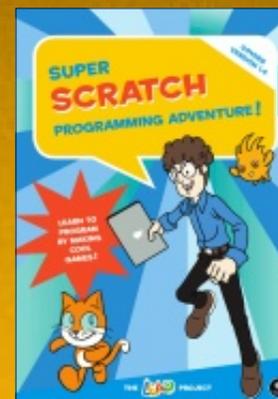




Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

ISSUE #68 - December 2012



BOOK REVIEW:
Super Scratch
Programming



BLENDER SERIES

LEARN 3D MODELLING AND ANIMATION

HowTo



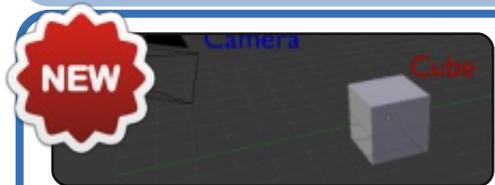
Programming Python 39 p.08



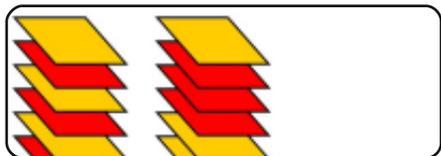
LibreOffice - Part 21 p.12



Install 12.10 On EEE PC p.15



Blender - Part 1 p.19



Inkscape - Part 8 p.22

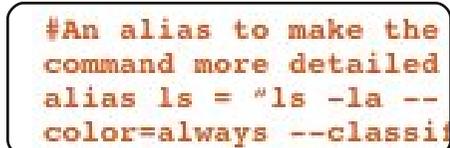
Graphics



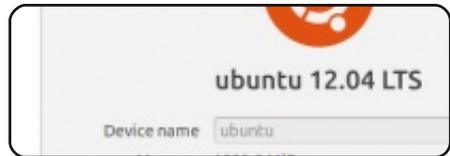
Full Circle

THE INDEPENDENT MAGAZINE FOR THE UBUNTU LINUX COMMUNITY

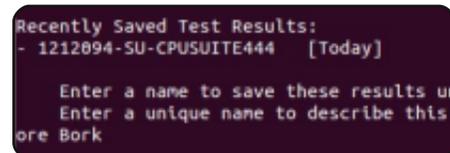
Columns



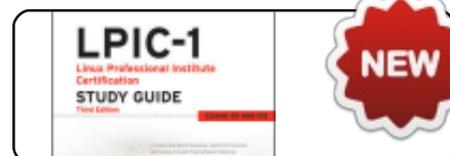
Command & Conquer p.06



Ask The New Guy p.31



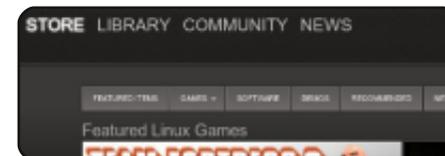
Linux Labs p.33



Ubuntu Certified p.61



Ubuntu News p.04



Ubuntu Games p.60



Q&A p.52



Closing Windows p.35

Opinions



My Story p.37



What Is... p.43



Review p.46



Letters p.50



Web Dev p.26

@ Web Dev



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

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



Welcome to the final issue of Full Circle!

Yes, that's right, this is indeed the last issue of Full Circle Magazine. The last of 2012*. We've got another jam-packed issue for you to close off 2012; Python and LibreOffice continue as ever, and we'll show you how to squeeze Ubuntu 12.10 into an old EEE PC. Lucas reviews the book *Super Scratch Programming*, and I take a look at *Linux Mint Debian Edition KDE*. If it's graphics you're after, then we've more Inkscape for you, and a new series where we look at Blender. Yes, we finally have a Blender series. Many thanks to Nicholas for stepping up to the plate.

Speaking of 'new series', we have Linux Certified – which discusses the trials and tribulations of learning Linux to become a Linux Certified Professional.

Gord ploughs his way through your questions in his Q&A, and we have a special Q&A this month with Didier Roche of Canonical. We asked you folks to send us questions you have about the making of Ubuntu, and you're certainly a quizzical bunch. Didier sat down and answered a slew of questions, and we present, to you, his unedited answers.

Wherever you are, and whatever you're doing, I want to wish you all a happy holiday season. I also want to thank all those who help me make Full Circle every month with their editing and writing, and you folks who submit articles.

All the best, keep in touch, and I'll see you in 2013!

Ronnie

ronnie@fullcirclemagazine.org



* Yes, I am going to keep using that joke at the end of every year.

This magazine was created using :



Full Circle Podcast

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

Hosts:

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



<http://fullcirclemagazine.org>



Download

CoC Version 2.0 has been released

Charles Profitt of the Ubuntu Community Council announces the new version of Ubuntu's Code of Conduct, an "important guiding document which forms the basis how organization operates internally and externally."

A summary of changes included in this revision of the Code of Conduct is as follows:

- merged CoC and LCoC, factored in diversity statement too
- generalised it, so other communities can more easily make use of it
- mention CC-BY-SA 3.0 license
- generally reduced duplication
- make it clear that everybody can take action, there's no need to wait for formal recognition
- explain governance principles
- concrete leadership guidelines: added paragraphs about courage, considerateness and decisiveness

The Code of Conduct is

available at the following address:
<https://launchpad.net/codeofconduct/2.0>

New community-announce mailing list!

There are all kinds of announcements in the Ubuntu community posted on the Ubuntu Fridge that are not strictly development or release-related - from the recent revision of the Code of Conduct, to upcoming community-wide events.

At the Ubuntu Developer Summit in October, the Community Council spoke with other members of the community and decided to launch a new community announcement mailing list so community members can receive these announcements by email as well.

The list is now live! Sign up here:
<https://lists.ubuntu.com/mailman/listinfo/community-announce>

Sputnik – Dell's Ubuntu-based developer laptop is here

Mark Murphy, Global Alliances Director at Canonical, announces that Sputnik (Dell's Ubuntu based laptop for developers) is available for purchase. "The Dell XPS 13 is a top spec, high-end ultra-mobile laptop, offering developers a complete client-to-cloud experience. It is the result of the Dell's bold Sputnik initiative which embraced the community and received terrific response from developers around the world." For now, Sputnik is available for purchase only in America and Canada. More information can be found at the following link:

<http://www.dell.com/us/soho/p/xps-13-linux/pd>

<http://blog.canonical.com/2012/11/29/sputnik-dells-ubuntu-based-developer-laptop-is-here/>

Jono Bacon: Ubuntu Community Appreciation Day

Jono Bacon unveils his Ubuntu appreciations a little bit later due to his paternal obligations. However, Bacon complement those mentioned in previous editions of the Ubuntu Weekly Newsletter. Bacon states: "Choosing people for Ubuntu Community Appreciation Day is always tough as we have so many wonderful people who actively participate in our community. From our developers to docs writers to translators to testers to advocates and more, everyone puts their brick in the wall to build a strong, competitive, and proficient Ubuntu. We would be nothing without your contributions."

<http://www.jonobacon.org/2012/11/28/ubuntu-community-appreciation-day/>

13.04 (Raring Ringtail) Alpha 1 Released!

Stéphane Graber, a member of the Ubuntu release team, announces the release of Ubuntu 13.04 (Raring Ringtail) Alpha 1 for Edubuntu and Kubuntu. The release announcement can be found at: <https://lists.ubuntu.com/archives/ubuntu-devel-announce/2012-December/000998.html>

Private Projects and Private Blueprints leave beta

As of the 4th of December, the Private Projects and Private Blueprints features in Launchpad leave their beta period, and are now stable. A commercial subscription is required to use these features, but anyone who creates a Private Project and/or a Private Blueprint will enter into a 30-day trial commercial subscription. <http://blog.launchpad.net/general/private-projects-and-private-blueprints-leave-beta>

Improving Community 'Getting Involved' Documentation

Jono Bacon writes about problems new members face when joining the community and improvements made in this regard. Bacon and his team fleshed out specific improvements at UDS-R, including refining information for new members and moving the documentation to a Wordpress platform. On Thursday, December 13th, Bacon and his team will be hosting a "Docs Jam" to get the process started. More information and guidelines can be found on Bacon's blog and following links: <https://wiki.ubuntu.com/CommunityWebsite> <http://daringfireball.net/projects/markdown/syntax> <http://www.jonobacon.org/2012/12/09/improving-community-getting-involved-documentation/>

Creating An Awesome LoCo Support Community

Jono Bacon gives advice to English-speaking LoCo teams, pointing them to an AskUbuntu.com page with detailed instructions for asking and answering Ubuntu LoCo questions on the site. This additional help is to make LoCo team work, and spreading the word about Ubuntu, more easy. <http://www.jonobacon.org/2012/12/15/creating-an-awesome-loco-support-community/>

Adding privacy setting support to your Unity Lens

Michael Hall shares how third-party developers can add support into their Unity Lens to check the new privacy settings. He shares two methods, including using the new Preferences Manager class in Unity 6 that allows developers access to the settings without going through GSettings/DConf. <http://mhall119.com/2012/12/adding-privacy-setting-support-to-your-unity-lens/>

Restarting the Ubuntu Learning project

Elizabeth Krumbach writes that the Ubuntu Learning project has been quiet for some time. To restart the project, she has been working with some people quietly for several months and now she unveils plans, the vision, and further tasks for this project. Krumbach also calls for volunteers who are willing to join and contribute to this project. <http://princessleia.com/journal/?p=7383>

Canonical Targets Unity Shell Maintenance for 13.04

Stephen M. Webb unveils a different approach to Unity development and fixing. According to Webb, Unity is going to be polished for the next release of Ubuntu. That means a bigger focus on what Unity is now, instead of adding new features. <http://bregmatter.wordpress.com/2012/12/19/canonical-targets-unity-shell-maintenance-for-13-04/>



This month, I thought I'd do something slightly different from the usual tutorial-style articles. Specifically, I'm going to share my experiences with the Synology DS213+ Network Attached Storage (NAS) that I recently purchased and set up, to work with Linux/Windows/Mac OS X, though I'll only be focusing on the Linux side of things. As such, if you already have a NAS set up (or don't intend on getting one), this article probably won't be for you.

The Product

The Synology DS213+ is the "small and medium business" option of the 2-bay NAS systems offered by Synology. As such, it's more expensive than the DS213/DS213j, but offers a few more features. For an exact list of what isn't included in the cheaper versions, you can compare the products on the Synology web page. At first glance it seems that the DS213+ simply offers features like LDAP and Windows AD,

hardware encryption, and slightly better performance. The NAS does not come with hard drives, so you will need to factor that in cost-wise. The model discussed in this article is using two Seagate ST3000DM001 (3TB) drives.

The Setup

Installing the drives into the device is painless and easy (even if it does require screws and a screwdriver). However, make certain to seat both drives properly, as not doing so can result in a drive not being recognized by the system.

Once you've plugged it in and inserted the drives, you'll need to start up the device by pressing the power button. For the software setup, you have two options: using the Synology Assistant software (available for Linux, Mac, and Windows), or use the web interface at <http://diskstation:5000>. Regardless of your choice, you'll need to let the software complete the first-run setup. This includes creating an administrator account.

It's a relatively painless process, if somewhat slow. Once you can log in, you'll need to decide on the storage manager volume you'd like to use. You have the choice between RAID 1 and RAID 0 (or to let the DSM system decide on its own, which generally means RAID 0 for single drive setups, or RAID 1 for dual drive setups). For those who don't know:

RAID 0 uses "striping". To explain what striping is, simply imagine you have a set of values you want to write (A1-A8), where the system writes A1 and A2 at the same time, but to the separate disks. This results in a high level of performance, but offers absolutely no failure recovery. If you're using

the NAS for temporary storage or a high throughput system, you'll probably want RAID 0. With RAID 0 you'll also have access to, in our case, 6TB of storage.

RAID 1 uses mirroring without parity or striping, meaning that Disk 1 is an exact copy of Disk 2 – this means that if one drive fails, you can simply replace it and keep going. The downside is that you have access to only one drive's worth of storage (3TB in this case). Generally you'll want to use RAID 1 or higher for backup systems. However, note that RAID 0 and 1 are the only options for two-drive systems.

The system discussed here uses RAID 1, due to the fact that it's primarily used for backups.

Diskstation Manager (DSM) also allows you to create multiple volumes by selecting specific drives, though it seems largely useless in a two-drive setup. Once your volume has been created, you're left with the biggest question of all: how do you want to use the space?



The Usage

Due to the fact that the test network included Mac, Windows and Linux computers, both the Mac File Service and the Windows File Service were enabled. Linux, however, can treat the Windows File Service as a simple Samba share (accessed by visiting `smb://<IP>/<share>/` in Nautilus). DSM also offers an NFS (Network File System) service, although the functions and features weren't tested.

In order to allow each person of my family's network access to a personal (and private) share, I had to create a separate user for each, disallowing access rights for all other users (save the admin account, in the case that IT help was required). Taking it one step further, a TimeMachine account was also created, due to the fact that we had multiple MacBooks which could all benefit from TimeMachine. To avoid TimeMachine filling up the entire 3TB available, a data cap was set in place (it's suggested to use 2-3 times the size of the actual Mac drive, which was impossible for 4 MacBooks – the total would be

about 4TB!). As such, it was set at 1TB. For anyone using something similar (snapshot backups), it would be recommended to exclude folders that won't change much (Videos/Music, for example), and instead manage a single backup/copy of these folders, to avoid unnecessary duplication of efforts.

Overall, the DSM interface is well designed and extremely easy to use. However, the clock and date do not, apparently, automatically synchronize with the network. As such, you'll need to change that setting before expecting scheduled tasks to run at the proper time.

The Noise

For anyone considering a NAS, it's important to consider location. If you're going to be keeping it in a bedroom, you'll probably want to invest in server-specific hard drives (they're quieter, although performance may be impacted), or set up a scheduled downtime, allowing you to sleep without the fans and disks whirring away. Also, the disk standby can result in loud spin-ups/spin-downs (at least with

the Seagate drives, although according to reviews these disks are simply a bit louder). Thus, you may want to balance power-saving with peace of mind. Another alternative would be to simply soundproof the device – keep it in a padded enclosure, suspend it somehow, or set up some sort of sound barrier between the device and the rest of the room. These are minor details, and also largely dependant on the specific drives and usage of the device. The score result disregards these minor annoyances, as the NAS enclosure itself isn't actually causing them. Indeed, the enclosure itself sports a suspended drive bay, and rubberized feet, which should help to minimize noise as well.

The Score

Overall, the NAS succeeds in offering an easy-to-use experience, while offering a great many features (both advanced and basic). It could easily be used as a personal server, as it offers web hosting as well as FTP options. Through 3rd party software, you can also easily install PHPMyAdmin and other common tools. The only downside to this is the cost,

especially due to the lack of drives. This is mostly offset by the wide range of features it offers, in comparison to similar NAS setups. Due to the form factor and design, along with the features and performance, this device gets a 4.5/5. The half a mark is simply because, while an excellent device, it could offer certain features to make life easier, such as hot swapping drive bays.



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.



Many, many months ago, we worked with API calls for Weather Underground.

Actually, it was in part 11 which was back in issue #37. Well, we are going to deal with APIs again, this time for a website named TVRage (<http://tvrage.com>). If you aren't familiar with this site, it deals with television shows. So far, every TV show that I could think of has been in their system. In this series of articles, we are going to revisit XML, APIs, and ElementTree to create a wrapper library that will allow us to create a small library which simplifies our retrieval of TV information on our favorite shows.

Now, I mentioned a wrapper library. What's that? In simple terms, when you create or use a wrapper library, you are using a set of code that "wraps" the complexity of the website's API into an easy-to-use library. Before we get started, I need to make a few things clear. First, this is a free service. However, they do request donations for use of their API. If you feel that this is a worthwhile

service, please consider donating \$10 US or more. Second, you should register at their website and get your own API key. It's free, so there's really no reason not to, especially if you are going to use the information provided here. In addition, you have access to a few other fields of information like series and episode summaries that are not included in the unregistered version. Third, they are hard at work at updating the API. This means that when you get to seeing this article, their API might have changed. We'll be using the public feeds, which are free for everyone to use as of December 2012. The API website is located at <http://services.tvrage.com/info.php?page=main> and shows a few examples of the types of information that are available.

Now, let's begin looking at the API and how we can use it.

Using their API, we can get very specific information about the show itself and/or we can get episode level information. There are basically three steps to finding

information about TV Shows. Here are the steps:

- Search their database looking for the show name to get the specific Show ID which must be used to get more data. Think of the showid value as a key directly into a record set in a database, which in this case it is.
- Once you have the Show ID, obtain the show level information.
- Finally, gather the information about a specific episode. This comes from a list of each and every episode that the show has had to date.

There are three basic web calls we will make to get this information. First is the search call, second the show information call, and finally the the episode list call.

Here are the base calls that we

will use...

- Search for ShowID based on a show name - <http://services.tvrage.com/feeds/search.php?show={SomeShow}>
- Pull the show level data based on the Show ID (sid) - <http://services.tvrage.com/feeds/showinfo.php?sid={SomeShowID}>
- Pull the episode list for Show ID (sid) - http://services.tvrage.com/feeds/episode_list.php?sid={SomeShowID}

What gets returned is a stream of data in XML format. Let's take a moment to review what XML looks like. The first line should always be similar to the one shown below to be considered a proper XML data stream (below).

```
<?xml version="1.0" encoding="UTF-8" ?>
<ROOT TAG>
  <PARENT TAG>
    <CHILD TAG 1>DATA</CLOSING CHILD TAG 1>
    <CHILD TAG 2>DATA</CLOSING CHILD TAG 2>
    <CHILD TAG 3>DATA</CLOSING CHILD TAG 3>
  </CLOSING PARENT TAG>
</CLOSING ROOT TAG>
```

Every piece of data is enclosed within a defining tag and end-tag. Sometimes you will have a child tag that is a parent tag in itself like this...

```
<CHILD PARENT TAG>
```

```
<CHILD TAG 1>DATA</CLOSING CHILD TAG 1>
```

```
</CLOSING CHILD PARENT TAG>
```

You also may see a tag that has an attribute associated with it:

```
<TAG INFORMATION = VALUE>
```

```
<CHILD TAG>DATA</CLOSING CHILD TAG>
```

```
</CLOSING TAG>
```

Sometimes, you might see a tag with no data associated with it. It would come across like this...

```
<prodnum/>
```

Sometimes, if there is no information for a specific tag, the tag itself just won't be there. Your program will have to deal with these possibilities.

So, when we go through and deal with the XML data, we start with the root tag, and parse each

tag – looking for the data we care about. In some instances we want everything; in others, we care about only certain pieces of the information.

Now, let's look at the first call and see what gets returned. Assume the show we are looking for is Buffy the Vampire Slayer. Our search call would look like this:

```
http://services.tvrage.com/feeds/search.php?show=buffy
```

The returned XML file would look like this:
<http://pastebin.com/Eh6ZtJ9N>.

Note that I put the indents in myself to make it easier for you to read. Now let's break down the XML file to see what we actually have.

<Results> - This is the ROOT of the XML data. The last line of the stream we get back should be the closing tag **</Results>**. Basically, this marks the beginning and end of the XML stream. There could be zero results or fifty results.

<show> This is the parent node that says "What follows (until the end show tag) is the information about a single tv show". Again, it's



ended by its end tag **</show>**. Anything within these two tags should be considered one show's worth of information.

<showid>2930</show> This is the showid tag. This holds the sid that we have to use to get the show information, in this case 2930.

<name>Buffy the Vampire Slayer</name> This is the name of the show

<link>...</link> This would be the link to the show itself (or, in the case of an episode, the episode information) on the TVRage website.

<country>...</country> The country of origin for the show.

...
</show>
</Results>

In the case of our program, we would be really interested in only the two fields **<showid>** and **<name>**. We might also consider paying attention to the **<started>** field as well. This is because we

rarely get back just one set of data, especially if we didn't give the absolutely complete show name. For example, if we were interested in the show "The Big Bang Theory," and searched using only the string "Big Bang", we would get twenty or so data sets back because anything that even remotely matched "big" or "bang" would be returned. However, if we were interested in the show "NCIS," and we searched for that, we would get back many responses. Some not what we would expect right away. Not only would we get "NCIS", "NCIS: Los Angeles", "The Real NCIS", but also "The Streets of San Francisco" and "Da Vinci's Inquest", and many more, since the letters "N" "C" "I" and "S" are in all of those, pretty much in that order.

Once we know the show id that we want, then we can request the show information for that ID. The data is similar to the data we just

got back in the search response, but more detailed. Again, using Buffy as our example request, here (next page, right) is an abbreviated version of the XML file.

You can see that much of the data is included in the original search response stream. However, things like network, network country, runtime, air day and time, are specific to this response set.

Next, we would request the episode list. If the show is only one season long and has/had only six episodes, this stream would be short. However, let's take the case of one of my favorite TV shows, Doctor Who. Doctor Who is a British TV show that, in its original form, started in 1963 and ran for 26 seasons ('series' for our friends in the UK) until 1989. Its first season alone had 42 episodes, while other seasons/series have around 24 episodes. You can see where you might have a HUGE stream to parse through.

What we get back from the episode list request is as shown on the next page (again using Buffy as our example); I'm going to just use part of the stream so you get a good idea of what comes back.

So to recap, the information we really want/need in the search for show id by name stream would be...

```
<showid>  
<name>  
<started>
```

In the Show Information stream we would (normally) want...

```
<seasons>  
<started>  
<start date>  
<origin_country>  
<status>  
<genres>  
<runtime>  
<network>  
<airtime>  
<airday>  
<timezone>
```

and from the episode list stream...

```
<Season>  
<episode number>  
<season number>  
<production number>  
<airdate>  
<link>  
<title>
```

A word of "warning" here. Season number and Episode number data are not what you might think right away. In the case of

```
<Showinfo>  
  <showid>2930</showid>  
  <showname>Buffy the Vampire Slayer</showname>  
  <showlink>http://tvrage.com/Buffy_The_Vampire_Slayer</showlink>  
  <seasons>7</seasons>  
  <started>1997</started>  
  <startdate>Mar/10/1997</startdate>  
  <ended>May/20/2003</ended>  
  <origin_country>US</origin_country>  
  <status>Canceled/Ended</status>  
  <classification>Scripted</classification>  
  <genres>  
    <genre>Action</genre>  
    <genre>Adventure</genre>  
    <genre>Comedy</genre>  
    <genre>Drama</genre>  
    <genre>Mystery</genre>  
    <genre>Sci-Fi</genre>  
  </genres>  
  <runtime>60</runtime>  
  <network country="US">UPN</network>  
  <airtime>20:00</airtime>  
  <airday>Tuesday</airday>  
  <timezone>GMT-5 -DST</timezone>  
  <akas>  
    <aka country="SE">Buffy & vampyrerna</aka>  
    <aka country="DE">Buffy - Im Bann der Dämonen</aka>  
    <aka country="NO">Buffy - Vampyrenes skrekk</aka>  
    <aka country="HU">Buffy a vámpírok réme</aka>  
    <aka country="FR">Buffy Contre les Vampires</aka>  
    <aka country="IT">Buffy l'Ammazza Vampiri</aka>  
    <aka country="PL">Buffy postrach wampirów</aka>  
    <aka country="BR">Buffy, a Caça-Vampiros</aka>  
    <aka country="PT">Buffy, a Caçadora de Vampiros</aka>  
    <aka country="ES">Buffy, Cazavampiros</aka>  
    <aka country="HR">Buffy, ubojica vampira</aka>  
    <aka country="FI">Buffy, vampyyrintappaja</aka>  
    <aka country="EE">Vampiiritapja Buffy</aka>  
    <aka country="IS">Vampírubaninn Buffy</aka>  
  </akas>  
</Showinfo>
```

the data from TVRage, the season number is the number of the episode within the season. The episode number is the number for that episode within the total life span of the series. The production number is a number that was used internally to the series, that, for many people, means little if anything.

Now that we have refreshed our memory on XML file structures and examined the TVRage API calls, we are ready to start our coding, but that will have to wait until next time.

Until then, have a good holiday season.



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

```
<Show>
  <name>Buffy the Vampire Slayer</name>
  <totalseasons>7</totalseasons>
  <Episodelist>
    <Season no="1">
      <episode>
        <epnum>1</epnum>
        <seasonnum>01</seasonnum>
        <prodnum>4V01</prodnum>
        <airdate>1997-03-10</airdate>
        <link>http://www.tvrage.com/Buffy_The_Vampire_Slayer/episodes/28077</link>
        <title>Welcome to the Hellmouth (1)</title>
      </episode>
      <episode>
        <epnum>2</epnum>
        <seasonnum>02</seasonnum>
        <prodnum>4V02</prodnum>
        <airdate>1997-03-10</airdate>
        <link>http://www.tvrage.com/Buffy_The_Vampire_Slayer/episodes/28078</link>
        <title>The Harvest (2)</title>
      </episode>
      <episode>
        <epnum>3</epnum>
        <seasonnum>03</seasonnum>
        <prodnum>4V03</prodnum>
        <airdate>1997-03-17</airdate>
        <link>http://www.tvrage.com/Buffy_The_Vampire_Slayer/episodes/28079</link>
        <title>Witch</title>
      </episode>
      ...
    </Season>
  </Episodelist>
</Show>
```



So far in our series on LibreOffice Base, we have planned the layout for our database, created the tables, and created the relationships between those tables. Now, we need to think about how we will input the data into our tables and link them all together. You can edit the tables manually in the tables section of the program, but the best way is through input forms. We will create three forms: Authors, Media, and Books. We will handle the forms for authors and media first as they are the easiest. The form for books will bring everything together in one place.

Create the Authors and Media Forms

Select forms from the Database pane on the left. You can create forms in two ways, through the wizard or through the Design View. We will use the wizard for Authors and Media. Click on Create from wizard... in the tasks pane. This starts the wizard.

On the first step of the wizard, select the Authors table from the dropdown, and move Name to the Fields to use. This is done by selecting the field and using the arrow keys. Click Next.

The second step deals with subforms, which we do not need for the Authors form. Click Next. You will notice that steps three and four are skipped. Those steps deal with setting up a subform. We will not use them in the our implementation of our database.

Step five asks us for a layout for our fields. Since we have only one field, we will choose the table layout. Select the table layout and click Next.

The sixth step is about how the form is used. We want the form to display all data, so select all data and click Next.

Step seven deals with the styles of the form. You can select different background layouts for your form. You can also select whether your controls are displayed with no border, 3D

borders, or flat. Click Next.

On the eighth step, we give the form a name, Authors, and choose whether we want to edit it more or use it. After naming the form, click Finish. The new form will pop up on your screen. Close it and save your work.

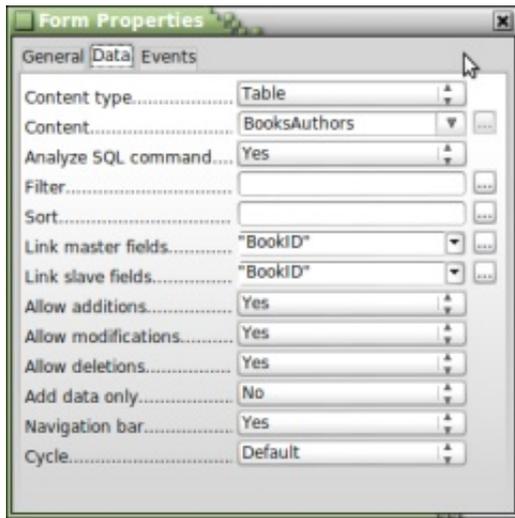
The Media form is created in the same way, only using the Media table instead of the Author table.

Create the Books Form

Although we will use the wizard to start our Books form, we will need to edit it afterwards to add the connections to authors and media. Go through the wizard again with the Books table, adding the fields Title and Published. You will skip the subforms again. We will add our subforms manually. For the layout, use either Columnar – Labels on Top, or Columnar – Labels Left. I used Columnar – Labels on top. After naming the form Books in step eight, select Modify the form, and click Finish.

This time the form is opened for editing. We will add two subforms to make a connection to the Authors and Media tables.

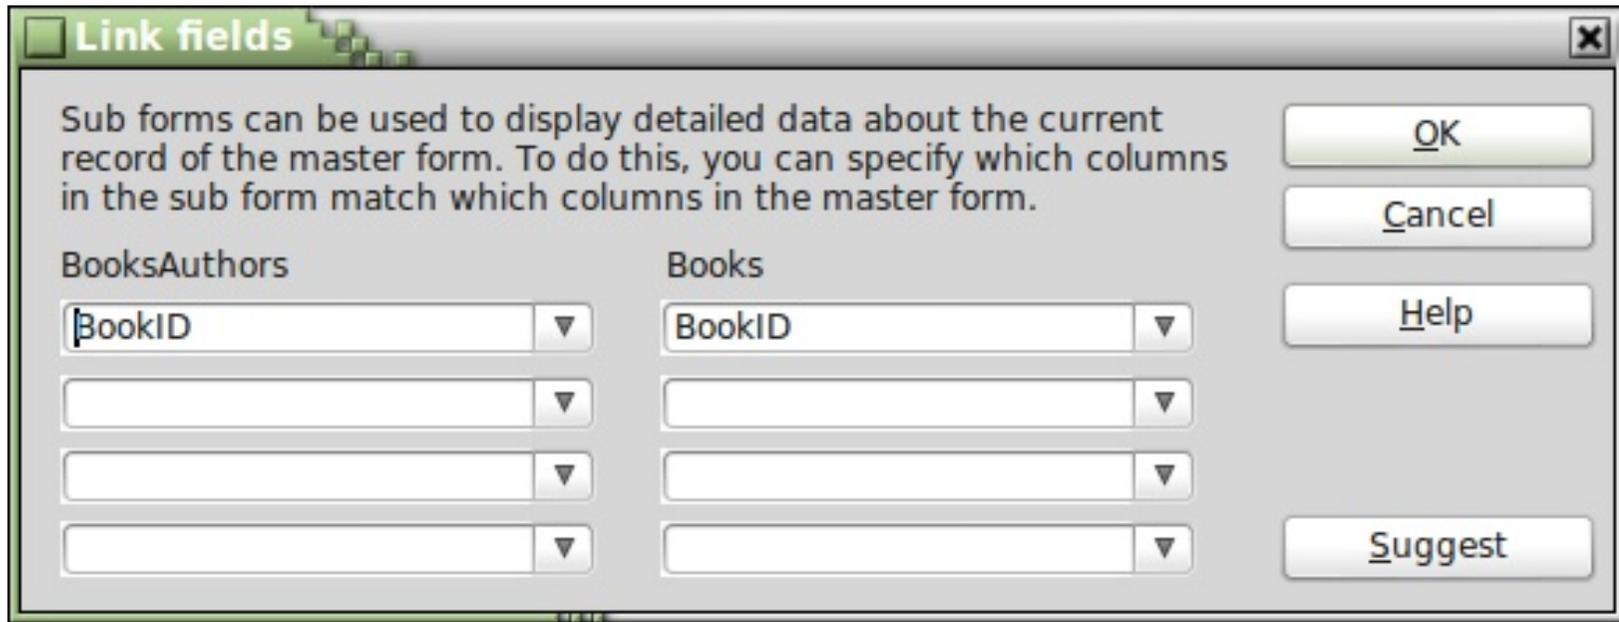
On the toolbar at the bottom of the form, select the Form Navigator. This opens a small dialog containing the elements of our form. Right-click MainForm, then New > Form. A new form is added to the list. Since it is already selected, just start typing to rename it FormAuthors. We will use this subform to create the link to our Authors table through the BooksAuthors table. Right-click the FormAuthors and select Properties. On the Data tab, select table for the Content type, and BooksAuthors for the Content. Now, click on the ellipse button next to List Master Field. A link-fields dialog pops up. This is where we create our link between the Books table and the Authors table. Under BooksAuthors, select BookID and under Books select BookID. Click OK, and you can close the Form Properties dialog.



Create another subform under MainForm named FormMedia. Edit the properties, setting the Content Type to table, and the Content to BooksMedia. Click the ellipse next to List Master Field, and select BookID under both BooksMedia and Books. Click OK and close the Form Properties dialog.

Now, we need to create our controls that will contain our authors and media. Remember, we need the ability to select more than one, so a table is our best choice for the control. Click on the More Controls button on the toolbar on the left side. A More Controls toolbar will pop up.

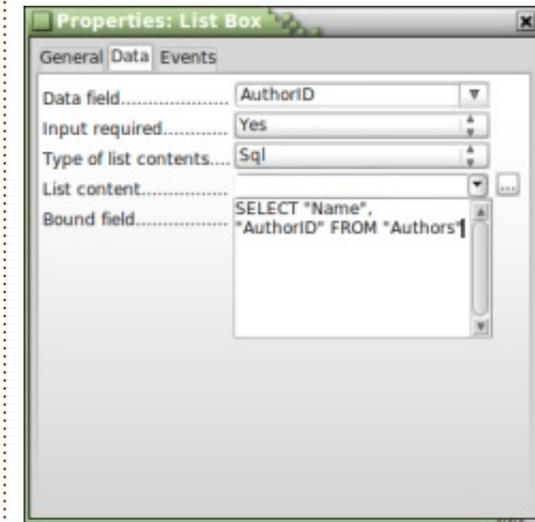
Make sure you have the



FormAuthors selected in the Form Navigator. Click on the Table button in the More Controls toolbar. Underneath the two input boxes for title and published, draw the table on the page. Right-click in the header section of the new table and select Insert Column > List Box. Right-click the newly created column and select Column. This displays the List Box properties. On the General tab, change the label to Authors. Switch to the Data tab. For the Data field, select AuhtorID. This tells Base we want to relate this to the AuthorID field in the BooksAuthors table. For Type of list content, select Sql. In List

Content, we will write a short SQL statement to fill our list box with the names from the Authors table. You can click on the down-arrow to give yourself a small edit box to work in. The SQL statement is:
`SELECT "Name", "AuthorID"
 FROM "Authors"`

This statement will select all the records from the Authors table. Set the Bound field to 1. The Bound field selects which field will fill the list box. Since we selected the Name field first, the bound field of 1 will fill the list box with the data from the Name fields. Close the list box properties and save your work.



We will create the Media control in the same way we did the Authors. Select the FormMedia in the Form Navigator. Draw the

table control to the right of your title and published controls. Create a column in the control. Open the column properties and change the label to Media. Set the Data field to MediaID, and the Type of list content to Sql. The SQL statement for the List Content is:

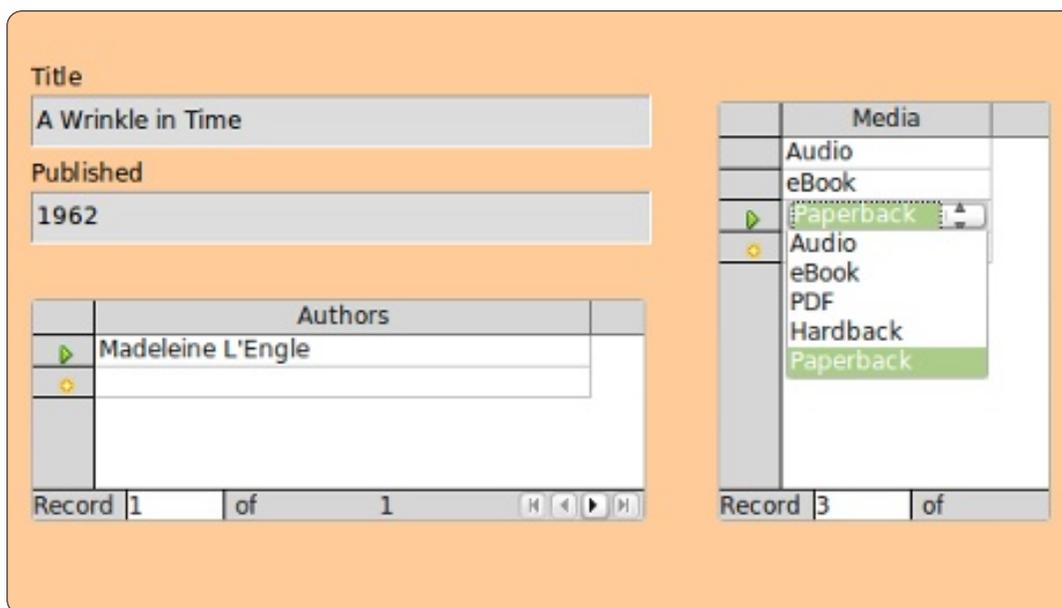
```
SELECT "Type", "MediaID"  
FROM "Media"
```

Set the Bound field to 1.

We are now finished with the form. Save and close it.

Using the Forms

The Authors and Media forms are simple to use. Just select an empty row and type in the name or type. You will want to add your authors and types before using the books form. The books form is easy to use as well. Fill in your Title and Published year (or Unknown if you don't know the year). The authors and media tables give you a list box from which you can select your authors and media types. Note that you can select more than one author and media type. To create a new record, click on the new record or next buttons



in the form control toolbar.

We now have a working method for entering data into our database. While it works, there is one disadvantage to this setup. If you find that you need an author or media type that is not there, you have to close the books form and open the authors or media forms. In a later How-To, we will attempt to overcome this inconvenience.

In the next HowTo, we will create a query and a report for extracting information from our database.



Elmer Perry's history of working, and programming, computers involves an Apple IIE, adding some Amiga, a generous helping of DOS and Windows, a dash of Unix, and blend well with Linux and Ubuntu.



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



HOW-TO

Written by Alan Ward

As some may remember, the Asus eeePC model 701 was one of the forerunners of the netbook movement. However, it was not known for speed even when it came out in 2007: a Celeron 900 MHz processor (underclocked down to 600 MHz), 512 MB or 1 GB RAM, and 4 GB internal flash-based hard drive were by no means impressive specifications. The 7" screen was tiny for a laptop by any standard, and even more so nowadays when many tablets offer more screen real estate. Battery life, at about 2 hours when new, could also be an issue.

Even so, it seems a shame to throw away such a nice piece of kit, especially when it is physically tougher than many other netbooks thanks to a hard plastic outer shell and a hard drive with no moving parts. But, can it use a recent version of our favorite operating system at credible speed? The short answer is a definite "yes!" - though installing it does need some (light) tweaking. For the longer answer, read on.

Since the eeePC has no optical drive, we need to use either a USB pendrive or an external CD unit to install Ubuntu, the pendrive being the better option for speed. The Escape key gives us access to the boot media menu, where most external USB units are readily recognized by the computer's BIOS. The Ubuntu installer boots

up nicely, either in live CD mode or just the installer itself. The live CD mode is actually best for our use, since we will need access to a terminal at some point.

First hurdle: the 4 GB hard disk is small!

The first difficulty we will

encounter is the installer complaining about the hard drive: it thinks 4 GB is just too small, and will not continue. Naturally, this is not at all true since a standard install of the Ubuntu 12.10 32-bit version takes up just about 2.5 GB. So it does fit, though not much space will be left for user documents.

This quirk of the installer can easily be circumvented: just plug in an external hard disk during installation, with enough free space on it to satisfy the installer. It does not seem to notice that Ubuntu is being installed on the internal 4 GB unit, while the free space is actually on another drive. In fact, if our pendrive is large enough, we shall not even need the external hard drive to convince the installer that enough free space is available.

Please do take note of drive identifications, though, if you try this. The internal 4 GB hard drive should be `/dev/sda`, while the pendrive or CD you are booting from is `/dev/sdb`, and a further



HOWTO - INSTALL 12.10 ONTO AN EEE PC

external hard drive (if present) would be `/dev/sdc`. Ensure yourself you are installing to `/dev/sda`, and the bootloader is also placed on that same drive.

The external drive can be removed along with the pendrive when installation is complete.

Setting up some swap

While I am lucky enough to own the 1 GB RAM version, other owners may be stuck with 512MB. This is not enough to run the system installer, so a bit of swap space comes in handy. We will set it up right away when the Live CD environment starts up, before firing up the installer.

First of all, we can use either a terminal or the `gnome-disks` graphical command to partition `/dev/sda`. Reserving about 3.5 GB for `/dev/sda1` (partition type 0x83, Linux) and the remaining 320 MB for `/dev/sda2` (partition type 0x82, Linux swap) seems a reasonable choice. And yes, 4 GB is a merely theoretical capacity that is considerably reduced in practice.

Once `/dev/sda` is partitioned,

we can start up swap inside a terminal using the following commands:

```
sudo mkswap /dev/sda2
```

```
sudo swapon /dev/sda2
```

Now, remaining in terminal-mode, we can issue the command:

```
free -m
```

which shows us both the amount of physical RAM (the line starting with "Mem:"), and the amount of virtual memory (starting with "Swap:"). If swap is correctly activated, this last line should indicate about 314 MB is available.

We can now proceed to start up the installer.

Small screen needs more real estate

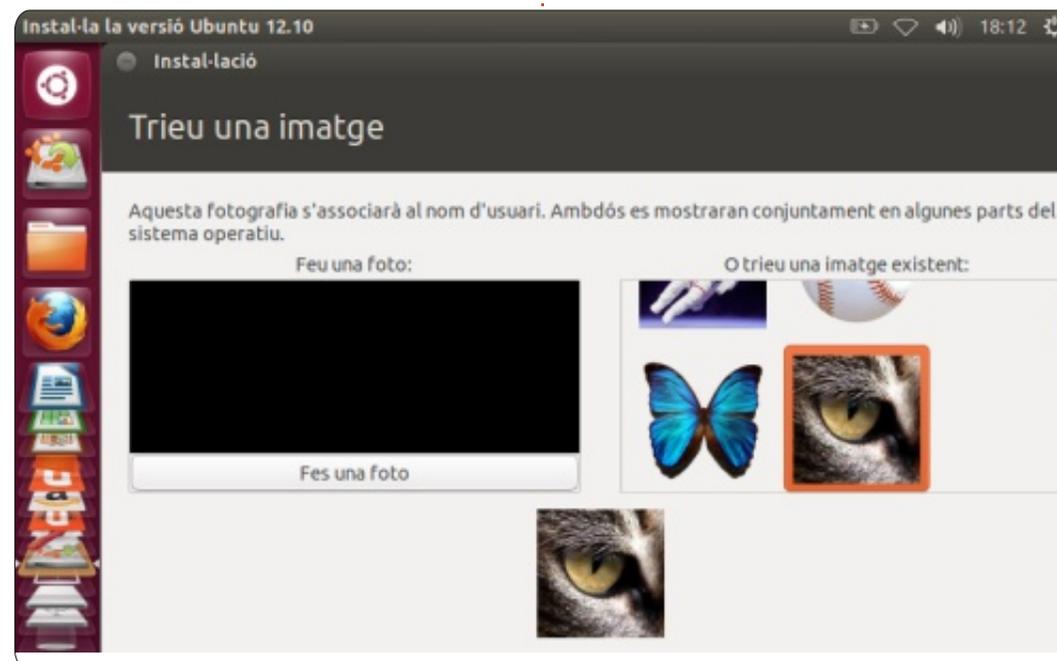
Once we start the installation process, we will notice straight away that the buttons at the bottom of the installer window are not in view. The eeePC's screen is simply too small, at 800x480 pixels. In most steps, this is not a problem, simply hitting Enter is enough to continue to the next stage.

However, for some screens, we may need to actually see those elusive buttons. So, when we need to move the installer window to a different position on screen, we can hold down the Alt key while we use the mouse to drag the window, holding it by any point - not just the window bar at the top.

Another possibility would be to plug in an external monitor just for installation, though things can get a bit awkward if you are not used to handling multihead setups. 800x480 pixels is what we will have during normal use, so we might as well get used to it, and moving windows around.

What about usability?

Once the installer finishes up, we reboot (removing the pendrive and eventual external hard drive as we go), and start up afresh. Boot-up times could have been considerably reduced if the constructor had used a different internal hard drive and connection. The particular flash controller chip included in the eeePC (a Siliconmotion SM223AC) is connected to the motherboard's parallel ATA interface, and



HOWTO - INSTALL 12.10 ONTO AN EEE PC

configured with a rather slow UDMA/66 connection. According to the manufacturer, this flash-based hard drive sustains 50 to 55 MB/s read speed - rather less than the 100 MB/s (laptop) or 140 MB/s (desktop) typical of modern platter hard disks, or the more than 300 MB/s read speed sustained by most recent solid state drives. So, all things being considered, the boot-up procedure is not too lengthy under Ubuntu 12.10 with a 65s wait until a usable desktop is ready.

Also on the bright side, graphics performance proves more than satisfactory. The integrated Intel 915 controller is very well supported under Linux, and the screen itself has a nice, crisp image. The matte screen finish does no harm, either. In any case, Unity's visual effects (dash lenses, changing workspaces ...) all work quite well, though, as could be expected, a tad more slowly than on a desktop with a dedicated graphics card.

As before, real estate is lacking on the 7" screen and some dash elements may not be visible. Luckily, the more important bits

are at the top left and readily accessible, so kudos to the Unity team for this one.

Now, onwards to the less rosy side of things. The weak CPU can be a bit of a bother. Some latency between clicking on an icon and it actually doing something is the norm, especially when the computer is rather loaded with various applications all running at the same time. In such situations, you should also be prepared to hear a continuous whine coming from the (small) internal fan. This being said, most light browsing or presentation work can be done with zero problems - and stability is

as good as could be expected from our favorite distribution.

Wireless connectivity is not the eeePC's most outstanding quality, to say the least. Short antennae in a convoluted position within the case probably don't help in acquiring a strong signal, so reception may at times be a matter of holding up the unit at arms length and assuming interesting positions while hunting for the best reception - which is invariably slow and choppy. However, if you are in the immediate vicinity of an access point, you can hope to browse the web with some fluidity. The F11 key may be useful to

expand the browser window to occupy the whole screen.

If wireless is a must for you, you may be more comfortable using an external (USB) wireless dongle. In that case, it is best to avoid loading the internal Atheros card's drivers at boot, to make sure Network Manager does not get confused about which card to use. To do so, acting as root or using the sudo command, add a file called `/etc/modprobe.d/blacklist-custom.conf` to your configuration, with the following content:

```
blacklist ath
```

```
blacklist ath5k
```

You will then need to reboot. This file can always be removed if or when you should wish to re-enable internal WiFi.

All in all, using this netbook as a production machine is not a complete fail. All hardware is recognized and correctly configured straight from the gate, no need to fiddle around with drivers. The screen is easier on the eyes than could be expected from its size, and the machine is





generally responsive enough to actually do some work on it. Even typing at some length - the kind of use I tend to avoid on a touchscreen tablet of the same size - is possible, though the keys are a shade small for comfort.

What is impressive is the fact that a five-year-old netbook, designed with all the compromises of speed and other characteristics that were both acceptable and necessary at that time, is still very much capable of running Ubuntu's latest and greatest. I think this speaks volumes both about the machine itself, and about the operating system. And, if all this

means one less computer ends up in a landfill before its time is truly over, then so much the better.



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

16x16 SUDOKU

Numbers 0 to 9 and letters A to F are to be filled into the 16x16 grid so that every row, every column, and every 4x4 box contains 0 to 9 and A - F.

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

Solutions are on the second last page.

Puzzles are copyright, and kindly provided by, **The Puzzle Club** - www.thepuzzleclub.com



HOW-TO

Written by Nicholas Kopakakis



Blender - Part 1

It's kind of odd to start an article with a Youtube link, but I believe that if you like what you'll see in there, you will love the rest of these tutorials.

Feel free to follow the link below and come back here in about 15 minutes, multiplied by the number of times you'll enjoy it: <http://www.youtube.com/watch?v=eRsGyueVLvQ>

Sintel (www.sintel.org) is not the reason I mess with Blender, but I believe that if I hadn't till then, after that I surely would have "started my journey".

So, lets talk about Blender (www.blender.org), the program behind Sintel and many similar movies. (Actually Blender is behind a lot of 3D things). Blender is available for Linux, Windows, Mac OS X and FreeBSD (32-bit and 64-bit systems).

When this article was written, the latest version was 2.64 but we'll work with version 2.62, which you can find in the Ubuntu

repositories. If you haven't already installed it, go ahead and install it from the repository.

In these tutorials we will use a 3-button mouse and a keyboard with number pad. Blender assumes that you have that. But, if you don't, you can change these setting under File > User preferences. Select the input tab and, on the left, mark Emulate 3-Button Mouse and Emulate Numpad.

You will encounter a splash screen showing the version of Blender at the top right corner and some links. Ignore them for now and press any mouse button outside that area.

If you are new to 3D animation software, don't worry. Give it some time and you'll get used to it: practice, experiment, and practice are the only three ways to excel in Blender. If you are familiar with other similar programs (Maya, 3DS Max, etc), you will find some

differences but the main philosophy is here.

So, we are in Blender and we see the 3D view editor, some buttons at the left, some buttons and numbers at the right, and a timeline at the bottom. Working in Blender mainly means "working in 3D view editor".

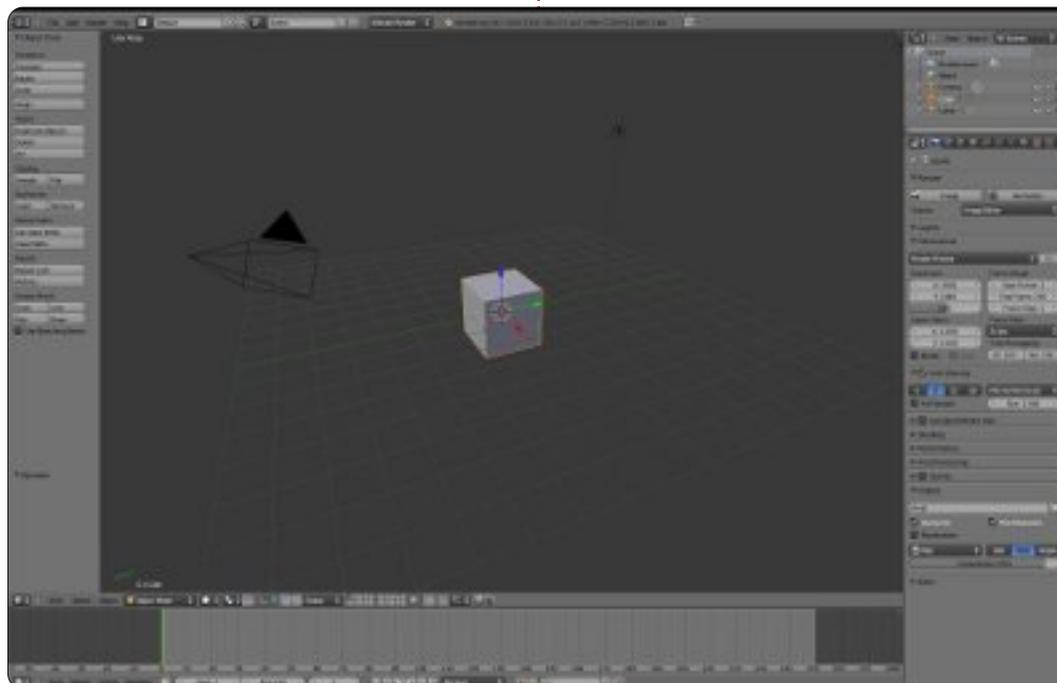
Working in Blender is modeling, animating, lighting, compositing, texturing (where we will use all editors), but the 3D view editor (or 3D view window) is still the place where everything is revealed.

The 3D view editor is the big window at the center of our screen. What you see there is a cube, a camera, and a light.

Press F12

What's going on? Well, you just rendered your first scene in blender!

TIP: You can inspect the previous rendered image anytime by pressing F11. Pressing F11 again



will return you at the 3D view editor.

Move the mouse on the image and press Esc or press F11 to exit the rendered image window and return to the 3D view editor.

All that you need to have a scene is a capture device (the camera), a theme to capture (the cube), and some light.

Now, let's focus on the bar on the bottom of the 3D view (shown below) called Header



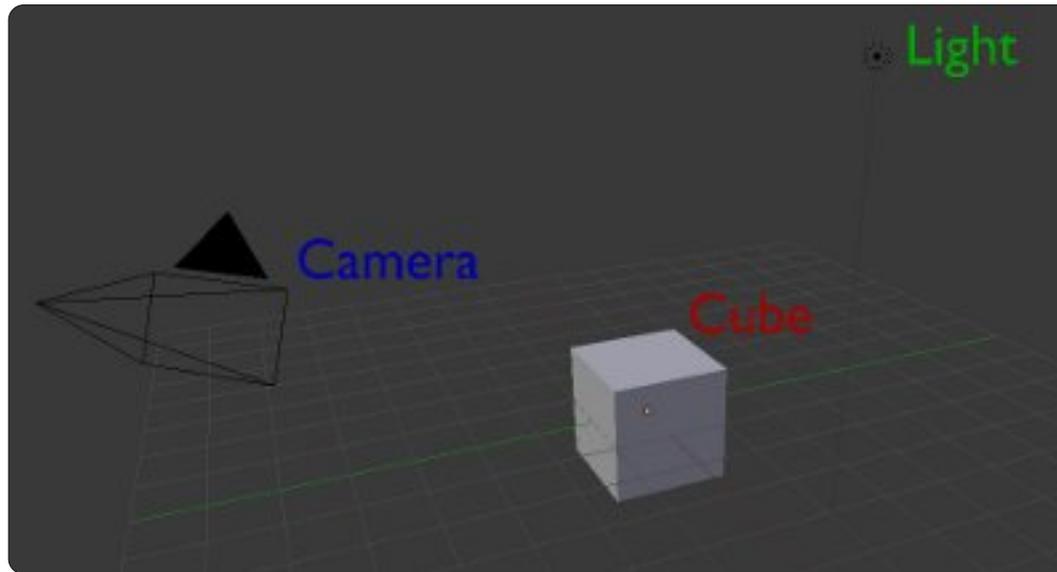
Notice the tiny little box (shown left) on the left-most side. This is the editor type selector.

Press your LMB (Left Mouse Button) on the little 3D box to change the editor type.

From the list select Python Console (the very first selection at the top).

Hey! Where is my box?

Don't worry, your cube is right



there where you left it. From the same menu select the 3D View. Voilà! Your box is there. See?

In Blender, you can change any window to any editor type window you want. This will be very handy in the future when we will want to be more productive. Go ahead and locate the four other window headers on your screen. (There are five in total.)

Notice that every window has a different header with different options, and every time you change editors, the header changes as well.

Now, change every editor to info (the header at the top of your screen is already in info so you don't have to change it.)

At any time, press File > New in info header to reload your default setup.

Ok. Let's start a new project; press File > New in info header.

Everything is back in place.

Move the mouse in the 3D view editor window, and press the MMB (Middle Mouse Button) while

moving the mouse around. You are rotating the view. Moving the mouse wheel up and down, you zoom in and out.

Alt+Shift+MouseWheel rotates the view up and down.

Alt+Control+MouseWheel rotates the view left and right.

Also you can pan the view. Shift+MMB and you can pan the view, with Shift+MouseWheel up and down and Control+MouseWheel for left and right.

You can zoom in and out without using the middle mouse wheel. Use Control+MMB and move the mouse up (zoom in) or down (zoom out).

Enough with rotation and panning. Let's go to selection!

Start a new project: File > New.

By default, when you start a new project, the cube is selected (indicated by the highlighted-orange edges of the cube)



Now press the A key. You deselect the cube (nothing in the scene is highlighted).

Pressing A key again will select All objects on-screen. The cube, the camera and the light. Press A one more time to deselect everything.

Now press the RMB (Right Mouse Button) on the cube. You select the cube and it's highlighted.

Press the RMB on the light. You have selected the light but you deselected the cube as well.

Press A to deselect the light.

Shift Selection

Press RMB on the light and then press Shift+RMB on the cube. You selected the light AND the cube (notice the more redish edges on the light).

Now Shift+RMB on the camera (notice the redish light). This is cumulative selection, meaning that you added the camera to the selection.

Shift+RMB to the cube.

The redish edges become orange. Shift+RMB again on the cube and you deselected the cube. Now you have in your selection the light and the camera. Press A to deselect all.

Now, using the MMB, rotate the view, in order to see all 3 objects, and deselect all with the A key. (You can instead load a brand new scene by clicking File->New or Control+N, and, on the new project, press A to deselect the cube.)

Box Selection Tool

Assuming that you have loaded a new project or you have a similar view, move the mouse over and right of the light and press the B key. A cross appeared. Now hold LMB and drag the mouse down and to the left in order to Box select all 3 objects.

Great! Everything is selected with a different and faster way than the Shift+RMB selection. Press A key again to deselect all.

Circle Selection Tool

Press C and a Circle appears.

Circle selection works in a cumulative manner like shift selection, so pressing LMB on the cube will add the cube to the selection.

Moving the mouse wheel up and down, you can adjust the circle radius making selection more precise or general.

Remember, any time you press LMB you add the object to the current selection.

Press MMB to remove an object from the selection.

Press Esc to exit the selection mode.

Lasso Selection Tool

Press and hold Control+LMB and drag your mouse to put the border around an object. You can select multiple objects in a more accurate way than the box selection.

You can add objects to the selection with the Control+LMB, and you can abstract objects from

the selection with the addition of the Shift key to the combination: Control+Shift+LMB. Press and hold Control+Shift+LMB and abstract objects from the selection.

Go to File->New or press Control+N. Start a new project.

Practice, Experiment and PRACTICE! Rotate the view, pan the screen, zoom in, zoom out, make multiple selections and deselections.

Next month, we will actually create something.

Go to http://www.youtube.com/watch?feature=player_embedded&v=R6MLUcmOul8

Enjoy "Tears of Steel" (www.tearsofsteel.org) by the Blender Foundation.



Nicholas lives and works in Greece. He is working for a post-production house (commercials - films) for several years. Three months ago he migrated to Ubuntu because "it renders faster". Blender found him two years ago.



HOW-TO

Written by Mark Crutch

Inkscape - Part 8

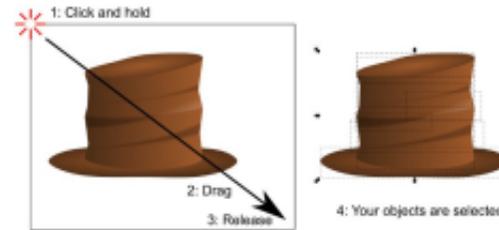
If you've been following this series from the start, you should now be quite comfortable with ellipses, rectangles and stars. You should be able to construct some quite complex paths using the Boolean operations, then manipulate them with the Node tool. You can give your objects colours, gradients and patterns, as well as apply markers, thickness and different end caps to their strokes. In short, you've got enough of a toolkit that you can create quite complex drawings, should you wish. In this instalment we will look at ways to manage that complexity.

As the number of objects in a drawing increases, it becomes less and less likely that manipulating them one at a time will be acceptable. If you want to scale or rotate our snowman's hat, for example, you will quickly become frustrated if you have to first manipulate the brim, then the side, then the top – and that's without considering the shadows and highlights we added last time. A far better approach is to select all

of the parts and perform your operations on them simultaneously.

Last time you learnt how to select two objects at once by holding SHIFT as you click on the second one. If you continue to hold SHIFT while clicking on more objects, they will each be added to the selection. Holding it while clicking on an already selected object will remove it from the selection. This approach is useful when you want to select just a few objects, or if they're widely spaced with other objects in between.

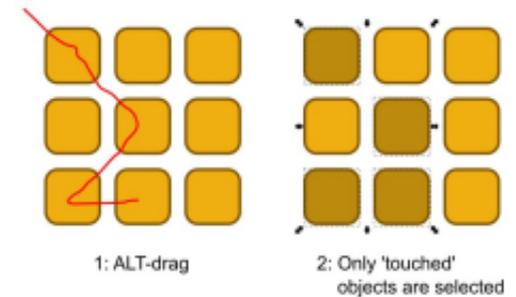
In the case of our hat, however, there's a much better way. The objects are clustered closely together which makes it easy to simply drag a selection box around them, as follows: click and hold on a blank area of the canvas, then drag the mouse diagonally away from the starting point. You'll see a rectangle – the selection box – which follows the mouse pointer. When you release the button, anything inside the rectangle will be selected.



Unfortunately, there's not always a blank area of the canvas nearby. In the case of our hat, if it's in situ on the snowman, then you'll probably have the background rectangle in the way. If you try to drag a selection box by starting on another object you'll just end up moving that object instead. The answer is to hold SHIFT as you start dragging your mouse, which will prevent Inkscape from selecting the object you started on. Remember, SHIFT-CLICK will add to your selection (or remove from it), but SHIFT-DRAG will prevent the clicked object being added.

One limitation with a selection box is that it will select only objects that are entirely enclosed within the box. This can be a problem if you're zoomed right in,

or your objects are too close together and you want to select only some of them. In these cases, you can use Inkscape's 'touch select' mode: just hold down the ALT key as you click and drag your mouse around, drawing a red line as you go. When you release the mouse button any objects that were touched by the red line will be selected. If you want to add to the existing selection, hold down SHIFT as well.



On many Linux systems you'll quickly discover a small issue if you try to use ALT-drag for the touch selection mode: often window managers use ALT-drag as a way to move the window around without having to drag the title bar, which prevents ALT-drag from working in Inkscape. There are three ways to



deal with this limitation:

- Disable the ALT-drag option in your window manager's preferences. This isn't always easy to do, and will usually have the side-effect of disabling it for all windows from all applications, not just Inkscape.
- Hold down the SUPER key (that's what it's nominally called in the Linux world, but you probably know it better as the WINDOWS key) in addition to ALT or SHIFT-ALT.
- Start dragging or SHIFT-dragging as though you are dragging a selection box, then press ALT to switch to touch select mode during the process. You can press and release ALT as often as you like during this process – it's only when the mouse button is released that Inkscape will check it to determine whether to use the selection box or touch select mode.

With several objects selected, it's time to combine them into a 'group' using the toolbox icon, the CTRL-G shortcut, or selecting the Object > Group menu item. Now you can move, scale, skew and rotate all the objects as though they were a single item. This makes it easy to combine the separate shapes that make up a single

entity, such as our snowman's hat. You can even create a group that contains only one object, but usually grouping is used to make it easier to work with several objects that form a single part of your image.

Often you will want to modify an object that is in your group. One approach is to select the group,



then use Object > Ungroup (CTRL-SHIFT-G) or the Ungroup toolbox button (shown left) in

order to break the group apart into its constituent parts. After modifying your object you can re-group them again. It's usually beneficial to edit an object without the rigmarole of ungrouping and then re-grouping afterwards, so Inkscape lets you 'enter' a group in order to work with its contents directly. The fastest way to enter a group is simply to double-click on it, though there is an option at the bottom of the right-click context menu for 'Enter Group g#20'. Inkscape assigns a unique identifier to a group, so this menu entry will be slightly different for each of them.

Once you have entered a group, you are free to edit its contents individually. The status bar will



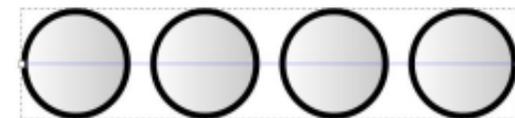
show you when you're inside a group by temporarily placing its ID into the layers pop-up, to the right of the fill, stroke and opacity settings (shown above).

Having entered a group, any objects that you paste from the clipboard will be added to the contents of that group, as will any new objects you create. There are various ways to exit a group, but the most common is simply to select an object that's outside the group. You can also double-click on a blank area of the canvas if you don't want to select something else, or select 'Go To Parent' from the context menu.

In addition to moving and transforming a group, you can also adjust its fill and stroke settings. Setting the fill or stroke will apply those values to every object in the group, as will a change to the stroke width. Unfortunately, none of the other line style attributes can be set like this, so if you want all the objects in your group to have a dashed stroke you'll need to enter the group and set each one

individually.

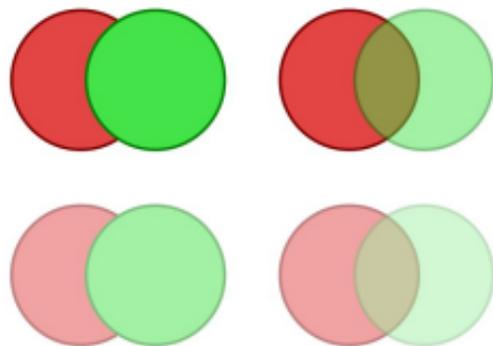
Setting a flat color for a group's fill or stroke does exactly what you might expect. But try setting a gradient or a pattern to the whole group and you might be a little surprised by the results. Suppose you want to apply a gradient from white to black across a number of objects: you might be tempted to group them and then apply the gradient to the group. This is the actual result you would get:



As you can see, although the gradient stops cover the whole width of the group, Inkscape has actually just used a small portion of it – the amount covered by the first object – and then repeated that small amount for every other object in the group. So instead of white to black across four objects, we get white to grey across one object, which is then used for the

other three as well. The same problem applies for patterns: each object has a copy of the first object's pattern applied, rather than there being one pattern that covers the whole group. There is an answer to both these problems in the form of clipping paths, but that's a more complex subject for another day.

The bottom section of the Fill and Stroke dialog does have an effect at the group level. If you set the opacity or blur for a group, it applies to the whole group as a single object. This is in addition to any opacity or blur that has been applied to individual objects, which lets you create complex combinations of effects. In this image, the top row shows two groups, one with no opacity on its objects, and one with the green circle set to 50% opacity. The bottom row shows what happens



when you then also apply an opacity of 50% to the group itself.

Rather than the opacity being applied to each individual object – as a fill color would be – the circles retain their individual opacities and then the group opacity is applied to the whole. The same rules apply for blurs: the individual objects are blurred first, then the group-level blur is applied to the whole collection.

It may seem counter-intuitive to have fill and stroke affect groups in one way, while blur and opacity affect them in a different way. In practice, it's something you get used to very quickly, and the artistic benefits of having multiple levels of opacity and blur easily outweigh any short-term confusion.

Groups are invaluable for gathering related objects together into a single easily-managed entity – such as the case of the snowman's hat at the start of this article. You could also create another group for the snowman's head, and a third for his body and arms. Moving him around on your drawing is now a lot simpler as there are only three objects to

select and move, rather than the dozens you had to contend with previously. But we can make things simpler still by creating a group that is made up of our three existing groups. Just select all three and create a group in the same way that we did earlier in the article.

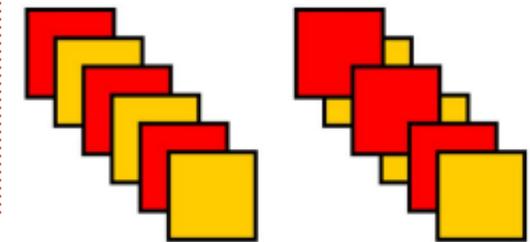
With a single group containing all the parts of our snowman, it's trivial to move him around. If you need to adjust the position or angle of his hat, just enter the group and you can interact with the three groups inside it. Select one of those and you can enter it again to get down to the individual objects. Inkscape lets you nest your groups as deeply as you like.

To exit a deeply nested group, you can use the same techniques as for a single group. To step back out of your groups one level at a time, you can double-click on a blank area of the canvas, or use the context menu's Go To Parent option. To jump directly to a particular level, you only need to click on another object or group that's at that level. Alternatively, you can use the layers pop-up on the status bar to jump straight to any ancestor group, or even right

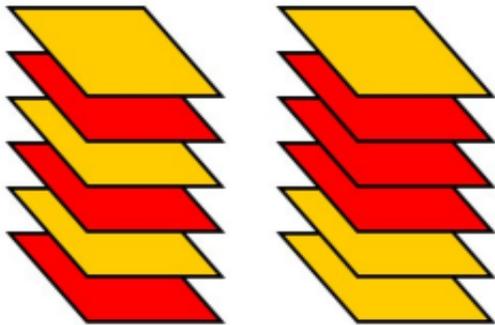
back up to the page level.

Nested groups follow the same rules as any other groups when it comes to fill and stroke colors, opacity and blur. Set a fill color on a nested group and all the objects, no matter how deeply nested, will be set to that color, but opacity and blur is applied to each object and group separately.

Although groups are an invaluable tool for managing complex drawings, they're not without their limitations. In particular, a group occupies a single 'slot' in the z-order stack – so you can't interleave the objects from one group with those from another. In this image, I've drawn a series of interleaved squares on the left. The right-hand image shows the result of combining just the red squares into a group: the entire group takes up just a single slot in the z-order, occupying the same slot as the topmost object in the group.



If you imagine these objects as a stack of paper sheets, it's a bit more obvious what's happened. The following image shows the interleaved sheets stacked on top of each other, followed by the rearrangement of the z-order once the red sheets are grouped together.



If you enter a group, you can rearrange the z-order of the objects, but only relative to one another. The group as a whole will always just take up a single slot. In this case, it means that the three red sheets are indivisible – you can't move one of the orange sheets between them without either ungrouping the red sheets, or moving the orange sheet inside the group as well. This limitation will prevent you from creating some groups that might otherwise

be useful – a rope that snakes around both the front and back of another object couldn't be grouped as a single object – but in many other cases groups will simplify your use of Inkscape significantly.



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/>

CODEWORD

Every number in the grid is 'code' for a letter of the alphabet. Thus the number '2' may correspond to the letter 'L', for instance. All - except the difficult codeword puzzles - come with a few letters to start you off

11	12	10	10	8	10	7	6		7	20	6	14													
	10		9		17		4		10		17														
26	17	8	14	10	7		22	14	6	10	8	26													
	3		24		8	10	14		10		3														
2	10	17	8		24		12	22	15	13	10	21													
	15				5				14																
14	6	10	14	6	10		4	22	6	14	10	21													
			10				10				15														
22	4	25	17	19	14		21		1	10	14	6													
	20		19		24	21	21		10		11														
17	19	23	20	22	14		8	22	4	25	22	21													
	18		25		10		10		4		15														
20	6	10	6		4	17	7	14	16	7	10	21													
1	2	3	4	5	6	7	8	9	10	11	12	13													
14	15	16	17	18	19	20	21	22	23	24	25	26													
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



Our application hasn't gotten very far as of yet. However, this month we are going to tie a lot of stuff together and finish the first instance of it up. Let's jump right into it, Shall we?

First, let's jump over to our index.html file in the root directory. We need to create a section for data to get written out. Right before the end of the article (</article>), make a new section containing an H3 and an unordered list. Give the section and the list unique ID's and the H3 some text that defines what will be displayed. Your code should end up something like the following:

```
<section
id="ubuntuVersionsDisplay">
  <h3>Ubuntu Versions</h3>
  <ul
id="ubuntuVersionsDisplayList"
"></ul>
</section>
```

That's it for the HTML. We need to jump back over to the main.js

file in our js directory. Find your processForm() function, and we will start out by examining a hint. Sometimes you will run into not knowing what you can do with certain objects. If you look back at your code, you will notice that we passed the id of the form itself to the processForm function. We don't really need it for getting the information because we also gave IDs to the input text fields, which makes it much easier to manage.

Using our ge function, we are going to get the object with the id's ubuVersNum and ubuVersName. We will put those into two different variables, one should be in ubuVersNumValue and the other in ubuVersNameValue accordingly. Simply throwing the objects of the text fields won't really help us though; we need to know what was written in them. We need the value. In the case of the version numbers, variable declaration should look something like this:

```
var ubuVersNumValue =
ge('ubuVersNum').value
```

```
var ubuVersNumValue = ge('ubuVersNum').value,
    ubuVersNameValue = ge('ubuVersName').value;
console.log(ubuVersNumValue);
console.log(ubuVersNameValue);
```

The only thing new in this statement is the .value piece. .value is actually a property of the text field object. After we get values into variables, let's console log them out and make sure we have everything working. Shown above is the code block.

Now you can fill out the form and submit it, and you should see whatever you put in the two form fields in the console. Although this is pretty cool at first, it isn't very productive visually to anyone other than ourselves. We are going to continually write to the console, but that is just for us, not our visitors/users. We need to display this stuff on our page. Remember that section we made a bit earlier? Now is the time to use it.

Continuing in the processForm function... We are going to continue right after the last block

of code where we made the contents of our form into variables. We are going to start out with a new if statement to make sure something got written in the number field.

```
if (ubuVersNumValue != '') {
}
```

If the ubuVersNumVariable does not equal nothing, then we want to display the information to the screen. We need to declare some new variables, these variables are going to hold some HTML elements. We need to start out by getting the element (ge) for the UL we created in our section. Then we need to create an element for an LI. We are also going to format the text (next page, top left) we want to enter into that list.

Notice we are using the document object model again to



```
if (ubuVersNumValue != '') {
  var list = ge('ubuntuVersionsDisplayList'),
      listItem = document.createElement('li'),
      listText = ubuVersNumValue + ": " +
        ubuVersNameValue;
}
```

create an li. That line actually uses DOM to create an element. So, by giving it the argument of 'li', that line will take it and make a `` and ready it for our needs. The next variable is just taking our variables and separating them via a colon. The first variable should have been no surprise by now, so we are going to move on. If you run that code, nothing will happen. Of course nothing will happen, we didn't put anything anywhere, we just made some variables.

Let's first put that text inside the listItem. The object's property for this is innerHTML. Using innerHTML will put anything inside of the `innerHTML`. The

next thing we are going to need to do (shown below) is append this child element to its parent, which is the UL. So, the list variable holds the parent object, and we are going to append its child, listItem: `list.appendChild(listItem)`.

Now, if you run this code, the screen doesn't refresh because we took away the default functionality of the form, and whatever we put in the form fields will just start displaying as a list under our heading for the Ubuntu Versions. You could also make the if statement (that checks for the ubuVersNumValue) to also check for the version name if you want to get real fancy. Maybe even an alert

```
if (ubuVersNumValue != '') {
  var list = ge('ubuntuVersionsDisplayList'),
      listItem = document.createElement('li'),
      listText = ubuVersNumValue + ": " + ubuVersNameValue;

  listItem.innerHTML = listText;
  list.appendChild(listItem);
}
```

to tell you what you need to fill out if you get it wrong.

This would be a good point to teach you a little about versioning and making sure you are utilizing the power of git. Like I said, git is versioning software that allows you to keep track of your app and where it has gone. It also allows you to use GitHub and share your code, as well as look at the history of your code easily. First open up a terminal and apt-get install git. While that is installing, go ahead and go to github.com and sign up for an account if you haven't already. After you get done signing up, create a new repository and call it FCM-UbuntuVers-Tut.

By now, your git software should be done installing. In your terminal, change your directory (cd) to where you have been coding your application. You can make sure you are in the right

directory by using the command:

```
pwd
```

It should display what directory you are currently in. Once you have confirmed you're in the right directory, we need to set up some global stuff and initialize git for the directory.

```
git config --global user.name "Full Name"
```

```
git config --global user.email email@address.com
```

```
git init
```

```
git add .
```

```
git commit -a -m "initial commit"
```

Future initializations will need only the last three commands. The first two were for setup purposes. Now let's get all that over to github.

```
git remote add origin git@github.com:UserName/FCM-UbuntuVers-Tut.git
```

```
git push origin master
```

Replace UserName with the username you used when you signed up for GitHub, and FCM-UbuntuVers-Tut should be the name you assigned to your repo. Repo is a short term for repository. You have essentially set up a repository for your code for this project. Each of your projects should get its own repository. When you push the repo to github, it is then public unless otherwise specified. So be careful about what you put up there.

Remember to check back next month. We will be going over stuff like local storage. You can also check out all the code at github: <https://github.com/aliendev/FCM-UbuntuVers>



Michael Youngblood has been in the industry of web design and development for 13 years. He has been working for a world wide wireless tech corp for six years and is working on his bachelor's of science in mobile development.



HOW-TO

Written by Ronnie Tucker

UPDATED!

Write For Full Circle Magazine

Guidelines

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

Rules

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

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

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

• In your article, please indicate where you would like a particular image to be placed by indicating the image name in a new paragraph or by embedding the

image in the ODT (Open Office) document.

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

• Do not use tables or any type of **bold** or *italic* formatting.

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

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

Translations

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

REVIEWS

Games/Applications

When reviewing games/applications please state clearly:

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

Hardware

When reviewing hardware please state clearly:

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

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



- ➔ 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

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



ASK THE NEW GUY

Written by Copil Yáñez

If you have a simple question and want an answer that doesn't involve taking sides in the "Amazon Lens" kerfuffle, contact me at copil.yanez@gmail.com.

Today's question is:

Q: How do I get help installing/using/tweaking Ubuntu?

A: You mean besides reading this magazine from cover to cover and memorizing everything in it, like it's the Torah and your Bat Mitzvah is next week?

Well, yes, there are other ways to get help. No, one of them is NOT asking your Uncle Leo, the one who says he can fix anything but actually just bangs on stuff with a hammer in the hope that it'll fix itself.

Below are a few suggestions on where to get help or find out more about setting up and using your Ubuntu install. But before you go, there are a few things you should

do to increase your likelihood of success. Happily, putting on pants is not one of them. Hurray for things you can do in pajamas!

As you may have already discovered, the Ubuntu community is very helpful. But even the most knowledgeable user can't guide you if she doesn't have some basic information about your specific situation. Let's gather that now.

Before you ask someone to help you with a problem:

1. Know what version of Ubuntu you're running. You can determine this in two ways. From the command line (which you can get to by pressing CTRL-ALT-T), type:

```
cat /etc/issue
```

This will tell you which version of Ubuntu is installed. If you prefer using menus, open the Dash search field and type System Settings. Click on the System Settings icon (looks like a wrench and a gear), and, in the window that pops up, look for the Details icon (just a



gear). Click that and you'll get a window like this. Same info, different presentation. Both work just fine.

2. Be aware of what you were doing when the issue arose. Does your monitor freeze after you click on a particular program icon? Does your network connection drop after doing a specific task? The more you know about what was happening when your problem or question arose, the more likely someone can help you.

3. Make a note of any error messages. Sometimes a window will pop up with strange hieroglyphics or incomprehensible tech babble. No, you haven't been

contacted by the Na'vi. That error code might not make sense to you, but it might be the key that helps someone else resolve your issue. Write it down or take a picture with your phone so you have it handy.

Okay, so you have a problem; you've determined what version of Ubuntu you have installed and you made a note of what was going on when the problem occurred. Now you're ready to get specific help.

The most obvious choice is simply to enter the message (or error code) into your favorite search engine, being sure to include which version of Ubuntu you're running. A lot of times this does the trick, especially if it's a pretty common problem. The search engine will return a number of viable solutions and you're on your way.

Of course, this method can return a lot of false positives, too. If that happens, you might feel like a sound engineer at a Nickelback jam session, wading through a lot



of muck to find anything worthwhile. Don't get too frustrated. Instead, head over to the largest Ubuntu forum <http://ubuntuforums.org/>. There you'll find FAQs, wikis and forums – with thousands of pages of information. If your problem is a common one, it's likely to be covered in a FAQ. If not, you can always ask in the forums. Before you do, though, be sure to do a thorough search to see if your question has been asked before. And if you do post a question, make sure it's posted in the correct sub-forum. Nothing gets a frequent contributor's knickers in a twist like seeing a post titled "How do I Change My Wallpaper?" in the Networking sub-forum.

Sometimes, your problem will be specific to your particular derivative (Lubuntu, Xubuntu, Edubuntu, etc.). When this happens, head over to the specific website dedicated to the version of Ubuntu you're running. In some cases, they'll have a better idea of how to fix your, say, Lubuntu-specific problem.

Another place you can turn for your Ubuntu questions is IRC channels. These are like party lines

where hundreds of users get together to discuss any topic you can imagine, each with its own category identified by a name that starts with a pound sign. One of the nice things about using IRC channels is that they allow for a back-and-forth with other users. If someone needs to know something about your specific machine to help fix your problem, but you don't know how to get the requested information, the other user can guide you and answer questions along the way. Even if you can't find your answer on an IRC channel, sometimes you will get additional information that you can use to post a question to a forum.

To visit IRC channels, you need an IRC client. Empathy is the default in Ubuntu, and you can find it by typing Empathy into the Dash search field. You can also run an IRC client from a browser by going here <http://webchat.freenode.net/>. In both cases, you'll be asked to choose a nickname (anything will do) and what channel you want to connect to. Try `#ubuntu` or `#ubuntu-beginners`. Don't get too distracted by all the text. Just type in your question and then keep an

eye out for any user answering you (usually referring to the nickname you entered).

Another community with a lot of active users is Reddit. If you've never used Reddit, it's like a newspaper run by capricious and pun-happy netizens. Anyone can post links, questions and comments anonymously, and the rest of the community (known as Redditors) can respond. Just like with forums, however, be sure to do a thorough search to see if your question has been answered already. If not, post your question in the appropriate "sub-reddit." Redditors can upvote or downvote both questions and answers. The best answers float to the top and less helpful answers move to the bottom (or at least that's how it's supposed to work - don't be surprised if the top answer to your question is a jpg of Sarcastic Willy Wonka).

Well, as you can see, there are a number of places to look for answers to your pressing Ubuntu questions. Never suffer alone with a problem you think is intractable. Likely, someone else has been in the same position as you, and found a way to resolve the issue –

leaving a trail of breadcrumbs for others to follow.

Your very last, most desperate move is to send ME an email with your problem. Seriously, I'm just as likely to get you arrested, I'm THAT ignorant.

Let me put it this way. If you think of helpful Ubuntu sites as the 17 people in line to be President of the United States, I'm the Secretary of Energy. Sure, I'm not as bad as the Secretary of Education, but still. You can do better.

Good luck and happy Ubunting!



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



How do you compare apples and oranges? Comparing Windows to Linux really is a bit like comparing an apple and an orange; they're members of the same thing (operating systems), but quite different once you work with them.

As a system builder who builds both, I've always wanted to find some way to measure the performance of both. Windows has lots of benchmarking tools, 3DMark, Prime95, PCMark, Passmark, Novabench and SiSoft Sandra to name a few. Linux too has several benchmarking tools, bonnie++, glmark, httest, etc. But comparing Windows benchmarks to different Linux benchmarks just doesn't work well.

Enter Michael Larabel's Phoronix Test Suite (PTS). Phoronix Test Suite is a cross-platform benchmark and testing suite that runs on Windows, Linux, Mac OS X and the BSD operating systems.

I used PTS to benchmark 2

systems, a HP 6710b dual core notebook running Linux Mint 14 MATE (64-bit edition), and a Phenom 4 core desktop running Ubuntu 12.04 (32-bit edition).

PTS is available in the Ubuntu software repositories, but the version in Ubuntu 12.04 was out-of-date, so I opted to download directly from the Phoronix Test Suite download site: <http://www.phoronix-test-suite.com/?k=downloads>

If you choose to open with Ubuntu Software Centre, you may get a message about the software being poorly packaged. It should run just fine if you continue to install it. Hopefully the packaging will be cleared up in the future.

The Phoronix Test Suite has a massive amount of tests covering everything from Disk to System to Graphics to Processor to Memory to Network. Tests can take from a few minutes to a few hours to download and run.

Part of the beauty of the

Phoronix Test Suite is its simplicity. It's a bit like the Drupal content management system in that you can pick what you want to install and run. PTS can also download and run a suite of tests (the cpu suite, for example), but be prepared for a long wait as some suites contain 25+ tests.

Installing from the PTS website puts an icon in the unity dash. That icon can be used to launch PTS, but since PTS is primarily a command line tool, it's easier just to run the command:

```
phoronix-test-suite
```

Running `phoronix-test-suite` lists the arguments PTS takes. To see the list of tests available run:

```
phoronix-test-suite list-available-tests
```

Be patient, PTS may take a minute or so to download the list of tests the first time.

Tests are not installed by default, so you'll need to install a test. In this example, we download

the bork file encryptor processor test:

```
phoronix-test-suite install bork
```

Be patient, downloading and installing some of the tests can take time. During that time the computer may appear to be doing nothing, but your hard drive lights should show activity.

To run a test, use the `run` argument along with the name of the test:

```
phoronix-test-suite run bork
```

In the case above, we run the `bork` test, which would have been previously downloaded.

`Bork` ran relatively quickly taking just over 4 of the estimated 7 minutes on our Phenom-based quad-core test machine. PTS outputs test results as a number which can be opened as a graph in a browser and uploaded to <http://openbenchmarking.org/> – the open, collaborating benchmarking site also developed

by Phoronix.

What makes PTS and openbenchmarking.org even more useful is the ability to compare results. Thousands of participants have uploaded their results from various tests to the openbenchmarking.org site. If you have PTS installed, you can compare your system with an already benchmarked system by running: **phoronix-test-suite benchmark** and the name of the system/bench following the result in the URL.

Earlier, I mentioned it was possible to run a suite of tests; **phoronix-test-suite list-available-suites** lists the suites of tests available. I installed and ran the cpu suite of tests on a dual core HP 6710b notebook running Linux Mint 14, then I ran **phoronix-test-suite benchmark 1212094-SU-CPUSUITE444** on the AMD Phenom based desktop to compare the exact same suite of tests. Running the command on a new system automatically downloads and installs the same tests installed on the old system. The results on the new system are then compared against the old system. 1212094-SU-CPUSUITE444

```
charm@eomie:~$ phoronix-test-suite list-available-tests

Phoronix Test Suite v4.0.1
Available Tests

pts/aio-stress          - AIO-Stress          Disk
pts/apache             - Apache Benchmark   System
pts/apitrace           - APITrace            Graphics
pts/battery-power-usage - Battery Power Usage System
pts/blogbench          - BlogBench           Disk
pts/bork               - Bork File Encrypter Processor
pts/botan              - Botan               Processor
pts/build-imagemagick  - Timed ImageMagick Compilation Processor
pts/build-linux-kernel - Timed Linux Kernel Compilation Processor
```

is the name that openbenchmarking.org assigned the initial test system. Part of the name, CPUSUITE, is based on the name I initially gave the test on the HP notebook. Better to name your tests something more useful like HP6710b2GBRAM, something that better reflects the system that the test is being run on.

In our comparison the dual core notebook actually beat the four core desktop on 3 of the 25 tests.

The full results of the tests and complete specifications for each system can be seen on the

openbenchmarking.org web site at: <http://openbenchmarking.org/result/1212094-SU-CPUSUITE444/>

Speaking about specifications, the phoronix-test-suite is also useful as a tool for getting information about your system. If all you want to do is find out system information, you can do it without downloading any additional tests - by running:

phoronix-test-suite system-info

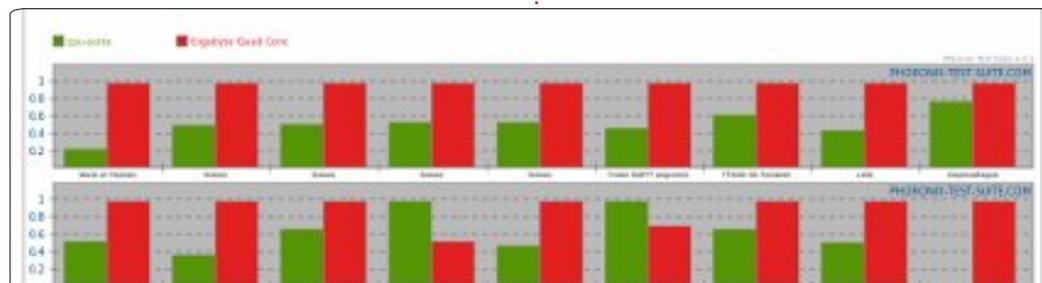
To get detailed system

information run:

phoronix-test-suite detailed-system-info

This covers most of the basics of the phoronix test suite. PTS is capable of much more if your needs are greater. If you're managing a lot of systems Phoronix also offers commercial support for the Phoronix test suite. More information about commercial support is available on the Phoronix Test Suite site: <http://www.phoronix-test-suite.com/?k=commercial>

Phoronix creator Michael Larabel also maintains <http://www.phoronix.com/> where he has written hundreds of articles about Linux support on the latest hardware.





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



CLOSING WINDOWS

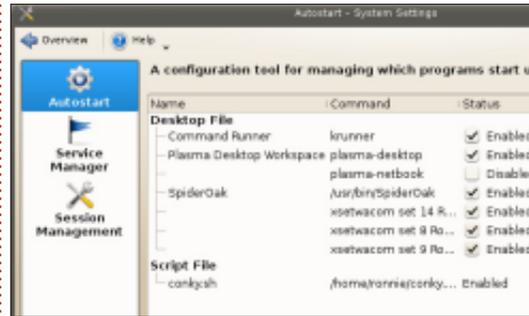
Written by:
Ronnie Tucker (KDE)
Jan Mussche (Gnome)
Elizabeth Krumbach (XFCE)
Mark Boyajian (LXDE)



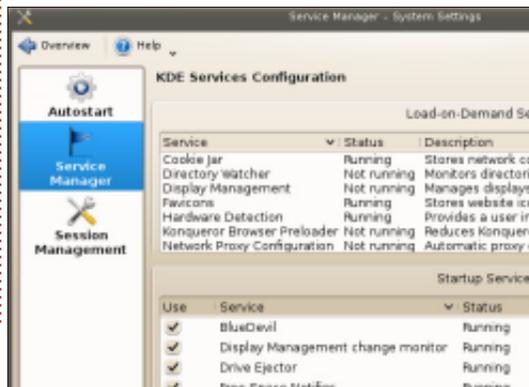
Having an application start, automatically, on startup is something akin to black magic – even in Windows XP. It's not exactly obvious, but dragging or copying an application (or shortcut to it) to the Start > Startup folder is how it's done.

Kubuntu

Like everything else in KDE, you can alter the startup applications by going into System Settings. The first option on the left of the window is 'Autostart', and it's here you add an application/script to your startup.



The Service Manager option on the left is for enabling/disabling actual KDE services such as bluetooth, power management, audio mixing and such like.



Running An App On Startup

Lastly, Session Management deals with how KDE will react when you choose to shut down the machine. Should it log out, restart the PC, or shut down the PC by default?

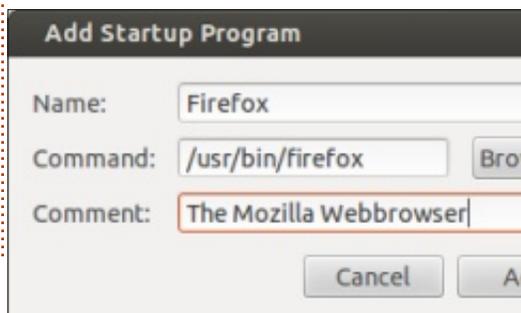
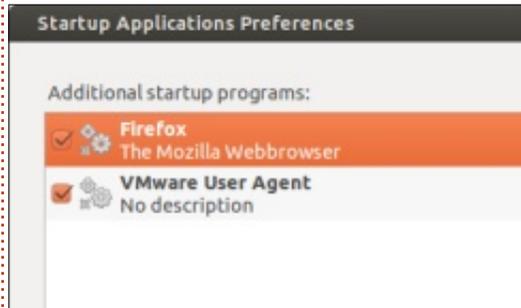
Gnome-Shell

In Gnome-Shell it's also very simple to add a program to the list of programs which should start when you boot the OS. In Applications > System Tools > Preferences > Startup Applications, you can add a program.

Click the Add button. Then, in the new window (Add Startup Program), give your program a name, browse for the file to start, and, if you like, you can also give some comment about it. Click Add again and the item will appear in the list as can be seen in the picture. After adding a program to the list, you can even select if you want it to start during the next boot, tick the box in front of the item, or, if not, untick the box. The item will remain in the list but it can be de-activated.

If you want to know what to put into the "Command" box, then have a look at the shortcut which is made for the program in the main menu. Copy-Paste the path and program name into the field and at next boot your program is running.

That's all there is to tell about autostart in the Gnome-Shell interface of Ubuntu. It's not much but it just shows how easy it is to use this OS.



Lubuntu

We have finally reached a topic for which the process in Lubuntu is the epitome of “simple and lightweight” ...and totally NOT at all intuitive! That “revelation” is in the context of MS Windows users migrating to Lubuntu, and in that context this statement is (IMHO) true.

I’ll admit it right now: there is no GUI application to add a program to the startup list. Also, you’ll need to be “root” (a.k.a. administrator) to alter this list. And, being Linux, there are many ways to accomplish this very simple task. All that said, I am going to present one way to accomplish this task; you may then alter it to your liking. Now that the requisite “confessions” are on the table, let’s see how incredibly easy it is to add startup applications to Lubuntu.

Here’s a summary of what you’re going to do:

- You are going to edit the “autostart” file that keeps the list of applications you want to start automatically when you start your system.

- This is a system file, so you will need to be “root” to edit this file.

Yes, that’s all you need to do. Let’s get started:

- Using the file manager, navigate to /etc/xdg/lxsession/lubuntu
- Typically, you will see two files in this directory:
 - autostart
 - desktop.conf
- From the main menu in the File Manager, select Tools > Open current folder as root
- You will be prompted for your password; enter your password and press Enter or click OK.
- A new window will open giving you root access to these files.
- Double-click the autostart file; it will open in a text editor

You will likely see a few lines of text in this file; leave them alone. These are applications that are already in the list to “autostart” when you log in to Lubuntu; you are going to add your desired applications to this list.

To add an application(s) to the list of programs to autostart, simply add the “launch name” (i.e. the name the system uses to launch the program) of a program you want to run when you start

Lubuntu, preceded by an “@” symbol. For example, if you want AbiWord to automatically start when you log in to Lubuntu, simply add this line to the end of the existing list (if any) in the autostart file:

```
@abiword
```

Save the file. That’s it - you’re done! The next time you log into Lubuntu, AbiWord will start and be waiting for you to use.

You can add as many applications as you want to this list to autostart. Likewise, if you no longer want an application to autostart, simply remove it from this list.

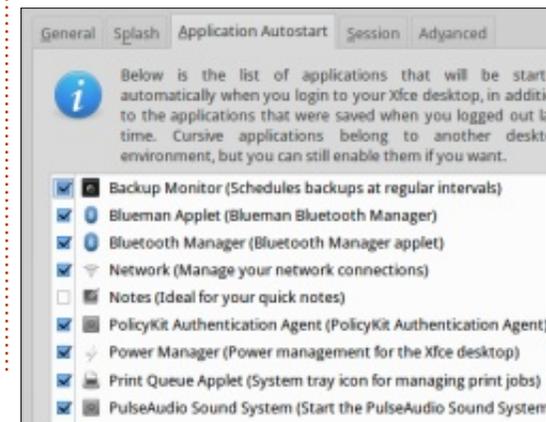
Xubuntu

In Xubuntu, having an application autostart when you log in is pretty straight-forward. You go to the Settings Manager and scroll down to the “System” section where you will want to open the icon called “Session and Startup.”

Session and Startup allows you to customize the Xfce splash screen, and other behavior related to logging in, and includes a tab for

“Application Autostart” which you will want to select. Inside this tab you’ll see a long list of predefined applications that you can choose to enable or disable upon log in.

If the application you wish to autostart is not in the list, you can click the “Add” button to define this information yourself. The fields you define are similar to what you would fill out for customizing your launch in Gnome-Shell, but in a different order. You can also select any application already in the list and select “Edit” to see how that is configured – if you’re looking for tips on setting something up. If you ever wish to stop auto-starting an application, you can unselect the checkbox next to it in this menu and that will disable it; you can also fully remove the definition with the “Remove” button.





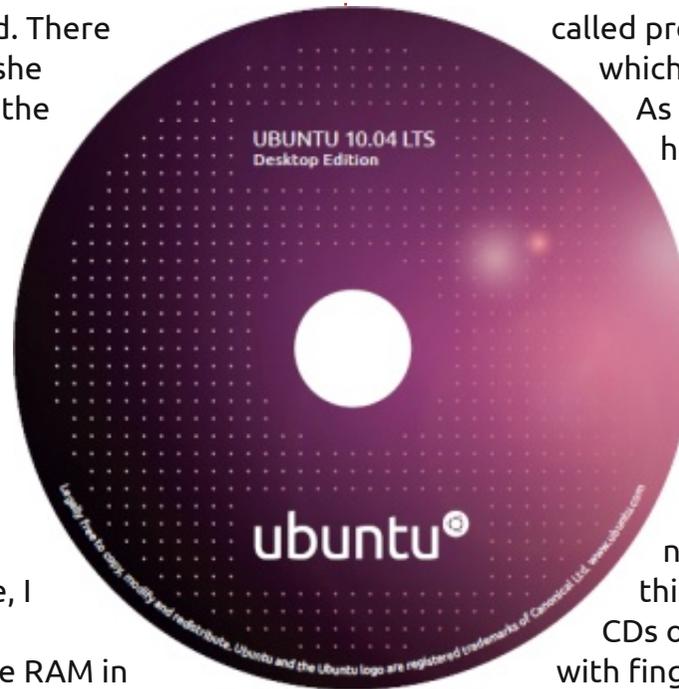
MY STORY

Written by Bill Pritchard

My first computer was a used Fujitsu-Siemens with 128 MB of RAM, loaded with Windows XP. It was 6 years old when I got it in 2005. Up to that point I had used machines at work using Windows 95 on a company network, using Lotus Notes. I had a bit of fun with this machine, which eventually died giving me a nice blank blue screen. I searched the support sites but eventually checked the hardware for faults and, lo and behold, poor contact of the RAM modules was rectified and I was back on the air. Up to this point I was working with dial-up, and then I got broadband. I then got a desktop built by the local shop based on an ASRock board with 1 GB of RAM and Windows XP. This was followed by a wireless router from my ISP which I used to set up a network with a Toshiba laptop with Win XP, which both my wife and I could use downstairs while my router and desktop were kept upstairs with the printer. So I got my wife using email! An ability to run iPlayer, and show photos from the laptop on TV, was also

well received. There was no way she would allow the living or dining room to be permanently cluttered up with a desktop computer and printer. To improve performance, I successfully upgraded the RAM in the machines.

So, now I had a home network, wireless, powerline ethernet, broadband, printing and laptop. My son's Fujitsu of the same vintage had by this time also died, so for the fun of it I brought it back from the dead with a new hard drive, which I successfully fitted. But what about the operating system? It was then I discovered that to get the machine back on, I had to buy a Windows XP CD at what I thought was a high price, and had things



called product codes which are a pain. As luck would have it, I found a magazine which offered a free CD with Ubuntu 10.04 LTS. I tentatively installed this, not knowing a thing about live CDs or Linux, and, with fingers crossed, got Ubuntu up and running.

It took a while to get a printer working with a suitable driver from the Internet; and then, eventually, I successfully added the machine to the network.

The AsRock desktop started crashing, so I loaded this with Ubuntu 10.04, but the problem persisted and was solved only when I replaced the hard drive and reloaded Ubuntu. This worked very well and was added to the network. I then decided to build up

an up-to-date machine based on a Asus motherboard package from Maplin with a SATA hard drive. (The local computer shop gave me a case free of charge.) This I also loaded with Ubuntu 10.04, and, when 12.04 was offered as an upgrade, I managed to get that running. I am now pretty familiar with the new desktop, and can switch from one machine to another without too much trouble. When I ran into troubles with the scanner, partitions, antivirus, and others, the help of Launchpad support with Ubuntu was invaluable.

So what to do next? The TV to computer interface looked interesting. I already had a basic Sky system, so, using a device called Neuros OSD, available from Amazon, I could record TV onto a USB stick or network. In addition, this is an open system which allowed access via telnet from a networked PC. I thought SSH looked interesting, but it was not installed. This has been a lot of fun to play with, but my wife did not agree and thought upgrading to



Sky+, which allowed recording, was a lot more user friendly. I had to agree, but Neuros remains connected to the TV because it plays movies from USB sticks and is a good backup to Sky.

By this time I was struggling up think up another feasible but interesting project, and started reading various Linux magazines. A lot of the content initially seemed beyond me, so I thought I would investigate C programming. It seemed a bit more academic than the projects I had already tried, and it was a job I could fit in at spare moments. A netbook seemed the answer, so I got a Toshiba unit – because my old Toshiba laptop still worked after 5 years of continuous use, although I had recently upgraded to a newer unit which runs Windows 7, which my wife was happy with. The Netbook came preloaded with Windows 7, so I dual loaded it with Ubuntu 12.04 LTS. I got most of my Windows programs running on it with the current exceptions of the chess program Fritz 8, which I use on the Netbook Windows 7 operating system, and Active sync for my pda, which I use on my old Win XP laptop. The difficulty which I have not solved is that I cannot

get Wi-Fi to work on Ubuntu 12.04. I have made do with a Netgear USB wireless dongle to date. I am coming to grips with the terminal and using the command line, and am gradually trying out some of the ideas in magazines such as Full Circle and Linux Format. I am studying the Cloud, Virtual Box, and LXC containers to see what I am best able to make use of.

As far as hardware is concerned, I have faulty sectors on my Asus desktop hard drive, so the

immediate task is to decide what to do regarding a backup and a new hard disk. I am very disappointed with this 1 TB Samsung hard disk. I have looked at tablets and phones, but apart from perhaps Android based devices, I see no real advantage for me at low cost. I am continuing with the programming tutorial and reading magazines, so perhaps some ideas will emerge from that work. I considered building a router on an old PC but the existing router supplied by my ISP

is likely to be far more reliable than one based on a hard drive in an old PC. In fact, I have never bought a new desktop, only revived old ones or built custom units.

In conclusion, I now have a total of 7 machines available to me: 2 running Win 7, 2 running Win XP, 2 running Ubuntu 10.04 and 2 running Ubuntu 12.04. I have one dual booted and one XP virtual machine, one is dedicated to printing and one I use for backup. I am reluctant to install other distros without long term support. I mainly use Linux, but my wife uses Windows 7, which was preloaded in the new Laptop at purchase. The jury is out as to where to go next, but the journey has been great fun.





MY OPINION

Written by Kevin B. O'Brien

Recently a Norwegian Kindle owner named Linn suddenly found her Kindle had been wiped clean of all the books she had purchased from Amazon. Whether this was a simple mistake or yet another act of evilness is not the main point of what I want to say here. It could have been either, or both, and it would not be relevant. The key is that Amazon could do this, and the "owner" could do nothing. According to Amazon's Kindle Store terms of use: "Kindle content is licensed, not sold". If you try to remove the DRM (not hard to do, actually), or transfer your purchase to another device, Amazon may legally "revoke your access to the Kindle Store and the Kindle Content without refund of any fees."

This particular type of evil is traceable to the software industry. They introduced the concept which you see expressed at the beginning of every EULA for proprietary software "This software is licensed, not sold". Note the identical language here?

This was a questionable practice when done by the software industry, and court cases have gone both ways on whether this is enforceable. And a well-crafted court case might overthrow Amazon's use of this tactic since they clearly say you are "Buying" when you are on the website or purchasing from within the app or Kindle device. But that is not what I rely on. And while I know how to break DRM, and anyone sufficiently motivated can find out how to do that with a little Googling, I don't think that is the optimal response either.

To me the optimal response would be one that punished companies that impose DRM and take ownership away from you – even though you give them your hard-earned money. And the only way to do that is not do business with them. You see, if you buy an e-book from Amazon, and then break the DRM, you haven't sent any kind of message to them. If instead you patronize a seller that does not impose DRM, you send a signal that you will pay for products that

respect your freedom. Fortunately, that is increasingly possible where e-books, audiobooks, and music are concerned. But this is a strategy that is not without its drawbacks, so you need to understand the trade-offs and go into this with your eyes open.

The first thing to understand is that it is primarily the publishers and rights-holders who insist on the DRM. Some of them have discovered that removing DRM does them no harm, and may do them some good. And if these companies start to see increased sales from dropping DRM and giving you back the rights you would normally have with a physical product, it might move the rest of the industry to stop being so evil. What rights do we mean? Well, if I buy a CD, I can lend it to my friend. If I buy a book I can sell it to a used book store when I am done. When I die, I can pass along my books and CDs to my heirs. They may just sell them all, but the point is that I have ownership rights to any physical product that allow me to own the product and

act accordingly. When those products become digital and you encounter the "this product is licensed not sold," what has just happened is that all of your rights have been removed.

Because some rights-holders have become enlightened, but others have not, the main trade-off you would encounter is that some products you might want to purchase are not available in ways that respect your freedom. For some people that might be a deal-breaker, but for me it is not. If I want to buy music, there is so much good stuff available to me that I could never get enough money to buy it all. But if I specifically want to buy the latest hit track that is on the top of the charts, I may not be able to do it if the record company is one of the troglodytes. As it happens, I am one of those curmudgeons who thinks most of the music that is "popular" today is crap anyway, so I don't really mind. As long as I can find lots of music I like to listen to, I am happy. Same thing with books (both e-books and audiobooks). I



can read or listen to only so many books in the time I have (for me, time is more of a limit than money for these things). I can find more books than I have time for, books that I really want to read/listen to, without giving up my rights. But, again, if I wanted to get the latest #1 book on the New York Times bestseller list, I might not be able to get it in a format that respects my freedom. For me, I don't care. I figure it is their loss when I don't buy their book. But this is the essential trade-off you will encounter if you go for freedom, you will have to occasionally accept that some products are just not available on those terms. My hope is that if enough people do value freedom enough to deliberately make those purchases, this will send a two-pronged market signal: publishers that do not respect your freedom will see sales go down, and publishers that do respect your freedom will see sales go up. As an example, the recent Humble e-book bundle, which was DRM-free files that do respect your freedom, sold so many copies that each book in the bundle would have qualified as a New York Times bestseller if the New York Times counted e-books. Which they will

do eventually when they solve their rectal-cranial insertion problem. So how do you do this? Fortunately, it is not that hard. I will mention some of the options, but the good news is that there are so many options available that I cannot list them all.

Music

The first question here is whether you are looking for traditionally sold-by-the-record-company tracks or the more indie Creative Commons self-published tracks. Both have their options, including some that are hybrids.

Major Label tracks, online storage and streaming

- **eMusic** - Monthly subscription lets you buy tracks priced at \$.49 - \$.79. No DRM. Tracks are from record labels, and a lot of back catalog is available.

- **Amazon** - Though they are bad with e-books and audiobooks, they sell tracks from the major labels without DRM. A little more expensive than eMusic, but more likely to have that one track you



must must have if eMusic doesn't have it. Also offers online storage and streaming of your tracks.

- **iTunes** - initially sold tracks with DRM, switched to selling tracks without DRM beginning in 2007. So this is an option as well. Does not currently offer online storage and streaming, but this may arrive in 2013 according to news reports.

- **Google Play** - Also offers DRM-free tracks from the record labels. Also offers online storage and streaming of your tracks.

- **Ubuntu One** - Yet another DRM-free source for major label tracks. Also offers online storage and streaming of your tracks.

Indie and Creative Commons

- **Soundcloud** - This is a music and audio sharing site, primarily.

- **Free Music Archive** - Lots of CC-licensed music.

- **Jamendo** - One of the premiere

CC music sites.

- **Bandcamp** - I just learned about this site from my friend Craig Maloney, who does the Open Metal Cast. This site has Creative Commons music from bands who want to build a relationship with their fans and sell them music. Good artists like Amanda Palmer are here.

e-Books

When it comes to books, you really are at the mercy of the individual publishers. Most music labels have finally come to accept that no DRM is the best way to go, but most book publishers are still being dragged into the 21st century kicking and screaming. But there are some good places to find e-books that respect your freedom.

- **Project Gutenberg** - This is the granddaddy of the DRM-free book sites. Project Gutenberg makes available books that are in the

Public Domain, i.e., where the copyright has run out. These are mainly older books, but a lot of classics are in here. They make books available in all of the major formats.

- **Baen Books** - This publisher specializes in the harder Science Fiction, but they really understand the new media landscape. They not only offer most of their books DRM-free and in multiple formats, but they also have the Baen Free Library, where they offer selected books free of charge. The hope is that with the first taste free, you will want to buy more. And it works. I went there to see what they had, discovered that they had the entire collected works of one of my favorite authors (James H. Schmitz) for sale, and bought the lot of them.

- **Tor/Forge** - A major publisher in the Science Fiction and Fantasy fields, they just moved to going DRM-free a few months ago. They did this because other publishers had been successful in so doing.

- **Angry Robot** - Along with Baen, a pioneer in selling DRM-free books in the Science Fiction and Fantasy fields.

- **Avon Romance** - A major publisher of romance novels, they just announced that they are

experimenting with DRM-free ebook sales.

- **O'Reilly Media** - The premiere publisher of technical books, they pretty much get everything right. They sell e-books without DRM. When a new edition of a book you already bought comes out, you can "upgrade" for a nominal fee (e.g. I upgraded my Kevin Purdy "Android" book for \$1). And with older books that they think are no longer worth being in print, they are removing the copyright and making them freely available.

- **ManyBooks.net** - This site has a lot of overlap with Project Gutenberg, but also has some newer works that have been made available, such as Charles Stross's *Accelerando*.

- **Fictionwise** - Although heavy on

the Science Fiction and Fantasy, has a lot of offerings in other genres as well. Reasonably priced and DRM-free.

- **Cory Doctorow** - Cory (pictured below) was one of the first authors to make a point of offering all of his works not only DRM-free but free of charge in e-book formats from his website. But, you know, when the book he co-authored with Charles Stross, *Rapture of the Nerds*, came out recently, I went to the Google Play store and bought it.

- **DriveThru Fiction** - An interesting site that also has Comics and RPG games available.

- **Apress** - A publisher of technical books that also offers reduced-price e-books if you have already purchased the print title. This is

something I'd like to see more of.

- **Packt Publishing** - Another technical book publisher with DRM-free books.

Audiobooks

This is where there is still a big disappointment. Audible, which is by any measure the clear leader here, insists on DRM on all of their books, which is why I refuse to get an account. Audible is now owned by Amazon, which sells music tracks as MP3 files without DRM, so there was hope when they bought Audible that we could get DRM-free audiobooks, but that was not the case. Fortunately, there are alternatives.

- **eMusic** - This is the same site I mentioned above for DRM-free music tracks. They also offer a subscription plan for audiobooks, \$10 a month gets you one book. Selection is not as good as Audible, but their list is growing all of the time, and I have had no trouble finding books there that interest me. I recently listened to Walter Isaacson's biography of Albert Einstein through a book I bought here.

- **Podiobooks** - This site offers



audiobooks in serialized form, much like podcasts offer you a file every week. Heavy on the Science Fiction and Fantasy at this point, but worth checking out. Scott Sigler and J.C. Hutchins are both available here, for instance.

• **Scott Sigler** - Scott (pictured below) used free content to get his name out, but still offers free audio versions on his website even though he now has a publisher.

• **Cory Doctorow** - Cory, in addition to offering free ebooks, also offers audiobooks that are DRM-free on a "name your own price" basis. Among the readers on his books are Neil Gaiman, Wil Wheaton, Spider Robinson, and Leo Laporte. He even sells files and CDs in Ogg format if you prefer to get your files that way. Dude seriously gets freedom, but if you know anything about Cory Doctorow you know that.

Conclusion

As I said above, the alternatives are not always perfect. Particularly with audiobooks, your selection is less than if you were willing to give your rights away, but there are enough that you can always find stuff you will enjoy. Second, I

probably missed many of the alternatives. The marketplace is changing rapidly and I don't always know everything going on. I just wanted to demonstrate that there are a sufficient number of viable alternatives that you don't have to sacrifice your freedom. So, if you do get caught by DRM, it is because you chose to. Third, while some of these offerings are free of charge, that is not the point. I selected items on the basis of respecting your rights and freedoms, and most of them require payment. In fact, that is the point. As Patrick Nielsen Hayden of Tor Books pointed out, it was the

commercial success of publishers (who offered DRM-free books) that got his company to try the experiment. We don't move the market by trying to find ways to not pay. We move the market by voting with our dollars for products that respect us. I hope I have given you enough ideas that you can help us to move to a DRM-free world.



Full Circle Podcast Episode 31, The Difficult Third Episode!!

NEW

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

Your hosts:

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



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

Olly & Tony bring everyone up to date with what's been going on in the "Summer Break", Olly talks about his Gnomebuntu install, Codecademy and Tony talks about his new Nexus 7.

Download



WHAT IS...

Written by TEAM LinQuestus

Raspberry Pi



The idea behind a tiny and cheap computer for kids came in 2006, when Eben Upton and his colleagues at the University of Cambridge's Computer Laboratory became concerned about the year-on-year decline in the numbers and skills levels of the A Level students applying to read Computer Science in each academic year.

From 2006 to 2008, Eben designed several versions of what

has now become the Raspberry Pi. By 2008, processors designed for mobile devices were becoming more affordable, and powerful enough to provide excellent multimedia, a feature that Eden felt would make the board desirable to kids who wouldn't initially be interested in a purely programming-oriented device.

Presently Pete Lomas is a trustee of the Raspberry Pi Foundation. He designed the final hardware that turned into the

Raspberry Pi.

Nowadays the Raspberry Pi is used in a majority of embedded devices from education to entertainment. Many researchers have started working on a Raspberry Pi to provide an efficient and economic computing Device.

About Raspberry Pi

The Raspberry Pi is a debit-card

sized computer. By plugging in your TV and a keyboard, it can work as your normal computer system. It is a small PC, which can be used for many of the things that your desktop PC does, like spreadsheets, word processing and games. You can watch Youtube video with HD Quality. You can program for your robot capture snapshots of different areas. By attaching an external display board, you can also design display boards for advertisements. Not only embedded devices, but one

WHAT IS

can also deploy a server with portability and plug-n-play characteristic.

With approximately \$25 cost, it is very cheap for students and researchers to work in the field of embedded and open source. There are two models, RPi-A and RPi-B, available in the market. Technical details of both are shown right.

Article compiled By:
TEAM LinQuestus

Prof. Chirag K. Gohel
(Assistant Professor Department of Computer & IT)

Mr. Himanshu Madhvani
(Final Year Student IT)

Mr. Sunil Lalvani
(Final Year Student IT)

Mr. Sagar Rangani
(Final Year Student IT)

Ms. Avni Savaliya
(Final Year Student IT)

<http://www.facebook.com/Linquestus>

<http://linquestus.in>

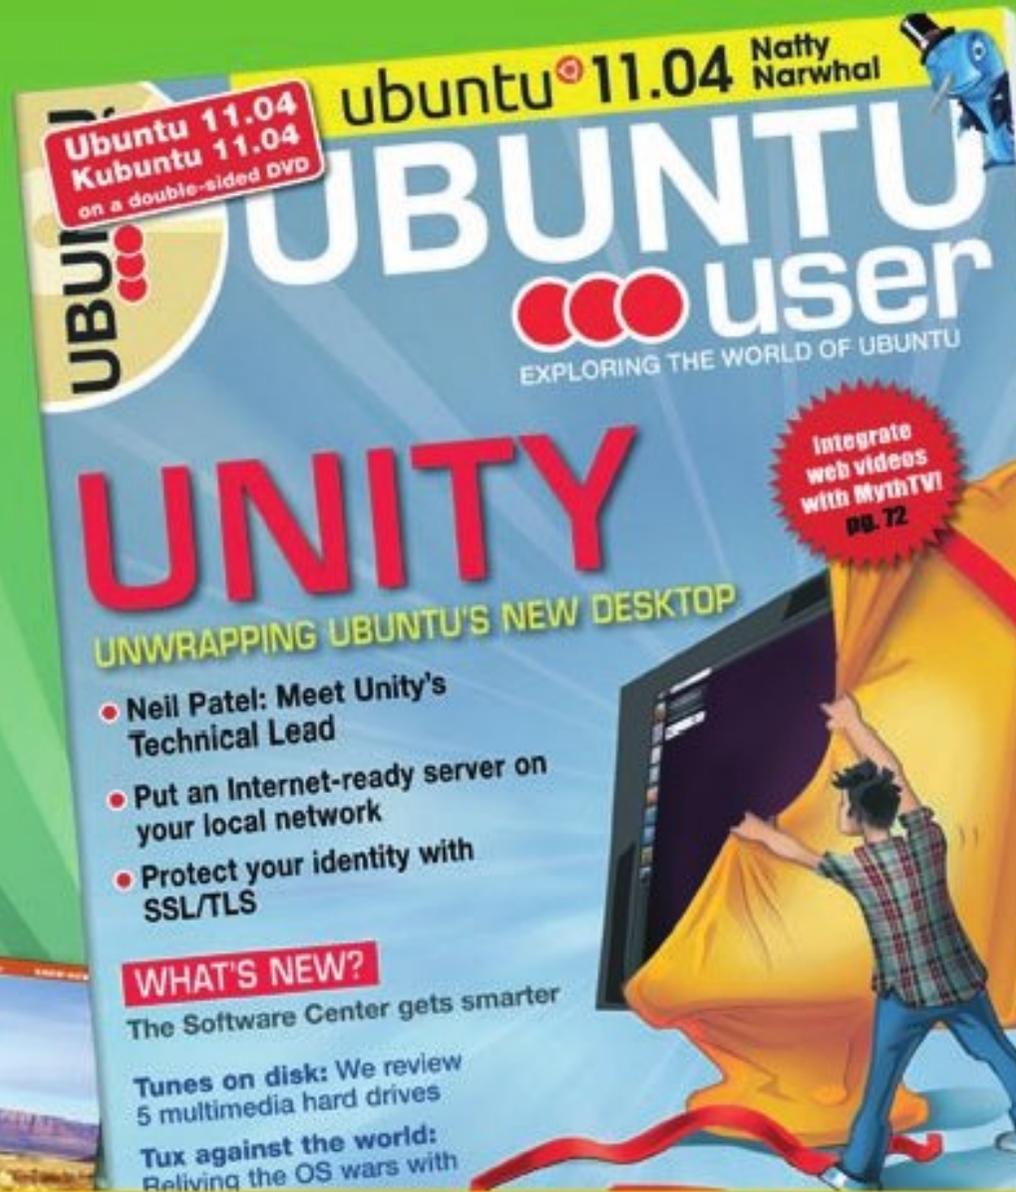
Specification	Model - A	Model - B
Approx Price	\$25	\$35
SOC	<u>Broadcom BCM2835</u> (CPU, GPU, DSP, SDRAM, and single USB port), <u>Broadcom Media Processor</u>	
CPU	700 <u>Mhz</u> ARM (11 Family)	
GPU	<u>Broadcom VideoCore IV</u>	
Memory	256 MB(Shared with GPU)	512 MB(Shared with GPU)
USB	1	2
Video Output	Composite RCA, 1.4 HDMI	
Audio Output	3.5 MM Jack, HDMI	
Storage	On Board SD/MMC Card	
Network	No	On Board 10/100 Ethernet
Power Ratings	300 <u>mA</u>	700 <u>mA</u>
Os	Debian GNU/Linux, Fedora, Arch Linux ARM	

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!

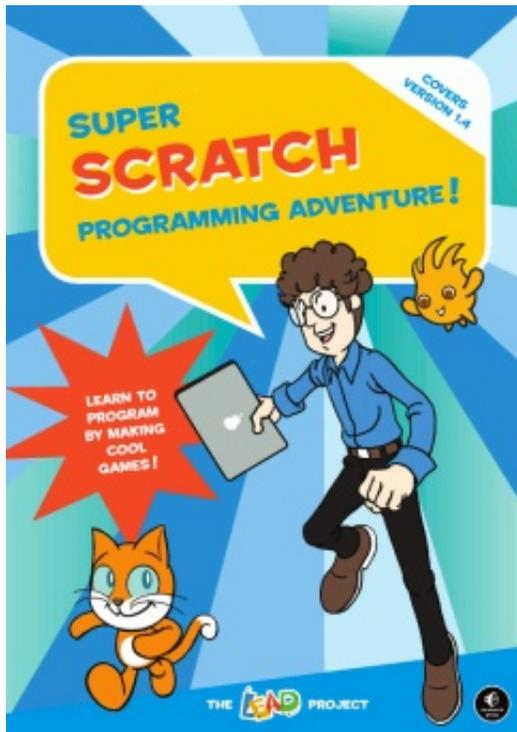


UBUNTU-USER.COM/SUBSCRIBE-NOW

FOLLOW US ON

TWITTER: UBUNTUSER

FACEBOOK: UBUNTUSERMAG



Super Scratch Programming Adventure is a comic-book styled programming primer. It's geared towards a younger audience, and covers version 1.4 of Scratch. It's about 158 pages long, and in full color. Each chapter of the comic develops the overall story, and gives the reader a reason behind the programming exercises (s)he's expected to complete. It definitely sounds like

an interesting idea that, if implemented correctly, could make game programming approachable for younger audiences. The question is: does it manage to do this?

The Explanations

The book is written well, and offers a screenshot for every step of the process. Due to the fact that the tools are introduced at the start of each chapter, this also allows the actual creation to flow fairly smoothly. The comic manages to convey the reasons for needing, for example, to plot a pentagon-like walking path, by making it relevant to the story. That being said – if you're not a big fan of comics (or find this book too childish), you may have trouble finding the motivation to follow the steps exactly. However, as the book is geared towards younger readers (age 8 and up), this should hardly be the case.

The explanations manage to cover the components and tools of Scratch well, to the point where

you feel comfortable using the interface, and, possibly, customizing certain steps of the exercises. The only possible issue that may arise is that, in certain “design” steps (for example, creating a space suit), the tools are explained, and then the next step is simply “draw your own suit!”. While they do offer the completed “astro-cat” sprite, it's hardly a sufficient way to help a child who may be struggling with that particular aspect. After all, some children will be adamant in managing to do it on their own, but designing anything like this can be frustrating beyond belief when you're not sure how to actually complete it.

That being said, the book is geared towards the programming side of things, which it does well. It gives you reasons behind the code that's implemented, while explaining it well enough for later usage in different projects. If you're looking for a complete guide into creating with Scratch, you'll probably want to find supplementary material for the

actual drawing/sprite creation steps.

The Language

There are certain words used in this book that I simply wouldn't expect an 8 year old to know. Stuff like “palette”, “magnifying”, “experiment”, “techniques”, etc. I may be misjudging our youth, but it seems to me like using simpler language (to some degree at least) would be beneficial. This same issue persists throughout the comic as well, using terms like “seven dimensional strings”. If your goal is to fascinate young children with the story in order to make your exercises seem relevant, it would be important to make the story flow smoothly for any reading level. This is something I find the book does not do sufficiently well.

If you don't mind answering questions like “what does this mean” or “what is a [...]”, then your kid should still thoroughly enjoy the book. Conversely, if your kid has access to a child-friendly

dictionary, the problem can also be averted. That being said, I feel that the book could have included a glossary of terms, and solved this problem completely.

Other than that, the language used is clear and patient. However, I don't feel the tone of the overall instruction is animated or all too witty, which could be a problem when trying to keep the attention of children. This could, however, be simply because I am far outside the intended age group, and as such may be misjudging younger readers' reactions.

The Art

This isn't usually a category with which one would score a technical book, but due to the fact that it's attempting to tell a story through a visual medium (a comic), and trying to teach the creation of video games, it is a very relevant category for this particular book.

The comic is well-drawn, and the characters are consistent in their appearance. That being said, the art style in particular didn't appeal to me. It was a combination of the color choices and the

backgrounds used for certain panels. This may be due to the fact that I am not their intended audience, but I would have preferred a different visual style. The selected art, however, is very consistent and well-drawn.

The sprites and other in-game visuals are also very well done. This can affect the reader in two ways: it can be daunting to try to compete with that style, especially as a beginner. It can also make a game seem cheap when swapping between two artists like that. It can be solved by using the supplied sprites, regardless of whether or not you could complete the drawing stages on your own. However, I find that this takes away from the overall experience of designing a game, if you're using other people's sprites. Some may be of the opinion that sprites at this resolution can't look too different in the end. Generally, I would agree. However, while I admit that I haven't explored all the aspects of sprite creation, it still seems to me that there are sufficient options and customization available that two sprites created by two artists can look extremely different.

The Conclusion

Does this book teach you to program using Scratch? Yes, definitely. It will impart enough knowledge for you to begin working on separate projects, if you so wish.

Does this book teach you to create complete games using Scratch? This is where I find the book falls just a bit short. Yes, it teaches you the tools to create sprites and environments. Yes, it teaches you to program these environments. However, it does not teach you how to use the tools to get results of the quality that they exhibit throughout the book. For anyone who plans to use this book to create completely original games, I highly recommend you look for a complimentary video guide, or book, to creating sprites in Scratch.

The other minor problems I mentioned (language, art style) will impact the score only a little – as I'm essentially guessing at what someone in the correct age group would think of it. Overall, I'd give this book a 3/5. While the book does claim to only teach you to

program, it also infers that you'll be creating games. While this is true in a technical definition of the word "create", most would assume this includes graphical creation as well. While that skill is covered by introducing the tools and expecting you to complete certain "drawing" steps, the book falls short of actually teaching you to use these. I find the book would be improved by ignoring the "drawing" steps completely and offering their ready-made sprites only, or else take the reader through step-by-step.

Full Circle would like to thank NoStarch Press for the review copy of this book. Support those who support FCM and buy from NoStarch.com



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.



REVIEW

Written by Ronnie Tucker

Linux Mint Debian Edition KDE

Linux Mint Debian Edition KDE (LMDE KDE) is an unofficial spin of Linux Mint Debian Edition (LMDE) which is, as the name suggests, based on Debian and uses Gnome. LMDE KDE uses, as you might have guessed, KDE.

Note, I say it's based on Debian. Kubuntu is based on Ubuntu which is based on Debian, so there are differences between using LMDE KDE and Kubuntu.

History

I'll let LMDE KDE creator **Schoelje** give you the background to his baby:
"The unofficial LMDE KDE, like most new things, was born out of frustration. I loved LMDE but I did not like Mate or Cinnamon. It felt like trying to revive an abandoned technology. As if Frankenstein had gone digital, but needed practice on its stitching. I liked Mint KDE, but I missed LMDE's stability. There's been many rumors that Clem [creator of Linux Mint] had

intentions to create LMDE KDE, but that never happened. So, I decided to create my own Boris Karloff: strong, stable and... beautiful. First, it was just a bunch of scripts and packages gathered from several Debian and Ubuntu repositories, but then I thought that there may be some people out there who are experiencing the same frustration of having no LMDE KDE, and that's when I decided to create the live DVD. Although, I, too, need some stitching practice, but I think the result has some merit."

Install

The first thing to note is that as of 12.10, the Ubuntu family now uses PAE as part of the kernel and booting process. This means that older hardware will not boot 12.10+. I noticed this when trying to install Kubuntu 12.10 on to an old EEE PC. Kubuntu 12.10 wouldn't even boot. The 32-bit LMDE KDE comes with two kernels that you can choose from on boot. The newer PAE version and also a non-PAE version. So this is handy for older hardware.



LMDE KDE installs pretty much the same as the Ubuntu family does.

A nice addition is that you can ping several repositories (after installing):

```
sudo mint-choose-debian-mirror
```

and the result is that you can choose, from a list, the fastest

repository. By default, LMDE KDE uses the LMDE repositories.

Welcome

On first boot, LMDE KDE will present you with the welcome window (shown above) which automatically detects third-party hardware with the option to install proprietary drivers. It's something that's taken for granted these days with the ease of use in K/Ubuntu,



but it's not the defacto standard in Debian.

Another nice addition to LMDE KDE is the Plymouth Manager which lets you tinker with the loading screen by choosing a resolution and relevant image from the lists provided.

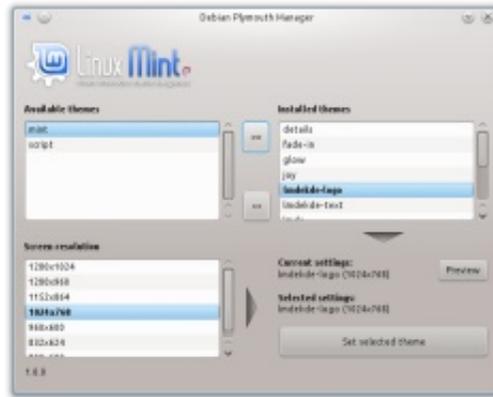
The desktop looks as you'd expect – default KDE with a Linux Mint wallpaper – so no surprises there.

Software

The usual great suite of KDE applications (as well as LibreOffice) are installed; K3B, Amarok, digiKam, Gwenview, and, if you like to import/export to photo sharing/hosting sites, you can install the extras package which will let you natively use Facebook, Picasa, et al from within digiKam and Gwenview. One unusual change is the lack of Kmail, but with the addition of Thunderbird. Might not suit all, but suits me fine and dandy.

Updates

Rather than having a dozen or



more updates every day, LMDE KDE will use the same system as LMDE with update 'packs' which will mean fewer updates, but this is to ensure extra stability.

Conclusion

Not having the Ubuntu repos/PPAs is a bit of a downer. I can't deny it. KDE 4.10 is around the corner and LMDE KDE is using 4.8, which makes me a bit sad. But, having said that, since it shares features with Debian, it may skip 4.9 altogether and use 4.10 (as they did with 4.5). With regards to PPAs, not having them does stop me trying more experimental applications, which means that, in the long run, my system is far more stable than previously. Yes, I'm not forced to install those unstable packages in Kubuntu, but when

they're there, and so easy to install... I'm sure you know how it is.

It's still a youngster, but I'm definitely a fan of LMDE KDE. UP6 isn't far off, and it'll be interesting to see what its creator, Schoelje, will introduce.

Pros:

Based on the granddaddy of them all.

Cons:

Can't use K/Ubuntu repos/PPAs

Download and info:

<http://forums.linuxmint.com/viewtopic.php?f=61&t=113571&p=633324>



UPDATE FROM SCHOELJE

In one of Clem's recent blog posts (<http://blog.linuxmint.com/?p=2237>), he had this to say about LMDE KDE and LMDE Xfce:

"Going forward, LMDE ISOs will be

released in two versions: MATE and Cinnamon. The decision was taken not to maintain a KDE version of LMDE and to stop maintaining the Xfce version. This was a painful decision to make considering the quality of LMDE Xfce and unofficial LMDE KDE, but one that made sense for Linux Mint since LMDE Xfce represented only 4% of the Linux Mint audience."

Together with the unofficial LMDE KDE, I (Schoelje) have decided to also take up maintenance of the unofficial LMDE Xfce. I have now started an "unofficial LMDE Xfce thread" here: <http://forums.linuxmint.com/viewtopic.php?f=61&t=118890>



I am **Schoelje** (www.schoelje.nl), business consultant by trade, and I live in the south of Holland with my wife, two kids and two cats. I started my career in IT a few years after I finished at the art academy. In those years, I got my experience in developing business solutions in VB6, C# and ASP.NET. Yes, that one. In January 2012, I started using Linux for the first time in my life which has brought me here.



goo.gl/FRTMI



facebook.com/fullcircle-magazine



twitter.com/#!/fullcirclemag



linkedin.com/company/full-circle-magazine



ubuntuforums.org/forum-display.php?f=270

I Like Syncs.

I have played with previous versions of Ubuntu and other Linux OS. I can say I enjoy using 12.04. After my wife had some kind of a virus downloaded and Windows 7 would no longer work, I made a dual boot system and she is now also becoming a fan. To the folks who work to make an OS that works, thank you.

I also like how I can sync Ubuntu One with my Android phone, But let's improve that a little more so if I have a file on the computer and not on the cloud drive, as long as I am on the same wifi I can access files from the computer through my phone.

K Green

Gord says: *put the file in a shared folder, then follow the steps in the second half of this page:*
<http://www.howtogeek.com/120882/how-to-access-shared-windows-folders-and-stream-videos-over-wifi-on-android/>

Flight Sims

I have recently become addicted to flight simulators. I wish they had been around when I was at school because they make geography come to life, require arithmetic skills for altitude, weight and currency calculations, teach meteorology and navigation, and the physics and science of designing and flying aircraft.

I own Microsoft 2004 and FSX, Lockheed Martin's Prepar3d, and X-Plane 10. All of them work well on 3 Windows desktops. However, I would prefer to use them on a Linux OS, preferably Mint. I googled for instructions on how to install them into Linux but, as is usual with Linux, there is little manual help, and usually that assumes that the reader can understand or guess missing technical details. So, can Full Circle provide a detailed guide to installing X-Plane under Linux Mint?

Dave Edwards

Motorized Telescope

Could somebody develop an article on how I could pilot my telescope from my laptop? Preferably with Python or C, and through USB (as parallel sockets no longer exist on laptops). Maybe explain how to move the telescope using up and down keys (for example).

Didier

Translating

I am really impressed by Full Circle. Actually, I am a new user – not just in Linux but even with computer science itself! You will, no doubt, be surprised to know that whatever I know today is the result of self-education over the last couple of years. Today, my system is fully Ubuntu. My next need was to learn Linux systematically, and for that I was in search of better literature which would not only help me to learn but also keep me abreast

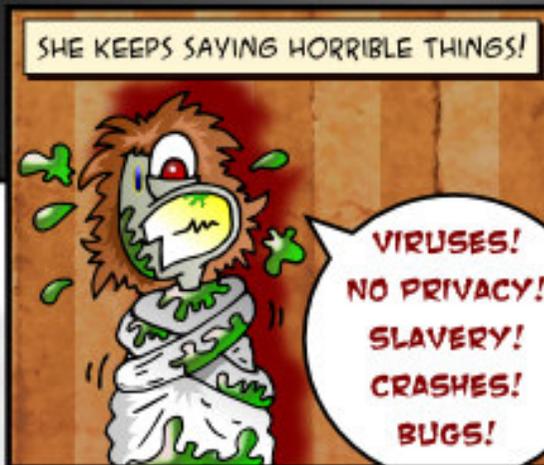
with the latest developments in the field. Through Ubuntu Software Centre in Ubuntu 12.04 and 12.10, I got an opportunity to find FULL CIRCLE. I liked the magazine very much. I downloaded all the previous issues and have started reading them.

I am from India, Gujarat state. Is there any team on the list of FULL CIRCLE for translating it into Gujarati language?

Navin Talati

Ronnie says: *I'm afraid there is no Gujarati translation of Full Circle. If anyone out there would like to translate Full Circle, please drop an email to ronnie@fullcirclemagazine.org.*

Tuxidermy



IT WAS SOMETHING LIKE "SUDO BEGONE".



Q&A

Compiled by Gord Campbell

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

Q I want to extract bits of FCM articles in plain text, to make my own synopsis in a small text file.

A To copy text from a PDF, position the cursor at the start of what you want, hold the mouse's left button down, and move to the end of what you want. When you release the button, the text should remain highlighted. Right-click and select "Copy." Now you should be able to paste it into any text editor. This works in Document Viewer and the Adobe Reader, and probably other PDF viewers. In Okular, the default PDF viewer for Kubuntu, you need to first go to the Tools menu and choose 'Text Select' to be able to select text, since the default 'tool' is the move/look tool to let you move through pages. It also appears that Okular does not work well if you select text with different formats, such as the drop-cap at the start of FCM articles.

Q I am trying to install Ubuntu 12.10. The computer is a Sony Vaio Pentium 4 1.8 GHz with 512 MB RAM.

A The minimum memory requirement for Ubuntu 12.10 is 768 MB. Try Xubuntu.

Q I am trying to take screenshots of selected areas of Firefox windows with Ctrl+Shift+PrtSrc, but I have no idea where these pictures get stored?

A They are saved in the clipboard. You can use GIMP, or something like that, and just paste them into a new image. You can also use the Screenshot program included in Ubuntu (hit the Super key and type "screens") which will afterwards ask you where you want to save the screenshot.

Using SHIFT+PRT-SCR (no ctrl) lets you select an area of the screen using your mouse, and, by default, saves the screenshot to the Pictures folder. Also, alt+prt-scr takes a picture of the currently-selected window, and prt-scr by itself takes a picture of the entire screen.

Q (Thanks to *thepreacherswife* in the Ubuntu Forums) As I was updating my system, including a new kernel version, I ran into a problem. When I rebooted, it went to a command-line shell and asked me to login.

A Reboot. At the grub prompt, choose an older kernel, and complete the system update.

Q Can I change the default paper size for printing?

A (Thanks to teamanx in the Ubuntu Forums) Yes, use this command:

```
gksudo gedit /etc/papersize
```

The following names are commonly understood by programs: a3, a4, a5, b5, letter, legal, executive, note and 11x17. For more info, use this command:

```
man papersize
```

Q I'm thinking of getting a [grossly expensive] laptop. The main thing I do is prepare presentations with HD images.

A (Thanks to *TheFu* in the Ubuntu Forums) You'd be better off with a moderately high-end desktop, using remote access back to it from any portable device. You could probably keep using your current laptop this way.



Q I have a wireless mouse that I want to use left-handed. However, on the left side of my laptop both USB ports are USB 3.0. If I plug the mouse into either of those ports, it does not work.

A (Thanks to *howefield* in the Ubuntu Forums) Plug in the mouse before you boot the computer.

Q According to VLC site, they are urging everyone to update to 2.0.4 due to a security vulnerability. Will Ubuntu update VLC soon? I'm uneasy using it with a serious security vulnerability.

A (Thanks to Cheesmill in the Ubuntu Forums) If you read the whole security advisory, then you can see how small of an issue this is. An attacker would have to trick you into opening a specially crafted PNG image with VLC. (VLC is an unusual choice for viewing image files.) Even if this happens they will succeed in only crashing VLC, they can't do anything else or

gain access to your system. These sorts of vulnerabilities are common, there is no point worrying yourself about them.

Q I am trying to install Eclipse on my desktop. However, the software center is requesting that I place the Ubuntu 12.04.1 CD in the /cdrom mount. I do not have a CD-ROM drive installed right now.

A The easiest option is to run Software Sources and uncheck the CD-ROM.

Q I have dual boot, 12.4 and Vista. Is it possible to have an icon on the desktop to switch to Linux if using Vista or vice-versa, as opposed to picking which OS to use at boot up?

A No. You must select the OS at boot time.

Q How can I choose a USB infrared receiver to use for remote control of a home theatre PC?

A The web site lirc.org is devoted to IR under Linux. It could stand some updating, but it should get you pointed in the right direction.

Q When I connect wirelessly from my Dell laptop in Ubuntu 12.04 to my wireless router, my connection is very spotty. The connection drops frequently from 54Mbps to 1Mbps, and often drops entirely. lspci shows:
09:00.0 Network controller: Intel Corporation Centrino Wireless-N 1030 (rev 34)

A Enter this command:

```
gksudo gedit /etc/modprobe.d/blacklist.conf
```

and add "blacklist acer_wmi" (without the quotes) to the end.

Tips and Techniques



How hot?

Regular readers will know that I have a temperature fetish: I want to know how hot the various parts of my computer are at all times.

I even get into discussions with people about how hot different components are. However, it seems that a lot of people think of computer temperatures as abstract numbers, and don't put them in context. Therefore, the following chart for degrees Celsius:

- 100 - water boils
- 95 - too-hot CPU
- 70 - busy CPU
- 40-50 - idle CPU
- 37 - body temperature
- 22 - room temperature
- 0 - water freezes

As I write this, my body is warmer than my hard drive. With any luck, my hard drive will last as long as I have!



Full Circle Magazine would like to thank all those who emailed in questions about the development of Ubuntu. We tried to use as many questions as possible, merging those which were similar, and adding those we thought relevant.

Huge thanks go to both Alan Pope, for putting us in touch with Didier Roche, and to Didier for taking the time to answer all of our questions.

Like any good story, let's start at the beginning...

DESIGN

Q How much of a say does Canonical (and the community) have in Ubuntu, and how do you decide what features will be added to the next Ubuntu?

A Both questions seem to be linked to me. Basically, developers are free to decide and

scratch their own itch on what they think is good for Ubuntu. Of course, for the core user experience, we are requesting the design team's feedback, and we try to align our goals with the main ones, like, in the past few iterations, going to a daily usable Ubuntu image or having an easier process for people making upstream development in Canonical to deliver on the platform.

In addition to that, we all have little pet projects dear to our hearts, and we try to do them as time permits. An example for me is OneConf, which I couldn't push as far as I wanted, but I guess it's still a working and nice feature for our users. We don't do them because of x or y, we just do them because we think it's good for Ubuntu. I guess before UDS (Ubuntu Developer Summit), we are all browsing forums, looking for ideas on brainstorm or similar website (or just talking to our LOCO) to gather feedback and see what we can add/enhance.

Q How do you decide what applications will be included/excluded from Ubuntu?

A Those discussion arise at UDS, where everyone can participate by coming (it's a free and open event) or remotely. We try to start the discussion first on mailing lists, like the Ubuntu-desktop one (<https://lists.ubuntu.com/mailman/listinfo/ubuntu-desktop>) to focus those passionate discussions.

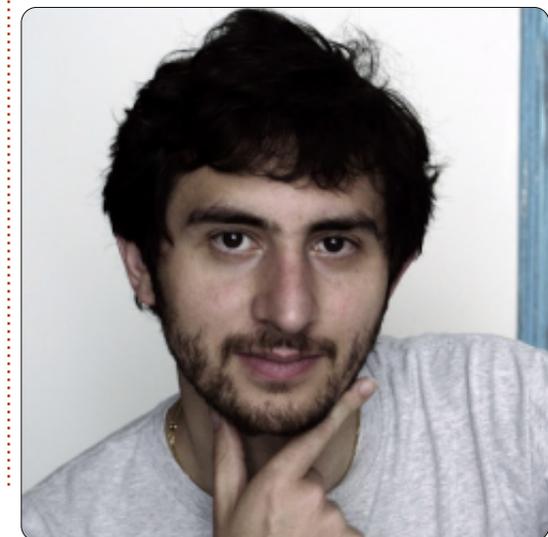
Q How much of Ubuntu is Debian, and what do you change?

A Lucas Nussbaum gave, a couple of years ago, a talk with some stats about this. I encourage you to read his insightful slides posted on <http://www.lucas-nussbaum.net/blog/?p=444>. The figures shouldn't have changed much, and I would say that we still have ~70% of packages untouched, directly imported from Debian,

15% of patched software compared to Debian to work with our toolchain and different set of dependencies (most of the times newer), and the rest is specific Ubuntu packages from our upstream and versions where we want to be ahead of Debian and pulled directly from them.

We have two main kind of differences:

- The user experience is focused, especially for Ubuntu (not speaking for different flavors, like Kubuntu, Lubuntu... here), and we have to make drastic choices to get to where we want the user



experience to be, even if this means diverging and modifying upstream components – this is the beauty of free software: if it doesn't fit your need, you can change it slightly.

- We also experiment with some advanced base support, like new stricter flags for the compiler, replacing bash by dash, adding some multi-arch support. Those kinds of advanced in-phase changes (taking a lot of time to have the whole set of packages in the archive working with them) are generally taken back in Debian. So Ubuntu endures the pain to do those transitions and Debian benefits from it.

Q How many teams construct Ubuntu (and who are they), and how is the work divided between the in-house/community teams?

A This is, surprisingly, a difficult question to answer. We have different teams in Canonical contributing to Ubuntu: the desktop team, foundations team, server team, kernel team, security team, community team, QA team, archive admins, release

team, but, in addition to that, we also have some upstreams like the team making the Unity or Ubuntu-one experience, who don't have upload rights to Ubuntu, but, by their code, really contribute to the distro. Also, there are a lot of focused teams by flavors, like Kubuntu, Lubuntu, Xubuntu, Edubuntu, Ubuntu Studio... Some of them have people overlapping as well. Not counting even the translations and documentation teams...

Also, if you count “constructing Ubuntu”, we shouldn't forget all other upstreams like Gnome, Xorg, OpenStack, as well as the Debian developers and the whole LOCO ecosystem. We all build Ubuntu in some way. Even people helping each other in a forum, or writing a document on a wiki, is building Ubuntu in their way. Also, maintaining the websites, even if they don't touch directly to Ubuntu, and contributing to launchpad or the main web site, they are contributing to Ubuntu.

DEVELOPMENT

Q What do the teams program in? C? Assembly? Amiga Basic (just kidding!)?

A No Amiga Basic AFAIK! Most of the development we are doing is mainly in C, C++, Python, Vala, Go, and shell of course. We need to have some fluency as well on autotools, cmake and makefile syntax. Sometimes, we have to touch perl when we really need to.

Q Do you have in-house specialists for networking, drivers, etc.?

A Indeed, even if we are “divided” across the teams, it doesn't mean everyone is doing everything. We all have some specialization (but not restricted to it). So, we have Mathieu in charge of network manager, the kernel team is responsible for the kernel drivers.

Q Do you use VMware/VirtualBox while developing Ubuntu?

A Not really, we do run it natively on our day-to-day laptop, which is way better, isn't it? Right now, my daily machine is on Raring (the development release), for instance.

Q How do all the programmers/teams communicate?

A We are using IRC (on freenode) for most of our discussions; others, needing less instant feedback but more development, are on the mailing lists (<https://lists.ubuntu.com/>). Everything is public, and you can even read the logs of the different channels ordered by date since Ubuntu's creation – available at <http://irclogs.ubuntu.com/>.

Q How do you keep up with all the changes in hardware which Ubuntu needs to run on?

A The kernel team is just awesome doing this task. Of course, most of the changes are directly

coming at the Linux kernel level and we benefit from those new supports, taking always the latest available kernel, but the hardware enablement team is working on enabling more and more hardware as well.

Q How is all the code organised between teams/individuals?

A Historically, we just had 2 divisions: core developers (<https://launchpad.net/~ubuntu-core-dev>) who can change any part of the archive, and MOTU (master of the universe: <https://launchpad.net/~motu>) who can modify only what is in universe/multiverse. Main is a smaller asset than universe which is self contained (meaning, we can build main with main components only), and basically (to really simplify, it's not a 1-to-1 match) contains what is supported officially and installed by default in Ubuntu (not flavors).

Nowadays, the landscape is more complicated. We have package sets, when people can earn their upload rights to only, for

instance, the “desktop components”; we also have per package upload rights for people interested in just one particular component. You can find all those detailed on <https://wiki.ubuntu.com/UbuntuDevelopers>.

Q How is the code compiled? What hardware is used in compiling, and how long does it take?

A We are building Raring on 4 different types of architectures: i386 (classic 32-bit machines, the default), amd64 (64-bit), powerpc (previous Apple machines processor architecture), and armhf.

This means that each package and sources are built on those 4 different kind of architectures, and, depending on what you install Ubuntu on, it will install packages built for your machine (this is a simplification again, we do have some packages like images assets, that are built only once, because they don't contain code that will produce different results from one architecture to another one).

What is happening is that Ubuntu developers are uploading (by ftp) a signed source package to launchpad. Then the build farm (<https://launchpad.net/builders>) will pick them up and split the load between the different machines. Once it's done, the binary package is published in the main archive, and then will get replicated to the Ubuntu mirrors.

Another source is the Debian repository, as we are syncing at the start of each release from Debian all packages we can sync (meaning there have been no Ubuntu changes compared to Debian – the 70% I was talking about earlier).

Of course, there are additional manual reviews for new packages, or packages needing to migrate from universe to main and so on...

Q Is the kernel tweaked specifically for Ubuntu?

A Indeed, the kernel is changed for Ubuntu. Most of the changes are reversed back to the Linux kernel itself, but always taking the latest kernel helps

doing some QA, spotting regressions and getting them fixed.

Q Is it possible for users to acquire parts of the Ubuntu code for modification? If so, where do they download it?

A That's really easy, and part of the default toolset. You have to enable the “source” repository (in `/etc/apt/sources.list`), or using Software Sources and enabling the “source code” checkbox. After a refresh from the repository, you can download any source you want from the command line using `apt-get source <package_name>`. Want the Unity source?

```
apt-get source unity.
```

REPOSITORIES

Q Are the repositories maintained by the community, or does Canonical look after them?

A The repository needs to have a level of security and high reliability. As the repositories are signing the binary packages with their private keys, we need to be accountable for what we deliver to our users and ensure there is no possibility to introduce evil content with the right signing. There are only very few people accessing those machines, with some Canonical IT guys and some old/long time Ubuntu developers (employed by Canonical). Also, you will notice that Canonical is funding the high cost of bandwidth and maintenance of those critical pieces.

Q Is there a person/team who looks after the Ubuntu Software Center side of things, or is it just a front end for Synaptic?

A There is a small team led by Michael Vogt delivering Ubuntu Software Center; you should also know that he's a long time apt contributor, maintainer of Synaptic, the previous gnome-app-install, and upgrade-manager. So you can see that those pieces are

all tied together by the same people (and it's quite funny to read on forums that people creating Synaptic are genius, and not the Ubuntu Software Center or the contrary).

Q How much hardware does it take to run the repositories?

A Wow, I have really no idea for that TBH – “a lot of bandwidth”. Don't forget that the repositories' content is mirrored to a lot of places to get lower latency. Those servers are not maintained by Canonical, but as told previously, the main archive contents are signed. We distribute the corresponding signatures to the user machines, and thus apt is able to check the integrity of the archive copy on the mirrors, and so we can ensure nothing was changed in between.

TESTING

Q How many people/teams test Ubuntu, and how often? Also, how is testing done? Is it per

package, or as a whole distro?

A There is the community team's awesome Nicholas Skaggs leading an effort to have more and more community participation on testing. In addition to that, we have the QA team, making regular install and upgrade testing.

Testing upgrades from release to release is a really difficult task, especially with the transitions and other things changing between releases. To address that, we have automated upgrade testing which installs an older version of Ubuntu on a machine, changes some configurations, and upgrades to the next release. We also have a daily automated latest ISO installer, ensuring that the latest produced ISO can be installed every day.

So, we do have multiple kinds of testing:

- unit testing that's enabled during package build. The build will fail if it doesn't pass.
- auto package testing, which is testing against the installed version of the component. Those won't be copied into the release pocket if they don't pass.

- ISO and upgrade testing, made every day and automatically.
- ISO and some components manual testing, on a regular basis (Nicholas is trying to get help on this, follow planet Ubuntu to get more info).
- some components like the whole Unity ecosystem (60 components) have additional tests that are run on a daily basis even before it's uploaded to distro.

Q How do you ensure that vital/necessary packages aren't missing?

A We are getting more and more automated tests enabling running a full session and testing applications as well as core experience. This is how we can see that a critical piece, after a daily ISO installation, is missing.

Note as well that packaging missing by accident will probably be on a mismatch component list (missing from main or being there for no reason) – this is another way of spotting this :) If the latest ISO couldn't build due to a missing components or a mismatch, we'll notice it quickly.

Mismatch shouldn't happen anymore with additional steps we introduced in Raring. Now all packages are going to a proposed pocket (similar to when we do Stable Release Updates), and are validated there before being copied to the release pocket which is the main archive. This validation ensures we don't get the archive into a non-working state.

Q Do you test on old machines? Eg: parallel port, floppy disk, etc.

A This is with manual testing most of the time. Martin Pitt is introducing some mock objects (fake objects used for testing target) to be able to spot regression for old configurations like that, but, let's be honest, Ubuntu's goal is not to be able to run on so old hardware, there are other distributions targeted for 10+ year-old machine :)

Q Is Ubuntu benchmarked during production, or prior to release?

A We have some limited automated benchmarks, but this is an area we are developing and are getting better and better with each release. This and automated testing are two historical weak focal points of free software. Nowadays, Ubuntu is changing this mindset, putting them at the core of the user experience. We are then growing on this and helping the whole ecosystem to evolve on it.

Q Is there an easy way (for users/testers) to follow a bug from reporting to fixing?

A Quite easy! Just find it on launchpad (normally, it's way more efficient to search directly on the component you think that is impacted by the bug), and click on the subscribe button on <https://bugs.launchpad.net/ubuntu/+bug/1>. You will get all notification (and numerous comments) on status changes. When the Ubuntu component of the bug is marked as "Fix Released", this means that the fix is now in the development version.

You can, on the launchpad interface, ask for backporting the fix to an older release and track its status here.

Q Is there a process for determining the origin of a defect in a test?

A This is based more on experience of the distro, knowing which components do what and dogfooding [\[http://en.wikipedia.org/wiki/Eating_your_own_dog_food\]](http://en.wikipedia.org/wiki/Eating_your_own_dog_food). Triaging bugs are a good start to get an overall overview on the components of a Ubuntu distribution and understanding what can be the root cause of a problem.

RELEASE

Q How is the final ISO compiled?

A The ISO is compiled every day, and you can find it at <http://cdimage.ubuntu.com/daily-live/current/> for Ubuntu itself. (changing the url, you can find the different flavors). A program called "germinate" (<https://wiki.ubuntu.com/Germinate>) takes some description/configuration files for what we install by default (called the "seeds", <https://wiki.ubuntu.com/SeedManagement>), and ensures that it's installing all the necessary dependencies. From this list, this will be installed in a live session with some additional packages to compress the image and produce a daily ISO.

So the final ISO is not that different from a daily one, it's just that we don't push any more packages unless their selected



fixes are near the end of the development cycle (processes are changing as well, so we also don't push any more new features), and, at some point, this ISO becomes the official one once the test results are in a good shape and we are proud of the quality to deliver it to our users.

We are trying to be more and more confident on features before pushing them to Ubuntu so that this stabilization period will become shorter and shorter, and even, why not, at some point, being able to say "each daily ISO can be the final one".

Q Who gives the final nod to release the ISO?

A The release team decides to call the ISO "finale". They traditionally gather in the London Bluefinn office for the release week to produce, and ensure everything is in one piece (and apparently drinking champagne once the release is announced), while the other remote employees and community are drinking water at home.

Q How is the ISO distributed to the various mirrors?

A I'm not familiar with the exact technical side, it's copied to a place that other mirrors are watching, and then rsynced as far as I know.

Q How are the various Ubuntu derivatives released? Do they get early access to the final ISO, or are they (X/L/Kubuntu) all separately compiled from Ubuntu?

A The flavors (not derivatives) are built from exactly the same archives as Ubuntu. So, when you install Shotwell on Ubuntu or Kubuntu, it's exactly the same binary package that is installed. The only differences are the set of packages that are installed and selected by default. For those, a separate ISO is produced by germinate taking a different seed; it uses a different configuration file for what is installed by default.

Other Info:

Let me just point to some links:

<https://wiki.ubuntu.com/UbuntuDevelopment>

for developing Ubuntu, a good start guide.

<http://developer.ubuntu.com/>
for developing apps ON Ubuntu

<https://wiki.ubuntu.com/RaringRingtail/ReleaseSchedule>

for the Ubuntu release schedule (for Raring)

<http://status.ubuntu.com/ubuntu-raring/>

to follow features and status on them that people working on Ubuntu are implementing

For testing and results, Nicholas is publishing a call for testing on planet.ubuntu.com, so watch this space.





The Steam for Linux beta officially launched on November 6th to a very limited set of participants. Since then, there have been a number of waves of additional testers, but, if you weren't one of the lucky ones selected, you can still get a chance to try the client before the final release.

Thanks to the clever folks over on Reddit, there is now a simple workaround to launch the beta without an invite, and here's how you can try it on your machine.

First, open up your terminal and enter the following:

```
wget
http://media.steampowered.com
/client/installer/steam.deb
```

```
sudo dpkg -i steam.deb &&
sudo apt-get install -f
```

Next, launch the Steam application from the dash. You'll then be prompted to login to your Steam account. You should receive an error reminding you that you

aren't enrolled in the beta, but just ignore it and close the dialog box.

Note: if you're running an AMD64 architecture, then you will also need to run the following command:

```
sudo apt-get install libjpeg-
turbo8:i386
libcurl3gnutls:i386
libogg0:i386 libpixmap-1-
0:i386 libSDL1.2debian:i386
libtheora0:i386
libvorbis0a:i386
libvorbisenc2:i386
libvorbisfile3:i386
libasound2:i386 libc6:i386
libgcc1:i386 libstdc++6:i386
libx11-6:i386 libxau6:i386
libxcb1:i386 libxdmcp6:i386
```

Once all of the dependencies are installed, you're finally ready to get going. To start the client just open the terminal and run:

```
steam steam://open/games
```

and you'll have access to your Steam library.

There are currently 29 officially supported games available, including titles Team Fortress 2,



Frozen Synapse and World of Goo.

My experience (using the above workaround) went pretty well on my modest laptop running Ubuntu 12.04 AMD64. Unfortunately, out of the games I tested, only two of the four games I launched successfully. While Braid and VVVVVV failed to find the executable file, Osmos and World of Goo looked beautiful and worked without a hitch.

In-game functions like Shift-Tab to access the steam Community (and other shortcuts) work as

expected. Visually, there are no real surprises - the client interface is sporting the same familiar look we all know from Windows and Mac. In short, some games may take a little tweaking to get up and running, but, overall, the Steam beta is an exciting look at what we can expect from Valve for the final version.



Jennifer is a fine arts student from the Chicagoland area. You can follow @missjendie on Twitter or visit her blog at missjendie.com.



I'm studying for exam 101, the first of two exams which comprise the first of three possible certifications from the Linux Professional Institute (LPI), or, simply, the LPIC-1. There will be six installments, give or take, charting my learning experience from layman to exam 101 certificate holder. In the articles, I'll share what I've learned about the exam itself, including syllabus content, resources and useful tips, as well as – where appropriate – my learning methods. In the future, I would also like to incorporate any information I can gather regarding my (hopefully!) improved employability as I (again hopefully!) progress through the curriculum and gain skills. In short, this article is ideal for anybody toying with the idea of investing time in learning for this qualification in the hope of landing a *buntu/Linux job.

The Exam

The first reality of exam 101,

which wasn't immediately obvious to me, is that the exam doesn't take place at the command-line. It's a multiple-choice test comprised of 60 questions to be completed in 90 minutes. This approach makes sense – for the examiners at least. Consider being presented with the following question testing the candidate's knowledge of how to redirect standard input, output and error streams:

Q: You want to store the standard output of the `ifconfig` command in a text file (`file.txt`) for future reference, and you want to wipe out any existing data in the file. How can you do so?

- a. `ifconfig < file.txt`
- b. `ifconfig >> file.txt`
- c. `ifconfig > file.txt`
- d. `ifconfig | file.txt`

(© Smith, R.W., Linux Professional Institute Certification: Study Guide, Second Edition. Wiley, Indiana: USA, 2009).

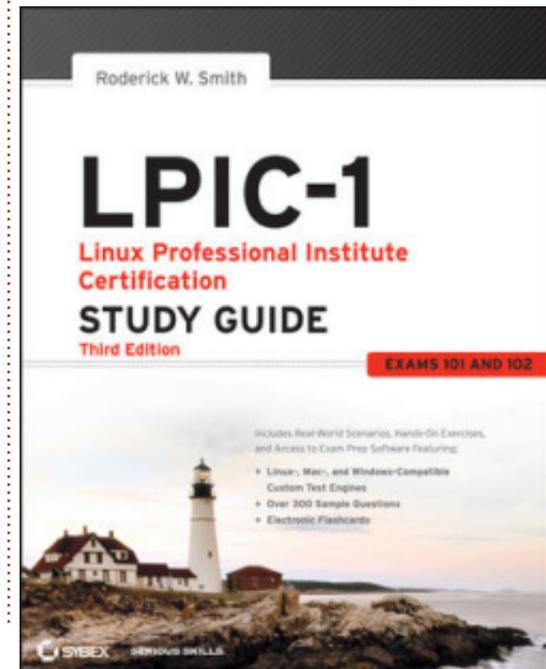
At the command-line, the candidate could simply execute each command in turn and examine

the results to arrive at the correct answer. In multiple-choice format, the candidate must be able to distinguish between the function of `>>` and `>`, which is key to answering this question correctly. If `file.txt` exists, `>` overwrites any existing content whereas `>>` would append to existing content. Both `>` and `>>` create `file.txt` if the file doesn't exist. (In this sense, the question is badly worded as it doesn't clarify explicitly if `file.txt` exists or not.) To sum up, if the testing took place at the command-line, no knowledge would be tested, yet, in multiple-choice format, the most nuanced differences are tested. Furthermore, the multiple-choice format eliminates the candidates' access to the man pages. Admittedly somewhat artificial, but effective in separating the wheat from the chaff.

Syllabus content and something useful

The first section of the book I'm working with, "Linux Professional

Institute Certification 1", explores Linux command-line tools, that is, text streams, filters, pipes, redirects and search, as well as search and replace with regular expressions. The commands covered in this section are: internal bash commands (`exec`, `exit/logout`, `pwd`, `cd`, `time`, `set`, `echo`); shell hotkeys (e.g. how to change lowercase to uppercase at the command-line – more on these later); `>`, `>>`, `<`, `<<`, `2>`, `2>>`, `&>` and `<>`; `|` (pipes); `cat`, `join`, `paste`, `od`, `sort`, `split`, `tr`, `unexpand`, `uniq`, `fmt`,



nl, pr, head, tail, less, more, cut, wc, grep and sed. Coverage of grep and sed is not exhaustive: as a case in point, an entire volume has been published on sed functionality alone. To finish, many of the commands I've learnt, I can't, as of yet, put to good use. For example, I don't have the need to study the structure of a graphics or audio file (usually non-ASCII) by using od (octal dump). But this is often the case with learning something new: one must be patient until one reaches a sort of critical mass, where knowledge suddenly becomes useful. And on that note, I would like to leave you with something useful I've learned: it's a shell hotkey for searching your command history. Hit Ctrl+R at the command-line, and start typing characters (the characters don't have to correspond to the first letters of the command or path you are searching). When you've found the command or path you want, hit ENTER to execute it, or hit Ctrl+G to return empty handed to the command-line.

Und tchuss.
(German for cheerio)

In my next article, I'll have a look at learning methods and

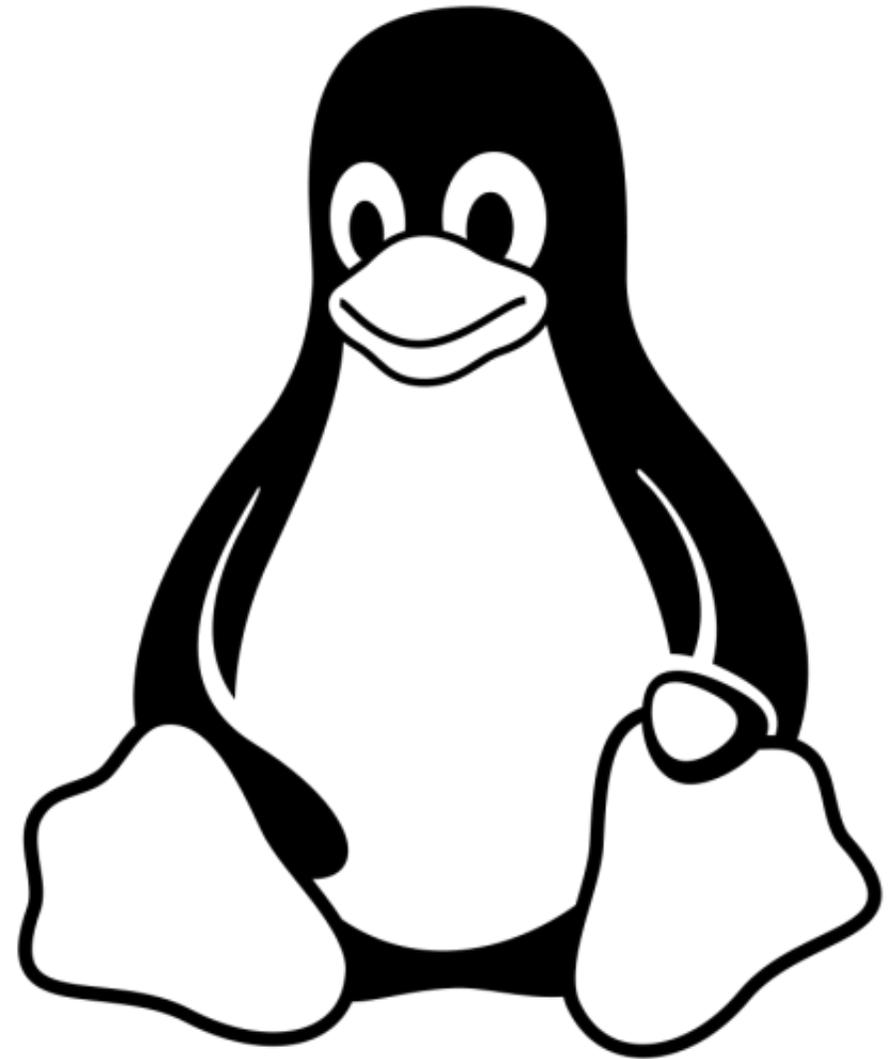
resources as well as more on syllabus content. That's it for now.



**Linux
Professional
Institute**



Richard Philip Witt, 30, living in Switzerland. I've been using Ubuntu desktop and server for roughly 2 years. But only for the basics. Now I want to discover its capabilities. Any mentors out there? Email me: chilledwinston2@hotmail.com.

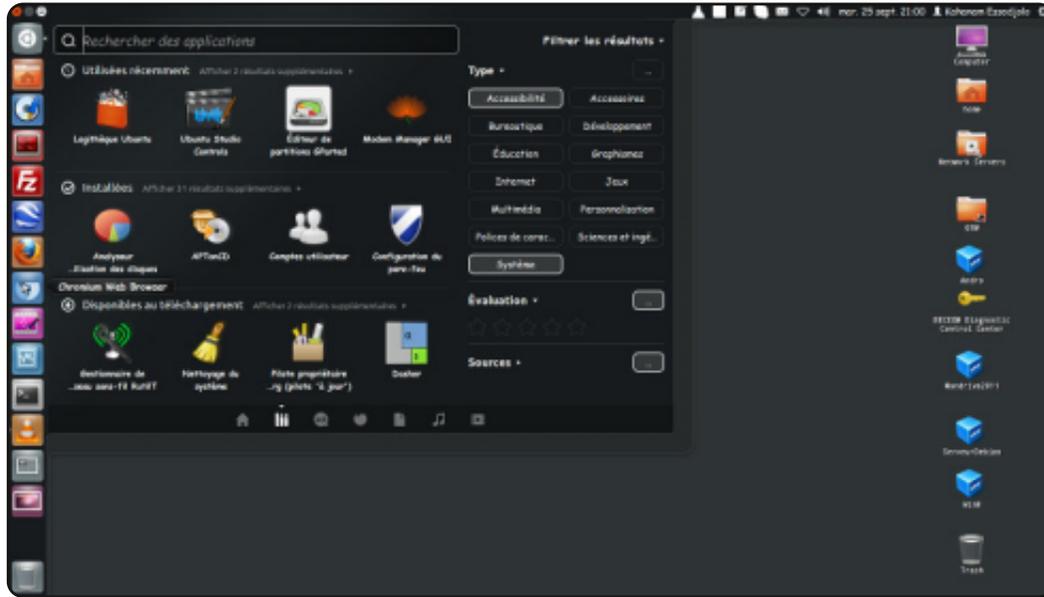


Linux



MY DESKTOP

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



Hi! I am Kahanam from Togo, and this is my desktop:

Operating System: Ubuntu 12.04 LTS Pangolin
CPU: Intel Core i3
RAM: 4 GB
HDD: 500 GB
Toshiba Satellite Pro.

I like Ubuntu and since the very beginning of this year (2012), I've chosen to use Ubuntu only. I like it. I enjoy it. It makes me happy.

Kahanam



I use Ubuntu 12.04, with icons for Mac OS; I love this robot as wallpaper. A part of my personal theme also results from the application, MyUnity. I also installed screenlets (notice the small notepad on the top right, Weather forecast, and processor monitor).

I have a Compaq Presario CQ57 PC, an AMD dual-core, 2 GB RAM, and a 250 GB HDD.

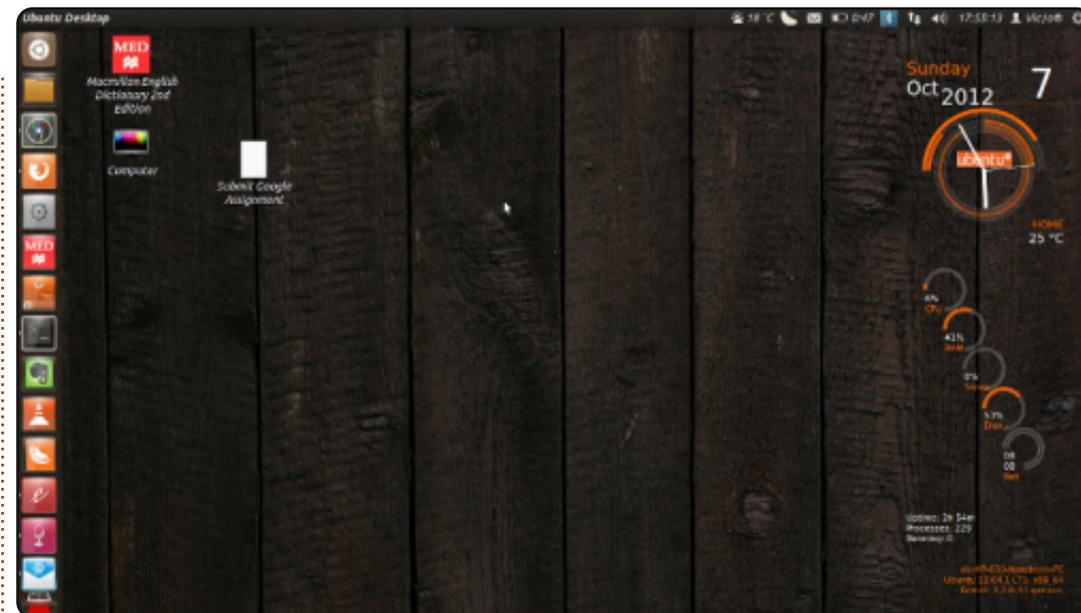
Shey Louis



This is my desktop, running on an Intel D2700MUD mobo (with an Intel Atom D2700 onboard CPU @ 2.13 GHz clock). Because of lack of graphics driver support, it doesn't run Compiz, so I can use only metacity --composite to run this "Cairo-dock + Unity panel" session. It also didn't play videos well, but, thanks to cedarview-graphics-drivers package, it plays videos smoothly now. I like using my native language for the system language. And I use wallpaper from Sability project, an Ubuntu derivative.

CPU: Intel Atom D2700 (onboard)
RAM: 2 GB (SODIMM)
Graphics: Intel GMA 3650
OS: Ubuntu (12.04)

Sabil Rashid



Attached is a screenshot from my Ubuntu 12.04 LTS 64-bit (Precise Pangolin) desktop (Unity) which I've upgraded to 12.04.1. I've dual-booted it with Windows-7, but I occasionally use it since I've gotten accustomed to Ubuntu.

I use Conky lua for system monitoring and providing me with useful information... and I really like its cool look. I've installed The Nitrox UMD theme which has really cool icons. I also use MyUnity, Ubuntu Tweak and Compiz for customization.

System Specs
HP-630-Notebook-PC
Intel core i3 @ 2.4Ghz
4 GB RAM
500 GB HDD sata

Victor Mumo

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

SUDOKU

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

CODE WORD



Codeword and 16x16 Suduko puzzles are copyright, and kindly provided by, **The Puzzle Club** - www.thepuzzleclub.com



HOW TO CONTRIBUTE

FULL CIRCLE NEEDS YOU!

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

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

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



FCM#69

Deadline:
Sunday 06th Jan. 2013.
Release:
Friday 25th Jan. 2013.



Full Circle Team



Editor - Ronnie Tucker
ronnie@fullcirclemagazine.org
Webmaster - Rob Kerfia
admin@fullcirclemagazine.org
Podcast - Les Pounder & Co.
podcast@fullcirclemagazine.org

Editing & Proofreading
Mike Kennedy, Lucas Westermann,
Gord Campbell, Robert Orsino,
Josh Hertel, Bert Jerred

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

Getting Full Circle Magazine:



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



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



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



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



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