NEW

The best Linux guide money can buy!

LINUX MADE SIMPLE

The ultimate guide to open source computing

148 pages of advice from the makers of **LINUX FORMAT**

Digital Edition

FUTURE

TENTH EDITION

LINUX
MADE SIMPLE

With the internet driving a new world of information exchange, business start-ups and online gaming, coders have suddenly become the gatekeepers to these new realms. Combine the huge interest in learning how to create and control these worlds with the surge of cool devices, such as the Raspberry Pi, and you've got a head of technological steam the world hasn't seen since the coding craze of the early 1980s. So no matter if you're looking to relive those heady '80s coding days or are a newbie looking to take your first steps into the coding world, you hold in your hands the ideal guide to start coding. Thanks to a new generation of open free software, we all can access operating systems, development tools, compilers and the programming languages needed to create professional programs, apps and tools. We'll show you how to get up and running with a Linux system, then access everything you need freely online. Coding is easy, exciting and fun. We'll explain the basics, move on to more advanced topics, explain how you can use the Raspberry Pi, and provide you with fully updated exciting and easy-to-follow projects. So what are you waiting for? Get coding!

FUTURE

LINUX MADE SIMPLE

Future PLC Quay House, The Ambury, Bath, BA1 1UA

Editorial
Editor **Katharine Marsh**
Designer **Perry Wardell-Wicks**
Compiled by **April Madden & Steve Dacombe**
Senior Art Editor **Andy Downes**
Head of Art & Design **Greg Whitaker**
Editorial Director **Jon White**
Managing Director **Grainne McKenna**

Photography
All copyrights and trademarks are recognised and respected

Advertising
Media packs are available on request
Commercial Director **Clare Dove**

International
Head of Print Licensing **Rachel Shaw**
licensing@futurenet.com
www.futurecontenthub.com

Circulation
Head of Newstrade **Tim Mathers**

Production
Head of Production **Mark Constance**
Production Project Manager **Matthew Eglinton**
Advertising Production Manager **Joanne Crosby**
Digital Editions Controller **Jason Hudson**
Production Managers **Keely Miller, Nola Cokely, Vivienne Calvert, Fran Twentyman**

Printed in the UK

Distributed by Marketforce – www.marketforce.co.uk
For enquiries, please email: mfcommunications@futurenet.com

GPSR EU RP (for authorities only)
eucomply OÜ Pärnu mnt 139b-14 11317, Tallinn, Estonia
hello@eucompliancepartner.com, +3375690241

Linux Made Simple Tenth Edition (TCB6862)
© 2025 Future Publishing Limited

We are committed to only using magazine paper which is derived from responsibly managed, certified forestry and chlorine-free manufacture. The paper in this bookazine was sourced and produced from sustainable managed forests, conforming to strict environmental and socioeconomic standards.

All contents © 2025 Future Publishing Limited or published under licence. All rights reserved. No part of this magazine may be used, stored, transmitted or reproduced in any way without the prior written permission of the publisher. Future Publishing Limited (company number 2008885) is registered in England and Wales. Registered office: Quay House, The Ambury, Bath BA1 1UA. All information contained in this publication is for information only and is, as far as we are aware, correct at the time of going to press. Future cannot accept any responsibility for errors or inaccuracies in such information. You are advised to contact manufacturers and retailers directly with regard to the price of products/services referred to in this publication. Apps and websites mentioned in this publication are not under our control. We are not responsible for their contents or any other changes or updates to them. This magazine is fully independent and not affiliated in any way with the companies mentioned herein.

FUTURE

Connectors. Creators. Experience Makers.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief Executive Officer **Kevin Li Ying**
Non-Executive Chairman **Richard Huntingford**
Chief Financial Officer **Sharjeel Suleman**
Tel +44 (0)1225 442 244

LINUX MADE SIMPLE

Contents

Get started

Getting started with Linux 10
Escape from Windows 10 14
Pick a distro 24
Top 100 Linux tools 34

Basics

Package management 44
Web browsers 46
Documents 48
Spreadsheets 50
PDF editing 52
Instant messaging 54
Email 56
Text editors 58

Fun

Playing music & video 62
Audio editing 64
Photo organisation 66
Photo editing 68
Video editing 70
Setting up a media centre 72
Playing games 74

Devices

Installing drivers 80
Setting up printers and scanners 82
Managing Linux drives 86
Using removable media 88
Understanding the Linux filesystem 90
Partitioning drives 92
Managing users 96

Advanced

File sharing 100
Alternative desktops 102
Remote desktop 104
Virtual machines 106
System repair 110
Fix startup issues 114
Two-factor authentication 118
Easy file encryption 120
Core commands 122
Troubleshooting 126
Terminal basics 136
Terminal files 138
Apt-get 140
Archiving files 142

Get Started

Introduction to Linux 10
Escape Windows 10 14
Pick a distro 24
Top 100 Linux tools 34

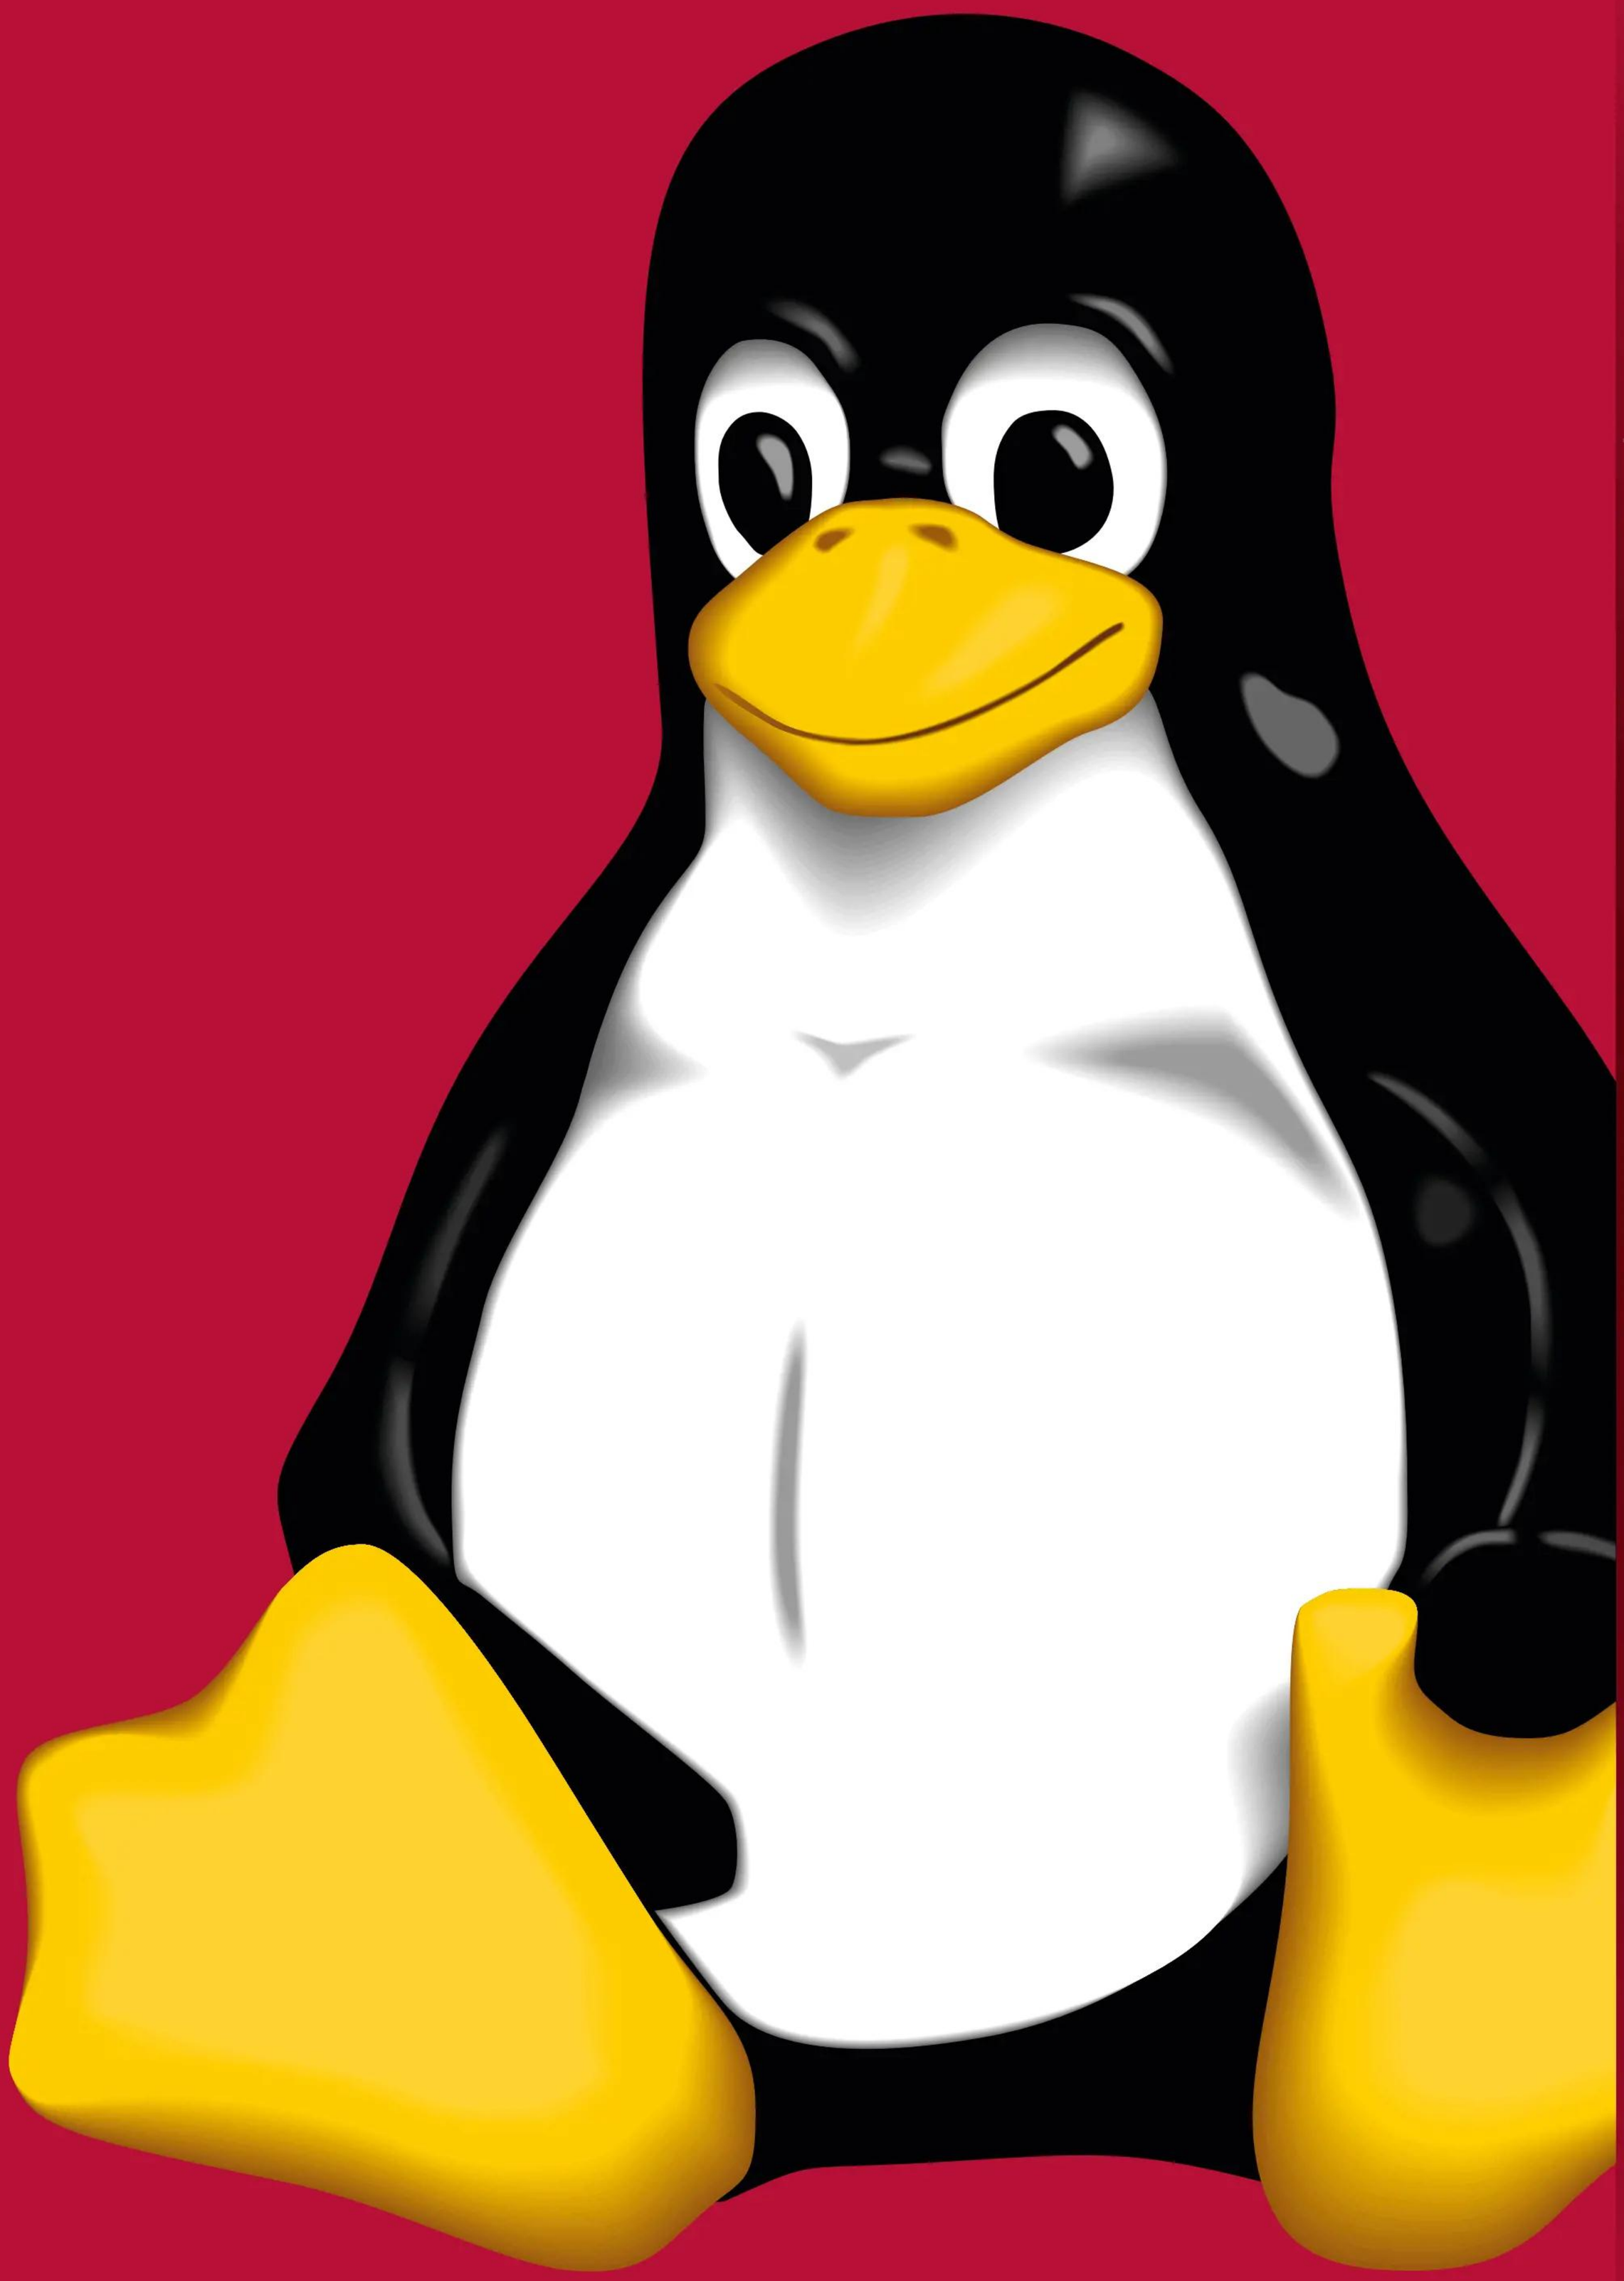

GET STARTED WITH LINUX

It runs on most desktop and laptop PCs, it's free and you don't even need to install it. Let's look at Linux!

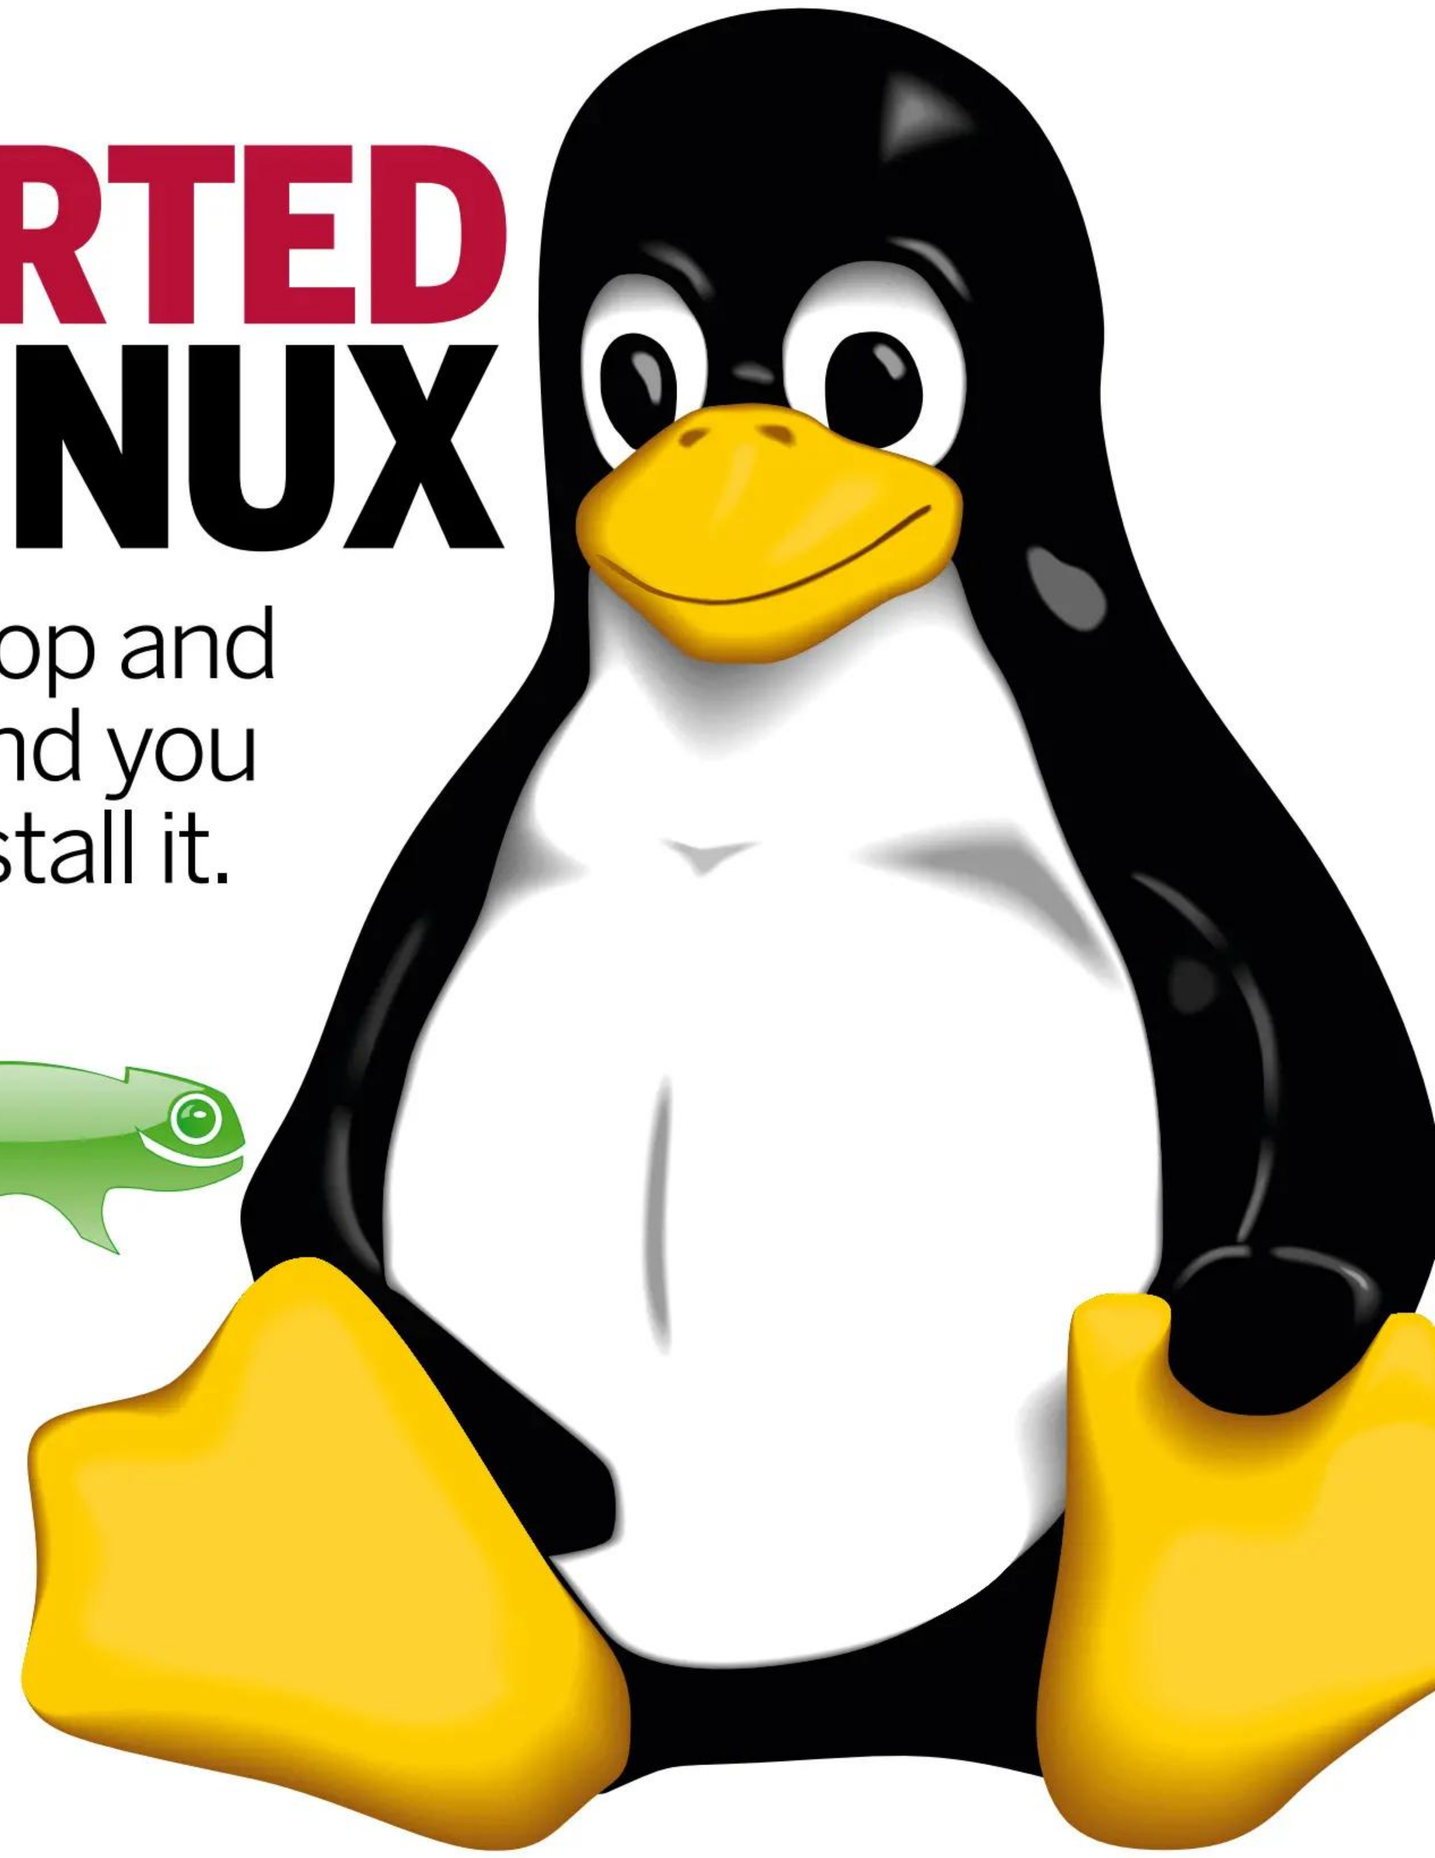

As you read through this *Linux Made Simple* bookazine, you'll notice that most of the screens don't look like Microsoft Windows or Apple mac OS. There's a good reason for that: they're not. Just as there are different car manufacturers or makes of TV, there's more than one operating system that can run your PC or Mac. It's just that Linux happens to be free because it's developed by thousands of coders around the world.

All the coding projects in this bookazine are based on someone running Linux on their desktop. You don't have to – the code works on Windows or mac OS – but Linux comes with many languages built in or ready to install from a central server. No scouting round dodgy sites; just run a command or fire up a software centre to get what you need.

That's the beauty of Linux – it's built by geeks for geeks to do geeky things. Once you're used to its slightly different interface and way of working, you'll find how easy it really is to get along with. We recommend that you use a version (known as a distro) called Linux Mint, and we're going to look at the ways you can get hold of Mint and then get it up and running on your PC.

> "It's not scary, it won't damage your PC and you don't even have to install anything."

It's not scary, it won't damage your PC and you don't even have to install anything if you don't want to. If you currently have a Windows PC, there are three options: run Linux in VirtualBox, run it off a DVD or USB drive, or install it on your PC dual-booting with Windows.

We're only going to look at the first two as they're the least likely to cause any damage. The walkthrough opposite explains how to run Mint within VirtualBox on top of Windows (or mac OS), while the next walkthrough shows you how to create a live image that you can boot and run from either a DVD or a USB flash drive.

VirtualBox

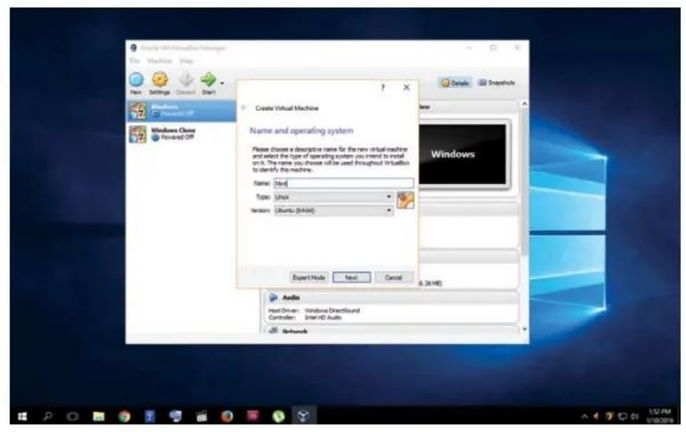

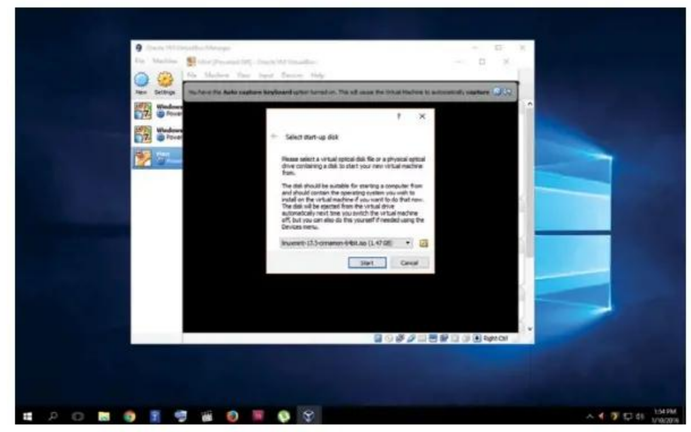

1 Get VirtualBox

Head to **www.virtualbox.org** and download Virtual Box 5 for your operating system, be that Windows or mac OS. Install it and be aware you'll need around 20GB of spare drive space to store the virtual OS file. You also need the Mint ISO file from **www.linuxmint.com**. Once installed, start VirtualBox, click New Machine, choose Linux and call it 'Mint'.

2 Create a machine

Choose Ubuntu and the bits should match the ISO you downloaded. Click Next. Under Memory, we recommend 2048, but if you have an 8GB PC, 4096 is best. You can leave all the rest as default settings, apart from the dynamic hard drive size. The default is 8GB, but we suggest 32GB just in case. Finish and click Start to get going.

3 Start virtual Mint

A prompt asks for a disc. Locate the Mint ISO file you downloaded and click Start. Linux Mint starts and, once loaded, you're free to try it out or use the Install icon to properly install Mint to the virtual machine. For extended use, in the virtual machine's settings under Display, you should enable 3D acceleration and allocate 16MB of memory.

If you've never tried Linux, we think you'll be surprised at how easy it is to use, and how good a development platform it makes. You should also know that Linux runs the majority of the internet, from its core servers to websites, alongside powering 97% of the world's supercomputers, and being used widely in science and industry. So it's not such a bad thing to understand, after all.

No matter which tutorial you follow, you'll need a copy of the latest Linux Mint Cinnamon distro. Head to **www.linuxmint.com/download.php** and download the 32-bit build, unless you know your system is 64-bit, in which case get that. If you have an older or slower machine, opt for the MATE edition, which is less graphically demanding.

Up and running

A big reason why Linux makes such a good development platform is that it was created by developers. In many ways, the key weakness of Linux has been its lack of ease of use (Mint is an exception to that rule) but then, when it's designed and created by developers, ease of use is going to come at the bottom of the to-do list. The upshot of this history is that there's a wealth of the most advanced tools freely available on all Linux platforms. As long as you're willing to seek out help yourself and contribute back to the community, you'll find a very welcoming and rich development ecosystem waiting for you.

Getting apps

With Windows, in the past you've been used to getting programs by downloading them from here, there and everywhere. More recently, the introduction of the Windows Store has at least centralised where software comes from and removed the worry of getting infected by viruses and malware. The fact is that with Linux, the key way of getting new tools and programs has always been from a central repository of software, protected and maintained by the distro's creators.

This is one of the reasons why Linux has remained so secure (it's not infallible) but people downloading dodgy software is a key way that machines become infected. With Linux Mint, there's the Software Center, which gives you access to hundreds of programs and all the programming development tools you »

Top tip

What's a distro?

Unlike Windows and mac OS, because Linux is free software, anyone can take it and effectively create their own OS to distribute. In the Linux world, these are called distros for short, and there are literally hundreds out there – not all good, not all maintained, but hundreds nonetheless.

The GNU of GNU/Linux

The GNU project (GNU stands for 'GNU's Not Unix') predates Linux by several years. It had created most of the basic tools needed by a computer of the early 1980s – compilers, text editors, file and directory manipulation commands, and much more – but did not have a usable kernel (some would say its kernel, GNU Hurd, is still not that usable). When Linus Torvalds started tinkering with his small project in 1991, he had a kernel without the tools to run on it. The two were put together and GNU/Linux was born – an operating system using the Linux kernel and the GNU toolset. It is not only the programs in **/bin** and **/usr/bin** that come from GNU; glibc is the core C library used in Linux and it also comes from GNU. So just about any time you do anything on your computer – every time you type a command or click an icon – GNU software is being run at some level.

No wonder the GNU die-hards get upset when we refer to our operating system as Linux and not GNU/Linux. It is worth mentioning that no one really denies the importance of the GNU aspect; calling the OS Linux rather than GNU/Linux has far more to do with convenience and laziness than politics – the full title is just too cumbersome.

Easy ways to run Linux

If you're not a big Linux user, then you'll probably not want to destroy your existing Windows or Mac system, and we don't blame you. The truth is, you don't need to – Linux is flexible enough that it can be run in a number of ways beside, on top of or alongside most other operating systems, and on most types of hardware – from virtual versions to versions running off spare USB drives, DVDs or on low-cost hardware such as the Raspberry Pi. The standard way, once you've got your hands on the ISO, is to burn it to a DVD. When you first turn on a PC, you can usually get it to boot from alternative media by pressing F11/F12, or hold down C on the Mac – some PCs boot from a suitable optical disc by default.

Another option is to install VirtualBox from **www.virtualbox.org** (see previous page). Install and run this – it looks complex, but creating a virtual PC is pretty easy if you stick to the default settings. The main stumbling block is ensuring under Storage that you add the ISO file to the virtual optical drive. There are more options available, including writing the ISO file to a suitable USB thumb drive and, following a similar boot process as discussed above, running Linux from this. To get this to work, you need to use a write tool such as UNetbootin from **http://unetbootin.github.io**.

For the brave-hearted, you can also install Linux on your system directly. Most versions of Linux create space alongside Windows and make a dual-boot system.

» could wish for. It's not why we're here but you can also download Valve's Steam gaming client and take advantage of over 1,700 Linux games it has to offer.

Drivers

We're not really here to talk about using Linux in every detail but there are a few standard questions that often crop up when people move over from Windows. One key one is where are all the drivers? The cool thing with Linux is that, on the whole, there's no need to worry about drivers – they're built into the Linux kernel. That isn't to say you can't add drivers, but they're generally not required. There are a couple of exceptions: certain more obscure laptop wireless cards can cause issues, while if you want maximum 3D gaming performance, you need to install the dedicated graphics driver from your card's manufacturer.

The Terminal

If you've heard of Linux, then one area you might fear or just wonder about is a thing called the Terminal. This can be called a number of different things, depending on the system, such as the command line, command prompt or command-line interface. It is a direct interface to the operating system and all of its tools, which you access through text commands. Going back to the early days of computers and Linux, as computers were so much slower, there weren't any graphical interfaces, so computers were controlled entirely through text commands.

Linux was developed originally in this type of environment, so all of its core tools are based on Terminal use. Many of them – or all the ones you'll care about – do have graphical interfaces these days. The fact is, the Terminal remains an efficient way of controlling Linux, and when it comes to troubleshooting, it offers a consistent set of tools and interfaces to resolve problems.

You don't need to know anything about the Terminal to use Linux on a day-to-day basis, but it's good to know that it's there just in case. However, we would advise you to at least open a Terminal and use the `ls` command to list a directory, and `cd` to change directory.

Linux isn't Windows

One thing to keep in mind when you first use Linux is that it's not Windows or Mac OS X. This largely means things you're used to in Windows won't work the same in Linux, or be in the same place, nor does it offer the same set of programs. So big commercial products such as Microsoft Office, Adobe Photoshop and development tools such as Microsoft Visual Studio aren't directly made for Linux – you can run them via a system called Wine – but the Linux open-source community has created its own tools such as LibreOffice, Gimp, Krita and a whole range of freely available and open-source products that offer the same capabilities.

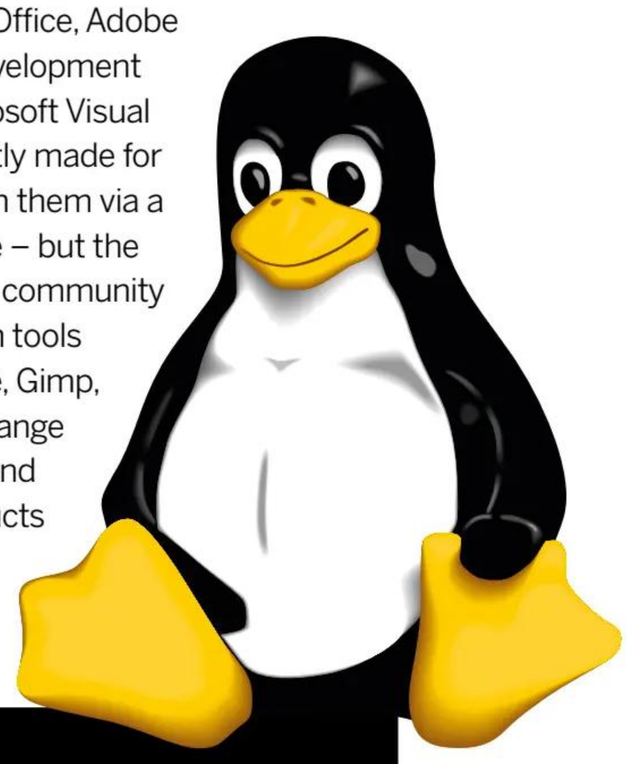

Mint on a USB drive

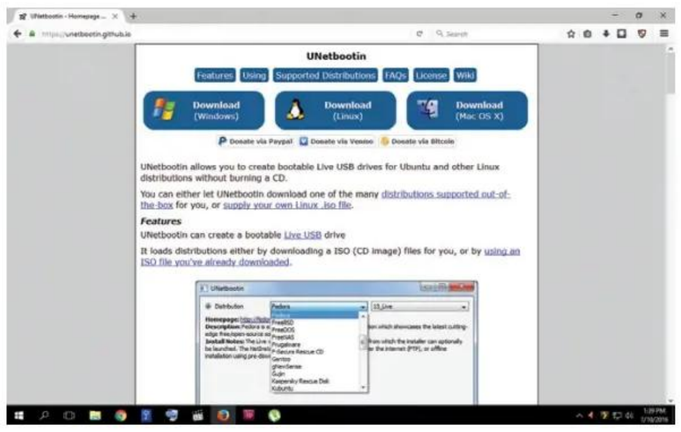

1 UNetbootin Linux

To run Mint from a USB stick, you first need a USB drive at least 4GB in size – 16GB would be ideal. You'll need the Mint ISO file from **www.linuxmint.com**, as discussed in the VirtualBox walkthrough, and we'll use the download tool UNetbootin from **http://unetbootin.github.io**. This installs the live disc ISO file directly to your USB drive.

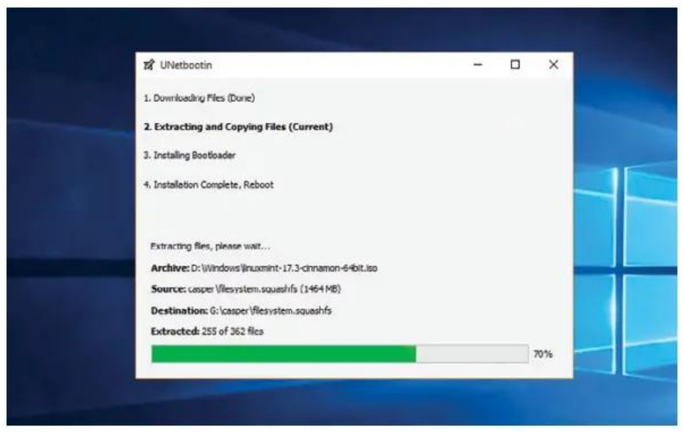

2 Install Mint

The tool can directly download the Mint ISO image for you, but it's best practice to do this yourself, then you have a local copy to hand if you want to create a DVD copy or redo the process. Select Diskimage and locate the file in the Download folder. Ensure you have the correct USB drive selected in the pull-down menu and click OK to create the drive.

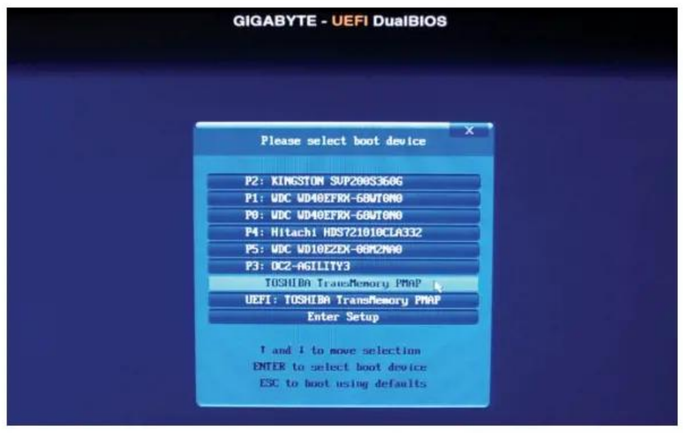

3 Boot and run

You can now boot your PC from the USB drive. However, you need to ensure your PC selects the USB drive as the boot device. Usually, when you first turn on your PC, a message says to press F11 or F12 to select the boot device. Some PCs have their own specific button – consult your manual or manufacturer. Linux Mint will now run.

Linux distro guide

A tantalising taste of the many different flavours of Linux...

In the Linux world, because it's free software that can be redistributed by absolutely anyone, this has lead to a proliferation of different versions of the Linux OS appearing. There are Linux distros tailored to a host of common uses, from hardcore Linux geeks to servers, science, media editing and so many more. The majority are general-use distros, which aim themselves at different types of users, such as experienced hands-on Terminal types, complete beginners, businesses, those with older PCs and people who want modern fancy interfaces. The choice can be bewildering, so to help you get started if you're interested in finding out more about Linux, here's our rundown of the major distro groups.

Debian

www.debian.org

One of the older distributions on the block, Debian is also technically the father of the most number of spin-off distros, including one of the most popular in the world – Ubuntu (see below). Debian itself can be thought of a bare-bones distro, because it comes with just enough to get it installed and up and running. It's really designed for servers and experts, but it's very stable and has complete repositories of software, which means that it is also very easy to extend.

Ubuntu

www.ubuntu.com

This is arguably the most popular – or at least the most widely known – Linux distro in the world. As we mentioned above, Ubuntu was spun out of the Debian project to create an easy-to-use distribution that would be suitable for anyone and everyone to use. Over the years, a huge number of spin-offs have been created – both official, such as Ubuntu Server, but also unofficial – because it became such an easy base to start from, with an easy installation, easy interface and easy software centre.

Mint

www.linuxmint.com

Based on Ubuntu, Mint took the crown of the most popular distro after Ubuntu moved to a more modern, touch-like desktop, whereas most Linux users wanted a traditional keyboard/mouse design, with a program menu and desktop icons. Mint offers the simplicity of installation, ease of use and the large software base of Ubuntu, but with optimised versions for older, slower computers, and a more fancy version that's known as Cinnamon.

Red Hat

www.redhat.com

It's highly unlikely you'll come across Red Hat as it's the big-business distro used by enterprises and corporations. It is, however, worth mentioning because it funds a couple of high-quality distros that can be freely used. Red Hat is a billion-dollar business and is, unusually, a paid-for Linux distro, but there are ways of effectively getting the same technology with a freely available distro.

Fedora

http://getfedora.com

The Fedora project was created by Red Hat as a test platform for cutting-edge Linux technology. Features are tested in Fedora and, when they're deemed stable enough, they are merged into the Red Hat distribution. This isn't as scary as it sounds, as Fedora only gets these features when they're stable. It's an excellent way of testing the latest technologies because Fedora tends to get them before most other distros.

OpenSUSE

www.opensuse.org

Another business-orientated distro that also has a long-standing pedigree. It might not be the sexiest distro on the block but it's widely used and often ranked in the top five distros, as it's super-stable and can be put to many uses, from standard desktop use to server. This is largely because it's another corporate-sponsored project and is much admired by developers and software vendors.

Mageia

www.mageia.org

Showing you how fluid the Linux distro world is, Mageia was created off the back of a long-running commercial Linux distro called Mandriva, which eventually went bust. From its ashes comes a fresh and powerful desktop distro, which offers a powerful but easy-to-use desktop and complete software library. It stands alongside Ubuntu and Mint for its ease of use yet manages to bring new features, too.

Arch

www.archlinux.org

A super-advanced version of Linux that you almost have to build from scratch. This means you need to be an expert to have any luck with it, but it also means you get an awesome OS that's exactly what you want. A big advantage is that Arch's active community of experts delivers the latest builds of software before anyone else gets them.

Manjaro

https://manjaro.github.io/

Arch Linux but without all the complexity. The Manjaro community has taken the base of Arch and created a pre-built Linux distro that is constantly updated, also called a rolling-release distro. While it does lean to the more technical side of distro releases, the developers have gone out of their way to try to make it accessible to all levels of users.

SteamOS

http://store.steampowered.com/steamos

Here's more of an example of how Linux is used in a fully commercial way. SteamOS is created by Valve Software, the company behind the PC digital distribution system called Steam. SteamOS is based on Debian and is a custom OS that drives a front for the Steam Big Picture mode, integrating controllers, home streaming and the store into one experience that you can enjoy from a box under your TV. ■

Escape from Windows 10

Had enough of Windows? Make the switch to Linux and discover the power and freedom it offers. Time to fling open the escape hatch!

The launch of Windows 10 was a great opportunity for those frustrated with the Microsoft ecosystem to find a better way. And what better way than Linux? If you've just bought a PC with Windows 10 pre-installed or upgraded from an earlier version and decided enough's enough, or have a family member or friend who's finally seen the light, this feature is for you.

The easiest way to migrate from Windows is through Mint, a user-friendly Linux distribution (distro) that's packed with all the features you need, wrapped up in an accessible user interface particularly suited to first-time switchers. You'll find the latest version at **www.linuxmint.com**. We'll be using version 17.3 to step you through the process, but most of the guide can easily be adapted to Ubuntu (either 15.10 or 14.04.3 LTS) if you'd prefer.

The process is straightforward: first, you'll back up your hard drive in its current state, so you know you can always roll back to the beginning of the process if necessary without data loss. Then we'll reveal how to install Mint alongside your existing Windows 10 installation, so you still have access to your old system while you make the transition. We'll sort out any dual-boot issues you have, plus expose the truth behind some myths that put people off switching in the first place.

> "The easiest way to migrate from Windows is through Mint, a user-friendly Linux distro"

We'll also take you on a tour of the new Cinnamon desktop, reveal the basics you need to get started, plus show you how to switch to it if you've installed Ubuntu. Last, but by no means least, we'll show you how to install apps on your new system, find open-source alternatives to the key programs you've been using previously and even reveal how to get some Windows programs running with the help of a clever tool called *Wine*.

Everything you need to make the switch is here, so now there's no excuse. Read on to discover how to escape Windows 10 and transform your computing experience.

Get backed up

Don't even think about installing Linux until you're backed up...

There's nothing worse than attempting to upgrade or install a new operating system only for something to go wrong. You can replace that gut-clenching feeling with one of minor irritation by making sure you've backed up your Windows installation (or just your data) before you begin, this will enable you to easily roll back to a working state if something goes awry.

There are numerous backup tools for the job, but we've plumped for *Redo Backup and Recovery*, which works beautifully with both Windows and Linux installations – go to **www.redobackup.org/download.php** and follow the instructions to download the Redo Backup live CD in ISO format. Burning it to disc is simple: pop a blank disc in the drive, browse to the ISO file in Windows 10, right-click it and choose 'Burn disc image'. If you'd rather create a bootable USB stick, make sure it's been formatted as FAT32, then download and run the latest version of Unetbootin from **http://sourceforge.net/projects/unetbootin** – select the Diskimage option, click '...' to select your ISO file and verify the correct USB drive has been selected before clicking 'OK'.

Before proceeding further, you should also create a Windows recovery disc or flash drive if you don't already have Windows 10 installation media – you may need this should you run into problems with dual-booting Ubuntu and Windows. Type 'recovery' into the Windows 10 Search bar and choose 'Create a recovery drive' to get started.

Once your recovery and *Redo Backup* media has been created, reboot your PC with the *Redo Backup* drive inserted, then tap the option at startup to select your boot device (often [F11] or [F12]). Select the drive and then choose 'Start Redo Backup' when prompted. If you end up back at the boot choice menu then it's likely you're trying to boot from a newer PC in UEFI rather than BIOS mode (*see the Roll back and Start Again box, below for help resolving this issue*). If you're unable to get *Redo Backup* working in any shape or form, try backing up using the built-in Windows disk imaging tool instead, which is part of the *Backup and Restore tool*.

Back up with Redo

When the main screen appears, click the 'Backup' button to get started. Select the drive Windows is currently installed on (typically drive 1) from the drop-down menu and click 'Next'. All of the drive's partitions will be selected – leave this as it is and click 'Next' again.

Next, you'll be prompted to choose your backup destination – typically the largest partition on an external hard drive directly plugged into your PC, so select it from the list, making sure there's adequate backup space available, and click 'Next' again. Once done, click 'Browse' to select or create a folder inside which your backup will be stored. Click 'Save Here' followed by 'Next'. Give your backup a suitable name – the default date is a little confusing, so feel free to change it to something like 'pre-Mint'. Finally, click 'Next' again, then sit back and wait while *Redo Backup* does the hard work of backing up your drive. This can take some time, so if you have another PC handy, why not use the time to prepare your Mint install media following the advice over the page? Once that's done, and the backup completes successfully, you're ready to install Linux.

If something goes wrong with your Mint installation, or you want to start again from scratch and try a different approach, then all you need to do is boot from your *Redo Backup* media and follow the step-by-step guide below, which should put things back exactly the way they were before you started. »

Quick tip

Live disc
A cool trick that Linux distros can pull is to enable you to run an entire Linux OS directly from a DVD called a Live Disc. This means you can test out the distro without installing.

Roll back and start again

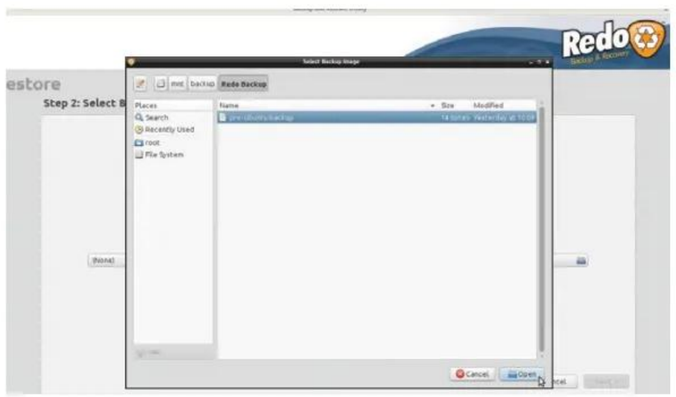

1 Boot and select backup

Boot from your Redo Backup CD or flash drive, clicking 'Start Redo Backup' when prompted. When the main screen appears, click the 'Restore' button. First, select the partition on the physical disk your backup is stored on, then click 'Next'. Click the box to open the 'Select Backup Image' dialog, then browse to the folder you placed your backup in, select the backup job file and click 'Open'.

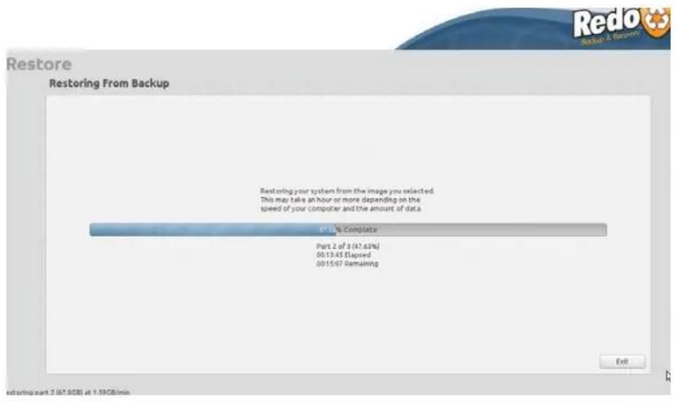

2 Restore PC

Click 'Next', then select the drive that contains your Windows installation. The entire drive will be overwritten with your backup, so read and digest the warning before clicking 'OK'. Wait while the drive is restored to the state it was in when you took the backup. As with the backup process, this can take some time, so be prepared for a long wait.

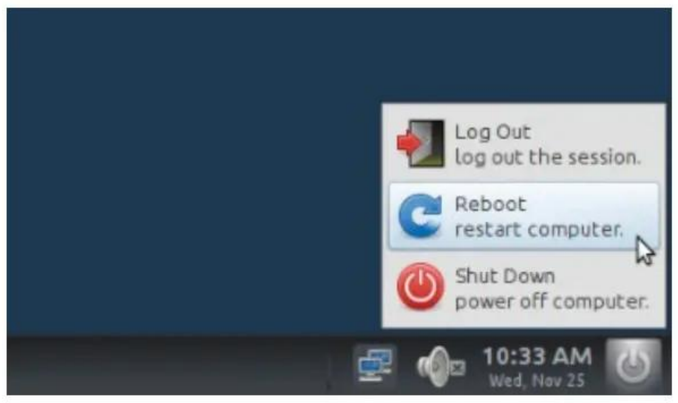

3 Boot and start again

Once the process completes, click 'OK' followed by 'Exit'. Click the power button in the bottom right-hand corner and choose 'Reboot'. Remove the *Redo Backup* boot media when prompted and press 'Enter'. Your system will reboot and you should find Windows loads as it did previously. You can restart the installation process now to correct a previous mistake.

Install Linux Mint

Preparations are complete, it's time to get Mint installed!

With your fail-safe backup in place, you're now ready to take the plunge and install Linux Mint. The following procedure assumes you want to install it alongside your existing Windows installation – it's a good route to follow for first-time switchers, as well as for those helping others to make the migration to Linux.

You'll find the latest version at **linuxmint.com**. Like all versions of Mint, it's a Long Term Support (LTS) release, with support guaranteed until April 2019. If you'd rather try Ubuntu, go to **www.ubuntu.com** and select 'Download'. Click 'Ubuntu Desktop' and download the correct version for your processor (64-bit in most cases) of either Ubuntu 14.04.3 LTS or for the latest release, 15.10. Creating your Mint or Ubuntu install disc or flash drive is the same process for *Redo Backup* (*on the previous page*) and the installation process is identical for both Mint and Ubuntu.

> **"The procedure assumes you want to install Mint alongside your existing Windows install"**

Quick tip

Full install
Mint and Ubuntu will offer to live beside Windows in total harmony, but they can also wipe out Windows for the full Linux experience.

When your PC boots, it makes use of a tiny piece of software to get itself up and running. In the past, this was known as the BIOS (Basic Input Output System), but it's been superseded by the more powerful and flexible UEFI (Unified Extensible Firmware Interface). In practical terms, UEFI offers a more user-friendly interface and better features, but it's not compatible with older hardware.

Most UEFI systems ship with a special Compatibility Support Module (CSM), which allows you to emulate the BIOS when attempting to use older hardware and operating systems. They also ship with Secure Boot enabled, a special feature designed to protect PCs from low-level malware.

Your PC will almost certainly have Secure Boot enabled, which isn't a problem for Mint, which fully supports it. However, another new feature – FastBoot (also called QuickBoot or FastStartup) – will almost certainly be enabled too, and this needs to be disabled before you can install Linux. (*See the Fix UEFI Boot Issues box, below, for how to do this.*)

Once you've edited your UEFI's settings, make sure your installation media is plugged in, then reboot your PC. Look for an option to access the boot menu when you start – typically a key like F11. This will bring up a list of available drives: select your DVD or USB flash drive and hit Enter. Don't panic if you see two entries – this happens when CSM has been enabled, with the non-UEFI option allowing you to boot using the older BIOS compatibility mode. However, Windows 10 will almost certainly have been installed using UEFI mode, so select that option to continue.

Partition your drive

When you arrive at the Mint desktop, double-click the 'Install Linux Mint' icon. Verify English is selected, click 'Continue' and then verify both boxes are checked before clicking 'Continue' again. Wait while the installer scans your hard drives. In a perfect world it should spot your Windows 10 install, but this isn't always the case. If it claims there's no Windows installation, the simplest thing to do is follow the step-by-step guide (*see the Manually Configure your Drive for Dual-Booting box, bottom, p17*) to set up your hard drive manually.

If the installer does detect Windows 10, leave 'Install Linux Mint alongside Windows 10' selected and click 'Next'. Your next step will be to hive off enough drive space to give Mint the room it needs. The installer will make a calculated guess, but you can click and drag the divider between the two to set this figure yourself. Note that if you have more than one hard drive, Mint may actually choose a different drive to install itself on, so verify which one it's picked and click the drop-down arrow next to 'Select drive:' to see if you can change this if necessary. If you can't, and you don't want to install it on the recommended drive, click 'Back' and follow the step-by-step guide to do things manually.

Fix UEFI boot issues

Switch on your PC and tap the key prompted to enter setup – usually [Del] or [F2]. Look for the FastBoot setting under Boot or Windows Configuration, then make sure it's set to disabled. If you have problems getting *Redo Backup* to launch, you should also temporarily disable Secure Boot and potentially tweak the CSM settings in order for *Redo Backup* to work correctly – look under Boot, Security or Windows Configuration for these settings.

When it comes to CSM, its settings may be hidden away in their own menu, where you'll find references to the likes of PXE (pre-boot environment), GOP (graphics) and Storage (or SATA), as well as USB and PS2 if applicable. This allows you to selectively switch CSM on and off for specific devices as required. If there's an option to run both – as is the case with some ASRock motherboards, select the appropriate option (such as 'Do not launch') to ensure maximum compatibility and performance.

Not all PCs, particularly off-the-shell models, will give you access to all of these settings – contact the manufacturer to query where these are, or ask about an update that brings them back. Also search for your PC's model and manufacturer and the words 'CSM', 'UEFI', 'Linux' or 'Mint' to see if there's a precedence set by other switchers you can follow to get the settings you need enabled.

Tweak your PC's CSM settings if necessary to get *Redo Backup and Recovery* working properly if it's a newer model.

Assuming Mint chooses the same drive as Windows 10 to install itself on, leave at least 10GB free space for Windows, and give Mint a minimum of 8GB drive space to work with – more if you plan to make it your primary OS. Remember that you can always repartition the drive later or remove the Windows partition.

Once you've set up the drive, click 'Install Now'. Make a note of the warning and click 'Continue'. You'll be asked to confirm the proposed changes – by default the Mint installer creates just two partitions (the Ext4 entry is your Linux partition, plus it creates a dedicated **swap** partition). Click 'Continue' to accept the change.

Next, the installer will attempt to detect your location, settling on London in most cases. You can accept this or type the name of your nearest large town to see if it's recognised. Click 'Continue', then set your keyboard layout.

The next stage is where you first type your name – you'll see the computer name is pre-filled for you along with your username as you type. You can change these and enter a password – you'll need to provide this when you log in, plus an administrator's password for performing certain tasks.

Only opt to log in automatically if your PC isn't shared, and if you're installing Ubuntu on a laptop tick 'Encrypt my home folder' to protect your data should you lose it. Click 'Continue' and you're done – Mint will install itself in the background. You can go off and make a standard cup of tea (or two) while you wait or sit and read some highlights about what features and applications are included in Ubuntu, and how to find more using the *Software Center*. When installation finishes, click 'Restart Now', remove the installation disc or drive when prompted, and hit Enter. Ubuntu is now installed, turn the page to find out how to start using it.

»

Manually configure your drive for dual-booting

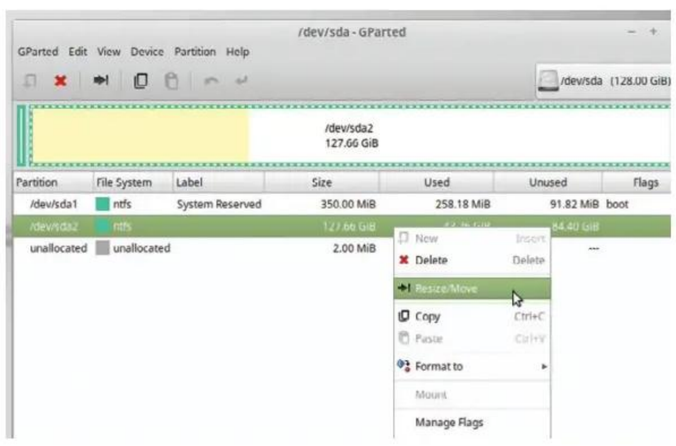

1 Launch GParted

Boot your Mint install media. When the desktop appears, click the 'Menu' button in the bottom left-hand corner and in the Search box type 'GParted'. Click the *GParted* entry when it appears to launch it. Verify your Windows drive is selected (**/dev/sda**) – you should see at least two partitions, the largest of which is your Windows partition. Right-click this and choose Resize/Move.

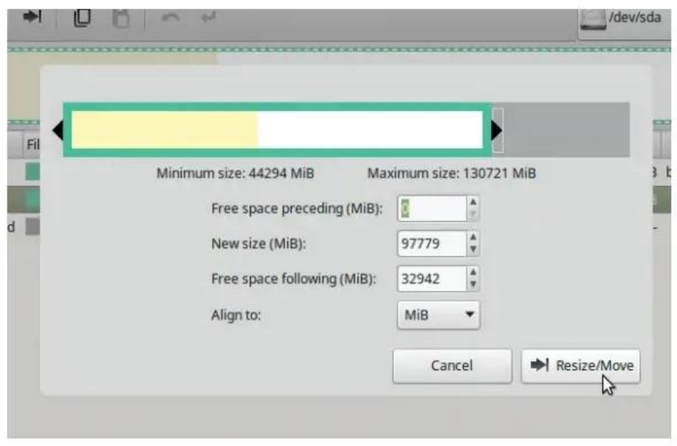

2 Free up space for Linux

Click on the right-hand slider and drag it left to free up enough space for your Ubuntu install – leave at least 10GB free space for Windows. Once done, click 'Resize/Move' followed by the 'Apply all operations' button, selecting 'Apply' when prompted. Wait for the partitioning to complete, then click 'Close'. Double-click the 'Install Linux Mint' shortcut. Select 'English' and click 'Continue'.

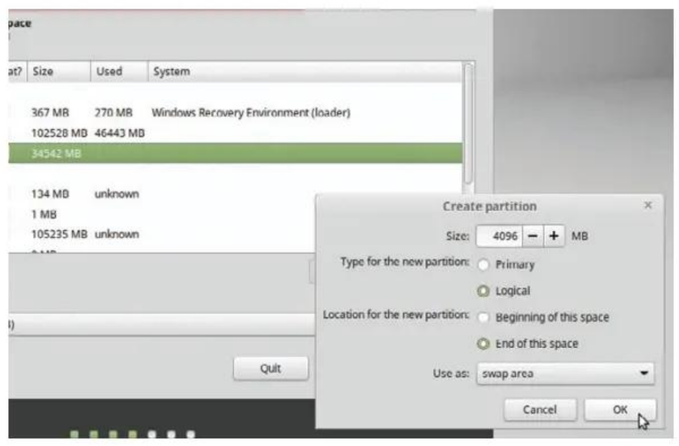

3 Start partitioning

Leave both boxes ticked and click 'Continue'. Select 'Something else' under 'Installation type' and click 'Continue'. Select the free space under **/dev/sda** and click the '+' button to set up your first partition. Set the size to 4096MB, leave Logical selected for the partition type, but select 'End of this space'. Click the 'Use as:' drop-down menu and choose 'swap area' and click 'OK'.

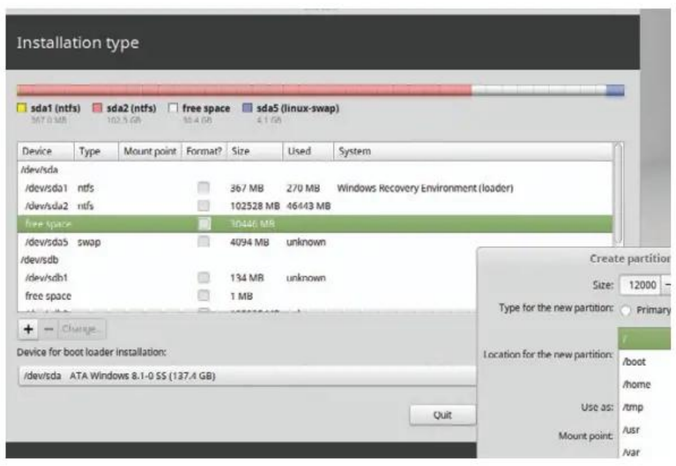

4 Set up Linux partition

Select the free space and click the '+' button again. This time, set the partition size to 12000MB (12GB), leave 'Logical', 'Beginning of this space' and 'Ext4 journaling file system' selected. Click the 'Mount point:' drop-down menu and choose '**/**' to make it the main partition for Ubuntu to reside on (if you have plenty of free space, make it bigger). Click 'OK' again.

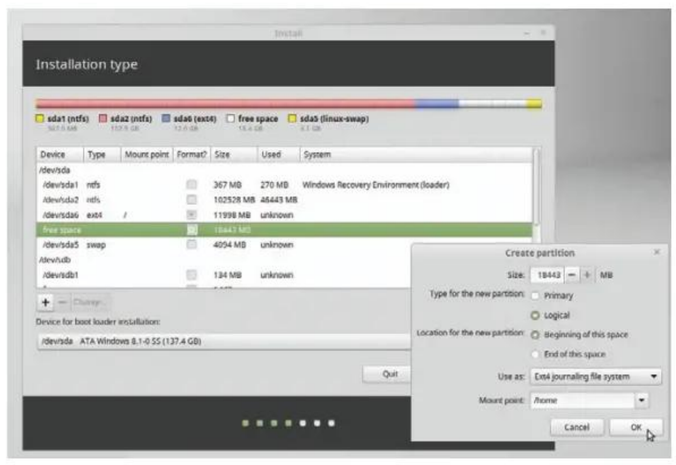

5 Create home partition

Select the remaining free space and click '+' again. Leave everything as it is, except the mount point, which you should point to **/home** before clicking 'OK'. The **/home** directory is where everyone's personal files and settings – including documents, pictures and programs – are stored, so make sure it's the biggest partition of the three. Once done, review your changes.

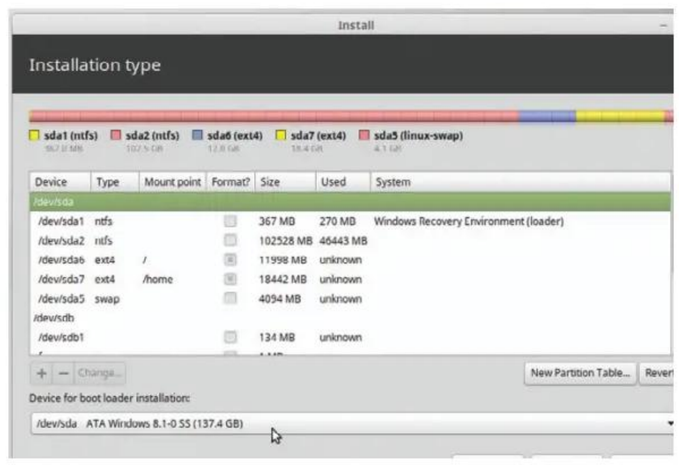

6 Set boot partition

If you need to set up your partitions again, select each of the new Linux partitions in turn and click – to remove them, then return to step three. The 'Device for boot loader installation' drop-down menu should point to your drive (**/dev/sda**), so verify this is the case before clicking 'Install Now'. Review your chosen partition settings and click 'Continue' to apply them and install Mint.

Tweak dual-boot settings

Discover how to switch between Linux and Windows using Grub.

» **When Mint installs itself, it changes the bootloader at the beginning of your hard drive to a program called *GNU GRand Unified Bootloader* (*Grub*). Put simply, *Grub* scans your hard drive for visible operating systems and then presents you with a boot menu, allowing you to switch between them. In most cases – specifically those where your Windows install is detected by the Mint installer – this should be set up to work perfectly from the off. You start your PC, *Grub* displays the boot menu, and if you don't make a choice within five seconds, it'll start to load Linux Mint.**

In this event, the only thing you'll want to know is how to make changes to the boot menu – you may prefer to make Windows the default OS, or you'd like to have a longer pause before Mint is selected by default. You can edit the *Grub* menu in the Terminal, but it's simpler to use a friendlier tool to do the job. (*See the Edit your Grub menu box, below.*)

Troubleshooting Grub

Unfortunately, *Grub* doesn't always do what you expect it to. If the Mint installer doesn't detect your Windows installation and you were forced to manually partition your hard drive, you may find your computer now boots directly to Mint with no sign of Windows. Don't worry, it's still there, but you'll need to perform some additional steps to make it accessible.

First, you'll need to boot from your Windows 10 rescue media: elect to repair your PC then select Troubleshoot > Advanced > Startup Repair to restore the original Windows bootloader. This restores Windows, but leaves you with no way to access Mint – to resolve this, boot from your Mint installation media, then press Ctrl+Alt+T to open Terminal and type the following commands:

```
sudo add-apt-repository ppa:yannubuntu/boot-repair
sudo apt-get update
sudo apt-get install -y boot-repair && boot-repair
```

When Boot Repair launches, choose the recommended repair and follow the steps. The process involves having to copy and paste some commands into the Terminal – use Shift+Ctrl+C to copy them, then right-click in the Terminal and choose Paste – as well as navigating a few additional tools. Press the Tab key to jump between options, then use the cursor keys and Space when selecting options, and finally Enter to select or proceed.

Once done, reboot your PC – hopefully both Windows 10 and Linux are now provided as working options in the boot menu. Note that Windows may not be referenced directly – if an entry refers to the Windows 10 bootloader, try it. This will likely work, in which case you can rename the entry using *Grub Customizer*.

If this doesn't happen (Grub may be present, but only offer to boot into Mint), then the repair tool may have placed boot entries directly into your PC's UEFI boot menu instead of into Grub. Start your PC in the usual way, but tap F11 or whatever key is needed to access the boot menu. You should see an additional option is present – either 'Windows Boot Manager' or 'linux' – selecting this should launch that OS, while choosing the original hard drive entry will boot the opposite OS instead. You can then choose which OS to boot into by default by entering your UEFI setup utility and examining the options under Boot (look for a 'Hard Drive Priorities' section).

If everything gets messed up completely, then simply restore your fail-safe backup using *Redo Backup* and start the installation process again.

Edit your Grub menu

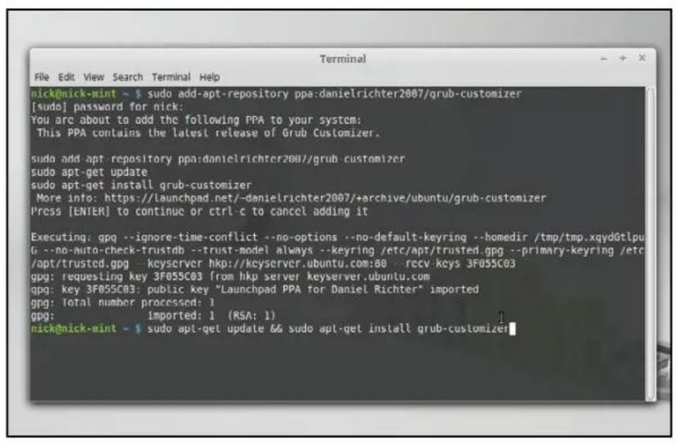

1 Install Grub Customizer

Open a Terminal window and type the following commands:

```
sudo add-apt-repository ppa:danielrichter2007/grub-customizer
sudo apt-get update && sudo apt-get install grub-customizer
```

Once done, launch *Grub Customizer* by clicking 'Menu' and searching for 'Grub'.

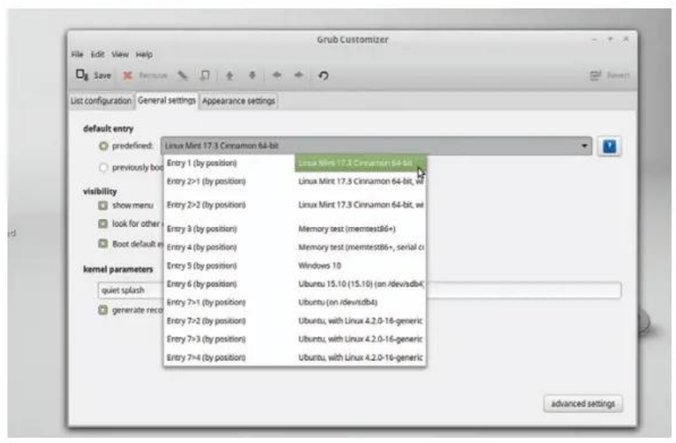

2 Change defaults

Switch to the 'General settings' tab. From here you can change the default OS via the predefined drop-down menu, or select 'previously booted entry'. You can also change the delay before the default entry is selected from this tab too – untick 'Boot default entry after' to pause *Grub* until an OS is selected.

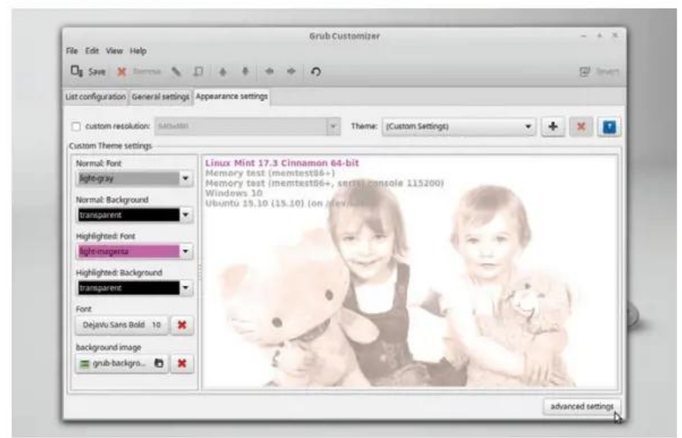

3 Change Grub's look

To make a few changes to how Grub looks, switch to the Appearance settings tab where you can prettify the Grub menu by changing its colour scheme, fonts, default resolution and even apply a nice background image (Note: the image needs to be a PNG). When you've finished making your changes, click the 'Save' button to update *Grub*.

Your questions answered

Still worried about switching to Linux? Let us put your mind at rest.

You're now ready to take your first steps into the Linux universe via the excellent Linux Mint distro, but before you get stuck in, let us explode some of the common myths and help you better understand why switching is a great move.

Q Isn't Linux just for geeks – newbies aren't welcome?
Utter rubbish. These days, Linux – and Mint in particular – is firmly aimed at normal human beings. It's easy to dive in and get started, as you've already found out with the installation process, and is used for a wide variety of professional and personal purposes.

Q Isn't Linux all about the command line?
Most Linux distros ship with a default desktop environment. You can even swap desktops to find the one that suits the way you work best. The underlying OS is still accessible via the command-line Terminal, which works in a similar way to the Windows Command Prompt.

Take the time to learn some of these commands and you'll find the Terminal is a fast and convenient means of getting stuff done, but Mint provides plenty of point-and-click tools, and in most cases you can avoid having to use the Terminal completely.

Q There aren't any decent programs on Linux, right?
Many developers provide native Linux builds of popular programs – check their websites before doing anything else. You'll also find plenty of free, open-source alternatives to favourite applications you've been using in Windows. Linux is renowned for its free, open-source roots, but don't worry, there's plenty of commercial software out there too. Finally, it's also possible to run many Windows programs directly in Linux through the *Wine* utility. We'll show you how later on in this feature (*see p23*).

Q Can I play games on Linux?
Sure you can. You'll find over 1,900 games are available through the Software Manager alone. Select Games from the main screen to browse. Look out for the free DOSBox Emulator, which you can use to play your old Windows games from the last millennium. Also check out Good Old Games (**www.gog.com**), which offers new titles and over 1,300 DRM-free classics made compatible with Linux. The Steam gaming platform (**http://store.steampowered.com**) is also available for Linux where you'll find over 1,500 games – many blockbuster titles – available to buy and download. In fact, Linux is a better platform for games than Windows thanks to the fact it makes fewer demands on your system resources, and developers have come to recognise this, providing native support for renowned gaming engines like the CryEngine.

› **Linux is the perfect platform for playing games on – *Steam for Linux* (and GoG.com) give you access to hundreds of titles, including blockbuster titles, for starters.**

Q Linux won't support my hardware, right?
In the vast majority of cases, you'll find your PC's core components work fine with Mint, while most hardware manufacturers provide dedicated drivers for external peripherals and hardware, from Wi-Fi adaptors to printers and scanners. Indeed, thanks to CUPS, installing and managing printers is a relatively painless.

One area where Linux has really improved is support for graphics cards. Mint installs generic drivers during the installation process, but you can replace these with optimised drivers from your card's hardware manufacturer. Click Menu > Administration > Driver Manager to switch from the open-source driver to a proprietary one if it's supported, or visit the manufacturer's website to manually download the drivers.

Q It's free, so I guess there's no support?
Take a trip to **www.linuxmint.com/links.php** and you'll find documentation and a vibrant, active community offering support and answers through forums and a chat room. Take the time to search on the internet, and you'll find the vast majority of your questions have already been answered.

Q Linux isn't reliable or secure, right?
Wrong – on both counts. Linux is no less reliable than Windows or OS X, and in many ways is more robust. It's also far more secure than its rivals, and not just because of the way it handles user privileges. Linux's open nature makes it better scrutinised, while the diverse number of distros make it less of a target for hackers. »

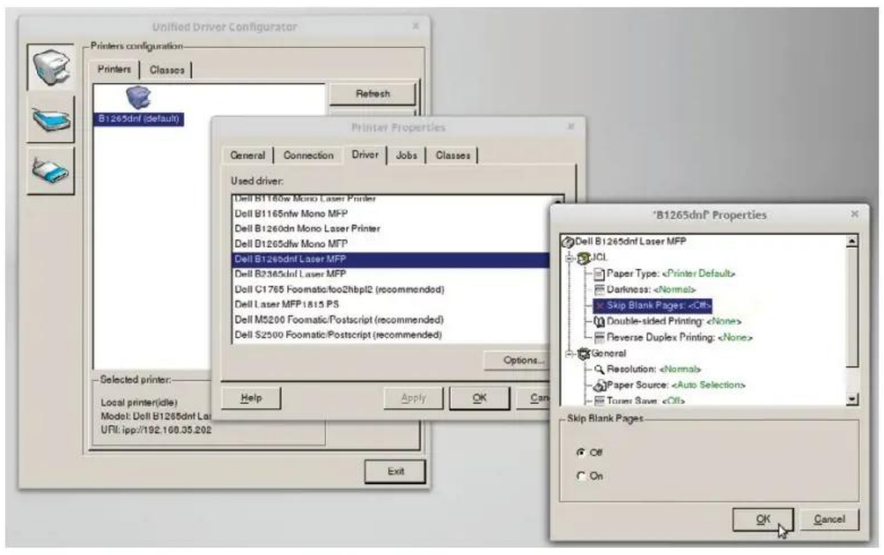

› **Setting up supported printers in Linux isn't tricky – and you have access to all the tools you need in Mint.**

Your first steps in Linux

Discover all the core skills you need to start using Linux Mint.

When you boot into Mint for the first time, you'll find it looks very similar to Windows. The default desktop is called Cinnamon, and is structured in a similar way to the Windows desktop: there's the equivalent of the Start menu, Taskbar and Notification area at the bottom of the screen, with desktop icons providing shortcuts to key locations.

Click the Menu button and up pops a menu with easy access to all your programs, grouped into categories. Speed things up by typing the first few letters of the program you're looking for into the Search box and it'll appear in the list of results, just one click away. On the left-hand side of the menu you'll see some other handy shortcuts – at the top key programs like *Firefox* (web browser), *Software Manager* (for finding and installing new applications), System Settings (the Mint equivalent of the Control Panel), Terminal and files manager. At the bottom are buttons for locking, logging out and shutting down or restarting your PC.

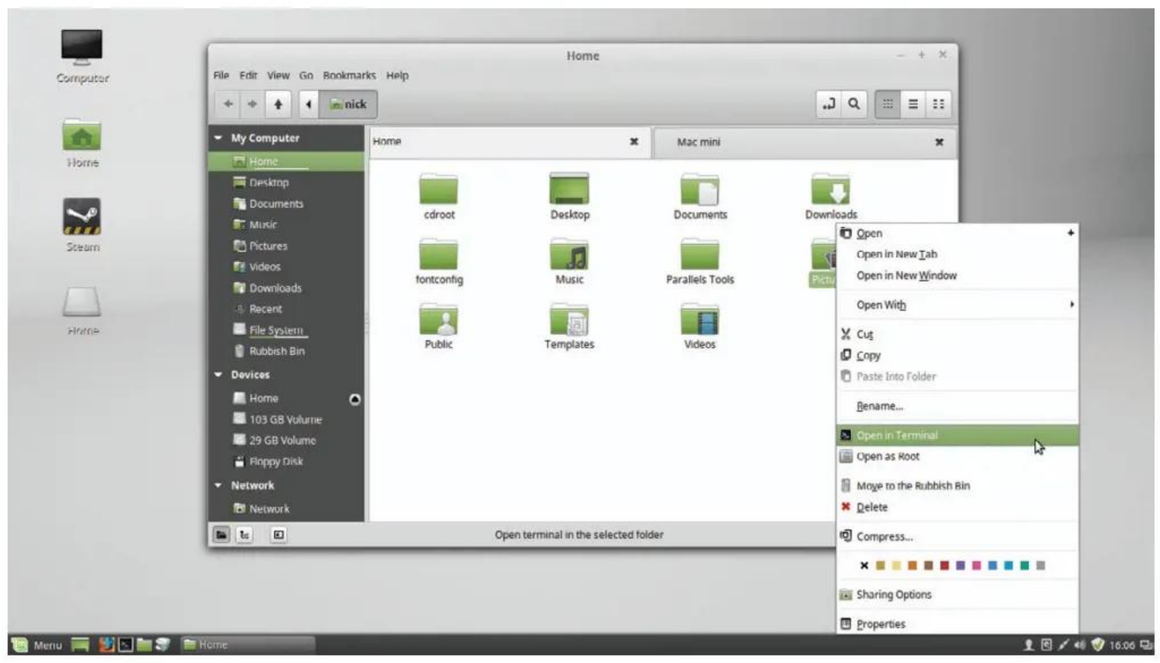

The Nemo file manager works in much the same way as File Explorer does in Windows 10, so you won't struggle to browse your hard drives.

Add your own shortcuts by dragging and dropping items into the list from the desktop or right-hand side of the menu. You can also pin shortcuts to the panel (Mint's equivalent of the taskbar) at the bottom of the screen too – just right-click and choose 'Add to panel' to do so.

Mint's panel is actually more configurable than Windows' Taskbar – you'll see lots of useful options in the bottom right-hand corner, and you can customise these, plus add new panels, simply by right-clicking the panel and selecting from the options on show.

Access drives

Mint works with a wide range of file systems, which means your existing NTFS and FAT32 partitions are visible from the moment you boot into the desktop. The best place to access your drives – and browse your files generally – is with Mint's equivalent of the File Explorer. It's named *Nemo*, and is accessible via the Files shortcuts you'll find on both the pane and the Menu. You can also open it by double-clicking the Computer or Home icons on the desktop.

The layout is similar to Windows 10's File Explorer, with your files displayed in the right-hand pane, and a series of shortcuts to key locations on the left. You can also bookmark favourite folders (select the folder in question and choose Bookmarks > Add to Bookmarks) for easy access from here too. *Nemo* also supports tabbed windows, which work in a similar way to your web browser, albeit with folders. Press Ctrl+T to add additional tabs, which make it easy to switch between multiple drives and folders without cluttering up your desktop with lots of open windows. You'll see an option for accessing your network in *Nemo* – click the Network shortcut in the left-hand pane to find other computers, and have your username and password ready if necessary.

You should find your PC's hardware has been – for the most part – automatically detected and set up. Click the

Share data with Windows

If you plan to continue using Windows and Ubuntu alongside each other, it pays to set up a dedicated **data** partition, then make it easily accessible to both operating systems. Use *GParted* from the Mint installer to set up a NTFS partition, then boot into Windows and browse to your user folder. Right-click each key data folder – Documents, Pictures etc – in turn and choose Properties > Location, then point them to the equivalent folder on your new data partition, moving the data across when prompted.

You should also disable hibernation in both Windows and Linux – data written to the shared partition while one OS is in hibernation may be wiped when that OS is next woken up. In Windows, open a Command Prompt (administrator), type `powercfg /hibernate off` and press Enter to enable this setting.

Once done, switch back to Mint. Disable hibernation by opening System Settings > Power Management and ensuring 'Never' is selected under the second drop-down menu. Next, open the *Disks* utility via the Dash. Select the disk containing your data partition in the left, then click on the partition on the right and make a note of where it's mounted (this tends to follow the format **/media/<username>/<partition name>**).

Finally, go to **www.ubuntu-tweak.com** and install the Ubuntu Tweak (yes, parts of it work in Linux Mint!). Open it from the Administration section of the Menu, then select Admins > User Folder to manually point your data folders to the folders stored on the data partition you created and identified earlier.

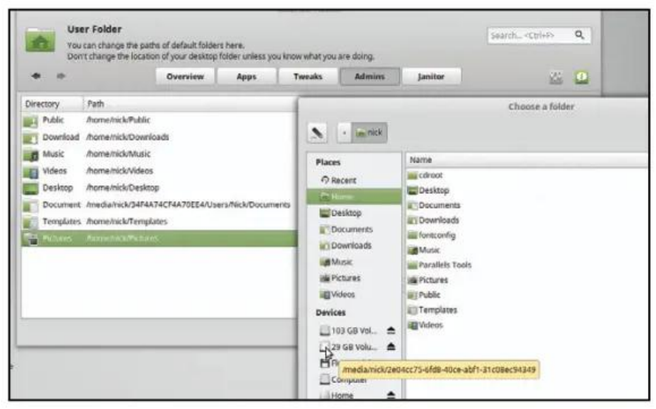

Share data between Windows and Linux using a dedicated NTFS partition.

Menu and choose System Settings to review current settings under the Hardware section. You can source missing drivers from the manufacturer's website if necessary.

Now you've dipped your toes in the water and discovered it to be not that deep (or cold), take your first system backup of your new dual-boot system using *Redo Backup*. Once complete, you can start to dig a little deeper into Linux. Let's start with the file system – Windows uses NTFS; there's the universal FAT32 file system, of course, and now you've got a third file system, Ext4, in place. When you boot back into Windows you'll find your Linux partitions are invisible – that's because Windows can't see Ext4. That's why your shared data partition – if you created one – was formatted as NTFS.

Open the *Nemo* file manager again and you'll find yourself inside the **Home** folder, which – if you manually configured your partitions when installing Mint – may be a separate partition. Like your Windows user folder, this contains all your personal data, but it also contains personal programs and settings too.

Click the Computer shortcut in the left-hand pane and you get to see Ubuntu in all its messy glory. Many of these folders can be ignored, while two worth noting are **/media** and **/mnt**. These are where shortcuts to all external drives and network folders are found. The key difference is that drives mounted to **/media** are done so on a per-user basis.

Key system tools

The following utilities help you tweak various parts of Mint without having to use the Terminal. Click System Settings to get a dialogue box similar to Windows' Control Panel. Here you'll find lots of useful tools for tweaking your installation: The Appearance section is where you customise Cinnamon's look, but also look under Preferences for areas where you can change the Cinnamon desktop's behaviour. If you select the Windows > Behaviour tab and flick the 'Attach dialogue windows to the parent window' switch 'Off' you can change the way Mint handles each separate dialogue box that's opened if you prefer the way Windows does it.

Mint also includes a number of handy built-in tools – there's the Disks tool (*which we touch on in the 'Share data with Windows' box, p20*), plus *Disk Usage Analyser*, which enables you to see a graphical breakdown of how your drive space is being used. The *System Monitor* tool is, essentially, the equivalent of Windows' Task Manager, and the *Update Manager* is where you go to keep Linux up to date.

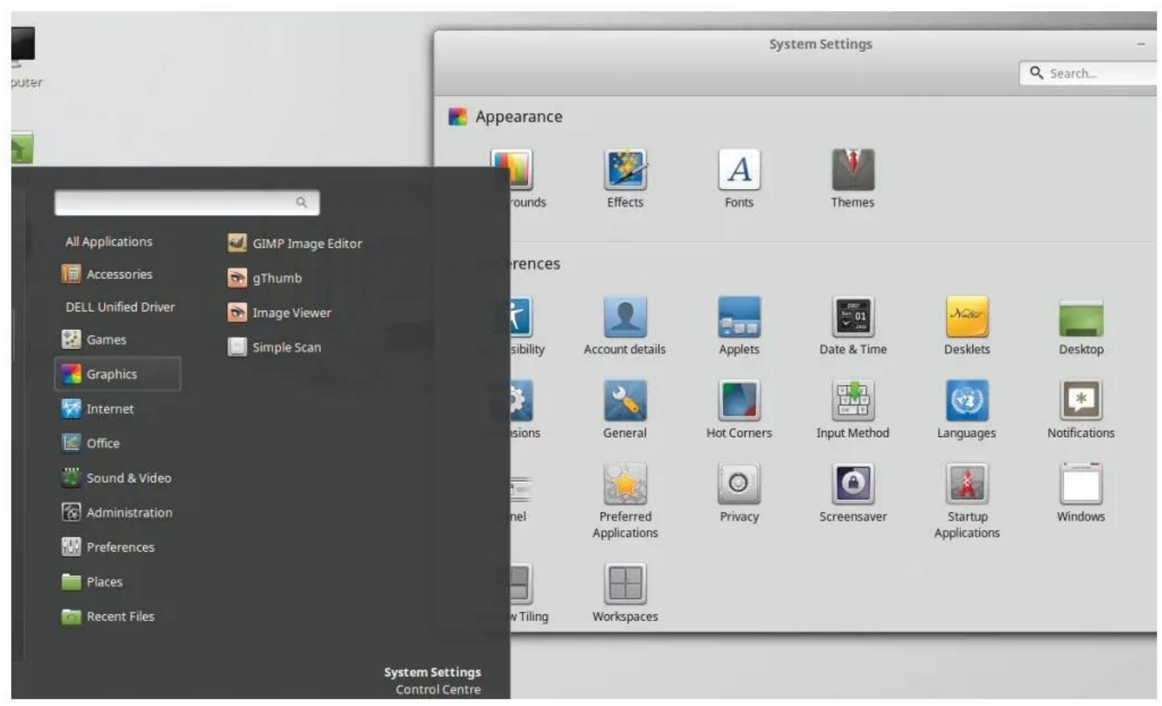

› **Take more control over your Cinnamon desktop using the System Settings tools**

The Terminal

At some point you'll need to access the Terminal. And once you've done so, you'll quickly realise it's the best way to do certain things. At the very least, familiarise yourself with the basics with this handy list of commands and you can press tab to complete commands and even file/folder names:

» `dir` – lists the contents of currently selected directory.

» `cd` – for changing directory – use `cd ..` to move up a level, `cd foldername` to move to a subfolder within the current directory or `cd /home/<username>/Downloads` to jump to a specific directory. Note that folder names are case sensitive. To jump back to your **home** directory, type `cd ~` at any time.

» `cp` – copy files. Other file-management commands include `mv` (move), `mkdir` (create folder) and `rm` (delete file).

» `sudo` – performs a task as an administrator. This is the equivalent of right-clicking a program in Windows and choosing 'Run as administrator'. The command precedes other commands (eg, `sudo apt-get install syncthing`), and you'll need to provide your password the first time you use it.

» `apt-get` is a tool for managing and installing packages. Use `apt-get install` to install a package, `apt-get upgrade` to update all existing software, and `apt-get check` to look for problems with software (known as broken dependencies).

» `wget` downloads files to your **Downloads** folder: eg `wget http://address.com/filename.gzip`.

To find out more about the Terminal, check out the beginner's guide at **www.ryanstutorials.net/linuxtutorial**. »

Quick tip

Drive names

A confusing point when moving to Linux is where have all the drive names gone, such as the C: drive and USB drives? Mint does a fine job of presenting removable media, but Linux doesn't use drive letters and has a more complex way of organising partitions and drive names.

Change desktops

One of Linux's strengths is that you're not tied to a single desktop. If you've installed Ubuntu rather than Mint you'll find it comes with a very different desktop environment in the form of Unity. It's still reasonably familiar to use, but you will have to make a bigger effort to master it.

If you're struggling to acclimatise to Unity, and would rather use Mint's Cinnamon desktop instead, then you can easily do so. First, open a Terminal window and type:

```
sudo add-apt-repository ppa:lestcape/cinnamon
sudo apt-get update && sudo apt-get install cinnamon
```

This will add the Cinnamon repository (repos) to Ubuntu, then download and install it. Once installed, you need to click the 'Settings' button in the top right-hand corner of the Unity desktop and choose 'Log out' to return to the login screen. You'll see a little Ubuntu icon appear next to your username – click this and a list of available desktops will appear, including Cinnamon. Select that to boot into your new desktop, complete with all the features we describe on the rest of this page. You can always switch back to Unity from the login screen.

If you subsequently outgrow the Cinnamon desktop, then you can easily switch desktops in Mint too. You can even install multiple desktops and move between them as circumstances dictate.

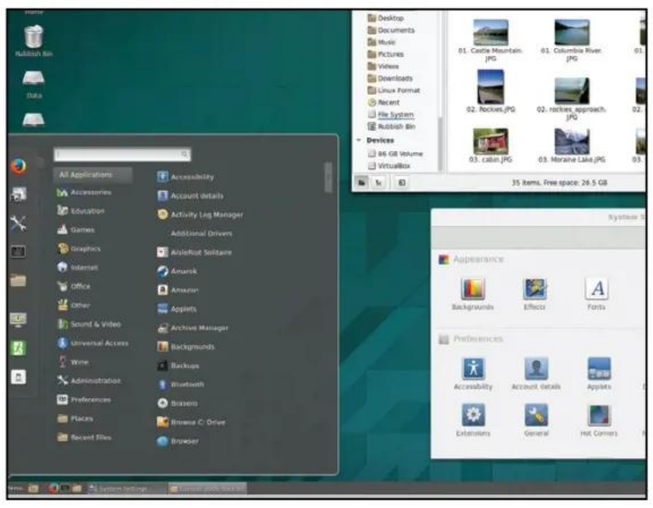

› **Cinnamon looks a little different in Ubuntu, but works exactly the same way.**

Install apps in Mint

How to find and add new programs to your distro.

Now that you're up and running in Mint, you'll be wanting to start making productive use of it. That naturally involves running applications and utilities. Start by taking an inventory of the programs you use in Windows and checking the website for each to see if a Linux build is available. If there isn't one, don't panic – you'll find the vast majority of Windows applications have an equivalent version for Linux. Better still, many of these are both free and open source, so you won't need to splash out to use them, but do of course donate to them.

The most obvious place to start is your office suite. When it comes to replacing your word processor, spreadsheet, database and presentation tool, you'll find everything you need is already built into Mint via the free LibreOffice software, which can be found under Menu > Office.

If you're looking for a specific open-source alternative to your original proprietary program, start your search at **http://bit.ly/SoftwareEquivalents** where you'll find links to other useful resources such as **www.osalt.com**. If you still can't find anything suitable, try **http://alternativeto.net**.

The simplest way to obtain and install software in Mint is through the *Software Manager*. It offers a neat, user-friendly front-end for finding and installing programs. Programs are organised into 'repositories' (or repos) and if your target program isn't in the repos supported by the *Software Manager* you'll have to find alternative ways of getting them.

First, visit the program's home page – here you may be offered a convenient downloadable installer in the form of a Deb package – choose the 32-bit or 64-bit download (if offered) to match the version of Ubuntu you installed. Once downloaded, navigate to your **Downloads** folder, double-click the file and follow the prompts to install it.

> "You'll find the vast majority of Windows applications have an equivalent version for Linux"

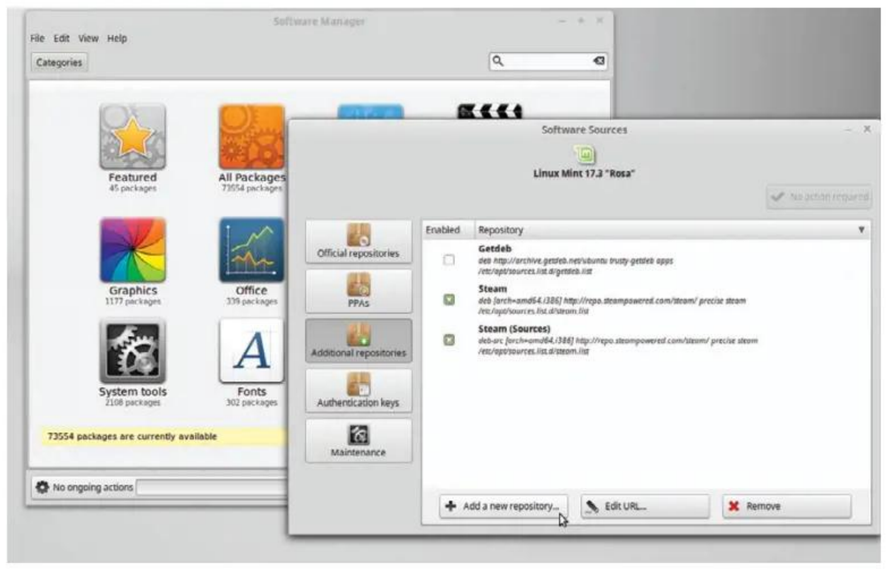

› Manage all your repositories through the Software & Updates tool, which you'll find in System Settings.

Other programs may require you to install them through the Terminal. This usually involves first adding the program's repos using the following command:

```
sudo add-apt-repository ppa:user/ppa-name
```

The `ppa:` stands for Personal Package Archive, which is a special software repos – generally focused on Ubuntu – that anyone, usually a software team, can establish for uploading source packages for anyone to download. Replace `ppa:user/ppa-name` with the PPA provided by the software producer. Once done, you can then install packages from that repos, either through *Software Manager* or – seeing as you're already in the Terminal – with the following command:

```
sudo apt-get update && sudo apt-get install <program>
```

Replace `<program>` with the name of your target program and it'll be downloaded and installed. Once added, you can view and manage all repos via System Settings > Software Sources.

Stay up to date

Programs you install through repos – either via the Terminal or *Software Manager* – are checked for updates on a regular basis, namely within 10 minutes of first booting up and then every two hours thereafter. That should be frequent enough for the vast majority of people, but you can check manually too using one of two methods. The first is done with Cinnamon: open Menu > Administration > Update Manager and click 'Refresh' to manually check here, and choose Edit > Preferences > Auto-Refresh tab to change the time between update checks (days, hours and minutes are supported).

To check for updates via the Terminal, type `sudo apt-get update` and hit Enter. You may find the application versions in the *Software Center* aren't as up to date as those provided on the program's website – by adding the relevant repos, you can ensure the latest version of the program is installed, and gain access to beta and other pre-release builds.

Popular alternative applications

Windows app	Alternative	Where to get it
Word, Excel, PowerPoint, Access	LibreOffice	Pre-installed
Publisher	Scribus	Software Manager (SM)
Photoshop	The GIMP	Pre-installed
Windows Movie Maker	OpenShot	Software Manager
Adobe Reader	Document Viewer	Pre-installed
iTunes	Amarok	Software Manager
Windows Media Player	VLC Media Player	Software Manager
Internet Explorer	Firefox, Chromium	Pre-installed or SM
Outlook/Windows Live Mail	Thunderbird, Evolution	Pre-installed or SM
Skype	Pidgin	Software Manager

Install Windows apps

If all else fails, use Wine to run Windows programs in Linux.

The best way to run applications in your new Linux Mint install is natively, but if there's a specific Windows program you can't do without then it may be possible to get it running. What you'll need is a special program called *Wine*, which uses some clever trickery to use parts of Windows, enabling you to run many Windows programs as if they were native Linux ones.

Wine can be a tricky tool to master, but thanks to the free *PlayOnLinux* tool, you can get many popular applications up and running with the minimum of fuss. Open the *Software Manager*, then search for 'PlayOnLinux'. Click 'Install' and it'll automatically download and install *Wine*.

Once installed, launch *PlayOnLinux* and allow it walk you through the first-use wizard. You'll be left with a blank window, so click the 'Install' button to get started. The Install menu allows you to search for programs or browse by category, with options including Games, Graphics and Office. Once you've found a program you want to install, select it and click the 'Install' button. You may need to supply the program's installation disc, while commercial applications will require a licence and – if applicable – activation. Just follow the wizard through – *PlayOnLinux* simplifies the process of installing any supported applications because it knows which additional components need downloading and how to configure them.

PlayOnLinux's supported applications database is relatively small – aside from *Microsoft Office*, notable examples include Spotify, *QuickTime Player* and *Adobe Photoshop* (CS4 or CS6 only). That's because it primarily focuses on games. If you find the application you want isn't listed, don't worry – *Wine* supports more programs than those listed by *PlayOnLinux*, but the install process is a more complicated. Start by checking **https://appdb.winehq.org** to see if the application is supported. If it's present, you'll see a list of tested versions, complete with rating ranging from Platinum all the way down to Garbage. Anything rated Bronze or better is likely to work in *Wine*, although you may find some glitches occur. Once done, follow the step-by-step guide (*see below*) to attempt installation through *PlayOnLinux*'s manual option. If you run into problems, consult the web pages for more advice, or try experimenting with different settings, but hopefully you'll end up with an application that does at least perform most key functions without problems.

Application configuration

Each program exists inside its own virtual space, which you can tweak by selecting it on the main *PlayOnLinux* screen and clicking 'Configure'. Most advanced configuration options are found under the Wine tab – click 'Configure Wine' to access a multi-tabbed dialogue box with lots of additional tweaks.

Wine should automatically map your **home** folders to their equivalent locations in the virtual drive, but make other locations more easily accessible from the Drive tab. Click 'Add', select a drive letter and then browse to the directory. Look under the **/media** folder for separate drives. Once selected, it'll now be accessible within the program via the drive letter you allocated to it.

If there's no mention of it at in the *Wine* HQ database, don't assume it won't work: you can't harm Linux, so feel free to try, using the program's website to determine what additional components (such as the .NET Framework) you might need. All of this should get you off to a solid and happy start to using Linux. Keep exploring and have fun!

Manually install an app using PlayOnLinux

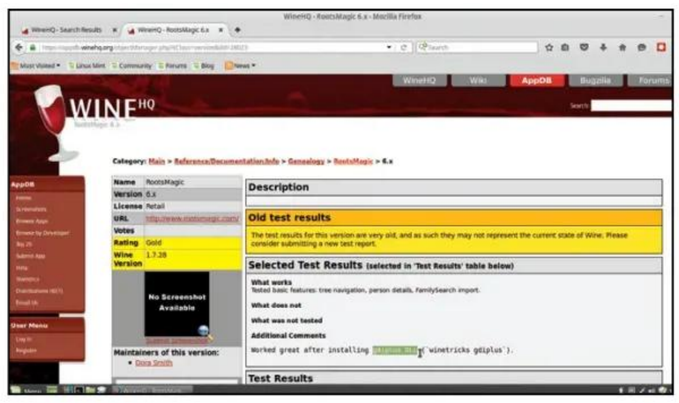

1 Take notes

Go to **https://appdb.winehq.org**, find the program you wish to install and verify the version you're installing is compatible (if it's not listed, but earlier versions are compatible – experiment!). Also make a note of any comments about additional installation requirements, such as gdiplus.dll. Now visit the program's website and download the program installer if you don't have an installation disc.

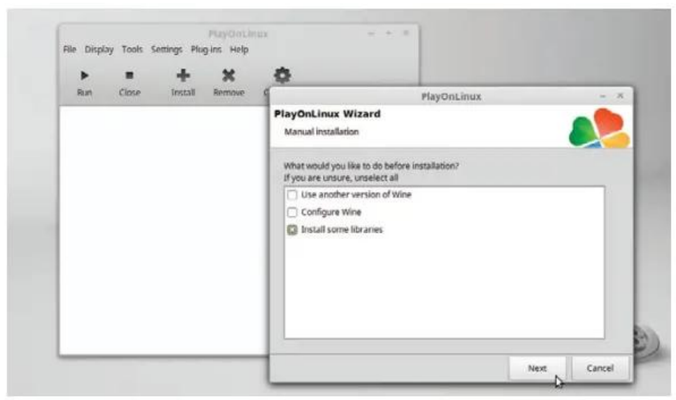

2 Set up virtual drive

Open *PlayOnLinux*, select Install and then click 'Install a non-listed program'. Click 'Next', select 'Install a program in a new virtual drive' and click 'Next' again. Name the drive after your program, click 'Next' and then tick the relevant options – in most cases just 'Install some libraries' before clicking 'Next'. Leave the 32-bits option selected for compatibility reasons and click 'Next' again to proceed with the next section.

3 Add libraries and install

Install any recommended packages when prompted. Once done, place a tick next to each required component (typically under POL_install) and click 'Next', then wait for them to download and install. Next, either select your CD drive or choose 'Select another file' to locate and run the installation program. Follow the setup wizard through, then launch the program to verify it works before adding shortcuts to finish.

PICK A DISTRO

If you're diving in to the world of Linux, make sure you pick the right package for you: different distros make a difference...

Your favourite Linux distribution isn't an individual unit in itself. On the inside, it is made up of various apps, libraries, modules and toolkits. On the outside, it's part of a much larger and very vibrant ecosystem that sustains several other distros.

Also part of this larger Linux ecosystem are very active user communities and various support infrastructures that help nourish the distros and other projects. Over the course of their lifetime, the different elements of the Linux ecosystem interact with each other as well as with their environment. They collaborate, exchange ideas and features, and sometimes swap resources for their mutual benefit as well as for the enhancement of the ecosystem.

The Linux ecosystem fosters the development of innovative projects and products. However, since the environment can only sustain only so many elements, the distros go through an evolutionary process of their own to weed out the weaklings and ensure the survival of the fittest. Through this process, the uninteresting, dull and unsustainable projects begin to perish. However, the strong ones survive, thrive and pass on their genetic code to the next generation of derivative distros.

> "The strong survive, thrive and pass on their genetic code to the next gen of derivative distros"

In this feature we will classify the popular distros as per their origins. We will analyse how they've evolved since their inception, and look at the unique traits they have passed on to derivatives that help distinguish them from their peers. We'll also look at the best distro from each distro genus* and then pit them against each other to help you pick the distro that is right for you. »

*Genus: A rank used in the biological classification of organisms.

Genus Debian

Made of free software and evolving...

» **The Debian project has played a significant role in the evolution of Linux and, in many ways, is the first real distribution created for the regular computer user. It was announced in August 1993 and had its first public release later that year, although its first stable release wasn't available until 1996. The project was even sponsored by the Free Software Foundation from November 1994 to November 1995.**

A key motivating factor that led Ian Murdock to create a new distro was the perceived poor maintenance and prevalence of bugs in the Softlanding Linux System (SLS) distro. Besides the software itself, Murdock's release included the Debian Linux Manifesto, which outlined his view for the new project, which prophesied that "distributions are essential to the future of Linux". In the Manifesto, he called for the distro to be maintained openly, in the spirit of Linux and GNU. One of the most significant goals for the distro was to "eliminate the need for the user to locate, download, compile, install and integrate a fairly large number of essential tools to assemble a working Linux system." In order to meet this goal, Debian developers made a significant contribution to the Linux world – the *dpkg* package manager.

This was originally written as a Perl program by Matt Welsh, Carl Streeter and Ian Murdock, and the main part of the tool was rewritten by Ian Jackson who became Debian Project Leader in 1998. It really is no surprise then that Debian is one of the most popular choices for derivative projects with over 130 active distros based on Debian (Source: **http://distrowatch.com**), including the likes of Ubuntu and a version of Linux Mint.

> **"The Debian Linux Manifesto prophesied that "distributions are essential to the future of Linux"**

The project also provides guidelines to help the derivative distros merge their work back into Debian. In addition to the derivatives there are several 'Pure Blends'; these are subsets of Debian configured to support a particular niche, such as Debian Edu, Debian Junior and Debian Med. Debian also supports a variety of platforms, including Intel i386 and above, Alpha, ARM, Intel IA-64, Motorola 68k, MIPS, PA-RISC, PowerPC, Sparc and more.

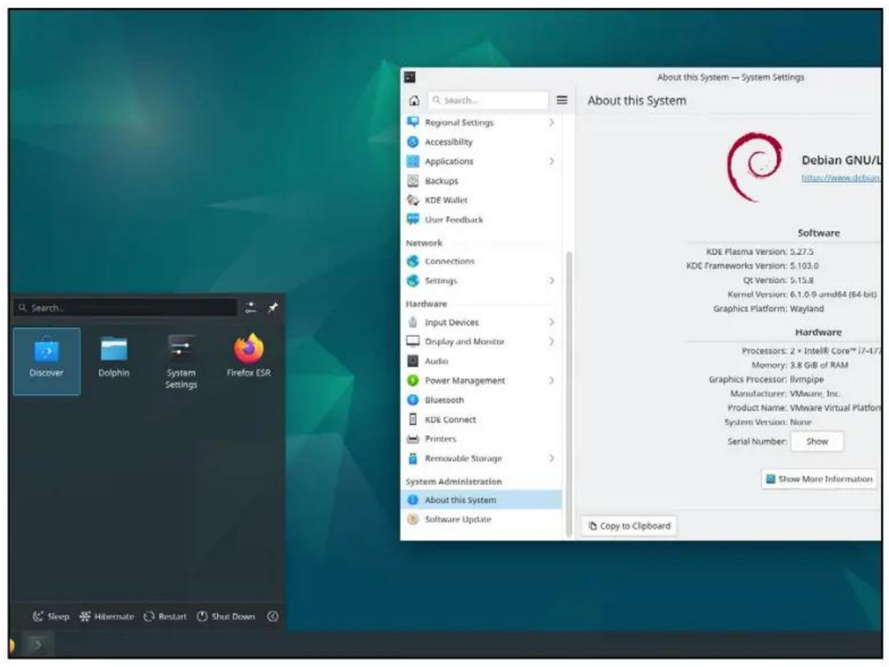

› **Even the most popular distribution for the Raspberry Pi called Raspbian, is based on the Debian Project.**

Rules of engagement

Another distinguishing aspect of Debian is that the distro is made entirely of free software. The project uses the Debian Free Software Guidelines (DFSG) to help determine whether a piece of software can be included. The DFSG is part of the Debian Social Contract which defines the moral agenda of the project.

The project produces three distros: Stable, Testing and Unstable. A Stable release is available every two years and is made by freezing the Testing release for a few months. Testing is designed to be the preview distro with newer packages and during the freeze any bugs are fixed and extremely buggy packages are removed. All releases are named after characters from the *Toy Story* films (the current Stable release is codenamed Jessie). All new packages are introduced in the Unstable release (codenamed Sid). This distro is for developers who require the latest packages and libraries. It's not intended to be used on a production machine and those interested must upgrade Debian Testing to get the latest Unstable. »

BEST OF BREED

Linux Mint Debian Edition

The Linux Mint Debian Edition (LMDE) is meant for users who wish to experience the best of Debian (directly rather than via Ubuntu) in an easy to use package. It's based on Debian Testing and is a semi-rolling release, which means it receives periodic updates via Update Packs. These are tested snapshots of Debian Testing to ensure stability and LMDE is binary compatible with Debian, which means you can switch to Debian Testing or Unstable for more frequent and bleeding-edge updates. However, LMDE isn't compatible with Linux Mint, so you can't use Ubuntu PPAs.

LMDE is designed to offer the same look and functionality of Linux Mint and is available as 32-bit and 64-bit Live DVD images with either the Mate or Cinnamon desktops. The distro ships with *Firefox*, *Thunderbird*, *VLC* media player and a plethora of other commonly used apps. Adobe Flash plugin and most other multimedia codecs are installed by default. The software repos and the underlying Deb package system makes software installation easy, thanks to tools, such as the Synaptic Package Manager.

Genus Ubuntu

Derivatives, they're coming outta the walls.

» **Ubuntu is, in many respects, the first distro to make a serious effort to bring in new users. The distro brought Linux into the mainstream, played a significant part in changing the notion and misconceptions about Linux and was able to successfully pitch itself as a viable OS alternative to Windows and mac OS.**

Ubuntu was started by Mark Shuttleworth. He formed Canonical after selling his security-firm, Thawte, to VeriSign. Shuttleworth was a huge fan of the Debian project. However, there were many things about Debian that didn't fit in with Shuttleworth's vision of an ideal OS. He therefore invited a dozen or so Debian developers he knew and respected to his flat in London in April 2004 and hashed out the groundwork for the Ubuntu project.

The group decided on a bunch of characteristics for the distro. For one, Ubuntu's packages would be based on those from Debian's unstable branch. However, unlike Debian, Ubuntu was to have a predictable cycle with frequent releases. To put the plan into action, it was decided that Ubuntu would release updated versions every six months and each release would receive free support for nine months. The plan was refined in later years and now every fourth release receives long-term support (LTS) for five years.

The group also decided to give emphasis to localisation and accessibility in order to appeal to users across the world. There was also a consensus on concentrating development efforts on ease of use and user-friendliness of the distro on the desktop. The first release of Ubuntu was in October 2004.

Ubuntu's development is funded by Shuttleworth's UK-based Canonical, which is a privately held computer software company. The company also supports development of other Ubuntu-related projects, for instance, Ubuntu's *Ubiquity* installer is one of the best tools for the job, and one of its distinguishing features is that it gives users the option to install closed source or patented third-party software, such as Fluendo's MP3 codec. Other useful user-centric projects that have tried to change the status quo are the *Ubuntu Software Center* and the recently discontinued Ubuntu One cloud hosting service.

Test by fire

But perhaps no other piece of technology has polarised the Linux community like Ubuntu's Unity desktop interface. The distro first introduced Unity with the Ubuntu Netbook Edition version 10.10. By the time 11.04 rolled off the press, the Netbook Edition had merged into the desktop edition and Unity became the default graphical interface for the Ubuntu distro. However, Shuttleworth has insisted that the Unity desktop plays a crucial role in Ubuntu's multi-device strategy. Unity will help standardise the display on smartphones, tablets, TV and other devices beyond the computer.

Thanks to its malleable nature, the distro has always been very popular with developers who want to create a custom distro for their particular niche. Ubuntu has perhaps seeded more distros than any other, and Ubuntu itself has several officially-supported spins: Kubuntu, Xubuntu, Ubuntu Gnome, Edubuntu and Ubuntu Studio. In addition to the main desktop edition, there's also a server edition that doesn't ship with a graphical desktop.

Ubuntu has helped give Linux mainstream coverage and has several celebrity users, including Cory Doctorow and Stephen Fry. However, pushing the envelope has its drawbacks and the award-winning distro has had its fair share of brickbats. It's still reeling under the Amazon controversy that arose when the distro included search results from the shopping giant in Unity's Dash whenever users searched for stuff on their computer.

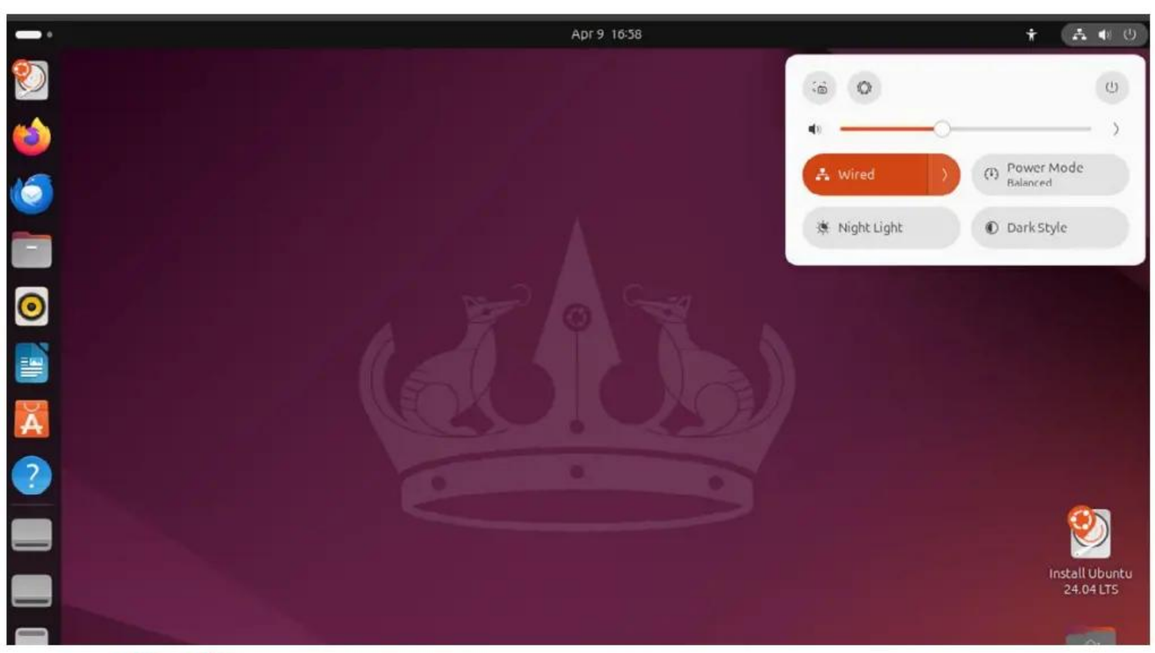

› **A number of vendors, such as Dell and Lenovo, offer computers pre-installed with Ubuntu.**

Trisquel GNU/Linux

Trisquel GNU/Linux goes to great lengths to do justice to its free software tag. Not only does the distro not include any proprietary software, it also strips out all non-free code from the components it inherits from Ubuntu, such as the kernel. Instead of the stock Ubuntu kernel, Trisquel uses the Linux-libre kernel that doesn't include any binary blobs. Thanks to its efforts, the distro has been endorsed by the Free Software Foundation.

There are several variants of the distro, the most common ones are the standard Trisquel release, which is available as a 700MB image with the Gnome desktop, and Trisquel mini, which is designed for older hardware and low-power systems, and uses LXDE, the lightweight desktop.

While the distro doesn't ship with any proprietary codecs, you can watch YouTube videos as it provides HTML5 support as well as Gnash, which is the free alternative to Adobe Flash. Trisquel includes all the usual desktop productivity apps, such as *LibreOffice*, *Evolution*, *Gwibber*, *Pidgin* and more. These are complemented by an impressive software repository.

Genus Red Hat

Millinery on a massive scale.

Another distribution that has played a crucial role in shaping Linux's DNA is Red Hat Linux, which was created in 1994 by Marc Ewing. Bob Young and his ACC Corporation bought Ewing's business and created Red Hat Software. The company went public in 1999 and achieved the eighth-biggest first-day gain in the history of Wall Street. It rode on the success of Red Hat Linux to became the first open source billion dollar company.

Over the years, some of the biggest and brightest Linux developers have worked with Red Hat. Soon after it went public, it acquired Michael Tiemann's Cygnus Solutions which had authored the GNU C++ Compiler and worked on the GNU C Compiler and the GNU Debugger.

One of Red Hat's most influential pieces of technology is its RPM packaging format. The file format is now the baseline package format of the Linux Standard Base (LSB), which aims to standardise the software system structure, including the filesystem hierarchy used in the Linux operating system. The LSB is a joint project by several Linux distros managed by the Linux Foundation. Red Hat was also one of the first Linux distros to support Executable and Linkable Format (ELF) instead of the older a.out format. ELF is the standard file format for executables, shared libraries and other files. Red Hat was also the first distro to attempt to unify the look of its Gnome and KDE desktop with the Bluecurve theme – which caused tension with the KDE developers. The distro has won laurels for its easily navigable graphical Anaconda installer.

Life after death

Initially, the Red Hat distro was offered as a free download and the company sustained itself by selling support packages. In 2003, however, Red Hat discontinued the Red Hat Linux distro and it now focuses solely on the Red Hat Enterprise Linux (RHEL) distro for enterprise environments. RHEL supports popular server architectures including x86, x86-64, Itanium, PowerPC and IBM System z. The lifecycle of newer RHEL releases spans 13 years, during which time the users get technical support, software updates, security updates and drivers for new hardware. Red Hat also has a very popular training and certification program called RHCP that's centred around RHEL.

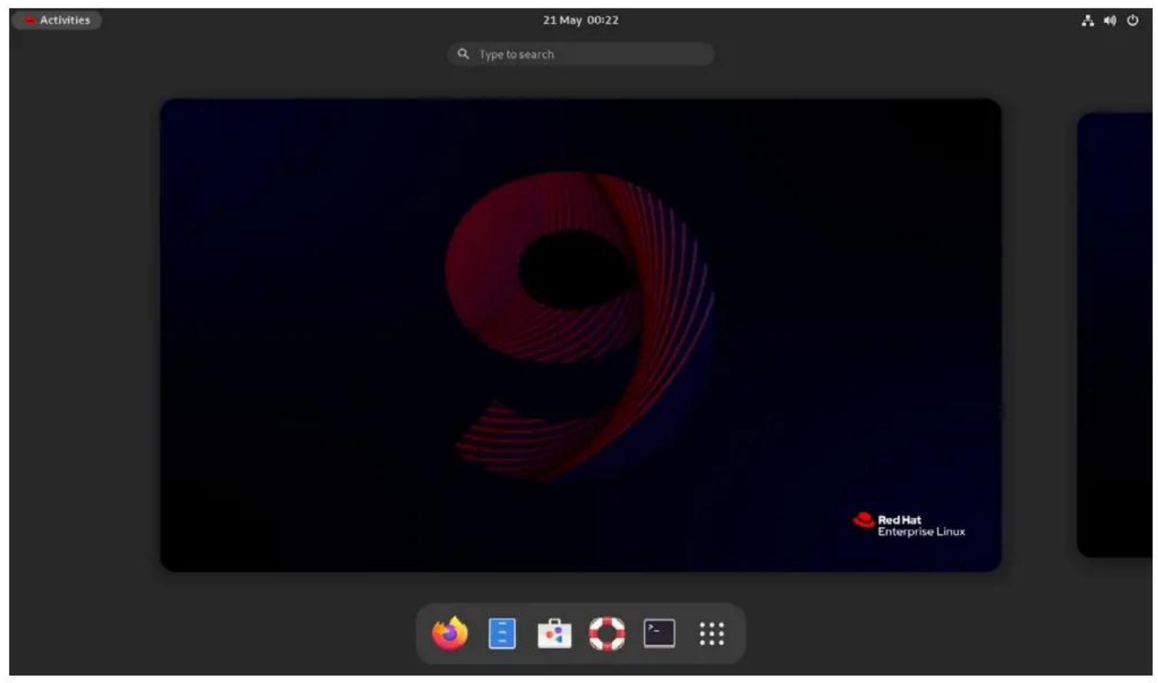

› Red Hat has served as the starting point for several other distros, such as Mandriva Linux.

When Red Hat Linux was discontinued, the company handed over development of the free distro to the community. The new project was called Fedora (*see p30*). The company steers the direction of the Fedora project and does so in order to use Fedora to incubate technologies that will eventually show up in RHEL.

Since the GPL prohibits it from restricting redistribution of RHEL, the company uses strict trademark rules to govern the redistribution. This has led to popular third-party derivatives that are built and redistributed after stripping away non-free components like Red Hat's trademarks. Distros such as CentOS, Scientific Linux and Oracle Linux offer 100% binary compatibility with RHEL.

> **"Some of the biggest and brightest Linux developers have worked with Red Hat."**

Red Hat has pioneered the professional open source business model, successfully mixing open source code and community development together with professional quality assurance, and a subscription-based support structure. The company also has employees working full-time on free and open source projects, such as Radeon, Nouveau and CentOS. »

CentOS

The CentOS distro has been the premier community-supported enterprise distro based on Red Hat Enterprise Linux (RHEL). The distro is built using the open source SRPMS from the RHEL distro. CentOS is one of the most popular server distros, suitable for all kinds of use cases, from web servers to enterprise desktops, and has been able to pitch itself as an ideal choice for anyone who wants to put together their own server but can't afford the RHEL subscription fees.

CentOS ships with RHEL's *Anaconda* installer and can run unattended installations across multiple machines thanks to *Kickstarter*. The installer provides various installation targets such as a web server, database server etc.

In January 2014, Red Hat announced that it would start to sponsor a bunch of core CentOS developers to work on the distro full-time. However, the developers and Red Hat have both insisted that the project remain independent of RHEL. The sponsorship ensures that all updates will be provided within 24 to 48 hours of upstream releases in RHEL.

Genus Fedora

You've been hit by a smooth distro...

» **Fedora has been around, in one form or another, since the early 1990s. The distro had its first release in 1995 and the early releases were named Red Hat Commercial Linux. During these early years, the distro was developed exclusively by Red Hat and the community was limited to contributing bug reports and contributing packages included in the distro. This changed in 2003 when the company shuttered Red Hat Linux in support of the Fedora Project and opened it up to contributions from the community.**

The aim of Fedora is to provide the latest packages while maintaining a completely free software system. The distro was initially called Fedora Core and was named after one of the two main software repositories – Core and Extras. The Fedora Core repo contained all the basic packages required by the distro as well as other packages distributed with the installation discs, and was maintained exclusively by Red Hat developers. The Fedora Extras repo was introduced with Fedora Core 3. It contained packages maintained by the community and was not distributed with the installation discs. This arrangement continued until version 7 in 2007 when the two repos were merged and the distro was renamed as Fedora.

Fedora's objective is to create a free software distribution with the help of the community. The development of the project is overseen and coordinated by the Fedora Project.

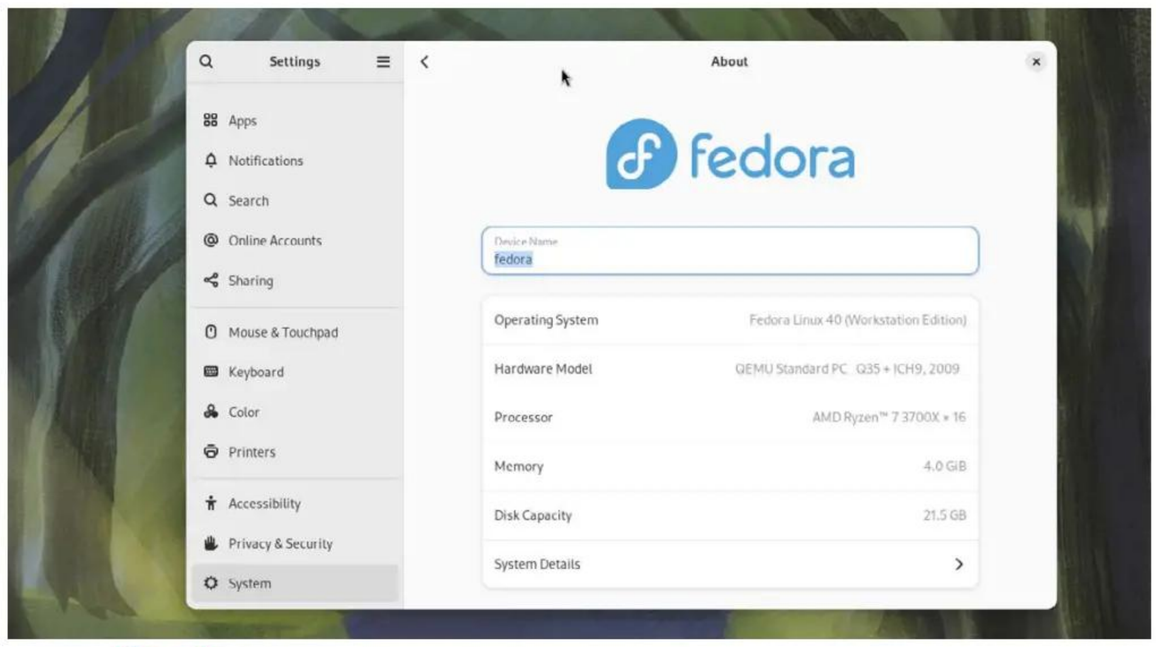

› **Fedora was one of the first distros to embrace the Security Enhanced Linux (SELinux) kernel module.**

It's made up of four Red Hat appointed members and five community elected members. The chairman of the board is appointed by Red Hat. Fedora strives to maintain a roughly six-month release cycle, with two releases each year. Every release is supported until the launch of the next two releases. The cycles are deliberately kept short so that developers can focus on innovation and introducing the latest technologies into the distro.

Feather in the cap

One way the community contributes is by hosting third-party repositories. In addition to its official software repos, there are several popular third-party software repos that usually contain software not included in the official repos – either because of the current laws of the country (such as multimedia codecs) or because the software doesn't meet Fedora's definition of free software. The Fedora project also produces the Extra Packages for Enterprise Linux (EPEL) repo, which contains packages for RHEL that are created by the community instead of Red Hat.

Apart from the main Fedora release, the project also ships various spins which are special-purpose distros aimed at specific interests, such as gaming, security, design, scientific computing etc. These are similar to Debian's Pure Blends. These and others are maintained by various Special Interest Groups (SIGs). The OLPC also runs a Fedora-based operating system. Fedora supports the x86 and ARM architectures and has also added support for PowerPC and IBM s390, starting with Fedora 20. Pidora is a Fedora Remix distro optimised for the Raspberry Pi.

Fedora's biggest contribution to the Linux ecosystem is its old command line package manager, *YUM* (Yellowdog Updater, Modified), which is based on *RPM* (Red Hat Package Manager). *YUM* enables automatic updates and dependency resolution, and works with the software repositories to manage the installation, upgrading and removal of packages. Since the release of Fedora 18 however, users have had the option to use the *dnf* tool which is a fork of *YUM*. The *dnf* tool has become the default package manager since Fedora 22, because it has better dependency resolution and is less memory intensive than other managers.

Korora

The Korora distribution started out as a way to ease the installation process of the Gentoo distro, but switched to using the Fedora distro as the base in 2010. The main aim of the distro is to make sure it works right out-of-the box for users.

Korora ships a live DVD, which includes a huge selection of apps that make it suitable for a large number of users and the distro offers five desktop choices – Gnome, KDE, Cinnamon, Xfce and Mate.

While Fedora only ships with open source software, Korora also includes some proprietary software, such as Adobe Flash, which are essential for catering to a wide user base. Korora also allow other software to be easily installed, such as *Google Chrome* and the proprietary graphics driver for Nvidia cards.

The distro has also eased a gripe for some Fedora users: graphical package management. Korora includes both *Apper* and *Yum Extender*, which are two of the most popular front-ends for YUM. Since it's based on Fedora, a new version of Korora is usually a few weeks behind a Fedora release.

Genus Mandrake

A distro which has a lot to scream about!

Until the release of Mandrake, Linux was generally thought of as a geek's OS. Mandrake was the first distribution that focused on the convenience of the user. The goal was to provide a distro that could be operated by regular computer users. It had features such as the ability to auto-mount CDs without messing around with configuration files, which brought greater convenience to the Linux desktop.

The Mandrake project has perhaps the most convoluted existence for a free software project. Over the years, the project has undergone various name changes, mergers and forks. However, it has spawned many distros and there are several major ones that are still active and can trace their lineage to Mandrake.

The distro has developed a bunch of custom tools – collectively known as drakxtools – to aid its users, who are called Drakes or Draks. One of the most distinguishing components created by the project is its *Mandrake Control Center* (MCC), which is now a centrepiece of all the derivatives. The MCC provides a single interface to access many different configuration tools. Using the *Control Center* in text mode is very useful in case of display problems or other serious issues, such as when the graphical server refuses to start. It's also interesting to note that all modules can be run as autonomous apps, without necessarily having to go through the *MCC*.

On life support

Mandrake Linux was first released in July 1998. It was based on Red Hat Linux 5.1 and featured the inaugural KDE desktop release. After the positive response, lead developer Gaël Duval, along with a bunch of others, created the company MandrakeSoft and in 2001 the company decided to go public. It faced its first major cash issue in late 2002 and asked its users to bail it out by subscribing to a paid service offering extra benefits, such as early access to releases and special editions. That wasn't enough and the company filed for bankruptcy protection in 2003. However, later that year MandrakeSoft announced its first quarterly profit and, in March 2004, a French court approved its plan to emerge from bankruptcy and return to normal operations.

› Rosa Linux is a Mageia fork that features some dramatic user interface mods such as this Simple Welcome app launcher.

The company also had to rename its product to Mandrakelinux, after losing a legal battle with the American Hearst Corporation over the Mandrake name and changed its name to Mandriva S.A. after acquiring the Brazilian company Conectiva in 2005. The distro became Mandriva Linux, but in 2006 Mandriva laid off several employees, including co-founder Duval. Amidst all the booing, the company continued putting out releases and created a niche for itself in the BRIC region (Brazil, Russia, India and China) as well as France and Italy.

Despite all their efforts, the company struggled to keep its balance sheet in the black. In 2010, Mandriva abandoned development of the community Linux distro to avoid bankruptcy. Immediately afterwards, former Mandriva employees announced Mageia, which has gone onto be one of the most popular Mandrake-derivatives. Mandriva S.A. transferred development to the community-driven Open Mandriva Association. The association's second release called OpenMandriva 2014.0 got a positive review from Duval. »

> **"One of the most distinguishing components created by the project is its Control Center"**

Salix Mageia

Mageia is one of the best-assembled community distros and does a wonderful job of carrying forward the Mandrake legacy. It has an expansive support infrastructure and very good documentation. The distro follows a nine-month release cycle and each is supported for 18 months. Mageia has installable live media as well as install-only DVD images.

Mageia boasts intuitive custom tools for managing various aspects of the distro. One of the best tools is the *Mageia Control Center*, which has modules for managing software, hardware peripherals and system services. Advanced users can employ it to share internet, set up a VPN and configure network shares. The distro uses the URPMI package management system and ships with three official repos. The Core repo contains open source packages, the Non-free hosts proprietary apps and drivers, and the Tainted repo includes patent-encumbered apps. The distro ships with several desktop environments, and the developers have made sure the user experience is consistent across all of them.

Genus SUSE

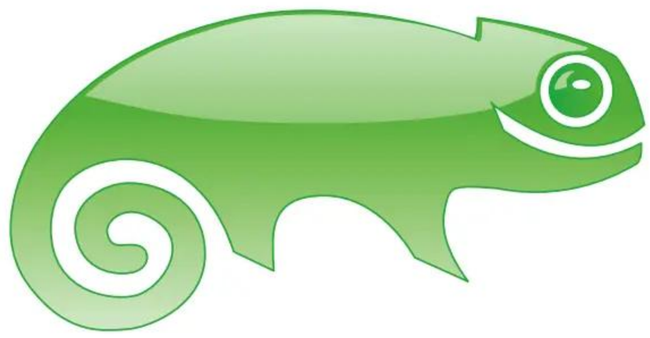

Nuga, nuga, nuga, gnu, nui*...

» **In 1992, Roland Dyroff, Burchard Steinbild, Hubert Mantel and Thomas Fehr founded Software und System Entwicklung (Software and Systems Development). The company started out as a service provider but the founders later decided to have a distro of their own that would cater to the enterprise user. The distro was named SUSE, after the acronym of their company, and it was a stock Slackware release that was translated into German and developed in close collaboration with Slackware's Patrick Volkerding.**

For building its very own distribution of Linux, SUSE used the now defunct Jurix distribution. Jurix was created by Florian La Roche, who subsequently joined the SUSE team and began to develop *YaST*, which is the distro's unique installer and configuration tool. The first SUSE distro that included *YaST* was released in May 1996 (*Yast* was rewritten in 1999, and was included for the first time in SUSE Linux 6.3 as an installer only).

Over time, SUSE Linux has incorporated many aspects of Red Hat Linux, such as its well-respected *RPM Package Manager*. In 1996, the first distribution under the name SUSE Linux was published as SUSE Linux 4.2. The confusing jump forward in version numbers was an intentional reference and homage to the answer to life, the universe and everything, as featured in Douglas Adams' *The Hitchhiker's Guide to the Galaxy*. *YaST*'s first version number, 0.42, was inspired by the same admiration for the author.

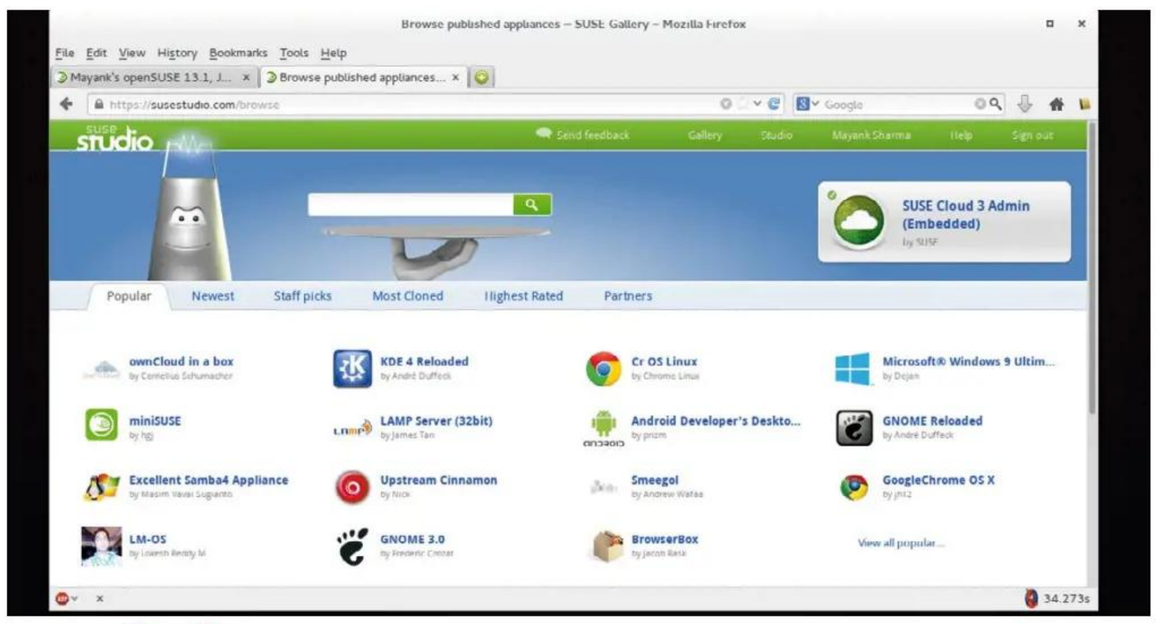

› **The SUSE Studio web service enables you to easily put together a customised OpenSUSE-based distro.**

SUSE's focus has always been on bringing open source to enterprise users. It introduced the SUSE Linux Enterprise Server in 2001, and changed the company name to SUSE Linux. After software and services company Novell acquired SUSE Linux in January 2004, the SUSE Linux Professional product was released as a 100% open source project and the OpenSUSE Project was launched – much like Red Hat did with Fedora. The software was always open source and now so was the process which enabled developers and users to test and evolve it.

Enterprising

The initial stable release from the OpenSUSE Project was SUSE Linux 10.0. It included both open source and proprietary applications, as well as retail boxed-set editions. This was also the first release which treated the Gnome desktop environment on a par with SUSE's default KDE desktop. As of version 10.2, the SUSE Linux distribution was officially rechristened as OpenSUSE.

In November 2006, Novell signed an agreement with Microsoft covering improvement of SUSE's inter-operability with Windows, cross-promotion and marketing of both products, and patent cross-licensing. This agreement is considered controversial by some of the FOSS community.

Novell was later acquired by The Attachmate Group in 2011, which then divided Novell and SUSE into two separate subsidiary companies. SUSE offers products and services around SUSE Linux Enterprise – a commercial offering that is based on OpenSUSE Linux.

SUSE develops multiple products for its enterprise business line. These products target corporate environments and have a longer lifecycle (seven years, extendable to 10), a longer development cycle (two to three years), technical support and certification by independent hardware and software vendors. SUSE Linux Enterprise products are only available for sale. There's also the SUSE Linux Enterprise Desktop (SLED) which is a desktop-oriented operating system designed for corporate environments. In contrast, OpenSUSE does not have separate distributions for servers, desktops and tablets, instead using various installation patterns for different types of installation.

*http://bit.ly/ChameleonSong

OpenSUSE

OpenSUSE is one of the best RPM-based distros. It comes in several editions for 32-bit and 64-bit architectures and also has ports for ARM v6, ARM v7, and the 64-bit ARM v8. Once known for its KDE desktop, OpenSUSE now looks good across all the major desktops. Besides KDE and Gnome, the distro also features Mate, Xfce, Enlightenment, and LXDE. You can download the distro either as a smaller live installable image or a install-only DVD image.

One of OpenSUSE's hallmarks is the distro's *YaST*, which is a setup and configuration utility that enables you to tweak many different aspects of the system. Another popular tool is *Snapper*, which enables you to revert to a previously created system snapshot.

The distro serves as the base for the SUSE Linux Enterprise products – much as Fedora does for RHEL – and is suitable for all types of users regardless of their skill set. The distro's installer is versatile and offers several customisation options. It can be navigated by new users and includes options to plug the installed system into a corporate directory server.

Genus Slackware

The tortoise distro that's outlasted many hares.

Slackware has the honour of being the oldest distro that's still actively maintained. It was created by Patrick Volkerding and had its first beta release in 1993. The project aims to create the most Unix-like Linux distribution. Slackware was originally derived from Softlanding Linux System (SLS), which was the first distro to provide TCP/IP and X Windows System in addition to the Linux kernel and basic utilities. SLS, however, was very buggy and the growing frustration of SLS users prompted Volkerding to release an SLS-like distro in July 1993.

Back then, in addition to being hosted on an anonymous FTP server at the Minnesota State University Moorhead, the distro was offered as 24 3.5-inch floppy disks. By the time version 2.1 was released in October 1994, the distro had swelled to 73 disks and Version 3 was released on CD-ROM.

The USP of the distro is that it makes very few changes to upstream packages. Unlike other distros that aim for a particular userbase or a wide variety of users, Slackware doesn't preclude user decisions and doesn't anticipate use cases. The user, therefore, has far greater control on the installed system with Slackware than most other distros.

Cut some slack

Unlike other distros, Slackware doesn't provide a graphical installation. It continues to use plain text files and only a small set of shell scripts for configuration and administration. The distro also doesn't provide an advanced graphical package management tool, relying instead on command line tools such as *pkgtool*, *installpkg*, *upgradepkg*, and *removepkg*. However, these native tools can't resolve dependency issues.

Slackware packages are just plain compressed TAR archives. The package contains the files that form part of the software being installed, as well as additional metadata files for the benefit of the Slackware package manager. As of Slackware 12.2, *slackpkg* has become the official tool for installing or upgrading packages automatically through a network or over the internet, complementing the traditional package tools suite that only operates locally. *Slackpkg* also doesn't resolve dependencies between packages. Traditionally, Slackware only offered a 32-bit release, and users had to rely on unofficial ports, such as slamd64 for 64-bit releases. Since Slackware 13, a 64-bit variant is also available and officially supported. In 2002, Stuart Winter started the ARMedslack project, a port of Slackware for ARM. In 2009, Volkerding knighted ARMedslack as an official port of Slackware. With the release of Slackware 14.0, the project has been completely renamed to Slackware ARM.

› In addition to Slackware-stable, the project also provides a testing-current branch for more bleeding-edge software.

It might sound surprising, but Slackware is a popular base for many distros. The derivatives include expansive desktop projects, live distros, security distros etc.

The Slackware project is also missing some of the common developer-friendly tools. For example, there's no official bug tracking system. Also, there is no official mechanism to become a contributor for Slackware. The final decision on what goes into the distribution is made by Volkerding – Slackware's 'Benevolent Dictator For Life'.

> **"Unlike other distros, Slackware doesn't preclude user decisions and doesn't anticipate use cases"**

In another departure from the norm, Slackware doesn't follow a fixed release schedule. The objective is to release a very stable system and so the project follows a release-when-ready philosophy, but still aims for one major release a year. »

Salix OS

Salix OS is one of the best Slackware-based distros: it's light, nimble and backwards compatible with Slackware. One of its salient features is that it minimises bloat by having only one application per task. The distro supports both 32-bit and 64-bit architectures and is available in five variants for the KDE, Mate, Xfce, Openbox, and Ratpoison desktops.

Salix offers three modes of installation – Full, Basic and Core. The Full option installs everything on the installation image; Basic provides a barebones system with just the graphical desktop and a few essential apps and the *Slapt* package manager; the Core option will only install a console-based system and is designed for users to custom-build their install.

The Full distro includes all the apps you'd expect on a desktop distro, and is often touted as Slackware with a graphical package manager. Its package manager, *Gslapt*, resembles the *Synaptic* package manager and also provides all the same functionality. Multimedia codecs aren't supplied out of the box, but that can be fixed with the distro's custom *Codecs Installer*.

The marsupials

Evolutionary masterpieces and mavericks.

» Gentoo Linux

The goal of the Gentoo project was to create a distro without pre-compiled binaries that was tuned to the hardware on which it was installed. Unlike a binary software distribution, the source code is compiled locally according to the user's preferences and is often optimised. It was initially called Enoch but Gentoo 1.0 was released in 2002.

Gentoo has the distinction of being one of the most configurable distros and appeals to Linux users who want full control of the software that's installed and running on their computer. Gentoo users get to create their system from the ground up: the distro encourages the user to build a Linux kernel tailored to their particular hardware. It allows very fine control of which services are installed and running. Memory usage can be reduced, compared to other distributions, by omitting unnecessary kernel features and services.

This distro is a rolling release and one of its notable features is its package management system called *Portage*. If you've never used it before, there's a steep learning curve to using Gentoo. Derivatives such as Funtoo can be a good starting point if you're not ready to dive straight in.

Arch Linux

Judd Vinet wanted to create a distro that was inspired by the simplicity of Crux, Slackware and BSD and thus created Arch Linux in 2002. Arch aims to provide a lightweight foundation on which the user can build according to their needs. In Vinet's words: "Arch is what you make it". A bit like life really.

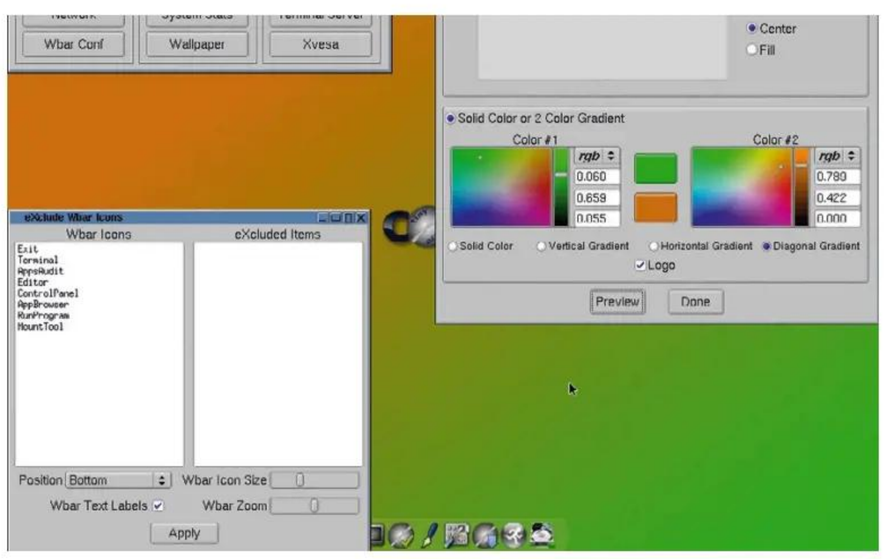

› **Tiny Core Linux is an ickle distro at 12MB. Ah, sweet!**

› **Arch Linux is ludicrously customisable, offering all the latest packages.**

The most impressive feature of the Arch distro is the *Pacman* package management tool. Arch is a rolling release that can be brought up to date with a single command. Installing Arch Linux is an involved process and although it is well-documented, it's still better suited for experienced Linux campaigners. However, Manjaro Linux is an Arch derivative and is more user-friendly and has a graphical installer.

Tiny Core Linux

If you can't invest time in creating an Arch or Gentoo installation, check out Tiny Core Linux. The distro installs the bare minimum software you need to boot into a very minimal X desktop. From this point on, you've complete control and can install apps from online repos or compile them manually.

The distro is a mere 12MB and bundles only a terminal, a text editor and an app launcher on top of the lightweight *FLWM* window manager. It has a control panel to manage bootup services and configure the launcher, but everything else needs to be pulled in from its manager, including the installer if you want Tiny Core on your hard disk. The distro also has a CorePlus variant, which has additional drivers for wireless cards, a remastering tool and internationalisation support.

Puppy Linux

One of our all-time favourites, Puppy Linux had its initial release in 2003 and the first stable one in 2005. The distro is built from the ground up and its initial goal was to support

Linux From Scratch

Rather than being a distribution itself, Linux From Scratch – popularly called LFS – is a freely available set of instructions to create your own custom distro from the ground up, entirely from source. The project was started in 1999 when its author, Gerard Beekmans, wanted to learn how a Linux distro works behind the scenes. While building his system from scratch, Beekmans wrote down the steps and released it as a HOWTO (*pictured, right*) thinking that there would probably be other people who would be interested.

LFS has grown quite a bit from its humble start, transforming from a single HOWTO to a multi-volume book. It has also spawned various sub-projects over time, such as BLFS or Beyond LFS which fleshes out the basic LFS system, and ALFS or Automated LFS, which is designed to help automate the process of creating an LFS system.

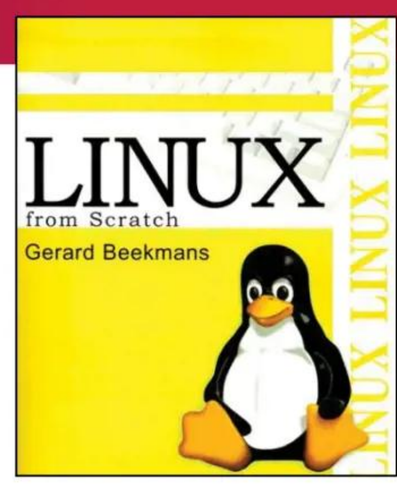

› Puppy Linux has become a handy distro for recovering data from PCs and removing malware from Windows.

older hardware that had been rendered useless due to lack of support in other distros.

The real power of the distro lies in its plethora of custom apps. There are custom apps to block website ads and add internet telephony, a podcast grabber, a secure downloader, an audio player, a DVD burning app and more. First-time users might be intimidated by Puppy's installer as it has no automatic partitioner, and fires up *Gparted* for you to format the disk. But each step in the installer is well-documented within the installer itself.

Packages for Puppy Linux are called pets, and have a .pet extension. You can install packages using its custom *Puppy Package Manager* tool, and you can configure it to download packages from other Puppy repos. The distro includes tools which can be used to easily churn out variants. Puppy Linux variants are called puplets. Popular puplets are WaryPuppy for supporting older hardware, RacyPuppy for newer hardware, the Slackware-based SlackoPuppy, and PrecisePuppy which is based on the Ubuntu LTS release.

SliTaz GNU/Linux

SliTaz stands for Simple Light Incredible Temporary Autonomous Zone and had its first stable release in 2008. The distro is built with home-brewed tools known as cookutils and uses *BusyBox* for many of its core functions. The distro includes a mixture of the LXDE and *OpenBox* window manager and is designed to perform on hardware with only 192MB of RAM. The distro weighs under 30MB and uses a mere 80MB of hard disk space.

The distro also has a bunch of custom tools such as the *Tazpkg* package manager and *SliTazPanel* for administering all aspects of the distro. SliTaz repos include over 3,000 packages for every popular open source app and it's a common option for powering low-powered machines.

PCLinuxOS

PCLinuxOS began life as a repository of RPM packages for the Mandrake distro in 2000 and became a distro in late 2003 as a branch of Mandrake Linux 9.2.

Although it retains a similar look and feel to Mandriva Linux, and its configuration tool and installer give away its Mandriva lineage, PCLinuxOS has diverged significantly.

The distro has replaced Mandrake's *URPMI* package management system, opting instead for APT-RPM. This is based on Debian's *APT* but uses RPM packages, together with the *Synaptic* package manager. PCLinuxOS is a KDE distro, but also has community spins around the LXDE and Mate desktops.

We're going to tread on that hallowed patch of earth where angels fear to tread (and not because a bushy-bearded Russell Crowe is eyeballing them) and attempt to pick an overall distro winner. This means we had to pick a criteria that allowed for ease of use alongside the ability to build in complexity for specific use cases and, we admit, the result is purely subjective. Don't agree with us? Why not email your top picks for each genus to Linux Format magazine at **lxf.letters@futurenet.com**.

Mageia

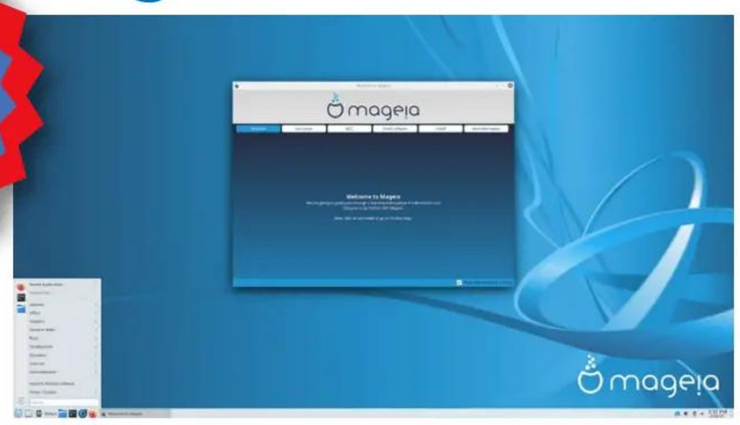

The community supported distro has everything you want from a modern Linux distribution – an active and vibrant user and developer community, a well-defined support structure, support for multiple desktops and install mechanisms.

OpenSUSE

Coming in at second place, the OpenSUSE distribution loses out because of recent activities of its corporate parent. Also, the distro still focuses on introducing changes that make it fit more snugly on the corporate desktop, rather than home user.

Korora

This is your best bet if you want a RPM-based distribution that works out of the box. However, Korora is still essentially a one-man show and inherits some of the less flattering features of its parent distro, Fedora.

TOP 100 LINUX TOOLS

Take a stroll through the open source garden as we pick the best apps, tools and utilities available to all Linux kind...

We all have our favourite open source apps that work for us better than any available alternative. But take a moment and step back from the *Emacs* vs *vim* type battles raging on in the Linux-verse and marvel at the sheer number of apps at our disposal. Your distros' software repositories give you access to thousands of apps, and you can install everything from fully featured app suites to nifty command-line utilities literally with the touch of a button.

There are open source apps and tools for all kinds of applications today. There's hardly any use case that isn't catered for by a community contributed app. Many of these apps have proved their mettle and offer features and performance benefits that surpass their proprietary counterparts. They have also proved themselves to be invaluable to home and business users in more than one sense of the word. According to rough estimates on **www.openhub.net**, some popular apps such as *LibreOffice*, *Firefox* and *Apache* would take several hundred person-years to develop and cost millions of pounds. Yet they are all available to you for no cost.

> **"Many of these apps have proved their mettle and surpass their proprietary counterparts."**

Open source apps come in many shapes and sizes and you can grade them based on their usability. There are feature-rich apps, task-oriented app suites, well put-together tools, and newfangled novelty apps and games. Some ship with well-designed graphical interfaces and others show their more versatile sides when operated from the command-line.

In this feature, we traverse this diverse and vast collection of open source gems on offer and pick the ones that are at the top of their game. In this list of the 100 best apps we've covered a wide range of categories. Whether you are a business owner, an educational institution, a developer, a home user, or a gamer, we've got something for everyone. While you'll be familiar with some of the most popular tools in this list, rest assured there are quite a few that might have missed your attention. If you've been unable to escape the clutches of commercial software, we're sure you'll find quite a few tools on this list that are suitable replacements.

Essential apps

A Linux desktop isn't complete without them!

LibreOffice

Forked from *OpenOffice.org*, *LibreOffice* has become one of the most popular office productivity suites. It includes programs for word processing, and can create spreadsheets, slide shows, diagrams and drawings, maintain databases, and compose mathematical formulae. It also offers good compatibility with documents in proprietary formats and has recently had a face lift.
www.libreoffice.org

Wine

Despite the increasing number of cross-platform apps that work on Linux, there are some that still only support Windows. This includes big third-party proprietary apps, such as *Adobe Photoshop* or just small niche home-grown tools that you can't do without. For such situations, you can use *Wine*, which generally run these Windows-only apps and games with ease. The project supports over 20,000 apps. Some work flawlessly out-of-the-box while others require minor configuration tweaks.
www.winehq.org

Remmina

With *Remmina* you can access a remote computer from the comforts of your desktop. It supports the widest range of protocols and will connect to all kinds of remote desktop servers. The app is easy to use, and has enough features that make it a viable option for occasional use.

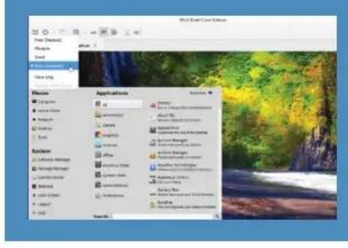

http://remmina.sourceforge.net

Thunderbird

Another gem from the Mozilla Foundation, *Thunderbird* is one of the best email clients, being easy to setup and is brimming with features. Simple setup wizards aid syncing with popular web-based email services and it can manage multiple accounts, supports encryption and is extended through add-ons.
www.mozilla.org/thunderbird

KeepassX

Trying to remember different passwords for the various services is a challenge for most humans (that don't count cards in Las Vegas for fun). You can defer this task to *KeePassX* which stores password in an encrypted database. It can fill in the password automatically and also includes a random password generator.
www.keepassx.org

BleachBit

A distro accumulates a lot of digital gunk over time. *BleachBit* helps you spring clean it and protect your privacy. It also removes temporary and other unnecessary files, and has tools to securely delete files or wipe them.
http://bleachbit.sourceforge.net

OpenSSH

When you need to interface with a remote computer, you cannot do without OpenSSH. It's a family of tools that provides secure tunnelling capabilities by encrypting all traffic and includes several authentication methods, and supports all SSH protocols.
www.openssh.org

Gufw

You may not be using a firewall currently, and if that's because they are difficult to set up then you need *Gufw*. It features an intuitive graphical interface for managing the inbound and outbound traffic rules for various apps and services and even individual ports. Its wizard-like graphical menus are designed especially for inexperienced users.
www.gufw.org

› Gufw has profiles and preconfigured rules to aid inexperienced users.

VirtualBox

When *Wine* doesn't cut it you can use *VirtualBox* to run an entire Windows installation inside a virtual machine. The software is also useful for installing experimental apps that you don't want to deploy on a real computer, and for testing other OSes without exposing it to real hardware.
www.virtualbox.org

Clonezilla

This is a cloning solution that's distributed as a live CD and is popular for doing bare metal backup and restoration of individual PCs. It can also deploy an image to multiple computers in a lab. Clonezilla can work with a large number of popular disks, partitions and filesystem types.
www.clonezilla.org

VLC

Distros ship with a functional video player. But if you need more control, there's no beating *VLC*. It supports virtually every video and audio format out there and includes handy CLI tools for advanced users.
www.videolan.org/vlc

PeaZip

PeaZip is a graphical archiving tool that can work with over 130 different types of archive files and can even create encrypted ones. It integrates with popular desktops and also has a CLI for advanced users.
http://bit.ly/PeaZipSF

Gparted

Use *Gparted* to restructure a disk on your computer. It's available as a live CD and can also be installed inside your distro. *Gparted* can create, resize, move, delete, reformat or check partitions and supports many filesystems.
www.gparted.org

ZuluCrypt

Create an encrypted disk within a file or within a non-system partition or USB disk. *ZuluCrypt* has an intuitive user interface and can be used to encrypt individual files with GPG.
http://bit.ly/zuluCrypt

HomeBank

This is a feature-rich finance app. It can import data from other apps and bank statements in popular formats. It can also detect duplicate transactions and features dynamic reports and is easy to use for budgeting.
http://homebank.free.fr

Internet apps

Get the best of the web with these tools.

Firefox

One of the most widely recognised pieces of open source software, Mozilla's *Firefox* web browser is the default browser on virtually every Linux distro. It's pretty responsive and known for its privacy features. You can customise it to the hilt and also extend it with an impressive number of extensions.
www.firefox.com

gFTP

The *gFTP* client is a feature-rich client that'll get the job done, if you need to download files via FTP occasionally. It has a simple two-pane interface that shows the content of the local and remote filesystem. Using *gFTP* you can also transfer files between two remote servers.
http://gftp.seul.org

Tox

Privacy conscious users should try the new decentralised IM and VoIP client called Tox. This relies on a distributed network, which uses P2P connections, the same technology used by BitTorrent to provide a direct connection, between users for chats and, unlike other Skype alternatives, Tox uses no centralised servers or supernodes, which could be compromised. All chats are also encrypted using the peer-audited NaCl crypto library.
https://tox.im

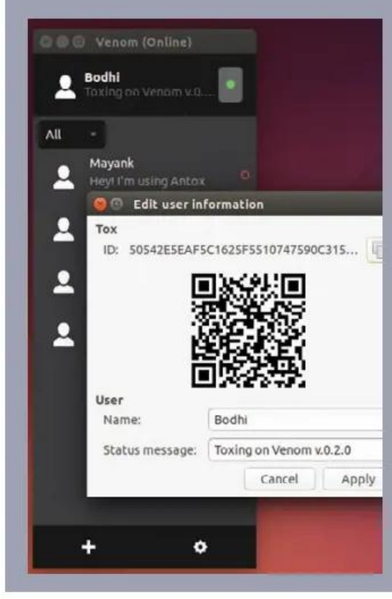

RSSOwl

An excellent desktop alternative to Google Reader, RSSOwl is a news aggregator for RSS and Atom News feeds that's easy to configure. The app gathers, organises, updates, and stores news in an easy to use, and saves selected items for offline viewing and sharing.
www.rssowl.org

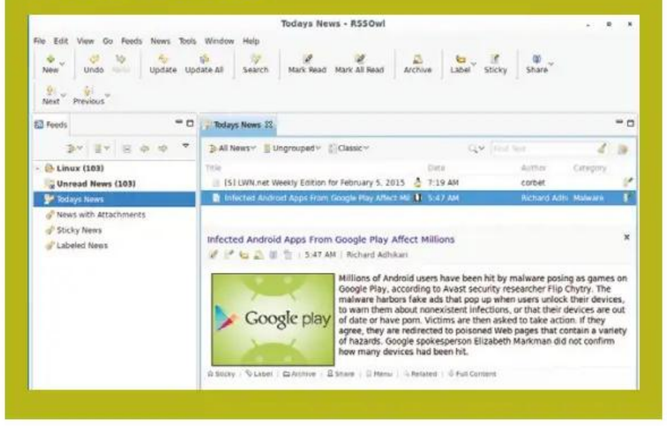

Jitsi

Jitsi is the best VoIP app, as long as you're not averse to Java apps. It supports IM and make one-to-one audio and video calls, as well as audio conference calls. It supports many of the widely used IM and telephony protocols, including SIP, XMPP, AIM, ICQ, MSN, etc. Jitsi has all the features you'd expect from a softphone, and more, such as encrypt text chats with OTR and voice and video by establishing a ZRTP session.
https://jitsi.org

Aria2

What makes *Aria2* a unique utility is that it can download the same file at the same time using different protocols. The lightweight CLI app can download via HTTP, FTP, BitTorrent and Metalink and can also open multiple connections to download the file faster.
http://aria2.sourceforge.net

Midori

The go-to browser for anyone concerned about resource consumption, *Midori* is popular with lightweight distros. Despite its lightweight nature and design, Midori has all the features you'd expect from a web browser including a speed dial, tabbed interface, bookmark management and configurable web search as well as an incognito mode.
www.midori-browser.org

FileZilla

For those who use FTP a lot, there's FileZilla. The client supports FTP, SFTP and FTPS protocols and has just about any configuration option you can imagine. It also has a tabbed interface so you can browse more than one server and even transfer files simultaneously between multiple servers.
https://filezilla-project.org

Deluge

BitTorrent is popular for downloading Linux distros and there are numerous download clients. One of the best is *Deluge* which has multiple front-ends, including a graphical and a web-interface. It has features that enable advanced users to tweak it to their liking and also has a nice library of plugins.
www.deluge-torrent.org

Pidgin

Pidgin is a wonderful app for instant messaging over many network protocols. You can sign in with multiple accounts in the single client and chat with many friends in different networks. You can use it to connect to AIM, MSN, Google Talk, Yahoo, Jabber, IRC and more chat networks all at once.
www.pidgin.im

Media

0 A.D.

This is a real-time civilisation-building strategy game that features impressive graphics and intense battle gameplay. It's yet to have a final release but has already won accolades in its current state.
http://play0ad.com

FreeCiv

A strategy game that challenges players to lead their tribe from 4,000 BCE to the space age.
www.freeciv.org

Alien Arena

A popular first person shooter with a sci-fi theme and the tournament style deathmatch of *Quake* and *Unreal Tournament*. The game has several game modes and over 60 maps, and is quite configurable.
http://red.planetarena.org

OpenMW

OpenMW is a new game engine that recreates the popular *Morrowind* RPG. The aim of the project isn't to improve game assets or add additional features but to provide gamers a more moddable edition of the game.
https://openmw.org

FlightGear

For fans of aircraft simulators there's *FlightGear* that aims to offer flight across real world terrain. It includes scenery for more than 20,000 airports, and can be extended with your own aircraft and locations.
www.flightgear.org

Office and productivity

Enhance your workflow with these apps.

Calligra

Unless you feel you need *LibreOffice*'s superior compatibility with proprietary formats, you may want to consider *Calligra*. It's a continuation of *KOffice* and unlike *LibreOffice*, *Calligra* has a modern-looking, modular design, and also uses Open Document as its native file format. It ships with a large clutch of apps. In addition to the *Words* word processor, *Tables* for spreadsheets, *Stage* for preparing presentations, and *Kexi* for managing databases, it also benefits from Krita for digital painting.
www.calligra.org

GnuCash

Gnome users have *GnuCash* which is similar to *KMyMoney* in terms of features, but also handles and categorises entries differently. *GnuCash* is a personal and small business accounting app that's based on double-entry for professional reporting and besides dealing with monetary transactions, it can track things such as stocks, bonds and mutual funds.
www.gnucash.org

Xournal

This app is very handy for when you need to scribble bits of information down for later. As well as typing out notes, you can use it with either a mouse or a stylus. It can also be used to add annotations to PDF files.
http://xournal.sourceforge.net

Zathura

This is a simple and a lightweight PDF reader that supports almost all the usual features you'd expect. You can search text strings, jump pages, zoom in and out, rotate pages, add bookmarks and more. In addition to PDFs, it can display DjVu and even encrypted documents.
https://pwmt.org/projects/zathura

AbiWord

The wide gap between rich text editors and word processors is occupied by *AbiWord*. It's lightweight but still offers commonly-used word processing features, which makes it a popular for lightweight distros. It also offers cloud-based collaboration capabilities via its **AbiCollab.net** service. **www.abiword.org**

ProjectLibre

A project management tool helps you stay on top of ongoing projects and *ProjectLibre* is one of the best. It's an award winning app that's used widely by many enterprises around the world. *ProjectLibre* has several useful features and can also visualise tasks with various charts and reports.
www.projectlibre.org

OpenLDAP

OpenLDAP is great for when you want to run a directory server. It implements the LDAP protocol and has all the expected features, including logging, replication, access control, user and group management etc. It also integrates with Active Directory.
www.openldap.org

Gnumeric

A *biWord* is usually paired with the lightweight *Gnumeric* spreadsheet app. However, the app isn't light on features and offers a lot more functionality than proprietary spreadsheet apps. *Gnumeric* will import data from Microsoft Excel files and there are import filters for other apps as well.
www.gnumeric.org

KMyMoney

Designed for KDE users, *KMyMoney* is a feature-rich accounting app. It supports different account types, such as Cash, Checking, Savings, etc and can categorise incomes and expenses, and can reconcile bank accounts. If your bank allows it, you can have *KMyMoney* connect to your bank directly to retrieve your account activity.
https://kmymoney.org

Calibre

You can use *Calibre* to manage your collection of ebooks, and supports a wide range of readers and smartphones. The app can import ebooks manually or, if you prefer, by syncing a reading device such as the Kindle. Any files imported can be sorted and grouped by metadata fields, which can be pulled from various online sources, such as **www.goodreads.com**.
www.calibre-ebook.com

Achievo

This is a web-based resource management tool with a simple interface for accessing its CRM, HRM and project management and planning tools. You can also track resources across multiple projects.
www.achievo.org

Okular

The default PDF viewer for KDE and includes a good number of useful features. Besides PDF it can also read a number of other file types, including Postscript, DjVu, CHM, XPS, ePub, TIFF, CBR, and others.
https://okular.kde.org

LaTex

LaTex is a document preparation system and document markup language based on TeX. Its purpose is to simplify TeX typesetting for documents containing mathematical formulae and is widely used in academia.
www.latex-project.org

Shutter

Besides capturing the full screen, Shutter can capture a specific area, or a window. You can also upload to a hosting service.
www.shutter-project.org

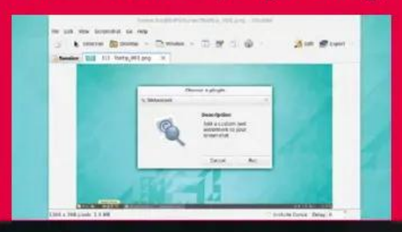

ClamAV

While most viruses and trojans will have no effect on Linux, you still can have infected files in your distro that can wreck havoc when accessed on a Windows machine. So be a good admin and use ClamAV to scan files.
www.clamav.net

Hobbyist

Follow your passion!

RawTherapee

Do you shoot with a digital single lens reflex camera (DSLR)? Then take a look at *RawTherapee* which includes a wide range of tools for processing and converting RAW files. In addition to basic manipulations, the app has extensive options for working with RAW files. Using the app you can adjust the colour and brightness values of your images, correct white balance, adjust tones, and a lot more. Besides RAW files you can also use *RawTherapee* for editing traditional image files, and it also includes Adobe Lens Correction profiles.

www.rawtherapee.com

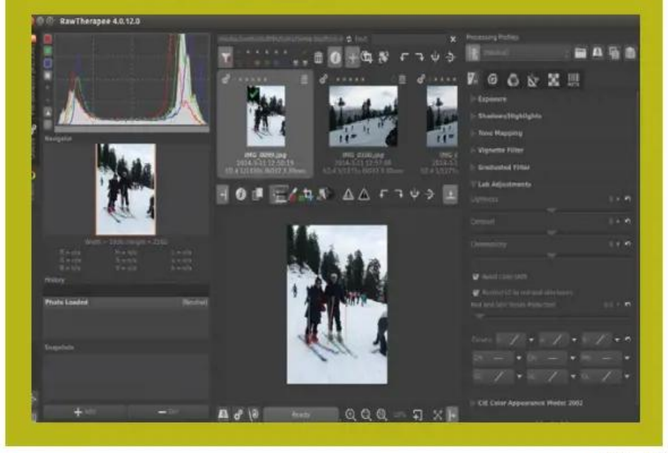

Scribus

A comprehensive desktop publishing program. *Scribus* can be used to create professional press-ready online and print documents including brochures, booklets, books and magazines. It has a feature-rich interface and has features, such as PostScript colour separations, support for CMYK and spot colours, ICC profiles, and printer marks. *Scribus* also includes a variety of templates and styles and you also get an array of settings and tools to precisely define and position the various layout elements you require.

www.scribus.net

Krita

Although *Krita* is part of the *Calligra* suite it needs a special mention of its own. *Krita* is a digital painting and illustration app that offers many expressive brushes, HDR painting, filters, perspective grids, painting assistants, and many other features you'd expect from such an app.

www.krita.org

Stellarium

Stellarium is a free open source planetarium for your computer. It calculates the positions of the Sun and Moon, planets and stars, and draws the sky as per the users location and time. It can also draw the constellations and simulate astronomical phenomena such as meteor showers, and eclipses.

www.stellarium.org

JOSM

Keen to contribute to the mapping project, OpenStreetMap? Then use *JOSM*. It's a Java-based offline map editor that can help you plot GPS traces. You can load GPS track-logs into JOSM and start adding streets to OpenStreetMap instantly. Although OpenStreetMap has several other editors available, most contributors use JOSM for their edits, as it lets them upload changes back to OSM quickly and easily enough. JOSM offers several features and can be extended with plugins and styles.

https://josm.openstreetmap.de

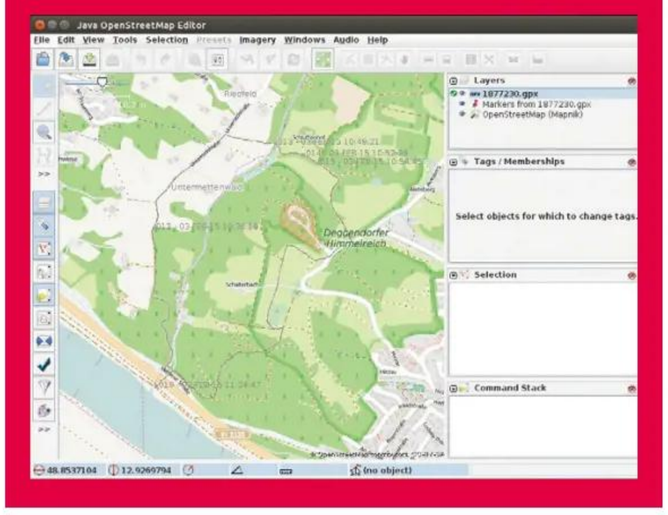

OpenShot

There are only a handful of video editors for Linux and *OpenShot* offers the best combination of features and ease of use for the home user. You can use it to combine videos, audio tracks, and still images together and add in captions, transitions, and more, and export the final product in a variety of formats. *Openshot* can also use *Blender* to create 3D animated titles.

www.openshot.org

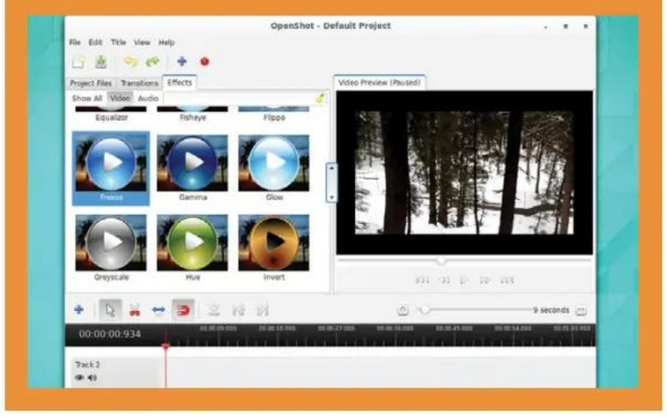

Inkscape

Another pro-quality tool, *Inkscape* offers advanced vector graphics editing and is popular for drawing vector art, line art, and designing logos and graphics. It's brimming with features, such as markers, clones, alpha blending and more, and is often compared to expensive proprietary apps such as *Illustrator* and *CorelDraw*.

www.inkscape.org

Cinelerra

Cinelerra is excellent if you need to edit more than home videos, as it's the most advanced non-linear video editor and compositor for Linux. It supports HiFi video and audio and is resolution and frame-rate independent, which enables it to edit videos of any size. The app has several advanced features, such as overlays, denoising, normalisation, time stretching, color balance, compositing, real time effects and a lot more. It also includes a compositing engine for performing tasks such as keying.

www.cinelerra.org

Media

Comix

Digital comics are distributed as comic book archive files that mainly consist of a series of image files, typically PNG or JPEG files, stored as a single archive file. *Comix* can read digital comics in virtually every format.

http://bit.ly/ComixApp

FontForge

FontForge is a feature-rich app for creating and editing fonts and supports all common font formats. It can extract information from a font file as well as convert from one format to another, and can be used for previews.

http://bit.ly/FontForge

CairoDock

CairoDock is a MacOS X dock-like app. One of its main advantages over other docks is that it doesn't require a compositing window manager to work and can add bling to older low-powered machines.

www.glx-dock.org

Audacity

If you need to work with audio, you should use the powerful *Audacity* sound editor. You can trim audio, combine tracks, and even stack multiple tracks, as well as export to a number of formats and quality settings.

http://bit.ly/AudacityApp

MPD

The Music Player Daemon is an audio player with a server-client architecture, which means you can control it remotely from another computer. It plays audio files, organises playlists and can maintain a music database.

www.musicpd.org

Development

Power tools and programs for power users.

jEdit

This is a text editor for programmers that supports auto indent, and syntax highlighting for more than 140 different programming languages. The app enables you to define complex macros and offers a powerful and user-friendly keyboard mapping system. It's highly configurable and customisable, and you can extend its functionality by adding plugins.
www.jedit.org

Eclipse

There's no beating *Eclipse*, the most feature-rich IDE. Although Java is its speciality, Eclipse supports a range of languages via plugins. In fact, its plugin marketplace is an indispensable resource. Eclipse does code refactoring and you can use it to extract the selection as a local variable or method. Since it can target multi-person installs, it handles version control very maturely.
www.eclipse.org

BlueFish

Do you develop for the web? *Bluefish* is a multi-language editor that's designed for web developers. It supports many programming and markup languages and focuses on dynamic and interactive websites. It supports code block folding, unlimited undo/redo, automatic tag closing, and syntax highlighting. Another useful feature is the snippets bar from where you can add the most common snippets of code for a variety of languages. *Bluefish* also has support for popular open source web apps such as MediaWiki and Wordpress.
http://bluefish.openoffice.nl

Meld

A graphical diff tool, *Meld* enables you to compare two or three files as well as whole directories. It includes features, such as syntax highlighting and direct file editing, and using the tool you can easily isolate and merge the differences. *Meld* can also be used to browse various popular version control systems such as *CVS* and *Subversion*.
www.meldmerge.org

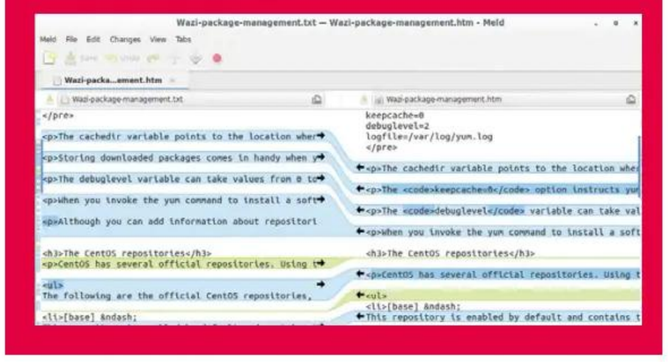

KompoZer

New and experienced HTML programmers will save a lot of amount of time and effort with the *KompoZer* editor. It has an intuitive interface and includes a colour picker, an FTP site manager, CSS editor, customisable toolbars, forms, spell checker, markup cleaner and can also validate code using W3C's HTML validator.
www.kompozer.net

Gimp

Despite its name, Gimp is a powerful, comprehensive image manipulation program. It offers a wide range of tools for professional-quality photo retouching and image manipulation capabilities for free. It also offers a huge list of features and supports all the common graphics file formats.
www.gimp.org

Blender

With *Blender* animators can create 3D printed models, visual effects, art, interactive 3D applications and video games. The app provides a wide range of features that can be used to create 3D animation films. It's a one-stop 3D package and includes a gaming engine, a video sequence editor, production-ready camera and object tracking, a large library of extensions, and an advanced physics engine. It can render fluid dynamics and simulate the movement of elastic objects and clothes.
www.blender.org

Geany

You don't need a full-blown IDE if you only program occasionally, which makes *Geany* a good choice. It's a cross between a plain text editor and an IDE with support for the popular languages and nifty features like a compile/run button, a listing of functions defined in the currently opened file, and much more.
www.geany.org

APTonCD

Suddenly realised that you need to move your Ubuntu installation or need to give a friend a copy of your setup? With *APTonCD* Ubuntu users can back up all of their installed packages to an ISO image, which can then be added as a software source on another installation. You can use this source to restore the packages on to the system or keep everything in the APT cache.
aptoncd.sourceforge.net

Clementine

Use Clementine to play locally stored music and streaming audio. The app has an attractive interface and it also helps organise and transfer music to various devices, and integrates well with popular cloud services.
www.clementine-player.org

Icecast

With *Icecast* you can stream music across the network. *Icecast* supports many audio streams simultaneously and listeners can access a stream via a remote media player and also configure MPD as a source.
www.icecast.org

Amarok

If you use KDE your distro may already include this music player, Amarok. It too integrates with several online audio services, and its features include creating dynamic playlists, bookmarks, scripting, context view.
https://amarok.kde.org

LMMS

LMMS is digital audio workstation that produces music by synthesising sounds, arranging samples, and playing them on a MIDI keyboard. It also has a song editor and plugins to simulate instruments and effects.
www.lmms.io

Kodi

Until recently *Kodi* was known as *XBMC*. It's an excellent option for users who wish to turn their PCs into media hubs. It plays most kinds of media files and works with TVs, IR and bluetooth remote controls.
www.kodi.tv

Utilities

Apps that let you do more with your computer.

Gnome Tweak Tool

Not satisfied with the stock Gnome desktop? Use the *Gnome Tweak Tool* to customise several aspects, including the appearance settings of the desktop. With this tweak app you can also change the behaviour of the Windows and Workspaces, manage extensions and you can even circumvent the design philosophy of Gnome 3 by placing icons, files and folders on the Gnome desktop.
http://bit.ly/GnomeTweakTool

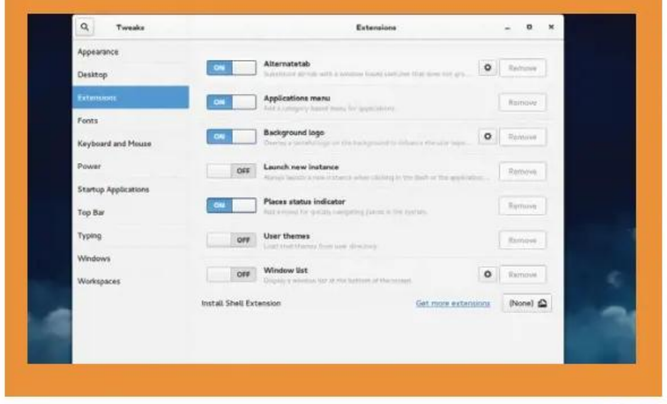

Grub Customizer

Grub 2 is the most popular Linux bootloader that's used by virtually all major distributions. It's an impressive piece of software with lots of options. The *Grub Customizer* is a simple to use graphical tool,which enables you to quickly customise all aspects the bootloader, including its appearance.
www.launchpad.net/grub-customizer

Handbrake

When the need to convert a video arises, *Handbrake*, the video transcoder app does a commendable job. It can convert nearly any format and supports a wide range of video codecs. One of its best features is built-in device profiles for popular devices that make the conversion process easier.
www.handbrake.fr

EasyStroke

Want to control your PC with the flick of the mouse? The *EasyStroke* app lets you define and manage gestures by recording the movements of your pointing device while holding down a specific mouse button. You can then configure actions that'll be executed when the app recognises the defined stroke.
https://easystroke.sourceforge.net

DOSBox

Relive the good ol' days with *DOSBox* and play your favourite classic DOS games that won't run on your modern hardware. This is an x86 PC emulator that creates an IBM PC compatible computer complete with compatible graphics and sound cards. The app can also simulate networking hardware for multiplayer games on the local network and even over the internet. The *Wine* project even uses code from *DOSBox* to bolster support for DOS apps.
www.dosbox.com

digiKam

One of the best photo management tools for Linux is *digiKam* and it has features that'll appeal to all kinds of users. It recognises all major image file formats and can organise and sort images based on metadata. The app also has plugins to export images to various online services.
www.digikam.org

Avidemux

Avidemux is a video editor and converter that can be used for basic cutting, filtering and encoding tasks. It supports many file types, including AVI, MPEG, and MP4. The app is designed for users who know what they want to do but also provides an intuitive interface so that tasks such as cutting and appending videos are pretty straightforward. The app has some presets and users can also save custom settings that make the app easier for new users to operate.
http://fixounet.free.fr/avidemux

K3b

Although it's designed for KDE, the *K3b* optical media burning utility is one of the finest for the job. The app can burn multiple El Torito boot images, audio CDs, VCDs, SVCDs, mixed-mode CDs, eMovix CDs, and DVDs. It can also rip DVDs and write ISO images.
www.k3b.org

Vokoscreen

A feature-rich screencasting app worthy of note is *Vokoscreen*, which is based on FFmpeg for handling multimedia data. *Vokoscreen* can capture both video and audio, with options to record the entire screen, window or a selected region, along with video from a webcam. The app supports MPEG4, x264, MP3 and Vorbis codecs and can save files in either .AVI and .MKV containers. The app offers some controls such as the ability to change the video quality and frames captured per second and can be used to make screencasts of games.
www.kohaupt-online.de/hp

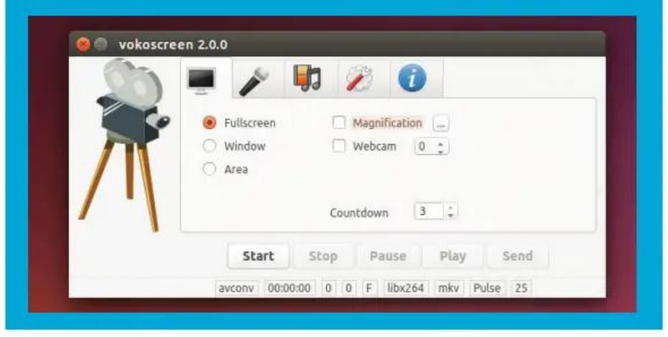

Terminal

Ncmpcpp

This is a command-line *MPD* client that's easy to use and customisable. It provides useful features such as the ability to sort playlists, song lyrics, item filtering, fetching artist's info from last.fm, tag editor and much more.
http://bit.ly/Ncmpcpp

Samba

Samba is a suite of programs that enables Linux users to access and use files, printers and other commonly shared resources on a Windows PC on a network and does this by supporting the SMB protocol.
www.samba.org

rTorrent

Here we have a command-line BitTorrent client with an ncurses interface. You can run it as a daemon and manage it with *screen* and since it supports SSH you can manage your torrents from any remote machine.
http://bit.ly/rTorrent

Links2

There are lightweight browsers and then there's *Links2*. This is a web browser that can render complex pages and even has a pull-down menu. It's also special because it's a CLI browser that you operate via the keyboard.
http://links.twibright.com

Midnight Commander

Before the days of graphical file managers, real hackers used *Midnight Commander*, known as *mc*. It's still your best option if you regularly find yourself in the console environment a lot.
http://bit.ly/MidnightCdr

Admin tools

Take charge of your distro with these power apps.

Redo Backup

We've mentioned the *Clonezilla* cloning solution earlier in the feature, but if all you need is a tool to swap out an old disk for a new one, then you use *Redo Backup and Recovery*. The tool is designed for inexperienced users and has the simplest of interfaces.

www.redobackup.org

XAMPP

The XAMPP stack gives you a single package that you can use as a sandbox to test and develop web apps. It includes all the necessary components such as *Apache*, *MySQL*, PHP, and Perl as well as several other libraries, modules and tools, such as *phpMyAdmin* and *FileZilla* for managing the stack components. Once installed, you can manage the various services via a graphical control panel.

www.apachefriends.org

Déjà Dup

The app's minimal GUI sets itself apart from the various other backup apps you'll find, and it lets you configure backups within a matter of minutes. *Déjà Dup* is based on *Duplicity* and provides just the right number of features for desktop users who aren't used to the ways of a backup tool.

http://live.gnome.org/DejaDup

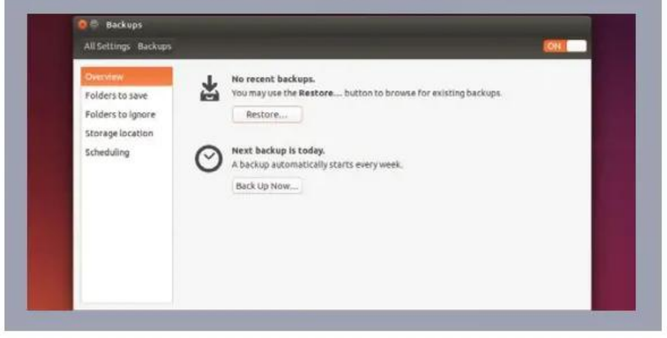

Qemu

It's is a feature-rich multi-purpose processor emulator and virtualiser. You can use it to create virtual machines and even emulate various hardware architectures. If you have the right hardware on tap (a processor with hardware virtualisation extensions), you can use *Qemu* with *KVM* in order to run virtual machines at near-native speed.

www.qemu.org

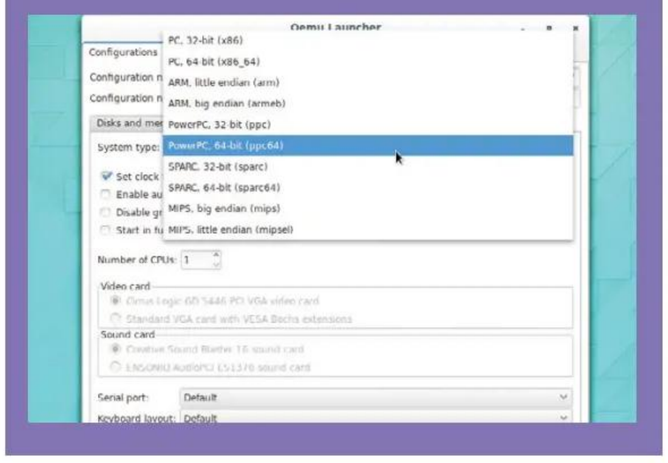

Mondo Rescue

Mondo is a unique backup solution that creates bootable backup and restoration disks customised for the system being used. *Mondo* has a text-driven interface and works with a wide range of file systems and can use a variety of media as backup mediums.

www.mondorescue.org

Open Media Vault

When you need more protection for your data than a simple backup then you need to deploy a NAS server. The *Open Media Vault* project is a Debian-based server that offers the power of commercial options in a way that's easy to setup and manage.

www.openmediavault.org

Conky

Concerned about the resource utilisation on your PC? *Conky* is a nifty little app that lets you keep an eye on your system. It can monitor and report on the states of various components. The tool is very flexible and highly configurable and can also display information from apps, such as weather updates.

http://conky.sourceforge.net

Turnkey Linux

The Turnkey project produces appliances which you can use to deploy a new server in a jiffy. A Turnkey appliance is a self-contained system that packs in a fully functional web app that runs on top of Just enough Operating System (JeOS) components required to power that particular app. All the appliances are based on Debian but are available in several formats depending on the hardware that you want to deploy it on. Once they're up and running you can manage each appliance using a browser-based interface.

www.turnkeylinux.org

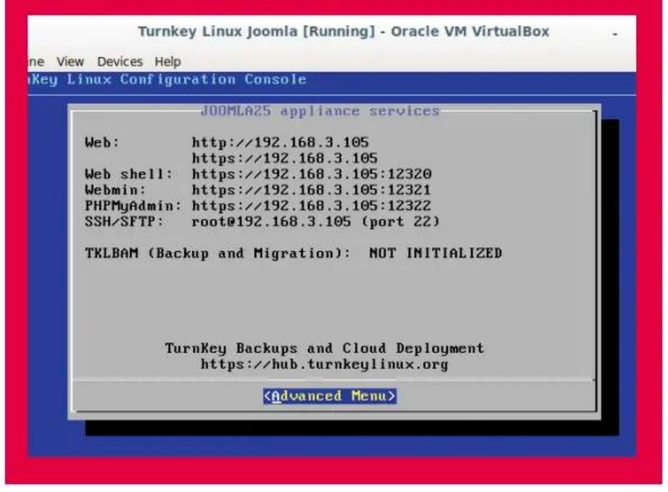

Zentyal

The Zentyal distro has all the components you need to run a gateway server. The distro simplifies the process of setting up, monitoring and controlling the components of the server with a host of custom management tools and helps you configure the servers without mucking about with config files.

www.zentyal.org

Mutt

Mutt is to email what *Links2* is to the web browser. It's a text-based mail client that is highly configurable and it supports both POP and IMAP protocols and has all the usual features you'd want from an email client.

www.mutt.org

Profanity

Profanity is a console-based client for the XMPP protocol that supports multi-user chats and OTR message encryption.

www.profanity.im

Canto

Want to do more from the command-line? Get the *Canto* CLI RSS feed reader. It supports RSS, Atom and RDF feeds and imports and exports feeds in OPML format. It has lots of customisation and even configure it with Python.

http://bit.ly/CantoRSS

mpg123

This is an MP3 audio player for the command-line that supports gapless playback. It's so good that its decoding library, libmpg123 is used by other audio players for MP3 playback

www.mpg123.de

FFmeg

One of the most versatile media conversion utilities, *FFmeg* can manipulate virtually any type of media file in various ways, such as changing bitrate, extract audio, record streams, extract stream and much more.

www.ffmpeg.org

Basics

Package management 44
Web browsers 46
Documents 48
Spreadsheets 50
PDF editing 52
Instant messaging 54
Email 56
Text editors 58

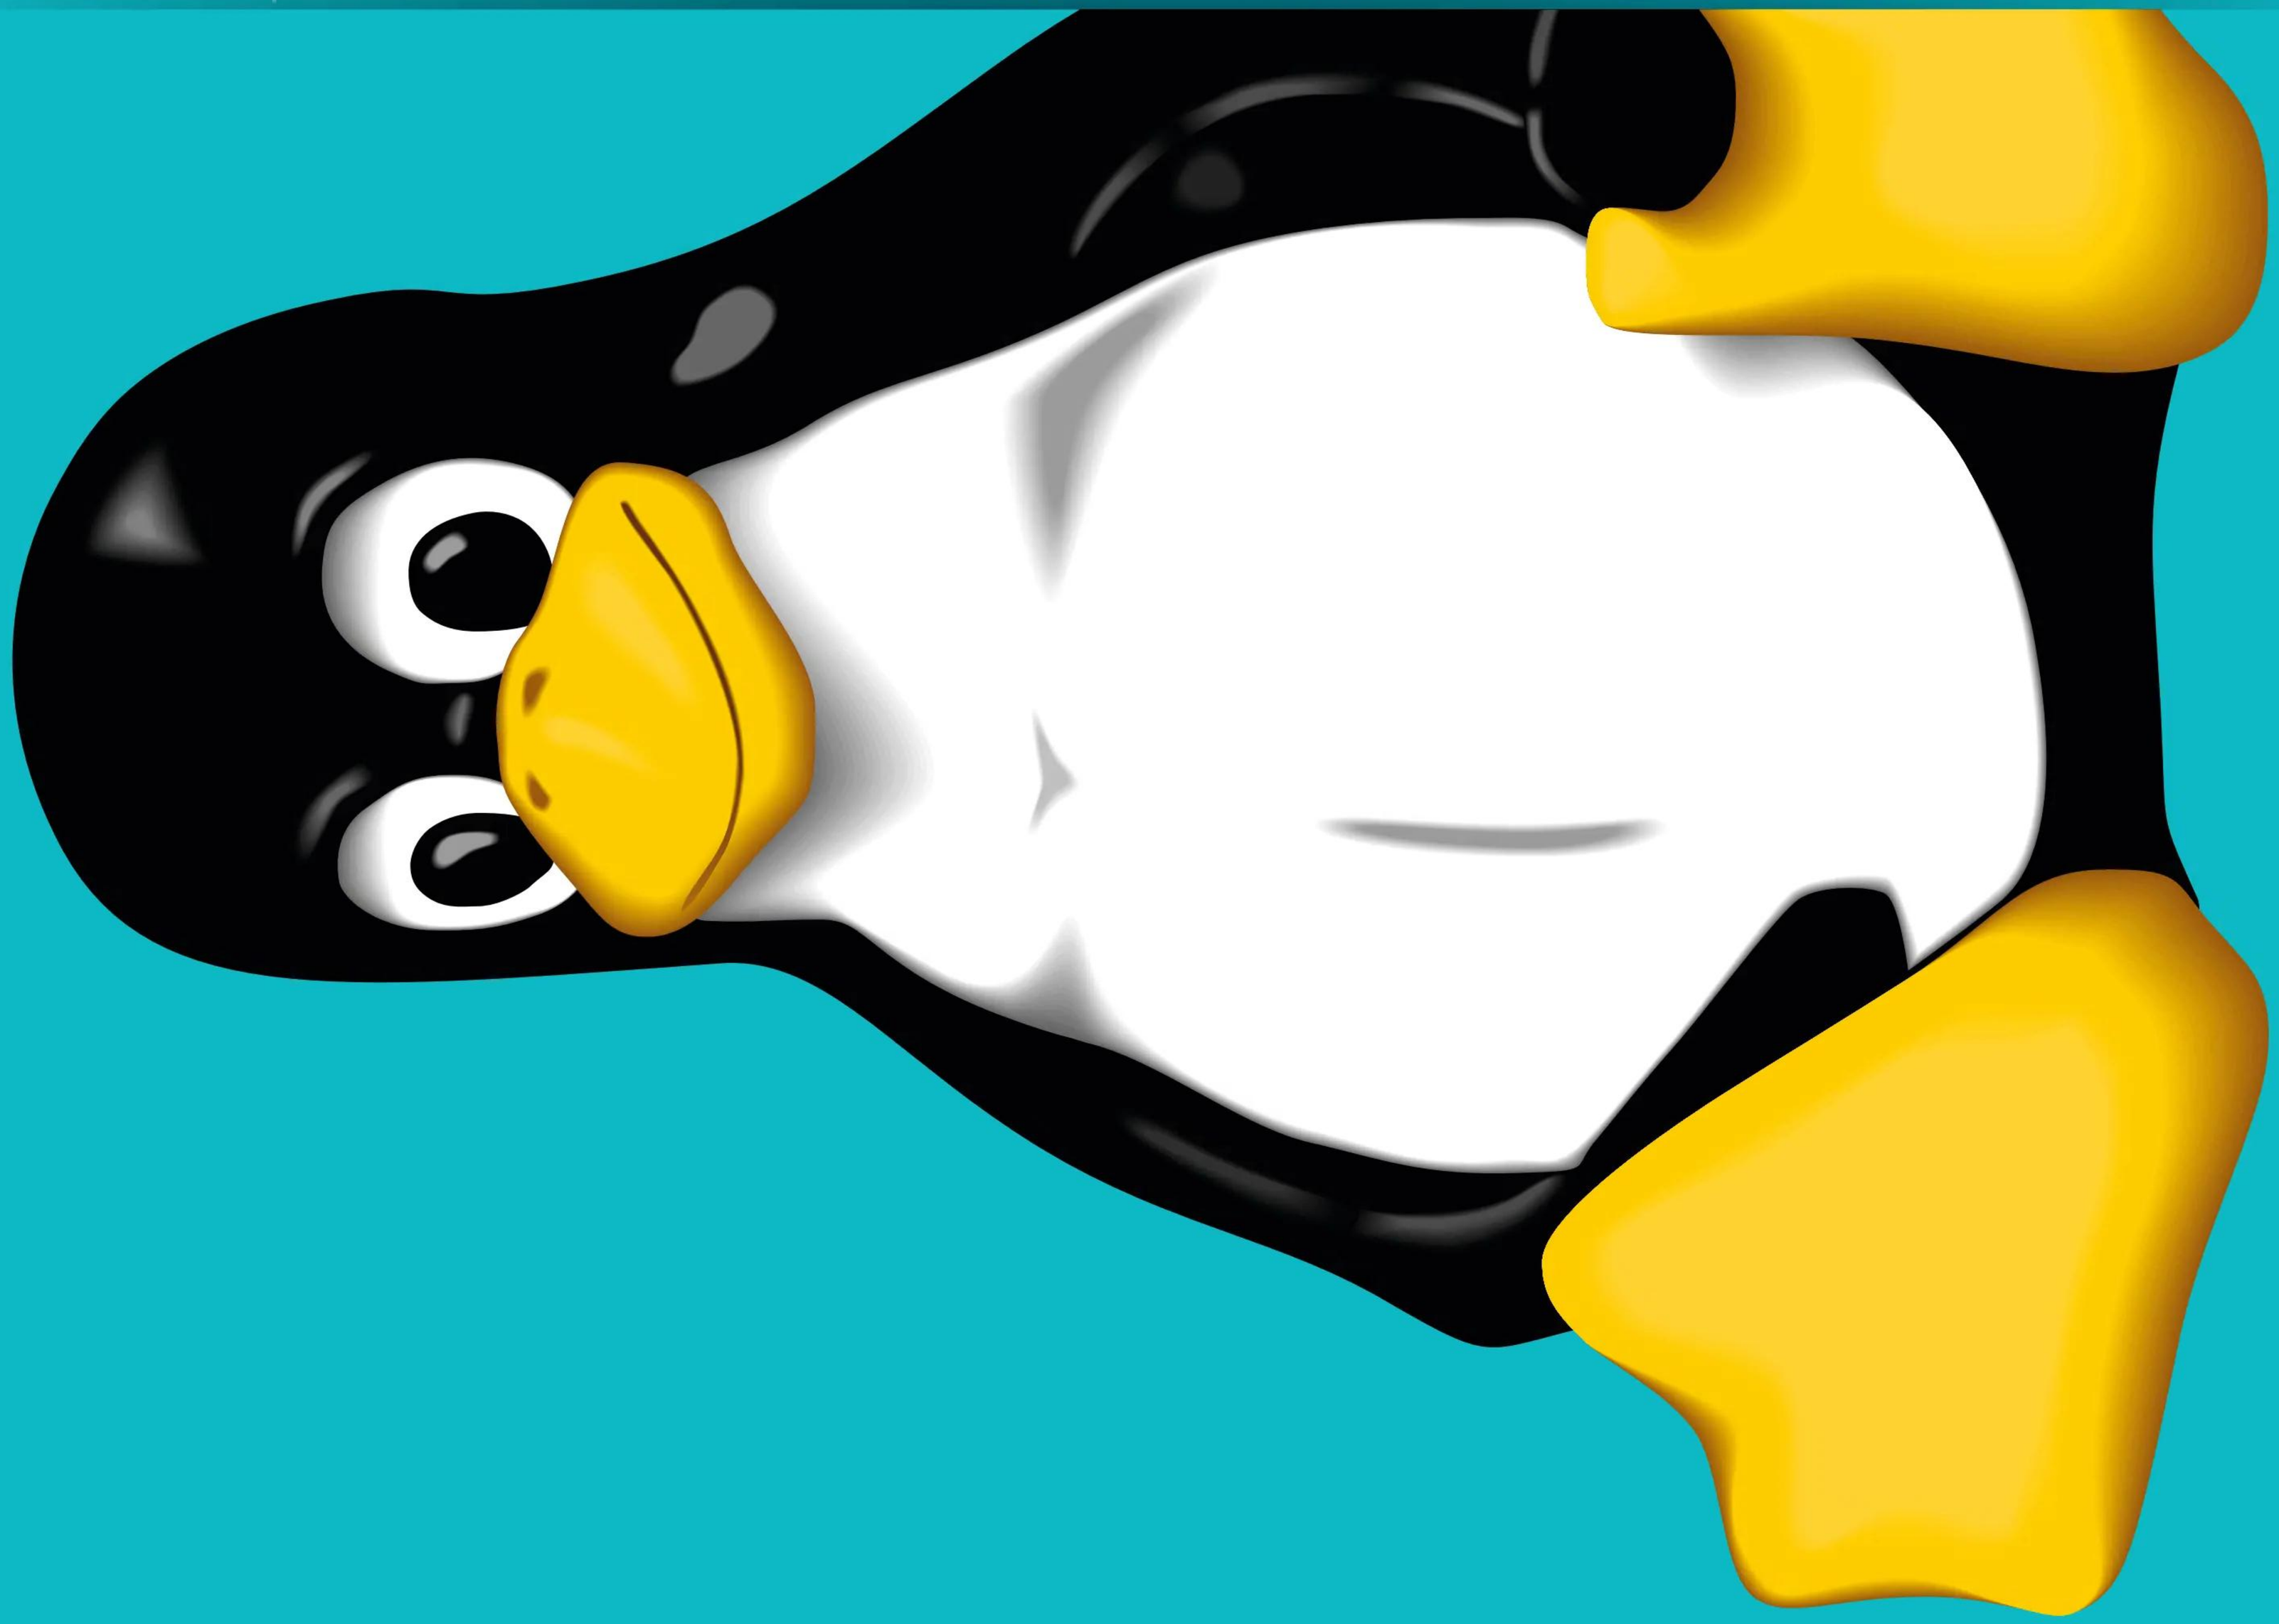

Packages: Installing new software

Getting fresh programs onto your Linux system is, despite what you may have heard, utterly straightforward and foolproof.

Quick tip

You can also install software from Ubuntu's dashboard. Click the search button at the top of the left hand bar, the apps icon at the bottom, and search for what you want.

Your Linux system comes pre-installed with a whole host of useful applications and software, but there will come a time – there always does – that you want something new. There are various complex ways of getting hold of new things, and perhaps they'll be routes you explore later on in your Linux life. For now, there's a much easier way to find new programs for your computer: a package manager.

Package managers deal with everything for you. They retrieve the very latest versions of your desired software from the internet. They sort out any dependencies – items your computer might be missing – that are required for you to be able to run that software. They install everything automatically. And, once your software is firmly ensconced in your PC, they deal with the process of upgrading or removing it later on. You never have to get your hands dirty. And, all being well, everything will work perfectly first time.

You might have come across something similar on your smartphone. Apple's App Store and the Android Play Store are each examples of package managers – it's unlikely you've ever had to manually download and install any apps on a modern phone. The whole process is seamless. Apple has brought the system to macOS, and Windows converts may know that Microsoft has been trying a similar scheme since Windows 8: the Microsoft App Store. Cynically, you might look at that as a way for Microsoft to control its software ecosystem, and you'd probably be right to be cynical, but it's also a way to get hold of (some) programs without the hassle – and potential danger – of downloading random things from the internet. It makes sense. It hasn't really worked, but the theory is sound.

Ubuntu's desktop package manager is called *Software Centre*. It accesses a number of sources – known as repositories – to grab the right version of your chosen program for your computer. Or at least it does for now – Ubuntu will soon be replacing it with an alternative, *GNOME Software*, in the 16.4 release of Ubuntu. But for now, let's look at how you might use *Software Centre* to download, install and remove software – *GNOME Software*'s interface and philosophy barely differ, so this will stand you in good stead.

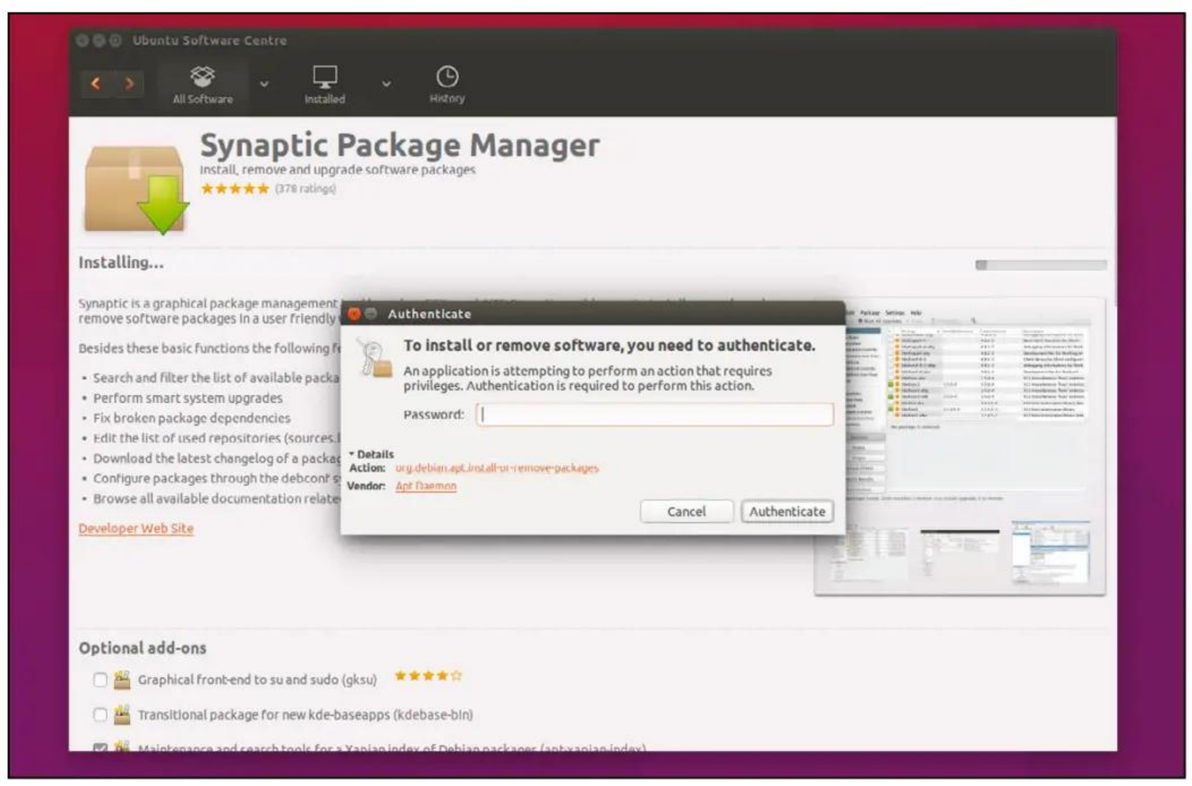

› Installing software requires that you have admin rights to your system – it's a good way of stopping other users from installing software you don't want!

Central reservation

As you can see from the annotation above right, there's not really an awful lot to *Ubuntu Software Centre*'s core interface. A big categorised list of software, buttons so you can see what's available, installed, what you've done and how downloads and installations are getting on. The key element for us is the search box at the top right of the window. It doesn't only search by software names – handy, considering the baffling monikers of most Linux packages – but by descriptions, too. So, for example, if we wanted to find a game, we might search, yes, 'game'. And everything with 'game' in its title, short description, and long description will pop up. Obviously that's a very broad example, but imagine you're looking for a new email client: it makes a lot more sense to type 'email' and find what you're looking for than it does to type 'MBoxImporter'.

Teeth-grindingly annoying nomenclature aside, then. Let's install something. And why don't we get a bit meta? Before *Software Centre*, Ubuntu featured a program called *Synaptic Package Manager*, which we'd wager will still be available after the move to *GNOME Software*. It does the same things (searching for packages, fixing dependencies, upgrading and

Inside Ubuntu Software Centre

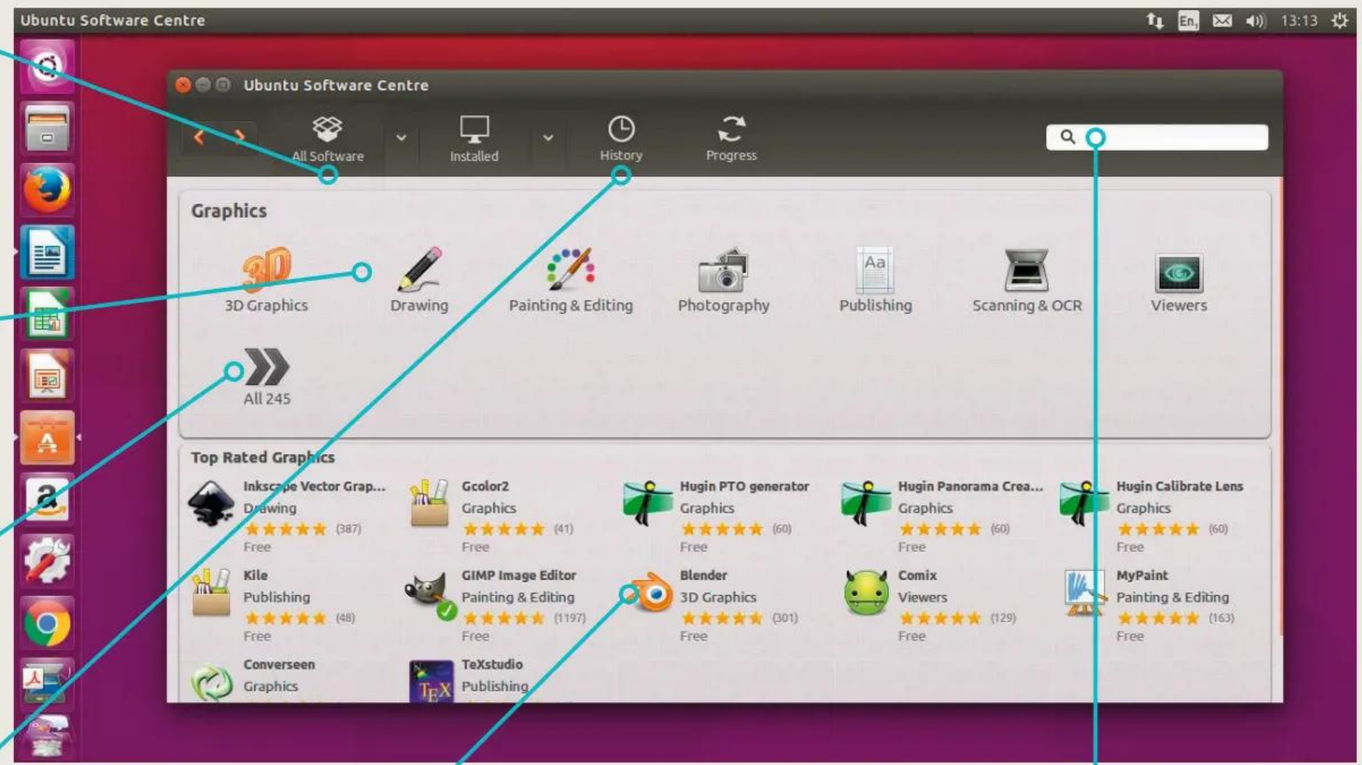

Everything
To return to the main menu of *Software Centre*, click this button. You can also use its menu to select where the software comes from – official Ubuntu releases or from Canonical's partners.

Categories
Once you've chosen a category in All Software, that category's subcategories are displayed, meaning you can drill down further to the software you really want.

See it all
Wondering just what there is to choose from in your current category? Click here to see the whole confusing list.

Past tense
What has *Software Centre* actually installed? Not just the big packages, but the little ones? The history page will let you know.

Top picks
The best (or at least the highest rated) apps in each category are listed here along with their average ratings and number of reviews (in brackets).

Search bar
Searching is vital if you're not sure of a package name, or not sure what category your chosen package is in. Just type in here and hit [Return].

the like) with a lot more information on hand. Perhaps it's not the best for Linux newcomers, but it's handy to have around when you take the next step.

Typing 'package manager' into the search box doesn't do us much good. Synaptic is in the list but it's not easy to find. So hit 'All Software' and we can find it in a more organic way. It's actually nestled in the typically opaque category 'Themes and Tweaks', so click that in the left hand column of the main page and you should see it listed in the top rated section of the window that follows. If not, click the All button at the top and you should see a much more manageable list.

Actually Installing

Click on *Synaptic Package Manager* and you'll see a host of information about it, a link to its web site, a screenshot (isn't it pretty?) and, on the right, an Install button. Before you go ahead and click it, scroll down a little – *Synaptic* comes with a host of add-ons that you might be interested in, and these can be installed simply by marking their individual check boxes. Scroll a bit further and you'll find reviews of the package – very useful if you're installing an unknown bit of software, as fellow users are rarely shy about sharing negative experiences – and *Software Centre* doesn't exactly have a stringent cruft filter – so these can be a big red flag in a lot of cases. When you're ready, hit Install. Type your password, hit Authenticate, and *Software Centre* will do the rest.

Head to the History panel to see what changes have actually been made. You should see that a package called synaptic has been installed, but look at items installed at a similar time; the list will include a host of *Synaptic*'s dependencies that have also made their way onto your system without you having to lift a finger. Beyond the finger you used to click the install button, that is.

Once removed

Let's reverse the process, while we're at it. Click the Installed button at the top of the window to find a list of software divided by category. We know that *Synaptic Package Manager* is in Themes and Tweaks, so click the arrow to the left of that category to expand it and locate *Synaptic* in the list. Click it once and you will see a pair of buttons appear – More Info, which is especially useful if you're weeding your system and can't remember what a particular bit of software actually does, and Remove. Click the latter.

If applicable, *Software Centre* will now warn you of any other things that also must be uninstalled at the same time. In this case, it's a package called update-manager, which we can easily reinstall later using the same technique we used to install *Synaptic Package Manager* in the first place. Click Remove All, enter your password, and then the removal will happen – but be sure to check the History tab for confirmation. Now go ahead and experiment with *Software Centre*. Install some games, remove some of the default apps and go crazy! ■

Quick tip

Most graphical package managers are just a front-end for terminal apps like apt-get, which you can use directly. There's more on this on p142.

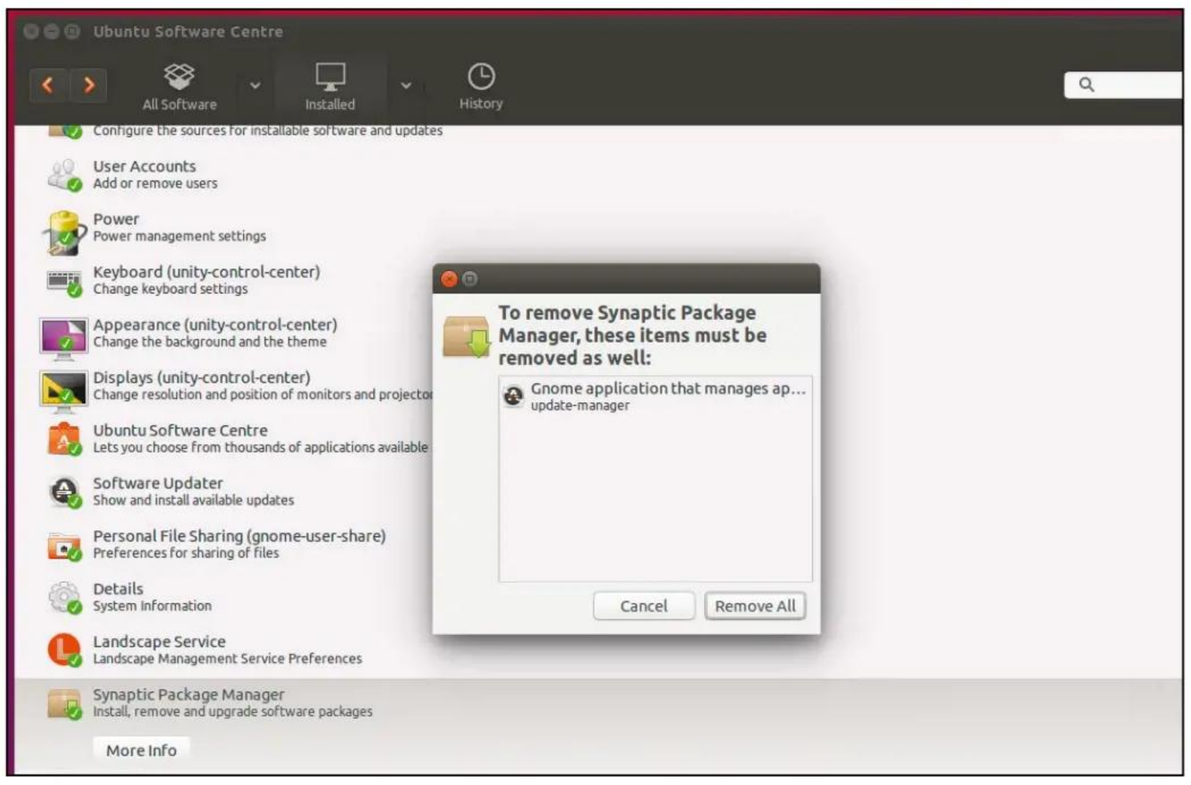

› Otherwise-redundant dependencies will be automatically removed when you uninstall packages – and sometimes you'll have to remove additional things too.

Web browsers: Pick the best

How do you see the web? You might be surprised at the difference when you check out the best browsers available.

Quick tip

Each browser has a different way of importing bookmarks – check your chosen browser's help file for more details on how to do it.

What is a computer if not a portal to the internet, these days? Accessing the online world is absolutely essential, and Linux is not short of web browser options. Many of them will be familiar if you're used to going online in other operating systems, some are variations (based on the software, but not exactly the same) and some are fairly unique, unusual browsers to try as well.

The good news if you're fairly established with your current browser is that it's usually rather easy to transfer your bookmarks over, so you won't lose those vital links. So before we look at your options, let's back these up from another system so that we can import them later on.

Backing up bookmarks

In *Google Chrome*, click the three-lined hamburger menu icon, go to Bookmarks, and click Bookmarks Manager. Click the 'Organize' drop-down menu and select Export Bookmarks. This will create an HTML file which you can later import into the browser of your choice.

If you're using Internet Explorer you'll find the setting in File/Import and Export/Export Favourites. *Microsoft Edge*, despite being newer, isn't quite as friendly, but thankfully *Firefox* can import *Edge* bookmarks. To get them (and the rest of your bookmarks) out of *Firefox*, click the bookmarks icon, Show All Bookmarks, and follow the steps within Import and Backup.

Sadly we're not made of space here, so users of other browsers will have to hit up their favourite search engine. Time to look at your options in Linux!

Browser history

The most established browser on Linux is probably *Firefox* which, although it's lost a little market share in recent years, displays the entire web perfectly and includes some neat features. Newer versions allow you to synchronise your tabs between devices if you sign in to all of them with the same Firefox Account, for example, so it's easy to pick up where you left off if you're jumping from machine to machine – or even user account to user account. Similarly, you can use the Firefox Hello system to browse the web simultaneously with other *Firefox* users, which is great for remote working, planning, or just hanging out online.

You may also find a browser named *Iceweasel* – this is, for all intents and purposes, *Firefox* in all but name. Certain minor elements and trademarks attached to *Firefox*, such as the logo, have been changed due to them violating the

Alternative Browsers

Lynx

Classic text-mode browsing, run from within a terminal window. *Lynx* is mostly a curiosity these days due to its archaic makeup, and it's only really useful on machines which absolutely cannot output anything other than text mode, but it can be interesting to see how the web looks without images or modern rendering techniques.

Midori

Using the GTKWebKit engine (a fork of the Blink rendering engine that powers *Chrome*, *Chromium* and *Opera*), *Midori* is the default browser of the xcfe window manager and features in a number of lightweight distributions (Bodhi, SystemRescueCD, Manjaro Linux). It does well on the tests we've listed in this article.

NetSurf

Originally coming from RISC OS – an ARM-based operating system that grew out of the Acorn Archimedes – *NetSurf* has been ported to a number of platforms since, including AmigaOS. It's not the absolute best browser, but it uses its own unique rendering engine, so if your Linux install struggles with others, give it a try.

Opera

So often the browser bridesmaid but never the browser bride, *Opera* is a clean, fast and quite wonderful program that deserves more recognition than it gets. The Opera Turbo feature, which compresses your web traffic, is particularly useful if you've got a slow internet connection or if your ISP blocks access to certain sites.

Pale Moon

A notable fork of *Firefox* that will please those who aren't too happy about the modern browser style, this clings to an old-school, fully-customisable user interface and uses its own rendering engine, Goanna, which forked from Gecko, the *Firefox* rendering engine. It's actively updated and high-performance, so well worth a look.

requirements of free software distribution. Also rebranded were email client *Thunderbird* (*Icedove*), 'internet application suite' *Seamonkey* (*Iceape*) and seemingly abandoned calendar app *Sunbird* (*Iceowl*). Expect these frigid alternatives to disappear soon, as distro creators seem to be working out their differences with *Firefox*'s parent company Mozilla.

Second to *Firefox* is *Chromium*, an open-source browser (available in a number of user-developed flavours) that forms the basis of *Google Chrome*. It doesn't include all of the features of its big trademarked brother – usually missing things which use proprietary non-open technology like Flash player – but it runs the same rendering engine, which means web sites look just as good. It also manages memory in the same way, allocating a portion for each of your open tabs, which is both a blessing and a curse. A blessing because if one tab has a problem it doesn't bring down the rest with it, but a curse because *Chromium* requires a relatively large amount of memory to run multiple tabs.

While you're unlikely to find big-boy *Chrome* in your distro's package manger, it's simple enough to get hold of the real thing. Just head to **google.com/chrome**, click the Download Now button, and select the appropriate file for your distro (usually 64bit deb, if you're using Ubuntu). Click Accept and Install, then open the resultant file with your package manager to get it installing. Annoyingly you'll need to run it from a terminal the first time; launch one by clicking the search button, searching for 'terminal' and selecting the Terminal application, then type `google-chrome-stable` to start it up for the first time. Once it's run once, it'll add its icon to your launcher so you can start it more easily in the future.

In both cases here we'd suggest it's best to run *Firefox* and *Chrome* over their open source alternatives. While this doesn't match the freedom philosophy of Linux, you're probably not in this for philosophical reasons. Having the absolute latest version of your chosen web browser is vital for security and compatibility reasons, and sometimes forked software just doesn't keep up with the times.

There are numerous other browsers available for Linux which all serve different purposes. Well, OK, usually they're there to provide you with a browsing experience that requires a smaller memory footprint for lower end machines. You'll find out more about these in the 'Alternative browsers' box.

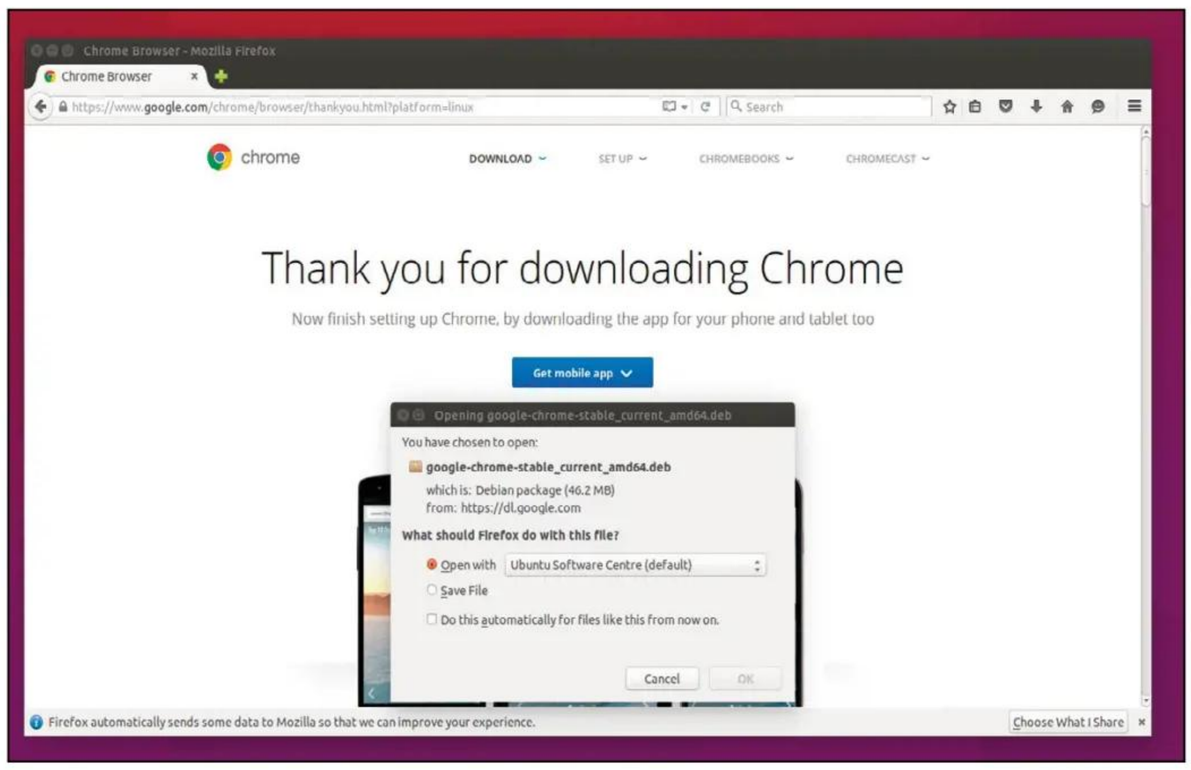

› Chrome probably isn't available in your package manager, so you'll need to head to its website (in a different browser) to get hold of it.

Testing capabilities

How well does your browser actually perform? Different browser options render online elements in different ways, so it's worth testing your chosen browser (and its competitors) thoroughly to make sure you're getting the best experience.

Head first to the Acid tests at **www.acidtests.org**, a long-standing series of rendering engine-stretching challenges which will prove how well your browser adheres to particular web standards or, at least, it shows how well your browser performs these particular tests.

A more modern benchmark is the HTML5 Test, found at **html5test.com**. While this doesn't animate any pretty pictures or give you a visual representation of your browser's flaws (boo) it does give you a detailed list of everything your browser can or can't do, with a score to match. Many Linux browsers are going to lose a few points for leaving out closed source technologies, so bear this in mind before you rush to uninstall your current option. The vast majority of websites include fallback options away from proprietary software anyway, and many of these are now becoming the norm rather than the alternative.

To the extreme

The most extreme rendering test we could find lives at **web.basemark.com**. Browsermark, for that is its name, runs a series of heavyweight tests that push your browser's performance in a number of ways, beyond just HTML5 or CSS, and might also give you some clues about the performance of your PC. Make sure your machine doesn't go to sleep while it's running, which should take about five minutes. When it's complete, you'll be given a final score, see which percentile your browser and desktop combined sits on, and even get a bit of information about the performance of your ISP.

While we've made some vague recommendations here, your choice of browser is really going to be down to personal taste and what works best for your Linux system. All computers and desktops perform differently, so you're encouraged to experiment, install, test, and find the best choice for you. Happy browsing. ■

Quick tip

Don't fret if your browser doesn't max out every test – if you're happy with the way it displays pages, that test result doesn't really matter!

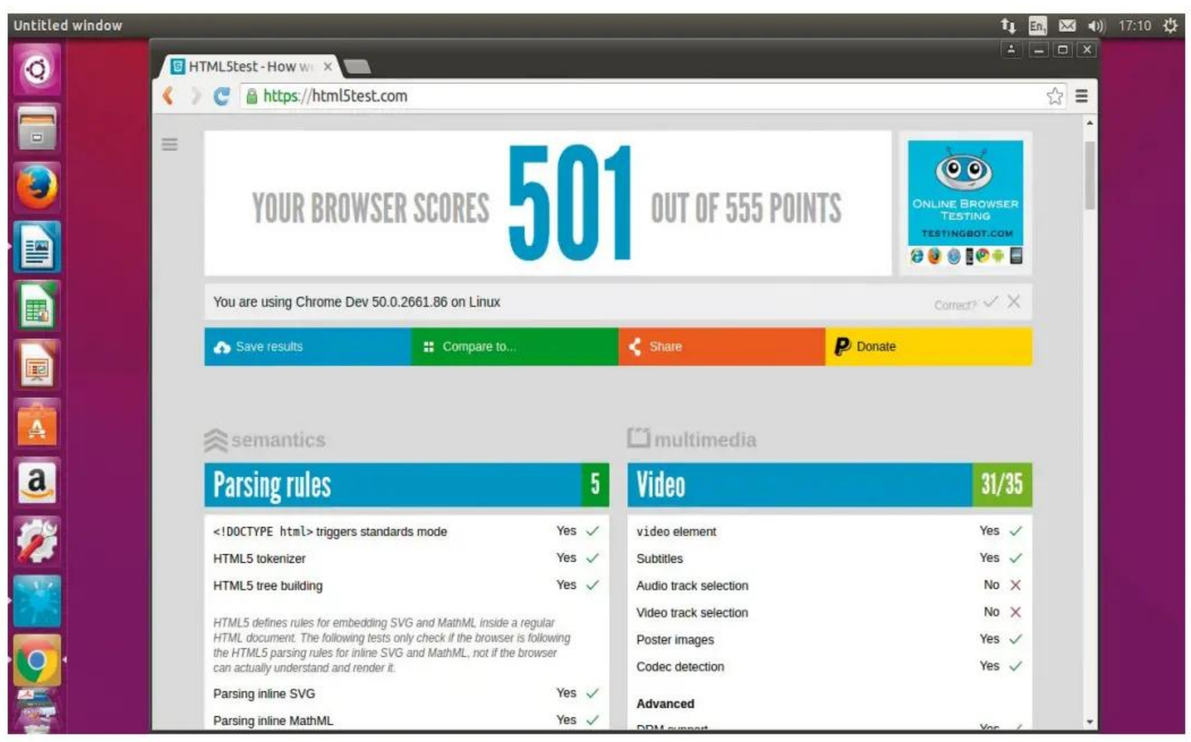

› The HTML5 test is a great way to see how muscular your browser truly is, but it's not the be all and end all of browser performance.

LibreOffice Writer: Create awesome docs

LibreOffice is a full-featured office suite that's completely free of charge. Let's see how Writer, its word processor, handles something complex.

Quick tip

Writer can pull in content from your *Calc* spreadsheets. Just copy the cells you want to import, then use the Edit > Paste Special function in *Writer* to bring them in.

Anything Windows can do, Linux can do… just as well. Not as catchy as the classic song, we'll admit, but it's the truth: as a mature desktop operating system, Linux is blessed with some solid software which can perform everyday desktop tasks with absolute aplomb.

Our recommendation in the office stakes is *LibreOffice*, a cross-platform collection of productivity applications which has been around, in many forms and under many names, since 1985. Despite its venerable legacy, it's still the most actively-developed office package out there.

The Write Stuff

LibreOffice 5 should come preinstalled with your Linux distribution - you can check that you have the latest version by running *Writer* and going to Help/About LibreOffice. If it isn't installed, open up your package manager and install it from there. You may be given the option between *LibreOffice Fresh* (the cutting-edge release) or *LibreOffice Still* (the stable, enterprise-ready version). Which you choose is entirely up to you; nothing that we cover here requires the very latest release. When it's all safely installed, run *LibreOffice Writer*.

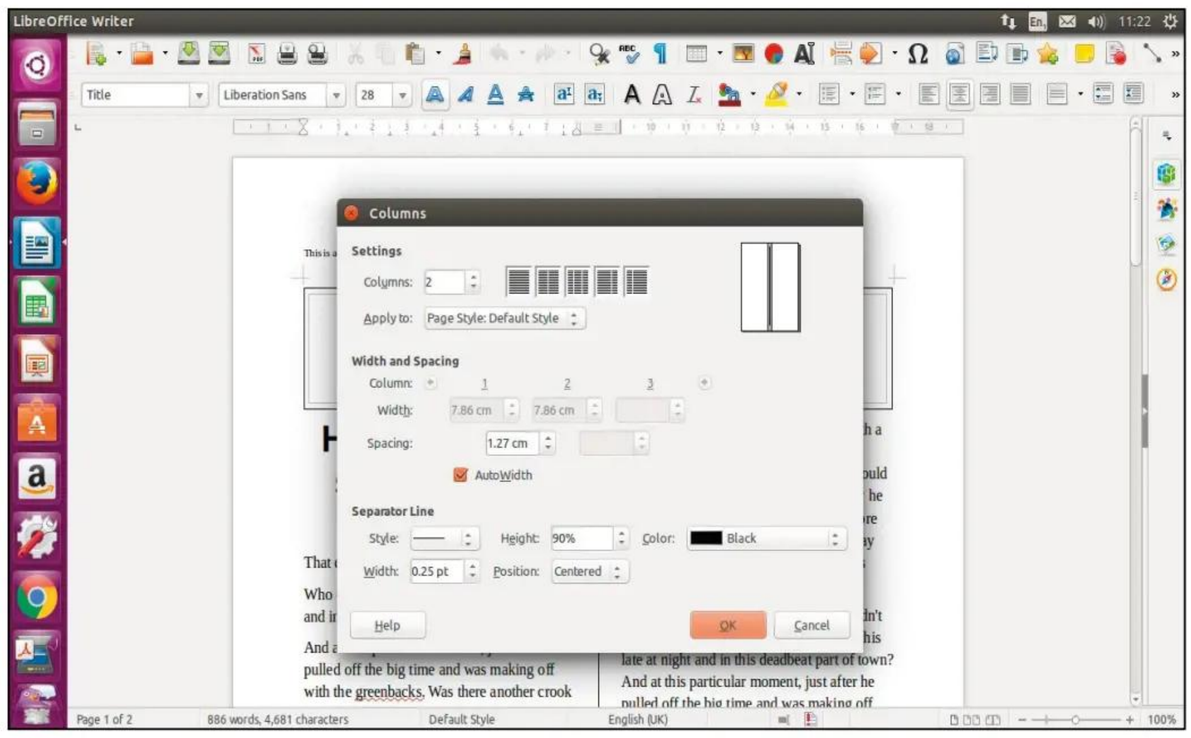

› The columns dialog box has a huge number of options in it, including, if you wish, columns of differing sizes on a single page.

Although Microsoft has changed *Microsoft Word*'s interface – for the worse, some say – in recent years, *Writer*'s layout should be instantly familiar to anyone who's used a word processor since the early '90s. There's a big white empty page. A pair of bars at the top of the screen offer quick access to text and document settings. There's a bar at the bottom of the window with information – a word count, character count, page count – and a zoom slider. And there's a flashing cursor, daring you to write something.

The act of typing is a little beyond the scope of this tutorial – we'll assume you can manage that – so let's create a basic newsletter to test a few of *Writer*'s features.

Layout: the facts

Start by splitting your page up into columns, into which you can type your text. In the menu bar at the top of the screen, click Format > Columns to bring up the columns dialog, then change the number of columns to two either by typing a '2' in the box, clicking the up arrow on the box, or clicking the two-column pictogram on the right of the box. You can now set a column spacing. You'll want this to be greater than zero, otherwise the text in your columns will merge into one, so set it to 0.50cm. If you prefer to use other units of measurement, you can type them in directly; entering 0.2in into the spacing box will automatically convert it to its metric equivalent.

Now let's add a separator line to break the columns up. Set the style to solid, the height to 90%, the colour to a dark grey and leave the width as it is. Again, you could put your own units of measurement in this box, although it's limited to a maximum width of 9pt. Because we changed the height, we can also change the line's position: centred looks best. Click OK to return to your document, and you'll see your changes have been made.

So it's time to type. Or it would be, but that sounds a bit too much like hard work. So instead we're going to fill our document with dummy text using *Writer*'s handy Auto Text feature, which can also give you quick templates for CVs, covering letters and the like. Go to Edit > Autotext, expand the Standard category, select Dummy Text, then click Insert to add your ream of filler text into your columns.

Linux Formatting

Back in your document, duplicate your dummy text a few times (select it all with [Ctrl]+[A], copy it with [Ctrl]+[C], then paste it with [Ctrl]+[V]) to fill up the page. We're going to style up the document now, using *Writer*'s default section styles. Click and drag your mouse to select the first sentence, click the drop-down box that currently says 'Default Style' and change its style to Title. You'll see that this applies a number of changes to the text at once: it centres it, changes its size to 28pt, and gives it a nice Bold sans-serif font. Select the rest of the text and give it the style 'Text body'.

As you hover over styles in the list, you might notice a menu button to their right. Hover over 'Title', hit that menu button then select Edit Style. In the popup that appears, tick the Autoupdate checkbox. We're now going to highlight the main benefit of using styles rather than changing fonts and sizes manually – you can edit the style once and your changes will be reflected throughout your document, meaning everything stays nice and consistent. Head to the font tab and change it to your liking, then hit Apply to see your changes reflected in your document. You can also use the tabs to set up perfectly formatted numbered lists, consistent highlights, drop caps and much more. For a slightly more visual way of doing this, manually tweak some text on the page until it looks just so, drag to highlight it, then hover over your intended style in the drop down box and select 'Update to match selection'.

Finishing touches

Our newsletter needs a header, so go to Insert > Header and select Default Style. This pops a seperate header section above our columns and automatically scales our seperator line to fit. Click your mouse in the header area to move the cursor there and type whatever text you want to include. If you'd like to add an auto-incrementing page number, put your text cursor where you'd like it to be and go to Insert > Fields > Page Number.

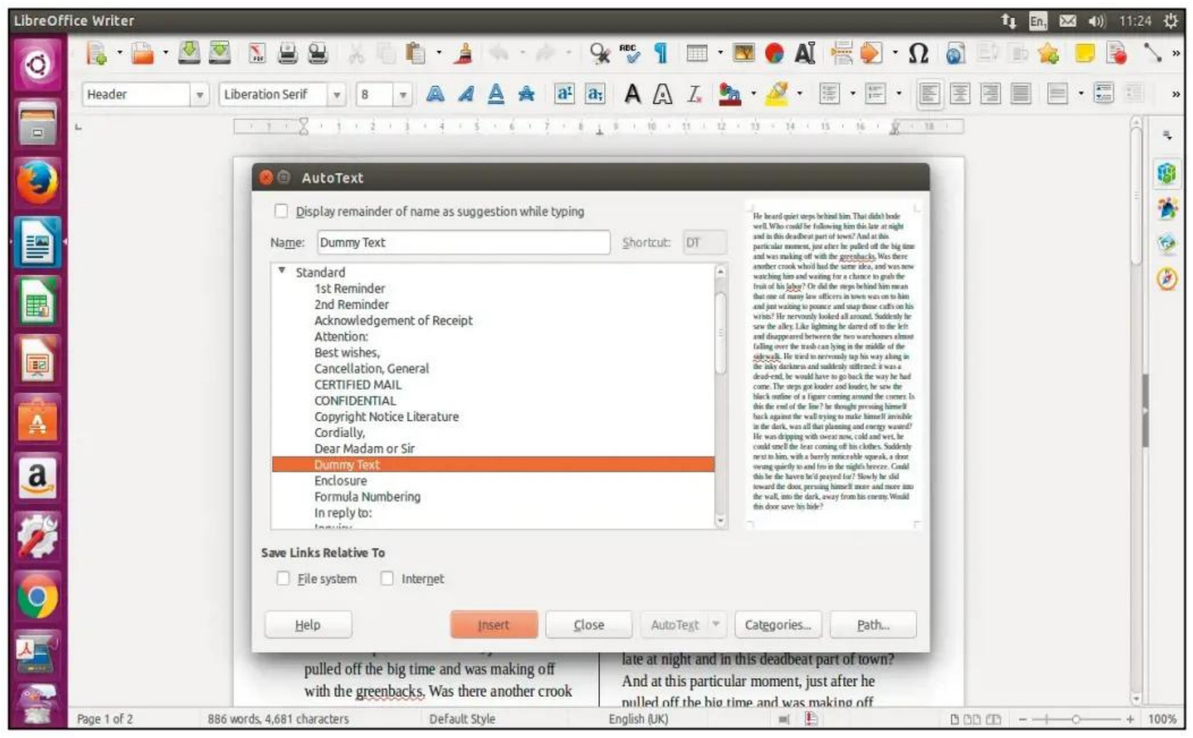

› Use the AutoText function to populate documents as you're designing them, even if you've got nothing written yet.

The header will appear on each page of our document, so we will need to use a slightly different tactic if we want to add a masthead to the first leaf – a frame. Click Insert > Frame to create one, and bear in mind that we're looking for 'frame' as opposed to 'floating frame' in this case, as the latter will sit on top of your text, whereas the former will push the page text out of the way. Stick with the defaults and you will see a small frame appear on the page. Now move it into place by hovering your mouse over its edge until the cursor turns into a fist and dragging it where you want it, and resize it by dragging the green anchor points around its edge until it's the perfect size. You can click inside and type as usual.

The rest, as they say, is up to you. Try using the image button to add pictures (make a frame first for ultimate control), and using the various paragraph tools to further style up your text. And don't forget to hit [F1] to see *LibreOffice*'s comprehensive help guide. ■

Quick tip

Save your documents in .docx format (or, alternatively, the older .doc format) and they'll be fully readable by people using Microsoft apps.

Getting around Writer

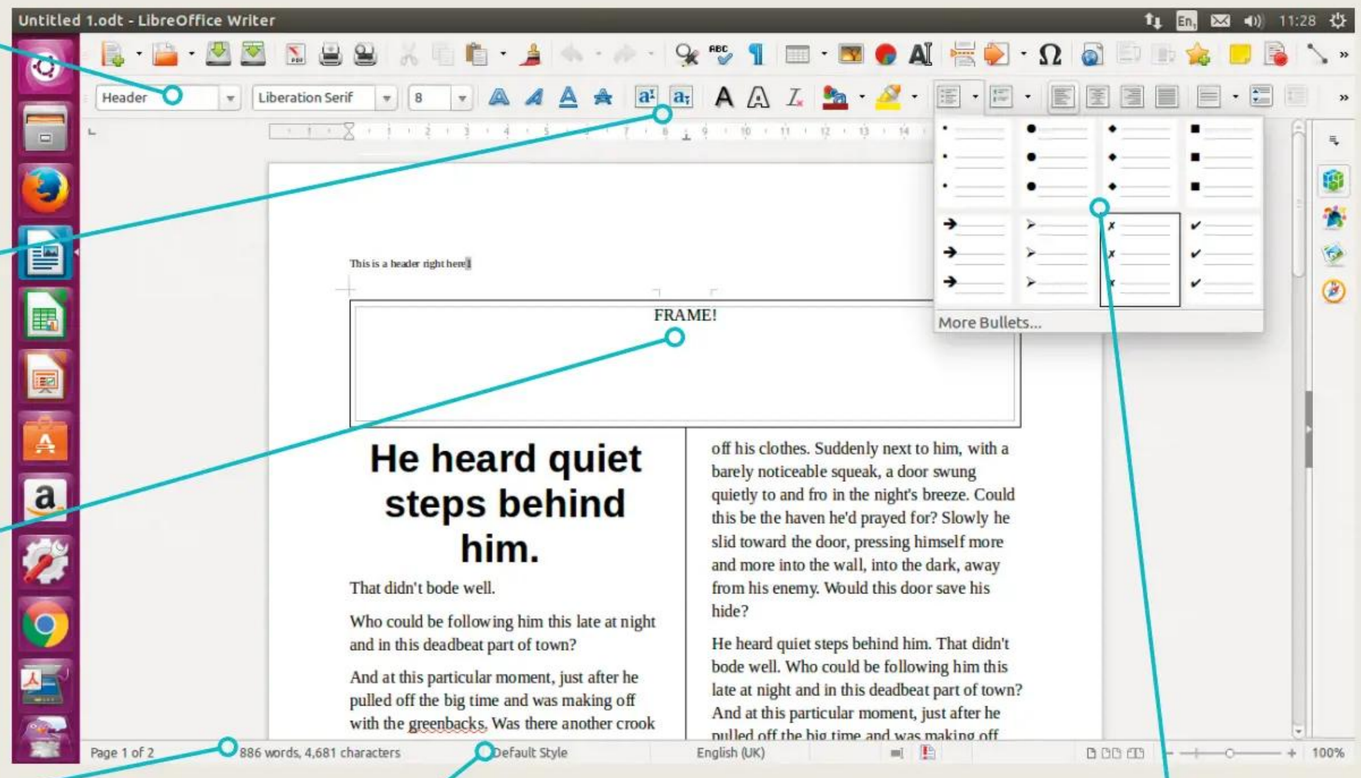

Style tips

Text styles, accessible through this menu, are vital if you want to create consistent documents without a lot of fiddly work.

Look sharp

Bold text, underlining and strikethrough probably don't warrant a text style of their own, so just highlight the text to change and hit these buttons to alter them.

Frame up

Text frames can help you place items (including images) that would otherwise get pushed around by page margins or columns.

How are you doing?

This quick counter keeps track of the number of words and characters you've written – if you're writing to a limit, it's essential.

Quite correct

Spelling and grammar errors are highlighted by a jagged underline, just as they are in Microsoft Word. Hit [F7] for a more thorough spellcheck, or [Shift]+[F7] to turn this feature off.

Extra menus

If you're adding bullet points or numbered lists, don't forget to choose your exact options by clicking the drop down list.

Calc: Super spreadsheets

We've worked with words, now it's time to keep your numbers in check using Linux's premier spreadsheet, LibreOffice Calc.

Quick tip

If something goes wrong, you may be presented with an unintelligible error code. Highlight it and look at the bottom right of the *Calc* window to see what it means.

Where would we be without spreadsheets? Either considerably poorer or considerably richer, most likely. Whichever is true, there comes a time when we all need to put numbers in a grid. You're probably familiar with Microsoft's *Excel*, which rules the spreadsheet roost over on Windows. On Linux there are, as ever, plenty of options, but we lean towards the well-honed spreadsheeting machine that is *LibreOffice Calc*. It should come preinstalled with your Linux distribution, but if it's nowhere to be found you'll need to install it – as part of the main *LibreOffice* package – from your distribution's package manager. In the case of Ubuntu 15.10, you'd use *Ubuntu Software Centre*. We're looking at *LibreOffice 5* for this tutorial, but if you have an earlier version you'll be fine to follow along.

As with *LibreOffice Writer*, running *Calc* for the first time should bring up a reasonably familiar environment. There's a standard grid of cells, menu bars at the top of the screen, below them the input line, and an area below the sheet where different pages can, once they're created, be accessed.

Realistically the way you use it is no different than you might have previously used *Excel*. Even the syntax is the same. But if you're unfamiliar, or if you need a refresher, let's create a sample spreadsheet to see what we can do with *Calc*.

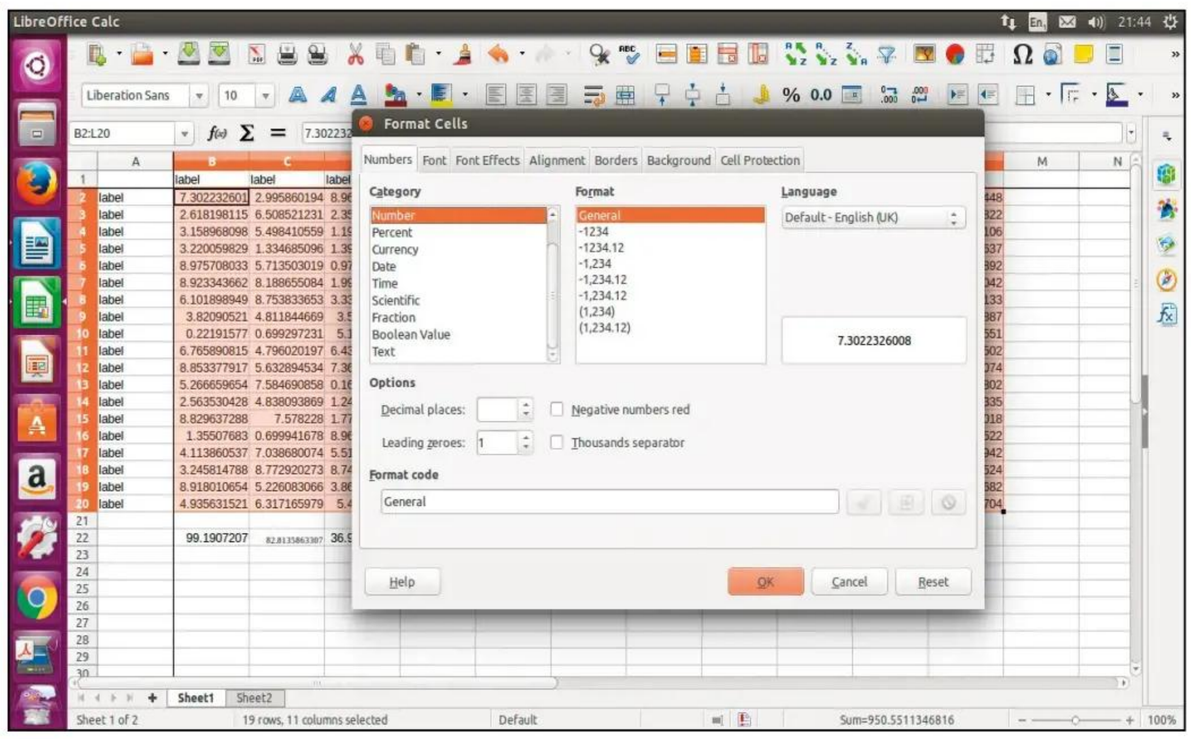

› You can format the numbers in your cells any way you like without affecting the original values. Try representing floating point numbers as fractions!

Numerous numbers

Let's start with a little housekeeping. Usually the leftmost and topmost cells – row 1 and column A – are used as labels to signify the contents of the cells they're next to. Fill in a bunch of sample labels – it doesn't matter what you type, we're merely playing with some of the features – by clicking in each of these edge cells and typing. You can use the arrow keys to navigate between them, which makes the process a bit quicker. You can also hover your mouse cursor over the bold dot in the corner of a cell and drag it to automatically copy the contents of that cell to those that you highlight. Fill in labels from A2 to A20 and from B1 to L1, leaving cell A1 empty – it's useless when we're laying out a spreadsheet like this.

Larger spreadsheets than ours may have content which goes further than your screen can easily display, so it's good practice to isolate the label columns in their own pane. On the right of the screen, just above the scrollbar, you'll see a small horizonal line – drag it downwards to until the top pane shows just the labels, and do the same for the small vertical line to the right of the bottom scrollbar. Hover your mouse over the icons on the right hand side of the top row of toolbars until you find the one labelled 'Freeze panes'. Click it, and your labels will be neatly locked in place.

Now it's time to fill in some data! But let's use one of Calc's features to generate our own random data to play with, rather than using anything useful. Click and drag over cells B2 to L20 to select them, then head to Edit > Fill > Random Number. You'll see your chosen cell range has been automatically inserted in the topmost box. We now need to choose what kind of random number to generate, and its range. By default it'll be a floating point number, although you can change the distribution to Uniform Normal if you want to make a whole number. We'll stick with standard Uniform for now, as it will allow us to demonstrate a further function later on. Leave the minimum at 0.1000 and change the maximum to 9.0000. Click 'Apply' and you'll have a nicely filled in spreadsheet.

Function room

Time to crunch those numbers a bit. The most basic spreadsheet function is a sum, which performs basic arithmetic on a group of cells. Let's start by adding the first column of numbers together. Click below it in cell B2, type

Cracking the Calc code

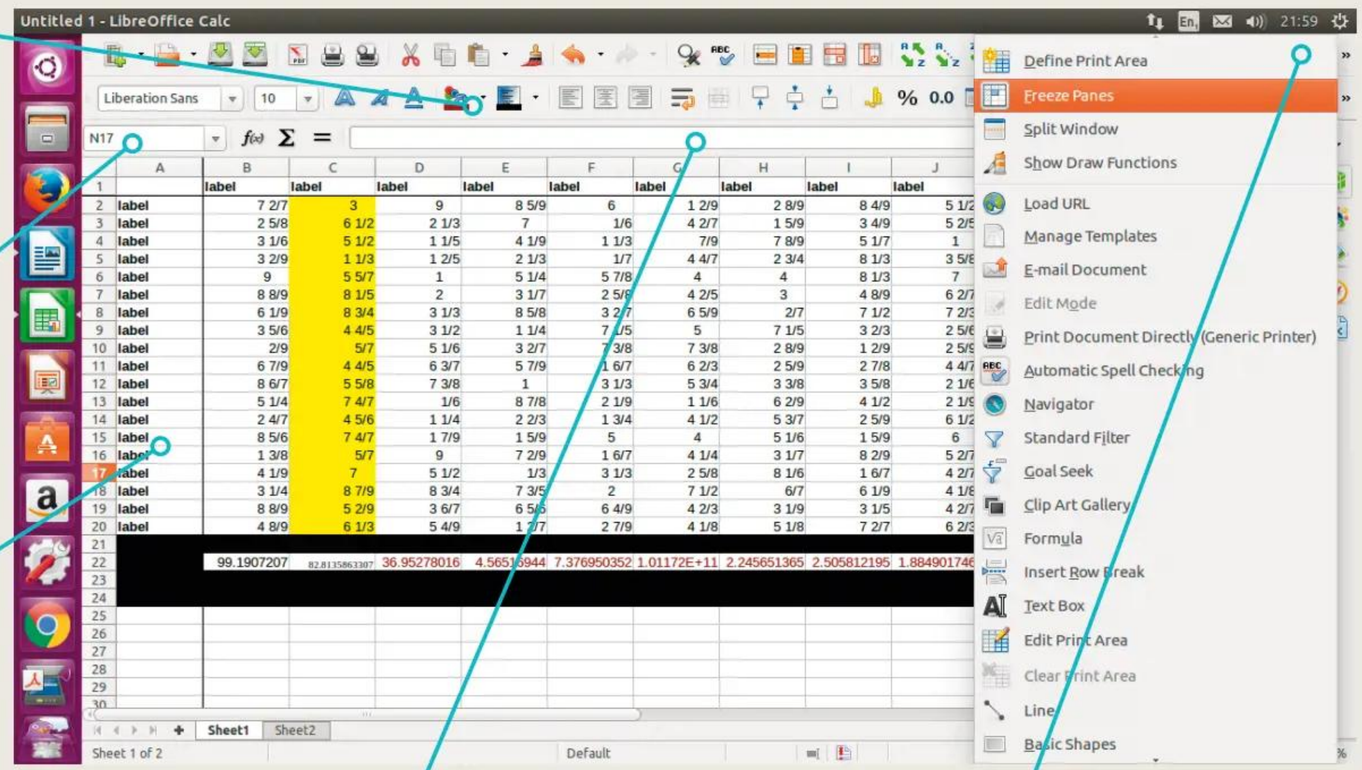

Colour up
Calc features standard the standard formatting tools you'd expect, so you can change fonts, sizes, colours and more.

Cell mate
The currently selected cell is always shown here – bear in mind that you can use the arrow keys to navigate between cells rather than clicking.

What a pane
Locking your panes like this gives a clean look and means your labels will always stay in place no matter where you are in the document.

Line up
You can type directly into cells, but it's often useful to use the input line, which always reflects the content of the cell.

Extended play
If all of the icons on your toolbar don't fit on the screen, click the small arrow double icon on the right to display an expanded menu like this one.

`=SUM(` then click and drag your mouse over the numbers in the column above. Pointers to the first and last cells will be added, with a colon between to indicate a cell range. Close the brackets with a `)` then hit [Enter] to see the total of all the numbers in that column.

Let's try some different functions. In the next cell to the right, we'll use SUMIF to add together all the numbers over 5. Type `=SUMIF(C2:C20;">5")` – the semicolon divides the cell range with the criteria for the sum, and remember to include quote marks around your criteria. In the next cell along, let's take this further. Typing `=SUMIF(D2:D20;">5";C2:C20)` adds an extra clause: it'll add together only the numbers in column D which lie next to numbers in column C which are greater than 5.

You can duplicate formulas (and automatically move their range of cells) in the same way that we earlier duplicated labels, by selecting their cell and dragging its bottom right corner. But for now, try replacing SUM and SUMIF with the following in subsequent columns: AVERAGE, MAX, PRODUCT, STDEV. And why not try creating functions using the totals we've already generated? The output of a function is just a value after all.

Simple science

Presentation is everything if you're handing off a spreadsheet to someone and you actually want them to understand it, so let's do a little sprucing up. We've filled our spreadsheet with a lot of long numbers, but there's a way of rounding them without losing the original granularity of the data. To prove it, select your original range of data (B2 to L20) and go to Format > Cells. Select 'Number' in the category box, and change the number of decimal places to '0' and hit OK. You'll see everything reduced to whole numbers visually, but the formula line below won't change to match because it's still working on the actual values stored within the cells. Clicking each cell will display its actual contents (rather than its formatted contents) in the input line at the top of the *Calc* window. You can also use cell formatting to turn numbers into currency values, force them to be text values (which can't be worked on with formulas) and apply fonts and borders.

A quick thing you should know before we sign off: you can create a new sheet within the same document by clicking the plus sign at the bottom left of the *Calc* window, and within it you can refer to cells from the first sheet by using the syntax <sheet name>.<cell>. So to sum the first column of the first sheet in a cell on the second, you'd use `=SUM(Sheet1.B2:Sheet1.B20)` – and you can, of course, use this technique to work with, and apply formulas to, numbers stored on multiple sheets. ■

Quick tip

Calc can, like *Writer*, save documents in a format compatible with Microsoft products. We'd recommend you choose .xls for maximum compatibility.

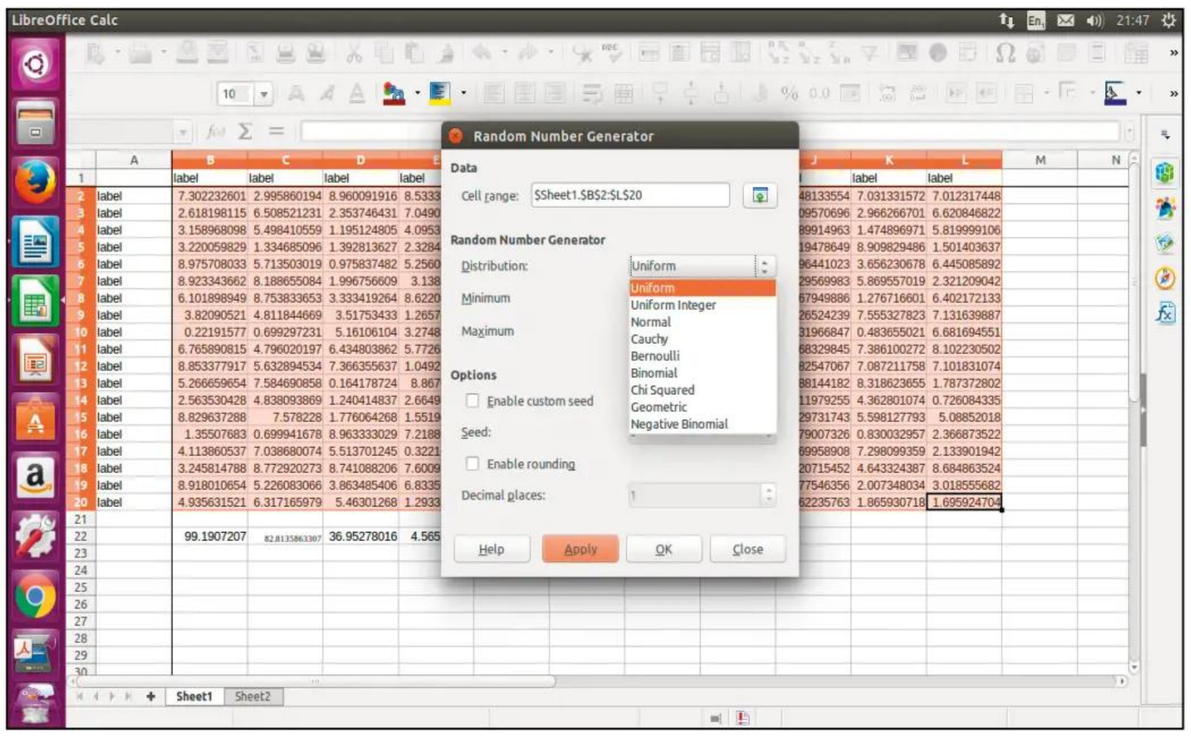

› **Random numbers can be inserted into cells, but you'll need to understand how, for example, Bernoulli numbers work to get them to work.**

PDFs: Edit and annotate

Portable they may be, but accessible? Not without a little hard work and determination.

Quick tip

LibreOffice Draw is capable of editing PDFs, but it's not always tip-top when it comes to saving them out. It's worth trying, if you have complex tasks to perform.

We've all seen PDFs, haven't we? They're neat little document packages which contain all the fonts and layout information required to display exactly what the creator intended, without requiring any of the software that was used to make them. They're also, usually, incredibly small. Hence the name Portable Document Format: universally supported on just about any system, almost always displayed correctly, and – at least these days – outputtable by just about any creative software.

But there's a problem: PDFs tend to be one-way streets, documents designed to exist in a read-only state. Editing a PDF is not an easy task. But, as we'll discover, there are ways around these restrictions; ways to get information out of PDFs, ways to tweak existing documents, and ways to create brand new PDFs from that source material.

Money matters

First, let's talk about the commercial options available to Linux users. There are two primary PDF editors on the market: *PDF Studio* and *Master PDF Editor*. The former comes in two flavours: Standard ($89), which can convert most documents to PDF, add annotations and comments to PDF documents, and split, merge or password protect your PDFs. For proper editing facilities you'll have to plump for the Pro version ($129) which also adds batch processing facilities and optical character recognition to extract text from images. There are free trials available of both versions at qoppa.com.

Slightly easier on the wallet at £36 is *Master PDF Editor* which, as its name implies, is pretty masterful at the task of editing PDF files, and also exporting them to various image formats, splitting and merging, document signing and much more. The reduced price, as far as we can tell, doesn't mean a great loss of functionality. There's a free evaluation version to download at code-industry.net so you can see if you like it before plunking down the cash.

There are other ways to perform the main functions of proper PDF editors without shelling out a penny, but they're all limited in their own way. So let's look at a few of the tasks you might want to perform, and how to get them done.

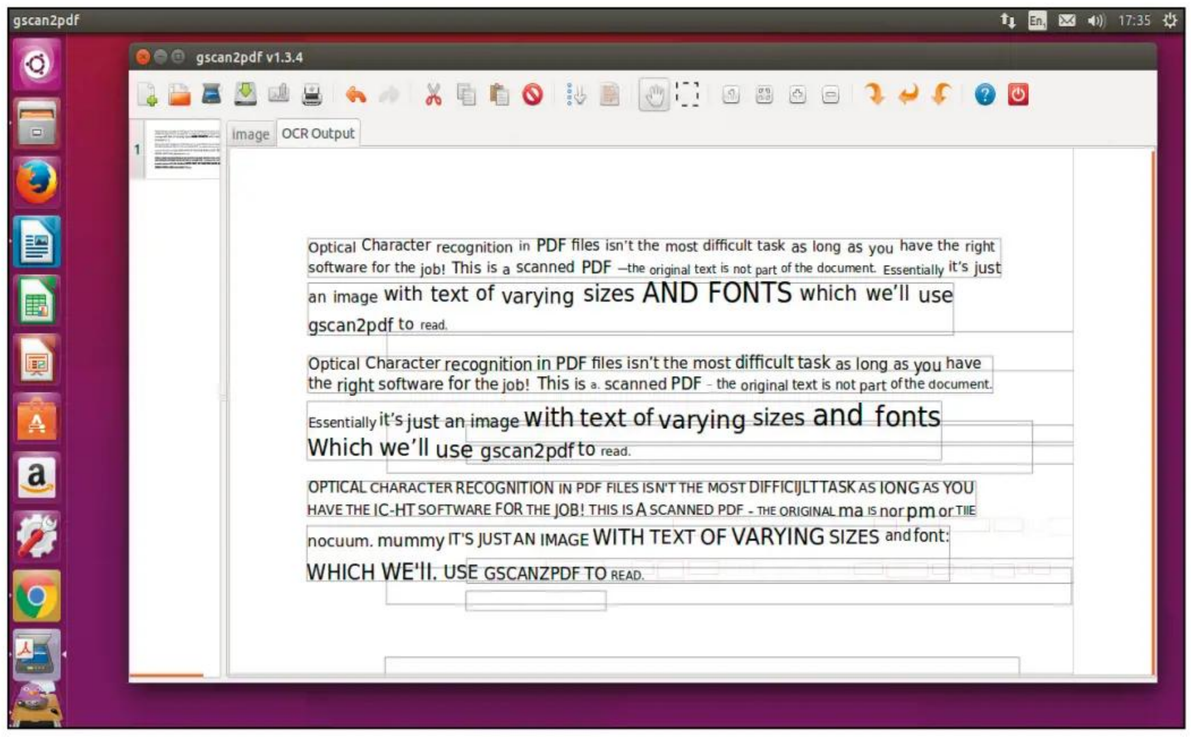

› Gscan2pdf, using the Tesseract OCR engine, does a pretty decent job of extracting text from images but struggles with more complex fonts.

Optical Character Recognition

OCR, as it's less tounge-twistingly known, is the process of getting words out of pictures. If someone's sent you a PDF of a scanned document, that's exactly what you're dealing with: a picture, just one in a nice package. In order to turn these printed characters back into editable text, your computer needs to, in base terms, trace over a bunch of pixels to work out what letters they might represent, then output that text.

We'd recommend a package called *gscan2pdf* for this task. It's primarily meant for turning images into PDFs – hey, another task ticked off the list – but it also has an OCR export feature built in. Install it from your package manager, run it, and hit [Ctrl]+[I] to import your PDF. If it's slightly off-kilter, you'll first want to use Tools > Clean Up to automatically adjust the text to be on a straight line, otherwise the OCR engine might get a bit confused. Make sure you only have a single edge checkbox selected, otherwise the deskew function won't work well. Click OK and your image should be adjusted.

Now let us run the actual OCR process. Go to Tools > OCR, and you'll see Tesseract listed as the OCR engine. This is also available on the command line, and it's a highly powerful tool, but rather difficult to run without some help – which is just what *gscan2pdf* is offering us. Check 'Threshold before OCR', and click 'Start OCR'. Click the OCR output tab, and you'll see the results. Trickier fonts might not read perfectly, so you can click problem words to correct them here if you spot any issues. Hit [Ctrl]+[S] to save, and open up your PDF in the viewer of your choice – you can now select the text with your mouse and copy and paste to your heart's desire.

Image editing

If you're pulling images out of PDFs, or you want to change their design, there's really only one sensible free choice in the the Linux domain: *Gimp*. That's not to say it's the only choice – there are plenty of image editors that can read PDFs – but *Gimp* is at the top of the heap. You likely already have it installed – it's in your package manager if not. Getting editing is just a case of opening the PDF using *Gimp*'s open command ([Ctrl]+[O]) and importing it as an image. You're then free to use Gimp's photo editing tools to manipulate it as if it were any other picture.

When you're finished, you'll need to use File > Export to save it out again as a PDF. This does have its downsides, though: *Gimp*'s standard import procedure flattens the PDF out, so you'll lose any text layers you might have previously had.

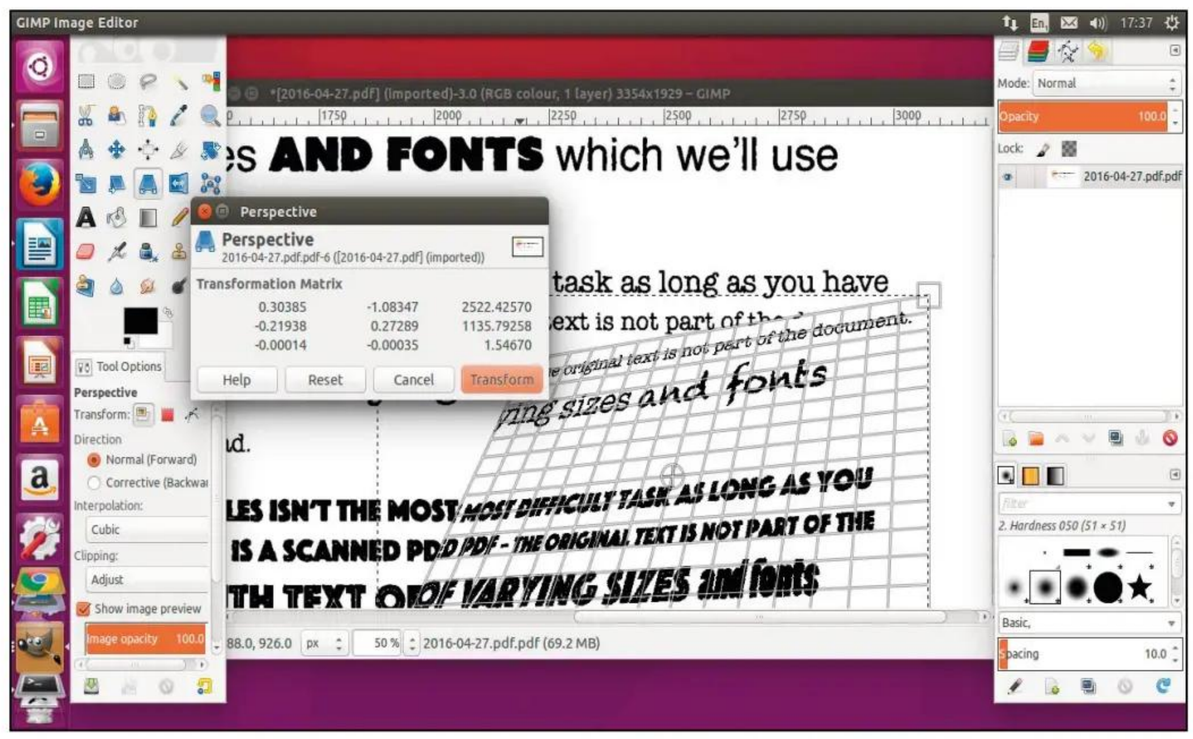

› Gimp can import PDFs and treat them just like image files – perfect if you want to extract the pictures from a scanned document.

Annotation

When you need to take notes on a PDF, you could do a lot worse than using *Xournal*. Ostensibly a note-taking application, it's actually the perfect alternative to printing out your PDF, scribbling on it with a Sharpie, then scanning it back in again – with a couple of other key benefits.

Start by downloading *Xournal* from your distro's package manager, then run it. Go to File > Annotate PDF, open your PDF, and start scribbling away with the left mouse button. If it doesn't work – we've had this issue on some machines – try toggling Options > Use Xinput to off.

Xournal offers a few handy tools. The pen, the default tool, lets you draw – or scrawl – pen marks on top of your document. The eraser, which you can automatically switch to and use by holding the right mouse button, scrubs out marks you've previously made. The highlighter is pretty useful, in that it works like a regular highlighter pen, smearing a bright colour while letting the text below shine through. When you're using it, try clicking the checkered colour box on the right hand side of the toolbar to alter its colour, and drag the opacity slider to select how bold you want this effect to be.

Next up is the text tool: just click anywhere on the document and type in your own notes. You can tweak the font with the rightmost element of the toolbar, and change any text you've already written by double clicking it with the text tool selected. You can insert images as notes with the next icon, toggle the following icon to automatically turn your crude scrawls into geometric shapes, and force straight line drawing by activating the ruler icon.

When you're finished, just head to File > Export To PDF and save your work as a new PDF file. Again, you'll lose any additional layers (like raw text, which will be tricky to OCR with your scribbles on top), but you can pass on your annotated PDF to colleagues, victims, or even friends. ■

Quick tip

You can convert many document formats to PDF using free online tools: we recommend trying **cloudconvert.com** first of all.

Inside Xournal

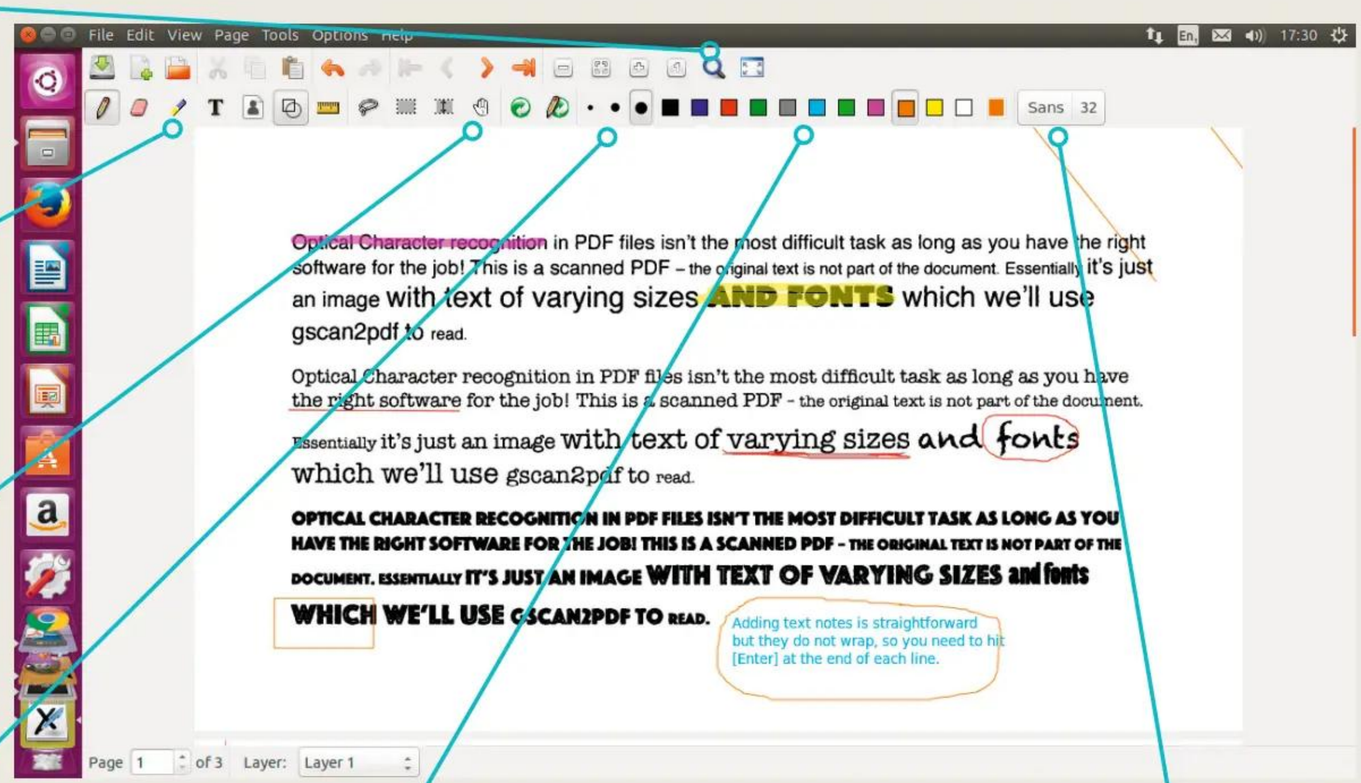

Zoom zoom

The magnifying glass is the only easy way to control your zoom level, particularly useful on high-resolution PDF files.

Drawing tools

Define your scribbles: will you use a pen? A highlighter? Infinitely neater text boxes? Pick your tools here and get annotating.

Move it around

The selection tools let you grab your scrawled notes and shift them around without affecting the background layer. Note that you'll need to.

Size it up

Xournal offers only a limited selection of tool sizes, so on high-res PDFs your notes may be a little small unless you use the biggest.

Colour theory

There are host of predefined and bold colours, perfect for highlighting, but click the rightmost icon to select an original colour of your own.

Face time

Use this button to select your text typeface – it won't affect notes you've already written, though.

Pidgin: Instant messaging

Connect with your friends in a network-neutral way and stay chatting to your heart's content.

Quick tip

By default *Pidgin* collects all of your chats as tabs in a single window. Just drag one of the tabs away to give it a window of its own.

Instant messaging has made a huge cultural impact over the past decade, but it's far from a modern concept. We can actually trace direct digital messaging back to the 1960s when multi-user systems at MIT contained their own internal messaging system. The first multi-user online client, as far as we can work out, was issued by Compuserve in 1980. Quaintly (and very much reflecting its place in time) they called it *CB Simulator* – and if you've ever used a CB radio, we're sure you can grasp the chaos and iniquity that ensued.

More recently, *ICQ* (released in 1996) really brought the concept of one-to-one or multi-user instant messaging to the attention of the public at a time when the internet was exploding in popularity, and it was been followed by a host of other IM clients: *AIM*, *Yahoo Messenger*, *MSN/ Windows Live Messenger* and so on.

Today instant messaging is almost ubiquitous and everywhere. There's likely an IM client in your email, one on your Facebook account, and one on your phone in the form of *Google Hangouts*, *Skype* or Apple's *iMessage*. There are a host more services running online. The problem is, they're almost universally incompatible with each other. If you really want to get connected with everyone, either you're going to have to be very convincing or you're going to have to support a whole host of different platforms. Good job we've got a solution, right?

› The core list of IM clients supported by Pidgin is long, but it can be expanded to support just about any platform.

Talk is sweet

First things first: a bit of bad news. Apple's *iMessage* is, sadly, out of reach on Linux unless you interface with a jailbroken Apple device using some rather roundabout methods. This is one of the cases where being convincing about changing IM systems comes in very handy.

The good news is that you can connect to just about every other major platform online, without needing to install a specific client. Facebook Messenger is on **facebook.com**, Google Hangouts can be accessed through Gmail or through Google+, Skype has its own web client, and so on. The fact that you're using Linux doesn't make this any different.

Of course, that's all a bit of a pain. Should you really have to keep a whole host of browser windows open on the off-chance that one of your chums messages you on that obscure service you signed up for a long time ago and never use but which they insist on mainlining for no good reason? No. You should not.

You can also get various instant messaging clients on Linux. A number of these are official releases. Skype, for example, has a well-seasoned Linux client despite being owned by Microsoft (the client existed before the MS purchase and has been updated since). Google offers an extension for *Google Chrome* that makes Hangouts work seamlessly, with desktop popup notifications, without forcing you to keep a particular browser window open.

There are also unofficial clients, and a stack of them, that purport to support various instant messaging protocols and generally do a reasonable job. Let's run down a few of the options! For Google Hangouts, install *yakyak* from your package manager. For WhatsApp, there's a client (about a year old) at **https://github.com/Aluxian/WhatsApp-Desktop/releases** which may or may not do the trick – download the appropriate .deb file for your system (32-bit or 64-bit) and install it through your package manager. For Facebook, use the deceptively-unaffiliated *Messenger for Desktop* at **www.messengerfordesktop.com**, from the same developer as the WhatsApp client.

Good words

As you've probably realised by now, we're not fans of all these individual platforms cluttering up our methods of communication. Think about your mobile phone – where would you be if all of the different mobile operators didn't allow communication between each other? You'd probably be at the phone shop demanding a handset which supported all

Peer inside Pidgin

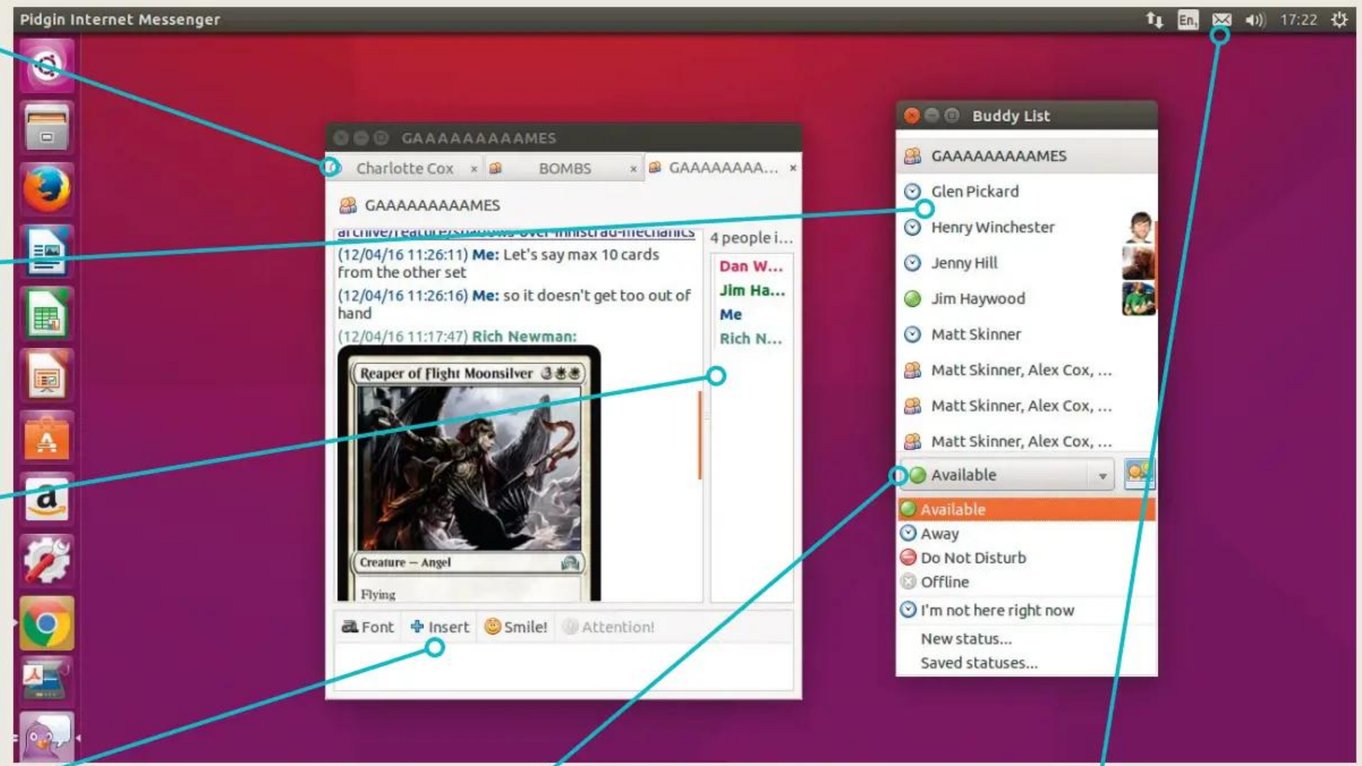

Talking tabs
Each of your chats will seamlessly sit in its own tab – chats with new messages will turn red so you can easily catch up.

Buddy up
The buddy list contains all of your friends and group chats. Try right-clicking a friend and selecting Get Info to find out more about them.

Who's in?
In a group chat or IRC chat, you'll see a sidebar listing everyone who's in the room at the current time. Right click their names to get more information.

Special delivery
The insert button will let you post items – usually pictures, but if your chat client supports it you may be able to add other files here too.

Status report
Click the drop down box in the buddy list to select your chat status – handy if you want some peace and quiet from contstant chatter!

Notify me
On Ubuntu, the envelope icon will turn blue when you have a new message, and you can interact with *Pidgin* here too.

the networks. And you'd be right to do so, which is where *Pidgin* comes in.

Pidgin is a multi-platform client which you can use to stay connected to all of your instant messaging accounts (well, all beyond the few it doesn't support) at one time. No need to run scads of desktop clients or keep multiple browser tabs open – one app, one place, many accounts.

First you'll need to install it – it should be in your package manager but if it isn't, download it from www.pidgin.im. On your first run it'll prompt you to add an IM account. We'll add a Google Chat account just to demonstrate. Click Add, use the dropdown box next to Protocol to select Google Talk, and add your Google account credentials in the username and password boxes. If you use Google's two-factor authentication (you should) then head to **https://security.google.com/settings/security/apppasswords** to generate a new password for *Pidgin*, and input this in the password box.

On the advanced tab, change 'connect server' to talk.google.com, and make sure Pidgin is set to require encryption in the drop down at the top of the window. Hit save, and click the check box to enable the account, and you should see it connect and display your buddy list. You can now add other accounts.

Radio chatter

There's a problem, though. Google doesn't use the XMPP protocol to handle its Hangouts group chats, so signing in using the default Google Talk protocol will net you private chats only. And *Pidgin*, by default, doesn't support Hangouts specifically. But there's hope: *Pidgin* includes a plugin system that allows you to extend its functionality, and there's a plugin that adds Hangouts support.

Head over to **https://bitbucket.org/EionRobb/purple-hangouts/** to find out how to install it. Basically it's just a case of copying and pasting a few lines of text into a terminal window: the first gets hold of the dependencies required, the second gets hold of the code for the plugin, and the third builds the code into something your OS can understand.

Once it's all run, close *Pidgin* (right-click its icon and select 'quit') then run it again. You'll need to re-add your account, but the method is slightly different this time. First, hit [Ctrl]+[A] to bring up the Manage Accounts screen, click on the Google account we'd previously added, and then hit Delete. Don't worry – this only deletes *Pidgin*'s reference to it, not the account itself. Click Add, pull down the protocol list, and you'll see a new entry for Hangouts. Enter the username of your Google account, the name you want to refer to yourself by in the alias box (otherwise you'll just see a long number) and hit Add. *Pidgin* will then open a browser window and a prompt for an OAuth code. In the browser window, sign in to your Google account, click 'Sign In' on the page that follows, and copy and paste the piece of code you're presented into the *Pidgin* dialog box. Hit OK and you'll have access to Hangouts.

There are many more *Pidgin* plugins that can extend its functionality. Try switching on a few of the built-in selection by going to Tools/Plug-ins. Happy chatting! ■

Quick tip

Pidgin also works as an IRC client, so you are able to keep in touch with real chat rooms as well as instant messages.

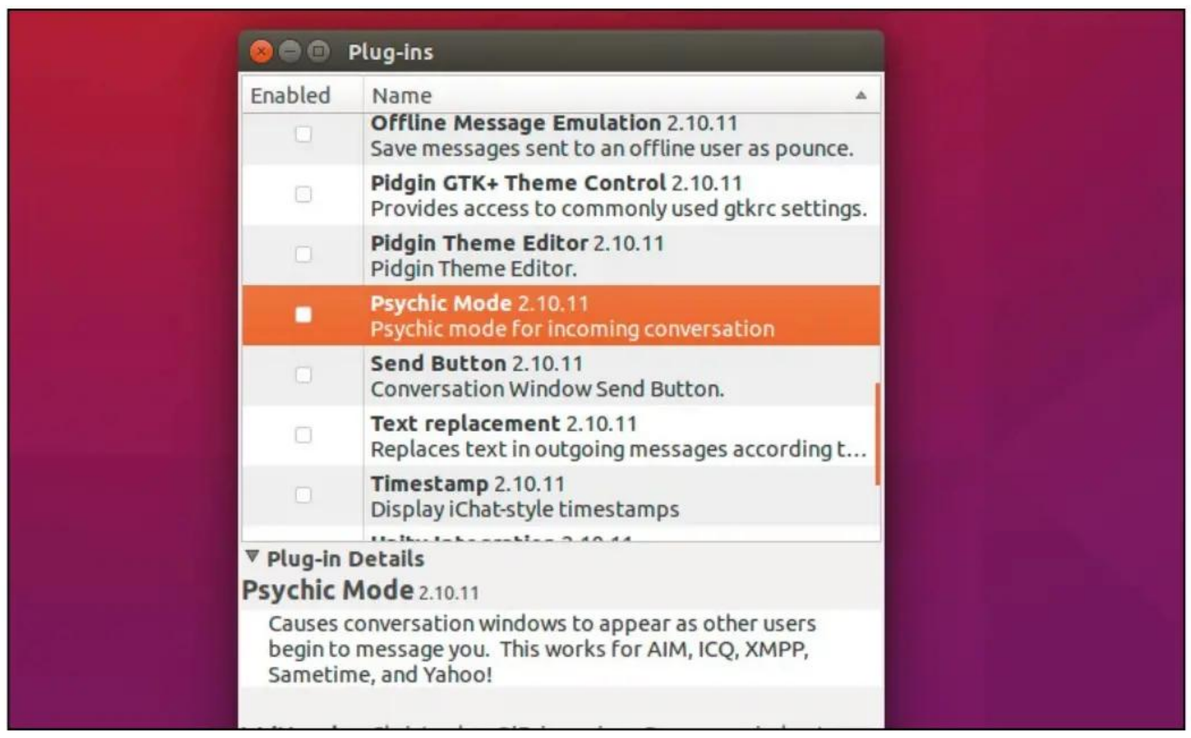

› Pidgin's plug-ins list contains a number of ways to customise instant messaging to your liking. We're fond of psychic mode. Spooky!

Thunderbird: Email made easy

Set up your very own desktop email client and keep your messages and contacts close at hand.

Quick tip

Thunderbird supports multiple email accounts at once – just go to File > New > Existing Mail Account in order to add another.

Email is one of the internet's oldest technologies. It outdates the web, it outdates instant messaging and it even outdates the internet's predecessor, Arpanet. MIT, so often the innovator in these fields, started using software called *Mailbox* and *SndMsg* to leave messages for each other in 1965. The @ symbol was introduced to route mail to specific computers during the development of Arpanet in the early 70s; by 1974 the US military was a heavy user, and by 1978 75% of all Arpanet traffic was email. You could even reasonably give credit to email for inspiring the earliest development of the internet: this was a closed-system technology that the public couldn't wait to use.

We'd stake money on you having at least one email address, and if you've been around for a while you probably have several to manage. While most email accounts now have some kind of web interface for accessing your inbox and sending new messages, it's arguable that this may not be the best way to manage your mail. Say you lose your internet connection: your inbox, stored entirely online, will be completely inaccessible. So many of us rely on old messages to remind us of information, this probably isn't acceptable.

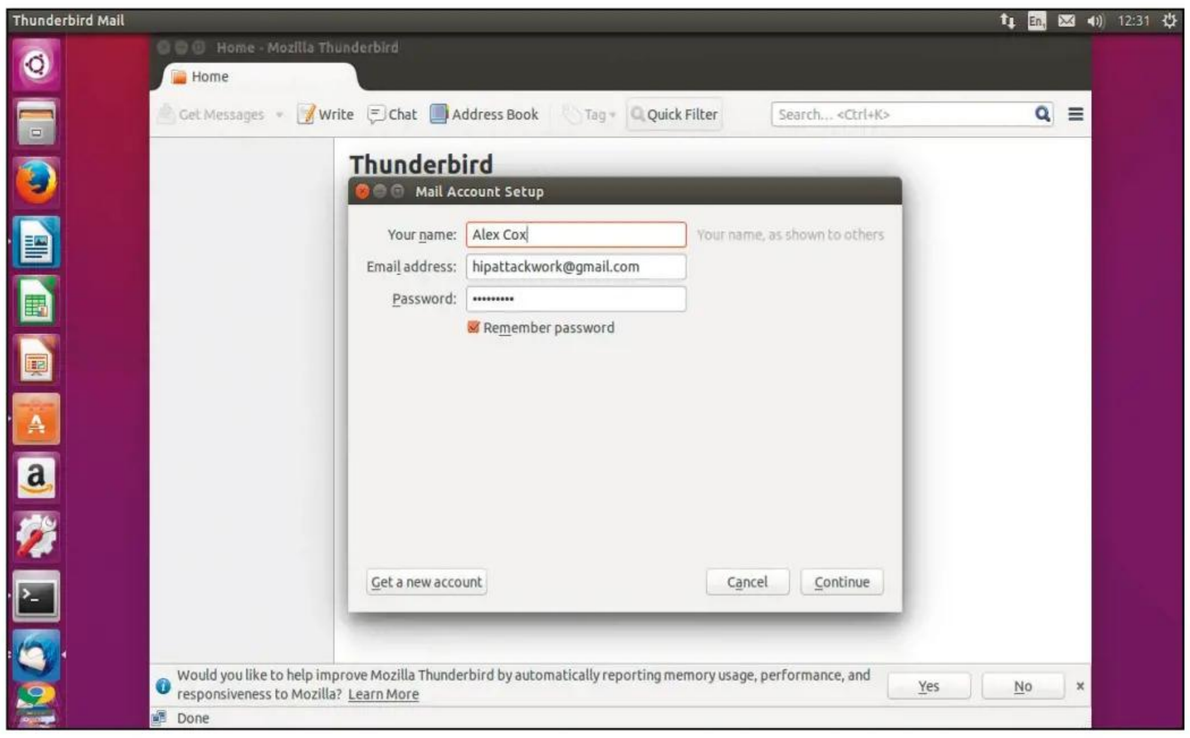

› Adding an account, if you're with a recognised provider, could not be easier – just add your email address and password.

Security, too, should be a concern. How well do you trust your webmail provider not to snoop on your messages? What happens to your email archive if that provider goes bust, or a hacker forces their way in and changes your password? These, and many other reasons, are why it makes sense to use a desktop email client that can be backed up, store your messages offline, and live – if you have such a facility set up – in an encrypted partition away from prying eyes.

Positive Outlook

If you've been around computers for a while, it's probable that you've used Microsoft's *Outlook* or *Outlook Express*, packages which do just the job we're looking for: downloading mail from servers, storing it, sending new mails and replies, and backing up your inbox. Funnily enough, Microsoft's flagship mail package isn't available on Linux platforms, but Mozilla's well-weathered *Thunderbird* is just as capable and should be quite familiar to former Outlook users. We'll take you through installing it, configuring it to connect to an account, and beyond.

Let's start by getting hold of it. Many distributions will have it installed by default – including Ubuntu, which we're using here – but if you don't have it head to your local friendly package manager and search for *Thunderbird Mail*. Install as usual (see p46 for more information on package managers) then run it from its icon.

Postmaster

When you run *Thunderbird* for the first time, you'll be given the option to sign up for a new free email address – we're going to skip this particular step for this tutorial, so click Skip. Next, fill in your current email details. Your email address and password are the important parts here, the 'Your name' field doesn't have to match the name you usually use for the account, but it's useful to use your real name here so recipients know who's emailing.

Brilliantly, *Thunderbird* now queries a database of server information. If you're using a recognised provider like **gmail.com** or **outlook.com**, it'll automatically grab the correct ingoing and outgoing server details so you don't have to worry about it. More obscure providers may not be in the

POP3 vs IMAP

There are two ways to receive email: POP3 (Post Office Protocol) and IMAP (Internet Message Access Protocol). POP downloads messages to your computer and (usually) deletes the originals from the server. It's the classic method, and good for those using a single machine or with a security-conscious bent: if the mails are gone from the server, it's going to be a lot more difficult for them to be traced or spied upon. There are big disadvantages when it comes to using POP3 on multiple machines – since it's a one-way transaction, there's no synchronisation between mailboxes, so you won't get an indication of the messages you've read or replied to, and your filing systems won't match up.

IMAP is a a two-way protocol, and probably the better choice if your provider offers it. Folders, read status and message deletions will all match up between the server and your mail client, meaning you can use the same account on multiple machines – and indeed use webmail and *Thunderbird* at the same time while keeping things synchronised.

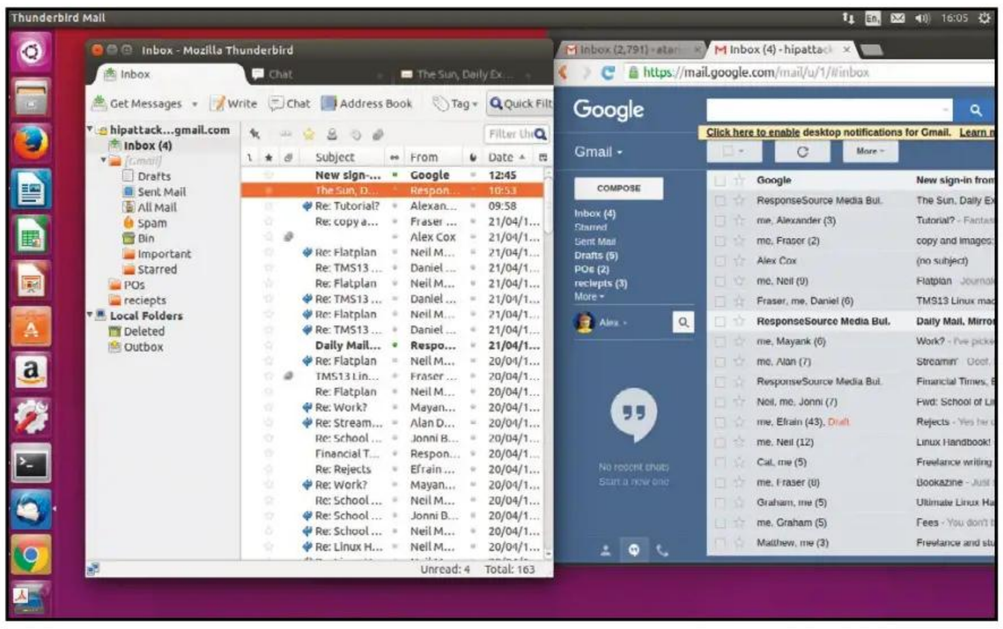

database, however, so you may have to fill in the blanks yourself. Your ISP or mailbox provider will provide the appropriate information; we can't tell you precisely what to put here, because very ISP is different. You can now choose whether to use POP3 or IMAP with your account – see the box above for more details on each – and hit Done. You may have to authorise the use of your email account in a further window, but you should only have to do this once.

A quick note for Gmail users: if you have Google's two-step authentication switched on (and you should do), you'll need to generate an app-specific password and use that, rather than your regular account password, to log in to your email through *Thunderbird*.

Split personality

Thunderbird's default main screen lists your email account and its subfolders on the left, and has a main window on the right split into two: the top pane contains the contents of your inbox, and the bottom pane contains the currently selected email. Try clicking on your inbox in the left column. *Thunderbird* will download all of your email from the server (this could take a while) and, if you've used IMAP, you should retain the read/unread status of your messages. By default, the inbox view displays messages in date order with the oldest at the top, but you can quickly re-sort your messages by clicking the column headers at the top.

Click on a message and you'll see it appear in the bottom pane of the main window. If this is a bit small, just hover your mouse over the bar between the panes and drag up to give yourself a bit more space. Still not enough room? Double-click the message and it'll open up in a tab of its own, giving you the full height and width of your *Thunderbird* window to play with. Click on the Inbox tab (or close the message tab) to get back where you were. Further layout options can be found in the View menu at the top of the screen.

In the book

While you have a message open, you'll see all the usual email functions (reply, reply all, forward etc.) lined up along the top of the window, and you can always hit the Write button to fire up a new mail. If an email includes embedded images linked from the internet, you'll also, at least initially, see a bar with a Preferences button. *Thunderbird* is restricting these images because, unlike plain old email, senders can use images to glean information about your computer, and whether you've read the message. Hit Preferences and you can choose to view images in just that email, for all messages from that address or domain, or head to the main Remote Content Preferences section where you can switch off this feature altogether.

When you're creating a new message, you'll probably discover that *Thunderbird* doesn't actually have an address book to start with, even if you've been using one with your webmail account. Manually importing your address book (Tools > Import) is the only way to get this done. The process is going to differ depending on your source, but it's not incredibly difficult.

Different versions of *Outlook* tuck away the export options in slightly different menus, but you should find them buried somewhere within the File menu. Make sure you export your address book as Comma Separated Values or .csv, as *Thunderbird* is happy to accept this format. In Gmail, select the contacts section from the top left of the main window, (just above the Compose button), click More then Export to find the option. ■

Quick tip

You can set up a signature to add to the bottom of all of your outgoing messages in the Edit > Account Settings menu.

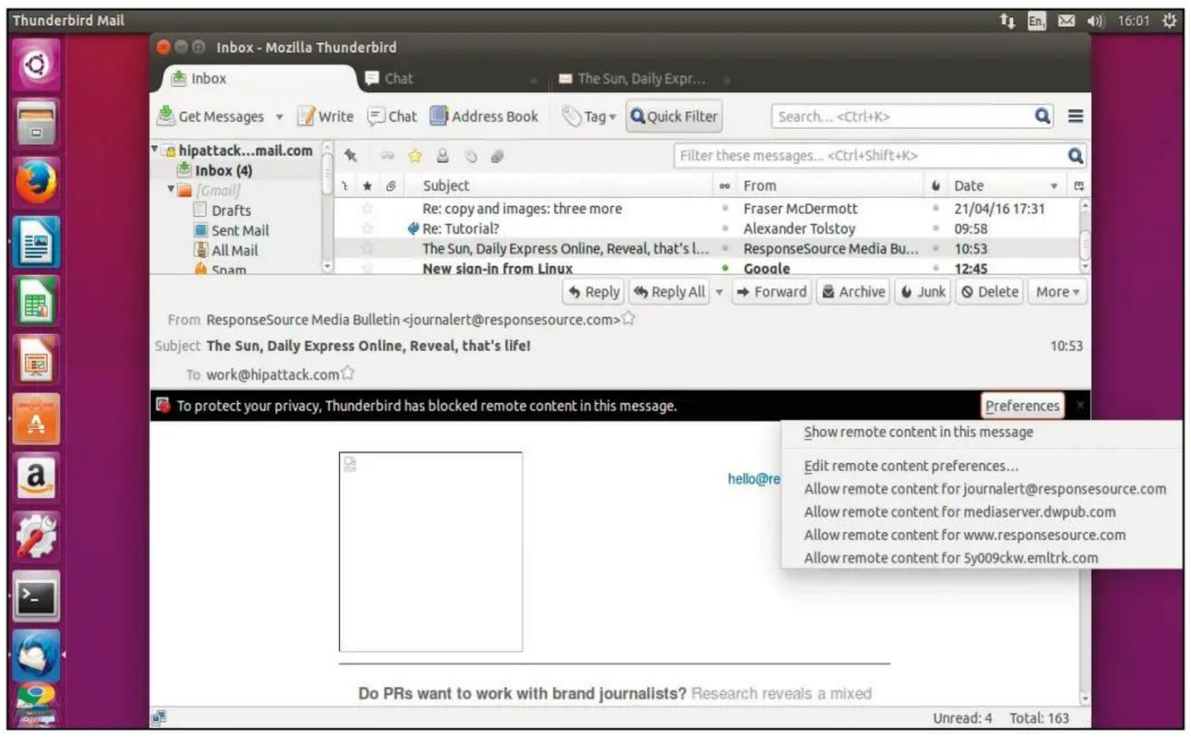

› Use the preferences button to determine which images you can see in emails you're sent – you can get quite specific.

Text editors: Word working

Wrangling the raw meat of words doesn't have to be torture. Here's a guide to the best editors around.

Quick tip

Ubuntu includes gedit by default, but it's been renamed 'text editor' in the applications list. You can still find it by searching for 'gedit'.

Word processors have their uses. You can certainly crack out *LibreOffice Writer* and its ilk if you want to construct a beautifully formatted document, but when it comes to constructing unformatted plain text – be it coding for the web, assembling traditional programs, or just writing in an environment unsullied by font frippery – traditional text editors come into their own. With a lot of the fuss and unnecessary elements removed, they're usually able to run with a much smaller memory footprint and therefore load faster, take up fewer system resources, and help you focus on the task at hand.

If you've used Windows, compare the likes of *Microsoft Word* with *Notepad*. The latter doesn't have many advanced features, but it certainly has its uses. On macOS, *TextEdit* fills the lightweight text editor role.

Many Linux text editors don't even need to be run in a graphical environment, meaning you can get your writing done on literally any machine. The usefulness of this aspect can't be understated as you progress in your Linux adventure. Connect to any remote machine in text mode and you can edit configuration files or write scripts straight from the command line. You can, of course, run these without dropping to text mode by using a program like Terminal to bring the command line into your GUI.

So let's break down a few of your options, see what can be done with them, and contrast and compare their features.

It should be noted that most of these editors include syntax highlighting – where the functions and variables of programming code are given different colours for ease of navigation – but it's not incredibly reliable on console-based software because different terminals have different colour display capabilities. If it's a feature you really need, opt for a GUI editor, or run the console apps through a program like *tmux*.

On the command line

We'll begin with the top-line text mode browsers, since many of these provided a blueprint for graphical editors that came later on.

emacs While many well-worn programmers will swear by this editor, we wouldn't recommend it to newcomers. Getting used to its somewhat unusual control scheme and outrageously large list of features takes some serious effort. Most of *emacs*' commands are hidden behind combinations of modifier keys ([Shift], [Ctrl], [Alt]) so you'll definitely need a cheat sheet.

vim Did we just say *emacs* took a lot of effort? Well *vim* is even more obtuse – but for a reason. It breaks its interface into several modes – insert, command and last-line to name three – which determine what's possible. When you're in command mode, for example, typing on the keyboard won't actually edit the document; instead, you type specific codes

Useful nano commands

[Ctrl]+[6] – Mark – Begins selecting text at the current cursor location. You can also select with the mouse.

[Ctrl]+[K] – Cut – Removes the selected text and puts it in the cut buffer.

[Alt]+[6] – Copy – Puts current selection in the cut buffer without removing it.

[Ctrl]+[U] – Uncut – Otherwise known as Paste. Inserts contents of the cutbuffer at cursor location.

[Ctrl]+[O] – Write Out – Saves the file. If a mark is set, saves the currently selected text as a file of its own.

[Ctrl]+[W] – Where Is – Search, essentially. Type the string you want to find and hit [Return].

[Ctrl]+[\] – Replace – Type the string that you want to replace, then its replacement.

[Ctrl]+[T] – To Spell – Sends your document to *spell* for spellchecking, if it's installed.

[Ctrl]+[J] – Justify – Tames unruly documents which have exceeded the screen width. Follow with [Ctrl]+[U] to undo this action. Note that you have to undo this immediately: if you perform any other actions, you'll have to manually unjustify the text.

[Ctrl]+[C] – Cursor position – Tells you where your cursor is. Also gives a handy line and character count.

[Ctrl]+[H] – Help – Accesses *nano*'s built in help system, which is more detailed than its man page.

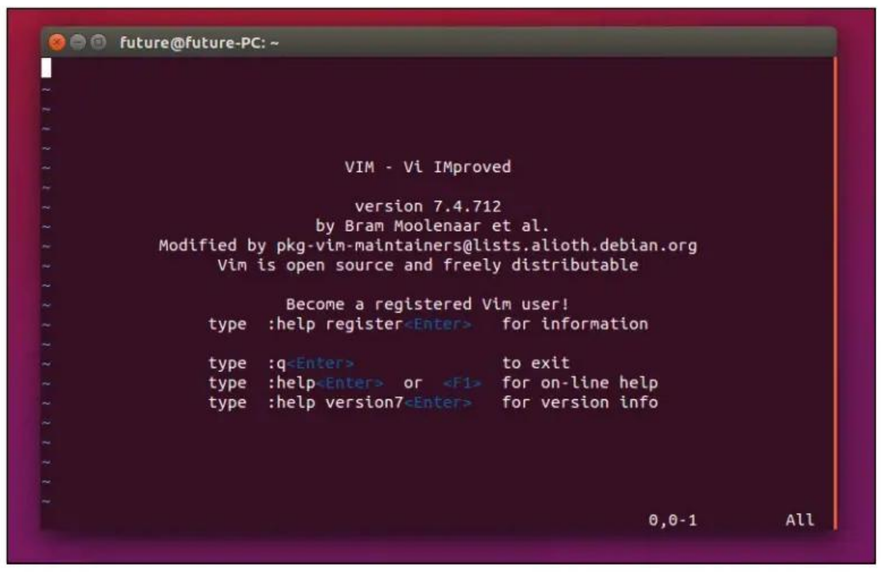

› **Vim can be tricky to get to grips with, but it does at least include detailed help on board.**

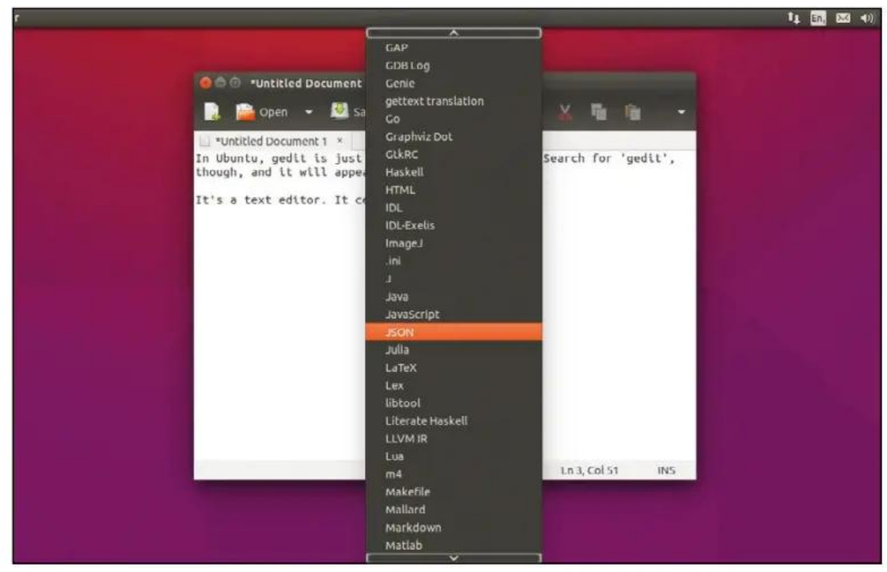

› **Gedit's syntax highlighting supports a comprehensive list of programming languages and document types.**

to do things like saving. And what's the code to save? It's `:w`. Which makes a lot more sense if you think in slightly different terminology: that's 'w' for 'write'.

nano Definitely a great choice both for beginners and those who don't want too much fuss, *nano* – and the non-GNU package it cloned, called *pico* – is a straightforward console text editor that retains the simplicity of most GUI editors. Don't be fooled into thinking there isn't complexity in there, because there's plenty, but much of it is hidden under the surface. We'll look at *nano* in more detail later on.

jed It's console-based. It works on plain text. But as long as your console supports mouse input (and even if it doesn't) *jed* is about as close as you could get to a GUI text editor, with proper menus and everything.

On the desktop

As you might expect, there's a little more deviation from the pure text-editor form on the desktop, but we'll look at some of the prime candidates and define Windows' *Wordpad* as our upper limit: feature rich, but still primarily about the words.

gedit This is the text editor that comes with most Gnome desktops, and probably our favourite. It expands on the basic text editor structure with syntax highlighting, full UTF-8 support, automatic indentation, a proper clipboard and more. *gedit* proves that if you have the choice to use a desktop editor, you probably should.

gvim The bonkers nature and terminology of *vim*, transported kicking and screaming onto the desktop. While it's a bit easier to navigate thanks to its graphical user interface, you may still need some vim skills to get the most out of this hardcore editor.

Sublime It's commercial software, and the thought of paying for a text editor might put you off. But *Sublime* is well respected, cross-platform, and for serious coders it's the absolute cream of the text editing crop. Take just one case: editing a variable name. In other editors you could use the standard find and replace function, but that may tweak other text in your code; *Sublime*'s intelligent handling of syntax makes renaming variables throughout your code easy and super-quick.

Beginning nano

Let's take a look at *nano* to see how easy it is to use. First you'll need to get it running, and it's important to do this as a super user, otherwise it's unlikely to have the relevant privileges. Start by opening a terminal window, and type:

```
sudo nano test.txt
```

You'll probably be asked for your administrator password after this. The command essentially says 'open nano (the second clause) as a super user (first clause) and edit the file test.txt (third clause)'. This doesn't create the file – it's not like create a new text document then opening it up – it just tells nano what you want to call it and, if you like, where you want to store it.

The *nano* interface appears, with your cursor in the main window. Type something in, then hit [Ctrl]+[X] (Ctrl is indicated by a caret ^ symbol in the help text at the bottom of the screen) to exit *nano*. You'll be asked if you want to save the modified buffer; essentially nano is asking if you want to write what's inside its buffer (the text in its interface) to the file. Hit [Y], hit enter to confirm the filename, and *nano* will exit.

Back at the command line, type ls and you will see your new file listed. Open it back up with the same command that we used earlier on – try hitting [Up] on your keyboard to move back through your previous commands to save you typing it again.

Now let's try a little cut and paste. Enter the phrase 'the quick brown fox jumps over the lazy dog' (good excuse to check your keyboard's working, eh?) and move your cursor to the start of 'jumps' with the arrow keys. Use [Ctrl]+[6] to mark the start of your cut, highlight the rest of the word and hit [Ctrl]+[K] to put it in your cut buffer. Now move to the end of the line and hit [Ctrl]+[U] to paste the contents of the cut buffer at the cursor position. ■

Quick tip

If you are using a console text editor in a terminal window, you will find that most of them offer mouse support for selecting characters and such.

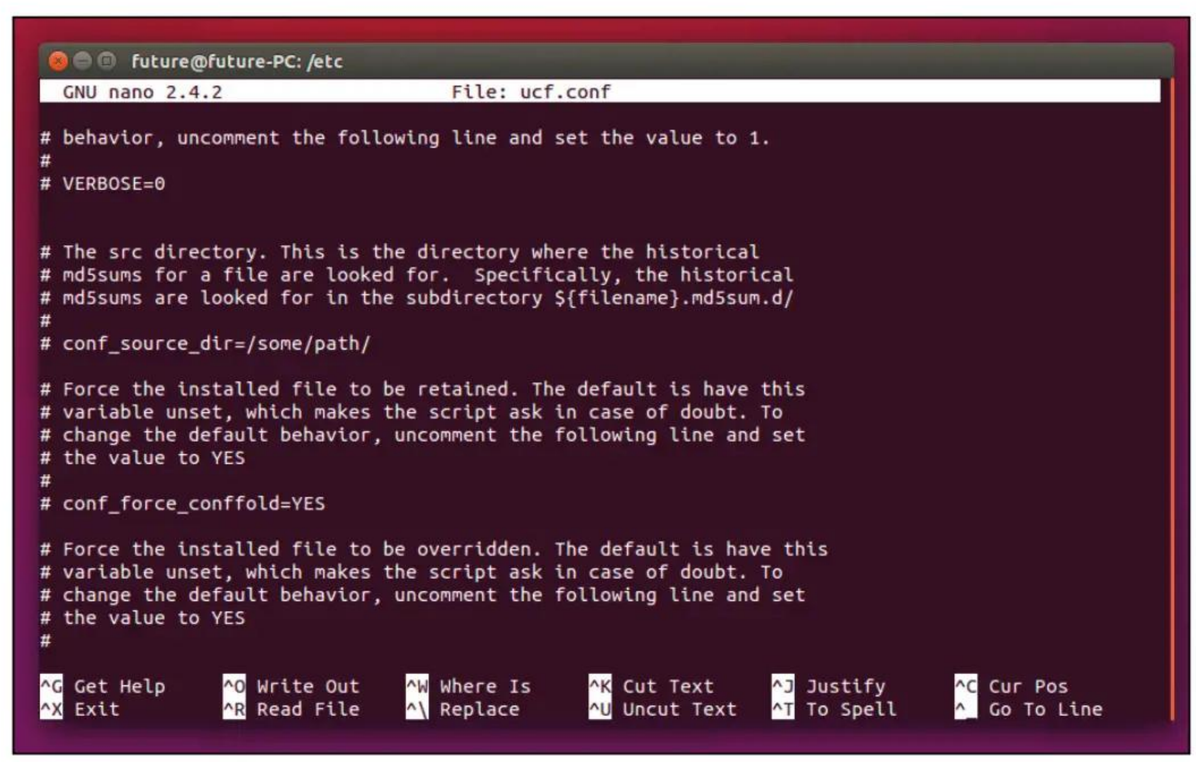

› **Nano sure ain't pretty, but if you only need to edit a file, it's the quickest and easiest option.**

Fun

Playing music & video 62
Audio editing 64
Photo organisation 66
Photo editing 68
Video editing 70
Setting up a media centre 72
Playing games 74

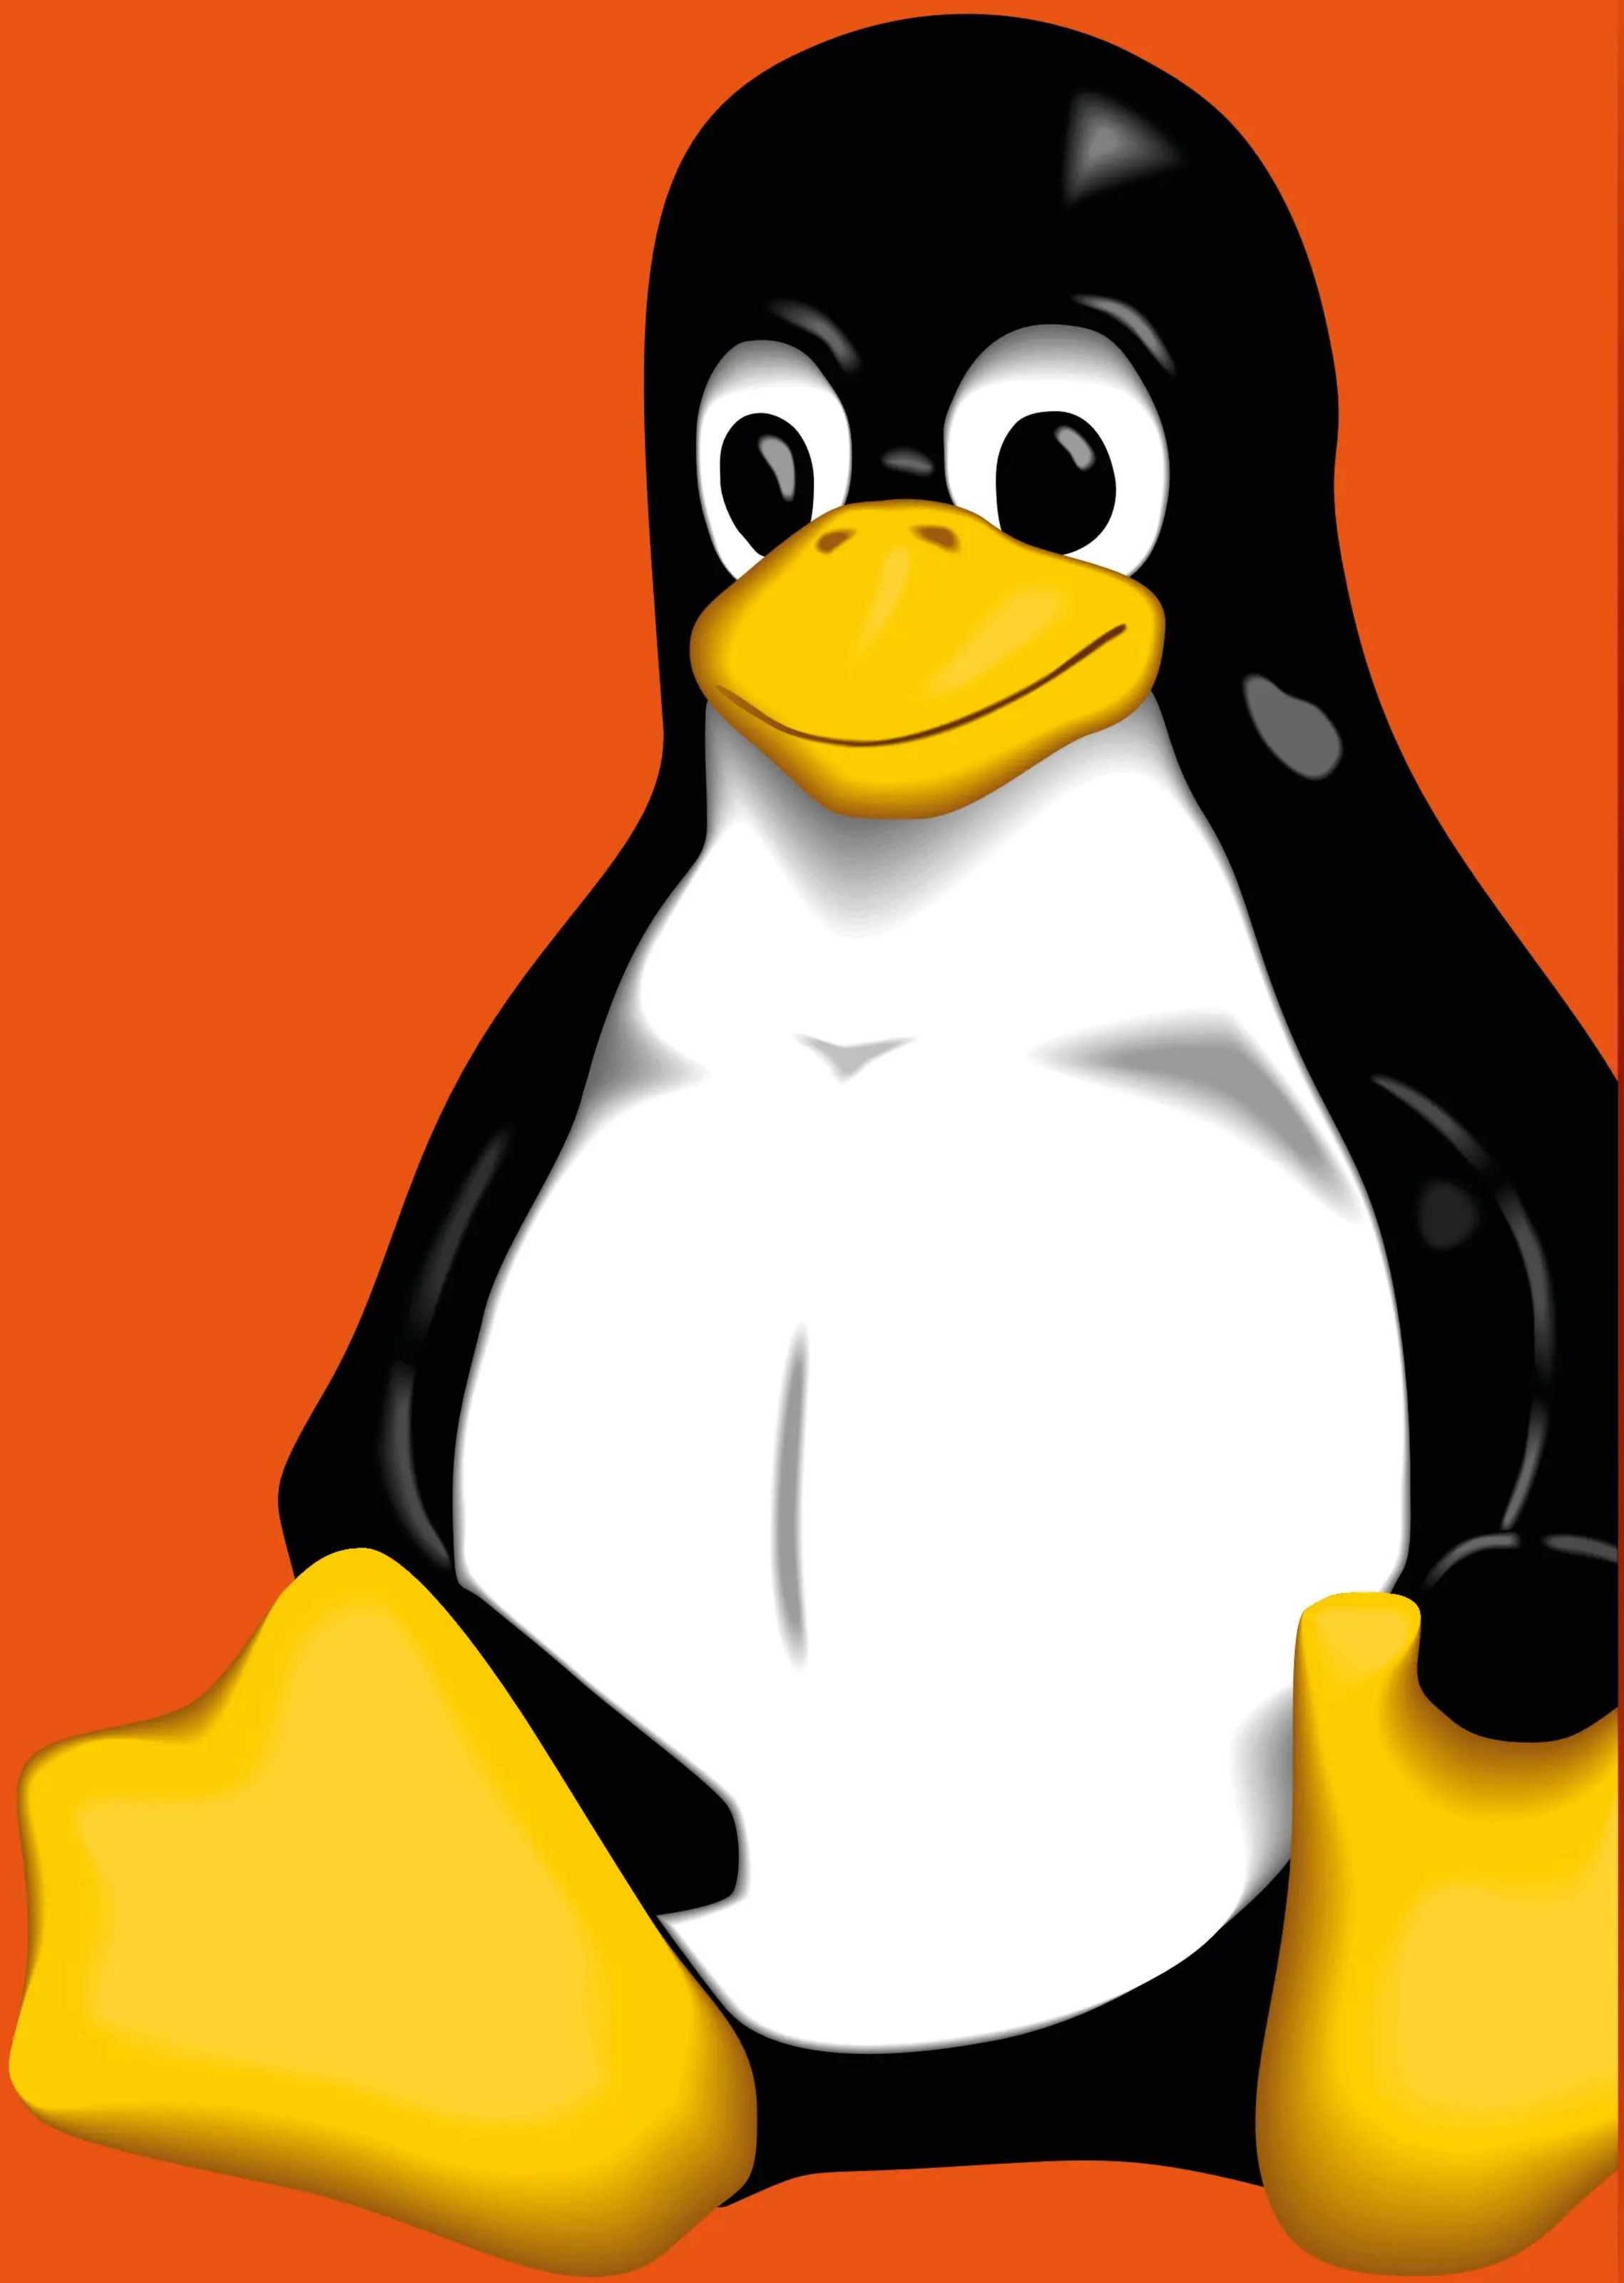

Rhythmbox: Playing media

Get your music collection blasting from your speakers and super-organised, then watch some videos too.

Quick tip

When it comes to online music, there's a desktop client for Spotify that can be added to Ubuntu and Google Play Music works through the browser. Netflix can be accessed if you install the Chrome browser, or use the User Agent Override add-on in Firefox.

Remember the days when we used to buy CD players with built-in radios, or cassette tapes for listening to our music? No? You're not alone. When it's possible to carry around a lifetime's worth of tunes on your phone, and stream audio wirelessly to your TV or HiFi system wirelessly, why on Earth would you invest in something so archaic?

We'll wager everyone needs a good media player on their desktop, however. Whether it's to enjoy music or internet radio while you work, or watch a movie on your laptop while travelling by train, digital media should be enjoyable whenever, wherever and however you need it.

And fortunately, Ubuntu has all the tools you need for playback built-in. If you're coming from Windows, there is one key adjustment you'll have to make. Microsoft's operating system comes with the all-encompassing *Windows Media Player* built-in, which is a hub for both music and video played back from a local source – like a hard drive or USB key – or a networked drive. For entertainment purposes, Ubuntu has separate applications for video and music playback.

For music, that application is *Rhythmbox*, which can manage your audio library, stream from remote services or a household server, and rip CDs directly to your hard drive. For video, the *Totem* player – which appears in the Unity launcher under the name Videos – serves similar purposes.

Living in a Rhythmbox

As far as audio goes, *Rhythmbox* is as fully featured as you need. It can be used to playback music from a local library, rip CDs, add album art to digital files automatically, purchase songs online and stream from another PC or network attached storage.

› **Probably the best Ubuntu video player, VLC is a must-have install.**

Before you start using *Rhythmbox*, however, there's one important thing you'll need to bear in mind. By default, Ubuntu doesn't include the software codecs required to playback the common MP3 format which many people use to store songs diigtally. Because many parts of the MP3 codec are still protected by patents, Linux distributions are unable to include it as standard without paying a licence fee.

Fortunately, the way around this is simple. Open up the *Ubuntu Software Centre* and search for "Ubuntu Restricted Extras". If it's not installed yet, click to download and there's little your PC will left unable to play.

Back to *Rhythmbox*, and the first time you open it up you'll want to point it at the place where you keep your song files stored in order to build up its library list. Open up the Edit menu and click Preferences, then Music. Click the Browse button to direct *Rhythmbox* to your Music library folder if it's not the default one, and then put a tick in the box "Watch my library for new files". Now any songs and albums in your library will be added to the *Rhythmbox* home screen, and any time you rip a CD, purchase music through *Rhythmbox* or add music from a portable player it'll be saved to the same physical place.

You'll also see options in this screen to select the format and quality for ripping music, and also for the naming convention you prefer in filenames that are automatically generated.

Once you've loaded some music up, the main view for *Rhythmbox* is fairly straightforward (see boxout above). You'll notice that the column on far left is pretty sparse initially, with some basic library options for selecting music, podcasts (see boxout) or internet radio stations. You can also see what's coming up in the play queue.

To start adding extra options here, you'll need to visit the Plugins menu, which you'll find if you click on Tools in the menu bar.

Rhythmbox comes with extra features which are disabled by default. In this menu you can turn extras such as Last FM scrobbling and the Jamendo music store. The selection of plugins is fairly limited initially, but you can add more with ease.

The one essential plugin to install is the *Grilo* suite, which will enable you to stream from other computers and libraries on your network. If *Grilo* isn't already installed, open up the *Software Centre* and search for grilo-plugins, then install the Base package from there. Enabling this in the *Rhythmbox* Plugins section should result in any shared libraries on your network popping up in the left-hand menu.

Exploring the Rhythmbox interface

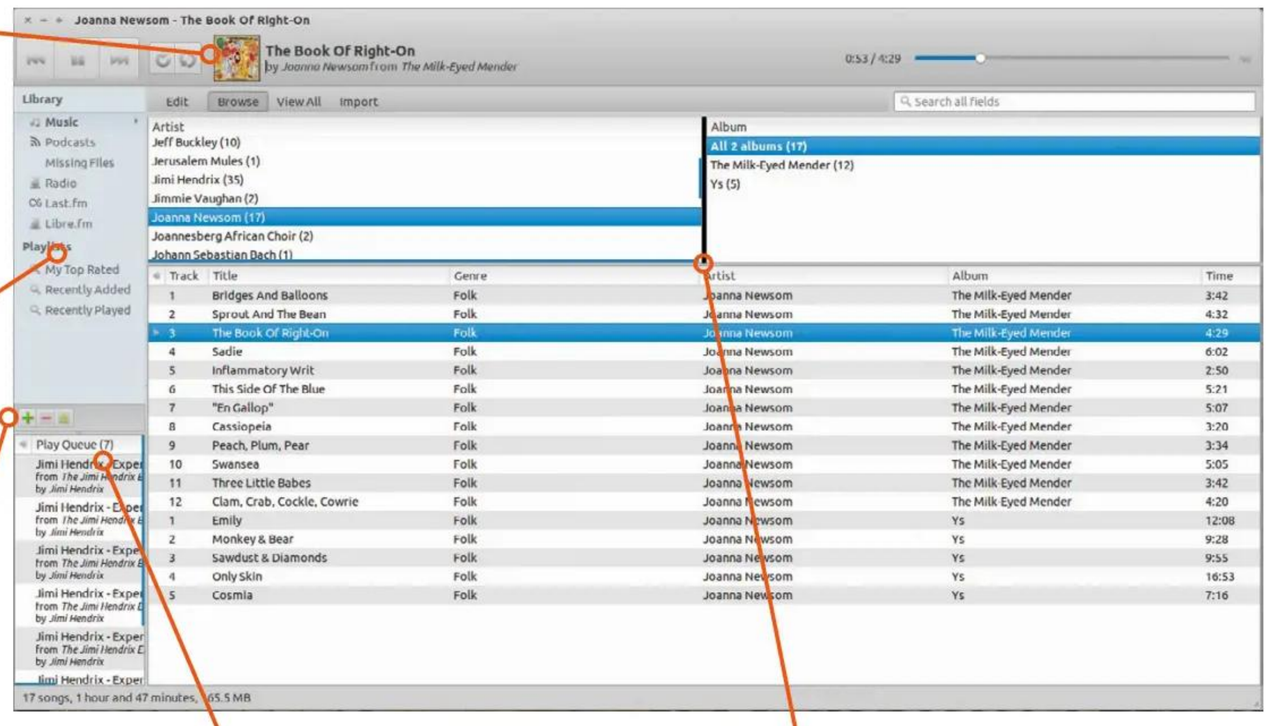

Album art
Rhythmbox will automatically pull album art from the net, but you might have to set-up the plug-in first.

Library
From here you can select local libraries, UPNP shares, internet radio or podcasts.

Add a source
The plus and minus buttons let you add a source for you library. This can be a folder on your hard drive, a URL or a shared drive.

Play queue
This list tells you what's up and coming on your playlist.

Artist list
On the left is the artist list, on the right their albums and in the main window a track listing.

Totem-ic

While *Rhythmbox* is reasonably well featured, *Totem* – the video player – is less so. It's fine for playing back videos selected directly from the file browser, but as a library manager it doesn't really cut it. You can add videos one by one, but there's no option to import a whole folder of videos or to automatically detect UPnP servers on your network.

Totem does have one very useful feature though. If you open up it up by selecting Videos from the Unity launcher and click the Channels option at the top of the screen, you'll get a selection of web videos including the latest movie trailers from Apple. You can add more from the Plugins menu.

For videos, we'd recommend installing *VLC*. It has a reasonably good library manager, and is also more bulletproof when it comes to playing back odd file formats than *Totem*. Plus, it handles remote libraries over UPnP well.

The best all-round media manager, though, is arguably *Banshee*. If you're moving from Windows and want something that looks and feels like Media Player, this is certainly worth downloading and trying out. It handles music, video and UPnP better than the default apps, plus it can connect directly to the Amazon music store.

One thing that Ubuntu is very good at, however, is controlling your music or video while the player app is in the background. Once a song is playing in *Rhythmbox* or *Banshee* (or even Spotify) just click on the volume icon in the task bar and you'll get per-app volume controls and full playback tools without having to re-open the app itself. ■

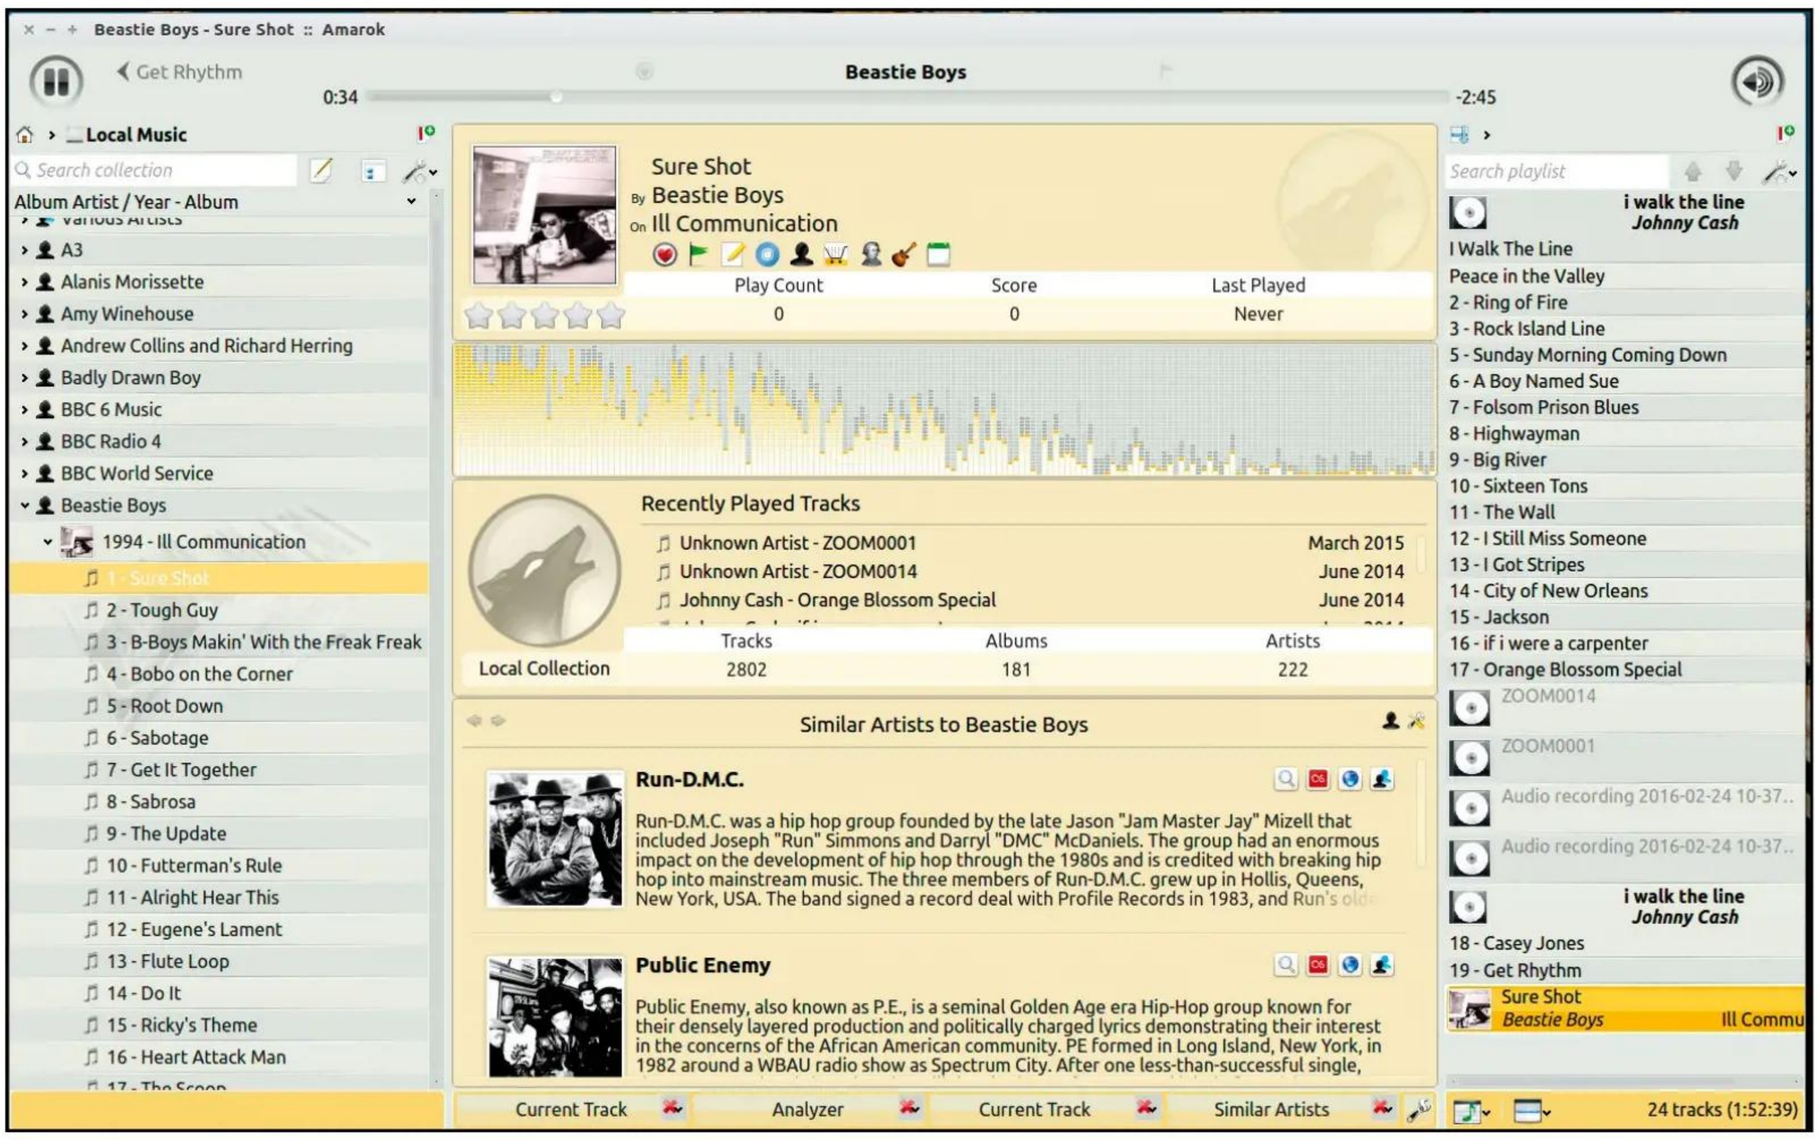

› **A good alternative to Rhythmbox, Amarok includes a more colourful interface and more add-ons.**

Quick tip

Rhythmbox will automatically check for and download your favourite podcasts. Go to Podcasts > Add and you can search the iTunes and Miroguide libraries, or add your own via RSS.

Audacity: Editing audio

Splicing, cutting and getting rid of awful noise are all within your grasp thanks to the free frills of Linux's top audio editor.

Quick tip

If you're trying to remove a section of audio, zoom in and make sure you're only the file at the exact point the waveform goes across the zero line. Anywhere else, and the edit will have an audible click when played back. Also remember you can insert silences and background noise to make edits to speech sound more natural.

While Linux has a sterling reputation as a server and desktop operating system, it's less well known that many major studios also use Linux for their creative workstations too. When it comes to high end audio production, for example, there are many powerful and stable packages that professional studios use for mastering music and film soundtracks all around the world.

Industry grade programs for audio editing include the excellent *Ardour* and *LMMS* (formerly known as the Linux Multimedia Studio). You'll find both of these suites in the *Ubuntu Software Centre* and they're well worth looking at just to understand the full potential of your OS. Both are completely free to download and use, and open source to boot.

Before you get too involved, however, it's worth pointing out that *Ardour* and *LMMS* are top-end Digital Audio Workstations (DAWs). They're a complete recording studio in software form, with virtual mixing desks, WAV editors, MIDI sequencers and more built in; *Ardour* or *LMMS* give you the kind of environment the late George Martin would have paid millions for, literally for free. DAWs support a huge variety of inputs and external hardware for audio production, and include virtualised version of equipment that's the mainstay of audio engineers.

For the newcomer, however, they're more than a little bit intimidating. They each have multiple windows and are stuffed with technical jargon. What if you're not scoring your latest composition for six mics and a string quartet, but just want to edit a podcast or a recording of your kids' school choir into something you can upload online? Don't worry, Ubuntu has you covered there too.

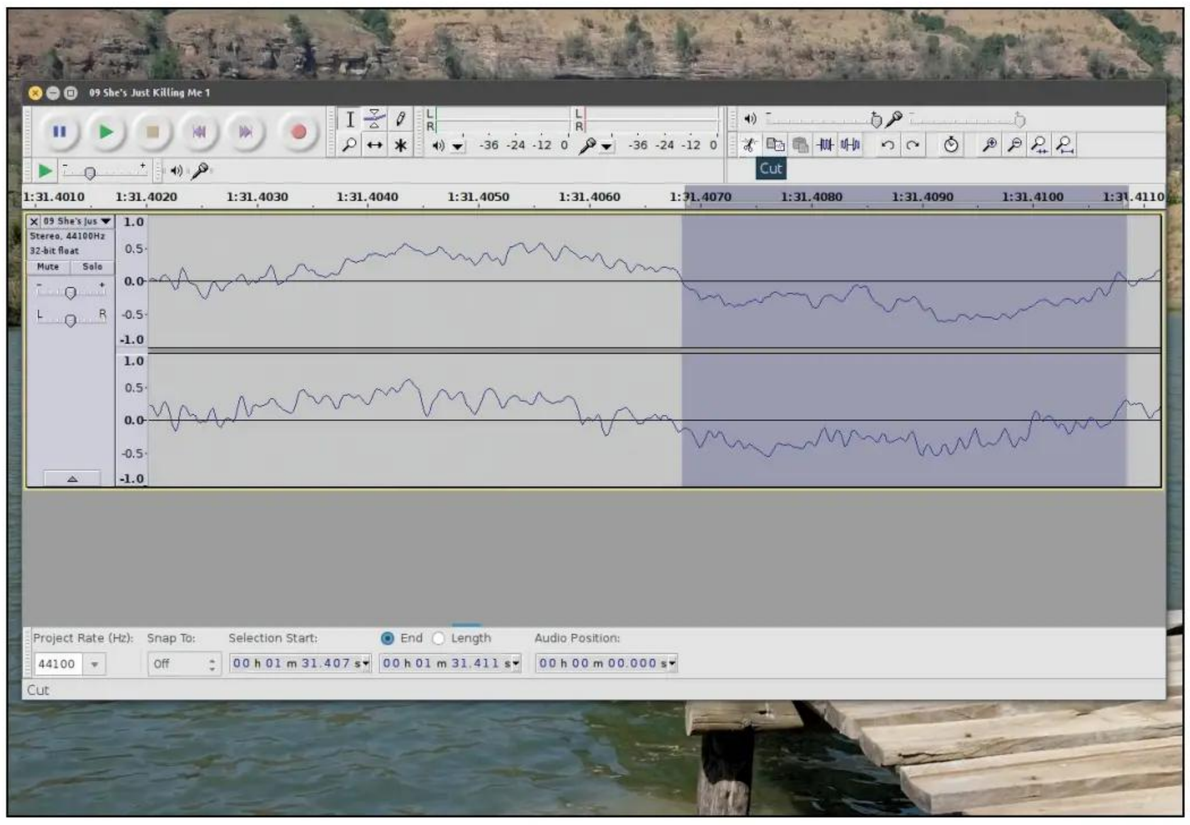

› If you need to remove segments of audio seamlessly, zoom in and cut at the exact point where it crosses the 0 line.

Sound it out

The "go to" software for audio editing in Ubuntu is *Audacity*. Actually, *Audacity* is the hero application for basic sound work on all platforms, it's really that good. In Ubuntu, you'll find it in the *Software Centre* just by searching for its name, so installing it is just a click and a password away.

By comparison to the DAWs mentioned above, *Audacity* is much more simple in its appearance. By default there's room in the window for a single stereo sample and a bunch of straightforward editing tools. Don't let that deceive you, however. *Audacity* has just about everything even a professional needs for preparing tracks.

First thing you'll need to do with *Audacity* is feed it some music, but before that you'll need to make sure that the music can be read. Just as with playing back music (*see page 66*), if you want to work with sounds stored in popular but closed formats like MP3, you'll need to install Ubuntu Restricted Extras from the *Software Centre*. This will add all the codecs for common media files to your PC.

From *Audacity*, there are two ways to load music into the editing window. Under the File menu, you'll see options for Open and for Import. If you select Open, or you drag and drop an audio file into the *Audacity* workspace or onto its icon in the Dash, your file will open up as a new project. If you select Import, it will load as a new track (or tracks) into the project you're currently working on.

Once you've done that, you'll see your file represented as a wave form in the editor. If you imported a file recorded in stereo, you'll see two tracks in your workspace, one for the right source and one for the left. You can add extra tracks using the Tracks menu at the top of the screen – they'll appear underneath the initial one or two.

If you've never seen an audio file as a wave form before, it's fairly easy to understand. The middle of each track, at the 0 point on the y-axis, represents silence. So if there's a part of the audio with a long flat line in the middle, there's nothing recorded there. Sounds are represented as spikes on either side of that line.

At this point, there are two keyboard shortcuts that you should learn by heart. The first is the spacebar, which starts the audio playing from wherever the current marker is and also stops playback. The second is the Control key. Hold this down and scroll the mouse wheel in order to zoom in and out of the waveform. If you zoom in close, you can see that what looks like triangular chunks on both sides of the 0 line is

All the Audacity

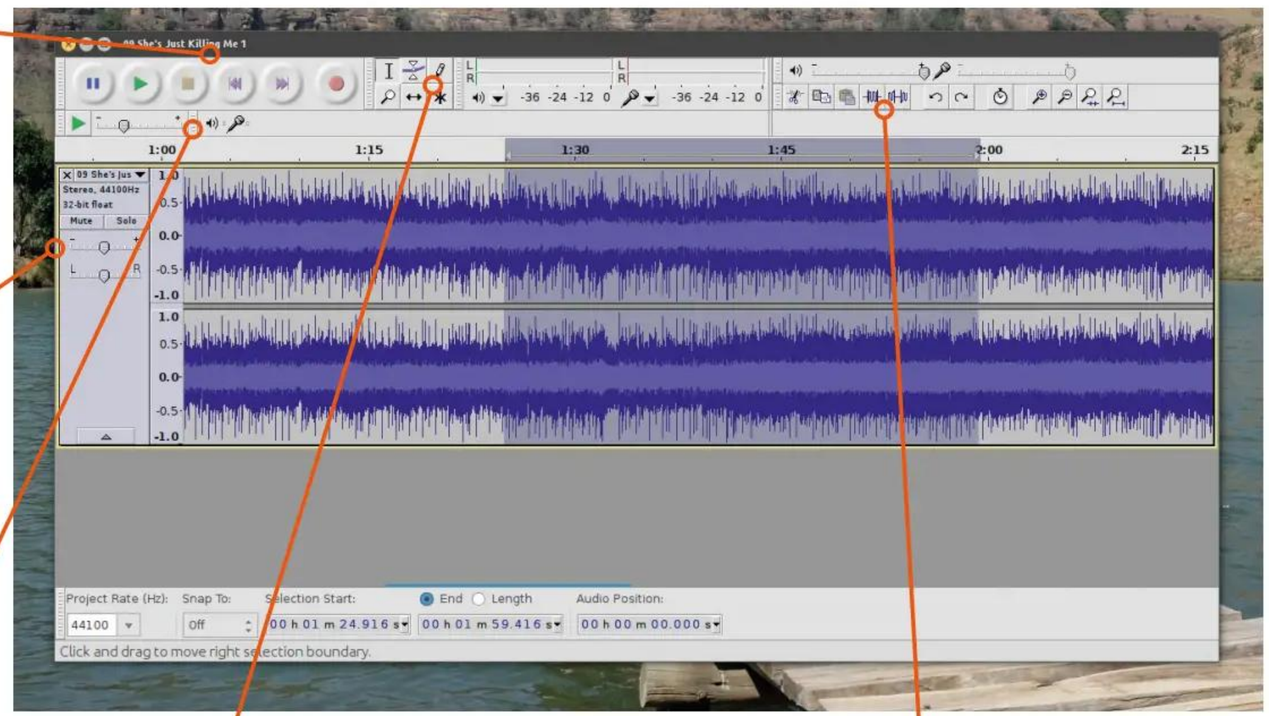

Playback controls
Simple playback controls let you listen to the tracks you're editing. Make sure you hit Stop before saving.

Track controls
You can mute a track or change the stereo balance between left and right from here.

Mic in
You can record directly from a microphone into Audacity – tweak the levels so you get good peaks.

Draw tool
The pencil icon is the draw tool, with which you can make manual edits to the wave form.

Trim and silence
These two wave icons let you remove the rest of the track and leave your selection or vice versa.

actually one long graph that moves rapidly above and below the line.

Using *Audacity* is simple. Move the marker by clicking your mouse on the time code, and you can play back from that point, mark it for editing later or slice a sample at the mark.

The main commands that you'll use to start with are Trim and Cut. These two are basically the inverse of each other: you begin by selecting an area of the file by click and dragging to highlight (or shift dragging from the current market position), then you can remove that section of a track using Cut or isolate it and remove everything else using Trim. Remember if you're working on multi-track audio trimming or cutting one track might leave it out of sync with others.

The next most important part of is the Effects menu. Here's where you'll find a suite of industry standard filters you can apply to your sounds. Select and area or an entire track and you can amplify it, compress the dynamic range (so the quiet bits get louder and loud bits get quieter) or add echoes, speed up the tempo and more.

There's far too many filters to describe in detail, suffice to say that the majority are software versions of those that have been in use in recording studios for decades.

Once you've trimmed your audio track, added your effects and edited in open and closing theme songs (assuming you're working on a podcast), exporting it to something useful is simple. By default, *Audacity* will save an uncompressed WAV file as part of a project. Choose Export Audio from the File menu to save as an Ogg Vorbis, MP3 or any other more common format suitable for playback on any device later. ■

› **Hugely powerful but also hugely complex, Ardour is a full recording studio on your laptop.**

Quick tip

If you want to try out *Ardour* or *LMMS*, head over to their respective websites, lmms.io and ardour.org. There's tons of documentation at both sites to help you get started on your professional mastering station.

Shotwell: Photo management

Looking for that special snapshot? Gather your photos and check out Ubuntu's built-in photo library app.

Quick tip

The Import function in *Shotwell*'s File menu will automatically copy images from a camera or USB drive to your main photo library, and index them in *Shotwell*'s database too.

Seven hundred and thirty billion. That's how many photos are uploaded to Facebook every year. And that's just a fraction of all the photos that are taken worldwide every 12 months. To put that in perspective, back in 1999, Kodak estimated humanity took 80bn photos a year.

We take a lot of photos. And even if every single one isn't a 'keeper', the chances are that your digital library is far more complex and vast than the cardboard box full of 'albums' that's mouldering in the attic. Fortunately, there are more than a few tools to help you sort through your library.

Shooting to Shotwell

By default, there are two ways of managing your photo library in Ubuntu. The first is to use the stock file browser. As long as your photos are organised into folders on your hard drive in a way that has meaning for you, such as by date shot, Files has a thumbnail view which can help you identify pics while looking through your hard drive. Just press Ctrl and spin the mousewheel to zoom in on a preview image for a larger view.

Ordering your photo library is essential, but there are better ways of browsing it. In Ubuntu, the default photo manager is *Shotwell*. It's beguilingly simple in appearance, but boasts top end features for organising your images. First of all, though, you'll need to tell it where to look. Open up the Preferences menu and you'll find an option to define your library. This is simply the folder on your hard drive in which your photos are stored. If you have a lot of photos, though, you might want to keep them on a separate drive.

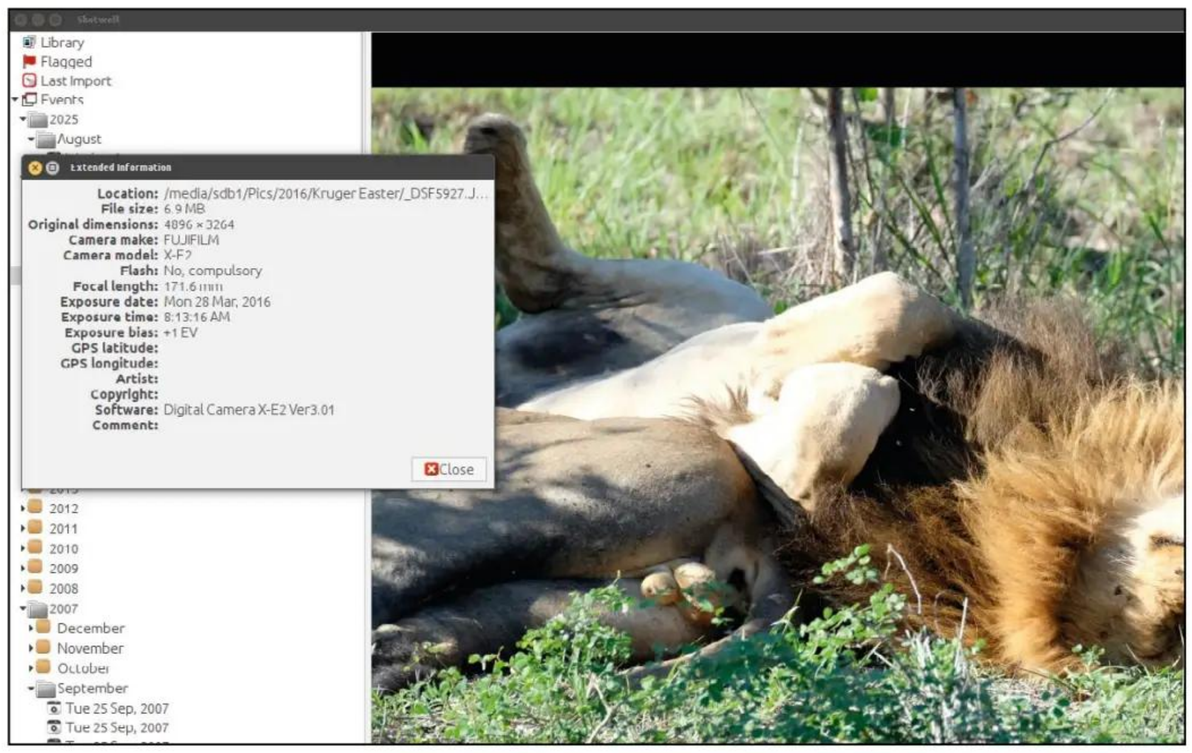

› Select Extended Information from the View menu to see the full detail of an image's EXIF file. This may even include location data for where the shot was taken.

Once you've set-up your library, *Shotwell* will begin the process of indexing it thoroughly. It will locate images images in folders that are inside the main library folder, so don't worry – if you've put a pic in there somewhere, *Shotwell* will locate it. Be warned, though – this can be a slow process, and *Shotwell* has been known to lock up if given too many files.

Navigating Shotwell

The way *Shotwell* creates its own index is by using the metadata stored in the photo files. The most basic metadata tag is the date an image was created. Once *Shotwell* has finished indexing your drive, you'll notice that the main index on the left is constructed to let you quickly zoom to a particular year, month or day to begin browsing your library contents.

Initially, you'll be presented with the date index on the left and thumbnails of images on the right. The images will be sorted into date order, with the newest first. If you scroll down, you'll go back through your library one sequence at a time.

Select a year in the left panel, and the thumbnails on the right will change. Instead of linking to photos, they'll now link to an album for a particular date. You can click through these to see all the images shot on that date. *Shotwell* will do its best to further segment your images up into "Events". So if it sees two groups of images shot on the same day but a couple of hours apart, it will create two albums for that day. You can right-click on the thumbnail that represents that album and either rename the Event into something meaningful rather than just a date or you can combine events into a single album if *Shotwell* has incorrectly split them up.

Shotwell will also index images by tags and location, if that data exists in the file. If it doesn't, you can add tags to a picture by selecting a single image or a group and pressing [Ctrl]+[T]. That'll make it much easier to find a pic later.

There are a few other features which are standard to most photo library applications, the most useful of which is the ability to tag photos with a star rating or a flag. You can do this by right-clicking on an image and selecting a rating or flag. This is particularly useful if you've taken a lot of photos on a single day, since you can make a quick first pass of the thumbnails to figure out which images you might want to keep. Then go to the View menu at the top of the screen and choose the Filter option – this lets you hide the unselected images without deleting them, so you can focus on further culling for "keepers".

Editing from within Shotwell

In a similar way, *Shotwell* also has a basic built-in editor which will allows you to add quick fixes to photos without changing

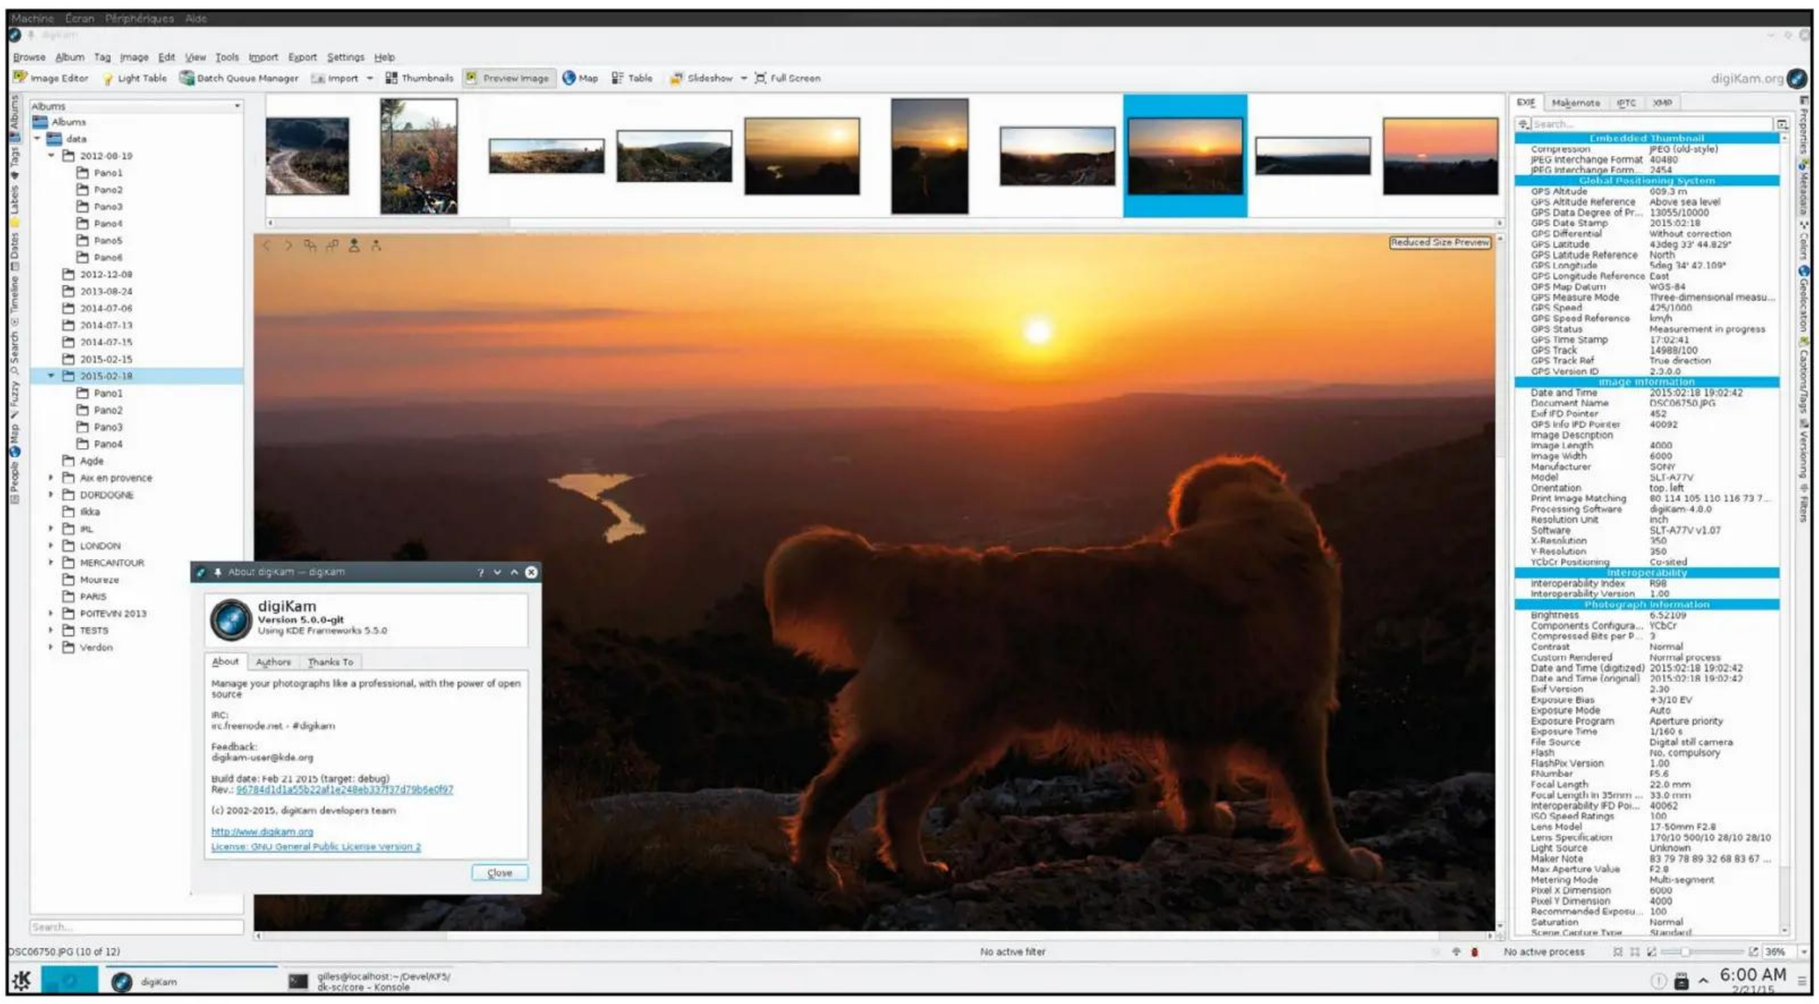

CC BY CAULIER GILLE

› HdigiKam is the default photo manager in KDE and is a very powerful, but complex, piece of software.

the original file. Once a photo has been selected, you'll see some simple tools at the bottom of the screen. These are mostly self-explanatory, Rotate, Crop, Straighten, Enhance and so on. Once you've made changes, you can export the image to a new file under the file menu.

If these basic tools are not enough, you can set-up *Shotwell* to automatically open your favourite photo-editing software through the Preferences menu. Usually, this will be *Gimp* (see page 68) and you are able to launch it just by right-clicking on an image and selecting "Open with external editor".

The final really useful feature baked into *Shotwell* is that you can upload directly from its interface to your favourite online photo-sharing site, like Flickr. You can add online service by opening Edit > Preferences > Plugins > Publishing, and adding account details in Ubuntu's "Online Accounts" manager that you'll find in the operating system Settings, accessible from the Dash.

Some of these references are a little out of date – for example Google's Picasa is listed although this has been subsumed by Google+ and Google Photos now – but uploading is easy nonetheless. From the album view, select the images you want to share and click Publish at the bottom of the window. Images will be sent along with metadata and captions as fast as your net connection allows.

For all that, though, *Shotwell* does have its limitations and can be buggy if you're working with large numbers of pics. If you want a really advanced library manager, and don't mind installing a lot of the KDE dependencies, try the excellent *digiKam* from the *Ubuntu Software Centre*. ■

Quick tip

You are able to add a tag to an image by simply dragging the picture thumbnail to a tag name in the left-hand panel.

Learn Shotwell well

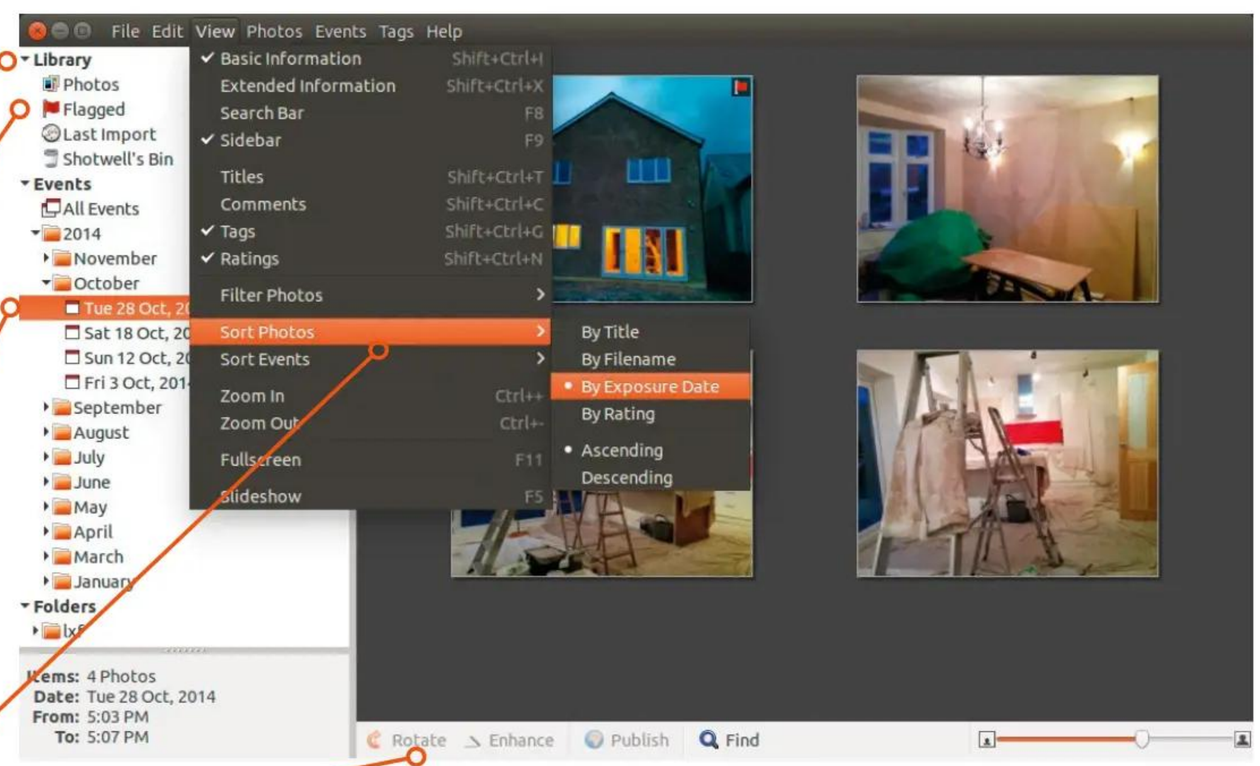

Get around

At the top of the list window, Library takes you back to the all photos view.

Flag day

Filter so that you only see your Flagged images by clicking here.

Dateline

The library is broken into a timeline, ordered by Year, Month and Date.

Sort it out

Shotwell sorts images based on date taken. You can change this information by selecting thumbnails and opening the Photos menu.

Photo filters

The Search bar has quick filters and a text box for hunting picture titles. You can turn it on and off by pressing F9.

See it all

Want to go full screen? Just press F9 and the image selected will blow up.

Gimp: Editing photographs

Think you need professional tools to do a professional job? Think again: Linux includes the best free photo editor around.

Quick tip

For those who like to shoot in RAW mode on their cameras, there's two excellent options for processing on Linux. *Darktable* is a fully open source RAW editor you'll find in the Ubuntu repositories, while Corel's *AfterShot Pro* is a well supported commercial contender.

There's one name which dominates digital photo editing all around the world. Adobe, and more specifically its *Photoshop* package, has been so popular for so long it's synonymous with the task and has become a verb. To "Photoshop" a picture is the same as conflating "to Google" for internet search or "to Hoover" for vacuuming the floor.

The bad news is that the vast majority of Adobe products aren't available on Linux operating systems. *Photoshop* certainly isn't and probably never will be.

Is that it then? The joys of Ubuntu and its ilk are not for photographers or those who like to tweak and digitally develop their snaps before sharing them with the world? Of course it doesn't.

Get your Gimp on

We'd wager that even on Windows most people who "Photoshop" their pictures aren't using Photoshop. It's an expensive piece of software after all and there's lots of alternatives. One of the best, which is available on Windows and Linux, is the wonderfully named *GNU Image Manipulation Program*.

Gimp is astonishingly powerful, so much so that it really is almost as feature rich as *Photoshop* and certainly as capable for almost every task. It's installed by default in Ubuntu, but if you don't have it you'll find it in the *Software Centre*.

When you first open *Gimp* up, it can be a little intimidating. This is a professional piece of software and designed to be used by professionals. Don't be put off though, it's easy to get a grasp of the basics.

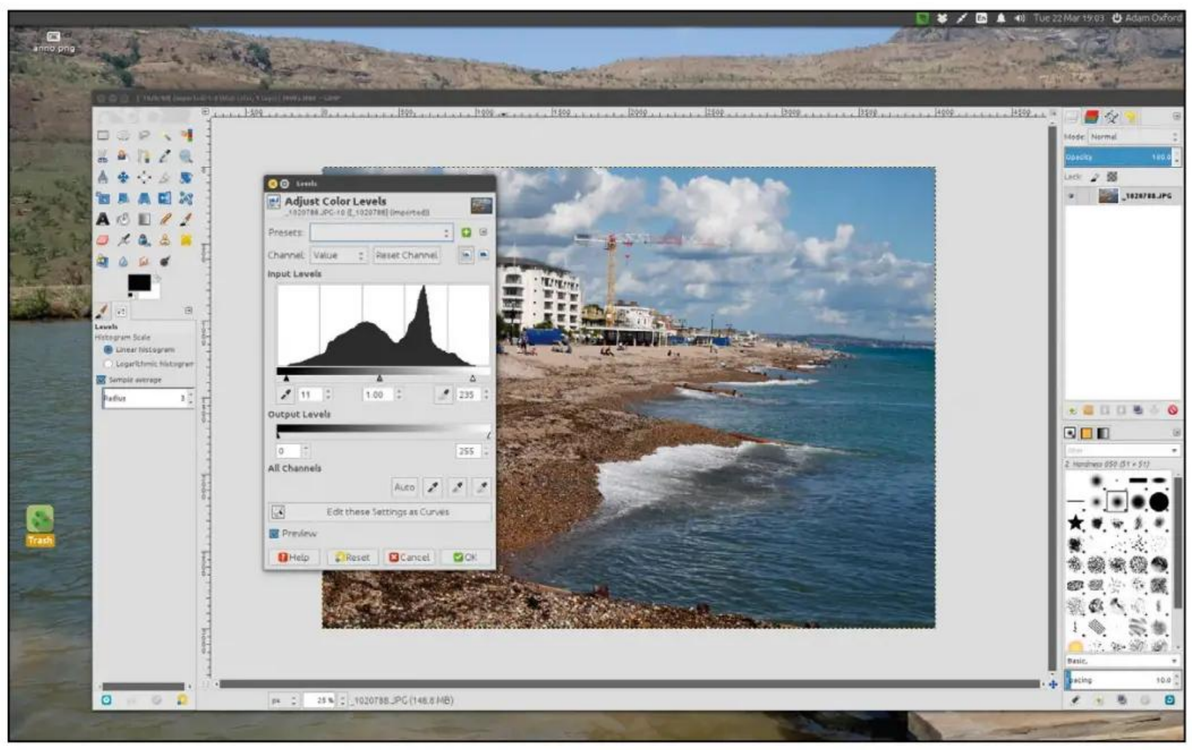

› **The Levels dialogue is the first port of call for photo fixing.**

The first thing you'll notice is that not one but three windows have appeared on your desktop. In the classic layout, you'll find the tools menu in a thin window on the left, the main workspace in the middle and a window for extra controls like layers, colour channels and brushes on the right.

It can be tricky at first to work in three different windows simultaneously, and there's an option to display everything in one Window if this is putting you off. Just go to Windows and tick Single Window Mode. This will give you a simpler layout that you may be more accustomed too.

If you can bear with it, however, we'd strongly advise sticking to separate windows for all the docks. It allows you to create a working area on your desktop that's exactly customised for your needs, and if you have more than one monitor you can put the tools on one screen with the image full width on another.

Colour correction

Gimp is, as we've said, a very powerful piece of software. It includes tools for cutting out parts of an image, creating layer masks, the classic clone and heal brushes loved by art editors for removing traces of humanity from a model's skin and a wide range of filters and plug-ins – plus support for more – for adding effects and punishing pixels into submission.

To get you started, though, here's an easy exercise in brightening up an underexposed photo. First, you'll need to get your photo into *Gimp*. A quick note – while *Gimp* does support RAW files from cameras, we wouldn't recommend it. Better to convert your image to JPG, TIF or PNG first.

Now you can open your pic up just by dragging its icon from the Files browser over the *Gimp* launcher or into the workspace itself if *Gimp* is already running. You can also open an image using the File menu if you prefer.

Our underexposed image is a little dark, but we're not just going to increase the brightness. That would be crude and have the effect of removing some of the dynamic range leaving you with a flat, washed out looking shot. We want to try and keep rich shadows and preserve as much detail without leaving the picture looking washed out.

The way to do this is to click on Colours and select Levels. This is the tool that professionals use to correct exposure. What you can see is a histogram which describes the "tonal range" of your image. The left side of the graph represents the amount of pixels in dark areas areas of the shot, and the right side light.

You'll notice three markers under the graph: one at either end and one in the middle.A well exposed shot will have readings all the way along the tonal range, depending on the contrast in the image captured. If a photo is underexposed,

Gimp key features

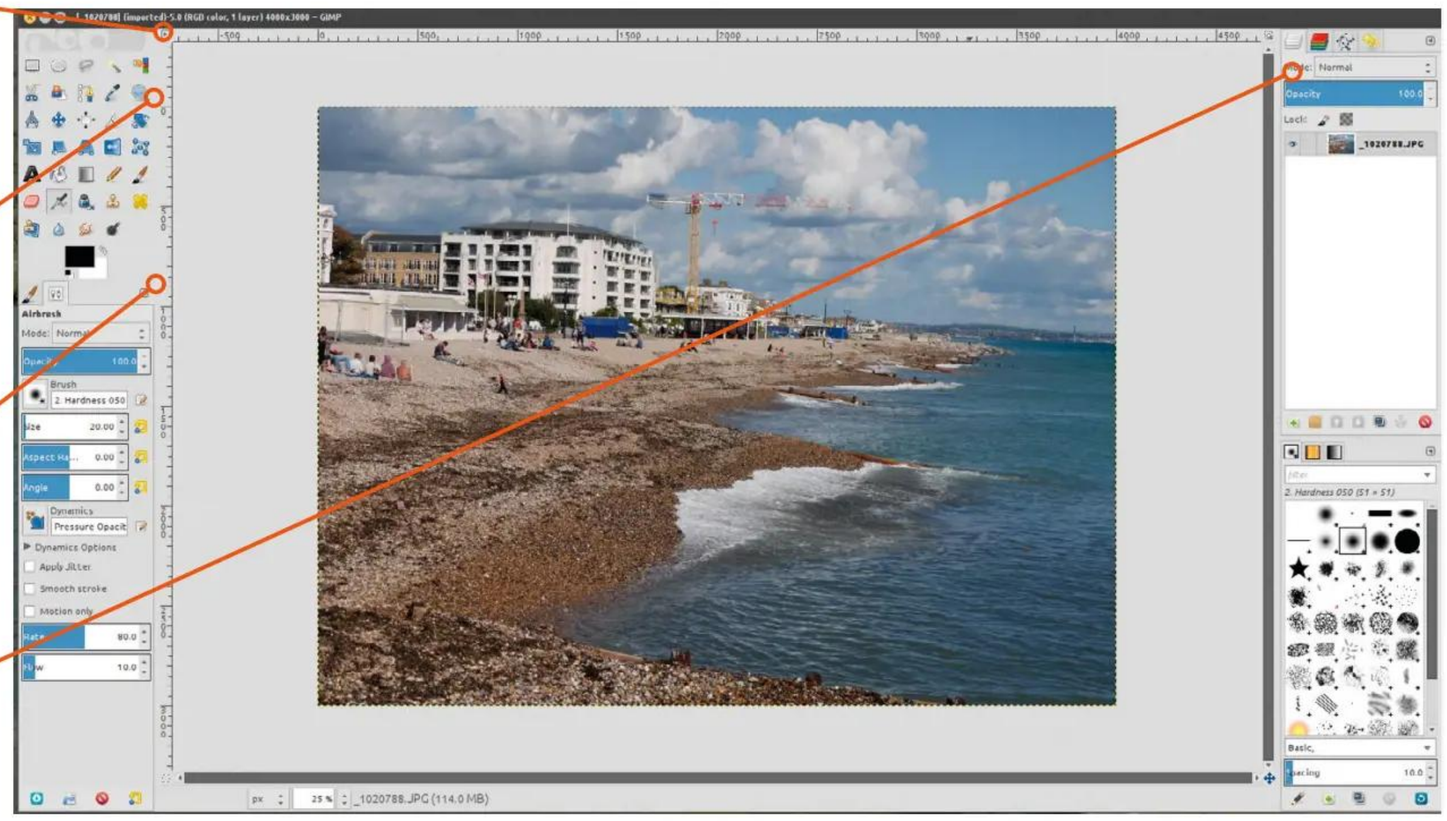

Filters
The Filters menu has a host of tools for everything from adding blur to cartoony frames.

Tools
The tools are a standard photo editing kitbag. Everything you'd expect is here.

Options
Beneath the tools themselves you'll find a full array of customisation options for tweaking their use.

Layers
The Layers window acts much the same as *Photoshop*. You can split an image, add extra layers and edit areas here.

however, all the readings will be bunched up at the left-hand side of the graph with a flatline towards the right. The best way to correct this is to drag the marker from the right towards the area where the tonal values kick in, reducing the tonal range of the image to those values that have readings.

If you're lucky, this will brighten the picture without destroying detail. Because the midpoint will also automatically adjust, it should retain a good contrast too, which merely changing the brightness would wash out.

As if that wasn't enough, this single tools can also act to correct the "white balance" of your picture. If you've taken a photo indoors, for example, you might have captured a yellow cast from the overhead lights. To remove this, you can use one of the three pipette symbols in the same Levels dialogue.

From left to right, these allow you to select the black point, gray point and white point of an image. *Gimp* will work out what the other colours should look like from there. The white point is the easiest to use – just click it then zoom in on a detail you know should be close to pure white. Teeth or eyes are are good, as are clothes or table clothes. Select that as your white point and watch the whole image adjust to something more natural looking.

And don't worry if you don't get it right first time. Just click on an different part of the image or the Reset button if all goes wrong.

That's really just scratching the surface of what *Gimp* can do. In the basic toolbox you can scale images, select areas, draw paths, rotate or add text. The only way to learn more is to try it out. ■

Quick tip

Remember that for quick photo fixes there are tons of online photo editors that work in a Linux browser just as well as they do in Windows. **Pixlr.com** is the best known, and great for quickly fixing red eye and so on.

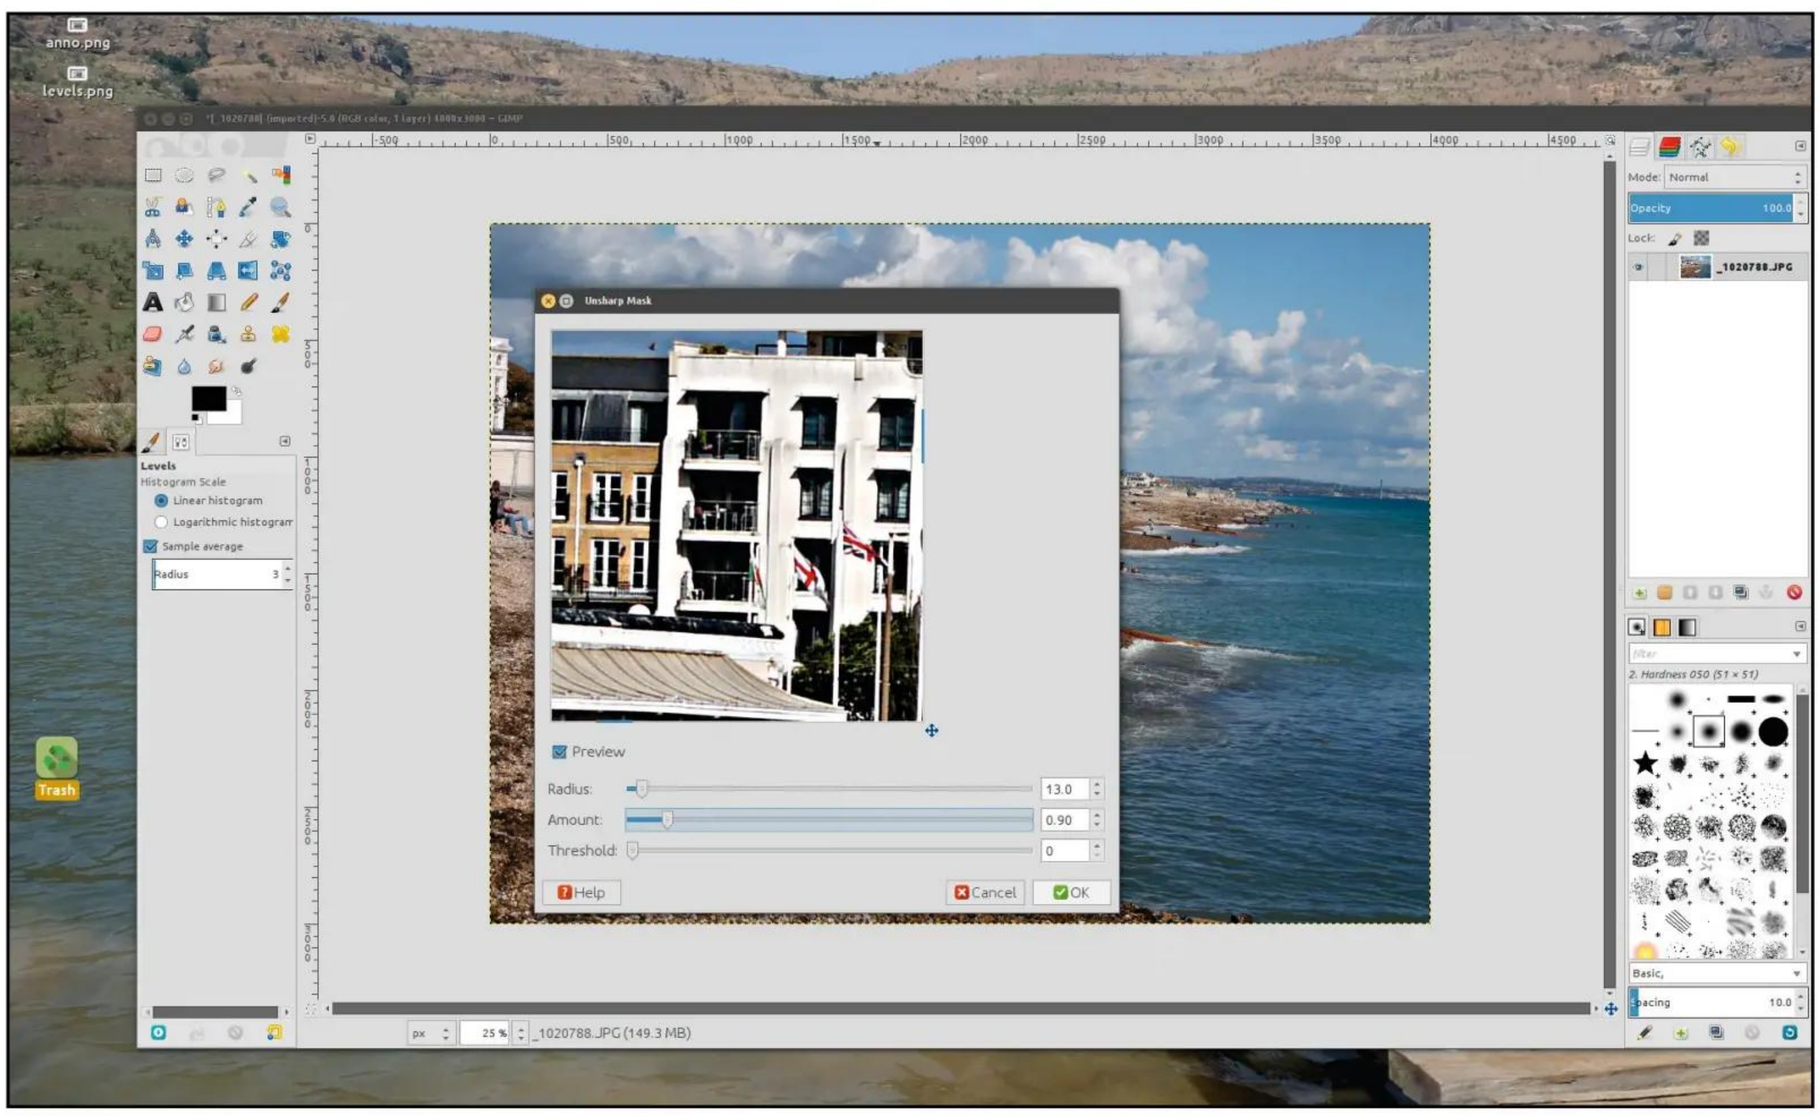

› **Gimp's Unsharp Mask is every bit as good at picking up detail in images as Photoshop's.**

Openshot: Video editing

You've got some lovely footage, and you've got some footage that must absolutely never be viewed by anyone. Time to splice!

Quick tip

In *Openshot*, you can right click on a clip for options such as fade in and out. You can also place several tracks on screen at once, in any of the four quarters of the playback window.

The likes of Instagram, Periscope and Vine may have made shooting and sharing short video clips much easier than in days of yore, but what if you want your family movie to be more than 15 seconds of quick gurning at the camera and gone? How do you stitch together multiple takes of little Betsy amusingly trying to eat the dog's dinner, and add captions, subtitles and a title page for her gran?

For cutting and stitching video with transitions and effects, there's no replacement, yet, for the traditional non-linear video editing desktop application. And fortunately for you, Linux has plenty to choose from.

Two to pick

For most users in the Ubuntu environment you'll likely settle quickly on one of two packages. The first is *Openshot* (**www.openshot.org**), which you can install directly from the *Ubuntu Software Centre*. *Openshot* is a straightforward, no-frills editor that makes sewing together clips easy. For those migrating from high-end video editing suites on other platforms and looking for something capable of cutting together broadcast quality footage, there's *Kdenline* (**www.kdenlive.org**).

You will find *Kdenlive* in the *Ubuntu Software Centre*, but it's best to install it using the instructions on its own homepage, as this will include the latest versions of some packages it depends upon. To do this, open a terminal by pressing [Alt]+[Ctrl]+[T] and type the following three lines.

```
$sudo add-apt-repository ppa:sunab/kdenlive-release
$sudo apt-get update
$sudo apt-get install kdenlive
```

Before you install either of these programs, make sure you have the multimedia codecs from the Ubuntu Restricted Extras package you'll find in the *Ubuntu Software Centre* first.

Your first project

The very first thing to do when you start *Openshot* is to go to the Edit > Preferences menu and click on the autosave tab. By default, *Openshot* doesn't back-up your files unless you tell it to, but as all video editing software can be temperamental at times and prone to crashing, we strongly recommend you set this to 15 minutes or so.

That done, you can get started on your first cut. Back in the main window you'll see three main elements. The white space in the top left will become a list of all the video clips you

Exploring the Openshot interface

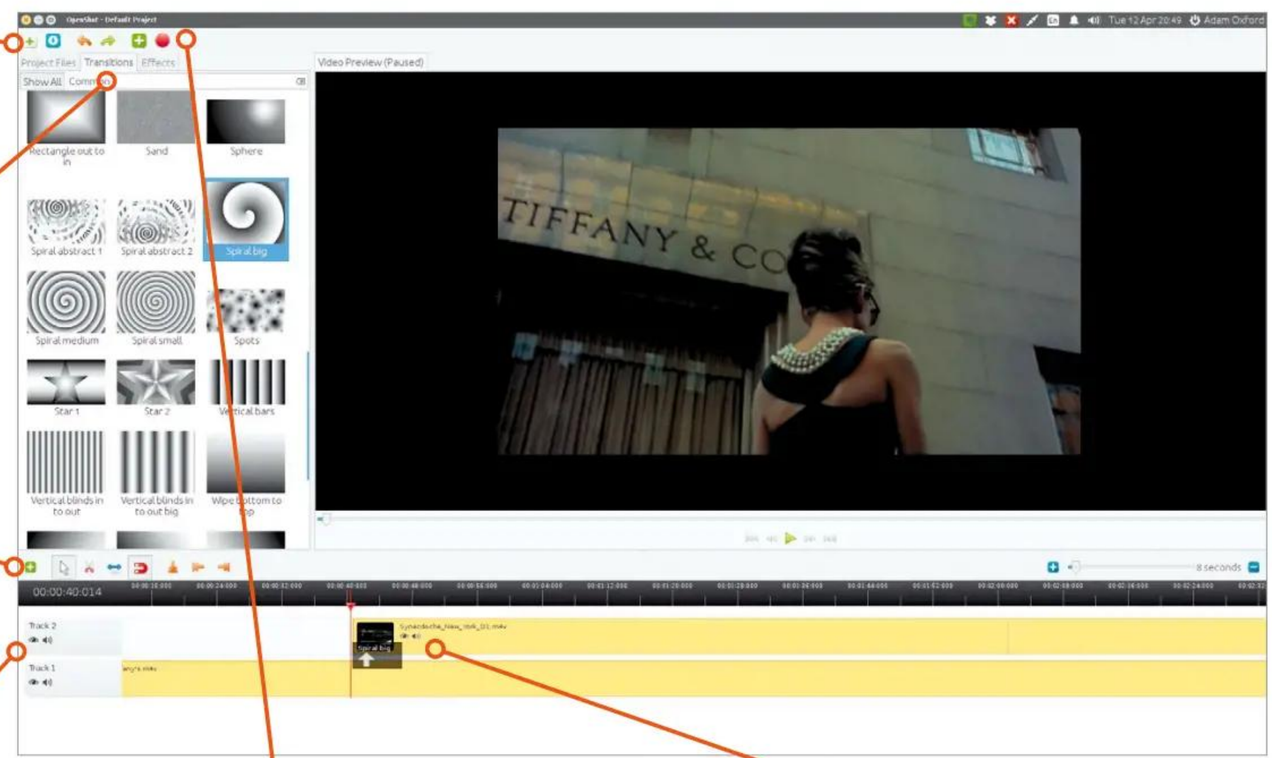

Bring 'em in
Import clips to the current project with the white Plus symbol here.

Mix it up
The Transitions tab is where you'll find the blend effects for splicing clips together.

Extra tracks
You can add extra tracks with the plus symbol in the middle.

Hide it away
The eye symbol hides the video in this track leaving just the audio (the volume symbol does the opposite.

Opening night
The red Render button exports your video to a file for playback elsewhere.

Press preview
The preview window on the right shows the current Timeline. Use it to inspect edits and Transitions.

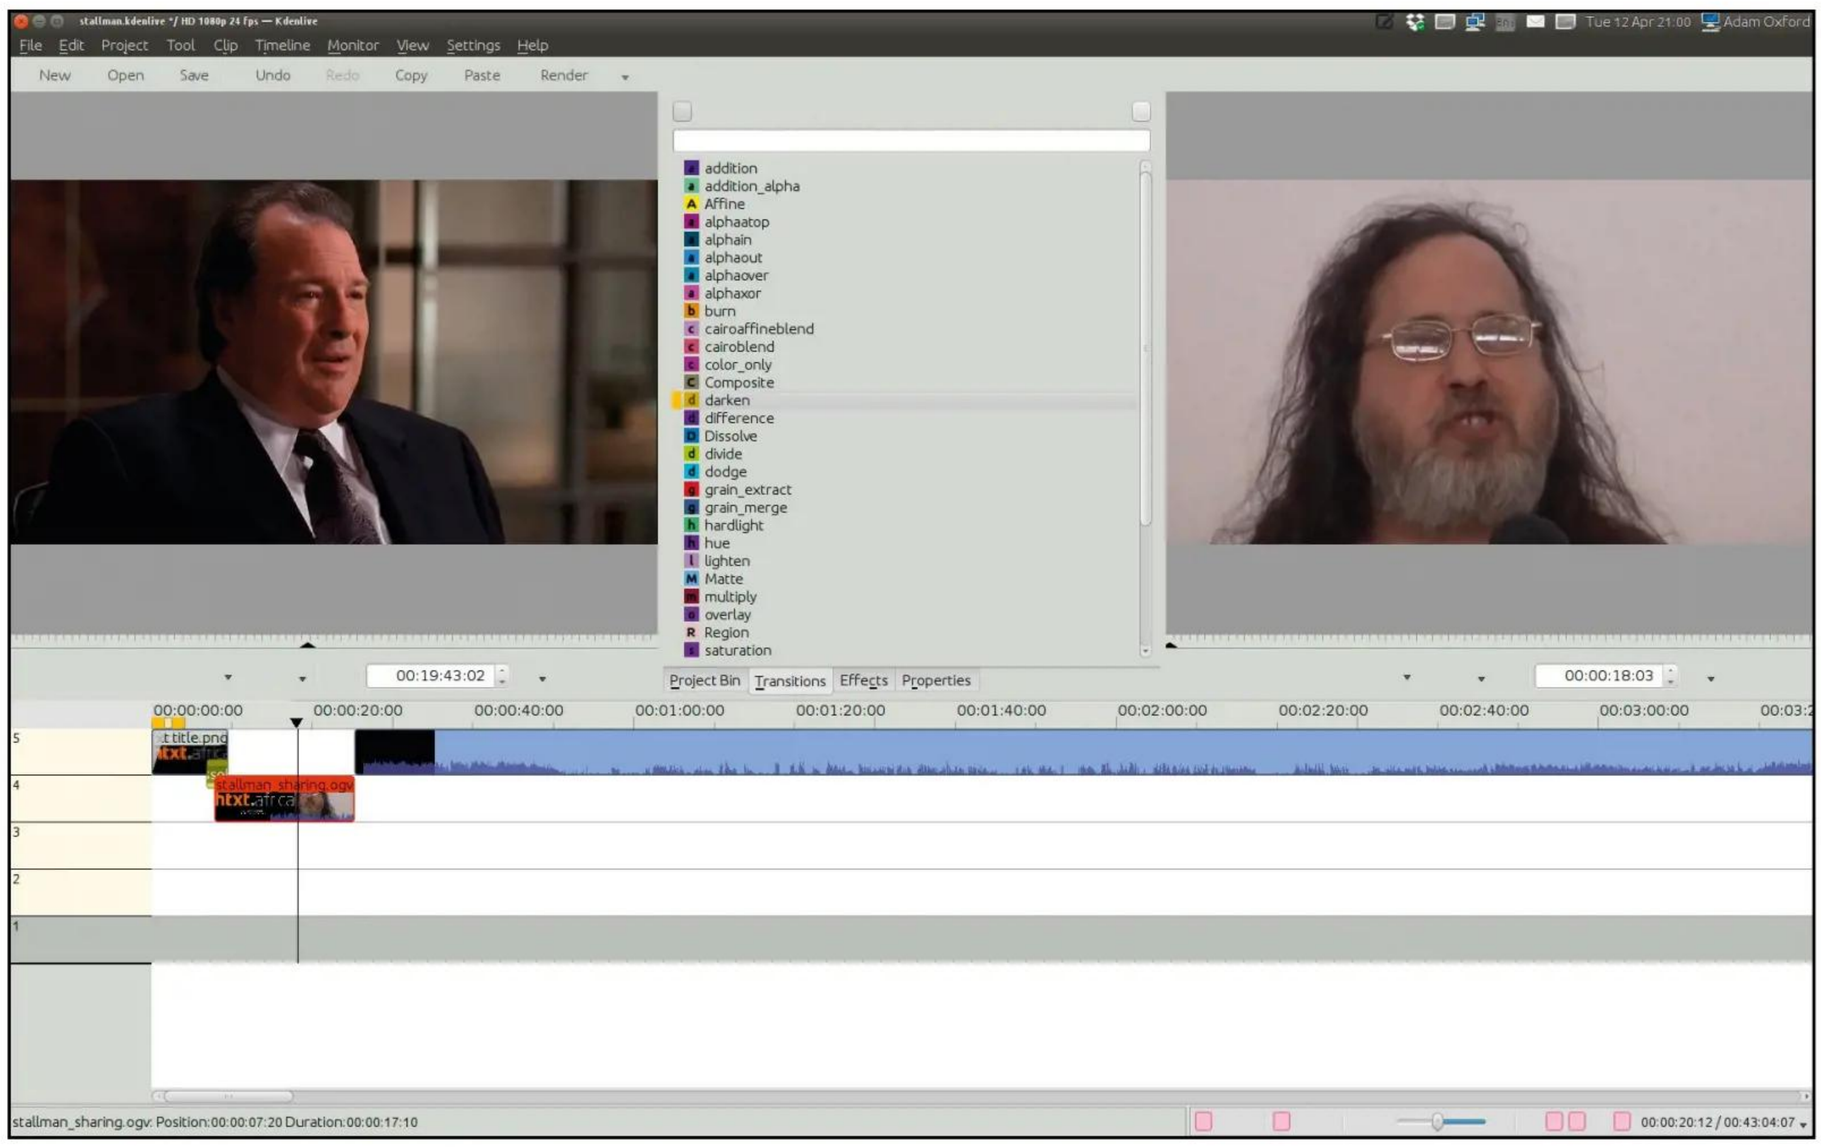

› Kdenlive looks more intimidating than Openshot, but it should be familiar enough to get started.

import into the current project. To the right of that is the video preview window and at the bottom are two Track timelines.

To import a clip, either drag and drop it form the file browser or desktop into the list space, or click the add clip icon at the top of the screen.

You will notice that *Openshot*'s toolbar, which is between the clip list and the timelines, is very small. In fact, there are only three real tools to pick from. The Arrow tool allows you to drag and drop clips around, the Razor tool (the scissors) slices a slip at the current point, and the resize tool automatically discards all video up to the current point. The downward arrow creates markers for reference when you are editing.

To get started, drag a clip from the list and put it onto one of the timelines. Now, you can start playback of the clip in the preview window with the spacebar, or the player controls under the preview window itself.

Editing is easy. Say you want to trim the first three seconds from your video – the bit where you're steadying the camera and issuing instructions to your actors. Stop playback or move the current playback marker to the point the real action starts and select the resize tool. Click on the point you want to start from the first few seconds will disappear.

Editing proper

Using the Razor tool is similar, but you keep the first three seconds as a seperate clip, which can now be manipulated independently on the timeline. This is useful if you want to take several short segments from a long clip, and stitch them together. Cut out the sequences you want to keep, and then select and delete the ones you don't.

The important part of video editing is how you stitch the clips back together again – and this is why there are two tracks in the timeline. To create an edit, you can either take two clips and put them together on one track, but the effect is likely to be jarring. Better to put align the clips, but put them on separate tracks with the end point of one clip slightly overlapping the start point of the other.

If you do that and press play, the edit will be the same as if they're together on the same track, but with the end point of the first clip overlapped. At this stage, you need to add a Transition – you'll see the Transition tab at the top of the screen. Transitions are effects that blend two clips together – either as a fade, a wipe, or something more dramatic.

To apply a Transition, drag and drop it over the point where the clips overlap. When you preview the footage you'll see a smooth blend with an added effect between the clips.

The third tab at the top, Effects, offers up a range of simple options for adjusting colours, brightness or applying other effects (like film grain or cinema-style lines) to a clip. You drag and drop these over the clip you want to alter just as with the Transitions.

Once you've stitched together a whole sequence, you can opt to add a separate third track with a separate audio, like background music. You can also use *Openshot* to make a photo slideshow: images can be imported into the playlist just like video clips. Just drag them onto the timeline and you can alter how ling they appear on screen for.

When you're done, click the Export Video button and tell *Openshot* how to encode the final video. This can be in a variety of resolutions and file format containers. Remember to save the project too before closing, so that you have the original timeline of edits to go back to at any time if you want to re-edit or make changes.

Taking it further

If you edit lots of video, you'll soon outgrow *Openshot*, and that's where *Kdenlive* kicks in. A full suite of editing tools that's used by broadcast media around the world, the basic principles of *Kdenlive* are exactly the shame. The key difference is that all the transitions and effects are far more customisable in *Kdenlive*, and there's support for more plugins and professional tools and techniques.

The other big difference is that you have two preview windows. The first is for editing a clip before adding it to the timeline, the second is a preview of the final work. ■

Quick tip

The second slider is for zooming in and out on the timelines. The most you can see at anyone time is a 200 second section of video.

Kodi: Media made easy

Kodi – formerly known as XBMC – is an awesome way to bring Linux right into your living room.

Quick tip

If you use *Kodi* via OSMC, it can also be used as a hotspot for other devices or wireless extender. Under internet settings you can turn a wired connection into an extra hotspot or, if *Kodi* connects to your router wirelessly, you can extend the connection through an ethernet port to a device that doesn't have wireless built-in.

There are plenty of great media players available for Linux, capable of finding, organising and accessing your digital media library wherever it is on your hard drive or around the house. But the ultimate media player, which turns your PC into a set-top box, is *Kodi*. Originally developed for the Xbox (and formerly known as Xbox Media Center), Kodi is an application that puts digital media into a "10 foot interface" for viewing with a remote on large TV. It's truly cross platform – there are versions of Kodi available for Windows or macOS, but Linux is its spiritual home.

Kodi is brilliant. It can be used to playback digital media, and can also play CDs, DVDs, Blu-ray discs and more, plus it has an extensive library of add-ons put together by fans that will access online resources like streaming BBC iPlayer or any one of a number of TV channels from around the world – including some high profile sports ones – and online archives of favourite TV shows, like *South Park*.

It also works with most DVB-T cards so you can watch live TV or turn your PC into a PVR system capable of recording shows for playback later.

Getting Kodi

Kodi isn't available in the standard Ubuntu repositories. To install it on your current machine, you'll need to open up a terminal and add the Kodi repository. Press [Alt]+[Ctrl]+[T] to open a terminal and type:

```
sudo apt-add-repository  ppa:team-xmbc/ppa
sudo apt-get update
sudo apt-get install kodi
```

Kodi is a small download, only around 75MB in its latest release form. After installing it, you'll notice that by default *Kodi* runs in full screen mode with no obvious way to get back to the desktop. Don't panic, just remember that it's designed for use on a big screen and ideally aimed at a media centre machine with a single purpose – not as a replacement for *Rythmbox* and *Totem*.

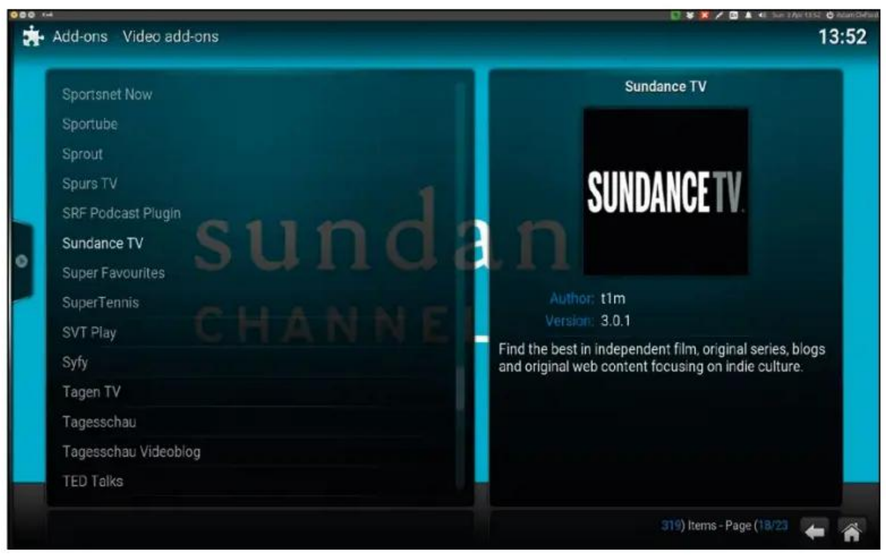

› Kodi has dozens of addons in its repositories, many of which allow access to content they probably shouldn't.

To get back to the desktop, simply click the power button and select Exit. To run *Kodi* in a window, go to System > Settings > System and click the Video Output tab. From here, you can then change the Display Mode to Windowed or Full Screen.

You can see from the above example that *Kodi*'s main weakness is that some of the menus are a little over complicated. The upside is that it's highly customisable: you can tweak lots of settings from subtitles to the kind of hardware accelerated decoding you prefer in the menus, so be careful not to change anything you don't fully understand.

Because *Kodi* is intended for use on a dedicated media centre/set-top box machine, you'll find several disk images at the kodi.tv website. These are customised operating systems (usually based on Ubuntu) that will boot a machine straight into *Kodi*, and have most desktop apps removed for maximum stability.

There are images of *Kodi* operating systems for non-traditional PCs too, like the Raspberry Pi (the best is *OSMC*) and there's an Android app too. The latter means it's available for dedicated devices like the Amazon Fire TV, which plugs into a spare HDMI port on your TV.

Getting started

Importing a media library into Kodi is fairly straightforward. Click on Videos, then Files and select "add Videos". This will bring up a file browser so you can point Kodi at the place where your videos are stored. It'll search through folders, so you can add a whole drive, but it will add all clips in the folder/drive selected. It's best, therefore, to point it at a dedicated video library unless you want a screen full of short clips sent via an instant messaging app or in your browser history.

The file browser isn't especially intuitive for newcomers if you're not used to the way Linux is organised under the hood. To find a hard drive other than the one your operating system is installed on, you'll need to go to Browse > Root file system > Media and find your drive, for example. The good news is that it's easy to add the contents of remote storage using the UPNP search tool in this same menu.

Adding music to your *Kodi* library is the same process from the Music tab on the homepage.

There's some really good news if you've got a large library of digital files built up over many years, with lots of different naming conventions and file types too. By default, *Kodi* will automatically cross reference files with online databases of TV shows and movies. That means it will display information about a show or film, and correctly number episodes and display the correct titles too.

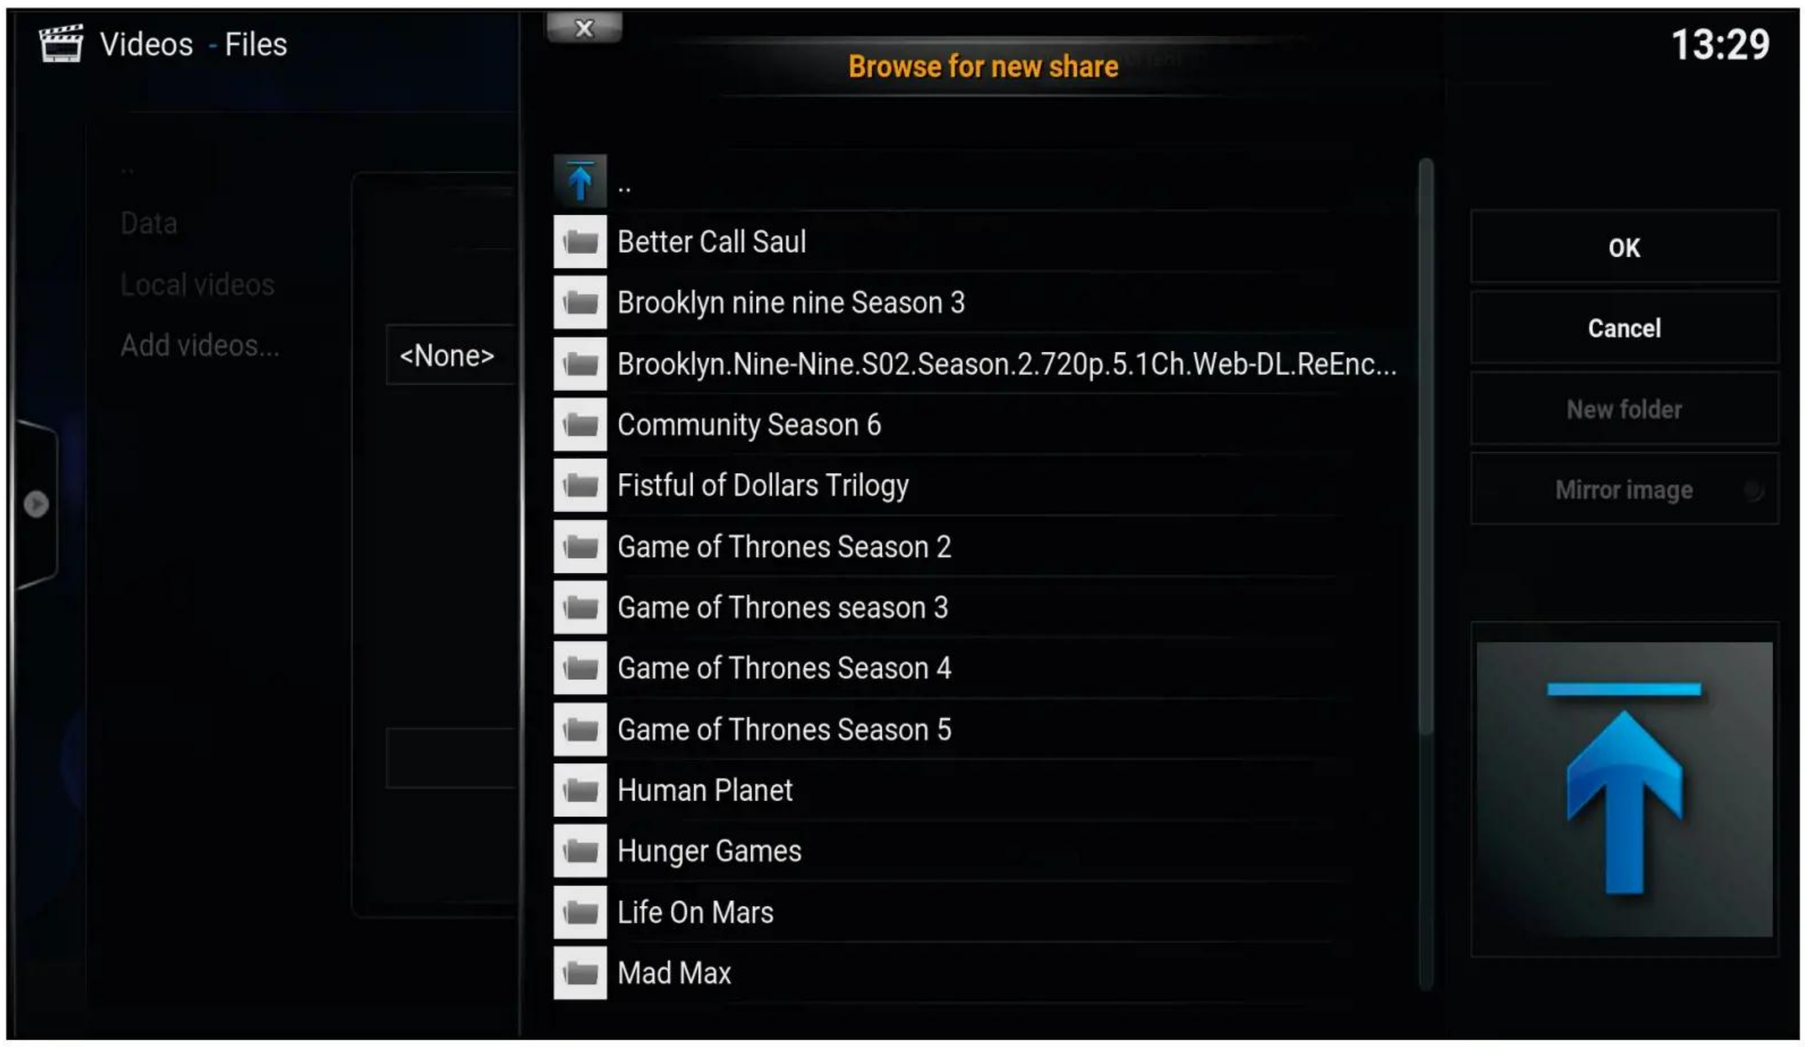

› **The add files dialogue gives you more options to define what kind of content you're adding.**

Quick tip

Under System > Settings > Appearance you can change the look and feel of *Kodi* by switching its skin. The default Confluence look is good, but give some others a go until you find one you like. Some may reorganise options on the homescreen a bit, but getting back to the default Confluence look is easy.

Controlling Kodi

There are tons of features inside *Kodi* to play around with, from gaming add-ons to weather information to new skins to try out. Most are well documented at **kodi.tv**, and sadly we only have room to scratch the surface of its capabilities here.

If you install Kodi as it's intended, on a low power PC or Raspberry Pi connected to your TV, you don't want to be faffing around with a keyboard and mouse to select, play and pause shows. You can add a traditional infra-red remote, if you have a USB receiver plugged in, but our favourite feature about *Kodi* is that it can be controlled from your phone.

Set-up from within *Kodi* is simple. Just go to System > Settings > Services. The Remote Control option you see here is not the one you want, but while you're in this menu, it's worth turning on UPNP so that other media players can take use *Kodi*'s library and display.

Open up the Web server tab, and select "Allow remote control via HTTP". Don't worry too much about the options here – and don't change them. They're for if you want to control *Kodi* via a web browser, but we're going to do something much easier than that.

On your phone, head to the App Store or Google Play. You'll find that there are official *Kodi* apps for remote control, but on Android at least we much prefer the unofficial *Yatse*. Download and install that to your phone.

When you first open up *Yatse*, ask it to search for local servers and it should find your *Kodi* machine. Select that and you have full control using cursors, an advanced menu system or the touch screen keyboard. ■

The Kodi code

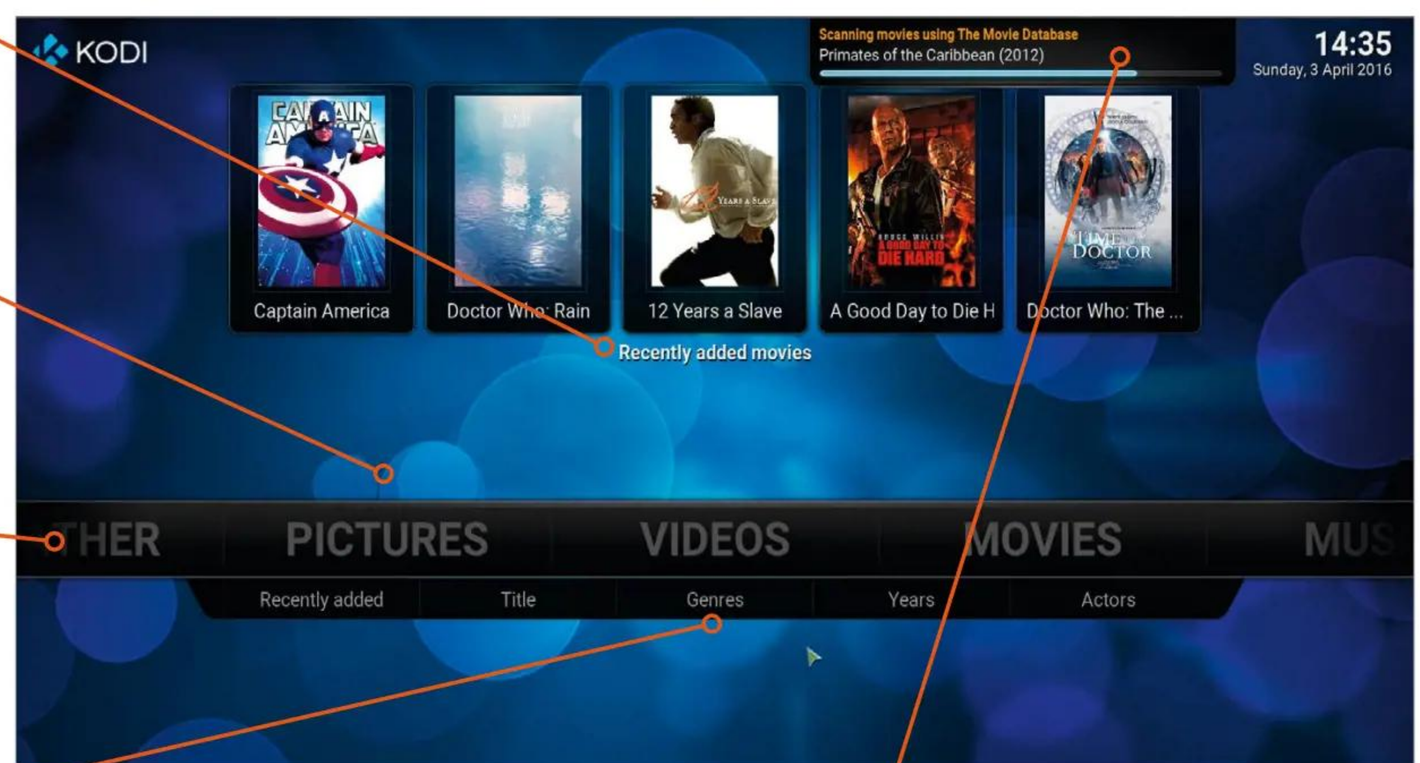

Recently added
Roll over this option and thumbnails of videos that have been updated or added will appear at the top.

Pictures
Kodi can also be used to browse images on your hard drive. This acts in the same way as the movie and music browsers.

Main options
Left and right on your controller (or clicking with your mouse) will take you to the different submenus.

Genres
You can search for media using metadata as well as by title.

Scanning dialogue
Kodi will cross reference files against online databases. The catch is it works best if you separate movies and TV shows into seperate folders.

Get gaming

Bored? Poppycock. Get that trigger finger ready and enjoy some games on Linux.

Gaming on Linux has never been better. With over 1,500 games available on Steam (store.steampowered.com) alone that support Ubuntu (officially from 12.04 LTS and, unofficially, many other distros), the breadth of games that you can get on Linux is astonishing: from casual card games and puzzles for whiling away a few hours, to indie games that might be low on budget but are high on inventiveness, and even big budget 'AAA' games are now available.

Games are also beginning to run better than ever on Linux distros now that many game engines used by some of the biggest names in the gaming industry – such as the powerful CryEngine.com – have native Linux support. Gamers switching from the bloated Windows OS to Linux can also see performance increases, as their system won't have to devote as much of its resources to background processes and instead concentrate on delivering the best gaming experience possible.

Not only is there a plethora of choice when it comes to gaming on Linux, it's now easier than ever to get started. Although Valve's proprietary store front has its issues – such as game DRM– it does make installing and running games an incredibly simple process. It's also not the only game in the Linux-verse, Good Old Games (**www.gog.com**) is another service that sells games for Linux. It offers a mix of classic games (that have been made compatible with modern computers), new games and has a no-DRM policy.

> **"We'll show you how you can prepare your system to handle any game you want to play."**

Hardware manufacturers have also been upping their game as well, making good strides to ensure their graphics drivers are Linux compatible, so we can fit the most powerful gaming components into our machines. Of course, the beauty of Linux is that if you're not too keen on using the hardware maker's proprietary drivers, there's a dedicated community creating third-party open source alternatives.

Thanks to the large variety of games available on Linux, the PC you use should be able to run at least some of them – though older machines without dedicated graphics cards may struggle to run the more graphically intensive and newer games – however you'll still be able to play some puzzle, indie and retro games.

If you want to give your PC a bit of a boost to handle more demanding games then you don't have to spend a fortune – a few modest upgrades here and there should give you a good experience when gaming. Most games – unless they are huge open world games – aren't particularly CPU-intensive, so if you have semi-recent dual- or quad-core processor, then you'll most likely be fine to leave your CPU as it is. Modern CPUs, such as Intel's Core i5 and Core i7 range, also include pretty decent integrated GPUs as well – which means you might not even need to buy a

› *Metro Last Light* is one of the more demanding games on Linux and also comes with a built-in benchmark tool.

› The Steam library lists all the games that you've purchased. Once bought, you can download and play a game on any machine as long as you've logged in once to your account.

dedicated graphics card to play many games. However, if your CPU is showing its age or you want to play more demanding games, then buying a dedicated graphics card is the best upgrade you can buy for your PC.

We have good news on two fronts here. First, most modern graphics cards play pretty nicely with Linux these days, so putting one in your system shouldn't cause any trouble. Both first-party and third-party drivers are quite easy to install too (more on that in a bit) and they generally behave themselves.

The second bit of good news is that you don't have to spend a lot of money to get a capable graphics card. Sure, there are some ludicrously expensive GPUs on the market, such as Nvidia's £1,149 Titan Xp, but these are aimed at gaming enthusiasts. For entry level and mid-range cards you can get something decent for between £60-£150, depending on the level of performance that you require. Despite these low prices you should be able to run most modern games smoothly thanks to most games being scalable for a variety of hardware configurations – so you might have to dial down an effect here or there, but you'll get smooth framerates.

Most modern games are also developed with the Xbox One and PlayStation 4 in mind, and despite being pretty new, these games consoles have relatively modest graphics capabilities compared to PCs. If you want to future proof your gaming machine a really general rule of thumb you can use when choosing a graphics card is that if you want to run games at resolutions higher than 1080p, or on more than one monitor, then we'd recommend going for a GPU with 4GB of GDDR memory. However, for most casual gamers this will be overkill, and any card equipped with 2GB should be plenty.

With this in mind, we'd suggest a graphics card, such as the Nvidia GTX 950, is a great choice if you want to play most of the latest games at 1080p. It comes with 2GB of memory and goes for around £120, which makes it excellent value for money. We got one in to test, in fact, and ran a number of benchmarks to see how good it performs.

Benchmarking

The first benchmark we ran was Heaven, which is a well-known benchmarking tool among gamers for seeing how well a graphics card performs. You can download the RUN file from **http://bit.ly/UnigineHeaven**. Once downloaded, open up the Terminal and type:

```
cd ~/Downloads
chmod 700 Unigine_Heaven-4.0.run
sudo ./Unigine_Heaven-4.0.run
```

To install the benchmark. To run it, change directory into the **Heaven** folder and type `./heaven` . A screen will appear asking you what type of benchmark you want to perform. Click 'OK' and an animated video will appear, putting your GPU through a workout. On the top-right hand corner of the screen you'll see a counter for frames per second – the higher it is the smoother the gameplay will be.

The test computer we're running the benchmarks on has an Intel Core i7-6700K processor, Intel's sixth-generation Core i7 flagship. It's one of the best CPUs around, and though it's a bit pricey at around £260-£300, »

Wine for games

Wine (which stands for Wine Is Not an Emulator) is a compatibility layer that enables Windows programs to run on Linux. Although it's usually used for running office suites and other programs that haven't made it to Linux (or don't have an open source alternative), it can also be used for running Windows games in Linux as well. Of course, there's a performance impact to running a game via *Wine* rather than natively, so it's not really recommended for graphically demanding games – unless you have one hell of a machine. However if you have CDs of some old Windows games lying around, or you want to try an indie that hasn't made it to Linux just yet – then *Wine* is a great possible solution.

To install, open up the Terminal and type

```
sudo add-apt-repository ppa:ubuntu-wine/ppa
sudo apt-get update
sudo apt-get install wine1.7
```

and a window will pop up asking you to read and agree to a EULA by Microsoft. You might have hoped you'd left all this behind, but it's a necessary step to using *Wine*. Once you've scrolled down and agreed, type in `$ winetricks` . To run the graphical user interface for *Wine*. You may have to install additional packages, and *Wine* will take you through the process. With 'Select the default wineprefix' selected click 'OK' and then you can begin installing required software to run Windows games in Linux.

» it's got good future proof qualities, with support for new DDR4 RAM and a powerful integrated Intel HD Graphics 530 GPU.

Without a dedicated graphics card, the Intel Core i7-6700K reached an average of 7.9 frames per second (fps), with a minimum fps of 5.6 in the *Heaven* benchmark. In comparison with a GTX 950 GPU installed, the *Heaven* reached an average of 39fps and a minimum of 20.4fps.

We also benchmarked the GTX 950 in *Metro Last Light*, a graphically demanding first person shooting game that was the first AAA game to be available for Linux. At full 1080p resolution the GTX 950 hovered around the 50 frames per second mark for most of the run, offering a smooth gameplay experience. However, we found there were sudden drops to 36fps when things got busy – though the graphics settings were set to maximum.

If you want a more consistent framerate then you can lower some of the graphical details, and it will still look lovely. It's certainly better than the integrated graphics that is on the Intel Core i7-6700K, which only managed 21fps in *Metro Last Light* – and that was while running the game at the lowest graphical settings, too.

› The *Heaven* benchmark will stress test your graphics card to see how well it performs.

Installing a GPU

It's clear, then, that if you want to play graphically demanding big budget games, a dedicated graphics card is a must-have on your shopping list. Graphics card are relatively straightforward to install. All you need to do is turn off your machine and pop off the side to get to your motherboard. The graphics card slots into the primary PCIe slot – check your motherboard's manual (or run a quick online search) if you're not sure, though most motherboards are clearly labelled up.

> **"If you want to play big budget games, a dedicated graphics card is a must-buy."**

One of the best things about putting together a PC is that it's pretty hard to plug something into the wrong port or slot – they simply won't fit. The PCIe slots have a little lever at one end of them, and pressing these down helps to eject the graphics card if you want to remove or replace it, so when installing you might have to angle the card so that the small cut out in the card fits under the lever. You'll also have to make sure that there's a gap in the chassis of your PC to allow you to reach the ports of the graphics card when it's installed. To create a gap most chassis allow you to unscrew a thin metal grille to give access to the card's ports. Find the grille that aligns with the PCIe slot you're going to be installing the card in, and use a screwdriver to unscrew the screws holding it in place – but make sure you keep hold of the screws. Remove the grille cover and then insert your graphics card into the PCIe slot. Now use the screws that you kept hold of and use them to secure the graphics card in place.

Many graphics cards – especially the more powerful ones – require that they are plugged in to your computer's PSU (power supply unit), so check to see if the card has one or even two of these PCIe six- and eight-pin sockets. Your PSU should have a cable running from it with matching connectors. If not, Molex adaptors should be provided with the card, however any good 500W PSU should provide at least one of these and we'd recommend upgrading if yours does not. You can now close up your PC and attach your monitor to your graphics card and power it on.

All going well, your PC should boot and you'll see the desktop environment of the distro that you are running. Steam, which has

› **Alienware is just one of the big names to get into the Steam Box business.**

Steam machines

Valve, the company behind the Steam platform, has been promising to disrupt the way we play games in our living rooms with Steam Machines. These are compact gaming PCs designed to replace gaming consoles, such as Xbox One and PlayStation 4. Running SteamOS, Valve's Debian-based distro, the first commercially available boxes should go on sale in November – hopefully by the time you read this magazine.

Steam Machines are being made by a variety of manufacturers, such as Asus and Alienware, and come in all different shapes and sizes. This means that if you don't want to build a gaming machine yourself, you should be able to find one that suits your needs – and budget. Prices start from $449 (around £292 – roughly the cost of a PlayStation 4 but with a heck of a lot more power) to a wallet numbing $4,999 (around £3,256) for Falcon Northwest's Tiki with its granite base.

You can find out more about Steam Machines and purchase them from Valve's own Steam website at **http://store.steampowered.com/sale/steam_machines**.

If you'd prefer to play your games with a controller rather than a keyboard and mouse, then many USB controllers will work plug-and-play with Linux and SteamOS, even Microsoft's Xbox 360 controller. Valve has also created its own Steam Controller (£40), with the unique device supposedly making it easy to control games that usually use mouse and keyboard controls with the gamepad.

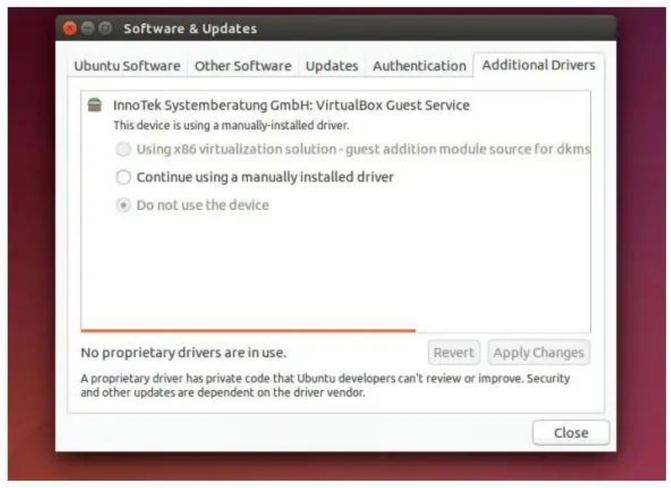

› The Additional Drivers system in Ubuntu makes it straightforward to install new open source or proprietary graphics drivers.

› The Nvidia GTX 950 is a very good graphics card for its price, though if you want a bit more oomph you'll need to spend a little extra.

the biggest collection of Linux games, runs on Debian-based systems, though only Ubuntu with Unity, Gnome or the KDE desktop are officially supported. We have run Steam on other Debian-based distros, Mageia and Arch with success, but for the ease of this feature we are running the latest version of Ubuntu.

Because of the growing number of Linux gamers both Canonical and Nvidia have worked together to make installing graphics drivers a lot easier. When Ubuntu is loaded you should see a message pop up saying that 'Restricted drivers are available'. Click this message and a window will open up listing a range of drivers that you can install. A recommended driver for your system will be highlighted, but it's worth checking some of the drivers as there will be information regarding whether the drivers are proprietary or open source. You can also install drivers from the terminal. Open it up and type `sudo Ubuntu-drivers devices`.

This gives you a list of possible drivers, eg for an Nvidia graphics card it could list:

```
driver   : xserver-xorg-video-nouveau - distro free builtin
driver   : nvidia-304-updates - distro non-free
driver   : nvidia-331 - distro non-free recommended
```

If we wanted to install the latest official proprietary driver from Nvidia – which Ubuntu recommends – then we would then type: `sudo apt-get install nvidia-331` to install.

You may notice that this isn't the bleeding edge proprietary driver from Nvidia, but it will likely run most indie games. For more graphically demanding games, such as *Middle-earth: Shadow of Mordor*, you'll need a more recent driver. In the past, this has meant either downloading and installing a driver manually from the Nvidia – which is probably too much *Terminal* exposure for most newcomers – or adding and enabling extra repositories (repos) to stay up to date. The latter option used to mean relying on third-party repos (eg **ppa:ubuntu-x-swat/x-updates** for the latest stable), but in an exciting and recent development Ubuntu has set up an official PPA especially for Linux gamers wanting the latest drivers. To fetch them, add and enable the PPA using:

```
$ sudo add-apt-repository ppa:graphics-drivers/ppa; sudo apt-get update
```

To upgrade to or install the latest driver use `$ sudo apt-get install nvidia-current`. If you're concerned about these drivers making your system bleed – most newly released Nvidia drivers can tend to have their teething problems – then we'd advise sticking to the x-swat PPA. To roll back changes made by the PPA use `sudo ppa-purge ppa_name`.

If you have an AMD graphics card and want to install the latest official drivers then you can download and install the AMD Catalyst program which gives an easy-to-use graphical interface for installing the latest drivers. If you have previously installed an AMD card you'll need to remove the old drivers by opening up the terminal and type in `sudo aticonfig –uninstall`. Now enter `sudosh ./amd-driver-installer-x86.x86_64.run` into the Terminal. Where it says 'amd-driver-installer-x86.x86_64.run' you'll want to double check that it's the most recent driver for your card on AMD's website (**http://support.amd.com**). A window will open allowing you to complete the installation. A word of warning, before updating your kernel down the line, remember to remove the proprietary driver.

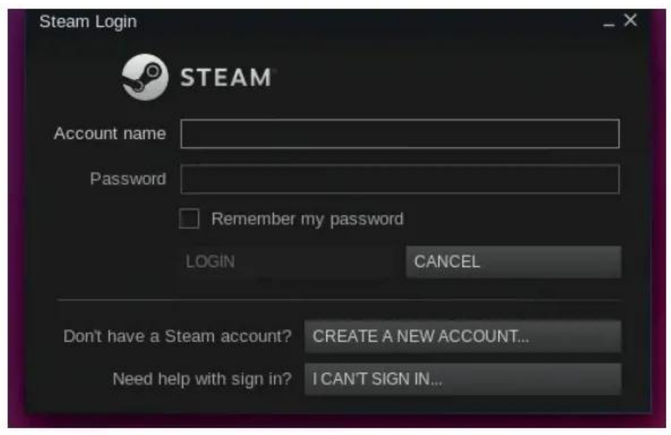

› You need to sign up for a free Steam account to log into the service.

Game services

Now that your Linux machine is ready to play games the big question is where can you get games to play? Steam is the obvious choice as it has a huge collection with a decent mix of genres and game types. It also features some nifty features such cross buy, so you can buy a Linux game, and if it's compatible with Windows or Mac, you can play it on those OSes as well without having to pay for another copy. Many games on Steam also support cloud saves for you saving your progress on one machine and switching to another and carry on playing where you left off.

To install Steam in Ubuntu you can either search for it in the *Software Center* or open up the terminal and type `sudo apt-get update` and then `sudo apt-get install steam` to install Steam. Once installed you will need to sign in with – or sign up for – a Steam account. After that buying, downloading and installing games through Steam is a piece of cake.

Good Old Games is another great service that's now supporting Linux. Although it's not quite as straightforward as buying through Steam, its GoG Galaxy application takes plenty of cues from the Steam client (but isn't currently available for Linux, you have to use the launcher), but it does come with its own pros, such as it only sells DRM-free games and it guarantees games to run on supported platforms or your money back. Also, as you can probably guess from the name, it started out specialising in old games, repackaging them so that they work on new hardware and OSes. Of course, it also features new games as well, and you can view and download games from **www.gog.com**.

We'd normally recommend taking a look at Desura, too (**www.desura.com**), which has a host of Linux indie games – both free and paid-for – but Desura's parent company filed for bankruptcy in June, and aside from Humble Bundle offering its assistance (**www.humblebundle.com**, and another great source of Linux games), there's been no official word on its future yet. The site is still up so the choice to buy is your own.

Regardless of the game service you decide to use, you should now have a Linux machine that's fully ready to play the latest and greatest games. ■

Devices

Installing drivers 80
Setting up printers and scanners 82
Managing Linux drives 86
Using removable media 88
Understanding the Linux filesystem 90
Partitioning drives 92
Managing users 96

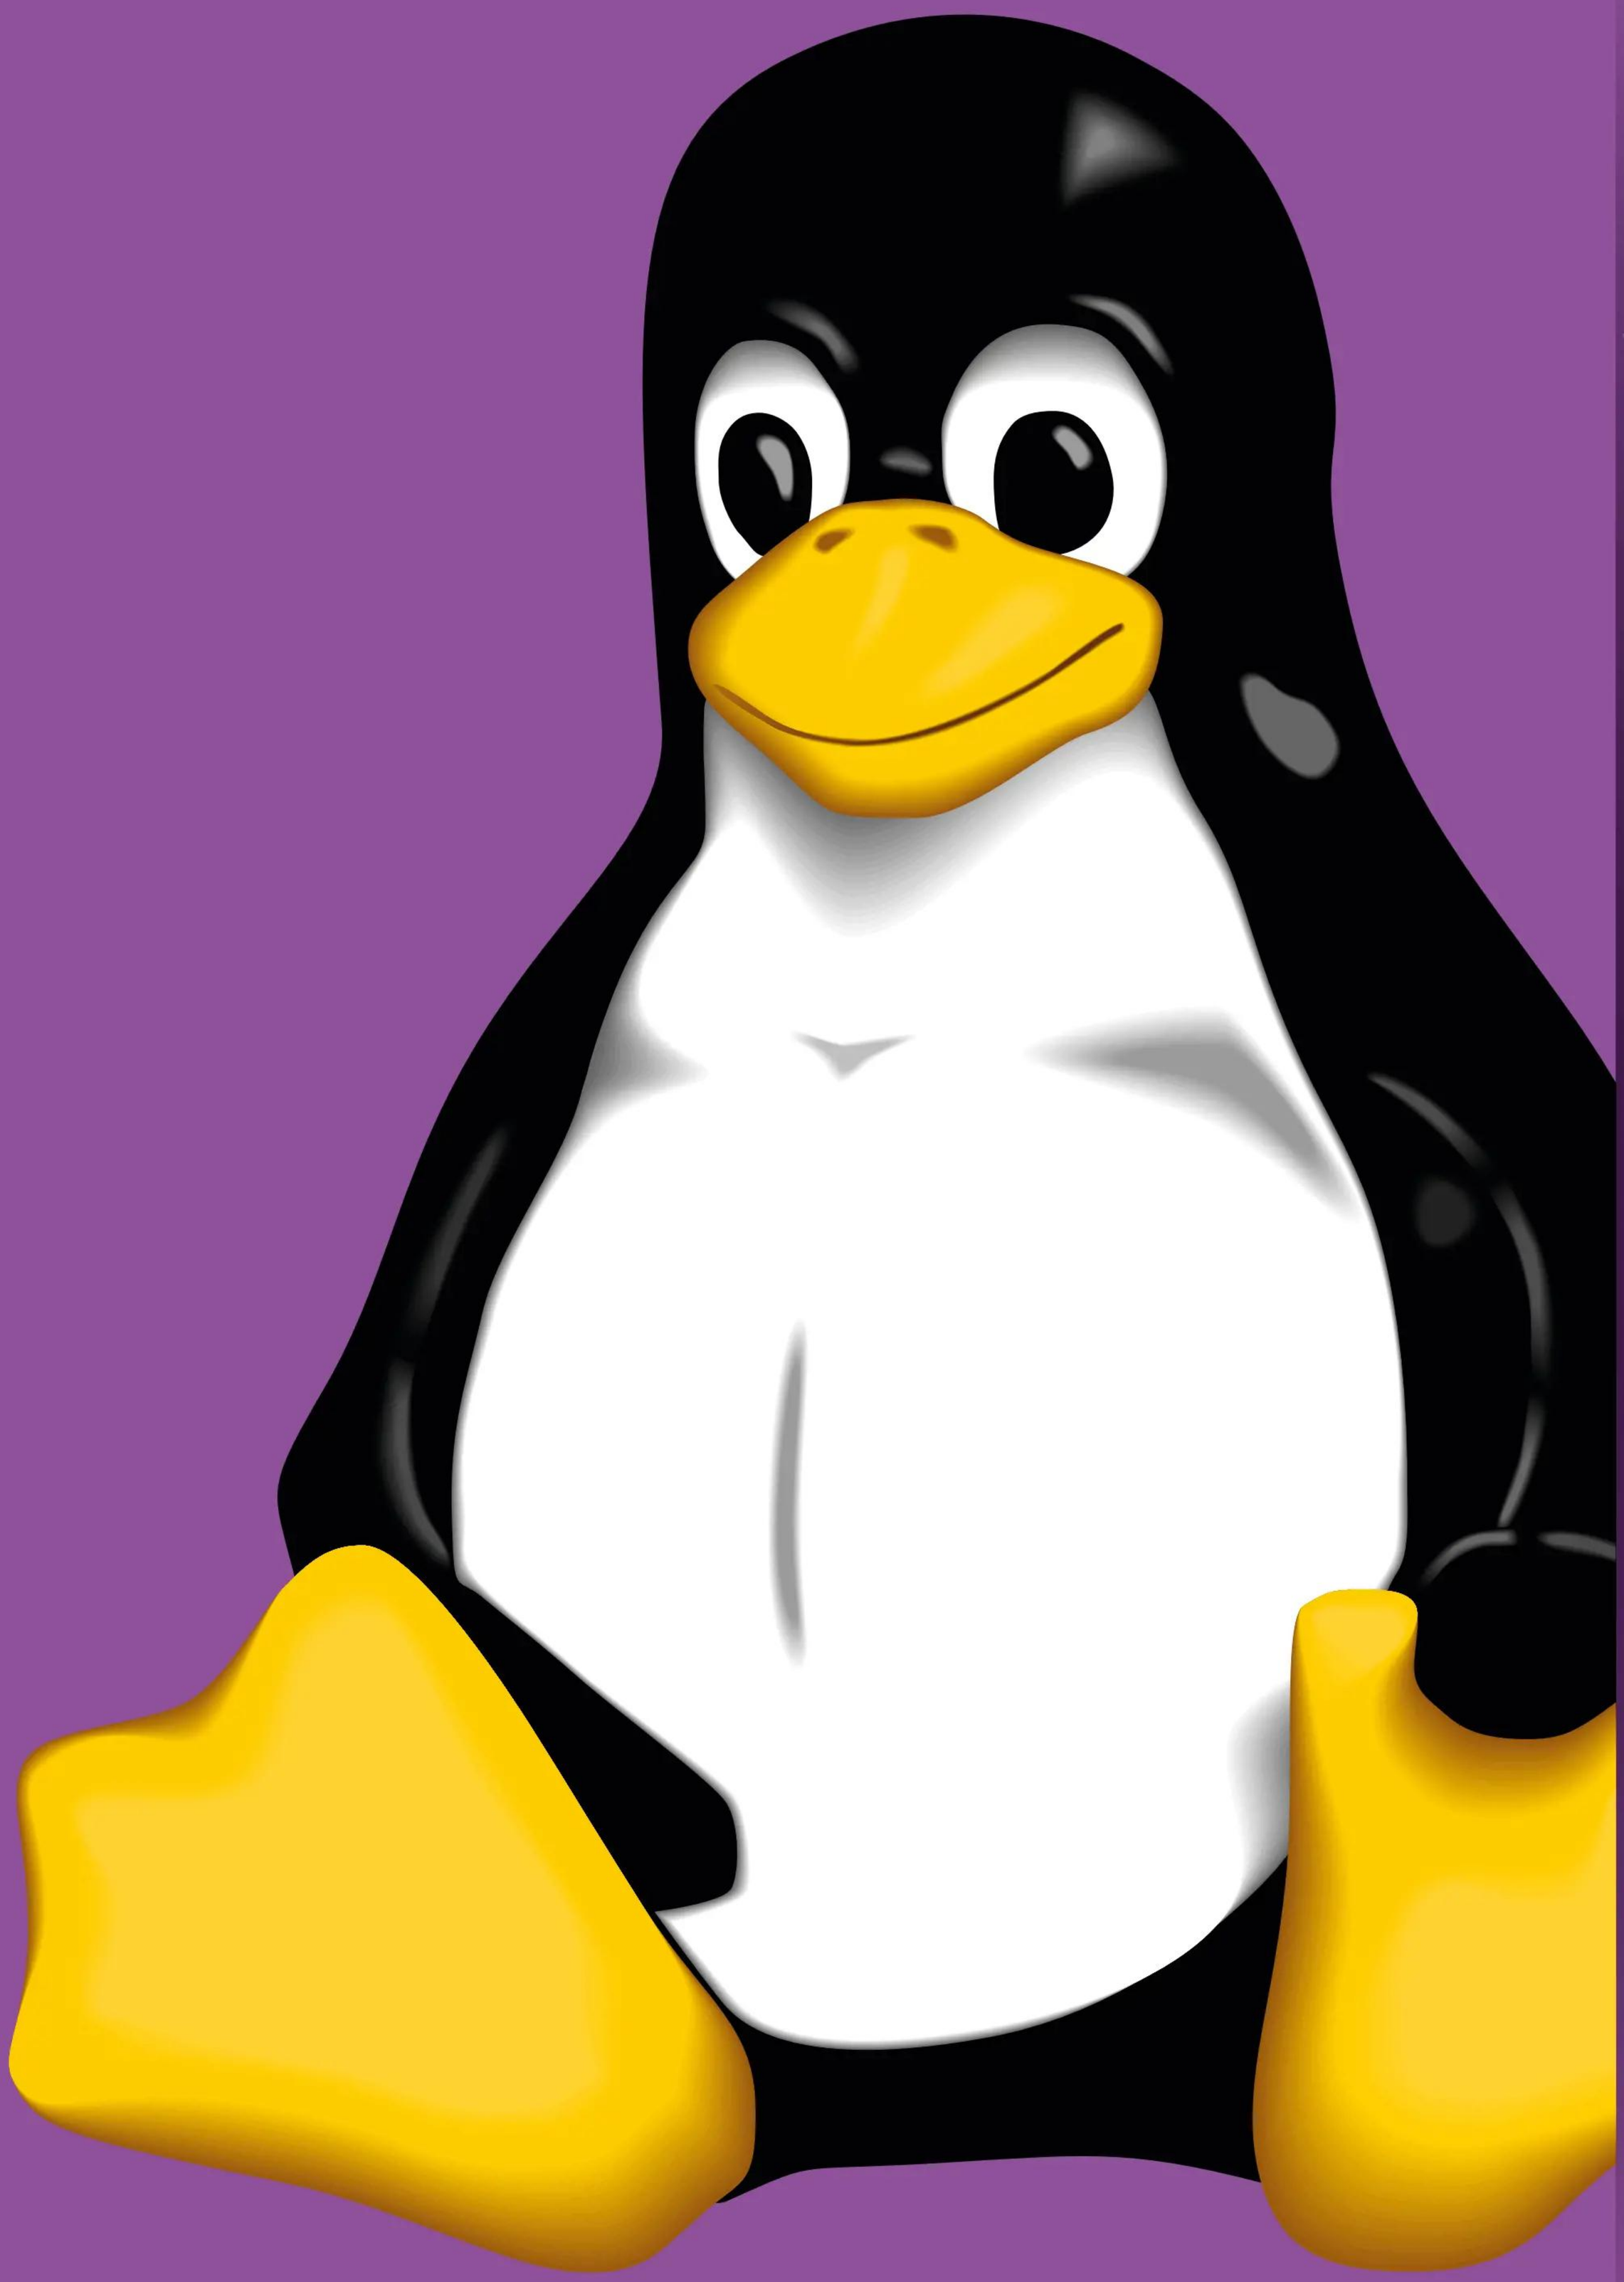

Hardware: Graphics card drivers

Let's take a long, hard look at the advantages and disadvantages of switching graphics card drivers.

Quick tip

To find out more about your graphics chipset – including its capabilities – install *mesa-utils* from the *Software Centre*. Once done, open a Terminal window and type the following: **glxinfo | grep render**.

One of the most confusing aspects of switching from Windows to Linux is understanding how drivers – and in particular graphics card drivers – work. When you install Ubuntu on your PC, you'll notice it appears to install the drivers you need for graphics to work perfectly, and indeed this is the case in the vast majority of cases. These drivers are open source, designed to provide the key functionality for any graphics card to work with Ubuntu.

Because they're designed to just work, many of these drivers aren't interested in providing you with cutting edge performance. For day to day use – browsing the internet, word processing and even editing images – they're perfectly adequate, but if you have a modern graphics card, and want to take full advantage of its capabilities – particularly for gaming – you'll need to investigate the possibility of replacing them with dedicated drivers provided by the manufacturers.

This second set of drivers – with the exception of drivers for Intel graphics – are closed source, which means that while they're still free to install, you're relying on the manufacturer for support. The open-source drivers that ship with Ubuntu are updated through Ubuntu's own update tool, so they'll update organically with the rest of your system.

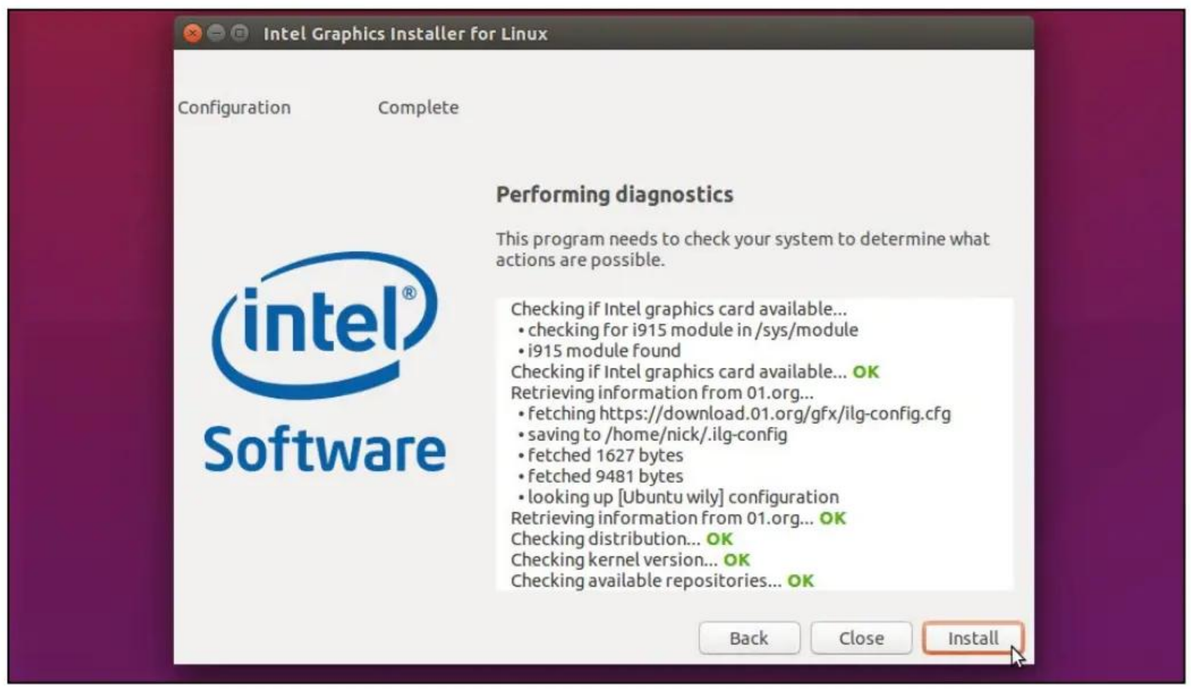

› **Intel's graphics drivers are open source, unlike those provided by NVIDIA and AMD, so need to be installed manually.**

Closed source – or proprietary – drivers rely on the manufacturer for updates to be delivered. The manufacturer may provide this functionality within the drivers themselves, or you may have to manually check the website to see if new drivers have been released.

What's available

The major three graphics card manufacturers – Intel, NVIDIA and AMD/ATI – all provide their own Ubuntu drivers. If you plan to upgrade to the closed-source graphics drive, the first thing to do is to identify your exact make and model of graphics card – particularly if you have no idea who the manufacturer is. If you head over to System Settings and select Details, you'll see a graphics section, but it won't tell you much. Instead, open a Terminal window and type:

```
lspci | grep vga
```

This will quickly list both manufacturer and model, which you'll need in order to determine if a proprietary driver is available for your card.

NVIDIA and AMD

If you have an NVIDIA or AMD graphics card, then the quickest way to see if a proprietary driver is available is by opening System Settings, selecting 'Software & Updates' and switching to the 'Additional Drivers' tab. If your card is compatible, you should see – after a short delay – a list of alternative drivers appear. Choose the latest version marked 'proprietary, tested' from the list and click 'Apply Changes'. When the process completes, reboot your PC.

Don't be surprised if your card isn't detected – if the card is more than five or six years old, then it's unlikely to be supported. AMD's support goes back to 2010 and the HD 5000 series, while NVIDIA cards date back to around the same time with GeForce 400 series. At first glance, you'll see that NVIDIA does still provide links to legacy drivers for older cards – see **www.nvidia.com/object/unix.html** – but these won't work with Ubuntu 15.10.

If your card is an older model, therefore, you're stuck with the open-source drivers – Nouveau for NVIDIA cards, and Radeon for AMD cards. Don't be too downhearted, both

Exploring the NVIDIA X Server interface

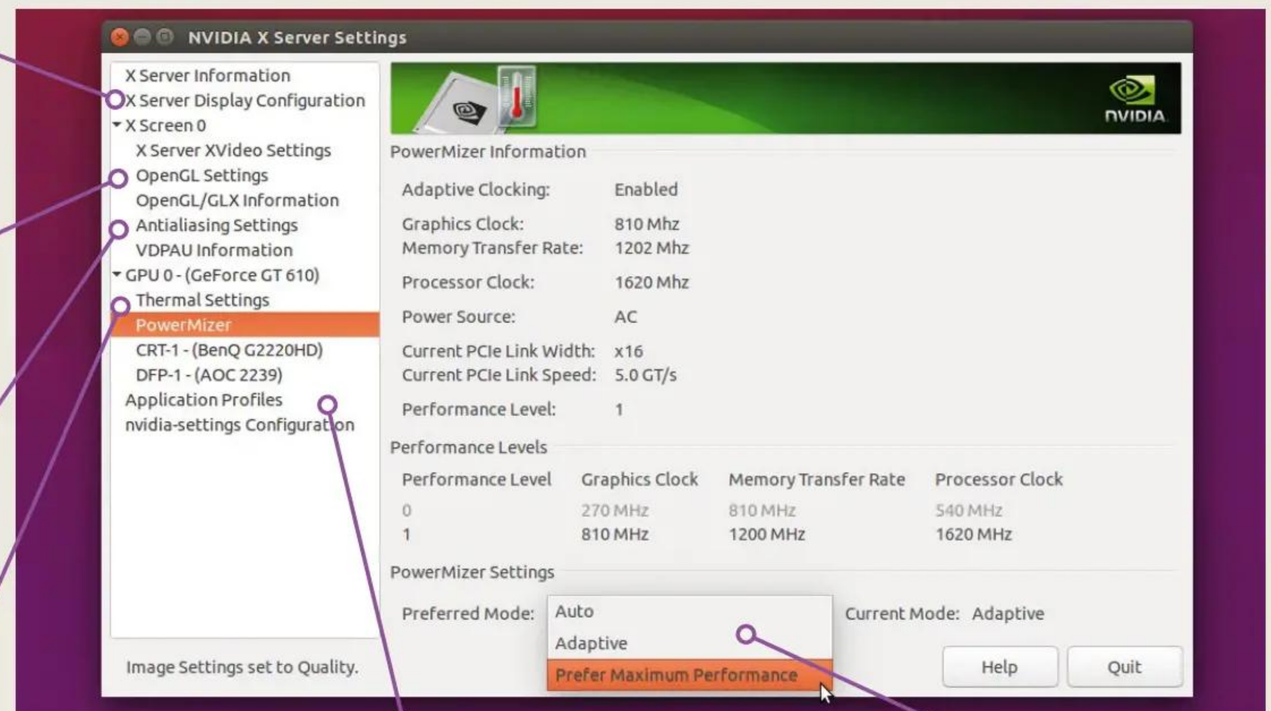

Display Configuration
Select this to tweak the resolution settings for each monitor attached to your PC – click a screen to select that monitor.

OpenGL settings
Access the 'Image Settings' slider from here, where you can try and strike a different balance between image quality and performance.

Antialiasing settings
This section allows you to configure the card's ability to smooth jagged edges and textures when rendering 3D objects using OpenGL.

Thermal Settings
This provides a temperature gauge for your graphics card – keep it open while you stress-test the card if you're worried about heat problems.

Application Profiles
Advanced users can fine-tune the various settings on offer for specific applications – typically games. Use the Help button to guide you.

PowerMizer
Some cards support different performance states – use this tab to view current settings and see if you can push the envelope if required.

provide more than basic functionality, with pretty decent 2D and 3D acceleration. They also support most older chipsets stretching back more than a decade. There's no need to install them either – if your card is correctly detected, the driver should already be in place. Visit **https://nouveau.freedesktop.org/wiki/** and **https://help.ubuntu.com/community/RadeonDriver** respectively for full details.

If you're lucky enough to be able to use the proprietary driver, then not only do you gain performance enhancements for gaming and other high-end purposes, you'll also be able to tweak your card using either the Catalyst Control Centre (AMD) or NVIDIA X Server utility – see the annotated screenshot for details of using the latter.

Intel chipsets

If your PC has an Intel processor and utilises its onboard graphics chip, you may be able to install the open-source Intel driver, which is found at **https://01.org/linuxgraphics**. At time of writing, the latest available driver is 1.4.0, which is designed for Ubuntu 15.10 and supports the following Intel graphics chips: 2nd-6th generation Intel Core processors, Bay Trail, Braswell and Hoxton Intel Celeron processors and the Intel Atom N450 (Pine Trail) chipset.

If your processor fits the bill, click the link under 'Latest download', then choose the 32-bit or 64-bit version depending on your Ubuntu build – if you don't know this, open up System Settings and click Details, then check the 'OS type'. Click the link to save the .deb file to your hard drive (or opt to open it in *Software Centre* if prompted). If necessary, double-click the file and click Install, then follow the prompts to install the utility. Once done, open the Dash and type 'Intel' to locate the graphics installer. Again, follow the prompts and if your chipset is correctly identified, click Install to add the driver.

Reboot when prompted, and you're done. There's no separate graphical configuration utility – instead, use Screen Display under System Settings to tweak your card's settings.

Troubleshooting

Sadly, not all driver updates are successful, and you can find yourself locked out of the Ubuntu desktop if one goes awry. If this happens with NVIDIA drivers, press [Ctrl] + [Alt] + [F1] to drop to the shell. Log in using your username and password, then enter the following command:

```
sudo ubuntu-drivers devices
```

This identifies which graphics drivers are installed. Assuming it reveals NVIDIA drivers, enter the following command next:

```
sudo apt-get autoremove –purge nvidia-*
```

Follow the prompts to purge the drivers from your system. Now add the following command:

```
supt apt-get install xserver-xorg-video-nouveau
```

Once done (don't worry if it tells you it's already present), type sudo reboot and you should hopefully regain access to your desktop.

AMD users should try the following lines instead:

```
sudo apt-get autoremove –purge fglrx-*
supt apt-get install xserver-xorg-video-ati
```

While they're the most likely solutions, we're sorry to say that these may not work – if this is the case, you'll need to Google your problem to try some of the other suggestions on offer. ■

Quick tip

To avoid any worry when considering installing new drivers, take a drive image following the before you install the drivers, allowing you to roll back should the worst happen.

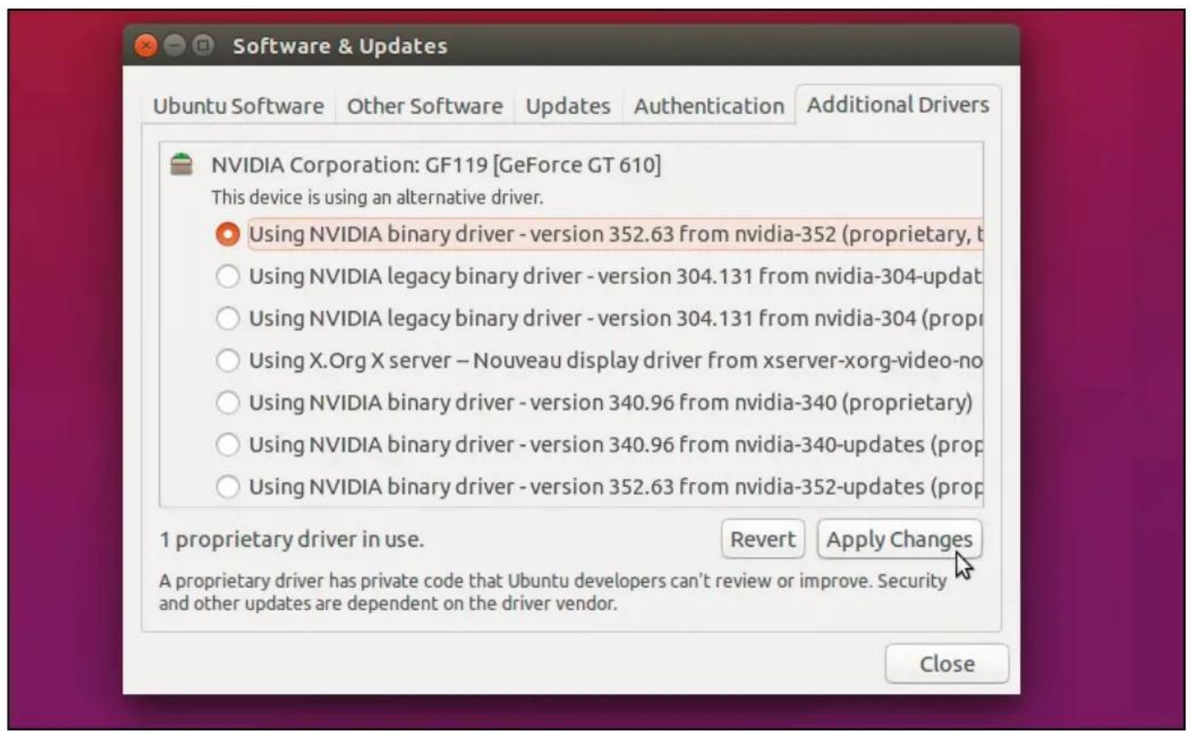

› Switching to proprietary graphics drivers is actually quite simple – assuming your card is supported, of course.

Printers: Get set up

Let's run through some tips to get your printers and scanners up and running using systems and tools such as CUPS, SANE and XSANE.

These days, it's almost mandatory to have a printer and scanner attached to your PC, but getting them set up in Linux can be a little tricky. Don't worry, though, because armed with the right drivers and tools you can quickly get things up and running. There are two core tools built into most Linux distributions (distros) that help simplify the process: *CUPS* for printers and *SANE* for scanners. *CUPS* provides a centralised management tool for adding and configuring your printers while *SANE* allows you to interface with your scanner in much the same way, as well as providing the tools you need to preview and scan images. We'll be focusing on these for this tutorial, coupled with some useful tools that mean you don't need to go into the terminal should you need to access or tweak them.

With other operating systems, you can usually plug and play your printer or scanner – the drivers will either already be pre-installed or you'll be prompted to download a driver, or insert a manufacturer's disc. While it's not quite that simple in Linux, the fact is that armed with the right procedure you shouldn't have to spend hours getting things set up.

The good news is that you don't need to start your search at your manufacturer's website when installing your printer. If you're lucky, you may find the drivers are already in place, particularly if your model is a network one, in which case all you need to do is open the *CUPS* front-end and see if it can detect your printer. *CUPS* should be pre-installed with later versions of Ubuntu – including 14.04.3 LTS – and you can access it from System Settings by clicking 'Printers'. If it's not there, install it through the *Software Centre* or via the terminal using `sudo apt-get install cups`.

› Network printers can often be set up even without drivers thanks to CUPS.

Once installed, click the 'Add' button to see if your printer is visible and supported. If your printer is connected via USB but not showing up – as is likely – then you'll first need to locate suitable drivers for it. Start by looking for Linux drivers on your manufacturer's website and for your particular distro (you may need to find a unified package). Hopefully your exact distro, including its version number, will be supported, but don't panic if it's not – we successfully installed a Dell B1265dnf all-in-one printer in Ubuntu 14.04.3 LTS using a unified package that only promised support up to 11.04.

If you're lucky, the driver may be supplied in a Deb package, in which case downloading and double-clicking it should launch the setup wizard. If your driver is provided as a tarball (with a .tar.gz extension), then it'll probably contain a shell script (.sh) installer, in which case follow the step-by-step guide opposite to discover how to get it up and running. Once done, re-open the *CUPS* front-end under System Settings and your new printer should be visible.

Network setup

If you're adding a networked printer then you should be able to spot it without having to first source drivers. Click 'Add', expand 'Network Printer' and wait or click 'Find Network Printer'. After a short pause, any detected printers should show up. Select yours and you'll see a choice of connection protocols appear. (*See What Connection Type? box, p84 for details of which one to choose if you're given the option.*) Select IPP if it's available, otherwise choose LPD or visit your manufacturer's website to see if it provides a driver that will likely offer an IPP connection if you install it.

If you're happy to proceed with the options offered by *CUPS*, click 'Forward' and it'll look for drivers, selecting them automatically if they're available. If it can't find any, you'll be given three choices: select from database, provide a PostScript Printer Description (PPD) file or search for a driver to download. Start by selecting your printer manufacturer to see if an exact match is available. If not, try a search – again, you may want to expand your search to your web browser, as it's limited in scope.

You can also use a PPD file if your printer is a PostScript model and you can source one (search for your printer name

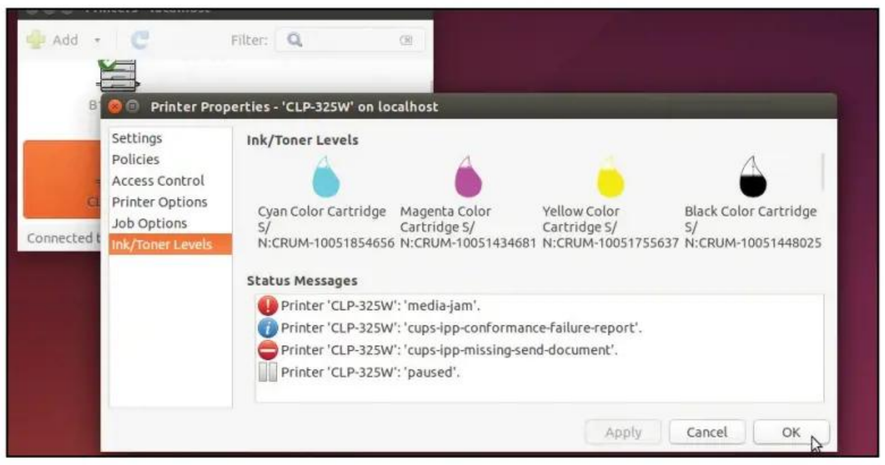

› Once installed, administer your printers using the Printers tool under System Settings.

and model and the word 'PPD'). This provides *CUPS* with a description of its capabilities, which it can use when outputting files.

If all else fails, don't worry: choose 'Select printer from database' again, but this time select 'Generic (recommended)' and click 'Forward'. A generic PCL driver will be pre-selected, so try this first: either the pxlmono or pxlcolor option depending on whether your printer is mono or colour. Click 'Forward' to install it.

Whichever route you go down when adding your printer, you'll be prompted to provide a name, description and location, then asked to print a test page – we recommend this to verify the connection is working. Once done, you'll be taken to the printer's configuration screen, where you can examine its settings and see what control you can wield over it.

Tweak printer settings

The main Settings tab allows you to change your printer's description and location, plus change its URL. If you click 'Change...' next to 'Device URI' you can attempt to connect using a different protocol. This is also where you'll find options for printing test pages and maintaining the printer.

The Policies tab allows you to change the printer's state in terms of whether it's enabled, accepting jobs or shared. You can also define what happens if the printer encounters an error and switch on authentication for its 'Operation Policy' if supported and required (if you go down this route, access the printer via the *CUPS* web interface to set things up).

Access Control allows you to deny (or allow) access to the printer on a user level. You should spend some time in both Printer Options and Job Options tab where you configure key printer defaults. Printer Options is mainly concerned with default quality and paper settings, while Job Options lets you configure things like the number of copies, page orientation, and individual elements such as text and images.

Finally, if your printer allows it, the Ink/Toner Levels provides a handy glance at your printer's core components, including elements like the fuser.

Now your printer is set up, it's time to use it. Open an application and when you're ready, choose File > Print. Your printer should be listed, so it's simply a case of tweaking whatever options are made available to you in the app's own dialog and clicking 'Print'. If you have more than one printer installed, the default one will always be pre-selected; change this via the System Settings > Printers dialog by right-clicking your target printer and choosing 'Set as default'.

In most cases, the *CUPS* tool provided under System Settings should have all the options you need, but *CUPS* can also be administered from your browser: type `localhost:631` to access it. There's a handy overview and links to the *CUPS* forum should you need it, but use the tabs at the top of the page to navigate the web-based utility. Go to Printers and click your printer to access its settings – here, navigation is done through a series of drop-down menus. If you've installed a dedicated driver and were asked to assign a user to it, you'll be prompted at this point to enter the username and password to access the required pages. »

Quick tip

Struggling to get your printer to function correctly? Visit **http://openprinting.org/printers** and see if your model is covered. If it is, scroll down to the comments to see if anyone else can help (or post your own question in the forums).

Install a printer driver tarball package

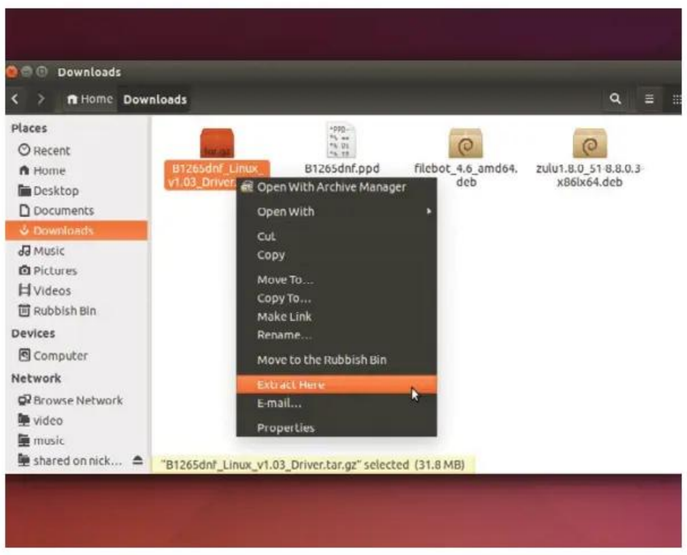

1 Obtain & extract installer

Browse your manufacturer's website and locate the required driver for your Linux distro, which may be packaged as a unified driver. Save the gar.tz file to your **Downloads** folder, then right-click it and choose 'Extract Here'. Once that's done, open the extracted folder and locate a file called **install.sh**. Right-click it and choose 'Properties' to make a note of its location.

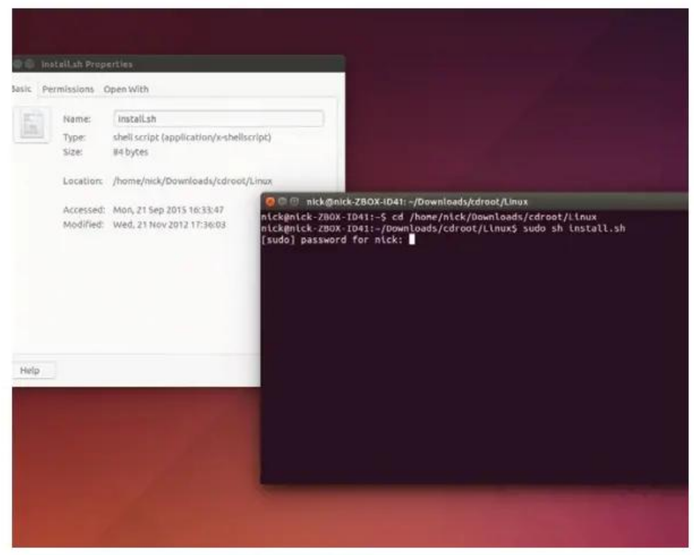

2 Start installation

Now you'll need to open up a terminal window and type in `cd <path>`. You'll need to replace <path> with the path listed under Location for the **install.sh** file that's found in its Properties window. Now you'll need to type the following command and press Enter, followed by your root password when prompted: `sudo sh install.sh`. After a short pause you should see an installer wizard pop up.

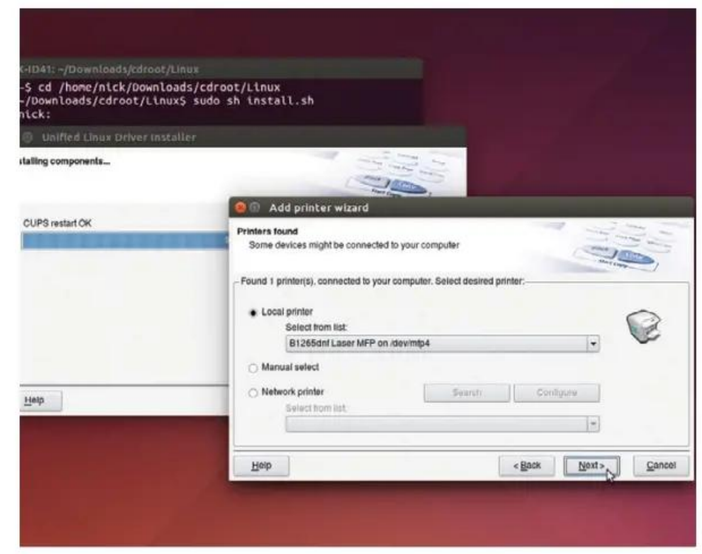

3 Follow install guide

You'll need to work your way through the installer wizard and make sure that you tick your username to allow you to access the printer. Once the wizard's complete, you may then be prompted to configure the printer for CUPS and need to follow its wizard to set it up, entering name, description and location, if required, before going ahead and testing the printer to finish the process off.

› **Some printers ship with their own utilities, which may provide you with more control over its settings.**

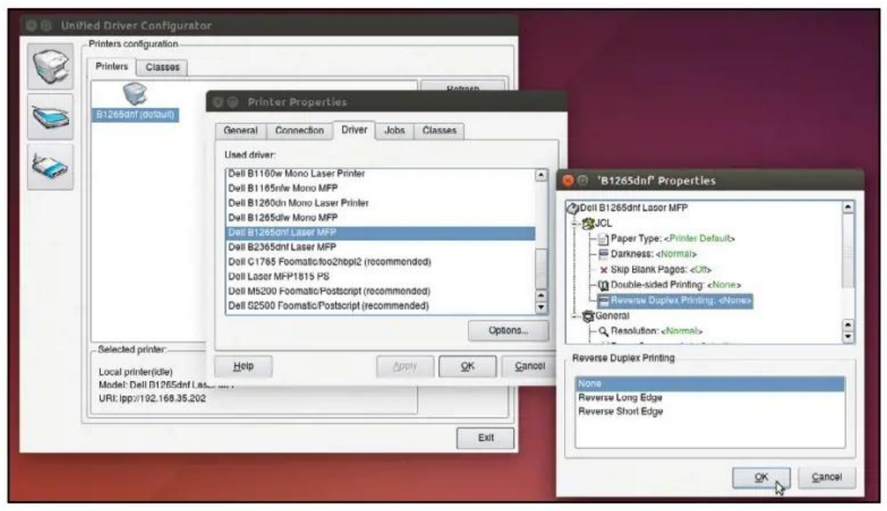

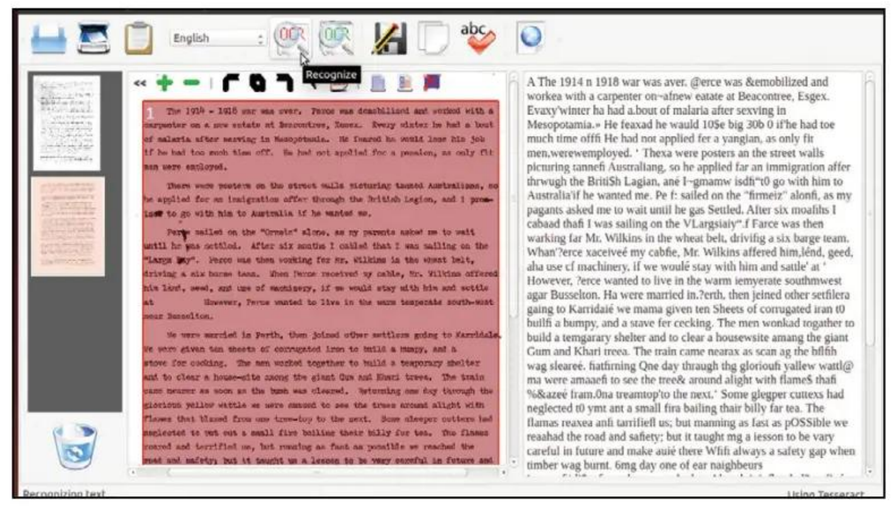

› **The success – or failure – of OCR is largely down to the quality of the original document.**

» If you've set up a network-connected printer, you should also be able to access machine-level settings directly – type in its IP address. If you don't know its IP address, it's under the printer's Settings – click 'Change' next to 'Device URI' and expand Network Printer. The IP address will be listed in brackets after your printer name. Enter the URL address and you'll access the printer itself, giving you even greater scope over how it functions (eg setting up duplex options for double-sided printing).

Quick tip

If you're trying to extract text from a multi-columned document, make sure you use YAGF's tools for correcting page skew and selecting multiple columns before you attempt recognition to maximise your chances of having a successful extraction.

Printer classes

CUPS allows you to group multiple printers into classes. This is mainly aimed at networks where lots of users want to share a limited number of printers. Each class has a hierarchy with those printers added first being preferred to those added later. When people come to print using a class, they're allocated a printer based on availability, so if one person is printing they'll always get access to the first printer, while if two are printing at the same time, the first person gets the preferred printer and the second gets the next available one.

With this in mind, it pays to work out how you want your printers classed: do you want faster ones made available first, or are you looking to steer people towards a high-volume printer such as a laser? Once done, open System Settings > Printers and click the '˘' button next to Add and select 'Class'.

You'll be prompted to provide a name, description and location in terms of a single printer, but it's better to think in terms of the class, eg, 'study-printers' or something similar. Once done, click 'Forward' and you'll be prompted to add existing printers to the class – remember to add them in order of preference. Once done, click 'Apply'. Once created, you can set this class as your default in place of a single printer – just right-click its entry and choose 'Set As Default'.

Linux also offers good built-in tools to help you access and make use of your scanner. In many cases – particularly with newer models – you can simply plug it in and there's a good chance it'll work. That's due in part to built-in support for SANE (Scanner Access Now Easy), which provides all the back-ends to scanning in Linux.

To see if you're lucky, plug in your scanner to a spare USB port and fire up *Simple Scan* to see if it detects it. If it does, you're in business – skip to the next section; if it doesn't – don't panic. First, check for Linux drivers with your manufacturer. Follow the instructions to download and install

What connection type?

If you're hooking up a network printer, then you'll have been given a choice of connections when setting your printer up. But what are the different protocols and which one should you choose if given the choice?

The simplest protocol is AppSocket (also referred to as HP JetDirect). It's usually used when no other connection is possible, but its simplicity also means it lacks any kind of fine control: eg there's no queue management so you can't abort print jobs once they've been sent. There's also no way to find out the printer's status, making it impossible to discover ink and toner levels, eg.

Next up is LPD, which stands for Line Printer Daemon (also known as LPR, or Line Printer Remote). This does implement queue management, and can tell the difference between different types of data, such as raw text or PostScript information. There's neither support for remote setup nor management, though, and you can't query the ink or toner (although these may be provided anyway) and your print jobs can't be encrypted.

The final protocol – IPP – is the most contemporary and functional. It stands for 'Internet Printing Protocol' and is what CUPS uses as the basis for managing both print jobs and queues. It's the preferred method of communication as you get more feedback from the printer and wield more control over it thanks to its bidirectional nature.

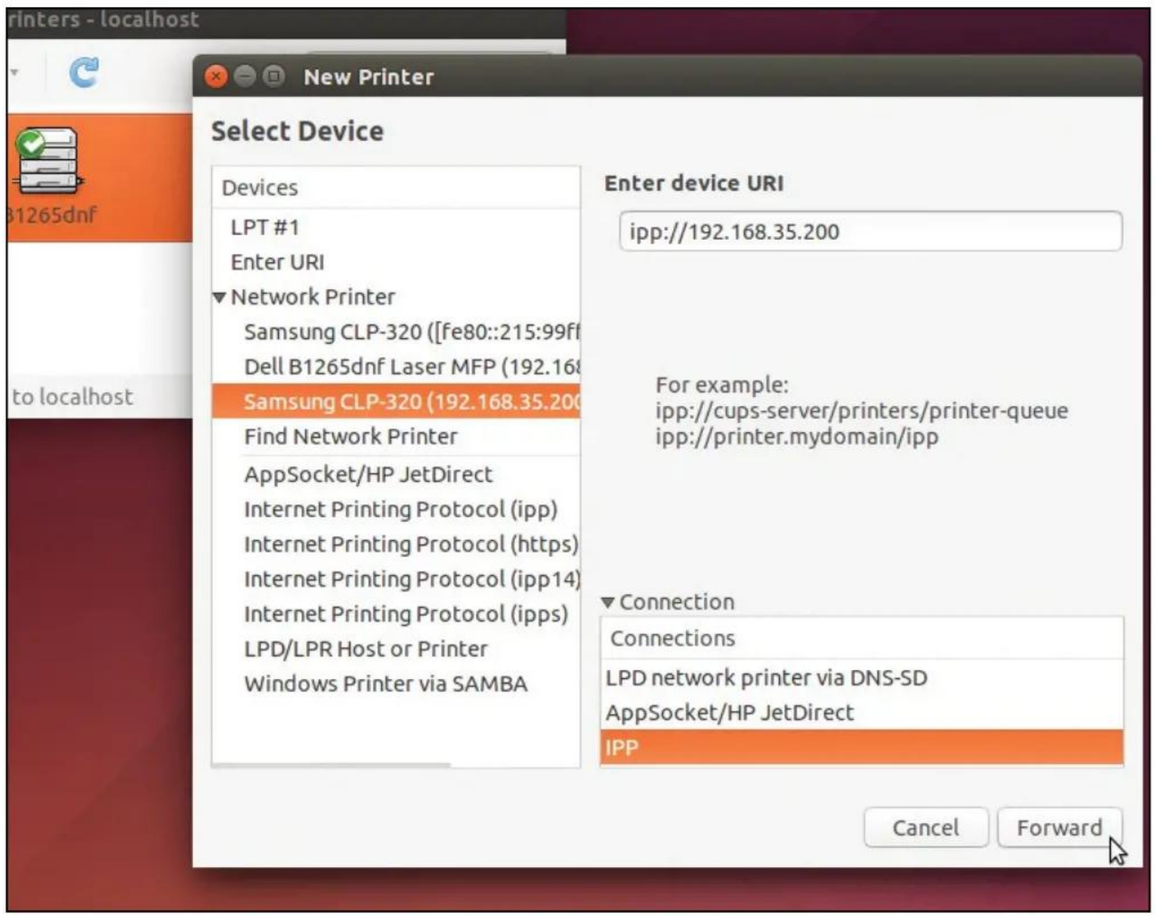

› **Select your protocol carefully when configuring your network printer using CUPS, we'd suggest that IPP is the best if it's supported.**

Get unsupported scanners working

If, after your best efforts, you're still unable to get your scanner set up using SANE, then take a trip to **www.hamrick.com** where you'll find a program called *VueScan*, which promises to support over 2,500 scanners.

Click the supported scanners link to browse by manufacturer – if your model is listed, then download the tarball and extract the **VueScan** folder, then double-click *vuescan* inside to launch the program proper. It should detect your scanner automatically, then provide you with a simple step-by-step wizard to preview and scan using the program.

VueScan provides a lot of useful tools to help you get the perfect scan, and it works with many scanners that are already supported in Linux too, so if you're struggling to get to grips with *XSane*, consider giving it a whirl. The only downside is that it's not free – the Standard Edition costs $39.95, while the more fully featured Pro Edition costs $99.95 (you can upgrade from Standard to Pro for the difference at any time, too). The good news is that *VueScan* is fully functional for the purposes of evaluation (other than a watermark across all your scans), so you can be sure it's right for you before making a purchase.

these – again, if they're provided with an **install.sh** script, follow the step-by-step guide (*see p71*) to get them installed.

If your scanner is part of a multi-function device than its drivers should be integrated with the printer's – if you added the printer component of a MFP without installing specific drivers, then you'll need to source and install these to get the scanner up and running. If you've installed drivers previously, but the scanner isn't detected, try installing them again over the top of the original drivers – we found this worked with our Dell B1265dnf after the scanner failed to show up. Once installed, your drivers may provide their own utility to help you verify the scanner is now recognised – eg Dell users should see a DELL Unified Driver Configurator shortcut on the desktop – or you'll need to fire up *Simple Scan* again at which point your scanner should now be visible.

For more troubleshooting advice on using SANE to detect your scanner, see **https://help.ubuntu.com/community/Scanners** or take a look at the box on *VueScan*.

As we've seen, Ubuntu comes with the *Simple Scan* front-end for SANE, which provides you with a simple, no-frills interface for your scanner. It's good for confirming your scanner works, but it lacks key tools like a preview function or the ability to fine-tune your scanner's settings.

Time to scan

However, if you really want to push the boat out, open the *Ubuntu Software Center* and search for 'xsane' to locate and install *XSane*, a more powerful scanning front-end. Once launched it'll display four windows – start by clicking 'Acquire preview' to create a preview of the image you're trying to scan. Once done, use the Standard Options window to select your image – choose 'Custom' under Page Format and then crop in using the four measurements beneath it.

Next, click the 1:1 magnifying glass button in the main window to zoom into the preview. Use the eyedropper tools to select white, black and grey parts of the image for colour correction purposes, you can then use the histogram's sliders in the top left-hand window to fine tune the colour balance. The right-hand window is where you choose where to save the file, what format to use and so on – choose 16 million colours for photos and ensure the scan resolution is increased from 75dpi to 300dpi or 600dpi if you want a hi-res scan. There are also controls for adjusting brightness, contrast and gamma.

By default, your scan will open in *XSane's Viewer* window when you click 'Scan' – it's a good idea to perform one scan like this, so you can review your image and perform further optimisation. Once you're happy with your changes, click 'Viewer' and change it to Save, then click 'Scan' one last time.

Another benefit of *XSane* is that it integrates with *Gimp* – this allows you to scan directly to it, and then fine-tune your image accordingly. Launch *Gimp*, choose File > Create > *XSane* to select your scanner or perform a search for your available devices. Everything's handled identically here except for one critical part: there's no output option, so once you click 'Scan' your image will automatically be transferred to open in *Gimp*.

XSane also links in with the *goocr* tool to a built-in Optical Character Recognition (OCR) engine. This allows you to scan in printed documents and convert them into editable text. It's a flexible solution, but if you want to try something else then install both cuneiform and *YAGF* from the *Ubuntu Software Center*. Cuneiform is an alternative OCR engine while *YAGF* provides a neat front-end to both cuneiform and tesseract and is an OCR engine that's renowned for its accuracy. *YAGF* integrates nicely with *XSane* too – launch *YAGF*, click the 'Scan' button and it'll open *XSane*. Set the colour to 'Black and White – Line Art' and set the dpi to 600. Click 'Scan' and the document will be scanned and sent back to *YAGF*. From here, review the scan quality, then use the controls above the scanned image to prepare it for recognition. Once done, click the red 'OCR' button next to the language drop-down menu. *YAGF* will use Tesseract by default to decode the page, with the editable text placed in the right-hand pane for you to review and correct. Don't like the results? Choose Settings > OCR Settings, select cuneiform and click OK, then try that instead – in our tests, Tesseract performed better, but not by much. Be prepared to try and fine-tune the original scan if necessary to make the text as clear and legible as possible.

Once you're done reviewing and editing the text, you can save it as a TXT file or copy it to the clipboard, ready for pasting into Writer or another editing program.

Congratulations, your scanner is now set up to perform both standard and text-based scans. ■

Quick tip

It's also possible to use CUPS to 'print' documents to PDF file where no native PDF option is provided. Install **cups-pdf** from the *Ubuntu Software Center* – once done, a new PDF printer will appear as an option when printing.

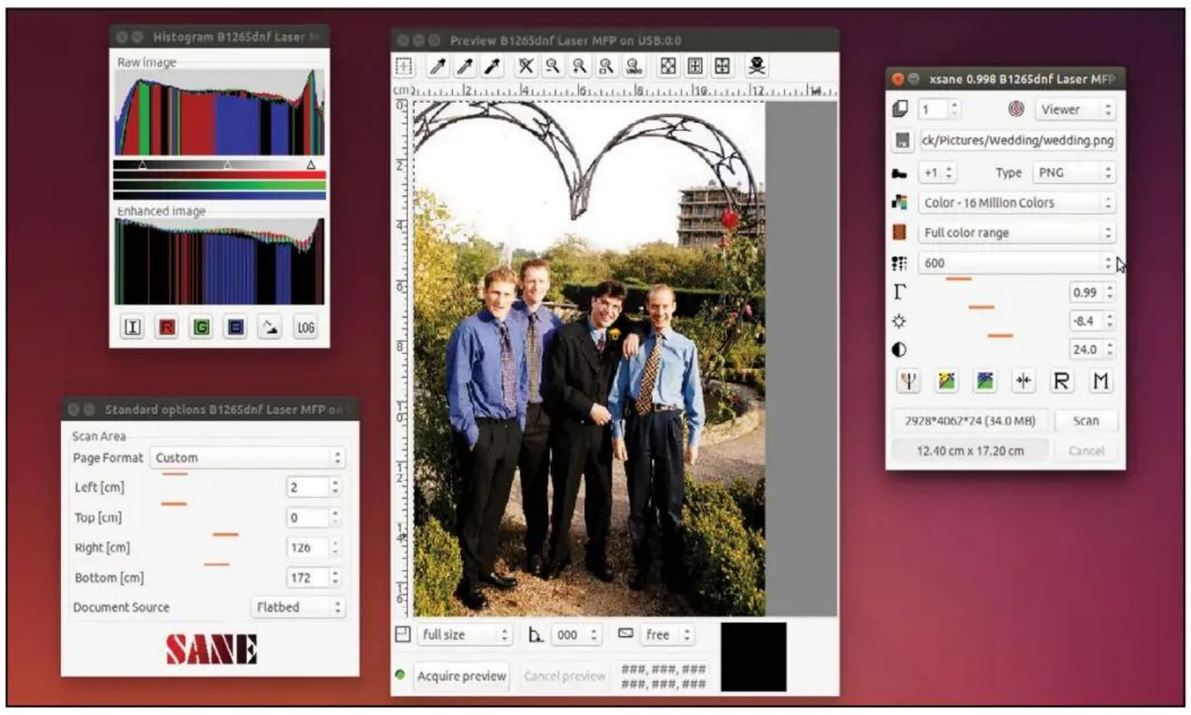

› ***XSane*** **provides a comprehensive set of scanning tools.**

Drives: Identify and locate

Where are those files? We look into how you can identify, access and locate the hard drives attached to your PC.

Quick tip

If you regularly switch USB drives around, they won't always be assigned the same identifier – so make sure you use the Disks utility to verify what's been assigned to each drive before.

If you're migrating from Windows, you may find the way Linux handles drives and partitions is bit of a culture shock. For starters, there are no drive letters involved – instead, you'll see lots of references to **/dev/sda1, /dev/sdc** and so on. It may look bewildering, but it doesn't take long to realise the way Linux treats drives is just as logical as Windows – more so in some cases. Read on to discover how to demystify the jargon and understand exactly how drives work within the Linux file system.

Identify your drives

Open the Disks utility from the Dash – use the annotation opposite to navigate it, making a note of how drives are referenced in Ubuntu. They appear to relate to a folder path beginning **/dev/** – and that's because they do just that. At the top (or root) of the Linux file system you'll find the dev folder, which stands for devices. The dev folder houses special files that Linux uses to identify the different pieces of hardware in your system.

Select each drive in turn in Disks and you'll notice the obvious pattern: they all refer to files 'sdx', with 'x' a letter that's unique to that drive – for example, **/dev/sda**. The '**sd**' part identifies the fact this is a supported storage device, and the final letter identifies which drive it is.

Linux is reasonably consistent in how it registers each device: 'sda' is always given to the drive on which Ubuntu is installed, for example. Then any other internal hard drives are registered before moving on to external drives, which include both hard drives and flash storage (more on the latter – as well as optical drives – over the page).

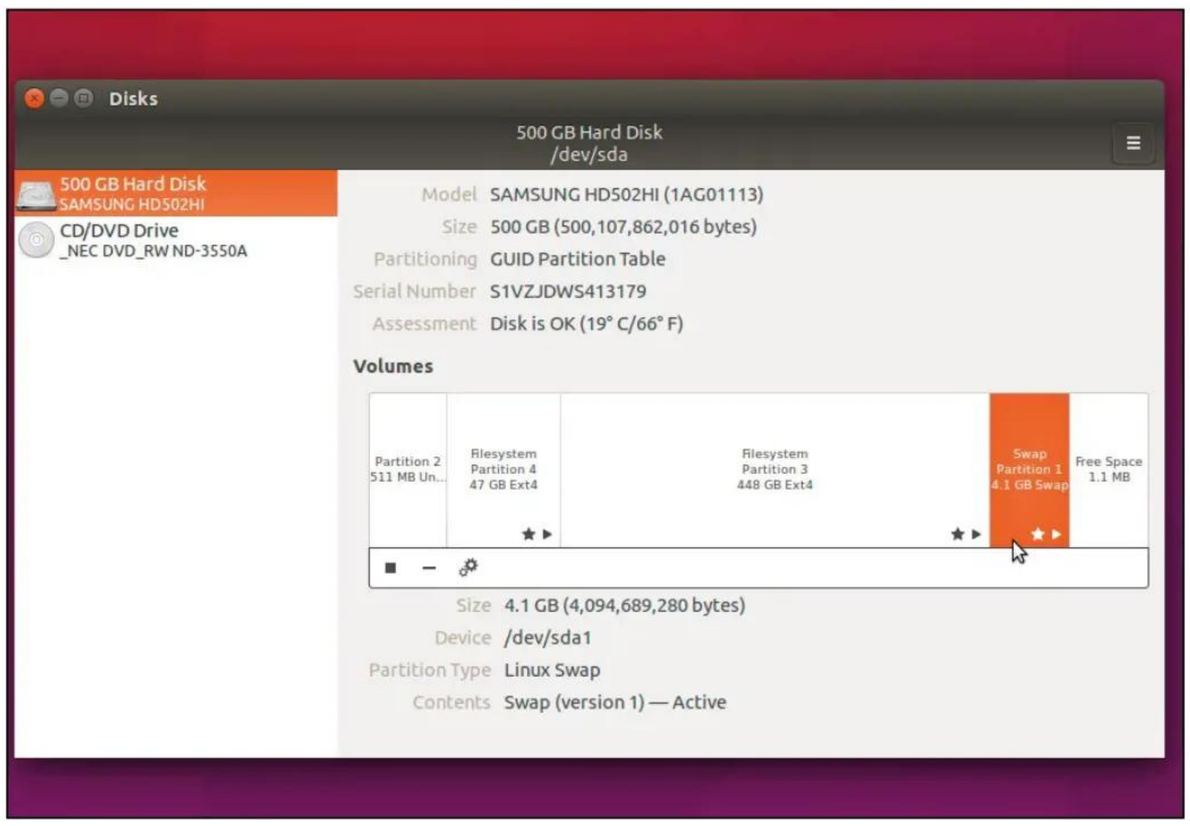

› Drives aren't always registered in the physical order in which they appear on your drive.

Partitioning

So far, so good – each drive is easy enough to identify. But many drives contain partitions – select drive sda and you should see a minimum of two partitions are displayed in the bar chart that provides a graphical representation of the disk's contents.

Partitioning is the act of splitting a single hard disk into multiple partitions, referred to as 'volumes'. If you've come from Windows, you're probably used to storing everything – programs, data and Windows itself – on a single volume, but in fact modern versions of Windows contain at least three partitions, with hidden recovery and boot partitions sitting at the beginning of the drive.

Windows identifies volumes by assigning them drive letters. Linux adopts a different approach, identifying each volume using a combination of the drive identifier (sdx), plus a unique number, so 'sda1' refers to the first partition on your first drive, while 'sdb3' refers to the third partition on the second drive, and so on.

On the surface this numbering makes perfect logical sense, but it doesn't always look quite so straightforward. Drives are partitioned according to 'schemes', and there are two principal schemes in operation: GPT (GUID Partition Table) are found in newer drives, but MBR (Master Boot Record) remains a popular choice among drives of all ages.

The key difference between them is that the GPT scheme has no meaningful limit to the number of partitions you can create, while MBR drives support a maximum of four partitions per drive. A workaround exists for MBR drives whereby one of these partitions can be an 'extended' partition, inside which you can bypass the four-drive limit by creating multiple 'logical' volumes.

GPT drives number drives according to the order in which they've been created or identified. This normally follows a logical pattern from left to right, but it's not always the case (for example, if you manually create the swap file to sit at the end of your drive before setting up the main system partition). It has no effect on how Linux operates, but can be confusing.

The MBR approach throws up its own mildly interesting anomaly when numbering volumes. Linux reserves 'sdx1' to 'sdx4' for physical partitions, with logical volumes

Exploring the Disks interface

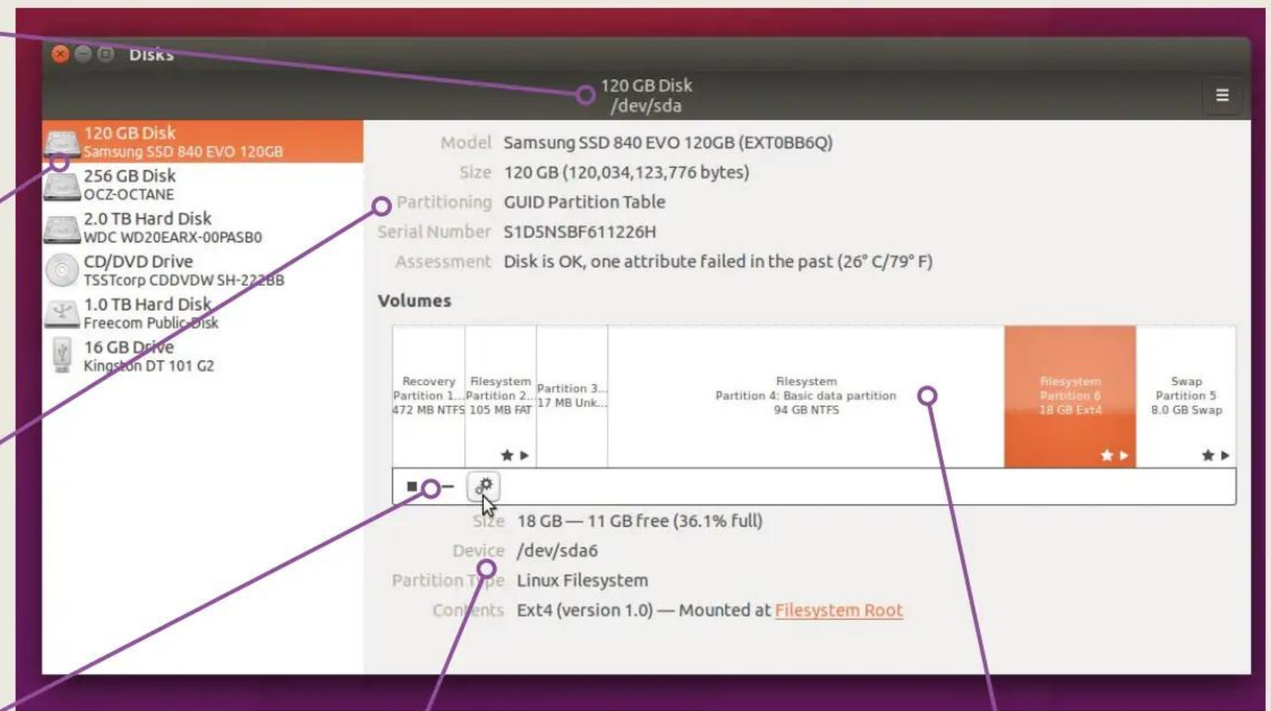

Drive identifier
The select drive's physical identifier (/dev/sda, /dev/sdb and so on) is clearly visible – click the menu button in the right-hand corner for options.

Drive list
All recognised drives are listed here, complete with helpful identify marks such as size and manufacturer. Select one to examine it in more detail.

Partition type
You can find out how the drive has been partitioned by examining the 'Partitioning' section of its summary. It'll be GUID or MBR – see the main text for more information.

Quick backup
Click this button and you'll see additional options, including one that allows you to take a drive image of the selected partition.

Volume info
Click a volume's entry in the bar chart to reveal more details about it, including its all-important identifier and a shortcut to where's it's mounted.

Volume chart
This bar chart provides a handy representation of how the drive is broken down – although it's not to scale when multiple smaller volumes exist.

assigned 'sdx5' onwards. When you install Ubuntu on a single, unpartitioned hard drive, it'll create two partitions: one primary partition for the main system volume (called sda1), and an extended partition inside which the swap file partition is stored (called sda2). You can examine this using the Disks utility and you'll see the swap file volume is allocated 'sda5' inside the extended partition. Again, it has no effect on how Linux operates, but it manages to explain any confusion.

How drives are accessed

We know Windows allocates drive letters to volumes and drives, making them easily visible and accessible. But how do you access these volumes in Linux? If we ignore the fact the Nautilus file manager provides quick and convenient shortcuts to volumes through its navigation pane, where are the drives situated?

At first glance, you might think the quickest way to access a drive or volume in Linux is through the **/dev** folder, but that won't work because the files inside – sda, sda1 and so on – aren't direct links or shortcuts.

Linux handles drives differently to Windows in that they're not treated as separate entities; instead, volumes are integrated into the main file system as part of a larger whole. This is done in Linux systems by 'mounting' each volume to a specific folder rather than allocating it a drive letter. You can then browse to the folder to access and interact with the files and folders stored on the volume in question.

When Linux boots, it mounts the main system partition to the root of the file system (**/**). The swap file partition isn't mounted, but other key folders can be stored on separate volumes too, such as the **/home** folder where all user data and settings are stored .

When you arrive at your desktop, only these key volumes are mounted. USB drives that you subsequently connect while Ubuntu is running will be mounted automatically, but while convenient shortcuts to other detected volumes appear in Nautilus' navigation pane, you'll need to click these to physically mount the drive in question.

The shortcuts are convenient, but where are these drives physically mounted? The good news is that Ubuntu picks a consistent location: open Nautilus and select the Computer shortcut to view the root of your drive. You'll see a 'media' folder – open this and you'll see a folder named after your username. Open this to reveal folders pointing to the relevant drives. Armed with this useful information you can quickly browse to the drive in Terminal using the 'cd' command – for example: **cd /media/nick/Data**.

Third-party access

Now you've learned how Linux identifies drives, plus – crucially – where they're found in your filesystem, you can start to make more sense of this information. It's helpful when repartitioning your drive using the Gparted utility, for example, plus you can now make sense of the semantics required when using the powerful 'dd' utility to back up partitions or drives. ■

Quick tip

Before disconnecting any drive from your PC, ensure you've first ejected it from the filesystem. To do this, open Files, locate the drive in the left-hand pane and click the eject button you see next to it.

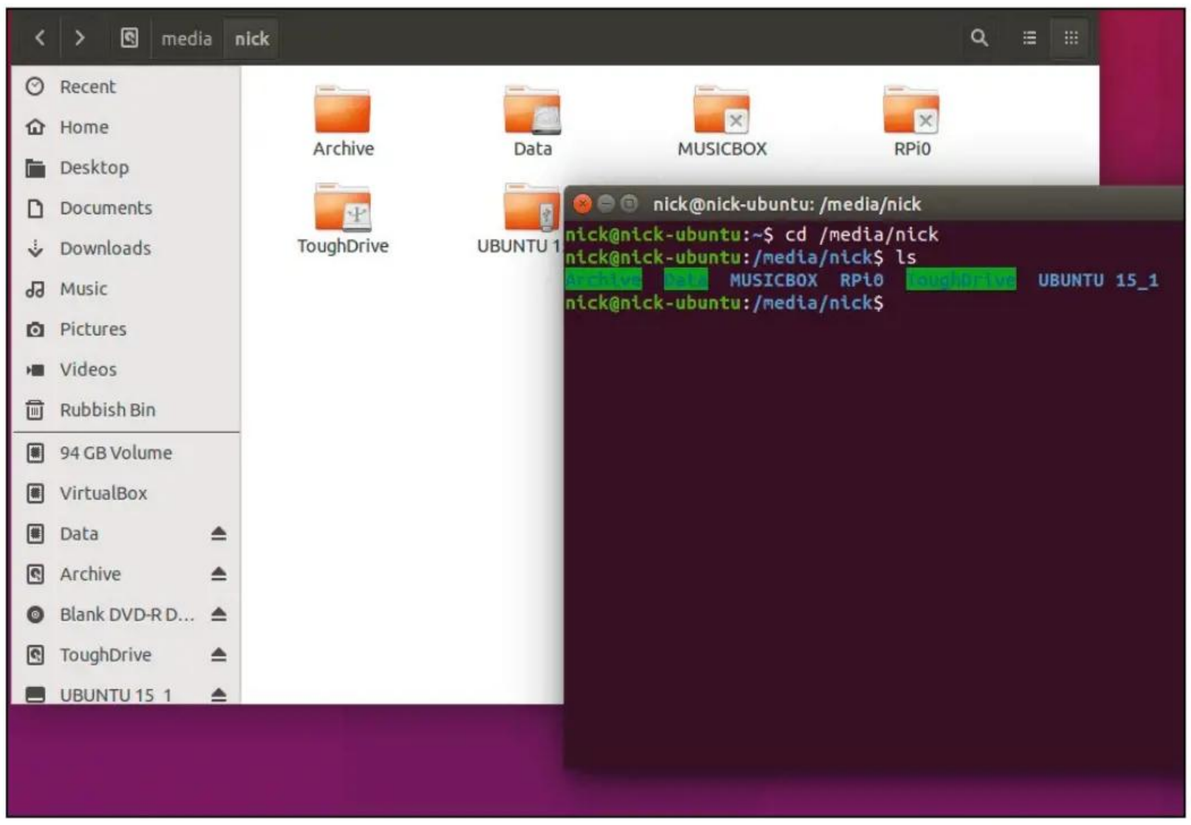

› Any drives you manually mount in Ubuntu will appear in the /media/<username> folder.

Drives: Access flash and optical media

Your Linux install is not an island. Here's how to access both flash storage and optical media using Ubuntu's built-in tools.

Quick tip

Verify what your CD or DVD drive is capable of doing in Ubuntu by opening a Terminal window, typing **wodim -prcap** and then hitting [Enter].

Linux handles drives differently to Windows, so if you've migrated from a Windows PC, we recommend you first read our tutorial on the previous page to ascertain how drives are labelled and accessed in Ubuntu. The good news is that – once you've over the initial shock – the transition isn't that big. And now you've understood how things work, you can go on to learn more about removable media, including USB flash storage as well as CD and DVD.

Working with USB drives

Ubuntu is more than happy to work with USB storage of any nature – hard drives and flash drives are treated equally. Just plug it in, and the drive will be automatically mounted and accessible, so long as it's formatted using a supported file system. Ubuntu ships with built-in support for both FAT32 and NTFS file systems as well as its own ext-based file system – once mounted, you can read and write to the drive like any other.

Larger drives – typically 64GB or bigger – are occasionally formatted in the exFAT format. Support for exFAT isn't included in Ubuntu by default, but you can quickly add full read/write support by opening the *Software Centre* and searching for 'exfat'. Install 'exfat-fuse' followed by 'exfat-utils' if you're running Ubuntu 64-bit, or choose the i386 labelled versions if you're running the 32-bit version. If you don't know which version you're running, go to System Settings > Details and examine the 'OS type' entry to find out. Once installed, full read/write support to any exFAT drive will be enabled.

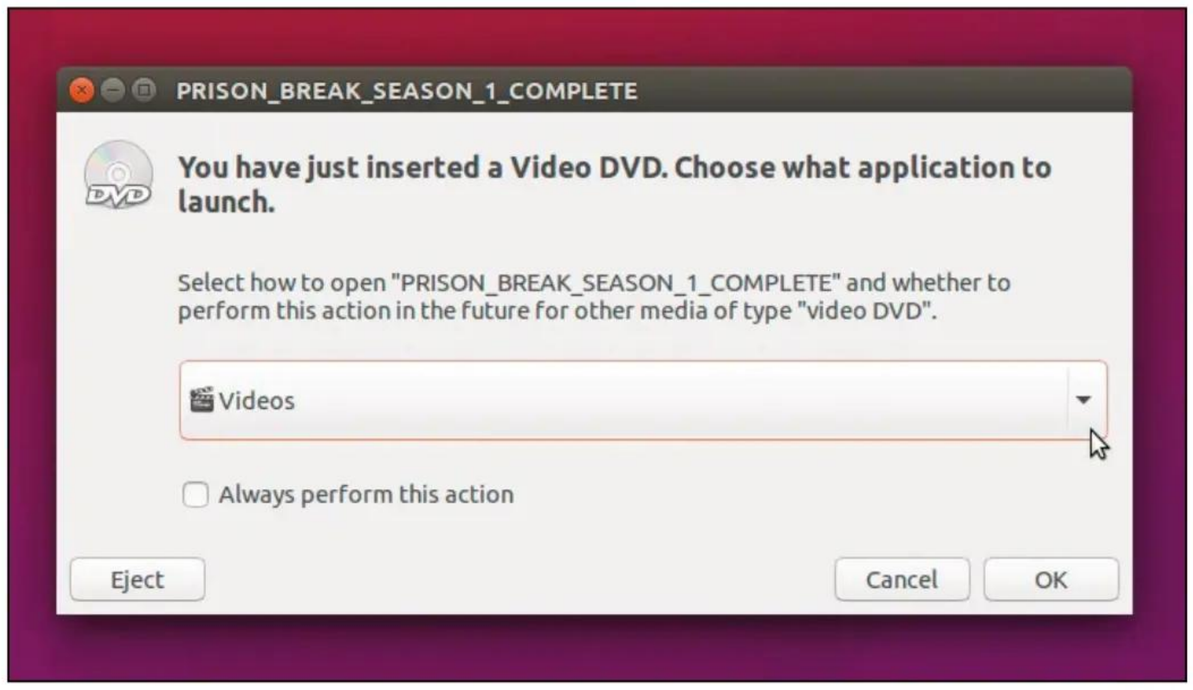

› **Pop a media or blank disc into your DVD drive and Ubuntu will pop up this helpful dialogue box.**

USB drives are registered in the same way as hard drives – open Disks to confirm the individual label for your drive. Select its volume and you'll see it's something like **/dev/sdb1** or **/dev/sde1** depending on how many other drives are present. Unlike permanently attached drives, these identifiers aren't constant. They're applied on a first-come, first-served basis, so if you were to unplug one drive, then plug it back in having first attached another drive, its label will change. Always verify the label in Disks before performing any tasks that require you to identify or address it by its identifier.

Once plugged in, your drive should show up in the Nautilus file manager – its actual mount point can be found under **/media/<username>/<volumelabel>**, where **<volumelabel>** is the name you gave it when the partition was created. Unlike the drive's identifier, this does stay the same.

When you've finished with the drive and want to remove it, click the eject button next to the drive name in Nautilus to unmount it before unplugging it. If you need to remove a drive from the Terminal, use the **umount** command. By default, the umount command is used to remove individual volumes, but by including the **-a** flag you can ensure all volumes for a particular drive are unmounted, allowing you to remove the drive safely even if it's been partitioned. Point umount towards a specific volume or its mountpoint using the following syntax:

```
umount -a /dev/sde1
umount -a /media/nick/RPi0
```

Once done, remove or eject the drive.

Read-only mount

Sometimes you may wish to access a USB drive in read-only mode, so data can't accidentally be written to the drive. By default, Ubuntu will mount any supported filesystem in read/write mode. The trick is to use the umount command with a special flag (**-r**) to unmount the selected volume, then

remount it in read-only mode – for example:

```
umount -r /media/nick/RPi0
```

If you find you can't unmount a drive using the **-a** command, you can use the **-r** flag instead on each mounted volume from that drive to remount them in read-only mode. This means you can then physically remove the drive without worrying about data loss.

Optical drives

CD and DVD drives are handled differently to regular hard drives. First, they're assigned different IDs - **/dev/sr0, /dev/sr1** and so on. You'll also see they're marked as read-only by default, even if the drive itself is a burner.

How you interact with discs depends on the disc itself. If the disc is recognised, it'll be mounted inside the **/media/<username>** folder like other removable drives, making it accessible through the Terminal as well as via Nautilus. If the disc is a straightforward data disc, it should open in Nautilus, enabling you to browse and access its contents.

Other types – multimedia and writable discs – are handled slightly different. A pop-up window will appear, similar to that which appears when you insert a disc in Windows. You'll be asked if you want to open a specific application, which varies on the type of disc you've inserted, plus there are options for opening the drive in Files to explore its contents, or do nothing.

Your app choice depends on what's installed – Ubuntu ships with four apps that can be used from the off: *Rhythmbox* (playing and ripping audio CDs), *Videos* (watching DVD movies), *CD/DVD Creator* (burning data discs) and *Brasero* (burning all kinds of discs, including DVD, SVCD and audio CD).

By default, Ubuntu can't play certain restricted formats, including encrypted DVDs. To play commercial DVDs, install the *libdvd-pkg* software from the *Software Centre* – once done, you'll find *Videos* is able to play them, as will other third-party tools, including *VLC Media Player*.

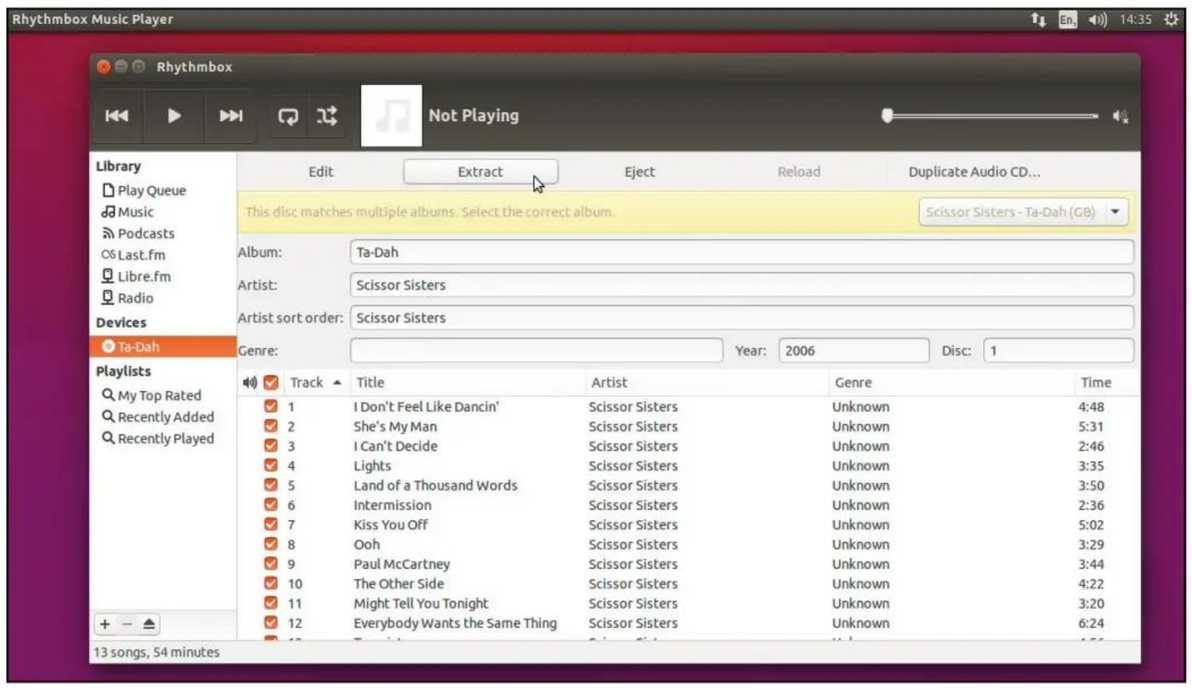

› **Rhythmbox can extract audio from your CDs as well as play them back to you.**

Burn discs

Insert a blank disc and you'll be prompted to open CD/DVD Creator. This provides a similar interface to Nautilus – drag files onto the window, give your disc a suitable title and click 'Write to Disc' to burn a data disc.

A more sophisticated tool for burning discs comes in the shape of *Brasero*. Launch this when prompted (or open it from the Dash) with a blank disc inserted into your drive, then select your project: you can create an audio CD, data disc, or video DVD/SVCD, plus there are tools for copying one disc to another as well as burning discs from support image formats (ISO, TOC and CUE). There's even a built-in cover editor (select Tools > Cover Editor) for designing and printing a cover for a regular CD jewel case, plus you can erase rewritable discs by selecting Tools > Blank...

Brasero is simple to use as the annotation below reveals – it's mostly just a case of dragging and dropping the files you wish to add to your disc in the correct order. There are also plenty of alternative burning tools out there – to create a video DVD with menus and try *DVDStyler*, for example, or check out *K3b* for a more general-purpose alternative. ■

Quick tip

You can also attach MMC/SD cards to Linux – if these are plugged directly into a card slot (as opposed to using a USB adapter), they'll be identified using **mmcblkx** rather than **sdx**. They work in the same way as other media.

Exploring Brasero

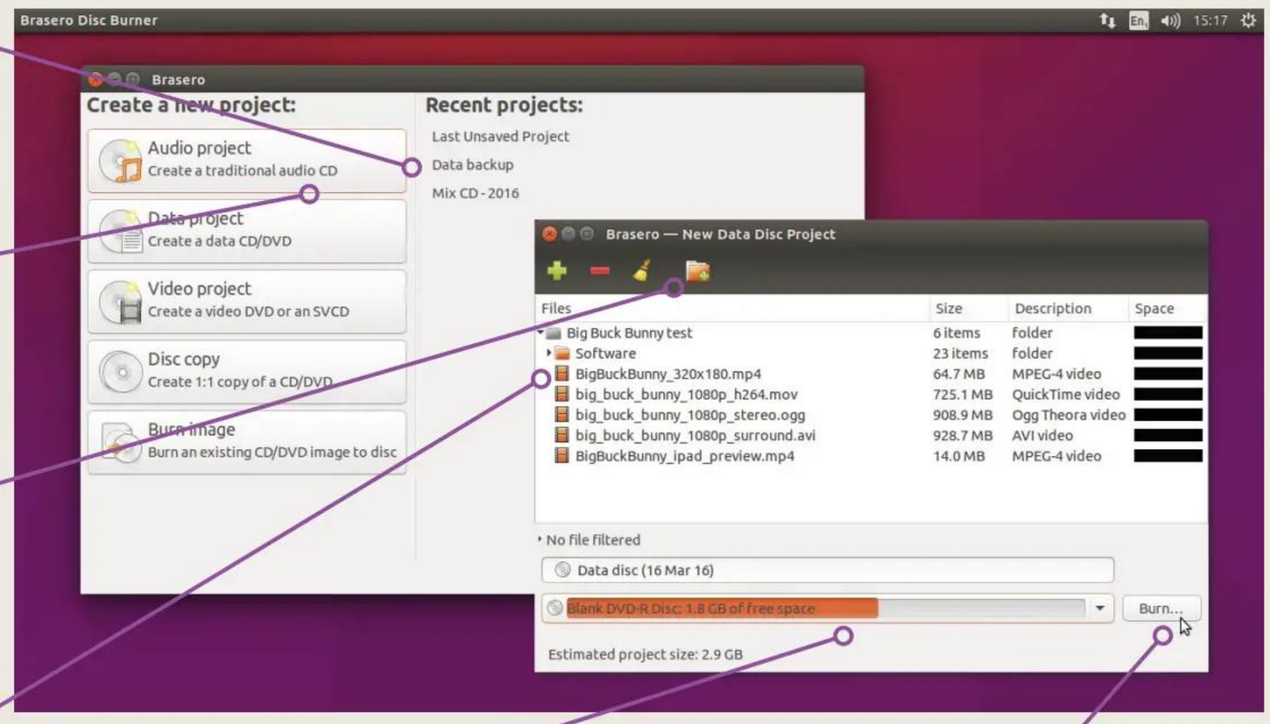

Access older projects

Choose File > Save to make a copy of your project before you're ready to burn it. You can then return to it for further tweaking.

Project type

Brasero allows you to burn all kinds of discs: audio, video and data, plus burn discs from image files and copy from one disc to another (two optical drives required).

File and folder controls

These buttons allow you to add and remove files to your project, plus perform other context-sensitive tweaks such as creating new folders for data discs.

File list

Review the current list of files, tracks or videos that you plan to burn – media tracks can be clicked and dragged to change their running order.

Choose media

Click this dropdown menu to switch between available drives or make a disc image. A handy guide tells you how much free space is left.

Burn disc

Once you're ready, click the Burn button. Set your speed, plus tweak options like leaving the disc open. You can even burn several copies if required.

Files: Navigation made easy

Puzzled about structure? Find out how the Linux file system works, plus the best way to browse it using Ubuntu's built-in file manager.

Quick tip

Don't be surprised to find lots of hidden files – even inside your home directory – after you've selected View > Show Hidden Files. Don't worry, this is normal, and the best thing to do is select the option again to hide them.

We've seen how Linux identifies the drives attached to your PC, and you've also discovered that Linux treats everything as a file – including folders (which is basically a file listing the files that make up that folder's contents). But how are these files organised on your hard drive, and how do you access them through Ubuntu's graphical file manager, Nautilus?

Let's start with the raw basics. Linux organises files in a tree-like structure that's a little bit similar to Windows, the major difference being that while Windows physically separates drives into individual trees, with the drive letter at the top, Linux lumps everything together in a single file tree, with a top-level root (**/**) directory and everything else – including drives – placed relative to that directory.

This can sound confusing, but in actual fact Linux provides a consistent, organised view of all the storage at your disposal. Once you understand where things are kept, it's relatively straightforward finding them again. And, in actual fact, for day-to-day use, you'll limit yourself to your personal user folder, which is always found inside the **/home** directory and is where your personal documents, photos, settings and other data is stored.

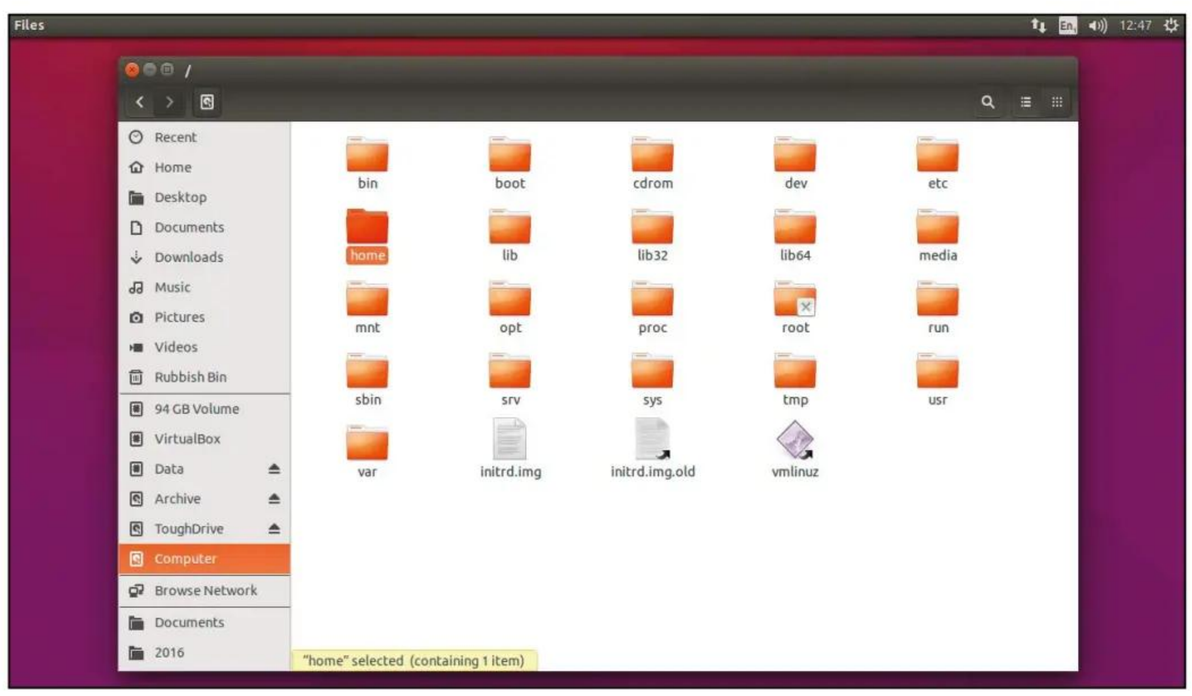

› You won't need to spend much – if any – time navigating the top level directory of Ubuntu's file system.

What makes things even easier to grasp is that Ubuntu ships with a user-friendly file manager in the form of Nautilus, which is accessible via the Files shortcut in the launcher. Nautilus will feel instantly familiar to anyone who's used Windows own file manager, Windows Explorer, and shares many of the same characteristics.

Files and folders are represented by icons, and double-clicking one opens it. You can right-click for more options, while files can be deleted to the Recycle Bin or removed immediately (select the file, hold [Shift] and press [Delete]) in exactly the same way. The annotation opposite reveals some of the key fundamentals in using Nautilus to access and find files within the Linux file system.

What lies beneath

Let's use Nautilus now to identify the key folders on your system, and find out what they're for. Open it using the Files shortcut on the launcher and by default it'll open to your personal home folder – **/home/<username>** – in the Linux file tree. Here you'll find familiar looking folders that mirror those found in other operating systems: they're all self-explanatory, aside perhaps from Public – a folder inside which you can store files that are accessible to others when they're logged on with their own user profile – and Examples. This is simply a shortcut to a folder containing the 'Ubuntu Free Culture Showcase', which aims to provide free wallpaper, music and video.

Next, click the Computer shortcut in the left-hand navigation pane and you'll be taken to the root directory where you'll see a large number of folders. Not all of these need explaining, but of those that do, the **/bin** folder is where common programs shared by all users (including root and the system itself) are installed – this includes all Terminal programs. Startup files – including the Linux kernel are held under **/boot**, while we've already seen how **/dev** contains references to your system's hardware including all attached storage devices.

The **/etc** folder contains many configuration files, while the various **/lib** folders contain program configuration data – similar to what's held in the Windows Registry. External

Exploring the Nautilus interface

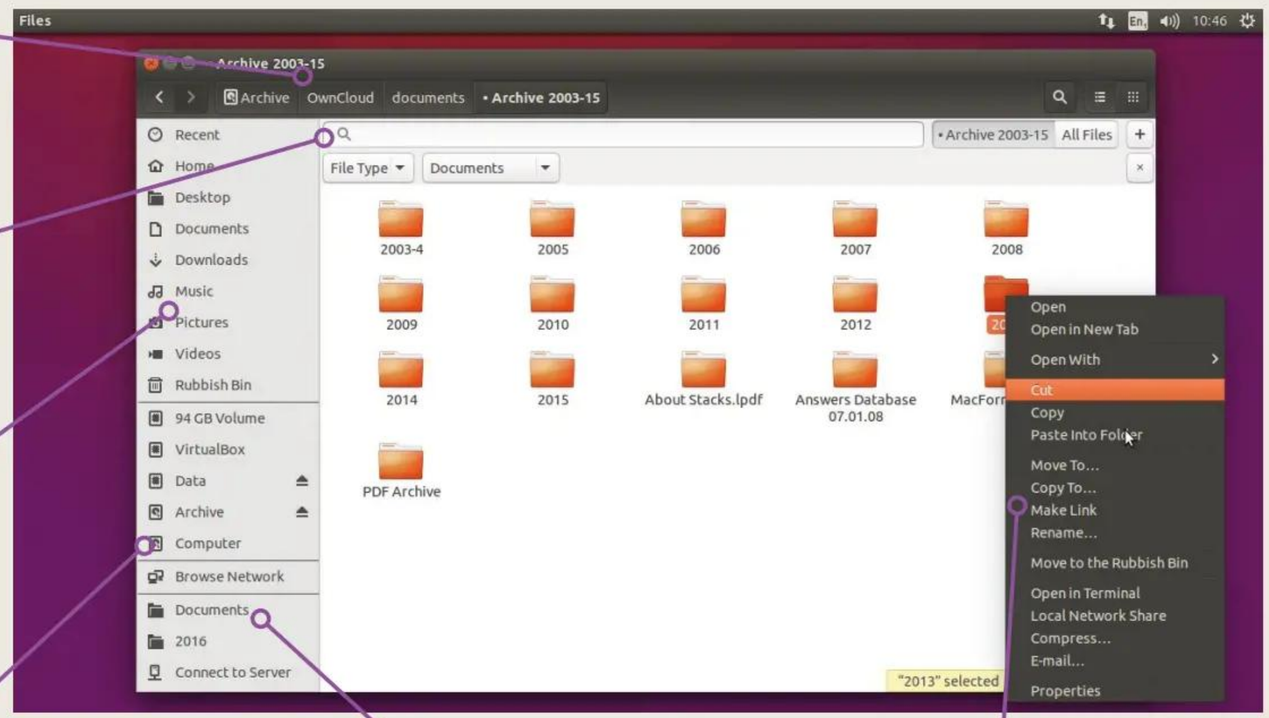

Breadcrumbs
As you drill through the folder tree, use these 'breadcrumb' folder icons to see where you are – click one to jump back to that folder.

Search
This provides convenient shortcuts to key parts of your system – the top section is self-explanatory, and includes a handy list of recently accessed files.

Navigation pane
You can find out how the drive has been partitioned by examining the 'Partitioning' section of its summary. It'll be GUID or MBR – see the main text for more information.

Attached drives
All recognised drives – internal and external – are listed here. The 'Computer' shortcut provides you with a link to the root of your Linux file system.

Bookmarks
Drag folders into this section to create bookmarks for one-click access in the future. They're the equivalent of Windows' Quick Access menu.

Right-click
Right-click a file or folder to reveal more options for interacting with it – 'move to' and 'copy to' options are particularly useful.

drives are mounted inside the **/media** folder, while **/mnt** serves as a 'temporary' mount point for devices like network folders.

The **/root** is the home directory for the root user. It's kept away from the main **/home** folder in case of boot problems and can only be accessed via the 'sudo' command. Peek inside the **/usr** folder and you'll find 'bin' and 'lib' folders – this is where programs you install under your own user profile are kept. Finally, the **/var** folder contains variable data that includes potentially helpful items such as logs for troubleshooting problems.

This is all interesting, and useful to know, but it's not something you should concern yourself with in day-to-day use.

Explore Nautilus

Return to your home directory. Now roll your mouse up to the menu bar where you'll see Nautilus has its own selection of menus. Browse these – File, Edit, View, Go, Bookmarks and Help – and you'll see a list of familiar-looking commands, but some should stand out.

Start with File > New Tab, which allows you to have multiple locations open without cluttering up your desktop with windows. You'll see the tabs appear above the folder pane and you can move (or copy by holding [Ctrl] as you click and drag) files between them by picking them up in one tab, dragging them over another tab and then waiting for that tab to open before dropping them.

Another handy tool can be found by choosing Edit > Select Items Matching... From here you can apply filters that ensure only particular files matching those filters (such as a file type – png, jpg and so on) are selected. While you're exploring the Edit menu, select Preferences to tweak Nautilus' behaviour – of particular interest will be the Display tab where you can add extra descriptive captions (such as file size or owner) under icons.

You'll also notice a reference to zoom here – by default, up to three captions can be displayed underneath icons depending on how large you've magnified the icons. You'll find these controls on the View menu, or you can use [Ctrl] + [+], [Ctrl] + [-] and [Ctrl] + [0] to quickly adjust the zoom level.

Hidden files

By default, Ubuntu will hide any file that's preceded with a period (**.**) mark – right-click a file in Nautilus and choose Rename to see how this works, then press [Ctrl] + [R] to refresh the view, when you'll see it disappear. Don't panic – select **View > Show Hidden Files** and it'll magically reappear, allowing you to rename it once again to remove the period mark and leave it permanently in view.

Struggling to find a particular file? Click the magnifying glass button and a search bar will pop up. Simply enter your search terms and a list of matching files will appear – if you want to widen your search to your entire file system, click the 'All Files' button.

You can also apply file type filters to your search, to restrict results to documents, photos or whatever criteria you choose. Click the '+' button, leave 'File Type' selected and then either pick a type from the drop-down menu (such as Illustration, Spreadsheet or Pdf/Postscript, or manually type a specific file extension (such as .png) to the filter. Nautilus supports multiple filters, so you could easily search for .png, .jpg and .tif photos without including .gif or other image types, for example. ■

Quick tip

If you don't like the way Nautilus uses breadcrumbs to display the current folder path, select Go > Enter Location... to replace it with a traditional view, allowing you to manually enter the path to jump straight to that location.

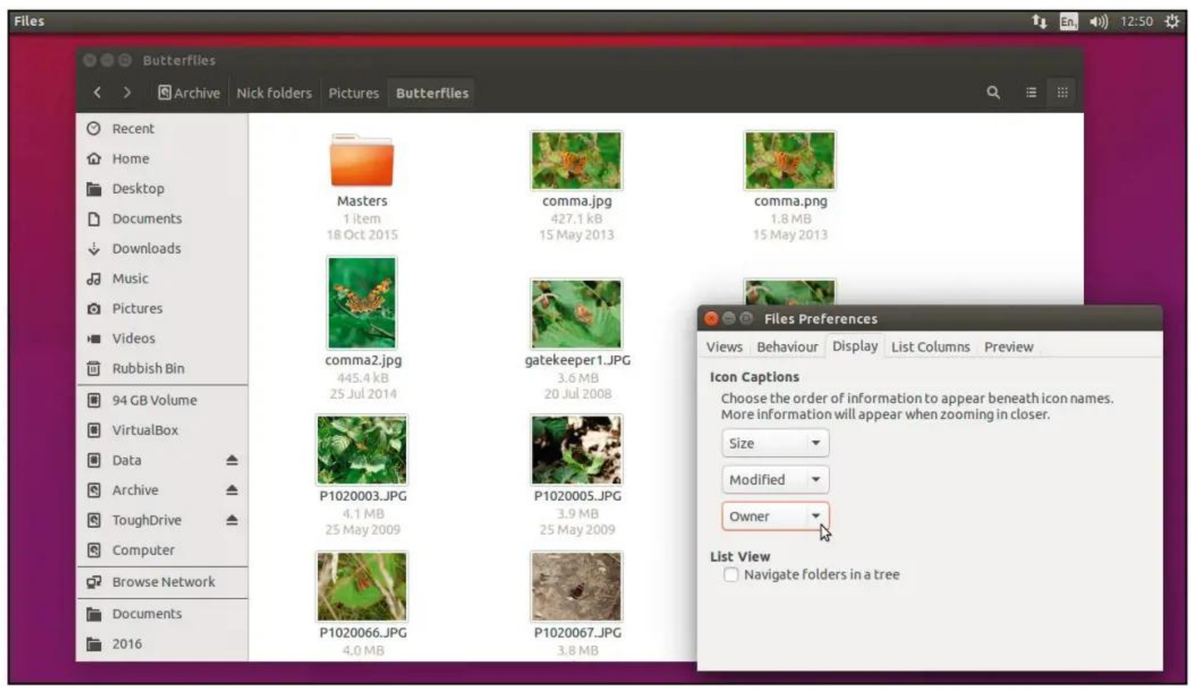

› Customise what gets shown underneath files when browsing in icon view – zoom into reveal more details.

GParted: Set your partitions

Divide up your hard drive using partitioning to protect your data and make Linux easier to work with.

Partitioning allows you to carve up a single physical hard drive into smaller, virtual drives known as partitions or volumes. Each volume acts independently of the others, providing a measure of redundancy for data spread across the drive. If one partition develops problems then you can restore it without affecting what's stored elsewhere.

This redundancy provides the first of a number of reasons to partition your drive: it allows you to move your personal data (typically your **home** folder) to another volume, so it's protected from any changes made to your system partition. You can then use your Linux installation as a sandbox, rolling back any unwanted changes without touching your data.

Another popular use for partitions is to run multiple operating systems on a single PC. You could run two versions of Linux side-by-side or set up a dual-boot Windows/Linux machine for compatibility purposes, eg. This is possible because each partition can be formatted using a different file system, so ext3 or ext4 for Linux, or NTFS for Windows, eg. You can then create a data partition that's visible to both OSes, ensuring you have the latest version of your documents and other files, whichever OS you happen to be running.

It's also possible to use partitions for other purposes, such as creating a dedicated partition for sharing over your network, or setting up independent partitions for temporary files, your swap file or even to run specific services like a web server. Whatever your needs, partitioning can play a crucial step in improving performance as well as protecting different parts of your drive from each other.

Quick tip

If you plan to share a partition with Windows then it needs to be set up as NTFS so your Windows install can read it. You'll need to verify the **ntfs-3g** file system add-on is installed.

Partition types

Hard drives are partitioned according to a scheme. There are two main types of scheme: the older Master Boot Record (MBR), which is limited to the first 2TB of a hard disk, and the newer GUID Partition Table (GPT). In both cases, partition information is recorded physically in the first sector of the drive in a partition table – GPT drives also store this information at the other end of the drive in the last sector for redundancy purposes.

One of the limitations of the older MBR scheme is that it only supports a maximum of four partitions, although one of these can be an extended partition, which in turn can house multiple logical volumes, bypassing this limit. GPT doesn't differentiate between primary and extended (logical) partitions and supports up to 128 partitions per drive.

GPT is designed in conjunction with modern UEFI firmware, but in some cases it can be backwards compatible with older legacy BIOS systems too. It's compulsory if your drive is over 2TB in size, but ultimately which partition scheme you use typically depends on what you already utilise. For simplicity's sake it's usually best practice to stick with what's already there, and if you're planning to wipe the drive completely clean and install Linux on it from scratch, let your distro decide which scheme is best based on your current setup and make sure that you stick with this setup post install.

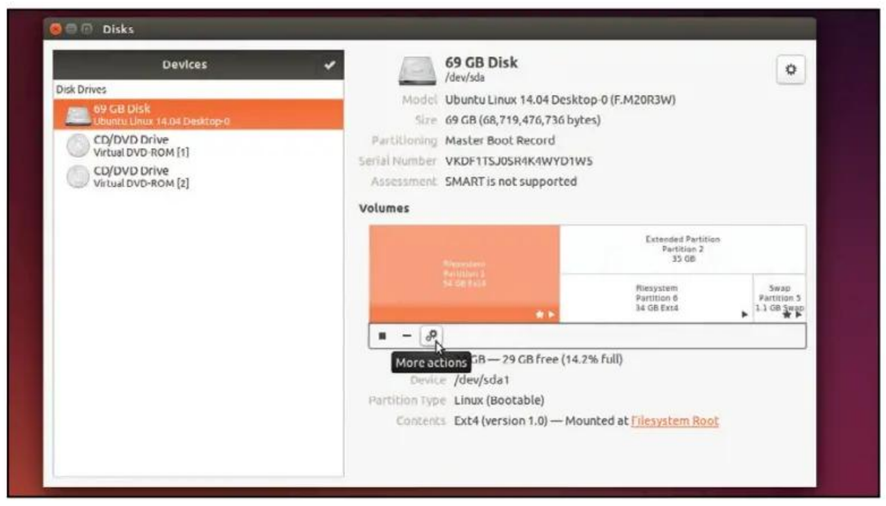

› The Disks utility provides a handy summary of how your partitions are currently set up.

Working with LVM

If you have set up your Linux install using Logical Volume Management then you'll find yourself frustrated should you attempt to repartition it using *GParted* – it simply won't work. That's not a problem though, because all you need is the right tool.

You can, of course, configure your partitions from the command line, but a better bet is to employ the user-friendly LVM utility. To do so, boot from your Live CD, then open Software Center and search for 'LVM'. Select 'Logical Volume Management' and click More Info > Use This Source > Install. Once installed, launch LVM and you'll find a more pleasant environment to work in.

Start by resizing your root partition to free up available space: expand Logical View, select 'root' and click 'Edit Properties'. Use the controls to shrink the partition to its desired size and click 'OK', then wait while the partition is resized. You can then switch to Logical View and click 'Create New Logical Volume' to set up any additional partitions you wish to create, including giving them a friendly name (such as 'home' for your **home** partition).

Don't bother configuring a mount point at this stage – instead, follow steps four through to six of the walkthrough (*see p101*) once the partition is in place to finish configuring it.

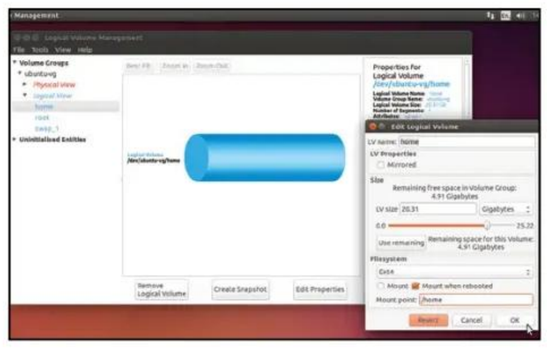

› The LVM utility (or as its package is called system-config-lvm) is the easiest partitioning tool to use on LVM-enabled setups.

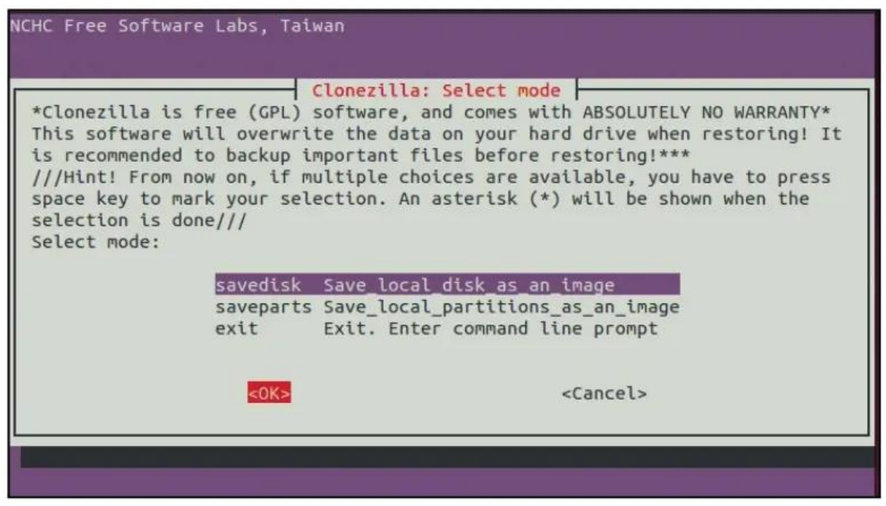

› Remember to take a full drive image of your hard disk before you begin partitioning it. Try Clonezilla.

In this tutorial we'll cover partitioning using MBR, but the process is a similar one for GPT drives. We're using Ubuntu 14.04.3 LTS, but you should be able to translate this to most Linux distros.

How your computer is currently partitioned depends on your individual circumstances, of course, as well as the distro you use. To see how your drive has been partitioned, as well as see what type of partition scheme (MBR or GPT) has been implemented, open the *Disks* utility from the Dash. Select your drive from the left-hand pane and check the 'Partitioning' entry to see what scheme it's been assigned.

Beneath this is the Volumes chart, which provides a graphical representation of your disk, revealing what partitions have been set up. The default Ubuntu 14.04 setup, eg, sees your drive partitioned into two: a main Linux partition that's bootable, with all your applications, data and settings in addition to your Linux distro. There's also a much smaller extended partition, inside which is the **swap** file volume.

Each partition on your hard drive is allocated a separate entry inside the **/dev** folder, which is where references to all the components that make up your PC are stored. Each physical hard drive is represented by three letters: hda, hdb and hdc etc for those drives attached to older IDE controllers, and sda, sdb and sdc etc for those attached to SCSI and newer SATA controllers.

Each physical drive's partitions are represented in numerical terms eg sda1 or sdb5. Look under Device to see which entry has been assigned to the currently selected volume in Disks. One complication: you may have been provided with an option to use LVM (Logical Volume Management) during install – if this is the case, then on traditional MBR-based setups, you'll see a small primary partition described as Linux (Bootable) and a much larger extended partition inside which is the Linux LVM volume. The primary partition exists because extended partitions aren't bootable, so the required boot files are placed on this small partition with the rest of your system left on the extended volume. (*See Working with LVM box, above*).

Partition structure

Partitions are accessed through the file system by mounting them at a specific level. Linux structures its files, folders and partitions as a tree. When you start Linux, your primary boot partition is mounted first at **/** at the root (or trunk) of the tree. Other partitions can then be mounted at specific folder points above it in the tree – the **/mnt** and **/media** folders are a good choice for a partition that's specifically been created for sharing, eg. It's also where external partitions – eg from another operating system – are automatically mounted if the partition's file system is recognised by Linux.

It's also possible to mount partitions directly to key directories like **/tmp** (temporary files) or **/home** (the home directories of each user on your PC) when applicable – this provides a seamless and consistent experience regardless of whether you decide to set up dedicated partitions for key folders or not. Mount points are stored in the **/etc/fstab** file.

The best time to partition your drive is when you first set up Linux on an empty drive. (*Check out the Partition from Scratch box, p94, for details on how easy this is.*)

You'll also encounter partitioning when setting up a dual-boot system. In this event, things are made much easier if your existing OS is detected during the installation process – if it is the hard work of partitioning is done for you because the installer suggests a suitable partition layout, which you can tweak to your personal needs based on how much free space is currently available on the drive.

That's not a practical route for most people to take – and thankfully there are third-party tools that can repartition your drive without data loss, although it's important to note there's always an element of risk involved. That's why it pays to take a full drive image of your current setup now, so if things go wrong you can reset and start again. Use a tool like Clonezilla (**http://clonezilla.org**) or create a snapshot if you're running your installation in a virtual machine. It also pays to sit down and work out what your requirements are and how much »

Quick tip

Another advantage of LVM over regular partitioning is the fact it allows you to add additional physical disks to your computer, then map the additional space on to existing volumes. This allows you to boost capacity without having to move or copy data between drives.

space you have to play with. In the first instance, decide what additional partitions you wish to create and what you're creating them for. If it's to protect them against a reinstall or possible corruption on the main drive, then they can be housed on the same partition; remember too that while this provides some protection against data loss, it's not a substitute for backing up.

If you're looking to boost performance by moving key folders – your **swap** file, perhaps, or the **usr** folder where programs are stored – then these will need to be housed on a separate physical drive to the one Linux is running on.

Other considerations include the size of each partition – this depends on its type, and how much spare space you have. Start by right-clicking the folder in question in your file manager and selecting Properties. This is technically the minimum amount of space you'll need to allocate the partition, but think what else might need to be stored in the folder and add on the extra space accordingly. Note: you'll be taking space from your primary partition for the new volume, which may limit how much space you can assign to it.

Other considerations: should the partition be primary or extended? If you only plan to create a single additional partition then simply resize your system partition and create a new primary partition in its place. If you plan to create a number of partitions on your main drive, then using the extended partition is advised (*see boxout, below*).

You'll also have to decide what file system to assign the new partition. If you're creating dedicated partitions for folders like **home** or **var**, then you'll want to make these the same as your primary Linux partition (typically ext3 or ext4). The same is true for partitions you plan to use as shared folders over your network. Indeed, the only time you'll want to change the file type is if you're creating a partition that you plan to share with a Windows installation on the same PC, in which case choose FAT32 or NTFS.

Quick tip

It's possible to add a new partition to **fstab** without editing it directly – open the *Disks* utility, select your partition and choose More Actions > Edit Mount Options... to do so. You can't, however, add the 'pass' field using this method, preventing the partition being checked for errors at startup.

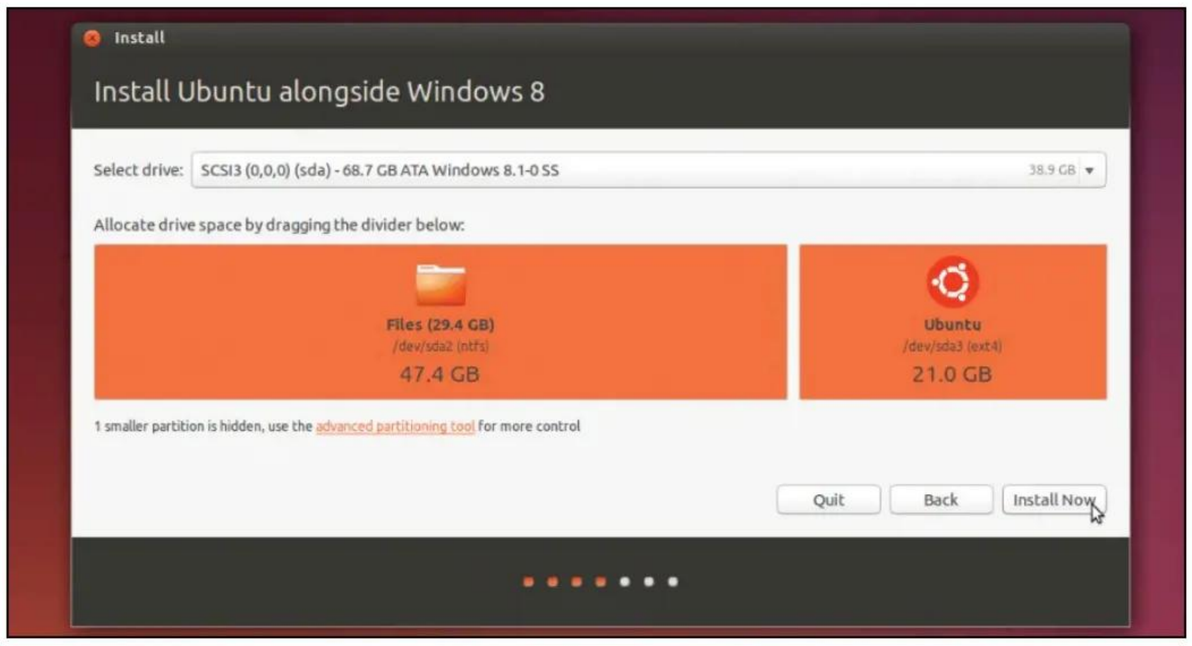

› As long as Ubuntu recognises the OS that you're trying to install it alongside, the partitioning process is easy.

When it comes to actually partitioning your drive, then the best tool to use already lives on your Ubuntu Live CD and it's *GParted*, a user-friendly graphical front end that makes resizing partitions simple (while also providing a handy graphical overview of the physical layout of your hard drive). You'll need to boot from this to resize your system partition anyway, so you might as well perform all your partitioning needs while outside your main Linux install.

Partitioning tools

You can also run *GParted* from its own standalone bootable install on CD or USB from **www.gparted.org/livecd.php**. Make sure you pick the right architecture, which is typically i586 for older, 32-bit computers, and amd64 for newer 64-bit machines with UEFI rather than legacy BIOS. Users of 32-bit can also experiment with the i686-PAE build if they find the i586 build is a bit sluggish.

The step-by-step guide (*see right*) reveals how to shrink your main partition before creating a second primary partition alongside it using *GParted* from the Ubuntu live CD. If you'd rather create your new partition or partitions inside the extended partition, then the procedure is slightly different. First, resize the system partition as outlined in the walkthrough. Once done, you need to extend the small extended partition that contains your **swap** file to take up all the remaining free space.

This can't be done while the **swap** file is in use, so right-click the swap partition in *GParted* and choose 'Swap off'. Once done, you'll be able to right-click the extended partition and choose Resize/Move. Type 0 (zero) into the 'Free space preceding' box and click 'Resize/Move' to quickly allocate all available space to the partition. You can now partition the free space as you wish without worrying about running out of available partitions – just make sure you calculate how much space each partition is likely to need and assign it accordingly, and don't forget to click the 'Apply' button when you're done.

The tools we've mentioned are pretty smart on the whole, but potentially dangerous task, which is why taking a full system backup before you begin is essential. If you do run into problems, check out the GParted help pages at **http://gparted.org/help.php**. You'll find handy links to an FAQ with, among other helpful things, a guide to fixing problems with *Grub*.

You may run into problems when attempting to mount new partitions into key folders such as your **home** folder. If Linux throws up an error on startup, press m to manually recover from it, then type `sudo nano /etc/fstab` to examine the file and check there are no errors preventing the partition from mounting correctly. If you can't find any problems, delete the line, save the file and restart before investigating. ■

Partition from scratch

The best time to partition is when you're setting up a fresh install. In Ubuntu, when prompted to erase the disk, select 'Something else' and click 'Continue'. You'll see a single device – **/dev/sda** – is set up. Click the 'New Partition Table...' button, read the warning and click 'Continue'. Now select the free space and click the '+' button next to Change to create a partition.

Let's begin with the main system partition. This can be as little as 15-20GB space, but make it larger if necessary. Once you've calculated how much you need in gigabytes, multiply it by 1,024 and enter the figure in the Size box. Leave 'Primary', 'Beginning of this space' and 'Ext4 journaling file system' selected, choose **/** under Mount point and click 'OK'.

Select the remaining free space and click '+' again. Click the 'Use as' drop-down menu and select 'swap area'. Set the partition type as Logical, but set its location to 'End of this space'. What size (in MB) it should be depends – a rule of thumb is to make it the same size as the RAM installed in your PC (make it double the size if you have 1GB or less). Click 'OK'.

Finally, select the remaining free space and click '+' again. Leave the size, type, location and 'Use as' settings as they are, and finally set the 'Mount point' to **/Home**. Click 'Install Now' followed by 'Continue' to set up your partitions and install Ubuntu.

Create a dedicated home partition with GParted

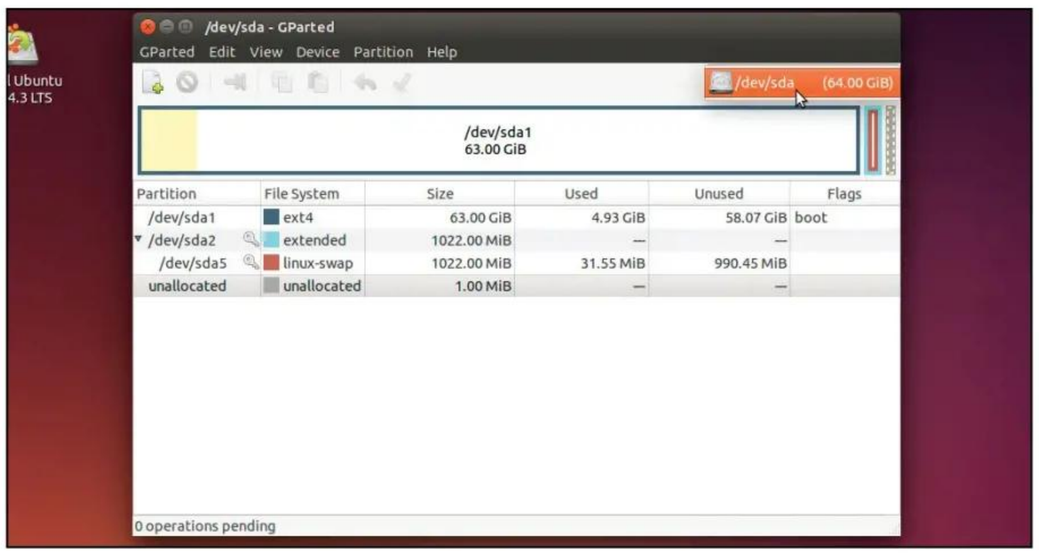

1 Launch GParted

Reboot from your Ubuntu Live CD or USB flash drive, opting to 'Try Ubuntu' when prompted. Once at the desktop, launch *GParted* by typing its name into the Dash. If you have more than one physical disk attached to your PC, make sure the correct one has been selected (typically **/dev/sda**) by clicking the button in the top right-hand corner of the *GParted* window.

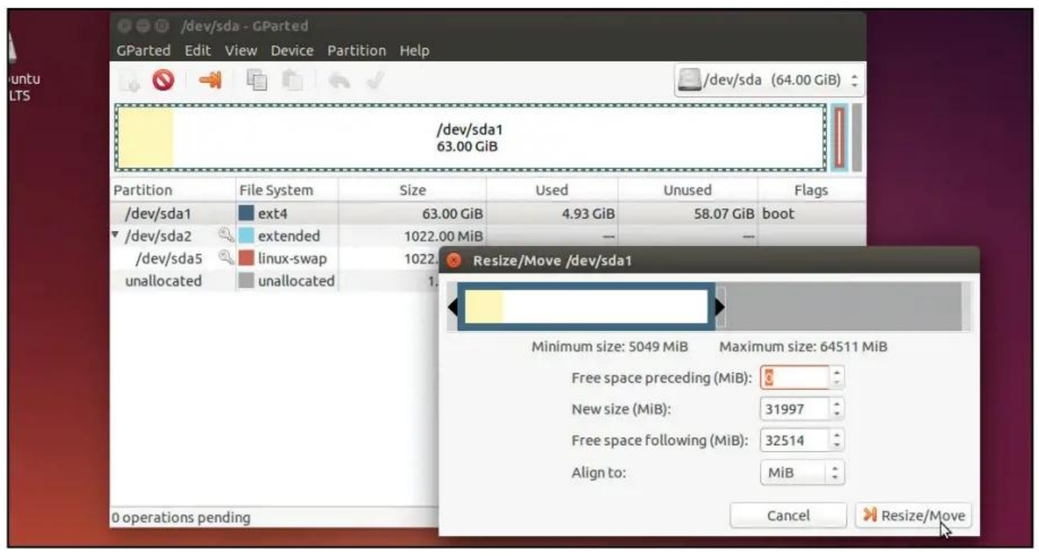

2 Resize system partition

Right-click the main partition (**/dev/sfa1**) graphic and choose 'Resize/Move'. Click and drag on the right-hand edge of the bar to reduce the partition's size roughly to where you want it to fall based on how much space you plan to allocate your **home** partition. If necessary, fine-tune the size using the 'Free space following (MiB)' to set an exact amount. Click 'Resize/Move'.

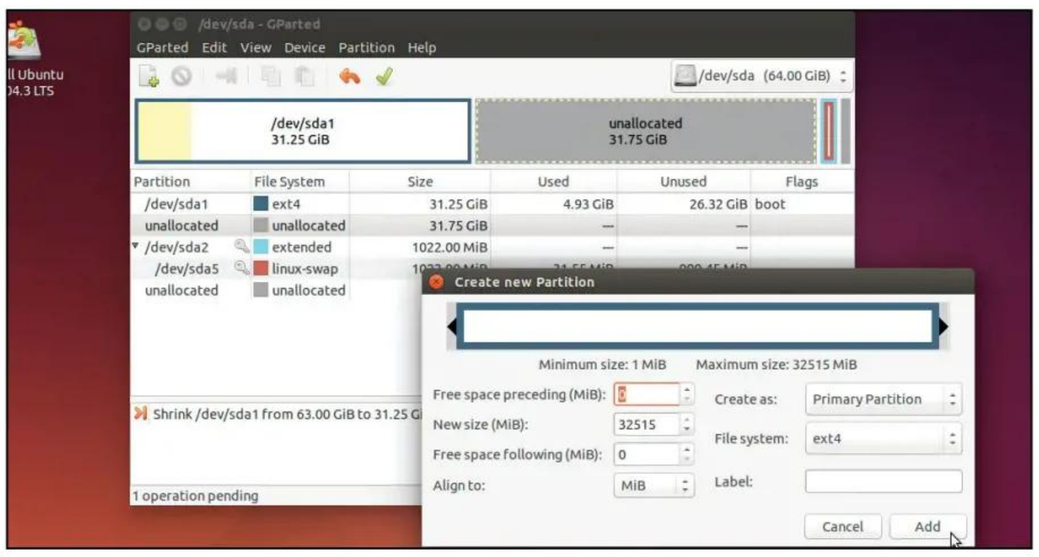

3 Create home partition

Next, right-click the unallocated space and choose 'New'. Assuming this is the only additional partition you wish to create, leave the default settings as they are and click 'Add'. Nothing has yet been done to your drive, so review your changes and if you're happy click the green tick button followed by 'Apply'. Wait while *GParted* first resizes your main partition and then creates the new one.

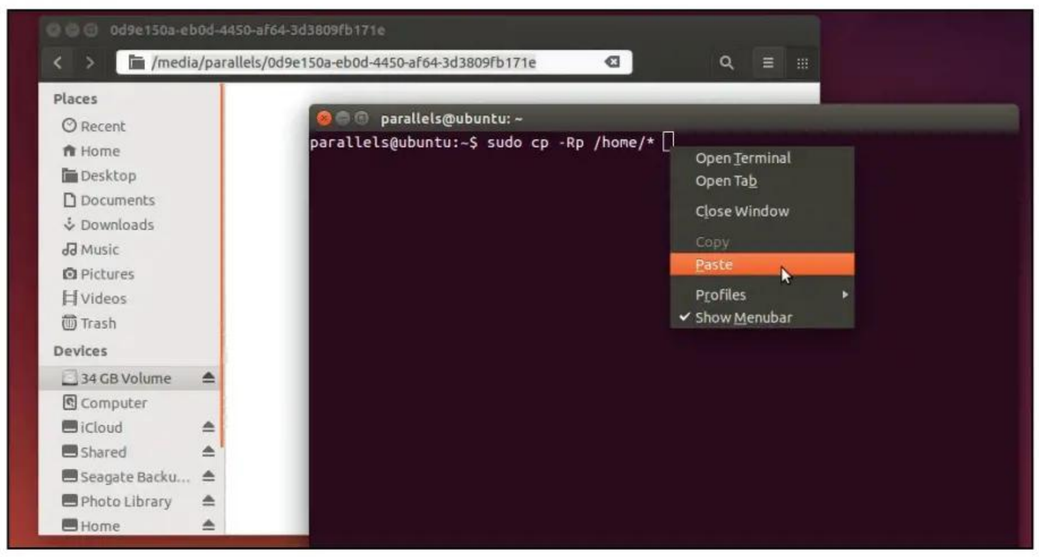

4 Copy files to new partition

Reboot into Ubuntu proper. Open Files, and click the newly visible volume that appears under Devices. Select Go > Enter Location, then select the location in the Address bar, right-click it and choose 'Copy'. Open Terminal, type `sudo cp -Rp /home/*` before right-clicking and choosing 'Paste' to enter the location of your new partition. Hit Enter and ignore any errors.

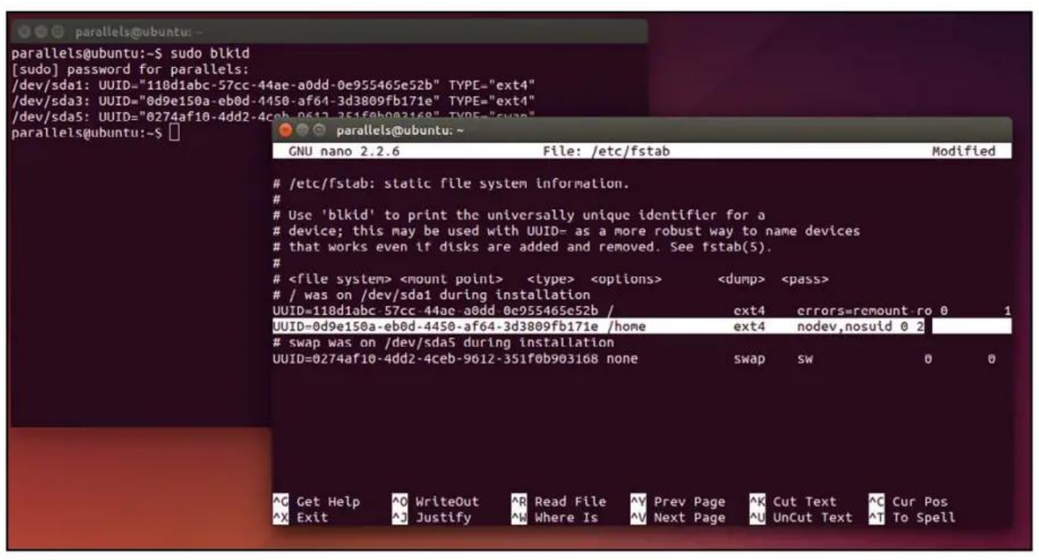

5 Update fstab

Once done, you should see the partition is now populated with the files from your **home** folder. Next, type `sudo blkid` and press Enter. Make a note of the UUID, then type `sudo nano /etc/fstab` and hit Enter. Type the following beneath the first line marked `UUID=`, replacing `<location>` with your partition's UUID: `UUID=<location> /home ext4 nodev,nosuid 0 2`.

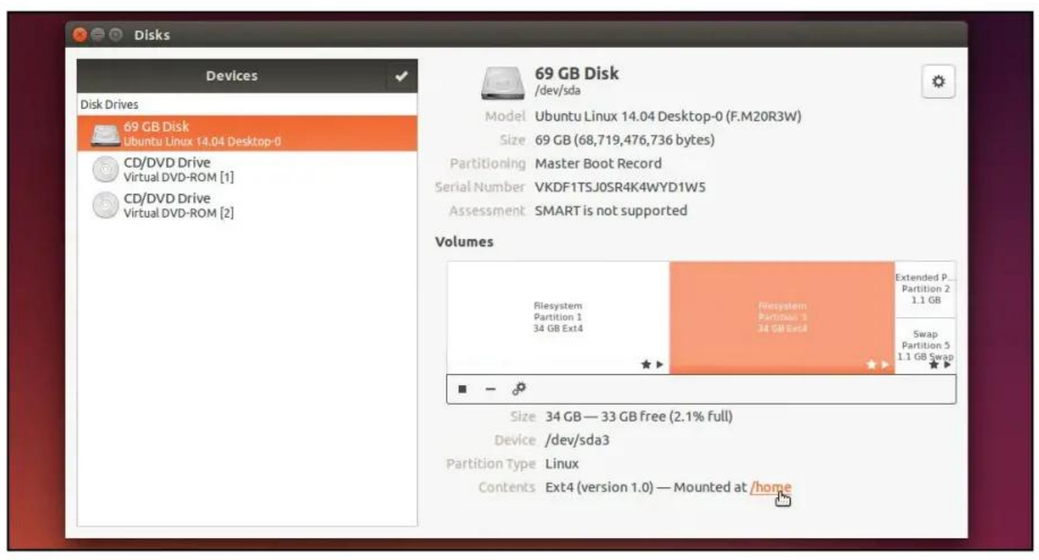

6 Mount new partition

Once you've made your changes to the **fstab** file, go ahead and save the updated file and close it, then type the following series of commands into the Terminal and press Enter:

```
cd / && sudo mv /home /home_old && sudo mkdir /home
```

You will need to reboot your computer and if you open *Disks* to verify that your new partition is correctly mounted at **/home**.

Files: File permissions uncovered

Take a deep breath and launch into the murky world of users, groups and files in Ubuntu.

Quick tip

If you share your PC with others, open System Settings and select User Accounts to give everyone their own, private user account and home folder. Share files with others using your Public folder.

If you've migrated to Ubuntu from Windows, you'll already have an inkling about file permissions and users, plus the effect one has on the other, but it can still be a confusing subject to grasp, particularly where Linux is concerned. Before you dive into the confusing world of who's got what access to which file, acquaint yourself with the essentials of file permissions.

Grasp the fundamentals

Despite appearances to the contrary, every single thing in Linux is a file, including folders and hardware devices. On top of this, the Linux file system – **ext2**, **ext3** or **ext4** – applies special permissions to files that ultimately determine how they're accessed, and by whom.

There are three levels of access to files in Linux, and while they differ slightly between files and folders, they follow the same broad principle: read, write and execute. Access to any or all of these permissions can be granted or blocked. Read access typically means viewing a file or looking at the contents of a folder, while write allows you to edit a file, or – in the case of folders – create or delete files that are inside that folder. Execute allows users to run files that are programs, or actually enter the folder in question.

These three levels are expressed as **r** (read), **w** (write) and **x** (execute). For example, a file with full access is expressed as **rwx**, while one that's read-only is **r--**, and another that has read and execute access is **r-x**. These terms – **r,w** and **x** – are important because they're displayed when you view extended

Exploring Nautilus's file permissions interface

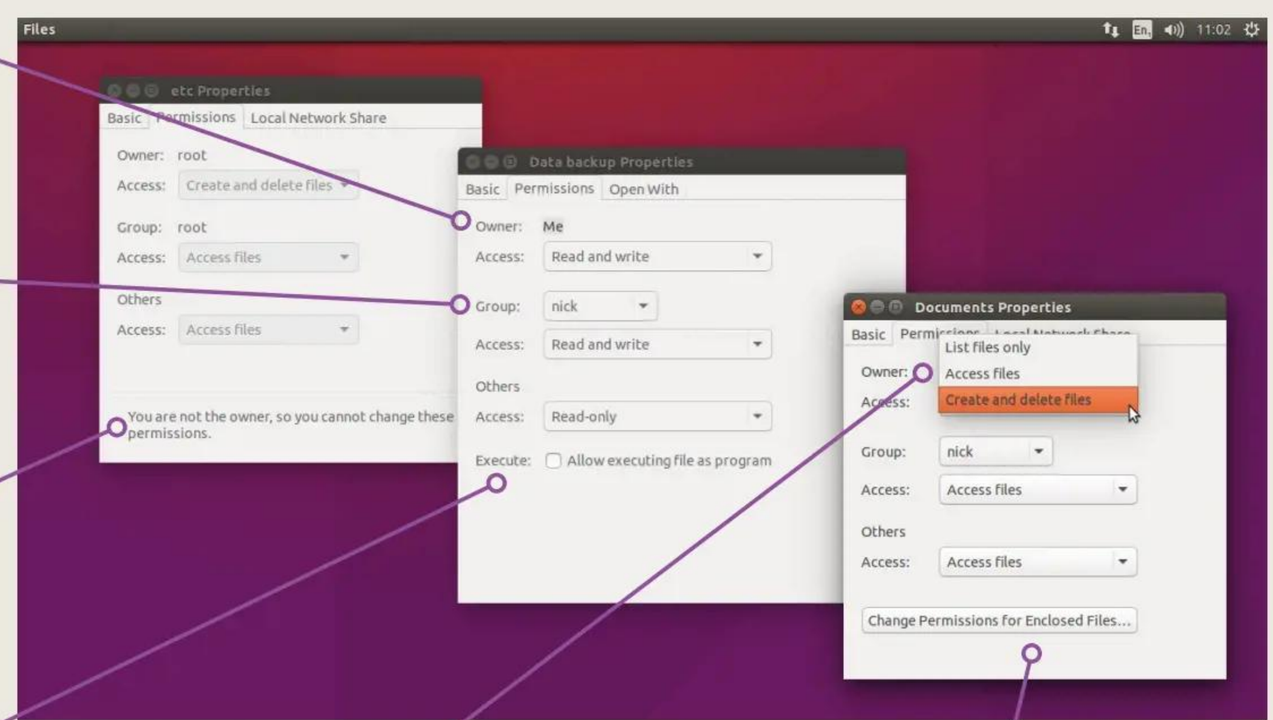

Owner

Another restriction when managing permissions through Nautilus is that you can only view a file or folder's owner through the Permissions tab.

Change group

Click the 'Group:' dropdown menu to view a list of available user groups on your system – most refer to hidden system users, so avoid these.

Access rights

If you're not the owner of a file or folder you can view permissions, but you can't change them – in Nautilus or via the Terminal.

Execute

This box only requires ticking if it has a .run or .bin file extension, or you want to run a file as a script. Note it affects the permissions for all users.

Change access

Click here to change the access level – you have fewer options in Nautilus. The three options shown for folders translate to r--, r-x and rwx respectively.

Enclosed files

This button will apply the same permissions you've set for the folder to all files inside it, ensuring they stay consistent.

file properties in the Terminal using the `ls –l` command to list the content of a folder.

Folder versus file

Things are complicated further because setting permissions on files is only half the story. The file's parent folder plays a crucial role too, because if you deny access to the folder by not granting the execute ('**x**') permission, then it doesn't matter what file permissions you set, the files themselves are kept out of reach, so you can't view, edit or execute them.

On the flip side of the coin, if you have a file that you've denied write access to, you might think that would protect it from being erased. In fact, the permission for deleting a file doesn't rest with the file itself, but rather the directory it's stored in. Set the folder to write ('**w**') access, and the file can be deleted, regardless of what permissions you set on it.

Owners, users and everyone else

Ubuntu is all about security, so rather than apply a single set of permissions to each file and folder on your system, Linux applies permissions to three separate categories, giving you more granular control over who has what level of control. The first is the 'owner' of the file, which by default is the user who created the file in question.

The second category refers to a single user group for whom specific permissions have been defined, and opens up a question about users and groups. Although it appears to apply to a specific user (often the same user as the file or folder's owner), this set of permissions actually applies to a user group.

When you create a new user, a group of the same name is also created, your user is added to that group and it's actually this group that Linux references here. It's also possible to add users to multiple groups, allowing advanced users to set up groups into which multiple users are added, giving all those users the same level of access to the selected file or folder based on the group they're part of.

The third and final category is a catch-all 'everyone else' group, which applies to all other users on your PC who aren't covered by the first two categories.

How are permissions set? When a file is created, it's automatically owned by the user that created it, whether that's your own user profile or – if you've created it through the sudo command – the '**root**' user. The owner is also automatically set as the default user or group. The owner typically has full access rights to the file or folder created, while 'everyone else' normally has more limited rights to files, and are usually blocked from folders.

This approach forms the cornerstone of security on your Linux file system. By default, all users have full ownership, access and control over their personal Home folder and its contents, while other users are blocked access. Outside the home folder, access is more restrictive – certain folders are accessible, but most are either read-only or off-limits completely, requiring you to access them via the 'root' superuser account.

View and change permissions

We've seen how you can use the '**ls -l**' command in the Terminal to view the permissions of a file or folder. Doing so not only reveals the permissions for each category, but also the name of the file owner and the group that has special access rights.

If you'd prefer a more user-friendly view, plus want limited access to changing permissions for any reason, browse to the folder or file in question, right-click it and choose Properties. Switch to the Security tab and use the annotation above to help guide you through what's what.

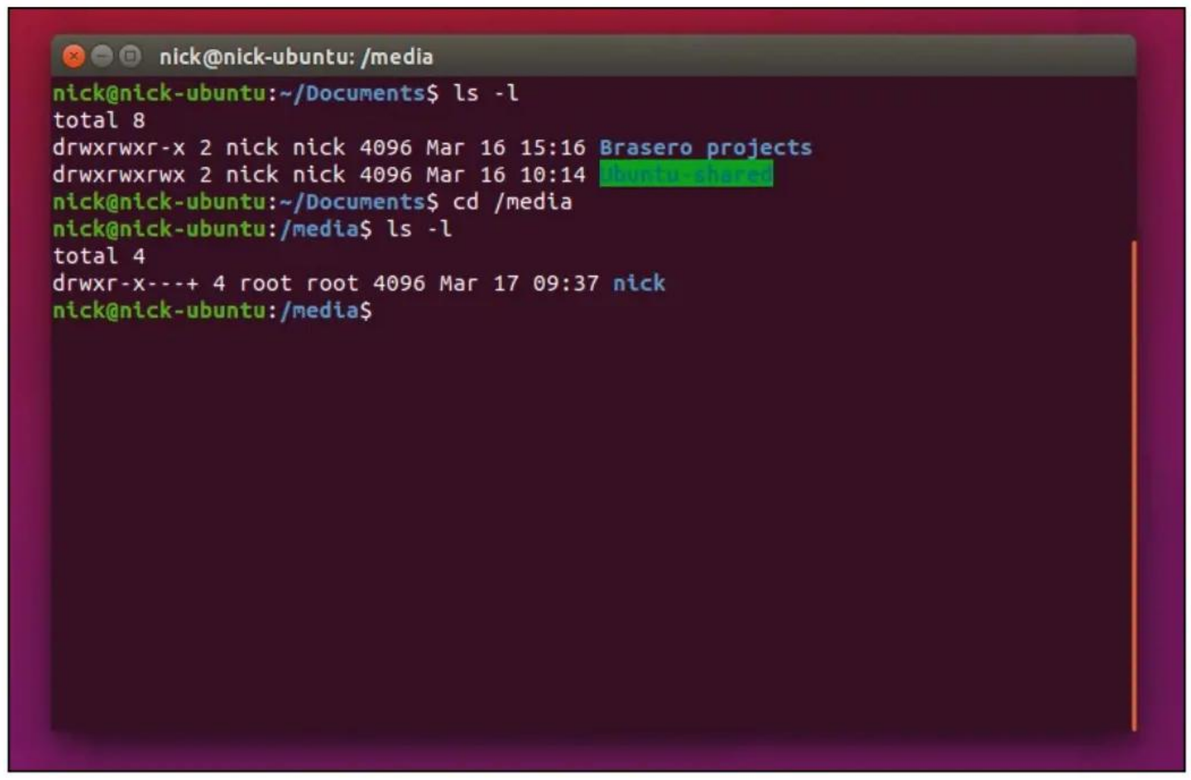

› **Check the permissions for any file or folder using the ls -l command in a Terminal window.**

It's worth noting a few things here. First, you can only change permissions for items you own, while the options on offer aren't as flexible as those available when you edit permissions using the Terminal. On the other hand, it's a much easier way to quickly check the permissions of an object and take remedial action if you think there's something awry.

Another thing: you can also change the file's group here. Click the dropdown menu and you'll see a list of users you won't recognise in addition to any users you've set up yourself – these are specially created system users, designed to do specific things without compromising on security. There's no reason to use these groups in the context of file and folder permissions, so leave them well alone.

In fact, it's probably a good idea to avoid messing about with permissions unless you know exactly what you're doing. Nautilus makes it impossible for you to lock yourself out of your own files by changing the file owner, but you can still quickly and easily get lost among the various options on offer.

That said, if you do want to learn about wielding complete control over permissions, users and groups – including the ability to assign users to different groups – you'll need to get your hands dirty in the Terminal. For more help and advice in this area, check out our tutorial on page 136. ■

Quick tip

If you connect to a drive that isn't formatted using the ext file system, then permissions for all of its files and folders are determine when you mount it: read/write by default, but read-only is also an option.

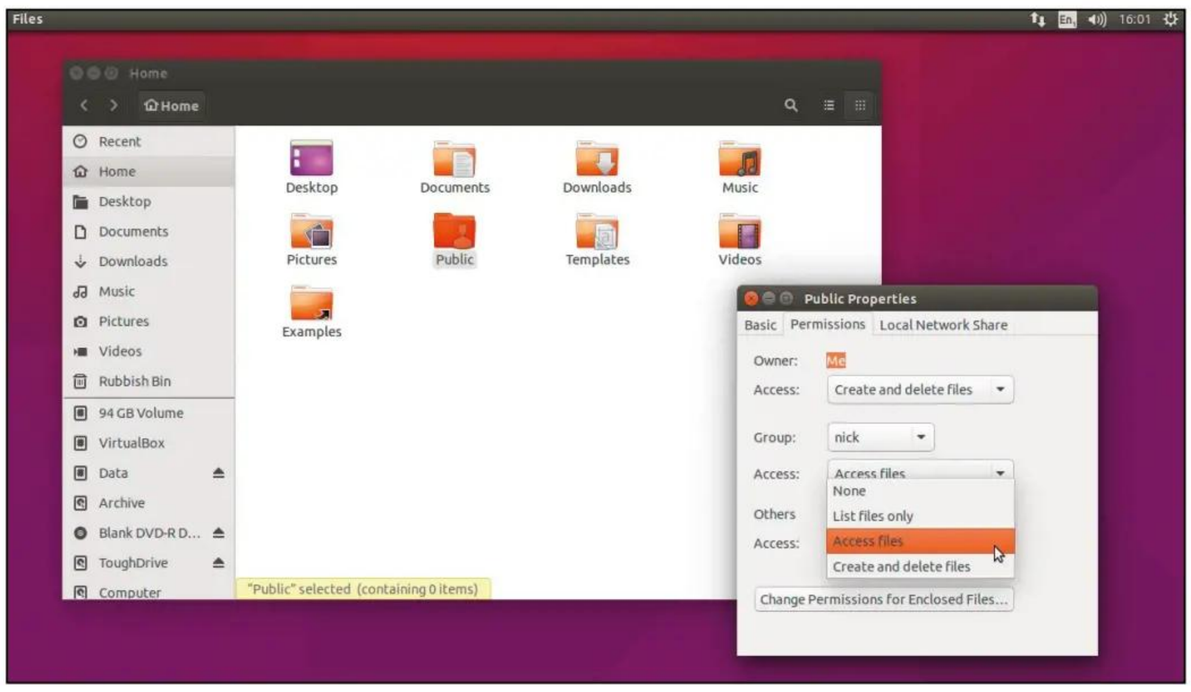

› **The Public folder is a good place to share files with other users of your PC on a read-only basis.**

Advanced

File sharing 100
Alternative desktops 102
Remote desktop 104
Virtual machines 106
System repair 110
Fix startup issues 114
Two-factor authentication 118
Easy file encryption 120
Core commands 122
Troubleshooting 126
Terminal basics 136
Terminal files 138
Apt-get 140
Archiving files 142

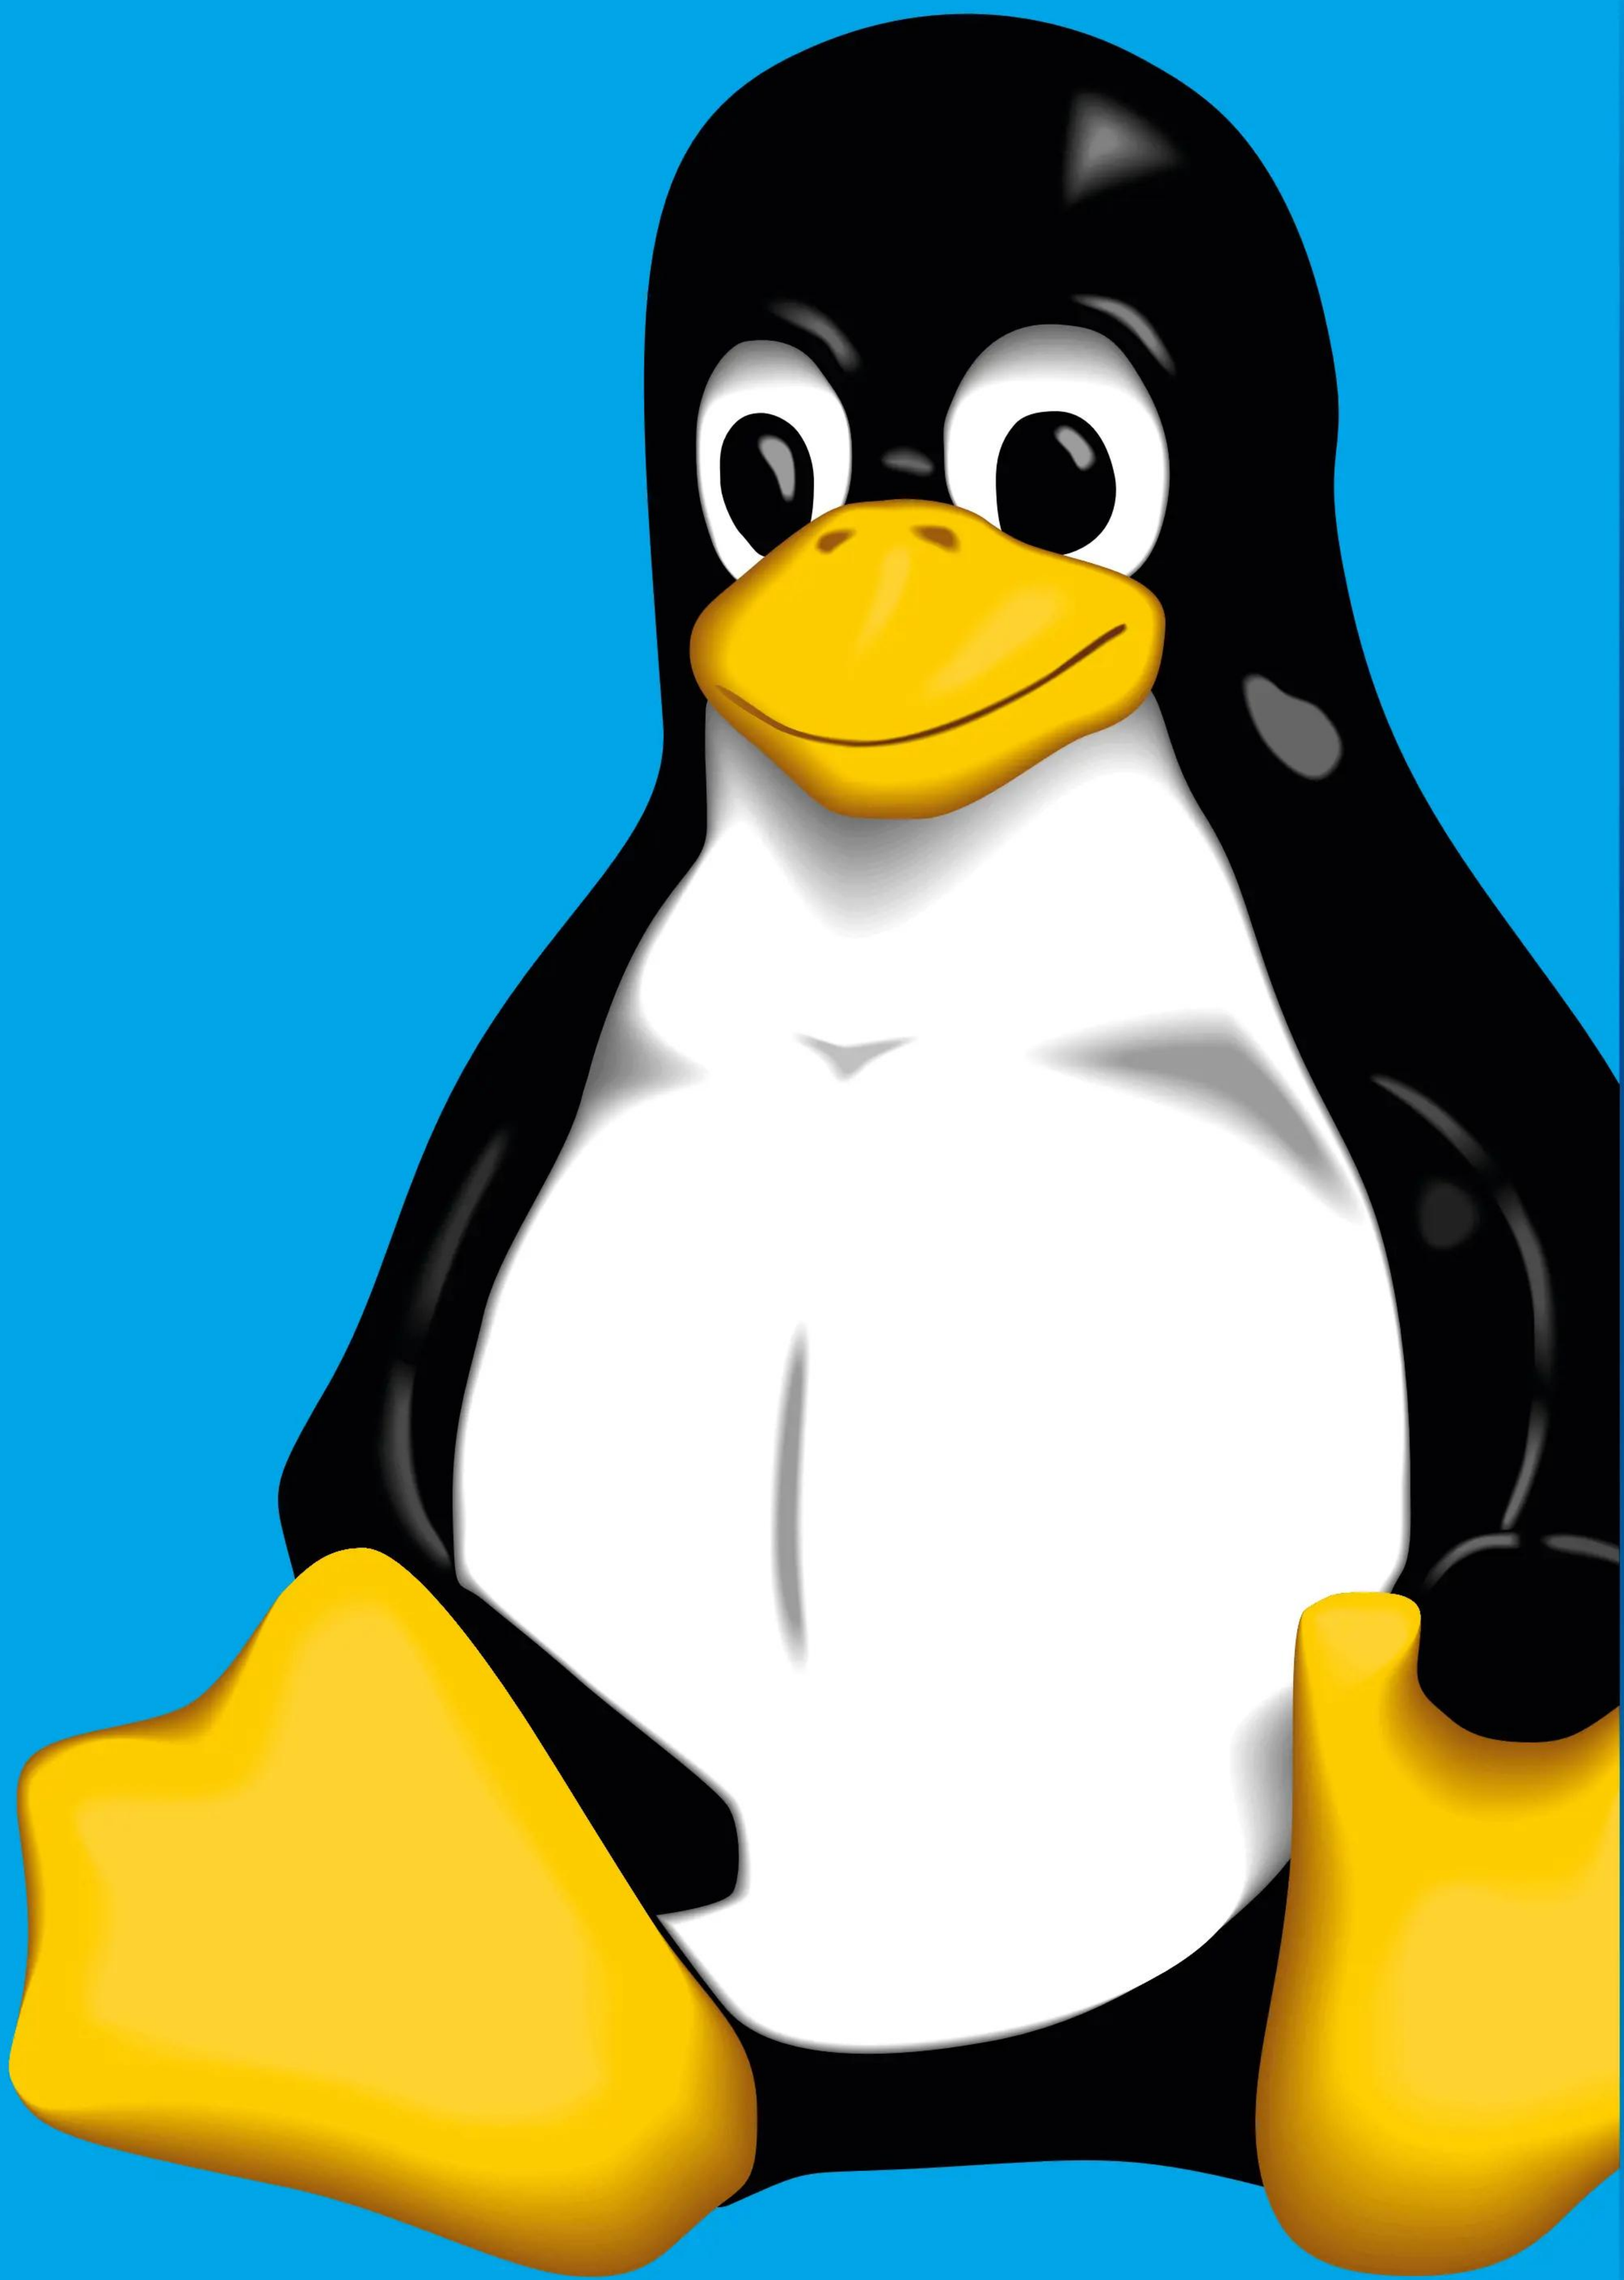

Networking: File-sharing made easy

Why should distance be a barrier? Here are all the steps to follow to get file sharing set up on your Ubuntu PC.

Quick tip

If you're on someone else's computer, make sure you don't choose the 'remember forever' option when connecting to a shared folder, otherwise anyone will able to reconnect using your credentials.

Sharing files between PCs on your network is a fundamental function of modern computing, and the good news is that Ubuntu makes file sharing relatively straightforward thanks to its support for the SMB (also known as SAMBA) sharing standard. Ubuntu also supports another standard – NFS – but that only works with other Linux machines, so for the purposes of this tutorial we're focussing on SMB file sharing, which allows you to both access shared folders and share files from your own PC with other computers – Linux, Windows and Mac – on your network.

Connecting to other computers is as easy as opening Ubuntu's Files tool from the dash and clicking 'Browse Network' – after a short pause you should see all visible computers on your network appear, along with a Windows Network folder that allows you to browse your network (by workgroup) for existing Windows shares.

Double-click on a computer and, if it has a shared folder set up, you'll be prompted to enter your username, domain and password to access it. The domain can be safely ignored in most user setups – the username and password will need to match a profile set up on your target computer (typically your own user profile on that computer). If you're connecting to a public folder, leave the username blank and click Connect. You'll be asked if you want to connect anonymously – leave this selected.

Before clicking Connect, you'll also see an option to save your user credentials to save you the bother of providing them the next time you connect – choose carefully here (*see the tip box*). Once connected, the shared folder works like any other: you'll be able to easily drag files from it to your PC, and – if the folder's shared permissions allow it – you can copy files into that folder too.

It's simple and incredibly straightforward, but note that the path to the shared folder is lost each you time you restart your PC, forcing you to manually browse to that folder again in order to connect to it. One way to save you some time when reconnecting to a frequently used shared folder is to create a bookmark to it in Files. Simply right-click the shared folder's entry in the left-hand pane underneath the 'Browse Network' button and choose 'Add Bookmark'. You'll see the bookmark appear at the bottom of the left-hand pane, and reconnecting in future is as simple as clicking the shortcut.

It's also possible to automatically connect to shared folders when Ubuntu launches, but this involves a trip to the Terminal. A simpler alternative involves installing an app called *Gigolo* from the Software Centre. Once installed, open it from its shortcut icon in the launcher.

Gigolo allows you to browse, connect and bookmark shared folders with ease – the annotation reveals how it works in more detail. As things stand, you'll need to launch *Gigolo* after starting Ubuntu to automatically connect to those shared folders. To solve this problem, configure *Gigolo* to start with Ubuntu: open dash and type 'startup', then launch Startup Applications. Click Add. Name it Gigolo, then type `/usr/bin/gigolo` into the Command box. Add a suitably descriptive comment such as 'Mount shares at startup' and click Add.

Gigolo will now launch when you log into Ubuntu – to prevent its main window from appearing, switch to the program and choose Edit > Preferences > Interface tab.

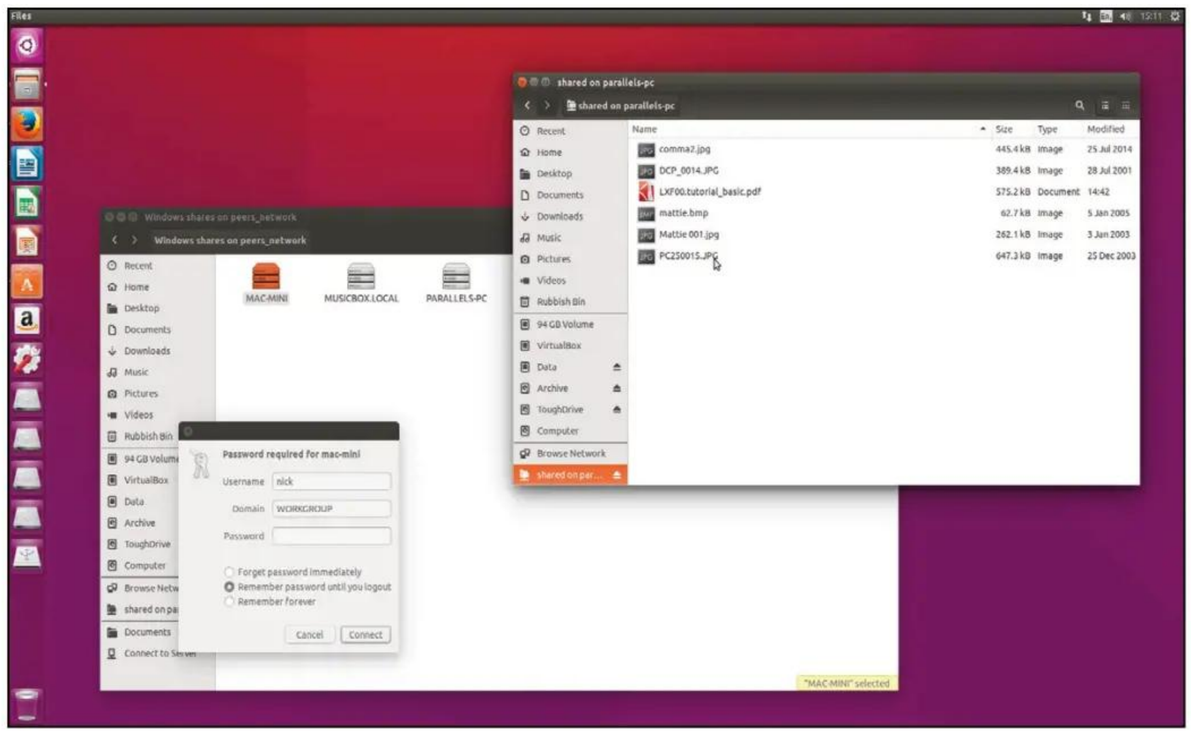

› **Use the 'Browse Network' shortcut to manually search your network for folders to connect to. It's utterly straightforward.**

Exploring the Gigolo interface

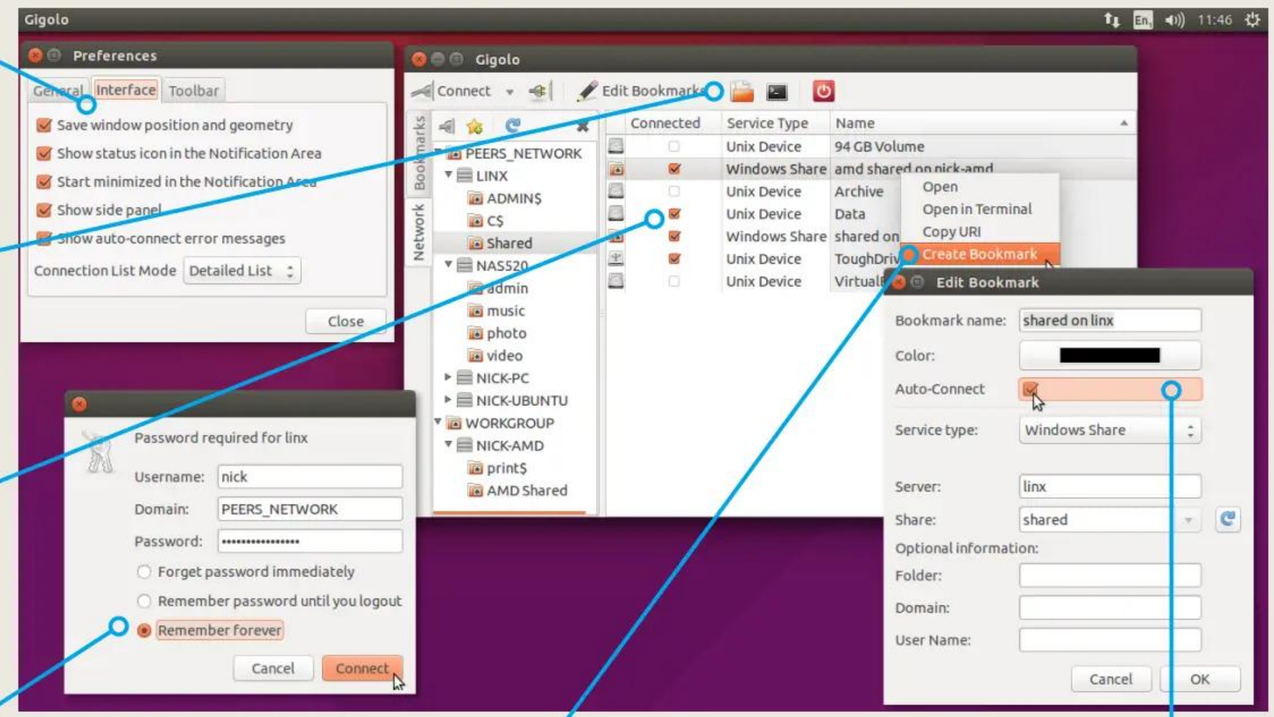

Change view
Gigolo's default view is a bit bare: select Edit > Preferences > Interface tab. Tick 'Show side panel' and set 'Connect List Mode' to 'Detailed List'..

Shortcut heaven
Click either of these two buttons to open the shared folder in Files or Terminal – a handy time saver when managing files.

Quick connection
Gigolo can also be used to quickly connect to non-mounted drives – just tick the box to make the connection..

Browse network
Select the Network tab and Gigolo will scan your local network – you'll be prompted to log in to each folder it finds (click Cancel to skip).

Add bookmark
Connect to your shared folder first so it appears in Gigolo's main window, then right-click its entry and choose 'Create Bookmark'.

Configure bookmark
You can configure shares to mount automatically with Gigolo (or when Ubuntu boots) – tick the 'Auto-Connect' box when setting it up.

Tick the 'Start minimized in the Notification Area' panel and click Close.

Share folders with others

If you want to set up a shared folder on your Ubuntu PC, the steps can be simple or complicated depending on your needs. Thanks to Ubuntu's granular use of file permissions, you can tie yourself in knots trying to set up folders that only specific people can access. Thankfully, if you want to set up a folder for public access (no password required), or you want to set up a folder that requires your own Ubuntu username and password to access, then things are a lot less complicated.

First, you'll need to install two additional packages before you can start sharing folders. Open a Terminal window and type the following two commands:

```
sudo apt-get install samba
sudo apt-get install libpam-smbpass
```

Reboot – when you log in, you may see a message telling you that your user profile has been added to Samba's file sharing group. Now browse to the folder you wish to share – for simplicity's sake (and to ensure you have all the right permissions), make sure this is a folder stored inside your personal Home folder. Right-click it and choose Properties, then switch to the 'Local Network Share' tab. Tick 'Share this folder', then – if required – choose a different share name for the folder to make it easier to identify on the network (you can also add a comment to the folder to make it easier to identify).

If you're happy to open up your folder to anyone on your network, then ticking the 'Guest access...' box under the 'Local Network Share' tab of the shared folder's properties is sufficient. If you want to restrict the shared folder to yourself, leave this unticked – anyone attempting to access this shared folder will need to know your Ubuntu username and password to access it.

Next, decide whether users can copy files to the folder as well as delete those within it (tick the 'Allow others to create and delete files in this folder' box), click 'Create Share' or 'Modify Share' followed by 'Add the permissions automatically' if prompted and you're done: the folder should now be visible and accessible over the network to all other connected computers, including Apple Macs and Windows PCs. You should also see a two-way arrow appear on the folder icon to indicate it's now being shared.

Beyond the basics

By default, your PC's Windows network workgroup is WORKGROUP. If your other computers are on a different workgroup you can change this by opening Terminal and typing the following command:

```
sudo nano /etc/samba/smb.conf
```

A text editor will open. Use the cursor keys to scroll down to the [global] section and change the following line to match your workgroup name:

```
workgroup = WORKGROUP
```

Press [Ctrl] + [O] and hit [Enter] to save the file, press [Ctrl] + [X] to quit *nano* and then reboot your PC. When you log back in, you'll be on the correct workgroup, which will make it easier for other PCs to see you. ■

Quick tip

If you're looking to take your network sharing to the next level, take a look at the GADMIN-SAMBA package, available through the Software Centre. This allows you to set up custom user profiles for sharing folders.

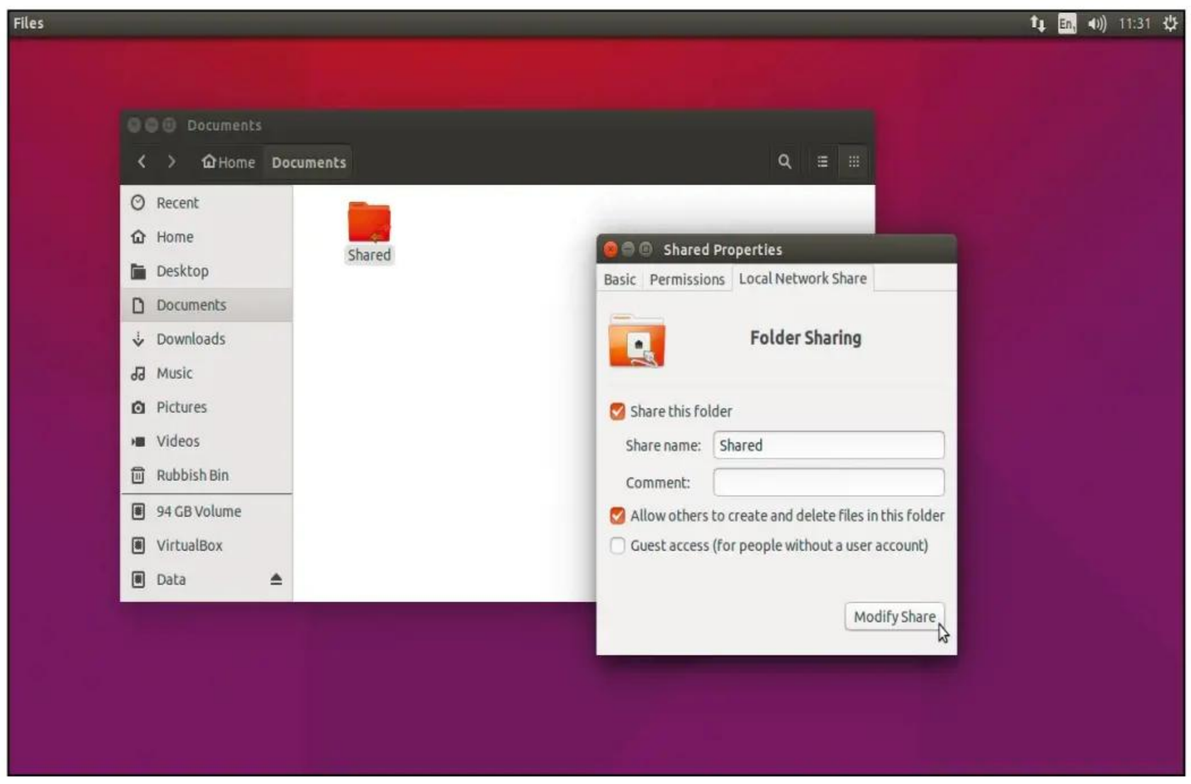

› Configuring your own folder shares is relatively simple, as long as you keep your ambitions in check somewhat.

Desktop: Install and switch DEs

Don't like your distro's desktop? No problem. Trade it in for an environment that's better suited to your needs.

Quick tip

Don't like the look of your new Xfce desktop? Ubuntu sets the default *GTK* theme for Xfce, which means missing icons and ugly looking panels. Change this to the Xfce style and Tango icons by selecting Applications Menu > Settings > Appearance.

One of Linux's great strengths is its versatility, and unlike other operating systems that ability to customise extends to your desktop environment (DE). Each distribution (distro) of Linux ships with its own DE, and while they share superficial similarities – specifically the point-and-click WIMP (Windows, Icons, Menus, Pointer) interface – there can be some quite radical differences in functionality as well as their look and feel.

The good news is that you're not locked into your distro's choice of desktop. Whether you're planning to switch to a different distro, but are put off by the desktop that it employs, or you're looking for a fresh way to interact with your desktop, Linux can accommodate you. In fact, you can even run multiple desktops side-by-side, allowing you to compare and contrast or use different desktops for different purposes.

Desktops are distributed as software packages, containing all the disparate elements required to make that desktop function (*see the Anatomy of a Desktop box, below*). Once desktops have been installed, you switch between them from the login screen by clicking the desktop icon that appears in the logon box next to your username.

So, if you're itching to try out some alternatives to your default distro, we'd say you can't go wrong starting with Cinnamon (**http://cinnamon.linuxmint.com**). It's the default desktop in Linux Mint, and is particularly popular with those switching from Windows as it utilises a similar style of desktop.

Spice things up

Cinnamon is highly configurable. Its main customisable elements – known as 'spices' – are its themes, applets, desklets (widgets you can place on the desktop) and extensions. These allow you to radically alter its capabilities as well as its look and feel – the step-by-step guide (*see bottom, p109*) reveals how to tweak many of these settings.

Four other desktop environments worth looking at include Xfce (**www.xfce.org**), Gnome (**www.gnome.org**), LXDE (**www.lxde.org**), and KDE Plasma 5 (**www.kde.org**). Even if you wanted to use Ubuntu's default desktop, Unity, isn't available as a standalone desktop for use in other systems.

There are lots of things to consider when choosing your desktop. How it looks is an obvious starting point, but does that translate into making it easier for you to use? Some desktops, eg, aren't as configurable as Cinnamon, so it's important you explore all the options there are before deciding if a desktop has the functionality and flexibility you

The anatomy of a desktop

Your desktop is the virtual environment in which you work, and provides a friendlier user interface with which to interact with your PC. The desktop consists of a number of different elements to produce the working environment you know and love – or hate, if you're switching.

The desktop typically contains icons, windows, toolbars, folders, wallpapers and – optionally – desktop widgets. Most of what you'll see is provided by the window manager. As its name implies it determines how the desktop's windowing system works by taking responsibility for the placement and appearance of windows and their component parts, such as menus, title bars and control buttons. The desktop's graphical elements – buttons, scrollbars, icons etc – are stored in special libraries. These include the widget toolkit, which is also utilised by applications so they can work seamlessly with the desktop. Two main toolkits exist: *Qt* and *GTK*, and while you can run applications made in one toolkit on a desktop built using another, they don't tend to look as good.

Many of these elements are stored in a theme, which makes it easy for users to change a desktop's look and appearance by defining how key elements look in terms of shape, colour and other elements.

Desktops also ship with a number of core tools and utilities, designed specifically to work well with that environment. Typical elements include a file manager, image viewer, text editor and terminal emulator, as well as a tweak tool that allows various aspects of the desktop to be customised.

It doesn't matter how sophisticated your DE is, they all share the same building blocks: toolbars, desktop, widget and windows etc.

Remove unwanted desktops

You can install as many desktop environments as you like – aside from taking up space on your hard drive, they don't consume any resources unless you're actually running one.

But there may come a time when you've finished experimenting with desktops, settled on your favourite (or favourites) and want to remove the rest. It goes without saying you shouldn't be logged into the DE that you're planning to remove.

With that prerequisite out of the way, removing the desktop is best done from the Terminal, particularly if you plan to remove all the software the DE installed alongside itself. Here's the two commands you'll need:

```
$ sudo apt-get remove packagename
$ sudo apt-get autoremove
```

Replace 'packagename' with the desktop's package name, which will be one of the following: **cinnamon**, **plasma-desktop** (KDE Plasma 5), **gnome**, **lxde** or **xfce4** (Xfce).

need in a package you're happy to look at. Despite their differences, you'll notice similarities between many Linux DEs. That's because they are basically variants of Gnome. Unity is the obvious example here, sharing many of its elements, but even the likes of Cinnamon betray their roots through the use of similar applications, built using the *GTK* toolkit that Gnome employs. The big rival to Gnome is KDE, and its differences extend beyond the fact it employs the *Qt* toolkit rather than *GTK*. For starters, it's one of the most configurable tools out there, which makes it more popular with advanced users.

There's also a question of performance and power consumption. Some desktops are more resource hungry than others, eg Unity and KDE Plasma 5. The increased demands do come with benefits, such as: a wider feature set, a bigger range of tools and flashier looks, which are missing from lightweight alternatives such as LXDE. But that means they're less suitable for older and low-powered machines.

The best thing to do is run a couple of them and see how they perform – you should notice those with a smaller footprint are the most responsive, particularly on slow PCs. They also place fewer demands on power consumption, which should benefit those running on batteries. Remember, thanks to the fact you can easily switch between desktops, you can mix and match – a low-powered desktop while on the road, with a more fully featured environment for when plugged into the mains.

You're ready to start experimenting with desktops, so how do you install them? In some cases you can simply search for the desktop using your distro's software centre, but for others you may have to manually add the repositories (repos) yourself before installing through Terminal.

In the case of Ubuntu, you'll find Xfce (type `xfce4` into the Search box to locate it), Gnome and LXDE are all available through the *Software Center*. Click 'More Info' when you've found one – because they usually ship with optional components added in, you might want to review these before clicking the 'Install' button. Cinnamon and KDE Plasma 5 require a trip to the Terminal. In the case of Cinnamon, enter:

```
$ sudo add-apt-repository ppa:lestcape/cinnamon
$ sudo apt-get update && sudo apt-get install cinnamon
```

For KDE Plasma 5, use these commands:

```
$ sudo add-apt-repository ppa:kubuntu-ppa/backports
$ sudo apt-get update && sudo apt-get install plasma-desktop
```

Working with desktops

Once you've installed a new desktop, log out of your account. When you find yourself back at the login screen you should see an icon next to your username. Click this and you'll see a list of available desktops appear – in some cases you'll see multiple options based on a single desktop. Select one, then enter your password and log in as normal to start using it.

When you log into a new desktop for the first time, none of your existing preferences will come across from your old DE (but don't worry, they're all still there, ready and waiting for you the next time you log back into your original desktop). One of the benefits of having to make changes to your new desktop early on is that it forces you to explore, giving you insight into how it works and what tweaks you can perform. ■

Customise your new desktop

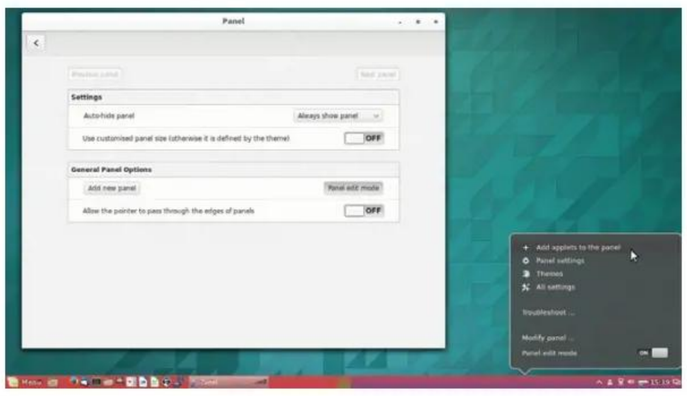

1 Tweak the taskbar

The Cinnamon's taskbar is called the panel and is split into three sections: menu button, app shortcuts and applets. You can pin app shortcuts to it by clicking 'Menu', right-clicking the app shortcut and choosing 'Pin to panel'. Right-clicking the panel reveals more options, like adding new applets or adding a second panel (choose 'Modify panel...').

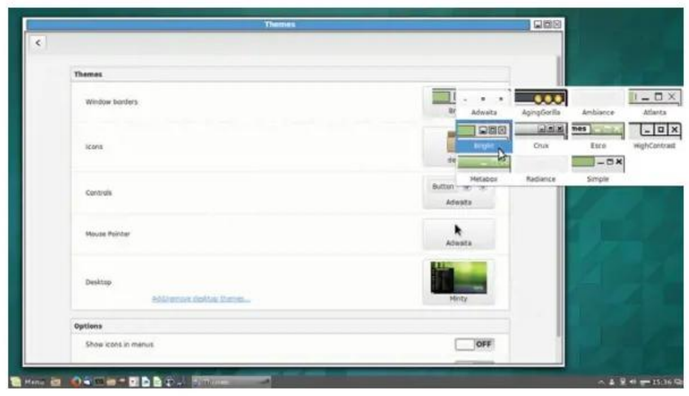

2 Customise look and feel

Click the 'Menu' button and select the 'System Settings' button to change the way Cinnamon looks. In the Appearance section, you can change the background and system font (you'll see options from other desktops you've installed too). Selecting Themes and clicking 'Add/remove desktop themes...' you can download more themes to pick from.

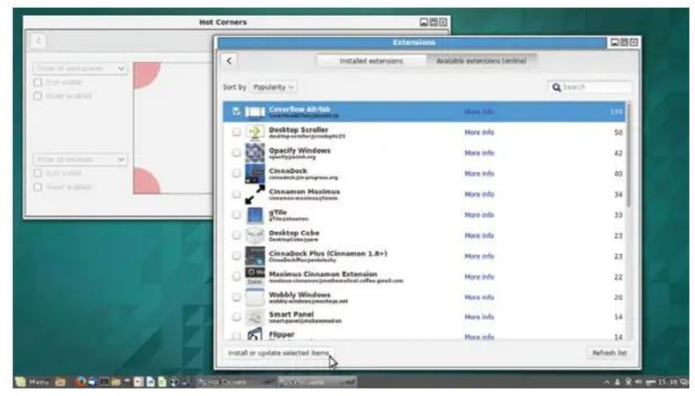

3 Change behaviour

Hot Corners enable you to specify actions that occur when your mouse goes into a corner of the screen, such as showing the desktop or displaying all workspaces. Windows is where you can edit the way windows behave and alter the display of the task switcher. Finally, the Extensions allows you to customise using third-party plugins.

Remote desktop: An easy guide

You're over here, but the Windows PC you want is over there. What to do? Well, we have a solution: a remote desktop connection.

Quick tip

Another tried-and-tested remote desktop solution is VNC, available on all major platforms and through a number of different programs – try *TightVNC* first.

One is never enough. These days most people – especially Linux-using enthusiasts – are likely to have more than one PC in their homes. Perhaps they even have home servers, media boxes, laptops and more. The annoyance is when you want to use more than one at once. Sure, you can pop a laptop next to your desktop. But what if you're slumped in front of the TV and need something from the desktop upstairs?

Enter remote desktop. This comes in a number of forms, but the basic concept is that it enables one remote PC to connect to another PC, over a network (even the internet) and take control of it. It's as if you were sat in front of the remote machine. Sound crazy? Gamers livestream their gameplay – effectively, remote desktop is just the same.

Isn't this a huge security risk? Well yes, if you don't use strong passwords and limit access. But we'll explain ways to minimise the risks. The first of which is not to open access to remote desktop over the internet without ensuring all the relevant user accounts are locked down. See the boxout for an option that we think is more secure and easier to set up.

For accessing PCs on your local home network we're going to use something called *rdesktop* [*could that stand for remote desktop, possibly? – Ed*]. Open a terminal and type:

```
sudo apt-get install rdesktop
```

We're ready to connect to your Windows PC. The client is sorted – the server needs work. Ensure your Windows box is ready to accept remote access. On the Windows PC log in, locate the Computer Start Menu entry, right-click on this and select Properties > Remote Access > Allow Connection.

› Viewing your Windows desktop from a Linux box is easy. Why get off your rear?

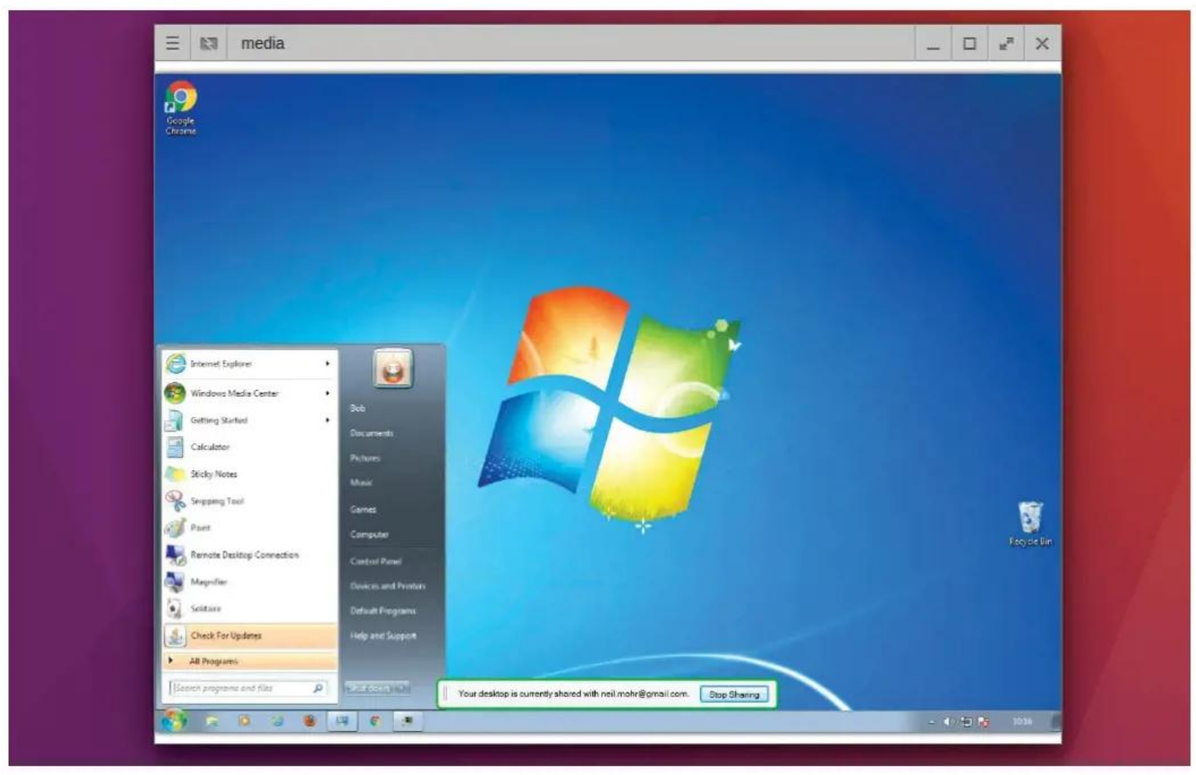

It's possible to tell Windows who can have access. Windows also uses user accounts and it's most sensible to create an account just for remote access with restricted access. However, it's likely you main account has administrator level access, so will by default have the relevant privileges for remote access. You can use the Select Users button to add additional accounts here.

To give a standard user Remote Access privileges select Start > Computer > Manage > Local Users and Groups > Users. A list of existing users is shown on the right. Double-click on the relevant user and select Member Of > Add > Advanced > Find Now and in the list select Remote Desktop User > OK > OK > OK. That user will now have remote access.

You'll also need to know the IP address of the Windows system. One way to find this is to open the Start Menu, type `CMD` and click OK. Type `ipconfig` and press return. Look up the IPv4 Address line.

We're now ready to connect from Linux to Windows remotely. Using *rdesktop* from the terminal takes the basic form of:

```
rdesktop <ip address>
```

However, it's useful to add a couple of options to include screen resolution through `-g <WxH>` plus a username with `-u "username"`. It's also worth using `-Pz`, which turns on bitmap caching and RDP compression. If you want better colour reproduction adding `-a 32` will attempt to force 32-bit colour, though it's likely this will fallback to 24-bit and frankly `-a 16` is fine. So putting this all together looks like this:

```
rdesktop -u "username" -g 1280x720 -a 32 -Pz 192.168.xxx.xxx
```

You can even redirect the sound, clipboard and other peripherals using the `-r` option. So for sound add `-r sound:local`. Getting technical, *rdesktop* is using the Windows RDP (Remote Desktop Protocol); by default Linux has no specific remote desktop system installed. As we now know how *rdesktop* works, we might as well stick with that when it comes to connecting two Linux machines together. Before a Linux PC can be connected to, you need to install an RDP server. To do this in a terminal type:

```
sudo apt-get install xrdp
sudo apt-get install lubuntu-desktop
echo lxsession -s Lubuntu -e LXDE > ~/.xsession
```

On Ubuntu this is something of a fudge, as the default 3D Unity desktop won't work with most remote desktop systems. To get around this we've installed an alternative desktop – this will look a little different, but it's the same system underneath (*for more on alternative desktop environments,*

Using Google Remote Desktop

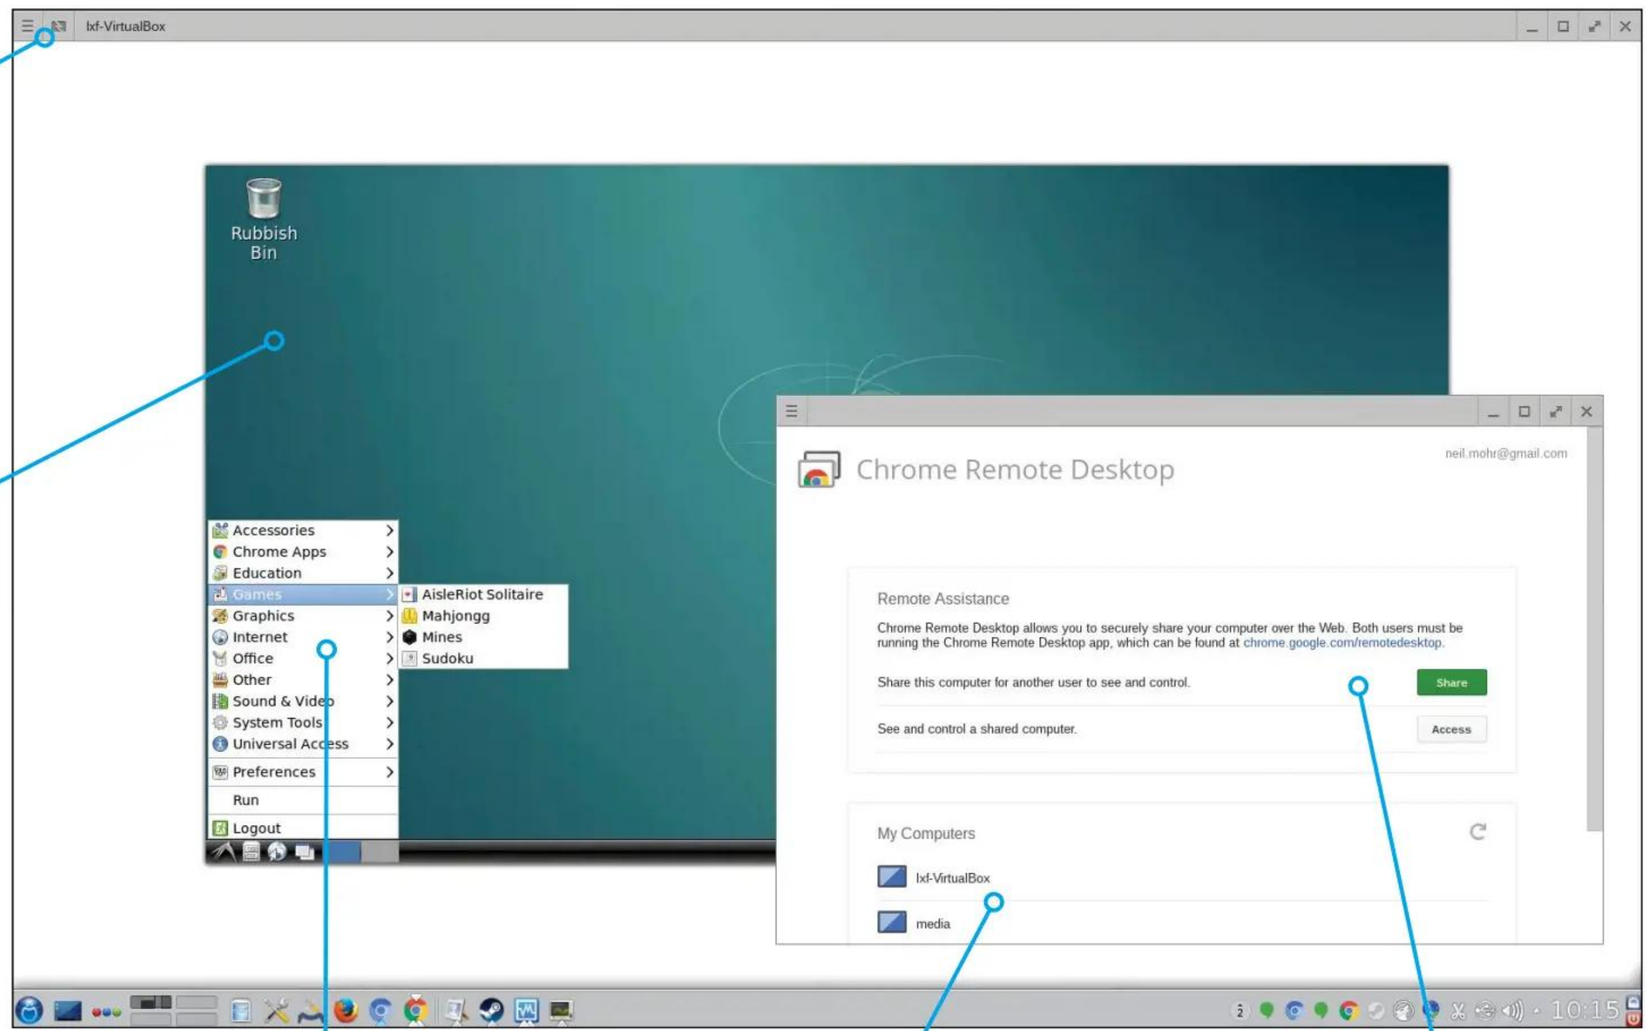

Close session
You can just close the window to end your sharing sessions, but it's better to click the official disconnection button here.

Remote session
Here you can see an in-progress remote session. We're accessing a Linux box using the replacement LXDE via Lubuntu.

Speed and quality
Be aware the upload speed on the remote session side dictates how quickly the screen can be updated. A slow connection will restrict the speed of the screen refresh/update.

Connect panel
A list of available systems appears here. If they're unavailable they'll be greyed out with a date and time of when they were last seen.

One-time share
You can request a lone-time connection using the Share button. A code is generated that the remote person types in after clicking the Access button.

check out page 102). At this point as before you can connect to this Linux box using its IP address. To find out the IP address, on the target PC, in a terminal window, you can type `ifconfig | grep inet` – the number after `inet addr:` is its IP, usually starting 192.168.

Just as when connecting to a Windows box, the *rdesktop* command is exactly the same for a Linux box:

```
rdesktop -u "username" -g 1280x720 -a 16 -Pz 192.168.xxx.xxx
```

You machine may feel different as it's using an alternative desktop environment, but your files and programs are all still there. If you feel the need there is a graphical interface for *rdesktop* called *grdesktop*, which will provide you with an interface similar to the Windows RDP program. You'll still need to know the IP address, username and password to connect, it just makes some of the options more obvious.

Google Desktop

The problem with enabling remote desktop over the internet is that it suddenly becomes possible for any old hacker to probe your internet facing ports for the RDP protocol. If they can spot an open port and brute force a weak password, they'll have full access to your desktop just as you would. That would be bad. Very bad.

If you require remote access over the internet then consider using the web browser *Google Chrome* with the Chrome Remote Desktop extension, available for free from its Web Store. This works on all systems that can run *Chrome*; you're able to connect between Windows, Mac and Linux PCs.

Install *Google Chrome* from **www.chrome.com** – on Linux ensure you get *Chrome* and not the open source version *Chromium* – switch to its Web Store and install the free Chrome Remote Desktop extension. This needs to be installed on each machine you want to access and to get access from.

On the target machine, open the extension through the Web App folder, open a blank tab and click the top-left icon. It requires you to have a Google account, and you'll need to log in on each machine, install the required remote desktop software using the 'enable remote connections' button and create a PIN to get access. We'd suggest you make this pin the length of a telephone number (but don't use a number that can be connected to you!) and protect your Google Account with two factor authentication, which requires a smartphone app to confirm your identity.

On Linux there are issues again with the Ubuntu Unity desktop, so as before we'll install an alternative desktop called Lubuntu and create a setting so the remote session knows which desktop to log into.

```
sudo apt-get install lubuntu-desktop
echo exec /usr/bin/startlxde > ~/.chrome-remote-desktop-session
```

Once set you can now connect from any other *Chrome*-running system with your account, even over the internet. Be aware, of course, that the remote system needs to be powered up and set not to fall asleep! Otherwise the Google system takes care of the complications of locating a PC on the internet for you. Handy. ■

VirtualBox: Virtualisation

Find out how virtualisation software can tap into your PC's unused processing power to help you run multiple operating systems.

Quick tip

To give your VMs a speed boost, enable VT-x/AMD-V acceleration. First, visit **http://bit.ly/1NFLGX2** to see if your processor is supported. If it is, make sure support is enabled in your PC's BIOS or UEFI – check your motherboard manual or website for instructions.

Today's multi-core PCs are built to run multiple tasks simultaneously, and what better way to tap into all that power than through virtualisation? Virtualisation, and in particular hardware virtualisation, is the process of splitting a single physical PC (known as the 'host') into multiple virtual PCs (referred to as 'guests'), each capable of working and acting independently of the other.

Virtualisation software allows the host to carve up its memory, processor, storage and other hardware resources in order to share individual parcels with one or more guests. If your PC is powerful enough, you can run multiple virtual machines in parallel, enabling you to effectively split your computer in two to perform different tasks without having to tie up multiple PCs.

Virtualisation isn't simply a means of dividing up computing power, though. It also enables you to easily run alternative operating systems in a safe, sandboxed environment – your guest PC can be isolated [*in theory – Ed*] from your host, making it safe to experiment with new software or simply try out a different flavour of Linux, for example. It can also be used for compatibility purposes – you may have switched from Windows, for instance, but want access to a virtual Windows machine to run old programs without having to use a dual-boot setup.

It goes without saying that the faster and more powerful your PC, the better equipped it is to run one or more virtual machines. That said, if performance isn't the be-all and end-all of your virtualisation experiments, then it's perfectly possible to run a single virtual machine in even relatively low-powered environments.

Choose VirtualBox

There are many virtualisation solutions available for Linux, but what better way to meet your needs (or even just dip your toes in the water) than with the open-source solution, *VirtualBox*? *VirtualBox* may be free, but it's still a powerful option that offers both a friendly graphical front-end for creating, launching and managing your virtual machines, plus a raft of command-line tools for those who need them.

An older version of *VirtualBox* is available through the *Ubuntu Software Centre*, but for the purposes of this tutorial we're going to focus on the newer version 5.x branch, which you can obtain from **www.virtualbox.org/wiki/Linux_Downloads**. You'll find that a variety of different builds exist, each one geared towards a specific distro (or distro version). Both 32-bit (i386) and 64-bit (AMD64) links are provided to downloadable and clickable Deb files, or you can follow the instructions provided to add the appropriate *VirtualBox* repository to your sources list.

Once it's installed, the quickest way to get started is to launch *VirtualBox* through the Dash. This opens the *Oracle VM VirtualBox Manager*, which is where all your virtual machines can be listed (and organised into groups). It's also where you create new VMs from scratch, but before you begin, select File > Preferences to change the default machine folder if you want to store your virtual machine settings somewhere other than your own home folder. This isn't a critical step, but as each guest may consume gigabytes of space for its own needs, you may prefer to choose a dedicated drive (or one with lots of free space). If you're looking to purchase a drive for your virtual machines, then consider an SSD to add zip to your VM's performance.

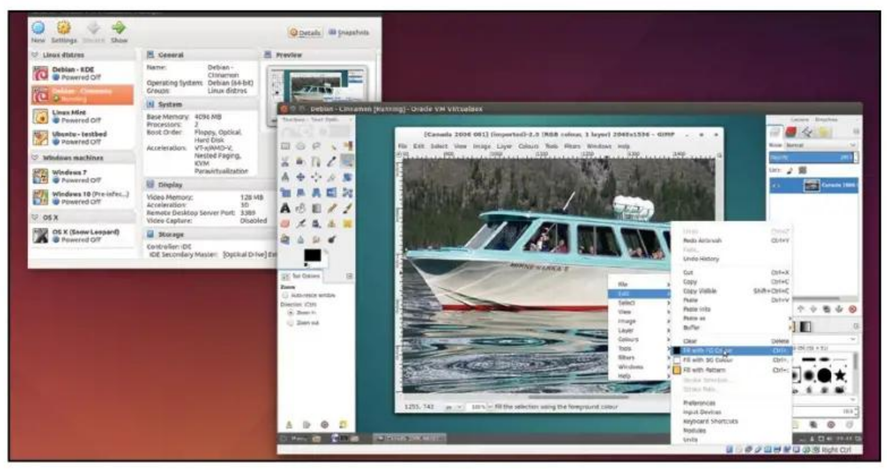

› ***VirtualBox*** **enables you to set up, manage and run multiple guest machines from the comfort of your desktop.**

Create your first VM

With your virtual machine folder set, click 'OK' and then click the 'New' button to create your first virtual machine. The Create Virtual Machine Wizard works in either of two ways, Guided or Expert, with the latter putting the three configuration steps in a single window. Start by selecting your chosen OS and version from the two drop-down menus – *VirtualBox* supports all the major OSes, including BSD, Solaris and IBM OS/2 in addition to Windows, OS X and – of course – Linux. The Version drop-down changes depending on your initial selection; all the major distros as well as Linux kernel versions from 2.2 onwards are available.

It's important to choose the right OS and version because this will ensure that other machine settings are set so they're compatible. You'll see this immediately when the 'Memory size' slider changes to match the OS. This will be set to a comfortable minimum setting, so feel free to alter it using the slider – it's colour-coded green, amber and red to help you set the memory to a level that's comfortable for your host PC. The figure you set is actual host RAM, not virtual memory, so

Headless setup

One way to maximise your host PC's resources is to run your virtual machine headless. This means there's no way of interacting with that VM on the host PC; instead, you access it remotely using the Remote Display Protocol (RDP). First, make sure you have the VirtualBox Extension Pack installed – this provides support for *VirtualBox's* implementation of RDP – then enable it on your VM via Settings > Display > Remote Display tab by ticking 'Enable Server'. You'll need to change the default port (3389) if you're setting up multiple VMs in this way – choose unique ports for each between 5000 and 5050.

Once it's configured, you can launch your VM from the Terminal via one of two commands:

```
VBoxHeadless --startvm <uuid|vmname>
```

```
VBoxManage startvm "VM name" --type headless
```

Alternatively, hold Shift as you click the VM in the VirtualBox Manager, and you'll be able to monitor its progress from the Preview window before switching to your remote computer.

When it comes to accessing your headless VM from another PC, the *rdesktop* client is built into most distros, but *VirtualBox* also ships with *rdesktop-vrdp*, which gives your guest access to any USB devices plugged into the PC you're sat at. Use the following command:

```
rdesktop-vrdp -r usb -a 16 -N 192.168.x.y:0000
```

Replace **.x.y** with your host PC's IP address, and **0000** with the port number you allocated (**3389** by default).

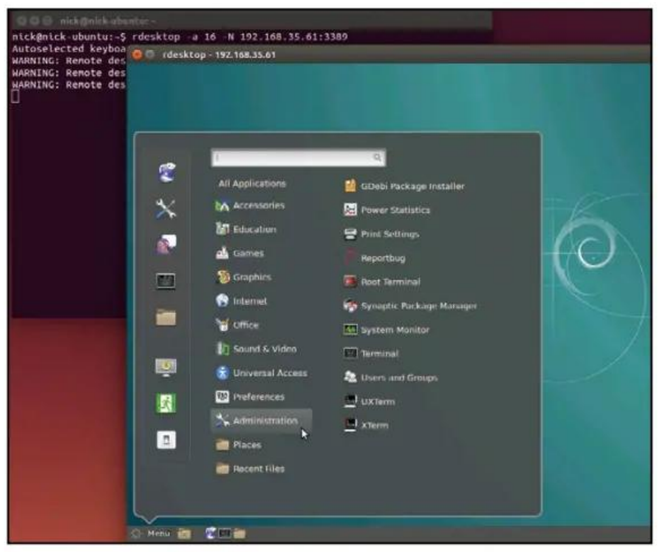

› Run your VM headless to cut resource usage if you plan to access it remotely.

be sure to leave enough for your PC's other tasks (including the running of *VirtualBox* itself).

The final option is to create a virtual hard disk. This basically starts out as a single file that represents your guest's hard drive, and will splinter off only when you start working with snapshots. In most cases, leave 'Create a virtual hard disk now' selected and click 'Create', at which point you'll need to set its size, location (click the little folder button to choose a different location from the default), file type and how the virtual file will behave. For these latter options, the defaults of 'VDI' and 'Dynamically allocated' usually work best; the latter ensures that the physical file containing your virtual hard drive's contents starts small and grows only as it's filled with data. Click 'Create' and your virtual machine is ready and waiting for action.

Virtual hardware tweaking

It's tempting to dive straight in and start using your new virtual machine, but while the basic hardware settings are in place, you should take the time to ensure it has all the power and resources it needs to function as you want it to. You can always tweak these settings later, but the best time to set it up is before you begin.

Select your new virtual machine and click the 'Settings' button. Switch to the System tab, where you'll find three tabs: Motherboard, Processor and Acceleration. You can tweak your VM's base memory from the Motherboard tab, as well as switch chipset, although unless you need PCI Express support the default PIIX3 should be fine in most cases. The Pointing Device is set to 'USB Tablet' by default, but there's a 'PS/2 Mouse' option for legacy purposes.

The Extended Features section should already be set up according to the OS you've chosen, but if you'd like your virtual machine to have a UEFI rather than a BIOS, tick 'Enable EFI' here. Note, however, that this works only for Linux and OS X; Windows guests aren't (yet) supported.

If you have a multi-core CPU installed, switch to the Processor tab to allocate more than a single core to your VM, making sure you don't attempt to allocate more cores than your processor physically possesses (Hyperthreading should be discounted). You may also need to tick 'Enable PAE/NX' if your virtual machine needs access to more than 4GB of RAM on a host PC with an older 32-bit processor.

The Acceleration tab allows you to tap into the processor's virtualisation features if they exist – see the tip for details.

Other key settings

Switch to the Display tab to configure your virtual graphics card. Start by allocating as much memory as you think you'll need, and also tick the 'Enable 3D Acceleration' box to improve performance across all your VMs. If you're running a Windows virtual machine, then tick the 2D option too. Switch to the Remote Display tab if you'd like to access your VM remotely. The Video Capture tab makes it possible to record your VM screen as a video should you want to do so – the former feature requires the VirtualBox Extension Pack, which we'll talk about shortly.

The Storage tab is where you can configure the internal storage of your virtual PC – by default your virtual hard drive is added to the SATA controller, from where you can add more drives. You'll also see that a single DVD drive is also added to the IDE controller. Select it and click the little disc button next to the Optical Drive drop-down to select a physical drive or mount an ISO disk image as a virtual drive instead. Tick the 'Passthrough' option if you'd like to be able to write discs, play audio CDs or watch encrypted DVDs.

The options in the Audio and Serial Ports tabs are largely self-explanatory, but if you plan to make your guest VM visible »

Quick tip

Make use of the *VirtualBox* Manager's new Group feature to organise your VMs into user-defined categories: right-click the first VM in the list and choose 'Group'. Right-click the group header and choose 'Rename', then create new machines directly from this group or drag other guests into it to assign them to the group.

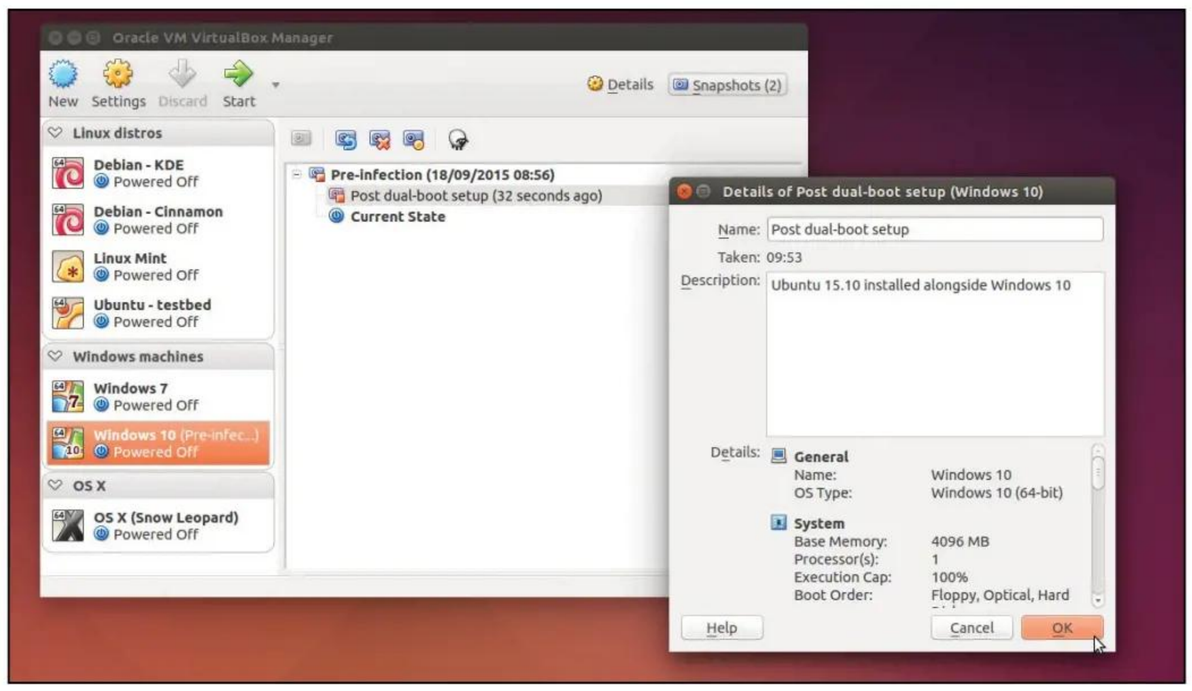

› The ability to take snapshots of your virtual machines makes them particularly suitable as test beds.

» over your local network for the purposes of sharing files and other resources, then select 'Network' and change the NAT setting to 'Bridged Adapter'. Other configurations are also available from here – 'NAT Network', eg, allows you to create a network of VMs that can see and interact with each other while remaining invisible to the host. NAT networks are configured independently via *VirtualBox's* File > Preferences menu (look under Network).

Working with USB peripherals

The USB tab is where you can capture specific USB devices for use in your VM. However, before you can use this feature, you need to make sure you add your username to the **vboxusers** group on your host PC using the following command in the Terminal:

```
sudo usermod -a -G vboxusers <username>
```

Once this is done, your USB devices will become visible to your *VirtualBox* guests. Note that *VirtualBox* supports only the older USB 1.1 implementation by default, but you can install the VirtualBox Extension Pack to add support for USB 2.0 and USB 3.0 among other extras (including PCI and host webcam passthrough). Download this Extension Pack from **www.virtualbox.org**, but note the licence restrictions: unlike *VirtualBox*, it's not open source and is free for 'personal evaluation' only.

You can easily connect to USB devices within your guest on the fly – click the USB button on the guest machine window and select your target peripheral from the list – but adding specific USB Device Filters here makes it possible to automatically capture specific devices when the VM boots. One example of where this could be handy is if you set up a VM as a headless TV server – it would allow the VM to take control of your USB TV stick the moment it starts. We cover the Shared Folders tab in the 'Share data' box below, while the User Interface tab allows you to specify which menu options are made available to this guest.

Quick tip

It's possible to port your virtual machines to different PCs – select File > Export Appliance to set up an archive in OVF (Open Virtualization Format) format, using the OVA extension to bundle everything into a single file. Be warned: it doesn't include snapshots and often changes the virtual hard disk from VDI to VMDK format.

Your first boot

With your VM's hardware set up, you're ready to go. You need to point your virtual CD/DVD drive towards an ISO file (or physical disc) containing the installer of the OS you wish to emulate, then start the VM and follow the prompts to get started. Once running, your virtual machine acts in exactly the same way your main PC does – click inside the main window and your mouse and keyboard may be 'captured' by the VM, allowing you to work inside it. To release these back to your host PC, press the right-hand Ctrl key.

Once you've installed your target OS in the guest machine you'll need to install the Guest Additions – a series of drivers and applications that enhance the VM's performance. Key additions include a better video driver supporting a wider range of resolutions and hardware acceleration, mouse pointer integration, which allows you to more easily move the mouse between host and VM without it being captured, and support for shared folders.

Installing these for Windows guests is as simple as selecting Devices > Insert Guest Additions CD image... After a short pause, the setup wizard should appear. Things are a bit more complicated for Linux guests – see chapter 4.2.2 under *VirtualBox's* Help > Contents menu for distro-by-distro guides. Once you've followed the prerequisites, open the file manager and browse to the root of the Guest Additions CD, then right-click inside the window and choose 'Open in Terminal'. Once the Terminal window opens, the following command should see the additions installed:

```
sudo sh ./VBoxLinuxAdditions.run
```

After rebooting you should be able to resize your VM window to the desired resolution simply by clicking and dragging on it – have the Displays panel open in your guest when you're doing this to verify the dimensions as you resize.

Share data

Getting data to and from your VM is a critical part of virtualisation, and *VirtualBox* makes this as simple as possible. The obvious way is to set up a bridged network as described earlier, then create shared folders with which you can swap data over your network, but there are other handy sharing tools provided too.

The Shared Folders feature works best with guests you don't want exposed to the wider network, and also allows you to make folders available from your host without sharing them on the network. Open your VM's settings and go to the Shared Folders tab and you can specify a folder on your host PC that's made available to your guest: click the plus ('+') button, select the folder you want to share and change its display name on your guest if necessary. You can also elect to make the folder read-only to the guest, have it mount automatically when the VM starts and, last but not least, choose 'Make Permanent' to have the shared folder persist beyond the current VM session.

Open the Devices menu and you'll find two other ways of sharing too: Shared Clipboard allows you to share the contents of the clipboard between host and guest (this can be limited to one-way sharing, or made bi-directional). You can also implement Drag-and-Drop, another way to quickly share files between host and guest by dragging files into and out of the guest machine window.

› **Make life (and file-sharing) easy: you can configure *VirtualBox* to allow you to quickly transfer files to and from your guest using drag-and-drop.**

Take a snapshot

Your VM is now set up and ready for action. It should work in exactly the same way as any physical machine, but it has one crucial advantage: snapshots. Snapshots let you take one-click backups of your guest at a specific point in time. You can then proceed secure in the knowledge you can roll back to the snapshot and undo all the changes you've made since.

You can create snapshots while your machine is powered off, or during use – just select Machine > Take Snapshot to do so. Give your snapshot an identifiable name, and also add a description if you wish, then click 'OK'.

When you take a snapshot, *VirtualBox* starts recording changes to the drive in a different file. If you delete a snapshot, those changes are merged back into the main file, while if you roll back to an earlier snapshot (or the base image), the snapshot's changes are lost unless you create an additional snapshot when prompted. VMs support multiple snapshots, and you can even move between them, allowing you to create multiple setups from within a single guest.

Terminal use

VirtualBox's user interface may be a convenient way to get started with virtualisation, but once you're up and running you'll be pleased to learn there are a number of command-line tools you can employ if that works better for you. You can even bypass the graphical VirtualBox Manager entirely if you're willing to learn the rather lengthy list of sub-commands for the *VBoxManage* tool, such as `createvm` and `startvm`, but even if you're happy with the point-and-click approach, there are a number of tools you should take a closer look at.

The first is *VBoxSDL* – if you'd like to launch your VM in a 'pure', distraction-free environment (so none of the controls offered by the default VM window), this is the tool for you. Its usage is pretty straightforward:

```
VBoxSDL --startvm <vmname>
```

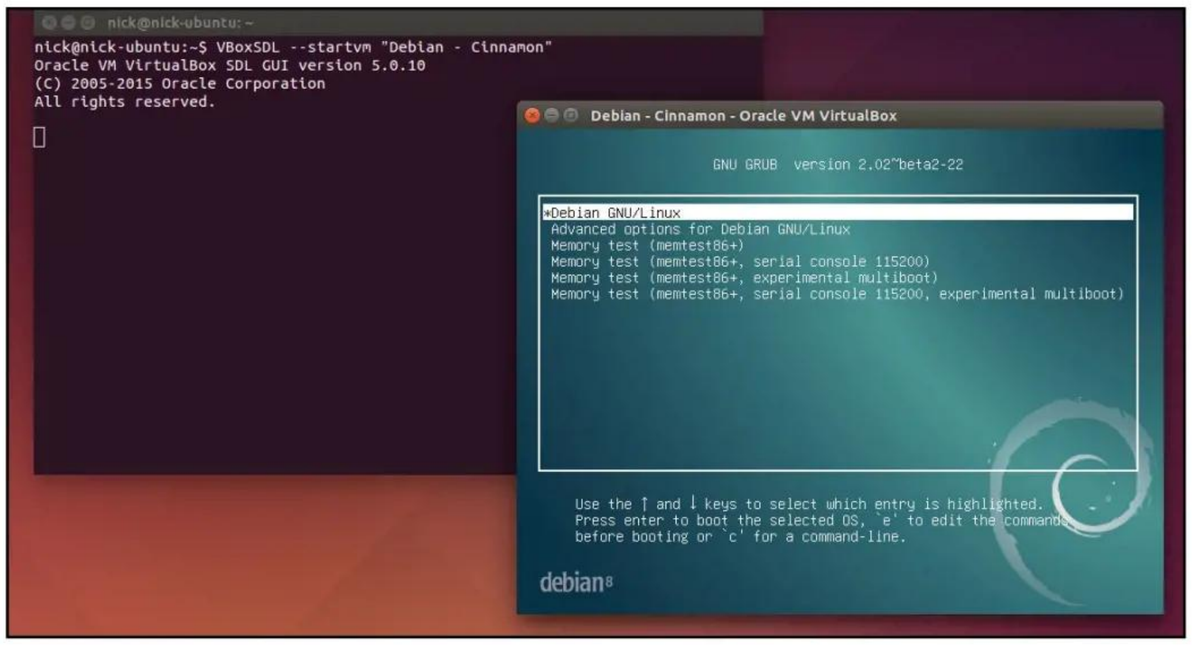

› **Remove all the desktop paraphernalia and run your guest in a lean, distraction-free window using *VBoxSDL*.**

Replace `<vmname>` with the name of your VM (or its UUID if you prefer). Once it's running, you'll not only have access to the menu commands offered by the main *VirtualBox* window, but some handy shortcuts you can employ while pressing the host key (the right [Ctrl] key by default): [f] toggles full-screen view on and off, while [n] takes a snapshot. Press [h] to press the ACPI power button, [p] to pause and resume, [q] to power off or [r] to reset. Finally, press [Del] in conjunction with the host key and you'll send a [Ctrl]+[Alt]+[Del] to the guest machine. Alternatively, shut down your VM using the *VBoxManage* tool – just type the following command to initiate the ACPI power button, eg:

```
VBoxManage controlvm "VM name" acpipowerbutton
```

Another handy command-line tool is *VBoxHeadless*, which enables you to run your virtual machine headless. To do this – and allow yourself to access it remotely from another computer (*check out our Headless setup box*).

Whether you plan to use *VirtualBox* from the command line or its GUI, you'll find it's packed with powerful and useful features that will convert you to the possibilities and power of virtualisation. You'll wonder how you ever coped before! ■

Extend the size of your VM drive

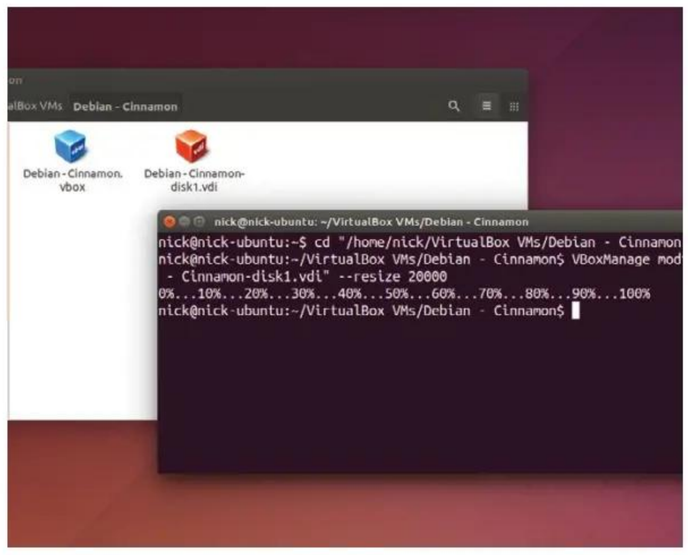

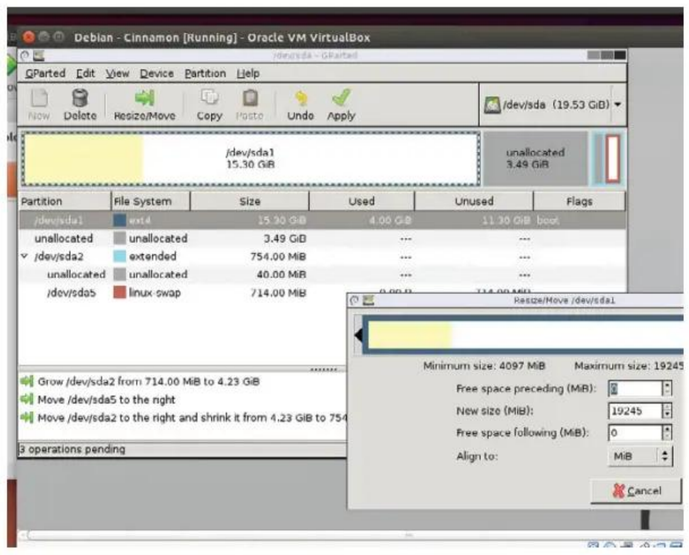

1 Consolidate snapshots

If your VM contains snapshots, the resizing process will affect only the original base image. To resolve this, right-click the VM and choose Settings, then append `-old` on to the end of its name. Click 'OK', right-click the VM again, but this time choose Clone. Click 'Expert Mode'', then rename it and verify that 'Full Clone' and 'Current machine state' are selected before clicking 'Clone'.

2 Resize virtual drive

Close *VirtualBox*, open Terminal and navigate to the folder containing your VDI file. Now type the following command, replacing **drivename.vdi** with the filename of your particular VDI file:

```
VBoxManage modifyhd "drivename.vdi" --resize 10000
```

The resize figure is in MB, so 10000 equals 10,000MB or 10GB.

3 Extend partition

The drive concerned has been resized, but you'll now need to repartition it. Boot your VM having attached an ISO of the *Gparted* Live CD and then use that to move partitions around to use the extra space – you may have to resize the extended partition first, then move the swap volume to the end before resizing the partition from the left to make the space available.

Rescatux: System repairs

All those commands feeling a bit complex? Repair common system problems without resorting to the command line.

Linux live CDs are a wonderful invention, they let you try new distros, show Linux off to your unenlightened friends, and fix broken systems. There are live distros aimed specifically at repairing damaged systems, but they have a common drawback. They all require a certain amount of expertise and most of us don't break our systems often enough to gain that sort of experience. When your computer sits there showing you nothing but a glum message from the bootloader, your main priority is fixing it as quickly as possible, not spending time using a search engine from a live CD to try and find the correct Grub incantation for your situation. I consider myself reasonably knowledgeable about bootloaders, but I still don't break them so often that I feel comfortable fixing them from the command line without at least a cursory RTFM to check my options.

Prep for live surgery

What we need is a live distro that is designed for fixing common problems without a great deal of background knowledge or research, and preferably one that doesn't require typing long commands where an error could make the situation worse. What we need is something like Rescatux.

Rescatux boots like a typical live CD to a lightweight LXDE desktop, but the window that opens on startup is the key difference. Called Rescapp, this is a one-stop centre for fixing various problems. Rather than simply list them, let's look at some of the problems that can arise when a computer system starts misbehaving at a low level, and how Rescatux can be used to fix them. This is not for dealing with minor user-level problems, a live CD such as Rescatux is usually brought out when things go seriously wrong.

Many system recovery operations require you to be booted from a live CD, either because normal booting is broken or because you need your root filesystem to be unmounted. You normally also need to use command line tools, and Rescatux provides all of this, but the Rescapp makes life much easier for many tasks.

When you press any of the operation buttons in Rescapp, it does not directly perform the operation. It displays a documentation page explaining how to use the option and, considering the low-level aspect of many of the operations, it's a good idea to read this. Then press the Run! button at the top right to perform the operation.

#1 Hard disk errors during boot

The first step when filesystem errors appear is to run **fsck** (that is short for filesystem check, not the expletive you use when you see the errors). This must be done while a filesystem is unmounted, hence the use of a live CD. Press the File System Check button. Rescapp temporarily mounts partitions in order to determine which distro they belong to. Of course, the corruption may prevent mounting (distro

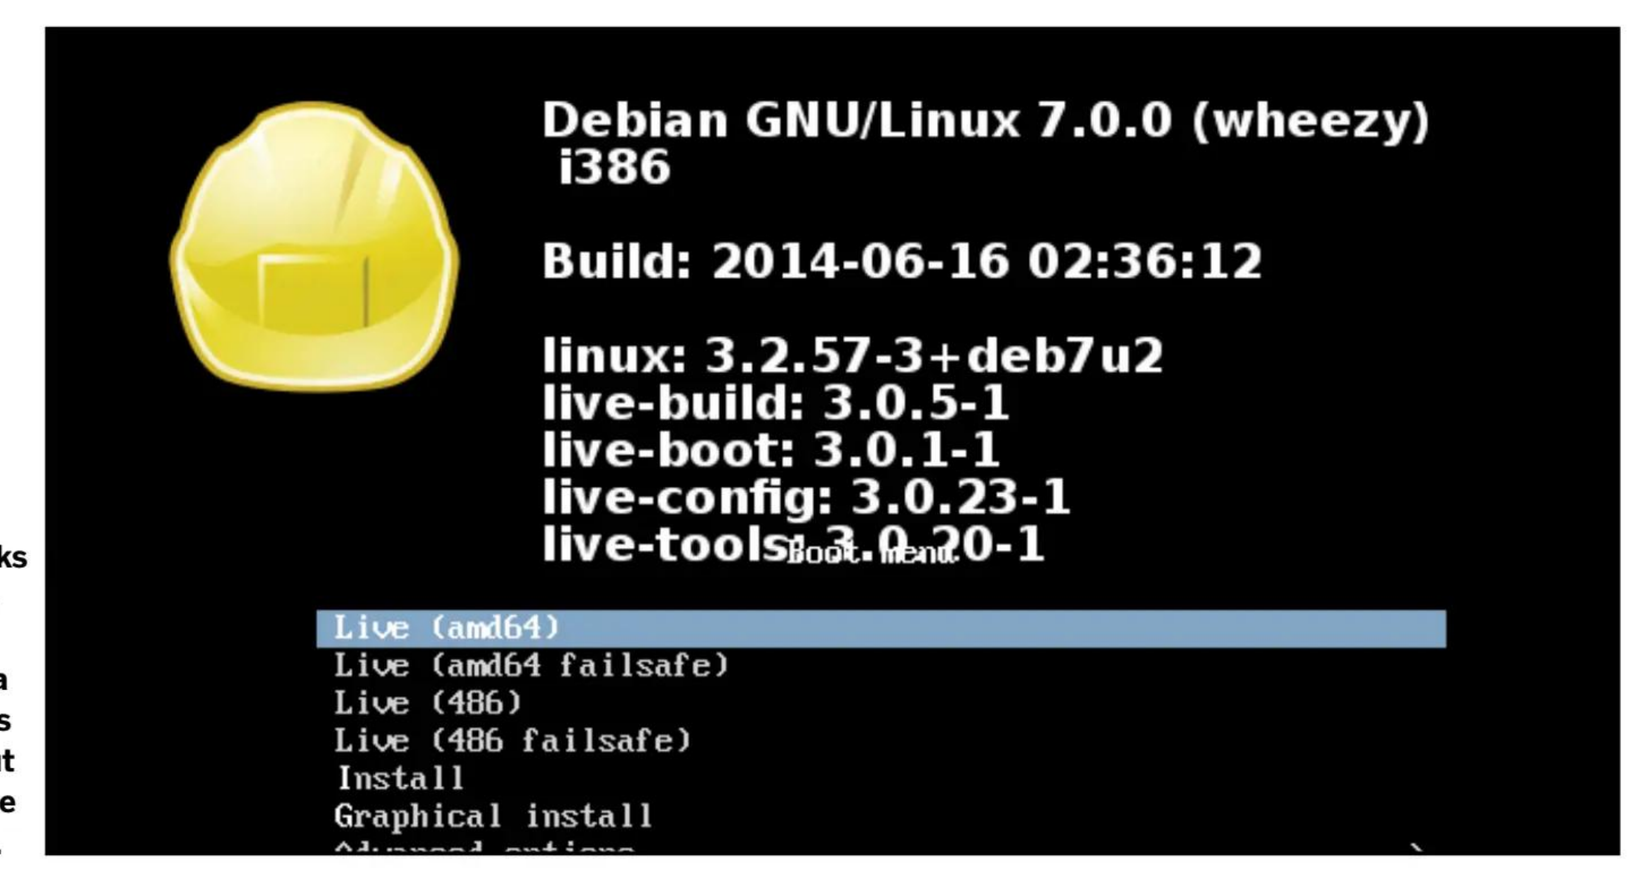

› Rescatux works with 32 and 64 bit systems. Booting 32 on a 64 bit system is usually safe, but not optimal. The reverse will fail.

Getting more help

While Rescapp is able to fix many problems with a couple of mouse clicks, sometimes you need more help. Rescapp has some tools to help with this; first of all there is a web browser to search for answers, possibly using error messages you received when trying to repair your system. Everything Rescapp does is logged, the Show Log button presents you with a list of log files – remember, this is a live CD so you will only see logs from the current session. You can also view the logs directly, they are saved in **rescapp/logs**. Looking through the log for the operation you are attempting may give information useful to you. If it does not help you understand the problem, it may help others, which is where the Share logs button comes in. After selecting a log to share, Rescapp will send the log to a pastebin on Debian's servers and give you the URL. Copy this somewhere safe and you can give it to anyone else for them to take a look at your logs. For a real time response, try the Chat button. This opens an IRC client to the **#rescatux** channel, where you can paste the URL of your pastebin and ask for help. You are not restricted to their own channel, of course, you could also try a channel dedicated to your distro for more specific advice. The 'Share log on forum' option works similarly, allowing you to ask for help on your favourite web forum.

But wait! Before you ask for help online, use the Boot Info Script button. This generates a file in **logs** containing information about your system, and you can share this with the 'Share log' option. Information about your system may be crucial to someone finding a solution to your problem.

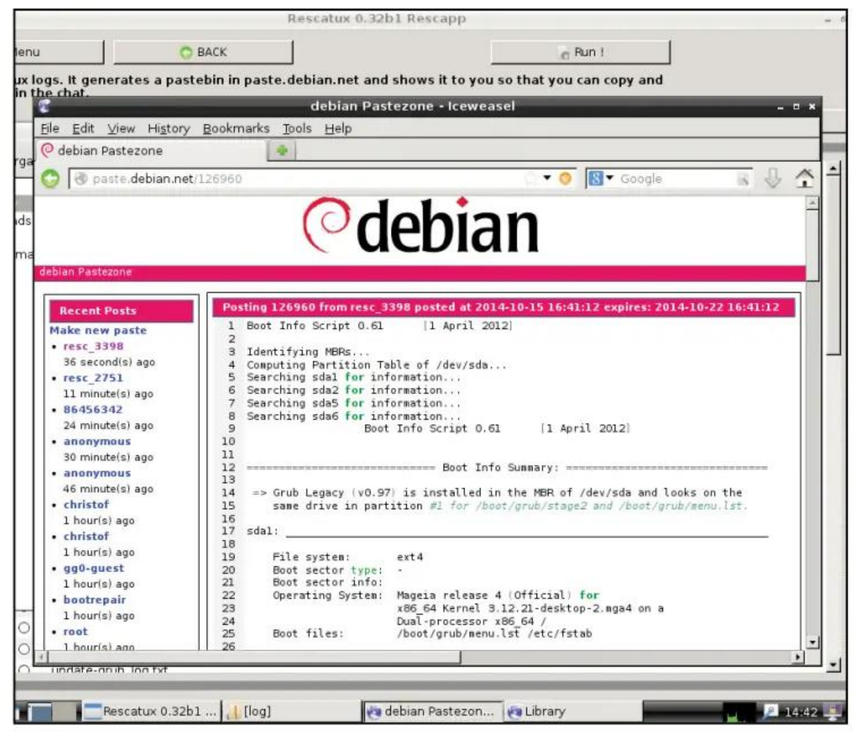

› The Share Log button sends a log file to a pastebin and gives you the URL so you can share it with anyone who wants to help you.

startup sequences are able to fix minor filesystem corruption transparently) so you may well be looking for one marked 'Cannot mount'. Only distro root directories can be identified, if you have a separate **home**, it will appear as 'Not detected' (or 'Cannot mount' if it is damaged). There may be other reasons for a partition being unmountable; it may be your swap partition or an extended partition, so choose carefully. If in doubt, the Boot Info Script log (covered later) lists your partitions and their types.

#2 My password is not recognised!

Aside from boot merely resulting in an unfriendly grub> prompt, this is one of the most scary moments of computer use. You checked that you typed it correctly and that the caps-lock is not on. You may have forgotten it or, on a new install, have mis-typed it on setup.

Resetting a password involves booting a live CD and messing around with **chroots** – you cannot simply edit a file – or you can use Rescapp. Press the 'Change Gnu/Linux password' button and, after reading the explanation, press Run!, pick the distro (there will always be at least two, your installed distro and Rescatux, which appears as Debian 7) and then select the user to change. Enter the new password and the job is done. Try not to forget this one! This button is only for Linux passwords. If you run a dual-boot system with Windows, there is a separate option to reset your Windows password.

#3 I deleted the wrong files

It is both amazing and terrifying how a simple typing mistake can cause so much damage. For example if you meant to type

```
rm -f *.txt
```

but typed

```
rm -f * .txt
```

instead. Or you wanted to reformat **/dev/sdc** but typed **sdb** instead. There is the odd filesystem-specific tool for recovering deleted files, such as **extundelete** – but a really determined typo can easily beat that, and it can't help if your partitions are gone. The first thing to do in such a situation is to stop writing to the disk – if the damage is on the partition containing your root filesystem you should shut down the computer with

```
sudo shutdown -n
```

This kills processes without using the usual **init** system, which reduces the number of disk writes. Now you can begin to boot Rescatux. If you partitioned a drive, you can use **testdisk** to search for the old partition boundaries and restore them. Repartitioning a drive only writes to the partition table, the actual data on the rest of the disk isn't touched until you format the new partitions. So if you can find the old partition boundaries and write them back into the partition table, everything should be back as it was. This is »

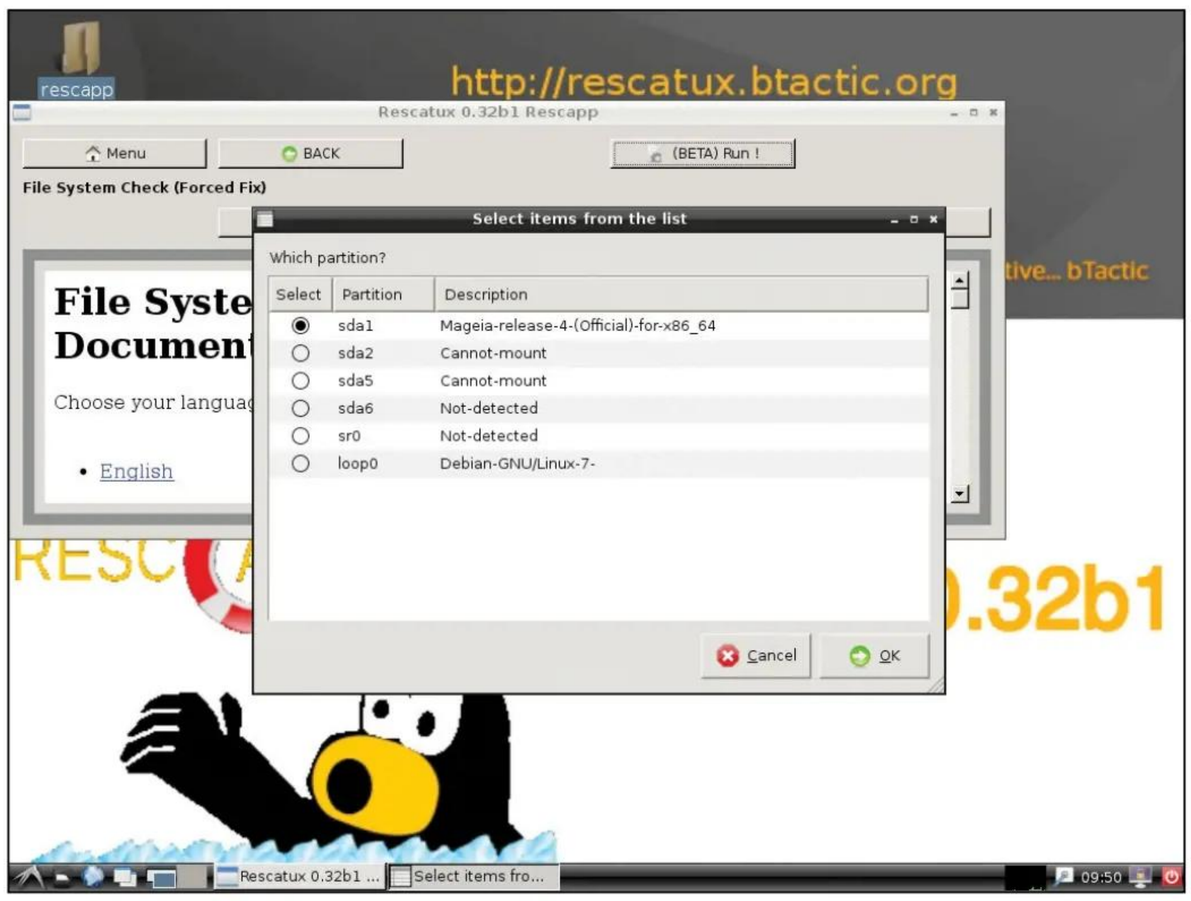

› If a filesystem needs repair, select it and Rescapp will do the rest. Cannot-mount may mean damage, but here it indicates swap and extended partitions.

» what **testdisk** does. After accepting the option to create a log file, that may be useful later, pick the disk to scan. The partition type should be Intel for the old-style MBR partition tables or EFI GPT for the newer GPT variant, the other choices are rather specialist. Removable drives occasionally use the None option, but they usually have a single partition with an MBR partition table. Select Analyse to scan the disk's partition table and then choose Deeper Scan to search for any lost partitions. If you find what you are looking for, Write will hopefully restore your lost settings. Although it is not mentioned in all of the menus, pressing the Q key usually takes you back to the previous menu. Testdisk is a very low-level tool, and its effects may not be reversible, where possible use **dd** to make a backup of your disk before proceeding with the process.

If you deleted files rather than partitions, the tool you want is PhotoRec. Photorec scans the disk for evidence of files and then attempts to reconstruct them – you will need another disk attached for saving these files to. Photorec can only find the contents of files, metadata such as ownerships, permissions and even the file name is not available to it.

So you end up with a lot of files with numeric names, although PhotoRec does give them a suitable extension based on the contents of the file. If the files are digital camera photos (PhotoRec was originally written to recover files from an erased memory card) or music files, you should find that any EXIF or ID3 tagging is preserved, making identification of the files relatively simple. Otherwise, you may have to spend some time trawling through the files to identify them, but that is better than losing your data altogether.

› **Forget your password? Rescapp lets you reset the password of any Linux or Windows users, including root, with a couple of mouse clicks.**

#4 I'm having problems with Windows

Go on, own up, some of you also use Windows, or have a "friend" who does. Rescapp also has options for repairing Windows systems, from resetting passwords to making users into administrators and other user management. It also has an option to restore the Windows MBR. The section on repairing Grub only applies if you still have at least one Linux distro installed. If you want to remove all Linux partitions from a drive, you will need to remove Grub from its boot sector and reinstall the Windows bootloader. Rescapp does this with the Restore Windows MBR button. Choose a disk, as with the Grub restore, and it will set up your hard disk to use the Windows bootloader.

#5 It's complicated

So far, we have looked at solutions to standard problems that can be deal with by a couple of mouse clicks. If things get more complicated, Rescatux contains much of the other software found on rescue discs, and you can just open a terminal and use it, but it will come as no surprise that you can also use more advanced tools from the Expert Tools section of Rescapp. These tools include:

Gparted – for (re)partitioning your hard disk.

Testdisk – to find partitions and filesystem on disks with a damaged partition table.

PhotoRec – to recover deleted or otherwise lost files from a disk, and not only photos.

OS Uninstaller – for removing extra distros from a multi-boot system

It is worth noting that the Expert Tools buttons still open a help page first, but this is generic help for Rescatux, not help for the individual programs.

Apart from those single-use programs, there is also Boot Repair which opens a window containing many options for altering the boot process.

This covers a number of operations, especially if you enable the Advanced options. It allows you to back up your partition tables and log files to a USB device, a wise step before you start changing things. Most of the other options in here let you tweak how the bootloader works. You can

Boot an ISO from Grub

Rescue discs are great, as long as you can find the thing when you really need it. You can copy an ISO image to a USB drive with **dd**

```
dd if=/lxfdvd/downloads/rescatux_cdrom_usb_hybrid_i386_amd64-486_0.32b1.iso of=/dev/sdX bs=4k
```

where **sdX** is your USB drive. That is more convenient, but Murphy's Law states that you won't be able to find the drive when you need it, so there is a more convenient option. To save you going through your pile of scratched up old DVD-Rs to find the one with Rescatux on it, here is how to boot it from your hard drive.

You need a working Grub2, so it is not an option in all cases, but if Grub and your boot partition are working it is more convenient. Put the ISO image (in this case we're using 0.32) into your **/boot** directory, then add the following to the bottom of the file **/etc/grub.d/40_custom** (do not modify the existing lines).

```
submenu "Rescatux 0.32" {
  set isofile=/Rescatux/rescatux_cdrom_usb_hybrid_i386_amd64-486_0.32b1.iso
  loopback loop $isofile

  menuentry "Rescatux 0.32 - 64 bit" {
    linux (loop)/live/vmlinuz1 findiso=$isofile boot=live config quiet splash
    initrd (loop)/live/initrd1.img
  }

  menuentry "Rescatux 0.32 - 32 bit" {
    linux (loop)/live/vmlinuz2 findiso=$isofile boot=live config quiet splash
    initrd (loop)/live/initrd2.img
  }
}
```

Now run **update-grub** or **grub-mkconfig**, or use the Update Grub Menus option in Rescapp, to update your menu. Then, when you reboot your system, you will have an option to boot into Rescatux.

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

If trying to boot from it gives an error that the ISO file could not be found, add this below the submenu line

```
set root='(hd0,1)'
```

where **hd0,1** refers to the first partition on the first disk (for some reason Grub counts disks from zero and partitions from one. Adjust to suit your system.

reinstall, as with the separate button covered elsewhere, but you can also change the location of the bootloader and the options it uses. How often have you searched for a solution to an issue only to be told to "add option xyz to Grub". You could go editing configuration files, but the Boot Repair window has a tab from which you can add various options without editing system critical files with the inherent risk of making things worse.

The current release of Rescatux is a beta and it does have some bugs. One or two of the options do not work properly, although we have not found anything dangerous. It's more a case of the option not doing anything rather than doing the wrong thing. We expect these bugs to be fixed in due course, but Rescatux is still worth keeping around as it can save a lot of heartache in many situations. ■

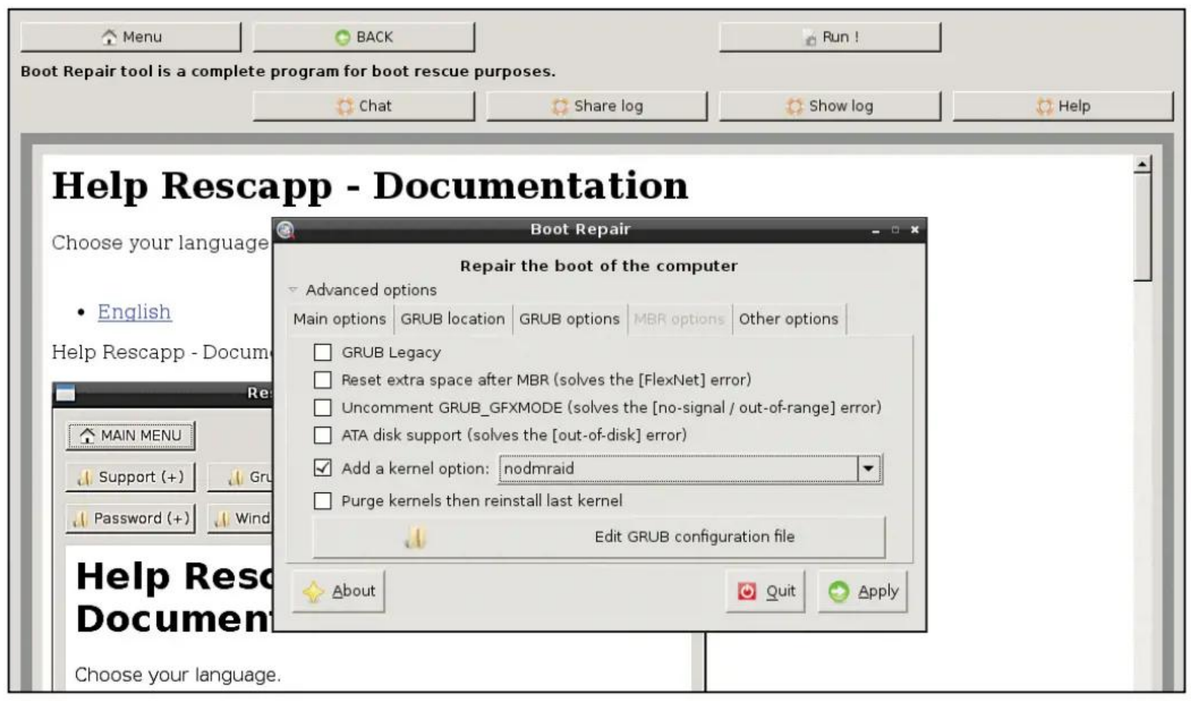

› **Tweak your Grub options as well as performing other Grub backup and repair operation from the Expert tools section.**

Fixing Grub

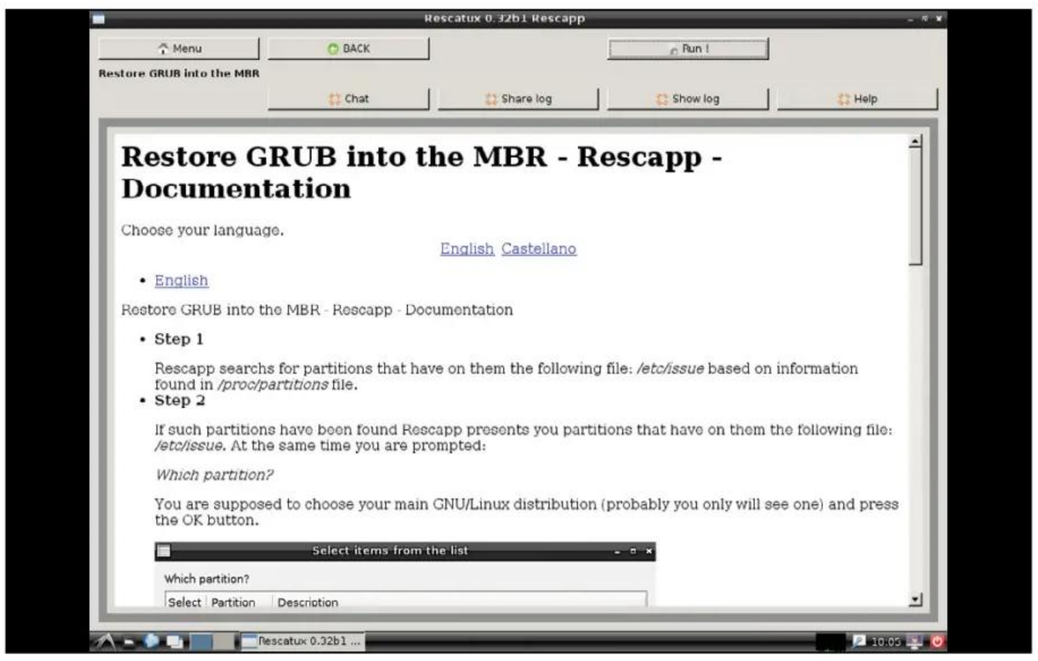

1 Restore Grub

Select Restore Grub from the main Rescapp window and read the help text to make sure you understand what is going on. Rescapp does most of the work, but you still need to know which distro and which disk you want to use for Grub. Press Run! when you are ready.

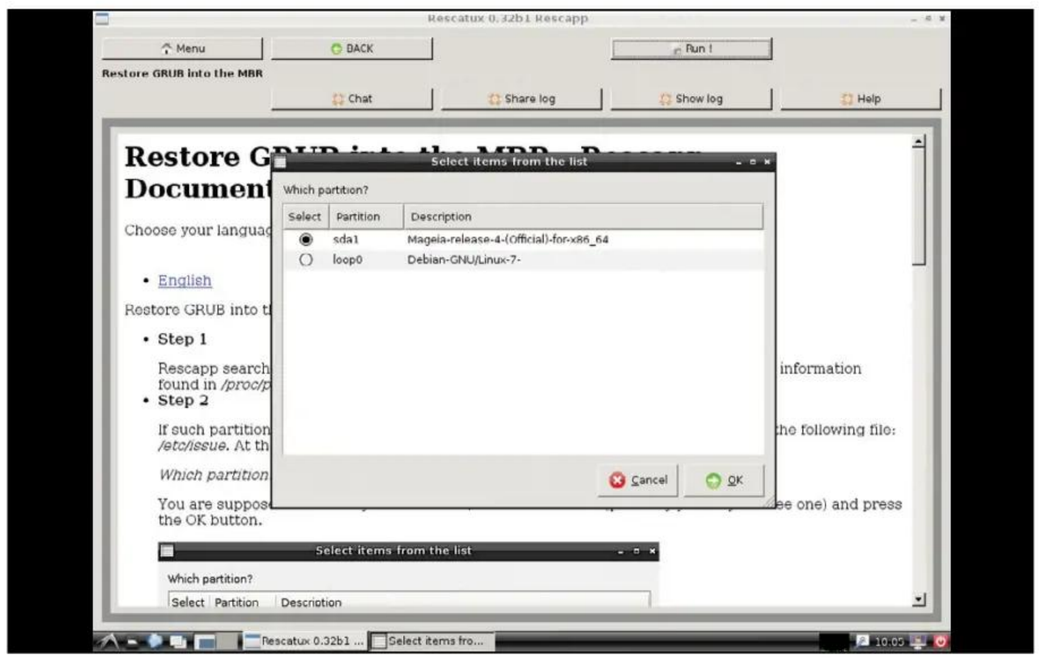

2 Select a distro

Rescapp will scan your disk partitions for installed distros, trying to recognise them by name when possible. Pick the one you want to use as your 'main' distro. This is the one that you are unlikely to want to remove, as that would involve repeating this process.

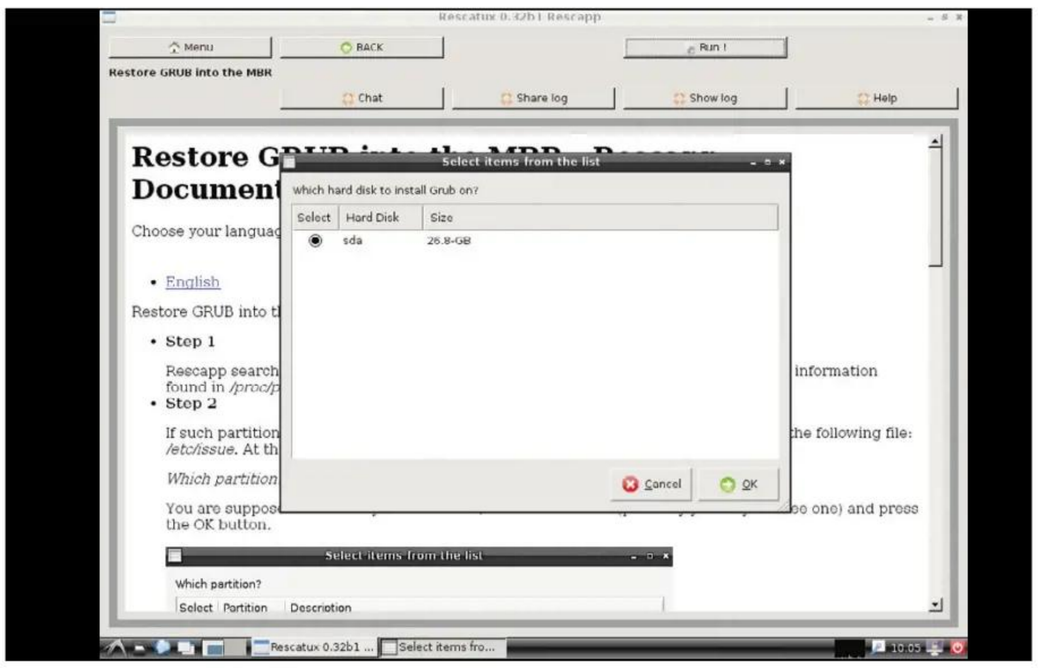

3 Grub selection

Now choose the disk to which you want to install Grub, usually sda. It doesn't have to be the one containing the distro you chose, but that is normally the case. If you boot from a USB stick, that may be recognised as sda with your hard disk on sdb.

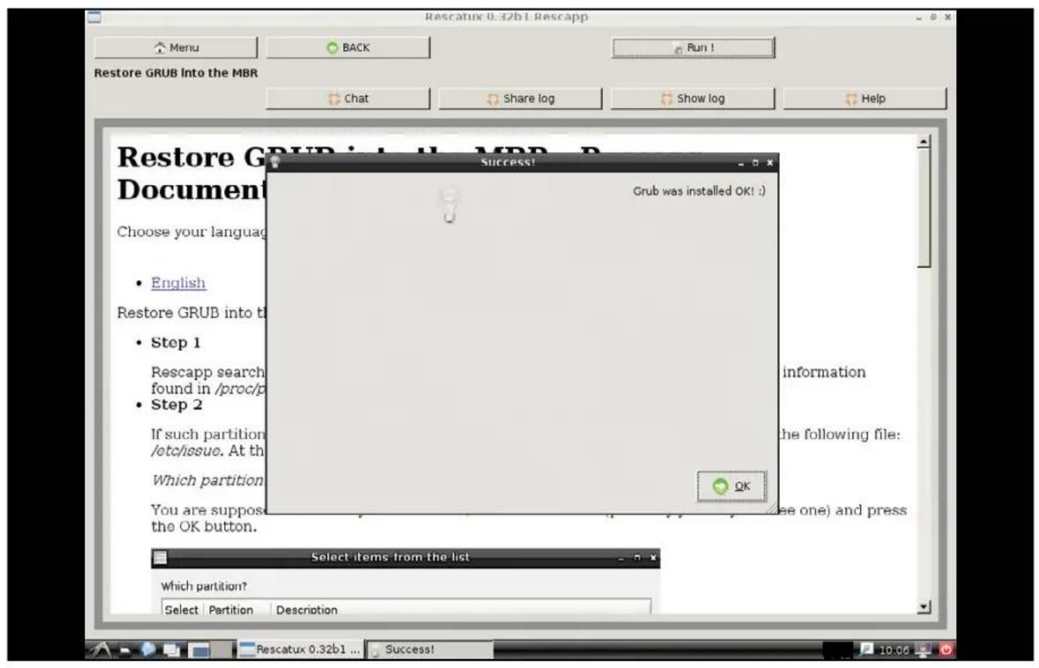

4 Auto-fix

Press OK, and Rescapp will perform the necessary steps, mounting and unmounting filesystems as needed and running **grub-install** with the correct **--boot-directory**, **--root-directory** and **device** arguments. Now you can reboot! It's as simple as that.

Ubuntu: Fix start-up issues

It's all gone wrong. Of course it has. Here's how to resolve issues with non-booting PCs without having to reach for the panic button (or a hammer).

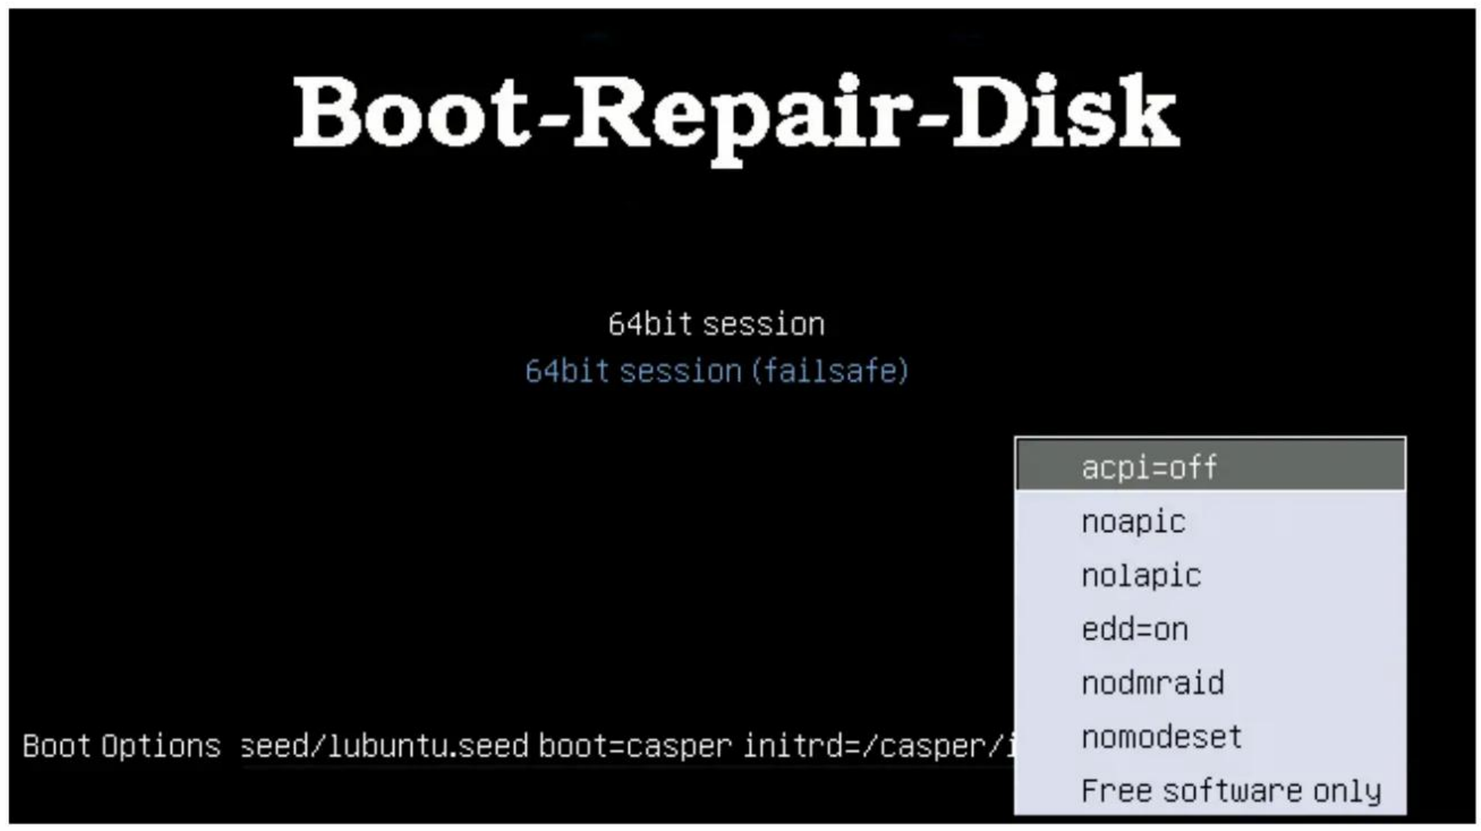

› **If you're struggling to get the Boot-Repair tool disc to boot, experiment with the various fail-safe boot options.**

Start-up problems. That moment when – having expected yourself to be getting on with your day's work or entertainment – you find yourself staring at a cryptic error message, or even worse, a blank screen. No matter how many times you press reset or restart, the same impenetrable barrier blocks your path. So, what can you do?

Start-up problems come in all shapes and sizes, and they can be difficult to track down. There are, however, some sound principles to use that will resolve many errors, and in this tutorial, we're going to look at the tools and techniques required to troubleshoot most start-up problems. You should start by examining how the boot process works (*see The Boot process box, p121*). This reveals that the boot process can be split into three broad stages centred around the *Grub 2* boot loader: pre-*Grub*, *Grub* and post-*Grub*. Knowing this allows you to focus your troubleshooting efforts based on where in the process the error or freeze occurs.

Let's start at the beginning. You switch on your PC. If power comes on, but nothing else happens, chances are you've a hardware issue to sort – if you recently poked around the innards of your PC, then check everything is connected correctly. If not, unplug all external devices except your keyboard and try again. If this doesn't work, open the case carefully and disconnect your internal drives too. If the computer now boots to the splash screen, you can try reconnecting the internal drives and trying again; if you're now able to boot to the login screen, shut down your PC and start reconnecting your external peripherals to see if the problem has cleared itself or can be targeted to a single device, in which case try a different cable, or go online and Google for known boot problems involving that device.

If you're lucky, your motherboard will emit a series of beeps or flashing lights you can use – again by enlisting the help of the internet – to identify the likely problem. This may involve replacing a component or something more drastic.

If you're able to get as far as your PC's splash screen, but then your computer hangs or a 'missing operating system' error message appears, then first think back to any recent changes. If you've overclocked your PC, eg, you should now enter the system EFI or BIOS and look for the option to load fail-safe defaults. Try rebooting again. If this fails, then the problem is likely to be with your hard drive, and so the first places to look are the MBR and *Grub*. If *Grub* isn't set to automatically appear when your PC starts, try rebooting while holding the Shift key or tapping Esc to bring up the *Grub* boot menu to confirm it's not able to even load itself. Jump to the *Boot-Repair* tool section once you've verified it's nowhere to be found. If *Grub* is able to load, but can't find any bootable OS you'll find yourself with a number of scenarios: you may be presented with a basic command prompt such as `grub>` or `grub rescue>`, which indicates one or more files required by *Grub* are missing or corrupt. You may get a specific error

The Boot process

When you press your system's power button, control is initially given to your PC's EFI or BIOS, which starts its various components, performs basic diagnostics tests and attempts to find a bootable device, which is typically the first hard drive. Once located, the BIOS or EFI looks for the Master Boot Record (MBR) at the very beginning of the drive, which has a tiny program inside that loads the next stage of the boot loader, reading a file (eg **e2fs_stage_1_5**), which in turn is able to load the *Grub* boot loader. A 'missing operating system' error at this point means you need your rescue disc for diagnostics as something is missing – either in *Grub*, the MBR or the drive itself.

Once *Grub* loads successfully it reads a file called **menu.lst**, which contains the list of choices you see in the boot menu. Each entry basically identifies the drive, partition and file that contains the Linux kernel, plus RAM disk file used by the kernel as it boots. The entry will also contain any additional parameters passed to the kernel.

Control is now passed to the kernel, which attempts to mount the root file system. This is a key moment, and if it fails you may get a kernel panic, or things might grind to a halt. If successful, it'll create a single process to run the **/sbin/upstart** file (other distros use init) – if this goes wrong, you'll get a panic, it may halt again or drop you into a root shell. At this point, upstart starts running scripts and upstart events to start other services and eventually bring you to the logon screen.

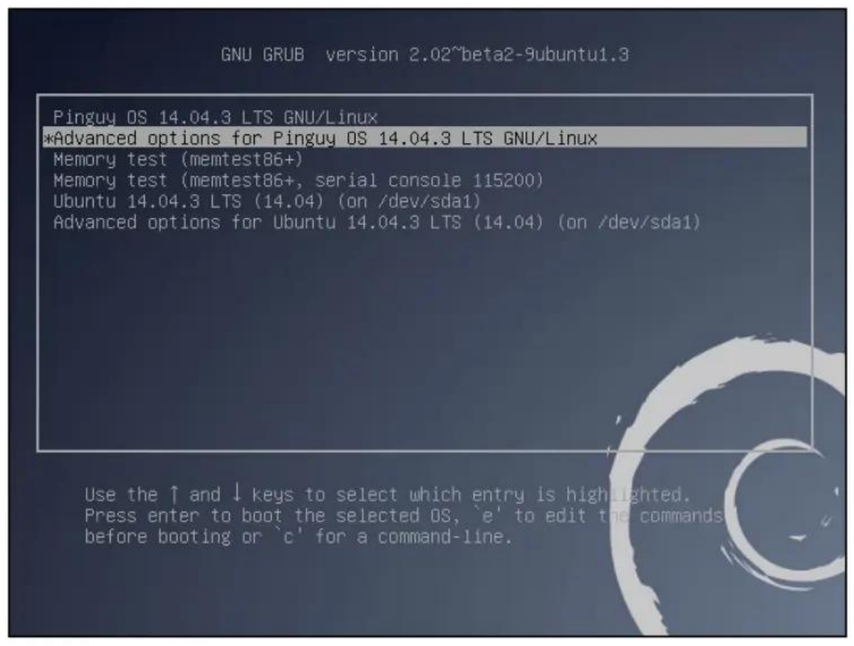

› The *Grub* menu's appearance is a critical point in the boot process – if your system gets this far your recovery options are greater.

message or frozen splash screen, or you may just see *Grub* and nothing else, indicating it can't even find the most basic information required to proceed.

If you press C you may be able to enter the *Grub* Terminal mode to perform basic checks and repairs – you can attempt to manually initiate the boot by pressing Ctrl+X or F10, eg, or use the `set` command to review current settings and change basic settings such as the graphics mode. Visit **http://bit.ly/Grub2Troubleshooting** for a detailed guide to using *Grub*'s own troubleshooting tools, but remember that in most cases the simplest fix is to use the *Boot-Repair* tool. If the *Grub* menu appears, then the issue may lie with its configuration file if things immediately grind to a halt after you select a menu option, but if Linux does start loading before grinding to a halt, the problem will lie with your operating system, in which case skip to the Post-*Grub* troubleshooting section (*see page 116*).

Boot-Repair tool

If you're struggling to fix *Grub* issues by hand, or there's no sign of *Grub* on your system at all, then you'll need to enlist the services of your rescue media and the *Boot-Repair* tool, which works with all Debian-based distros, including Ubuntu.

The *Boot-Repair* tool itself will launch automatically when you boot from a *Boot-Repair* tool disc, but if you're unable to create it, but have access to a Linux installation disc, use that in a live environment instead, then grab the *Boot-Repair* tool using the following commands:

```
$ sudo add-apt-repository ppa:yannubuntu/boot-repair
$ sudo apt-get update
$ sudo apt-get install boot-repair
$ boot-repair
```

The *Boot-Repair* tool is focussed on those early boot problems caused by the hard drive's boot sector, MBR and *Grub*. It basically provides a convenient and user-friendly graphical front-end to the tools required to fix many problems. The tool offers a 'Recommended repair' option that promises to fix most frequent problems, or you can click 'Advanced options' to see what it can do and manually select specific fixes without getting your hands dirty in the Terminal. The step-by-step guide (*see Tweak Boot-Repair tool Settings, page 117*) reveals what repairs and tweaks are possible, but note the tool is context-sensitive, and some options may be greyed out or missing depending on your setup. The tool automatically generates a log of your system and what it attempts to do, which you can then share on the Ubuntu user forums if necessary. Before attempting any advanced tweaks on your own, it pays to try the recommended option first, then ask for help on the forums using the output logs generated – this will ensure you choose the right option and don't cause more damage.

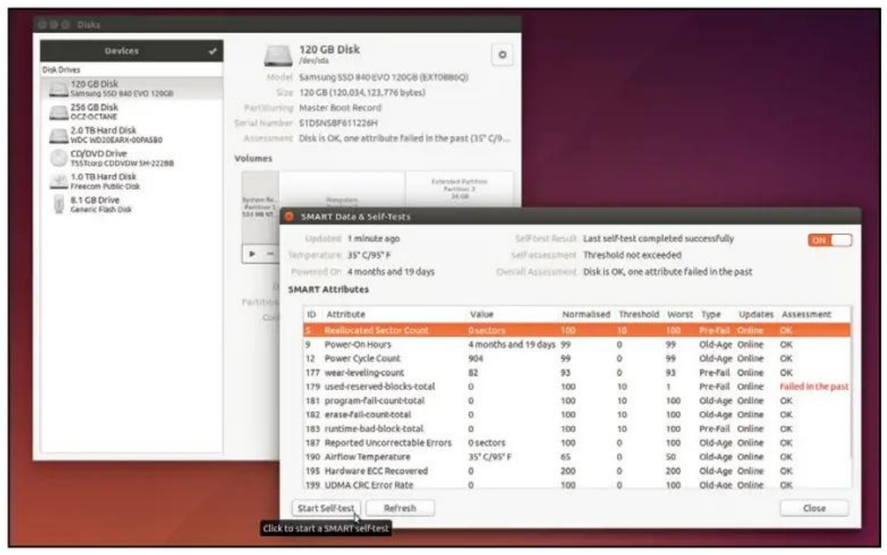

› Give your hard drive the once-over when booting from your rescue disc to check it's working as it should.

Non-Grub boot issues

Your rescue disc will also come in handy should you not even get as far as *Grub* loading. Once booted, verify the existence and state of your hard drive. Open the file manager and see if your partitions are visible and if you can access the files on them – this is a good time to back up any precious files before you proceed further. If nothing shows up, check whether the hard drive has been detected by opening the Disks utility from your Ubuntu Live CD – if you're using a *Boot-Repair* tool CD, you'll need to install the **gnome-disk-utility** through the Synaptic Package Manager under System Tools. Once installed, open it via the Accessibility menu. The *Disks* tool lists all physically attached drives – if yours isn't visible, you may find the drive has failed, in which case you'll be reaching for your latest backup after shelling out for a »

Quick tip

When you boot with your rescue media inserted your PC may automatically detect and boot from it; if it doesn't, look for an option to bring up a boot menu when your PC starts (typically a key like F11) to select it manually. Failing that, enter the EFI or BIOS configuration to set it as the first boot device.

» drive replacement, or starting again from scratch with a fresh Ubuntu installation and new-found love of backing up your system. Assuming your drive does show up, select it from the left-hand menu where you can examine the partition table plus check its physical health via its SMART [*Not all that smart – Ed*] attributes. Don't panic unless the drive is deemed on its last legs, but do focus your next check on the partitions themselves. If you run the *Boot-Repair* tool, its recommended settings will include a full disk check, but you can manually perform this check yourself using *GParted*, which is on both rescue discs.

GParted enables you to see how your partitions are arranged, as well as revealing which one is the boot. Right-click this and verify Mount is greyed out before choosing 'Check to schedule a disk check using the *fsck* tool'. This will check for and attempt to repair any problems it finds as soon as you click 'Apply', but it's important the partition isn't mounted before the check is run. Also give it as long as it needs to complete – this could take hours or even days in some extreme cases, and cancelling or aborting will almost certainly corrupt the partition. Make sure the check is run on all partitions on the boot drive.

In most cases, assuming the drive isn't physically damaged or corrupt beyond repair, running these tests should ensure you're able to at least get *Grub* working again.

Post-Grub troubleshooting

If you find that *Grub* appears to be working fine, but your problems begin when you attempt to load Linux itself. Try switching to verbose mode during boot by pressing the Esc key to see if any clues appear in the messages that scroll past (or if it hangs at a certain point). Make a note of these and do a search online for them for more advice. If this doesn't happen, hold Shift at boot to bring up the *Grub* menu if necessary, then select 'Advanced options' followed by '(recovery mode)', which will launch Ubuntu in a minimal state, plus mount the file system in read-only mode. If this is successful, after a succession of scrolling messages you should find yourself presented with the Recovery Menu, offering nine options.

The options are all pretty much self-explanatory – the clean option may be of use if your hard drive is full, which can cause boot problems. If your problems started because a package failed to install properly, then *dpkg* will repair it and hopefully get things working again.

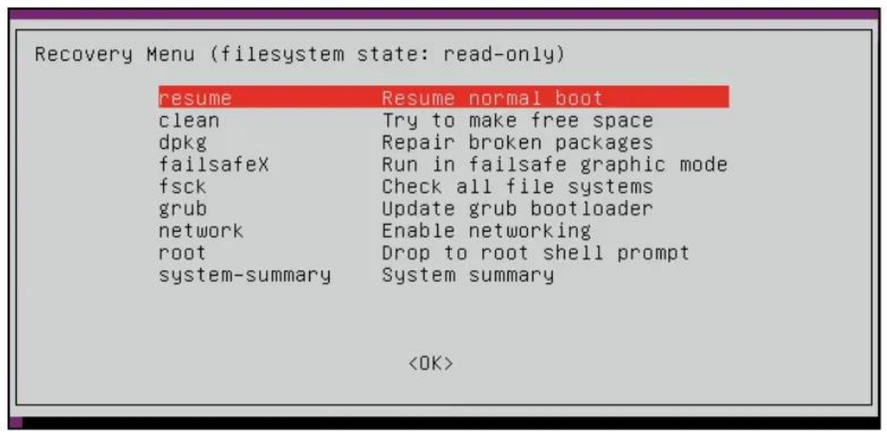

› Ubuntu's recovery mode lets you try various fixes when *Grub* works, but Ubuntu no longer wants to.

The failsafeX option is a useful if you find yourself booting to a black screen or the graphical desktop doesn't appear to be working correctly – it basically bypasses problems with your graphics drivers or X server to give you a failsafe graphics mode to troubleshoot your problem from.

We've touched on *fsck* already – this will check the drive for corrupt files, which can clear many errors, particularly if your PC crashed and has failed to boot since. The grub option isn't relevant unless you've used *Grub*'s own recovery tools in place of the Boot-Repair tool to get this far in the boot process – selecting this will make your changes permanent.

Use the network option to re-enable networking, and the root option to drop to the shell prompt, allowing you to troubleshoot directly from there. If doing so, be sure to mount the file system in read/write mode using the following command: `mount -o remount,rw /`

You can also pass temporary kernel parameters to Ubuntu during the boot process, which may help in some scenarios. With your chosen operating system selected in *Grub*, press the e key to edit the kernel file. Scroll down to the line beginning `linux` – parameters are added to the end of this line after `quiet splash`. You'll need to make sure that you leave a space between each parameter. Once done, press Ctrl+x to boot with those parameters.

Note that any parameters you add here are temporary – in other words, they're removed the next time you boot, so you can experiment until you find a solution that works, then – if necessary – make it permanent by editing the *Grub* configuration file (**sudo nano /etc/default/grub**).

You can also pass parameters from the live CD environment using the *Boot-Repair* tool using the 'Add a kernel option', which includes 15 common parameters that

Quick tip

System logs are a valuable source of troubleshooting information – and you can access these from the **/var/log** directory using your rescue disc's file manager or *nano* in a shell. Look out in particular for syslog, and investigate the dmesg shell command too.

Take a fail-safe backup

It may seem strange, but if you're struggling with start-up issues, you should attempt to take a backup of your hard drive before you perform any repairs – this means if you mess things up completely you can always roll back your system to the state it was in when the start-up problem first manifested itself.

Of course, if you're diligent and you back up your system regularly, you could always simply roll things back now to a working state, although bear in mind there may be data loss involved if your **home** folder is on the same partition as your Linux installation (as is the case with default Ubuntu installs).

You'll need a suitable backup device – typically a USB-connected hard drive – and a tool that takes a complete drive image of your system. The *dd* command-line tool can be used from both Ubuntu and Boot Repair Tool live CD environments, but the backup drive needs to be at least the same size – and preferably – bigger than the drive you're copying.

At the other end of the complexity scale is *Redo Backup & Recovery*. You'll need a blank CD or DVD to burn its 261MB ISO file to, but it provides an easy to follow graphical UI. You'll find it at **www.redobackup.org** – check out the feature on page 14 for a guide to using it.

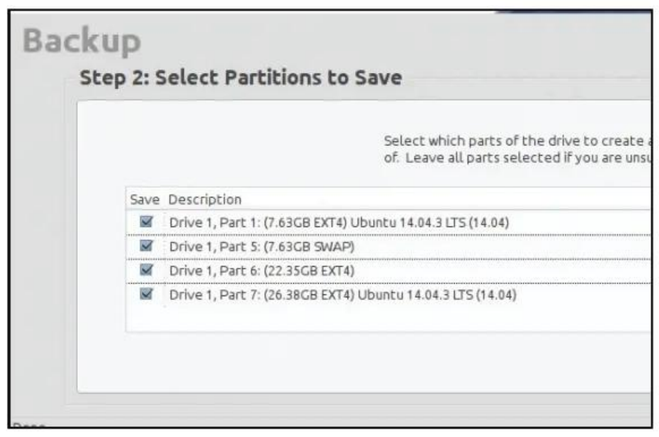

***Redo Backup & Recovery* offers by far the simplest way to make a fail-safe backup of your whole hard drive.**

can help with troubleshooting. Examples of these include acpi=off, which disables the ACPI system that's known to cause random reboots or system freezes on certain PCs, and nomodeset, which instructs Ubuntu to only load graphics drivers after the X environment has been loaded, and not before. These temporary parameters can be passed to your rescue disc too, in case you're having problems getting that working. Press [F6] at the initial boot screen to choose from the options on show. For more information on specific parameters, do an online search for the parameter or visit **http://bit.ly/KernelParametersList** for a complete list.

Repair install

There's one last thing you can try from the *Grub* boot menu – if your kernel has been upgraded, it's possible to boot using an older version of the kernel from the Advanced options screen under *Grub*. You'll see each version of the kernel listed – try the previous version if you believe your boot problem is linked to the latest kernel. If this works, you can make the version you've used permanent by editing the *Grub* configuration file – the simplest way to do this is by using the *Boot-Repair* tool.

If things look particularly bleak, then you may have luck reinstalling Ubuntu over the top of itself. Boot from the Ubuntu Live CD and choose the option to 'Install Ubuntu' when prompted. When you get to the 'Installation type' screen you'll be presented with a new option, pre-selected by default: 'Reinstall Ubuntu…'

This option basically reinstalls Ubuntu without touching your **home** folder or partition, which means not only should your documents and other files be preserved, but key settings and many programs may be left alone too. It'll also leave entries in your boot menu alone, ensuring you won't lose access to other operating systems.

What will be replaced are system-wide files, which will hopefully root out any corrupt ones and get your PC up and running again. Although it doesn't affect your files, it's still good practice to back up the drive – or at least your **home** folder or partition – before you begin.

To ensure you don't lose anything from your system, make sure you recreate all user accounts with the same login and password, including – of course – your own during the install process. ■

Tweak Boot-Repair tool settings

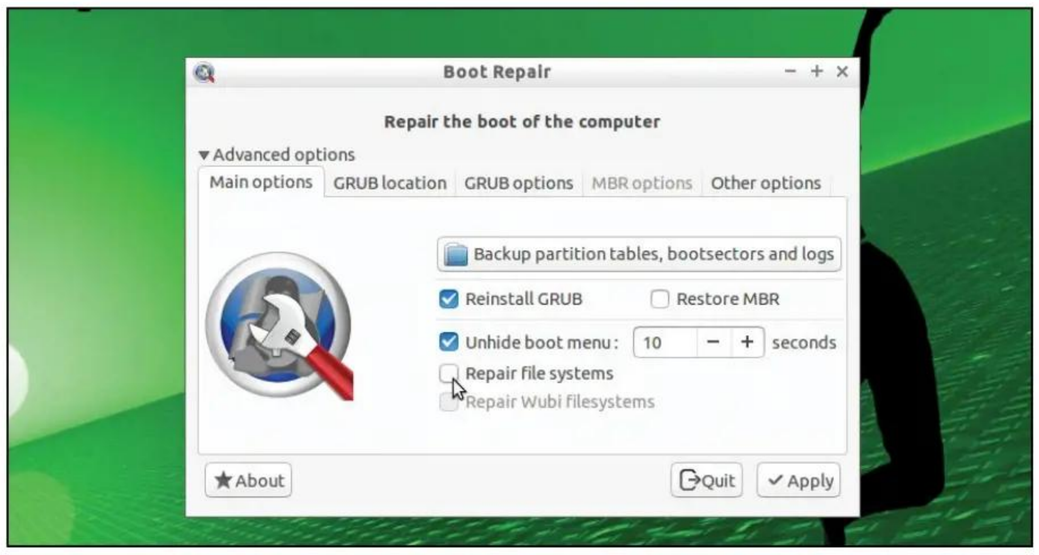

1 Main options

The first tab offers a convenient button for backing up your current partition table, boot sector and log – click this to copy this key information. It's also where you can reinstall *Grub*, restore the MBR and choose whether to hide the *Grub* menu. If you think your filesystem is corrupt, tick 'Repair file systems' to have it checked and fixed.

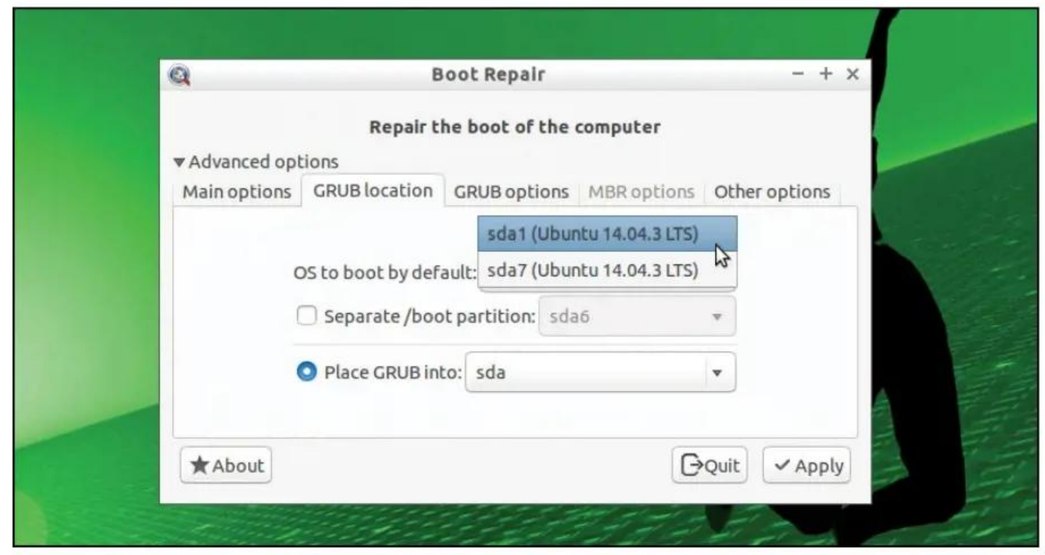

2 Grub location

This tab allows you to specify which OS to boot by default in a multi-boot setup. You can also choose to place *Grub* in its own separate **/boot** partition if you wish – typically this is only needed on encrypted disks, drives with LVM set up or some older PCs. The final option specifies which drive *Grub* itself will be placed (**sda** by default).

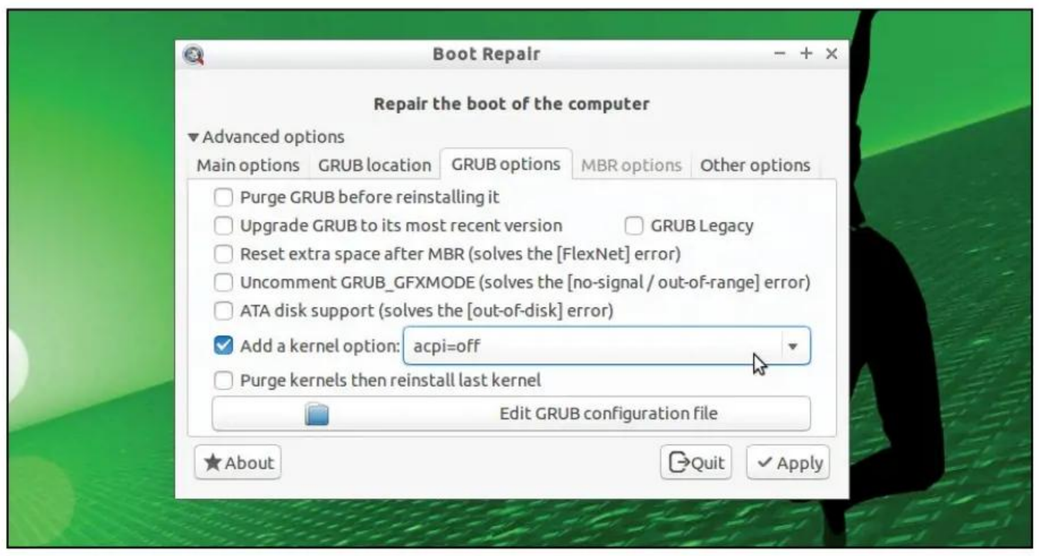

3 Grub options

This section opens with options for making sure *Grub* is updated to its latest version. There's also three specific error fixes offered. You can also add new kernel options to the *Grub* menu here, or purge all previous kernels before reinstalling the last one. You may even see an option allowing you to edit the *Grub* configuration file directly.

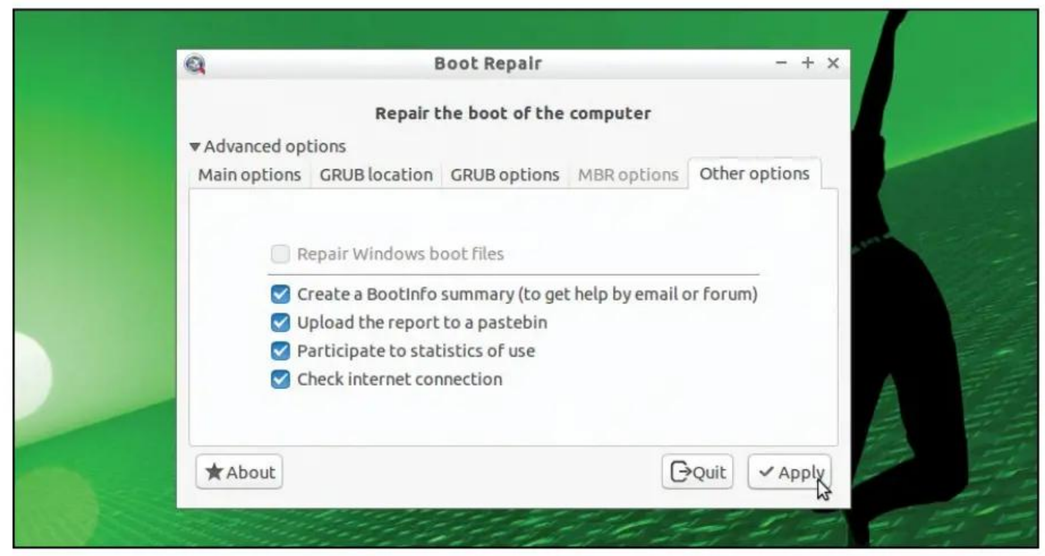

4 Other tweaks

If the MBR options tab isn't greyed out, use it to restore your MBR from a backup and choose which partition gets booted from it. The final Other options tab offers an opportunity for repairing Windows files (irrelevant in most cases) and provides options for pasting a summary of your settings online for reference.

Ubuntu: Google authenticator

We demonstrate how to put two authentications into your system account.

Locking access to your computer (and the data it holds) with just a password doesn't cut the mustard any more in this age of multi-core processors and easy to use password cracking tools. If you're really concerned about unauthorised access to your computer, you should definitely add an additional layer of authentication.

One of the easiest mechanisms to implement such a two-step verification is the Google Authenticator service, which issues a time-based authentication token to supplement the existing password challenge. Once you've integrated the service with your Ubuntu login, in addition to your user password, you'll be prompted for one of the quickly expiring tokens before being allowed to log in. Google Authenticator will generate these OTPs on your Android device once it's been configured for every user on your Ubuntu machine.

To implement multi-factor authentication, you'll need the Google Authenticator PAM (Pluggable Authentication Module). A PAM is a mechanism used to plug different forms of authentication into a Linux computer.

The Google Authenticator PAM module is available in the official Ubuntu's software repositories. To install the package on Ubuntu, head to the terminal and type:

```
$ sudo apt-get install libpam-google-authenticator
```

Once the package has been installed make sure you're logged in as the user you want to protect with the two-factor authentication. Now in the terminal window, type:

```
$ google-authenticator
```

This will initiate the process of creating a secret key for the user by asking you a bunch of questions. While it's safe to answer yes to all of them, it's a good idea to understand each one of them before making your final choice as these choices help balance security with ease-of-use. The first question is a pretty safe one and you should allow the command to update your Google Authenticator file by answering yes.

You'll then be asked if you'd like to restrict the use of a token which forces you to wait for 30 seconds between login. While it might seem inconvenient at first, you should agree to this limitation for maximum protection. The next question asks for permission to increase the time window that tokens can be use for from the default 1:30 minutes to 4:00 minutes. Although you can answer yes to this question to avoid any issues, type no for maximum security. If you notice any issues later on, rerun the command and increase the expiration time as suggested. The fourth and the last question asks you to limit the number of

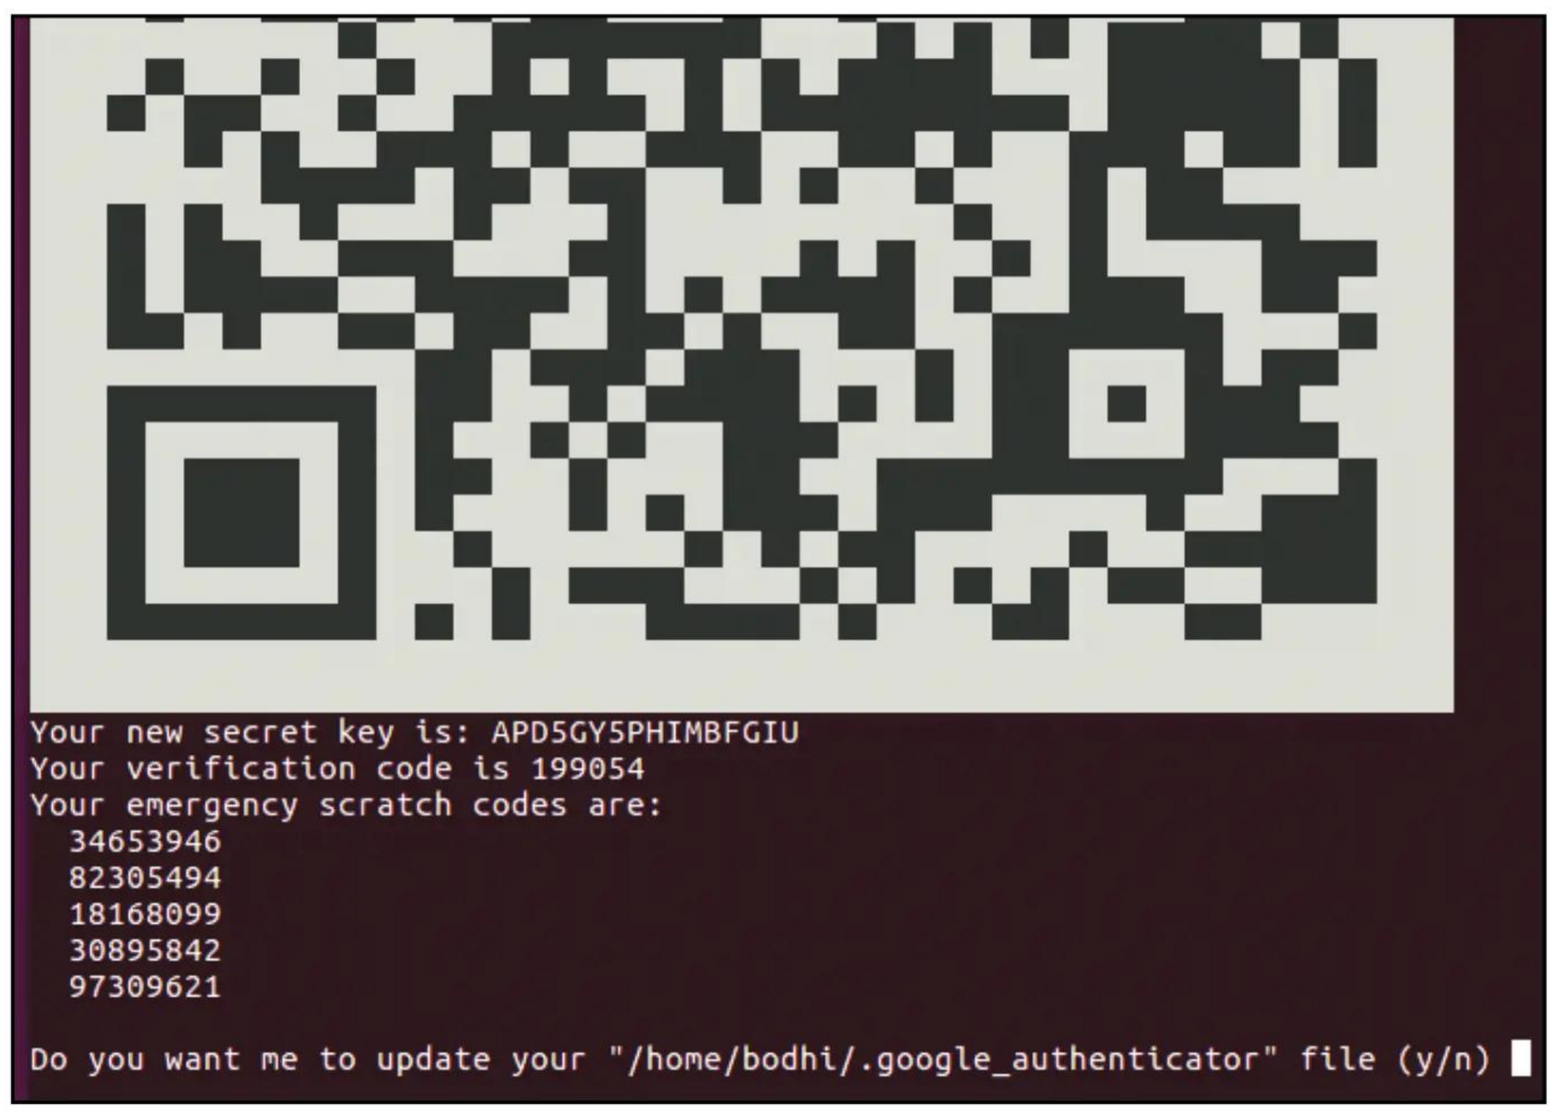

› After the initial setup, the two factor authentication will work even if your system and Android device aren't online.

Quick tip

To disable two-factor authentication, simply edit the **/etc/pam.d/common-auth** file and comment the line that we've added in the tutorial.

attempts for entering the authentication code. You should definitely enable this option as it helps prevent brute-force login attacks.

Book of Codex

When it's done, Google Authenticator will present you with a secret key and several emergency scratch codes. You'll need to make sure you note down these emergency scratch codes somewhere safe. They'll help you log in if you misplace the Android phone which generates the OTP. Each code can only be used once. The `google-authenticator` command will also generate a QR code that you can scan with your Android phone. Since we haven't installed the app yet, for the time being just note down the 16-digit code.

Now repeat this process for each user account that uses your computer. Ask everyone you share the computer with to log into their account, run google-authenticator, and make a note of their respective emergency scratch codes along with the 16-digit code. After you've generated the authentication code for all users, it's time to configure the login process to work with Google Authenticator. All you need to do is edit one file to add two-step authentication for all login attempts. Again, fire up the terminal and type:

```
$ sudo nano /etc/pam.d/common-auth
```

Scroll to the end of the file and add the following line:

```
auth required pam_google_authenticator.so nullok
```

Then save the file and exit. Here we've asked Ubuntu to use the Google Authenticator PAM module for all login attempts. The `nullok` bit at the end asks Ubuntu to let a user log in even if they haven't run the `google-authenticator` command to set up two-factor authentication. So let's assume you have two users, amber and megan, and have setup Google Authentication only for amber. Thanks to nullok, while amber will have to enter the OTP, megan will be able to log in with just her password. Note however that while this is a useful flexibility to have while you're testing Google Authenticator, once everything works smoothly and you have no issues logging in with the two factor authentication, it's advisable to force all users to log in through Google Authenticator only by removing the nullok bit for this command.

Go go gadget

Your Ubuntu installation is now all set up for two-factor authentication. To receive the OTPs, you will need to have installed the Google Authenticator app on your Android mobile phone from the Google Play Store. After installing the app, you will have to add an account for all of the users that you've run the `google-authenticator` command for on your Ubuntu installation. To do this, open the app and from the main window, tap the menu button (the three vertical dots in the upper right hand corner). Here tap 'Set up account' and then select the Enter key provided option. Now enter the 16-digit secret key that you have noted earlier after you had run through the google-authenticator tool. Give the account a name (a good idea is to use the username of the account this is for) and tap the 'Add' button. You've now set up two-factor authentication on your computer. The Android app will generate a new six-digit code for you every 30 seconds. When you log into your account or enter a `sudo` command, Ubuntu will prompt you for your password and you'll then be asked to enter the authentication code. At this point, enter the digits currently on display in the Android app. Once you've logged in successfully, make sure you edit the **etc/pam.d/common-auth** file and remove the `nullok` option to force login through Google authenticator. Also remember to create an account in the Android app for all the users on your Ubuntu installation. Going through the additional security prompt might seem like a hassle at first, especially when you need to switch to `sudo` to quickly edit a configuration file. However, if you're using the computer in a public place, you'll quickly learn to appreciate the benefits of the two-factor authentication mechanism.

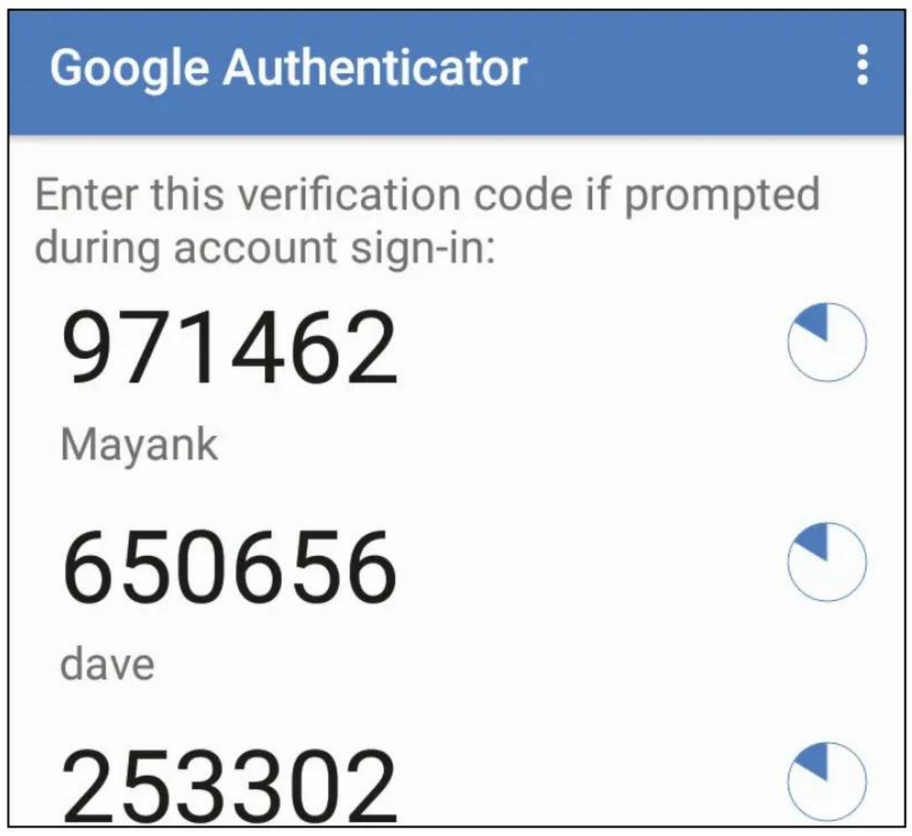

› Google Play Store has a number of other OTP algorithm compatible apps that you can use instead of the official Google app to generate your OTPs.

Quick tip

Even if you don't have multiple users, it's a good idea to create one to have a fallback in case there's a problem with the set up of the two-factor authentication mechanism.

Secure SSH with Google Authenticator

Instead of local logins, some people prefer to enable two-factor authentication only for remote SSH logins. For this, begin by making sure you've generated a code for the user you want to log in as the remote user. Next, make sure you use Google Authenticator for SSH logins by editing SSH's PAM configuration file with `sudo nano /etc/pam.d/sshd`. Scroll down to the bottom and add the following line to the file:

```
auth required pam_google_authenticator.so nullok
```

Just as in the main tutorial, the `nullok` word at the end tells PAM that this authentication method is optional. This allows users without a Google authenticator key to still log in using their SSH key. Remember that this is just a failsafe to prevent you from being locked out in case something goes wrong with the setup process. However, once you've tested it successfully, generate a key for all SSH users and delete `nullok` from the end of this line to make logins via OTP mandatory.

After editing the SSH's PAM file, it's time to configure SSH to support this kind of authentication. Open the SSH configuration file for editing with `sudo nano /etc/ssh/sshd_config`. Look for the line that reads `ChallengeResponseAuthentication` and set its value to yes. If the line doesn't exist, make sure you add it manually. Save and close the file, then restart SSH to reload the config files with `sudo service ssh restart`. You'll now be prompted for both your password and Google Authenticator code whenever you attempt to log in via SSH.

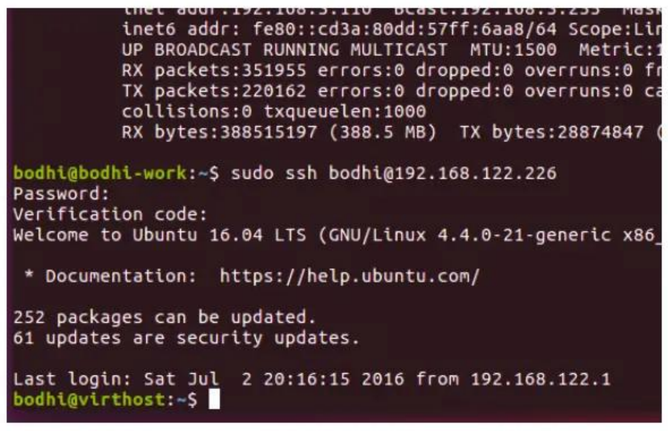

TruPax: Easy file encryption

Introducing the versatile Java utility that can encrypt your files in just three simple steps in a large container file.

For non-crypto nerds, *VeraCrypt* is a program that enables you to create encrypted containers of any size, inside which you can place your personal photos, bank details or any other data for which the world isn't ready. *VeraCrypt* mounts these containers with a password, which means only you can access your files. For the uber-paranoid there are even options to combine multiple encryption ciphers and use key files in addition to a password.

VeraCrypt does require you to have some idea in advance of the size of all the files you need. Once an encrypted container has been created, it's not possible to resize it. What's more, if you create a container in the full belief that there's enough room for all your files and then find you don't have enough space, there's nothing to do but start the volume creation process all over again.

This is a nuisance at the best of times but is particularly frustrating if you choose to store your encrypted container with a cloud service such as Dropbox, as the upload process has to be restarted.

Trupax posits a solution to creating redundant, huge blocks of data on your machine. This Java app can be used to select files or folders before creating a volume and will create one of exactly the right size. There is even an option to add some more free space if you think you'll need it.

In this tutorial, we will explore how to use *TruPax* to create an encrypted container of exactly the right size, then open it in *VeraCrypt*. In order to proceed, you'll need Java installed on your machine (either the vanilla variety from Oracle or OpenJDK – keep reading for more info). Once the volume is created, you'll also need *VeraCrypt* pre-installed on your system to mount it.

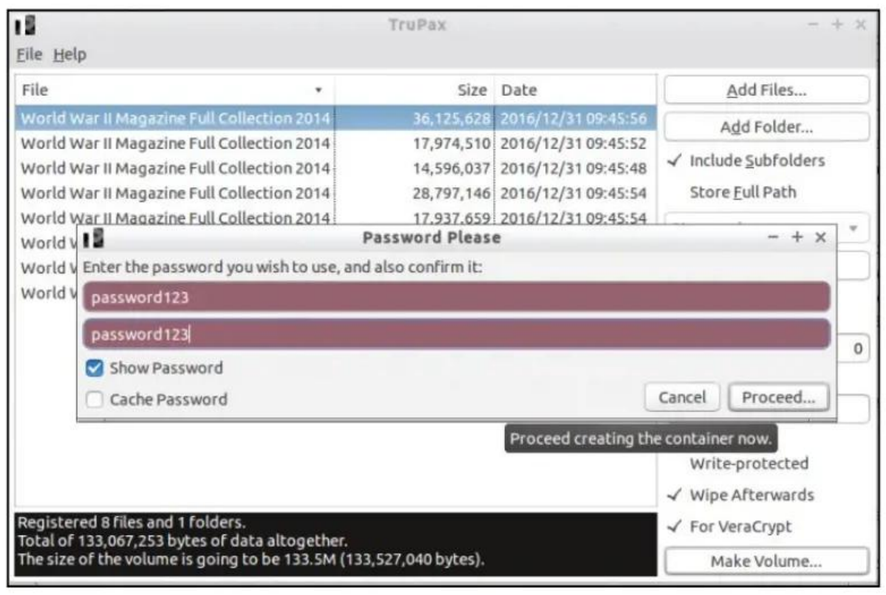

› If the world isn't ready to know you're a military history buff, create a secure container for your magazine collection. Check 'show password' to display as you type.

VeraCrypt can be downloaded from **http://veracrypt.codeplex.com**. The website itself has some excellent documentation on getting started or you can see our previous tutorial in **Linux Format issue 218**.

This tutorial was written for Ubuntu Linux but both Java and *VeraCrypt* are compatible with all Linux versions, so you should have no trouble running *TruPax* regardless.

Another great advantage of *TruPax* is that admin privileges are not required simply to create a volume or

Quick tip

Before clicking 'Add Folder', be sure to check the box marked 'Include Subfolders' to copy any directories inside it into a new container too.

TruPax Plus

As convenient as *TruPax* is for creating volumes of just the right size, this comes at the price of extra security.

If you wish you can also add a key file. This means in order to mount your *TruPax* volume, *VeraCrypt* would use both your password and a file. This is a form of Two Factor Authentication (something you have and something you know) and hugely reduces the chance that an attacker can access your volume.

VeraCrypt also has a random keyfile generator. Go to Tools > Keyfiles to launch this.

You can use one or a number of files. Remember that the first 64 kilobytes of the file must not change, or the container can't be opened. This is why images and music files work better. *VeraCrypt* doesn't alter the file, so if for instance you choose an MP3 of Aqua's Doctor Jones as your key file, there's no way from examining the file itself to know that it's been used this way.

In *VeraCrypt* choose Volumes > Change Volume password to get started. Enter your password in the field in the Current section, then again in the New section. Check 'Use Keyfiles' then click the 'Keyfiles' button to select the actual files. The order in which you select them doesn't matter at all.

Before clicking 'OK' at the top right, move to the dropdown menu marked 'PKCS-5 PRF' in the New section, and choose 'HMAC-Whirlpool' to use a hash that has been developed entirely independently of our friends in the NSA.

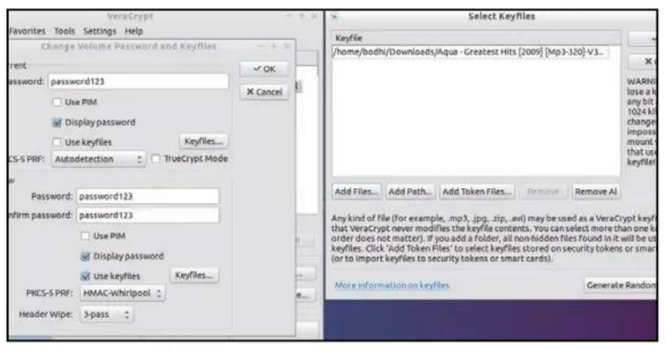

extract its contents elsewhere on the system. *VeraCrypt* will, however, require your admin password to mount the container so that you can edit your files.

One of the advantages *TruPax* has over *VeraCrypt* is that it's extremely simple to use. Once the program has launched, simply click 'Add Files' or 'Add Folder' to load your data into the main window.

Once the files are added, you'll see a notification at the bottom of the window telling you how large your container will be. If you want space for more files, use the 'Free Space' box on the right-hand side of the window. This is quite intuitive and will recognise values such as '500m' or '2g'.

Optionally you can give your volume a label. Before clicking the 'Make Volume' button at the bottom right, make sure to check 'Wipe Afterwards' if you want to securely delete the original files (see below).

You'll be asked to set a password for the volume, then *TruPax* will begin to generate your container.

Wipe and extract

Unless you specifically ask it not to, by unchecking the box 'Wipe Afterwards', *TruPax* will securely erase all of the original files after copying them into a secure container. Make sure before you begin that all of your files are backed up to a safe and secure place.

You can use File > Extract in the menu at the top left to make an unencrypted copy of your files, provided you know the password. As such, if you only want to encrypt files for long term storage, such as for backups, you may find that you don't need to use *VeraCrypt* at all.

If, which is more likely, you want to be able to edit these files from time to time and add more to your container, you can still use *TruPax* to create the container initially and then actually mount it using *VeraCrypt*.

TruPax is written in Java and therefore needs a JRE (Java Runtime Environment) installed on your machine. The tutorial lists how to add the repository for this on your machine as well as how to install Java 8, which is the minimum required for *TruPax* to run.

Security conscious people and/or those committed to truly free and open source software may prefer to install OpenJDK which is also developed by the good people of Oracle but contains no closed-source code. If you prefer to do this, ignore Step One of the tutorial and run the commands:

```
sudo add-apt-repository ppa:openjdk-r/ppa
sudo apt-get update
sudo apt-get install openjdk-8-jre
```

If you are one of those people who goes straight to the tutorial before reading the main article (you know who you are), rest assured both Oracle's Java RunTime Environment and OpenJDK can live on the same machine happily together. Run the command `sudo update-alternatives --config java` if you want to change the default Java program to OpenJDK. Select the correct number, press 'Enter', then restart your machine to apply your changes.

Tru lies

TruPax in itself can only create encrypted containers for files and folders, then extract them all at once elsewhere. This isn't very handy for day to day use, which is why we recommend using it in tandem with *VeraCrypt*.

Unlike *VeraCrypt*, *TruPax* cannot create hidden volumes, whereby you can have two passwords for the container – one which you can give to an adversary and leads to harmless data and another which leads to your real files. You can of course use *VeraCrypt* to create a hidden volume and copy your *TruPax* container there if you wish.

If you follow the steps in 'TruPlax Plus' to increase the security of your container through adding key files, you won't be able to extract the contents using *TruPax*.

For more details about configuring *TruPax*, make sure to read the **trupax_EN.html** that you downloaded along with the program, or take yourself along to the developer's website at **www.coderslagoon.com**.

Quick tip

By default, *TruPax* creates containers that are compatible with *VeraCrypt*'s predecessor *TrueCrypt*, which comes with a number of security flaws. So for safety, make sure to check 'For Veracrypt' when creating a volume.

Encrypting with TruPax

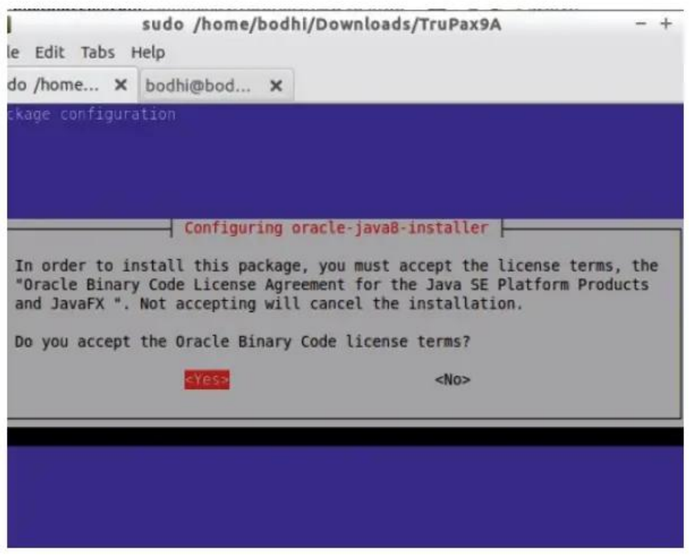

1 Download JRE & TruPax

Open Terminal and add the PPA for Java with `sudo add-apt-repository ppa:webupd8team/java`. Next run `sudo apt-get update` then `sudo apt-get install default-jre oracle-java8-installer`. Head to **http://coders lagoon.com** to download *TruPax* for your machine. Right click the file, choose 'Extract', then use `cd` to navigate to its folder, eg, `cd /home/bodhi/Downloads/TruPax9A`.

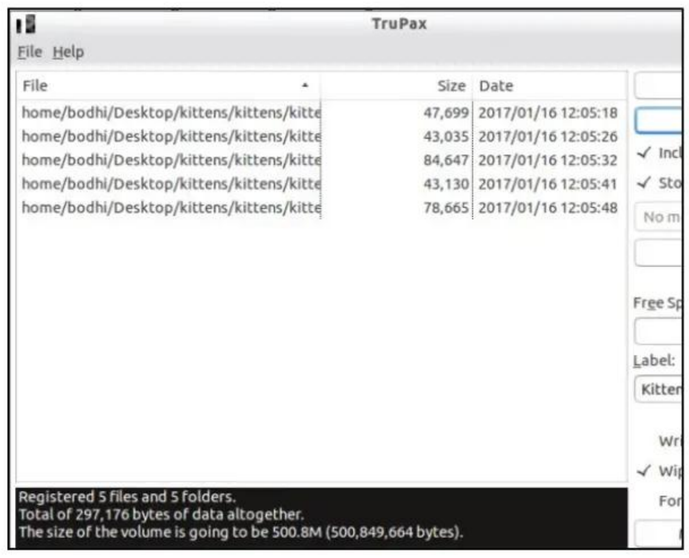

2 Launch TruPax

Run `java -jar trupax.jar` to launch *TruPax*. Choose your language (German or English are the choices). Click 'Add Files' or 'Add Folder' to select data for encryption. Click 'Make Volume' when you are ready to set a password. *TruPax* will ask you where to save the file and to set a name, such as **kittens.hc**. You can use any name that takes your fancy as an extension.

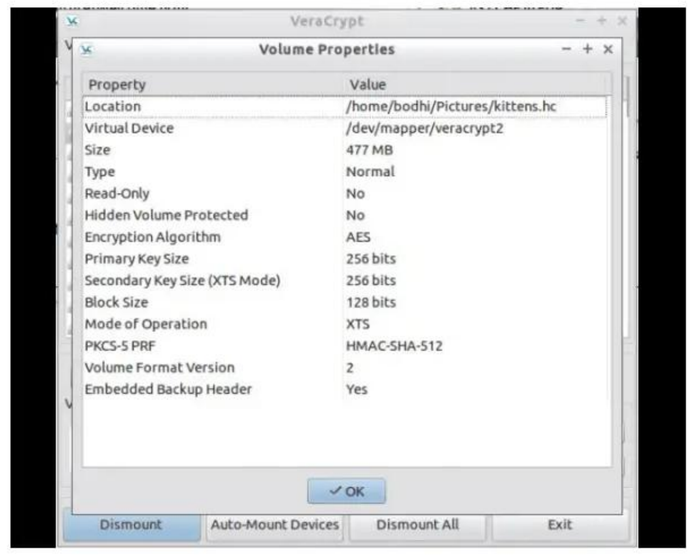

3 Mount in VeraCrypt

Launch *VeraCrypt*, then click 'Select File'. Navigate to your *TruPax* volume, then click 'Mount'. Enter the password you chose earlier. Find more information about your new volume by choosing Volumes > Volume Properties. This will show where the container is mounted, eg, **/dev/mapper/veracrypt2**. At this stage, you may wish to increase the security (see 'TruPax Plus').

Admin: Core commands

20 terminal commands that all Linux web server admins should know.

Are you an 'accidental admin'? Someone who realised, too late, that they were responsible for the workings of a Linux server and – because something has gone wrong – finds themselves lost in a world of terminals and command lines that make little sense to normal humans?

What is SSH, you may be asking yourself. Do those letters after 'tar' actually mean anything real? How do I apply security patches to my server? Don't worry, you're not alone. And to help you out, we've put together this quick guide with essential Linux commands that every accidental admin should know.

Becoming an accidental admin

While we'd argue that they should, not everyone who starts using Linux as an operating system does so through choice. We suspect that most people's first interaction with Linux happens somewhat unwittingly. You click a button on your ISP's account page to set up a personal or business web server – for a website, email address or online application – and suddenly you're a Linux admin. Even though you don't know it yet.

When you're starting out with your web server, things are usually straightforward. Nearly all hosting providers will give you a web interface such as Cpanel or Plesk to manage your server. These are powerful pieces of software that give you quick an easy access to logs, mail services and one-click installations of popular applications such as Wordpress or forums. But the first time you have to do something that isn't straightforward to do through the graphical control panel, you're suddenly out of the world of icons and explanatory tooltips and into the world of the text-only Terminal.

To make things worse, for a lot of people the first time they have to deal with the Terminal is when something has gone wrong and can't be fixed through the control panel. Or perhaps you've just read that there's a major security flaw sweeping the web and all Linux servers *must* be updated at once (it happens – search for 'Heartbleed' to find out more). Suddenly you realise that your nice control panel hasn't actually been updating your server's operating system with security patches and your small personal blog may well be part of a massive international botnet used to launch DDOS attacks against others. Not only are you a stranger in a strange land, you're probably trying to recover or fix something that was really important to you, but which you never gave much thought to while it was being hosted for a couple of pounds a month and seemed hassle-free.

You are an 'accidental admin'. Someone who is responsible for keeping a Linux webserver running and secure—but you didn't even realise it. You thought all that was included in your couple of pounds a month you pay to your ISP – and only found out it's not when it was too late.

Since most webservers are running Ubuntu, this guide is based on that particular distribution. And all the commands here are just as applicable to a Linux desktop as they are to a web server, of course.

```
670  df
671  man df
672  df -h
673  top
674  w
675  nano
676  dir
677  history
678  diff -r /media/studiopc/theRaptorRaid/Pics /media/sdb1/Pics
679  sudo su
680  ssh root@crazy8ball.htxt.co.zax
681  ssh root@crazy8ball.htxt.co.za
682  ssh root@htxt.co.za
683  ssh root@41.79.76.130
684  ifconfig
685  ls | less
686  uit
687  ls | less
688  apt-get update
689  sudo apt-get update
690  history
studiopc@studiopc:~$
```

› Can't remember that really clever thing you did last week? History is your friend.

1 sudo

The most fundamental thing to know about Linux's approach to administration is that there are two types of accounts that can be logged in: a regular user or an administrator (aka 'superuser'). Regular users aren't allowed to make changes to files or directories that they don't own—and in particular this applies to the core operating system files which are owned by an admin called 'root'.

Root or admin privileges can be temporarily granted to a regular user by typing `sudo` in front of any Linux command. So to edit the configuration file that controls which disks are mounted using the text editor, *nano*, you might type `sudo`

Connecting to the server

As an accidental admin, your first challenge is going to be connecting to your server in the first place. In your web control panel, you might see an option to open a Terminal or console in your web browser, but this tends to be quite a laggy way of doing things.

It's better to open up a Terminal window on your own machine (if you're running Ubuntu just press Alt+Ctrl+t, if you're on Windows you'll need an application like *PuTTY*). Now, at your command prompt, type `ssh username@yourserver.com` (or you can replace `yourserver.com` with an IP address).

The `ssh` command will open a secure shell on the target machine with the specified username. You should get a password prompt before the connection is allowed and you will end up in a text interface that starts in the **home** folder of the username.

If you're going to be connecting regularly, there's an even more secure way of using `ssh` and that's to bypass the password prompt all together and use encrypted keys for access instead. To follow this approach, you'll need to create a public/private SSH keypair on your machine (for example, Ubuntu users can type something like `ssh-keygen -t rsa -b 4096 -C "your_email@example.com"`) and copy the public part of the key into the **.ssh** folder on the target server.

You will find some full instructions for doing this here: **https://help.github.com/articles/generating-an-ssh-key**.

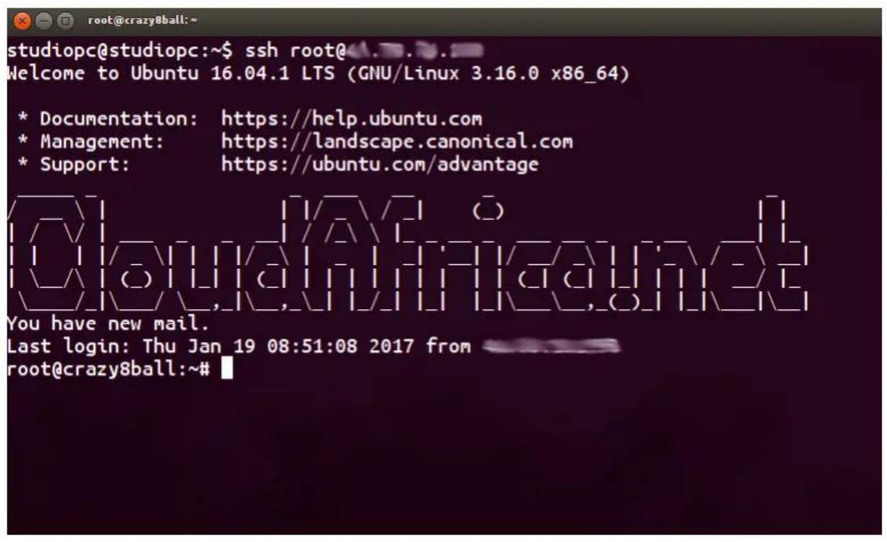

› Even if someone copies your key, they'll still need a password to unlock it.

`nano /etc/fstab` (we really don't recommend this unless you know what you're doing). After entering `sudo`, you'll be asked for your user password. On a desktop PC, this is the same one that you use to log in. If you're logging into your own webserver, however, there's a good chance that you'll already be the root user and won't need a password to make important changes.

If you can't execute `sudo` commands, your web host has restricted your level of access and it probably can't be changed. User accounts can be part of 'groups' in Linux and only members of the sudoers groups can use the `sudo` command to temporarily grant themselves admin privileges.

2 su

While `sudo` gives you great power, it still has limitations. Most of all, if you've got a whole bunch of commands to enter, you don't want to have to type it out at the start of every single line [*at least the password has a 5 minute timeout–Ed*]. This is where `su` comes in, which will give you superuser powers until you close the terminal window. Type `sudo su` followed by your password, and you'll see the prompt change from `yourname@yourserver` to `root@yourserver`. You might think `su` stands for superuser, but It's actually a command to change to any user on the system and if it's used without an account name after it, `su` assumes you want to be root. However, using `su myname` will switch you back to your original, non-super, login.

3 ifconfig

Since you're troubleshooting a web server, it's probably a good idea to get as many details about its actual connection as possible noted down. The `ifconfig` command can be run without `sudo` privileges and tells you details about every live network connection, physical or virtual. Often this is just for checking your IP address, which it reports under the name of the adaptor, but it's also useful to see if you're connected to a VPN or not. If a connection is described as eth0, for example, it's an Ethernet cable meanwhile tun0 is a VPN tunnel.

4 chown

There's tons more you can learn about `chmod` and we strongly recommend that you do, but it has a sister command that's even more powerful. While `chmod` dictates what users who aren't the owner of a file can do, the `chown` command changes the file owner and group that it belongs to completely. Again, you'll probably need to put `sudo` in front of anything you `chown`, but the syntax is again simple. An example might be `chown myname:mygroup filename.file`.

5 service restart

No, we're not telling you to 'try turning it off and on again', but sometimes it's a good place to start (and sometimes its essential to load changes into memory). It's possible you might be used to start and stop background processes on a Windows desktop through the graphical System Monitor or Task Manager in Windows. However, in the command line Terminal to a server it's a little more tricky, but not by much.

Confusingly, because many Linux distributions have changed the way they manage startup services (by switching to *systemd*) there are two ways of doing this. The old way, which still works a lot of the time, is to just type `service myservice restart`, preceded with `sudo`, when it's necessary. The new, correct, way is a little bit more verbose: `systemctl`

»

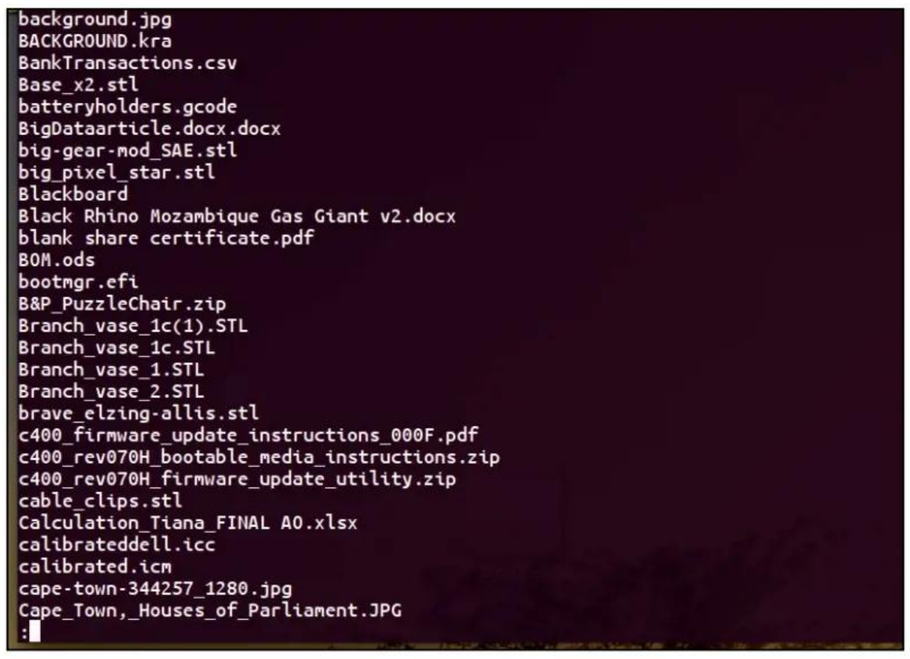

› Unless you can read 1,000 lines a second, you'll need to use `ls | less` to explore folders.

Quick tip

If you're changing names, permissions or ownership most commands have a -R or -r option, which stands for 'recursive'. Essentially, this changes the attributes of all files inside a folder, rather than just the folder itself.

» `restart myservice.service`. So if you want to restart Apache, for example, the core software that turns a mere computer into a web server, it would be `sudo systemctl restart apache2.service`.

6 ls

The key to understanding the console is all in the path (*see Path To box, below*), which tells you whereabouts you are in the folder structure at any given time. But how do you know what else is in your current location? Easy: you use `ls`. The `ls` command lists all the files within the folder that you're currently browsing. If there's a lot of files to list, use `ls | less` to pause at the end of each page of filenames.

7 cat

A command you'll often see if you're following instructions you've found online – and aren't always sure what you're doing – `cat` is short for concatenate and is used to combine files together. In its simplest form it can be used to take file1.txt and file2.txt and turn them into file3.txt, but it can also be combined with other commands to create a new file based on searching for patterns or words in the original.

Quite often you'll see `cat` used simply to explore a single file – if you don't specify an output filename, `cat` just writes what it finds to the screen. So online walkthroughs often use `cat` as a way of searching for text within a file and displaying the results in the terminal. This is because `cat` is non-destructive—it's very hard to accidentally use `cat` to change the original file where other commands might do.

8 find

A useful and under used command, the `find` command is pretty self-explanatory. It can be used to find stuff. Typing it by itself is much like `ls`, except that it lists all of the files within sub-directories of your current location as well as those in your current directory. You can use it to search for filenames using the format `find -name "filename.txt"`. By inserting a path before the `-name` option, you can point it at specific starting folders to speed things up. By changing the `-name` option you can search by days since last accessed (`-atime`) or more.

9 df

Maybe your server problems are to do with disk space? Type `df` and you'll get a full breakdown of the size and usage of every volume currently mounted on your system. By default it'll give you big numbers in bytes, but if you run `df -h` (which stands for 'human readable' the volume sizes will be reported in megabytes, gigabytes or whatever is appropriate.

10 apt-get update && upgrade

Probably the single most important command to know and fear. We all know that to keep a computer system secure you need to keep it updated, but if you've got control of a Linux box the chances are that it isn't doing that automatically.

A simple `sudo apt-get update` will order your system to check for the latest versions of any applications it's running, and `sudo apt-get upgrade` will download and install them. For the most part these are safe commands to use and should be run regularly—but occasionally updating one piece of software can break another, so back-up first...

11 grep

As computer commands go there are few more fantastically named for the newcomer than the `grep` [it's a real verb!–Ed] command. How on earth are you ever going to master this Linux stuff if it just makes words up? But `grep` is a great utility for looking for patterns within files. Want to find every line that talks about cheddar in a book about cheeses? `grep "cheddar" bookofcheese.txt` will do it for you. Even better you can use it to search within multiple files using wildcards. So `grep "cheddar" *.txt` will find every text file in which cheddar is reference. So now you grok `grep`, right?

12 top

When you're working in a graphical user interface such as a Linux desktop environment or Windows desktop, there's always an application like System Monitor or Task Manager which will call up a list of running applications and give you details about how many CPU cycles, memory or storage they're using. It's a vital troubleshooting tool if you have a program that's misbehaving and you don't know what it is.

In a similar way, you can bring up a table of running applications in the Linux Terminal that does the same thing by typing `top`.

Like a lot of command line utilities, it's not immediately obvious how you can close `top` once you're finished with it without closing the terminal window itself—the almost universal command to get back to a prompt is Ctrl+c.

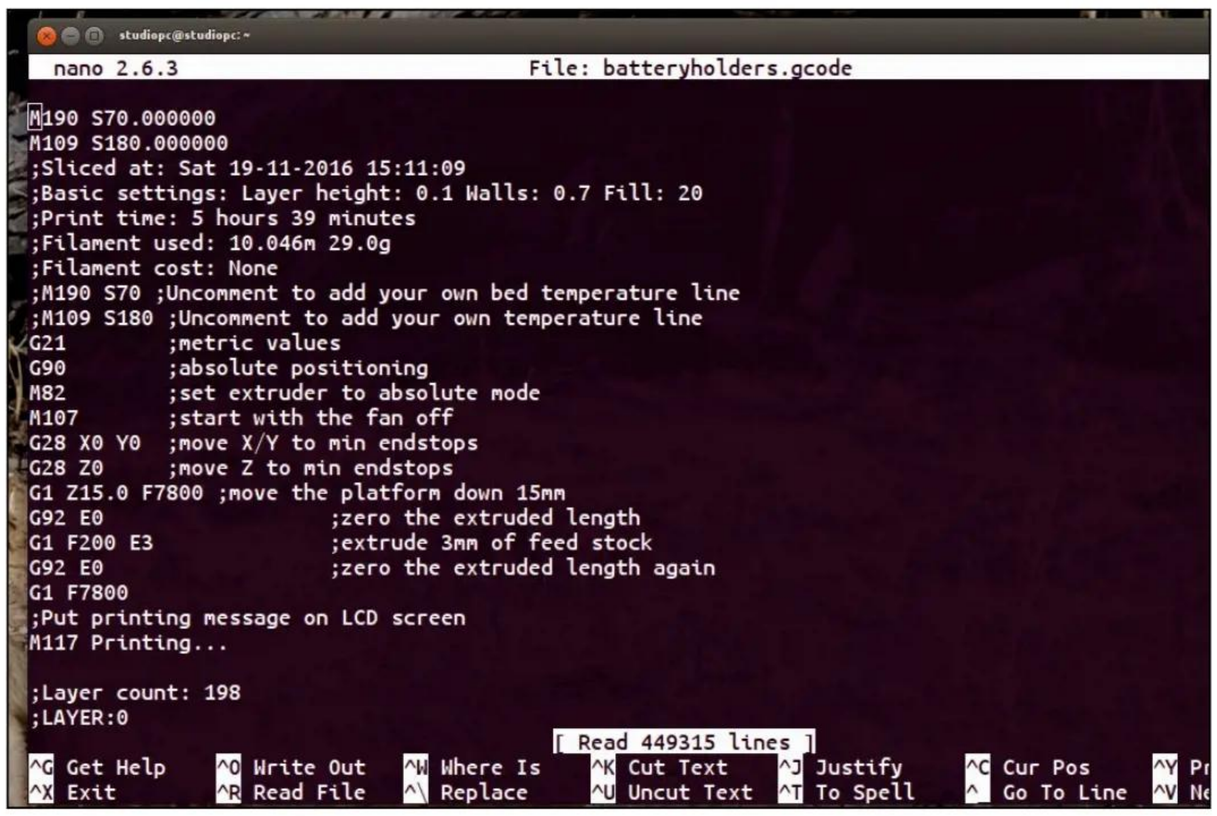

› Nano isn't the only terminal text editor, but it's the easiest to use.

Path to

When you open a Terminal window within Linux, it can be a bit disorientating. But the words that sit in front of the flashing cursor will tell you where you are.

The first word is the name of the user you're logged in on, and it's followed by an '@' sign. The second word is the hostname of the machine you're logged into. If you open up a Terminal on your desktop, usually the username and hostname are the same. So you'll see 'myname@myname'. When you log into a remote server, though, they'll be very different.

This information is followed by a colon which is followed by the path to the directory you're in, followed by a dollar sign. When you first open a Terminal, it will usually print yourname@yourname:~$. The tilde '~' indicates you're in the **home** folder for your username. If the dollar sign is replaced with a '#', you're using the machine as a root user. See `cd` for moving around and watch how the path changes as you do.

20. chmod

User permissions are one of the most important parts of Linux security to understand. Every file has a set of permissions which defines who can see a file; who can read and write to a file; and who can execute a file as a program.

A file which can be seen by web visitors, but can only be changed by a specific user, is just about as basic as it gets when it comes to locking down a server. The problem is that some files need to be changeable and some don't—think of a Wordpress installation for a blog. You want Wordpress to be able to write some files so it can update them, but there's also a lot of files you don't want it to be able to change—and you really don't want to give it power to execute code unless you have to. The flipside is that problems with web servers can be traced back to incorrect file permissions, when an app needs to be able to modify a file but has been locked out by default.

Your friend in this area is `chmod`. It changes permissions for which users and groups can read, write or execute files. It's usually followed by three digits to indicate what the owner, members of its group and everyone else can do. Each digit from 0-7, where 7 allows for read, write and execute and 1 is execute only. If your user 'owns' the file in question, the syntax is simple.

`chmod 777 filename`, for example, will give all users the ability to read and write to a file. It's good practice not to leave files in this state on a webserver—for obvious reasons. If you don't own the file, you'll need to add `sudo` to the front of that command.

13 kill, killall

Using `top` you can figure out which application is using all your CPU cycles, but how do you stop it without a right-click > End process menu? You use the command `kill` followed by the process name. If you want to be sure and kill every process with a name that contains that application name, you use `killall`. So `kill firefox` will close down a web browser on a Linux desktop.

14 w

From the weirdness of `grep` to the elegance of the `w` command, a whole command in a single letter. If you think another user is logged into your system, this is an important command to know. You can use `w` to list all currently active users, although don't rely on it too much as it's not hard for a hacker to be hidden.

15 passwd

You must use `passwd` with extreme care. Ultra extreme care. Because the next word you write after it will become your login password, so if you type it incorrectly or forget it, you're going to find yourself in serious trouble. You can only change your own user's password by default, but if you grant yourself sudo powers you can change any user's credentials by including their username after the password itself. Typing `sudo passwd`, meanwhile, will change the password for root.Check out the manual (`man passwd`) page for some useful options to expire passwords after a certain period of time and so on.

16 cd

If you have a graphical interface and file browser, it's pretty easy to move to new locations on your hard drive just by clicking on them. In the Terminal, we know where we are because of the path (*see Path To box, left*), and switch location using `cd` which stands for 'change directory'.

The `cd` command in mainly used in three ways:

1 **cd foldername** This will move you to that folder, provided it exists within the folder you're currently browsing (use `ls` if you're not sure).

2 **cd ~/path/to/folder** This will take you to a specific location within your **home** folder (the `~` character tells `cd` to start looking in your **home** folder). Starting with a `/` will tell `cd` to start the path at the **root** folder of your hard drive.

3 **cd ..** This final useful command simply takes you up one level in the folder structure.

17 mv & rm & cp

When you get the hang of it, using a terminal as a file manager becomes pretty simple and quite a joyful experience. As well as `cd`, the three fundamental commands are `mv`, `rm` and `cp`. The `mv` command is used to move a file from one location to another, `rm` is used to remove or delete a file and `cp` will copy files and folders.

Just as with `cd`, you can either enter a filename to operate on a file in the directory you're working in or a full path starting from the root of the drive with `~`. For `mv` the syntax is `mv ~/location1/file1.file ~/location2/location`.

The big thing to remember is that in the Terminal there's no undo or undelete function: if you `rm` a file and it's gone forever (or at least will require very specialist skills to retrieve) and in a similar fashion, if you `mv` or `cp` a file you'd better make a note of where it went.

```
studiopc@studiopc:/etc$ cd /etc/X11/Xsession.d
studiopc@studiopc:/etc/X11/Xsession.d$ ls
00upstart                       60x11-common_localhost       90atk-adaptor
20dbus_xdg-runtime              60x11-common_xdg_path        90consolekit
20x11-common_process-args       60xbrlapi                    90gpg-agent
30x11-common_xresources         60xdg-user-dirs-update       90qt5-opengl
35x11-common_xhost-local        65compiz_profile-on-session  90qt-a11y
40x11-common_xsessionrc         65snappy                     90x11-common_ssh-agent
50_check_unity_support          70gconfd_path-on-session     95dbus_update-activation-en
50x11-common_determine-startup  70im-config_launch           99upstart
55cinnamon-session_gnomerc      75dbus_dbus-launch           99x11-common_start
55gnome-session_gnomerc         81overlay-scrollbar
studiopc@studiopc:/etc/X11/Xsession.d$ cd ..
studiopc@studiopc:/etc/X11$ cd ..
studiopc@studiopc:/etc$ cd ..
studiopc@studiopc:/$ ls
bin   dev  include        lib    libGL.so    media  proc  sbin  swap  ubiquity-apt-cl
boot  etc  initrd.img     lib32  libx32      mnt    root  snap  sys   usr
cdrom home initrd.img.old lib64  lost+found  opt    run   srv   tmp   var
studiopc@studiopc:/$
```

› **Keep an eye on the directory path in front of the command line to figure out where you are.**

18 nano

It might seem odd if you've spent your life in graphical applications and utilities, but complex programs run in the text terminal, too. There are several text editors that normally come as part of the whole package – notably *nano* and *vi*. You can open a blank document by typing `nano`, or you can edit an existing one by typing `nano ~path/to/text.txt` (and you can do the same with *vi*). Some of the terminology may seem odd, though – to write out (Ctrl+O) means save, for example, and so forth.

19 history

And finally, if you've been copying and pasting commands from the internet all day, you might want to take a moment to check up on what you've actually done. You can use `history` to give you a list of all the terminal commands entered going back a long, long way. Execute specific numbered commands with `!<num>`, or you can go back through recent commands just by using the up and down arrows and then re-issue them by tapping Enter. You are also able to search for commands by pressing Ctrl+R.

Quick tip

One command that's invaluable is **man** which is short for 'manual'. This will open up the help file for any other command. So if you want to know all the options for the **ls** command, simply type **man ls** and see what comes up.

YOUR LINUX PROBLEMS SOLVED!

We tackle the most common issues and presents a collection of solutions that's sure to keep your system running smooth and trouble free. Well, mostly.

Stalled boot

Q My computer won't boot. The splash screen comes up and then it just sits there.

A The first thing to do is be patient – the boot process may be trying to initialise a piece of hardware, or access the network before it's available. These can time out after a few minutes, either allowing boot to proceed or at least giving you an error message to work with. Unfortunately, the splash screen can hide any error messages. You can disable the splash screen at the boot menu: if your computer starts booting without showing a menu, hold down [Esc] when it boots to call up the menu. If there's a recovery or safe option, pick that, as it disables the splash screen and often sets safer (if potentially less efficient) boot options. Otherwise, put the highlight on your normal boot option and press [E] for edit. Find the line that starts with **linux** and look for options like quiet, splash or theme. Delete these options and press [F10] to continue booting.

This won't fix your problem (unless the problem is the splash screen), but it will let you see what is going on and either get an error message or an idea of where the boot is failing. If it appears to be a hardware detection issue, try disconnecting all unnecessary hardware, such as printers, scanners or USB rocket launchers. Once you can get past the problem, you can start to fix it.

Need help?
Visit us for friendly advice at:
linuxformat.com /forums

Boot errors

Q When I boot I see error messages along the lines of

```
Pci Express Device Error. Severity NN. Uncorrected non Fatal. Pci Bus Error
Type:N Transaction Layer. Flow Control Protocol N:First. Receiver 1
```

This is a problem I'm experiencing on a newly bought motherboard that works fine with Windows.

A Hardware working with Windows is no indication that it's fault free. The "It works with Windows, ship it and we'll fix any other problems with an update" attitude is not uncommon. With anything motherboard related, the first step is to check whether there's a BIOS update available. Even with a brand new motherboard, by the time it has travelled through the shipping and distribution process there is usually a new BIOS available. Flashing a new BIOS can fix boot problems but if yours persist, there is more you can do.

You can change the way the kernel addresses the hardware with a number of parameters at boot time. At the *Grub* boot menu (press [Esc] if it does not show up by default), highlight your normal boot option and press [E] (for edit). Edit the line starting with **kernel** by adding some or all of these options to the end of the line, separated by a space:

```
noapic
```

or

```
noapictimer, acpi=off
```

or

```
acpi=force
```

```
irqpoll
pci=nommconf
pci=nomsi
```

Note the difference in spelling between the second and third set of options – about all that **acpi** and **apic** have in common is the letters in their names. Press [F10] to boot with the options you added. Once you find a combination that has the kernel and your hardware co-operating with one another, make the change permanent by adding the options to GRUB_CMDLINE_LINUX_DEFAULT and GRUB_CMDLINE_LINUX in **/etc/default/grub**, saving the file and updating *Grub* with one of these commands. Which one depends on your distro – if the first one fails, try the second and then the third.

```
sudo update-grub2
sudo grub2-mkconfig -o /boot/grub2/grub.cfg
sudo grub2-mkconfig -o /boot/grub/grub.cfg
```

Re-installing Windows broke Linux

Q I installed Linux onto my Windows computer and it gave me the choice of which to use when I turned it on. I recently had to reinstall Windows and now Linux has disappeared. I don't want to reinstall Linux as I would lose all my documents and settings. Is there any way to get my Linux back?

A Linux uses the *Grub* bootloader to choose between different boot options, but Windows uses its own bootloader. A bootloader is a small piece of code embedded at the start of your hard drive that kickstarts the process of loading your operating system. When you installed Windows, it installed its own bootloader, overwriting your existing *Grub*. Don't worry, it's only the bootloader code that's gone – your Linux system and files are still there, and all you need to do is restore *Grub*. Some distro install discs have an option to repair the bootloader, or there's a more generic option in *Rescatux* (**www.supergrubdisk.org/rescatux**). Boot from the *Rescatux* CD and select the Restore Grub option. This will search for your Linux partitions and ask you to choose the partition that contains the root filesystem of your Linux installation and then ask which hard disk you want to install *Grub* on – usually **/dev/sda**, the first disk. The window may go blank for a worrying amount of time but do not panic or reboot: wait until it tells you *Grub* was re-installed before rebooting, and you should then get your old boot menu back, with access to your untouched Linux system.

Dual-booting more than one distro

Q I know I can dual-boot between Linux and Windows, but can I do the same with more than one Linux distro? If so, how do I share my data between them? How should I partition my disk?

A Yes, it is possible. Each distro needs its own **root** partition but you can share the same **swap** between them all. Having a separate partition for **/home** is recommended – you don't want to tie your data to a single distro you may later decide to stop using. While a separate **/home** partition is desirable, don't use the same **home** directory for each distro as there may be incompatibilities between different versions of the software used. I prefer to put the distro name in the username, as in "fred-ubuntu", "fred-debian", etc. As long as each user has the same UID – most installers make the first user 1000 – you'll have write access to all of them, then you can create symlinks for important data directories, like documents, music, photos and so on, to share these files between all distros while keeping the configuration files separate.

As far as the boot menu is concerned, pick one distro as your main one and let that handle booting. Every time you install a new distro, skip the option to install a bootloader if possible. Then reboot into the main one and run:

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

to update the boot menu.

```
  GNU nano 2.2.6                File: /etc/default/grub                          Modified

# If you change this file, run 'update-grub' afterwards to update
# /boot/grub/grub.cfg.
# For full documentation of the options in this file, see:
#   info -f grub -n 'Simple configuration'

GRUB_DEFAULT=0
GRUB_HIDDEN_TIMEOUT=0
GRUB_HIDDEN_TIMEOUT_QUIET=true
GRUB_TIMEOUT=10
GRUB_DISTRIBUTOR=`lsb_release -i -s 2> /dev/null || echo Debian`
GRUB_CMDLINE_LINUX_DEFAULT="quiet splash add stuff here"
GRUB_CMDLINE_LINUX="and here"

# Uncomment to enable BadRAM filtering, modify to suit your needs
# This works with Linux (no patch required) and with any kernel that obtains
# the memory map information from GRUB (GNU Mach, kernel of FreeBSD ...)
#GRUB_BADRAM="0x01234567,0xfefefefe,0x89abcdef,0xefefefef"

# Uncomment to disable graphical terminal (grub-pc only)
#GRUB_TERMINAL=console

# The resolution used on graphical terminal
# note that you can use only modes which your graphic card supports via VBE
# you can see them in real GRUB with the command `vbeinfo'

^G Get Help    ^O WriteOut    ^R Read File   ^Y Prev Page   ^K Cut Text    ^C Cur Pos
^X Exit        ^J Justify     ^W Where Is    ^V Next Page   ^U UnCut Text  ^T To Spell
```

› To make permanent changes to your Grub menu, edit /etc/default/grub and run grub2-mkconfig or update-grub.

Removing a distro broke everything

Q I deleted Linux partitions from my computer, in preparation for installing something else. Now I cannot boot it. All I get is a message that *Grub* cannot find files and a prompt saying grub rescue>**. I can't even boot into Windows.**

A While the initial boot code for *Grub* is stored in the Master Boot Record (MBR) at the start of your hard disk, this is a very small space. So the initial *Grub* code loads files and configuration information from the **grub** or **grub2** directory on your Linux boot partition. When you removed that partition, you left *Grub* nowhere to go. If you are planning on installing another distro in the now-empty space, simply run the installer and it will set up *Grub* for you, detecting your Windows system too, and all will be well after a reboot.

If you are using the space for an existing Windows or Linux installation, you need to either set up *Grub* for the other Linux system or reinstall the Windows bootloader. The easiest way to reset *Grub* is use the Restore Grub option on the *Rescatux* live CD – choose the Linux installation you want to be responsible for holding the *Grub* files. Restoring the Windows bootloader involves booting the Windows disc (or rescue partition) and running one of these commands, depending on the Windows version. For Windows XP it's the following:

```
fixboot
fixmbr
```

For Vista and Windows 7 and 8 it's:

```
bootrec /fixboot
bootrec /fixmbr
bootrec /rebuildbcd
```

When your computer won't even boot, a live CD is a godsend, so it pays to always have one handy. A live CD is a complete Linux environment that boots from a CD and needs no installation. Apart from being very handy, a live CD is ideal for a rescue system. There are some dedicated rescue CDs, such as System Rescue CD, Rescatux and Grml, but most distro installer discs nowadays boot into a live environment that contains enough software to fix many problems. »

Hardware difficulties

Solving problems with recalcitrant bits of kit.

» Finding hardware drivers

Q Where do I get drivers for my hardware? The accompanying disc includes Windows drivers only.

A Most hardware doesn't need separate drivers with Linux: the kernel includes drivers for a massive range of hardware. Nine times out of ten, the kernel recognises the hardware when you plug it in and loads the appropriate driver without any input from you. Some hardware does have separate drivers, often because they use proprietary code that cannot be included with the kernel. In such cases, a live CD, such as Knoppix or System Rescue CD, will often recognise the device – these tend to include a lot of extra drivers. If the hardware is an internal device, it will usually show up in the output from **lspci**, and adding the **-k** option shows which driver is being used.

Once you know that, you can search for which package provides that driver, either using your package manager or a web search. If it's in a standard package that's not installed by default, as is the case with many wireless drivers, installing the package should be all you need to get the hardware working. The process is similar with USB devices except you use the **lsusb** command, and there is no handy **-k** equivalent. Note that on some systems, these can only be run as root.

If there's no driver support by your kernel or package manager, you have to enter the murky waters of downloading and installing drivers directly from a manufacturer's website. If they provide a Deb or RPM package suitable for your system, this isn't too bad. Otherwise, make sure you get the right driver for your architecture and follow the installation instructions carefully. Drivers have to be linked into the kernel, so you will probably need the kernel source and build-essential packages installed.

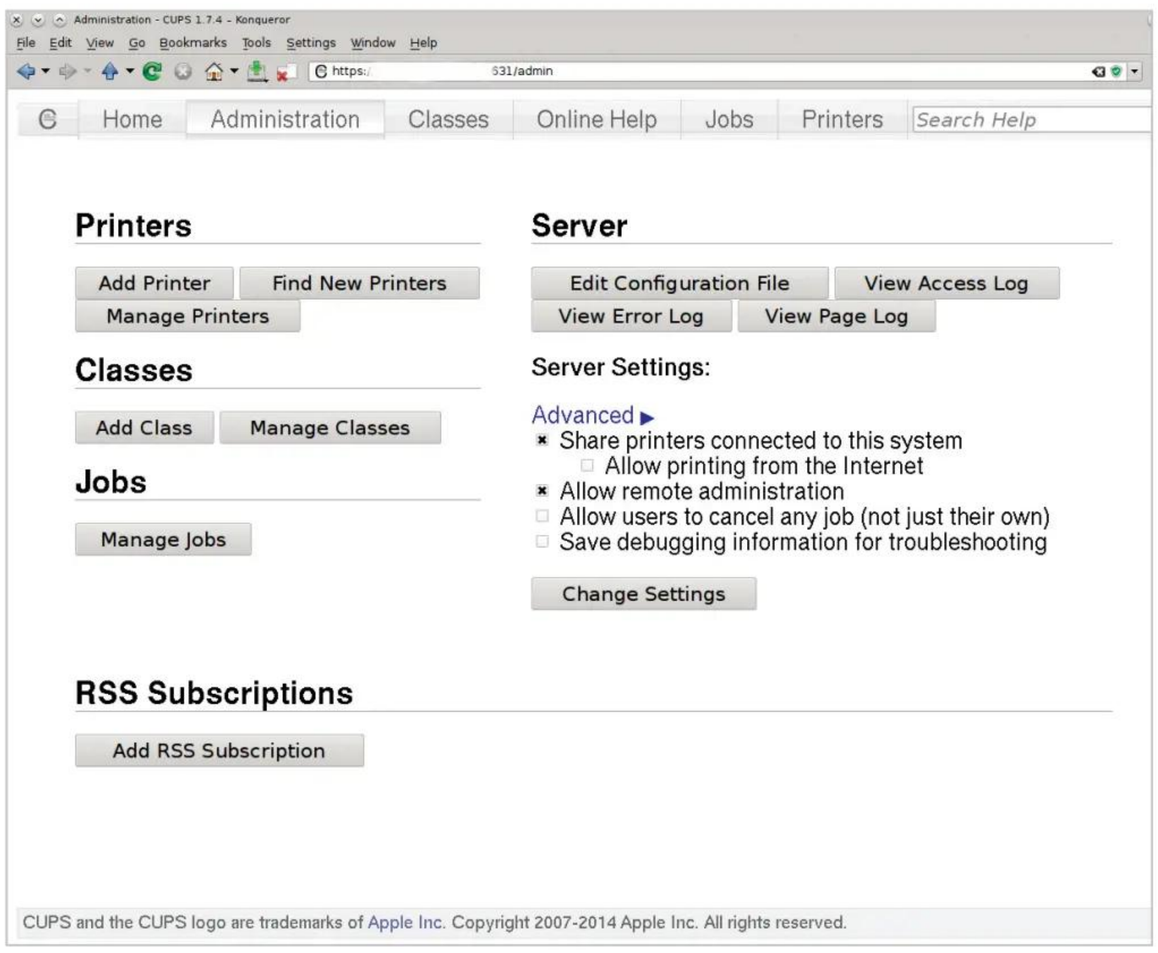

› While most distributions have a GUI for managing printers, the CUPS user interface is universal and works well.

Random crashes

Q My computer is suffering from random crashes and lockups – it doesn't seem to be related to any particular piece of software.

A This may well be a hardware issue and there are a number of common causes. Faulty memory can cause this and, because of the way Linux makes use of all available memory, is more likely to do so in Linux that Windows. Install *Memtester* from the terminal with **sudo apt-get install memtester**. Let this run for at least a couple of passes, preferably overnight. Overheating is also a common cause: try running a system monitor or Taskbar applet that shows your system temperature. Make sure all vents are clear and, if you are comfortable with taking a screwdriver to your case, check that your CPU heatsink and fan are not clogged.

A third cause, for desktop computers, is a substandard power supply. Cheap units can often fail gradually, giving just this sort of behaviour and damaging other components on their way down. The only reliable test is to try a different PSU.

Printer problems

Q My distribution hasn't spotted my printer. How do I install the drivers?

A How you install drivers – or even if there are any available – depends on your printer. The first place to go is **http://linuxprinting.org** to check their database. If you do this before buying a printer, it will save you a lot of trouble. Enter the make and model of printer and their database will tell you if it is supported, and which drivers to use. CUPS, the printing system used by Linux, comes with many drivers, so there's often no need to install anything. There are further

Why use the terminal?

Much of the advice we give here is in the form of terminal commands. While most distros have their own configuration programs, they vary considerably while the underlying commands they call remain constant across all distros. By cutting out the middle man and running those commands directly, the solutions we give are portable across all distros. It also means they work even if your problem is preventing the desktop from running.

Some commands need to be run as the root user, which is done either by prefixing the command with **sudo** (when you will be prompted to give your password) or by running **su** in the terminal first, which will ask you for the root password. We will use the sudo method throughout this article, as it is the only option with some distros. If you have full root access through **su**, simply run the command without the **sudo**. Commands given without **sudo** should be run as your normal user.

drivers in the gutenprint package – **http://linuxprinting.org** will tell you if you need this. HP provides open source drivers for its printers, so if you have an HP model, install the *hplip* package to make all HP printer and scanner drivers available. Other manufacturers provide closed source binary packages. Check whether your distro supplies these as a package, or otherwise, as a last resort, you will need to download and install them by following the instructions on the printer manufacturer's website.

Once you have the proper drivers installed, you can set up the printer using either your distribution's printer administration tool or the CUPS web interface. To use the latter, which will work on any distro, point your browser at **https://localhost:631**, log in if asked, click on the Administration tab, then Add Printer and follow the prompts, bearing in mind the information you already have about the best driver to use.

The solid option

Q I'd like to replace my hard drive with an SSD. How do I set it up and is there anything I need to be careful of? I'm thinking about the problems caused by writing to flash memory too often. What about TRIM?

A There isn't much different to be done with setting up an SSD compared with a spinning drive these days. Disk partitioning tools default to values that are safe for both types of drive now. Most of the data around now suggests that an SSD is capable of outlasting a spinning drive. However, it won't hurt to add **noatime** to your mount options in **/etc/fstab**, if you don't already: it gives a speed boost too.

You have a couple of options for copying the data. You could back up your **home** directory and do a fresh installation on the SSD, or you could copy everything across. You should boot from a live CD to do this, and use a SATA to USB adapter if you don't have a spare SATA port. Mount each partition, apart from **swap**, in both drives, and use **rsync** to copy everything across for each partition:

```
rsync -a /mnt/old_partition/ /mnt/new_partition/
```

You'll also need to reinstall *Grub* to the bootloader, by booting something like *Rescatux* with only the new drive in the computer. TRIM, cleaning up no-longer-used areas of the storage, can be done in two ways. The easiest is to add **relatime,discard** to the mount options in **/etc/fstab**.

The only drawback, however, is that this can slow the drive down while trim operations are ongoing, so many prefer to run *fstrim* when the computer is lightly loaded. A *Cron* task that runs **fstrim --all** once a week should be enough for most usage. The **--all** option trims all filesystems that support it and silently ignores others. Add the **--verbose** option and *Cron* will mail you a summary of what was done.

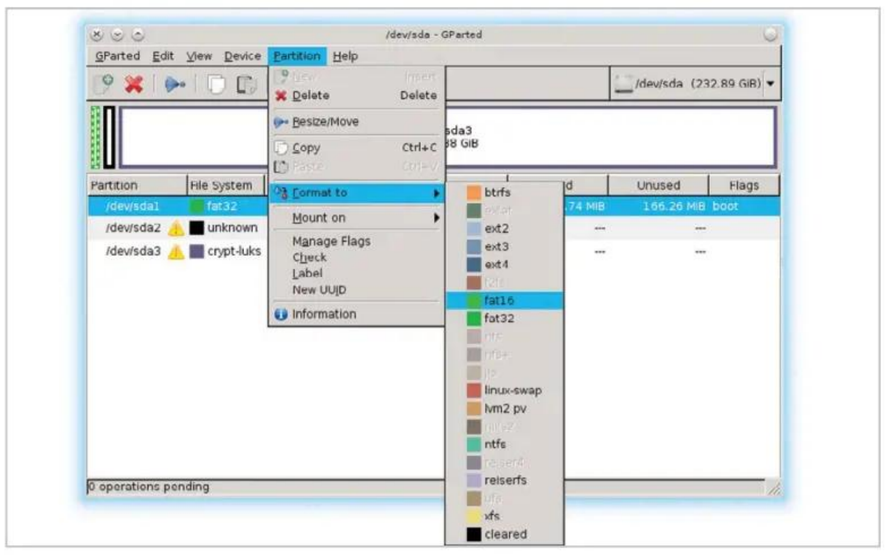

› If you want to wipe a hard disk beyond all hope of recovery, *Darik's Boot and Nuke* (*DBAN*) is a good option.

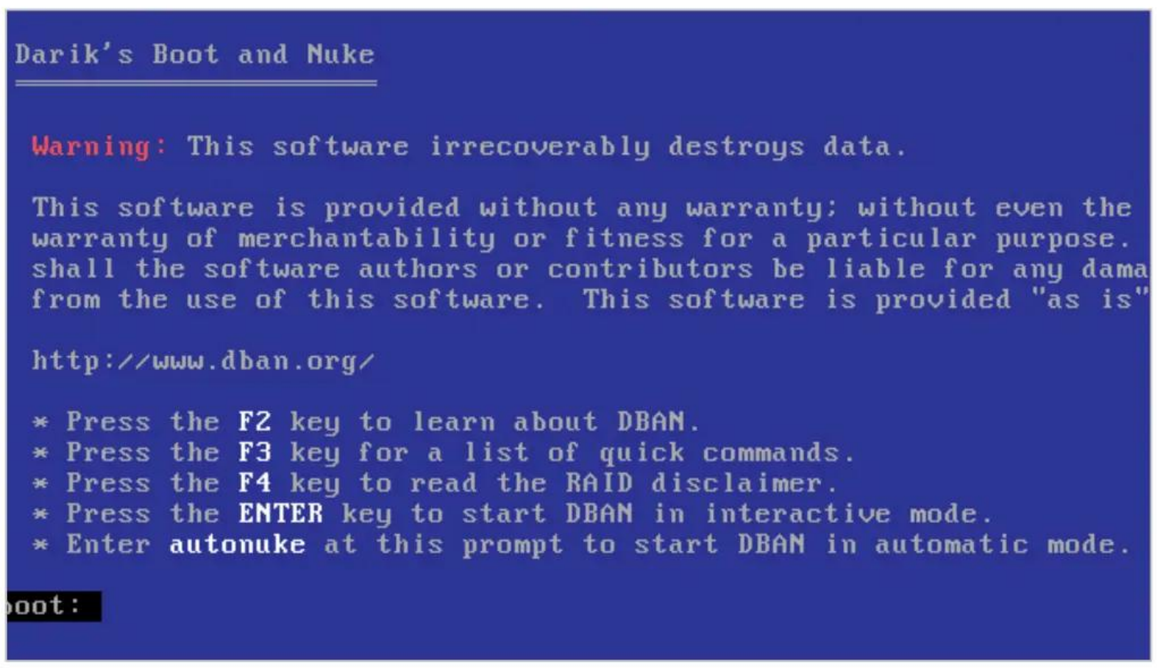

› Partitioning a new SSD, or any type of hard disk, is easy with *Gparted* – most distros and live CDs include it.

Deleting files for real

Q I would like to make sure my hard drive is completely wiped before disposing of it. I understand that simply deleting all files is not enough, so what should I do to clear all trace of my data from the drive?

A It depends on how you intend to dispose of the drive. If you are throwing it away, a large hammer is pretty terminal to your data, but there are more sophisticated – if decidedly less satisfying – ways of securely clearing a drive, methods that also leave the drive in a state suitable for resale.

> **"Most of the data around now suggests that an SSD is capable of outlasting a spinning drive."**

You are right in thinking that deleting files isn't enough, as that only removes their entries from the directory table, and the data is still there. You need to overwrite all the data on the drive. Conventional wisdom states that you need to do this several times, overwriting with different patterns, or random data, to prevent even the most determined of data recovery attempts. For most purposes, overwriting the drive with zeroes will do:

```
sudo dd if=/dev/zero of=/dev/sdb bs=4k
```

This cannot be run on the system drive, so if that's the one that you wish to wipe, you'll need to run the command from a live CD instead. You can give yourself further peace of mind by adding in some random data, by using **/dev/urandom** instead of **/dev/zero**. Be aware that the **dd** command doesn't ask if you are sure or double-check to see if you've given it the right device before pressing Enter – there are no second chances here.

If you still want to go down the route of multiple overwrites for peace of mind, try the *Darik's Boot and Nuke* (DBAN) live CD from **http://www.dban.org**, which is the next most effective thing to that large hammer you're waving about. »

System issues

Fix common difficulties with general system admin.

» Drive filling up

Q I am running out of space on my hard drive – the space just seems to be disappearing. Where did it go and how can I get it back?

A If you have more than one partition, the first step is to see which is affected, with the **df** command:

```
df -h
```

This shows the total, used and available space on each partition. If you have a separate **/home**, you now know whether it's your home directory or the system root that's filling up. The steps taken are the same in each case, although you need to be root to clean up the system directories.

For the system files, the usual culprits are **/tmp** and **/var**. The first is used for temporary files, but sometimes old files get left in here, especially when processes fail without cleaning up. **/var/log** contains system log files and can fill up, particularly if a process is logging lots of errors. Look at the contents of any large files in here before you deal with them. It's worth installing *logrotate*, which keeps your log files under control. Its default setup starts a new log file each week and compresses the last one, keeping the last four weeks of logs. This can save a huge amount of space.

In your **home** directory, check your **downloads** folder, your browser and mailer caches and also the **.cache** directory. You can use **du** to check the sizes of directories:

```
du -sh * | sort -h
```

lists everything in size order. If you want to check on hidden directories, use:

```
du -sh .[0-9a-z]* | sort -h
```

There are also graphical tools to visualise disk usage – *Filelight* for KDE and *Disk Usage Analyser* for Gnome – while *Bleachbit* is good for identifying and cleaning up unnecessary files, like browser caches and other temporary files.

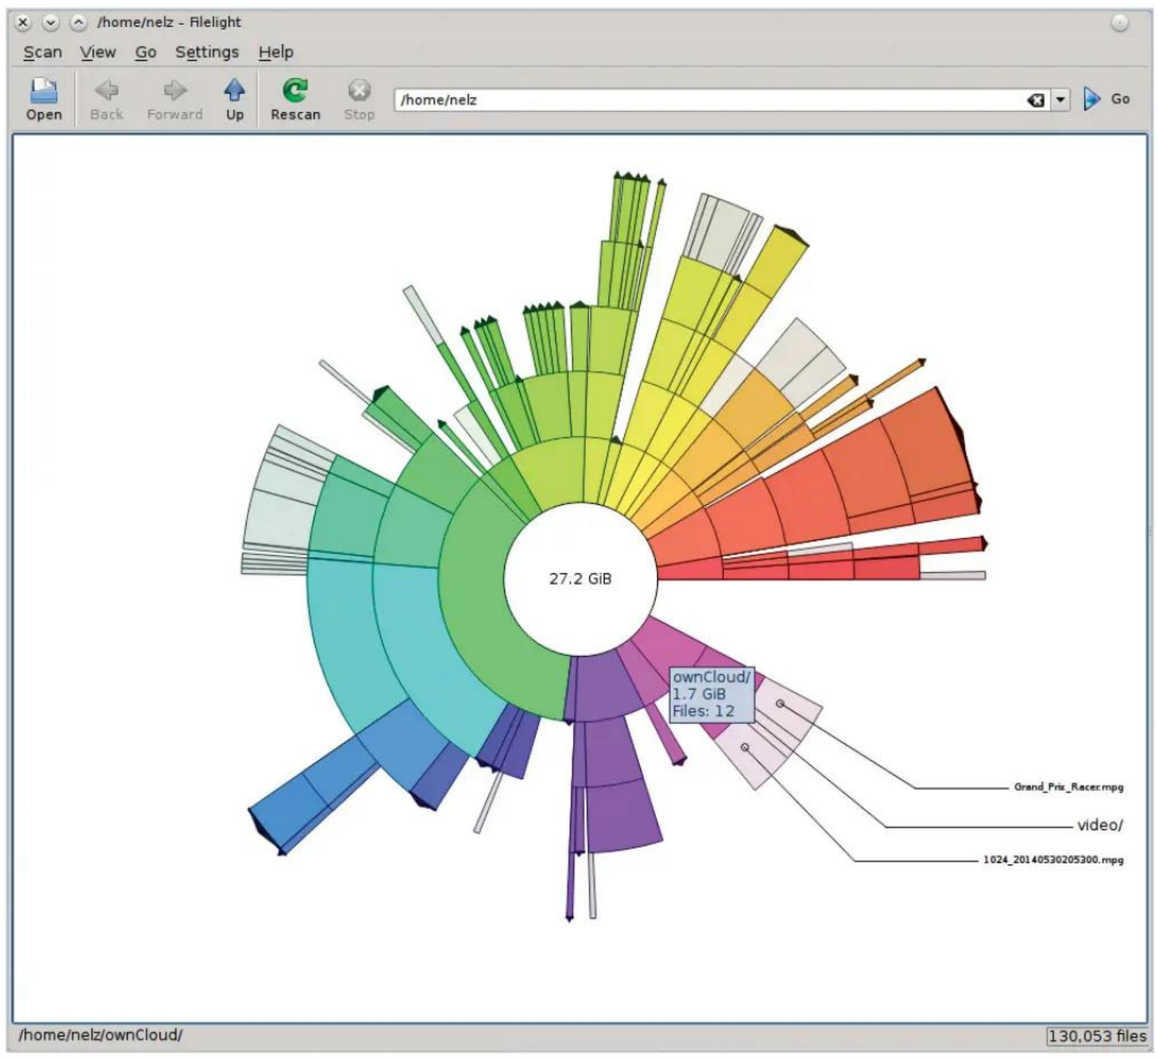

› ***Filelight*, like *Gnome Disk Usage Analyser*, can give you a good graphical representation of how you managed to fill that hard drive.**

Where has all my memory gone?

Q I have noticed that after running Linux for a while, I've hardly any memory left. Where is it going? How can I find the culprit and what can I do about it?

A This is actually perfectly normal and the way Linux is designed to work. Free memory is memory that's doing nothing when it could be helping the performance of your computer. The Linux kernel allocates memory that isn't in active use by processes to things, such as filesystem caches in order to improve performance. You can see this by looking at the output from **free -h**.

```
             total   used   free   shared  buffers  cached
Mem:         15G     15G    128M   982M    396K     13G
-/+ buffers/cache:   2.3G    13G
```

On a computer fitted with 16GB of physical memory, only 128MB is reported as free. Of the memory in use, 13GB of it is in use as buffers. These greatly speed up disk access by caching written data in memory and reading in more than you need in anticipation of your next read. If this memory is needed for anything more important, it will be released immediately.

Looking for trouble

Linux logs just about everything. When something goes wrong, these logs are the places to look. The system log is in **/var/log** and usually called messages or syslog, depending on the logger used by your distro. If program *foo* is misbehaving, try looking for messages with:

```
grep foo /var/log/messages
```

You can also look at the log messages in real time with

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

which will display new entries to the log as they appear (**-f** means follow). Run this in a terminal and then do whatever is causing your problems to see what it says. The kernel buffer can be read with the **dmesg** command, which is useful if you're having difficulties with hardware. It too has a "follow" option

```
dmesg --follow
```

which is useful when dealing with hot-plugged hardware. To see how built-in hardware is detected and handled at boot time, run

```
dmesg | less
```

after booting, then either scroll through the messages or, as there are a lot, press **/** to search for something relevant. If your distro uses *systemd*, log entries can be displayed with **journalctl**. You can do the same follow stuff with

```
journalctl --follow
```

or look for errors since the last boot with:

```
journalctl -b -p err
```

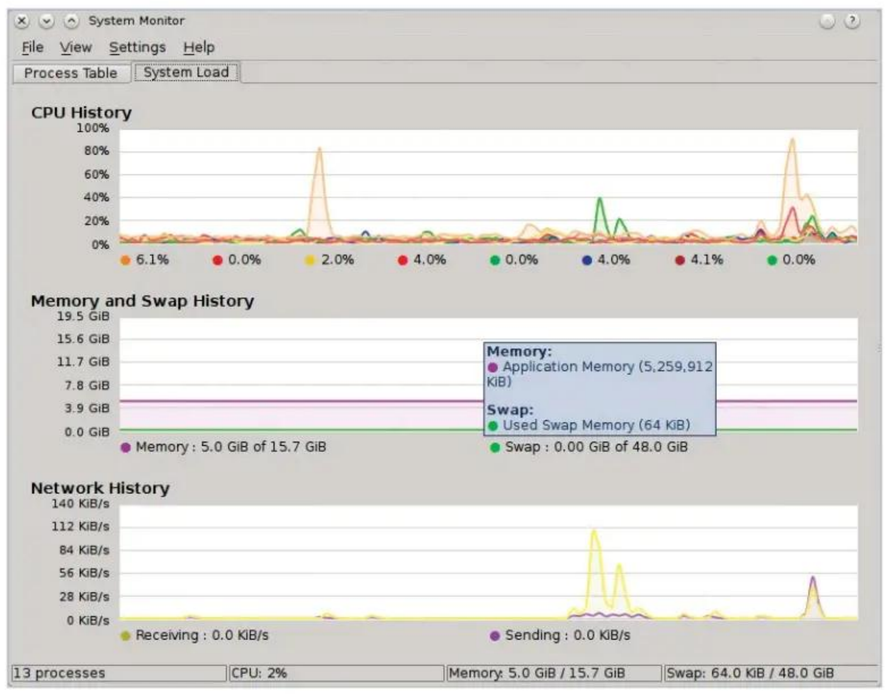

› When using system monitors, make sure you look at figures that exclude memory used by filesystem buffers.

It's not just filesystem data that remains in memory – try checking your free memory, then starting and stopping a large program like *Firefox* and you will see that your free memory decreases. Now start *Firefox* again and see how much quicker it starts the second time. This is because the kernel keeps data in memory ready for the next use. The same applies to shared libraries, which is one reason why sticking to programs from one desktop environment can be faster than running a mixture, especially with limited memory.

Moving home

Q I have heard that it's better to have my home directory on a separate partition. Why is this and how do I do it?

A First of all, it's not normally your individual home directory that has its own filesystem but the complete **/home** directory, in which your home resides. There are a number of reasons for doing this, such as simplifying backup procedures, but the main reason is to keep your data, and that of any other users, separate from the operating system. That means you can upgrade or even replace your OS and still keep all your personal settings and data.

The simplest way to do it is with *Gparted* and a separate hard drive. Make sure the hard drive is formatted with a Linux filesystem, to preserve all the information saved with your files, then use **rsync** to copy everything over.

```
sudo mkdir /media/extdrive/home/
sudo rsync -a --delete /home/ /media/extdrive/
```

This assumes your separate drive is mounted at **/media/extdrive**. The trailing slashes are important. There are other tools you can use for the copy, but **rsync** has the advantage that it can carry on where it left off if the process is interrupted. This means you can carry on using your PC while you do this, then boot from a live CD and repeat the process to update the files altered while you were using the system.

A live CD is needed for the next stage, where you delete the contents of **/home** and use *Gparted* to resize the root partition to a reasonable size – 20GB should be enough, but make it around 50% larger than the space currently used. Now create a new partition after the root partition and format it, using **ext4** unless you have a reason to do otherwise.

Load the **/etc/fstab** file from the root partition into an editor and add an extra line:

```
/dev/sdaN  /home  ext4  defaults,noatime  0 0
```

where **sdaN** is the partition you just created. Finally, copy the data back from the external drive into the new **/home** and reboot. It is possible to do this without access to an external drive, but it is a more fiddly process.

Universally Unreadable IDs

Q I find the use of UUIDs in /etc/fstab **annoying. What's wrong with the old method of using the device nodes, like** /dev/sda1**, and can I change back?**

A The trouble with device nodes is that they can change. If you add or remove a drive, including USB flash drives, or change the BIOS boot order, **sda** can become **sdb**. UUIDs on the other hand are (fairly) unique; it doesn't matter if the drive order changes, that unfriendly string stays the same. And that's the problem: UUIDs are good for computers as they give a simple, reliable identification, but they are not human-friendly. A compromise is to use filesystem labels. If your root partition is on **/dev/sda1** you could run

> **"Free memory is memory that's doing nothing when it could be helping performance."**

```
sudo e2label /dev/sda1 root
```

Then change the device specification in **/etc/fstab** from **"UUID=youllneverreadorrememberthis"** to **"LABEL=root"**. Unlike many operation on the root filesystem, this one is safe to perform on a running system. It will take effect when you reboot. If you are ever likely to move drives between computers, use less generic label than 'root'.

Hard drive copying

Q How do I clone a complete hard drive? I am using Ubuntu but the drive also contains Windows.

A Cloning a drive that the system is running on is a risky proposition. The OS could, and probably will, change files on the drive as you are trying to clone it. The simplest and safest way of doing this is with *Clonezilla*, a dedicated live CD distro. Burn the ISO to a CD and boot from it. Then use the menus in *Clonezilla* to make a clone of your hard drive, either directly to another drive or to an image file. »

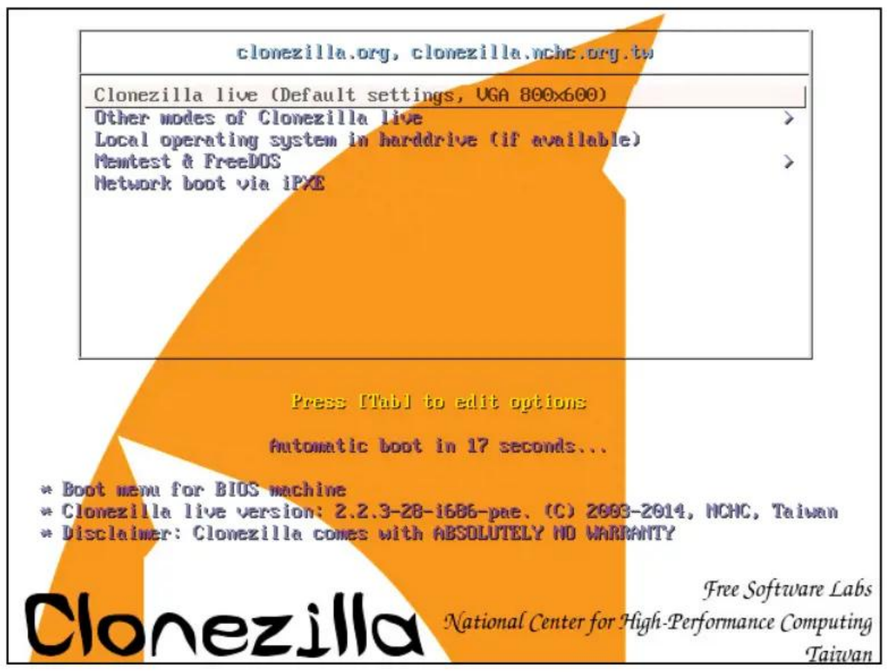

› If you want to make a copy of a complete hard disk, *Clonezilla* is the usual first choice.

Software glitches

Software misbehaving? Or missing? We try to help.

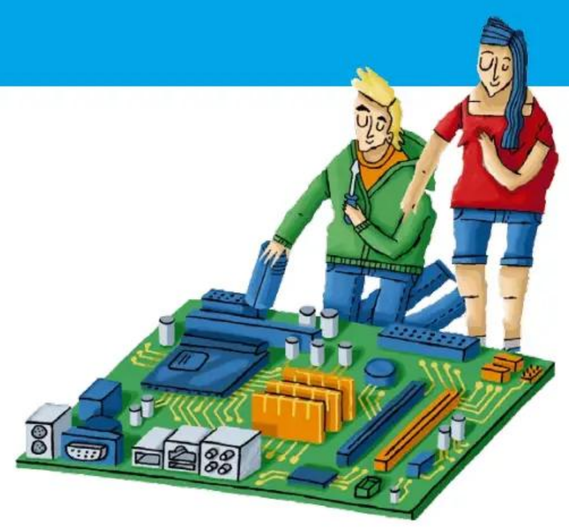

» Installing software

Q I'm new to Linux and cannot see how to install software. On Windows, I go to the program's website, download an .exe and run it. What do I do?

A Software installation works differently in Linux. Programs are generally provided as source code and the individual distros create program packages from this. That means you know the program is compatible with the rest of the software from your distro. It also means you use the distro itself to install. Look for a program called *Software Manager* or similar: this will let you search for and install any software you need. It will also take care of any dependencies (other programs needed by the one you are installing) and notify you when updates are available. There is no need for you to check back on web pages or for programs to 'phone home'.

Use the source

Q I want to install a program but there's no package for my distro and the program's website only provides source code. How do I install it?

A The precise method varies from one program to another. The first step is always to unpack the tarball or ZIP file and look for a file called **Readme**, **Install** or similar. This will explain how to install the software. For 90% of projects, the standard *autotools* system is used, which requires three commands to build and install the software. These must be run from inside the directory created when you unpacked the tarball:

```
./configure
make
sudo make install
```

The first command checks that your system has all the required dependencies and sets up the build environment to suit your system. The second compiles the source code into programs to run on your computer, and the third command installs those programs to the correct locations and handles any final configuration.

You need a standard build environment installed to do this. Most distros have a package called *build-essentials*, or something similar, that installs everything you need to build and install programs from source.

Don't throw away the source directory after installation. If you want to remove the software at a later date, go back in there and run:

```
sudo make uninstall
```

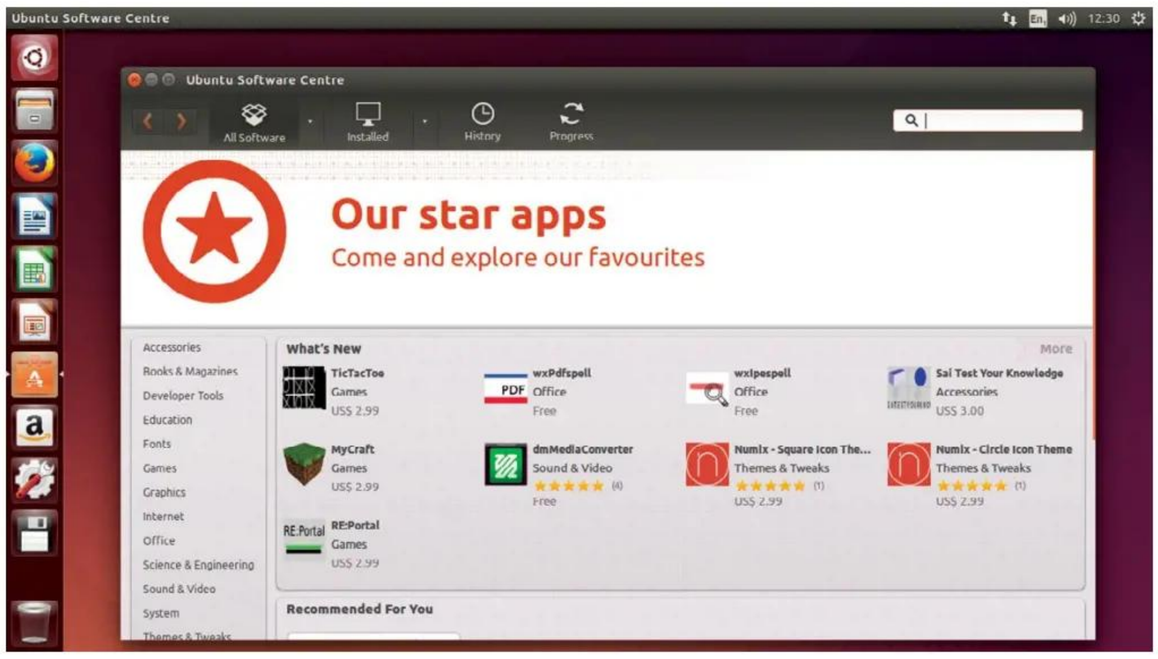

› Don't go trawling the web for software from dubious sources. Most software you need can be found in your distro's software manager.

No DVD drive

Q I have an ISO image that I can burn to a CD/DVD, but I want to use it with a slimline laptop with no DVD drive. Can I use a USB stick instead?

A You can do this. The exact approach depends on the type of ISO image. Some distros now release what are called hybrid ISO images that work with either medium. You copy them to a DVD using whatever DVD burning software you prefer, or you can copy the file to a USB stick with **dd**

```
sudo dd if=distro.iso of=/dev/sdb bs=4k
```

where **if** is the file to copy and **of** is the destination (**bs=4k** speeds up the process quite dramatically, but isn't compulsory). Warning: **dd** does exactly what you tell it to, with no 'are you sure' niceties/annoyances, so make sure you have the correct destination drive.

How do you know you have a hybrid ISO? The distro's website may well say so, but you can test with **fdisk**:

```
sudo fdisk -l distro.iso
```

If the line under the Device line includes the name of the ISO image, you have a hybrid. If you don't, you can convert it with *isohybrid*. You may need to install the *syslinux* package first, then run:

Run it from a shell

If a program fails to run from the desktop launcher, or runs but later crashes or gives other problems, try running it from a terminal. This will not magically make it work correctly, but most programs send output to the terminal when things go wrong, so it should give you a clue to the problem or at least something to enter into a search engine. This assumes you know the name of the command to run it from the terminal, which is often a lower case version of the program title, so to run *Firefox* you would type **firef**[TAB] to get the correct name. If you cannot guess the name this way, you can list which programs the package installed into your command path. With Debian derivatives the command is, for example:

```
dpkg -L firefox | grep bin/
```

That's an upper case L – lower case is used for searching package files instead of installed packages. The RPM equivalent is:

```
rpm -ql | grep bin/
```

The **grep** part shows only those files installed into one of the ***bin** directories, instead of the full package list.

```
isohybrid distro.iso
```

Now **fdisk** should show it as a hybrid and you can **dd** it to your USB stick.

Go slow

Q My computer seems to be running with one of its CPU cores tied behind its back. What is going on?

A The most likely cause is that a particular process is hogging a lot of your resources, and the program to investigate this is called *top*. Run it in a terminal and it will show you a summary of overall system usage and a list of programs in order of their CPU usage. The third line will show you the overall CPU usages, which will be something like this (on a lightly loaded system):

```
%Cpu(s): 0.7 us, 0.2 sy, 0.0 ni, 99.1 id, 0.0 wa, 0.0 hi, 0.0 si, 0.0 st
```

The first three figures show how much of your CPU capacity is being used for user, system and nice tasks. A nice task is one running at a lower priority, so it uses the CPU only when nothing else wants it. The fourth figure is 'idle' – the amount of CPU time not being used. If this is close to zero, you now know why your computer is running slowly – look at the list of tasks to see which is eating the most CPU time. If you have a multi-core CPU, pressing **1** switches between showing an overall usage and separate figures for each core. That **wa** (waiting) figure is also significant: it shows how much time the CPU is waiting for I/O and can be high on disk intensive operations. If it is routinely high, you may have a disk problem – either read errors or a nearly full filesystem that has become heavily fragmented.

Look at the memory line. If memory is full and swap is in use, you will lose performance. The task list is ordered by CPU usage, but you can press F to pop up a list of fields. Move the cursor over the field you want to sort on and press S followed by Q. It sounds fiddly but the window opened by F allows for a lot more than changing sort orders. If memory is low, this will help you find the greedy programs.

Lost password

Q I reinstalled Linux recently and thought I kept the same password, but I cannot log in with it. Is there a way I can find or reset my password?

A You can reset your password by booting from a live CD/DVD – either a specific rescue disc like System Rescue CD or your distro's own disc. The first step is to identify the partition where your distro is installed (this is the root filesystem for the distro; your **home** directory is not needed).

```
top - 12:32:02 up 1 day,  3:55, 10 users,  load average: 0.88, 0.81, 0.64
Tasks: 329 total,   4 running, 325 sleeping,   0 stopped,   0 zombie
%Cpu(s): 11.9 us, 13.6 sy,  0.0 ni, 74.5 id,  0.0 wa,  0.0 hi,  0.0 si,  0.0 st
KiB Mem:  16426484 total, 16298072 used,   128412 free,     1064 buffers
KiB Swap: 50331636 total,       64 used, 50331572 free.  9982416 cached Mem

  PID USER      PR  NI    VIRT    RES    SHR S  %CPU %MEM     TIME+ COMMAND
 1501 nelz      20   0  527856  61644  42968 R  91.0  0.4   4:37.09 konsole
 9612 nelz      20   0    6008    640    560 R  52.5  0.0   0:09.47 yes
 4692 root      20   0       0      0      0 S  16.9  0.0   0:01.85 kworker/0:1
 9618 root      20   0       0      0      0 S  12.0  0.0   0:00.60 kworker/1:0
 1876 root      20   0       0      0      0 S   5.6  0.0   0:01.19 kworker/3:28
  900 root      20   0  328960 139884  88972 S   4.3  0.9  22:40.02 X
 8958 root      20   0       0      0      0 S   3.0  0.0   0:00.48 kworker/5:9
 9480 nelz      20   0 2581628 1.356g  17316 S   2.3  8.7   1:09.18 qemu-system-x86
 8567 root      20   0       0      0      0 S   2.0  0.0   0:00.16 kworker/7:10
 9616 nelz      20   0  264480  37280  31612 S   2.0  0.2   0:00.15 ksnapshot
  769 root      20   0  451260  39244  11724 S   1.7  0.2  45:53.79 motion
 1087 nelz      20   0 6528944 269324  99928 R   1.7  1.6  25:54.12 plasma-desktop
 4676 root      20   0       0      0      0 S   1.7  0.0   0:01.31 kworker/2:2
 6851 root      20   0       0      0      0 S   1.7  0.0   0:00.14 kworker/3:1
 6909 root      20   0       0      0      0 S   1.3  0.0   0:00.07 kworker/2:11
 8027 root      20   0       0      0      0 S   1.3  0.0   0:00.17 kworker/1:5
 9447 root      20   0       0      0      0 S   1.3  0.0   0:00.19 kworker/4:0
 1726 root      20   0       0      0      0 S   1.0  0.0   0:05.68 kworker/0:7
16622 nelz      20   0  917700 205456  53264 S   1.0  1.3  14:20.68 chrome
16633 nelz      20   0  934616 212564  53792 S   0.7  1.3  15:04.58 chrome
16685 nelz      20   0  882012  30004  21852 S   0.7  0.2   5:15.12 vlc
 1494 nelz      20   0  218124  29236  24160 S   0.3  0.2   1:41.05 kfilebox
```

› ***Top*** **gives an overview of resources used by your system as well as a list of the programs using the most.**

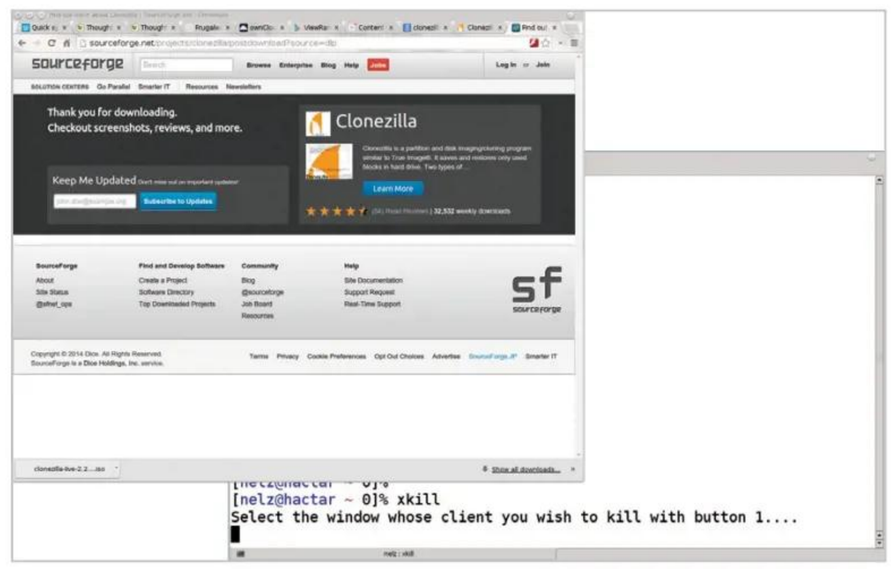

› ***Xkill*** **will kill the first window you click on, and the process behind it, so be careful.**

If you have booted to a graphical desktop, you can browse the available drives, looking at those of a suitable size until you find the one that contains directories such as **bin**, **etc** and **lib**. Once you have the correct partition mounted, say in **/mnt/myroot**, you can use **chroot** to fix it:

```
sudo chroot /mnt/myroot /bin/bash
```

This changes the root directory to the path you give and runs the command following. So this switches you to the root of your installed system and runs the *Bash* shell as root. You are now inside your installation and can change your password with:

```
passwd yourusername
```

It will ask you for the new password twice and then write it to disk. Press [Ctrl]+[D] to exit the **chroot**, reboot from your hard disk and you should now be able to log in.

Killing windows

Q I run some bleeding edge versions of software and occasionally they crash, leaving the window still open. How do I close a window and program like this so that I can run it again?

> **"Most distros have a package that installs everything you need to build and install programs."**

A For graphical programs, the simplest answer is *xkill*, which is usually installed by default. When you run *xkill*, your mouse pointer changes to a skull-and-crossbones, the next window you click on will be forcibly closed, along with its program, so be careful where you click.

There are also command line alternatives. **kill** takes a process ID as an argument and kills it. **Killall** does the same but works with program names, so be careful if you may have more than one copy of the program running. Both send a TERM signal by default, which asks the program to shut down cleanly. If this fails, add **-KILL** to the arguments to send the more forceful KILL signal:

```
killall -KILL someprogram
```

If you want to use **kill**, there are various ways to find the PID. If the process is using a lot of CPU or memory it will show up in top, and you can kill it from there by pressing K. Otherwise, use **pgrep**:

```
pgrep -fl programname
```

The **-fl** options mean you see the full command line, so you can be sure you are killing the right process. »

Getting connected

Now let's tackle some connectivity and Internet hurdles.

» Going slower

Q My Internet connection seems really slow – downloading files takes forever. Do I need to fix this in Linux or change my router's settings?

A Stay away from your router settings until you have tried everything else – you could end up breaking a fully working setup. The first step is to check that it really is slow, by using a service like **http://speedtest.net**. Next, check that nothing is sucking up your bandwidth in the background. A good program for this is *nethogs*, which you may have to install from your package manager:

```
sudo nethogs eth0
```

This gives a *top*-like display of what's using your bandwidth. If you're connecting to your router by wireless, try a wired connection to rule out Wi-Fi issues. If that makes a difference, try testing in different locations.

There's also the possibility of congestion with other wireless routers from your neighbours, which can be avoided by switching your router to a different wireless channel. You can see the wireless channels in use with this command

```
sudo iwlist wlan0 scan | grep -e Address: -e Channel: -e ESSID:
```

Iwlist produces a lot of output about each access point it discovers, hence the need to filter it with **grep**. You can also use something like *Wi-Fi Analyser* on Android, which gives a graphical display of the channel and strength of all access points in range, enabling you to map any points where the signal from your router is weak as well as identifying competition for bandwidth.

There's also the highly scientific test of picking up the phone: if the line is noisy, you will suffer reduced speeds and you'll need to report the problem to your line provider. If all the above checks show nothing, check with your ISP before tampering with previously working router settings.

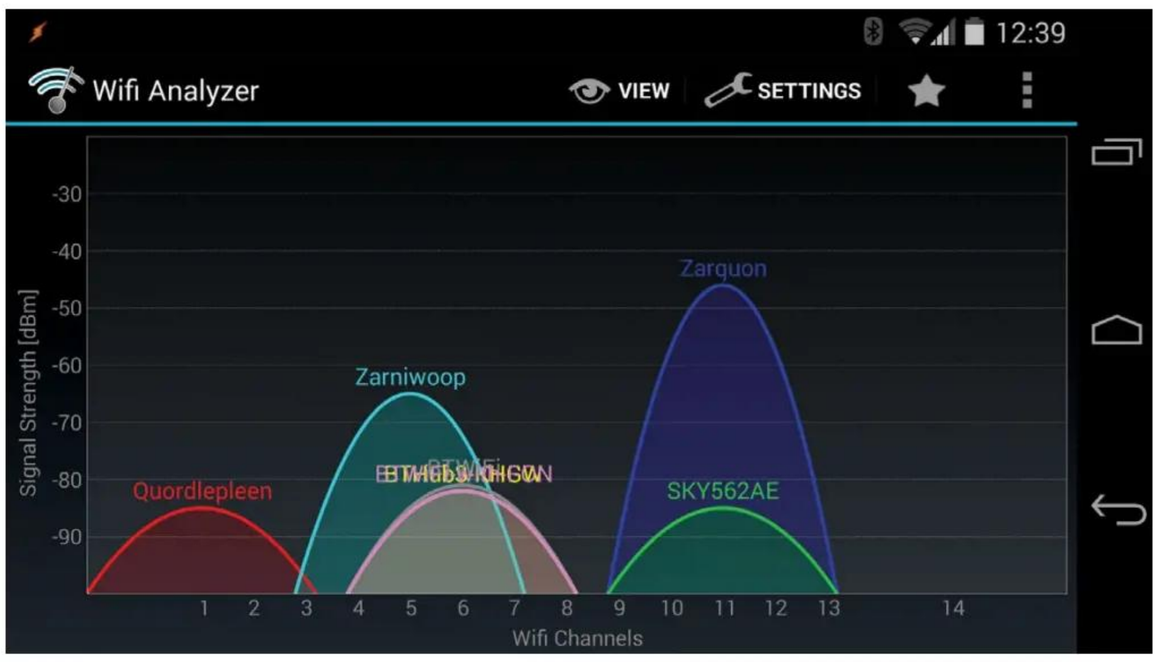

› ***Wi-Fi Analyzer* on Android shows access points in range and which channels they use, so you can set yours to avoid competition.**

Streaming video

Q When I try to watch videos on Amazon Prime or Netflix (I have accounts with both) I get a message saying my computer is unsupported. Can I fool them into thinking I am running Windows? Will that let me watch the services I have paid for?

A These services use *Silverlight*, a proprietary technology from Microsoft, to render the video and apply DRM. There's no *Silverlight* for Linux but there's a solution called *Pipelight*, which is a browser plugin that in turn uses *Wine* to run *Silverlight*. It requires a patched *Wine* but all the packages you need are available for most distros. Users of Ubuntu-based distros can add a PPA and install everything with:

```
sudo add-apt-repository ppa:pipelight/stable
sudo apt-get update
sudo apt-get install pipelight-multi
```

This adds the *Pipelight* PPA and then installs the software. Instructions for other distributions are at **http://bit.ly/PipelightInstall**. Set it up with:

```
sudo pipelight-plugin --create-mozilla-plugins
pipelight-plugin --enable silverlight
```

This will download the required *Silverlight* software and set up plugins for your browser. Installing *Pipelight* isn't enough to watch your videos; you will still be denied access because your browser tells the website you are running Linux. Install one of the User-Agent control extensions for your browser: such as **user-agent-overrider** or **uacontrol** on *Firefox* or **user-agent-switcher-for-chrome** if you use *Chrome*. Then use it to set your User Agent to "Firefox for Windows" and enjoy your streaming video.

The last but one resort

When your computer seems totally locked up and no amount of key bashing has any effect whatsoever, the last resort is to cut the power, with all the risks of data corruption this entails. But there's one more thing you can try: the Magic SysReq keys. These are special key combinations that can be used to shut down a locked-up system reasonably gracefully. Their name comes from the fact that they use the SysReq key, sometimes marked PrintScr, and they are truly magic.

Even if your keyboard does nothing else, these work because the kernel sees these key presses directly with no userspace software getting in the way.

While holding down Alt and SysReq, press each of these keys in turn:

» **R** – resets keyboard.
» **E** – exits all remaining running programs with a KILL signal.
» **I** – interrupts all running programs with a TERM signal.
» **S** – syncs all filesystems.
» **U** – unmounts all filesystems and remounts them read-only.
» **B** – reboots.

There are various mnemonics to help remember this key sequence, some of them harder to remember than the keypress sequence itself – there's one to do with elephants. The one I find easiest to remember is that it is BUSIER backwards, but the one I like most is 'Reboot Even If System Utterly Broken'.

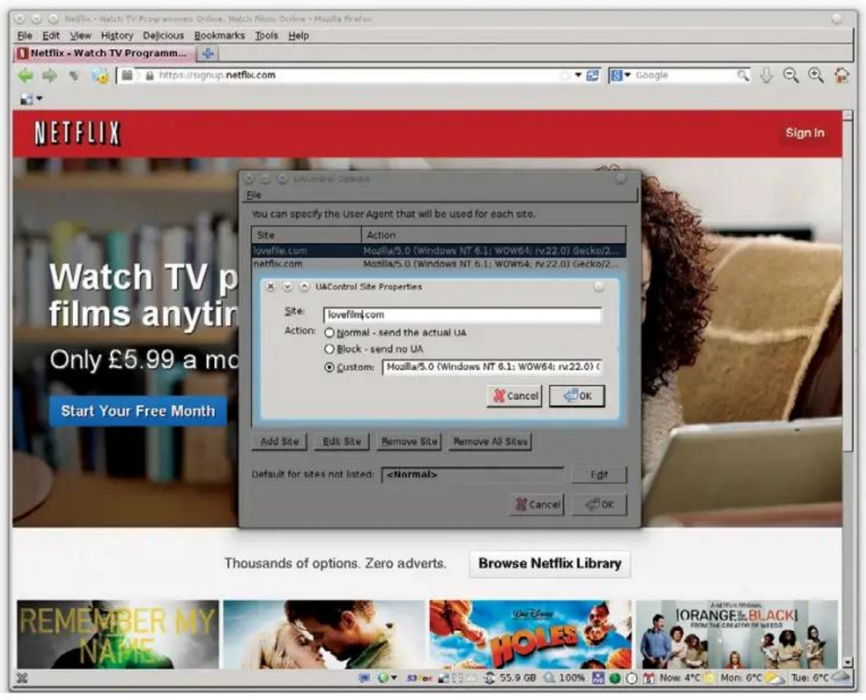

› **Get Netflix running by installing *Pipelight* and telling your browser to pretend that it is running on Windows.**

Mounting Android

Q I have bought myself a Nexus 7 Android tablet and I'm very pleased with it. But – there's always a but, isn't there? – I can't mount it as a USB device. My Android phone pops up a dialog when I plug it into the USB port of my computer and lets me mount it as a storage device, but this tablet doesn't seem to work like that. Why have they got rid of a perfectly reasonable system and how do I get my files on there now?

A The trouble with the USB storage system is that a filesystem can be mounted by only one OS at a time. That means that before you can mount it on your computer as a mass storage device, it has to be unmounted on the phone/tablet. That's why you get the warning message on your phone about apps that use the SD card having to be stopped. It's even worse on something like the Nexus 7, where there's no removable storage – you can't unmount a filesystem currently in use by the OS. Newer Android devices avoid this by using MTP (Media Transfer Protocol). There are a few options for mounting MTP devices under Linux, all implemented as FUSE filesystems. I prefer **jmtpfs**. To mount the first device found – you would normally only have one connected at a time – pass *jmtpfs* the name of the mount point. This directory must exist and be writable by the user running the command:

```
jmtpfs ~/nexus
```

Then you can access the contents of the device while still using it. When you have finished, unmount the device with:

```
fusermount -u ~/nexus
```

Remote desktops

Q I want to be able to access the desktop of one computer from another, like the remote desktop feature of Windows.

A You have two options here. For a complete desktop you can use *VNC*. This is often installed by default; if not, install either the *TightVNC* or *VNC* package (the former should give better speed on slower networks). You run the server on the computer you want to view; the first time you will be asked to set up a password. Then run the client on the other computer, giving it the address of the first computer.

However, if you only want to run a single program from the remote computer on your local desktop, *VNC* can be overkill and the X display of Linux already allows network forwarding. You need to set up SSH on the remote computer – just install the *ssh-server* package in the usual way. Then test that you can connect from the local computer with:

```
ssh user@hostname
```

This will ask for your password and then give you a shell session on the remote computer – you can check you are on the remote computer by running **uname -a**. Press Ctrl-D to log out, then try running a graphical program like this:

```
ssh -Y user@hostname someprogram
```

It should open its window in front of you. Depending on how your distro set up its SSH server, you may find **-X** is better than **-Y**. If neither works, check **/etc/ssh/sshd_config** on the computer running the server and make sure it contains

```
X11Forwarding yes
```

Incidentally, if you are using the same username on both computers, you can leave that out and just do

```
ssh -Y hostname progname
```

What's my address?

Q I'm running a Raspberry Pi without keyboard or screen. I plugged it into my network but I don't know what its IP address is, so can't SSH into it.

A If your device is getting its IP address from a DHCP server, you can often configure the server to give a specific address to a hostname. Plug the Pi's SD card into your computer and edit **/etc/hostname** on it to set your hostname, then set your DHCP server, usually in your router, to give that hostname a set address. Alternatively, install *netdiscover* (**http://bit.ly/Netdiscover**) and run it like this:

```
sudo netdiscover -i eth0 -r 192.168.1.0/24
```

where **-i** is your network interface and **-r** the IP range used by your network. This will list all the hosts on your network. Unfortunately the Pi shows up only as 'Unknown vendor', so run it with and without the Pi to see the address that appears. If you have an Android device, install *Fing*, which does the same thing graphically and recognises the Pi. Once you have the address, you can SSH in and run **ipconfig** to get the MAC address, which you can use to tell your DHCP server which address to give it. Raspbian prints the IP address from **/etc/rc.local** when it boots. You could modify this to email you instead. »

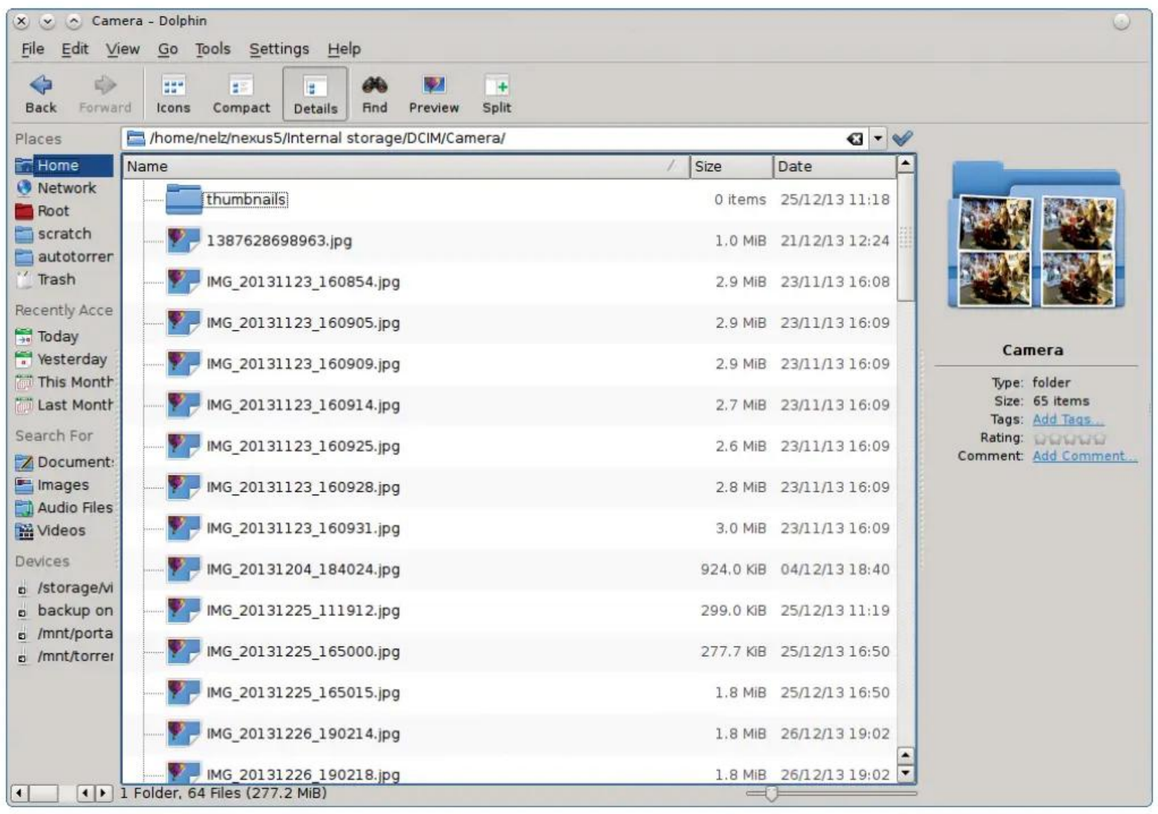

› **MTP is the new way of connecting to Android devices, and it means you can keep using the device while it is mounted on your desktop.**

Terminal: How to get started

Flex those fingers and dive head first into the inky darkness of the terminal. Let us show you how to start handling its mysterious commands.

Quick tip

If you're struggling to type the right command, you can wipe all previous actions from view simply by typing **clear** and hitting Enter. Note this won't affect your command history.

The terminal is an incredibly important part of your Linux desktop. It doesn't matter how wedded you are to point and click over the command line, at some point you're going to have to dip your toe in the terminal's dark expanse and use it. Don't worry, though, because the terminal isn't as scary as it might appear, and if you take the time to learn the basics you'll discover it can be a far quicker and more effective way of getting certain tasks done.

As you'd expect, a terminal effectively gives you access to your Linux shell, which means it works in exactly the same way using the same language (*Bash*). This means you can do anything in the terminal you'd normally do at the command line, all without leaving the relative comfort of your desktop. That makes learning how to use the terminal – and *Bash* – doubly advantageous as it gives you your first glimpse into working with the underlying Linux shell.

We're basing this tutorial on Ubuntu, so start by opening the Dash and typing 'terminal' into the search box. You'll find the terminal of course, but you'll also see two entries called *UXTerm* and *XTerm* too. This highlights the fact there are multiple terminal emulators that you can run in order to interact with the shell. There are differences between them, of course, but fundamentally they do the same thing.

For the purposes of this tutorial we're sticking with the default terminal, which is basically the *gnome-terminal* emulator – technically it's emulating a TeleTYpe (TTY) session. It has all the functionality you'll need, but both *XTerm* and *UXTerm* are worth noting because although they are more minimalist tools and neither require any dependencies to run. This means if anything stops the main terminal from running, you can use either as a backup. As an aside, the only difference between the two is that *UXTerm* supports the expanded Unicode character set.

› The --help flag can be used with any command to find out what it does, plus what arguments to use.

How Bash works

The Linux shell uses the *Bash* shell and command language to perform tasks, and it uses a relatively straightforward syntax for each command: `utility command -option`.

First, the 'utility' portion of the command is the tool you wish to run, such as `ls` for listing the contents of a directory, or `apt-get` to trigger the *APT* package management tool. Secondly, the `command` section is where you specify exactly what you want the utility to do, for example, typing `apt-get install` instructs the package management utility to install the named package, eg: `apt-get install vlc`.

And finally, the `-option` section of the command is where one or more 'flags' can be set to specify certain preferences. Each flag is preceded by one or two dashes (--) and the most useful of all is the `--help` option, which provides a brief

Speed up text entry

It doesn't matter how fleet of hand your typing skills are, the command line can still be a time-consuming, frustrating experience. Thankfully the terminal comes equipped with lots of handy time-saving shortcuts. This issue let's take a look at how you can easily access previously used commands and view suggestions:

» **Up/down arrows** Browse your command history.

» `history` Use this to view your command history

» **Ctrl+r** Search command history. Type letters to narrow down search, with the most recent match displayed, and keep pressing Ctrl+r to view other matches.

» **Tab** View suggestions or auto-complete a word or path if only one suggestion exists. Press **~+Tab** to autofill your username, **@+Tab** to autofill your host name and **$+Tab** to autofill a variable.

Your first terminal commands

While it's possible to install and manage software using a combination of the *Software Center* and Ubuntu's *Software & Updates* setting panel, it's often quicker to make use of the *Advanced Package Tool* (*APT*) family of tools. Here's some key ways that they can be used (see *sudo* use below):

» `$ apt-cache pkgnames` Lists all available packages from sources listed in the **/etc/apt/sources.list** file.

» `$ sudo add-apt-repository ppa:<repository name>` Adds a specific Launchpad PPA repository to the sources list.

» `$ sudo apt-get update` Gets the latest package lists (including updated versions) from all listed repositories.

» `$ sudo apt-get install <package>` Installs all the named package. This will also download and install any required dependencies for the packages.

» `$ apt-get remove <package>` Use this to remove an installed package. Use `apt-get purge <package>` to also remove all its configuration files, and `apt-get autoremove` to remove packages installed by other packages that are no longer needed.

» `$ sudo apt-get upgrade` Upgrades all installed software – run `sudo apt-get update` before running this. Other useful `apt-get` commands include `apt-get check` a diagnostic tool that checks for broken dependencies, `apt-get autoclean`, which removes Deb files from removed packages.

```
nick@nick-ubuntu: ~
nick@nick-ubuntu:~$ apt-cache show vlc
Package: vlc
Priority: optional
Section: universe/graphics
Installed-Size: 3765
Maintainer: Ubuntu Developers <ubuntu-devel-discuss@lists.ubuntu.com>
Original-Maintainer: Debian Multimedia Maintainers <pkg-multimedia-maintainers@lists.alioth.debian.org
>
Architecture: amd64
Version: 2.1.6-0ubuntu14.04.1
Replaces: vlc-data (<< 1.1.5), vlc-nox (<< 2.0.2)
Provides: mp3-decoder
Depends: fonts-freefont-ttf, vlc-nox (= 2.1.6-0ubuntu14.04.1), libaa1 (>= 1.4p5), libc6 (>= 2.15), lib
caca0 (>= 0.99.beta17-1), libfreetype6 (>= 2.2.1), libfribidi0 (>= 0.19.2), libgcc1 (>= 1:4.1.1), libg
l1-mesa-glx | libgl1, libqtcore4 (>= 4:4.8.0), libqtgui4 (>= 4:4.8.0), libsdl-image1.2 (>= 1.2.10), li
bsdl1.2debian (>= 1.2.11), libstdc++6 (>= 4.6), libtar0, libva-x11-1 (>> 1.3.0~), libva1 (>> 1.3.0~),
libvlccore7 (>= 2.1.0), libx11-6, libxcb-composite0, libxcb-keysyms1 (>= 0.3.9), libxcb-randr0 (>= 1.1
), libxcb-shm0, libxcb-xv0 (>= 1.2), libxcb1 (>= 1.6), libxext6, libxinerama1, libxpm4, zlib1g (>= 1:1
.2.3.3)
Pre-Depends: dpkg (>= 1.15.6~)
Recommends: vlc-plugin-notify (= 2.1.6-0ubuntu14.04.1), vlc-plugin-pulse (= 2.1.6-0ubuntu14.04.1), xdg
-utils
Suggests: videolan-doc
Breaks: vlc-data (<< 1.1.5), vlc-nox (<< 2.0.2)
Filename: pool/universe/v/vlc/vlc_2.1.6-0ubuntu14.04.1_amd64.deb
Size: 1212144
MD5sum: fbb2933ada01d9ccddd319ddea21bd89
SHA1: 4bb0e71315956d97ce18b67d7eeb27ee1b968fbe
SHA256: 636992ae393297dd5afdba39b9cb3a0958b3d1f246e51a47660cc3f5993ca35f
Description-en: multimedia player and streamer
```

› **The apt-cache package can also be used to search for specific packages or reveal a package's dependencies.**

description of the utility, plus lists all available commands and options, eg `ls -l`.

The `-l` flag tells the list directory tool to provide detailed information about the contents of the folder it's listing, including: permissions; who owns the file; the date it was last modified; and its size in bytes. Utilities can be run without any commands or options – eg `ls` on its own provides a basic list of all folders and files in a directory. You can also run utilities with a combination of commands and/or options.

Restricted access

Open the terminal and you'll see something like this appear: `username@pc-name:~$`. This indicates that you're logged on to the shell as your own user account. This means that you have access to a limited number of commands – you can run `ls` directly, eg, but not to install a package using `apt-get`, because the command in question requires root access. This is achieved one of two ways – if you're an administrative user, as the default user in Ubuntu is, then you can precede your command with the `sudo` command, eg `sudo apt-get install vlc`. You'll be prompted for your account password, and then the command will run. You should find that you can run more `sudo`-based commands without being re-prompted for your password (for five minutes) while the terminal is open. On some distros you can log on to the terminal as the root user with `su` – you'll be prompted for the root password at which point you'll see the following prompt: `root@pc-name:~$`.

Once logged in, you can enter commands with no restrictions. We recommend you use the `sudo` command rather than this approach and if you're running Ubuntu then you'll find `su` won't work because the root account password is locked for security reasons.

When installing some distros or adding new users to Ubuntu, you may find your user account isn't added to the `sudo` group by default. To resolve this, you need to open the terminal in an account that does have root access (or use `su` command if supported) and type `sudo adduser <username> sudo`. You can also add the user to other groups with the command by listing all the groups you wish to add, eg: `sudo adduser <username> adm sudo lpadmin sambashare`.

Another handy tool is *gksudo*, which allows you to launch desktop applications with root privileges. It's of most use when wanting to use the file manager to browse your system with root access: `gksudo nautilus`. Make sure you leave the terminal open while the application is running, otherwise it'll close when the terminal does. When you're done, close the application window, then press Ctrl+c in the terminal, which interrupts the currently running program and returns you to the command line.

We've already discussed the `--help` flag, but there are other help-related tools you can use too. First, there's `whatis` – which you can type with any command to get a brief description of it and any specified elements, eg `whatis apt-get install vlc` will describe the *apt-get* tool, the install argument and what package **vlc** is. Flags are ignored.

If you're looking for a full-blown manual, then the *man* tool provides access to your distro's online reference manual, which is started with `man intro`. This provides you with a long and detailed intro to the command line. Once done press q to quit back to the terminal. For more advice on navigating the manual, type `man man` or pair it with a tool, eg `man ls`.

Now you've taken your first steps into the world of the terminal, check out the box (*Your First Terminal Commands, above*) for some useful package management commands you can work with. Next, we'll look at how to navigate your filesystem from the terminal, plus launch programs and delve into more useful shortcuts to help speed up the way you interact with the command line. ■

Terminal: Work with files

Time to turn our attention to navigating your file system and manipulating files and folders from the beloved Terminal.

Previously we introduced you to the basics of using the Terminal. We opened by revealing it works in the same way as your Linux shell; how commands are structured (utility command -option); plus gave you the tools to manage software packages and get further help. This time, we're going to look at how you can navigate your file system, work with files and folders and learn some more time-saving shortcuts in the bargain.

When you open a new Terminal window, the command prompt automatically places you in your own personal **home** folder. You can verify this using the `ls` command, which lists the contents of the current folder. The default Terminal application displays folder names in blue, and filenames in white, helping you differentiate between them. The `ls` command can be used in other ways too. Start by typing `ls -a` to display all files, including those that begin with a period mark (`.`), which are normally hidden from view. Then try `ls --recursive` , the `--recursive` option basically means that the contents of sub-folders are also displayed.

If you want more detail about the folder's contents – permissions settings, user and group owners, plus file size (in bytes) and date last modified, use `ls -l` . If you'd prefer to list file sizes in kilobytes, megabytes or even gigabytes depending on their size, add the `-h` option—so use `lh -h -l` instead. There are many more options for `ls` and you can use the `--help` option to list them all.

Navigating your file system is done using the `cd` command – to move down one level to a sub-folder that's inside the current directory use `cd <subfolder>` , replacing <subfolder> with the name of the folder you wish to access. Remember that folder and filenames are case sensitive, so if the folder begins with a capital letter – as your personal **Documents** folder does, eg – you'll get an error about the folder not existing if you type it all in lower case, eg, `cd documents` . You can also move down several levels at once using the following syntax: `cd subfolder/subfolder2` . To move back up to the previous level, use `cd ..` , you can also use the `/` character to move up multiple levels at once, eg `cd ../..` moves up two levels.

› Make good use of ' and \ characters when folder paths contain spaces and other special characters.

What if you want to go somewhere completely different? Use `cd /` to place yourself in the root directory, or navigate anywhere on your system by entering the exact path, including that preceding `/` character to indicate you're navigating from the top level, eg `cd /media/username` .

The `~` character works in a similar way to `/` , except this places you in your home directory. So typing `cd ~/Documents` is the same as typing `cd /home/username/`

Speedier navigation

In last part we revealed some handy keyboard shortcuts to help you enter commands more quickly, but the following keys will help you navigate the Terminal itself more efficiently:

» **Home/End** Press these to jump to the beginning or end of the current line.

» **Ctrl+left/right cursor** Move quickly between arguments.

» **Ctrl+u** Clear the entire line to start again.

» **Ctrl+k** Delete everything from the cursor's position onwards.

» **Ctrl+w** Delete the word before the cursor.

Accidentally omitted `sudo` from your command? Just type `sudo !!` and hit Enter to repeat the last command with `sudo` applied to it. And if you make a typo when entering a command, instead of retyping the entire command again, just use the following syntax to correct the mistyped word (in the following example, `dpkg` was originally mistyped as dkpg):

```
^dkpg^dpkg
```

Boost your learning

Now you're starting to flex your muscles in Terminal, how about expanding your knowledge by instructing it to display information about a random command each time you open it? To do this, you need to edit a file, so open the Terminal and type the following:

```
nano ~/.bashrc
```

This opens the file in the *nano* text editor. Use the cursor keys to scroll down to the bottom of the file, then add the following line to it:

```
echo "Did you know that:"; whatis $(ls /bin | shuf -n 1)
```

Press [Ctrl]+[o] to save the file (just hit Enter to overwrite it), then [Ctrl]+[x] to exit *nano*. Now close the Terminal window and open a new one to get a brief description of a command. Just type the following, with the actual command listed for a longer description: `<command> --help`.

`Documents`. One final trick —you've jumped to another directory, but how do you go back to the previous directory quickly? Simple, just type `cd -` to do so.

Working with files and folders

You can now list directories and navigate your file system, but what about doing something practical, like moving and copying files? You'll find a range of different commands exist, and the tricks you've learned about navigation will hold you in good stead here too.

Let's start by looking at commands for copying (`cp`) and moving (`mv`) files and folders. The same options apply to both commands. The basic syntax is cp/mv <source> <target>. The source and target can be complete paths following the same rules for the `cd` command, but it's generally good practice to first navigate to the folder containing the file or folder you wish to copy or move. Once done, you can simply specify the file or folder name as the source, like so `cp invoice.odt ~/Documents/Backup`.

This creates a copy of the file with the same name. The following copies the file to the specified directory and renames it too: `cp invoice.odt ~/Documents/Backup/invoice-backup.odt`. If you want to create a copy of the file within the same file, simply use `cp invoice.odt invoice-backup.odt`.

Substitute `mv` for `cp` in any of the above commands, and the file is moved, moved and renamed or simply renamed. What happens if there's already a file called **invoice-backup.odt** in existence? It'll be overwritten without as much as a by your leave, so make sure you're asked if you want to overwrite it by adding the `-i` flag like this `mv -i invoice.odt invoice-backup.odt`.

You can also copy folders using the `cp` or `mv` commands. Here, you need to include the recursive option, which ensures the folder is copied across with all its contents and correctly arranged in their original locations relative to the parent folder: `cp -r ~/Documents /mnt/sdb1/Backup/`.

If the **Backup** folder exists, then the **Documents** folder will be recreated inside it; if not, then the **Backup** folder is created and the contents of the **Documents** folder are copied into it instead.

Use the `rm` command to delete a single file, eg `rm invoice.odt`. The `rmdir` command deletes folders, but only empty ones. If you want to delete a folder and all its contents, use the command `rm -r foldername`.

You can also create new folders with `mkdir` command – simply type `mkdir` folder, replacing folder with your chosen folder name. Use the `touch` command to create an empty file, such as `touch config.sys`.

Wildcards are often used to speed things up in searches, and can also be applied to file commands too – the asterisk (`*`) character can be used to quickly access a folder with a long name, eg `cd Doc*`.

This works fine if there's only one folder beginning with Doc, but if there are two (say Doctor and Documents), then the command would open the first matching folder, which is Doctor in this instance. To avoid this, use `cd Doc*ts` instead (unless you have a folder called Documents and Doctorists).

Two characters that are more useful when navigating are the single quotation mark (`'`) and backslash (`\`) characters. Use single quotation marks around files or file paths that contain spaces, such as `cd ~\Documents\Doctor Who`.

You should also use quotation marks when creating folders in this way, eg simply typing `mkdir Doctor Who` will actually create two separate folders called **Doctor** and **Who**, so type `mkdir 'Doctor Who'` to get the folder you want.

You can also use the `\` character to get around this too, eg `mkdir Doctor\ Who` works in the same way, because the `\` character instructs `mkdir` to treat the following character (in this instance the space) as 'special'.

We finish off by revealing some handy characters that allow you to run multiple commands on a single line. The `&&` argument does just that, so you can do the following to quickly update your repos and update all available software:

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

`&&` is like the AND command in that the second command will only be performed if the first completes successfully. If you wanted the second command to only run if the first command failed then you'd use `||` instead. If you want the second command to run after the first regardless of what happens, then use the `;` eg,

```
sudo apt-get update ; sudo apt-get remove appname
```

instead: of `&&`. ■

Quick tip

Some file managers allow you to right-click a folder and open the Terminal at that location, but you have to manually add this option to Ubuntu's *Nautilus* file manager. Install **nautilus-open-terminal** from the Software Center, then open a Terminal window, type **nautilus -q** and press Enter. The option will now appear.

```
media                                   VirtualBox VMs
Music                                   wget-log
nick@nick-ubuntu:~$ ls -h -l
total 2.5M
drwxrwxr-x 2 nick nick 4.0K Feb 15 15:08 deja-dup
drwxr-xr-x 2 nick nick 4.0K Feb 28 17:21 Desktop
drwxrwxr-x 2 nick nick 4.0K Feb 26 17:35 Doctor
drwxr-xr-x 3 nick nick 4.0K Feb 26 17:44 Documents
drwxr-xr-x 4 nick nick 4.0K Feb 27 13:23 Downloads
-rw-r--r-- 1 nick nick 8.8K Nov 26 12:51 examples.desktop
drwxr-xr-x 2 root root 4.0K Feb  4 19:49 fedora
-rw-rw-r-- 1 nick nick 446K Jan 20 12:35 linuxdesktops-enlightenment.png
drwxr-xr-x 3 root root 4.0K Feb  4 19:50 media
drwxr-xr-x 2 nick nick 4.0K Nov 26 13:10 Music
drwx------ 2 nick nick 4.0K Feb 10 16:43 NoMachine
drwxr-xr-x 2 nick nick 4.0K Nov 26 13:10 Pictures
lrwxrwxrwx 1 nick nick   36 Dec 10 13:35 PlayOnLinux's virtual drives -> /home/n
ick/.PlayOnLinux//wineprefix/
drwxr-xr-x 2 nick nick 4.0K Nov 26 13:10 Public
-rw-rw-r-- 1 nick nick 871K Jan 13 10:22 steamos1.png
drwxr-xr-x 2 nick nick 4.0K Nov 26 13:10 Templates
-rw-rw-r-- 1 nick nick 1.1M Dec 24 22:09 tigervncserver_1.6.0-3ubuntu1_amd64.deb
drwxr-xr-x 2 nick nick 4.0K Nov 26 13:10 Videos
drwxrwxr-x 6 nick nick 4.0K Nov 30 09:43 VirtualBox VMs
-rw-rw-r-- 1 nick nick 3.5K Jan 21 11:39 wget-log
nick@nick-ubuntu:~$
```

› **Use ls to find out more about the files and folders in a current directory.**

Terminal: Apt-get in action

New to Linux? Then allow us to guide you through your first steps with **apt-get**, the powerful command line tool.

One of the biggest changes that catches Windows users moving to Linux is the way that software is installed. Instead of downloading an executable file from some website or other, running it and hoping it doesn't clobber your existing library files (DLLs) or install some dubious adware or malware, Linux distributions maintain repositories of software, which are all packaged up for that distro and tested for compatibility with the rest of the distro.

In this tutorial, we will look at how this is done by distros that use the *Advanced Packaging Tool* (apt) software management system, as developed by Debian and used by distros from Ubuntu to Raspbian on the Raspberry Pi.

Repositories

A repository is a collection of software packages for a distro. Each major distro release will have its own repositories, and the packages will have been built for and tested with that release, but it is more than a collection of files. Each repo (as they are usually called) is indexed, making it easy to find what you want. It can also be checked for updates for your package manager without any need to visit websites to check for updates, or the need for software to 'phone home' to check.

More importantly, each package in a repo is signed with the repository's GPG (encryption) key, which is checked when installing packages. This means you can trust the software installed from there to be what it says it is, and not some infected trojan that's been uploaded maliciously.

A repository also makes dependency handling simple. A dependency is a program that the program you want to install needs to run, such as a library. Instead of bundling everything in the package and ending up with multiple copies of the same library on your computer (which is what Windows does), a package simply lists its dependencies so that your package manager can check whether they are already installed, and grab them from the repo if not.

In addition to the default repositories provided by the distro, there are several third-party ones that can be added to your package manager. These are not guaranteed to be tested to the same standards as the official repos, but many of them are very good, and if you stick to the popularly recommended repos for your distro, you won't go far wrong.

Ubuntu has also introduced the concept of the PPA, or Personal Package Archive, which are small repositories for individual projects. These may each be added individually to your package manager, but be careful about adding any untrusted sources.

Package management

We have used the term 'package manager' a few times now but what is it? Basically, this is a program that enables you to

Package management

```
Terminal
File Edit View Search Terminal Help
nelz@mintvm ~ > sudo apt-get install --dry-run thunderbird
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following extra packages will be installed:
  thunderbird-gnome-support thunderbird-locale-en
Suggested packages:
  ttf-lyx
The following packages will be upgraded:
  thunderbird thunderbird-gnome-support thunderbird-locale-en
3 upgraded, 0 newly installed, 0 to remove and 88 not upgraded.
Inst thunderbird-locale-en [1:24.2.0+build1-0ubuntu0.13.10.1] (1:24.3.0+build2-0ubuntu0
.13.10.1 Ubuntu:13.10/saucy-updates [i386]) []
Inst thunderbird [1:24.2.0+build1-0ubuntu0.13.10.1] (1:24.3.0+build2-0ubuntu0.13.10.1 U
buntu:13.10/saucy-updates [i386]) [thunderbird-gnome-support:i386 ]
Inst thunderbird-gnome-support [1:24.2.0+build1-0ubuntu0.13.10.1] (1:24.3.0+build2-0ubu
ntu0.13.10.1 Ubuntu:13.10/saucy-updates [i386])
Conf thunderbird (1:24.3.0+build2-0ubuntu0.13.10.1 Ubuntu:13.10/saucy-updates [i386])
Conf thunderbird-locale-en (1:24.3.0+build2-0ubuntu0.13.10.1 Ubuntu:13.10/saucy-updates
 [i386])
Conf thunderbird-gnome-support (1:24.3.0+build2-0ubuntu0.13.10.1 Ubuntu:13.10/saucy-upd
ates [i386])
nelz@mintvm ~ >
```

1 Install

Using **apt-get install** will check the dependencies of the packages you want and install any that are needed. Adding **--dry-run** to **apt-get install** enables you to see what would be done, without actually writing anything to your hard drive. If you are happy, run the command again without **--dry-run**.

```
Terminal
File Edit View Search Terminal Help
thunderbird - Email, RSS and newsgroup client with integrated spam filter
thunderbird-dbg - Email, RSS and newsgroup client - debug symbols
thunderbird-dev - Email, RSS and newsgroup client - development files
thunderbird-globalmenu - Email, RSS and newsgroup client (transitional package)
thunderbird-gnome-support - Email, RSS and newsgroup client - GNOME support
thunderbird-gnome-support-dbg - Email, RSS and newsgroup client - transitional package
thunderbird-locale-af - Transitional package for unavailable language
thunderbird-locale-ar - Arabic language pack for Thunderbird
thunderbird-locale-ast - Asturian language pack for Thunderbird
thunderbird-locale-be - Belarusian language pack for Thunderbird
thunderbird-locale-bg - Bulgarian language pack for Thunderbird
thunderbird-locale-bn - Bengali language pack for Thunderbird
thunderbird-locale-bn-bd - Transitional Bengali language pack for Thunderbird
thunderbird-locale-br - Breton language pack for Thunderbird
thunderbird-locale-ca - Catalan; Valencian language pack for Thunderbird
thunderbird-locale-cs - Czech language pack for Thunderbird
thunderbird-locale-da - Danish language pack for Thunderbird
thunderbird-locale-de - German language pack for Thunderbird
thunderbird-locale-el - Greek language pack for Thunderbird
thunderbird-locale-en - English language pack for Thunderbird
thunderbird-locale-en-gb - Transitional English language pack for Thunderbird
thunderbird-locale-en-us - Transitional English language pack for Thunderbird
thunderbird-locale-es - Spanish; Castilian language pack for Thunderbird
thunderbird-locale-es-ar - Transitional Spanish language pack for Thunderbird
thunderbird-locale-es-es - Transitional Spanish language pack for Thunderbird
:
```

2 Search

Use **apt-cache search** to find what's available. The **--names-only** option can give a more manageable set of results if you know the program's name. Otherwise let **apt-cache search** go through the descriptions, too, and view the results in **less**. You don't need to use **sudo** as **search** doesn't write to your drive.

```
File Edit View Search Terminal Help
  hplip-data libavcodec53 libavformat53 libavutil51 libbind9-90 libcups2
  libcupscgi1 libcupsimage2 libcupsmime1 libcupsppdc1 libcurl3-gnutls
  libdns99 libespeak1 libgnutls-openssl27 libgnutls26 libhpmud0 libisc95
  libisccc90 libisccfg90 liblwres90 libmagic1 libnautilus-extension1a
  libnspr4 libnss3 libnss3-1d libpostproc52 libpurple0 libpython2.7
  libpython2.7-minimal libpython2.7-stdlib libpython3.3
  libpython3.3-minimal libpython3.3-stdlib libsane-hpaio
  libsoup-gnome2.4-1 libsoup2.4-1 libssl1.0.0 libswscale2 libxfont1
  linux-firmware linux-libc-dev mint-mirrors mintbackup nautilus-data
  openjdk-7-jre openjdk-7-jre-headless openjdk-7-jre-lib openssl pidgin
  pidgin-data printer-driver-hpcups printer-driver-postscript-hp
  python-aptdaemon python-aptdaemon.gtk3widgets python-cupshelpers
  python2.7 python2.7-minimal python3-aptdaemon
  python3-aptdaemon.gtk3widgets python3-aptdaemon.pkcompat python3.3
  python3.3-minimal simple-scan system-config-printer-common
  system-config-printer-gnome system-config-printer-udev thunderbird
  thunderbird-gnome-support thunderbird-locale-en thunderbird-locale-en-us
  xserver-common xserver-xephyr xserver-xorg-core xserver-xorg-video-intel
91 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
Need to get 171 MB of archives.
After this operation, 3,569 kB of additional disk space will be used.
Do you want to continue [Y/n]?
```

3 Update

Run **apt-get update** to update all your package lists, followed by **apt-get upgrade** to update all your installed software to the latest versions. In our case, it's well overdue. Then **apt** will show you what needs to be updated, and how much needs to be downloaded, before asking whether you want to proceed.

```
Terminal
File  Edit  View  Search  Terminal  Help

                    SUMMARY OF LESS COMMANDS

      Commands marked with * may be preceded by a number, N.
      Notes in parentheses indicate the behavior if N is given.
      A key preceded by a caret indicates the Ctrl key; thus ^K is ctrl-K.

  h  H                 Display this help.
  q  :q  Q  :Q  ZZ     Exit.
 ---------------------------------------------------------------------------

                           MOVING

  e  ^E  j  ^N  CR  *  Forward  one line   (or N lines).
  y  ^Y  k  ^K  ^P  *  Backward one line   (or N lines).
  f  ^F  ^V  SPACE  *  Forward  one window (or N lines).
  b  ^B  ESC-v      *  Backward one window (or N lines).
  z                 *  Forward  one window (and set window to N).
  w                 *  Backward one window (and set window to N).
  ESC-SPACE         *  Forward  one window, but don't stop at end-of-file.
  d  ^D             *  Forward  one half-window (and set half-window to N).
HELP -- Press RETURN for more, or q when done
```

› Less displays text from any source – from a file, the output of another program or its built-in help if you manage to get stuck.

install, update and remove software, including taking care of dependencies. It also enables you to search for programs of interest, as well as performing other functions.

All distros will have command line package management tools. You can access them either by using your system's search and looking for **terminal** or using Ctrl+Alt+T in desktops such as Unity, Gnome or Xfce, even if they also provide a fancy graphical front end. The main commands are as follows:

» **apt-get** Installs, upgrades and uninstalls packages.
» **apt-cache** This works with the repository index files, such as searching for packages.
» **add-apt-repository** Adds extra repositories to the system.
» **dpkg** A lower level package manipulation command.

These commands generally require root (superuser) access, so should be run at the root user or with **sudo** – we will stick with the **sudo** approach here. We've already mentioned that repos are indexed, so the first thing to do is update your index files to match the current contents of the repositories with:

```
sudo apt-get update
```

Then you probably want to make sure that your system is up to date:

```
sudo apt-get upgrade
```

This will list the packages it wants to install, tell you how much space it needs for the download, and then get on with it when you tell it to. When you want to install some new software, unless you have been told the exact name to install, you may want to search for it first, like this:

```
apt-cache search gimp
```

This will spit out a long list of packages, because it searches both name and description, and lists anything mentioning gimp, and there are a lot of them. To search only the names, use the **-n** or **--names-only** option:

```
apt-cache search -n gimp
```

This often gives a more manageable output, but still a lot in this case, perhaps too much to fit in your terminal window. The solution to this is to pipe the output from this command to the program **less**:

```
apt-cache search -n gimp | less
```

The **less** command is a pager – it lets you read text page by page and scroll through it. It can be used with any program that generates lots of terminal output to make it easier to read (see the 'Package management' walkthrough opposite for more details). Once you have found the package you want, installation is as simple as:

```
sudo apt-get install gimp
```

You can install multiple programs by giving them all to **apt-get** at once:

```
sudo apt-get install program1 program2...
```

Not every program you try will be what you want, so you can tidy up your hard drive by uninstalling it with:

```
sudo apt-get remove program1
```

Or you can use:

```
sudo apt-get purge program1
```

Both commands remove the program, but **remove** leaves its configuration files in place while **purge** deletes those, too.

There are a number of extra options you can use with **apt-get**. The **man** page lists them all (type **man apt-get** in the terminal), but one of the most useful is **--dry-run**. This has **apt-get** show you what it would do without actually doing it – a useful chance to check that you are giving the right command. Remember, computers do what you tell them to, not what you want them to do! Finally, you don't normally need to use **dpkg**, but it is useful for listing everything you have installed with **dpkg -L**. ■

Tar: Archiving your files

Archiving contains many mysteries, such as zips and tarballs, but we are on hand to expand your knowledge of their uses.

Most people reading this are probably pretty familiar with creating, sending or receiving ZIP files. *Zip* takes a collection of files and stores them in a ZIP archive file, compressing the data in the process. As well as storing the contents of the files, it also stores metadata, which is extra information associated with an object. In the case of files, it includes the modification times, their owners and permissions, and, of course, the name of each file. When you unzip an archive, all of this information is extracted, which recreates the original set of files exactly as they were.

Archives have several uses, but the most popular include bundling up a set of files for download – a single file is easier to handle and the compression makes it faster to download – and creating backups. As *Zip* has been around for ages, everyone can use it and all OSes can handle it. Nonetheless, *Zip* does have a number of drawbacks. The main issue is that its compression is poor by modern standards. Over the past 25 years, compression technology has moved on and even though there have been improvements to *Zip*, there are better alternatives available. Another disadvantage of *Zip* is that it is intended for archiving to a file, whereas sometimes we want to send the data to another device or service.

The standard archival program for Unix-like operating systems is *Tar*, so called because *Tar* was originally used to store backups on tape drives (Tape ARchive). It works in a different way to *Zip* because it sends all the archived data to its standard output and doesn't compress the data by default. This is because many tape drives already had hardware compression built in. The lack of compression code may seem like a disadvantage, but it's actually a convenience. As *Tar* is able to pipe its data via an external compression program, it can use any compressor it likes – even one that wasn't in existence when *Tar* was developed.

Compression programs work on one file or stream of data and produce one compressed file or stream, so this splits the job into two parts: archival and compression. While this may seem more complex, *Tar* is capable of handling the details itself. Let's say we have a directory that is called **foo**. We want to create an archive of it, which is often referred to as a tarball. We can use one of these options:

```
tar cf foo.tar foo
tar czf foo.tar.gz foo
tar cjf foo.tar.bz2 foo
tar cJf foo.tar.xz foo
```

The **c** option tells the *Tar* program that we are creating an archive, while **f** tells it that we are storing the archive in a file using the given name. Therefore, the first command creates an uncompressed archive that is called **foo.tar**. The subsequent commands add an extra option that tells *Tar* which particular type of compression to use: **z** uses gzip

Pipes and streams

A core part of the philosophy of Unix-like operating systems is for tools to be specialised – each tool should do one job and do it well. That means we need a way to chain programs together to achieve more complex tasks, and this is done with pipes. Every terminal program has streams for standard input, standard output and standard error (usually referred to as **stdin**, **stdout** and **stderr**). Standard input is how the program accepts information and it defaults to the keyboard when running a command from a terminal. Standard output is the output from the command that defaults to printing to the terminal window. Standard error goes to the same place, but all of these can be changed. Consider the example of:

```
tar c foo | sdc >foo.tar.sdc
```

The **|** character is the pipe and it connects the standard output of one program to the standard input of another. So *Tar* archives the contents of **foo** and sends it to **stdout**, which is then piped to the input of **sdc**. Most compressors default to working on a file:

```
gzip somefile
```

This compresses a file to **somefile.gz**. When no file is given, it compresses standard input and sends the result to standard output. The **>** character performs redirection and diverts standard output from the terminal to a file, so you get a compressed archive in **foo.tar.sdc**.

Standard error is not redirected by **>**. If something goes wrong, error messages are printed to the terminal – if they were sent to the file, you wouldn't know something was wrong.

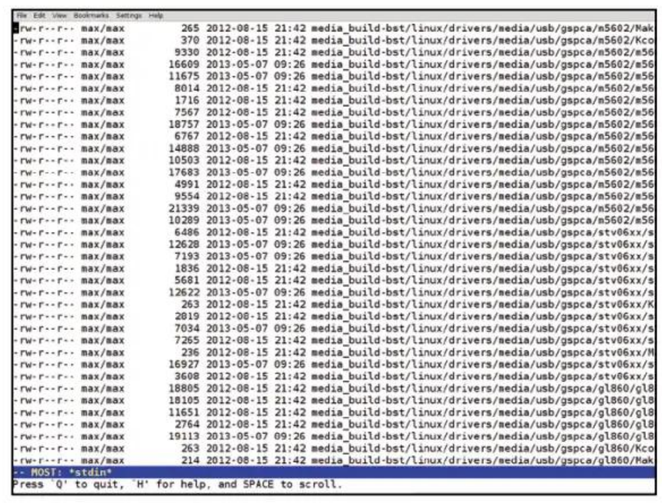

› Running *Tar* with the t (test) option not only checks that the archive is correct, but also shows you the full contents – usually enough for you to need to pipe it through a pager such as *Less*.

Compression types

There are two types of compression: lossless and lossy. Lossy compression achieves a greater reduction in file sizes by carefully discarding data that doesn't greatly affect the results. JPEG and MP3 files, for example, use lossy compression. Archiving uses lossless compression and uncompresses the result to get you back exactly where you started. There are a number of lossless compression methods available, here are some of the most common.

» **Deflate** This is the default compression used by *Zip*; it is old and not particularly effective, but is fast and well supported.

» **Compress** This is an old Unix compression program. Its files have a .Z extension, but you are unlikely to see many in the wild. There are also patent issues with it.

» **Gzip** This *Compress* replacement is completely open and still in use today. *Gzip* uses Deflate, but generally produces smaller archives than *Zip* with its default settings. It does not give the best compression, but it compresses and decompresses quickly, making it a good choice when size is not the most vital consideration.

» **Bzip2** This is a more efficient compressor, but at the expense of speed. Compression in particular can be slower, but it produces more compact results.

» **xz** This more recent program uses the LZMA2 compression algorithm also used by *7-Zip*. It is fast – particularly for decompression – and works brilliantly. It is the compressor of choice for many key Linux projects, including Coreutils and the kernel itself.

```
[nelz@hactar TerminalTutorials/Part4 0]% l test.tar.*
-rw-r--r-- 1 nelz users 6.6M Apr  4 12:17 test.tar.bz2
-rw-r--r-- 1 nelz users 8.2M Apr  4 12:17 test.tar.gz
-rw-r--r-- 1 nelz users 5.4M Apr  4 12:17 test.tar.xz
[nelz@hactar TerminalTutorials/Part4 0]%
[nelz@hactar TerminalTutorials/Part4 0]% gzip -l test.tar.gz
         compressed        uncompressed  ratio uncompressed_name
            8536353            33525760  74.5% test.tar
[nelz@hactar TerminalTutorials/Part4 0]%
[nelz@hactar TerminalTutorials/Part4 0]% xz -l test.tar.xz
Strms  Blocks   Compressed Uncompressed  Ratio  Check   Filename
    1       1  5,505.8 KiB     32.0 MiB  0.168  CRC64   test.tar.xz
[nelz@hactar TerminalTutorials/Part4 0]%
```

› **Here is the same archive using different compression programs. The *Gzip* version is significantly larger, but faster to compress and decompress.**

compression, **j** uses bzip2 compression and **J** uses xz compression. (Watch the capitalisation!)

There are also long versions of these arguments that make the commands more readable, but most of us are lazy and use the version that is shorter to type. However, we could also have used this command line if we wanted to:

```
tar --create --gzip --file foo.tar.gz foo
```

The file extension is not required, but it's a convention that makes it easier for people to see what type of archive it is – the system itself needs no such help as it can work all this out for itself. Unpacking an archive is simply a matter of replacing **c** with **x**, or **--create** with **--extract**. However, you don't need to give the compression type, as *Tar* figures it out:

```
tar xf foo.tar.gz
```

Another option you may want to add is **v** or **--verbose**, to show you what *Tar* is doing. If you have been given a tarball, you may want to see what is inside it without unpacking it. If you have created an archive, particularly a backup, you may want to check it's correct before relying on it. The test option checks the integrity and lists the contents of the archive.

```
tar tvf foo.tar.gz
```

Those are the main *Tar* options, but the program has many more, such as **A** or **--concatenate** to add files to an existing archive instead of creating a new one.

Future proofing

We mentioned that *Tar* can handle any new compression format that comes along because it passes compression to another program. There are command line options to do this automatically for gzip, bzip2 and xz, but what if someone comes up with a new compressor? Say something like **sdc** – super-duper compressor? You could create an uncompressed tarball and then use **sdc** to compress it, but that's wasteful and slow, so instead use a pipe:

```
tar c foo | sdc >foo.tar.sdc
unsdc foo.tar.sdc | tar xv
```

Here, we have used only the **--create** option with *Tar*. The lack of a destination causes *Tar* to send the archive data to standard output, which is then piped to the **sdc** compressor program. The second command reverses the whole process, decompressing the archive before sending it to *Tar* for extraction. ■

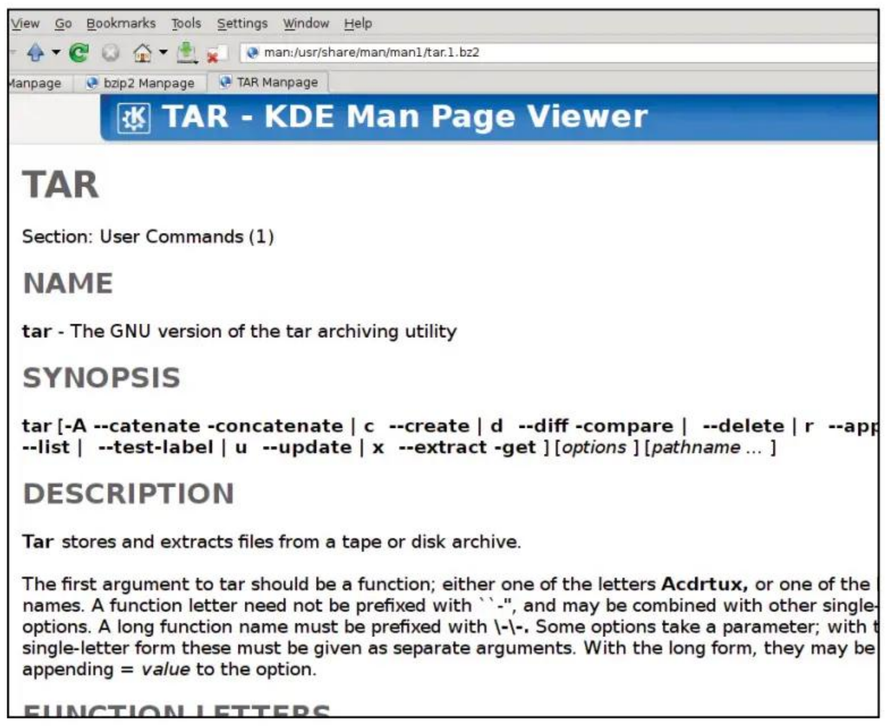

› ***Tar* and the compressors have man pages, with many options. You'll usually only need the ones covered here.**

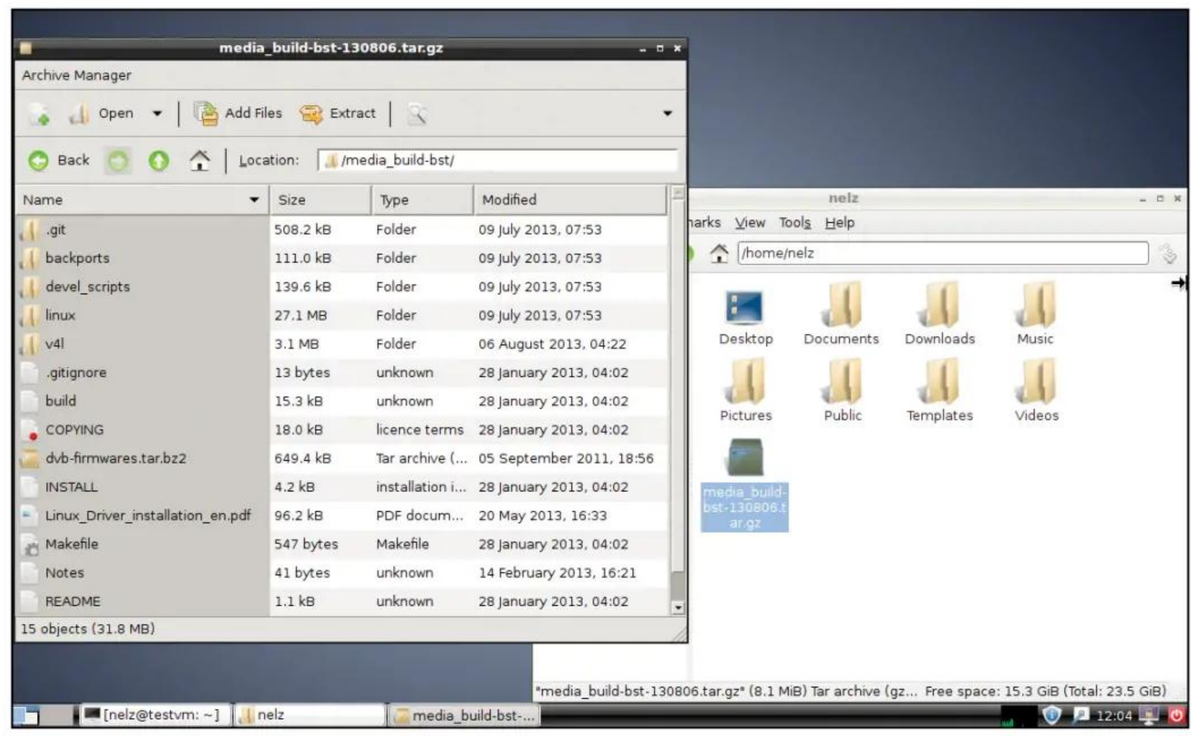

› **Most environments can view the archive contents. Shown above is the result of double-clicking a tarball in LXDE (the preferred desktop for Raspberry Pi).**

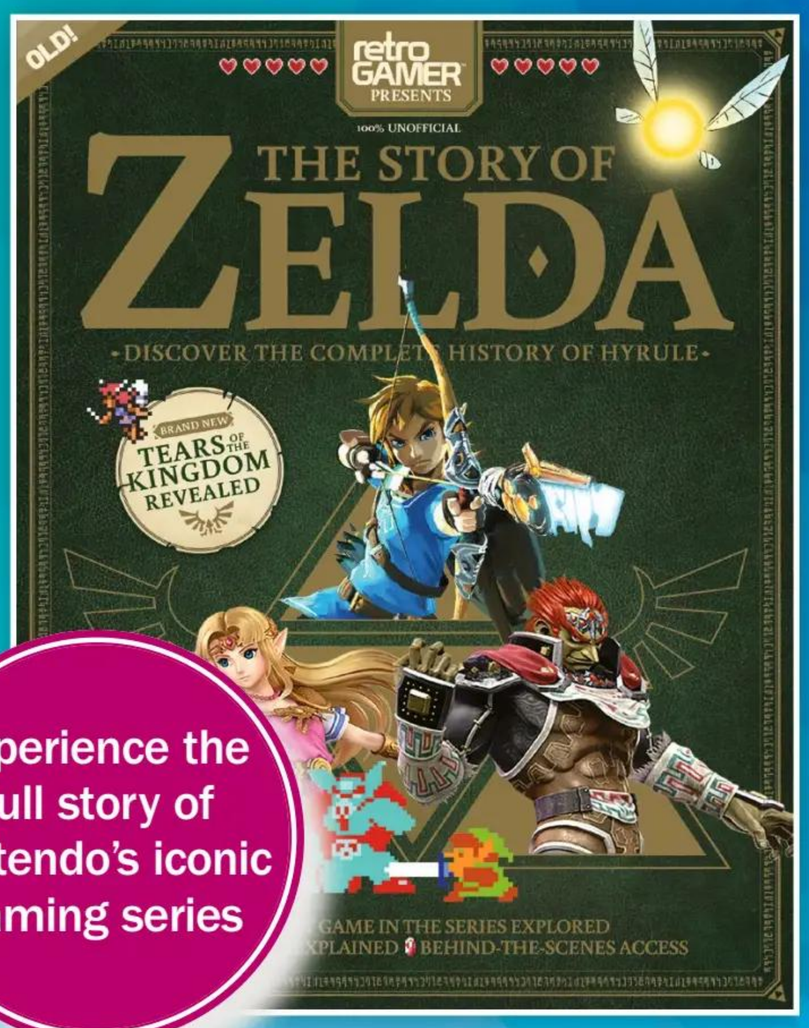

Experience the full story of Nintendo's iconic gaming series

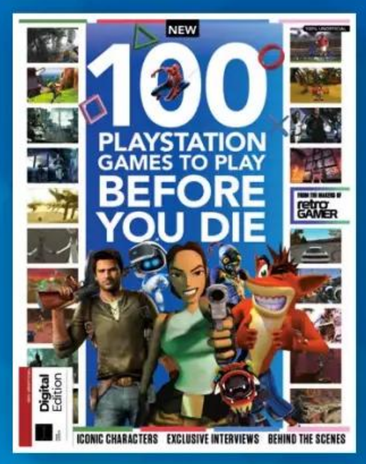

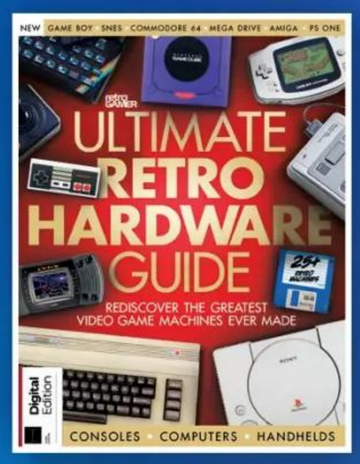

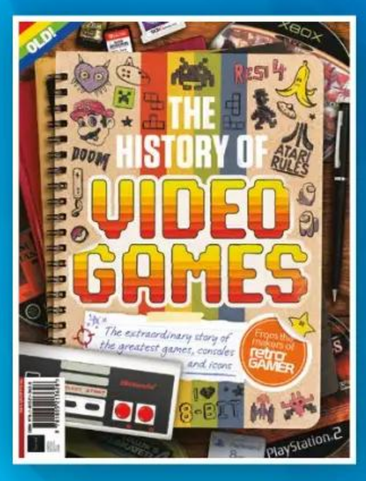

Check out the best sci-fi films, games, TV shows and books from yesteryear

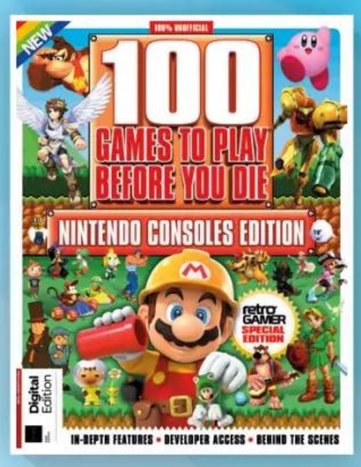

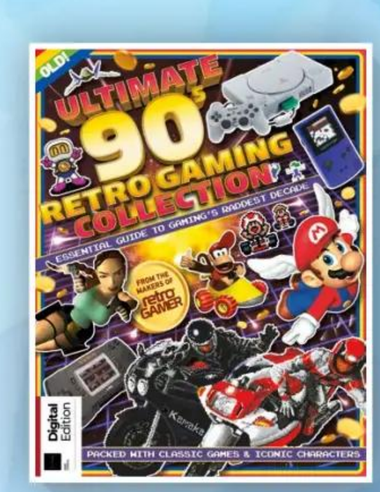

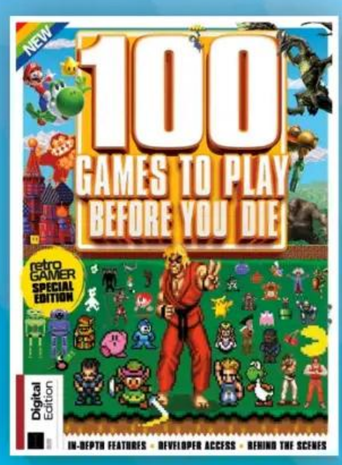

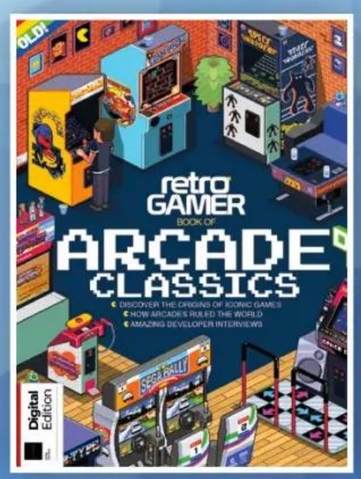

Grab a pack of batteries and explore the world of handheld retro gaming

Get great savings when you buy direct from us

1000s of great titles, many not available anywhere else

World-wide delivery and super-safe ordering

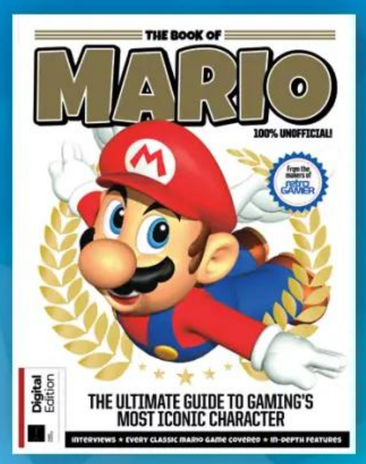

THERE'S A BOOKAZINE FOR EVERYONE

Whether you're a film buff, an avid gamer, or love all things sci-fi, we've got you covered with our range of entertainment titles

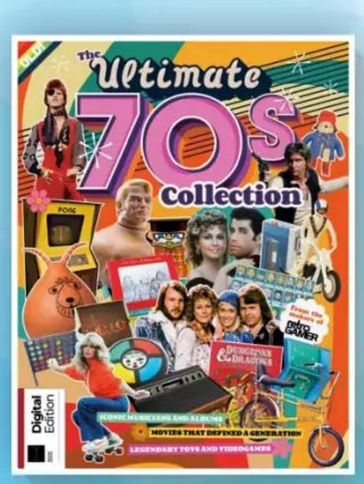

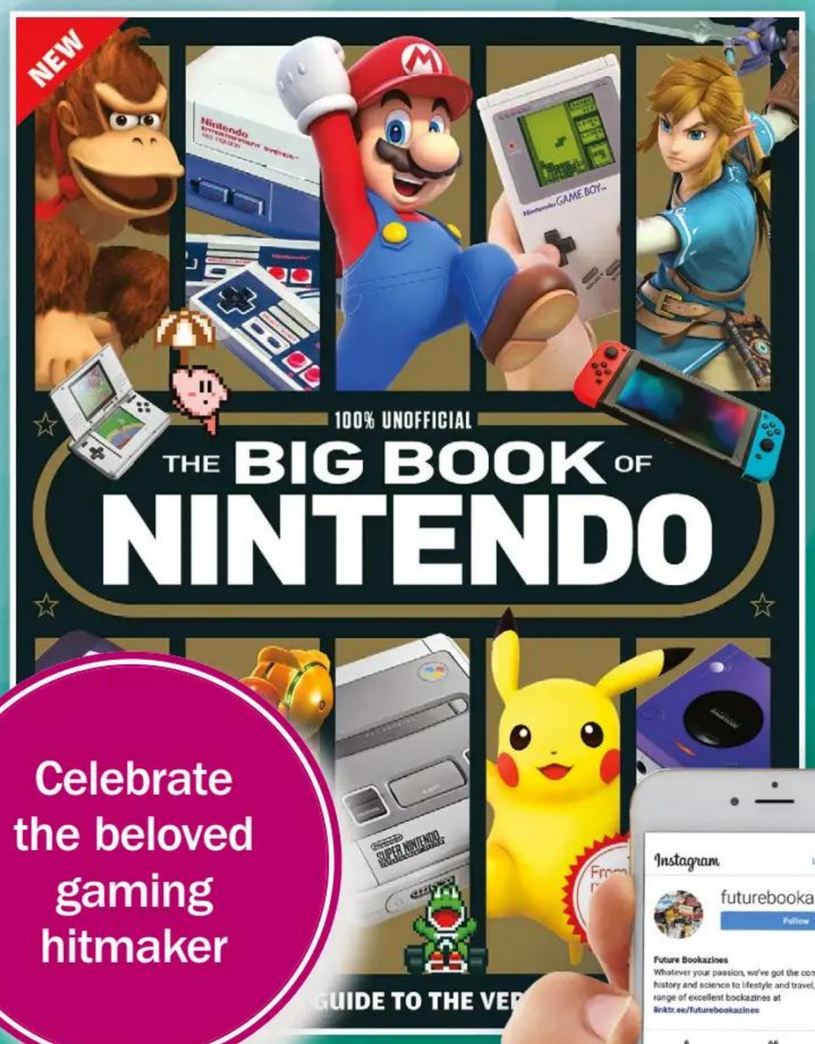

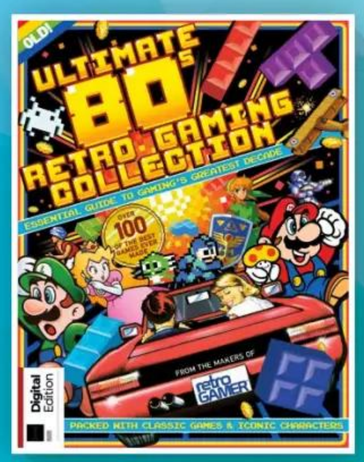

Celebrate the beloved gaming hitmaker

Follow us on Instagram @futurebookazines

FUTURE

www.magazinesdirect.com/filmandgaming

Magazines, back issues & bookazines.

SUBSCRIBE & SAVE UP TO 61%

Delivered direct to your door or straight to your device

Choose from over 80 magazines and make great savings off the store price!

Binders, books and back issues also available

Simply visit www.magazinesdirect.com

No hidden costs · Shipping included in all prices · We deliver to over 100 countries · Secure online payment

FUTURE

magazinesdirect.com

Official Magazine Subscription Store

LINUX MADE SIMPLE

Enjoy a hassle-free transition to Linux with our expert advice

Learn how to upgrade, customise and master the free OS with a collection of tutorials and guides from the Linux experts

148 pages of advice from the makers of **LINUX FORMAT**

Dozens of easy tutorials

- Try out a new Linux OS, risk free!
- Troubleshoot problems the easy way
- Connect to Chromecast, Android and more

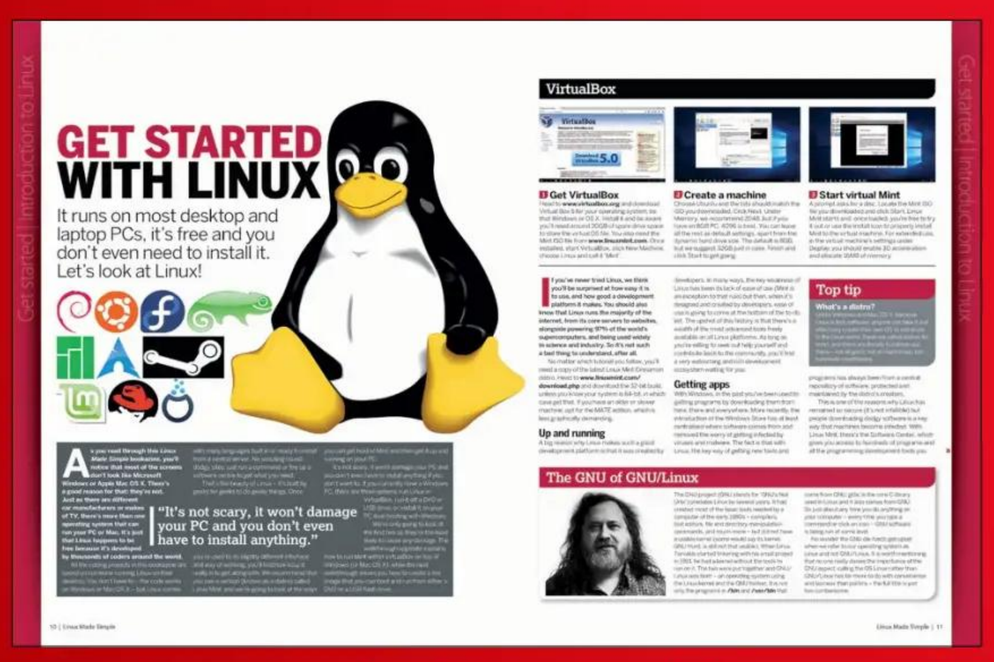

Get started with Linux and discover the right distro for you

Do everything you could do with Windows or macOS – all for free

Enjoy photos, videos, streaming audio and much more

Rescue any broken computer with a super-powered Linux toolkit

9021
9000
BOOKAZINE