Full Circle

Photo: Auntie P (Flickr.com)

TRUECRYPT
IS THERE A SAFE ALTERNATIVE?
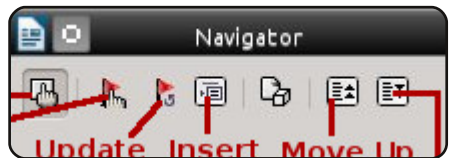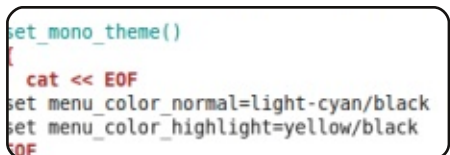
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

Navigator
Update Insert Move Up

```
#An alias to make the ls
command more detailed
alias ls = "ls -la --
color=always --classify"
```

```
set_mono_theme()
  cat << EOF
set menu_color_normal=light-cyan/black
set menu_color_highlight=yellow/black
OF
```

Desktop
Ubuntu Kylin
Alternative downloads

NEW

Puppet Reporting and Monitoring
Create insightful reports for your server infrastructure using Puppet

EDITORIAL

WELCOME TO ANOTHER ISSUE OF FULL CIRCLE.

Almost a full house again this month. There's Python, LibreOffice, part three of the GRUB boot menu series with the last part coming next month. Inkscape is here, but no Blender this month I'm afraid.

As well as Charles' regular Linux Labs I've added another Linux Labs article of my own on using BitTorrent Sync. BitTorrent always seems to get a bad rap due to it being used for illegal means, but this new Sync thing really works quite well for sharing files between computers and even mobile devices. If you're thinking of moving from DropBox, or the now defunct Ubuntu One, but want something that's not hosted then I'd recommend giving BitTorrent Sync a look. I believe there is a similar fully F/OSS alternative out there, but I haven't tried it as yet.

The main focus though, this month, is a look at alternatives to TrueCrypt. It seems that TrueCrypt is (depending on who you listen to) either insecure, and/or not being developed further. This leaves a lot of encrypted drives either without support or possibly being insecure, so many will need to find something else. Iain Mackeand has tested several alternatives and shows you how each is used. Although, I have to say, not many look as friendly, and GUI based, as TrueCrypt is/was.

All the best, and keep in touch!
Ronnie
ronnie@fullcirclemagazine.org

Full Circle Podcast

Released monthly, each episode covers all the latest Ubuntu news, opinions, reviews, interviews and listener feedback. The Side-Pod is a new addition, it's an extra (irregular) short-form podcast which is intended to be a branch of the main podcast. It's somewhere to put all the general technology and non-Ubuntu stuff that doesn't fit in the main podcast.

Hosts:
• Les Pounder
• Tony Hughes
• Jon Chamberlain
• Oliver Clark

http://fullcirclemagazine.org

AUDIO MP3   AUDIO OGG

Download

THE NSA THINKS YOU ARE AN EXTREMIST IF YOU CARE ABOUT PRIVACY

You don't have to perform a search about pressure cooker bombs or terrorism to get on the National Security Agency's radar. Just being a little privacy-conscious and curious about Tor is enough.

Last summer, SecurityWatch wrote about the NSA's X-Keyscore, a program which sucked in large amounts of Web data—from emails, Facebook activity, and online surfing, to name a few—for analysts to search and find patterns. A group of journalists and activists have analyzed X-Keyscore's source code, and discovered the rules used to determine which individual should be subject to deep surveillance are broader than we anticipated, according to a report from German site Tagesschau. Jacob Appelbaum, a member of the Tor Project and one of the authors, also posted an English version on Der Este.

Source: http://securitywatch.pcmag.com/privacy/325273-the-nsa-thinks-you-are-an-extremist-if-you-care-about-privacy

Submitted by: **Arnfried Walbrecht**

KOREAN GOVT TO TURN ITS BACK ON MICROSOFT... AND USE WHAT INSTEAD?

As has been reported widely in Asia, South Korea is using the fact that Windows XP is no longer supported as a reason to walk away from Microsoft completely. While the rationale for this stance makes sense; updating to new versions of the Microsoft software can be expensive and time-consuming, the government should also immediately discontinue its use of Hangul. Hangul is a Korean software package similar to MS Office that is used as the default (and required) file format for most government departments. It is a major frustration for anyone submitting documents to the government, who don't have the peculiar Korean software installed.

Source: http://e27.co/korean-government-to-turn-its-back-on-microsoft-and-use-what-instead-hangul-20140703/

Submitted by: **Arnfried Walbrecht**

NSA TARGETS LINUX JOURNAL AS 'EXTREMIST FORUM': REPORT

The Linux Journal, a Linux user community website, has been flagged as an "extremist forum" by the United States' National Security Agency (NSA), while its users have been flagged as "extremists" under the agency's XKeyscore program, according to leaked source code.

The source code, which was published this week by German public broadcaster, ARD, also identified at least two German Tor Directory Authority servers — one in Berlin, the other in Nuremberg — as being under surveillance by the NSA.

The Tor Project is an independent, open-source anonymous software and browsing network that directs traffic through its free, worldwide, volunteer network to conceal users' locations and usage from surveillance.

Source: http://www.zdnet.com/nsa-targets-linux-journal-as-extremist-forum-report-7000031241/

Submitted by: **Arnfried Walbrecht**

THE LINUX FOUNDATION ANNOUNCES EARLY KEYNOTE SPEAKER LINE UP FOR LINUXCON + CLOUDOPEN EUROPE

The Linux Foundation, the nonprofit organisation dedicated to accelerating the growth of Linux and collaborative

development, today announced the keynote speakers for LinuxCon + CloudOpen + Embedded Linux Conference Europe, taking place 13-15 October at the Congress Centre Düsseldorf.

LinuxCon Europe is the place to learn from the best and the brightest, delivering content from the leading maintainers, developers and project leads in the Linux community. There's simply no other event in Europe where developers, sysadmins, architects, and all types and levels of technical talent gather together under one roof for education, collaboration and problem-solving to further knowledge of Linux and open source collaboration in general. LinuxCon will feature over 100 sessions with innovative content on the latest kernel updates, storage technologies and interfaces, security, Internet of Things, and talks related to open source collaboration and best practices.

CloudOpen Europe features technical content that addresses open cloud platforms and developer tools. It is the only venue where various open cloud technologies come together under one roof to advance Apache Stratos, AWS, Chef, CloudStack, Docker, Gluster, Hadoop, Juju, KVM, Linux, MariaDB, MySQL, OpenDaylight, OpenStack, oVirt, Puppet, SaltStack, the Xen Project, and more.

Source:
http://www.koreaittimes.com/story/38605/linux-foundation-announces-early-keynote-speaker-line-linuxcon-cloudopen-europe

Submitted by: **Arnfried Walbrecht**

VALVE UPDATES STEAMOS WITH THE LATEST NVIDIA, AMD, AND INTEL DRIVERS

The Beta version of SteamOS, a Debian-based distribution developed by Valve to be used in its hybrid PC/console, has just received an update and numerous packages.

Valve has two builds for SteamOS. One is a stable version (sort of) and the other one is a Beta (Alchemist). The two versions are not all that different from one another, but the Valve developers are using the Beta release to test some of the new updates before they hit the stable branch.

This is just the Beta version of SteamOS and not all of the packages included are stable. It will take a while until all these changes will be added to the Stable branch. The system requirements for Steam OS haven't changed and have been pretty much the same since the beginning: an Intel or AMD 64-bit capable processor, 4GB or more memory, a 250GB or larger disk, NVIDIA, Intel, or AMD graphics card, and a USB port or DVD drive for installation. Check the official announcement for more details about this release.

Source:
http://news.softpedia.com/news/Valve-Updates-SteamOS-With-the-Latest-NVIDIA-AMD-and-Intel-Driver-448769.shtml

Submitted by: **Silviu Stahie**

LINUX KERNEL 3.15.3 IS NOW AVAILABLE FOR DOWNLOAD

Greg Kroah-Hartman had the pleasure of announcing earlier today, July 1, that the third maintenance release for the current stable 3.15 branch of the Linux kernel is available for download, urging users to upgrade as soon as their Linux distributions update the respective packages on the official software repositories.

The Linux kernel 3.15.3 is a pretty standard release that introduces various updated drivers, some filesystem improvements, especially for Btrfs and EXT4, random mm and Bluetooth fixes, and the usual architecture enhancements (ARM, ARM64, IA64, SPARC, PowerPC, s390, and x86).

Be aware, though, that upgrading to a new Linux kernel package might break some things on your system, so it is preferable to wait a few days and see if anyone complains about it on the official channels of your distribution.

Source:
http://news.softpedia.com/news/Linux-Kernel-3-15-3-Is-Now-Available-for-Download-448998.shtml

Submitted by: **Marius Nestor**

XFCE APP LAUNCHER `WHISKER MENU` SEES NEW RELEASE

Whisker Menu is an application menu / launcher for Xfce that features a search function so you can easily find the application you want to launch. The menu supports browsing apps by category, you can add applications to favorites and more. The tool is used as the default Xubuntu application menu starting with the latest 14.04 release and in Linux Mint Xfce starting with version 15 (Olivia).

The Whisker Menu PPA was updated to the latest 1.4.0 version recently and you can use it to both upgrade to the latest version obviously, as well as to install the tool in (X)Ubuntu versions for which Whisker Menu isn't available in the official repositories (supported versions: Ubuntu 14.04, 13.10 and 12.04, and the corresponding Linux Mint versions). To see what is different from the previous release, see the changelog in its main website.

Source:
http://www.webupd8.org/2014/06/xfce-app-launcher-whisker-menu-sees-new.html

Submitted by: **Andrew**

WHY DID MICROSOFT JOIN THE LINUX FOUNDATION'S ALLSEEN ALLIANCE?

When people think of open source, they don't usually associate Microsoft with it. But the company recently surprised many when it joined the Linux Foundation's open source AllSeen Alliance. The AllSeen Alliance's mission is to create a standard for device communications.

Has Microsoft changed its attitude toward open source in general, or is there another reason for its uncharacteristic behavior? Computerworld speculates on what might have motivated Microsoft to join the AllSeen Alliance.

Source:
http://www.itworld.com/open-source/425651/why-did-microsoft-join-linux-foundations-allseen-alliance

Submitted by: Jim Lynch

PHP FIXES OPENSSL FLAWS IN NEW RELEASES

The PHP Group has released new versions of the popular scripting language that fix a number of bugs, including two in OpenSSL. The flaws fixed in OpenSSL don't rise to the level of the major bugs such as Heartbleed that have popped up in the last few months. But PHP 5.5.14 and 5.4.30 both contain fixes for the two vulnerabilities, one related to the way that OpenSSL handles timestamps on some certificates, and the other also involves timestamps, but in a different way.

Source:
http://threatpost.com/php-fixes-openssl-flaws-in-new-releases/106908

Submitted by: **Dennis Fisher**

CHROME REMOTE DESKTOP FREE DOWNLOAD FOR WINDOWS PC

Chrome remote desktop is a remote access tool and can be used to access other desktops from your computer.

The prerequisite to using this remote desktop software is that both the computer used for controlling and the computer being controlled must have the Google Chrome browser installed on them. You must also have an extension installed on both the machines. This extension can be downloaded from Chrome Web Store. It is a good idea to first install the extension on the computer being controlled. An interesting feature of the Chrome remote desktop tool is that it can be used to control any of your computers if you have more than one. Similarly, it can also be used to access a friend's computer.

Source:
http://thefusejoplin.com/2014/07/chrome-remote-desktop-free-download-windows-pc/

Submitted by: **Arnfried Walbrecht**

Build Your Own Handheld "Game Girl" For All Your Retro Gaming Needs

Calling all makers! Do you like to play with electronics? Do you have access to a 3D printer? If so, you can make your very own DIY handheld gaming device, thanks to the mad geniuses at Adafruit.

The Game Girl (also known as the PiGRRL) is, in essence, a Raspberry Pi Linux computer in Game Boy-style housing, made with some scavenged components from a SNES controller. Though an 8-bit NES emulator is the obvious software choice here, you can run whatever you fancy on it. Have a look at the video below for an overview of the project, or head to Adafruit for detailed instructions.

I should also mention that this project was created in celebration of the Game Boy's 25th anniversary. I feel old.

Source:
http://www.themarysue.com/diy-game-girl/

Submitted by: **Arnfried Walbrecht**

CoreOS Linux Released as 'World's First OS as a Service'

A Linux distribution that updates itself. That's the idea behind the new commercial Linux product from CoreOS, called CoreOS Managed Linux.

Source:
http://www.toptechnews.com/article/index.php?story_id=132007RMRPWO

Submitted by: **Peter Odigie**

Open Xchange: The internet wouldn't have happened without Linux

RAFAEL LAGUNA, CEO of open source cloud productivity suite Open Xchange (OX), has told The INQUIRER about his hopes for a more open internet.

Open Xchange is a cross-platform service that provides a range of web apps, cloud storage, encrypted email and instant messaging. The service can be downloaded - white labeled, free of charge, and self-configured - with the company's business model based on optional support and configuration. It is licensed to a range of ISPs and telcos across Europe who are running on Linux systems.

Laguna asks, "What would have happened to the internet if there had been no Linux? What if it had been set up based on Windows servers and Unix servers from Sun or IBM?

"The internet simply wouldn't have happened. It would have been way too expensive. Companies would have tried to control who had the internet and who didn't. It would have prohibited the internet from ever coming into existence."

Source:
http://www.theinquirer.net/inquirer/feature/2354003/open-xchange-the-internet-wouldnt-have-happened-without-linux

Submitted by: **Arnfried Walbrecht**

Cisco: Linux is the "Single and Best" Tech for IoT

Cisco earlier this year unveiled its plans to build smarter routers and switches to help manage the massive flows of data expected between Internet-connected devices and the data center. This re-architecting of the Internet to bring computing capabilities to the edge of the network is what the company calls "fog computing" and it could help alleviate the data center strain that Gartner analysts predict will come from 26 billion installed units in the Internet of Things by 2020.

"Tens of billions of ever-smarter edge devices create data (we call it big data now) that it is impossible to move fast enough through the network," said Michael Enescu, CTO of Open Source Initiatives at Cisco. "We have to deal with this, otherwise it will present huge scale and security issues in the core."

Source:
http://www.linux.com/news/featured-blogs/200-libby-clark/779505-cisco-linux-is-the-single-and-best-tech-for-iot

NEWS

Submitted by: **Arnfried Walbrecht**

LINUX FOUNDATION INTRODUCES LINUX FOR CARS

The connected car is shifting into high gear, and the Linux Foundation wants an open-source platform in the pole position. The non-profit consortium recently announced the debut of Automotive Grade Linux (AGL), a customizable, open-source automotive software stack with Linux at its core.

There have been Linux-based car systems before. The difference here is that the Linux Foundation is pushing AGL as the core automotive system for developers – the groundwork, as opposed to a production-ready system like Ford Sync or QNX. As Debian Linux is to Ubuntu, so the AGL aims to be for a future in-car system from the Hyundais and Toyotas of the world.

The AGL does come with a core set of capabilities. The platform is based on the Tizen In-Vehicle Infotainment (IVI) Project (yes, that Tizen) and includes a variety of in-car software for climate control, maps, dashboard displays, media playback, and smartphone links. The Linux Foundation says the AGL can also support mission-critical systems, such as roadway instrumentation.

Source:
http://www.itworld.com/operating-systems/425696/linux-foundation-introduces-linux-cars

Submitted by: **Arnfried Walbrecht**

MINIX CREATOR TANNENBAUM TO RETIRE

Professor Andrew Tannenbaum, creator of the Minix operating system, has announced that he will retire from the University of Vrije in the Netherlands, where he has worked for 43 years.

Tannenbaum wrote a clone of UNIX called Minix in 1987 and published a book about the design and implementation of operating systems, listing the source code in the appendix. His book served to help Finnish developer Linus Torvalds create the Linux kernel which initially used the Minix file system.

Soon after the release of Linux in 1991, Tannenbaum got involved in a debate about the relative merits of Minix and Linux; prominent among those who responded to his comments was Torvalds himself.

Tannenbaum was born in New York but moved to the Netherlands with his Dutch wife after he received his doctorate. At the Vrije University, he taught courses about computer organisation and operating systems and also supervised students who were working towards doctorates.

He has written a number of widely used books on computer science, some of which have been translated into as many as 21 languages.

Source:
http://www.itwire.com/business-it-news/open-source/64719-minix-creator-tannenbaum-to-retire

Submitted by: **Arnfried Walbrecht**

CROSSOVER 13.2: INSTALLING WINDOWS PROGRAM ON LINUX JUST GOT EASIER

Sometimes, you really need a Windows desktop application, even if you're a die-hard Linux user. For those times, I recommend CodeWeavers' CrossOver. Better still, the latest version makes installing Windows programs on Linux systems easier than ever.

Microsoft Office 2010 running on Linux? Sure, it's easy to do with CodeWeavers' CrossOver.

For those of you who haven't used CrossOver before, it enables you to run some, but not all, Windows programs on both Linux PCs and Macs. The program is based on the open source project Wine, an implementation of the Windows application programming interface (API) on top of the Unix/Linux operating system family. For each application, CrossOver/Wine creates a small Windows virtual container for just that program.

You actually don't need

CrossOver to run Windows applications on Linux or a Mac. You can do it with Wine alone — if you know precisely what you're doing. But CrossOver gives you an easy, automated installation of Windows applications, and technical support.

Source:
http://www.zdnet.com/crossover-13-2-installing-windows-program-on-linux-just-got-easier-7000031505/

Submitted by: **Arnfried Walbrecht**

Intel to ship Galileo Gen2 open-source computer in August for $60

Intel's answer to the popular $25 Raspberry Pi credit-card sized PC, the Galileo Gen2, is set to be available in August for around $60.

Like the Pi, the Galileo Gen2 computer is an uncased board with all the components necessary to build a PC. Only external peripherals need to be attached. Galileo is targeted at the community of do-it-yourselfers and enthusiasts who make robots, small electronics, wearable devices and even PCs.

The Gen2 succeeds the first Galileo computer, which started shipping late last year for around $70. The new Galileo has a low-power x86 processor called Quark, while the Raspberry Pi has an ARMv7 processor.

Source:
http://www.computerworld.com/s/article/9249741/Intel_to_ship_Galileo_Gen2_open_source_computer_in_August_for_60

Submitted by: **Arnfried Walbrecht**

Android, Arduino development board goes on sale

The UDOO prototyping board for Android, Linux, Arduino and Google ADK applications is available from German distributor Conrad Business Supplies.

Created by SECO and AIDILAB, the development board is based on a dual or quad core ARM cortex-A9 CPU with a dedicated ARM processor for IO management.

Expected to be used by students and development engineers, the board runs an ARM-optimised version of Linux Ubuntu called Linaro. This means developers can use the Arduino IDE to build and upload sketches to the Arduino-compatible embedded board, without additional or external cable connections.

Source:
http://www.electronicsweekly.com/news/design/embedded-systems/android-arduino-development-board-goes-sale-2014-07/

Submitted by: **Arnfried Walbrecht**

New Mayhem malware targets Linux, UNIX servers

A new malware that runs on UNIX-like servers even with restricted privileges has already infected machines in Australia and is actively hunting for more targets, a new research paper has shown.

Three researchers from Russian web provider Yandex - Andrej Kovalev, Konstantin Ostrashkevich and Evgeny Sidorov - said in the technical analysis of the malware, published on security and anti-virus specialist publication Virus Bulletin, that Mayhem functions like a traditional Windows bot.

Mayhem was discovered in April this year and does not require a privilege escalation vulnerability - it does not have to run as the root super user - to work on Linux-based systems, or on FreeBSD servers.

Source:
http://www.itnews.com.au/News/390053,new-mayhem-malware-targets-linux-unix-servers.aspx?utm_source=feed&utm_medium=rss&utm_campaign=editors_picks

Submitted by: **Arnfried Walbrecht**

This month I thought I'd tie-in with our series of Git articles by referring to something I generally store in a Git repository – websites and CMS templates for work. This way I can iterate without worrying about losing previous revisions, and can host the repository on my NAS at home, or my shared web host, making it easier to grab a copy if I'm on a new computer, or in case something goes horribly wrong with my hardware. Instead of running through the same steps on how to create a Git repository from last month, I'm going to focus on setting up SASS (a CSS pre-processor) using Node.js and Grunt. This can be done locally, or on a web host (provided you can install Node.js and npm – check the documentation).

DEFINITIONS

SASS – stands for "Syntactically Awesome Style Sheets", and is a format you can use to generate typical CSS. However, it provides extra features which makes it easier for the programmer, while still compiling into normal CSS. Such as: functions, variables (i.e. defining colors in the header, so changing the color requires only one change), nesting of selectors, and mixins (mini functions to convert one line of code into many). Also referred to as SCSS.

Node.js – "Is a platform built on Chrome's JavaScript runtime for easily building fast, scalable network applications" (their Website). It's essentially a platform for running JavaScript programs without requiring something like Apache (or, at least, that's how I'd describe it).

NPM – Node Packaged Modules (essentially a repository for Node applications).

Grunt - "The JavaScript Task Runner" - this is essentially a file and folder watcher, that executes user-defined commands when something changes. I typically use it only for compiling SASS into CSS files, but it could do much more.

PREREQUISITES (TESTED IN 14.04)

Install Node.js and NPM:

```
sudo apt-get install nodejs
npm ruby-sass ruby-compass
```

Install grunt:

```
sudo npm install -g grunt
```

Connect Node to the expected location (this is caused by the Ubuntu package):

```
sudo ln -s /usr/bin/nodejs
/usr/bin/node
```

Note: You may need to run the npm commands as root. Doing so should not cause any issues.

Second note: Grunt is not the only way to compile SASS files, so if you're using a framework or starter theme that doesn't use Grunt, you'll either need to configure it from scratch, or else follow that framework's specific documentation.

GETTING STARTED (FROM SCRATCH)

Make sure you're currently working from the root directory of your project or template (i.e. the folder where the css and sass folders are).

Create a file called package.json that contains the following (or run npm init):
http://pastebin.com/qPMrRMNW

Naturally, you'll want to adjust some of the fields for your project. If you want to add dependencies while installing modules, run the command:

```
npm install <module> --save-dev
```

Replace <module> with the actual name of the module (i.e. grunt).

Install the modules required in packages.json by running the following in the project folder:

```
npm install
```

For SASS, you'll also need to install grunt-contrib-sass and grunt-contrib-watch.

Create your Gruntfile (either Gruntfile.js or Gruntfile.coffee). The file I typically use, is: http://pastebin.com/dWZUsvJn

The format of the folders then falls into: app/assets/css, app/assets/src/scss, app/assets/src/js. Any other folders can be fairly easily read from the Gruntfile (or added).

The way the Gruntfile works is that you define objects for various sections of your project – for example the project object contains the paths to each section of your app folder (and they can be referenced with project.app, for example). The banner is typically used to create those commented headers in open source projects – where the author, project and license is listed. These are pulled from the package.json file (stored in the object pkg). If the information doesn't exist, it will simply leave it blank. As such, the banner in this file contains all the information a full project would typically have, as it will work fine

with less information – that way I have one file for any kind of project.

The sass section is relatively standard – this defines where the css files will go, that it is using the expanded style when in development (this means the file is not minified and all line breaks and spacing are intact). The dist (for distribution) uses the compressed style. Compass is enabled in both (Compass is a CSS authoring framework). You can leave it disabled, and avoid installing ruby-compass at the beginning of this file. However, you may eventually run into errors.

You can also define an uglify object using grunt-contrib-uglify to minify JavaScript files. More information can be found here: https://github.com/gruntjs/grunt-contrib-uglify

As the Gruntfile by default runs sass:dev (with expanded styles), you can run sass:dist yourself with the following command:

```
grunt sass:dist
```

This will then create a compressed css file, for use in a

production environment.

GETTING STARTED (USING A FRAMEWORK OR STARTER THEME THAT USES GRUNT)

Go to the directory of the Framework/Theme.

Run: npm install

Run: grunt

Start editing SASS files and let grunt do its work!

THE STYLES

As an example of what SASS can do, I've created this file: http://pastebin.com/Wr0LmTXF

This is then compiled into normal css, as seen in this file: http://pastebin.com/ruUSfVKs

What you might notice is that it even converted the specified hex values for the colors into the CSS keywords "black" and "white", and that the mixin created all the lines of code for all browsers using the correct values and format. This means, if you need to use various

gradients, it cuts down on how much you need to repeat yourself.

It illustrates what a mixin is (linear-gradient), how to use it (@include), and defining variables in SASS. As well as a normal body selector. You can also nest items inside a body, or define different versions of it, such as :hover.

```
body {
    a {
        &:hover {
        }
    }
}
```

Hopefully this article has piqued your interest in SASS. If there is any interest in another article on SASS specifically, feel free to email me at lswest34+fcm@gmail.com. If you have any questions, corrections, or requests, you are also more than welcome to send me a quick email.

Lucas has learned all he knows from repeatedly breaking his system, then having no other option but to discover how to fix it. You can email Lucas at: lswest34@gmail.com.

W e've been working on a Cross Stitch pattern generator. Last month we did the UI portion, and now it's time to do the code that does the most of the work. Next month we will start working on the PDF file output portion.

We'll work on the menu items first. The code is shown below.

The global ReadyToProcess variable is used to make sure that if the user presses the Process button, the system doesn't try to process things without anything to process. We use the tkFileDialog askopenfilename built-in dialog routine to get the filename of the original image. We then get the number of colors in the original image as well as the width and height. We save those values and display them in the GUI. We then open the image and create a thumbnail image to display in the left image in the bottom frame. See the text box to the right.

Next we do the ShowHideGrid function. This simply exchanges two images in the right image label based on the global variable ShowGrid. If False, we change the text on the show/hide button, then set the ShowGrid variable to true and set the image to the one with the grid. Otherwise we change the text on the show/hide button to "Show Grid", set the ShowGrid variable to False and put up the ungridded image. Code is on the next page, top left.

The StitchSizeSelect function is fired whenever the stitch size combobox is changed. We get the value from the combo box and assign it to a local variable.

```
def StitchSizeSelect(self,p):

selection = ComboStitch.get()
```

```
        OriginalFilename.set(fileName)
        OriginalColorCount.set(self.GetColorCount(fileName))
        OriginalSize.set(self.GetHW(fileName))
        masterimage=Image.open(fileName)
        masterimage.thumbnail((400,400))
        self.img = ImageTk.PhotoImage(masterimage)
        self.lblImageL['image'] = self.img
        ReadyToProcess = True
```

The FileSave menu option will simply call the CreatePDF routine, once it's finished.

```
    def FileSave(self):
        self.CreatePDF()
```

We'll stub out the ShowHelp and ShowAbout routines with a dialog box saying that those options are not yet available.

```
    def ShowHelp(self):
        tkMessageBox.showinfo(title="Help",message='Sorry,
but help is not yet available.')

    def ShowAbout(self):
        tkMessageBox.showinfo(title="About",message='Sorry,
but the About function is not yet available.')
```

We've written the OpenDB routine a dozen times before, so you should know what it does.

```
    def OpenDB(self):
        global connection
        global cursor
        #---------------------------------
        connection = apsw.Connection("floss.db3")
        cursor = connection.cursor()
```

```
        def GetFileName(self):
            global ReadyToProcess
            #---------------------------------
            fileName = tkFileDialog.askopenfilename(parent=root,filetypes=self.picFormats ,title="Select File to open...")
```

```python
def ShowHideGrid(self):
    global ShowGrid
    #---------------------------------
    if ShowGrid == False:
        self.btnShowGrid['text'] = 'Hide Grid'
        ShowGrid = True
        self.im2=Image.open(self.GridImage)
        self.im2.thumbnail((400,400))
        self.img3 = ImageTk.PhotoImage(self.im2)
        self.lblImageR['image'] = self.img3
    else:
        self.btnShowGrid['text'] = 'Show Grid'
        ShowGrid = False
        self.im2=Image.open(self.ProcessedImage)
        self.im2.thumbnail((400,400))
        self.img3 = ImageTk.PhotoImage(self.im2)
        self.lblImageR['image'] = self.img3
```

```python
def AidaSizeSelect(self,p):
    selection = ComboSize.get()
    if selection != "30":
        pos = selection.find("x")
        width = int(selection[:pos])
        height=int(selection[pos+1:])
    else:
        width = 30
        height = 30
    FabricWidth.set(width)
    FabricHeight.set(height)
```

The AidaSizeSelect function (top right) is very similar to the StitchSizeSelect function. We set the FabricWidth and FabricHeight globals based on the selection on the combo box. We also default to 30x30 if they select 30.

We have a variable called ReadyToProcess (below) just in case the user tries to run the process function before the image is loaded.

We pixelate the original file to a 5x5 pixel matrix This allows us to group that 5x5 matrix to a single color. We then reduce the colors, get the width and height of the processed image and set the size so the user can see how big the resulting image will be.

```python
# Place image

self.im2=Image.open(Reduced)

self.im2.thumbnail((400,400))

self.img3 =
ImageTk.PhotoImage(self.im2)

self.lblImageR['image'] =
self.img3

self.ProcessedImage =
'im1.png'
```

The above set of code places the processed image into the image that will hold the processed image. The next set of code will create a grid so that the user will have the grid to do the cross stitching.

```python
self.MakeLines(Reduced,5)

self.MakeLines2('output.png',
50)

self.im2 =
Image.open('output2.png')

self.im2.thumbnail((400,400))

self.img3 =
ImageTk.PhotoImage(self.im2)

self.lblImageR['image'] =
self.img3

self.FillScrolledList('output
.png')
```

```python
def Process(self):
    global ReadyToProcess
    #---------------------------------
    if ReadyToProcess == False:
        tkMessageBox.showinfo(title="ERROR...",message='You must load an original imaage first.')
    else:
        newimage = self.Pixelate(OriginalFilename.get(),5)
        Reduced = self.ReduceColors(newimage)
        W,H = self.GetHW2(Reduced)
        siz = "{0}x{1}".format(W/5,H/5)
        ProcessedSize.set(siz)
```

```
self.GridImage =
'output2.png'
```

We stub the CreatePDF function until we finish the PDF function next month.

```
    def CreatePDF(self):

tkMessageBox.showinfo(title="
Create PDF",message='Sorry,
but the Create PDF function
is not yet available.')
```

The OriginalInfo() routine gets and sets variables based on the original image format, size and mode.

```
def OriginalInfo(self,file):
    im = Image.open(file)
    imFormat = im.format
    imSize = im.size
    imMode = im.mode

    self.size = imSize
    self.imformat = imFormat
    self.immode = imMode
```

The GetColorCount function uses the .getcolors method to get the number of colors in the image file. We have to use 1600000 as the maxcolors parameter because if the image contains more than 256 colors (or whatever is in the parameter, the method returns 'None'. This function is similar to the GetColors function except the GetColors works with an already

```
def Pixelate(self,im,pixelSize):
    image = Image.open(im)
    self.GetColors(image)
    image = image.resize(((image.size[0]/pixelSize, image.size[1]/pixelSize), Image.NEAREST)
    image = image.resize(((image.size[0]*pixelSize, image.size[1]*pixelSize), Image.NEAREST)
    self.GetColors(image)
    #image.show()
    image.save('newimage.png')
    return 'newimage.png'
```

opened image file. If you use GetColorCount, you have to pass an unopened file.

```
def GetColorCount(self,file):
    im = Image.open(file)
    numColors =
im.getcolors(1600000)
    self.colors =
len(numColors)
    return self.colors
```

The next two functions return the height and width of the image file in pixels. The difference between the two is that GetHW returns a string like 1024x768 and GetHW2 returns two integers.

```
def GetHW(self,file):
```

```
    im = Image.open(file)
    tmp =
"{0}x{1}".format(im.size[0],i
m.size[1])
    return tmp

def GetHW2(self,file):
    im = Image.open(file)
    return
im.size[0],im.size[1]
```

GetColors will get the number of colors in the passed image file. We use 1.6 million colors as the parameter, because the image.getcolors() routine defaults to 0 over color count over 256.

```
def GetColors(self,image):
    numColors =
image.getcolors(1600000)
    colors = len(numColors)
```

The Pixelate function (above) takes two parameters, image filename (im) and the size of pixels you want. The work is done by the image.resize method. I found this routine on the web in a number of places. In this instance we will be passing a pixel size of 5, which works well for Cross Stitch projects. We also tell the method to take the color of the nearest neighbor. This returns a new image, which we save as a file and return the filename.

The ReduceColors routine (below) basically uses the Image.ADAPTIVE pallet so we can

```
def ReduceColors(self,ImageName):
    #Reduce colors
    numcolors=MaxColors.get()
    image = Image.open(ImageName)
    output = image.convert('P', palette=Image.ADAPTIVE, colors=numcolors)
    x = output.convert("RGB")
    self.GetColors(x)
    numcolors = x.getcolors()
    ProcessedColors.set(len(numcolors))
    x.save('im1.png')
    return 'im1.png'
```

get a much smaller number of colors.

There are two MakeLines (top right) routines. They create the grid we spoke of earlier.

Rgb2Hex() returns a hex value of the RGB value that is passed in. We will use this to try to compare the colors in the database with the colors in the image.

```python
def Rgb2Hex(self,rgb):
    return '#%02x%02x%02x' %
rgb
```

The ScrollList (below) on the right side holds the colors that will

be used to get the proper floss colors. We simply create labels to hold the colors (visual) and text.

This (next page) is the routine that we use to try to find the closest match between the color in the image and the color in the database. There are many different algorithms on the web that you can look at and try to understand the logic behind it. It gets rather complicated.

Ok. That's all for this month. Next time, we will start creating the PDF output file so the cross stitcher has something to work

```python
def MakeLines(self,im,pixelSize):
    global backgroundColor1
    #--------------------------------
    image = Image.open(im)
    pixel = image.load()
    for i in range(0,image.size[0],pixelSize):
        for j in range(0,image.size[1],pixelSize):
            for r in range(pixelSize):
                pixel[i+r,j] = backgroundColor1
                pixel[i,j+r] = backgroundColor1
    image.save('output.png')

def MakeLines2(self,im,pixelSize):
    global backgroundColor2
    #--------------------------------
    image = Image.open(im)
    pixel = image.load()
    for i in range(0,image.size[0],pixelSize):
        for j in range(0,image.size[1],pixelSize):
            for r in range(pixelSize):
                try:
                    pixel[i+r,j] = backgroundColor2
                    pixel[i,j+r] = backgroundColor2
                except:
                    pass
    image.save('output2.png')
```

```python
    def FillScrolledList(self,filename):
        im = Image.open(filename)
        numColors = im.getcolors()
        colors = len(numColors)
        cntr = 1
        for c in numColors:
            hexcolor = self.Rgb2Hex(c[1])
            lblColor=Label(self.sfFrame,text="            ",bg=hexcolor,relief=GROOVE)
            lblColor.grid(row = cntr, column = 0, sticky = 'nsew',padx=10,pady=5)
            pkID = self.GetBestDistance(c[1][0],c[1][1],c[1][2])
            sql = "SELECT * FROM DMC WHERE pkID = {0}".format(pkID)
            rset = cursor.execute(sql)
            for r in rset:
                hexcolor2 = r[6]
                dmcnum = r[1]
                colorname = r[2]
            lblColor2=Label(self.sfFrame,text="            ",bg="#" + hexcolor2,relief=GROOVE)
            lblColor2.grid(row = cntr,column = 1,sticky = 'w',padx=5,pady=5)
            lblColor3=Label(self.sfFrame,text = str(dmcnum) + "-" + colorname,justify=LEFT)
            DmcColor.set(dmcnum)
            lblColor3.grid(row = cntr, column = 2,sticky = "w",padx=1,pady=5)
            cntr += 1
```

with.

As always, the code is available on PasteBin at http://pastebin.com/DmQ1GeUx. We will continue in the next month or so. I'm facing some surgery soon so I'm not sure how soon I will be able to sit for any long periods of time. Until then, enjoy.

Greg Walters is owner of RainyDay Solutions, LLC, a consulting company in Aurora, Colorado, and has been programming since 1972. He enjoys cooking, hiking, music, and spending time with his family. His website is www.thedesignatedgeek.net.

```python
def GetBestDistance(self,r1,g1,b1):
    # dist = math.sqrt(((r1-r2)**2) + ((g1-g2)**2) + ((b1-b2)**2))
    sql = "SELECT * FROM DMC"
    rset = cursor.execute(sql)
    BestDist = 10000.0
    for r in rset:
        pkID = r[0]
        r2 = r[3]
        g2 = r[4]
        b2 = r[5]
        dist = math.sqrt(((r1-r2)**2) + ((g1-g2)**2) + ((b1-b2)**2))
        if dist < BestDist:
            BestDist = dist
            BestpkID = pkID
    return BestpkID
```

When I was in school, I was very focused on writing, and I didn't see the value in many other subjects, especially mathematics. As I grew older, I began to appreciate and enjoy mathematics. Through the years, I have studied mathematics on my own. While I studied, I created many documents containing my mathematics notes. I decided it was time to put them all into one document with chapters, a table of contents, and an index. But how does one combine a bunch of documents into one?

I guess I could have just copied the text from one document into another big document, but this could get messy, especially since I wasn't sure about the best order for the documents. I wanted a method where I could work on the documents individually then combine them into one, and I wanted the ability to rearrange the order of the documents. In LibreOffice, you can do this with 'master documents'. You can think of a master document as a container that joins together separate Writer documents. Just

the kind of thing I needed. So, I decided to use a master document.

PREPARING THE DOCUMENTS

Since the documents were scattered through several years, they were created with different versions of OpenOffice and LibreOffice. They had no style or template in common. I needed to create a template for consistency throughout all the documents and the master document.

To create my template, I started with a new blank document. While creating my template, there were a few things I took into consideration. I knew I wanted to create a table of contents, so I would need to edit the "Outline Numbering." I also wanted each chapter to start on a new page. Finally, I would need a unique title

page style.

To edit the "Outline Numbering", I went to Tools > Outline Numbering. I edited the top level to put "Chapter" and the chapter number before the top level headings, just as I discussed in Part 38 of this series (Full Circle issue 85). I decided to leave the other levels bare, although I would probably include some of them in the table of contents.

To make each chapter start on a new page, I edited the outline top level style, "Heading 1." I opened the "Styles and Formatting" tab in the sidebar, right-clicked on "Heading 1" in the paragraph styles, and selected "Modify." On the "Text Flow" tab under "Breaks", I checked "Insert". For the type I selected "Page," and for the position, I selected "Before." OK saved the changes.
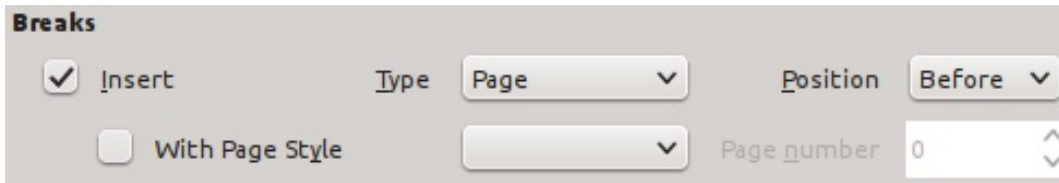

NOTE: You can also open the "Styles and Formatting" window from its button on the Formatting toolbar, the menus (Format > Styles and Formatting), or by pressing F11.

I created a page style named "Title Page" based off the "First Page" style. The only real change I made was to set the top to halfway down the page to center the title on the page vertically. I modified the "Title" paragraph style to a font and size of my liking. I also created the paragraph style "Byline" based on the "Subtitle" style. Not many changes here, just selected a font to go with the title font.

To save the template, File > Save as Template opened the Template Manager. I clicked on Save, and the program prompted me for a name for the template. I named it "Math Reference." I clicked OK, and I had a new template.

When you are creating a template, you may have other

changes you want to make. These were the ones that I felt were a good starting point for me. You can go as far as you want with a template. In the end it is up to you. If you decide, while working on sub-documents, that you need to make more changes to styles, make the changes to the template rather than the individual document.
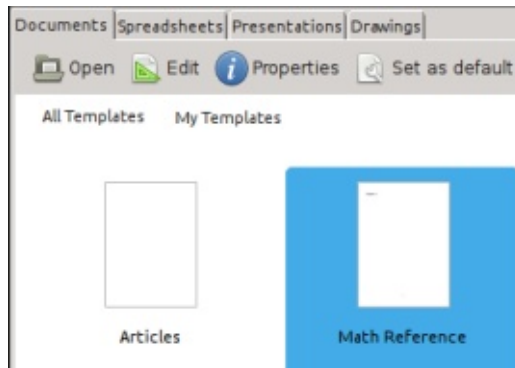
To make changes to my mathematics template, File > New > Templates. I selected the template and clicked the Edit button. This opened the template rather than a new document using the template. I made my changes and then saved it as though it was a normal document (Click the save button on the main toolbar, File > Save, or CTRL-S). When I opened one of the documents that uses the template, LibreOffice notified me that the template had changed and asked me to update the document with the new version of the template. I clicked "Update Styles", and the styles updated in the document.

Now, I needed to apply my new template to my existing documents. I opened each of the documents and Edit > Select All. I then created a new document using my template, File > New > Templates. I selected my template and clicked Open. Edit > Paste and the text and objects from the original file were pasted into the new file. I closed the old document because I wanted to save it with the same filename. I saved the new document and I had the old file contents using the new template. I repeated this procedure for all the existing documents.

CREATING THE MASTER DOCUMENT

To create my master document, I opened a new document using my "Math Reference" template. Selecting the "Title" style from the paragraph styles, I typed in my title. I then created the subtitle and byline. Once I had something in the document, I saved it as a master document, File > Send > Create Master Document. I gave the document a name and clicked Save. I then went to the page styles in the Styles and Formatting dialog and changed the page style to "Title Page."

INSERTING DOCUMENTS

Once I had a master document, I could start adding files to it. When I saved the master document, it opened the Navigator in a floating window automatically. I could have worked in this floating window, but I closed it and used the Navigator panel in the sidebar. They are both the same, so it's your choice which one you use. The Navigator window is opened by pressing F5 on the keyboard, or from the menus, View > Navigator.

The Navigator in master documents is different from other documents. By default it is in document list mode. The first button in the toolbar is a toggle button that allowed me to toggle between the document list and a standard Navigator panel. While working in the master document, I did not see any reason for switching to normal mode, but it is there should you have a reason for navigating to an object or

subheading in one of the sub-documents.

I noticed that one document was already in the list of files. This was the Text of the master document itself, at this point, my title page. I would add other Text blocks as needed in the master document using the insert button. To add my documents to the master document, I clicked and held on the Insert button <image>. Dragging down to the File selection, I released the mouse button. An insert file dialog appeared for me to find and select my file. I clicked open and the file was added to my master document. When adding a file to the master document, Writer always adds it above the currently selected file. I used the "Move Up" and "Move Down" buttons <image> to get my documents in the order I wanted them. All references and chapter numbers adjusted automatically as I moved the documents up and down in the list.

EDITING FROM THE MASTER DOCUMENT

All documents linked to the
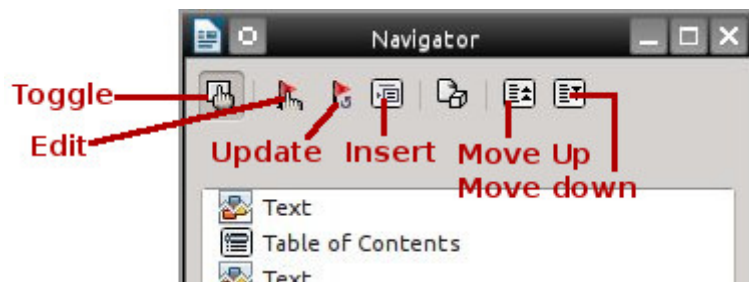
master document are read only in the master document. To edit a document from the master document, I selected the document in the Navigator list and clicked the edit button. A new window containing the document I wanted to edit opened. I made my changes to the document, saved, and closed it. In the master document, I clicked and held the Update button. Dragging I released on Links. A message window appeared asking me whether I wanted to update all the links in the document. I clicked Yes, and the document I just edited updated in the master document. I was also able to open the document through normal methods and edit it that way. The changes still appeared when I updated the master document's links.

I inserted a table of contents by inserting a page break after my title page and using the same methods discussed in Part 38 of this series (Full Circle issue 85). I did the same for an index, but I had to insert a Text block at the end of the master document through the insert button. This text block was needed to create an editable block at the end of the master document.

Master documents are a great way to create large documents. In my case, I wanted the flexibility to work on the individual documents and the ability to rearrange the order of the documents. A master document works best with a template shared by all the documents. Just like normal Writer documents, you can add a table of contents and indexes. While a master document is not always the best choice, it is a good choice if you need the ability to move the different parts around or have different people authoring the different parts. This was just one example of using a master document. If you plan things well, you can start from the very beginning before you even write the first document.

Elmer Perry's history of working, and programming, computers involves an Apple][E, adding some Amiga, a generous helping of DOS and Windows, a dash of Unix, and blend well with Linux and Ubuntu.
He blogs at
http://eeperry.wordpress.com

In previous parts I explained how GRUB (version 2) works and how you can tune its behaviour. The workings of the GRUB menu are determined by the configuration file /boot/grub.grub.cfg. This file is created when you type 'sudo update-grub' in a terminal. One of the files that is read by update-grub is /etc/default/grub, which we discussed last time. Further, we learned that the executable scripts under /etc/grub.d/ determine other aspects of /boot/grub/grub.cfg, and these will be our focus now.

The names of the scripts in /etc/grub.d/ all start with a number, and (only) those which are actually set as 'executable' are read by update-grub in the order of their number. What we may want to do here is to disable the executable flag of the scripts according to our choice (with the chmod command) – especially the script '40_custom'. You can use 40_custom to add custom menu entries.

This is how a default 40_custom script looks like:

```
#!/bin/sh
exec tail -n +3 $0
# This file provides an easy way to add custom menu entries. Simply type the
# menu entries you want to add after this comment. Be careful not to change
# the 'exec tail' line above.
```

So, except for the exec line, this script is empty (# is for comments). Simply add your custom entries below these lines. 'Simply' you say? Well, we'll do that in a minute. The custom entries you provide here are directly copied into grub.cfg when running 'sudo update-grub'. And custom entries is what we will need – we still need a better look at how menu entries in grub.cfg are composed.

Let's consider two types of menu entries which you will encounter in a typical grub.cfg on a pc that contains a Linux distribution and Windows (I still have XP). (For Mac, see links at the end.) Copy the entries to 40_custom to experiment with these: these custom menu entries will just appear as extra selectable lines in your GRUB menu, so there's no harm in doing this. The extracts below are from an unmanipulated grub.cfg that was created by the grub-pc package on a pc. During installation, its contents were automagically tailored for my specific hardware (e.g. taking into account the type of partitioning of the hard disk, e.g. msdos-partitioning if Windows was there first). Also, its specifics will partly depend on the version of the grub-pc package. Consequently, I advise to copy the corresponding parts from your grub.cfg, as this is already tested and proved for your machine.

The menu entry for a Linux distribution (e.g. Bodhi Linux) is shown on the next page (top).

Let's have a look at some essentials here:

• **set root=(hdx, y)** sets the device and root directory on which grub has to look for the primary files to load when booting this OS (i.e. after this menu entry has been selected in the GRUB menu). Beware of the odd numbering of devices and partitions! Number x refers to the first, second, third, … hard drive and starts at 0. Number y refers to the partition on hdx but starts counting at 1. So the first partition on the first hard drive (/dev/sda1) is (hd0,1), the second partition (/dev/sda2) is (hd0,2) and so on. In Windows partitioning layout, an msdos prefix is used as in my case.

• In the case of a Linux distribution, the files that GRUB must look for on this partition are the kernel and the initial ramdisk image (initrd.img). The initial ramdisk image is loaded by the initrd command (the initial ramdisk is a scheme for loading a temporary root file system into memory in the boot process of the Linux kernel). linux loads the Linux kernel. Two alternatives exist here. The first, as presented, is to provide the direct path to the files themselves, which means their names will need to be changed when a more recent kernel has been installed. This adjustment of grub.cfg is

```
menuentry 'Bodhi Linux, with Linux 3.7.0-7-generic' --class bodhi --class gnu-linux --class gnu --class os {
        recordfail
        gfxmode $linux_gfx_mode
        insmod gzio
        insmod part_msdos
        insmod ext2
        set root='(hd0,msdos7)'
        search --no-floppy --fs-uuid --set=root 57509291-b751-4131-aee7-0be040e93ce6
        linux /boot/vmlinuz-3.7.0-7-generic root=UUID=57509291-b751-4131-aee7-0be040e93ce6 ro   quiet
        initrd     /boot/initrd.img-3.7.0-7-generic
}

Menu entry for Windows (XP in my case):
menuentry "Windows xp" {
        insmod part_msdos
        insmod ntfs
        set root='(hd0,msdos1)'
        search --no-floppy --fs-uuid --set C600FB5E00FB53C1
        drivemap -s (hd0) ${root}
        chainloader +1
```

automatically taken care of when a new kernel is installed. The alternative is that you provide the symlinks to your latest kernel and ramdisk (which are often present in /boot/). The command 'sudo update-grub' automatically takes care of these things, with separate entries that point to older installed kernels. Check it out!

• In all cases, it is also possible to 'chainload' another boot loader whenever one is installed on the chosen partition. Indeed, boot loaders can also be installed at the beginning of a partition. Chainloading is the loading of a

new boot loader from the boot loader you are in. Chainloading a bootloader that resides on the (hdx,y) that is set, is accomplished with the GRUB command 'chainloader +1'. As Windows needs its own boot loader to boot, chainloading is obligatory to boot Windows. A Linux installation may or may not have a (second) GRUB installation on its own partition (it behaves exactly the same: it refers to /boot/grub/grub.cfg to display its menu). In case it does have a separate GRUB, in the MBR GRUB's configuration file you can choose whether you want to chainload that distro's GRUB or directly boot

its kernel.

Except for the insmod statements, you may want to experiment by omitting lines in the menuentries that you copied and see if things still work (they may). You can even do this for one session only, without altering grub.cfg. Just press 'e' after highlighting the menuentry in the GRUB menu that you want to inspect or change. GRUB will show you the commands from the menu entry; you can safely edit these and execute with Ctrl-X (nothing will be saved to grub.cfg). Also try omitting the 'quiet' parameter in

the linux statement to see what happens. Those wishing to learn more are referred to the links at the end. But with the above ingredients you will get along.

You may also want to change some text colors in the GRUB menu, especially if you will use a background 'splash' image (see previous time). Color settings are present in the script 05_debian_theme or possibly in an extra script like 06_mint_theme (under/etc/grub.d/). You will use such things as set color_normal, color_highlight, menu_color_normal and
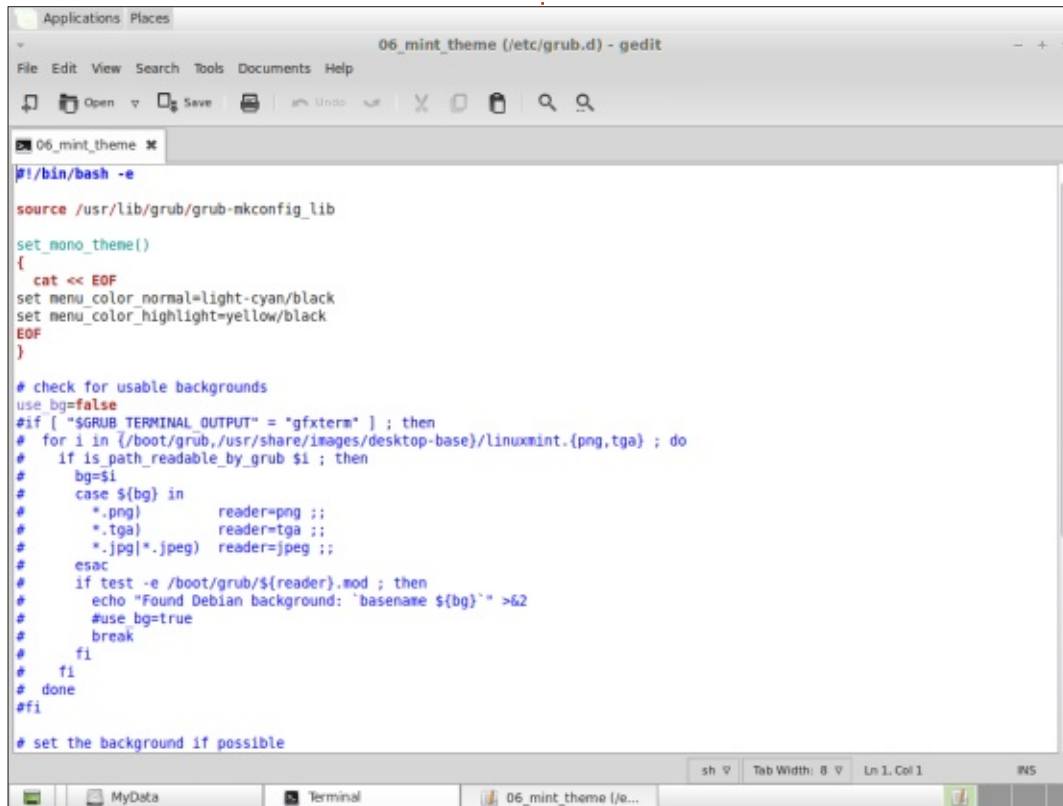
menu_color_highlight; inspect the online resources for more information. In short, the defaults are black/light-gray for (menu_)color_highlight and white/black for (menu_)color_normal. 'black/light-gray' means black text on a light-gray background (for the selected menu entry); 'white/black' is white text on a transparent background (black as background color is not black!). That will be important when you use a background image. The following colors are available: black, blue, brown, cyan, dark-gray, green, light-cyan, light-blue, light-green, light-gray, light-magenta, light-red, magenta, red, white and yellow. Have fun with these things. Beforehand, you can test the effect in real-time by running the appropriate GRUB commands in the GRUB terminal (e.g. set color_normal=blue/white), but you will need to insert these settings in 05_debian_theme to get it set automatically through grub.cfg. You can also experiment with grub.cfg and insert the 'set … color=' statements directly in there, but beware that grub.cfg is overwritten after update-grub is run.

So, we made alterations to one or two files in /etc/grub.d/ and to /etc/default/grub. Now we update /boot/grub/grub.cfg correspondingly (backup your existing grub.cfg). So run:

```
sudo update-grub
```

Don't forget to have a look at grub.cfg to check if everything looks OK. Your custom entries should be there. Reboot and try them out. GRUB is yours now; enjoy!

Next time will be the final part of this mini-series. I will tell some basics about partitioning – and then we build our multibooting machine…

INTERESTING LINKS

• Ubuntu documentation on GRUB2 - plus related links at the page bottom, on GRUB terminal ('troubleshooting'), menu displays and custom menu entries: https://help.ubuntu.com/community/Grub2

• Dedoimedo GRUB2 full tutorial, with very useful hints on GRUB configuration: http://www.dedoimedo.com/computers/grub-2.html

• A short list of kernel parameters, with links to full lists: https://wiki.archlinux.org/index.php/kernel_parameters#Parameter_list

• Booting Mac OS from GRUB: http://leveluplinux.blogspot.be/2012/09/dual-booting-linux-and-os-x-using-grub2.html or http://www.maketecheasier.com/create-a-mac-entry-in-grub2 (not tested)

• The official GRUB manual: http://www.gnu.org/software/grub/manual/

Floris Vanderhaeghe became a Linux fan through Ubuntu 10.10. You can email him at tux7546@gmail.com.

Before I dive headlong into the last couple of items in Inkscape's tool palette, I have a quick update on the previous article. Last time I drew attention to the misleading placement of the tablet pressure button on the Tweak tool's control bar; but as well as mentioning it in the article, I did my duty as a user and filed a bug report. I'm pleased to say that the bug was quickly triaged, then fixed in just a couple of days, so the next major release of Inkscape should no longer suffer from this minor anomaly.

Now, onto those last two tools – the ones I use least in my work with Inkscape – starting with the Eraser. You can switch to this tool using either the tool palette icon, or by pressing Shift-E. There are barely any options available via the tool control bar, and the reason I rarely use this tool is that its features are all available via other tools anyway, albeit a little less directly. Let's take a look at that minimal tool control bar.

Once again the widgets are not very intuitively positioned – the two buttons on the right switch between two different eraser modes, but the Width slider on the left actually applies only if the second mode is selected. Time to file another bug report...

The first mode is used to delete entire objects from your drawing. With this button selected, drawing on the canvas produces a red line. On releasing the mouse button, any objects that the red line touches will be removed.
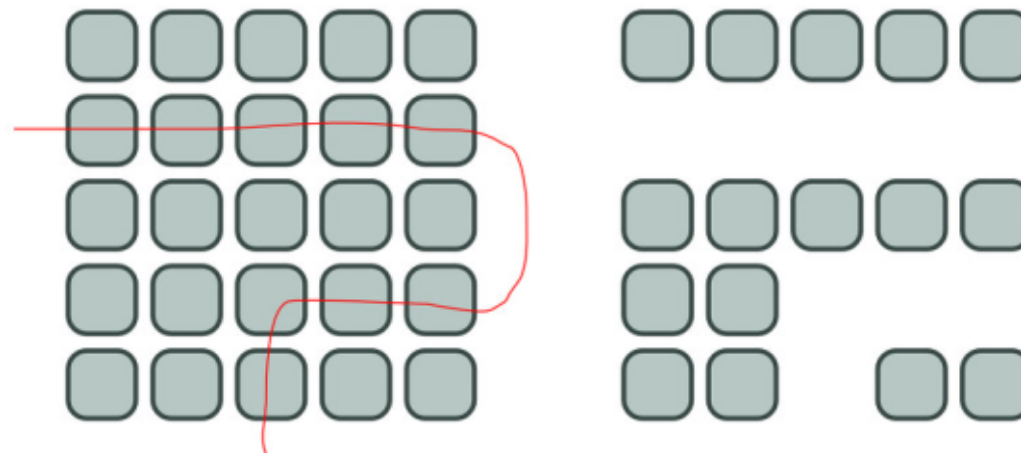
If this sounds slightly familiar, it's because similar functionality is available using the Select tool. By holding the Alt key whilst drawing on the canvas with the Select tool active, you can draw a thin red line, with any objects the line touches being selected when the mouse button is released. From there, it's just a quick tap of the Delete key to remove them. Although the Eraser tool is slightly more direct, in that there is no need to press the Delete key, performing this task using the Select tool instead offers you the opportunity to add more objects to your selection – or remove some from it – before finally committing your deletion.

Another reason to be wary of this mode is that it's easy to accidentally delete too many objects. You might think that, if you pre-select some objects before enabling the tool, you would be able to selectively remove only those that the line touches and which were also selected. In practice all of the selected items will be removed as soon as you use the tool – even if you don't touch any of them! You can imagine how disruptive this could be if you still have objects selected that are outside of the visible canvas area. Ideally you should de-select everything (Edit > Deselect) before using this mode in order to prevent such unexpected behaviour.

The second mode of the Eraser tool is perhaps more useful: it lets you remove parts of paths and objects by drawing over them. The Width slider determines the size of the swathe that will be cut through your path, and objects will be automatically converted to paths as required. Unlike the other mode, pre-selecting objects can be useful here as only those in the selection
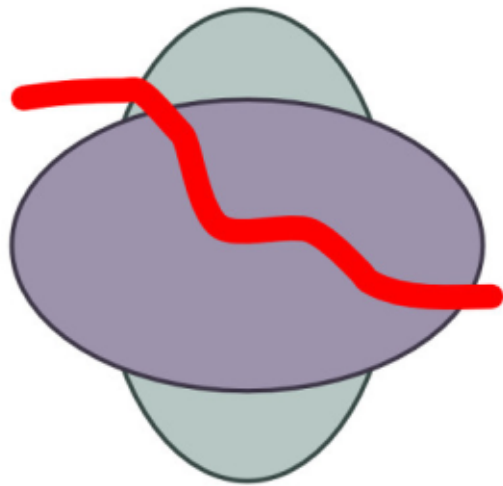
will be converted and carved. Alternatively, ensure that nothing is selected in order to cut through all the visible objects.

Although this mode is handy for quickly hacking an object into several paths, or just carving a little slice off the edge, the fixed width and profile of the cutting line limits its potential. A more flexible approach is to use the Calligraphy tool to create a path, then the Boolean Difference operation to remove it from your object. For cutting through several objects the Calligraphy tool can also be used whilst holding the Alt key to switch it to "subtract mode", which has a similar effect to the Eraser tool. The only caveat is that the Calligraphy tool doesn't always
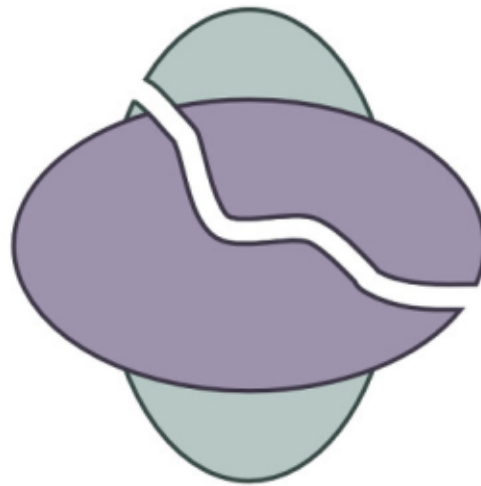
seem to convert primitive objects to paths when it should, so you may have to manually convert them first.

That's all there is to the Eraser tool: a "touch" mode that saves a single keypress compared with the Selector tool approach, and a "path" mode that is far less functional than the Calligraphy tool equivalent.

On to the last of the main Inkscape tools – one which I rarely use simply because I don't need it for the comics I create: the Connector tool.

The Connector tool is used to draw lines between objects, with

the lines being "connected" to those objects, at least in the sense that moving the object also moves the end of the connector line. This makes it a useful tool for creating flowcharts, diagrams and annotations, although its capabilities are somewhat limited when compared with dedicated applications such as Dia, Calligra Flow (formerly Kivio) and yEd.

To get started with the Connector tool, it helps to have a few objects already created. Activate the tool via the icon or using the "o" or Ctrl-F2 shortcuts. As you move the mouse over objects, you'll notice that they gain a square handle in the center; click and drag a line between the handle in one object and the handle in another to create a connector. You can also click, then move the mouse, then click again, but dragging the line tends to be more reliable. You've now connected two objects and if you press the Spacebar to switch back to the Selector tool, you should be able to move the objects around and see that the connector line remains attached.

The first thing to notice is that the connector appears to emanate

from the edge of your object, even though the connection was made to the center. Unlike other charting tools, Inkscape doesn't give you any option about where on an object the connector links to – you always create connections using the handle at the center, and they visibly connect to a point on the edge that you have no control over. The next thing to note is that, by default, your connector line will happily trample right across any objects that might lie in its path.
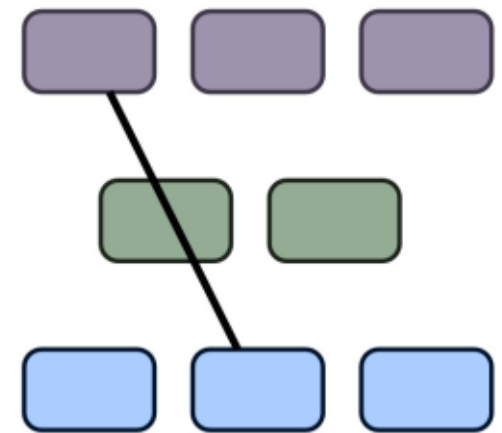

Inkscape offers a solution to this by allowing you to selectively choose objects that the connector should avoid. First you need to select the connector itself (you can select more than one), then select the objects to avoid. Now switch to the Connector tool and click the
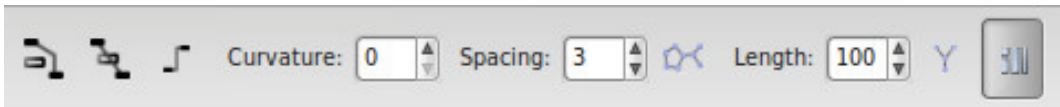
first button on the tool control bar.

You should find that the connector is now re-routed to avoid your selected objects, and you can use the Spacing control to adjust the amount of clearance that's added. If you change your mind, you can use the second button on the bar to specifically allow the connector to pass through the selected objects. With some avoidance in place, my sample connector now looks like this.
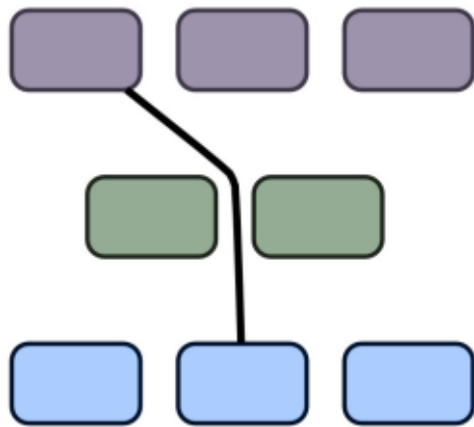
It's an improvement, but it doesn't look all that tidy. Perhaps restricting it to only perpendicular angles would help. With the

connector selected, the third button – which is actually a toggle – switches that connector between "orthogonal" (allowing only horizontal and vertical lines) and "polyline" (any angle is allowed). Enabling the toggle certainly makes things neater, but now there's another problem.
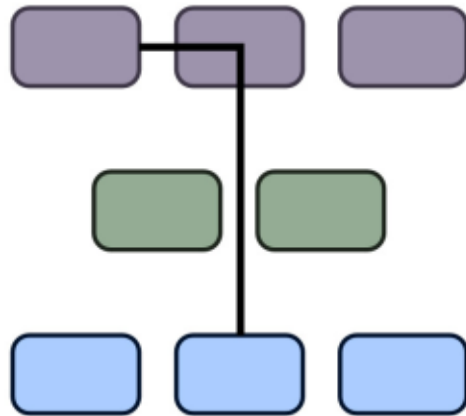
Because I'd only told the connector to avoid the green box, it's now passing right through one of the others. I could select the connector and this other box, then click the "avoid" button again, but I usually find that a better solution is actually to select all the connectors and all the nearby objects, so that you are pre-emptively telling

Inkscape to avoid them even if you move the objects around later. This makes it easier to draw your diagram without constantly having to adjust the avoidance rules. You may need to tweak the Spacing value to get a reasonable layout. With an orthogonal connector and Spacing set to zero, my chart is starting to look more acceptable.
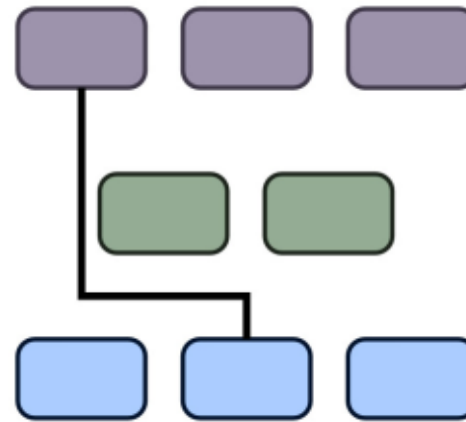
It's still not much of a chart though, so I'll add a few more random connections. Connectors can be styled in the same way as any other path in Inkscape, so I've changed line colors, widths and dashes as well as adding some End markers to provide arrow-heads. To get the arrow-heads to match the colors of the lines, I've used Extensions > Modify Path > Color Markers to Match Stroke, and added some shape to the lines

using the Curvature setting in the Connector's tool control bar. Now that's what I call a chart!
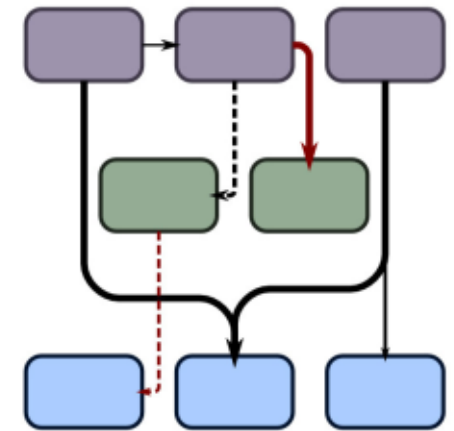
Once you've used the Connectors tool to forge links between objects in your chart or diagram you can also let Inkscape have a try at arranging it for you. The button on the control bar between Spacing and Length triggers this layout algorithm, with the Length value and the two remaining buttons affecting the final result. These latter controls aren't interactive, so you need to re-trigger the layout algorithm by pressing the button each time you change any of them. The layout algorithm only on only selected connectors and objects, which can be handy if you just want to tidy up part of a complex diagram, but can also result in a complete mess so

make sure you save first. By selecting all of my chart and pressing the trigger button you can see that it's not really as complex as it initially appears.
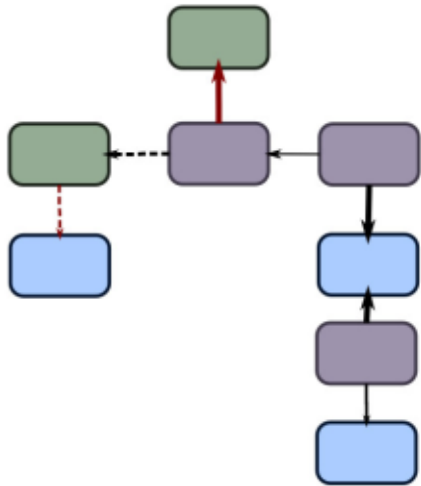
Don't be too impressed by that result: I had to cheat to get the chart looking that neat. Inkscape's algorithm spread the three arms of the diagram out quite widely, creating an unnecessarily sprawling image. Nevertheless it provided a clearer view of the relationships I had randomly drawn, making it easy to move the boxes around to give this more professional looking result.

The three controls that you have over Inkscape's own algorithm don't really amount to much. The Length parameter lets you adjust the preferred length of the connectors in the rearranged drawing. Unfortunately it's easy to go too low with this value, as it doesn't seem to follow a very linear scale, which results in boxes nestled far too close to each other. Just increase the value and hit the trigger button again if that happens to you.

The first of the control buttons tries to adjust the algorithm so that lines with end markers will always tend downwards. This doesn't work for start markers, so make sure you draw from source to destination if you want to use this mode. Enabling this and pressing the trigger results in this layout for my chart... after a little manual adjustment, once again.
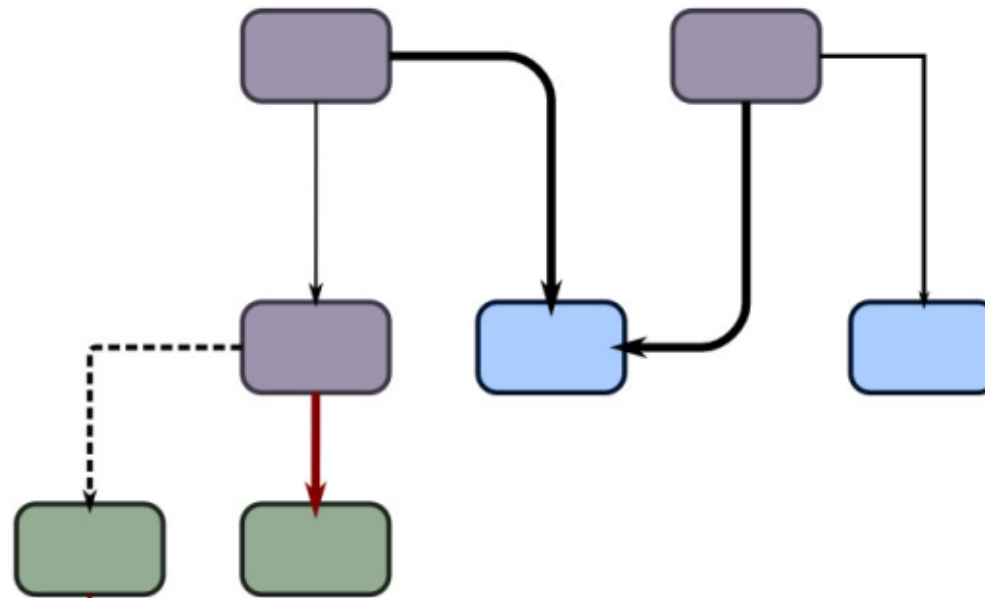
The final button tells the algorithm that shapes aren't allowed to overlap. This didn't prove to be a problem with my test chart, but if you prefer to not have the shapes in your diagram overlapping one another, it's probably best to enable this option anyway.

There are a couple of final things to note about the Connector tool: first, it's one of the more unstable parts of Inkscape. It can cause crashes, mess up the undo buffer, and produce results that are quite removed from your intentions! Make sure you save frequently if you use it. Secondly, you can't convert the connectors into "normal" paths using Edit > Object to Path, as you might expect. You can still manipulate the path's nodes manually by selecting a connector then switching to the Node tool, but any changes can be lost if you move the linked object or click the trigger button.

You may find the Connector tool useful for simple diagrams, but I usually find that such simple connections are just as easy to draw and modify manually, with no need for automatic layout algorithms. Once you reach a level of complexity where such facilities are useful, it's probably time to switch to a dedicated application instead.

Mark uses Inkscape to create three webcomics, 'The Greys', 'Monsters, Inked' and 'Elvie', which can all be found at http://www.peppertop.com/
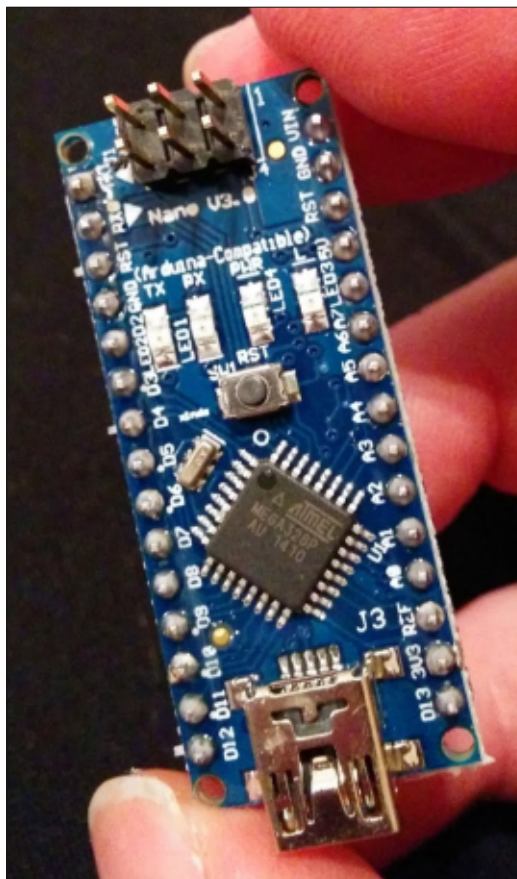
Several months ago, I took the plunge and decided to risk buying a couple of cheap Arduino Nanos from China. These are still Arduinos – but just not from the original designers. They were listed on eBay for £3 each so I went for three of them. Apparently keeping it under £15 means no import tax, which turned out to be true.

The Nanos arrived complete with a dinky USB cable. I hooked it up to my laptop and tried running the blinking LED sketch. Error. Now I was worried. I feared the worst in that either they were damaged or just plain didn't work. Then I realised that I hadn't changed the Arduino model in the IDE menu. Oops!

With the LED sketch uploaded and running I now knew the Nanos were working. Great!

The plan for one Nano was to realise my dream of having indicators on my bicycle. First I dug out the old sketch for the spaceship panel which covered holding down a button to flash an LED. I removed the unnecessary code and comments and now had a Nano with one flashing LED when a button was held down. I now duplicated the button, LED and wiring and voila! Bicycle indicators!

Note that the Nano slots directly into the breadboard unlike the Uno which sits separately.

I did think about modifying the code so that pressing once would turn the LED on then pressing again to turn it off, but that might mean me cycling along with indicators flashing not realising I've left them on and, ultimately, draining the battery (which I'll hook up later).

Obviously, this would now leave the breadboard and be physically soldered onto veroboard with wires running the length of the bicycle. But that's for me to figure out later.

My (butchered) code is at:
http://pastebin.com/urJPfr9C
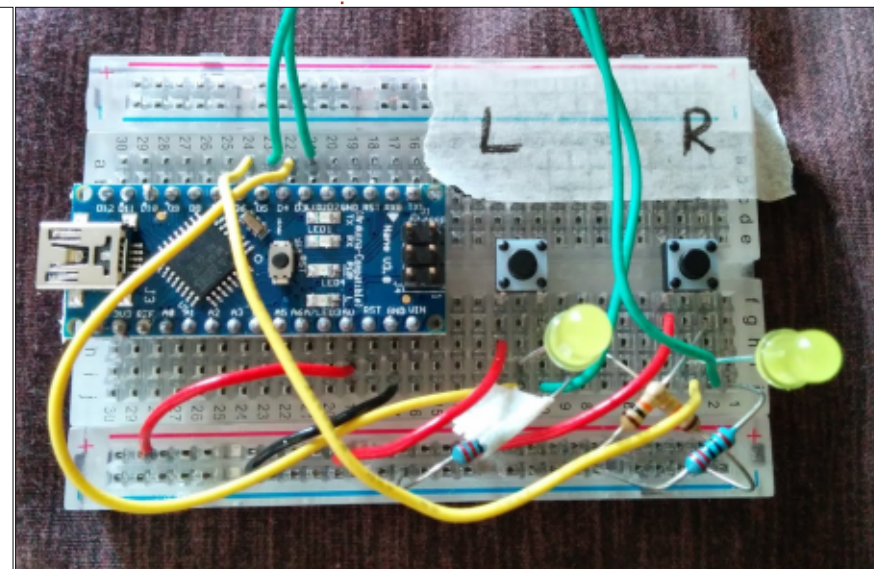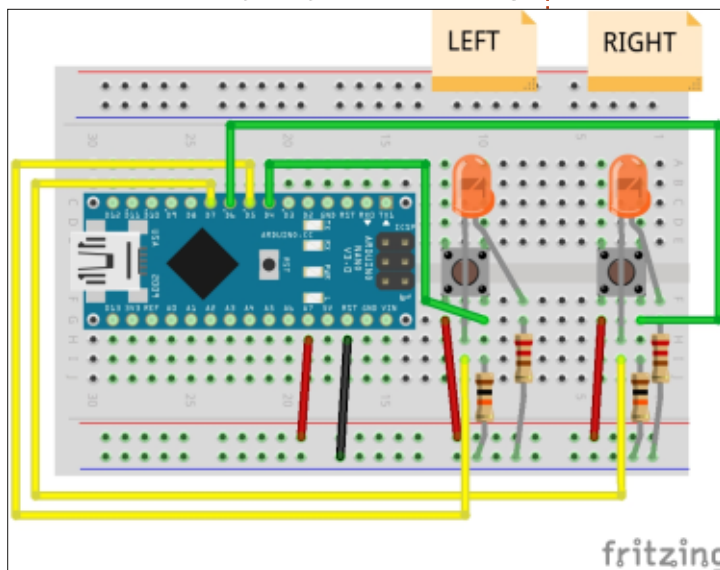
GUIDELINES

The single rule for an article is that **it must somehow be linked to Ubuntu or one of the many derivatives of Ubuntu (Kubuntu, Xubuntu, Lubuntu, etc).**

RULES

• There is no word limit for articles, but be advised that long articles may be split across several issues.

• For advice, please refer to the **Official Full Circle Style Guide**: http://url.fullcirclemagazine.org/75d471

• Write your article in whichever software you choose, I would recommend LibreOffice, but most importantly - **PLEASE SPELL AND GRAMMAR CHECK IT!**

• In your article, please indicate where you would like a particular image to be placed by indicating the image name in a new paragraph or by embedding the image in the ODT (Open Office) document.

• Images should be JPG, no wider than 800 pixels, and use low compression.

• Do <u>not</u> use tables or any type of **bold** or *italic* formatting.

If you are writing a review, please follow these guidelines :

When you are ready to submit your article please email it to: articles@fullcirclemagazine.org

TRANSLATIONS

If you would like to translate Full Circle into your native language please send an email to ronnie@fullcirclemagazine.org and we will either put you in touch with an existing team, or give you access to the raw text to translate from. With a completed PDF, you will be able to upload your file to the main Full Circle site.

REVIEWS

GAMES/APPLICATIONS
When reviewing games/applications please state clearly:

• title of the game
• who makes the game
• is it free, or a paid download?
• where to get it from (give download/homepage URL)
• is it Linux native, or did you use Wine?
• your marks out of five
• a summary with positive and negative points

HARDWARE
When reviewing hardware please state clearly:

• make and model of the hardware
• what category would you put this hardware into?
• any glitches that you may have had while using the hardware?
• easy to get the hardware working in Linux?
• did you have to use Windows drivers?
• marks out of five
• a summary with positive and negative points

You <u>don't</u> need to be an expert to write an article - write about the games, applications and hardware that you use every day.

One of the ways Linux has been promoted in the past was as a replacement operating system for older computers. As Linux distributions evolve and change, they often demand more resources. Last month Les, Olly and Tony of the Full Circle Podcast banded together to install Ubuntu on 3 systems: one which barely met the 700MHz, 768MB RAM minimum requirements, a Pentium 4 1.6GHz-based system, and a Core 2 Duo-based system – with poor results

for the first two installs. Check out their experience by downloading Full Circle Podcast #41: http://fullcirclemagazine.org/2014/06/08/full-circle-podcast-episode-41-trusted-to-fail/

At our computer reuse project, our minimum build is a Pentium 4 class computer running at 2.4GHz or better with at least 1GB of RAM and an 80GB hard drive. We set this minimum rather arbitrarily, a 2GHz computer would work just as well, especially if you increased the RAM to 2GB. For this article, I randomly picked a Pentium 4 class machine

from our back shelf. The machine I picked was a Dell Dimension 3100 with the following specifications: Pentium 4HT 2.8GHz, 2 x 256MB DDR2 533MHz (512MB total), Samsung HD080 HJ/P 80GB Hard Drive, Sony CD-RW CRx217E, 8MB onboard video RAM.

When Les, Olly and Tony approached their build, it was to test full-blown Ubuntu with Unity on all 3 platforms. This is a great method for modern hardware, but it's a bit like trying to put a 1.6L engine in a 2-tonne truck - you just don't have enough power to drive

the interface. On "other" platforms, we all know the gulf between the minimum requirements and recommended requirements for a program can be enormous.

The solution is a two-pronged approach - upgrade what hardware you can and use software/an interface that works with your hardware. Of course you have to keep the end user in mind when changing the interface, but the great thing about Linux is that it's so customizable. With a bit of customization you can ease someone's transition to Linux. Take the screenshot of Xubuntu 14.04, shown left, for example.

Normally when we build a refurbished computer, the first step is evaluating the computer. Our Dell Dimension 3100 has several shortcomings:

First there is no PCIe 16x or AGP expansion slot in this particular model of computer. PCIe 16x video cards are the current common standard. AGP, while old, is
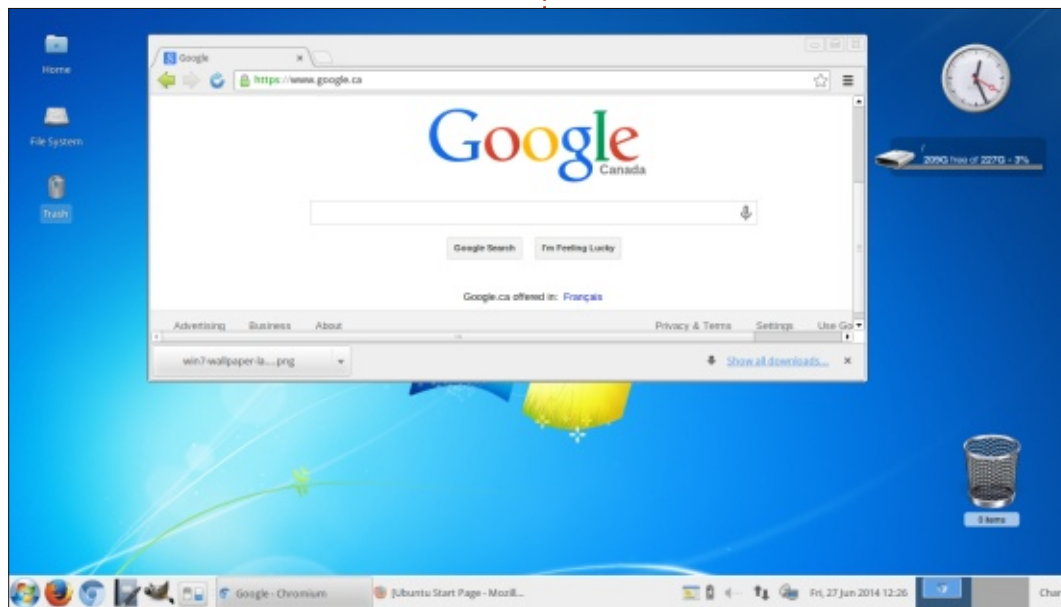
relatively easy to find on Kijiji, eBay or from most places that sell used computer parts. The Dimension 3100 actually has a PCIe 1x slot, so it's possible a better video card could be found, but in our computer refurbishing project we've seen exactly one PCIe 1x video card over the past year.

Our Dimension 3100 also had only two memory slots, thereby limiting RAM options. On the positive side of things, the RAM type was DDR2. While DDR2 isn't as readily available as DDR3, it's a lot easier to get than older DDR.

Our machine also had a CD-RW drive. As a standard practice at our computer recycling project, we always replace CD-R or CD-RW drives with at least a DVD-R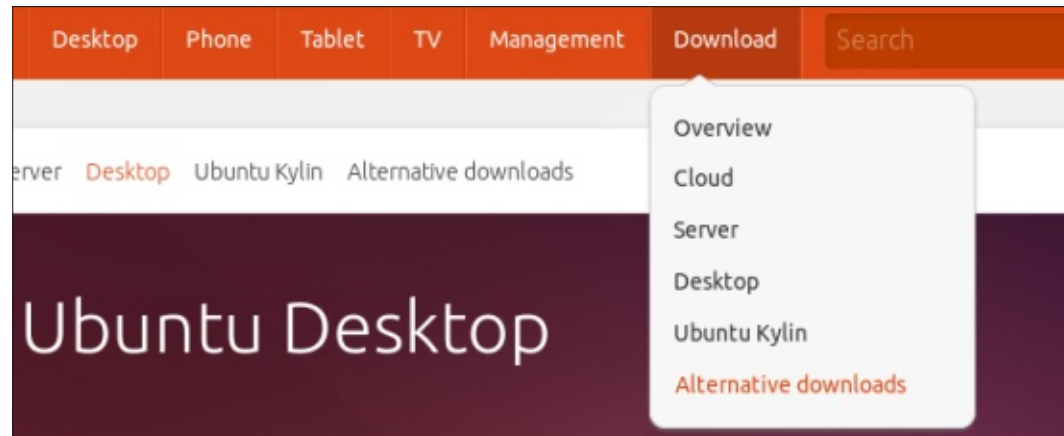 or DVD-RW drive. One of the problems the Full Circle Podcast crew mentioned was the limiting size of the CD drive in one of their machines; this won't actually be a problem for our install since we won't be installing using the same ISO image.

Back in the "old days" of Ubuntu 4.10, the installer was text based. Graphical installers are all the rage these days because they're usually part of a live DVD so you can try the environment before installing. Graphical installers are nice, but text-menu-based installers are often faster and give you more control over your install. Instead of grabbing the main 32-bit or 64-bit installer from http://www.ubuntu.com/ I grabbed the network installer (a mini.iso) from the alternative downloads link: http://www.ubuntu.com/download/alternative-downloads

Depending on the machine, and the person who wants the machine, I'll often dig a little deeper into the machine specifications. Our Dimension 3100, while containing only a single 2.8GHz Hyper Threaded core CPU, actually supports 64bit instructions. However, since we have only 2GB of RAM available, we went with a 32-bit version of the network installer.

The fastest method for installing the mini network ISO would be to use Unetbootin to create a bootable network install USB key: http://unetbootin.sourceforge.net. Our Dimension 3100 supported booting from USB key and we set it to boot from our key, but it didn't work. We double checked the BIOS settings to make sure the 'USB no boot' option wasn't set (it wasn't). USB was enabled in the BIOS. Our USB keyboard and mouse worked fine in all ports.

The culprit could have been our USB key. If a key isn't unmounted correctly then data corruption can occur. But in this case the culprit was the USB KVM (Keyboard Video Mouse) switch we were using to switch between 2 computers. Our experience has been that USB-based KVM switches tend to draw off enough power to cause problems with USB keys. The switches appear to work fine because USB Keyboards and Mice work when plugged in. But anything that needs a bit more power doesn't work with the USB KVM switch (which is the reason some USB KVM switches have a power connector for external power from the wall, of course they almost never come with an external power adapter). Once we set up the computer independent of the KVM switch, it booted to the USB key and we were able to start the network installer.

Earlier I mentioned that the network installer is pretty fast, but, if you have a slow connection, your install may be quite slow. To install the base system (command prompt) it took around 25 minutes on our single core Pentium 4 (start to finish). This time includes the time it took to respond to several questions the text installer asks. While fast, the text installer also asks more questions than the graphical installer (questions like

which proxy to use).

Once the install was finished we were left with a console login. To make the computer more useful to general computer users, we needed to install a desktop environment. Unity is a great desktop environment – the more I use it, the more I like it, but it doesn't work well with all video cards, and is a bit of a resource hog compared to XFCE, LXDE and Enlightenment.

Normally we install the Xubuntu desktop environment for all the Linux computers we build. Standardizing on one desktop environment makes it easy to support someone when they're having problems. To install XFCE plus most of the packages found in the Xubuntu desktop it takes only one command:

```
sudo apt-get install xubuntu-desktop
```

Enlightenment is a wicked fast desktop that works well on just about any video card. Initially we installed this because it installs almost 1,000 packages fewer than xubuntu-desktop. The problem we ran into was that we also needed

to install the xserver (xserver-org) and a login manager (lightdm). Installing these packages took only another 10 minutes, but we really didn't have a good selection of software afterwards, so we just ended up installing xubuntu-desktop.

Xubuntu worked surprisingly well considering the 8MB of video RAM. We've worked with 1MB and 8MB video on older Small Form Factor Compaq EVOs (D51) and

found all kinds of problems: screen tearing, choppy video, stuttering. Even with a 32MB video card the Compaq EVO never worked very well for us. And on other generic systems which had only 1MB-8MB of video, an instance where the location bar of Firefox would sometimes blacken so you couldn't see the text until you clicked back in the location bar.

We watched a youtube video on our Dimension 3100. The video

displayed with slight pauses but it wasn't stuttering enough to make it unwatchable. Firefox felt a bit sluggish, but there were no display problems and no screen tearing. Programs generally loaded fast (of course nothing like a high-end quad core, but with enough speed to get work done).

The takeaway from all this is that it's possible to install a working Linux-based desktop environment on low-end hardware that would suit the "average computer user." Video card upgrades (if possible) and some extra RAM go a long way.

The AGP video architecture might be old by today's standards but a 256MB or 512MB AGP can still give your system a big boost. If you are looking for video cards of this vintage, we found it best to go with NVidia-based cards since you can normally find proprietary drivers in the "Additional Drivers" program (in Xubuntu 14.04 you can find it bundled in Settings Manager). We recently put a 512MB AMD Radeon HD 2600 AGP in a system and it performed worse than a 128MB NVidia card simply because of the drivers available.
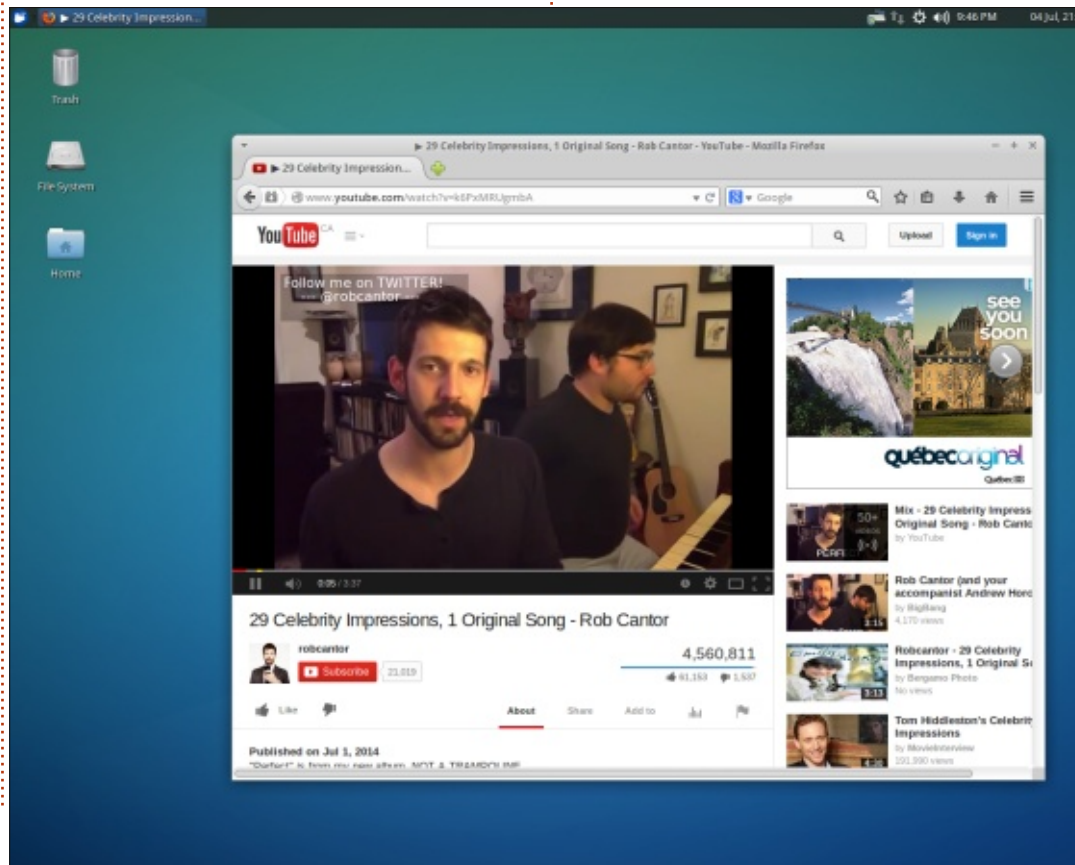
For extra parts, check to see if there's a local computer refurbishing or recycling project in your area; parts can be a lot less than you'll find on eBay, you save yourself shipping, and you have somewhere you can go back to if a part doesn't work. Sites like kijiji are awash with computers and parts. If you buy a second-hand computer from Kijiji there are a few things to look for (if you can, not all sellers will allow you to fully examine equipment):

Check the USB ports, are some damaged, are they filled with dust, or tar from cigarettes? Don't buy anything you don't feel is well maintained.

If you can examine the inside of the computer. Look for bulging or blown capacitors on the motherboard. Check there are no obviously missing parts (battery, heat sinks, hard drive, RAM, etc) and that things look in good order. A look inside will also tell you if the computer has been well maintained (dust, tar or worse). Systems often can turn on with blown capacitors, but might behave erratically, it's best to stay away from these systems, even if you intend to replace the

capacitors. If capacitors are blown, there might be something else that has electrical damage that's less visible.

Again if the seller is willing, watch the system boot, even if it's just to the BIOS. As a computer refurbisher, we used to recommend that people who buy on Kijiji also ensure they have a Windows COA.

Check the drives. Does the CD/DVD drive even open? We see lots of computers come into our project that have nice DVD writers that have been jammed with broken CDs, pencils, money, and other things that should never go

in a DVD drive.

If you're buying an older Dell computer, many Dell systems have 4 small lights on the back. Check to see that all the lights are green, if a couple are orange there is a problem with the computer. Dell provides documentation on their website on the diagnostic lights for each model of computer they sell. I'm not endorsing them here, just mentioning a useful tip that leads to my next point.

When you can, get as much documentation about the model of system (or motherboard if it's a generic aka "whitebox" system). Good documentation will tell you about expansion options (how

much RAM a system will take, which version of PCIe the system supports (2.0, 3.0, etc.) and can help when troubleshooting problems.

And if all else fails, try a different distribution. The great thing about Linux is there's always choice. Choice gives us a wonderful array of desktop managers: KDE, Gnome, Unity, WindowMaker, Enlightenment, LXDE, BlackBox, etc. Choice gives us a wonderful array of software too, so if something's just too heavy for your hardware try a lighter option: use Abiword instead of LibreOffice Writer (or even something web based like Google Docs). Choice is the reason why Linux is still great on older hardware.

I'll let the BitTorrent Sync website explain (in several quick quotes) what the product does:

"Sync never stores your files on servers, so they stay safe from data breaches and prying eyes. Create a Sync folder for your stuff. Open it on any computer, phone, or tablet. Access docs, share photos, and start working from anywhere. BitTorrent Sync skips the cloud to deliver your files at lightning speed. No matter where you are. Sync, send and share as much as you want. There are no file size limits, or caps on creativity. It's as simple as that.

First you'll want to go to http://www.bittorrent.com/sync/downloads and download either the x86 (32-bit) or x64 (64-bit) version. You'll get an archive file which contains two files; a license (which we don't need) and an executable file.

Unarchive the btsync_x64 (or btsync_x86) .tar.gz file. I have my btsync file (the executable) in /home/ronnie/BitTorrent Sync.

Run it by opening a terminal and typing (in my case):

```
cd BitTorrent\ Sync/

./btsync
```

There's no GUI, it's web based, so you can now close the terminal. Open your web browser and go to http://localhost:8888. You'll need to agree to their terms and conditions. After that you'll need to create an account (shown below).
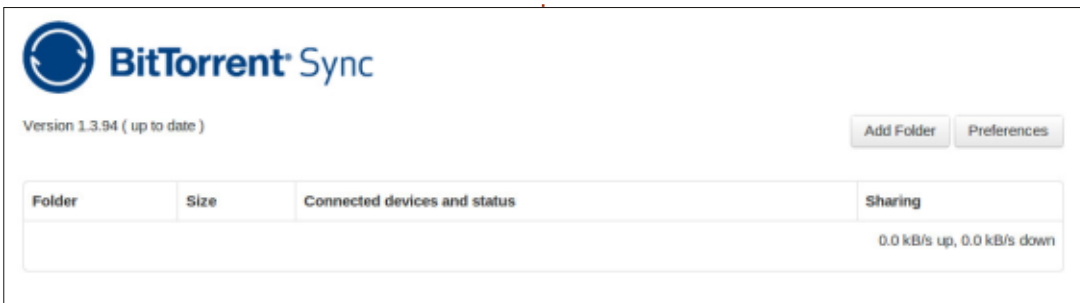
Enter a name and password and click OK.

With the account created it'll ask you to login again. Congratulations, you're in

BitTorrent Sync (shown above).

Now create a folder on your machine that will be synced to other machines. Do not put the sync folder inside the BitTorrent Sync folder.

In the BitTorrent Sync WebUI you click 'Add Folder' (shown on the next page, top left).

Here you give it the folder you want to sync (in my case it's 'BitTorrent Sync Folder') then click 'Generate'. This will create a 'secret' (or a code) that is unique and identifies that folder. Finally, click 'Add'.

Your folder is now ready to send, but it's empty. For this initial test, I've added an image to my sync folder. After a few seconds
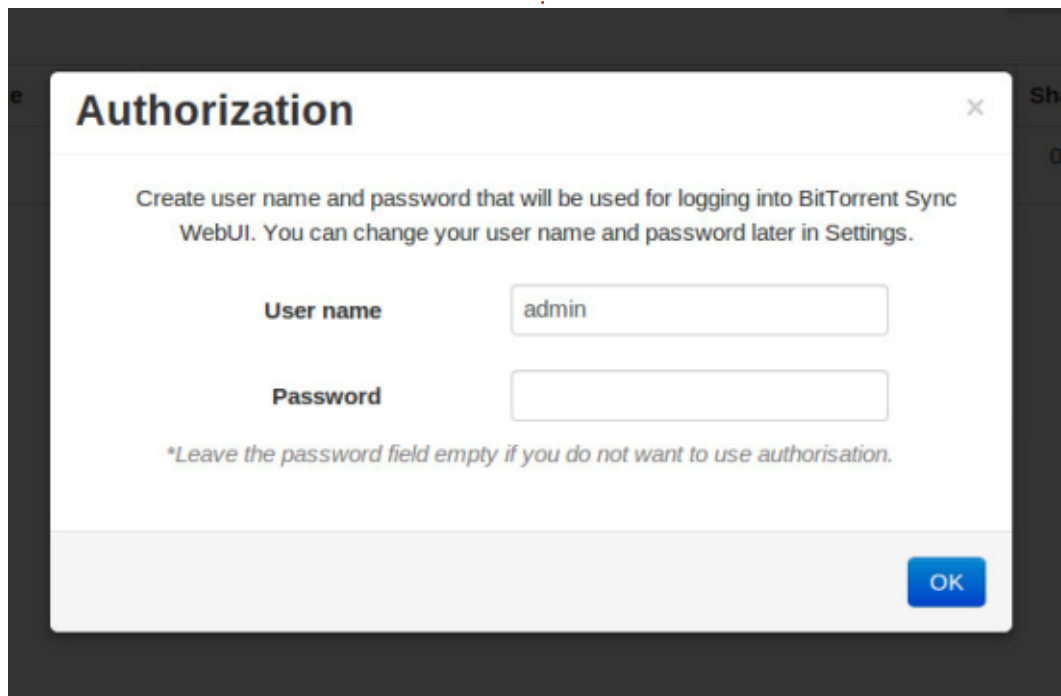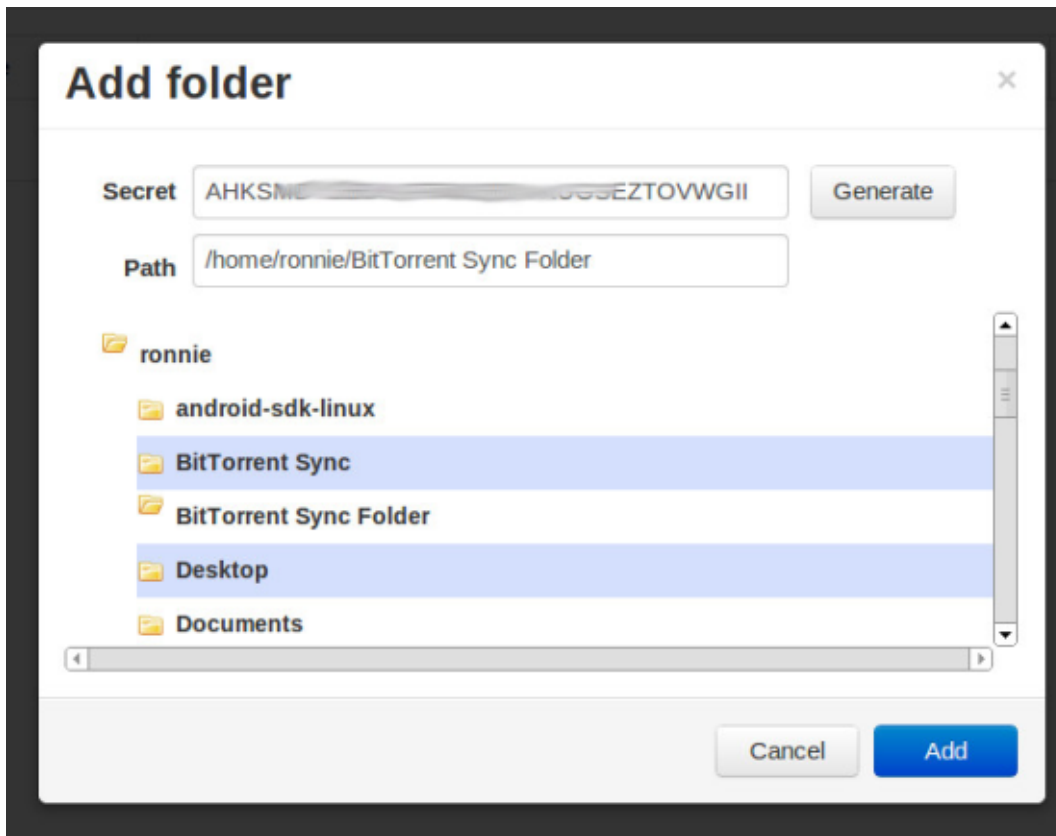
the 'Size' for that folder quickly updates and says 198 KB in 1 files. It has indexed it and is ready to send, but it has nowhere to send it. I'll now add BitTorrent Sync to my laptop using the same method as above (download the archive, unarchive it, run the btsync file, and open it in my web browser).

To get the folder on my laptop, I first create a folder on my laptop (you can call it anything, but I gave it the same name as on my desktop machine), then click 'Add Folder'. In the popup I point it to the folder I just created but do NOT click Generate. Here I type in the 'secret' that was generated on my desktop machine. Finally, I click OK.

My laptop now shows that it synced my one file and is linked to ronnie-desktop (shown below left).

The screenshot of my laptop screen was saved into my sync folder, and, when I went to my desktop machine, the screenshot was in my sync folder and I could see that it was connected to my laptop and that I now had two files in my folder.

One thing to watch out for, though, with a laptop: beware of closing the lid and the laptop closing the internet connection and going into hibernation or low power. That will stop the sync between machines. If possible, change your settings to keep the internet connection alive if nothing else.

Oops!

So let's say I accidentally delete the screenshot I took on my laptop. My webUI now shows 1 file left. I could probably recover the file from my rubbish bin, but what if I'm on my laptop? Well, if you show hidden files, and go into the sync folder you'll see that there's a .SyncArchive folder. That folder holds the file I deleted, but it's only stored on the machine that it was created on. So while it doesn't show up at all on my desktop machine, the deleted file is in the .SyncArchive folder on my laptop.

Mobile Devices

If you'd like to send your sync folder to your mobile device, then you need to install the BitTorrent Sync app and a barcode/QR scanning app. I'm using Android and got the app via the Google Play store.
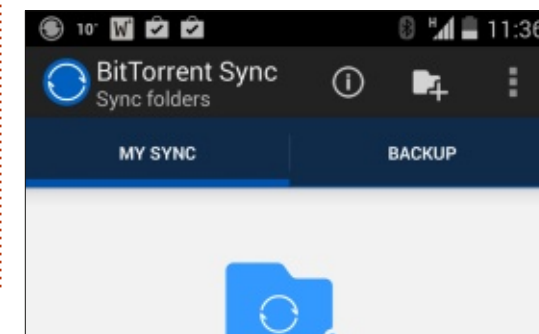
Once installed and running you simply tap the screen to add a folder.
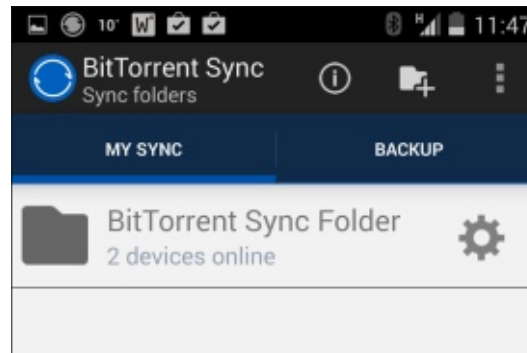
You're now asked to choose a folder on your device. This is where it will save the synced files. Next, you can either type in the big long 'secret' (best of luck there!), or you can scan the associated QR code. To get the QR code you click the 'Secret / QR' button which is on your webUI.
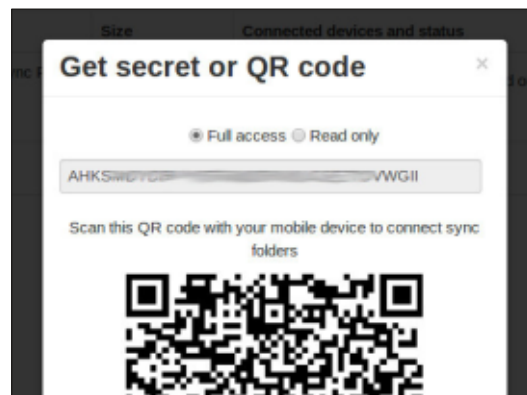
Finally, I ticked the 'Automatically' sync' box as I have unlimited data on my mobile phone. If you're using mobile data, then you'll need to go into the app settings and enable it to use mobile data. This is off by default to save you from suddenly getting a large download which may cost you money on bandwidth.

One thing to watch here is that you don't copy large files into the sync folder which then have to be downloaded to your mobile device.

You can also assign folders on your device to be backed up which is handy for photo folders and such like.

Using a file manager (on your device), you can place files into your sync folder and they will, of course, appear on your desktop/laptop machines. I used that idea to get the screenshots from my phone to my desktop (without using my wifi connection as a test) and that let me easily insert the screens into this article.

## STARTUP

You'll obviously want to add the btsync application to your startup. How to do this varies by desktop, but in Mint it's done by going into the System Settings > Start-up Programs and clicking 'Add', giving it a name and the command to run (using the Browse button).

READ-ONLY

BitTorrent Sync also has a nice 'Read-Only' feature that you can use. When generating a 'secret', you can choose to have it read-only. Handy for sharing with a friend who just needs to grab files, but not make changes. You can see the two 'secrets' for each folder by clicking the 'cog' icon (beside the red X icon) and clicking the 'Advanced' tab. In there you can also generate a 'one-time secret' which is valid for only 24 hours.

CONCLUSION

The only downside that I can see with BitTorrent Sync is that there's no server (like Drive or Dropbox) that holds your files should your desktop/laptop have an accident. Having said that, I think the whole idea behind BitTorrent Sync is that you don't rely on a company to hold your files as they may be up to no good with your stuff.

The fact that it lets you sync to mobile devices is a great idea, and handy for quickly moving files from device to desktop. I usually use AirDroid for getting files to my desktop, but BitTorrent Sync may make things much faster and easier.

Ronnie is the founder and (still!) editor of Full Circle. He's a part-time arts and crafts sort of guy, and now an Arduino tinkerer.

I have been a long time user of Ubuntu and a personal computer enthusiast since the Tandy 1000 with the operating system built into rom. I typically stay with the LTS releases because I need my computer to work well with the programs (apps, for the young readers out there) that I need to get things done, and sometimes upgrading versions can introduce unexpected results. I am currently running Ubuntu 12.04 and Unity.

I thought I might chronicle my move to a solid state drive or SSD. Most PC enthusiasts say that moving over to an SSD is one of the most noticeable speed upgrades one can make to a computer. When SSDs first came out, they were small in capacity and high in price, and reliability was questionable. Today, SSDs, while still more expensive per gigabyte than traditional hard drives, seem ready for prime time. With Ubuntu 14.04 LTS shipping with trim support enabled by default, it felt like it was time to give an SSD a try. My computer is far from new. My motherboard is an older socket 775 using a core 2 quad cpu and 4 gigabytes of RAM. It has SATA connections but only the slower 3Gb/s connections. So while it should be faster than my hard drive, it won't be as fast as one on a new modern motherboard.

I have a USB hard drive that I back up to regularly using the Deja Dup backup tool. It's good practice to also back up your data without any compression. I'll copy and paste a second copy of everything into a different folder just before I'm ready to pull out the drive. (Don't forget your bookmarks!) I am also going to install the new SSD on its own at first. Although I intend to ultimately use the old hard drive as storage, I am not willing to risk everything on just one backup device holding all my data. Before I do anything, I'll write down all the programs that I need or want to reinstall on the new system, and verify my backup. My plan is to install the SSD, do a clean install of Ubuntu 14.04 LTS, and update it. At that point I'll install the programs I need and test them. If all seems fine, I'll move my data from the USB drive back into my new operating system. The Linux experts out there are, no doubt, saying right now that there are easier ways of getting this done, and they would be right. But I have ruined a perfectly good Ubuntu installation too many times to trust myself to untried command-line entries. The old hard drive will stay in the computer, disconnected for a couple of months until I am confident everything works. Ok, plan in place. Lets see if it works.

Choosing an SSD is harder than you might expect. Choosing a size that I can live with for a while was first on the list. If I'm going to spend hard earned money on an SSD, I want to use that SSD for almost everything. I currently have a 500 GB hard drive and have used 133 GB of space on that drive. A 256 or 500 GB drive should be fine for my needs, but the 500 GB would be a little more future proof. Price and warranty are the next factors to consider. I have found that price typically reflects the capacity of the drive, the type of memory and controller firmware used in the device, and the length of warranty the manufacturer provides. With my current setup, the speed of the drive isn't necessarily a deciding factor as most any SSD will be faster than what my motherboard can utilize. But, if a shiny new computer is on the horizon, and you want to migrate your new SSD to that computer, the speed of the device should be considered. Based on my needs and budget, I chose a Crucial M500 240 gigabyte drive.

Attaching the new SSD to the SATA cable and power was no problem but my old case doesn't have the proper bays for attaching the drive to the case. The solution is to order adapter brackets with the drive, or break out the zip ties and duct tape. A word on firmware. SSD makers are always refining and tweaking the firmware (or instruction set) that SSD's use to operate. Updating the firmware on your new SSD should be first on the list of things to do before you spend all that time setting up your new operating system. Although the process is usually safe to do

anytime, why risk it? A quick trip to the manufacturer's web site showed that there was an update available. Now a bit of panic set in. There is a file for Windows and a file for Mac, but no Linux. Fortunately the Mac file was an iso file to make a bootable cd. Once downloaded, I burned the iso to a CD with Brasero and rebooted with the cd in the drive. My computer is set up to boot from a CD but you may have to watch the prompts closely when you boot – to hit a function key and select your boot preference. I followed the prompts from the firmware update CD and, in moments, I had completed the update process. With the drive in place, and firmware updated, it's time to boot the computer with the Ubuntu live DVD and see what it finds.

Ubuntu found the new SSD and in a few minutes installation was done. Had it not been for my slow internet connection pulling the updates and installing the restricted extras, it would have finished even faster. The reboot was where I first noticed the speed of this drive. A program I use regularly is LibreOffice writer. It opened almost instantly. Everything is fast on this SSD. It's

pure joy to turn on and play with. Click on something and there it is.

Ok, back to business. After purchasing my SSD, I found that Ubuntu 14.04 LTS is only enabling trim on Samsung and Intel SSD's out of the box. Without trim, my drive will surely slow as it ages. A search of the web turned up several sites detailing how to enable trim. I chose to follow an article on the Webupd8.org site. The article can be found here: http://www.webupd8.org/2013/01/enable-trim-on-SSD-solid-state-drives.html. Two methods are outlined. I used the recommended way and so far I have had no problems. Their write-up is excellent and I just copied and pasted the necessary commands into the terminal.

After installing all the programs I need and copying my data back to the new SSD, the speed still amazes me. Everything is working great except for one small problem. I use Audacity occasionally and the ffmpeg codec I need to convert files in that program is not available for 14.04 in the repositories. Because of this, I still need to keep my 12.04 installation handy. It's not very

convenient swapping cables to access the old hard drive, maybe a virtual machine with 12.04 will be on the short-list. Early adoption isn't always the best choice. This put a snag in using the old hard drive as storage for now. I found a Western Digital 1 terabyte drive on sale and a SATA cable at a local big box store that was too good to pass up. Here is where my second problem came up. As soon as I plugged my new storage drive in, the computer would no longer boot. An error message indicated that there was no operating system present. I unplugged the SSD and plugged in the old drive with 12.04. The machine booted right up. At this point my first thought was that my new SSD was dead. Don't panic I thought. Put everything back the way it was before it quit. The machine booted right up. Now add the new hard drive back and no boot. As it turned out my motherboard would favor the hard drive over the SSD in the bios. Setting the SSD as the primary drive and making sure the boot order was set correctly solved the problem. Using the disks program in Ubuntu, I formatted my new hard drive to Ext4. To do this, start by selecting the hard drive in the devices column and click on the

gear icon in the right hand corner, then choose format. Give your drive a name. I called mine "storage". Not very creative, but I know what it is. With this second drive now ready for storage and backup, I can now make a new separate backup on the USB hard drive and take it off-site in case of a major catastrophe. The only other side effect to this upgrade is that any other computer I use now feels so slow.

This project brought back that fun factor again. Knowing that I have good backups of my data, the fear one feels when doing an operating system overhaul was gone. If everything went belly up, I had my old hard drive with 12.04 on hand to go back to for answers. I can say without any reservation that an SSD is a great way to put that snap back into your rig. Take your time, do your homework, back up and back up again. It's like having a new computer without the big cost and you do it yourself. Just make sure you have some fun with it too!

On pages 10 and 11 of FCM#84 (congratulations to Ronnie with the 7th birthday) Arnfried Walbrecht wrote about Linux taking over the desktop and laptop world. I see things a little different.

Yes, I know he was talking about Chrome-OS and the Chromebooks, which are a great solution for people who are doing (almost) nothing else with their computer than browsing the net.

But seriously, Linux taking over the desktop and laptop? I don't think so.

The reason for my negativity is the fact that the Linux world is so divided. A phrase you can read everywhere is: Linux is all about choices. That's great. But I think nobody sees that those choices are what is keeping Linux a small player.

There is one Microsoft and there is one Apple, while there are how many Linux distributions? Distrowatch shows the Top 100 in its list. Each distro mentioned there has a couple of different versions to choose from, making the total much more.

Does anybody have an idea how many different Linux distributions there are? Everybody who has some degree in programming starts his own distro, based on another, which is based on a third already. A few changes, an extra program added to the base installation, a new theme to make it look good and a new star is born.

Why? To make a living? Does it pay that well to make a product which you give away? Can't imagine that.

Wouldn't it be better if all these programmers would join into a large group of programmers to not make a different distro but a better one? To focus on what is really needed to make a rock-solid, fast, safe distro.

One which you install without any problems, one which detects all of your hardware without the need of having to search for drivers which have to be installed later, one which is rock solid so you can enjoy it for years, one which is fast so it makes it fun to work with, one which is safe.

Isn't that more important than having yet another distro which is 99.9% the same as the one it is derived from? The 0.1% being a color green instead of a color purple, some extra codecs installed to make it ready for use out of the box, or in other words: peanuts.

I mean look at what a Linux distribution is. You have:
• a Linux kernel, which all distros use
• a desktop environment, of which I think there are too many already
• programs made by 3rd parties, like an office suite, media players, web-browsers, etc
• a company theme which many people throw overboard immediately after installing to have their own choice of theme with matching desktop background

What can you make different in a distro? Not much. Sure, you can choose another DE with the matching programs like a file manager, a CD/DVD burning program, etc, but since many companies supply different versions, all these DEs are used already. So, no differences there.

Why do people keep trying to re-invent the wheel? In case you didn't know it yet, it is there already.

Why not join and make fewer different products, but make them faster, more stable and more safe. Make a few different distros so there is still something to choose. Since the beginning there have been rpm-based distros and there have been deb-based distros. Some 3rd party software appears in rpm and not in deb format, others the other way around. Conclusion: you can't get all the software you need without being a computer nerd and change an rpm into a deb (or vice versa).

Winning the world when you still have to do things like this is impossible. The world is bigger

than a group of nerds. Having to use the terminal to do things "easier" and "faster" is also something which, I think, should change.

When Linux wants to take over the software world things should really change, change big time.

I have used different distros since 2008 and I can manage myself pretty well nowadays. Still I need help, so I Google for answers, or I use the forums. Happily I can say that almost all of the time I get the solution I need. I am not a nerd, nor will I ever be, I am a computer user, not a programmer, not an IT specialist. The computer has to do what I want it to do, not more, not less.

I would love to see Linux taking over the software world, I really believe in it, but, as long as programmers and the companies behind them are fighting each other instead of working together to get the perfect OS, this will not happen.

EXTRA! EXTRA! READ ALL ABOUT IT!

Our glorious news reporters are now posting regular news updates to the main Full Circle site.

Click the NEWS link, in the site menu at the top of the page, and you'll see the news headlines.

Alternatively, look on the right side of any page on the site, and you'll see the five latest news posts.

Feel free to discuss the news items. It's maybe something that can spill back from the site into the magazine. **Enjoy!**

PYTHON SPECIAL EDITIONS:

http://fullcirclemagazine.org/issue-py01/

http://fullcirclemagazine.org/issue-py02/

http://fullcirclemagazine.org/python-special-edition-issue-three/

http://fullcirclemagazine.org/python-special-edition-volume-four/

http://fullcirclemagazine.org/python-special-edition-volume-five/

http://fullcirclemagazine.org/python-special-edition-volume-six/

Puppet Reporting and Monitoring
by Michael Duffy
http://www.packtpub.com/puppet-reporting-and-monitoring/book

Michael Duffy is an experienced DevOps engineer, and works at Sky.com. Monitoring is not bad, but is old and costs too much time and too much money. Too much time for your administrators and programmers and too much money for you.

Puppet is possibly the fastest growing configuration management tool on the planet, and this is in no small part due to its combination of power and accessibility.

This is a shame; Puppet's reporting capability is one of its most overlooked yet powerful features.

If used correctly, its built-in reporting abilities can give you stunning levels of detail about your infrastructure, from the amount of hardware used and networking details to details about how and when resources were changed.

This book is designed for anyone who wants to learn more about the fundamental components of Puppet reporting.

To get the most out of this book, you should already be familiar with Puppet and be comfortable with its major components such as the Puppet master and Puppet agent. You should also be comfortable with reading code, and in particular, you should be at least passingly familiar with Ruby. Finally, you should be happy working on the command-line in the Linux/Unix flavor of your choice.

The metrics that the Puppet agent passes to the Puppet master are very granular and offer a fantastic insight into where Puppet is spending its time, be it fetching, processing, or applying changes.

The dashboards can be used with Puppet and take a whistle-stop tour of some of the major features that each of them has. The dashboards can offer some quick and easy reporting options but also have some limitations.

In chapter 6 we can explore the PuppetDB query API. Data is useful only if you have some means to access it, and yet this is a truism that many systems seem to have forgotten, relying instead on developers to come along and fill whatever egregious gaps in data exploration the original product left out. Fortunately, Puppet offers a rich data discovery tool in the form of the PuppetDB API and its associated query language.

In the next chapter you can learn how to write "Custom Reports" with PuppetDB and create a menu-driven application.

In chapter 9, the final one, we find a recap of Puppet dashboards and integration with third-party components, looking back at the alerting feature and integration with external alerting systems, analyzing metrics and changes with Graphite, and anomaly detection with Etsy Skyline.

Hopefully, by now you are looking at the Puppet reporting tools as a gateway that allows Puppet to communicate with the wealth of systems that you are already using to both monitor and report, and thinking of new ways to use these tools with the additional data that Puppet provides.

I've read some Puppet books, this is probably one of the best written. The language is clear and

fresh, topics are well expressed and explained. Even if the contents revolve around Puppet's reporting functionalities and PuppetDB, the author takes, at times, the occasion to explain more general Puppet features and functionalities.

In this book, you can find an easy-to-follow guide with extensive examples to explain Puppet's reporting capabilities to facilitate effective implementation of Puppet as a reporting tool in the real world.

Experienced users may be familiar with most of the contents, but, even for them, in the very last pages of the book, a very interesting concept is expressed which might lead to quite useful use cases: the possibility to use Puppet's reporting as the drive to infrastructure orchestration activities.

Calogero is a consultant in business processes and "collective intelligence". He uses Puppet at SEATPG, which operates the largest local ad platform in Italy, connecting consumers to businesses and helping businesses and communities grow.

LibreOffice Special Editions:

http://fullcirclemagazine.org/libreoffice-special-edition-volume-one/

http://fullcirclemagazine.org/libreoffice-special-edition-volume-two/

http://fullcirclemagazine.org/libreoffice-special-edition-volume-three/

Inkscape Special Editions:

http://fullcirclemagazine.org/inkscape-special-edition-volume-one/

http://fullcirclemagazine.org/inkscape-special-edition-volume-two/

Let's face it, mining digital currency nowadays is hardly worth it. It'll cost you more in electricity than you'll make in [insert favourite coin here]. Having said that, ASIC miners are all the rage these days for Bitcoin mining (which has a difficulty that is sky-high and will take supercomputers to get anything from it). Not so good for folks who want to try the lower end stuff (such as Litecoin) which use scrypt cryptography. Enter, stage left, the new fad of USB miners which use hardly any electricity.

The first wave of USB miners were Bitcoin only, but now they're releasing dualminers which can mine both Bitcoin and scrypt coins.

To use one of these, you need to get the latest video drivers first (Nvidia for me).

Next, we need to get some software to run with the Dualminer. Unfortunately, none are available 'off the shelf', so we'll need to compile a special version of cgminer ourselves. Download the git repository in a terminal with:

```
git clone
https://github.com/dualminer/
dualminer-cgminer
```

After the download is complete:

```
cd dualminer-cgminer
```


Before trying anything, make the auto and configure script executable. Open a window, go into the dualminer-cgminer folder. Right-click on configure and choose Properties. Click the Permissions tab and tick the box beside 'allow executing file as program'.

Well need a couple of dependencies, so use your package manager to install libtool, libudev-dev and libncurses5-dev. Or use the terminal if you're up to it.

Time to get compiling. In a terminal:

```
sudo ./autogen.sh

sudo ./configure --enable-
dualminer --enable-scrypt --
disable-opencl --
prefix=/home/USERNAME/dualmin
der-cgminer
```

(obviously replace USERNAME with your username)

Continue with:

```
make && make install
```

Wait a little while until it's all compiled, then run cgminer with:

```
sudo ./cgminer --scrypt -o
stratum+tcp://ltc.give-me-
coins.com:3333 -u
wakuangdenongmin.1 -p 1
```

NOTE: for me I had to run it with sudo otherwise it wouldn't detect my device

Now wait. Probably quite some time. It takes me days to get about 0.01 of a Litecoin. The problem is exacerbated by the pool getting new blocks which displays:

```
Stratum from pool 0 detected
new block
```

and the whole mining process grinds to a halt for a short time.

LETTERS

If you would like to submit a letter for publication, compliment or complaint, please email it to: letters@fullcirclemagazine.org. PLEASE NOTE: some letters may be edited for space.

FIRESTARTER

I have a question about Firestarter. I was wondering why they got rid of it? Now they say to download GUFW. I don't know how to use this one as all I did with firestarter was download it and use it. According to the people who said to download GUFW, it has to be set up. I don't know how to do this. Could you help me with this?

Lou

Ronnie says: *Any GUFW experts out there want to write an article on how to set up GUFW?*

UBUNTU + MATE

Regarding FCM#86, "Ubuntu MATE Flavor Could Arrive Soon, Prototype Looks Great Already".

It is already possible to easily add the MATE Desktop to Ubuntu 14.04. I have recently installed MATE 1.8 on UbuntuStudio 14.04, which already had the standard UbuntuStudio and XFCE Desktops installed. MATE appears to be working satisfactorily. There are a few minor quirks to get used to, but nothing major.

This link gives all the details. A built-in version of MATE might be somewhat better.

http://ubuntuportal.com/2014/06/how-to-add-mate-desktop-environment-1-8-to-ubuntu-14-04-lts.html

jl

ROBOLINUX

I just wanted to write in to give my response to "Beware the RoboLinux" in the last issue, and share my experiences with the RoboLinux people.

I had been running a file server in my office that used Windows XP. I really wanted to build a new server that could take over almost all the functions of the old XP machine, but would run Ubuntu instead. I found the software from RoboLinux extremely helpful. My servers do not dual-boot. They run only one OS. Understand that the RoboLinux people actually have a few different projects going on at once. The first is RoboLinux itself, which is a Linux distro in its own right. I can't comment on it because I have been happily using Ubuntu since 4.10 (almost 10 years now), and wanted to stay with Ubuntu. The other Robolinux project involves taking an existing Windows XP or Windows 7 installation and creating an image that will import without problems into Virtualbox. I had read online for a while about doing that, and although you can do it yourself using all free tools, it is not without problems, takes a long time, and seems fairly confusing. The videos on the Robolinux site are a little vague, and short on the actual nuts and bolts of how to do this, but they promised easy, step-by-step instructions once you made a "contribution" to them.

As it turns out, the instructions

were very good (with one small exception that I will mention later). They told you how to create the image, convert it to a vbox compatible format, and install it in virtualbox with just a few mouse clicks. They include many nice screen-shots showing exactly how everything should look throughout the process. Mostly, you are just waiting for your computer to do the work for you. When it was done, the virtual machine started right up, and everything was right where it should have been. I did make one small change in the virtualbox settings for the Windows XP machine. I turned on the bridging network adapter so it would communicate with everything on my office network exactly as it had before.

The only thing missing from the instructions was the need to install the dkms package in Ubuntu first. The dkms package is maintained by Canonical and is part of the regular repositories, but is considered "optional" and is not installed by default. Otherwise, when you get a kernel update and reboot, Virtualbox will give you an error that you need to recompile the vbox drivers. It does give you the command to do it, which is

accomplished in a terminal, preceded by sudo of course. Not a big deal, but if you had the dkms package installed, it would do that for you automatically any time there is a kernel update. You can find the dkms package in software center, or synaptic, or just a good old "sudo apt-get install dkms" should install the package and all its dependencies.

I am very happy with the software and feel it was well worth the few dollars charged. I believe it saved me several hours of work at the very least and made my new Ubuntu 14.04 64-bit file server project everything I hoped it would be.

PS: I wrote to the RoboLinux people and told them about the dkms packages needed for Ubuntu and this was the reply from John at Robolinux:

"I'm going to slip a dkms install cli into our Ubuntu stealth vm software as I feel it will help users and make our software a more pleasant and complete solution for the Ubuntu community. I can do that dynamically since our scripts are specifically isolated for Ubuntu users and always accessible via our

dedicated Robolinux Debian servers."

They are quite responsive and supportive of the Ubuntu community. Overall, they were a pleasure to deal with. Highly recommended.
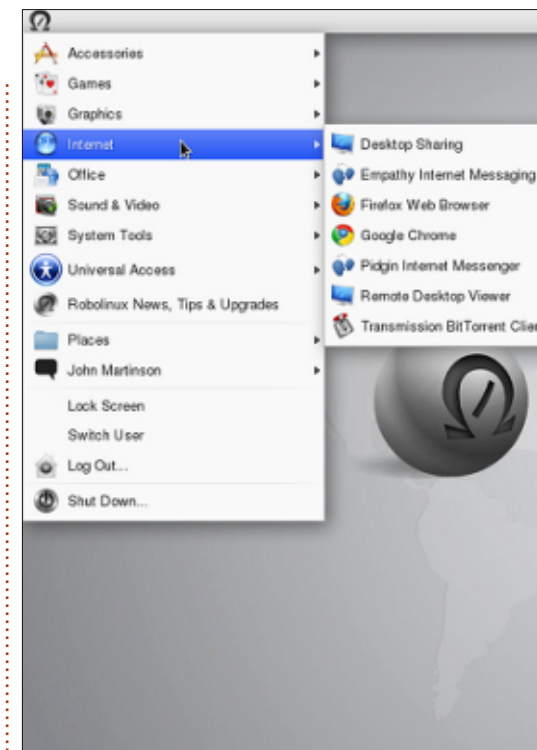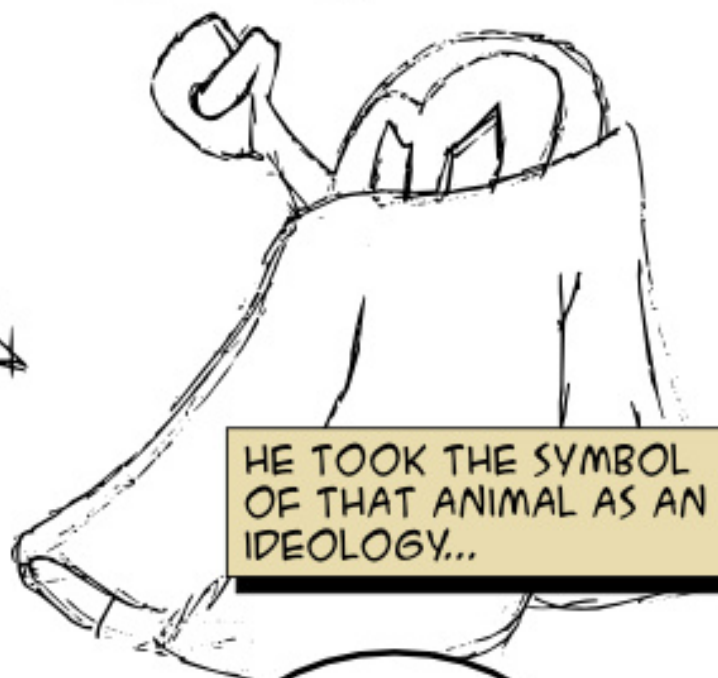
Martin G Miller

Q&A
Compiled by Gord Campbell

If you have a Linux question, email it to: misc@fullcirclemagazine.org, and Gord will answer them in a future issue. Please include as much information as you can about your query.

Q I'm trying to create a print server for an MFC-7860DW. Printing works properly, but after installing "cloudprint" using PIP, I get an error message.

A Sharing a printer is a lot easier than you thought. See: https://help.ubuntu.com/community/NetworkPrintingWithUbuntu

Q I am using Linux Mint Qiana with Kernel 3.13, and LMDE with Kernel 3.11.2, so why is V3.10 newsworthy?

A (Thanks to **Ronnie**) Kernel 3.10.41 is an LTS release, so it will receive support for several years. See https://www.kernel.org/category/releases.html

Q How do you do a simple sum in LibreOffice calc?

A Put the cursor where you want the sum, click the sum icon, press enter.

Q I have done many updates, and a couple of version upgrades, but Grub still shows me only the original kernel versions.

A (Thanks to RobertKH in the Ubuntu Forums) Run these commands:

```
sudo grub-install /dev/sda

sudo update-grub
```

Q I installed NS2, but I can't find where the files were installed in order to modify some of them. Specifically I want to modify aodv.h and aodv.cc for a project of mine.

A If you install Synaptic Package Manager and search for ns2, you can right-click and select "properties". One of the tabs is "installed files," which shows you all the files, and their locations. It works only for installed packages.

Q How can I make an application open where I want it in Kubuntu?

A (Thanks to **AnotherKevin** in the Ubuntu Forums.) Position the window to the place on your desktop you want it to go when you start the application. Next, click the icon in the upper left corner of the app's border. Choose More Actions >> Special Window Settings. In the box that follows, tick the boxes for Position and Size, then, using the drop-down lists, choose Remember.

Q How can I delay a startup application in 14.04?

A (Thanks to **mc4man** in the Ubuntu Forums) In your home folder, go to the hidden folder .config/autostart, and locate the .desktop file for your application. Edit the file, adding a line to the end:

`X-GNOME-Autostart-Delay=30`

where 30 is the number of seconds of delay you want.

It won't work for Dropbox, because Dropbox rewrites its .desktop file every time it runs.

Q I have a new laptop with Windows 8.1 installed. I plan to dual-boot with Ubuntu 14.04. Should I turn off UEFI and secure boot? When should I set up the partitions?

A No, leave UEFI and secure boot enabled. Use Windows to shrink the existing partition(s) so there is space for Ubuntu. Always select "something else" during installation, and explicitly set up the partitions as you want them.

TOP NEW QUESTIONS AT ASKUBUNTU

* Make it so you can't stop command with CTRL+C
http://goo.gl/KTWyYC

* When you uninstall a program on Ubuntu, does the software leave any residue?
http://goo.gl/1f6FQH

Q&A

* Does Ubuntu Touch consume less power than Android?
http://goo.gl/0hxFQX

* Is Ubuntu MAAS free?
http://goo.gl/iZR3vV

* What does TTY stand for?
http://goo.gl/yEdGMp

* Why doesn't cd work in a shell script?
http://goo.gl/JMupZu

* Unable to remove a file which has a name like a command argument
http://goo.gl/2tAaVo

* Install Ubuntu without CD and USB, how?
http://goo.gl/Pcoqt2

* How to write superscript characters in gedit?
http://goo.gl/4nCEGz

TIPS AND TECHNIQUES

GIVING HELP

I'm depressed. In order to produce this column, I spend a significant amount of time browsing the Ubuntu Forums, as well as other resources. I see a lot of questions which are incomplete, contradictory, unclear and ungrammatical -- even from people who are obviously English speakers. And that's OK.

However, at least a third of the "answers" demonstrate that the author has not read the question. One person had a question about Wi-Fi, and a forum moderator (!) asked him to provide additional information -- except that none of it was relevant to a wireless adapter. People who have written thousands of posts make stupid suggestions. I would love to ridicule them, but that would get me banned from the forums.

It's great to see so many people trying to be helpful. It would be even better if they would carefully read the question, so they don't write nonsense.

Gord had a long career in the computer industry, then retired for several years. More recently, he somehow found himself "The IT Guy" at a 15-person accounting firm in downtown Toronto.

Full Circle Podcast Episode 41, Trusted To Fail!!

Welcome to our new format show, there are several changes from the previous format, the most important being we are now recording together at the Blackpool Makerspace in the office. This Episode we Test Ubuntu 14.04, Review of Official Ubuntu Server Book.

Your hosts:
• Les Pounder
• Tony Hughes
• Oliver Clark

from the
Blackpool (UK) LUG
http://blackpool.lug.org.uk

Download

With the possible demise of Truecrypt, it seemed like a good time to cast around for a similar encryption application. After a little time, it became clear that Truecrypt is unusual in that it is, as far as I can see, the only cross platform application which is able to encrypt containers and drives or partitions. I have settled on using the following methods of encrypting which may look difficult but I promise are really quite simple. I am not very good at scripting, and am probably lazy, so this stuff will work for anyone.

Ubuntu is very good at encrypting drives and partitions with built-in encryption tools. ENCFS is useful for encrypting folders and files. AESCrypt makes encrypting a file for sending to someone else very easy, and Tomb creates encrypted containers as you could do easily with Truecrypt. The downside to all of these applications is that they are command-line tools. Please don't move to the next article just yet. They are quite simple to master – I am not a lover of the command-

line really, and managed to use all of them successfully in an afternoon.

I should also mention that not all of these are open source, but all are free to use.

AESCRYPT

AESCrypt requires you to go to www.aescrypt.com/download, scroll down to Linux, and download the 32-bit or 64-bit version you need. Although it is called AES Crypt – Gui, and the description says that it is a GUI Application, I have not been able to find one in Xubuntu. Anyway, download the file, and unzip it to your chosen folder. Make the file executable (right mouse click, select properties, and on the permissions tab check "Allow this file to run as a program" and click OK). Run the program either from the command-line or right mouse click, and select Execute. You will immediately be asked for a password. This is the password you would provide for a sudo command. Select your language, and if you wish to

proceed, select Yes. Click Next, accept the default folder selections if you wish (I did) by clicking Next, and then click Finished.

I have pretty much copied the commands and description from the relevant part of the manual on the website. Suppose you have a file called "passwords.txt" that you would like to encrypt using the password "apples". You would enter the following command:

```
aescrypt -e -p apples
passwords.txt
```

That's it! The program will create a file with the name "passwords.txt.aes". When you want to later decrypt the file "passwords.txt.aes", you would enter the following command:

```
aescrypt -d -p apples
passwords.txt.aes
```

What if you want to decrypt a file, but just want to have it displayed on the screen and not stored in a plain-text file? You can

do that using this syntax:

```
aescrypt -d -o -
passwords.txt.aes
```

You can do all sorts of whizzy things with piping and scripts but there are other tools which, it seems to me me, are better at encrypting multiple files. I should also mention that when you encrypt a file using AESCrypt, it does not delete the original but creates an encrypted file of the same name. If you would like more information then please do go to the above website where there is more information than you can shake a stick at.

ENCFS

ENCFS is in the Ubuntu repositories so "sudo apt-get install ENCFS" gets it installed in short order. This works differently in that it uses two folders. One folder is the source folder where the encrypted files are stored. The other folder is the mount point. The mount point provides the unencrypted view of the files and

SECURITY

folders in the source directory. Folder and file names in the source are encrypted, so, although the folder structure is visible, the types of files and their names are obscured. The best thing about ENCFS is that you do not have to create an encrypted container; so you are constrained only by the size of the hard drive your source folder is on. There are very good user instructions at http://www.howtoforge.com/encrypt-your-data-with-encfs-ubuntu-13.04. I could type all of that out again but it really is quite self explanatory. Once the folders are set up, all that you ever need to do is run the encfs command below. Oh yes and make sure you remember the key or you have locked yourself out of the data. I guess this is true for all of the tools mentioned here.

Simple example follows; change the file path names as required:

```
$ mkdir
/home/username/sourcefolder

$ mkdir
/home/username/mountpoint

$ encfs
/home/username/sourcefolder
/home/username/mountpoint
```

```
Volume key not found,
creating new encrypted
volume.

Password: [password entered
here]

Verify: [password entered
here]
```

The mount point folder may be used through your file manager in the normal way.

To lock down the mount point when you are done with it, the following command should be used;

```
$ fusermount -u
/home/username/mountpoint
```

TOMB

Tomb can be downloaded from www.dyne.org/software/tomb. Tomb is a bit funny and probably is closest to Truecrypt in that it creates an encrypted container into which you put files. It has a quirk to it in that you have to open a file-manager from the terminal as superuser. I guess because the commands are all run prefixed by "sudo". Other than that, it does work well.

To business. Once you have

downloaded the compressed file, you need to unzip the contents somewhere. I keep my applications like this in a folder called applications. Imaginative or what? Open a terminal window and cd to where you decompressed the folder Tomb-x.x.x. I have cribbed the following instructions from the Tomb website as I cannot simplify them any more than this. To create a 100MB tomb called "secret" do:

```
cd to the Tomb directory

sudo tomb dig -s 100
secret.tomb

sudo tomb forge
secret.tomb.key

sudo tomb lock secret.tomb -k
secret.tomb.key
```

To open it, do:

```
sudo tomb open secret.tomb -k
secret.tomb.key
```

and after you are done:

```
sudo tomb close
```

or, if you are in a hurry:

```
sudo tomb slam all
```

I have found that if you want to have secret.tomb in another

location you have to use the full file path. The same goes for the key. For example; If you are using Office365 (Outlook 2013), click on the three blue dots at bottom left of the window. Click on Folders. You will now see everything in the left-hand pane in a folder view. If you expand the folder which your Inbox was recovered into, you will see your old calendar and contacts there. Drag contacts from this contacts folder and drop them into your current contacts folder.

If you are still using Outlook 2010, the folders view icon is a little yellow folder at the bottom of the left hand window.

```
tomb open
/home/username/private/secret
.tomb -k
/home/username/private/secret
.tomb.key
```

There is much more on the website about Tomb which I recommend anyone interested to read:
www.dyne.org/software/tomb.

This next bit is not really anything to do with encryption. As I said at the outset, I am quite lazy, I do not remember long commands very well and am always looking to

SECURITY

```
color:desktop:blue:blue
color:title:blue:white
color:base:blue:white
menu:main:Main Menu:Main menu

     exec:open tomb:p:sudo tomb open /home/username/private/secret.tomb -k /home/username/private/secret.tomb.key
     exec:close open tomb:P:sudo tomb close
     nop
exit:E_xit
```

find easy ways to do things. So I found a really useful utility called *pdmenu*. You can use it to create a menu of commands in a terminal window. Scroll up and down the menu using the cursor keys and the Enter key to run a command. It is available in the repositories and:

```
sudo apt-get-install pdmenu
```

will have it for you in a minute or two. It generates a default pdmenurc file called /etc/pdmenu. Create a file in your favourite text editor called .pdmenurc, modify it to your taste, and save in your home directory. When you run pdmenu in a terminal window it will find this file first and use it. I will show you (top) what the file would look like if it contained the commands to open and close the tomb file above.

The first four lines set up how the window will look.

exec – tells pdmenu that it is going to run a command

: are spacers between parts of the instruction

open tomb is the name that will be displayed

p tells pdmenu to wait for the user to input something

the command is the last part of the instruction.

```
exec:_Tetris for
Terminals::/usr/games/tt

exec:_Adventure:pause:/usr/ga
mes/adventure
```

The two above commands will display Tetris for Terminals with no flag so /usr/games/tt will launch when it is selected. The underscore makes the following letter a hotkey. So pressing t will take you to that selection in the menu. Adventure has a pause flag where

you could be required to enter some information before being returned to the menu. The man pdmenu command will give you lots of useful information, and if you really want to get into it, there is quite a lot of information on the Internet.

I have found that pdmenu is versatile; it will run nested menus, and it removes the tedium of remembering long command syntax.

The foregoing is not an exhaustive list of encryption methods but the uncertainty surrounding Truecrypt made me revisit what was available. I am certain that my computer is not impregnable but I am also certain that I have made the data on my computer available to only a very few people who might have some specialised skills. If my computer were to be stolen or lost, although someone may use the hardware, I

can be relatively sure that they will not use the data stored on it.

Iain is a full time IT professional in Oxford for an international consultancy business. Most of his work is Windows related but he has managed to squeeze in an Ubuntu server, and has been a desktop user of *buntu for many years.

Elizabeth K. Joseph: Can you tell us a little about yourself?

Svetlana Belkin: I am Svetlana Belkin, an active Ubuntu Member since July 2013, and I gained my Membership on February 6, 2014. This month will mark my first year of working in the Ubuntu Community.

I am not a developer, I cannot code to save my life!

I am a biology major with a focus on Cellular and Molecular Biology who uses Ubuntu because it and the FOSS world match how I think.

EKJ: What inspired you to get involved in the Ubuntu community?

SB: An idea for a multi-player online game that is based on Mario Party but instead of mini-games, players use cards that are either attack, defense, or traps to get coins. The one with the most coins wins but everyone can keep the coins that they gained to shop for more cards and avatar items.

This was about one year ago, and I wanted to find someone who could help develop it. Since I am a woman, I joined Ubuntu Women to seek one out. But I quickly found out that it was a bad choice and I started to work on improving the Ubuntu Women Wiki to have it up-to-date. That's what led me into doing other things within the Ubuntu Community.

EKJ: What are your roles within the Ubuntu community and what plans do you have for the future?

SB: My main role within the Ubuntu Community is to help newcomers to find their place in the Community and to network with women (Ubuntu Women) and scientists (Ubuntu Scientists) alike to improve the FOSS world.

I also help the Ubuntu Documentation team to keep the Ubuntu Community Help Wiki up-to-date.

My future plans are to train new leaders within the Community so they know how to lead.

EKJ: Have you hit any barriers with getting involved and what can you recommend to newcomers?

SB: Newcomers need to remember that they do not need to be a developer to get involved - that's the barrier that I hit.

I would recommend to newcomers that they should not think that they need to be developers, and they should take these steps: they should start out small, join the team/project and its mailing-list, make sure to read all of the documentation for that project/team, and introduce themselves to the team via the mailing-lists. The best route – if they do not know what skills they have or what teams/projects to join – is to go to their Local Community and ask on the mailing list or their IRC channel.

EKJ: Is there anything you feel the Ubuntu project could improve on when it comes to new folks coming to the project?

SB: The main thing is the lack of Ubuntu Recruitment/Promo/Comms teams where the new folks can join and ask what teams/projects they can put their skills into. The other flavors have these teams but Ubuntu does not.

EKJ: What other things are you interested in outside of open source and Ubuntu?

I make art from time to time, and play my favorite and the only Multi-User Dungeon, Armageddon MUD.

The Humble Indie Bundle has been one of the biggest blessings for Linux gaming as well as a very crucial positive change catalyst in the Linux gaming revolution of the past three years. I have been a fan of the Humble Indie Bundle since its early humble beginnings (pun intended). On May 27 2014, the video game Symphony made its debut on Linux and Mac OS X as part of the Humble Indie Bundle. Symphony had originally been released for Microsoft Windows in 2012, and almost immediately became a big indie hit and received rave reviews from both critics and fans alike. The Humble Indie Bundle folks wasted no time and got to work to bring Symphony to Linux as well as Mac OS X.

Symphony is a shoot 'em up video game produced by Empty Clip Studios. In Symphony, you control a ship which must shoot other enemy ships sort of reminiscent of classic video games like Asteroids and Galaga. The main concept that sets Symphony apart from other shoot 'em up games is

that the enemy ships, as well as the layout of the levels, are based on patterns found in the game's music. In addition, you can add your entire music library to the video game, or, if you prefer, you can add only select songs from your library. As the game states, "your music is under attack" and it's up to you to protect it. A slow, classical song will generate a slow-paced level, whereas a fast upbeat dance tune will give you a really fast-paced level with more enemy ships that

you'll need to shoot down. The game concept is not entirely original, as there's already been other games that have used music in a similar fashion. Regardless, Symphony is a great game to play.

To play Symphony, you can go to symphonygame.com which is the official game website, or you can also get it from Desura and Steam. Symphony is small in size, it needs only about ½GB (500MB) of hard drive space, which makes the

downloading and installation a breeze. You can control your ship with your mouse, and supposedly with a controller, even though my controller hasn't been able to work with this game.

The game is fun to play, but, because it's music driven, and especially for having a name like Symphony, I would expect music to play a leading role in this game. Contrary to what the name implies, you can actually turn the volume all the way down and probably get the same scores you get if you're actually listening to the music and moving according to what you hear. A music driven game should be easier to play if you're paying attention to the music, which is something that this game fails to do. Besides the music playing a rather minor role in this game, there were some graphics problems I encountered which seem to be related to the proprietary AMD driver that Steam requires me to use in order to play games. At random times, the entire screen, or part of it, would turn white even though the game would

contents ^

go on as if nothing wrong had happened. I researched my problem and found it to be actually quite common among gamers who use AMD graphics cards with proprietary drivers. I hope this is addressed in the future and the game becomes more enjoyable to play.

Minimum System Requirements
OS: Major Linux distributions from 2012
Processor: Dual core, 2 GHz or faster
Memory: 2 GB RAM
Graphics: OpenGL 2.1 (with shader model 3 video card or higher)
Not recommended for any integrated graphics GPUs
Hard drive: 500 MB available space

CONCLUSION

Here's my Pro's and Con's for the game Symphony as well as my rating for this game.

Pro's
• Interesting concept of using your own music as the princess that needs to be rescued.
• Very easy to play, yet it doesn't get boring quickly. I could keep playing the game again and again.

Con's
• The way the music is used, unfortunately, doesn't play a more prominent role, and in fact the game could very well be played with the volume off.
• Glitches with graphics on computers with AMD graphics cards.
• Although it claims to have controller support, my two controllers did not work with this game.

Unfortunately, I have to give this game a mere 2½ out of 5 stars. The con's just flat outweigh the pro's; especially the glitch with AMD graphics cards is a deal buster (if you happen to have such kind of graphics card).

★★⯪☆☆

Oscar graduated from CSUN, is a Music Director/Teacher, beta tester, Wikipedia editor, and Ubuntu Forums contributor. You can contact him via: www.gplus.to/7bluehand or email: www.7bluehand@gmail.com

Ave Maria (Schubert)
Gisella
Ave Maria

POWER UP!

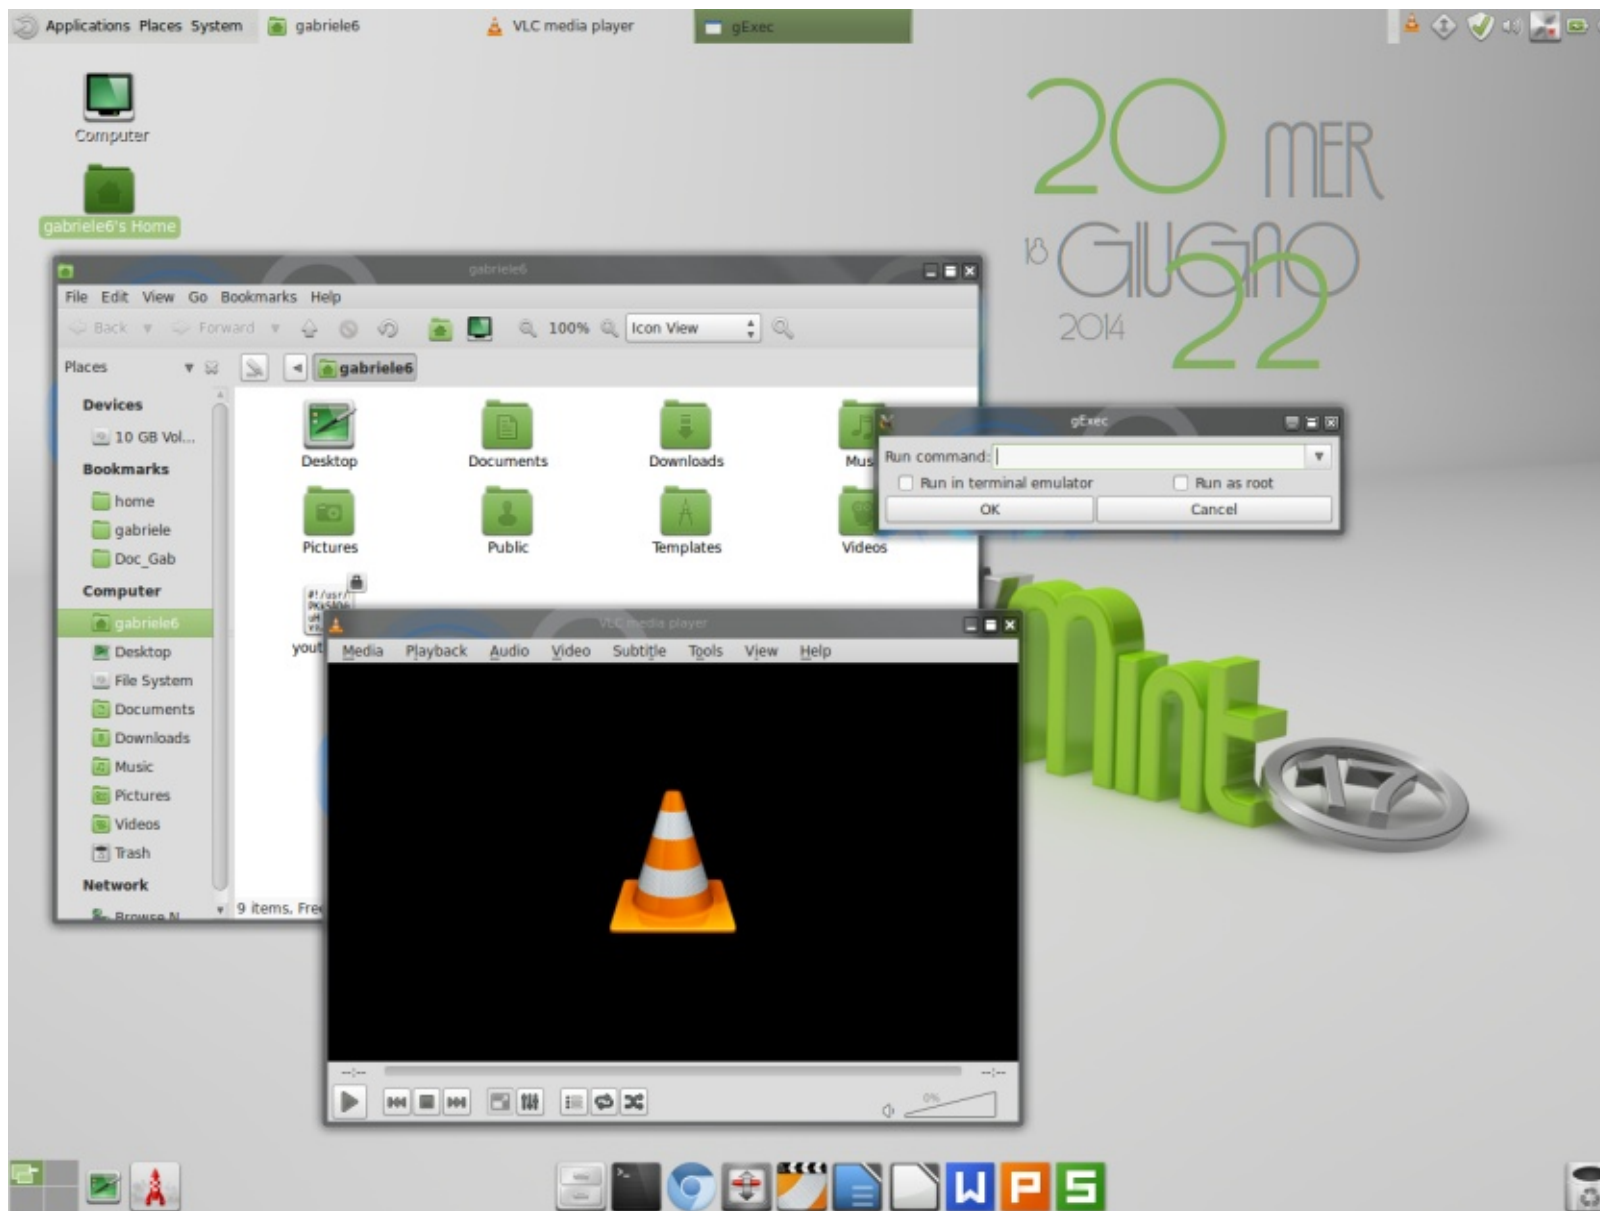♪ 68,997
ALL TARGETS REACHED

00:06

MY DESKTOP

Your chance to show the world your desktop or PC. Email your screenshots and photos to: misc@fullcirclemagazine.org and include a brief paragraph about your desktop, your PC's specs and any other interesting tidbits about your setup.

The system here features an Acer 1280 X 1024 75Hz monitor and Gateway/Acer SX-2800 micro desktop with 2.33 GHz Intel Core 2 Quad processor, 4 GB RAM and WDC WD640 drive. The OS is Linux Mint 17 (LTS) with MATE 1.8.0 desktop. A pretty vanilla system these days, but it's more than I need. Note my frequently used apps appear on the left-hand panel which is normally hidden. I found the uncredited image on the web somewhere. It is highly appropriate for a Linux system isn't it?

The Ancestral Quest and G4FON shortcuts launch the respective Windows programs on Wine 1.6.2 from the MINT repo. Note also that MINT 17 has an app to support Dropbox, which I use to copy pictures and books to and from my iPad.

Dave Rowell

few seconds more while booting, however this way it's very easy turning it off.

Here are the details of appearance settings: the icon theme is Mate Faenza Dark, available in the official repos, the theme is Menta, one out of the Mate official themes package, also stored in Ubuntu repos, and the window decoration is glowglass, a KWin theme decoration from the web.

The elegant time widget is a customized Frozen Cherry Desktop widget I've found in one of the bundles I installed; by the way, in order to keep the cpu load low, I raised the Conky update interval from 1 to 10 seconds.

My PC specs: Lenovo T43p, 2.13GHz Pentium M, 2GB RAM, 80GB HD, Ati Fire GL3200.

Gabriele Tettamanzi

This is my Linux Mint 17 KMate: I've pimped up Mate using the KWin window manager from KDE project and Conky.

I've installed kde-window-manager and systemsettings, the latter because I need to tweak the style fine tuning in Application

appearance to low cpu-high resolution mode. Kwin is launched by the command kwin –replace, I set it as startup application: it takes

HOW TO CONTRIBUTE

FULL CIRCLE NEEDS YOU!

A magazine isn't a magazine without articles and Full Circle is no exception. We need your opinions, desktops, stories, how-to's, reviews, and anything else you want to tell your fellow *buntu users. Send your articles to: articles@fullcirclemagazine.org

We are always looking for new articles to include in Full Circle. For help and advice please see the **Official Full Circle Style Guide**: http://url.fullcirclemagazine.org/75d471

Send your **comments** or Linux experiences to: letters@fullcirclemagazine.org
Hardware/software **reviews** should be sent to: reviews@fullcirclemagazine.org
Questions for Q&A should go to: questions@fullcirclemagazine.org
Desktop screens should be emailed to: misc@fullcirclemagazine.org
... or you can visit our **forum** via: fullcirclemagazine.org

FCM#88
Deadline:
Sunday 10th Aug. 2014.
Release:
Friday 29th Aug. 2014.

Getting Full Circle Magazine:

EPUB Format - Recent editions of Full Circle have a link to the epub file on the downloads page. If you have any problems with the epub file, you can drop an email to: mobile@fullcirclemagazine.org

Issuu - You can read Full Circle online via Issuu: http://issuu.com/fullcirclemagazine. Please share and rate FCM as it helps to spread the word about FCM and Ubuntu Linux.