



# Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #61 - May 2012



**AUDIO FLUX**  
NEW SECTION  
ON FREE CC MUSIC



## DVD RIPPING & ENCODING FOUR SYSTEMS: TIMED AND TESTED

full circle magazine is neither affiliated with, nor endorsed by, Canonical Ltd.

## HowTo

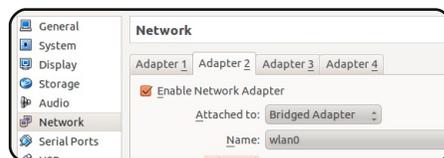


Python - Part 33 p.07

```
op - 15:52:05 up 16 days, 17:33, 2 users, load a
Tasks: 243 total, 4 running, 238 sleeping, 0 stc
Cpu(s): 32.0%us, 25.2%sy, 16.2%ni, 26.3%id, 0.2%wa
Mem: 15949272k total, 15835844k used, 113428k fre
Swap: 6048436k total, 1704k used, 6046732k fre
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM
1980	kevin	20	0	4469m	3.7g	3.6g	S	66	24.6
1050	root	20	0	222m	17k	7000k	B	25	0.2

Use The TOP Command p.10



VirtualBox Networking p.15

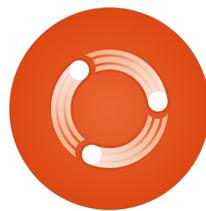


GIMP - Beanstalk Pt2 p.21



Inkscape - Part 1 p.24

## Graphics



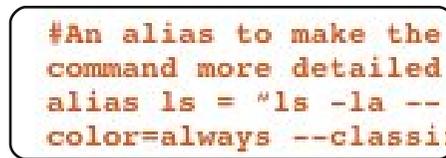
# Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

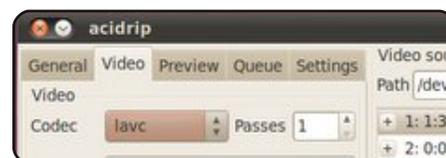
## Columns



Linux News p.04



Command & Conquer p.05



Linux Labs p.29



Ubuntu Women p.XX



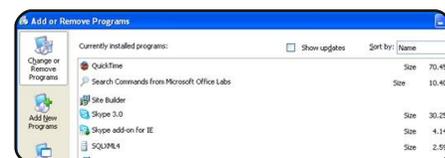
Audio Flux p.52



Ubuntu Games p.53



Q&A p.50



Closing Windows p.32

## Opinions

My Story - Full Circle Magazine

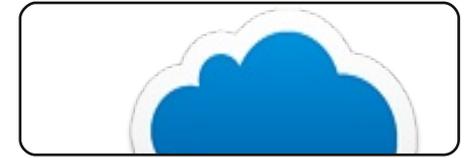
by Ronnie Tucker

Since I can't give away free party hats, or slices of cake, I thought I'd spend a page (or three) showing the process of how an issue of FCM is made.

My Story p.38



My Opinion p.42



I Think... p.43



Book Review p.45



Letters p.46

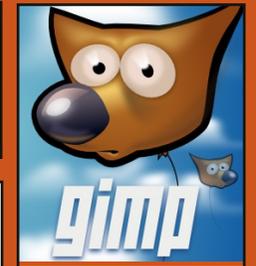


The articles contained in this magazine are released under the Creative Commons Attribution-Share Alike 3.0 Unported license. This means you can adapt, copy, distribute and transmit the articles but only under the following conditions: you must attribute the work to the original author in some way (at least a name, email or URL) and to this magazine by name ('Full Circle Magazine') and the URL [www.fullcirclemagazine.org](http://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.**



This magazine was created using :



## Welcome to another issue of Full Circle!

**H**ave you all finished playing with 12.04 yet? I upgraded my Kubuntu 11.10 to 12.04 a couple of days after the release and, I have to admit, the upgrade was flawless. I had to leave it downloading overnight though, as I assume the servers were being hammered. Top marks to the Kubuntu team. I hope your upgrades went just as well.

No LibreOffice this month, Elmer is a bit busy in the real-world. Instead, you have an excellent tutorial on the TOP command. Top is a task manager app which you run from the command line. Robin Catling (he of the missing podcast) has written a four-part article on VirtualBox networking. Rather than run it across four issues, I've crammed it into this one issue. He's already called me 'nuts' (among other names) for running it in one article. We're adding a village and the beanstalk in my GIMP tutorial, and, while it's a bit early in the year, you'll be making a snowman in the Inkscape tutorial. Instead of a software review, we have a book review this month. Have a read through my review, then click the link to get a discount. Exclusive to Full Circle readers!

Returning like a blast from the past - which he is - is Nicola Cappellini. Nicola was a proofreader many moons ago, and actually came up with, and wrote, our writers' guidelines. He has returned with a music column - where he'll suggest Creative Commons songs (next month), sites (this month) and software.

**All the best, and keep in touch!**

*Ronnie*

[ronnie@fullcirclemagazine.org](mailto:ronnie@fullcirclemagazine.org)

## Full Circle Podcast

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

### **Hosts:**

Robin Catling

Ed Hewitt

Dave Wilkins



<http://fullcirclemagazine.org>



## Ubuntu 12.10 named as "Quantal Quetzal"



The challenge of coming up with a code name for the successor to Ubuntu 12.04 LTS, "Precise Pangolin" which is due to be released this week, has been met by Mark Shuttleworth announcing the new name for 12.10. "Quantal Quetzal" is the name announced in a blog posting by Shuttleworth, who writes that it would be redundant to use "Quality" as part of the code name, as *"Every release has quality first these days"*.

The Quetzal is a family of birds found in western Mexico with mostly iridescent green plumage.

Source: h-online.com

## Tiny PC Runs Linux and Android 4.0--and Costs Just \$74

First we saw the Raspberry Pi and the Cotton Candy devices emerge. The latest to appear [is] the MK802 micro-PC, a USB-sized device priced at \$74.

Featuring a single-core 1.5GHz AllWinner A10 Cortex A8 ARM processor, Android 4.0, 512MB of DDR3 high-capacity memory, and WiFi connectivity.

With a MALI400 graphics processing unit, the device [...] features 4GB Flash storage, a microSD slot, and two USB ports: one full-sized and one micro. Video output is via 1080p HDMI and users can tap either an Android virtual keyboard or add a wireless mouse and keyboard. Perhaps best of all is that users can run Ubuntu, Debian, or another Linux distribution of their choice via microSD card.

Source: pcadvisor.co.uk

## Linux Mint 13 "Maya" RC released!



Available in two editions, Linux Mint 13 features the choice between a productive, stable and mature MATE 1.2 desktop and the brand new modern-looking and exciting Cinnamon 1.4. These two desktops are among the best available, they're perfectly integrated within Linux Mint and represent great alternatives to Gnome 2 users. Linux Mint 13 is also an LTS (Long Term Support) release and it will be supported until April 2017.

Source: linuxmint.com



## EPUB

Finally, we have mobile editions of **Full Circle** on the downloads page. At the moment, there are only a couple of issues online, but we're hoping to have back issues online shortly. If you've any problems with the epub file/formatting, you can drop an email to Jens at:

[mobile@fullcirclemagazine.org](mailto:mobile@fullcirclemagazine.org)

Big thanks to Jens, and the beta testers, for making them a reality.

## Google Currents



Google Currents has been released worldwide, so install it on your Android/Apple devices, search for 'full circle' (in the app) and add issues 55-60 to your app. Or, you can click the links on the FCM download page for those issues.



In last month's article, I posed the question whether or not readers would be interested in a thorough article covering my configuration files. The response I had to this question made me feel like there may be some interest in the files, and so I will be posting links to my configuration files, and explanations of what they do and why I need them. The scripts are in my github here: <https://github.com/lswest/dotfiles/tree/master/C&C> (this folder is not going to be updated unless I have to make corrections – my actual current files will always be in master/).

An example screenshot of my system: <http://lswest.deviantart.com/art/March-2012-Screenshot-289550803>

I won't be going in-depth about all the options I use in my files, since it would result in a very, very long article, and most of it wouldn't apply to every reader. Instead, I feel most of the files are clear (and, where I have discussed them in previous articles, I have

referred back to them). Any specific questions can always be directed to me at the email address listed below.

## Window Manager

The `.xmonad.hs` is the configuration file for my window manager of choice (XMonad), and is unimportant for anyone who isn't using it. I combine this with the `.conkyrc_dwm_bar` file to create my task bar (it was originally used in DWM, hence the name).

### General terminal settings:

The rest of the files are more general. The `.Xdefaults` file contains my colors and settings for `rxvt-unicode`. I also use Zshell, so my `.zshrc` file is a bit more detailed than a similar `.bashrc` file. I do, however, use the Vim-Powerline symbols (see issue #59), so be aware of that if you carry over my PS1 prompt.

### Conky files

(all lua scripts are in `scripts/`): The `.conkyl_mpd` file creates a lua-based conky instance with

Music Player Daemon information. This could be adapted, but it's essentially useless to anyone who doesn't use MPD (though, I don't see a reason why you wouldn't use MPD!).

The `.conkyrc` file is my “main” conky instance with a bit of system information in it. It uses a lua file to create a smooth semi-transparent background, but could easily be used without it.

The `.conkyrc_reminders` file contains a front-end for my python To Do scripts (see issue #46). It formats text documents into a bulleted list using the file names for the section names.

## Further scripts include:

`Wallpapers.sh`, which selects a random wallpaper from a specific folder every 15 minutes (easily extendable).

**Mpd-info:** a script to yank music state & song information, and format them for a single-line conky.

**Dvol:** a script to adjust volume (including an on screen display using `dzen`)

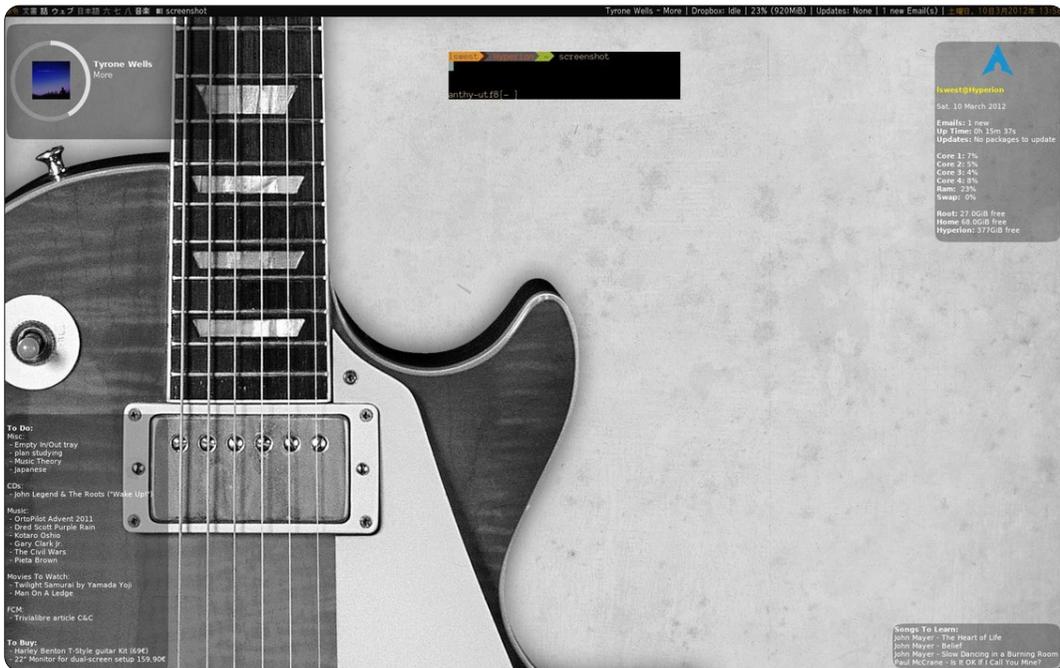
**volstate:** A volume state script that pulls information about the volume (muted, or a percentage of total volume if unmuted).

**Mpd-cover:** A script (not created by me) that downloads album art for a song using information from MPD, and saves it into a `.covers` folder. Works fairly well, unless you have weird symbols or exotic music.

**Packages-short:** An indicator script for new updates (using `pacman`, an alternative with `apt-get` would be:)

**jDate:** a customized Japanese-language date and clock that I use with conky and `dzen` as a taskbar.

That essentially sums up my list of configuration files. It may seem like a lot (or very little, depending on what you're used to), but it covers a rather wide range of settings. Occasionally, I find myself



tweaking a single file or adding a small script to solve a particular issue I have, but otherwise I'm quite content with the current setup. I also find this combination to be the most efficient I've used to date. So, for anyone who's looking for a way to work more efficiently, I highly recommend using conky to display relevant information on the desktop, and a tiling window manager for actual work.

If you have any questions, comments, or suggestions, feel free to email me at [lswest34@gmail.com](mailto:lswest34@gmail.com). If you do

email me, please include "C&C" or "FCM" in the subject line, so that I don't overlook it. For anyone who does use one of my configuration files, consider letting me know if you find a new use for it (or adjust it any way). Who knows? Maybe someone else will find it useful – or maybe I've been trying to do something similar for a long time.



**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](mailto:lswest34@gmail.com).

## CODEWORD

Every number in the codewords grid is 'code' for a letter of the alphabet. At the end you should have a different letter in each numbered box, and a word in English in each of the horizontal and vertical runs on the codeword grid.

22	5	3	19	25	2	26	1		24	22	11	8													
	3		25		24		19		10		3														
2	18	14	3	19	21		22	3	11	24	2	9													
	11		21		19	6	19		21		15														
22	11	18	14		14		21	26	26	20	24	2													
	24				24				2																
25	2	19	11	24	25		7	2	24	24	22	14													
			18				21				14														
22	19	21	19	19	16		19		19	13	26	17													
	16		1		19	22	22		15		2														
22	11	4	18	1	10		4	3	2	2	19	8													
	21		22		18		24		24		15														
23	24	22	14		16	19	22	14	24	2	24	25													
1	2	3	4	5	6	7	8	9	10	11	12	13													
				Q					X																
14	15	16	17	18	19	20	21	22	23	24	25	26													
		M																							
A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z

Solutions are on the second last page.

puzzles are copyright, and kindly provided by, **The Puzzle Club** - [www.thepuzzleclub.com](http://www.thepuzzleclub.com)



This time, we'll set up the Android SDK on our Linux desktop. We'll also create a virtual Android device, install SL4A and python on it, and do a quick test.

Please be aware, this is not something you would want to do for machines that have less than 1 GB of ram. The emulator eats up a huge amount of memory. I've tried it on a laptop running Ubuntu with only 512 MB of ram. It WILL work, but it is REALLY slow.

Here's a quick list of what we'll do. We'll go step-by-step in a minute.

- Install the Java JDK6.
- Install the Android SDK starter pack.
- Create and setup AVDs.
- Test AVD, and install SL4A and Python.

In reality, we should also install Eclipse and the Android ADT plugin for Eclipse, but, since we won't be dealing with Eclipse in this set of articles, we can bypass that. If you want to include those steps, head

over to <http://developer.android.com/sdk/installing.html> to see all the steps in the suggested order. Let's get started.

## STEP 1 - Java JDK 6

From everything I've read and tried, it must be the actual Sun release. OpenJDK is not supposed to work. You can find information on this on the web, but here's the steps that I did. In a terminal, type the following...

```
sudo add-apt-repository  
ppa:ferramroberto/java
```

```
sudo apt-get update
```

```
sudo apt-get install sun-  
java6-jdk
```

Once everything here is done, you will want to edit your .bashrc file to set "JAVA\_HOME" so everything runs correctly. I used gedit and, at the bottom of the file, I added the following line...

```
export  
JAVA_HOME="/usr/lib/jvm/java-  
6-sun-1.6.0.06"
```

Save the file and move on to step 2.

## STEP 2 - Android SDK Starter Pack

Now the actual "fun" begins. You'll want to go to [developer.android.com/sdk/index.html](http://developer.android.com/sdk/index.html). This is where the SDK is located. Download the latest version for Linux, which, at the time of this writing, is android-sdk\_r18-linux.tgz. Using Archive Manager, unpack it somewhere convenient. I put it in my home directory. Everything runs directly from this folder, so you really don't have to install anything. So the path for me is /home/greg/android-sdk-linux. Navigate to this folder, then go to the tools folder. There you will find a file called "android". This is what runs the actual SDK. I created a launcher on my desktop to make it easy to get to.

Now the boring part. Run the android file, and the Android SDK Manager will start. It will go out

and update the platforms that are available. I will warn you now that this process will take some time, so don't bother if you don't have a lot of time to deal with it. For the sake of brevity, I would suggest you get only one platform to start. A good one to begin with is the Android 2.1 platform, since, for the most part, if you develop for an older platform, there should be no problem running on a newer platform. You also need to get the Tools set as well. Simply check the box next to those two items, then click on the install button. Once you get the platform of your choice, and the tool set, you are almost ready to create your first virtual machine.

## STEP 3 - Create and set up your first AVD

Back in the Android SDK Manager, select Tools from the main menu, then select Manage AVDs. This will open a new window. Since this is the first time, there won't be any virtual devices set up. Click on the "New" button. This opens yet another window where



we define the properties of the virtual Android device. Here's the steps that you should use to set up a simple Android emulator device:

- Set the name of the device. This is important if you have more than one device.
- Set the target platform level.
- Set the size of the SD card (see below).
- Set the skin resolution.
- Create the device.

So, In the name text box, type "Test1". Under the target combo-box, select Android 2.1 - API Level 7. In the text box for "SD Card:" enter 512 and make sure the dropdown shows "MiB". Under "Skin", set the resolution to 800x600. (You can play with the other built-in sizes on your own.) Finally, click the "Create AVD" button. Soon, you'll see a message box saying that the AVD was created.

### STEP 4 - Testing the AVD and installing SL4A and Python

Now, finally, we can have a bit of fun. Highlight the AVD you just created and click on the Start

button. In the dialog box that pops up, simply click the "Launch" button. Now, you have to wait a few minutes for the virtual device to be created in memory, and the Android platform to be loaded and started. (We'll talk about speeding this process up in later runs.)

Once the AVD starts up and you have the "home" screen up, you will install SL4A. Using the browser or the google web search box on the home screen, search for "sl4a". Go to the downloads page, and you'll eventually find the web page for the downloads at <http://code.google.com/p/android-scripting/downloads/list>.

Scroll down the page until you get to the sl4a\_r5 link. Open the link and tap on the "sl4a\_r5.apk" link. Notice I said "tap" rather than "click". Start thinking about using your finger to tap the screen rather than clicking the mouse. It will

make your programming transition easier. You'll see the download start. You may have to pull down the notification bar at the top to get to the downloaded file. Tap on that, then tap the install button.



Once the file is downloaded, you'll be presented with the option to open the downloaded app or to tap "Done" to exit the installer. Here we will want to tap "Open".

Now SL4A will start.

You'll probably see a dialog asking if you will agree to usage tracking. Either accept or refuse this - it's up to you. Before we go any farther, you should know some keyboard shortcuts that will help you move around. Since we don't have a "real" Android device, buttons like Back, Home, and Menu, aren't available. You'll need them to navigate around. Here's a few important shortcuts.

Back - Escape

Home - Home  
Menu - F2

Now we will want to download and install python into SL4A. To do this, first tap Menu (press F2). Select "View" from the menu. Now select "Interpreters". It looks like nothing happened, but tap Menu again (F2), then select "Add" from the popup. Now scroll down and select "Python 2.6.2". This will download the base package for Python for Android. Install the package, then open it. You will be presented with four options. Install, Import Modules, Browse Modules, and Uninstall Module. Tap on Install. This will start downloading and installing all the pieces of the latest Python for Android. This can take a few minutes.

Once everything is done, tap Back (escape key) until you get to the SL4A Interpreters screen. Now everything is loaded for us to play in Python on Android. Tap Python 2.6.2, and you'll be in the "standard" Python shell. This is just like the shell on your desktop. Type the following three lines, one at a time, into the shell. Be sure to wait for the ">>>" prompt each time.

```
import android

droid = android.Android()

droid.makeToast("Hello from
Python on Android")
```

After you type the last line and press Enter, you'll see a rounded corner box at the center bottom of the shell that says "Hello from Python on Android". That's what the "droid.makeToast" command does.

You've written your first Python script for Android. Neat, huh?

Now let's create a shortcut on the Android home screen. Tap the Home key (Home button). If you chose the 2.1 platform, you should see a slider bar on the far right of the screen. If you chose another platform, it might be a square or rectangle consisting of small squares. Either way, this gets you to the Apps screen. Tap that, and find the SL4A icon. Now perform a "long tap" (long click), which will create a shortcut on the Home screen. Move the shortcut wherever you want it.

Next, we will create our first saved script. Go back into SL4A.

You should be presented with the sample scripts that come with Python 4 Android. Tap the Menu button and select "Add". Select "Python 2.6.2" from the list. You'll be presented with the script editor. At the top is the filename box with ".py" already filled out. Below that is the editor window that already has the first two lines of our program entered for us. (I included them below in italics so you can check it. We also used these two lines in our first sample.)

```
import android

droid = android.Android()
```

Now, enter the following two lines to the python script.

```
uname =
droid.dialogGetInput("What's
your name?")

droid.makeToast("Hello %s
from Python on Android") %
uname.result
```

The first new line will create a dialog box (droid.dialogGetInput()) that asks for the user's name. The response is returned to our program in uname.result. We've already used the droid.makeToast() function.

Name the file andtest1.py, then tap Done, and tap "Save & Run". If everything worked, you should see a dialog box asking for your name. After you enter it, you should see the alert at the bottom of the screen saying "Hello Your Name from Python on Android".

That's all for this time. For now, there's a TON of documentation about SL4A for free on the web. You can play a bit on your own until next time. I'd suggest that you start by going to <http://code.google.com/p/android-scripting/wiki/Tutorials>.



**Greg** is the owner of RainyDay Solutions, LLC, a consulting company in Aurora, Colorado, and has been programming since 1972. He enjoys cooking, hiking, music, and spending time with his family. His website is [www.thedesignedgeek.net](http://www.thedesignedgeek.net).



O'Reilly are looking forward to celebrating Velocity's 5th Year with you **June 25-27**, at the **Santa Clara Convention Center**. You'll meet the smartest people working in web performance and operations at the O'Reilly Velocity Conference. Web and mobile users expect better performance than ever before. To meet, and exceed, their expectations, you need to master a daunting array of web performance, operations, and mobile performance issues. Velocity offers the best opportunity to learn the newest info on what you need to know to build a faster and stronger web.

Take advantage of this rare opportunity to meet face-to-face with a cadre of industry leaders who are taking web performance and operations to the next level. Velocity packs a wealth of big ideas, know-how, and connections into three concentrated days. You'll be able to apply what you've learned immediately and you'll be well prepared for what lies ahead with four in-depth tracks covering the key aspects of web performance, operations, mobile performance, and Velocity culture.

Velocity has sold out the last two years, so if you want to reserve your spot at Velocity 2012, register now and save an additional 20% with code **FULLCIR**.



# HOW-TO

Written by Kevin O'Brien

# Use The TOP Command

One of the great advantages of using Linux is that there are some great tools available to help you understand what is going on with your computer and diagnose possible problems. One of the most useful is the top command. I am going to cover some of the things you can do, and maybe mention one or two other commands as well.

First of all, just to get it out of the way, there is an alternative called htop, and I do plan to cover it later. But htop generally needs to be installed before you can use it, while top should already be on your system, making it a good starting point. Usage of the command is simple: just open a terminal/console and type top. The result will be something like the image shown right.

There is a lot of information on this screen, so it will take us a little time to go through all of the options. What you can see right away is that this is listing processes running on your

```
top - 15:52:05 up 16 days, 17:33, 2 users, load average: 1.87, 1.75, 1.73
Tasks: 243 total, 4 running, 238 sleeping, 0 stopped, 1 zombie
Cpu(s): 32.0%us, 25.2%sy, 16.2%ni, 26.3%id, 0.2%wa, 0.0%hi, 0.2%si, 0.0%st
Mem: 15949272k total, 15835844k used, 113428k free, 308704k buffers
Swap: 6048436k total, 1704k used, 6046732k free, 8081796k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1980	kevin	20	0	4469m	3.7g	3.6g	S	66	24.6	16951:08	VirtualBox
1058	root	20	0	232m	47m	7232	R	25	0.3	822:53.24	Xorg
14712	kevin	20	0	371m	35m	16m	R	16	0.2	3886:56	gtk-gnash
11545	kevin	20	0	371m	35m	15m	S	16	0.2	359:39.98	gtk-gnash
11442	kevin	20	0	722m	239m	32m	S	9	1.5	170:03.41	firefox-bin
11719	kevin	20	0	344m	19m	14m	S	8	0.1	176:54.61	gtk-gnash
4293	kevin	20	0	975m	97m	28m	S	6	0.6	786:30.65	plasma-desktop
8514	kevin	20	0	265m	22m	17m	S	2	0.1	0:00.26	ksnapshot
6745	kevin	20	0	423m	28m	18m	S	1	0.2	0:01.40	konsole
11585	kevin	20	0	341m	17m	12m	S	1	0.1	15:00.51	gtk-gnash
11608	kevin	20	0	343m	17m	12m	S	1	0.1	18:46.56	gtk-gnash
11657	kevin	20	0	342m	17m	12m	S	1	0.1	14:49.67	gtk-gnash
11718	kevin	20	0	343m	18m	13m	S	1	0.1	14:31.60	gtk-gnash
12488	root	20	0	451m	42m	25m	S	1	0.2	70:40.60	krusader

computer, and you can see the Process ID for each one, how much RAM each one is using, what percentage of the CPU each one is using, the owner of each process, etc. Then you can see all sorts of cryptic numbers above this listing. We will cover all of it either in this article or one to follow, but to get there we need to get going!

## CPU Usage

By default, top lists processes in order of the amount of CPU each one is using, expressed as a

percentage of the total available. This is important to know, since if your CPU is maxing out you will see degraded performance. This can show up as lags in responding to keyboard and/or mouse input, jerkiness on audio or video playback, etc. On my Kubuntu desktop, I have a side panel set up with monitors for CPU usage, CPU temperature, Memory usage, Swap usage, and network traffic - so that I can monitor these critical functions and prevent problems from getting out of hand. I have seen situations where the CPU

usage maxed out at 100% and stayed there (usually as a result of Flash, which cannot die soon enough, but that is a rant for another day). When that happens, the top command lets me quickly check and see what application is problematic so I can kill it.

One of the nice things about the top command is that it is interactive as long as you have it up in the terminal. So, you can kill a process quite easily by simply typing a k with the terminal open and top running. This will open a prompt above the process list asking you which process to kill. Just enter the Process ID of the misbehaving application and it will be gone.

Now, if you are looking at the screenshot of top running on my computer you may have noticed something. I said that it gave the CPU usage as a percentage of the total. And if you looked carefully you might have seen that the percentages add up to more than 100%. How can that be, you ask? Well, the answer is that it is

looking at these as percentages of the core that the process is running on. Since this computer is a dual-core machine, it has two processors and can distribute individual processes to whichever core it wishes. So I could theoretically see up to 200% if I added up the numbers here (though that would be bad since it would indicate I was maxed out). If I had a quad-core, I could have up to 400%, etc.

## Priority and Niceness

The idea of niceness is to determine which processes should get more goodies when running, and which should be put in the background. In other words, to set some priorities of access to the CPU. This is done by using a

niceness number, which appears in the column NI. In the screenshot you see that all of these processes are running at a niceness number of 0. What that means is that they are running at a default priority which has not been altered in any way. Niceness numbers run from -20 to +19, with -20 being the highest priority. I said that 0 is the default, but you can check it on your system by running the command `nice` without any arguments. What is returned is the default niceness level. I will leave this topic here for now, but if you want to know more there is a good web page on this topic at <http://www.ibm.com/developerworks/linux/library/l-lpic1-v3-103-6/>. This article will explain how you can change niceness levels for certain processes if you wish to do so.

## Memory

Next to the column on CPU usage in the screenshot is a column for memory usage, again expressed as a percentage of the total available. In this case, it happens that the process using the most CPU is also the one using the most memory, which is not unusual. But suppose you wanted to see your processes sorted in the order of the memory they consume? Well, as I mentioned before, the top command is interactive. To change the sort order, just press an upper-case letter O while the command is running in the terminal. This brings up a very useful screenful of sort options (shown below).

With this screen open you can

press a lower case letter n, then enter, and get a listing in order of memory percentage used. Or you can sort in other ways if needed.

Again, this is useful if you find you are running out of memory and need to know where it is going. If one process is using a lot of memory unexpectedly, that would definitely be an indication. That does imply you have some idea of what constitutes normal in these situations. The best way to build a sense of that is to check periodically, and observe what is going on. In my case, I build my machines with 16GB of RAM these days, so I don't expect to see very high percentage usage in most cases. For instance, right now I have my bottom panel filled with program icons for programs I have open (18 of them right now), and a

```
Current Sort Field: K for window L:Def
Select sort field via field letter, type any other key to return 

a: PID          = Process Id          v: nDRT        = Dirty Pages count
b: PPID         = Parent Process Pid  w: S           = Process Status
c: RUSER       = Real user name      x: COMMAND    = Command name/line
d: UID         = User Id             y: WCHAN      = Sleeping in Function
e: USER       = User Name           z: Flags      = Task Flags <sched.h>
f: GROUP       = Group Name
g: TTY        = Controlling Tty
h: PR         = Priority
i: NI         = Nice value
j: P          = Last used cpu (SMP)
* K: %CPU     = CPU usage
l: TIME      = CPU Time
m: TIME+    = CPU Time, hundredths
n: %MEM     = Memory usage (RES)
o: VIRT     = Virtual Image (kb)

Note1:
If a selected sort field can't be
shown due to screen width or your
field order, the '<' and '>' keys
will be unavailable until a field
within viewable range is chosen.

Note2:
Field sorting uses internal values.
```

```
top - 15:52:05 up 16 days, 17:33, 2 users, load average: 1.87, 1.75, 1.73
Tasks: 243 total, 4 running, 238 sleeping, 0 stopped, 1 zombie
Cpu(s): 32.0%us, 25.2%sy, 16.2%ni, 26.3%id, 0.2%wa, 0.0%hi, 0.2%si, 0.0%st
Mem: 15949272k total, 15835844k used, 113428k free, 308704k buffers
Swap: 6048436k total, 1704k used, 6046732k free, 8081796k cached

  PID USER      PR  NI  VIRT  RES  SHR  S  %CPU  %MEM    TIME+  COMMAND
 1980 kevin    20   0 4469m 3.7g 3.6g S   66 24.6 16951:08 VirtualBox
 1058 root       20   0 232m  47m 7232 R   25  0.3  822:53.24 Xorg
14712 kevin    20   0 371m  35m 16m R   16  0.2   3886:56 gtk-gnash
11545 kevin    20   0 371m  35m 15m S   16  0.2   359:39.98 gtk-gnash
11442 kevin    20   0 722m 239m 32m S    9  1.5 170:03.41 firefox-bin
11719 kevin    20   0 344m  19m 14m S    8  0.1 176:54.61 gtk-gnash
 4293 kevin    20   0 975m  97m 28m S    6  0.6 786:30.65 plasma-desktop
 8514 kevin    20   0 265m  22m 17m S    2  0.1   0:00.26 ksnapshot
 6745 kevin    20   0 423m  28m 18m S    1  0.2   0:01.40 konsole
```



# HOWTO - USE THE TOP COMMAND

quick scan of the output of top shows I am using somewhere in the 35-40% range of my total memory.

In the screenshot, you do see one big memory hog, but that is actually expected. I had VirtualBox open and running a virtual machine at the time, and I had configured it to use 4GB of RAM (plus other resources, of course.) So, in this case, I did see what I expected to see. But if I saw Firefox using that much memory, I would know it was a problem and I would shut it down promptly.

## q for Quit

If you have top running in your terminal, you might want to know how to get out of it, and that is simple, just press the letter q (for Quit), and you will be back at your terminal prompt. You can get more information in either of two ways: the old-school way is to type man top in the terminal, but the new, improved, way is to type info top. Though I think you will find the same result either way. The point is that this is a rich command with a lot of options

## Interpretation of

```
kevin@kimball:~$ uptime
14:50:28 up 17 days, 16:32,  2 users,  load average: 1.63, 1.85, 1.92
kevin@kimball:~$ █
```

## System Data

We've looked at some basics of the top command, and focused on looking at the process listings to spot and correct possible problems. This is still a very useful thing, of course, but there is also a lot of system data in the output that is useful. Recall the screenshot we used last time to display the output of this command (shown below)

Now we want to focus on those numbers on the top that are presenting some very useful system data. So let's start at the top (literally):

```
top - 15:52:05 up 16 days, 17:33,  2 users,  load average: 1.87, 1.75, 1.73
Tasks: 243 total,  4 running, 238 sleeping,  0 stopped,  1 zombie
Cpu(s): 32.0%us, 25.2%sy, 16.2%ni, 26.3%id,  0.2%wa,  0.0%hi,  0.2%si,  0.0%st
Mem: 15949272k total, 15835844k used,  113428k free,  308704k buffers
Swap: 6048436k total,  1704k used,  6046732k free,  8081796k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
1980	kevin	20	0	4469m	3.7g	3.6g	S	66	24.6	16951.08	VirtualBox
1058	root	20	0	232m	47m	7232	R	25	0.3	822.53.24	Xorg
14712	kevin	20	0	371m	35m	16m	R	16	0.2	3886:56	gtk-gnash
11545	kevin	20	0	371m	35m	15m	S	16	0.2	359:39.98	gtk-gnash
11442	kevin	20	0	722m	239m	32m	S	9	1.5	170:03.41	firefox-bin
11719	kevin	20	0	344m	19m	14m	S	8	0.1	176:54.61	gtk-gnash
4293	kevin	20	0	975m	97m	28m	S	6	0.6	786:30.65	plasma-desktop
8514	kevin	20	0	265m	22m	17m	S	2	0.1	0:00.26	ksnapshot
6745	kevin	20	0	423m	28m	18m	S	1	0.2	0:01.40	konsole
11585	kevin	20	0	341m	17m	12m	S	1	0.1	15:00.51	gtk-gnash

## Line 1, the Top line

On the first line (above), we have the uptime. Actually, this is information you could get using the uptime command as well:

So this is a clue that the top command is gathering information that is available individually from other commands and bringing it together in one package of awesome goodness. Very convenient that is.

So in order we see that the time is 15:52:05, and the computer has been up over 17 days. It has two users right now, which is normal.

One of the users is root, but you should never run as root for ordinary activities. That is a dangerous and insecure practice. As long as you are logged in as root, any software that runs on your system has root-level privileges. The preferred way to run is to create an ordinary user who does not have quite so high a level of rights, and run as that user. In this case, that user is kevin. By the way, Ubuntu makes it very difficult for you to do anything other than this procedure. If you need higher rights, you use the sudo command to give yourself temporary privileges.

The last part of this line is called load. These three numbers are giving the load for the previous 1, 5, and 15 minutes. But what is load? It is the average number of processes that are runnable, or are uninterruptible. Basically, without getting too technical, it is how occupied the CPU is most of the time. But the wrinkle is that it is not adjusted (normalized) for the number of CPUs. What this means is that a single CPU system with a load of 1 is loaded all of the time. But on my dual core system, I never got to 2, so I am OK. If you had a quad-core, the magic number

# HOWTO - USE THE TOP COMMAND

would be 4, and so on.

## Line 2, Tasks

There is nothing interesting to see here. That last category, zombie, sounds like it ought to be at least interesting, but it really isn't. Zombie processes are runs that have finished running, and will shortly be closed.

## Line 3, CPU(s)

This is worth a look or two. Last time we looked at the processes on the bottom of the top command's output to see if any one process was hogging things. On this line, instead of looking at the individual processes, we are looking at the total picture of what is going on. And here we don't need to worry about how many cores we have, these numbers aggregate all of the data for all cores.

The first statistic is %us, which in this case is 32.0%. This is the percentage of the CPU cycles that are taken up by user processes. This does not necessarily mean ones that a person started, they can mean processes kicked off by

Apache, MySQL, etc. If this percentage is very high, it can be an indication of a problem, since we have other demands to consider. For example, the next statistic is %sy, which is the percentage of CPU cycles that are taken up by the kernel and by other system processes. Obviously you need to have some cycles available for this or you won't have a functioning computer. The third one, %id, is percentage of time the CPU is idle, and the higher the better here (within reason, you need to actually use the computer!). As long as you have some reasonable idle time available, you probably don't have a problem. You can double-check this by looking at the fourth statistic, %wa. This is the percentage of time that a process had to wait for access to the CPU. In this case, .2% is good. You won't be likely to see this at 0.0% too much, since, by the nature of computing, processes are competing for CPU time, but a high number here would definitely indicate a problem.

The rest of the statistics are pretty ignorable, as they deal with really obscure issues, but you can look them in the man page for top.

## Lines 4 & 5, Memory and Swap

These two lines are best addressed together, since you need to combine this information to tell a complete story. What we need to know is how much memory is being used, and how much is available, at any one time. This is important because lack of RAM is the most common cause of a slow, sluggish computer. This can sometimes look like a different problem altogether, which is why it is important to look at the actual data. For instance, if you noticed your hard drive was constantly "chattering" (known as thrashing), you might think you had a hard drive or I/O problem, but in fact this is most commonly caused by a lack of RAM. When there is not enough RAM to hold all of the program code and data currently in use, some of it gets copied out to the hard drive (called paging) to free up space for other code and data. The place where this data gets copied is called the swap area. So when your hard drive is constantly thrashing, it usually means that code and data is constantly being written to and

read from the swap area, and more RAM would eliminate this problem.

Now, one of the things you need to understand to interpret this data is that writing to the hard drive and reading from it is approximately 4 gazillion times slower than reading and writing to RAM. So you want to minimize the use of swap for performance reasons. But because RAM is so much faster than the hard drive, the operating system will prefer to use it whenever possible. One way to speed things up is to keep code in memory even when you have closed the program. After all, you might open it up again, and pulling it from RAM will speed it up a lot. So the operating system caches a lot of code in RAM that is not currently being actively used. Because of this, the reported RAM usage will look like you are on the verge of running out, but this may not be the case. You need to look at all of the data to assess this.

In this case, we start off by noting that this machine has 15,949,272k of RAM. In other words, 16GB, which I knew because that is what I installed in this box. And the next number says that practically all of this 16GB is being

# HOWTO - USE THE TOP COMMAND

used. Is this a problem? Not really. If you look at the second line, you see that I have 6GB of swap space, but hardly any of it is being used (I am using just under 2MB of RAM here). And the last number tells the story. Of my 16GB of RAM, fully half of it, 8GB, is being used to cache code. If I wanted to open a program that was already in the cache, great, the code is already there and it will open quickly. If I want to open some other program, the operating system will delete some of the code that is in cache to free up the space, so there is no problem.

## Htop, the Alternative

I actually prefer top, but some people like htop better, and I think you will see why. For some things it is easier to work with, particularly if you need to do some tasks related to processes. But note that it is not installed by default on many systems, so you will need to install it first. On Ubuntu machines, install it with

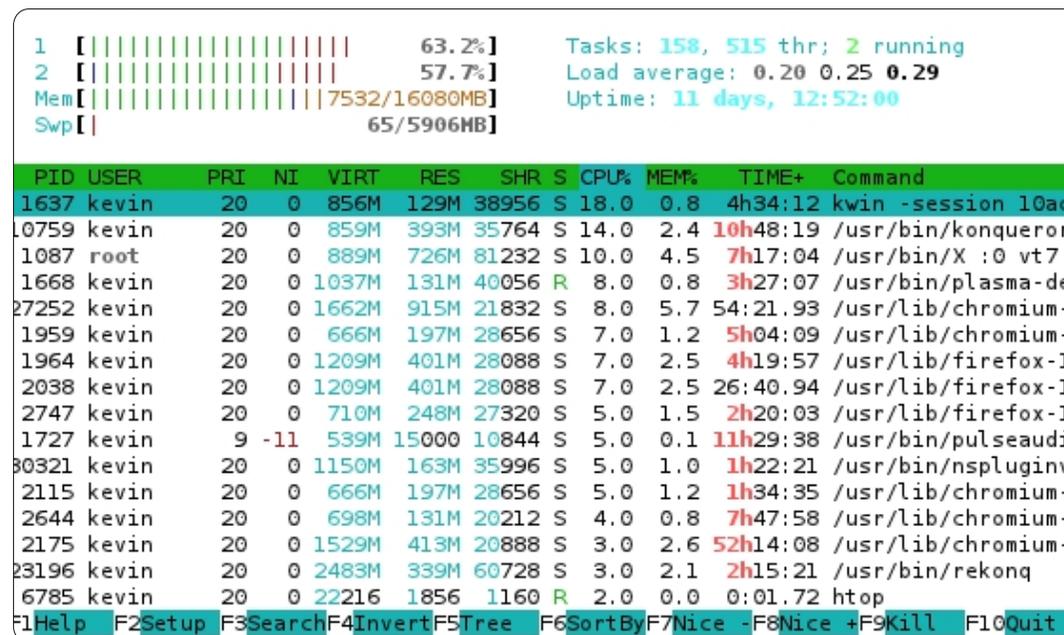
```
sudo apt-get install htop
```

First, you can see that it presents much the same data on

individual processes as the top command. Processes are still listed in order of CPU usage by default, you still see the process ID, User, CPU%, and MEM%, just as before. You can see the command that launched the process, instead of just the program name. Unlike top, htop lets you scroll horizontally using the arrow keys.

On htop, you do have one interesting addition, which is a separate graphical display of the CPU usage for each CPU or core that you have, in this case 1 and 2 since it is a dual-core machine. And you can see the memory and swap usage in ways that you might find easier to read. Uptime, Loads, and Tasks, are shown on the top right.

The real advantage of htop comes when you want to do something to one or more of your processes. You simply use the up and down arrow to highlight the process, then use one of the function keys shown on the bottom. For instance, if you highlight a process and then press F9 you will kill the process. Pressing F7 (Nice -) will lower the nice number, thus increasing the priority (yes, this is not intuitive). And pressing F8 will reduce the



priority by raising the nice number. But be aware that to give a really high priority to a process you would need to have root access, perhaps by using the command

```
sudo htop
```

Personally, I don't have much reason to mess with this, but the worst that could happen is that you would need to reboot your computer if you really screw it up.

Other function keys let you quickly change the sort order, change the field to sort on, and so on.

In summary, I think htop is very useful, but I tend to use top more often for two reasons. First is that I like the more detailed information it gives me. And second is that I know it will be available on any system I am likely to sit down to, while htop will need to be installed, and that means a working Internet connection, which I might not have. But, in general, these two commands do much the same thing, and are a crucial addition to your Linux tool kit.



This all started when I needed to run a sandboxed instance of Wordpress. The simple part was creating a virtual Ubuntu 11.10 server as a guest inside my VirtualBox host. Then I discovered I knew much less about networking than I thought, a topic full of packets and switches and routing. I know even less about networking in VirtualBox, where all that stuff is virtualized in software. Several attempts and a networking crash-course later, mission accomplished, and, to save you my pain, here's what I found.

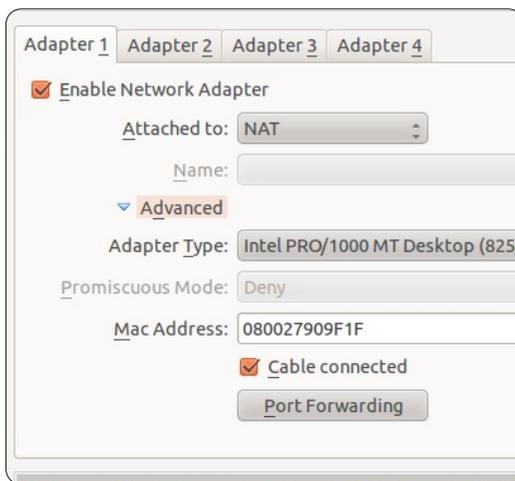
## VirtualBox Network Options

Install any version of VirtualBox from 3.0 onwards, and you'll find similar layouts and capabilities.

For any of your installed virtual machines, click on Settings > Network. This is the default view. Any virtual machine can have up to four network adapters - which you can enable as needed. Most of the

time, you'll need just one. Usually when you create a virtual machine, you can leave VirtualBox to get on with creating a single network adapter - allowing you to see out onto the Internet.

You might need multiple adapters of different types, or the same type with different settings, according to the use of your virtual machine and the networks - physical and virtual - to which it is connected.



For my virtual web-server running Wordpress, things get a bit more complex, so here we go... The Adapter Type defines the virtual

hardware available to your virtual machine (VM). VirtualBox will do its best to mediate the virtual network card with whatever physical adapters you have on your host machine. Open the Advanced pane and you get additional options for setting up your network adapter. I'm going to run through the settings deliberately out of sequence, starting with the Adapter type.

**Adapter Type** : VirtualBox will attempt to emulate some common types of network cards for which it has the drivers and protocols. PCnet-FAST III is the default selection, although I often select the Intel PRO/1000MT adapters if I'm looking for better compatibility with Intel hardware. There are two server versions for higher throughput if you intend your VM to run hot in a production environment. Desktop users needn't worry. If you have a problem with one of the Adapter types, you can try another, including a PCnet-FAST II for older machines.

**Mode** : the rather racy sounding 'Promiscuous Mode' has to do with the behaviour of the virtual switch used by Internal Networking; Bridged and Host Only. You can make a VM's network port accept traffic targeted for other virtual machines, or even the host, in addition to traffic addressed specifically to it. Promiscuous Mode is something 99% of users should never touch, and is there in cases where you need to diagnose the cause of certain network problems.

**MAC Address** : short for Media Access Control, the MAC address is a hardware address that uniquely identifies each network resource or node of a network. It is a unique identifying number for something that is usually hard coded into physical network adapters. VirtualBox generates the MAC addresses of virtual machine network cards.

There is a little button on the right to regenerate the MAC address; should you clone a virtual machine, it will need its own

unique MAC to run on your virtual network alongside the original.

The Cable connected check-box does the same thing as plugging or unplugging a physical cable. This is the setting that connects your virtual network adapter to your virtual network. It is different from the top-most check-box for Enable Network Adapter, which allows you to effectively insert or pull the network adapter from the machine.

Finally, the Port Forwarding button opens another dialog allowing you to define the traffic rules for the adapter; how traffic of certain types is routed between host and guest. This applies only to certain network attachments which we will look at next. This is determined by the Attached to option (with its little pal, the Name option), which is the sharp end of networking in VirtualBox, and the thing that gave me most trouble at the start of my experiment.

## 'Here be dragons.'

There are four types of network attachment available, and many possible combinations of settings

across the other controls. This is where VirtualBox networking can seem like medieval sorcery. Following through the right selection for your given purpose can be simple. There may even be multiple right answers to achieve your goal. But everything else looks doomed to failure: I know...

## Attached to Network types

We have four workable options here:

- Network Address Translation, which is the default,
- Bridged,
- Host Only, and
- Internal Network.

'Not attached' is also a type, but used for the purpose of keeping an adapter in place for troubleshooting. In this mode, VirtualBox reports to the guest that a network card is present, but that there is no connection.

## Network Address Translation (NAT)

NAT enables the guest machine to see out onto the Internet, but

via a private IP address that cannot be seen from the host, or indeed, the rest of your physical network. It will allow you to browse the web, download files, and view e-mail inside the guest, but the outside world will never be able to communicate with the guest machine directly.

When a guest machine sends an IP packet to some remote machine, the VirtualBox NAT service will intercept the packet, extract the TCP/IP segments, change the IP address to the IP address of the host machine, then send it. The outside world only sees the IP address of the host machine. Replies are received by the host machine and sent on to the Guest Machine.

For example, on your home network, your host and other physical machines will typically have addresses starting in the 192.168.x.x range. In VirtualBox, NAT adapters will be begin at 10.0.2.1, incrementing addresses up to 10.0.2.24 in what's called a sub-net. This is not usually routed onto the main network, so this sub-net will be inaccessible from your host. Your guest is able to see out onto the Internet for software

updates and web-surfing, but is invisible to the rest of your network.

The VirtualBox manual is a little more explicit:

"In NAT mode, the guest network interface is assigned to the IPv4 range 10.0.x.0/24 by default, where x corresponds to the instance of the NAT interface +2. So x is 2 when there is only one NAT instance active. In that case, the guest is assigned to the address 10.0.2.15, the gateway is set to 10.0.2.2 and the name server can be found at 10.0.2.3." (Oracle Corporation, 2012, Chapter 9)

NAT is useful when you don't really care what IP addresses your guest machines use, so long as they are unique. You will need to do more configuration when you need to forward traffic or expose services like a web-server to the outside world. Additionally, you cannot enable file and folder sharing over the network.

## Bridged Adapter

Under the Bridged Adapter, your virtual machines behave as any other computer on the

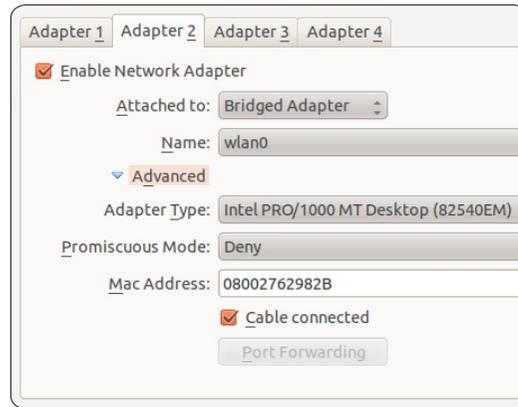
# HOWTO - VIRTUALBOX NETWORKING

network where the hosting system resides; the bridged adapter bridges the virtual and physical networks. The outside world can directly communicate with the guest machine.

The Bridged Adapter connects through the host to your default network device that allocates IP addresses for your physical network. VirtualBox connects to one of your installed network cards and exchanges network packets directly; it bridges the virtual and physical networks. In normal use, it will try to get a standard 192.168.x.x IP address from your router, so that your virtual machine can look like a perfectly respectable physical device alongside all the other ironmongery on your network.

You may have more than one active network device available to your host; for example, my laptop has a wired ethernet port (called eth0) and a wireless card (called wlan0). The Name options allows you to choose which one you want to bridge with VirtualBox.

In my case, I use the Wireless adapter wlan0, because that is connected to the router, while



eth0 doesn't even have a cable.

So my host connects as 192.168.0.2, allocated by the router. My bridged virtual machine pops up and requests its own IP address which the router allocates as 192.168.2.6, neither knowing nor caring that VirtualBox is handling the pass-through of traffic via the host. My virtual machine becomes just another device on the local network. If I count my host and three VM's under Bridged, then I have four machines visible on my physical network.

## Come here, there's more...

NAT is useful because it protects our guest systems from the Internet at large, but in order

to access them, we need to set up port forwarding to access the guests from the host (I have servers on some of my guests). Bridged Networking allows access, but the guests are exposed.

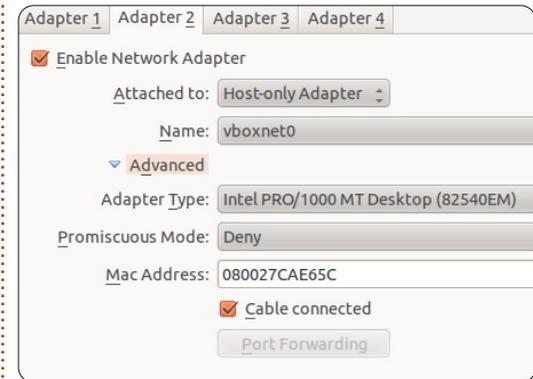
If your network access device (be it a router, switch, or hard configuration by your ISP) will not allocate more than one IP address, you probably can't use Bridged networking.

## Host-only Adapter

Under the Host-only adapter, virtual machines can communicate between each other and with the hosting system but not outside. The Host-only adapter uses its own dedicated network device, called vboxnet0, to set up a sub-net and allocates IP addresses to guest machines. The guest machines cannot communicate with the outside world, since they are not connected via a physical interface. Host-only provides restricted services, useful for creating private networks under a VirtualBox host for its guests.

Unlike other virtualization products, the VirtualBox NAT adapter doesn't bridge the default

network device on your hosts, so there is no direct network access to NAT-ed machines and no access to applications and data on the host itself. Let's look at an example.



Typically your host has its own network address, the one it uses to access the Internet - commonly 192.168.0.101. Under Host-only, the host machine also becomes the VirtualBox router, with the default IP address of 192.168.56.1. The Host creates an internal local area network serving all the guest machines set-up for Host-only, visible to the rest of your network. The vboxnet0 adapter starts issuing IP addresses from 192.168.56.101 onward, but you can change the default IP address allocation, if you want.

Similar to the Bridged adapter, Host-only uses different address

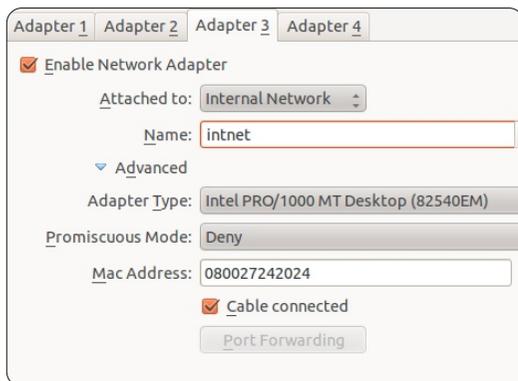
ranges. You can easily allow guest machines to obtain addresses using the DHCP (dynamic allocation, which will likely be a different address per session) that VirtualBox provides.

Additionally, the Host-only network for the host and guest does not have a gateway out to the Internet. It's used to just connect the host and the guest, much like a network switch. Therefore Host-only adapter does not provide guest machines with Internet access; vboxnet0 has no default gateway, so while the addition of vboxnet0 greatly simplifies networking between the Host and guests using it, you don't get external access or port forwarding, so you may still need a second NAT or Bridged adapter attached to your guest to achieve full access.

## Internal network

If you want several guests machines to communicate with each other on one host, but with nothing else, then we can use this Internal network mode. Although you can use Bridged Networking for this, Internal Network is more secure. In Bridged networking, all

the packets are sent/received from the physical network adapter in the Host machine, and the traffic can be tapped (say, by attaching a packet sniffer to the Host).



The Internal network option creates, according to the VirtualBox manual, "a software-based network which is visible to selected virtual machines, but not to applications running on the host or to the outside world." This provides a network containing the host and a set of virtual machines, but none of it goes through the host's physical network interface - it is entirely virtual, with VirtualBox acting as a network switch. What you get is a private LAN for your guest machines only, without any access to the external world, which makes it very secure. Possible uses might be running a top-secret development server and clients,

conducting penetration testing or otherwise creating a secure Intranet for a team or organisation. It's an ideal way to lock down an environment against unauthorised software installs, downloads, uploads and Facebook-ing during work time.

This is where you begin to see the different types of network setups come into their own, being fit for different purposes. Which brings me back to my start-point; creating a virtual server for Wordpress development and testing.

## Guest Access to the Internet

First, I need my guest machine to be able to access the internet for installing updates, downloading packages, and the like. I need to access it from my host machine, but I don't want the server to be accessible from anywhere else on the network.

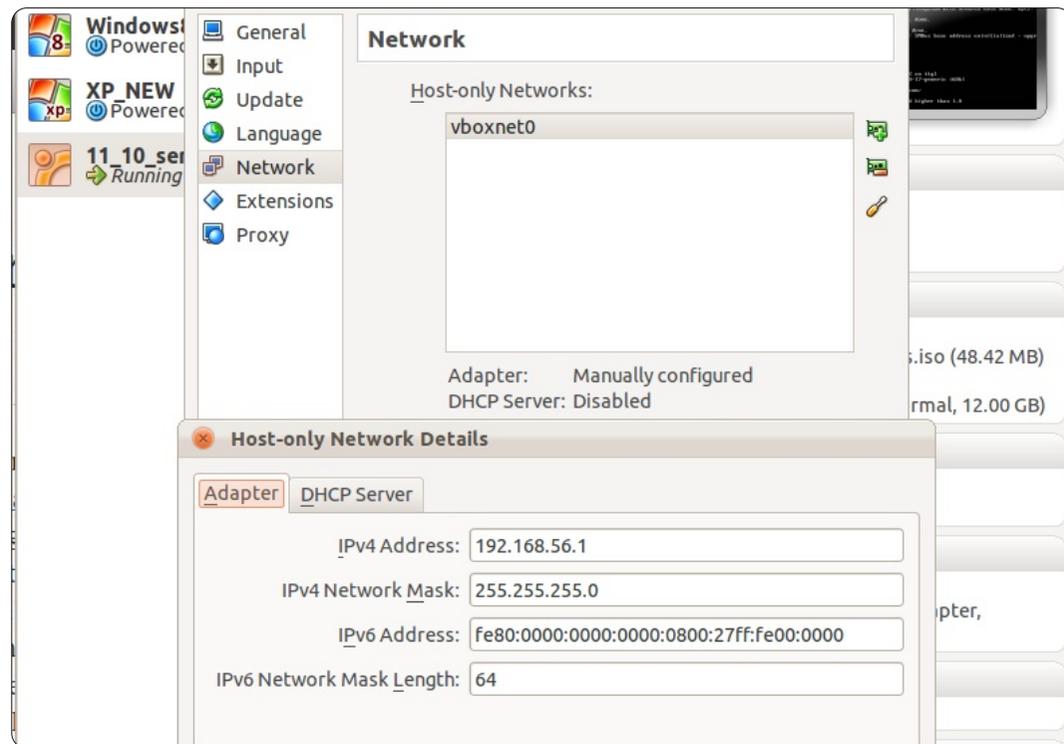
Selecting my guest machine from the VirtualBox Manager window, I then select Settings > Network.

I begin by keeping the default networking NAT adapter. This allows the guest machine to access the Internet through my host's connection. The guest isn't visible to anyone else on my network; I cannot access any guest resources from my host machine either, nor can any guest machines access each other.

## Configure Host-only Adapter in VirtualBox Manager

The easy method is to use the default Host-only adapter provided by VirtualBox, this is the one named vboxnet0, but you can add more if you want several separate Host-only networks. If you go to the VirtualBox Manager main window and select File > Preferences > Network, you can configure vboxnet0 or add new adapters.

Click on the "Edit" button for your Host-only network, that's the screwdriver icon titled "Edit host only network," and you'll find the default settings. The IP address 192.168.56.1 is the address at which the guests can access the host. I'll leave this as-is.



By default, vboxnet0 has a DHCP server set up to dynamically allocate IP addresses on a per session basis. I want my Wordpress virtual server to keep static IP addresses, so on the DHCP tab, I uncheck the “Enable Server” box for DHCP.

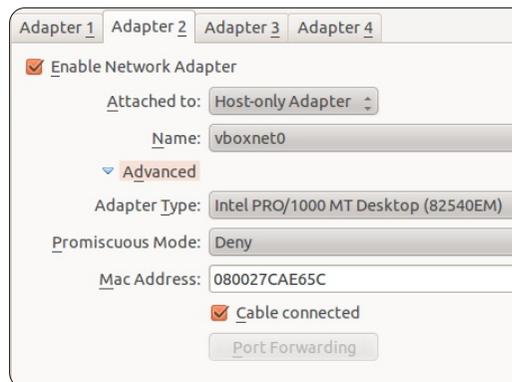
## Add Host-only Adapter to Guest

I'll add another network adapter to my guest machine, a Host-Only Adapter, that enables a self-contained, private, virtual

network consisting solely of my host and any guests using host-only. Any of the member machines can access each other, but nothing outside it can get in.

So on my Adapter 2 tab for my guest, I set the Attached to option to Host-only, and the name is the default vboxnet0. Promiscuous mode doesn't matter, but I do need the Cable Connected checkbox enabled.

## Configure the Guest



I need my guest's virtual server to have a static IP address on the host-only network, otherwise the address changes for every session and I have to work out what it is before I can access it from the host. Logging into the guest, I open a terminal and issue the following command:

```
ifconfig eth1 192.168.56.101 netmask 255.255.255.0 up
```

This sets the IP address on my host-only adapter for this guest, and brings up the network interface. It uses the 192 network prefix for domestic networks, followed by .168 for the start of the host ID, .56 as the host-only sub-net range, and I could use any final identifier above .1 (my host) on my host-only network. The NAT adapter is sitting on eth0, so logically the Host-only adapter is bound to eth1. I can now secure

shell (SSH) or browse to the guest from the host using this IP address as a test.

However, this is just temporary; if I reboot, this configuration will disappear. To make it permanent, I need to add it to the /etc/network/interfaces file (as root) in the form:

```
# The host-only network interface
auto eth1
iface eth1 inet static
address 192.168.56.101
netmask 255.255.255.0
network 192.168.56.0
broadcast 192.168.56.255
```

Once I reboot, this interface will be raised automatically. I can check that it's present using the ifconfig command in a terminal on the guest.

## Naming of parts

Since I'm not so good with numbers and IP addresses, I'd rather use the guest machine name over the IP address. To do this I can add an alias to the /etc/hosts file on my Host machine, so that I can browse by name to my running guest.

I edit the file (as root) adding

the line:

```
192.168.56.101 ocelotsvr
```

If I add more guests to the host-only network, I can do the same so that I can address any of them, from any of them, by name.

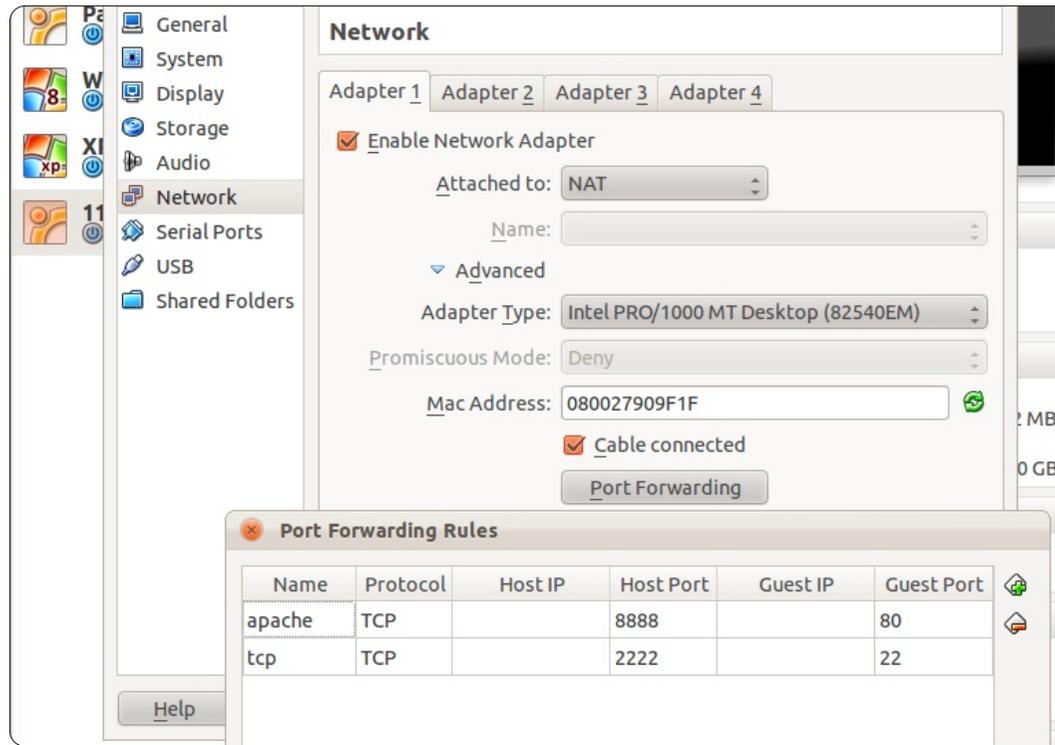
## Any Port in a Storm?

The networking experts out there will have spotted an alternate route to accessing my virtual server.

Using my default NAT adapter, I could run some extra configuration there to let me access my virtual server from the host - without using Host-only networking.

I can set up port forwarding in VirtualBox, using the guest's Network Adapter settings panel. Going to the NAT adapter, the big button at the bottom is for Port forwarding. This opens a further configuration dialog box so I can set the port forwarding rules for this network adapter for this guest only.

I need two rules; one to access the Apache web-server on my



guest, the other to handle all the other TCP traffic (http requests, mostly).

Keeping it simple, I'll name them Apache and TCP; and both use the TCP protocol. If I stick to the common port numbers, 8888 on the host forwarding to port 80 on the guest for Apache; then 2222 on the host forwarding to port 22 on the guest; this gives me specific access on the guest for those service requests. Anything else will be rejected by VirtualBox as there's no viable route.

This also means any other guests I wish to run can't access this virtual server, as no network route exists under NAT.



The Ubuntu Podcast covers all the latest news and issues facing Ubuntu Linux users and Free Software fans in general. The show appeals to the newest user and the oldest coder. Our discussions cover the development of Ubuntu but aren't overly technical. We are lucky enough to have some great guests on the show, telling us first hand about the latest exciting developments they are working on, in a way that we can all understand! We also talk about the Ubuntu community and what it gets up to.

The show is presented by members of the UK's Ubuntu Linux community. Because it is covered by the Ubuntu Code of Conduct it is suitable for all.

The show is broadcast live every fortnight on a Tuesday evening (British time) and is available for download the following day.

[podcast.ubuntu-uk.org](http://podcast.ubuntu-uk.org)



# HOW-TO

Written by Ronnie Tucker

# GIMP - The Beanstalk Pt2

**O**K, so, we've got our sky recoloured and fading into the background; next, we want to have a little village below the sky.

## Sources:

Village:

<http://www.sxc.hu/browse.phtml?f=view&id=1215281>

Tree:

[http://alfoart.com/flash/beanstalk\\_tutorial/625100\\_53361668.jpg?http://www.2textured.com/main.php?q2\\_itemId=281](http://alfoart.com/flash/beanstalk_tutorial/625100_53361668.jpg?http://www.2textured.com/main.php?q2_itemId=281)

## YouTube:

<http://www.youtube.com/watch?v=KTmehu8x2j4>

## The Village

Bring the village photo into your scene. How? This is where I test if you've read part one or not. Same idea: open the village image and drag it into our main scene. You'll probably have to resize it to about 750 pixels wide, though.

Should you see something like this, don't panic



If you look at the list of layers, you'll see that the village layer is sandwiched between the sky and background images. We need the village to be at the top of the list. Simply drag it up above the sky layer.



Click the move icon, and drag the village layer down to the

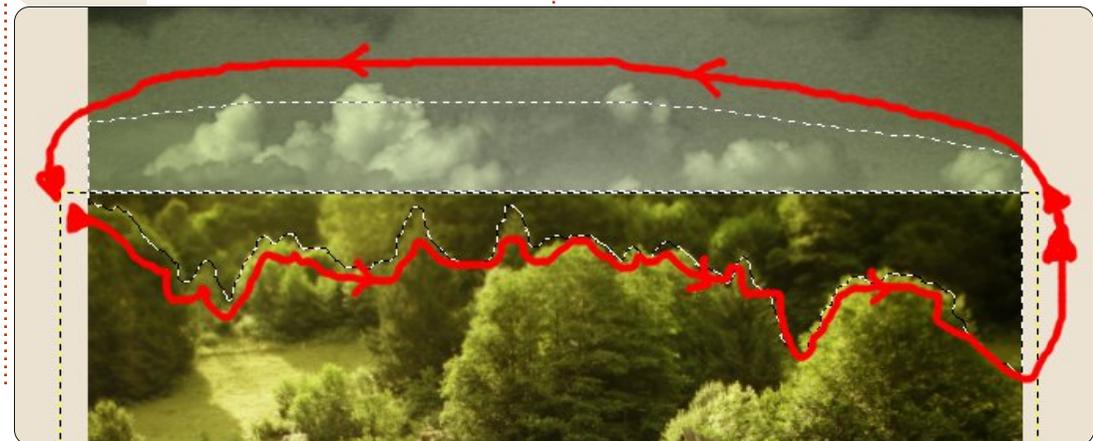
bottom of the image.

We're going to chop away some excess trees that are behind the village to give us a nice landscape.

## Selections

The selection tools are some of the most important tools in GIMP as they allow you to be as detailed as you need to be when selecting an outline. For this, our first big tutorial, we'll go with a quick and dirty selection.

Click the 'Free Select Tool' (shown left) and draw around the treeline keeping only full trees. Draw out the side of the image, and



loop back to where you started. Press the Enter key on your keyboard to complete the selection.

You can, of course, go clockwise if you prefer.

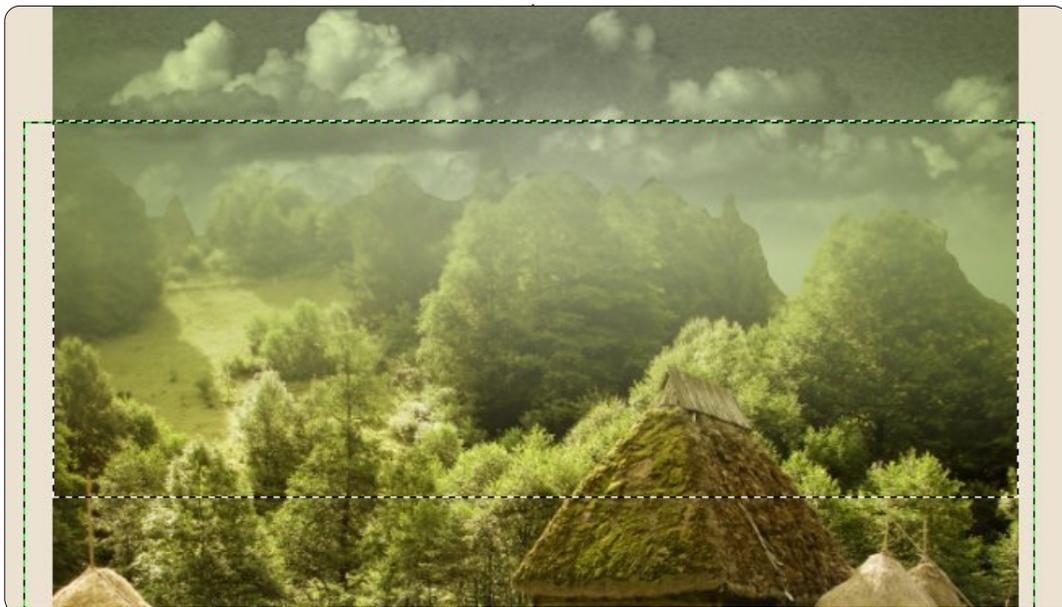
Press the Delete key on your keyboard, and anything inside that selection will be removed.

**TIP:** If you press the Delete key and the selected area shows black, then you'll have to go to the menu, click Edit > Undo. Right click on the layers thumbnail, and choose 'Add Alpha Channel'. Now you can press Delete to remove your selection. This extra 'Add Alpha Channel' step isn't always necessary.



You can, if you like, go around the treeline and tweak it with the eraser, but for now we'll cover it up with a layer mask. You do

remember how to create a layer mask, don't you? Yep, right clicking on the village layer, and choosing to add a white layer mask.



Use the rectangle select tool to select the top half of the village (below left) and, like last time, use a black and white linear gradient to fade the top part of the village.



Remember last time how I spoke about the layer mask as being non-destructive? Well, here's your chance to try it out. We created a white layer which was completely transparent, so choose a foreground colour of white and click the 'Paintbrush' icon (shown left).

**TIP:** If you need to enlarge/shrink the paintbrush you can use the square bracket keys (that's [ and ]).

The idea here (shown below right) is to (on the layer mask!) paint white over some of the foreground trees to remove them from the fog effect in the background.

OK, let's get our whopping great tree inserted and we'll finish up part two.

I'll show you another quick way of inserting an image. Click the link above for the tree source. In your browser, right click the image and copy the image to the clipboard. Go to your main image in GIMP, and, in the menu, click Edit > Paste As > New Layer. Voila!



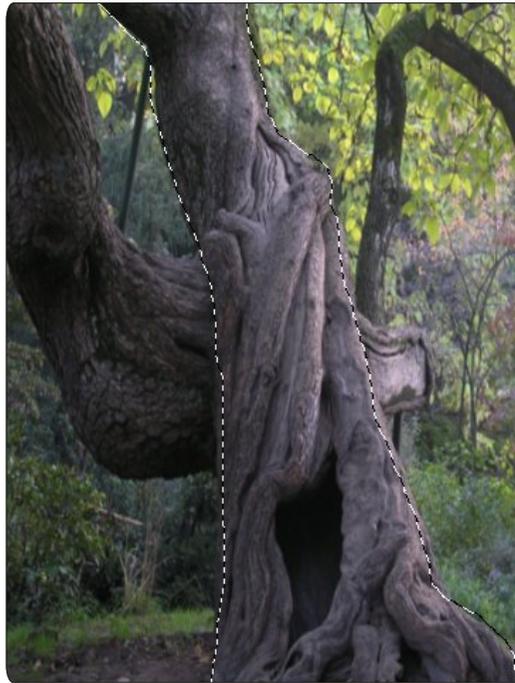
## HOWTO - GIMP: THE BEANSTALK Pt2

You'll have to resize the layer to about 600 pixels wide, and move it down to have the tree roots halfway down the grassy part of the village. Clicking resize, and then on the tree, I'm unlinking the width and height numbers as I just want to stretch the tree vertically to about 650 high.



Like we did with the village, it's time to trim out the excess background, we just want to keep the tree. Time to click the free select tool and get to work.

This time we want to keep what's inside the selection, so, in



the menu, click Select > Invert and press Delete. One tree. But I think I'd like it growing up to the right, so click Layer > Transform > Flip Horizontally. I'd also like it to have a tint of green like the rest of the image, so click Colors > Colorize, and move the 'Hue' slider until you get a greenish tint to the tree.

Finally, apply a layer mask to the tree layer, and select the top third of the tree and use the blend tool to fade the tree into the clouds.

One last thing, your homework for this lesson: use the Dodge/Burn tool on the village



layer to paint a shadow from the tree across the grass.

**In the final part of the Beanstalk image we'll add some pizzazz to the image.**



**Ronnie** is the founder, and editor, of Full Circle, an official Ubuntu member, and part-time artist who's work can be seen at:

<http://ronnietucker.co.uk>



# HOW-TO

Written by Mark Crutch

# Inkscape - Pt1

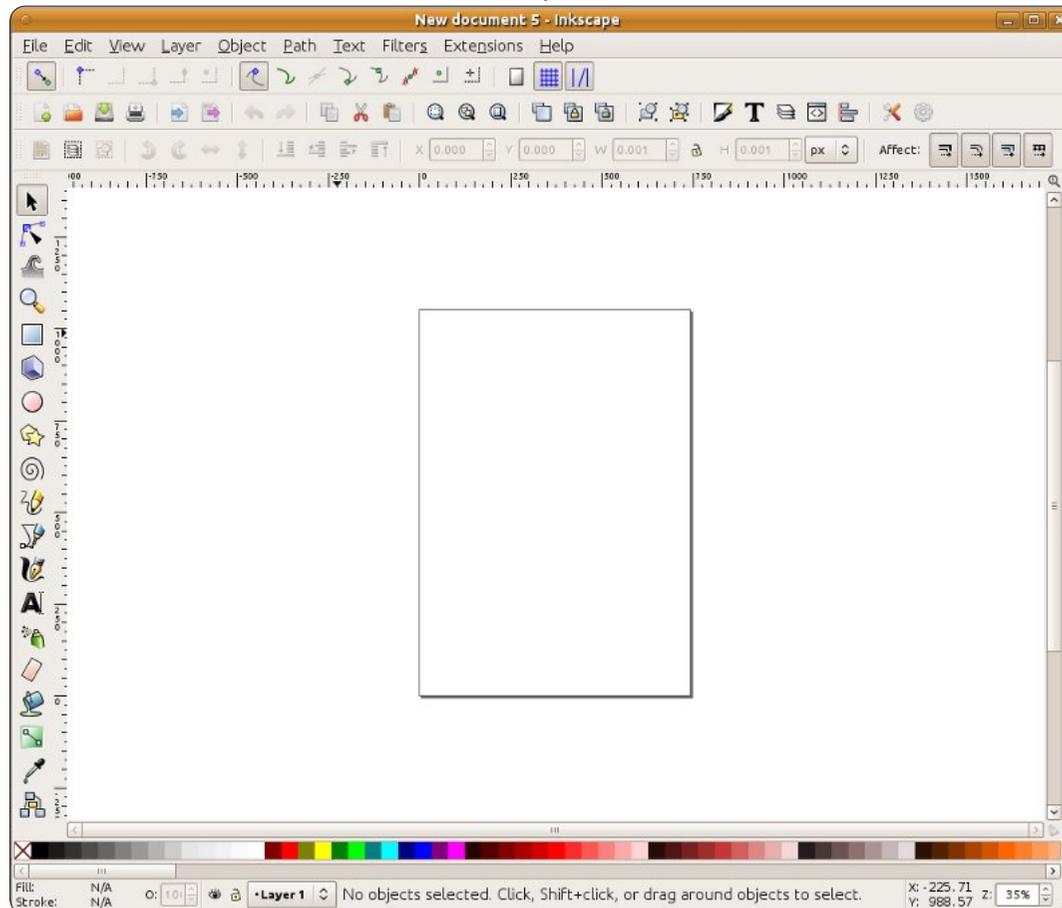
**B**efore you skip on past this tutorial thinking it's not for you, I'm going to let you in on a little secret: I can't draw. If you give me a pencil and paper, or sit me in front of The GIMP with an expensive graphics tablet, then I'll easily be bested by a four-year-old. Yet, somehow, I'm one of the artists for a regular webcomic. So, how does a fat-fingered cave painter like me manage to produce artistic works? I use Inkscape.

Inkscape is no replacement for an artistic eye, or three years at art school, but, if your problem is largely one of hand-eye coordination, then it offers you something that pencils and The GIMP don't. It gives you the opportunity to tweak and modify every line and shape that you draw - until you're happy with the result. If your hand-drawn lines are close, but not quite what you'd imagined, Inkscape gives you the time and tools to alter them, rather than having to throw them away and try again. So, before you dismiss this tutorial because you're not a

natural artist, why not have a try anyway - you might surprise yourself. If art comes naturally to you, then you've already got a head start, but, hopefully, you'll be able to pick up some Inkscape-specific tricks and tips along the way.

Let's begin by getting a copy of Inkscape. It's in the repositories of most desktop Linux distributions, so just use your normal package installer. These tutorials will be based on the 0.48 series, which has been the stable version for some time now.

On first launching Inkscape, you'll be presented with a blank working area, surrounded by a variety of toolbars, and a menu bar at the top of the window (or at the top of the screen if you're using Unity on a recent Ubuntu release). The exact arrangement of the toolbars can be modified to a limited degree: you can show and hide them using the View > Show/Hide menu, and you can switch between three predefined layouts using the Default, Custom, and Wide options at the bottom of the View menu. For these tutorials, I'll be using the Custom layout (which is badly named, as it doesn't let you customise it at all), with all the toolbars visible, and resulting in a main window that looks something like that shown left.

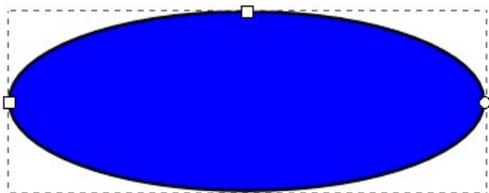


Now, let's draw something. From the Toolbox (which runs down the left regardless of your screen layout), click on the Circles and Ellipses tool (left).

Move your cursor back into the working area and you should

## HOWTO - INKSCAPE

notice that the pointer has changed to indicate that you're using the Circle tool. In the working area is an outline that represents your page – although Inkscape will happily let you draw outside its boundaries - which provides you with a large area for rough work, reference images or anything else that you don't want to appear on your printed page or exported image. Click within the page, and drag to the right and downwards to create an ellipse, releasing the mouse button when you're happy with its size and shape. Yours will probably be a different color, but we'll sort that out shortly.

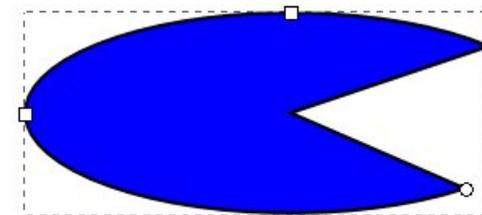
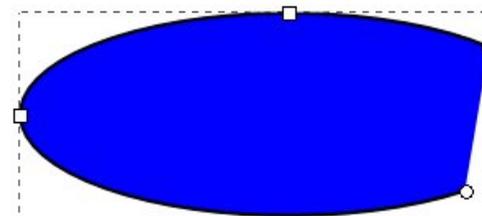


Your ellipse should have a dotted line around it, indicating that it's currently selected. Many operations in Inkscape only work on the objects that are currently selected. The quickest way to de-select everything is to click on a blank section of the working area, away from any of the objects

you've drawn. To re-select your ellipse just click on it with the Circles and Ellipsis tool still active.

As well as the dotted line, you should see two small squares and a small circle on the edge of your ellipse. These are referred to as "handles", and are a graphical way for you to modify some of the properties of your objects. Try moving one of the square handles by dragging it with the mouse: it changes the radius of the ellipse in one direction. The second square handle changes on the other radius.

The circle is a little different – it's actually a pair of circles on top of each other. They allow you to turn your circle into a segment or an arc. Try dragging one round the outline of the ellipse. Now drag the other. Inkscape tries to guess whether you want a segment or an arc, so will flip between the two modes as you move the handles. You can switch between them manually using the buttons on the Tool Control Bar, just above the drawing area. Clicking on the third button will put the handles on top of each other once more, returning you to a full ellipse.



Let's give your ellipse a little color. Select it and look down to the bottom-left of the Inkscape window. There you will find a pair of colored swatches labelled Fill and Stroke. Fill shows the color that is used to draw the inside of your ellipse, while Stroke shows the color that is used to draw its outline. You can right-click on the number next to the Stroke swatch to pick from a few standard sizes if your stroke is too thick or thin.



Above those two swatches, you'll find a whole load more swatches stretching the width of the window. This is referred to as the color palette. Click on a swatch in the palette to set the fill color of the currently selected object, and SHIFT-click to set the stroke color. At the far left of the palette is a swatch with a cross through it which can be used to set the fill or stroke to transparent. Try setting the fill to transparent and giving the stroke a color. Now play with the circular handles and the Segment and Arc buttons – it should be a bit clearer why the latter is so named.

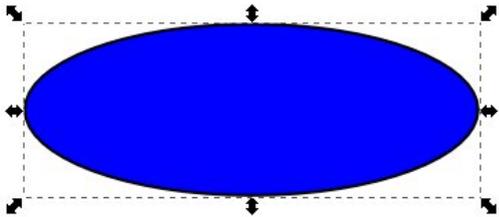


Finally let's move the ellipse around on the page. For this you'll need to click on the Select tool, which is the first button in the toolbox and looks like an arrow (left).

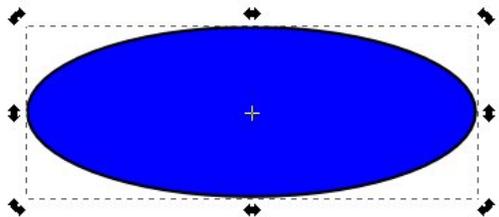
With this tool enabled, you can move an object simply by clicking and dragging it. When an object is selected using the Select tool, you'll see a dotted line around it, but rather than the small square and circular handles we saw when the Circle tool was active, you'll

## HOWTO - INKSCAPE

see some arrow-shaped handles arranged around the outside of the dotted line.



Dragging these arrows will let you change the width and height of the ellipse. Although they may seem to have the same effect as the small square handles you used earlier, they're subtly different. Try clicking on the ellipse a second time, as though you're trying to select it again. The handles change to a different collection of arrows, and a small cross appears in the center of the object.



By dragging the arrows at the corners you can rotate your object. The handles at the sides let you skew it. Try playing with them to get a feel for how they work. You

can drag the little cross around to change the center of rotation; SHIFT-click on it to return it to the middle of the object. A single click on the ellipse will switch back and forth between the resize and rotate modes.

A double-click will switch you to the Circle tool, letting you gain access to the small square and circular handles once more. Try rotating an ellipse using the Select tool, then double-click to access the square handles and change the ellipse using them. Back to the Select tool and change the width and height using the arrow handles. See, I told you they were different!

If you draw a few more circles and ellipses and move them around, you'll soon find that they can overlap or obscure each other as though they're stacked on top of one another. When you select an object with the Select tool, you'll find a group of four icons which let you move your objects up and down so that they appear in front of or behind other objects. The first and last buttons will make your object jump to the back or front of the stack, whereas the second and third let you nudge the

object down and up by one position at a time.



Sometimes you might need to get a closer view of the objects you create – or zoom out for an overview. We'll cover that in detail in a later article, but for now you can simply use the plus (+) and minus (-) keys on your keyboard to zoom in and out, and use the scrollbars to move your view around.

Now you know how to create ellipses, segments, and arcs, change their colors, and adjust the thickness of the stroke. You can move them around on the page and nudge them up and down in the stack of objects you've created. You can zoom in and out and move around the canvas. Next time we'll introduce some other shapes, but even with just ellipses, segments, and arcs, you can start to make some simple images: although it's not really the right season for it, how about drawing a snowman.

Finally, you should save your

work using the File > Save (or Save As...) menu. There's a pop-up menu in the save dialogue that lets you choose between various formats. For now just choose Inkscape SVG, the first item on the list – I'll cover some of the other options in a future instalment.



**Mark** has been using Linux since 1994, and uses Inkscape to create two webcomics, 'The Greys' and 'Monsters, Inked' which can both be found at:

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



## Guidelines

**T**he 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).

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

## Writing

There is no word limit for articles, but be advised that long articles may be split across several issues. In your article, please place where you would like a particular image to be. Please do not use any formatting in your document.

## Images

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

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

For a more detailed list of the style rules and common pitfalls please refer to: <https://wiki.ubuntu.com/UbuntuMagazine/Style> - in short: US spelling, no l33t speak and no smilies.

When you are ready to submit your article please email it to: [articles@fullcirclemagazine.org](mailto:articles@fullcirclemagazine.org)

If you can't write articles, but hang out in Ubuntu Forums, send us interesting forum threads that we could print.

## Non-English Writers

If your native language isn't English, don't worry. Write your article, and the proof-readers will read it for you and correct any grammatical or spelling errors. Not only are you helping the magazine and the community, but we'll help you with your English!

## 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.**





- ➔ 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>

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.

Download mobile clients  
for **iOS & Android**

**JOIN SPIDEROAK NOW**  
Get 2 Free GBs



**Warning:** Ripping a Commercial DVD, even for backup purposes, is illegal in many countries. However, ripping your own home videos is not illegal.



Whether you're filling out your home-built media centre, or transferring video from DVD to another format, at some point you're likely to need to rip a DVD and encode it to a different format. Depending on your media and your system, the process can be really fast or really slow.

So, what do you need to rip and encode video, and just how quick is the process? I set out to find the answers to these questions using 4 test systems. The systems ranged from a single core Athlon XP 2800+ with 512MB of RAM to an Intel Core 2 Quad CPU Q8300 with 4GB of RAM. For the experiment, I used 2 AMD-based systems and 2 Intel-based systems. Some of the results were quite surprising. The exact

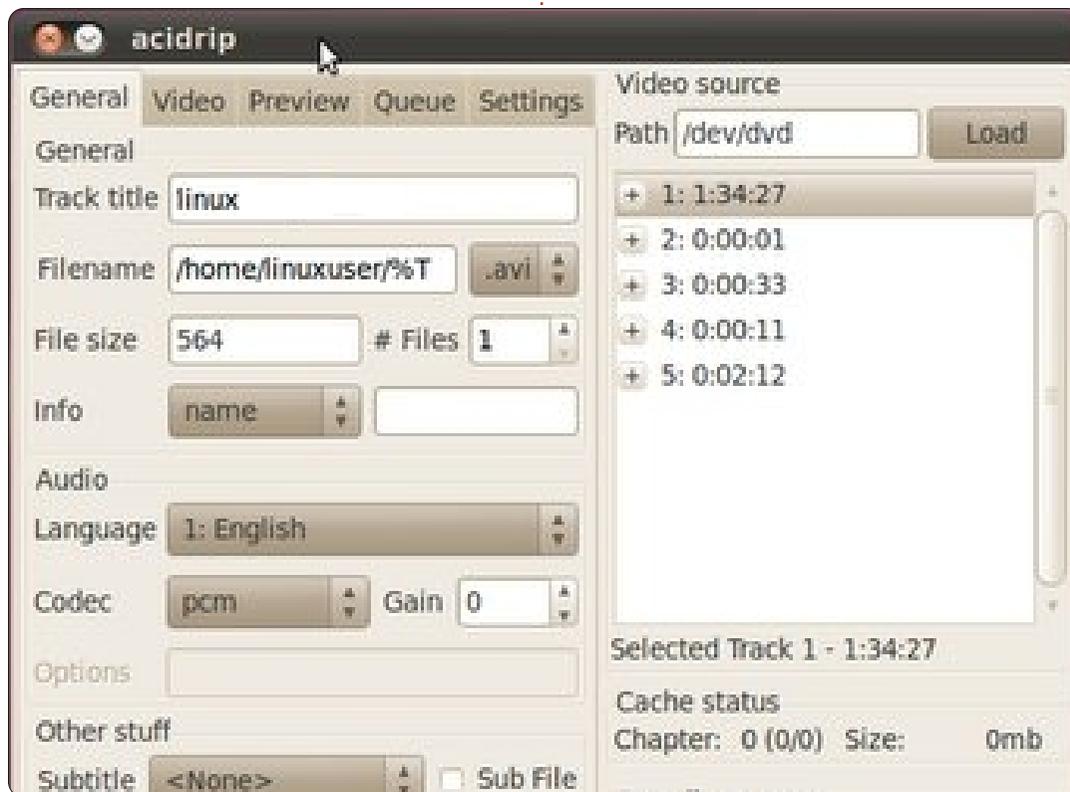
specifications for each system used are as follows:

• **pluto** - AMD Athlon XP 2800+ (2.1GHz) with 512MB RAM (2 x 256MB), onboard SiS video (128MB), LG GSA-H55N DVD burner, and a Maxtor 6L040J2 hard drive.

• **venus** - Intel Core 2 Duo (3.2GHz) with 1GB RAM (2 x 512MB), onboard ATI Radeon X1300 video (128MB), a Benq DW1650 DVD burner, and a 250GB Western Digital WD2500JS-60N hard drive.

• **saturn** - AMD Phenom II X4 905e (2.5GHz) with 4GB RAM (4 x 1GB), an ATI Radeon HD3450 video card (256MB), a TSSTCorp SH-S183L DVD burner, and a 2TB Seagate ST2000DL003-9VT1 hard drive.

• **jupiter** - a stock Dell Inspiron 545, Intel Core 2 Quad Q8300 (2.5GHz), 4GB RAM (4 x 1GB), ATI Radeon X800 onboard video (128MB), a TSSTCorp TS-H653G DVD burner, and a 500GB Western Digital WD5000AAKS-7 hard drive.



At the end of the experiment I also changed up the first system, pluto, doubling the RAM, changing out the DVD burner, and adding a 256MB video card.

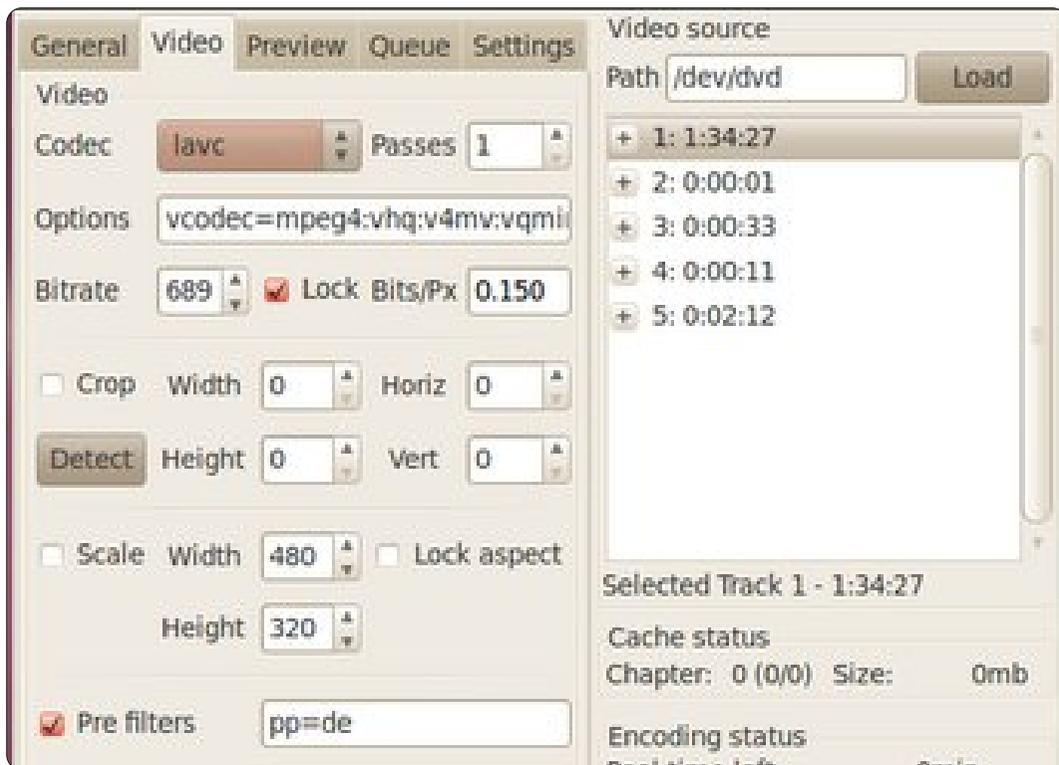
For software I used Ubuntu 10.04-LTS on pluto, venus, and jupiter, and 12.04-LTS on saturn. For ripping and encoding I'm fond of Acidrip, I find it synchronizes audio and video well. To keep the

time for the total process down I made some adjustments to Acidrip's default settings.

On the general screen I used the settings shown above.

The settings here are not much different from the defaults - other than the fact that I chose to encode with the pcm Codec, and selected English. The file size





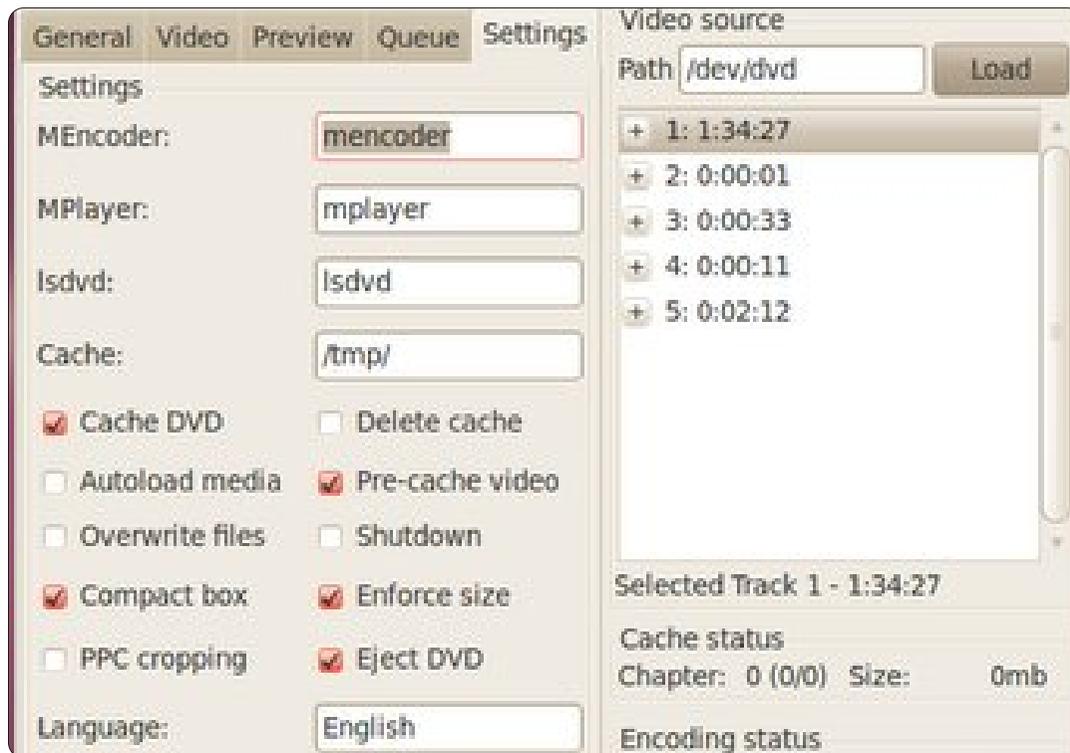
adjusts when I adjust the settings on the Video settings tab. The Video settings tab of Acidrip is shown above.

Most of the changes I made to the default settings of Acidrip are on the Video settings tab. It's important to note here that I did not set the width and height to 480x320, more specifically note that I unchecked the scale and crop check boxes. What I did check was the Lock Bits/Px box. Once the box is adjusted, I adjust the bitrate up or down so that the Lock Bits/Px

reads between 0.150 and 0.200. The higher the number, the larger the file size on the general tab. I also set the Video codec to lavc. The last changes I made were on the Settings tab which is shown above right.

On the settings screen, I set Acidrip to Cache DVD and to Eject the DVD when it's done caching.

The whole process is really 2 steps: ripping the DVD, which consists of copying the DVD contents to the hard drive, and



encoding the DVD contents to a compressed format. The DVD ripping process can be very quick depending on the hardware used. The encoding process is normally much longer. Interestingly, wrong or bad hardware can actually make the ripping process take almost as long as the encoding process (more on this later).

For the experiment I chose a DVD that was 91 minutes long, the length of a typical feature movie.

As expected, pluto, the single

core, slowest system, performed the slowest - it took 12 minutes and 9 seconds to rip the DVD, and a whopping 2 hours, 39 minutes and 27 seconds to complete both the ripping and encoding process.

Venus, the dual core system, produced some very surprising results - clocking in hands down with the fastest DVD rip at 4 minutes and 52 seconds. The entire process, ripping and encoding, took 51 minutes and 33 seconds.

Another surprise came from saturn which beat all the systems by completing the whole process in 37 minutes and 14 seconds, but took 11 minutes and 37 seconds to rip the DVD.

Jupiter, the stock quad core Dell, ripped the DVD slightly slower than saturn, with 3 seconds difference between the two, but took over 3 minutes longer for the whole process at 40 minutes and 25 seconds.

From these results we can draw some interesting conclusions: a good DVD burner does matter in the large scheme of things. Despite having a slower CPU and less memory, venus, the dual core system out-ripped both quad core systems with less than half the time. The Benq DW1650 literally destroyed the other DVD burners when it came to ripping

performance. If you're only ripping one DVD, this might not be much of a concern, but multiply it by 10, 20 or 50 DVDs, and you're talking about a fair amount of time.

I was also surprised that the AMD quad core outperformed the Intel quad core. I chocked up the win to the fact that I had a 256MB video card in the AMD system - until I ran one more experiment.

For the last test I added a 256MB AGP video card to pluto, the single core AMD Athlon XP 2800+ system. I also added 1GB of RAM (2GB actually, but the system only recognized 1GB despite the motherboard manual indicating the system would recognize 2GB), and a Memorex 3202-3269 DVD burner. I used the same DVD and settings.

Unfortunately, the time I had

access to this system ran out, but I was able to get some results. The DVD ripping process took a whopping 15 minutes and 7 seconds. Whether the DVD burner was just slow or bad I'm not sure, but it lagged far behind all the others. After 1 hour, this revised pluto had encoded just 28 minutes and 7 seconds of video. The slow DVD burner no doubt slowed down the whole process, but what's really telling here is that the better video card doesn't seem to help as much as having more CPU cores. Even if we took off the 15 minutes for ripping, it's still only 28 minutes of video (of 91) ripped in 45 minutes.

Given the choice between a better video card and more CPU cores, I'd buy more CPU cores. Whether the AMD quad core won because of the video card or not I'll leave up to you to judge, but I'd

like to think the AMD system was simply better than the stock Dell Intel-based quad core. The one sure thing is that if you're planning on ripping a large DVD collection, you can save a significant amount of time by picking up a DVD burner like the one mentioned in this article.

System Name	Cores / CPU	Rip Time	Encoding Time
pluto	1 / AMD Athlon XP 2800+	00:12.09:475	02:39.27:443
venus	2 / Core Duo 3.2GHz	00:4.52:400	00:51.33:852
saturn	4 / Phenom X4 905e 2.5GHz	00:11.37:524	00:37.14:647
jupiter	4 / Intel Core 2 Quad 2.5GHz	00:11.44:477	00:40.25:240



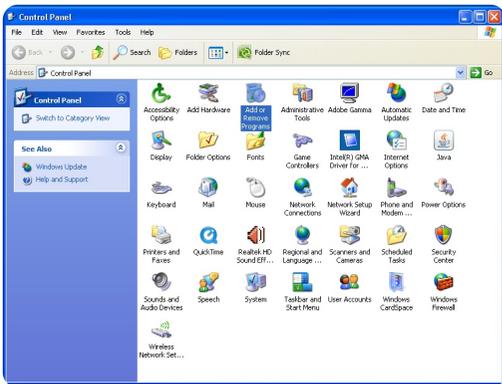
**Charles** is a step-father, husband, and Linux fan who runs a not-for-profit computer refurbishing project. When not breaking hardware/servers he maintains a blog at <http://www.charlesmccolm.com/>.



# CLOSING WINDOWS

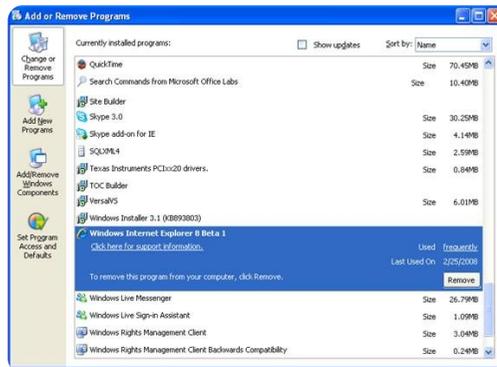
Written by:  
**Ronnie Tucker (KDE)**  
**Jan Mussche (Gnome)**  
**Elizabeth Krumbach (XFCE)**  
**Mark Boyajian (LXDE)**  
**David Tigue (Unity)**

**M**ost Windows applications will have an Uninstall option that can be seen either in the appropriate menu on the Start button, or it'll be listed in the Add/Remove Programs function. The latter is started from the Control Panel:



You are then presented with a list of all applications installed in

## Windows.



As you can see from the screenshot, you can uninstall and add new programs, and even Windows components. Obviously care should be taken when uninstalling Windows components.

## Adding/Removing Software

### Kubuntu

Adding/removing software in Kubuntu is done using an application called Muon Software Centre. This is located in K > Applications > System.

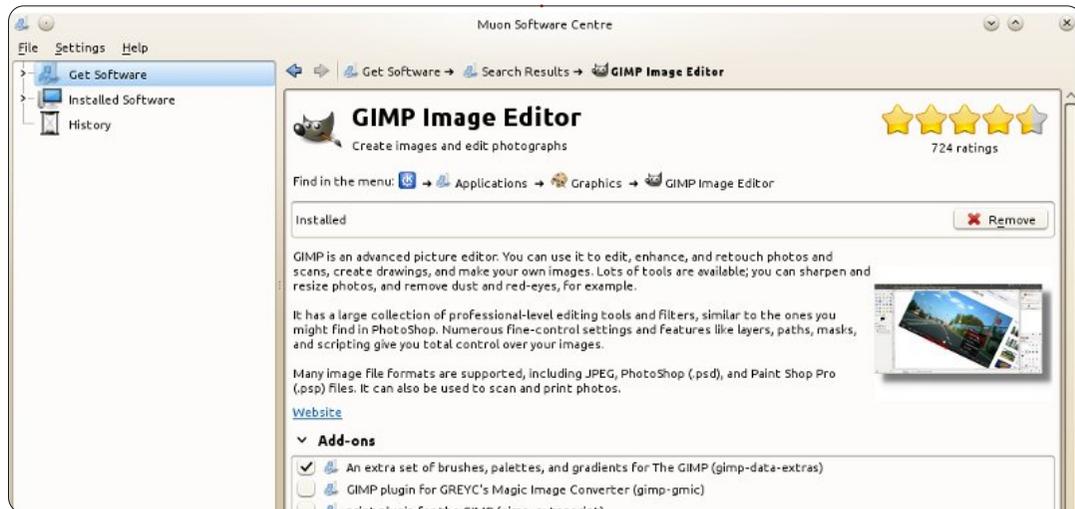
It's far from elegant, but it does the job of listing applications you

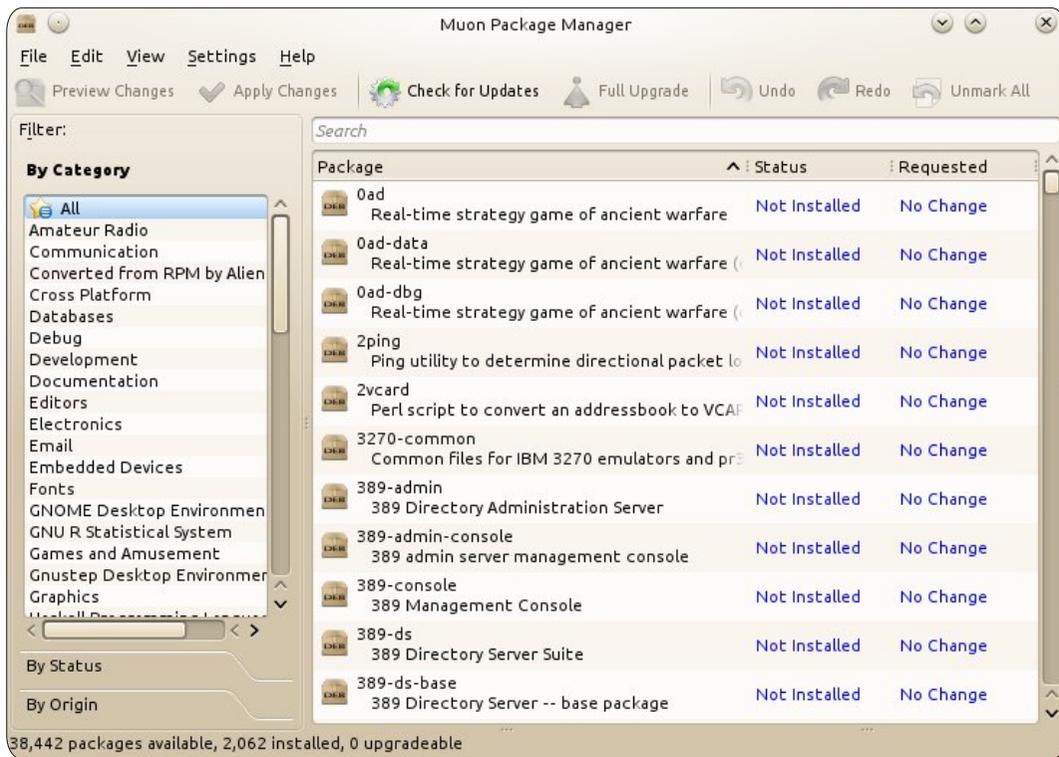
# Add/Remove Software



can get, or have (left panel), and categorizing software you can get (right panel). You can, if you wish, search for applications by entering some text in the search box at the top right of the window.

Clicking the application name, then the 'info' button will show you more detailed information about the application such as a description, add-ons, and even user reviews. Clicking the install/remove button at the top of the info screen





will do the relevant action.

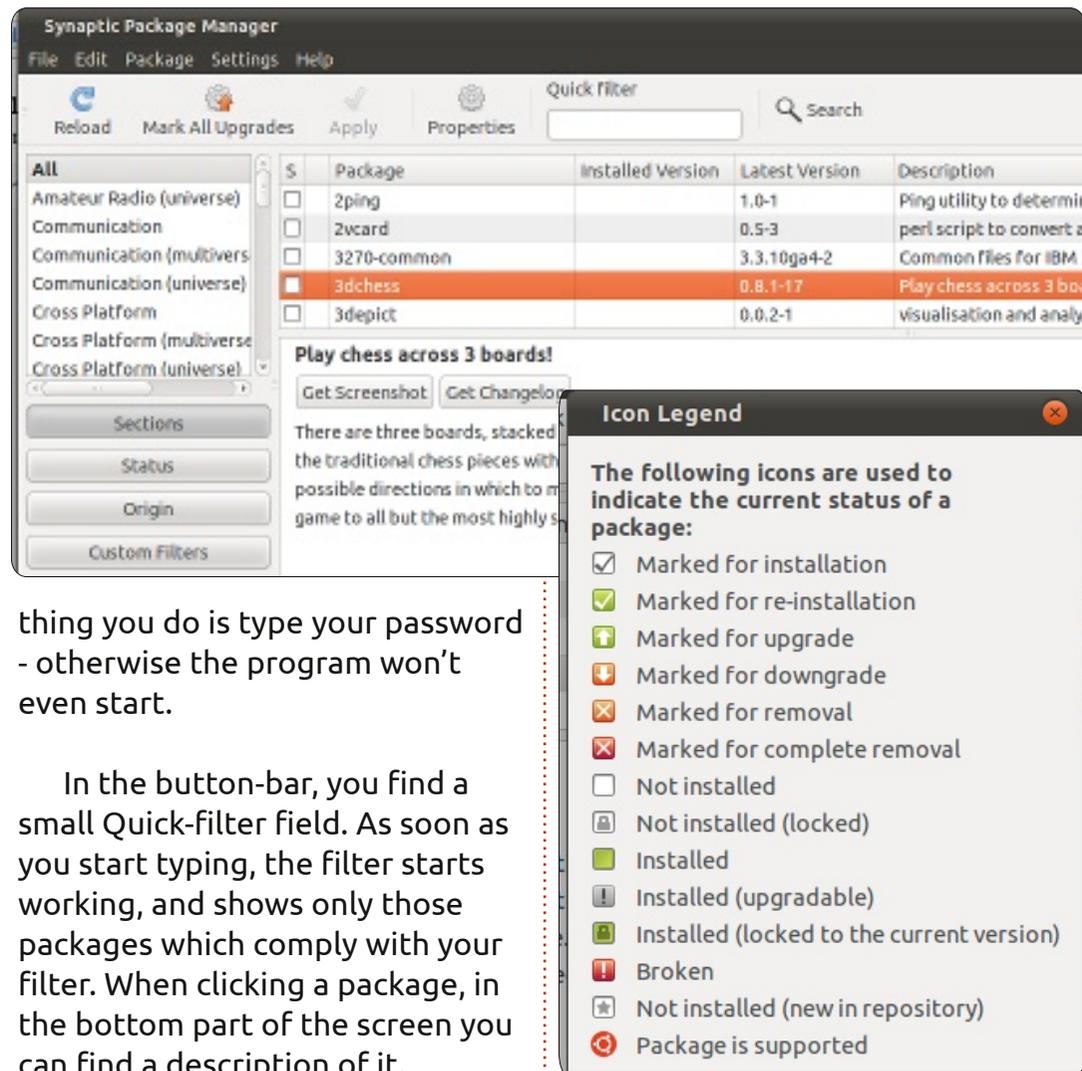
Another way of adding/removing software is via the Muon Package Manager. It's also in K > Applications > System, and is more like the old Ubuntu Package Manager. This Muon Package Manager is for fine tuning of files such as libraries, broken files, and such like.

## Gnome-Shell

As was so nicely put in the previous section, the Gnome

version of Ubuntu has the old Ubuntu Package manager called Synaptic. Synaptic is a one-stop place for installing and uninstalling software, adding/removing repositories (handled in the next section), and fixing broken packages.

It works pretty much the same as Muon, although I always find one difference very striking: in Muon you can choose your desired software, and prepare it for installing, and then type in your password; in Synaptic, the first



thing you do is type your password - otherwise the program won't even start.

In the button-bar, you find a small Quick-filter field. As soon as you start typing, the filter starts working, and shows only those packages which comply with your filter. When clicking a package, in the bottom part of the screen you can find a description of it.

When you want to install a package you simply click the small box in front of the package name, and select Mark for installation. In the menu bar you find the Apply option - which will start the installation process. Should any dependency have to be installed as

well, the program will tell you now.

After installation, the little box in front of the package name is colored green to indicate this package is installed. To find the meaning of the different colors, choose Help menu > Icon legend to

# CLOSING WINDOWS

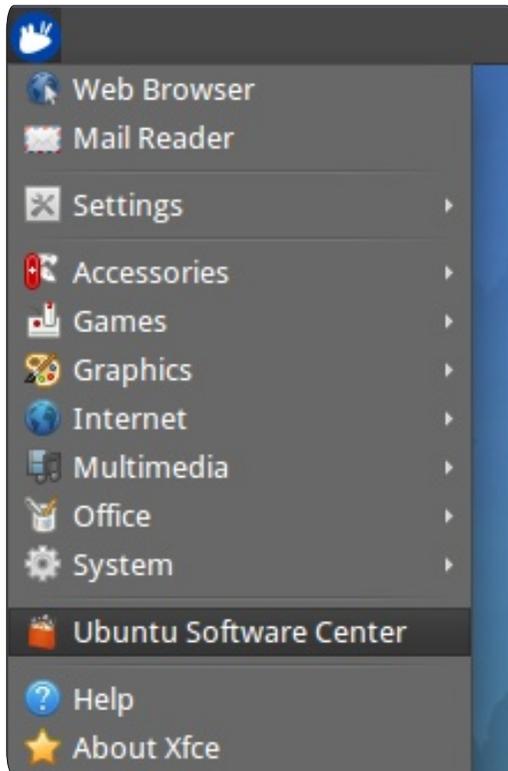
get the full list:

The menu Edit holds the Fix broken packages option. Broken packages make it impossible to install anything, so you need to fix them first.

Another option is Repository, but this will be described in the next section of this issue.

## XFCE

In Xubuntu, just like Ubuntu, you will use the Ubuntu Software



Center to add and remove software.

The Ubuntu Software Center is launched simply via Mouse menu > Ubuntu Software Center

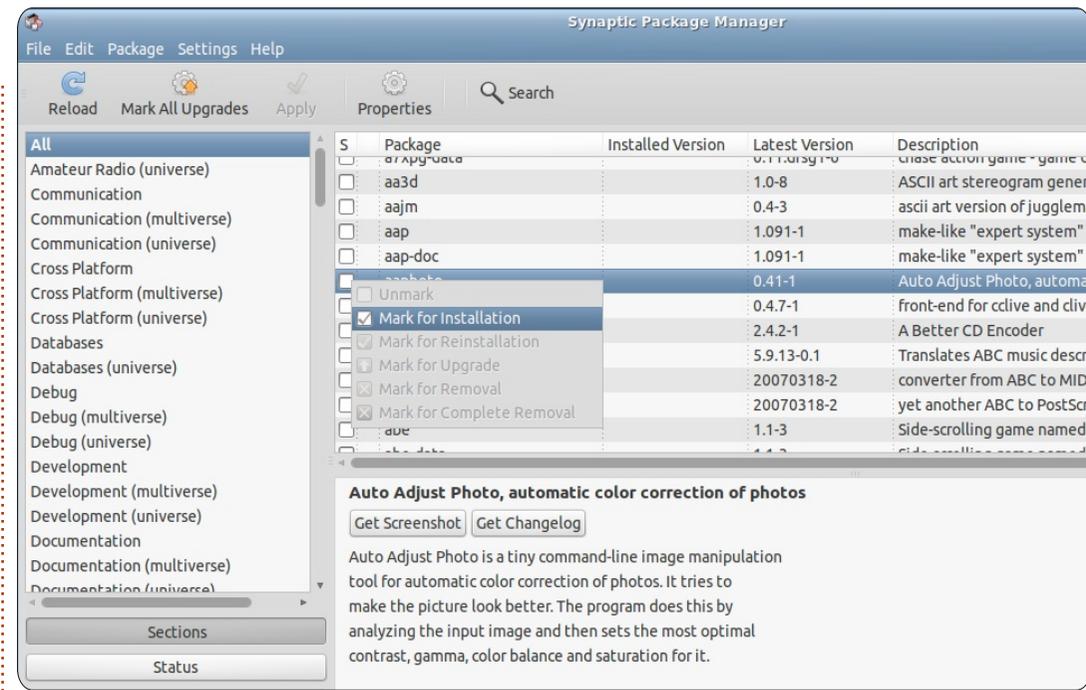
The Ubuntu Software Center features thousands of applications from the Ubuntu repositories and beyond, including featured free and paid applications, and content including books and magazines.

When you click on a piece of software or content in the Ubuntu Software Center, you will be presented with a page of information with a description, and possibly add-ons, a screenshot, what people have “also installed,” and user reviews. There will also be an Install button on the right which will allow you to easily install the software.

See the Unity section for more about the Ubuntu Software Center.

## LXDE

As with the other “\*buntu’s” described above, LXDE also uses the Synaptic Package Manager as the default software management



application. You open it by clicking on the Main Menu icon and selecting System > Synaptic Package Manager.

You will immediately be prompted to enter your “root” (administrator) password; Synaptic won’t open without it. Once you’ve entered your “root” password, Synaptic will open.

The previous screenshot shows an application selected; the pop-up menu displays the options which, in this case, is only to “Mark for Installation”. If this application was already installed, then the other

(greyed out) options would be available for selection.

Also, once you have selected an application (as indicated by the “blue highlight” - part of the default Theme in LXDE), you not only get a description of the program in the lower pane (the pane below the list of available applications), but the “Properties” icon in the toolbar activates. Clicking the “Properties” icon opens a window with more information about the selected program.

As you can see, each of the five tabs provides more detailed



information than what is available from the main Synaptic window, including all program dependencies. It is very handy to know, prior to beginning the install process, what other software dependencies must be met in order for the selected program to work on your system. You would, in fact, be able to review this list prior to committing the install, as Synaptic identifies and displays and marks all required dependencies prior to starting the actual software installation; however, it is sometimes handy to review these requirements without actually initiating the install process.

Again, I must emphasize that the primary goal of LXDE is to provide a “lean” distribution with minimal resource requirements, and this drives the development of everything in this “distro” right down to the default installed software applications. That said, and as mentioned in previous installments of this series, while you’re in Synaptic, you can search for and install other programs to your liking. If your computer resources will accommodate more resource-hungry programs, then you can certainly choose to install

them. For example, if you would like to use the “Ubuntu Software Center” application that is discussed above for XFCE, simply search for “Software Center” in Synaptic, and it will find the “Ubuntu Software Center. As described above, mark it for installation and use it alongside or instead of Synaptic.

This is one of the great strengths and pleasures of open source software: having the freedom to explore and use software not only for free (no cost), but freely (no compromising licenses). And applications like Synaptic, the Ubuntu Software Center, and others, make exploring new applications fun and easy.

## Unity

In Unity, just like Xfce/Xubuntu, you will use the Ubuntu Software Center application to add/remove software. The Ubuntu Software Center is an excellent application that has been in development for a long time now, and the developers have continuously made it better with each release. Ok, you want to install some software? Let’s do it. First open the Ubuntu Software Center by clicking on the the

button in your launcher on the right-hand-side of the screen that looks like an orange shopping bag with bubbles coming out of it. This will open a window that will show you a list of some of the more popular titles, and a picture of them at the top. Below that you will see a list of some of the newer software added to the Software Center. Let’s say you want to install “Lord of Ultima”. You will need to search and find the application. You have several options to do this. First you can type in ‘Lords of Ultima’ in the search bar in the upper right-hand corner of the window. Another way to do this is to click on the ‘Games’ label on the left-hand side. Once you find the application you want to install, select it and click on the ‘More Info’ button. Doing this will show you the overview page of the application. You will see a short description of the application, how much the application costs, and a button on the right-hand side that will either say ‘Install’ or ‘Buy’. Don’t get frightened, most applications are free, but some do cost a little bit of cash. The game we are installing in this example, “Lords of Ultima”, has a ‘Buy’ button, but it is still free. If you look to the left of the ‘Buy’ button,

you will notice that the cost of this application is ‘US \$0.00’. So click ‘Buy/Install’, type in your password when prompted, and watch your application being installed. Once the Ubuntu Software Center shows the application as being installed you will be able to find it in the Dash.

Now, you may be saying, “How do I remove an application?”. Well if you are asking yourself that, no worries, I’ll explain. The Ubuntu Software Center has made things pretty easy for you as far as uninstalling applications go. To get a list of all the applications installed on your system click the button near the top center of the Ubuntu Software Center labelled ‘Installed’. Doing this will show a list of all the software you have installed. It gives you this list broken down into sections. (i.e., Accessories, Games, Internet, etc.) So, let’s uninstall ‘Lords of Ultima’. Since you should have the list in front of you now, click on the arrow next to ‘Games’ to expand it. Scroll down until you find ‘Lords of Ultima,’ and select it. You will see two buttons when you highlight the installed application, ‘More Info’ and ‘Remove’. If all you want to do is uninstall the application,

# CLOSING WINDOWS

then click on the 'Remove' button. You will be prompted to type in your password; do that. Once you have typed in your password, and hit the 'Enter' button, the Ubuntu Software Center does the rest. Before you know it, the application has been removed, and you will be able to go on to the next application you want to remove.

The thing I have noticed with the Ubuntu Software Center is that it is set up to be pretty much 'fool proof'. So don't be afraid to mess around with it. If you try to uninstall an important application, then the Ubuntu Software Center will let you know about it. If you try to install an application that will cause a conflict with other applications, again you will be told about it and given the option to cancel the install. My best advice is to just 'play' with it. Get to know the application, and don't be afraid to try different things with it. You will notice there is a 'History' button; it will show you all applications that have been installed or removed. Also there is an 'All Software' button that will show you all the software available. Beside those buttons are small triangles, if clicked you will see a small drop down menu

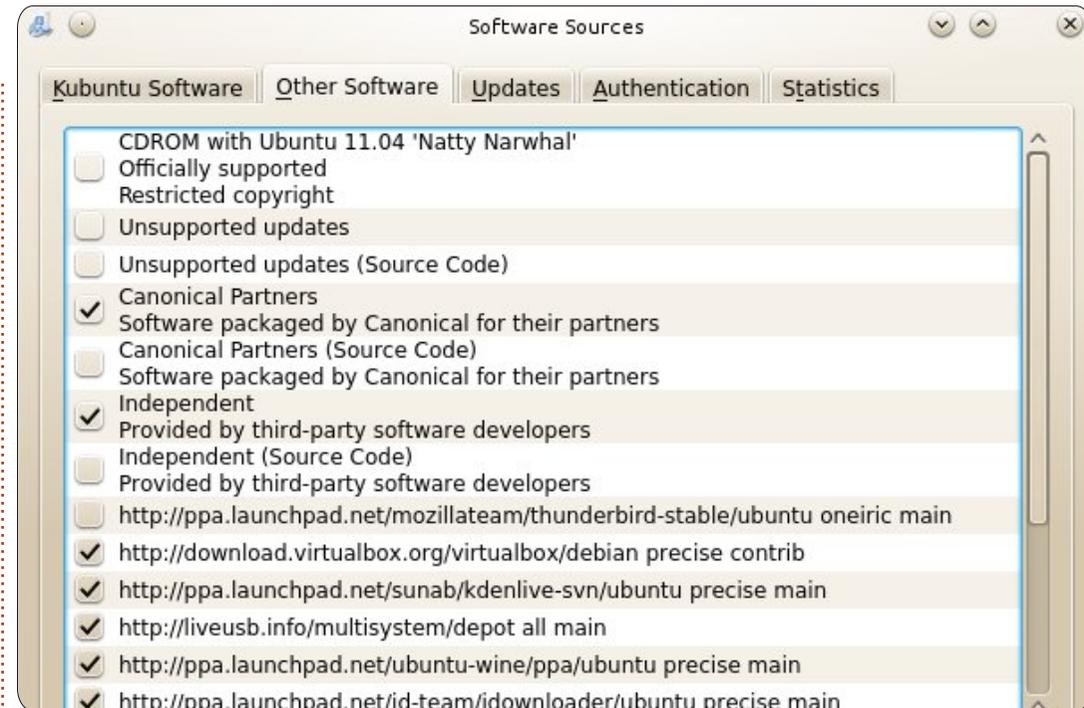
allowing you to better refine what you are searching for. Just remember, you can always remove an application if you decide that you don't like it, so why not try out any application that sounds like it might be interesting to you.

## Adding/Removing Repositories

The one thing that Windows lacks, which we have in Linux, is repositories. Think of repositories as being libraries that you can borrow books from. If a new library opens up in your neighborhood you can go and browse the books they have, and they might have something the other library didn't have. It's the same with repositories. On installing your chosen \*buntu, you'll get access to the Ubuntu repositories, but you can (if you wish) add repositories from other folks which will grant you access to newer versions of software, or software that isn't in the Ubuntu repos.

## Kubuntu

Both the Muon Software Centre and Muon Package Manager let you view, add, and remove



repositories - by going to the menu and clicking Settings > Configure Software Sources.

Before you are even allowed to see the software sources, you'll need to enter your admin password.

From here, you can add a repo, edit a current repo, or remove a current repo. Clicking 'Add' will present you with a window where you paste in the repo as given by the software supplier. So, for example, if I were to add the Wine repo, as given on their site at:

<http://www.winehq.org/download/ubuntu> I'd click 'Add,' and then paste in:

**ppa:ubuntu-wine/ppa**

**NOTE:** the old format for repos (for example: *deb http://ppa.launchpad.net/ubuntu-wine/ppa/ubuntu precise main*) is still valid.

Then follow the prompts to allow the Software Centre (or Package Manager) to grab the contents of the new repo, and reload its list of available

applications.

## Gnome-Shell

Synaptic also makes use of repositories as described above. Click Settings > Repositories to see which ones are installed already and what can be done with them.

The first tab shows you which main parts are enabled and which server you normally use when using Synaptic.

The second Tab, Other Software, shows you which repositories are enabled to install from. Here you can add new repos manually. Look at the example in the Kubuntu section; in Synaptic it works just the same.

The Updates tab shows you how the program will handle updates of installed packages. When you want to keep your OS in shape then be careful which update sections you enable. Proposed and Backports might have versions of software that are not fully tested and safe, which might endanger your installation.

To make sure you install only

from safe sources, many repo's come with an authentication key which goes in here. When adding a new repo using the ppa:ubuntu-wine/ppa line, the key will automatically be added.

In Statistics, you can enable a tickbox to upload info about your system.

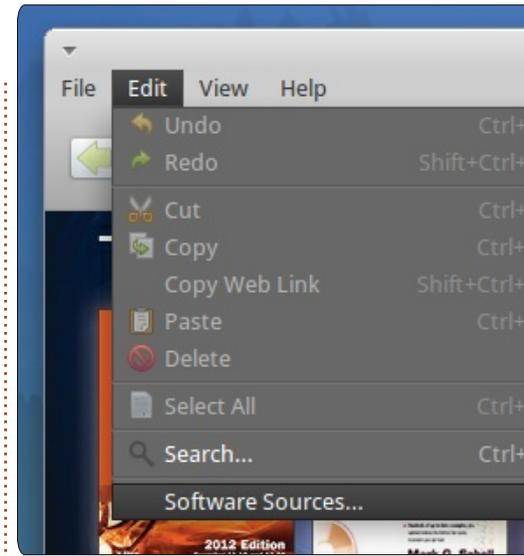
After adding a new repo, or removing one, make sure to use the Reload button in the main screen to synchronize the list of packages with the latest list of repos.

## XFCE

Since Xubuntu uses the Ubuntu Software Center, you add additional repositories by going to Edit > Software Sources... in the Ubuntu Software Center.

This will open up a window with a series of tabs for different software repository types that are available. There are many repositories already available to select and deselect, and you can also add your own repositories.

See the Unity section for more



about handling software repositories.

## LXDE

Thanks to Jan Mussche for the excellent description of managing repositories in the section above on Gnome-Shell; the same procedures apply to LXDE, so they do not bear repeating. This is a great opportunity to emphasize something you have already learned from this series (Closing Windows): knowledge and experience gained in one Linux distribution (distro) can go a long way to familiarizing you with other distros. We've all read from detractors [of Linux and open source] that "there are too many

Linuxes; all this choice is too confusing - users don't know which one to use." That's like saying you have too much freedom or too many flavors of ice cream! But, to concretely address this accusation, I must point out that if you are inclined to explore/test different distros, much with what you've become familiar in one distro will translate robustly to others. As a contributing author to this series, I am struck by the richness this kind of knowledge and experience brings to the user, and the world of options and opportunities it yields compared with other technology "monocultures." So, do be encouraged to explore and learn, and use this series as a great starting point for your journey!

## Unity

With Unity using the Ubuntu Software Center just like Xubuntu/Xfce, you can refer to the Xfce section to see how easy it is to add extra repositories. The instructions in that section are exactly the same in Unity.



**A**s Ronnie mentioned already in the last issue, I'm responsible for ensuring that you, dear readers, can read our Full Circle Magazine on your tablet and Android devices

But, what is needed to get you the finished product on your devices? Well, in short we need four different things:

- IrfanView
- a text-based Web editor
- Sigil
- Calibre.

You will forgive me if I do not include all the details, for that I would need a completely separate issue. However, I would like to offer you a little insight into my work. Therefore, I will describe an example in this article, the emergence of an article within the EPUB file. And which article is to be better than Ronnie's Birthday article?

OK! Let's get it on! Once I got the invitation from Ronnie that I may download the articles from

<input type="checkbox"/>	TITEL	EIGENTÜMER	ZULETZT GEÄNDERT
<input type="checkbox"/>	☆ <a href="#">FCM60 - How To - Beginning Python 32</a> Freigegeben	Greg Walters	11. Apr. Gord Campbell
<input type="checkbox"/>	☆ <a href="#">FCM60 - HowTo - Prey.odt</a> Freigegeben	Ronnie Tucker	11. Apr. Michael Kennedy
<input type="checkbox"/>	☆ <a href="#">FCM60 - Q&amp;A</a> Freigegeben	Gord Campbell	11. Apr. Michael Kennedy
<input type="checkbox"/>	☆ <a href="#">FCM60 - Linux Labs</a> Freigegeben	Ronnie Tucker	11. Apr. Michael Kennedy
<input type="checkbox"/>	☆ <a href="#">FCM60 - My Story - FCM</a> Freigegeben	Ronnie Tucker	10. Apr. Gord Campbell
<input type="checkbox"/>	☆ <a href="#">FCM60 - Letters</a> Freigegeben	Ronnie Tucker	10. Apr. Gord Campbell
<input type="checkbox"/>	☆ <a href="#">FCM 60 - Ubuntu Games - SNES Emulation Pt2</a> Freigegeben	Riku Järvinen	10. Apr. Gord Campbell
<input checked="" type="checkbox"/>	☆ <a href="#">FCM60 - My Story - Ubuntu</a> Freigegeben	Ronnie Tucker	10. Apr. Michael Kennedy
<input type="checkbox"/>	☆ <a href="#">FCM60 - Command &amp; Conquer - tricks and tips II</a> Freigegeben	Lucas Westermann	9. Apr. Michael Kennedy
<input type="checkbox"/>	☆ <a href="#">FCM60 - Ubuntu Women</a> Freigegeben	Lyz	9. Apr. Michael Kennedy

the G-Docs. I do so normally in two versions: ODT, and HTML.

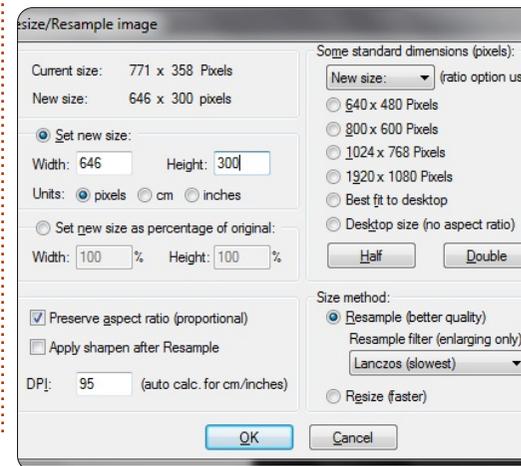
But why in two variants you ask? A legitimate question, but one that we will be answering this within the article.

Ok, the two downloads have been completed. Now we can proceed to the next step: the editing of the photos. This is necessary because the authors hand in their screenshots in JPG files as well as in PNG files. That is why I need the HTML download. Each article is located in a separate folder, so I can easily get the pictures from the "images" folder.

Because we want our Full Circle Magazine Mobile Edition to be as

small as possible, we have decided to minimize the images within the articles to a maximum height of 300 pixels.

To achieve this, I use the software "IrfanView". I open the file, use ctrl + r for resizing, and thereafter ctrl + s for saving as PNG.



When I'm done, I must make one final preparation before I can begin the "real" work, the layout of the EPUB edition.

But before I continue here, I would like to make a short detour. As mentioned earlier, we create our issues of FCM with a program called Sigil. A short visit with our friend Google will tell you that this is a so-called WYSIWYG (What you see, is what you get) editor. That means that we could – theoretically - create our issues solely with this software. But, if you look closer, you will recognize that the document markup source code is far from good.

So I've decided to create the source code in a web editor (I use the editor Scriptly in Windows, and Blue Fish in Linux). For those who are not familiar with HTML Markup Language, there is also Kompozer, available in the Ubuntu Software Center.

But enough of theory. Let's get back to the practice.

As I addressed earlier, we need one last little thing before we begin the real work. We create a small web project on our computer. What do we need? Three folders and a CSS file: the three folders bearing the following names: a) Images, b) Styles, and c) Texts. The CSS file, inside the Styles folder, we name "fcm-design.css".

Now we have it all together. Let us now have a new file in our Web Editor, which includes the following basic structure:

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN"
"http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html
xmlns="http://www.w3.org/1999/xhtml" xml:lang="de"
lang="de">
<head>
<title>Title of the FCM Article</title>
<link
href=" ../Styles/fcm-design.css" type="text/css"
rel="stylesheet" />
</head>
<body>
</body>
</html>
```

As you can see, we link the CSS file directly in HTML style. That's

```
Quelltext von: file:///C:/Users/SilverLion/Downloads/FCM60MyStoryFCM.html - SeaMonkey
Datei Bearbeiten Ansicht Extras Fenster Hilfe
<html><head><title>FCM60 - My Story - FCM</title><style type="text/css">ol{margin:0;padding:0}
72pt}.c4{margin:5px;border:1px solid black}.c5{line-height:1.0}.c3{height:11pt}.c0{font-style:
top:24pt;line-height:1.15;text-align:left;color:#000000;font-size:36pt;font-family:Arial;font-
height:1.15;text-align:left;color:#666666;font-style:italic;font-size:24pt;font-family:Georgia
family:Arial}h1{padding-top:24pt;line-height:1.15;text-align:left;color:#000000;font-size:18pt
top:18pt;line-height:1.15;text-align:left;color:#000000;font-size:14pt;font-family:Arial;font-
height:1.15;text-align:left;color:#666666;font-size:12pt;font-family:Arial;font-weight:bold;pa
align:left;color:#666666;font-style:italic;font-size:11pt;font-family:Arial;padding-bottom:2pt
align:left;color:#666666;font-size:10pt;font-family:Arial;font-weight:bold;padding-bottom:2pt}
style:italic;font-size:10pt;font-family:Arial;padding-bottom:2pt}</style></head><body class="c
</p><p class="c2"><span class="c0">by Ronnie Tucker</span></p><p class="c3 c2"><span class="c0">
party hats, or slices of cake, I thought I'd spend a page (or three) showing the process
</span></p><p class="c2"><span class="c1">Week 1 - The Deadline</span></p><p class="c2"><span>
deadline for the regular folks. But, roughly two weeks before the deadline, I send out my first
it's ignored, but I do it anyway.</span></p><p class="c3 c2"><span></span></p><p class="
```

no problem for Sigil. And it makes our job much easier.

Ok, then. That is enough theory. Let's layout an article!

## Step # 1

Open your folder with the



created "webproject", your preferred web-editor (like bluefish), and the Article you want to layout in ODT. Why ODT and not the HTML Version? Well, I'll show you.

As you can see from the image above, that is one hell of a mess to

go through and format into proper HTML source!

So Instead we take a clean ODT like the one shown below.

That will save us a lot of time. Believe you me!

So, with these three things open, we could create a new CSS-File with the Code given in the box below, and start our work.

```
h1.headline {
text-align: center;
font-size: xx-large;
font-family: monospace;
font-weight: bold;
text-decoration: underline;
}
p.text {
text-align: justify;
font-family: Tahoma;
font-size: 12pt;
}
div.photos {
text-align:center;
}
```

These are the most important things you should have in your EPUB-Layout: The Headlines (h1.headline), the Paragraphs (p.text), and a box for screenshots / pictures you would like to insert. I'll skip the details for now, but, if you are interested in learning more about CSS / HTML, Google will

show you a lot of good tutorials like SelfHTML.org or W3C.

Ok, all things are set: We have our ODT, our editor, our CSS-Design, and our pictures ready.

## Step # 2

Now the actual coding begins: for every paragraph in the ODT, I prepare a line with `<p class="text"></p>` and duplicate it several times, then I check on how many pictures there are, and insert `<div class="photos"></div>` as often as needed.

Now we can switch between the ODT and our editor to copy / paste the article into the HTML Source-Code we prepared.

And when we are done, it looks like that shown below left.

Once I am done copying and have the article ready in HTML Source Code, I start Sigil which looks very similar (right).

Let me give you a short introduction into the software: As you can see, we can find our three folders Texts (in the picture referred to as "Text" because I am writing this article from a German system), Styles ("Formatierung"), and Images ("Bilder") - on the left side. On the Upper Side, you see two buttons marked: the red one is for the source-code-view in the software. This is where you need to copy and paste the HTML source you created in your editor. But before you do this, you'll need to



create a new "Section.xhtml" by clicking the blue marked button.

## Step #3

Ok, we have inserted our HTML Code into Sigil. But we forgot one – well, actually two – things that are rather important: our pictures AND our css-file which gives the layout-design to our article. So how do we get those things into Sigil? That is easy: Sigil has a function to include already existing files into the epub-project you are currently working in. Therefore you need to click on the folder "Styles" ("Formatierungen") or "Images" ("Bilder") once with your right mouse button, and say "include existing files". That will open a popup window to enable you to browse to the location of the web project.

So! WE ARE DONE! You just finished your first Ebook! Well, not a whole Ebook, but an article inside an Ebook. Just click File => Save (as) now, and give your Ebook a Title.

## Conclusion

For those of you who want a short list to work with, here are the most important things:

- Resize your pictures to max. 300 px of height (more is not recommended here!)
- Collect all Content in Plain Text ("Bilder") once with your right mouse button, and say "include existing files".
- Create a "web-project" with three folders named: "Images"; "Styles" and "Texts"
- Create a CSS-Design for your Ebook (For Help, see the W3C





# MY OPINION

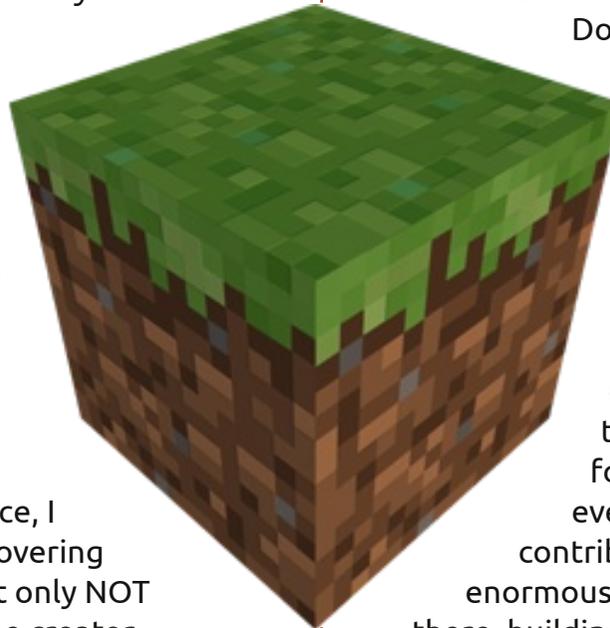
Written by Mattias Nykkel

It's interesting that people prefer to use the wrong tool rather than realize what is best in each situation, and this goes for both Microsoftish guys and Unixish guys. What's the big deal? Why not just ask each other what fun stuff you guys are doing? You answer us truthfully, both the advantages and the disadvantages you might have when compared with the Microsoft community. Then you do the same. Ask us what we see as the big benefit using Microsoft stuff. I am not a purist, rather the opposite. I have just never heard a Linux/Java/Mac guy tell me anything in terms other than to brag. This is not helping me to understand the good stuff in your world.

I can make a qualified guess, but that is not good enough. Let's start talking.

I am 40 years old, living just outside Gothenburg in Sweden. A month ago, my son got into the wonderful world of Minecraft - what I would relate to as what LEGO was for me when I was 8.

Just playing around with it for a couple of days, he discovered something called mods. I realized what it was, but initially thought someone unlawfully had hacked the game. He is 8 years old, and extremely eager to learn new stuff - one could argue too eager! But, with a great deal of patience, I survived. Discovering that it was not only NOT hacked, but the creator approved development kit (MCP - MineCraft Coder Pack) was released! Not only did I want to install some funny mods, but I wanted to create them! This is the first time I have gotten in contact with Java as a programmer. I am a .NET programmer, and had no idea where to start, apart from what was available via the MCP wiki site [http://mcp.ocean-](http://mcp.ocean-labs.de/index.php/MCP_Mod_System)



[labs.de/index.php/MCP\\_Mod\\_System](http://labs.de/index.php/MCP_Mod_System)

I started here, and followed the instructions by the letter. Downloaded jre, the jdk, Eclipse, and the MCP source which contained everything required. People had done this on their spare time, for free. Just so everyone can contribute to the enormous creativity out there, building bigger worlds, more fun stuff, more sick stuff!, but it is fun as hell. I continued following the instructions, decompiling the source code, deobfuscating it, and opened it in Eclipse. A new world opened up... or not.. This is almost exactly the same as working in C#. Some differences like base is called super, etc, but that is just semantics. I could write code just

as if I was building it in Visual Studio. There were some advantages like the simplicity of addressing a folder instead of a Solution file, and some disadvantages like the complexity of addressing a folder instead of a Solution file. If we talked a bit, these problems could be solved! Easily! Hmm....

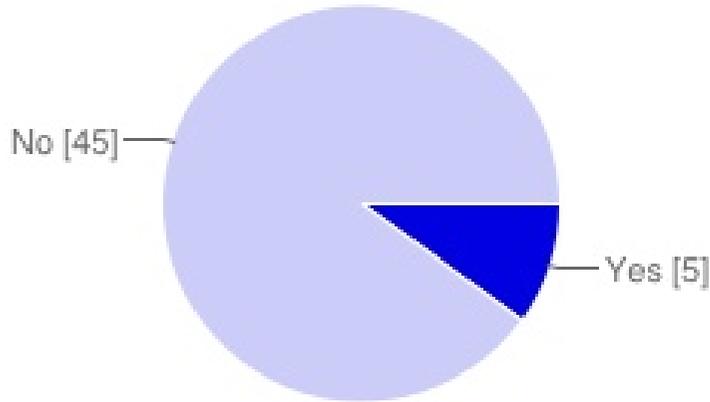
Now, I have created a Christmas mod with Reindeers that explode when killed, Christmas presents that can be crafted and give you lots of random items, but they can explode as well, sometimes. Before Christmas, I will have a Santa flying in with his reindeers and sled every evening when night falls, dropping lots of presents that can be harvested. I just have to figure out how to create the same code in multiplayer mode... Anyone..? The code base is different.

So, to my real question. Can't we just stop sulking, and begin sharing our knowledge? Please...?



# I THINK...

Last month's question was:  
**Have you helped fund a  
Kickstarter game?**



# KICK STARTER .COM

<b>Yes</b>	<b>5</b>	<b>10%</b>
<b>No</b>	<b>45</b>	<b>90%</b>

**T**his month, we wanted to know how many of you have funded a Kickstarter game. The vast majority of those who responded either don't play (computer) games, and/or weren't aware of what a Kickstarter project was. Some of the gamers who answered "No" are the cautious type and prefer to wait and see if the project is worth the time and money. As far as specific titles go, a few of our generous Ubuntu gamers cited *Leisure Suit Larry* and *Steel Storm 2*

as their Kickstarter picks.

**Thanks to everyone who responded, and don't forget to check out next month's poll!**

This month I'd like to ask:

**After the current issue, which [of the two options] would you like to see converted to epub?**

To give your tuppence worth go to: <http://goo.gl/i9CeY>

Closing date for this question is **Sunday 10th June 2012.**



## Coming Soon! The Fullcircle Podcast Returns!

It may be a new team of podcasters, but the format will be the same.

We'll be talking about Full Circle Magazine, news, reviews and interviews.

Your new team are:

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

All are members of the Blackpool (UK) LUG  
<http://blackpool.lug.org.uk>

Keep an eye out for our debut episode, coming **soon.**

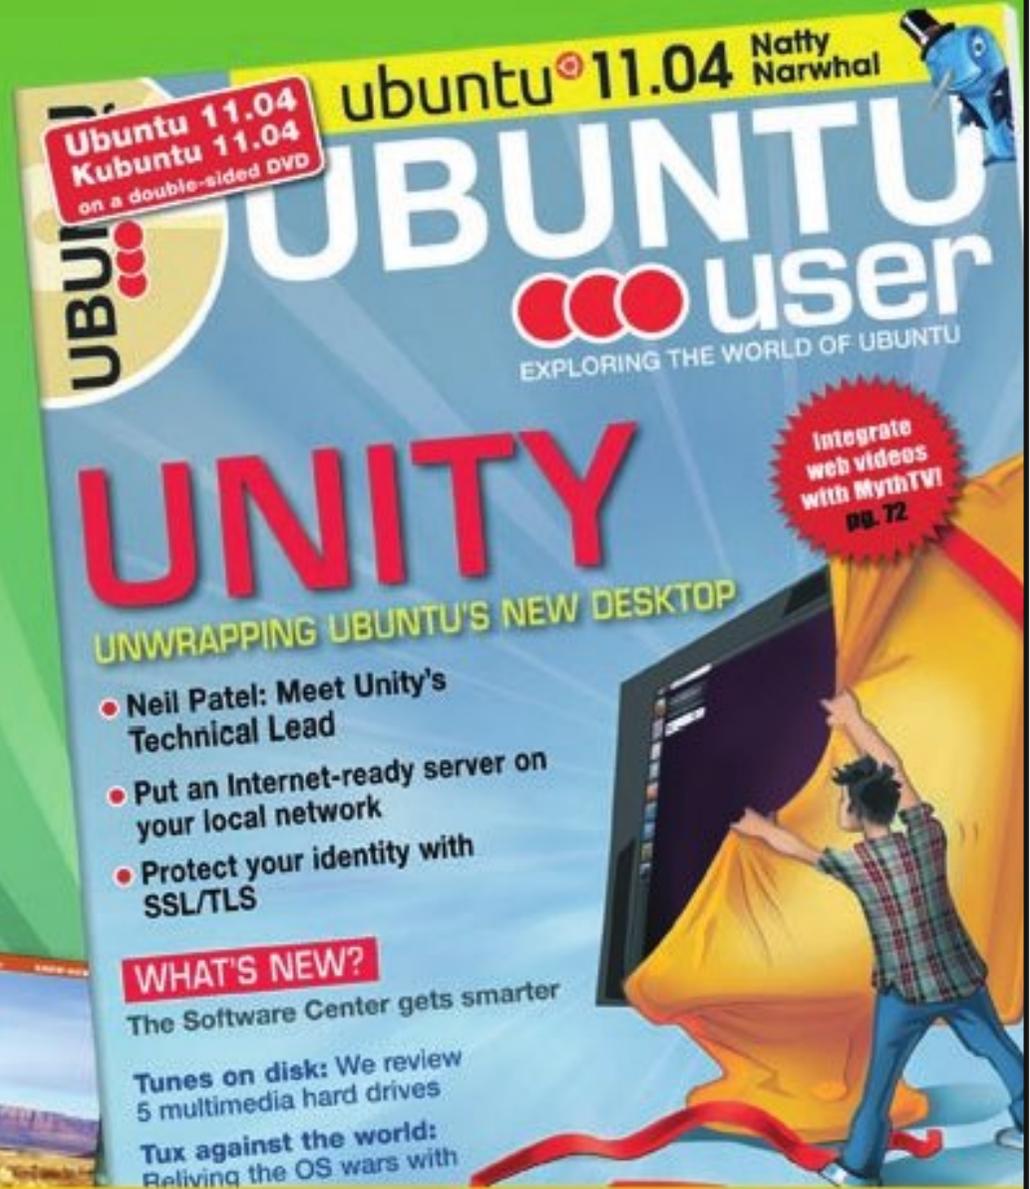


# MORE UBUNTU!

Can't get enough Ubuntu?  
We've got a whole lot more!

**Ubuntu User** is your roadmap to the Ubuntu community. In the pages of **Ubuntu User**, you'll learn about the latest tools, best tricks, and newest developments in the Ubuntu story.

**DON'T MISS ANOTHER ISSUE!**



[UBUNTU-USER.COM/SUBSCRIBE-NOW](http://UBUNTU-USER.COM/SUBSCRIBE-NOW)

FOLLOW US ON

TWITTER: UBUNTUSER

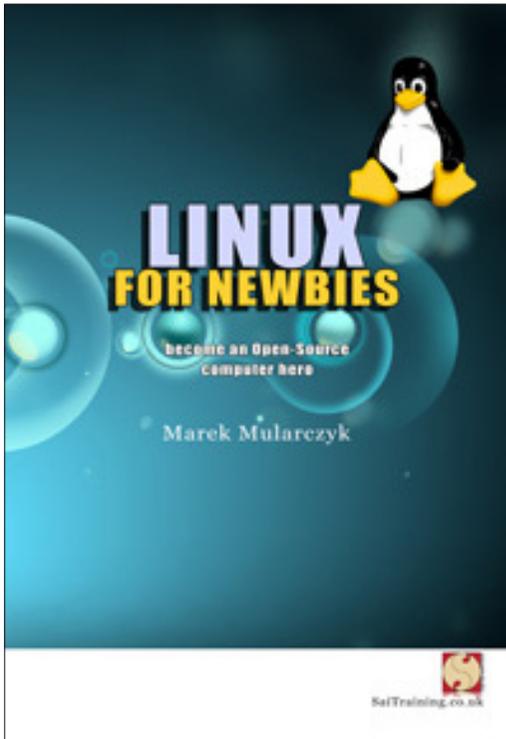
FACEBOOK: UBUNTUSERMAG



# BOOK REVIEW

Written by Ronnie Tucker

# Linux For Newbies



## Linux for Newbies become an Open-Source computer hero

by Marek Mularczyk

210 pages, available in print, PDF  
and EPUB from:

<http://saitraining.co.uk/bookLinux.html>

Printed: £12.99

Printed & PDF/EPUB: £17.99

PDF/EPUB: £7.99

**T**he book begins with a brief history/description of Linux, and then we're into preparing a computer for an install. And it's the installation that pleasantly surprised me about the book. Marek focuses entirely on Kubuntu. My favorite distro! In this edition, he's covering 11.04, but I'm told that in future updates he will be using a more recent version. Nonetheless, it's good to see a book that doesn't use Ubuntu.

Now that you've got Kubuntu installed, the book takes you through the basics of virtual desktops, installing software and such things. Getting you up to speed, Marek dives into showing you various email applications, as well as browsers, telephony, and compression/extraction.

Productivity is discussed using LibreOffice, and a quick runthrough of Dockey - which Marek shows in several of his screenshots. For those not in the know, Dockey is that bar at the bottom of the screen to make your OS look more

Mac-like.

Next come graphics and multimedia. Using brief examples, we are shown GIMP, DigiKam, and Amarok. The book finishes up with a piece on Konsole, and installing more applications.

The last chapter is very handy for those used to Windows; it's here that Marek gives Kubuntu equivalents for Outlook, Internet Explorer, MSN, Skype, NotePad, Nero, Illustrator, Photoshop, and more. Very handy.

All in all, this is a good book. It does have a couple of downsides - the main one being that it's based on an old version (11.04), but I'm told future versions will address this. Marek also switches between standard desktop screens (what you see after install) to his own desktop (with Dockey) which might confuse some, but it's a minor quibble.

If you want a nice introduction to Kubuntu, but don't want to mortgage your home to buy it, nor

have it two inches thick and injure your back lifting it, this is an excellent choice for beginners. Not only are you getting a helpful book, but you're helping a self publishing author.



## SPECIAL OFFER

As a Full Circle reader you can buy the EPUB/PDF of Marek's book for just £4.99 by clicking this link: <http://goo.gl/Cyo2l>

## COMPETITION

To enter win a copy of the book, simply answer this question: **which distro is used in the book?**

Email your answer to: [competition@fullcirclemagazine.org](mailto:competition@fullcirclemagazine.org)

**The competition will close on Friday 15th June, and the winners will be chosen at random. Good luck!**



THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY.



## Join us on:



[goo.gl/FRTMI](http://goo.gl/FRTMI)



[facebook.com/fullcirclemagazine](https://facebook.com/fullcirclemagazine)



[twitter.com/#!/fullcirclemag](https://twitter.com/#!/fullcirclemag)



[linkedin.com/company/full-circle-magazine](https://linkedin.com/company/full-circle-magazine)



[ubuntuforums.org/forumdisplay.php?f=270](https://ubuntuforums.org/forumdisplay.php?f=270)

## Table Of Contents

I am looking for an easy way to find articles in your magazines without having to pull out a dvd and look through each issue. I love all the articles and tutorials, but its so hard to remember which stories are in which issue. I am running Natty (11.04) with gnome classic desktop on an old inspiron 1545 laptop.

**Howard**

Ronnie says: *If you look at the top of the FCM site (<http://fullcirclemagazine.org>), you'll see a link to our 'Table of Contents' which is quite regularly updated by Robin Catling.*

## Tune Out, Drop Out

Regarding FCM#58 (Tune out, Drop Out, Get things Done, by Allan J Smithie), he is absolutely right - Tune out, Drop Out, Get things Done - is the only way to be more productive in doing and learning to make better use of computers, be it at home or at work! The main problem I've found is getting used to the sudden silence! lol.

**Rodney Shinkfield**

## No Java

The live-CD/DVD of Ubuntu has no Sun/Oracle Java; so it is impossible to contact my bank (in

Europe) and other official offices with such a disc. Although it works fine with Linux Mint Debian 64-bit Live CD. Maybe with other Linux distros it works?

That is a very great disadvantage when using the live-edition with banking - especially for security reasons.

**René Gruneisen**

## 24hr Rule

First off, I enjoyed FCM#58 of Full Circle. I would like to point out something that Riku Järvinen stated in his article regarding bsnes. I don't know about his country, but

some caution when relaying information on said subject matter.

**Thomas Holbrook II**

## Tablet Trouble

I typically try to take "My Opinion" with a grain of salt, but I felt really compelled to comment after reading the "My Opinion" by Art Schreckengost in FCM #60.

First, I'd like to say thanks for stepping up and writing an article for Full Circle. I know that it is a struggle to get articles composed, and every article is appreciated. Secondly, I appreciate Art's experience and years of wisdom. I think we can benefit from Art's experience in the Linux world, and we all can appreciate hearing about his experiences with tablets. My complaint however is that the article was written as an opinion piece on "tablet devices," yet Art purposely avoids the largest player in this market space, which is Apple.

I understand the point as it is not an "Android" device ... but,

none-the-less, it is a tablet, and actually does quite well to counter all (most) of the issues that Art has with Android. I would have been fine with this argument if nearly eight paragraphs of this article had not be devoted to discussing the Windows tablet offerings. I thought we were sticking to Android for the definition of tablet? So why are we discussing the Windows tablet for so many paragraphs?

Art asks the question, "will tablets progress far enough in the evolutionary chain to survive, or will they go the way of previous but now doomed technology wizardry?". I think by excluding Apple and Microsoft from the discussion, then, yes, maybe "tablets" do seem doomed but when you look at the analysts discussing the full market, you see that the tablet market is quite large and quite lucrative.

Ok, so onto my last issue, Unity [not Unity again! - Ed]. So, Art's gripe with Unity is that it "may be great for tablets ... but it has severe shortcomings when used on computers with a lot of programs". Aren't we discussing tablets? I think Art countered his own

argument in saying that "Unity may be great for tablets". Great, that's what we are talking about, and, for the Ubuntu fans, that's what we were hoping to hear.

So, that being said, I doubt we are looking at a near "curtain call" for tablets. I think it is fair to question Canonical's play in the space (although they haven't officially made it to the game yet) and I think it would be wise to consider the entire market - and not just one small segment - when trying to decide the fate of a booming technology market.

**Brad Tomy**

## f.lux and Redshift

In the Q&A section of FCM#60 the subject "Two Tips and Techniques - Day and Night" mentioned f.lux



Much better and especially designed for Linux is Redshift: <http://jonls.dk/redshift/>

**Ludo Beckers**



## My Sad Inbox

My inbox is looking rather empty again folks. Surely you have something on your mind that you want to say. Doesn't need to be just about Ubuntu. It can be about any of the other \*buntu's, in fact, it can be about anything Linux.

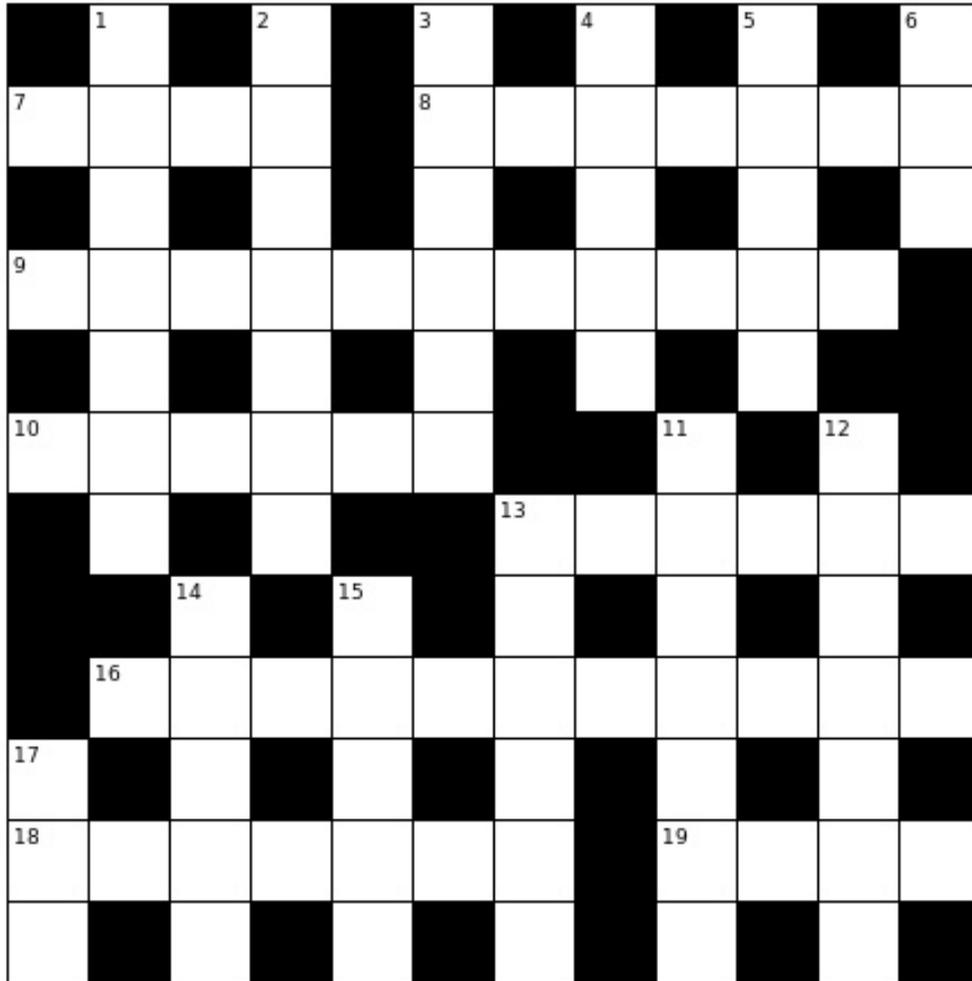
Whether you've written an app, need help with something, have a quick (non-technical) question, or just want to show us something you've done, this is the place for it.

Send it to: [letters@fullcirclemagazine.org](mailto:letters@fullcirclemagazine.org)



# PUZZLES

If you would like to submit a puzzle for publication, please email it to: [letters@fullcirclemagazine.org](mailto:letters@fullcirclemagazine.org). Solutions are on the second last page of this issue. No peeking!



### Across

- 7 Cartridge, say, filled first and last with double magenta, initially (4)
- 8 Recipes arranged for 13 release (7)
- 9 Already bundled in sent message involving concerning group (11)
- 10 The first woman included to solve kernel progress (6)
- 13 Rectified a mistaken 12 (6)
- 16 Optician pal ordered something from Software Centre (11)
- 18 User friendly environment - in soft English notation - beginning to develop Linux icon (7)
- 19 disarray("neon") returns NULL (4)

### Down

- 1 Make better daemon to perform random walk (7)
- 2 Create object code from small computer stack (7)
- 3 A cheap sort of server (6)
- 4 Immature Greek characters? (5)
- 5 Text Editor, German, has love for moving pictures (5)
- 6 Document network taken up centrally from, maybe, worldwide (3)
- 11 Changing when one metal guitarist initially supports head journalist (7)
- 12 Correcting mistakes leads to downfall (7)
- 13 Young, or about 100? (6)
- 14 How the best kind of source begins? (5)
- 15 Programming language has strong support - not limited to the desktop (5)
- 17 Fitting package manager (3)

Crossword by Marlon Mann

The solutions to all the puzzles are on the second last page of this issue. **No peeking!**





O'REILLY®  
**Velocity**  
Web Performance  
and Operations  
CONFERENCE

June 25–27, 2012 | Santa Clara, CA  
Register Now & Save 20% with code FULLCIR

2–4 October, 2012 | London, England  
Registration opens June 2012

See [velocityconf.com](http://velocityconf.com) for more details.



*"Velocity is the conference where people talk about how to get things done in the real world—if you want to know how the best in the world handle their operations, Velocity is the place to learn."*

—ADAM JACOB, OPSCODE

**Velocity is much more than a conference;** it's become the essential training event for web professionals from companies of all sizes.





# Q&A

Compiled by Gord Campbell

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

**Q** I want to try Ubuntu, but I'm worried about the valuable pictures, documents, music, and videos on my computer.

**A** Before you do anything, go and buy an external drive, and copy all your valuable files onto the external drive. Or buy a spindle of DVD-R discs, and burn copies of your files onto DVD.

**Q** Where can I find information about what's going on when Ubuntu boots?

**A** <http://upstart.ubuntu.com/cookbook/>

**Q** If I do `gksudo nautilus` in 12.04, the desktop background changes from what I have selected to the Ubuntu default wallpaper, and the file manager is

then handling drawing of the desktop.

**A** Try:  

```
gksudo 'nautilus --no-desktop'
```

**Q** I'm not getting any sound from VLC in Mint 12. All the other sound apps work fine.

**A** Open VLC, select Tools, Preferences, the Audio tab. For "Output module," select ALSA. For device, select the appropriate one. Click on Save and close the window. Or, run software manager and install the VLC "pulse" plugin.

**Q** The boot process stops at a screen that says "40GB drive missing. Continue to wait or press S to skip mounting, or M for manual recovery". How can I stop

this message from appearing?

**A** (Thanks to *audiomick* in the Ubuntu Forums)  
Open a terminal and enter this command:

```
gksudo gedit /etc/fstab
```

Find the line which includes "40GB" and turn it into a comment by inserting a number-sign as the first character. Save the file.

**Q** When I run Disk Utility, I get the message, "*The partition is misaligned by 1024 bytes. This may result in very poor performance. Repartitioning is suggested.*" I have a new "Advanced Format" drive.

**A** Use the latest GPartEd Live to partition the HDD to the nearest MB. The partition needs to begin on a 4K boundary, which is 8 of the old 512-byte sectors. It really does affect system performance.

**Q** I have a Lenovo laptop. The function-keys default to controlling brightness, volume etc. I would rather have them be F1, F2 etc.

**A** To change this in the BIOS, see this page: <http://forums.lenovo.com/t5/ThinkPad-Edge-Knowledge-Base/Fn-Fx-remapping/ta-p/237588>

**Q** Can I use Skydrive from Ubuntu?

**A** Yes, there is a web interface which allows dragging files into Skydrive, and you can download from Skydrive. You will need a Windows Live ID, such as a Hotmail account.



**Q** I have a Belkin Components F5D8053 N Wireless USB Adapter v3000, and when I upgraded to 12.04, it doesn't connect to my wireless router.

**A** (Thanks to **Chili555** in the Ubuntu Forums) Open Terminal, and enter this command:

```
sudo modprobe rt2800usb
```

When that works, add the driver permanently with these commands:

```
sudo su
echo rt2800usb >> /etc/modules
exit
```

**Q** I am trying to use Remote for VLC on my Android phone, but I get "connection refused."

**A** (Thanks to **Steve Beresford**) You need to edit VLC's "hosts" file, to include the IP address of your Android phone:

```
gksudo gedit /etc/vlc/http/.hosts
```

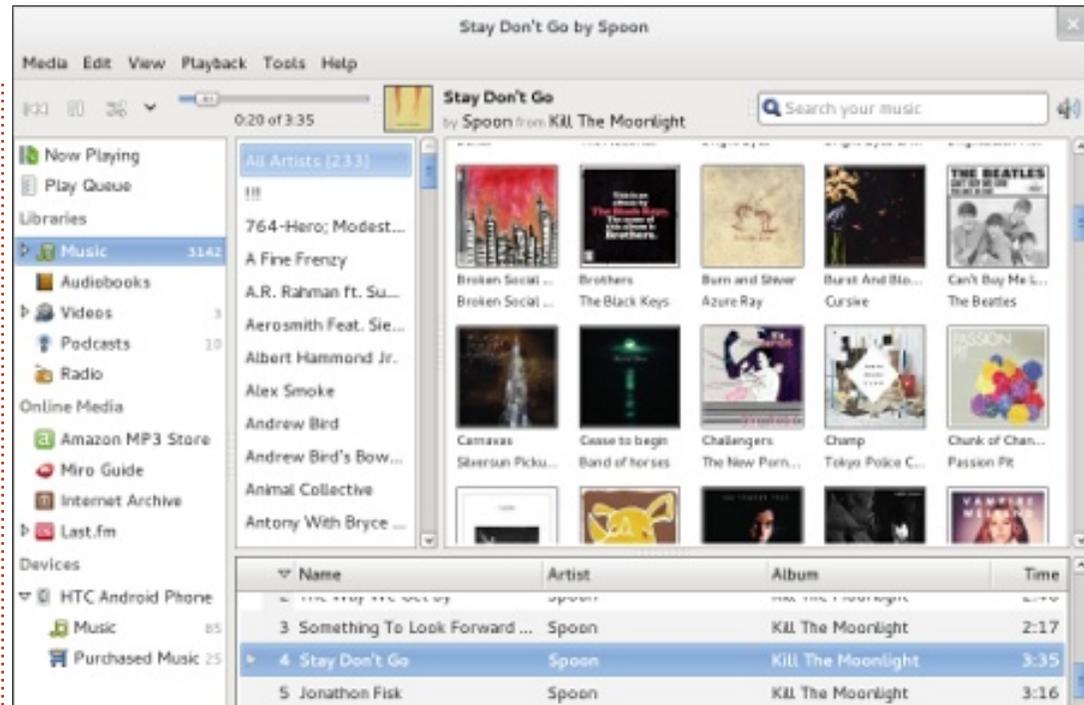


**Q** I have the Nvidia driver 295.40 installed in Ubuntu 12.04. Using Banshee, video playback stutters randomly and uses about 50-60% of CPU, and VLC is even worse.

**A** (Thanks to **Christopher Suttles**) I ran:

```
sudo apt-get install vdpau-va-driver vainfo libva
```

CPU usage is the same with Banshee, but video playback is now flawless.



**Q** Is there a weather applet for Ubuntu 12.04?

**A** Yes, install indicator-weather from the software center, then run it. Provide your location and other preferences the first time you run it.

**Q** Please give me step-by-step instructions on how to connect my Beagleboard to my

Ubuntu computer using an ad-hoc wireless connection.

**A** Step one: go to the Google web site

2. search for: *ubuntu ad-hoc wireless network*

3. go to the first link: <https://help.ubuntu.com/community/WifiDocs/Adhoc>



In this new section, Nicola Cappellini will introduce you to not only music sites, but also recommend some tracks across a wide range of musical genres.

There's lots of Creative Commons music out there, here is a list of some of the more interesting sites.



## CcMixer

([ccmixter.org](http://ccmixter.org))

A treasure trove of samples and sound clips. Browse through the library and download whatever you want, listen to other users being creative. This is a wonderful resource to practice mixing and mastering. The featured samples are excellent quality, and many of the non-featured samples are also great. How to make a remix: download the audio, have your way with it, upload it back to ccMixer, and be a part of the community.



## AudioTool

([audiotool.com](http://audiotool.com))

Launch the app on the site, and start making music right away. Drag loops, synths, and drum sounds into an effects chain of your making. Does the flange effect sound better before or after the tube distortion? You make the call. The interface is really slick, arrange the boxes however you like and take control over every parameter in the mix. When you're done export the audio and see your work on the audio tool site.



## WolframTones

([wolframtones.com](http://wolframtones.com))

Choose a starting pattern and an algorithm. Then play with the instrumentation and musical mode, and hear the result back instantly. Electronic works created on the site can be distributed under an Attribution-Noncommercial-No Derivative Works policy. Admittedly restrictive, I wanted to

include this in the list because it's so interesting! Midi sounds can be saved, emailed, and even exported as cellphone ringtones. This "compositional engine" is a great tool for inspiring philosophical discussions about the nature of music, compositional methods, musical perception and intentionality. John Cage would certainly approve!



## SoundCloud

([soundcloud.com](http://soundcloud.com))

The waveform playbox-- you've probably seen it before. Browse through the uploads or explore the tags to listen to the music. Users can publish under any creative commons license they wish, and can also choose to make the tracks available for download or not. Listening is free even for non-users, premium accounts allow for more uploads (measured in minutes), downloads of your music, stats, and more tools for sharing.



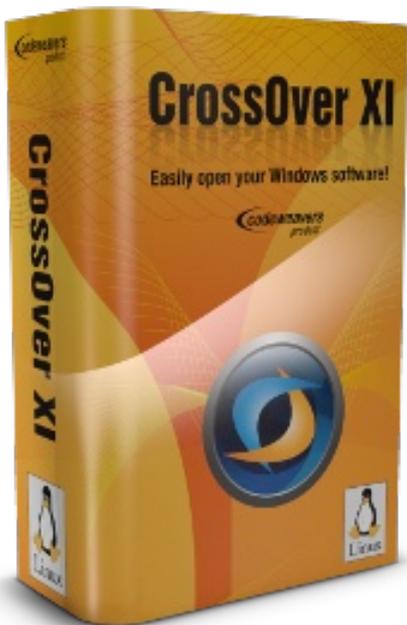
## Free Music Archive

([freemusicarchive.org](http://freemusicarchive.org))

Just as the radio made music available to anybody with a radio, the Free Music Archive provides the same service for internet users. Browse through the interactive library, and discover new tracks meticulously organized by usage trends, genre, as well as by curator. Information about the tracks and artists is good for getting an idea of the context in which the music was made.



**Nicola** is a freelance musician, translator, and web-whiz. He can be seen on the stage, behind the scenes, and in the zone - for whatever task is at hand.



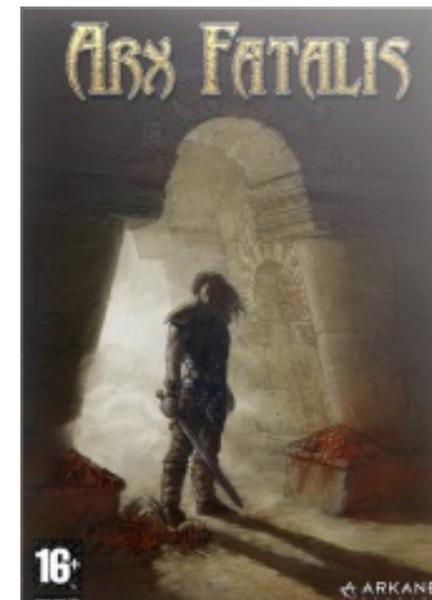
renewing the subscription of the app and the purchase fee. Of course, if all you bought was the old Crossover Games app, then the price has gone up for you, I believe, and come down for everyone else who bought Crossover Standard or Crossover Pro: the starting price is US\$39.95. The price increases as you add more support/upgrade time, maxing out at US\$59.95 - which gives you 12 months, I believe, <http://www.codeweavers.com/>

**C**odeweavers, the people responsible for the **Crossover** application and the Open Source WINE project, have discontinued support for Crossover Professional, Crossover Standard, and Crossover Games (basically the payware versions of WINE), and have now combined all three tiers into one program called Crossover (catchy name eh?), and have reduced the price of both

# ARX LIBERTATIS

**T**here has been a fan-port of a Windows-only game to Linux recently; the Arx Libertatis Project has successfully ported and updated the graphics of the game **Arx Fatalis** to Linux, Windows,

FreeBSD, and Mac (<http://arx-libertatis.org/>). The game SDK was officially released as open source, but not the game data, so you will need a copy of the game/demo to play it. The game can be purchased at [Gog.com](http://Gog.com), and follow the instructions on how to extract it. No need to compile Arx Libertatis yourself as there is a PPA for it.





This month, we take a brief look at 0AD, a free open-source real-time strategy game of ancient warfare. It has been under development by Wildfire games (<http://wildfiregames.com>) since 2001, while the latest stable release was announced about a month ago. Focusing on the history of Western civilization, the game might become a noteworthy competitor for commercial titles such as the Age of Empires and Rise of Nations.

## Installation & Features

Since Precise 12.04 LTS, 0AD is available directly from the Universe repository. Using the command line, simply type:

```
sudo apt-get update && sudo apt-get install 0ad.
```

If you're interested in development versions, check out the PPA from the 0AD homepage (<http://wildfiregames.com/0ad/>).

For an open-source project, and

by a team of hobbyist game developers, 0AD has a load of features - including unique civilizations, versatile units, and extended moddability (with GPL-available source code and a scenario editor). In addition, a great deal of time and effort has been spent on 3D graphics modelling, and research in pursuit of historical realism. One thing I especially liked about the units is that you can use them interchangeably for collecting goods and combat, a feature not readily available in many RTS titles.

## Gameplay

I was pleasantly surprised by the fact that 0AD runs very smoothly for an alpha version. It shows that the game has been under development for quite some time since there are no major bugs present. This, however, means the choices you can make are somewhat limited, i.e. there are no Campaigns, and the Options panel does not yet work at all, so you just have to use the default settings. For lower-end machines, this might



cause a problem; I noticed a slow-down on my rig (E6750, 4GB DDR2, GTS 8800) when tens of units were battling each other. The minimum requirements (1 GHz processor, 512 MB RAM and GeForce 3) stated on the homepage probably haven't been updated for some time.

## Conclusions

All in all, 0AD is a promising candidate which has great potential as a standard RTS, but doesn't offer anything very special to the genre. Beautiful graphics, historical realism, smooth and versatile gameplay, and all combined with cross-platform support, are likely to attract players who enjoy history and strategy. However, because the project is still in alpha, there is

much to do for a fully-fledged RTS. While there are quite a few scenarios and the possibility for customization and multiplayer (peer to peer), the lack of campaigns and interesting storylines doesn't really do the game justice. Nevertheless, this seems to be the core problem with open source games in general, and I truly hope that developers would start to pay more attention on storylines rather than graphics, etc. I probably over-emphasize this since I'm a fan of great stories.



**Riku Järvinen** (rierjarv) is a CS major student from Finland who delves into the Linux and Open Source gaming world once in a while.



# UBUNTU GAMES

Written by Dougn Redhammer

## Oil Rush

**O**il Rush (<http://oilrush-game.com/>) by Unigine, a real-time strategy game first announced in the fall of 2010, was released this January for the Linux, Mac and Windows operating systems.

Is Oil Rush worth your time after an almost two-year wait?

The Game features three gaming options:

- Campaign, aka the Story Mode
- Quick Game (for those who want to just wet their beaks with a quick play of the game)
- Multiplayer, consisting of LAN and Online.

This review will focus on the Campaign Mode of Oil Rush. The Campaign Mode has a total of four chapters - with each chapter containing three to five missions, and, depending on your playing style and at what difficulty the game is set, you will have around fifteen to twenty-three hours of game play.

First off, the game looks great

and runs great - provided you have a modern pc; anything more than three to four years old that doesn't have a hardware upgrade in RAM/CPU/video card will struggle to run this game.

At its core, Oil Rush is a very basic RTS game, but, being basic doesn't mean that it's not fun to play.



As the game progresses, and as with all RTS games, you get new types of units to command. However, unlike other RTS games, you don't get to carry over anything you have learned via research, nor do your units keep their upgrades. Every new mission in Oil Rush sees your units starting anew from scratch. Each mission

requires the player to spend research points differently; one mission may see you spend points in beefing up the defense of your units and bases, while another may call for you to spend points in weapons research, sabotage and speed upgrades, and yet another mission may call for you to mix all your research points into both defense and offense. However, your research of both offense and defense will not be as in-depth as if you had focused on just one.

Research Points are earned via combat and/or successfully capturing bases/facilities. Since there is a limited number of bases and enemy units, you have to carefully think about where you wish to spend your points; invest too heavily in defending your bases and providing support for your units and you may find that, with no weapons research done, you are able to only hold the enemy to a stalemate. On the other hand, invest too heavily in weapons and the speed of your units, and you may lose your bases as you didn't invest in defense.

As was said before, Oil Rush is a basic RTS, so, as a result, there is no true base building within the game. Instead, there are bases that you can capture, and then upgrade their defenses. Each base represents a different unit that you can capture, categorized into heavy, medium and light builds, and whether they are good or vulnerable against other units.

As the name of the game implies, Oil is key to your survival - it is the resource that you depend on for the upkeep of your units and bases. Without oil, you will be unable to create defenses or use any of the active skills you learn via research. Some of these skills include sabotaging enemy bases, propaganda to increase the defense of your own base/units, deploying support vehicles that can repair your units, increasing the production speed of your units, seeing what the enemy is up to on an unexplored region of the map, or even nuking your enemies. Oil can be obtained via the capturing of oil rigs. However, defense turrets, etc, cannot be set around oil rigs, so it's up to the player to decide what units to deploy around them - you can



easily lose ownership of an oil rig within the blink of an eye.

All in all, the Campaign Mode of Oil Rush is a fun ride, good for the asking price of US\$19.99. Which, if purchased via the Unigine Store, nets you the Linux, Mac, and Windows DRM-free versions, and also the Steam version.

### Pros :

- Good Game to get your feet wet with if you have never played an RTS game before.
- Beautiful Graphics.

### Cons :

- Requires a modern PC/Mac; not for an old computer more than three to four years old.
- For RTS veteran Gamers, the lack of options will be off-putting.

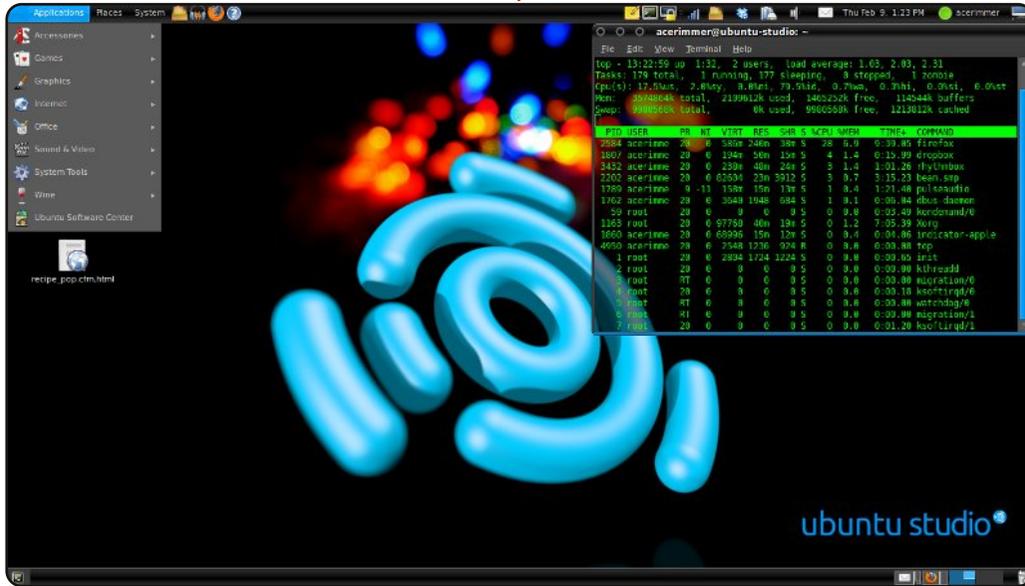
**Oil Rush** is available for purchase via the Ubuntu Software Center, Desura (<http://www.desura.com/games/oil-rush>), and The Unigine Store (<http://unigine.com/products/oilrush/>).





# MY DESKTOP

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

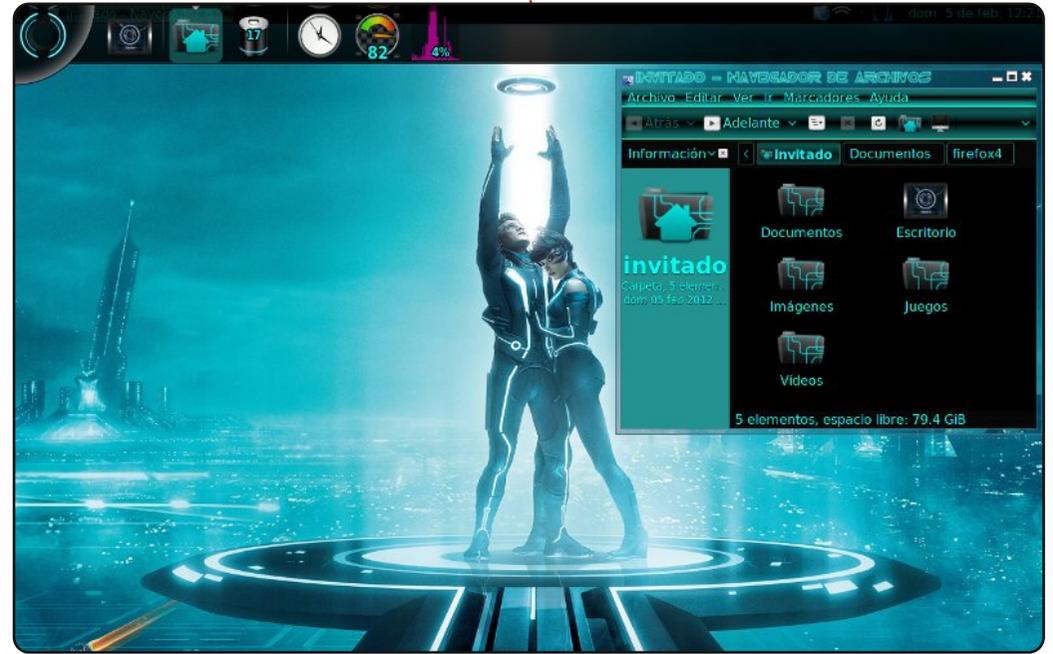


My laptop runs Ubuntu Studio 10.04. I create podcasts in Audacity, edit videos with Openshot, and edit graphics with Gimp.

I'm using the Ubuntu Studio theme and icons. The wallpaper is called "eternalstudio," and was originally created in pov-raytrace by Garry Parker of ubuntuatanic fame. With Mr. Parker's permission, I updated this and his ubuntu logo wallpaper with new Ubuntu logos. I've made these wallpapers available for download at: <http://tinyurl.com/Ubuntustudio-Wallpaper>

My computer is a Dell 1545 Inspiron with an Intel dual-core at 2.3GHz, 4 GB of ram and an Intel GMA 450 gpu

**Charles F. Howlett**



I'm a fan of the movie TRON:Legacy, so I decided to theme up my desktop around the movie. As can be seen, the icons, the controls, and the window borders are according to the movie.

The font in the window border is not part of the original emerald theme, so I changed it. Also I'm using Avant Window Navigator as a panel, so I can quickly access my files.

My OS is Ubuntu 10.04, and I'm using a Compaq Presario C700 with 1GB of RAM and a 120 GB hard disk.

**Christian Ali Morales Hernández (aka: Kerochris)**



I have handmade wallpaper with the Ubuntu logo and metal-like colored background.

I have used Ubuntu since 11.04, and, after a few weeks, it is now my main System. I like this system. It is free, and, best of all, it doesn't need any tweaks. It works right out of the box.

### System:

Ubuntu 11.10 with default Unity theme

Dell Inspiron M5010

AMD Turion x64 2.5 GHz

Mobility Radeon v550 1GB

RAM 4GB

HDD 500GB

**Evgeny Kozlov**



I have used Linux since 7.10. So, here I am showing you my Linux Mint 11 Desktop.

I love using Mint, since it's an Ubuntu derivative and has many codecs included. I love simplicity and love Ubuntu Classic desktop, but not Unity or Gnome Shell. I'm using Nautilus Elementary, replacing default Nautilus, it's much simpler and nicer.

The screenshot was taken on my ASUS A43E, 2G RAM, CPU Intel Pentium B950 2.1GHz. I've also installed this Mint on my ACER AOP 531 netbook, and on my PC (RAM 2G, CPU Dual Core 1.6G, Mainboard ASUS P5G41T).

**thom**

4	6	A	C	2	E	0	1	5	8	3	F	7	B	D	9
1	3	B	9	4	6	8	C	E	7	2	D	0	A	5	F
0	8	F	E	3	7	D	5	C	B	A	9	6	2	4	1
7	5	D	2	B	9	A	F	4	1	0	6	E	3	8	C
D	A	6	1	7	B	5	4	F	C	8	E	3	0	9	2
9	E	2	F	A	8	C	D	B	3	4	0	5	6	1	7
C	7	5	3	9	F	6	0	D	A	1	2	B	8	E	4
8	0	4	B	E	1	3	2	9	5	6	7	F	C	A	D
A	9	0	8	C	2	E	B	1	D	F	3	4	5	7	6
5	4	3	D	0	A	7	9	2	6	C	B	1	E	F	8
2	B	E	6	F	5	1	8	0	9	7	4	C	D	3	A
F	1	C	7	D	3	4	6	8	E	5	A	2	9	0	B
E	F	8	A	6	C	B	7	3	0	9	1	D	4	2	5
3	2	7	5	1	D	9	A	6	4	E	C	8	F	B	0
6	D	1	4	5	0	F	3	A	2	B	8	9	7	C	E
B	C	9	0	8	4	2	E	7	F	D	5	A	1	6	3

16X16 SUDUKO



	I		C		A		B		V		W				
A	M	M	O		P	R	E	C	I	S	E				
	P		M		A		T		D		B				
P	R	E	P	A	C	K	A	G	E	D					
	O		I		H		S		O						
E	V	O	L	V	E				E		U				
	E		E				R	E	D	O	N	E			
		O		C			E		I		D				
	A	P	P	L	I	C	A	T	I	O	N				
A		E		O			E		I		I				
P	E	N	G	U	I	N			N	O	N	E			
T		S		D		T			G		G				

CROSSWORD

S	Q	U	A	D	R	O	N			E	S	P	Y		
	U		D		E		A		X		U				
R	I	T	U	A	L		S	U	P	E	R	B			
	P		L		A	J	A		L		G				
S	P	I	T		T			L	O	O	K	E	R		
	E				E				R						
D	R	A	P	E	D			F	R	E	E	S	T		
			I					L				T			
S	A	L	A	A	M		A		A	V	O	W			
	M		N		A	S	S		G		R				
S	P	H	I	N	X			H	U	R	R	A	Y		
	L		S		I			E		E		G			
Z	E	S	T			M	A	S	T	E	R	E	D		
1	2	3	4	5	6	7	8	9	10	11	12	13			
N	R	U	H	Q	J	F	Y	B	X	P	C	V			
14	15	16	17	18	19	20	21	22	23	24	25	26			
T	G	M	W	I	A	K	L	S	Z	E	D	O			

CODE WORD

ABCDEFGHIJKLMNOPQRSTUVWXYZ

Codeword and 16x16 Suduko puzzles are copyright, and kindly provided by, **The Puzzle Club** - [www.thepuzzleclub.com](http://www.thepuzzleclub.com)



# HOW TO CONTRIBUTE

We are always looking for new articles to include in Full Circle. For article guidelines, ideas, and for issue translation, please see our wiki:

<http://wiki.ubuntu.com/UbuntuMagazine>

Please email your articles to: [articles@fullcirclemagazine.org](mailto:articles@fullcirclemagazine.org)



If you would like to submit **news**, email it to: [news@fullcirclemagazine.org](mailto:news@fullcirclemagazine.org)

Send your **comments** or Linux experiences to: [letters@fullcirclemagazine.org](mailto:letters@fullcirclemagazine.org)

Hardware/software **reviews** should be sent to: [reviews@fullcirclemagazine.org](mailto:reviews@fullcirclemagazine.org)

**Questions** for Q&A should go to: [questions@fullcirclemagazine.org](mailto:questions@fullcirclemagazine.org)

**Desktop** screens should be emailed to: [misc@fullcirclemagazine.org](mailto:misc@fullcirclemagazine.org)

... or you can visit our **forum** via: [www.fullcirclemagazine.org](http://www.fullcirclemagazine.org)

## FULL CIRCLE NEEDS YOU!

A magazine isn't a magazine without articles and Full Circle is no exception. We need your Opinions, Desktops and Stories. We also need Reviews (games, apps & hardware), How-To articles (on any K/X/Ubuntu subject) and any questions, or suggestions, you may have.

Send them to: [articles@fullcirclemagazine.org](mailto:articles@fullcirclemagazine.org)

## Full Circle Team



**Editor** - Ronnie Tucker

[ronnie@fullcirclemagazine.org](mailto:ronnie@fullcirclemagazine.org)

**Webmaster** - Rob Kerfia

[admin@fullcirclemagazine.org](mailto:admin@fullcirclemagazine.org)

**Comms Mgr** - Robert Clipsham

[mrmonday@fullcirclemagazine.org](mailto:mrmonday@fullcirclemagazine.org)

**Podcast** - Robin Catling

[podcast@fullcirclemagazine.org](mailto:podcast@fullcirclemagazine.org)

## Editing & Proofreading

Mike Kennedy

Lucas Westermann

Gord Campbell

Robert Orsino

Josh Hertel

Our thanks go out to Canonical, the many translation teams around the world and to **Thorsten Wilms** for the current Full Circle logo.

**Deadline for FCM#62:**  
**Sunday 10th June 2012.**



**Release date for FCM#62:**  
**Friday 29th June 2012.**

