

Genius Guide




A comprehensive masterclass
to becoming an instant expert

Linux & Open Source

Volume 3

The essential guide to mastering open source software and operating systems






Genius Guide

A comprehensive masterclass to
becoming an instant expert

Welcome to Linux & Open Source

As the world of Linux and open source continues to grow, so do the possibilities and opportunities open to developers, coders and everyday users. As new distros launch and current ones improve, there's a wealth of features and functions waiting to be implemented. In this book you'll find tutorials and guides on how to become a Linux master, from building faster web servers and replacing Ubuntu's Unity, to creating a blog with Django and developing Android apps. On top of that, we run through some of the best distros and software to help you get more from your torrent client, media centre and more. And if that wasn't enough, on the free DVD at the back of the book we've included four auto-booting distros for you to enjoy, including Linux Mint and openSUSE.



Genius Guide Linux & Open Source

Imagine Publishing Ltd
Richmond House
33 Richmond Hill
Bournemouth
Dorset BH2 6EZ
☎ +44 (0) 1202 586200

Website: www.imagine-publishing.co.uk
Twitter: @Books_Imagine

Facebook: www.facebook.com/ImagineBookazines

Editor in Chief
Dave Harfield

Production Editor
Jon White
Dan Collins

Senior Art Editor
Danielle Dixon

Design
Anne-Claire Pickard
Kerry Dorsey

Printed by
William Gibbons, 26 Planetary Road, Willenhall, West Midlands, WV13 3XT

Distributed in the UK & Eire by
Imagine Publishing Ltd, www.imagineshop.co.uk. Tel 01202 586200

Distributed in Australia by
Gordon & Gotch, Equinox Centre, 18 Rodborough Road, Frenchs Forest,
NSW 2086. Tel + 61 2 9972 8800

Distributed in the Rest of the World by
Marketforce, Blue Fin Building, 110 Southwark Street, London, SE1 0SU

Disclaimer

The publisher cannot accept responsibility for any unsolicited material lost or damaged in the post. All text and layout is the copyright of Imagine Publishing Ltd. Nothing in this bookazine may be reproduced in whole or part without the written permission of the publisher. All copyrights are recognised and used specifically for the purpose of criticism and review. Although the bookazine has endeavoured to ensure all information is correct at time of print, prices and availability may change. This bookazine is fully independent and not affiliated in any way with the companies mentioned herein.

Linux & Open Source Genius Guide Volume 3 © 2012 Imagine Publishing Ltd

ISBN 978-1908955319

Part of the

Linux User & Developer

bookazine series



Contents

Your guide to what's inside



Page 8
50 essential
tools to help
you get more
from Linux

Tips & Tricks

- 22 Better remote desktop with NX
- 26 Master your Linux terminal sessions
- 30 Manage your network with Zeroshell
- 34 Set up an all-in-one home and office server
- 38 Discover how to build a faster web server
- 42 Make your desktop more productive with Xmonad
- 46 An introduction to penetration testing using BackTrack 5
- 50 Build your own Linux distribution with SUSE Studio
- 54 Make your own Wikipedia
- 58 How to distribute Linux desktops to thin clients
- 62 Replace Ubuntu's Unity desktop environment
- 66 Professional podcasting with Linux
- 70 Monitor and graph your network statistics
- 74 Create reports and charts with Pentaho BI
- 78 Supercharge your R experience
- 82 Create Btrfs snapshots in openSUSE 12.1
- 86 Streaming media to your Android device
- 90 Beginner's guide to Arduino
- 94 How to interface Android with Arduino
- 98 Getting Things Done – improve workflow and productivity

Masterclasses

- 106 Back up & restore – the complete guide
- 112 Discover the best way to a perfect dual boot
- 120 Learn how to keep your system safe
- 126 Start building a blog with Django
- 130 Add content to your Django blog
- 134 Customise your blog in Django
- 138 Begin a bug hunt in LibreOffice
- 142 Find and fix known bugs in LibreOffice
- 146 Become an open-source mentor in LibreOffice





Developer Guides

- 152 System logging in the cloud with Papertrail
- 156 Better PHP on the command line
- 160 Server-side JavaScript with Node.js
- 164 Create an HTML photo editor
- 168 Cross-platform mobile app development
- 172 Cross-platform mobile web app development
- 176 Using Hadoop tasks on Amazon's Elastic MapReduce
- 180 Develop on your Android smartphone with AIDE
- 184 Building LibreOffice extensions



Essential Software

- 190 The best desktop distros
- 196 Torrent clients
- 200 Media centres
- 204 Mind-mapping software
- 208 Project managers
- 212 CD rippers
- 216 Email clients
- 220 Password managers

On the disc

- 226 Four live distros, tutorial files and much more



50 essential tools





50

ESSENTIAL TOOLS TO HELP YOU GET MORE FROM LINUX

Squeeze every last bit of usability from your Linux system with these amazing tools

In a world where the focus on a lot of Linux distributions is usability, you're usually supplied with a small suite of useful tools, applications, and libraries to get you going the moment your system is installed. With popular software choices like Firefox and LibreOffice on just about every distro, you can get the basics done without going hunting for any more packages.

Of course one of the great things about Linux is that you don't have to settle for the basics. With a staggering amount of applications, tools, and other packages available, you can both finely streamline and

greatly enhance your day-to-day activities. You can create your own custom user experience that does everything you'd ever want your computer to do.

In this feature, we're going to highlight some of the top tools from the open-source community that can help with exactly that. Whether you want to make better use of the internet for both work and play, make your software development more efficient and helpful, or just add some utilities to better manage your system, we have a selection of tools that are a must-have for any Linux power user.



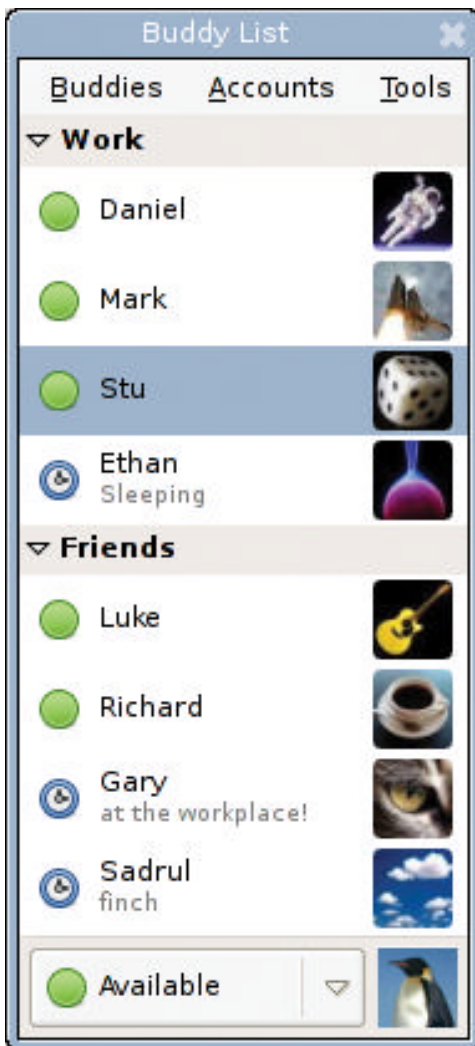
Web & social media

The internet is continually evolving, with more bandwidth and more users meaning we can do more and find out more on the web every single day. We use it to converse with friends, engage in global conversations, share our views, and consume media of any kind. The open-source community has developed plenty of tools to make using today's internet as effective as possible.

Turpial

turpial.org/ve/

Social networks can be a major part of your day-to-day life, and the browser interface is not always the best way to keep up with them. Turpial is a very lightweight Twitter client that integrates seamlessly with desktop notification services and offers more functionality. You can temporarily silence users, choose alternative URL shorteners, and create extra columns for mentions, direct messages and searches.



Pidgin

www.pidgin.im/

Pidgin is an instant messaging client that consolidates the most popular chat services into one simple application. It has support for Google Talk, Windows Live Messenger, AIM, Facebook, Yahoo, IRC, ICQ, and any other IM service that supports XMPP. Not only does it have this large selection of connection types, it's also fully extensible, with a selection of plug-ins included with any Pidgin install.

AWStats

awstats.sourceforge.net/


AWStats is a fully featured tool that creates advanced web, streaming, FTP or mail server statistics. This log analyser can be set up in both CGI and via command line, and displays all the relevant information in graphical web pages. It can analyse log files from all major server tools and a lot of other web, proxy, WAP, streaming servers, mail servers and some FTP servers.



FileZilla

filezilla-project.org/

When you have web space, and you want an easy way to upload files quickly, an FTP client is of course the best option to do so. Instead of making your way through your web hostings user panel, you can use FileZilla to instantly connect to your webspace and upload your files with minimum hassle. Sporting a straightforward graphical interface, you can easily manage connections, network speeds, and remote edit while also uploading files.



WordPress

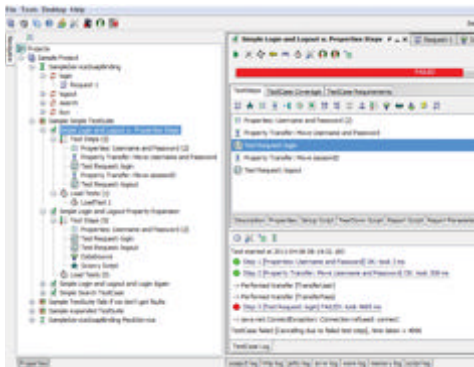
wordpress.org/

A lot of different, free blogging services have been and gone, and the only thing that's stuck around is an actual blogging tool itself, WordPress. Built in PHP, WordPress is a complete software package that can help you create any kind of content delivery service on the web.

Having your own blog rather than relying on free services has a lot of advantages. Firstly, you have absolute control over it, from the aesthetics and what media it can display to being able to give it a proper web address. Secondly, with WordPress, you have an editable and highly extendible interface.

WordPress also makes installing and setting up a PHP Blog very easy. Simply uploading the files to some webspace creates a front page that guides you through a straightforward setup process, creating appropriate database tables, administrators, and generally getting your new site ready to post blogs immediately.

WordPress by default allows you to create custom menus, set up RSS feeds, have categories and tags, and create and manage a multitude of users with different permissions. It really is the ultimate blogging package.



SoapUI

www.soapui.org/

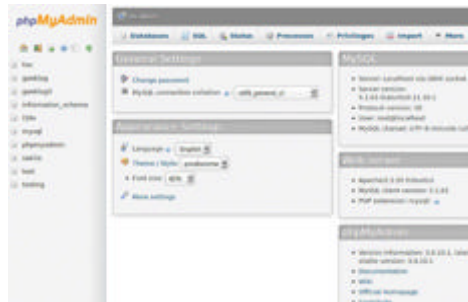
If you need a way to test an online service or site you've been working on, then you can't go wrong with SoapUI. It's a very functional testing solution, with an easy-to-use graphical interface and enterprise-class features. SoapUI enables you to easily and quickly create automated functional, regression, compliance and load tests that can be executed immediately. In a single test environment, SoapUI provides complete test coverage and supports a lot of major, standard protocols and technologies. The interface is simple to navigate, and allows you to keep tabs on all your ongoing and completed tests. It's also highly extendible with various plugins, and you can also plug SoapUI directly into various IDEs.



TorrentFlux

sourceforge.net/projects/torrentflux/

If you have a home server or separate system that handles your torrents, you may not have found a good way to manage what you're actually torrenting. TorrentFlux adds an easy-to-use PHP front end to BitTorrent and BitTornado that allows you to add, remove, and generally look after your torrenting from anywhere that has an internet connection.



phpMyAdmin

www.phpmyadmin.net/

phpMyAdmin is a tool that is designed to handle the administration of MySQL databases over the internet with a great web interface. It supports a wide range of operations with MySQL, with the most frequently used functions supported by the user interface, such as managing databases, tables, fields, relations, indexes, users, permissions, and more. You also have the ability to directly execute any SQL statement if you need that extra control.



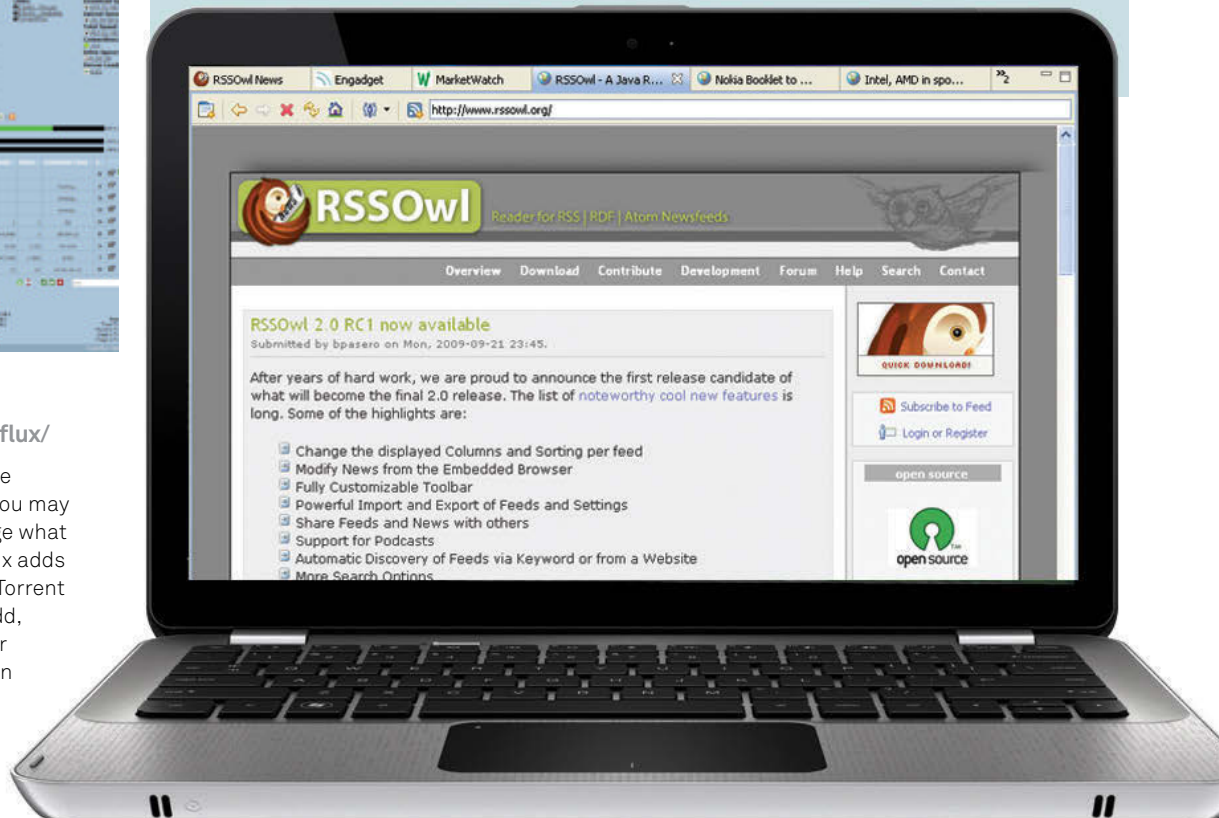
MediaWiki

www.mediawiki.org/

MediaWiki is the software base that powers Wikipedia. MediaWiki can be used as a powerful collaboration tool to help document and develop projects, or just to create a knowledge base on any subject. All the great features of Wikipedia are part of the MediaWiki package, such as intuitive editing that allows for simple formatting and inline links, discussion pages, and a version control system so you can roll back to any previous iteration of a page.

RSSOwl www.rssowl.org/

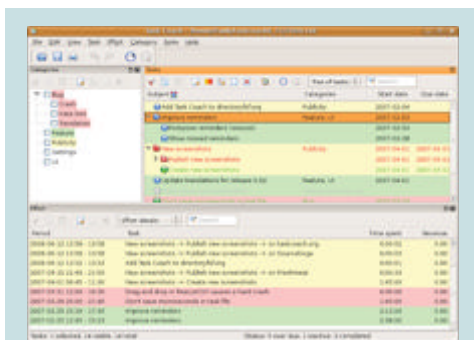
RSSOwl is a handy RSS client that lets you aggregate various feeds, as well as the ability to sync with Google Reader. It has an embedded browser, a newspaper view mode, and it can even create automatic feeds based on keywords. The fantastic interface is highly customisable, and even includes tabbed browsing to emulate the web for those that prefer it.



50 essential tools

Office

Office tools are the backbone of any operating system, and with Linux you're absolutely spoiled for choice. From complete office suites to simple task managers, there's something for everyone, from those in high-level corporate environments to students writing up their final projects. In this section we're going to cover some of the best applications that help you work more efficiently.



Task Coach

taskcoach.org/

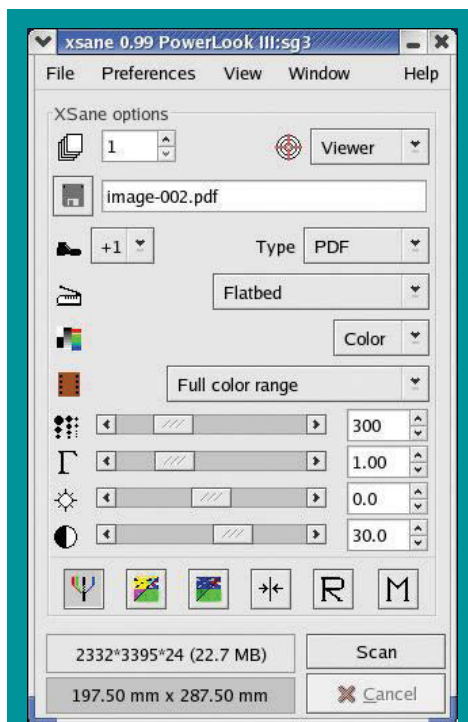
A simple yet powerful to-do manager, Task Coach helps you to keep track of personal tasks and to-do lists. It is designed for composite tasks, and also offers effort tracking, categories, notes and more. Tasks can be created by various methods, such as dragging emails with task details into the Task Coach window, and your to-do list can be exported to HTML, CSV, and even iCal.



GoldenDict

goldendict.org/


GoldenDict is a feature-rich dictionary look-up application, with support for multiple dictionary file formats, and with the ability to search online sources such as any Wiktionary and any MediaWiki-powered site. You can also set up text-to-speech to read out pronunciations of words, and use arbitrary websites using specific, templated URL patterns as a dictionary source.



XSane

www.xsane.org/

XSane is a graphical frontend for SANE-library, which is an interface for scanners. Using XSane, you can scan to file, photocopy, create a fax, create and send email, and use it in GIMP as a plug-in.



PDF Split and Merge

www.pdfsam.org/

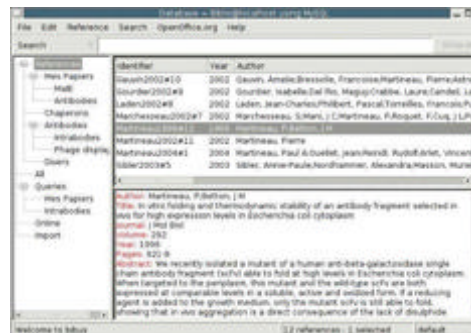
PDF Split and Merge is a nice and simple, easy-to-use utility to split, merge, and mix PDF files. It can be used in a graphical interface, or in the command line.



FreeMind

freemind.sourceforge.net

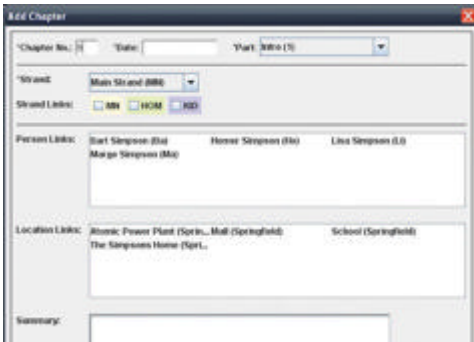
FreeMind is a Java-based mind-mapping tool, which allows people to edit a hierarchical set of ideas around a central concept. The non-linear approach of this tool can assist in brainstorming new outlines and projects as ideas are added around the mind map. FreeMind has a great, understandable interface which allows you to easily note down your brainstorm. You can create nodes graphically, but you can also code them in HTML, or paste web content and have FreeMind smartly arrange it. FreeMind can also export to a number of formats, such as ODF, PDF, XML, and SVG images so you can share your mind map with anyone. It's a great productivity tool that can help you get the most out of a project.



Bibus

bibus-biblio.sourceforge.net/

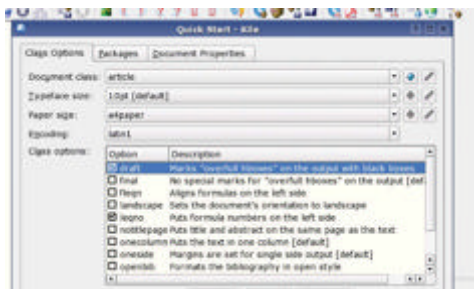
Keeping track of references for longer projects can get tricky, especially when your bibliography list has long since outnumbered your fingers and toes. The Bibus Bibliographic Database is bibliographic and reference management software, with the ability to search, edit, and sort records. Bibus can be used in conjunction with both LibreOffice and OpenOffice to insert and format bibliographies properly, so there's no need to slowly copy and paste everything in.



Storybook

www.novelist.ch/joomla/

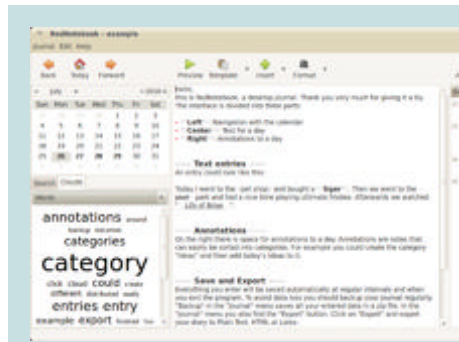
Storybook is novel-writing software for creative writers, novelists, and authors, for use when a word processor or text editor just doesn't cut it. Storybook lets you create an overarching plot that you can refer to whenever you wish without scrolling up or down pages, allowing you to never lose overview if you're deep into story writing, while also aiding you in the structuring of your final product. You can also manage data on characters, locations, scenes, items, tags and ideas in one place. A simple interface is provided to enable you to assign your defined objects to each scene and to keep an overview of your work with user-friendly chart tools. There are even separate views so you can keep track of chronology and chapters.



Kile

kile.sourceforge.net/

Kile is an Integrated LaTeX Environment – a user-friendly TeX/LaTeX editor using Qt that has a rich feature set enabling you to easily create and edit high-quality TeX and LaTeX files for papers, projects, and presentations. Kile is a powerful tool in this regard, letting you compile, convert and view your documents with just one click, collecting documents together into one project, a QuickPreview for parts of documents, and an auto-completion of relevant commands.



RedNotebook

rednotebook.sourceforge.net/

Taking notes throughout the day is usually a good habit to get into. RedNotebook is the equivalent of a journal, with all the enhancements of modern software. With calendar navigation, tagging entries, text search, and the ability to export to PDF, HTML, and LaTeX, you'll be able to easily find any specific past activity.

LibreOffice www.libreoffice.org/

Everything we've featured so far has had a feature set focused on one or a few things. LibreOffice, though, is an entire suite of tools that make up just about the rest of everything you'd need in office software, and is the best office suite on Linux. A fork of OpenOffice, LibreOffice was started to fix the thousands of neglected bugs in the OO codebase. Now at version 3.5, a huge amount of these bugs have been squashed, and a lot of great new features have found their way into the suite.

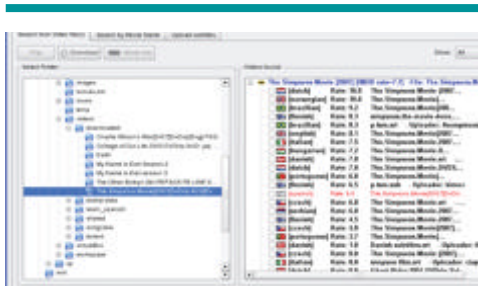
LibreOffice comes with five main applications. Writer is the word processor that's simple enough for a quick note, but powerful enough to create full project reports with contents, diagrams, indexes, and more. For spreadsheets you can use Calc, which includes all the mathematical and organising tools you'd expect from a spreadsheet program. Impress is a fantastic tool for creating presentations, with full multimedia support and the usual selection of effects, transitions, and animations. Base is a fully featured, desktop database front end that is accessible by any level of user for any task they require. There's also Draw and Math, two tools to create graphs and presentable math equations respectively.

In recent updates it's gained a more user-friendly UI that is still customisable, a more robust grammar detection system that works more naturally for the English language, and an update checker that notifies you of newer versions and allows for manual download of them. It's getting better all the time, with community-driven bug fixing, and more features being added. It really is an amazing suite of software, and a true open-source success story.



Entertainment


Using Linux isn't just about working, and most distributions will come with some way to listen to your music or watch a DVD. As we're talking about the open-source community though, there is a huge amount of software that can help you consume and curate your media library whether you're in your living room or at work.



SubDownloader

subdownloader.net/

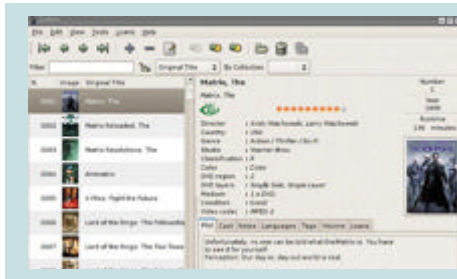
Need subtitles for a film? SubDownloader allows for automatic downloading and uploading of subs for video files using fast hashing. It can autodetect languages and can scan 27GB of movies in seven seconds.



HandBrake

handbrake.fr/


With HandBrake, you'll never have to find your copy of your favourite DVD again, as it is a powerful DVD/Blu-ray ripper. Whether you want a plain video file, or to maintain the chapters and subtitles, HandBrake has you covered.



Griffith

griffith.cc/

If you've ever used XBMC or Boxee, you'll know that they scrape information from the internet to add metadata to the shows you're watching. Griffith is an application that enables you to grab this info for your own database, where you can manage your entire media collection.



DigiKam

www.digikam.org/

DigiKam is an advanced digital photo management application. Your photos are organized in albums which can be sorted chronologically, by folder layout, or by custom collections.

VLC

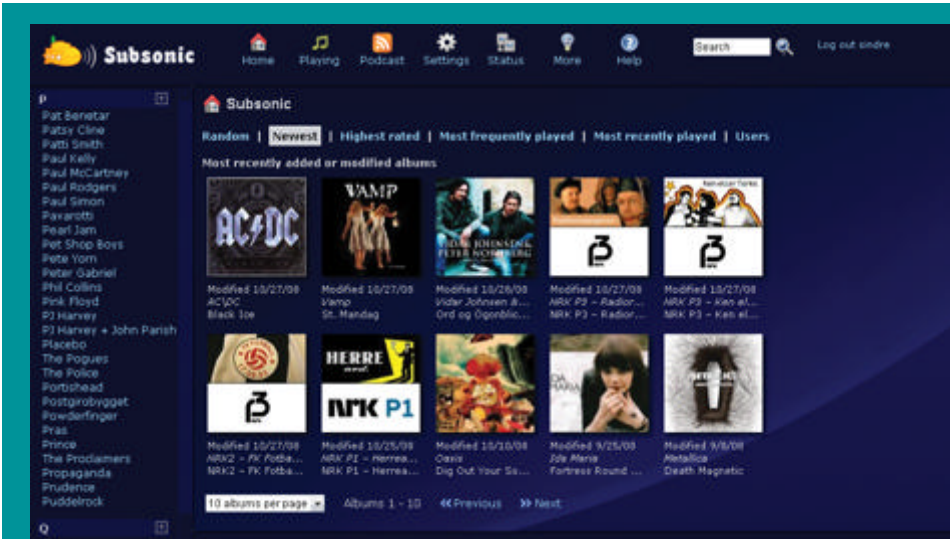
www.videolan.org/vlc/index.html

VLC isn't just a powerful and lightweight media player that can decode just about any type of file format – it's actually a lot more than that. Starting off as the Client to VideoLAN Server (VLS), at the very core of VLC is the ability to view any type of web or network stream. Over the past few years, VLC has evolved past the need for VLS, even being able to broadcast streams itself.

VLC can send out streams of any available media, video files, music files, and DVD video if you've configured it to play DVDs. You can also stream video of the current desktop, with complete control over the sample rate and quality of the video. It can also record the live video of your desktop, much like it can do with a network or internet stream, encoding compressed video on the fly if you have a powerful machine. You can even set a local file as the stream, as long as you have a setup that's able to convert video or audio from one file type to another.

This shouldn't diminish from its standard uses as a media player though, especially as one of the very few Linux media players that can actually play Blu-ray discs. It also has rich playlist creation and editing, is incredibly extensible, and can perform multi-core decoding of high-definition content, as well as playback 10-bit h264 video. It remains one of the most powerful all-in-one media playing, streaming, and recording solutions.



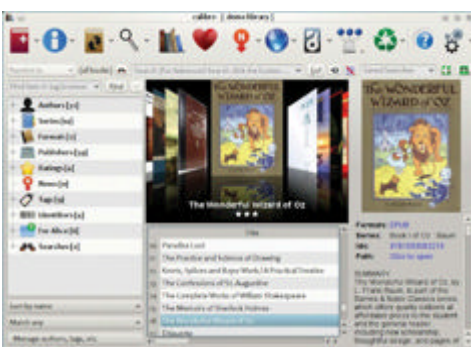


Subsonic www.subsonic.org

A web-based media streamer, Subsonic proves complete access to your music on the host machine anywhere in the world. You can use it to listen to your own music while at work without carrying around extra storage, or even have it play throughout a house with its ability to stream to multiple players simultaneously.


It's designed to handle very large music collections, and although optimised for MP3 streaming, it can be used with any media format that can be streamed over HTTP, such as OGG. Subsonic also supports transcoding and streaming of virtually any audio format, including FLAC, APE, WMA, Musepack, WavPack and Shorten.

Subsonic allows you to set bitrate limits in the event that you have limited upspeed or general bandwidth, in which case it resamples the music for the stream. In addition to being a streaming media server, it also works very well as a sort of jukebox, as the web interface, which includes search and indexing facilities, is optimised for efficient browsing through large media libraries.



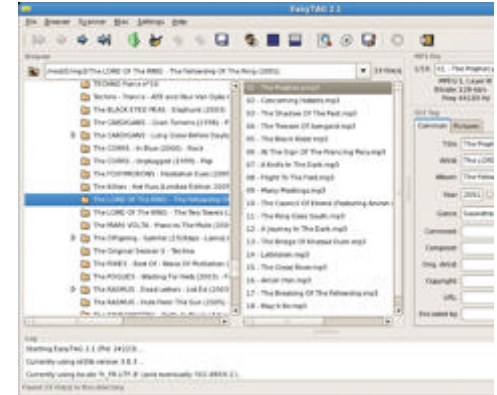
Calibre calibre-ebook.com/

Ebooks are a fantastic way to keep a veritable library on something no larger than a USB stick. Calibre is one of the best ebook management applications available, helping you properly organise what you're reading. It also contains a great converter, enabling you to convert a huge number of file types, including PDFs and ePubs, to almost any format you'd need.



gPodder gpodder.org/

gPodder is a podcast aggregator, finding and downloading the latest episodes in your subscribed list with the added ability of managing podcasts on a large selection of audio devices. It includes in-built playback software for both video and audio casts, along with a searchable database of shows in case you're running low, or want to add something new.



EasyTag easytag.sourceforge.net/

EasyTag is a remarkably powerful id3 tag editor for just about any kind of audio file that supports it. With an advanced search and index function, you can find which files are not properly formatted, edit id3 information, create batch formatting operations, and even add an image tag.



Clementine www.clementine-player.org/

Clementine isn't just a standard audio player; it comes with a whole host of tools to make your music listening experience better, such as the ability to create and maintain dynamic playlists. These playlists are updated as you rate and listen to your music, modifying the song selection based on what you're currently playing and liking. You can also use it to stream music from online services such as Spotify, Last.fm, Grooveshark, and more, meaning you can keep all your music playing in one application. To keep with the trend of staying in the software, portable music devices can also be synced and managed with Clementine, without the need for an external application. On top of all that, it's a great audio player that integrates with desktop notification systems, and can even be controlled by a Wii Remote.

Development

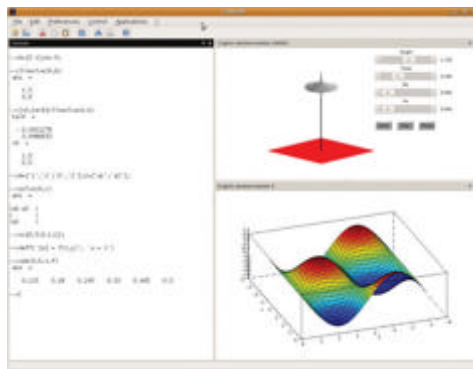
All the amazing software we've talked about so far would not have existed if it wasn't for the development efforts of the Linux community. Whether you're a software engineer dealing in cloud computing, or a hobbyist wanting to fix a bug in your favourite application, the tools for creating, editing, and refining are mostly the same. Here are ten of the best.



Valgrind

valgrind.org/


Valgrind is an instrumentation framework for building dynamic analysis tools. There are Valgrind tools that can automatically detect a lot of memory management and threading bugs, while profiling your programs in detail. Valgrind comes with six great tools, including a memory error detector, two thread error detectors, a cache and branch-prediction profiler, a call-graph generating cache and branch-prediction profiler, and a heap profiler. It's essential if you want to make sure your programs are as stable as possible.



Scilab

www.scilab.org/

Scilab is an open-source alternative to MATLAB, with similar features and coding syntax. Like MATLAB, Scilab is a numerically-orientated programming language that can be used in conjunction with other languages to process high-level problems involving a lot of metrics.



Gtranslator

projects.gnome.org/gtranslator/


If you're creating a new GNOME fork, or want to help in the translation of GNOME desktop itself, then Gtranslator is a tool you need to use to in the translation process. It handles the gettext po files, and includes some standard features like Find/Replace, Translation Memory, and the ability to have different Translator Profiles. It's also extensible through plug-ins.



Meld

meldmerge.org/

Meld is a visually orientated diff and merge tool developed with coding in mind. It can help you compare files, directories, and version controlled projects so that you can keep on top of any changes made to the code. As well as handling comparison of two files, Meld also has the ability to make three-way comparisons of both files and directories. It can filter unnecessary text, highlight syntax, and updates the comparisons on the fly as you edit.




Geany www.geany.org/

Of course the most important tool for any developer to have is a great multi-purpose development environment, and Geany is one of the best available on the Linux platform. Geany is a lightweight Integrated Development Environment, which was specifically developed as a small and fast IDE. Geany was also created in such a way where it has only a few dependencies from other packages, and is completely desktop environment agnostic, only needing GTK2 runtime libraries.

Geany comes with a number of features (some basic, some more advanced), and supports almost 50 programming languages, including C, Java, PHP, HTML, Python, Perl, and Pascal. Of course it contains syntax highlighting, allowing you to not only determine different functions and variables in your code, but also if you haven't correctly closed a bracket or case. It has code folding, allowing you to highlight two parts of code separate from each, and then compare them side by side by hiding the irrelevant code between them. Geany even has its own build system, which lets you compile and execute your code without having to open a separate application.


Geany has a rich plug-in system as well, and comes bundled with a few that can help you convert files from one format to another, add a tab for file browsing, or split it into two windows. There's a dedicated repository of third-party plug-ins for Geany available from its website that add a staggering amount of different features.



TEA

tea-editor.sourceforge.net/

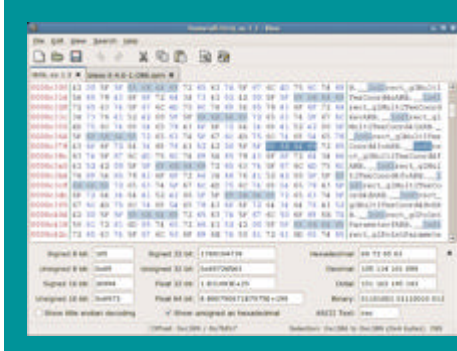
Every system needs a text editor, and while gedit and nano are good, TEA is the best of the bunch. The reason for this is because TEA isn't just a plain text editor, although it can of course be used that way from the terminal. TEA is about as close to an IDE that a text editor can get, not only because of its great syntax support, but also because it acts as a file manager. You can browse your file system in quite some depth before opening multiple files to work on, and there's even a to-do list to prioritise work.



gitg

freecode.com/projects/gitg


It's no secret that Git is one of the best ways to manage source code, with fantastic revision control features. Outside of a web browser, one of the best ways to visualise a Git repository is by using Gitg. With Gitg you can view the source and revision history through a straightforward interface, and it allows you to make commits.



Bless Hex Editor

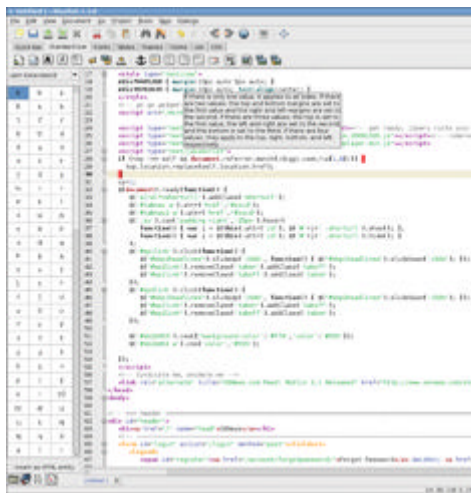
home.gna.org/bless/

In case you ever need to start editing Hex, Bless has you covered. It comes with a lot modern text editor functions, such as tabbed files, find and replace options, copy and pasting, and the ability to highlight selection patterns. As well as all the UI stuff, you can edit large data files and block devices, and export data to text and HTML thanks to a plug-in system.



Arduino arduino.cc

Arduino is of course the open-source micro-controller that can be used to power and control a whole manner of physical devices. You can also get a great IDE designed specifically to create code, or sketches, for Arduino from the Arduino website. Sketches can be written and uploaded on the fly to the board, with an onboard facility to check for errors, and an area in the IDE that shows messages from the board.



Bluefish

bluefish.openoffice.nl/index.html

While other IDEs will allow you to easily code in web-based languages, it can be useful to have a separate application for creating webcode. Bluefish is a lightweight, graphical editor and IDE mainly aimed towards web development, with advanced support for web-based languages and scripts such as auto-completion on JavaScript.

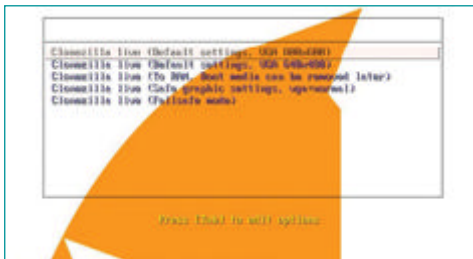
It also includes a specialised, customisable toolbar for creating HTML code, which includes dialogs and wizards for HTML tags, with information on all their attributes. It also has a fully featured image insert dialog that allows for thumbnail creation and automatic linking with the original image, and multi-thumbnail generation for easy creation of photo albums or screenshot pages.

Bluefish can be used for other types of programming, with syntax highlighting for a large number of languages. This includes C/C++, Python, Perl, etc, and has some degree of code awareness, allowing for automatic closing of tags, and spell checking for text inside comments. It's also optimised for having multiple files and projects open at once, with support for remote files using GVFS, if you're set up correctly.

50 essential tools

Utilities

Thanks to the total customisability of Linux, you can do just about anything to your system with the right know-how. But even if you don't have the know-how, there are plenty of tools you can get to guide you through just about every part of your distribution. In this section we'll be covering the utilities to maintain and manage your system.




Clonezilla

clonezilla.org/

You'll probably hear us repeat ourselves again before this book is over, but backing up is very important. One of the most complete ways to back up your system is to make an image of your hard drive, enabling you to completely restore it at a later date. Clonezilla is a tool that helps you do just that, being an open-source clone of Norton's Ghost software.

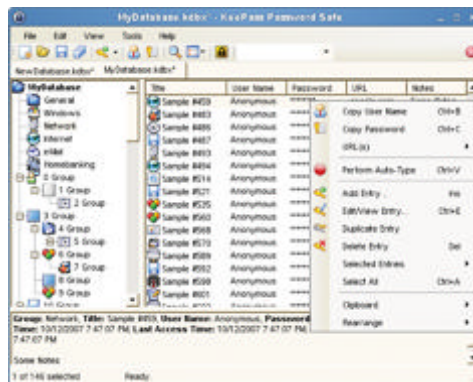
Clonezilla comes in two forms, a live style image that you need to burn to disc, or as a program on some system rescue distros such as Parted Magic or SystemRescueCD. It supports all major file system types, and you can save whole discs or single partitions. These can be saved to a local disk, or over the network to SSH, NFS, and SMB directories.



FreeFileSync

freefilesync.sourceforge.net/

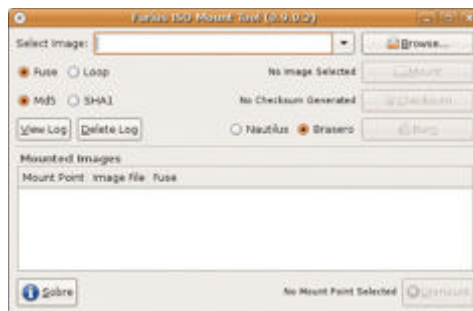
FreeFileSync is a visual folder comparison and synchronisation tool designed for keeping the contents of specific directories the same. It can detect moved and renamed folders so that there are no duplicates. FreeFileSync is a great way to copy files over multiple installs, or sync music to external storage.



KeePass

keepass.info/

KeePass is a fantastic password manager that helps you manage your passwords both easily and securely. You can put all your passwords in one database, which is locked with a master key or key file so that you only have to remember one master password or select the key file to unlock the whole database – not just the passwords. You can also export to a number of formats, including txt, HTML and XML.



Furius ISO Mount

launchpad.net/furiusisomount/

Furius ISO Mount is an image manager that can perform two main different tasks. Firstly, it allows you to mount image files to a virtual drive so that you can access them without having to burn a disc. Secondly, you can actually use the software to burn images to disc. It's also available in several languages other than English.




Conky

conky.sourceforge.net/

Having a good idea of what's going on in your system can be an easy way to diagnose faults and issues with packages and other software. It's especially handy if you're developing new programs and want to keep an eye on the system load. You can of course just run top; however, that only gives a small part of the story.

Here's where Conky can help. Conky is a highly configurable application that allows you to have unprecedented access to a wide array of information regarding your system. With a huge library of available parameters that can be displayed in text, numbers, or bars, it can be customised to display accurate information of the exact data you want. You can also completely change the look, location, and behaviour of Conky to suit your needs.


It's not just system settings it can track either. It has built-in support for tracking your music, with song title, time elapsed, and next song in playlist able to be displayed. You can also keep tabs on the Weather, including the use of symbols if set up correctly. It can be a little tricky to set up at first, but there are plenty of configuration files from folks online that you can borrow or copy.



Remmina

remmina.sourceforge.net/


When it comes to remote desktop clients, Remmina is one of the best applications available. It gives you complete control over how you connect to different machines, and includes VNC and RDP protocols for connecting to both Linux and Windows machines alike. Remmina gives you a lot of options to get the best result for your connection, with the ability to lower the quality or colour depth to allow for a smoother experience. Remote desktop connections are a great way to manage and maintain computers from a distance, and Remmina lets you save connections to easily and quickly access any system.



Bacula

www.bacula.org/


We can never stress enough how important it is to back up your software on as regular a basis as possible. We also know though that it's not something you always remember to do, and maybe you don't always have time to get it set up. This is why we like applications like Bacula, which not only offers advanced backup procedures, but also includes scheduling. With a schedule properly set up, you barely have to interact with the software for it to do its job. While Bacula will work fine on a single machine, it can also manage networked systems.



GParted

gparted.sourceforge.net/

GParted is a graphical partition editor that can be used to manage your hard drives and their partitions. GParted has support for just about every file system, and uses this to display a list of partitions for any selected drive. The interface will show you the file system of each partition, where it's mounted, and there's a graphical representation of their order in the hard drive, including how much space has been used.



MyUnity

launchpad.net/~myunity

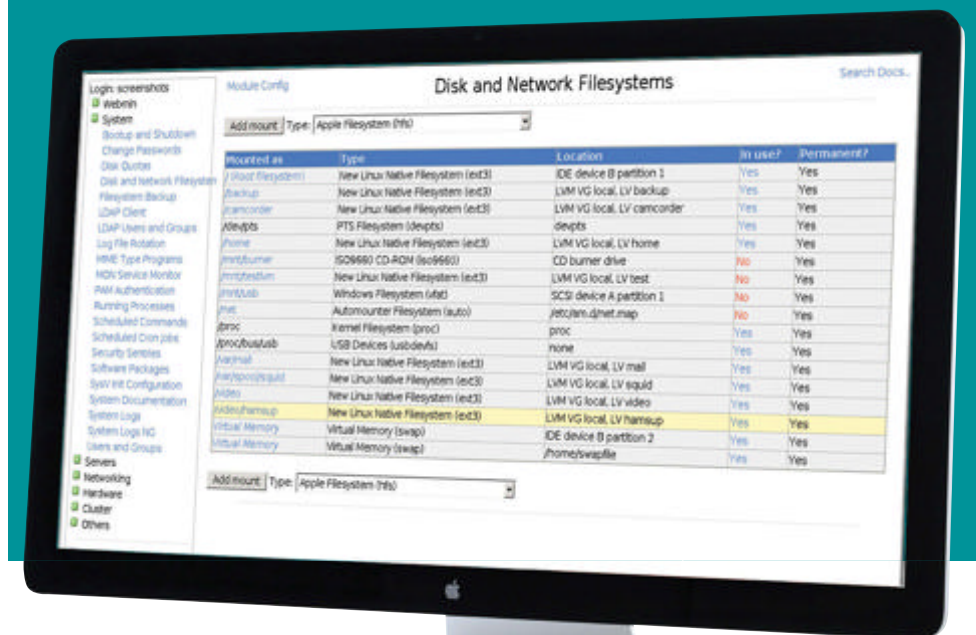
Even though customisation of Unity is becoming easier and easier in each successive version of Ubuntu, there are still a lot of options unavailable in the standard software settings. MyUnity seeks to fix that by giving you every option in a graphical program. You can use it to change icon sizes, gain more control over the sidebar's behaviour, include and remove more system icons and more features.

Webmin www.webmin.com/

You can't always be at a system you need to be at to do some maintenance, even if you've got a highly configured remote desktop solution. Even then, it can be tedious to update and manage multiple machines at once if you're having to connect to one at a time. This is where Webmin comes in – a web-based interface for system administration for Linux systems. Using any browser that supports tables, forms, and Java, you can set up user accounts, Apache, DNS, file sharing and so on.

Webmin consists of a simple web server, and a number of CGI programs which directly update system files. There's a huge number of additional modules that can be added to Webmin, adding abilities such as manual and scheduled backups, the ability to burn images to CDs from the systems, changing passwords, scheduling commands, and much more.

The website warns over the software only being optimised for specific operating systems, but don't worry – this includes all major Linux distributions. It's very powerful, with an advanced and extensive feature set that makes it the perfect way to manage all your home computers, or maybe even all the computers in your office. While it may not replace you physically being at the systems at all times, it can reduce your time away from yours.



Tips & Tricks

Improve the way you work

22 Better remote desktop with NX

Connect to your remote desktop using the NX protocol

26 Master your terminal sessions

Learn to detach and reattach command-line programs to their terminals

30 Manage your network with Zeroshell

Prioritise and administer your network traffic using this router distro

34 Set up an all-in-one home and office server

Discover how to install and configure Superb Mini Server

38 Build a faster web server

Switch out Apache for the lightweight Nginx HTTP server

42 Make your desktop more productive with Xmonad

Maximise your screen space with this tiling window manager

46 An introduction to penetration testing using BackTrack 5

Test your network's security using this essential distro

50 Build your own Linux distribution with SUSE Studio

Create your very own distribution using the popular SUSE Studio

54 Make your own Wikipedia

Design your own wiki site using the same software that powers the real thing

58 How to distribute Linux desktops to thin clients

Use the Linux Terminal Server Project to serve Linux desktops to thin clients

62 Replace Ubuntu's Unity desktop environment

Configure your system to run MATE, GNOME, Shell or Mint's Cinnamon

66 Professional podcasting with Linux

Create podcasts like the pros using the latest version of Audacity

70 Monitor and graph your network statistics

Create graphical representations of network data statistics in easy steps

74 Create reports and charts with Pentaho BI

Use this open source intelligence suite to gain an insight into your business

78 Supercharge your R experience

Analyse large amounts of data by running your R jobs in parallel

82 Create Btrfs snapshots in openSUSE 12.1

Reverse all your mistakes using the revolutionary Snapper tool

86 Streaming media to your Android device

Set up a full media and file-serving solution using FreeNAS

90 The beginner's guide to Arduino

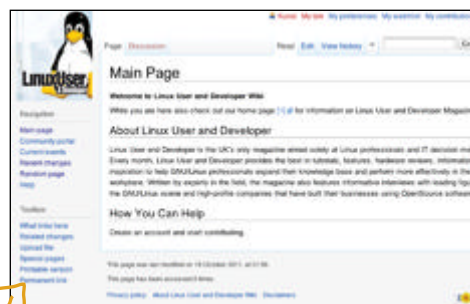
Take your first steps using the Arduino prototyping board

94 How to interface Android with Arduino

Use your Android device to trigger actions on an Arduino board

98 Getting Things Done

Improve your productivity and improve workflow using open-source software



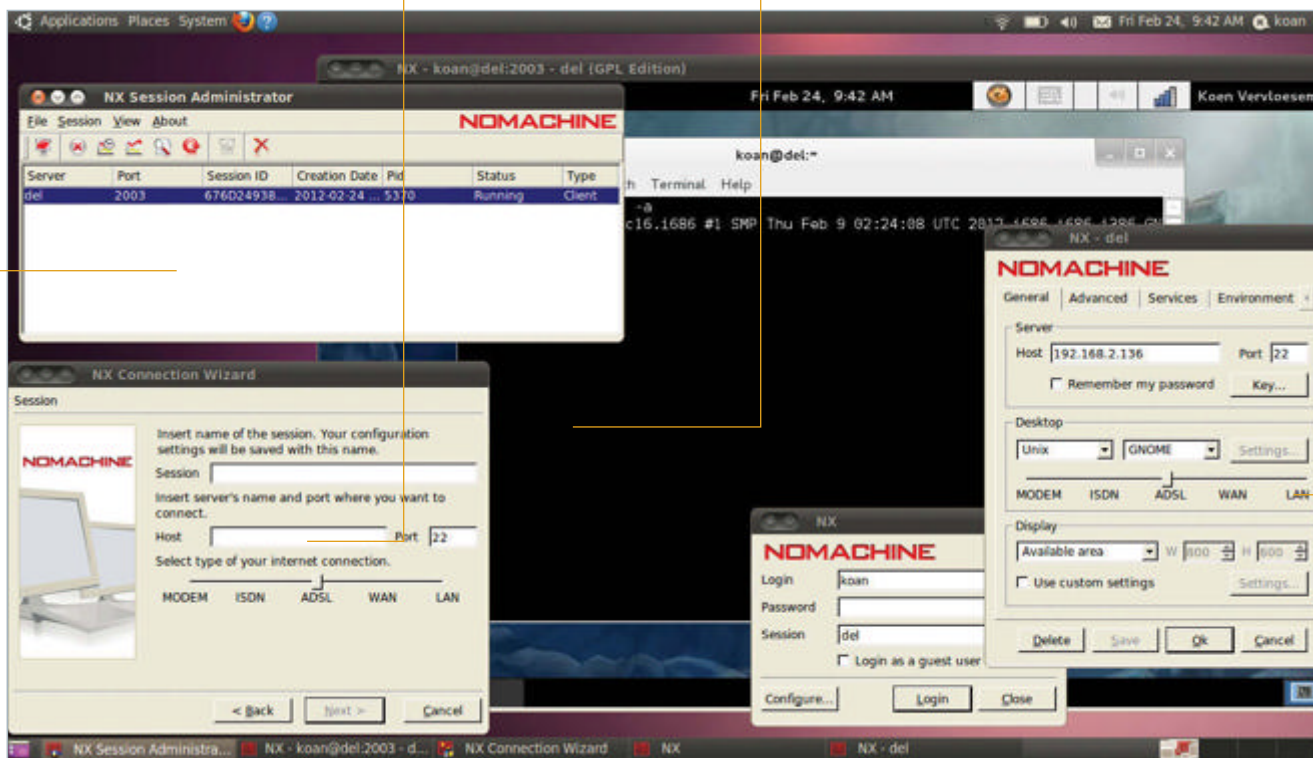


With the NX Session Administrator you can administer all NX sessions and look at the logs for troubleshooting

The NX Connection Wizard is an easy-to-use way to set up a new NX connection

Connect to a desktop on a remote machine, including the option to suspend and resume desktop sessions

The NoMachine NX client lets you configure a lot of settings for your NX sessions



Better remote desktop with NX

If you don't like the rigidity of VNC, RDP and other remote desktop protocols, take a look at NX – it even works over the internet

Resources

FreeNX <http://freenx.berlios.de/>

NoMachine NX <http://www.nomachine.com/>

Advisor

Koen Vervloesem has been writing about free and open source software, and IT in general, since 2000. He has master's degrees in computer science and philosophy and can be reached at koen@vervloesem.eu



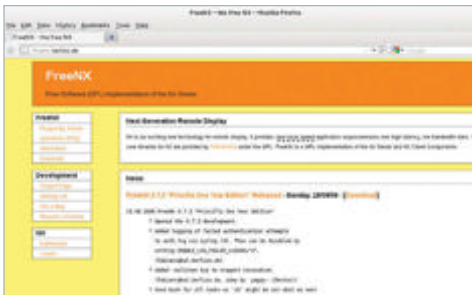
If you need to connect to a remote desktop, chances are you're using VNC or RDP, but another interesting solution is NX. It uses SSH as its

transport protocol, which gives it encryption and authentication for free. This also makes user management simple: every user that can log into a server with SSH is able to connect to a remote desktop with NX.

Another interesting feature of the NX protocol are the sessions: you can not only log into a graphical desktop remotely, but you can also leave the session without terminating it. So your programs keep running on your desktop and


you can reattach to the session later, even from another computer. It's even possible to leave multiple sessions running and choose which session you are reconnecting to.

The NX protocol is developed by NoMachine, who offers clients and server software for Linux, Windows, Mac OS X and Solaris. There are also some independent open source projects, both client and server. For instance there's the FreeNX server and Google's Neatx server, and on the client side there's the open source programs QtNX and OpenNX. Some general remote desktop programs also support the NX protocol: Remmina for instance.



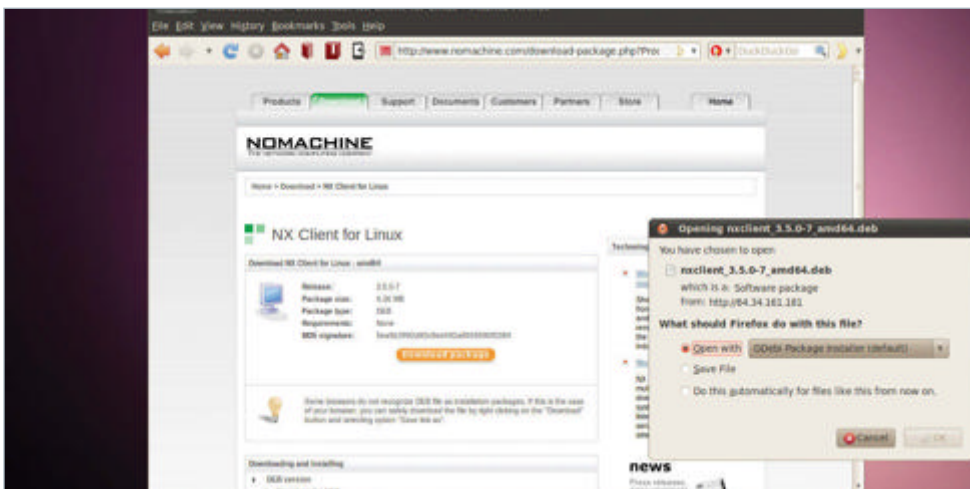
01 FreeNX server

On the server side, we'll install FreeNX, as it's the most mature of the open source NX servers. On our Fedora machine, we installed it with 'yum install freenx-server'. You don't have to start a FreeNX service, but make sure that SSH is running.



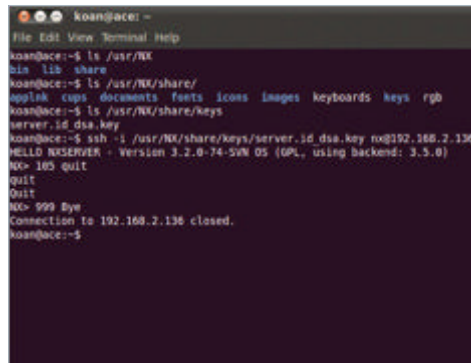
02 Initial setup

After installing the FreeNX server, it needs an initial configuration. Fortunately, FreeNX comes with a setup program that does this for you. Enter nxsetup as root, which creates the nx user, an initial configuration and the default keypair of NoMachine. This is the minimal server setup you need to test NX.



03 Install the NoMachine client

Now download NoMachine's official client for Linux on your client machine. You can find rpm, deb and tar.gz versions for



04 Test

Try to SSH into the machine as the nx user to test if the server setup is done right. Using the default NoMachine keypair, try this on the client: 'ssh -i /usr/NX/share/keys/server.id_dsa.key nx@NXSERVERIP'. You should get the response 'HELLO NXSERVER'. Enter 'quit' to close the connection.

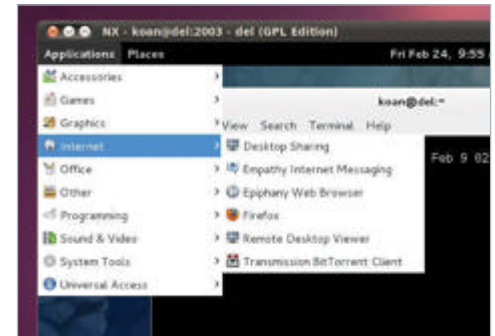


05 Client configuration

Start the NoMachine client (in /usr/NX/bin/nxclient). Enter a name for your session and configure the settings. The most important one is the hostname of the machine you're running the FreeNX server on. Try 'testdrive.millinux.com' to test the client on a demo server.

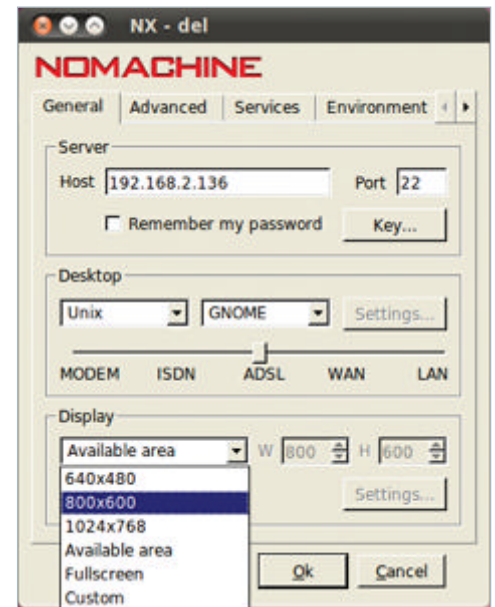
06 Login

Enter your login and password or check 'Login as a guest user' for the demo server. Although NX uses SSH with the nx user for the connection, you should enter your username and corresponding password on the server here, which you normally use for local logins as well.



07 Remote desktop

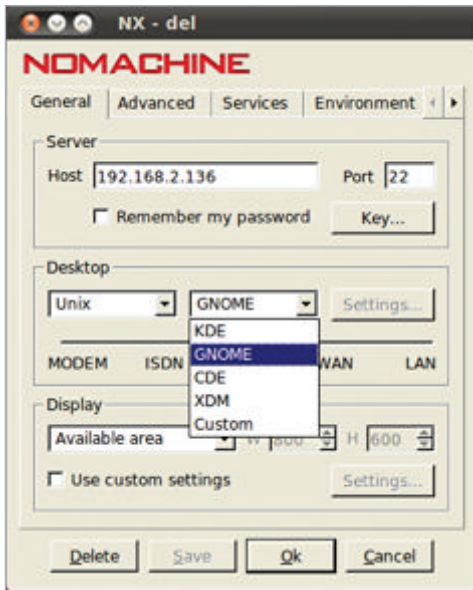
Click on Login, after which the NX client starts up the remote desktop connection. After a few seconds you'll see the desktop from the remote computer appearing on your local screen and you can start working right away.



08 Display settings

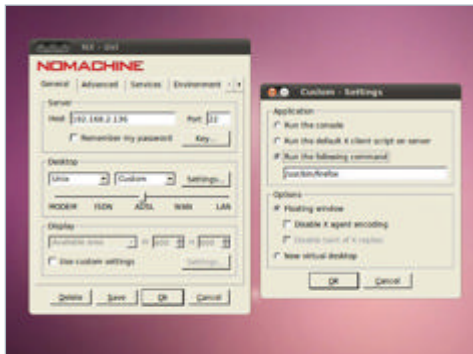
However, you will probably want to change the default display settings. Start up a new NoMachine client and click on 'Configure...' for your session. By default, the Display setting will have 'Available area' as its value, but you can change the preferred display size for the session to a specific resolution or even to full-screen.

i386 and x86_64 in the Download section of www.nomachine.com. Download the right package for your system and install it with your package manager.



09 Desktop environment

By default, the NoMachine client configures its connection for UNIX as the display type, and KDE as the desktop environment. Change this setting if you're connecting to another desktop environment, or choose Custom and specify in the settings to run the default X client script.



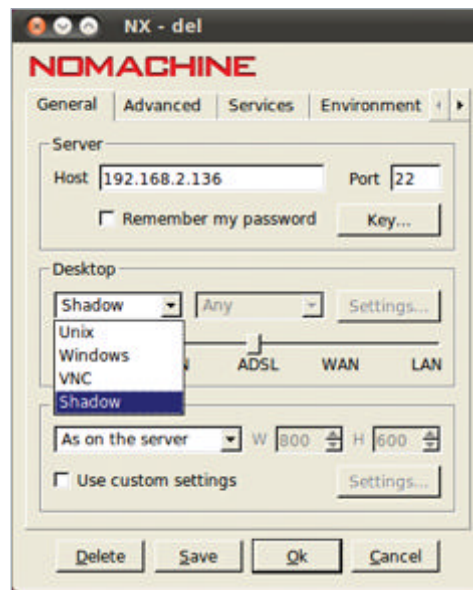
10 A single application

If you choose Custom in the Desktop settings, you can also specify to run a particular application, which then runs remotely in its own window without the complete remote desktop around it. This is great if you want to be able to suspend and resume a remote application.

“NX offers better performance for remote desktops than X”

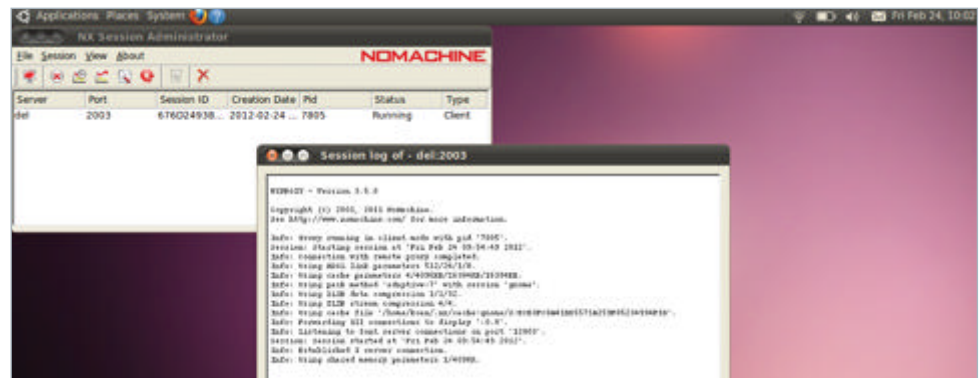
11 Better performance

The NX protocol makes use of various techniques to offer better performance for remote desktops than X (for instance with ssh -X). You can guide the use of these techniques by selecting your network type in the slider in the General tab. If you choose ADSL (the default), NX efficiently compresses the X traffic and defers screen updates with lazy encoding.



12 Desktop sharing

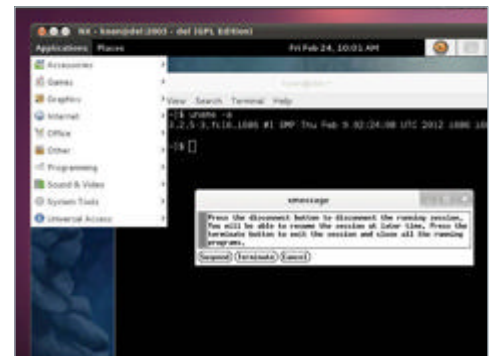
If you choose Shadow as the display type in the General tab, you can remotely connect to the X display running on the server. This is a form of desktop sharing: both on the NX client and on the local display of the server you interact with the same X session.



15 Administer your sessions

If you face some issues with NX sessions, start the NX Session Administrator (/usr/NX/bin/nxclient -admin), which shows you all current sessions, their status, process ID, session ID and creation date. You can disconnect a recalcitrant

“You can specify to run a particular application, which runs remotely in its own window”



13 Suspend

When you close the window of an NX session, you get the choice between suspending and terminating the session. If you terminate it, the session ends and all running programs are closed. If you suspend it, all programs keep running and you can resume the session later.

14 Resume

If there's a session running in the background and you connect to the NX server again, it automatically resumes your suspended session. This even works when connecting from another machine, so it's ideal when you need to continue your work on another computer.


```

koan@del:~
File Edit View Search Terminal Help
# This directives controls the verbosity of the server-wide log.
# 0: No Logging
# 1: Errors
# 2: Warnings
# 3: Important information
# 4: Server - Client communication
# 5: Information
# 6: Debugging information
# 7: stderr of some applications
#NX_LOG_LEVEL=0

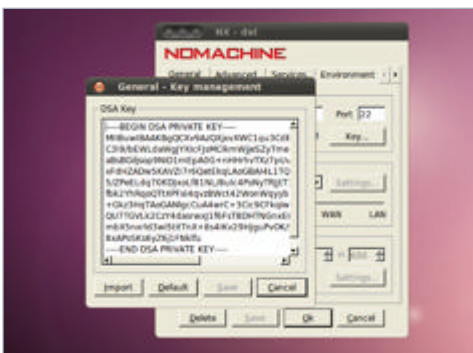
# By setting this to 0 the nxserver might be a bit faster, but passwords can
found in the log files.
#NX_LOG_SECURE=1

# Before turning logging on, please make sure that NX_LOGFILE is
# writeable for the "nx" user
#NX_LOGFILE=/var/log/nx/nxserver.log

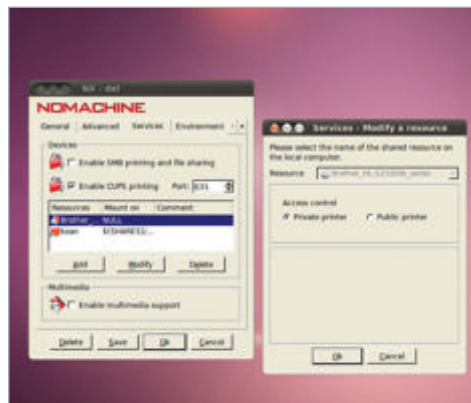
# This directive controls if the temporary session directory
# ($HOME/.nx/C-<hostname>-<display>-<session_id>) should be kept after a
# session has ended. A successfully terminated session will be saved as
# T-C-<hostname>-<display>-<session_id> while a failed session will be saved
# as F-C-<hostname>-<display>-<session_id>.
# The default is to cleanup the directories.
#SESSION_LOG_CLEAN=1
279 1

```

16 Log Unfortunately, FreeNX disables logging by default, which complicates matters during troubleshooting. So if you face some problems, edit `/etc/nxserver/node.conf` and set `NX_LOG_LEVEL=3` to log important information and `SESSION_LOG_CLEAN=0` to keep log files after a session has ended.



17 Custom keys Until now we used the default NoMachine keypair for testing purposes. For security, enter `nxsetup` as root but now answer Yes to the question to use your own custom keypair. Then copy the file `/var/lib/nxserver/home/.ssh/client.id_dsa` key to your client machine and import it into the NoMachine client by choosing 'Key...' in the General tab, Import and then Save.



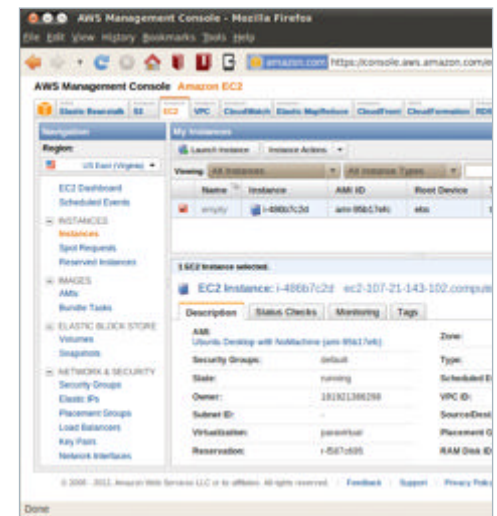
18 Printer sharing If you want to be able to print to a printer that is connected to your local machine in programs that you are running in the NX session on the server, this is possible: just `chmod 755 /usr/lib64/cups/backend/ipp`, check 'Enable CUPS printing' in the Services tab of your NX session's settings and add your printer.

19 File sharing In the same way, you can share files (and printers) from your local disk using SMB. Install Samba, check 'Enable SMB printing and sharing' and click Add. Select the name of your home directory or printer and supply the accompanying username and password.

20 For Windows clients NoMachine also offers an NX client for Windows, which is great if you want to connect to an NX session on your Linux machine when you're working on your Windows machine from time to time.

21 QtNX If you don't like the NoMachine client, try QtNX - which is, for instance, available in Ubuntu's package repository. It doesn't have as many settings as NoMachine's client and it doesn't support some advanced features, but it does the job and it has a log window for troubleshooting.

22 Remmina Another client that supports NX is the remote desktop client Remmina, at least since version 0.8. The advantage of this is that you can use the same client for your remote desktop connections with various protocols, including VNC and RDP.

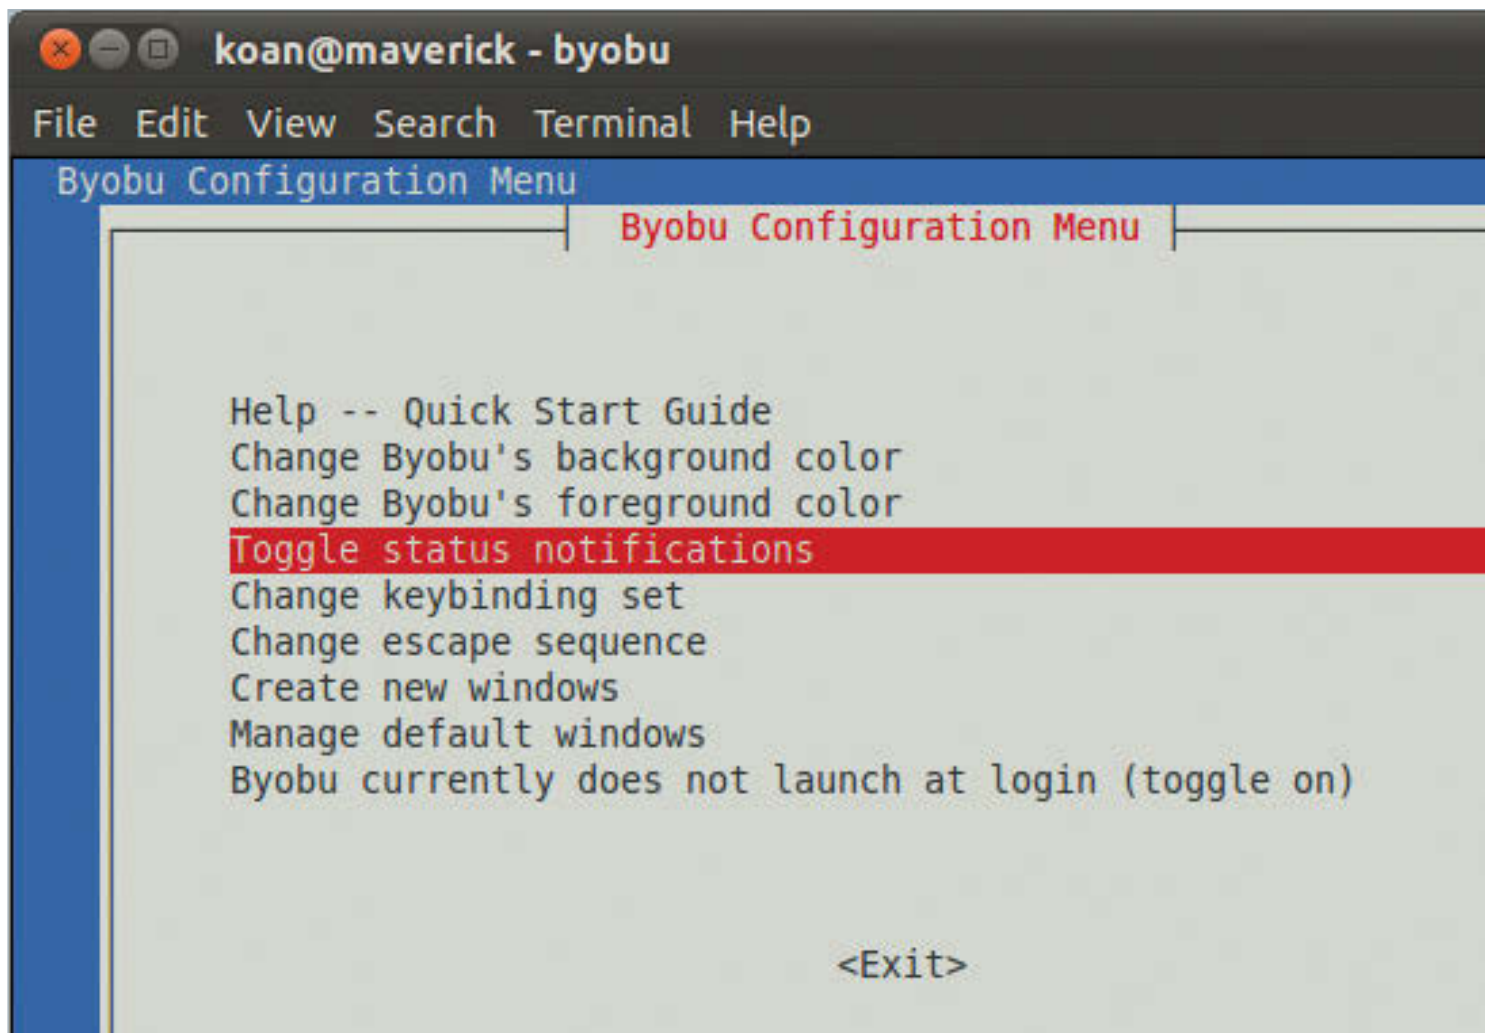


23 Ubuntu desktop in the cloud To show that the NX protocol is really working well over internet, try connecting to an Ubuntu desktop in the cloud. You can find Ubuntu Desktop images for Amazon EC2 with Google's NX server Neatx or FreeNX installed. Fire up an instance in the AWS Management Console.

24 Cloud session Now log into your EC2 instance with `ssh`, add a new user with `sudo adduser <username>` and give it a password. After this, you can connect to your EC2 instance by entering its public DNS address as the hostname in your NX client. Now experience the cloud magic...

Master your terminal sessions

Command-line programs are not glued to their terminals. Learn the tricks on how to detach and reattach them



■ Changing Byobu's configuration is easy

Advisor

Koen Vervloesem has been writing about free and open source software, and IT in general, since 2000. He has master's degrees in computer science and philosophy and can be reached at koen@vervloesem.eu.



Even in this world of modern graphical desktops where everyone is excited about KDE 4, GNOME 3 and Unity, much work still happens inside terminals – not only graphical terminals like GNOME Terminal, but also in virtual consoles outside X. If you're a command-line aficionado, you'll probably juggle with multiple terminals simultaneously to run all your commands. Fortunately, some programs

exist to help you manage working with multiple terminal sessions. In this tutorial we'll show you some handy tips, and at the end you'll be able to detach from sessions, reattach to sessions from another terminal or even from another computer. And you'll see how you can use these possibilities for staying online with a command-line IRC client or for session sharing where you do tech support and the person you're helping can see what commands you're typing.

```
koan@maverick: ~
File Edit View Search Terminal Help
Need to get 66,1kB of archives.
After this operation, 217kB of additional disk space will be used.
Get:1 http://be.archive.ubuntu.com/ubuntu/ maverick/universe htop amd64 0.8.
buntu1 [66,1kB]
Fetched 66,1kB in 0s (261kB/s)
Selecting previously deselected package htop.
(Reading database ... 130902 files and directories currently installed.)
Unpacking htop (from ../htop_0.8.3-lubuntu1_amd64.deb) ...
Processing triggers for desktop-file-utils ...
Processing triggers for python-gmenu ...
Rebuilding /usr/share/applications/desktop.nl_BE.utf8.cache...
Processing triggers for man-db ...
Processing triggers for python-support ...
Setting up htop (0.8.3-lubuntu1) ...
koan@maverick:~$ htop

[3]+  Stopped                  htop
koan@maverick:~$ jobs
[1]  Stopped                  man man
[2]- Stopped                  top
[3]+ Stopped                  htop
```

“Every program you start on a terminal is associated with this terminal session”

Job control

Before we start with the fancy stuff about detaching and reattaching, let's get back to the basics. Every program you start on a terminal is associated with this terminal session. If it's an interactive program, it prints its output to this terminal and asks you for input on this same terminal, effectively preventing you from doing any other work in this terminal. However, by pressing Ctrl+Z you can pause the program, which gives you back control of the terminal and lets you do other stuff. Note that the program is then suspended, so it stops doing anything! If the program is an editor or a graphical program you started from the terminal, you probably don't mind, but if it's a program downloading a big file, that's probably not what you want. To remedy this, type 'bg', which resumes the program in the background. If at any time you want to resume the program in the foreground and be able to read its output and give it input, type 'fg'. And if you know beforehand that you want to run a program in the background, instead of starting it, type Ctrl+Z and then bg and you can just append an ampersand (&) to the command.

If you're used to the functionality of Ctrl+Z, fg and bg, there will probably be some time when you want to use it for more than one program simultaneously in the same terminal. This is possible, but then you need the concept of 'job IDs'. Have you noticed that every time you suspend a program, the terminal shows a number between

brackets, like [1] or [2], followed by the command you suspended? This number is associated with each program you suspend. You can show a list of all suspended commands by typing 'jobs'. Now when you want to resume any of these jobs, you just have to add the job ID to the bg or fg command – for example 'fg 2'.

Don't terminate it

You also have to realise that the fact that a program is suspended or is running in the background doesn't mean that it isn't associated any more with the terminal. You can check this with the command pstree, which shows the dependencies of all running commands. Even when you have suspended a command, it is still shown as a child process of the shell (for instance Bash or zsh) in your terminal. What this also means is that when you close the terminal, for instance by logging out of the shell, the command will be terminated, because all child processes are sent the SIGHUP signal. Luckily, your shell will warn you that you have active tasks when you try to log out or exit your terminal, but a second logout or exit command kills the processes ruthlessly.

There are two ways to remedy this behaviour. When you have started a command and you want to log out without terminating the command, suspend it, resume it in the background with bg, and then run disown followed by % and the job ID – for instance 'disown -h %1' (or just without the %1 parameter if you have only one job). The disown command

removes the job from the table of active jobs for the current terminal session, and connects it as a child process directly to the init process, which only terminates when you shut down your computer. As a result of this, you can safely close the terminal without terminating the command. You can also check this with the pstree command after typing the disown command. If you know beforehand that you want to keep a program running after logging out, launch it preceded with the nohup program. Nohup is to disown what the ampersand is to Ctrl+Z and bg.

Note that disown and nohup disconnect the input and output of the program completely from the terminal, so the program's output goes nowhere, and it doesn't have any way to ask you for input. So this is not suitable for interactive programs. And even after the program has done its work and exits, there's no way for you to know whether the program has exited successfully or with an error. But it can be interesting for that curl or wget command, to download a big ISO file of your favourite Linux distro, that you can let run without any user intervention.

Breaking out of the session

Job control limits you to one terminal: jobs are numbered starting from 1 for each terminal, so you can't just suspend a program in one terminal and resume it in another one. Moreover, if you have disconnected a program from your terminal with disown or nohup, there's no way to interact with it

any more, even in the same terminal: it won't appear in the list of active jobs and you can't resume it in the foreground with fg.

A more flexible solution is reptyr, which takes an existing running program and attaches it to a new terminal. This is really simple because you don't even have to care about job IDs. The only number you have to know is the process ID, which you can ask by typing 'pgrep <program>' or 'pidof <program>'. Now go to another terminal and enter 'reptyr <pid>', with <pid> the process ID of your program. After this, the program is detached completely from the original terminal and appears in the current terminal, where you can see its output and enter your input. This even works if you started a program locally and then went away, after which you can log into your computer with SSH and then 're-pty' the program to your SSH terminal session.

Reptyr is Linux-only because it uses ptrace and is highly dependent on Linux's system calls, so if you're a BSD user you're out of luck. On Ubuntu 10.10 and higher, ptrace is disabled by default for security reasons, but you can enable it permanently by editing the file /etc/sysctl.d/10-pttrace.conf and putting the following line in it:

```
kernel.yama.ptrace_scope = 0
```

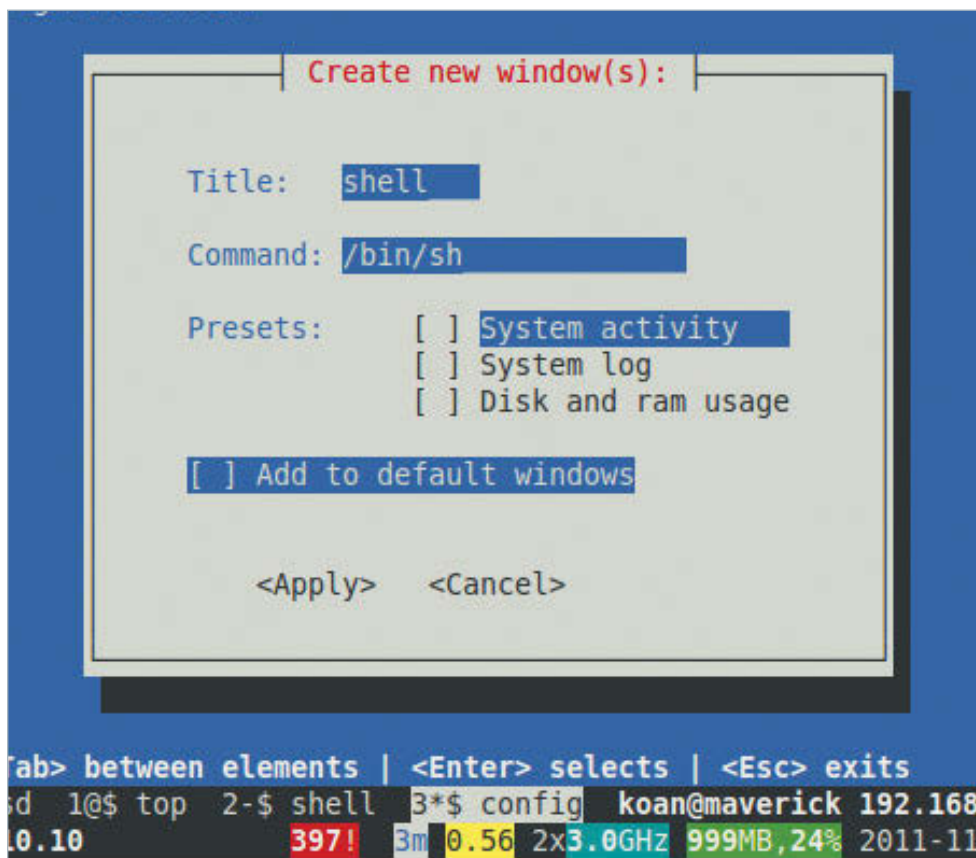
If you want to test this without rebooting, enter the following command:

```
echo 0 | sudo tee /proc/sys/kernel/yama/ptrace_scope
```

Screen

If you know beforehand that you want to detach a program from its terminal and reattach it to another one, you can use the popular program, GNU Screen. But Screen is much more: it's a terminal multiplexer. It gives you access to multiple separate terminal sessions inside a single terminal session. This makes Screen essentially a sort of window manager on the console. You can run various command-line programs inside a Screen session, detach from the session, log out, leave all your programs running, and reattach to Screen later (from another terminal or even from another computer using SSH) to continue.

If you want to run some programs in Screen, just start a new Screen instance with the command 'screen', which starts your default shell in a first window. If, instead, you directly want to start a specific program inside Screen, just enter 'screen' appended by the command (you can add other programs to the same Screen session later), after which the command is run inside a Screen window. When you're inside a Screen session, the Screen program recognises some key combinations, all preceded by Ctrl+A. With 'Ctrl+A C', you create a new window and switch to it (it opens your default shell); with 'Ctrl+A N' you switch to the next window; with



■ Create a new window in Byobu

'Ctrl+A P' you switch to the previous window. 'Ctrl+A "' gives a list of the available windows and lets you select one to switch to. You can switch to a specific window number with 'Ctrl+A <number>'. And when you want to log out, you can just leave your programs running by 'Ctrl+A D', which detaches Screen from your terminal.

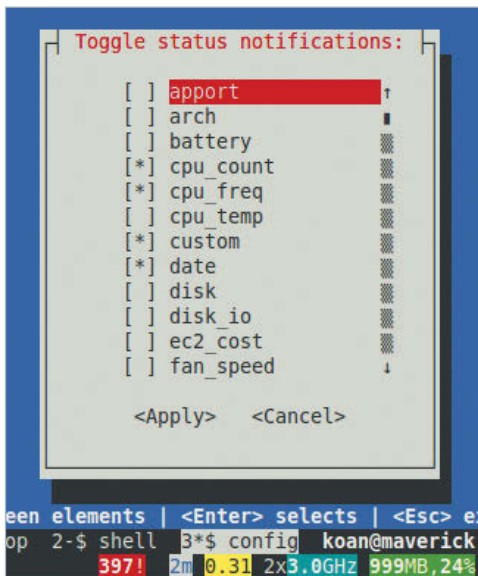
Now when you want to go back to your programs, just run 'screen -r' on a terminal, after which Screen starts up and attaches your running Screen session to your terminal. It's also possible to run multiple Screen instances, which you can list with the command 'screen -ls'. By default, your sessions are named pid.tty.host, but you can change the tty.host part to a more sensible name if you start the Screen instance with the option '-S <session name>'. You can attach to a specific Screen session with 'screen -r <session name>'. However, 'screen -r' only attaches to a detached session, so if you have left your

Screen session attached and want to attach to it from another computer using SSH, you first have to detach it, which is possible with 'screen -d <session name>'. You can also reattach a session and if necessary detach it first with one command: 'screen -d -r <session name>'.

Screen on speed

'Ctrl+A ?' shows you a concise help table with some shortcut keys, and the man page gives you the full details of Screen's endless possibilities. For instance, you can specify a lot of Screen's behaviour in a file, .screenrc. However, Screen is difficult to tweak, and therefore some Ubuntu Server developers decided to create some out-of-the-box Screen profiles that make use of Screen's advanced features. This program was first called screen-profiles but is now known as Byobu. If it's installed (which it is by default in recent Ubuntu releases), just run 'byobu', after which you'll see a status bar

“All in all, Screen and Byobu can make your life in terminals much easier and much more productive”



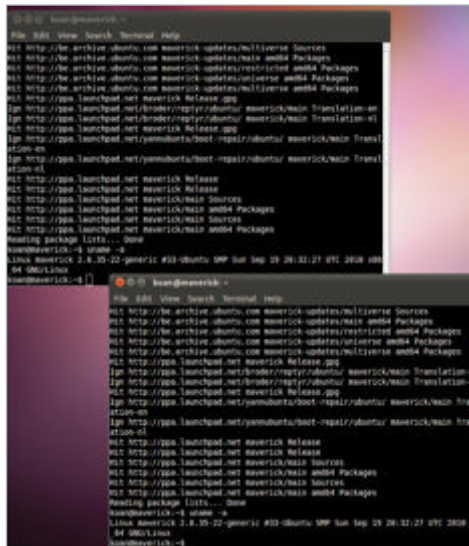
at the bottom of your Screen sessions, with some information such as Ubuntu's release number, the number of available package updates, the system load, the number of running processes, memory use, the date and time, and so on. Above this status bar you'll see a list of the opened Screen windows.

But Byobu goes much further: by pressing F9 you get extensive configuration possibilities. For instance, in the menu 'Toggle status notifications' you can choose what appears in the status bar. Apart from the things we already mentioned, you can add information like uptime, battery status, upload and download speed, the Wi-Fi signal quality and much more. Moreover, you can choose to automatically start up Byobu when you log in, which can be handy when you forget this, and you can change the key bindings from the menu. You can also set some windows that are created by default when you start a new Byobu session.

Byobu has some easy key bindings: instead of Ctrl+A and another key, you can do some tasks with a press on a function key. By default F12 locks the terminal, F9 opens Byobu's options, F8 lets you enter a name for the current window, and F7 enters 'copy/scrollback mode', which opens a vi-like editor that allows you to copy text from the current window and its history into the paste buffer. F6 detaches from a session, F4 moves to the next and F3 to the previous window, and F2 creates a new window. Shift+F2 splits the screen horizontally, and Ctrl+F2 splits it vertically.


All in all, Screen and Byobu can make your life in terminals much easier and much more productive. If you prefer the BSD way, maybe tmux is something for you, and the Ubuntu developers are working on tmux support in Byobu. Moreover, when you forgot to use Screen or tmux when you really had to, reptyr comes to the rescue.

Session sharing




1. Multiple displays

You can use 'screen -x' to attach to a not-detached window, and this leaves both windows visible, which you can use as a poor man's session sharing mechanism for tech support: log into the person's computer with his account and password and create a named Screen session. Then ask him to type into a terminal window 'screen -x <session name>'. Now you both see the same Screen session, and he can see what you type to help him.



2. Multi-user mode

Unfortunately the previous tip only works when you both log into the same user account. However, Screen has a multi-user mode. Create a Screen session and enable multi-user mode with 'Ctrl+A :multiuser on'. Then add the other user to the access control list with 'Ctrl+A :acladd <other user>'. Optionally only give him read access with 'Ctrl+A :aclchg <other user> +rx'. And finally, ask the other user to attach to your Screen session with 'screen -x <your user/session name>'.
 koan@maverick:~\$ screen -x koan/helpdesk
 koan@maverick:~\$ screen -x koan/helpdesk
 koan@maverick:~\$



3. SUID root

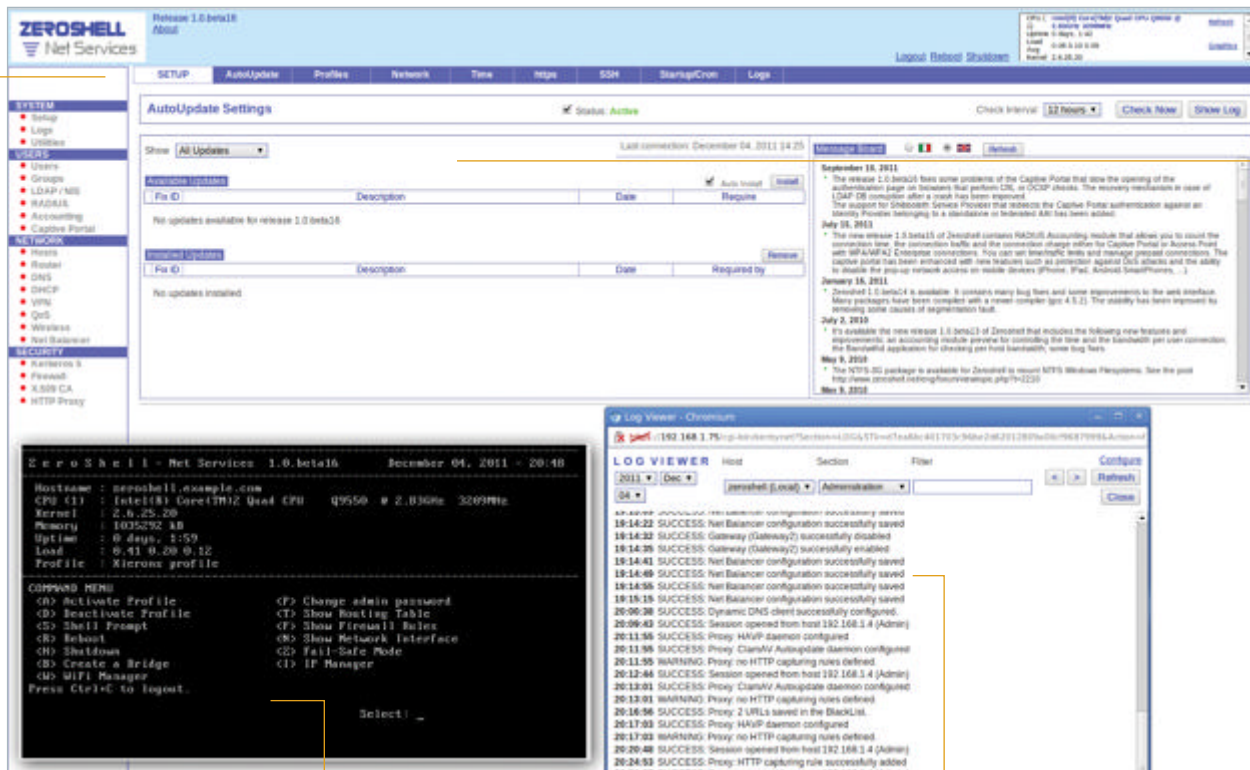
If the other user gets an error message about multi-user support, you have to enable SUID root for the Screen program. This is done by typing 'sudo chmod +s \$(which screen)'. Also, delete lingering Screen files with 'sudo rm -fr /var/run/screen/*' and then try your session sharing again. Note that making the Screen command SUID root is a potential security issue...

Tmux

Screen is not the only terminal multiplexer in town. If you don't like Screen, have a try at tmux

A lesser known but equally powerful terminal multiplexer is tmux. Because of its BSD licence, it's more popular in the BSD world, and it's part of OpenBSD's default installation since version 4.6. Tmux has some advantages over Screen. For instance, its configuration file is in general much clearer, as well as its documentation in the man page. Moreover, in contrast to Screen, tmux commands are the same whether you enter them in the configuration file or interactively in tmux itself. Another interesting feature is that tmux has a flexible client/server model: windows are independent of each other, can be attached to multiple sessions and viewed in multiple clients simultaneously, and they can even be moved from one session to another one on the same tmux server.

If you want to get to know tmux, a good place to start is OpenBSD's documentation (www.openbsd.org/faq/faq7.html#tmux) and the tmux man page. We'll list you some shortcuts to get going: by default tmux uses the prefix Ctrl+B instead of Screen's Ctrl+A, followed by a key for a command, such as C to create a new window, N and P to go to the next/previous window, D to detach your session, and so on. Attaching to a detached session is possible with the command 'tmux attach'.



While not pretty, the web interface is straightforward and very powerful. Features are accessed through the top and side menus

The host machine just uses a console interface, which is mainly only needed during initial setup; the rest is done through the web interface on another machine

Zeroshell offers extensive logging options, which can be very granular, allowing for detailed network use analysis

Although still labelled as a beta, Zeroshell is very stable and used in many production environments; updates are frequent and can be automatically installed

Manage your network with Zeroshell

Prioritise and administer your network traffic using this specialist Linux router distro

Advisor

Kieron Howard First introduced to Linux eight years ago when installing Debian on his Xbox, Kieron has continued to enjoy putting Linux on devices such as iPods, PS3s and various phones



Resources

Zeroshell: Available as a ISO, VM image or USB image: <http://www.zeroshell.net/eng/download/>

Host machine requirements: Pentium 233MHz, 96MB RAM, CD drive and compatible NIC

Extra NIC if you use a standalone modem

HDD or USB stick to store settings
Switch/wireless access point

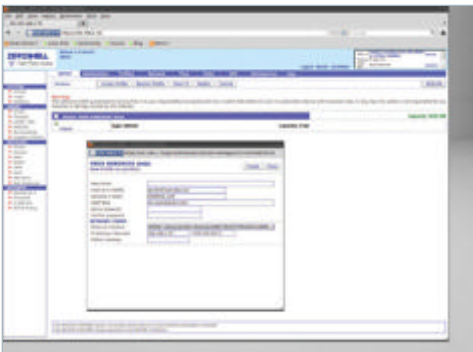


It's now uncommon for even small offices to not have a LAN of some sort, and many of us have home networks with which we share our internet connection. Most home routers are fairly basic and are designed for ease of use and low cost, as most users never even log into it after the initial setup, but there is a lot you can do to improve your network's performance by managing traffic. A fantastic, free, open


source solution is Zeroshell. The distro offers all the standard router options such as firewall, DHCP server, and RADIUS server for WPA/WPA2 authentication, but also a host of features such as: QoS (quality of service) and traffic shaping management, a dynamic DNS client that's great for those who use an ISP that assigns dynamic IPs, the ability to run cron'd scripts and more. We'll take you through the process of setting up the router and some basic traffic shaping.



01 Start up Zeroshell
 Boot into Zeroshell using your chosen method and you will see the main interface. Zeroshell is designed to be managed through a web interface, but some settings can be modified in the console on the host, such as IP and username/password.

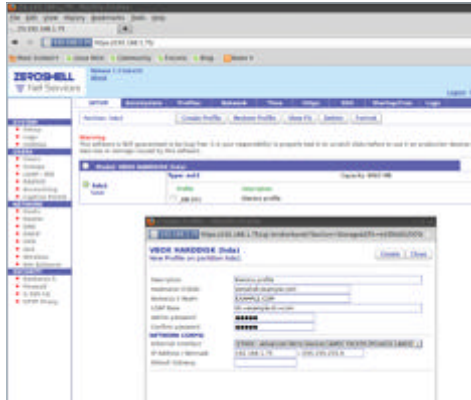


02 Login
 First we need to set up a profile to store our settings. Log into the web interface by entering the IP listed in the console and the username and password. Your browser may throw a certificate warning, but you can add an exception for this.



03 Create working partition
 We need a place to persist our settings. In the menu, choose Profiles and then the drive you want to save to. Then select New Partition. Choose a size and file system (ext3 is fine), then click Create. After a few seconds, you should see the partition listed.

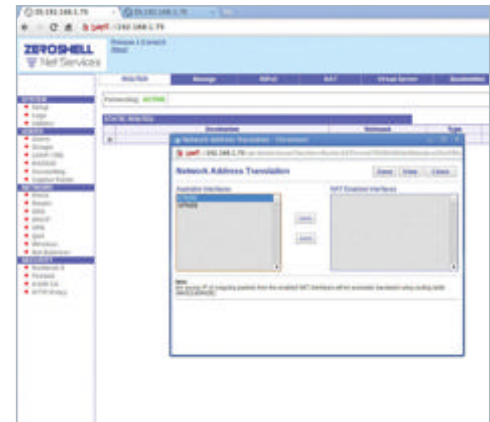
04 Create profile
 Click the radio button next to the partition we just made, then Create Profile. Enter a description and a password and click Create. Then check the radio next to the profile and click Activate. Zeroshell will then reboot, so you'll need to log in again.



05 Add gateway
 Depending on your network setup, you may need to define the gateway so that Zeroshell has internet access. In the Network menu, choose Gateway and enter the IP. To test this is working, head back to the homepage and check the news has updated.




06 Start DHCP server
 Now let's set up a DHCP server to make our install start behaving like a real router. In the left menu under Network, choose DHCP. Next to Subnet select 'new'. Choose your interface and then save. Add in some IP ranges (eg 192.168.1.10 – 192.168.1.90) and save once more.



07 Configure NAT
 Network address translation (NAT) is another common requirement, and is easily configured in Networking>Routing and then in the top menu, NAT. Add the interfaces that you want to use NAT and click Save.



08 Enable QoS
 Now if you connect a new piece of hardware to your network, it should get assigned an IP using DHCP and be able to connect to the internet. Let's now start doing some more interesting work by enabling some traffic shaping. Under Network, click QoS.

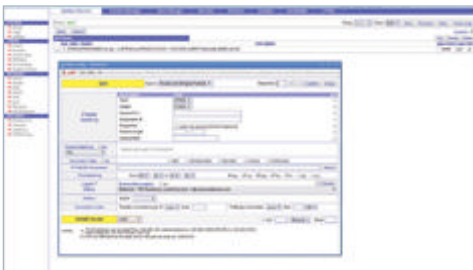


09 QoS classes
 We can create multiple classes for each priority of traffic we want. Click Class Manager and New. Choose High priority and set some Max and guaranteed values. You can also add in a lower priority class and then add some values for that.



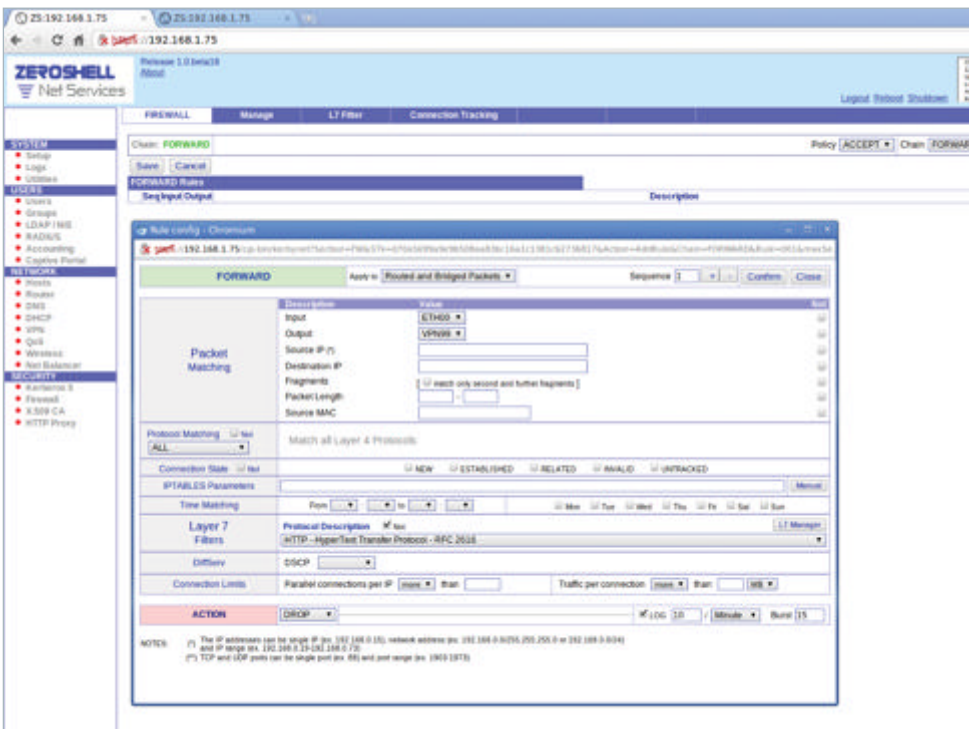
10 Associate class with rules

Let's now assign a protocol to our 'high' class. Click on Classifier and then 'add'. Choose Adaptors, then in this example we will say that FTP traffic gets our 'high' class so will take priority over the rest. Make sure to save changes.



11 Throttle P2P during the day

Create another classification; this time we will choose our 'low' class and match to BitTorrent traffic. Then set the Time Matching to business hours in during weekdays, ensuring our line stays unclogged when it's most needed.



12 Install Dansguardian content filter


Dansguardian is an open source content filter – handy if you have young children. **In the console, type 'S' to get a shell, then: cd /Database**

<http://www.zeroshell.net/listing/BA13-DansGuardian-Web-Content-Filter-2.8.0.6-1.0.beta13.tar.bz2>

`tar xvfj BA13-DansGuardian-Web-Content-Filter-2.8.0.6-1.0.beta13.tar.bz2`
`cd BA13`


`./install.sh`

Full details on configuring can be found at <http://dansguardian.org/>, but it'll now automatically filter adult content on your network.




13 Setting up the basic net balancing with failover

Setting up the router to automatically use another connection if one of them dies can be very handy in large networks. To do this, in the Netbalancer menu add in all the gateways on your network, giving them a weight value to prioritise.



14 Adding in failover monitor settings


Zeroshell will check up to three address to make sure that they are reachable, to manage its failover settings. Enter an IP with a high uptime like Google.com and then Enable the monitoring and IP.



15 Setting up a firewall

Zeroshell has an easy-to-use firewall manager for blocking ports and traffic that might not be desirable. As an example, apply a strict policy of only accepting HTTP traffic, open the Firewall setting, enter your interfaces, then tick the 'not' box on the Layer 7 filter and choose HTTP. Under the Action part, you can choose to Drop or Reject. Dropping traffic returns no error packet, while rejecting will remove the packet and return a 'Communication administratively prohibited' error.


“In this example we will say that FTP traffic gets priority over the rest”



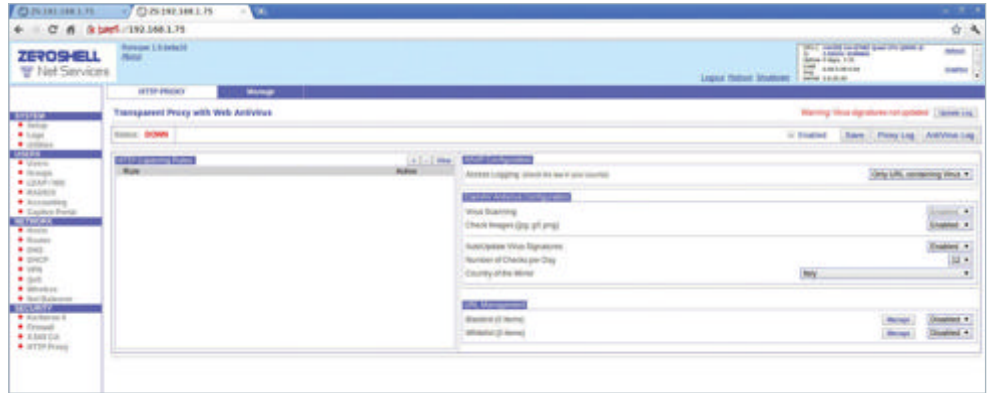
16 Add a dynamic DNS client
 Many home ISPs will assign a different IP every time your connection is reset. This makes it hard to use VNC remotely, for instance. To solve this we can add a dynamic DNS client, which will submit our IP and translate to an easy-to-remember address.




17 Set up DynDNS account
 There are many dynamic DNS providers to be found around the internet, but in this example we shall use www.dyn.com. Set up a new account and make a note of your address, username and password. Basic usage is free, but there is a pro tier as well.



18 Add DynDNS details
 Go to DNS and Dynamic DNS, then New, add in your account details, and then click the Enabled box. You can now access your Zeroshell machine using the easy-to-remember address that will stay static.



19 Add in anti-virus protection
 It's always a good idea to run some anti-virus of sorts on your network, and Zeroshell has the option of checking traffic for unwanted programs using the open source ClamAV. To enable it, under Security click HTTP Proxy.



20 Enable anti-virus
 Firstly, update the virus signatures database by clicking 'enable'. You can check its progress in the update log. Once done, the status should be 'active'. You can adjust the schedule for updates as well as the location. Here it's also possible to log any access to the router.




21 Blacklist / whitelist URLs
 Depending on the nature of your network, you might wish to black- or whitelist some URLs. This is done in the HTTP Proxy page: simply click Manage and then add in the URLs you wish to white- or blacklist.



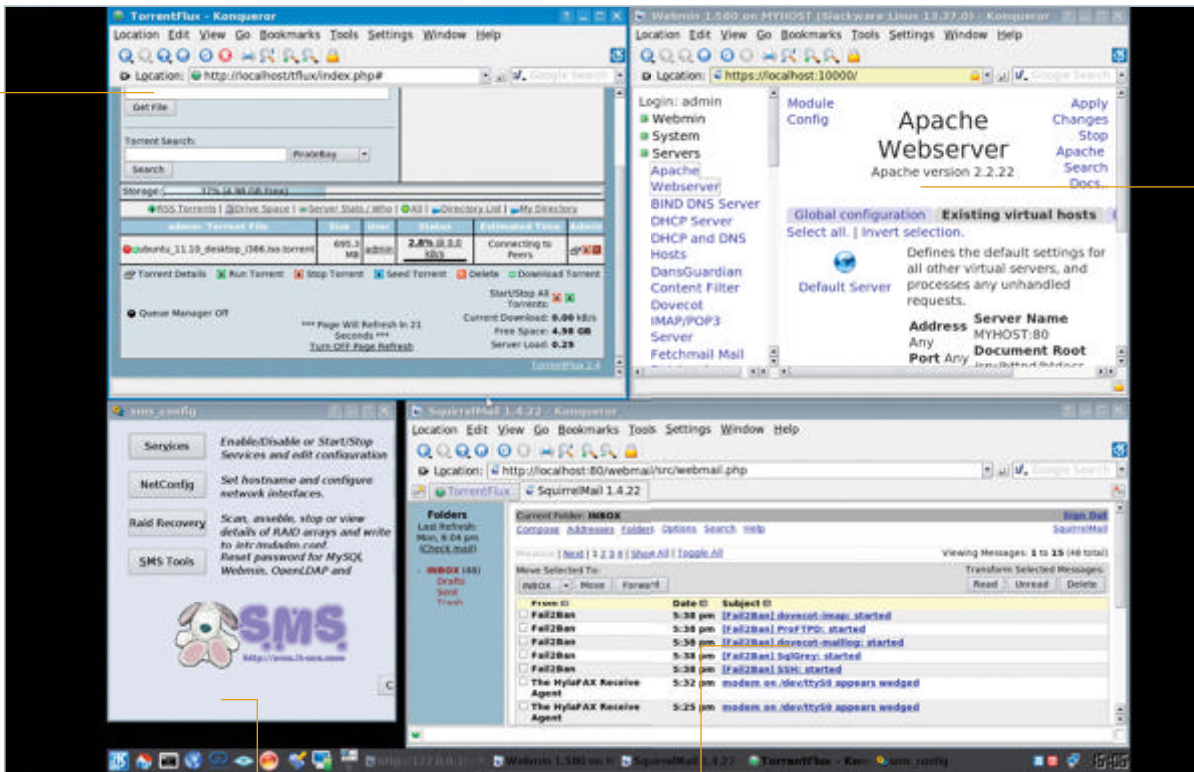
22 Route traffic through transparent proxy
 If you wish to route your internet traffic through a transparent proxy, you can add rules for capturing traffic on port 80. Click on the '+' icon and then add in the source and destination IPs for the proxy.



23 Enable encrypted Wi-Fi
 Secure Wi-Fi is vital for any business, and important at home too. Zeroshell features a powerful RADIUS server for encryption. To set it up, under Users click RADIUS and Enabled. Then click Trusted CAs and export as DER or PEM format.



24 Set up users
 You'll need to copy the certificate to each user's machine. Under Users and X509, export the certificate for each user you wish to give access; check that under Authentication Protocol, RADIUS is ticked. They can then log in using their password.



SMS comes with TorrentFlux, which is a handy web interface for managing the download of torrents

The GUI comes with setup options for many of the services bundled with SMS – great for quick tweaks

Postfix Mail is the email server used in SMS, and SquirrelMail is the web interface; both of these are widely used by many organisations

SMS comes with a full LAMP stack, with Apache configured from the word go. This makes setting up a web server easy in just a few minutes

Set up an all-in-one home & office server

Install and configure SMS – Superb Mini Server – the Swiss army knife of the server world

Advisor
Kieron Howard First introduced to Linux eight years ago when installing Debian on his Xbox, web developer Kieron has continued to enjoy putting Linux on all sorts of devices



Resources

SMS install ISO <http://sms.it-ccs.com/downloads.html>

Host machine or VM (500MHz Pentium 128MB RAM min spec)

Client machine to access server




The area of IT where Linux-based operating systems really dominate is the server market. The stability of the Linux kernel along with the flexibility due to its open source nature have made it a big hit, especially in the web server market. The fact that most distros are free is the icing on the cake. Setting up a server for someone new to Linux can be intimidating, especially with the necessity to use the command line. Luckily there is a distro called SMS (Superb Mini Server) that

aims to do many common tasks 'out of the box'. The inclusion of mini in the name is misleading: it comes with DHCP server, full LAMP stack, OpenSSH, SMB server, two email servers, FTP, VNC, anti-virus, content firewall, Perl, Python and more, most of which is configured to run without too much setup from the user. This tutorial will take you through installation and setup of some of the features of SMS, as well as installing a desktop environment for those not au fait with using a terminal emulator.



01 Choose installation version

SMS is available as a bootable live CD or a regular install CD. To test it out in a non-destructive way, choose the live CD – which can be installed onto USB using LiLi. This guide will be using a full install, though.



02 Boot and install

Put the CD or USB stick into your machine, making sure the device is set as boot in the BIOS. You will then see the first 'boot:' prompt. Hit Return here. The installation process will follow. Enter a keyboard mapping to start with, or press Return to use US.

03 Partition drives

At the next prompt, log in using 'root'. You can list your drives using 'fdisk -l'. Run 'fdisk [your drive]' then 'n' to make a new partition; 'p' to make primary. Enter the start and finish sectors (the default is fine if you have a blank drive).



04 Write changes

You can make a separate swap partition – which can speed things up, especially if its on another drive. To save your changes in fdisk, use 'w'. This will then exit fdisk. Enter 'setup'. You will then see the SMS Setup screen.



05 Select Target

Scroll down to 'Select Target' to choose which drive to install to. Choose the partition we just made and press Return. You can then format it and choose a file system. A safe option is to use the common ext4 type.




06 Choose installation source

You can install SMS from a wide range of locations – we used the CD method, so choose that. The next screen gives you some package options; all except the development packages are selected by default. As we want as much functionality as possible, we will leave everything else ticked.



07 Install options

You can set the install to run in a more verbose way by using 'Terse' – great if you are having issues – or you can tailor the install using 'menu'. If this is your first time using SMS, just choose 'full'.



08 Installation process

The system unpacks to around 2.5GB from the CD-sized ISO, so depending on drive and processor speed, the install may take a little while. It will ask whether you want to make a USB flash drive to boot from, which is optional.

“SMS is available as a live CD or a regular install CD”

Tips & Tricks



09 Select hostname

To help with accessing your SMS server from other services, give your server a hostname and domain that's easier to remember than just an IP, and means the server can use a DHCP-assigned address from your router – if you want, you can still use a static IP.



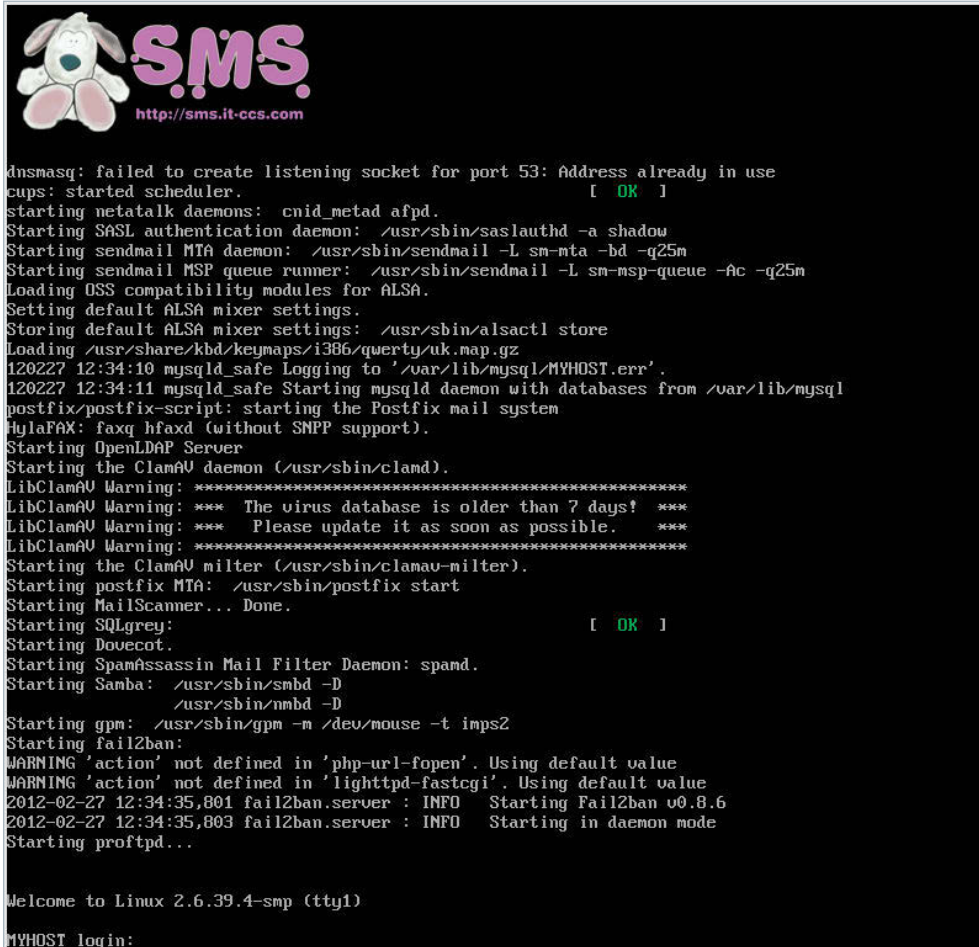
10 Startup services

The next screen tells SMS what services we want to run on boot. For this install, as we are building a 'Swiss army knife', let's enable everything. Scroll down using the cursor and use the space bar to select anything not ticked. Press Enter to confirm.



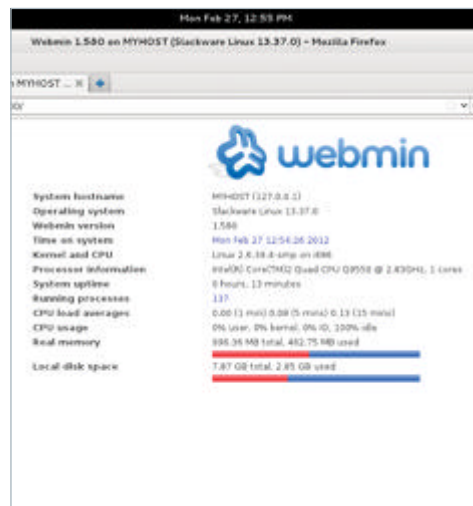
11 Set up desktop environment

You will then need to set up the time and font settings, and after that choose your desktop environment. The default is KDE, but you can install others. As we will be accessing the server using a web interface, you don't need to worry too much about this for now.




12 Reboot

After setting up a new root password, you can then reboot the machine. After a minute or so you will see the login prompt. We can now switch to another machine on our network and enter `https://[MYHOST]:1000`, or by IP in a browser.



13 Admin page

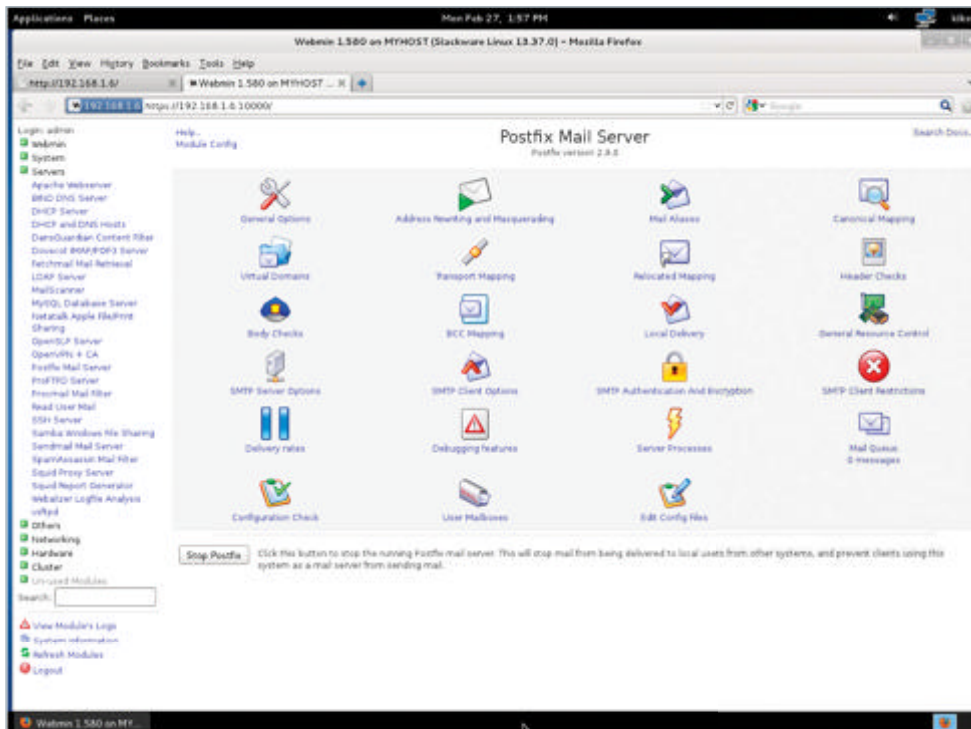
Log into the web administration page using username 'admin' and the 'admin' password. This page is where we control all the services that our server runs. If you click on the 'servers' option in the left menu, you will see the huge amount of features SMS has.



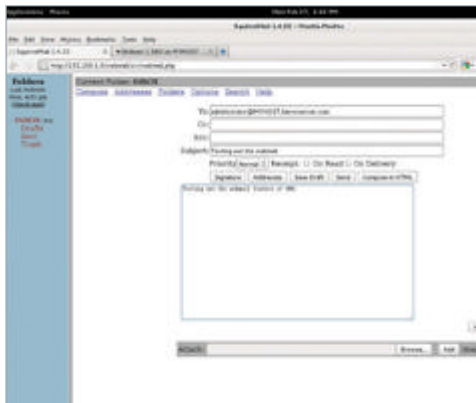
14 Secure the server

We now need to secure the server. Change the web admin password by going to 'WebMin>webminusers' and then click on 'admin' and then 'settings' under password. Change the password. Back on the server, log in and enter 'passwd administrator'.


“If you fancy hooking up a monitor to the server and accessing it directly, you can install KDE”



15 Set up email server
 Back in the web interface, under servers click 'PostFix Mail Server'. The service should work locally without much need to edit settings; to receive external emails you'll need to set up an MX record with your host. Click on 'User MailBoxes' and the 'Administrator'. It's likely you'll have a few messages from some services which are failing. One of these will probably be the Fax server hylafax. To disable it in the terminal, enter 'chmod -x /etc/rc.d/rc.hylafax'. You'll then need to restart the server.



16 Webmail login
 SMS comes with SquirrelMail as its webmail interface. To access it, enter: http://[My Host or IP]/webmail/src/login.php and then enter your administrator username and password to access that mailbox. Other accounts can be set up from the 'Webmin Users'.



17 SSH
 An SSH server is enabled by default. Configuration is done using the SSH server menu. Here you can control access. To SSH into your server remotely from a terminal window, enter 'ssh root@[MY HOST]' – you can replace 'root' with your username.



18 Set up and use TorrentFlux
 Being able to start/stop and manage torrents remotely is a handy feature. You can login to TorrentFlux from http://[MY HOST]/tflux, with username 'admin' and password 'admin'. Click on 'profile' to change these.

19 Add torrent
 To add a torrent, you can either use the web interface to upload one from wherever you are, or just enter the URL for one. You can then click on the green arrow to start off the download.

20 View completed torrents
 Completed torrents are put into Desktop/tflux_downloads/[user]. By default, the completed files are shared using SMB under /var/smb/tflux_downloads/ – so media streamers can mount this for instant access.

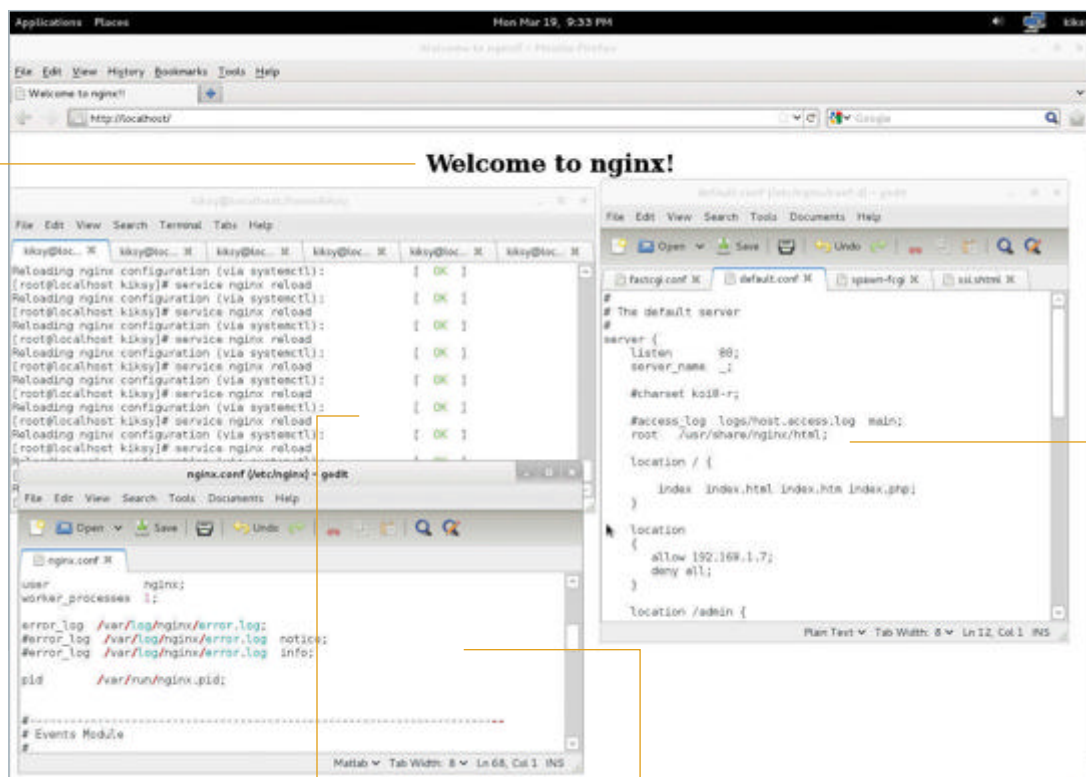


21 Set up GUI
 If you fancy hooking up a monitor to the server and accessing it directly, you can install KDE or another desktop environment. In the server's terminal make sure you're logged in as root and type 'xorgsetup'. It will attempt to auto-detect video and mouse settings.

22 Select colour depth and keyboard
 You can then select the colour depth. Setting a lower value will result in a poorer visual experience, but will save some processing power. After that it will ask whether you want any special keyboard settings. Once the setup process has finished, type 'startx' to load up the GUI.

23 Using KDE
 Having a proper desktop environment means your server can be used as a normal computer, great for browsing the web and other common tasks, while still acting as a server. It also reduces the amount of command-line work needed if you are new to the CLI.

24 Editing settings using GUI
 The GUI comes with 'sms_config_tool', which when clicked allows you to start/stop services as well as quickly edit their .conf files. If your host machine doesn't have lots of RAM or a fast CPU, it's best to drop out of KDE once you've finished editing.



The 'Hello World' for Nginx proves the software is serving HTML files fine; once you have seen this, you can start to delve deeper into the config files

We will set Nginx to run as a service, so we can use 'service nginx reload' to restart the server

You can include config files within others; here the nginx.conf controls our base processes

The default.conf file here is what sets up our site's location and settings; from that we can then apply rules based on the URIs

Build a faster web server

Switch out Apache for the lightweight Nginx HTTP server – it's faster and easy to use, as we explain...

Advisor

Kieron Howard First introduced to Linux eight years ago when installing Debian on his Xbox, web developer Kieron has continued to enjoy putting Linux on all sorts of devices



Chances are if you or your company run a web server, it's built using a LAMP stack – and while your database or scripting language may differ, by far the most popular HTTP server for the last decade or so has been Apache. Apache is reliable, feature-rich and has a large community. Apache isn't known for its speed,

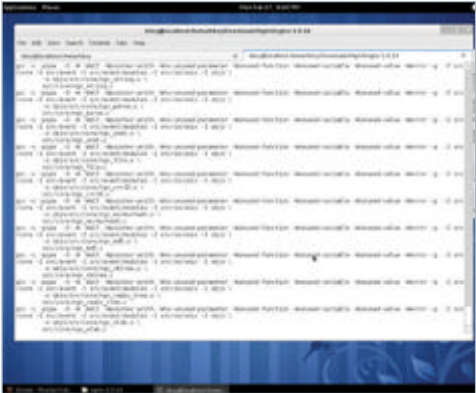
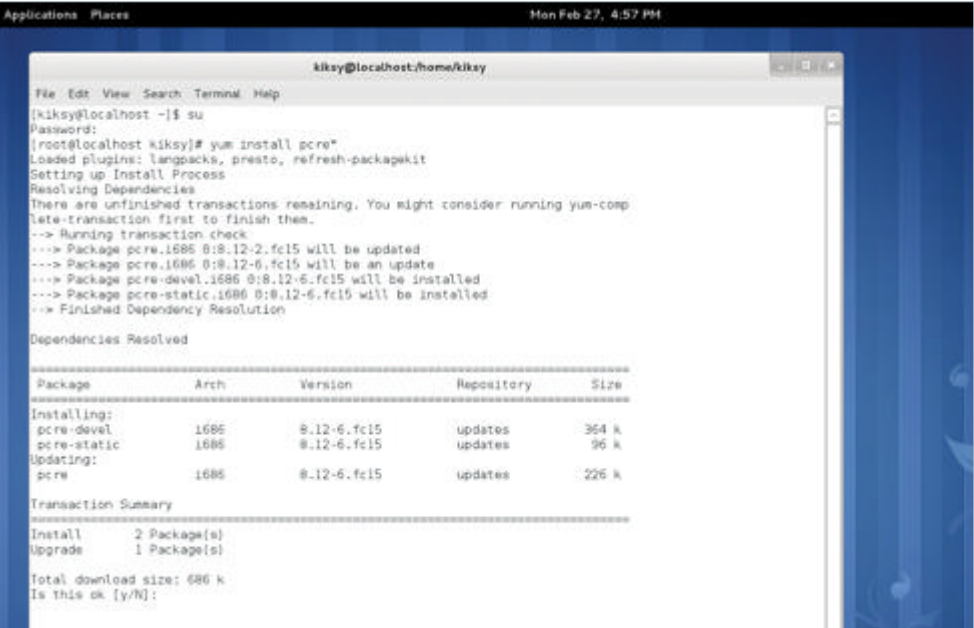
though, and while not slow, numerous other alternatives have arisen such as Lighttpd and Varnish, which aim to be much more focused on performance. In recent years more interest has been shown in Nginx, written by Igor Sysoev. He initially wrote it nearly ten years ago and it has been open source since 2004. In January 2012 Netcraft reported that 12 per cent of websites were hosted using Nginx. This tutorial will take you through the installation process, from building Nginx from source to making a basic working setup, then benchmarking it so you can compare it against your server. Finally we'll run through adding in some modules such as enabling HTTP authentication.

Resources

Suitable init.d script
– example here:

<http://articles.slicehost.com/2009/2/2/centos-adding-an-nginx-init-script>

http://wiki.nginx.org/Install#Building_Nginx_From_Source

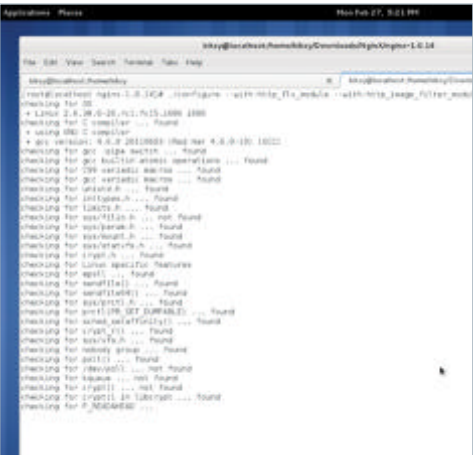


04 Build from source
Once the makefile has finished being written, all you need to do next is enter: `make` in the same directory to start the build process. This should take a few minutes. Once it's completed, enter `make install` to install Nginx.

01 Download dependencies for compiling
Before we can compile Nginx from source we need a few dependencies. To install Perl Compatible Regular Expression (PCRE) using YUM, you should enter:

```

yum install pcre*
or...
apt |apt-get install libpcre3 libpcre3-dev
Next we need zlib, so:
yum install zlib zlib-devel
or...
apt-get install zlib1g zlib1g-dev
    
```



05 Start Nginx
To start the service running, cd to the `/usr/local/nginx/sbin` directory and then just enter:

```

./nginx
    
```

There is no success message, but you can test if it has started okay by running the command again. You should then see an error message saying the ports are already in use.



02 install Open SSL
So that we can enable serving of secure HTTP, we need openssl:

```

yum install openssl openssl-devel
    
```

or...
`apt-get install openssl openssl-dev`
Now extract the tarball containing the Nginx source and in the terminal, cd to the directory containing the files.

03 Set build options
Nginx is designed to be highly modular, and so we need to add in some common features to our build that are not enabled by default.

```

./configure --with-http_ssl_module
--with-http_image_filter_module with-
http_flv_module
    
```

These add in SSL, image manipulation with GDLibrary and Flash video support.

06 Set as service
In the terminal, `run:`

```

sudo gedit /etc/init.d/nginx
    
```

to create a new init script. Then copy in the example from the link above, or modify it to suit your needs, although this should cover all the default settings. Then make the file executable `using:`

```

chmod +x /etc/init.d/nginx
    
```

You can now use:

```

service nginx stop
or 'start'.
    
```

“In January 2012 Netcraft reported that 12 per cent of websites were hosted using Nginx, and that number is steadily growing”



07 Set on boot

For Fedora and similar distros, use:

```
update-rc.d f nginx defaults
```

to set to launch Nginx at boot. For Ubuntu and other Debian spin-offs, you can use:

```
update-rc.d f nginx defaults
```

Reboot the machine and Nginx should now be running.



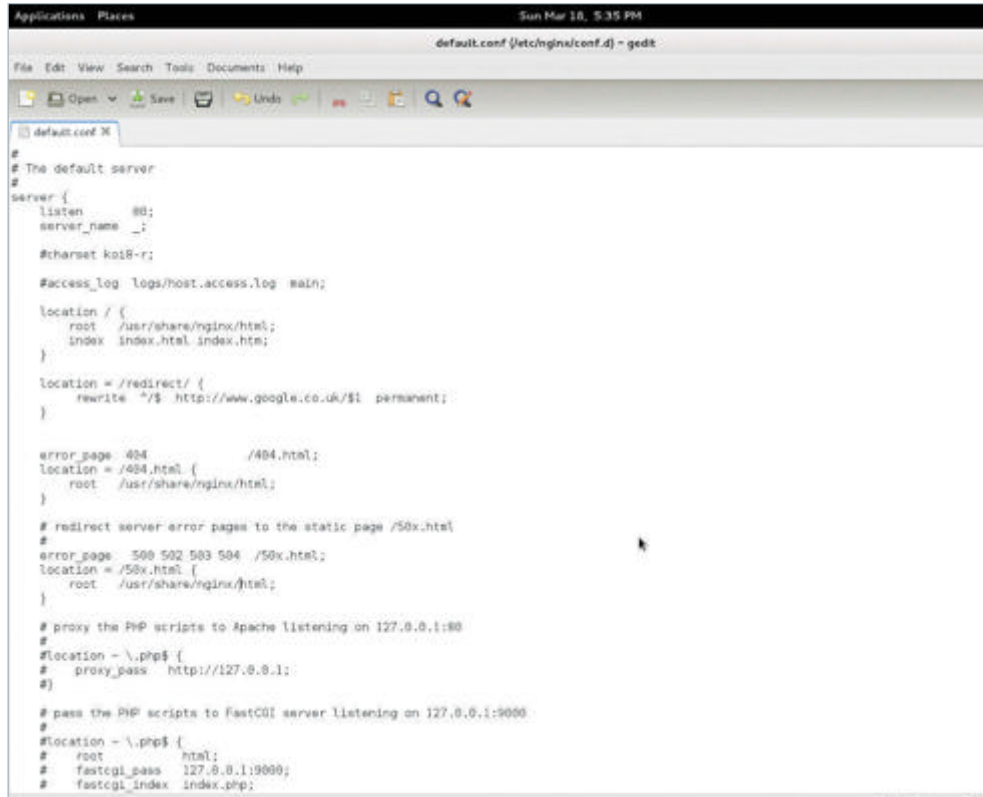
08 Test the setup

Navigate to 'http://localhost' and you should see the Nginx 'Hello World' message, which is 'Welcome to Nginx'. If you can see this then everything is running correctly. If not, check your dependencies as listed above.



09 Redirects

Nginx URL rewriting works a little differently to Apache's, but the basics are the same. When we built our Nginx from source we also made sure PCRE was included, so we use regular expressions for our rewrites. Open up 'usr/local/nginx/conf/conf' again.




10 Redirecting

The redirect aspect of Nginx is very powerful. As a basic example, **enter**:

```
location /redirect/ {  
    rewrite ^/$ http://www.google.  
co.uk/$1 permanent  
}
```

This example will then send the browser to www.google.co.uk/redirect when 'localhost/redirect' is entered.



11 Error pages

The error page controlling works on a similar principle. It uses the 'error_page' directive. So, for example, if the server returns a 403, you can redirect to your chosen forbidden message **using**:

```
error_page 403 /errors/forbidden.html
```

12 Conditional statements


You can use conditional statements within your directives as well. So, to block any non GET requests to a route, **you could do**:

```
if ($request_method !~ ^(GET)$ ) {  
    return 403;  
}
```

Which would then show our forbidden message from above.

13 Tuning performance

Depending on your situation and requirements you may need to configure your nginx setup from the default. This is especially likely if you are running on a more powerful machine. By default only one 'worker process' will run, which is good for single core machines.



14 Set worker processes

For every core on your nginx server, you can add one worker process, so a quad core machine would have 4 worker processes. You can also increase the number of worker_connections from the default 1024 if you are on a higher spec server.

“Try it on your Apache server to compare results”

```

kiksy@localhost ~$ httpperf --server 192.168.1.7 --port 80 --uri /index.html --rate 400 num-conn 30000 --timeout 5
httpperf --timeout=5 --client=0 --server=192.168.1.7 --port=80 --uri=/index.html --rate=400 --send-buffer=4096 --recv-buffer=16
384 --num-conn=1 --num-call=1
Maximum connect burst length: 0

Total: connections 1 requests 1 replies 1 test-duration 0.003 s
Connection rate: 330.0 conn/s (3.0 ms/conn, <=1 concurrent connections)
Connection time [ms]: min 3.0 avg 3.0 max 3.0 median 3.5 stdev 0.0
Connection time [ms]: connect 1.7
Connection length [replies/conn]: 1.000

Request rate: 330.0 req/s (3.0 ms/req)
Request size [B]: 74.0

Reply rate [replies/s]: min 0.0 avg 0.0 max 0.0 stdev 0.0 (0 samples)
Reply time [ms]: response 1.3 transfer 0.0
Reply size [B]: header 216.0 content 151.0 footer 0.0 (total 367.0)
Reply status: 1xx=0 2xx=1 3xx=0 4xx=0 5xx=0

CPU time [s]: user 0.00 system 0.00 (user 0.0% system 33.1% total 33.1%)
Net I/O: 142.5 KB/s (1.2*10% bps)

Errors: total 0 client-time 0 socket-time 0 connrefused 0 connreset 0
Errors: fs-unavail 0 adminavail 0 trab-full 0 other 0
kiksy@localhost ~$

```

15 Run speed tests

Install 'httpperf' using:

```
yum httpperf
```

for Fedora and then open up a terminal. Httpperf allows you to automatically open thousands of connections for a single file to simulate a high traffic load. **Enter:**

```

httpperf --server [YOUR SERVER] --port
[YOUR PORT] --uri /
index.html --rate 500 --num-conn 32000
--num-call 1 --timeout 10

```

For example:

```

httpperf --server 192.168.1.7 --port 80
--uri /
index.html --rate 500 --num-conn 32000
--num-call 1 --timeout 10

```

It will then output the CPU load. Try it on your Apache server to compare results.



16 SSI module

Server Side Includes are supported from within Nginx, allowing you to insert files into other HTML files. To include a test file, create a new .shtml file with some dummy text in it and save it to your site root.

17 Open conf

Next, open up your 'default.conf' file and **add in:**

```
location ~* \.shtml$ {
```

```
ssi on;
```

```
}
```

...within the server block. Save this and then open up the index.html file you made earlier.

18 Add in the include

You can then add in the include statement, which is **like this:**

```
<!--# include file="myinclude.html" -->
```

If you create a template using SSIs then you can create basic dynamic sites without the need to use PHP or equivalent.

19 SSI blocks

As with most template engines, you can use blocks to act as holders for other file parts. **For example:**

```

<!--# block name="headerblock" -->This
is my header<!--# endblock -->
Then if another file is not present, you
can revert to the default using
<!--# include virtual="header.html"
stub="headerblock" -->

```



20 SSI variables

Basic variables can be employed, by using something **akin to:**

```

<!--# set var="THIS_IS_MY_VAR"
value="Nginx says hello" -->

```

And you can then use that in other statements like:

```

<!--# echo var="THIS_IS_MY_VAR" -->
or within values using $. eg.
<!--# include virtual="$THIS_IS_MY_VAR"
stub="headerblock" -->

```

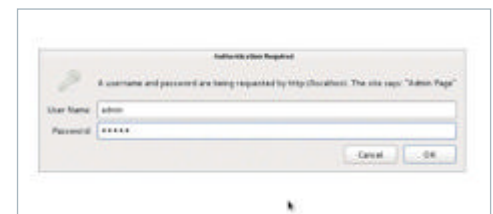
21 Conditional blocks

The SSI module also supports basic conditional statements. To compare a variable to a value, **you can use:**

```

<!--# if expr="$THIS_IS_MY_VAR = Nginx
says Hello"
//do something
<!--# else
//do something else
<!--# endif -->

```



22 Basic HTTP authentication

To enable a login for a route, you just need to add a location **and then:**

```

location /admin/ {
    auth_basic "Admin Page";
    auth_basic_user_file password_file;
}

```

You'll also then need to create a password file.

23 Generating password file

To create a password file you can use Apache's htpasswd; or if Apache is not installed then use something **similar to:**

```
printf "Admin:${openssl passwd -crypt
admin}\n" >> .password_file
```

You then need to place the file relative to the default.conf file, not site root.

24 IP Blacklist/Whitelist

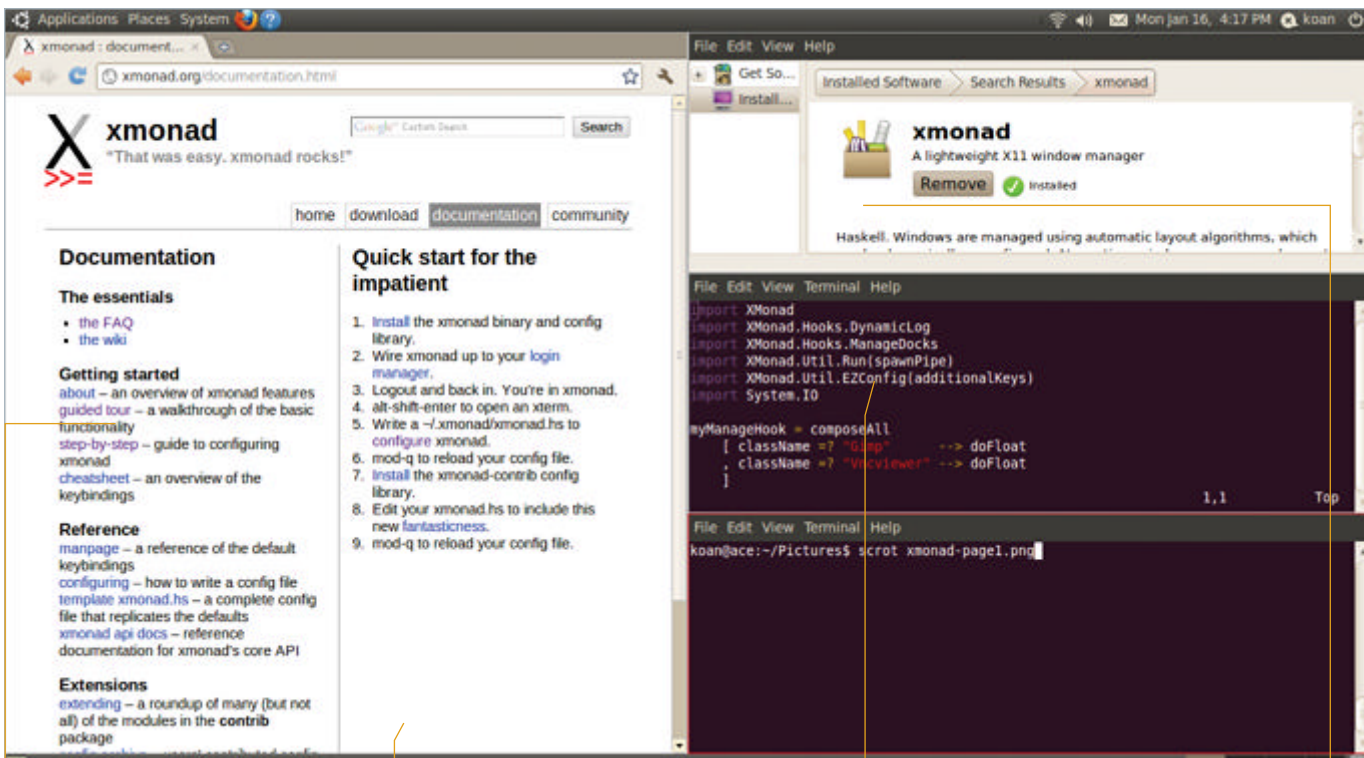
Another handy feature is to block or allow specific IP's to a route.

```

Location
{
    allow 192.168.1.7;
    deny all;
}

```

...would just allow the machine on 192.168.1.7 access, but refuse anyone else – which might be handy for development, or admin areas.



Xmonad can be used independently or within a GNOME or KDE session. The latter gives you the benefits of the GNOME panel and other familiar components

Xmonad is not easy to use, but it makes up for this by providing excellent documentation on its official website

Xmonad is completely customisable, thanks to the configuration file `~/xmonad/xmonad.hs`, which is written in the Haskell programming language

Xmonad automatically tiles all windows, so it's one of the few window managers that actually manages your windows instead of delegating this task to you

Make your desktop more productive with Xmonad

Maximise your screen space with this useful tiling window manager



Modern desktop environments like GNOME 3, KDE 4 and Unity come with a lot of bells and whistles and all seem to take the Mac's visually pleasing user interface as an example.

A couple of lesser-known window managers, however, are swimming against the tide by offering a minimal, no-nonsense desktop environment.

One of these is xmonad, a tiling window manager written in the Haskell language. Tiling means that xmonad automatically arranges your windows in a non-overlapping way (tiled), so you don't have to spend your precious time

aligning and searching for windows. Moreover, xmonad offers a quite minimal desktop experience: by default there are no window decorations, no status bar, only your windows.

But you're free to add all of these features and more if you want it, and xmonad supports advanced features such as multi-monitor support, and all actions can be done with your keyboard. And last but not least, partially thanks to Haskell, xmonad is rock solid. The result of all of this: xmonad makes you more productive because you can concentrate on getting things done.

Advisor

Koen Vervloesem has been writing about free and open source software, and IT in general, since 2000. He has master's degrees in computer science and philosophy and can be reached at koen@vervloesem.eu

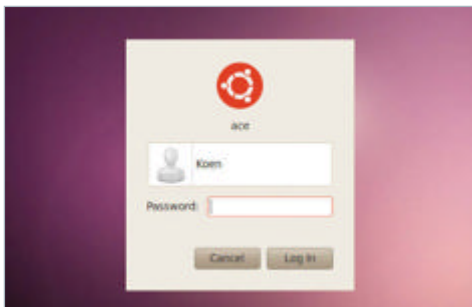


```
koan@ace:~$ apt-cache show xmonad
Package: xmonad
Priority: optional
Section: universe/x11
Installed-Size: 2364
Maintainer: Ubuntu Developers <ubuntu-devel-discuss@lists.ubuntu.com>
Original-Maintainer: Joachim Breitner <nomeata@debian.org>
Architecture: amd64
Version: 0.9.1-2
Replaces: xmonad-rebuild
Depends: libc6 (>= 2.4), libffi5 (>= 3.0.4), libgmp3c2, libx11-6 (>= 0), libxext6,
 libxineramal, x11-utils
Recommends: libghc6-xmonad-dev, libghc6-xmonad-doc
Suggests: dmenu
Conflicts: xmonad-rebuild
Filename: pool/universe/x/xmonad/xmonad_0.9.1-2_amd64.deb
Size: 531566
MD5sum: d7a0046c72bab7998fb7f5075135d786
SHA1: 9450464c2f0a2510872302136ee604ae6345ecfa
SHA256: d038a5196df4109da688dd28fcf65e220bc67e47d8e906ed91ca77ca9f0293e0
Description: A lightweight X11 window manager
Haskell. Windows are managed using automatic layout algorithms,
which can be dynamically reconfigured. At any time windows are
arranged so as to maximise the use of screen real estate. All
features of the window manager are accessible purely from the
keyboard: a mouse is entirely optional. Xmonad is configured in
Haskell, and custom layout algorithms may be implemented by the user
in config files. A principle of Xmonad is predictability: the user
should know in advance precisely the window arrangement that will
result from any action.
.
This package comes pre-configured with the default configuration. If
you want to build your custom-configured version, make sure that
libghc6-xmonad-dev is installed and put your configuration in
~/.xmonad/xmonad.hs
Homepage: http://xmonad.org/
```

“You don’t have to spend precious time aligning and searching for windows”

01 Install xmonad
 Xmonad is included in many distros’ package repositories. On Debian and Ubuntu, xmonad is split into three packages: xmonad which runs the window manager in its default configuration, libghc6-xmonad-dev for custom configurations, and libghc6-xmonad-contrib-dev which includes extra modules.

02 Login
 The Ubuntu package for xmonad puts its own entry in the GDM login screen, which makes it very easy to try out xmonad: just log out and log in again, but change the session from the default one (such as GNOME or Unity) to xmonad.



03 Alternative login
 On many distributions there’s an alternative way which gives you more flexibility: create a file ~/.xsession in your home directory, make it executable with `chmod +x ~/.xsession` and enter ‘exec xmonad’ in it. When logging in at the GDM screen, select the entry that says ‘xsession’ or ‘default session’ from the menu.

04 Blank slate
 After you are logged into your xmonad session for the first time, you’re greeted by an empty screen, a testament to xmonad’s minimalism. Now hit `Alt+Shift+Return` to launch a terminal, which will fill the complete screen. When you open other windows, they all automatically rearrange.

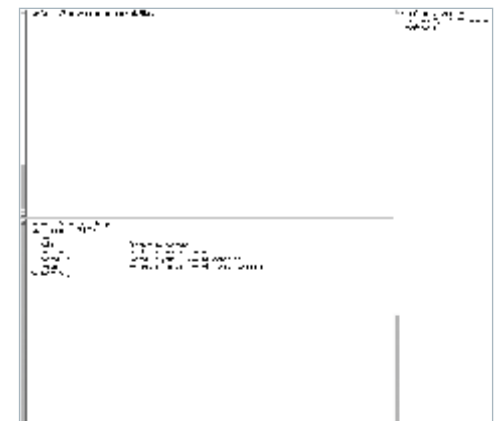


05 Focus
 You have probably noticed that there’s a small red border around the focused window. You can focus another window by moving your mouse above it, but you can also do this with the keyboard by pressing `Alt+J` and `Alt+K`, to cycle through all windows. `Alt+Tab` also works.

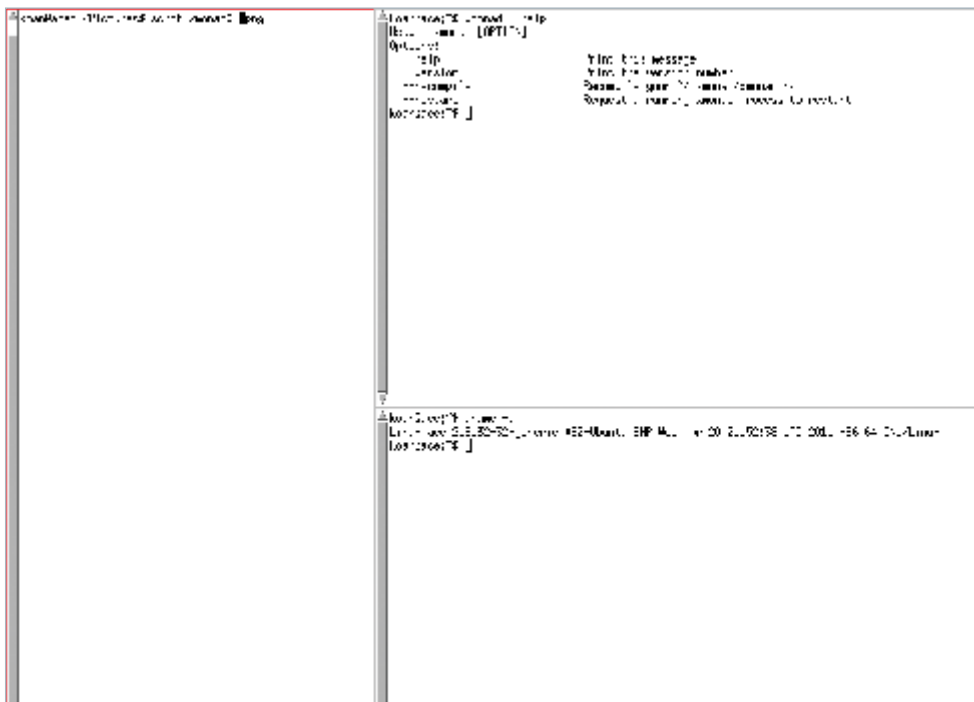
06 Shuffling windows
 While xmonad automatically tiles all windows you open, this doesn’t mean that you don’t have any influence on where they are placed: you can swap the currently focused window with its neighbour before or after with `Alt+Shift+J` and `Alt+Shift+K`.

07 Swapping masters
 As you have seen, xmonad divides the screen into two panes, and any new window you open will be placed left and as high as the whole screen. This is called the ‘master pane’. You can swap the focused window with the window in the master pane using `Alt+Return`, and you can switch to the master window with `Alt+M`.

08 More masters
 By default xmonad shows just one window in the master pane, but this is completely configurable: `alt+,` (comma) increases the number of windows in the master pane, while `alt+.` (period/full stop) decreases this number and moves the other windows back to the subordinate pane.



“As the windows have no decoration with a close button, you’ll face the problem that you can’t exit some programs. However, in xmonad there’s a shortcut key for that”

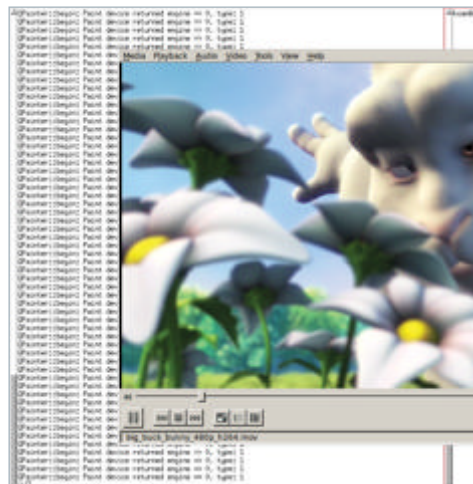


09 Resizing windows Another thing that you can change dynamically is the ratio between the master and the subordinate pane. Both get half of the screen by default, but hitting Alt+H and Alt+L a few times changes this ratio.

10 Layout Until now the layout of our windows has been fundamentally the same: a master pane at the left and a subordinate pane at the right. With Alt+space you cycle through two other tiling algorithms: the wide mode with a master pane above, and the maximised mode which maximises the currently focused window.



11 Floating windows You can easily make a window ‘floating’ (placed in front of all the tiled windows) by dragging it with the left mouse button while holding Alt. You can resize a floating window by dragging it with the right mouse button while holding Alt. Pressing Alt+T pushes a focused floating window back in the tiling layout.



12 Launcher menu Of course, terminal windows are not the only ones you want to work with. First install dmenu (in Ubuntu in the package dmenu), and then press Alt+P. Type the first few letters of a program name and hit Enter, after which the program is launched and placed above the currently focused window.

13 Workspaces You can switch to other workspaces with Alt+1 to Alt+9, and open other windows there. Each workspace can have its own layout mode. Moving a window from one workspace to another one is as easy as focusing on the window and press Alt+Shift+number (the number of the workspace).

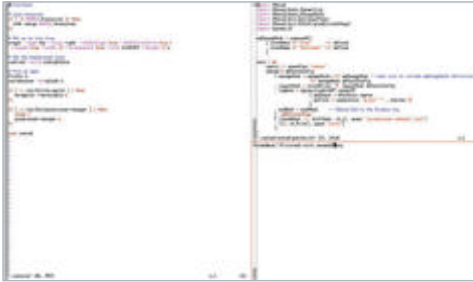
14 Exit Before we move on to some more advanced stuff, we need to tell you about some handy shortcut keys. As the windows have no decoration with a close button, you’ll face the problem that you can’t exit some programs. However, in xmonad there’s a shortcut key for that: Alt+Shift+C. To quit xmonad itself (and go back to the login screen), press Alt+Shift+Q. Alt+Q reloads the xmonad configuration, which comes in handy in the following page.

15 Custom configuration Xmonad is configured by the file ~/.xmonad/xmonad.hs, which is actually a Haskell program. You can completely customise xmonad’s behaviour in this file. You can start with a simple configuration which reuses the default configuration but changes some minor parameters, for instance the terminal program that is started when hitting Alt+Shift+Enter, or you can instruct xmonad to use the Windows key as the modifier key for all shortcut keys instead of Alt. After you have changed the configuration, run ‘xmonad --recompile’ to identify any syntax errors, and if all goes well you can simply let xmonad reload its configuration with Alt+Q.



16 Other apps

Xmonad is just a window manager, so when you are logged in, you don't have a status bar, and probably you don't even have network connectivity because Network Manager normally sets this up. So you have to fire up programs like nm-applet, trayler (for applets) and so on in your .xsession file.



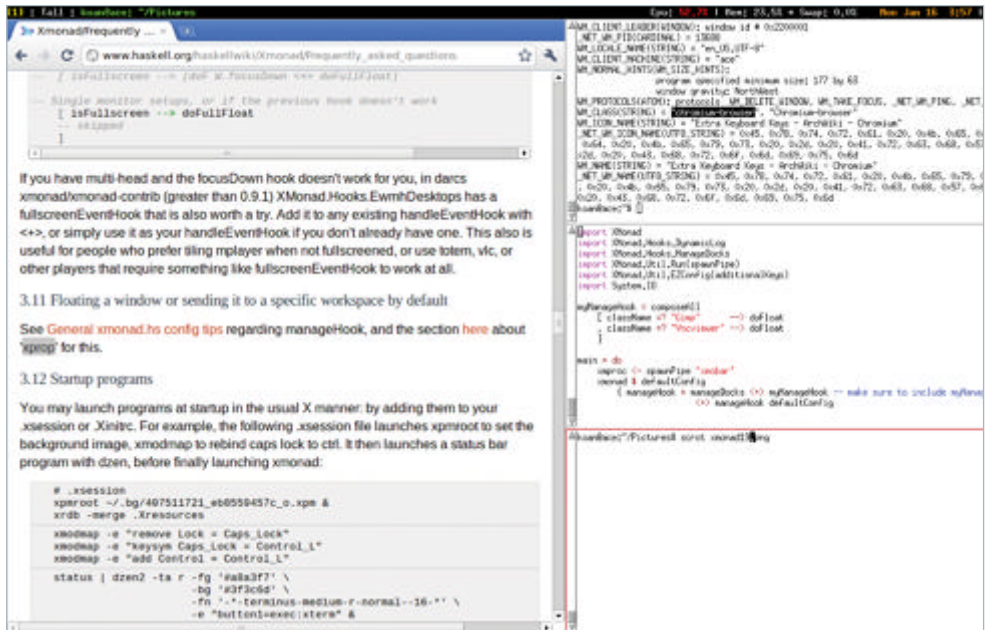
17 Xmobar

Another interesting addition is xmobar, a program that accepts input on stdin and displays the input or any other information at an arbitrary position on the screen. You can configure it in ~/.xmobarrc to show a kind of status bar with the date and time and some system information.



18 Bind your keys

You can bind specific programs to shortcut keys in ~/.xmonad/xmonad.hs. For instance, a convenient key binding to use is the PrintScreen key to launch scrot, a lightweight screen capture command. Another handy command to bind to a key is 'xscreensaver-command -lock' to lock your screen. Use xev to know which key symbols to use.



19 Float automatically

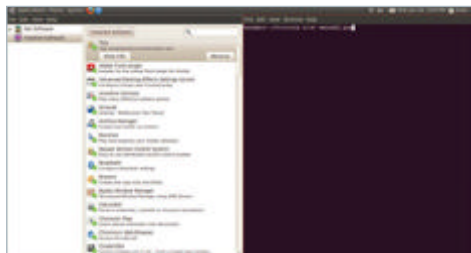
For some programs, you don't want xmonad to tile them, for instance because the program has too many windows you want to position yourself or because you prefer to manage its size yourself. Fortunately, it's easy to float specific programs automatically in ~/.xmonad/xmonad.hs by specifying their window class name or title, which can be found with xprop.

20 Compositing

If you like a little bit of eye candy with xmonad, you'll be delighted to hear that you can use compositing features like transparency. Just enable compositing in your X server, confirm with 'xdpyinfo | grep Composite' that your hardware supports it and fire up 'xcompmgr -c &' in your ~/.xsession.

21 Use xmonad in GNOME

If you really like your current desktop environment – such as GNOME, Unity or KDE – but you only want to replace its window manager with xmonad, that's also possible. There are some modules in xmonad-contrib that integrate xmonad with the pagers, status bars, tray apps, run dialog, logout window and so on from GNOME, KDE or Xfce.



22 Extensions

The beauty of xmonad is that it can be extended (although you have to know Haskell for this...). The xmonad-contrib package contains tons of third-party extensions to use in your ~/.xmonad/xmonad.hs, such as extra layouts, hooks, actions and a lot of utility functions. All of these are extensively documented on <http://xmonad.org/xmonad-docs/xmonad-contrib/>.

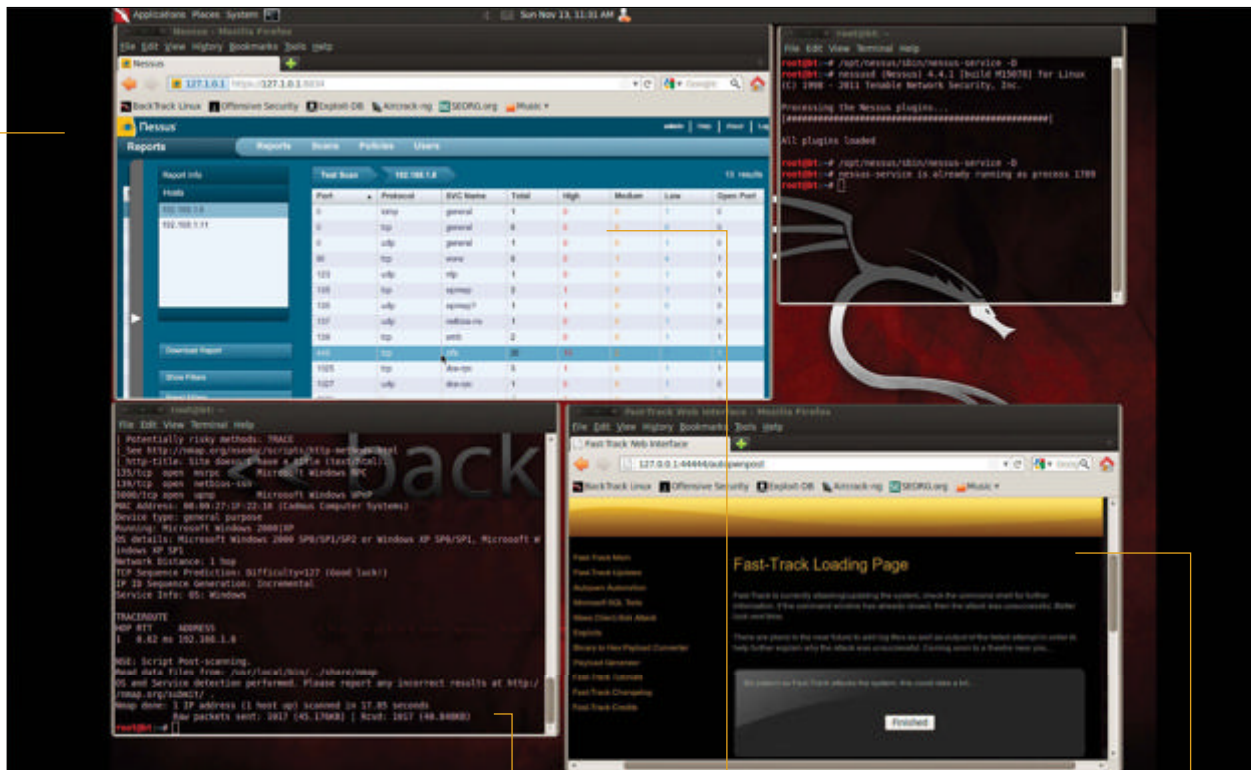
23 Graphical problems

Most programs are window manager agnostic, but some of them are problematic when used with xmonad. For instance, Java applications may result in grey blobs being shown. Export the environment variable _JAVA_AWT_WM_NONREParenting=1 to solve this. Also, OpenOffice.org looks bad in xmonad; export OOo_FORCE_DESKTOP=gnome.

24 Trouble?

Xmonad is not easy to configure, so you'll surely encounter some problems. Always check ~/.xsession-errors first for error messages, and have a look at the project's FAQ for some common problems and solutions. After a while you'll learn the xmonad way of doing things, which will make your work on the desktop much more productive!

“After a while you'll learn the xmonad way of doing things”



Nessus web interface, showing the vulnerabilities found in our target machine. The interface can be accessed remotely and there is even an Android app

The results of an Nmap port scan. Here it has shown us the OS, as well as the open ports and services

BackTrack looks and acts like other common distros. This version is using GNOME, but a KDE version is also available

Autopwn Automation is built from the Metasploit framework, but automates the tasks, making it incredibly simple

An introduction to penetration testing using BackTrack 5

A beginner's guide to testing your network's security using the popular BackTrack 5 distro

Advisor

Kieron Howard First introduced to Linux eight years ago when installing Debian on his Xbox, Kieron has continued to enjoy putting Linux on devices such as iPods, PS3s and various phones



Resources

BackTrack 5: <http://www.backtrack-linux.org/backtrack/backtrack-5-release/>

Nessus: <http://www.tenable.com/products/nessus> (Free for home use)




It's now common for SMEs to have complex network setups, with a variety of web, email, remote access, file servers and more.

If the business has grown rapidly or had a rapid turnover of IT administrators, it can be hard to keep track of exactly how secure the network is.

Many admins like to take the path of least resistance – if it ain't broke – so some systems can remain unpatched and updated for months or even years. Being able to test the network yourself, while obviously not as thorough as a proper audit by a security firm, can at least

give you some preliminary information on the robustness of your setup. Penetration testing is the act of intentionally trying to break into a network (that you are authorised to) and gain control of it, or part of it. While a huge subject, which is ever evolving, there are some tools which help simplify the process and act as a gentle introduction.

The Linux distro BackTrack 5 is one of the most well known and features a plethora of tools for pen-testing, all bundled up into a live CD. This guide will take you through the first few steps of pen-testing...



01 Network setup

For this tutorial, we are going to assume you have the IP of each machine that you wish to test on your network. Here we will test a mixture of Windows and Linux machines (the targets). In a terminal, use ifconfig (Linux) or ipconfig (Win) to get the IP.



02 Boot up BackTrack

Log into BackTrack using the username 'root' and password 'toor'. Then we will need to enable networking. Enter 'start networking' and then 'ifconfig eth0 up'. Replace eth0 with eth1 or wlan0 for wireless adaptors, depending on which you wish to use. Test the connection with 'ping google.com'.




03 Run GNOME/KDE

It's possible and often necessary to run BackTrack commands through a terminal. However, the distro also comes with a GUI, either KDE or GNOME, which you will have chosen when downloading. Start up the GUI with 'startx'. The distro is fairly stripped out, save for all the pentesting tools!



04 Port scanning

In this example, one of our targets is an unpatched Windows XP SP1 machine. The first thing we shall do is to find out which ports are open; from that we can make a guess as to what services are running. Start Nmap from BackTrack>Information Gathering>Network Analysis>Network Scanners




05 Run Nmap

Enter 'nmap -sT -p- -PN 192.168.1.6' or whatever your target IP is. The -T switch runs a TCP scan, but you can also run UDP; -p means scan all ports. After a few seconds, Nmap will return the results. Here we see 1025 is open, so possibly running a web server.

06 Run SYN scan

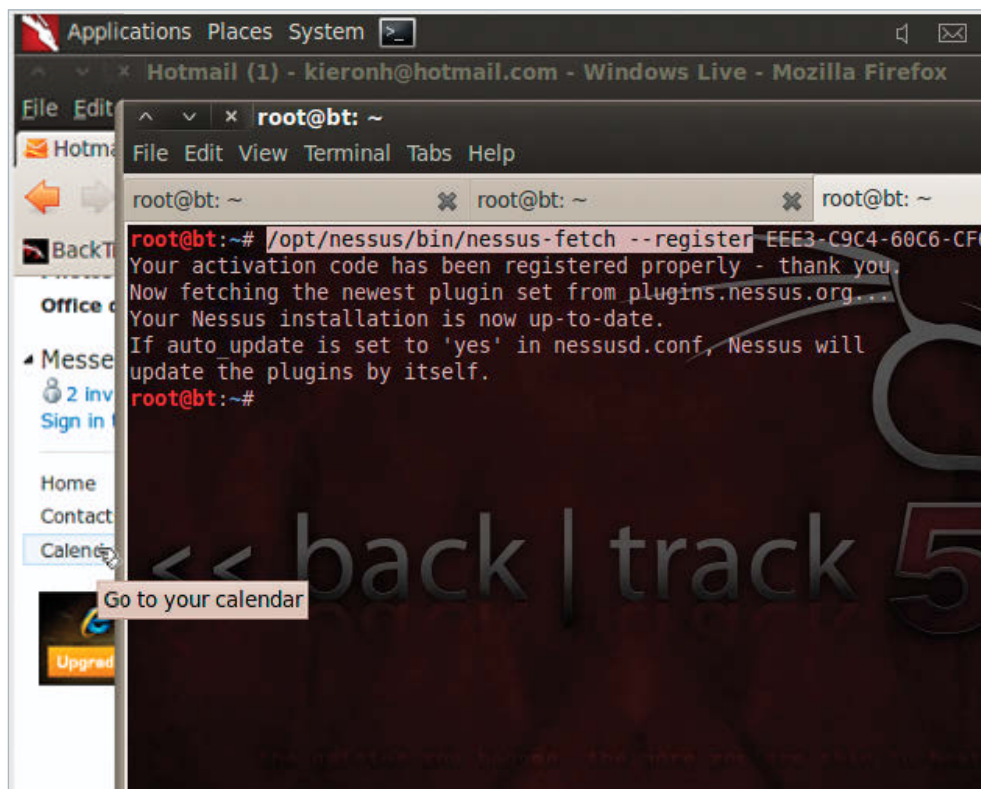
Now let's run another scan on our Linux target. This time we'll use a SYN scan, which is harder to trace, as it only performs the first part of the handshake. Enter 'nmap -sS -p -PN 192.168.1.11'. Here we see an SSH server running.



07 Attempt to SSH into the target

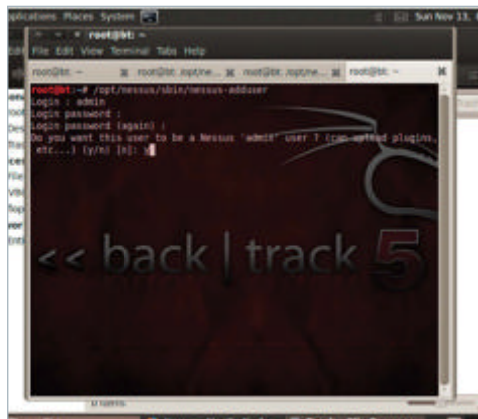
Now we have a possible entry point, we can start to try to gain entry. A simple way now would be to try to SSH as root into the target. Many people use lazy guessable passwords, and this simple port scan has shown how dangerous that is.

“Enabling you to find vulnerabilities, Nessus is a powerful tool in helping with network security”




08 Setup Nessus

Let's assume we couldn't guess the password, and try to scan for vulnerability on the XP machine. Nessus is the most common tool for this. Register on the site and then once you have a key, enter `/opt/nessus/bin/nessus-fetch --register MY-KEY-HERE`.




09 Start Nessus server

Start the Nessus server by entering `/opt/nessus/sbin/nessus-service -D`. Nessus will update the plug-ins and then start. Add a user with `/opt/nessus/sbin/nessus-adduser`. In a browser you can then enter `https://127.0.1.01:8834` and log in using the credentials you just set up. You will then see the main Nessus dashboard.



10 Add Nessus policy

A scan policy defines what is performed during our test. Click on 'Policy' then 'Add'. Make sure the 'Safe Check' box is checked. Unticking this will result in Nessus trying to run exploits on the target, and may cause problems.




11 Set up vulnerability scan

To set up a new scan, click on the 'Scans' button, then 'Add'. Enter a name for the scan, then choose the policy that we set up a minute ago. After that, enter the IP addresses of the targets you wish to scan.

12 Launch scan

Press 'Launch scan' to start the check off. You will see the status of the scan in the menu. Clicking on 'Browse' will allow you to see the progress of the scan. The number of possible vulnerabilities is shown, along with a grading, from 'Low' to 'High'.

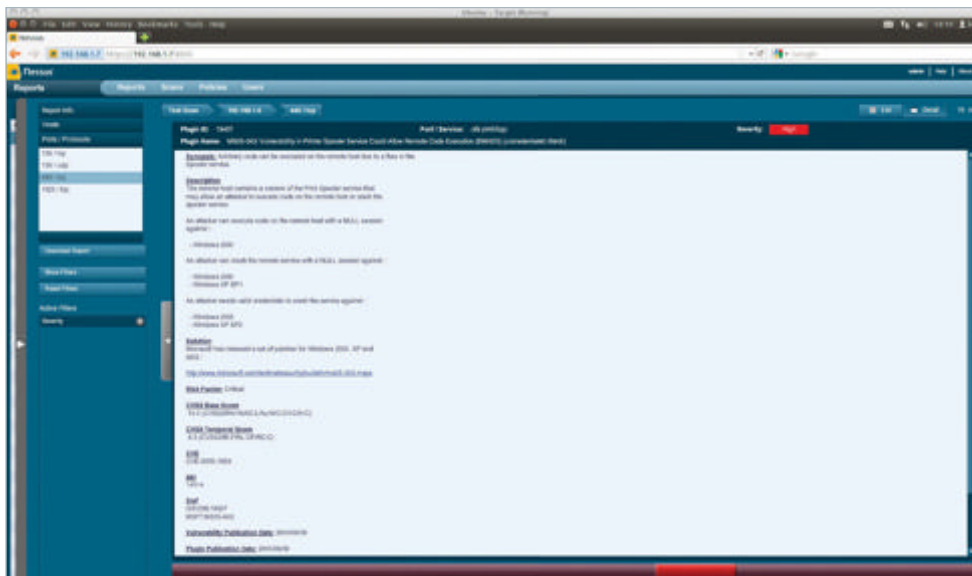


13 Analyse scan results

Click on the IP and you will be taken into a detailed breakdown of the results of the report. Here we see a list of all the ports and services which have a possible vulnerability. In this example our CIFS server has 10 high, two medium and seven low.

14 Find out more information on the vulnerability

Clicking on the number of a vulnerability will then bring up a more in-depth description of it. Ones of particular interest are those marked 'Allow Remote Code Execution'. From a security point of view, it is vital to patch these as soon as possible.



15 What's needed to fix the vulnerability
Clicking on the vulnerability will then bring up a report on the issue. The CIFS service running on our Windows XP machine is riddled with security holes; the one shown here is marked as 'Critical' and on Windows 2000 allows for code execution without even having a user session. There is a link to the patch to fix this, along with the date the vulnerability was discovered and patch date. Here they are both on the same day. As you can see, Nessus is a powerful tool in helping with network security.

16 Fast-Track exploitation tool
To highlight the severity of these vulnerabilities, let's now try and exploit them. Fast-Track is an incredibly easy tool to use, and really hammers home the importance of pen-testing and keeping your machines up to date. Run it from BackTrack>Exploitation Tools>Network Exploitation.



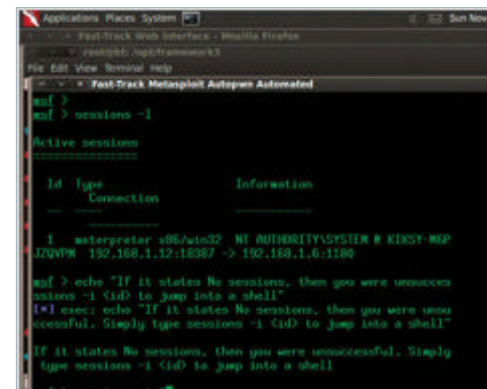
17 Fast-Track interface
Once the web-gui is running, in a browser enter 127.0.0.1:44444. You will then see the main page. Be warned: Fast-Track can be destructive, so it might be worth cloning any target machines into VMs to test them first.




18 Set up Autopwn Automation
Running FastTrack's automation option is stunningly easy. Just enter the IP address of the target machine, choose whether you want to connect to the victim on payload deployment, or if you want the target to connect to you, and hit 'Metasploit Autopwn'.

19 Set up MySQL database driver (optional)
To set up MySQL as our database driver we need to open a terminal and update the Metasploit framework. Enter 'apt-get update' then 'apt-get dist-upgrade', 'service mysql start' (assuming you have MySQL installed already) and 'msfconsole'. Choose MySQL with 'db_driver mysql'.

20 Update config
In a terminal, 'open nano /pentest/exploits/fasttrack/config/fasttrack_config' and edit the PATH to /opt/framework3/app/. Save, then 'nano /pentest/exploits/fasttrack/bin/ftsrc/autopwn.py'. Edit the line 'child1.sendline ('db_driver sqlite3')' to 'child1.sendline ('db_driver postgresql)', and 'child1.send ('db_create pentest')' to 'child1.send ('db_connect msf3: YOURPASSWORD@127.0.0.1:7175/msf3)'. The password is in /opt/framework3/database.yaml.



21 Run exploits
In the Fast-Track web interface, run the Autopwn feature. This will open up an MSF console and run through all known exploits for the available ports and services. It may take some time but by the end of it, if any exploits were successful, you will have a console session running.



22 Test access
To access the session on the target, type 'sessions -i 1', replacing the number with whichever session you want to access. This will open the Meterpreter prompt. Type 'ls' to show the contents of directory. Here you can see a list of Windows .dlls.

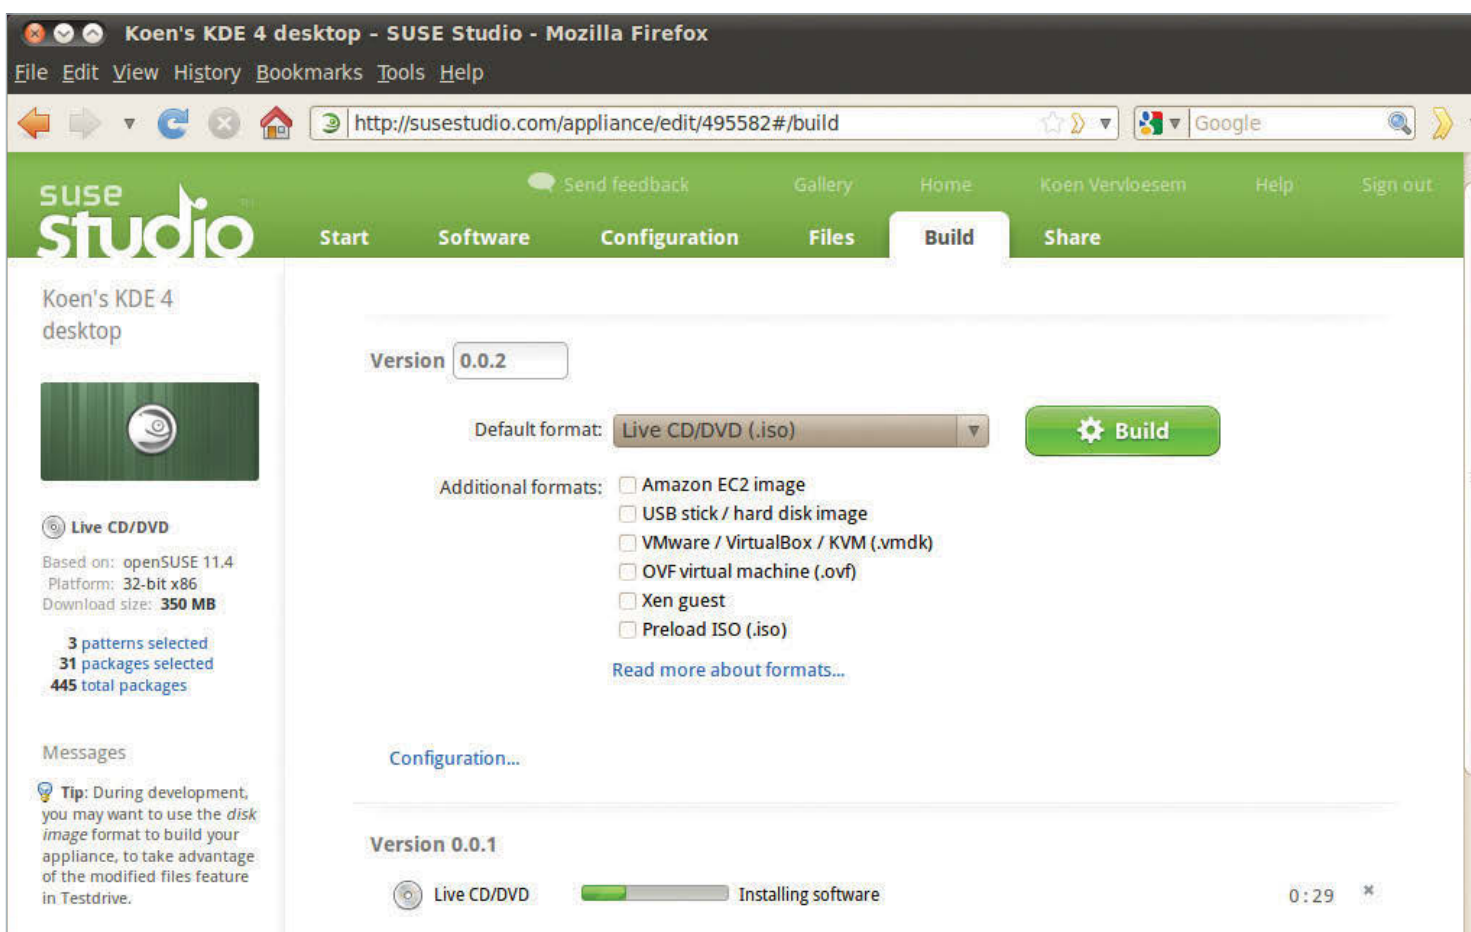
23 Install VNC
At the Meterpreter prompt, type 'run vnc'. This will install a VNC server on the target then output the IP and port it's running on. In the BackTrack terminal, enter 'apt-get install xtightvncviewer' to install our VNC Viewer; you can then view the target machine's desktop.

24 Take screenshot
To grab a screenshot as proof of entry, in Meterpreter type 'ps' to list processes running on the target. Then change to explorer.exe by typing 'migrate <Process ID>' (eg 'migrate 1460'), then 'use espia' and 'screengrab'. Now go and update the machine to make sure that no one else can do it!

Build your own Linux distribution with SUSE Studio

Have you always wanted to create your own custom Linux distribution? Look no further than SUSE Studio...

■ Build your appliance in one of various formats



Linux is known for its extreme customisability, so you can always tailor your own Linux desktop to your taste by installing your favourite software, changing the default configuration and extending the distribution with custom shell scripts. But perhaps you want something more: a custom Linux distribution, ready to be distributed and installed by yourself or others.

However, creating your own Linux distribution is no small feat, as it is a complex and time-consuming operation. There are a couple of projects that can help you with this endeavour,

and SUSE Studio is without doubt the most easy-to-use, the most well-engineered and the most efficient way to build your own Linux distribution. Even if you have no previous experience whatsoever, you'll have made your own Linux distribution in less than 15 minutes.

SUSE Studio is a web-based solution, so you just browse to susestudio.com, create an account and you can start configuring your Linux distribution (SUSE Studio calls it an 'appliance'). The process is really easy to follow, and you get a lot of guidance and tips along the way. At the end, you choose whether

Advisor

Koen Vervloesem has been writing about free and open source software, and IT in general, since 2000. He has master's degrees in computer science and philosophy and can be reached at koen@vervloesem.eu



you want to create an ISO image for a live CD/DVD or some other disk formats, and you can even push your appliance image to Amazon EC2 to launch an instance in the cloud. All this takes just a few clicks in your web browser. The only obvious downside is that appliances you create with SUSE Studio are based on openSUSE or SUSE Linux Enterprise. So if you don't like the SUSE family, you're out of luck.

Why build your own appliance?

There could be many reasons to build your own appliance. Maybe you want an easy way to install the same customised distribution on all your systems. Then you can customise this installation just once on the SUSE Studio website, generate the ISO image and install this image on all your computers. After this, you have the same set of your favourite applications and configuration tweaks on all your machines.

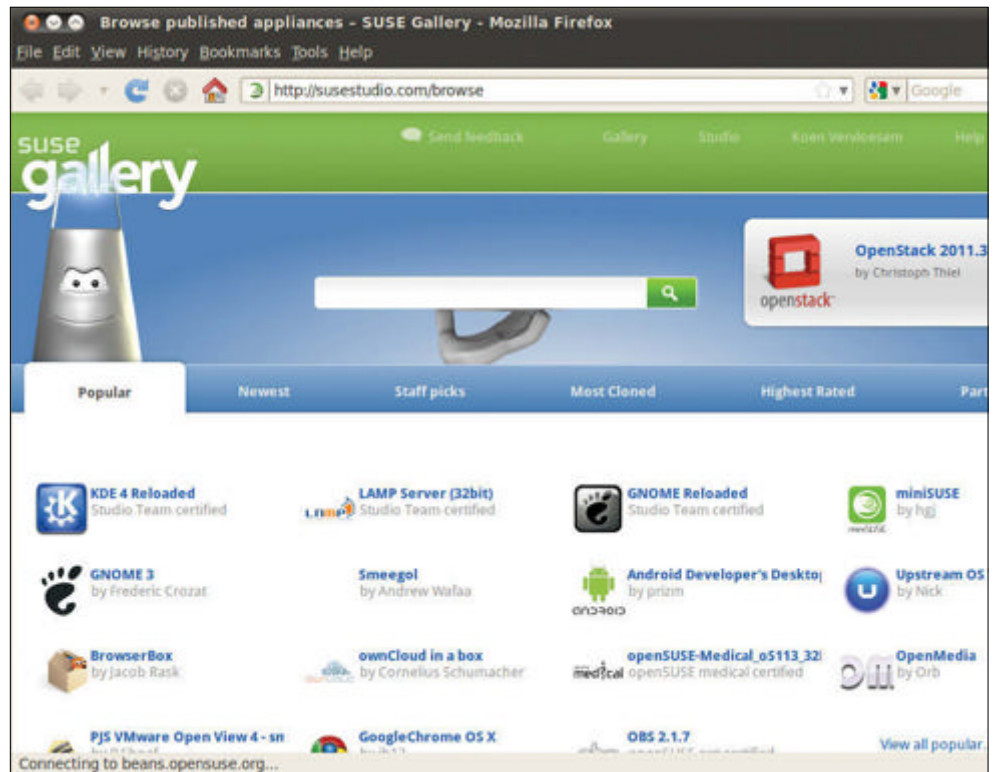
But there are also other use cases: maybe you always wanted to create and distribute your own Linux distribution aimed at a specific niche. With SUSE Studio you can do this very easily, including your own artwork. And you can publish the result on SUSE Gallery, where others can download it, comment, submit ratings, or even use your appliance as the base of their own custom distribution.

Or maybe you want to demo some application during a trade show or in a class you're teaching. Just create a live CD with this application and some demo content and off you go. Or maybe you want a distribution with your company's branding. There could be a lot of reasons to build your own Linux distribution, and in many cases SUSE Studio is the right tool for the job.

Getting started

Just browse to susestudio.com and click on 'Sign in or create an account'. The website accepts logins from Novell, Google, Yahoo!, Twitter, Facebook and OpenID. So you use your account on one of these other websites to log into susestudio.com, but without sharing your password on these websites to SUSE Studio. If you don't have an account on one of these websites, SUSE Studio asks you to create an OpenID or Novell account.

After the first login, the system offers some templates to build an appliance upon. After later logins, you can also start the same sequence with a click on 'Create new appliance...'. You have the choice of openSUSE or SUSE Linux Enterprise for the base system. At this moment SUSE Studio lists openSUSE 11.4, SUSE Linux Enterprise 11 SP1 and SUSE Linux Enterprise 10 SP4, but when you're reading this article you'll probably see the newest openSUSE 12.1 in the list. For each operating system, you have the choice between various templates. Just Enough OS (JeOS) is a tiny minimalistic appliance with the bare minimum components to run a server



■ Browse the appliances that other users have made

application. Server is a text-only base that can be used to run a complete server operating system. Minimal X is a lightweight graphical system, so not with GNOME or KDE but using the IceWM window manager, which is especially useful for making a kiosk-like appliance. Two other base templates offer GNOME and KDE 4 respectively, and last but not least you can import a Kiwi or AutoYaST configuration file that specifies the contents of your appliance. After you have made your choice, you simply have to select your architecture (32- or 64-bit), choose a name for your appliance and click on 'Create appliance'.

After this, you navigate through the rest of the steps by clicking on some tabs: Start, Software, Configuration, Files, Build and Share. At any time, you can see the current disk footprint of the appliance in a sidebar at the left, which is useful while building an image that has to fit on a CD or a USB stick. The sidebar also shows helpful messages and tips, depending on the tab you have opened and even depending on the packages you have installed. Now click on 'Switch to the Software tab to continue'.

Adding your own software

In the Software tab, you can add and remove packages and repositories at will. If needed, you can even upload your own RPM files. In the sidebar you can see the number of patterns selected (a pattern

is a group of related packages; eg base, x11, office, laptop, kde4...), the number of explicit packages selected, and the total number of packages. By clicking on any of these numbers, you get a list of the relevant packages or patterns in the Software tab. You can also search for packages and patterns or click on some icons to get a list of all patterns, recommended packages, packages in specific software groups, and so on. This interface clearly shows that SUSE Studio is well engineered and has undergone a lot of usability tweaking. For example, the Recommended list of applications is different depending on the template you have chosen.

Click on the name of a package to get more information. This shows a pop-up with some details, such as the repository, version, size, licence and a short description. Just click on the '+add' button at the left of the package name to add it, or on the '-remove' button when the package is selected and you don't want it. If you want to install a package that is not available, or a newer version than the one that is available, click on 'Add repositories...'. After which you can search for all available resources, including popular ones like PackMan and many repositories from the Build Service. And if your favourite repository isn't in the list, just click on 'Import new repository...' and enter the repository URL, after which the repository is available to all SUSE Studio users.

Configuring your new appliance

After you have chosen the software packages you want to have in your appliance, the Configuration tab allows you to configure a lot of options on various pages. For instance, on the General page you can enter the default locale and time zone (or specify that these values are asked for on the first boot). You can also configure the network (DHCP by default) and the firewall, and add users and specify their shell.

The Personalize page allows you to select or upload a logo that is shown on boot, as well as a background picture that is used for the boot menu, graphical boot, display manager and desktop wallpaper. On the Startup page you choose the default runlevel, such as 5 for a graphical login and 3 for a text-based login. You can even add an end-user licence agreement (EULA) that is shown to the user when your appliance boots for the first time.

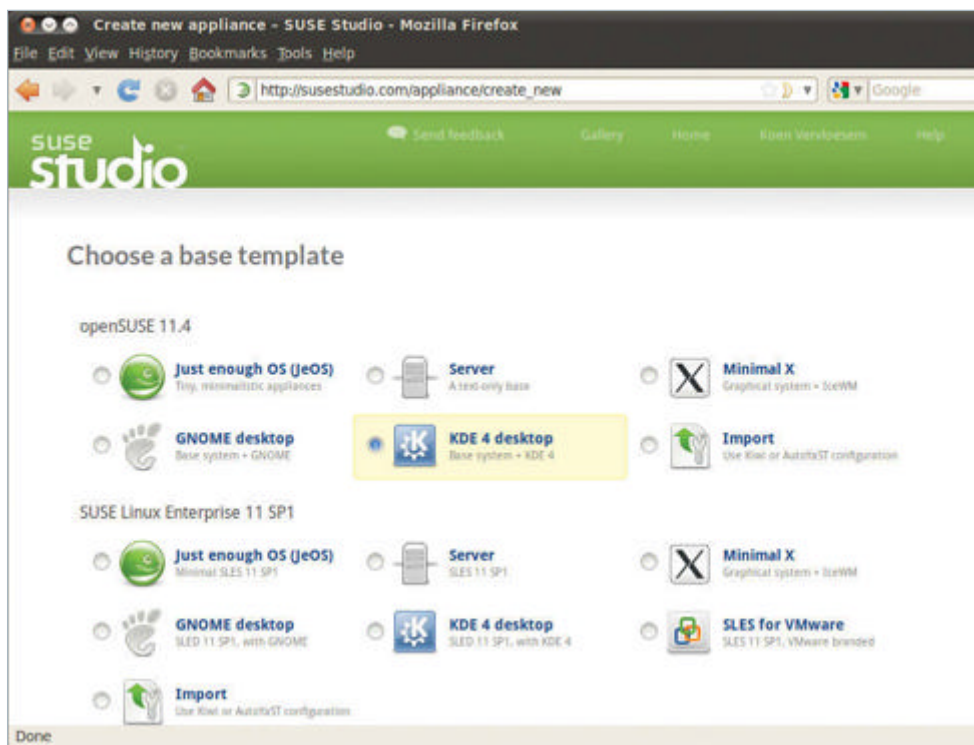
On the Server page, you can choose to set up PostgreSQL or MySQL, after which you're asked to upload a database dump file and configure MySQL users and permissions. When you don't have MySQL installed but are setting up MySQL, SUSE Studio complains with a big red message in the sidebar, and you just have to click on it to add the MySQL server package.

On the Desktop page, you can choose to automatically log in a specific user, and you can add desktop programs that are automatically started for specific users after they are logged into their X session. On the Appliance page you can set some technical parameters, like the available RAM and disk size for virtual machines and the amount of swap space for a disk image. You can also enable extended memory (PAE) to access more than 4GB and the Logical Volume Manager (LVM), as well as the live installer, Xen host mode, VMware CD-ROM support and WebYaST. And finally, on the Scripts page, you can also add custom Bash scripts in order to run specific commands at the end of the build or whenever the appliance boots.

You're almost at the end now. The Files tab lets you add specific files that will be copied into the appliance after packages are installed. This is an optional step and it's surely not needed in many cases. But you can use this, for instance, if you want to populate the user's home directory with some example content or some dot files for the configuration of the installed packages. If you upload a single file, it will be copied to the specified directory, and if you upload an archive, it will be extracted into the directory.

Building the finished product

That's it for the configuration; now you can start building. Enter a version number for your appliance and choose a disk format: 'USB stick / hard disk image' if you want to copy the image to



■ Choose a base template for your appliance

a USB stick or hard disk with the `dd` command, 'Amazon EC2 image' if you want to run your appliance on Amazon's cloud, or 'Live CD/DVD' if you want to burn the image to a disc to run it as a live OS – and there are also various formats to run your appliance inside hypervisors such as VMware, VirtualBox, KVM and Xen. You can also enable more than one disk image format.

It's possible that you need to add some packages to build a specific format. For instance, when you choose to build an Amazon EC2 image, you need the `suse-ami-tools` package from the `Virtualization:Cloud:EC2` repository, so you'll get a warning in the sidebar and can install the required package with one click. Afterwards, click on Build to start building the appliance. You can also view the configuration of your appliance in an HTML version or export it to a `tar.gz` archive with a recipe to build the same appliance locally using the command-line tool `Kiwi`.

Test, test and test some more...

The build process itself takes less than five minutes most of the time, even for a complete desktop

system. After this, you can download the result. However, there's an even neater way: click on Testdrive. This allows you to boot and test your distribution in your web browser without the need to download it first. The appliance runs in a virtual machine on Novell's servers and exposes its screen via VNC to a Flash applet running in your browser. If you have problems with Flash, you can even connect using a normal VNC viewer: SUSE Studio shows you the command and the password then. There are even buttons to switch to a different virtual console, to press `Ctrl-Alt-Del`, to press `Ctrl-Alt-Backspace`, or to change the keyboard layout. Each Testdrive instance gets 512MB of RAM and an hour running time on the server.

All in all, SUSE Studio is very easy to use, has a very polished interface and good documentation. Moreover, it's completely web-based, so you don't need to install any tools to create your own distro. And Testdrive is the killer feature: being able to test your own distro without having to download it is a huge time-saver. So what are you waiting for? Go play with SUSE Studio and share your work with the world.

“Testdrive allows you to boot and test your distribution in your web browser without the need to download it first”

Launch your appliance on Amazon EC2



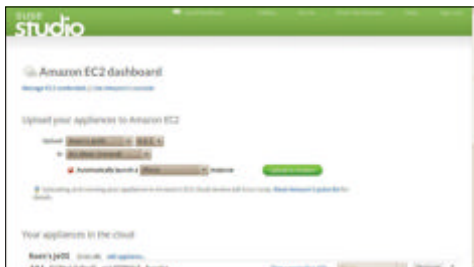
1. Configure EC2

Click on 'Manage EC2...' on the homepage of SUSE Studio. Then enter your Access Key ID and Secret Access Key (both of which you can find in your Amazon AWS account page), as well as the default region you want to run your appliance in.



2. Create an EC2-ready appliance

Create an appliance, and in the Build tab choose 'USB stick / hard disk image' for testing and 'Amazon EC2 image' as an additional format. Click on Build, then add the suse-ami-tools package from the Virtualization:Cloud:EC2 repository when you get the warning that you have forgotten it.

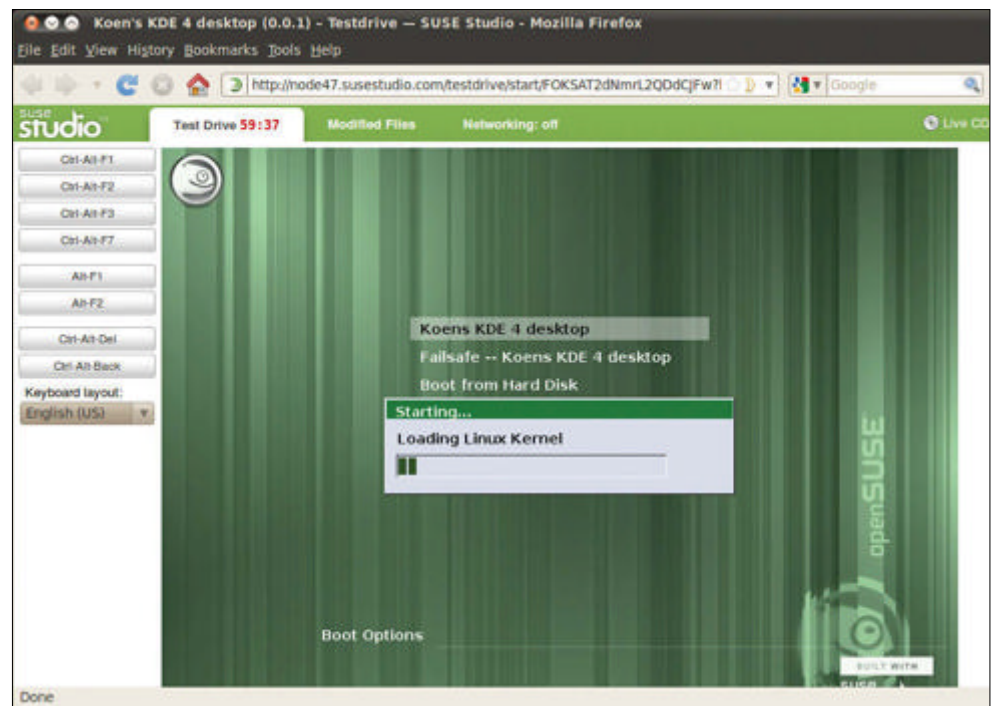


3. Launch your appliance on EC2

Test the image with Testdrive and if it works, click on 'Build additional', after which you should click on 'Upload to EC2'. If you have enabled the right checkbox, the appliance even automatically launches on EC2 after it has been uploaded. Instructions about how to connect are shown when you click on 'Connection info...'

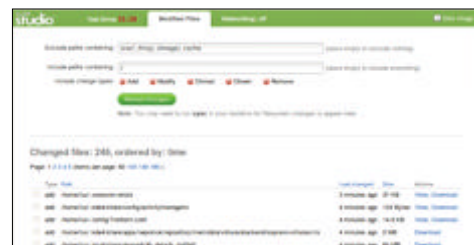
“If you have enabled the right checkbox, the appliance even automatically launches on EC2 after it has been uploaded”

Test your appliance in your browser



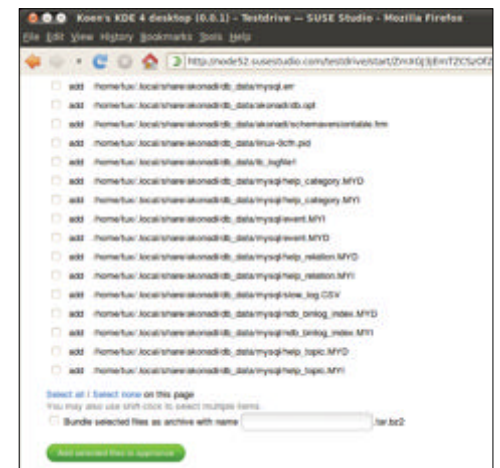
1. Start Testdrive

After you have built an image for your appliance, click on Testdrive. To take advantage of all the features of Testdrive, build your appliance with the 'USB stick / hard disk image' format for testing, and only build your desired image format after testing with Testdrive.



2. View modified files

Click on the Modified Files tab to get a list of all files that have been modified during the test drive. You can view the contents of each file or download them individually. When a file is changed from the original file in the appliance, you can even ask to see the differences.



3. Add modified files

Testdrive not only allows you to try out your appliance, but you can also tweak some of its behaviour by changing some configuration files. Afterwards, go to the Modified Files tab, select the configuration files you changed and click on 'Add selected files to appliance'.

Make your own Wikipedia

Create your own wiki site using the same software that powers the world-famous Wikipedia

Advisor

Kunal Deo is a veteran open source developer. Currently he is leading two open source projects: WinOpen64 and KUN Wiki. He is also a KDE developer. He has contributed to many open source projects, including KDE-Solaris, Belenix and Openmoko



Wiki is a web content delivery system that allows users to be the creators as well as consumers of the content. There

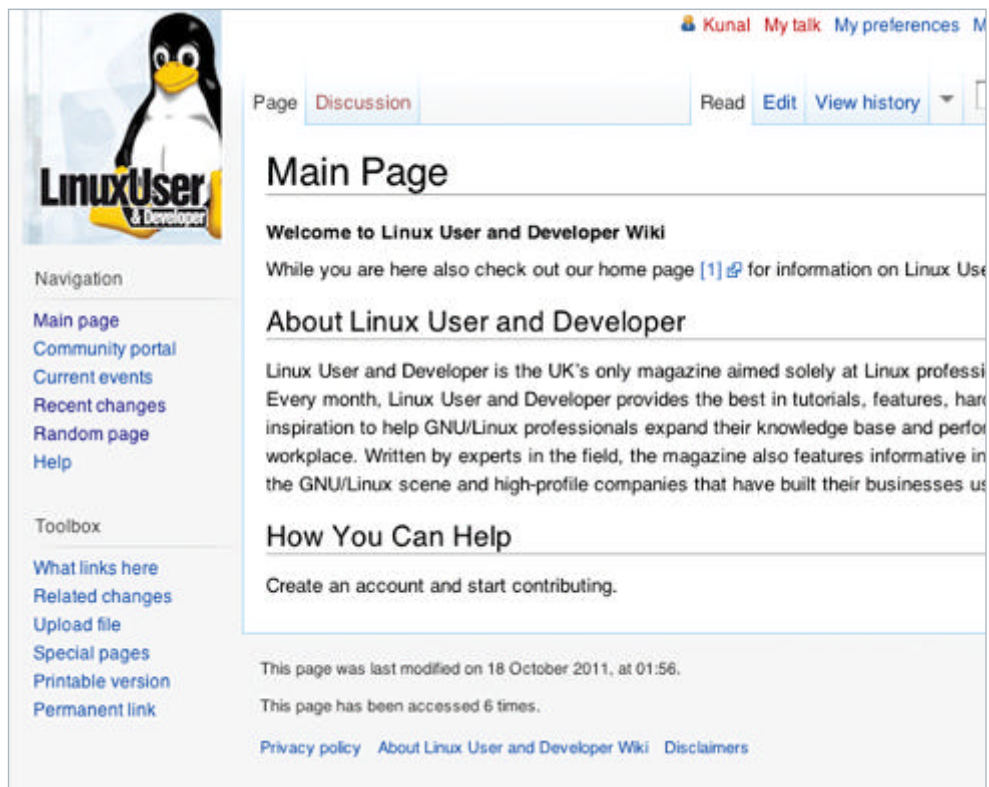
is a good chance that you have already used a wiki, one of the most popular websites in the world – Wikipedia. Wikis aren't just useful for building user driven encyclopedias, though – they are also incredibly useful for building mainstream sites such as project, documentation and support sites. A few examples are the openSUSE wiki (<http://en.opensuse.org/Portal:Wiki>), the OpenOffice.org documentation project (<http://wiki.services.openoffice.org/wiki/Documentation>), the CloudKick Support wiki (https://support.cloudkick.com/Main_Page) and there are many, many more.

The wiki method stands apart from the traditional content management system by providing the following benefits...

Ease of use: Creating a wiki site does not require you to have an expertise in web technologies such as HTML, JavaScript etc. The only thing you will need to learn is wiki syntax, which is in turn used by the Wiki Engine to create impressive webpages. In fact, even that is optional. With the advent of WYSIWYG wiki editors, you will just need to create content like you would in any word processor.

Collaborative: Wikis also provide the best collaboration platform. Your users will themselves be able to add content. Sometimes this is the only way to get up-to-date and diverse information on your site. This also makes a lot of sense for a public domain or open source project as it gives you an open environment to maintain its documents. Some also call wikis the open source document repository.

Full history: Wikis tend to keep the complete history of a page. So if you're not happy with a certain edit of a wiki page, you can easily look up its history and restore the page to its original state. Not only that,



but it also gives you a perspective on what has changed over a period of time.

Broad range of applications: A wiki can be used to do almost anything. You can even use a wiki to power up the whole website. As already mentioned, it is a very capable content management system which can be used to power encyclopaedias, documentation, support information etc.

MediaWiki

MediaWiki is an open source wiki program developed originally for use on Wikipedia. We will be using MediaWiki to create the wiki site for this tutorial. We have chosen to go with MediaWiki not just because it is open source, but also because it is powerful, extensible and has the biggest community of all the wiki software packages available.

MediaWiki is written in PHP, which uses a MySQL (or PostgreSQL, Oracle, SQLite) database

server as the back-end. MediaWiki software is made to scale with the proven scalability of Wikipedia; it is also one of the easiest wiki systems to deploy and use, but its biggest strength comes in the form of extensions.

Extensions are used to provide additional functionalities (for example extended wiki markup, enhanced security, a different user interface) without modifying the core source code.

Resources

A modern distribution like Mandriva, Ubuntu, Fedora or openSUSE

LAMP Stack: [Linux](#), [Apache](#), [MySQL](#), [PHP version 5.2.3 or later](#)
[PHP PCRE \(Perl Compatible Regular Expressions\) Library](#)
[MySQL Version 4.0 or later](#)
[PHP Configuration File \(often php.ini\)](#)

Latest version of MediaWiki (1.19.1 at time of writing): mediawiki.org/wiki/Download

1. Installing the MediaWiki archive

This step will require a write access to your web server directory. Extract the mediawiki-x.xx.x.tar.gz archive and copy it under the web server root directory. You may want to rename the extracted directory before copying it into the web server root directory, as the directory name may play a part in the wiki server URL. For this tutorial we are calling it LUDWiki.



■ MediaWiki installation screen

2. Starting the MediaWiki web-based installer

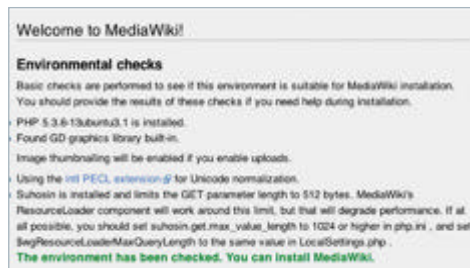
Before we move on, start the web server and the database server. Point your browser to the wiki server. The server will report that the LocalSettings.php was not found and will ask you to set up the wiki. Click on the link to start the installation wizard.



3. Selecting the language

MediaWiki is one of the most translated wiki systems in the world. This is where you get to choose the language for your wiki. You'll need to make two selections here. First you must select your language, which will configure the language for the MediaWiki admin and installation interfaces; then the wiki language, which configures the language of the content to be placed in the wiki. We set both as English.

“A wiki can be used to do almost anything. You can even use a wiki to power up the whole website”



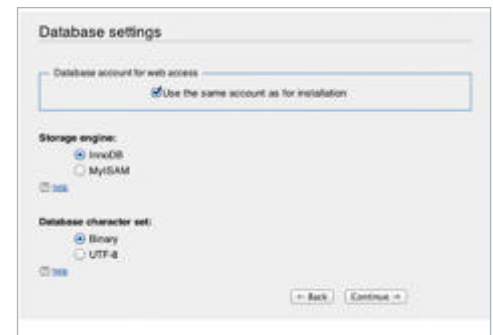
4. Installation environment check

In the next step, MediaWiki will do a little sanity check and report back with results. If you see some green text saying 'The environment has been checked. You can install MediaWiki', you are ready to continue. But wait: if you see any warnings it will be a good idea to get rid of them first (like any good programmer would do). Most times it will report things related to PHP, such as the missing extensions or non-optimised configuration. You should read these messages carefully and take action to rectify them.



5. Database connection

MediaWiki can be used with a wide variety of databases. In this step MediaWiki checks with the PHP to see what supported database extensions are enabled and provides you with a list to select from. In this case we have both SQLite and MySQL enabled. Since MySQL is the primary target for MediaWiki, we'll go with that. For MySQL, fill in the database hostname (or the server IP address) where the server is running, database name and other details.



6. Database settings

In this step you will need to select some advanced database properties. Default values are acceptable in most scenarios, but depending upon your needs you can modify these options. For example, if you are creating a small-scale wiki installation you may want to select MyISAM as the storage engine as it provides faster performance for such small-scale wiki deployments. Or, if you want the wiki text conversion but are able to live with limited Unicode support, you can select UTF-8 as the database character set.



7. Name

This is a very basic step. Here you can configure the database name and the administrator account. At this point you can choose to continue with further configuration or just start the installation process. For this tutorial we would like to configure a few more options. Select 'Ask me more questions' and continue.

User rights profile:

- Traditional wiki
- Account creation required
- Authorized editors only
- Private wiki

8. Options

In this step you get to configure few of the advanced settings for the wiki...

User Rights Profile: While Wikipedia is editable by anyone without even requiring them to log in, this situation may not be preferable for everybody. In fact, most of the third-party MediaWiki-based sites require you to at least log in before you can make edits. In the same spirit, MediaWiki provides the following types of user rights profile...

- Traditional Wiki: Allows anyone to edit without even logging in, Wikipedia style.
- Account Creation Required: Still open for anyone to edit, but will be required to register and log in with the wiki site.
- Authorized Editors Only: Pages can be viewed by everybody, but can be edited only by the approved users.
- Private Wiki: Only approved users can edit or even view the wiki.

Copyright and license:

- Creative Commons Attribution Share Alike
- Creative Commons Attribution Non-Commercial Share
- Creative Commons Zero
- Public Domain
- GNU Free Documentation License 1.2
- GNU Free Documentation License 1.3 or later
- No license footer
- Select a custom Creative Commons license

Copyright and License: This is where you can select the appropriate licence for your content. This will appear in the wiki's footer. If you are doing a large corporate deployment, seek the advice of the legal department to find out which option is right for you.

Email Options: This is where you can select the appropriate options for the email settings. Most of the default options are acceptable here.

Images and File Uploads: Here you can enable file uploads. You will also need to change the images' (subdirectory in the MediaWiki root directory) permissions to be writable by the web server. For a logo you need to specify a 135x160 pixel image.

Object Caching: If you are expecting high traffic for your wiki site you should select a proper caching system. For standard deployments you should get by using no caching.

9. Installation

In the next step you will be asked for the installation confirmation. Click continue. It will set up the database and tables, create a default wiki, create an admin account etc. Once you reach the completion screen, the installation wizard will ask you to download a configuration file called LocalSettings.php, which is created specifically for your installation. Download this file and place it in the root folder of the MediaWiki installation directory.

Complete!

Congratulations! You have successfully installed MediaWiki. The installer has generated a LocalSettings.php configuration file.

You will need to download it and put it in the base of the installation directory as index.php. The download should have a link below. If the download was not offered, or if you cancelled it, you can click the link below:

[Download LocalSettings.php](#)

Note: If you do not do this now, this generated configuration file will be deleted. You can download it later if you exit the installation without downloading it. When that has been done, you can enter your wiki site.

After placing the LocalSettings.php you can click the link to enter your wiki.

Your MediaWiki-based site is now ready for consumption. Users can now start creating pages on this wiki. Since we have selected 'Accounts Creation Required' as the user rights profile, users will have to register first to create or edit wiki pages.

As a first step you may want to log into your wiki site and modify the first page. Everything you know and love about Wikipedia is right there – for example, User Talk page, Discussion Page, History, My Contributions etc.

To create a new page just navigate to the URL – for example, mysite/LUDwiki/index.php/Main_Page/New_Page. You will be informed that there is no text for that page and will be given an opportunity to add text to it.

“MediaWiki has a large community based around it which builds a lot of cool extensions”

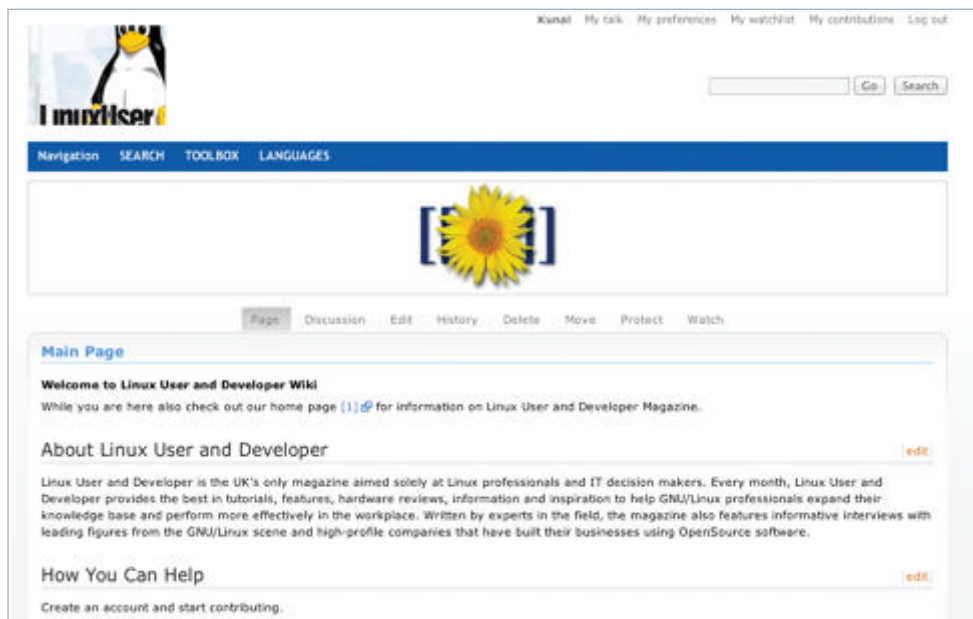
Customising MediaWiki

As mentioned earlier, one of the biggest strengths of MediaWiki lies in the extensibility of it. MediaWiki has a large community based around it which builds a lot of cool extensions.

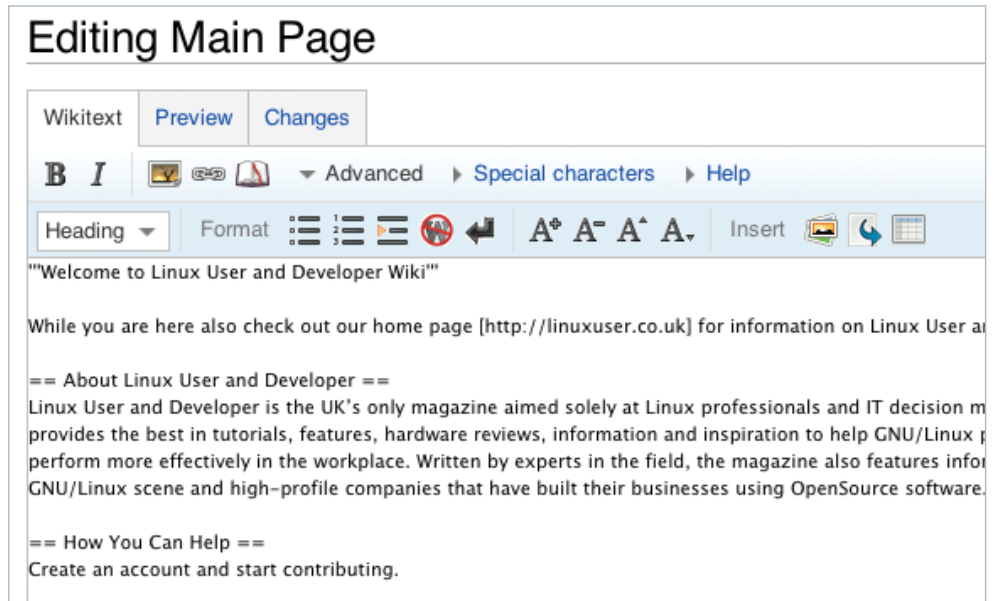
10. Using skins

Skins collectively drive the look and feel of the wiki site. MediaWiki ships with a few themes which are installed in the skins directory. To use any of the installed skins, you need to edit the LocalSettings.php file and change the value of \$wgDefaultSkin. You can also download MediWiki skins from various websites. Installation of a third-party skin is very simple. Download the ZIP archive, extract it in MediaWiki's skins folder, then change the \$wgDefaultSkin value to reflect the name of the downloaded skin. Also, this goes without saying: make sure you are adhering to the licence agreement of the skin you are installing.

Notice the logo – each skin has its own specification for the images and logo, which will need to be modified on a skin-to-skin basis. You should check the skin developer's website for more information.



■ LUDWiki with the gumaxdd skin applied to it



11. Using extensions

What is cooler than changing the look and feel of Wiki is adding actual functionality to it. For example, the standard MediaWiki installation comes with a very standard editor to enter text into it. There is not a WYSIWYG way of adding or editing text into it. Let's just change that.

A MediaWiki extension called WikiEditor provides a nice WYSIWYG editor interface for editing and creating wiki pages. You can download it from the Extension page located at www.mediawiki.org/wiki/Extension:WikiEditor. The Extension page will also

provide some important information such as the licence and the MediaWiki version it is compatible with. This page should also contain information on the installation of the extension.

Let's go through the extension installation now. Download and extract the archive into the extensions directory of the MediaWiki install directory. Now add the following line to the LocalSettings.php:

```
require_once( "$IP/extensions/WikiEditor/WikiEditor.php" );
```

To enable a preference by default but still allow users to disable it in preferences, add the following lines to LocalSettings.php:

```
$wgDefaultUserOptions[ 'usebetatoolbar' ] = 1;
```

```
$wgDefaultUserOptions[ 'usebetatoolbar-cgd' ] = 1;
```

```
$wgDefaultUserOptions[ 'wikieditor-preview' ] = 1;
```

The WikiEditor extension should be enabled in MediaWiki by default now. Try editing a page now and you will be presented with a more elegant way of editing text inside your wiki. That's just one extension, there are many more available at www.mediawiki.org/wiki/Extension_Matrix.

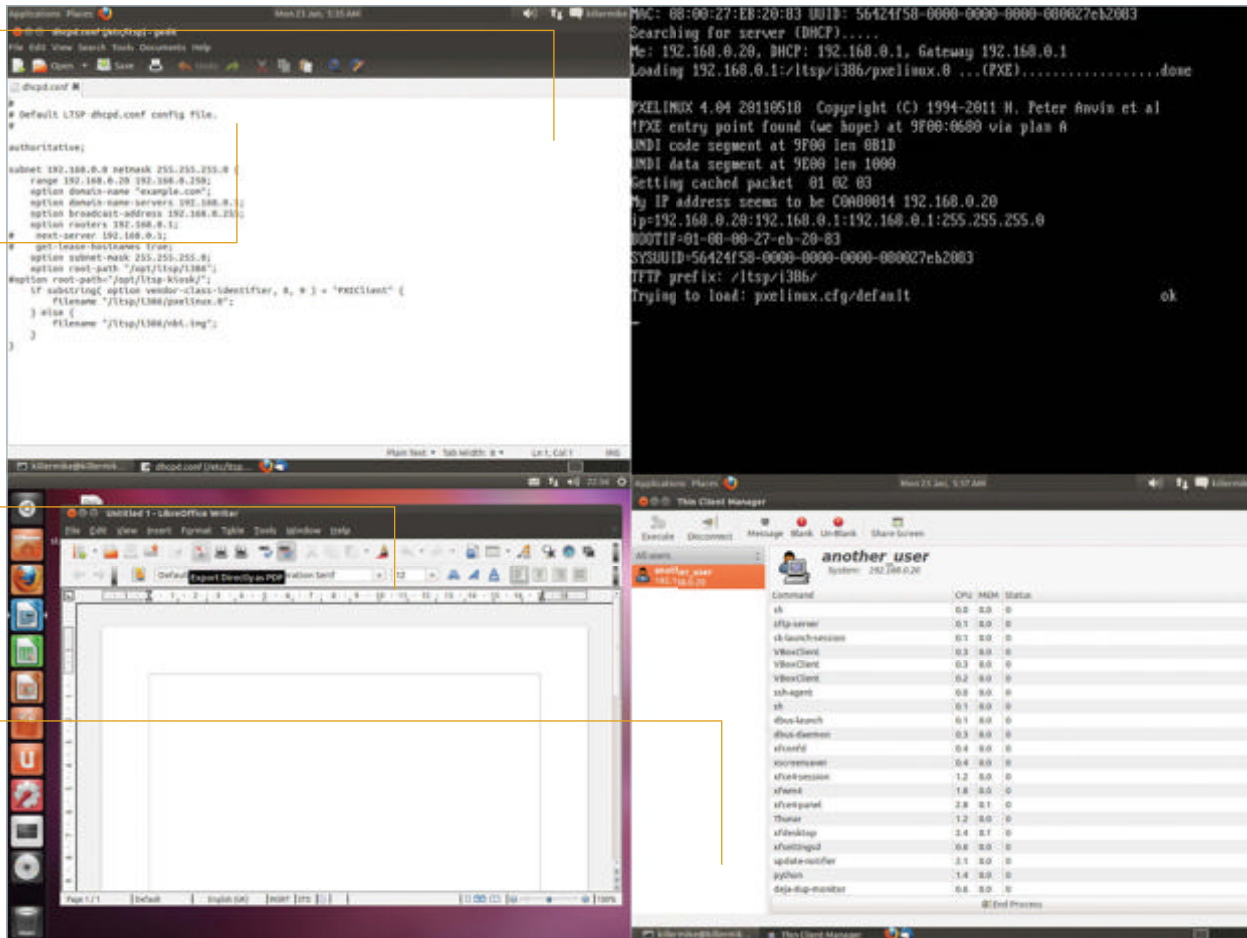
MediaWiki is not just the wiki that powers the world's biggest site, Wikipedia, but is also open source and free. But it does not end there. It is also the most extensible wiki software in this universe. If you are planning to roll out a wiki or even a fully fledged site, give MediaWiki a chance – it will blow your mind.

There are only four main text files that you'll have to edit, and we'll take you step-by-step through every stage of the process

Once everything is in place, we'll try booting over the network using PXE. DHCP tells the client where to look and NBD hands over the initial files

For everyday tasks, you can put together a client desktop that works in much the same way as a standard Linux one

Once you've got the basic desktop up and running, it's time to begin customising it and adding extra admin tools



How to distribute Linux desktops to thin clients

Ever wanted to serve Linux desktops to thin clients using open standards and free software? The Linux Terminal Server Project holds the key...



The process works as the server uses DHCP to assign an IP address to a client. DHCP also gives the PXE firmware on the network card the name of the files it needs to request from the server in order to boot. These are fetched using a file transfer protocol called TFTP. The first file is PXELINUX, the Linux network bootloader. Next, the Linux kernel is fetched along with a boot image that's decompressed into a RAM disk. Once the kernel has booted, a basic Linux

system is up and running. Subsequent file transfer to and from the client is accomplished via a more fully featured protocol called NBD.

An LTSP server contains a stripped-down Linux distro in a subdirectory of opt/ltsp/. Applications don't execute on the clients, they run on the server and pass graphics, sound and user input over the network. It's not hard to see why LTSP is finding particular favour in education these days, but it's also suitable for scenarios such as office installations.

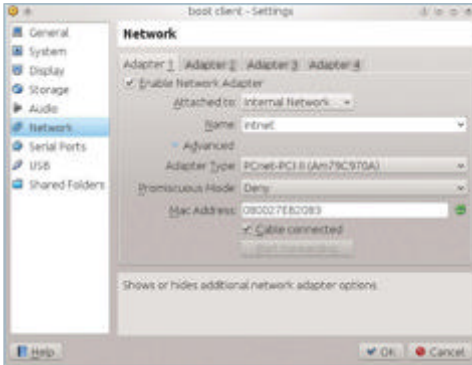
Advisor

Michael Reed is a technology writer and he's been hacking away at Linux for over 15 years. He specialises in desktop Linux solutions.



01 Start with a VM

If you're new to LTSP, you might want to start experimenting with a virtualiser. Most such software offers the option of setting up an internal network and booting one virtual machine from another. See the documentation for your virtualiser.



02 Set up the network

It's possible to get LTSP working with a single network card, but these examples presume that the server has two network cards so that you don't have to disable DHCP on your router. The server runs its own DHCP server for the thin client network.



03 Build a server

The LTSP docs recommend that a server should have at least 256MB with 192MB per client that it supports. With a 2GHz processor and 100BASE-TX Ethernet, a server with 4GB of RAM should be able to handle 20 clients for a mixture of web and office tasks.



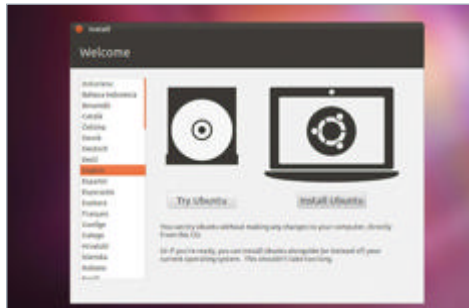
04 Select clients

Although the server needs to be fairly substantial, LTSP is a great way of breathing life into older hardware such as 800MHz PCs that are otherwise going to be binned. 256MB of RAM is recommended. You can also use LTSP with purpose-built thin clients.



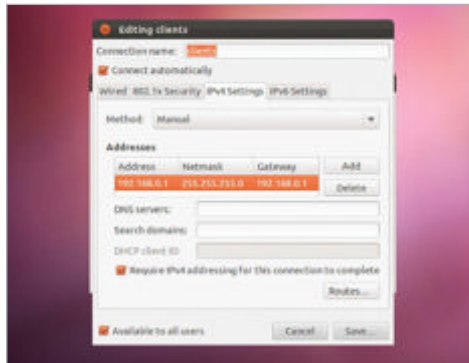
05 Install Linux

Most Linux distros can support LTSP, but their implementation of it varies. Debian uses NFS as the file transfer protocol whereas Ubuntu uses NBD for performance reasons, for example. These examples use Ubuntu 10.11. Check the documentation for your distribution.



06 Configure your network


Manually assign an IP address to your second network card (eth1). If your first card is on the 10.x.x network, put the second adaptor onto IP address 192.168.0.1 with netmask 255.255.255.0. Use the UI or edit the /etc/network/interfaces file.



“It’s usually better to get a basic system working at first and to add the extra tools after”

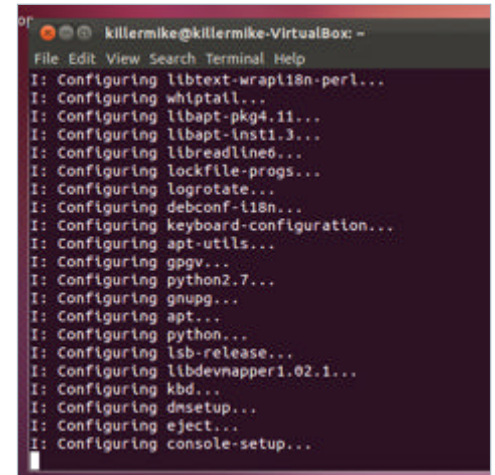
07 Install LTSP

Install LTSP itself by typing 'sudo apt-get install ltsp-server-standalone ltsp-docs' or by using a package manager. This will fetch all the parts you need in order to begin work. It's usually better to get a basic system working at first and to add the extra tools afterwards.



08 Build the LTSP environment

Type 'ltsp-build-client'. This fetches the packages that constitute the small Linux distro that is served to the clients. Use the --arch parameter to specify different client architectures. Without this switch, the script will fetch the files for the same architecture as the server. Use the 'man ltsp-build-client' page for all options.



“There are myriad Linux administrative tools to do things such as setting up whole groups of users or restricting their privileges”

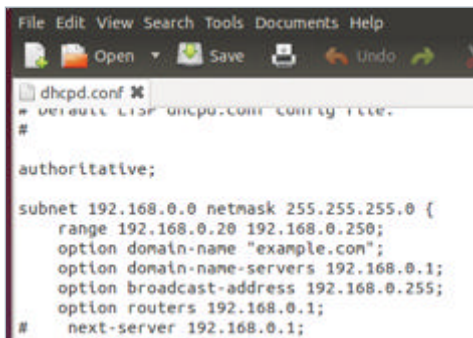
09 Edit /etc/dhcp/dhcpd.conf

There are two versions of dhcpd.conf you'll need to edit. Start by typing 'sudo gedit /etc/dhcp/dhcpd.conf' to edit the first file and add the line 'include "/etc/ltsp/dhcpd.conf"' to the bottom.



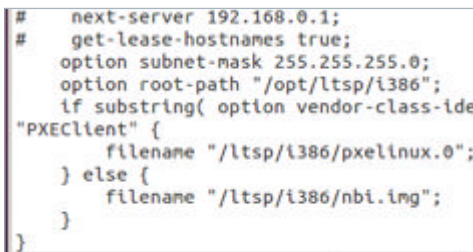
10 Edit /etc/ltsp/dhcpd.conf part 1

Now edit '/etc/ltsp/dhcpd.conf'. Towards the top of the file you have to make sure that everything points to your second network card. For example, if your card was on 192.168.0.1, make sure that the subnet is 192.168.0.0.



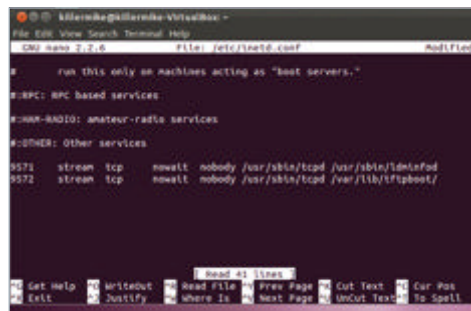
11 Edit /etc/ltsp/dhcpd.conf part 2

The lower half of this file concerns the pathnames which are going to be used to boot the client. The two files that we're interested in, pxelinux.0 and nbi.img, are both located in (as far as TFTP is concerned) /tftp/i386/.



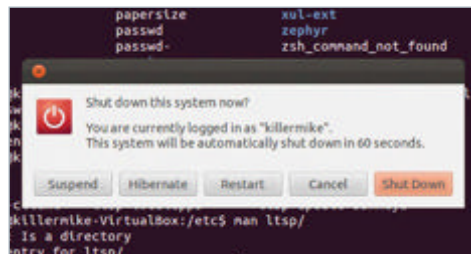
12 Configure TFTP

The file that you need to edit is '/etc/inetd.conf'. Change the last entry in the final line to point to your TFTP root – which, on Ubuntu, is '/var/lib/tftpboot/'.



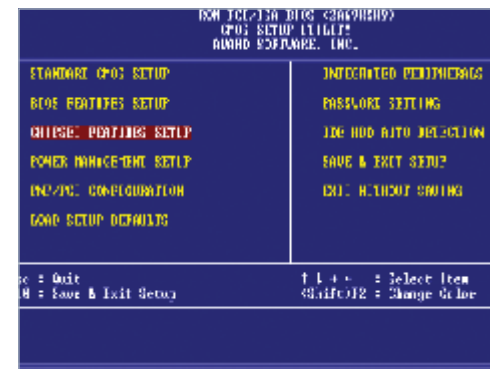
13 Reboot the server

Everything should now be configured for a basic LTSP setup, and all of the services that you have configured need to be restarted. It's easiest to reboot the machine. Do 'sudo ltsp-update-sshkeys' any time you alter the IP settings of the server.



14 Configure the client

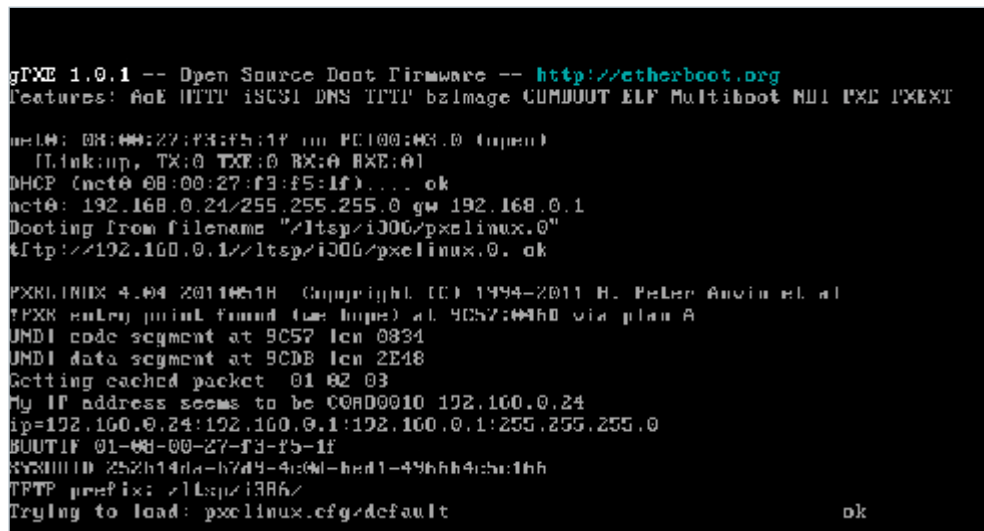
Hopefully, the client itself won't need much configuration as all of the options are selected on the server. The one thing you do need to do is enter the BIOS setup (usually by pressing Del or F2 on startup) and making network boot the default.



15 Troubleshooting (if the client doesn't boot)

DHCPD – Most clients will give verbose feedback when beginning to netboot. Check the reported paths. Type 'ifconfig -a' on the server and see if the client is discovering the MAC address of the server.

TFTP – Execute 'tftp localhost' on the server and try to fetch the PXE Linux loader manually by issuing 'get pxelinux.0' to see if TFTP is working. If the connection is made but you can't fetch the file, check the paths in '/etc/inetd.conf'. Now try fetching from a machine on the client network.



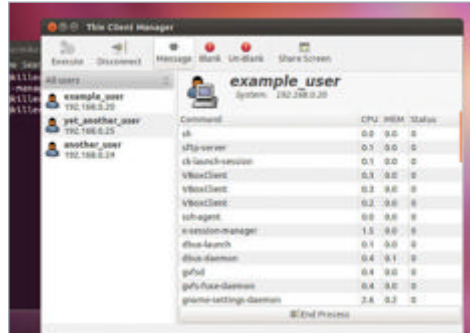
16 Log in from a client

If everything has gone according to plan you should now be able to log into your client. To test that everything is working, use the user name and password that you normally log in to the server with. Notice that it defaults to LDM, the LTSP login manager.



19 Add GUI management

Consider adding Thin Client Manager via the package manager. It allows you to monitor logged-in users and keep an eye on the processes that they are running. You can also message multiple users (useful for system-wide notices), disconnect them, lock the screen or disconnect them.



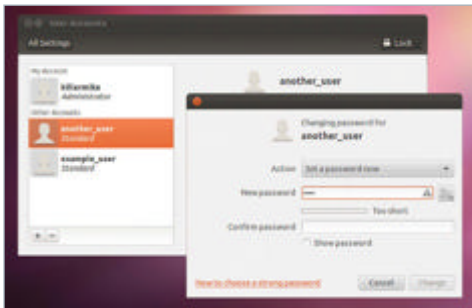
22 Evaluate ready-made LTSP distros

Edubuntu, an official member of the Ubuntu family, is an education-orientated distro and it contains a ready-to-run LTSP server. Distros like this one and K12LTSP come into their own for fast deployment and demoing what LTSP can do.



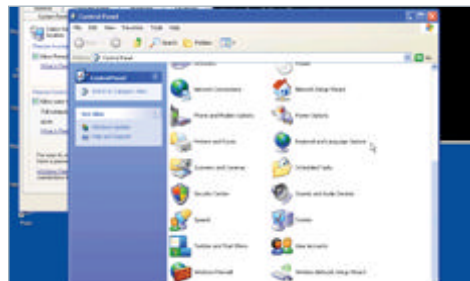
17 Set up additional users

Add and administrate users on the server machine in the same way that you would with any Linux box. Of course, there are myriad Linux administrative tools to do things such as setting up whole groups of users, restricting their privileges or customising application access.



20 Test remote desktop support

The LDSP client can also function as an RDP client. Start by creating the client configuration file: 'sudo nano /var/lib/tftpboot/ltsf/i386/lts.conf'. Simply add a section header '[DEFAULT]' and then add the line 'RDP_SERVER="[ip address of server]"'.




23 Test web browsing

As the programs are executing on the server rather than the client, they have the same web access possibilities as the server. Flash works too, but bear in mind that watching video over the web is a resource-intensive activity.




18 Add multiple architectures (AMD64, PowerPC Mac etc)

Build a separate environment with, for example, 'sudo ltsf-build-client --arch ppc'. Then open up '/etc/ltsf/dhcpd.conf' and add 'if substring (option vendor-class-identifier, 20, 3) = "ppc" { option root-path "/opt/ltsf/ppc"; }' to add support for PowerPC Macintosh computers.



21 Test multimedia

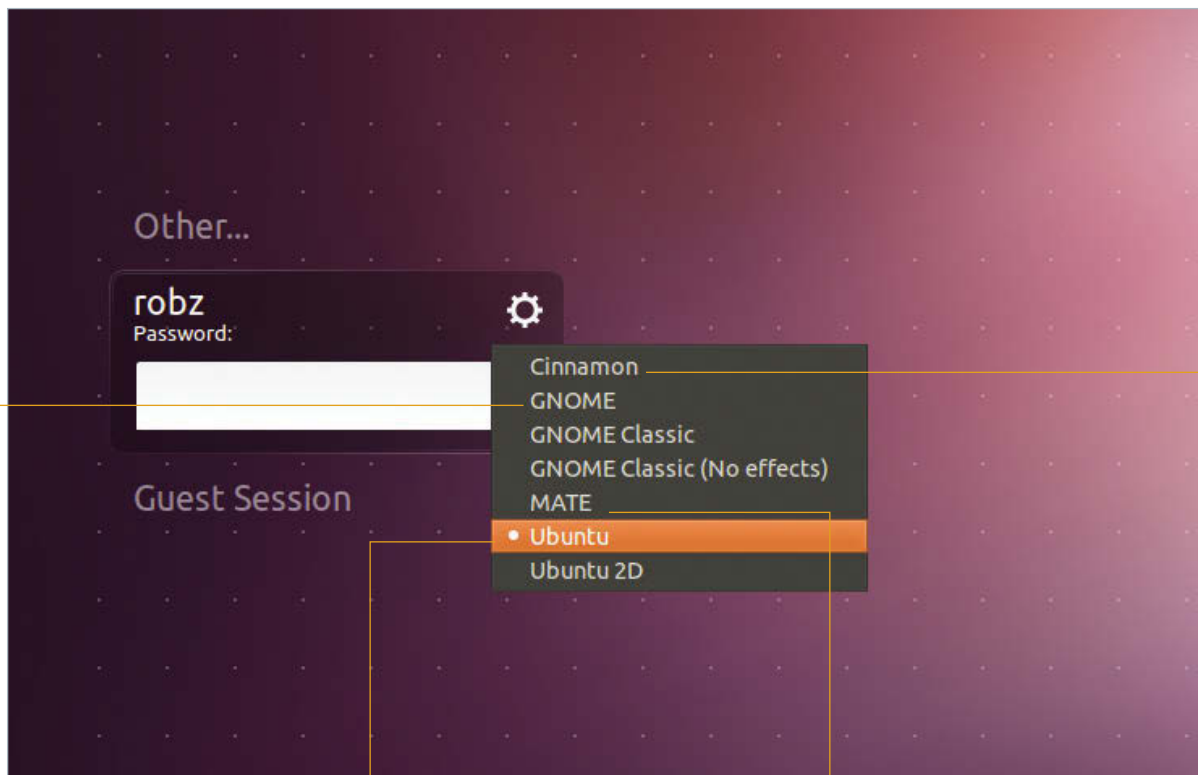
Multimedia performance is an area in which a thin client solution will never be competitive with a dedicated, standalone PC. On the plus side, LTSP gives most of the same facilities that standalone machines would have, while PulseAudio provides sound support to clients.



24 Tweak the client

A lightweight desktop environment works far better over LTSP than a heavyweight such as the latest GNOME or KDE desktops. Install a replacement DE such as Xfce or LXDE on the server in the normal way and set it as the default for clients (type 'man "lts.conf"').





GNOME Shell is the evolution of the popular GNOME desktop, dropping the bottom bar and creating an overlay for running and available programs

Unity is the standard Ubuntu desktop. Canonical's design philosophy for Unity is an environment that can be used for any application

MATE is a project to maintain and update GNOME 2 for those who prefer the veteran desktop environment

Cinnamon was built by the Linux Mint team as an alternative to GNOME 3 for those who prefer the older GNOME 2

Replace Ubuntu's Unity desktop environment

Love Ubuntu, but can't stand Unity? Here's how to configure your system to run MATE, GNOME Shell or Mint's Cinnamon at login instead

Advisor

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

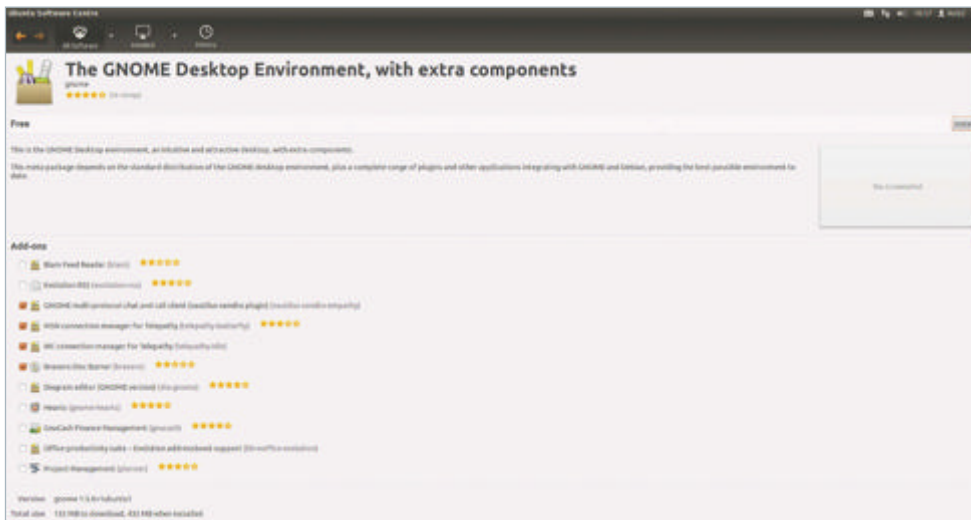


One of the greatest parts about using Linux operating systems is the sheer amount of choice you have in almost every aspect of your user experience. The desktop is one example of this, with each environment offering a different layout or promising to be a lightweight alternative.

Vanilla Ubuntu of course comes with Unity, Canonical's 'one size fits all' UI, and the main

Ubuntu spin-offs are notable for their different desktop environments. Kubuntu has KDE, Xubuntu has Xfce, and so on. There are a few popular desktop types that don't have their own punny distro name, though, and you can get them installed onto Ubuntu without losing Unity.

In this guide we'll show you how to set up and configure some of the more popular desktop environments installed – such as the post-modern GNOME, the nostalgic MATE and newcomer Cinnamon – and how to use them in conjunction with Ubuntu and Unity for a seamless user experience.




01 Install GNOME Shell

GNOME is already in the Ubuntu repository, and you can either install it from the Ubuntu Software Centre or open a terminal **and type:**

```
sudo apt-get install gnome-shell
```

During the install it may ask which Display Manager to use, gdm or lightdm. Select lightdm.



02 Boot into GNOME Shell

After the installation has finished, log out or restart your computer. At the login screen, click your username. A drop-down menu will contain the options for which desktop to use, such as Ubuntu (Unity), GNOME (GNOME 3) or GNOME Classic (GNOME 2-esque).



03 Uninstall GNOME Shell

If you decide that GNOME Shell is not for you, you can always remove it. Boot into Unity from the login screen and **open the terminal:**

```
sudo apt-get remove gnome-shell gnome-session-fallback
```


And to remove any leftover packages:
`sudo apt-get autoremove`

04 Install GNOME Classic only

There's also a version of GNOME Shell that emulates GNOME 2-style menus. To use this, go into the terminal **and type:**

```
sudo apt-get install gnome-session-fallback
```

Once installed, log out and select GNOME Classic from the drop-down menu.




05 Uninstall GNOME Classic

GNOME Classic can be uninstalled whenever you want; just remember to remove it from a different desktop environment. **In the terminal, type:**

```
sudo apt-get remove gnome-session-fallback
```

And to remove any leftover packages:
`sudo apt-get autoremove`



06 Add the MATE repository

To install MATE, you will first need to add a relevant repository. These can be added in the terminal using the **following commands.**

```
sudo add-apt-repository "deb http://packages.mate-desktop.org/repo/ubuntu oneiric main"
```

Follow up with 'sudo apt-get update' to get MATE added to the packages list. MATE is currently only available for 11.10.

07 Install MATE


Once the update is complete, it will take two commands in the terminal for MATE to properly install:

```
sudo apt-get install mate-archive-keyring
```

followed by

```
sudo apt-get install mate-core
```

The keyring allows the whole of mate-core to install on the system, which includes about 80 packages.



08 Boot into MATE

To boot into MATE, first log off. When selecting your username, find the drop-down menu with the Desktop Environments and change it from Ubuntu to MATE. The last desktop environment you selected will be automatically used the next time you log in.

Unity tweaking

Get extra control over your Unity environment with MyUnity

A lot of the extra customisation features found in the other desktop environments might be attractive to some users; however, these kind of settings can be accessed for Unity using the MyUnity tool. As well as basic appearance tools, the program includes options to change the autohide behaviour, how mountable devices show on the sidebar, along with transparency and colour settings.

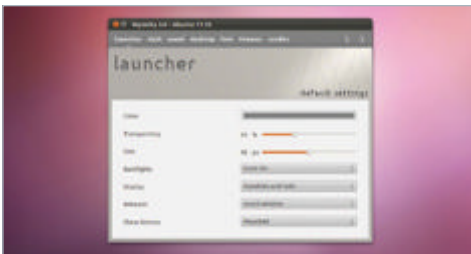
MyUnity is readily available in 12.04, but to install it on 11.10, first add the repository with the following command:

```
sudo add-apt-repository ppa:myunity/ppa
```

To get MyUnity added to the package list, type 'sudo apt-get update' and then install with:

```
sudo apt-get install myunity
```

There are a lot of options in MyUnity to fiddle with, allowing you to have Unity the way you want it.



■ MyUnity is a powerful tool that opens up Unity



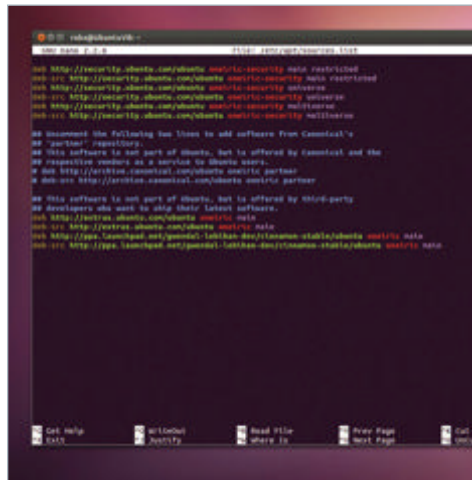
09 Remove MATE

Uninstalling MATE is simple. Boot into Unity, and open up the terminal:

```
sudo apt-get remove mate-core
```

APT will alert you to the packages that are no longer required now mate-core is gone; to uninstall these, use:

```
sudo apt-get autoremove
```

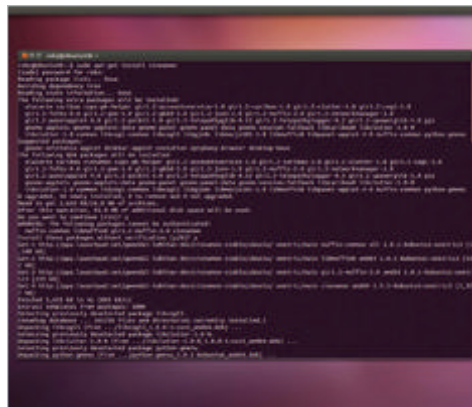


10 Adding the Cinnamon repository

Cinnamon is distributed via a PPA for Ubuntu 11.10 and 12.04; to add it to the Ubuntu repository, type into the terminal:

```
sudo add-apt-repository "deb http://ppa.launchpad.net/gwendal-lebihan-dev/cinnamon-stable/ubuntu oneiric main"
```

If you're using 12.04, replace 'oneiric' with 'precise'. Type 'sudo apt-get update' to update the list of software ready for install.



11 Install Cinnamon


To add Cinnamon, type into the terminal:

```
sudo apt-get install cinnamon
```

It may ask if it's okay to install the packages depending on the state of the PPA, but this is the official repository so allow the install.

12 Boot Cinnamon

To boot into Cinnamon, you'll need to log out. At the login screen, click the gear next to your username and select Cinnamon before logging back in. You may also notice that GNOME Classic is available; this is installed with Cinnamon.



13 Uninstall Cinnamon

To completely uninstall Cinnamon, you'll need to add an extra step to the process. Start by booting into Unity Desktop, and opening the terminal:

```
sudo apt-get remove cinnamon
```

Then remove GNOME Classic:

```
sudo apt-get remove gnome-session-fallback
```

And finally clean up your system:


```
sudo apt-get autoremove
```

14 Preparation for Mint GNOME Shell extensions

MGSE first requires GNOME 3 to be installed, so follow the previous steps for that. You'll need to add the Linux Mint repository like so:

```
sudo add-apt-repository "deb http://packages.linuxmint.com/lisa main upstream import"
```

And then 'sudo apt-get update' to add the new packages.



15 Install MGSE

Open the terminal and type:

```
sudo apt-get install linuxmint-keyring
```

```
sudo apt-get install mint-meta-mgse
```

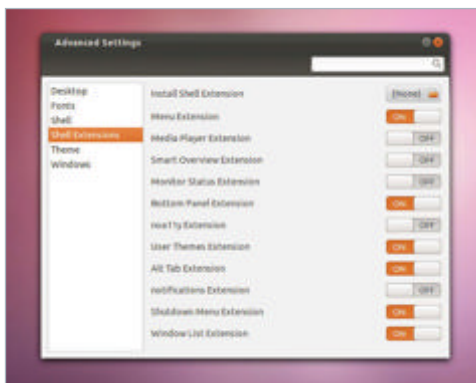
Boot into GNOME 3 desktop, or restart GNOME if you're already in it, and open Advanced Settings from



“The last desktop environment you selected will be automatically used next time”

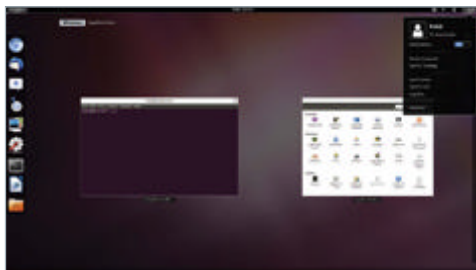
the applications list. Go to Shell Extensions to find all the extra options to activate MGSE.

The three main extensions to activate are Menu Extension, Bottom Panel Extension and Window List Extension. The new features will activate on the fly; however, to get them to work fully you will need to restart GNOME again.



16 Other MGSE extensions

There are other great features included in MGSE, such as a more classic Alt Tab extension, an icon to help manage multiple monitors, and the ability to add a 'shut down' option to the user menu. These and any other extensions can be turned off at any time.

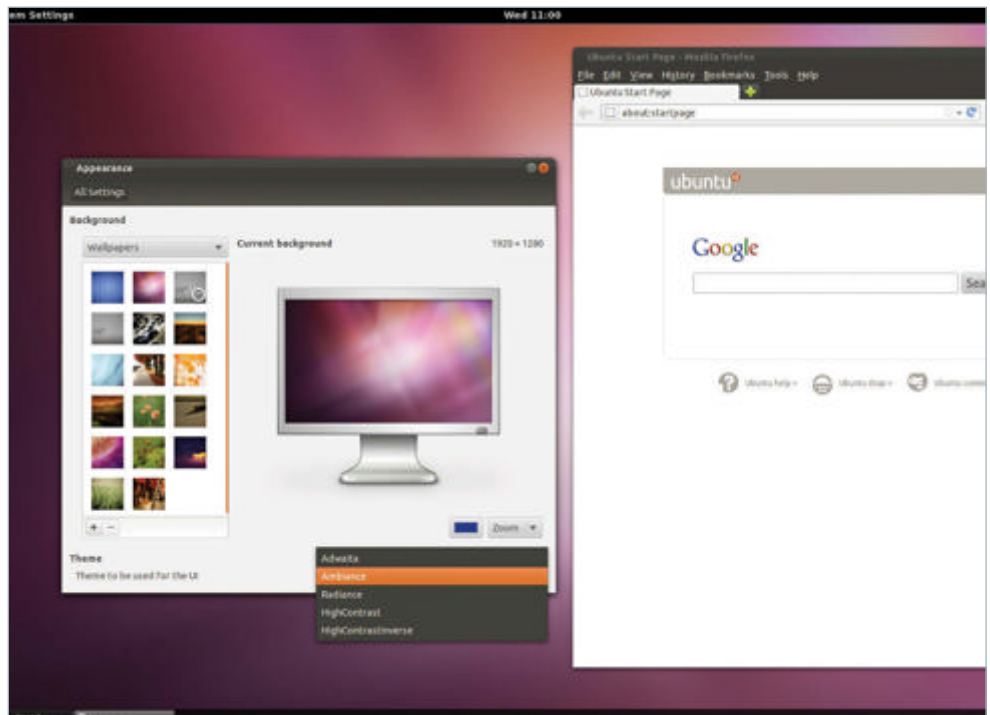


17 Remove MGSE

If vanilla GNOME 3 is more your thing, you can always remove MGSE. Turn off the extensions and reboot GNOME, then go into the terminal and type:

```
sudo apt-get remove mint-meta-mgse
```

Follow this up with 'sudo apt-get autoremove' to lose the unnecessary extra packages.

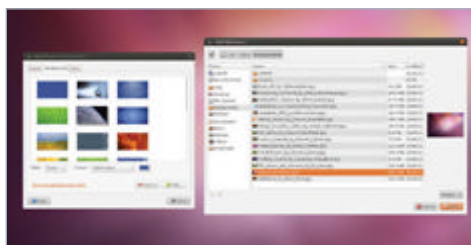


18 Ubuntu theme – basics

All the desktops will use their own standard theme, but if you miss the orange and black icons adorning your windows, the 'Ambiance' Ubuntu theme is available in the Appearance menu in System Settings. You may have to log out and back in for the changes to take effect.

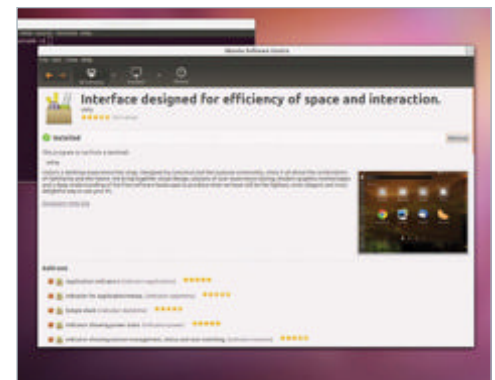
19 Ubuntu theme – GNOME and Cinnamon extras

The theme in Cinnamon can be accessed via the Cinnamon Settings menu, in Themes. To add the minimise and maximise buttons to GNOME, go to Advanced Settings>Shell and click on the Arrangement drop-down menu to change from Close Only to All.



20 Ubuntu theme – wallpaper

The normal Ubuntu wallpapers are kept in /usr/share/backgrounds. MATE will launch by standard with its own wallpaper, and you'll need to add the standard background from the Appearance Preferences menu if you want the desktop to match the rest of the theme.



21 Remove Unity

If you don't plan to use Unity again, you can remove it from your machine. Either do this from the Software Centre, or go into the terminal and type:

```
sudo apt-get remove unity unity-2d
```

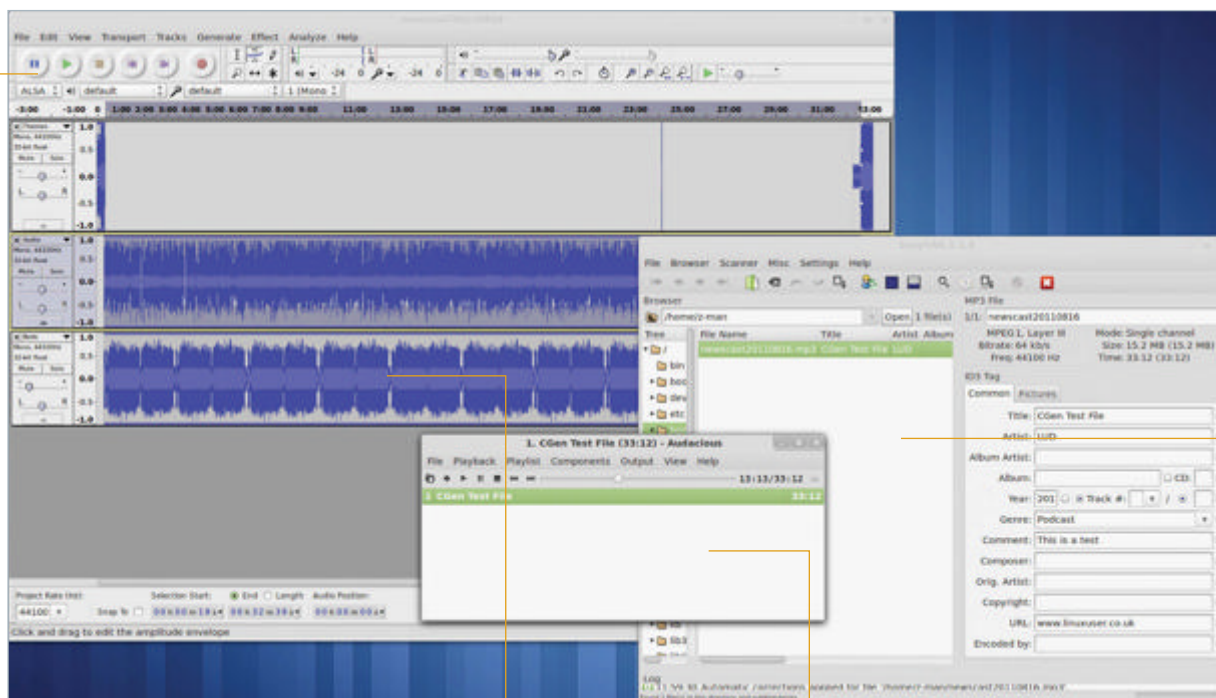
This will remove the Unity and Unity 2D packages, and 'sudo apt-get autoremove' will get rid of the now unused dependencies.

22 Recover Unity

Unity is easily reinstalled – and while you can do it in the Software Centre, the best way is via the terminal.

```
sudo apt-get install unity unity-2d
```

You don't have to install Unity 2D. Also, as the settings are saved separately, this will restore your menu as well.



Audacity is free and open source, and a very powerful piece of software. The newly released Audacity 2.0 is fantastic

Audacity has all the right tools to create a professional-sounding podcast, with plenty of great and intuitive visual editors

Having your say in the form of a podcast is a great way to express yourself, and Linux has all the tools to make this happen

EasyTAG gives you complete control over ID3 tags, including the addition of image tags to the files

Professional podcasting with Linux

You too can podcast like a pro using open source software on Linux, thanks to the new version of Audacity

Advisor

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

Resources

Audacity: <http://audacity.sourceforge.net/>
EasyTAG: <http://easytag.sourceforge.net/>



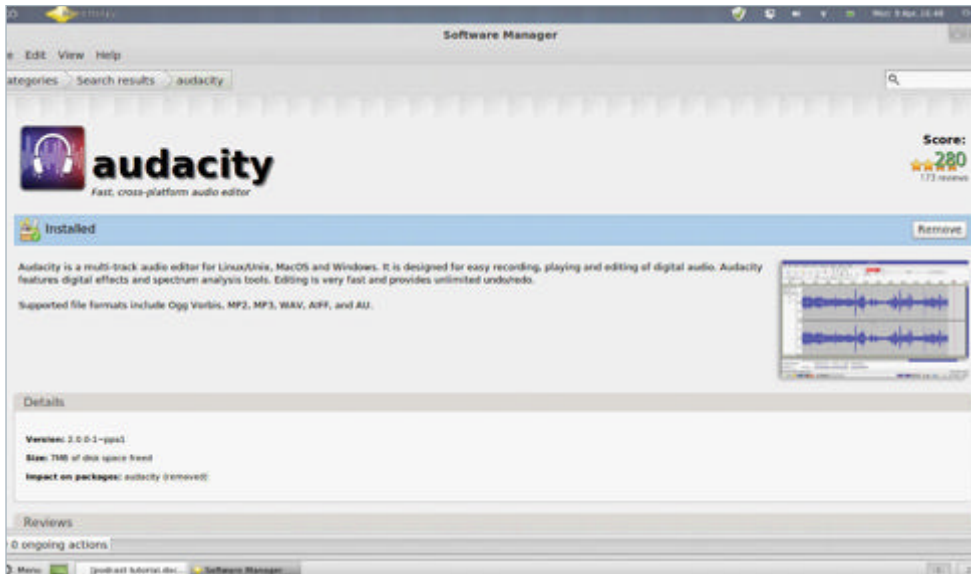
Recently, an article on the **BBC website put forward the question: is podcasting a fad that is already over?**

It pointed out how Ricky Gervais's podcast was talked about in the mainstream media quite often a few years ago, but you barely hear about podcasts any more. However, it concluded, quite rightly, that just because you weren't hearing about them, it didn't mean that they weren't there.

We live in this awesome future of interconnected, mobile PCs that can

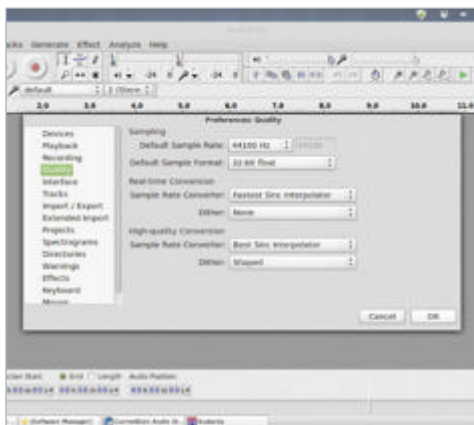
automatically organise these apparently unheard-of podcasts into a playlist of non-stop aural content. If you're already on the podcast bandwagon, you might be thinking about how to start your own. The number one question, though, is how?

We'll give credit where credit's due, GarageBand on Mac will do a fine job of a podcast. In fact, most of the podcasts you listen to are probably created on it. However, for both an open source and superior method, we can turn to Audacity.



01 Install Audacity

Either visit the Audacity website for the source code, or check your package repository. Audacity recently upgraded from 1.3.14 Beta to a full 2.0 stable release, although essentially they're the same. However, you want one of these and not the older 1.2.



02 Set up Audacity

When making a podcast, it helps to consider the widest possible audience. Luckily, you can do this without making huge sacrifices. To start, go to Preferences, then Quality and change sample format to 16-bit. Keep the sample rate at 44100Hz.

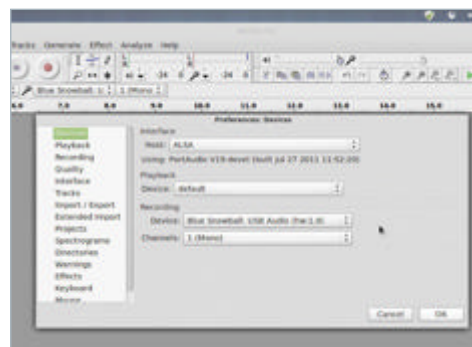
03 Adding libraries

While some package repositories will include Audacity with LAME MP3 and FFmpeg libraries, compiling from source will not. In Preferences, locate the Library tab and follow the instructions to install and link both to Audacity. Now we can import relevant files, and output usable files.



04 Get a microphone

There are two types of microphones: dynamic and condenser. Dynamic mics are used in noisy crowded situations like a concert or conference, while condensers are suitable for a studio environment with little outside noise. We recommend USB mics, such as the Blue Snowball, for many reasons.

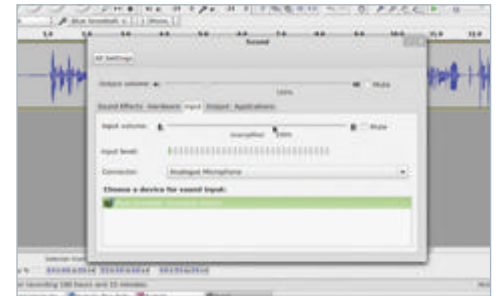


05 Connecting a microphone

Once you have the mic connected, select it from the drop-down menu (or set it as default in Preferences) and set the input channels to mono. Using monaural helps keep the quality up without wasting file size on two channels of the same audio.

06 Select a level

It's a good idea to test the levels of the microphone. This can either be done by the volume slider at the top right, or in the system-level sound settings. Test it by talking at a normal tone; the peaks should be just visible.




07 Optional extras

At this point you might want to consider pop screens – thin fabric meshes that reduce the popping and hissing from certain vocal sounds – or a shock absorber. Vibrations can affect quality and cause extra noise on an audio track that a shock absorber can help negate.

08 Ready to record

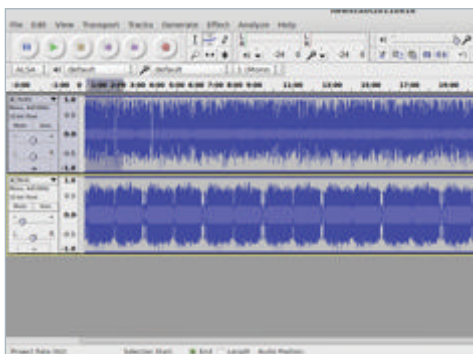
Get yourself comfy and hit Record. Make sure you keep roughly the same distance away from your mic while recording, so that the volume stays fairly level. If you can grab a co-host, it'll help keep the conversation more lively.



09 Recording controls

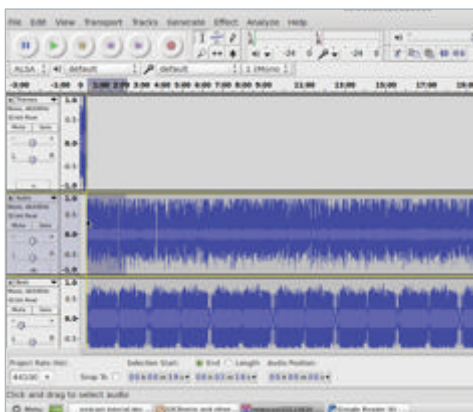
While recording, if you hit Stop and then start recording again, a new track will be created. If you hit Pause, you will be able to keep the recording going on the same track once you unpause. The bar at the top will help you keep track of time.

“If you can grab a co-host, it'll help keep the conversation more lively”




10 Edit the podcast

Once you've finished recording, it's time to edit. Audacity has plenty of great visual tools to aid in this, starting with the Selection tool. A highlighted section can be deleted, separated, or have effects applied. A yellow bar will indicate the boundaries of other tracks.



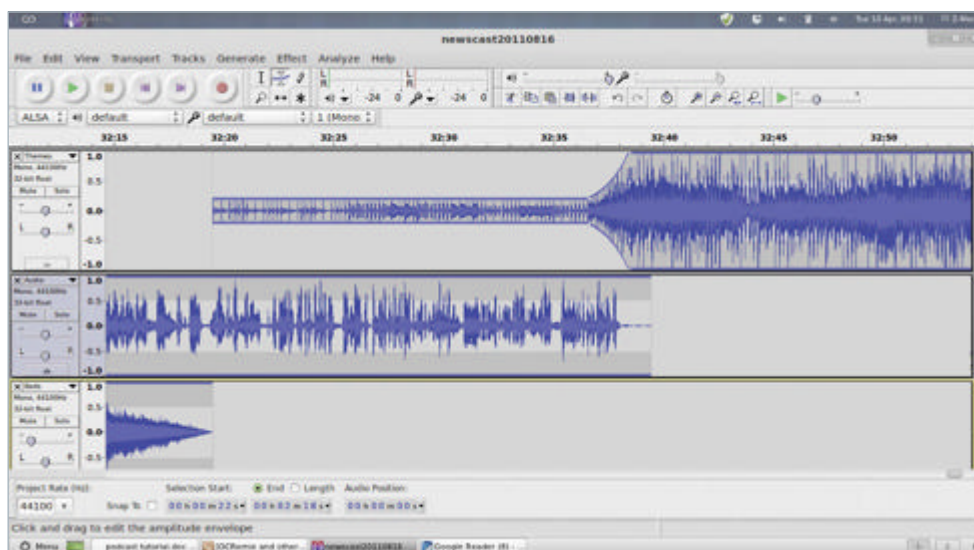
11 Post-processing

Using File>Import>Audio, you can add beds (background music), theme tunes or other sounds to your project. The Timeshift tool can be used to move audio up and down a timeline, with the yellow bars again indicating when you're at the boundaries of other audio.




12 Consolidate tracks

You can have more than one piece of audio on a track – you can copy and paste selections to any point in another track, or drag the audio using Timeshift. Try to keep as few tracks as possible; this aids in exporting later on.




13 Track fades

You can use the Envelope tool to modify specific parts of the audio's volume, setting lows and highs to create volume-changing effects. You can also apply the Fade Out effect to a selection of audio, or change the overall volume of the entire track from on the left of the track.



14 Special effects

Audacity comes with effects that can be useful for podcasting, like the self-explanatory Repeat and Amplify. Noise Removal can be used to analyse a silent section of audio and then apply a filter to a track to remove unwanted background noise.

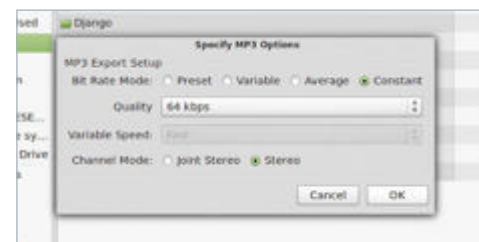


15 Testing, 1, 2

After editing, give it one last listen – to catch any interference, unnecessary pauses and quiet parts to be fixed. It's also good to get a gauge on how the podcast works in case you need to change anything the next time you record. You can individually mute or single out tracks, and start from any point in the timeline if you want to double-check anything in particular.

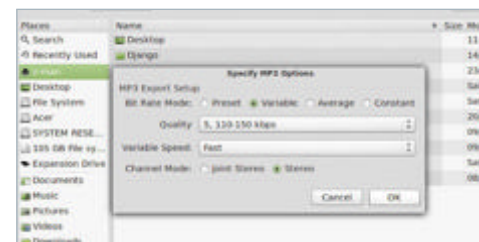
16 Prepare to export

If you haven't yet, make sure all the tracks are mono. Select the track and then go Tracks>Stereo Track to Mono. Audacity will not output a mono file unless all the tracks are already monaural. This helps keep down file size.



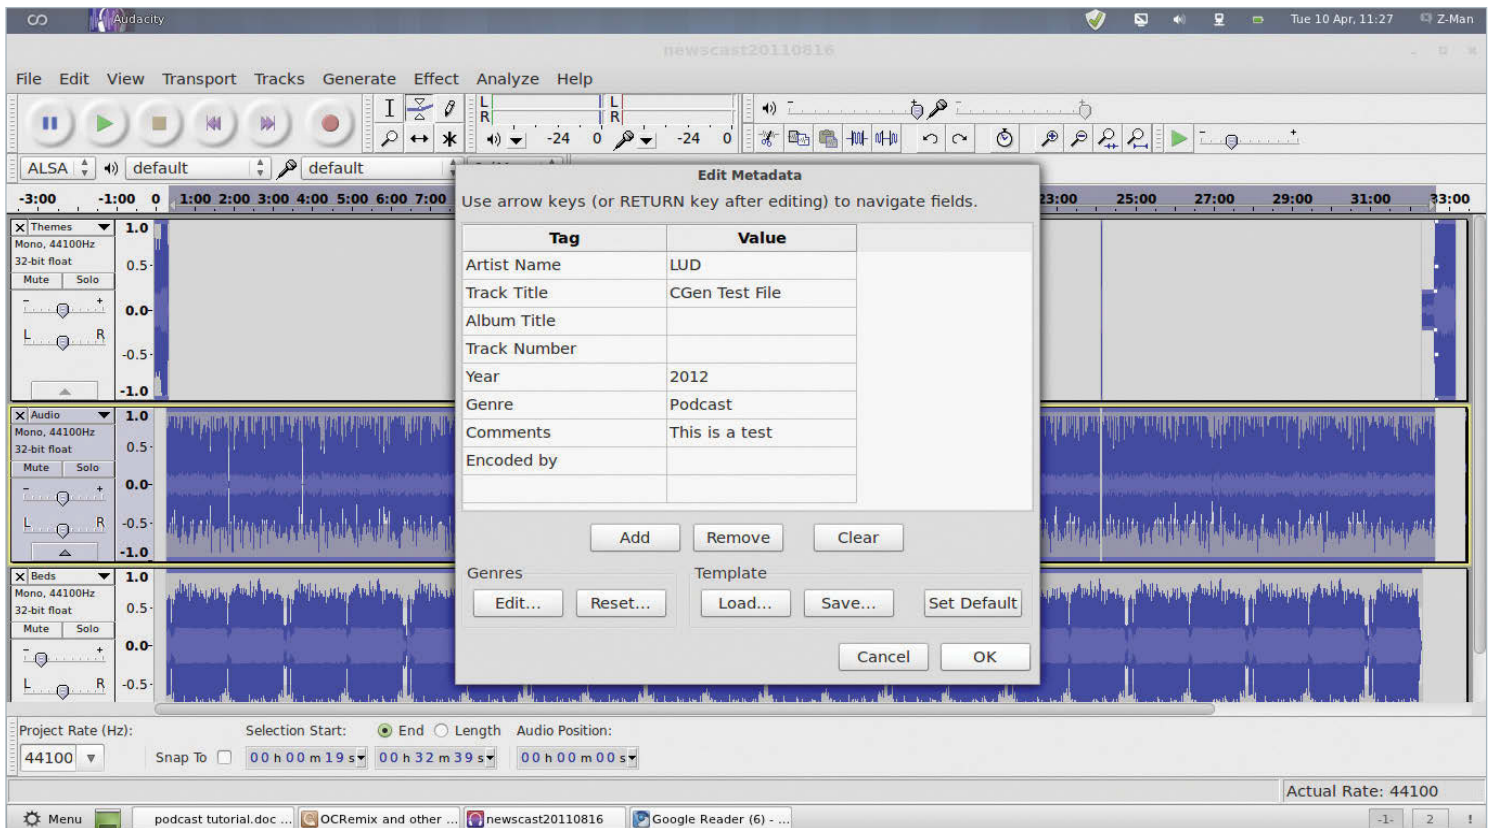
17 Export options

Click File>Export to begin the process. We recommend exporting as an MP3 file, again to make sure it reaches the widest audience. A 64kbps mono file will have the same quality as a 128kbps stereo file, and output at about 30MB per hour of audio.



18 Advanced exporting

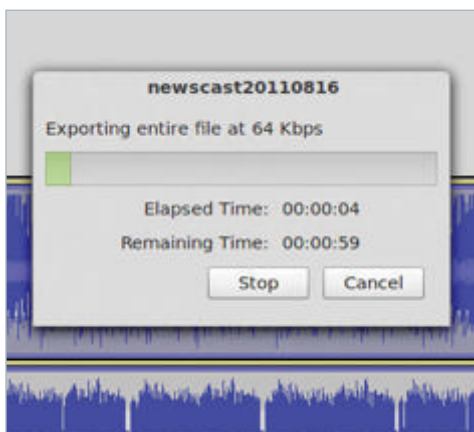
If you want to play with increasing the audio quality and keeping the file size down, you can try using a variable bitrate (VBR). This will change the quality so that it's lower during silence, and higher when the sound is complicated. Some audio players cannot track through VBR properly, though.



19 Meta

When you've chosen a file type and a location to save to, Audacity will throw up a window to add some metadata, such as title, author, genre etc. A lot of portable media players benefit from having this relevant data available to help sort files.

“Audacity comes with effects that can be useful for podcasting”



20 Encode

If all has been done correctly up until this point, Audacity will let you know that it will be exporting as a mono file. The export time is affected by number of tracks, length of the podcast and available computer resources.




21 Tag, you're it

Audio files can be tagged with an image that shows up in audio players. Usually it's a logo, or a picture relevant to the discussion. EasyTAG is the best way to add this image to a file, so install it from the listed website or your package manager.

22 Image tag

Select the podcast to view ID3 details and go to the Pictures tab. Click the '+' button to browse for an image and then click the small square to the right of the image to add it to the tag, then save.



23 Advanced meta options

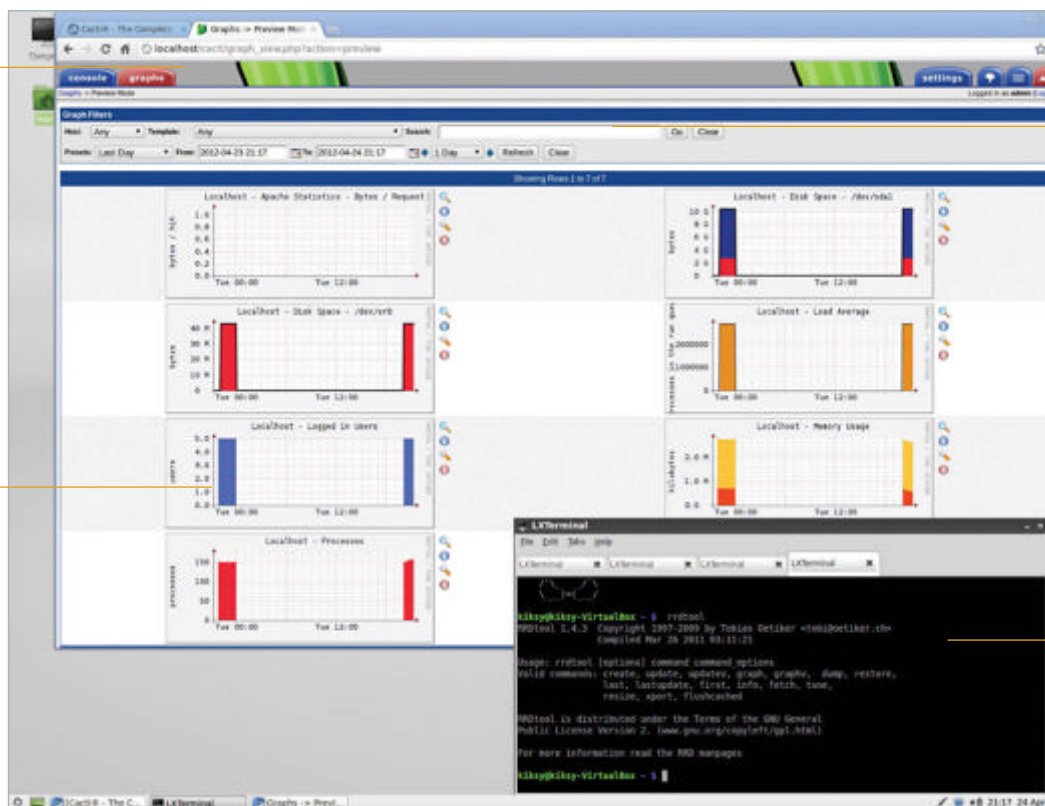
EasyTAG will sanitise your metadata by default, but you can edit or add more data to the common tags. Check your podcast in a media player, file browser, or restart EasyTAG to ensure the tags and images have been correctly saved.

24 Hello, and welcome...

You're done! Upload your podcast to a website, set up some feeds and submit it to any aggregation site you choose. WordPress has some great podcast plug-ins to aid in this as well. Tell your friends and work colleagues – and, most importantly, listen to feedback.

All you need to get Cacti up and running is a basic PHP/MySQL install; the main Cacti package has an easy-to-follow wizard

Graphs can be updated at any interval you choose, although 1 – 5 and 15 minutes are the default



Graphs can be automatically cron'd to upload to an FTP, or exported as CSVs

Cacti runs as an easy-to-use web interface for the popular RRDTool, allowing easy viewing of graphs and data

Monitor and graph your network statistics

Create graphical representations of network data statistics in easy steps

Advisor

Kieron Howard First introduced to Linux eight years ago when installing Debian on his Xbox, web developer Kieron has continued to enjoy putting Linux on all sorts of devices



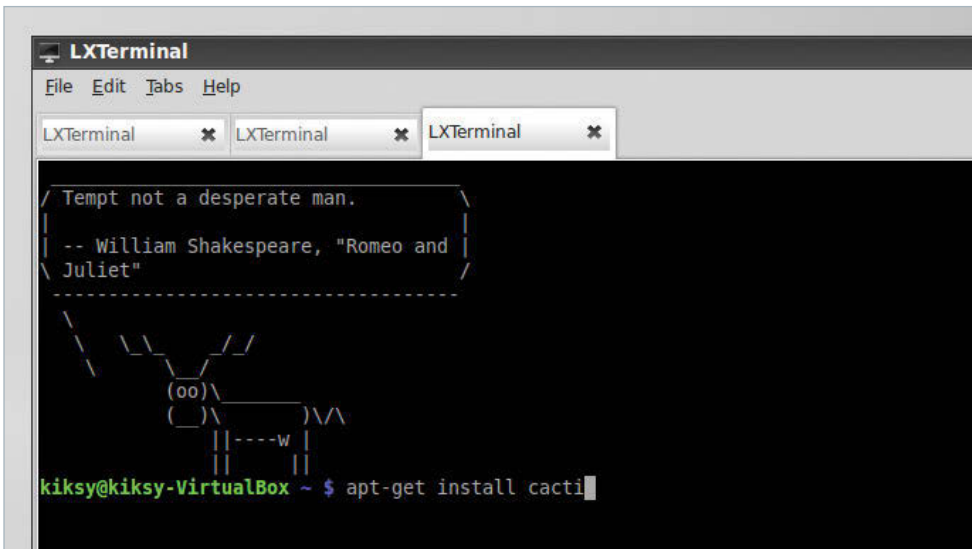
Network monitoring is something from which even administrators of fairly small LANs can benefit.

Being able to gather and display data on network and device traffic can provide information on how best to expand a network. It can also help to identify problems, sometimes


even before they occur. RRDTool has been around since 1999 and is widely used up to enterprise level for systems analysis. Cacti is a web-based front-end for RRDTool. While it's been around for over ten years, it has only reached version 0.8, but don't think it isn't ready for production – it's stable and has thousands of users. This tutorial will take you through setting up Cacti from scratch, then adding a device to be polled. Then we will add in a new template and output some graphs. Don't worry if you don't have any other devices that support SNMP, as we can use the local machine to test things out.

Resources

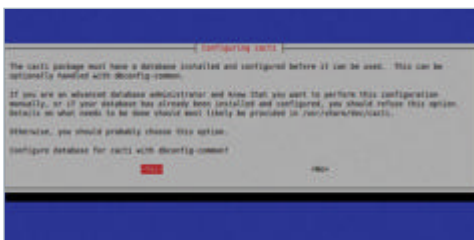
- Cacti : <http://www.cacti.net/>
- LAMP stack or equivalent
- ALAN
- Other SNMP devices (optional)



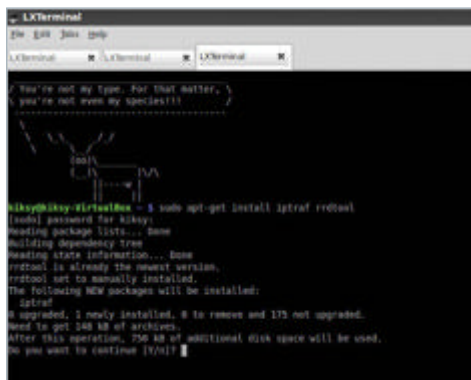
01 Install Cacti
 To start off, if you are running a Debian derivative, enter 'sudo apt-get update' to update your package list, then to get Cacti 'sudo apt-get install cacti'. For Fedora it's 'yum install cacti'. We will assume you already have a running Apache/PHP/MySQL server.



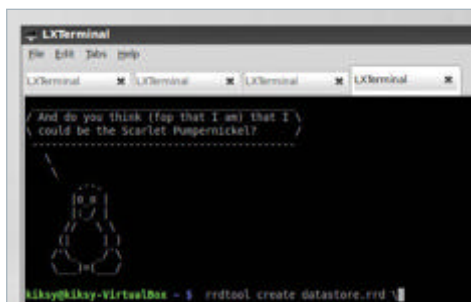
02 Set up Cacti
 Next you will need to enter some configuration options. Choose your Apache installation from the menu, or you can configure it yourself if you have a special version or wish to use something else.



03 Configure database
 Cacti needs a database to store all its records; you can configure this yourself or Cacti can automate it. Enter your MySQL root password when asked for it. Then you can enter a password for Cacti's db access.



04 Install RRDTool
 Next up we need to install RRDTool and any other dependencies. 'sudo apt-get install rrdtool iptraf' to start the process off. Once installed, open up a terminal window to create an example RRD file.



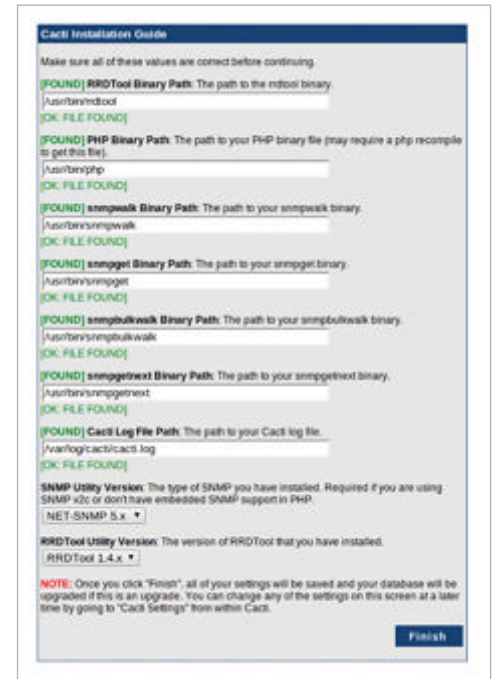
05 Create new RRDTool
 Type: 'rrdtool create datafile.rrd \
 DS:mysource:ABSOLUTE \
 TE:900:0:10000000 \
 RRA:AVERAGE:0.5:1:9600 \
 RRA:AVERAGE:0.5:4:9600 \
 RRA:AVERAGE:0.5:24:6000'.

This then sets up a basic RRD file.

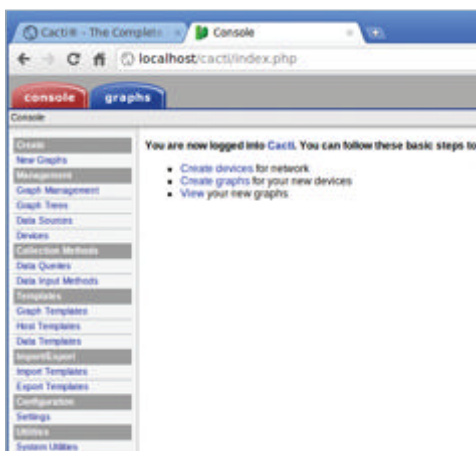
06 Create new RRDTool – part 2
 To explain that file... the DS: is the datasource; the other important element of this file is the intervals for polling. Here we have set 3 archives for 1 measurement, then the average of 4, then the average of 24. So the first is 15 minutes, (1), then an hour (15mins x4) etc.



07 Enter the Cacti interface
 Open a browser and navigate to 'localhost/cacti' and you should see the first welcome page. If not, then make sure Apache is running okay and that Cacti has installed into the correct directory and that it has the correct permissions.



08 Check dependencies
 The installer will then look for all the required files needed to start the setup. As long as all the files are found, you can then click 'next'; otherwise go back to terminal or your package manager and install them as required.




09 Log into Cacti

After that, you'll then be able to log into Cacti using username 'admin' and password 'admin', forced to change the password, then taken to the home screen. Now the fun starts.



10 Create new device

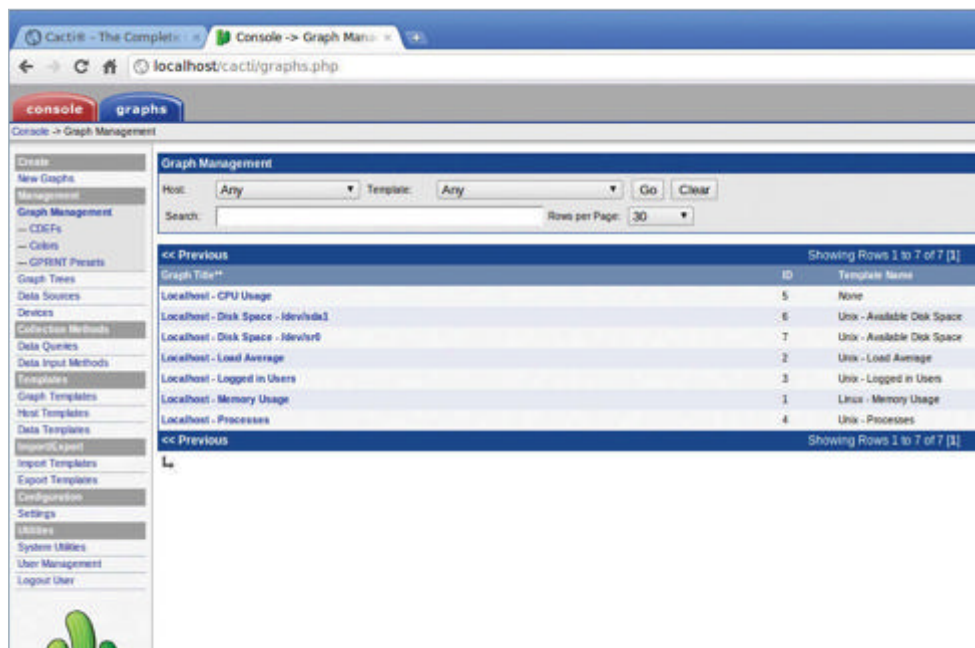
Click on 'Create Devices for Network' and you should see listed 'LocalHost' along with some stats about this. As we are currently on localhost, that's not terribly exciting so let's add a new device by clicking on 'Add'.



11 Ping the device


You can enter an IP or hostname for the device you want to gain data about. Then you can set Cacti to ping the device to make sure that it's reachable, using the 'Downed Device Detection' drop-down menu.

“Change the data sources to cross-reference different items that are polled”



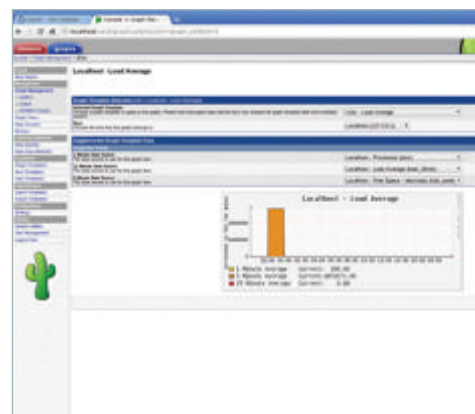
12 Create your first graph

Depending on the device you've added, you can create lots of interesting graphs. You'll need to make sure the device supports SNMP and it's configured – but assuming it's a Linux box, let's make a CPU load graph.



13 Set up graph

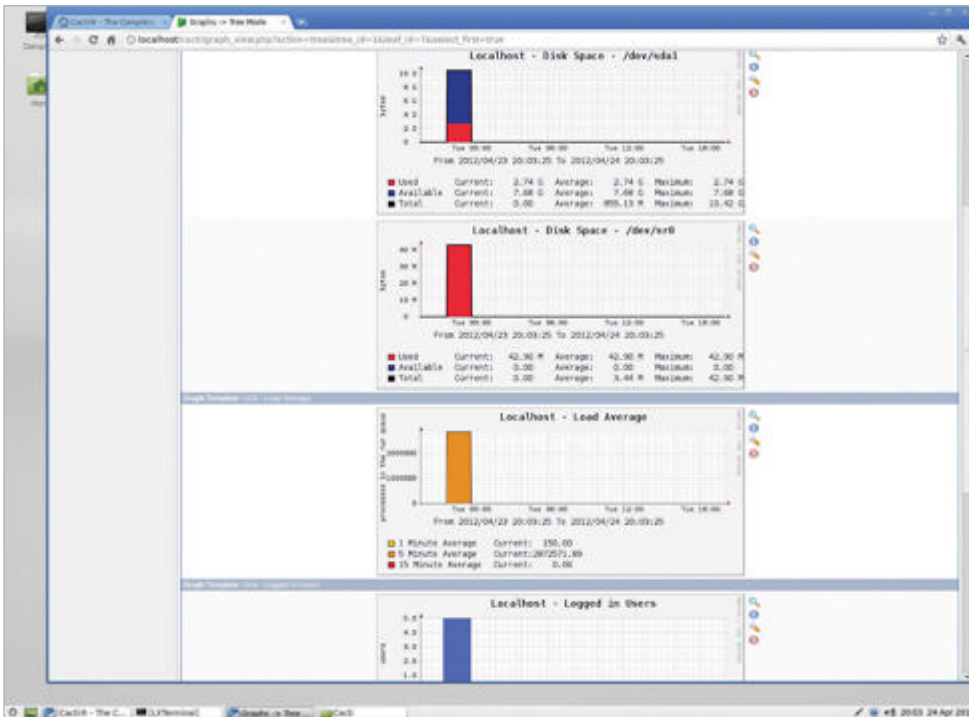
In the left-hand menu, click 'Graph Management' and then choose one of your devices. Localhost is fine to test out some of the options, as lots of preconfigured templates are already applied. Next click on 'Load Average' to display the CPU load over time.



14 Changing data sources


You can change the data sources to cross-reference different items that are polled using the drop-down under 'Supplemental Graph Template Data'. You'll need to save and reload to make the changes. As our minimum value is 1 min, you might need to wait a while before you see any interesting results.

“You can create lots of visually interesting graphs”




15 See all graphs

To see all the graphs you have made, click on the 'Graph' tab next to the 'Console' one at the top of the page. Here you can then filter them by date as well as search through them by keyword. There is also a button to export each graph's data as a CSV file, to be able to import into another source and manipulate. If you click on the spanner icon, you can see the graph's properties.



16 Edit graph properties

Click on 'Graph Templates' to see a list of currently created templates that are being used to poll for data. Choose one to open up its properties. Click 'Add' next to 'Graph Input Items' to add another field.



17 Add graph field

Let's say we want to map the max memory load over time – in this case we can add these items using the drop-down menu and then click on Create, then on Save. You will then need to refresh.

18 Load a pre-made template

There are many pre-made templates and scripts available online to help you out, and a great resource is www.debianhelp.co.uk/cactitemplates.htm. To import them, download and click on 'Import template', then select the file, or you can copy and paste the XML.

19 Import the Perl script

Next we need to import the associated script that actually contains the information on how and what to retrieve. Go to 'Data Sources' and 'Add', then choose the template we just uploaded and a device.



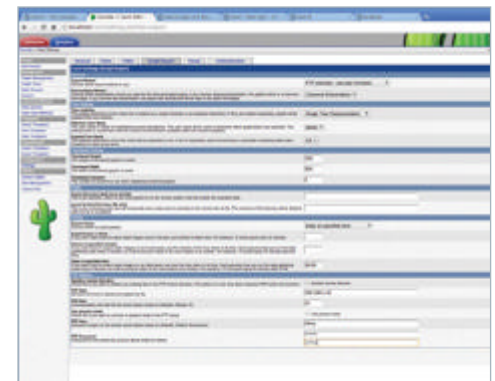
20 Link to the script

Click 'Create' and then in the data source field enter the location of the script. It might be a good idea to store all your scripts together; for example, 'home/user/cacti/scripts/routerstats.pl'. Then click on 'Create'.

“There is also a button to export each graph's data as a CSV file”


21 View new graph

Click on the graph tab again at the top and you should see at the bottom the new graph we just created from the imported template and script. Here we are tracking our local Apache server's activity.



22 Schedule a graph export

In the 'Settings' menu you can schedule a cron to run to export all the graphs. You can enter FTP details to automatically upload to a central location, use SSH or just save them locally for easy retrieval.



23 Running SNMP client

Setting up SNMP on another Linux box is fairly simple. You'll need the net-SNMP package installed. Edit the snmpd.conf; the minimum value to work is 'rocommunity public', which receives any request with no security.

24 SNMP on Android

There are numerous SNMP client and manager apps for Android, meaning you can monitor you phone or tablet's traffic, usage and even location using Cacti. If you have a large number of devices deployed on your site, this could provide valuable statistical information.

Tips & Tricks

Browse your folders with interactive reports, analyzer reports and dashboards

You can create your own interactive reports, analyzer reports and dashboards to gain insights about your business data

View your reports in tabs. You can even change some parameters and generate customised reports on the fly

Each folder contains some files – for instance, the various reports you can open

The screenshot shows the Pentaho User Console in a Mozilla Firefox browser. The interface includes a navigation pane on the left with folders like 'Steel Wheels', 'Analysis', 'Dashboards', 'Widget Library', and 'Reporting'. The main area displays a report titled 'Steel Wheels' with a 'Report Parameters' section containing fields for Product Line (Classic Cars), Region (NA), Year (2003), Top N Customers (3), and Output Type (HTML (Paginated)). Below the parameters is a 'View Report' button and an 'Auto-Submit' checkbox. The report content shows a table of 'Top 3 Customers' for 'Steel Wheels'.

Customer	Sales (US \$)	Volume (Units)	% of Total
Muscle Machine Inc.	\$ 95,657	774	16.28%
Mixi Gifts Distributors Ltd.	\$ 94,211	791	16.04%
Online Diecast Creations Co.	\$ 63,981	585	10.89%
Other Customers	\$ 333,579	2,809	56.79%

Create reports and charts with Pentaho BI

Gain essential insights into your business with this open source business intelligence suite

Resources

Pentaho BI Suite 4.1:

<http://www.pentaho.com/>

Advisor

Koen Vervloesem has been writing about free and open source software, and IT in general, since 2000. He has master's degrees in computer science and philosophy and can be reached at koen@vervloesem.eu



In 2004, a couple of veterans in the business intelligence (BI) industry started the company Pentaho. Their

ambitious goal was to develop an open source BI solution and to offer services for this solution that would be better than those from other BI vendors. Now in 2012, the Pentaho BI Suite has reached version 4.1. It has become a powerful open source integrated BI solution with many reporting and analysis components.

With its many possibilities, the Pentaho BI Suite is a solution you should evaluate if you're searching for new insights into your business.

It has a nicely polished web interface with graphical wizards to create new reports, charts and dashboards. The downside of its comprehensiveness is that you may get easily lost, but the excellent documentation makes up for this.

There are two versions of the software: the Community Edition and the Enterprise Edition. The latter adds certified versions of the software, support and maintenance, as well as some advanced capabilities. For this tutorial, we'll use an evaluation version of the Enterprise Edition of Pentaho BI Suite 4.1.



01 Download Pentaho Business Analytics
Go to www.pentaho.com/download/ to download a free 30-day evaluation version of Pentaho Business Analytics. Click on 'Download 32 Bit' or 'Download 64 Bit', as appropriate. Enter 'uname -m' in a terminal window if you're not sure whether you have a 32- or 64-bit Linux distribution.




02 Enter your details
Now you have to enter your contact details, such as your name, company name, role and for which purpose you're considering Pentaho. After this, your web browser downloads a huge 652MB .bin installer file that bundles all necessary components to run the Pentaho BI Suite.

03 Start the installer
Make the installer file executable with 'chmod +x pentaho-bi-suite-4.1.0-GA-x64.bin' and then start it with './pentaho-bi-suite-4.1.0-GA-x64.bin'. A wizard will guide you through the installation. Choose the 'Default' setup type, which installs all Pentaho applications with predefined settings.



04 Master password
There are a lot of users defined for the various Pentaho applications, and each of them can get a separate password. However, in the 'Default' setup type, you're asked to enter a master password that is used for all these users. Don't forget it!



05 Starting all components
At the end of the installation program, all Pentaho components are started (or you can do it manually later with './ctlscrip.sh start' in Pentaho's directory). This can take some time if you don't have a powerful machine. Among these components are an Apache Tomcat application server, a MySQL database server and more.



06 Sample users
The installation finishes with some important information about two sample users that are created: the administrator joe and the system user suzy, both with 'password' as their password. When you click on Finish, the Pentaho User Console is opened in your default web browser.

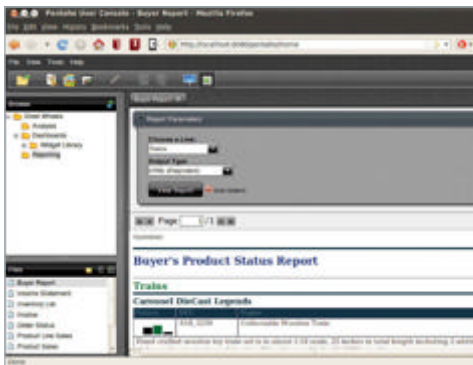


07 User Console
The User Console is the place where you do all your stuff in Pentaho. You can also open it by manually navigating your web browser to <http://localhost:8080/>. Enter the user name and password for one of the sample users and click on Login.



08 Wizards
When you are logged in, the User Console offers some wizards to create a new interactive report, analyzer report or dashboard. You also have the possibility to create a new data source from a CSV file or a database, or to manage an existing data source.

Tips & Tricks



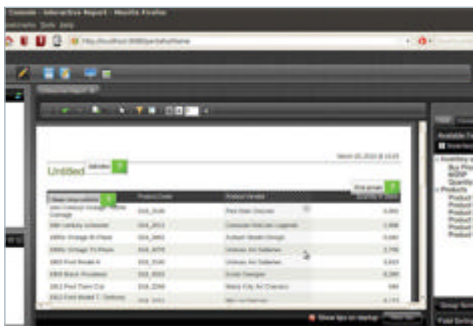
09 Sample reports

First take a look at some sample reports to see what's possible. Click on the 'Toggle Browser' icon and then browse the Steel Wheels folder at the left. For instance, select the Buyer Report file in the Reporting folder and click on the Open icon to view the report.



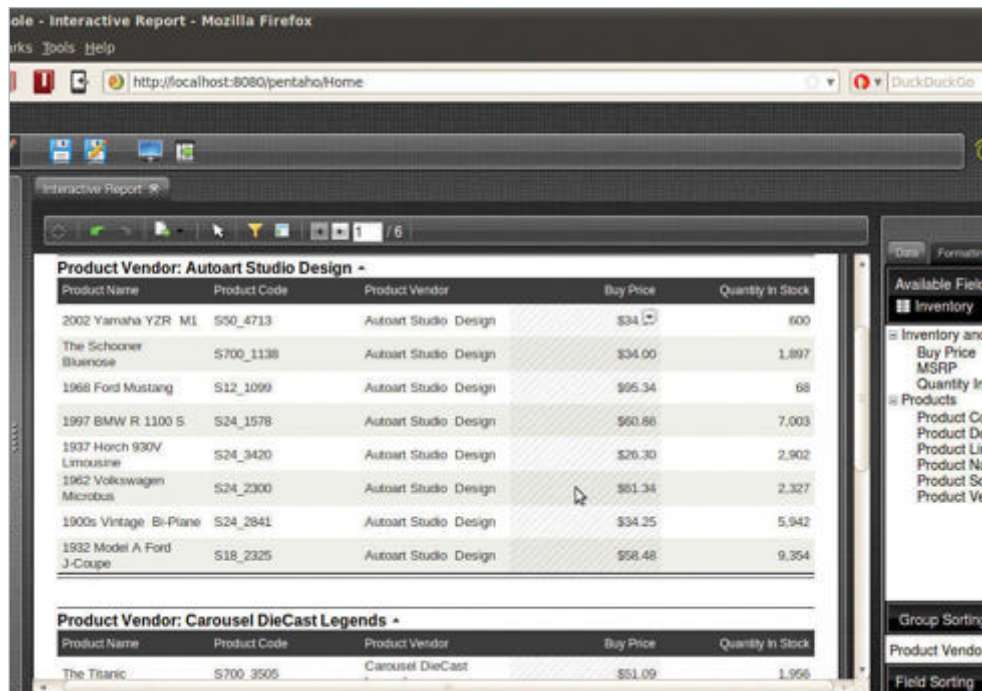
10 New interactive report

Click on the 'New Interactive Report' icon to create your own report. First, you have to specify the data source you want to use for your report. If you click on the green plus sign, you can add an external data source from a CSV file or SQL database.



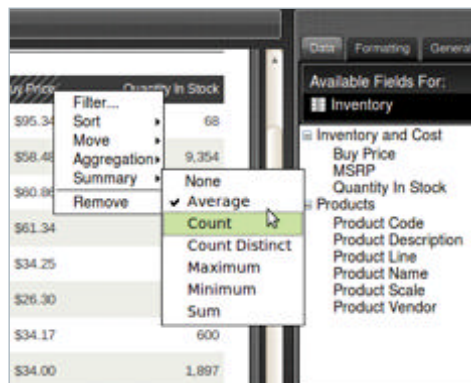
11 Add columns

You see an empty report now, but editing it is easy: just drag and drop fields in the Data tab at the right onto the form to add columns. You can also drag and drop these columns to arrange them in a particular order.



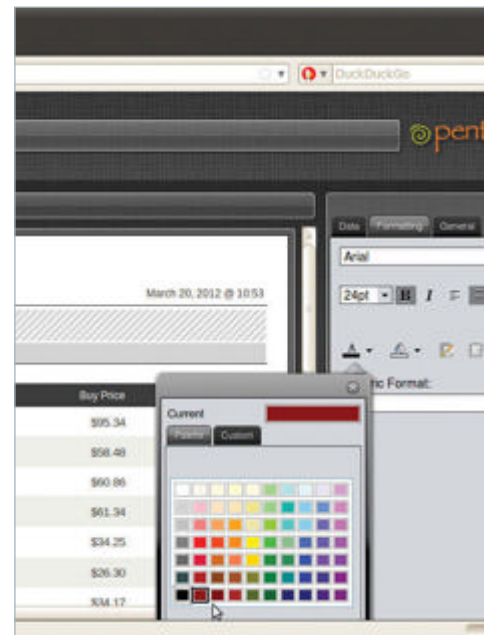
12 Grouping

Grouping items in a table is just as easy. For instance, when we drag the Product Vendor field from the Data tab and drop it right above the column heading, all products of the same vendor are grouped in a separate table with this vendor's name as the heading.



13 Column properties

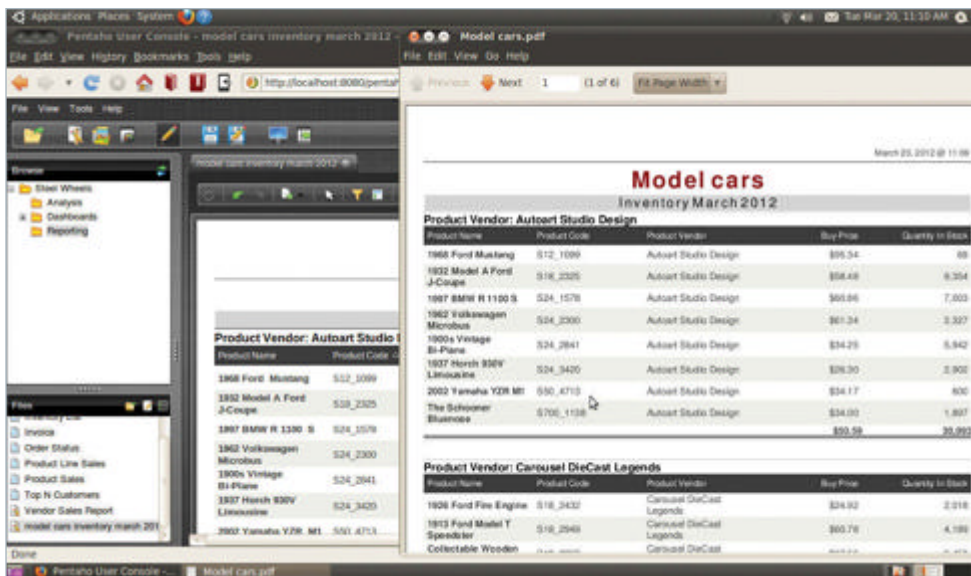
When you click on the triangle next to a column heading, a menu appears that enables you to change some properties. For instance, you can sort the entries in the menu in ascending or descending order, or you can show an average or a sum of the values at the bottom of the column.



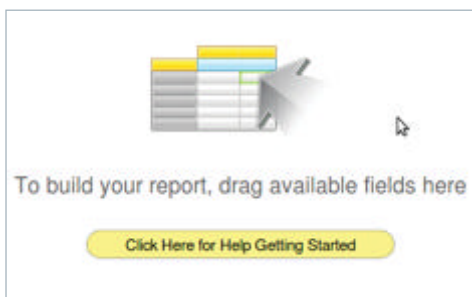
14 Formatting

Changing the formatting is also straightforward: select a column or a heading and then choose the formatting in the Formatting tab, such as the font, font size, colour and background colour and so on.

“You can change the format – such as a pie chart, bar chart, line chart and so on”



15 Export
 When you're happy with the way your report looks, it's time to click on the Save icon to save your previous work. But you're not finished yet: if you want to show your report to other people, you have to export it to be able to send them the file. Click on the Export icon and choose the format: PDF, Excel, CSV or HTML. Excel and CSV are meant for people who want to analyse and transform the data themselves, while PDF and HTML are fine for static purposes.



16 Analyse your data
 Interactive reports are nice for an overview, but if you want to drill down more into the details, click on 'New Analyzer Report'. Choose a data source and click OK. Available fields are shown at the right and you can drag and drop them to your report.

17 Formatting
 Now drag the fields that you want to analyse – for instance, Years, Months, Line and Sales – and drop them to the report. If you double-click on a column heading, you can choose the number format (eg currency), and by right-clicking you get a menu for conditional formatting and other tasks.



18 Charts
 If you click on the Charts icon, your report instantly changes into a chart. By clicking on the arrow next to the icon, you can change the chart format – such as a pie chart, bar chart, line chart and so on.



19 Filters
 If you want to drill down into a part of the data, you can use a filter. Click on the Filter icon and drag and drop a field to the filter area. For instance, you can limit the data to a specific year.

20 Bird's-eye view
 If you have created a lot of interactive and analyzer reports, you have seen many details but lack a bird's-eye view of your company as a whole. For this purpose, Pentaho lets you create your own dashboards. Click on the New Dashboard icon to create an empty dashboard template.

21 Inserting content
 Click on the Insert Content icon of one of the dashboard widgets. Choose a chart, data table, URL or file. For instance, we can create a dashboard that shows a table with the products in our inventory with less than 1,000 units in stock.

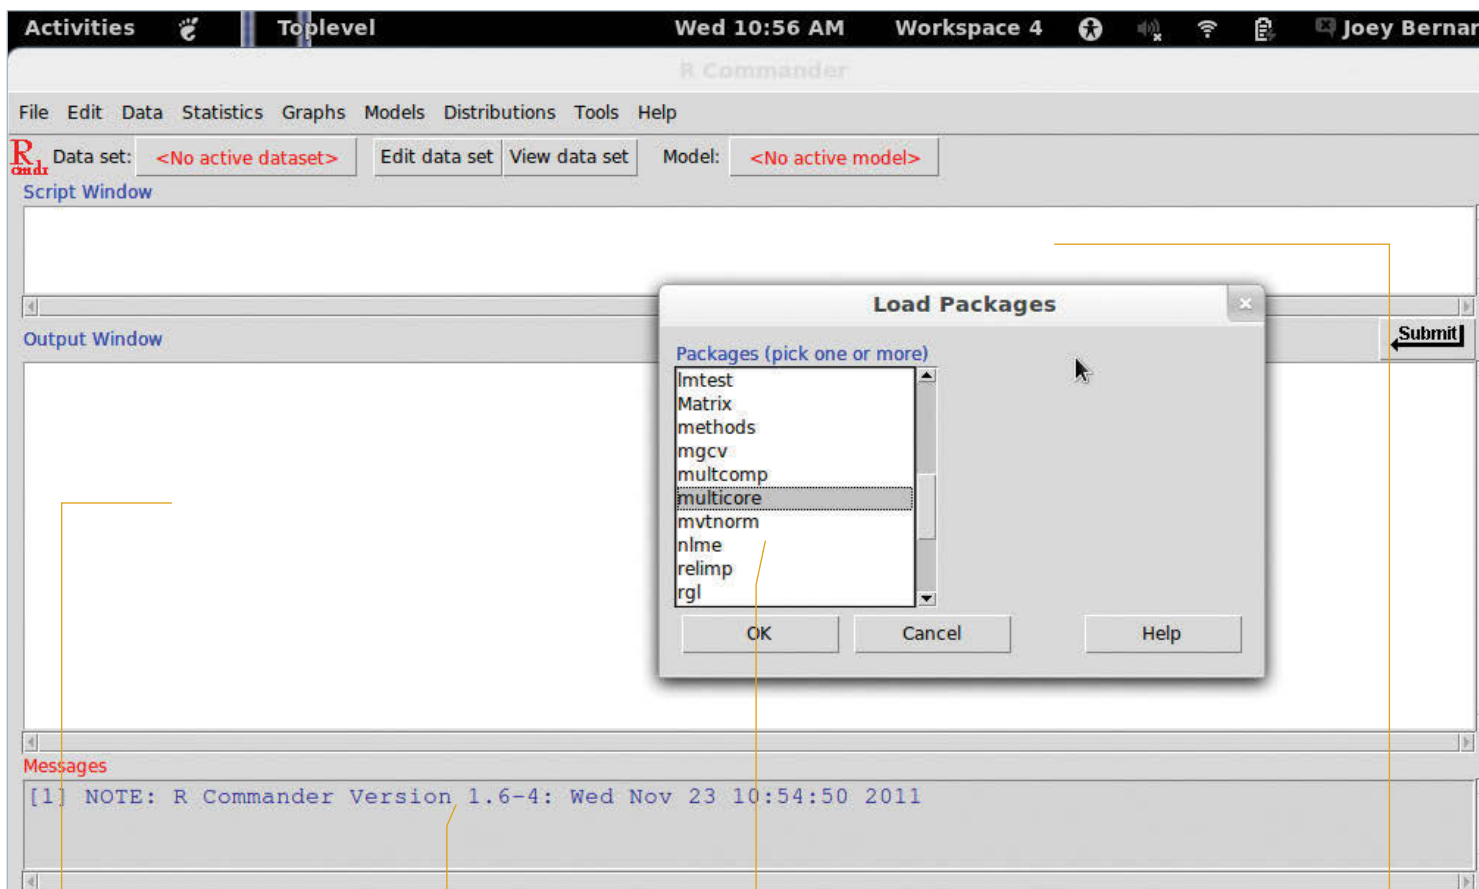
22 Files
 You can also insert an arbitrary file, which is an interesting way to reuse your interactive or analyzer reports that you have already made. For instance, you can load the analyzer report about the sales in a widget and show its chart.



23 Charts
 If you choose a chart in the Insert Content menu, you can choose the fields you want to show, as well as the chart type and other data. The chart designer builds up a chart based on your input; when you click OK, the chart is added to your dashboard widget.



24 Open source business intelligence
 We have merely scratched the surface in this tutorial, but it's important to point out that the Pentaho BI Suite offers a comprehensive solution for business intelligence, reporting, data mining and even Big Data. Consult the website for some documentation, including data sheets, white papers and videos about the possibilities.



All of the output from your R commands show up here. This can be logged out to a file so that you can reference it later

R also has to print out messages to the user. To keep your output clean, they show up here

Loading libraries becomes a point-and-click operation when you use one of the R graphical interfaces

This is the main script window. You can enter your R commands here to get your work done

Supercharge your R experience

R is one of the heavyweights in open source science. Learn how to run your R jobs in parallel to analyse large amounts of data

Advisor

Joey Bernard Using Linux as a desktop since the 1990s, Joey now works as a research consultant, helping university researchers do scientific computing



R is considered one of the best statistics packages around. One of its strengths is its flexibility, with a repository full of modules giving any number of extra features. The one we will look at here is running your analysis in a parallel environment. There are two general forms of parallelism: shared-memory parallelism running on one machine and message-

passing parallelism running across a cluster of networked machines. Many distros will contain packages for some of the features discussed here; however, we'll assume your distro doesn't contain these features, so we'll tell you how to install them from CRAN. The only assumption is that you have the core of R installed, either from your distro's package manager or from a download from the R project website.

```

jbernard@stargate: ~
File Edit View Search Terminal Help
jbernard@stargate:~$ R
R version 2.13.1 (2011-07-08)
Copyright (C) 2011 The R Foundation for Statistical Computing
ISBN 3-900051-07-0
Platform: i686-pc-linux-gnu (32-bit)

R is free software and comes with ABSOLUTELY NO WARRANTY.
You are welcome to redistribute it under certain conditions.
Type 'license()' or 'licence()' for distribution details.

Natural language support but running in an English locale

R is a collaborative project with many contributors.
Type 'contributors()' for more information and
'citation()' on how to cite R or R packages in publications.

Type 'demo()' for some demos, 'help()' for on-line help, or
'help.start()' for an HTML browser interface to help.
Type 'q()' to quit R.

```

■ Starting R on the console

01 Starting R

The default interface for R is a text-based console interface. You normally start it by executing 'R' in a terminal window or on the console. There are GUI options available, the two most popular being R Commander and RGUI. All of the following steps will assume that you are working on the console.

02 Checking what you have

One of the first things to do is to check what modules are already installed. The command to do this is simply 'library()'. When you run this with no options in the brackets, you get a list of all of the modules installed on your system. Along with the name, you will get a short description of what the module provides. You can find out more about a particular module with the command 'help(commandname)' - for example, 'help(grid)'.

```

jbernard@stargate: ~
File Edit View Search Terminal Help
Packages in library '/usr/lib/R/library':
base          The R Base Package
boot          Bootstrap Functions (originally by Angelo Canty
              for S)
class         Functions for Classification
cluster       Cluster Analysis Extended Rousseeuw et al.
codetools     Code Analysis Tools for R
compiler      The R Compiler Package
datasets      The R Datasets Package
foreign       Read Data Stored by Minitab, S, SAS, SPSS,
              Stata, Systat, dBase, ...
graphics      The R Graphics Package
grDevices     The R Graphics Devices and Support for Colours
              and Fonts
grid          The Grid Graphics Package
KernSmooth   Functions for kernel smoothing for Wand & Jones
              (1995)
lattice       Lattice Graphics
MASS          Support Functions and Datasets for Venables and
              Ripley's MASS
Matrix        Sparse and Dense Matrix Classes and Methods

```


“An even more general tool to do parallel calculations is **SNOW**. This lets you connect machines over several different transport mechanisms”

```
jbernard@stargate: ~  
File Edit View Search Terminal Help  
grid package:graphics R Documentation  
Add Grid to a Plot  
Description:  
    'grid' adds an 'nx' by 'ny' rectangular grid to an existing plot.  
Usage:  
    grid(nx = NULL, ny = nx, col = "lightgray", lty = "dotted",  
         lwd = par("lwd"), equilog = TRUE)  
Arguments:  
nx,ny: number of cells of the grid in x and y direction. When  
        'NULL', as per default, the grid aligns with the tick marks  
        on the corresponding _default_ axis (i.e., tickmarks as  
        computed by 'axTicks'). When 'NA', no grid lines are drawn  
        in the corresponding direction.  
col: character or (integer) numeric; color of the grid lines.  
lty: character or (integer) numeric; line type of the grid lines.
```

03 Installing packages

Installing packages has been made relatively easy in R, with a whole suite of built-in utilities to manage modules. The most basic command is `install.packages(modulename)`. It will try to install this module into the default library, which will fail if you aren't root. It'll then ask if it can install it into the R subdirectory of your home directory. It will then ask you to pick a CRAN mirror to download from. It'll then compile and install the downloaded module.

```
gcc -I/usr/share/R/include -fPIC -std=gnu99 -O3 -pipe -g -c fork.o -o fo  
gcc -I/usr/share/R/include -fPIC -std=gnu99 -O3 -pipe -g -c fork.o -o f  
gcc -I/usr/share/R/include -fPIC -std=gnu99 -O3 -pipe -g -c fork.o -o f  
gcc -I/usr/share/R/include -fPIC -std=gnu99 -O3 -pipe -g -c winfix.o -o  
gcc -shared -o multicore.so fork.o fork.o fork.o winfix.o /usr/lib/R/lib-  
installing to /home/jbernard/R/x86_64-linux-gnu-library/2.13/multicore/libs  
** R  
** preparing package for lazy loading  
** installing help indices  
** building package indices ...  
** testing if installed package can be loaded  
* DONE (multicore)  
The downloaded packages are in  
  "/home/jbernard/R/x86_64-linux-gnu-library/2.13/downloaded_packages"
```

04 Installing multicore

In order to install the multicore module, you'll need to make sure you have a complete GNU toolchain installed first, since there are C source files that need to be compiled. As an example, you should be fine installing the gcc package in Ubuntu. It should install all of the dependencies that you will need.

■ Installing a package

05 Setting up multicore

Multicore is relatively easy to set up. You load the module code with the command `library("multicore")`. Once it gets loaded, you have access to the help files and all of the functions. By default, multicore will try to figure out how many cores are available to run on.

06 Running a multicore job

Multicore provides a series of functions to help spread your work across multiple cores. The first command is `mclapply`. This function applies a function to each element of a vector. For example, let's say you have a vector of numbers called `x` and you want to find the square roots of all these numbers. You can do this **with**

```
mclapply(x, sqrt)
```

If you don't want to actually run on all of the available cores, you can set the total number **with**

```
mclapply(x, sqrt, mc.cores=4)
```

This will run on four of your machine's cores.

07 General parallel code

You can run an expression across multiple cores by using the `parallel` command. You can run it simply **with**

```
parallel(expression)
```

This will start a background process in order to execute the given expression. This lets you get on with other work during this background process. You can use the `collect` command to force the main R environment to wait until the processes run with `parallel` are done.

08 Using multiple machines

Multicore is limited to using a single machine. This is a serious limitation. If you want to use more than one machine, the most common method is to use MPI. The most common module in R to do this is `Rmpi`. You will need to have an MPI implementation installed on your system, such as `MPICH` or `openMPI`. If it gets installed under your distribution's package manager, it should be found without any problems by R.

09 Installing Rmpi

Installing `Rmpi` is simple, once the support libraries are installed. Just **execute**

```
install.packages("Rmpi")
```

This will download the source files, compile it and install it. Once it gets installed, you can load it into your R session **with the command**

```
library("Rmpi")
```

From here, you can start to actually use it.

10 Starting Rmpi

To start up `Rmpi`, you need to select how many slots to use, and which hosts to run on. If your MPI implementation already knows which hosts to run on, you can simply **execute**

```
mpi.spawn.Rslaves(nslaves=4)
```

to get four slots. If you need to set the hosts, you can hand them in through the `hosts` option as a list.

11 Using Rmpi

Now that you have some slaves created, you need to get them to do some work. You do have access to all of the low level commands you may be used to if you have already done MPI coding in C or FORTRAN. The simplest way to get work done is to **use**

```
mpi.remote.exec(cmd)
```

This command takes the script `cmd` and sends it out to all of the slaves to run. When you are done, you can **use**

```
mpi.close.Rslaves()
```

This goes out and cleanly shuts down all of the slaves that you created previously.

12 Other MPI options

`Rmpi` is not your only option. It actually gets used by other modules to provide parallel computational functions. One of these is `doMPI`, which provides an easier interface to MPI. You can start a cluster process with `openMPIcluster` and set a worker pool with `dompiWorkerLoop`. You can then use the `%dopar%` option of the `foreach` command. This takes the work being done in the `foreach` loop and spreads it out across the previously created worker pool.

13 SNOW (Simple Network Of Workstations)

An even more general tool to do parallel calculations is `SNOW`. This lets you connect machines over several different transport mechanisms. You can connect over MPI, PVM, or even raw network socket connections. Again, you **install SNOW with**

```
install.packages("snow")
```

You can then load it into your R session **with**

```
library("snow")
```

14 Getting set up

You need to create a cluster object to run on. If you need four MPI slots for a job, you can **execute**

```
cl = makeCluster(4, type="MPI")
```

You can then use this cluster object to run jobs. If you want to specify which hosts to run on, you can

pass these in as a list:

```
cl = makeCluster(c("host1", "host2"), type="SOCK")
```

This will create two nodes on `host1` and `host2`, connected through raw network sockets.

15 Running on SNOW

Once you have your cluster, you can **run code by executing**

```
clusterCall(cl, myfunc)
```

where `myfunc` is a function that you want to run across these nodes.

If, instead, you have some operation that you want to apply to a vector of values, you can **use**

```
clusterApply(cl, x, myfunc)
```

This command takes the function `myfunc` and applies it on each element of `x`. There is also a load-balancing version called `clusterApplyLB`, which tries to balance the work across the given nodes more evenly.

16 When you're done

When you're done running your job through `SNOW`, you need to clean up. This is done rather simply **with**

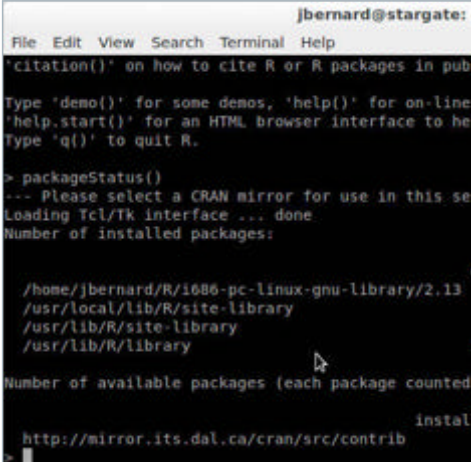
```
stopCluster(cl)
```

This is necessary, especially if you are running across multiple machines. If you neglect this step, you may need to go around and kill off these processes manually.

17 Keeping packages updated

After all this work, you'll want to make sure you keep yourself updated. You can check the status of your installed packages with `packageStatus()`. If you see that some packages need to be updated, you can use `update.packages()`. This will run through each of the packages that have updated versions and ask you which to update.

■ Updating packages in R



```
jbernard@stargate:
File Edit View Search Terminal Help
'citation()' on how to cite R or R packages in pub
Type 'demo()' for some demos, 'help()' for on-line
'help.start()' for an HTML browser interface to he
Type 'q()' to quit R.

> packageStatus()
--- Please select a CRAN mirror for use in this se
Loading Tcl/Tk interface ... done
Number of installed packages:

/home/jbernard/R/i686-pc-linux-gnu-library/2.13
/usr/local/lib/R/site-library
/usr/lib/R/site-library
/usr/lib/R/library

Number of available packages (each package counted
instal
http://mirror.its.dal.ca/cran/src/contrib
```

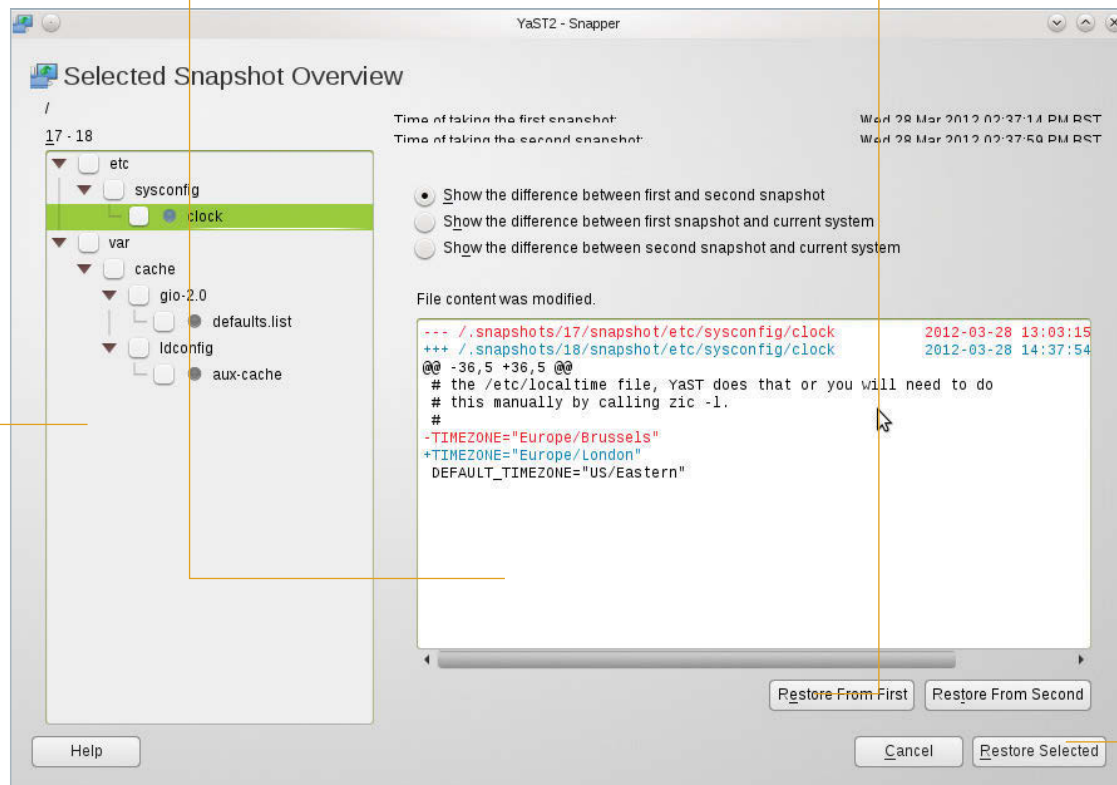
“If you want to use more than one machine the most common method is to use MPI”

Snapper shows you a list of all files that have been modified between two snapshots

Snapper can show the difference between various snapshots for specific files you're interested in

Restore the file with its original content

Restore more than one file at a time by ticking all the checkboxes and pressing the 'Restore Selected' button



Create Btrfs snapshots in openSUSE 12.1

Reverse all your mistakes using openSUSE's revolutionary Snapper tool

Resources

openSUSE 12.1 or SUSE Linux Enterprise 11 SP2 or higher:

<http://www.opensuse.org/en/>

A Btrfs root file system:

<https://btrfs.wiki.kernel.org/>

Snapper: <http://en.opensuse.org/Portal:Snapper>



Btrfs is the future file system for Linux, using the 'copy-on-write' concept, which improves performance and reliability. It's

already available in Arch Linux, in openSUSE, SUSE Linux Enterprise, Ubuntu, Sabayon Linux, Red Hat Enterprise Linux, Fedora, Debian and Slackware, although there's no mainstream Linux distro using Btrfs by default for the root file system.

OpenSUSE has full support for the Btrfs file system since the 12.1 release. You just have to tick a checkbox in the installer to create a Btrfs root file system instead of the default ext4 one.

On top of this, openSUSE 12.1 offers Snapper, a tool for managing Btrfs snapshots. The basic idea of Snapper is that it automatically creates a snapshot before and after running YaST or Zypper, compares the two snapshots and therefore provides the means to revert the differences between these two snapshots.

Snapper's configuration uses sensible defaults, such as hourly snapshots. Moreover, you can manage your snapshots with a user-friendly YaST module or with the command-line snapper tool. You can even compare the contents of a single file between two snapshots or roll back changes to a single file.

Advisor


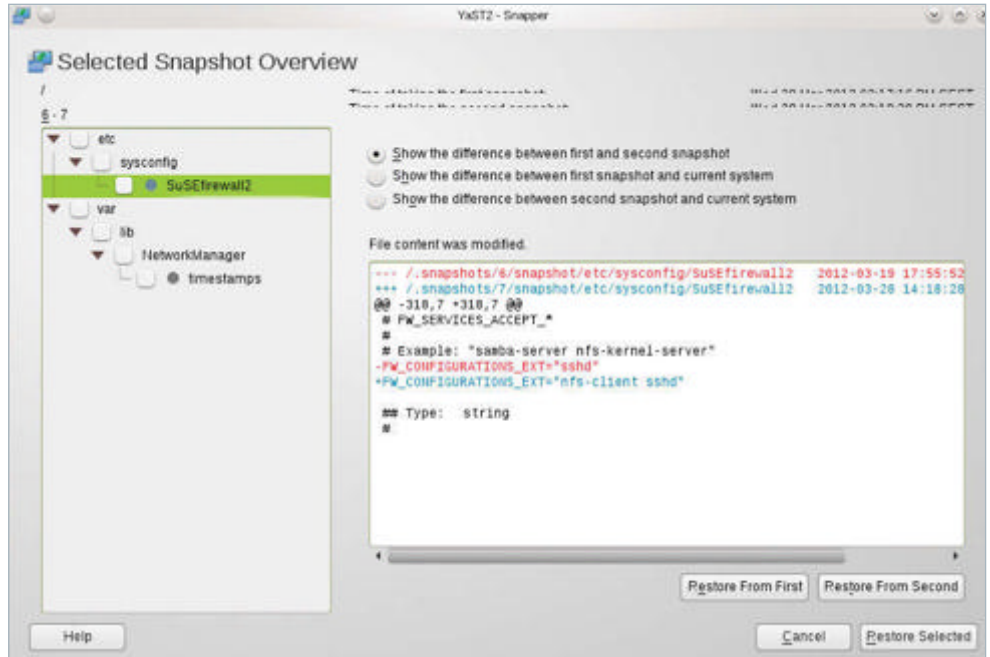
Koen Vervloesem has been writing about free and open source software, and IT in general, since 2000. He has master's degrees in computer science and philosophy and can be reached at koen@vervloesem.eu





01 Requirements

Make sure that your system fulfils the requirements: your root file system should use Btrfs; the packages snapper, snapper-zypp-plugin and yast2-snapper should be installed; and 'sudo snapper list-configs' should show a configuration for /. If you chose Btrfs for your root file system in openSUSE's installer, this should all be okay.

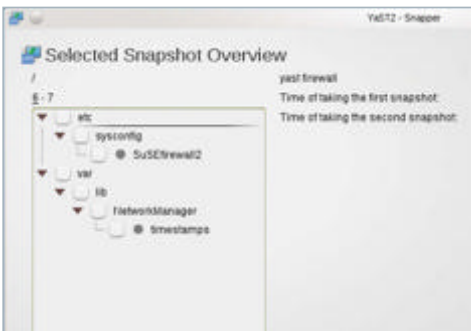



02 List the snapshots

Open the Snapper module in YaST or by entering 'yast2 snapper', which shows you a list of available snapshots. The snapshots that you created manually or that are automatically made hourly have their type set to 'Single'; the ones made by Zypper and YaST are called 'Pre & Post'.

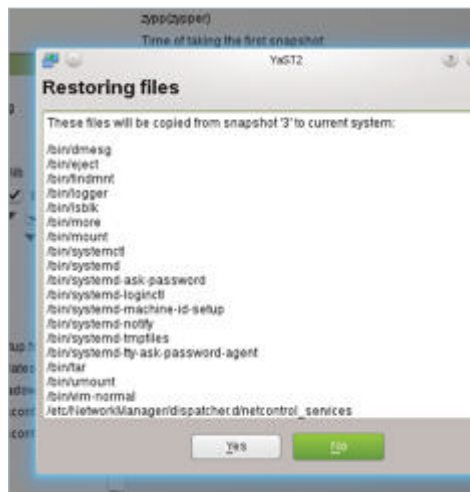
04 Peek into changed files

But there's more: you can display a 'diff' between the pre and post version of a file by selecting it from the list. We can see here that we have added nfs-client to the allowed services of openSUSE's firewall. If you want to revert this file, click on 'Restore From First' and confirm.



03 List changes

Select a 'Pre & Post' snapshot and click on 'Show Changes' to get an overview of the list of files that differ between the two snapshots. In this example, we opened a port in the firewall, which apparently changed only two files.




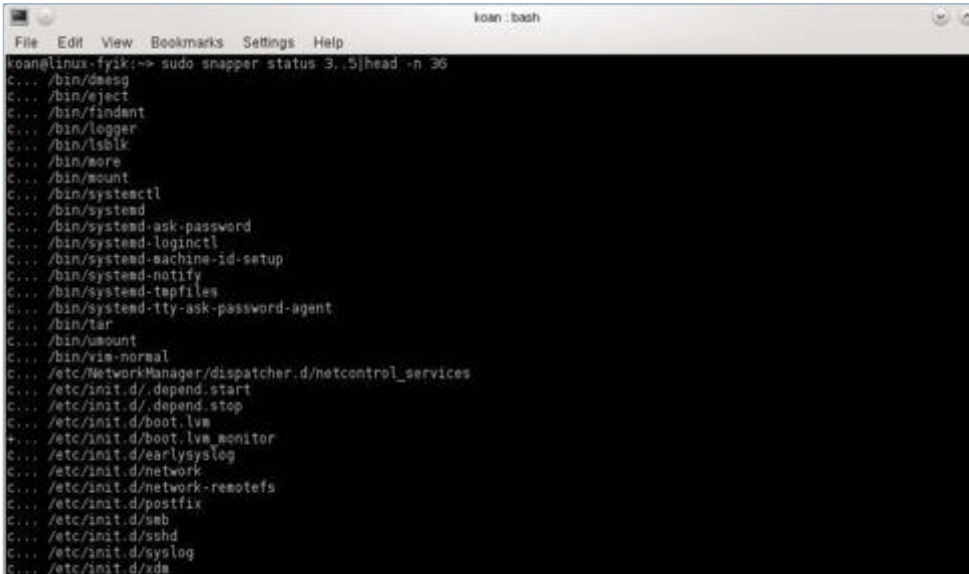
05 Revert more files

If you want to revert more than one file, just tick the checkboxes of the respective files or directories, click Restore Selected and confirm your choice by clicking Yes. Note that you should closely review all changes between two snapshots before starting a complete rollback!

06 On the command line

If you prefer the command line, you can get a list of Snapper snapshots with the command 'sudo snapper list'. Limit the results to pre and post snapshots with the option '-t pre-post', or to single snapshots with the option '-t single'.

“Snapper automatically creates a snapshot before and after running YaST or Zypper”



07 Snapper status

A 'sudo snapper status PRE..POST' command – where PRE and POST are the IDs of the pre and post snapshots – will list all the files that have been changed between the snapshots. Changed files are marked with a 'c', added files are marked with a '+', and deleted files get a '-' before the filename.

09 Snapper undochange

Reverting one or more files is also easy on the command line: use the command 'sudo snapper -v undochange PRE..POST FILENAMES'. Without a filename, this command reverts all changed files between the pre and post snapshots. Attention: this command doesn't ask for confirmation, so be absolutely sure what you're doing.

11 Other configurations


You can create your own configurations for other partitions formatted with Btrfs. To use Snapper with your home directory, add a configuration with the command 'sudo snapper -c home create-config /home', which creates a new configuration file /etc/snapper/configs/home. Verify the existence of a new configuration with the 'sudo snapper list-configs' command.

12 Adjusting the configuration

These configuration files let you customise the behaviour of Snapper on a per-file-system basis. For instance, by default old snapshots are cleaned up daily by a cron job. The variable NUMBER_LIMIT defines how many snapshots are kept and NUMBER_MIN_AGE defines how old a snapshot has to be (in seconds) to be deleted.

08 Snapper diff

If you want to peek into the contents of the pre and post versions of one or more specific files, use the command 'sudo snapper diff PRE..POST FILENAMES'. Without a filename, this shows a diff of all the changed files between the snapshots. A tip: pipe the output to colordiff to get a coloured diff.




13 Timeline cleanup

By default, Snapper creates hourly snapshots and also deletes old ones (TIMELINE_CREATE and TIMELINE_CLEANUP are set to 'yes'). Other variables specify that the last ten snapshots are kept, as well as the first daily snapshot of the last ten days, and the same for monthly and yearly snapshots.

14 Empty pre-post pairs

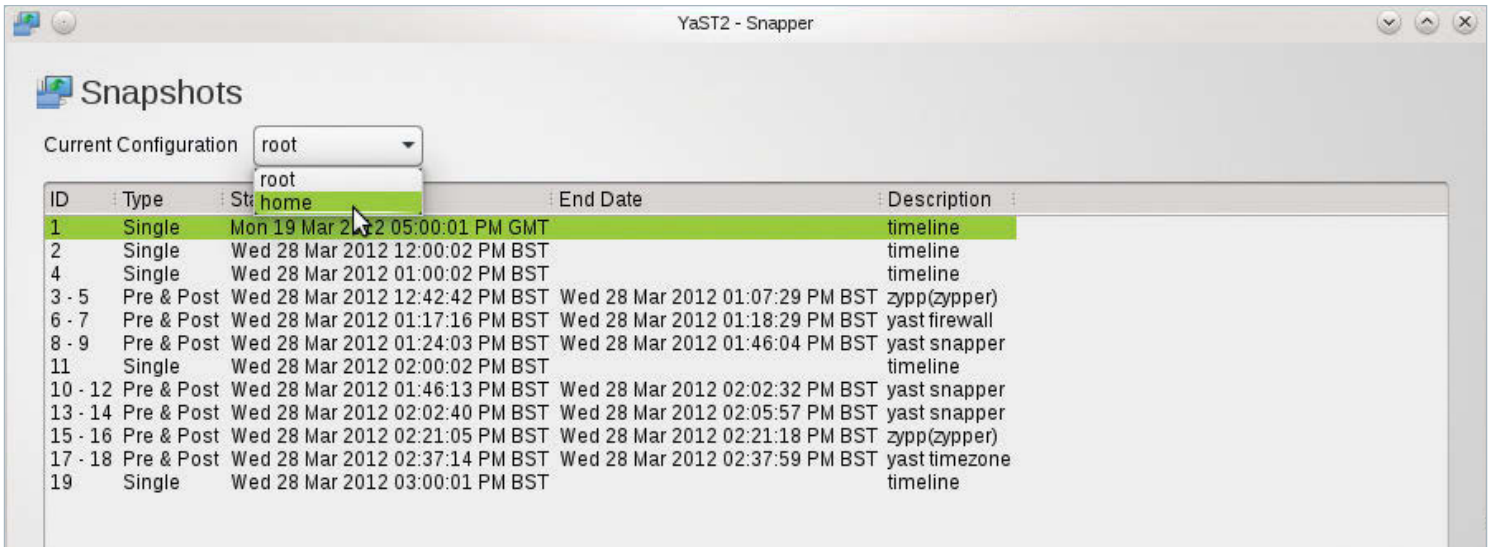
Other variables are for cleaning up empty pre-post pairs: EMPTY_PRE_POST_CLEANUP is by default set to 'yes' and EMPTY_PRE_POST_MIN_AGE is set to 30 minutes (1800 seconds). You can adjust all variables in the configuration files to change Snapper's behaviour, except for SUBVOLUME and FSTYPE.



10 Configuration

By default, openSUSE only creates a Snapper configuration for your Btrfs-formatted root file system. You can verify this with the 'sudo snapper list-configs' command, which only shows root by default. Snapper's behaviour for '/' is listed in the configuration file /etc/snapper/configs/root.

“Pipe snapper diff's output to colordiff to get a coloured diff”




15 Working with configurations

Now that you have more than one Snapper configuration, you can choose your configuration in the drop-down menu of the Snapper module for YaST if you want to manage the snapshots of another configuration. If you want to use another configuration with the command-line snapper command, you have to use the option '-c CONFIGNAME' before the subcommands. If you don't specify a configuration, the snapper command will use the default configuration, which is for your root file system. Each configuration has its own set of snapshots and corresponding IDs, so always be sure that you're working with the right configuration.



16 Manually creating snapshot

The 'snapper create' command lets you manually create snapshots, which can be interesting if you'd like to use the handy diff, status and undochange commands for your own tasks. The '-d DESCRIPTION' option lets you give a description and the '-u USERDATA' is for arbitrary key-value pairs separated by a comma.



17 Snapshot types

By default, Snapper creates a snapshot of type 'single'. You can create a pre or post snapshot with the option '-t pre' or '-t post'. When creating a post snapshot, you have to provide the ID of the corresponding pre snapshot with the option '--pre-number ID'.

18 Cleanup algorithms

By default, any snapshot that you create manually will never be deleted automatically. If you prefer that Snapper cleans up your old manually created snapshots, you have to specify a cleanup algorithm with the option '-c ALGORITHM' - where ALGORITHM is one of number, timeline or empty-pre-post, which are set up in the configuration file.

19 Modifying snapshot metadata

After a snapshot has been made, it cannot be modified, but you can change some of its metadata with 'sudo snapper modify [options] ID'. Use '-d DESCRIPTION' to change the description, '-u USERDATA' to change the userdata, and '-c ALGORITHM' to change the cleanup algorithm.

20 Deleting snapshots

You can manually delete a snapshot with 'sudo snapper delete ID'. When deleting a pre snapshot, you should also delete its corresponding post snapshot, and vice versa. Also make sure to delete old snapshots first, as these occupy the most disk space because of the copy-on-write approach of Btrfs.

21 Manual cleanup

Note that the automatically made snapshots are already cleaned up by default using a daily cron job, so you shouldn't have to delete them. If, however, you want to run

the cleanup algorithm manually, for instance after changing the cleanup configuration, you can do this with the command 'sudo snapper cleanup ALGORITHM'.



22 No boot files

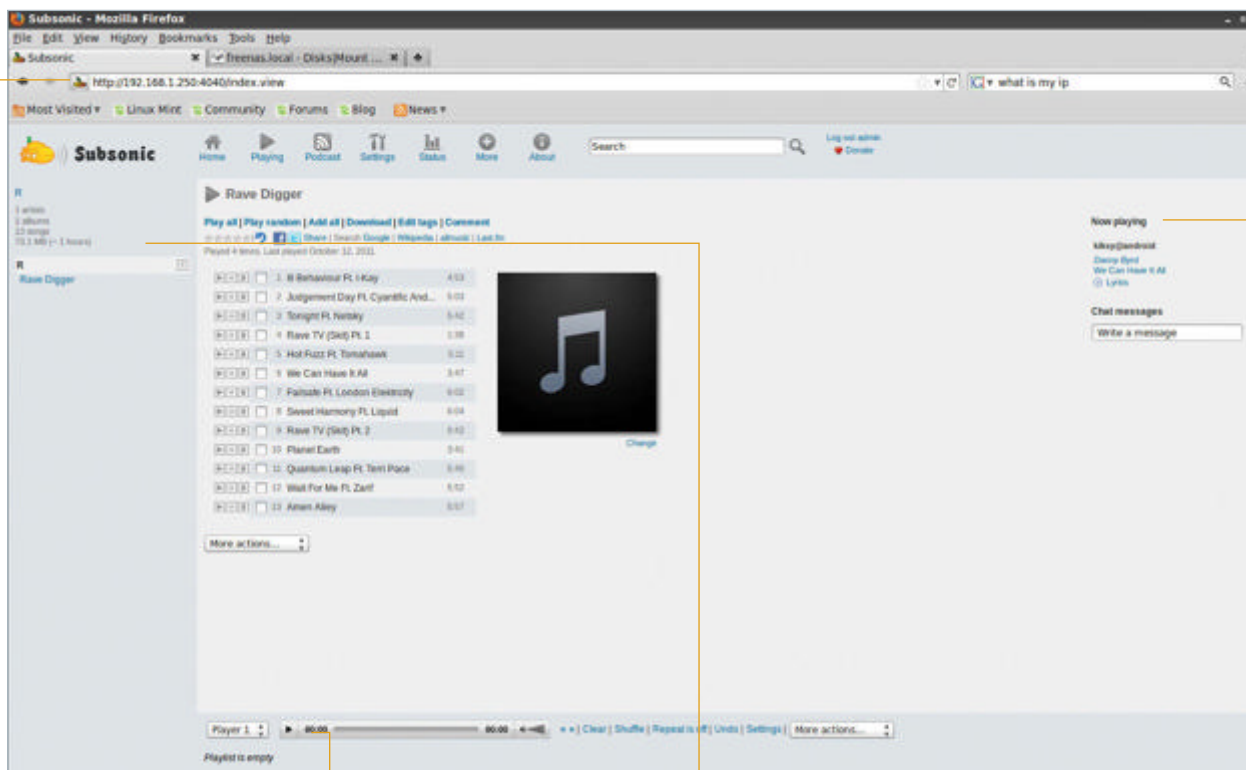
In openSUSE 12.1, the /boot partition isn't Btrfs-formatted, so Snapper doesn't support it. This means that whenever you do a rollback of a kernel installation, you need to manually remove the kernel and its initrd from /boot and remove its boot entry in /boot/grub/menu.lst.

23 Help

If you can't memorise the snapper commands or if you want to delve deeper into the specifics of a particular snapper command, the man page is an excellent resource. The 'snapper help' command gives a more concise overview of the commands.

24 Disabling Snapper

If you're not convinced that you need the automatic Snapper snapshots, you can always disable them. Pre and post snapshots made by YaST are disabled by setting USE_SNAPPER to no in /etc/sysconfig/yast2, and pre and post snapshots made by Zypper are disabled by uninstalling the snapper-zypp-plugin package.



As the interface uses a web browser to access the server, you can listen to your music library from any machine in the world

The player doesn't just support music; videos can also be watched using Subsonic, even on your mobile device

The menu on the left gives you a quick list of your media library, organised alphabetically. There is theoretically no limit on how big your collection can be!

The 'Now Playing' auto-updates to show you what others are listening to. You can even have live chats with them

Streaming media to your Android device

Amazon, Apple and Google all offer media streaming, but can it be done using only open source?

Resources

FreeNAS version 7: www.freenas.org/

Subsonic Standalone: www.subsonic.org/pages/index.jsp

Subsonic Android App: <https://market.android.com/details?id=net.sourceforge.subsonic.androidapp&hl=en>



As our smartphones have replaced iPods and their capacious hard drives, we have been left with a dilemma when wanting media on the move. Most phones only come with a few gigabytes of flash storage, which is great for a few albums or a couple of films, but what about if you go on long trip?

Amazon, Canonical, Google and Apple all offer music streaming services from the cloud, which are great, paid-for services. Each one, however, has its downsides, be it cost, space or file format limitations. Setting up your own home server

means you can stream as much media as hard drive space you have, can share it with friends and have greater control of format support.

If you have an old PC or laptop lying around, put it to good use by setting up a full media and file-serving solution using FreeNAS. This guide will take you through the steps needed to get it up and running. Note that the current version of FreeNAS (8) doesn't support changing the size of the /var directory, and has limited support for installing external packages, so installing Subsonic is not possible. This should be fixed in a later release.

Advisor

Kieron Howard First introduced to Linux eight years ago when installing Debian on his Xbox, Kieron has continued to enjoy putting Linux on devices such as iPods, PS3s and various phones




01 Choose FreeNAS version

If you are going to install FreeNAS from a CD drive, download the ISO file. If your machine doesn't have a CD-ROM drive, then download the full install version and follow these instructions: http://doc.freenas.org/index.php/Installing_from_the_Image. Also, be sure to choose the version for your architecture.


02 Install FreeNAS

Put your CD or USB drive into the machine, making sure the device is set as first boot in the BIOS. Also plug in your 4GB+ USB stick. You will then see the menu. Choose 9 and then choose Full OS on HDD + DATA + SWAP. Make the OS partition around 2,048MB.




03 Boot FreeNAS

After installing, remove the CD or USB stick and reboot the machine. You will then be presented with the FreeNAS config screen. The machine should automatically connect to your network and obtain an IP via DHCP. If not, press 1 to set up network connections.



04 View the web interface

The FreeNAS server is designed to be headless, so setting up shares etc is done through a browser on another computer. Enter the address shown on the config screen, such as 192.168.1.250. You will then be presented with the main FreeNAS admin page. Log in with 'admin' and 'freenas'.



05 Setup Admin user

As we are going to be exposing our server to the world, one of the first things we should do is set up a new admin password. In the left menu go to Access>Users and Groups. You can also change the Admin user as well clicking the '+' button.



06 Add some storage

A NAS isn't much use without any media attached. To add a drive, click on Disks>Management and '+'. This will give you a drop-down of attached storage and a prompt for a name and file system.

07 Add ZFS storage

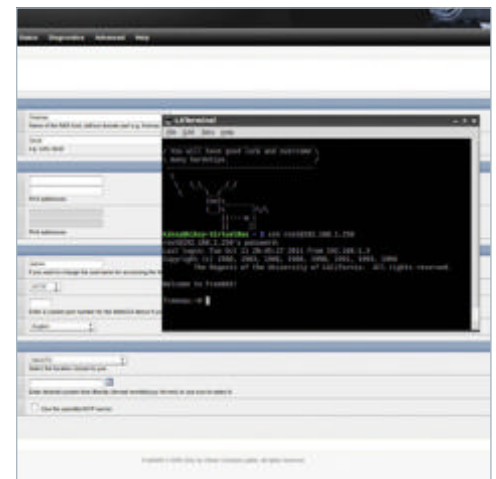
FreeNAS works best with the ZFS file system; and if you have a blank drive, it's best to set it up in FreeNAS. Click on Disks>Format and 'ZFS' as file type. Then choose the drive you wish to format. Check ZFS and hit OK.

08 Enable SSH

To be able to work on the setup from other machines, let's enable SSH. In the Services menu, click on the SSH option. Check 'Permit Root Login' and 'Password Authentication'. Then press 'Enable' and 'Save and Restart'.

09 SSH into FreeNAS

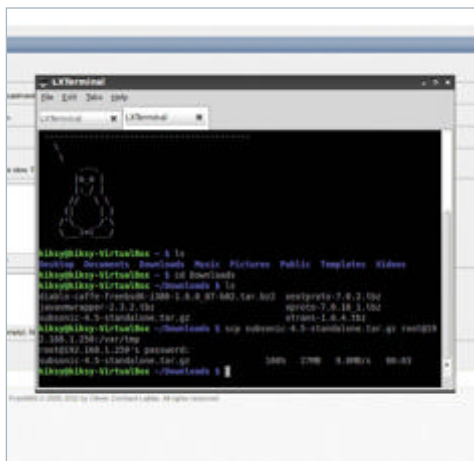
Open up a Terminal window on another machine and enter 'ssh root@192.168.1.250' (replace the IP with your own) and then 'YES' when it asks about the key. Enter your admin password you set earlier. You should then be logged in and see the '#freenas' prompt.



“Now we just need to add some music and videos to the server and we will be rocking in no time”

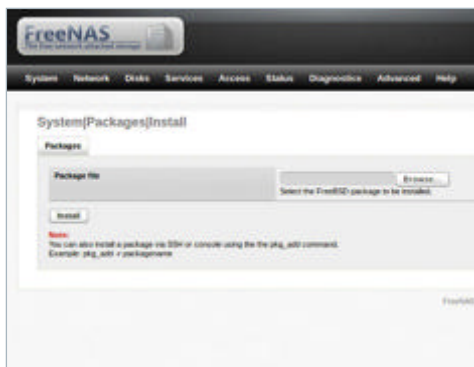
10 SCP Subsonic to FreeNAS

We need to copy the Subsonic files across to FreeNAS with 'scp subsonic-4.5-standalone.tar.gz root@192.168.1.250:/var/tmp' and then the packages that we downloaded. Use 'Diablo Latte JRE 1.6.0-7 FreeBSD 7.x/i386 End-User 27M diablo-jre-freebsd7.i386.1.6.0.07.02.tbz', replacing the filename with the other .tbz package and encoder files.



11 Install packages

Move to the upload directory using 'cd /var/tmp' and then install each package in turn: pkg_add -v pkg-config-0.25_1.tbz pkg_add -v xtrans-1.0.4.tbz pkg_add -v xproto-7.0.10_1.tbz pkg_add -v xextproto-7.0.2.tbz pkg_add -v javavmwrapper-2.3.2.tbz pkg_add -v diablo-jre-freebsd7.i386.1.6.0.07.02.tbz If you have any problems, you can also install packages using the web interface by going to System>Packages.



12 Install Subsonic

Change directory to /var and then make some new directories:

```
mkdir subsonic
cd subsonic
mkdir standalone
cd standalone
```

Move the uploaded Subsonic to the directory:

```
mv /var/tmp/subsonic-4.0.1-standalone.tar.gz /var/subsonic/standalone/subsonic-4.0.1-standalone.tar.gz
```

Extract the tarball:

```
tar xvfz subsonic-4.0.1-standalone.tar.gz
```

This should unpack the files in that directory (/var/subsonic/standalone).



13 Install encoders

In the 'Subsonic' directory, make a new one named 'transcoders' with mkdir. Then copy across the encoders MAC, LAME and FLAC that we downloaded earlier: 'scp lame.tbz root@192.168.1.250:/var/subsonic/transcode'. Install the encoders using 'pkg_add -v lame.tbz' for each encoder.

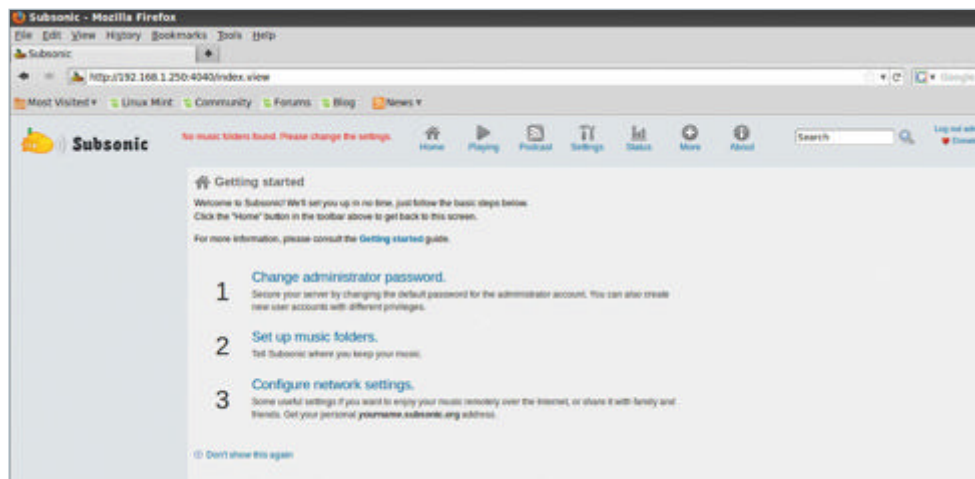


14 Test Subsonic

To make sure everything is working, let's test out Subsonic. Cd into 'var/subsonic/standalone' and run 'sh subsonic.sh'. You should see the message 'Started Subsonic [PID 1234] /var/subsonic/subsonic_sh.log'. In your browser, navigate to your FreeNAS IP and port 4040. For example, 192.168.1.250:4040.

15 Log into Subsonic


You should then see the Subsonic login page. Use 'admin' for the user name and 'admin' for password. Subsonic's main menu will then greet you! Congratulations, the hardest work is done. Now we just need to add some music and videos to the server and we will be rocking in no time. But first let's secure the admin account. Go to Settings>Users and change the admin password. While you are here you can also set up other user accounts if you plan on sharing the server with other people.



■ Secure your Subsonic admin account by changing the password

16 Add music to Subsonic

Back in FreeNAS, go to Disks>Mount Point and choose the drives you attached earlier. Set a mount location and file system type. Give each disk a name and then apply the changes. The status in the menu should say 'OK'.




17 Update the search index

Subsonic won't automatically recognise new files as you add them. The search is usually set to update every 24 hours, but you can schedule this how you want. Go to Settings>Search and then click 'Update Search Now'. Go to 'About' to view the log and progress.

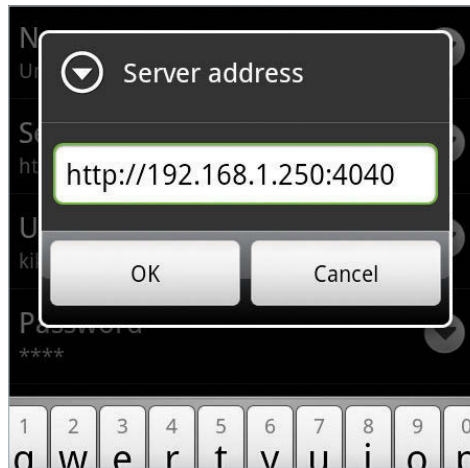
18 Download the Subsonic app

On your Android device, install the Subsonic Music Streamer app from the Market. Open it up and, for now, make sure you are on the same Wi-Fi network as your FreeNAS server so we can test it out. In the app, choose Select Server>Unused 1.




19 Subsonic app

Press the Menu button and 'Unused 1' in the settings. Enter the IP of the server and the port. Then enter your username and password and 'Test Connection'. You should get the 'Connection OK' message. If not, make sure you are on the same Wi-Fi network and not 3G and your user is set up correctly.




20 Play music

In the main screen, press the second icon in the bottom menu to bring up the artist list. Tap on one you wish to play then choose a song. After a few seconds of buffering depending on your connection speed, the music should play.



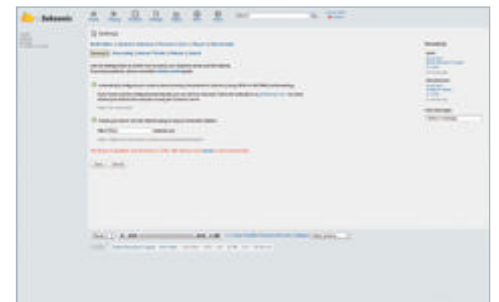
21 Automatically download podcasts

You can set Subsonic to automatically download podcasts for you. Go to 'Podcast' then add the .xml link in the box. You can set Subsonic to check for new episodes in Podcast>Episodes and also where to download them to.



22 Access Subsonic on the go easily

Subsonic offers a service giving you an easy-to-remember URL rather than having to use your external IP (which may change every time you reset your router). Go to Settings>Network and choose your address. Use of this service may require a donation to Subsonic.



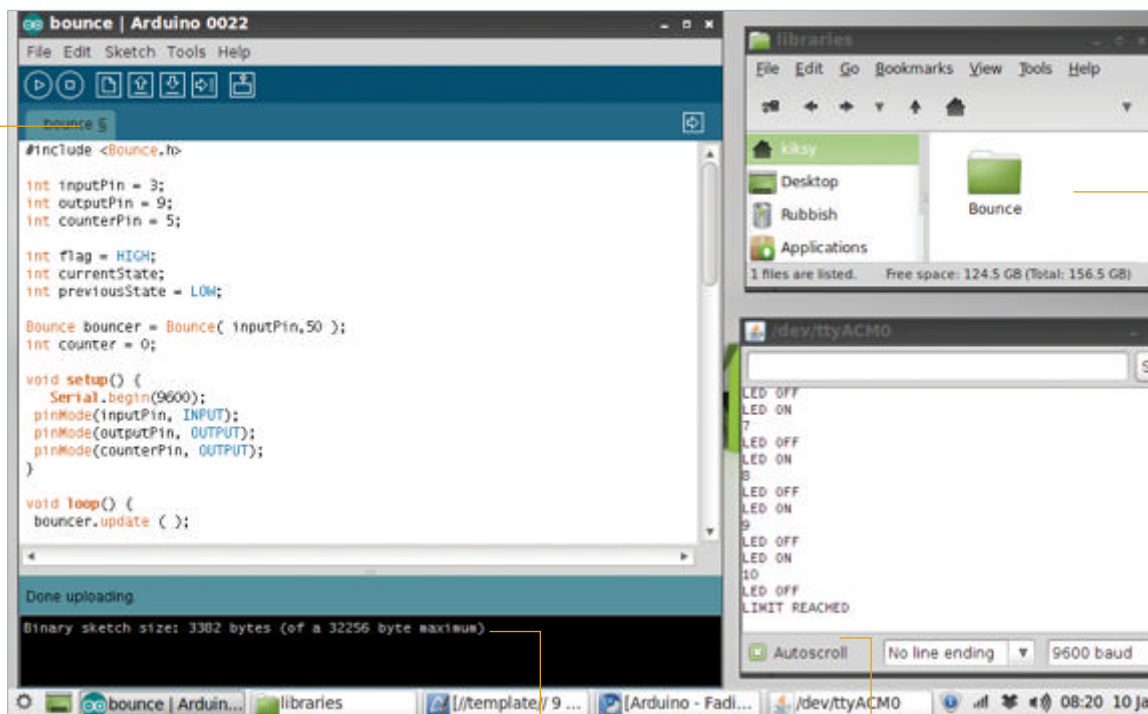
23 Access Subsonic on the move – manual setup

If the automatic router UPnP settings don't work for you, then in your router's settings you will need to forward the port to Subsonic. Each router is different but it can usually be found in 'Port Mapping', 'Port Forwarding' or similar.



24 Forward port

Enter the local IP as your Subsonic IP, eg 192.168.1.250, the local port as 4040 and the external port as 4040 (or whichever port you wish to connect on). Now, in the Android app, set up the 'Unused 2' server as your external connection.



The interface is simply laid out and easy to use. Code highlighting helps with those not so familiar with programming

When compiling your sketch, the status is shown here, and failures can be debugged from the output. The LED TX/RX lights also flash on the board

The serial monitor allows for easy debugging on your host machine. You can also send inputs to the Arduino

The IDE comes with a few examples and some pre-installed libraries, there are lots more available online to speed up development

Beginner's guide to Arduino

Taking your first steps using the Arduino prototyping board

Resources

Arduino Uno board: <http://arduino.cc/en/Main/Buy>

Breadboard: <http://www.maplin.co.uk/ad-100-breadboard-5195?c=froogle&u=5195&t=module>

Jumper wires: <http://www.hotsolder.co.uk/jumper-wires-pack-of-10-115mm-264-p.asp>

Switch: <http://proto-pic.co.uk/mini-push-button-switch/>

Selection of LEDs: <http://www.maplin.co.uk/multi-colour-led-613168>

USB A-B cable: <http://www.amazon.co.uk/HP-USB-Cable-Metre-C6518A/dp/B00004Z5DD>

Resistors: <http://www.ecrater.co.uk/p/6796434/30pcs-1k-ohm-resistors-1-2w-5>

Bounce library: <http://www.arduino.cc/playground/Code/Bounce>

Advisor

Kieron Howard First introduced to Linux eight years ago when installing Debian on his Xbox, Kieron has continued to enjoy putting Linux on devices such as iPods, PS3s and various phones



There are plenty of reasons to start experimenting with Arduino. Maybe you have an idea for a product and need to produce a proof of concept; or perhaps you have a cool idea for a piece of interactive furniture or just want to play around controlling real-world objects using code. A quick search around the internet will reveal a

wealth of exciting examples of projects built using Arduino and you really are only limited by your imagination.

With the Arduino units themselves being fairly cheap, and the fact that small electronic components are inexpensive, it doesn't require a huge outlay to start producing some clever and inspirational projects. The Arduino language is based on wiring, which in turn extends on processing. This tutorial will assume no prior coding knowledge, but if you do have any experience with C/C++ you should be right at home. The guide will show you how to take an input from a physical source, then do something with that data and output the result to an LED.

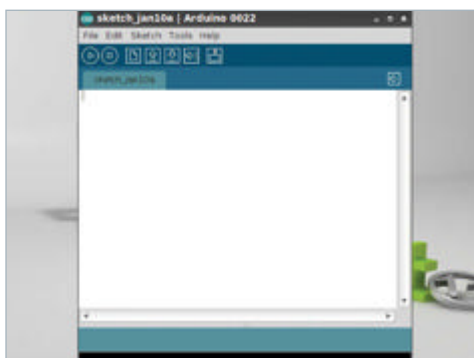


01 Install the Arduino IDE

The first thing we need to do is set up our dev environment. The Arduino packages are available in Synaptic apt or 'sudo apt-get install arduino'; Fedora users can just 'sudo yum install arduino' or you can download from <http://arduino.cc/en/Main/Software>.

02 Plug in Arduino

Plug the Arduino into your machine using the USB cable, and then open up the Arduino IDE. You should see 'Arduino on Dev/ttyACM0' or similar in the bottom right of the interface. If not then you may need to check if any other dependencies need to be installed.



03 Power check

The most basic 'Hello World' program is to flash the built-in LED on the board. This just makes sure we have everything working and set up correctly. You should already see an illuminated green LED by the 'ON' label, indicating that it is receiving power.



04 Hello World

In your IDE, go to File>New to make a new sketch. The first thing is to write our 'Setup'. Setup will run once when the device is first turned on or reset, and is used for declaring which pins are to be input or output.

Enter:

```
void setup()
{
  pinMode(13, OUTPUT);
}
```

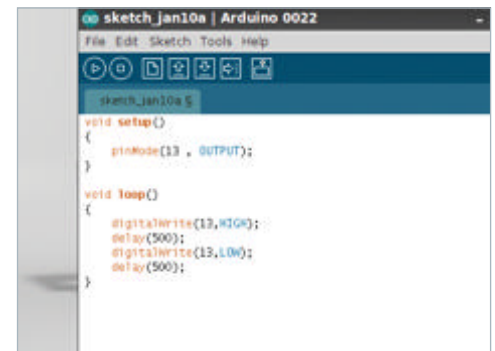


05 Hello World part 2

The previous code sets our 13 pin as an output. Pin 13 on a Uno board is the built-in LED.

Next enter:

```
void loop()
{
  digitalWrite(13, HIGH);
  delay(500);
  digitalWrite(13, LOW); // set the LED off
  delay(500);
}
```



06 Hello World part 3

The previous code sets pin 13 to ON (High), then waits half a second, then sets it to OFF (LOW). As it's a loop with no exit, this will continue indefinitely. Click on the 'Play' icon to verify your code for errors.

07 Hello World part 4

Assuming the code compiled okay, click the 'Upload' button. This will then send our little sketch to the Arduino. You should see the RX LED flash briefly while it's uploading, then the orange 13 LED should start to flash on and off.



08 Output to an external LED

Next up, let's add in one of our external LEDs. Place the longer leg of the LED into your breadboard, and then the resistor next to it. Add the other leg of the resistor to the ground - 'GND' - connection on the Arduino, and then the other LED leg to pin 9.

09 Change pin in code

In our code, we then need to change the output from pin 13 to 9, so our sketch should look like:

```
void setup()
{
  pinMode(9, OUTPUT);
}
```

```
void loop()
{
  digitalWrite(9, HIGH);
  delay(500);
  digitalWrite(9, LOW);
  delay(500);
}
```



10 Add in an external input

Now let's add in an external input to control the LED. On pin 3, connect a jumper wire and run it to the breadboard, then place the legs of the switch next to it. We will set pin 3 to be an input, and control our output on pin 9.

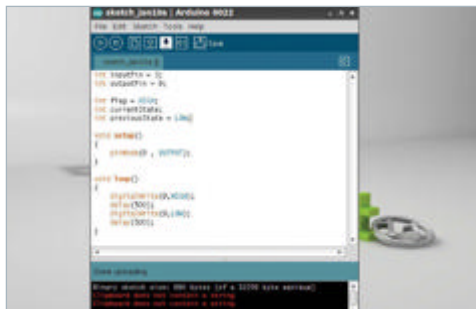
11 Set up our variables

At the top of our sketch, **enter:**

```
int inputPin = 3;
int outputPin = 9;
```

```
int flag = HIGH;
int currentState;
int previousState = LOW;
```

These are the variables we will be using in our loop to set the LED to on or off.



12 The setup

This is very similar to before except we've placed our out and in pins into variables rather than having them constant, which **we use here:**

```
void setup()
{
  pinMode(inputPin, INPUT);
  pinMode(outputPin, OUTPUT);
}
```



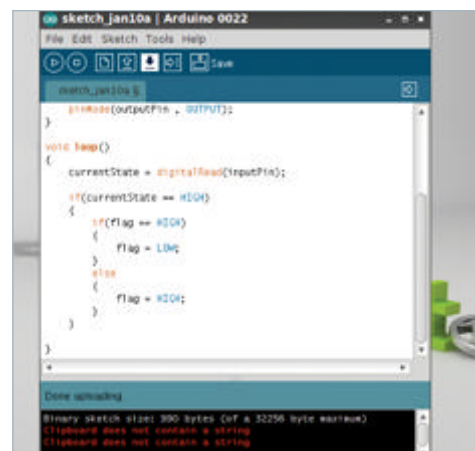
13 Read the current state of output

In our loop, the first line **should be:** `currentState = digitalRead(inputPin);` This reads the current position of our inputPin (HIGH or LOW) and assigns it to the currentState variable. As we have used digitalRead, it will only ever be one of two values: HIGH or LOW.

14 First IF statement

Next we need to find out what our LED is doing and perform an action **based on that.**

```
if (currentState == HIGH)
{
  if (flag == HIGH)
  {
    flag = LOW;
  }
  else
  {
    flag = HIGH;
  }
}
```



15 The completed sketch

The last part of the sketch is this:

```
digitalWrite(outputPin, flag);
previousState = currentState;
```

...which sets our output to the opposite of what it was before and defines our state flag for the next time the switch is pressed. The complete sketch should **look like:**

```
int inputPin = 3;
int outputPin = 9;
```

```
int flag = HIGH;
int currentState;
int previousState = LOW;
```

```
void setup()
{
  pinMode(inputPin, INPUT);
  pinMode(outputPin, OUTPUT);
}
```

```
void loop()
{
  currentState = digitalRead(inputPin);
```

```
  if (currentState == HIGH) {
    if (flag == HIGH)
    {
      flag = LOW;
    }
    else
    {

```

“Maybe you have a product idea and need to produce a proof of concept”

```

    flag = HIGH;
  }
}

digitalWrite(outputPin, flag);

previousState = currentState;
}

```

16 Output the state to the serial monitor

Arduino has a handy serial monitor function built in, allowing you to output to the attached computer. In the setup add 'Serial.begin(9600);'

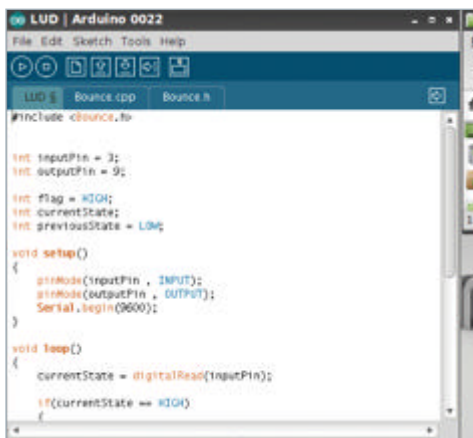
Under the 'flag = LOW;' and 'flag = HIGH;' lines, add in 'Serial.println("LED ON");' and 'Serial.println("LED OFF");' Upload the sketch and click Tools>Serial Monitor.

17 Debounce the input

You may have noticed the LED flickering, and the serial monitor producing more than one 'ON/OFF' message per click. This is due to the fact that some switches don't produce reliable solid 0 or 1 states. Luckily we can solve this by using one of the many libraries available.

18 Import Bounce library

Unzip the Bounce library and put it in your Sketch folder. This can be found by going to Preferences. Next go to Sketch>Import Library; at the bottom you should see Bounce. Add it and you will see '#include <Bounce.h>' at the top of your sketch.



```

LUD | Arduino 0022
File Edit Sketch Tools Help
LUD | Bounce.cpp Bounce.h
#include <bounce.h>

int inputPin = 3;
int outputPin = 9;

int flag = HIGH;
int currentState;
int previousState = LOW;

void setup()
{
  pinMode(inputPin, INPUT);
  pinMode(outputPin, OUTPUT);
  Serial.begin(9600);
}

void loop()
{
  currentState = digitalRead(inputPin);
  if(currentState == HIGH)

```

19 Set up a Bounce object

Above your setup, **add in**
Bounce bouncer = Bounce(inputPin,50);

And in the loop function, replace **currentState =digitalRead(inputPin);** with **bouncer.update ();**
int currentState = bouncer.read();

This should stabilise your input. If it doesn't do so then you can try adding in a delay to the inputs.

20 Count inputs

Let's simply process the number of inputs and outputs to another LED once the count has reached 10. **Add in:**

```

int counterPin = 5;
int counter = 0;

```

before the setup, and in the setup add **pinMode(counterPin, OUTPUT);**



```

LUD | Arduino 0022
File Edit Sketch Tools Help
LUD | Bounce.cpp Bounce.h
int flag = HIGH;
int currentState;
int previousState = LOW;
Bounce bouncer = Bounce( inputPin,50 );
int counterPin = 5;
int counter = 0;

void setup()
{
  pinMode(inputPin, INPUT);
  pinMode(outputPin, OUTPUT);
  pinMode(counterPin, OUTPUT);
  Serial.begin(9600);
}

void loop()
{
  bouncer.update ( );
  int currentState = bouncer.read();

```

21 Start counting on iterations

Then we can add 1 to our count variable each time the switch is turned on.

Underneath our:
Serial.println("LED ON");
 add
counter++;

We can also output the count to our serial monitor:
Serial.println(counter);

22 Add in our ELSE IF

Then we can add in our ELSE IF:
else if (counter >= 10)

```

{
  Serial.println("LIMIT REACHED");
  digitalWrite(counterPin, HIGH);
  delay(1000);
  digitalWrite(counterPin, LOW);
  delay(1000);
  counter = 0;
}

```

```

digitalWrite(outputPin, flag);

```

This will flash our other LED on pin 5 twice and then reset the counter to 0 to start the process over again.

23 Total code

The inner loop code should look like:

```

if (currentState == HIGH)
{
  if (flag == HIGH)
  {
    flag = LOW;
    Serial.println("LED OFF");
  }
  else
  {
    flag = HIGH;
    Serial.println("LED ON");
    Serial.println(counter);
    counter++;
  }
}
else if (counter >= 10)
{
  Serial.println("LIMIT REACHED");
  digitalWrite(counterPin, HIGH);
  delay(1000);
  digitalWrite(counterPin, LOW);
  delay(1000);
  counter = 0;
}
}
}

```

24 Analog output

The Arduino board can also output analogue signals instead of just TRUE/FALSE. To enable this, in your ELSE IF replace the digitalWrite section with:

```

for(int fadeValue = 0 ; fadeValue <= 255;
fadeValue +=5)
{
  analogWrite(counterPin, fadeValue);
  delay(10);
}

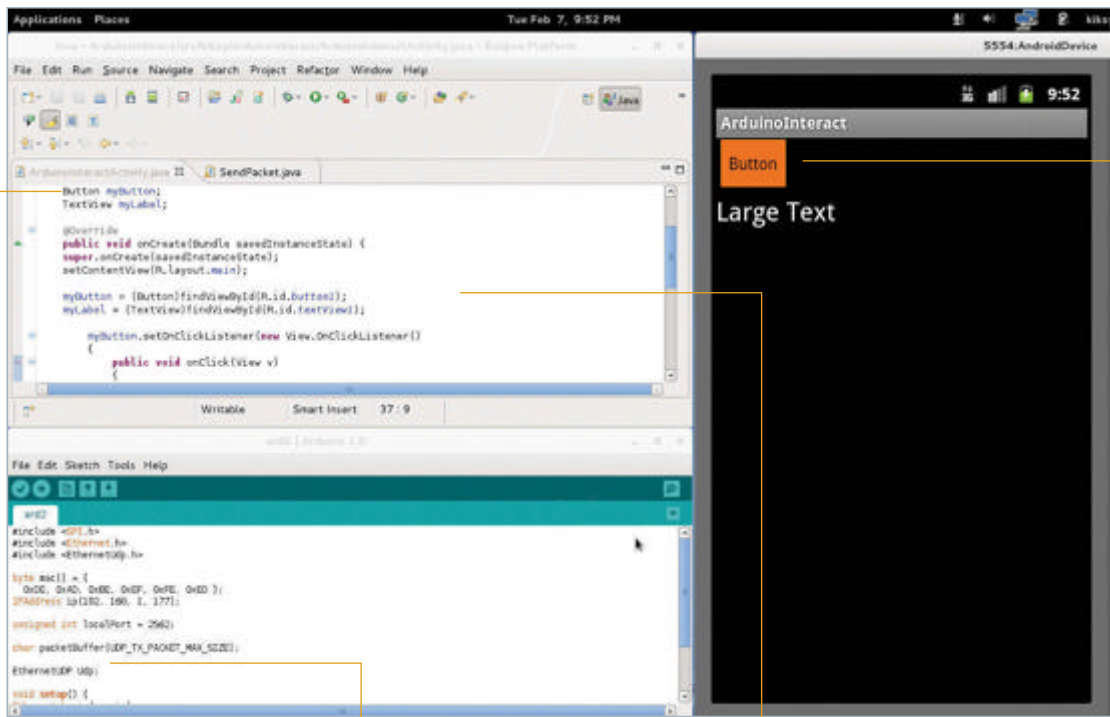
```

```

for(int fadeValue = 255 ; fadeValue >=
0; fadeValue -=5)
{
  analogWrite(counterPin, fadeValue);
  delay(10);
}

```

“The Arduino board can also output analogue signals”



The Android app consists of a simple button, which fires a UDP packet across your LAN to the listening Arduino

The app is just two classes: the view with the button, and then the class which deals with network settings and sending the packet

The Arduino IDE is simple and very easy to use

Although you can use the app in the simulator, it's much more satisfying using it on a real device, especially from another room

How to interface Android with Arduino

How to use your Android device to control and trigger actions on an Arduino board

Resources

Arduino Uno: <http://arduino.cc/en/Main/arduinoBoardUno>

Arduino Ethernet Shield: <http://arduino.cc/en/Main/ArduinoEthernetShield>

Ethernet cable

A-B USB cable

Android device

Android SDK and suitable development environment

EthernetUdp library: <http://tinyurl.com/ctlztp>

Android Easy Packet Blast

app (for testing UDP packets): <https://market.android.com/details?id=com.hunterdavis.easypacketblast>

Advisor

Kieron Howard is a web developer with a penchant for all things Linux. He is hugely excited by the possibilities that Android can bring to the mobile world



In last month's Android tutorial we went through some basic Arduino controls to get used to the IDE and language. This time we will take it a step further and interact with some real-world objects via our Android device, using the Arduino as a controller to process input and output. You'll need to have an Android development

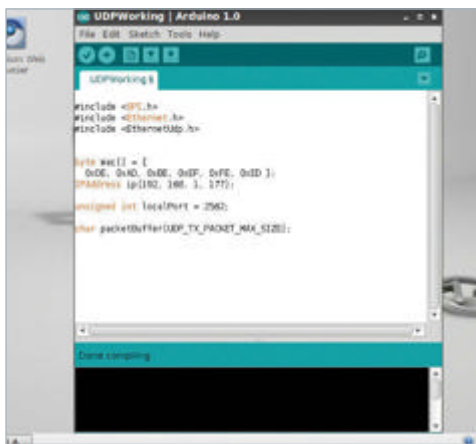
environment set up, and it's assumed that you already have your Arduino IDE ready to go. We'll be using the Arduino Ethernet Shield: a small board that attaches to the top of a compatible Arduino and allows you to attach it to your network and the internet. The tutorial will show you how to produce a simple Android app that broadcasts a UDP message across your LAN, and we'll go through setting up an Arduino sketch which listens for UDP packets. Once the Arduino receives the packet, it will then turn on one of the pins. It would be simple to expand on this tutorial, adding more buttons and then send a different message with each one, which then triggers different activity on the Arduino.

01 Import dependencies

To start off, we need to import the required Ethernet libraries **using**:

```
#include <SPI.h>
#include <Ethernet.h>
#include <EthernetUdp.h>
```

The EthernetUdp library needs to be added to your Sketchbook/Libraries directory. These libraries enable the networking functions of the EthernetShield.



02 Set up network interface

Assign an IP, MAC and port on which to listen to our Arduino:

```
byte mac[] = {
  0xDE, 0xAD, 0xBE, 0xEF, 0xFE, 0xED };
IPAddress ip(192, 168, 1, 177);
unsigned int localPort = 2562;
```

Add a buffer for our received UDP packet:
`char packetBuffer[UDP_TX_PACKET_MAX_SIZE];`

Initiate a new EthernetUDP class:
`EthernetUDP Udp;`

03 Set up our 'Setup'

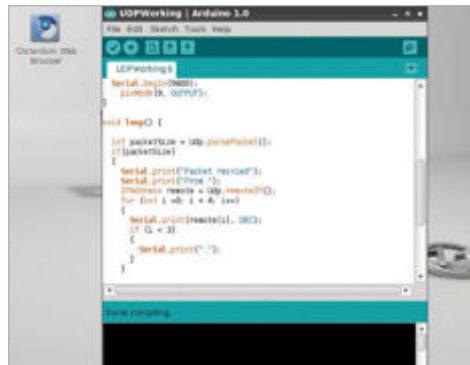
Every Arduino sketch has a setup which runs once when the device is turned on or reset.

```
void setup() {

Ethernet.begin(mac, ip);
Udp.begin(localPort);

Serial.begin(9600);
pinMode(9, OUTPUT);
}
```

This fills the Ethernet and UDP classes with the relative variables we set earlier, then starts our serial monitor and sets pin 9 on our Arduino as an output.



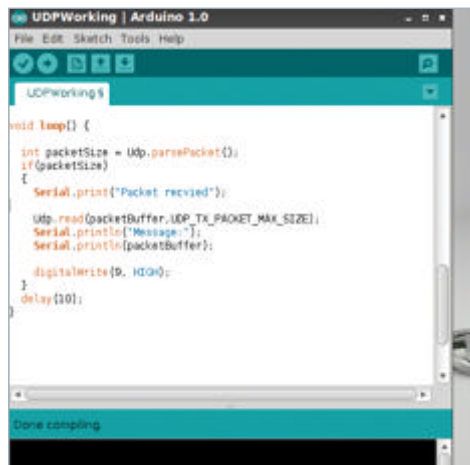
04 Start our loop

Now let's add in our main loop.

```
void loop() {

int packetSize = Udp.parsePacket();
if(packetSize)
{
  Serial.print("Packet received");
}
}
```

This checks if the variable packetSize has value, then the serial monitor will print a confirmation. We can then store and parse the values of the packet into our buffer.



05 Add UDP message to buffer

We use our UDP library to process the contents of our packet buffer, and then print out the message to our serial monitor.

```
Udp.read(packetBuffer, UDP_TX_PACKET_MAX_SIZE);
Serial.println("Message:");
Serial.println(packetBuffer);
```

```
digitalWrite(9, HIGH);
```

Lastly we then set our pin 9 to ON.

06 Add delay

Add a delay of 10ms to give the Arduino time to breathe. Your loop should look like:

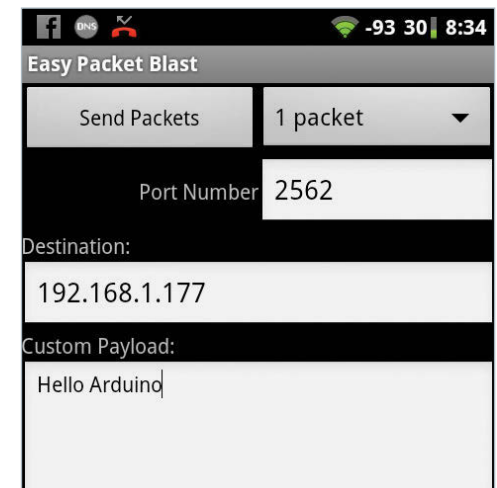
```
void loop()
{
int packetSize = Udp.parsePacket();
if(packetSize)
{
  Serial.print("Packet received");

  Udp.read(packetBuffer, UDP_TX_PACKET_MAX_SIZE);
  Serial.println("Message:");
  Serial.println(packetBuffer);

  digitalWrite(9, HIGH);
}
delay(10);
}
```

07 Plug into network

If you want, you can add a LED on pin 9 for some real-world output. Plug your Ethernet cable into the Shield, then into your network's router or switch. Upload the sketch and start the serial monitor by clicking on the magnifying glass in the top right of the IDE.

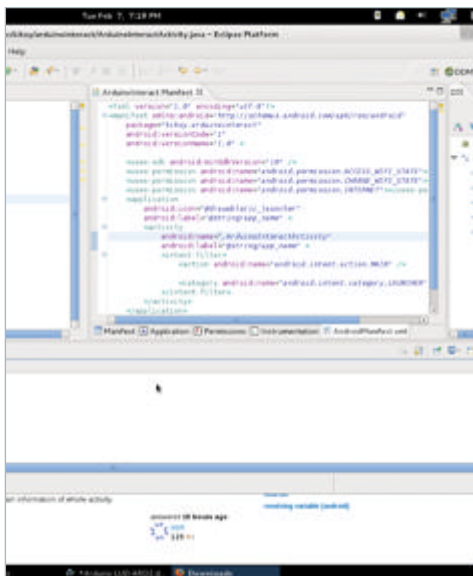


08 Test UDP listener

On your Android device, run the Easy Packet Blast app. Enter the appropriate port and IP, then add in a little message. Make sure your Android device is on the same network and click 'Send Packets'. You should see the message appear in the serial monitor.

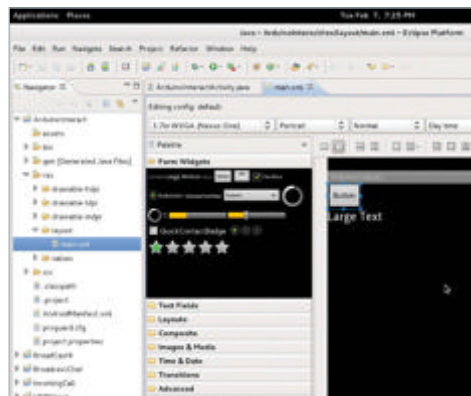
09 Start the Android app

Open up Eclipse and set up a new Android project from File>New>Android Project. Enter in the details and choose your target. Once you've done that, we can start our simple button app to send UDP packets to the Arduino to tell it to do something.



12 Add in the graphical button

In the graphical window tab, you can see your main app view. Delete anything currently on the screen, then replace with a Button and Large Text from the 'Form Widgets' tab on the left. Open the 'main.xml' view and make a note of the button and textview IDs.



14 Add button clickListener

As we want our button to do something when pressed, we need to add a clickListener to it and set up an action for when it's clicked.

```
myButton.setOnClickListener(new View.OnClickListener() {
    public void onClick(View v)
    {
        // Action to perform when button is clicked
    }
});
```

15 Tell the app to send packet

```
Add:
myLabel.setText("Packet Sent");

SendPacket send = new SendPacket();

send.onReceive(getApplicationContext()
, getIntent());
```

to the onClick method. This will change the button label to 'Packet Sent' when you click it, then we initiate a new SendPacket class (which we will build in a minute) and tell it to run the onReceive method. In your src>name>'project' folder add a new class File>New>Class and call it SendPacket. This class will handle all our networking, generate our UDP packet and broadcast it on our LAN.

10 Main interface class

First, import the required dependencies:

```
import android.app.Activity;

import android.os.Bundle;

import android.view.View;
import android.widget.Button;
import android.widget.TextView;
```

and then add these permissions to the AndroidManifest.xml...

```
<uses-permission
android:name="android.permission.ACCESS_WIFI_STATE"></uses-permission>
<uses-permission
android:name="android.permission.CHANGE_WIFI_STATE"></uses-permission>
<uses-permission
android:name="android.permission.INTERNET"></uses-permission>
```

11 Add in the button reference

Then we start our main class:

```
public class ArduinoInteractActivity
extends Activity
{
    Button myButton;
    TextView myLabel;
}
```

And then declare our button we are going to make in a second, along with a text label. Switch to the res>layout>main.xml where we will add in our button image.

13 Set up onCreate()

The onCreate method will run as soon as the app starts. Here we find our button using the IDs we just noted down and label and assign them variables.

```
@Override
public void onCreate(Bundle
savedInstanceState)
{
    super.
onCreate(savedInstanceState);
    setContentView(R.layout.main);

    myButton = (Button)findViewById(R.
id.button1);
    myLabel = (TextView)
findViewById(R.id.textView1);
}
```

16 Import class dependencies

```
import android.content.BroadcastReceiver;
import android.content.Context;
import android.content.Intent;

import java.io.IOException;
import java.net.DatagramPacket;
import java.net.DatagramSocket;
import java.net.InetAddress;
import java.net.SocketException;
import java.net.UnknownHostException;

import android.net.DhcpInfo;
import android.net.wifi.WifiInfo;
import android.net.wifi.WifiManager;
```

“This will handle all our networking, generate our UDP packet and broadcast it”

“It would be simple to expand on this tutorial, triggering different activities on the Arduino”

17 Declare our class

Add in the new class:
public class SendPacket extends BroadcastReceiver {

```
private static final int UDP_SERVER_PORT = 2562;
Context mContext;
DatagramSocket mSocket;
InetAddress myBcastIP, myLocalIP;
}
```

Then we add in a port – make sure this is the same as you set earlier in your Arduino sketch.

18 onReceive method

We call this method when the button in the app is clicked:

```
public void onReceive(Context context, Intent intent) {
    DatagramSocket ds = null;
    mContext = context;
}
```

The DatagramSocket is specifically for sending and receiving UDP packets.

19 Set up our sockets

Now add in the sockets, and get the broadcast address using the appropriate getter (all needs to be in a try-catch block):

```
try {
    ds = new DatagramSocket();

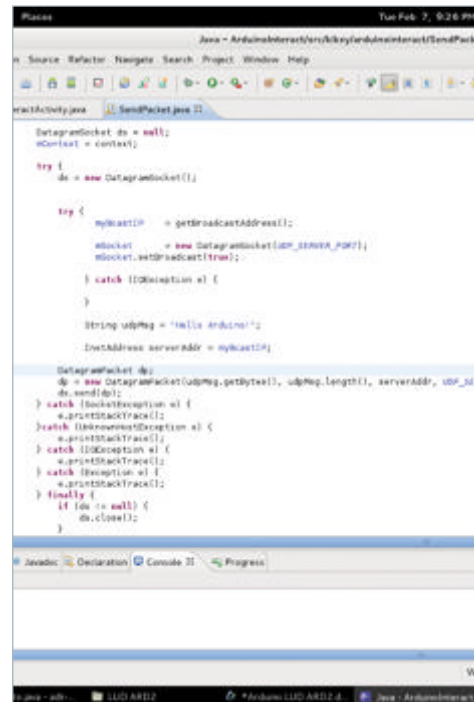
    try {
        myBcastIP = getBroadcastAddress();

        mSocket = new DatagramSocket(UDP_SERVER_PORT);
        mSocket.setBroadcast(true);
    } catch (IOException e) {
    }
}
```

20 Add in a message

Then we add in a message to send, specify the broadcast address and then send the packet. This needs to be in a try/catch block too.

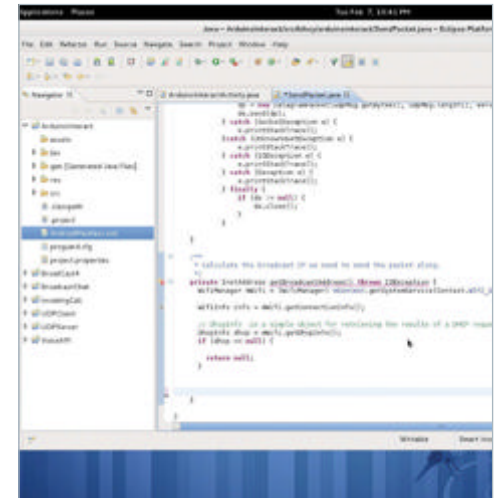
```
String udpMsg = "Hello Arduino!";
InetAddress serverAddr = myBcastIP;
DatagramPacket dp = new DatagramPacket(udpMsg.getBytes(), udpMsg.length(), serverAddr, UDP_SERVER_PORT);
ds.send(dp);
```



21 Try/catch the send

Then we need to catch our possible thrown exception:

```
} catch (SocketException e) {
e.printStackTrace();
} catch (UnknownHostException e) {
e.printStackTrace();
} catch (IOException e) {
e.printStackTrace();
} catch (Exception e) {
e.printStackTrace();
} finally {
if (ds != null) {
ds.close();
}
}
```



22 Add in our address getter getBroadcastAddress

Then we add our private method called in onReceive():

```
private InetAddress getBroadcastAddress() throws IOException {
    WifiManager mWifi = (WifiManager) mContext.getSystemService(Context.WIFI_SERVICE);

    WifiInfo info = mWifi.getConnectionInfo();

    DhcpInfo dhcp = mWifi.getDhcpInfo();
    if (dhcp == null) {
        return null;
    }
}
```

23 getBroadcastAddress part 2

The last part gets our IP from our DHCP server...

```
int broadcast = (dhcp.ipAddress & dhcp.netmask) | ~dhcp.netmask;
byte[] quads = new byte[4];
for (int k = 0; k < 4; k++)
    quads[k] = (byte) ((broadcast >> k * 8) & 0xFF);
return InetAddress.getByAddress(quads);
}
```

24 Compile and run

Run the app on your phone, making sure you are connected to your LAN, press the button and you should see the LED light up if you have one. Here we have instead connected the Arduino to the remote shutter release on a DSLR to give you an idea of the possibilities.

Getting Things Done



■ The Hipster PDA – go analogue: no battery worries with this mobile list app

Advisor

Richard Smedley is a father and self-confessed geek, who wants all kids to try the creative magic of coding. He set up teachyourkidstocode.org to find the best ways of starting this off



Discover how open-source software offers ways for you to streamline your workflow and improve your productivity



Ten years ago, David Allen developed **Getting Things Done (GTD)**, a methodology he calls “**Stress Free Productivity.**” The system works

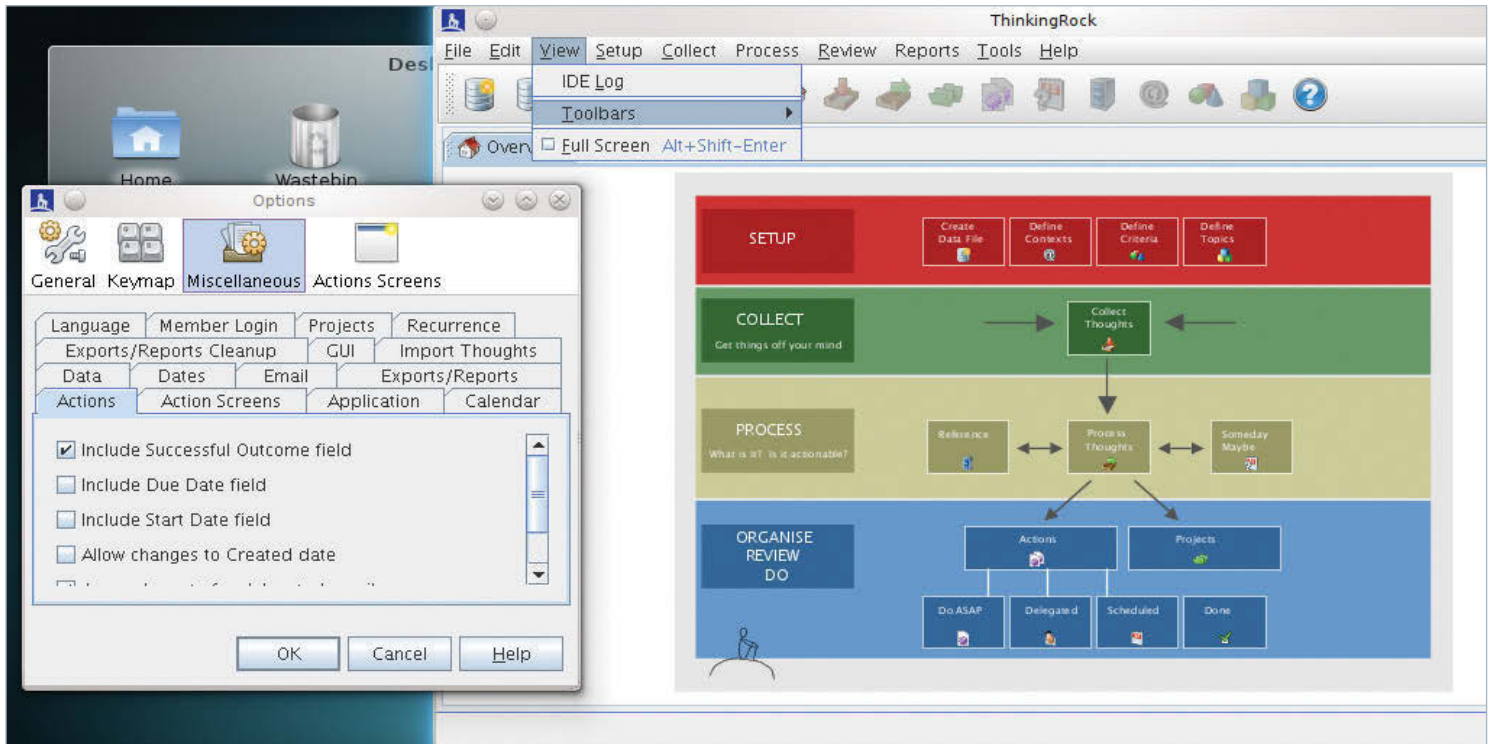
by relieving the brain not just of remembering what’s going on and how to deal with it but, more importantly, distinguishing between what’s urgent and what’s actually important.

Wired called GTD “A new cult for the info age”, but forget the hype – GTD is a system that works for busy people who have too much to keep track of. Let’s give you a high-speed run-through of the methodology before looking at the software.

It’s all about clearing the mind’s inbox of ‘stuff. This is accomplished by turning all items

into ‘next actions’ and placing them in a trusted system that enables you to stop fretting over remembering things, and therefore get on with getting things done.

Next actions are then stored in whatever form suits you best. You can even use a paper-based system, but we’re looking here at various free software solutions. Many of these apps are cross-platform. Of course, there are cloud-based alternatives (and implementing GTD in Gmail is well documented), but with something this crucial to your day-to-day productivity, you may not wish to surrender control of your system to somebody else’s servers and the continued beneficence of their business model.



■ Configurable, comprehensive, and very faithful to GTD methodology – ThinkingRock wants to run your life

The transformation of Stuff

Initially, you capture all of the ‘stuff’ floating round in your head, your inbox, on your desk, wherever. In the first chapter of GTD, Allen defines Stuff as: “anything you have allowed into your psychological or physical world that doesn’t belong where it is, but for which you haven’t yet determined the desired outcome and the next action step”. So, the lawn needs cutting, you’re nearly out of milk, you have a report to write for Friday morning, you have to renew a prescription for your elderly father, the front door needs repainting before winter. The mind is a jumble of tasks, and the subconscious cannot distinguish them by urgency or importance, so it shouts to you about all of them, leaving you feeling overburdened.

Setting up a GTD system, it can take you *hours* to pile all of this up, writing things down on slips of paper as they occur to you and throwing them into an inbox. Now everything there – and everything that enters your purview in future – must be processed mercilessly.

First, can you get rid of it? If it doesn’t demand an action, can it be junked or archived? If it’s for an indefinite future date, add it to your ‘someday/maybe’ list, which will hold it pending review. What’s left is what demands action. Separate out

‘projects’, things that need working through in a number of steps: you will do these in the extra time GTD creates for you, and working through them will produce more next actions.

Can the next action be done in an arbitrarily short time, say two minutes? If so, do it. Job done. If not, is it somebody else’s task? If so, hand it off, but record it on your Waiting list if you’re waiting for a reply or result. Now work through the rest of the pile – anything that needs doing at a specific time goes on the calendar, every other next action goes on your to-do list. Group items by opportunity – so phone calls are listed together. Write the numbers there on the list with them. Now when you have a spare 15 minutes between meetings, you can work your way down the list.

Notice the one list you haven’t given yourself? No daily to-do list: each day you do the calendar items, and then the appropriate Next Action items, supplemented by larger project work tasks. The Next Action, Waiting, and calendar should get a weekly review, along with the project plans. The GTD book has more on this, and on project planning, and is the recommended starting point, but here we turn to software to manage your system once you start implementing GTD.

The software

Free and open source software programmers have scratched their collective itch to produce a choice

of GTD apps for everybody’s situation. To be usable, the software must make it easy to capture to-do items and put them in the right place for timely retrieval and review. Some portability is close to essential, too. Let’s start with what you may already got installed on your system, and look at what’s available within GNOME.

Tomboy – GNOME’s little sticky notes – has a hypertext system of autolinking, and recognising email addresses, which brings Next Actions to just a click away. The necessity for a Mono runtime may discourage non-GNOME users from trying it, but they’ll miss out on simple wiki-like linking, instant access from the Panel, and a streamlined but powerful editing interface that may come as a surprise to those who’ve just used it for the odd Post-It note-style reminder.

Tomboy plug-ins (known as add-ins) enable back-link following, Evolution import (drag an email onto a note), and export to HTML, LaTeX or graph. The website has a walkthrough tutorial on creating an add-in, which may be of interest to anyone learning C#. Tomboy isn’t confined to GNOME – it runs on Windows and Macs – and Tomboy GTD tutorials across the web testify to its popularity.


Anyone holding onto the old version of GNOME desktop might like to try Getting Things GNOME, whose special tricks include a script to turn Evolution or Mutt emails into tasks. Another GNOME option in development is Tasque –

Getting GTD Done

Want to get productive with GTD? Start here...

While David Allen jets around the world charging CEOs and VPs thousands of dollars for personal 'training' in GTD, he's also made the whole thing available in the form of the book, *Getting Things Done* (ISBN 978-0749922641), which gives you everything you need to know. The Wikipedia page on GTD is a fair introduction if you're curious but not ready to invest in a paperback. You can also read more at Allen's website, www.davidco.com, but it's hard to begrudge paying a tenner for the book – even if you don't choose to follow all of the GTD methodology, it'll get you thinking about how you work.

Once you're ready to roll your mental inbox into a series of 'next actions', there are countless GTD resources on the web. While we round up the FOSS solutions, you'll pick up tips into optimising the system at 43folders.com and elsewhere. Just remember that reading about productivity improvements is not a substitute for sorting out your system and applying it.



```
richard@netbk10:~$ task add Install Debian on new netbook
A configuration file could not be found in /home/richard

Would you like a sample /home/richard/.taskrc created, so taskwarrior can proceed? (y/n) y
Created task 1.

richard@netbk10:~$ task 1 annotate Use the netinstall CD
Annotated 1 with 'Use the netinstall CD'.
richard@netbk10:~$ task 1 annotate Manually partition: 4GB swap; 10GB left free for distro trials
Annotated 1 with 'Manually partition: 4GB swap'.
bash: 10GB: command not found
richard@netbk10:~$ task 1 denotate Manually partition: 4GB swap;
Found annotation 'Manually partition: 4GB swap' and deleted it.
richard@netbk10:~$ task 1 annotate "Manually partition: 4GB swap; 10GB left free for distro trial
Annotated 1 with 'Manually partition: 4GB swap; 10GB left free for distro trials'.
richard@netbk10:~$ task add "Add http://emac.s.naquadah.org/ repository for Emacs24 snapshot"
Created task 2.

richard@netbk10:~$ task add Copy home dotfiles from the blue memory stick
Created task 3.

richard@netbk10:~$ task 1,2 +online
Modified 2 tasks.
richard@netbk10:~$ task long

ID Project Pri Added Started Due Recur Countdown Age Deps Tags Description
2 11/25/2011 2 mins online Add http://emac.s.naquadah
3 11/25/2011 1 min Copy home dotfiles from
1 11/25/2011 7 mins online Install Debian on new ne
11/25/2011 Use the netin
11/25/2011 Manually part
3 tasks
```

■ Taskwarrior is simple to start with, and its deeper abilities are well documented

available in the openSUSE and Fedora repositories, whose multiple back-ends include the popular Remember-the-Milk web service.

Command line

Mikaël Navarro was looking for something lightweight and simple to carry round on a USB stick and use on different platforms. Existing solutions were missing features or were hard work to look after, so he made his own: yaGTD – Yet Another Getting Things Done. It combines iKog's powerful command-line interface and the Stephen 'Seven Habits of Highly Effective People' Covey categories of urgency seen in pyGTD.

It runs on any platform with Python 2.3 or above, with:

```
python yagtd.py
```

There's a version for the Symbian S60 platform, and an Emacs major mode for reading yaGTD lists. Development has slowed recently, but for a lot of people this already does everything that they need.

Sticking with our lovely plain text files, the people at todotxt.com have given us the tools to easily work with a file saved as `todo.txt` – whether on the command line, or from a smartphone. As the authors say, "A `todo.txt` file future-proofs your task list and stores it in a text file YOU control. That's why a `todo.txt` file is the discerning nerd's task manager of choice." Cross-platform use is possible with Cygwin on Windows.

`todo.sh` is a neat little script that makes it simple to add, update and re-prioritise Next Actions on your list. Tagging enables you to group items in a GTD way, such as phone calls to make. Lacking reminders, due dates and other time-sensitive components, it's not pretending to run your life for you, but it's a great way to manage your list if you're often at the command line – or on a mobile device: `Todo.txt Touch` integrates with Dropbox to provide your Android phone (iPhone version in development) with the same management of your list. If you don't mind using Dropbox (which does have command-line tools for your PC and servers), it could be the solution for you.

Support ninjas

Taskwarrior is inspired by `todo.sh`, but seeks to improve upon it in features and layout. Easy to install – it's in almost every repository, and consists of simple code depending on only the most basic of libraries – it will run on any *NIX system.

One possibly unique feature is its dedicated support ninjas – earlier this year writer Frankie Sachs blogged that he'd posted on Twitter that he couldn't figure the way to set up a daily recurring task on @taskwarrior, and the next morning he'd had an email explaining how – Taskwarrior support had tracked him down to help fix his problem!

The new version has a server-based option, with mobile clients on the way. Time tracking



■ Get things done with Wunderlist



“Small things done consistently over time have major impact” David Allen

is in development, with further features added in response to a recent survey and other user feedback.

The Emacs way

Staying with the command line, but moving to something a little more heavyweight, it's time to try things the Emacs way. For those who missed our Real-World Emacs series (check out **Linux User & Developer** magazine issues 91-3), Org-mode builds on Outline mode's ability to build hierarchical lists from simple text, and adds in connections to calendar (use Planner.el and a PHP hack to publish your calendars), address book (BBDB), emails (a large choice of client software in Emacs) and a powerful editor with version control and SSH, meaning you can maintain your lists easily across machines. In fact Org-mode can do just about anything – following the absorbing links at <http://orgmode.org/worg/org-tutorials/> brings volumes of info, including Natural Project Planning with Org-mode – using Allen's five-step process of 'The Natural Planning Model'.

Org-mode can even be used away from Emacs, in its Android form, and the lists are text-based, so can be edited anywhere (including vim!), so take a look even if you're not prepared to

embrace the Emacs view of the universe. Whereas most GTD apps start from tasks (or at least Next Actions), Org-mode builds up from notes, from which the tasks emerge, allowing for free-thinking and development – Org-mode notes can also be exported to FreeMind to produce mindmaps. We find it a winner, combining the simplicity of text with the flexible, powerful views built on Outline mode, and good integration with other Emacs features – but we're prepared to admit it's not for everybody.

GUI goodness

For those resistant to the command line – or those needing cross-platform tools, Java-based GTD-Free can (like yaGTD) be run anywhere from a USB stick. Its clean interface is designed to draw GTD newbies through task workflow with minimal training. It organises actions by projects, remind dates and priority, and features some nice touches, such as export to PDF, HTML and XML documents. This is a good bet for those who are comfortable with traditional Windows desktop apps.

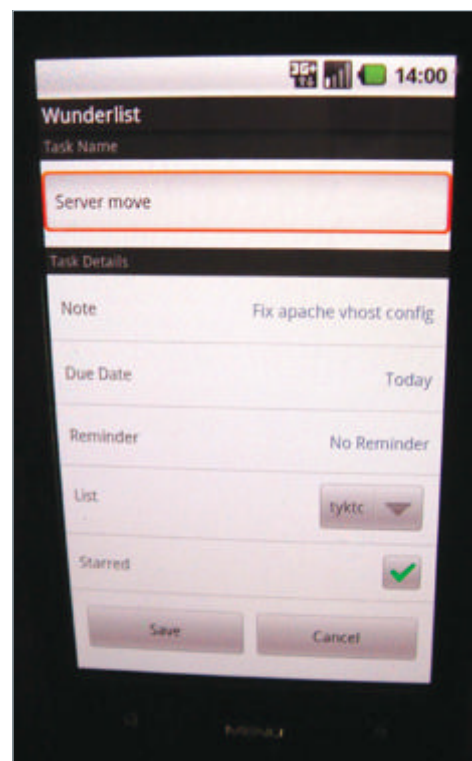
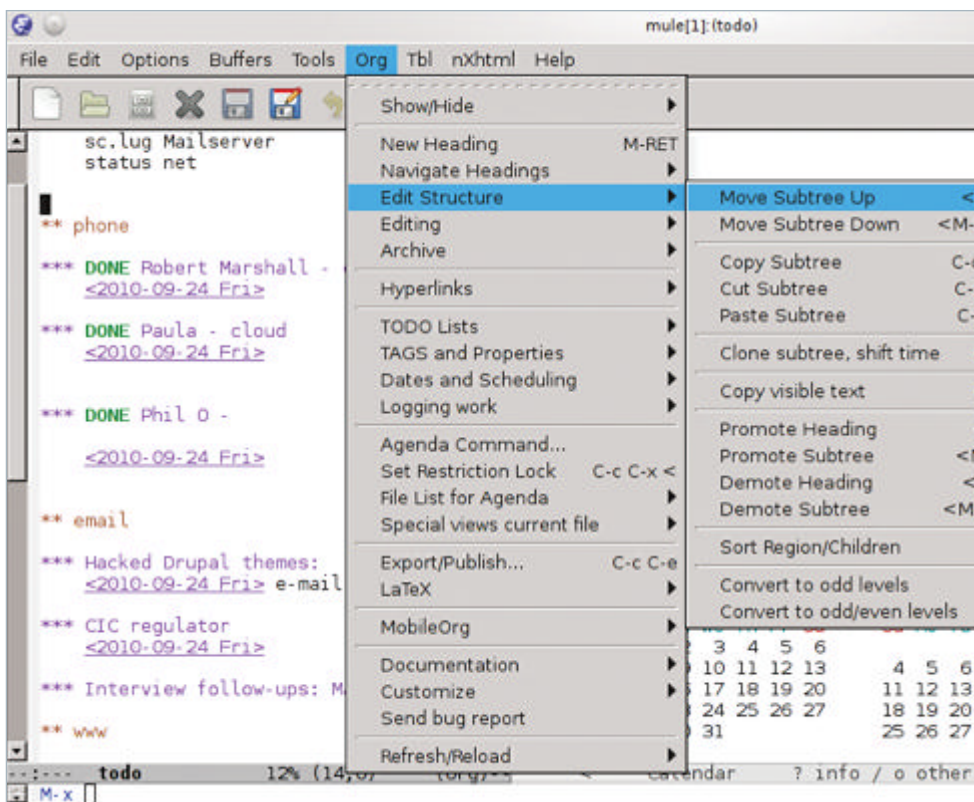
Thinking Rock, also Java, is the most comprehensively GTD of all the apps covered here, in sticking to the letter of Allen's workflow

model. Although you have to purchase the latest version on the website, you can download old versions under the Common Development and Distribution License (CDDL). The nearly three-year-old 2.2.1 is actually a good app, if you don't mind the overhead of Java as the trade off for cross-platform compatibility.

Although GTD-Free and Thinking Rock are both Java, they're different in approach, with Thinking Rock being something of an overwhelming experience, wanting to take over your life and put your every movement in a GTD system, while GTD-Free is eager to please, but simpler, and very quick to work with. Occasional freezes in both could be put down to Java Virtual Machine problems on Linux, but that's small comfort to those affected.

'S wonderful

Apache-licensed Wunderlist is barely a year old, yet has drawn many plaudits and gathered more than a million users, and already looks a mature and capable program. From the install, which takes you to the simple tutorial, through to entering data, the user is shown a friendly interface that doesn't bog them down in unnecessary detail. At its simplest Wunderlist enables you to create sharable lists with no learning a system of doing



■ Sync your to-do lists from Android with Wunderlist

“Building your GTD system around your email client makes a lot of sense”

things on the software. Implementing GTD is well supported by the app's features, however, and its understanding of human-readable input like 'in a month', various shortcuts, and features like a daily email of all overdue tasks, will win you over.

If you want to hand out tasks to others, you can print or email from Wunderlist, or publish in HTML in the cloud. You can share the list location through social media – a special touch that makes it great to use on mobile devices. Sharing with other Wunderlist users is also easy, and you can invite others to join from the same interface. You can even export diary snippets as PDFs.

Once you register a Wunderlist account, your data is backed up automatically to Amazon S3 cloud storage, and available to you from any version of the app you use, as well as through a web interface. Available in 64- or 32-bit, as well as versions for Windows, Mac OS X, iPod, iPad, and a very nice Android app, so while the registration and the business model of a company buying cloud services and giving them to you for free (they have a premium product which should fund all of this) may cause you to hesitate, the quality of the software as well as its availability across multiple

devices will put it near the top of the evaluation pile for many. If you're a GTD obsessive you'll find it doesn't have every last touch that Thinking Rock has, but for those who want something that lets you quickly work on your lists without jumping through hoops, Wunderlist is perfect.

On the web

Web-based GTD apps, such as NextAction (built on Google Gears) have come and gone, but with even personal software becoming web-based, thanks to the proliferation of mobile devices, expect to see far more development here. One of the most popular is Rails-based Tracks, which presents a clean, well-considered interface. For example, tasks change colour through red, orange and green as completion date looms. Shipping with its own web server, you can run it on your desktop, or stick it on a server where any of your devices can see it – it has a lightweight mobile interface for your smartphone. Tracks is multi-user but not collaborative. The AJAX interface and autocompletion score bonus points.

Tracks isn't limited to web interactions: the API allows for shell script, Ruby script or AppleScript

so you could write something to get data in and out to interact with the rest of your software. Tracks also provides feeds of your appointments and lists in various formats, and exports data in YAML, CSV or XML form. The interface is very configurable, with context views for different office tasks, and the usual tag features. Add in all the pretty progress graphs on the statistics page and you may find you're too busy playing with Tracks to actually get anything done.

Other web solutions include LAMP-based TaskFreak, which is slowly evolving time-tracking and multi-user features – but one piece of web software that suggests itself above all else for editable lists is the wiki. While any wiki could be used for GTD – the simpler the better, both in file storage and in system overhead – TiddlyWiki is favoured. In fact, GTD TiddlyWiki Plus is TiddlyWiki with page templates and stylesheet for editing and producing lists to be printed on 3-inch by 5-inch cards for the famed GTD Hipster PDA, or 'Personal Analogue Assistant'.

Another TiddlyWiki offspring is MonkeyGTD, last year renamed mGSD (as in Getting Stuff Done) following a legal letter about using the trademarked term GTD in its name. Installation is as simple as downloading an empty MonkeyGTD file, then opening it in your browser. Instead of pages, TiddlyWiki has tiddlers, and multiple tiddlers can be open on the same page

simultaneously. The documentation covers using mGSD with all areas of GTD workflow, and this app can be extended by writing plug-ins in JavaScript.

Despite a recent slowdown in development, we recommend mGSD as one to try for anyone who's got familiar with GTD and needs to quickly get going with something simple yet flexible. Further support is available in the GTD TiddyWiki Google Group – which covers both the above and any other variation of that wiki in a GTD system.

You have mail

Many people use the workflow of their email system, whether on Evolution, Thunderbird or, if forced to at work, Outlook. These can be synced to mobile devices, and a number of plug-ins for Evolution (now integrated into the latest release) make GTD even easier, and you can open GNOME Planner files as Evolution tasks. KDE's PIM also makes a good GTD tool if that's your desktop of choice.

Years before Android, some people used jPilot with the Palm-based Treo smartphones to carry GTD lists on the move, syncing with Outlook in the office if required. Indeed, with email still the centre of communications and task allocation in most offices, building your GTD system around your email client makes a lot of sense.

IMAP means you can carry this off from any internet-connected device. Plug-ins for Thunderbird, or simply using its semantic tags, makes it an ideal solution already on your desktop, possibly. Nowadays, social media is as big a distraction from work as email used to be, so reduce your distractions by collecting all of the notices of activity automatically issued by Twitter, Facebook, Google+, Diaspora*, LinkedIn, and whoever, in a mail folder where they quietly await attention at an appropriate point in the working day.

Group GTD

Can you take this GTD methodology to a whole company? Plainly, if you are a project manager, you can incorporate the essence of GTD into your team's work – but what about software that works across a team for GTD? Here there are few perfect fits, but most collaboration software, such as FengOffice (formerly OpenGoo – see LUD issue 90), can be adapted – however, we should make special mention of Chandler.

Chandler is well known to many in the FOSS community as Mitch Kapoor's big folly – millions of dollars blown on a project to produce an Outlook/Exchange replacement that never came to fruition. However, when the money ran out, a core team brought the software to version 1.0 in a form repurposed for GTD, and now describe Chandler as a "Note-to-self organiser". Like Wunderlist, Chandler offers online storage (the Chandler Hub),

Inbox Zero

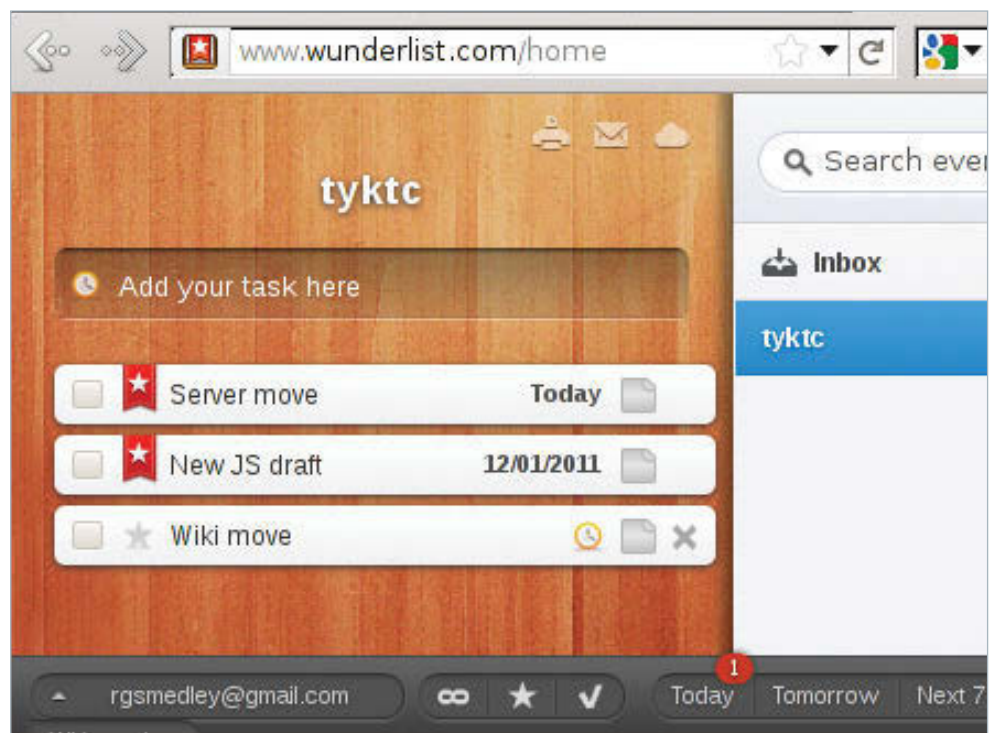
Declutter that inbox and free your mind

While GTD was conceived in a largely paper-based world, its spread to the virtual world, and evolution by numerous adherents, has produced many tweaks. Merlin Mann's 43folders.com site developed Inbox Zero to "reclaim your email, your attention, and your life."

Essentially, you limit your email reading to a small number of sessions – say two or three a day – and quickly triage the email using the GTD filtering system: does it require an action, or can it be junked or archived? Can it be answered in two minutes (a great discipline for much

correspondence) Can it be delegated? Is it time sensitive? Move it to a calendar event. Add it to a next-action-driven folder (or add the action to the list and archive or delete the mail).

This process is not about having no email in your physical inbox, but about having none there nagging at the back of your mind. If your inbox is too big for you to start on it right now, move everything to a folder called DMZ (demilitarised zone!), and start with emails that come in today. If you don't get a chance to process the contents of the DMZ, perhaps none of it was that important?



accessible through a web interface as well as the Chandler desktop client, and the ChandlerQE client for Android and iPhone. You can also configure your own hub with the Chandler server.

GTD features include categorising lists into collections to group say, all of the tasks you have to carry out in the office, or all your phone calls. It's never taken off, and can be hard to get running on modern distros, but the familiar email-based system of workflow and three-paneled window demands less contortion to fit most people's work patterns, and should be seriously considered by anyone looking for

a collaboration solution, despite recent lack of development.

A final word: there are hundreds of apps, and thousands of webpages, to help you implement the ideal GTD system for you, and to tweak until perfect. Just remember that unnecessary hours tweaking GTD is just another form of procrastination, so start simple: read the GTD book and try an email or text list or wiki-based solution to get you going. Once you have a working system, and you've used it long enough to really know its advantages and shortcomings, return here and consider if other options are necessary.

Masterclass

Become a Linux power user

106 Back up & restore

Safeguard your system with our guide to saving and recovering files

112 Perfect dual boot

Manage your system's multiple personalities and master dual booting

120 Keep your system safe

Discover the holes in your system and find out how to deal with them

126 Start building a blog with Django

Use this Python-based web framework to create a blog

130 Add content to your Django blog

Continue building your blog and master front-end content delivery

134 Customise your blog in Django

Learn about some of the advanced features you can utilise with Django

138 Begin a bug hunt in LibreOffice


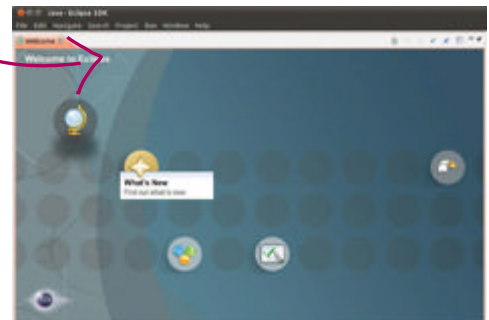
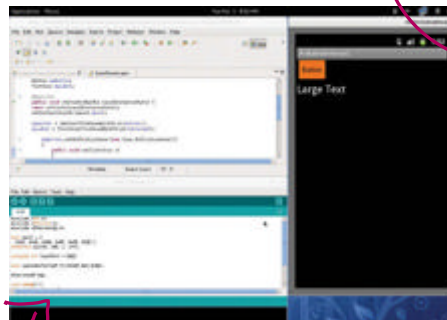
Get involved with the open source community and lead a big hunt

142 Find and fix known bugs in LibreOffice

Find a problem, resolve the issue and submit your bug fixes


146 Become an open-source mentor in LibreOffice

Share your knowledge with others and help build a community






es



```
#!/BIN/SH -E
ECHO "ADDING ANOTHER OS"
CAT << EOF
MENUENTRY "ANOTHER OS"
SET ROOT=(HDD,X)
LINUX /BOOT/VMLINUX-Y RO QUIET
INITRD /BOOT/INITRD.IMG-Y
}
EOF
```





Backup & restore

Safeguard your system from disaster with our definitive guide to saving and recovering your files



It's probably an old wives' tale, but we've heard it said that folks with more technical know-how and more data are less likely to back up their

files and systems in case the worst happens.

It's definitely a lot more common to see robust and thorough methods of backup in business, though, as compared to home use. However, it's never been easier to keep your data secure from unfortunate incidents.

Whether you just want to save a few important files or create a complete image of your hard drive, there are plenty of ways to achieve both. With the popularity of cloud storage, the capacity of modern hard drives and the refinement of software, there are also more options than ever before to keep your system safe.

In this guide we will cover some of the best ways to perform system backups, from command-line operations to powerful graphical tools and finally creating a full-blown disk image of your hard drive if necessary. We'll even throw in some tips on various types of storage, and a small guide on how you can make a live disc of your OS.

Backing up is only half the battle, though, and of course just performing a backup doesn't automatically mean that any future errors will be instantly rectified. Knowing how to restore your system is as important as being able to back it up. We will show you how to make the most of your backups if the time ever comes to use them, employing the same or similar software as what we're using for the backup to make it as simple and hassle free as possible.

Advisor

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

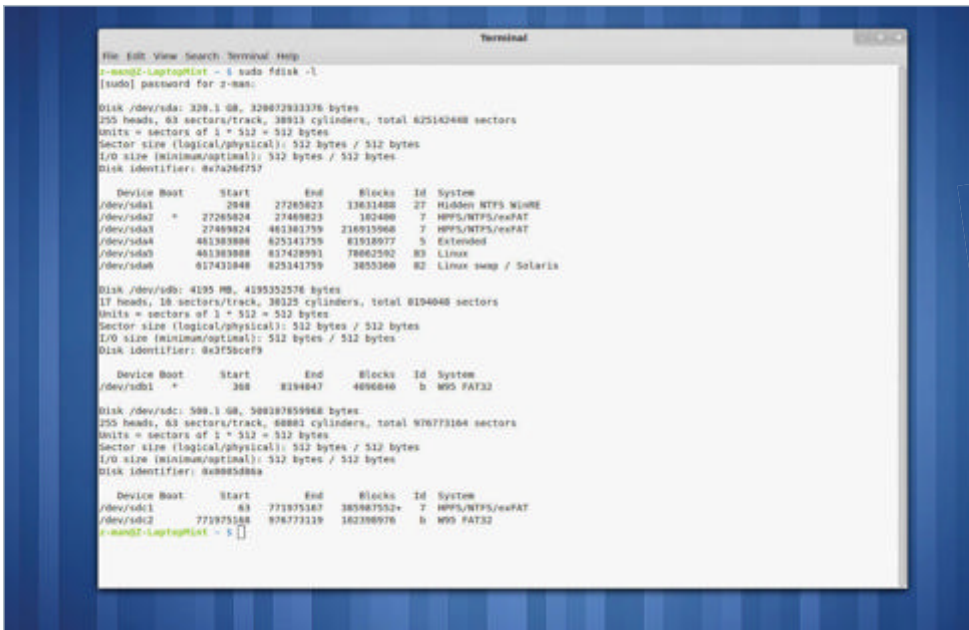


Play it safe

We know, it sounds a little ridiculous; this guide is to help you out in case problems arise. However, if you're at all doubtful about any of the methods we employ in the following pages, test them on non-important files and folders just to make sure you're not putting a sudo in the wrong place or clicking on the wrong directory. It happens to the best of us.

File backup quick-start

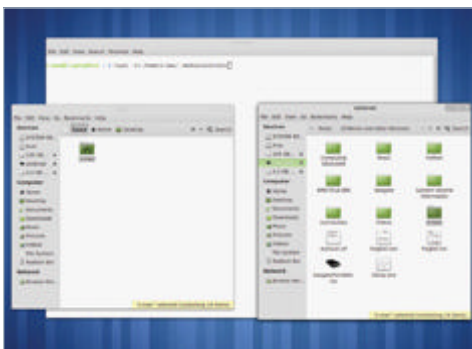
The simplest way to back up files is using the rsync command – in most Linux distros by default, and easily installable otherwise



01 Source and destination

You will need to know the relevant folder paths for the source and destination of your backup. The default home folder in Linux is found

at `/home/user/`, where 'user' is of course your username. You can also look at the properties of a folder in a window manager to see the path.



02 Use rsync

Rsync commands work as such:
rsync options source destination

We are going to use options `r` and `z`; `r` is recursive, and `z` will compress files. In our example, we're saving to an external hard drive with this command:
rsync -rz /home/user /media/external

03 Better rsync

There are a few more options for rsync that you may want to consider. Using `-v` will show you each file rsync processes, so you know it's actually working. Using `-a` will preserve the file timestamps, meaning that future backups will only change newer versions of files.

Location, location, location

Where should you store your data?

You're backing up data to survive many problems, including hard drive corruption or failure, so storing backup files locally is not a foolproof plan. Storing remotely is the best solution, and there are many ways to do this.



External hard drive

A great and simple way to store your backups that you can easily move to any location. External hard drives come in a range of sizes and prices, without any need for subscription costs or maintenance. Some people even store them in safe deposit boxes, keeping them secure from burglary or property damage. If you are keeping it off site, though, there's no way to automatically schedule a backup like the other methods.



Networked or external server

A tried and tested method, favoured by businesses. Servers have the upside of being optimised for storage space rather than speed, and as they're usually always on and connected via a network connection, scheduled backups can be done at times when you're least likely to be using a PC. Servers do require maintenance, though, and a house fire may kill both your PC and server. Off-site servers can be very expensive.



Cloud storage

Rising in popularity. With an abundance of available repositories at reasonable prices, a lot of services will automatically sync folders straight to your account after any changes are made. Being synced from local folders, it's easier to set up a schedule, and the off-site nature means that it's secure from incidents that would affect local hardware. The only limiting factor is transfer speed, upload speeds being significantly slower than network or USB.



Back up & restore your home folder

Get your important files secured in no time at all with easy-to-use graphical tools

Different users are of course going to have different backup requirements, and a lot of people may only need to save specific files and documents. Whether it's because you switch distribution on a weekly basis or just use a laptop for work, backing up your home folder may be enough for you.

Some distros may include a file backup program of their own, like Ubuntu which has backup tool Déjà Dup installed as standard. These graphical tools are great, usually acting as a front-end to the rsync tool we used over the page. Some

of them, like Déjà Dup, also add features to schedule a time to perform backups. Scheduling done right will update only the files that have been changed since the last operation.

In this part of the guide, we are going to talk about another great piece of open source software called luckyBackup that is available in most repositories. This is a popular graphical front-end for rsync, which is well worth having a look at if you want to get a quick and solid back up of your computer.

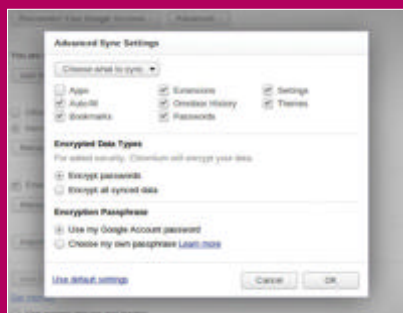
“Backing up your home folder may well be enough for you”

Back up your browsing habits

Save your cookies, bookmarks and more with ease

Your whole life may be on the web. With native tools in the most popular browsers, you can not only keep these

precious customisations safe, you can even use them across multiple systems and distros.



Google Chromium

Google's Chromium and Chrome will hook straight into your Google account, allowing you to sync bookmarks, settings, cookies, form data, apps etc straight to the cloud. You can even selectively choose which data syncs to and from specific systems, and it does work between Chromium and Chrome as well.

Mozilla Firefox

The open source favourite from Mozilla has its own Firefox Sync, syncing a lot of the same data as Chromium does – passwords, bookmarks, add-ons, tabs etc. This is linked to a specific Firefox account you need to set up for it; however, you will be able to browse tabs that were open on another connected system.

Manually

A lot of browsers will let you export data anyway, and the best ones let you select what to export. This has to be done manually every time, though. However, as this data will be kept in specific folders on your system, you could even back those up with everything else.


Back up your files the easy way

Use luckyBackup to quickly & simply back up your computer



01 Set up luckyBackup

Find luckyBackup in most repositories, or compile from the source on the disc. Once you've done that, just run luckyBackup and it will set up a default profile for you.




02 Create a task

Below Task on the right, click Add. You can select the source of what you wish to back up, and the destination for the backed up files. Selecting a folder will copy the entire contents over into a folder of the same name at the destination.



03 Back up

Check the task and click Run – it'll show a progress window with detailed info. The first time will take a while, but successive backups will skip files if they haven't been modified.

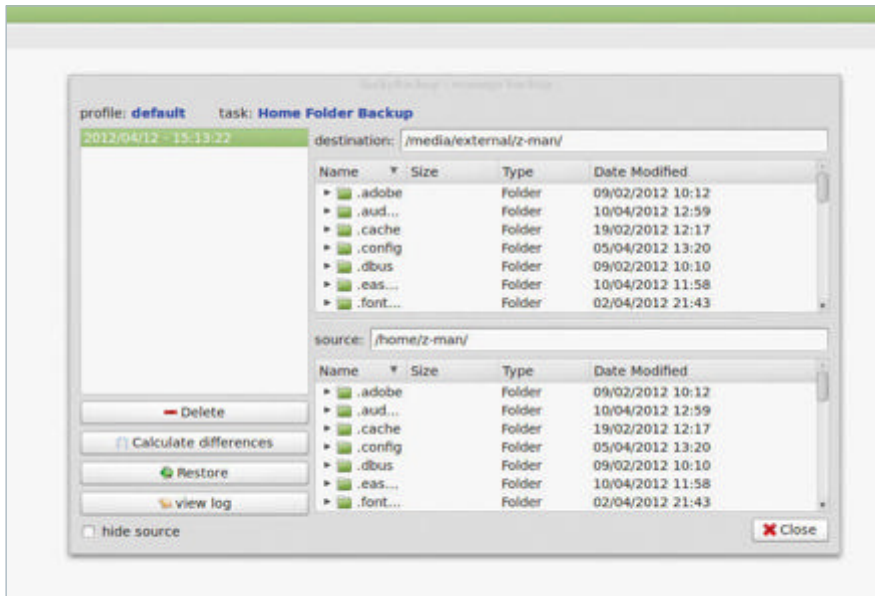


04 Sync files from the backup

If you're backing up to cloud storage like Dropbox, you may have edited the files from another PC. You can add a 'sync' task to make sure your files are always up to date at both ends.


Restore your files after disaster

Recover from disaster with luckyBackup's restore feature





01 Prepare to restore

Firstly make sure you that have luckyBackup installed on your system. You will also need to make sure that the previous destination for the backup is connected and available to be read.




02 Import your profile

If you're on a newly installed system, you'll need to import the profile so that the restoration process knows what to do. Go to Profile>Import and find it in the destination directory in the folder .luckybackup-snapshots.



03 Manage your backups

Click the task you used to create the backup, then Task> ManageTask. This shows all the snapshots, the complete backups from specific dates, which you can selectively restore from in case you don't want the very latest.



04 Restore

Once you've selected a snapshot, click Restore to get things started. If you encounter any errors while restoring, you may have to run luckyBackup as a superuser, so open a terminal and type 'sudo luckybackup'.

05 Reinstate backup

If you've recovered from a disaster, make sure your backup procedures are still in place – recovering the profile should do this, but it doesn't hurt to run a simulation just in case. Check the dry option before running.

Back up to the cloud

Save to the cloud and avoid rainy computer days

At its most basic level, cloud storage is very good for saving documents and other small files. Typical free storage of 2-5GB isn't great for backing up extensive music collections, but perfect for work documents. Here are some recommended services:



Dropbox

Dropbox offers a free 2GB (2.5GB if you complete the tutorials), and the standout thing about this service is how it integrates with file browsers such as Nautilus, 'syncing' a specified folder to and from your account with newer versions of all files.



Ubuntu One

Canonical's cross-platform service is separate from Ubuntu, and offers 5GB of free space plus music streaming. While it works in a similar way to Dropbox, it can be a pain to set up on non-Debian distros.



ownCloud

A great open source cloud storage solution with 5GB for free, ownCloud also includes some basic web apps for viewing photos, managing contacts, and writing text. It's mainly web based for now, but synchronisation tools and smartphone apps are on their way.



Back up & restore a disk image

Ghost your entire hard drive so you can completely restore your system

While a basic backup of files and settings will suffice for a lot of users, it won't be enough for everyone. Your system might be more than something you just work on, from a highly customised operating system that caters to your specific needs, to a carefully partitioned hard drive with multiple distributions. This is where a hard drive image comes in.

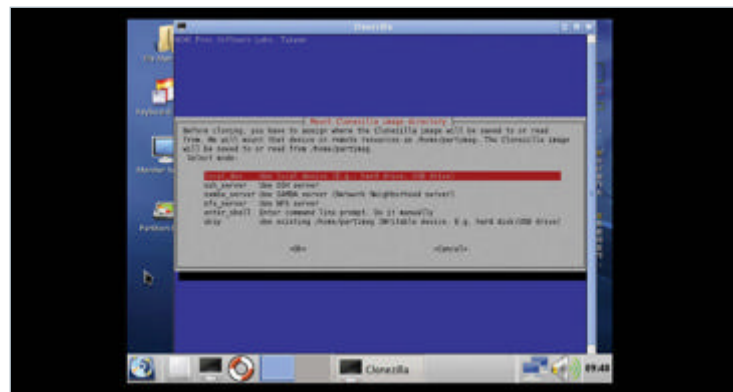
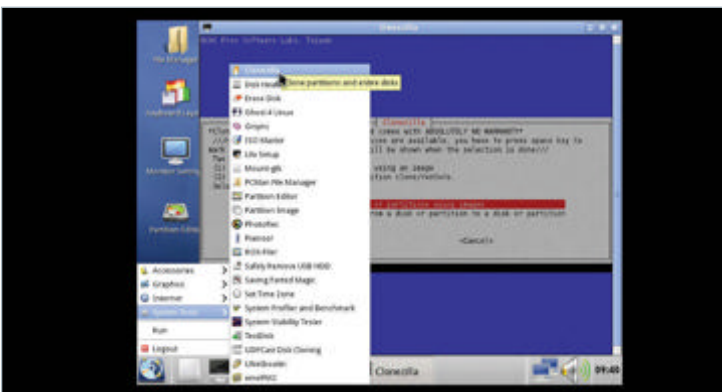
Taking an image of your hard drive is a much more complete way to back up your computer. It's not something you'll be able to perform every day, though – while the previous method will update any new or updated files to your previous backup, disk imaging will have to create an entire clone of your disk every time. It's still something we recommend doing, whether it's once a week

or once a month, and there are some specialised tools to help you do it.

These tools come in the form of live distros, such as Parted Magic. You'll find the latter on the disc at the back of the book and it's a great, all-purpose, system management tool. It also contains the ghosting software we desire for this guide, Clonezilla.

Create a disk image the easy way

Use the free disc and Clonezilla to completely back up your system


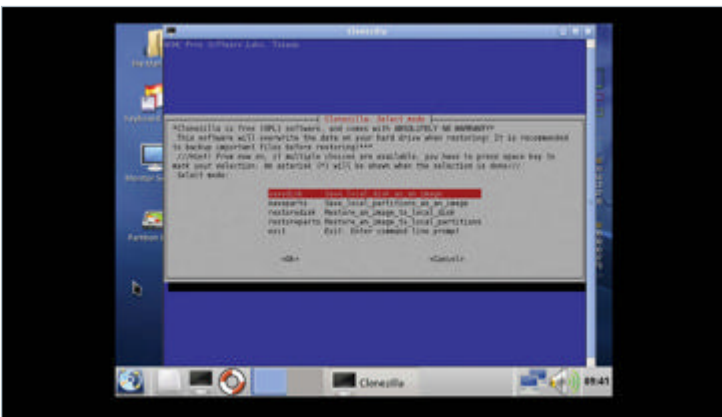


01 Boot into Parted Magic and Clonezilla

Take the free disc, or your own burnt disc if you wish, and reboot your PC. Boot from the disc and choose default settings for Parted Magic. Click on the app menu at the bottom right, and then select Clonezilla from System Tools.

02 Choose a destination

We want to select device-image to create an image. Then choose where you want the image to go – it gives choices of local devices or network shares, and then directories in the destination to choose from. You'll have to set up a directory beforehand if you want a new one.



03 Choose an operation

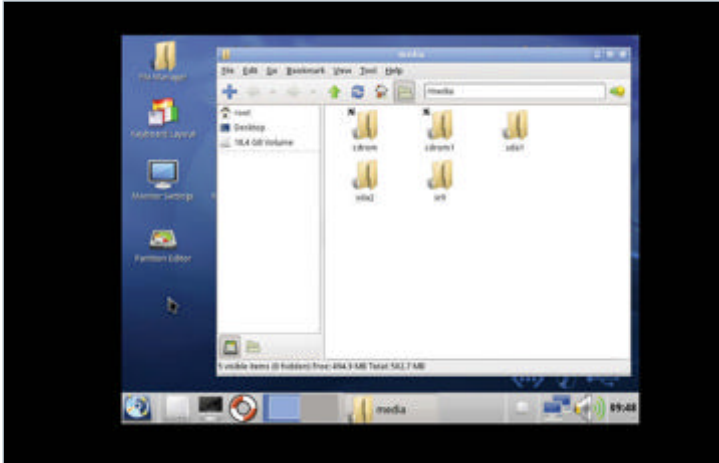
Select beginner mode: it has all the tools we'll need. You can either choose savedisk to create an entire image of the hard drive with all of its partitions, or saveparts for specific partitions. You'll be asked what to name the image, and the source you want to back up.

04 Create the image

Clonezilla will now create an image of your system's entire hard disk – naturally, this can take some time. Afterwards, it will make sure the image can be used for restoration purposes so you have no doubt whether it worked or not.

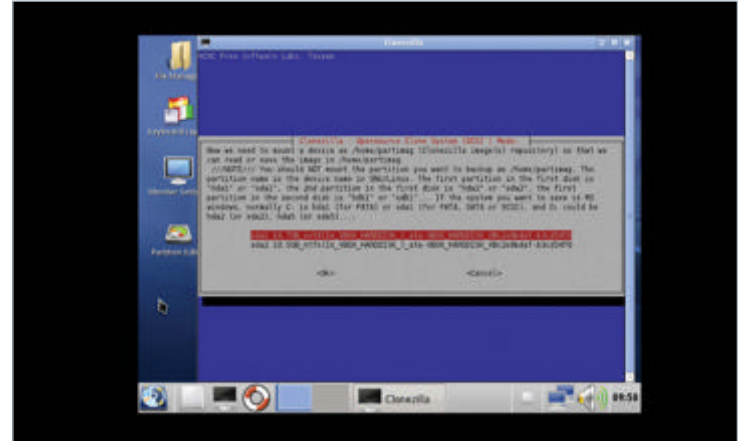
Restoring your disk image after disaster

Bring back your machine byte for byte in easy steps



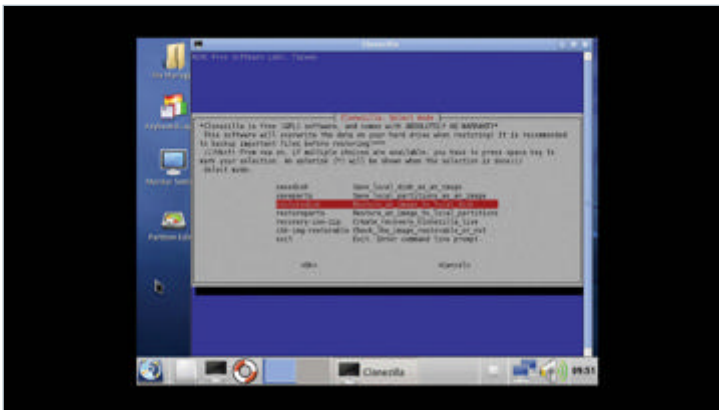
01 Prepare your system

If you're recovering from disaster, such as a dead hard drive, make sure everything has been fixed or replaced. Also be aware that recovering the disk image will completely wipe the hard drive, so ensure you use the right one!




02 Parted Magic and Clonezilla

Follow the steps from the backup section to get to Clonezilla. We need to start setting up Clonezilla as if we're backing up the hard drive we wish to restore to, with the repository and the directory the same as previously.



03 Choose an operation

From the beginner menu again, you can choose to restore disk or restore parts. It will then scan the available images on the destination for you to choose. By default, all the images will be dated, to make the process easier. Then choose the disk or partition to restore to.



04 Restore the image

Clonezilla will ask you twice if it's fine for the image to be restored. After you've confirmed that, your computer will be restored to its former glory. If you're now on a bigger hard drive, make sure to create a new image to make future backups easier.

Create an ISO of your OS

Make a live disc of your system for easy installation on new and fixed machines

One somewhat alternative backup technique is to create an ISO image of your customised distro, allowing you to install it on other systems without the usual setup process. While this has applications in terms of setting up office or special-use PCs, it's also useful as a snapshot of your system if it dies. There are a couple of major programs that help you do this.

Remastersys

For Debian/Ubuntu-based distributions, Remastersys enables you to create an ISO completely from your current install, and also select if you want to save or ignore your documents and files. Remastersys is available from the project's website at www.remastersys.com.

Revisor

If you use a Fedora-based distro, the Revisor program will guide you through creating an ISO based on your system. You can even add extra repositories and packages to the ISO if you also plan to use the image for other applications. You can obtain Revisor from the Fedora repositories, or <http://revisor.fedoraunity.org>.

Perfect Dual Boot

The only guide you'll ever need to manage your system's multiple personalities

Advisor

Rob Zwetsloot A Linux user who is also an avid PC gamer, Rob is well versed in the methods of dual booting after years of practice on many machines



The ability to dual-boot is a very powerful feature in computing – it allows developers to create cross-platform applications, it can help separate work and play, and it also lets adventurous enthusiasts test out new operating systems without losing what they know and love. It can be a scary prospect, though, for some: without proper knowledge it's all too easy to erase or corrupt important documents and system files.

Usually, modern distros will allow you to install alongside other operating systems, and even automatically add the option to boot to the other OS if it uses GRUB 2. However,

this is quick and dirty and has a habit of not always working.

Fear not though, as we have put together a comprehensive guide for everything you'll need to successfully dual-boot. Our step-by-step guides will take you from the simple methods of setting up a hard drive and changing the order of a boot menu, to the more advanced features of adding custom menu entries and tweaking the aesthetics to suit your own personal preferences. We'll even throw in some tips and tricks for Windows and OS X.

BACK UP NOW!

A friendly reminder before we begin



It should almost go without saying, but now would be a great time to back up your system. Of course, we should all have a regularly updated backup just in case anything happens, but let's just cover some sure-fire ways now. A lot of OSs will have their own backup system in place, but there is always luckyBackup. We recommend keeping the backup on an external hard drive or system, although it's highly unlikely that you will break the internal hard drive that badly. You could even create an image of the hard drive using Parted Magic, which is included on the disc at the back of the book. This allows you to restore your hard drive to the exact way it was before adding an OS.

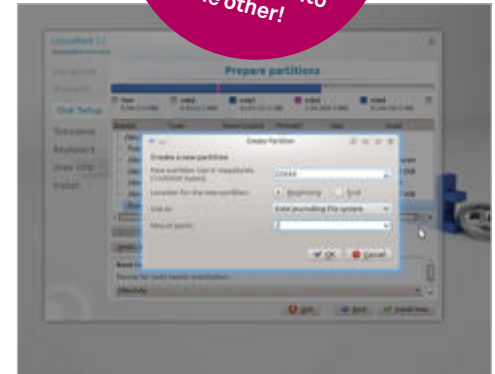
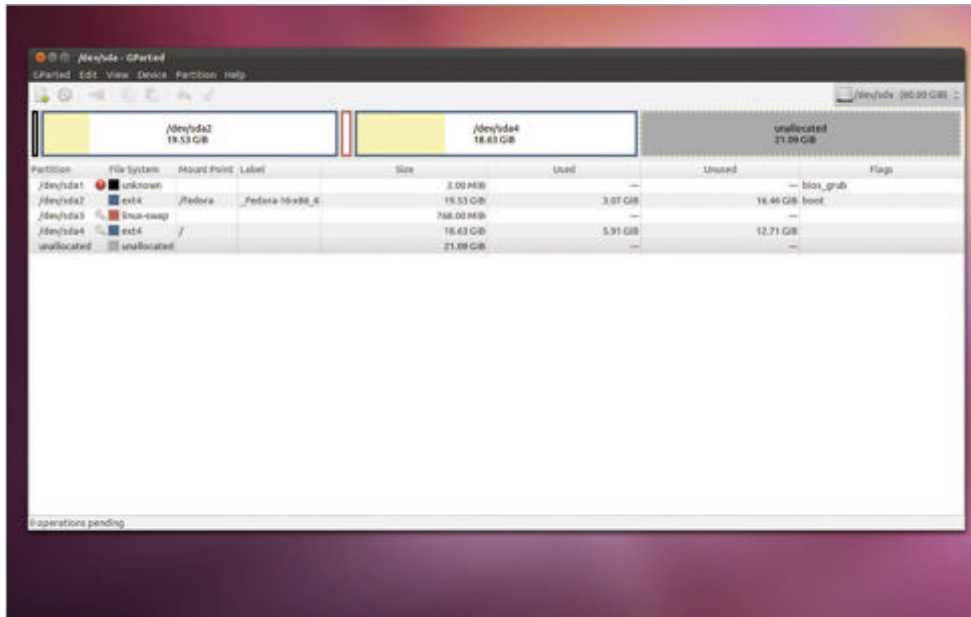




Dual-boot the easy way

So you're thinking of adding another OS to your system? Then grab your image and you'll be dual-booting in no time

You only need one swap partition as it can be used by any Linux system. Just don't suspend one and boot into the other!



01 Prepare your hard drive
Make sure that there's at least 20GB of free space on your hard drive. Use GParted to shrink the root partition so this 20GB

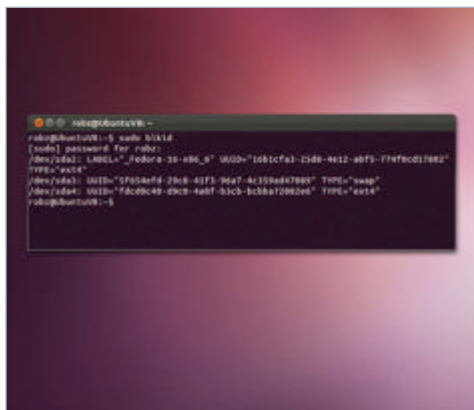
becomes unallocated space, by right-clicking on the partition and selecting 'Move/Resize'. Once this is done, click 'Edit' then 'Apply All Operations'.

02 Install the second OS
Put in your installable medium and reboot the PC. Ubuntu, Fedora and their derivatives will give you options to install 'alongside' the other OS. This will use the available free space you just made. Otherwise, when asked, create an ext4 partition out of the free space and set that as root (/).


If you're installing Ubuntu, Fedora 16, or anything based on them, then you're already done! They come with GRUB 2, which will automatically add your other OS to the boot menu. However, if the OS you just installed uses GRUB, you still have a few steps to go...

Getting GRUB Legacy to dual-boot

If you were using GRUB 2, you'd be done by now!



01 Prepare to add an OS
Boot into Linux and open the terminal. Type 'sudo blkid' to return the UUID of the partitions on your hard drive, and make a note of the UUID for the partition with the OS you need to add.



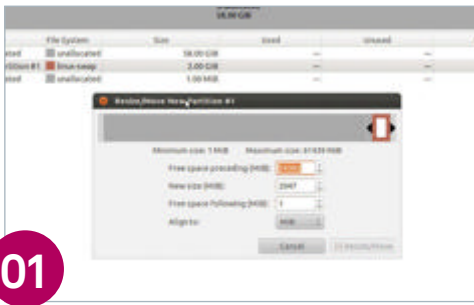
02 The kernel image
Mount the root partition of the other OS, and use ls in the terminal to list the contents of that partition's /boot folder. Make a note of the name of the kernel image, which starts with 'vmlinuz'. Also note the accompanying init file in this folder, which ends '.img-' followed by the same string of numbers as the kernel image.



03 Adding the menu entry
In the terminal, type 'sudo nano /boot/grub/menu.lst' for Debian/Ubuntu, or 'sudo nano /boot/grub/grub.conf' for Red Hat/Fedora. Use the arrow keys to scroll down to the entries at the end of the file. You need to enter the new item as follows, with your UUID being X, kernel as vmlinuz-Y, and the init as initrd.img-Y:
title Another OS
uuid X
kernel /boot/vmlinuz-Y root=UUID=X ro quiet splash
initrd /boot/initrd.img-Y
quiet
Save, then 'sudo update-grub'.

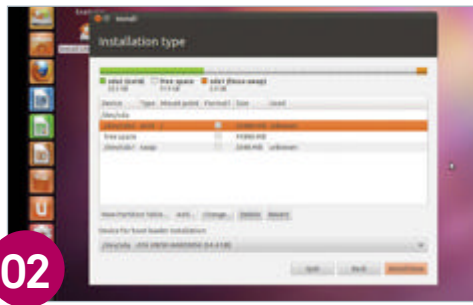
Partitioning your hard drive

We've shown you in our quick guide how to shrink an existing partition, but here's how to set up your hard drive from scratch with the necessary partitions



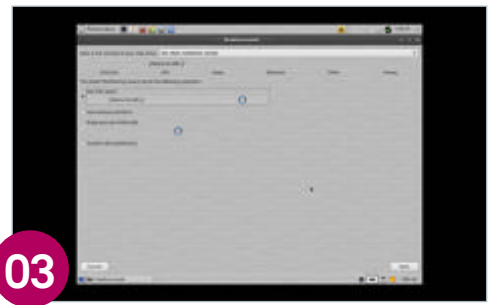
Setting up the swap partition

In GParted, delete any existing partitions so all that remains is unallocated space. Go to Partition>New. Select linux-swap in the file system drop-down menu, and make it the same size as the RAM your system has. Set the MB following value to 0 to put it at the end of the hard drive.



Installing the first OS

During install, select the option to set up custom partitions if it doesn't do it automatically. Create a boot partition (such as a bios_grub file system) with at least 2MB, as well as a root partition (/) in the ext4 file system. This should be at least 20GB in size, but also leave at least 20GB of free space.



Installing the second OS

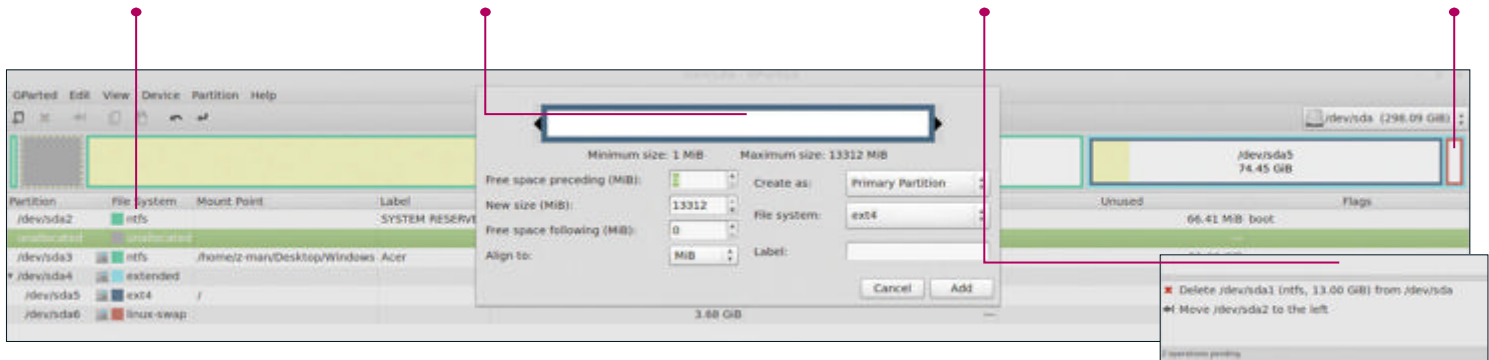
Put in your other installable medium and reboot the PC. Ubuntu, Fedora and their derivatives will give you options to install 'alongside' the other OS. This will use the available free space you just made. Otherwise, when asked, create an ext4 partition out of the free space and set that as root (/).

List of partitions: GParted organises partitions by name, with information on size, file system etc per entry

Adding a new partition is dead easy, giving you precise control over the size and location as well as the file system

Changes aren't applied until you say so, allowing you to play around with sizing before committing

The linux-swap should preferably be at the end; however, it will work anywhere else




Restoring your PC

The worst has happened. During the partitioning process, something went horribly wrong. Now, somehow nothing is working, nothing is booting, and you don't have time to troubleshoot. This is where the backups or disk image we told you to make will save the day...



Restoring your backup

Boot into a live disc and connect the external storage with your backup. For restoring the backup, we'd suggest luckyBackup again. The quick and easy wizard really pays off.



Restoring an image

Parted Magic is probably the best for doing this, so boot into it with a live disc. Use the disk cloning tool Partimage to restore the image you made previously.


Master GRUB

The boot menu is the key to managing your dual-boot system – and if you’ve just installed a second OS, there might be some changes you want to make...

01 Getting started with GRUB
Options for GRUB Legacy can be found at '/boot/grub/menu.lst' in Debian/Ubuntu, and '/boot/grub/grub.conf' in Red Hat/Fedora. GRUB 2 is in '/etc/default/grub'. Access them by typing 'sudo nano /boot/grub/menu.lst', 'sudo nano /boot/grub/grub.conf', or 'sudo nano /etc/default/grub/menu.lst'.



02 Changing the default selection in GRUB 2
In the terminal type 'grep menuentry /boot/grub/grub.cfg' to show the order of the boot menu entries. Note the number of the one you'd like to be default, then change the number following GRUB_DEFAULT in '/etc/default/grub'. Save and 'sudo update-grub'.



03 Or, if you're changing the default selection in GRUB Legacy:
Go to the end of menu.lst/grub.conf, where it shows all the boot time selections. With the first entry being 0, note the number of the entry you wish to be default, then change the number following default at the beginning of the file. Save and 'sudo update-grub'.

04 Altering the timeout duration in GRUB
The number of seconds until the default option is selected can be altered in the settings file. In menu.lst/grub.conf this option is called timeout, while in '/etc/default/grub' this is shown as GRUB_TIMEOUT. Change the number to your liking, save and then 'sudo update-grub'. Alternatively, you can rearrange the entire boot menu...

USING GRUB LEGACY

Rearranging the boot order in GRUB Legacy is easy

Go to 'menu.lst/grub.conf' and use the down arrow to navigate to the end of the file to find the boot options. Rearrange these into the order you want, including the details about the root and kernel etc in the move. Save and then 'sudo update-grub'.



GRUB 2 With GRUB 2, rearranging your boot menu is as easy as 1-2-3

01 Understanding the boot order in GRUB 2
Type 'ls /etc/grub.d/' into the terminal to see the individual files that govern the order of the boot menu. By default, these are labelled 10_linux, 20_memtest86+ and 30_os-prober. 10_linux holds the info for the first entry, while 30_os-prober finds the other operating systems on your machine.

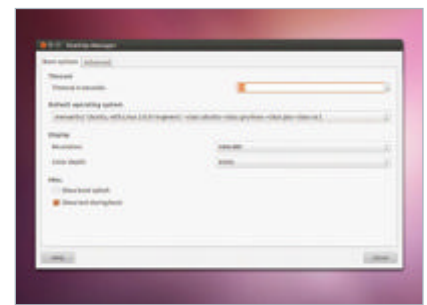
02 Rearrange the boot order in GRUB 2
The order is based on the number preceding the filename. So, to have the other OS appear before the first entry, you can change the 30 to a 09. To do this, type 'sudo mv 30_os-prober 09_os-prober' into the terminal, followed by 'sudo update-grub'.

“The boot order is based on the number preceding the filename”

DOING IT THE GUI WAY

Some alternate UI tools to try

There are a few UI tools you can use to edit the boot menu, like Startup-Manager that's bundled with Ubuntu. These give you a similar amount of control as you get by manually editing the boot menus; however, it's not available for every platform and development was halted last year. There are also a few graphical editors for Windows: one also named Startup Manager and another called easyBCD can be useful if you want to use the Windows bootloader; however, they only add the ability to chainload into the Linux bootloader anyway.



■ You get some limited control with Startup-Manager

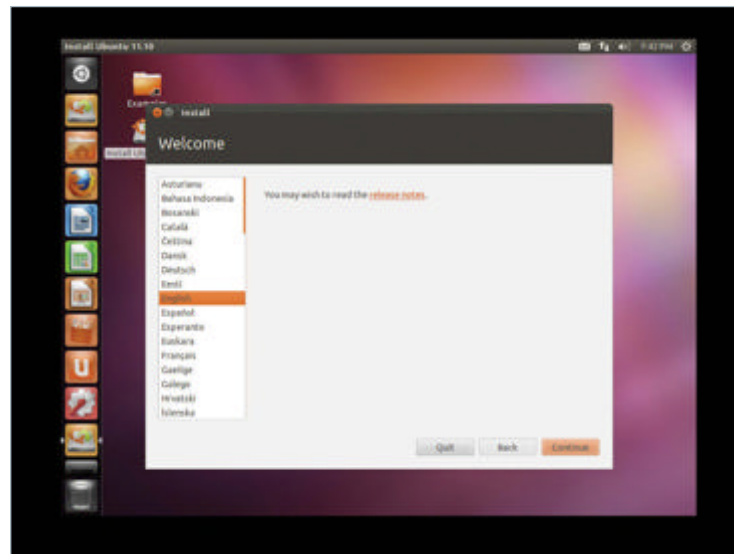
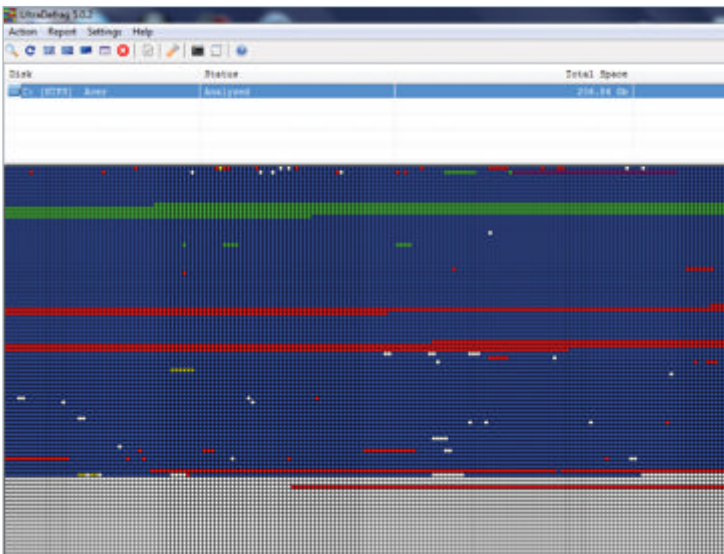
Dual-booting Windows and Linux

With Windows in the mix, setting up hard drive partitions is trickier – especially if Windows is already installed...

The main difference in partitioning your drive with Windows comes with the file systems – NTFS has bad fragmentation, which could cause problems

with shrinking the drive if files have been pushed to the end of the partition. Also, Windows can't read ext4 partitions that Linux uses by default. Thus, we

recommend having a larger NTFS partition if you plan to store media, as it can then be used across both operating systems.



01 Getting started

Use a live disc, such as the one in the back of this mag, containing Parted Magic, to shrink the NTFS partition. As well as the minimum of at least 20GB of free

space to install Linux, create some extra free space to accommodate the required swap partition. This only needs to be as big as the available RAM on your system.

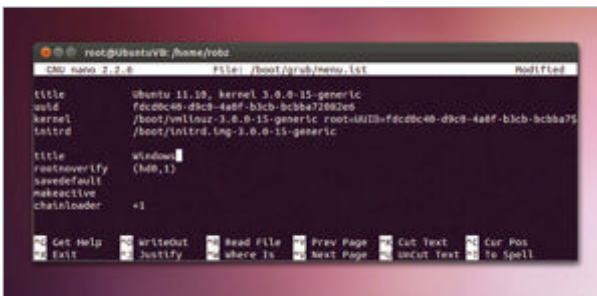
02 Installing your distro

Install your preferred Linux distribution to the remaining free space, creating the swap partition to be the same size as your system memory, and fill the rest of the space with an ext4 partition and

set it as root (/). On the latest Ubuntu or Fedora builds, GRUB 2 should configure your other OS automatically. If not, type 'sudo update-grub' in the terminal. Refer to our earlier step-by-step to change the boot order.

Adding Windows to GRUB Legacy

Windows and Linux working in perfect harmony



First, a little preparation...

Use GParted to find out the partition of the Windows install. For our example, we will use sda2, which is the second partition on sda. As GRUB counts from 0, this will become (hd0,1), where hd0 is the hard drive, and 1 is the second partition.

Next, add your menu entry

As before, use nano to go to 'menu.lst/grub.conf', and then use the arrow keys to scroll down to the menu entries at the end of the file. As in the last step, we will use the example (hd0,1):

```
title Windows
rootnoverify (hd0,1)
savedefault
makeactive
chainloader +1
```

Save, then 'sudo update-grub'.

“The main difference in partitioning your drive with Windows comes with the file systems – NTFS has bad fragmentation”

What about adding Windows to a Linux box?

The process is a little different, but it's made much easier with the following steps...

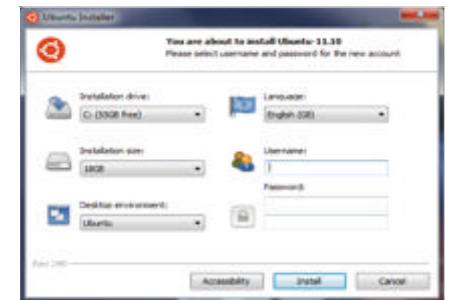
01 Create some space on your drive
 As mentioned previously, we recommend giving Windows as much hard drive space as you can spare if you plan to use files between the two operating systems. Follow the steps to shrink a partition as before, but this time create an NTFS partition out of the remaining space.

02 Kick-start the Windows installation
 Windows should detect the NTFS partition you created and ask to format it. Let it, and finish the rest of the installation. You will notice there is no way to boot into Linux. Either you've just installed Windows, or something has gone wrong! Grab a live disc and follow either Legacy or GRUB 2 guides below to get it sorted...

WUBI

Install Ubuntu from Windows

WUBI is very simple way to install Ubuntu alongside Windows, and officially supported by Canonical. The download can be found on the Ubuntu website under Windows Installer, and it installs to your Windows Partition and sorts out the dual booting automatically. It comes at a price, though – you can't create the same custom partitions as you could otherwise, and fragmentation of the NTFS partition can affect performance of Ubuntu. Still, it's an easy way for interested Windows users to test out Ubuntu and Linux if they're interested, without the worry of losing files. Check out www.ubuntu.com for more details.



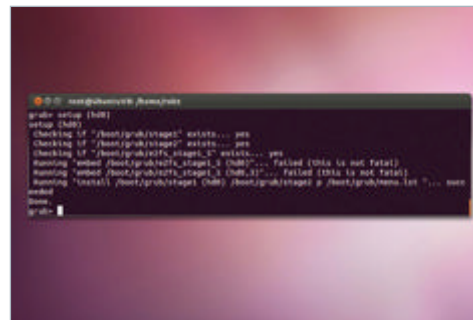
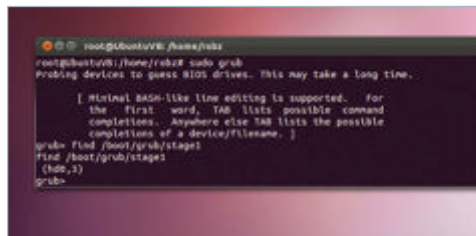
■ WUBI allows you to install Ubuntu with minimal preparation

How to recover GRUB Legacy

Save your boot menu from disaster...

01 Boot into a live CD distro
 Put in your preferred live disc, or use the one in the back of this book, and boot into a Live Linux Environment. Go to the terminal and type 'sudo grub' to enter the GRUB command line. Type 'find /boot/grub/stage1' to list all bootable images, designated with a partition number like '(hd0,1)'. Note yours down – we will use hd0,1 as an example.

02 Then, reinstall GRUB
 From the same command line, type 'root (hd0,1)', then 'setup (hd0)'. This will reinstall GRUB to the MBR. Type 'quit' to exit the GRUB command line, and then reboot your machine.



What about GRUB 2?

Simple – just reinstall it...

Boot into a Live Linux Environment (you can use the disc at the back of the book for this), and open the terminal. Use GParted to find the name of the partition Linux is installed in – for this example, 'sda'. Type 'sudo grub-install /dev/sda', then reboot your system.

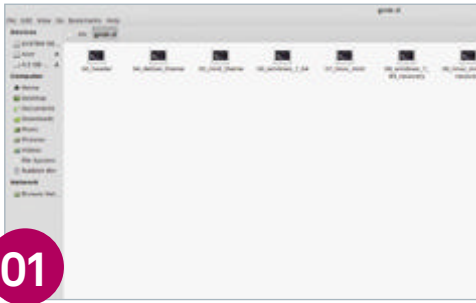
“We recommend having a larger NTFS partition if you plan to store media”

■ The finished product doesn't have to look quite so dull. Turn the page to get tweaking...




Create amazing GRUB 2 menus

There's no denying that GRUB 2 is a vast improvement over Legacy, but it's still not pretty. Here's how to take full control of your GRUB 2 boot menu...



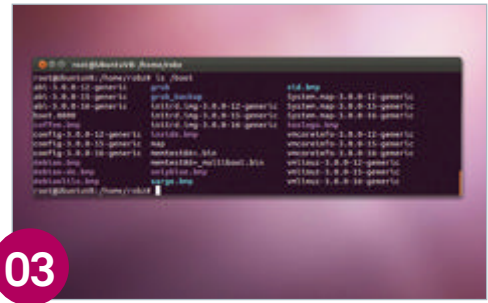
01 Create and modify your scripts

It's a good idea to create your scripts in the order you want them, with a number preceding the label, eg 12_fedora16. The info for memtest and your default distro will already exist, so you just need to change the numbers attached to the filename. You can even rename them, as long as the scripts are numbered in your desired order.



02 Preparing to create your menu item

Similar to GRUB Legacy, you will need to know the root partition, as well as the kernel image and initrd if booting into another Linux. Using GParted, you can figure out the root partition on the different OSs, ie sda2. GRUB 2 counts from 1, so in this example the root partition would be (hd0,2).



03 Kernel image for Linux

Mount the root partition of the other OS, and use ls in the terminal to list the contents of that partition's /boot folder. Make a note of the name of the kernel image, which starts with 'vmlinuz'. Also note the accompanying init file in this folder, which ends '.img-' followed by the same string of numbers as the kernel image.

DUAL BOOT WITH MAC

The future of dual booting with the forbidden fruit...

Intel-based Macs use EFI, the proposed replacement for the BIOS, and this causes problems with current dual-boot techniques for Linux (and Windows). To dual-boot with Linux you need to first use the OS X Disk Utility tool to shrink the current partition by creating a new one out of the free space in any file format. Again, we recommend at least 20GB plus the amount of RAM your system has for the swap partition. Then install rEFIt on OS X and boot into a live disc. Only GRUB 2 will work for this, so make sure you have the latest Fedora, Ubuntu or a derivative. Use GParted to erase the partition you just created in OS X (it should be at the end). Install your OS by following our step-by-step guide in this article. Once completed, reboot and select the partition tool in rEFIt to sync the partition tables. Completely shut down the Mac afterwards, and when you reboot you will have the option to boot to Linux and OS X.



04 Adding Linux menu entry

The file must contain the following code (X is partition number of hard drive, and Y is kernel image number):

```
#!/bin/sh -e
echo "Adding Another OS"
cat << EOF
menuentry "Another OS" {
set root=(hd0,X)
linux /boot/vmlinuz-Y ro quiet splash
initrd /boot/initrd.img-Y
}
EOF
```

05 Adding Windows menu entry

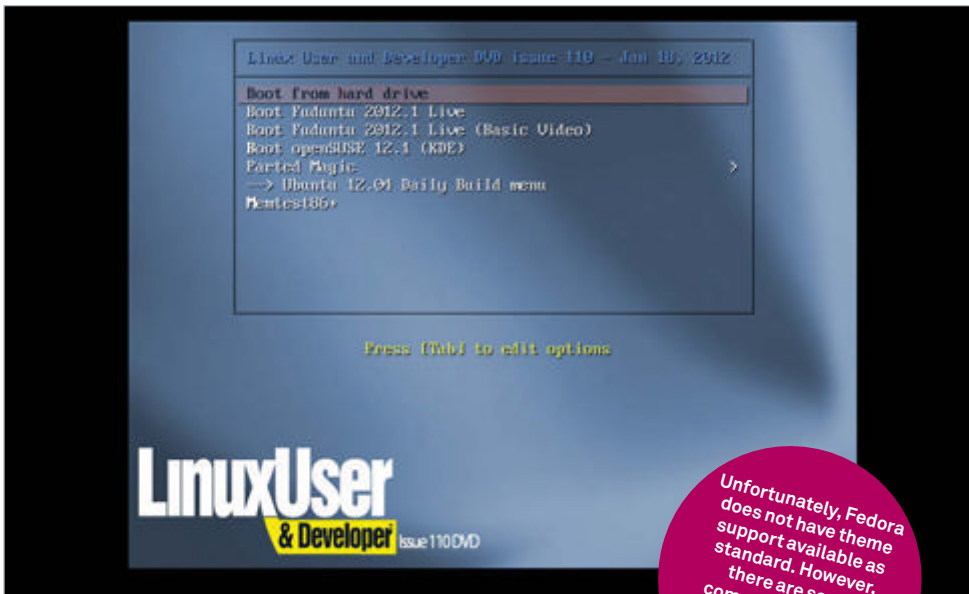
Use the following code for the Windows menu entry (where X is the partition number of your hard drive):

```
#!/bin/sh -e
echo "Adding Windows"
cat << EOF
menuentry "Windows" {
set root=(hd0,X)
chainloader (hd0,X)+1
}
EOF
```

06 Finalise the boot menu

There are two final steps to perform to update the boot menu – first you need to make the scripts you've just written executable by typing 'sudo chmod +x' followed by the name of the file in the terminal. After you have done that with all your new scripts, type 'sudo update-grub'.

“It's a good idea to create your scripts in the order you want them with a number preceding the label”



■ With enough tweaking, you can achieve this and much more

Unfortunately, Fedora does not have theme support available as standard. However, there are some community themes available for download, as with any other OS with GRUB 2

Customise GRUB


If you're finding the standard monochrome colour scheme for the boot menu a bit dreary, you can liven it up...

Prettify GRUB Legacy

Give GRUB a makeover in two easy steps

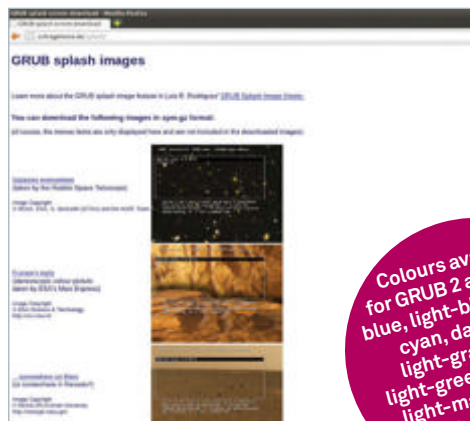
Tweaking the colour scheme

Use 'sudo nano /boot/grub/menu.lst' (or 'sudo nano /boot/grub/grub.conf'). Before the menu entries, start a line with 'color', then specify the normal colours, then the highlight colours. For example, 'color black/blue black/green' will have black text on a blue background, then when highlighted it will be black text on a green background for contrast.



Add an eye-catching background

Use 'sudo mkdir /boot/grub/images' – where the splash image will live. GRUB Legacy background images must be in an '.xpm.gz' format to work, but there are plenty of repositories online for splash images. Copy the image into the new folder you just created. In 'menu.lst/grub.conf', add the line 'splashimage' and then point towards the image file, including the partition it's on; eg 'splashimage (hd0,1)/boot/grub/images/splash.xpm.gz'.



Colours available for GRUB 2 are black, blue, light-blue, brown, cyan, dark-gray, light-gray, green, light-green, magenta, light-magenta, red, light-red, white, and yellow

Tweaking GRUB 2

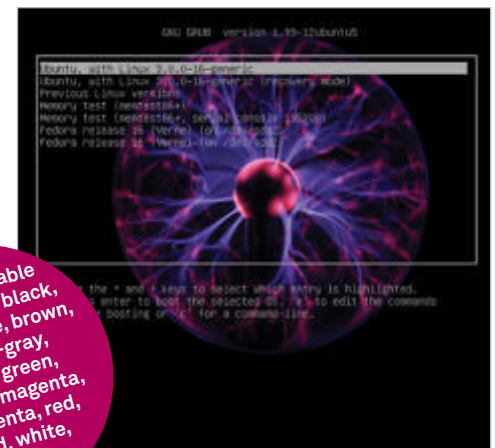
GRUB 2 is much more pliable...

01 Change your colour scheme
Use 'ls /etc/grub.d' to list the files that modify the GRUB 2 menu. Before the menu entries will be a theme file, eg '05_debian-theme'. 'sudo nano 05_debian-theme' will open it up for editing, and navigating down the file you will find 'menu_color_normal' and 'menu_color_highlight'. The first colour listed is the foreground, followed by the background colour. For example, 'black/blue and black/green' will have black text on a blue background, then when highlighted it will be black text on a green background.



02 Beautify your backdrop
Images must be an 8-bit, RGB image with the extensions .tga, .png or .jpg. Simply place the image you desire in '/boot/grub' and when update-grub is run, it will grab the image. If there are multiple images in the folder, it will choose them in the order of jpg/png/tga, and then by alphanumeric order if there is more than one with that file extension.

03 And don't forget!
Always remember to save and 'sudo update-grub' after any changes are made.



Keep your system safe

Swayam Prakasha explains various types of security holes that could exist in your system and how to deal with them



Computer security is like any other form of security.

It plays a key role in today's internet age, especially in organisations like financial institutions. The field of information security has grown and evolved significantly in recent years. For any organisation, security has become the key component to focus on. The main focus of any system administrator is to make sure that your network has fewer security holes. Most of the external unauthorised people (aka hackers) normally operate by using the well-known security tools and techniques to find out how they can get into your network/systems. Identification of security holes in your system is an important aspect for any organisation. This activity needs to be carried out on a regular basis and once identified, appropriate measures need to be taken so that such security loopholes can be fixed.

Of the major internet outages caused by computer crimes, many of them have come from the exploitation of security holes that been known for some time before the incident. You need to make sure that all your systems are up-to-date as and when new patches are available. This will reduce the probability of attacks by external hackers. There are plenty of reasons why we have systems with vulnerabilities. **Some of these reasons include:**

- Firms will often outsource their IT function to other external companies. The latter will focus mainly on keeping the network up and running, therefore security will be compromised as they need to keep the users happy.
- There will be too many patches to apply. And while fixing one thing, it is not surprising if another thing gets broken. This will pose a severe problem in the case of mission-critical systems.
- All system administrators may not have complete awareness of various vulnerabilities. Therefore it is up to the organisation (or up to the individuals) to make sure that all the concerned individuals are properly trained.

It is important to understand here that there are usually several different ways to get into a system. This makes the job of external hackers much easier. When we have a network with various components, there might be several potential windows that could allow an unauthorised individual to gain access to a system. And normally, when one type of attack doesn't work, they can always try another.

In short, information security systems are expected to provide confidentiality, data integrity and availability.

In the following sections, we will take a detailed look at some of the potential ways that someone with the right knowledge can cause havoc on your company's system. Once we have identified various vulnerabilities, it is up to us to go ahead and come up with remedial mechanisms. Vulnerability management is the practical way of identifying, classifying, remediating and mitigating vulnerabilities. And as expected, this mechanism generally refers to software vulnerabilities in your network or computing systems.

Web security has gained a lot of focus nowadays and a recent study has revealed that as many as 70 per cent of websites have vulnerabilities that could lead to the theft of sensitive corporate data such as credit card information and customer lists. Hackers will normally concentrate their efforts on web-based applications such as shopping carts, forms, login pages and dynamic content. As we all know, web applications are accessible 24/7 from anywhere in the world and therefore insecure web apps provide easy access to back-end corporate databases. A victim's website can be used to launch criminal activities such as hosting phishing sites or to transfer illicit content, while abusing the website's bandwidth and making its owner liable for these unlawful acts. Therefore we can conclude that website security is possibly today's most overlooked aspect of securing the enterprise and should be a priority in any organisation.

As far as application security is concerned, we need to come up with proactive measures by constantly analysing network threats and identifying the potential attacks or new attack vectors in your infrastructure. We can consider the application security process to consist of four stages (**Fig 1**), as described on page 121.

In the initial stage of gathering information, one will collect various information such as domain names, range of IP addresses and information about various hosts in the network. In the second

Advisor

Swayam Prakasha has a master's degree in computer engineering. He has been working in IT for several years, concentrating on areas such as operating systems, networking, network security, eCommerce, and LDAP and web servers



“Identification of holes in your system is an important aspect for any organisation”

Whenever we talk about security, we all know that we need to focus on three characteristics, as clearly defined by the CIA model:

Confidentiality – That is, we are able to prevent the disclosure of the information to unauthorised individuals.

Integrity – Here, we are able to make sure that data will not be modified during the transit.

Availability – Availability makes sure that information will be available whenever it is needed. High-availability systems are expected to be available all the time. Therefore they need to prevent service disruptions due to power outages, hardware failures and system upgrades. There is one important attack known as denial-of-service (DoS) that will disturb the availability of the systems.



Fig 1 A quick look at the application security process

“Vulnerability management is the practical way of identifying, classifying, remediating and mitigating vulnerabilities in your network or computer systems”

phase, one will carry out penetration tests to see if there are hosts with any weaknesses. Then in the third stage, one can use manual as well as automatic processes to determine if there are any security holes in the system under consideration. Therefore vulnerability scanning is a part of this stage. Attack vectors specifically refer to improper configurations in the servers as well as weaknesses in the services. Once we are able to identify these vectors, one can exploit the system for any sensitive data. In the final analysis stage, a thorough study of the system vulnerabilities will be carried out. And focus will be given in coming up with strategic recommendations. The recommendations can then be classified and prioritised based on people, technology and process. You can close this application security process by implementing the proposed recommendations.

It is important to understand here that apart from penetration tests for networks, one can also carry out security tests for specific applications, eg for mail platforms. Application penetration testing serves the identification and objective evaluation of potential security deficiencies and risks of publicly accessible applications.

You will usually come across various vulnerabilities. It is very important to detect and correct these vulnerabilities, failing which we will need to compromise on the sensitive data. When it comes to web application security, two of the most important vulnerabilities are:

SQL injection

Cross-site scripting (XSS)

In simple words, SQL injection can be defined as a hacking technique which modifies SQL queries

in order to gain access to data in the database. In other words, SQL injection is a term describing the act of passing SQL code into an application that was not intended by the developer. Some of the common vulnerabilities exploited by SQL injection include:

- Poor input validation in your web applications.
- Over-privileged application logins to the database.
- Weak permissions that fail to restrict the application's login to the database.

We may note here that many of the problems that allow SQL injection are not the fault of the database server but rather are due to poor input validation and coding at other code layers. Therefore SQL injection can usually happen when a developer uses poor input validation in his/her code. In other words, when we have an SQL query that expects a user input, we always assume that the user will not enter something malicious.

In most cases, single quotes will be the main culprit. There is an easy way to check if your application is subject to SQL injection vulnerabilities or not. You need to briefly disable the error handling so that ODBC errors or SQL server errors are displayed. Then, try inputting single quotes into your application to see if you can cause it to fail. A failure is indicative of poor validation and corruption of the SQL string. One way to prevent this type of vulnerability is to always use parameterised queries for all data access in your code.

Another very important vulnerability with web applications is cross-site scripting (XSS). These days, websites are more dynamic in nature. When we say dynamic, what it means is that they are

Vulnerability scanners

A vulnerability can be defined as a weakness that will allow a hacker to gain access to a system, thereby he can reduce a system's information assurance. Vulnerability scanners can be used to fix the discovered security loopholes.

Once the vulnerabilities are identified, it is important to come up with a process to address these vulnerabilities. Vulnerability scanners help us a lot in this aspect. Vulnerability scanners attempt to identify vulnerabilities in the hosts scanned. They also help in identifying out-of-date software versions, applicable patches and so on. Therefore vulnerability scanners will provide information about vulnerable services or ports on hosts. In most cases, one can find patches or updates to cure the problems. It is interesting to note that there are both commercial as well open source vulnerability scanners.

Masterclasses

able to deliver different output to a user depending on their settings and requirements. Though dynamic websites provide more user friendliness, they basically suffer from this major threat of cross-site scripting.

Cross-site scripting typically happens when a web application gathers malicious data from a user. As we know, the data is normally gathered in the form of a hyperlink which contains malicious content within it. Once the data is collected by the web application, it creates an output page for the user containing the malicious data that was originally sent to it, but in a manner to make it appear as valid content from the website. Cross-site scripting can be considered as a hacking technique that leverages vulnerabilities in the code of a web application to allow an attacker to send malicious content from an end-user and collect some type of data from the victim.

It is important to note here that cross-site scripting is a potentially dangerous security exposure that should be considered when designing a secure web-based application. It has been found that cross-site scripting carried out on websites represents roughly 80 per cent of all security

vulnerabilities. And cross-site scripting is gaining a lot of popularity among attackers as an easy exposure to find in websites. There are many risks associated with this cross-site scripting. It is not surprising to see that users can unknowingly execute malicious scripts when viewing dynamically generated pages based on content provided by an attacker. So, the cross-site scripting's success can be attributed to the fact that it requires the victim to execute a malicious URL which may be crafted in such a manner that it appears to be legitimate at first glance.

So how can you prevent cross-site scripting vulnerability? To reduce the chances of your site becoming a victim of a cross-site scripting attack, it's essential that you develop any web application using some form of security development life-cycle (SDL). A critical rule we need to understand here is to assume that whenever we develop secure applications, all data received by the application is from an untrusted source. Not trusting user input means validating it for type, length, format and range whenever data passes through a trust boundary. In other words, we need to review every point on our site where user-supplied data is handled and processed and ensure that, before being passed back to the user, any values accepted from the client side are checked, filtered and encoded.

Let us have a look at another popular vulnerability – web server exploits. It is common these days for companies to run a web server. Web servers can be considered as favourites for having various bugs and security loopholes in them. The large number of holes is due to the ever-expanding number and types of protocols and commands that web servers have to deal with.

An external hacker can access databases and other confidential information by getting through the web server. Organisations should make every attempt to implement appropriate security management practices and controls when maintaining and operating a secure web server. It is important to make sure that the system administrators carry out various activities so as to make them much more secure:

- **Installing and configuring systems in compliance with the organisational security policies and standard system and network configurations.**

“Cross-site scripting carried out on websites represents roughly 80 per cent of all security vulnerabilities”

- **Maintaining systems in a secure manner. This includes taking frequent backups and timely application of patches.**
- **Monitoring system integrity, protection levels and security-related events.**
- **Conducting security tests as required.**

We can conclude that protecting a web server involves hardening the underlying operating system, the web server application and the network to prevent malicious entities from directly attacking the web server. It is possible to avoid many security issues if the operating systems underlying the web servers are configured appropriately. For web servers, the authorised users who can configure the operating systems are limited to a small number of designated web server administrators.

On the other hand, the users who can access the public web server will normally range from unrestricted to restricted subsets of the internet community. The web server administrator should be able to enforce restrictions on the policy by configuring the operating system in such a way that it has to authenticate a prospective user by requiring proof that the user is authorised for such access.

The bulk of the time we see an exploited server it has been exploited in one of two ways – either someone has guessed a password and logged in as a user via SSH, or a web application has a security hole that has been exploited. A basic step to prevent this is to change the SSH port from 22 to something else and for this you need to edit the `/etc/ssh/sshd_config` file. To make the changes effective, restart SSH with the `'/etc/init.d/sshd restart'` command.

Email server exploits

As we all know, email has been one of the most important and widely used internet applications. Therefore it is only natural that an email server can be a source of many security problems.

Let us try to understand a few common attacks against SMTP servers.

a. EXPN command – The SMTP command EXPN provides details about users on a system. A potential hacker can use this information to set up an attack later.

b. Too many recipients – An attacker is trying to send a single email message with a very large number of recipients. In a way, this can be considered as a denial-of-service (DoS) attack and many email systems cannot handle too many recipients, thereby causing the system to crash. In order to prevent this, make sure that you set the maximum number of recipients for a single email. If the number of recipients exceeds this count, one can trigger intrusion detection.

c. A very long email name – With this, an attacker will try to gain control of the email service through a buffer overflow. There needs to be a configuration – maximum length of an email address – for this item properly set.

“Many of the problems that allow SQL injection are not the fault of the database server but rather are due to poor input validation and coding”

Another vulnerability that we need to focus on is database exploits. We have seen that several company websites offer external access into their databases. For instance, you might allow customers to place and check the status of orders online, allow employees to get information on benefits programs via the web etc. It is natural that all of these functions access an internal company database. We can understand that by doing this, we are opening up a potential source of vulnerabilities.

There are many ways to attack a database. External attacks may exploit configuration weaknesses that expose the database server. An insecure web application may also be used to exploit the database. As seen earlier, SQL injection is one of the main threats to a database server. You need to understand and minimise the other threats such as unauthorised server access and password cracking.

“Organisations should make every attempt to implement appropriate security management practices and controls when maintaining and operating a secure web server”

Let us understand an interesting vulnerability – buffer overflow. Buffer overflow can be considered as an exploit that takes advantage of a program that is waiting for an input from a user. Assume that a program is waiting for a user to enter his name. Instead of entering his name, an attacker may enter an executable command that normally exceeds the stack size. In order

for this command to be executed, the attacker needs to specify a return address pointing to the malicious command.

Naturally, the program partially crashes because the stack overflowed. It then tries to recover by going to the return address, but the return address has been changed to point to the command specified by the hacker. You need to

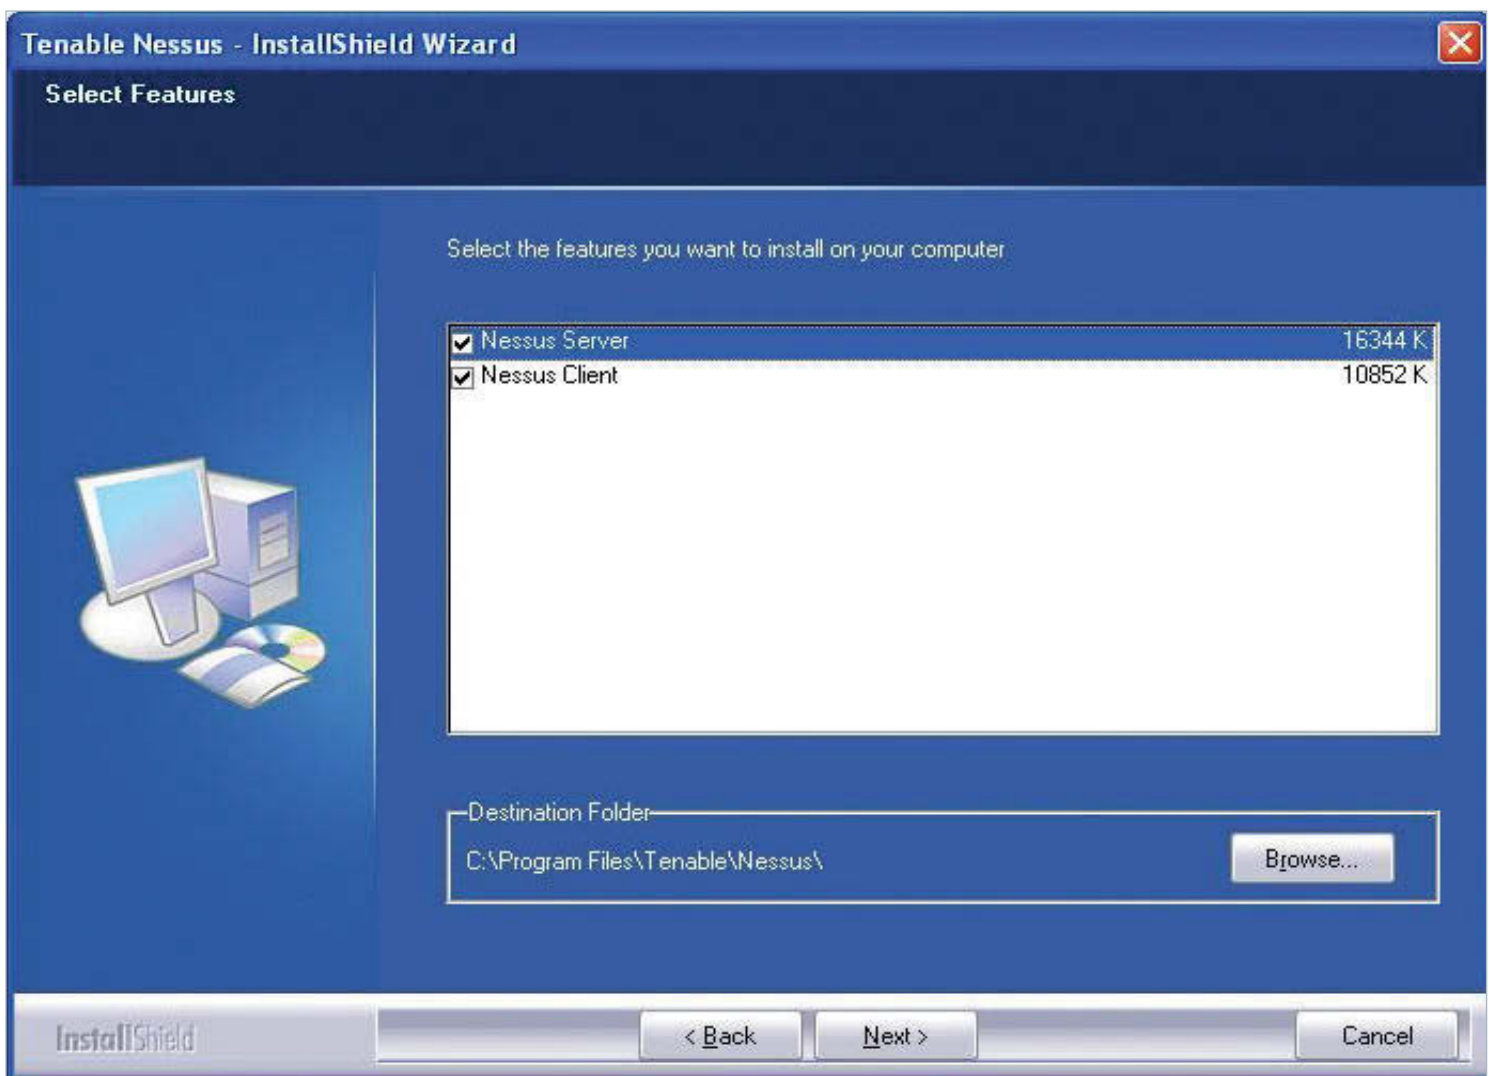


Fig 2 Downloading and installing Nessus

Masterclasses

Phishing

Phishing is a major internet security threat. With it, an external hacker can steal information such as your username and password with ease.

We've all had them – they usually work by sending messages to your email account pretending to be

trusted sites such as eBay, Paypal or even your bank. The links will automatically redirect you to a fake site that mimics the trusted site. When you type in your personal information, the hacker will use it for making online transactions with your information or your money.

The best way to avoid becoming a phishing scam victim is to use your best judgement (and to avoid clicking links in emails, but type in URLs in your browser instead). And make it a practice to use internet security software packages that have great phishing detection systems.


PHISHER



OPPORTUNITIES TO BLOCK

1. INITIAL WEBSITE COMPROMISE
2. MASS PHISHING EMAIL
3. VICTIM CLICKS ON MISLEADING URL
4. PHISH WEBSITE IS DISPLAYED
5. VICTIM SUBMITS ACCOUNT INFORMATION

COMPROMISE A HOST AND INSTALL A
PHISH WEBSITE AND MASS-MAILER



VICTIM WEB SERVER

SENDS OUT PHISHING EMAIL



VICTIM CLICKS A PHISH URL



PHISH WEBSITE IS VIEWED



VICTIM SUBMITS INFORMATION



VICTIM USERS

“Nessus scans the entire network for potential security risks and provides detailed reporting that enables the remediation of gaps”

follow the following steps in order to prevent buffer overflow attacks.

Minimise the use of library files included with the compiler – If an attacker finds a weakness within a specific library file, any application that includes this library file will also have the weakness.

Validate all user inputs – Make sure that when a user enters a string, the length will be compared

“System administrators need to be prepared to detect and mitigate vulnerabilities... Various players in the market have come up with their own solutions to help avoid them”

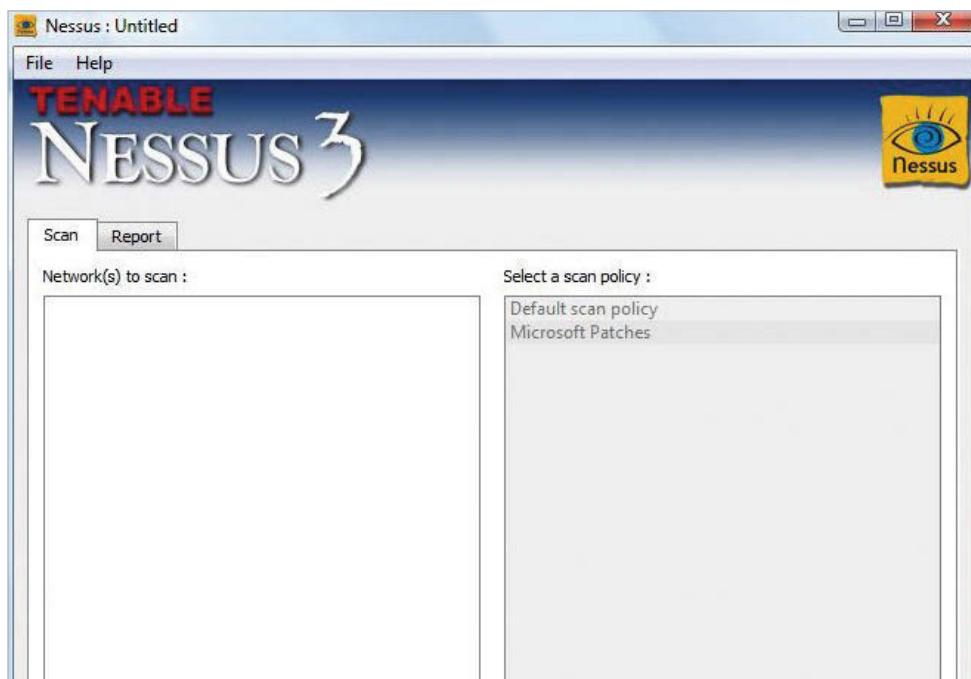


Fig 3 Scanning and reporting using Nessus

against the maximum allowed and truncated if needed.

Filter potential malicious input – Make sure that you have a filtering mechanism to filter out reserved symbols. If such symbols are included within a user input, they will cause the application to crash.

Test your application – Make sure that you thoroughly test any application before it gets deployed. Have a group of people go through the program with a fine-toothed comb and try to crash the program.

If you are looking for a less expensive vulnerability scanner, try Nessus. Nessus scans the entire network for potential security risks and provides detailed reporting that enables the remediation of gaps (Fig 3). It is good to understand the entire process of downloading and installing the free version (Fig 2, p123). Follow the process outlined at www.nessus.org/products/nessus/select-your-operating-system and you can install it without any issues. We have installed it on our machine and were able to run it successfully.

We can use Nessus to conduct vulnerability scans and it presents scan results in a report. We need to navigate through this to view the various alerts and then proper corrective actions can be taken.

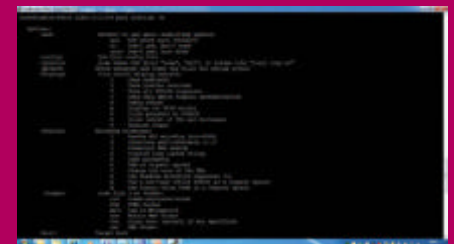
Last few words

Vulnerabilities have become a part of life these days and system administrators need to be prepared to detect and mitigate them. In this feature, we took a look at some of the most common vulnerabilities such as SQL injection, cross-site scripting and vulnerabilities related to web servers and databases. Various players in the market have come up with their own solutions that will help the organisations avoid these vulnerabilities. And not to forget, there are a large number of open source solutions and we had a quick look at one such solution – Nessus. With this information, it is now time for system administrators to go and gain knowledge of various commercial as well as open source solutions and use them in their day-to-day network/system-related activities.

An open source vulnerability scanner – Nikto


Nikto is an advanced web vulnerability scanner which can help you expose the potential holes in your web server and therefore allow you to fix them before malicious users attempt to exploit them.

For Nikto to run effectively, the operating system needs to have Perl installed on it. After downloading and extraction, Nikto can be used for scanning a host (either using an IP or a hostname). It has a detailed help page and one can get useful information on its various options.



Nikto comes with a configuration file and this file is used not only for some scanning configurations, but also for user options which can be used with every scan (such as a proxy, or the -generic parameter) One can also add the -evasion flag and this option will help you in getting an idea of what holes may exist in your configuration.

■ A detailed look at Nikto options

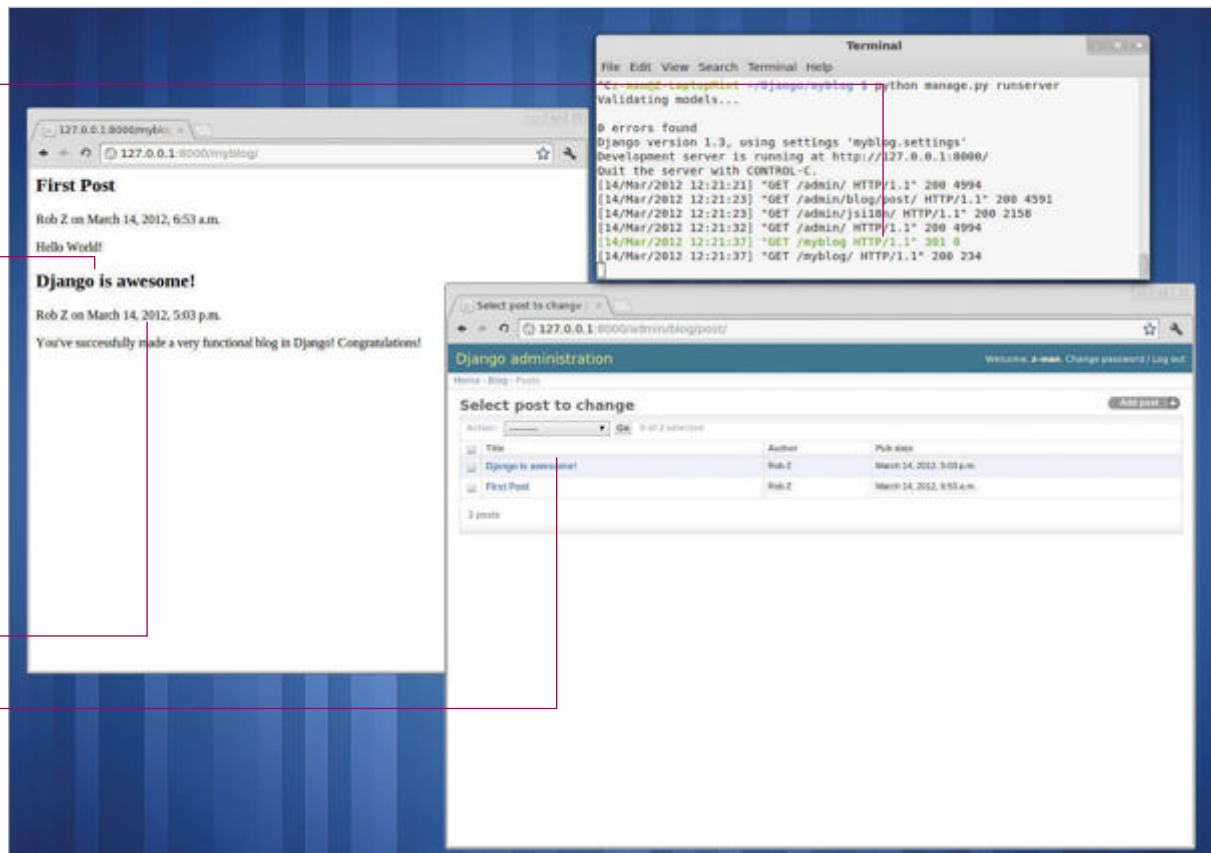


Django comes with a lightweight development server so you can test all your work locally

Django is of course able to read and write to SQL databases, but it needs very little prior knowledge to succeed in doing so

Using HTML and CSS in conjunction with Django is clear and straightforward; it's much easier to bug-fix than PHP

Django comes with a generic back-end site that is set up in seconds, and easily customisable after that



Resources

Python Source Code:
www.python.org/download/releases/2.7.2
Django Source Code:
www.djangoproject.com/download

Start building a blog with Django

Learn how to use this extremely powerful Python-based web framework to create a complete blog from scratch in record time



Creating your own blog always feels like a great accomplishment. Sure, you could use WordPress if you need a complete blog with every feature you'd ever need right now. And Tumblr exists for people who just want to write something or post pictures of corgis in space.

You don't have full control from start to finish with a prefabricated blog, though, and neither of our examples are written in the fantastic Django.

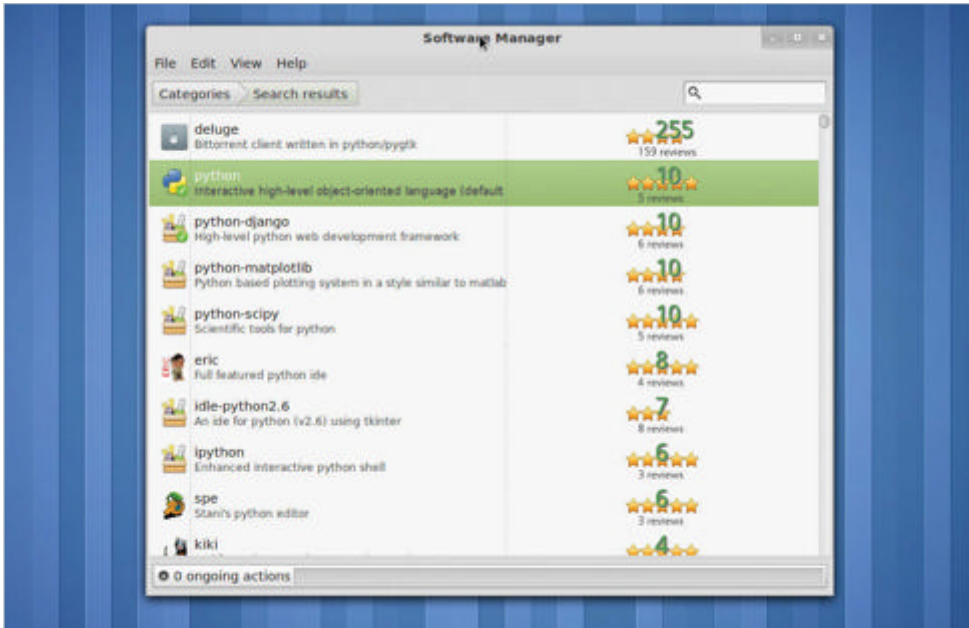
Django is of course based on Python, the object-orientated programming language designed to have clearly readable syntax. Due to its Python base, it's an incredibly powerful and simple-to-use language for web development with a vast array of applications.

So let's use it to make a blog. In this first part we will explore how to set up Django, writing and reading to a database, creating a front and back end, and some interactions with HTML.

Advisor

Rob Zwetsloot is a habitual blogger and sometime web designer who loves Python and Django more than any man really should






01 Install Python

Django is based on Python, and requires it to be installed to develop on. Python 2.7 is the recommended version, and this is installed with the python package. If you want to check your version, start the Python shell by typing 'python' into the terminal.



02 Install Django

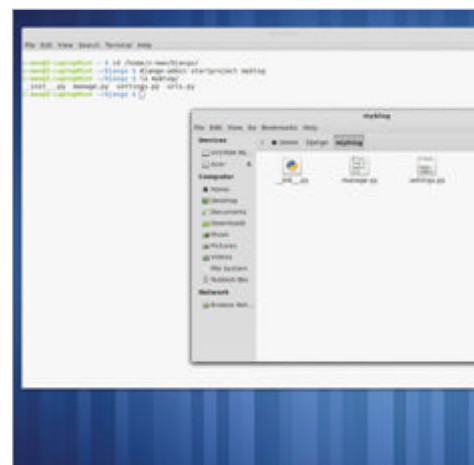
Most operating systems will have a Django package available in the repository, like python-django in Debian. The Django website has a list if you have trouble finding it, or you could build it from source. Make sure you install version 1.3.



03 Verify your Django

To make sure Django is installed properly, and that you have the right version, enter the Python shell by typing 'python' and **enter the following:**
`import django`
`print django.get_version()`

It will return a version number if it has installed correctly, which should be 1.3.



04 Start a new project

In the terminal, cd to the folder you want to develop the blog in, and then **run the next command:**
`django-admin startproject myblog`

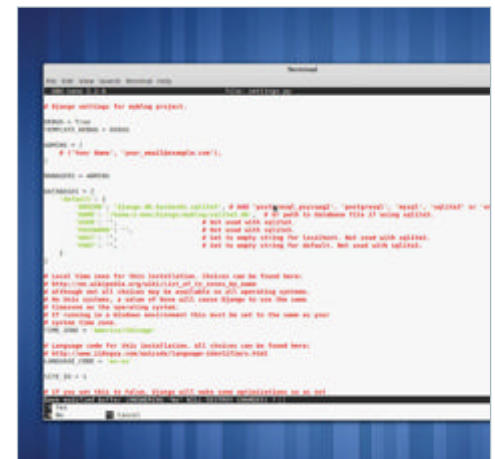
Here, 'myblog' can be replaced by whatever you wish to name the project, but we'll use it for the upcoming examples.



05 Start the development server

Django comes with a lightweight development server to test out work locally. We can also use it to check our work, so cd to the myblog folder and **then use:**
`python manage.py runserver`

If all goes well, it should return zero errors. Use Ctrl+C to exit the server.



06 Configure the database

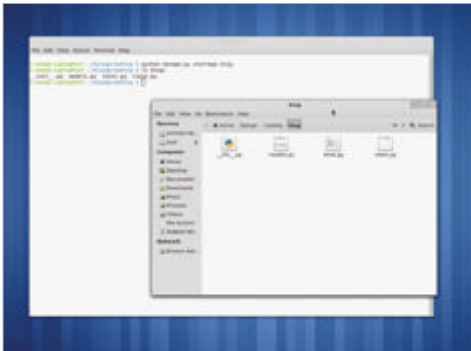
The database settings are kept in the settings.py file. Open it up with your favourite editor and go to the Databases section. **Change ENGINE to:**
`'ENGINE': 'django.db.backends.sqlite3',`
And in NAME, put the absolute path – for example:
`'NAME': '/home/user/projects/myblog/sqlite.db',`
Save and exit.

07 Create the database

The database file will be generated by **using the command:**
`python manage.py syncdb`

During the creation, it will ask you to set up a superuser, which you can do now.

The SQLite database file will be created in your myblog folder.



08 Create your blog

Now it's time to create a blog app in your project. **Type:**
`python manage.py startapp blog`

This creates the models file which is where all your data lives. You can change 'blog' to another name, but we'll use it in our examples.



09 Start your blog model

We can now take the first steps in creating our blog model. Open `models.py` and change it so it says the following:

```
from django.db import models
class Post(models.Model):
    post = models.TextField()
```

This creates the `Post` class, which has a subclass that contains your blog text.



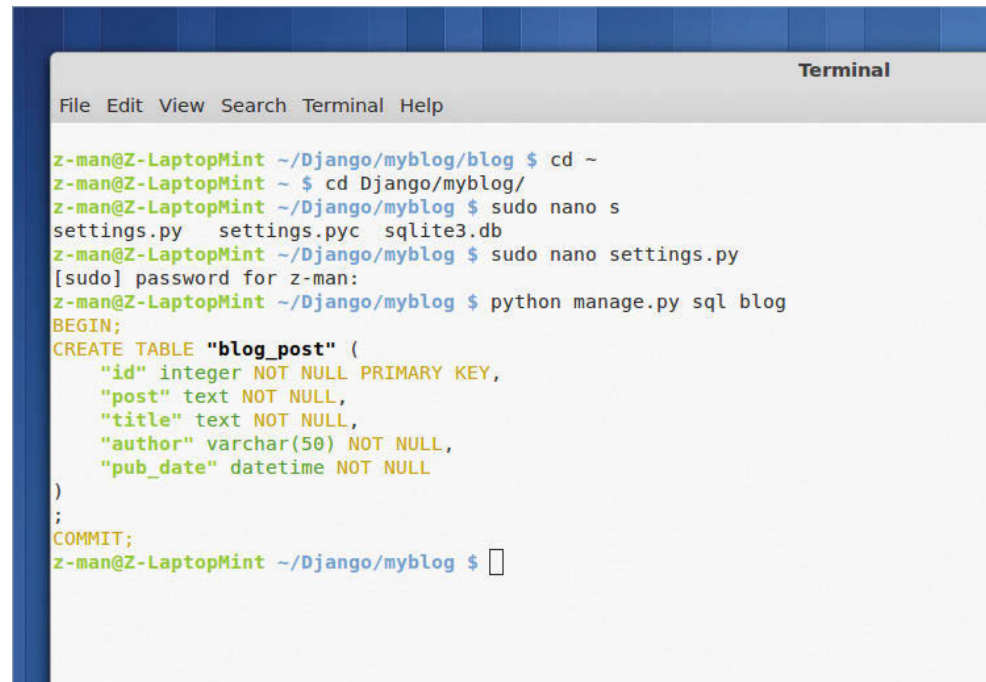
10 Customise your blog

Let's now expand the blog model a bit so it resembles a more classic blog:

```
class Post(models.Model):
    post = models.TextField()
    title = models.TextField()
    author = models.CharField(max_length=50)
    pub_date = models.DateTimeField()
```

A `CharField` needs to have a character limit defined, and `DateTimeField` holds the time values.

“You don't have full control from start to finish with a prefabricated blog – but you will with Django”



11 Install your app

Your app needs to be installed to your project, which is very simple. Open the `settings.py` file again, go to the `INSTALLED_APPS` section and add: 'blog',

Then run the following to create the database tables:

```
python manage.py sql blog
```

And finally:

```
python manage.py syncdb
```



12 Set up to post

Now we can create a post and test out our code. First though, **enter the Python shell:**
`python manage.py shell`

Then execute these commands to add all the necessary fields and data:

```
from blog.models import Post
import datetime
```

13 Let's blog

Create the post. **For this example, we will call it `test_post`:**
`test_post = Post()`

Now let's add the blog content:

```
test_post.post = 'Hello World!'
test_post.title = 'First Post'
test_post.author = 'Me'
test_post.pub_date = datetime.datetime.now()
```

And then save it with:

```
test_post.save()
```

14 Start the site back-end


To create the admin site, edit `urls.py` from the `myblog` directory, and uncomment or **add the following lines:**

```
from django.contrib import admin
admin.autodiscover()
url(r'^admin/', include(admin.site.urls)),
```

Save and exit, then edit `settings.py` and uncomment this line from `INSTALLED_APPS`:

```
'django.contrib.admin',
```

The admin site is now at `127.0.0.1:8000/admin/`.



15 Setup the admin page

The admin page has a generic, usable template, but you need to configure it to view, edit, create and delete posts. First, create a new file `admin.py` in the blog directory **and enter:**

```
from blog.models import Post
from django.contrib import admin

admin.site.register(Post)
```

To have the posts display nicely on the site, edit `models.py` **and add:**

```
class Post (models.Model):
    ...
    def __unicode__(self):
        return self.title
```

Save, and run:

```
python manage.py syncdb
```

The admin page is now usable! You should be able to see the other posts, and it's now a lot easier to add more.

16 Activate the front-end

Open up `urls.py` from the `myblog` directory in your editor and **add the following to the `urlpatterns` section:**

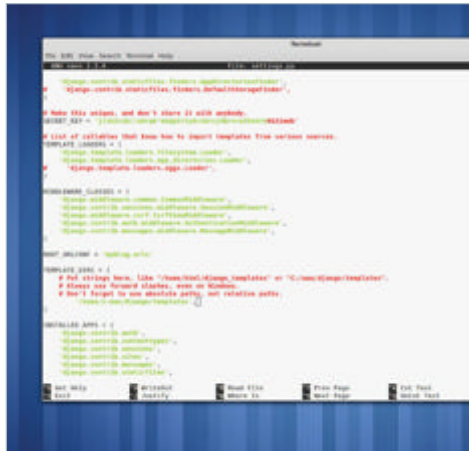
```
url(r'^myblog/', 'blog.urls.index'),
```

One of the examples in the file can be uncommented and edited to this as well. It points to a model we will now create.

17 Create another urls file

You need to create another `urls` file in the app directory, in our case `blog/urls.py`. Create it and **add the following:**

```
from django.template import Context, loader
from blog.models import Post
from django.http import HttpResponse
def index(request):
    post_list = Post.objects.all()
    t = loader.get_template('blog/index.html')
    c = Context({
        'post_list': poll_list,
    })
    return HttpResponse(t.render(c))
```



18 Start the template

The code we've just written looks for a template that currently doesn't exist. We first need to tell Django where templates are to be looked for in `settings.py`:

```
TEMPLATE_DIRS = (
    '/home/user/projects/templates',
)
```


You can put the template directory wherever you want, as long as it's referenced here.

19 Write a template

Now to write the site template. In our example, we're using `index.html`:

```
{% for post in post_list %}
  {{ post.title }}
  {{ post.author }}
  {{ post.pub_date }}
  {{ post.post }}
{% endfor %}
```

This needs to be located in a folder with the same name as your app within the template directory.



20 View your handiwork

Let's make sure this worked. Start the developer **server with:**

```
python manage.py runserver
```

And navigate to `127.0.0.1:8000/myblog/`.

It's not pretty, but you should have successfully called upon your stored data. We'll spend the next steps tidying it up a bit.

“Django is an incredibly powerful and simple-to-use language for web development”

21 Format the front page

Go back into the template file, `index.html`, and **add the following html tags:**

```
{% for post in post_list %}
  <h2>{{ post.title }}</h2>
  {{ post.author }} on {{ post.pub_date }}
<p>{{ post.post }}</p>
{% endfor %}
```

This is just an example – the post can be in any order with any tags.

22 Spruce up the admin list

We'll do this in the `admin.py` file in our blog directory; open it in your editor and make the **following changes:**

```
from blog.models import Post
from django.contrib import admin
class Admin(admin.ModelAdmin):
    list_display = ['title', 'author', 'pub_date']
admin.site.register(Post, Admin)
```

In this case `'list_display'` is a fixed variable name.

23 A logical post page

The new post page on the site might not be in an order you're comfortable with. We'll change that now in `admin.py` with the **following additions:**

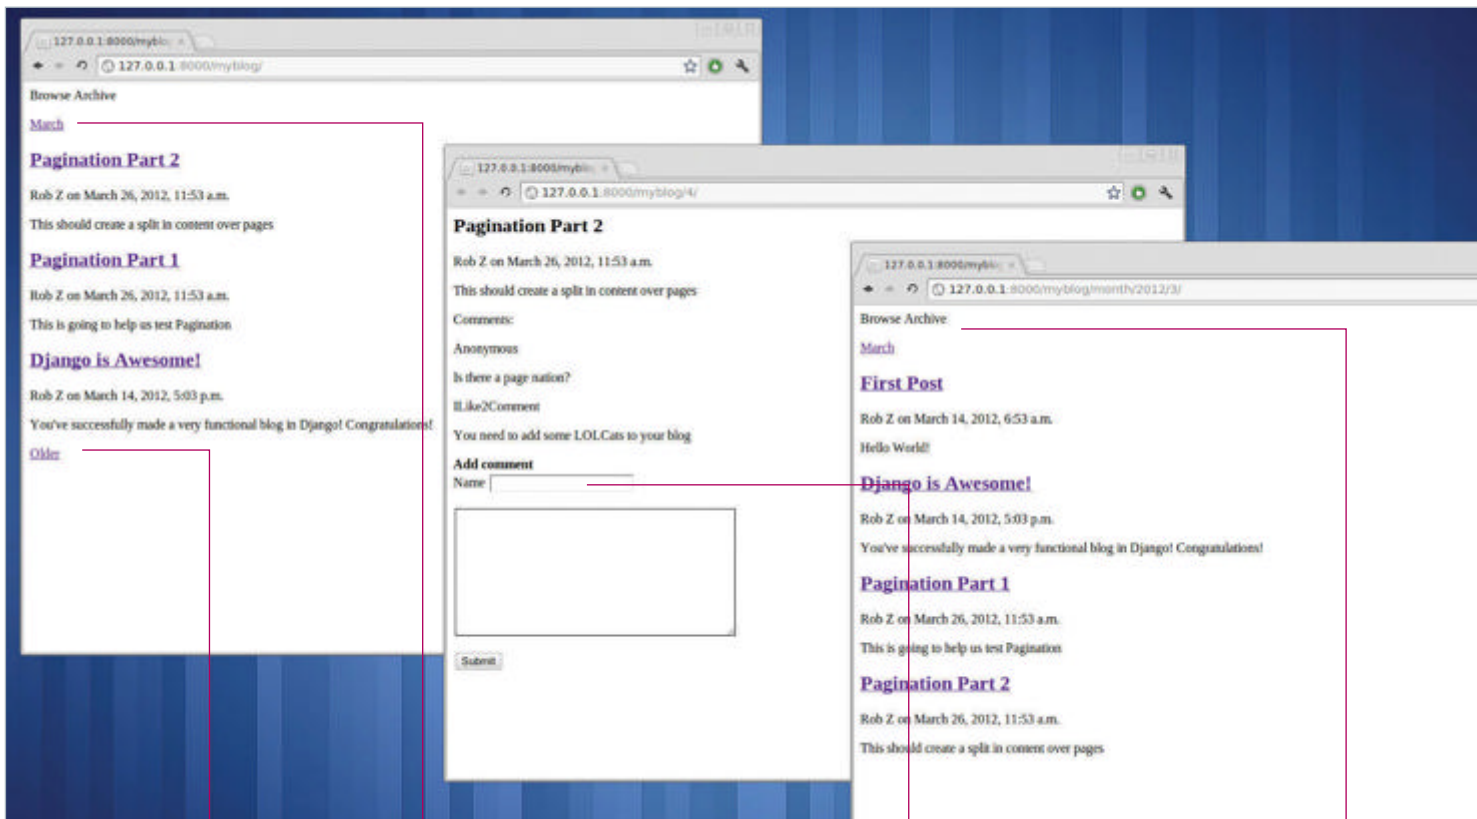
```
class Admin(admin.ModelAdmin):
    list_display = ['title', 'author', 'pub_date']
    fields = ['title', 'pub_date', 'author', 'post']
admin.site.register(Post, Admin)
```

Remember to save!

24 A functional blog

So there you have it! Navigating to `127.0.0.1:8000/admin/` or `127.0.0.1:8000/myblog/` will show off the fine work you've created. Django is dead easy to use once you know how, and there are plenty of tweaks you should be able to make after this tutorial.

Masterclasses



Django has built-in code to deal with pagination very cleanly and effectively

With Django we can make simple sidebars that list archives by month

Allow your readers to give you feedback, and moderate them in the admin panel

With minimal extra code, our template can display the month archive from the sidebar

Add content to your Django blog

We continue building an awesome blog using the powerful Django, and this part is all about the front-end content delivery

Advisor

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



In our last tutorial we showed you how to build the most basic of blogs, and learned how to use a bit of Django in the process. We can now

set up a new project, create a database and

write code to read and write to the database. All simple stuff, but it's core to building websites where Django might be called upon.

This time we will give the front end of the site an overhaul, making it more of the standard you would expect from a modern blog. This will include a sidebar, pages, post pages and the ability to add and moderate comments. In the process we will learn some more of the benefits that come with using Django to develop websites.

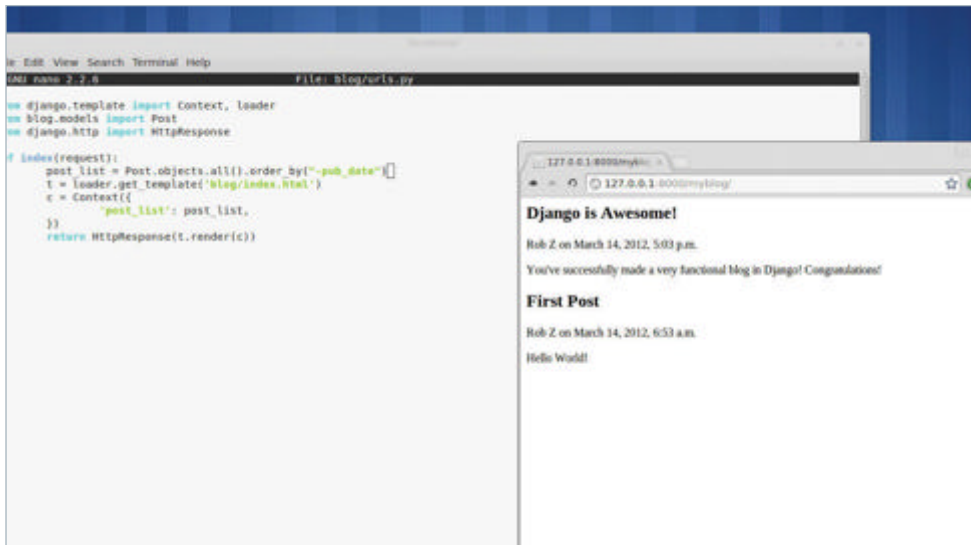
You should keep using Django 1.3 for this tutorial, as discussed in part one.

Resources

Python base:

<http://www.python.org/download/>

Django source: <https://www.djangoproject.com/download/>



01 New blog order

We left off last time with the blog displaying posts in chronological order, which isn't very helpful to readers. To correct this, open up `urls.py` in the blog folder and edit the following line:

```
post_list = Post.objects.all().order_by("pub-date")
```

This makes sure that posts are displayed in reverse order (newest first).

02 A view to a page

You'll want to be able to link specific pages, of course, and to do that we first have to define what goes into these pages in the `urls.py` file in the **blog folder**:

```
def post_page(request, post_id):
    post_page = Post.objects.get(pk=post_id)
    return render_to_response('blog/post.html', {'post_page': post_page})
```

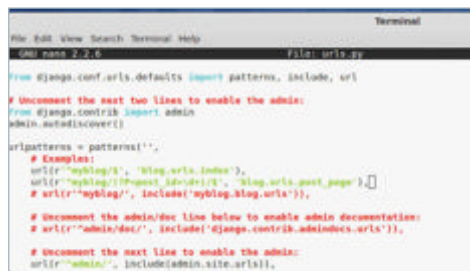


03 Clean up your code

You may notice that we used a different return command to the index definition – this is a shortcut that makes writing the code a bit easier. To get it working, **add**:

```
from django.shortcuts import render_to_response
```

We recommend that you edit the index code to match `post_page`.




04 Edit URLs

In `urls.py` in `myblog` we need to make some additions and modifications for the website to direct to the **post correctly**:

```
url(r'^myblog/$', 'blog.urls.index'),
url(r'^myblog/(?P<post_id>\d+)/$', 'blog.urls.post_page'),
```

The `post_id` is the number of the post, which is auto-generated. The '\$' is important to make the redirection work.



05 A post template

We told the `post_page` to point towards a template we now need to create. In the same location as `index.html`, create `post.html` with the following formatting to resemble the **front page**:

```
<h2>{{ post_page.title }}</h2>
{{ post_page.author }} on {{ post_page.pub_date }}
<p>{{ post_page.post }}</p>
```



06 Link to the page

Let's get these links working from the main page. Open up the `index.html` file and make the following change:

```
<h2><a href="/myblog/{{ post.pk }}"/>{{ post.title }}</a></h2>
```

This is a very simple addition using an absolute link, and requires no fiddling with the views or model.



07 Pagination

To get blog posts split up over pages, we need to make some additions to `urls.py` in the **blog folder**:

```
post_list = Post.objects.all().order_by("pub-date")
paginator = Paginator(post_list, 3)
try: list_page = request.GET.get("list_page", '1')
except ValueError: list_page = 1
post_list = paginator.page(list_page)
return render_to_response('blog/index.html', {'post_list': post_list})
```



08 Please turn over

Now we need to add the navigation links to the blog, so open the `index template` for editing:

```
{% if post_list.has_previous %}
  <a href="?list_page={{ post_list.previous_page_number }}">Newer </a>
{% endif %}
{% if post_list.has_next %}
  <a href="?list_page={{ post_list.next_page_number }}">Older </a>
{% endif %}
```

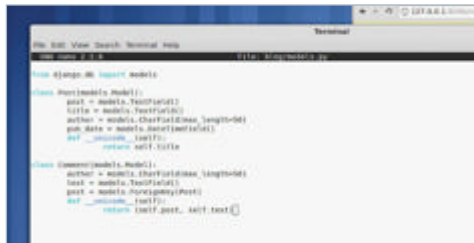
Masterclasses

09 Wrong page

Let's add a quick bit of code to return somebody to the previous page if they get the URL wrong:

```
from django.core.paginator import  
Paginator, EmptyPage, InvalidPage
```

```
try:  
    post_list = paginator.page(list_page)  
except (EmptyPage, InvalidPage):  
    post_list = paginator.page(paginator.  
num_pages)  
The last part replaces 'post_list = paginator.  
page(list_page)'.
```

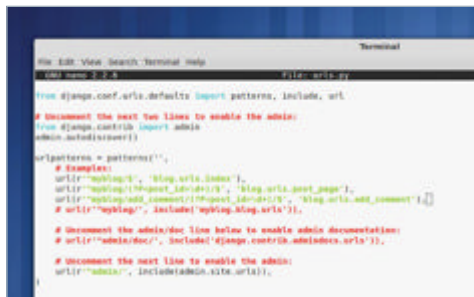


10 Have your say

Everyone has their opinion on the internet. You can give your readers the ability to comment, and we'll start **by editing models.py**:

```
class Comment(models.Model):  
    author = models.CharField(max_  
length=50)  
    text = models.TextField()  
    post = models.ForeignKey(Post)  
    def __unicode__(self):  
        return (self.post, self.text)
```

We've made it so they can put their name with a comment.



11 Back to the comment

We now need to add a small line to the urls.py file in myblog so the comment can be posted then sent **back to the original page**:

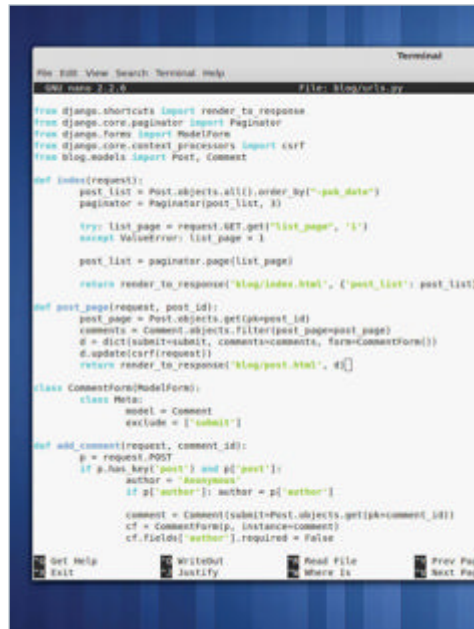
```
url(r'^myblog/add_comment/(\d+)/$', 'blog.  
urls.add_comment'),  
This URL pattern calls the ID of the page that you're on.
```

“We need to be able to process the data and metadata in the forms”

12 Form a comment

We need to be able to process the data and metadata in the forms, so let's add a class to urls.py in the blog folder with the **following additions**:

```
from django.forms import ModelForm  
from blog.models import Post, Comment  
class CommentForm(ModelForm):  
    class Meta:  
        model = Comment  
        exclude = ['post']
```



13 In the post

We need to attribute the comments to the post they're being made on, so update the **post_page definition**:

```
from django.core.context_processors import  
csrf  
def post_page(request, post_id):  
    post_page = Post.objects.get(pk=post_  
id)  
    comments = Comment.objects.  
filter(post=post_page)  
    d = dict(post_page=post_page,  
comments=comments, form=CommentForm())  
    d.update(csrf(request))  
    return render_to_response('blog/post.  
html', d)
```

The CSRF tag is to prevent cross-site request forgery.



14 Comment template

Let's get the post page ready for comments by **adding this to post.html**:

```
<p>Comments:</p>  
{% for comment in comments %}  
    {{ comment.author }}  
    <p>{{ comment.text }}</p>  
{% endfor %}  
<strong>Add comment</strong>  
<form action="{% url blog.urls.add_comment  
post_page.id %}" method="POST">{% csrf_  
token %}  
    Name {{ form.author }}  
    <p>{{ form.text }}</p>  
    <input type="submit" value="Submit">  
</form>
```

15 Define your comments

The final step is defining the comments in blog/urls.py, **and it's a big one**:

```
def add_comment(request, comment_id):  
    p = request.POST  
    if p.has_key('text') and p['text']:  
        author = 'Anonymous'  
        if p['author']: author =  
p['author']  
        comment = Comment(post=Post.  
objects.get(pk=comment_id))  
        cf = CommentForm(p,  
instance=comment)  
        cf.fields['author'].required =  
False  
        comment = cf.save(commit=False)  
        comment.author = author  
        comment.save()  
        return HttpResponseRedirect(reverse('b  
og.urls.post_page', args=[comment_id]))
```

This ensures text has been entered, and if not specified author is 'Anonymous'. Before testing, run syncdb so comment tables can be created.

```

File Edit View Search Terminal Help
GNU nano 2.2.6 File: blog/admin.py

from blog.models import Post, Comment
from django.contrib import admin

class PostAdmin(admin.ModelAdmin):
    list_display = ['title', 'author', 'pub_date']
    fields = ['title', 'pub_date', 'author', 'post']

admin.site.register(Post, PostAdmin)

class CommentAdmin(admin.ModelAdmin):
    list_display = ['text', 'author', 'post']

admin.site.register(Comment, CommentAdmin)

```


16 Administrate

Like the posts, we can get the Admin page to see comments. Start editing `blogs/admin.py` to get this **feature added:**

```

from blog.models import Post, Comment
from django.contrib import admin
class PostAdmin(admin.ModelAdmin):
    list_display = ['title', 'author',
'pub_date']
    fields = ['title', 'pub_date',
'author', 'post']
admin.site.register(Post, PostAdmin)

```



17 Comment-specific admin features

Now we can add the comment-specific admin features without causing any clashes:

```

class CommentAdmin(admin.ModelAdmin):
    list_display = ['text', 'author',
'post']
admin.site.register(Comment, CommentAdmin)

```

This will show the comments on the admin site, and you can see the comment, the author and the post it's connected to.

18 Sidebar beginnings

Django makes it pretty easy to order posts by years and months, but first we need to import some new models into `blog/urls.py`:

```

import time
from calendar import month_name

```

We're going to define two new functions, `month_timeline` and `month`, to make the sidebar.

```

class CommentAdmin(admin.ModelAdmin):
    list_display = ['text', 'author', 'post']
    fields = ['text', 'author', 'post']
admin.site.register(Comment, CommentAdmin)

def month_timeline():
    year, month = time.localtime()[2:]
    begin = Post.objects.order_by('pub_date')[0]
    month_begin = begin.pub_date.month
    year_begin = begin.pub_date.year
    month_list = []
    for y in range(year, year_begin-1, -1):
        start, end = 12, 0
        if y == year: start = month
        if y == year_begin: end = month_begin-1
        for m in range(start, end, -1):
            month_list.append((y, m, month_name[m]))
        return month_list

```

19 Start to define month_timeline

First we need to get all the information from the posts:

```

def month_timeline():
    year, month = time.localtime()[2:]
    begin = Post.objects.order_by('pub_date')[0]
    month_begin = begin.pub_date.month
    year_begin = begin.pub_date.year
    month_list = []

```

The `[:2]` makes sure we only get the time information we need.

```

class CommentAdmin(admin.ModelAdmin):
    list_display = ['text', 'author', 'post']
    fields = ['text', 'author', 'post']
admin.site.register(Comment, CommentAdmin)

def month_timeline():
    year, month = time.localtime()[2:]
    begin = Post.objects.order_by('pub_date')[0]
    month_begin = begin.pub_date.month
    year_begin = begin.pub_date.year
    month_list = []
    for y in range(year, year_begin-1, -1):
        start, end = 12, 0
        if y == year: start = month
        if y == year_begin: end = month_begin-1
        for m in range(start, end, -1):
            month_list.append((y, m, month_name[m]))
        return month_list

```

20 Finish your definition

Now we will order the posts by month and year starting from our first month.

```

for y in range(year, year_begin-1, -1):
    start, end = 12, 0
    if y == year: start = month
    if y == year_begin: end = month_begin-1
    for m in range(start, end, -1):
        month_list.append((y, m, month_name[m]))
    return month_list

```

21 Return to reader

With the list organised, we can now define month so we can display it on the blog:

```

def month(request, year, month):
    post_list = Post.objects.filter(pub_date__year=year, pub_date__month=month)
    return render_to_response('blog/index.html', dict(sidebar_list=post_list, month_list=month_timeline()))

```

Now we need to link it up to the index template.

22 Finalise your sidebar definition

Edit the return command on the index function to include the sidebar information:

```

return render_to_response('blog/index.html', dict(post_list=post_list, sidebar_list=post_list.object_list, month_list=month_timeline()))

```

Then add this line to `urls.py` in myblog so a month page can be rendered:

```

url(r'^myblog/month/(\d+)/(\d+)/$', 'blog.urls.month'),

```

All we need to do now is display the information on the site.

23 Sidebar on the web

Go to the index template. First of all, change the first line of the post forloop to:

```

{% for post in sidebar_list %}
Simple enough. Now we need to add the sidebar information:
{% for month in month_list %}
<p><a href="{% url blog.urls.month month.0 month.1 %}">{{ month.2 }}</a></p>
{% endfor %}

```



24 Sidebar finale

Obviously it's not at the side right now – that's a job for the HTML and CSS. The info is there, though, and you can manipulate it any way you want. However, your blog is now a lot more friendly to your readers.

Advisor

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



Customise your blog in Django

In the last of our series of blog tutorials, we'll cover some of the more advanced features you can utilise with the power of Django

Resources

Python base:

<http://www.python.org/download/>

Django source: <https://www.djangoproject.com/download/>

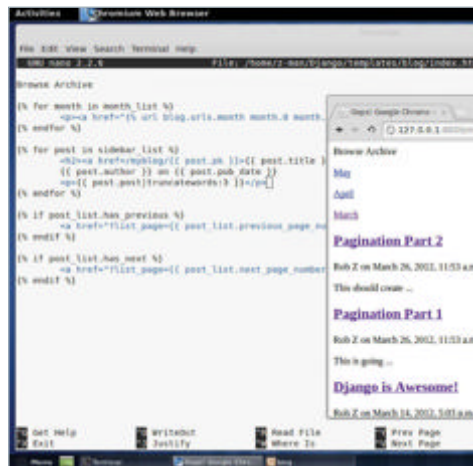


We're drawing to the end of our Django tutorial, and so far we've built the blog up to create and display posts, allow people to make comments, and filter posts by month like a classic blog sidebar. We still have a bit to go until it becomes more like a classic blog, though.

Here we're going to add summaries, excerpts, categories and finally an RSS feed. This allows us to look at a few things – firstly we'll get a better

understanding of cross-model referencing and how that works in the admin site. We will also go through how to make changes to the database, and how Django helps when creating an SQL query.

Finally, the RSS feed is part of a standard feed library in Django itself. We will learn how to import and use it to create a simple list of the latest entries that click through to the posts. By the end of the tutorial your blog will be finally finished!



01 Summarise

On a normal blog we're going to have much longer articles. We can generate a summary of each of these on the index page template like so:

```
<p>{{ post.post|truncatewords:3 }}</p>
```

This automatically takes the first three words of the post – of course, you can use any number.



02 Manual excerpt

If you don't want an automatic summary, we can add an excerpt field to our post model so you can craft one manually:

```
excerpt = models.TextField()
```

To limit the characters in your excerpt, use a CharField like for our author section.



03 Write an excerpt

To write the excerpt, or append it to the previous posts, we'll have to add it to the admin page. Open up admin.py and edit the fields section of the AdminPost class to add excerpt:

```
fields = ['title', 'pub_date', 'author', 'post', 'excerpt']
```

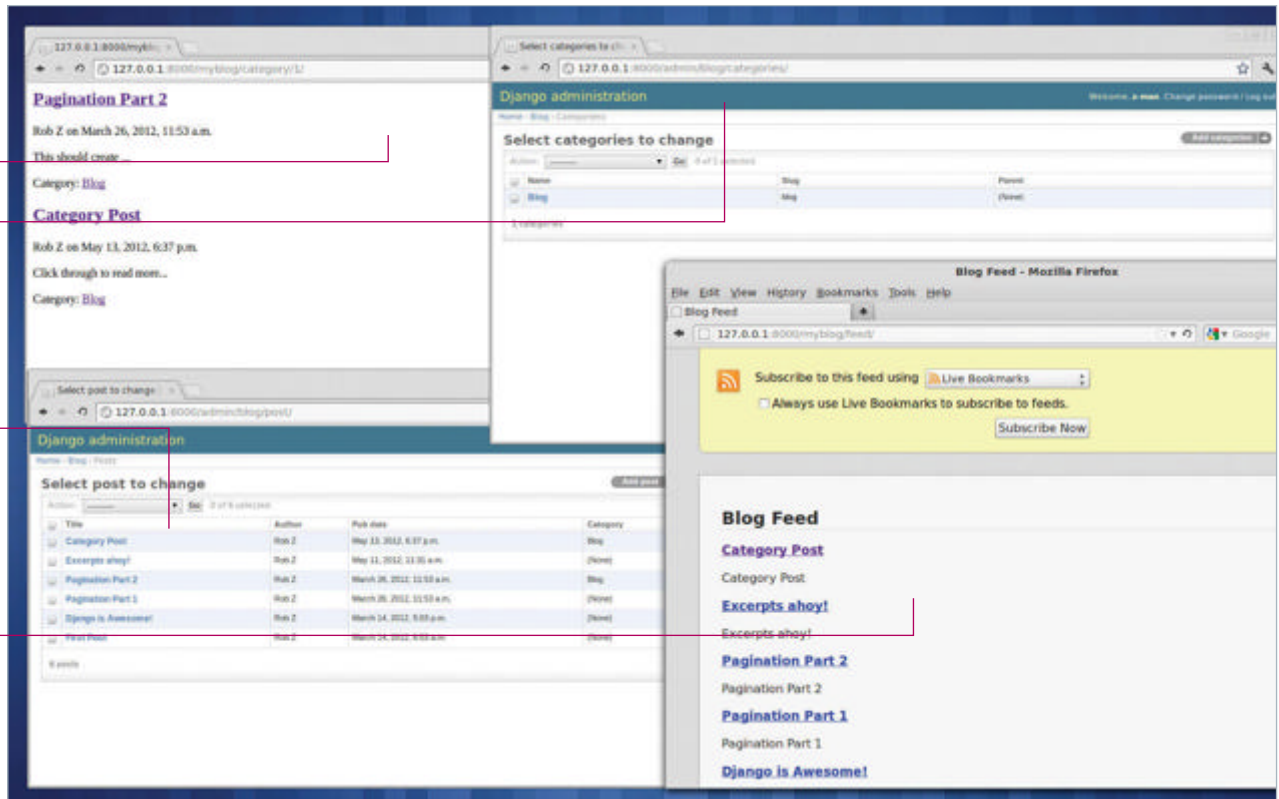
“We're going to add summaries, excerpts, categories and finally an RSS feed”

Have automatic summaries or manually crafted excerpts for your blog posts

Create and manage parent and child categories as a separate function of the blog

Learn how to alter the database to create posts with categories, and add them to other posts

Create custom RSS feeds using built-in Django functions



```

25 You want to create a post:
26 post = models.ForeignKey(post, on_delete=models.CASCADE)
27 post.save()
28
29 You want to create a post:
30 post = models.ForeignKey(post, on_delete=models.CASCADE)
31 post.save()
32
33 You want to create a post:
34 post = models.ForeignKey(post, on_delete=models.CASCADE)
35 post.save()
36
37 You want to create a post:
38 post = models.ForeignKey(post, on_delete=models.CASCADE)
39 post.save()
40
41 You want to create a post:
42 post = models.ForeignKey(post, on_delete=models.CASCADE)
43 post.save()
44
45 You want to create a post:
46 post = models.ForeignKey(post, on_delete=models.CASCADE)
47 post.save()
48
49 You want to create a post:
50 post = models.ForeignKey(post, on_delete=models.CASCADE)
51 post.save()
52
53 You want to create a post:
54 post = models.ForeignKey(post, on_delete=models.CASCADE)
55 post.save()
56
57 You want to create a post:
58 post = models.ForeignKey(post, on_delete=models.CASCADE)
59 post.save()
60
61 You want to create a post:
62 post = models.ForeignKey(post, on_delete=models.CASCADE)
63 post.save()
64
65 You want to create a post:
66 post = models.ForeignKey(post, on_delete=models.CASCADE)
67 post.save()
68
69 You want to create a post:
70 post = models.ForeignKey(post, on_delete=models.CASCADE)
71 post.save()
72
73 You want to create a post:
74 post = models.ForeignKey(post, on_delete=models.CASCADE)
75 post.save()
76
77 You want to create a post:
78 post = models.ForeignKey(post, on_delete=models.CASCADE)
79 post.save()
80
81 You want to create a post:
82 post = models.ForeignKey(post, on_delete=models.CASCADE)
83 post.save()
84
85 You want to create a post:
86 post = models.ForeignKey(post, on_delete=models.CASCADE)
87 post.save()
88
89 You want to create a post:
90 post = models.ForeignKey(post, on_delete=models.CASCADE)
91 post.save()
92
93 You want to create a post:
94 post = models.ForeignKey(post, on_delete=models.CASCADE)
95 post.save()
96
97 You want to create a post:
98 post = models.ForeignKey(post, on_delete=models.CASCADE)
99 post.save()
100

```

04 Excerpt or summary You can replace the post content in the index template with the excerpt, but we can keep it as a backup for if the excerpt is empty:

```
{% if post.excerpt %} <p>{{ post.excerpt }}</p> {% else %} <p>{{ post.post|truncatewords:3 }}</p> {% endif %}
```

```

1 # Django ORM
2 # Django ORM
3 # Django ORM
4 # Django ORM
5 # Django ORM
6 # Django ORM
7 # Django ORM
8 # Django ORM
9 # Django ORM
10 # Django ORM
11 # Django ORM
12 # Django ORM
13 # Django ORM
14 # Django ORM
15 # Django ORM
16 # Django ORM
17 # Django ORM
18 # Django ORM
19 # Django ORM
20 # Django ORM
21 # Django ORM
22 # Django ORM
23 # Django ORM
24 # Django ORM
25 # Django ORM
26 # Django ORM
27 # Django ORM
28 # Django ORM
29 # Django ORM
30 # Django ORM
31 # Django ORM
32 # Django ORM
33 # Django ORM
34 # Django ORM
35 # Django ORM
36 # Django ORM
37 # Django ORM
38 # Django ORM
39 # Django ORM
40 # Django ORM
41 # Django ORM
42 # Django ORM
43 # Django ORM
44 # Django ORM
45 # Django ORM
46 # Django ORM
47 # Django ORM
48 # Django ORM
49 # Django ORM
50 # Django ORM
51 # Django ORM
52 # Django ORM
53 # Django ORM
54 # Django ORM
55 # Django ORM
56 # Django ORM
57 # Django ORM
58 # Django ORM
59 # Django ORM
60 # Django ORM
61 # Django ORM
62 # Django ORM
63 # Django ORM
64 # Django ORM
65 # Django ORM
66 # Django ORM
67 # Django ORM
68 # Django ORM
69 # Django ORM
70 # Django ORM
71 # Django ORM
72 # Django ORM
73 # Django ORM
74 # Django ORM
75 # Django ORM
76 # Django ORM
77 # Django ORM
78 # Django ORM
79 # Django ORM
80 # Django ORM
81 # Django ORM
82 # Django ORM
83 # Django ORM
84 # Django ORM
85 # Django ORM
86 # Django ORM
87 # Django ORM
88 # Django ORM
89 # Django ORM
90 # Django ORM
91 # Django ORM
92 # Django ORM
93 # Django ORM
94 # Django ORM
95 # Django ORM
96 # Django ORM
97 # Django ORM
98 # Django ORM
99 # Django ORM
100 # Django ORM

```

05 Database error If you've decided to test the changes, you'll

have noticed our web server has stopped working. This is because there is no excerpt column in our database. Therefore we need to add the excerpt column. To find out how, run:

```
$ python manage.py sqlall blog
```

06 Database query The output will show you what the SQL code is to add the models to the database. We want to add the excerpt field specifically, which should look something like this:

```
"excerpt" text NOT NULL
```

Make a note of it.

```

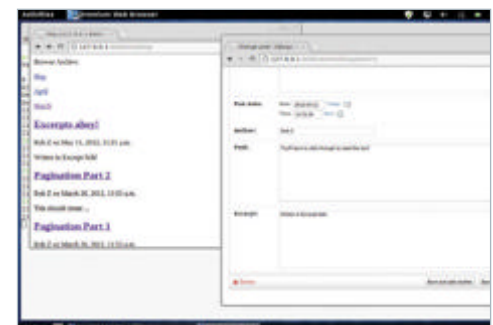
z-maw@LaptopMint ~/Django/myblog $ python manage.py dbshell
SQLite version 3.7.7 2011-06-23 19:49:22
Enter ".help" for instructions
Enter SQL statements terminated with a ";"
sqlite> ALTER TABLE "blog_post"
...> ADD "excerpt" text
...>

```

07 Alter table To get into the database shell and add the field, run: `$ python manage.py dbshell` Then we need to use an ALTER TABLE query:

ALTER TABLE "blog_post". And then enter the code we noted down like so: ADD "excerpt" text;

08 Save the changes We've removed NOT NULL as we already have entries that won't have an excerpt, and want to make it so an auto summary can be made. Save the changes with: COMMIT; and then exit the shell with: .quit



09 Test it out Now we can test out the excerpt code – create a new post or edit an existing one to have an excerpt. If you've followed our steps correctly it should work; if not, you may need to do a bit of bug fixing.

Masterclasses

```
class Categories(models.Model):
    name = models.CharField(unique=True,
                           max_length=200)
    slug = models.SlugField(unique=True,
                           max_length=100)
    parent = models.ForeignKey('self',
                              blank=True, null=True,
                              related_name='child')
    def __unicode__(self):
        return (self.name)
```

10 Category model

We can add a model for blog categories:

```
class Categories(models.Model):
    name = models.CharField(unique=True,
                           max_length=200)
    slug = models.SlugField(unique=True,
                           max_length=100)
    parent = models.ForeignKey('self',
                              blank=True, null=True,
                              related_name='child')
    def __unicode__(self):
        return (self.name)
```

This allows for parent and child categories.

```
from django.contrib import admin
from blog.models import Post, Comment, Categories

class PostAdmin(admin.ModelAdmin):
    list_display = ['title', 'author', 'pub_date']
    fields = ['title', 'pub_date', 'author', 'post', 'excerpt']
admin.site.register(Post, PostAdmin)

class CommentAdmin(admin.ModelAdmin):
    list_display = ['text', 'author', 'post']
admin.site.register(Comment, CommentAdmin)

class CategoriesAdmin(admin.ModelAdmin):
    list_display = ['name', 'slug', 'parent']
    fields = ['name', 'slug', 'parent']
admin.site.register(Categories, CategoriesAdmin)
```

11 Administrate categories

We can add it to the admin site by creating a Categories section in admin.py:

```
class CategoriesAdmin(admin.ModelAdmin):
    list_display = ['name', 'slug', 'parent']
    fields = ['name', 'slug', 'parent']
admin.site.register(Categories,
                   CategoriesAdmin)
```

Before we can make categories, though, we need to create the database table:

```
$ python manage.py syncdb
```

```
class Categories(models.Model):
    name = models.CharField(unique=True,
                           max_length=200)
    slug = models.SlugField(unique=True,
                           max_length=100)
    parent = models.ForeignKey('self',
                              blank=True, null=True,
                              related_name='child')
    def __unicode__(self):
        return (self.name)
```

12 Categorise the posts

Similarly to what we did with the comments,

we want to add a ForeignKey to the Post model so we can attribute a post to a category. Add this line:

```
category = models.ForeignKey(Categories)
```

And move Categories to the top of models.py.

```
class Post(models.Model):
    title = models.CharField(max_length=200)
    author = models.CharField(max_length=100)
    pub_date = models.DateTimeField()
    excerpt = models.CharField(max_length=200)
    category = models.ForeignKey(Categories)
    def __unicode__(self):
        return (self.title)
```

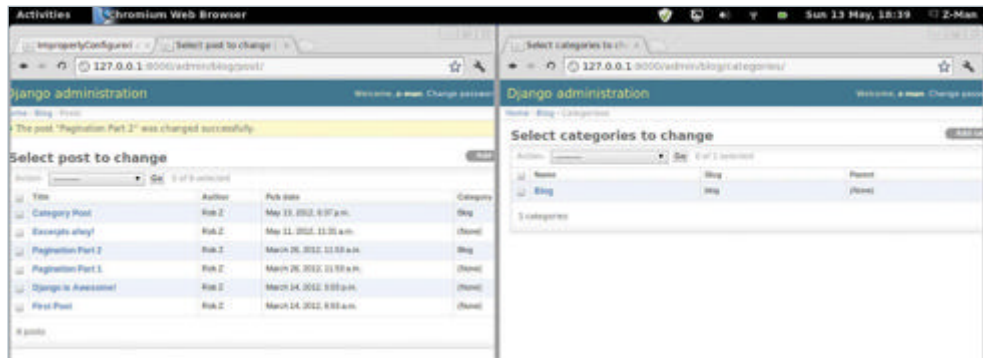
13 Database category

Like before, we'll find out the SQL needed to alter the table: `$ python manage.py sqlall blog` Which for our example returns a somewhat different code than before: "category_id" integer NOT NULL REFERENCES "blog_categories" ("id") It's an ID we're getting, not text, from the categories table.

```
ALTER TABLE "blog_post"
ADD "category_id" integer NOT NULL REFERENCES "blog_categories" ("id")
```

14 Alter table – part 2

Again let's enter the database shell: `python manage.py dbshell` We'll continue much like before, but with the new code: `ALTER TABLE "blog_post" ADD "category_id" integer REFERENCES "blog_categories" ("id");` And finally, to save: `COMMIT;`



15 Administrate categories – part 2

Now we can go back to admin.py and add the new category fields to the PostAdmin model: `list_display = ['title', 'author', 'pub_date', 'category']` fields = ['title', 'pub_date', 'author', 'post', 'excerpt', 'category'] Our previous blog posts with no category have disappeared! To fix this, go back to models.py and make this change to the Post model: `category = models.ForeignKey(Categories, blank=True, null=True)` So we can now create categories separately, assign them to posts, and view posts without a category.

“We can now create categories separately”



16 Category display

As our urls.py in the blog directory gets all the post fields, to the index template we just add: `<p>Category: {{ post.category }}</p>` And to the post template: `<p>Category: {{ post_list.category }}</p>`

```
<p>Category: {{ post.category }}</p>
```

17 Category page

First we need to define our category in blog/urls.py. Import Categories and then add: `def blog_categories(request, category_id):` categories = Categories.objects.get(pk=category_id) We need the category_id to call the corresponding posts.


```

comment.save()
return HttpResponseRedirect(reverse('blog.urls.post_page', args=[comment_id]))

def month_timeline():
    year, month = time.localtime()[:2]
    begin = Post.objects.order_by('pub_date')[0]
    month_begin = begin.pub_date.month
    year_begin = begin.pub_date.year
    month_list = []

    for y in range(year, year_begin-1, -1):
        start, end = 12, 0
        if y == year: start = month
        if y == year_begin: end = month_begin-1
        for m in range(start, end, -1):
            month_list.append((y, m, month_name[m]))
    return month_list

def month(request, year, month):
    post_list = Post.objects.filter(pub_date_year=year, pub_date_month=month)
    return render_to_response('blog/index.html', dict(sidebar_list=post_list, month_list=month_timeline()))

def blog_categories(request, category_id):
    categories = Categories.objects.get(pk=category_id)
    category_posts = Post.objects.filter(category=categories)
    return render_to_response('blog/categories.html', dict(category_posts=category_posts, categories=categories))

```

18 Category definition

Finish the definition by using the parent_id to filter the correct Posts, then render the response: `category_posts = Post.objects.filter(category=categories)` return `render_to_response('blog/categories.html', dict(category_posts=category_posts, categories=categories))`

Again we're calling a new template that we'll construct shortly.

```

from django.conf.urls.defaults import patterns, include, url

# Uncomment the next two lines to enable the admin:
from django.contrib import admin
admin.autodiscover()

urlpatterns = patterns('',
    # Examples:
    url(r'^myblog/$', 'blog.urls.index'),
    url(r'^myblog/(?P<post_id>[0-9]+)$', 'blog.urls.post_page'),
    url(r'^myblog/add_comment/(?P<id>[0-9]+)$', 'blog.urls.add_comment'),
    url(r'^myblog/month/(?P<id>[0-9]+)$', 'blog.urls.month'),
    url(r'^myblog/category/(?P<category_id>[0-9]+)$', 'blog.urls.blog_categories'),
    url(r'^myblog/', include('myblog.blog.urls')),
    # Uncomment the admin/doc line below to enable admin documentation:
    # url(r'^admin/doc/', include('django.contrib.admin.views.urls')),
    # Uncomment the next line to enable the admin:
    url(r'^admin/', include(admin.site.urls)),
)

```

19 Category URLs

We'll create the URL in `urls.py` as for the post page, only it'll give the slug of the category instead of an ID in the link: `url(r'^myblog/category/(?P<category_id>[0-9]+)$', 'blog.urls.blog_categories')`,

“Finally, let's make the categories click through to the relevant page”


```

{% for post in category_posts %}
<div class="blog-post">
<h2><a href="/myblog/{{ post.pk }}"/>{{ post.title }}</a></h2>
{{ post.author }} on {{ post.pub_date }}
{{ post.excerpt }}
</div>
</for>

```

20 Category template

We'll use something similar to the Index and Post template to create a category page template: `{% for post in category_posts %} <h2>{{ post.title }}</h2> {{ post.author }} on {{ post.pub_date }} % if post.excerpt % <p>{{ post.excerpt }}</p> {% else %} <p>{{ post.post|truncatewords:3 }}</p> {% endif %} <p>Category: {{ post.category }}</p> {% endfor %}`



21 Category clickthrough

Finally, let's make the categories click through to the relevant page by changing the

category display to be: `<p>Category: {{ post.category }}</p>` This can go on the categories, post and index template.

```

return month_list

def month(request, year, month):
    post_list = Post.objects.filter(pub_date_year=year, pub_date_month=month)
    return render_to_response('blog/index.html', dict(sidebar_list=post_list, month_timeline()))

def blog_categories(request, category_id):
    categories = Categories.objects.get(pk=category_id)
    category_posts = Post.objects.filter(category=categories)
    return render_to_response('blog/categories.html', dict(category_posts=category_posts, categories=categories))

class BlogFeed(Feed):
    title = "Blog Feed"
    link = "/"

    def items(self):
        return Post.objects.order_by('-pub_date')

    def item_link(self, post):
        return post.get_absolute_url()

```

22 RSS

Django has a built-in RSS framework. In `blog/urls.py` add: `from django.contrib.syndication.views import Feed class BlogFeed(Feed): title = "Blog Feed" link = "/" def items(self): return Post.objects.order_by("-pub_date") def item_title(self, post): return post.title`

```

return month_list

def month(request, year, month):
    post_list = Post.objects.filter(pub_date_year=year, pub_date_month=month)
    return render_to_response('blog/index.html', dict(sidebar_list=post_list, month_timeline()))

def blog_categories(request, category_id):
    categories = Categories.objects.get(pk=category_id)
    category_posts = Post.objects.filter(category=categories)
    return render_to_response('blog/categories.html', dict(category_posts=category_posts, categories=categories))

class BlogFeed(Feed):
    title = "Blog Feed"
    link = "/"

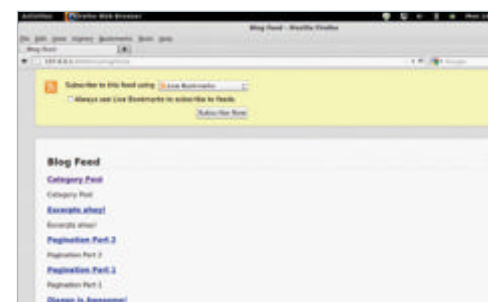
    def items(self):
        return Post.objects.order_by('-pub_date')

    def item_link(self, post):
        return post.get_absolute_url()

```

23 RSS links

We need to define `item_link` for the feed so that the feed items can link to the right place. We have to give the complete URL and the post ID for it work: `def item_link(self, post): link = "http://127.0.0.1:8000/myblog/" + str(post.pk) return link`



24 RSS URLs

The final step is adding the feed URL to `urls.py`: `url(r'^myblog/feed/$', BlogFeed(),)` And now your blog is now fully functional. With a bit more tweaking and theming, you can get it online and blog away!

ID	Severity	Priority	OS	Assignee	Status	Resolution	Summary
30711	crit	medium	All	libreoffice-bugs	REOP	---	[EasyHack]: LibreOffice saves input fields in a non-standard ODF way
30732	enh	high	All	libreoffice-bugs	NEW	---	[EasyHack]: Character formatting not retained in entries of TOC, table lists, etc.
30873	nor	medium	Win	giuseppe.castagno	NEW	---	[EasyHack]: Field variables break with additional formats
31005	maj	high	Win	cedric.bossinat.ooo	NEW	---	[EasyHack]: Table Autoformats does not save (and of course apply) all properties (Writer and Calc)
31022	min	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: "About dialog" wording and layout are suboptimal
31251	enh	medium	All	sebastien-libreoffice	ASSI	---	[EasyHack]: Make the default page look better
31609	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: Support different smart quote settings by language.
31616	enh	medium	All	kohei.yoshida	NEW	---	[EasyHack]: [Chart]: Allow drag and drop for Data labels
31853	nor	medium	All	lemoyne.castle	ASSI	---	[EasyHack]: test bug for testing needinfo.greasemonkey script
32364	enh	medium	Linu	libreoffice-bugs	NEW	---	[EasyHack]: Add a "filter" to Insert > Field > Other > Cross reference tab
32500	enh	medium	All	libreoffice-bugs	REOP	---	[EasyHack]: GTK style doesn't draw some elements via GTK
32506	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: LibreOffice Base - Find and replace
32562	nor	medium	Linu	aruz	ASSI	---	[EasyHack]: Support dbusmenu on Ubuntu or other Linuxen
32664	nor	medium	Win	libreoffice-bugs	NEW	---	[EasyHack]: Keyboard volume keys don't work when LibreOffice is in focus
32781	nor	medium	All	andre.schnabel	REOP	---	[EasyHack]: Extraneous spaces in outline mode.
33831	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: Media.wiki export pictures as files
33960	nor	medium	All	libreoffice-bugs	REOP	---	[EasyHack]: cross reference to a list number dot bug
34324	nor	medium	Linu	prladek	NEW	---	[EasyHack]: Formula icon in quickstarter menu looks inconsistent.
34465	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: get rid of all calls to virtual const SfxPosition* PwZ const SfxPosition, USHORT rWbich)
34591	nor	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: HTML import: table borders set to paragraphs instead of cells
34692	enh	medium	All	tsung1112	NEW	---	[EasyHack]: Print Current Page in Print dialog - Calc, Draw, Impress (Writer is finished)
34965	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: Make Impress slides editable during presentation
35090	nor	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: Subsequenttest failures
35862	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: "Increase font" and "decrease font" could work even if the selection contains differing font sizes
35973	enh	medium	All	programming	ASSI	---	[EasyHack]: Remember the state of the sidebar pane in Impress

Don't be fooled by the name: plenty of 'Easy Hacks' are of a high priority

We're using LibreOffice as our example, but many other projects work in a similar way

Some bugs call for more than just a code fix, with design or translation skills required

There are always more bugs than developers – your help will certainly be welcomed!

Begin a bug hunt in LibreOffice

Ever wondered how to start fixing bugs? Read on...

Advisor

Gareth Halfacree Having spent most of his professional life alternately praising and swearing at Linux, there are few problems that haven't caused Gareth some heartache at one time or another



We all make use of free, libre and open source software. Between the Linux kernel and the GNU packages, it makes up the very heart of every Linux-based operating system around – and the vast majority of the client applications, too, from web browsers like Chromium and Firefox to productivity suites like LibreOffice.

The FLOSS methodology is one of giving back, however. While not every user can be expected to contribute to the community, the more that do the better the project becomes. It's a naturally beneficial arrangement.

We've teamed up with three members of the LibreOffice development community – Michael Meeks, Markus Mohrhard and André Schnabel – to offer a quick-start guide to giving back to your favourite project, regardless of your level of expertise.

Over next few tutorials, we'll detail exactly how to make yourself a valuable member of any open source or free software community, how to find bugs, how to communicate with your new, global team and how to fix problems in a maintainable manner which will have the entire community singing your praises in no time flat.

Resources

A web browser: www.getfirefox.com

A copy of Git: www.git-scm.com

A text editor or IDE: www.eclipse.org

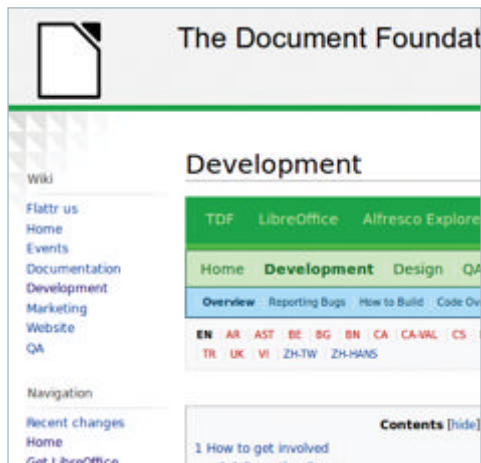
An analytical mind

Passing knowledge of C++

(or another suitable language)

01 Join the community

The first thing to remember: the development community actively wants your help. “Contributors really are welcome, and they should feel that,” explains André Schnabel, a relative newcomer to the community himself. A helping hand is always welcome, so don’t be afraid to get stuck in.



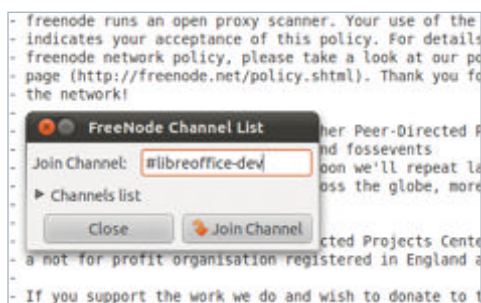
02 Browse the wiki

Before you leap into the codebase, learn about the project you’ve chosen to join. The LibreOffice project has a useful development wiki: <http://j.mp/developerwiki>.

You don’t need to read the whole thing at once, but at least give it a quick run-through.

03 Join the mailing list

Developing as part of a team is all about communication. The LibreOffice community runs several mailing lists, but the one you should be interested in is the development list. Instructions for joining and contributing can be found in the wiki here: <http://j.mp/developerlist>.



04 Join the IRC channel

Mailing lists are a great way of creating a permanent record of communication and changes, but sometimes you need something more immediate. Like many projects, LibreOffice runs an IRC channel for real-time chat between developers. Head to [#libreoffice-dev](http://irc.freenode.net) on irc.freenode.net to join in.

05 Soak up the atmosphere

As friendly as open development communities can be, it’s always worth taking your time to figure out how things work. Rather than making an immediate contribution, hang back for a while and learn the community’s foibles. Do your homework and you’ll slide right into place.

Selection of Open Tasks

Easy Hacks

Improving EasyHacks Bugzilla-Wiki integration

The pages:

- http://wiki.documentfoundation.org/Development/Easy_Hacks_by_Difficulty
- http://wiki.documentfoundation.org/Development/Easy_Hacks_by_required_Skill

are currently manually created because transclusion arguments do not seem to be exported to the feed uris. Manually creating these pages is tiresome and errorprone. It would be great to have a Script to recreate these (and maybe more: `by_Topic`, `by_Component`) pages. The pages should be stored at: <http://cgit.freedesktop.org/libreoffice/contrib> in the end.

It would be even more cool to instead of having to write a huge feed-URL like:

```
<feed url="https://bugs.freedesktop.org/buglist.cgi?bug_status=NEW&bug_status=ASSIGNED&bug_status=..."
```

06 Check out the Easy Hacks

The LibreOffice developers maintain a list of so-called ‘Easy Hacks’ – fixes suitable for newcomers to the community. Many are medium or even high severity, but they make a good starting point and come with the offer of help. The list: <http://j.mp/easyhacks>.



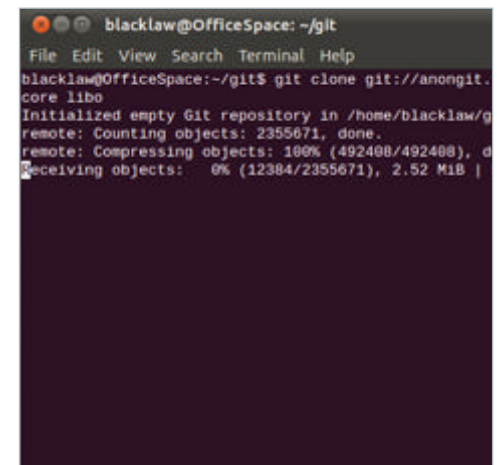
07 Gather your equipment

If you’re already a developer, the chances are good that you’ll have everything you need already to hand. If not, install and familiarise yourself with an IDE and the Git version control system – see the links at the start of this article for suggestions.

“A helping hand is always welcome, so get stuck in”

08 Sign up to Bugzilla

LibreOffice bugs are monitored using an issue-tracking system called ‘Bugzilla’ – and if you want to comment, you’ll need an account. Head over to the Free Desktop bug-tracking system – <http://j.mp/freedesktop> – and sign up. Don’t worry, it’s completely free.



09 Download the codebase

Before you can start hacking on a bug, you’ll need a copy of the codebase. This can be cloned from the Git repository with a simple command. Just make a new directory and **type the following:**

```
git clone git://anongit.freedesktop.org/libreoffice/core/libo
```


“Not all bugs require the entire codebase to be compiled from scratch, but it’s good practice. LibreOffice is a very big project so you’ll need to put aside around two hours”

Preparation

Dependencies

The first step is to make sure that you have all the dependencies needed for your platform to build the product. There are pages that details what they are and how to install them for [Windows](#), [Linux](#), and [MacOS](#).

Disk Space

At minimum, with all the git repositories, and the product build and packaged, you will need anywhere from 7 to 10 GB of disk space, depending on the platform, the build options and the specific autogen options you use, few Gb more if you make a debug build.

Bear in mind that if you want to do more than just building the product once for your own enjoyment, you most likely will want to have more than one working directory (we will see [later](#) how to set them up). A workdir allow you to have a separate build environment without having to duplicate the git history, but it will still cost you 6 to 9 Gb per workdir.

10 Build the project

Not all bugs require the entire codebase to be compiled from scratch, but it’s good practice. LibreOffice is a very big project so you’ll need to put aside around two hours – on an average PC – and read this: <http://j.mp/buildlo>.

11 Tailor to your interests

Although a subsection of the overall bug list, the Easy Hacks are surprisingly vast. Start narrowing down your options according to your interests: if you’re a numbers whiz, look for things in Calc; if you’re more of a designer, check the UI and UX bugs.

Contents [hide]

- 1 Introduction
- 2 List of Easy Hacks by required Skill
 - 2.1 Easy Hacks requiring C++ Skills
 - 2.1.1 Skill Level: Beginner
 - 2.1.2 Skill Level: Interesting
 - 2.1.3 Skill Level: Unknown
 - 2.2 Easy Hacks requiring C++ Debugging Skill
 - 2.2.1 Skill Level: Beginner
 - 2.2.2 Skill Level: Interesting
 - 2.2.3 Skill Level: Unknown
 - 2.3 Easy Hacks requiring Project Infrastructure
 - 2.3.1 Skill Level: Beginner
 - 2.3.2 Skill Level: Interesting
 - 2.3.3 Skill Level: Unknown
 - 2.4 Easy Hacks requiring Java Skills
 - 2.4.1 Skill Level: Beginner
 - 2.4.2 Skill Level: Interesting
 - 2.4.3 Skill Level: Unknown
 - 2.5 Easy Hacks requiring Python Skills
 - 2.5.1 Skill Level: Beginner
 - 2.5.2 Skill Level: Interesting
 - 2.5.3 Skill Level: Unknown

Bug 31616 - EasyHack: [Chart]: Allow dr

Status: NEW

Product: LibreOffice

Component: Chart

Version: unspecified

Platform: All All

Importance: medium enhancement

Assigned To: Kohel Yoshida

QA Contact:

URL:

Whiteboard: EasyHack feature

Keywords:

Depends on:

Blocks:

[Show dependency tree / graph](#)

12 Pick a bug

Don’t try to save the world in a single bound: find a bug that matches your interests and that you think you can reasonably tackle in your free time, and make it yours. Always read the original bug report carefully before committing yourself, though.



13 Duplicate the bug

Before you can fix the bug, you need to make sure it exists in your installation. Check the bug report and make sure you’re running the same version of LibreOffice, then attempt to duplicate the problem. If you can’t, you’ll never know if it’s fixed.

14 Take ownership of the bug

Remember when we emphasised communication? We weren’t kidding. Post a comment to the bug tracker and the mailing list, and tag the entry in the Easy Hack wiki page – if it’s an Easy Hack – to make sure people know not to duplicate effort.

“Remember when we emphasised communication? We weren’t kidding”

```
FreeNode: #libreoffice-dev
IRC Edit Network Discussion View Help
This is the LibreOffice DEVELOPER channel. Please ask USER questions on #libreoffice ONLY. | master is used for LO-3.5 | 3.3.4 and 3.4.4 are out | http://wiki.documentfoundation.org/Development | Use http://libreoffice.pastebin.ca/ for build logs | dev list: libreoffice@lists.freedesktop.org

[14:53] --> You are now talking on #libreoffice-dev
[14:53] --- Topic for #libreoffice-dev is This is the LibreOffice DEVELOPER channel. Please ask USER questions on #libreoffice ONLY. | master is used for LO-3.5 | 3.3.4 and 3.4.4 are out | http://wiki.documentfoundation.org/Development | Use http://libreoffice.pastebin.ca/ for build logs | dev list: libreoffice@lists.freedesktop.org
[14:53] Topic for #libreoffice-dev set by dtardon|dtardon@nat/redhat/x-zxvdgttxsahmbisg at Thu Nov 10 05:59:08 2011
[14:53] [freenode-info] if you're at a conference and other people are having trouble connecting, please mention it to staff: http://freenode.net/faq.shtml#gettinghelp
[14:53] #libreoffice-dev :http://libreoffice.org

blacklaw Hi, guys. I'm going to be working on Bug 31616 - drag and dropping data chart labels - if anyone has any tips.
```

15 Communicate with other developers

Yes, it's another communication step, and this one's just as important. It's the biggest change coders have to make when moving from solo work to team work.

"Get the name of the developer who already evaluated the hack," Schnabel advises. "This is the person who may give more information on the mailing list or via IRC."

While the original developer who declared the hack an easy one is a great source for tips, remember the rest of the community: even other newcomers may have something to offer, so ask questions.



16 Check with the UX advisors

If your chosen bug involves any user interface changes, you'll need to – surprise – communicate with another group. The ux-advise mailing list is where the user experience types hang out, so discuss your bug with them before making any changes: <http://j.mp/uxadvise>.

```
Rainer Bielefeld 2010-11-14 22:53:03 P
[Reproducible] with "LibreOffice 3.3.0 - 13.2.99.2)". Currently several formatting the chart data label dialogs, but there is a label to a new position.

We had <http://www.openoffice.org/issues/problem, may be some work already has been done.

I believe <http://www.openoffice.org/issues/bugzilla/show_bug.cgi?id=31616 should be handled together with this issue.

tbehrens 2010-11-17 09:41:13 PST
Kohei, sounds like a feature request - I'd like to see it in the next version. http://wiki.documentfoundation.org/Development/EasyHacks ?
```

17 Read the bug report again

Always re-read the bug report. Check for new comments and see if there are any hints as to where to start looking in the codebase. "Developers often leave a comment about source files or classes where you can start analysing the issue," Schnabel points out.

18 Stay up to date

While you're analysing your chosen bug, keep abreast of what other developers are doing and post updates on your progress with the bug. "Markus gave me some warning that he is reviewing some code in the same area just to prevent merging troubles," Schnabel remembers.

19 Pinpoint the affected code

This step can eat up time, but it's absolutely critical: using any hints from the bug report and/or by communicating with the bug's filer, find the precise area of the codebase that you should be working on and make sure you understand it well.

20 Examine the files

While it's important to have focus, you need a broader awareness of how things work too. Look at the files affected by your bug, plus associated files, and get an overall feel for the coding style and conventions. Don't worry about specifics for now, though.



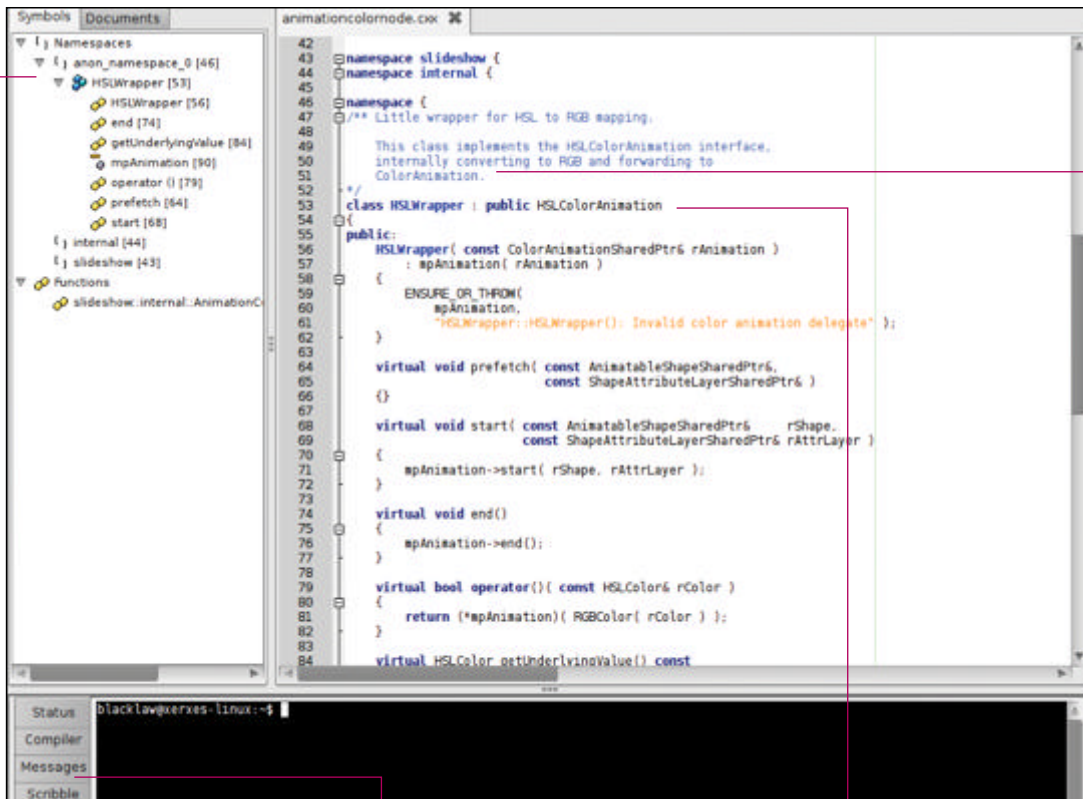
21 Learn the coding style

Looking at the files is a great first step, but before you commit code you need to be aware of the project's conventions. For LibreOffice, there are two documents: the Coding Standards and the Code Conventions – at <http://j.mp/codingstandards> and <http://j.mp/codeconventions>.

As you may have noticed, being an effective contributor to a free, libre or open source project requires preparation – but it is well worth the effort. Without volunteers willing to root out bugs and write new features, FOSS software simply wouldn't exist as we know it today.

By offering your time – whether you're a seasoned coder or a newcomer wanting to learn more – you'll help keep your favourite projects alive, make the software better for all and earn the thanks and respect of the existing community. It's a truly rewarding endeavour, and the fact that you've read through this means you're well on your way. In the next article, we'll pick an actual bug and work through a real-life fix, from diagnosing the problem to submitting the patch.

Masterclasses



Remember that changes you make to the source – however small – will impact other areas of the package

Keep notes as you're going along – if your IDE has a notepad function, be sure to make good use of it

Make sure you understand and abide by code style and naming conventions for your chosen project

If you need an explanation of what you've done, put it in separate documentation – not the comments

Find and fix known bugs in LibreOffice

Our journey into fixing bugs in open source software continues

Advisor Gareth Halfacree



A free software enthusiast should always be willing to give back, so Gareth's learning to hunt and fix bugs.



It's all too easy to take from open source projects; but for such efforts to survive, people need to give back. We've just looked at

what you need to do in order to become a valuable member of an open source project – in this case, the LibreOffice project as run by The Document Foundation.

We focused the entire first part on the communication between the contributor and the community as a whole. It's a key lesson to learn: you can be the best coder out there, but if you 'go rogue' and don't communicate your work effectively you can be more of a hindrance than a help.

While some open source contributors do so as part of their day job, others are not as lucky – but don't let that put you off. Without volunteers helping with development, many projects would stagnate and die – including, almost certainly, packages you rely on every day.

Here, we'll be taking a look at an actual fix – André Schnabel's work on a display bug in LibreOffice Calc – in order to highlight how to actually go about finding a problem, resolving the issue and submitting a fix. While some of the steps here are specific to LibreOffice, the overall effort is the same for any open source project. Follow our guide and you'll soon become an asset to any development community.

Resources

A web browser: www.getfirefox.com

A copy of Git: www.git-scm.com

A text editor or IDE: www.eclipse.org

An analytical mind

Passing knowledge of C++

(or another suitable language)

Masterclasses

```
FreeNode: #libreoffice-dev
IRC Edit Network Discussion View Help
▼ This is the LibreOffice DEVELOPER channel. Please ask USER questions on #libreoffice ONLY. | master is used for LO-3.5 | 3.3.4 and 3.4.4 are out | http://wiki.documentfoundation.org/Development | Use http://libreoffice.pastebin.ca/ for build logs | dev list: libreoffice@lists.freedesktop.org

[14:53] --> You are now talking on #libreoffice-dev
[14:53] --- Topic for #libreoffice-dev is This is the LibreOffice DEVELOPER channel. Please ask USER questions on #libreoffice ONLY. | master is used for LO-3.5 | 3.3.4 and 3.4.4 are out | http://wiki.documentfoundation.org/Development | Use http://libreoffice.pastebin.ca/ for build logs | dev list: libreoffice@lists.freedesktop.org
[14:53] Topic for #libreoffice-dev set by dtardon!dtardon@nat/redhat/x-zxvdgttxsahmbisg at Thu Nov 10 05:59:08 2011
[14:53] [freenode-info] if you're at a conference and other people are having trouble connecting, please mention it to staff: http://freenode.net/faq.shtml#gettinghelp
[14:53] #libreoffice-dev :http://libreoffice.org
```

10 Communicate

Yes, we're back to this old chestnut again, but it's a mantra worth repeating. Show your prototypes – both in mock-up form and as working code – to other developers, as the feedback you receive at this stage will be crucial to the quality of the fix.

11 Implement fix

With the feedback you receive on your prototypes, pick the best method of fixing the flaw or implementing the feature and make it happen. If substantial time has passed, refresh your local source tree – via another git pull – in case other changes could affect your work.

12 Document changes

Before submitting your fix, thoroughly document every change you made. Although a fix is always welcome, a fix which comes with comprehensive documentation is doubly welcome – and avoids problems in the future when work is done in the same area.

13 Test fix

Once you've implemented your fix, it's time to test. Compile the codebase by following the guidelines offered by The Document Foundation, and make sure that your fix both works and hasn't thrown any related functionality off at the same time.

“Before submitting a fix, thoroughly document every change you made”

14 Submit fix

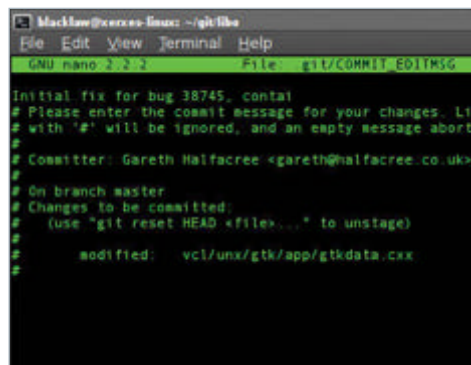
To submit your fix for review, commit the changes locally:

```
git commit -a
```

Once done, you'll need to create a patch file for submission to the mailing list:

```
git format-patch HEAD~1
```

This patch file should then be submitted to the developer mailing list with a summary of the changes you've made.



```
MacLlaw@xxxxx-linuz: ~/git/libo
File Edit View Terminal Help
GNU nano 2.2.2 File: git/COMMIT_EDITMSG
Initial fix for bug 38745, contal
# Please enter the commit message for your changes. Lines
# with '#' will be ignored, and an empty message aborts
#
#
# Committer: Gareth Halfacree <gareth@halfacree.co.uk>
#
# On branch master
# Changes to be committed:
#   (use "git reset HEAD <file>..." to unstage)
#
#       modified:   vcl/unx/gtk/app/gtkdata.cxx
#
```

15 Document fix

Yes, it's time for more documentation. Make sure you include the changes you made as a part of your fix, but also try to write as much supporting documentation as possible.

If your fix makes changes in the user interface, be sure to highlight that too. “My change triggered a bug in the column header display in Calc,” Schnabel remembers. While this secondary bug was soon fixed by developers working on Calc itself, it was thanks to Schnabel's comprehensive documentation that it could be discovered and rectified so quickly.



“When you’ve received feedback from the reviewer, alter your code accordingly”

16 Wait for review
 “Once you send the patch to the mailing list, it will be reviewed by experienced developers,” Schnabel explains, an important step in developing a patch. “For my fix, the review took about one week. Our developers always are busy, so be patient.”

17 Read feedback
 The feedback received from the mailing list will be vital. “I actually had two or three errors in my fix,” Schnabel remembers. “The reviewer very kindly pointed me to those errors, and offered to fix them.” Even if your code is perfect, it may need editing for style or clarity.

18 Make changes as required
 When you’ve received feedback from the reviewer, alter your code accordingly. Depending on the changes, your next patch may need another review process – or it could be accepted into the codebase right away as the ‘official’ fix. Either way, refresh your codebase with git pull.

Selection of Open Tasks

Easy Hacks

Improving EasyHacks Bugzilla-Wiki integration

The pages:

- http://wiki.documentfoundation.org/Development/Easy_Hacks_by_Difficulty
- http://wiki.documentfoundation.org/Development/Easy_Hacks_by_required_Skill

are currently manually created because transclusion arguments do not seem to be expanded in the feed urls. Manually creating these pages is tiresome and errorprone. It would be great to have a Script to recreate these (and maybe more: by_Topic, by_Component) pages. The scri should be stored at: <http://cgit.freedesktop.org/libreoffice/contrib> in the end.

It would be even more cool to instead of having to write a huge feed-URL like:

```
<feed url="https://bugs.freedesktop.org/buglist.cgi?bug_status=NEW&bug_status=ASSIGNED&bug_status=
```

19 Pick another Easy Hack
 For LibreOffice developers, now’s the time to get stuck into another Easy Hack from the list. Once you’ve got a few of those under your belt, you’ll be ready to take on something meatier – and will have made a major contribution to the project’s success.

20 Spread your wings
 As an officially recognised contributor to an open source project, the world is now your oyster. Either stick with your chosen project and help make it the best it can be, or take a look around for other concerns that might need your help – and bask in your success.

André Schnabel 2011-07-03 11:35:43 PDT

patch is currently in review. Implementation follows this spec:
http://wiki.documentfoundation.org/User:Andreschnabel/Spec_Calc_grid_lines_on_colored_background

Katarina Machalkova 2011-07-06 02:28:56 PDT

Just pushed this to master. For the next major release (3.5)

Josef.Latt 2011-07-31 03:56:42 PDT

Why can the patch not be integrated in the next bugfixrelease?

André Schnabel 2011-07-31 23:14:39 PDT

This should not be included in 3.4.x, as it is not just a bugfix - functionality including new UI strings and help texts. We would break localizations if we would include all this.

Contributing a fix to a large-scale open source project – even one as minor as a single-line change to the source code – can be a lengthy procedure. Between prototyping, testing, communicating, documenting and waiting for the review process to complete, it’s certainly not a weekend job.

Neither does it need to take over your life, however: if you can only spare a little time, you’ll still be welcomed into the community. You can either take on lower-priority projects or collaborate on larger tasks with other part-time coders.

Schnabel’s example bug took, he estimates, around four days to fix, of which only 15 per cent was spent coding. The rest of the time was taken up with learning, experimentation and discussion – as well as the all-important documentation. Far from being an afterthought, these parts of the process are vital to the success of any large-scale open source project.

“Don’t worry, if you do make some mistakes in your first bug fix,” Schnabel advises. “This is what the Easy Hacks are there for.”

For a more in-depth look at Schnabel’s fix, check out his user page on The Document Foundation’s wiki:
<http://j.mp/logrids>



Become an open-source mentor in LibreOffice

As we wrap up our series, we take a look at how to move from bug-hunting beginner to seasoned pro

Advisor Gareth Halfacree



A free software enthusiast should always be willing to give back, so Gareth's learning to hunt and fix bugs

Timing is everything

Understand the development cycle to better time your efforts

What you contribute to a project is important, but when you contribute it can be almost as vital. Every project has its own development cycle, and understanding what phase you're in is a key part of being an effective contributor.

"You need to adapt to the bug fix/feature development cycle," explains LibreOffice developer Markus Mohrhard. "Between the first beta and the point-one release, focus on bug fixing with special attention to your own bugs."

By focusing your efforts on bug fixing at a time when, thanks to the release of the beta, a massive influx of bug reports arrives, it's possible to iron out as many issues as possible before a final release.

Once the beta portion of the development cycle has finished, it's possible to go back to other tasks – including mentoring and code cleanup – without leaving dozens of bug reports unanswered.



If you've followed our last couple of tutorials, you should now be an accomplished contributor to the LibreOffice project. You'll likely have

personally resolved several of the entries on the 'Easy Hacks' list – LibreOffice's collection of non-showstopping issues that nevertheless require resolution, and by far the best place for a newcomer to the community to focus their efforts. But despite having a list of accepted patches to your name, there's still plenty to do, and as your skill increases it's time to think about moving on to the next level: mentorship.

"Becoming a mentor is something that is a non-step," explains Michael Meeks, founding member of The Document Foundation and passionate LibreOffice contributor. "Everyone should be a mentor of someone less skilled than them – there is no need to get a badge first.

"The more people you mentor and develop, the faster both the community grows and your position in it as a helpful person," Meeks adds by way of incentive.

The idea of mentoring others when you're a relative newcomer to the community yourself can seem daunting, and even unwelcome; after all, nobody likes the guy that comes storming in to an

established community to tell them they're doing it all wrong.

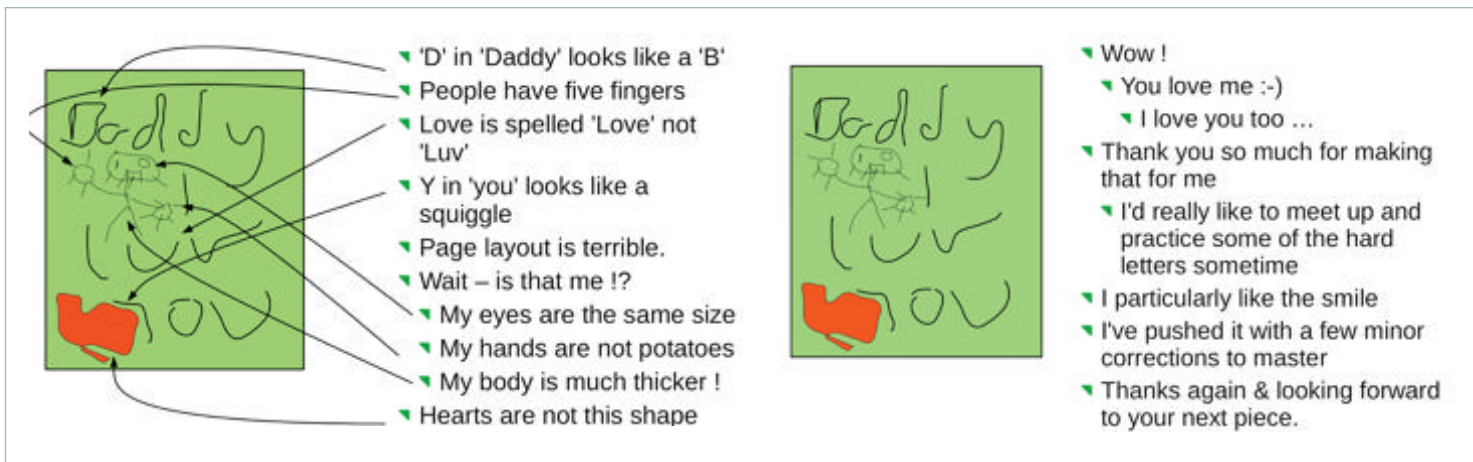
However, this is misunderstanding what mentoring is all about: taking what your mentors have taught you and passing that down to people in the position that you were a few months ago.

Attitude is key

If things have gone according to plan, your entry into the development community was smooth. You were made to feel welcome, and where mistakes were – inevitably – made, you were corrected without being made to feel as though you had done wrong.

That's no accident, LibreOffice developer Markus Mohrhard explains, but a key ethos that community members must embrace. "Be friendly to beginners," Mohrhard warns, "even if they did something wrong. That should be always the case, but it is much more important in contact with beginners."

In this, you're at an advantage: it's not that long ago that you yourself were a beginner. Remember your welcome to the community, and the patience that met your first stumbling questions; remember the help you were given during your time fixing the bug you chose from the Easy Hacks list; remember the way your



■ Michael Meeks' demonstration of how not to criticise a newcomer. ■ Instead, he advises a more nurturing approach for mentors to take.

patch was received and, quite likely, modified by a more experienced community member before, hopefully, being merged into the main source code as an official fix for a vexatious flaw in the software.

There are key points to remember when mentoring, Meeks explains: a nurturing and encouraging attitude is a must, and where criticism is required it should always be delivered in a way that doesn't make the contributor feel small – using, Meeks suggests, a “praise/criticism/praise sandwich” – in order for them to feel their efforts are appreciated by the community.

As with our previous guides in this series, communication is key: if you're mentoring a new user, be sure to respond to them promptly; if you can't, send them a quick email explaining why and – if necessary – ask another developer to take over. Leaving a newcomer hanging in limbo can make them feel as though their input is considered worthless, when nothing could be farther from the truth.

There are key attitudes to avoid, too: don't participate in extensive public criticism of newcomers, no matter what their own attitude or expertise may be. You were a beginner yourself not too long ago, and while it can seem tempting to score points off a new contributor in order to demonstrate your own increasing knowledge, it will simply scare the newcomer off – and win you no friends in the existing community.

“The better and kinder you are as a mentor, the better your ‘mentees’ will mentor others,” Meeks explains. “Bad responses drive people away. Bad responses encourage other list participants to respond badly too,” encouraging what Meeks describes as a “cycle of badness” that can quickly discourage people from contributing and eventually lead to the inexorable death of an open source project.

```

FreeNode: #libreoffice-dev
IRC Edit Network Discussion View Help
▼ This is the LibreOffice DEVELOPER channel. Please ask USER questions on #libreoffice ONLY. | master is used for LO-3.5 | 3.3.4 and 3.4.4 are out | http://wiki.documentfoundation.org/Development | Use http://libreoffice.pastebin.ca/ for build logs | dev list: libreoffice@lists.freedesktop.org

[14:53] --> You are now talking on #libreoffice-dev
[14:53] --- Topic for #libreoffice-dev is This is the LibreOffice DEVELOPER channel. Please ask USER questions on #libreoffice ONLY. | master is used for LO-3.5 | 3.3.4 and 3.4.4 are out | http://wiki.documentfoundation.org/Development | Use http://libreoffice.pastebin.ca/ for build logs | dev list: libreoffice@lists.freedesktop.org
[14:53] Topic for #libreoffice-dev set by dtardon!dtardon@nat/redhat/x-zxvdgttxsahmbisg at Thu Nov 10 05:59:08 2011
[14:53] [freenode-info] if you're at a conference and other people are having trouble connecting, please mention it to staff: http://freenode.net/faq.shtml#gettinghelp
[14:53] #libreoffice-dev :http://libreoffice.org
  
```

■ Always stay in close communication with other developers

“If you're mentoring a new user, be sure to respond to them promptly”

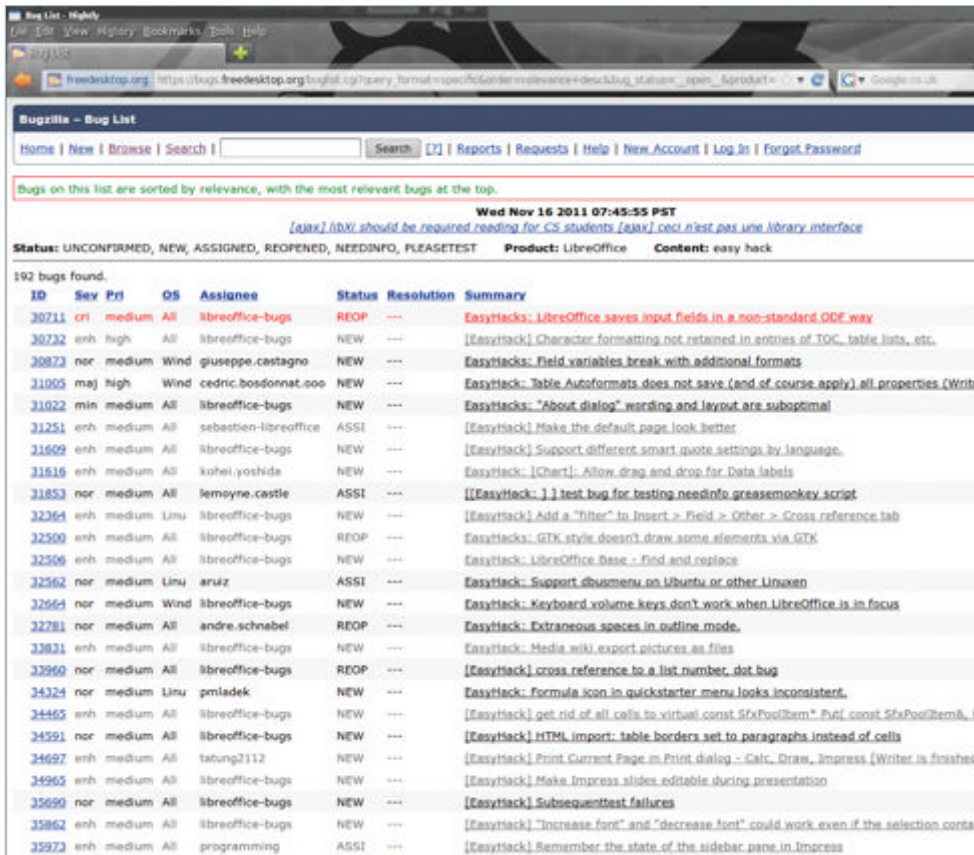
The praise/criticism/praise sandwich Letting the newcomers down gently

If you choose to become a mentor to new developers who are not quite as far along the learning curve as yourself, it's important to behave in an encouraging way. As when you first started, each newcomer is simply eager to help; being too blunt when discussing a problem with a submitted patch – or, worse, denigrating them in public as a way of scoring points – will likely discourage them from ever contributing to the project, losing your team a valuable ally.

One key technique for communicating criticism in a helpful manner is what LibreOffice founder Michael Meeks describes as the

“praise/criticism/praise sandwich”. As the name suggests, it's a way of acknowledging the effort and time an individual has placed into a patch while also covering any changes that need to be made.

Don't dwell on the negatives; if a coder has missed something obvious, start by telling them something they did right. Thank them for their time, praise their efforts, and only then gently address whatever issue needs resolving for the patch to be viable before again finding something to praise; never end a communication on a negative.



ID	Sex	Pri	OS	Assignee	Status	Resolution	Summary
302711	cri	medium	All	libreoffice-bugs	REOP	---	[EasyHack] LibreOffice saves input fields in a non-standard ODF way
302732	enh	high	All	libreoffice-bugs	NEW	---	[EasyHack] Character formatting not retained in entries of TOC, table lists, etc.
308723	nor	medium	Wind	giuseppe.castagno	NEW	---	[EasyHack] Field variables break with additional formats
31005	maj	high	Wind	cedric.bosdonnat.ooo	NEW	---	[EasyHack] Table Autoformats does not save (and of course apply) all properties (Write
31022	min	medium	All	libreoffice-bugs	NEW	---	[EasyHack] "About dialog" wording and layout are suboptimal
31251	enh	medium	All	sebastian-libreoffice	ASSI	---	[EasyHack] Make the default page look better
31609	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack] Support different smart quote settings by language.
31818	enh	medium	All	kohel.yoshida	NEW	---	[EasyHack] [Chart]: Allow drag and drop for Data labels
31853	nor	medium	All	lemoyne.castle	ASSI	---	[EasyHack] [] test bug for testing needinfo greasemonkey script
32364	enh	medium	Linu	libreoffice-bugs	NEW	---	[EasyHack] Add a "filter" to Insert > Field > Other > Cross reference tab
32500	enh	medium	All	libreoffice-bugs	REOP	---	[EasyHack]: GTK style doesn't draw some elements via GTK
32506	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: LibreOffice Base - Find and replace
32562	nor	medium	Linu	aruz	ASSI	---	[EasyHack]: Support dsh menu on Ubuntu or other Linuxes
32664	nor	medium	Wind	libreoffice-bugs	NEW	---	[EasyHack]: Keyboard volume keys don't work when LibreOffice is in focus
32781	nor	medium	All	andre.schnabel	REOP	---	[EasyHack]: Extraneous spaces in outline mode.
33821	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack]: Media miki export pictures as files
33960	nor	medium	All	libreoffice-bugs	REOP	---	[EasyHack] cross reference to a list number. dot bug
34324	nor	medium	Linu	gmladek	NEW	---	[EasyHack]: Formule icon in quickstarter menu looks inconsistent.
34465	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack] get rid of all calls to virtual const SfxPoolItem* PwF const SfxPoolItem&, U
34591	nor	medium	All	libreoffice-bugs	NEW	---	[EasyHack] HTML import: table borders set to paragraphs instead of cells
34697	enh	medium	All	stahng2112	NEW	---	[EasyHack] Print Current Page in Print dialog - Calc, Draw, Impress (Writer is finished)
34965	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack] Make Impress slides editable during presentation
35090	nor	medium	All	libreoffice-bugs	NEW	---	[EasyHack] Subsequent failures
35862	enh	medium	All	libreoffice-bugs	NEW	---	[EasyHack] "Increase font" and "decrease font" could work even if the selection contain
35973	enh	medium	All	programming	ASSI	---	[EasyHack] Remember the state of the sidebar pane in Impress

“Split your time between mentoring and seeking more knowledge”

■ Keep abreast of the latest bugs, and offer to mentor new contributors.

you first started on your journey to becoming a valued contributor. Therefore, go back over your communications with the community. Check the emails you sent to the mailing list, have a look at chat transcripts from the IRC server. Look for the questions you asked, and remember the answer you were given. These answers are valuable; no matter how basic the question may have been, it's a given that a new contributor to the project will soon ask the exact same things. If you're ready, you can be there to help and further cement your reputation as a truly valued member of the community.

Every question you field, and every answer you offer, is freeing people further up the curve to be addressing more complicated issues; issues that you yourself can learn from in order to progress your own knowledge and, in turn, that of those project members you choose to mentor.

Climbing the curve

As Meeks explains, mentorship is a 'non-step,' something that everyone can do for those who are not as far along the learning curve – or simply on a tangential path – as themselves. While becoming a mentor is a rewarding and useful undertaking for a community member, it can also feel restricting; doubly so if you find yourself always answering the same questions over and over again.

The solution, therefore, is to work your way up the curve; find new challenges, ask new questions and continue to work at the problems faced by your chosen project. That's not to say you should stop answering queries to which you know the answer; but rather that you should split your time between mentoring newcomers and seeking more knowledge yourself.

How best to go about this? Mohrhard claims it's simple, and goes back to a recurring theme throughout our tutorials: communication. "Keep in contact with the other developers, and especially the devs that work in the same area as yourself," he advises. "Take part in discussions on the mailing list or developer chat, and keep in contact with the UX [user experience] guys; I think that's an extremely

Not just another bug-hunt

A more focused beta-testing approach can pay dividends

One of the principal problem-finding techniques used in the development of LibreOffice is the Virtual Bug-Hunt, a two-day event which typically takes place in the beta portion of the development cycle.

The concept is simple: for two days, volunteers – both developers and end-users – come together in a virtual environment to rigorously test the beta software, both by simply playing around with new and existing features and by addressing test cases assigned by the development teams.

Although such a focused beta-testing process is relatively uncommon in the open source world, it's a technique that The Document Foundation is finding highly successful: a Bug-Hunt session in December saw an extra 150 people join the development IRC channel and a total of 70 bugs filed for review; most of which came from Gustavo Pacheco, for which he was named the official LibreOffice Bug-Hunt Hero for the session.

But I'm a beginner myself

It's easy to dismiss this part of the series. After all, the chances are that up until recently you'd never even thought about contributing to a major open source project; what could you possibly have to offer as a mentor?

There's plenty you could offer, in fact. Think back to your participation in open-source projects. Has there ever been a time when you feel you've learned something, either about the codebase itself or the desires and attitudes of the community working upon it? If the answer is yes, then that's what you have to offer: you're further along a never-ending learning curve than those who have started after you, and that experience – however scant it may seem – is an incredibly valuable thing.

There's no getting away from the fact that, in any established open source project, there will be contributors much farther along that learning curve than you. Rather than a failing, however, this represents an opportunity: if those long-serving members spend their time fielding the same questions from newcomers, their time is not being well spent. Instead, volunteer your own help: it's almost certain that any given newcomer to the project will spend much of his or her time asking the same questions that you asked when

Reviewing patches

No matter where on the curve you are, you have something to offer

One of the key considerations for keeping a new contributor happy is instant gratification: a submitted patch should be reviewed as quickly as possible, to avoid the contributor becoming disillusioned.

While each patch will have an official reviewer, if you're not the named party it's still possible to help. If you're becoming familiar with a particular section of code, keep checking the mailing list and the bug reporting system for new patches.

Each time you come across a patch that piques your interest, download and test it out on your own system.

Even if you're not in charge of the patch, it's a useful contribution to make; if the patch works, you might be able to validate the contributor's efforts sooner than the official patch watcher; if it doesn't, you may be able to identify why and offer a resolution that will see the patch merged into the codebase in record time.

```

André Schnabel 2011-07-03 11:35:43 PDT
patch is currently in review. Implementation follows th
http://wiki.documentfoundation.org/User:Andreschnabel
/Spec_Calc_grid_lines_on_colored_background

Katarina Machalkova 2011-07-06 02:28:56 PDT
Just pushed this to master. For the next major release

Josef.Latt 2011-07-31 03:56:42 PDT
Why can the patch not be integrated in the next bugfixr

André Schnabel 2011-07-31 23:14:39 PDT
This should not be included in 3.4.x, as it is not just
functionality including new UI strings and help texts.
localizations if we would include all this.
    
```

important point for finding out about possible new features and the need to rework dialogues.”

Specialising can also help rapidly increase both your knowledge and your utility to the project, Mohrhard claims. By focusing your efforts on one particular part of what is, after all, an incredibly diverse and complex codebase, you can reach the deeper understanding that is required to tackle more complicated bugs.

When choosing your focus, try to pick something that gels with your particular area of interest or expertise: “As I used to give user support and to do translations, I focused on UI [user interface] related stuff,” explains relatively new LibreOffice community member André Schnabel. “As Markus [Mohrhard] is studying mathematics, his focus is more on the core implementation in Calc.”

Cleaning up

As your knowledge of the codebase grows, there's another task to consider adopting; and it's one that beginners simply can't tackle on their own: code

cleanup. “If you just want to ‘improve some code,’ that's perfect,” explains Schnabel. “We have a lot of things to clean up.”

Any complex project finds itself in a similar state, and LibreOffice is no exception: its organic growth through the years, coupled with its origins in the OpenOffice project and the sometimes difficult procedure required to submit code patches to corporate owner Oracle, mean that portions of the code require significant cleanup to make future maintenance easier. Other portions are entirely obsolete, having been bypassed in later patches without ever having been removed from the source tree.

When LibreOffice split from OpenOffice under the auspices of The Document Foundation, one of the first tasks the project had to undertake was code cleanup. While much work has been done in this regard, there's still plenty to do; and the task is just as vital as that of fixing reported bugs, as without clean code future maintenance of the codebase becomes near impossible.

“Always try to clean up code while fixing bugs or developing new features,” exhorts Mohrhard. “This is difficult as a beginner, but you get used to it if you understand the code better.”

It's important to reiterate that LibreOffice is no exception here; any sufficiently complex software project will have its own issues with code cleanliness, and it's a problem that needs to be kept under control. If you're using the lessons from this series to contribute to a different project, be sure to seek out code guidelines; if your project doesn't have any, suggest collaborating with other developers to create a guideline that can be used as the benchmark for all future code corrections, to ensure a coherent approach to coding that'll be readable to project newcomers in the future.

Keep on coding

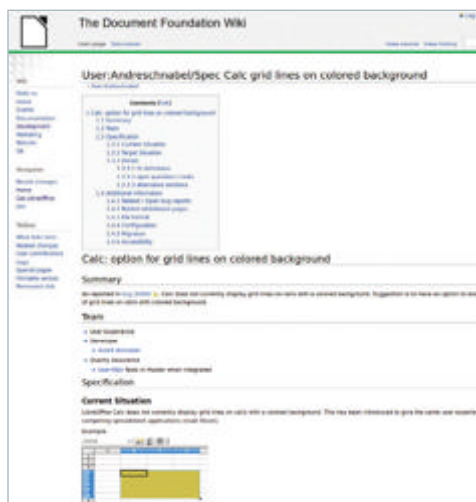
Finally, at some point you'll need to address a very real problem: burnout. The chances are that, like the vast majority of contributors, you're working on an open source project out of the goodness of your heart. While some are lucky enough to work for a company like Novell, SUSE, Red Hat or IBM that pays full-time developers to make open source projects better, you will almost certainly be doing so in addition to earning a crust with a day job.

As a result, it's all too easy for what started as a fun hobby to become overwhelming; as your knowledge increases, so too do the demands on your time. It's important to recognise when this is happening, and to step back.

Don't worry about losing face in the community; the other project contributors want you to be at your best, and if you're finding yourself stressed and harassed that simply isn't going to happen. Take a break, spread your workload across other volunteers and come back when you're feeling refreshed.

Don't let temporary setbacks put you off contributing. As a newcomer, if someone ignores the guidelines on encouragement it can be, as you'd expect, discouraging; likewise as a more seasoned member of the community, there will be times when things seem bleak. Always work to resolve issues, and resist the urge to throw in the towel. Remember that your input into the project is appreciated, both explicitly by other community members who see the difference you make and implicitly by the many users of the software who benefit from changes and fixes that you yourself have written, tested, implemented, reviewed or mentored.

Becoming an open source contributor is an infinitely rewarding task, and your efforts are never in vain.



■ Think about helping with documentation tasks, too.

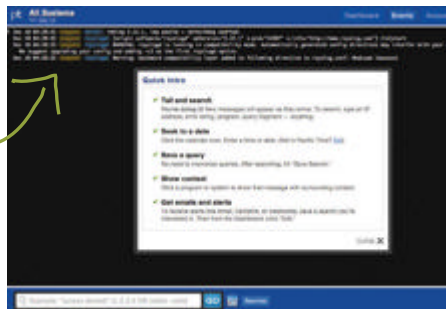
“Any complex software project will have its own issues with code cleanliness”

Developer Gu

Essential tools for coders

152 System logging in the cloud with Papertrail

Access your system logs from anywhere in the world



156 Better PHP on the command line


Build a flexible micro-framework to tame PHP command-line programs

160 Server-side JavaScript with Node.js

Learn how to power applications server-side with JavaScript

164 Create an HTML photo editor

Use the <canvas> element to load and manipulate images



168 Cross-platform mobile app development

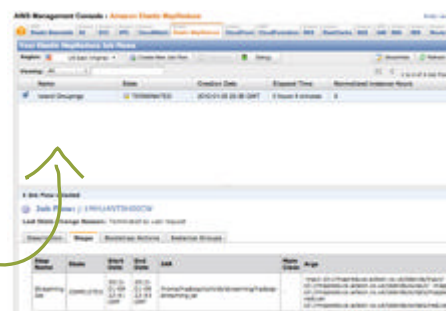
Design apps that will work across multiple platforms

172 Cross-platform mobile web app development

Discover the best way to make the mobile web work for you

176 Using Hadoop tasks on Amazon's Elastic MapReduce

Accomplish huge tasks using the power of Hadoop and command-line tools




180 Develop on your Android smartphone with AIDE

Create your own Hello World app to run on Android

184 Building LibreOffice extensions

Extend the functionality of LibreOffice with these top tips



“The ability to reach out to a large segment of the market is very important, especially if you are a new developer”





ides



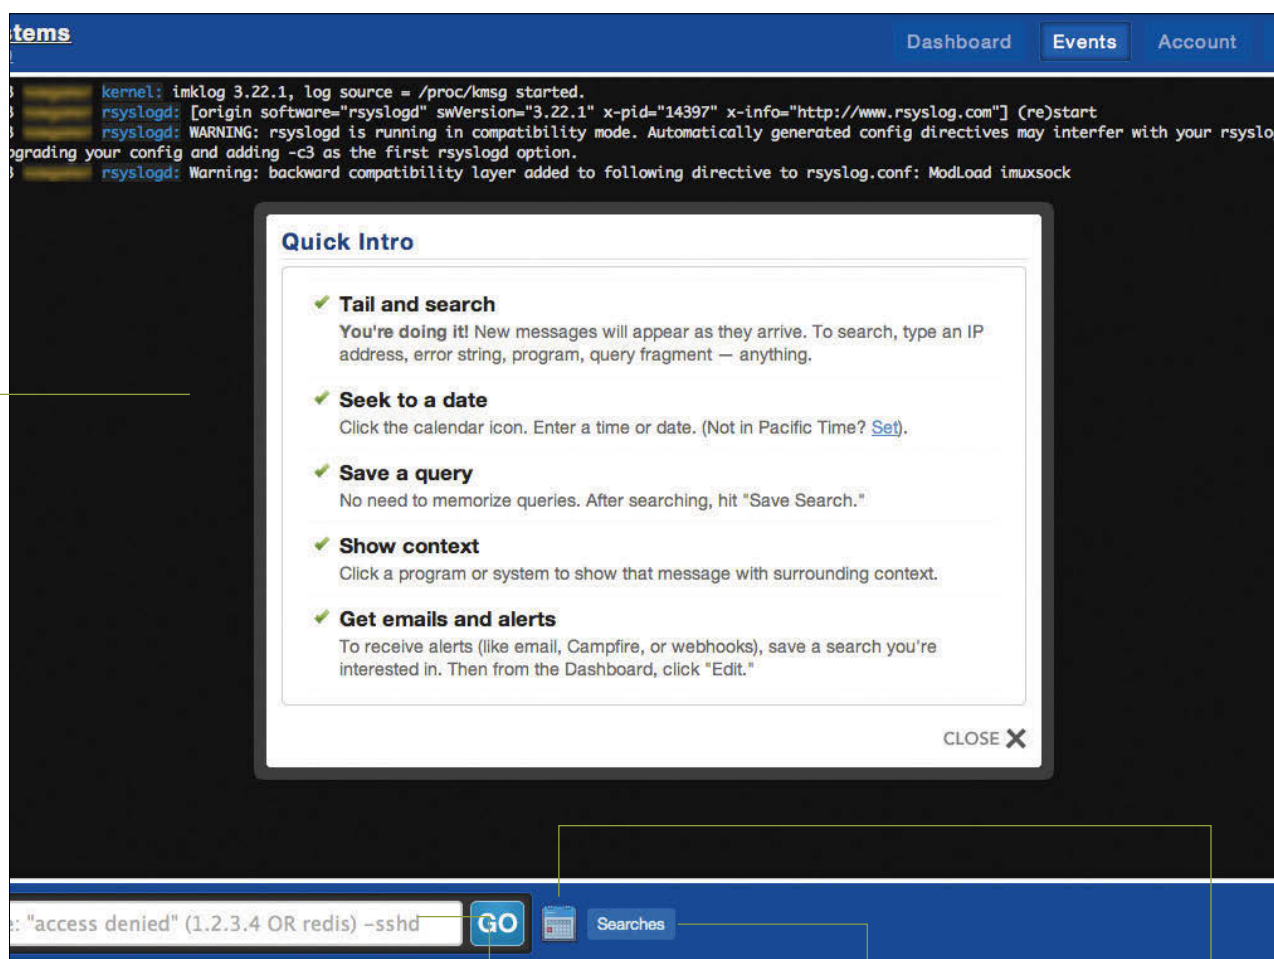
System logging in the cloud with Papertrail

Advisor

Andy Leon is a web developer with Java and PHP as his languages of choice, and a fair amount of server-wrangling as a secondary skill



Get all your system logs in one place, accessible anywhere in the world via the web in easy steps...



Log messages appear in this main window in real-time

Messages can be filtered easily by typing into the search box. Common search syntaxes are supported

Common searches can be saved and actions such as email alerts or web hooks set up for them

Restrict the display to a certain date range by clicking on the calendar icon

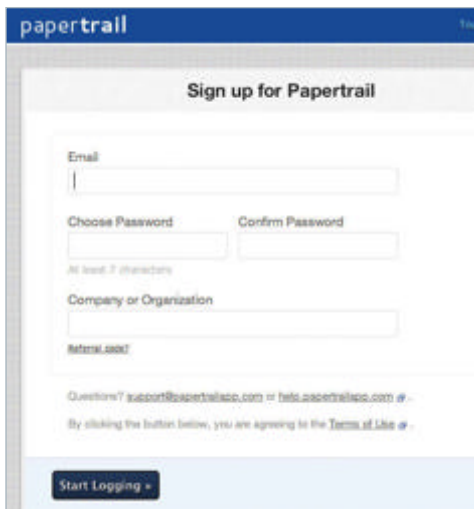


Nobody will argue that logging is not important, especially when dealing with multiple systems, but it can fast become a very unwanted chore.

Dealing with a multitude of files on the command-line is cumbersome and aggregating logs from several machines can soon start to seem chaotic.

Papertrail aims to alleviate these problems and a lot more besides, by providing cloud-based log file management. Any logging that takes advantage of the standard system logger can be configured in seconds and Papertrail's own daemon can take care of anything else. Once data is being collected, it can be searched directly from a web-based

interface, by anybody you care to create an account for. Developers no longer need to SSH into servers in order to query log files and with web hooks and email notifications, everyone can be kept informed of any potential problems. It's easy to get started, there's a free tier available and your regular text-file logging can continue uninterrupted.

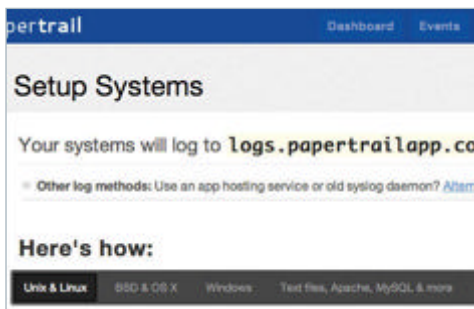


01 Creating an account
Signing up to the Papertrail service just requires an email address, the name of your organisation and for you to choose a password for your account. Head over to www.papertrailapp.com to get started.

You can add new users from the Account tab, once set up.

Plan
» Free
» 5 GB first month
» 100 MB/month
» 24 hours search
» 7 days archive
» Unlimited systems
» Unlimited users

02 Service plans
Papertrail provides a free-to-user plan, which is restricted to 100MB of logging a month (5GB in the first month, though), 24 hours of searchable logs and seven days' worth of archives. There's no restrictions on the number of systems that can log data or the number of users who can access it.



“Papertrail works by your existing system logger forwarding messages”

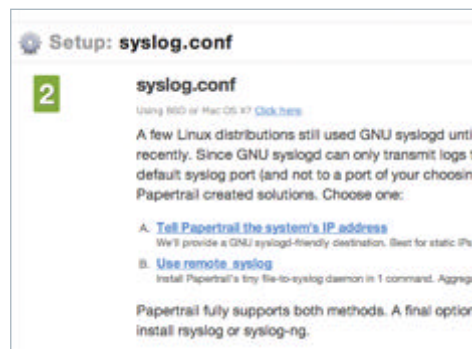
03 Logging address
With your account created, you'll be given an address to use for logging. Papertrail works by your existing system logger forwarding messages onto logs.papertrailapp.com. Different accounts use different port numbers for their logging – you'll be able to see yours here at the end of the address.

```
$ ls -d /etc/*syslog*
/etc/syslog.conf

$ cat /etc/issue
CentOS release 5.6 (Final)
Kernel \r on an \m

$
```

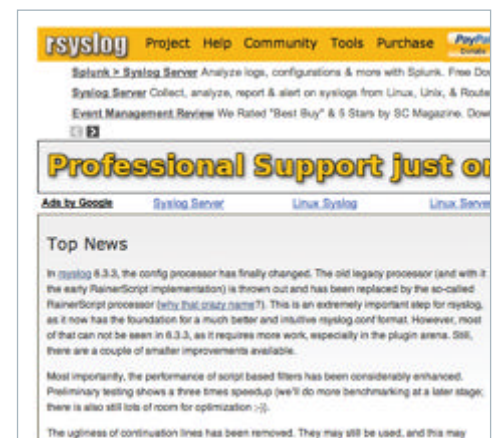
04 Our current logger
To configure this forwarding of log messages, we first need to identify which logger our system is running. We can do by running `ls -d /etc/*syslog*` on the command line to see what config files are installed. If there are multiple loggers installed, check for what's running using `ps ax | grep syslog`.



05 Syslog doesn't cut it
Here we can see this system is running the older 'syslogd' service (still common on CentOS and some others). While this logger is capable of forwarding messages between systems, it will only do so on a default port. Papertrail has workarounds if you wish to keep syslogd, but it would be better to upgrade the logger.



06 Finding a better logger
There are two other major syslog-compatible loggers: syslog-ng and rsyslog. Syslog-ng was the first of these to be introduced, back in 1998. It incorporates many extensions to the original BSD syslog protocol and overall is very feature-rich. The configuration file suffers as a result, though, becoming a little over-complex.



07 The new wave
For our logger, though, we're going to opt for rsyslog. This program has been gaining popularity and is now the default for distributions such as Fedora, openSUSE and Debian. It has also been adopted in RHEL 5 and later, meaning it will probably become default in a future version of CentOS as well.

08 Installing rsyslog
Even if it's not the default logger for your system, you're almost guaranteed to have rsyslog available through your distribution's package manager. On our CentOS system, we run `yum install rsyslog` to install the logger and any dependencies.

```
$ sudo /sbin/service syslog stop
Shutting down kernel logger:
Shutting down system logger:
```

```
$ sudo /sbin/chkconfig syslog off
```

```
$ █
```

09 Removing syslog

We need to stop the existing syslog daemon and then remove it from the various runlevels. Use `/sbin/service syslog stop` to stop the immediate process. Then use `/sbin/chkconfig syslog off` to switch off the corresponding runlevel scripts – this is so that it doesn't start back up again if the system reboots.

```
$ ls -l /etc/init.d/rsyslog
-rwxr-xr-x 1 root root 2454 Nov 30 2009 /etc/init.d/rsyslog

$ sudo /sbin/chkconfig --add rsyslog

$ sudo /sbin/chkconfig rsyslog on
```

10 Bringing rsyslog into play


If we look in `/etc/init.d/` we can see that the package manager has added a runlevel script for rsyslog. To bring rsyslog in as a replacement for syslog, we can run `/sbin/chkconfig -add rsyslog`, then `/sbin/chkconfig rsyslog on` to add the control script and then turn on their use.

```
$ sudo /sbin/service rsyslog start
Starting system logger:
```

```
$ █
```

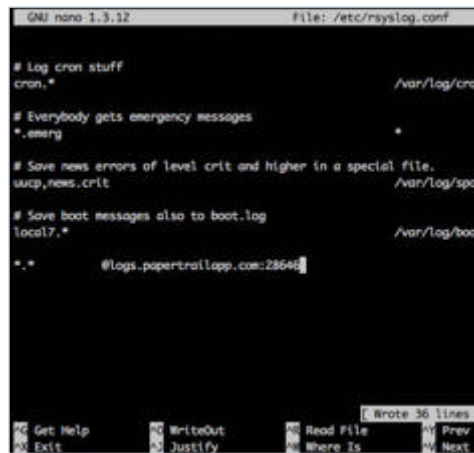
11 Starting the service

This, of course, does not mean that rsyslog is running – it just means it will be told to start or stop at the next runlevel change, depending on the configuration. We can bring it online with the service command `/sbin/service rsyslog start`.



12 New config file

If we run `ls -d /etc/*syslog*` again, we will see that a new config file has been added for rsyslog. If you haven't made any changes to the standard syslog.conf, then the new file should suffice, as the files are broadly compatible. If you have made changes, try using diff to merge the changes.



“Email alerts or web hooks can set up for significant events”

13 Forwarding to Papertrail

To set up the forwarding to Papertrail, add `*.* @logs.papertrailapp.com:28646` to the rsyslog configuration file, making sure that you're using the port number that Papertrail has assigned to your account. If you don't want to forward all log messages, then the `*.*` filter can be adjusted just as it normally would be.

```
$ sudo /sbin/service rsyslog restart
Shutting down system logger:
Starting system logger:
```

```
$ █
```

14 Restart the daemon

Restart the rsyslog for this to take effect. Any new messages that are now sent to the logging daemon and match our filter will be forwarded on to Papertrail as well as being funnelled into local files as they have previously been.



15 The Events window

Restarting rsyslog does, in itself, generate some messages that consequently are sent to Papertrail. By pulling up the Events window, we can immediately see these messages displayed. While the Events window is open, it will automatically scroll as new messages arrive.

Messages can be filtered by typing into the search box at the bottom. It supports a fairly standard search syntax such as quoting phrases and negating terms. The Help button at the top has more info.

To restrict the messages to within a particular date or time, use the calendar button next to the search box. Remember that on a free plan, you can only search back a maximum of 24 hours.


```
$ sudo gem install remote_syslog
Building native extensions. This could take a while...
Successfully installed eventmachine-0.12.10
Successfully installed eventmachine-tail-0.6.3
Successfully installed syslog_protocol-0.9.1
Successfully installed em-resolv-replace-1.1.1
Successfully installed remote_syslog-1.4.2
5 gems installed
Installing ri documentation for eventmachine-0.12.10...
Installing ri documentation for eventmachine-tail-0.6.3...
Installing ri documentation for syslog_protocol-0.9.1...
Installing ri documentation for em-resolv-replace-1.1.1...
Installing ri documentation for remote_syslog-1.4.2...
Installing RDoc documentation for eventmachine-0.12.10...
Could not find main page README
Could not find main page README
Could not find main page README
Could not find main page README
Installing RDoc documentation for eventmachine-tail-0.6.3...
Installing RDoc documentation for syslog_protocol-0.9.1...
Installing RDoc documentation for em-resolv-replace-1.1.1...
Installing RDoc documentation for remote_syslog-1.4.2...
```

16 Aggregating other logs

With our logger daemon configured to forward on messages, we can turn our attention to other logs on the system. The Papertrail team have created a small aggregation utility to monitor log files and forward messages. It's available as a Ruby Gem, so make sure Ruby and Gem are available on your system, then run 'gem install remote_syslog'.

```
$ remote_syslog
Usage: remote_syslog [options] --path to add'l log to ... --path to add'l log to
Example: remote_syslog -c config/remote_syslog.yml -p 12345 /var/log/apache2.log

Options (default):
  -c, --configfile PATH          Path to config (/etc/log_files.yml)
  -d, --dest-host HOSTNAME      Destination syslog hostname or IP (logs.papertrailapp.com)
  -p, --port PORT                Destination syslog port (514)
  -s, --source SYSLOG            Don't daemonize and detach from the terminal
  -f, --facility FACILITY        Facility (user)
  -l, --logfile LOGFILE         Local filename to read from
  -L, --log-dir DIRECTORY       Directory to write .pid file in (/var/run/)
  -F, --pid-file FILENAME        PID filename (program name.pid)
  -r, --remote-syslog            Parse files as syslog-formatted file
  -S, --severity SEVERITY       Severity (notice)
  -c, --color                    Strip color codes
  -t, --tcp                       Connect via TCP with TLS
  -h, --help                       Show this message
```

17 Using remote_syslog

This will install the remote_syslog binary on your system. We can use it from the command line by giving it a port number using the '-p' switch and specifying a log file for it to work with. For example, 'remote_syslog -p 28646 /var/log/httpd/access_log' would start forwarding anything written to Apache's access log to Papertrail.

```
$ sudo remote_syslog -p 28646 /var/log/httpd/access_log
remote_syslog: process with pid 15545 started.

$ ps ax | grep remote_syslog
15510 ?        S      0:00 remote_syslog
15545 ?        S      0:00 remote_syslog
15646 pts/2    S+    0:00 grep remote_syslog
```

18 Multiple processes

Notice how in the previous step, the remote_syslog command has spawned itself as a separate process. This process will reside in memory, continuing to monitor the access_log file and send any new messages on to Papertrail. It does mean, however, that we cannot monitor another log file in the same way.

“The Papertrail team have created a small aggregation utility to monitor log files and forward messages”

```
$ sudo remote_syslog -p 28646 /var/log/httpd/error_log
ERROR: there is already one or more instance(s) of the prog
$
```

19 Configuration file

To overcome this, we can use the remote_syslog configuration file. By default, remote_syslog looks for this at /etc/log_files.yml, but you can configure this through the '--configfile' option if need be. For now, let's create /etc/log_files.yml and open it for editing.

```
GNU nano 1.3.12 File: /etc/log_files.yml
files: [ /var/log/httpd/access_log, /var/log/httpd/error_log ]
destination:
  host: logs.papertrailapp.com
  port: 28646
```

20 YAML format

The config file is in YAML format – a simple line-based text format that uses colons to designate keys, and square brackets / commas for lists (for more info, see <http://yaml.org/>). This example file sets up remote_syslog to forward both the Apache access and errors logs to Papertrail on our port number.

```
$ sudo killall remote_syslog

$ ps ax | grep remote_syslog
21801 pts/1    S+    0:00 grep remote_syslog

$ sudo remote_syslog
remote_syslog: process with pid 21842 started.
```

21 Using the config file

To get our config file running, we start by dropping all existing remote_syslog processes using 'killall'. Check that there's nothing still running, then start it without any arguments. The log files to watch will be read from the config file, as will the port, so we no longer have to specify that as the argument.



22 The CLI interface

With all our logging information going into Papertrail, we might like to get it back out again in a format we can use in other scripts. Papertrail has a CLI program to do exactly this. Install it using 'gem install papertrail-cli' and then retrieve your API key from your profile page.

23 Your API key

The API key must be added to a papertrail.yml file inside your home folder, as a 'token' line. Echo 'token: YSzpVM26aajpJJxxxx' and redirect the output to this file, changing it to match your actual API key.

```
$ papertrail -h
papertrail - command-line tail and search for Papertrail log management service
  -h, --help                Show usage
  -f, --follow              Continue running and print new events (off)
  -d, --delay SECONDS      Delay between refresh (3)
  -c, --configfile PATH    Path to config (/c:/papertrail.yml)

Usage: papertrail [-F] [-d SECONDS] [-c PATH/to/papertrail.yml] [query]

Examples:
papertrail -f
papertrail something
papertrail -f 1.2.3 feature
papertrail -f "(user OR db) Origin OR ipaddr" accepted
papertrail -f -d 10 "not OR connection refused"
```

24 Sample commands

You can now run 'papertrail-cli' from the command line, passing it any query that you would use in the search box in the Events window. Results are returned line-based, so they can be used with other tools such as grep or awk.

Better PHP on the command line

Build a flexible micro-framework to tame PHP command-line programs

Advisor

Andy Leon is a web developer with Java and PHP as his languages of choice, and a fair amount of server-wrangling as a secondary skill



Resources

Some PHP experience
Tutorial files (on the DVD)



The UNIX style of command line has persisted for over 30 years. Built on the concept that small, powerful tools can communicate universally through common plain text interfaces, it gives the ultimate in flexibility. Any number of these small programs can be linked through pipes and filters to achieve a bigger task.

This also allows the traditional UNIX utilities, written in C, to be mixed with other languages, be they compiled or interpreted, or somewhere in between. Dynamic, interpreted languages like PHP

do not fare well in these situations. A PHP program generally consists of multiple files that don't lend themselves to portability.

In this tutorial, we're going to show you how to create a micro-sized framework for working with PHP on the command line. We're going to build a directory in our include path to hold a bootstrap file and any supporting files our command-line programs will need. Combine this with a powerful class-loading component and we'll be able to write compact and portable PHP files that can be used to accomplish any task from within the shell.

```
$ which php
/opt/bitnami/php/bin/php
```

```
$ php -v
PHP 5.3.8 (cli) (built: Aug 26 2011 03:59:50)
Copyright (c) 1997-2011 The PHP Group
Zend Engine v2.3.0, Copyright (c) 1998-2011 Zend Technologies
```

```
$
```

01 Checking PHP

Make sure that you've got a command-line version of PHP installed. We can check this with 'which php'. Then run 'php -v' to check that you've got 5.3 or above installed, as we're going to be incorporating components that make use of namespaces. Note that the CLI version and Apache module can differ.

02 Installing the CLI version

If it's not there, or it's out of date, we can use a package manager to install it. In Ubuntu, for example, we can use 'sudo apt-get install php5' in order to install common PHP files, then 'sudo apt-get install php5-cli' to install the command-line interface.

03 Find the config

With PHP CLI installed, we can run 'php -i' to display information about our PHP version. For now, we just want to find out where our main config file is, so filter the output by running 'php -i | grep "Loaded Configuration"'.



```
$ php -i | grep "Loaded Configuration"
Loaded Configuration File => /opt/bitnami/php/etc/php.ini
```

```
$
```

```
default_mimetype = "text/html"

; PHP's default character set is set to empty.
; http://php.net/default-charset
default_charset = "iso-8859-1"

; Always populate the $HTTP_RAW_POST_DATA variable. PHP's default behavior is
; to disable this feature.
; http://php.net/always-populate-raw-post-data
always_populate_raw_post_data = 0

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Paths and Directories ;
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

;INI: "/path2:/path2"
include_path = ".:/opt/bitnami/php/lib/php"

; Windows: "/path2:/path2"
include_path = ".:/opt/bitnami/php/lib/php"

; PHP's default setting for include_path is ".:/path2:/php/pear"
http://php.net/include-path

; The root of the PHP pages, used only if nonempty.
; If PHP was not compiled with FORCE_REDIRECT, you SHOULD set doc_root.
; If you are running php as a CGI under any web server (other than IIS)
; see documentation for security issues.  The alternate is to use the
; cgi.force_redirect configuration below
; http://php.net/doc-root
doc_root =

; The directory under which PHP opens the script using /-username used only
; if nonempty.
```

04 Locate the include path
We want to make our framework available to PHP from anywhere, so we need to investigate where PHP will look for files. Open the configuration file and search for the 'include_path' directive. If there's a suitable directory in here to base our files, use that; if not, create one and add it here.

```
$ pwd
/opt/bitnami/php/lib/php

$ mkdir cli

$ mkdir cli/src cli/vendor

$
```

05 Create our base
Change into this directory and create a folder called 'cli' to hold our framework. Inside the 'cli' directory, create a 'src' directory to hold code that we'll write, and a 'vendor', which we'll use to hold any third-party libraries that we need. You may need to set some less restrictive permissions on the directories.

```
$ sudo apt-get install git-core
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following extra packages will be installed:
  libdigest-sha1-perl liberror-perl
Suggested packages:
  git-doc git-arch git-cvs git-email git-gnome-run git-gui gitk gitweb
The following NEW packages will be installed:
  git-core libdigest-sha1-perl liberror-perl
0 upgraded, 3 newly installed, 0 to remove and 31 not upgraded.
Need to get 6,193B of archives.
```

06 Installing Git
We're going to use the Git version control system to install individual PHP components and libraries from GitHub. Use 'which git' to check whether it's installed and 'sudo apt-get install git-core' if it's not. There are other, graphical Git packages, but we'll just use the command line.

“The first function of the bootstrap is to start an instance of the class loader”

```
$ git clone https://github.com/symfony/ClassLoader.git vendor/Symfony/Component/ClassLoader
Initialized empty Git repository in /opt/bitnami/php/lib/php/vendor/Symfony/Component/ClassLoader
remote: Counting objects: 185, done.
remote: Compressing objects: 100% (87/87), done.
remote: Total 185 (delta 116), reused 167 (delta 98)
Receiving objects: 100% (185/185), 29.33 KiB, done.
Resolving deltas: 100% (116/116), done.

$
```

07 The ClassLoader component
The first component we'll need is the Symfony ClassLoader. This is used to autoload other components when we need them by maintaining a list of namespaces and the directories associated with them. From the 'cli' directory, run 'git clone https://github.com/symfony/ClassLoader.git vendor/Symfony/Component/ClassLoader' to add it to our framework.

```
$ git clone https://github.com/symfony/Console.git vendor/Symfony/Component/Console
Initialized empty Git repository in /opt/bitnami/php/lib/php/vendor/Symfony/Component/Console
remote: Counting objects: 683, done.
remote: Compressing objects: 100% (334/334), done.
remote: Total 683 (delta 467), reused 565 (delta 349)
Receiving objects: 100% (683/683), 123.35 KiB, done.
Resolving deltas: 100% (467/467), done.

$
```

08 The Console component
We're also going to utilise Symfony's Console component. This will provide a lot of the text processing we need for our command to handle arguments and options, as well as an easy cross-platform way to handle formatted output. Again, from the 'cli' directory, run 'git clone https://github.com/symfony/Console.git vendor/Symfony/Component/Console'.

```
1 <?php
2
3 /*
4 * bootstrap.php
5 *
6 */
```

09 The bootstrap file
Create a 'bootstrap' file to do a lot of the legwork for our commands. Because we've placed this within PHP's include path, it will be available to any PHP script, regardless of where we store it. Start by creating 'bootstrap.php' in our 'cli' directory.

```
1 <?php
2
3 /*
4 * bootstrap.php
5 *
6 */
7
8 require __DIR__.'vendor/Symfony/Component/ClassLoader/ClassLoader.php';
9
10 $loader = new Symfony\Component\ClassLoader\ClassLoader();
```

10 Loading classes
The first function of the bootstrap is to start an instance of the class loader. Use a require statement to pull in the ClassLoader class, and then create a new instance of it. As this is responsible for loading classes, we can't issue a 'use' statement until this is up and running, so we need to use the long name.

11 Component namespaces
At the moment, the only components we've got available are Symfony components, which are both under the same Symfony namespace. We add this namespace to the ClassLoader and tell it which directory to find it in. This is an array, so we can easily add more namespaces here for different components.


```
1 <?php
2
3 /*
4  * bootstrap.php
5  *
6  */
7
8 require __DIR__ . '/vendor/Symfony/Component/ClassLoader/UniversalClassLoader.php';
9
10 $loader = new Symfony\Component\ClassLoader\UniversalClassLoader();
11
12 $loader->registerNamespaces(
13     array(
14         'Symfony' => __DIR__ . '/vendor'
15     )
16 );
17
18 $loader->register();
```

12 Register the loader

It's worth noting that this method is only good for namespaced components (PHP 5.2 and upwards). If we need to dynamically load components with the older style of class names (eg `Vendor_Product_Class()`), then we can use the similar `registerPrefixes()` method. Either way, with our configuration done, we finalise the class loader with the `register()` method.

```
1 <?php
2
3 /*
4  * console.php
5  *
6  */
7
8 require 'cli/bootstrap.php';
```

13 Create a test script

To run a test, we want to run a test script somewhere other than in the 'cli' directory. Create a file in your home directory called 'console.php' and add a single require line to it, calling in 'cli/bootstrap.php'. Run this script using 'php console.php'. If you get an error, check your 'include_path' setting again.

```
1 <?php
2
3 /*
4  * console.php
5  *
6  */
7
8 require 'cli/bootstrap.php';
9
10 use Symfony\Component\Console\Application;
11 use Symfony\Component\Console\Input\InputInterface;
12 use Symfony\Component\Console\Output\OutputInterface;
13 use Symfony\Component\Console\Input\InputOption;
14 use Symfony\Component\Console\Input\InputArgument;
```

14 Cleaner code

Now to explore the Symfony console component a bit. It consists of an Application class that can have multiple commands registered to it. Each of these commands deals with an InputInterface and OutputInterface. To make our code more readable, we add use statements for these and a couple of others to our test command.

15 Creating a command

We can now create a new Application instance, passing in a name and a version number, and create a command for that application, called 'test'. We can add a description to our command, as well as some help text. The `<info>` tags help us with formatting our output. They provide an easy-to-use and cross-platform shortcut to achieving common formatting. You can use other tags, such as `<comment>`, `<question>` and `<error>` to achieve different layouts.

16 Providing output

The code that gives our command legs is set in the form of a closure passed to the `setCode()` method. The parameters are an InputInterface and an OutputInterface. In this simple test, we'll use the OutputInterface to write a line to the console. Call the `run()` method of the Application object to put it all into action.

“The code that gives our command legs is set in the form of a closure passed to the `setCode()` method”

```
$ php console.php test
Does not test anything yet

$
```

17 Running the command The command can now be run using 'php console.php test'. If we had registered multiple commands to the same Application object, then we could change 'test' to any of these in order to have different behaviour. We could also use 'php test.php help test' to display our help text.

```
1 <php
2 /*
3  * console.php
4  *
5  */
6
7
8 require 'cli/bootstrap.php';
9
10 use Symfony\Component\Console\Application;
11 use Symfony\Component\Console\Input\InputInterface;
12 use Symfony\Component\Console\Output\OutputInterface;
13 use Symfony\Component\Console\Input\InputOption;
14 use Symfony\Component\Console\Input\InputArgument;
15
16 $app = new Application('Test Console App', '0.1');
17
18 $command = $app->register('test')
19     ->setDescription('A command to test things')
20     ->setHelp('Usage: <info>php console.php test</info>');
21
22 $command->setDefinition(
23     array(
24         new InputArgument('site', InputArgument::REQUIRED, 'Site to check')
25     )
26 );
27
28 $command->setCode(
29     function(InputInterface $input, OutputInterface $output)
30     {
31         $output->writeln('Does not test anything yet');
32     }
33 );
```

18 Retrieving input To do something more useful with our application, we need to be able to accept input from the user. For this, we must first define what input we expect. Call the setDefinition() method on our command, passing in an array of InputArgument (or InputOption) objects. Here we specify a required 'site' argument.

```
$ php console.php test
[RuntimeException]
Not enough arguments.

test site

$ php console.php test site=http://kernel.org
Does not test anything yet
```

19 Passing arguments If we try to run our command now, we'll get an error message in the form of a RuntimeException. To satisfy this requirement, we can add a site address after the 'test' command. This can be a raw string provided it doesn't contain any spaces. Otherwise, it will need to be quoted or escaped.

“Simple output tags help us create cross-platform formatting of our output easily”

```
2
3 /*
4  * console.php
5  *
6  */
7
8 require 'cli/bootstrap.php';
9
10 use Symfony\Component\Console\Application;
11 use Symfony\Component\Console\Input\InputInterface;
12 use Symfony\Component\Console\Output\OutputInterface;
13 use Symfony\Component\Console\Input\InputOption;
14 use Symfony\Component\Console\Input\InputArgument;
15
16 $app = new Application('Test Console App', '0.1');
17
18 $command = $app->register('test')
19     ->setDescription('A command to test things')
20     ->setHelp('Usage: <info>php console.php test</info>');
21
22 $command->setDefinition(
23     array(
24         new InputArgument('site', InputArgument::REQUIRED, 'Site to check')
25     )
26 );
```

20 Sending input back out Inside our main command code, we can now access this argument as part of the InputInterface, through the getArguments() method. Here, we simply echo the value back out to the console, through the OutputInterface.

```
11 use Symfony\Component\Console\Input\InputOption;
12 use Symfony\Component\Console\Input\InputArgument;
13
14 $app = new Application('Test Console App', '0.1');
15
16 $command = $app->register('test')
17     ->setDescription('A command to test things')
18     ->setHelp('Usage: <info>php console.php test</info>');
19
20 $command->setDefinition(
21     array(
22         new InputArgument('site', InputArgument::REQUIRED, 'Site to check')
23     )
24 );
25
26 $command->setCode(
27     function(InputInterface $input, OutputInterface $output)
28     {
```

21 Searching through HTML With this incoming argument, we can do a very rudimentary test of whether a site is working properly. By retrieving the contents of the index page into a string, we can search for the closing </html> tag – this is a good indicator that the page has not run into any errors.

22 Testing our test Run the test on http://kernel.org and, all being well, we should see the confirmation that the page is OK. Google famously does not close its <html> tag, so this is a good failure assertion. Run the test on http://google.com and it will return a failure.

23 (Backwards) return values Our command is a little verbose at the moment and does not return proper values. Commands should not be 'chatty', so we should remove the 'Page is OK' message and instead just return 0 for success or a positive integer to indicate an error. Remember, command-line return values are counter-intuitive to a PHP programmer.

```
$ php console.php test http://kernel.org

$ echo $?
0

$ php console.php test http://google.com
Site may be down

$ echo $?
1

$
```

24 Our finished command If we rerun the tests, there's now less output, but we can retrieve proper command-line return values to establish whether there was an error or not. With no output on success, we can pipe input to a log file, so that the log file only gets appended to when there's an error.

Many different modules can be loaded in using a simple require() method. Some modules are being standardised as part of CommonJS

Node is more pure networking language than others that run through a web server. We have fine control over HTTP through the ability to create listening servers

Using standard request and response objects we can build dynamic routes in our application for different URLs

We have to deliver static files ourselves though, as well as generate any errors that might occur

```
1 /* test.js - Test Application */
2
3 var util = require('util'),
4     http = require('http'),
5     fs = require('fs'),
6     path = require('path');
7
8 http.createServer(function(request, response) {
9     if(request.url === '/test') {
10        response.statusCode = 200;
11        response.setHeader('Content-type', 'text/html');
12        response.write('<p>Sample output</p>');
13    }
14    else {
15        staticFile(request, response);
16    }
17 }).listen(8080);
18
19 function staticFile(request, response) {
20     var filePath = path.join(process.cwd(), 'web', request.url);
21
22     path.exists(filePath, function(exists) {
23         if(exists) {
24             fs.readFile(filePath, 'binary', function(error, file) {
25                 if(error) {
26                     response.statusCode = 500;
27                     response.write(error);
28                     response.end();
29                 }
30                 else {
31                     response.statusCode = 200;
32                     response.write(file, 'binary');
33                     response.end();
34                 }
35             });
36         }
37         else {
38             response.statusCode = 404;
39             response.end();
40         }
41     });
42 }
```

Server-side JavaScript with Node.js

JavaScript is not just confined to the browser – it can power applications server-side as well. Here's how...



Thanks to Netscape, JavaScript has become ubiquitous in the world of web work for adding interactivity to pages. Due to the actions of many others, it's also become a hated language for this very reason. However, much of what people dislike about JavaScript is not the language, but the creaking document object model they have to control to get anything useful from a webpage.

JavaScript is very interesting – the most used prototype-based language by quite a long way – and, outside of its native environment, sheds a lot of the awkwardness often associated with it.

Recent speed increases of the JavaScript implementations in modern web browsers have piqued interest in using it as a back-end language. Since every web developer is exposed to JavaScript, why learn another back-end language when you could use just the one?

Borrowing the V8 JavaScript engine used in the Google Chrome browser, Node.js provides a way to build complex network applications using JavaScript. It provides libraries for areas JavaScript normally doesn't go, such as file system access or standard IO. Let's get it up and running and see what it can do.

Advisor

Andy Leon is a web developer with Java and PHP as his languages of choice, and a fair amount of server-wrangling as a secondary skill





```
$ curl http://nodejs.org/dist/v0.6.10/node-v0.6.10.tar.gz -o nodejs.tar.gz
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload  Total   Spent    Left   Speed
100 10.0M  100 10.0M    0     0  583k      0  0:00:17  0:00:17 --:--:-- 734k

$ tar -xzf nodejs.tar.gz

$ cd node-v0.6.10/
```

01 Downloading source

Download the most up-to-date source package from <http://nodejs.org> (current 0.6.10), unzip and extract the archive and then change into the 'node-v0.6.10/' directory, replacing the version number if need be. There are RPMs and DEBs available for binary installs as well.

```
Checking for gcc           : ok
Checking for library dl    : yes
Checking for openssl       : not found
Checking for function SSL_library_init : yes
Checking for header openssl/crypto.h : yes
Checking for library util  : yes
Checking for library rt    : not found
Checking for fdotasync(2) with c++   : no
'configure' finished successfully (4.882s)

$
```

02 Building and Installing

In the source directory, it's the usual procedure of './configure', then 'make', then finally 'sudo make install'. As long as the configure phase passes, do not worry about anything reported 'missing'. Once complete, we can confirm the installation by running 'which node'.

```
/* test.js - Test Application */

var util = require('util');|
```

03 Our application file

Start by creating a file called 'test.js' and opening it up in an editor. Similar to Java, a Node script starts able to do very little – almost all functionality is contained within modules, which must be included using require.

```
/* test.js - Test Application */

var util = require('util');

util.puts('Test Application');
```

04 Loading modules

We've loaded the 'util' module into a variable called util. This variable could have been anything, but for convenience we've given it the same name as the module. We can now call a method of that module, whose function is to output some text to the standard output.

```
$ node test.js
Test Application

$
```

05 First test run

Save the file down and then run it from the command line using the 'node' binary. If there are any errors they will be shown in the terminal – otherwise it will echo 'Test Application' on its own line and then exit.

06 Finding help

The 'puts' method is obviously only one of a number of methods that can be called for this module, which itself is just one of a number of different modules. Node's documentation (<http://nodejs.org/docs/latest/api/index.html>) is updated regularly with the module and method names.

```
/* test.js - Test Application */

var util = require('util');
    http = require('http');|

util.puts('Test Application');
```

07 Adding HTTP

Our first app runs, but it's not a web app. We could use Node as part of a CGI setup, calling the binary for each request, but we'd lose out on a lot of the benefits Node offers. Instead we'll use the 'http' module, now included at the start.

```
/* test.js - Test Application */

var util = require('util');
    http = require('http');|
    http.createServer(function(request, response) {
}).listen(8080);

util.puts('Test Application');
```

08 Create a listening server

The 'http' module lets us create a server, passing in a function for handling requests and delivering responses. Leave the function blank for now, then call the 'listen' method with an available port number to get the server running.

09 Why won't it stop?

If we run this program now, we notice that 'Test Application' gets printed to the terminal but that the application never terminates. This highlights how the 'createServer' method has started its own thread, which runs constantly but allows the rest of the application to continue.

The screenshot shows the Node.js website navigation menu on the left with options: HOME, DOWNLOAD, ABOUT, NPM REGISTRY, DOCS, BLOG, COMMUNITY, LOGOS, JOBS, and @nodejs. The main content area features the Node.js logo, the title 'Node.js v0.6.10 Manual & Document', a link to 'Index | View on single page', and a 'Table of Contents' section with links to: Synopsis, Globals, STUDIO, Timers, Modules, and C/C++ Addons.

“Node.js provides libraries for areas JavaScript normally doesn't go”

```
/* test.js - Test Application */  
  
var util = require('util'),  
    http = require('http');  
  
http.createServer(function(request,  
    response.end());  
}).listen(8080);  
  
util.puts('Test Application');
```

10 Closing requests

If we point our browser at `http://localhost:8080`, we'll get a similar effect of the browser doing nothing until it times out. This is because our function doesn't do anything. If we close the request with the `end()` method then the browser will not time out, it will just receive a blank response.

```
/* test.js - Test Application */  
  
var util = require('util'),  
    http = require('http');  
  
http.createServer(function(request, response) {  
    // Scripts that operate in a Web server  
    // through SAPI behind here...  
  
    // ... and would end here.  
    response.end();  
}).listen(8080);  
  
util.puts('Test Application');
```

11 No server assistance

There's an important distinction here. Because we're starting a server session from within our code, we have to do the HTTP legwork that other languages might handle for us. In PHP, for example, the request is automatically terminated when the end of the script is reached.

```
/* test.js - Test Application */  
  
var util = require('util');  
    http = require('http');  
  
http.createServer(function(request, response) {  
    response.write('<p>Sample output</p>');  
    response.end();  
}).listen(8080);  
  
util.puts('Test Application');
```

12 Displaying output

To do something more useful than just dropping connections, we can write output to the response using the `write()` method. After restarting the application, the browser will now receive our paragraph text for its request, but it will display it in plain text.

“Node.js provides a way to build complex network applications using JavaScript”

13 Sending headers

This is because of another thing we've taken for granted: the default headers that languages running through a web server send along with the output. In Node we have to specify these, so we should at least send a status code ('200' for success) and tell the browser that our content is HTML.

```
/* test.js - Test Application */  
  
var util = require('util');  
    http = require('http');  
  
http.createServer(function(request, re  
    response.statusCode = 200;  
    response.setHeader('Content-type',  
    response.write('<p>Sample output<  
    response.end();  
}).listen(8080);  
  
util.puts('Test Application');
```

14 One-for-all

With the headers in place, the browser now recognises the content as HTML and gives us a proper paragraph. The next quirk we notice is that at no point has the browser requested 'test.js'. In fact, we can request any URL on our server and get the same response.

```
$ curl http://localhost:8080  
<p>Sample output</p>  
$ curl http://localhost:8080/any_random_url  
<p>Sample output</p>  
$
```

```
/* test.js - Test Application */  
  
var util = require('util'),  
    http = require('http');  
  
http.createServer(function(request, response) {  
    if(request.url === '/test') {  
        response.statusCode = 200;  
        response.setHeader('Content-type', 'text/html');  
        response.write('<p>Sample output</p>');  
    }  
    else {  
        response.statusCode = 404;  
    }  
    response.end();  
}).listen(8080);
```

15 Adding routes

Looking at our code, we can see why this is. For every single request we're returning the same response. To be more functional we need to be able to set up some routing. By testing our request object we can make sure we only return a response for a given URL (in this case '/test'). Our else condition is to set the status code to 404 ('not found').

In addition to our application pages, we need to be able to serve static files such as images, stylesheets and front-end JavaScript. Create a directory called 'web/' and add a 'favicon.ico' file to it.

```
/* test.js - Test Application */  
  
var util = require('util'),  
    http = require('http');  
  
http.createServer(function(request, response) {  
    if(request.url === '/test') {  
        response.statusCode = 200;  
        response.setHeader('Content-type', 'text/html');  
        response.write('<p>Sample output</p>');  
    }  
    else {  
        staticFile(request, response);  
    }  
    response.end();  
}).listen(8080);
```

16 Static files

Instead of immediately returning a 404 error, we're going to see if there's a static file available in our 'web' directory that matches the URL that's been requested. We could do this in line, but it's neater if we pass the request and response to a function.

“Because we're starting a server session from within our code, we have to do the HTTP legwork”

```
/* test.js - Test Application */
var util = require('util'),
    http = require('http'),
    fs = require('fs'),
    path = require('path');

http.createServer(function(request, response) {
  if(request.url === '/test') {
    response.statusCode = 200;
    response.setHeader('Content-type', 'text/html');
    response.write('<p>Sample output</p>');
  }
  else {
    staticFile(request, response);
  }
}

response.end();

}).listen(8080);
```

17 Adding more modules

Before we build this function, there's a couple of additional modules that we need to include the functionality from. At the top of our script, add `require()` calls for the 'fs' module (file system) and 'path' module (working with file paths).

```
response.write('<p>Sample output</p>');
}
else {
  staticFile(request, response);
}

response.end();

}).listen(8080);

function staticFile(request, response) {
  var filePath = path.join(process.cwd(), 'web', request.url);
}
```

18 Getting file paths

We can now declare our function and establish a path in the file system to search for the URL requested. We use the `cwd()` method of the 'process' object to get the current working directory, then we join it with our 'web/' directory and the request URL.

```
response.write('<p>Sample output</p>');
}
else {
  staticFile(request, response);
}

response.end();

}).listen(8080);

function staticFile(request, response) {
  var filePath = path.join(process.cwd(), 'web', request.url);
  path.exists(filePath, function(exists) {
  });
});
```

19 Finding a file

Inside our function we use the `path.exists()` method to check whether the file requested exists in our 'web/' directory. This does not use a standard true/false return value as you might expect, but a callback function with the result.

20 File not found

We can test the result of this inside the callback function. If the file does not exist then we set the response status to 404 as we did previously. We don't need to close the request as the flow will pass back to where we first called `staticFile()`.

```
function staticFile(request, response) {
  var filePath = path.join(process.cwd(), 'web', request.url);
  path.exists(filePath, function(exists) {
    if(exists) {
    }
    else {
      response.statusCode = 404;
      response.end();
    }
  });
}
```

```
function staticFile(request, response) {
  var filePath = path.join(process.cwd(), 'web', request.url);
  path.exists(filePath, function(exists) {
    if(exists) {
      fs.readFile(filePath, 'binary', function(error, file) {
        if(error) {
        }
        else {
        }
      });
    }
    else {
      response.statusCode = 404;
      response.end();
    }
  });
}
```

21 Reading a file

If the file exists, we can attempt to open the file. We use 'binary' as the mode as we're unsure of the type and, once again, we use a callback function to establish whether we've been successful, rather than a true/false return value.

```
function staticFile(request, response) {
  var filePath = path.join(process.cwd(), 'web', request.url);
  path.exists(filePath, function(exists) {
    if(exists) {
      fs.readFile(filePath, 'binary', function(error, file) {
        if(error) {
        }
        else {
          response.statusCode = 200;
          response.write(file, 'binary');
          response.end();
        }
      });
    }
    else {
      response.statusCode = 404;
      response.end();
    }
  });
}
```

22 Sending the file back

Instead of testing for success here, we're actually testing for an error. If there's no error then we're fine to send the file to the browser. We change the status code of the response to 200 and then we output the file in binary form.

```
function staticFile(request, response) {
  var filePath = path.join(process.cwd(), 'web', request.url);
  path.exists(filePath, function(exists) {
    if(exists) {
      fs.readFile(filePath, 'binary', function(error, file) {
        if(error) {
          response.statusCode = 500;
          response.write(error);
          response.end();
        }
        else {
          response.statusCode = 200;
          response.write(file, 'binary');
          response.end();
        }
      });
    }
    else {
      response.statusCode = 404;
      response.end();
    }
  });
}
```

23 Handling errors

If there was an error, we change the status code to 500 ('Internal Server Error') and then we output the text of the error message we received. In a more thorough implementation, you might check what the error was and return a different status code (eg 403 for 'forbidden').

```
/* test.js - Test Application */
var util = require('util'),
    http = require('http'),
    fs = require('fs'),
    path = require('path');

http.createServer(function(request, response) {
  if(request.url === '/test') {
    response.statusCode = 200;
    response.setHeader('Content-type', 'text/html');
    response.write('<p>Sample output</p>');
  }
  else {
    staticFile(request, response);
  }
}

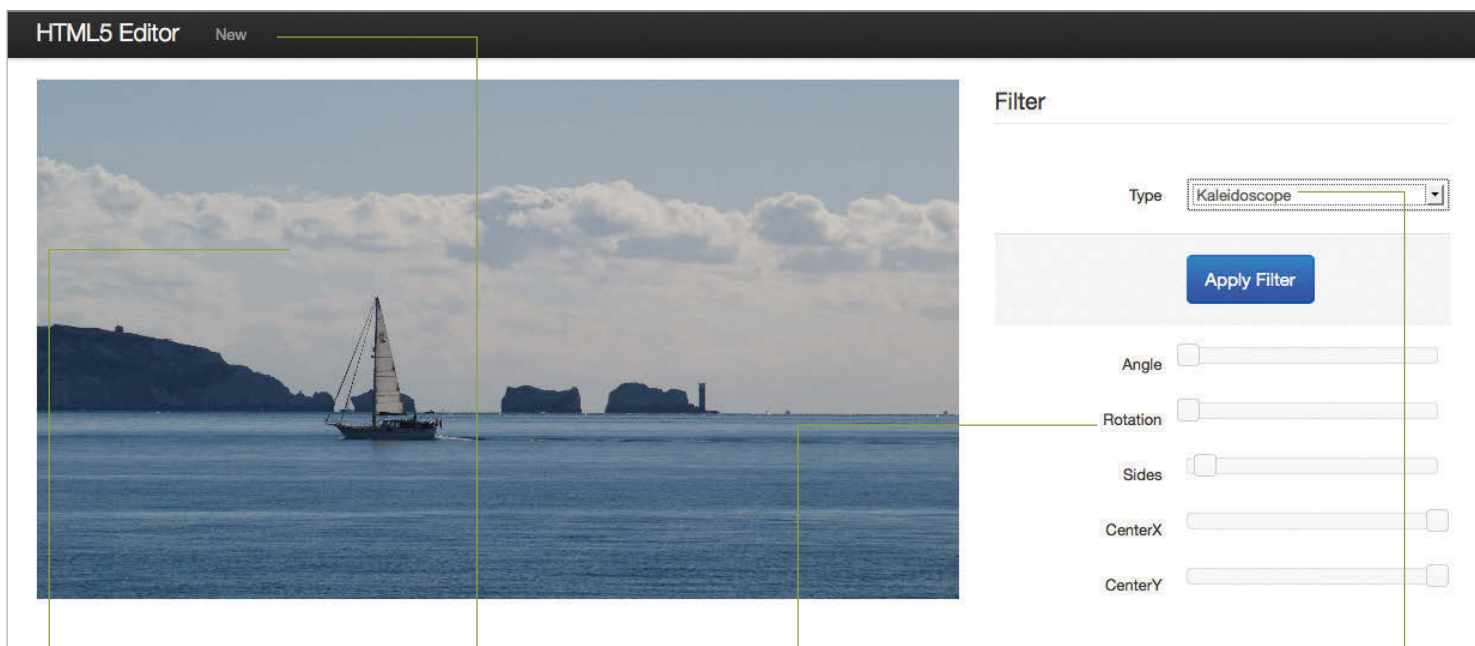
response.end();

}).listen(8080);

function staticFile(request, response) {
  var filePath = path.join(process.cwd(), 'web', request.url);
  path.exists(filePath, function(exists) {
    if(exists) {
      fs.readFile(filePath, 'binary', function(error, file) {
        if(error) {
          response.statusCode = 500;
          response.write(error);
          response.end();
        }
        else {
          response.statusCode = 200;
          response.write(file, 'binary');
          response.end();
        }
      });
    }
    else {
      response.statusCode = 404;
      response.end();
    }
  });
}
```

24 Our Node server

Our program should now be capable of fulfilling a single dynamic request, sending back static files and handling 'file not found' errors - all in under 50 lines of code and written in the same language we would use in the browser.



Our `<canvas>` element can load photos from local disk into the browser through a drag-and-drop interface

Click the 'New' link at the top of the page to start over with another photo

If the filter takes options, they will appear here in the form of sliders, so they can be adjusted easily

We can select from over 40 filters to apply to the image. Our editor will dynamically load the required code when needed

Create an HTML5 photo editor

Resources

Firefox source code:

<https://github.com/acleon/MaxE> (also on the disc)

Use the HTML5 `<canvas>` element to load and manipulate your images

Advisor

Andy Leon is a web developer with Java and PHP as his languages of choice, and a fair amount of server-wrangling as a secondary skill



HTML5 is a broad specification, bringing improvements to many areas. Of all this new technology, though, it seems the `<canvas>` element has captured the most interest and presented the best opportunities for building web apps. The `<canvas>` element is a simple concept: give the developer a space they can draw to using vector and raster functions.

HTML5 has spurred the need for better JavaScript tooling and over the past year we've

been spoilt with a range of new libraries to accomplish various tasks. In this tutorial we will combine some JavaScript image filters with a dynamic script loader, a templating library and a trusty jQueryUI component to build a complete web app.

We'll use HTML5's drag-and-drop support and FileReader API to load images from the user's machine, render them to a `<canvas>` tag and then manipulate them in real-time using filters you'd find in any desktop image editor.



```
<!DOCTYPE HTML>
<html lang="en">
<head>
  <meta charset="UTF-8" />
  <title>HTML5 Image Editor</title>

  <link rel="stylesheet" href="css/bootstrap.min.css" />

  <link rel="stylesheet" href="css/editor.css" />
  <script src="js/editor.js"></script>

</head>
<body>
</body>
</html>
```

01 Blank HTML5 document

Create a blank HTML5 document with the appropriate doctype declaration. Use `<link>` and `<script>` tags to add a default stylesheet and script file. We'll also add Twitter's Bootstrap CSS (<http://twitter.github.com/bootstrap/>) to help with the styling of the page.

```
<!DOCTYPE HTML>
<html lang="en">
<head>
  <meta charset="UTF-8" />
  <title>HTML5 Image Editor</title>

  <link rel="stylesheet" href="css/bootstrap.min.css" />

  <link rel="stylesheet" href="css/js/jquery/jquery-ui-1.8.16.custom.css" />
  <script src="http://ajax.googleapis.com/ajax/libs/jquery/1.7.2/jquery.min.js"></script>
  <script src="js/jquery-ui.min.js"></script>

  <link rel="stylesheet" href="css/editor.css" />
  <script src="js/editor.js"></script>
</head>
```

02 jQuery and jQuery UI

We'll be using jQuery for most of our scripting and we'll need parts of the jQuery UI for working with sliders. We can link against a hosted version of jQuery and use a custom build of jQuery UI (<http://jqueryui.com/download>) – we only need the 'Core', 'Widget', 'Mouse' and 'Slider' included.

```
<!DOCTYPE HTML>
<html lang="en">
<head>
  <meta charset="UTF-8" />
  <title>HTML5 Image Editor</title>

  <link rel="stylesheet" href="css/bootstrap.min.css" />

  <link rel="stylesheet" href="css/js/jquery/jquery-ui-1.8.16.custom.css" />
  <script src="http://ajax.googleapis.com/ajax/libs/jquery/1.7.2/jquery.min.js"></script>
  <script src="js/jquery-ui.min.js"></script>

  <link rel="stylesheet" href="css/editor.css" />
  <script src="js/editor.js"></script>

  <script src="js/libraries/mustache.js"></script>
</head>
```

03 JSManipulate filters

Download the individual JavaScript image filters created by Joel Basada (<https://github.com/JoelBasada/JSManipulate>). Put these into a folder called 'filters' and then include the core 'filterutils.js' in the <head>. The other files will be loaded as we need them.

```
<!DOCTYPE HTML>
<html lang="en">
<head>
  <meta charset="UTF-8" />
  <title>HTML5 Image Editor</title>

  <link rel="stylesheet" href="css/bootstrap.min.css" />

  <link rel="stylesheet" href="css/js/jquery/jquery-ui-1.8.16.custom.css" />
  <script src="http://ajax.googleapis.com/ajax/libs/jquery/1.7.2/jquery.min.js"></script>
  <script src="js/jquery-ui.min.js"></script>

  <link rel="stylesheet" href="css/editor.css" />
  <script src="js/editor.js"></script>

  <script src="js/libraries/mustache.js"></script>
  <script src="js/libraries/filterutils.js"></script>
  <script src="js/libraries/blur.js"></script>
  <script src="js/libraries/contrast.js"></script>
  <script src="js/libraries/crop.js"></script>
  <script src="js/libraries/drop.js"></script>
  <script src="js/libraries/flip.js"></script>
  <script src="js/libraries/gaussian.js"></script>
  <script src="js/libraries/rotate.js"></script>
  <script src="js/libraries/sharpen.js"></script>
  <script src="js/libraries/sepia.js"></script>
  <script src="js/libraries/sun.js"></script>
  <script src="js/libraries/undo.js"></script>
  <script src="js/libraries/vignette.js"></script>
  <script src="js/libraries/watermark.js"></script>
  <script src="js/libraries/white.js"></script>
  <script src="js/libraries/yepnope.js"></script>
</head>
```

04 YepNope

For our autoloading, we'll be using YepNope (<http://yepnopejs.com>) – a remarkably simple JavaScript library. We can give it a test to run and then provide functions to execute or files to load based on the outcome of that test – either 'yep' or 'nope'.

```
<script src="http://yepnope.com/yepnope.js"></script>
<script src="js/libraries/yepnope.js"></script>
<script src="js/libraries/yepnope-test.js"></script>
<script src="js/libraries/yepnope-load.js"></script>
```

05 Templating with Mustache

We'll also need a bit of light templating in order to generate the markup for our sliders. For this we will be using Mustache – a logic-less template language using the familiar syntax of curly braces. Download the JavaScript version from <https://github.com/janl/mustache.js>.

06 Canvas and drop target

Create a <div> to be used as a drop target for our files (give it the ID 'fileDrop') and a <canvas> element that will hold our image data (give it the ID 'canvas'). Use CSS to style the drop target appropriately and hide the <canvas> element for the moment using 'display: none;'.

```
(function($){
$(document).ready(function(){
$('#fileDrop').on('dragover', function(e){
e.preventDefault();
});
$.event.props.push('dataTransfer');
$('#fileDrop').on('drop', function(e){
e.preventDefault();
processUploads(e.dataTransfer.files);
});
})(jQuery);
```

07 Drop target events

Attach functions to the 'dragover' and 'drop' events on our drop target. In each case, we prevent the default action from running. For the 'drop' event, we pass the received files to a 'processUploads()' function (this is separated so that we could have different upload methods).

```
function processUploads(files){
if(files instanceof FileList && files.length){
if(!files[0].type.match(/image.*$/)){
var img = document.createElement('img');
var reader = new FileReader();
img.onload = function(){
imageToCanvas(img, document.getElementById('canvas'));
reader.onload = function(){
img.src = e.target.result;
reader.readAsDataURL(files[0]);
};
} else {
alert('That is not a valid image file.');
```

08 Processing the uploads

In our 'processUploads' function we check for an image MIME type before creating a FileReader object to load the contents into an tag. When the tag has finished loading, we pass it and the <canvas> element into a function called 'imageToCanvas()'.

```
function imageToCanvas(img, canvas){
canvas.width = img.width;
canvas.height = img.height;
var ctx = canvas.getContext('2d');
ctx.drawImage(img, 0, 0);
}
```

09 Putting the image on the canvas

Our 'imageToCanvas()' function resizes our <canvas> tag using the width and height properties of the tag. We then retrieve the '2D' context of the <canvas> tag and use the 'drawImage()' function to recreate the image on the canvas.

```
$('#fileDrop').on('drop', function(e){
e.preventDefault();
processUploads(e.dataTransfer.files);
$(this).fadeOut('slow', function(){
$('#canvas').fadeIn('slow');
});
});
```

10 Showing the canvas

Back in our 'drop' event handler, we've now loaded the image into the <canvas> tag and no longer need the drop target. We use jQuery's animations to fade out the drop target and replace it with our <canvas> tag.

“Of all the new HTML5 tech, <canvas> has captured the most interest”

```
<select id="filterSelect">
  <option data-file="blur.js" data-function="BlurFilter">Blur</option>
  <option data-file="brightness.js" data-function="BrightnessFilter">Brightness</option>
  <option data-file="bump.js" data-function="BumpFilter">Bump</option>
  <option data-file="circlesmear.js" data-function="CircleSmearFilter">Circle Smear</option>
  <option data-file="contrast.js" data-function="ContrastFilter">Contrast</option>
  <option data-file="crosssmear.js" data-function="CrossSmearFilter">Cross Smear</option>
  <option data-file="diffusion.js" data-function="DiffusionFilter">Diffusion</option>
  <option data-file="dither.js" data-function="DitherFilter">Dither</option>
  <option data-file="edge.js" data-function="EdgeFilter">Edge</option>
  <option data-file="emboss.js" data-function="EmbossFilter">Emboss</option>
  <option data-file="exposure.js" data-function="ExposureFilter">Exposure</option>
</select>
```

11 Filter selector

Create a `<select>` element to hold a list of our filters. For each `<option>` tag, we add the name of the filter function and the name of the file which contains that filter. We get these from the JavaScript files in our 'filters/' folder.

```
$('#filterSelect').change(function() {
  var filterFunction = $(this).children('option:selected').data('function');
  var filterFile = 'filters/' + $(this).children('option:selected').data('file');

  yepnope({
    test: window[filterFunction],
  });
});
```

12 Selecting a filter

Attach a function to the 'change' effect of our filter selector that starts by fetching these custom properties and placing them in the 'filterFunction' and 'filterFile' variables. We add the 'filters/' folder name to our 'filterFile' variable as well.

13 Dynamically loading JavaScript

We can now use YepNope to check whether the filter function we want to use is defined. We pass in our 'filterFile' variable to the 'nope' property so that if the function does not exist yet, YepNope will load the file in order to define it.

```
$('#filterSelect').change(function() {
  var filterFunction = $(this).children('option:selected').data('function');
  var filterFile = 'filters/' + $(this).children('option:selected').data('file');

  yepnope({
    test: window[filterFunction],
    nope: filterFile,
    complete: function() {
      // ...
    }
  });
});
```

14 YepNope complete

We can attach a function to the 'complete' property of YepNope, which will run immediately if the filter function is defined, or after the 'nope' JavaScript file has loaded. This is the perfect way to safely run code that has dependencies.

```
function KaleidoscopeFilter() {
  this.name = "Kaleidoscope";
  this.defaultValues = {
    angle : 0,
    rotation : 0,
    sides : 3,
    centerX : 0.5,
    centerY : 0.5
  };
  this.valueRanges = {
    angle : {min: 0, max: 360},
    rotation : {min: 0, max: 360},
    sides : {min: 1, max: 30},
    centerX : {min: 0.0, max: 1.0},
    centerY : {min: 0.0, max: 1.0}
  };
  if(!FilterUtils){
```

15 Looking at the filter files

If we open up some of the JavaScript files from the 'filters/' folder, we can see that the filters all follow a similar pattern. Some filters need arguments passed to them to make them work (for example, RGBAdjust requires an RGB value to work with).

Each filter object defines a function called 'valueRanges()' for retrieving a list of options, including 'min' and 'max' boundaries for the value. They also define a 'defaultValues()' function that gives a standard value for each of these options.

```
complete: function() {
  var selectedFilter = eval('new ' + filterFunction);

  for(var property in selectedFilter.valueRanges) {
    var label = property.charAt(0).toUpperCase() + property.slice(1);
    selectedFilter.valueRanges[property].label = label;
  }
}
```

16 Accessing filter options

Inside our YepNope 'complete' function, we can create a new object from the filter name and then loop through all the option properties returned from the 'valueRanges()' function. We also add a 'name' and 'label' to each of these, based on the property name.

```
complete: function() {
  var selectedFilter = eval('new ' + filterFunction);

  for(var property in selectedFilter.valueRanges) {
    var label = property.charAt(0).toUpperCase() + property.slice(1);
    selectedFilter.valueRanges[property].name = property;
    selectedFilter.valueRanges[property].label = label;

    var defaultVal = selectedFilter.defaultValues[property];
    selectedFilter.valueRanges[property].value = defaultVal;
  }
}
```

17 Getting the default value

Still in our loop, we can use the 'defaultValues()' function together with the name of the property to retrieve a suitable value for our slider to start on. We add this to the property that we're looping through so that all the data is in one place.

```
<form>
</form>
</div>
</div>
<script type="text/mustache" id="property-template">
</script>
</body>
</html>
```

18 Mustache template

To get all this data into the DOM, we're going to create a simple Mustache template in our HTML file. We enclose the template in a `<script>` tag to stop it from being rendered as HTML. It's given a type of 'text/mustache' to stop it from being interpreted as JavaScript.

“This is the perfect way to safely run code that has dependencies”


```
<script type="text/mustache" id="property-template">
  <div class="control-group filter-property">
    <label class="control-label" for="filter-property-{{ name }}">{{ label }}</label>

    <div class="controls">
      <div id="filter-property-{{ name }}"
        class="slider"
        data-name="{{ name }}"
        data-min="{{ min }}"
        data-max="{{ max }}"
        data-default="{{ value }}">
    </div>
  </div>
</script>
```

“With these options available we can apply an almost unending number of filter combinations to our images”

19 Slider markup

Our Mustache template creates the basic markup for a jQuery UI slider (a <div> is all that's required). We use the HTML5 'data-' attributes to add the various option properties (min, max, name and default value) to this element so they can be accessed later.

```
complete: function() {
  var selectedFilter = eval('new ' + filterFunction);
  var propertyTemplate = $('#property-template').html();

  // Remove existing filter properties
  $('div.filter-property').remove();

  for(var property in selectedFilter.valueRanges) {
    var label = property.charAt(0).toUpperCase() + property.slice(1);
```

20 Retrieving the template

Back in our YepNope function, we can now load this template from the <script> tag by targeting the ID and using the 'html()' function. We also add a line of jQuery to remove any existing 'filter-property' markup so we don't mix options from different filters.

```
selectedFilter.valueRanges[property].name = property;
selectedFilter.valueRanges[property].label = label;

var defaultVal = selectedFilter.defaultValues[property];
selectedFilter.valueRanges[property].value = defaultVal;

var html = Mustache.render(propertyTemplate, selectedFilter.valueRanges[property]);
$('#form').append(html);
```

21 Rendering HTML

Inside the loop, we can use Mustache's render() function to apply the template to the current option property that we're working with. We then append this to the DOM at a suitable place.

```
$('#div.slider').each(function() {
  $(this).slider({
    min : $(this).data('min'),
    max : $(this).data('max'),
    value : $(this).data('default')
  });
});
```

22 Enabling sliders

After the loop, we can select all the 'slider' divs that we've added to the DOM and enable them as jQuery UI sliders. The 'data-' properties that we defined are used here to give each slider its individual minimum, maximum and default values.

```
$('#filterApply').click(function(e) {
  e.preventDefault();

  var filterFunction = $('#filterSelect').children('option:selected').data('filterFunction');
  var canvas = document.getElementById('canvas');
  var context = canvas.getContext('2d');
  var data = context.getImageData(0, 0, canvas.width, canvas.height);
  var filter = eval('new ' + filterFunction);

  var options = getFilterOptions();
  filter.filter(data, options);
  context.putImageData(data, 0, 0);
});
```

23 Applying the filter

When applying the filter, we retrieve the 'filterFunction' of the currently selected <option>, create a new object from it and then call the 'filter' function on the canvas's '2D' context. We pass in options that we get from a 'getFilterOptions()' function.

```
function getFilterOptions() {
  var options = {};

  $('#div.slider').each(function() {
    var key = $(this).data('name');
    var value = $(this).slider('value');
    options[key] = value;
  });

  console.log(options);

  return options;
}
```

24 The filter options

This function loops through all slider elements and builds a map of the option name (held in a 'data-' attribute) and the current value of that slider. With these options available we can apply an almost unending number of filter combinations to our images.

Cross-platform mobile application development

Don't aim for a percentage share of the mobile market with your apps – take on every major platform to truly maximise your earning potential. Here's how...

Advisor

Kunal Deo is a veteran open source developer. Currently he is leading two open source projects: WinOpen64 and KUN Wiki. He is also a KDE developer. He has contributed to many open source projects, including KDE-Solaris, Belenix and Openmoko



As consumers we generally hate to see any market segment ruled by a monopoly. The more competing platforms there are in the market, the more

innovation and choice can be found. As developers, though, this can be something of a mixed blessing. On the positive side, you are not tied to a single platform and you can easily move over to the 'next big thing'. The downside is that you will never be able to reach 100 per cent of market unless you put many, many extra hours into developing your application for all those platforms concurrently.

The ability to reach out to a large segment of the market is often very important, especially if you are new developer. A sound strategy might not be to support every single platform out there (it's largely impossible anyway), but to get yourself onto all the major players and grab as big a slice of the pie as is realistic.

In the mobile space the three biggest players are Android, iOS and Windows Phone 7. Even if you can market your app to the first two platforms, you are in an incredibly strong position. Going the extra mile to extend it to the third platform will give you diminishing returns (because Android and iOS have such a firm grip on the market), but it puts you in a commanding position with market reach that many developers simply can't achieve.

While a few of you may be outraged not to see your favoured mobile platforms listed here (such as Symbian, Windows Mobile, BlackBerry, JavaME etc), please note that we have limited scope in this article.

Basic guidelines

There are a few things that you need to pay attention to when building cross-platform mobile

applications. These aren't platform specific and are even valid when you're targeting just one platform. In fact, most can also be seen as just good development practices.

Modular application design: Almost all the modern mobile platforms provide object-oriented APIs, so there is no reason for not building a modular application. The key to modern application development is reusability. You should design your application in a way so that each of the main parts of the application can be separately developed and, when required, can be used in other projects as well. A very popular software design pattern that is extremely popular with all the major mobile platforms (ie Android, iOS and Windows Phone) is MVC (model-view-controller). MVC allows you divide your application into core tiers so that it can be developed and reused easily. We will look further into MVC later in this article.

Know your screen: This is not a problem with iOS and Windows Phone as they both offer only one resolution throughout their phone platforms. So, even though the screen size may alter, the resolution for the most part stays the same. With Android it is a different situation: you have got different screen sizes, different display resolutions and different pixel densities. So if you are planning to bring your application to the Android platform, you need to plan out in advance how you want to deal with less or more screen real estate. Nobody likes interfaces that are too zoomed out or zoomed in.

Syncing is the key: If you have worked hard and developed a really great app for all the platforms but they cannot talk to each other, then your app is only half useful. You need to make sure that your apps on all the platforms are able to sync data so that if a user is accessing your application from different devices their personal data is available everywhere, irrespective of the device they are using. This goes a long way in providing a seamless experience across multiple devices.

Resources

Official SDKs:

Google Android SDK: <http://developer.android.com/sdk/index.html>

Apple iOS SDK: <http://developer.apple.com/devcenter/ios/index.action>

Microsoft Windows Phone SDK: <http://create.msdn.com/en-us/resources/downloads>

Cross-platform SDKs:

MonoTouch: <http://xamarin.com/monotouch>

Mono for Android: <http://android.xamarin.com/>

Titanium Mobile: <http://www.appcelerator.com/products/titanium-mobile-application-development/>

“If you are planning to bring your app to Android, you need to plan how you want to deal with less or more screen real estate”

Find out the highest common denominator: When implementing special features such as hardware-accelerated graphics, always-on background data or multi-threaded features, your app should run well on all platforms. This does not mean dumbing down the app, but making sure that you use relevant hardware features on all the supported platforms to make your application perform great on all the supported devices.

Don't be an alien: When developing UIs across multiple devices, you do not need to make them look alike. Each mobile platform comes with its own set of UI widgets and UI design guidelines, so make use of them. This will make your application user experience look good and immersive instead of being a cheap port.

Abstraction is the key: Your app should be as abstract as possible – not only in terms of development (explained in first point) but also with its secondary portions, such as images, strings, data etc. You should be able to replace these different parts when moving from platform to platform. Most of the time, implementing abstraction only takes a careful application design; other times you may have to cleverly rely on other technologies such as Lua,

“A sound strategy might not be to support every single platform out there, but to get yourself onto all the major players”

which can be easily embedded into any application to handle some crucial tasks for which you will not need to make any changes when you want it to port to other platforms.

Approaches

The general rule of life is also true in mobile application development. The more the difficult path is, the better the results will be. This is not to discourage you, but just to set your expectations straight. In this section we will look at the two approaches for building cross-platform applications.

Approach 1: Adopting model-view-controller architecture

Model-view-controller (MVC) is not an SDK or toolkit that will help you with porting your app.

It is a software architecture principle. The basic principle of MVC is the isolation of user interface, application logic and data logic from each other so that each can be developed, tested, maintained and ported separately. Therefore, with this approach you will have to develop each application natively using the official SDKs, but the good news is that implementing MVC architecture is easy and it is supported on all the major mobile platforms: Android, iOS and Windows Phone.

Components of MVC architecture

Model: The model is aware of all data that needs to be displayed. It also knows all the operations that can be applied to transform objects. The model represents the data and the business logic that facilitates the flow of data into an application. The key point is that the model is not aware about the presentation of the data.

View: The view represents the presentation of an application. It asks the model for the data but it is not dependent on it; ie the view continues to exist even if the business logic changes.

Controller: The controller is responsible for intercepting the requests from the view and passes it to the model for the appropriate action. After the action has been taken on the data, the controller is responsible for directing the appropriate view to the user. In GUIs, the views and the controllers often work very closely together.

MVC implementation in Android

Android's MVC components can be loosely connected to the following:

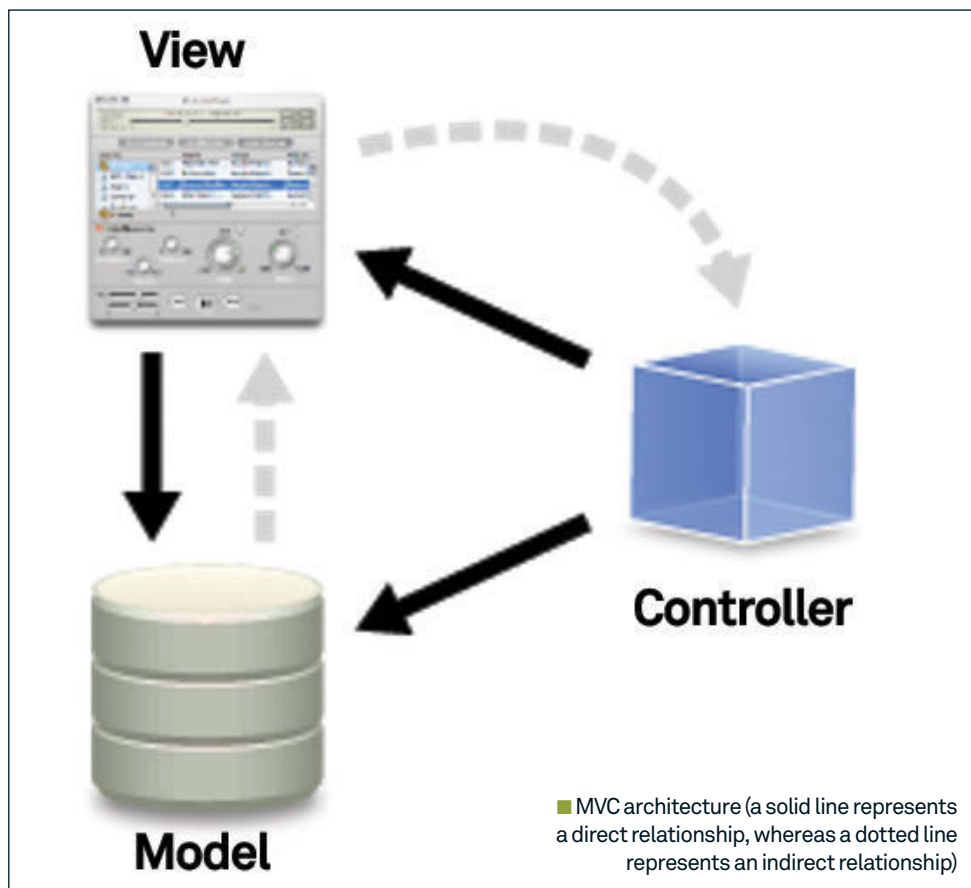
Model -> Subclass of Application

View -> Subclasses of View

Controller -> Activity

MVC implementation in iOS

Cocoa Touch (and Cocoa) uses the MVC design throughout its implementation. By default most of the iOS project templates a view (Interface files) and controller files (files with names ending with ViewController). You will need to create a model file separately.



MVC pros and cons

Pros:

- Your application will be native on each platform.
- Because of the MVC architecture, parts of your application can be easily used on other platforms or in other applications.
- Common data and controller logic results into a more stable and high-performance application.
- You will get to use native tools and SDKs on each of the supported platforms.

Cons:

- Not really a porting technique. You will still need to develop code separately for each platform.
- MVC architecture cannot be applied effectively to certain types of application.

MVC implementation in Windows Phone

MVC pattern implementation in Windows Phone is provided via third-party frameworks. One of the most popular frameworks is called Windows Phone MVC (<http://windowsphonemvc.codeplex.com/>). Windows Phone MVC uses commonly identified file suffixes to differentiate between various MVC components. Windows Phone MVC follows the following conventions:

Default Controller – **Home**


Default View – **Main**

View Location – **Namespace must end with ControllerName.ActionName**

Windows Phone MVC is tolerant to common suffixes, so `MainView`, `MainDialog`, `MainPage` etc will all be resolved from the `Main` action.

`ViewModels` are passed into the `Action` result methods, so there is no convention needed for them – you can put them wherever you want.

MVC is very helpful in porting code as you have a clear idea of which parts of the code (or logic) need to be rewritten for the new platform. It may not help you in targeting multiple platforms from one source tree, but it is very helpful when you are porting your applications natively on all the major platforms. Also, remember that native applications are always the best applications.



■ Example of an iOS project; notice the MVC pattern right in the project filenames

Approach 2: Using native cross-platform SDKs

If using MVC and writing native code on each platform is too much work for you, using a native cross-platform SDK is your next best bet. Native cross-platform SDKs generate native code for the supported platform. This means your apps will be able to make use of native UI elements and certain hardware features. One of the biggest advantages is that your apps will be sold in the same marketplace where the

native applications are sold. In this section we will take a look at such SDKs.

MonoTouch and Mono for Android

Mono is a popular open source implementation of the Microsoft .NET framework and runs on a wide variety of platforms including Linux, Mac and Windows. MonoTouch and Mono for Android are implementations of Mono on the iPhone and Android platforms. MonoTouch and Mono for Android are commercial products developed by Xamarin

“Native cross-platform SDKs generate native code for the supported platform. This means your apps will be able to make use of native UI elements and certain hardware features”

“MonoTouch and Mono for Android are capable of converting C# code into native iOS or Android objects”

■ Mono for Android website

■ Titanium project wizard

languages. It later converts HTML and JavaScript into native iOS or Android code. Unlike other web-based SDKs, Titanium does not rely on Web UI (or other web rendering engines); instead it generates mapping for native UI components, which provides a far better user experience and overall look and feel. One of the best features of Titanium is its extensibility. It allows developers to code modules in Objective-C or Java to extend the core functionality of the application.

If you are already an expert in HTML and JavaScript and are interested in developing native applications for both Android and iOS, Titanium Mobile is a good choice.

Pros and cons of native cross-platform SDK

Pros:

- Near native application performance.
- Choice of alternative programming languages.
- Choice of alternative integrated development environment.
- Single sourcing for multiple platforms.

Cons:

- Features incomplete compared to the platform-specific SDK.
- Risk of being invalid if banned by the platform vendor.
- Difficult to debug.
- Costly. Sometimes you may need to pay for both the native and the cross-platform SDK.

(co-founded by Miguel de Icaza, also the co-founder of the GNOME Desktop project). MonoTouch and Mono for Android allow developers to reuse their engine and business logic code across all mobile platforms, swapping out the user interface code for a platform-specific look and feel.

In its current state, MonoTouch is a fairly updated product with the support of iOS 5, while the Mono for Android still depends upon the older Android 2.2 SDK. Nevertheless, both versions are capable of converting C# code into native iOS or Android objects. Since they both support C# as a primary language, which is also the native programming language for the Windows Phone platform, it is very easy to share the code with that platform as well.

Titanium Mobile

Titanium is a mobile application development platform which allows developers to produce native applications for iOS and Android from a single source. Titanium uses web technologies such as HTML and JavaScript as the primary programming

The market for smartphone (and other mobile device) applications is growing day by day. The chances are that it will even become bigger than that for traditional computing devices. By making your application available on all the major mobile platforms, you will secure a decent piece of the pie.

You might be wondering why we didn't cover HTML5 in this article. Well, while HTML5 is gathering a lot of traction as the next-generation web platform, it is still a web platform. That's not to say it is no good for mobile applications, but we wanted to cover just the native application platforms for this feature. HTML elements are covered elsewhere in the book though.

Cross-platform mobile web application development

Advisor

Kunal Deo is a veteran open source developer. Currently he is leading two open source projects: WinOpen64 and KUN Wiki. He is also a KDE developer. He has contributed to many open source projects, including KDE-Solaris, Belenix and Openmoko



The mobile web is very powerful. Read on to find out how to leverage it



Not too long ago the mobile web was just about how dumb you could make your existing apps so that they'd work on the small, under-powered devices

which people used to have. To serve the very purpose of creating simple but stupid pages, WML and WAP were created, which meant your gorgeous website would just become a small white page with a collection of links. A lot has changed since then, a hell of a lot. So much so, in fact, that the mobile web is now one of the leading lights in the innovation of web standards.

The mobile web is currently dominated by WebKit-based web browsers. Both Android and iOS use WebKit-powered browsers, accounting for more than 90 per cent of the total market share of internet surfers on mobile devices. WebKit itself is an open source product which was originally created for the KDE project (as KHTML, part of the KDE's web browser Konqueror) and was later used by Apple to build the Safari web browser. Currently WebKit is developed by its initial creators, including KDE and Apple as well as other companies such as Google and Nokia. Since WebKit has an extremely large market share on mobile devices, this presents a standard platform to target. Do not worry, it is not Internet Explorer. WebKit is open source and is developed by multiple organisations (even competing ones) so it will always be on the cutting

edge when it comes to supporting standards. In fact, as of today html5test.com (an HTML5 test site) rates WebKit-based browsers higher on HTML5 standards than other web browsers.

Why a web app?

Before we get into this, let's get one thing clear. Native applications are not necessarily better than web apps. Native applications are the best way to create applications if you care deeply about the performance and the overall user experience. It takes more time, talent and resources to develop native applications and for each platform you develop for you'll have to do most of the work again (but you can be smart about it and save yourself a lot of time and hassle by planning ahead). Having said that, web apps are important too. Not everybody wants to learn a new programming language each time they develop a mobile application. Web apps provide a unique set of advantages that make web app development more feasible. Web apps are perfect in the following scenarios:

- You already have skills in web technologies and are now interested in building applications for these hot new smartphones.
- You do not have the time and resources to develop separate apps for each platform and want a consistent and easy way to develop apps for multiple platforms from a single codebase.
- You do not want to go through the slow and sometimes painful app store review process to make your app available in each platform's store and you want full control on how your apps get distributed and monetised.
- You already have a lot of web-based apps for the desktop and now you are trying to build a mobile version and interfaces for those services.

Are web applications really that powerful?

Yes they are. HTML5 has redefined what a web app can do. HTML5 (and its related technologies) has brought many new features which were previously only possible on desktop applications. Let's take a look at a few of the important features of HTML5 the make web apps not only powerful but also easy to program:

Databases: HTML5 provides a database API called IndexedDB (also known as Indexed Database API). This API is used for creating and working with a database of records holding simple values and hierarchical objects. Each record consists of a key and some values. Moreover, the database maintains indexes of the records it stores.

Local Storage: Local Storage provides a Key-Value based data storage system. It typically allows you to store 5MB of data on the client side, but this can be increased by asking permission from the user.

Offline Support: People expect apps to work even when they are not connected to the internet. HTML5 helps developers to write web apps which can work without a network connection. It often uses Manifest files for caching and syncing data for offline use.

Web Workers: Web Workers is a way of introducing multithreading inside a web application. It allows Java Script code to execute in the background independently of other, user-interface scripts that may be executing. This helps keeping the main part of the app working without having an unresponsive UI.

Web Sockets: This is defined as a full-duplex single socket connection over which messages can be sent between client and server. The Web Sockets standard simplifies much of the complexity around

Resources

HTML5 W3C spec (developing):

<http://www.w3.org/TR/html5/>

WebKit open source project:

<http://www.webkit.org/>

HTML5 online test:

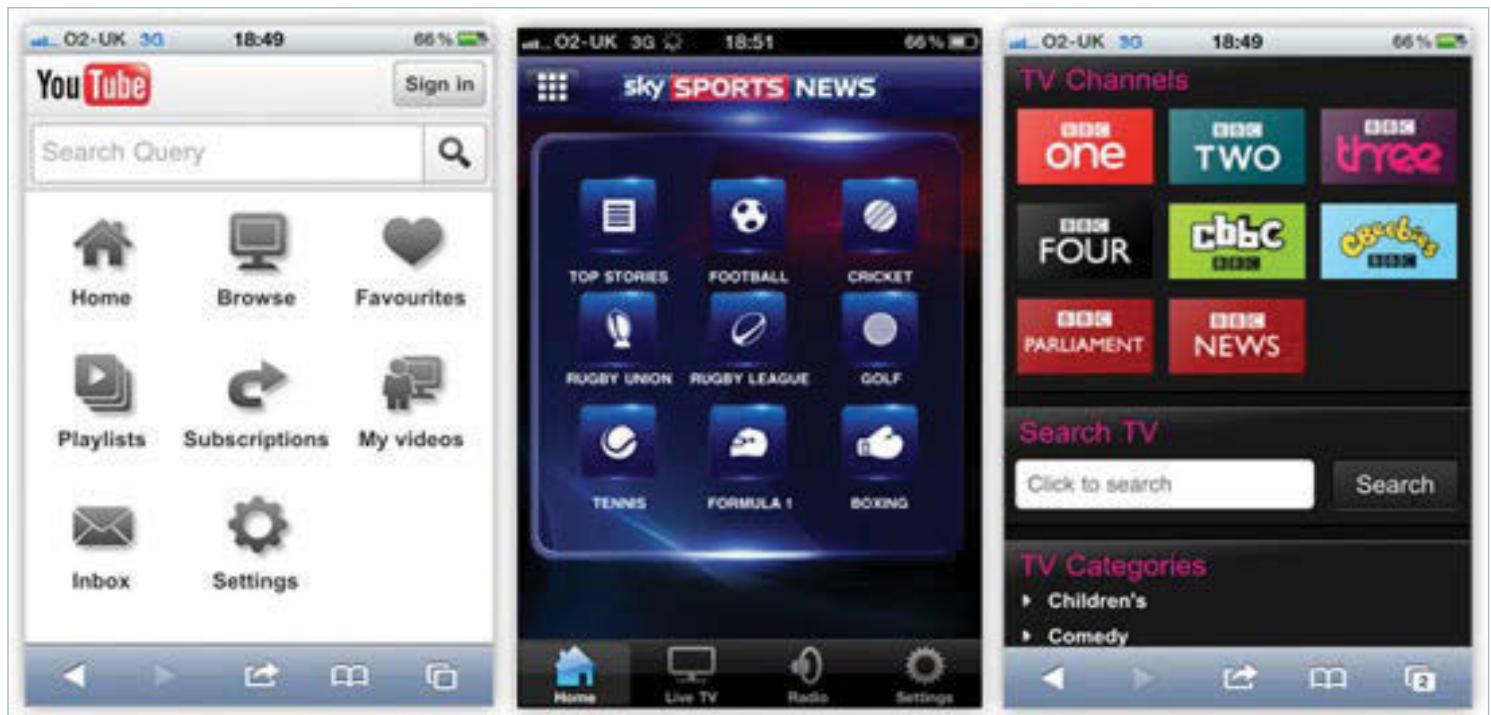
<http://www.html5test.com/>

Modernizr JavaScript library:

<http://www.modernizr.com/>

Sencha Touch: <http://www.sencha.com/products/touch>

“Make sure your app continues to work even when the device is offline”



bi-directional web communication and connection management. Web Sockets allows web apps to communicate more efficiently.

Multimedia Support: HTML5 implements many audio and video tags. This enables the playback of various media files without the need for third-party plug-ins.

Canvas: Canvas allows developers to create dynamic, scriptable rendering of 2D shapes and bitmap images. This enables an intuitive and beautiful user interface for web apps.

Basic guidelines for developing cross-platform mobile web apps

There are a few things that you should keep in mind while developing cross-platform web apps. The following are a few guidelines which will help you develop better apps:

Standards are not enough: Even though HTML5 has gained wide adoption in modern platforms, it is still far from being finalised. It is evolving at a very fast pace and so features change all the time. Furthermore, browsers can choose which standards they will implement. Choose the target platforms wisely; if your target platforms are Android and iOS, you are safe as they both use WebKit-based browsers and have minor feature differences.

Find out what works and what doesn't: As already mentioned, HTML5 is an evolving standard. Therefore it is important to know beforehand about the features that your app will be using. You can use <http://html5test.com> to test overall capability of



the target browser. You can also use an API called Modernizr (modernizr.com) to test for a specific set of features from your app itself.

Develop an app, not a site: While the line between a good HTML5 site and a web app is becoming blurred these days, it is important to keep in mind that you are competing with native applications. So keep the perspective of an app while developing applications.

Gear up for offline: A network connection is not always available on mobile devices, so make sure your app continues to work even in the absence of an active internet connection.

In addition to the guidelines mentioned here, you should also follow those mentioned in the feature 'Cross-platform mobile application development' in issue 109. Guidelines in that article include Modular

Application Design, Know Your Screen, Syncing is the Key, Find out the Highest Common Denominator, Consistent UI and Abstraction.

Building a cross-platform web app

While developing full web apps is certainly possible, the task of porting one to different platforms is a tiresome job. You will also need to track various platform differences yourself and write a lot of boilerplate code. This is where mobile web app frameworks come into the picture. They make developing web apps fun, as they abstract most of underlying API into easy-to-use features. You can also use mobile web app frameworks as bridges which enable some of the native features of the platform to be used inside the web apps. Some mobile web app frameworks also allow you to wrap your web apps into native applications so that they can coexist with the native applications inside the platform's app store.

In this section we will be building a simple Twitter search app called LUDTSearch using Sencha Touch. This application is based on the examples shipped by the Sencha Touch framework.

Sencha Touch

Sencha Touch is an open source (GPL v3) MVC-compatible web framework for mobile applications. It makes extensive use of HTML5 to deliver a native-like application experience on modern touch-enabled mobile devices.

The following assumes that you have a working web server set up and know the web server document directory. If you do not have a working web server set up, please consult your web server documentation on how to do so.

Okay, that's enough talk - let's build the LUDTSearch app.

1. Setting up the project directories

LUDTSearch is an MVC (model-view-controller) app. Accordingly, the app consists of the following directories:

```
LUDTSearch -----> Root Directory of the
App
app ----> Main App Directory
  |----controllers -----> Controller
Directory
  |----models-----> Model Directory
  |----stores-----> Store
Directory
  |----views-----> Views
Directory
lib ----> Library directory, contains
custom and third party libraries
resources ----> Contains CSS, Images and
other resources
```

The primary library is contained in the file `sencha-touch.js`. This file and the related resources can be obtained from www.sencha.com/products/touch/download/. Copy `sencha-touch.js` and the `resources` directory from the downloaded `sencha-touch-1.1.1.zip` file and put them inside the `lib` folder of the app.

For this app we are using a simple Twitter Search API to search Twitter. A proxy to this API is contained in the `TwitterProxy.js` inside the `lib` folder of the app. (Check this book's disc for full source code).

2. Bootstrap

This is where we will set up our app. This file will have import declarations for the library and will define the application's model, view, controller and stores.

@code excerpt: index.html

```
.....
<!-- Place your view files here
-->
<div id="sencha-views">
```

```
<script type="text/javascript"
src="app/views/Viewport.js"></script>
.....
<!-- Place your model files here
-->
<div id="sencha-models">
  <script type="text/javascript"
src="app/models/Search.js"></script>
  <script type="text/javascript"
src="app/models/Tweet.js"></script>
.....
<div id="sencha-controllers">
  <script type="text/javascript"
src="app/controllers/searches.js"></
script>
  </div>
  <!-- Place your store files here
-->
  <div id="sencha-controllers">
    <script type="text/javascript"
src="app/stores/Searches.js"></script>
.....
```

3. Defining the app

This is where we register the application name, default URL and target. Registering the application name also creates the global variable of the same name.

@code: app.js

```
Ext.regApplication({
  name      : "LUDTSearch",
  defaultUrl : 'searches/first',
  defaultTarget: "viewport",
  icon: 'resources/images/icon.png',
  glossOnIcon: false,
  phoneStartupScreen: 'resources/images/
phone_startup.png',
  tabletStartupScreen: 'resources/
images/tablet_startup.png',
  /**
   * This is called automatically when
  the page loads. Here we set up the main
  component on the page - the Viewport
  */
  launch: function() {
    this.viewport = new LUDTSearch.
```

```
Viewport({
  application: this,
  listeners: {
    selectionchange: this.
onSelectionChange
  }
});
},
/**
 * @private
 * Handles taps on the saved searches
list
*/
onSelectionChange: function(selModel,
records) {
  var search = records[0];
  if (search) {
    Ext.defer(function() {
      Ext.dispatch({
        controller:
'searches',
        action : 'show',
        instance : search,
        historyUrl:
'searches/' + search.get('query')
      }, 10, Ext);
    }
  }
});
```

4. Setting up the URL routes

Similar to Rails, `routes.js` maps the URL to a controller action.

@code: routes.js

```
Ext.Router.draw(function(map) {
  map.connect("searches/first",
{controller: 'searches', action:
'first'});
  map.connect("searches/:query",
{controller: 'searches', action: 'show'});
});
```

5. Setting up the controller

As evident from the above code excerpt, we only have one controller in our app called 'searches'. Here we have two public actions called 'first' and 'show'. 'First' is called when a user simply navigates to the application URL; 'show' is called when the user has tapped on the search item.

@code excerpt: app/controllers/searches.js

```
first: function() {
  var store = Ext.
getStore('Searches'),
  first = store.first();
```

“We will be building a simple Twitter search app using Sencha Touch, an open source MVC-compatible web framework for mobile apps”

```

    if (first) {
        Ext.dispatch({
            controller: "searches",
            action: "show",
            instance: first,
            historyUrl: "searches/" +
first.get('query')
        });
    } else {
        this.noSearches();
    }
},
.....
....

```

6. Creating application models

Our app consists of two models: Search and Tweet. The Search model uses the device's local storage to save the user's search items. The Tweet model relies on the Twitter proxy (lib/TwitterProxy.js) to search Twitter.

```

@code excerpt: app/models/Search.js
Ext.regModel("Search", {
    fields: [
        {name: "id", type: "int"},
        {name: "query", type: "string"}
    ],

```

“HTML5 has made it possible to provide a native application experience from a web app. So if you are working on a new app, give web apps a try”

```

    hasMany: {
        model: "Tweet",
        name: 'tweets',
        filterProperty: 'query',
        storeConfig: {
            pageSize: 15,
            remoteFilter: true,
            clearOnPageLoad: false
        }
    },
    proxy: {
        type: 'localstorage',
        id: 'twitter-searches'
    }
});

```

7. Registering data store

In this step we will be defining the data store which uses the previously saved searches

and then loads them up automatically on the search bar.

```

@code excerpt: app/store/Searches.js
Ext.regStore('Searches', {
    model: 'Search',
    autoLoad: true
});

```

This store is also used during the application start-up (in app/controllers/searches.js):

```

    first: function() {
        var store = Ext.
getStore('Searches'),
        first = store.first();

```

Lastly, we need to configure views. Views in the app consist of the following files:

app/views/Viewport.js : This defines the application's shell. This shell will be used further to populate with relevant data. The shell itself doesn't change but acts as a container.

app/views/SearchBar.js : This contains a Search List instance which lists all of the saved Searches.

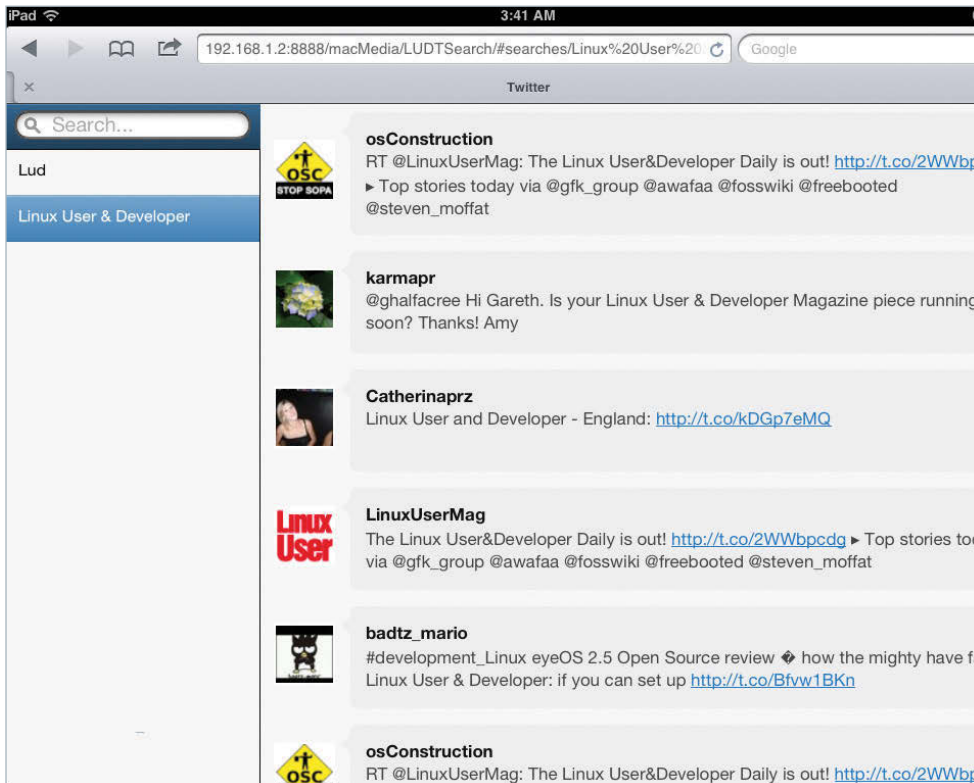
app/views/Searches.js : Shows all of the searches that the user has saved. It also contains a new search form docked to the top.

app/views/Tweets.js: This renders the tweet search results using the Tweet model.

Refer to the code directory on the book disc for full source of the application.

To run this application, copy the LUDTwitter directory to the web server document root directory. Then browse to `http://<WebserverIpAddress>/LUDTwitter` from your favourite device to run the app.

■ LUDTwitter in action on iPad 2



So that was a quick introduction to building cross-platform web apps. We also took a quick look at the Sencha Touch library – an MVC-enabled HTML5 web framework for creating web apps. As you can see, HTML5 has made it possible to provide a native application experience from a web app. So if you are working on a new app, give web apps a try; we are confident that frameworks like Sencha Touch will make the development process fun and productive.

Using Hadoop tasks on Amazon's Elastic MapReduce

Accomplish huge tasks using the power of Hadoop cluster and some home-made command-line tools...

Resources

[Amazon Web Services account](#)

Advisor

Andy Leon is a web developer with Java and PHP as his languages of choice, and a fair amount of server-wrangling as a secondary skill



The concept of MapReduce was introduced by Google back in 2004 and it's one of brilliant simplicity. By splitting data-processing tasks down into two simple functions and ensuring that each one can run independently of the other, these tasks can be scaled infinitely across many machines.

Apache Hadoop is the de facto open source implementation of MapReduce, used by the likes of Facebook, Last.fm and StumbleUpon to analyse giant datasets. Amazon's Elastic MapReduce takes Hadoop and makes it available over the firm's EC2

cloud platform, automatically provisioning machines for use with a cluster.

We're not required to use a cluster though. Hadoop Streaming jobs can be built in such a way that they can be run from the command line on a single machine. We can accomplish entire tasks with this method or use it for test runs on small sets of data.

Here we'll run through an example job flow that takes geographical data from Freebase and organises islands by where they're located. We'll run this on a single machine and then demonstrate how to run it across a cluster.

Step Name	State	Start Date	End Date	JAR	Main Class	Args
Streaming Job	COMPLETED	2012-01-09 22:41 GMT	2012-01-09 22:43 GMT	/home/hadoop/contrib/streaming/hadoop-streaming.jar	-	-input s3://mapreduce.acleon.co.uk/islands/input/ -output s3://mapreduce.acleon.co.uk/islands/output/ -mapper s3://mapreduce.acleon.co.uk/islands/scripts/mapper.php -reducer

See all your currently running job flows in one place

Create new job flows on-the-fly

Job flows run on Hadoop and so are compatible with standard options

Leave a job running and feed more data in using S3 as it becomes available

01 Getting data

To get some sample data to work with, head over to Freebase's data dumps at <http://download.freebase.com/datadumps/latest/browse/> and download 'island.tsv' from under the 'geography' folder. This contains information on all the world islands, line-by-line with tab-separated fields.

Index of /datadumps/latest/browse/geog

Name	Last modified	Size	Description
Parent Directory	-	-	-
glacier_scrinium.tsv	09-Jan-2012 02:14	81	
mountain_range.tsv	09-Jan-2012 02:14	150K	
body_of_water.tsv	09-Jan-2012 02:14	1.1M	
waterfall.tsv	09-Jan-2012 02:14	20K	
mountain.tsv	09-Jan-2012 02:14	24K	
lake.tsv	09-Jan-2012 02:14	570K	
glacier_type.tsv	09-Jan-2012 02:14	190	
mountain_pass.tsv	09-Jan-2012 02:14	31K	
lake_type.tsv	09-Jan-2012 02:14	43K	
mountain.tsv	09-Jan-2012 02:14	492K	
mountain_type.tsv	09-Jan-2012 02:14	35K	
island_group.tsv	09-Jan-2012 02:14	31K	
glacier_status.tsv	09-Jan-2012 02:14	631	
glacier.tsv	09-Jan-2012 02:14	57K	
geographical_feature_category.tsv	09-Jan-2012 02:14	44K	
geographical_feature.tsv	09-Jan-2012 02:14	1.7M	
island.tsv	09-Jan-2012 02:14	265K	

02 Finding fields

Running the 'head' utility to get the first few lines, we can see how the fields are laid out within the file. The one we're interested in is 'island_group', the fifth field in the header row (number four when zero-indexed). Use 'sed' to trim the header row.

```
$ head island.tsv
name    id      max_length  max_width  island_group  body_of_wate
Shishulok Island  /w/94azgvw  Canadian Arctic Arch
Novgorod Island  /w/94y_jw
Senja        /w/94fBoc
Two Tree Island /w/94z11q
Grand Bahama  /w/93st9j
Azaki Island  /w/93y8fz1
Nagit Island  /w/94c_y8
Pronke Island /w/948651m
Achemar Island /w/93y93f3

$ sed -i id island.tsv > island_data.tsv
$
```

03 The PHP files

With our data in place, we can start to create the job files. Touch two PHP files called 'mapper.php' and 'reducer.php'. We want to be able to run both of these from the command line as easily as possible, so set the execution permissions on them as well.

```
$ touch mapper.php reducer.php
$ chmod u+x mapper.php reducer.php
$
```

“With our data in place, we can start to create the job files”

```
1 #!/usr/bin/env php
2 <?php
3
4
```

04 Shebang

Open up 'mapper.php' and add a PHP shebang at the top of the file, and the PHP opening tag. We want our script to be as portable as possible, so we're using the 'env' version of the shebang. This means there's less chance the PHP binary will not be found.

05 Streaming data

Rather than pass a file into our script, Hadoop will split the file into parts and send those parts into the script through the standard input. So whereas you would normally expect to open the TSV file directly, we're going to open the pseudo-file 'php://stdin' to get our input.

```
1 #!/usr/bin/env php
2 <?php
3
4 $inputStream = fopen('php://stdin', 'r');
5
```

07 Checking for groups

We know from our header row that we're expecting the island groups to be in the fifth column, which is the fourth when we start from zero. We can check whether this column is empty – and if not, we call a function called emit with the value and a 1.

```
1 #!/usr/bin/env php
2 <?php
3
4 $inputStream = fopen('php://stdin', 'r');
5
6 while($line = fgets($inputStream))
7 {
8     $line = trim($line);
9
10    $values = explode("\t", $line);
11
12    if(!empty($values[4]))
13    {
14        emit($values[4], 1);
15    }
16 }
17
18
```

06 Reading lines

A while() can then be used to go through all of the available lines of data Hadoop will be supplying. Inside the loop, we can do any cleanup tasks we need to (trim in this case) and then we split the line into an array using the tab character.

```
1 #!/usr/bin/env php
2 <?php
3
4 $inputStream = fopen('php://stdin', 'r');
5
6 while($line = fgets($inputStream))
7 {
8     $line = trim($line);
9
10    $values = explode("\t", $line);
11
12    |
13 }
```

08 The emit function

Emit is not a PHP function, it's one we're going to write. Its simple job is to take any number of arguments, join all of those arguments together with a tab character and echo them followed by a new line. It outputs data in the same format we received it.

```
12 if(!empty($values[4]))
13 {
14     emit($values[4], 1);
15 }
16
17 }
18
19 function emit()
20 {
21     $args = func_get_args();
22
23     echo implode("\t", $args)."\n";
24 }
```

```
Windmill Islands 1
Windmill Islands 1
Windmill Islands 1
Windmill Islands 1
Windmill Islands 1
Windmill Islands 1
Windward Islands,Society Islands 1
Windward Islands,Society Islands 1
Zed Islands 1
Zed Islands 1
Zed Islands 1
Zhoushan 1
Zhoushan 1
Åland Islands 1
Îles d'Hyères 1
Îles des Saintes 1
Îles des Saintes 1
Öckerö Kommun,Gothenburg Archipelago 1
Öckerö Kommun,Gothenburg Archipelago 1
$
```

09 Testing the mapper

Our mapper can now be tested. This is not a large dataset, so we can run the whole job from the command line using `./mapper.php < island_data.tsv | sort`. Written to the console will be an ordered list of keys and values – one for every time a group value was found.

10 Starting reduction

The reducer is a little more complex. We start in the same way as the mapper, using standard input again to get our lines. After that, we declare a couple of variables – the current key, null as we haven't started yet; and an empty array to hold our reduced data.

```
1 #!/usr/bin/env php
2 <?php
3
4 $inputStream = fopen('php://stdin')
5
6 $currentKey = null;
7 $reducer = array();
8
9 |
```

11 Key and count

We'll use the same `while()` loop, `trim()` and `explode()` functions, but as we're only expecting two values, we can use `list()` to put them into `$key` and `$count`. The `$count` is redundant as it is always 1, but this will not be the case for every mapper.

```
4 $inputStream = fopen('php://stdin', 'r');
5
6 $currentKey = null;
7 $reducer = array();
8
9 while($line = fgets($inputStream))
10 {
11     $line = trim($line);
12     $values = explode("\t", $line);
13
14     list($key, $count) = $values;
15
16
17 }
```

12 The reducer array

To reduce the data, we look at what the key is and whether we've encountered it before. If not, we add it and the `$count` to our reducer array. If it's already in `$reducer`, we increment it by the number of the `$count`.

```
13
14 list($key, $count) = $values;
15
16 if(array_key_exists($currentKey, $reducer))
17 {
18     $reducer[$currentKey] += $count;
19 }
20 else
21 {
22     $reducer[$currentKey] = $count;
23 }
24 }
```

13 Key boundaries

Above this we add a check as to whether the new `$key` is the same as the one before. If not, we've come to the end of that particular batch, so we output the key from our `$reducer` array, using the familiar line-based output.

```
20 while($line = fgets($inputStream))
21 {
22     $line = trim($line);
23     $values = explode("\t", $line);
24
25     list($key, $count) = $values;
26
27     if($key != $currentKey)
28     {
29         if(!is_null($currentKey))
30         {
31             echo $currentKey, "\t", $reducer[$currentKey], "\n";
32         }
33
34         $currentKey = $key;
35     }
36
37     if(array_key_exists($currentKey, $reducer))
```

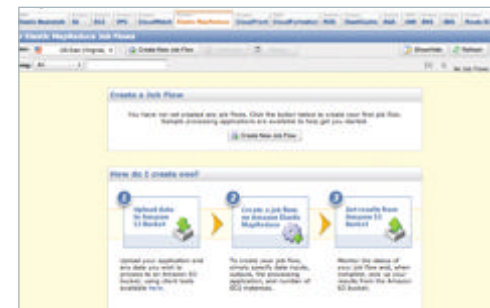
14 Final test run

Putting this all together, we can take the output from the sort, pass it to the reducer and get our processed data. Run `./mapper.php < island_data.tsv | sort | ./reducer.php` and what comes out is a list of all the different groups and how many islands are in each.

```
Virgin Islands 17
Virgin Islands,Leeward Islands 1
Virgin Islands,United States Virgin Islands,Leeward Islands 1
Viti Levu Group 9
Volcano Islands 1
Wellesley Islands 2
Windmill Islands 26
Windward Islands,Society Islands 2
Zed Islands 3
Zhoushan 2
Åland Islands 1
Îles d'Hyères 1
Îles des Saintes 2
$
```


15 Taking it to the Cloud

If we had our own Hadoop cluster, we could add our PHP files and data to a streaming job and set it running. Alternatively, we can use Amazon's Elastic MapReduce to provision a cluster for our work.

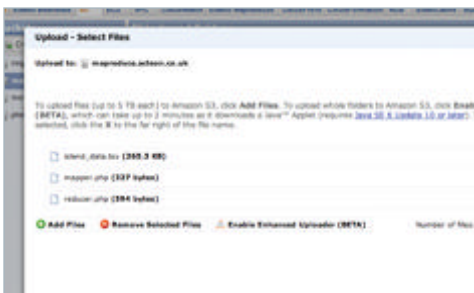


“To reduce the data, we look at what the key is and whether we've encountered it before”

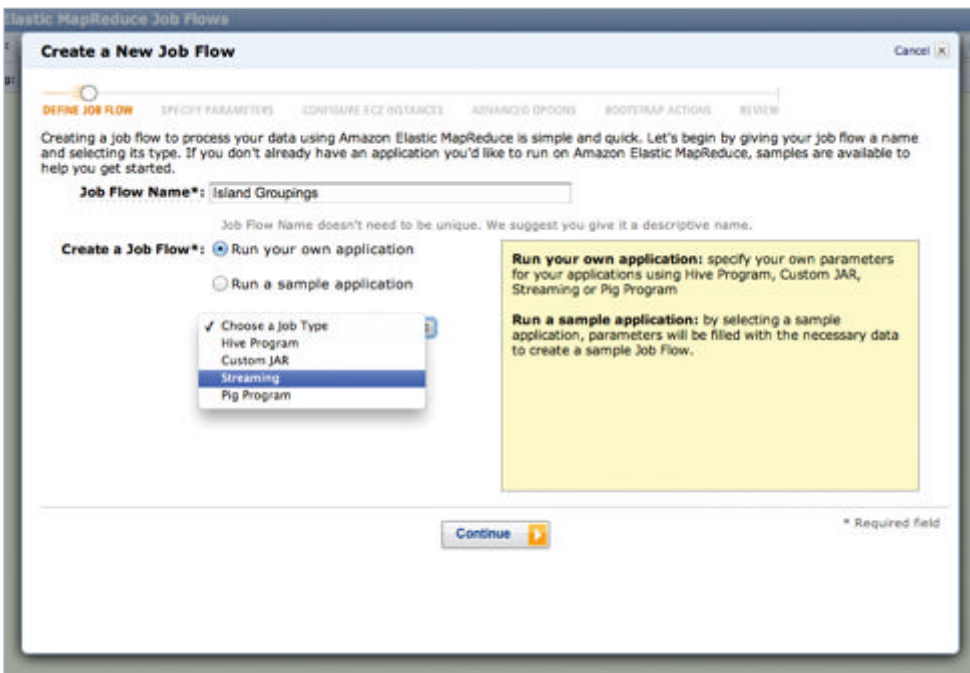
16 Uploading – part 1
Log into S3 and create a new bucket to hold your MapReduce job. In here create two folders called 'input' and 'scripts'.



17 Uploading – part 2
Upload the 'mapper.php' and 'reducer.php' files to the 'scripts' folder, and 'island_data.tsv' to the 'input' folder.

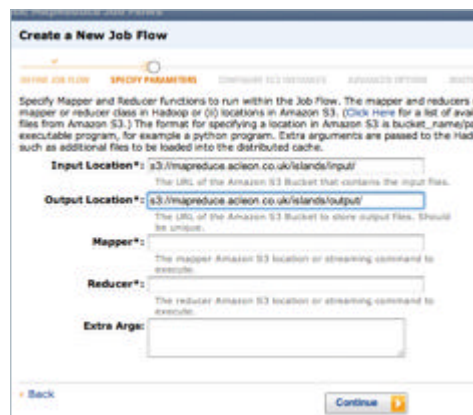


18 Creating the job
Go back to the MapReduce tab, click the 'Create Job Flow' button and enter a name for it. Select to run your own application rather than an example one, and then select Streaming as the job type.



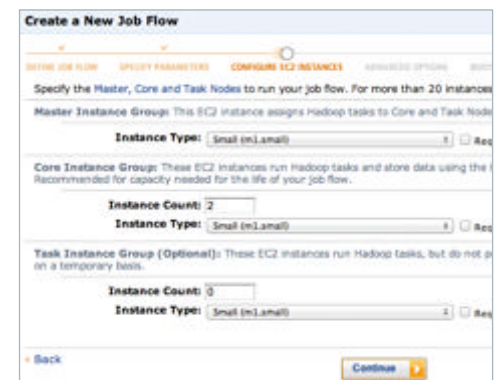
“You can secure a job flow by selecting a key pair to use with it or by attaching it to an existing Virtual Private Cloud”

19 S3 folders
On the next screen, enter the path to the input folder in your S3 bucket. Also, add the path to a folder on S3 that doesn't currently exist, to hold the output of the job flow.



20 PHP scripts
Enter the paths to the 'mapper.php' and 'reducer.php' files, uploaded to the 'scripts' folder in the S3 bucket. Double-check all of these paths as the job will fail if they're not right and need to be recreated.

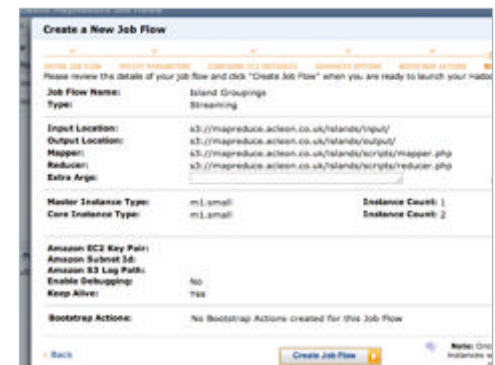
21 Instances
We can now select the type of instance to act as the master, and the type and number of core instances used for processing. There's no need to start up any Task Instances.



22 Security and logging
On the next screen, you can secure a job flow by selecting a key pair to use with it or by attaching it to an existing Virtual Private Cloud. You can also enable logging by providing an S3 path to use as a log file.

23 Bootstrap actions
Not necessary for this particular job, but if we had used external libraries not normally available on the MapReduce instances (for example, CURL), we can upload a shell script to be run before the job flow.

24 Job confirmation
On the final screen, we can review all our job flow information before starting it. Once we confirm, we'll return to the MapReduce job screen. When the job is done, our data will be waiting in the S3 output folder.



Advisor Joey Bernard



Joey has been using Linux at home for almost 20 years. Now he's getting paid for it as a computational research consultant with ACEnet

Resources

Download AIDE: <https://play.google.com/store/apps/details?id=com.aide.ui&hl=en>

Community: <https://plus.google.com/101304250883271700981/posts>

Develop on your Android smartphone with AIDE

With AIDE, you can now develop on your Android phone or tablet. Here we will walk you through writing a Hello World app on the move...



Previously, if you were interested in developing apps for the Android platform, you could only edit the source files while you were on the go. To compile and package your app, you had to move these files back onto a PC. With the introduction of AIDE, you can now do all of the steps while you're on the move. If you have Wi-Fi access, you can even handle source code revision control through Git.

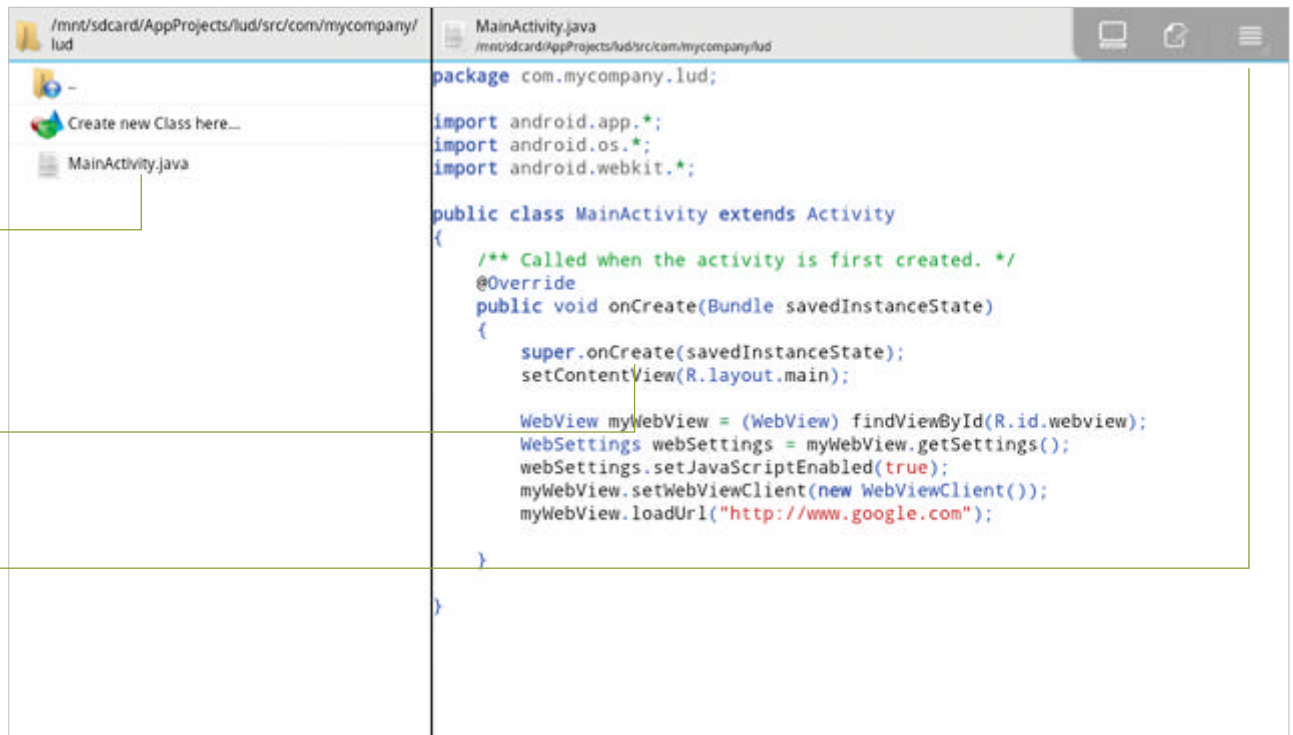
In this tutorial we'll get AIDE installed on a device, look at some of the options and go through the steps required to get your first Hello World app edited, compiled, packaged and installed. You'll see how the IDE works, and some of the great features that are provided. The team behind AIDE have been working very hard to add new features and to fix bugs, so be sure to keep updated once you start using it.



“The team behind AIDE have been working very hard to add new features and to fix bugs”

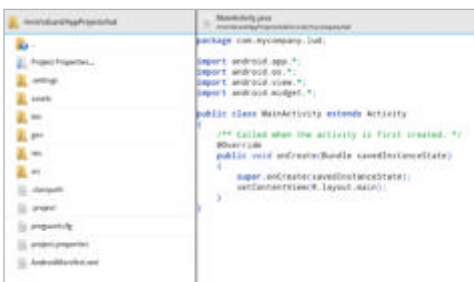
01 Install AIDE

AIDE is available through the Google App Store, recently renamed Play. AIDE uses your SD card as storage for the project files. Amazingly for a full IDE, it only takes up a little more than 10MB of space.



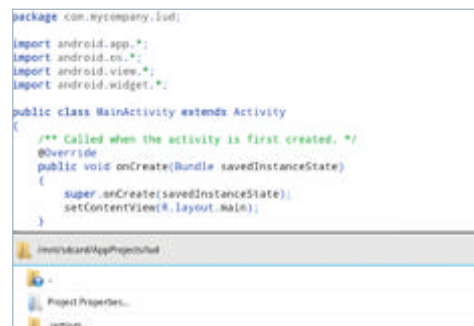
02 Create a new project

To create a new project, you start at the file listing. Clicking on 'Create new App Project here...' will bring up an option dialog where you can enter the app name and package name for your new app. There is also an option to select the type of app, but it is currently fixed to be 'activity', which is a standard app.




03 Open MainActivity.java

The default name of the main class is MainActivity. When the creation process finishes, it starts up with the main class open in the editor. You are started off with some decent boilerplate code on which to build. The main function is onCreate, which sets the contents of the main view with the function call 'setContentView(R.layout.main)'.



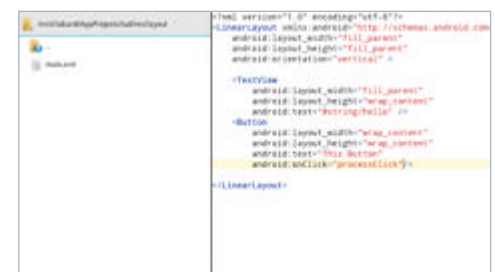
04 Change the screen layout

In the top right-hand corner, there are several icons. The first one looks like a desktop computer. Clicking this icon changes the layout of the IDE panels, from the default arrangement with a folder listing on the left and the editor on the right, to just the editor window, to the editor window at the top and the folder listing at the bottom.



05 Look at Main Layout

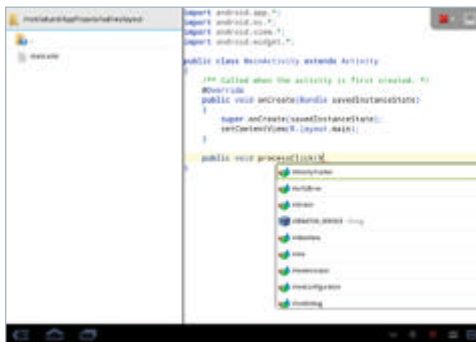
Screen layouts are defined in XML files. These files are stored in the subdirectory 'res/layout'. If you navigate there and open main.xml, you can see that the boilerplate code gives you a TextView object that displays some text. This text is stored in a string resource (res/values/strings.xml). You can edit the text that is displayed by editing this file.



06 Add a new button

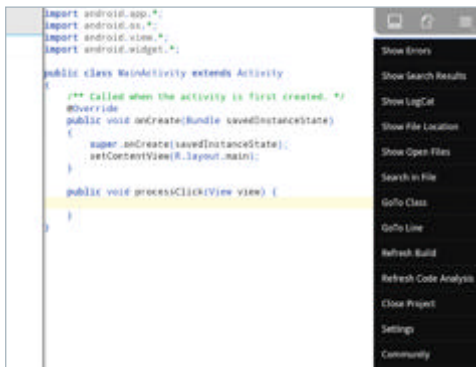
Adding a new button is as simple as adding a new entry in the main.xml layout file. You can change the layout width and height. You can set the button text here explicitly.

A better option is to keep all of your text values stored in the strings.xml file. This will make it easier if you decide to make your application multilingual.



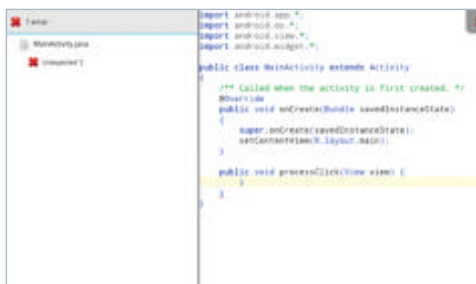
07 Autocomplete

Like any good IDE, AIDE provides autocomplete wherever possible. All of the default classes and methods should be handled correctly. You simply need to start typing and a drop-down will appear, listing all of the possible completions that AIDE knows about. You can scroll down and make your selection.



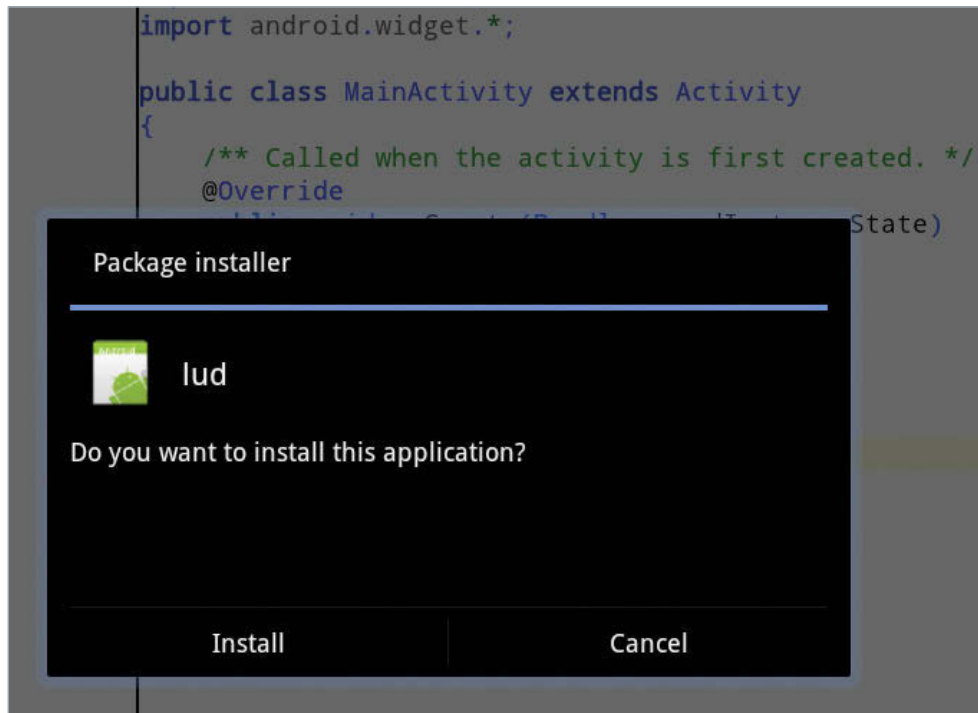
08 Get help on a function

There are several ways to get help. From the menu option, there is an entry called 'Community'. Clicking this gives you links to AIDE's Google+ page, Twitter account, Facebook page and email address. In this way you can ask the community for help. If you long-press in the editor window, you get a list of options. One of them, Help, pulls up the developer documentation from the Android site.



09 Checking out errors

AIDE is constantly checking for syntax errors. As soon as any show up, they are indicated in the top left corner by a large red X and a count of the detected errors. Clicking on the large red X will bring up a list of the errors. Clicking on an error will take you to the location in the files. You can then go ahead and make the required edits.



10 Compiling and running your app

The compile phase handles compiling all of your Java code, and bundling of all of the various resource files into a single APK file. If it runs into any compile errors, it will stop and give you an error message. Assuming that everything is successful, you'll get a dialog box asking you to install the new application. Once it is installed, give it a try and see your handiwork.

11 Where is your app?

By default, all of the files for your app are stored in the folder /mnt/sdcard/AppProjects. Each app is stored in its own subdirectory. Once you have built your application and are ready to share, you can find the APK file in the 'bin' subdirectory of your app's subdirectory.

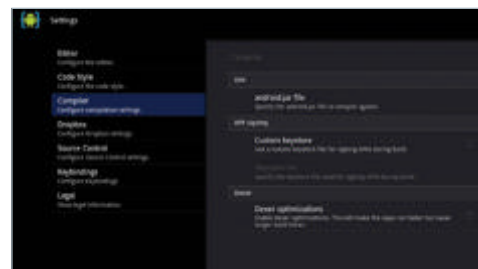
12 Compiling against a different version

By default, AIDE compiles your app against the version of Android running on your device. If you want to compile for a different version, you will need to copy over the associated android.jar file. Once you do so, you can go to the compiler settings in AIDE and set the path to that file.

“Like any good IDE, AIDE provides autocomplete wherever possible”

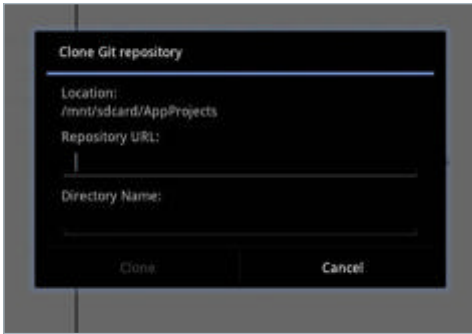
13 Signing your APK

APKs (Android application packages) need to be signed in order to validate where they come from. You can choose to use a custom keystore in order to sign the APK files that you are creating. When doing so, you can set the path to where the keystore file is stored on your device's SD card.




14 Optimise your build

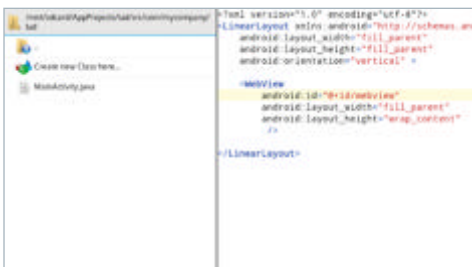
By default, AIDE performs no optimisations. This is so that you can do quick rebuilds and check your app out as you make changes. In the compiler setting, there is an option to enable dex optimisations; this will make your app faster at the cost of a longer compile run.



15 Cloning a Git repository
 One of the new features added is the ability to clone Git repositories. When you have the folder display up, there is an option to 'Clone Git Repository here...'. Clicking on this will pop up a display where you can put in the repository URL and directory name. This way, you can work on code hosted at sites like GitHub.

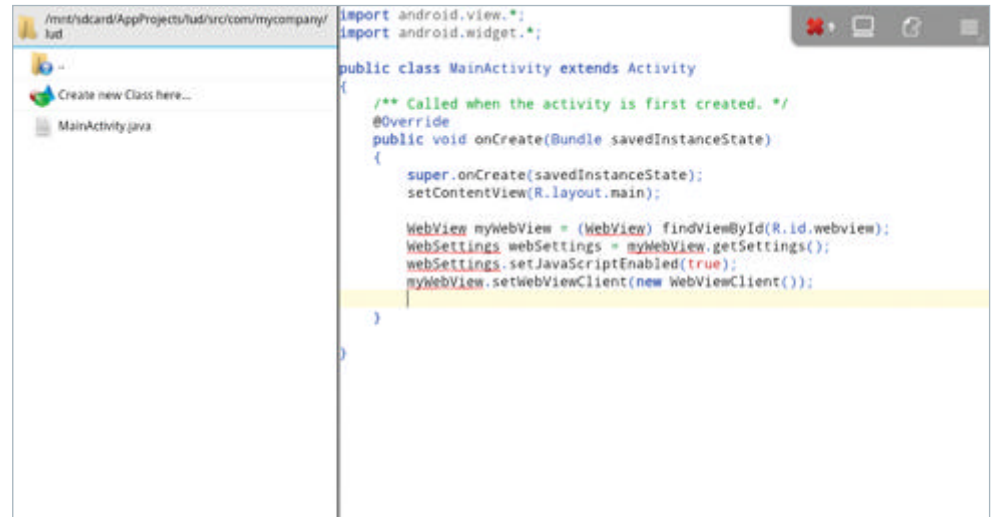


16 Settings for Git
 In the settings area for AIDE, there is a category for source control options. While in here, you will be able to set the username and user email for commits. You can also set a path for '.ssh', to store your credentials. You can also select whether or not to keep your SSH information in memory.

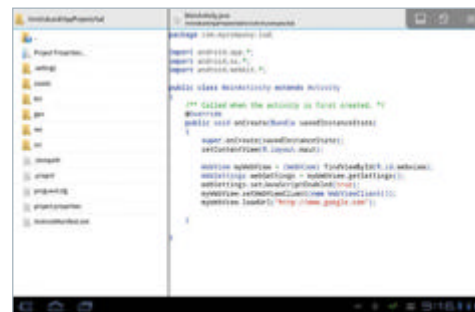


17 Adding a web object to the layout
 You can write your own mini web browser. To start off, you need to open the main.xml file and change the TextView entry to WebView. This will create a web object that will fill the complete main layout. This isn't enough, though...

“One of the new features added is the ability to clone Git repositories”



18 Adding to onCreate
 You need to add code to the onCreate method of the main class. This code creates a WebView object to display the webpage. You can change settings, like whether JavaScript is enabled or not. Check the API docs to find out what settings can be managed. If you want clicked links to open in this WebView, you'll also need to add a WebViewClient object.



19 Fixing import statements
 There is a very helpful function, 'Organize Imports'. When you start adding new code, you may forget to add the appropriate import statements. If you long-press the editor window, you will get a pop-up where you can 'Organize Imports'. This will tell AIDE to analyse your code and add in any import statements that are missing from it.



20 Loading a URL
 The WebView object by default doesn't have any buttons for navigation or URL entry. Luckily, you can add in an initial URL by using the 'loadUrl' method. You will also need to add a request for the INTERNET permission in your manifest file. This file is named AndroidManifest.xml and is located in the main directory of your app. You need to add a 'uses-permission ... />' section. You can also set the minimum version and the targeted version of Android.

Building LibreOffice extensions

How to build new features into everybody's favourite open source office suite

Advisor

Kunal Deo is a veteran open source developer. Currently he is leading two open source projects: WinOpen64 and KUN Wiki. He is also a KDE developer. He has contributed to many open source projects, including KDE-Solaris, Belenix and Openmoko



In this tutorial we will learn how to build extensions for LibreOffice. These extensions allow us to add new features to LibreOffice without changing the core source code. We will be using the LibreOffice SDK (also identified as the dev package in most of the Linux distributions) to build the extensions. We will be making use of a Python library called Python UNO to program LibreOffice.

Note: For this tutorial we are using openSUSE 12.1. But this can be adapted to any Linux distribution. You can install the SDK using the following command:

```
$ sudo zypper install libreoffice-sdk
```

This command will take care of all the dependencies in one go.

LibreOffice is the most popular open source productivity suite and comes pre-installed on most major Linux distributions. LibreOffice itself is not a new product, but a fork of the now obsolete OpenOffice. Although a raft of new developments have been integrated into LibreOffice since it forked, it is still fairly compatible with OpenOffice.

One of the biggest advantages of LibreOffice over its competition is its extensibility. Users can add features to LibreOffice without changing the core code, by using LibreOffice SDK, and it's very flexible in terms of programming language support. You can use C, C++, C#, Visual Basic or Python to build LibreOffice extensions. They say getting a practical and hands-on perspective on the technology you're trying to learn is often the best route to success, so in this tutorial we'll be using Python to develop an example extension for LibreOffice which is actually a really useful open source application that's also available to download from GitHub.

What are we building?

In the spirit of Valentine's Day, we will be building a LUD Love Letter Writer called Py LibreOffice Love Letter Writer. This truly useful tool will let you create beautiful love letters for your loved ones in no time at all. Now, the sceptics may say that programs have no heart so can't create touching love letters, but this is no ordinary program, as you'll soon see. Of course, this will be an open source program, which we've hosted on github.com, so everybody can contribute

and make it better. Share the love! Okay, let's get started...

1. Extension design

Although the Love Letter Writer (LLW) is a LibreOffice extension, it will have a separate user interface and it will be able to start and use LibreOffice as required. We will be using PyQt to design our user interface and, as such, we will focus more on Python and LibreOffice connectivity (there's a brief introduction to PyQt on [LinuxUser.co.uk](http://bit.ly/yGJywq) - <http://bit.ly/yGJywq>).

So how will this extension work? Simply speaking, the user will use our extension's user interface to fill in some basic info like the names involved and that data will be processed in Python. After that we will use the Python-UNO bridge to compose the document in LibreOffice and once that process is complete, the user can either modify the document further or print it and seal it in an envelope with a tender loving kiss (or whatever floats your boat).

2. User interface

The user interface has got to be simple and, for convenience, it would be good if the UI were to be bootable from outside LibreOffice. The extension also contains just one window to keep things simple. This is where we'll collect all the required information, such as the sender's name, recipient's name and some other fun things. In case you're wondering, the interface is being designed using Qt Designer.

3. Building core functions (coreFunctions.py)

This is where we'll see the heart of the application. Here we will be building core functions that we will use in LLW. Please note that these function names are not actually valid Python function names – they just represent the broad functionalities. You can try entering Python lines directly into a Python shell to experiment, but just keep track of objects you are creating.

startLibreOffice ()

We need to start LibreOffice in such a way that it's able to make a connection using URP (UNO Remote Protocol). This protocol is used to communicate UNO (Universal Network Objects) calls across processes.

Resources

LibreOffice core packages: Including Writer, Calc, Impress, Draw, Base, Math etc. We may not need all of them, but it is nice to have all the libs installed at once

LibreOffice SDK: Often identified as dev packages. If you are not able to locate the dev or sdk package, you can download install it from <http://www.libreoffice.org/download/>. You may need to download additional packages in order for the packages to work successfully, so installing from the package manager is way easier

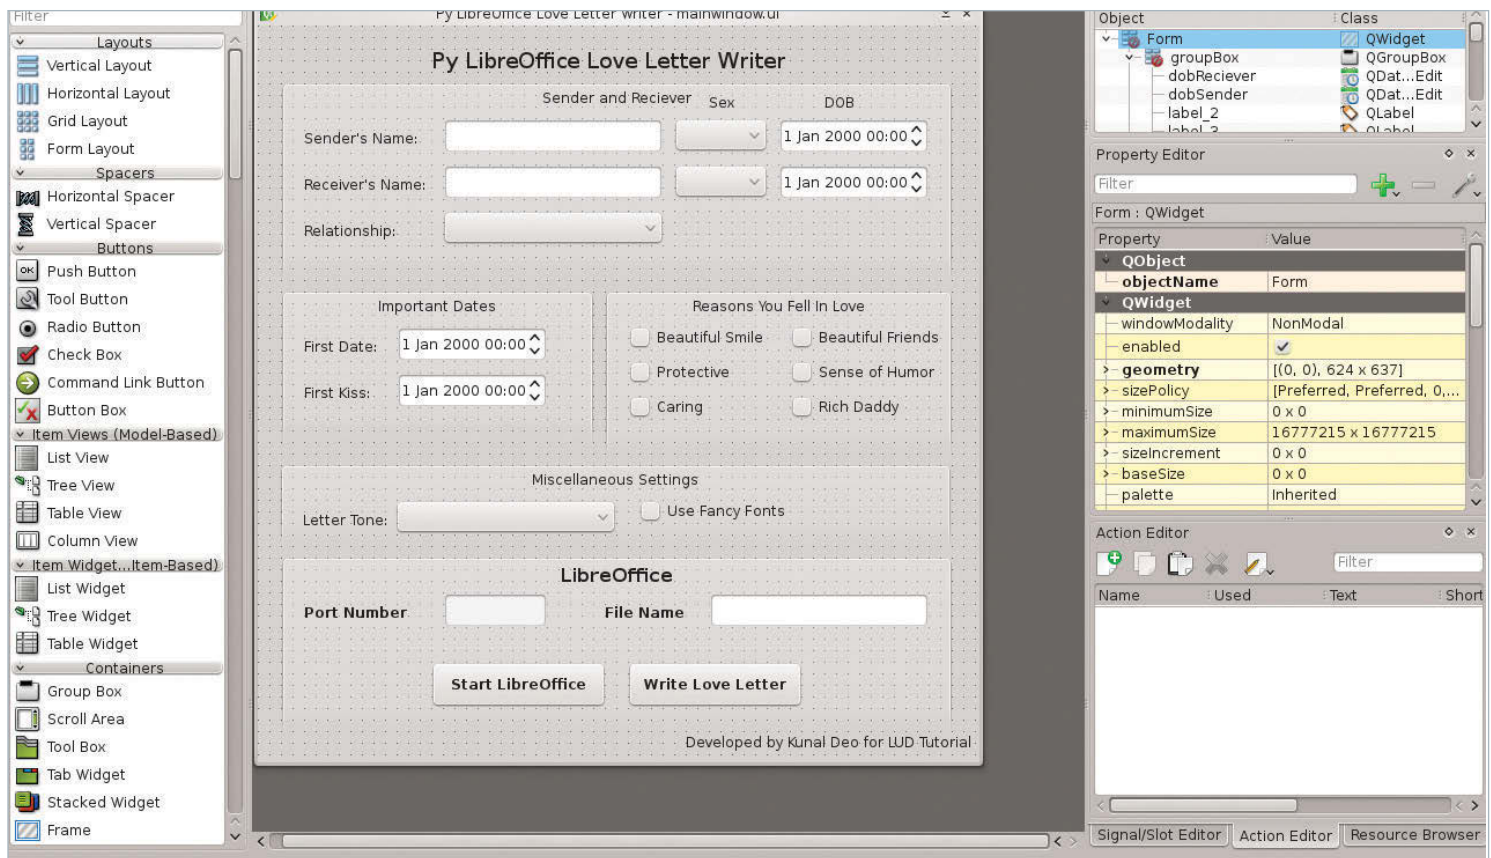
Mandatory development tools: These include GCC, G++, Make, Python, JDK etc

Python-UNO bridge for

LibreOffice: This allows us to use the standard LibreOffice API from Python

PyQt4 development libraries

“We'll be using Python to develop an example extension for LibreOffice”



The following command will start the LibreOffice **in URP mode:**

```
$ soffice "--accept=socket,port=2002;urp;"
initializeLibreOfficeConnection ()
```

Now we need to connect to the urp service that we have just started and set up the overall API.

```
import uno
# If this fails, you may need to set
PYTHONPATH environment variable as
# /usr/lib/libreoffice/basis3.4/program
(modify the path as per your installation)
local = uno.getComponentContext()
# here we are obtaining the component
context which we will use to access other
un
# components
resolver = local.ServiceManager.
createInstanceWithContext("com.sun.star.
bridge.UnoUrlResolver", local)
# here we are accessing the UnoUrlResolver
service. This will allow us to make
connect
# to LibreOffice.
```

```
context = resolver.resolve("uno:socket,h
ost=localhost,port=2002;urp;StarOffice.
ComponentContext")
#here we are hooking up the context and
connection to LibreOffice.
#remember port number should be the same
as of startLibreOffice().
#Now we are ready to make API calls
desktop = context.ServiceManager.
createInstanceWithContext("com.sun.star.
frame.Desktop", context)
```

Here we are loading the service com.sun.star.frame.Desktop which represents the environment for components. Think of it as the default LibreOffice service which we will use to create, load and manage documents.

```
createNewDocument()
This function will be used to create a new
blank document.
document = desktop.loadComponentFromURL("p
rivate:factory/swriter", "_blank", 0, ())
```

saveDocument (filepath)

This function will be used to save the document.

```
def
document.storeAsURL(filepath,())
# file path should be a valid URI, for
example a local file can be represented as
# file:///home/kunal/myfile.odt
```

openDocument (filepath)

This function will be used to open a file.

```
document = desktop.
loadComponentFromURL(filepath,"_blank", 0,
())
getTextCursor ( )
This function will return a cursor that we can use to
write/edit text and modify text properties.
cursor = document.Text.createTextCursor()
# this call will fail if you do not have a
document open.
```

insertText (text)

This function will insert text into the document at the current cursor location.

```
document.Text.insertString(cursor, text,
0)
insertParaSpacing ( )
```

You can use C-style character sequences as well. For example, 'n' for new line, 't' for tab spacing etc.

“This will be an FOSS program, which we’ll be hosting on github.com”

```
document.Text.insertString(cursor,
"\n\n\t", 0)
```

changeFont (fontName)

This function will be used to change fontName. Note that this and other cursor-related functions (mentioned here) will only set the property for text to be created. If you want to modify existing text, you will need to select it first and then change the property.

```
cursor.setPropertyValue("CharFontName",
fontName)
```

To get a property you can use `getPropertyValue`. For example, `'cursor.getPropertyValue("CharFontName")'` will return the current font name.

changeFontSize (size)

On the same lines we will be using `'cursor.setPropertyValue'` to set the font size.

```
cursor.setPropertyValue("CharHeight",
size)
```

For more granular control over font appearance, you may want to set the `CharWeight` property as well.

Based on the above discussion we can now come up with `coreFunctions.py`.

```
def initializeLibreOfficeConnection():
    local = uno.getComponentContext()
    resolver = local.ServiceManager.
createInstanceWithContext("com.sun.star.
bridge.UnoUrlResolver", local)
    context = resolver.resolve("uno:socket
,host=localhost,port=2002;urp;StarOffice.
ComponentContext")
    desktop = context.ServiceManager.
createInstanceWithContext("com.sun.star.
frame.Desktop", context)
    return desktop
def createNewDocument(desktop):
    document = desktop.loadComponentFrom
URL("private:factory/swriter", "_blank",
0, ())
    return document
def saveDocument(document, filepath):
    document.storeAsURL(filepath, ())
    def openDocument(desktop, filepath):
        document = desktop.
loadComponentFromURL(filepath, "_blank", 0,
())
        return document
def getTextCursor(document):
    cursor = document.Text.
createTextCursor()
    return cursor
```

```
def insertText(text, document, cursor):
    document.Text.insertString(cursor,
text, 0)
```

```
def insertParaSpacing(document, cursor):
    document.Text.insertString(cursor,
"\n\n\t", 0)
```

```
def changeFont(fontName, cursor):
    cursor.setPropertyValue("CharFontNa
me", fontName)
```

```
def changeFontSize(size, cursor):
    cursor.setPropertyValue("CharHeight",
size)
```

4. Main module (main.py)

`main.py` is the entry point for the extension. The primary purpose of this is to initialise the main window with initial data. It also serves as two primary purposes: accepting the user data and sending it to another module called `loveLetter.Writer`. `Write ()` to prepare the LibreOffice document.

First you need to convert the `.ui` file (created using Qt Designer – on the disc it is named as `mainwindow.ui`)

to Python:

```
$ pyuic4 mainwindow.ui -o mainwindow.py
```

In the same directory, create a file named `loveLetterWriter.py` with the following contents:

```
from PyQt4 import QtCore, QtGui
import sys
import os
import subprocess
import loveLetterWriter
#import UI
from mainwindow import Ui_Form
class Main(QtGui.QMainWindow):
    def __init__(self, parent=None):
        QtGui.QWidget.__init__(self, parent)
        self.ui = Ui_Form()
        self.ui.setupUi(self)

        # Populate default values
        self.ui.portNumber.setText("2002")
        self.ui.relationship.addItem('Wife
Husband Boyfriend Girlfriend'.split())
        self.ui.senderGender.addItem('Male
Female'.split())
        self.ui.receiverGender.
addItem('Female Male'.split())
        self.ui.letterTone.
addItem('Romantic Casual Comedy
Apologetic'.split())
```

```
self.ui.fileName.
setText("loveLetter.odt")
#Connect Buttons
QtCore.QObject.connect (self.ui
.startLibreOfficeButton,QtCore.SIGNAL
("clicked()"),self.startLibreOffice)
QtCore.QObject.connect (self.
ui.writeLoveLetterButton,QtCore.SIGNAL
("clicked()"),self.writeLoveLetter)
def startLibreOffice(self):
    portNumber = str(self.ui.portNumber.
text())
    #commands.getoutput('soffice "--acce
pt=socket,port='+portNumber+';urp;")
    os.system('soffice "--accept=socket,
port='+portNumber+';urp;" &')
    def writeLoveLetter(self):
        #Get current values
        senderName = str(self.ui.senderName.
text())
        receiverName = str(self.
ui.receiverName.text())
        senderGender = str(self.
ui.senderGender.currentText())
        receiverGender = str(self.
ui.receiverGender.currentText())
        relationship = str(self.
ui.relationship.currentText())
        fileName = str(self.ui.fileName.
text())
        loveLetterWriter.write(senderName,receiverName,receiverGender,rel
ationship,fileName)
    def main():
        app = QtGui.QApplication(sys.argv)
        window = Main()
        window.show()
        sys.exit(app.exec_())
if __name__ == "__main__":
    main()
```

5. Love Letter Writer module (loveLetterWriter.py)

This module is called by the Main module to write the letter. It uses `coreFunctions.py` to connect to LibreOffice and perform its functions. To keep things simple we are only building a simple letter. But as with all open source software, you can always **extend it to suit your needs.**

```
import coreFunctions
from os.path import expanduser
from PyQt4 import QtGui
def write(self, senderName, receiverName, se
nderGender, receiverGender, relationship, fi
leName):
    desktop = coreFunctions.
initializeLibreOfficeConnection()
    document = coreFunctions.
createNewDocument(desktop)
```

“Building extensions for LibreOffice is surprisingly easy”

```

    cursor = coreFunctions.
getTextCursor(document)
    coreFunctions.
changeFont('Courier',cursor)
    coreFunctions.
changeFontSize(15,cursor)
    coreFunctions.insertText("My
Dear "+relationship+"
"+receiverName,document,cursor)
    coreFunctions.insertParaSpacing(docume
nt,cursor)
    coreFunctions.insertText("I am looking
forward to holding you once again in my
loving arms. You are beautiful. You are
mine.",document,cursor)
    coreFunctions.insertText("\n\
n",document,cursor)
    coreFunctions.insertText("From a
loving and longing heart",document,cursor)
    coreFunctions.insertText("\n\
n",document,cursor)
    coreFunctions.insertText(senderName,do
cument,cursor)

#Save the file
file = 'file://'+expanduser('~')+'/' +
fileName
    coreFunctions.
saveDocument(document,file)

    QtGui.QMessageBox.warning(self,"I
nformation","Letter has been saved at
"+file,QtGui.QMessageBox.Ok)

```

6. Running the extension

You extension is finished now. You can now run the following command to **start the extension:**

```

$ cd <project dir>
$ python main.py

```

Enter all the details. Click on 'Start LibreOffice'. Once LibreOffice has started, enter a filename and click 'Write Love Letter'. Your letter will be displayed inside LibreOffice. Congratulations, you have made it this far. Now you can modify the code (especially LoveLetterWriter.py) to generate more intelligent letters.

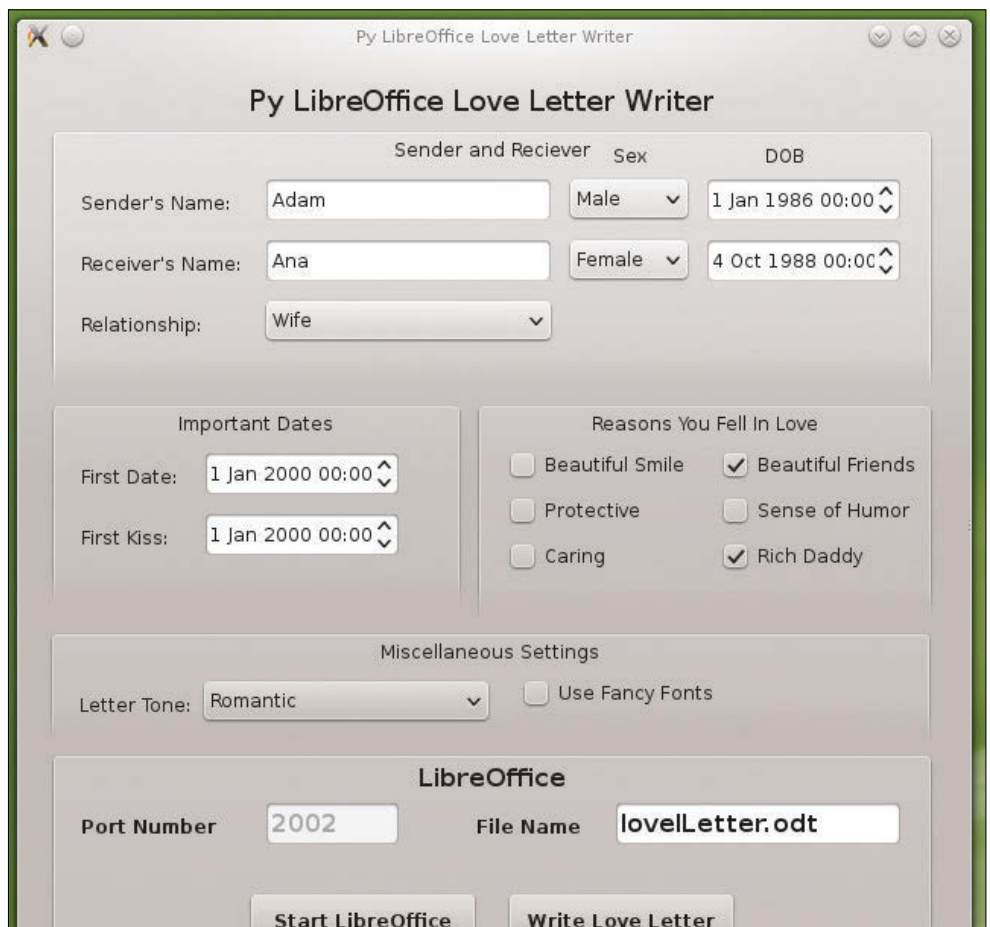
LibreOffice is an excellent office productivity suite and building extensions for it is surprisingly easy. Don't forget, this is now a real-world open source extension which is available for everyone to contribute to. You are welcome to participate via GitHub at the following address: <https://github.com/kunaldeo/Py-LibreOffice-Love-Letter-Writer>.

“We're only building a simple letter”



■ LLW output

■ LLW in action



Essential Sof

The best distros and projects

190 The best desktop distros

A look at the ten essential distros that could change the way you work

196 Torrent clients

Investigating the best way to get the most out of your file sharing

200 Media centres

Featuring top programs to help you enjoy your media

204 Mind-mapping software

Plan your thoughts with these four mind-mapping kits

208 Project managers

Handy business tools that will help improve your productivity

212 CD rippers

Copy your discs onto your PC with no trouble at all

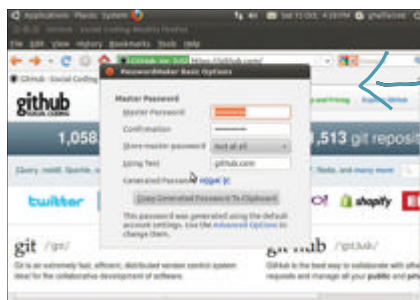
216 Email clients

Find out which open-source email clients offers the best features

220 Password managers

Keep track of passwords and let your computer do all the hard work for you

“Open-source software continues to go from strength to strength, with some great distros out there”



tware



THE BEST DESKTOP DISTROS



Join us as we look at the movers and shakers of the Linux world, to see which represent the brightest lights for the future...



Linux and open-source software continue to go from strength to strength. But over the past few years or so, a lot has changed – and for some, it hasn't always been for the better.

There are many familiar friends to be found in our list of hottest distros, but there are also a couple of newcomers that you may not have tried. Thanks to

some interesting changes at the top of the charts, there's also an unexpected champion – but saying any more would spoil the surprise.

A Linux distribution is, as always, a personal thing. While we'd like to think we've been fair and even-handed, it's perfectly possible that you think we've overlooked something important in a reviewed distribution – or even missed out your

favourite distro entirely. If so, please do let us know! We're always open to feedback, and we'd love to hear from you if you've got a strong opinion about what makes a good distro great. If we've missed a particularly impressive distribution, we'll be sure to give it some love in a future book – so don't be shy to put your opinions forward over at www.linuxuser.co.uk.



Mageia

A newcomer to our charts, Mageia's popularity is growing rapidly

10

Pros

For users accustomed to the way Mandriva does things, switching to Mageia is straightforward

Cons

Thus far, the project has done little to distinguish itself from its parent distribution

A fork of Mandriva Linux, born in 2010 when Edge-IT went through liquidation and many of the project's developers were laid off, Mageia is proving to be an increasingly popular option for those who want something familiar yet with a new and vibrant community behind it.

Mageia has attracted a large number of Mandriva developers to its ranks, and many more have committed financially to the future of the distribution – at the last count, around 270 people had donated to the running of Mageia.

In the nine months that it took to create Mageia 1, the first full release, the team has achieved much: KDE 4, GNOME 2, Xfce 4 and LXDE spins are all readily available, and the project's repositories contain much of the common software required on a desktop or laptop machine.

“Forking an existing open source project is never an easy decision to make, and forking Mandriva Linux is a huge task,” the project's founders

admitted when Mageia was first announced, but the increasing popularity shows that it's an approach which has paid off.

The Mageia project still has a lot of work ahead, however, not least of which will be committing to a regular and predictable release cycle to ensure its users aren't caught high and dry.



■ Mageia 1 KDE spin running a calculator, playing a video and showing application menu

Key features


- Mageia retains the familiar Mandriva tools, like urpmi and drakxtools.
- Spins of most common desktop environments are already available, including GNOME2 and KDE4.
- The project's founders hope to provide a more stable, predictable fork of Mandriva.
- Mageia's goals include the targeting of new architectures – with ARM support due soon – and improving integration between desktop environments and applications.
- Mageia is already proving a popular choice for Mandrake/Mandriva fans.

The facts

Mageia was born from Mandriva, when its parent company was going through a period of flux.

The first release, Mageia 1, was very popular, paving the way for the release of Mageia 2.

'Mageia' – more properly written 'μαγεία' – is the Greek word for 'magic.'



PCLinuxOS

Despite being held back by outdated package management, it remains popular

9

Pros

A good-looking distro with some excellent software and great support material

Cons

Poor naming conventions, confusing choices and difficult package management may put users off

Bill 'Texstar' Reynolds may have started PCLinuxOS as a spin-off of Mandrake, but the distribution has grown rapidly since 2003 into a fully fledged distro in its own right.

The latest version, PCLinuxOS 2012.02, brings a range of improvements along with one or two disappointments. In the latter camp is the insistence on keeping on the now outdated

2.6 kernel tree rather than the more up-to-date 3.x – although the inclusion of a BFS scheduler and PAE support as standard is welcome – along with the use of the slightly clunky Synaptic package manager.

That's not to say that the past year hasn't brought a raft of improvements: the new brushed-metal theme is modern and sleek, and the use of KDE 4 in the standard spin has also given birth to the 'Fullmonty' edition – a special spin featuring virtual



desktops dedicated to particular tasks, including internet use, office work, multimedia and games.

The PCLinuxOS community is still small, but tight-knit and humorous – Reynolds once told us “unfortunately, we can't come up with anyone sane enough to talk to a reporter” when asked for a quote – making it a great distribution to get involved in.

The facts


PCLinuxOS, aka PCLOS, was born out of Mandrake Linux by Bill Reynolds in 2003.

While primarily based on KDE 4, various spins are available – including the 'Fullmonty' one.

Unlike its fellow Mandrake spin-offs, PCLOS uses APT package management via Synaptic.

Key features

- The 'mylivecd' script takes a snapshot of the current environment, making it easy to create custom live CDs.
- The APT-based repository includes 12,000 RPM packages covering most common requirements.
- Like the smaller Puppy Linux, the PCLOS live CD can be copied into RAM for diskless use.
- PCLOS includes impressive localisation features, supporting over 85 languages.
- While the default version uses KDE 4, spins with Enlightenment, Xfce, LXDE, GNOME ZenMini, OpenBox, and GNOME are available.



In review

PCLinuxOS is packed with some excellent software choices and a decent array of support options. However, some of the design decisions spoil any hope of it properly catering for beginners, effectively making the 'Radically Simple' catch-line slightly misleading and damaging its chances of breaking into the mainstream. There's a wonderful distro hiding in there, but still much work to be done.
Linux User & Developer 9/13/5



Pardus Linux

8

Growing in popularity, this Turkish distribution goes from strength to strength

Pros

The developers of Pardus go out of their way to make things easy for newcomers

Cons

The repositories continue to grow, but still lag behind the more well-known distributions in quantity

While Pardus was originally developed as a native Turkish Linux distribution, its popularity outside its home country – where it finds use as the de facto distribution for both government and military projects – is growing rapidly.

The Pardus Migration Partner project, launched in 2008, continues to push large-scale migrations in its native country,

while those turned off by recent changes to the user interfaces of more famous distributions will likely be pleased by the familiar yet polished KDE desktop.

Although the current stable release of Pardus dates back to February, the system feels up to date. The inclusion of non-free packages like Flash and MP3 codecs also makes the system easy to use for newcomers to Linux, but does bring the drawback of making the distribution unsuitable for free software purists.

The Pardus 2011 release features many improvements over the early beta version, including

expanded hardware support, the option to load either proprietary or open graphics drivers for AMD and Nvidia users, and simple command-line configuration of both wired and wireless networking.

Pardus would rank higher in our list if it weren't for one thing: with the majority of its developers working in Turkish, it can be a struggle for English-speaking contributors to integrate fully with the Pardus development community.



Key features

- The PiSi package manager supports ratings, screenshots, and package deltas for fast updating.
- A better than average installer – YALI, is a dream to work with, if a trifle slow.
- The Kaptan configuration tool, launched on first

login, makes desktop customisation easy for beginners.

- The KDE-based interface is clean and tidy, with plenty of scope for customisation.
- Many proprietary drivers and packages, such as Java and Flash, are installed as standard.

The facts

A project of the Turkish National Research Institute of Electronics and Cryptology.

It was originally designed as a native Turkish fork of Gentoo, but is now a discrete distro.

'Pardus' is the Latin name for the Anatolian leopard, which is pictured on the distro's logo.



Despite its Turkish origins, Pardus is well localised for others

In review

Pardus 2011 is shaping up to be one of the most exciting distro releases of the year. It's a must-try distribution for distro hoppers, but Pardus needs more software packages in its repositories to gain traction with regular Linux users. *Linux User & Developer* 95 4/5



Puppy Linux

7

Always a great choice for older systems, Puppy Linux continues to improve

Pros

Puppy is great for older hardware, and its simple UI is easy to use even for beginners

Cons


Even when you've decided to use Puppy, you then need to choose which of several versions to install

Puppy Linux is a truly astounding project, booting from a sub-130MB ISO directly into RAM to create one of the most feature-rich distributions around for older hardware – or those who just like to get the most speed possible out of their systems.

With project founder Barry Kauler still working on one of Puppy's facets – the long-term support version, known as Wary Puppy – the distribution shows no clear sign of slowing, and offers a snappy experience on modern hardware while still running more than acceptably on even outdated systems.

Puppy's organisation is somewhat chaotic. "Puppy has never been run as a 'proper' project," Kauler has admitted. "It has always been ad-hoc." Despite this, the distro continues to receive frequent updates and to grow its fan base.

Puppy Linux is unique in our list by being, in fact, three very different distributions: users are asked



to decide between Wary Puppy, the Slackware-compatible Slacko and the Ubuntu-compatible LuPu. The two binary-compatible distributions can install software from the mainstream parent's repositories, while all three support Puppy's own package format.

While some see this split as a disadvantage, the level of choice offered is often one of the biggest draws towards Puppy as a distribution.

The facts

Puppy was originally developed by Barry Kauler and first released in 2003.

Built to run entirely from RAM, some versions will boot on a system with as little as 48MB.

The name 'Puppy' is in memory of Kauler's dog, a surprisingly brave chihuahua.

Key features

- Designed to run entirely in RAM, Puppy can be used to boot diskless workstations.
- Built-in write-caching technology means that Puppy won't excessively write to flash storage devices if used as a live USB.
- If packages aren't available for Puppy directly, LuPu and Slacko

can use Ubuntu and Slackware repositories respectively.

- Wary, the long-term support version, will run on almost any x86 system you could reasonably still want to use.
- Puppy is one of the smallest distros to include a full graphical user interface.



Supports packages from many different distributions



In review

Despite substantial changes under the hood, the new version of Puppy Linux still remains the lightweight, fast and user-friendly distro we've come to love. Although the Puppy Linux 5.0 desktop has not changed much, it does sport a few improvements here and there. Add to this the ability to install software from Ubuntu repositories and you have yet another excellent release of Puppy Linux. *Linux User & Developer* 95 4/5



Arch Linux

Not the easiest distro around, but it is one of the most flexible and powerful

6

Pros

Arch is fast, flexible and allows those with the know-how to create a lightweight, custom platform for their needs


Cons

If you're not a Linux pro, you might find the lack of GUI and relatively high learning curve a struggle

Since its creation in 2002, Arch Linux has stayed true to its roots as a minimalist distribution designed for the more knowledgeable Linux user – and while it's not the most user-friendly distribution around, it is one of the most flexible.

The most recent release, Arch 2012.07.15, adds a wide variety of improvements including the Linux 3.x kernel tree, support for Btrfs and NILFS2 file systems, and Syslinux bootloader support. The project's rolling release schedule also keeps things up to date over time, but often comes as a shock to those used to distributions with a more regimented release schedule.

Arch isn't for everyone, however: the distribution's insistence on providing a bare-bones console-based environment which can be customised by the user – defended by project lead Aaron Griffin as being “better to be technically elegant with a higher learning curve than to be easy to use and technically



■ You don't have to love the console to use Arch: graphical tools are available to install

crap” – makes it a poor choice for newcomers to Linux.

Despite this, it's a clear choice for inclusion in our top ten simply because of its power and flexibility: if you know what you're doing, Arch can be shaped into precisely what you require from a distribution with no preconceptions from the project's maintainers getting in the way.

Key features

- Arch is a completely independent distro, designed to offer a minimalist platform for customisation.
- A custom package manager called Pacman takes care of both binary packages and package builds.
- While its minimalist approach makes it popular for servers, Arch

has a growing following of desktop users.

- A 'rolling release' update cycle means that an Arch system can always be updated, and never gets retired.
- Most configuration is done at the shell, although third-party graphical interfaces for tasks like package management exist.

The facts

Arch was created as a minimalist distro in 2002 by Judd Vinet, and is currently led by Aaron Griffin.

Arch is designed as the basis for a custom platform – the only distro in our top ten not to include a default GUI.

The distro's fans take pride in the fact that Arch isn't the most user-friendly distribution out there.



Debian

A true Linux granddaddy, it is the base on which many other distros are built

5

Pros

A clear focus on free software – though a non-free repository exists – gives Debian the moral high ground

Cons

The distribution feels a little 'outdated' compared to some of its derivatives

Although its numerous child distributions – which include the popular Ubuntu and Linux Mint distributions – have become some of the most popular around, plain old Debian still has its fans.

Launched in 1994 by Ian Murdock, Debian is still going strong. The most recent release, Debian 'Squeeze' 6.0.5, marks

a drastic change for the project: in addition to the use of the Linux kernel, Squeeze introduces support for the BSD kernel for the first time, with Hurd microkernel support expected to follow in short order. Those opting for the Linux kernel will find that it is also completely free for the first time, which can cause issues with certain proprietary hardware which expects to find non-free 'binary blob' modules.

Many of the criticisms brought against Debian when the Ubuntu project was founded still hold true

– in particular a somewhat slow release schedule. It's something the team is looking to improve, however: rather than the traditional two-year schedule used from 'Potato' to 'Squeeze,' the next release – named 'Wheezy' – is due out in 2012, just one year after Squeeze was released.

"Debian is committed to 100 per cent free software," project lead Stefano Zacchiroli asserts – a focus that has directly resulted in a disagreement with the Mozilla Foundation, causing the Debian project to fork off the browser and email client projects Iceweasel and IceDove from Mozilla's Firefox and Thunderbird respectively.



Key features

- Debian's main focus is on truly free software, offering non-free packages only under sufferance.
- Uses the popular APT package management system, with Synaptic for GUI users.
- Squeeze includes an optional BSD kernel

– and future releases will feature the Hurd microkernel.

- One of the oldest-running distributions, with a long and distinguished history.
- Directly responsible for the creation of around 130 spin-off distributions.



■ The live CD lets you check compatibility

In review

A good choice if you want to run a server or a desktop with stable packages and a couple of years of security patches, even more if your hardware is from an exotic architecture. The downside is that some packages are a bit too old...
Linux User & Developer 95
3/5

The facts

Debian gets its name from project founder Ian Murdock and his then-girlfriend Debra Lynn – 'Deb-lan.'

First released in 1994, Debian is one of the oldest distributions around and has over a hundred spin-offs.

Includes its own set of guidelines for devs: the Debian Social Contract and the Free Software Guidelines.



openSUSE

The one-time Novell-sponsored distro now has a new home under Attachmate

4

Pros

Tumbleweed offers a way to keep on the bleeding edge between releases, and openSUSE's KDE spin is undeniably pretty

Cons

Certain high-profile packages like WordPress and Dropbox aren't available in the default repositories

Along with Debian, openSUSE is one of the oldest distributions in our list - and it's seen plenty of change in that time.


For years, openSUSE was sponsored by Novell, which sold a commercial version known as SUSE Linux Enterprise. When Attachmate Corp bought Novell in late 2010, however, it decided to sever the ties and give the project the support it needed to become a business unit in its own right.

The release of openSUSE version 12.1, promised to bring some serious improvements over its predecessor: support for the Btrfs file system is included as standard, while the Xorg configuration tool SaX 3 looks promising - if a little buggy in its beta format.

"OpenSUSE is a tool that lets me do most things without having to get my hands too dirty at the command line," project member Andrew Wafaa told us. "I'm now a fan not only because of the freedom and power I have, but also because of this

green lizard," he quipped, referring to the distro's gecko mascot.

"Why do I use openSUSE? Because of all the power behind a few funny words like Zypper, KIWI, WebYaST, LibreOffice and more," added openSUSE chairperson Alan Clark. "Every release just keeps getting easier to install with greater stability and new features."



■ Package management via YaST2 isn't particularly pretty, but it works well

Key features

- Tumbleweed makes it easy to stay on the bleeding edge, with regularly updated packages.
- The KDE Plasma interface is one of the best we've seen, with great visuals.
- Advanced users are backed up with a great toolchain, including a

Build Service and the SUSE Studio site.



- 32-bit users have support for 4GB of RAM straight out of the box, with no PAE installation required.
- Customisation is easy, with the ability to roll your own live CDs with custom content if you want.

The facts

Originally a restricted commercial release, it's one of the oldest distros around.

Under Attachmate, openSUSE is a business unit that's separate from Novell.

Features a green gecko, known as the Geeko, as its mascot.



In review

An exciting release if you want to play with the newest Linux tech. The Btrfs file system and snapshots - to revert accidental changes - are interesting features. The new Xorg configuration tool SaX3 also looks interesting, but the developers still have to iron out some kinks.

Linux User & Developer 95
4/5



Fedora

Fedora has come a long way from its Red Hat roots, and stays at the bleeding edge

3

Pros

Fedora includes excellent security, the latest GNOME Shell, and up-to-date packages in a slick distribution

Cons


Some areas can still be a little technical for Linux newcomers, and fans of GNOME 2 are out of luck

Fedora (once known as Fedora Core) was born out of the ashes of the open source Red Hat project, and while it retains its focus on more technical users its latest releases are suitable for all.

Fedora 17 - codenamed Beefy Miracle - includes cutting-edge software packages such as the latest build of GNOME 3 as its standard desktop environment, along with the newest Linux 3.x kernel branch. Despite this, it's a very traditional operating system - Evolution remains the default email client, for example, despite others moving to Thunderbird.

The newest release also converts more services to systemd - first introduced back in Fedora 15 - natively, and improves the distribution's support for virtualisation systems, including new support for redirection of USB 2.0-connected devices.

Perhaps most impressive of all is Fedora 17's inclusion of OpenStack as standard, making it a



great distribution for those looking to play with cloud-connected software stacks.

While previous releases of Fedora have been notoriously finicky - with even hardcore fans admitting that it can take a little finessing to get the system working to its full potential - newer versions have gone some way to addressing that concern, and Fedora 17 leaves us hopeful for the future.

The facts

First released in 2003, Fedora Core was to take over from the retired Red Hat Linux.

The distro grew out of a volunteer project providing additional software for Red Hat Linux.


Fedora is so named after the type of hat worn by the 'Shadowman' in Red Hat's logo.

Key features

- Fedora's installer includes the option to create a totally encrypted installation.
- SELinux, the NSA-approved access control system, is included with Fedora and enabled by default.
- The distro is slowly but surely making the move to a systemd init

process and native Btrfs support.

- While the 'official' Fedora spin features GNOME 3 and the GNOME Shell, there's also a KDE version.
- Fedora has a reputation for being one of the first 'mainstream' distributions to implement new technologies.



In review

All of the new features and updates do add up, and even with the sometimes limiting selection of packages, it's a very usable and stable distribution that is great for working on current and future open-source projects.

Linux User & Developer 115
4/5



Ubuntu Linux

For years the king of the hill, Canonical's Ubuntu has slipped down the ranks

2

Pros

Ubuntu's historical popularity means it's often supported by commercial software and hardware firms with drivers etc

Cons

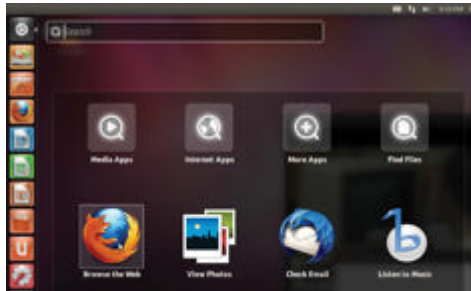
Canonical's use of the Unity desktop environment hasn't proven popular, and has cost it the top spot here

You might be surprised to see Ubuntu in second place: Canonical's easy-to-use distribution is typically at the top of any list of Linux distros that concentrates on general-purpose suitability.

We haven't made a mistake, however: Ubuntu's popularity is now waning, thanks largely to what the community sees as a heavy-handed and dictatorial approach from owner Canonical.

The main reason for this is Unity, Canonical's dedicated desktop environment. Built on GNOME 3 and used in place of GNOME Shell, Unity was first introduced in Ubuntu 11.04 to help push the distro towards the growing number of touch-enabled devices.

It's a move that has cost the company a chunk of the community's support – a trend which only got worse when Ubuntu 11.10 dropped the 'Ubuntu Classic' mode in favour of a 2D version of Unity that doesn't require hardware acceleration support.



Support might for Ubuntu be waning, but Canonical isn't wavering. "We are committed to Unity, we see Unity as – for a variety of reasons, and across a variety of form-factors – the principal driving force for where Ubuntu is going," Canonical's Gerry Carr told us at the launch of Ubuntu 11.10. "Are we concerned? Yes. Do we hope that people will come with us? Yes. Are we going to change our minds on that? No. We know it's an argument that we have to win, we know it's an argument we have to win over time, we know it's an argument we're not going to win with absolutely everybody, but we think it's the right thing to do."

Key features

- Since Ubuntu 11.04, the distribution has used a custom desktop environment dubbed Unity.
- A 'Software Center' makes finding new packages – and even buying commercial packages – easy.
- Ubuntu is designed to offer a single distribution for tablets, desktop PCs

and laptops, and netbooks.


- Backed commercially by Canonical, which offers for-cost enterprise support and add-on packages.
- Ubuntu's popularity in the server world is growing, with Amazon offering Ubuntu cloud implementations on its EC2 service.

The facts

Ubuntu is named after the South African philosophy of 'ubuntu,' which means 'humanity towards others.'

Although still very popular, Ubuntu's reputation has been damaged by a move to the Unity desktop.

Ubuntu was born from the Debian distribution after concerns were raised about the latter's long release cycles.



In review

Where the open source ecosystem is concerned, it's the most recognised brand alongside Tux himself. Does Canonical play entirely by the rules? No, and this is a very important problem...

Linux User & Developer 95
4/5



Linux Mint

Always present in our top-ten lists, Mint has finally overtaken Ubuntu

1

Pros

If you're a fan of GNOME 2, Linux Mint could rapidly become your only option if you want to stay up to date

Cons

The bundled software is a bit conservative; it has few extra features over Ubuntu

A spin-off of Ubuntu founded in 2006 with a more loose approach to the ideals of free software than its Debian grandparent, Linux Mint has been by far the biggest winner of Ubuntu's loss of community support.


The Linux Mint 13 release offers an easy explanation why: despite being built from

the latest GNOME 3 packages, Mint tries to keep things as easy as possible by packing a new feature dubbed MATE which forks the GNOME 2.32 code and combines it with GNOME 3. This in conjunction with their in-house built Cinnamon, designed to emulate the GNOME 2 appearance while providing some of the better features from GNOME 3 and easy customization, is turning heads in the Linux community.

The result of his team's effort, project lead Clement Lefebvre hopes, will be an up-to-date, fully

supported desktop environment that combines the best of both worlds without forcing its users into a touch-centric environment which it has been shown many people simply don't want: since distributions like Fedora and Ubuntu moved to GNOME 3, Linux Mint has seen its popularity rise by around 40 per cent.

It's a bold move, and one which is largely unproven: the Linux Mint team has plenty of hard work ahead of it, but having achieved the seemingly impossible – overtaking Ubuntu at the top of popularity metrics like DistroWatch – already, the signs are good that Lefebvre and team are on the right track.



Key features

- MATE and Cinnamon, introduced in 13, are designed to offer different takes of the GNOME 2 experience.
- The MintMenu launcher highlights new packages, and includes integrated search functionality.
- A 'sensible defaults' approach makes it easy

to include proprietary packages like Flash and MP3 codecs.

• Always known for its visual flair, the newest Mint releases continue to offer an attractive and well-thought-out interface.

• Linux Mint is kept up to date with six-monthly releases.



In review

Mint has continued the tradition of taking all the good parts of Ubuntu, and replacing all the bad bits with fantastic alternatives. Cinnamon and MATE are both great desktops.

Linux User & Developer 115
5/5

The facts

Linux Mint's explosion in popularity is largely thanks to its continued support for GNOME 2.

Project founder Clement Lefebvre aims to create a distro with 'sensible defaults, like Mac OS'.

Two new tools – MATE and Cinnamon – may offer a solution to any who dislike GNOME 3 and Unity.

Torrent clients

We investigate the best way to get the most out of your file sharing by putting four of the best BitTorrent clients to the test

Transmission

This classic torrent client is a standard for a lot of distros, but is it the best?



Let's be honest, no one would blame you for just using Transmission. It's the default BitTorrent application for a lot of Linux distributions, and is great at giving you the very basic information you need to download and maintain your torrents. It requires minimal setup outside of letting it run in the background, and has a good set of default settings for peer connections and queuing.

While it does have a simple and clean interface, it wouldn't be fair to call it a simple client. The minimal UI belies a fairly comprehensive set of options available in the preferences menu. It's very customisable, starting from the basics of having multiple bandwidth caps. You can have multiple cap settings, normal and alternate speeds. The normal setting is mainly designed for standard use, which in most cases would be the maximum bandwidth available; however, you can of course enter any

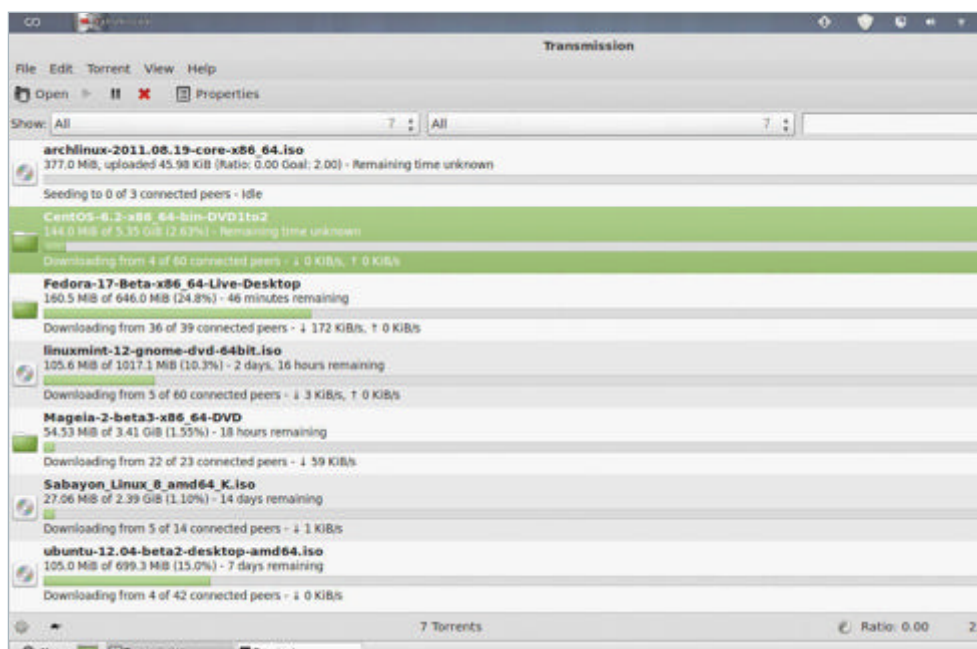
values you wish. While the alternate bandwidth caps would be used mostly for reducing speed, you are able to make it unlimited in case you want the normal speeds to be lower.

There's also a web client for remote access. This can be done locally if you prefer to do as much as possible in the browser, or perhaps find switching windows a bit cumbersome in GNOME 3 or Unity, and the interface gives you just about all the control you would need to maintain and add torrents. You can assign it to different ports, and give it a username and password so that people can't start randomly adding malicious torrents from the outside.

Speed-wise, there's no problem really. And obviously there shouldn't be, as the download and upload rates can only be as good as your connection and the torrent's seeds. However, we have in the past come across clients that are slower than others for whatever reasons, but Transmission doesn't fall foul

of this. This may have something to do with it being a very low-resource program, one of the lowest in our tests.

If you've been using Transmission for a while, none of this should be so surprising – it's a great torrent client initially presented as simple and easy to use, but with plenty of depth for the more advanced user.



SCORES

Setup

Does not run in the background as default, but that's easily fixed along with other basic settings

8

Ease of use

Due to its simple design, it's very easy to use, with a well-laid-out interface

9

Customisation

Nicely customisable, with separate bandwidth caps, UI behaviour modifications etc

8

Features

Not the largest feature set, but it does include a web interface for remote access

7

Overall

Transmission does more than it needs to, and without sacrificing any of its more basic functions to do so. A rock-solid and reliable option

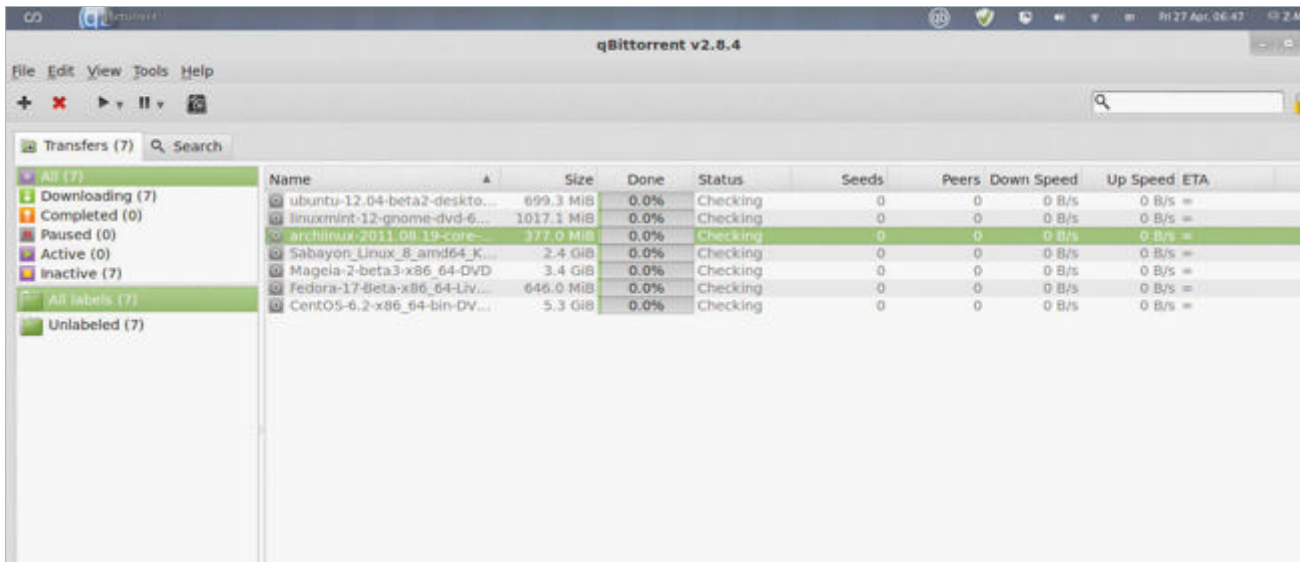
8

More information
www.transmissionbt.com



qBittorrent

A Qt-based client with a UI apparently inspired by the popular μ Torrent



■ Similar but not the same as μ Torrent, the customisable interface can show you all the information you'd want



Named after the great cross-platform μ Torrent, the mainly Linux-based qBittorrent aims to be its open source equivalent.

The interface while supposedly taken from μ Torrent, looks somewhat different to the original's, although it still maintains a well-set-out and logical interface.

By default it's a very simplistic setup, offering no information window when adding torrents, and does not run in the background by default. It still, of course, gives you enough information to properly keep an eye on what you're downloading, but Transmission is roughly the same, so you might wonder why you would continue with qBittorrent.

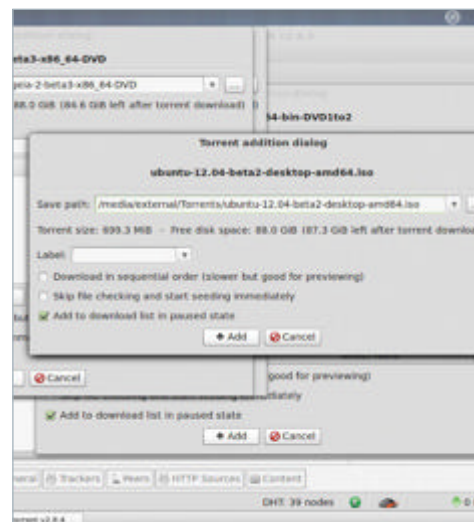
The standard, post-install qBittorrent is only the tip of the iceberg. The torrent client boasts a selection of additional features and options that cover almost any desire while downloading and uploading with torrents. As well as the ability to run programs after torrents are finished, or have a web interface for remote control, you can even have it send you an email when torrents are done.

There's built-in, advanced scheduling to allow you to use alternate bandwidths; for example, setting it to throttle your speed at specific times on weekdays. As well as setting overall active torrents, you can set individual limits for downloads and uploads. Unlike other clients, when you meet your target seeding ratio you can have them completely halt, freeing up bandwidth to seed other torrents. It can even guide you through creating and then seeding a torrent to share with others.

As well as all the other features, and vast ways to customise them, there's a great search feature.

Unlike P2P software of old, it doesn't connect to a network of other qBittorrent users, but instead uses established search engines. You can add and remove whatever search locations you want as well, with further customisation options.

There's so much to like about qBittorrent, and on top of all the ways to use and automate it, it's a well-put-together piece of software. While it can use more resources than something like Transmission, the trade-off is more than worth it for the extra functionality that it adds to the experience.



■ There are plenty of unique features in even the most basic functions of the client

SCORES

Setup

The basic setup includes very little of what you'd expect from other torrent clients

7

Ease of use

The expansive menu is very straightforward and plainly labelled, albeit a touch cluttered

9

Customisation

Pretty much every feature and function can be customised in some way

10

Features

A comprehensive feature set has anything we'd need, and more we didn't know we needed

10

Overall

An amazing piece of software – easy for light users, and fully featured for serious seeders

9

More information
www.qbittorrent.org



Deluge

A Python-based client that leans on its plug-ins

■ It does actually have a nice way to add multiple torrents, despite other flaws

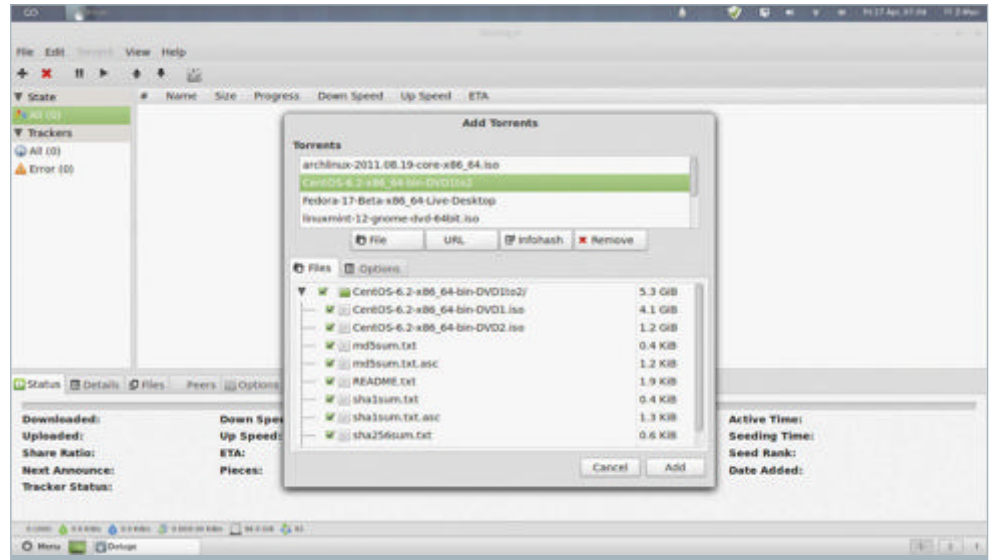


Deluge is a very basic torrent client that mimics the interface layout of µTorrent, more so even than qBittorrent. Deluge's simplistic presentation is meant to be augmented by its plug-in system, with first- and third-party plug-ins allowing you to add in the 'full functionality' of other torrent clients.

At first glance it does have the most features activated when initially installed – with the three panes of category, torrents and information along with a pop-up for adding torrents, and the ability to run in the background. It's a good start, but sadly it goes a bit downhill from here.

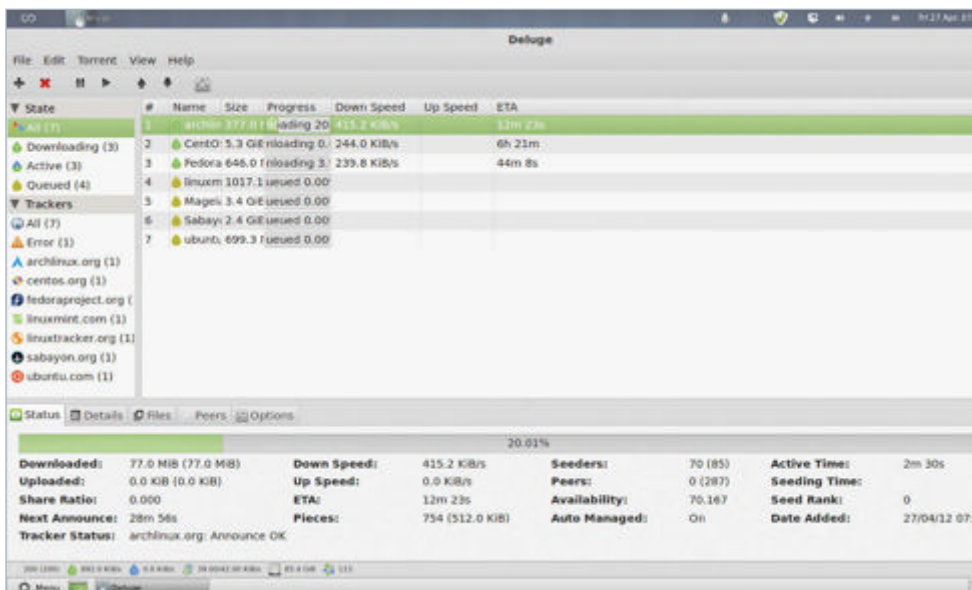
When the Deluge website mentioned that it needs plug-ins to make it a more fully featured client, it wasn't lying. The same kind of bare-bones options are present as the command-line-only rTorrent, with possibly even less control over the customisation of these features. The options and preferences have no kind of plug-in search function, or a link to the website with the plug-ins available. While the website does inform users of this, the Deluge software in the package manager has nothing to indicate this. So to the informed user, that's not a problem: you know to go to the website to get the plug-ins. Unfortunately, the selection of first- and third-party additions are not hugely varied, and a couple of these plug-ins even crashed Deluge after we took the slightly roundabout way to add them. It's not a very pleasant experience.

It's an unfortunate side effect of its lack of popularity compared to the other clients – technically due to the plug-in system it could be



overly customised to your specific preferences, only with the features you wish to have, leading to a slightly less resource-intensive program overall. It falls fairly short of that promise, though, with the lacklustre offering.

Out of all the clients tested, Deluge is the one we'd have the most trouble recommending then, at least out of the box. Maybe down the line the community behind it will improve, making better plug-ins. However, the core dev team need to make finding and adding them a lot easier before our concerns can begin to be addressed.



SCORES

Setup

Very little setup needed for the client, although a lack of info for further customisation is frustrating

8

Ease of use

Finding and adding plug-ins is not immediately apparent, and can crash the software

6

Customisation

The plug-in library should be used for customisation. However, it currently has a poor selection

7

Features

Few features are included by standard, since it relies on the underdeveloped plug-ins

6

Overall

Deluge could be a lot better, if only its plug-in library was richer. As it stands, though, it pales in comparison to the other graphical clients

7

More information
<http://deluge-torrent.org>

Media centres

In a world where the Boxee Box and Roku exist, Linux HTPC software can sometimes be overlooked, but not today...

Boxee 1.5 Out of its box, can this interconnected software reign supreme?



Boxee is quite famous these days.

As a major competitor in the world of set-top boxes, going up against the Apple juggernaut and cult favourite Roku, the Boxee Box has managed to so far outlive Google TV and is far more preferable to a lot of smart TV software. The heart of it all has been the open source, XBMC-derived Boxee software: a powerful, socially connected media centre with hundreds of apps and an easy-to-use interface.

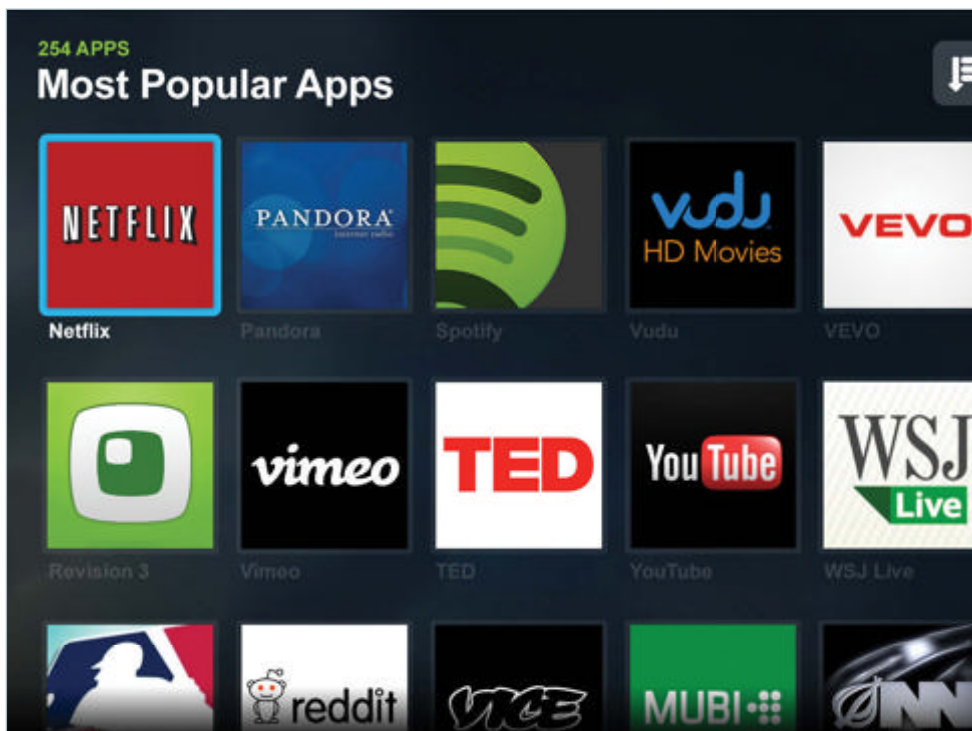
Unfortunately, Boxee's days on the desktop support have come to an end, with the recently release Boxee 1.5 as the last version of the software to be made available for Linux, Windows and Mac OS X. Oh well, it's still usable, right? Well, sort of.

The 1.5 packages were available from the Boxee website until the end of January, so by now they've disappeared. In Ubuntu it should be available in the Software Centre, but it doesn't seem to be there any more. No matter, the source code is freely available on the website... after you register for an account on the site, though. Already that's a hurdle to a lot of users, even if there are handy Readme files to guide you through compiling.

Except a makefile is in the wrong place, and you need to create a specific folder for it to go in, or change a line to make sure it looks in the right place. And finally, after the source is all ready for installing, it turns out you can't install on a 64-bit system. That's why the package was unavailable: it's only in the

32-bit version. Oh, and by the way, the 32-bit package needs to be repackaged as it doesn't install correctly.

It's a horrible mess that shouldn't exist, and a bad way for Boxee to leave the desktop space. Of course, the software – once you finally get it installed, if at all – is amazing. It's intuitive, well labelled and the feeds for your friends and queue is great. Scraping the metadata is nice and quick, and it's displayed nicely in the interface. It's just a shame that there are so many hoops to jump through.



■ Boxee has so many apps to augment your TV experience

SCORES

Setup

A huge hassle to install, with no support for 64-bit, buggy source code and broken packages

3

Ease of use

Fantastic, incredibly intuitive and very responsive

9

Customisation

You can easily add shares manually, or use them with UPnP

8

Features

Hundreds of apps for almost every online media service you can think of

10

Overall

Boxee has evolved into such an incredible piece of software, it's a shame installing it is a horrible mess

8


More information
www.boxee.tv

XBMC Media Center 11.0 Eden

The old dog has never looked so good, or worked so well



■ XBMC gives you control over where the scraped data comes from

 **Don't let XBMC Media Center's humble beginnings as an original Xbox hack dissuade you; now in its eleventh iteration, XBMC has earned its status as one of the most renowned media centre platforms.**

Installing and setting up XBMC is a very easy process, with packages available for all major distributions. If it's not included in the standard repositories, the XBMC wiki has a simple guide to point you in the right direction, and the source code is easily obtainable from the website. Once it's installed, there is an app available in the program menus; however, all this does is kick you back to the log screen. Whereas once this

used to be an application you launched within your desktop environment, XBMC is now its own desktop environment. There are plenty of benefits to this; if you have a dedicated HTPC, it can boot straight into XBMC without having to load a whole desktop environment you'll never see.

Setting up XBMC is nice and easy, with plenty of advanced customisation options for people with very specific setups. Options for scan, refresh rate and surround sound join the customary resolution and base volume, as well as plenty of skins to play around with. Adding media sources is very simple, with a wide array of options for

local sources, connected devices and every type of network connection you can think of. Adding a UPnP share from both Windows and Linux systems over the network was met with little to no hassle, and indexing was pretty speedy as well. You can even choose which scrapers to use to obtain metadata, although this of course works better with local content. When it does work, though, the backgrounds, icons and programming information are very well presented.

The playback experience is of course smooth, and when tracking through even high-definition files we didn't experience any lag or buffering. There are a lot of apps available for XBMC, including add-ons for Netflix, Tested and social media; however, the range isn't as good as Boxee, or as accessible.

XBMC is still a fantastic, functional media centre, and it's used as the basis for a lot of other media centres such as Boxee itself.

SCORES

Setup

Nice and simple, although it might confuse the more novice user

9

Ease of use

Well laid out, optimised for big screens, scrapes smartly

10

Customisation

A little complicated if you don't know exactly how you're sharing

8

Features

While it has the basic features you would need, there's not much app variety

8

Overall

Not quite perfect, but it's a fine choice for any HTPC. Some user-friendly tweaks would go a long way

9

More information
www.xbmc.org

Moovida 1.09

Fluendo's HTPC offering may be not up to date on Linux, but does it need to be?



Moovida is media centre software that comes from the guys over at Fluendo, and while they're a big multimedia company, a lot of their software and codecs are proprietary. Moovida on Linux, however, is open source, although it is also a complete version behind the supported release, which is only available on Windows for free.

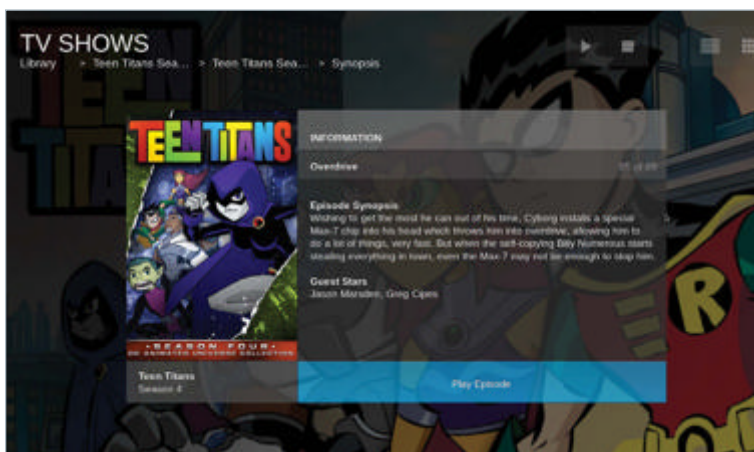
Despite these warning signs, this three-year-old software is certainly not bad. Once installed, you can run it on top of your desktop environment and it automatically scales to your monitor's – or TV's – resolution. It looks

fine in 1080p resolutions, and is well laid out for easy navigation.

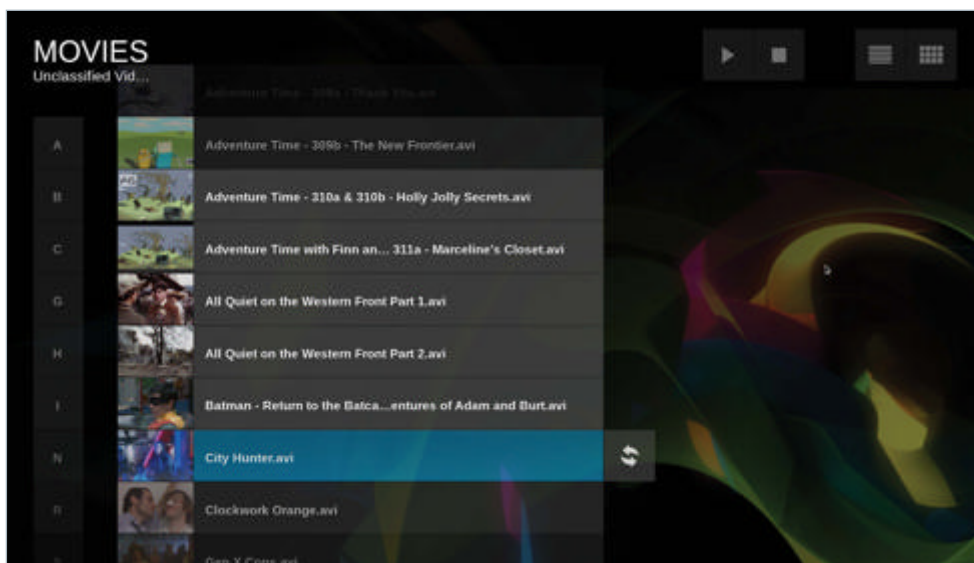
Sadly, the package is not readily available on all major distributions' package installers, so compiling from source may be the only way to get it. There's also very little support from Fluendo, so problems with the software will not really be fixed, and there are a few of them.

Moovida is not easily configurable. Actually finding and watching media on the software can be slightly unintuitive, not looking in the basic places you might have videos or music by default. Folder locations are browsable in a separate part of the interface, and these can

■ Moovida scrapes about as well as XBMC, and displays the results well



■ Moovida's interface is nice and functional



SCORES

Setup

Some distros may need to compile from source, and adding media can be a pain

7

Ease of use

A somewhat inconsistent user experience that can be randomly great and bad

8

Customisation

While there is a menu option for networked computers, it doesn't work

0

Features

Asides from lack of networking, there are a few very basic apps and add-ons

6

Overall

After you've overcome the difficult setup, and made sure all your media is locally stored, it can be a great viewing experience

7

More information
www.moovida.com

then be added and given a category. It's not the most straightforward method of adding media, unfortunately.

Networked shares didn't work at all either. There is an option to look for networked computers, but it just doesn't locate any other systems over a LAN. This seems like a massive oversight, and while it can be easily overcome by adding media to the HTPC itself, that may not be an option for everyone. Having large storage capacity locally is going to be either expensive, increase hardware size, or both.

That might seem a lot of negative comments. However, once you get past the annoying setup, maybe delve into the well-hidden conf file, the rest of the viewing experience is actually really good. Moovida does the usual metadata scraping for TV shows and films, with backgrounds, information and cover art popping up after a quick scan of the available content. There are also add-ons for Moovida, although the selection is fairly basic and limited to the likes of online clip sites like YouTube and not much else. Overall it just shoots short of its XBMC-based contemporaries, or at least this outdated version available does.

Freevo 1.9 An oldie, but a goodie, this HTPC software packs a few extra features



Freevo is not standard HTPC software, although it can be used as such. As well as including PVR functionality and support for multiple back-ends, it can also be used for home automation. Unlike the more popular MythTV, it has support to play local and networked files out of the box, so to speak. The current build of Freevo is about three years old, much like Moovida. The team over at Freevo have been developing Freevo 2 for about three years, with a beta still a way off.

Freevo is readily available from a lot of repositories and the Freevo website, with packages for all major distributions and, of course, source code. The installer will guide you through some setup stuff for your TV, such as the resolution, NTSC or PAL, and what kind of TV aerial it might expect for its PVR functionality. It's nice to have all this stuff sorted so in theory you can get started straight away.

Unfortunately this is not the case. While it'll optimise itself for use on your TV, that's about it. To add any media from local storage or shared network folders, you need to manually edit a conf file in Python. The file itself has plenty of documentation

into how you need to format the additions; however, the locations to add these short lines of code are a long way down the file. Even if they weren't, this is not a user-friendly way to get this media added.

All the major settings are kept in this and another file, and while this frees up the viewable interface, it really limits the user base. Sure, if you have the technical know-how, patience and a couple of days free, you can create an incredibly powerful and automated system spread across several back- and front-ends in every TV in the house.

On top of that, the basic themes are ugly. While the software knows the resolution of the displays it's running on, fonts look horribly low-res. Aesthetics aren't everything, of course, and navigation is simple and well labelled. Likewise, playback is great, with support for every file type we could throw at it. There are also a few plug-ins, but mainly to aid the PVR functionality. Generally though, it looks and feels outdated to both use and set up.



Freevo will scrape a little data for its functional menus

SCORES

Setup

Easy to install, but adding sources manually to a Python file is not ideal in any circumstance

5

Ease of use

A very basic and uninspired interface will get you where you need to go

7

Customisation

Solid streaming, but there's no support to just browse UPnP shares

7

Features

There's a wide variety of random features, but a lack of online content viewers

7

Overall

Freevo has its place, or several different ones around a house if you wish. It's just too much in an era of plug-and-play devices

6

More information
<http://freevo.sourceforge.net>

Mind-mapping software

Knowing your mind is one thing, but mapping it is something else entirely. We pit the four best mind-mapping tools for Linux against each other to see how they fare

FreeMind

FreeMind is a powerful tool for mind mapping, that comes packed with features but can frustrate



Let's get this out of the way first: FreeMind is a Java application, designed for cross-platform use on any system that has a Java runtime environment. If

you dislike Java apps in general, you're probably not going to like this one.

With that out of the way, FreeMind is a great mind-mapping package. Its Java base means that load times can be long on older hardware, and memory usage somewhat inefficient, but it's simple enough that it runs well on anything with more than a handful of RAM.

The user interface in FreeMind is quick and easy to pick up. Icons have text labels, the documentation is impressive and there's a great community behind the software. One thing that can catch you out, however, is the lack of drag-and-drop: in the default edit mode, it's impossible to simply pick a node up from the map and slide it elsewhere for aesthetic purposes.

This issue aside, using FreeMind is a pleasure. It's possible to do most things with the mouse, although learning the keyboard shortcuts makes things a

lot faster. Inserting a new child node is as simple as pressing 'Insert,' for example, while entering a new sibling node – one next to the node which is currently active – is as fast as pressing 'Enter'.


FreeMind's clean interface and relative simplicity hides some powerful functionality: it's possible to import data from a variety of formats quickly and easily, as but one example. Copy data from a webpage and past it into FreeMind and it'll automatically work out the best way to represent what you copied as a mind map – although it'll drop any special formatting in the process.

Nodes can contain icons, hyperlinks and can even be encrypted – either as individual nodes or as an entire encrypted FreeMind file. FreeMind is the only tool on test to offer this functionality, and while we'd raise questions about the robustness of the encryption on offer it's certainly a way to keep nosy coworkers or family members from snooping.

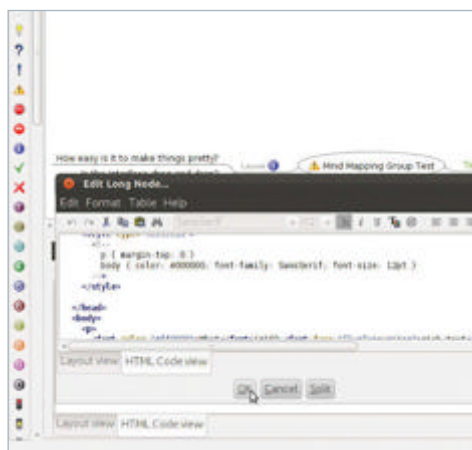
FreeMind's export functions are also excellent, with a vast variety of formats supported – including

interactive websites through a Java- or Flash-based client package. More prosaic formats include Open Document Format, PDF, XML, XHTML and SVG images for easy poster-sized printing.

A fork of FreeMind, dubbed Freeplane, adds new features but isn't yet available in most software repositories.



■ Formatting in FreeMind can be simple, or you can tweak individual node's HTML source



■ For the advanced user, almost everything in a FreeMind map can be edited

SCORES

Ease of use

Despite the irritation that comes as you try to drag and drop a node, FreeMind is pretty easy to pick up

8

Performance

FreeMind's power comes at an a cost for those on older hardware, with long load times

7

Features

FreeMind has plenty of features, with pretty much every requirement easily met

9

Compatibility

If FreeMind can't import your data, somebody's probably written a script that can

10

Overall

It falls short of perfect, but for the majority of use cases FreeMind is the obvious choice for a mind-mapping tool


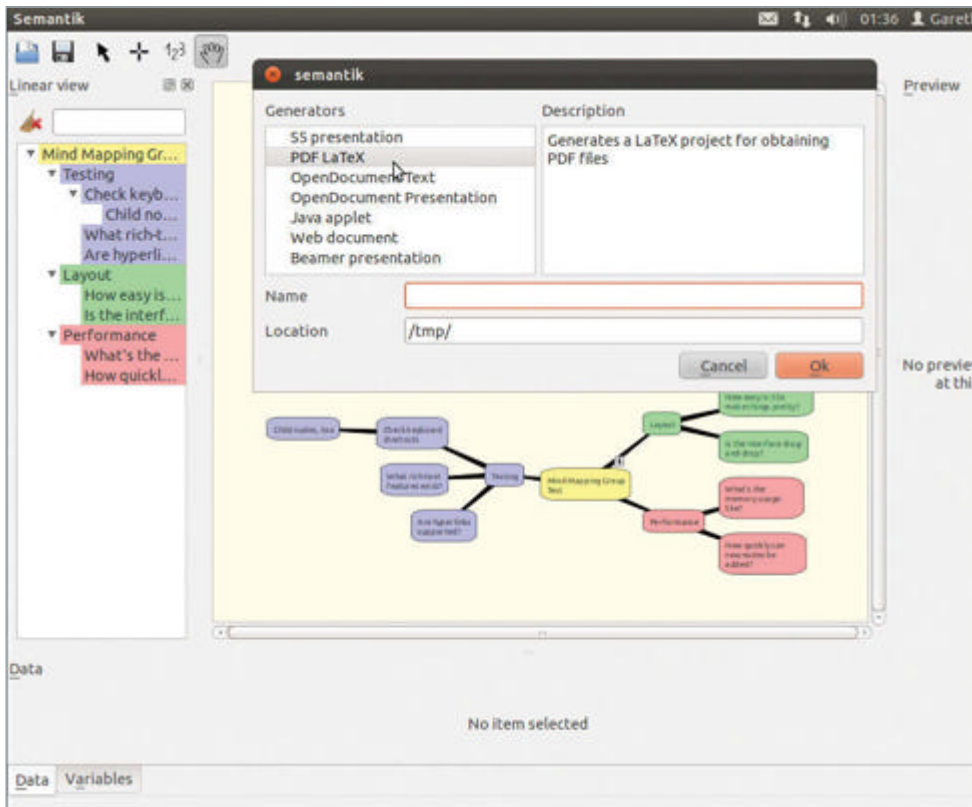
8

More information

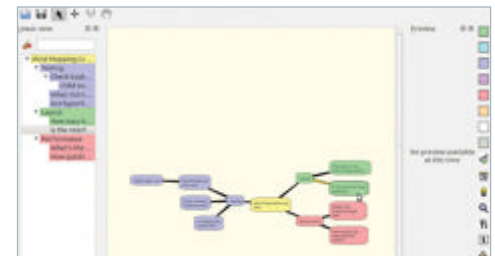
<http://freemind.sourceforge.net/>

Semantik

Designed for KDE, Semantik is attractive but comes with a few gremlins under the bonnet



■ Sadly, the results aren't always great. Semantik's HTML export, in particular, is basic



■ Semantik's export options are plentiful – even including an interactive Java applet



Semantik is a bit of an oddball in this list: while it certainly conforms to the general functionality of a mind-mapping application, its original purpose was for use as a document layout tool. As a result, there are certain things it doesn't do so well, and others that it does absolutely brilliantly.

The feature most likely to grab a user's attention is the automatic zoom. As nodes are added to the map, it starts to shrink so the entire map is kept on screen. It's a neat feature, but one that comes with a slight issue: as the map zooms, and as you override the zoom with the keyboard or mouse wheel, scrolling can become broken, leaving you with your map sat at the bottom corner of the screen. It's by no means a game-breaker, but it can be irritating. Where Semantik differs from other tools on test is the 'Data' view. This allows additional info to be added for the document generation process – perhaps the most obvious sign that Semantik isn't a mind-mapping tool in the traditional sense.

Another sign is in poor data import functionality. Where applications like FreeMind and VYM assume that you've got existing data you'll want to bring into

your mind map – in a similar way that cut-outs from magazines are used for design boards – Semantik assumes you'll be creating all your data from scratch.

As befits a document generation utility, however, exporting the data is quick and easy. There are plenty of formats supported, including OpenDocument, LaTeX, and even a customised Java applet for embedding into webpages. If you're creating mind maps to use in publications, Semantik is certainly worth a look.

Sadly, it's let down by its lack of certain basic functions. Rich-text handling is hit and miss, while the editing can be awkward compared to other tools on test. Semantik is certainly usable, but we'd suggest that there are better options for those not tied to the software's admittedly excellent export functionality.

Perhaps the biggest issue with Semantik's feature set is the inability to simply add hyperlinks to a node. While mind-mapping, being able to link nodes to local or remote resources as a memory-jogger is invaluable, and this simple omission makes Semantik hard to recommend compared with other tools on test.

SCORES

Ease of use

The automatic zoom is handy, but sometimes breaks scrolling in a confusing way

7

Performance

For KDE users, performance will be fine – everyone else, expect delays as libraries are loaded

7

Features

Most features are present and correct, although better handling of rich text would be nice

7

Compatibility

There are plenty of ways to get data out of Semantik, but getting data in can be awkward

6

Overall

Could be good for KDE users, but isn't the best for full-time mind-mapping use

6

More information

<http://code.google.com/p/semantik/>

Labyrinth The fastest tool on test, is Labyrinth a little too simple?

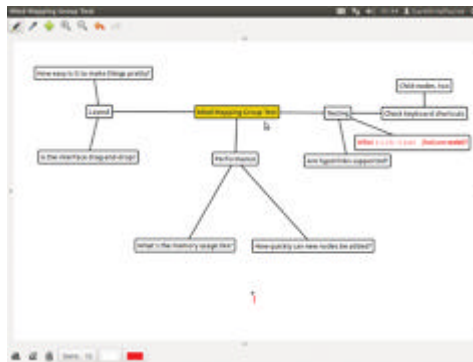


Labyrinth is simple, and that's both a good thing and a bad thing.

The good part is that it loads incredibly quickly: load times for the application itself and files created therein are by far the fastest in Labyrinth compared to any other tool on test. If you're on restricted hardware, that will probably be enough to sway your decision.

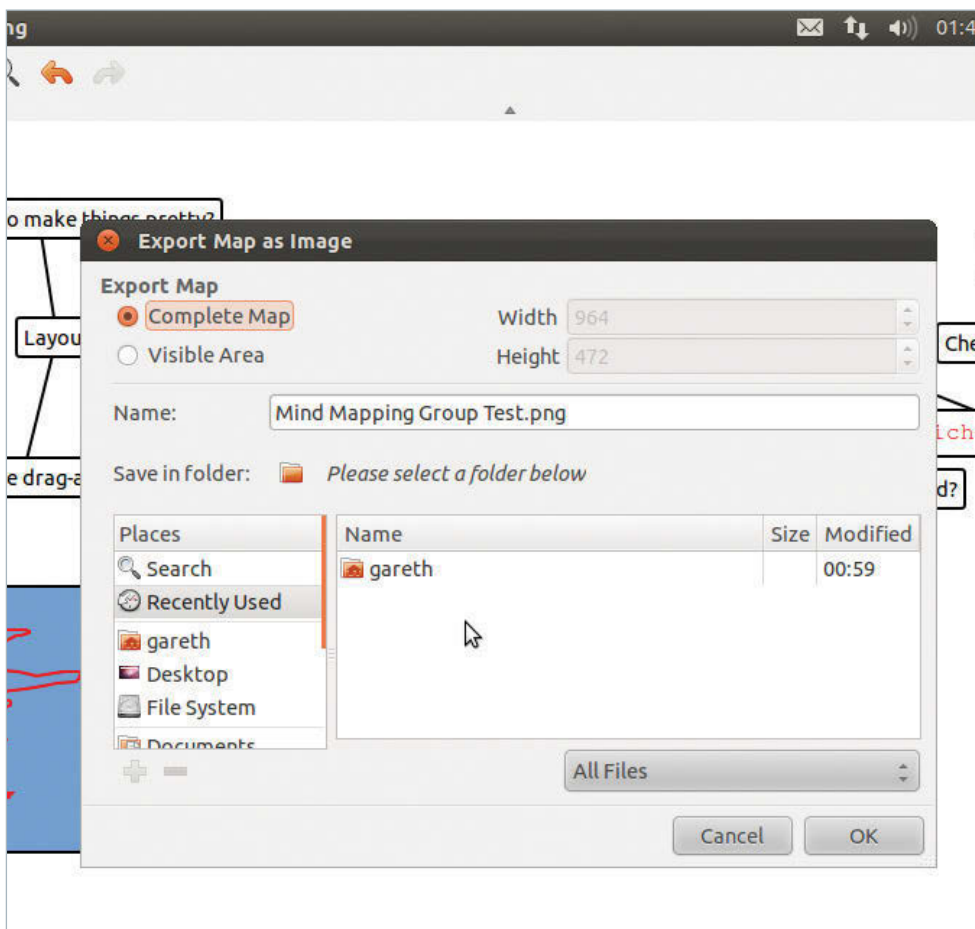
The first screen seen is one of its handiest features: all files created, regardless of where they are saved, appear as if in a virtual filing cabinet. Files are also automatically saved, even when freshly created – something the auto-save functionality of other apps forgets, with the feature only kicking in after the file is saved for the first time.

The simplicity has a dark side too, however. Labyrinth's interface is spartan in the extreme, and for a newcomer it can be difficult to figure out exactly how to get started. Creating a node is as simple as clicking anywhere on the page: a node becomes a child node if you click on an existing node first.



Once you've got the hang of things, Labyrinth starts to show its potential. Its rich-text editing is basic but handy, while it's easy to reposition existing nodes with the drag-and-drop interface.

Labyrinth also packs a 'Drawing' tool, which creates a new child node to hold freehand scribbles created with a mouse, touch-screen, digitiser or



SCORES

Ease of use

The 'filing cabinet' of maps is neat, but the Labyrinth interface is overly simplified

6

Performance

Easily the fastest tool on test, Labyrinth is great for older hardware or those in a hurry

9

Features

Sadly, Labyrinth's speed comes at a cost to its feature set. It works, but it's basic

6

Compatibility

You can export as a map or an image, and that's about it. Importing is a pain, too

4

Overall

Lightweight and simple, Labyrinth is tempting but you'll probably run into its limitations sooner rather than later

5

More information

<http://code.google.com/p/labyrinth/>

graphics tablet. It's a neat idea, but one that grows old fast: the editing functionality is extremely basic, and if you're looking to embed an image you're probably going to be better off creating it in your favourite illustration package and importing it into one of the other software packages on test instead.

Sadly, that's where Labyrinth's good points run out. Importing existing data is either a pain or impossible depending on the format, while the software only supports two export formats by default: you can save a Labyrinth-compatible map file, or a non-interactive image. With rival packages offering multiple formats and even interactive webpage widgets as export options, it's a real shame.

It's this import/export aspect that really lets Labyrinth down. With a little more work on the user interface and the ability to export and import a larger number of formats, it could really shine as a fast and lightweight alternative to the likes of FreeMind. As it is, however, it's something that you'll likely come to regret installing before long.

■ Labyrinth's export options are poor: only map files and PNGs are supported

View Your Mind (VYM) An attractive and powerful utility



If you're a KDE Bookmarks user, install VYM. It's really that simple: out of all the packages on test, VYM is the only one to offer native compatibility to handle KDE3 and KDE4 Bookmarks files. If you have a lot of existing data you'd like to see visualised in this format, it's an easy choice to make.

That's not to say that VYM is a one-trick pony. Its import and export facilities are excellent: files created in Mind Manager and FreeMind can be imported as-is, while export formats include OpenDocument, webpage, LaTeX, XML, CSV and even TaskJuggler files for those dealing with the visualisation of project management tasks.

VYM's icon-based interface is clean, bearing a similar layout to Freemind. Some icons can be a trifle oblique, though, and it takes a little while to become fully acquainted with the software.

Once you're settled, however, there's little VYM can't do – with the notable exception being its poor

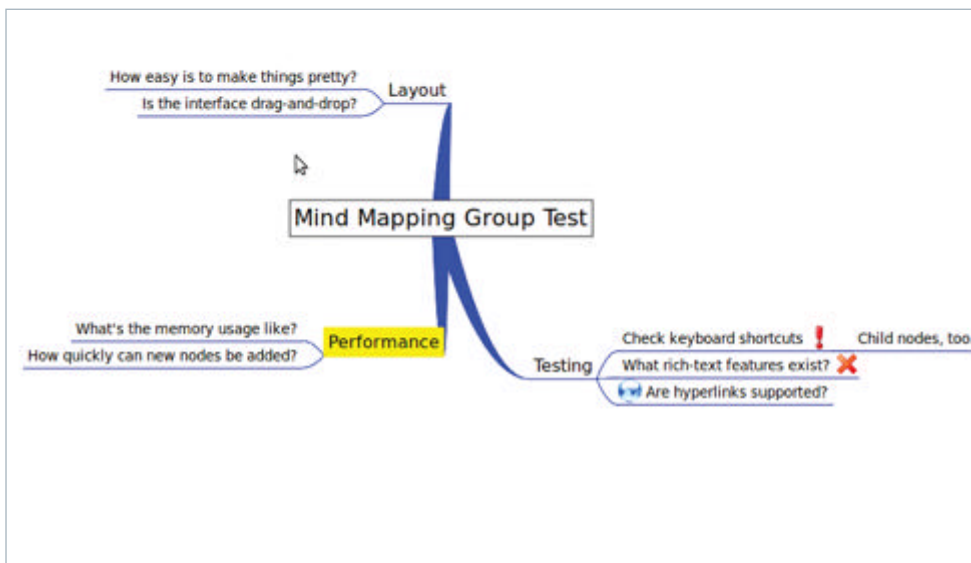
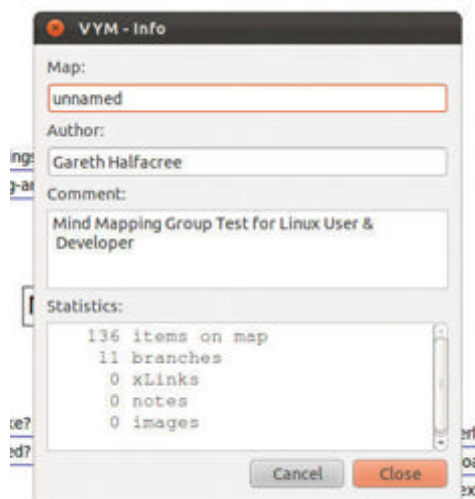
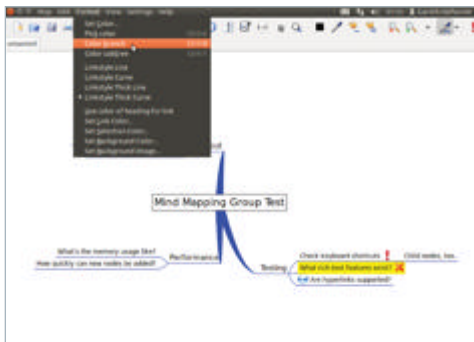
rich-text handling facilities. On a more positive note, features like node sorting can make dealing with complex mind maps a breeze. If you would rather organise your thoughts manually, VYM's drag-and-drop interface is quick and re-routes the links between nodes quickly and efficiently.

VYM also comes with a large number of icons that can be applied to each node. While adding icons isn't for everyone, it can help bring a bit of visual distinction to a map – and give you an at-a-glance index for rapidly finding particular categories of thought.

As with FreeMind, it's also possible to add hyperlinks to individual nodes. When you first get into mind mapping, this can seem like a feature you'd not likely use – but it can prove invaluable for providing context or background to short notations within the mind map itself.

The performance of VYM is roughly on a par with that of FreeMind, with the exception of loading: as a native application, rather than a Java app, VYM loads significantly faster. During use, however, it's hard to draw a distinction between the two – even when loading large mind maps from disk.

Where VYM falls down, however, is in its extended feature set: where FreeMind allows its users to delve as deep as the HTML and XML that make up a particular node for true customisation, VYM relies on predefined node styles. There's plenty to choose from, but if you're after customisation you'd probably be better off with FreeMind.



■ VYM's interface is clean, but some of the icons are a trifle unclear for newcomers

SCORES

Ease of use

VYM's icon-based interface can be oblique, although its drag-and-drop layout system works well

8

Performance

Performance is generally good, although VYM can take a while to load initially

8

Features

VYM is rich in features, although its rich-text capabilities lag behind those of FreeMind

8

Compatibility

VYM's import functionality is great, and it exports to a wide variety of formats

9

Overall

If you're a KDE Bookmarks user, VYM's support for them will come in handy. For others, FreeMind is a better choice

8

More information

<http://sourceforge.net/projects/vym/>

Project managers

If you use your Linux system for business purposes, eventually you're going to need a project management tool. Here we will take a look at four of the best...

Planner Part of the GNOME desktop project, Planner covers the basics well



If you're a GNOME user, the chances are that – even before we delve into its features and restrictions – Planner should be your weapon of choice. While

it's not as powerful as some tools on test, it covers the basics required of a decent project management suite well, and as part of the GNOME desktop project integration is – as you would expect – excellent.

While there are no default templates for various projects, the main interface of Planner is clean and uncluttered. Split into four main sections, Planner gives the user quick access to a Gantt chart view, a table of tasks for the project in question, a table of resources and a schedule of which resources are in use at any given time.

As with other applications on test, it takes a while to get used to Planner's particular vocabulary if you're not familiar with project management terminology. 'Resources', in this instance, refers either to staff working on the project or physical items that are required: for a software development project, for example, you might create entries for

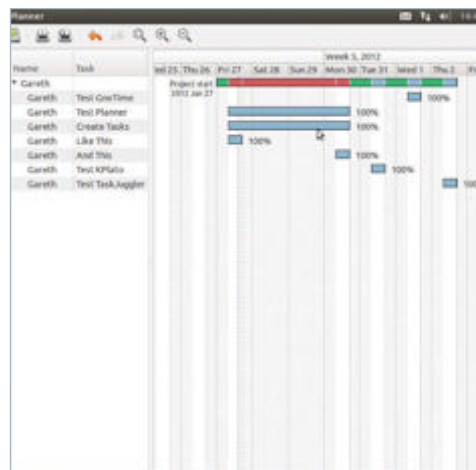
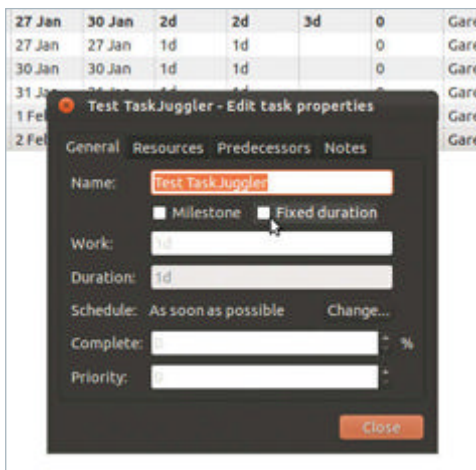
your programmers and testers, while a building project would require entries for labourers and – where supply is limited – tools like drills and routers.

It can be a pain to set up the resource list at first, but it's a vital part of the process: when scheduling tasks, the user needs to take into account the availability of the resources in order to avoid double-booking.

The part that will really trip up newcomers, however, is the 'Predecessors' option. Appearing as a tab in the task edit dialog, Predecessors allows the user to specify which tasks rely on which other tasks being completed: a testing phase, for example, must logically follow a bug-fixing phase, which in turn follows the coding phase.

Sadly, this is one of Planner's weaker areas: if you don't specify predecessors for each task, the software will attempt to schedule them simultaneously; and with the Predecessors tab requiring use of an awkward drop-down box to select the order tasks should be completed in, it can be

painful when organising larger projects. For those coming from a Windows environment, however, there's a last feature that should put Planner head and shoulders above the competition: it's the only tool on test that can natively import existing Microsoft Project files.



■ Editing tasks in Planner is straightforward, but watch for 'Predecessors'

SCORES

Features

It might not have KPlato's power, but Planner covers the project management basics well

6

Ease of use

Once you've got the hang of things like 'Predecessors,' Planner is easy to use

8

Import/Export

If you're looking to use existing Microsoft Project files, Planner could be your only choice

8

Compatibility

GTK+ based, Planner integrates well with GNOME and works fine elsewhere

8

Overall

Lightweight and simple, Labyrinth is tempting but you'll probably run into its limitations sooner rather than later

8

More information
<http://live.gnome.org/Planner>

KPlato

Part of KOffice, KPlato is an immensely powerful tool in the right hands



As with GNOME users and Planner, our advice to KDE users is this: use KPlato.

It's powerful, it's flexible, and as part of the KOffice suite the chances are good you've either already got it installed or it's the work of mere moments to grab a copy from your distribution's repository.

That's the good news. The bad news is that, unlike GNOME users and Planner, you're in for a bit of a steep learning curve. KPlato is an incredibly powerful piece of software, offering by far the most functionality of any tool on test – save perhaps, in the hands of an experienced user, TaskJuggler.


Thankfully, the team behind KPlato has worked hard to ease users in to the package's many features and functions: a splash screen offers a choice of three default templates, from a blank project ripe for customisation through a project with basic settings already defined and up to a regimented project based on a standard eight-hour work day and 40-hour week.

Loading up a template gets you a friendly welcoming screen, offering access to tutorials that can help you get to grips with the software; and you're going to need them. KPlato offers some extremely impressive functionality rarely seen outside expensive commercial packages, including

a handy automated scheduler: simply add your tasks, use the drag-and-drop dependencies editor – the equivalent to Planner's Predecessors tab – to define which tasks rely on which other tasks, assign resources and hit the 'Calculate' button for KPlato to automatically work out the optimal schedule while avoiding resource conflicts or out-of-order execution. Sadly, it can be a difficult road to get from A to B for a newcomer: with so much power on offer, there are times – such as when faced with the 'Optimistic/Pessimistic' choices, or the default estimate measurement of 'Effort' – when you're left scratching your head in confusion.

It's these often oblique terms – to anyone who hasn't had training in project management, anyway – that can make KPlato seem awkward and obtuse. A look through the documentation helps to clear things up, but it can be a slow process for a newcomer to the app to become proficient enough to the point where the software begins to save time, rather than eat it.

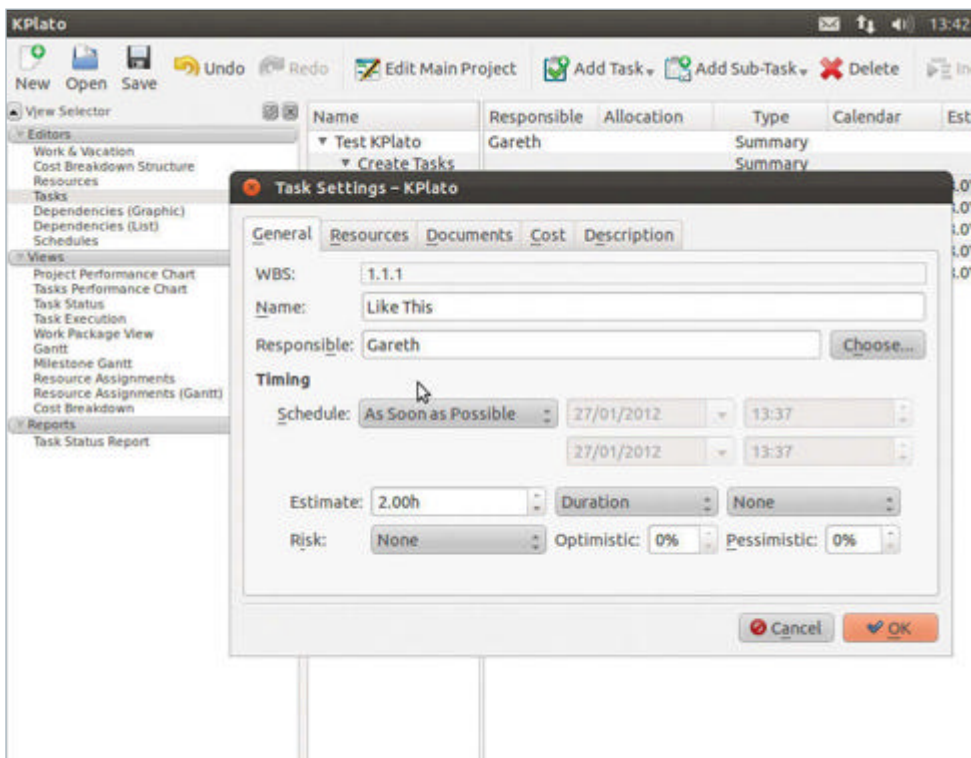
Our advice: stick with it. If you put the effort into learning KPlato, you will be rewarded with an incredibly powerful project management tool. And while it might be overkill for smaller projects, it can really pay dividends on more complicated endeavours.



■ KPlato's Gantt chart view, populated by the automated scheduler, is great



■ Default templates and a good tutorial get you started with KPlato



SCORES

Features

It's hard to think of anything that KPlato doesn't offer the budding project manager

10

Ease of use

Some aspects of KPlato can be confusing, but work through the tutorials and you'll be fine

7

Import/Export

Exporting to PDF and ICS files is handy, but there's no import functionality for other formats

7

Compatibility

For KDE users, KPlato works brilliantly; for other desktops, beware the dependencies

8

Overall

For those willing to learn, it's hard to find a better project management app

8

More information

www.kde.org/applications/office/kplato/

GnoTime More of a time tracker than a project manager, GnoTime has its place – if you can get it working



If talk of predecessors, optimism and Gantt charts all seems too confusing, GnoTime – the Gnome Time Tracker, to give it its full title

– could be an alternative. Designed more for freelancers working on billable projects than managers looking at controlling staff, it offers some neat functionality in a relatively simplistic user interface.

Thanks to its focus on time tracking, GnoTime is by far the simplest tool on test to get started with: create a new project, add in tasks and hit the 'Timer' button to begin tracking the time you've spent on each section.

For more detailed reports, you can add diary entries with lengthy notes – a handy feature when it comes time to invoice and the manager wants to know where all the time was spent.

There's another feature where GnoTime comes in handy, too: an HTML template, complete with replaceable header graphic, allows the user to automatically generate an invoice which can be printed out as a PDF and sent to the client or project manager. Customisable 'value' fields even add the cost up automatically, although it's questionable

whether you'd want to send an invoice which is literally accurate to the second.

GnoTime's reporting functionality is its key feature: based around HTML, each report – whether it's an invoice, a journal, or the diary entries you've been keeping for each task – can be published or saved as a file for external publication. It's a clever way of keeping things transparent, even if the default appearance for each report is a trifle spartan.

Sadly, there are plenty of caveats with using GnoTime, too. The first issue we encountered was with getting the software up and running on our test machine: under Canonical's Ubuntu 11.10, featuring the Unity desktop, GnoTime installed but refused to create new projects or tasks. It wasn't until we switched to a testbed system running the older GNOME 2-based Ubuntu 10.04 that we were able to get GnoTime operational. The software is also somewhat buggy, which couples with its lack of functionality to make it hard to recommend. If all you're concerned about is your individual input into a project, it certainly has its uses; but if you're trying to manage multiple resources or a more complicated and lengthy project, it's a poor choice – even if you can get it to work properly on your system.

SCORES

Features

More of a time tracker than a project management tool, GnoTime lacks a lot of vital features

4

Ease of use

A lack of features does, however, translate directly into pleasant simplicity

8

Import/Export

The ability to export both tasks and projects to tab-delimited files is handy

5

Compatibility

We ran into trouble on Ubuntu 11.10, having to revert to an 11.04 test system to create new tasks

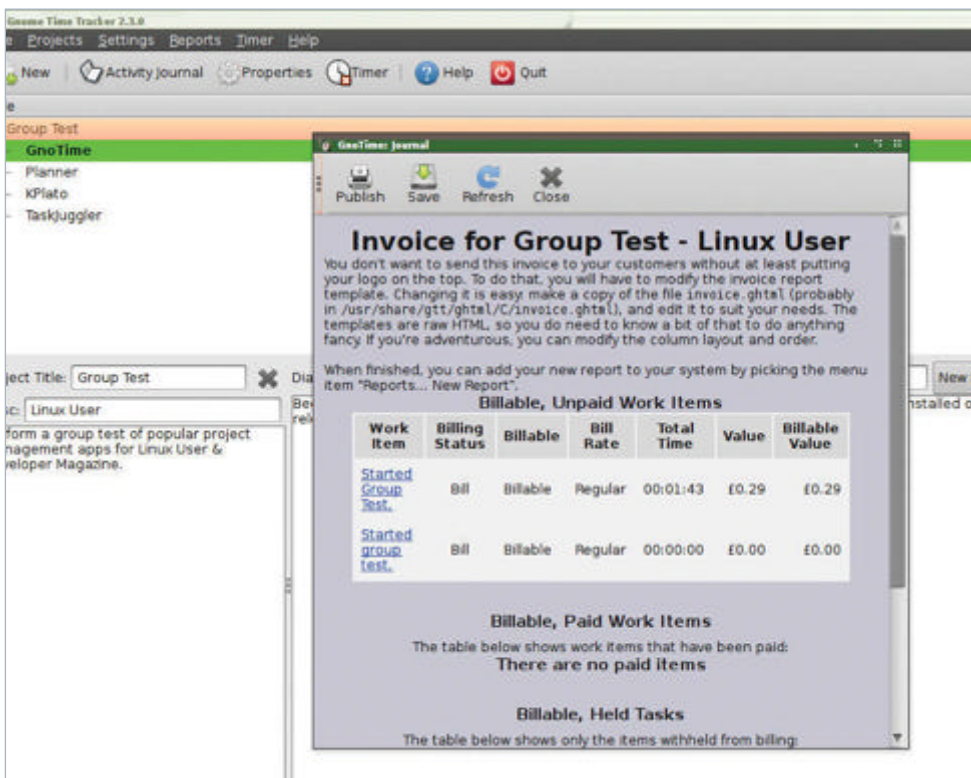
4

Overall


If all you need is a time tracker, GnoTime's worth a look – but there are better tools out there

5

More information
<http://gtr.sourceforge.net>



■ The invoice generator in GnoTime is a handy feature for freelancers



■ GnoTime's task editor is basic, but does the job

TaskJuggler

If you like a challenge, learning TaskJuggler is certainly something to consider



Sometimes it's easy to tell what sort of a person came up with the idea for a particular application, and with TaskJuggler it's a no-brainer: based entirely around extensible markup language (XML), TaskJuggler is the sort of software that won't appeal to anyone except programmers and hackers.

If KPlato's interface is so packed with features so as to be confusing, TaskJuggler's is actively hostile to anyone who hasn't seen an integrated development environment before. Loading a template results in a syntax highlighted view that will be instantly familiar to programmers and alien to everyone else.

That's not to say that TaskJuggler lacks in graphical niceties, however: browsing through the

system will show reports that include HTML and CSV along with the project manager's favourite, the Gantt chart. These all operate in much the same way as in KPlato or Planner, but it's getting the data into the system in the first place that can prove a problem.


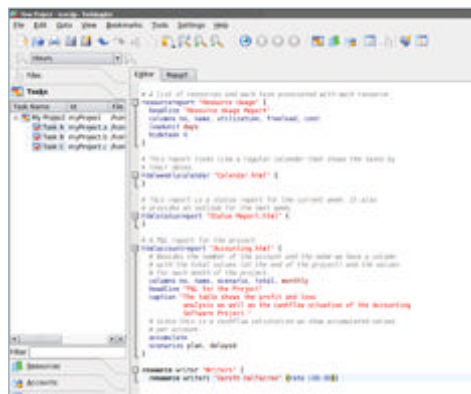
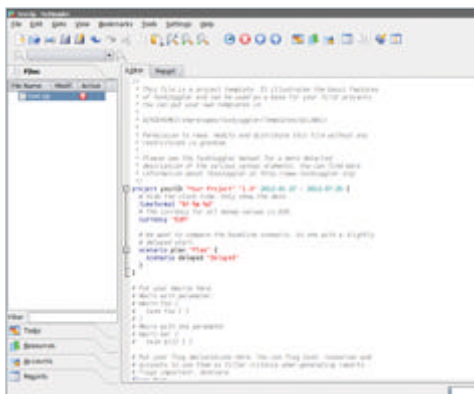
Loading up the default template makes things a bit clearer, thanks to good commenting, but it's still an extremely confusing environment for anyone except a coder: all projects, tasks, resources and schedules need to be defined by hand in the XML before any of the TaskJuggler features can be used. If you're not used to coding, it's a task that will immediately send you running for one of the alternative programs on test here.

If you're a coder willing to put in the time, however, TaskJuggler proves to be an extremely flexible tool. It's even possible to add macros directly into the

project file itself, or include other sub-project files in much the same way as a C programmer might include a header file in his or her project.

Sadly, as with GnoTime, there's a catch: TaskJuggler is built around KDE 3.5, and has been removed from the repositories of most distributions as part of the move to KDE 4. While it's possible to compile the program and have it run under KDE 4, we found the software too buggy to use until we dropped back to a system with the KDE 3.5 libraries installed.

For the hardcore hacker, it might be worth the effort; for most users, however, we'd suggest going with Planner or KPlato instead and spending the time you would have invested trying to get TaskJuggler to play nicely in a modern Linux environment on actually managing the project.



■ While TaskJuggler includes a Gantt chart mode, it's not easy to configure for newcomers

SCORES

Features

If you don't mind a bit of XML hacking, TaskJuggler will likely do everything you want and more

9

Ease of use

Programmers and hackers will feel at home with TaskJuggler, but it's hostile to newcomers

4

Import/Export

The XML files are human readable, and there are HTML, CSV and iCal report options

5

Compatibility

TaskJuggler is built against KDE 3.5, and it's difficult to get running on a KDE 4 system

4

Overall

Undoubtedly powerful, but unless you like a challenge it's worth investing your time in learning KPlato instead


5

More information
www.taskjuggler.org

CD rippers

Getting your CDs onto your PC needn't be a chore: we test four of the best audio CD ripping packages around for your convenience

Asunder This GTK+-based ripper is lightweight yet powerful

 Originally developed in 2005, but then abandoned for two years before being picked up again in 2007, Asunder is a surprisingly powerful ripper with remarkably few dependencies: based on GTK+, it doesn't require any GNOME libraries and is happy working in almost any Linux desktop environment.

Despite its lightweight nature, Asunder offers a powerful list of features. The most recent version can rip audio CD tracks to uncompressed WAV files, lossy MP3, Monkey's Audio, AAC, Musepack or WavPack and Ogg Vorbis files, or lossless FLAC format.

Choosing FLAC – the Free Lossless Audio Codec – set to compression level five as our test format, we set about ripping a raft of audio CDs to put Asunder through its paces.

If we were expecting problems, we were pleasantly surprised: despite throwing some difficult

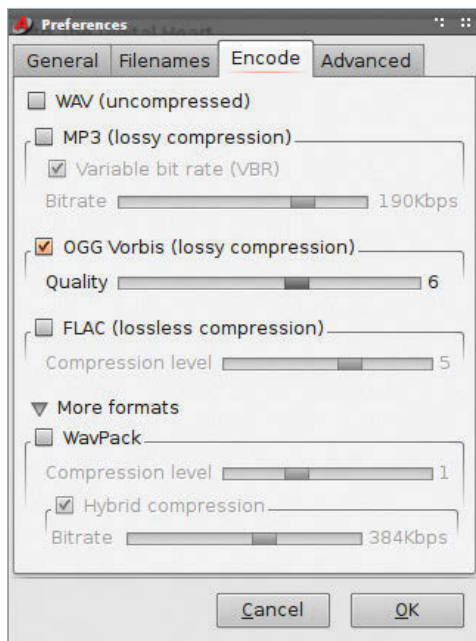
tasks at the package, including several scratched CDs and some which bore the dreaded 'will not play on PC/Mac' label which indicates intrusive and awkward digital rights management (DRM) technology, Asunder didn't balk once.

That's not to say everything went smoothly, however. While Asunder was happy to rip our scratched CDs, several tracks showed signs of skipping when played back – indicating that, despite its use of the cdparanoia ripping engine, the system was skipping sectors after a number of failed reads.

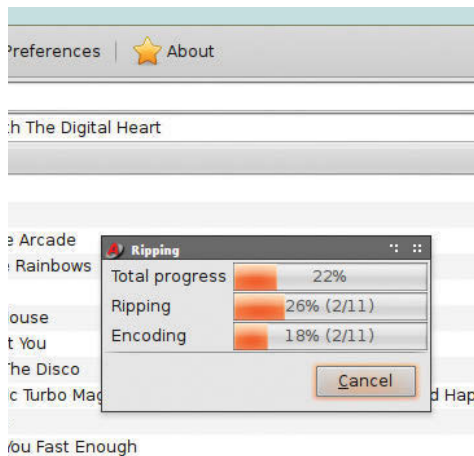
Unlike Audex, Asunder's simplicity became awkward at this point: using the GUI, there is no real way to get at the underlying functionality of cdparanoia to force more re-read attempts or tweak any other settings. As a result, Asunder might not be the best choice for those who don't look after their discs particularly well; although for non-scratched discs the default settings are acceptable, if a little slow.

That aside, Asunder proved itself an extremely capable package. Even adjusting the configuration of its various codecs was straightforward, although again some of the features more advanced users may like to toggle – such as joint stereo or CRC mode for MP3s – are abstracted away from the main interface.

The only real lack we found in Asunder was its inability to download cover art for ripped albums. With the majority of playback software supporting cover art – either stored within the metadata of individual files for some format, or as a JPEG image in the album's folder – it's a shame that more rippers don't include this functionality, which is only present in Audex from our test group.



■ Asunder's codec configuration is extremely user-friendly, if basic



SCORES

Ease of use

Asunder is friendly, but its simplicity can sometimes be awkward for power users

9

Features

Some more advanced features are abstracted away, while there's no cover image support

6

Speed

Asunder isn't fast, but it gets the job done so long as the CDs aren't too scratched

6

Compatibility

While no CDs failed to rip, some scratched examples did exhibit signs of skipping

9

Overall

For a lightweight, easy-to-use ripper, Asunder is hard to beat – but the option to download cover art would be a nice improvement

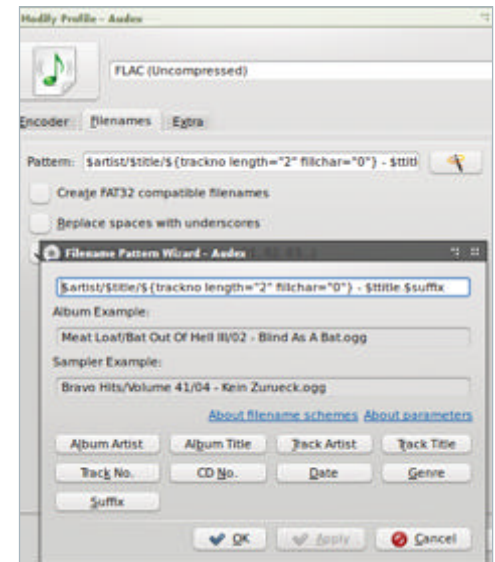
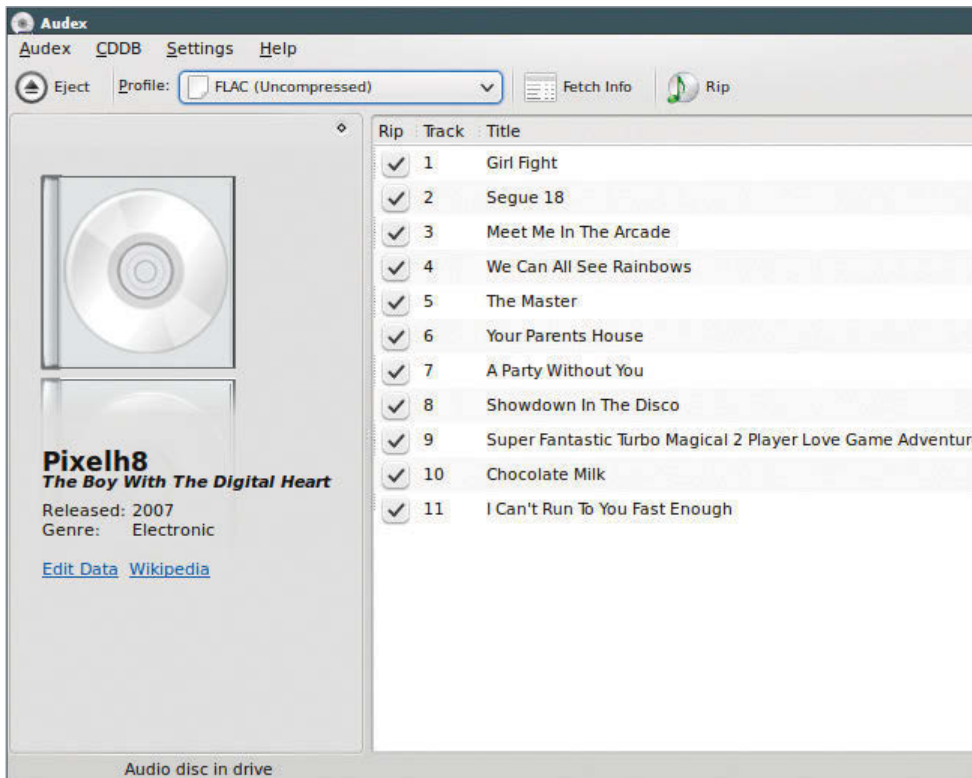
8

More information
<http://littlesvr.ca/asunder/>



Audex

Designed for KDE, Audex has had a somewhat troubled development history



■ Audex is an attractive application, but needs the KDE libraries



For a long time, it appeared that the Audex project had died: when maintainer Marco Nelles failed to hit a self-imposed release deadline for Audex

0.80 in January 2010 few worried, but when a new deadline of March was also missed concerns were raised. Nelles' silence was broken in January this year, however, with the release of Audex 0.74 Beta 1 – the version on test.

While installation can be awkward for those not on KDE, requiring the installation of numerous KDE libraries, initial impressions are worth it: Audex has by far the most polished user interface of any ripper on test, based around KDE 4. While this does result in a slow loading time, the application is fast and responsive during use; and extremely attractive.

As with the other packages on test, we set Audex to FLAC format at compression level five for the benchmark rip – using the slightly misnamed 'FLAC (Uncompressed)' preset – and were amazed when it finished its task in just two minutes and 44 seconds, nearly half that of the other packages on trial.

The software also flew through the other discs on test, including the DRM-encumbered samples which proudly boast of their incompatibility with

PC CD drives, until we came to the scratched discs. Here, sadly, Audex ran into its first real problem.

As with the other tools on test, Audex employs the popular `cdparanoia` system for the actual ripping process. This uses a variety of techniques in order to get the cleanest rip possible – in many cases successfully pulling audio data off a CD too scratched to play any other way. For most discs, Audex worked fine; on our most scratched example, however, the package hit an infinite loop whereby it would attempt to re-read a damaged sector over and over again, without even the ability to abort the process. A quick 'killall audex' got us out of the package, and disabling the on-by-default 'Never skip on read error' option resulted in a successful rip of the whole CD with only a few skips evident on the most damaged tracks.

Despite this, there's plenty to recommend Audex: fine-grained configuration is possible, if sometimes a little complex, while useful tools like filename repair functionality are included as standard. The album art downloader is also impressive, if a little hit-and-miss on less common CDs; with no other package on test even attempting to offer the feature, however, it's a welcome inclusion.

SCORES

Ease of use

By far the most attractive package on test, Audex can sometimes be confusing

8

Features

It's hard to fault Audex on the feature front, with album art support and plenty of codecs

10

Speed

Audex is blisteringly fast for most discs, until it hits a read error and hangs entirely

8

Compatibility

DRM is no barrier to Audex, although scratches can cause its default settings a few problems

8

Overall

For KDE users, Audex is a great choice. For others, the weight of the libraries may be a problem, but its feature set is second to none

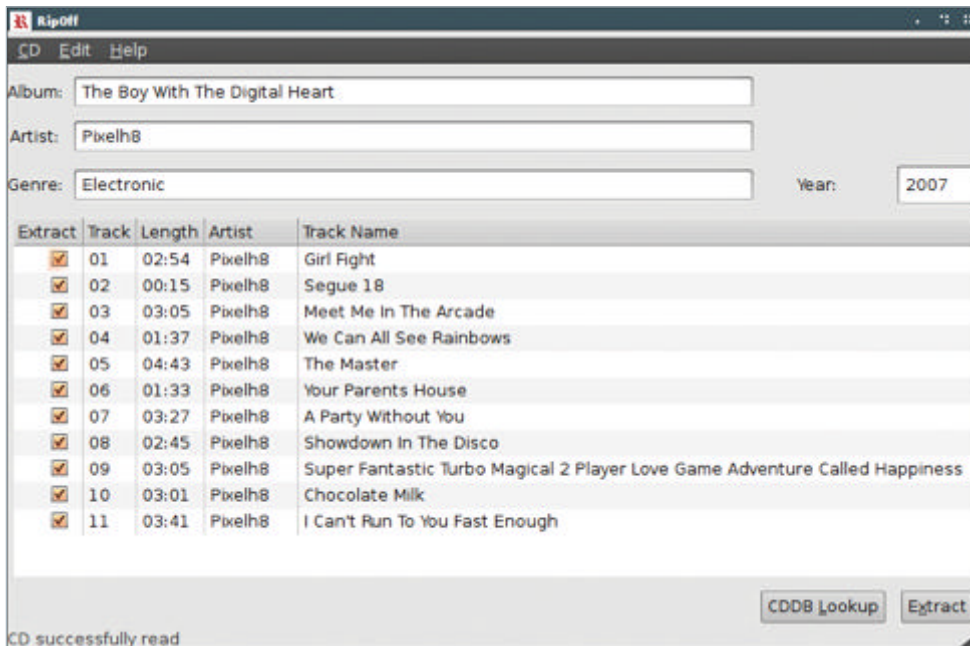
8

More information

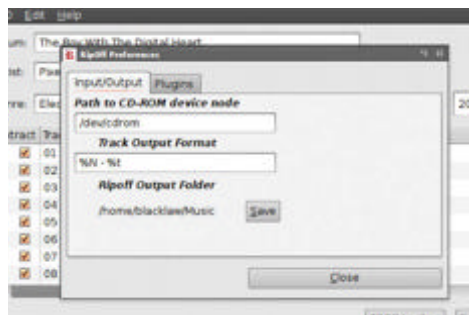
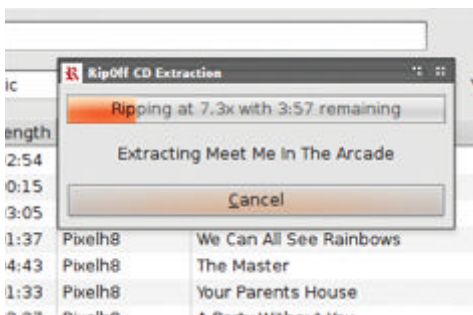
<http://kde.maniatek.com/audex/>

RipOff

It might be old, but RipOff remains a popular option for many



■ A spartan interface makes RipOff feel somewhat dated



While once a promising project, RipOff has been stagnating for a while: still officially a beta, the last release – RipOff 0.8.3 – was uploaded to SourceForge back in 2007, and there has been little work on addressing the extant bugs since; with one particular outstanding issue, caused by a loss of internet connection during a query of the CD DataBase (CDDb), making the application hang, that's a real shame.

That's not to say that RipOff doesn't work: despite a somewhat lightweight list of features and a very spartan interface, RipOff was happy with almost every CD we threw at it during testing. Scratched CDs ripped fine – although the worst did exhibit skipping at the usual points in the resultant FLAC

files, but RipOff is hardly alone here – while the benchmark test disc completed in a respectable four minutes and 12 seconds to give RipOff second place in the speed stakes.

Sadly, there is a problem at the heart of RipOff: unlike the other packages on test, it singularly failed to rip any of our DRM-encumbered 'copy protected' discs. While this will come as welcome news to those in the record industry who would rather we bought everything multiple times in multiple formats, it's bad news for anyone whose collection includes discs bearing the dreaded signs of digital rights management technology.

Codec support in RipOff is slim but bearable: the basics are all accounted for, while a plug-in infrastructure allows developers to easily code new formats should they so choose.

SCORES

Ease of use

A spartan interface looks simple enough, but configuration can be counter-intuitive

6

Features

RipOff is pretty basic, although developers may want to look at the plug-in support

4

Speed

Speed is RipOff's strong point, coming in second during the benchmark rip

7

Compatibility

Most CDs are supported as standard, but RipOff struggles with DRM

5

Overall

Once a stalwart tool, RipOff is showing its age. Unless you're nostalgic for simpler times, it's not recommended

5

More information
<http://ripoffc.sourceforge.net/>

From an end-user's perspective, however, RipOff isn't the easiest thing around: its basic interface hides things in bizarre places, with the codec selection choice being made in the plug-in tab of the configuration screen; once highlighted, a plug-in's particular codec is used for ripping, but there's nothing obvious in the UI to tell you that this is the case.

Configuring the output filename and directory options is equally hit-and-miss, with no explanation within the program itself of what the various flags do in the filename template configuration screen.

For those who need a particularly lightweight package, or developers looking to add support for an esoteric codec of their own devising, RipOff could be worth a look. The age of the package and its lack of features, coupled with its complete inability to retrieve audio tracks from the DRM-laden discs that are becoming all too common, make it a difficult program to recommend for general use, however.



RipperX

The oldest package on test, RipperX does a good job of keeping up with the kids



Registered with SourceForge back in 2000, RipperX is the oldest CD ripping tool to make it into our test group. Unlike RipOff, however, it has changed

with the times: while the most recent version dates back to November 2010, its icon-driven user interface is a lot more attractive than that of its spartan friend.

Sadly, positive initial impressions soon gave way to annoyance: while the basic user interface is attractive enough, attempting to configure various settings soon sees one run into a number of user experience confusions. Changing the codec settings to conform to our FLAC compression level five standard for the benchmark run, for example, saw us going into a tab marked 'MP3' - a counter-intuitive holdover from when the MP3 format was the only choice available for those ripping audio CDs to a compressed file.


Once configured, RipperX dealt with every CD we could find admirably; while its benchmark showing of four minutes and 59 seconds was the slowest of all the rippers on test, it was only 11 seconds slower than its nearest rival Asunder.

The poor speed showing of RipperX is something of a surprise: the software first rips to uncompressed WAV files, before running these through the chosen codec at the same time as the next track on the disc is being ripped. While it's a clever way of doing things, it clearly doesn't have as much impact on speed while using the lightweight FLAC encoder as it may do on more intensive lossy algorithms like OggVorbis or MP3.

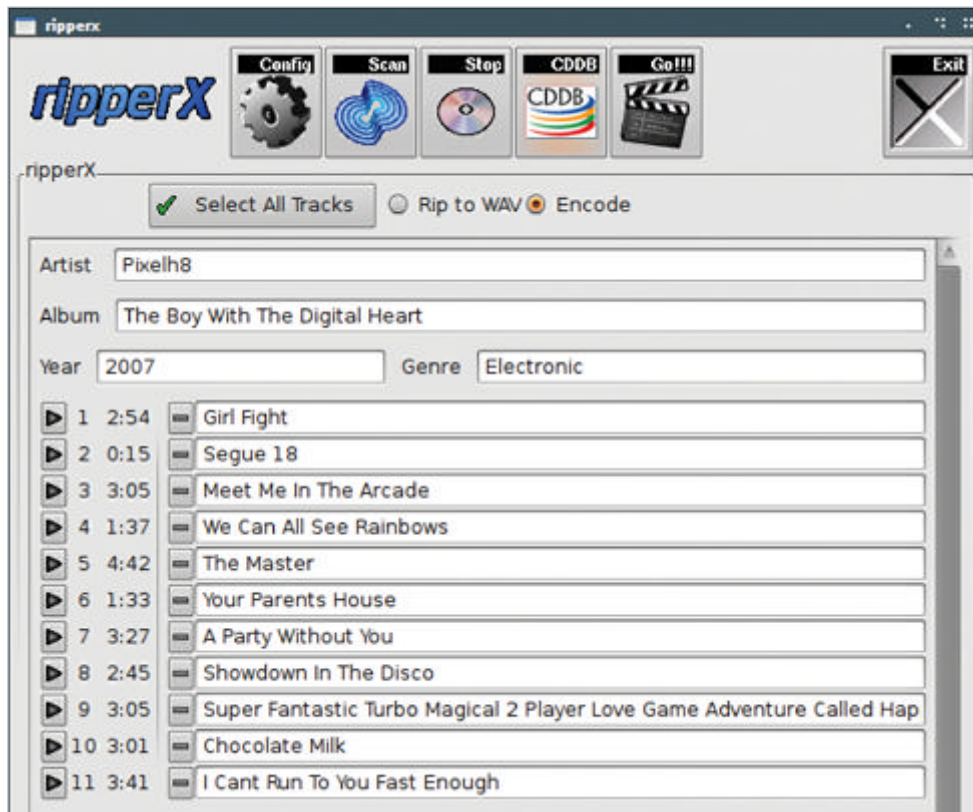
Despite being over a year since the last release, RipperX's three-year gap over RipOff shows: every CD we could lay our hands on with DRM technology included ripped perfectly, with the only issue on show a failure to load data from the CD DataBase. While manual data entry for the track listings is a pain, it's better than not being able to rip the discs at all.

As with RipOff, RipperX is fairly feature-light. There's no support for downloading or display of cover art imagery, for starters. But unlike RipOff it's relatively straightforward to customise the output filename format in RipperX - albeit it nowhere near as simple as the wizard-based option found in Audex. The wealth of configuration options

available for those using the Lame MP3 encoder is impressive, but support for other formats feels like an afterthought; the presence of a meaningless bitrate selector when FLAC is chosen, in particular, is likely to confuse many users.



■ RipperX includes a cdparanoia read-out at the bottom left during ripping



SCORES

Ease of use

The icon-based UI is attractive, but configuration can be confusing in the extreme

6

Features

Clearly designed for basic MP3 ripping, RipperX is somewhat light on features

5

Speed

RipperX was the slowest package on trial - albeit only by 11 seconds

5

Compatibility

Scratched discs and DRM pose little problem for RipperX

9

Overall

Poor user experience choices let an otherwise impressive package down, with configuration proving confusing for non-MP3 ripping

6


More information

<http://sourceforge.net/projects/ripperx/>

Email clients

Despite an increase in webmail use, desktop email clients are still proving popular – but which open source option offers the most flexibility?

Claws Launched in 2001 from the Sylpheed project, is Claws still leading the way?

 Just a few years ago, it wouldn't have made sense to have both Claws and Sylpheed in the same group test: the former was merely the development version of the latter, offering a sneak preview of technologies that would make it to the stable codebase over time.

Since 2005, Claws has been a standalone project in its own right and has diverged enough from its parent package to deserve an entry of its own in this group test.

There's certainly a lot to recommend Claws: compact and lightweight – the Debian packages are less than 2MB to download, and take up less than 4MB when installed – it's great for older systems, while the GTK2-based interface is quite happy in most desktop environments.

Sadly, certain parts of the program show their age: the setup process isn't very welcoming to newcomers, and it can be hard to figure out how to change existing account settings if you make a mistake on the initial wizard.

Once set up, however, Claws is still pretty sharp: the traditional three-pane layout – accounts and folders on the left, and emails plus a 'preview pane' on the right – is clear and easy to use, and the application uses minimal RAM.

Unfortunately, some of this is due to the package's lightweight feature set: by default, Claws will only read and compose plain-text emails – likely to be seen as a blessing by some and a curse by others –



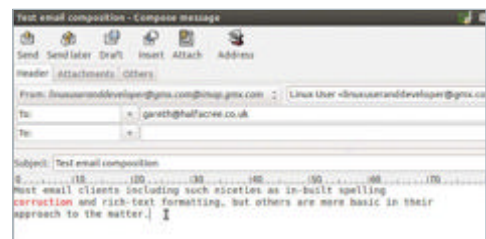
■ Without a plug-in, Claws doesn't render HTML emails

although plug-in packages are available to add this and other functionality to the client.

Despite this, Claws includes some nice features by default: emails written in its built-in editor come complete with a 'ruler' marking off columns at the top, while a spellcheck routine works in the background to highlight typos in an easy-to-read red colour.

The handling of attachments in Claws is also worthy of note: in addition to placing viewable attachments – such as JPEG images – below the message in both the preview pane and message windows, all attachments are seen as icons to the right of the message for one-click download or opening.

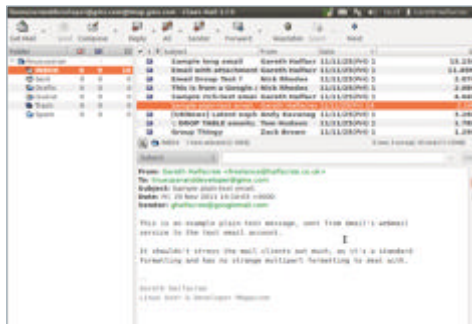
During testing, Claws didn't balk once at long emails and was the fastest by far to load our oversized 11.49MB message. Although configuration can be awkward, once running it works well – and any gaps in its feature set can be rounded out with optional plug-in packages.



■ Although basic, the plain-text Claws editor is clean and tidy



■ The icon-based approach to attachment handling makes Claws feel sleek



■ Claws handles plain-text emails well, ensuring useful headers are displayed

SCORES

Compatibility

Although initial setup can be awkward, POP3, IMAP and mbox accounts work great

7

Writing

It's a clean layout, but if you want rich-text handling you'll need a plug-in

6

Reading

Again, you'll need a plug-in for HTML emails, but the attachment handling system is attractive

7

Extras

The basic client is, understandably, basic, but a wealth of plug-ins add in new features

7

Overall

If you're running on older hardware or want a plain-text email client that flies, there's little to compare with the lightweight Claws

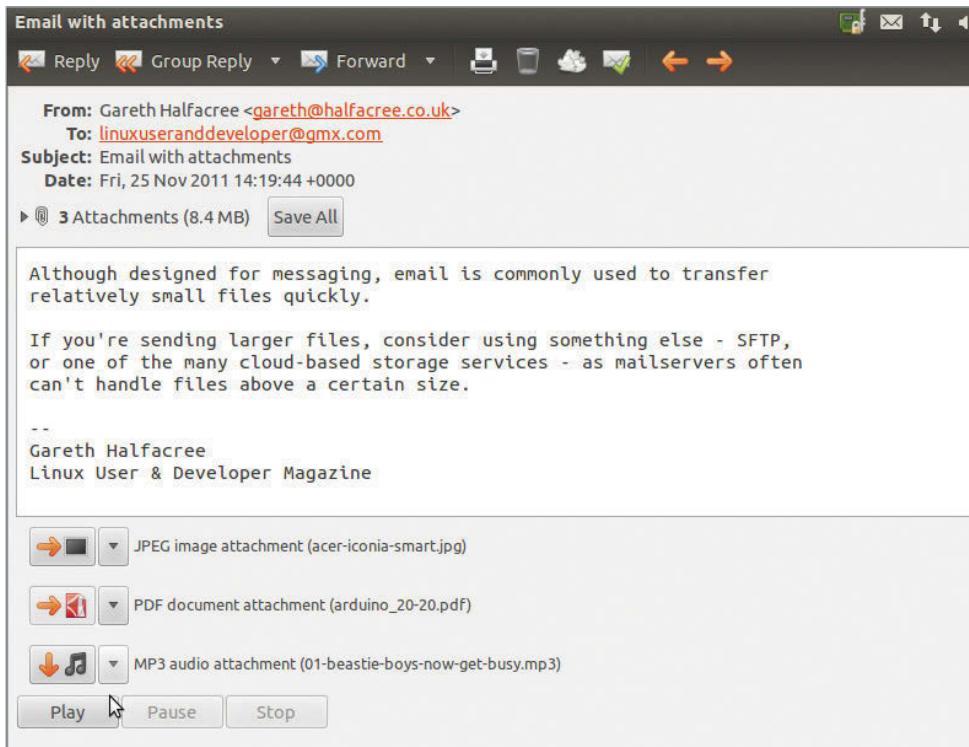
7

More information
<http://ripoffc.sourceforge.net/>



Evolution

Originally the default client for several distros, Evolution has fallen from grace



- There's no inline image viewing by default in Evolution, unlike other clients



It's not hard to see from where the creators of Evolution drew their inspiration: between the included calendar, which integrates nicely with the desktop on many platforms, to the slick interface and focus on impressive groupware facilities, the whole package screams Outlook.

The result is a mixed bag: there's no denying that Evolution is the most feature-rich client on test this month, but it comes at a heavy cost. Literally: although it takes up a mere 4MB of disk space, the Evolution client can cause older hardware to chug.

Once running, however, there's much to recommend the package – and for users in a corporate environment, Evolution's plug-in for Exchange server handling (older releases only, sadly) makes it a great alternative to Microsoft's Outlook.

Account setup is quick and easy, and Evolution supports more account types than any other client on test: whether you need Exchange, IMAP, IMAP+, POP3, Maildir, MH or mbox handling, Evolution steps up to the task with gusto. Full support for SSL and TLS security is, naturally, included as standard.


Sadly, things aren't quite as perfect once you get into the package proper. The three-pane view

is as clean and usable as any other client on test, but the email composition editor chooses some odd defaults: the built-in spelling checker opts for a yellow-on-white colour scheme for highlighting, which can make errors difficult to see on overly bright, washed-out or high-resolution displays – something from which other packages don't suffer, due to the use of red-on-white instead.


Attachment handling could be improved, too: by default, image attachments aren't shown inline, although a drop-down menu at the top of the email does mean you can access all attached files without having to scroll to the bottom of a message.

For some reason, Evolution was the slowest of all clients on test to load our 11.4MB email with several attachments. Although we're not talking hours – the email downloaded from the remote IMAP4 server and displayed in around 25 seconds – it was a delay compared to faster clients like Claws.

This is unlikely to be a concern for those choosing Evolution for its groupware facilities or Exchange support, but it's worth bearing in mind if all you're after is a simple email client, with the perceived 'bloat' directly leading to Canonical's decision to move to Mozilla's Thunderbird as the default client in Ubuntu from 11.10 onwards.



- The Evolution editor is clean, but spelling errors can be hard to spot



- Evolution's three-pane view is clean, if a little uninspiring

SCORES

Compatibility

Regardless of how your email is stored or accessed, you'll likely find support here

10

Writing

Aside from a daft yellow-on-white colour scheme for the spelling checker, composition is good

8

Reading

Attachment handling could be improved, but support for HTML emails is included as standard

8

Extras

If you need groupware facilities or a calendar, Evolution is an excellent choice

8

Overall

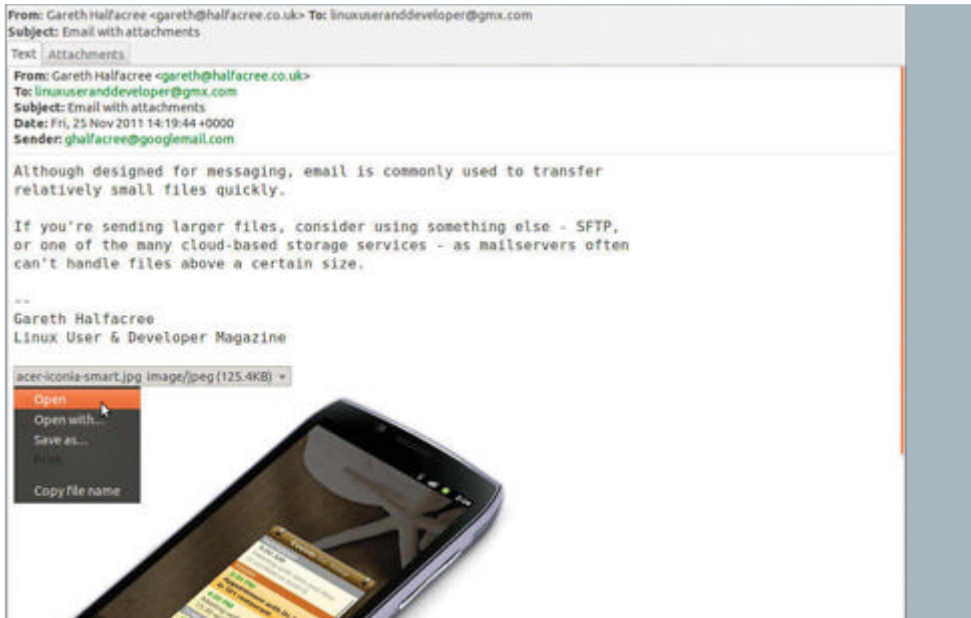
Evolution is doubtless impressive, but as increasing numbers of features are included it becomes less attractive as a pure email client

8

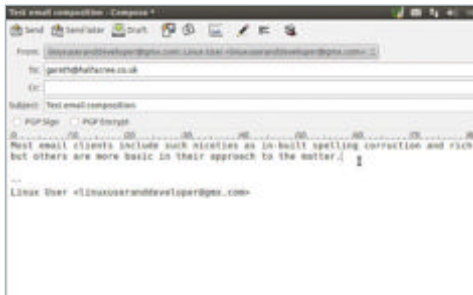
More information

<http://projects.gnome.org/evolution/>

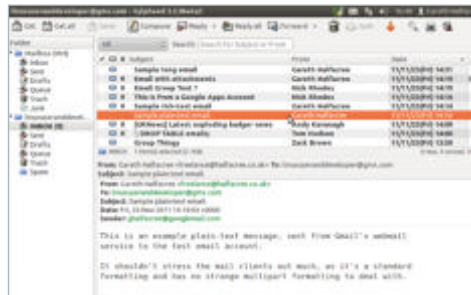
Sylpheed Can the venerable old project still hold its own against the newcomers?




■ Sylpheed lacks the icon-based attachment handling of its successor, Claws



■ With no built-in spelling check by default, you need to compose your emails carefully!



■ Sylpheed is well laid out and includes all relevant headers in the standard view

 **Sylpheed is the oldest open source mail client on test, with its first public release dating back all the way to the year 2000.**

In the intervening years, the project has seen plenty of competition, but – with the most recent stable release appearing in September this year – is still going strong.

The forking of the Sylpheed-Claws development release in 2004 did, it must be said, cost the main Sylpheed project dearly: with the departure of several developers keen to add new features to the codebase, Sylpheed is sometimes seen as the Miss Havisham (“shut up in these rooms a long time”) of email clients – but that’s decidedly unfair.

While Sylpheed may lack some of the flair of its offspring, nobody can deny that it’s lightweight and blazingly fast, even on older hardware. Although

only supporting POP3 and IMAP mail accounts, the developers have thoughtfully included Gmail-specific settings for those looking to use a desktop client to access their Google-provided email accounts.

Sadly, Sylpheed is definitely showing its age now: although similar in appearance to Claws, many of that package’s more impressive additions – such as the spelling check facility that operates as you type an email, and the icon-based attachment handling system – are missing in action.

Like Claws, Sylpheed includes a plug-in system to add functionality to its sparse feature set, but there isn’t much to see: plug-ins are thin on the ground, with Claws easily outdistancing its predecessor in this area. If you’re on old or space-restricted hardware, however, Sylpheed becomes a more

SCORES

Compatibility

If you use anything other than POP3 or IMAP, you’re out of luck in Sylpheed

6

Writing

Similar in appearance to Claws, you’ll soon find the lack of a spellcheck feature grating

4

Reading

Inline display of images is included, but again Claws outshines its predecessor.

5

Extras

Some useful plug-ins are available, but nowhere near as many as for Claws or Thunderbird

5

Overall

It’s the oldest client on test, and also the most dated feeling – but if you want a no-nonsense email client, still worth a look

5

More information
<http://projects.gnome.org/evolution/>



Thunderbird

Born out of the Mozilla Foundation's Firefox project, Thunderbird is still popular



The Mozilla Foundation's Thunderbird project – first released in beta form in 2003 and now up to version 8.0 as it adopts the same high-speed incremental numbering system as the Firefox browser project – is one of the most popular email clients around, and it's easy to see why when you first load it up.

While the project has come in for criticism in recent versions – with some pointing to the software's growing size and complexity as evidence of 'bloat' – it's an amazingly user-friendly experience.

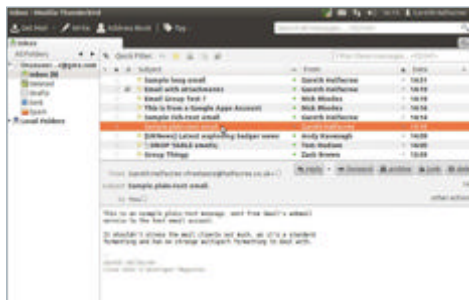
The Mozilla team have developed a method of determining the server settings required, either through a direct probe of the mail server or via a look-up database populated with configurations for almost every hosting provider and ISP in the world. These allow a newcomer to get up and running with the best – and the most secure – access possible, even if they don't know their IMAP from their TLS.

Although the by-now-familiar three-pane default layout is present and correct, Thunderbird includes some neat tricks up its sleeve for organising your mail. A powerful tagging system combined with search and filtering capabilities make it easy to keep things in order.

It's not all plain sailing, however: unlike Evolution, Thunderbird has little support for local mail accounts and no calendar and PIM (personal


information management) facilities without the use of a third-party plug-in.

That last point adds a new feature to the 'pros' column, however: Thunderbird shares the same plug-in engine as the Firefox browser, meaning that the client has a wealth of powerful add-ons to choose from, including support for GPG-signed messages – something the software lacks by default. Thunderbird was quick to load our test emails, and the layout works well. A potentially divisive feature that opens email in a new tab, rather than window, when double-clicked could lead to confusion, but helps to keep down on taskbar clutter when referring to multiple messages – and can be switched off if desired.



■ The standard three-pane view is unencumbered by dozens of icons at the top

The impressive user interface and list of features has led to Canonical adopting Thunderbird as the default email client in Ubuntu 11.10 and above – due, it claims, to user demand – but the rapid turnaround in major version numbers can cause system administrators headaches. Despite this – and following some glitches that turned users off previously – Thunderbird remains a powerful, client well suited to most users' requirements.



■ Thunderbird's composition window is clean and uncluttered; it also includes HTML support as standard

SCORES

Compatibility

The automated setup process is a joy, but local account handling leaves a lot to be desired

9

Writing

The composition window is neat and uncluttered, and supports HTML email by default

9

Reading

Tabbed 'browsing' is likely to be divisive, but format support can't be faulted

9

Extras

With the same plug-in engine as Firefox under the hood, there's plenty of scope for expansion

9

Overall

It's not quite perfect, and older hardware may struggle, but for many users Thunderbird is the perfect choice of email client

9

More information

www.mozilla.org/en-GB/thunderbird/




■ Thunderbird's attachment handling is good, but doesn't feel as slick as the icons in Claws

Password managers

Keeping track of your passwords needn't be a hassle, but which password manager should you trust with your most precious secrets, and why choose one above another?

KeePassX Originally a port of a Windows program, KeePassX is a handy cross-platform application

 While KeePass was originally developed for the Windows operating system, it proved useful enough to get ported to Linux as KeePass/L. Now available as a cross-platform application, it's been renamed to KeePassX – and if you're interested in security, it's well worth a look.

Although the project has been languishing – the last official update was in August 2010 and while the project's maintainer promised to be working on a 2.0 update, it has yet to materialise – the codebase is at a mature enough stage that it's ready for general use.

KeePassX has two features which help it stand out from other password-saving tools on test: an incredibly rich password generation system can be used to create secure passwords based on a variety of different schemes, with a built-in entropy gathering system to ensure that the passwords can't be guessed; and users have full control over where the password database is stored, along with the encryption strength used to protect it. When creating a database – which is held as a single file,

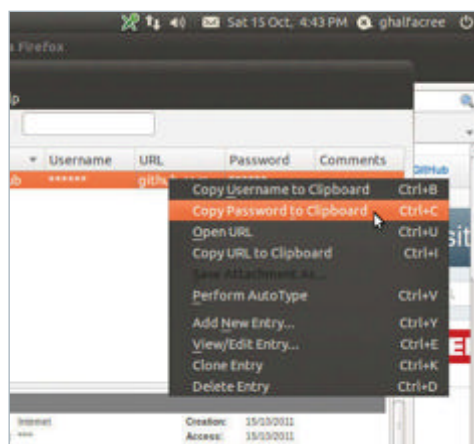
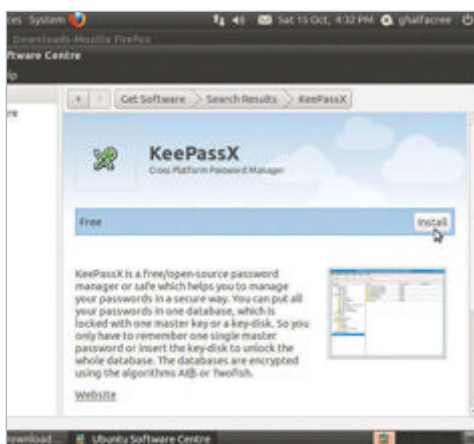
and can be copied to an external drive for added security – it's possible to choose the encryption scheme from AES or Twofish, along with the number of 'rounds' used. In addition, it's possible to use a keyfile which is required along with the password to unlock the database.

Although KeePassX offers impressive security, its features lag behind some of the other tools on test. There's no integration with browsers – one of the most common applications you'll be entering passwords into after system login – and while it includes an 'auto-type' system to fill in forms automatically, it's not straightforward to configure.

Instead, you'll find yourself spending most of your time copying and pasting passwords from the database. This isn't too onerous: KeePassX includes two keyboard shortcuts – Ctrl+C and Ctrl+B – for copying passwords and usernames respectively, and the copied password is cleared from the clipboard after use.

While other tools on test offer quicker ways to log you in, it's hard to fault KeePassX's approach – and when the database is kept on an external drive like

a USB memory stick, security is high. The cross-platform nature of the software also helps if you find yourself needing to log in from an unfamiliar system, and there are standalone binaries for many operating systems that will allow you to include the entire package on the same storage device.



■ KeePassX is available in the repositories of many distributions

SCORES

Features

The password generation functionality is great, but it lacks additional features

6

Ease of use

Simple copying and pasting is fine, but the 'auto-type' system can be awkward

6

Flexibility

Cross-platform support is a bonus, but there's little in the way of integration

5

Security

Good encryption and full control over the database location makes KeePassX pretty secure

9

Overall

There are better tools on test, but KeePassX is a great choice if you want full control over your passwords

7

More information
www.keepassx.org

LastPass

Is it really a good idea to trust your passwords to a cloud-based system?



If you're of the paranoid persuasion you may want to skip this part of the group test, as there's one integral feature of LastPass that is unlikely to sit well with your desire for security: the system stores your passwords remotely, on a server outside your control.

While there's no denying that such a move has a serious impact on security, LastPass uses hefty encryption and a 'master password' to lock down its data stores – and, despite a security scare in May this year when unauthorised traffic was noticed on several LastPass servers, that layer of protection has never been known to have been breached.

If you can get over the fact that your passwords are stored permanently on a remote system connected to the internet, there's a lot to like about LastPass: plug-ins are available for most common browsers, as are standalone programs and a web-based interface.

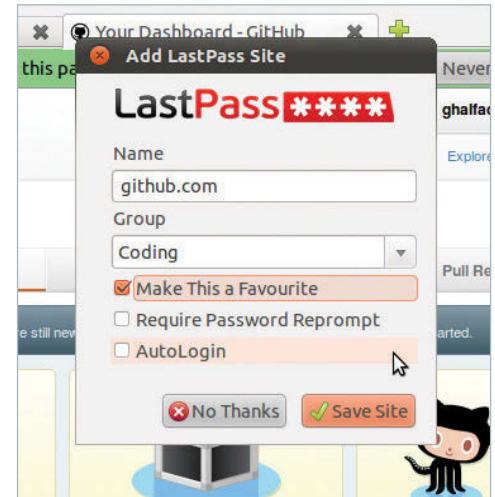
If you're willing to pay, you can even get LastPass support for mobiles, including iOS and Android-based devices – and there we come across another difference between LastPass and the other entries in this group test: it's a 'freemium' package, with extra facilities available to those who are willing to shell out \$1 a month for the privilege. Thankfully,

basic functionality – secure password storage and syncing – is always free.

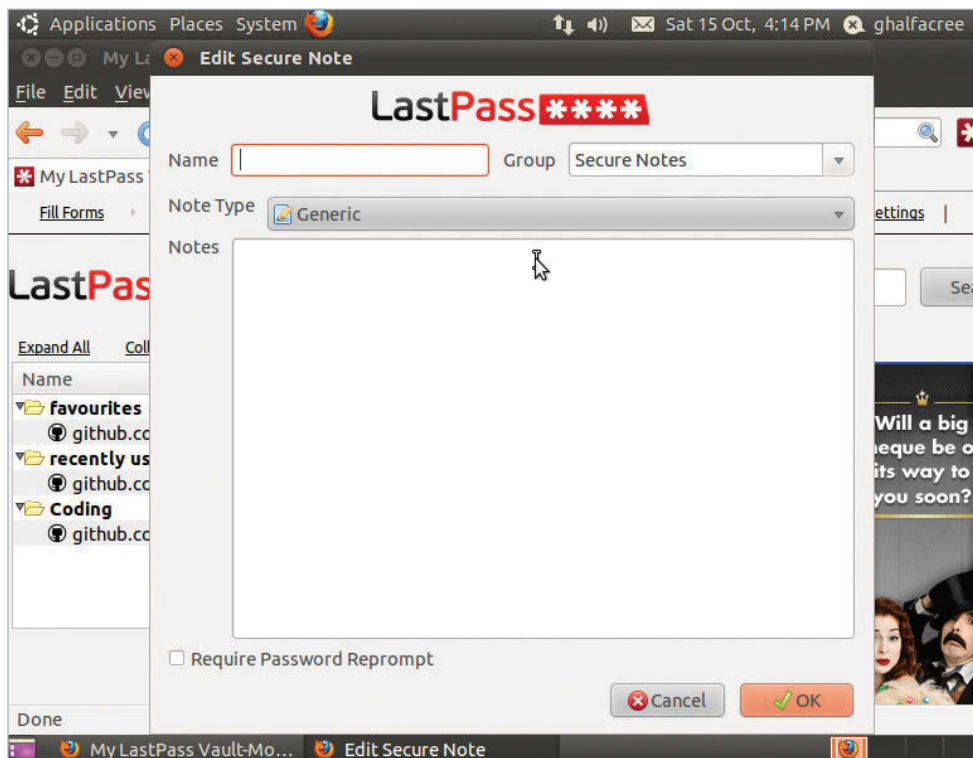
LastPass has another interesting feature: it completely replaces the internal password store in many browsers, removing what the company claims is an 'insecure' password storage system and replacing it with cloud-powered LastPass technology.

The result is a major win for usability: when the plug-in is installed in a compatible browser, users are able to see the same familiar interface and password-saving dialogs as before, but with improved encryption and cloud-based storage for retrieval on any device.

This latter feature is an absolute life-saver for anyone who regularly uses more than one system: remember a password on the browser of one computer, and that password is then available on any other computer you link to your LastPass account. With that said, there's still the worry that – however remote the chance might be – someone could one day successfully attack the LastPass server, reverse the salted encryption and make off with your precious passwords. For those looking for features, it's a great choice – but if you want guaranteed security, you should likely look elsewhere.



■ LastPass-stored passwords can be used to automatically log into sites



SCORES

Features

Password synchronisation is handy, and can be used to make secure notes too

8

Ease of use

If all you want is something that will work with your favourite browser, LastPass is perfect

8

Flexibility

LastPass is extremely flexible, but advanced features like mobile support require payment

7

Security

With passwords stored on a remote server, LastPass will always be at risk of attack

5

Overall

If you can trust the cloud to store your passwords, LastPass is a great choice

7

More information
<https://lastpass.com>

PasswordMaker

While PasswordMaker doesn't actually store passwords, it's a handy tool



As the name suggests, PasswordMaker is something of a wildcard entry in this group test: unlike every other tool, it doesn't actually contain any facility for saving your existing passwords.

Instead, it attempts to remove the need to ever save your passwords anywhere – and, as a result, it offers higher security than even the most comprehensive encrypted database can offer.

The principle behind PasswordMaker is simple: taking a master password – the only password you will need to remember – and the URL of the site you're visiting, the software uses a hashing function to generate a unique password for that particular site. The result: despite remembering only a single password, you have access to an infinite supply of long, secure passwords.

It's a clever tactic, and one which works well in practice: the browser plug-ins feature a 'Cool Key' quick-fill which allows you to enter your password and have it find and automatically populate password fields in web forms – and if there are two, such as when you're signing up for a new account, it will even fill both in at the same time.

The passwords generated can be customised to use different character sets and be different lengths, while it's also possible to switch from the default hashing algorithm – and you may need to do

so: on slower hardware, the hashing process can be slow, and with the algorithm kicking in each time you type a letter of your master password it can be awkward to enter on older systems.

It's also possible to set up different profiles with different rules, using less-secure passwords for some sites and longer ones for others. The hashing system also means that passwords are 'synced' between systems like LastPass, but without the security implications of storing them on a server somewhere: using the same rules, master password and URL, each system will generate the same unique password.

PasswordMaker does have its faults, however: if you have existing passwords, you'll need to replace them with PasswordMaker-generated versions if you want to make the switch, as there's no facility for entering your own passwords for the software to store. There's no faulting the developers for cross-platform support, however: versions of PasswordMaker are available for many mobile platforms, command-line use, multiple programming languages, most browsers – and for everything else, a downloadable HTML file that works in almost any browser on any device.

SCORES

Features

Although it doesn't support password storage, PasswordMaker is feature-rich and cross-platform

8

Ease of use

The autofill functionality of PasswordMaker is extremely quick to use

8

Flexibility

PasswordMaker does what it's designed to do very well, but manages little extra

5

Security

With passwords never stored locally or remotely, PasswordMaker is supremely secure

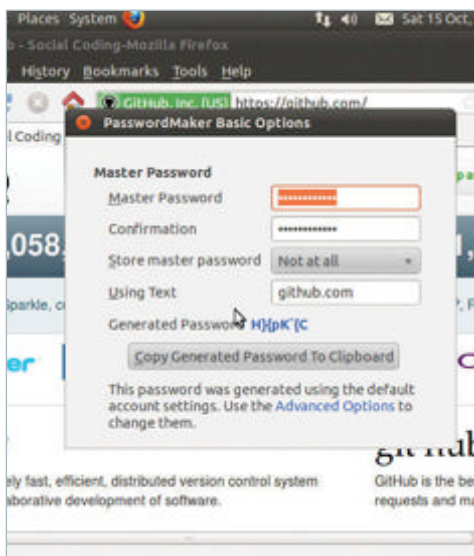
10

Overall

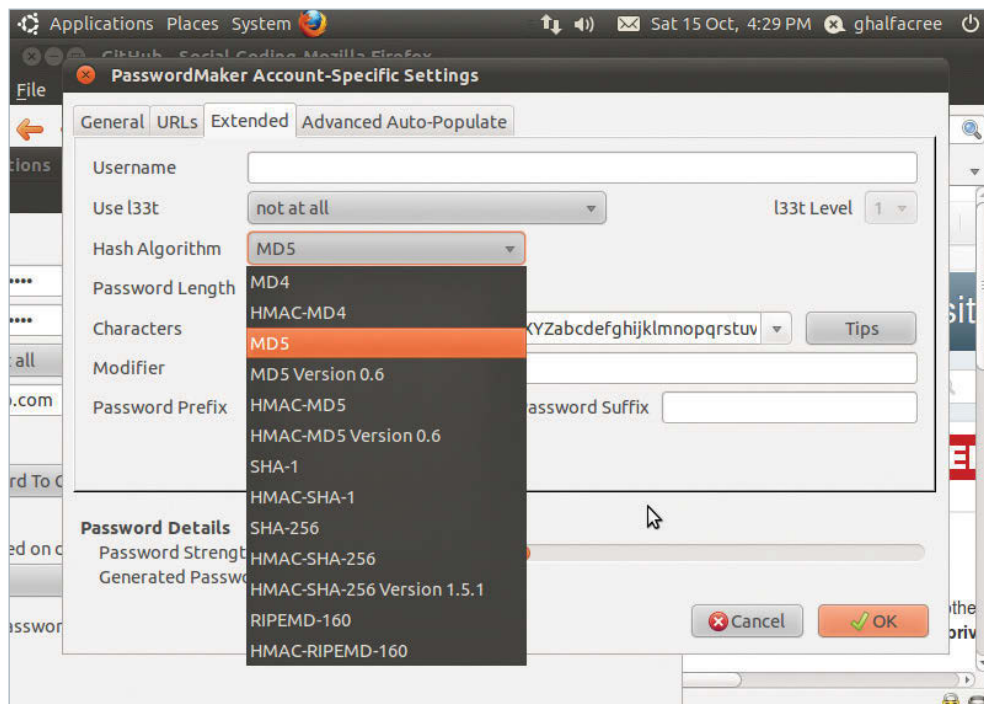
If you don't have a massive quantity of existing passwords, a move to PasswordMaker is well worthwhile

9

More information
<http://passwordmaker.org>



■ Passwords can be copied and pasted, or autofilled into forms



Seahorse

Originally developed for encryption keys, Seahorse also stores passwords



While it's true that Seahorse is primarily a tool for managing GPG and SSH keys, it's also possible to store your passwords in the software – and it comes with the advantage that if you're running a system based on GNOME, it's almost certainly already available.

Going under different names on different systems – and usually found under 'Preferences' as 'Passwords and Encryption Keys' – Seahorse is GNOME's password and key management system.

The first thing to be aware of with Seahorse: it's basic. Unlike other tools on test, it's not easy to customise with different encryption systems for storage, and neither does it connect to a remote server to synchronise passwords between client devices. There's no autofill functionality, and it isn't

particularly user-friendly in and of itself. That's not to say that it's not useful, however: as an integral part of GNOME, Seahorse ties in nicely to Nautilus and many web browsers, using its internal storage system in place of any standalone systems that existed before.

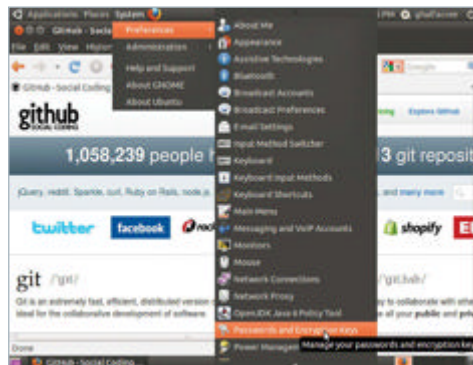
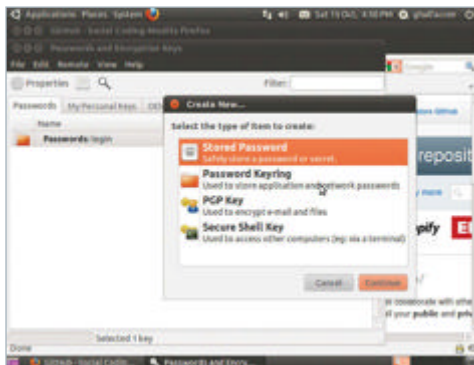
If you're a GNOME user, open up Seahorse now: you'll almost certainly find that it's already populated with your Wi-Fi network passwords, server passwords, instant messaging login details, and even your web passwords from Firefox and other browsers.

As a result, it's slightly hard to score Seahorse on usability: while the application itself is somewhat obtuse, it's designed to tie into other programs' own user interfaces – and the form-filling functionality of Firefox works fine, pulling information from Seahorse's keyring to log you into your favourite sites.

There's no getting away from the basic nature of Seahorse, though. There's little in the way of extras, and if you're not a GNOME user you'll either have to look elsewhere or install a whole lot of libraries to get it up and running.

The main benefit of Seahorse, however, is that it's almost certainly already there: if you're running GNOME, you already have it – and if you're willing to learn, it's possible to take keyrings and store them centrally, creating a homebrew LastPass-style synchronisation system for your passwords.

If you don't already – knowingly – use Seahorse, then it's not a good choice to start now: alternatives like KeePassX and LastPass offer far greater functionality or improved security. With that said, if all you need is an encrypted data store for usernames and passwords, Seahorse will deliver admirably.



■ Passwords can be viewed and copied directly from within Seahorse

SCORES

Features

There's not a lot in the way of extras, but Seahorse does its primary job well

5

Ease of use

Seahorse largely relies on individual applications to provide a helpful interface

5

Flexibility

For a tool originally designed for GPG/SSH keys, Seahorse is surprisingly flexible

7

Security

With an encrypted database stored locally, Seahorse is pretty secure

8

Overall

It's basic, but it does the job well – and if you're running GNOME, you've already got a copy


7

More information

<http://live.gnome.org/Seahorse>

Special
trial offer

Enjoyed
this book?



Exclusive offer for new



Try
3 issues
for just
£5*

*This offer entitles new UK Direct Debit subscribers to receive their first three issues for £5. After these issues, subscribers will then pay £25.15 every six issues. Subscribers can cancel this subscription at any time. New subscriptions will start from the next available issue. Offer code 'ZGGZIN' must be quoted to receive this special subscriptions price. Direct Debit guarantee available on request.

** This is a US subscription offer. The USA issue rate is based on an annual subscription price of £65 for 13 issues, which is equivalent to \$102 at the time of writing compared with the newsstand price of \$15.99 for 13 issues, being \$207.85. Your subscription will start from the next available issue.



About
the
mag



The magazine for the GNU generation

Written for you

Linux User is the only magazine dedicated to advanced users, developers & IT professionals

In-depth guides & features


Written by grass-roots developers & industry experts

4.5GB DVD every issue

Four of the hottest distros feature every month - insert the disc, reboot your PC & test them all!

subscribers to...

Linux User & Developer™



Try three issues for **£5 in the UK***
or just **\$7.85 per issue in the USA****
(saving 51% off the newsstand price)

For amazing offers please visit
www.imaginesubs.co.uk/lud

Quote code ZGGZIN

Or telephone UK 0844 249 0282 overseas +44 (0) 1795 418 661



✓ iPad ✓ iPhone ✓ Android phone ✓ Android tablet ✓ Apple Mac ✓ Windows PC



360 Magazine

3D Artist


No Disc. No Problem

Many of the files you're looking for can be found on the magazine's website

Imagine digital editions are a new and exciting way to experience our world-leading magazines and bookazines.

To get the most out of your digital editions, be sure to enjoy all of our fantastic features, including:

- Zoomable text and pictures
- In-app browsing
- Searchable text
- Take your collection with you
- Read offline

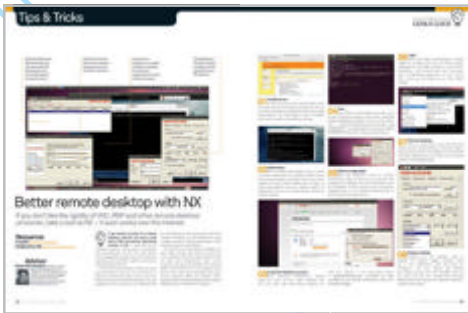


To buy more Imagine digital editions and for the latest issues and best offers, please go to

www.GreatDigitalMags.com

Linux for Professionals

The complete guide to using and developing for Linux, written by respected experts in the open source scene



App development

> Create your own Android app

Pro tips & tricks

> Tutorials for Django, Unity and more

Back up & restore

> Learn to keep your system secure

For all Linux users

> Work with Ubuntu, Linux Mint & Fedora

Software guides

> Discover the best media centres, email clients and password managers



django



Parted Magic 6.7



Fedora 17



Linux Mint 13



openSUSE 12.1

Volume 3

£14.99



IMAGINE PUBLISHING

ISBN-13: 978-1908955319



www.imaginebookshop.co.uk