



Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #80 - December 2013



Photo: comedy_nose (Flickr.com)

MONEY MANAGEMENT GNUCASH, GRISBI & KMYMONEY

Full Circle Magazine is neither affiliated with, nor endorsed by, Canonical Ltd.



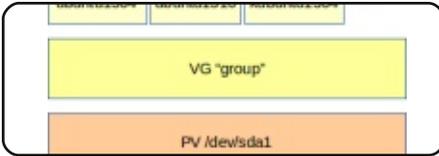
HowTo



Python p.08



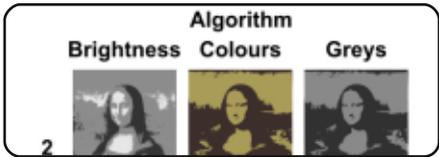
LibreOffice p.10



Using LVM p.12



JPG to PDF p.19



Inkscape p.20



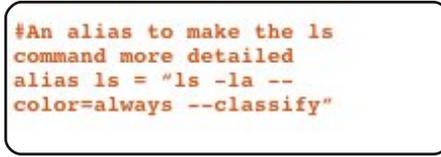
Graphics



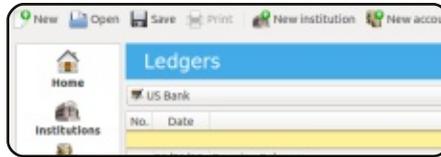
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

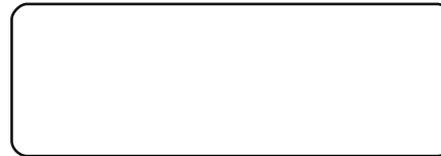
Columns



Command & Conquer p.06



Ask The New Guy p.26



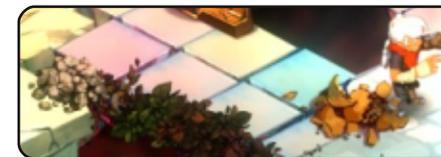
Linux Labs p.30



My Desktop p.44



Ubuntu News p.04



Ubuntu Games p.42

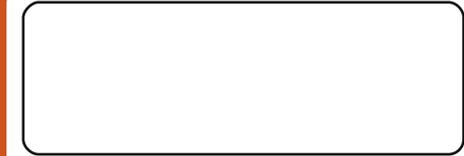


Q&A p.40



Ubuntu Women p.45

Opinions



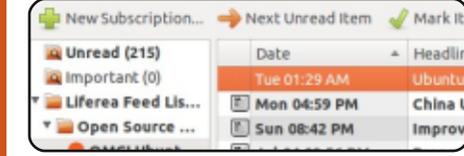
My Story p.32



My Opinion p.XX



Review p.33



Software Showdown p.34



Letters p.38



The articles contained in this magazine are released under the Creative Commons Attribution-Share Alike 3.0 Unported license. This means you can adapt, copy, distribute and transmit the articles but only under the following conditions: you must attribute the work to the original author in some way (at least a name, email or URL) and to this magazine by name ('Full Circle Magazine') and the URL www.fullcirclemagazine.org (but not attribute the article(s) in any way that suggests that they endorse you or your use of the work). If you alter, transform, or build upon this work, you must distribute the resulting work under the same, similar or a compatible license.

Full Circle magazine is entirely independent of Canonical, the sponsor of the Ubuntu projects, and the views and opinions in the magazine should in no way be assumed to have Canonical endorsement.



WELCOME TO A PROPER ISSUE OF FULL CIRCLE!

I have to start this month with an apology about the state of the previous issue (FCM#79). To cut a long story short: I ran a dd command without properly checking drive letters, and completely annihilated my 1TB drive partitions and lost all my data. Including, of course, the Scribus file for FCM#79. Needless to say, I've written up an article on how to backup your data and a brief article on data recovery. Take it from me, don't sit there and think (like I did) that it'll never happen to you. It can!

Anyway, enough doom mongering. We have a full house of articles for you again. If you're a Blender fan, don't panic, Nicholas will be back next month. With Python, and LibreOffice, and there's a great article on using the Logical Volume Manager. I'm not even going to pretend I know much about it. Give it a read and see what you think.

Needing a review, I wrote up a quick article about my USB Microscope that I bought a couple of years ago. It's great fun and recently came in handy when I was cataloging a bundle of old coins I have. It may seem like a cheap USB gimmick (I always get asked about it when I post images from it on Google+), but this is one instance where it really did come in handy.

For you gamers, we have a review by Oscar of Bastion. At the moment, I'm trying out Starbound which is fun. It looks (and plays) very like Terraria, but has its own little quirks. It's still in beta, but I may review it next month.

All the best, and keep in touch!

Ronnie

ronnie@fullcirclemagazine.org



This magazine was created using :



Full Circle Podcast

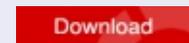
Released monthly, each episode covers all the latest Ubuntu news, opinions, reviews, interviews and listener feedback. The Side-Pod is a new addition, it's an extra (irregular) short-form podcast which is intended to be a branch of the main podcast. It's somewhere to put all the general technology and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

Hosts:

- Les Pounder
- Tony Hughes
- Jon Chamberlain
- Oliver Clark



<http://fullcirclemagazine.org>





NEW DDoS BOTNET INFECTS WINDOWS ALONG WITH LINUX COMPUTERS; POLAND CERT

According to Poland Computer Emergency Response Team (CERT), its security researchers recently discovered one fresh DDoS (distributed denial-of-service) botnet that spread a kind of malware which contaminated Linux as well as Windows computers, published softpedia.com dated December 18, 2013.

The researchers state that the botnet has been created for executing solely DDoS assaults, the majority being DNS Amplification assaults. At the time the report by Poland CERT was released, nearly all anti-virus programs had detected the malware targeting Windows computers, whereas, merely a few AV engines could detect the malware for Linux. What's more, the Linux malware attempted at linking up with the command and control server

through one high-powered Transmission Control Protocol (TCP) port.

The researchers explained that the port and IP of the C&C server were encrypted. When the bot became active, it transmitted information about operating systems unencrypted while looked forward to receiving commands from the server. Help Net Security published this dated December 18, 2013.

Upon studying the malicious program, it was concluded that the threat could execute 4 kinds of DDoS assaults. Besides, it had other capabilities which hadn't yet been carried out.

It was further found that while the Linux variant of the bot communicated with the C&C system through an Internet Protocol address, the Windows version communicated through one domain name as also over another port. However, both variants of the threat used the same command-and-control infrastructure,

suggesting that an identical group of criminals were behind both, as per analysis by Polish CERT.

And because the malware executed just DDoS assaults the researchers think the attackers possibly focused on compromising PCs using lot of network bandwidth such as servers, which they had. That possibly was also responsible for the two malware variants -Linux OS was the preferred one for sever systems.

Meanwhile, similar as the malicious program which was detected recently targeting Linux, another bot lately came in the notice of Andre DiMino, a George Washington University security researcher who had let attackers hijack his honeypot Linux-based computers. The bot was developed in Perl language, DiMino discovered.

Source:

<http://www.spamfighter.com/New-s-18732-New-DDoS-Botnet-Infests-Windows-along-with-Linux-Computers-Poland-CERT.htm>

MOVING A CITY TO LINUX NEEDS POLITICAL BACKING

This year saw the completion of the city of Munich's switch to Linux, a move that began about ten years ago. "One of the biggest lessons learned was that you can't do such a project without continued political backing," said Peter Hofmann, the leader of the LiMux project, summing up the experience.

The Munich city authority migrated around 14,800 of the 15,000 or so PCs on its network to LiMux, its own Linux distribution based on Ubuntu, exceeding its initial goal of migrating 12,000 desktops.

Munich decided to migrate its IT systems when Microsoft said it planned to discontinue support for the operating system the city then relied on, Windows NT 4.0. The city was forced to choose between moving to a newer version of Windows, or finding an alternative platform, as new software and new

versions of existing software would not be available on Windows NT. The city council decided to go with Linux to become more independent from software vendors.

Continued political backing was key to the success of the migration, said Hofmann.

"We had it from the start and it never failed. We had to treat our politicians as stakeholders and keep them informed," he said.

By doing this, the politicians never lost interest and always knew what the people involved in the project were doing, he said. *"I saw a lot of other open source projects going down the sink,"* because they didn't have that backing, or lost it, he said.

Source:
<http://www.pcworld.com/article/2082460/moving-a-city-to-linux-needs-political-backing-says-munich-project-leader.html>

VALVE JOINS THE LINUX FOUNDATION

Valve is joining the Linux Foundation in the run up to the gaming giant's launch of its Steam Machine, a console based on the Linux open source operating system which the foundation exists to promote.

In September, Valve, the developer of the Half Life and Portal series of games as well as the Steam online gaming platform, announced it would be using Linux as the basis of its Steam Machine line of games consoles. The system takes a radically different approach to TV gaming from that advocated by the Xbox One and PS4, allowing any manufacturer to build a compatible system and install the Linux-based Steam OS on top of it.

While the code which underpins Linux is free for anyone to use, the Linux Foundation controls the trademark itself, and co-ordinates development of the operating system between the numerous companies which contribute programming time.

By joining the foundation, Valve will be able to feed the improvements it makes to Linux for the Steam Machine back into the

main body of code, and also more directly benefit from the network of developers which work on Linux itself. *"Through these efforts we hope to contribute tools for developers building new experiences on Linux, compel hardware manufacturers to prioritize support for Linux, and ultimately deliver an elegant and open platform for Linux users,"* says Mike Sartain, a key member of the Linux team at Valve.

Source:
<http://www.theguardian.com/technology/2013/dec/04/valve-joins-linux-steam-machine-gaming>

LINUX DEVS MAY 'CRY INTO OUR LONELY BEERS' AT CHRISTMAS

Linus Torvalds has let release candidate five for version 3.13 of the Linux kernel into the wild for some festive footling.

The Linux Lord let the new release candidate loose in this post that declares "Nothing really exciting stands out" which is "just how I want it."

"It's the 'how did that ever even pass cursory testing' bugs that make me upset," he continues, "and if those existed, people were appropriately ashamed and quiet about them ;)".

The post also outlines his plan for at least three more release candidates in the 3.13 development process, suggesting that "LCA" (probably a reference to January 6th's Linux.conf.au in Perth, Australia) will keep Torvalds on the road for a while and slow development.

Torvalds also offers Christmas greetings in his own unique tone, to wit:

Ho ho ho, Christmas is almost upon us, and -rc5 is the last rc before most of us gorge ourselves into insensibility. Or cry into our lonely beers. Or go out for Chinese food. Or whatever you happen to do.

Source:
http://www.theregister.co.uk/2013/12/24/torvalds_linux_devs_may_cry_into_our_lonely_beers_at_christmas/



Early in December I was spending some time working on a website that will eventually become my website for my freelance IT business. I started off using Bolt CMS and ended up with a design I quite liked, and overall, Bolt did everything I wanted. However, it did a lot more than I needed as well, and I have slight reservations for using a blog-oriented CMS to manage the content of a web page. Instead, I decided to give Ruby on Rails a shot, to see which option I would ultimately go with. I had previously spent a while learning both Ruby and Ruby on Rails, so jumping in wasn't quite as daunting. This month's article is for anyone hesitating about diving in, and for anyone interested in web design who would like to pick up some new skills.

PREREQUISITES

The ruby package must be installed before using Rails. The current one in Ubuntu is version 1.9.2 (package: ruby1.9.1). The

current version is actually 2.0.0-p353, so it may be better to install ruby-rvm (ruby version manager), and use that to install/manage your ruby versions. For anyone who needs to install it in other operating systems too, there are good instructions here:

<https://www.ruby-lang.org/en/downloads/>

Rails installation instructions:
<http://rubyonrails.org/download>

RUBY ON RAILS – WHAT IS IT?

Ruby on Rails is a framework designed in Ruby to make the generation of complex websites much easier. It's designed with the idea in mind that a programmer can focus on convention over configuration, while still designing the website they originally had in mind.

Ruby on Rails – Why should I try it?

- By default it uses SCSS for its stylesheets, allowing for cleaner

formatting in your stylesheets (it's automatically parsed to css when running the server).

- Integrates a server setup, which can be started with "rails server" or "rails s", and polls for changes to the files, while serving up your website on localhost:3000 (port 3000, by default). Naturally, this can also be used in a production environment with an actual domain.

- Allows for easy extensibility by editing the Gemfile.

- Has modules for certain tasks (i.e. User authentication) which you can easily add to your own project, and helps to avoid feature overload.

- You can create your own classes and controllers, so that only items you actually use are saved in your database.

- Error handling – This is a bit subjective, but I find that Rails' approach to throwing errors and debug information is a lot more helpful/cleaner than most other web-based languages.

- Large community – odds are you'll be able to find a solution to most of your problems online, and some

areas may even have Rails/Ruby meet-ups.

- Integrates with Git very well, allowing for clean revision control of your website's code. It can also be used to upload your site to hosts like Heroku, without the need to use FTP.

- Includes a routing table, allowing you immense control over the URL formatting and restrictions.

Possible hardships:

- Very programming-focused – if this isn't your forte and you're not interested in learning, you may very well struggle.

- Requires a basic knowledge of Ruby – since Rails uses the Ruby syntax for controllers, models and helpers, it's highly recommended you have a working knowledge of it. However, something like <http://tryruby.org/> could be sufficient, if you're willing to refine your knowledge as you work in Rails.

- Unless you follow a tutorial for your first project, you may run into issues where you're unsure what to Google, due to the errors being specific to your custom models.

COMMAND & CONQUER

- The rails generate command can be slightly confusing at first, especially when it results in rake db:migrate errors. However, once you work your way through one error, you should have a good understanding of how to fix future ones.
- If you don't know about SCSS, you will lose out on the benefit of it, but it won't impact your CSS generation.

YOU'VE SOLD ME! WHERE DO I START?

The first step I'd recommend is <http://tryruby.org>, in order to learn/freshen up on Ruby. Once you've completed that, Code School also offers a free course in Rails, called Rails for Zombies (<http://railsforzombies.org>). While it may seem silly at points, it is actually a great resource for an introduction to Rails. There are also countless books that cover Ruby on Rails (admittedly, I haven't read any, and as such don't feel I can make a recommendation).

Once you've read up on Rails and have done a tutorial of some sort, be sure to pick out a specific thing you'd like to achieve (i.e.

Design a class/controller to handle pages and their content). Then look into the rails generate scaffold command. My only tip here is to make sure you integrate all the variables you'll need from the get-go in the scaffold command. An example command could be:

```
rails generate scaffold Page
title:string content:string
slug:string
```

This will generate the controller, model and views required for Page, and configure it to have a title, content, and slug attribute in the database, which you can input through the edit page. Since this article isn't intended to be a detailed how-to, I will leave you with just two more hints:

- In the above example, after running the server, you can see your pages by going to <http://localhost:3000/pages>
- Once you've used the scaffold command and have the files, it works best if you then generate specific controllers or models as you need them. To see a full list, type rails generate -h

IT STILL SEEMS KIND OF DAUNTING - CAN I REALLY DO THIS?

If you want a success story, the one that first jumps to mind is the creator of Freelancify, who learned Ruby on Rails in 12 weeks (specifically, from the time he began learning Ruby on Rails and the time he put Freelancify online was 12 weeks). According to his blog post, he wasn't a programmer at all before jumping into Ruby on Rails.

Source:

<http://www.jamesfend.com/learned-ruby-rails-12-weeks-launch-freelancify>

WILL YOU BE DOING A TUTORIAL ON HOW TO SET UP RAILS?

While I find the Ruby on Rails documentation is excellent, and there are some guides online as well, I can understand if some people prefer an up-to-date step-by-step guide from me. If that's the case, let me know in an email at lswest34+fc@gmail.com and I will happily cover setting up Rails 4.0.2

and getting a project started.

As always, I hope you've found this article interesting, and that I have enticed one or two of you into trying Ruby on Rails. If you have any comments, questions, suggestions or requests, you're welcome to contact me at lswest34+fc@gmail.com. I'd also be happy to hear about any experiences with Ruby or Ruby on Rails you've had.



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.



This month, I thought I'd talk about a couple of lesser known functions, maketrans and translate.

We'll start with the translate method. The translate method returns a copy of a string – with all characters in the translate table replaced, or has the characters in the optional parameter deletechars removed from the string. Here's the syntax.

```
s =
str.translate(table[, deletecharacters])
```

Before we get to the table portion of the method, let's look at the delete portion. Let's say that you have the string "The time has come". And you want to delete all the vowels (for some weird reason) from that string. You can code it like this:

```
astr = "The time has come"
astr.translate(None, 'aeiou')
```

will return:

```
"Th tm hs cm"
```

Notice that we included None as the translate table. While this part is cool, it gets better. There is a function called maketrans. It takes an input string and an output string as parameters and returns a table that is used as the first parameter into the translate method. Here (top right) is a very simple example.

It returns:

```
"Th2 t3m2 h1s c4m2"
```

Let's look at what this does. We assign intable to a string of vowels as before. outtable is assigned the numbers 1,2,3,4,5 as a string. When we make the call to maketrans, our actual trantable is as follows (shown below. The "\x" means that it is hexadecimal char):

```
\x54\x68\x65\x20\x74\x69\x6d\x65\x20\x68\x61\x73\x20\x63\x6f\x6d\x65
T h e t i m e h a s c o m e
```

```
'\x00\x01\x02\x03\x04\x05\x06\x07\x08\t\n\x0b\x0c\r\x0e\x0f\x10\x11\x12\x13\x14\x15\x16\x17\x18\x19\x1a\x1b\x1c\x1d\x1e\x1f !"#%&\'()*+,-./0123456789:;<=>?@ABC
DEFGHIJKLMNOPQRSTUVWXYZ[\ ]^_`1bcd2fgh3jklmn4pqrst5vwxyz{|}~\x7f\x80\x81\x82\x83\x84\x85\x86\x87\x88\x89\x8a\x8b\x8c\x8d\x8e\x8f\x90\x91\x92\x93\x94\x95\x96\x97\x98\x99\x9a\x9b\x9c\x9d\x9e\x9f\xa0\xa1\xa2\xa3\xa4\xa5\xa6\xa7\xa8\xa9\xaa\xab\xac\xad\xae\xaf\xb0\xb1\xb2\xb3\xb4\xb5\xb6\xb7\xb8\xb9\xba\xbb\xbc\xbd\xbe\xbf\xca\xcb\xcc\xcd\xce\xcf\xda\xdb\xdc\xdd\xde\xdf\xe0\xe1\xe2\xe3\xe4\xe5\xe6\xe7\xe8\xe9\xea\xeb\xec\xed\xee\xef\xf0\xf1\xf2\xf3\xf4\xf5\xf6\xf7\xf8\xf9\xfa\xfb\xfc\xfd\xfe\xff'
```

```
intable = 'aeiou'
outtable = '12345'
trantable = maketrans(intable, outtable)
astr = "The time has come"
astr.translate(trantable)
```

If you look at it carefully, you'll see that the lowercase vowel letters are replaced with the numbers we specified:

```
1bcd2fgh3jklmn4pqrst5vwxyz
```

If you look even closer, you'll see that there actually 256 entries starting with "\x00" and ending with "\xff". So the table contains the entire 256 possible ascii character set. So, when the translate method gets the table, it iterates (or walks through) each

character, getting that characters value in Hex, and then finds that value in the translate table and substitutes it in the output string. The Hex representation of our original astr string ('The time has come') is shown below.

So now it should be making sense.

Now the purpose of this whole thing. Think back to your schooling where you learned about Julius Ceasar. Whenever he wanted to send a message of a confidential

matter, he would use a cipher that would shift all the letters of the alphabet three characters to the right. So, using today's English alphabet:

```
ABCDEFGHIJKLMNOPQRSTUVWXYZabc  
defghijklmnopqrstuvwxyz
```

becomes:

```
DEFGHIJKLMNOPQRSTUVWXYZabcde  
fghijklmnopqrstuvwxyzABC
```

While this seems very simple by today's standards, when I was a school kid, we used this all the time to send messages to each other. We used a different index into the string to start the encryption string, the logic behind it was the same.

No one knows how effective this actually was for good old Julius. One would think that if someone intercepted the message, they would have thought that it was in some foreign language. We can only speculate.

We can easily use the translate method and the maketrans helper function to allow us to have fun with this. Let's say we want to make a simple program that allows us to enter a string of "plain text" and get back an encrypted string using the same side right method

that Caesar used. For simplicity sake, let's only use uppercase characters (shown top right).

Everything in the above code is pretty much what we've covered above or in earlier Python articles, but I'll go over it quickly.

The first two lines are the in and out strings. We've just shifted the characters and wrapped around to create the out string. The next two lines create a table for encoding and one for decoding. Line 5 prompts the user to enter a string to encode. We then encode that string (EncString) in the next line. To decode it, we simply use the translate method on the encoded string to get the plain text back. Finally we print both strings out. Here's the output of the program.

```
Enter the plaintext string ->  
THE TIME HAS COME  
Encoded string is -  
WKH WLPK KDV FRPH  
Decoded string is -  
THE TIME HAS COME
```

Just like back in school. But let's flesh it out just a bit to make it a bit more usable. The code is almost the same with a few exceptions. First, we have added a space to the end of the intab string and in

```
from string import maketrans  
#-----  
intab = "ABCDEFGHIJKLMNOPQRSTUVWXYZ"  
outtab = "DEFGHIJKLMNOPQRSTUVWXYZABC"  
EncTrantab = maketrans(intab, outtab) #Encode  
DecTrantab = maketrans(outtab, intab) #Decode  
instr = raw_input("Enter the plaintext string -> ")  
EncString = instr.translate(EncTrantab)  
DecString = EncString.translate(DecTrantab)  
print("Encoded string is - %s" % EncString)  
print("Decoded string is - %s" % DecString)
```

between the "Z" and the "A" in the outtab string. This helps keep the actual words from being too obvious in the encrypted string. The next change is where we ask if the user wants to encode or decode the string. Finally we added an if statement to control what we print (shown bottom right).

The output from the program is:

```
Encode or Decode (E or D) -> E  
Enter the string -> THE TIME HAS
```

```
COME  
Encoded string is -  
WKHCWLPKCKDVCFRPH
```

And to test the decode side of things:

```
Encode or Decode (E or D) -> D  
Enter the string ->  
WKHCWLPKCKDVCFRPH  
Decoded string is - THE TIME HAS  
COME
```

Well, hopefully you are starting to get ideas about how to use this new information in your own code.

```
from string import maketrans  
  
#Be sure to include the space character in the strings  
intab = "ABCDEFGHIJKLMNOPQRSTUVWXYZ "  
outtab = "DEFGHIJKLMNOPQRSTUVWXYZ ABC"  
EncTrantab = maketrans(intab, outtab) #Encode  
DecTrantab = maketrans(outtab, intab) #Decode  
  
which = raw_input("Encode or Decode (E or D) -> ")  
instr = raw_input("Enter the string -> ")  
EncString = instr.translate(EncTrantab)  
DecString = instr.translate(DecTrantab)  
  
if which == "E":  
    print("Encoded string is - %s" % EncString)  
else:  
    print("Decoded string is - %s" % DecString)
```



Using the default HSQLDB database management system (DMS) in LibreOffice Base is great for creating new databases, but sometimes you need to connect to an existing database on another DMS like MySQL, dBase, JDBC, Oracle, PostgreSQL, and others. I showed you how to connect to a simple spreadsheet database back in Issue 52. Today, I will show you how to connect to a database created in MySQL.

MySQL is a DMS used on many websites. Wordpress and other content management systems use MySQL for their databases. You might have many reasons to access a database created in MySQL. MySQL is open source, and you can install it from your software manager.

INSTALL MYSQL

Installing MySQL from your software manager requires the installation of two packages. The first package is "mysql-client" – the client program for MySQL. It allows you to interact with MySQL.

If you are installing MySQL on your current machine, you will need to install the MySQL server package as well: "mysql-server". When you install the server package, it prompts you for a "root" password and asks for confirmation. This is your master password for MySQL. Jot it down somewhere. You will need it later. If you are connecting to a database on a different machine, you don't need this package.

The last package we need to install is "libreoffice-mysql-connector". This package will allow us to connect to MySQL without using an ODBC or JDBC wrapper.

CREATE A TEST MYSQL DATABASE

If this is the first time you are

connecting to an external database using LibreOffice Base, I recommend you create a test database in MySQL where you can make mistakes and learn. You can install a program like PHP MyAdmin or EMMA, but I think it is easier just to use the MySQL command-line client.

To connect to MySQL on your local machine, use the command:

```
mysql -h localhost -u root -p
```

MySQL will prompt you for the "root" password. Enter it and you are greeted with a welcome message and the "mysql>" prompt. To connect to a MySQL server on another computer, replace "localhost" with the server name or IP address. Keep in mind that you will need to set up MySQL to accept remote connections. For

details on setting up MySQL to accept remote connections, read this tutorial:

<http://www.cyberciti.biz/tips/how-do-i-enable-remote-access-to-mysql-database-server.html>

To create a database use the following command at the MySQL prompt:

```
CREATE DATABASE <dbname>;
```

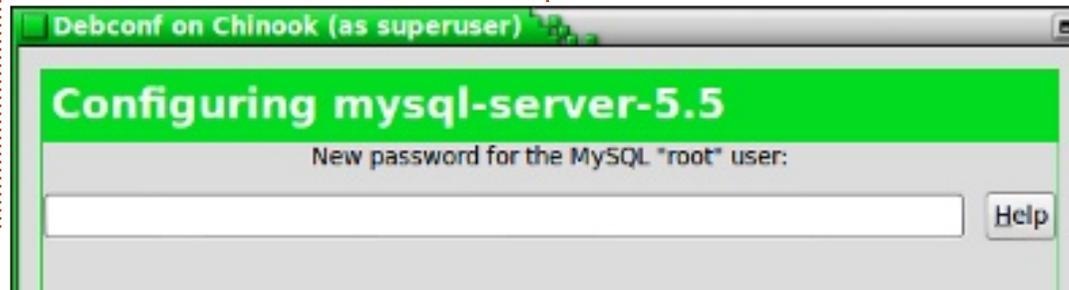
where <dbname> is the name you want to give the database. To create a databases name "myTestDB" enter:

```
CREATE DATABASE myTestDB;
```

MySQL responds with a message that the query was executed. Enter

```
SHOW DATABASES;
```

to generate a list of all the databases. You should see your new database among those listed. An empty database is all we need to connect to MySQL from Base, but you can also connect to a fully functional database as well; but,

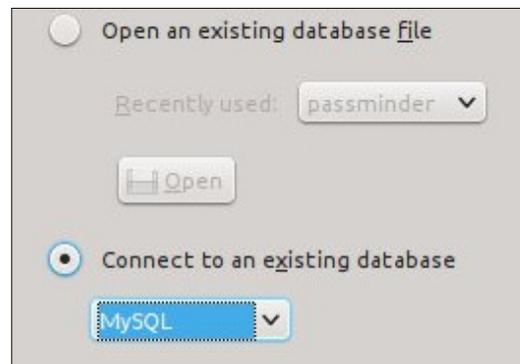


like I said, it is better to practice on a test database first. Type "quit" and press Enter to exit your MySQL session.

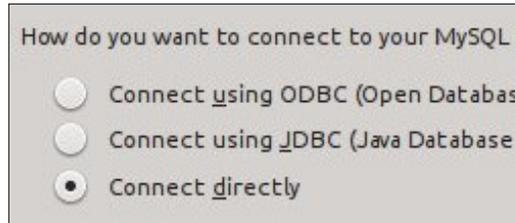
CONNECT BASE TO MYSQL DATABASE

Now for the part we have been working for – actually connecting Base to the MySQL database. Start by creating a new database document in LibreOffice, File > New > Database. We will use this document to connect to the database we created before. Using the database wizard (recommended), it is a five step process.

Step 1: We want to connect to an existing database. Don't confuse this with connecting to an existing database file. You want the last selection. From the drop down list select MySQL. Click Next.



Step 2: Since we installed the "libreoffice-mysql-connector", we have no use for the ODBC or JDBC connections. Select "Connect directly". Click Next.

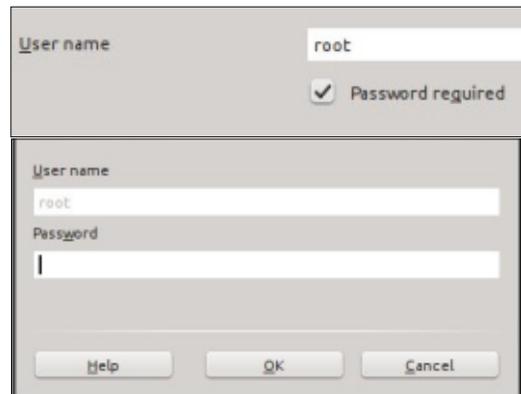


Step 3: The database is the name of the database we just created in MySQL. Our example database was "myTestDB". The server is "localhost" for the local computer, or the server name or IP address of another computer. Leave the Port as the default unless you changed it in your MySQL settings or have been told to do so by your IT department. Click Next.

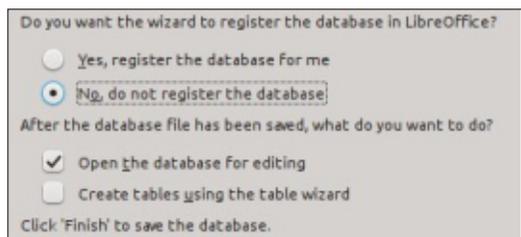


Step 4: The user name you need to connect to MySQL from your

computer. Use "root" if you are connecting to MySQL on your local computer. Check "Password required". Click "Test Connection". You are prompted for your password. Enter your password and click OK. If you set everything correctly, you will get a message back saying Base successfully connected to the database. OK on the message and click Next.



Step 5: Select Register or Don't Register according to your needs for the database. A registered database is available in other LibreOffice modules. Check "Open database for editing". Click Finish. Base will prompt you for a file name. Enter a name and click Save.



You can now interact with the MySQL database in the same way you would an HSQLDB database. You can create tables, queries, reports, and forms just like you would with the default DMS. You will need to pay attention to the data types as they are slightly different in MySQL. For an overview of the HSQLDB data types see my article at <http://wp.me/pvwht-8Q>. For a list of MySQL data types see <http://dev.mysql.com/doc/refman/5.0/en/data-type-overview.html>.

While using the default DMS for Base is good for new databases, sometimes it becomes necessary to connect to existing databases. In this article, we showed how this is done by connecting to an external database we created in MySQL. We installed MySQL and created a blank database. Using the database wizard in Base, we connected to that database where we can create tables, forms, queries, and reports to interact with the database. Base can also connect to other database types like dBase, JDBC, Oracle, PostgreSQL, and others.



Though I have been using and studying GNU/Linux for quite some time now, LVM is a new beast in my menagerie. Perhaps other aspects of system administration seemed more interesting, or maybe I didn't feel the need for LVM. But now that I know more about it, it is fast becoming an essential part of my toolkit.

If we take a look at back issues of our favorite magazine (Full Circle Mag, what other?), we actually find that LVM is mentioned very few times. It did appear in the 2011 "Special Edition 01: the Perfect Server", and it was also mentioned briefly in issues 31 and 67. But there have been no in-depth articles on the matter.

For today, after giving a simple example to put all the pieces in place, we will look at three different scenarios, covering some real-world, hands-on applications for LVM. In the first, we will use LVM to give us some flexibility when installing a new system. This will be a system for which we do not really know from day one

exactly what disk space requirements will arise for our system and user document partitions. Should we retain a Windows partition, or will we want to remove it later on and combine this space with our Ubuntu partitions? In the second scenario, we will see how the LVM snapshot capability may be a life-saver if we should want to upgrade an existing system, without risking losing it all. Finally, in the third case we will be creating a bootable USB drive with several Ubuntu (or other) distributions all at once, so we can boot a computer from any of them.

WHAT LVM IS

Normal disk management under GNU/Linux - and also under Ubuntu - consists of working with disk partitions on the one hand, and file-systems on the other. Disk partitions are the low-level divisions of our disks. If you start out with a disk detected by the system as `/dev/sda`, it may contain partitions `/dev/sda1`, `/dev/sda2`, etc. Each of these partitions can then be formatted using a file-

system: extended for GNU/Linux, NTFS or VFAT for Windows, HFS for Mac-OS-X, etc. The limitation with this scheme is that each partition can support one single file-system.

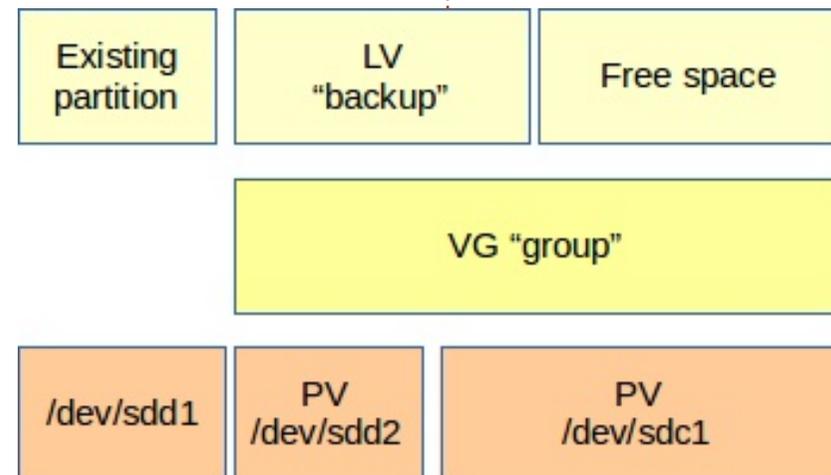
The main idea behind the Logical Volume Manager (LVM) is to insert a third layer between the physical partitions on one hand, and file-systems on the other. This is an abstraction layer that makes file-systems more independent from the physical disks beneath. Physical disks are grouped together into this abstraction layer on one side, and then file-systems are created on the other that are completely independent from the disks.

A SIMPLE EXAMPLE

To use LVM, in the first place one should make sure we have the appropriate programs installed. They come as standard on most modern Ubuntu versions (13.04, 13.10, ...) and derivatives (e.g. Linux Mint 15). They can also be installed on earlier versions - look for package `lvm2` in your favorite software manager.

Once the software is installed, we need to go through three steps:

1. We convert our physical disk partitions to LVM Physical Volumes (PV). Each partition converted will



HOWTO - USE LOGICAL VOLUME MANAGER

no longer be able to support a file-system directly, so we must be sure they do not contain any data we wish to keep.

2. We group our Physical Volumes together into a Volume Group (VG). There is no actual need to have more than one PV in a particular VG, but if needed we may do so. PVs are then combined within the VG, and their sizes add up. The group then be used as a single pool of disk space we can split into chunks and use as we wish.

3. We create one or more Logical Volumes out of our Volume Group.

To take a practice example, we could start out with a computer with two drives, /dev/sdc and /dev/sdd, both being 2 TB volumes. sdc holds a single partition occupying the whole disk and noted sdc1, while sdd is divided into sdd1 and sdd2, both with 1 TB. We want to use sdc1 and sdd2 with LVM, conserving sdd1 without making any changes to it.

Step 1 would be:

```
sudo bash
```

(we need to be root to alter LVM volumes)

```
pvcreate /dev/sdc1
```

```
pvcreate /dev/sdd2
```

Once we have the two PVs created, we can create our volume group to host them. The VG may be called just about anything (but no spaces or “weird” characters in the name), so in this case we will simply call it “group”. So we will create it first with one volume, then add the second:

```
vgcreate group /dev/sdc1
```

```
vgextend group /dev/sdd2
```

We now obtain a single volume group, with total capacity 3 TB. We can now proceed to step 3, and “carve” out one or more logical volumes as necessary - and yes, the technical term used is indeed “to carve”. For example, if we wished to create a 1.5 TB volume we will be using as a backup, we could do it this way:

```
lvcreate group -L 1500M -n backup
```

Giving the LV a significant name with the -n option will make it easier for us to recognize and use.

For example:

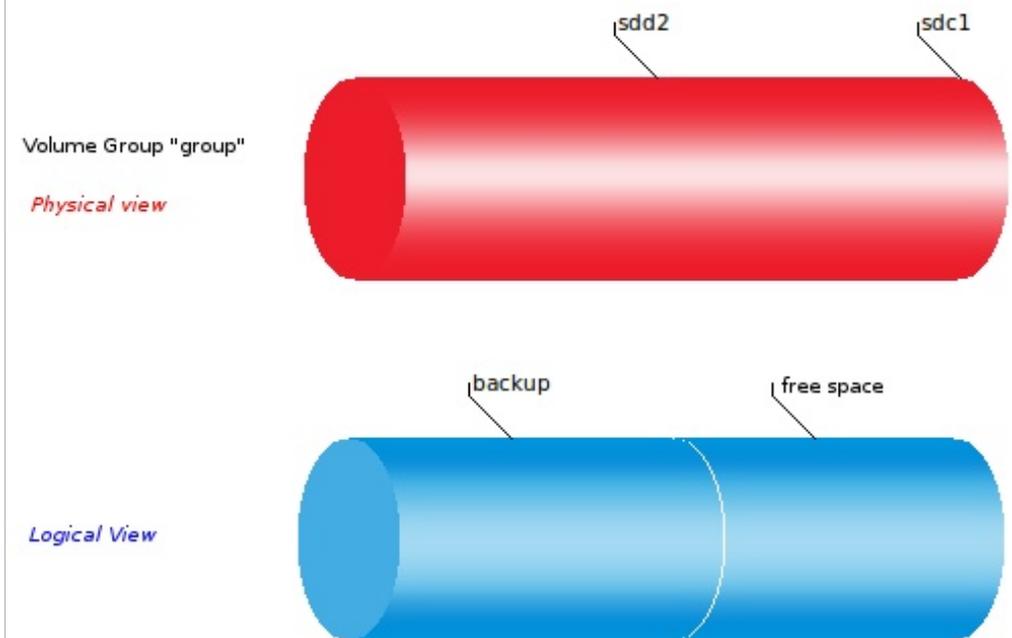
```
mkfs.ext4 /dev/group/backup
```

```
mount /dev/group/backup /mnt
```

Note that LV divisions borders do not necessarily coincide with disk partitions; in this example our backup LV may in actual fact be completely contained within /dev/sdc1, or it could just as well be divided up into chunks physically located on both /dev/sdc1 and /dev/sdd2. We do

not have to manage it. On the other hand, we do know we still have some 1.5 TB of free space we can further carve up – into one or several more LVs. So this simple example already shows us some of LVM’s potential flexibility.

Though I personally prefer to use the console, not everybody feels that way. So a graphical front-end for LVM has been developed, called system-config-lvm. If you feel so inclined, by all means install it on top of lvm and give it a test. It



HOWTO - USE LOGICAL VOLUME MANAGER

also outputs some pretty graphics such as those above. However, it does have some issues with recent versions of Linux Mint and Cinnamon (but not with Unity).

SCENARIO 1: JUGGLING DISKS WHEN YOU DO NOT KNOW WHAT THE FINAL SYSTEM WILL LOOK LIKE

Most of us have installed Ubuntu or one of its derivatives on a disk that already contained a Windows partition. At least at the very beginning, we often wanted to retain at least part of the Windows installation, just in case we wanted to go back - for whatever reason. At the end of the day, however, this "going back" somehow rarely happens. What does happen is that we run out of space on the Linux side, and end up thinking of adding the Windows partition or partitions to Ubuntu. LVM makes this easier.

Let us set out with a typical laptop hard drive, that often comes from the factory preset with three partitions:

- /dev/sda1 as C:, the main Windows system disk;

- /dev/sda2 as D:, where we place our files, often set up as the My Documents folder;
- /dev/sda3, an invisible partition under Windows that holds a system image that can be used to reinstall Windows to its factory settings.

A simple installation choice is to replace the NTFS D: partition with our Ubuntu system. If we simply discard the existing D: and give the installer free reign we will probably end up with something a tad complicated. It loves to create two separate partitions, one for the system itself, and another as swap memory. So we could end up with something like this:

- /dev/sda1 as the existing C:
- /dev/sda2 as an ext4 partition
- /dev/sda4, a primary partition containing
 - /dev/sda5, a secondary partition used as swap
- /dev/sda3, the existing invisible image partition

This is already a bit involved, and will become even more so when, later on, we decide to get rid of Windows once and for all, and reuse /dev/sda1 and /dev/sda3 for Ubuntu as well.

So, in comes LVM to help us

simplify things. When installing Ubuntu for the first time, instead of going directly to the installer, we will fire up a terminal, become root and:

```
sudo bash
```

(no password is necessary with the LiveCD)

```
pvcreate /dev/sda2
```

```
vgcreate group /dev/sda2
```

```
vgcreate group -n system -L <system size>
```

```
vgcreate group -n swap -L <swap size>
```

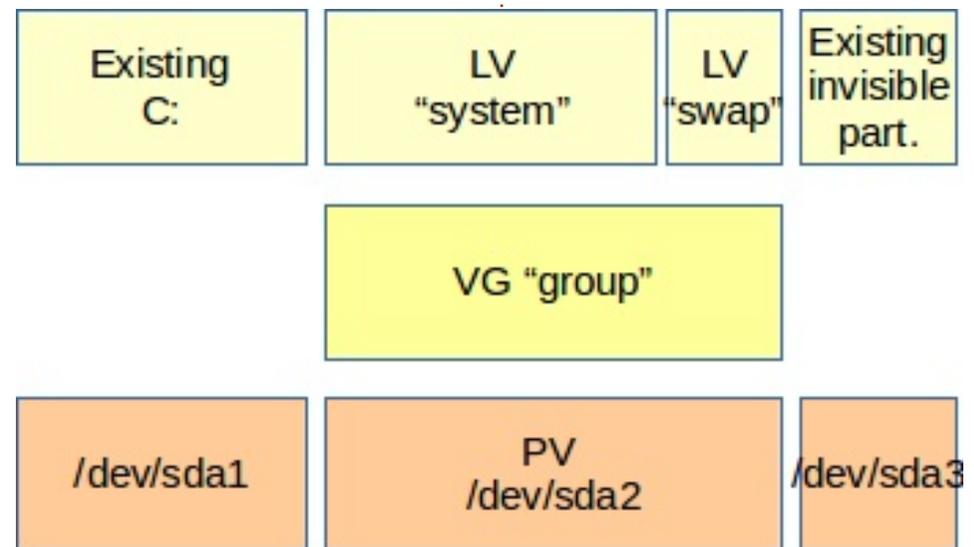
We now have:

- /dev/sda1 as the existing C:
- /dev/sda2 as a PV, used by VG "group" and containing

- /dev/group/system, a partition we will use as the system root
- /dev/group/swap, a partition used as swap
- /dev/sda3, the existing invisible image partition

We can now proceed to the Ubuntu installer, and when presented with the screen to choose disk management, go to "manual partitioning". From there, we get /dev/group/system formatted as an ext4 file-system and mounted on /, while /dev/group/swap gets formatted and used as swap. The rest of the system installation process should go on as usual.

A word of warning: earlier versions of Ubuntu were not LVM-



HOWTO - USE LOGICAL VOLUME MANAGER

aware. What is described here should work out-of-the-box for Raring (13.04) and Saucy (13.10), as well as Linux Mint 15. Earlier versions need some heavy-duty console work to mount the target partition, bind-mount accessory directories such as /tmp and /dev; go to it in a chroot jail and install the lvm2 package. This is certainly doable, but perhaps not for the faint of heart.

Now, what happens when we wish to add the Windows partitions to our Ubuntu system? Let us say, for example, that we simply wish to reform /dev/sda1 (C:), retaining the image partition for the time being. Inside a terminal, we create a new PV from /dev/sda1, and add it to VG "ubuntu":

```
sudo bash
pvcreate /dev/sda1
vgextend group /dev/sda1
```

If we now use the vgsdisplay command, it can be seen that the space available on VG "ubuntu" is larger than before, comprising both /dev/sda2 and /dev/sda1. We can then make our LV "system" and root file-system larger. This cannot be done while it is mounted, so we

will need to boot from the Ubuntu LiveCD once more. Supposing our /dev/sda1 partition was 40 GBytes in size, and /dev/sda2 was 60 GBytes, we could:

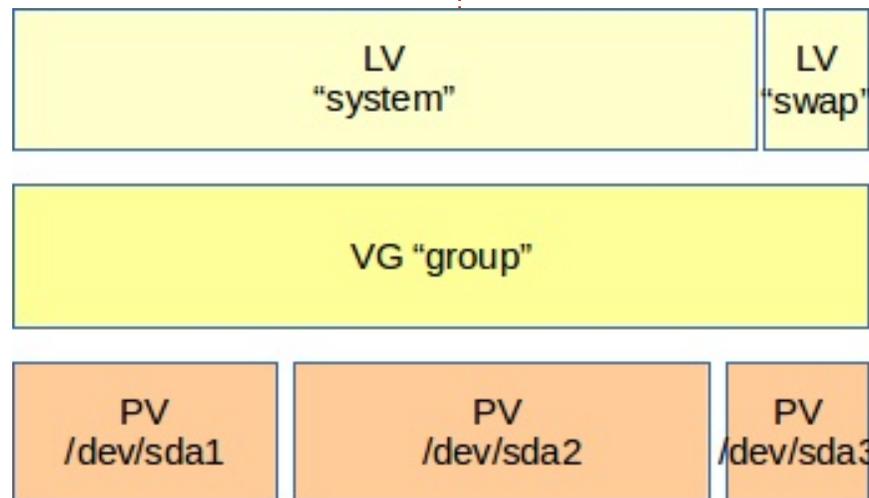
```
sudo bash
lvextend /dev/ubuntu/system
+40G
```

(adds the newly available 40 GB to LV "system")

```
resize2fs /dev/ubuntu/system
100G
```

(makes the "system" ext4 filesystem a total of 100 GB large)

We should then be able to reboot as usual from the hard disk, and test drive our new, enlarged, system. Our system now looks like this:



- /dev/sda1 and /dev/sda2 as PVs, used by VG "group" and containing
 - /dev/group/system, a partition we will use as the system root
 - /dev/group/swap, a partition used as swap
- /dev/sda3, the existing invisible image partition

If we do not need the image partition any more, we could also go one step further, and add it as well to VG "group".

The usual warnings apply: we are making extensive changes to the operating system, so things could eventually go very wrong. Naturally, you should do this for the first time only on a system which you are prepared to reinstall from scratch if things go bad. And, of course, make a backup of your

data before starting!

SCENARIO 2: A SAFETY NET WHEN UPGRADING AN EXISTING SYSTEM

One of the nice things about Ubuntu and GNU/Linux systems in general is that new, upgraded versions come up very often, in our case every six months. So there is a natural tendency to upgrade systems that may or may not really need it - just to stay on top of things. But strange things do sometimes happen when upgrading, even with Linux. For example, in theory it should be possible to upgrade an existing Ubuntu-based system to a newer version by simply changing the repositories in /etc/apt/sources.list and then executing

```
sudo bash
aptitude update ; aptitude
safe-upgrade -y ; shutdown -r
now
```

In actual fact, this is not a good idea for many distributions. Some old cruft always gets carried over to the new version. More importantly, new versions may need newer versions of their

HOWTO - USE LOGICAL VOLUME MANAGER

configuration files, that might even not be in the same directories any more. This is particularly true of Linux Mint: except for their Debian-based distributions (LMDE), upgrading in place has specifically been discouraged - see Clément Lefebvre's "How to upgrade to a newer release"

(<http://community.linuxmint.com/tutorial/view/2>).

Anyhow, if downloading, installing and re-configuring a new system is just too much hassle and you do want to go down the path of in-place upgrading, it is always nice to have a Plan B ready: a way to get your system back up and working as it was before trying to upgrade. Once more, LVM comes to the rescue with its snapshot capability.

The way this works is, if we have enough free space in our VG, we create a second LV that is essentially a copy of the first, though in essence "frozen in time". We can then do whatever we want to the first LV, and can always copy our files back from the snapshot if needed.

First, we create the snapshot of LV "system". In my own case, I tend

to use very small 10 GByte system partitions, so I will create a snapshot reserving double this space - just in case it fills up for whatever reason. We will call the snapshot "system-snap":

```
sudo bash
```

```
lvcreate -s -n system-snap -L 20G /dev/group/system
```

That's it, all done. Now we can proceed to upgrade our system to whatever newer version is available (Ubuntu 13.10, perhaps?), or perform any other kind of black magic that could eventually bork things in a big way.

Let us suppose that it has indeed done that, and that our original system (in LV "system") is in a very bad shape. We need to reboot from the trusty Ubuntu LiveCD, and in a terminal make a copy of the snapshot we will work from:

```
sudo bash
```

```
lvrename /dev/group/system system-old
```

```
lvcreate group -n system -L 10G
```

(put in here whatever size you used

for the system partition)

```
dd if=/dev/group/system-snap of=/dev/group/system bs=1M
```

At this point, LVs should be this way:

```
lvscan
```

```
ACTIVE Original  
'/dev/group/system-old'  
[10,00 GiB] inherit
```

```
ACTIVE Snapshot  
'/dev/group/system-snap'  
[20,00 GiB] inherit
```

```
ACTIVE '/dev/group/system'  
[10,00 GiB] inherit
```

The "system-old" LV is the old, wrecked system. "system-snap" is the snapshot containing a copy of the system as it was before making the fatal changes, and "system" is the current, recreated, system.

We can clean up a bit by removing both the old system and the snapshot, regaining the space they occupied:

```
lvremove /dev/group/system-snap
```

```
lvremove /dev/group/system-old
```

Depending on what happened while upgrading/breaking your

system, your GRUB installation may also be off-course. It is perhaps a good idea to recreate it while you are still booted into the LiveCD.

Mount your new /dev/group/system, for example on /media/ubuntu/system, and reinstall GRUB on /dev/sda with:

```
grub-install --boot-directory /media/ubuntu/system/boot /dev/sda
```

SCENARIO 3: MAKING YOUR "SWISS ARMY KNIFE" USB BOOT DISK

In this day and age of external USB hard drives, we have at our disposal a simple means of booting ailing computers, and fixing whatever is wrong or even cloning a system image. The standard Ubuntu LiveCD image can help us, naturally, but another way of going about it is to simply install Ubuntu on an external drive. That way, we can add our very own collection of tools and keep it set up just the way we want it.

However, some tasks need simple tools such as booting a very old computer off a lightweight Xubuntu, while others may need a

HOWTO - USE LOGICAL VOLUME MANAGER

more complete solution, such as cloning a school lab's worth of computers all to the same spec with Ubuntu Studio. Since modern hard drives have plenty of space, why not put several different distributions on the same external drive?

However, we may encounter several practical problems. If we want to use the disk to boot older computers, we will need to use the classical Master Boot Record (MBR) partitioning system. This means we are limited to four principal partitions - slightly more if we use secondary partitions, but still very limited in number. So, even if the hard disk itself has more than enough capacity for several dozen Ubuntu installations - perhaps of varying ages and flavors - the partition map will not take them.

Even worse is that unless we plan very carefully beforehand how much space to assign to each installation, we may end up with some partitions running out of space, while there is plenty left over on the disk. A case in point would be if we have a backup partition on disk, as well as the system installations.

At this point, it should be clear that we would be better off using LVM on the external disk.

Setting it up is simplicity itself. Once we have booted from the LiveCD for the first Ubuntu system we wish to install, we use the terminal or the gparted utility to create a first ext4 partition occupying the entire external disk: /dev/sda1 or similar. Once that is done, we can proceed to the terminal, create a PV from the partition, and a VG from that. Then we can carve out a first LV from the VG, format it and install Ubuntu on it.

For successive installations of other versions of *buntu, we

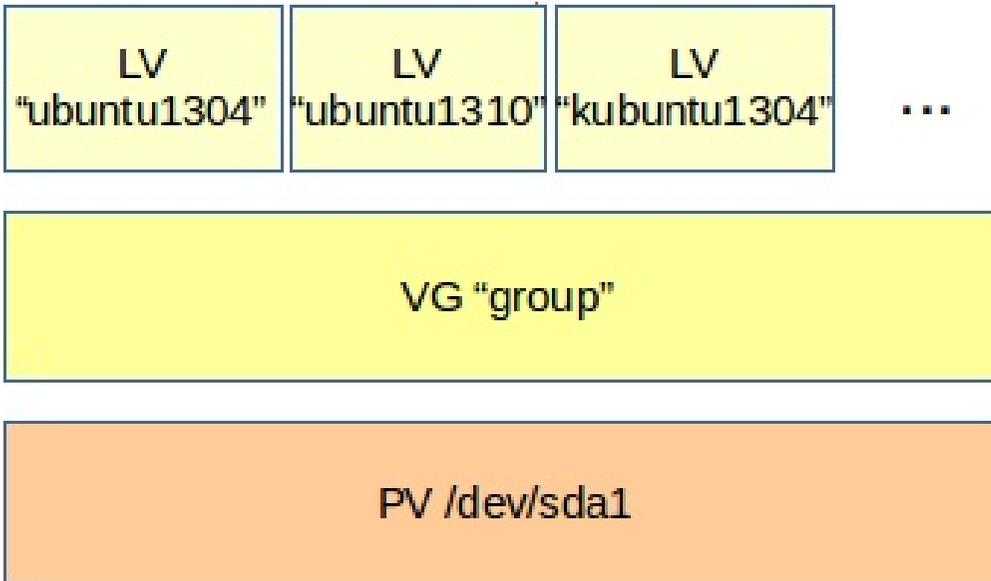
proceed in the same way: boot from the appropriate LiveCD, use the terminal to create a new LV (use significant volume labels!), format the new LV with an ext4 file-system, an install in the traditional way. Each time, however, the GRUB installation will probably get overwritten by the installer so, whatever you do, make sure beforehand the version of Ubuntu you are installing does actually support lvm out of the box! Just run any of the commands - lvsdisplay will suffice - from the terminal.

A helpful tip: it is best to practice installing Ubuntu on an external disk on a computer with its internal hard drive

disconnected. Doing things this way can help avoid wee mistakes, such as overwriting the main hard drive's system inadvertently...

CONCLUSIONS

Though LVM may seem to be a bit of overkill for a simple Ubuntu installation, as we have seen in these three scenarios - installing



HOWTO - USE LOGICAL VOLUME MANAGER

GNU/Linux on a (former) Windows machine, using a snapshot to keep our backs covered when upgrading in place, and making an external USB boot disk - it does give us a lot of flexibility. Without LVM, one partition must correspond to precisely one file-system. With LVM, we can play separately on two different levels: with partitions on the one hand, and with Logical Volumes on the other. While I hope these three scenarios may help the reader in real-life situations, this is by no means the end of LVM applications. Further possibilities include using it to encrypt a complete partition - as set out in Xavier Berger's "Create a thief-proof PC", in Full Circle issue 67 (<http://fullcirclemagazine.org/issue-67/>). If we want to set up a data repository or NAS box on our network, LVM mirroring gets us data replication (two copies of each item) without having to fiddle around with RAID volumes.

In any case, I would like to encourage everybody to play around with LVM, at least to explore its various possibilities. But please stay safe: do it on a computer you do not need for work, that you are happy to format if need be, and obviously with all

data on it correctly backed up. Even small units such as USB flash pen-drives are enough to try out LVM, though the system may or may not regain access to the LVs if you pop them out and back in again (it usually does). Perhaps the best way to test it out would be with virtual machines under Virtualbox. If your main computer has enough RAM (say 2 GB), enough CPU horsepower (any dual-core should suffice) and some space disk space left, this would definitely be the way to go.



Alan teaches computer science at Escola Andorrana de Batxillerat (high-school). He has given GNU/Linux courses at the University of Andorra and currently teaches GNU/Linux systems administration at the Open University of Catalunya (UOC).

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/>



HOW-TO

Written by Robin Catling

JPG To PDF Using ImageMagick

Fear not Blender fans – Nicholas will return next month. Well, assuming his employer will actually give him a day off and some free time that is...

It's a common enough task, trying to convert multiple jpg files into one pdf, particularly when I don't need these to be converted with such high quality, I just want the black and white text readable.

I can scan, crop, and monochrome in a graphics program, but compiling them into a single PDF booklet was always tricky. Which is how I started using ImageMagick.

ImageMagick is a command-line conversion program that is capable of so many more batch operations than this - resizing, compression, format conversion - and it's available on all platforms - Linux, Windows and Mac.

Bear in mind that creating a PDF document from multiple JPEG images can take some time, and you may want to trial different

settings for size and quality of output, so I suggest you make a copy of the JPEG files in a temporary folder to play around with and use Imagemagick on those – NOT on your originals. Reducing the resolution as a first step will also make things much quicker.

Originally I was just using the basic:

```
convert *.jpg output.pdf
and/or convert *.jpg -adjoin
output.pdf
```

which works most of the time. However, there's a bug in the convert routine which can in some versions give a segmentation fault when converting a number of JPEG files to one PDF file.

What this command does is take all the .jpg's (or format of your choice) in a folder, and convert them to a single PDF - you can name it whatever you like.

You can avoid the segmentation fault bug and do the compression at the same time if you use:

```
convert *.JPG -compress Zip
output.pdf
```

but the zip compression appears quite inefficient and results in huge file sizes.

You could resize and lower the quality of the images using:

```
mogrify -resize 50% -quality
25
```

Which overwrites the originals. You can combine resizing and conversion using:

```
convert -quality 25 -resize
50% *.jpg -adjoin output.pdf
```

which works, but takes longer as you're combining batch operations.

And yes, you can resizing images without overwriting the originals by specifying a new file name:

```
convert '*.JPG' -resize
640x480 newfile%03d.jpg
```

which outputs the converted images as newfile001.jpg, newfile002.jpg, and so on.

Alternatively, if you want to retain the original filename and prepend "new", you could use a bit more code:

```
for file in *.JPG ; do
convert "$file" -resize
640x480 "new-${file}" ; done
```

This is just a sample of what Imagemagick can do if you are prepared to experiment at the command-line.



After a few years messing about in the TV and film industries, **Robin Catling** now has a proper job once more, doing worthwhile things as Web Systems Coordinator at Dartinton Trust's Research in Practice.

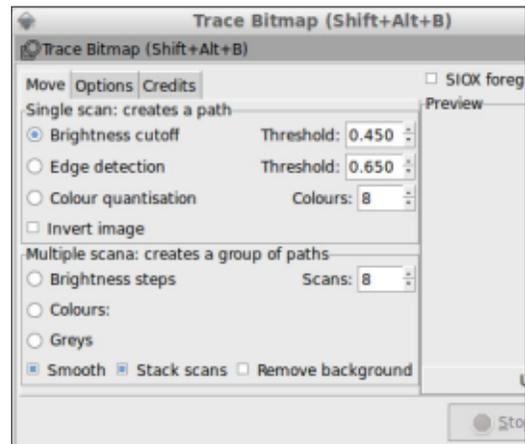


HOW-TO

Written by Mark Crutch

Inkscape - Part 20

In this instalment I'll be continuing to look at the Trace Bitmap dialog. Previously I covered the Single Scan algorithms that are used to create a single path from a bitmap image, and which tend to work best on simple line art. This time it's the turn of the Multiple Scans section of the dialog, which creates more than one path, and can often be a better option when dealing with colored images and logos.



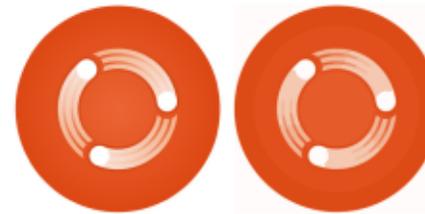
You may recall from last time that the Potrace code within Inkscape only traces a simple black and white image. The three Single Scan algorithms represent different ways to reduce an image down to a single collection of black

and white pixels. The Multiple Scans option, on the other hand, creates a collection of different black and white images, and then passes each of them to the Potrace algorithm individually. The multiple paths that are created are then assembled into a single group before being inserted into your Inkscape document.

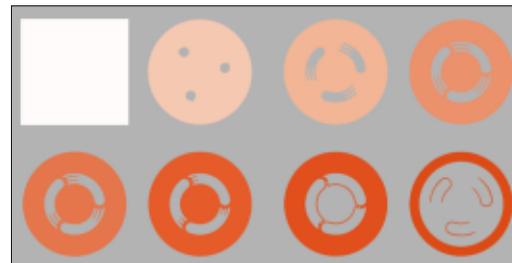
This part of the dialog has radio buttons for selecting one of three algorithms. Based on the top half of the dialog, you would be forgiven for thinking that the Scans spinbox applied to only the Brightness Steps algorithm, but it actually sets the number of paths that will be created regardless of which algorithm you choose. Similarly, the checkboxes at the bottom apply to all three algorithms.

Let's begin by tracing the Full Circle Magazine logo. For this example, I'll use the "Colours" algorithm, with 8 scans and the "Smooth" and "Stack scans" checkboxes enabled – I'll explain all these options a little later.

Selecting the logo and clicking the Update button suggests that the final result will be okay, so I can go ahead and click on the OK button to perform the trace. The result looks like this, with the original PNG image on the left and the traced version on the right.



It's not a perfect trace – though that was never to be expected, given that we've reduced it to 8 colors. But it's not too bad, and does represent the original image quite well. By ungrouping the 8 paths and separating them out against a gray background we can get a clearer view of the paths that have been created.



There are a couple of things to note about these paths. The first is that there's a square path which is used as the background for the final trace. This obviously results in our round logo becoming square, even though the original image was a PNG with transparent corners. The second observation is that the paths stack up on top of each other: each scan sits on top of the one before it, obscuring the lower one and just letting it show through where holes have been left – the final image is created by successively hiding parts of the lower objects.

Consider the three white circles in the final logo. As you can see, there is no single trace that contains three white circles. What you see in the final trace is the white background showing through the holes that have been left in each subsequent layer in the stack. Suppose you wanted to tidy up those circles a little. You would potentially have to modify seven of the eight paths!

And what of the square

background? You can remove it by checking the “Remove Background” option in the Trace Bitmap dialog, but that just creates 7 paths rather than 8. You get all the same paths as before, but the background path is removed.

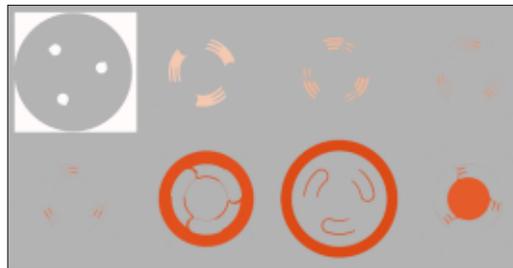
However, because the three white circles are the result of the background path showing through, we actually end up with three holes in the traced image instead.



In this case, it's not too tricky to fix up the results, regardless of whether you've chosen to automatically remove the background or not. As a general rule, if you're tracing a rectangular image with no transparency then there's little to be gained from removing the background.

However, for tracing images where transparency is important, it's usually best to either check that box, or to manually remove the background layer afterwards. Be aware that you might have some fixing up to do, though.

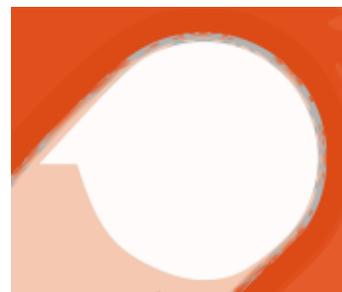
The way in which the paths sit on top of each other, with later paths obscuring those below, tends to give the best visual results. But not all Inkscape users are using it for artistic reasons; there's a vibrant community of people who use it to create files for vinyl cutters, plotters and laser cutters. For these users the stacked traces would result in some lines being cut and re-cut up to eight times. In the previous example the circular outline of the FCM logo is repeated across most of the paths. Not only is this inefficient, but it could also result in damage to the work being produced, or even to the machine itself. By un-checking the “Stack Scans” option, you will end up with a series of paths that fit inside each other without overlapping.



There's no hiding of lower paths by those further up the z-index. In fact you could re-order the z-indexes and still get the same visual result. The circular outline of our logo now appears only twice –

once as an interior path of the background trace, and once as the outside of the large orange path.

For some tasks – even those that don't involve blades or lasers – a trace made in this way can be easier to edit than one made with “Stack Scans” enabled. As always, the only way to know for sure is to try both approaches and see which works best with your particular source image. Be aware, however, that disabling this option can leave small gaps, blobs and other artefacts between the traced paths. A close up of one of the white circles in the FCM logo reveals this issue all too clearly (the gray part is the background color showing through gaps between the paths).



The last of the three checkboxes, “Smooth”, specifies whether the bitmap image is traced in its original form, or after slightly blurring it. The purpose of blurring

it is to smooth out some of the minor color changes that often appear in an image, usually resulting in a trace that is less broken up and has fewer nodes. Because the Potrace algorithm produces only monochrome paths, the resultant trace will still be crisp and sharp, even with this option enabled. Usually I find it's better to leave this checked and benefit from simpler paths, but, if you want to preserve the fidelity of the original image as much as possible, then you may prefer to uncheck it.

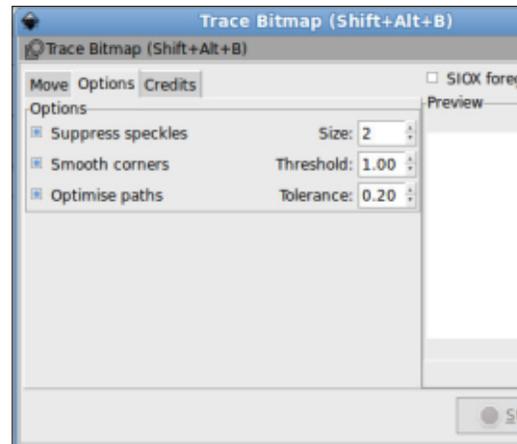
Now that I've explained the options, it's time to take a look at the other two algorithms. All the examples so far have used the “Colours” algorithm. The “Greys” algorithm is a trivial one to explain: it's exactly the same as the “Colours” algorithm except that the resultant paths are converted to shades of gray.

The “Brightness steps” algorithm separates the bitmap image based on the brightness of each pixel, and produces a group of grayscale paths. The number of paths created is always one more than the “Scans” value, giving you a minimum of three paths. I've found this algorithm to be more

problematic than the other two: on my machine it will create a preview for the FCM logo, but won't actually trace it at all! Fortunately it seems happy to trace renaissance paintings, so I've used an image of the Mona Lisa to create a sample grid of the same image traced using each of the three algorithms, with three different values for the "Scans" spinbox. I've only gone as far as 32 scans here – larger values will give more accurate traces but can take a long time to complete and result in a lot of paths and nodes to deal with.



There are some other options in the Trace Bitmap dialog that apply to all the tracing algorithms, whether they are of the Single Scan or Multiple Scan variety. These are gathered together into their own tab in the dialog.



Each can be enabled or disabled using the checkboxes to the left, and each also takes a single parameter which is set in the corresponding spinbox to the right. I tend to leave these all enabled with the default settings, however you may wish to disable or edit them for particular traces.

"Suppress speckles" removes all paths with a size that is less than the number specified. Turning this off can lead to some very long tracing times. Setting it to a higher

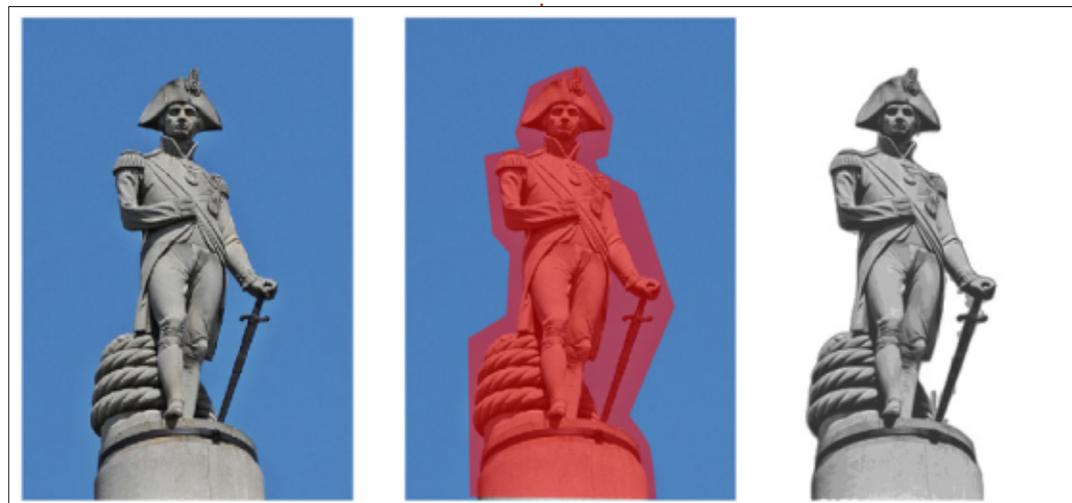
value can speed up tracing, but lose fine detail in the process.

"Smooth corners" creates paths with rounded corners. Disabling this, or setting the threshold to 0 results in sharp corners instead. The threshold value determines the amount of rounding that is allowed. The effect of this option is most visible on images with sharp corners – if you want to trace a pixel image or a QR code and preserve the individual squares then you may want to disable this option.

"Optimise paths" allows Inkscape to reduce the number of nodes by simplifying the traced paths in a similar manner to the Ctrl-L shortcut. The resultant paths aren't a precise match for the original, but

are usually close enough for artistic purposes. The Tolerance spinbox sets the amount of variation from the original path that is allowed, with higher values allowing more variation and resulting in fewer nodes.

There's one final option in the Trace Bitmap dialog that sits on its own in the top-right corner: "SIOX foreground selection". This uses the Simple Interactive Object eXtraction algorithm (<http://www.siox.org/>) to separate a foreground object from the background prior to tracing. With this enabled, the dialog requires two objects to be selected: the image to be traced, and a filled path that roughly covers the foreground object to be extracted. You can see the effect with this



trace of a picture of Horatio Nelson's statue on top of his eponymous column in London.

Although this mode might appear useful, in practice it can be difficult to get a good result. The previous example only worked well when I reduced the original high-resolution image down to something with fewer pixels. In my experience, you would be better off separating out the foreground using The GIMP (which has its own, more interactive, implementation of the SIOX algorithm) and then tracing the result.

Automatic tracing of images is no magic bullet. It's not a practical way to convert a bitmap to an infinitely scalable vector unless you're also prepared to accept a loss of color depth and fine detail. And whilst it can be useful for tracing logos and line art, the resultant paths are likely to still require additional editing. Often the time spent trying to tidy up the results of an automatic trace would have been better spent manually tracing instead: it's better to manually trace a round logo as a circle than to automatically trace it and end up with an approximately circular path, for example. Don't

rule out the possibility of mixing both methods, though. As always, you really need to experiment with your own images in order to decide which approach works best for you.

Image Credits

"La Gioconda" (aka "Mona Lisa") by Leonardo da Vinci

http://en.wikipedia.org/wiki/File:Mona_Lisa,_by_Leonardo_da_Vinci,_from_C2RMF_retouched.jpg

Statue of Admiral Nelson, at the top of Nelson's Column, London.

http://commons.wikimedia.org/wiki/File:Admiral_Horatio_Nelson,_Nelson's_Column,_Trafalgar_Square,_London.JPG



Mark's Inkscape created webcomic, 'Monsters, Inked' is now available to buy as a book from <http://www.peppertop.com/shop/>

LIBREOFFICE SPECIAL EDITIONS:



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

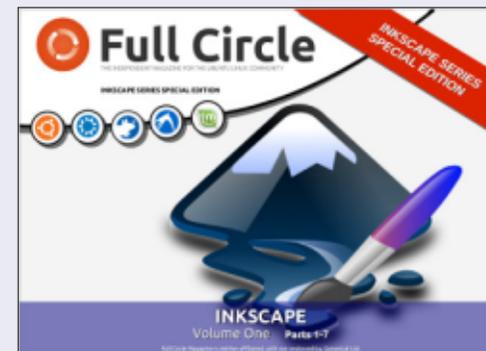


<http://fullcirclemagazine.org/libreoffice-special-edition-volume-two/>



<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/>



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





Online
BACKUP

Secure
SYNC

Easy
SHARING

Whether you need to access a document you have stored on a remote server, synchronize data between a Mac, Windows or Linux device, share important business documents with your clients, or just rest easy knowing all of your data is safely, securely, and automatically backed up - SpiderOak's free online backup, online sync and online sharing solution can handle all your needs!

SpiderOak offers a different approach to online backup by combining a suite of services into one consolidated tool - free online backup, synchronization, sharing, remote access, and storage. This difference is further measured in our zero-knowledge privacy policy - the first one ever employed in this setting. Our flexible design allows you to handle data from any operating system (Mac, Windows and Linux) or location (external drives, network volumes, USB keys, etc...) using just one centralized account.

- ➔ Access all your data in one de-duplicated location
- ➔ Configurable multi-platform synchronization
- ➔ Preserve all historical versions & deleted files
- ➔ Share folders instantly in web ShareRooms w / RSS
- ➔ Retrieve files from any internet-connected device
- ➔ Comprehensive 'zero-knowledge' data encryption
- ➔ 2 GBs Free / \$10 per 100 GBs / Unlimited devices

<https://spideroak.com>

Download mobile clients
for **iOS & Android**

JOIN SPIDEROAK NOW
Get 2 Free GBs

Get 25% off any SpiderOak package
with the code: **FullcirclemagFans**



ASK THE NEW GUY

Written by Copil Yáñez

Hi, everyone! Welcome back to Ask the New Guy!

If you have a simple question, contact me at copil.yanez@gmail.com.

Today's question is:

Q: Geepez101 asks: I am looking for a basic checkbook register program. The ones in the software center for Linux Mint 15 are too hard for a 65 year-old man to figure out. I need something simple to use, do you know of any simple programs like this?

A: Thanks for getting in touch, **Geepez101**. I am not sure if you're a 65 year-old man, or if that's just the metric by which you judge things. Like, I sort of imagine you as a 22 year-old woman walking into a hardware store and saying, "Hi, I'm looking for a table saw, preferably one that's good for cross-cutting. Oh, and make sure a 65 year-old man can use it." I think this is actually a brilliant new certification program. Anything

you can buy should have a tag on it labeling who can use it. Apple products would be 21H (21 year-old hipsters), Michael Bay Movies would be 18MC (18 year-olds plus Man Children) and Justin Bieber would be S13 (screechy 13 year-olds).

By this metric, I'd like to think Ubuntu would be labeled AA (All Ages). Let's see if that's true as we look for some software that might work for you. Linux Mint and Ubuntu, the two most popular beginner distros, are similar enough that any solution for one should work for the other.

The first, most obvious place to look for any kind of software is in the Ubuntu Software Center. You call it up by clicking on its icon along the left (it looks like a paper grocery sack), or type the word Software in the Unity search field.

Once you've got the Software Center open, you can type the description of what you're looking for in the top-right. I tried Checkbook Register first, and

found nothing. Then I tried Finance. Bingo!

Once I got a list of finance apps, I sorted them by Top Rated to see what others think are the best programs out there. One thing that sometimes slows me down is spending too much time deciding which program to install. I look at ratings, reviews, go out to find blog posts and comparisons. But there's no cost to installing any of these, so I've tried to train myself to just install a few of the most popular ones and then test them out myself. If one of them seems to flow for me, I use it. If none do, I can always remove them (from the same Software Center) or just leave them installed, no worries.

I installed GnuCash, Grisbi, and KMyMoney, because, why not? They installed quickly and appeared along the left side of my screen. Let's take a quick drive and see if any of these might work for you, Geepez101.

I clicked on **GnuCash** and it asked me if I wanted to import any

existing files. That would assume I had any existing files which, in turn, assumes I have any money. I have neither. So I clicked no and loaded the tutorial.



The tutorial started with a list of about ten million things I could do, none of which sounded remotely interesting. I'm guessing this is not the sort of thing a 65 year-old man wants to see. But there was a listing for Checkbook, let's try that.

Wall of text. Booooooring. I clicked a few things, tried setting up a ledger (whatever the hell that is), but couldn't easily find a simple checkbook register showing Bank Account, Deposits, Withdrawals, and a Total. I could feel my 65 year-old man's blood pressure rising. He



worked hard to get where he is, he doesn't have time to mess around with stuff that doesn't have what he wants.

GnuCash gets the 35AF label (35 year-old accounting fetishist).



Let's try **Grisbi**, which doesn't sound like accounting software at all. It sounds more like a music streaming app, where you type in spoken word erotica, and it spits out a stream of sexy poetry set to spa music. I'm willing to give it a chance either way.

Okay, so I start Grisbi and it feels friendly, even though it's taking me through a bunch of general configuration options. It's holding my hand, trying not to scare me off as it takes me through the setup screens. Grisbi feels like my buddy down the road who just got new solar panels installed. He's

all, "they'll save you money!", and "they're environmentally friendly!", and "you get a tax subsidy!" And he's smiling the whole time, you know, and it's going to break his heart when I tell him I just don't need something so complex for my 600 sq. ft. ground-floor apartment. I click through all the screens quickly, ignoring most of the text like I imagine my 65 year-old dad used to do when he had instructions for building a cabinet. I finally get to a screen with a menu item that says Bank Account. Okay, that might be the one. Does it matter I've apparently set it up in Italian Lira?

I open the Bank Account screen and whoa, there, tiger. If I had categories and sub-categories of payees, I might be smart enough to have made some money in the first place. But once again, my inner 65 year-old is a busy man who just wants a simple check register. Ins and outs with a running total. Is that so hard?

Sorry, Grisbi, you are trying too hard. I'm breaking up with you. And no matter what anyone else says, it's you, not me.

Verdit? Grisbi gets the 20SOAS

label (20-somethings who like Overly Attached Software).



On to **KMyMoney**. Again with the crazy names, what the heck? Did someone forget to finish naming it? What's the K for? Couldn't you remove it completely and have a pretty decent name for your program? I'm beginning to think the developer has a thing for the letter K, and weird fetishes like that, while entertaining in my everyday life, aren't so endearing when I JUST. WANT. A. CHECK. REGISTER!

Still, I'm willing to cut it some slack. Setup takes less than fifteen seconds as I repeatedly click Continue until the screens stop moving. This time, there's an intuitive screen that allows me to set up a bank account with a starting balance, which seems promising.

I set myself up with a US\$100,000,000 bank account. In my mind, 65 year-old me is a brilliant but ruthless industrialist who single-handedly introduced sweet-and-salty popcorn to Central America in the 80's and made a fortune, which he now spends trying to genetically engineer tiny elephants the size of a housecat. How cute would that be? A little herd of tiny elephants, running through your house, knocking over appliances and washing themselves from the toilet bowl. Sixty-five year-old me is like John Galt without an Ayn Rand fetish.

I am going to go ahead and acknowledge that this is the third time in this column I've mentioned fetishes, and accept that there might be more going on at home than meets the eye.

Anyhoo, I set up my bank account and then use the New option to add some debits. As is the case with all these programs, each option and screen has wayyy more flags, fields and toggles than I really need for my simple check register. But as long as my random clicking through screens doesn't break anything and still gives me

something that looks familiar, I'm happy.

Below is what I build after about three minutes of playing with KMyMoney.

I've got a starting balance, several transactions, and a running total. Not bad. As you can tell from all the options along the left, there's a lot I don't need. But as far as getting something set up quickly and painlessly, KMyMoney is KForMe.

Verdict? KMyMoney gets labeled 65MWCLPGN (65 year-old men who can look past goofy nomenclature).

Now, before any fans of GnuCash and Grisbi come at me with the argument that both those

programs can do exactly the same thing as KMyMoney, let me agree with you. The issue isn't what can this software do? It's what can I do with this software?

This is the the main reason I chose to answer Geepez101's question this month. Not because I think a lot of you might be looking for checkbook registers. But because Geepez101's search is the same one a lot of beginners undertake for whatever it is they need. And because the world of Linux is so huge, deep, and complex, sometimes we get lost in the embarrassment of riches afforded us.

Rather than necessarily look for a less complex program, start with a search for anything that meets your most basic criteria.

Remember, you're not paying for any of this. Install the program, which couldn't be easier, in either Ubuntu or Mint, and play around. If you can't get what you want within sixty seconds, move on. No one will judge you for not giving a piece of software a full 360 peer review. Kick the tires, see if it's what you need. It takes less time to load, test, rinse and repeat, than it does to check reviews or forums for just the right piece of software. And once you find something that works and you become familiar with it, you might find having more options will allow you to expand how you use it. What's more, using one program could eventually lead to realizing why one of the others you dropped initially might be a better choice.

It's true that we don't need a

Saturn V to get us to the grocery store and back. But if one's available, does the job, and doesn't blow you up on the launchpad, why not give it a try?

Now before I abandon you, Geepez101, let me give you two more options you may not have considered. One relates to something I've argued here before. If you can find a cloud solution that works for you, it makes system hangups, crashes and breakdowns non-events. Just replace your computer or reinstall Ubuntu, and all your data is still available from the cloud.

Financial sites take a bit more vetting, of course, but I feel pretty comfortable putting some of my most trusted information online, mainly because those sites tend to do a better job of backing up my information than I do.

I went looking for an online check register and found ClearCheckbook.com. It doesn't link to your accounts or ask you for bank details, and I was able to set up the following in under a minute.

Simple, clean, easy. And if your system crashes some day, you'll

The screenshot shows the KMyMoney interface with a ledger for 'US Bank'. The ledger is titled 'Ledgers' and shows a search bar for 'US Bank' and a status dropdown set to 'Any status'. The table below shows transactions for the current fiscal year and today. The transactions are as follows:

| No. | Date | Details | C | Payment | Deposit | Balance |
|---------------------|----------|---|---|---------------|----------------|----------------|
| Current fiscal year | | | | | | |
| | 01/01/13 | Opening Balances | | | 100,000,000.00 | 100,000,000.00 |
| Today | | | | | | |
| 7 | 12/08/13 | 7-11 Sushi | ⚠ | 4.25 | | 99,999,995.75 |
| 3 | 12/08/13 | Three year supply of Soylent | ⚠ | 3,000.00 | | 99,996,995.75 |
| 5 | 12/08/13 | Aicor Life Extension full body cryonics account | ⚠ | 100,000.00 | | 99,896,995.75 |
| 1 | 12/08/13 | Lanborghini | ⚠ | 1,500,000.00 | | 98,396,995.75 |
| 2 | 12/08/13 | Elephant Genetic Testing | ⚠ | 10,000,000.00 | | 88,396,995.75 |

ASK THE NEW GUY

still know exactly how much you spent on Tiny Elephant Litter.

One final option you may or may not have considered, Geepez101. You probably already have LibreOffice loaded on your computer. If so, you could set up a spreadsheet and add a few columns. Very simple, but plenty of options to make it more complex if you need it.

You've been a good sport, Geepez101. Seriously, thank you for asking the question.

And whether you're a 65 year-old man, a 22 year-old woman, or anything in-between, I hope something I've covered here will help you find what you're looking for.

Good luck and happy Ubuntuing!



Copil is an Aztec name that roughly translates to "you need my heart for what again?" His love of women's shoes is chronicled at yaconfidential.blogspot.com. You can also watch him embarrass himself on Twitter (@copil).

ClearCheckbook.com
Money management made easy

Overviews: Your Account Dashboard
Register: Balance Your Checkbook
Reports: Where's Your Money Going?
Budget: Set Some Limits
Tools: More Tools To Help!

Account Dashboard + Add New Gadget

You can move your gadgets around in any order you would like. If you want to remove a gadget from this page, click the 'X' in the upper right of the gadget. You can minimize and maximize the gadgets by clicking the '-' or '+' icons and you can access a gadgets properties by clicking the arrow icon. To add new gadgets, [click here](#)

Add Transaction

Date: Dec 8 2013
Amount:
Description:
Transaction: Withdrawal
Account: Checking
Category: --
 Jive Transaction
Add Transaction Split
+Add Account +Add Category

Latest Transactions

| Date | Amount | Account | Category |
|---|----------------|----------|----------|
| 12/8/2013 | -1.95 | Checking | -- |
| Watercress & cream cheese finger sandwich | | | |
| 12/8/2013 | -1,000,000.00 | Checking | -- |
| Tiny Elephan shoes | | | |
| 12/8/2013 | 100,000,000.00 | Checking | -- |
| Initial Deposit | | | |

[View Account Summary](#)

Account Overviews

| Account: | Jived: | Balance: |
|-----------------------|---------------|------------------------|
| Checking | 0.00 | 98,999,998.05 |
| Total Balance: | \$0.00 | \$98,999,998.05 |

Budget Watch

Budget List:

| Budget | Spent | Limit | Percent |
|--------------|-------------|-------------|-----------|
| Total | 0.00 | 0.00 | 0% |

Reminders / Recurring Transactions

December 2013

| Sun | Mon | Tue | Wed | Thr | Fri | Sat |
|-----|-----|-----|-----|-----|-----|-----|
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| 8 | 9 | 10 | 11 | 12 | 13 | 14 |
| 15 | 16 | 17 | 18 | 19 | 20 | 21 |
| 22 | 23 | 24 | 25 | 26 | 27 | 28 |



When the 1 terabyte hard drive in my wife's PC failed almost 2 years ago we bought a 2 terabyte hard drive to replace it. We were very fortunate because 2 terabyte hard drives were at an all time low price and our 1 terabyte drive hadn't completely failed yet (smart reported errors) so we were able to copy all the data over to the new hard drive. Even if the drive had completely failed we had regular backups of her most important data. Although we had backups, it was still several days before I was able to buy a new hard drive from our favourite computer store.

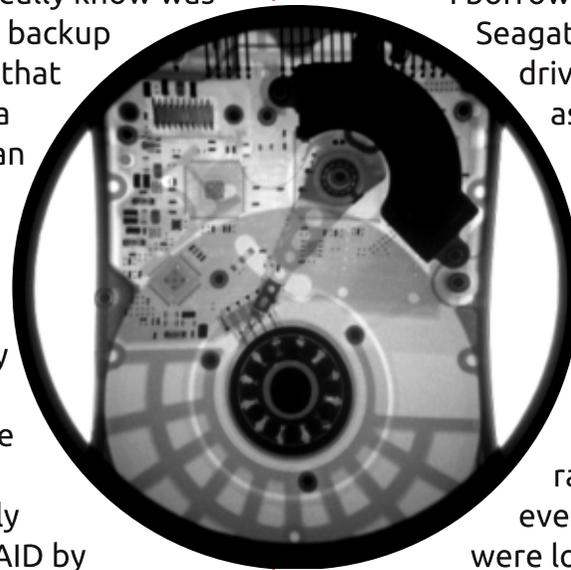
Ever since that drive failure I intended to investigate RAID, Redundant Array of Independent Disks. The key for us is redundancy, if one hard drive fails, we want the system to keep humming until we have a replacement drive.

Before looking into RAID I knew a few things about RAID: it involves multiple hard drives, sometimes the drives are used in a stripe, combined together for more space

and speed, and sometimes for redundancy, when one drive mirrors the contents of another, and sometimes a combination of both. I also knew that there were several levels of RAID, RAID 0, 1, 2, 3, 4, 5, 6, and nested RAID levels (0+1/10, 100, 30, 50 and 60).

What I didn't really know was that RAID is not a backup solution. It's true that when you mirror a drive you do get an identical copy of the drive, but therein lies the backup problem - if you accidentally delete a file from one drive it's gone from all mirrored drives. If you really want a backup, RAID by itself isn't going to help you retrieve lost data.

I also wasn't sure about was which level to use, disks to buy, or how to best implement RAID. Luckily the organization I work for has a great sysadmin I can bug



from time to time (thanks Paul!). If there's one lesson I can impart: plan your RAID array before buying equipment. In our case we already had the computer the RAID array would go in, but I wished we had other equipment as I was building the array.

I borrowed 4 x 250GB Seagate STxxxxxx drives from work to as a test before going out and shelling out lots of money for 4 new drives. As I went to put the drives into the computer I ran into 2 issues: even though there were lots of drive bays the case felt cramped and the new power supply I bought didn't have enough Serial ATA ends to power all the drives plus the SATA Blu-ray drive.

All four hard drives fit into the case, however the top drive was

getting dangerously close to blocking the SATA headers on the motherboard. When connecting the SATA power connectors it felt like there wasn't enough space between drives. I've seen my fair share of broken SATA connections on both hard drives and optical media and the cramped nature of the power connections made me think that over time the drive power headers might break. The case I was working with was a mid-tower sized case but it really felt like I should have bought a roomier gaming case with a bit more space both horizontally and vertically.

It also dawned on me that I might not have enough SATA data headers on the motherboard, I had 4 hard drives plus a SATA Blu-ray drive, requiring 5 headers (luckily the board had enough). Some motherboards might only have a couple of SATA headers, so check before you buy a bunch of drives. If you don't have enough drives you can always buy PCI or PCIe SATA card, though if you're going to buy an expensive card you might look at a cheaper NAS box (at least for

home) since the price is almost the same as an expensive RAID card.

My wife doesn't like to wait on her computer, so while redundancy was important speed was also equally important - I couldn't have the RAID array slowing down her system. RAID 0, known as striping, stripes your data across multiple drives giving you better performance, but it provides no mirroring of data or parity error checking. If one drive in a RAID 0 configuration dies your entire stripe is rendered useless, thus RAID 0 really isn't RAID (redundancy). Because we want redundancy we also need to mirror the drive. RAID 1 provides mirroring, but no parity error checking and no striping. In a RAID 1 configuration as long as you have 1 good drive functioning the RAID array will continue to function.

For this month we'll be creating a RAID 10 configuration (RAID 1+0), we'll mirror a couple of striped drives (250GB+250GB=500GB) to another set of drives. Although we'll be using 1TB of drives only 500GB will be available to us. RAID 10 needs a minimum of 4 drives (Linux can do

RAID 10 with only 2 drives using the Linux MD driver, but this isn't official "industry standard" RAID 10), 2 get striped together and the other two mirror the data.

When I started connecting up all the drives I noticed RAID features in the motherboard BIOS. Initially this was how I thought I'd set up the RAID array. Having lived through software replacements of hardware (winmodems, video cards stealing from motherboard RAM, etc) I've always subscribed to the idea that a hardware solution is better than a software one. I didn't really want to go out and buy a couple of hardware RAID controllers for home because good hardware RAID controllers are costly and you need an identical controller for backup (even revision matters) if one fails. After spending a bunch of time fiddling with the RAID BIOS I did a bit more research and found a lot of advice against using motherboard RAID. Some research suggested that if the motherboard dies you lose your data, other research said it might be possible to recover the data using an identical board but indicated most motherboard RAID took more cycles away from your

CPU than dedicated RAID hardware.

Ultimately I decided to use the Linux software MD-RAID setup. While it also draws cycles from the CPU it's less expensive to implement and easier to recover (buying an identical motherboard would have been pretty difficult since the motherboard I'm using is no longer manufactured).

TO BE CONTINUED....



Charles McColm is the author of *Instant XBMC*, and the project manager of a not-for-profit computer reuse project. When not building PCs, removing malware, encouraging people to use Linux, and hosting local Ubuntu hours, Charles blogs at <http://www.charlesmccolm.com/>.

Full Circle Podcast Episode 38, Just The Two Of Us

Your hosts:

- Les Pounder
- Tony Hughes
- Jon Chamberlain
- Oliver Clark
and Freaky Clown



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

In this Episode we announce the new format for the shows, talk about our hardware, review, issue 76 of the magazine and we have an interview from the STEM York Raspberry Jam.

Download



MY STORY

Written by Michal Urbanski

I started with Linux in late 90's. Those were times when Ubuntu didn't exist yet, and the most popular distribution was Red Hat Linux. I was then curious about what Linux was, and how it worked. Although Windows had gained total desktop market share at that time, I felt something 'magical' about Linux and its terminal. It was just different and reminded me of the good MS-DOS times when GUI – as it is known nowadays – didn't exist yet for standard users.

But I never got into Linux as much as I could. Many distributions appeared, and I tried many of them, but, from the beginning, my favorite was Ubuntu. Mainly because it was so user friendly. However, with almost all distributions and versions, I had some problems. Some of them didn't work correctly after installation; there were issues with drivers; some even wouldn't install on my computers because of hardware mismatches – they just didn't work. So my passion for Linux decreased; however I'm always optimistic about each new

release.

Just to clarify – I'm a Microsoft guy. I work with Microsoft products every day and I live from programming in .NET. I'm not planning to change this in the near future; however I'm an open-minded person, so I decided that I would like to improve my various programming skills and Ubuntu would be a great platform in order to achieve this goal. So I installed Ubuntu 13.04 on a virtual machine – because I don't like using my main operating system (which should be fast, and without any additions installed) for development tasks, and I'm happy with it. It works smoothly and fast without any significant issues. The evolution of virtualization software makes such decisions easier. Frequently, I have installed with dual- or multi-boot using custom partitioning, but it just does not feel right these days.

Last week I had a pretty scary moment. I was doing maintenance tasks on moving files between internal and external drives, which I use as a kind of backup. Suddenly,

the most important external drive – with the most of my important data – stopped working. It was spinning up at the power on, but Windows didn't manage to mount them into the operating system. I was pretty scared seeing this behavior, because I had already lost one drive in a similar way. There were two possibilities – either the disk was down, or the external enclosure was broken in some way.

I had to quickly find a way to retrieve files from the disk. They were important for me and I needed them fast. But the problem was that this external drive is an old ATA and I had no possibility of plugging it into my current PC. I had to find the reason why the disk wasn't mounting on my computer. Then, I thought that, instead of spending some money and losing time to find, buy and wait for new external enclosure, I could use my really old PC which was put in the closet for a couple of years. I grabbed it, put an empty hard drive inside, and powered on. It has single core processor, 1 GB of RAM, makes strange noises, but it's

working. I needed an operation system only in order to retrieve files. So my natural first thought was to install Ubuntu 13.04 which I had burned on CD earlier. Installation was quite fast; it took about 20 minutes and I had Ubuntu running – with its full Unity experience. I have to say that I didn't expect such a good performance with such an old hardware. It is not super fast but it runs quite well.

Having prepared this computer with a working Ubuntu, I plugged my external drive and... it worked. All data was there without any broken partitions or corrupted files. Quick connect to Windows share, and I was able to copy all files into my working SATA internal hard drive. Thanks Ubuntu - you saved my money and my time, and restored my old computer to life.





REVIEW

Written by Ronnie Tucker

USB Microscope

Two years ago this month (Dec. 2013), I bought a dinky little USB microscope from Amazon (<http://www.amazon.co.uk/dp/B002JCSCYI>). It was on sale for £35 and I bought it thinking it'd be a bit of a laugh and either be pretty cheap and useless, or not work at all in Linux. I was wrong on all accounts!



The build quality is pretty OK for the price. Even after two years of (granted, infrequent) use, it's still holding up well.

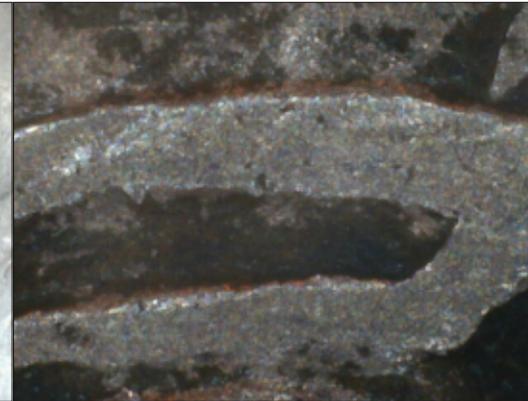
On the side of the microscope you have a large rotating dial which



is your focus. Really you have only two focus settings: near and far. Everything in between is a tweak to get sharper focus. Also on the side is a button for taking a snapshot (although most software has a snapshot facility which will give you a more stable shot), and a dial for turning the light on and dimming (if required). That's basically it. You plug the microscope into your USB port and you're off.

Obviously you'll want some software to use it with. I use *Cheese* (in Mint 16) but previously I used *Kamosa* in KDE.

Both acknowledge the microscope as a webcam and let you take still photos or videos. The microscope's highest resolution is



640 x 480 which is more than adequate for tinkering about with.

I recently started going through a box of old coins that I've had for years (that I got from my dad) and this little microscope was indispensable for finding mint marks and reading signatures.

Just watch when you're buying it though as this is the x400 version I'm reviewing. There is a cheaper x200 version which I have not tried.

PROS:

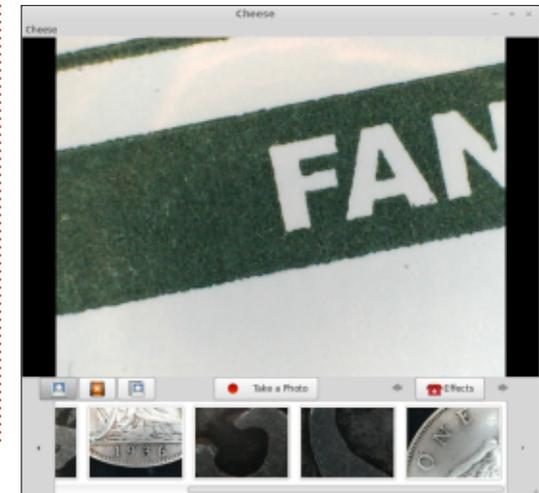
- It's cheap and does the job
- It takes pretty good quality photos considering its low price
- It has a light which comes in handy

CONS:

- The stand is a bit flimsy



I purchased mine from Amazon: <http://www.amazon.co.uk/dp/B002JCSCYI>





SOFTWARE SHOWDOWN

Written by Tushar Bhargava

The Old Guard

In July 2012, while the Ubuntu App Showdown was in full swing, an article appeared on OMG! Ubuntu!. The article entitled 'Rise of the RSS Clients' listed 4 new RSS readers - LightRead, Una reader, Feedivore and Pressboy - that were being developed as a part of the competition. The community's excitement was palpable, many users wanted a slicker and more modern newsreader than their current one.

Now, more than a year later, I went back to see how these new RSS clients were faring.

LightRead, with its beautiful UI, had won a large fan following, many referred to it as the 'Best Linux RSS Client'. However, the project was dealt a crippling blow when Google Reader shut down in July, 2013. LightRead was essentially an 'offline Google Reader'. As one user succinctly put it, without Google Reader, LightRead was 'useless'. The LightRead team responded by saying they've applied for API access for Feedly, but, at the time

of writing, no further update has been released.

Una reader had a rough start, the project lead expressed his fears on a comment thread: "I don't believe Una will be ready in time for the contest". His fears turned out to be well founded, and Una didn't make the Ubuntu App Showdown list. Later in the year, Chris McCutcheon made an announcement - "Due to the lack of free time, the Una project is being put on hold." However, mid 2013 Chris made another announcement, one that promised a brighter future: "Una Reader development will be starting up again but this time in multiple flavors."

And how is Pressboy doing? Well, it's available for download; however, I suspect development is not all that active. In an embarrassing turn of events, I became one of the top contributors by simply filing a bug about Pressboy not working for my version of Ubuntu

Finally Feedivore.

Well, the news is bad; I tried hard to locate the app; however, all I found was a year old Reddit thread and a screenshot.

I also looked through the list of submitted apps for the Showdown but failed to locate Feedivore. To put it simply, Feedivore has disappeared.

So, while LightRead finds a way to work around the Google Reader shutdown, Una completes its resurrection, and Pressboy finds new contributors (and don't get me wrong, I'm really looking forward to all of that happening), we have to fall back on the Old Guard - Akregator, Liferea and RSSOwl.

ONCE UPON A TIME

Akregator has been a part of KDE since the 3.4 release circa 2005.

Liferea was created on July 3, 2003. Lars Windolf, the project lead, explained the need for such a

RSS client. *"Problem was that Straw and Syndicator were too buggy, and back then installing Python and its relevant libraries was hell. Solving the dependency once at compile time seemed to be the better choice."*

RSSOwl, by strange coincidence, also began in July 2003. It was created by Benjamin Pasero. In January 2005, it was chosen as the SourceForge Project of the Month, marking its arrival as a major Open Source newsreader.

USER INTERFACE

Most RSS readers have a traditional three-pane user interface (UI) that resembles email readers. The first column is inevitably the browser that helps you select the feed source. The second column split into 2 panes, one gives a bird eye view of all the news items while the second shows the specific entry. Such a UI, while rather monotonous, has the advantage of being tried and tested and simply works.

SOFTWARE SHOWDOWN



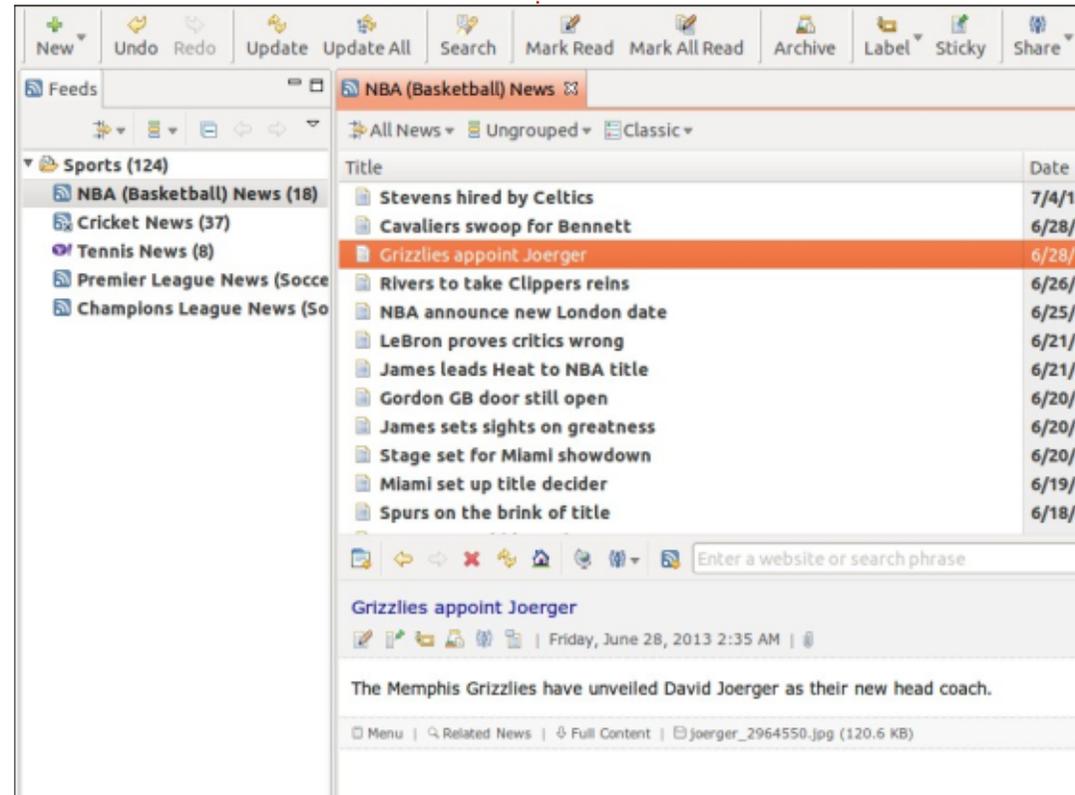
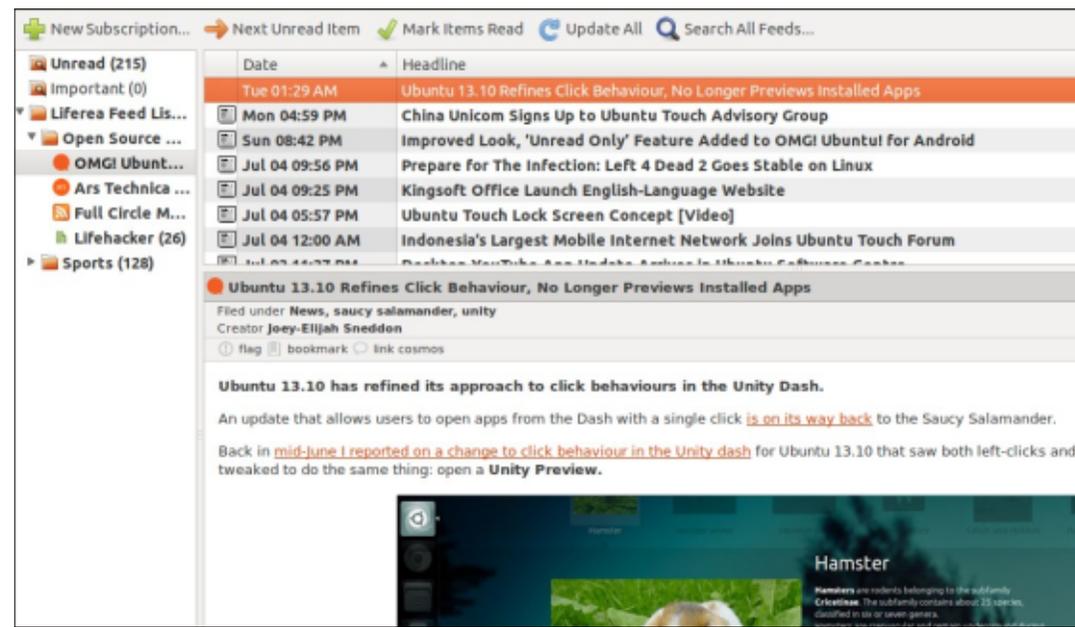
Akregator (above left), Liferea (above right) and RSSOwl (bottom right) all have a three-pane GUI but with subtle differences that I have outlined below.

Akregator has a simple toolbar with options to fetch feeds for the current source or to fetch all feeds. You can also mark a feed as read and move back and forth. The inbuilt browser supports tabbed browsing and is unobtrusive – allowing you to view the news comfortably. The search bar is also straightforward and allows you to easily find what you’re looking for.

Liferea offers more

comprehensive controls from the toolbar. Its toolbar includes options to create a new subscription, see the next unread item, mark items as read, update all feeds, and even search through the feeds. For most users these controls are sufficient and they will seldom have to navigate through the menu. Smart design. The inbuilt browser also supports tabs; however, it’s not as simple or fast as Akregator’s browser.

RSSOwl offers perhaps the most complex UI. As is clear from the screenshot, its toolbar is stuffed with buttons and drop-down menus; however, it still



remains fairly intuitive. Such a crowded toolbar has its advantage too; you can perform almost all (not just common) actions without ever going through the confusing menu hierarchy. If you prefer a more minimalistic toolbar, you can customize it (View -> Customize Toolbar). In the end, if I had to sum up RSSOwl's GUI in two words, I would say 'Power User'.

READING EXPERIENCE AND OTHER FEATURES

Akregator has a wonderful embedded browser. The tabs allow you to keep several stories open at once, and the minimal browser controls mean the news item gets prominence. Akregator aims to be a simple RSS reader; it doesn't have many more features. There are no smart folders or sharing options. But most users probably would not need such features under normal circumstances. The lack of features doesn't take away from Akregator; in fact it helps to underscore the fact that Akregator is a minimal easy-to-use RSS reader that is perfect for casual users.

Liferea offers a great viewing experience as well. Its browser is

not as good as Akregator, but still functions nicely. The ability to add bookmarks is appreciated. Liferea also allows you to define 'Search Folders' and has advanced search options that include searching in Twitter or Reddit.com. On the whole, Liferea does a good job of providing all the basic features users want in a simple UI – with more advanced features concealed under the hood.

Finally, RSSOwl – I've mentioned this before but, for emphasis, I'll repeat it – RSSOwl is a newsreader for power users. Sure it offers a great viewing experience, a fully functional browser, and sharing options (including Facebook, Twitter, Email) that will be appreciated by all users. However, the majority of its offerings such as the ability to archive news items, a labelling system, multi-field search, news bins coupled with news filters that make them automatic, are clearly targeted at advanced users.

CONCLUSION

The choice of an RSS reader is an extremely personal one. However, I think most users will find Liferea to be a good fit. With

its simple UI and comprehensive toolbar, Liferea can be operated by even the most non-tech savvy. However, features like 'search folders' and news bins will be greatly appreciated by the more experienced users.

Akregator is a great choice for those who wish to have a smoother browsing experience, and do not need any of the above mentioned advanced features. RSSOwl is perfect for advanced users who wish to set up complicated systems in place for staying on top of the news.

THE NEAR FUTURE

With LightRead, Pressboy and Una reader at the vanguard of the Linux RSS client revolution, we will, hopefully, soon have even more options. When I asked Mr. Windolf, project lead of Liferea, whether he was worried about the competition, he replied, "They are competition in the sense of trying to create even better feed readers and matching other use cases. But as we are all open source projects, it is about the fun and learning new things." To a more open future.

SUMMARY - AKREGATOR



The Good

- Simple UI
- Great inbuilt browser

The Bad

- Lack of advanced features such as automatic folders, sharing options, etc.
- The toolbar doesn't have some common options like 'Add New Subscription'

Website:

<http://kde.org/applications/interne/t/akregator/>

SUMMARY - LIFEREA



The Good

- Simple and intuitive UI with comprehensive toolbar options
- Advanced features like creating smart folders and news bins

The Bad

- No sharing options
- Inbuilt browser has inconvenient tabs

Website: <http://lzone.de/liferea/>

Note: At the time of writing the new release of Liferea (1.10.0) had not come out; I based my article on an older version and many blemishes may now have been fixed.

SUMMARY - RSSOwl



The Good

- Plethora of options including an archiving system, advanced search, etc.
- Sharing options that are perfect for more social users
- Flexible and customizable UI

The Bad

- UI a bit complicated compared to the other two
- Consumes a lot of memory while running, sometimes crashes

Website: <http://www.rssowl.org/>

THE WINNER OF THIS SOFTWARE SHOWDOWN IS

Liferea!



Tushar is a 17-year-old Indian who loves Ubuntu/FOSS. He programs in Java and C++, enjoys writing and, recently, making Android apps. If you enjoyed this article, his blog is at tusharbhargava.wordpress.com for more articles.

QUICK MY STORY

by Andrea Fiorina

I've followed Full Circle since FCM#10 (or so), and it's always a pleasure when a new issue is released.

I began using Linux in 2004 from the first version of Ubuntu (I've still got the CD with Ubuntu 4.10 on it!) but since I got married I cannot follow my curiosity to try and break the system and learn what should not be done. And when you have your first daughter, free time is always less and less. But I'm still curious and whenever I can I follow a question and try to learn something.

Recently, I planned to buy a Raspberry PI to setup a little FTP server in my office (we're a very little factory of only three people, and we try to do everything by ourselves if possible, and by using open source software for my personal, and economic, reasons), buy PCduino to set up a media center at home, setup a VPN from my house and my office, and find a solution to plan the backup in my office between the little server with Mint and the other PC's which have Windows XP, 7 and Linux.

I'm waiting until Christmas to buy anything (my wife always told me that I'm not allowed buy anything on Christmas Eve). I was trying to backup my PC's, but I cannot manage the rsync/Windows and I stopped trying after several nights with lack of sleep and very little result.

I remembered that Full Circle was look for some material to publish (even letters and request), and this is my first step



Join us on:



goo.gl/FRTMI



facebook.com/fullcirclemagazine



twitter.com/#!/fullcirclemag



linkedin.com/company/full-circle-magazine



ubuntuforums.org/forumdisplay.php?f=270

YOUR NEW YEARS RESOLUTION FOR 2014

I can't force you, but I want you all to promise me that you'll take the time to write at least one article for FCM in 2014. At the time of writing this (mid-December 2013) I have very few articles left in reserve for 2014 (certainly not enough for half the year) and I'm getting tired of begging for articles. I really love creating FCM for you folks out there, but if no articles come in there's nothing to publish. Please don't rely on the regular writers. Without them this magazine would never have lasted this long.

I'd like to thank all the regular (and sometime) writers, proof readers and translation teams who help make this possible.

Best wishes to you all for 2014

Ronnie

FULL CIRCLE NEEDS YOU!



Without reader input **Full Circle** would be an empty PDF file (which I don't think many people would find particularly interesting). We are always looking for articles, reviews, anything! Even small things like letters and desktop screens help fill the magazine.

See the article **Writing for Full Circle** in this issue to read our basic guidelines.

Have a look at the last page of any issue to get the details of where to send your contributions.

Tuxidermy

THERE'S AN OLD LEGEND ABOUT A MAGICAL NIGHT.

WHEN A MYSTICAL BEING CROSSES THE SKY SPREADING GIFTS.

ZOOOOOOOOOM!



SPECIAL GIFTS.
FREEDOM GIFTS.

AND THE WORLD THEN UNDERSTANDS THAT IT'S ALL ABOUT CARING, AND SHARING, AND HELPING.

AND EVERYBODY'S HAPPY.

C'MON, SANTA!
I ASKED YOU FOR A NEW LAWYER AND I GOT THIS?!

WELL, MAYBE NOT EVERYBODY.

MERRY CHRISTMAS!





Q&A

Compiled by Gord Campbell

If you have Ubuntu-related questions, email them to: questions@fullcirclemagazine.org, and Gord will answer them in a future issue. Please include as much information as you can about your problem.

Q How I can safely delete Windows 8 and install Ubuntu?

A You don't need to delete Windows, just install Ubuntu on top of it.

At the same time, it might be prudent to keep Windows in case you might sell the machine in the future. There are tools in Windows to shrink the partition, so you could dual-boot. Ubuntu can use any free space in the Windows partition, so all you lose is the space actually used by Windows.

Q What software would you suggest for copying dvds and also for ripping/shrinking dvds?

A Handbrake.

Q I accidentally uninstalled my network manager, and now I can't reinstall it. How can I install network manager?

A See this thread: <http://ubuntuforums.org/showthread.php?t=2191294>

TOP NEW QUESTIONS AT ASKUBUNTU

* How to add an event in Ubuntu calender?
<http://goo.gl/wsDyjf>

* Which distribution to choose while downloading Skype from its official site for Ubuntu 13.10?
<http://goo.gl/ljjjEL>

* How to easily start a webserver in any folder?
<http://goo.gl/uAVNBm>

* How can I recursively delete all files of a specific extension in the current directory?
<http://goo.gl/ytxfo6>

* Confusion with printf command?
<http://goo.gl/dsOJRT>

* Convey some message to all

users

<http://goo.gl/VdUP1v>

* Auto switch workspaces
<http://goo.gl/TkSZRb>

* How can I optimize Ubuntu for minimal energy usage?
<http://goo.gl/zfljWT>

* Command to convert an upper-case string to lower-case?
<http://goo.gl/Glpf0r>

* What is the difference with all the different types of .tar archives?
<http://goo.gl/WEobdy>

* Should I be concerned about the Linux.Darlloz worm?
<http://goo.gl/R3h5j0>

* Does Ubuntu 13.10 still encounter issues with UEFI and Windows 8?
<http://goo.gl/dVeesb>

TIPS AND TECHNIQUES



Android Connection

I recently got my first Android smartphone. After figuring out how to make and receive phone calls, the next thing I wanted to do was put some music on the phone so I can listen while riding public transit.

For some reason, the Android developers decided a phone shouldn't look like a flash drive when it's plugged into a computer. After lots of Google searches, I found a couple of pages which held out hope that I could cope with "mtp" for transferring files. Then The Right Answer popped up: Airdroid.

I enable the phone's Wi-Fi, then install Airdroid from the Play Store. Start Airdroid. Connect to web.airdroid.com on my computer. Point the phone's camera at the QR code on the web page and tap, "scan QR code." Wow!

The web page shows me how much





space is being used of the built-in storage and the SD card. On the left are icons for music, videos and pictures, and several others. I can easily move music, pictures and videos back and forth between the computer and the phone, just what I wanted.

The goodness doesn't end there: there's an icon for Contacts, where I can add and edit contacts using my PC's keyboard instead of the crummy little touch-screen keyboard on the phone.

Airdroid has more features which I haven't explored. The things I have used are enough to make me say, "this is the best app of the decade."



After a long career in the computer industry, including a stint as editor of Computing Canada and Computer Dealer News, **Gord** is now more-or-less retired.



STARTING NEXT MONTH: SECURITY ISSUES

Beginning with FCM#81 **Michael Boelen**, open source developer and author of *Rootkit Hunter* and *Lynis* (<http://cisofy.com>), will answer your questions on security related issues.

In **FCM#81** he'll show you how to improve the security of your Ubuntu install with Lynis then answer a couple of your questions.

Don't worry, Q&A will continue as normal!



What will you do now that you've survived The Calamity? Well, I wouldn't blame you if you answer by saying you'll be playing Bastion on Ubuntu.

Bastion is a one-player, action role-playing video game developed by Supergiant Games and published by Warner Bros. Interactive Entertainment. Originally released for "Xbox Live Arcade" in July 2011, Bastion was later released for Windows in August 2011. Bastion was even released as a web app for Google Chrome in December 2011. Then, in May of 2012, Bastion was released for Linux and Mac OS X. Ever since its original release, Bastion has been very well received by critics, and it has had very strong sales on every platform for which it's been released.

To play Bastion, you just need to download and install it either from the Ubuntu Software Center, or through Steam. If you get it through the Ubuntu Software Center, then it will cost you \$18.99, but through Steam the current selling price is \$14.99. Having saved a little money by buying it through

Steam, you can add the Bastion Soundtrack for \$9.99 which is also highly recommended. Having installed the game, you're ready to play.

One aspect I really like about Bastion is that you can play with either a mouse-keyboard setup or with a game controller. Originally I was playing the game with a mouse and keyboard. If you use a mouse and keyboard, then you will be using the standard W-A-S-D movement bindings, the mouse for attacking, and a few other keys for other miscellaneous things such as

drinking a health bottle or a magic potion for special moves. When Steam became available for Ubuntu in February 2013, I re-installed the game and started from scratch. I also started using the Razer Onza RZ06 game controller, which by the way, works very well through Ubuntu. If you use a game controller, then you will be using the left analog stick for movement, and the A-B-X-Y buttons for attacking and such. No matter which option you go with, the actions will be the same; you can carry two weapons at a time in addition to one special move. You

can change your choice of weapons and/or special moves any time as long as it's not during a battle level. To perform the special moves, you must consume an entire magic potion. You can carry up to three magic potions and three health bottles. The health bottles can be refilled in several fountains, but it's always wise to be conservative and know when to use a magic potion or a health bottle.

In Bastion you play the role of The Kid, who wakes up from The Calamity – only to find Caeldonia, his homeland, gone to waste. In order to re-store Caeldonia to its former glory, The Kid must go to various colorful floating lands and collect Cores that give power to the Bastion so that it can re-create the great Caeldonia that once existed. Along the way, you must collect different kinds of weapons, powerful & exotic elixirs, and learn different special moves that improve your fighting skills so you can defeat a wide range of enemies who try to prevent you from collecting Cores. You also meet a few interesting people along the



way who return to the Bastion and interact with The Kid along his quest to save Caeldonia. Perhaps the most important of these characters is an old man named Rucks who also happens to be the narrator of the story as well as The Kid's mentor who helps him along the way.

As you move through the game, the floating landmasses build walkways for you through which you can walk. Throughout the game you will encounter some of the most colorful and beautiful scenic graphics you have ever seen. The graphics are hand painted and truly unique. In the meantime, you are listening to the narrator tell your story with a play by play commentary of your every move with a most wicked sense of humor, while in the background plays one of the best soundtracks I've heard from a game in the last few years.

The game auto-saves effortlessly, and, if you're playing it through Steam, then you've got that extra sense of security knowing that your progress is saved with their "Steam Cloud" feature, which means that no matter what computer you use,

you can freely download and play the game while also maintaining your previous progress.

The minimum requirements are:

OS: Ubuntu 12.04
CPU: 1.7GHz or better
Memory: 2 GB or more
Video Card: 512 MB VRAM or better
Hard Drive: 1.3 GB

CONCLUSION

I definitely recommend this game to anyone. Since I first played this game, I was very impressed by the artwork. It can be enjoyed by casual and hardcore gamers alike,

no matter if they're young or old.

Pro's:

- Exquisitely colorful hand-painted artwork
- Easy to play
- Made by a small independent developer
- You can use a mouse-keyboard or a game controller
- Relatively low minimum requirements
- If you like short games, this one can be completed in one day
- Very pleasant music
- Narrator is funny and will even make fun of your actual playing style

Con's:

- If you're a FLOSS purist, this game is not proprietary, but at least you can play it with free graphics drivers
- If you take into consideration that the game can be completed in less than five hours, paying more than \$10 might be too much. The lowest price, through Steam, is \$14.99

Overall, I give the game 4 out of 5 stars. The high price and short completion time, unfortunately, make me take out one full star. Other than that, it's a great game. If it were a little longer, or if it cost around \$10, I would give it a perfect five out of five.



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



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.



I am from Bangladesh. I have been using Ubuntu since 7.04. I love it and all its derivatives (Mint, Bodhi, Xubuntu...). I use Ubuntu on my desktop. But I install Linux

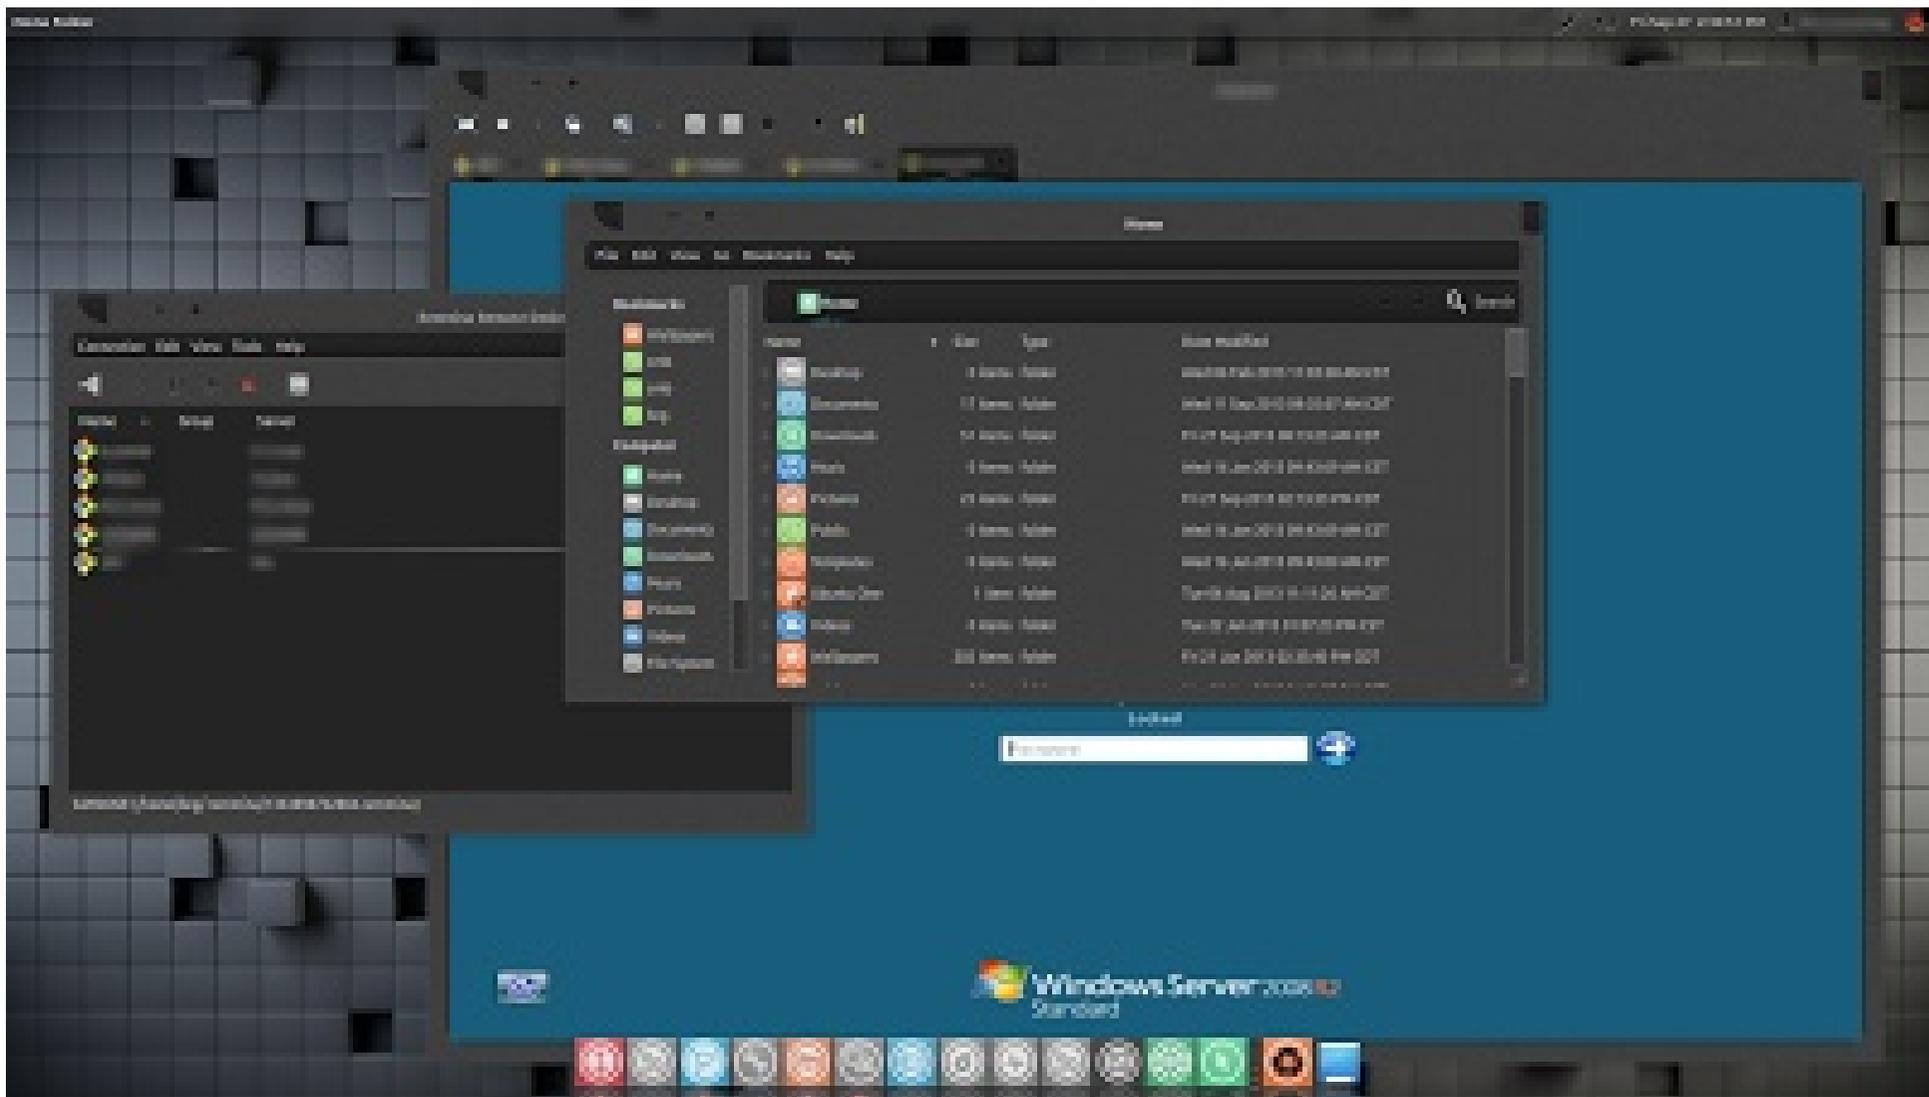
Mint15 Mate in my laptop.

I create an additional panel for right side. I add some applications shortcut here which help me in my

everyday work. I like conky. I think it is a great tool – which is available only in linux.

Specs: Toshiba Satellite C800D. Dual-core AMD E2-1800 APU with Radeon HD Graphics. 2 GB RAM.

Susanta Barman



Been an Ubuntu user at home and work since 8.04 Hardy Heron.

At my office, I recycled an old HP Compaq Deskpro DC5100 Intel Pentium 4, 3.0 GHz with 3GB RAM, Intel 915G graphics – that was an old XP workstation on our Active

Directory domain, and it is now running Ubuntu 12.04 with Unity.

GTK theme: malys-rough-right
Windows theme: malys-rough-right
Icon theme: malys-ex
Cursor theme: Pulse-Glass
Dock: Dockey

I made the workstation a member of our Active Directory domain using Likewise Open. I use Remmina to remotely manage our Windows servers.

Ubuntu is great for getting

more life out of what would be redundant Windows workstations.

Baz Greenaway



Asus F5VL Laptop
OS - Ubuntu 13.04
Resolution - 1280x800
Processor - Intel(R) Core(TM)2 Duo
CPU T5750

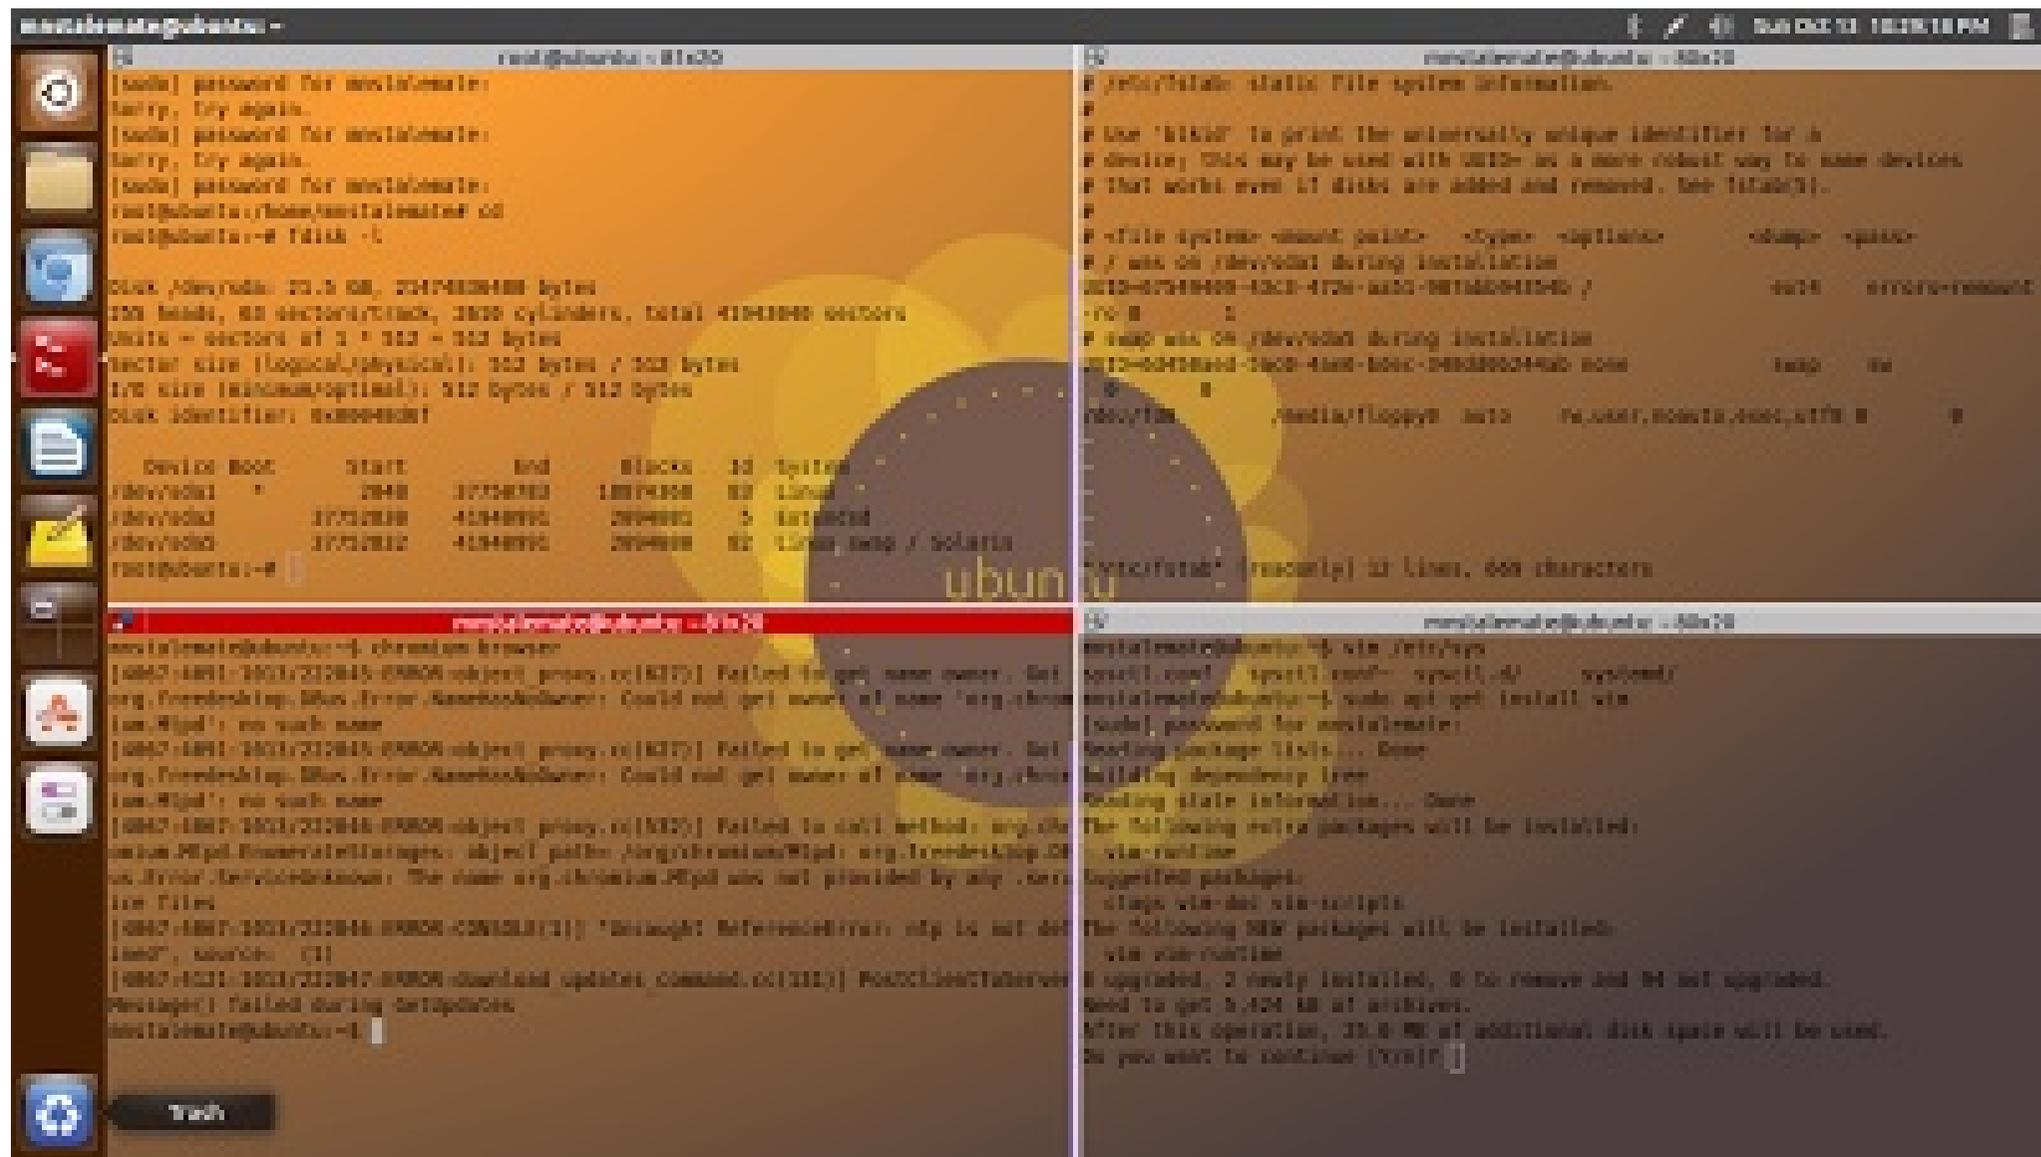
Graphics - ATI Radeon X2300
Ram - 2 GB
Drive - 320 GB

I got this laptop for free with no

HDD on board. I have an iMac with
Mac OS X10.7.5, and my girlfriend's
desktop PC with Windows 8. My
decision to install Ubuntu was
pretty obvious cause we both love

working in Ubuntu very much

Vitaly Rudenko



I have a Dell XPS L502X. I have been using Ubuntu since 2011. Installed theme is Moka with the same icon pack. On the desktop, terminator is running four split terminals.

Display : 15.6" (39.6cm) HD (1366x768) WLED display with TrueLife™
 Processor : 2nd generation Intel® Core™ i7-2670QM processor 2.20 GHz with Turbo Boost up to 3.10

GHZ
 RAM : 4GB 1333MHz DDR3 SDRAM
 HDD : 750GB 7200RPM Hard Drive
 GRAPHICS : 2GB NVIDIA® GeForce® GT 540M graphics with Optimus

O.S : Ubuntu 13.04 Raring Ringtail 64-bit

Muddassir Nazir



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#81

Deadline:
Sunday 12th Jan. 2014.
Release:
Friday 31st Jan. 2014.



Full Circle Team



Editor - Ronnie Tucker
ronnie@fullcirclemagazine.org

Webmaster - Rob Kerfia
admin@fullcirclemagazine.org

Podcast - Les Pounder & Co.
podcast@fullcirclemagazine.org

Editing & Proofreading

Mike Kennedy, Gord Campbell, Robert Orsino, Josh Hertel, Bert Jerred

Our thanks go to Canonical, the many translation teams around the world and **Thorsten Wilms** for the FCM logo.

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



Google Currents - Install the Google Currents app on your Android/Apple devices, search for 'full circle' (within the app) and you'll be able to add issues 55+. Or, you can click the links on the FCM download pages.



Ubuntu Software Centre - You can get FCM via the Ubuntu Software Centre: <https://apps.ubuntu.com/cat/>. Search for 'full circle', choose an issue, and click the download button.



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.



Ubuntu One - You can now have an issue delivered to your free Ubuntu One space by clicking the 'Send to Ubuntu One' button which is available on issues 51+.