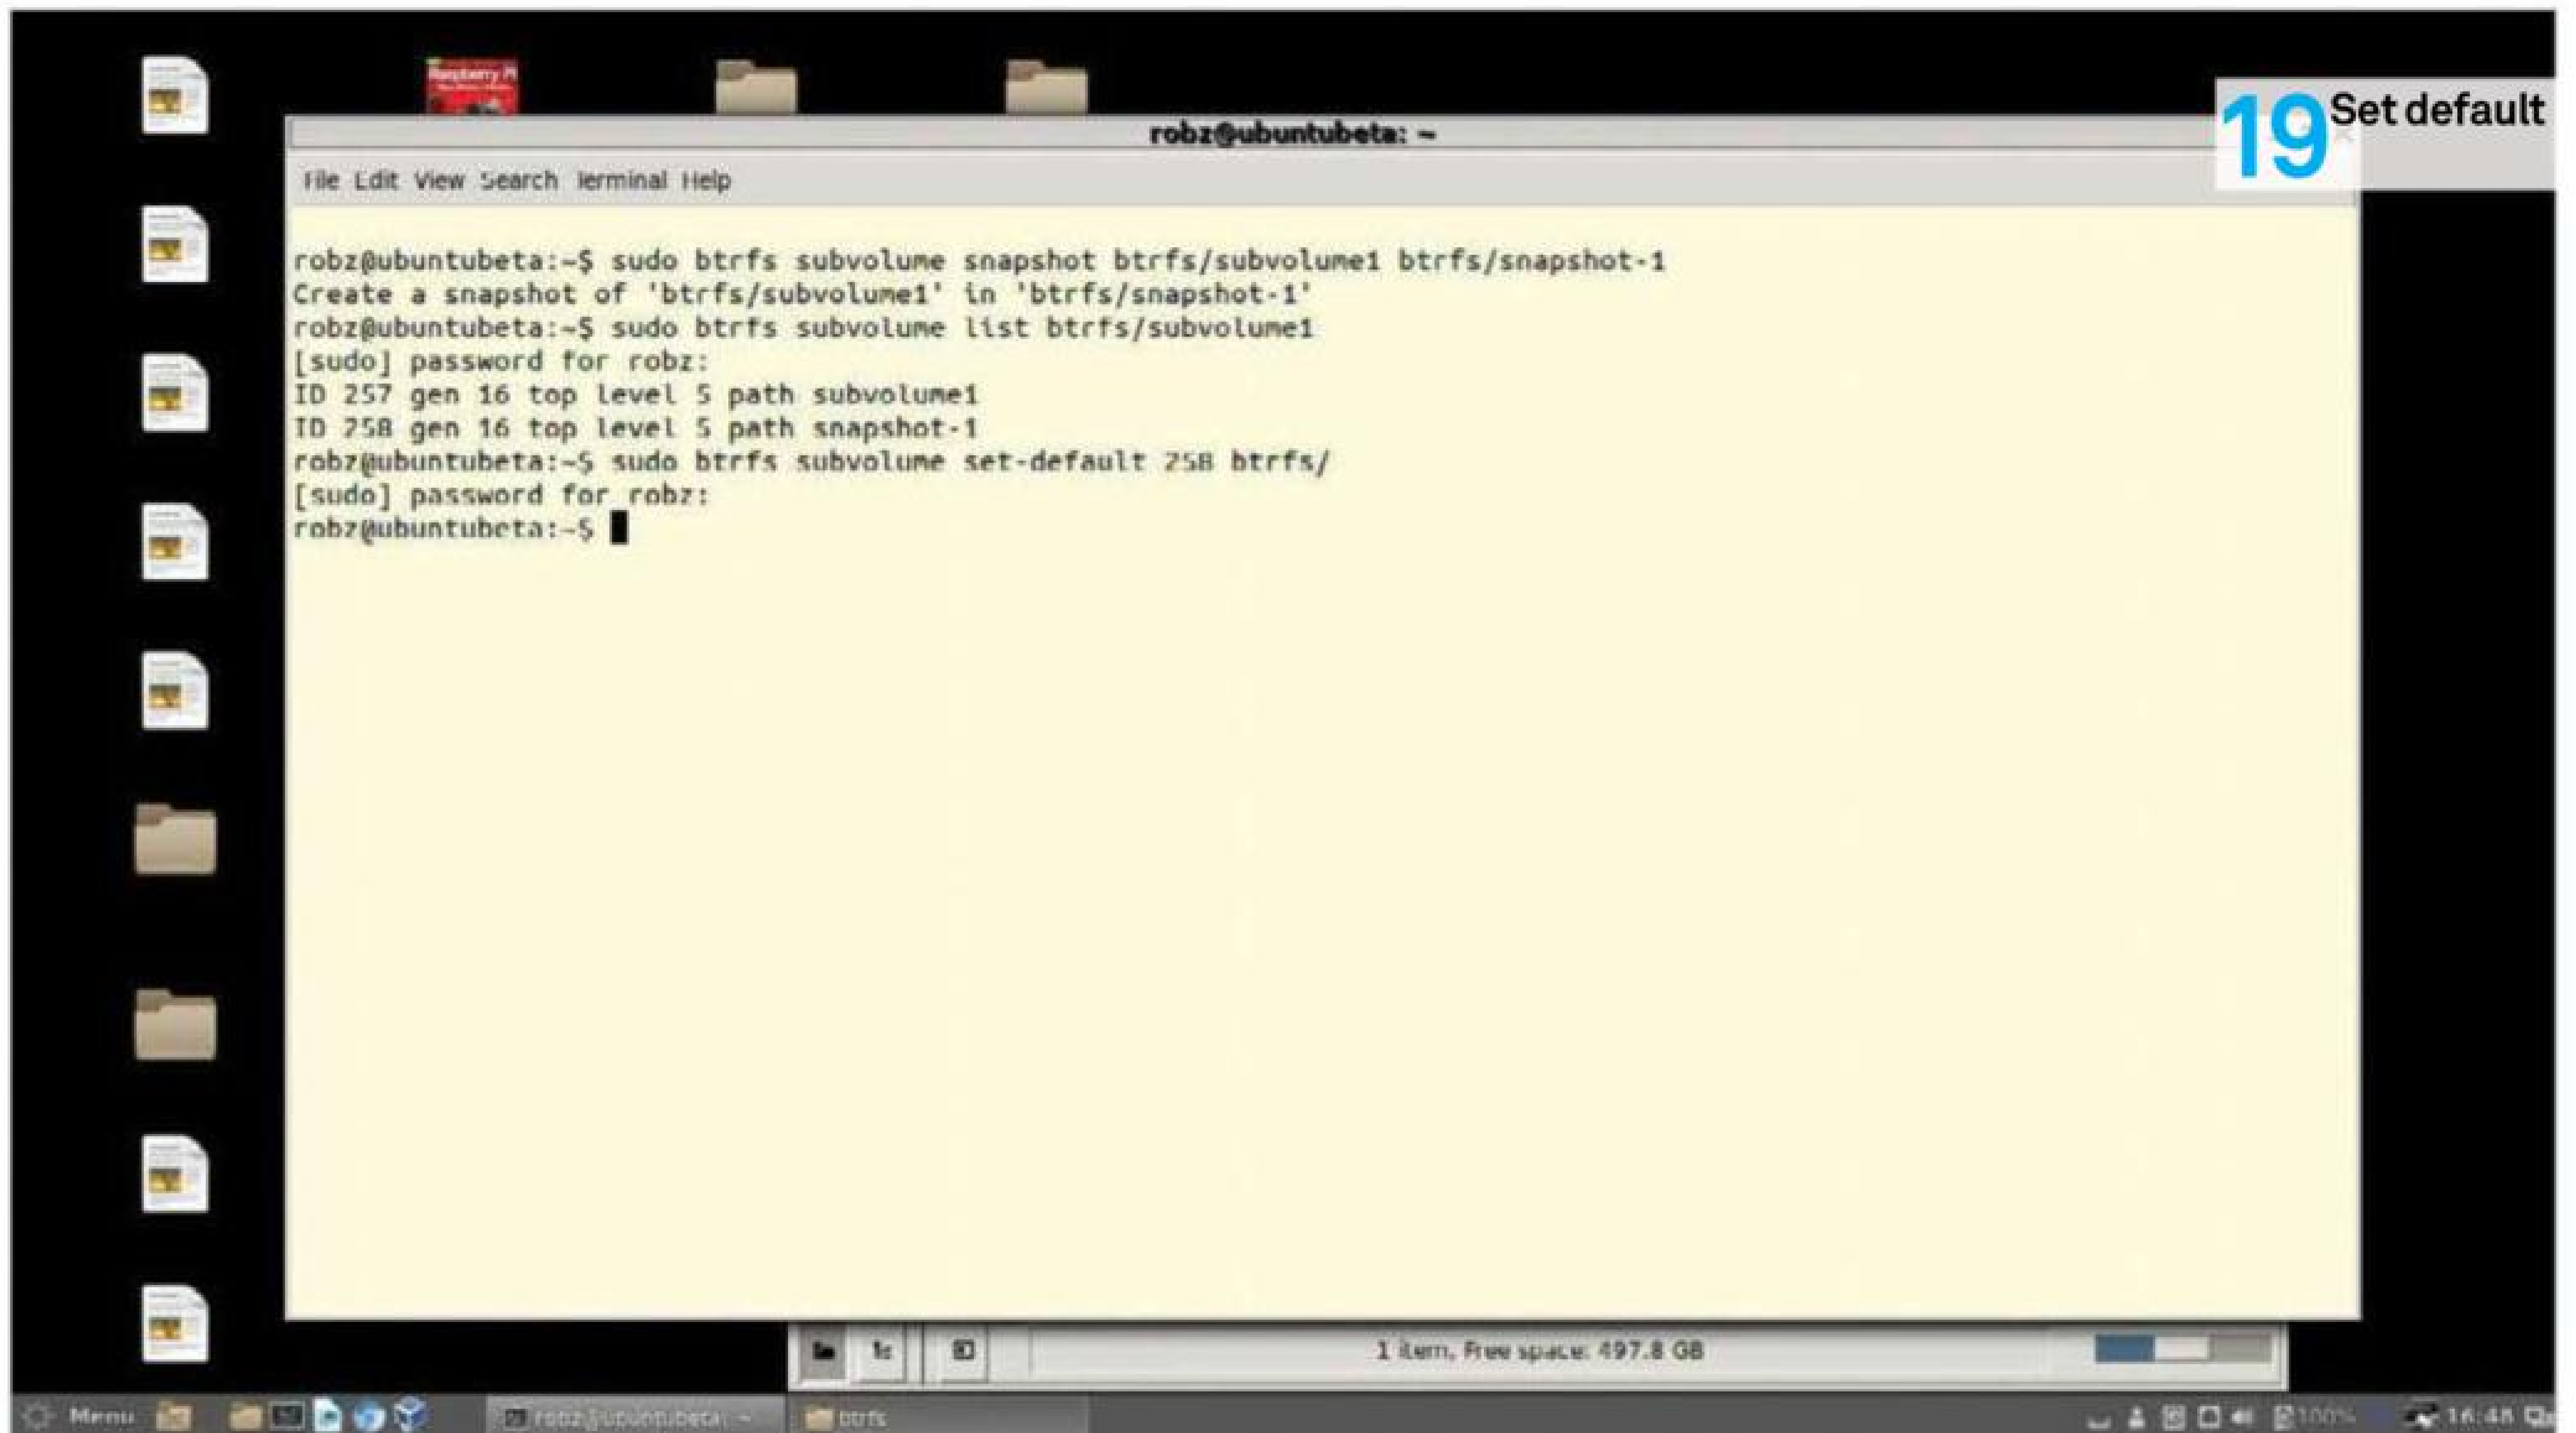
13 RAID 10

Basic mirroring and striping with RAID 10 is good for large data repositories and redundancies. We can use both options `-m` and `-d` to specify that we want metadata and data to be properly made redundant with:

```
$ mkfs.btrfs -m raid10 -d raid10 /dev/sdb /dev/sdc /dev/sdd /dev/sde
```

14 Look at the different options

The options for `-m` and `-d` allow you to do some fancy setups. You can create a btrfs JBOD (Just a Bunch Of Drives) by doing `-d single /dev/sdb /dev/sdc`. You can also have metadata be non-redundant while data is redundant with `-m raid0 -d raid1 /dev/sdb /dev/sdc /dev/sdd`. There are many more ways to do it as you see fit.



15 RAID 5 and 6

RAID 5 and RAID 6 are excellent concepts – using a mixture of striping and mirroring across several drives so that up to 75 per cent can be used for storage while still being safe from hard drive failure. This offers more storage than a standard RAID 10 with the same redundancies (in theory), and RAID 6 increases the redundancy but decreases the overall storage of five drives to 60 per cent of their full capacity.

While it's supported in btrfs there are still bugs, and even at the best of times 5 and 6 can be a hassle. For more details on RAID 5 and 6 in btrfs, visit the documentation <https://btrfs.wiki.kernel.org/index.php/RAID56>.

16 Create snapshots

The way btrfs writes data means that the old version of modified data still exists as long as there's free space – this means you can 'roll back' files up to a certain point. However, you can also create snapshots of subvolumes that are more permanent.

17 Add a subvolume backup

To create this snapshot, head back into the terminal and make sure you know where the subvolume lies. The snapshot will be created in the btrfs partition, but it is best

to put it in a separate location to the original subvolume1 to save confusion:

```
$ btrfs subvolume snapshot btrfs/subvolume1 btrfs/snapshot-1
```



18 Include a roll back snapshot

The individual snapshots can be mounted on their own and you can also replace a newer version with an older version. First of all you will need to list the snapshots that exist for a subvolume, and you can do this in the terminal with:

```
$ btrfs subvolume list btrfs/subvolume1
```

19 Set default

Note down the id for the snapshot (in our case, 258) and then set it as the default for the subvolume1 by using:

```
$ btrfs subvolume set-default 258 btrfs/
```

Then unmount and remount the subvolume to make the changes occur.




20 Remedy defragmentation

Due to the way btrfs writes files, which we outlined in an earlier step, this can cause significant fragmentation. This is a well-known issue though, so btrfs has an in-built defragmentation tool that works even while the volume is mounted. You can use it with:

```
$ btrfs file system defragment btrfs/
```

21 The future file system

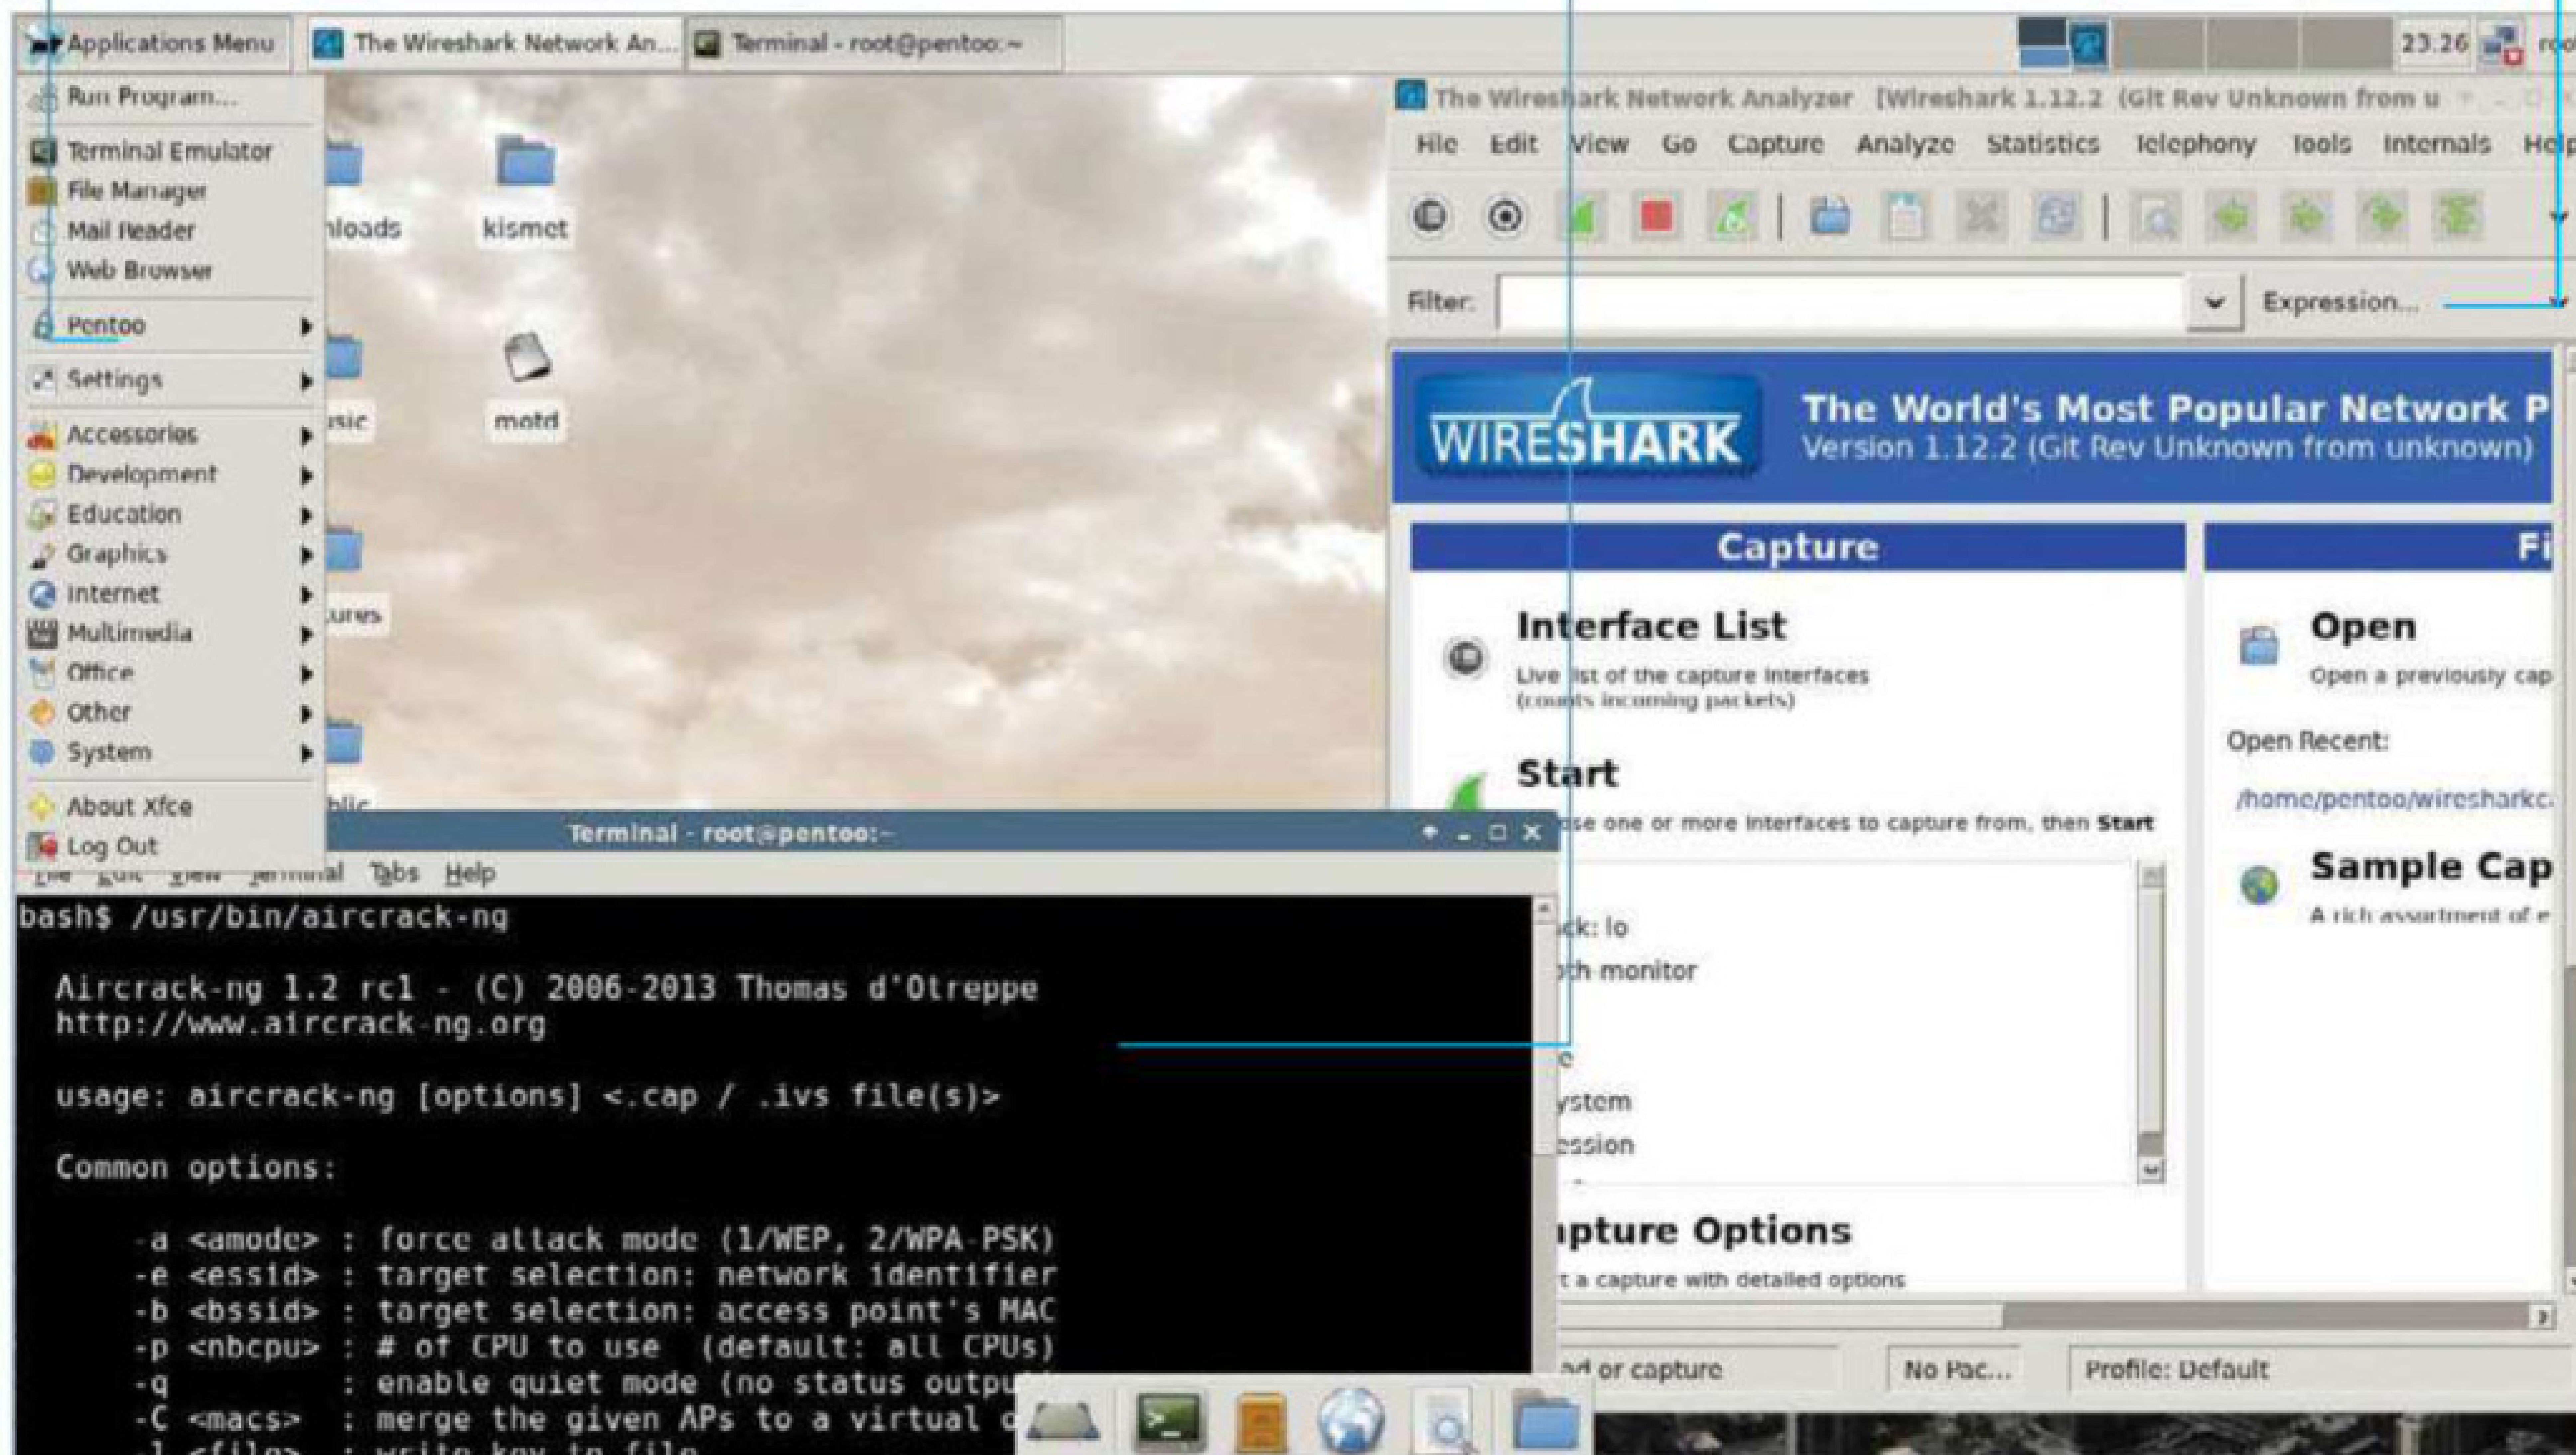
With all these tools you should be well on your way to working with btrfs as a replacement for your other file systems. It is a truly excellent concept for the next generations of fs and it's a great idea to get on board with it now.



There are a range of excellent pentesting apps that are all installed by default

Use aircrack-ng to figure out just how secure your wireless network really is

Wireshark can keep track of the traffic in your system so you can find out if there are any problems



Network penetration testing with Pentoo

Advisor



Rob Zwetsloot models complex systems and is a web developer proficient in Python, Django and PHP. He loves to experiment with computing

The Gentoo-based utility distro can find vulnerabilities in your network

Resources

Pentoo www.pentoo.ch

Over the past few issues we've had tutorials and features on increasing your privacy and security. It's important to know how to shield yourself and your systems from attacks, and now we're going to look at the problem from the other angle: how someone would attempt to attack your system.

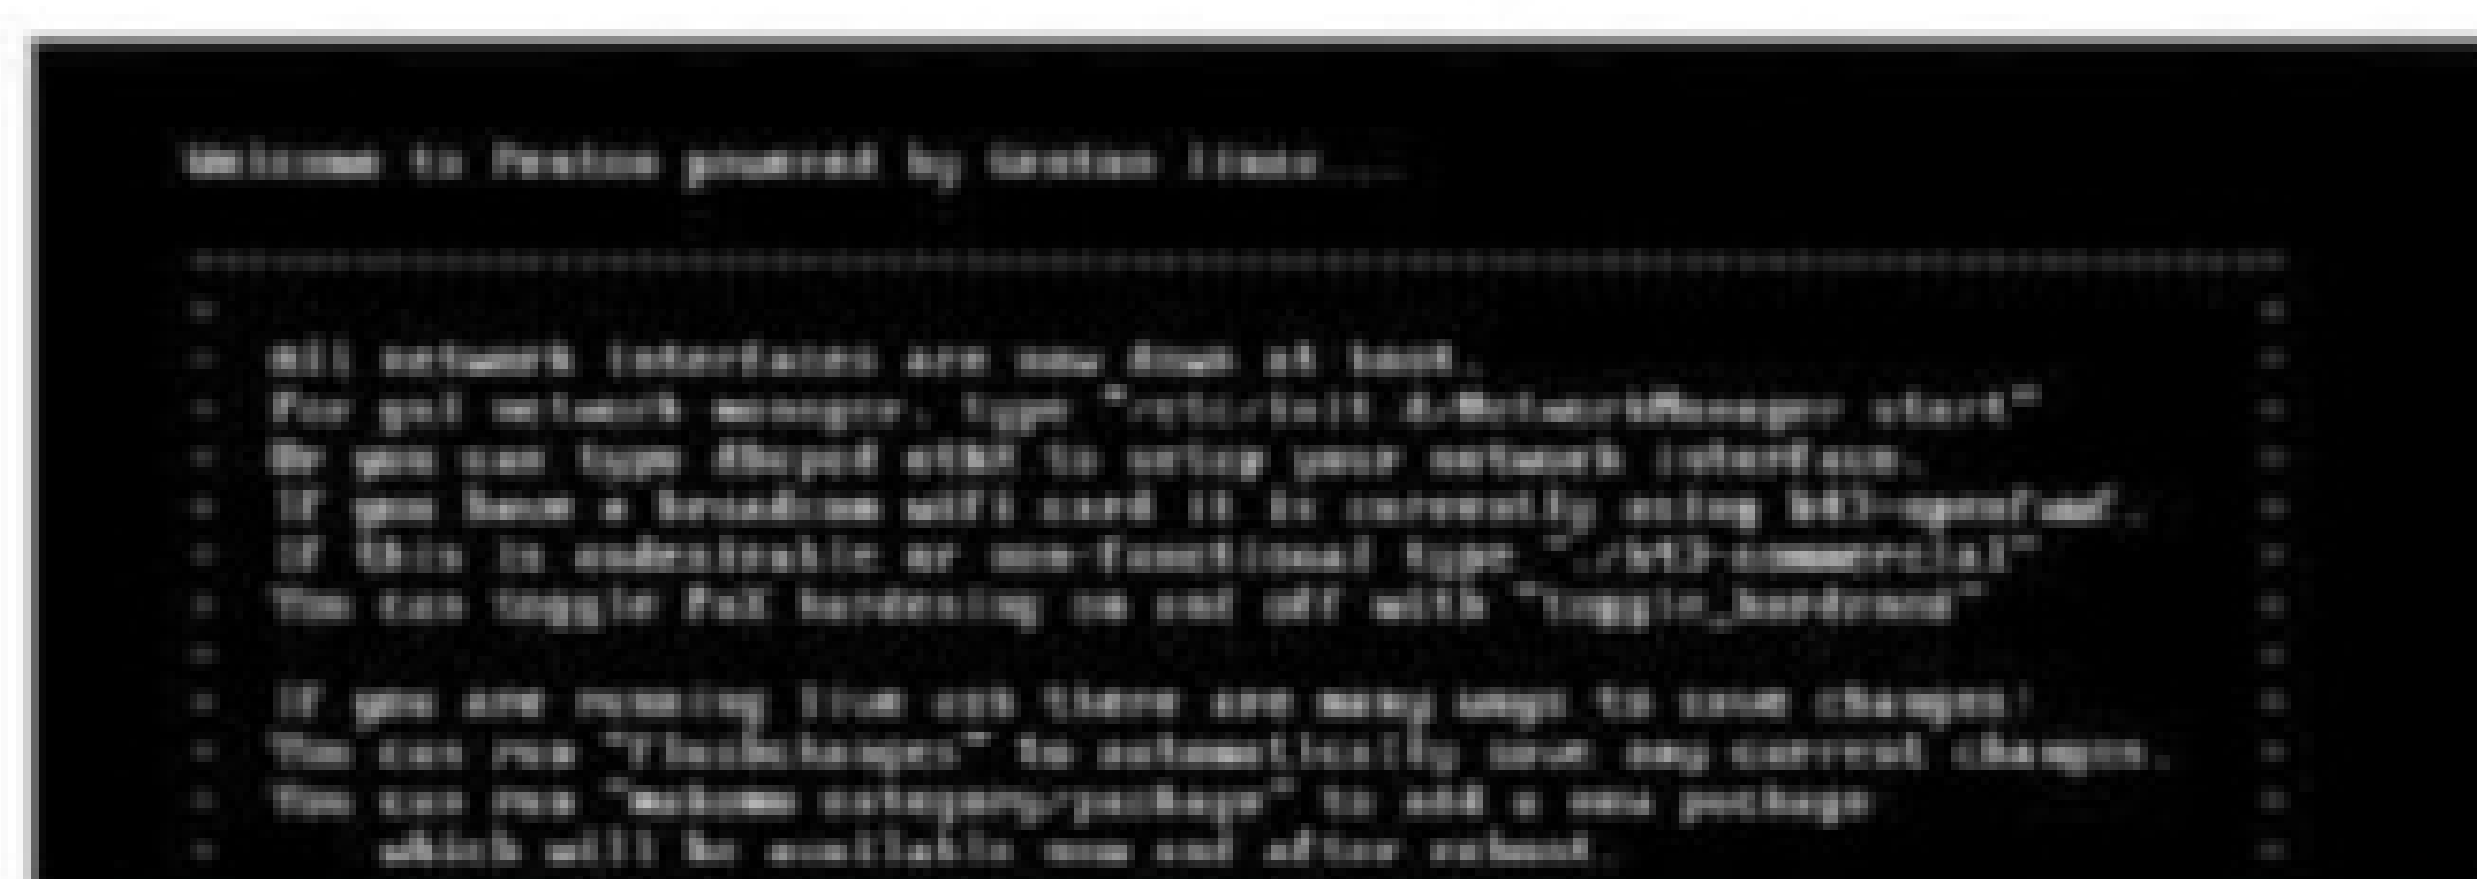
White hat or ethical hacking is an important step in any security process, and we're going to take a quick look at probing your own local network to try and figure out any problems you

might have. There are a few different tools and distros that can aid you in this, but today we are going to focus on Pentoo.

Pentoo is less known than Kali or Backtrack Linux, but it's by no means a poorer product. Built on Gentoo, it can be used perfectly well live and can be installed for a more permanent testing rig. While it has a full suite of pentesting software, we'll be concentrating on some of the network testing apps for the purpose of exposing any vulnerabilities in your network.

Network penetration testing with Pentoo

Find vulnerabilities in your network with this Gentoo-based utility distro



01 Grab and load up Pentoo

Get the ISO of Pentoo from the distro's website (www.pentoo.ch) and write it to disc or USB. Reboot into the distro and either press Enter if you use the Colony's English or type 40 for the Queen's variety. Use `startx` at the command prompt to then launch the desktop.

02 Explore applications

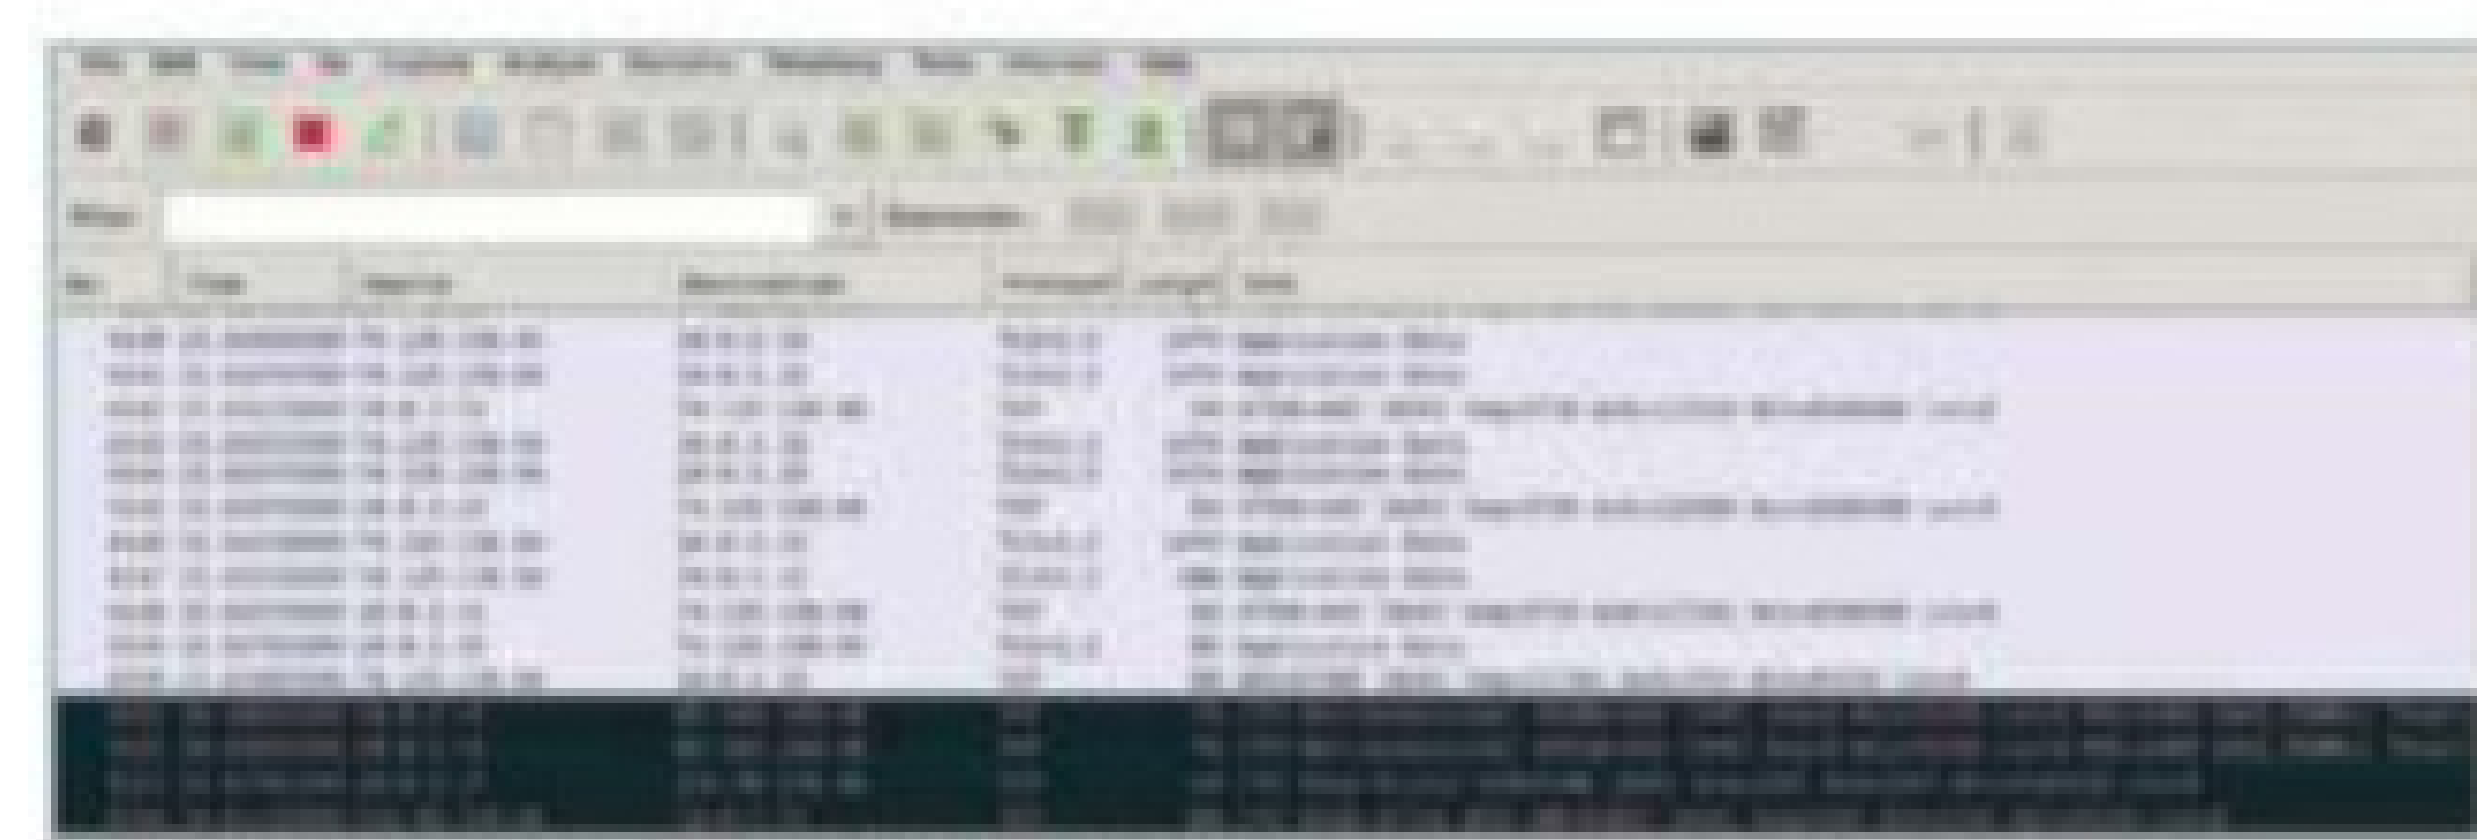
All the analysis and pentesting packages are kept in the Pentoo section of the main menu. Pentoo can be used for all manner of testing from computers to networks to websites and more. For now, we're interested in Wireshark, which is in the Analyzer submenu.

03 Use the interface

Wireshark is presented on a web-like interface as you interact with the backend of the software. You can begin capturing the traffic from any number of interface devices to see what exactly can be seen within the network.

04 More options

Before we start looking at the data, we can also make some more advanced settings on what the capture will grab. The basic settings will get a lot of data while it snoops the network, but you can set where it stores the files and if they should be split up. For now though, click Start.



05 Capture basic data

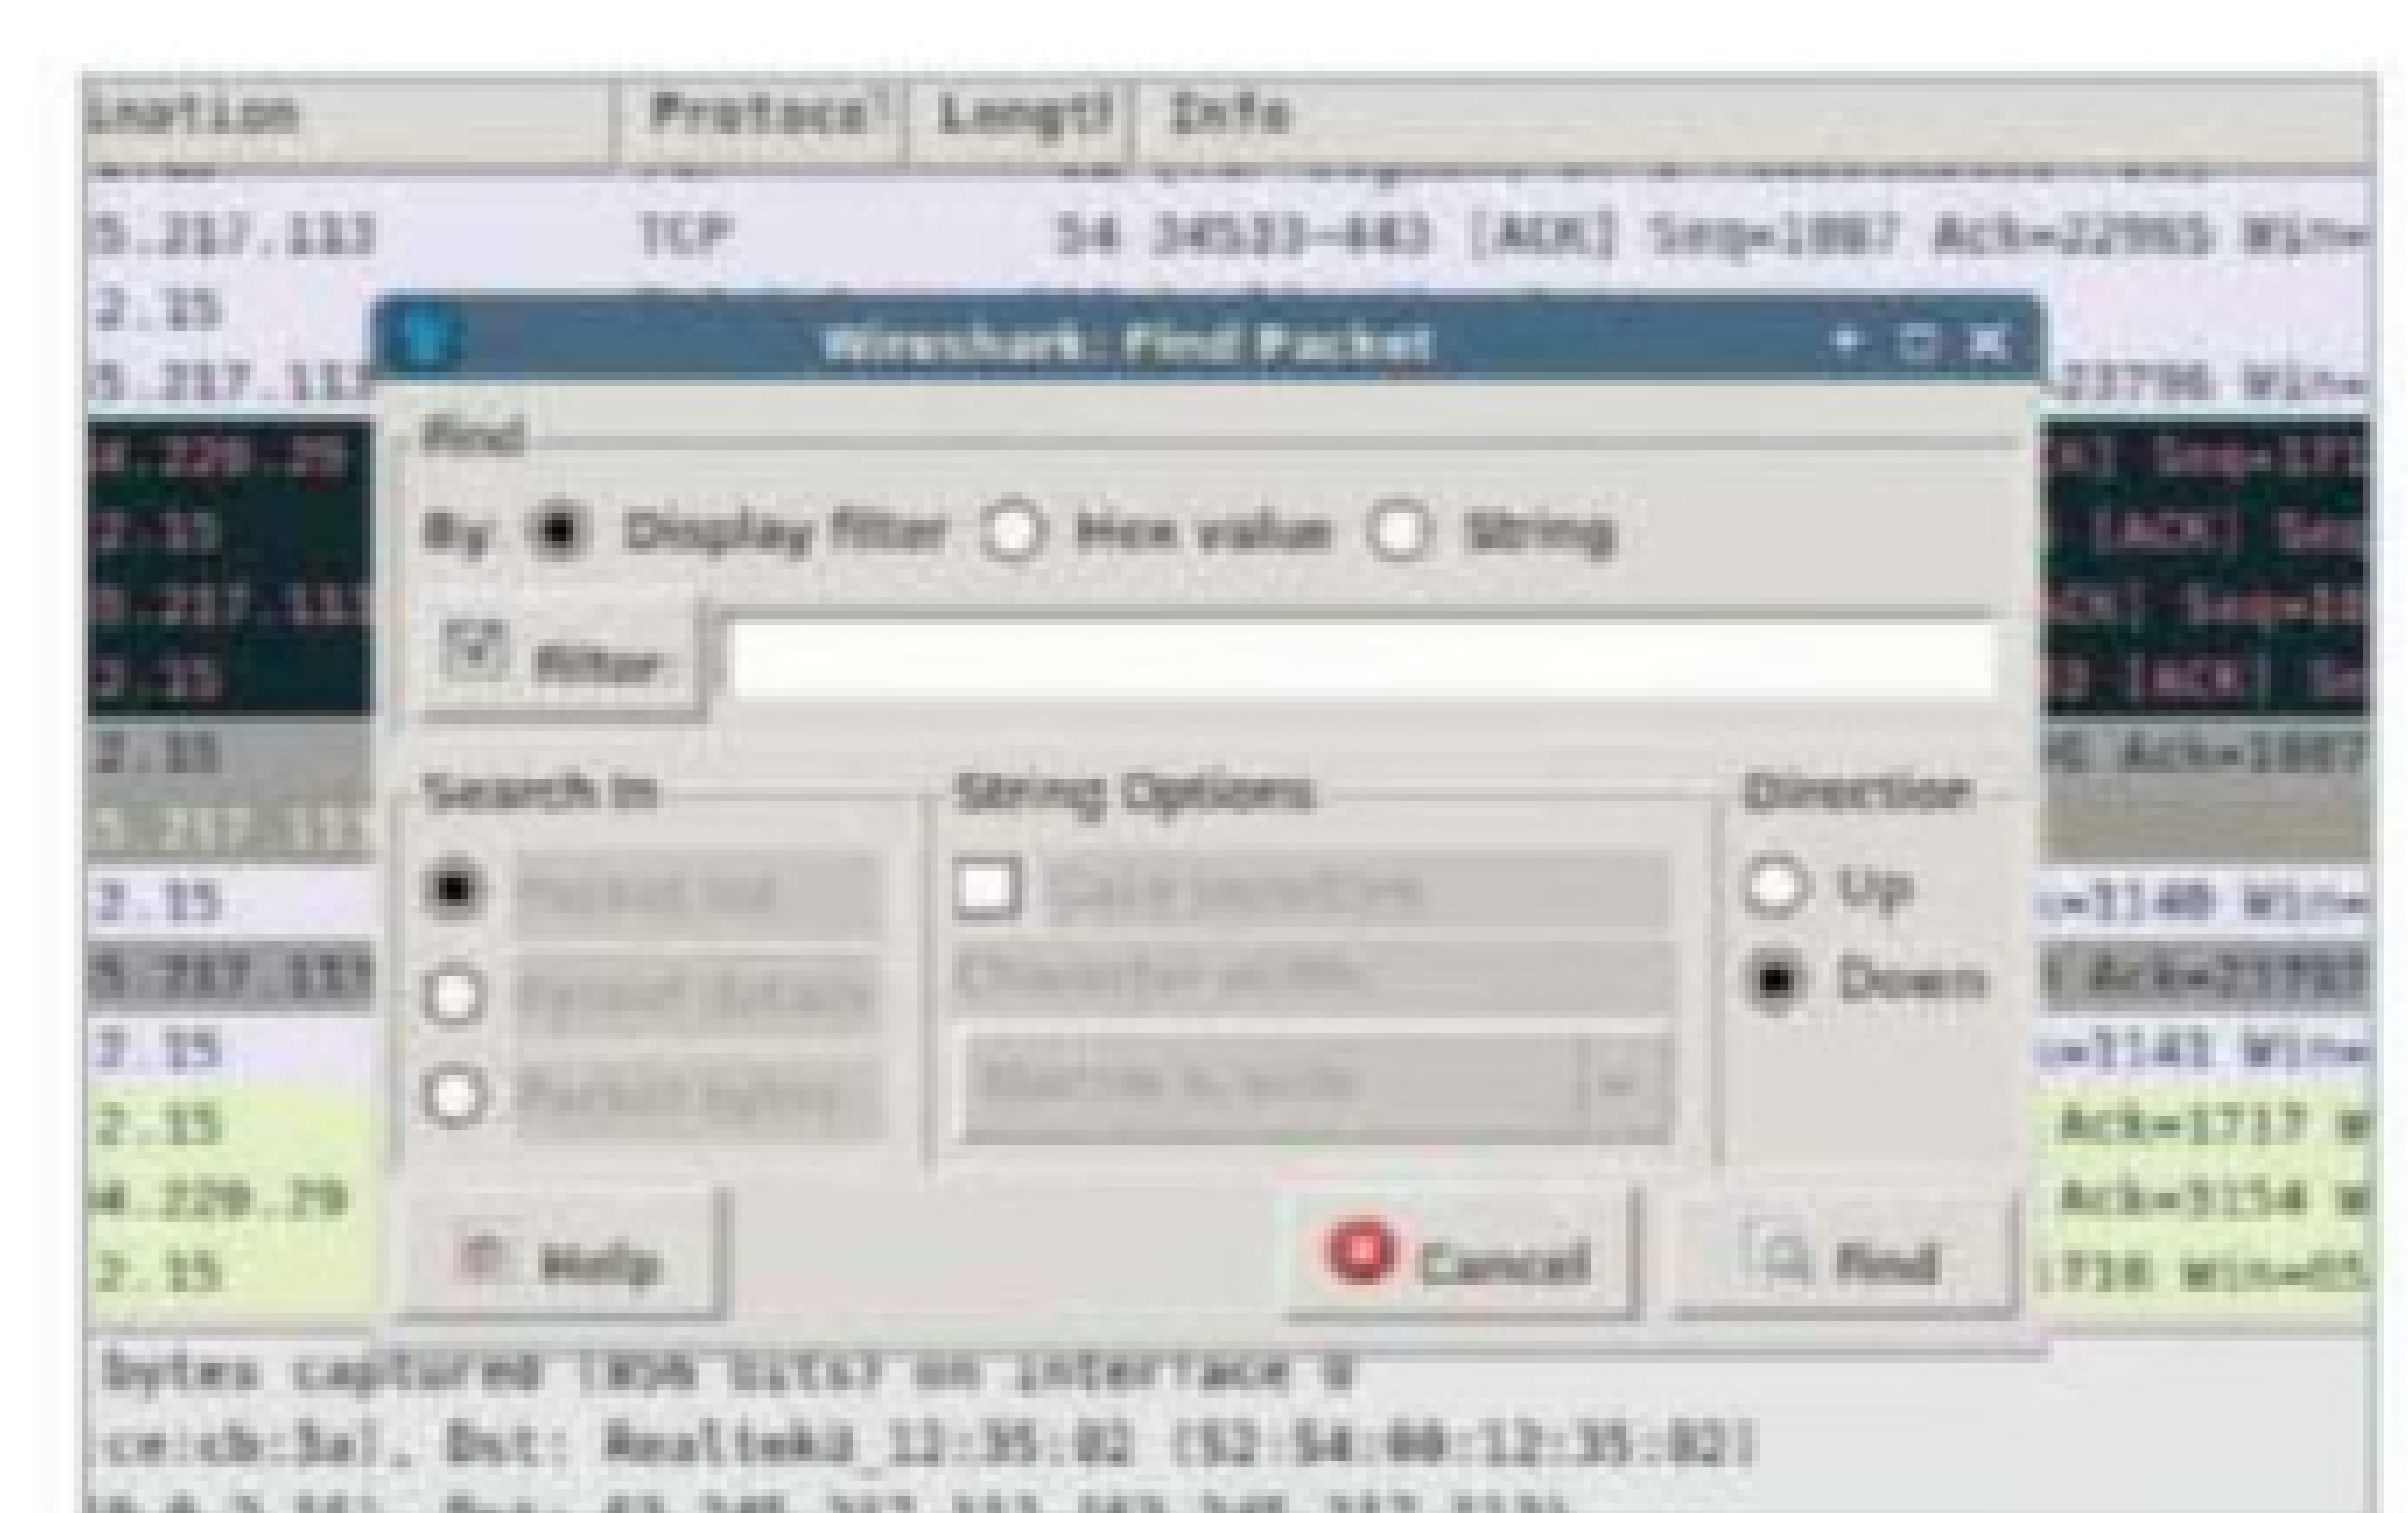
You'll start to see any data that crosses the interfaces with plenty of information on the packets, their origin and their destination. At a basic level you can see if there are any untoward packets being sent in or out – good for checking your own sites and LAN sites as well as the system in general.

06 Filter your data

For basic capture data like this, there will be a lot to sift through by default. You can filter it while it's running or stop it to inspect anything you might have been trying to do during a test. Apply regular expressions to narrow anything down.

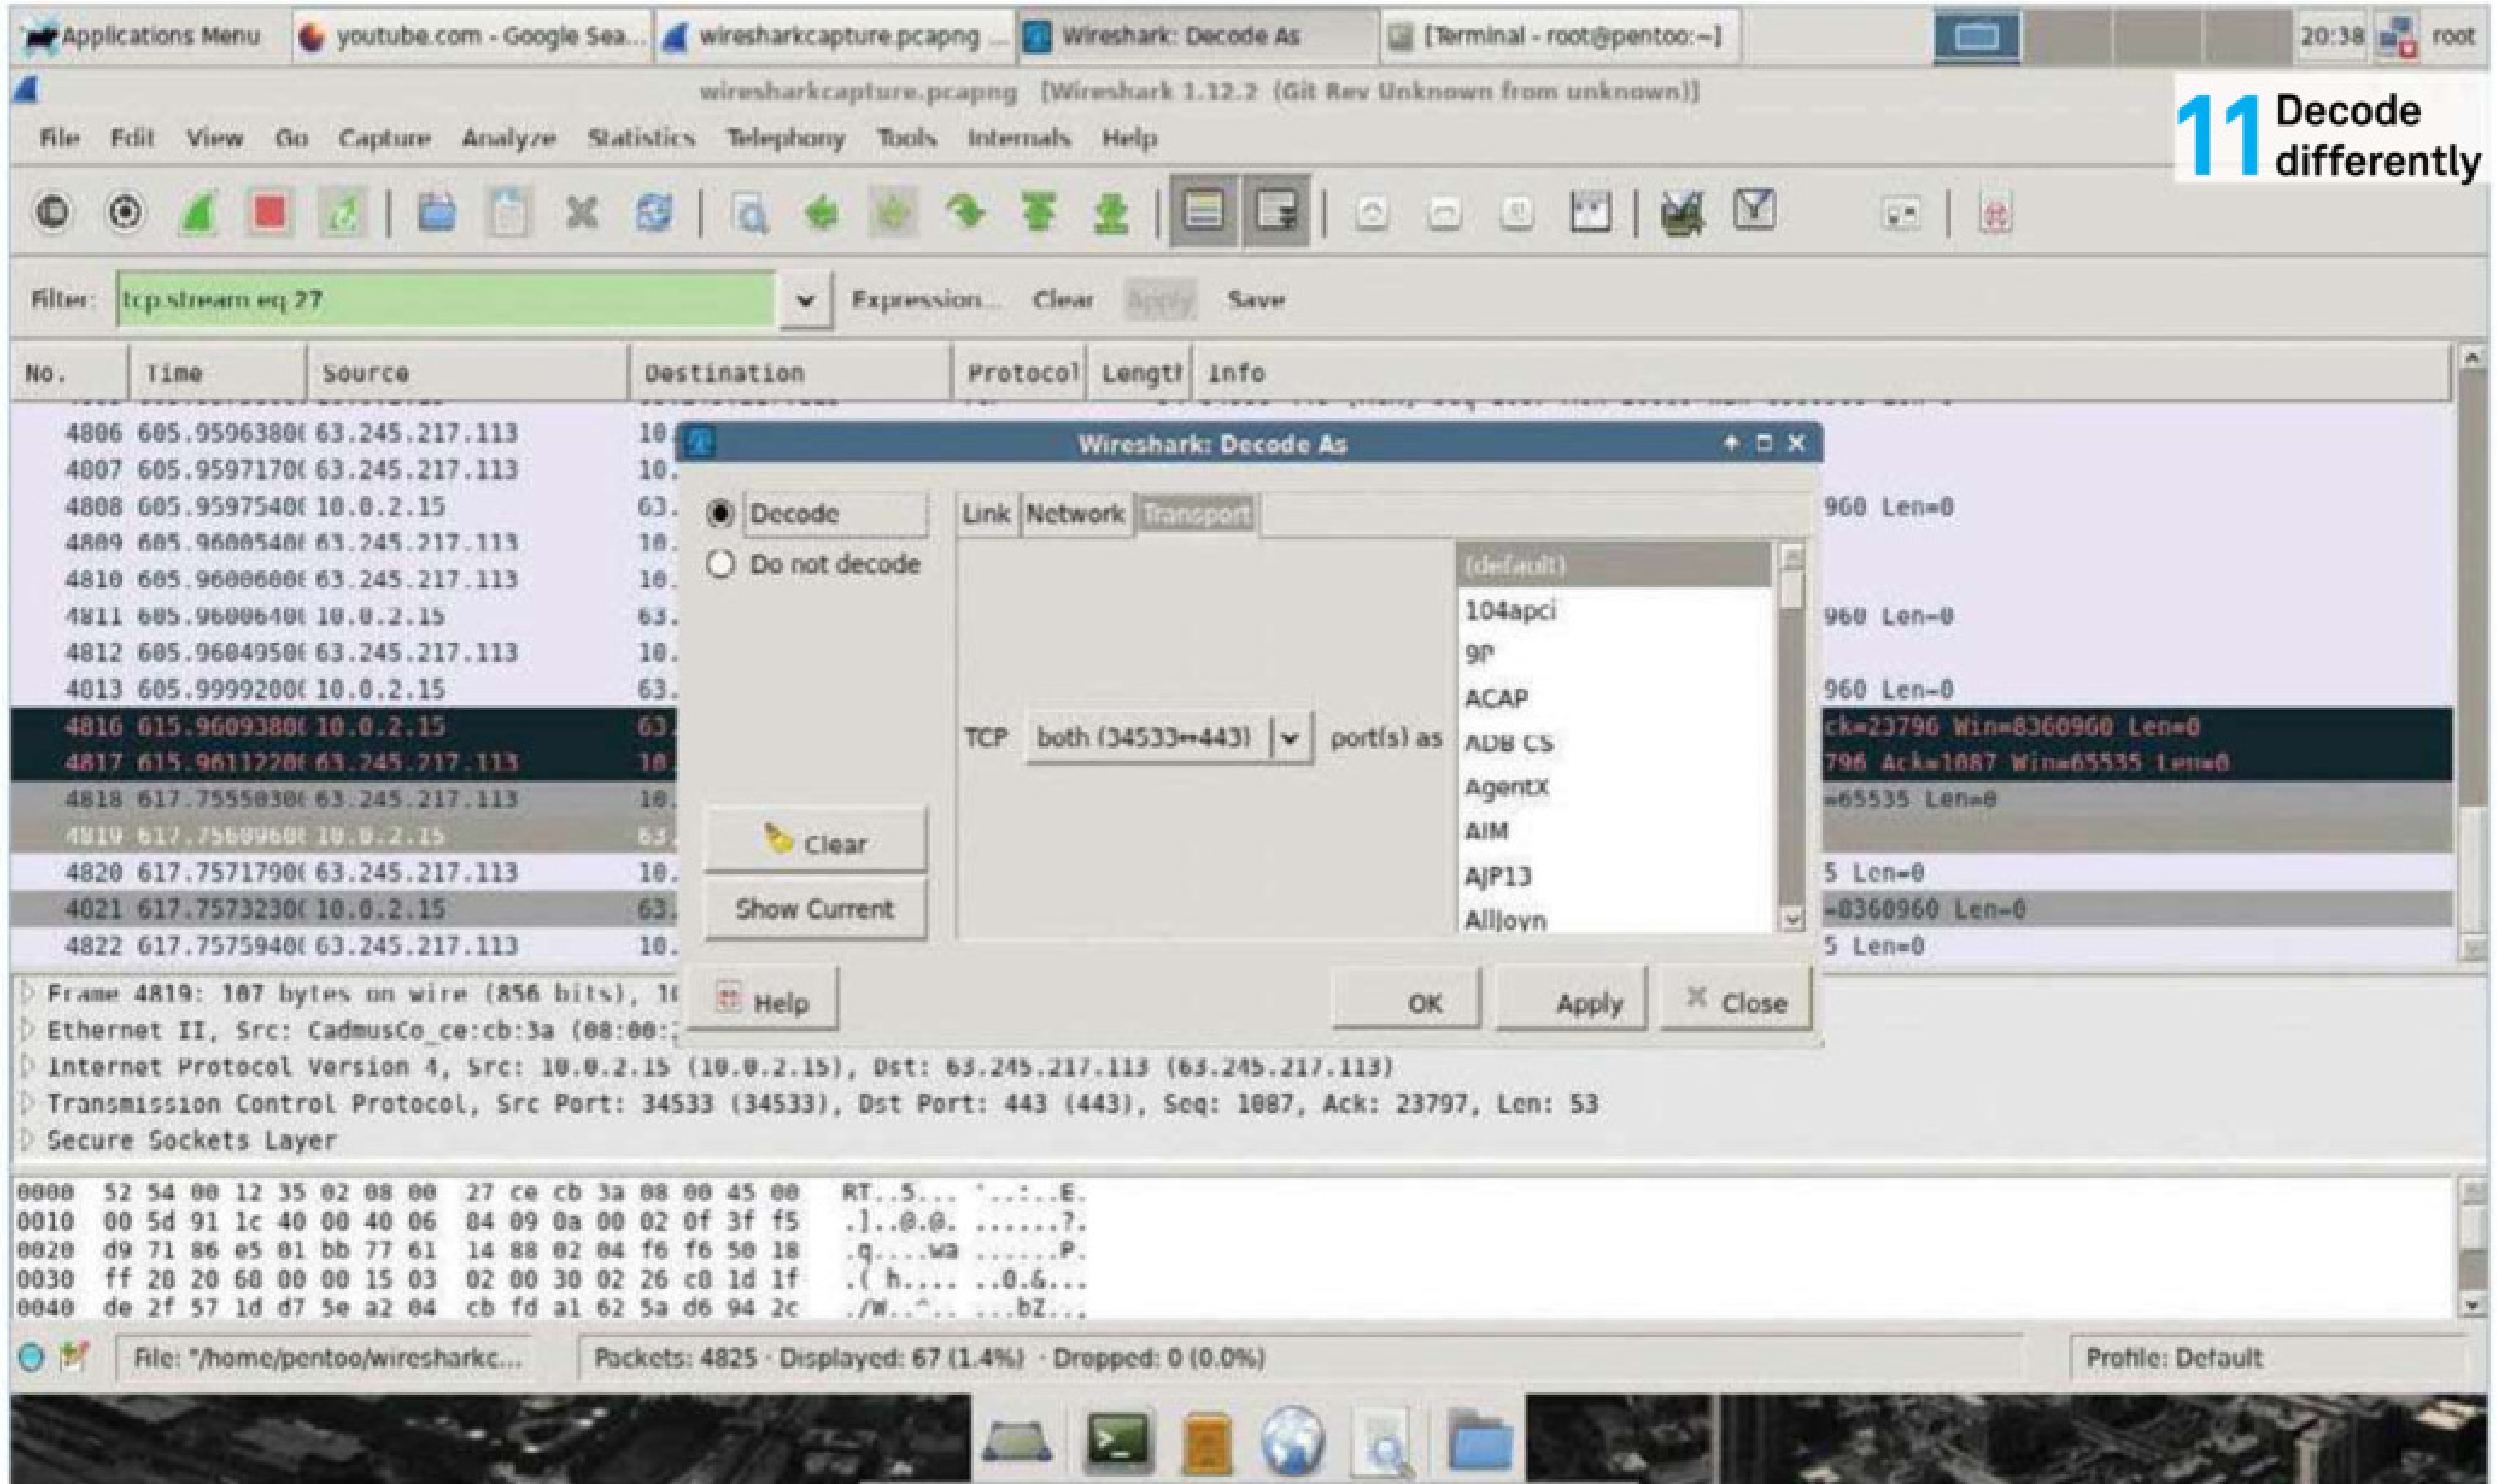
07 Save and retrieve

Captures you've made can then be saved to file as you would with any normal type of file. The benefit of this is that you can reopen it in the future in the same capture window, to then search through in case you need to either compare something or remember a packet.



08 Find individual packets

Sent a very specific packet during or as part of your test? You can search for a known packet in the logs using the Find Packet tool to look further into the records. It requires accurate knowledge though, so you'll need to be well aware of what you sent.



11 Decode differently

09 Use VoIP tracking
You can also see the VoIP traffic on a system under the Telephony menu on the tool bar. This will show the caller and receiver, what protocol they're using what kind of data is being transmitted and more.



10 Endpoint detection
One of the weaknesses in a network can be remote machines that access a network. When going through a capture, you can see what endpoints have been connecting to the network and filter out which ones are authorised or not. You can also set up a remote attack and see if it registers.

11 Decode differently
Sometimes odd packets and messages may occur on a different port that Wireshark has a problem translating. This is simple to fix, you just need to have a look at the package using a different protocol with the Decode As tool. It may just be a random connection but it could be something else which will definitely need sorting.

12 Wireless defence test
Pentoo has an excellent suite of software called aircrack-ng that can attempt to discover and break a wireless network. This software tries to sniff packets on the network to figure it out. To start we need to make our wireless card able to see the packets, so we need to use airmon-ng.

13 Employ promiscuous mode
Open up the terminal to use airmon-ng. Check ifconfig if you need to figure out the interface name of your card, and then turn it to the correct mode using something like:
`$ airmon-ng start wlan0`

14 Find an access point
Airmmon will change your wlan to now be mon0, which is what we're going to use to try and find an access point. This is useful if you've tried to make your Wi-Fi hidden but it's also useful to see just how powerful your network is around a building. Start it with:

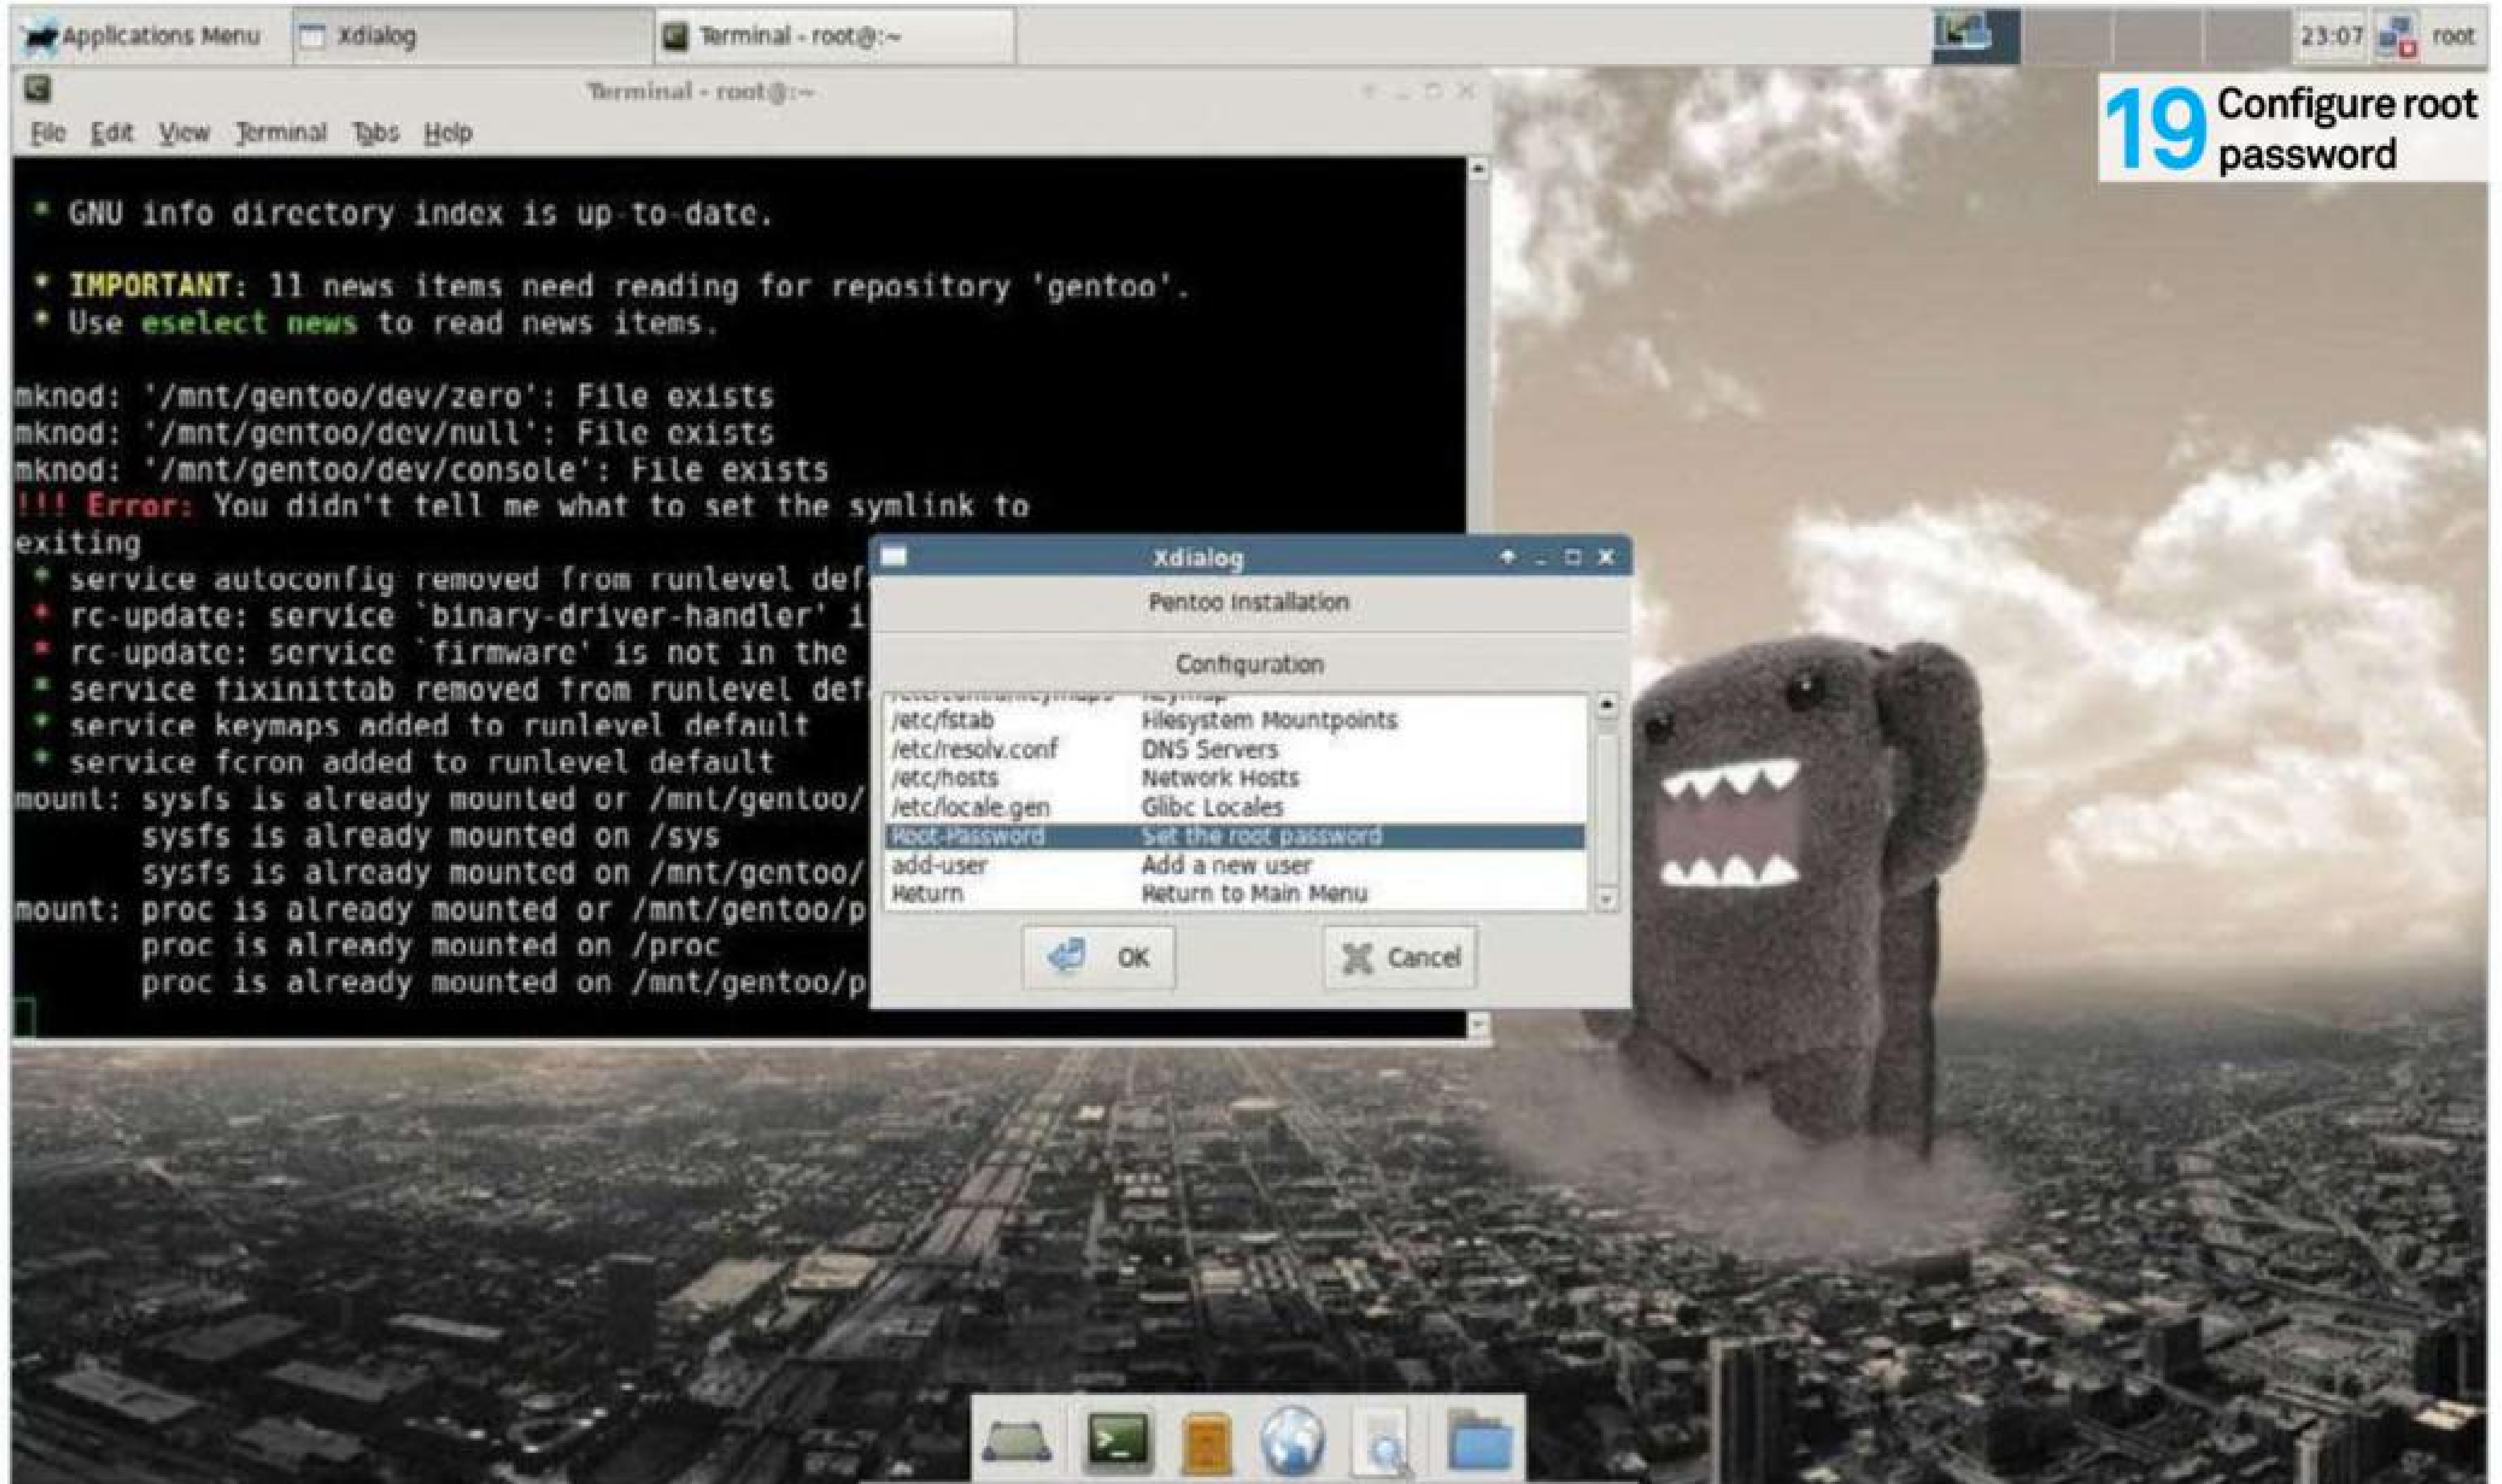
```
$ airodump-ng mon0
```



15 Test your passwords
You can now try and use aircrack-ng to brute-force your passwords and see if they're strong enough. The basic way of doing this is to use something like:

```
$ aircrack-ng -b [MAC address/BSSID]
```

This will try and crack the password on the access point. There are various options you

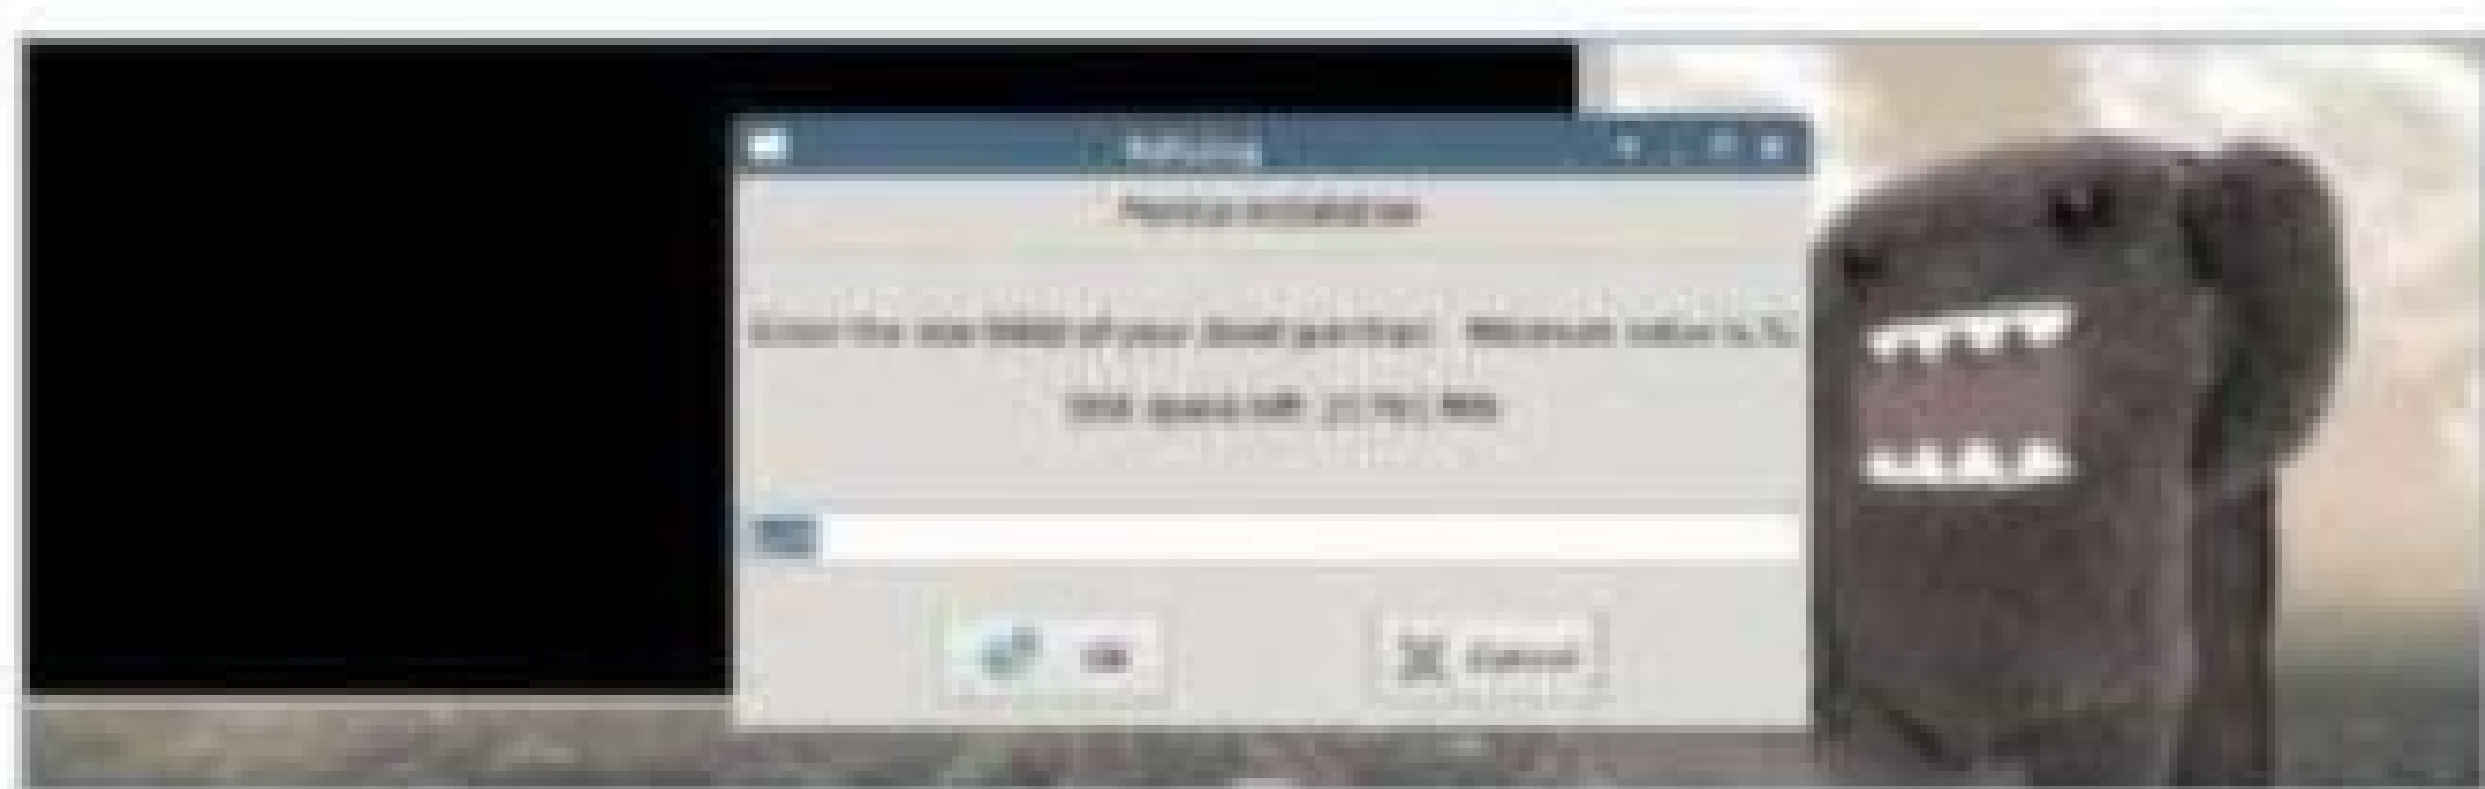


19 Configure root password

can use that will limit CPUs, focus on a specific password types and even bring in a list of dictionary words and common passwords. The latter is useful if you're checking out someone else's network for them. You can see more options in the main file.

16 Install Pentoo

Need a more permanent version of Pentoo that you can plug into your system at a moment's notice? You can install the software to your computer with the installer, however it's not quite as simple as something like Ubuntu. Click on the installer to start.



17 Set hard drive

Option one has you set the hard drive options. It's a bit more involved as you directly set the boot size, swap and the storage. Its recommendations can be good but it won't enable you to do any advanced partitioning.

“ Pentoo has an excellent suite called aircrack-ng ”

18 Copy the distribution

Working on the hard drive is the first step to the installation of Pentoo. So once the hard drive has been fully worked on, you will be able to perform the automatic install. Select option two and wait a while – it will unpack the files and install them and also ask if you would like to transfer any settings from your current setup.

19 Configure root password

Go to configure the system and start adding users and configuring the root password. Without the root password set, you could have great difficulty logging in once the installation has finished. Go back to the main menu, install GRUB as the bootloader and then reboot.



20 New Pentoo


You will boot into your fresh copy of Pentoo, able to be used to diagnose issues with your networks and other people's networks with great efficiency. As it's Gentoo, you'll have to understand how portage works to start updating, but it's very easy once you know how.



Program a client-server application

Each operating system provides its own set of networking APIs, but Qt gives a solution to simplify creating chatty applications

Advisor



Tam Hanna has been in the IT business since the days of the Palm IIIc. Serving as journalist, tutor, speaker and author of scientific books, he has seen every aspect of the mobile market more than once

Modern computer networks are composed of multiple protocols. Creating well-formed packets by hand is almost impossible – the protocol stack and its underlying network hardware affect the format.

Engineers working on Unix solved this problem by introducing an abstraction layer – so-called sockets acting as endpoints between two local or remote applications. Developers use the standardised interface in order to provide commands to the socket, which are then translated to the network.

When dealing with sockets, a basic understanding of networking is beneficial. For now, it shall suffice to define that messages can be transmitted by two high-level protocols. TCP – short for Transmission Control Protocol – provides confirmation when data has been delivered successfully. This is accomplished by a relatively complex sequence of packets which add overhead to the communication process. However, some applications don't require this. For them, UDP provides a sleeker communication protocol that forgoes delivery confirmation. On a network, computers are identified by their IP address. Individual services are then identified by their port numbers – consider them apartment door numbers inside a high-rise building.

Developers working on communication software tend to face a chicken-egg situation – if no server has been written, the client can't connect. On the other hand, the lack of a client means that the server can't be tested.

In practical projects you should start out by creating a rough mock-up of the messages that the client and server will exchange. You don't need to get this 100 per cent right on the first try; most protocols change as they get implemented.

Our finger exercise-level applications don't require this level of sophistication. It will suffice to begin by creating a command line project called ImagineTCPServer. Qt Creator does not add the networking module to newly created applications. This can easily be remedied by opening its .pro file and adjusting the QT directive so that network gets included at compile time:

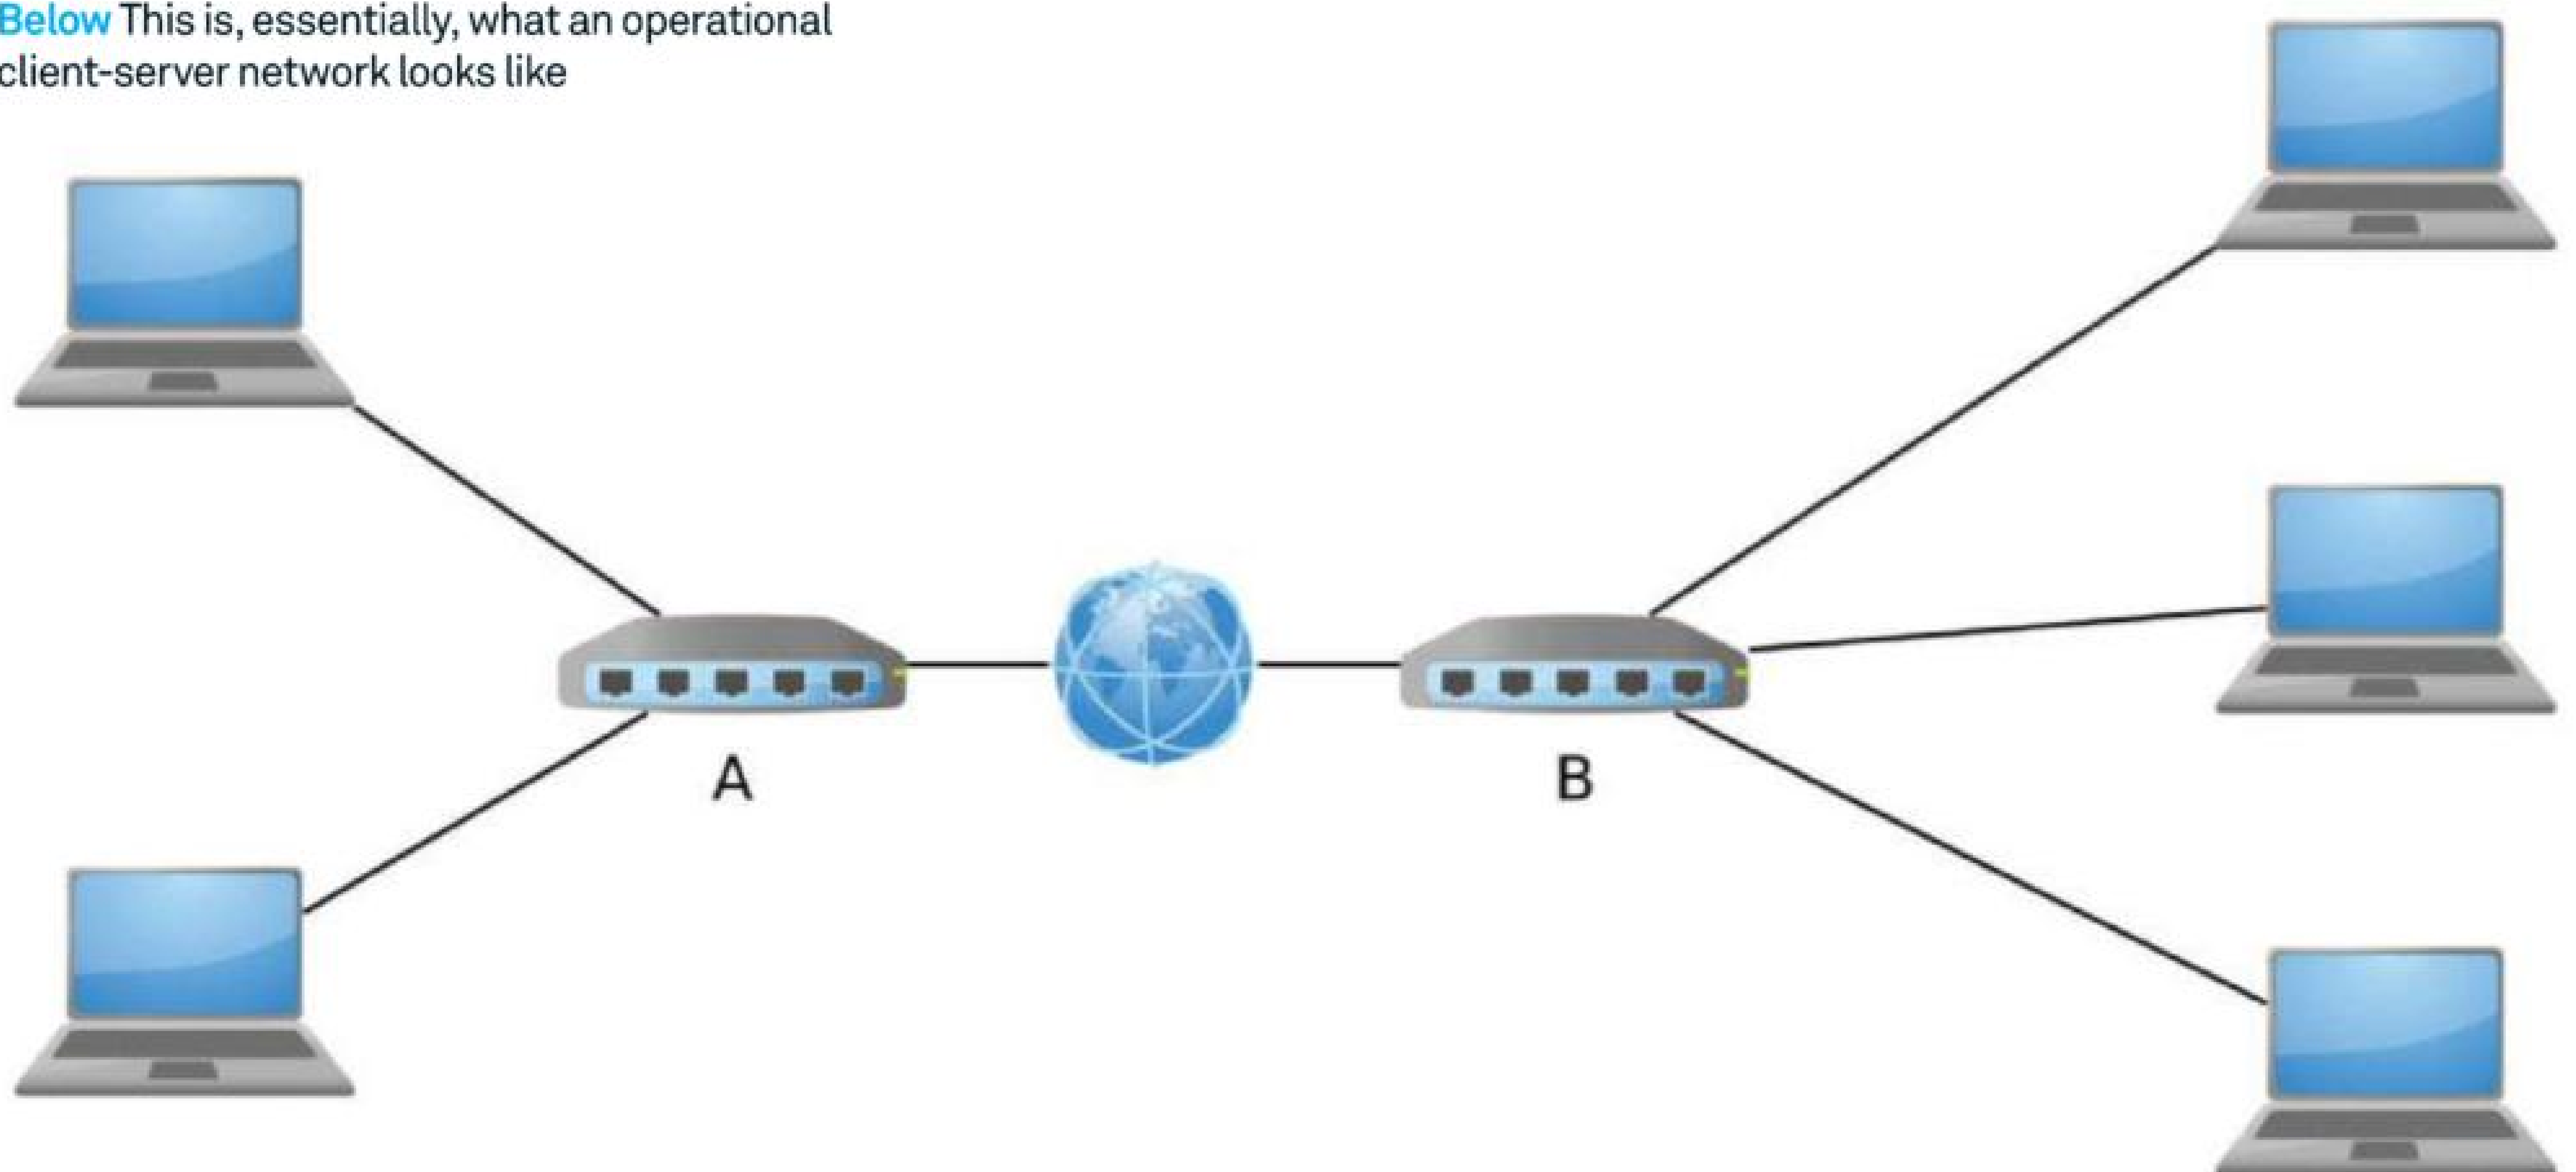
```
QT += network
```

With that, it is time to create the server object. Its declaration can be placed inside the main() function responsible for program bring-up:

```
myServer=new QTcpServer();
if(myServer->listen(QHostAddress::Any,1025)
)==false)
{
    qDebug() << "Server error" <<
myServer->serverError();
    return 0;
}
else
{
    qDebug() << "Server up" << myServer-
>serverAddress() << " " << myServer-
>serverPort();
}
```

In classic Berkeley socket programming, server and client logic are handled by instances of the socket class. Qt provides a dedicated server that contains helper methods for connectivity management. After start-up, we invoke the listen() method in order to make our QTcpServer wait for incoming connections.

Below This is, essentially, what an operational client-server network looks like



Passing in `QHostAddress::Any` makes the class listen on any network interfaces – you can constrain the code to a specific IP address by passing it in. Parameter number two is responsible for selecting the port that needs to be opened.

Accessing ports below the magic number 1024 requires root privileges on many Unix systems – simply running the program from the IDE would result in an access error. Avoid this problem by using port number 1025. Finally, the `exec()` method of `QCoreApplication` is invoked to start the main event loop.

Debugging and deploying network-related programs can be quite complex. For example, `ImagineTCPServer` will fail to start if another application is already connected to the port in question. However, this can be remedied by choosing a different port or by closing the offending application, which can be found via the `netstat` tool.

Signals and slots

Most network operations run asynchronously. Methods like `listen()` return immediately, delegating the actual work to a background thread. In the case of `QTcpServer`, incoming connections are announced by the emission of the `newConnection` signal.

Signals can only be handled by classes derived from `QObject`. Click `File>New File or Project` to open the creation wizard and select `C++ → C++ Class` in order to start the class generator. `ImagineListenerClass.h` looks like the following:

“ Signals can only be handled by classes derived from QObject ”

```
class ImagineListenerClass : public
QObject
{
    Q_OBJECT
public:
    explicit ImagineListenerClass(QObject
*parent = 0);
public slots:
    void connectionIncoming();
};
```

As discussed in previous issues of the magazine, slots are implemented via normal member functions. Add the following routine to `imaginelistenerclass.cpp`:

```
void ImagineListenerClass::connectionInc
oming()
{
    qDebug() << "Client connection
attempt";
}
```

Signal-Slot connections are established via the `connect` method. It requires pointers to the sending and the receiving objects, the actual signal and slot are to be specified via `SIGNAL()` and `SLOT()` macros:

```
else
{
    qDebug() << myServer-
>serverAddress() << " " << myServer-
>serverPort();
    myListener=new
ImagineListenerClass();
    QObject::connect(myServer,
SIGNAL(newConnection()), myListener,
SLOT(connectionIncoming()));
}
```

Client time

The work on the server is completed, but only for now. Start it up by clicking the Play button inside Qt Creator and then make sure that you keep your console window open. You can now check its functionality by opening a terminal window and entering `netstat -tlnp`, and the `ImagineTCPServer` will show up. Open a second session of Qt Creator in order to develop the client. Proceed to creating another Qt Console Application named `ImagineTCPClient` and modify its `.pro` file to enable networking as we have outlined above.



Qt coding series

Missed one of our Qt tutorials? Here's the full run of previous guides: Hello World (issue 133); Event-driven programming (134); Responsive development (135); Reading and writing files (136); Reading log files (137); Solving real-world problems with simulations (138); List Model widgets (145); Custom Model widgets (146).

You can order any of these back issues from the Imagine Shop – even better, all of them except issues 145 and 146 are currently on sale at half price: imagineshop.co.uk/magazines/linuxuser.html.

“ Multithreaded programs can show all kinds of odd behaviour ”

Streams of data

Real network applications exchange data, which is accomplished by using the data streams embedded into the socket instances. First, modify the constructor of `ImagineListenerClass` to provide the `myServer` instance to the method responsible for handling the `newConnection` signal. Then adjust its body as per **Fig. 01**.

`QTcpServer` can happily handle multiple pending connections at the same time. `NextPendingConnection` returns a socket representing the connection to the selected client. We'll wire its `disconnected` signal to the `deleteLater` slot to ensure unneeded sockets will be eliminated automatically. In the next step a byte array is populated and sent to the client.

Qt's underlying `DataStream` class gets improved as the framework evolves. Specifying a data format ensures that binaries built with later versions of Qt can interpret the data provided in the socket. Many applications use Qt 4.0's format. If your program is limited to Qt 5 and above due to other dependencies, you can also define this version as baseline.

The next step involves modifications to the client. Insert a class similar to the `ImagineListener` to handle the signals emitted from the socket class. Our application has to deal with three events – we need to notify the user of successful and failed connections and must receive data sent from the server.

Status messages are handled by emitting the corresponding statements to the console via `QDebug()` (see **Fig 02**).

Receiving data is more complex. Our simple protocol doesn't inform the client in advance about how much data is to be transmitted. Due to that, we simply read whatever amount of data has been received:

```
void ImagineReceiver::dataReady()
{
    QDataStream in(mySocket);
    in.setVersion(QDataStream::Qt_4_0);
    QString thatsIt;
    in >> thatsIt;

    qDebug() << thatsIt;
}

```

Each of the slots must be connected to the socket. This should be done before the

`connectToHost` method is invoked. On some operating systems the connected signal is emitted before `connectToHost` returns (**Fig. 03**).

Run the program in its current state to display the greeting message from the server – it will show up in the client's window. `ErrorOccurred` will be invoked to inform our program that the server terminated its connection.

Further interaction

Our server currently kicks out clients after providing them with a friendly greeting message. A polite server would stick around and wait for a response. Even though we could keep the sockets in our current example, this would lead to issues once more than one client is connected – we connect all signals to one slot which makes keeping the different connections apart difficult.

A thread is a convenient solution for this problem. Threads are subroutines that run in parallel to the main application code – spawning threads permits your code to do multiple things at the same time. Adding a thread is as easy as adding a new class whose header looks like this:

```
class ImagineThread:public QThread
{
    Q_OBJECT
public:
    ImagineThread(QTcpSocket* _aSocket);
    void run();
public:
    QTcpSocket* mySocket;
};

```

The actual payload can be found in the `run()` method which will be executed in the background after the thread was spawned:

```
void ImagineThread::run()
{
    exec(); //Start event loop
}
void ImagineThread::dataReady()
{
    QDataStream in(mySocket);
    in.setVersion(QDataStream::Qt_4_0);
    QString thatsIt;
    in >> thatsIt;
    qDebug() << thatsIt;
    exit();
}

```

«

Connecting to servers requires the use of a `QTcpSocket`. Its most basic implementation is made up of but two lines:

```
QTcpSocket *mySocket;
int main(int argc, char *argv[])
{
    QApplication a(argc, argv);
    mySocket=new QTcpSocket();
    mySocket->connectToHost(QHostAddress::
LocalHost, 1025);
    return a.exec();
}

```

`ConnectToHost()` is another asynchronous method which takes two parameters. The first one designates the IP to use, while the second specifies the port. Successful connections will be indicated by the emission of a signal which we will handle in the next step. For now, a running instance of `ImagineTCPServer` will display a connection attempted message when our client is started.

At the first glance, using TCP/IP to connect two processes on the same machine sounds like overkill. However, many programs are implemented in this fashion – handling inter-process communication via sockets provides a safe and simple way to de-couple components from one another.

ImagineThread's sole role involves staying until a signal is emitted. We accomplish that by invoking the `exec()` method which creates an event loop. It keeps idling until `exit` is called from the `dataReady` handler.

Our server must be modified in order to start the thread when a new client connects (Fig. 04).

A classic beginner's mistake involves calling the `run()` method of a thread directly. Background execution can be accomplished only when `start()` is called – it invokes the thread-spawning logic and then proceeds to executing the payload.

Finally, change the client so that it sends some data to its master:

```
void ImagineReceiver::dataReady()
{
    . . .

    QDataStream out(&block,
                   QIODevice::WriteOnly);
    out.setVersion(QDataStream::Qt_4_0);
    out << QString("Hello, server!");
    mySocket->write(block);
}
```

Multithreaded programs can show all kinds of odd behaviour due to a situation called race condition. If two routines access a shared element at the same time, havoc is guaranteed. Qt provides a variety of methods, such as mutexes, which can mitigate such problems.


Advanced considerations

In the beginnings of the internet, every computer had its own public IP address. Nowadays many, if not most, systems find themselves behind layers of routers providing network address translation and/or firewall services.

The image on page 47 shows a network made up of two local networks connected to one another via the Internet. If a networked application on system A wants to connect to system B, router B must be configured to expose the port in question. This can usually be accomplished in the backend of the router – since every manufacturer provides a slightly different user interface, it's best to refer to your router's (online) manual.

Conclusion

Well-versed developers can implement the most complex of protocols. Fortunately, Qt provides ready-made classes that implement commonly-used protocols such as HTTP. Next time, we will use one of these to create a small RSS reader.



```
void ImagineListenerClass::connectionIncoming()
{
    QTcpSocket *clientConnection = myServer->nextPendingConnection();
    connect(clientConnection, SIGNAL(disconnected()), clientConnection,
           SLOT(deleteLater()));

    QByteArray block;
    QDataStream out(&block, QIODevice::WriteOnly);
    out.setVersion(QDataStream::Qt_4_0);

    out << QString("Hello, client!");

    clientConnection->write(block);
    clientConnection->disconnectFromHost();
}
```

Fig 01

```
void ImagineReceiver::amConnected()
{
    qDebug() << "Connection successful";
}
void ImagineReceiver::errorOccurred(QAbstractSocket::SocketError anError)
{
    qDebug() << "Error: " << anError;
}
```

Fig 02

```
int main(int argc, char *argv[])
{
    QCoreApplication a(argc, argv);
    mySocket=new QTcpSocket();
    ImagineReceiver *myListener=new ImagineReceiver(mySocket);
    mySocket->connect(mySocket, SIGNAL(connected()),myListener,SLOT(amConnected()));
    mySocket->connect(mySocket, SIGNAL(readyRead()),myListener,SLOT(dataReady()));
    mySocket->connect(mySocket, SIGNAL(error(QAbstractSocket::SocketError)),
                    myListener,SLOT(errorOccurred(QAbstractSocket::SocketError)));
    mySocket->connectToHost(QHostAddress::LocalHost, 1025);
    return a.exec();
}
```

Fig 03

```
void ImagineListenerClass::connectionIncoming()
{
    QTcpSocket *clientConnection = myServer->nextPendingConnection();
    connect(clientConnection, SIGNAL(disconnected()), clientConnection,
           SLOT(deleteLater()));

    clientConnection->setParent(0);
    ImagineThread* myThread=new ImagineThread(clientConnection);
    connect(clientConnection,SIGNAL(readyRead()), myThread,SLOT(dataReady()));
    myThread->start();
    QByteArray block;
    QDataStream out(&block, QIODevice::WriteOnly);
    out.setVersion(QDataStream::Qt_4_0);
    out << QString("Hello, client!");
    clientConnection->write(block);
}
```

Fig 04

SUBSCRIBE TO
AND

SAVE 30%



PAY ONLY
£12.50
EVERY 3 ISSUES



ORDER HOTLINE | ONLINE AT

0844 249 0282

www.imaginesubs.co.uk/lud

EVERY ISSUE PACKED WITH...

- Programming guides from **Linux experts**
- Breaking news from **enterprise vendors**
- Creative **Raspberry Pi** projects and advice
- Best Linux **distros** and **hardware** reviewed

WHY YOU SHOULD SUBSCRIBE...

- Save 30% off the cover price – just £12.50 every 3 issues on Direct Debit
- FREE delivery direct to your door
- Never miss an issue

ORDER BY 30TH APRIL

Use code **PQ15** for this offer.

BY POST

Send your completed form to:
Linux User & Developer Subscriptions, 800 Guillat Avenue,
Kent Science Park, Sittingbourne, Kent ME9 8GU

YOUR DETAILS


Title _____ First name _____
Surname _____
Address _____

Postcode _____ Country _____
Telephone number _____
Mobile number _____
Email address _____

DIRECT DEBIT PAYMENT

UK Direct Debit Payment

I will pay just £12.50 every 3 issues

		Instruction to your Bank or Building Society to pay by Direct Debit		
<small>IMAGINE PUBLISHING Please fill in the form and send it to: Imagine Publishing Limited, 800 Guillat Avenue, Kent Science Park, Sittingbourne, Kent, ME9 8GU</small>				
<small>Name and full postal address of your Bank or Building Society</small>		<small>Originator's Identification Number</small>		
To: The Manager _____ Bank/Building Society		5 0 1 8 8 4		
Address _____		<small>Reference Number</small>		

_____ Postcode _____		<small>Instructions to your Bank or Building Society</small>		
Name(s) of account holder(s) _____		<small>Please pay Imagine Publishing Limited Direct Debits from the account detailed in this instruction subject to the safeguards assured by the Direct Debit guarantee. I understand that this instruction may remain with Imagine Publishing Limited and, if so, details will be passed on electronically to my Bank/Building Society</small>		
Branch sort code _____		Signature(s) _____		
Bank/Building Society account number _____		Date _____		
<small>Banks and Building Societies may not accept Direct Debit instructions for some types of account</small>				
<small>A6 instruction form</small>				

PAYMENT DETAILS

YOUR EXCLUSIVE READER PRICE, 1 YEAR (13 ISSUES)

UK £62.40 (Save 20%) **Europe – £70** **World – £80** **USA – £80**

Cheque

I enclose a cheque for £ _____
(made payable to Imagine Publishing Ltd)

Credit/Debit Card

Visa Mastercard Amex Maestro

Card number _____ Expiry date _____
| | | | | | | | | | | | | | | | | | | | | |

Issue number (if Maestro)

Signed _____

Date _____

Please tick if you do not wish to receive any promotional material from Imagine Publishing Ltd by post
by telephone via email

Please tick if you do not wish to receive any promotional material from other companies by post
by telephone Please tick if you DO wish to receive such information via email

TERMS & CONDITIONS

Terms and Conditions: This offer entitles new UK direct debit subscribers to pay only £12.50 every 3 issues. New subscriptions will start from the next available issue. Offer code PQ15 must be quoted to receive this special subscription price. Details of the Direct Debit guarantee are available on request. This offer expires 30 April 2015. Imagine Publishing reserves the right to limit this type of offer to one per household. Subscribers can cancel this subscription at any time.

Storemags

PRACTICAL Raspberry Pi


68 “As long as there’s a constant wireless signal, you can see and direct the Rapiro as you wish”



Contents



60 Controlling a robot arm



62 Draw circuits with BC paint



64 Secure your Raspberry Pi



68 Add sight and control to Rapiro



Raspberry Pi Car Computer

Make your own touchscreen navigation system that gives directions, local weather reports and plays music



Cars are getting clever. These days, with smart navigation interfaces built into new cars, you don't need to go out and buy yourself a TomTom to get help with directions. But if you've got a Raspberry Pi then you don't even need to buy that – let alone a new car!

In this project we will show you how to build your own car computer with your Pi, a quality touchscreen like the 9-inch model from SainSmart that we're using here, and a few other bits like a GPS module and USB 3G modem. Your CarPi will be able to use open source navigation software Navit to show your route map on screen, plus speech synthesis to read out directions, and it will also be able to check your location and give you weather reports. It'll work as a music player too, of course.

It's an ambitious project, but you will gain a solid understanding of custom-made interfaces, navigation software and geolocation data, touchscreen calibration, speech synthesis and more. While you don't have to use the same SainSmart screen as us, we do recommend it for this project as it is one of the few large touchscreens out there for the Pi. There are more improvements at the end too, so check the components list, make sure you've got everything and let's get started!





Above We're using Adafruit's excellent GPS Breakout kit here: bit.ly/1G8X2gw

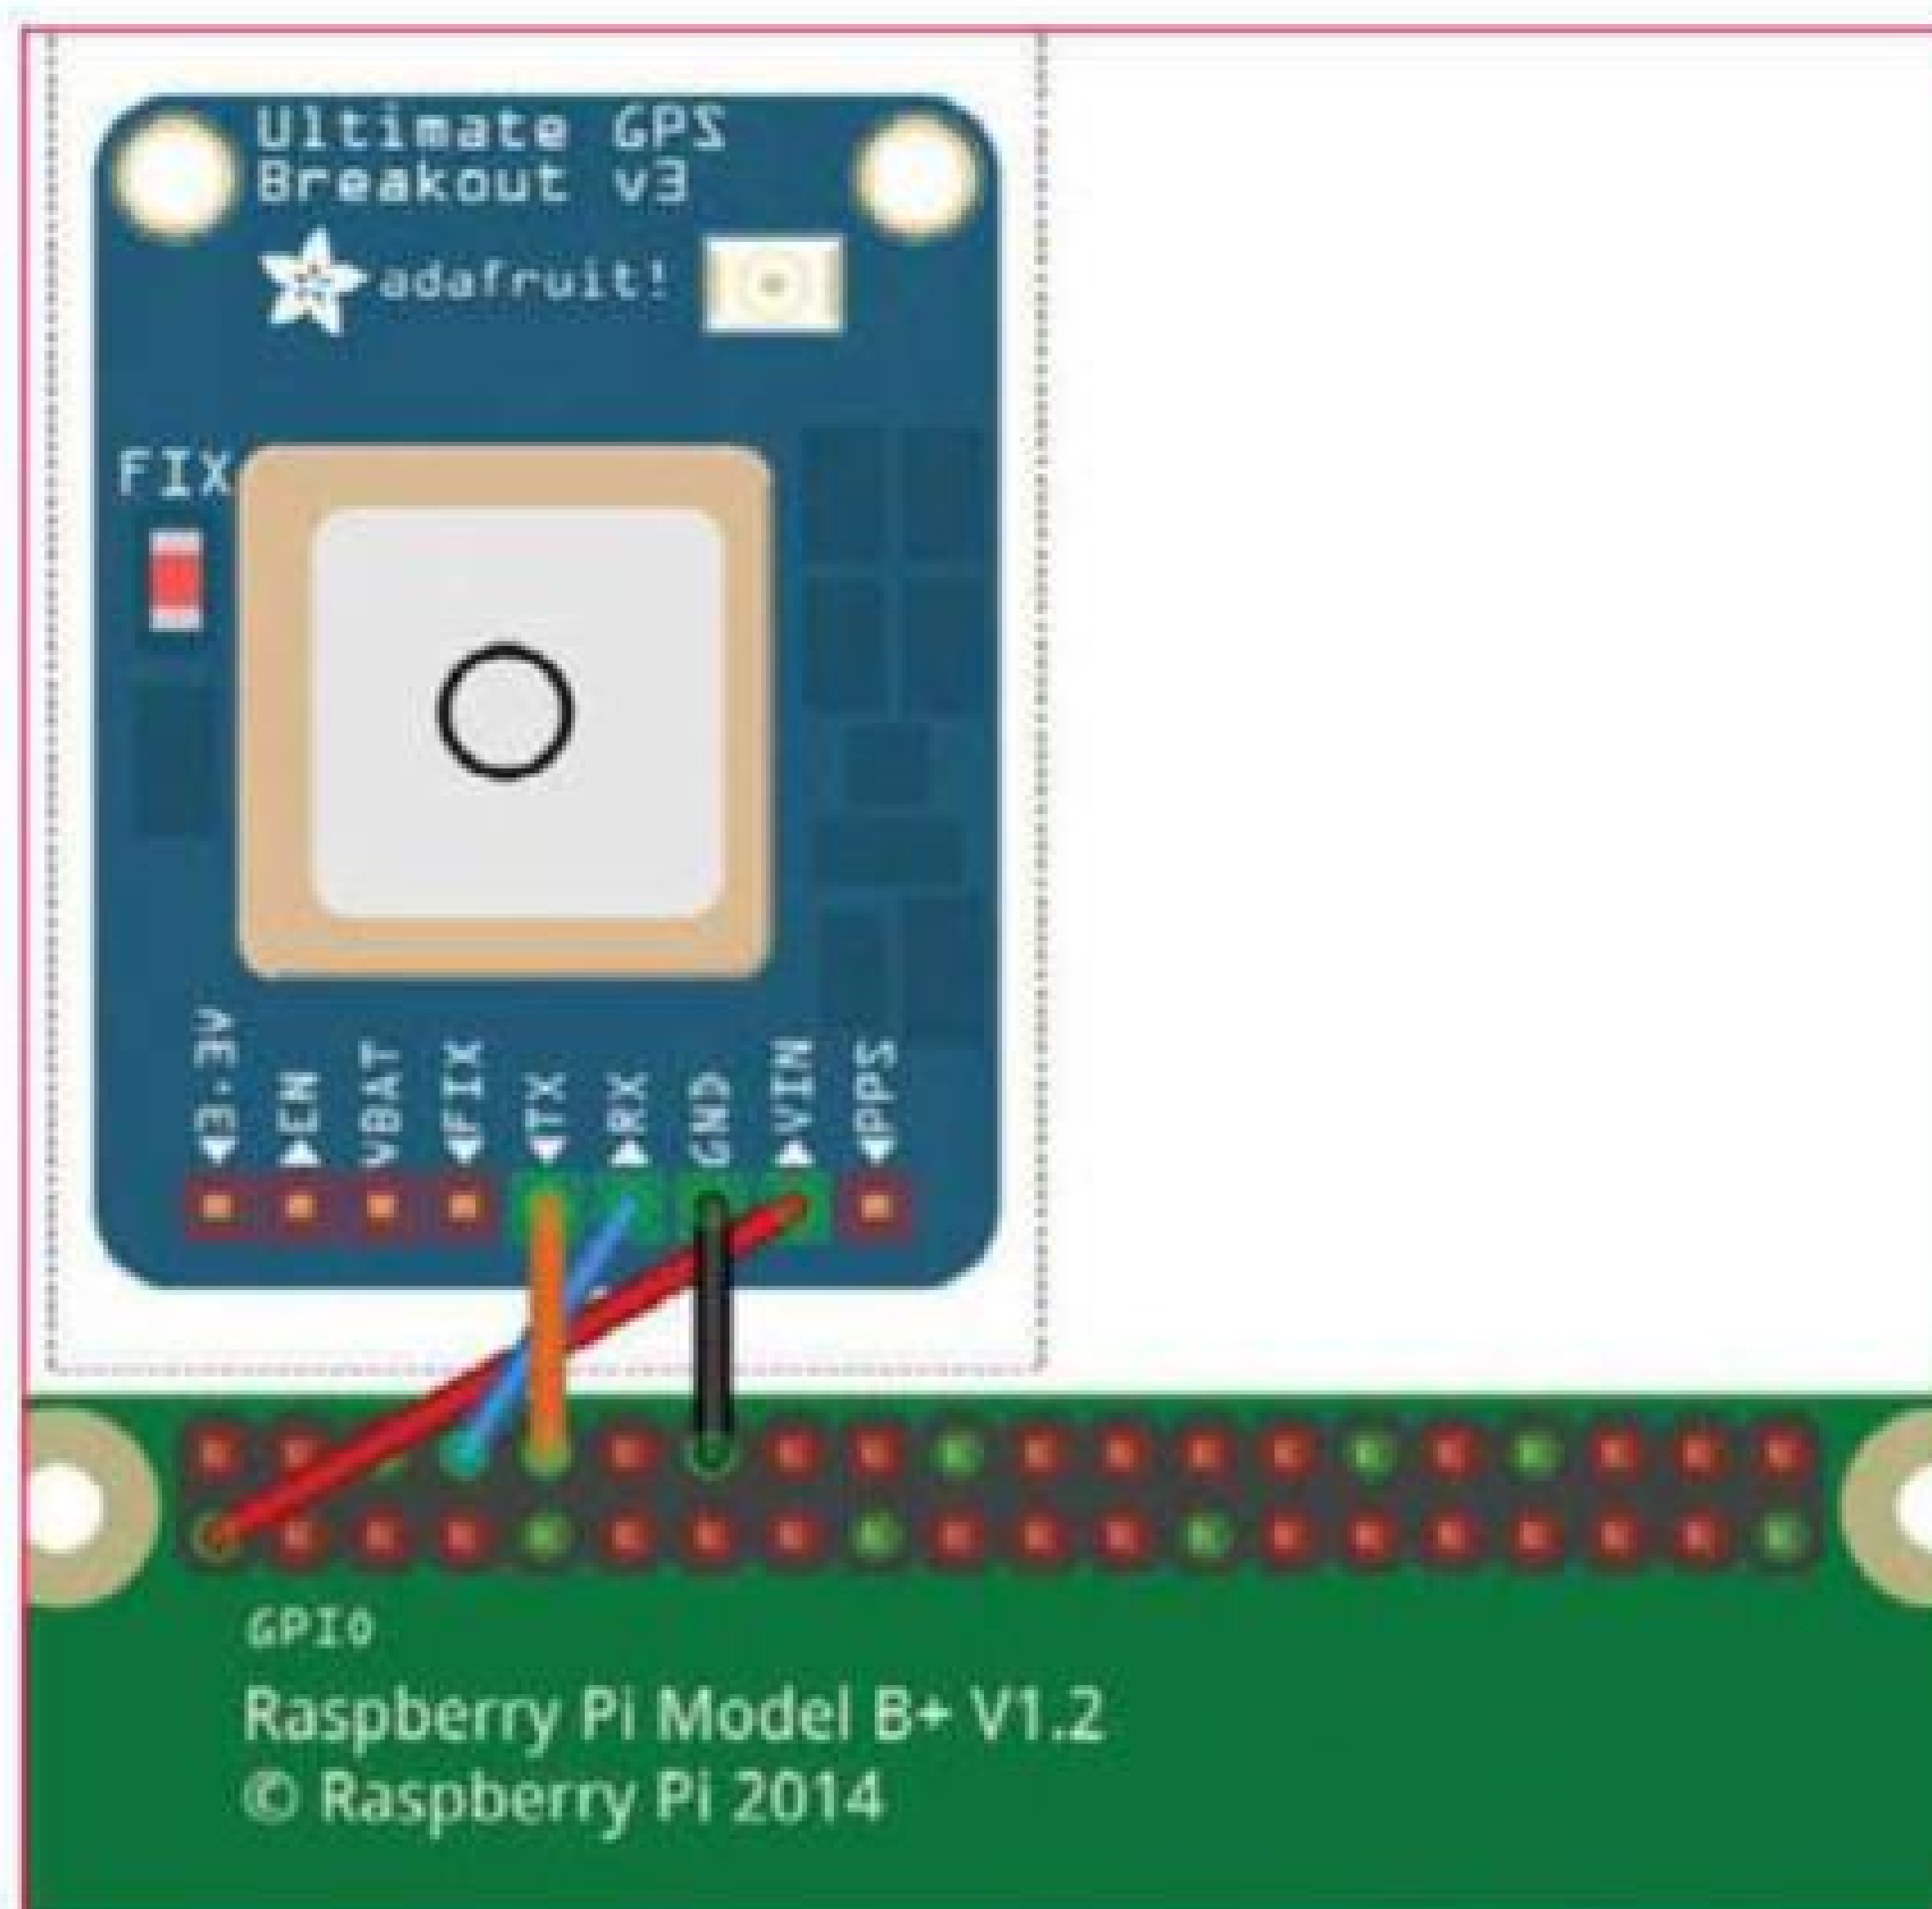


Liam Fraser is the creator of the Raspberry Pi Tutorials YouTube series and volunteers as a Linux server administrator for the Raspberry Pi Foundation. He also studies Computer Science at the University of York.

01 Basic configuration

Boot up your Raspberry Pi and expand the filesystem using `raspi-config`. Go to Advanced Options and disable the Serial connection – you'll need this to talk to the GPS module later. In `raspi-config`, enable X at boot as the pi user. Say Yes to reboot. Once rebooted, ensure your packages are up to date with:

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



02 Connect GPS module

Solder the pin headers onto the Adafruit GPS module. You can also solder the battery connector which is used to keep the device partially active, giving a faster fix. You only need to use 4 pins: 3.3V, ground, serial transmit and serial receive. Power the Pi off again before connecting anything.

As we are using GPS, the antenna will have to go outside or under a window to gain signal. Connect the antenna to the board and power everything back on. The light on the GPS module will flash frequently while finding a fix. Once it has one, it will blink every 15 seconds.

03 Install navigation software

Begin to install the Navit navigation software by entering:

```
sudo apt-get install navit gpsd gpsd-clients espeak
sudo nano /etc/default/gpsd
set START_DAEMON="true"
```

...and set:

```
DEVICES="/dev/ttyAMA0"
```

Start the GPS daemon with:

```
sudo /etc/init.d/gpsd start
```

You can check it's working by looking at the GPS data with:

```
cgps -s
```

04 Connect the screen

The SainSmart screen doesn't come with any written instructions. Instead there is a YouTube video on their website with details about how to put it together: bit.ly/1DF6eJJ. The important part is that the DC power supply should be 12V.

05 Set the screen resolution

We will have to force the correct resolution (1024x600) for the screen by editing `/boot/config.txt` with `sudo`. To do so, add the following options:

```
framebuffer_width=1024
framebuffer_height=600
hdmiforce_hotplug=1
hdmicvt=1024 600 60 3 0 0 0
hdmigroup=2
hdmimode=87
```

For the changes to properly take effect you will need to reboot with `sudo reboot`.

06 Download kernel source

To start the touchscreen, you need to compile an extra kernel module to support it. The program `rpi-source` (github.com/notro/rpi-source/wiki) will find the source of your kernel. Install `rpi-source` with:

```
sudo wget https://raw.githubusercontent.com/notro/rpi-source/master/rpi-source -O /usr/bin/rpi-source
&& sudo chmod +x /usr/bin/rpi-source && /usr/bin/rpi-source -q -tag-update
```

Then run `rpi-source` to get the source of the running kernel.

07 Update GCC

Recent Raspberry Pi kernels are compiled with GCC 4.8. Raspbian only comes with 4.6 so you will have to install 4.8 to continue with the following steps. Do this by entering:

```
sudo apt-get install -y gcc-4.8  
g++-4.8 ncurses-dev
```

Then you have to set GCC 4.8 as the default:

```
sudo update-alternatives  
--install /usr/bin/gcc gcc /usr/  
bin/gcc-4.6 20  
sudo update-alternatives  
--install /usr/bin/gcc gcc /usr/  
bin/gcc-4.8 50  
sudo update-alternatives  
--install /usr/bin/g++ g++ /usr/  
bin/g++-4.6 20  
sudo update-alternatives  
--install /usr/bin/g++ g++ /usr/  
bin/g++-4.8 50
```

08 Pick the module to compile

Rpi-source puts the kernel source in a folder called 'linux'. To choose the USB Touchscreen Driver, enter the following:

```
cd linux  
make menuconfig  
Device Drivers -> Input device  
support -> Generic input layer  
(needed for keyboard, mouse,  
...) -> Touchscreens (press space  
to include) -> USB Touchscreen  
Driver (press M to make module)
```

Once you've done that, you then need to make sure you save your changes as '.config' and run scripts/diffconfig to see the differences.

09 Compile and install the module

Now you need to compile and install the module. Do so by entering:

```
make prepare  
make SUBDIRS=drivers/input/  
touchscreen modules  
sudo make SUBDIRS=drivers/input/  
touchscreen modules_install  
sudo depmod
```

If you unplug and reconnect the touchscreen, it should work fine but it will probably need calibrating.


Full code listing

```
#!/usr/bin/env python2  
  
import os, sys, requests, pygame  
from gps import *  
from pygame.locals import *  
  
class WeatherClient:  
    apikey = "7232a1f6857090f33b9d1c7a74721"  
  
    @staticmethod  
    def latlon():  
        gpsd = gps(mode=WATCH_ENABLE)  
  
        # Needs better error handling  
        try:  
            while True:  
                report = gpsd.next()  
                if report['class'] == 'TPV':  
                    gpsd.close()  
                    return report['lat'], report['lon']  
  
        except:  
            return None, None  
  
    @staticmethod  
    def usefuldata(j):  
        # Returns a string of useful weather data from a LOT of json  
        d = j['data']['current_condition'][0]  
        out = "Now - Temp: {0}C, Feels Like: {1}C, Description: {2}\n"\  
            .format(d['temp_C'],  
                  d['FeelsLikeC'],  
                  d['weatherDesc'][0]['value'])  
  
        hourly = j['data']['weather'][0]['hourly']  
        hour_count = 1  
        for h in hourly:  
            out += (" +{0}hr - Temp: {1}C, Feels Like: {2}C, Chance of Rain:"  
                  "{3}%, Description: {4}\n")\  
                .format(hour_count,  
                        h['tempC'],  
                        h['FeelsLikeC'],  
                        h['chanceofrain'],  
                        h['weatherDesc'][0]['value'])  
  
        hour_count += 1  
  
        # Rstrip removes trailing newline  
        return out.rstrip()  
  
    @staticmethod  
    def update():  
        errstr = "Error getting weather data"  
  
        lat, lon = WeatherClient.latlon()  
        if lat == None or lon == None:  
            return errstr  
  
        api_req = ("http://api.worldweatheronline.com/free/v2/weather.ashx"  
                  "?q={0}%2C{1}&format=json&key={2}").format(lat, lon,  
                                                             WeatherClient.apikey)  
  
        r = None
```


Tutorial

SainSmart's 9-inch HDMI/VGA touchscreen (bit.ly/1Ciu4H9) has a fantastic display and is perfect for all sorts of Pi projects

Adafruit's Ultimate GPS Breakout kit provides Navit and the weather function with the location data that they require



The screen control panel that comes with the SainSmart screen enables you to easily change the display settings (i.e. brightness, contrast, etc) as well as the input (i.e. HDMI, VGA, AV1, etc)

As well as the main controller board, the touch screen is connected to a four-line USB controller which then plugs into the Pi's USB port

Embed the screen

We've looked at the PiTFT and the HDMIPi before, but the SainSmart touchscreen we're using here is uniquely suited to many embedded projects. It's larger than the PiTFT but also without the large bezels of the HDMIPi – and it's incredibly thin – so it's the kind of thing that is really useful for installation projects, whether that's something as simple as a photo slideshow in a real picture frame or a home automation control interface embedded into a cupboard door.

10 Calibrate the touchscreen

At this point, you can easily calibrate the touchscreen by entering the following:

```
cd /etc/X11
sudo mkdir xorg.conf.d
cd xorg.conf.d
sudo nano 99-calibration.conf
```

...with the following content:

```
Section "InputClass"
    Identifier "calibration"
    MatchProduct "eGalax Inc. USB TouchController"
    Option "SwapAxes" "1"
    Option "InvertX" "1"
EndSection
```

Invert X actually inverts Y because the axes have been swapped around. Reboot again for these changes to occur. Now the calibration is roughly correct, download an input calibrator that Adafruit have packaged already.

```
wget http://adafruit-download.s3.amazonaws.com/
xinput-calibrator_0.7.5-1_armhf.deb
sudo dpkg -i xinput-calibrator_0.7.5-1_armhf.deb
DISPLAY=:0.0 xinput_calibrator
```

DISPLAY=:0.0 is useful because you can run the program from any terminal (including an SSH session) and have it appear on the touchscreen. Touch the points on the screen as prompted. Once the program is finished, you should get an output that is similar to the following:

```
Option "Calibration" "84 1957 270 1830"
```

Add it to the '99-calibration.conf' file that we created earlier just below the other Option entries.

11 Download maps

Navit needs maps; download them from maps.navit-project.org. You can either use the web browser on the Pi or download the map from another machine and copy it using `scp`. Use the predefined area option to select where you live. The smaller the area that you pick, the less data you will have to process. Here the UK has a map size of 608 MB. Now move the map to the navit folder:

```
mkdir -p /home/pi/.navit/maps
mv /home/pi/Downloads/$your_map /home/pi/.
navit/$country.bin
```

For example:

```
mv /home/pi/Downloads/osm_bbox_-9.7,49.6,2.2,61.2.bin
/home/pi/.navit/maps/UK.bin
```

12 Navit configuration

Sudo-edit `/etc/navit/navit.xml` with your favourite editor. Search for `openstreetmaps`. Now disable the sample map above, enable the `openstreetmap` mapset and set the data variable to where you just moved your map. In this case it looks like this:

```
<!-- Mapset template for openstreetmaps -->
<mapset enabled="yes">
<map type="binfile" enabled="yes" data="/home/
pi/.navit/maps/UK.bin"/>
</mapset>
```

Then search for `osd` entries similar to:

```
<osd enabled="yes" type="compass"/>
```

...and enable the ones you want – we recommend enabling them all. You may want to zoom in closer than the default map layout. A zoom value of 64 is useful.

13 Sound configuration

Before configuring speech support for Navit, configure the external sound card. You have to stop the Broadcom module from loading and remove some Raspberry Pi-specific ALSA (Advanced Linux Sound Architecture). To do this, sudo-edit /etc/modprobe and comment out (i.e. prefix with a #):

```
■ snd-bcm2835
```

Then run:

```
■ sudo rm /etc/modprobe.d/alsa*
```

Reboot for the changes to take effect. Use alsamixer to set the volume on the if it's too quiet.

14 Download a voice

The speech synthesis software needs a voice and a proprietary binary. You can get both by completing the following steps:

```
■ sudo mkdir -p /usr/share/mbrola/voices/
```

```
■ wget http://www.tcts.fpms.ac.be/synthesis/mbrola/dba/en1/en1-980910.zip
```

```
■ unzip en1-980910.zip
```

```
■ sudo cp en1/en1 /usr/share/mbrola/voices
```

```
■ wget http://www.tcts.fpms.ac.be/synthesis/mbrola/bin/raspberri_pi/mbrola.tgz
```

```
■ tar zxvf mbrola.tgz
```

```
■ sudo mv mbrola /usr/local/bin/
```

15 Create speech script

Navit supports speech by running an external script and passing the text to speak as an argument. Create one using:

```
■ cd /home/pi/.navit
```

```
■ wget http://liamfraser.co.uk/lud/carpi/chime.wav
```

```
■ touch speech.sh
```

```
■ chmod +x speech.sh
```

Now edit speech.sh:

```
■ #!/bin/bash
```

```
■ aplay -r 44100 /home/pi/.navit/chime.wav
```

```
■ espeak -vmb-en1 -s 110 -a 150 -p 50 "$1"
```

Finally, test it with:

```
■ ./speech.sh "Hello World"
```

Full code listing

```
try:
    r = requests.get(api_req)
except requests.exceptions.RequestException as e:
    return errstr

return WeatherClient.usefuldata(r.json())

class CarLauncher:
    def __init__(self):
        pygame.init()
        pygame.mixer.quit() # Don't need sound
        screen_info = pygame.display.Info()
        self.screen = pygame.display.set_mode((screen_info.current_w,
                                                screen_info.current_h))

        pygame.display.set_caption('Car Launcher')
        self.titlefont = pygame.font.Font(None, 100)
        self.wfont = pygame.font.Font(None, 30)
        self.w_text = None # Weather text

    def clean_background(self):
        background = pygame.Surface(self.screen.get_size())
        self.background = background.convert()
        self.background.fill((0, 0, 0))

        # Render title centered
        text = self.titlefont.render("CarPi Launcher", 1, (255, 255, 255))
        textpos = text.get_rect()
        textpos.centerx = self.background.get_rect().centerx
        self.background.blit(text, textpos)

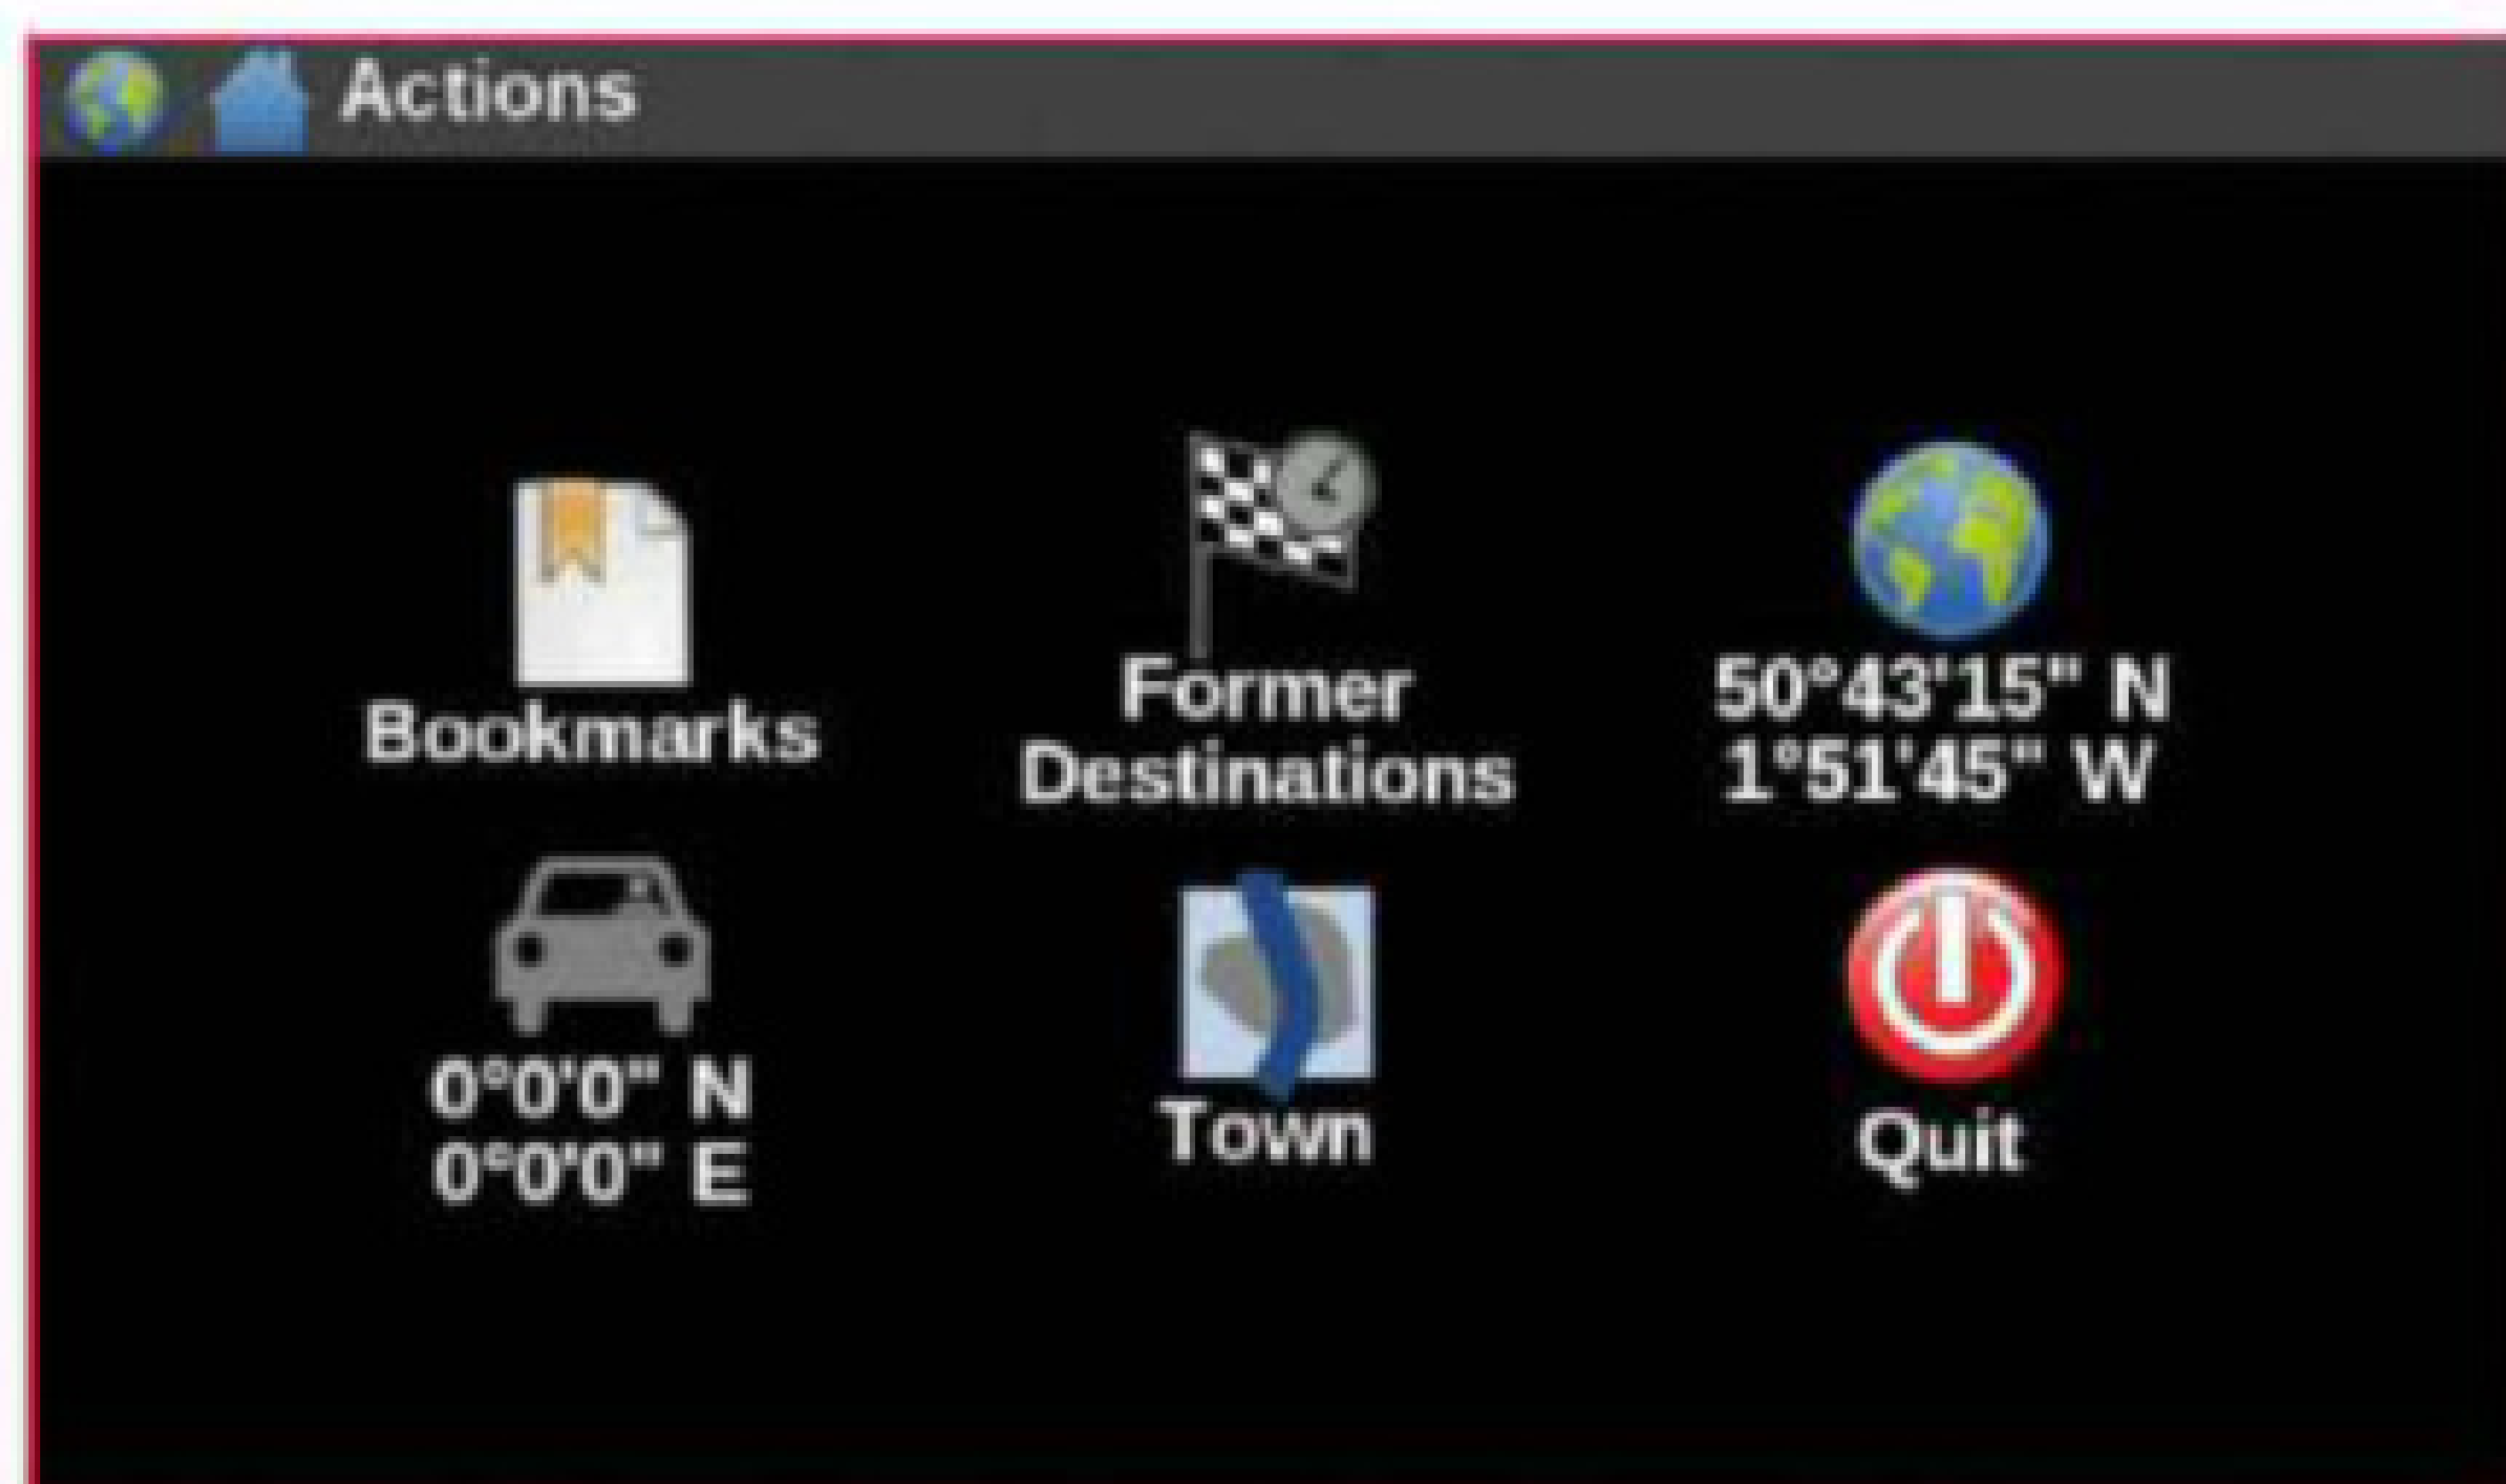
        self.screen.blit(self.background, (0,0))
        pygame.display.flip()

    def main_menu(self):
        # btns maps Text -> Rectangles we can do collision detection on
        self.btns = {'Music' : None, 'NAV' : None, 'Weather' : None}

        item_num = 1
        for key in self.btns:
            text = self.titlefont.render(key, 1, (255,255,255))
            textpos = text.get_rect()
            max_width = self.background.get_rect().width / len(self.btns)
            center_offset = max_width * 0.5
            # This y pos puts buttons just below title
            textpos.centery = self.background.get_rect().centery / 2
            textpos.centerx = (max_width * item_num) - center_offset
            self.btns[key] = textpos
            self.screen.blit(text, textpos)
            item_num += 1

        pygame.display.flip()

    def select_rect(self, rect, text):
        # Colour a rect the user has clicked in green
        surface = pygame.Surface((rect.w, rect.h))
        surface.fill((0, 255, 0))
        # Now we have to draw the text over it again
        t = self.titlefont.render(text, 1, (255,255,255))
        surface.blit(t, (0,0))
        self.screen.blit(surface, rect)
        pygame.display.flip()
```

Above The Navit software comes with a host of options built into its menu hierarchy



Above The pypmptouchgui front-end for the music player is surprisingly featureful

You will need to write your own launcher for CarPi



Make it mobile

It is definitely best to put this project together in a clean workspace so that you can clearly see what you're working with and ensure everything is correctly wired and soldered, but the point of the project is to make this setup portable so that you can put it in your car and use it on the road. You could install everything into a single, hand-made enclosure or customise a large bought one, or you could secure the various parts inside, for example, your glovebox or car doors. You'll also need to power both the screen and your Pi with a power pack and ensure that the GPS antenna is fastened into a good spot for signal.

16 Configure Navit for speech

The last part is simple. Edit the Navit config file again (`/etc/navit/navit.xml`) and replace the following line:

```
<speech type="cmdline" data="echo 'Fix the speech tag in navit.xml to let navit say:' '%s'" cps="15"/>
```

...with:

```
<speech type="cmdline" data="/home/pi/.navit/speech.sh %s" cps="10" />
```

Now you can run Navit with `DISPLAY=:0.0 navit` and have fun experimenting.

17 Install the music player

MPD is the music player back-end and pypmptouchgui is the front-end that needs installing manually:

```
sudo apt-get install mpd ncmpcpp
wget http://www.spida.net/projects/software/pypmptouchgui/pypmptouchgui-0.320.tgz
tar zxvf pypmptouchgui-0.320.tgz
cd pypmptouchgui-0.320/
sudo python setup.py install
# Fix hard coded path in software
sudo ln -s /usr/local/share/pypmptouchgui/ /usr/share/pypmptouchgui
```

18 Copy music

Scp (secure copy protocol) was used here to copy music. First get the Pi's IP address by running `ip addr`. Then

run `sudo passwd` to set a password for root. From a computer with music on, run:

```
scp -r music_folder root@pi_ip_address:/var/lib/mpd/music/
```

Then on the Pi, change the ownership of the music that you just copied:

```
sudo chown -R mpd:audio /var/lib/mpd/music
```

19 Update mpd music library

Ncmpcpp is a command line client for mpd. Type `ncmpcpp` and press `U` to update the library. Press `3` to browse the library and check the music is there, and press `Q` to quit. Pressing `1` will select the help screen if you want to do more.

20 Install awesome window manager

Now you will need to write your own launcher for CarPi, which will run full-screen. To ensure every application is forced to full-screen, use awesome window manager in full-screen mode.

```
sudo apt-get install awesome
sudo rm /etc/alternatives/x-session-manager
sudo ln -s /usr/bin/awesome /etc/alternatives/x-session-manager
```

When changing the default x-session-manager, awesome will be auto-started at boot instead of LXDE. If you reboot the Pi, awesome should then load up automatically.


21 Install the requirements for your launcher

The launcher is going to use a weather API combined with data location data from the GPS receiver to give weather updates when requested. The nicest HTTP API for Python is requests, which you can install by doing the following:

```
sudo apt-get install python-pip
sudo pip install requests
```

22 Write the launcher code

Creating the code itself is pretty self explanatory, but you can use our ready-made version by downloading the CarPi package from FileSilo.co.uk and extracting carlauncher/carlauncher.py.




23 Start the launcher automatically

Sudo-edit /etc/xdg/awesome/rc.lua and move awful.layout.suit.max.fullscreen to the top of the layouts list. Add the following to the bottom of the file:

```
awful.util.spawn_with_shell("/home/pi/carlauncher/carlauncher.py")
```

Now reboot again and the launcher should come up automatically.



24 Future improvements

There are a number of improvements that could be made to the base project at this point:

- Make the launcher switch between applications rather than start them again each time
- Make the launcher look better aesthetically with icons
- Use Mopidy instead of MPD so you can use Spotify
- Further Navit configuration to make it more featureful
- An SSD or USB flash drive for storage to make things quicker

Full code listing

```
def reset(self):
    self.clean_background()
    self.main_menu()
    self.render_weather()

def execute(self, path):
    os.system(path)
    # os.system blocks so by the time we get here application
    # has finished
    self.reset()

def render_weather(self):
    if self.w_text == None:
        return

    # Get y starting at the bottom of the nav button
    margin = 10
    y = self.btns['NAV'].bottomleft[1] + margin

    for t in self.w_text.split("\n"):
        line = self.wfont.render(t.rstrip(), 1, (255,255,255))
        line_rect = line.get_rect()
        line_rect.centerx = self.background.get_rect().centerx
        line_rect.y = y
        self.screen.blit(line, line_rect)
        y += margin + line_rect.height

pygame.display.flip()

def handle_events(self, events):
    for e in events:
        if e.type == QUIT:
            sys.exit()
        elif e.type == MOUSEBUTTONDOWN:
            pos = pygame.mouse.get_pos()
            # Check if it collides with any of the buttons
            for btn_text, rect in self.btns.iteritems():
                if rect.collidepoint(pos):
                    self.select_rect(rect, btn_text)
                    if btn_text == "NAV":
                        self.execute("/usr/bin/navit")
                    elif btn_text == "Music":
                        self.execute("/usr/local/bin/pympdtouchgui")
                    elif btn_text == "Weather":
                        self.w_text = WeatherClient.update()
                        # Reset will render weather if string is populated
                        self.reset()

def loop(self):
    clock = pygame.time.Clock()
    self.reset()

    while 1:
        self.handle_events(pygame.event.get())
        # 5 fps is plenty
        clock.tick(5)

if __name__ == "__main__":
    cl = CarLauncher()
    cl.loop()
```


MPU-6050 Containing a MEMS accelerometer and a MEMS gyroscope, this sensor reads the x, y and z axis channels with 16-bit ADC conversion

Robot Arm Available from Maplin Electronics and OWI Robotics, the arm comes with software for control even before you get the MPU-6050 involved



Veroboard Veroboard is great to tidy up wires in projects like this, where they get in the way but can't really be run through a breadboard

Components list

- Raspberry Pi Model B
- Maplin Robotic Arm Kit With USB PC Interface
- MPU-6050 Six-Axis Gyro and Accelerometer
- 3 Mini Push Button Switches
- Veroboard
- Velcro strap
- 1m Ribbon Cable

Left This robotic arm is one of the most used ones and there are tonnes of guides for it

Below One of these buttons controls the light on the end of the robotic arm, while the other two open and close its gripper





Robot Arm

Joseph Thomas gets to grips with natural motion control for a Raspberry Pi-powered robot arm

What first inspired you to begin your robot arm project?

The robot arm itself was one I'd seen years ago and I really wanted it because it's something that you can control yourself – it really captured my young imagination. I was volunteering at a science museum down here in Harlow and this club based around the Raspberry Pi sprung up, and I bought the robot arm because I wanted it. So then I had the Raspberry Pi thing going on at the same time and thought, why not meld the two?

I had this complicated system of key presses to get it to do anything, which was a bit boring, and then James Dali (one of the people who helps out with the club) gave me the idea of shoving an accelerometer on the top of it to give an idea of where it is. I took that and thought, 'What if I had the accelerometer on me and sort of used it to mirror the motion of my hand?' So I looked around, searched up the accelerometer he was using (the MPU-6050) and then found it for about £5 on eBay – it's normally about £30 from SparkFun but I'm on a student budget... A lot of the code I've used is borrowed but open source, and people have said it's fine, so then I went through and had two programs – one that could control the arm, one that took the input in from the accelerometer – and kind of just smashed them together. It's not that nice to look at, but it works and that's all that really matters.

So what exactly are you reading with that MPU-6050?

There's the gyroscope and the accelerometer in the code I'd found – you can use one or the other, but the gyroscope is very good for degrees over time and it tends to drift, while the accelerometer is good for sudden turns and for measuring gravity. If you compare the two to each other then you can get a rough angle all of the time, so it's essentially the accelerometer and the gyroscope used together to correct the faults with one or the other. It's got two axes of motion – pitch and roll.

Take us through the code itself.

So in the first bit it finds where the actual I2C interface is and there's a quick setup – I've got three buttons on there to control the gripper and the lights, so it sets those up – and then there's a bit which is using the USB library to find the robot arm, then spitting it out if that's an issue. There are a couple of definitions for some functions to actually move the arm, so it's a little bit easier – each motor direction is a different binary number – and then there are more definitions for setting up reading data from the accelerometer and a bit of maths for making sure the gyro and the accelerometer are both giving the correct angle. Then there's this while loop with a try inside it that is just pulling the accelerometer for data, spitting out the maths stuff, before just checking that the angle given is within a certain range. If it is, move this motor left (for example), or if a button is pressed then it turns a light on. The only problem I've had with it is that to actually move it, it requires a change in angle – so there's not a continuous thing. I have to wave my hand a little bit, but there's that degree angle and if I trip it then it'll move around.

Have you considered adding any more forms of control?

Yeah, I've done a lot of research into this. In terms of other ways to control it, I quite like the intuitiveness of it – to rotate and move this arm you are moving your own arm, so that's something I've been focussing on and trying to get even more intuitive. Trying to get some sort of – I bought an Arduino at some point – trying to build an actual robotic hand and then spreading out from there. Eventually, my big plan – many, many years in the future – is to have an entire sort of human body that is controlled by the movements of the

user, but that's a very large plan which I haven't put too much into just yet! But essentially, the prototype that people have done before is sort of having pot sensors – potentiometers – on the fingers just to measure the actual rotation and closing of the fist, then having that represented with servos and then possibly doing that with actual pieces of string to sort of emulate the tendons. So you'd have a single servo, or a couple of servos, in an arm bit that would pull string which would close each finger in turn.

Another idea, which seems to be one of the most viable, is having it completely brain controlled... There's a fair amount of interest in reading brain activity – you can do it with the NeuroSky, for example. There's quite a nice open source project which I might end up using because it has four inputs, so you can measure at least two things at once and that seems to be a fairly interesting place to go. It's expensive though, and if you're going open source then they have a lot of warnings on the websites saying that you do this at your own risk, this is not a medical product, you may fry your brain...

What is the next step then?

Further projects would probably be replacing the motors. Because it's motor-driven, it's timing-based, so having something with servos instead where I can have a definite angle would be a lot more useful, a lot more precise and wouldn't tend to go... one of the problems with it is that if you tell it to keep going in one direction, it will keep going in one direction whether it wants to or not, and there's this awful grinding of gears as it attempts to go in one direction and can't. So that will probably be a new arm, a new robot, trying to get it a bit more nice-looking and more precise.



Joseph Thomas is a student helping to run a Raspberry Pi club from a science museum in Harlow, where they have worked on projects ranging from a robot arm to a portable Pi.

Like it?

The robot arm that Joseph is using can be bought from Maplins in the UK (bit.ly/1Da9BrT) or ordered from Adafruit elsewhere in the world (bit.ly/1yXIDQt). There are many guides online to get you up and running, such as this one: bit.ly/1AKd00U.

Further reading

NeuroSky has a whole product family dedicated to EEG and ECG biosensors, including the popular MindWave headsets (neurosky.com), and there are a few hacks available too (bit.ly/1C7w0SP). OpenBCI is a burgeoning open source project dedicated to brain-computer interfaces (openbci.com).

Another idea, which seems to be one of the most viable, is having it completely brain controlled



What you'll need

- Bare Conductive paint (pen or tub)
- Male to female jumper wires
- An assortment of LEDs, switches and resistors (optional)

Draw circuits with Bare Conductive paint



Aaron Shaw
Aaron volunteers at The MagPi (www.themagpi.com) and has been involved with Raspberry Pi since the start, finding himself lucky enough to use Pis at work and for play!

Assembling circuits has never been so easy with the joys of conductive paint, enabling you to bring together art and electronics in a whole new way

Playing with electronics and physical computing is a very rewarding task – it's capable of causing huge grins to develop on the faces of people of all ages. For a beginner though, the mess of wires and components can become very confusing quite quickly and things like soldering can be a safety concern when children are involved. Bare Conductive has taken the joy of electronics and made it far safer, easier and more versatile with their conductive paint. You can literally draw wires on paper with a paintbrush, use it for cold-soldering or a conductive adhesive and much, much more. There are not a great deal of boundaries to what you can do – if you can do it with normal paint and normal electronics, then you can do it with Bare Conductive paint (even multi-layer circuits are possible).

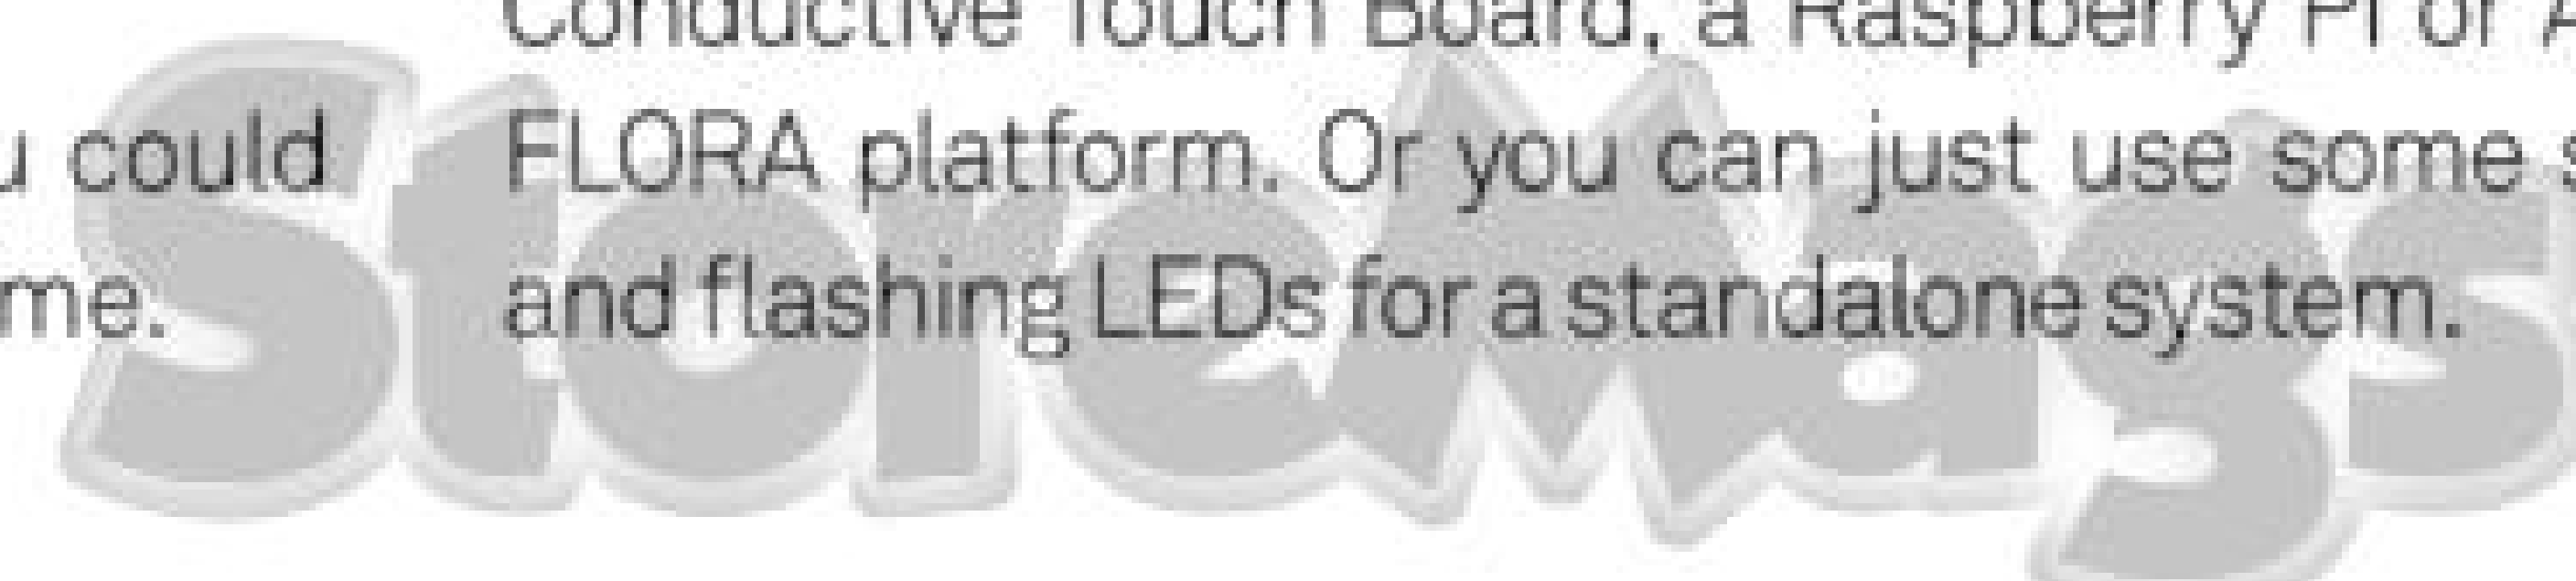
Pair this paint with a microcontroller board and you could be creating interactive art, clothing and projects in no time.

01 Get your tools


Paint and a paintbrush aren't the first items that come to mind when you think about electronics, so you may be wondering where to get them from. Bare Conductive stock the paint and a selection of components in their shop (bareconductive.com/shop) but you will need to go somewhere else for art supplies. We would recommend trying a high street craft shop such as Hobbycraft (hobbycraft.co.uk) or a local independent.

02 Pick your platform

The great thing about Bare Conductive paint is that, when dry, it works just like normal wiring! That means you can use it with any of your favourite microcontrollers like the Bare Conductive Touch Board, a Raspberry Pi or Adafruit's wearable FLORA platform. Or you can just use some small pin batteries and flashing LEDs for a standalone system.




Draw circuits with Bare Conductive paint



03 Start to paint

You can paint Bare Conductive paint onto pretty much any surface – paper, fabric, walls, clothing, wood, plastic and much more. For really accurate shapes and results, the best idea is to create or purchase a stencil (paper stencils are easiest to make at home but use vinyl for the best edge finish).



04 Connect it up

There are plenty of ways to connect to the conductive paint (from battery packs or microcontrollers for example) no matter what surface it's on, because once it is dry it acts just like an uninsulated wire. Therefore you can use wires glued on with the paint, paper clips, bulldog clips, alligator clips or even sewn-in conductive snaps for wearables projects.

05 Make repairs


The conductive paint is thick and when it's dry it becomes quite strong. These means you can use it to cold solder things together and repair any breakages. In other words, you could glue components into a circuit board or glue wires together and they would still function electrically. You can even use it to repair damaged tracks on circuit boards.

06 Clean up

A lot of you are probably thinking that something as cool as conductive paint is going to be nasty stuff. Actually Bare Conductive paint is non-toxic, water-based and water-soluble, and can therefore be cleaned easily with soap and water.

07 Make it waterproof

This paint only comes in black and is not waterproof. However, the great thing is that you can use it underneath or alongside any regular paints, varnishes and waterproofing sprays in order to act as insulation – or just to add some colour into your designs!



08 Touch and sound

Bare Conductive paint can also be used as a capacitive surface, meaning you can use it for touch, gesture or proximity controls when it is paired with a suitable control board. Bare Conductive make their own called the Touch Board which has everything you need to start experimenting with touch and sound. It can even act as a MIDI controller, an interface or an instrument!

Above Cut custom templates to suit your project's style and requirements

Touch Board

Bare Conductive's Touch Board is an Arduino device compatible with any existing shields and code you might have. It works with any conductive material as well as the Conductive Paint – you can even wire it up to a metal ruler. Using the Conductive Paint, you can also create touchless sensors, for example drawing and programming an electric drum kit that responds to waves and passes over your custom shapes.



Left Running a private cloud? Make sure no one can break into it...

Secure your Raspberry Pi



Christian Cawley

is a former IT and software support engineer and since 2010 he has written for computer and smartphone users, both online and in magazines

Concerned about the security of data stored on your Raspberry Pi? Protect yourself with passwords, firewalls and some physical security

There is a distinct security risk around your Raspberry Pi. Storing anything from passwords to firewalls, this important saved data can be stolen or pocketed with minimal effort if someone knows how.

Therefore it's a relief to learn that several tools, tricks and methods can be applied to keep your device and data away from prying eyes. You might, for example, be running a home security cam with images uploaded to a cloud account. These images would

be visible to anyone who possesses your Raspberry Pi's login details if you haven't bothered to change the defaults. Such a project (and many others) also demands that a firewall is installed for further improved security on a network.

Whether you're simply changing passwords, keeping your Pi under lock and key or installing a firewall, you'll be surprised at how easy it is to secure your Raspberry Pi and protect all of your important information and files.

What you'll need

- Velcro
- Adhesive putty
- Lockable cupboard, strongbox, etc.


```
login as: pi
pi@192.168.0.1's password:
Linux raspberrypi 3.12.30+ #730 PREEMPT Fri Dec 19 10:31:24 GMT 2014 armv6l

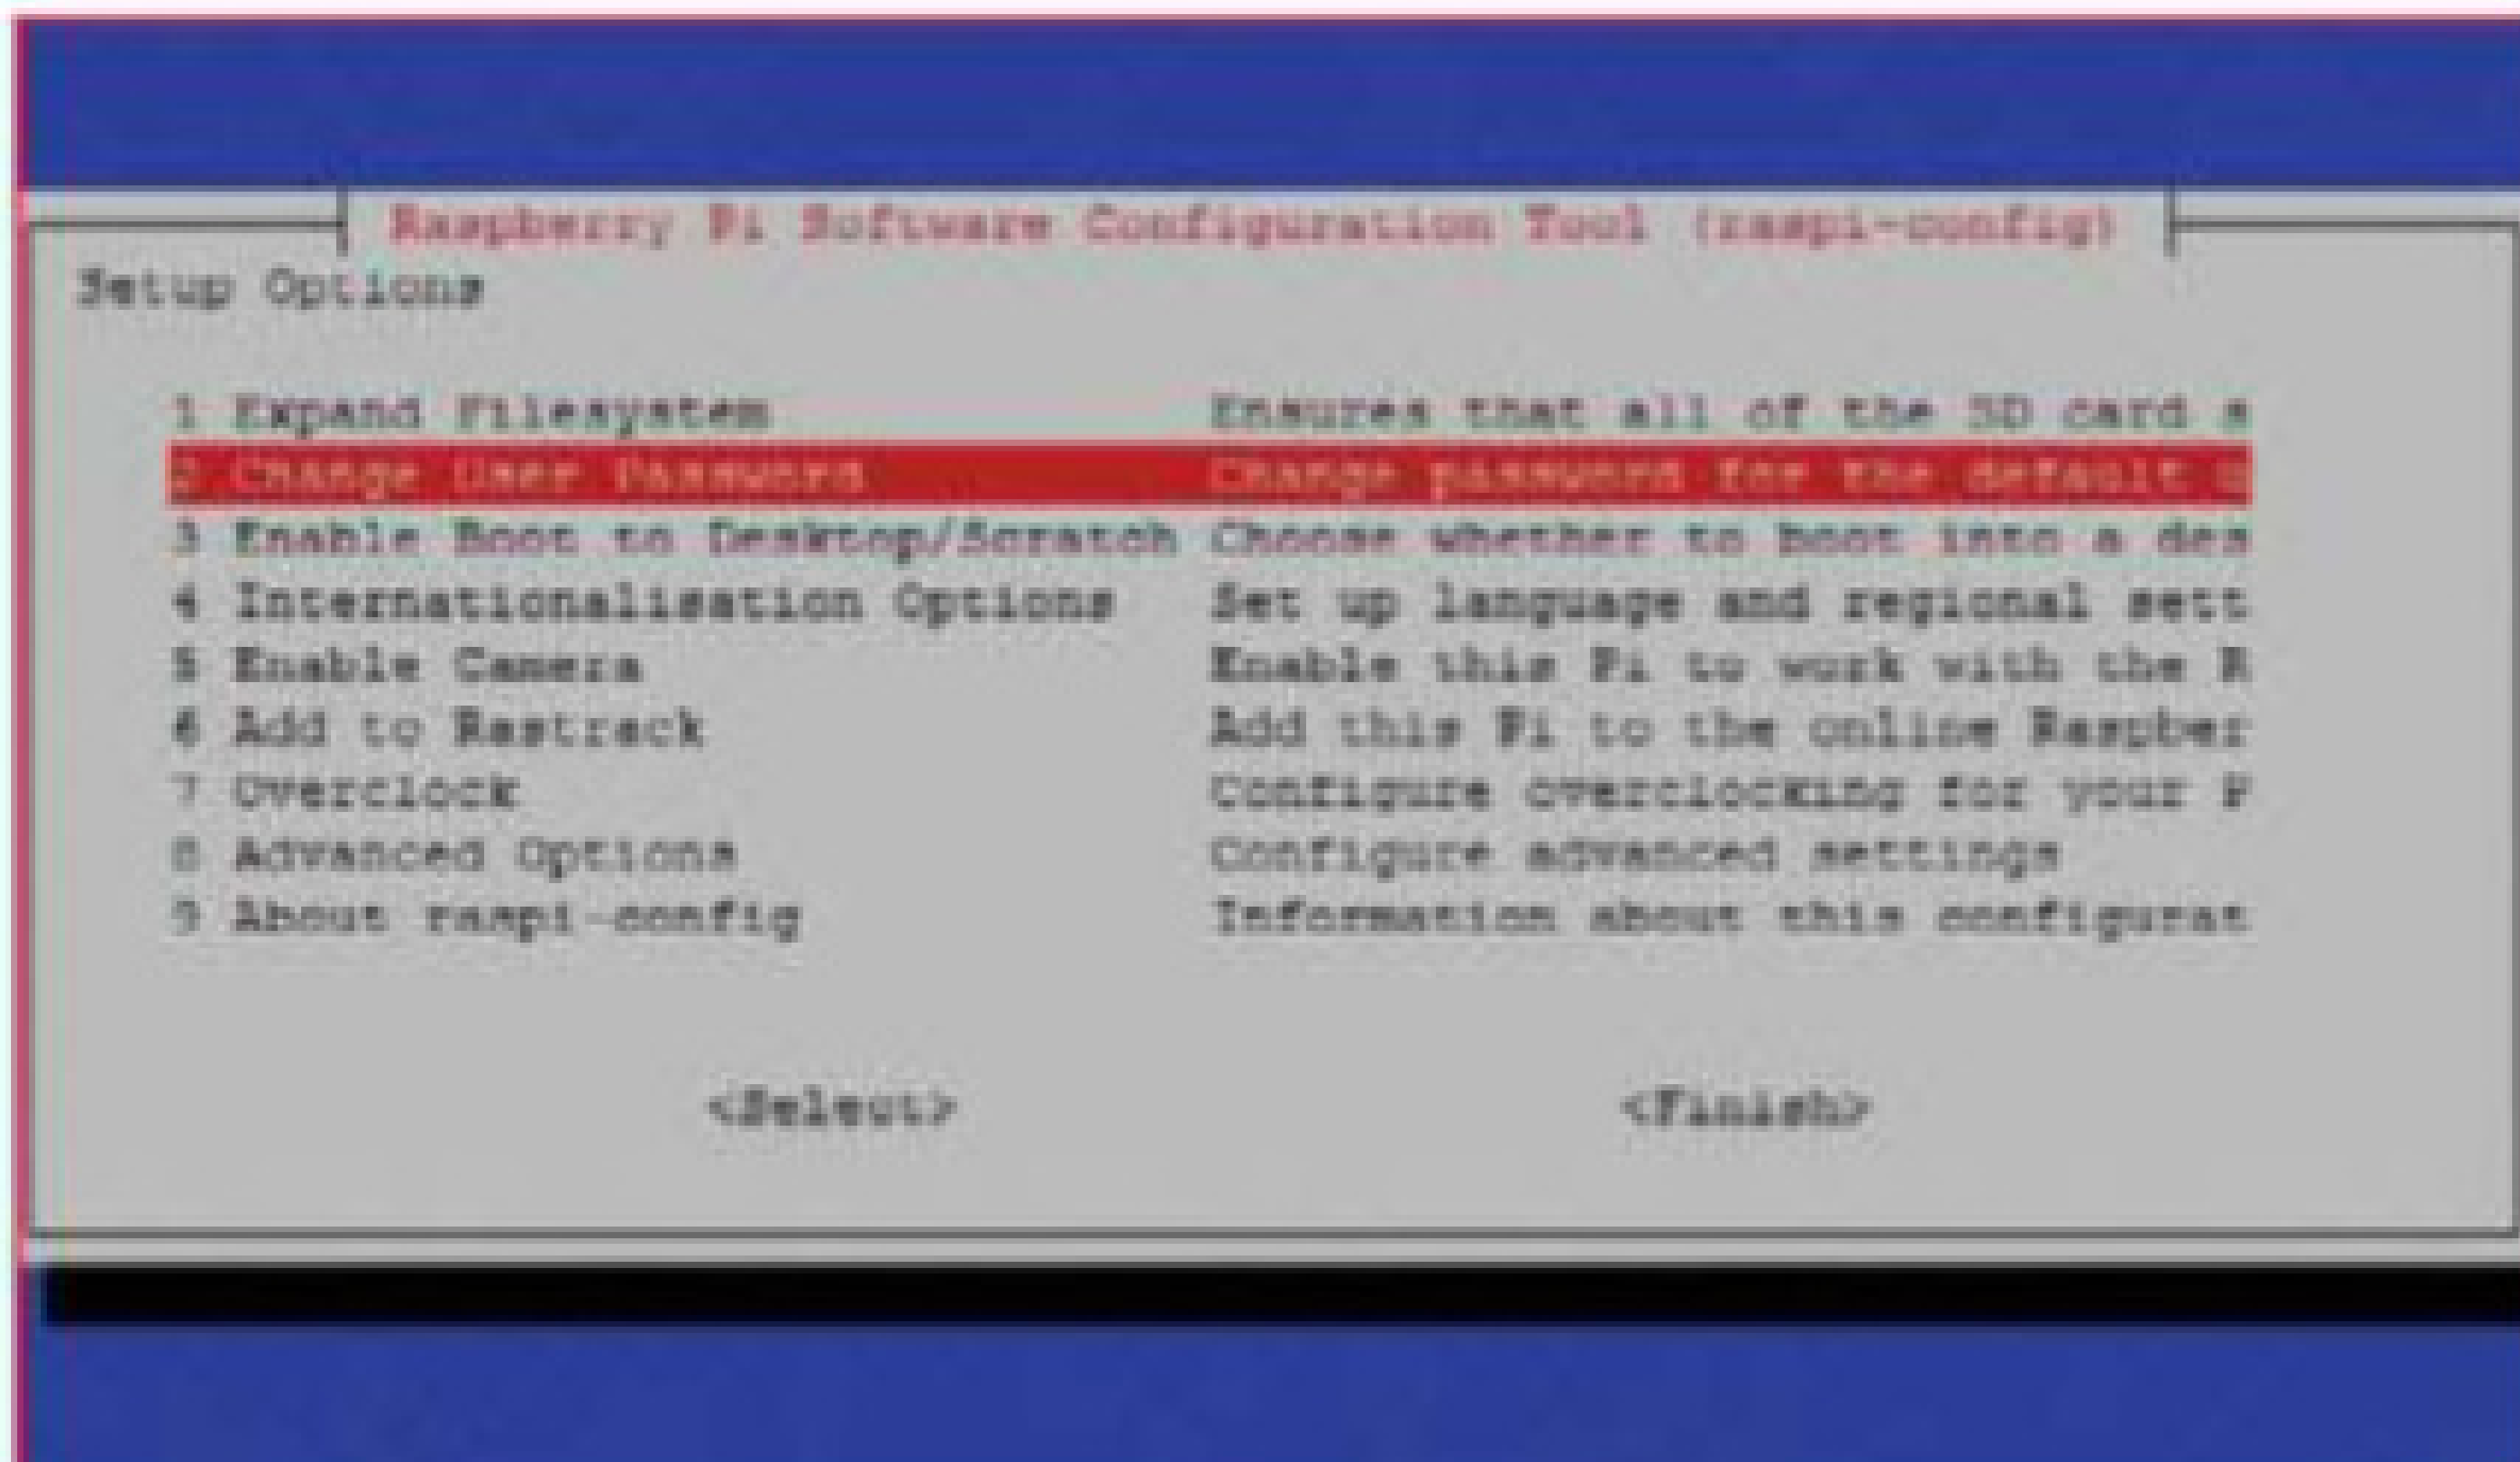
The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Wed Jan 21 10:30:02 2015
pi@raspberrypi ~$ passwd
Changing password for pi.
(current) UNIX password: █
```

01 Stop using the default password

Everyone who uses a Raspberry Pi knows that the default Raspbian credentials are 'pi' and 'raspberry'. Naturally, this means that anyone can sign into your computer if you haven't changed these defaults – something you'll need to do as a matter of urgency. After signing in, open the terminal and set a new password with:

```
passwd
```



02 Change password with raspi-config

If you're setting up a new installation of Raspbian, changing the password is one of the first things that you should do. With a new install, the first boot will automatically run the raspi-config screen.

Here, use the arrow keys to find the second option, change User Password and then follow the on-screen prompts to set yourself a new passcode.

```
pi@raspberrypi ~$ sudo useradd -m atomickarma -G sudo
pi@raspberrypi ~$ sudo passwd atomickarma
```

03 Create a new user account

To completely baffle anyone attempting to gain access using default credentials, take the most secure option and create a new user account. In the command line, enter:

```
sudo useradd -m username -G sudo
```

The `-m` switch creates a new home directory, while the second `sudo` adds the new account to the superuser group.

With a desktop computer and SD card reader, there is a way that you can recover your password

```
pi@raspberrypi ~$ sudo useradd -m atomickarma -G sudo
pi@raspberrypi ~$ sudo passwd atomickarma
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
pi@raspberrypi ~$ █
```

04 Give the new account a password

With the new account set up, the next step is to set a password. As you're not signed into the account at this stage, you won't be using the `passwd` command. Instead, enter:

```
sudo passwd username
```

With the new account ready to use, you should be ready to remove the default pi account from Raspbian altogether.

```
GNU nano 2.2.6 File: /etc/sudoers.tmp Modified
Defaults    secure_path="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin"

# Host alias specification

# User alias specification

# Cmnd alias specification

# User privilege specification
root    ALL=(ALL:ALL) ALL

# Allow members of group sudo to execute any command
sudo    ALL=(ALL:ALL) ALL

# See sudoers(8) for more information on "#include" directives:

#includedir /etc/sudoers.d
pi ALL=(ALL) NOPASSWD: ALL
atomickarma ALL=(ALL) NOPASSWD: ALL
```

05 Delete the default Raspbian account

You no longer need the default user account, pi. Sign out and login to your new account, and confirm it is correctly set up by opening:

```
sudo visudo
```

...and adding...

```
username ALL=(ALL) NOPASSWD: ALL
```

...to the final line. Save and exit with `Ctrl+X`. Now that's done, simply delete the old account with:

```
sudo deluser pi
```

Then remove the home directory:

```
sudo deluser --remove-home pi
```

06 Recover a lost password

If you've somehow forgotten your Raspberry Pi user account password or suspect that someone has changed it, what can you do?

With a desktop computer and SD card reader, there is a way that you can recover your password. Begin by inserting the Pi's SD card into your PC's card reader.

Use a proximity sensor

If you're genuinely concerned about your Raspberry Pi's physical security, you may consider employing some additional hardware to make it less of a target.

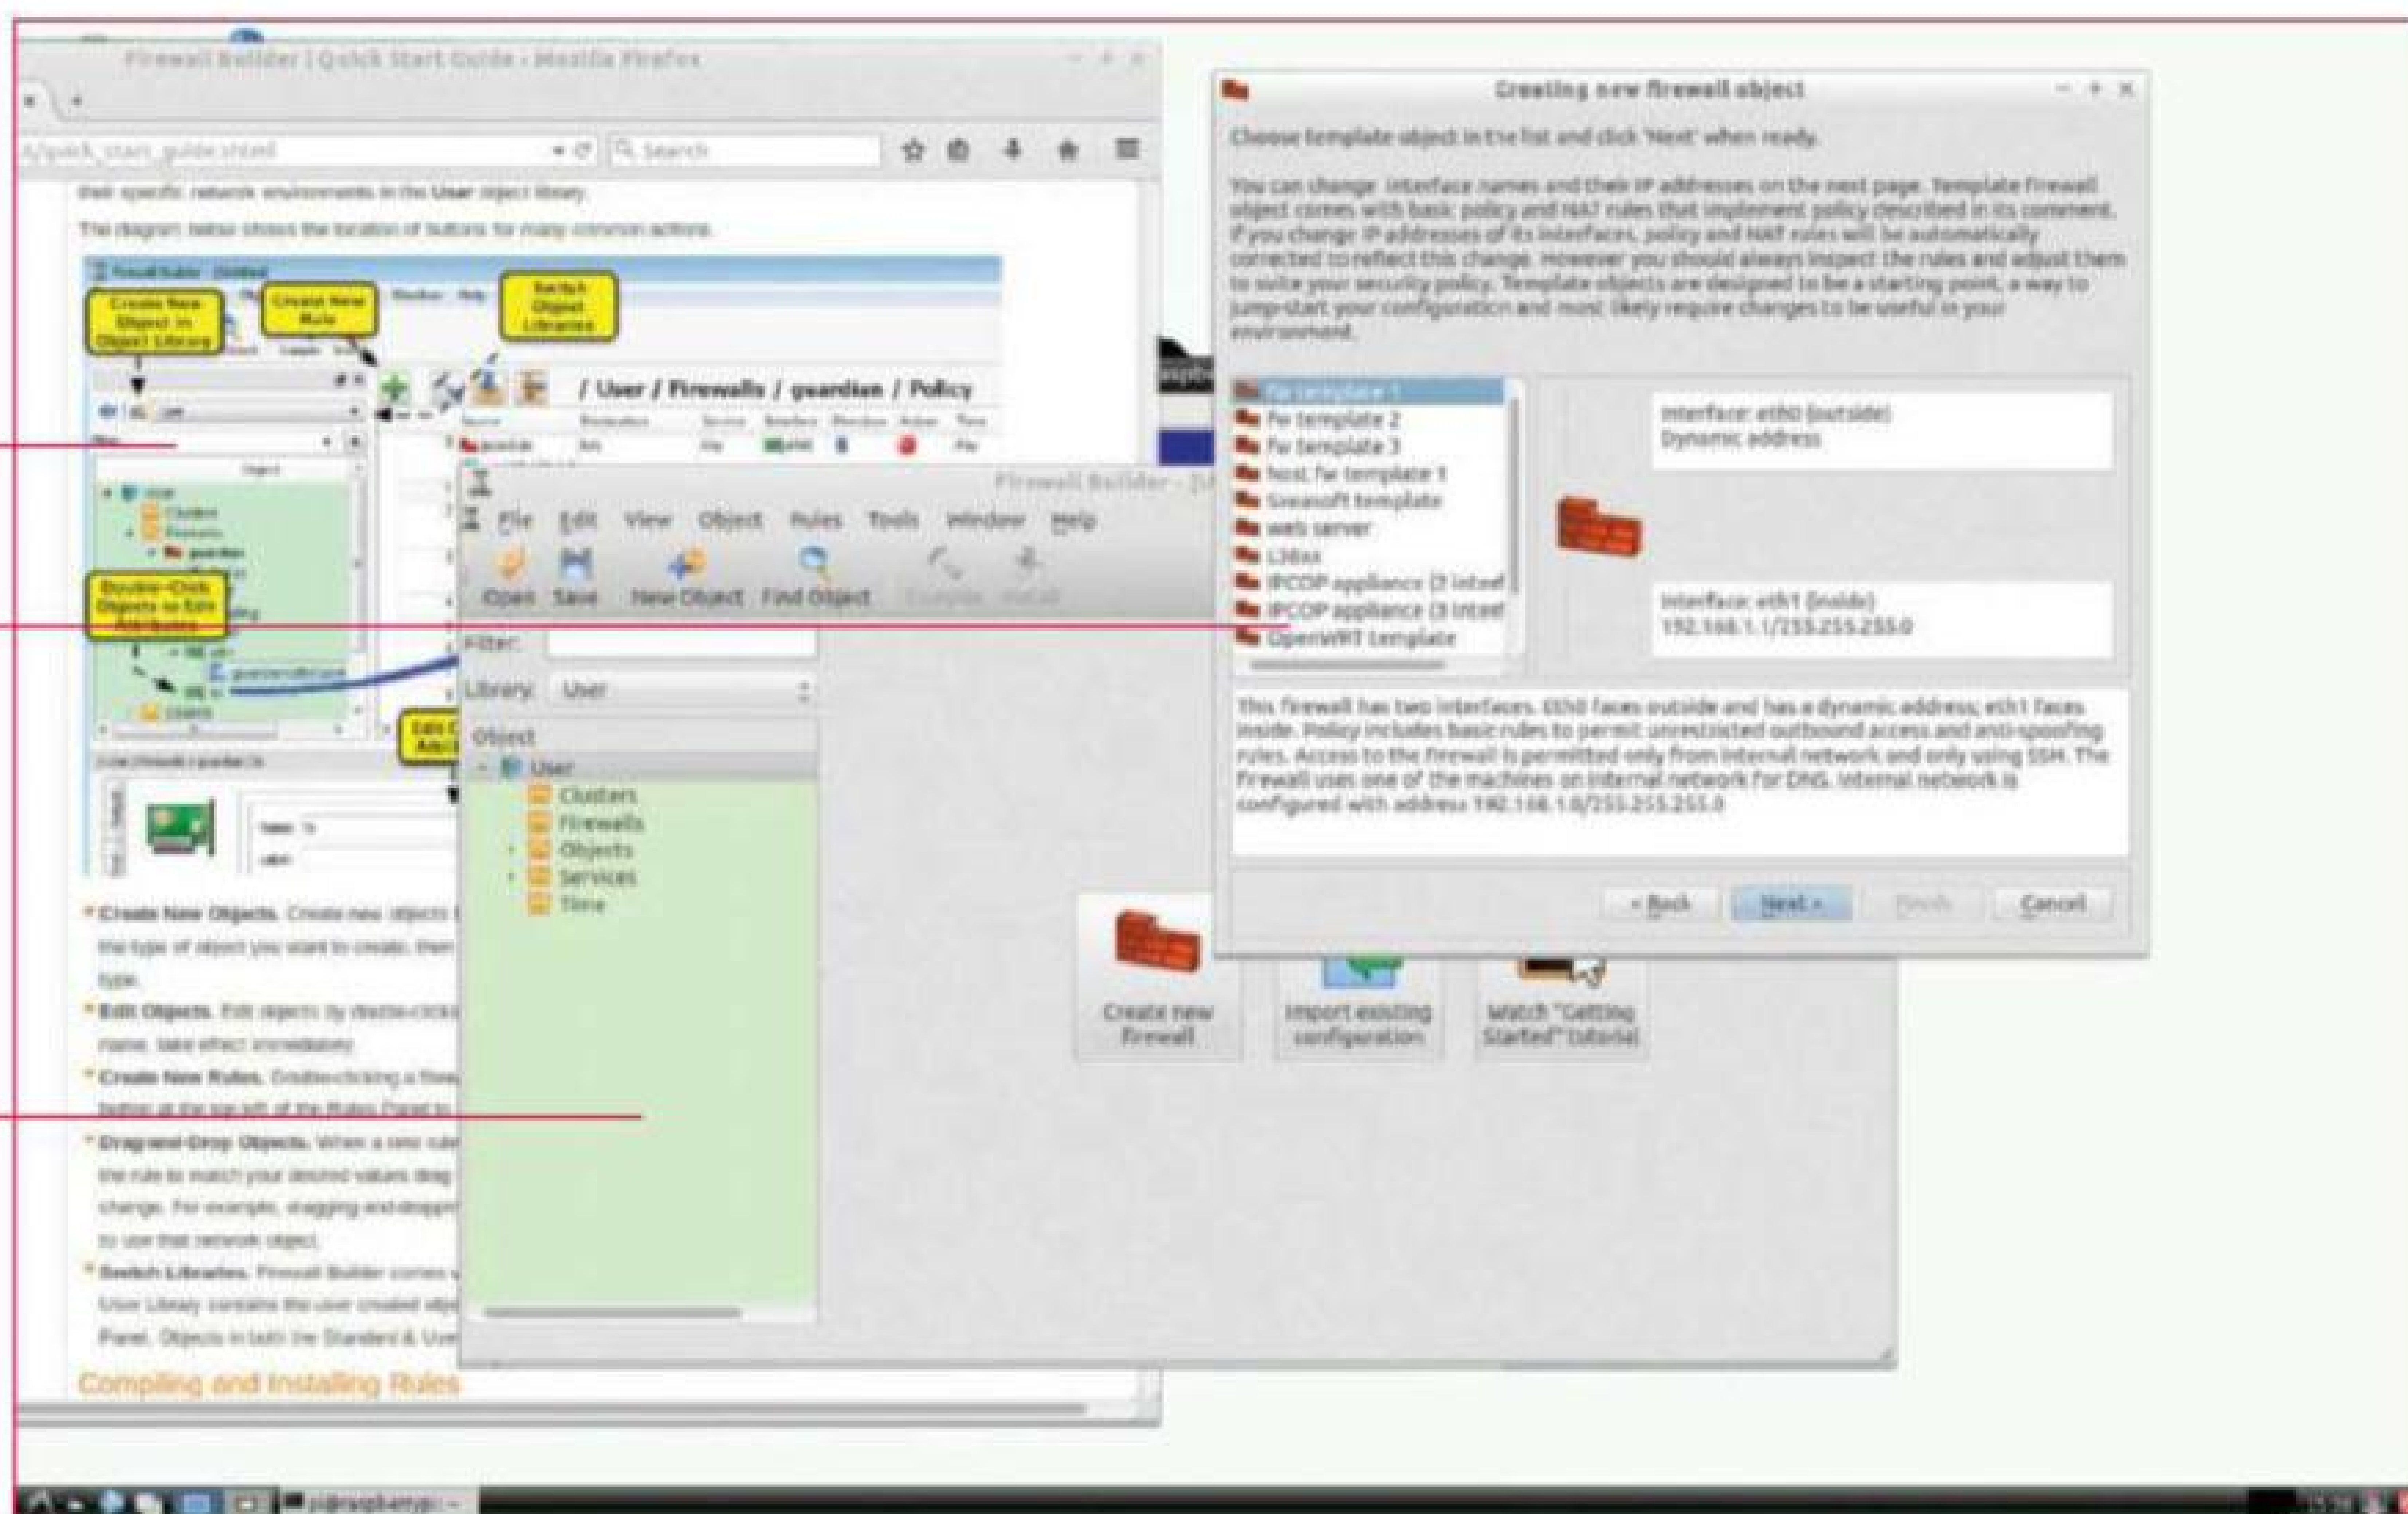
Your best option here is probably a proximity sensor configured to detect an unauthorised presence. When coupled with a buzzer, this can detect the presence of an intruder and alert you. It's even possible to configure such an alert as an email message if you're likely to be elsewhere, and it makes for a great side project.

Tutorial

Fwbuilder has a great quick-start guide that handily annotates the entire interface

Several firewall templates are available for the most common types of setup

The objects in this panel can be dragged out into the rules panel on the right-hand side



Hiding hardware

Putting your hardware out of sight and/or out of reach is a good option for security, and for something as small as the Raspberry Pi and an SD card you have quite a few options.

For instance, using Velcro or some adhesive putty you might attach the computer to the back of a cupboard or unit, kitchen kickboards or even under a car seat. The SD card, meanwhile, is so compact that you could easily place it under a carpet or even make a home for it in a cushion or shelf – just don't forget where you put it!

07 Edit cmdline.txt

Find the file `cmdline.txt` and open it in your Linux desktop text editor. Add the following to the end of the last line of the file:

```
init=/bin/sh
```

As the Raspberry Pi boots, this command will be read, enabling us to access a screen to reset the password. Save and eject the card.

08 Change the lost password

Unfortunately you won't be able to use SSH to recover the password, so instead connect a monitor and keyboard to your Raspberry Pi. Boot the Pi and wait for the prompt, at which point you should enter:

```
passwd username
```

Type the password, hit Enter and type it again to confirm.


09 Initialise the Raspbian boot

Thanks to the added code, we have changed the standard Raspbian boot to display a new prompt that will let us change the password.

When this is done, enter the following command to put things back in order:

```
sync  
exec /sbin/init
```

The Pi will now boot Raspbian normally, enabling you to sign in with the new password.



10 Revert cmdline.txt

We are not done yet though. Safely shutdown your Raspberry Pi with:

```
sudo shutdown -h now
```

With the Pi powered down, remove the SD card and insert it into the card reader again. Open `cmdline.txt` in your text editor once again and remove `init=/bin/sh`, then save and exit. This stops anyone else from resetting your password.

11 Physically secure your Raspberry Pi

Keeping digital intruders out of your Raspberry Pi with firewalls and secure account passwords is only part of the story. To fully protect your Pi you need to think outside of the box.

Barely larger than a credit card, the Raspberry Pi computer can easily be picked up and palmed. Physical security is paramount, but a genuinely secure Raspberry Pi case – for example, one compatible with Kensington locks – has yet to be released. However the ProtoArmour aluminium case from www.mobileappsystems.com can be screwed to a secure surface, which is great for more permanent project setups.

Below If you tend to access your Pi remotely via Wi-Fi, consider keeping it locked away



12 Lock it in a drawer

Probably the best way to keep your Raspberry Pi secure is to make sure you keep it locked in a drawer or cabinet – particularly useful if you use the device as part of a security cam system or as a cloud server storing valuable documents.

If no lockable storage is available and you're taking some time away from home where it isn't practical to take the Pi with you, another solution is needed. This might be to travel with your Pi's SD card in your wallet, perhaps leaving the computer attached to the back of a wardrobe with Velcro.

13 Add a firewall

Regardless of which operating system you're using, adding a firewall is a guaranteed way to improve your computer's security. While the Raspberry Pi has a built-in firewall, it is tricky to configure.

Thankfully, some other people have noticed this too and released fwbuilder, an interface to the otherwise complex iptables firewall that comes with Raspbian.

14 Install fwbuilder in Raspbian

Because iptables is a bit fiddly and errors can leave you with no network connection, fwbuilder has been developed to make firewall configuration quick and painless.

We'll use the **apt-get** command to first check for updates and then install fwbuilder:

```
sudo apt-get update
sudo apt-get install fwbuilder
```

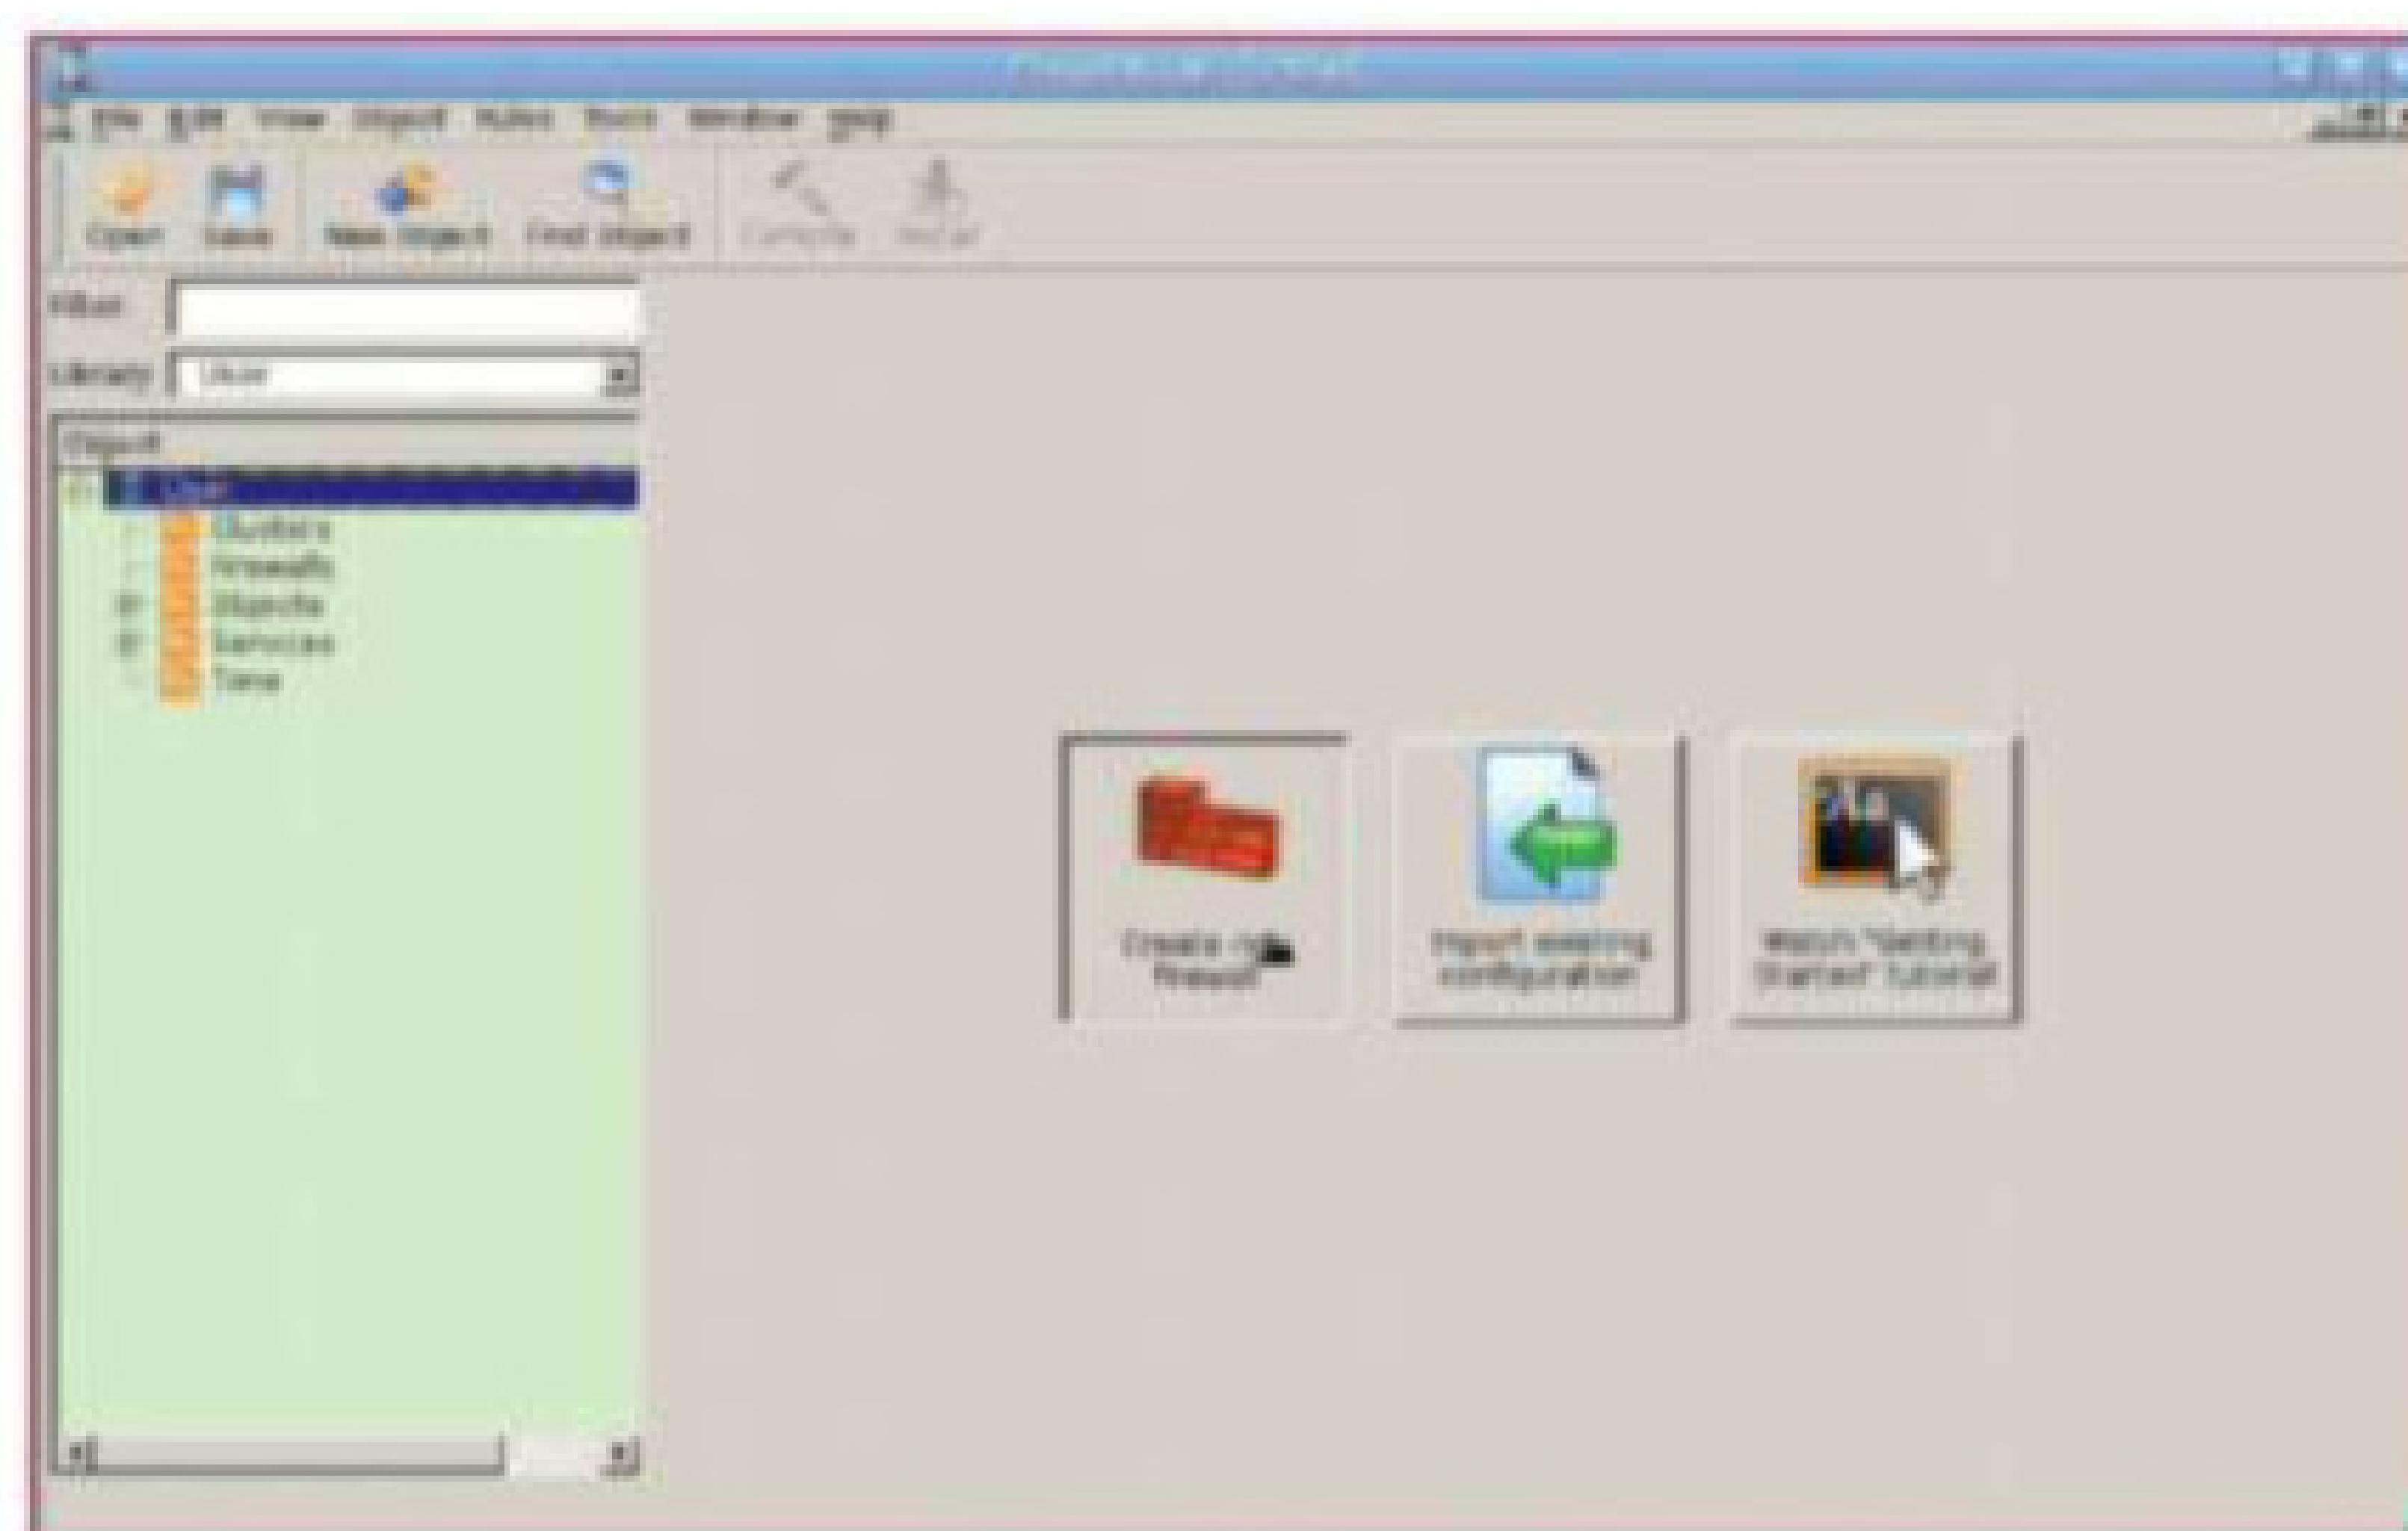
Follow the prompts to install and, once complete, switch to the Raspberry Pi GUI by entering:

```
startx
```

In the Pi's mouse-driven desktop, launch fwbuilder from the Internet menu. Upon launching fwbuilder, follow the given steps to set up your Raspberry Pi firewall and save the resulting script.

We're nearly done but some adjustments are still required before your Pi fully connects to the network.

A firewall is guaranteed to improve your security



15 Complete firewall configuration

Launch the `/etc/network/interfaces` script in your text editor and complete configuration by adding

```
pre-up /home/pi/fwbuilder/firewall.fw
```

Next, find the section labelled "Epilog" and add

```
route add default gw [YOUR.ROUTER.IP.HERE] eth0
```

If you're using a wireless card, add the same line but switch the last characters to wlan0:

```
route add default gw [YOUR.ROUTER.IP.HERE] wlan0
```

16 Consider Raspberry Pi theft

While losing your Raspberry Pi or the data on it, might initially seem like a disaster, don't be disheartened. As long as you have taken steps to backup data or clone your SD card, you at least have continuity when you resume the project. You can also check our boxouts for methods to help you deal with physical theft.

Pocket your Pi

If you're still concerned with your Pi's safety, put yourself in the place of a potential thief. Where would you stash it? Probably in your pocket. The Raspberry Pi is small enough to take with you, so why leave it lying around? Any security questions relating to your Raspberry Pi can be addressed by keeping it close whenever necessary.



**Rob
Zwetsloot**

models complex systems and is a web developer proficient in Python, Django and PHP. He loves to experiment with computing

Add web control and a camera to Rapiro

Make the Rapiro more wireless by installing a remote web interface and see where it's going in the process

What you'll need

- Raspberry Pi Model B
- Raspberry Pi Camera
- Latest Raspbian image
raspberrypi.org/downloads
- A wireless dongle





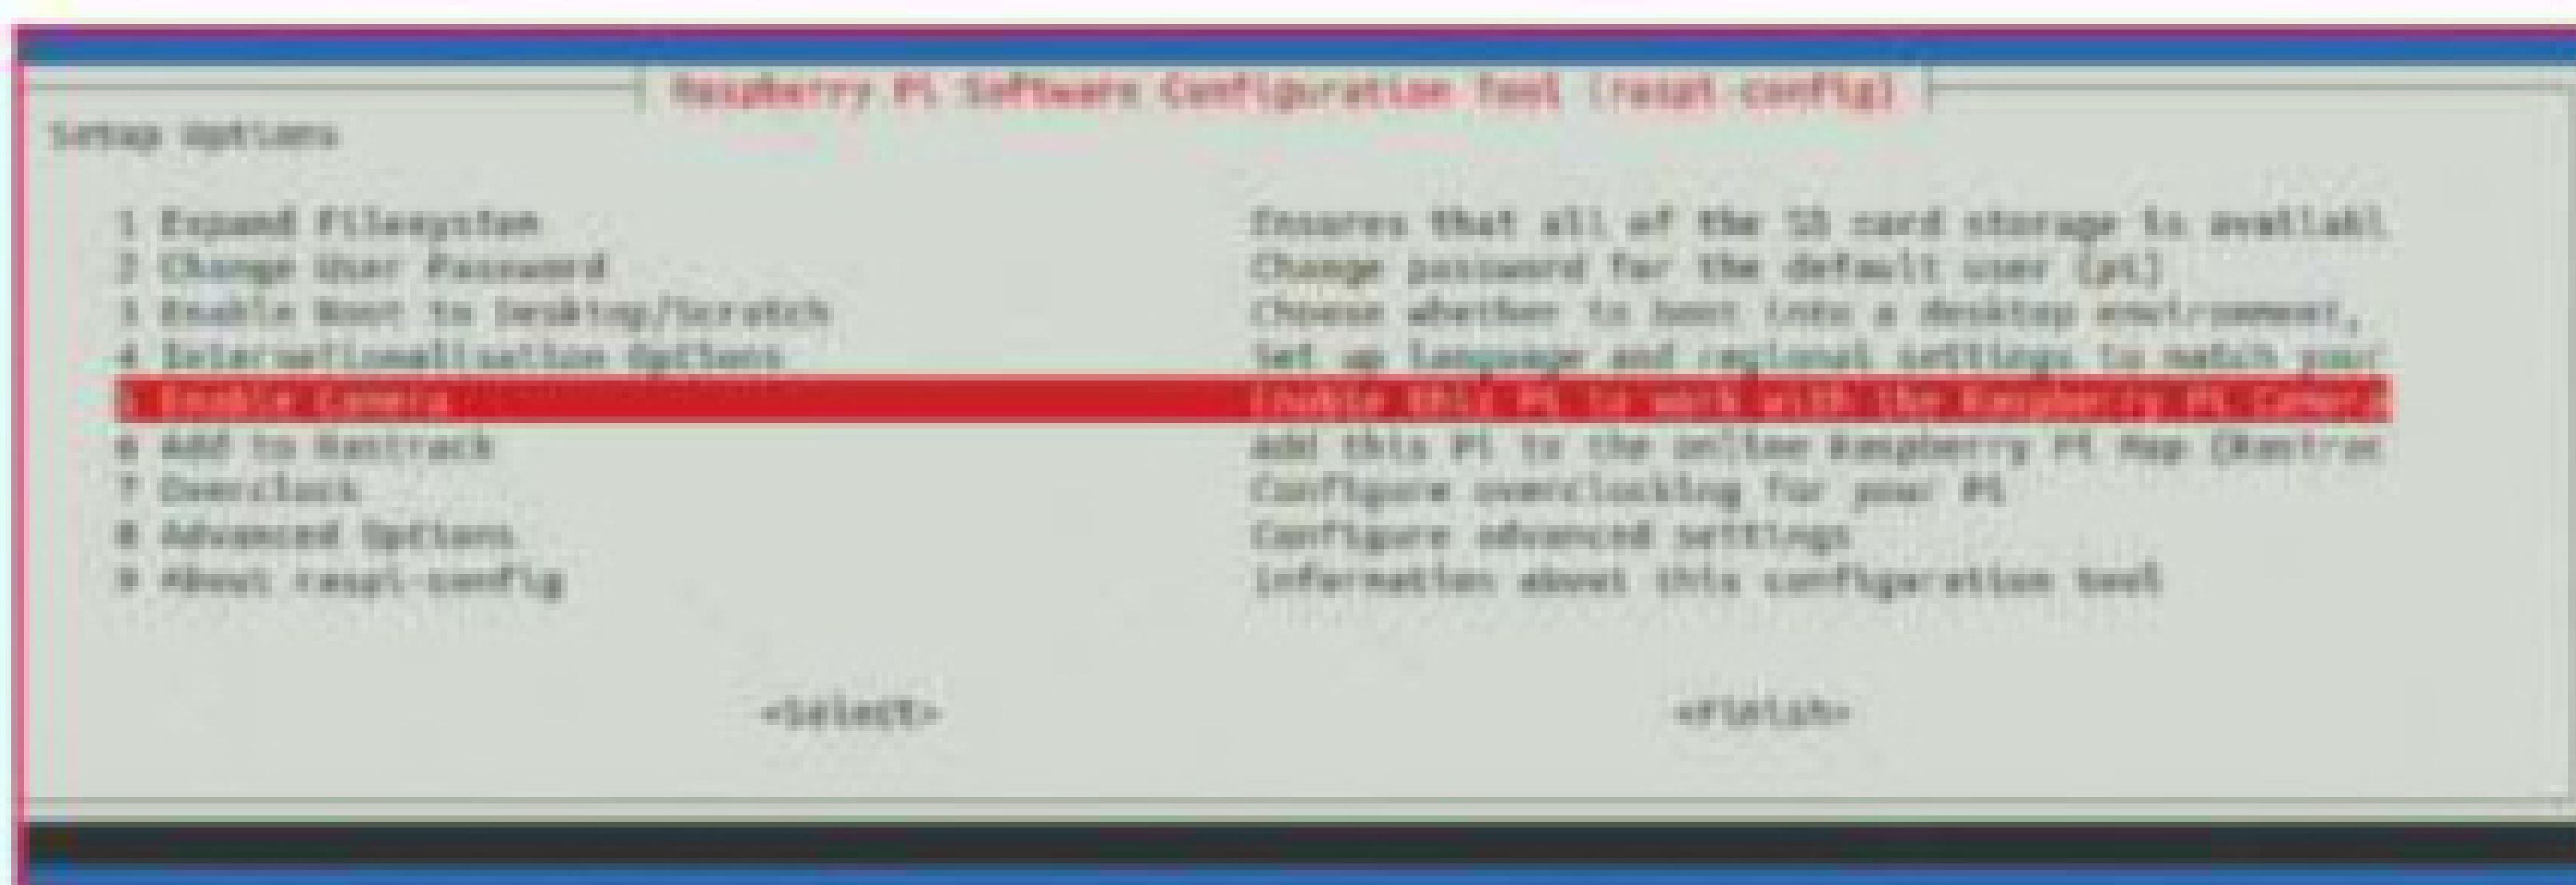
Left The design of the Rapiro means that it's quite simple to figure out where everything should sit

Last time, we added the ability to control the Rapiro using a PS3 controller. We want to give even better remote control to the user, so this time we're adding a web interface and installing the Pi camera into the head of the Rapiro so you can see what it's looking at.

Theoretically, as long as there's a constant wireless signal, you can see and direct the Rapiro as you wish. The only real limit is battery power. Part of this tutorial will show you how to set up a web server on the Raspberry Pi so that it can handle all the robotics and commands, and also stream the camera from the Rapiro.

01 Set up Raspbian

First, let's make sure we get Raspbian properly set up. Update your SD card or write the latest version to the card itself. For the update do the usual `sudo apt-get update && sudo apt-get upgrade`. If you've written a new card, insert it now.



02 Further set up

With a new card, turn on the Pi and get to the configuration screen. On an existing Raspbian, use `sudo raspi-config` in the terminal. Set the desktop to CLI to conserve the power, extend the file system and enable the Raspberry Pi camera module.

03 Connect to the Wi-Fi

Give it a quick reboot once you've finished configuring. If you haven't set up the Wi-Fi yet, the easiest way to do so is to go into the desktop with `startx` and use the graphical wireless network tool to set it up. It will save these settings even when booting into the command line. Make sure to also get the IP address assigned to the Pi using `ifconfig`.



04 Install the Raspberry Pi

Now that everything is set up, turn the Pi off and open up your Rapiro. Put the Pi in its slot and connect up the Arduino cable from the base. Take the front piece of the head and remove the plastic part above the eyes that plugs into the camera hole.

05 Attach the camera

Use the same screws that attached the plug to install the Pi camera – do it upside down so that the cable can fit. Once it's affixed securely, attach it to the Raspberry Pi in the camera port with the silver of the cable facing towards the front of the Rapiro. Reattach the head to the body.

06 Log in via ssh

Now your Rapiro is reassembled, plug in your Wi-Fi adapter and turn the Rapiro on. Give it a few moments and then try and log into the Raspberry Pi via ssh. From the terminal on another computer, use:

```
$ ssh pi@[IP address]
```

07 Disable getty


We need to disable the normal login screen on the Raspberry Pi so we can use the serial port. Access the init file using `sudo nano /etc/inittab` and find the following line:

```
T0:23:respawn:/sbin/getty -L ttyAMA0 115200 vt100
```

Comment it out by adding a # in front of the line.

3D-print Rapiro

As promised by creator Shota Ishiwatari during its initial Kickstarter campaign, the 3D models for Rapiro are now available to download for free from Thingiverse, so if you or a friend have a 3D printer, or you're happy paying a visit to one of the many 3D printing shops now appearing on the high streets, you can 3D-print your very own Rapiro. These days you can even print in colour, as well as all apply all sorts of finishing effects like metal and wood. Download the files from: thingiverse.com/thing:309466.

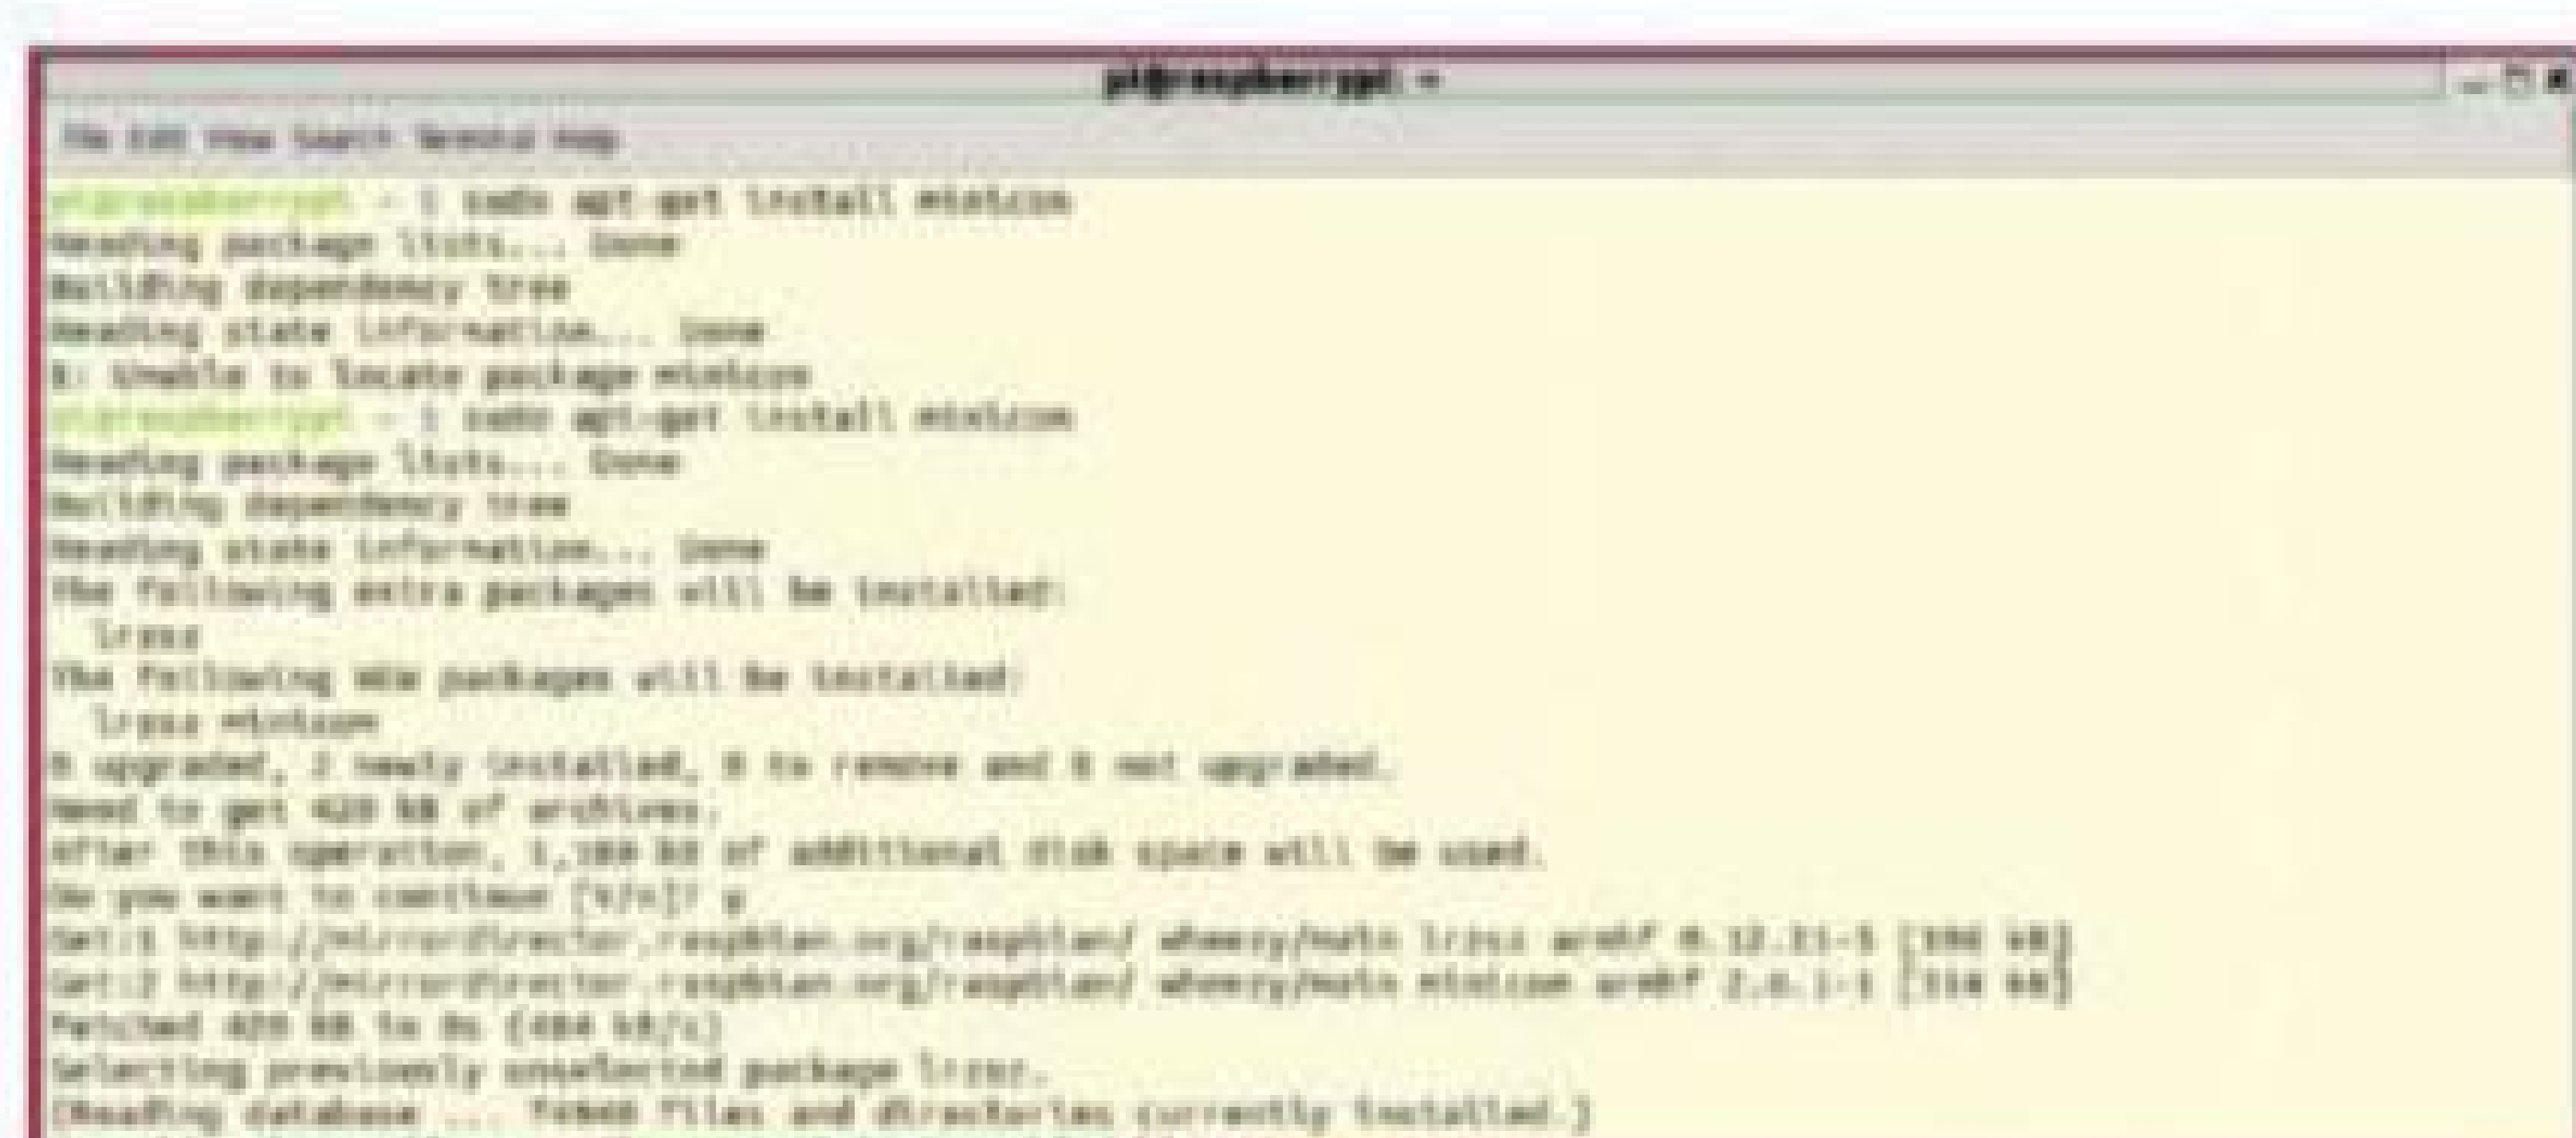


Above Be careful not to knock the camera module's ribbon out of place when you reattach the face

08 Stop serial ports at boot

Once the inittab is saved, we need to stop the Pi from trying to send out data to the serial port when it boots. Use `sudo nano /boot/cmdline.txt` and delete the following from the file before saving and rebooting the Pi:

```
console=ttyAMA0,115200 kgdboc=ttyAMA0,115200
```



09 Install command software

Using the Rapiro in the Pi means that you need extra software – if you did the remote control tutorial last time, you probably already have it. Just to be safe though, install it with:

```
$ sudo apt-get install minicom python-serial
```



10 Test the ports

Before we start testing what we've done, you'll need to change some permissions to a most foul number. Once it's done though, we can write a little Python script for testing:

```
$ sudo chmod 666 /dev/ttyAMA0
```

11 Do a script test

Create a new Python file in the terminal with `nano remotetest.py` and enter the following:

```
import serial
import time
ser = serial.Serial('/dev/ttyAMA0', 57600, timeout=1)
ser.open()

ser.write("#PR255G000B000T050")
time.sleep(2)
ser.write("#PR000G000B255T050")
time.sleep(2)
ser.write("#H")

ser.close()
```

12 Run the test

Save the file. This should change the eye colour on the Rapiro to red then back to blue. Run it and any other custom python scripts for Rapiro by using:

```
$ python remotetest.py
```



13 Install Apache

To get the web server working, we now need to install Apache onto the Pi. Do that and then configure it with:

```
$ sudo apt-get install apache2 libapache2-mod-python
$ sudo usermod -aG dialout www-data
```




Above The web interface has useful options for video recording, including motion detection

14 Download the server files

We now have three sets of files to download to two separate locations. First, use `cd` to move to `/usr/lib/cgi-bin/` and download the first set with:

```
$ sudo wget https://dl.dropboxusercontent.com/u/54426040/test.py
```

Change its permissions with `sudo chmod 755 /usr/lib/cgi-bin/test.py` then `cd` to `/var/www/` to download files two and three:

```
$ sudo wget https://dl.dropboxusercontent.com/u/54426040/control.css
```

```
$ sudo wget https://dl.dropboxusercontent.com/u/54426040/Remote.Control.Icons.zip
```

Unzip the last file with `sudo unzip Remote.Control.Icons.zip`.

15 Launch the interface

Finally, you can launch and control your Rapiro from `http://[IP address]/cgi-bin/test.py`. It's quite a basic interface that you can easily customise yourself to make it a little neater. This interface was created by eLinux user JaixBly.

16 Download the camera software

There is great camera software from Silvan Melchior on the Raspberry Pi forums that we can use to stream the Pi camera. Back in the terminal, download the files with:

```
$ git clone https://github.com/silvanmelchior/RPi_Cam_Web_Interface.git
```

Use `cd` to move to `Rpi_Cam_Web_Interface` and make the installer executable with:

```
$ chmod u+x RPi_Cam_Web_Interface_Installer.sh
```

17 Install the camera software

With the software downloaded and set up, it's time to properly install it. Do this using the following command and then reboot the Pi/Rapiro:

```
$ ./RPi_Cam_Web_Interface_Installer.sh install
```



18 Test out the camera

After the reboot, the red LED on the camera should light up to indicate it's on. To see the stream, go to `172.25.12.84` in the browser to see a stream with lots of details. It may be upside down so rotate the image.

19 Add files to controls

Back in the terminal, download some extra files for the web interface with:

```
$ wget https://www.dropbox.com/s/jt6sn1o3rytau7t/RaPiRo.zip
```

Unzip it with `unzip RaPiRo.zip` and copy the replacement files into their respective spots, do a `sudo chmod 755 /usr/lib/cgi-bin/rapiro.py` and then reboot.



Joey Bernard
As a true renaissance man, he splits his time between building furniture, helping researchers with scientific computing problems, and writing Android apps

Creating your own digital assistant

Everyone would like to tell their computers exactly what to do. Well with Python and a Raspberry Pi, now you can

Voice control

In this and further issues, we will look at the parts needed to make your own voice control software for your projects. If you want a virtual assistant, one project is the Jasper system (jasperproject.github.io). The documentation on the main website has a description of hardware to attach to your Raspberry Pi and a full set of instructions for installation and configuration. There is a set of standard modules included to allow interaction with various services. Use the time, Gmail or even the joke module, and there are also third-party modules for you to access. There is even a developer API and documentation to help you add your own functionality to Jasper.

Everyone who has watched the *Iron Man* movies has probably dreamt of having their own artificially intelligent computer system to do their every bid and call. While Jarvis has massive amounts of computing power behind him, you can construct the front-end with very modest resources. With a Raspberry Pi and the Python programming language, you can build your own personal digital assistant that can be used as a front-end to whatever massive supercomputing resources that you use in your day-to-day life as a playboy, philanthropist genius. We will go over the basics that you will need to know over the next few issues, so that by the end of the series you should be able to build your own rudimentary, customised agent.

The first step to interacting with the humans around us is to listen for verbal commands so that we know what we need to process. You have several options available to handle this task. To keep things simple, we will be dealing only with devices that are plugged into one of the USB ports. With that stipulation you can talk directly with the USB device at the

lowest level. This might be necessary if you are trying to use something that is rather unusual to do the listening, but you will probably be better off using something that is a bit more common. In this case you can use the Python module PyAudio. PyAudio provides a Python wrapper around the low level cross-platform library PortAudio. Assuming that you are using something like Raspbian for your distribution, you can easily install the required software with the command:

```
sudo apt-get install python-pyaudio
```

If you need the latest version you can always grab and build it from source. PyAudio provides functionality to both read in audio data from a microphone, along with the ability to play audio data out to headphones or speakers. So we will use it as our main form of interaction with the computer.

The first step is to be able to read in some audio commands from the humans who happen to be nearby. You will need to import the 'pyaudio' module

before you can start interacting with the microphone. The way PyAudio works is similar to working with files, so it should seem familiar to most programmers. You start by creating a new PyAudio object with the statement `p = pyaudio.PyAudio()`. You can then open an input stream with the function `p.open(...)`, with several parameters. You can set the data format for the recording; in the example code we used `format=pyaudio.paInt16`. You can set the rate in Hertz for sampling. For example, we are using `rate=44100`, which is the standard 44.1KHz sampling rate. You also need to say how big a buffer to use for the recording - we used `frames_per_buffer=1024`. Since we want to record, you will need to use `input=True`. The last parameter is to select the number of channels to record on, in this case we will use `channels=2`. Now that the stream has been opened, you can start to read from it. You will need to read the audio data in using the same chunk size that you used when you created the stream - it will look like `stream.read(1024)`. You can then simply loop and read until you are done. There are then two commands to shutdown the input stream. You need to call `stream.stop_stream()` and then `stream.close()`. If you are completely done, you can now call `p.terminate()` to shutdown the connection to the audio devices on your Raspberry Pi.

The next step is to be able to send audio output so that Jarvis can talk to you as well. For this you can use PyAudio, so we won't have to look at another Python module. To make things simple, let's say that you have a WAVE file that you want to play. You can use the 'wave' Python module to load it. Once again, you will create a PyAudio object and open a stream. The parameter 'output' should be set to true. The format, the number of channels and the rate is all information that will be derived from the audio data stored in your WAVE file. To actually hear



Right Check out the documentation to see what Jasper can do: bit.ly/1MCdDh4

Full code listing

```
# You need to import the pyaudio module
import pyaudio

# First, we will listen
# We need to set some parameters
# Buffer chunk size in bytes
CHUNK = 1024
# The audio format
FORMAT = pyaudio.paInt16
# The number of channels to record on
CHANNELS = 2
# The sample rate, 44.1KHz
RATE = 44100
# The number of seconds to record for
RECORD_SECS = 5

# Next, we create a PyAudio object
p = pyaudio.PyAudio()

# We need a stream to record from
stream = p.open(format=FORMAT, channels=CHANNELS,
                rate=RATE, input=True, frames_per_buffer=CHUNK)

# We can now record into a temporary buffer
frames = []
for i in range(0, int(RATE / CHUNK * RECORD_SECS)):
    data = stream.read(CHUNK)
    frames.append(data)

# We can now shut everything down
stream.stop_stream()
stream.close()
p.terminate()

# If we want to play a wave file, we will need the wave module
import wave

# We can open it, give a filename
wf = wave.open("filename.wav", "rb")

# We need a new PyAudio object
p = pyaudio.PyAudio()

# We will open a stream, using the settings from the wave file
stream = p.open(format=p.get_format_from_width(wf.getsampwidth()),
                channels=wf.getnchannels(), rate=wf.getframerate(),
                output=True)

# We can now read from the file and play it out
data = wf.readframes(CHUNK)
while data != '':
    stream.write(data)
    data = wf.readframes(CHUNK)

# Don't forget to shut everything down again
stream.stop_stream()
stream.close()
p.terminate()
```

the audio you can simply loop through, reading one chunk of data from the WAVE file at a time and immediately writing out to the PyAudio stream. Once you're done you can stop the stream and close it, as you did above.

In both of the above cases, the functions block when you call them until they have completed. What are the options if you want still be able to do processing while you are either recording audio or outputting audio? There are non-blocking versions that take a callback function as an extra parameter called `stream_callback`. This callback function takes four parameters, named `in_data`, `frame_count`, `time_info`, and `status`. The `in_data` parameter will contain the recorded audio if `input` is true. The callback function needs to return a tuple with the values `out_data` and `flag`. `Out_data` contains the data to be outputted if `output` is true in the call to the function `open`. If the `input` is true instead, then `out_data` should be equal to `None`. The `flag` can be any of `paContinue`, `paComplete` or `paAbort`, with obvious meanings. One thing to be aware of is that you cannot call, read or write functions when you wish to use a callback function. Once the stream is opened, you simply call the function `stream.start_stream()`. This starts a separate thread to handle this stream processing. You can use `stream.is_active()` to check on the current status. Once the stream processing is done, you can call `stream.stop_stream()` to stop the secondary thread.

Now that we have covered how to get audio information into and out of your Raspberry Pi, you can start by adding this functionality to your next project. In the next issue, we will look at how to convert this audio information into something usable by the computer by using voice recognition modules. We will also look at the different ways to turn text into audio output using TTS modules.

20
Raspberry Pi 2
 cases to be won
bit.ly/1CSrBG2



Rob Zwetsloot models complex systems and is a web developer proficient in Python, Django and PHP. He loves to experiment with computing

If you like this...

Check out our unboxing video where our senior staff writer Rob Zwetsloot goes through the changes on the Pi 2 with visual aids and pointing: bit.ly/1JrKc2i

Further reading

Check out last issue for our full review of the Raspberry Pi 2, as well as the Raspberry Pi Foundation website for other details: bit.ly/1LeTyKd

What is the Raspberry Pi 2?

Still confused over what the new version of the Raspberry Pi means to you? Missed our review last issue? We're here to fill you in

There is another Raspberry Pi then? Is it another update, a brand new thing or something entirely different? I'm getting a bit confused now!

This time it's the Raspberry Pi 2 Model B – basically a brand new Raspberry Pi.

Basically a brand new Raspberry Pi?

Well the thing that truly counts, the thing that makes it a Raspberry Pi, is the chip that powers it. The Broadcom BCM2835 chip is literally the core of the original Raspberry Pi, the system-on-a-chip that had the ARMv6 700 MHz processor and the VideoCore graphics chip. This was what was on the compute module from a while back, along with the 512 MB of RAM present

in the Pi Model B v2 and B+. Anyway, this chip has been updated and the RAM increased, which is what really makes it the Pi 2.

What is the new chip about then?

The BCM2836 is a slightly modified version of the 2835 – the main difference is that it now has an ARMv7 quad-core, 900 MHz chip and 512 Kb of cache. This is paired with an increased RAM of 1 GB and results in a huge power boost to the Raspberry Pi.

How much more powerful?

The Raspberry Pi Foundation estimates that it's about six times more powerful than the original Pi. From the tests

that we did, it means that we don't experience slow-down on conventional tasks like updating software or browsing the web for tutorial info, for example.

I've just seen a picture of it and it looks no different from the Model B+?

Yep, the main and only real noticeable difference is that the RAM is now on the underside of the board instead of the top. The redesign itself was actually made to accommodate the Pi 2 hardware, releasing a sort-of stop-gap device that also meant that manufacturers were able to get cases out into the wild before the 2 was a actually released.



Above This little processor here is what's boosting up the Raspberry Pi 2



Above It wasn't always the case, but the Raspberry Pi is now made in Wales here in the UK

Does it have Wi-Fi?

No, there is no Wi-Fi built in. Apparently that would have been difficult to accomplish. Instead, they've got a Pi-branded Wi-Fi dongle that is very cheap and very powerful.

Does everything work on it that worked on the original Raspberry Pi?

Basically, yes. It's best to grab the updated Raspbian that works better on ARMv7, but as the software itself uses the same packages and the same ports, there should be no problem with the way it works. The only issues you might come across are code snippets that take into account a slower Raspberry Pi, but those can be fixed.

Will the old Raspberry Pi's not be supported anymore?

They absolutely will be – the Raspbians and such will still work on the old Pis

for the time being, with no plans to end support. New software may require the Pi 2, but the Foundation can't really control that.

What about the Model A – will there be a new one of those as well?

At the moment, the Pi Foundation say there will not be a Model A. But they also said a Raspberry Pi 2 probably wouldn't be until 2017 when they were working on it for the February release this year, so we'll see.

Give me the damage then – what's the new price going to be?

It's staying at the same price – \$35 or about £25 – and for a much more powerful piece of kit as well. It means that the upgrade is cheap and affordable. You can grab it from the same places as well, like element14 and RS components.

RasPi magazine issue eight out now

Make a hand-held video player using your Raspberry Pi and a touchscreen

Did you know that we have a special digital magazine dedicated to the Raspberry Pi? **RasPi** magazine is available in the Google Play Store (bit.ly/1nRozP6) as well as the iTunes Newsstand (bit.ly/1vnUvHQ), and this month we take on the PiTFT, Twitter and VPNs.

Video player

The PiTFT is a great little screen that's a perfect fit for cases such as those sold by Adafruit, making it ideal for portable projects like a hand-held video player. Our complete guide shows you exactly how it's done, and we also show how the PiTFT can be used to make your own Raspberry Pi phone.

Master Twitter

Twitter really is everywhere these days – but how can you make use of the platform as the base for an interesting Pi project? We answer this question in two ways, walking you through the process of setting up first a wireless flood sensor that will tweet when water levels rise to a certain point – a great educational project as well as one with a range of practical uses. We also show you how to code your own Twitter bot that will trawl the timelines and automatically retweet from preset accounts for you.



Subscribe today & start saving
Non-USA readers see page 50

Subscribe



USA special offer

Subscribe today and get 5 free issues*



Your exclusive offer code: **USA**

Exclusive subscriber benefits

- Subscribe today and pay just \$130 for 13 issues*
- Pay as little as \$10 an issue – usually \$16.99 in stores

For more information and to order visit www.imaginesubs.co.uk/lud
Or call +44 (0) 1795 418661

Imagine Publishing publishes more than **20 monthly magazines**, some of which have been running for over **10 years**, reaching over **4 million readers** every year



Reviews

77 Group Test | 82 MinnowBoard Max | 84 OpenPi



GROUP TEST

Music players

What's the best software for playing and maintaining your music collection, as well as discovering new tunes online?

Nightingale

Forked from the Songbird project that was discontinued in 2013, Nightingale is based on the XULRunner framework and GStreamer libraries, offering a customisable panel-based layout with extensions for online services. Development is progressing well and a version 2.0 is in the works.

Download: bit.ly/1I27qv9

Tomahawk

Heavily influenced by iTunes and also OS X design in general, Tomahawk is a crisp, modern-looking music player with a core focus on hooking into online streaming, scrobbling and radio services. The software aims to provide users with a single access point for the multitude of online services and communities.

Download: bit.ly/1EGyqhd

CMus

A console music player that takes after Vi-style softwares, CMus (C* Music Player) is lightweight and fast. As well as being able to take advantage of scores of commands, hotkeys can be remapped to suit the user and the program is fully scriptable, so it can be hooked into other services with a little hacking.

Download: bit.ly/1j16Onu

GMusicBrowser

Built with GTK+ 2 and written in Perl, GMusicBrowser is a hugely customisable graphical application that is well established in the community and specifically focuses on taming massive music collections. There are tonnes of configuration options and tools for making your library more accessible.

Download: bit.ly/16A1MQq

Nightingale

Uncluttered and smart, Nightingale is a good choice for everyday listening



■ Album art, lyrics and bio info are scraped and displayed in the side panels

Playback

There are just the standard playback controls at the top-left of the interface, with no keyboard shortcuts beyond a Spacebar play/pause. There is an equaliser but it's completely manual with no presets to use; you'll need to adjust the frequency sliders by hand (and have an idea of how to do this).

Interface

Nightingale has a smart panel-based layout by default, which is customisable with different skins – or 'feathers' – and a traditional three-pane view (customisable to genre, artist or album), and then hideable panels for add-ons at the edges. These add-ons are well balanced, with album art in the corner, lyrics in a sidebar and extra info at the bottom.

Management

Metadata is pulled in from the tracks themselves, so if your library isn't well maintained then it will look gappy. However, you can batch edit metadata to fix this and locally scan for album artwork. Smart playlists can be generated and you can synchronise your library with set folders.

Online

The MLyrics and MashTape add-ons work really nicely, filling the gap from missing metadata in the panels, while the Soundcloud add-on lets you both play and download tracks and podcasts. Unfortunately, the SHOUTcast add-on for web radio was broken for us in v1.12. There's a good community effort here and dozens more homemade plug-ins.

Overall

Nightingale is probably the most balanced program in this group test, with a smart interface and decent playback, management and online tools. It is well-rounded, making it a great choice for everyday listening.

7

Tomahawk

Tomhawk has the best-looking interface and a unique approach



■ The wide selection of high-quality plug-ins is quite astounding

Playback

Playback itself is limited to the standard controls at the foot of the interface, with no keyboard shortcuts beyond the Spacebar play/pause. There are almost no options in the settings either. The contextual menu has some nice tools though, like finding similar tracks, queuing songs and sharing to friends.

Interface

One of those apps that seems made for Mac, and in this case very much made for OS X Yosemite, Tomahawk looks amazing. It's crisp, clean and modern, brilliantly structured and with the best handling of layout and album artwork of them all. It's also a very readable program and you can tell that a lot of thought went into the design.

Management

The side of the software dealing with your imported collection is hugely basic. There's a simple column view for your library and a search bar at the top. Even the playlist is a simple drag-to-add with nothing in terms of scanning through your library to certain criteria.

Online

Any gripes with local music management are swept aside by the online aspect. There are lots of plugins: Spotify, Google Play Music, Rdio, Soundcloud, Grooveshark, Beats, Deezer, Last.fm – the list goes on. You can search them all at once, set up custom radio stations and keep up with your in-app friends via the Feed and Inbox.

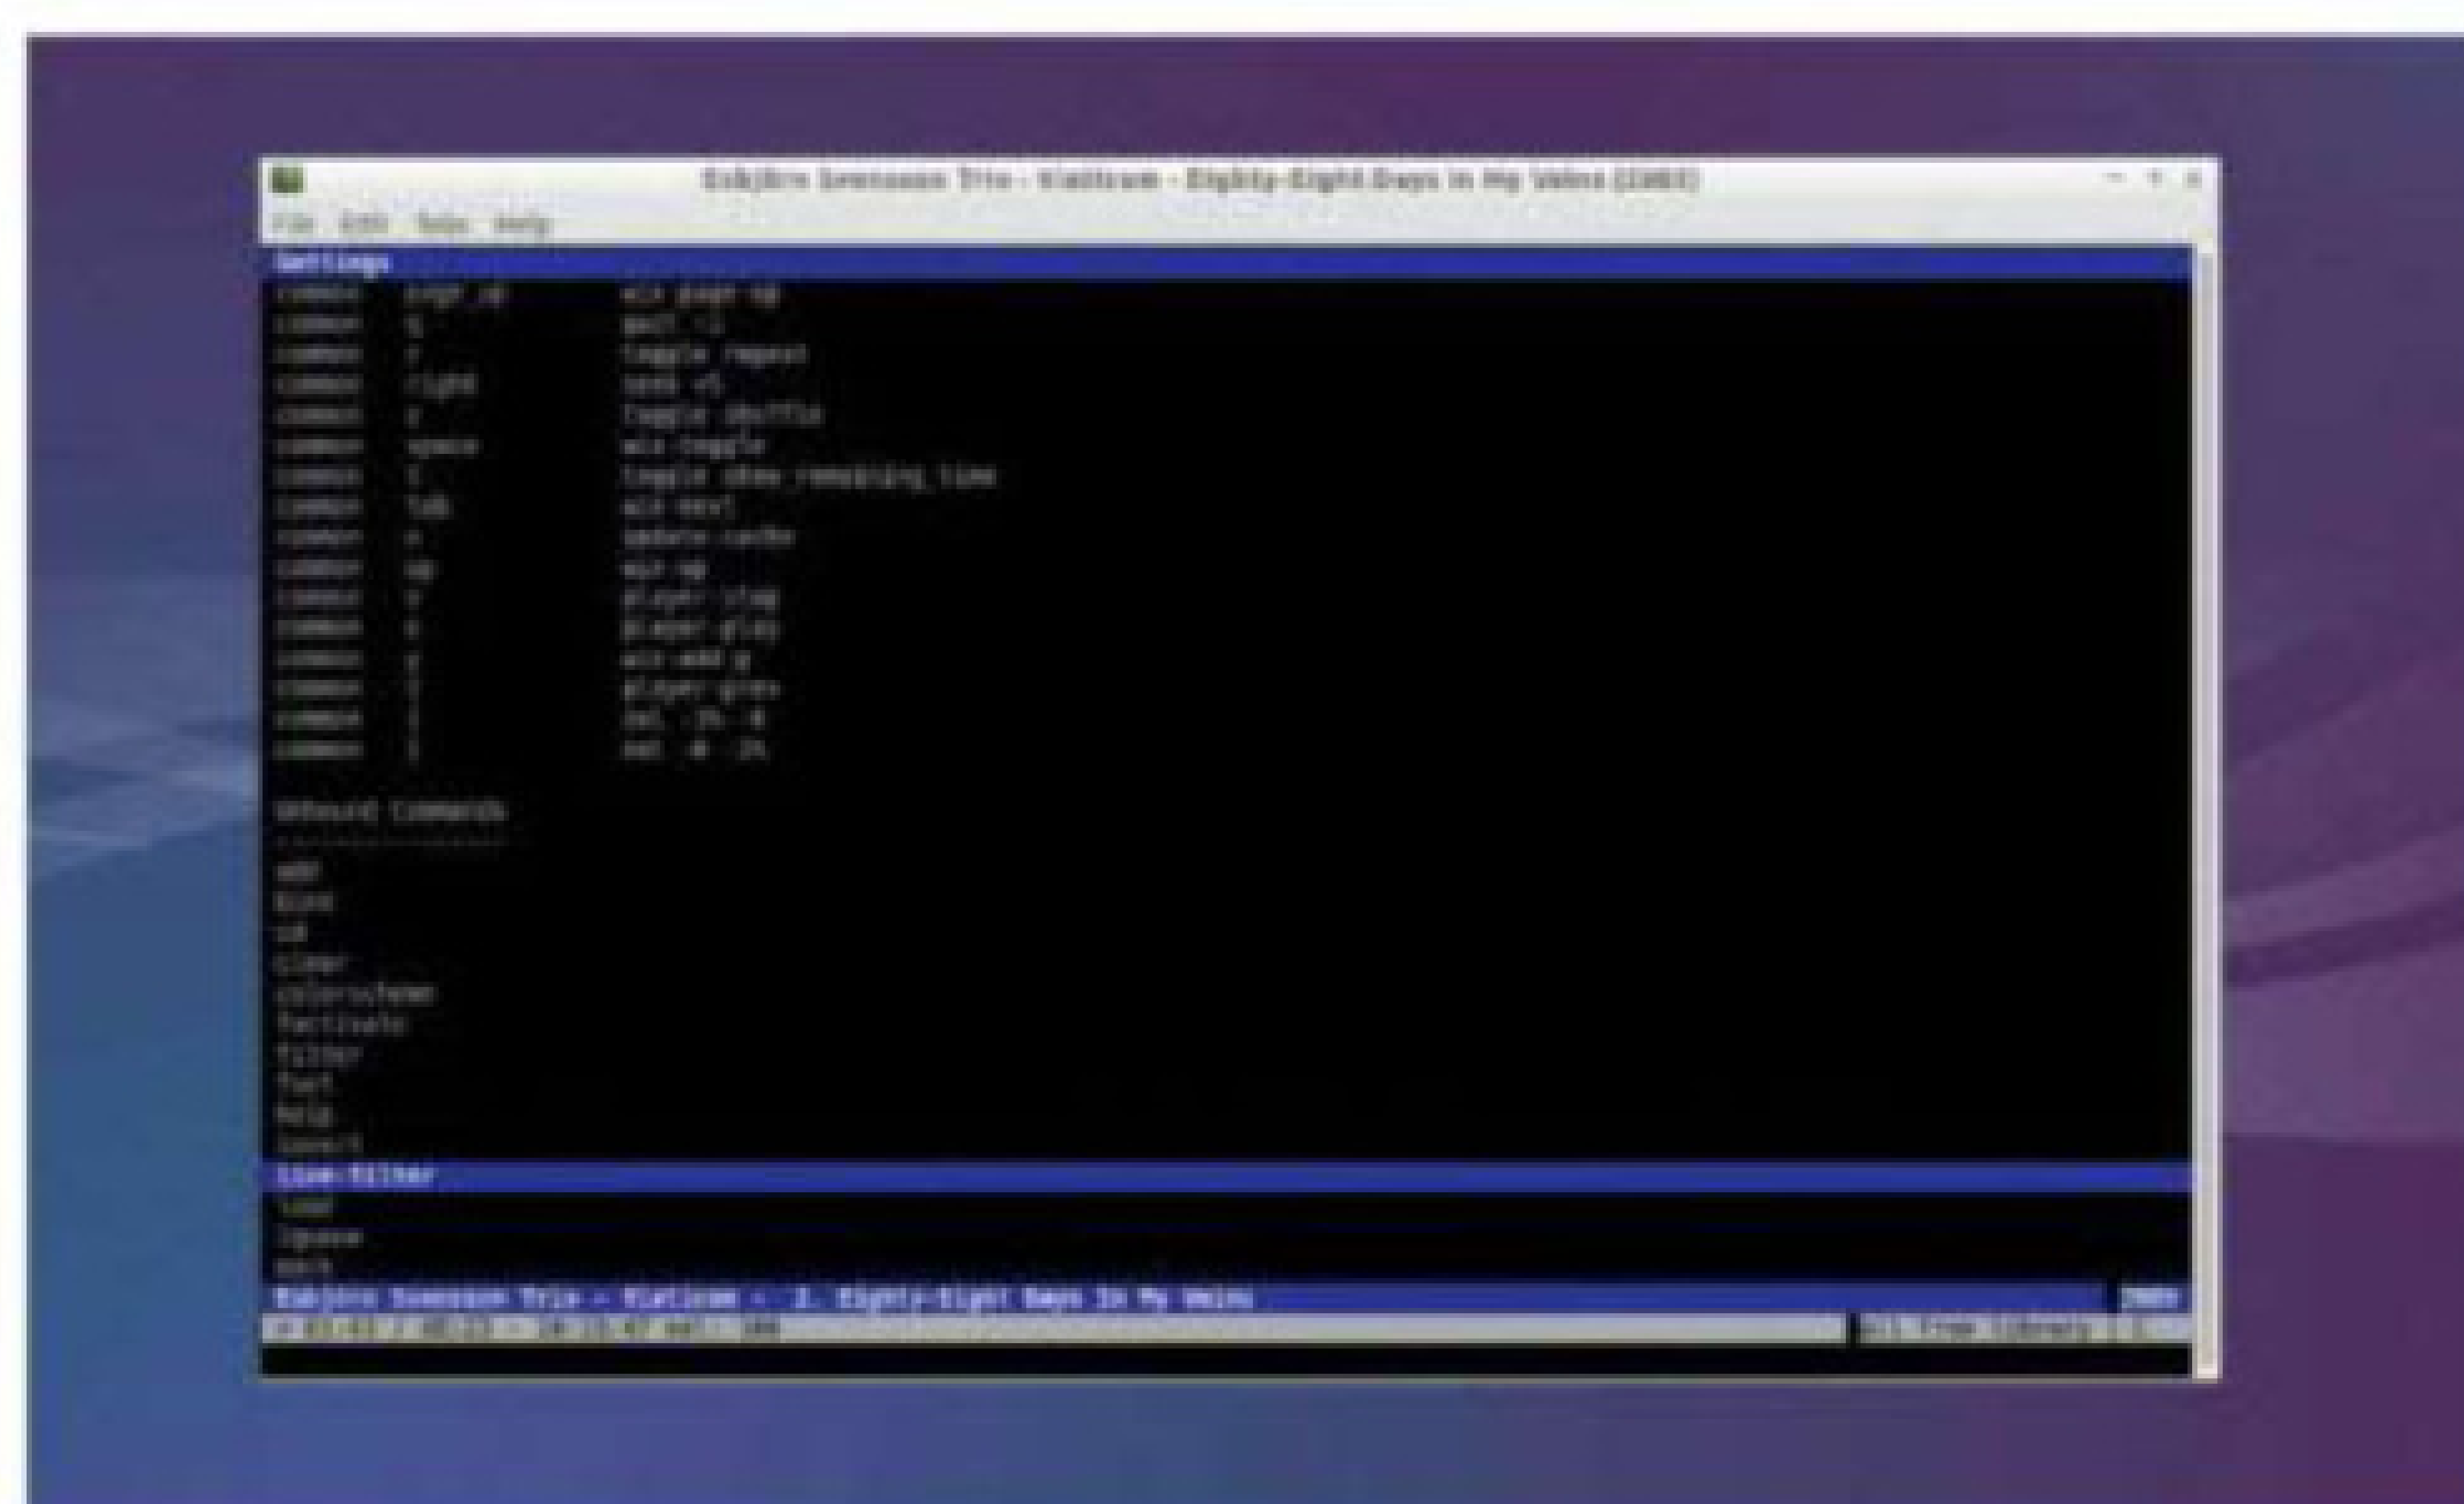
Overall

Playing music isn't the focus here – you hit play; job done – but that's not the objective. The idea is to provide a control centre to hook into the web services you use and make it social, and in these terms it's a fantastic program.

8

CMus

Prefer to keep everything inside the terminal? Give CMus your time



■ Every key is a shortcut and there are plenty of other unbound commands

Playback

CMus uses Vi-style search and command modes – you just need to remember them all! It works well, with commands responding instantly and enabling you to build up muscle memory. There are some nice configuration options too, such as resuming playback on boot and mapping a shuffle key.

Interface

As a command line interface CMus is beautifully simple, with the 1-7 keys displaying the seven views. Spend a little time getting to know the defaults or mapping your own commands and you'll have blazing fast speed in a no-frills environment. CMus is fully scriptable too, and you can also use cmus-remote to hook into it via other applications.

Management

As well as instant responses from commands, your library loads instantly every time. You can view metadata but not edit it, though this may be an advantage as some players can mess up your metadata. Searches are powerful and you can set up Filters that are stored between sessions.

Online

SHOUTcast and Icecast streaming is supported, and scripts are available to enable scrobbling to Last.fm/Libre.fm. By default there's nothing else really in terms of online services, but again because of its scriptable nature you can build your own tools or find those others have made. There are scripts out in the wild for pulling in lyrics, for example.

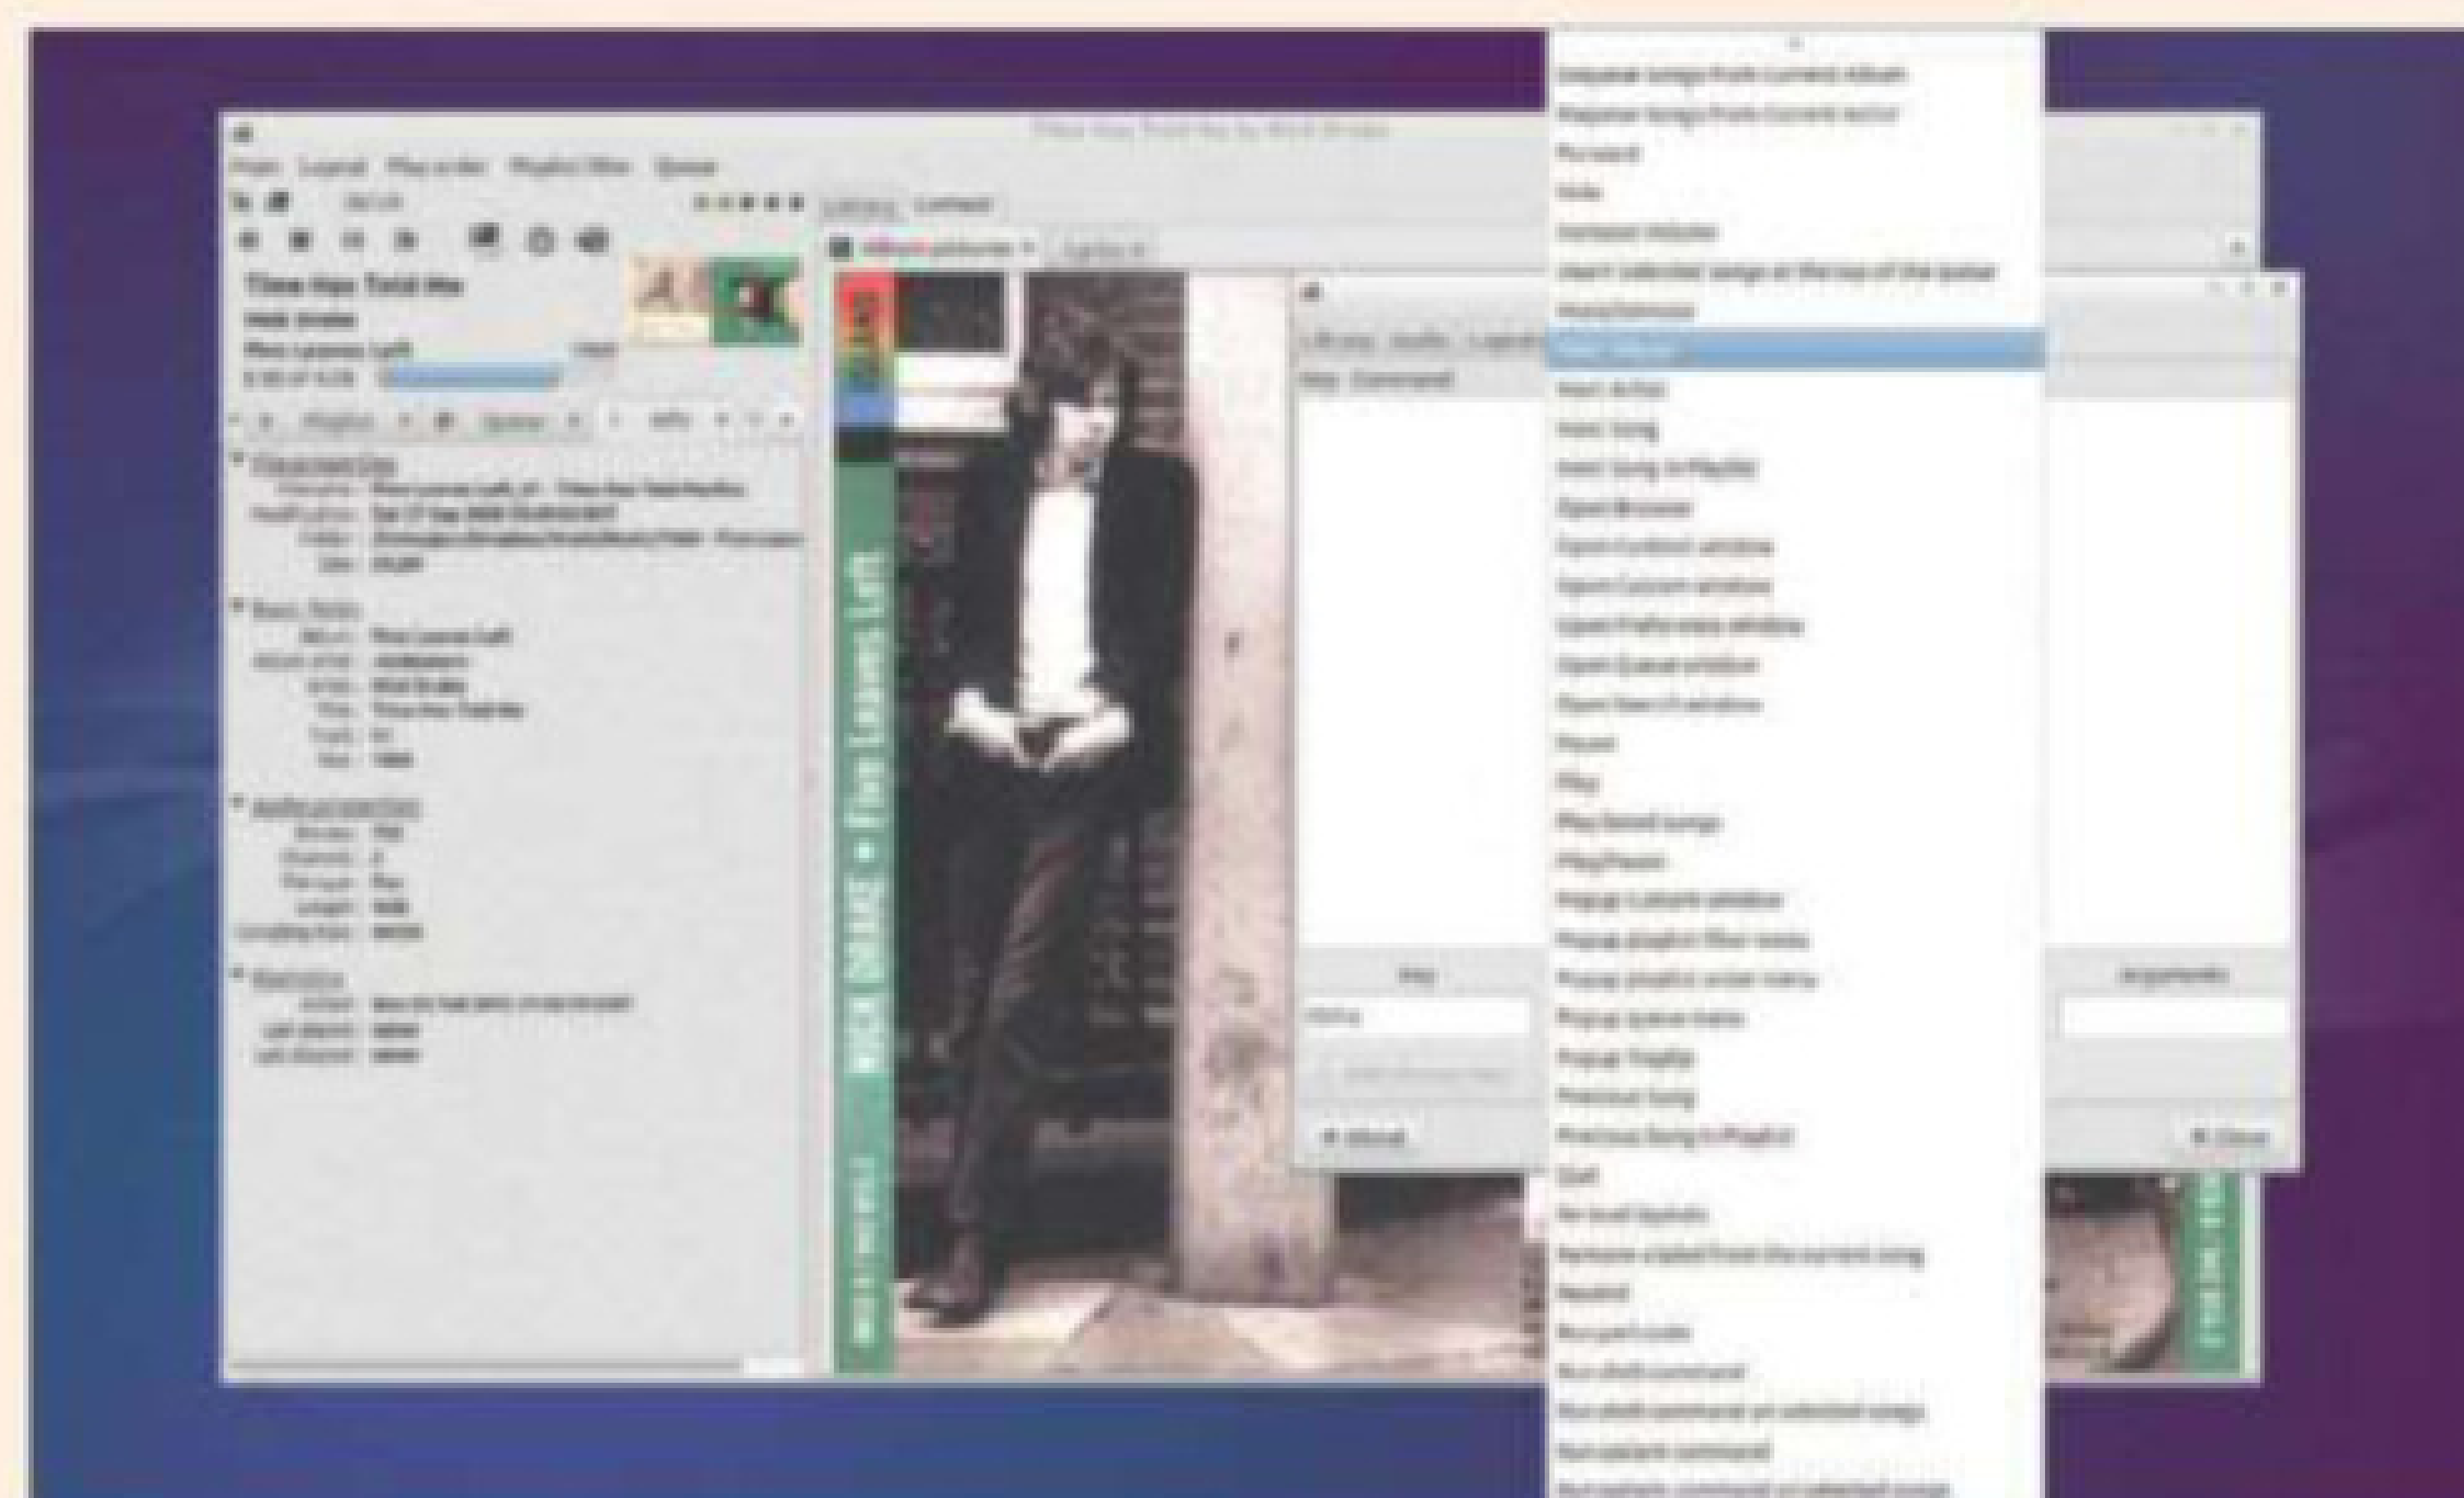
Overall

The focus on speed and customisation in this lightweight music player make it great for those wanting to stick to the CLI. It falls down on online services, so you may prefer to use a simpler program when deciding what to listen to next.

6

GMusicBrowser

GMusicBrowser is deceptively powerful despite its simple appearance



■ Hot keys can be configured for almost every feature inside the program

Playback

The play order, playlist filter and queue menus are powerful. Play order lets you select weighted shuffles that bias the selection to a set criteria, while the queue can be set to auto-fill as it empties or shutdown on completion. Custom keyboard shortcuts can be set for pretty much everything.

Interface

The standard three-pane view isn't quite as pretty as Tomahawk, but it's customisable and more useful. Similarly, the pane on the left can be switched from a detailed info view to show your queue, playlist or album art. A Context tab pulls in lyrics and album art, and there are dozens of skins – even a 'Make it look like [e.g. RhythmBox]' option.

Management

Customisable layouts, powerful search/filter options (unlimited nesting) and the ability to batch edit the metadata for tracks make this a great tool for getting huge libraries into shape. Editing data, you can get GMB to autosuggest track numbers and titles based on file names.

Online

GMusicBrowser is a little limited in this regard. Accessing online services isn't really the point of GMB as it focuses on your local collection, so about all you can do here is enable a plugin for Last.fm scrobbling. As said though, it does pull in lyrics and album art automatically and there's also support for Icecast servers.

Overall

Powerful and configurable, GMB trounces Nightingale in a couple of key categories. The fully-featured software lacks online services though, so if you don't care about those then it's definitely worth tweaking GMB to your liking.

7

In brief: Compare and contrast our verdicts

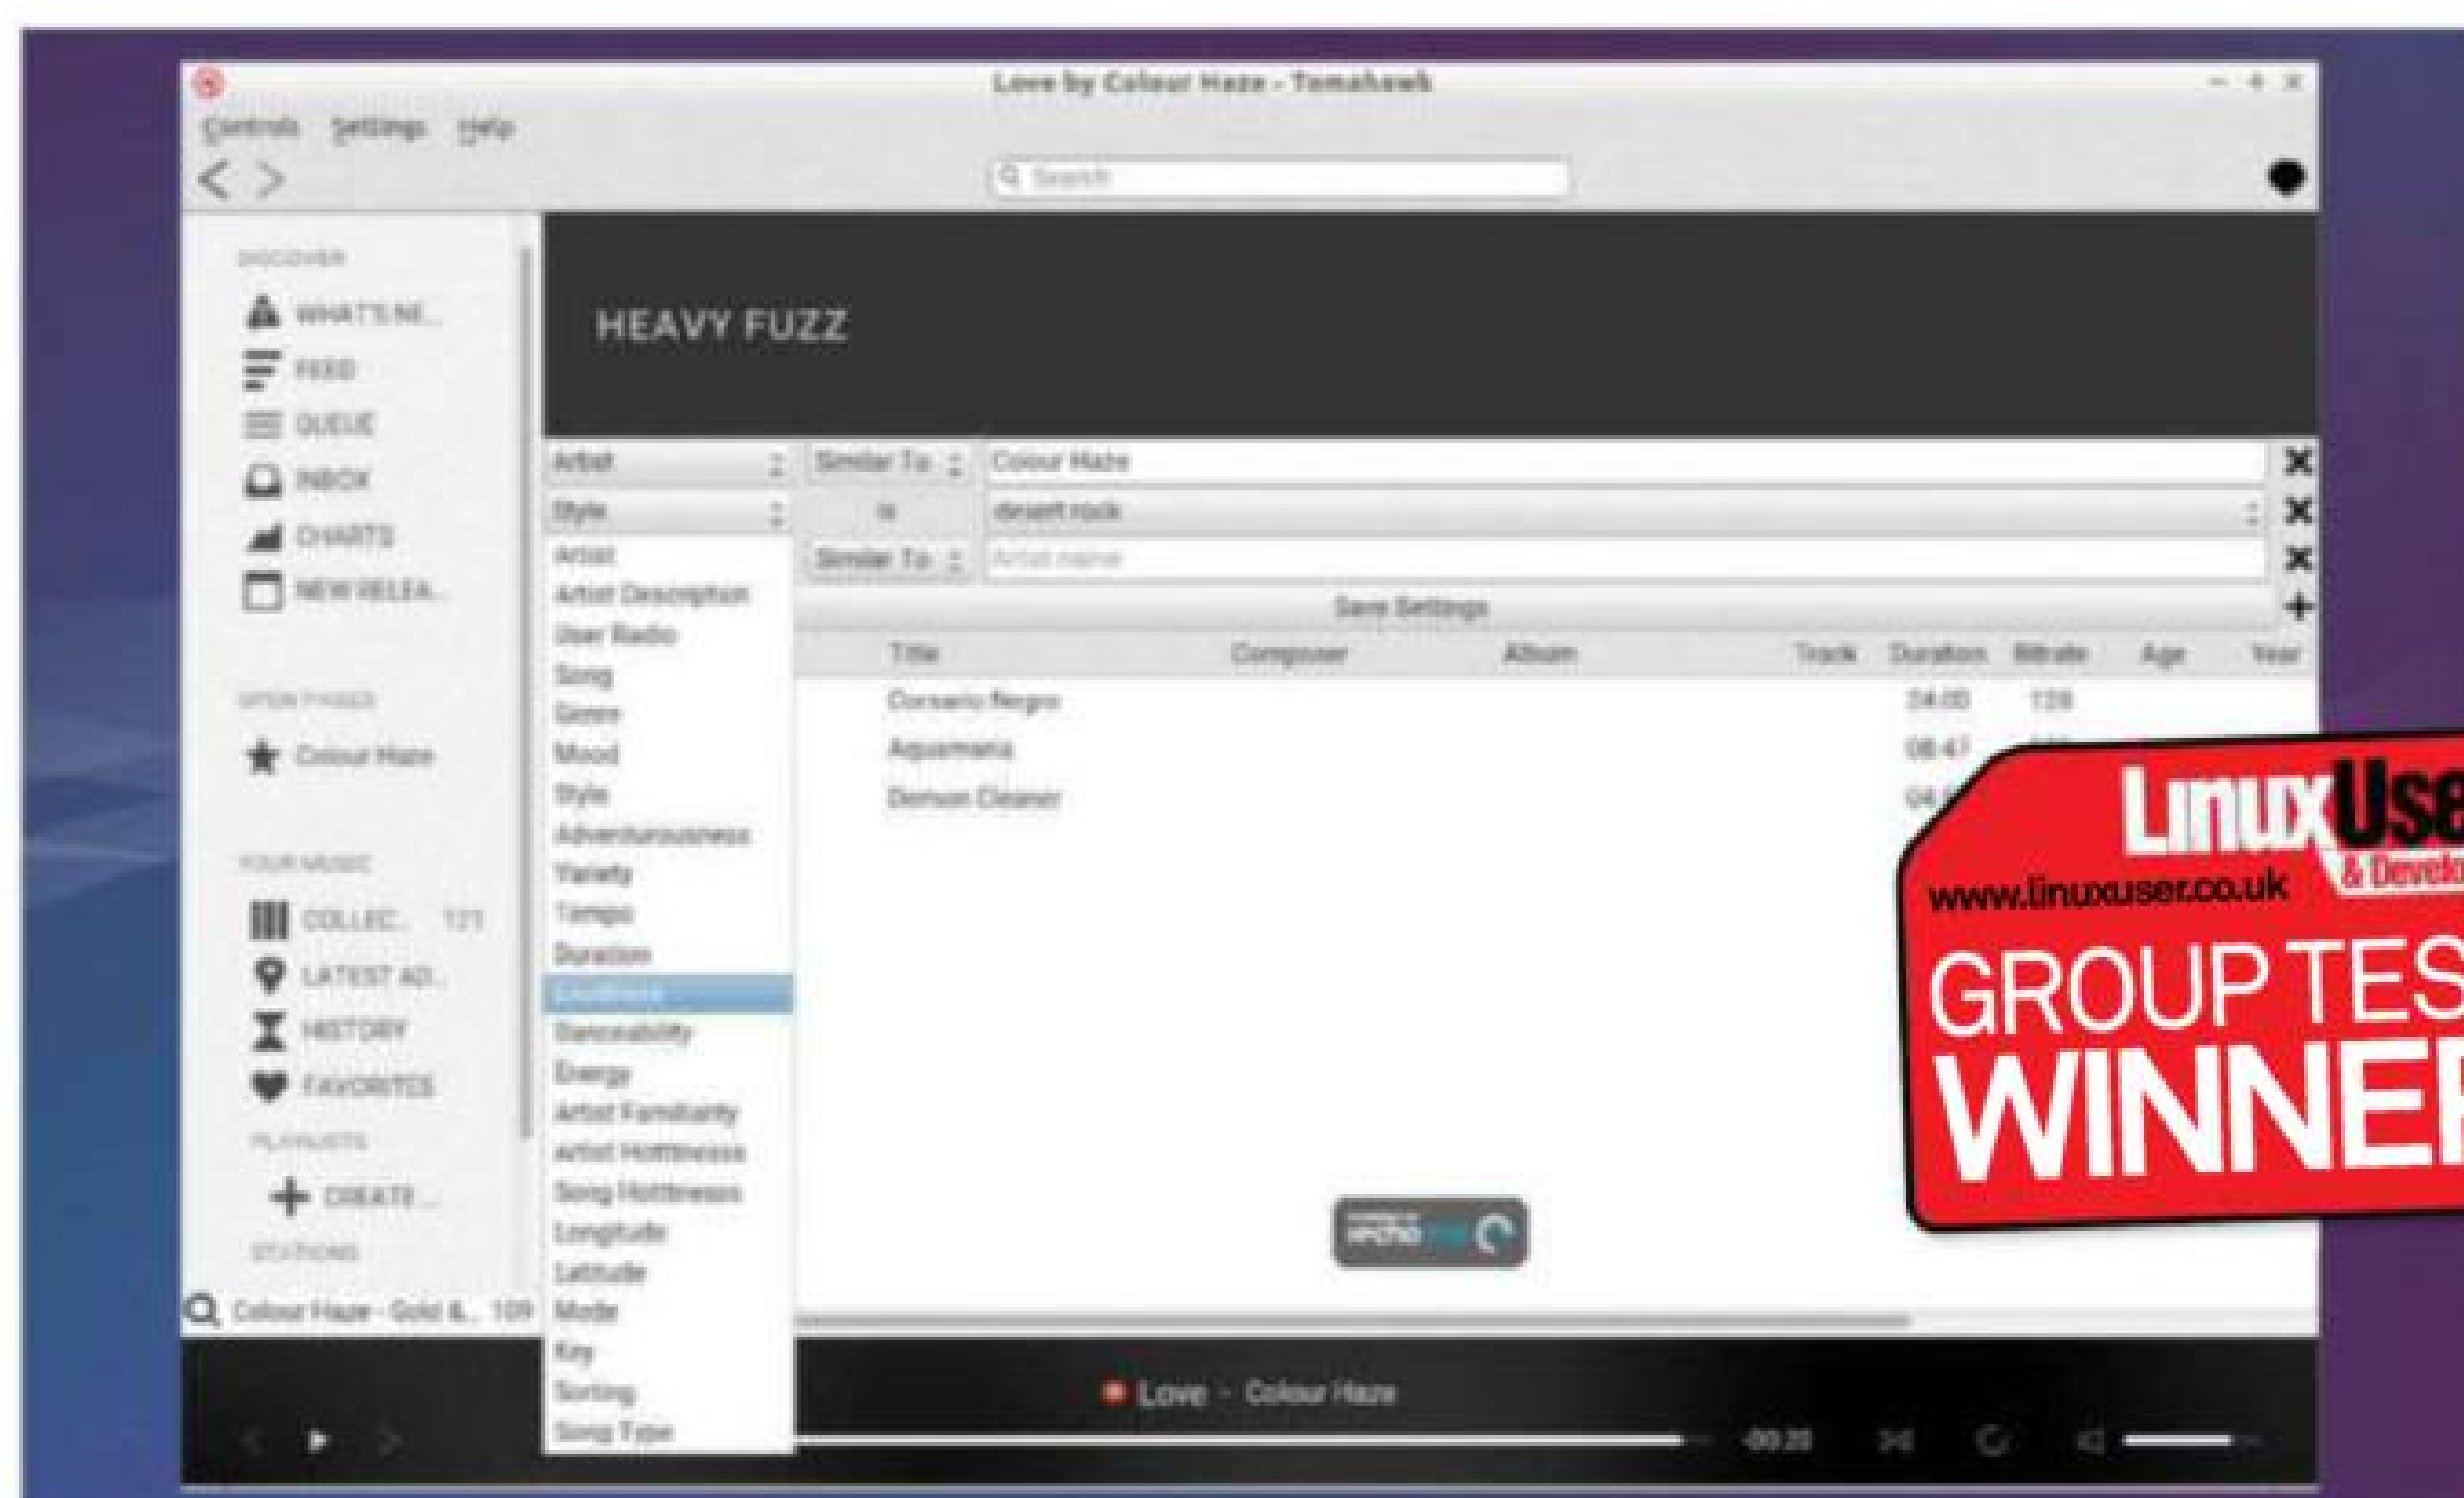
	Nightingale	Tomahawk	CMus	GMusicBrowser
Playback	Basic and with no shortcuts or options to speak of 5	Basic – no shortcuts or options, but good contextual menus 6	Fast and complete control with a separate command per key 8	Fantastic range of options and shortcuts to cover everything 9
Interface	Well-designed and good-looking. Nice to use 8	Clean, modern and appealing. It will really grow on you 9	Simple, but that is its strength. Highly configurable 6	Highly useful, customisable and has a range of skins 7
Management	Can batch edit metadata and generate playlists 6	Basic, so not suited to properly managing large libraries 5	Great filters but sadly no metadata edit option is available 5	Autosuggest tools speed up metadata batch editing 8
Online	Art, lyrics and bios are smartly integrated to aid easy listening 7	Superb. Can search all your favourite streams at the same time 10	Limited services, although custom scripts are supported 4	Limited to Last.fm and Icecast; not the point of GMB 3
Overall	Very balanced. A good all-rounder, perfect for everyday listening 7	Great design and perfect for online, social listeners 8	Focused on speed and customisation, it's the player for CLI pros 6	Powerful, configurable, but lacks online services 7

AND THE WINNER IS...

Tomahawk

Despite the fact that it performed relatively poorly against the competition in terms of playback options and the ability to manage your local library, Tomahawk grew on us throughout this group test. While Nightingale is without question the most balanced application here, scoring decently across all the categories, and GMusicBrowser – despite its somewhat bland and dated appearance on first glance – is arguably the most powerful GUI-based software, we found ourselves appreciating Tomahawk more and more as we used it. CMus was a hugely fun outing but we can't see its appeal extending to more than the most zealous terminal users.

While you can't really tweak and polish away at your locally stored files in the same way that you can with the latter programs, the whole point of Tomahawk is that it embraces the way that music is most often consumed and shared these days – online. If you're paying for Spotify or Google Play Music, you can add those sources to the mix. If you're a Soundcloud or Rdio hopper, add those. Add a whole bunch of sources and you can suddenly find any track in the world, all with the full artist info and gorgeous artwork



■ Custom radio stations can really be finely controlled

instantly displayed in the interface. The custom radio stations are a highlight here, scanning through all your sources to find you songs that fit any detailed criteria, as are the smart playlists. While Tomahawk would need a larger userbase and more notoriety to really sell the community aspect of the software, it is definitely in place with features like the Feed to see your friends' activity and integrated sharing options. Aside from the online side of things, the clincher

was the usability of the program. The modern design and structure makes Tomahawk lovely to use, especially if you're not quite sure what you're in the mood to listen to and fancy having a browse. The downside to all this is, of course, the lack of more complex options and even basic configurability, though fortunately in this case the team has struck the right balance between simplicity and functionality.

■ **Gavin Thomas**

StoreMags.com

RasPi

Digital magazine for Raspberry Pi makers



PROJECTS



INTERVIEWS




TUTORIALS



EXPERT ADVICE



Download the **free** app
on Google Play and iTunes



MINI PC

MinnowBoard Max

Operating System

Debian (Android 4.4 and others available)

Processor

Dual-Core 1.33GHz 64-bit Intel Atom E3825

Graphics Processor

Integrated Intel HD Graphics

Memory

2GB DDR3, 8MB SPI Flash

Dimensions

74mm x 104mm x 22mm

Storage

1x Micro-SD (SDIO supported), 1x SATA 2

Connectivity

1x Micro-HDMI, 1x Gigabit Ethernet, 1x USB 3.0, 1x USB 2.0, 3x UART, 26-pin GPIO (8x general, 2 supporting PWM), I2C, SPI

Price

€178.80 (approx £133)

It's the follow-up to Intel's original MinnowBoard, but can the Max take on its ARM-based rivals?

That Intel is pushing its products into the embedded space is no secret. The company's presentations at the Consumer Electronics Show in Las Vegas this year concentrated on its efforts to turn its chip design know-how into low-power products suitable for everything from smartphones to Internet of Things (IoT) sensor networks. The MinnowBoard project is a big part of that.

The original MinnowBoard launched back in 2013 as a partnership between Intel and CircuitCo, the organisation which designed and developed the ARM-based BeagleBoard development platform. Like the BeagleBoard, the MinnowBoard, reviewed in issue 131, was an open-hardware development platform designed for everyone from engineers testing out new product ideas to hobbyists looking for a new toy.

Adoption of the MinnowBoard was slow, however. Use of a 32-bit UEFI firmware meant that operating system



compatibility was poor, despite Intel producing a Yocto-compliant board support package (BSP). Coupled with relatively poor performance from its single-core 32-bit processor, it's fair to say the MinnowBoard was not a commercial success.

Success is relative, however. When you have as much money as Intel to throw at these projects, failure is simply an opportunity to learn – and learn Intel has. The MinnowBoard Max arrives as the successor to the MinnowBoard, production of which has now ceased. As well as being considerably smaller, a quick glance at the specifications demonstrates that Intel and its CircuitCo partner have been listening intently to feedback from users.

The biggest change is a shift from a 32-bit processor to a more modern 64-bit chip, complete with rewritten 64-bit UEFI firmware. As well as boosting performance, this gives the MinnowBoard Max software compatibility

MinnowBoard Max

Below The 60-pin Lure header anticipates the success of Intel's Lure add-on boards



Pros

An improvement on the original MinnowBoard, there's plenty to recommend the Max variant for embedded x86 development

Cons

Its price is high compared to rival ARM boards, while the Lure concept has yet to gain traction with third-party manufacturers

“The biggest change is a shift from a 32-bit processor to a 64-bit chip”

of which its predecessor could have only dreamed: theoretically, any operating system which can boot on a general-purpose x86 PC can boot on the MinnowBoard Max with a minimum of effort. Microsoft has, naturally, made much of its support for Windows 8.1; of more interest to its target market is support for Linux distributions ranging from Debian and Red Hat all the way through to Google's semi-closed Android platform.

The new processor, Intel's Atom E3825 system-on-chip (SoC), brings improved performance, boosted clock speed to 1.33GHz and two physical cores alongside an increased allowance of 2GB of DDR3 memory. A lower-cost variant of the MinnowBoard Max featuring a single-core 1.46GHz Atom E3815 and 1GB of RAM, is planned to launch in the near future but was not available at the time of writing.

A quick benchmark shows just how different the new Atom processor is to its predecessor: running a single-core test on the SysBench application showed a 95th percentile completion time of 3.22ms – far faster than the 11.49ms of the first MinnowBoard – and the Max boasts a second core to help it streak further away from the original.

As a development board, it's no surprise to find that the MinnowBoard Max has a number of features lacking from general-purpose PCs like Intel's Next Unit of Computing devices, which, development features aside, offer significantly higher specifications at a similar price. A 26-pin GPIO header immediately makes one think of the non-plus variants of the Raspberry Pi, although unlike its undoubted inspiration the MinnowBoard Max features pulse width modulation (PWM) control over two of the eight user-accessible pins for control of servos, dimming of LEDs and similar functions.

The biggest feature, however, is the MinnowBoard's Lure header. Properly known as the High Speed Expansion Connector and located on the underside of the board, this 60-pin header includes additional GPIO pins along with mSATA, mPCIe, USB, I2C and more. Adoption of the standard has been slow, but add-on Lures are beginning to appear with everything from Arduino compatibility to mPCIe headers for future expansion.

Gareth Halfacree

Summary

The MinnowBoard Max is an undeniable improvement on the original MinnowBoard design in both software compatibility and performance. It offers flexibility, but for those who don't need open hardware or GPIO capabilities. It's beaten in value by Intel's own general-purpose Next Unit of Computing (NUC) machines, while its ARM-based rivals typically sell for a fraction of the MinnowBoard's initial launch price.





MINI PC

OpenPi

The base for custom Raspberry Pi-powered hardware utilises the compute module

The OpenPi on first look is a curious device – a nondescript black box with merely an HDMI and a microUSB slot. There's no real indication of what it might be, however cracking it open reveals a custom board connected to a Raspberry Pi compute module. Inside as standard is a wireless dongle and a bluetooth receiver for a mini-wireless keyboard/mouse combo. It seems quite simple and to be fair in this state it is – it's basically just a (fully-functioning) Raspberry Pi.

That's actually the point of it though. With the compute module and the OpenPi board, you have full access to the usual Raspberry Pi power and settings. The selling point of the OpenPi though is that you can then take this board – which is completely open hardware – and modify the plans yourself to make a custom board that fits your needs. Wireless Things thinks of it as an easier way to create an Internet of Things and they've succeeded in creating the platform to do this.

As we've hinted at, using it as a 'normal' Raspberry Pi works. It's fully functional and you can easily access it via SSH thanks to the included Wi-Fi adapter and pre-

installed software. This makes it great off the bat as a Wi-Fi repeater or file server. It comes with standard Raspbian installed that can easily be updated and upgraded as you go to keep up with the latest and greatest software. While there's limited ways to hack the board as it is, the open hardware aspect means the entire PCB specs are there for modification and Wireless Things will even make you custom cases for it as well.

Is it for normal users though? Probably not – it's £300 more expensive than the Pi itself and aimed at a different userbase. Is it for makers? To a degree, yes. The PCB plans being open can help make something very custom for your project, however the standard board doesn't offer much extra to begin with. It's excellent inspiration and a good start for makers, but it's aimed more towards the business and start-up market for creating systems in offices that can benefit from an Internet of Things.

■ Rob Zwetsloot



Operating System

Raspbian

Processor

BCM2835, ARMv6 700 MHz

Memory

512 MB

Dimensions

104mm x 79mm x 33mm

PCB

Wireless Things OpenPi open hardware PCB

Price

£335

Summary

A unique box with a huge amount of potential thanks to an open PCB design and a tiny Compute Module that lets you to perform any Raspberry Pi operation with no limits other than your imagination. It's for a very specific market though, so normal folks need not apply.



Storemags.com

BUILD A BETTER WEB

www.webdesignermag.co.uk



web designer

Available from all good newsagents and supermarkets

ON SALE NOW

jQuery special effects | CSS animations | HTML5 games | Get started with Ruby

DESIGN INSPIRATION

PRACTICAL TIPS

BEHIND THE SCENES

STEP-BY-STEP ADVICE

INDUSTRY OPINION



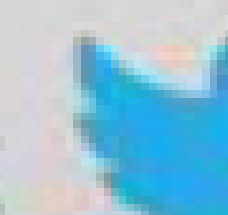
BUY YOUR ISSUE TODAY

Print edition available at www.imagineshop.co.uk

Digital edition available at www.greatdigitalmags.com

Available on the following platforms



 facebook.com/webdesignermag  twitter.com/webdesignermag



FRESH. FAST. FREE.
LIBRE.



The Document Foundation
proudly presents:

LibreOffice

the world's most advanced free software office suite

www.libreoffice.org



UK SERVERS

DEDICATED SERVERS

FROM
JUST

£39


/MONTH
PLUS VAT

SAME DAY
SETUP24/7 UK
FREEPHONE
SUPPORTNO
LONG TERM
CONTRACTS

Speak to a member of our team, or visit our website to find out more
0800 610 1 620 | www.ukservers.com

15% DISCOUNT

WITH PROMO CODE
LINUX1



I spy with my little eye domains ending in...
.actor .buzz **.club** .email **.guru**
.ink .host **.ninja** **.photos** **.press**
.recipes **.pub** **.reviews** .support
.tips **.website** .webcam **.wiki**

Think
Different

Think
Context

Think
Available



www.blacknight.host sales@blacknight.host @blacknight



Domains

.COM just £5.80/pa



Hosting

From £2.99/pm + FREE Domain



Cloud

SSD Server from £5/pm



Servers

Save 50% for 3 months



Datacentre

Rackspace from £35/pm



0800 808 5450



netcetera.co.uk/lud

Questions & answers

Contact us...

Web: www.linuxuser.co.uk

Email: linuxuser@imagine-publishing.co.uk



Your questions answered

Send us your questions and we'll do our best to answer them!

Dial it up

I am somewhat desperate. I've been searching for a solution for days and haven't been able to find one (I can't believe that I'm the only person in need of one for this).

This is what I am looking for: how can I make my smartphone (Android) dial a phone number that I see on the screen of my computer (running Ubuntu). Let's say some number of a signature of an email in Thunderbird or a number I googled in Firefox or the like. There are apps/software to

do exactly that for Windows, but for Linux? I type numbers I see on the screen into my smartphone a few times a day. This is such a waste of concentration and could be done so much faster with the right app!

Is there an app around, or a quick and simple DIY solution at all?

Thanks for any answers!

Henning

To be honest there should be a better method really – there are so many ways to connect to Android over the network that it seems odd and slightly unfair that there isn't just a simple Firefox plugin and corresponding app that you are able to use.

There are solutions though! We like to use Airdroid on Android to transfer files and such at times, but you can also connect to messaging services with it – including calls. This means you can easily copy and paste the number directly into the dialler and it will make the call automatically.

Airdroid requires a fairly active connection though, so something like Pushbullet may work better for you. It makes way for a more passive connection with pop-ups and such built into the browser and you can also call through it as well. Grab it here: play.google.com/store/apps/details?id=com.pushbullet.android.

Hope that one of these helps you out!

FAQ: Compiling software

Q: What are the benefits of compiling your own software?

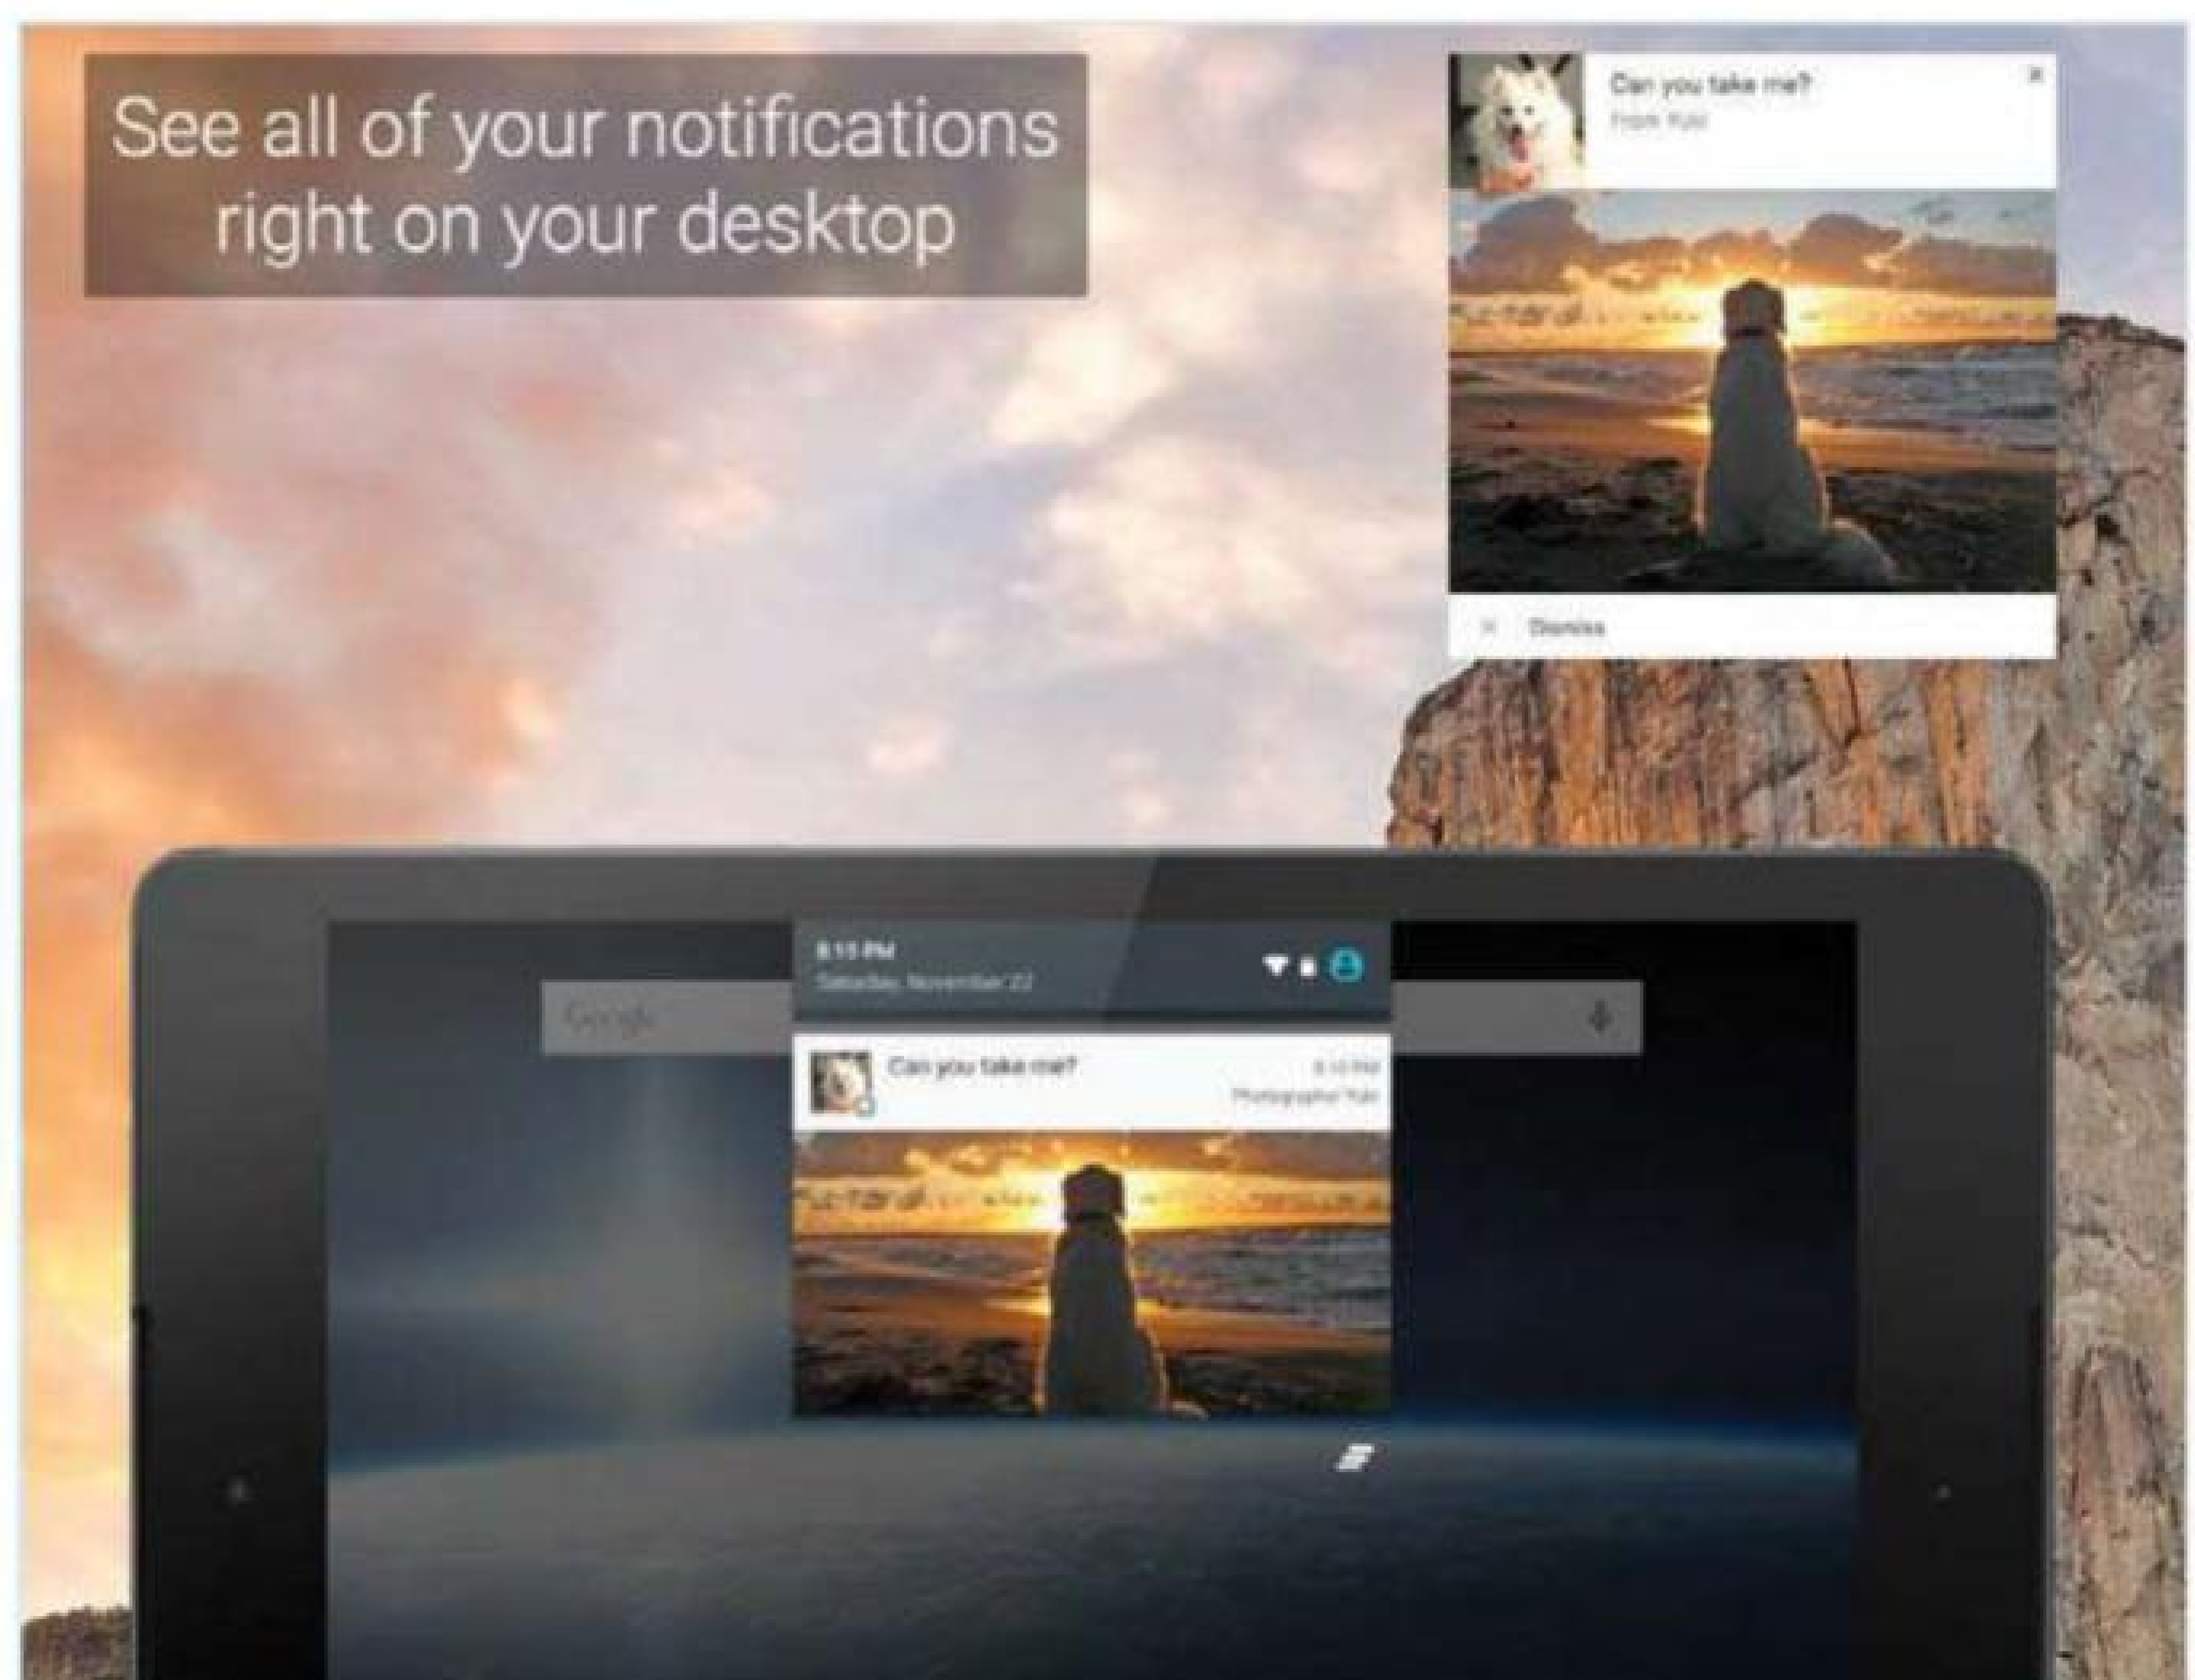
A: The repos may not have the latest version of the software you want and there may not be any binaries to use for a manual install. Compiling also lets you to have the most up-to-date software and you can build it using newer libraries so it works on your system.

Q: What do I need to compile?

A: The source files and any dependencies installed – you'll see if you're missing any when you go to compile. Otherwise, the compiling software is gcc which comes under different packages in some distros – your distro will tell you what you need for it.

Q: How do I compile?

A: Usually it's a case of using the build, make and make-install commands, or a variation or selection of all three. The source files should come with some instructions that will help.



Above Connecting to your phone from a Linux PC shouldn't be hard

storemags.com

Apache doku

I am currently running dokuwiki and rancid on the same host as part of my DR solution for documentation. I would like to create a symbolic link from rancid (/var/lib/rancid/*group*/configs/*current configs*) to my dokuwiki pages folder (/opt/dokuwiki/data/pages/*pages*) so that the currently running config for any switch, router or wireless AP can be looked up from within dokuwiki itself.

Apache is set to follow symbolic links, I gave www-data:www-data permissions to the files using setfacl but dokuwiki will not follow the links. Any guidance would be much appreciated.

Thank you.

Alex H

You basically need to go into the config and change some of the settings for permissions and aliases with something like:

```
<Directory /var/lib/rancid/*whatever*>
Order deny,allow
Allow from all
</Directory>
</VirtualHost>
```

Then you can create the symbolic links to put the files from /var/lib/rancid/*whatever* into the /opt/dokuwiki/data/pages/ directory. Hope this solves your problem!


You must recover

I've made a terrible error and managed to lose all my data on my hard drive! Luckily it's just the NTFS hard drive that had all the data from my home directory stored inside, so I haven't lost my entire system but my backups aren't very recent, and I'd like to recover them if possible! I haven't really touched the disc since then in case it overwrites stuff, but my computer still sees it and so does parted.

Hope you can help me out in my time of need!

Alex Clarke

You're in luck as there's specific Linux software called TestDisk that will help you in this exact situation. Install it via your repos or download it from the website (bit.ly/1crc2Vu) and then run it in the terminal. It will let you probe the disc



Windows XP died, again

I have been using VirtualBox on 64-bit Ubuntu 14.04 for some time and believe that the updates are set to be automatic.

Recently, I went to use one of the virtual machines running my old Windows XP. I had 'Saved' this VM. When I went to start it, the screen showed up just fine during the restore process but then died when the progress bar reached the end.

The message I received was:

```
Unsupported version 11 of data unit
'vga' (instance #0, pass
0xffffffff (VERR_SSM_UNSUPPORTED_DATA_
UNIT_VERSION)).
```

When asking for more details, I got given the following three lines:

```
Result Code: NS_ERROR_FAILURE
(0x80004005)
Component: Console
Interface: IConsole {8ab7c520-2442-4b66-
8d74-4ff1e195d2b6}
```

Apparently I need an updated data unit 'vga' but I don't know whether that's a VirtualBox thing or an Ubuntu thing and, in any case, I have no idea where to find such a thing.

Can anyone tell me how or where to start troubleshooting this?

Many thanks,

Martin G

The problem here seems to be with the saved state you have – we're assuming a snapshot? You don't need to update or change anything to do with the video adapter settings; your best bet is to just delete the snapshot and start is as usual.

In the future, it's probably best just to shut down normally or create a snapshot of a powered-down state. If it's to capture a bit of work you are doing then you can either pause the machine or save it the old-fashioned way – Ctrl+S!

Left DokuWiki is an excellent open source wiki implementation you can host yourself



Questions & answers



MythTV myth

I am having issues exiting the mythtv front-end since around the last update. The back-end seems to be operating well but the fc20 front-end won't exit if I do anything with any video (either watch TV or watch anything that has been recorded).

At first I thought it was the something with fc20 so I put fc21 on and had the same issues. I do not have the problem with the front-end on my other PC. I checked the hard drives and the RAM and those don't seem to have issues. I have also reinstalled fc20 and still had the problem, so I'm pretty sure it is something with the program. I have also turned off AirPlay. I saw that was an issue last year but it still persists. Any input would be much appreciated because it is literally unusable right now. Once it freezes I have to hard reset the PC to get it off because I can't kill the process, and it even hangs when I just use reboot. I found nothing in the logs.

Carson Taylor

If it's hanging but the entire system is not down, then this sounds like it could be one of many things. First of all though, dial in remotely and see if it's the whole system or just parts of the front-end.

After that check the device drivers, especially the graphics drivers. Maybe use the official drivers if you've just been using nouveau so far.

Finally, make sure that your hard drives are all in working order. If it's trying to record to a dodgy hard drive then it may just hang the entire system while it keeps trying and buffering the video.


Hope one of these helps you fix it.



Above MythTV is a powerful all-in-one media solution, but it can go wrong



Above MilkyTracker works just like FastTracker, and supports ProTracker compatibility



and look for broken partitions, data or anything else between. Refer to the documents to it and compare it to your exact situation to get the most out of it.

Elephant in the room

I'm in the process of installing PostgreSQL. Per the instructions on Ubuntu's site, I ran:

```
sudo apt-get install postgresql
postgresql-contrib
sudo apt-get install pgadmin3
sudo -u postgres psql postgres
```

...and encountered the following error:

could not change directory to "/home/bev":
Permission denied

I know enough about Linux to pause before attempting to change file/directory permissions for something like this. What should be my next step?

Bev

Right, the issue may be that PostgreSQL can't access the folder you're in for whatever reason. Permission issues can always be annoying, unfortunately. We suggest the best thing you can do is move to another folder, say tmp, with:

```
$ cd /tmp
```

It should then build just fine and also should work well without you having to do anything with changing location either.

Another tracker

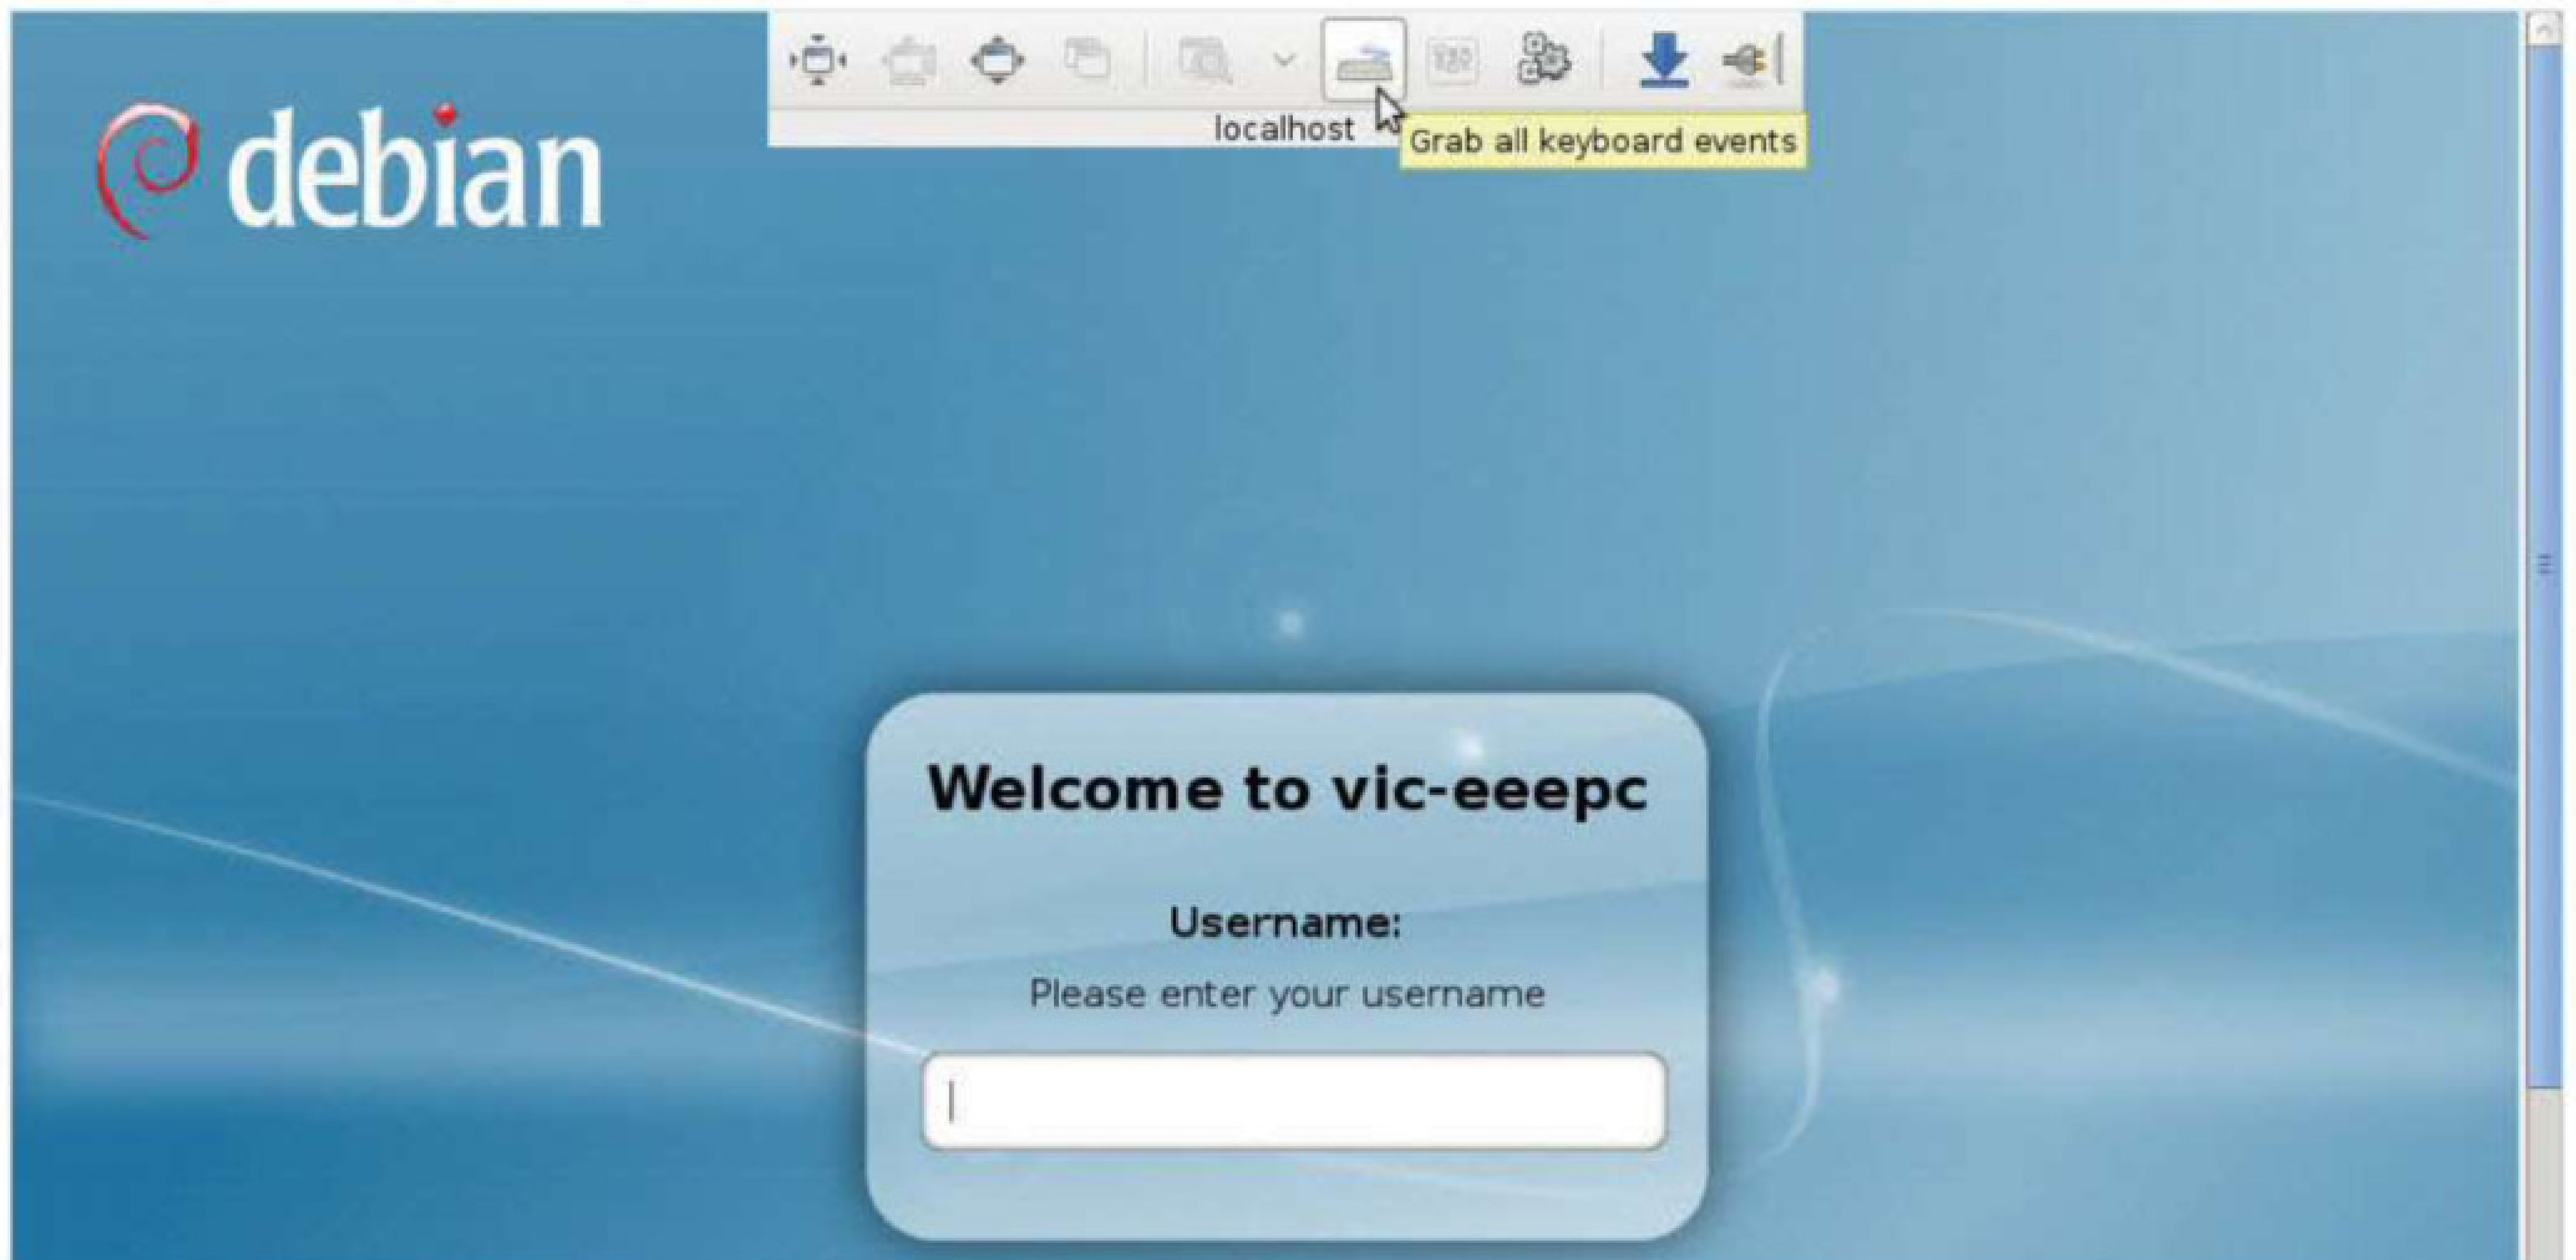
I used to love using ProTracker on Amiga – it was a great little piece of software for tracker creation and I've been wondering recently if there are any Linux alternatives to this tracker? I haven't been able to properly find any – do you know of any that are around at all?

Thanks a lot.

Dan

Thanks for writing in Dan. MilkyTracker is a popular open source take on apps like ProTracker, although it's a bit more based on FastTracker II, the DOS program. However, Fast and ProTracker work very similarly, so it sounds like it would be your best bet for getting back into tracking.





Above When connecting remotely, some settings are different for displays

VNC issues

Sometimes I need to do work in the terminal on my system. However, when I try and open the terminal emulator, it opens briefly and then automatically closes.

The system in question is CentOS, which I have to access remotely via VNC. It's using XFCE as a desktop environment too. It works fine if I'm logged in as root, so I can only assume it's a problem with the user type. I don't want to make some of the normal users have full root privileges and I'd rather not log in as root every time. Have you heard of anything like this before? Thanks!

Thomas Reid


Try this, Thomas: open the terminal logged in as root, then change to another user with:

```
$ su -l <user>
```

This enables you to open another shell to log in with and then try and open the terminal. Any errors will show up in the root terminal. It's likely that you can't load into the shell because the user won't have access to connect to the X Server, as you're using it via VNC. To solve this, add the relevant users to the list of those allowed to make this connection. For xhost, type:

```
# xhost local:<user>
```

...as root. Give the system a reboot and it should now be working fine.



For everything you need
to build your own
amazing robot



ROBOTBITS.com
BUILD SOMETHING AMAZING

Hosting listings


Sponsorship opportunity

Bring attention to your brand by sponsoring this section. Contact Alex Carnegie on +44(0)1202 586430

Got a deal you think we should list?

Whether you're a hosting firm or a happy customer who wants a favourite provider listed, drop us a line with the details!
linuxuser@imagine-publishing.co.uk

Dedicated server listings

NAME AND URL	PACKAGE	PHONE NUMBER	COST PER MONTH	MINIMUM CONTRACT TERM	CPU CORES / SPEED	DISK SPACE	1GBPS INTERNET CONNECTION	HARDWARE RAID	REMOTE POWER REBOOT	PERMANENT KVM	UPTIME GUARANTEE	NETWORK BACKUP STORAGE	PRIVATE SUBNET	24/7 PHONE SUPPORT
 Netcetera www.netcetera.co.uk/linux	2200DC	0800 8085450	£25	1 month	Dual Core 2.2GHz	160GB	N/A	Raid 1	✓	✓	✓	✓	✓	✓
	3000DC	0800 8085450	£40	1 month	Dual Core 3GHz	2 x 250GB	N/A	Raid 1	✓	✓	✓	✓	✓	✓
	2660QC	0800 8085450	£65	1 month	Intel 2.66GHz Quad Core Xeon Processor	2 x 500GB	N/A	Raid 1	✓	✓	✓	✓	✓	✓
	Developer	0800 8085450	£2.99	1 month	N/A	1GB	N/A	✓	✓	✓	✓	✓	✓	✓
	One	0800 8085450	£9.99	1 month	N/A	5GB	N/A	✓	✓	✓	✓	✓	✓	✓
	Reseller	0800 8085450	£24.99	1 month	N/A	Unlimited	N/A	✓	✓	✓	✓	✓	✓	✓
Bravo14 (http://bravo14.co.uk)	Starter Linux	N/A	£20	N/A	N/A	2,000MB	N/A	✓	✓	✓	✓	✗	✓	✓
Bravo14 (http://bravo14.co.uk)	Starter Windows	N/A	£20	N/A	N/A	2,000MB	N/A	✓	✓	✓	✓	✗	✓	✓
Bravo14 (http://bravo14.co.uk)	Business Linux	N/A	£45	N/A	N/A	4,000MB	N/A	✓	✓	✓	✓	✗	✓	✓
Bravo14 (http://bravo14.co.uk)	Business Windows	N/A	£45	N/A	N/A	4,000MB	N/A	✓	✓	✓	✓	✗	✓	✓
Bravo14 (http://bravo14.co.uk)	Ultimate Linux	N/A	£60	N/A	N/A	Unlimited	N/A	✓	✓	✓	✓	✗	✓	✓
Bravo14 (http://bravo14.co.uk)	Ultimate Windows	N/A	£60	N/A	N/A	Unlimited	N/A	✓	✓	✓	✓	✗	✓	✓
catalyst2 (www.catalyst2.com)	Bronze Managed Dedicated Server	0800 107 79 79	£199	1 month	1x 2.4GHz vCPU	50GB	✓	✓	✓	✓	99.90%	✓	✓	✓
catalyst2 (www.catalyst2.com)	Silver Managed Dedicated Server	0800 107 79 79	£299	1 month	1x 2.4GHz vCPU	80GB	✓	✓	✓	✓	99.90%	✓	✓	✓
catalyst2 (www.catalyst2.com)	Gold Managed Dedicated Server	0800 107 79 79	£399	1 month	2x 2.4GHz vCPU	150GB	✓	✓	✓	✓	99.90%	✓	✓	✓
123-Reg (www.123-reg.co.uk)	Dell PowerEdge R200 (Ubuntu Linux)	0871 230 9525	£69.99	12 months	4x 2.13GHz	2x 160GB	10Mbit	✓	✓	✗	99.99%	0	✗	✓
123-Reg (www.123-reg.co.uk)	Dell PowerEdge R200 (Windows Web Edition)	0871 230 9525	£79.99	12 months	4x 2.13GHz	2x 160GB	10Mbit	✓	✓	✗	99.99%	0	✗	✓
Daily (www.daily.co.uk)	Linux VPS Pro	0845 466 2100	£29.99	1 month	2.27 Intel Quad Core	60GB	100Mbps	✓	✓	✗	✗*	✓ - full backup	✗	✗**
Daily (www.daily.co.uk)	Linux VPS Max	0845 466 2100	£59.99	1 month	2.27 Intel Quad Core	100GB	100Mbps	✓	✓	✗	✗*	✓ - full backup	✗	✗**
Daily (www.daily.co.uk)	Windows VPS Pro	0845 466 2100	£34.99	1 month	2.27 Intel Quad Core	60GB	100Mbps	✓	✓	✗	✗*	✓ - full backup	✗	✗**
Daily (www.daily.co.uk)	Windows VPS Max	0845 466 2100	£64.99	1 month	2.27 Intel Quad Core	100GB	100Mbps	✓	✓	✗	✗*	✓ - full backup	✗	✗**
Daily (www.daily.co.uk)	VPS Pro Hyper-V	0845 466 2100	£44.99	1 month	2.27 Intel Quad Core	60GB	100Mbps	✓	✓	✗	✗*	✓ - 1GB	✗	✗**
Daily (www.daily.co.uk)	VPS Max Hyper-V	0845 466 2100	£74.99	1 month	2.27 Intel Quad Core	100GB	100Mbps	✓	✓	✗	✗*	✓ - 1GB	✗	✗**
Daily (www.daily.co.uk)	VPS Ultra Hyper-V	0845 466 2100	£139.99	1 month	2.27 Intel Quad Core	200GB	100Mbps	✓	✓	✗	✗*	✓ - 1GB	✗	✗**
Heart Internet (www.heartinternet.co.uk/dedicated-servers)	Linux Dual Core	0845 644 7750	£79.99	12 months	Dual Core Xeon 2.33GHz	160GB	✓	✓	✓	✗	99.99%	✓	✗	24/7 Ticket support
Heart Internet (www.heartinternet.co.uk/dedicated-servers)	Windows Dual Core	0845 644 7750	£89.99	12 months	Dual Core Xeon 2.33GHz	160GB	✓	✓	✓	✗	99.99%	✓	✗	24/7 Ticket support
Heart Internet (www.heartinternet.co.uk/dedicated-servers)	Linux Quad Core	0845 644 7750	£129.99	12 months	Quad Core Xeon 2.5GHz	250GB	✓	✓	✓	✗	99.99%	✓	✗	24/7 Ticket support
Webfusion (www.webfusion.co.uk)	Dell PowerEdge R210	0845 130 1602	£79.99	12 months	2x 3.06GHz	250GB	Up to 100Mbit	✗	✓	✗	99.99%	Free	0	✓
Webfusion (www.webfusion.co.uk)	Dell PowerEdge R210	0845 130 1602	£119.99	12 months	4x 2.66GHz	2x 250GB	Up to 100Mbit	✓	✓	✗	99.99%	Free	0	✓
Webfusion (www.webfusion.co.uk)	Dell PowerEdge R210	0845 130 1602	£149.99	12 months	4x 2.66GHz	2x 500GB	Up to 100Mbit	✓	✓	✗	99.99%	Free	0	✓

0 = Option

GET YOUR LISTING **HIGHLIGHTED!** CONTACT DANIEL

alex.carnegie@imagine-publishing.co.uk

+44(0)1202 586430

Dedicated and Shared server listings

NAME AND URL	PACKAGE	PHONE NUMBER	COST	WEB SPACE	MONTHLY BANDWIDTH	POP3 ACCOUNTS	DATABASE SUPPORT	SHOPPING CART	VIRUS FILTER	FIREWALL	PHONE SUPPORT	EMAIL SUPPORT	WEB CONTROL PANEL	SERVICE LEVEL AGREEMENT
Blacknight (www.blacknight.com)	Minimus	+44 (0)845 5280242	€49.95	10GB	150GB	1,500	✓	✓	✓	✓	✓	✓	✓	✗
Blacknight (www.blacknight.com)	Medius	+44 (0)845 5280242	€89.95	20GB	300GB	5,000	✓	✓	✓	✓	✓	✓	✓	✗
Blacknight (www.blacknight.com)	Maximus	+44 (0)845 5280242	€149.95	30GB	600GB	Unlimited	✓	✓	✓	✓	✓	✓	✓	✗
Digital Gibbon Ltd (www.digitalgibbon.com)	Cheeky Chimp	N/A	Free	500MB	Unlimited	5	✓	✗	✓	✓	✗	✓	✓	✗
Digital Gibbon Ltd (www.digitalgibbon.com)	Digital Gibbon	N/A	£12	5GB	Unlimited	10	✓	✗	✓	✓	✗	✓	✓	✗
Digital Gibbon Ltd (www.digitalgibbon.com)	Silverback	N/A	£24	Unlimited	Unlimited	Unlimited	✓	✓	✓	✓	✗	✓	✓	✗
Digital Gibbon Ltd (www.digitalgibbon.com)	WordPress hosting	N/A	£12	5GB	Unlimited	10	✓	✗	✓	✓	✗	✓	✓	✗
eHosting (www.ehosting.com)	Starter	0844 999 4100	£23.88	1GB	25GB	10	✗	✗	✗	✗	✗	✓	✓	✓
eHosting (www.ehosting.com)	Personal	0844 999 4100	£59.88	2.5GB	Unlimited	50	✓	✗	✗	✗	✗	✓	✓	✓
eHosting (www.ehosting.com)	Expert	0844 999 4100	£95.88	5GB	Unlimited	250	✓	✗	✗	✗	✗	✓	✓	✓
eHosting (www.ehosting.com)	Virtual	0844 999 4100	£227.88	50GB	Unlimited	Unlimited	✓	✗	✗	✗	✓	✓	✓	✓
Equiphase (www.equiphase.net)	Bronze	0121 314 4865	£30	200MB	2GB	10	✓	✓	✗	✓	✗	✓	✓	✓
Equiphase (www.equiphase.net)	Silver	0121 314 4865	£42	400MB	5GB	20	✓	✓	✗	✓	✗	✓	✓	✓
Equiphase (www.equiphase.net)	Gold	0121 314 4865	£72	800MB	10GB	100	✓	✓	✗	✓	✗	✓	✓	✓
Equiphase (www.equiphase.net)	Platinum	0121 314 4865	£114	1,200MB	40GB	200	✓	✓	✓	✓	✓	✓	✓	✓
Eurofasthost.com (www.eurofasthost.com)	Email Only	02380 249 823	£40	1GB	2GB	10	✗	✗	✓	✓	✓	✓	✓	✓
Eurofasthost.com (www.eurofasthost.com)	Essential	02380 249 823	£75	2GB	5GB	10	✗	✗	✓	✓	✓	✓	✓	✓
Eurofasthost.com (www.eurofasthost.com)	Superior	02380 249 823	£140	5GB	10GB	25	✓	✓	✓	✓	✓	✓	✓	✓
Eurofasthost.com (www.eurofasthost.com)	Premium	02380 249 823	£250	10GB	25GB	100	✓	✓	✓	✓	✓	✓	✓	✓
Evohosting (www.evohosting.co.uk)	Starter	N/A	£29.99	500MB	1GB	3	✓	✓	✓	✓	✗	✓	✓	✓
Evohosting (www.evohosting.co.uk)	Home	N/A	£54.99	2.5GB	30GB	50	✓	✓	✓	✓	✗	✓	✓	✓
Evohosting (www.evohosting.co.uk)	Business	N/A	£79.99	6.5GB	Unlimited	Unlimited	✓	✓	✓	✓	✗	✓	✓	✓
Evohosting (www.evohosting.co.uk)	eCommerce	N/A	£159.99	30GB	Unlimited	Unlimited	✓	✓	✓	✓	✗	✓	✓	✓
Giacom (www.giacom.com)	Business Pro	0800 542 7500	£199	100MB	2GB	100	✓	✓	✓	✓	✓	✓	✓	✓
Heart Internet (www.heartinternet.co.uk)	Starter Professional	0845 644 7750	£29.80	2.5GB	10GB	1,000	✗	✗	✗	✓	✗	✓	✓	✓
Heart Internet (www.heartinternet.co.uk)	Home Professional	0845 644 7750	£89.99	10GB	50GB	10,000	✓	✓	✗	✓	✗	✓	✓	✓
Heart Internet (www.heartinternet.co.uk)	Business Professional	0845 644 7750	£129.99	Unlimited	Unlimited	Unlimited	✓	✓	✗	✓	✗	✓	✓	✓
Heart Internet (www.heartinternet.co.uk)	Reseller Professional	0845 644 7750	£299.99	Unlimited	Unlimited	Unlimited	✓	✓	✗	✓	✗	✓	✓	✓
Hostway (www.hostway.co.uk)	Silver	0808 180 1880	£79.50	150MB	3GB	5	✗	0	✓	✓	✗	✓	✓	✗
NameHog (www.namehog.net)	Email Only	0845 612 0330	£11.75	25MB	1GB	5	✗	✗	✓	✓	✓	✓	✓	✓
NameHog (www.namehog.net)	Standard Package	0845 612 0330	£35.25	100MB	4.5GB	10	✗	✗	✓	✓	✓	✓	✓	✓
NameHog (www.namehog.net)	Professional Package	0845 612 0330	£58.75	250MB	8GB	25	✓	✓	✓	✓	✓	✓	✓	✓
NameHog (www.namehog.net)	Expert Package	0845 612 0330	£105.75	500MB	15GB	75	✓	✓	✓	✓	✓	✓	✓	✓
Skymarket (www.skymarket.co.uk)	Standard 1	0800 321 7788	£49	10MB	2GB	1	✓	✗	✓	✓	✓	✓	✓	✓
Skymarket (www.skymarket.co.uk)	Standard 2	0800 321 7788	£69	20MB	2GB	1	✓	✗	✓	✓	✓	✓	✓	✓
Skymarket (www.skymarket.co.uk)	Premium 1	0800 321 7788	£99	25MB	2GB	1	✓	✗	✓	✓	✓	✓	✓	✓
ServWise (https://www.servwise.com)	Personal	0800 520 0716	£25.20	1GB	10GB	Unlimited	✓	✓	✓	✓	✓	✓	✓	✓
ServWise (https://www.servwise.com)	Business	0800 520 0716	£50.40	2GB	20GB	Unlimited	✓	✓	✓	✓	✓	✓	✓	✓
ServWise (https://www.servwise.com)	Reseller	0800 520 0716	£126	4GB	40GB	Unlimited	✓	✓	✓	✓	✓	✓	✓	✓
WebFusion (www.webfusion.co.uk)	Fusion Professional	0845 130 1602	£107.40	5GB	50GB	1,000	✓	✗	✓	✓	✓	✓	✓	✗
WebFusion (www.webfusion.co.uk)	Fusion Business	0845 130 1602	£179.40	10GB	150GB	1,500	✓	✗	✓	✓	✓	✓	✓	✗
WebFusion (www.webfusion.co.uk)	Fusion Developer	0845 130 1602	£227.40	20GB	300GB	5,000	✓	✗	✓	✓	✓	✓	✓	✗
WebFusion (www.webfusion.co.uk)	Fusion Reseller	0845 130 1602	£329.99	Unlimited	Unlimited	Unlimited	✓	✗	✓	✓	✓	✓	✓	✗

0 = Option

YOUR VIEW

Linux User Letters

Your opinions about the magazine, Linux and open source

FCC change

Hi there, I'm a bit confused by the net neutrality stories and hoping you can explain what's going on. I heard that the chairman of the FCC changed his mind recently and is in favour of supporting net neutrality with Title II – but what is that?

Niall Yealland


Hi Niall. Yes, that's right – Tom Wheeler, chairman of the FCC (Federal Communications Commission), recently wrote an opinion editorial for *Wired* where he declared that he's put new proposals forward to the board on how to “preserve the Internet as an open platform for innovation and free expression”. And the crux of it is that he wants to modernise Title II and then use it to “implement and enforce open internet protections”.

Basically, Title II is the second section of The Communications Act of 1934 and it deals with common carriers, which are people or companies that transport people or goods – this covers telecommunications providers. Until recently, the FCC's own Telecommunications Act of 1996 has been used to regulate telecommunications in the United States, but Wheeler wants to use an updated version of Title II instead with a view to banning “paid prioritisation, and the blocking and throttling of lawful content and services” (the much-derided Internet ‘fast lanes’, for example). By the time you read this, the FCC should have voted on the proposal at its open meeting on February 26th.

No GNOME love

Can you tell me what the fuss is with GNOME 3? I've only just started using Linux and I've been trying to read around and test virtual boxes to learn more about this world. I get the impression that everyone hates GNOME 3 but I seem to see it in lots of the distros I look at on distrowatch.com. Why is it everywhere if no one likes it?

Max Aldridge



Left Big names like Fedora and openSUSE make use of GNOME 3

The thing with GNOME 3 is quite complicated. We'll try and give you the gist of it but popularity issues like this are never simple. One major strand of it is that it was such a radical change. As you'll know from your research, GNOME 3 isn't the familiar desktop metaphor that you see with something like Linux Mint, with its application launcher off to the left, programs group together as ‘Activities’ in the top bar and then all kinds of oddities like the way Alt-Tabbing works. Compared to GNOME 2, which was more traditional, GNOME 3 was essentially a brand new desktop – not necessarily a bad thing, but when you've been using GNOME 2 for years and know your way around it perfectly, GNOME 3 comes as a bit of a surprise.

The second strand of it was that people started to give feedback to the developers and then came to believe that these developers weren't interested in listening, with a typical recommendation being to use a shell extension to work around a fundamental feature that the user didn't like. Generally speaking, it fed into the growing unpopularity of GNOME 3. Hope that helps!

Android Magazine



With more and more Android features being merged back into the Linux kernel, it's becoming much easier to develop for one of the most popular mobile

operating systems around. With over 25 billion app downloads and over 100 million Android devices worldwide, there's a wide audience of folks ready and willing to consume apps. For a more Android-driven editorial, you can look to our sister mag **Android Magazine**, the only publication dedicated to the platform. Along with news and reviews to keep you up-to-date on everything Android, there are also tutorials and advice on developing and hacking your hardware. Find out more at www.littlegreenrobot.co.uk.

storemags

Look for
issue 151
on 9 April
Want it sooner?
**Subscribe
today!**

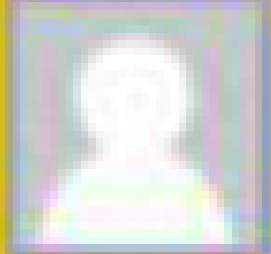
Linux User
THE MAGAZINE FOR
THE GNU GENERATION  **& Developer**

THREAD BARE

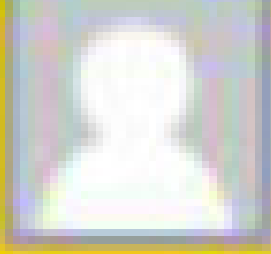
Systemd split

The controversial init replacement proved as divisive on the linuxuser.co.uk Disqus threads as it has done across Linuxland – read the rest at bit.ly/1zVygqs.

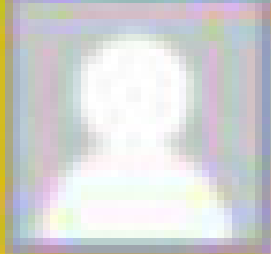
Robert Pogson said:

 The thing that bugs me about systemd is that the thing is still being written and is not debugged, yet distros like Debian are making it the default init system. That's just not right. The particular issue which causes me not to use systemd on my main machine is that I have a bunch of processes running on it and systemd wants all of them to have priority over X, my interface, what I want to be up as soon as possible. I kid you not. Apache, DHCP, NFS, MySQL, PostgreSQL, and some others all have to finish starting up before systemd will even begin to start X. [...] My system that used to boot in ~40s now takes 2min. I reboot a lot because I build and install kernels with it. Systemd is wasting my time.

Daniel Sandman said:


 My 10-year old machine went from a minute starting up to about 10 seconds. The thing with systemd is that it starts services in parallel instead of serial. So it will not do as the old init that had to wait for services before it could start the next one. Systemd can theoretically start them all at once. Another good thing is that it work with dependencies... Some stuff like DHCP needs to have a working Internet connection or it will fail. What systemd does here is to make sure you have that connection before it starts DHCP. If you have an issue with your Internet connection you can easily troubleshoot with 'systemctl status' and similar. I use the inbuilt capability of systemd... the systemd-networkd. It improved the connection time from 10 seconds to 0.250 seconds.

lucius-cornelius said:


 "They" tried, for years, to destroy Linux. "Only hackers use it", "only hippies use it", "only communists or terrorists use it", "we own patents for most of it" and each one failed. Now they're attacking it from within and it's worked beautifully. One community torn asunder over systemd. Most distros now

firmly in the palm of Red Hat and thus under their control. The modularity and control that distinguished Linux from other OS's, now mostly gone and by the time Poettering has finished, it will all be gone. And then it will be too late. Thankfully there are still some distros holding out – Slackware, Crux, Pisi, Manjaro OpenRC and Devuan if it gets off the ground. Long may they continue to resist. But I don't hold out much hope in the long run. This is Corporate takeover 101 and so few even see what's happening that the chances of stopping it are next to zero. Sad.

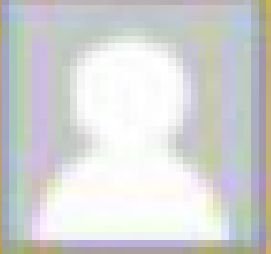
average linux user said:

 Evil corporations are at it again. Is that what you are saying? Well, I am glad you are sad. Red Hat does not control anything, they just develop stuff. The other Linux distros are FREE to use Red Hat code or not use it.

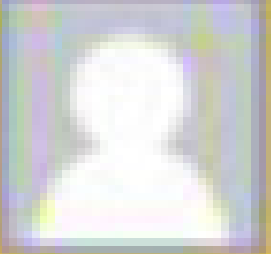
IJK said:

 What worries me about systemd is the dependencies that it will introduce. I don't care if Gnome depends on it, for Gnome is a pile of manure and I won't use it anyway (I still think it is a gross mistake to make it dependent on systemd though.) However, if things like Gimp are dependent on systemd my blood starts to boil. If that dependency exists, systemd has clearly exceeded its role and is on its way of becoming a lock-in module – my way or the highway. If that comes to pass that will be the end of Linux for me.

TheRealTachyon said:

 The thing that is most disturbing about systemd is the organised and active group of defenders and trolls that scour the Internet for articles and discussions on systemd and actively go after detractors. I haven't the stomach to read all the comments on this article, but I guarantee you'll find what I'm talking about in the comments.

Glenn Holmer said:

 The complaint is that systemd is monolithic, and monolithic control systems with one point of failure have never been the Unix, or Linux, way of doing things: bit.ly/1AlgwZc.

Imagine Publishing Ltd
Richmond House, 33 Richmond Hill
Bournemouth, Dorset, BH2 6EZ
☎ +44 (0) 1202 586200
Web: www.imagine-publishing.co.uk
www.linuxuser.co.uk
www.greatdigitalmags.com

Magazine team

Deputy Editor Gavin Thomas

gavin.thomas@imagine-publishing.co.uk
☎ 01202 586257

Production Editor Siobhan Maguire

Designer Sam Ribbits

Senior Staff Writer Rob Zwetsloot

Photographer James Sheppard

Senior Art Editor Will Shum

Editor in Chief Dan Hutchinson

Publishing Director Aaron Asadi

Head of Design Ross Andrews

Contributors

Joey Bernard, Christian Cawley, Liam Fraser, Gareth Halfacree, Tam Hanna, Richard Hillesley, Jon Masters, Aaron Shaw, Richard Smedley, Nitish Tiwari, Mihalis Tsoukalos.

Advertising

Digital or printed media packs are available on request.

Head of Sales Hang Deretz

☎ 01202 586442

hang.deretz@imagine-publishing.co.uk

Advertising Manager Alex Carnegie

☎ 01202 586430

alex.carnegie@imagine-publishing.co.uk

FileSilo.co.uk

Assets and resource files for this magazine can now be found on this website.

Support filesilohelp@imagine-publishing.co.uk

International

Linux User & Developer is available for licensing.

Head of International Licensing Cathy Blackman

☎ +44 (0) 1202 586401

licensing@imagine-publishing.co.uk

Subscriptions

For all subscriptions enquiries

☎ 0844 249 0282 (UK)

☎ +44 (0) 1795 418661 (Overseas)

Email: LUD@servicehelpline.co.uk

6 issue subscription (UK) – £25.15

13 issue subscription (Europe) – £70 (ROW) – £80

Circulation

Head of Circulation Darren Pearce

☎ 01202 586200

Production

Production Director Jane Hawkins

☎ 01202 586200

Founder

Group Managing Director Damian Butt

Printing & Distribution

Printed by William Gibbons, 26 Planetary Road,
Willenhall, West Midlands, WV13 3XT

Distributed in the UK, Eire & the Rest of the World by:
Marketforce, Blue Fin Building, 110 Southwark Street
London, SE1 0SU
☎ 0203 148 3300
www.marketforce.co.uk

Distributed in Australia by:

Network Services (a division of Bauer Media Group)

Level 21 Civic Tower, 66-68 Goulburn Street

Sydney, New South Wales 2000, Australia

☎ +61 2 8667 5288

Disclaimer

The publisher cannot accept responsibility for any unsolicited material lost or damaged in the post. All text and layout is the copyright of Imagine Publishing Ltd. Nothing in this magazine may be reproduced in whole or part without the written permission of the publisher. All copyrights are recognised and used specifically for the purpose of criticism and review. Although the magazine has endeavoured to ensure all information is correct at time of print, prices and availability may change. This magazine is fully independent and not affiliated in any way with the companies mentioned herein.

If you submit material to Imagine Publishing via post, email, social network or any other means, you automatically grant Imagine Publishing an irrevocable, perpetual, royalty-free license to use the material across its entire portfolio, in print, online and digital, and to deliver the material to existing and future clients, including but not limited to international licensees for reproduction in international, licensed editions of Imagine products. Any material you submit is sent at your risk and, although every care is taken, neither Imagine Publishing nor its employees, agents or subcontractors shall be liable for the loss or damage.


**IMAGINE
PUBLISHING**

© Imagine Publishing Ltd 2015
ISSN 2041-3270

 **recycle**

When you have finished with
this magazine please recycle it.

YOUR FREE RESOURCES

LOG IN TO WWW.FILESILO.CO.UK/LINUXUSER AND DOWNLOAD THE LATEST DISTROS AND FREE SOFTWARE TODAY

LATEST DISTROS



Korora 21



Manjaro 0.8.12



GParted Live 0.21.0-1



Raspbian '15

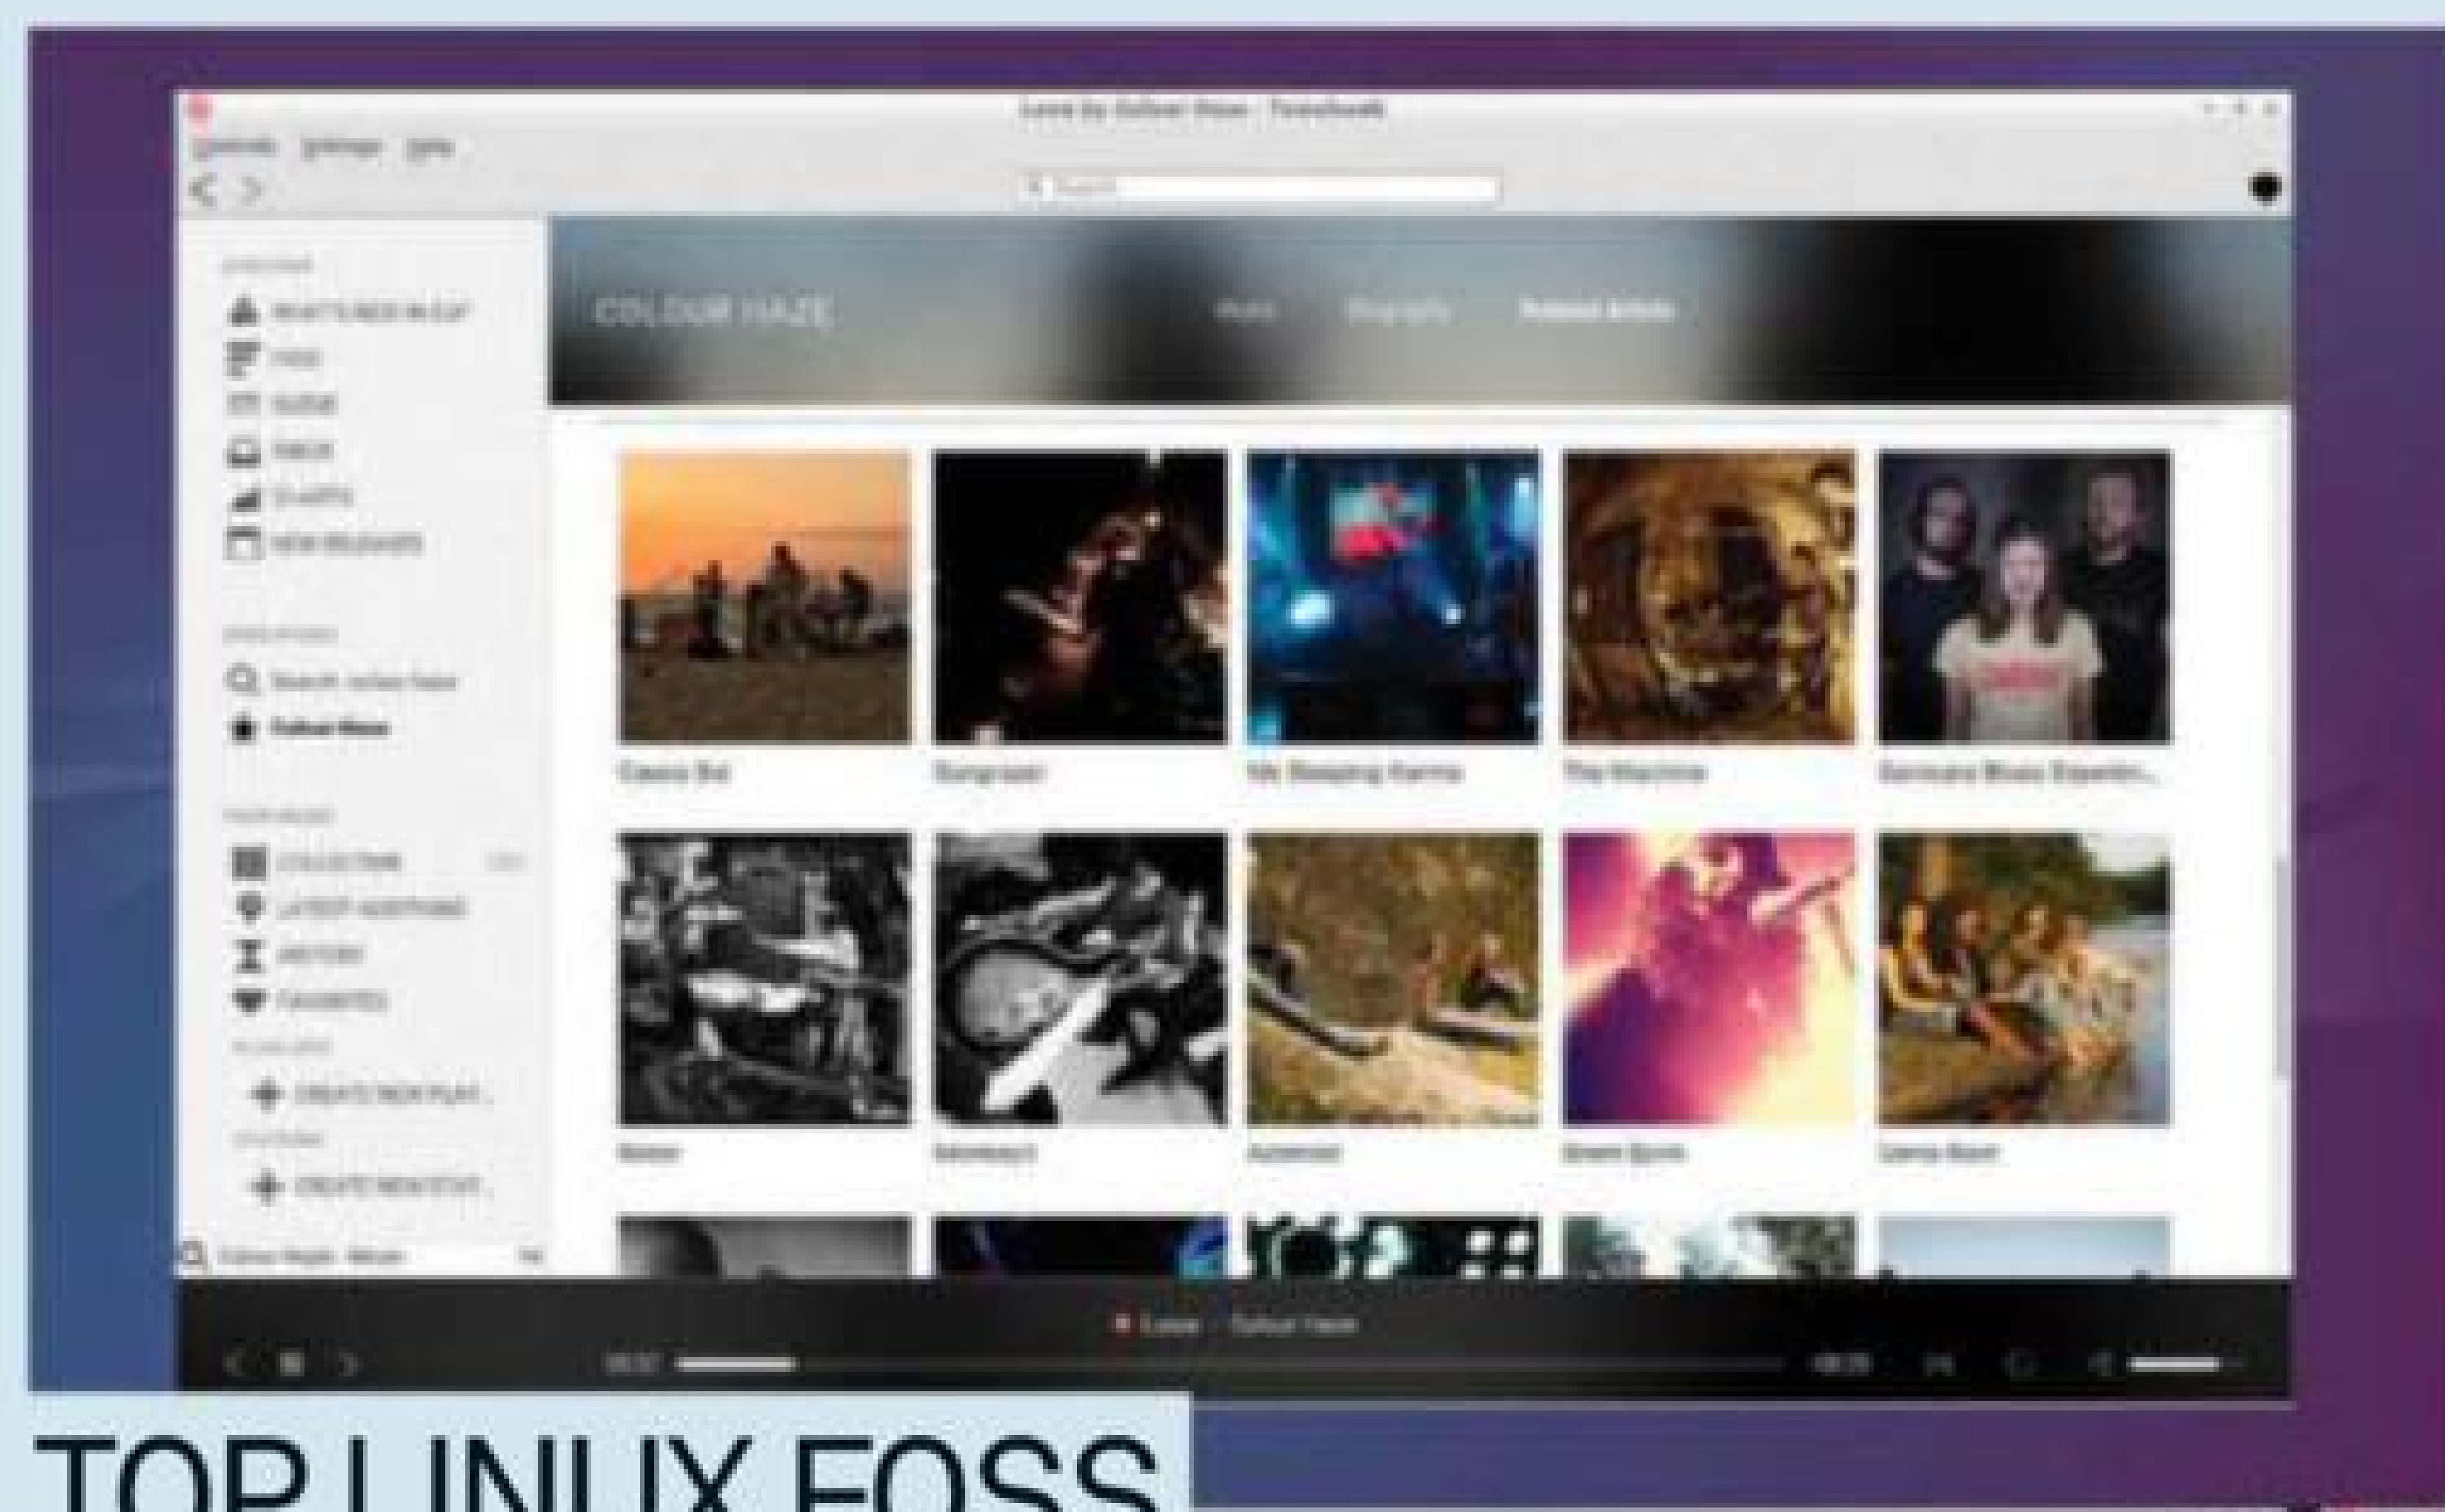
YOUR BONUS RESOURCES

ON FILESIL0 THIS ISSUE, FREE AND EXCLUSIVE FOR **LINUX USER & DEVELOPER** READERS, YOU'LL FIND THESE GREAT RESOURCES...

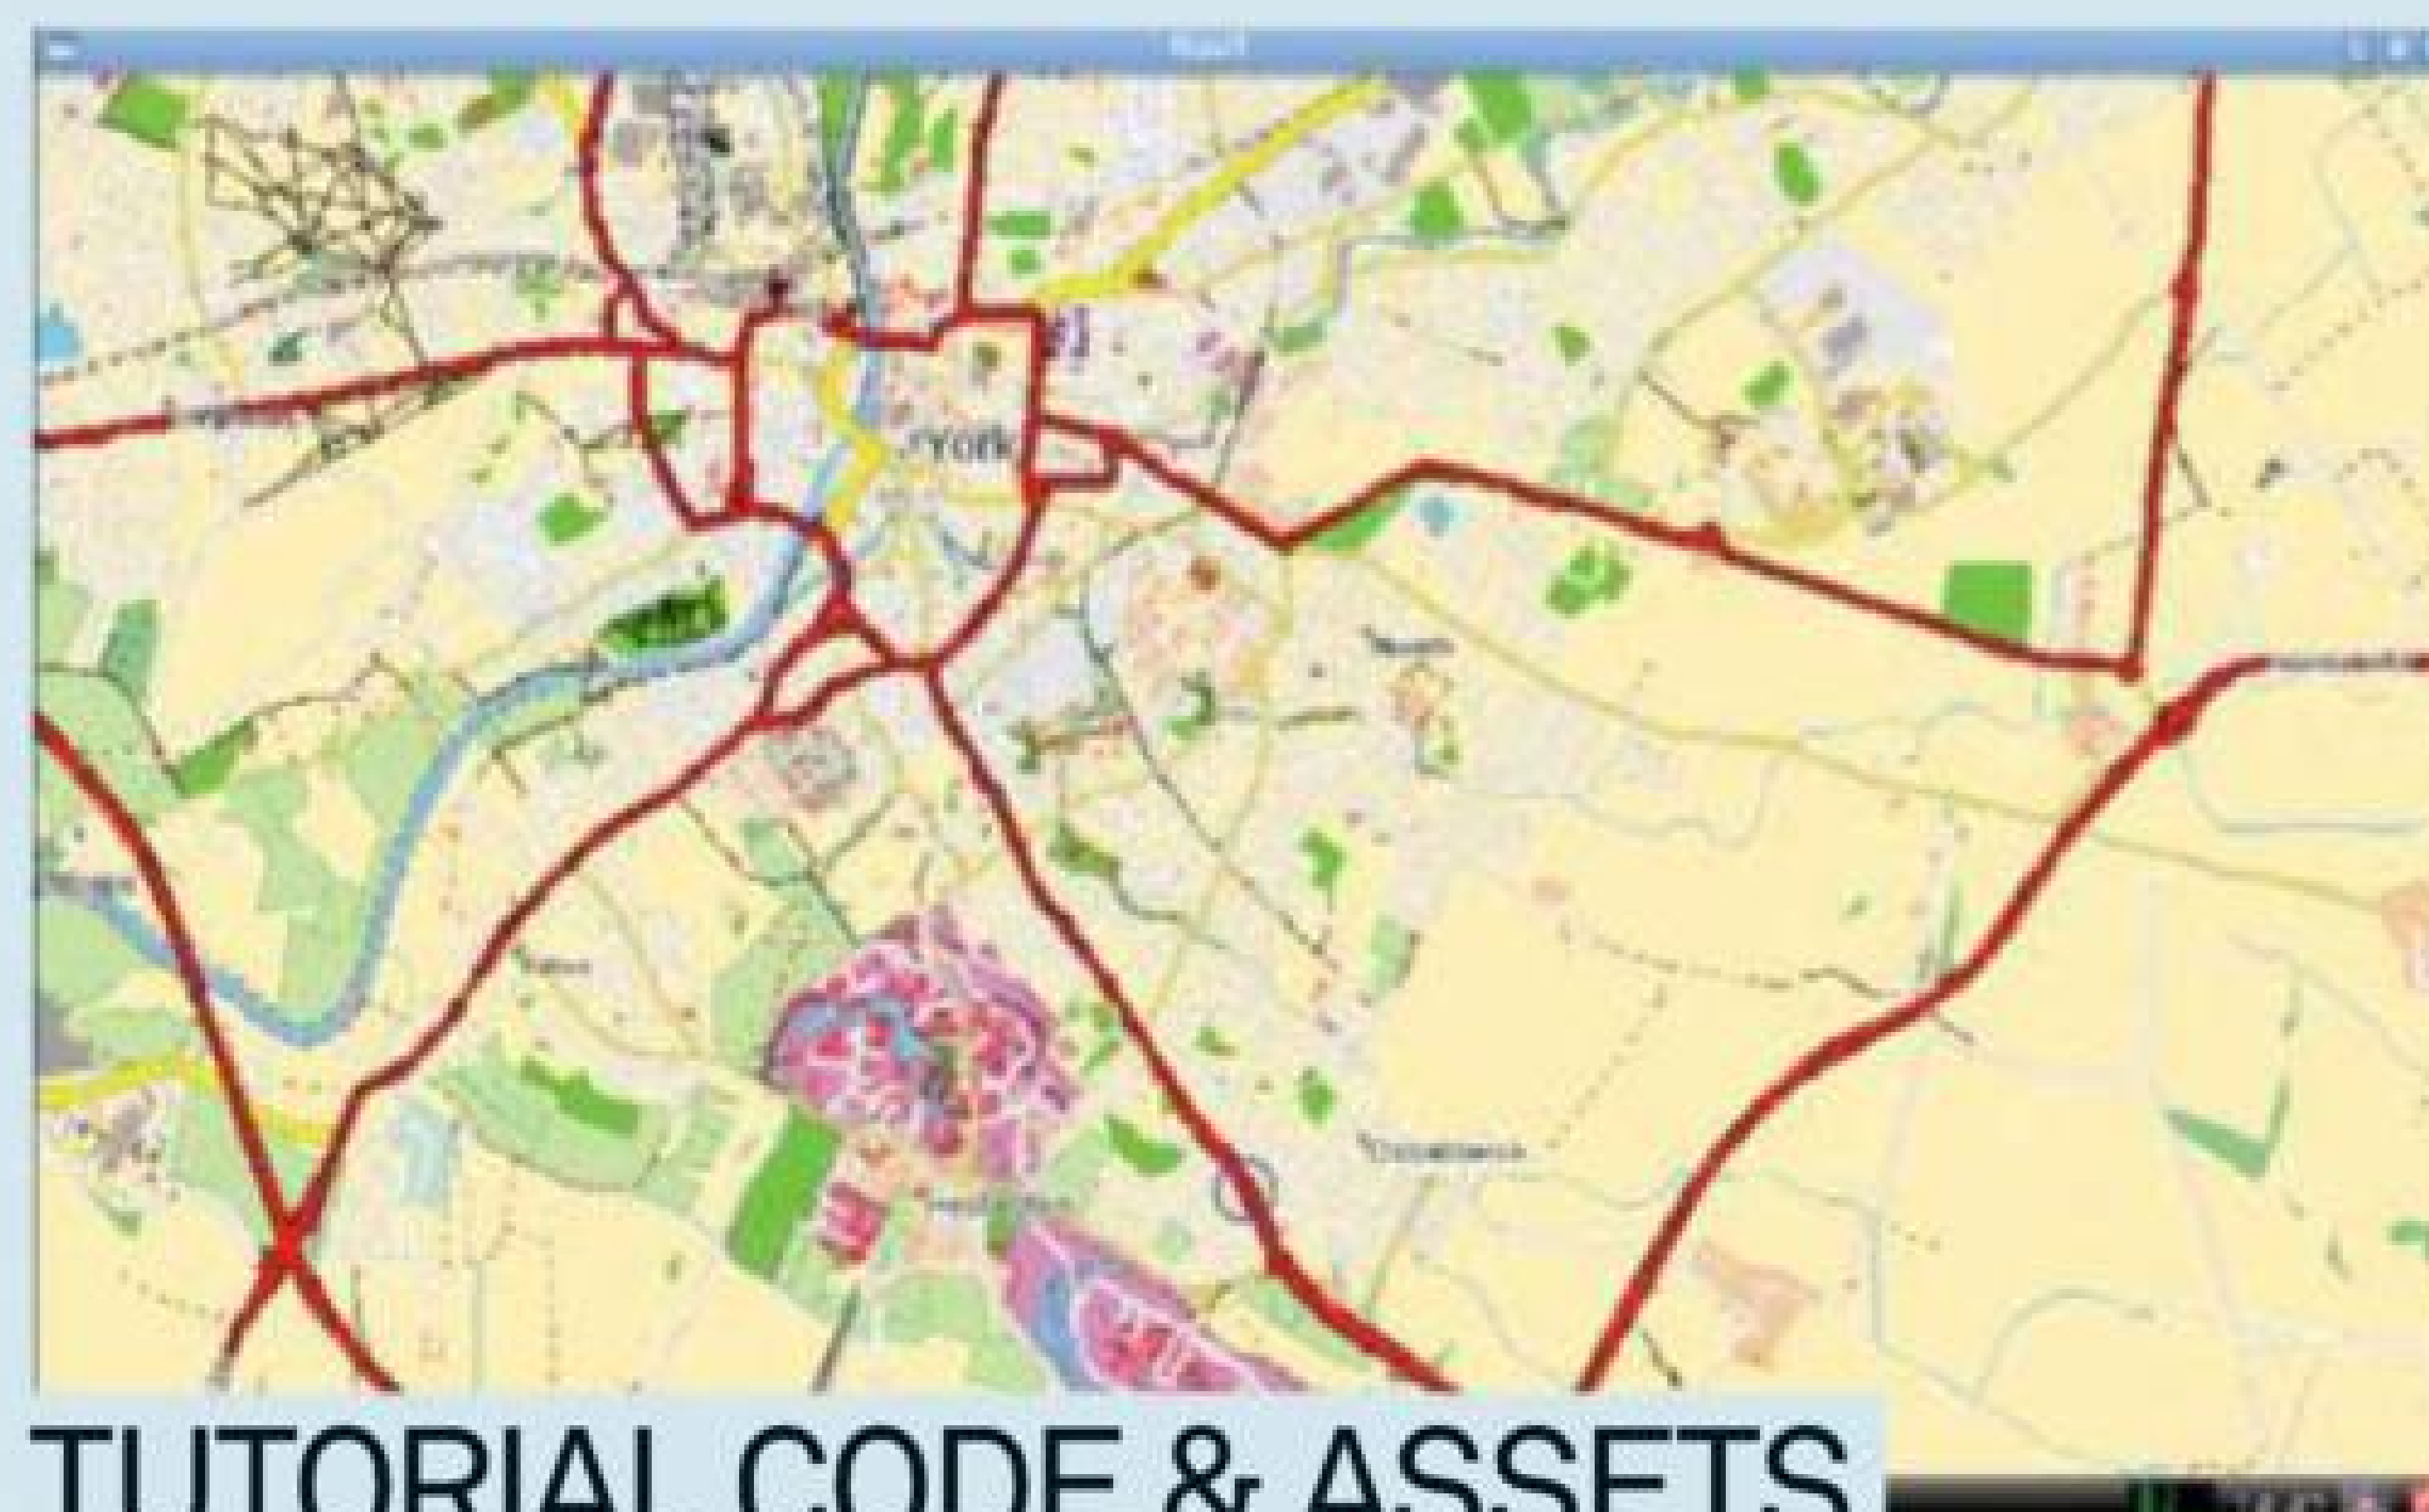
- » **4 Latest distros**, including Korora 21, Manjaro 0.8.12, GParted Live 0.21.0-1 and Raspbian 2015-01-31.
- » **16 Best FOSS packages**, including BTRFS, UNetbootin, Tomahawk, Navit and more.
- » The Linux Foundation, Red Hat and Raspberry Pi video guides, tutorials and webinars.
- » **Code and assets** for the Raspberry Pi Car Computer, LaTeX and Rapiro tutorials.



LENGTH OF VIDEO TUTORIALS: **17** HOURS



TOP LINUX FOSS



TUTORIAL CODE & ASSETS

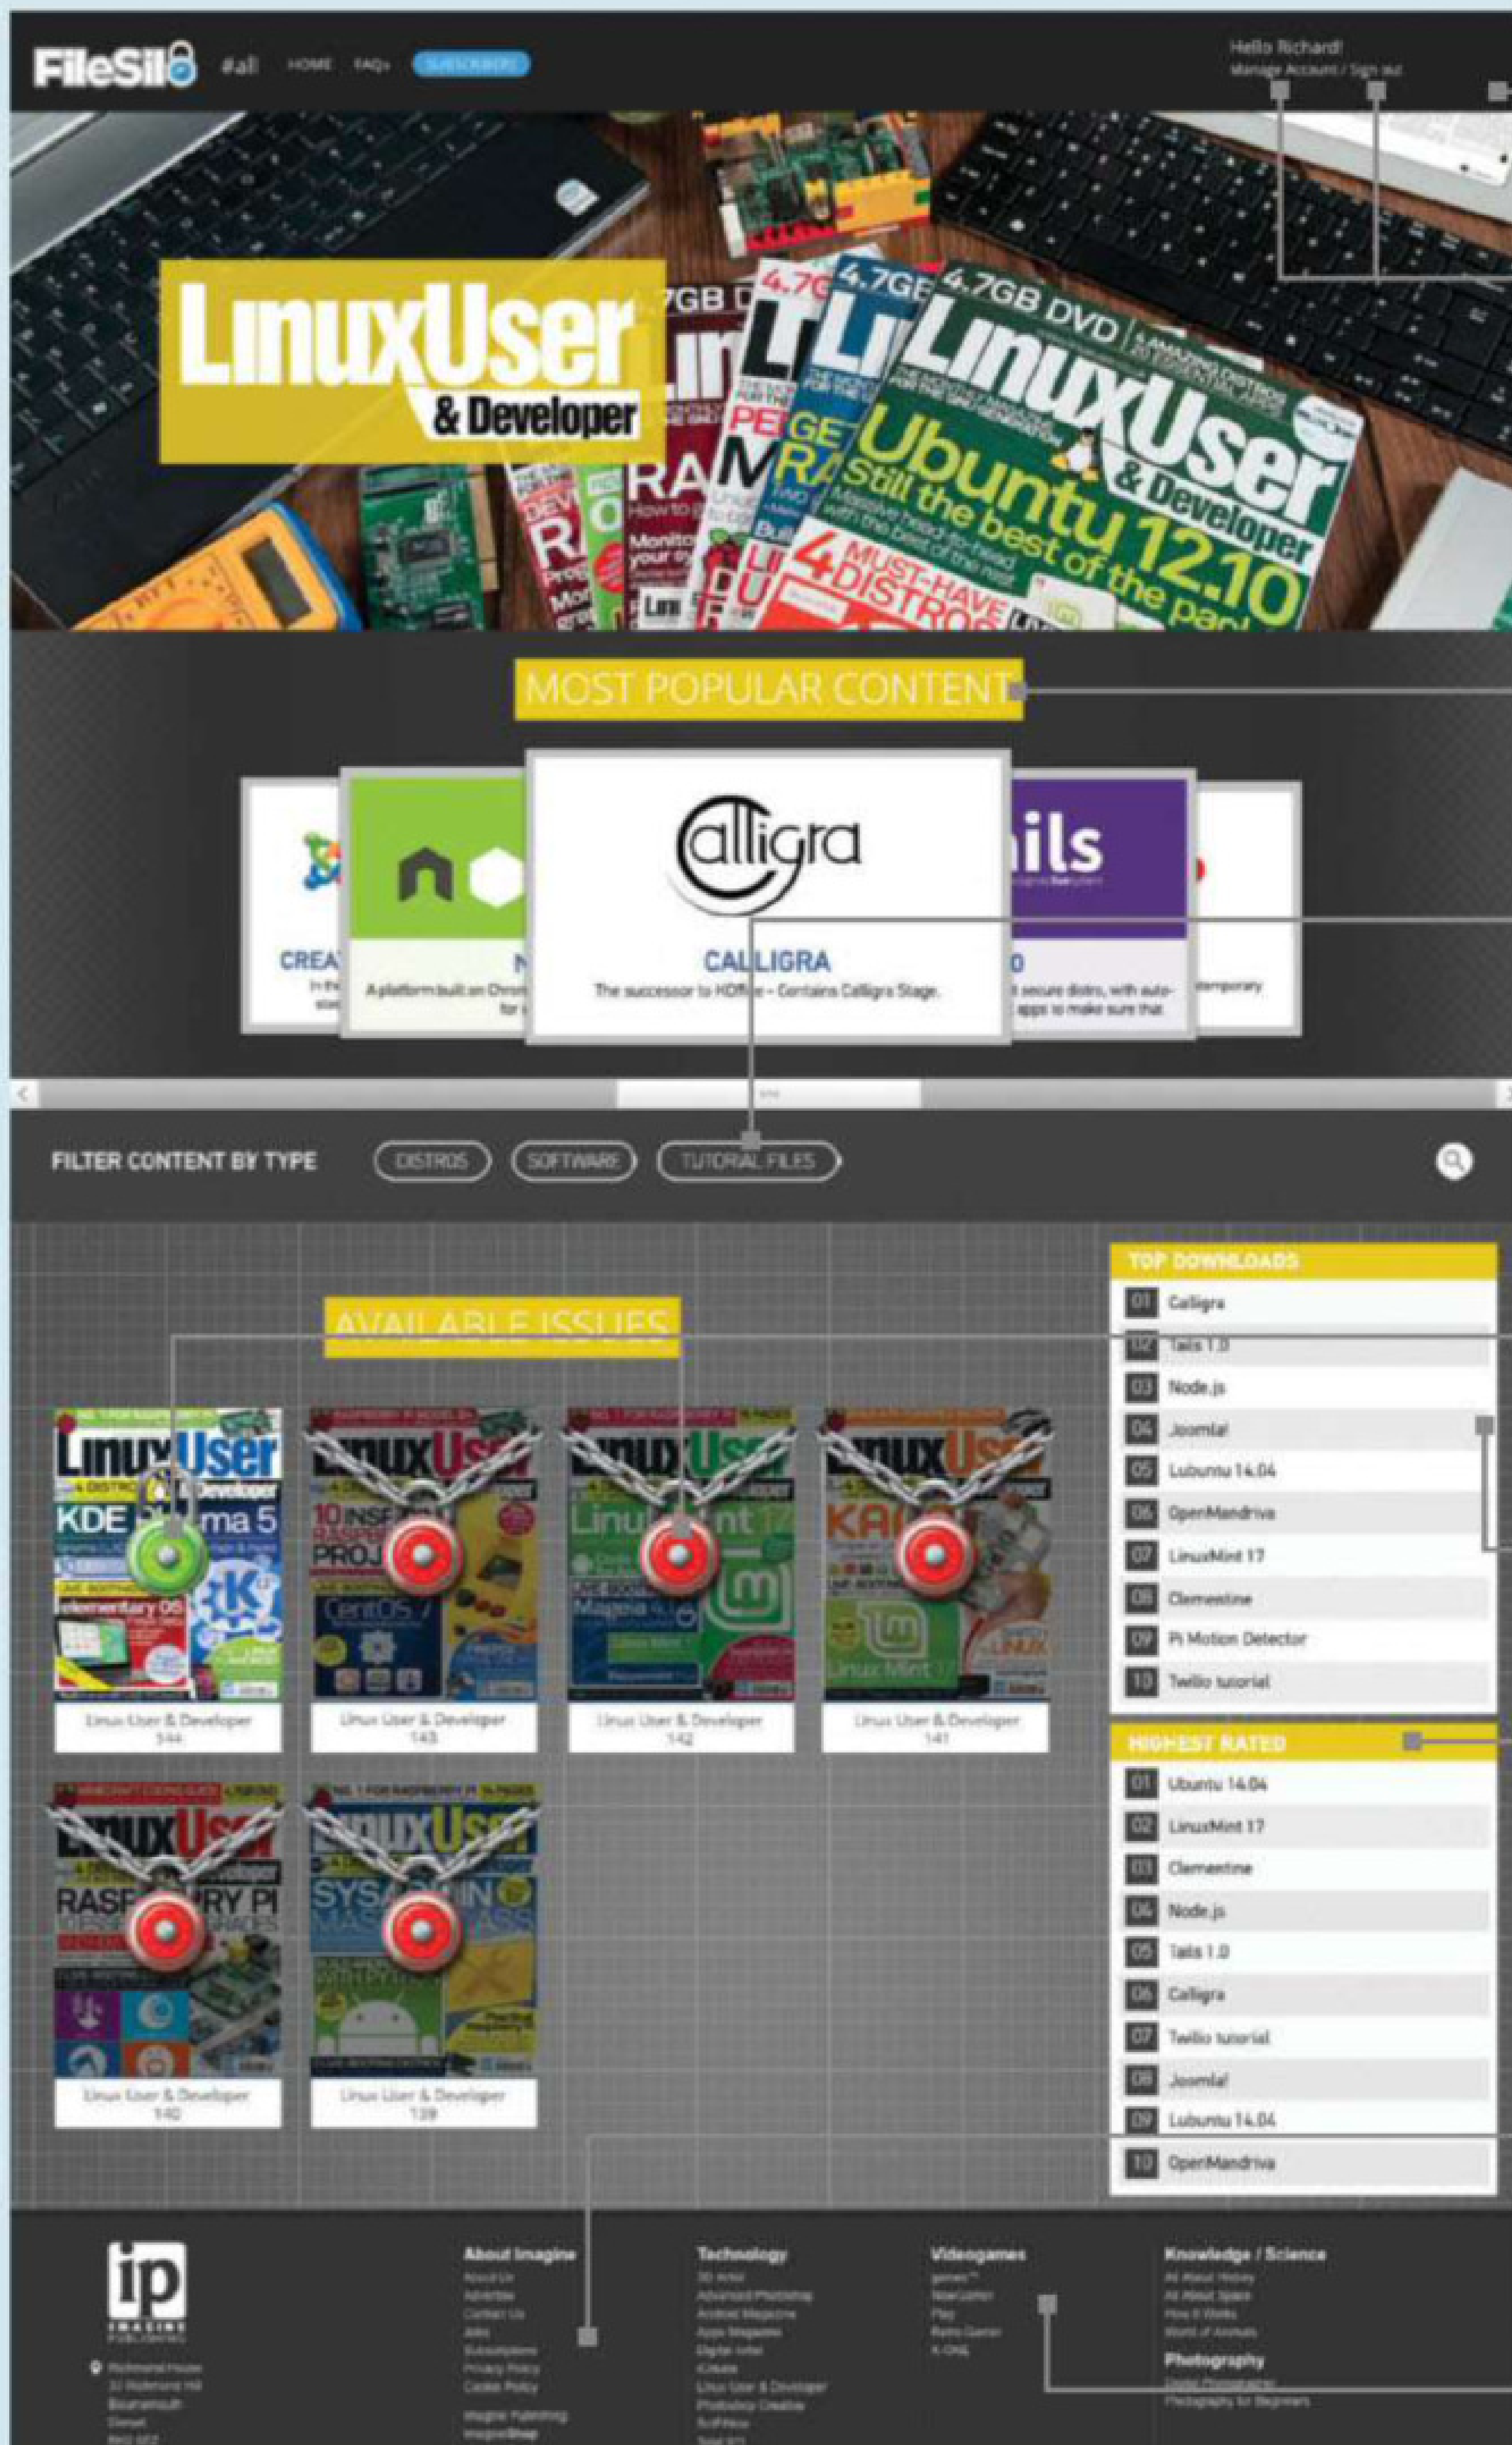
www.filesilo.co.uk/linuxuser

Storemags

FILESILO – THE HOME OF PRO RESOURCES

DISCOVER YOUR FREE ONLINE ASSETS

- 🔒 A rapidly growing library
- 🔒 Updated continually with cool resources
- 🔒 Lets you keep your downloads organised
- 🔒 Browse and access your content from anywhere
- 🔒 No more torn disc pages to ruin your magazines
- 🔒 No more broken discs
- 🔒 Print subscribers get all the content
- 🔒 Digital magazine owners get all the content too!
- 🔒 Each issue's content is free with your magazine
- 🔒 Secure online access to your free resources



This is the new FileSilo site that replaces your disc. You'll find it by visiting the link on the following page.

The first time you use FileSilo you'll need to register. After that, you can use the email address and password you provided to log in.

The most popular downloads are shown in this carousel, so see what your fellow readers are enjoying!

If you're looking for a particular type of content like distros or Python files, use these filters to refine your search.

Green open padlocks show the issues you have accessed. Red closed padlocks show the ones you need to buy or unlock.

Top Downloads are listed here, so you can get an instant look at the most popular downloaded content.

Check out the Highest Rated list to see the resources that other readers have voted for as the best!

Find out more about our online stores, and useful FAQs like our cookie and privacy policies and contact details.

Discover our amazing sister magazines and the wealth of content and information that they provide.

HOW TO USE **FileSilo**

EVERYTHING YOU NEED TO KNOW ABOUT ACCESSING YOUR NEW DIGITAL REPOSITORY

To access FileSilo, please visit www.filesilo.co.uk/linuxuser

01 Follow the instructions on-screen to create an account with our secure FileSilo system, then log in and unlock the issue by answering a simple question about the magazine.

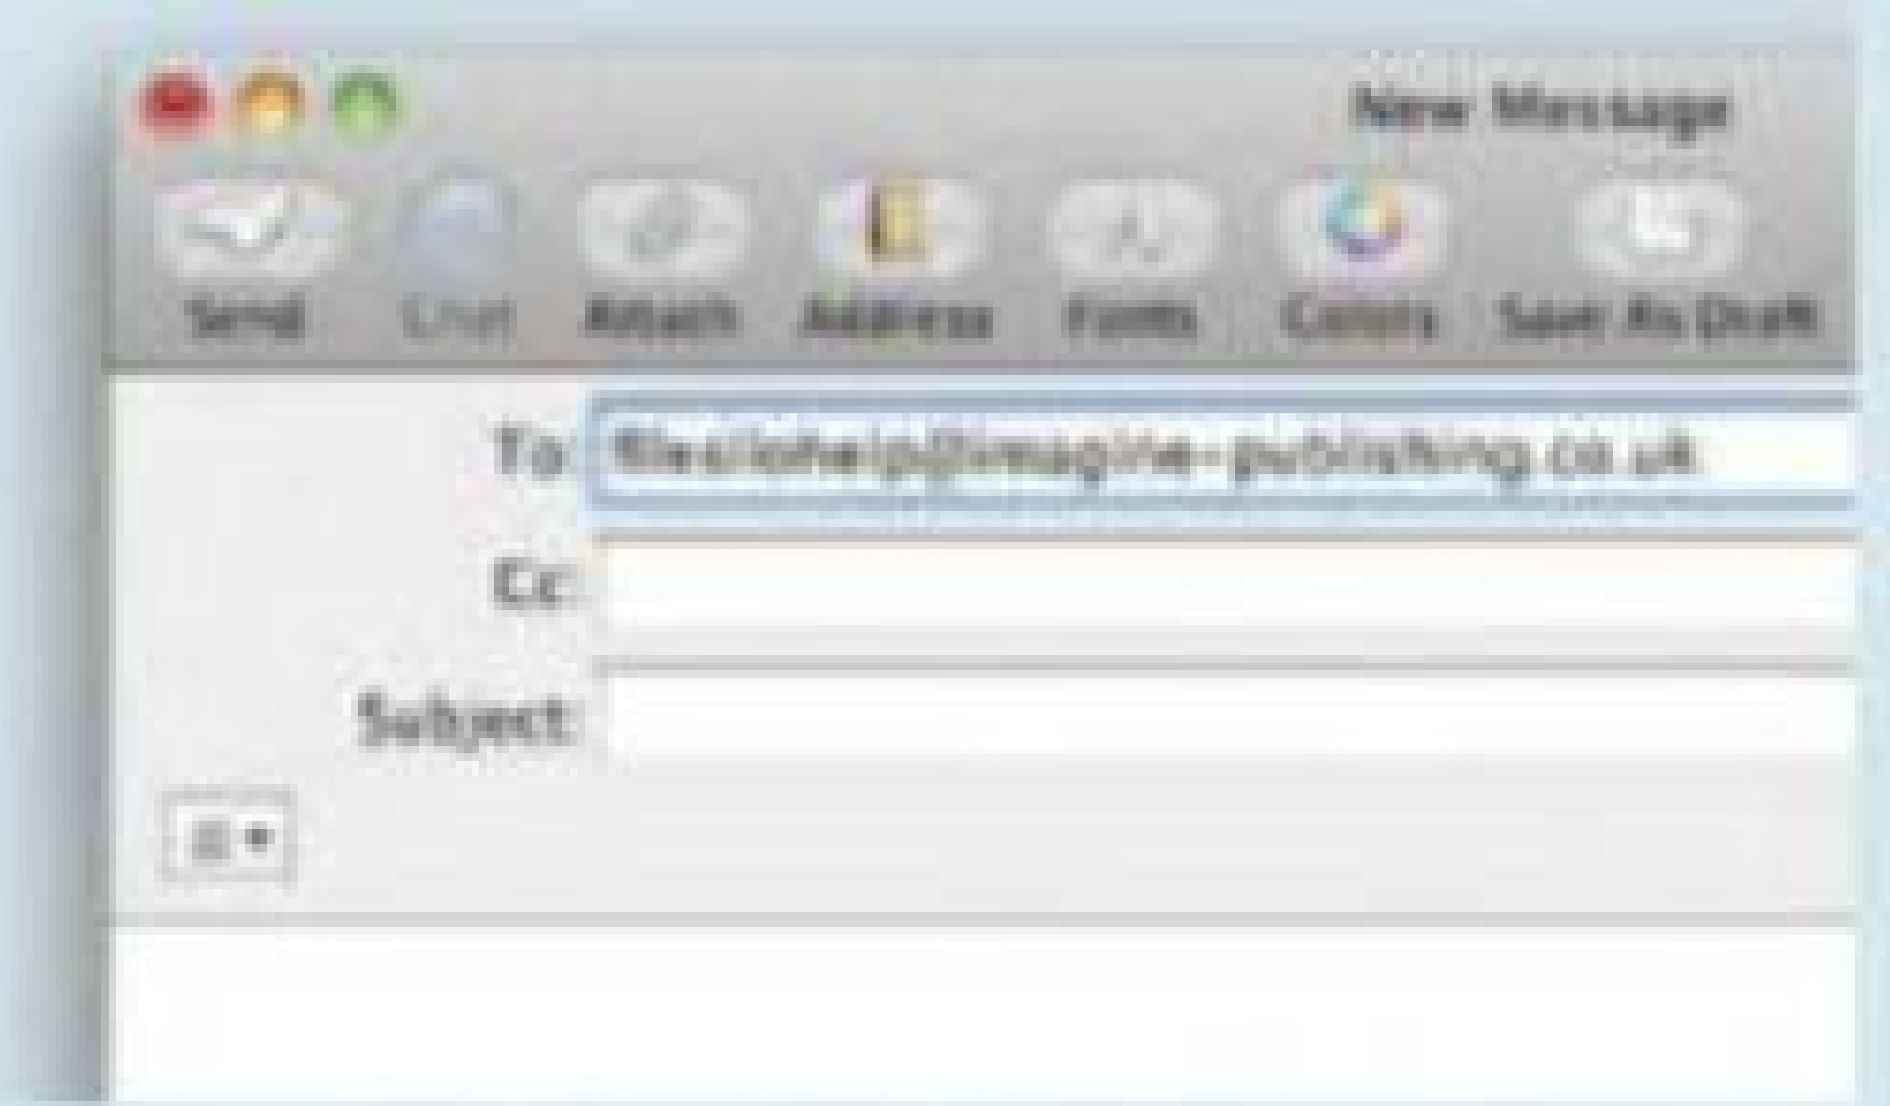


You can access the content for free with your issue.

02 If you're a print subscriber, you can easily unlock all the content by entering your unique Web ID. Your Web ID is the eight-digit alphanumeric code printed above your address details on the mailing label of your subscription copies. It can also be found on your renewal letters.

03 You can access FileSilo on any desktop, tablet or smartphone device using any popular browser (such as Firefox, Chrome or Safari). However, we recommend that you use a desktop to download content, as you may not be able to download files to your phone or tablet.

04 If you have any problems with accessing content on FileSilo, or with the registration process, take a look at the FAQs online or email filesilohelp@imagine-publishing.co.uk



MORE TUTORIALS AND INSPIRATION


Finished reading this issue? There's plenty more free and open source goodness waiting for you on the Linux User & Developer website. Features, tutorials, reviews, opinion pieces and the best open source news are uploaded on a daily basis, covering Linux kernel development, the hottest new distros and FOSS, Raspberry Pi projects and interviews, programming guides and more. Join our burgeoning community of Linux users and developers and discover new Linux tools today.



www.linuxuser.co.uk

Issue 151 of **LinuxUser** is on sale 9 April 2015 from GreatDigitalMags.com

wirelessthings.net/radios



USB Mac Linux Windows Android - Raspberry Pi & Pi2 - Arduino - XBee - Surface Mount

Dedicated, efficient & powerful

Ben Wallace, Director of drpdigital at drp,
Fasthosts customer since 2010


Dedicated to business excellence

For sheer flexibility and reliability, nothing compares with your own Dedicated Server.

Fasthosts offers the performance, technology and security that can take you to the next level – with UK-based service that keeps you online (99.99% uptime guarantee).

Give your business a server that meets your own budget and needs:

- UK data centres
- 24/7 UK support
- Full remote KVM
- Up to 128GB RAM
- Unlimited bandwidth
- NEW cPanel and WHM
- SSD technology
- Massive storage



 TRUSTPILOT



Trustpilot rating correct at
time of going to print 02/2015



Dedicated Servers from:

£19.00
per month

Per month ex VAT. Normal price £39.00 per month ex VAT.
12 month minimum term contract. One-off set-up fee £49.00

Call **0333 0142 709**
or visit **fasthosts.co.uk**

SERVERS • WEB HOSTING • DOMAIN NAMES • EXCHANGE EMAIL

