The *NEW* PCLinuxOS Magazine

Table Of Contents

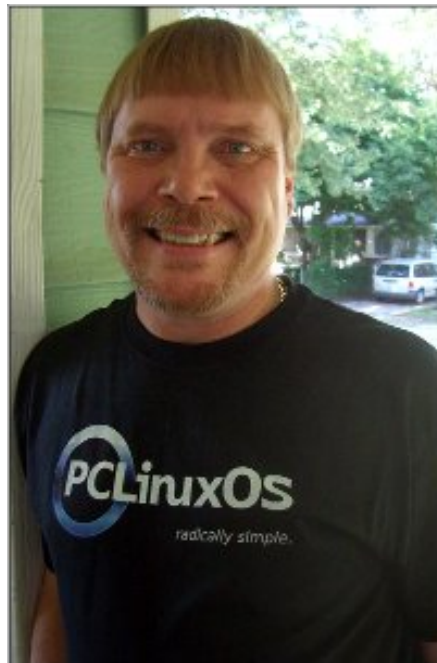# *Welcome From The Chief Editor*

by Paul Arnote (parnote)

The holidays have finally come and gone, the packages have all been unwrapped, the Christmas tree and other holiday decorations are coming down, and a new year is upon us. **Texstar** and the **PCLinuxOS Packaging Crew** are busy putting the new tool chain to good use, working on getting the PCLinuxOS 2010 release nearer to completion. The updates continue to roll out at an amazing pace, with literally hundreds of new and updated packages becoming available every week.

This month's magazine cover features snow covered photos from **ms_meme**. On the inside, the contents are hot enough to melt that snow. To start off, we take a **look back** at notable news of **2009**, and all that has happened with computing, PCLinuxOS, Linux, and FOSS. Peter Kelly continues his **Command Line Interface Intro** column, with part four of his continuing series. Gary L. Ratliff, Sr. continues his march through the alphabet of computer languages, taking a look at **Guile** in his **Computer Languages A to Z** series. **Hootiegibbon** chats with PCLinuxOS Magazine, in the **Behind The Scenes** column, where we continue to get to know the folks who work so diligently behind the scenes to make PCLinuxOS the quality Linux distro that it is. We take a look at the recent e-book explosion, where many of the e-book readers out on the market run on Linux, in the **Book Worms Unite: An E-Book Explosion**. You may be surprised at just how much free e-book content you can find out in the digital realm, if you just look for it.

Meemaw has transcribed **Sproggy's Glass Panel Tutorial** video for the PCLinuxOS Magazine, creating a step-by-step tutorial of how to create a glass panel background for your desktop. Mark Szorady provides another dose of his cartoon, **Double Take**, along with another **Gimp Tip**. We have a couple of new **testimonials** this month, as well. This month's **Gadgets & Gears** column takes a look at the recently opened **Official PCLinuxOS Store** on CafePress.com, where you can buy a wide range of products emblazoned with the PCLinuxOS logo. Patrick G. Horneker takes a look at some of the less-obvious ways (other than VirtualBox and VMWare) to run other operating systems on a PCLinuxOS system, with his **You Can Have Your Cake & Eat It Too** article. Malcolm Ripley reprises his **Archiving Movies With dvd::rip** article, with further tips on how to use dvd::rip to archive your DVD movie collection.

Meanwhile, I hope each and every one of you have had a wonderful holiday season, and I hope that the New Year brings you happiness, prosperity, good health, and bundles of good fortune.

2009: A Look Back

by Paul Arnote (parnote)

With the arrival of 2010, it's a good time to take a look back at all that has gone on, both with PCLinuxOS and Linux, during 2009. Without a doubt, 2009 had seen many events related to both, large and small. Let's take a look, month by month, at all that has transpired. My list is not all-inclusive, and only scratches the surface. However, I've included a lot of "highlights" of development that's flying along with many popular packages. It does seem that development is alive and well with PCLinuxOS, and with FOSS in general.

JANUARY, 2009

• Moblin 2.0 Alpha released by Intel, based on an Ubuntu core.
• OpenOffice 3.0.1 released.
• Mozilla backs open video for the web.
• Linux Foundation kicks off "We're Linux" video contest.
• VirtualBox 2.1.2 released.

FEBRUARY, 2009

• PCLinuxOS N1PTT TR6 Released
• Microsoft files suit against TomTom for patent infringement over the use of VFAT in their GPS systems.
• World of Goo for Linux released.
• VirtualBox 2.1.4 released.
• Android 1.1 released.

MARCH, 2009

• Texstar returns to the helm, after taking a leave to deal with personal business.
• PCLinuxOS 2009.1 Released
• TomTom and Microsoft settle patent dispute, after TomTom counter-sues Microsoft for violating 4 of it's patents with MS Streets & Maps.
• Linux kernel 2.6.29 released.
• Gnome 2.26.0 released.
• KDE 4.2.1 released.
• Qt 4.5 released.
• XFCE 4.6.1 released.

APRIL, 2009

• VirtualBox 2.2.2 released.
• XFCE 4.6.1 released.
• Oracle buys Sun.
• Gnome 2.26.1 released.
• VirtualBox 2.2.0 released.
• KDE 4.2.2 released.

MAY, 2009

• Moblin 2.0 Beta released by Intel. This time, it's based off a Fedora core, instead of an Ubuntu core. Specifically designed for Intel's Atom processor, and for use on netbooks using said processor.
• Gnome 2.26.2 released.
• SourceForge's 2009 Community Choice Awards announced.
• Nokia opens Qt tool kit source code repositories to the public.
• OpenOffice 3.1 released.
• KDE 4.2.3 released.
• Linux achieves 1% market share on desktop.

JUNE, 2009

• PCLinuxOS AppStore announced.
• Addlocale released for PCLinuxOS, allowing users to run PCLinuxOS in 82 different languages.
• GetOpenOffice released for PCLinuxOS, allowing easy, automated installation of OpenOffice in PCLinuxOS.
• Firefox 3.5 released.
• VirtualBox 3.0.0 released.
• Python 3.1 released.
• Qt 4.5.2 released.
• Intel develops a simpler alternative to ACPI for Linux, called Simple Firmware Interface, or SFI. A

patch has been developed and submitted for the 2.6.32 kernel.

• First USB 3.0 drivers emerge for Linux.

• Linux kernel 2.6.30 released.

• Amarok 2.1 released.

• KDE 4.2.4 released.

JULY, 2009

• PCLinuxOS 2009.2 Released

• PCLinuxOS Minime KDE3 2009.1 Released

• The NEW PCLinuxOS Magazine is revived and reborn.

• Paul Arnote named Chief Editor of the magazine.

• Firefox reaches 1 billion downloads.

• PortableApps.com reaches 100 million downloads.

• BBC begins work on a 4-part, 4-hour documentary on the World Wide Web. The working title is "Digital Revolution," and is slated to air on BBC Two sometime in 2010. Included in the documentary will be an interview with Sir Tim Berners-Lee, the inventor of HTML and the World Wide Web.

• VLC 1.0.0 released.

• Gnome 2.26.3 released.

AUGUST, 2009

• PCLXDE PCLinuxOS LXDE Remaster, 2009.4 Released

• Andrew Strick named Assistant Editor of the magazine.

• OpenOffice 3.1.1 released.

• KDE 4.3.1 released.

SEPTEMBER, 2009

• Linux Kernel 2.6.31 released.

• Phoenix, PCLinuxOS XFCE Remaster, 2009.4 Released

• Linus Torvalds, at LinuxCon 2009 in Portland, OR: "Linux 'Bloated and Huge.'"

• Moblin 2.0 released.

• Gnome 2.28 released, with release of Gnome 3.0 announced for September, 2010.

• Open Invention Network (OIN), led by IBM and RedHat, purchase 22 software patents to help protect open source software.

OCTOBER, 2009

• Windows 7 released.

• KDE 4.3.2 released.

• Linux saves Aussie electric grid when Windows computers infected with virus.

• RedHat files amicus brief with U.S. Supreme Court to exclude software from patentability.

• Official KDE 3.5 support comes to an end.

NOVEMBER, 2009

• Google Chrome OS Beta released.

• KDE 4.3.3 released.

• VirtualBox 3.1.0 released.

• KOffice 2.1 released.

• Android 2.0 source code released.

• Microsoft admits to violation of the GPL license by including GPL'd code in Windows 7 USB tool.

DECEMBER, 2009

• KDE 4.3.4 released.

• Linux Kernel 2.6.32 released.

• Google Chrome (Browser) Linux Beta released.

• KDE 4.4 Beta 1 released.

Visit Us On IRC

• Launch your favorite IRC Chat Client software (xchat, pidgin, kopete, etc.)

• Go to freenode.net

• Type "/join #pclinuxos-mag" (without the quotes)

OGG

MP3

What a year this has been on my face there's a grin
I've almost told Windows "Goodbye"
I never had any fun too many scans I did run
We never saw eye to eye

No more hours of up dating for me
At last I will be set free
That's all in the past I knew it couldn't last
I'm glad I'm telling Windows "Goodbye"

HAPPY NEW YEAR

What a year this has been now the fun will begin
I've almost told Windows "Goodbye"
There's a glow on my screen hear those sounds from Kaffeine
I'm glad I'm telling Windows "Go fly"

Now with Linux my life seems to be
A part of the bourgeoisie
And after I boot and log in as root
I don't regret rebelling I don't fret the yelling
I don't sweat telling Windows "Goodbye"

Behind The Scenes: A Chat With Hootiegibbon

by Paul Arnote (parnote)

In the continuing series of articles taking a look at the people "behind the scenes" that help keep PCLinuxOS and the forum running smoothly, we get the chance to learn more about Hootiegibbon, a Global Moderator in the main forum.

Would you introduce yourself, and tell us more about you (marital status, "real" job, full name, where you live, family, etc.)?

My "real" name is Jason Wilkinson, or Jase for short – although I often get called many other, less printable names though. I have been married for 10 years, and have 3 children. The eldest is a soon to be 9 year old boy, my daughter is 4 years old, and my youngest son has just turned 1 in December. We live in a town close to Cannock Chase, in Staffordshire England

I have worked in banking for the last 5-6 years, in a non-IT position. Previous to that, I was in various types of sales environments, from retail to door to door.

How long have you been using computers?

I first 'got' into computing after about 6 months of my wife getting her first computer, which was in 2002. So I was a 'late starter' in computing, and have no formal or informal qualifications, beyond any reputation I may have.

What was your first Linux distro?

My very first Distro was Slax – which was, and still is, an amazing operating system. I used it on the computer mentioned above, and learned some basic Linux stuff to help trouble shoot the hardware. This led to my being a Distro-Hopper of LiveCD's. Then, when I managed to get my own PC (it was a secondhand Dell Celeron 1100), I started to try out

various installed Linux distro's, like SuSe, Mandrake, Bunty, etc.

When did you switch to PCLinuxOS?

I found PCLinuxOS at the 0.91 release, although it was not until the 0.92 release that I had a permanent install of it. At that time, if I recall correctly, thare were 3 ISO's released: one plain, one with fglrx, and – the one that blew me away – the NVidia ISO release. I registered with the forums a couple of months later in January 2006, although it was not until March or April that i made my first post.

What is it about PCLinuxOS that made you stick around?

The ideology and community. Both shine brighter with each release that Texstar serves out.

What "flavors" of PCLinuxOS do you use?

Most of my installs are spun from MiniMe, although I also have installs of all the versions that I use fairly regularly.

Do you run any other OS's? If so, which ones, and what attracts you to them?

Yes, there are 3 other OS's that I admire.

OpenSolaris - its rock solid, the ZFS file system is a

marvel and it has a bearable Gnome Desktop environment.

PC-BSD (I say PC-BSD but it covers all of the flavors) - it has had a usable KDE4 desktop frontend for quite a while, and the BSD backend is scalable and secure. PCBSD has a great benefit in the .pbi apps that install into userspace, while maintaining compatibility with ports.

Finally, there is Haiku, which is an astonishing OS. It's still in alpha stage, although it's got a lot of potential.

I also like to look at the other 'fringe' OS's available like minuetOS and syllable via VBox

You will note that no Linux kernel based OS are named above. This is because I feel that PCLinuxOS is the best, and I tend not to keep installs of other Linux OS's.

You have quite a reputation in the PCLinuxOS Forum for being one of the resident WiFi gurus. What is the best advice you can give users who may be trying to sort out their own WiFi connectivity problems?

Google is your friend, (that and the ability to start over and reinstall if things go real bad). For the most part, problems that you have, other people have (or have had). It's just a matter of finding one of those people. With selective searching, the fix or a hint can be found.

With the exception of sorting out WiFi and network connectivity problems, what would you say is your biggest strength? You biggest weakness?

My strength? Thats a good'un. I really cant think of it. The WiFi thing just came about because of my own mistakes and problems that I faced, so I guess trouble shooting? As for weakness - anything involving packaging, or samba, as I don't have one of those windows boxes some have. There are too many weaknesses to list.

As a global moderator in the PCLinuxOS Forum, what do you find to be the most challenging part of your role there?

Judging at which point to pull a thread or post due to content. Sometimes this is easy, others its all very grey. That, and remaining diplomatic with those that fail to use 'rule one' of the forum - search first. I tend to become a little sarcastic (although I try not to be). I have had for some time now a link to the search function in my Sig.

What do you find to be the most rewarding part of your role as a forum global moderator?

The same as being a member of the forum – getting to know some great people, even if it is 'virtually.' All the global moderator bit means is that if things go astray in a thread, I can get it tidied up.

On average, how much time (per week) would you say you put into your duties as a forum global moderator?

Whenever I am online basically, normally for about 2.5hrs each week night, and for about 5 hrs on Fridays and Saturdays. So, on average, about 22 hourrs a week, sometimes more though.

What bird would best symbolize your personality, and why?

That is a difficult question. Possibly a stuffed and mounted Owl, as my Mrs tells me I am nocturnal and don't move much.

What song would best describe your style in dealing with issues in the forum?

Friends - by LedZeppelin

How did you come up with the name "Hootiegibbon?"

Hootiegibbon is a name my wife came up with when I started with computers and first went online. I don't actually know what a hootiegibbon is, but it has stuck with me for all this time.

What parting words of wisdom would you like to leave us with?

You are unique, just like everyone else.

Screenshot Showcase

Posted by nightowl512, December 11, 2009, running KDE 4.3.4

by Brian Blanchette (redjama)

As one of the most critically acclaimed shooters of all time, Call of Duty 4: Modern Warfare is a prime example of a tough act to follow. Yet, amidst a raging storm of anticipation and expectation, Call of Duty: Modern Warfare 2 has done it. The new campaign is chock-full of intense action and dramatic moments, and though it is more muddled than its predecessor (in more ways than one), it's still an absolute blast. The new Special Ops mode allows you to experience some campaign-inspired thrills with a friend, and it's an engaging challenge to coordinate your maneuvers and tackle the varied objectives.

Last, but not least, the competitive multiplayer that took the online shooter community by storm two years ago is back. Though the addictive action remains the same at its core, there are a host of new elements that make matches more accessible, more strategic, and more rewarding. Unfortunately, these improvements are marred by limited online flexibility that may leave the PC shooter community out in the cold.

System Requirements:

OS: Microsoft Windows XP, Windows Vista and Windows 7 (Windows 95/98/ME/2000 are unsupported) or Wine ver. 1.1.34
Processor: Intel Pentium 4 3.2 GHz or AMD Athlon 64 3200+ processor or better supported Memory: 1 GB RAM
**Graphics:** 256 MB NVIDIA GeForce 6600GT or better or ATI Radeon 1600XT or better

DirectX®: Microsoft DirectX(R) 9.0c
Hard Drive: 12GB of free hard drive space
Sound: 100% DirectX 9.0c compatible sound card
Internet: Broadband connection and service required for Multiplayer Connectivity. Internet Connection required for activation.
Co-op/Multiplayer Hosting: To host Co-op or MP matches, a 2Ghz dual-core or better processor is recommended.

How To Install Modern Warfare 2 Under PCLinuxOS using Wine ver. 1.1.34

You will need a few programs, listed below.

First get winetricks. To download it, open up terminal and type:
"`wget www.kegel.com/wine/winetricks`" (without the quotes).
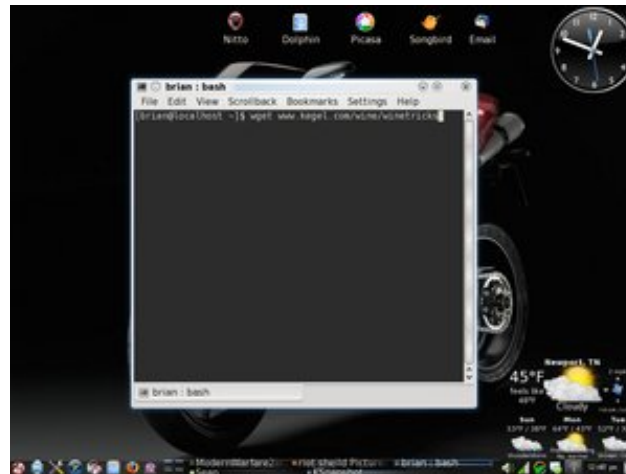
And then, after that's done, open up terminal and type"sh winetricks" (without the quotation marks), nd check the following files to install.


Mark all of them at the same time (hold down the Ctrl key while selecting them), and it will highlight them blue: d3dx9, corefonts, dotnet20, and ie6.
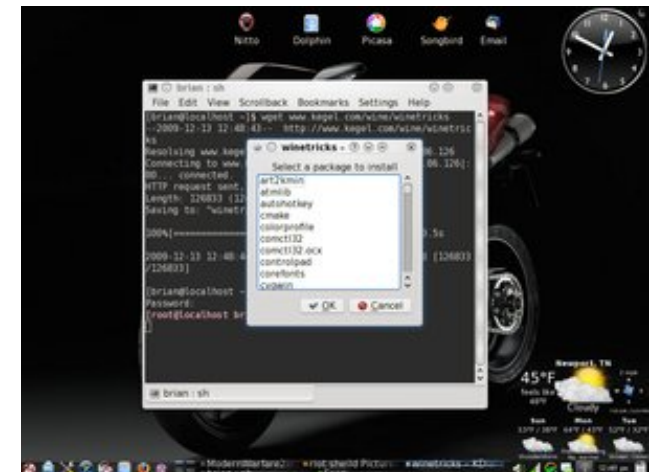
Then click OK. The process will take about 30 minutes to download and install the needed files. Once that is done, you need a program called Steam. You can download Steam from here: storefront.steampowered.com/download/ SteamInstall.msi

Once that is done downloading, right click the .msi file, and select "Open with windows program loader." If you have Wine installed correctly, the installation should proceed normally.
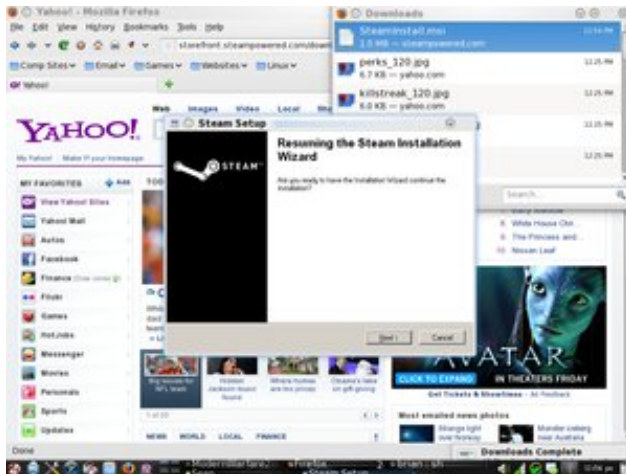
Then, log in with your Steam account. If you don't have an account, make one. After that, login, then move to the game installation part.

What your going to want to do is insert Disc 1 into your optical drive, browse the disc, and launch the file named "Setup.exe." Then, when you are asked to insert Disc 2, don't do anything until you open up a terminal session, and type "wine eject" (without the

quotes). This will cause Disc 1 to be ejected from your optical drive. Now, you can insert Disc 2. Click "OK" on the Steam installer. Once that is complete, it's time to have fun.

Modern Warfare 2 Tips and Tricks

While easy enough to play, all the new customization options in Modern Warfare 2 could throw new soldiers for a loop in terms of strategy. So, whether you opt for the single player or multi-player game, you're going to need a little help. For those that haven't made it through boot camp, or been handpicked by a five-star general, these Modern Warfare 2 tips and tricks should be exactly the training you need.

Slow and Steady Wins the Campaign

On harder difficulties, running and gunning will do nothing but get you killed. And now, with the new dynamic spawning system, the AI enemies are more deadly than ever. So slow it down and take advantage of your surroundings. Look for combustibles, like cars or fuel tanks, to help clear out the bad guys, and always leverage your allies to flush them out. When they die, they come back. You don't have the same luxury.

It's a Riot!

Riot shields attract a lot of attention and can take a lot of damage, and even though you're not shooting, you'll still gain experience points. Perfect for soaking up killstreak reward fire, like sentry guns or helicopters. Just don't travel alone, as all it takes is one stray shot in your exposed back and it's game over.

Be an Effective Counter-Counter Terrorist
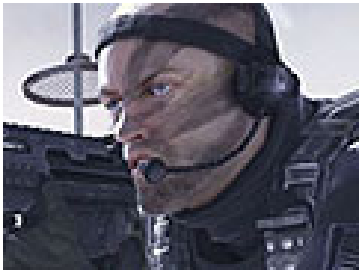
Thwart your opponents' attempts at dominance by stocking up on the right gear and perks. Tactical insertion will let you return behind enemy lines to re-capture points or flank them, and the Scrambler perk will disable their minimap. Remember that all the killstreak rewards are in the game environment, so being that guy who takes out enemy UAVs, choppers and planes can help not just you but the whole team.

Get in the Killstreak Zone

Unless you're a master player, achieving some of those higher killstreak rewards are an uphill challenge. So make sure you pick your chain effectively. Sentry guns are excellent ways to rack up early kills versus Predator missiles, and then you're next best bet is an automated reward like the Pave Low, instead of a controllable AC-130 or helicopter gunner. Care package rewards do not contribute to the streak, so be careful there, and you'll be on your way to nukes in no time!

Perk Up, Soldier!
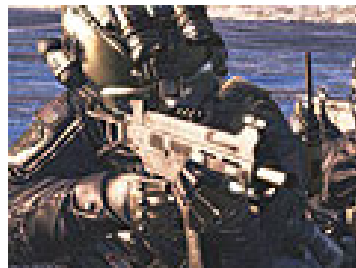
Picking the right mix is the first step towards kickin' tail. If you prefer to fly under the radar, opt for perks like Cold Blooded, which hides you from enemy UAVs and kill streak reward vehicles. Attachments have replaced popular perks Deep Impact and Double Tap, so leverage those for added firepower. And last but never least, the One Man Army perk lets you select from other classes, which can have their own unique perks. Stack One Man Army into those, and you'll have a variety of classes you can keep changing to.

Actual Gamplay Screenshots:

I hope everyone has fun with this. I know I am.

RedJama

Computer Languages A to Z: Guile

by Gary L. Ratliff, Sr. (eronstuc)

Guile stands for: GNU's Ubiquitous Intelligent Language for Extensions. It is an outgrowth of the flame war started in 1994, when Richard M. Stallman wrote an article about why Tcl should not be used for programming. The system comes with a standard install of the various versions of PCLinuxOS. The current version is 1.8.6. I found this on the other computer, which has version 2007, and also on my Gnome 2009-2 install of the system. This acts as an interpreter for the Scheme dialect of Lisp. It also has a library which may be linked into the compile of various languages to allow it to be used to extend the language.

Because Scheme is a dialect of Lisp, I will take this opportunity to mention some Lisp items which directly relate to Common Lisp. I found and installed these systems when I was working on the Elisp segment of this series. The first item is called Lisp in a Box. This provides a Lisp Development system based on Slime. It provides an Emacs editor, which is separate from your emacs editor, and is for use with the "in the box" system. There are several code segments from the book "Practical Common Lisp," written by Peter Seibel. The text of the book is available online at http://www.gigamonkeys.com. The other is the free trial of LispWorks Personal Edition. This is version 5.1.1, and it is installed by completing a form and agreeing to its terms. This is available from http://www.lispworks.com. Once installed you will need to rename one of the libraries. However, these details will be covered once we reach lisp. Or by writing them, they will help you solve any problems you might have with the installation. If you enjoy the fact that lisp is most likely completely different than any other language

with which you have worked, you will enjoy obtaining these items.

Also on the system, I found that the package umb-scheme seems to come with the default install. The system reports that it is using version 3.2, and there is also a message that there was a Fatal Error. However, the system appears to function fine, regardless of this statement. So either of these should be availavle to test the use of the scheme code you wish to try. To start, use a terminal and enter **$guile** to enter into the guile interpreter.

Or enter: **$umb<tab>**

The tab key will expand this to the full name, "umb-scheme." These will begin their respective programs. Now, to exit guile, use: (quit), while to exit the umb-scheme system, press the Control key and then the d key. This was explained on the startup of the umb-scheme program. Offhand, umb stands for University of Massachusetts Boston. Both of these are based on the R5RS standard of the Scheme Lisp Dialect. Now the standard was recently updated, and the R6RS standard was adopted.

Learning to use the Scheme Dialect of Lisp and the new standard.

The newest standard of the dialect was adopted in July, 2009, and texts which are being published now are using this standard. The current versions of the system will be using the R5RS standard. I consider the biggest change to be that the new standard allows the case sensitivity of symbols, which the R5RS does not. This means that if you have a function called foo in your system, you could type in FOO, foo, FoO, etc., and they would all call the

same function. Code based upon the newer standard would allow these to be completely different functions. One of the books I recommend using to learn Scheme is the book from MIT Press, The Scheme Programming Language, 4th Edition, or use The Scheme Programming Language, 3rd Edition. Both are written by R. Kent Dybvig. If you use the 4th Edition, and key in its examples while using either the guile or the umb-scheme system currently installed, you will receive several syntax errors which you will not experience while using the 3rd Edition, which uses the older standard. If you wish to try the latest features in the new standard, there is a free version of Chez Scheme available. This is available in a threaded and non threaded version, for both 32bit and 64bit Intel systems. The threaded version for 32 bit is: pcsv7.9.3-ti3le.tar.gz, while the non-threaded version is: pcsv7.9.3-i3le.tar.gz. They also stated that an rpm version was available. However, trying to install the rpm version yielded that the dependency to libinfo.so.5 was not meet. My idea is that the guile system will soon be upgraded to use the new standard anyways.

There are some interesting tools available. For example, let's explore the familiar factorial function. This is defined as 1 for factorial 0, and by multiplying all the integers from 1 to n for successive numbers. So the screen shot will show how to define the function in Scheme, and then show how to trace this function.

This is done by entering the following into guile:

```
(define fact
  (lambda (x)
    (if (zero? x)
      1
```

```
        (* x (fact (1- x))))))

(fact 5)
(trace fact)
(fact 5)
(fact 10)
```


This shows the steps which the process goes through to arrive at the answer. Also note that 1- is itself a separate function. If you wish to learn what

the item is, you may enter it at the keyboard without any parentheses surrounding it. For example, + will yield <primitive-generic +>, lambda will yeild <primitive-builtin-macro! Lambda>, and 1- will yield <primitive-procedure 1->

Now one of the functions of the supplied library is to allow the ease of adding full functionality of different languages to be added to an existing language. Thus a compiled C++ or C program could have the Scheme interpreter provided by the libguile.so library to be added to the program. An example of how easy this is, is provided in the tutorial and sited below:

```
    /* test the new libgh.a (Guile High-level
library) with a trivial
    program */

#include <stdio.h>

#include <guile/gh.h>

void main_prog(int argc, char *argv[]);

main(int argc, char *argv[])
{
  gh_enter(argc, argv, main_prog);
}

void main_prog(int argc, char *argv[])
{
  int done;
  char input_str[200];

  gh_eval_str("(display \"hello Guile\")");
  gh_eval_str("(newline)");
```

```
    /* for fun, evaluate some simple Scheme
expressions here */
    gh_eval_str("(define (square x) (* x x))");
    gh_eval_str("(define (fact n) (if (= n 1) 1 (* n
(fact (- n 1)))))");
    gh_eval_str("(square 9)");

    /* now sit in a Scheme eval loop: I input the
expressions, have
      Guile evaluate them, and then get another
expression. */
    done = 0;
    fputs("learn0> ", stdout);
    while (fgets(input_str, 199, stdin) != NULL) {
      gh_eval_str(input_str);
      fputs("\nlearn0> ", stdout);
    }

    exit(0);
    }
```

This program should be named learn0.c, and it is compiled with these two passes of the gcc compiler:

$gcc -c learn0.c -o learn0.o
$gcc -o learn0 learn0.o -lguile -lm.

The -c flag is used to tell the compiler to stop after the object file has been created, instead of continuing till the executable file is created. The -o flag is used to tell the compiler to name the file produced learn0.o. In the second run of the compiler, this file is compiled, and the output named learn0, and is linked to two libraries the guile library and the math library. Once this has been done, you will find that ./learn0 will start a file which will function exactly as if you had entered the guile command to start guile itself.

Here is another excellent source for learning the Scheme dialect of Lisp. This book was developed at MIT Press, and is used in an introductory course in computer programming at many universities throughout the country. If you ever watched the movie Hackers, you may recall where the computer jock mentioned a litany of different named colored computer manuals. Well, this is known as the Purple Book.
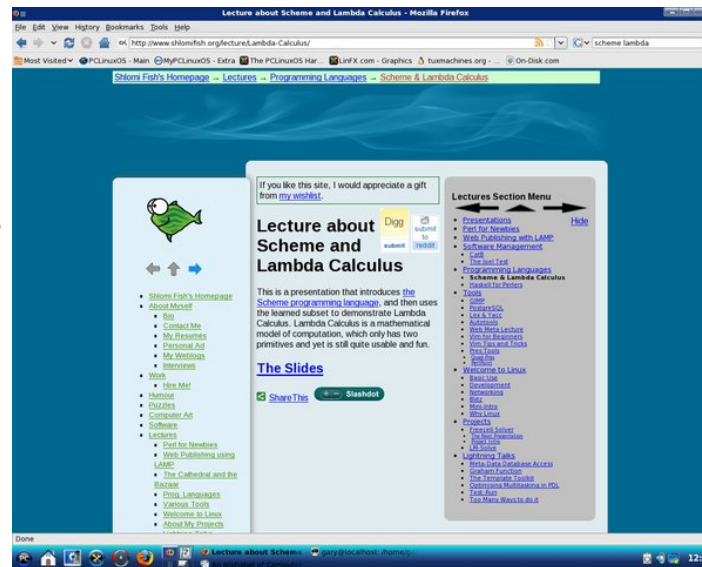
By visiting this site and clicking on the slides link, you will be taken to a complete lecture that will introduce all the features of the scheme dialect of lisp. You do have quite a bit of power, as the top level definitions can easily be changed. For example: You learned that 2 + 3 = 5 and that this would be expressed in scheme or guile as: (+ 2 3.) However it is possible to enter (define + -) and now

when you ask for (+ 2 3) the answer will be given as -1!! This would be the answer if you had entered (- 2 3) Also, if you now used the system to do your math homework, you would most likely flunk if you were performing a drill on addition.

The only remedy is to use (exit), or Control D, to exit the system and restart, which will kill your idea of the new addition primitive.

One of the major uses for guile is to allow programs to become extensible. However, the examples sited in the documentation for this do not compile correctly. However, this should be enough to wet your appetite for learning more about this system. The very latest on the language blogs was that Scheme would be split into two different languages. Only time will tell on this move.

Screenshot Showcase

Posted by Linuxera, December 17, 2009, running e-17

PCLinuxOS
Radically Simple

Double Take
by Mark Szorady

I'M RESOLVING TO STOP WEARING FUNNY HATS.

I'M RESOLVING TO START WEARING FUNNY HATS.

HAPPY NEW YEAR!

©2010 Mark Szorady. Distributed by georgetoon.com

Find at least seven differences between cartoons.

Answers on Page 30

Mark's Quick Gimp Tip

Happy New Year! It's a new year and a new beginning. And, as tradition dictates, we make resolutions to meet new goals and better our lives. Most of these resolutions deal with our personal well being. We resolve to lose weight, exercise more, cut out fatty foods, eat healthier foods, etc.

And, like me, I'm sure you've compiled a list of resolutions that you hope to keep in the coming year.

 But please add one resolution to this list (especially if you're reading PCLinuxOS magazine for the first time). Resolve to give Linux and The Gimp a try.

The Gimp is the best, most complete, and most powerful graphics program on the face of the planet that comes absolutely free of charge!

That's right! The Gimp is absolutely 100 percent free!

How do you get The Gimp? You can download The Gimp and install it to your computer system by visiting their web site at www.gimp.org.

 But the easiest way to get The Gimp is to simply install PCLinuxOS to your computer!

(Besides, it works best in Linux!) You'll not only get The Gimp already installed and ready to run, but you'll also have what I consider to be the best distribution of the Linux OS, PCLinuxOS! Start the new year right! Get PCLinuxOS and get The Gimp!

-Mark Szorady is a nationally syndicated cartoonist. His work is distributed by georgetoon.com. Email Mark at georgetoon@gmail.com.

You Can Have Your Cake & Eat It Too
(or running other operating systems on PCLinuxOS)

by Patrick G. Horneker (phorneker)

Last month, I discussed running PCLinuxOS on Slackware. This is one example of the use of virtual machines to run multiple operating systems on a single physical machine. (I mention this as PCLinuxOS also runs on Intel-based Macintoshes.)

While you can run PCLinuxOS on Slackware, OpenSUSE, Fedora, Ubuntu, or any other Linux you wish, You can also run any Linux distribution on PCLinuxOS, using VirtualBox and/or QEMU.

I mentioned VirtualBox as it does the best job of utilizing network and hardware resources between Linux distributions.

QEMU

QEMU is a good virtual machine for running older Linux distributions, Windows 9x/ME, PC/MS-DOS (with and without Windows 3.x), FreeDOS with OpenGEM, MikeOS (created by Mike Saunders of Linux Format fame) and other operating systems that run on Intel x86 processors (as well as the AMD equivalents).

As an Intel emulator, QEMU emulates a x86 machine with a Cirrus 34xx video chipset for video, your choice of an Adlib, Sound Blaster 16, or ESS1860 for audio, NE2000 Ethernet emulation, with TUN capability for networking, virtual serial and parallel ports, which can be redirected to disk files (for later printing), or to actual devices such as a USB Modem (for serial port access) or a USB printer (for the parallel port).

On my laptop, I have a copy of IBM PC-DOS 6.3 with DJGPP installed (for DOS development) running inside a QEMU virtual machine. (I had that copy of PC-DOS stored in the closet for about ten years.)

Hercules

This software package emulates a IBM System 370, System 390, or a z-Series mainframe. You can use this to see how Linux runs on an IBM mainframe. However, to understand how Hercules works, you need to understand how to install and maintain these classic IBM mainframes.

You will need to have the Hercules package installed, a 3270 terminal emulator, and some packages containing the operating system, and supporting software packages. Fortunately, the Hercules website has links to sites for downloading of these packages.

Android

This is Google's newest platform for smart phone development. You start by

downloading the basic SDK to an empty directory on your hard drive. While you are still connected, unpack the archive and run ./tools/android to launch the system manager.

You can use any development environment such as Eclipse, KDevelop or Anjuta, or even use a text editor such as Emacs for Android development.

Within the system manager, you can install development packages, including API, toolkits, and
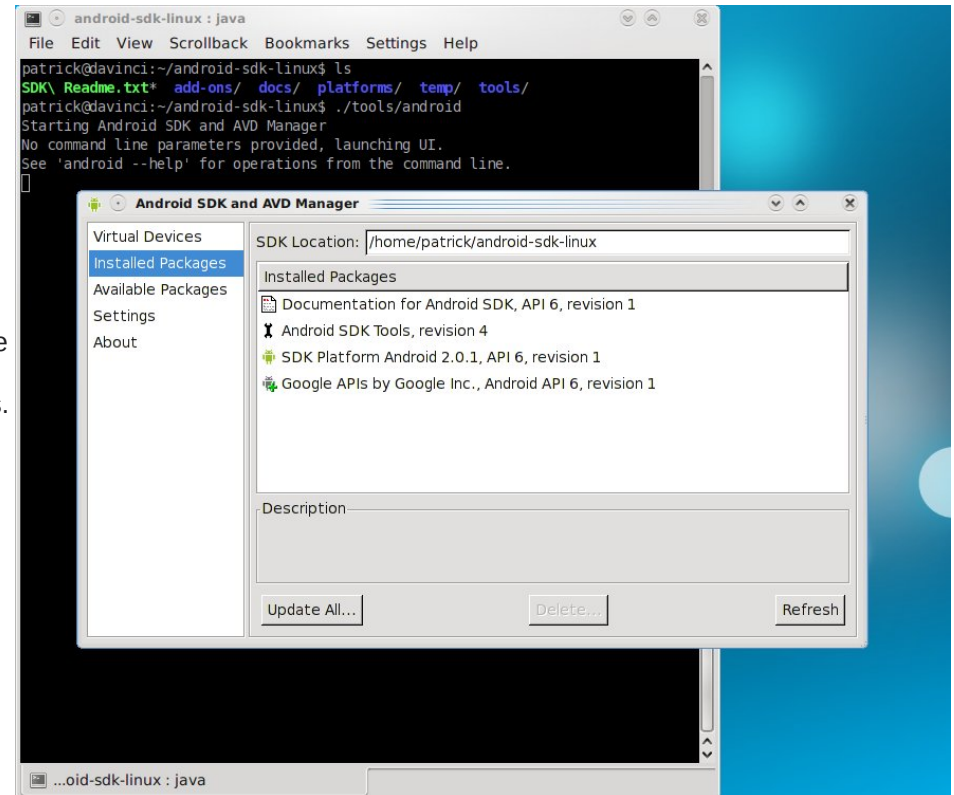
documentation. You can also create virtual Android devices to test your applications you create before deploying them to smart phones.

Most of the development for Android takes place on the command line, or through your favorite development environment.

In this example, I have the latest versions of the Android SDK tools, the SDK Platform, and the application APIs from Google installed.

I also have a virtual smart phone with the latest Android (version 2.0.1 as of this writing) installed.

It is here where I have a virtual Motorola Droid phone (as advertised by Verizon Wireless) installed.

Note: To create virtual Android phones and to download the installed packages, you must be connected to the Internet as the tools are downloaded directly from Google servers.

Typing `./tools/android --help` in the terminal window provides a list of command line options:

Usage:
```
 android [global options] action
[action options]
```

```
Global options:
  -v --verbose  Verbose mode: errors,
warnings and informational messages
are printed.

  -h --help     This help.
```

```
-s --silent    Silent mode: only errors
are printed out.
```

```
Valid actions are composed of a verb
and an optional direct object:
-   list                : Lists existing
targets or virtual devices.
-   list avd          : Lists existing
Android Virtual Devices.
-   list target     : Lists existing
targets.
- create avd        : Creates a new
Android Virtual Device.
-   move avd          : Moves or
renames an Android Virtual Device.
- delete avd        : Deletes an
Android Virtual Device.
- update avd        : Updates an
Android Virtual Device to match the
folders of a new SDK.

- create project     : Creates a new
Android Project.

- update project     : Updates an
Android Project (must have an
AndroidManifest.xml).

- create test-project: Creates a new
Android Test Project.

- update test-project: Updates an
Android Test Project (must have an
AndroidManifest.xml).

- update adb         : Updates adb to
support the USB devices declared in
the SDK add-ons.
```

```
- update sdk          : Updates the SDK
by suggesting new platforms to install
if available.

Action "list ":
  Lists existing targets or virtual
devices.
Options:

  No options

Action "list avd":
  Lists existing Android Virtual
Devices.
Options:

  No options

Action "list target":
  Lists existing targets.
Options:
  No options

Action "create avd":
  Creates a new Android Virtual
Device.
Options:

  -c --sdcard   Path to a shared SD
card image, or size of a new sdcard
for the new AVD

  -t --target   Target id of the new
AVD [required]

  -n --name       Name of the new AVD
[required]
```

```
  -p --path      Location path of the
directory where the new AVD will be
created
  -f --force     Force creation
(override an existing AVD)

  -s --skin      Skin of the new AVD

Action "move avd":
  Moves or renames an Android Virtual
Device.
Options:

  -p --path      New location path of
the directory where to move the AVD
  -n --name      Name of the AVD to
move or rename [required]
  -r --rename    New name of the AVD to
rename

Action "delete avd":
  Deletes an Android Virtual Device.
Options:
  -n --name      Name of the AVD to
delete [required]

Action "update avd":
  Updates an Android Virtual Device to
match the folders of a new SDK.
Options:

  -n --name      Name of the AVD to
update [required]

Action "create project":
  Creates a new Android Project.
Options:
```

```
  -n --name      Project name
  -t --target    Target id of the new
project [required]
  -p --path      Location path of new
project [required]
  -k --package   Package name
[required]
  -a --activity Activity name
[required]


Action "update project":
  Updates an Android Project (must
have an AndroidManifest.xml).
Options:
  -p --path      Location path of the
project [required]
  -n --name      Project name
  -t --target    Target id to set for
the project
  -s --subprojects Also update any
projects in sub-folders, such as test
projects.


Action "create test-project":
  Creates a new Android Test Project.
Options:
  -p --path      Location path of new
project [required]
  -m --main      Location path of the
project to test, relative to the new
project [required]
  -n --name      Project name


Action "update test-project":
  Updates an Android Test Project
(must have an AndroidManifest.xml).
Options:
```

```
  -m --main      Location path of the
project to test, relative to the new
project [required]
  -p --path      Location path of the
project [required]


Action "update adb":
  Updates adb to support the USB
devices declared in the SDK add-ons.
Options:
  No options


Action "update sdk":
  Updates the SDK by suggesting new
platforms to install if available.
Options:
  No options
```

Note that these options can also be accessed through the user interface that launches when **--help** is not provided on the command line.

iPhone Development

Unfortunately, the same could not be said for development for the iPhone. Just learning Objective-C is not enough. According to the Apple website, you need to have registered your iPhone (that is, you must own an iPhone, not just download the simulator) to develop applications for the iPhone. Registration allows you to obtain a user account for which you login to download the SDK. Also, you must be running Mac OS-X to download the SDK.

However, Jonathan Zdziarski has written a book called iPhone Open Application Development that

allows you to work around the SDK restrictions. The link on this page allows you to order the book from Amazon.com, and is provided for your convenience.

The Future is Mobile

While we have emulators such as VirtualBox and QEMU to run other operating systems, the future of computing is clearly moving in the direction of mobile devices. While PCLinuxOS can run on the current generation of netbooks, and more traditional laptops (of which I recommend the latter for everyday use), it is the smart phone where there is a substantial growth in applications development.

Besides Android (which is itself a Linux derivative, and requires Linux for development) and iPhone (as you would guess is a Mac OS-X derivative), we have the next generation Palm devices, which do not resemble the PalmOS devices of yesteryear as far as functionality and usability.

The good news here is that the current generation of Palm devices are Linux based! Hence, the older PalmOS tools, such as KPilot and Gnome-Pilot, are not needed. Instead, you can use the Bluetooth capabilities to sync these devices.

Screenshot Showcase

Posted by Crow, December 4, 2009, runnin KDE 4.3.4

Gadgets & Gear: Official PCLinuxOS Store Opens

by Paul Arnote (parnote)

This month, we definitely take a look at gear. PCLinuxOS gear, to be specific. Just in time for the holiday season, the Official PCLinuxOS Store opened on CafePress.com. Now, you can support your favorite Linux distro by purchasing, wearing, and using items emblazoned with the PCLinuxOS logo. In the process, you'll be helping to spread the word to the "un-informed" about the best OS for PC's – anywhere!

There are a wide range of products available. Here, PCLinuxOS user and e-17 theme master Agust shows off his PCLinuxOS Golf Shirt ($19.99) and PCLinuxOS Trucker Hat ($14.99).

Various styles of t-shirts are available, ranging in price from $12.99 to $23.99, and in various colors. You can even purchase t-shirts for your littlest PCLinuxOS users. Shirts for toddlers and infants sell for $10.99, while the shirts for the kids sell for $15.99.

Three styles of hats are available, ranging in price from $14.99 to $18.99. You can buy a thong ($9.99) for that special PCLinuxOS user in your life, or boxer shorts ($14.99) printed with the PCLinuxOS logo.

If clothing and hats aren't your thing, then you can purchase other, more utilitarian items. You can purchase a PCLinuxOS wall clock for only $13.99. There is also a choice of two different sized coffee mugs (Neal approved), for $14.99 and $15.99. You can choose from a tile coaster with the PCLinuxOS logo ($6.50), or a PCLinuxOS mouse pad ($12.99). There are also 3" x 5" stickers that you can place on the side of your computer case, or on your car's bumper, that are printed with the PCLinuxOS logo ($3.99). You have a choice of either an oval or rectangular sticker.

All profits from the sale of the items will go to help support continued development of PCLinuxOS.
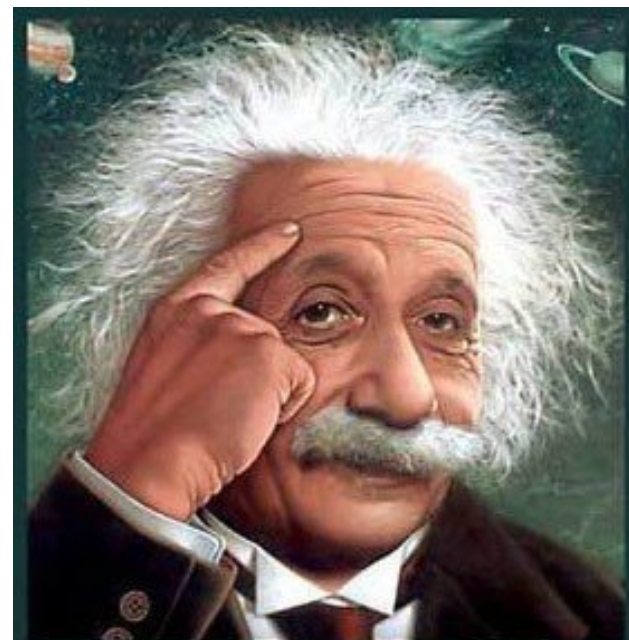
It's easier than $e=mc^2$
It's elemental
It's light years ahead
It's a wise choice
It's Radically Simple
It's ...

Book Worms Unite: An E-Book Explosion

by Paul Arnote (parnote)

It's a cold, wintery day. The snow is piling up outside. A log, or two, blazes in the fireplace. Or, the sound of the pounding surf beats against the shore, as the sun beats down with equal vigor on your bathing-suit-clad body on the beach. In either case, for many people, it represents the perfect time to curl up with a good book.

Today, curling up with a good book can take on a digital flavor. The market has seen the introduction of quite a few e-book readers, and many of them are running Linux. An added benefit is that many books – as many as 200 or more (depending on the amount of memory) – can be loaded onto a digital e-book reader. Literally, that means a whole bookcase full of books can comfortably fit into one tiny, handheld device.

Virtually all the e-book readers on the market today have high contrast black-on-white screens, most capable of displaying 16 levels of gray scale. This makes most of the e-book readers capable of being read, even in bright sunlight. Designed to literally sip power, they can run for hours on a single charge of their built-in rechargeable batteries. Some can even display JPG photos, and can double as very capable MP3 media players.
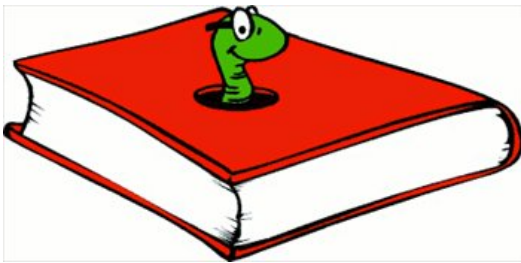
The Commercial Players

Amazon Kindle 2

Probably one of the better known e-book readers on the market, the Amazon Kindle 2 has recently experienced a price drop (from $299 to $259), presumably to better compete with the increased offering of e-book readers coming onto the market.

As one of the earlier e-book readers, the Amazon Kindle 2 is a second generation e-book reader, and is packed with features, the least of which is the fact that it uses Linux as its operating system. The Amazon Kindle 2 features a 6" (15.24 cm) high contrast black-on-white screen that can be read, even in direct sunlight. E-books can be purchased on Amazon.com, as well as other retailers. Purchases from Amazon.com are wirelessly transmitted to your Kindle 2, via the built in 3G wireless (over the AT&T wireless network). If you are not in an area serviced by AT&T's 3G coverage, you can, alternatively, download the e-book content to your PC, and transfer it to the Kindle 2 via a USB 2.0 connection. There are no wireless charges, since they are built-in to the price of the Kindle 2 content purchased from Amazon.com.

Charging the Kindle 2's built-in rechargeable batteries can be

accomplished by plugging it into a USB 2.0 connection on your computer, or by using the supplied AC charger. With the wireless turned on, the Kindle is reported to have a battery life of one week, and two weeks with the wireless turned off. Of course, the battery life will vary, depending on wireless signal strength and availability.

There is also a built-in dictionary, to allow you to look up English words without interruption of your reading. Using experimental text-to-speech technology, the Kindle 2 will also read your documents to you. Additionally, the Kindle 2 supports PDF files, in addition to the Kindle 2's native format (AZW), TXT files, Audible (Audible Enhanced (AA, AAX)), MP3, unprotected MOBI, PRC natively, and HTML, DOC, JPEG, GIF, PNG, BMP through conversion. The screen allows the user to zoom the displayed content, as well as changing the size of the text displayed on the Kindle 2 screen.

The Kindle 2 comes with 2 GB of internal memory, of which 1.4 GB is available to store user content. There currently are no memory expansion options. This model of the Kindle 2 is reported to hold up to 1,500 titles. There is a larger model, the Kindle DX, which is reported to hold 3,500 titles, as well as sporting a 9" screen. The larger model sells for $489.
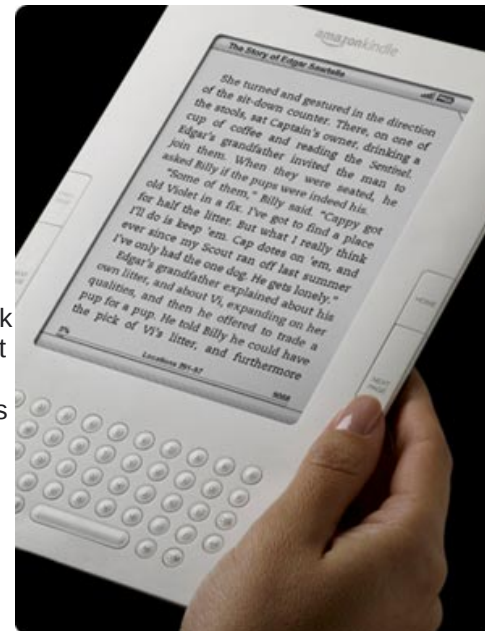
On the downside, there are reports of problems with Amazon's Kindle service. Some of the problems are incompatibilities between the formats used with the original Kindle and the Kindle 2 (Amazon.com has to convert your Kindle 1 files to Kindle 2 compatible files, and send them to you), Amazon.com not retaining your purchases so you can re-download them, and some of your purchases "disappearing" because Amazon.com has "pulled the title." The last two problems can be circumvented by downloading your purchases to your computer, and loading them onto your Kindle 2 via the USB 2.0 cable. This way, YOU have a backup, which can be burned to a CD or DVD for later re-installation.

Sony Reader

Depending on which model of the Sony Reader you purchase, it can be either one of the most expensive, or the least expensive of the e-book readers on the market. It is unclear from the Sony website which OS powers the Sony reader, but despite Sony's lack of support for Linux in the past, the Sony Reader's OS is Linux based.

At the lower end of the spectrum is the Sony Reader Pocket Edition ($199.99). One of the smaller e-book readers, it has a 5" screen. But it is also the most "stripped down" of any of the e-book readers reviewed here. The Pocket Edition model comes with 512 MB of built-in memory, of which 440 MB is available for user content. There is no option, on this model, for memory expansion. It is capable of reading DRM-protected content (DRM-PDF, DRM-ePub, and DRM-BBeB), as well as unprotected

ePub, BBeB Book, PDF, TXT, RTF, Micrsoft® Word (the latter after conversion). The built-in rechargeable battery will last for approximately 7,000 page turns per charge. The Pocket Reader content is loaded via the supplied USB 2.0 cable, from your computer. The screen is also only capable of displaying 8 levels of gray scale.

In the middle of the offerings by Sony, is the Sony Reader Touch Edition ($299.99). Like it's smaller cousin, it's 6" screen is only capable of displaying 8 levels of gray scale. Besides reading all of the same file formats as the Pocket Reader, it also plays unprotected (read that DRM free) MP3 and AAC audio files, and displays JPEG, PNG, GIF, and BMP files. It also has 512 MB of installed memory, 380 MB of which is available for content. Unlike it's smaller cousin, it does allow for memory expansion, utilizing either Memory Stick or SD card media up to 16 GB. Additionally, it utilizes a touch screen, where you swipe your finger across the screen to flip "pages" of the e-book.

At the high end is the Sony Reader Daily Edition ($399.99). It is also Sony's newest e-book reader in their lineup. The Daily Edition has all the features of it's two, lower priced cousins, plus a 7" touch screen capable of displaying 16 shades of gray scale. It also has free 3G connectivity, so you can connect to the Sony Reader Store to purchase more e-books. The claims are that you can connect to the Sony Reader Store from most places within the U.S. Also, with the built-in 3G wireless, you can download your daily newspaper, or any other periodical, to your Daily Edition. It's built-in rechargeable lithium-ion battery is reported to last for up to a week with wireless turned on, and for up to two and a half weeks with the wireless turned off. It comes with 1.6 GB of built-in memory, which can be expanded, either via Memory Stick or SD card media, up to 32 GB. Finally, you can read your digital e-book material with the screen in either landscape or portrait orientation. When you change to the landscape orientation (depending on the content you are reading), the Daily Edition displays two pages at a time. There is also a built-in dictionary, to allow your to look up words you may come across, without leaving your e-book content. Definitely a lot of features, but then they do come at the highest price of any of the e-book readers.

COOL-ER eReader

The COOL-ER eReader ($249) is a fairly new upstart, created by Interead.com, Ltd. out of the UK in 2008. The COOL-ER eReader runs Linux as it's operating system, on a Samsung® S3C2440 ARM 400MHz processor. The operating system runs on the internal memory space of 128 KB, and provides

COOL-ER™ is a light, small, and elegant electronic book reader
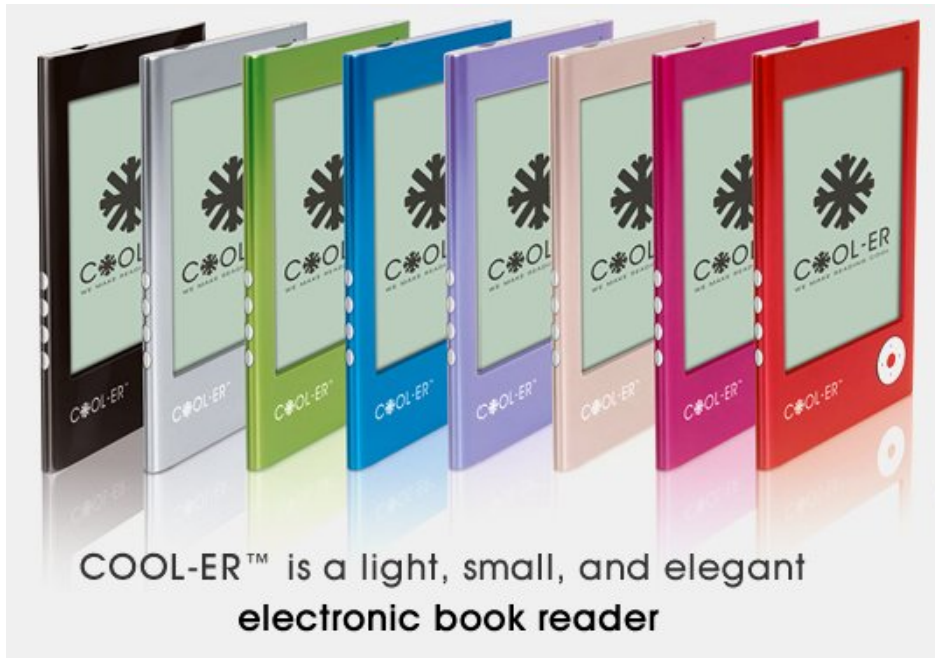
a separate 1 GB of memory for storage. That storage is expandable, via SD card media, up to an additional 4 GB.

It's built-in rechargeable lithium polymer battery will hold a charge for up to 8,000 page reads. It also sports a 6" high contrast screen, that mimics ink on paper, and is visible even in bright sunlight. The screen is capable of displaying 8 shades of gray scale. It also allows for easy switching between portrait and landscape viewing orientation.

Your e-book content is loaded onto the unit via a supplied USB cable, plugged into any available USB port. The COOL-ER eReader supports PDF, EPUB,

FB2, RTF, TXT, HTML, PRC, JPG AND MP3 unprotected content, as well as PDF and EPUB DRM content.

Additionally, the COOL-ER eReader can display your e-book content in 8 different languages: English, German, French, Spanish, Portuguese, Russian, Traditional Chinese, and Simplified Chinese.

The COOL-ER eReader has it's own e-book store, called COOLER eBOOKS. There, you can purchase DRM-protected content, or access over 1,000,000 free titles from Google. Prices for protected content range from around $2, to just under $50 for the entire Twilight series.

Foxit eSlick Reader

The Foxit eSlick Reader ($259) is another e-book reader that runs using embedded Linux as it's operating system. The embedded Linux OS is GPL'd, but the Foxit Reader software on the eSlick Reader is not, and uses no GPL code. And, in case you are wondering, this is the same Foxit that makes the Foxit PDF Reader software that was recently added to the PCLinuxOS repository.

The eSlick Reader has a 6" screen, that unfortunately, has the lowest gray scale display capability, only capable of displaying 4 shades of gray scale. It is capable of utilizing content in the PDF, EPUB, PDB, TXT, HTML, GIF, BMP, JPEG, and PNG formats, as well as playing unprotected MP3 files.

The Foxit eSlick Reader web site claims that the built-in rechargeable lithium ion battery allows for reading 8,000 pages before needing a recharge. However, there are quite a few reports in the Foxit forums about poor battery life, with some reports of only being able to read between 80 and 90 pages before the battery charge is exhausted.

There are also reports of poor performance from the

eSlick Reader in it's capability to display files in the HTML format. On the other hand, much praise has been showered on the eSlick Reader's ability to display PDF files.

There are ongoing attempts by the Foxit team to address the shortcomings of the eSlick Reader. Currently, there is a firmware upgrade, but even the initial reports on the Foxit forums indicate that it is

coming up short in addressing the shortcomings of the device.

The eSlick Reader does come with a lot of accessories in the package. These include the AC charger, a USB cable, a 2 GB SD card, earphones, and earphone adapter (that allows use of earphones with a 3.5 mm earphone plug with the eSlick Reader's 2.5 mm earphone jack), a CD that includes the User Manual, Foxit Reader Pro Pack, PDF Creator, PDF Editor (Trial) & PDF Page Organizer Pro(Trial). Unfortunately, all of the included software is Windows-based.

As much as I really want to like this e-book reader, I cannot recommend it. You can get any of the other e-book readers listed here, that have more features and better specifications, and spend less money. Hopefully, if the Foxit team will get their act together and really fix the issues that this e-book reader has, they will have a good and viable device to compete in the e-book reader market. It has the potential to be a contender, but not in it's current state.

txtr Reader

Produced by txtr GmbH, a Berlin-based company that started up in 2008, the txtr (299€) is currently available for "pre-order" from their web site. So while actual product has not yet rolled off of their production line, the specs for this e-book reader do look good – at least on paper. Let's hope that the implementation is as good.

It is primarily aimed at the German market, but there are nearly 3,000 English language titles in their online e-book store, txtr.

The txtr features a 6", high contrast screen. Details on how many levels of gray scale can be displayed was not available. It runs a Linux kernel, on the Freescale ARM11 CPU, with a 532MHz clock speed. It is slated to have 128 MB SDRAM of built-in memory, and will come with an 8 GB MicroSD memory card.

The e-book reader will also have a touch screen, and page flips will occur with a slide of a finger. The txtr also has a built-in lithium-ion rechargeable battery, but no details on battery life were available. The battery is recharged via the USB cable. It has native support for PDF and EPUB e-book formats. It will also have wireless support, via SIM-enabled EDGE/GPRS, and content can also be loaded, via the USB 2.0 cable connection.

Barnes & Noble Nook

The Barnes & Noble Nook ($259) is a late entry in this article. I didn't even know it existed, until I stumbled upon it five days before Christmas, after I already had the article finished. For me, the Nook is like a "best kept secret." It is a well-thought-out e-book reader, packed with features. And having the backing of Barnes & Noble, one of the most well-known bookstores in the world, makes this a perfect marriage.

To start off with, the Nook is the first e-book reader to run on Google's Android. It features a 6" high-contrast, glare-free screen that can be read, even in bright or direct sunlight. The main reader screen is capable of displaying 16 shades of gray scale. Additionally, the Nook has a 3.5 inch full-color touch navigation screen below the main reader screen, to select the content that you want to read. This full-color touch navigation screen also doubles as a touch screen keyboard, which you can use to type in your search criteria for content.

The Nook has built-in wireless connectivity, with both free AT&T 3G wireless AND 802.11 b/g WiFi. Additionally, the Nook can connect to the free WiFi connection at Barnes & Noble bookstores, where you can gain access to exclusive content for the Nook. The Nook's software can also be updated via the wireless connection.

The 2 GB of built-in memory allows for the storage of up to 1,500 e-books. But the Nook also allows you to use microSD and microSDHC memory cards, for virtually unlimited storage potential. You can also

load your e-book content onto the Nook, via the USB cable, connected to your computer via an empty USB 2.0 port, which can also be used to charge your Nook. The built-in rechargeable lithium-ion battery allows for reading the Nook for 10 days (with wireless turned off) before needing a recharge. And, should you need a new battery, the Nook's battery can be replaced by the end user.

Barnes & Noble claims to have more content available for the Nook, than Amazon has available for it's Kindle 2. Both support wireless delivery of eNewspapers, eMagazines, and other content. The Nook can read DRM (protected) content, in either the PDF or EPUB formats. It can also read content in the unprotected PDF, EPUB, PDB, MP3, GIF, JPG, PNG, and BMP formats. It will not, however,

read content in the DOC, LIT, TXT, AMZ (Amazon), or LRZ/LRX (Sony) e-book formats.

Just as with Amazon's Kindle 2, the Nook has a built-in dictionary to allow you to look up words without interrupting your reading session.

The Nook also allows the user to set up screen savers, personalized with the user's personal photos in any of the supported graphic file formats. One other unique feature of the Nook is that you can lend your digital content – one time – to a friend, much as you would loan a book to a friend, had you purchased the traditional paper book content. While your e-book content is on "loan" to a friend, you do not have access to that content. Unlike with traditional paper books, which you may loan to a friend and never get back (or that you may get back looking like it had been in the trench fighting of WWI), the loaned digital content automatically reverts back to you after 14 days, regardless. It is somewhat discouraging that the digital content may be loaned out only one time, unlike a traditional paper book, which may be loaned out many, many times over it's life. While it's understandable that e-book content sellers want to profit from the sale of that said content, the restriction makes no sense, since the original purchaser has no access to the content while it is on loan – just as with the traditional paper version. Plus, the limit appears, to me anyways, to be rather restrictive and severely limits the sharing and flow of information.

Overall, the Nook represents the best-of-the-best of the e-book readers. Is it a perfect solution? Of course not. Like anything else, it has it's shortcomings (the one-time-only lending, lack of support for TXT, DOC, and HTML files). But, the Nook provides the most bang for the buck, delivers the most features that e-book users want, and delivers a feature packed e-book reader at a price that is very competitive. This is why I give it the PCLinuxOS Magazine Editor's Choice distinction. Just don't be in a hurry – the Barnes & Noble website lists the anticipated shipping date as February 1st.

Do I Really Need A Reader?

That is a great question. Undoubtedly, e-book readers do have some advantages. They are small, and for the most part, they literally sip the power from their rechargeable batteries very slowly. And, they don't require you to fire up your computer to read.

But if you are a "road warrior," and rarely travel without your laptop or netbook, you can also join in on the e-book revolution. All you need is a reader. If you wish to read the DRM-protected content, you will have to use software from a commercial software vendor, such as Adobe. But if you love to read classic literature, you will need nothing more than a PDF reader, or a text editor. Certainly, the PDF reader is a bit easier, but plain text files do work well, too.

In PCLinuxOS, you have quite a few choices for reading either file format. For plain text files, you can view them in Kate, KWrite, Leafpad, Mousepad, AbiWord, OpenOffice, vi, or any other text editor of your choice. For PDF files, you can view them in Evince, Kpdf, Okular (on KDE 4), Adobe Acrobat Reader, or any other document viewer that supports PDF files.

If you have a preference for PDF files, you even have tools at your disposal to convert your plain text e-book content into PDF format. Probably the easiest is to load your plain text e-book content into OpenOffice Writer, dress it up with the fonts you like, and export it in the PDF format. You can also use Scribus, the Linux desktop publishing program (and the same program used to assemble the PDF version of the PCLinuxOS Magazine) to create your PDF file. One advantage of using Scribus is that you can make links within the PDF document that allow you to easily jump to each chapter in the e-book, with ease. There is another choice for simple exporting of files to PDF format. From any program that can print, choose "Print To File," and choose to save that file as a PDF file. It may be simple and crude, but it's also very easy and fast.

Sometimes, you may find e-book content that is in HTML format. You can easily create PDF files from it, as well. Simply save the HTML file from your browser. Then, use HTMLDOC to convert the HTML file into a PDF file. Both the GUI and CLI version of HTMLDOC are in the PCLinuxOS repository.

If you are going to use OpenOffice to export directly to a PDF file – something it can do very quickly – you may want to change your page size so that it can be easily displayed on the commercial e-book readers. Change your page size to 3.6" x 4.7". You will then have a PDF file perfectly formatted to display on the 6" screens of many of the commercial e-book readers, and you can share with the rest of the e-book community.

So Where's The Content?

If you go with the commercial e-book readers, most of them have their own online stores where you can purchase content. Most of that content is DRM-protected, however. But if you like classical literature (read that literature that is in the public domain), then you have lots of options on where to obtain your e-book content. It is the free e-book content that I will focus on here. And you might be surprised at how recent some of that free e-book material is. I was surprised, as I found the entire series of books about vampires and werewolves series available, along with all of Dan Brown's books (more on that in just a bit).

Probably one of the best, and best known, places to obtain content is at Project Gutenberg. Started in 1971 by it's founder, Michael Hart, the goal of Project Gutenberg is to place public domain literature within the digital realm, and make those digital copies accessible, free of charge, to anyone who wanted access to them. Today, Project Gutenberg has over 20,000 titles in it's collection, and it is growing every year. They are also always seeking volunteers to help digitize public domain works. Most of the content there is as plain text files, or ZIP files of plain text files. So if you are a H.G.

Wells fan, or a George Orwell fan, or an Edgar Rice Burroughs fan, or a Mary Shelley fan, or the fan of any other author of classical literature, Project Gutenberg will be like a well-equipped playground to you.

Another good site is Scribd. You can find all sorts of recent material there, ranging from entire books to book excerpts to press releases to ... well, the sky is the limit. I was able to find a free copy of Dan Brown's newest book, The Lost Symbol. If you aren't familiar with Dan Brown, he's the famous author of The DaVinci Code, a best seller from a few years back.

I was also able to download all of Dan Brown's books, including Angels & Demons, the book upon which the popular movie of the same name is based upon and that is currently in the theaters, from the Colors and Harmony blog site. It was here that I also found the entire Twilight series of books, along with all the books about a popular boy wizard – whatever his name is.

Another free e-book site is Free-eBooks.net. As a free member, you can download five free e-books per month. The site features classical literature, as well as fan fiction. If you want to download more than five free e-books per month, you can become a "premium member" for less than $5 per month.

Some other free e-book sites, in no particular order:

getfreeebooks.com

Manybooks

Planet PDF Free E-books

Feedbooks

Planet eBook

BookYards

FreeBookSpot

Baen Free Library

PDFGenie.com

Summary

As we move ever deeper into a digital world, it's inevitable that our other forms of media follow, lest we lose them and all they do to enrich our lives. E-books are a natural bedfellow for that digital world. Despite all of our advances into an ever increasing digital world, there is little that can replace the joy that many receive by curling up with a good book. Except now, you can curl up with a few thousand at a time, and read one after another, where ever you may be, without pause.

Answers to Mark Szorady's Double Take: (1) Table different; (2) Shirt stripes different; (3) Confetti missing; (4) Balloon missing; (5) Mouth different; (6) "Stop" changed to "Start"; (7) Glasses stem missing

PCLinuxOnline

PCLinuxOS News and Information

Screenshot Showcase

Posted by ramchu, December 01, 2009, running KDE 4.3.4

**From Sproggy's Video, "Glass Panel Tutorial" Transcribed by Meemaw**

(Note: a few times Sproggy navigated to items that may not be known by those who are just starting out in Inkscape, so I put a few extra comments inside brackets like this [Comment] I hope it helps!)

Hello, my name is Kori, pet name is Sproggy. This is a tutorial for Inkscape, my favorite application for artwork. The version I used for the tutorial was 0.46. Let's just give it a check. [You can check yours by clicking Help > About Inkscape].
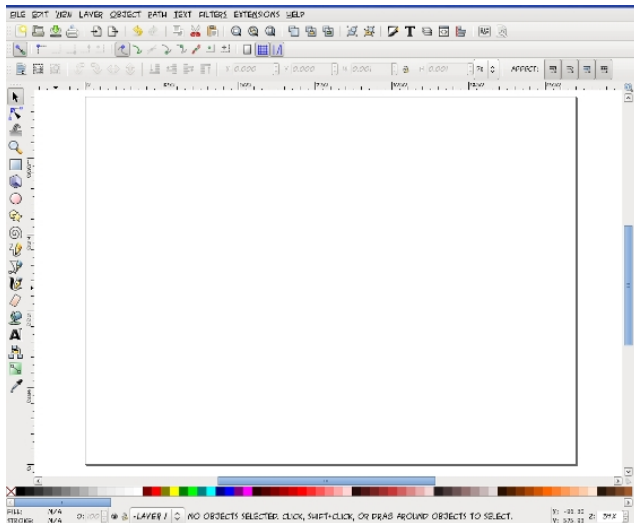

This is my drawing template I created – [Inkscape opens up a new document when it starts – go to File > Document Properties] Set custom size 1600 pixels width and 1200 pixels height – make sure there is a check in the box before "Show Page Border" - save as DrawingTemplate.svg, and you can just open it up when you want to start a basic project.
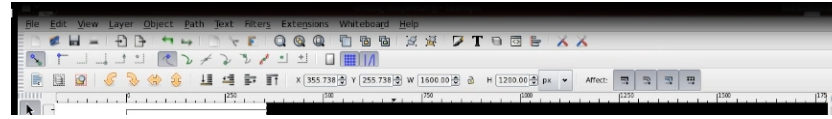

Inkscape Main Toolbars

To start a glass panel wallpaper, draw a rectangle in the center of the page. Up above, change the width to 1600 and height to 1200 [above your drawing is a tool bar with numbers you can edit – X and Y settings show where on the page your object is measuring from the bottom left corner of the page. W and H settings are the width and height of your object, so you don't have to strain with the mouse to get something a certain size – you can put in the size manually.] Change W to 1600 and H to 1200. Then we're going to change the position, so click on the toolbar button for **Align and Distribute** and center the rectangle on the page.
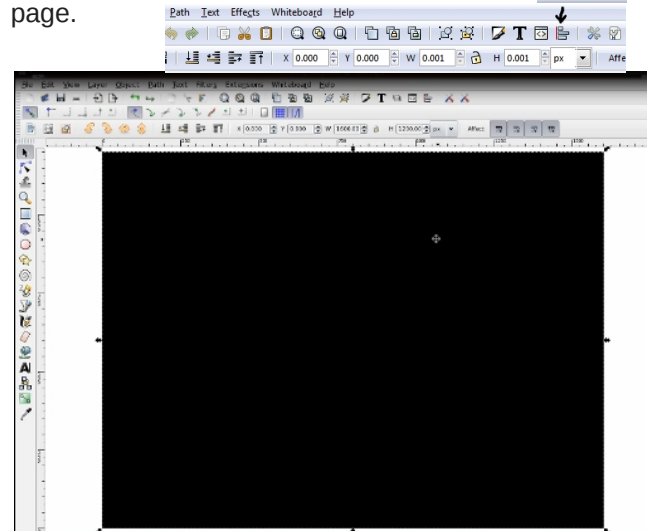

I know this may be a bit silly, but I always add a base background to what I create so should there be any transparency I can see if there are any errors with it. I'm actually going to use this object for my background but I'm going to click on **Edit Object** and change a few settings on this rectangle area. [In Stroke Paint he chose none by clicking on the X.] In Fill, I'm going to choose a radial gradient using these boxes. As you can se, it automatically puts a gradient here, but that's not the gradient I want for this image. I'm going to duplicate it and then I'm going to edit it. [Click Duplicate, then click Edit....] This is how the gradient looks now, dark to light from center to outside. We're going to reverse that. Click on the stops [you should get a drop-box with 2 stops.] Choose the second stop – go to the bar at the bottom and move the slides over to the right so the stop is completely black. Then choose the first stop and change the number to the right of the slides to 25. [Now your gradient should have the gray oval shape in the center and be dark at the outside.]
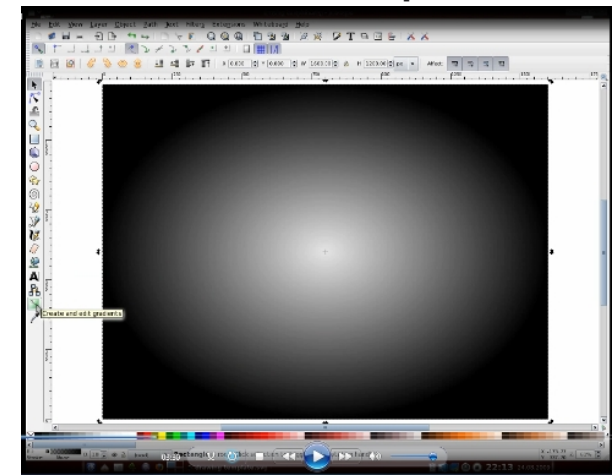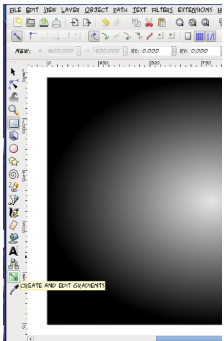
Over here is the **Create and Edit Gradients** tool. [towards the bottom of the left side of your page] I'm going to slightly zoom out.... Click on that and you'll see lines inside your gradient. Hold down CTRL and click on one of the endpoints and slide it away from the center to make the lighter part of the gradient larger. Do the same with the other endpoint as well.
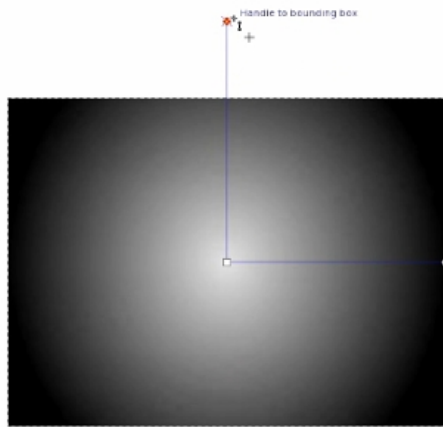

[Save your work.]

Click on draw rectangle again and draw another. This one will be 1000 x 500. Once again I'm going to align this to the center of the page. I'm going to duplicate this rectangle by clicking on it, holding down CTRL and pressing D. Then I'm going to click on the white [the color swatches at the bottom] to change the color of the top rectangle. Now I'm going to make the white rectangle lie behind the black one,
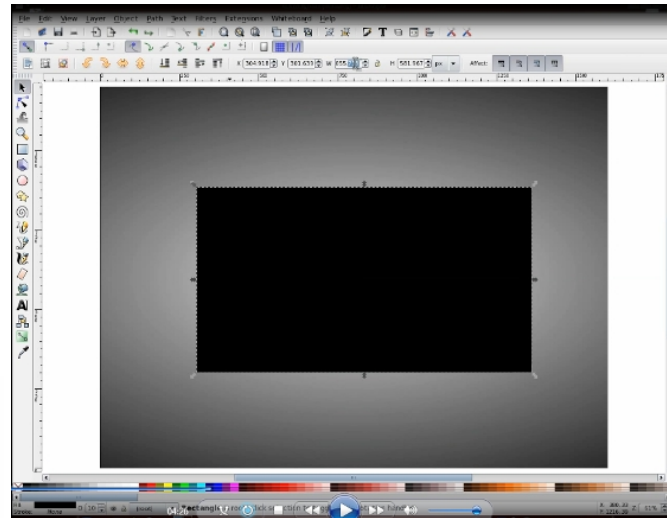

so up here we have four little buttons that we can click to position the layers. I'm going to click the "lower one level" button to move the white rectangle behind the black.

Now I'm going to make the white rectangle a little larger than the black one. Click on **Path > Dynamic Offset**, then grab the top left corner [there should be a 'grabber' there] and pull out to about there [gives it a kind of wide border.] Then I'm going to duplicate it [CTRL + D] Go back and click 'Select' [arrow at top left] because I'm going to drop it down as well. Then change it to black. [It should be behind the white rectangle.] Click on **Path > Dynamic Offset** again, but when you grab it, just pull it out enough so that it looks like a black line border. Click on **Select** again and go to **Edit Object** again. When you see the fill, click that off [click on the X for none] because I don't want fill on this object. You'll see it disappears. Now I want **Stroke Paint** and choose to have it as a flat
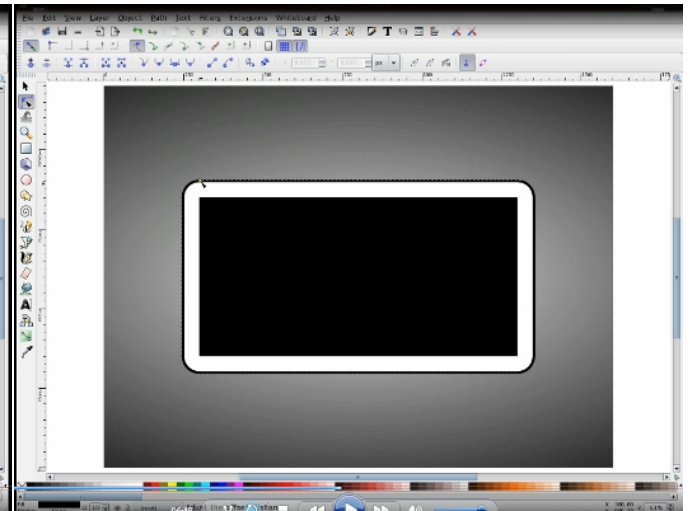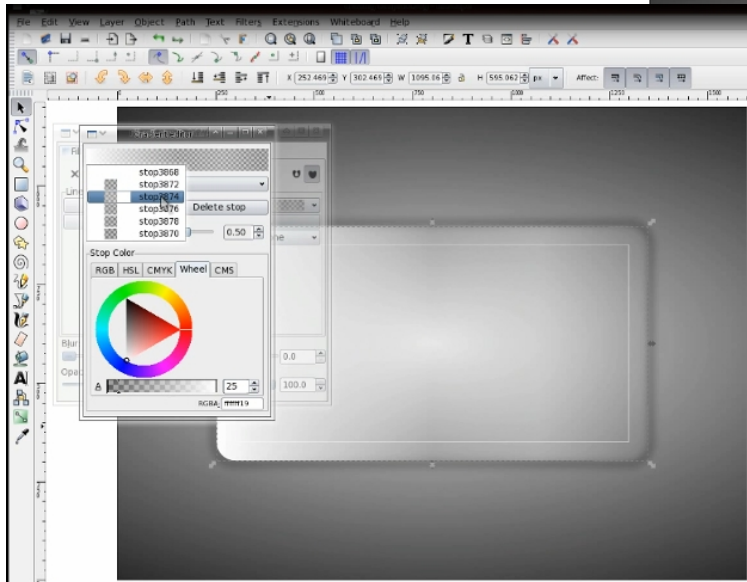

color [black] – then I need to click on **Stroke Style** and change the width of that to 5 [highlight the 1.000 in width and type in 5] – Now the width of the line alone is not too exciting, but if I put in a blur of 2.5, that slightly changes the image and it looks like you now have a complete drop shadow around the white border.   [Save your work.]
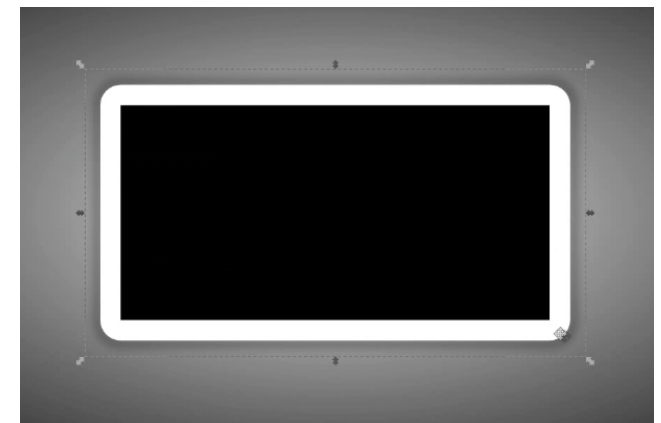
Now that we've gotten that far we're going to click back into the center rectangle again. We're going to click on **Edit Object**, change the fill to none, the **Stroke Paint** to white and check that the **Stroke Style** width is 1.000. Now we have a white rectangle and a smaller rectangle with a white border. Now click on the white rectangle. We need to add a gradient to this to make it "glassified" - we don't want a stroke on it but we do want a fill. Choose the linear gradient and click on that. I don't want the default one that shows up first, so I'm going to duplicate it,

then edit it. As in the first we see there are two stops here, so we're going to add five more stops to it [click Add Stop five times.] We should now have seven. We're going to go in and adjust the values for these stops. We'll change the second one to 25, and so our two gradient values will be 25 [light gray] and 255 [white]. The third one will be 255 and the fourth

will be 25 and so on. As you can see, it doesn't look quite like glass yet, so what I'm going to do is click on **Create and Edit Gradients** [there will be a line across your rectangle with a 'handle/grabber' at each stop.] Drag the grabbers – slide them and even them out across the rectangle to give it a slightly better, glassier look. Try to even them out as much as possible.
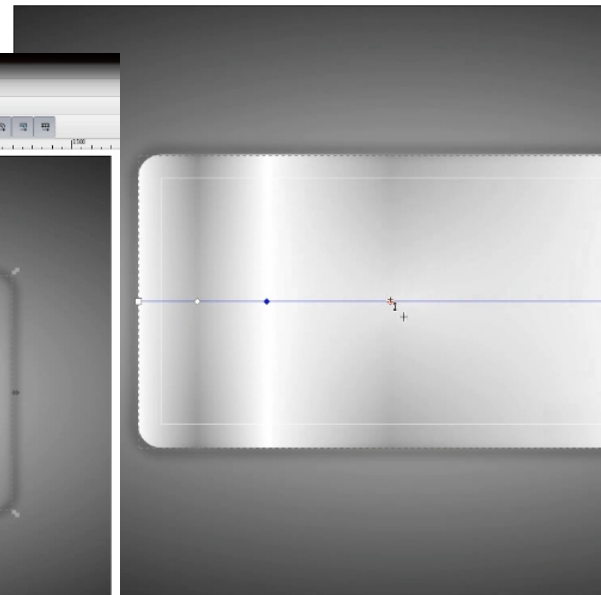
[Save your work.]

Now I'm going to take the end grabbers and move them diagonally from corner to corner. We're starting to get a glassy fill but it's still too heavy on the color, so we're going to select the object and edit the fill, changing it to 25%.
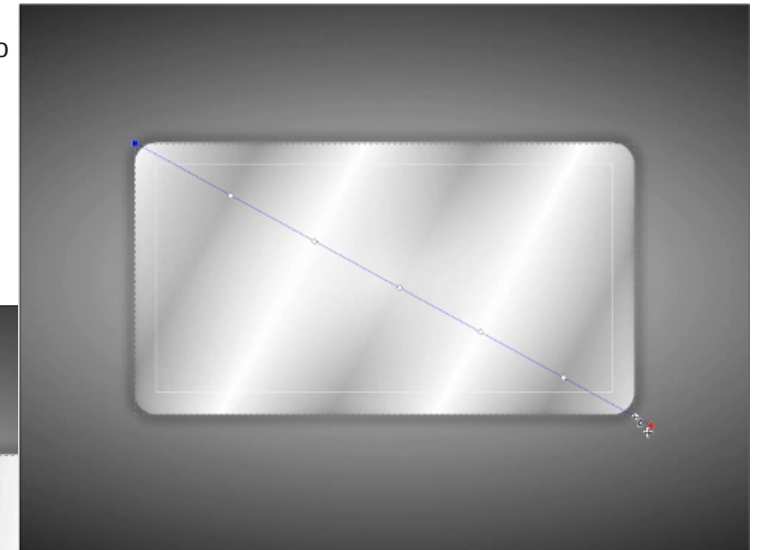
We are nearly there.

As with the white line, we're going to put a line around the edge, so duplicate the white rectangle. You'll notice it looks slightly different, but in this case, I'm going to select it and remove the fill, then set the stroke at 1 and leave the opacity at 25%.
   [Save your work.]

There you have the glass panel. You can't really appreciate its full clarity or transparency here, but I have one I made earlier, so I'm going to minimize Inkscape and the one I made earlier is on my desktop. It has my favorite Linux distribution's name emblazoned across it. As you can see, it's very effective. With it being transparent when you export it as a png, you can set it as your desktop background with no color, then all you need

Screenshot Showcase

to do is adjust the background color to the one you want. [Started red, then changed to blue, then to green.] See? No need to change the image, just change the background color.
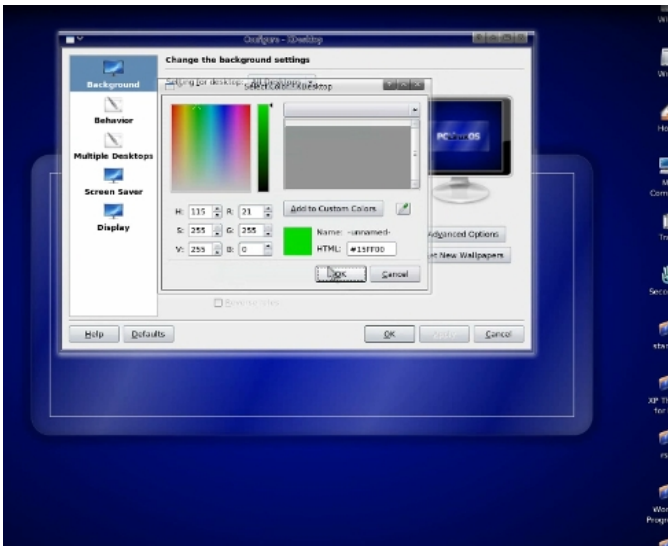
Posted by Agust on December 8, 2009 using E17

Forum Foibles: New Year's Resolutions

I resolve to spend all the 300+ days without logging onto windows!!!! even once! That could be a tall order, considering. **taelti**

I'm thinking that 2010, I should shed that crusty shell a bit, when in the presence of close friends. **sammy2fish**

I resolve to stay smoke free, and to work hard to get into better physical shape. **Paul Arnote (parnote)**

To get a job that gives me more time to do my art and phoenix. **Sproggy**

Understand the questions before I post responses. **rudge**

Learn more of programming. Contribute more to PCLinuxOS. Be a better friend. **Neal**

Rest as much as possible.
I've broken all the others I've made.
Maybe reverse psychology will work this time? **jaydot**

The New Year is now upon us
Oh how the time goes by
Members have made their resolutions
Will they keep them....or just try?

Thanks to critter's magazine tutorials I resolve to learn and **use** a new command each wee**k ...**for at least 4 weeks. **ms_meme**

I resolve to never use my main computer for experimentation..... lest it have to be re-installed...... but to only use my old laptop that is the extra 'quick-internet-look-up' machine that everyone uses. **MeeMaw**

My New Years Resolution: To be more like Neal. **steve161**

Resolution? I'm just gonna wait for the year to unfold. **Archie**

My resolution is to take better care of myself and be kinder to my fellow humans.. I know that last one is kinda up in the air...but will give it my best. **Scoundrel**

I resolve to learn packaging this coming year and spend more time with the grand kids! **Wildman**

Never question the police. Always answer the police. **hornerm**

I resolve to get my piano out of storage and start playing again. It's been in storage since a recent move. **gseaman**
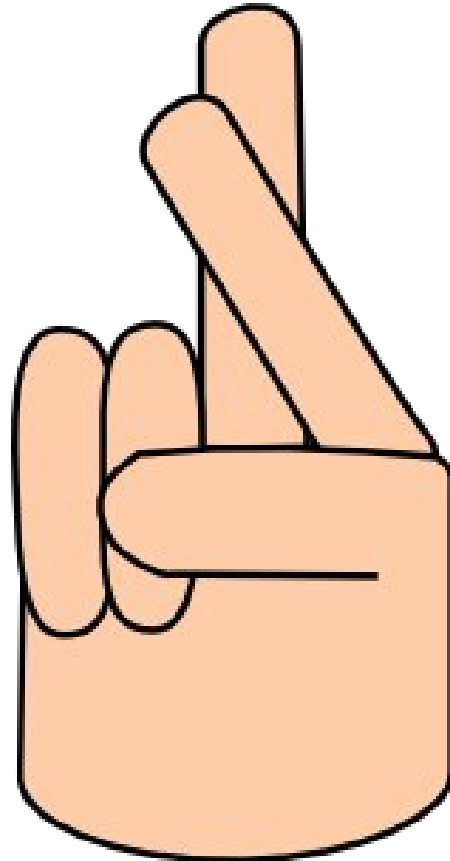
**You too can make a committment
Promise PCLOS to be true
Renew your pledge to help
As a way of saying "Thank You"**

ms_meme

I resolve to get a webcam and set up Skype under PCLinuxOS to be able to see and talk with my daughter living in Utah. Hey...I don't have any vices. **MGBguy**

by Peter Kelly (critter)

Processes

When you start an executable file, either by clicking on its icon or by entering its name in a terminal, you are actually starting what Linux knows as a process. This may, in turn, start other processes, or series of processes. Other processes are started during the boot up sequence, or automatically as required by the system.

A process is what acts on files and the data that those files contain. As Linux is a multitasking, multi-user operating system, you and other users on the system may start many applications that have their own sets of processes. Obviously, Linux has to have a method of controlling and keeping track of all these processes.

Each process is given a unique identity when it is started, known as the process id number, or PID, is allocated space in memory and assigned certain other system resources that needn't concern us here. When booting the system, the first process to be started is always init, and is given the PID number 1. To get a graphical representation of the processes currently active on your system type in a console:

pstree -p

The -p option shows the PIDs.

The initial process started by an application is known as the **parent**, and any processes spawned from that process are known as **child processes**. Modern systems use threads, or Light Weight Processes (LWP), as child processes where practical, which share some of the resources of the parent. Communication between the system and processes is done by signals, and the whole show is orchestrated by the scheduler.

Usually all this happens under the hood, and we don't have to bother about it. But just occasionally, something goes wrong, the system starts to misbehave, and user intervention is required. This is where familiarity with the command line comes in.

Open a terminal and type:

glxgears > /dev/null &

This a fairly resource intensive application (usually used for benchmarking systems) that we don't want to see the output from in this case, so we dump it to nowhere. The & character puts this process running in the background so that we get control of our terminal back. More about background tasks in a moment.

Now type:

top

That's a lot of information! Don't worry, you don't need all of this. Look at the line representing the glxgears application (the next to last line in the screen shot). The first column is labeled PID and has a value of 8534, that's what we need.

Suppose this application refused to respond and you couldn't close it. In top type **k** and you will be prompted "PID to kill:" Enter 8534. This gives the prompt **Kill PID 8534 with signal [15]:**

If you enter **n** here, the command is canceled. Enter **y**, and and the application is sent the default signal 15. Signal 15 is known as SIGTERM, and is the normal termination signal. On receipt of this signal, the application will close gracefully, handing back any resources to the system, kind of like asking somebody to 'please leave'. Occasionally, you meet a troublemaker of a process that just will not go. Then, you can use the strongest signal of all, signal 9, known as SIGKILL. This is more of a bullet in the head approach – sure it works, but it might just make more trouble in the long run, and is best avoided if possible.

There are lots of signals available, but those are the only two that you are likely to have to issue manually. Top is a very comprehensive program that uses single letter commands. **h** brings up the help screen, **A** for the alternate display using field groups, **z** toggles the color of the current group, and **W**

saves your settings. There is so much more to this utility that you really do need to use the 'h' (help) command to get the most out of it. Remember that Linux is very literal, and uppercase or lowercase commands give very different results.

```
4:Usr - 08:02:45 up 1:09, 1 user, load average: 0.19, 0.11, 0.15
Tasks: 91 total, 2 running, 89 sleeping, 0 stopped, 0 zombie
Cpu(s): 1.0%us, 4.7%sy, 0.0%ni, 94.3%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 1035628k total, 335784k used, 699844k free, 65320k buffers
Swap: 533192k total, 0k used, 533192k free, 169604k cached

1 PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM   TIME+  COMMAND
 2766 root      20   0 45712  23m 4420 S  1.3  2.3 24:59.92 X
 3234 jane      20   0 35272  16m  13m S  1.3  1.6  0:01.74 kdesktop
2 PID  PPID    TIME+  %CPU %MEM PR  NI S  VIRT SWAP  RES  UID COMMAND
12469 3215   0:00.10 0.0  0.9  20   0 S 31444  22m 8912  501 kio_trash
12455 3215   0:00.10 0.0  0.8  20   0 S 31360  22m 8360  501 kio system
3 PID %MEM  VIRT SWAP  RES CODE DATA  SHR nFLT nDRT S  PR  NI %CPU COMMAND
 3282 2.6 40072 12m  26m    4 16m 9976  123    0 S  20   0  1.3 net_applet
 2766 2.3 45712 21m 23m 1728 20m 4420  69     0 S  20   0  1.3 X
4 PID  PPID  UID USER    RUSER   TTY        TIME+  %CPU %MEM S COMMAND
 2646  1   73 xfs     xfs     ?       0:00.06 0.0  0.2 S xfs
 2766 2730   0 root    root    tty7   24:59.92 1.3  2.3 S X
 2751  1    0 root    root    ?       0:00.06 0.0  0.3 S smbd
```

Another useful tool for manipulating processes is **ps** (process status). To see a list of all processes owned by the user issuing the command **ps ux** (this will be you if used as a normal user, root or maybe someone else if you have used the su command).

Note that the leading hyphen usually required when supplying options to a command, is optional for this command.

Even when you think that you aren't doing anything, you will get a fairly long list. Scroll up to the top of the list and you will see a header line describing the various columns of information. The second column here is the PID, and the last column is the command that was used to start the process. If, as in the previous example, we were trying to find the pid of glxgears we would use the command

ps ux | grep glxgears

```
jane@home > ~ $ ps ux | grep glxgears
jane   11730 1.6 0.1  3780 1316 pts/1    S    07:10   0:02 glxgears
jane   11931 0.0 0.0  2996  704 pts/1    S+   07:13   0:00 grep --color glxgears
jane@home > ~ $
```

This gives us two results. The first one, with the PID of 11730, is the one that we are looking for. The second one is the grep command we used to filter the results. Why grep ---color? Where did that come from? PCLinuxOS provides an alias for the grep command so that matches are highlighted in color (or, if you live in the UK, they will be in full glorious colour). Type alias | grep grep on the command line to see it.

```
jane@home > ~ $ alias | grep grep
alias grep='grep --color'
jane@home > ~ $
```

The grep command has also filtered out the header line, as that did not contain the expression 'glxgears'.

To get rid of the errant process, we use the kill command:

kill 11730

We could specify a signal, as in kill -s 9 11730, but I think that the default, unstated signal 15 is powerful enough magic.

A signal may be sent by name or by number. If you want the complete list of signal names and numbers type

kill -l

You may only terminate any process that you own,

so be especially careful when working with root privileges.

When working in a GUI, if you want to start another application, then clicking on its icon will open it in a new window. For command line work, you can open another terminal emulator. But what if you are locked into a single terminal, as you may be if the system has crashed? You may occasionally start a process which takes a long time to complete and need to execute another command.

One way around this is to follow the command with an ampersand (**&**), as we did with the glxgears application. This puts the process into the background and returns control of the terminal to the user. If the process is already running, it can be put into the background using the command **control + z**, which suspends the process running in the foreground, and then use the command **bg**, which causes the process to resume execution in the background.

To list the processes running in the current shell, use the command **jobs**.

```
jane@home > ~ $ glxgears > /dev/null &
[1] 8319
jane@home > ~ $ jobs
[1]+ Running                 glxgears >/dev/null &
jane@home > ~ $
```

Each job started in the shell is given a number, which is displayed in square brackets, along with the PID. Here, job 1 has PID 8319. The jobs command displays the job number, its status (running, Stopped or Terminated), and the command that initiated the process. To kill a job, we use the **kill** command, like this:

kill %N && fg where N is the job number.

The **&&** part is used to tell the shell to wait until the kill command has finished, and then to move to the foreground. If we don't add that bit, the the job will be terminated, but not removed from the job list. Just a bit of tidying up.

I know this all sounds rather complicated but it can all be summarized like this.

Command &: Start a job in the background

Control + z: Suspend the job currently running in the background

bg N: Continue suspended job N in the background

fg N: Move suspended job N to the foreground

jobs: List all jobs

kill %N && fg: Kill job N

For the glxgears example, we dumped the output to /dev/null. When a job is running in the background, it will still produce output, which it will happily spew out to the screen as you are trying to work on another command. So, it is usually a good idea to redirect the output of a back-grounded job to a log file, or similar, and to dump any error messages.

e.g., **find /usr -iname *.ogg > musicfiles 2>/dev/null &**

This will put the names of any files found into the file 'musicfiles', ignoring the case of the file name, and discard any error messages, such as trying to enter directories for which you don't have access. If you don't want to keep any output at all from the application, then there is a special construct that will grab everything, and send it to where ever you like, usually /dev/null.

command 2>&1 /dev/null

All this means is append **STDERR** to **STDOUT**, and send them both to the same place. Don't worry if this doesn't seem very intuitive. It is a very commonly used expression, but you would be amazed at how many experienced users who use it don't understand it!

Backing up and Archiving

The problem with data backup is that most people don't do it. They may mean to, they may forget, they may find it too complicated, or they may not do it regularly enough. Even those who doggedly back up their data regularly, rarely bother to check that the data can actually be restored, which rather defeats the object. If you are one of the very small majority who don't fall into this category, then you can skip this next part.

For the rest of us, there is some good news. There is some software available that will do all our backups for us, will never forget, and will check the integrity of the backup automatically. And best of all, it is free! It's called Linux.

Before starting a backup strategy, you should consider what you want to backup, how often, and to where. There are various types of backup.

Full - What it says

Differential - backup only what has changed since the last full backup

Incremental - backup only what has changed since the last backup of any type

Clone - Usually used for the backup of entire file systems, mindlessly copying everything block by block

Synchronized - An exact copy of all data in a location, optionally deleting data in the target destination that no longer exists in the source location.

Each has it's pros and cons, and there are many dedicated applications that will do backups in any way that you want. The Linux/Unix ethos is to use multiple small applications to perform more complex tasks, and using this, we can produce a tailor made backup system and automate it. First off then, let's take a look at some of the applications that are available. All of the following are either installed by default, or available in the PCLinuxOS repositories.

cp: Yes, the simple copy command, when used with the -a (archive) option, is a simple (but not very efficient) backup utility.

tar: One of the oldest utilities, its name means 'tape archive', and it simply takes a list of files and lumps them all together. For this reason, the resulting file is often referred to as a 'tarball'. It is often used in conjunction with a compression utility, like gzip or bzip2, and it can do this automatically. To create a compressed archive of all files in the current directory (here I assume that jane is in her Pictures directory) and write it to a folder in janes home directory named junk, I might use the command

tar -czf ~/junk/mypics.tar.gz *

-c create an archive

-z filter the output through gzip to compress it

-f use the following file name for the archive

Although Linux doesn't need dos-like file name extentions such as .exe or .zip to determine the file type (the information is in the meta data pointed to by the files inode), it is helpful and traditional to do so with the tar command. If I had substituted **-j** for **-z**, then tar would have used the bzip2 utility. Then, the extension .tar.bz is usually used.

To list the files in an archive use:

tar -tvf mypics.tar.gz

To extract the files use:

tar -xf mypics.tar.gz {list of file names}

If no file names are given, then all files are extracted. Tar is really only suitable for backing up a small number of files at a time. For large or full backups, there are better tools available. When using compression utilities, be aware that they are not very fault tolerant, and the loss of even one byte can render the entire archive unreadable. This is a very real danger when using media such as floppies r CDs.

```
jane@daisy > ~ $ cd Pictures/
jane@daisy > Pictures $ tar -czf ~/junk/mypics.tar.gz *
jane@daisy > Pictures $ cd ~/junk/
jane@daisy > junk $ ls -l
total 44
-rw-r--r-- 1 jane jane 42933 Dec  6 11:35 mypics.tar.gz
jane@daisy > junk $ tar -tvf mypics.tar.gz
-rw-rw-r-- jane/jane      13165 2009-12-01 13:23 snapshot1.png
-rw-rw-r-- jane/jane       7460 2009-12-01 13:22 snapshot2.png
-rw-rw-r-- jane/jane       2511 2009-09-05 18:39 snapshot3.png
-rw-rw-r-- jane/jane       2481 2009-09-05 18:58 snapshot4.png
-rw-rw-r-- jane/jane       4059 2009-09-05 19:08 snapshot5.png
-rw-rw-r-- jane/jane       8435 2009-09-05 19:39 snapshot6.png
-rw-rw-r-- jane/jane       7138 2009-12-01 12:33 snapshot7.png
jane@daisy > junk $ tar -xf mypics.tar.gz snapshot5.png
jane@daisy > junk $ ls -l
total 48
-rw-r--r-- 1 jane jane 42933 Dec  6 11:35 mypics.tar.gz
-rw-r--r-- 1 jane jane  4059 Sep  5 19:08 snapshot5.png
jane@daisy > junk $
```

dd falls into the 'clone' category, and is an extremely useful command to know how to use.

There are lots of options for this command, but most of the time you can get by with just these four:

if={name of the input file}

of={name of the output file}

bs=(block size}

count={number of blocks to transfer}

Notice the **=** sign in the options. This is mandatory. The important thing to remember when using dd is not to mix up the input file and the output file, as you will get no prompt from dd – just blind obedience. So, if you wanted to copy a partition to an empty partition and you mixed them up, then you would overwrite the first partition with blank data from the empty partition losing all data, no questions asked.

To copy a single file:

dd if=/home/jane/contacts2 of=/home/jane/junk/contacts

```
jane@daisy > ~ $ dd if=/home/jane/contacts2 of=/home/jane/junk/contacts
0+1 records in
0+1 records out
60 bytes (60 B) copied, 0.000218269 seconds, 275 kB/s
jane@daisy > ~ $ ls -l junk/
total 52
-rw-r--r-- 1 jane jane     60 Dec  6 11:43 contacts
-rw-r--r-- 1 jane jane 42933 Dec  6 11:35 mypics.tar.gz
-rw-r--r-- 1 jane jane  4059 Sep  5 19:08 snapshot5.png
jane@daisy > ~ $
```

To copy an entire partition (make sure that the destination is large enough):

dd if=/dev/hda1 of=/dev/hdf1 to do a partition to partition copy

or

dd if=/dev/hda1 of=/backup/hda1.img will make an image file of the partition.

This way, entire file systems can be cloned and then restored with

dd if=/backup/hda1.img of=/dev/hda1

PCLinuxOS
Radically Simple

Here, /dev/hda1 should not be mounted, and not be part of the currently running operating system.

If you have a 1TB drive, then be prepared for a long wait.

The first 512 bytes of a hard disk contains the MBR, or Master Boot Record, and the Partition table. If this gets corrupted, then your system may not know how to boot up, or how your drive is partitioned. This can be fixed, if you are patient and have some in depth knowledge of file systems (or some understanding friends), but life can be so much easier if you have a backup of that data. Here dd is ideal. su to root, and type:

dd if=/dev/hda of=/backup/hda_mbr.img bs=512 count=1

```
jane@daisy > ~ $ su
Password:
[root@daisy jane]# dd if=/dev/hda of=/backup/hda_mbr.img bs=512  count=1
1+0 records in
1+0 records out
512 bytes (512 B) copied, 0.000261765 seconds, 2.0 MB/s
[root@daisy jane]# cd /backup
[root@daisy backup]# ls -l
total 64
-rw-r--r--  1 root root   512 Dec  6 12:07 hda_mbr.img
drwx------ 24 jane jane  4096 Dec  5 10:49 jane/
-rw-r--r--  1 jane jane 52622 Dec  5 12:03 jane-log
[root@daisy backup]#
```

This means copy one block of 512 bytes from /dev/hda to the image file /backup/hda_mbr.img.

(Note: no partition number is entered as the MBR relates to the whole drive, not to any particular partition).

Now if you do have problems, you can simply restore those 512 bytes. To restore the MBR, the drive should not be running the current operating system.

Boot up from the Live CD.

In a terminal, su to get root privileges.

Make a directory, e.g **mkdir /mnt/hda**,

Mount the drive there, **mount /dev/hda /mnt/hda**

Make sure that the image file is where you think that it is, i.e **/backup/hda_mbr.img**, as running from the Live CD, the directory /backup may not exist.

dd if=/backup/hda_mbr.img of=/mnt/hda bs=512 count=1

Stop and re-read this command at least twice before pressing enter.

You should now be able to reboot the original system.

If you want to see what a MBR looks like, (it won't mean much) type:

cat /backup/hda_mbr.img

Then you will get a load of gibberish on the terminal, and the chances are that your prompt and anything you type on the command line are now gibberish. If that is the case then don't panic. Just type:

reset

and all will be well.

Mondo archive is an excellent and extremely reliable full system backup tool which is run from the command line, and includes a graphical interface in the terminal. There is an excellent step by step tutorial in the PCLinuxOS forums by catlord17 (http://www.pclinuxos.com/forum/index.php/topic,597 05.0.html).

Rsync is one of my all time favorites. It is used to synchronize two sets of files, either locally or between remote machines, even if they are thousands of miles apart. It includes options for compression, can delete files on the destination that no longer exist in the source location, and can be used over ssh (secure shell). Best of all, it only transfers the differences between files on the source and destination, and if interrupted, will resume from the break point when restarted. To give you an idea of the power of this command, I maintain a copy of the PCLinuxOS repositories on a local hard drive, and automatically update it in the early hours, when network traffic is lightest. The full repository is currently around 17GB, but I download only what has changed and delete obsolete packages. The command to do this is quite long, but as it is automated, I don't have to remember it. If you're interested it looks like this:

rsync -aPvz --stats --delete --log-file=/home/pete/update.log --exclude=SRPM* ftp.heanet.ie::pub/pclinuxos/apt/pclinuxos/2007/ /data/repo/2007/

partimage is another command line utility that includes a graphical interface and is quite easy to use. It will make an image of an unmounted partition, using compression if required, and will only

copy used blocks. It allows you to include comments that can be useful when restoring. For example, you could include the size and type of file system of the original partition. It will also allow you to split up the image so that it may be saved to multiple CDs or DVDs.

Creating a backup

Now that we have a few tools to play with, it is time to decide what to back up and where to save it. The size of the backup may determine where to back up to, but ideally, you want to back up to a different device than the one that holds the original data. If a hard drive dies, then you will be relieved that your back up didn't go with it. If you are doing a full system backup for disaster recovery, then mondo archive or partimage are probably the way to go. If you want to incrementally backup the files from a particular location that have changed over a set time period, then first you have to get a list of those files. The find utility includes an option to locate files based on the time they were last modified. To find all files in a directory that have changed in the last week

find /home/jane/ -mtime -7 \! -type d

-mtime looks for the file modification time.

Here I have excluded files of type directory with **\! type d**.

The ! means 'not', and must be preceded with a back slash to prevent the shell from interpreting it as a special character. This is known as 'escaping' the character.

For files changed in the last week use:

find /home/jane/ -mtime -1 \! -type d

If we redirect the output of this command to a file, then we can use that file to feed tar:

find /home/jane/ -mtime -1 \! -type d > /tmp/newfiles
tar -cv -T /tmp/newfiles -f /backup/newfiles.tar

The **-T** option tells the utility to get its list of file names from the following file.

Reinstalling a Linux system these days is a relatively quick and trivial task, but getting everything back exactly as you like it can be more of a headache. All of your personal settings are stored in your /home/{user_name} folder. This is convenient, but has the disadvantage that as you own those files, then you can do whatever you like (or dislike) to them, and this is what new users usually do and experienced users, who should know better, still often do. This directory then is a prime candidate to have a backed up copy somewhere safe that is fully updated automatically and regularly. Sounds like a job for rsync.

First, we need to collect a bit of information.

It is necessary to preserve ownership, groups, links and modification times.

We want to back up /home/jane but not all of it.

We don't want to copy the contents of junk/, a folder jane uses to store temporary stuff.

There is a second hard drive mounted at /backup, and this is where we want our backup.

As this is going to be an automatic process, we should save the output of the program to a log.

No compression – data integrity is paramount, and the destination has several hundred GB free.

Obsolete files should be deleted.

The backup should be run daily, 1:30 AM, and also weekly at 3:00AM Sunday under a different name. This gives a week to have a change of mind about some settings. The odd times are chosen to be when the system is likely to be under a low load, and to avoid any conflicts due to daylight saving time changes.

Recovery should be a trivial task.

The command then should look like this:

rsync --archive --exclude=junk/* --delete --log-file=/backup/daily/jane-log /home/jane /backup/daily/

This must be all on one line. If your command includes more options and your editor can't cope then you will need to use the line continuation sequence

\enter.

It works like this. When you reach the end of a line and press the enter key bash inserts an invisible character known as a 'newline'. If we type \ before pressing the enter key, then this newline character is 'escaped' or in effect ignored. The effect of this is that you can continue typing extremely long commands and bash will see it as one long line. I think the limit is 2048 characters, but I have never reached it, and I have typed in some pretty long commands!

Here, I have used the long version for all the options, as this is going to be saved to run automatically, and this format makes the command easier to follow in the future.

The trailing slash on directory names is important.

/home/jane backs up the directory jane and all the files and directories under it into the directory /backup/daily.

/home/jane/ would not create the directory jane but would copy all of its contents into /backup/daily.

Similarly we want the directory junk/ to be created so that it is still available if we do a restore, but we don't want its contents. So, we exclude only the contents with junk/*.

When you have built a long command, you don't want to have to retype it every time. You could create an alias but a better, and much more powerful method, is to include it in a script. Scripts can be very long and complicated, or short and to the point like this one.

Make two directories in the destination – here the destination is /backup:

mkdir /backup/{daily,weekly}

The curly braces include a comma separated list if directories to create.

Open your text editor of choice, and type in the following (copy and paste from the magazine doesn't always work as it carries over some of those 'invisible' characters that are used to typeset the article).

#!/bin/bash
#Script to back up my home directory
rsync --archive --exclude=junk/* --delete --log-file=/backup/daily/jane-log /home/jane /backup/daily/

Don't forget that the long rsync command must all be on one line.

I'll Explain all the weird stuff at the beginning when we get to more involved scripts.

Save it to your home directory under whatever file name takes your fancy (e.g. home-bkup). Having it

in your home directory has the advantage that it also gets backed up – self preservation.

Change the files permissions to make it executable.

```
jane@daisy > ~ $ chmod +x home-bkup
jane@daisy > ~ $ ls -l home-bkup
-rwxr-xr-x 1 jane jane 151 Dec  6 12:57 home-bkup*
jane@daisy > ~ $
```

Type ./home-bkup and the script should execute – check it out. ./ is needed to tell bash to look first here for the file, which is not where it would normally look for commands.

ls -al /backup/daily

```
jane@daisy > ~ $ ls -al /backup/daily
total 140
drwxr-xr-x  3 jane jane    4096 Dec  6 12:57 ./
drwxrwxrwx  4 root root    4096 Dec  6 12:53 ../
drwx------ 26 jane jane    4096 Dec  6 12:57 jane/
-rw-r--r--  1 jane jane 124517 Dec  6 12:57 jane-log
jane@daisy > ~ $
```

Make a copy of the file with a new name for the weekly backup.

cp home-bkup home-bkup-weekly

Open this new file in an editor, and change the destinations in the command to read:

rsync --archive --exclude=junk/* --delete --log-file=/backup/weekly/jane-log /home/jane /backup/weekly/

Save the file.

Check that this one also worked. You should have a copy of your home directory and a log in each of the

two new directories that we created in **/backup/ – daily/ and weekly/**.

Now, to backup your home directory, all you need to do is to run that script. This, however, is where a lot of backup strategies fail, as you have to remember to run it. So let's automate It. To do this, we use a utility called cron. This utility is used by the system to run certain things at a particular time, like updating the database that is used by the locate command.

To set up cron, **su** to root. We need to use:

crontab -e

The **-e** puts us in edit mode in the vi editor.

The top line of the file shows the required format for the file. The first five columns are for when we want the command executed. An asterisk in any column means 'from the first to the last'. In the day of week column, **0** and **7** both represent Sunday.

So to execute our two scripts at 1:30AM every morning and 3AM Sunday respectively we add lines as shown.

Type o to open a new line below the top one and type in these two lines.

30 1 * * * /home/jane/home-bkup

0 3 * * 7 /home/jane/home-bkup-weekly

Press the escape key, followed by a colon, and you will get a prompt at the bottom of the screen. Type **wq** (write then quit) to save the file and then quit.

```
# min(0-59)  hours(0-23) day(1-31) month(1-12) dow(0-7)    command
30 1 * * * /home/jane/home-bkup
0 3 * * 7 /home/jane/home-bkup-weekly
~
~
~
~
~
~
:wq
```

Type **crontab -l** to see a listing of the file and check your typing.

Tomorrow you can check to see if it worked!

Screenshot Showcase
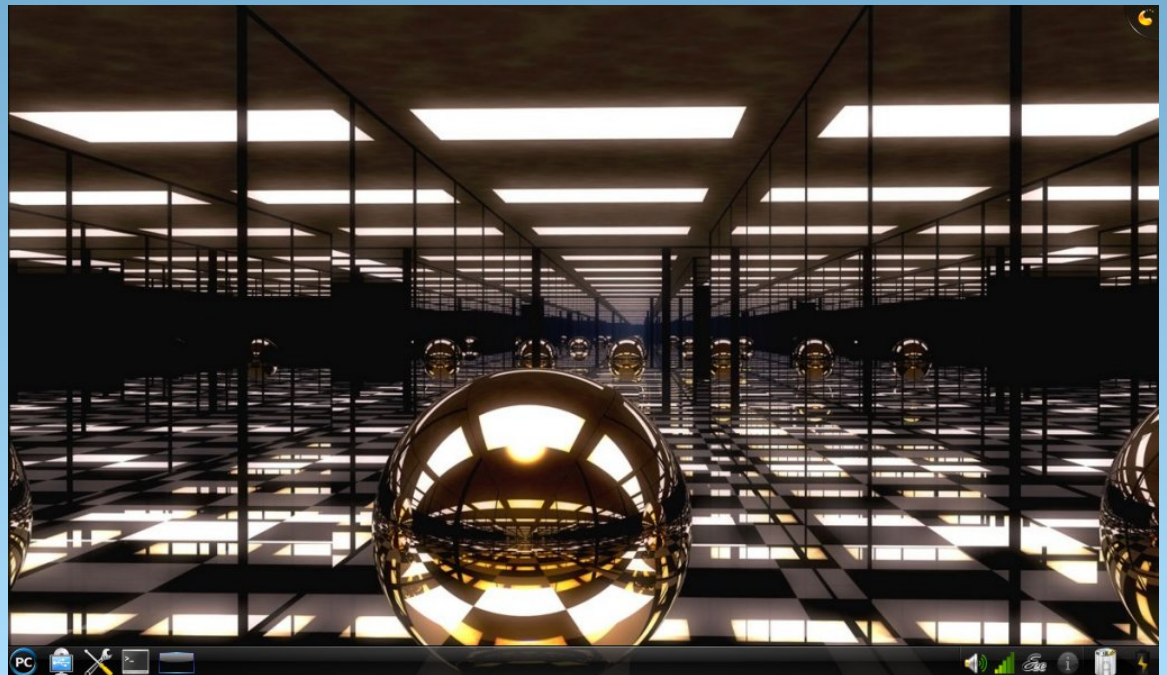
Posted by coffeetime, December 3, 2009, running KDE 4.3.4 on eeePC

Archiving Movies With dvd::rip – An Update

by Malcolm Ripley (malc)

In the September, 2009 issue of the PCLinuxOS Magazine, Malcolm covered how to archive your DVD movies using dvd::rip (available from the PCLinuxOS repository). This month, Malcolm updates his approach with dvd::rip, giving us additional tips on how to speed up the process, on how to re-sync audio and control the file size, on how to improve the audio, and finally, on how to rip subtitles. – Paul Arnote, PCLinuxOS Magazine Chief Editor.

The notes given here describe some minor changes to the overall process to improve the quality, or rather consistency, of the ripped videos. In addition, subtitles have now been assessed.

Preparation

Install *lsdvd*. This does two things. First, it allows dvd::rip subtitling to work. Secondly, it makes the reading of DVD titles and chapters much faster.

According to the dvd::rip website, you need *subtitle2pgm* for subtitles, which is not available for download via Synaptic. Yet quite by accident (and subsequent logical deduction), I discovered that *lsdvd* enables this functionality.

Out of Synch Audio and Smaller videos

I originally recommended entering a *Frame Range* start of 250. However, even this is not enough for some videos. A more consistent approach, which also clips the ending titles, and thus saving 20-30Mb, is given below:

First, after ripping the selected title, view it and note where the video starts, taking an approximate note of how many seconds it is.

Second, on the *Clip & Zoom* tab, set the frame size as described in the original article.

Third, enter a value in *Grab preview frame #* starting with 25 time the number of seconds to the start of the movie, and press the Grab frame button. This is a rough guide, since I actually round to the nearest 100. Confirm that this is at the start of the video. If it is not at the start of the video, adjust the value and repeat. Note the number.

Fourth, enter a frame number near the end of the video (I round this to the nearest 1000), and confirm that the frame is at the start of the ending titles. I have found that it is best to set this where the white titles are against a black background. With a colored background you can often have some end of video comments (e.g. "Joe Bloggs spent 10 years in jail before release"), stills or bloopers and you don't want to clip those. Note the number.

Fifth, and finally, on the *Transcode* page, enter the two frame values for the start and end of the video.

Dark Video Frames

Be aware that on the *Clip & Zoom* tab, the example frame comes from the middle of the ripped title. Sometimes, this can be a little dark and the *Apply Preset Values* action can be fooled. To avoid this, modify the preview frame to be one that clearly has visible pixels at all edges.

Improved Audio

I now select MP3(stereo) for dramas, comedies, TV shows, etc. For action, sci-fi, horror, etc., I select AC3 to get the full surround sound. Be careful, as AC3 increases the size of the video quite a bit.

Subtitles

RIP Title **Tab**

Before ripping the title, ensure the correct subtitle language has been selected.

***Clip & Zoom* Tab**

First, set the frame size, as originally described.

Second, you will need to set where you want the subtitles to appear. If you want the subtitles to appear in the black border at the bottom of the screen, adjust the second clipping bottom value to be zero, or as close as possible, ensuring that the total of both top+bottom is a multiple of 16. For example 56 top and 56 bottom, becomes 56 top and 8 bottom. Alternatively, if you prefer to embed the subtitles over the film image, simply adjust the vertical offset in the Subtitles tab (below).

Third, click *generate preview images*.

Fourth, you could now determine the start of the video for audio synch problems. But, if its not in your native language and you are reading subtitles, would you notice?

The above alternate steps is down to personal choice. Personally I prefer the bordered subtitles for a film with subtitles all the way through (foreign language). I generally use the offset option for films with occasional subtitles (e.g. labels).

***Subtitles* Tab**

Ensure that *DVDtitle* is the correct one.

Check that the correct subtitles are selected by viewing some of the panes. Video commentaries

are also subtitled, and you usually don't want those. There are also videos, like "The Mummy," that have subtitles for the hard of hearing, as well as subtitles for the non-English parts. You need to select the correct set. Also take note of the type of text, solid or outline.

Tick : *Active this subtitle for rendering*.

Tick : *Colors enable manipulation*.

For Subtitle images with solid black text, select *Gray A/B* values of 255 and 0, and *Index A/B* values of 3 and 1.

For Subtitle images with outline text, select *Gray A/B* values of 255 and 0, and *Index A/B* values of 1 and 3.

***Transcode* Tab**

As originally described.

Final Word

I can't think of anything new to add until I start ripping blu-ray. Just remember to keep within the copyright laws of your location.

Testimonial: Easy Printer Setup

-retired

My old Epson printer had hiccups, so I needed a new one quickly, and based on remarks generally around the forum, I selected an HP - the PhotoSmart b109n. I got it home, took it out of the box, but the set-up instructions were lousy. So, I though I'd set it up on Windows first, to make sure all was well. And because there were lots of suggestions around about downloading specific Linux drivers, etc. Then, more than 1/2 hour later, after waiting for various things to run, I eventually got it printing. There were no problems, understand, just vague directions and time to run various bits of 'setup.'

Then I thought, "OK now for PCLOS." One minute later, it is printing out a test page! Then the 'Biggy.' This printer has WiFi built in. In 2 minutes, the WEP key was entered, and it was printing wirelessly!! Can this be that easy? So I switched on my wife's machine..... same result. No need for any driver downloads, or any fiddling at all - it just works!

The other thing was that in each case the scanner worked flawlessly at first attempt. I'm not too sure how the system sees the scanner. At first try, I used sane, and everything went OK. Then, I tried OpenOffice and "Insert, picture, select source", then "Insert, picture, scan." To be honest, I was expecting trouble, and didn't take much notice of what I was trying. Then it basically worked and there was I just gawping! There are a few things I'll have to chase up: It only lets me choose 75 or 1200 resolution in OpenOffice. But, if I use Gimp, there's no problem - anything I like, more or less.

Also, it seems a bit slow to start from OpenOffice. Still, it works!

My only criticism of the HP is that the ink costs are much higher than the old Epson. If only they would follow Epson's lead.

The NEW PCLinuxOS Magazine

Created with Scribus 1.3.5

Clockwise, from upper left:

Posted by T6, December 1, 2009, running KDE 4.3.4.

Posted by Pirate, December 18, 2009, running KDE 4.3.4.

Posted by Sproggy, December 23, 2009, running KDE 4.3.4.

Posted by parnote, December 17, 2009, running KDE 4.3.4.