

The PCLinuxOS magazine

Volume 87

April, 2014

*Send A Gmail Message
Without Opening Gmail*

*Inkscape Tutorial: Create A
Paisley Pattern, Part One*

Keep Your System COOL

*Back To School With
Free Linux Intro Course*

*Handy Utilities To Organize
Your Life, Part One*

*PC Remote Control
With TeamViewer*

Hard Drive Care & Repair

Hard Drive Failure – And Recovery

*Programming With
Gtkdialog, Part One*

Making Files Immutable

KDenLive, Part Five

And more inside!

PCLinuxOS
Goes
To
Work



Table Of Contents

- 3 Welcome From The Chief Editor**
- 4 Hard Drive Care & Repair**
- 6 Screenshot Showcase**
- 7 Hard Drive Failure – And Recovery**
- 12 Screenshot Showcase**
- 13 PCLinuxOS Recipe Corner**
- 14 ms_meme's Nook: Happy PCLOS Wanderer**
- 15 PCLinuxOS Family Member Spotlight: OnlyHuman**
- 16 Computer Hardware 101, Part Two: Cooling Systems**
- 33 Inkscape Tutorial: Create A Paisley Pattern, Part One**
- 35 Screenshot Showcase**
- 36 Send A Gmail Message Without Opening Gmail**
- 44 Making Files Immutable**
- 45 PCLinuxOS Goes To Work**
- 48 Screenshot Showcase**
- 49 PCLinuxOS Puzzled Partitions**
- 52 Programming With Gtkdialog: Part One**
- 56 PC Remote Control With TeamViewer**
- 58 Handy Utilities To Organize Your Life, Part One**
- 62 Screenshot Showcase**
- 63 KDenLive, Part Five**
- 65 Screenshot Showcase**
- 66 Back To School With Free Linux Intro Course**
- 67 More Screenshot Showcase**

Disclaimer

1. All the contents of The PCLinuxOS Magazine are only for general information and/or use. Such contents do not constitute advice and should not be relied upon in making (or refraining from making) any decision. Any specific advice or replies to queries in any part of the magazine is/are the person opinion of such experts/consultants/persons and are not subscribed to by The PCLinuxOS Magazine.
2. The information in The PCLinuxOS Magazine is provided on an "AS IS" basis, and all warranties, expressed or implied of any kind, regarding any matter pertaining to any information, advice or replies are disclaimed and excluded.
3. The PCLinuxOS Magazine and its associates shall not be liable, at any time, for damages (including, but not limited to, without limitation, damages of any kind) arising in contract, tort or otherwise, from the use of or inability to use the magazine, or any of its contents, or from any action taken (or refrained from being taken) as a result of using the magazine or any such contents or for any failure of performance, error, omission, interruption, deletion, defect, delay in operation or transmission, computer virus, communications line failure, theft or destruction or unauthorized access to, alteration of, or use of information contained on the magazine.
4. No representations, warranties or guarantees whatsoever are made as to the accuracy, adequacy, reliability, completeness, suitability, or applicability of the information to a particular situation. All trademarks are the property of their respective owners.
5. Certain links on the magazine lead to resources located on servers maintained by third parties over whom The PCLinuxOS Magazine has no control or connection, business or otherwise. These sites are external to The PCLinuxOS Magazine and by visiting these, you are doing so of your own accord and assume all responsibility and liability for such action.

Material Submitted by Users

A majority of sections in the magazine contain materials submitted by users. The PCLinuxOS Magazine accepts no responsibility for the content, accuracy, conformity to applicable laws of such material.


Entire Agreement

These terms constitute the entire agreement between the parties with respect to the subject matter hereof and supersedes and replaces all prior or contemporaneous understandings or agreements, written or oral, regarding such subject matter.



Welcome From The Chief Editor

During the last half of March, news about a new Linux malware exploit lit up the IT newswires. It isn't taking long for details about the exploit to come forward, either. Estimates place the number of infected web servers at between 10,000 and 25,000.



The exploit, called Windigo, has been active since 2011. Some experts claim that it utilizes vulnerabilities on web servers to inject some malicious JavaScript into websites. That malicious code sends 35 million spam emails per day from the infected servers, and exposes Windows users to drive-by malware attacks. This is reported to have infested the Linux Foundation's kernel.org servers. There are also reports that users of Linux kernel 2.6.x are most vulnerable.

Instead of relying on technical vulnerabilities, Windigo uses stolen credentials to gain access to servers. If you run a web server, simple password authentication is inadequate, and administrators should rely on two-factor authentication.

Windigo is comprised of three separate attacks. The first one, known as Linux/Ebury, which creates an OpenSSH backdoor, through which credentials can be stolen and take control of the servers. The second one, known as Linux/Cdorked, is a HTTP backdoor that redirects visitors to malicious software exploits and fraudulent content. The third one, known as Perl/Calfbot, causes the infected machines to send spam.

If you run a web server and want to know if it is infected, run the following code at the command line:

```
ssh -G 2>&1 | grep -e illegal -e unknown > /dev/null && echo "System clean" || echo "System infected"
```

If your server is infected, it's highly recommended that you completely reinstall the system, given how difficult this malware has been proven to remove. Also, it's recommended that you consider any currently used credentials compromised, and change them on the new installation.

If your server isn't yet infected, it's highly recommended that you update to the latest stable kernel that is available to you.

Lessons?

Certainly, we all can learn something by this latest attack. First, it emphasizes why you should never run your machine under the root account. Running under the root account, you expose yourself to any attacks, as rootkits and malware have an easy route to install themselves on your system. Second, keep your kernel updated to the latest version that runs on your computer hardware configuration. Running the latest kernel ensures that you have the latest security updates, plugging any "holes" where malware may gain entry. Third, avail yourself of some of the security tools that Linux provides, such as rkhunter, to seek out any rootkits that may try to infect your machine.

It's no doubt that Linux is infinitely more secure than Windows. That's something that every Linux user already knows. But no one, no where, can ever make an operating system that's 100% secure. There will always be vulnerabilities. I'm just happy that I run Linux, where viruses, malware, spyware and all the other types of "wares" are the exception, rather than the rule.

Until next month, I bid you peace, serenity, happiness, prosperity ... and security.

The PCLinuxOS magazine

The PCLinuxOS name, logo and colors are the trademark of Texstar.

The PCLinuxOS Magazine is a monthly online publication containing PCLinuxOS-related materials. It is published primarily for members of the PCLinuxOS community. The magazine staff is comprised of volunteers from the PCLinuxOS community.

Visit us online at <http://www.pclosmag.com>

This release was made possible by the following volunteers:

Chief Editor: Paul Arnote (parnote)
Assistant Editor: Meemaw
Artwork: Timeth, ms_meme, Meemaw
Magazine Layout: Paul Arnote, Meemaw, ms_meme
HTML Layout: YouCanToo

Staff:

ms_meme	loudog
Meemaw	YouCanToo
Gary L. Ratliff, Sr.	Pete Kelly
Daniel Meiß-Wilhelm	Antonis Komis
daiashi	Patrick Horneker
Smileeb	

Contributors:

horusfalcon	Phil
-------------	------

The PCLinuxOS Magazine is released under the Creative Commons Attribution-NonCommercial-Share-Alike 3.0 Unported license. Some rights are reserved.
Copyright © 2014.



Hard Drive Care & Repair

by Phil

The hard drive is a modern mechanical marvel. The disk platters are coated with ferromagnetic material and rotate at around 7,000 rpm. The head clearance to the disk is 20-30 nanometres, and they can extract files at around 1Mb per second. Because they are mechanical devices, they will certainly fail. If you are lucky, you will get warnings and can do something. Otherwise, you may have a catastrophic, instant fail and everything on that drive may be lost.

Hard drives are a key component to a system. It's lucky for Linux users that the hard drive is regularly checked and maintained by the system. You can expect a typical hard drive to last around 5 years. Some will fail quickly, and others seem to last for decades (there are plenty of 10 year old drives in use).

Your hard drive will fail. If you have valuable files, you should back them up to another drive, or elsewhere. If they are really valuable, back them up multiple times, and make sure you check that your

backed up files really do exist. Furthermore make sure important backups are completely separate from your system, as malicious crackers and pedantic machines can and will wantonly obliterate backups for their amusement.

What to Look for from a Failing Drive

Your system will boot up slower than usual, errors may be shown when doing so, you may hear strange noises, or no noise at all. Files may no longer be readable, or you may not be able save a file. The disk may not be mountable, or the system may say the disk does not exist. It may take much longer than usual to read or save a file, or the system starts freezing/crashing.

Monitoring and Checking Your Drive

SMART

You can check the wellbeing of your disk drives by enabling SMART. This is a monitoring system built into modern drives which keeps an eye on them, and can give a warning of a failing drive. It is very simple to do this. In Synaptic apply the following items:

```
smartmontools
gsmartcontrol
gnome-disk-utility
task-mate (the above work best in a Gnomish
desktop)
```

If you login to a MATE desktop by changing the session type on the login screen, you can review the integrity of your drives and read and order their



health reports. Early warning may be evident, or maybe your drive will crash anyway with no warning.

fsck

Your system automatically checks your drives every so often. You can check your drives and partitions at anytime with fsck, especially if you are worried or say there has been a power outage.

NOTE: The drives and partitions you are checking should be unmounted. The best way is to use a live disk, maybe one of the light ones.

Change to a root terminal (remembering to take care now that you are root).

fdisk -l (list) shows your mounted and unmounted drives, and partitions

fsck /dev/sda1 (check a partition)

fsck -fy /dev/sda1 (NOTE: check and repair which will amend the file structure without intervention. If there are a lot of errors on a partition, you may not want to fix the drive, and switch to file backup and disk recovery)

Disk/Partition Full?

Disk full? A golden oldie. It is not uncommon to inadvertently fill up your root / partition, after which your desktop or system may not boot. If, for example, your KDE desktop will not boot, change to a different session, maybe an LXDE desktop which may work, or switch to a live disk. A partition should be no more than 90% full.

On KDE More Applications > Monitoring > KDiskFree is worth a look from time to time.

To find what is filling up a disk (large file?) change to a root terminal (remembering to take care now that you are the root user).

df -al (This shows your mounted drives and how much space they have used)

If your / disk is full, it is time to play hunt the rabbit:

cd /
du -hsx * | sort -rh | head -10

Keep changing directories to trace the culprit(s). If it's just one rogue file delete it, and your job is done. If the partition is too small, ask on the forum about resizing a partition.

Looking After Your Drive

If you have a laptop, do not walk around with it or move it when it is working. Gyroscopic forces will apply strain to the drive, and with clearances in nanometres, that is not sensible.

Make sure you shutdown your system correctly. Do not just pull the plug.

Have some surge protection to protect against power spikes. Dirty power is a killer of drives, along with many other components of your system.

If you have a power outage, a brownout, or lightning strike nearby, check your drives.

Do not bang your machine, if it is vibrating, to fix the issue.

Backup your data, as your drive will eventually fail. If your files are not backed up when this happens, they will most likely be lost.

Recovery and File Extraction

If you suspect your drive is failing, run some diagnostic tests on partitions.

NOTE: when dealing with disk drives and partitions, it is imperative that they are unmounted, particularly the root / partition. Use a live disk with a light desktop.

Here is a rough checklist:

For a list of your disks, enter **fdisk -l** at a command prompt to list them.

For each suspect partition, enter the following command at a command prompt:

fsck /dev/sdNn (eg /sda1)

If it has a lot of issues, then immediately back up the files, if you can.

If there are only a few issues, backup your files and try to fix the drive errors, entering the following command at a command prompt:

fsck -fy /dev/sdNn (eg /sda1)

If you get an error on "superblocks," enter the following command at a command prompt:

fsck -fy -b superblock /dev/sdNn

Superblock backups stored on blocks:

```
32768, 98304, 163840, 229376, 294912,
819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000,
23887872
```

If the partition is badly mangled and you cannot get any files off, then more invasive measures are needed.

Make an image of the partition, using clonezilla, dd, or the Gnome disk utility.

Check the cable connections to the drive, and perhaps change them. Check the memory with memtest, which is one of the boot options.

In a root terminal run testdisk:

<http://www.cgsecurity.org/wiki/TestDisk>

This will attempt to salvage partitions on a disk. It will make changes to the entire drive if you allow it, so be warned. If you are lucky, you may get your drive back, in which case extract all your valuable files now.

If you have a failed drive and wish to try to extract any files from it, your last resort is photorec. This will search through a broken drive for files (you can refine its search). It dumps everything it finds into a folder of your choice. The result is a mess of unnamed files which you need to sort through (filtering by size and extension helps), and then rename the file to whatever you want.

You can try to format and partition the broken drive to put it back into working order. PCC may work, otherwise fdisk and mkfs.ext4 will help. A disk which has failed will always be deemed to be dubious, so do not rely on any such disk.

Looking for an old article?
Can't find what you want? Try the

PCLinuxOS Magazine's
searchable index!

The **PCLinuxOS** magazine




linuxfordummies.org

There Are No Stupid Questions

PCLinuxOS All your PCLinuxOS
Connect connections in one convenient location!

Screenshot Showcase



Posted by weirdwolf, March 16, 2014, running LXDE.

Hard Drive Failure – And Recovery

by Phil

I have a typical desktop box running PCLinuxOS. It has two drives: sda is the main workaday drive, and sdb is the backup drive which is brought up maybe once a day via a cron job. Both are about 3 years old and are the same size, both at 500GB. Of note, sdb was partitioned into four separate partitions, on the basis (hope?) that if one partition failed the others would remain functional.

One day, while logging out of a second KDE desktop session to my main session, I noticed the session had crashed and lots of error messages were scrolling up the screen. The desktop was unavailable, so I rebooted using the REISUB technique.

The machine rebooted to the login screen, but KDE would not boot, only getting as far as the globe. After several attempts, I changed the session to LXDE, which booted without issue (Lesson - have another desktop installed).

Having a non-booting desktop triggered a recollection that this might be a full root / partition. I have done that before. So:

df -al

(Alternative is More Applications > Monitoring > KDiskFree)

This said / was 100% full.

The solution is to open a root terminal and then hunt for the problem file(s) with:

du -hsx * | sort -rh | head -10

There were many directories copied over from elsewhere, I am guessing sda. I do not recall copying over all those directories, and the source was from different partitions. I am guessing that as I also have a failing mind-of-its-own-mouse, I may have accidentally copied directories over in a root Dolphin session.

I purged all files that should not have been in the / sda1 partition and logged back into a fully functional KDE session.

From here, I decided to check all my partitions and drives. First I used Dolphin to look at partitions. Drive sda was in the clear. When trying to look at sdb partitions with Dolphin, this error message was displayed:

An error occurred while accessing '122.5 GiB Hard Drive', the system responded: The requested operation has failed: Error mounting /dev/sdb5 at /media/044df901-45ee-4986-b59c-ecbb6b3ac508: Command-line `mount -t "ext4" -o "uhelper=udisks2,nodev,nosuid" "/dev/sdb5" "/media/044df901-45ee-4986-b59c-ecbb6b3ac508"' exited with non-zero exit status 32: wrong fs type, bad option, bad superblock on /dev/sdb5, missing codepage or helper program, or other error In some cases useful info is found in syslog - try dmesg | tail

I then tried **fdisk -l** (list, as root), and extracted the information for sdb:

```
Disk /dev/sdb: 500.1 GB, 500107862016 bytes, 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

I then decided to check and try to repair all partitions with fsck. To do this, all partitions being checked must be unmounted. Therefore I used a "live" LXDE system (from a Live CD or Live USB).

The fsck command is as below. Drive sda was fine, but all partitions in sdb had many inode errors:

fsck /dev/sdN (This checks a whole drive, no changes unless instructed)

fsck /dev/sdNx (This checks a partition, no changes unless instructed)

fsck -fy /dev/sdNx (This checks and repairs a partition without referral, so exercise caution)

Normally, fsck takes very little time to check a partition of drive. This happens at boot from time to time, which is why your drives are usually in good condition. Of particular note, two of the sdb partitions took a few minutes to complete, and the two other partitions took well over an hour each, with a vast number of errors. When done:

fdisk -l /dev/sdb

```
Disk /dev/sdb: 500.1 GB, 500107862016 bytes, 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1	*		63	256558049	128278993+	83 Linux
/dev/sdb2		256558050	976768064	360105007+		5 Extended
/dev/sdb5		256558113	513405269	128423578+		83 Linux
/dev/sdb6		513405333	770204294	128399481		83 Linux
/dev/sdb7		770204358	976768064	103281853+		83 Linux

However, the error message "an error has occurred" persisted when trying to look at the partitions with Dolphin.

Next was to run testdisk as root on sdb. It is an excellent utility:

<http://www.cgsecurity.org/wiki/TestDisk>

This resulted in sdb becoming even more mangled. Fdisk said:

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		63125685	976768064	456821190	f	W95 Ext'd (LBA)
/dev/sdb5		63125701	78122817	7498558+	83	Linux

Next, try to repair by using a different superblock. From the forum:

When a file system gets corrupted, it uses the superblock to repair the filesystem. If the main superblock is itself corrupted, there are backup superblocks one can use to do the repair. In a default ext4 filesystem, the superblocks are in these locations.

Superblock backups stored on blocks:

```
32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632, 2654208,
4096000, 7962624, 11239424, 20480000, 23887872
```

To use these in a filesystem check, the command would be:

```
fsck -fy -b 32768 /dev/sd<whatever> <Enter>
```

Replace <whatever> with the actual partition designation, and 32768 with each backup superblock listed, until you find one that isn't corrupted and works to restore the filesystem. Part of the restoration is the repair of any damaged superblocks, including those you may have also tried, but didn't work. You have multiple chances to fix your filesystem, on each affected partition. Don't give up until you've tried them all.

The result was no change. All superblocks were tried.

Next I tried to use photorec to recover a clonezilla image of a machine which I keep an eye on. Photorec is an excellent utility which will recover files off a broken drive if it can find them, and can be tweaked to search out particular file types. However, the recovered files have a numeric number as a name, and you have to go through recovered files to rename them. A painful and long process if you have a lot of files to go through.

I installed SMART tools and associated programmes, and from a MATE session, SMART says sdb is good.

```
SMART Error Log Version: 1
No Errors Logged
```

```
SMART Self-test log structure revision number 1
```

Num	Test_Description	Status	Remaining	LifeTime(hours)
	LBA_of_first_error			
# 1	Extended offline	Completed without error	00%	11303
# 2	Short offline	Completed without error	00%	11302

I now gave up on the recovery process, having damaged the file system, and try to partition and format the drive with PCC, gparted, and fsck (with mkfs to format a created partition).

The result is one partition which will not mount.

Summary

I have attempted to repair sdb using anything I could think of, with some suggestions from the forum applied. (I am always amazed and grateful for the expertise available on tap at the forum). The net result of my meddling was a broken sdb drive and no recovered files. All data except one directory and clonezilla image files are on the primary drive, so there is a loss but not a great loss. At this stage, I am convinced the drive is broken and needs replacing.

Recovery (Under Remote Control)

From here, I reached out to the forum for help on how to partition a drive from scratch. Under remote control, I applied commands as directed, posted back the result, and awaited orders. I have to sit on my hands and not do anything, which is very hard to do.

I check the memory for errors using memtest from the grub screen options at boot. The memory is good.

I zero the MBR:

```
dd if=/dev/zero of=/dev/sdb bs=512 count=1
```

```
1+0 records in
1+0 records out
512 bytes (512 B) copied, 0.000417889 s, 1.2 MB/s
```

(Addition note for information):

"I only wish to add that if the drive was using GUID (GPT) at some point, then it has a second instance of the partition scheme information (at the end of the drive), and this too needs to be wiped."

Zero the drive:

```
dd if=/dev/zero of=/dev/sda bs=4096 <Enter>
```

This takes many hours to complete. Leave it to run.

Open the box, replace the cables (Hardware fault?).

Run `fsck -fyc /dev/sdb1`

`fdisk -l /dev/sdb`

```
Disk /dev/sdb: 500.1 GB, 500107862016 bytes, 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x68b42369
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		2048	976773167	488385560	83	Linux

```
fsck -fyc /dev/sdb1
fsck from util-linux 2.22.2
e2fsck 1.42.9 (28-Dec-2013)
Superblock has an invalid journal (inode 8).
Clear? yes
```

```
*** ext3 journal has been deleted - filesystem is now ext2 only
***
```

```
Resize inode not valid.  Recreate? yes
```

```
ext2fs_block_iterate: Ext2 file too big while sanity checking the
bad blocks inode
```

```
/dev/sdb1: ***** FILE SYSTEM WAS MODIFIED *****
```

Try this command:

```
mke2fs -t ext4 /dev/sdb1 <Enter>
```

Result:

```
mke2fs -t ext4 /dev/sdb1
mke2fs 1.42.9 (28-Dec-2013)
Could not stat /dev/sdb1 --- No such file or directory
```

`Fdisk - Device sdb1 no longer exists.`

`Fdisk following day:`

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		2048	976773167	488385560	83	Linux

Instruction - To test for hardware fault remove the drive from the case and reattach as a USB device. Run `fdisk -l`:

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		2048	976773167	488385560	83	Linux

Instruction

I want you to try creating a new filesystem on `/dev/sdb1`, now that it's entirely free of the original computer.

```
mke2fs -t ext4 /dev/sdb1 <Enter>
```

Once the process completes, do these commands;

```
mkdir -p /mnt/here <Enter>
```

```
mount /dev/sdb1 /mnt/here <Enter>
```

...and if no errors up to this point:

```
ls -l /mnt/here <Enter>
```

Post your results.

Results:

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		2048	976773167	488385560	83	Linux

```
[root@localhost philip]# mke2fs -t ext4 /dev/sdb1
mke2fs 1.42.9 (28-Dec-2013)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
30531584 inodes, 122096390 blocks
6104819 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=4294967296
3727 block groups
32768 blocks per group, 32768 fragments per group
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736,
    1605632, 2654208, 4096000, 7962624, 11239424, 20480000,
    23887872, 71663616, 78675968, 102400000
```

```
Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done
```

```
[root@localhost philip]# mkdir -p /mnt/here
```

```
[root@localhost philip]# mount /dev/sdb1 /mnt/here
mount: wrong fs type, bad option, bad superblock on /dev/sdb1,
missing codepage or helper program, or other error
In some cases useful info is found in syslog - try
dmesg | tail or so
```

Instruction

Try:

```
dmesg | tail <Enter>
```

Result:

```
# dmesg | tail
sd 13:0:0:0: [sdb] No Caching mode page found
sd 13:0:0:0: [sdb] Assuming drive cache: write through
sd 13:0:0:0: [sdb] No Caching mode page found
sd 13:0:0:0: [sdb] Assuming drive cache: write through
sdb: sdb1
sd 13:0:0:0: [sdb] No Caching mode page found
sd 13:0:0:0: [sdb] Assuming drive cache: write through
sd 13:0:0:0: [sdb] Attached SCSI disk
EXT4-fs (sdb1): ext4_check_descriptors: Checksum for group 0
failed (2659!=31464)
EXT4-fs (sdb1): group descriptors corrupted!
```

Instruction:

Open fdisk, and use the d command to delete the sdb1 partition, then enter the n command to create a new partition. Accept the primary designation, and partition number, but for a start sector, use 4096. then accept the default for the last sector. Finish with the w command, then do an fdisk -l again.

Run the commands:

```
mke2fs -t ext4 -c /dev/sdb1 <Enter>
```

The -c argument does a bad block check before actually writing the file system.

```
mount /dev/sdb1 /mnt/here <Enter>
```

Starting the partition at a different location on the drive avoids the location of the main superblock that has constantly failed, without losing an appreciable amount of overall space. If that is the only truly bad block causing all your problems, this should work. If the drive is riddled with bad blocks, and continues to fail to create usable filesystems, I'd be looking for a new one.

If the partition actually mounts:

```
ls -l /mnt/here <Enter>
```

Input:

fdisk /dev/sdb

```
Command (m for help): d
Selected partition 1
Partition 1 is deleted
```

```
Command (m for help): p
```

```
Disk /dev/sdb: 500.1 GB, 500107862016 bytes, 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x68b42369
```

Device	Boot	Start	End	Blocks	Id	System
--------	------	-------	-----	--------	----	--------

```
Command (m for help):
```

```
Create new partition:
```

```
Command (m for help): n
```

```
Partition type:
```

```
  p primary (0 primary, 0 extended, 4 free)
  e extended
```

```
Select (default p):
```

```
Using default response p
```

```
Partition number (1-4, default 1): 1
```

```
First sector (2048-976773167, default 2048): 4096
```

```
Last sector, +sectors or +size{K,M,G} (4096-976773167, default
976773167):
```

```
Using default value 976773167
```

```
Partition 1 of type Linux and of size 465.8 GiB is set
```

```
Command (m for help): w
```

```
The partition table has been altered!
```

```
Calling ioctl() to re-read partition table.
```

```
Syncing disks.
```

fdisk -l /dev/sdb

```
Disk /dev/sdb: 500.1 GB, 500107862016 bytes, 976773168 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x68b42369
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		4096	976773167	488384536	83	Linux

NOTE - This next part took many hours

```
mke2fs -t ext4 -c /dev/sdb1
```

```
mke2fs 1.42.9 (28-Dec-2013)
```

```
Filesystem label=
```

```
OS type: Linux
```

```
Block size=4096 (log=2)
```

```
Fragment size=4096 (log=2)
```

```
Stride=0 blocks, Stripe width=0 blocks
```

```
30531584 inodes, 122096134 blocks
```

```
6104806 blocks (5.00%) reserved for the super user
```

```
First data block=0
```

```
Maximum filesystem blocks=4294967296
```

```
3727 block groups
```

```
32768 blocks per group, 32768 fragments per group
```

```
8192 inodes per group
```

```
Superblock backups stored on blocks:
```

```
    32768, 98304, 163840, 229376, 294912, 819200, 884736,
    1605632, 2654208, 4096000, 7962624, 11239424, 20480000,
    23887872, 71663616, 78675968, 102400000
```

```
Checking for bad blocks (read-only test): 0.00% done, 0:00
elapsed. (0/0/0 errorsdone)
```

```
Allocating group tables: done
```

```
Writing inode tables: done
```

```
Creating journal (32768 blocks): done
```

```
Disk /dev/sdb: 500.1 GB, 500107862016 bytes, 976773168 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

```
Sector size (logical/physical): 512 bytes / 512 bytes
```

```
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

```
Disk identifier: 0x68b42369
```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		4096	976773167	488384536	83	Linux

```
mount /dev/sdc1 /mnt/here
```

```
ls -l /mnt/here
```

```
total 16
```

```
drwx----- 2 root root 16384 Mar 10 00:12 lost+found/
```

SUCCESS!

Additional information:

In order to be able to write to that partition, with the greatest ease, while the partition is mounted on /mnt/here:

```
chown -R user:user /mnt/here <Enter>
```

Synopsis

A hard drive is comprised of sectors of 512 bytes (also called units by fdisk) onto which your data is saved. The sectors are organized into blocks, which may be of 4096 bytes which are organised by the onboard drive firmware.

Somewhere between unit 2048 and 4096 something has been damaged so that a read/write to this part of the disk fails, due to a bad block. By avoiding the bad block, the disk now works.



Screenshot Showcase



Posted by deckard313, March 29, 2014, running MATE.

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 ...

PCLinuxOS
Radically Simple

PCLinuxOS Recipe Corner



From The Kitchen of
You Can Too

Caramel Apple Cookie Bread

Ingredients:

½ cup + 2 Tbsp. softened butter. The 2 tablespoons are for the caramel apples...
2 small to medium apples peeled and diced
(I used gala apples)
¼ cup brown sugar
1 tsp. vanilla
2 large eggs
1 cup white sugar
1¾ cup all-purpose flour
½ tsp. salt
1 tsp. baking soda

Recipe Directions:

Preheat oven to 350 degrees. In a small sauce pan place 2 tablespoons of butter ¼ cup sugar and diced apples. On medium to low heat simmer apples for 3-5 minutes or until apples are tender but firm. In large bowl cream the rest of the sugar and butter, slowly add eggs and vanilla to the mixture. Once the mixture has been creamed sift flour, baking soda and salt; mix dry ingredient into the wet ingredients.

Fold caramel apples into the batter so the batter becomes moist and sticky. Once done place batter into a greased loaf pan and put in the oven. The bread should bake for an hour and twenty minutes (I have a skinny loaf pan so it takes longer to bake, if your loaf pan is wider than mine it make not take as long) or stick a toothpick in the middle. If it comes out clean the bread is done. Let cool completely. Once cool take out the loaf pan and serve.




A collection of PCLinuxOS merchandise. On the left side, there is a white polo shirt, a white mug with the PCLinuxOS logo and the slogan 'radically simple', a circular logo with the PCLinuxOS logo and slogan, and a square logo with the PCLinuxOS logo and slogan. On the right side, there is a black t-shirt, a tan baseball cap with the PCLinuxOS logo and slogan, white shorts with the PCLinuxOS logo, and a white t-shirt with the PCLinuxOS logo and slogan.

Support PCLinuxOS! Get Your Official
PCLinuxOS
Merchandise Today!

ms_meme's Nook: Happy PCLOS Wanderer

I love to go a wandering along the internet
And as I surf I love to sing about the best OS yet

PCL PCL PCL PC LO LO LO
PCL PCL PCLOS




My hat goes off to Texstar he is just the most
He's made our lives rich and full PCLOS is grandiose

PCL PCL PCL PC LO LO LO
PCL PCL PCLOS

MP3

I love to wander to the Forum a happy place to be
To see the folks and read the jokes they all tickle me

PCL PCL PCL PC LO LO LO
PCL PCL PCLOS



May I learn the Linux code before the day I die
They all say it's good to know but I sure don't know why

PCL PCL PCL PC LO LO LO
PCL PCL PCLOS

OGG



PCLinuxOS Family Member Spotlight: OnlyHuman

as told to Smileeb

This month, we get the chance to learn more about OnlyHuman.

How old are you?

I am 61.

Married, single or what?

I'm divorced.

Children, grandchildren?

I have one child.

Retired or working and for how long and at what?

Retired because of ill health. My condition limits me in doing many things I would like to do.

What is the area you live in like. Weather, Quietness, Scenery?

I live in the UK England in the southeast, it is country-like, quite close.

Are you handy with your hands and have any hobbies?


I used to be fairly handy, but not now. Linux on the PC keeps me busy.

What is your education level?

Nothing much to speak of.

Do you like to travel, go camping?

It's fairly quiet here, but I don't go out much. I play a little on guitar just for fun.



What caused you to try Linux and join this forum?

I joined the PCLinuxOS forum years ago, in January 2009, as I fancied a change from using Windows. I tried a few distros very briefly, but PCLinuxOS seems to detect my hardware best, so stuck with PCLinuxOS. I made my personal PCLinuxOS e17 remasters.



I stay in mostly, though in summer I sometimes strum my guitar sitting in my garden.




Donate To PCLinuxOS

*Community Supported.
No Billionaires/Millionaires.
No Corporate Backing Or Funding.*

Click [here](#) to make a one-time donation through Google Checkout.

Or, click one of the amounts down below to make a monthly, recurring donation.



Computer Hardware 101, Part Two: Cooling Systems

by horusfalcon

SAFETY FIRST!

When working inside a computer chassis, ALL work except for testing and observation should be done with **POWER OFF!** Observe common-sense safety precautions and wear eye protection whenever a computer is running for observation or testing with a cover removed.

Overview

Computer systems generate heat during operation - it is unavoidable. Removing that heat from critical components keeps a system running at peak efficiency and preserves the life of the system. Failing to cool these components properly can result in degraded performance, possibly also leading to damage to the system, or even system failure.

For these reasons, it is wise for a computer user to have at least a basic understanding of how cooling systems for computers work. This article will cover:

- An overview of the types of cooling systems commonly available,
- Common system hardware profiles, and how they affect cooling system designs,
- Design considerations for new or existing cooling systems,
- Hardware selection, installation, and maintenance tips, including some brief remarks on system modifications (modding) for improved performance.

Types of Cooling Systems

The three basic categories to be considered here are:

- Liquid Cooling,
- Solid-State Cooling, and,
- Air Cooling

Cryogenic cooling, that is, the use of liquified gases to achieve super-cold temperatures, is not within the scope of this article. It requires advanced methods, materials, and techniques. Attempting it is fraught with personal risk and is not recommended without proper training and equipment.

Author's Note:

I have limited exposure to liquid and solid-state cooling systems, and, so, will only be discussing them to the extent of my knowledge. Air cooling systems are more common, and will be discussed to greater depth primarily because I know more about them, and because they have greater applicability to more types of systems.

Liquid Cooling

This type of cooling system uses a special heatsink called a *cooling block* (or *water block*) in which some sort of liquid (usually water-based) *coolant* is circulated to remove heat. The coolant is then carried by a pump to a reservoir and heat exchanger (a radiator or chiller) where it is cooled and returned to the system again to remove more heat.

Benefits: The benefits of liquid cooling are quieter (possibly even silent) operation, high heat removal

capacity (which results in lower operating temperatures at the cooled devices), and separation of airflow from the case (resulting in a cleaner case interior with far less dust).

Drawbacks: So why aren't liquid cooled systems more popular? To gain these benefits, there are some drawbacks to be accepted. Liquid cooling systems are more complex, with more moving parts which can fail. Failures in a liquid cooling system tend to have more rapid onset and more serious consequences, so they require closer monitoring and more maintenance to keep them working properly. The water pump, in particular, represents a single point of failure in most designs.

They are also more expensive, and require more care to install them properly to prevent leaks (which can cause system damage).

Finally, the heat exchanger is typically a fairly large item, and requires more power to operate. This raises the system's overall electric power consumption, and increases its physical footprint.

For these reasons, and primarily to control costs, most computer vendors will not opt for liquid cooling on any but their most powerful systems.

Unless your system (or your local climate) truly requires it, liquid cooling will cost more than it is typically worth. For systems that do require it, nothing else works quite as well. The decision as to whether to use liquid cooling should be based *on a serious evaluation of the system under consideration*, else the money spent will be largely wasted.

Solid-State Cooling

Solid-state cooling relies mainly on something called the [Peltier Effect](#). This was first observed in 1823 by Jean Charles Peltier, a French physicist. As he applied voltage across a junction of dissimilar materials, he observed that one side got hot while the other cooled down.

So much for the historical background. A Peltier Effect Device (PED) is, in effect, an array of thermocouples being used “in reverse” to pump heat away from a heat source.

Benefits: The main advantages of using a PED for cooling are, firstly, that it can achieve temperatures at the cooling junction *much lower* than any commonly used method, be it air or liquid cooling, and, secondly, that it has no moving parts. This can be of benefit to those who choose to overclock their systems. In combination with a big enough secondary heatsink, a PED can also be used to create a nearly silent cooling system.

Drawbacks: So, what's not to like? PEDs are not cheap nor are they particularly efficient. They must be closely matched to the equipment they are cooling for anything resembling economical operation.

In addition, because the devices can produce very low temperatures at the cooled device, water can condense on the board if voltage to the unit is not properly controlled (leading to electrical problems). This requires either some provisions to supply a controlled voltage to the unit in response to changes in temperature at the cooled device to make it self-regulating, and/or a method to remove condensation. This adds expense and complexity to the system. Some systems also resort to using desiccants to control humidity, which means an additional maintenance commitment.

As the hot side of the device rises in temperature, it must be cooled or it will become overwhelmed and

the PED will be damaged. This, in turn, will lead to damage to the chip beneath it as the failed device will actually continue to transfer residual heat into it after it has shut down. Close temperature monitoring of solid-state cooling systems is, thus, essential for predicting and intercepting failures. More attention to maintenance of secondary cooling systems is also required to ensure their continued reliability.

Air Cooling Systems

Air cooling systems should be more familiar to most readers for the simple reason that computer vendors and builders usually default to this method of cooling. The core components of an air cooling system are the heatsinks and cooling fans. What could be simpler?

Benefits: Simplicity is one of the main advantages to air cooling. The components are relatively simple and few. Fans, heatsinks, some fastening hardware, and some sort of thermal transfer medium are really all that is required.

All these components are inexpensive and readily available at a variety of vendors online, and even locally in some cases. It does not require much more than ordinary care, and the ability to read and follow simple directions to install those components correctly.

If upgrades or modifications to a system become necessary to improve its cooling capacity, these are relatively easy to do in stepwise fashion, so cost can be contained, and each step can be evaluated before moving on.

Air cooling systems are versatile and flexible. Components can be easily configured and reconfigured, airflows diverted via inexpensive shrouds and baffles, and small, noisy fans can be easily replaced with larger, more slowly turning fans that move more air quietly. A wide variety of hardware is available to meet any system need.

Drawbacks: So what are the drawbacks? The greatest of these is the tendency for air cooled systems to collect dust, smoke, pet dander, and other airborne debris. *These are eventual death to a computer*, as they coat fan blades, heatsink fins, and other exposed surfaces inside the computer. Periodic maintenance to remove foreign matter *is a must* to bust the rust and dust.

Air cooling does not have nearly the heat capacity liquid cooling does, and cannot achieve temperatures as low. Some advanced designs (phase change heatsinks, air chiller units, and others) can get very close, but, as with most things thermodynamical, sooner or later one reaches a point of diminishing returns, and liquid cooling or adding solid-state cooling becomes a more attractive option.

The use of fans also brings some drawbacks: we've already covered how dust and tar from smoke coat fan blades (which then must be cleaned). Fans also represent the greatest noise sources in a computer system. There are ways to mitigate the noise to acceptable levels by proper sizing, mounting, selection, and control, but truly silent operation is more difficult to achieve while still cooling properly.

Fans also represent single points of failure for the devices they cool. For this reason, it is wise to monitor fan conditions closely for signs of impending failure, and replace them when needed.

Finally, axial fans (the type usually used in a computer chassis) have a dead spot in their airflow just downstream of the fan blade hub. Using ducts, shrouds, and baffles can mitigate these dead spots, though.

System Profiles

The type of system a user has will, in large part, determine how free that user is to modify or upgrade their cooling system. It will also help determine a

best path forward for keeping their system's cooling capabilities up to snuff. For these reasons, we will consider the following system profiles.

Profile: Laptop

- Least capable of supporting modifications, hardest to maintain or upgrade.
- Heat Load: relative to removal capacity, can be very high.
- Best path:
 - + Keep system clean as possible to minimize need for maintenance or repairs. Internal inspection and cleaning not generally recommended except by trained and qualified technicians.
 - + Blow out the airflow path frequently to keep it clear.
 - + Use monitoring tools to keep up with system condition.
 - + Perform maintenance when indicated by monitoring tools.
 - + Use OEM or OEM-approved components for best results - DO NOT "go cheap" on laptop components, and do not attempt to "refurbish" damaged components

Profile: Desktop

- Moderate capability for modification or upgrade, and upgrades are relatively easy to do, depending on level of modification.
- Heat Load: relative to removal capacity, moderate to high.
- Best path:
 - + Evaluate existing system to form a baseline. If found inadequate, add capacity to improve system behavior.
 - + Keep system clean as possible and clean periodically to maintain at peak efficiency.
 - + Repair/replace failing components during early phases of failure. This is greatly preferable to waiting until they fail completely (*system damage can result!*)
 - + Select quality cooling system components - resist the urge to "go cheap" here.

Profile: Workstation

- Capability for upgrade or modification varies by vendor, and by the ingenuity of the technician. Factory systems are usually well-designed and will perform well enough if properly maintained.
- Heat Load: relative to removal capacity, usually moderate to very high, depending on particulars of the use case.
- Best Path:
 - + Evaluate the existing system. If found inadequate, consider cooling system upgrade through OEM vendor first before going to the aftermarket.
 - + Keep system as clean as possible to maintain system at peak efficiency. Periodic inspection and cleaning is recommended.
 - + Repair/replace failing components during early phases of failure - minimize possibility of system damage to high-end components.
 - + Use durable, reliable components (spend more for parts here). OEM or OEM-approved solutions are best in most cases.

Profile: Server

- Very little capability for upgrade or modification in most cases. Factory system is usually well-designed and will perform well if maintained properly.
- Heat Load: relative to heat removal capacity, usually moderate to very high, depending on particulars of the use case. Most designs presume operation in a controlled environment.
- Best Path:
 - + Evaluate the existing system. Seek OEM or approved components first for any required repair or upgrade.
 - + Keep system as clean as possible to maintain system at peak efficiency. Periodic inspection and cleaning is recommended.
 - + Utilize system monitoring tools and remote system monitoring capability if in an Enterprise environment to track system conditions.
 - + Perform repairs using OEM or approved

components during initial phases of failure. Be aware of any failover or outage provisions necessary in an Enterprise environment.

Profile: Gaming Desktop System

- Upgradability depends on many factors. In all cases, use high-quality and properly rated system components whenever possible.
- Heat Load: higher than most systems relative to usual heat removal capacity.
- Best Path:
 - + Ensure system design is adequate before placing system under heavy load. If upgrade or repair to system is required, seek high-quality, high capacity, high reliability components and install properly for best results.
 - + Keep system as clean as possible to maintain system at peak efficiency. Periodic inspection and cleaning is recommended.
 - + Utilize system monitoring tools and investigate at first sign of adverse change in system operating profile.
 - + Perform any repairs needed during the initial phases of failure.

System Design Considerations

Next we examine those factors which can influence the type and design of a cooling system for a computer. We'll start by looking at the factors contributing to a system's heat load.

Estimated Heat Load

There is typically a "short list" of components which contribute to the heat load of all systems:

- **CPU Type, Quantity, & Rated Heat Dissipation requirements**

Regardless of what type CPU, or who makes it, the best place to find information about it is from the

manufacturer's literature. The rated heat dissipation is usually expressed as a number of watts. The acronym used on most data sheets is TDP (Total Dissipated Power).

- **Support Chip Cooling (North/Southbridge, regulators, etc.)**

The same holds true here - the chip's data sheet is the best place to find out how best to cool a support chip.

- **RAM Type, arrangement, and cooling**

This can be a bit more of an art than a science. Most RAM vendors don't publish readily available data on heat dissipation for their products. Heat spreaders (specialized heat sinks made just for RAM modules) should be selected with an eye toward physical compatibility with the RAM modules in use, else more harm than good will come of their use.

RAM modules can be purchased with heat spreaders pre-installed. This is the best way to ensure the combination is properly installed, matched, and effective.

Ducting or shrouds may be used to direct airflow over RAM modules operating at high temperatures. Much of the airflow path will depend on the internal layout of the components.

- **Video Card Cooling Requirements**

A great many video cards come with pre-designed, pre-built cooling systems. Data on these, available from the manufacturers, is useful primarily in seeing how much total airflow is needed to keep GPUs cool. Custom cooling solutions for these are generally expensive. It's usually easier and simpler to ensure the video gear used has decent cooling built in, and to keep it clean and in good operating condition.

That said, certain very high performance video cards are good candidates for water cooling. Don't overlook this option if it presents the best protection for high-end gear.

- **Power Supply Unit Contributions to Heat Load**

A high quality PSU that produces clean, well-conditioned power can do much to lower the overall heat load of a system. The PSU also contributes much to heat removal capacity since its fans move hot air out of the case, pulling in cooler outside air.

Conversely, a cheap noisy PSU will be a constant source of trouble as attempts are made to extract maximum performance from a system.

Estimating Total Heat Load

Estimating Total Heat Load is a matter of adding up all the various heat loads from each component in the system. Generally, a high degree of precision in this estimate is not essential to be able to properly choose components which will effectively cool a given system, but the better your estimate, the more confidence you may have in your design.

Consider a simple example: a system with a single AMD Phenom II 4100 CPU chip on a Gigabyte GA-990FXA-UD3 motherboard, 8 GB of DDR3-1600 MHz RAM (2 X 4GB modules), an MSI N610 GT video card, a single 500 GB 7200 RPM SATA hard drive, and a SATA DVD-RW drive. All of this is powered by a 550 Watt ATX12V power supply.

The estimated heat loads for such a system might be something like:

CPU: 95W
(actual value from AMD's website)
Motherboard: 100W
(guesswork - no hard data available)
RAM: 9W
(estimate based on [this page](#).)

Video Card: 29W
(based on power consumption)
Hard Drive: 5W
(based on power consumption)
Optical Drive: 15W
(guesswork - no hard data available)
Power Supply:
(not factored - PSU provides airflow)
TOTAL EST. LOAD: 253W
worth of heat to be removed by all cooling.

Now, before anybody jumps my case about confusing input power consumption with dissipated power (which produces heat), no. Here's why I'm using power consumption as a general guide: in many cases, no matter how hard you might search for it, data on heat dissipation for a component is simply not given - you have to guess. A simple way to do this is to account for all the power the device consumes. This will give a somewhat padded estimate, but that's better than lowballing it. If you wanted to be more conservative, estimate that at least 45% of input power will be converted to waste heat.

So, what does this number mean? First, it's good to make sure we understand that high heat load does not always translate to high temperatures everywhere in the system. Some portions of the system will get hotter than others because a relatively high amount of heat is produced in a small volume.

The CPU in a system provide a good example of this - there's a lot going on "under the hood" in a CPU, a good deal of power being consumed in a relatively small volume, so without a large heatsink to remove the heat, temperatures can climb very quickly.

The same is generally true for GPU (Graphics Processor Units), too. Generally, a more modern multi-core, massively parallel GPU will generate more waste heat since it is doing more faster to render images to the display very quickly, sharply, and in millions of colors.

Why did I *not* factor in the heat generated by the Power Supply Unit (PSU) into the total estimated heat load? I made a simplification: The PSU *provides more than enough airflow to remove its own heat load from the system*, and helps pull warm air from the interior of the case as well, so it is actually a net negative in terms of its heat contribution to the system.

If the PSU in use in a system is fanless and contributes heat to the interior, a guess based on its power output (or, better yet, data from the manufacturer about heat dissipation from the unit) will have to be factored in to ensure heat removal is adequate.

In some very early ATX designs, the fan on the PSU *pulled air into the case*. This is a poor design, and should be replaced with a more modern unit that exhausts heat from the case.

System Airflow

The cooling airflow for a system represents the ultimate heat sink for it. This is true even in most liquid cooling systems, for they use airflow over a radiator or heat exchanger to remove heat from the liquid returning to the system.

So... how much cooling airflow is needed to remove a given amount of heat? The physics (and the math) of heat removal can get complicated. We won't be getting into that here to any great depth because it really isn't necessary. For now realize that higher total estimated heat load translates to a need for higher airflow inside a computer's interior.

The basic factors to consider in designing the airflow in a system's case are: total air volume in the case, net airflow rate from all case and PSU fans, air exchange rate in that volume, and airflow design.

• Total Interior Air Volume

Again we enter the realm of estimation, as the interior of any given computer case is usually home to a fair assortment of components and case hardware that take up some (possibly even most) of that volume.

Start with interior case measurements. If, for instance, a case measures 18 cm wide by 46 cm deep by 60 cm tall, the total possible volume is 18 X 46 X 60 or 49,680 cubic cm. Now, subtract at least 10% of that volume for interior brackets, drive shelves, and other fixtures. This leaves us, in this example, with 49,680 - 4968 or 44,712 cubic cm. This value may be translated to liters by dividing by 1,000. 44,712 cubic cm equates to 44.712 liters.

Look up or measure the external dimensions of the PSU and calculate how much space it occupies. Say our example PSU is a standard ATX12V unit from Cooler Master, the V550, which measures 15 X 14 X 8.6 cm or 1,806 cubic cm, 1.806 liters. Subtract this from your estimated interior volume of 44.712 liters to leave 42.906 liters.

Depending on how precise you want to be with this, you could continue measuring the volumes of all the various components that get put in the case and subtracting. Standard measurements are available for hard drives, optical drives, floppies, etc., but it is easy to measure them prior to installation yourself and do the arithmetic.

For purposes of our example here, we're just going to make a crude estimate of how much volume the motherboard, drives, and other fixtures take up. The motherboard, even with its heatsinks, fans, and I/O cards will only take up an additional 10% or so of the remaining volume, leaving 38.615 liters. The drives (one HDD and one optical drive) may take up 5 cm by 12 cm by 8 cm or so for another .48 liters, leaving 38.135 liters.

So, this represents the air volume available inside the case to provide cooling for all the various loads. This brings us to the next factors under consideration: Airflow Design and Total Interior Air Flow.

• Case Airflow Design and Impacts on Total Interior Airflow

We will separate case airflow designs into three basic broad categories. Those in which more air comes into a case than goes out are termed as **positive pressure** designs. Conversely, those in which more air goes out than comes in are **negative pressure** designs. A case in which nearly equal amounts of airflow enter and leave the case is considered to be a **balanced** or **neutral pressure** design.

• Negative Pressure Designs

Negative pressure designs pull more air to the outside than they take in, and so tend to have a slightly higher exchange rate than positive pressure designs.

• Positive Pressure Designs

Positive pressure designs move more air into the case than can be exhausted. This raises pressure inside the case slightly.

There are proponents of positive pressure who believe that it results in less dust. Doubtless, with airflow exiting more slowly than incoming air enters, the air becomes slightly more dense as pressure rises, and so cooling might *theoretically* be improved. The operative question is, "By how much?". This will vary greatly with different designs.

My own preference is to build for a balanced design in which the airflow entering is very near the airflow leaving, but that's just me. Any computer, regardless of airflow design, will require periodic inspection and cleaning.

• Total Interior Airflow and Air Exchange Rate

Case air exchange rate is a function of total available interior air volume, air flow into the case, and air flow out of the case. A basic formula is:

$$R_E = ((F_{IN} + F_{OUT})/2) \times 28.32 / V$$

where R_E represents Exchange Rate, F_{IN} is incoming flow, F_{OUT} is outgoing flow and V represents the case interior air volume. This number represents the number of times per minute the air in the case is exchanged with fresh air. This formula is a rather over-simplified version of what actually happens, but will yield a number useful for comparisons.

Unhappily, *most fan airflow specifications are given in CFM* (cubic feet per minute). These will have to be converted to liters/minute. The conversion factor for this is 1.0 CFM = 28.32 liters/minute. Add up all the CFMs from the fans, take the average, then convert once by multiplying by 28.32. (We could have measured the air volume in cubic inches, converted to cubic feet, etc., ad nauseum. The method presented here seems less complicated to me, anyway.)

How do we determine incoming and outgoing flow rates? Look at the specifications for all the fans, then add up all the rated airflows of fans pulling air out of a case to find total outgoing flow. In a similar manner, add up all those fans pushing air into a case for total incoming flow. The average of incoming and outgoing represents the average case airflow.

Now someone is bound to ask the question: *If airflow in is significantly more than airflow out, won't the case pressure rise and airflow be reduced?* That's true, and it's why we're working with the average.

In designs where all the fans (the *active* airflow) are pushing or all the fans are pulling, the other end (the *passive* airflow) will be governed by how much

available area there is to admit passive airflow.

Here's an example of a system with active airflows for both incoming and outgoing. Consider our example case with a volume, V , of 38.135 liters. Mounting two 140 mm fans (one pulling in, and one pulling out, each at 66 CFM) and a PSU with a rated flow outward of 40 CFM, we have:

$$\begin{aligned} R_E &= ((F_{IN} + F_{OUT})/2) \times 28.32 / V \\ &= ((66 + 106)/2) \times 28.32 \text{ liters per minute} / (38.135 \text{ liters/case vols.}) \\ &= 63.84 \text{ case vols./minute} \end{aligned}$$

Now that's a substantial amount of air exchange. Actual air exchange may be somewhat less than this (from drag and other factors), but the number is useful in getting an idea of how much air we can move. Note that our example, since it has more air going out than coming in, is a negative pressure design.

If the CPU had a duct mounted to pull air from outside onto its fan, that fan would contribute to incoming airflow and thus alter the result by a corresponding amount. If it had a fan with 40 CFM of flow, this could be considered a balanced design.

Case Design: Baffles, and Blowholes, and Shrouds (Oh, My!)

All jokes about not being in Kansas anymore aside, the use of internal case features to help direct cool air to critical components is nothing new, but can be a challenge for those not familiar with the methods. What follows is a survey of basic airflow management features.

Blowholes: Blow holes or large vents in a case will also affect exchange rate, and case operating pressure. There are endless ways to direct airflow inside a case to optimize cooling to critical

components. These can be factory-made or done as case modifications. (A bit of skill with tools and metalworking will be needed for a DIY vent or blowhole.) If these are air inlets to a case, consider installing an air filter to trap dust before it get into the case.



Sentey CS1-1410 case exterior showing side vent and top blowhole

Blowholes differ from most vents in that they are usually in the top of a case, and use a fan to pull warm air up and out. The Sentey case pictured above is an unusual design in that the PSU mount is in the bottom, and the motherboard is mounted nearer the top.

The Intel Design Guide for Thermally Advantaged Chassis, which can be found [here](#), covers vented side covers as relate to Intel CPUs. This guide also incorporates specifications and references data from <http://www.formfactors.org/formfactor.asp> which are also good reading.

CPU Ducts (aka Chassis Air Guides): We mentioned CPU ducts above. These are typically sheet metal, or (preferably) flexible mylar or plastic ducts which connect the fan on a CPU heatsink directly to outside air via a grille or a vent hole in the

side of the case. The idea is a simple one: maximize the CPU fan's effectiveness by feeding it air that has not been warmed by other loads in the system. In this arrangement, typically the CPU fan is pulling fresh air into the case interior and pushing it down over the fins of the heatsink to cool them. That way, a larger volume of cool air removes heat without a great rise in temperature, and can go on to contribute to total case cool air flow.



a typical CPU duct shown mounted to the side panel of a case

Ductwork for CPU, GPU and other heat loads can be improvised from dryer duct material, cardboard and duct tape, foam core board, or any of several other suitable materials. Metallic ductwork is not recommended due to the risk of short circuits should it work loose and touch electronic components! A very nice CPU duct made from PVC pipe is shown in an instructable [here](#).

Fans for ductwork can be used to blow air into the duct, replacing the CPU Fan, the CPU Fan can be used by itself to draw in cool air through the duct, or fans can be used at both ends for particularly hot-running gear.

Intel published some specifications for chassis air guides [here](#). These include some very useful mechanical drawings with dimensions.

Case Shrouds: Similar in concept to ductwork is the use of a shroud. Shrouds are like ducts as they are used to divert airflow to a particular component. Where they differ is that a shroud is open on one or more of its sides. Typically shrouds are used to turn a case fan into a CPU fan for a very large passive heatsink.



Precision 690 Workstation CPU/RAM shroud with fan - helpowl.com

The Precision 690, pictured here, also uses the same shroud to direct air to the RAM risers to cool the memory modules. Note the very large fan to the left of the shroud, mounted to push air through the shroud and out the back of the case.

Shrouds are more often seen as features of OEM designs but can be easily improvised for any given design limited only by the ingenuity of the builder. I've had best results with foam core board glued together with Elmer's Carpenter's Wood Glue or Devcon 5-Minute Epoxy (High-melt Hot Glue also works). If in a hurry to prototype a design, I'll use packing tape or duct tape to put everything together, then go back later with a more finished design after proving its effectiveness.

A properly designed shroud can be nearly as effective as a duct at concentrating cooling airflow

where it needs to be, but easier to remove and clean, as all of its internals will be accessible after it is removed.

Baffles: Baffles are simply sheets of plastic, cardboard, or foam core that are placed inside a case to route the airflow to avoid thermal "shorting" or to isolate or divert a portion of the airflow to a particular section of the case. These are typically less efficient than shrouds or ducts, but permit more liberal routing of cables, etc. and are easier to remove or work around. Some case designs have baffles built in. If you see this in a case you are considering buying, check to be sure whether or not the baffles are fixed or removable first, and also check for cable pass-throughs. This might save a headache or three later on during your build.

With the advent of 3d printing, it may already be possible to devise more interesting custom airflow management solutions for computers. (If I ever get my [MendelMax](#) printer built and running, I'll let you know.)

Mounting Case Fans

Most cases locate their power supply units in the upper rear of the interior, with the PSU exhausting air out of vents in the rear of the case. Some newer designs locate the PSU in the bottom of the case pulling air down and out the rear at the bottom. For these case types, the rear of the case should be reserved for fans which pull air out of the case.


There is a common-sense reason for this: Two fans in the same panel pulling airflow in opposing directions will create an airflow "short". The airflow coming into the case will not have opportunity to circulate and pick up heat before the outgoing fan pulls it back out of the case.

This leaves the front and sides as possible inlets for air, with the rear of the case acting as the exhaust.

Case pressure is managed by selecting fans with the proper airflows to support the desired pressure design.

Installing Fan Guards and Filters

For reasons of safety, if a grille is not provided by the case to cover a fan's blades, one should be provided by the builder. These are reasonably inexpensive and can be procured from any number of reputable vendors. Alternately, a protective grille can be easily made from sheet metal (by drilling enough holes in it) or from open mesh wire screen. It's generally faster and more finished looking to just buy a nice fan guard.



steel mesh and more traditional chrome fan guards - google images

Filtration units can also serve as fan guards, as they usually have a filter support that prevents the filter media from being pulled into the fan's blades.



black plastic filter unit with plastic filter media - google images

I have improvised filters "on the cheap" from HVAC "hammock" roll filter media by simply attaching velcro strips to hold it in place on the exterior of the fan guard, but this is a rather gnarly-looking solution. For a more finished appearance, buy or build a filter housing.

Filters for computer fans are available in a wide variety of media types: metallic and plastic screens usually only catch the big chunks, but are washable and reusable. Plastic foam filters are better at excluding fine dust, but must be washed and changed regularly or they become brittle and start producing debris themselves. My personal favorite is HVAC roll filter media. This is a synthetic fiber mat that usually has an appearance similar to fiberglass insulation, but much more open and much less itchy. It can easily be cut with scissors or a pocketknife to any needed shape and size to cover any inlet which needs filtering. In bulk it is relatively cheap.



"hammock" filter media roll - homedepot.com

Most bi-color media (like the blue and white pictured above) can be easily split to fit into thin filter housings. This media is not reusable, and must be discarded in the trash after being replaced with a fresh piece. The real advantage it has (besides being a very effective filter medium) is that less than \$15.00 will buy a roll big enough to make several hundred fan filters. They had me at "hello".

No matter what sort of filters are employed, they must be checked regularly and frequently for signs of obstructions which will reduce airflow. Clean or replace as needed as part of system maintenance.

GPU Fans and Contributions to Case Airflow

This is actually pretty simple. Since GPU fans usually take airflow from inside the case, and exhaust heat back into the case interior, they are not truly contributors to case air exchange rate, and merely add their heat load to the interior. A way to change this would be to install a duct or shroud to divert outside air to the GPU fan's intake. If this is done, then the GPU fan can be considered a contributor to airflow coming into the case. Most video cards do not make provisions for this, but I've learned never to discount the ingenuity of a determined builder.

(I would strongly recommend an inlet filter be used for such a modification. GPU Fans are remarkably low profile in most designs, and should be kept as clean as possible for longer service life.)

Cable Management and Impacts on Airflow

Here again, the concepts are simple. Flat ribbon cables present a high profile for drag or obstruction to airflow within a case. SATA drives, with their narrower cable profiles, have done much to improve airflow in modern systems. If I were building a new system today, I would not use any devices requiring PATA/IDE interfacing and their attendant ribbon cables.

If PATA/IDE devices are still in use in a case, there are alternatives. Round cables are not exceptionally expensive (around \$5 or less), but are generally of very high quality construction and worth the price.

Less expensive (and somewhat less effective) is the alternative of modifying the flat ribbon cables to be narrower. *Be advised before considering this option that modern IDE cables with 80 conductors are easily damaged by what is about to be attempted.*

By using an X-acto knife or similar tool to slice several times between the conductors in a ribbon

cable, it can be bundled into a much smaller profile with cable ties or spiral wrap. As mentioned already, the potential exists to damage the cable in the process of modification, and round cables are cheap enough now that it hardly seems worth the effort.

Power Supply cables tend to festoon themselves all over a case's interior if not properly dressed out of the way, also impeding airflow. Modular power supply units, which allow a builder to only use the cables needed for a build, and to add others later as the system is upgraded, can help clean up cable routings to minimize their impact on system airflow. (They make the interior of the case look neater and more professionally built, too.)

Generally, the cleaner and neater the case's interior is, the less impediment it presents to airflow. Cases being sold these days can have much more in the way of cable management features (pass-throughs, cable anchoring points, etc.) than in times past. Choose well when buying a new case.

Going High-Tech: Smart Airflow Monitoring and Controls

So far we've talked about how to design for cooling in terms of structural and electromechanical methods. Now we get into the automation and controls used for regulating temperatures in a computer system. We'll start this discussion with some basics of instrumentation and control.

A Basic Temperature Control Loop: Most temperature control systems are comprised of a sensor which sends input signals to a controller. This controller, which is usually composed of software and hardware, sends output signals to control one or more end devices like fans to respond to changes in temperature. It acts much like a thermostat for a home might, and acts to ramp up fan speeds as the temperatures being controlled (the process value) rises above the desired value (or setpoint).

Temperature Sensors: The most common sensors used for monitoring and control functions are diodes built into the various chips used in a computer. CPU, GPU, and support chipsets may all have sensors depending on the design and implementation of a particular motherboard or graphics card.

These sensors monitor the current across the diode junctions, and can be quite accurate (I've seen claims of 0.5 deg C or 0.8 deg F as a tolerance in the literature).

These are the best sensors for measurement of chip temperatures, as they are built into the chip and are small enough to be very responsive to changes.

Outboard sensors (those not built into the various chips) are available in a wide variety of types depending on how much one might wish to spend. The two most common types for these uses are Thermistors and Thermocouples. We won't be discussing specifics of these other than to say they are capable of measurements more than accurate and repeatable enough for any purpose inside a computer case if properly installed and maintained.

Controllers: We can divide controllers into two very basic types: manual and automatic. The automatic controls can be further subdivided into those built into a system's motherboard and those designed for mounting in a drive bay. Let's start by talking about built-in systems first, since they frequently represent the least expensive and least complicated option.

Onboard Controllers: Back when vendors began providing temperature monitoring and control features as part of a motherboard's Basic Input/Output System (or BIOS), the designs were rather crude. Sensor accuracies were impaired by firmware that didn't have enough resolution to properly represent their input to the right precision. Controls were rudimentary alarm and shutdown functions without any real control over fans or other devices.

More modern motherboard designs have come a good way since then, but it still pays to shop around, as vendors' implementations are most certainly not equal across the industry.

Even modern boards from such manufacturers as Gigabyte, MSI, and Asus may use sensor scalings that are grossly inaccurate (anything more than plus or minus 5 deg C from actual). These can still control fans adequately enough to keep a system running well, though, and so should not be overlooked as the least expensive option.

There are also variations within a manufacturer's product line, so it is wise to do some reading in the vendor's literature for a given board and to check in with the tech blogosphere at places like motherboards.org or maybe Tom's Hardware or Phoronix (which is more geared to folks who use Linux - hint, hint!). One might also discuss prospective motherboards on the PCLinuxOS Forums prior to purchase to benefit from the knowledge and experience of members of the community.

Drive Bay Fan Controllers: There are numerous products designed for mounting in one or more adjacent 5.25" drive bays which will control fans in a computer. These range from simple manual speed controls (which provide a voltage to the fans controlled by knobs on the front of the controller) to more sophisticated electronic units which provide front-panel readouts of sensor inputs and automated control of fans to actually control temperatures.

Manual fan controller units typically do not have a software component, and are thus easier for Linux users to implement. They typically do not provide monitoring at all, though, so users should rely on the built-in temperature monitoring software in Linux to track temperatures and adjust fan speeds accordingly.

Akust FS-01-0107-AKS 6-channel manual fan controller - akust.com



The more sophisticated and better automated units on the market bring all the whiz-bang and bling into temperature controls. These units typically feature LCD or OLED (organic LED) displays reading out temperatures, possibly touchscreens, and automated control of up to six fans based on the temperature readings. Some may even provide a wireless remote! Linux users should beware of units which depend on software - such software is typically Windows-based, and such units will be of no use at all or only limited use in a Linux system.



NZXT Sentry LX 5-channel Fan Controller - www.nzxt.com

Some of the things worth paying attention to when shopping for one of these units are:

- **Size:** does the unit occupy one bay or two? The two-bay units provide lots more display area, but

take two adjacent bays. How much room is available in the case to mount the device?

- **Sensors:** these units provide their own (outboard) sensors. Follow the vendor's guidelines for mounting sensors as set forth in their literature for best results.

- **Outputs:** how many fan outputs are provided? How much power can each output provide? (This should be matched to how much power a controlled fan will draw.)

- **Touchpanel:** some of these units will provide touchpanel controls. These can be neat, but not everyone gets on well with touchpanels.

- **Software:** be aware whether or not a controller requires software to function. If so, make sure the vendor provides software for your operating system (hopefully that's PCLinuxOS, by gum!).

- **Cables:** typically cables are provided for all fan outputs and sensor inputs. Dress these well out of the way in the case during installation. Make sure they are of adequate length or extendable before buying. Determine whether two-pin fans are supported - many controllers only support three-pin fans.

- **Alarming/Shutdown Functions:** does the unit provide alarming and/or shutdown functions, and how are these implemented?

An electronic fan controller can be a real asset in controlling noise levels, and, finally, it can make quite a statement as a highly visible part of a computer build.

Hardware

As was mentioned way back at the beginning of the article, the basic hardware for an air cooled system is very simple: a fan, a heatsink, and some thermal couplant to interface between the heatsink and the load.

Here, in this final section, we explore how to put the hardware together.

Thermal Transfer Media

There are several types of thermal transfer media. Before we get into types, though, let's quickly mention what their purpose is. All the surfaces involved - chip die covers, heatsink mating surfaces, etc., all have mechanical imperfections. What thermal transfer media do is provide a filler at the interface that fills the voids where surfaces are not completely flat. This helps conduct heat across the two surface boundaries, lowering the thermal impedance of the interface.

Thermal Pastes: The most popular type of medium is paste. This is available in a wide variety of formulations which all have their advocates in the literature. Whatever is used, it should be *thermally* conductive without being *electrically* conductive. This is so that, if it migrates onto circuit surfaces it will not short them out.

Speaking of migration, paste should have some thickness and should not migrate easily with rises in temperature. It should also resist drying out or hardening with age.

The most common paste is made from zinc oxide powder suspended in a thick silicone grease. This is very reliable in its thermal and mechanical properties, and is absolutely non-conductive. A tube like the one pictured here (6.5 grams) is less than \$3.00 retail.



GC Electronics 10-8108 Zinc Oxide Heatsink Compound - jameco.com

Arctic Silver and pastes similar to it tout much higher thermal conductivities, (and are priced accordingly). I have used Arctic Silver 5 with good results, and found it to migrate less than the zinc oxide under high temperatures. Radio Shack sells a twin-pack of Arctic Silver 5 (a total of 7 grams) for about \$15.00 last time I checked. It can be purchased less expensively from online sources if one is not in a hurry. More about Arctic Silver can be read at the manufacturer's website:

<http://www.arcticsilver.com/as5.htm>



Arctic Silver in 3.5 and 12 gram dispensers - arcticsilver.com

Be aware that Arctic Silver has a tendency to cling aggressively once surfaces are mated. It may be necessary to use a bit of "diplomacy" to avoid pulling the CPU with the heatsink if service is needed. Take it slow and easy if the heatsink is stuck, and use a thin non-conductive object like a plastic scribe to attempt to pry the sink up from the chip die.

Arctic Silver is also marketed as being "non-conductive but capacitive" by its manufacturer, and caution is advised when applying it to ensure none gets on circuit traces.


There are several other formulations available. I've used a grease similar to Arctic Silver made by Dynex that is not as aggressive but seems to work well. If in

doubt as to suitability, read up at the manufacturer's website before using.

Arctic Silver also sells a paste called Ceramique 2 which is billed as a phase-changing paste. I've never used it, but from what I'm seeing it is quite well liked in tech circles.

Do not use pastes or grease compounds not made for thermal coupling, or that are electrically conductive. System damage can and will result.

Thermal Pads: These are little pre-cut purple, pink, or light blue pads that look like sticks of chewing gum, but are actually composed of ceramic microbeads in an elastomeric matrix.



Ziotek P1958 thermal pads - ziotek.com


Pads are generally a headache to use, are not as thermally conductive as pastes, and have to be stored in a cool and dry environment prior to use.

Why would anyone use these troublesome critters, then? They are typically used by OEMs for secondary heatsink elements on laptops, and in other applications where a sizeable gap may exist between the heatsink and the chip die. They are able to fill that gap without migrating whatsoever, and are at least adequately conductive enough to protect the load.

Phase Change Pads: These differ from the elastomeric pads above in that they are thinner, and

of a different composition. They also tend to adhere quite aggressively when they are heated for the first time after being applied.

The place these are most often seen is on OEM-supplied CPU heatsinks for desktop or workstation use. The appearance is generally that of a wet-looking grey or beige colored putty spread very thinly over the mating surface of the heatsink's slug. If touched, it would have a slightly tacky feel, but it's not recommended to touch the surface as oils from the skin may impede heat transfer across the interface.



Rosewill RZX-100 heatsink with phase change pad in place - rosewill.com

The operating principle is that heat causes a phase change in the material which allows it to flow to fill voids in the interface, enhancing heat transfer.

If your heatsink is provided with this type of pad, use it unless you have worries about being able to remove the heatsink later. In that case, carefully clean the surface of the slug with alcohol or acetone, and apply a good quality thermal paste.

If you want to use this type of pad in a custom design, 3M supplies it in rolls or large pads, but it's not cheap. Check their [website](#) for more info. Aftermarket vendors sell this material pre-packaged in common sizes that are much more affordable.

Heatsinks: A heatsink is basically a big block of metal which is finned and slotted to help it dissipate heat from a load to prevent that load from suffering thermal failure.

Metals are generally good conductors of heat. Of the materials in use for heatsinks, three stand out as better conductors than the rest: silver, copper, and aluminum.

Silver is rapidly appreciating in price per ounce at this writing, and so is not usually considered for use in heatsinks except in the most demanding situations.


Next in thermal conductivity is *Copper*. For much less than silver, copper can provide over 90% of silver's ability to dissipate heat. Copper heatsinks are, thus, the choice where small size relative to load is desired at an affordable cost. Copper does oxidize, but this oxide does not usually impede the function of a heatsink unless allowed to degenerate unchecked.

The most common metal used for heatsinks, and by far the least expensive, is *Aluminum*. Aluminum oxidizes almost immediately on contact with air, but the oxidation forms a very tough and conductive skin on the metal which helps protect it from further oxidation. Roughly speaking, the thermal conductivity is about half that of copper. This means that an aluminum heatsink will be much larger for a given load than one made from copper.

Lately, newer heatsinks of a metallic composite have become more popular. These use a copper slug to contact the load bonded to an aluminum fin body which dissipates the heat from the slug. This combination can be highly effective if done properly. (See the [Rosewill RZX-100](#) pictured earlier for an example of this type.)

Heatpipe Type Heatsinks: More advanced heatsinks make use of *heatpipe* technology to

improve their effectiveness. A heatpipe is a sealed and evacuated tube or tubing loop which contains a liquid which boils on contact with the heat of the load, becoming a vapor which travels up the tube to the finned end where it is cooled to condense back into a liquid which returns to the hot end by gravity, capillary action, or centrifugal force to begin the cycle over. Heatpipes can transfer truly incredible amounts of heat from a load.



a heatpipe type heatsink - [wikimedia](#)

Wikipedia has an [excellent in-depth article](#) on heatpipe technology, which is included in the reference section for those interested in further reading on the subject.

The major drawbacks to using a heatpipe style cooler are two: they generally require more room and more airflow (a bigger fan) than a simple heatsink. For crowded cases, check clearances carefully before attempting to install a unit of this type. This brings us to our next topic.

Selecting and Preparing to Install a Heatsink

The factors which drive selection of a heatsink are:

- Heat from the Load.
- Available case volume and clearances

- Need for a fan or fans to cool the heatsink, and their sizes
- Whether ducts or shrouds will be part of the design.

We've already covered where to look for the amount of radiated heat from the load - it's in the vendor's literature in most cases, especially for CPUs and support chips.

Case Volume and Clearances

CPUs are the major loads in a system, and typically pose the most challenges for design in terms of case volume and clearances. If building a multi-CPU system, more case volume and clearance is better than less.

If upgrading from stock CPU coolers, check carefully how much clearance there is in between the CPUs and from the CPUs to the video slot on the motherboard. These dimensions may be given in the motherboard manual. If so, they will generally be reliable. If not given, or if you're not in a trusting mood, consider removing the motherboard from the case to make taking measurements easier.

If your intended video card has any real thickness to it (as many modern ones do) account for this in determining your clearances to the nearest CPU sockets.

Now, when shopping for heatsinks and fans for the CPUs in the system, look at the dimensions of those and compare them to your measurements, possibly even sketching things out to get a feel for the mechanical arrangement. If the clearances are adequate, proceed.

Selecting Fans to Match the Heatsink

Typically vendors will include a fan of appropriate size with any CPU cooler they sell, often mounting it

prior to shipment. Quality and effectiveness of these will vary by vendor and even by differing products in a vendor's line.

If clearances permit it, mount the fan on an adapter that moves it a short distance away from the heatsink. This will minimize or eliminate the dead spot at the hub of the fan, and cool the entire heatsink (and, thus, the load) more evenly.

Fans will be discussed in greater detail further on, but for now, look for adequate airflow, good construction, and low noise levels in all the fans used for a build.

Using Shrouds and Case Air Guides

If using a shrouded or ducted design, the most direct route is the best route for airflow. Any twists, turns or offsets in the airflow path represent a source of turbulence, drag, and loss. In the case of a shroud, the cooling fan can be located far enough away to entirely eliminate the fan's dead spot by the time air reaches the heatsink.

If using a CPU duct (a Case Air Guide) to bring in outside air, leave a small gap between the CPU fan and the end of the duct. This will prevent any rubbing or pressure on the fan from the duct. A gap as small as 5 mm is more than adequate. Most CAGs are adjustable easily enough to accommodate this without much difficulty.

Preparing a Heatsink for Installation

When preparing to install a heatsink, first look at its condition. Inspect it for any broken fins, cracked heatpipes, and surface imperfections on its mating surface. Inventory the fixtures and fasteners provided with the unit, and inspect the provided fan for clearance from the fins and general condition if it is the stock fan is to be used. Any significant

damage to any component is cause to return the unit for an exchange or refund.

Condition of the mating surface is especially important. If scratches or pits exist in the mating surface of a new unit, return it for an exchange. Examine any phase change material or other heat transfer medium provided "from the factory". If this has been disturbed in any way, either clean it off and apply new or return it to the vendor for an exchange.

Resurfacing an Existing Heatsink

If this is a maintenance event on an existing system, and it has been removed from its CPU, first clean and then inspect the mating surface. (Alcohol or acetone make excellent cleaners, but do NOT clean over the computer chassis with either - move away from the computer to do the cleaning to avoid fouling the motherboard with solvents).

Any defect in the mating surface should be buffed out. A good method for ensuring the mating surface stays flat is to use wet-or-dry sandpaper of varying grades on a piece of plate glass. Start with wet 600 grit paper, and move up through about two or three grades to 1200 as surface defects are removed and a flat, smooth surface begins to present itself. Always keep the surface of the heatsink pressed flat against the glass and paper and use moderate, even pressure. The resulting finish will be more than adequate to accept the heat transfer medium and mate up to the CPU die.

Note: smaller heatsinks for support chips and other small loads in a system are usually provided by the vendor and are more or less permanently attached and should not require much maintenance other than keeping them free of dust or other contaminants. These should only be removed and inspected if a problem with overheating develops. The reason for this is that removal of these heatsinks can damage a motherboard, even if done

properly because they are frequently mounted with an adhesive heat transfer medium.

Preparing the Motherboard for Installation

Presuming the CPU is already correctly installed, next inspect its upper surface, and clean off any residual heat transfer medium or foreign matter clinging to it. I typically will use a small alcohol prep pad that has been squeezed to remove excess liquid, and gently scrub the top surface of the chip. Phase change media can be stubborn to remove, and will require patience and persistence.

Once cleaning is completed, use aerosol duster "canned air" (or clean, dry low-pressure compressed air) to blow the motherboard dry of any residual alcohol that may have contacted the motherboard or the CPU socket. Ensure the area is entirely dry before going further.

Most modern heatsink designs are built to use a "universal" mount that is already provided by the motherboard vendor. If the instructions included with your heatsink require it, *install any fixtures that came with the heatsink for retaining it in place*. This may require removing the motherboard from the machine (or, at the least, opening the back side cover of the case) to access the back side of the board.

Applying Heat Transfer Medium

Once any needed fixtures are in place, and the motherboard is dry, it is time to prep the mating surfaces with heat transfer medium. If using a medium pre-installed on the heatsink, inspect it one final time for any signs of disturbance, and, if all is well, mount the heatsink in accordance with its instructions.

If applying fresh heat transfer paste, spread a small amount onto the CPU's mating surface and tool it

into an even layer with a wooden or plastic spreading tool. (I make these from old Wal-Mart gift cards.) *A typical mistake for the beginner to make is using too much paste! More than a thin even layer applied to one surface is too much.* Clean up any excess paste that migrates onto the motherboard prior to installing the heatsink. (I use q-tips lightly dabbed with alcohol, and allow it to dry completely before proceeding.)

TAKE CARE here not to overpower any spring clips and slip out of them during mounting. Damage to the motherboard can result from gouging its surface with a hard metal tool.

Also take care not to move the heatsink excessively during installation, especially if using phase change material.

Once all clips are securely attached to their mounting points and any latches are secured, connect the CPU Fan to its socket on the motherboard. This completes the installation of the CPU heatsink.

Fans

Fans for use with computers are generally of two types: axial flow and radial flow. Radial flow are by far the least common, so we'll discuss them first.

Radial flow fans take air in through their centers and discharge it from an opening in the radius of the fan, much like a squirrel-cage blower in an air conditioning system.

These are most often seen in laptop computers and as so-called "slot coolers". Their main usefulness in a desktop system is to provide localized airflow to supplement that provided by larger capacity axial fans in the case.



Antec Cyclone Blower slot cooling fan - newegg.com

Axial flow fans are the much more common type used for computer cooling applications. In this type, fan blades rotate to pull air along the axis of rotation. They typically have greater airflow than the axial fans available for computer uses, and so are more generally useful.



typical axial flow fans - cooler-master.com

Axial fans have one major drawback: the motor at the hub of the fan blade creates a "dead spot" just downstream of which airflow is severely reduced. Newer designs use a tapered hub (sort of like the spinner on an aircraft propeller) to reduce this dead

spot, and, if it is a big concern, a fan shroud or adapter can be used to put distance between the fan and what it's cooling to eliminate the dead spot altogether.

Fans in general have a few specifications other than physical size which are useful in helping to select them for a given purpose: airflow, noise level, current, voltage, bearing type, and materials used in their construction.

Since fans tend to collect dust, consider using a fan with sealed magnetic levitation bearings. These are virtually frictionless and fans so equipped tend to last a very long time before needing replacement. Somewhat less expensive are the sealed ball or needle bearing designs. Last in quality are those fans which use sleeve bearings - these tend to fail much sooner, but are generally much cheaper. There are other bearing types out there, but I am not so familiar with those and so won't comment further.

Look carefully at rated airflow and operating current specifications when shopping for fans to go on a CPU cooler. Bigger fans can turn more slowly and deliver more airflow more quietly than smaller ones.

Most fans will specify airflow in CFM, but some (Noctua, notably) use cubic meters per hour. To convert CFM to cubic meters per hour, multiply by 0.588577770211 or just use the converter found [here](#). To go the other way and convert cubic meters per hour to CFM, multiply by 1.69901082 or use [this converter](#).

If the power consumption in Watts is needed for a fan (say, to match it to a fan controller), multiply its rated current by its operating voltage. For example, a 12vdc fan that draws .33 amperes (or 330 milliamps) will consume 4 Watts at maximum speed and airflow.

Lastly, but not least in importance, look at the sound pressure level rating for any fan being considered. (I will never forget the 3-bladed, 4500 RPM, 92 mm fan I put in my old Precision 220 Workstation - it moved over 100 CFM, but at 66 dB! - it sounded like a jet aircraft on the runway!)

A good noise figure for most fans will be anywhere below 30 dB, and the lower the better. If mounting a fan to a case opening, consider using a soft rubber or vinyl sound isolating mount. This will reduce the amount of fan noise being coupled to the metal case and amplified.



80mm acoustic dampening fan mount - xoxide.com

When looking at a fan, also take note of the number of blades it has. More blades generally move more air more quietly. Fewer blades each have to take a bigger bite and spin faster to move the same amount of air, which generates more noise.

An excellent introduction to noise levels and decibels sound pressure level can be found at endpcnoise.com.

We'll end the section on fans with a few tips on how to mount them for desired airflow. Most axial fans will have arrows on them which indicate two things: direction of rotation, and direction of airflow.



flow arrows on an axial fan - google images

Find these arrows whenever mounting the fan, and mount it facing the direction which will produce flow in the desired direction. Don't forget to plug all the fans in after mounting!

Case Modifications

Case modification (or modding) is a sport engaged in by those who are seeking geek creds, or those who just want a system to have an extra performance edge.

While I have done some mods for looks in the past, my major focus has always been modding for improved system performance in one way or another. It can be a great way to improve cooling for a system inexpensively.

We'll discuss two general classes of modifications: internal and external.

External mods affect the outer surfaces of the case. These include such things as cutting a blowhole, or adding a window to a side panel, or even a custom airbrushed paint job.

Internal mods include such things as installing extra case fans, building and installing shrouds, baffles, or custom ductwork to direct airflow, adding cold cathode or LED lighting to the inside of a case, installing liquid or solid-state cooling, and my favorite, wire and cable management for improved airflow.

What modifications you might wish to undertake are limited only by your imagination, your willingness to design, plan, and execute the mod, and your willingness to experiment and recover gracefully from the mistakes you will make. As the old saying goes, "Make no mistake: sooner or later something will go wrong or you're just not doing it right!"

Case modders are a strange bunch. They have their own language, and their own ways of doing things. There are as many opinions on how best to do something as there are modders. What all modders share, though, is a common fascination with tools.

Basic Tools for Modding

Most any modder will agree: there's no such thing as having too many tools, especially power tools. There are some very basic tools that will show up in almost every modder's kit, though. Here's a basic list:

- **Screwdrivers:** these can range from plain to fancy. I favor the Great Neck multi-bladed screwdriver (it features Phillips, Flat, and Torx in one tool), but there are some applications for which a long straight-shank screwdriver will be needed. Top of the line quality is not essential, but good tools tend to last longer and work better.
 - **Precision Screwdrivers:** go high quality here, with Phillips as well as Cabinet Tip flat-bladed sets. Get a set with handles that give a good grip and have a spinner cap for quick run-up during assembly.
 - **Bit Sets:** these are sets of hexagonal bits which end in various screwdriver, hex key, torx, secure

torx, and other handy bits. I got mine at Advance Auto Parts for \$9.95. Use with a multi-bit screwdriver handle or a power screwdriver for those jobs with lots of fasteners.

- **Electric Drill:** a good quality $\frac{3}{8}$ " drill motor and a drill index filled with various assorted drill bits always comes in handy when a hole is needed in a cover or bracket.

- **Pliers:** various pliers such as slip-joint, needle-nosed, diagonal cutters, and others perform a wide variety of tasks. Smaller miniature pliers sets are useful in the close confines of a computer case.

- **Files:** no, not the digital kind, the ones used for removal and shaping of metal. A good assortment of fine and rough cut files in a variety of sizes with safe handles comes in handy for all those little detail touches.

- **Dremel Tool:** a good quality variable-speed rotary tool with a wide assortment of bits can make cutting, sanding, polishing, and other tasks go very quickly. Even so, slow and easy does it - mistakes made at high speeds tend to be big mistakes.

- **Adel Nibbler:** this specialty tool amounts to nothing more than a handheld sheet metal broaching tool. Each time the handles are squeezed together, a punch is pulled through the metal to cut out a small slug. Slower than cutting with a Dremel or with Tin Snips, the nibbler can cut complex shapes without deforming the edges. I highly recommend this tool for those special jobs.

- **Heat Gun:** resembling an industrial blow dryer, these get much hotter, and are useful for heating glued surfaces to ease separation, or for shrinking heat-shrinkable tubing. Do not use near a motherboard or where heat sensitive chips or other devices are present which may be damaged.

- **Pencil Soldering Iron and Desoldering Tools:** nothing else will quite do when wires or terminals need to be soldered. Most all solder being made these days is lead free. I still prefer lead solders, though, as they just make a better joint and flow much more freely when heated properly. Desoldering tools help when defective components have to be removed to be replaced.

- **Combination Square with Scribe:** a combination square with square, mitre, and protractor attachments, a nice steel rule, and a scribe housed in the square make a nice way to measure and mark off metal case elements for cutting or other modifications. Find a set that has a rule marked in both inches and centimeters.

Certainly there are other tools which could prove very useful. These basics will be a great start on your modding toolkit. Whichever tools you choose to use, wear the protective gear required to keep you safe from injury. Happy Modding!

Conclusions

The various "pieces-parts" of a computer cooling system all work together when the system is properly designed to protect and enhance the performance of the computer. When properly designed, installed, and maintained, these systems can provide years of reliable service.

Hopefully, this article will serve as a jumping-off point to spur the readers' curiosity, inventiveness, and willingness to explore this vital aspect of computer technology.

References

Computer Cooling in System Design: one of several highly technical and very lavishly illustrated articles by the technical staff at NMB Minibea, a manufacturer of cooling fans for computer and other uses.

Digital Technology Arts: a source of tightly-distilled and useful information on a wide range of computing topics.

Xoxide's Water Cooling How-To: Dave Melchior's excellent article on liquid cooling.

Computer Water Cooling System: a fascinating Instructable in which a nice fellow from Korea home-brews a very professional looking liquid cooling system from an oil cooler for a Hyundai, some Foamex sheets, and various off-the-shelf components. A good example of what can be done by a creative inventor.

Wikipedia Article - Thermoelectric Effect: an excellent introduction to Peltier Effect and Seebeck Effect with sound theoretical and mathematical underpinnings.

Wikipedia Article - Heat Pipe: another excellent article from Wikipedia, this one explaining how heat pipes work.

The Peltier Effect: an excellent concise article at tech-faq.com.

The Heatsink Guide - Peltier Guide: One of several insightful articles at heatsink-guide.com.

Quick-Cool.com's Peltier Cooling Library Page: several very well written white papers on Peltier Effect and Peltier cooling.

Tech Powerup's CPU Database Page: a compendium of massive amounts of CPU data from various vendors.

CPU-world.com: an excellent informational resource.

CPU-DB at Stanford University: Lots of good stuff here, including data for hardware produced over the last 40 years!

Computer Cooling Duct Instructable: A very cool instructable showing a great use for some cheap PVC pipe, a small square of thin plywood, and some clever cutting and gluing.

Intel Ark: Intel's source for technical data on their integrated circuits and electronic products, including CPUs and support chipsets.

Thermally Advantaged Chassis Design Guide: one of Intel's very technical specifications with some quite useful mechanical drawings and information on computer case design.

Chassis Air Guide Design Guide: another useful design spec from Intel.

Decibels dBA - a good introduction to what decibels of sound pressure level really mean from endpcnoise.com.

Convert CFM to Cubic Meters per Hour: one of many excellent unit conversion pages from <http://www.convertunits.com>

interfacebus.com: an excellent source of a wide variety of technical information concerning computers, their hardware, and how they communicate.

Vendors

Disclaimer: listing of a vendor here does not constitute an endorsement. Do your own research as to the suitability of any components for a particular purpose.

Koolance: a source for a wide variety of liquid cooling system products and components.

Xoxide.com: another vendor of cooling system and other products, and a good source of general information about cooling systems and modding.

Tellurex FAQ: product information with a good bit of theory on Peltier Effect from Tellurex.

International Community
PCLinuxOS Sites



Netherlands



Turkey



Denmark



Czechoslovakia



Italy



Poland



FINLAND



Brazil



Does your computer run slow?

Are you tired of all the "Blue Screens of Death" computer crashes?



Are viruses, adware, malware & spyware slowing you down?

Get your PC back to good health TODAY!

Get




Download your copy today! FREE!

Inkscape Tutorial: Create A Paisley Pattern, Part One

by Meemaw

I grew up in the 1960's, and sometimes we wore some pretty wild stuff. One of my high school friends, very talented artistically, painted her yellow car with hundreds of paisley patterns on our last school day before graduation. Go to Google, search for paisley, and you can find tons of patterns!




This sort of pattern seems to be making a comeback. I've found it on everything from clothing to tablecloths and from tablet and phone covers to tattoos. As you can see, this is basically a curved teardrop pattern with a lot of extra patterns inside it to make it fancier. How can we make one of these? We can use Inkscape.


Last month we made several small patterns to use in a project, and we will use some of them now to make our paisley. We will start drawing the teardrop using the Bezier tool. Then we will "pretty it up" with the small patterns we made last month. Let's make one similar to this one:




(Wow! I don't want to have to place all those little circles by hand!) Don't worry, we'll use a different method. Let's draw our teardrop first. Open a page in Inkscape and choose your **Bezier** tool. Using that tool, make a kind of teardrop shape, clicking with your mouse at least 10 or 12 times. That will give us nodes to use to manipulate our shape. Set the stroke black or gray, and the fill white.



When we were making our basic shapes, we used the Auto-Smooth tool to round off our leaf. This time we want to smooth out this teardrop, so select all of your nodes by drawing a frame around them with the Node tool, BUT, before you hit Auto-Smooth, press Shift and de-select the node on the very end. We want it to have a point on it.




Now click Auto-Smooth. Each of your line sections will become rounded and the curve will be much smoother. You may still feel like adjusting a few of the nodes separately. Go ahead! It's your project. I adjusted some of them in the example below.




Now, instead of trying to draw another one of these, we are


going to duplicate this one and adjust the size. Click on your object and duplicate it. Remember that the duplicate is placed directly on top of the original, and it is now selected. Click on **Path > Outset** to increase the size of your object, or press **<CTRL> +)**. Each time you press it, your new curve will grow a bit. My preferences are set to 2 px, and we want it 20 px bigger all the way around, so we should do this ten times. You can use the menu, but the keyboard shortcut is much quicker!



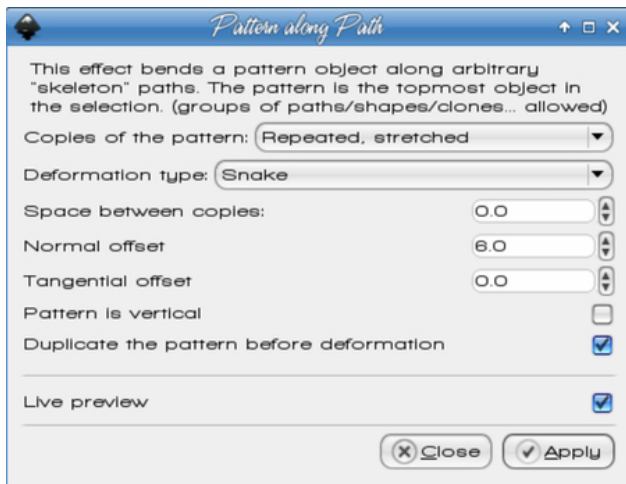
Now, choose the original object and duplicate it again. This time use the **Path > Inset** menu item or **<CTRL> + (**. This one will shrink inside of the original so you will have three separate curves.



Now we want to place some decorations on our object. Remember those decorative patterns we made last month? Import them into your new project. It's easier to import the svg file so changes can be made to them if necessary. After you import it, you should click on Ungroup so you can move each of them around freely.



Now we are going to decorate our curves. Notice that I have duplicated and resized one of the circular shapes and it is on my page above my paisley curve. Click on the circular shape and, holding the **<Shift>** key, click on the outer curve. Now click on **Extensions > Generate from Path > Pattern along Path**. The following window will appear:



The following options are available, and I will explain each:


Single; Repeated; Single, stretched or Repeated, stretched - This choice determines how the objects are arranged around your path. Single means that only a single copy is used, while repeated means that (obviously) the pattern is repeated all around the path. Stretched means that the pattern is stretched around whatever curve you use.

Snake or Ribbon - Snake means that the copies are arranged in a flat manner, while ribbon means that they have a bit of perspective to them. Depending on your project, you may need the ribbon setting instead, but we're using Snake here.

Space between copies - You can set your group of objects to be spaced out, or touching at the edges. This will require a number setting. A setting of zero lets them touch, while a setting of 3 pixels will space them out 3 pixels.


Normal Offset - This will set your group of objects directly on your curve, but setting the offset with a positive or negative number can set them inside or outside your curve.

Tangential Offset - This setting has the effect of moving your objects sideways on your curve, which can be helpful if you see a corner where your objects seem to have a result you don't like. For example, my diamond might show up as on the left, but shifting it a bit would give the result on the right.



Each pattern is different, so you would have to play with the settings to get what you want. If making the setting positive doesn't seem to work, try a negative setting.


Pattern is vertical - I haven't used this except to experiment, but I found that it orients all of the objects at a 90 degree angle. I had my diamond shape oriented "longways", but when I checked the box, the extension oriented the diamonds with the short edges touching. In the example below, the outer line results if Vertical isn't checked, and the inside line results if it is checked.



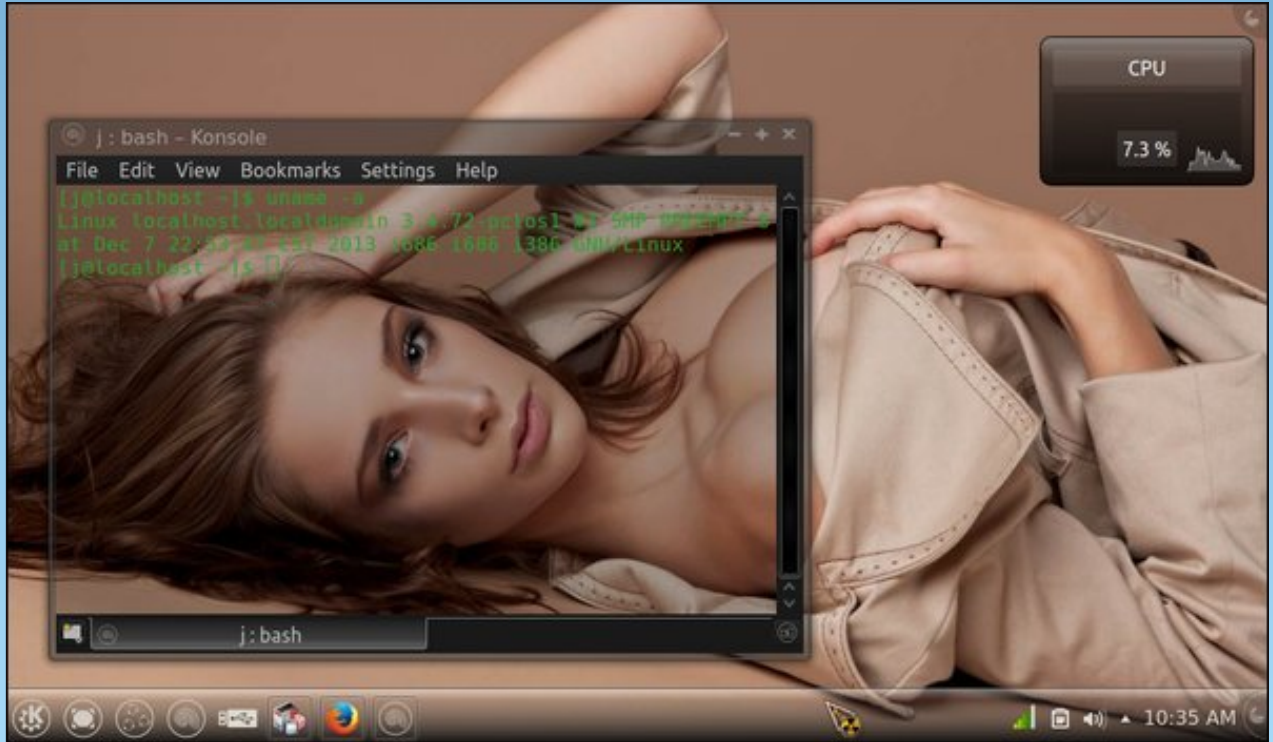
Duplicate pattern before deformation - Duplicates the original pattern you are using before you close the window. If it isn't checked, the single pattern above the project will be deleted. It might be OK to leave it unchecked, but it's always nice to have a copy of the pattern until you are completely finished with your project.

Live Preview - Lets you see your choices as you alter them. If you change one of the number settings, you move to another settings box before you see the change.

In the example (next page, top left), the pink & purple circles and the outside purple circle combination were added to the outermost curve, the blue & purple ellipses were added to the next curve inside, then the diamonds were added to the innermost curve.







Screenshot Showcase



Posted by μT6, March 14, 2014, running KDE.

The settings are as follows:
All are Repeated, stretched ... Snake, then for the Space Between Copies, Normal Offset and Tangential Offset settings;

-  1st; 0, 6, 0 (6 for offset put it inside the curve I used)
-  2nd; 1,-8,0 (-8 for offset put it outside the curve)
-  3rd; 0, 0, 0 then tangential of 2.5 (to fix the corner)
-  4th: 0, 0, 0

All have Duplicate pattern before deformation checked. Make sure you have saved your work.

This is a good start! I'm sure your selected patterns and colors are very different from mine, although I made a couple of different colored patterns. Next month we will finish this design.



Send A Gmail Message Without Opening Gmail

by Paul Arnote (parnote)

As the editor of The PCLinuxOS Magazine, I find myself emailing some files back and forth with members of the magazine staff quite frequently. Oh sure, we also use Google Docs, but there are some times that I need to send an individual file to individual staff members. For that task, email is a better solution, as far as I'm concerned.

It is a pain, however, to have to always go to my browser (or email client ... I do not use a standalone email client. I prefer to use Gmail's web interface) and send the file from within Gmail's web interface. If I'm already looking at the file in my file manager, why shouldn't I be able to just send it from there? Well, I can. And so can you.



Before we get started, I don't want to get into a discussion about Google and their privacy reputation. I am fully aware of Google's reputation in this arena, and I choose to use their service despite it. It works for my uses, and I don't have anything to hide (at least, I don't think I do). With that stated, this article is about using Gmail's SMTP servers to send email from your Gmail account, without having to go into, or "open," Gmail. You can set this up to use any other email service or provider that you want, but you're on your own for figuring out how to set it up – although I suspect your ISP or an internet search will be a valuable resource.

Sendmail or Postfix?

Just in case you're not aware, there are two main mail transport programs in the PCLinuxOS repository: sendmail and postfix. In the which-came-first-the-chicken-

or-the-egg equation, sendmail has been around longer. Postfix came along later, aiming to make configuration easier while maintaining compatibility with the commands that users have grown accustomed to using with sendmail. Plus, postfix is faster and more secure.

I started out to use sendmail. But, whoa! Not only was it complicated to set up, it was fickle as well. Sometimes it would work, sometimes it wouldn't. After writing my script and getting everything just how I wanted it, something prevented sendmail from completing the email transaction. I tinkered with a few things, and managed to get sendmail to work again. However, I never did figure out exactly what I did that caused sendmail to finally start sending emails again.

So, so I decided to cut my losses, installed postfix and tried again. Note that you can only have one or the other – either sendmail OR postfix – installed. Installing postfix in Synaptic caused sendmail to be uninstalled. Similarly, installing sendmail will cause postfix to be uninstalled. Despite more steps, postfix seemed easier to setup. It also "felt" more stable, and worked whenever I called it into action. Thankfully, I didn't have to rewrite my script, since both sendmail and postfix use the same commands to send mail.



Send A Gmail Message Without Opening Gmail

Abbreviated instructions for both sendmail and postfix are in the comments of the send-gmail script, which is [downloadable](#) from the magazine's website, along with links to the original articles where I found the initial instructions. Those comments have been stripped from the version of the script published here in the pages of the magazine, for space considerations. We will, however, go over those instructions within this article.

Feel free to attempt the sendmail route, if you wish – and if you're brave enough. However, although it has a fewer number of steps, it's a far more complex program than postfix. The latter originally set out to simplify the entire process, and succeeded. I've included the instructions for sendmail below, for your convenience. Your mileage may vary.

The directions for the setup of either sendmail or postfix should be the same, regardless of which desktop environment you are using.

Laying the foundation with sendmail

Step 1: Install sendmail from Synaptic

Step 2: As the root user, enter the following commands on the command line:

```
mkdir -m 700 /etc/mail/authinfo
```

```
cd /etc/mail/authinfo
```

```
echo 'AuthInfo: "U:root" "I:YourEmailAddress" "P:YourPassword"' > gmail-auth
```

(the I: must precede your email address, and P: must precede your password)

```
makemap hash gmail-auth < gmail-auth
```

Step 3: Still as root, edit your sendmail.mc file (in /etc/mail). Place the following lines RIGHT ABOVE the first line that says MAILER. (DO NOT place the following lines at the top/beginning of your sendmail.mc file!).

```
define(`SMART_HOST', `[smtp.gmail.com]')dn1
define(`RELAY_MAILER_ARGS', `TCP $h 587')dn1
define(`ESMTP_MAILER_ARGS', `TCP $h 587')dn1
define(`confAUTH_OPTIONS', `A p')dn1
TRUST_AUTH_MECH(`EXTERNAL DIGEST-MD5 CRAM-MD5 LOGIN PLAIN')dn1
define(`confAUTH_MECHANISMS', `EXTERNAL GSSAPI DIGEST-MD5 CRAM-MD5 LOGIN PLAIN')dn1
FEATURE(`authinfo', `hash -o /etc/mail/authinfo/gmail-auth.db')dn1
```

Step 4: Still as root, rebuild sendmail's configuration. Run the following command:

```
make -C /etc/mail
```

Step 5: As root, restart the sendmail service:

```
/etc/init.d/sendmail reload
```

Step 6: IF you've done everything correctly, test it. You should receive an email from yourself by entering the following in a terminal:

```
echo "This is a test" | mail -s "Testing" [your-email-address@yourdomain.com]
```

Remove the brackets from around the email address!

Laying the foundation with postfix

Step 1: Install postfix from Synaptic.

The next nine (9) steps need to be performed as the root user.

Step 2: Open /etc/postfix/main.cf and add the following commands to the file:

```
relayhost = [smtp.gmail.com]:587
smtp_sasl_auth_enable = yes
smtp_sasl_password_maps = hash:/etc/postfix/sasl_passwd
smtp_sasl_security_options = noanonymous
smtp_tls_CAfile = /etc/postfix/cacert.pem
smtp_use_tls = yes
```

Step 3: Create the file /etc/postfix/sasl_passwd.

Step 4: Your sasl_passwd account should contain the following:

```
[smtp.gmail.com]:587 USERNAME@gmail.com:PASSWORD
```

Step 5: Fix permissions and update postfix to use our sasl_passwd file:

```
chmod 400 /etc/postfix/sasl_passwd
postmap /etc/postfix/sasl_passwd
```

Step 6: Download cacerts.pem from <http://curl.haxx.se/ca/cacert.pem> and place the file in /etc/postfix. Right click on the link and select "Save file as..." from the

browser menu. Otherwise, your browser may attempt to open it for what it is (a certificate file) and not go any farther.

Step 7: Open the `/etc/postfix/aliases` file. Find the line that reads `postmaster: root`, and replace "root" with a REAL user (your username on your computer). This is on line 15 of my aliases file. Do the same thing on line 80, changing the second instance of "root" to a real user.

Step 8: Run the command `newaliases`, then restart postfix with the command `service postfix restart`.

Step 9: RESTART YOUR COMPUTER!

Step 10: IF you've done everything correctly, test it. You should receive an email from yourself by entering the following in a terminal:

```
echo "This is a test" | mail -s "Testing" [your-email-address@yourdomain.com]
```

Remove the brackets from around the email address!

The send-gmail script

So now that we have a foundation, – the backend (either sendmail or postfix) in place to communicate with the Gmail SMTP servers – we can create a script to handle all of the user end input. I chose to use Zenity to provide dialog boxes for data input, due to its simplicity and relative ease of use.

First, here's the script:

```
1. #!/bin/bash
2. 3. # send-gmail.sh
4. #
5. # A script that uses sendmail (mail) to send a file to the
   specified
6. # user from the command line, or from a Thunar Custom Action. The
7. # script is enhanced with Zenity text entry dialog boxes to walk
   the
8. # user through sending the email.
9. #
10. # Postfix portion based on an article from:
11. # https://rtcamp.com/tutorials/linux/ubuntu-postfix-gmail-smtp/
12. # and info from footstep11 in the PCLinuxOS forum
13. # http://www.pclinuxos.com/forum/index.php/topic,124750.msg1043028.html#msg1043028
14. #
15. # Sendmail portion based on an article from:
```

```
16. # http://linuxconfig.org/configuring-gmail-as-sendmail-email-relay
17. #
18. # CLI Usage: send-gmail.sh [path/filename of file to attach]
19. # Thunar Custom Action Usage: send-gmail.sh %f or send-gmail.sh %n
20. #
21. # Copyright © 2014, Paul Arnote. All Rights Reserved.
22. # As originally published in The PCLinuxOS Magazine, April 2014.
23. # Read the article for more complete information on how to set up.
24. # http://pclosmag.com/download.php?f=2014-04.pdf
25. # Released under the GPLv2 license.
26. #
27. # You are free to alter, change or distribute this work, provided
28. # you do so while retaining this notice, and so long as your
   derivative
29. # work is also released to the public under any of the GPL
   licenses. Any
30. # distribution of this work must also comply with the terms of the
   GPLv2
31. # license, and no charge may be made for this work, whatsoever.
32. #
33. # The postfix and sendmail setup information has been removed from
   the print
34. # version of this script, to save space in the magazine. However,
   the
35. # setup information is available in the version of the script that
   can be
36. # downloaded from the magazine website here,
37. # http://pclosmag.com/Misc-Downloads/send-gmail.sh.txt
38. # Don't forget to strip the file of the .txt file extension, and
   make the
39. # file executable (chmod +x send-gmail.sh).
40. #
41.
42. hasCC=0
43. hasBCC=0
44.
45. function getSubject {
46. Subject=`zenity --title="Send Gmail" --entry --text="Enter your
   email subject:"`
47.   if [ $? == 1 ]; then
48.     exit
49.   fi
50.   if [ `echo ${#Subject}` -lt 3 ]; then
51.     zenity --title="Send Gmail" --question --text="No
   subject specified.\n\nDo you want to provide a subject?" --ok-
   label="Yes" --cancel-label="No"
52.     if [ $? == 1 ]; then
53.       zenity --title="Send Gmail" --info --
   text="Exiting. bject specified."
54.       exit
55.     else
56.       getSubject
57.     fi
58.   fi
```

Send A GMail Message Without Opening Gmail

```
59. }
60.
61. function getMessage {
62. Message=`zenity --title="Send Gmail" --entry --text="Enter your
email text:"`
63.   if [ $? == 1 ]; then
64.     exit
65.   fi
66.   if [ `echo ${#Message}` -lt 5 ]; then
67.     zenity --title="Send Gmail" --question --text="No
message text specified.\n\nDo you want to provide message text?" --ok-
label="Yes" --cancel-label="No"
68.     if [ $? == 1 ]; then
69.       zenity --title="Send Gmail" --info --
text="Exiting. No message text specified."
70.       exit
71.     else
72.       getMessage
73.     fi
74.   fi
75. }
76.
77. function getRecipient {
78. TO=`zenity --title="Send Gmail" --entry --text="Enter your
recipient's address:"`
79.   if [ $? == 1 ]; then
80.     exit
81.   fi
82.   if [ `echo ${#TO}` -lt 10 ]; then
83.     zenity --title="Send Gmail" --question --text="Improper
or no primary recipient specified.\n\nDo you want to provide a
recipient?" --ok-label="Yes" --cancel-label="No"
84.     if [ $? == 1 ]; then
85.       zenity --title="Send Gmail" --info --
text="Exiting. No recipient specified."
86.       exit
87.     else
88.       getRecipient
89.     fi
90.   fi
91. }
92.
93. function getCC {
94.   zenity --title="Send Gmail" --question --text="Do you want to
send a CC\n\to another recipient?" --ok-label="Yes" --cancel-label="No"
95.   if [ $? == 1 ]; then
96.     let 'hasCC = 0'
97.   else
98.     CC=`zenity --title="Send Gmail" --entry --
text="Enter your CC recipient's address:"`
99.     let 'hasCC = 1'
100.   fi
101.   if [ "$hasCC" == "1" ]; then
102.     if [ `echo ${#CC}` -lt 10 ]; then
```

```
103.       zenity --title="Send Gmail" --question --
text="Improper or no CC recipient specified.\n\nDo you want to provide
a recipient?" --ok-label="Yes" --cancel-label="No"
104.       if [ $? == 1 ]; then
105.         let 'hasCC = 0'
106.       else
107.         getCC
108.       fi
109.     fi
110. fi
111. }
112.
113. function getBCC {
114.   zenity --title="Send Gmail" --question --text="Do you want to
send a BCC\n\to another recipient?" --ok-label="Yes" --cancel-
label="No"
115.   if [ $? == 1 ]; then
116.     let 'hasBCC = 0'
117.   else
118.     BCC=`zenity --title="Send Gmail" --entry --
text="Enter your BCC recipient's address:"`
119.     let 'hasBCC = 1'
120.   fi
121.   if [ "$hasBCC" == "1" ]; then
122.     if [ `echo ${#BCC}` -lt 10 ]; then
123.       zenity --title="Send Gmail" --question --
text="Improper or no BCC recipient specified.\n\nDo you want to
provide a recipient?" --ok-label="Yes" --cancel-label="No"
124.       if [ $? == 1 ]; then
125.         let 'hasBCC = 0'
126.       else
127.         getBCC
128.       fi
129.     fi
130. fi
131. }
132.
133.   if [ $# -lt 1 ]; then
134.     zenity --title="Send Gmail" --info --
text="<b>Usage:</b>\n\tsend-gmail.sh [path/filename]\n\n<b>Thunar
Custom Action:</b><i>\n\t<i>send-gmail.sh %f</i>\n\t\tOR\n\t<i>send-
gmail.sh %n</i>\n\n\nNo file specified. Exiting."
135.     # Xdialog --title "Send Gmail" --infobox "Usage:\nsend-
gmail.sh [path/filename]\n\nThunar Custom Action:\nsend-gmail.sh
%f\nOR\nsend-gmail.sh %n\n\n\nNo file specified. Exiting." 0 0 0
136.     exit
137.   fi
138.
139. FILE1=`zenity --title="Send Gmail" --entry --entry-text=$1 --
text="Enter the path and filename\nof the file to attach:"`
140.   if [ $? == 1 ]; then
141.     exit
142.   fi
143.   if [ -f $1 ]; then
```

Send A GMail Message Without Opening Gmail

```
144.         FILE=$FILE1
145.     fi
146.     if [ -z $1 ]; then
147.         zenity --title="Send Gmail" --info --text="Exiting. No
file attachment specified."
148.         exit
149.     fi
150.
151. # If you want the compressed file placed inside the directory
being compressed,
152. # use $1/$1.tar.gz in both the tar command and the declaration
of the FILE variable.
153. # This is the default behavior.
154. #
155. # If you want the compressed file placed in the parent directory
of the directory
156. # being compressed, use $1.tar.gz in both the tar command and
the declaration
157. # of the FILE variable (or just remove the first $1/ from the
tar command and
158. # the FILE variable).
159.
160.     if [ -d $1 ]; then
161.         tar -zcvf $1/$1.tar.gz $1
162.         FILE=$1/$1.tar.gz
163.     fi
164.     if [ ! -e $1 ]; then
165.         zenity --title="Send Gmail" --info --text="File cannot
be found. Exiting."
166.         exit
167.     fi
168.
169. getSubject
170. getMessage
171. getRecipient
172. getCC
173. getBCC
174.
175. if [ "$hasCC" == "0" ] && [ "$hasBCC" == "0" ]; then
176.     echo ${Message} | mailx -s "$Subject" -a $FILE $TO
177.     exit
178. fi
179.
180. if [ "$hasCC" == "1" ] && [ "$hasBCC" == "0" ]; then
181.     echo ${Message} | mailx -s "$Subject" -a $FILE -c $CC $TO
182.     exit
183. fi
184.
185. if [ "$hasCC" = "0" ] && [ "$hasBCC" == "1" ]; then
186.     echo ${Message} | mailx -s "$Subject" -a $FILE -b $BCC $TO
187.     exit
188. fi
189.
190. if [ "$hasCC" == "1" ] && [ "$hasBCC" == "1" ]; then
```

```
191.     echo ${Message} | mailx -s "$Subject" -a $FILE -b $BCC -c $CC
$TO
192.     exit
193.     fi
194.
195. exit 0
```

To start off with, the script defines two variables, `hasCC` and `hasBCC`, and sets them to zero. It then uses a series of functions to define the data entry dialog boxes, controlled by Zenity. These dialog boxes allow you to specify the text of your email message, the recipient's address, the subject of your email, plus any CC (carbon copy) or BCC (blind carbon copy) recipients. The first three data items are required to send an email, but not every email is sent with CC or BCC recipients.

Since CC and BCC recipients are optional, they present a special case. They also alter the command we use to send email. This is where the two variables come into play. If there are no CC or BCC recipients, the value of the `hasCC` and `hasBCC` variables remain zero. However, if the end user decides to send the email to CC or BCC recipients, or both, the value of the appropriate variable is changed to 1. It's the value of these variables which dictates which command to use to send the email (hence, the last set of four conditional if statements at the end of the script).

One other thing that is required to send an email is a filename or path and filename, passed as a command line parameter. Remember, my original intent is to be able to send an email directly from my file manager, without having to open or switch to my browser or email client. In fact, the script starts off with showing the dialog box where you can specify the filename, or the full path and filename. If you have selected the file in your file manager, and if you have the script set up to function from your file manager's right-click context menu, the filename should already be displayed in the filename entry window. But more on that a little later on.

It then calls each function, one at a time in a modular fashion, giving the user a chance to enter the required information in each dialog box. When you respond to the last dialog box (regarding BCC recipients), the file will be emailed to your designated recipients.

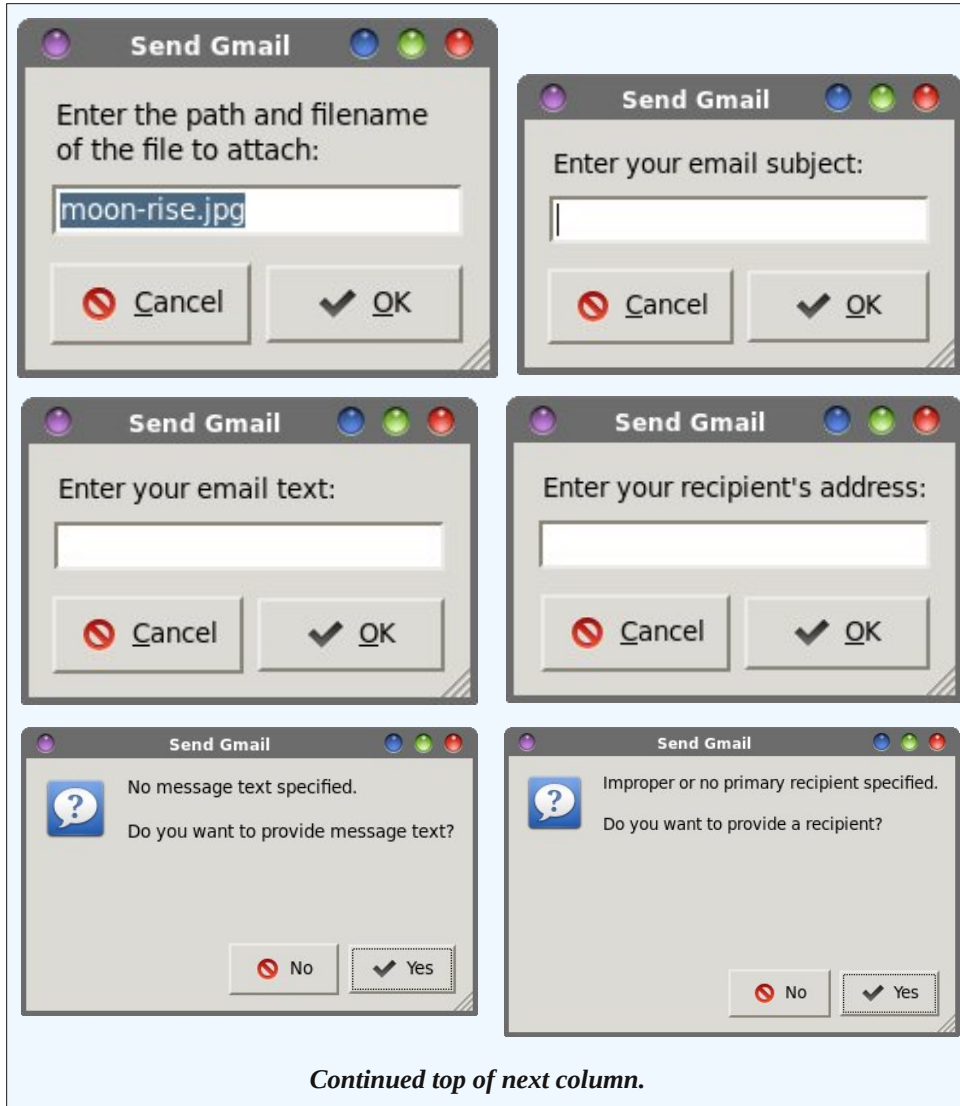
What you get

The `send-gmail` script has the potential to display 19 different Zenity dialog boxes, but it would be pointless and redundant to display all the different dialog boxes. While the script has very little error checking, it does check for the presence or absence of data. The presence of certain data (like a CC or BCC

Send A Gmail Message Without Opening Gmail

recipient) changes the command to send email. Similarly, if you enter an email address of less than 10 characters long (this limitation is in effect for the primary recipient, CC recipient and BCC recipient email addresses), a subject of less than three characters, or a message text of less than five characters, you will be given an opportunity to either re-enter the information or exit.

Here are nine of the more common dialog boxes you will see:



How to use

The send-gmail script is a fully functioning bash script that can be used from the command line. When executed from the command line, you will still get the graphical dialog boxes into which you input data. Remember that you must specify a filename (or full path and filename) as a command line parameter.

But then remember our original intent – to be able to use it from your current file manager. Fortunately, many of the modern file managers give you the ability to extend their functionality, via their right-click context menu. It would be best to store the script in a directory that is in your \$PATH statement, to make it easier to implement. Otherwise, you will have to supply the full path and filename to the script when you call it into action. The examples below assume that you have stored the script in a directory in your path.

Under **KDE**, you can set up a KDE service menu. We've covered how to do that in a previous issue of The PCLinuxOS Magazine. Because it is a lengthy description of the process (although not a hard one to implement), I'll refer you to the previous [article](#) that explains the process. The command to execute should be `send-gmail.sh %f`.

But basically, your entry should look something like what I have listed below. Be sure to save it as `send-gmail.desktop`, in `~/.kde4/share/kde4/services` if you only want the script to be accessed by a particular user, or in `/usr/share/kde4/services` if you want the script to be accessible to everyone who uses the computer.

```
[Desktop Entry]
Type=Service
Name=Send Gmail
Encoding=UTF-8
ServiceTypes=KongqPopupMenu/Plugin
Actions=SendGmail
MimeType=application/octet-stream
Icon=gmail
```

```
[Desktop Action SendGmail]
Name=SendGmail
Exec=send-gmail.sh %f
Icon=gmail
```

Under **Xfce**, you can set up a Thunar Custom Action. The command will be either `send-gmail.sh %f` or `send-gmail.sh %n` (either one will work ... `%f` is the full path and filename, while `%n` is just the filename). Under the “Appearance Conditions” tab, leave the file pattern set to the default value of `*`, and place a checkmark in front of every file type. Now, when you right click on a file, it will be attached to your outgoing email and sent to your specified recipient(s).

Under **Mate**, which is essentially a continuation of Gnome 2.x, you can set up the script for use from Caja's right-click context menu. Simply place a copy of the `send-gmail.sh` script in `~/.config/caja/scripts`. Make sure to mark the file as “executable.” If you did it right, “`send-gmail.sh`” should show up on Caja's right-click context menu, under Scripts. Now, when you right click on a file and select

Scripts > `send-gmail.sh`, the selected file attachment will be sent to the specified recipient(s).

Under **Cinnamon**, the file manager is Nemo, which is a fork of Nautilus. You can set up the script to use Nemo's right-click context menu. Just place a copy of the `send-mail.sh` script in the `~/gnome2/nemo-scripts` folder, and mark the file as executable. Now, the `send-gmail.sh` script should show up under Nemo's right-click context menu, as an entry under the Scripts menu item. Select the file you want to send as an email attachment, right click on it, and select our new menu entry. Your selected file will be whisked off to your recipient(s).

LXDE users, along with others who use PCFileMan file manager, are pretty much out of luck for using this script from within PCFileMan. The developers of PCFileMan only recently started allowing for custom extensions to the file manager. However, those extensions are written in C and have to be compiled with gcc. As far as I can tell, the only thing it allows are extensions written in C and compiled. Even then, the details are a bit sparse. There is hope (hope springs eternal, right?) that the developers will build in easier, user-accessible extensions into PCFileMan when LXDE makes the switch from being based on Gtk+ to being based on Qt.

There IS one option for PCFileMan users, though (which is more than they've ever had previously). Right click on a file or directory you want to send. Select Open With... from the context menu. Click on the “Custom Command Line” tab in the dialog box that appears. Enter `send-gmail.sh %f` (provided you've saved the script into a directory in your path; otherwise you will have to prefix the script filename with the full path) into the command entry box, then click on OK. I know that it's a lot of steps, and it may just be easier to send the file from your browser or email client.

Aside from that convoluted solution, the other solution is to open a terminal in the directory that contains the file/directory you want to send, and run the `send-gmail` script from the command line. Be sure to specify the filename, or directory name, of the file/directory you want to send. In fact, that might be way easier than using the “Custom Command Line” method we previously discussed.

In any of the recipient entry dialog boxes (primary recipient, CC recipient or BCC recipient), you can enter multiple email addresses simply by separating each address with a comma. For example, `pclinuxos.mag@gmail.com, yourfriend@somedomain.net, someoneyouknow@thisdomain.com` would email the file to all three addresses.

By default, the script is set up to send single file attachments. It certainly can be extended to handle multiple file attachments, but this avoids the chore of listing through multiple files and adding each attachment to the outgoing email. If you

want to send multiple files as an email attachment, it would be best to package them all up in a compressed archive file (.tar.gz, .zip, .bz2, .xz, etc.) and send the archive file to your recipient(s).

However, there is one exception. If you select a directory of files to send to your recipient(s), the send-gmail.sh script will automagically create a .tar.gz file, with the same name as the directory, store that archive file within the selected directory, then send that same archive file to your recipient(s). By storing the archive file within the directory being archived, we keep things nice and tidy. If, though, you want the archive file to be saved in the parent directory of the directory being archived, you can easily change the script. Simply alter lines 161 and 162 in the listing above to read as follows:

```
tar -zcvf $1.tar.gz $1
FILE=$1.tar.gz
```

Similarly, you can alter the tar command to produce an archive file type other than a .tar.gz file. Changing the “z” in the tar command to “j” will create .tar.bz2 files, and changing the “z” to “J” will create .tar.xz files. Be sure to change the file extension to match the compression algorithm you choose to use. Also, if you want to use ZIP archives (for sharing with your Windows friends ... even though many Windows utilities exist that allow them to read a .tar.gz file, such as WinZIP and PeaZip), you can change the commands to read like this:

```
zip -r -x $1.zip $1.zip $1
FILE=$1.zip
```

One other quick point: the commands that deal with filenames (which is pretty much the whole script, since it's centered around sending an email with a file attachment) work very well with filenames that **don't** have spaces in them. If you are in the habit of using spaces in your filenames and directories ... STOP! Linux is not Windows, and it was Windows that introduced this abominable practice. We've covered this topic many times before here in the

magazine, and it's a debate that goes on ad nauseum. You should NOT use spaces in filenames on a Linux file system. Spaces are used to separate command line parameters, so spaces in filenames and directory names are a BAD idea, since the command line – and bash scripts – see a space as a delimiter for the command line parameters. It's a foregone conclusion that we won't solve that debate here.

If you find the spaces-in-your-filenames-and-directories habit too hard to break, though, simply put double quotes around the filenames in the script. I'll leave that for you to do on your own. Call it “diversional therapy.” If you have a need for this script and you still put spaces in filenames (a dash or underscore works just as well, if you must), maybe forcing you to make the changes yourself will help underscore the reasons why you need to stop using spaces in your filenames.

Caveats

Remember that the send-gmail script has very little error checking built in. As such, send-gmail has no idea if what you entered as an email address is a proper email address or not.

There is a daily 500 email limit on the number of emails that you can send using Gmail's SMTP servers. That limit is applied over a rolling 24 hour time period. So, if you reach the limit at 1 p.m., your Gmail account will be unable to send emails via Gmail's SMTP server until 1 p.m. the next day. I've never reached that limit (I don't send THAT many emails), but that may be an important consideration for users who send a LOT of emails. Google does this to prevent their SMTP servers from abuse by email spammers and mass marketers. In the event that you need to send more emails than the applied limit on Gmail, it might be better to use your own SMTP server to handle the task.

Summary

As you can see (if you made it this far), the send-gmail script can save you steps and time when you need to send a file as an email attachment. Certainly, there are areas where it can be further extended. For example, it wouldn't be a terribly hard task to display a file selection dialog box in the event that a filename wasn't specified. It also wouldn't be a terribly hard task to allow the user to select multiple files, create an archive file from them, then attach that archive file to the outgoing email. Nor would it be too difficult to allow the user to select which type of archive file to create. You could get as fancy as you like, adding in the ability to check if the entry is a bona fide email address.

Already, this script has streamlined my ability to send files as email attachments. I hope you find that it does the same for you. I've attempted to make the use of this script as desktop-environment agnostic as I could. If you come up with some enhancements to this script, I'd love to hear from you. Drop me a PM in the PCLinuxOS forum, or send me an email at pclinuxos.mag@gmail.com.



Making Files Immutable

by YouCanToo

I have a web hosting customer who keeps deleting random files and then complaining things are not working as they should. Linux comes to the rescue with the **chattr** command.

Making a file Immutable, which even the root user cannot delete.

im·mu·ta·ble

adjective

adjective: **immutable**

1. unchanging over time or unable to be changes.
2. "an immutable fact:
3. fixed, set, rigid, inflexible, permanent, established, carved in stone.

Here is a cool tip on how you can make files on your system immutable. By immutable, I mean even root can't delete the files if he chooses to. Linux ships with a tool called **chattr** which can be used for the purpose. 'chattr' is similar to the 'attrib' DOS equivalent tool but much more powerful and flexible.


To make your file immutable, open a console window as the root user and enter the following command

```
[root@localhost dwmoar]# chattr +i filename
```

NOTE... You can only do it logged in as root.

Here the **+i** option sets the immutable bit for the file. Once this bit is set, even root can't delete or tamper with the file.

If you want to unset the immutable flag, just run the following command:



```
[root@localhost dwmoar]# chattr -i filename
```

You can check what are the attributes of a file by using the following command:

```
[root@localhost dwmoar]# lsattr filename --i----- filename
```

If the immutable flag is set, there will be an 'i' in the listing. This command is used by system administrators to restrict the users from changing a

file in a particular way or even the administrator can by mistake delete a critical file because of a mis-typed command. But if the immutable flag is set, these mistakes can be avoided.

chattr can be used to set/unset many more file attributes. If you want to allow everybody to just append data to a file and not change already entered data, you can set the append bit as follows:

```
[root@localhost dwmoar]# chattr +a filename
```

Now the filename can only be opened in append mode for writing data. You can unset the append attribute as follows:

```
[root@localhost dwmoar]# chattr -a filename
```



PCLinuxOS Goes To Work

by Paul Arnote (parnote)

Most of us use PCLinuxOS in our daily, personal lives. We're comforted by the security and lack of viruses, malware, spyware, crapware and all of the other kinds of evil "wares" that infest that "other" commercially available operating system.

But there are some PCLinuxOS users who depend on it to help them earn their income. Yes, they have "taken PCLinuxOS to work." Here are the stories of some of those users who use PCLinuxOS in their work.

Georgetoon

Technically, I'm 50/50 in using PCLinuxOS for business. I use gLabels for business cards. As you know, I use GIMP to create all my comic features for my business.



For bookkeeping, I use Quickbooks. But I run Windows in Virtualbox within PCLinuxOS to accomplish this.

I am looking at possibly switching to GnuCash. GnuCash appears to do the same things as Quickbooks, but the GUI is quite different. It's unfamiliar. Because I've used Quickbooks since 1994, there's a steep learning curve to GnuCash.

From what I've seen and test-driven, it's a good business program. I'm hoping it's as complete as Quickbooks. I'm hoping for a situation to arise which would allow me to make the switch. Right now, there's too much accounts history with QB. I need a situation where I'd be making a fresh start from scratch.

When I first launched GeorgeToon, Inc., I used Kompozer to build the website. When my nephew graduated with a degree in computer science, I hired him to design the web site. He used WordPress, so I can now just do all my blog updates and page updates within Chrome in PCLinuxOS.

kc1di

I use PCLinuxOS in my Ministry. I'm a minister at a small church in Maine. I use LibreOffice for most documents and letters, and Thunderbird for all emails and prayer letter updating. I use Kompozer and Seamonkey for web pages content and updating, and Filezilla for Webpage transfers. I use Okular to read PDF files.

I use PCLinuxOS for my home office networking, and use Simple Scan to scan documents and photos. I use Mirage and Gimp to edit photos, and Skype for video conferencing. To top it all off, I use

Xiphos for Bible Study, along with E-Swords, which I run under Wine.




I've been using PCLinuxOS exclusively for the past two years, and before that, off and on. I have it installed on four of our computers here. And, it just works most of the time. Only one computer does not have PCLinuxOS installed on it, and that is my wife's laptop, which has Ubuntu LTS on it. I installed it because I couldn't get PCLinuxOS to play well with the UEFI/Secureboot on it. And, she likes Unity. Go figure.

YouCanToo

I run a small web hosting and design service using PCLinuxOS. I have about 400 customers currently. You might have heard of one of them before: Habitat for Humanity, Oregon.

I started using PCLinuxOS for my web server in May 21, 2007, with PCLinuxOS 0.93a, a.k.a. Big Daddy. Up to that point in time, I had been using RedHat. Currently, I have eight servers, and I am using a mix of PCLinuxOS and Centos 6. Of the eight web servers I'm running, six of them are running PCLinuxOS. The other two, running Centos, are

located out of house. One of them is a dedicated server, and the other a VPS (Virtual Private Server). Shortly, both of these will be changed over to PCLinuxOS servers.



All my invoicing is done using customized software running on a PCLinuxOS box. All business cards, mailing labels, letterheads and website layout was designed on PCLinuxOS with software found in the repository, such as gLabels, Bluefish and a good old text editor.

Unlike what a few might want or have you think, PCLinuxOS have been a very stable platform on which to run not only your desktop/laptop, but also a web server.

ghostbunny

I'm employee in a small IT business. The primary focus of the business is software development, in the area of electronic signature, identity management and smart cards. We have a "bring your own device" policy there, so I always take my PCLinuxOS laptop with me. I really do everything with it.

I develop applications with Netbeans and GlassFish. They are not in the repository, but they work well locally installed in my home directory. Netbeans uses the Java SDK from the repo also. I use Thunderbird for the business mails, and a lot of other repository stuff like PCSC for our smart card related work.

I use Texlive, in combination with Texmaker, to write documents, as well as LibreOffice and a very buggy MS Office, under Wine, for general writing. Other applications in the company are web based and can be accessed via a normal browser, so there is no really need of an MS operating system.

Ramchu

I use PCLinuxOS as my everyday OS, so I guess that I can honestly say that I do use it in my HVAC business.



I used LibreOffice to design and fill in my billing invoices. I also use it to create a clickable PDF work order for the mobile home dealers that I do business with, so that they can quickly fill it in and email to my office.

I have also used PCLinuxOS with LibreOffice to create and print out my business cards.

I also use the Kraft small business package software installed in PCLinuxOS to create detailed bid estimates and presentations. I tried to get my wife to use it for our bookkeeping, but unfortunately she is a diehard QuickBooks user in Windows.

dm+

I am an electrical engineer and I have been using GNU/Linux - exclusively - on the desktop since 1999. PCLinuxOS found its way to my desktop in 2007 and hasn't left since. The preference is mainly due to the package mix in the standard distribution, as well as aesthetics.

OpenOffice, and lately, LibreOffice is my main tool. I use it to create, edit, print, and share technical documents, drawings, plans, reports and presentations, including calculations and charts in more than one language, and in more than one format.

Batch text and image processing tools, ranging from individual applications (Gimp, the ImageMagick suite, awk, sed) to office macros and command line scripts, have seen heavy use at times. When some programming is needed, I use whatever editor is available (Kate, Kwrite, mc, vi) to write my code. I then use diff and kdiff3 to track changes and maintain versions and backups.

Sylpheed is my preferred email client, and various web browsers (Firefox, Konqueror, and the Lynx family) are used to keep in touch with the world and my colleagues. Encryption is also used, mainly to safeguard the laptop and the removable disk contents from unintended disclosure (CryptSetup, shred).

Part of my work involves communications with various embedded devices and/or networked

systems. ftp, ftp, ssh, SMB/CIFS (samba), ntfs, vnc, but also snmp clients and/or servers are often used to that end. Wireshark/tshark have been a great help in troubleshooting such "open bench" connections.


To be productive, you need a pleasant environment. Enter the audio players. XMMS and DeaDBeef liven up the ambiance as tolerated by the circumstances.

loudog

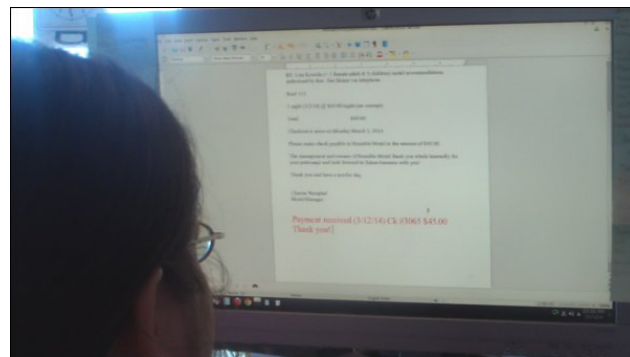
Managing a budget Motel and RV Park can be very satisfying, especially when it's located at the base of some very pristine mountains in the least populated state in the US. You get to meet many interesting people from many walks of life and different cultures who come here for the tranquility of sparse development, and traditional values. I enjoy pitching PCLinuxOS when I can, usually with a patron who is having trouble logging into the motel's wireless LAN. At other times, it can be very, very stressful, dealing with things ranging from stopped up commodes, frozen pipes, extinguished pilot lights, leaking faucets, to – ugh – shoveling and plowing snow. From the unruly and flat out intoxicated and menacing guests (a personal favorite), to the nice people who just can't pay today but promise, promise, promise to pay at some later date, it is always an enjoyment at the end of the day to log into PCLinuxOS and it just works, for business and pleasure.

We use PCLinuxOS for everything here, business wise and personal. From doing the billing (the secretary Charise, my wife, engaged in billing, above), the production of business cards, recording in the daily log, creating warnings, weather required by the local government or newly discovered by staff, notices, guest records, no trespassing orders and the big nastygram: eviction notices.

Notice the sweet little devilish smile on the secretary's face as she writes an eviction notice to some exceptionally problem tenants. Producing some of these documents is unfortunate, but a necessary part of Motel / RV park management. I mean really folks, how many times do you need to see these? They usually receive two to three notices, per guest, before eviction. We've been through this 47 times total in 2 years 10 months according to the log, so it's no wonder all the emergency response staff in town know my wife and me by sight.



The firefighters are also frequent visitors, and their visit is always associated with cooking foolishness or jury rigged electrical and or heating of a Recreational Vehicle, often resulting in serious injury (9 times).



Ambulance personnel are also on a first name basis, almost exclusively the result of drug overdose (whether prescribed or illicit drugs), and/or alcohol and fights (29 times, not including the coroner).



This is me, Lou, the manager, reading the nastygrams just before delivery to the lucky recipients. It's not just a job, it's an adventure. Thankfully, Tux is always there, never causing quite enough trouble to require a written notice himself, but certainly glad to help with producing them when necessary. Yes my friends, PCLinuxOS is definitely a secure and deeply appreciated part of the business.



Screenshot Showcase



Posted by nymira, March 15, 2014, running KDE.

PCLinuxOS Puzzled Partitions

			8		7	1	4	
4	5		1	6				
		7		5		2		
9							8	
	1	4				6	2	
	3							9
		3		8		4		
				2	4		3	6
	4	9	5		1			

SUDOKU RULES: There is only one valid solution to each Sudoku puzzle. The only way the puzzle can be considered solved correctly is when all 81 boxes contain numbers and the other Sudoku rules have been followed.

When you start a game of Sudoku, some blocks will be prefilled for you. You cannot change these numbers in the course of the game.



Each column must contain all of the numbers 1 through 9 and no two numbers in the same column of a Sudoku puzzle can be the same. Each row must contain all of the numbers 1 through 9 and no two numbers in the same row of a Sudoku puzzle can be the same.

Each block must contain all of the numbers 1 through 9 and no two numbers in the same block of a Sudoku puzzle can be the same.



SCRAPPLER RULES:

1. Follow the rules of Scrabble®. You can view them [here](#). You have seven (7) letter tiles with which to make as long of a word as you possibly can. Words are based on the English language. Non-English language words are NOT allowed.
2. Red letters are scored double points. Green letters are scored triple points.
3. Add up the score of all the letters that you used. Unused letters are not scored. For red or green letters, apply the multiplier when tallying up your score. Next, apply any additional scoring multipliers, such as double or triple word score.
4. An additional 50 points is added for using all seven (7) of your tiles in a set to make your word. You will not necessarily be able to use all seven (7) of the letters in your set to form a "legal" word.
5. In case you are having difficulty seeing the point value on the letter tiles, here is a list of how they are scored:
 - 0 points: 2 blank tiles
 - 1 point: E, A, I, O, N, R, T, L, S, U
 - 2 points: D, G
 - 3 points: B, C, M, P
 - 4 points: F, H, V, W, Y
 - 5 points: K
 - 8 points: J, X
 - 10 points: Q, Z
6. Optionally, a time limit of 60 minutes should apply to the game, averaging to 12 minutes per letter tile set.
7. Have fun! It's only a game!




Possible score 204, average score 143.

Download Puzzle Solutions Here

PCLinuxOS Crossword Puzzle: April 2014

Occupations



1. Runs the cash register at the store
2. Assists in scientific research or medical diagnosis
3. Uses special metal rods & electricity to bind metal together
4. Doctor for your teeth
5. Educates our children
6. Grows the food we need
7. Fills your medicine prescriptions
8. Repairs vehicles
9. Builds furniture, houses & other buildings
10. Raises animals
11. Doctor for animals
12. Repairs and programs computer systems

[Download Puzzle Solutions Here](#)

Occupations Word Find

X P Z G W D L D B V S D W C U Z R L N W N Z B K S S O E F J
 U Z G Y L Q C A Y S E C R E T A R Y Q U Q W E T O D H D D Q
 A P G Q X O J M X R U U Y E O P B Q B Y N D D N A D D K H S
 R C S Y X O Z E R S C K P R G B T P L Y V G O K R I S R I F
 E D H P D J P S V A E D P K S M I Y N N G C K C O K U O Y V
 S K Q X F P W V R M P S P A Q P L K X H L Y A D T A X U S E
 P N J S R X L P E N T P R E Q B T Q H Y B L E I A O V A K Z
 I F G T S A E I R F L R F U K M A N N O U N C E R P R X Q X
 R C W O L N Y B G G J M F T N R S B S P H B V R T V Q D B P
 A T R R T E U T C O M P U T E R S P E C I A L I S T L I R S
 T S D E F O J V E X M B T H Z I N T B T G I Q K I R A F E X
 O I R C I K O F H C A F C L N K S N K U B C O C N P B J K X
 R C L L M H C U Y W H N H W D I N E S T H O Z X I S O O R I
 Y A X E Q A S G B N A N F N T D Z D S F Q I Q M M R R J O V
 T M E R Q E U A B R P V I N I E A I P X T C E M D F A N W D
 H R W K E K H A C L R S E C G N G S O T S I T R A Z T W N N
 E A T K Y K N P Q A H D P K I W J E L Q Y G O L E Q O E O A
 R H E N R K Y G D Z C K M O B A A R I Y I Z T T N U R L I M
 A P E D E K K D K X U O A G H R N P C E E J L E E E Y D T S
 P Q G R M G W W D M B V W C H E R Y E Y D B K I V M W E C E
 I C D E R M A D Y J N W S Z I E Z X O C O U N S E L O R U L
 S K K I A E L E N H W I G I G N M R F U A D O R U I R S R A
 T H U V F Y Y B C V T X B X O I A W F E G U D D T F K L T S
 G J Z H H A Z W Q N V G X I H G V H I U H M P B E R E N S H
 H I O L T E K C A O A B J I D N Z G C H T V P C A J R P N Y
 Z W G X L C U D V L J R E N T E A T E E H M V W C X W H O R
 V E T E R I N A R I A N U P U K P I R O M C X C H Z D I C E
 Y E Q L D X G M L U S R L S Y R K S V S K K O L E I F J L D
 P I Q X L E Q D C H W V F B N X C N R K G B K N R M R B D E
 P S K I Q V R J V C G U Z T F I H L W H X R D B P G G C B Z

- Administrator
- Announcer
- Artist
- Banker
- Carpenter
- Cashier
- Computer Specialist
- Construction Worker
- Counselor
- Dentist
- Doctor
- Engineer
- Farmer
- Insurance Agent
- Laboratory Worker
- Lawyer
- Mechanic
- Nurse
- Pharmacist
- Police Officer
- President
- Rancher
- Respiratory Therapist
- Salesman
- Secretary
- Store Clerk
- Teacher
- Veterinarian
- Welder
- XRy Technician

[Download Puzzle Solutions Here](#)



Programming With Gtkdialog: Part One

by Pete Kelly (critter)

With a standard installation of PCLinuxOS, you get hundreds of applications. With access to the repositories, you have access to thousands more. This is nice, but I actually only use about a couple dozen on a regular basis, despite having very many more installed.


To get at any of my favourite applications meant opening a menu and clicking through several sub-menus to find it, or to create keyboard shortcuts, which I could never remember. I could, of course, have put shortcuts on the desktop, but I prefer an uncluttered desktop. Plus, the desktop is usually covered by the current application, anyways.

To solve the problem I wrote a little utility using a program called **gtkdialog**. This utility pops up a little dialog box with one button for each application. Each button has a suitable icon, and will display a tooltip describing the application when you hover the mouse over it. As an added advantage, the dialog can be made to appear centred around wherever the mouse happens to be. Clicking one of the buttons closes the utility and launches the chosen application.



All of the items in this dialog, that's buttons, titles, dividers etc., are known as widgets. To make this program, you need a bash script, which you don't need to understand – although it would be helpful if you did, along with some plain text which describes, in a series of stanzas (groups of lines), each of the widgets that you want to display. The text in the stanzas uses markup tags to communicate with **gtkdialog**. Once you have seen a few examples of this, these tags are pretty much self-explanatory.

To get the layout correct, you also need to understand the structure of the text in the script. Let's start with a simpler, 2x2 button dialog with a cancel button. The process is the same, no matter what size you want. The contents of the dialog are a series of horizontal and vertical boxes, which you join together like Legos. Into the boxes you put widgets, such as buttons. Your text stanzas then describe the appearance and actions of the widgets.



In the dialog window, we start with a vertical box (blue Lego brick), add 3 horizontal boxes (red bricks), separated by 2 horizontal separators (cyan bricks, only one is highlighted in the image). In the top 2 horizontal boxes we add frames (green bricks), which help to describe and group the contents. In these frames we add our user-defined buttons (brown bricks), and then finally, in the bottom horizontal box, we add a **gtkdialog** predefined button which enables us to abort without actually selecting anything.

The code to display a **gtkdialog** can be as simple as this:

```
#!/bin/bash
export MY_DIALOG='<vbox><button cancel></button></vbox>'
gtkdialog --program MY_DIALOG
```



The simple script displays the simple dialog shown above. Clicking the button makes it go away. What did you expect?

You can write the text in any plain text editor, but a programmers editor such as kate, SciTE or vim might be more suitable. After saving the text, the file needs to be made executable. Right click on the file in your file manager, select properties, permissions and check the 'is executable' box. Alternatively, in a terminal, type `chmod u+x my_dialog`, or whatever you called it.

Having a closer look at this short script we find that:

Line 1 tells the system that we want bash to interpret the script.

Line 2 puts the text that is our program into a variable, or container, called MY_DIALOG, and exports it to make it available to the gtkdialog application. This could actually be called anything you like.

The text tells gtkdialog to put a vertical box (that's done by the `<vbox>` tag) into the dialog, and to put a cancel button into the vbox (this is done with the `<button>` tag). The cancel button is one of several pre-defined gtkdialog buttons. The others are yes, no, ok and help. The `</button>` and `</vbox>` tags terminate the button and vbox definitions, respectively.

Line 3 finally calls the gtkdialog application, telling it to use as its program instructions the text we put in the variable MY_DIALOG. Note the two dashes before the word program.

You could simply put more and more text into line 2 between the quotes to make a much more complex script, but this has two disadvantages. First, it would be very difficult to maintain or change, and second, the shell script part is just a little too basic.

Below is the same script rewritten to address these issues. You don't need to understand the shell script part, in grey text, as this is simply used as a 'wrapper' for gtkdialog. Only the text between the quotes is important to us here, and this we can keep simple.

```
#!/bin/bash
DISPLAY=":0"
exec 9>/tmp/pid
if ! flock -n 9 ; then
    zenity --warning --text="GTK-Launcher is already running";
    exit 1
fi
```

```
export MY_DIALOG='
<vbox>
  <button cancel>
</button>
</vbox>'
gtkdialog --program=MY_DIALOG
rm -f /tmp/pid
Exec 9>&-
```

I have included indentation in the text to separate the different widgets, this makes following the code a little easier. An even better idea is to use a text editor that supports syntax highlighting. Most modern text editors such as kate, scite or gvim do support this, although I have yet to find one that natively supports gtkdialog type code. As a work-around, I usually change the highlighting support to a language such as HTML or XML. This is not perfect, but it does help.

The shell wrapper script stuff, for those who are interested, works like this:

DISPLAY=":0" lets the X windows system know on which display it should display the dialog, Linux is after all a multi user system. This is sometimes necessary if you want to launch the script from another GUI application or another login session.

The rest of the code is to intercept an already running instance of the script by looking for a temporary file opened when the script starts and act accordingly. I called zenity to display the warning dialog here where I could have created another gtkdialog. The reason is that it's simple, it works and I use the same style header in many other scripts which makes for good programming practice.

Plus a bit of clean up stuff.

Okay, looking at the code to produce the 2x2 launcher utility, ignoring the grey bits, we have this (I have changed the MY_DIALOG variable name to LAUNCHER). The text colors correspond to the colors in the dialog diagram.

```
export LAUNCHER='
<window window_position="2" title="GTK-Launcher" icon-name="fork"
resizable="false">

<vbox>
  <hbox>
    <frame Office>
      <hbox>
        <button tooltip-text="Word Processor">
```

```

<input file>/usr/share/icons/
  libreoffice4.2-writer.png</input>
<action>libreoffice4.2 --writer
  &</action>
<action>rm -f /tmp/pid</action>
<action>EXIT:ok</action>
</button>
<button tooltip-text="Spreadsheet">
  <input file>/usr/share/icons/
    libreoffice4.2-calc.png</input>
  <action>libreoffice4.2 --calc
    &</action>
  <action>rm -f /tmp/pid</action>
  <action>EXIT:ok</action>
</button>
</hbox>
</frame>
</hbox>
<hseparator width-request="140"></hseparator>
</hbox>
<hbox>
  <frame Info>
    <hbox>
      <button tooltip-text="Address Book">
        <input file>/usr/share/icons/
          kaddressbook.png</input>
        <height>40</height>
        <action>kaddressbook &</action>
        <action>rm -f /tmp/pid</action>
        <action>EXIT:ok</action>
      </button>
      <button tooltip-text="PCLinuxOS
        Documentation Portal">
        <input file>/usr/share/icons/
          bookcase.png</input>
        <height>40</height>
        <action>pclos-docuportal &</action>
        <action>rm -f /tmp/pid</action>
        <action>EXIT:ok</action>
      </button>
    </hbox>
  </frame>
</hbox>
<hseparator width-request="140"></hseparator>
</hbox>
<hbox>
  <button cancel>
    <action>rm -f /tmp/pid</action>

```

```

      <action>EXIT:Cancel</action>
    </button>
  </hbox>
</vbox>
</window>'
gtkdialog -program=LAUNCHER

```

I cheated a bit here to make the text fit the page but you must give the complete url (file location address) when specifying external items such as icons.

You can see in the code listing how strict indentation and syntax highlighting help to make the code easier to read. Using the Lego analogy again we have red `<hbox>` bricks, brown `<button>` bricks, green `<frame>` bricks and so forth. We just click them together to assemble our script.

At the top we have a line that starts `<window>...`

This describes the `gtkdialog` window that will appear on screen.

`window_position="2"` Position 2 means start centered wherever the cursor is.*

`title="GTK-Launcher"` The title that will appear in the window title bar.

`icon-name="fork"` The icon that will appear in the window title bar and the taskbar.

`resizable="false"` The dialog cannot be re-sized.

**Window_position 1 is centred in the (current) screen. Position 0 is bottom right of the default screen but this can be adjusted by passing X and Y offsets if you need an absolute position.*

The window widget can contain only one child widget. For this reason you will generally want to start with a `vbox` or `hbox` widget, each of which can hold multiple child widgets.

Next we have a `vbox` containing a `hbox` which contains a `frame` with the title 'Office'. Inside that is another `hbox` containing two buttons.

I have highlighted the stanza that completely describes the first button.

Hovering the mouse over this button raises the tooltip "Word Processor." The icon used on this button is `/usr/share/icons/libreoffice4.2-writer.png`. Clicking on the button will initiate three actions:

1. Launch LibreOffice writer The ampersand (&) tells the script to continue this script after launching writer, not wait until writer is terminated which would prevent any further use of the launcher.
2. Delete the temporary file we created in the bash code to prevent multiple instances of the launcher being run. This is necessary if you use the bash code that I included, or the existence of the temporary file will prevent further runs of the application.
3. Exit the application returning a status of 'ok.'

Once you have your script complete, you will want an easy method to launch it. On my Thinkpad laptop, I have a large blue button which is only used by the system at boot up (it was originally also used in windows before I exorcised it). This is an ideal candidate. I simply hit the button, and the utility appears where the mouse is. My desktop keyboard has a similarly useless key - the calculator key. With a bit of imagination, the calculator logo on the key-top looks a little like the launcher utility, so I use that. Most desktop environments have a way of associating a key press with starting an application, or you could set it up using xbindkeys, see [here](#).

You can increase the complexity of the utility by simply adding similar stanzas, vboxes, frames, hboxes and dividers as required. Add them one by one, from the outside in and test after each addition, as any imbalance, such as a missing tag, will cause the script not to run. Want another row of icons? Click in a red brick, a green brick inside that then another red containing some brown ones goes into the green one, edit the text stanzas and you're done.

Why the extra hbox inside the frame? Try removing it and see what difference it makes.

The size of items is, for the most part, taken care of automatically by gtkdialog but you may override this as I have done in the 'Info' frame to make the buttons better fit with their allotted icons, and also with the horizontal separators to control the overall aspect of the utility.

You may have heard of or seen scripts that use Zenity and YAD, instead of Gtkdialog. Zenity and YAD are actually two programs that put all the Gtkdialog commands in simpler to use "wrappers." While you can use either of these wrapper programs, you will have finer control and more options by using Gtkdialog directly.

Of course you also need to know what to put in the stanzas, although you could just copy these and change the tool tip, icon and the application you want to

execute. That would work, but there are other options you could use to personalize your dialogs. The documentation for gtkdialog is rather fragmented. There are examples of code to demonstrate all of the widgets in `/usr/share/doc/gtkdialog/examples`, although to access it, you may need to re-install gtkdialog from the repository. By far the most complete reference I have found is online at <http://code.google.com/p/gtkdialog/w/list>.

There is also some excellent documentation in the Puppy Linux forums at <http://murga-linux.com/puppy/viewtopic.php?t=38608>.

You may also find these additional [documentation](#) and [examples](#) particularly useful and helpful.

This little utility demonstrates just some of the power and functionality that you can achieve by using what is, for the most part, plain text and common sense. While it is unlikely that you will replace any of the major applications such as Gimp or LibreOffice, there is a whole lot more that can be achieved to make life easier using utilities such as gtkdialog.

This short demonstration has only scratched the surface. There is so much more, but then you knew that. This is, after all, Linux.



PC Remote Control With TeamViewer

by YouCanToo



What is TeamViewer?

TeamViewer is a software program that establishes connections between your computer (at your home or office) and any other PC all around the world within just a few seconds. You can remotely control your friend's or partner's PC as if you were sitting right in front of it.

How do I install Teamviewer?

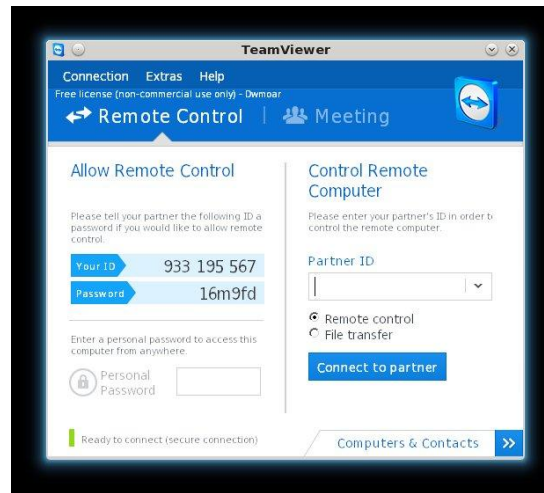
Open Synaptic, click on Search. Type Teamviewer in the search box and click on the search button. Mark Teamviewer for installation then click the "Apply" and "Apply" again and wait for it to be installed on your system. Yes, It is really that easy.

Always make sure to reload your package list in Synaptic before installing any software.

Running Teamviewer

Actually, once you have Teamviewer installed, you can run it without messing with any setting at all. You will find Teamviewer in the Internet section of your menu. The most current version of Teamviewer is version 9.0.24147-1 as of this writing.

The first time you run Teamviewer, it will ask you to accept its EULA. After agreeing to it, you will be presented with the main window.



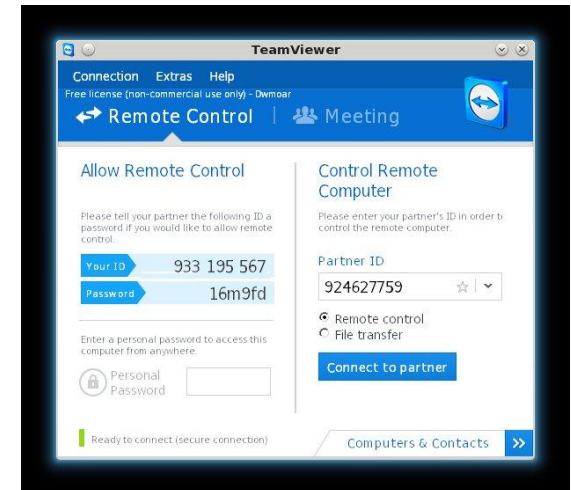
This is the window that you will use when connecting to a remote user's machine. The screen is divided into two main sections. On the right side will be the client's (remote) machine, and on the left side will be the local machine. Which side of the window you need to be using will depend on who is helping whom.

Let's say that "Slick Willie" has asked me to help him on his machine. Slick Willie lives across town, country, or on the other side of this great planet. Because it is not always convenient to drop everything, hop in the car/plane/boat and go to his place, I suggest that we use Teamviewer and take care of his problem.

Let's not forget that in order to use Teamviewer to help Slick Willie, both Slick Willie and I have to be running Teamviewer. If Slick Willie was running that

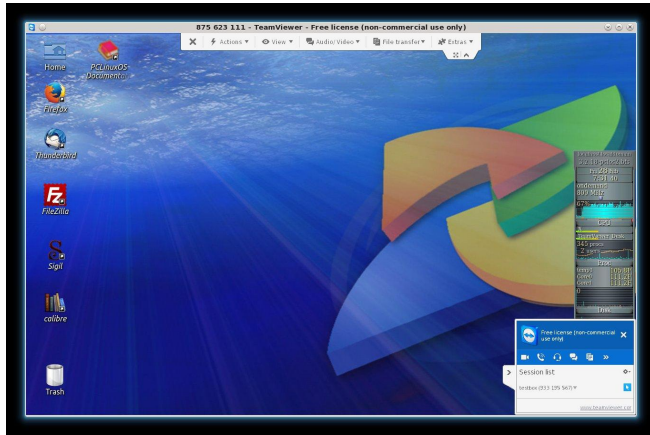
other OS, commonly named after a glass product, then he should head over to http://download.teamviewer.com/download/TeamViewer_Setup_en-ckj.exe and download it for himself. Teamviewer runs on Linux, Mac, Windows and mobile devices Android, iOS, and Windows Phone 8. Lucky for us, Slick Willie is using our favorite Linux Distro.

For us to log into someone else's system, we have to get some information from them. We will need their **ID** and **Password** from the left side of the Teamviewer window. Once we have received this information, we can set to work getting our person of need some help. In the case of Slick Willie, his **ID** 924 627 759 and his **Password** 8796. Now we can get him the help he needs. We want to put his ID number into the right side of the Teamviewer window where it says **Partner ID**,

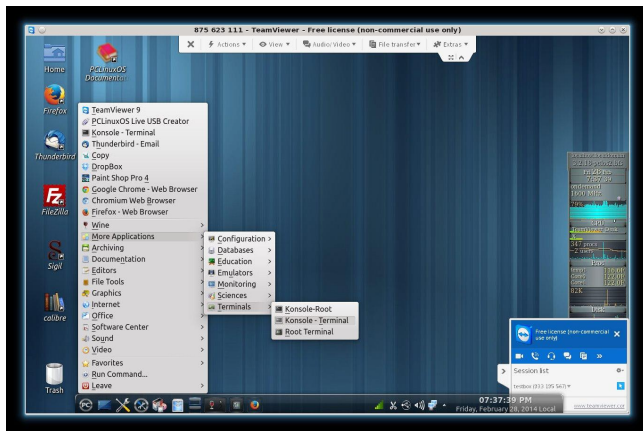


and then click the blue "Connect to partner" button. After trying to connect to his machine, it will open a password dialog box. Now type in the password he

gave you and press "log on". After a few moments you should be greeted with a remote desktop display.

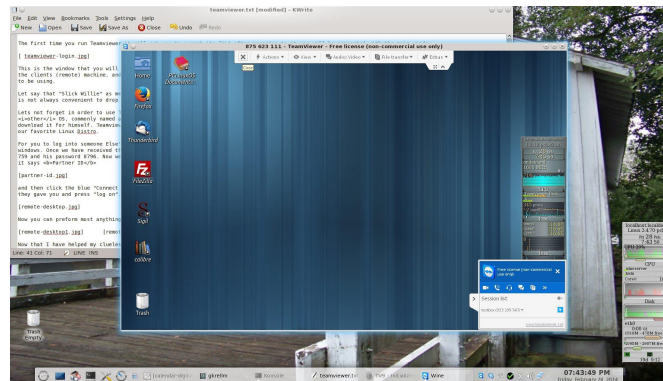
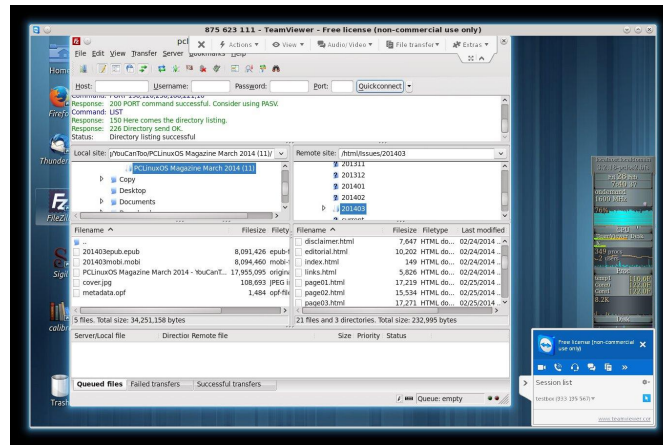


Now you can perform most anything that you could if you were sitting directly in front of his machine.



Now that I have helped my clueless friend, Slick Willie, I can log out. At the top center of the remote desktop window just click the "X" (center, right).

The session window will close and you will be greeted with the Sponsored session notice (bottom, right).



Just click on the "OK" button and you will be returned back to where you were before.

Final Thoughts

What we have done in this brief article is a very small sample of what Teamviewer can do, straight out of the box with no tweaking or fiddling around. Some places to get additional help with Teamviewer are;

Teamviewer Forum <http://teamviewerforums.com/>
 User Manuals and First Steps <http://www.teamviewer.com/en/help/firststeps.aspx>

Misc Information: It appears that many Teamviewer users are having issues when running Teamviewer 9 hanging at the initializing display parameters screen. This appears to be a Teamviewer bug. I also had one of my many machines hang at the "initializing display parameters" window. My solution was to simply revert that machine back to Teamviewer 8. After that, it worked like it should have. YMMV

Special Note: You can download a copy of the PCLinuxOS Teamviewer 8 rpm at archives.pclusers.com/teamviewer8-8.0.17147-4pclos2013.i586.rpm. Remember that is is an older 32 bit version of Teamviewer that is no longer being supported by the PCLinuxOS developers. Please do not ask them for help if you decide to use this version.

Visit Us On IRC

- Launch your favorite IRC Chat Client software (xchat, pidgin, kopete, etc.)
- Go to freenode.net
- Type `/join #pclosmag` (without the quotes)



Handy Utilities To Organize Your Life, Part One


by Meemaw

I don't know about you, but I need to document things I'm doing and things I need to remember. Most of us take notes in meetings to remember what the speaker said, especially if it's something new you need to learn. If I have my phone, I use the notepad on it, but if I'm on the computer, I like to have some sort of notes program to keep things recorded. In Synaptic, I have found several sticky-note-type programs to use.

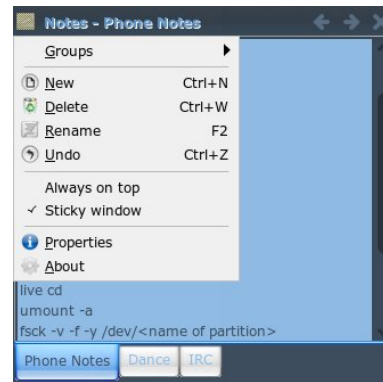
Xfce notes panel widget

I use Xfce as a desktop environment, so the first thing I looked for was a program within Xfce. That way, I wouldn't have to install anything extra. It turns out that there is already a nice sticky-note widget that can be added to my panel.

Right-clicking the panel and choosing **Panel > Panel Preferences**, a window appears with the preferences, plus a tab that says **Items**. Choosing this tab will show you the items on your panel. Clicking the plus sign will give you a list of items you can add to the panel. Choose the **Notes** panel widget and it will be added to the end of the panel. It looks like a pad of yellow paper. Right-clicking, choose **Move** and drag it where you want it on the panel. A simple left click will open a note.

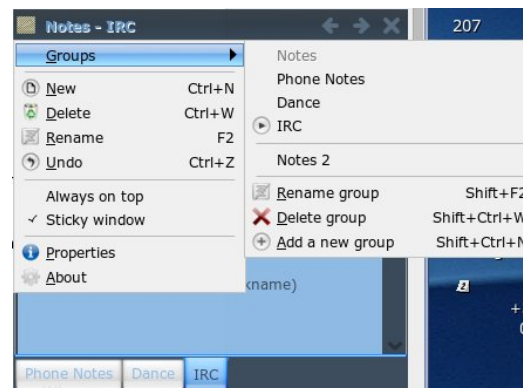


The yellow icon is there below the note window. Clicking the icon at the top left corner of the note, you will see the notes menu.



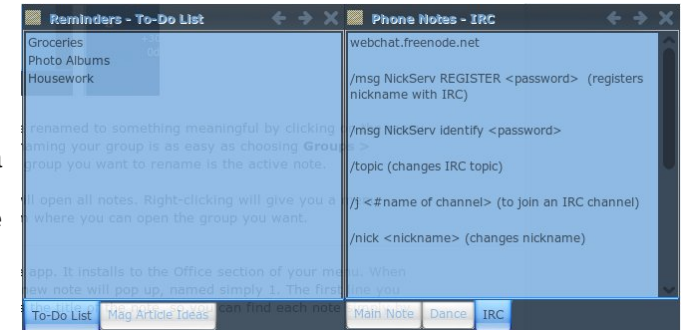
The **Properties** (preferences) menu is opened by a right-click on the panel icon, or, as you see above, from the main menu as well. There you can decide the color, font and tab position of your notes.

If you have several notes you want to keep together, choose **New** from the menu. A new tab will appear in your first note window, and you can use it. If your note is a totally different subject, choose **Groups > Add A New Group**. When you choose to add a new group from the main menu, a new notes window pops up for your use.



The note you are viewing can be renamed to something meaningful by clicking on the menu icon and choosing **Rename**. Renaming your group is as easy as choosing **Groups > Rename Group**. Make sure the group you want to rename is the active group.

Left-clicking on the panel icon will open all notes. Right-clicking will give you a menu containing the item **Groups**, and from there you can open the group you want.

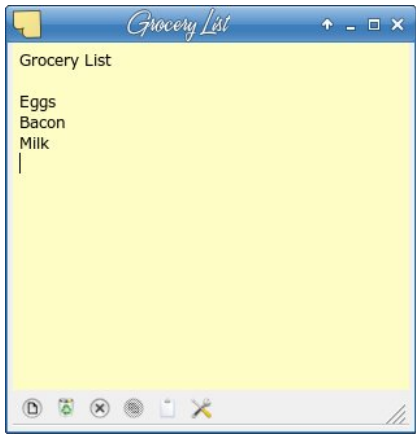


Here, I have renamed the groups and notes to suit myself. Notes can be copied and pasted to text documents.

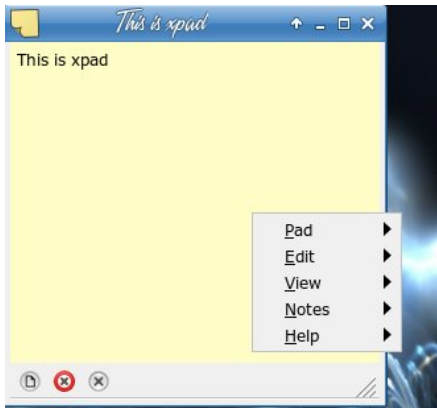
XPad

xpad is another nice sticky note app. It installs to the Office section of your menu. When you open it for the first time, a new note will pop up. The first line you type into the note will remain as the title of the note, so you can find each note simply by the title. Also visible will be an xpad icon in your system tray (next page, top left).

The main menu for xpad is easily accessed, simply by right-clicking the active note. The preferences menu is in the **Edit** menu or in the pop-up menu that



appears when you right-click the tray icon. The **View** menu has only four items: View Toolbar, Autohide Toolbar, Scrollbar and Window Decorations.

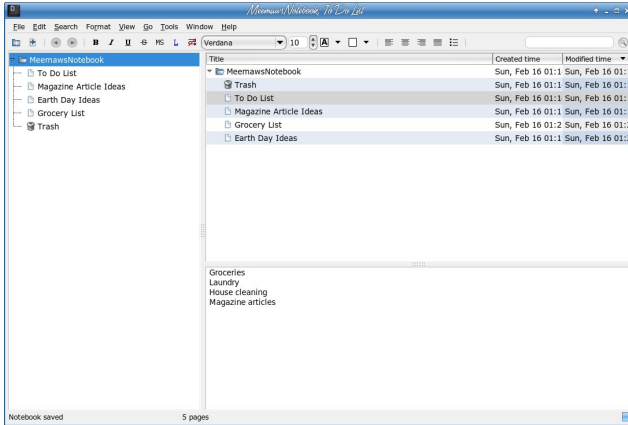


The toolbar across the bottom can be configured by right-clicking the toolbar and adding or removing your desired items.

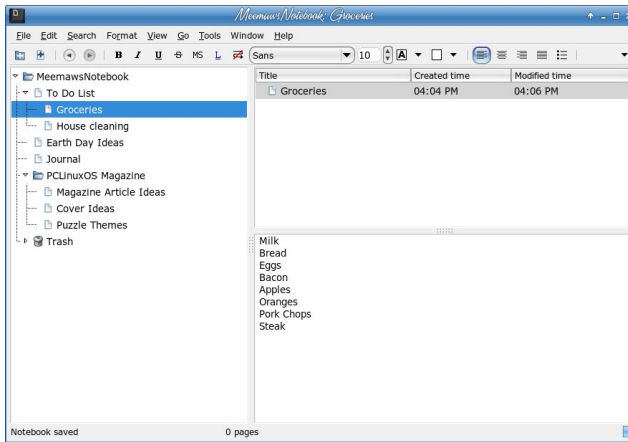
This app doesn't seem to have a group feature like the Xfce version, but it's still a nice program. Notes cannot be printed but can be copied and pasted to text documents.

KeepNote

KeepNote is more of a notebook than a sticky. When you open it, you see a window with three sections, looking kind of like a file manager. You are asked to add a notebook and save it. Mine is Meemaws Notebook. The file saved to a folder by that name in my /home, but I'm sure if you wanted it in your Documents folder you could save it there.



The tree at left will eventually show all the “pages” in your notebook. You can add pages to it, or you can add folders and “child pages”. Pages can be added to the main notebook, and child pages can be added to pages or folders. I have expanded and rearranged mine a bit now.



KeepNote has many features. You have a text formatting bar above your windows, and also have several other tools you can use. One of my pages is a journal, so I want to separate the entries. From the Edit menu, I can insert today's date before I write, and at the end I can insert a horizontal line. A screenshot or an image can be inserted, or a file can be attached.

Your notebook is searchable, as well. The easiest way is to widen your window until the search box appears at the top right, type in your desired search parameter, and click the magnifying glass.

Your notebook can be viewed in a text editor or a web browser.

One thing I'm not sure I like is that the formatting bar, apparently, can't be edited. All the formatting aids are available in the Format menu, and there are a few I use only seldom, so I could make my window smaller if I could remove them (or at least make it so the search box was visible without widening the window).

Tux Cards

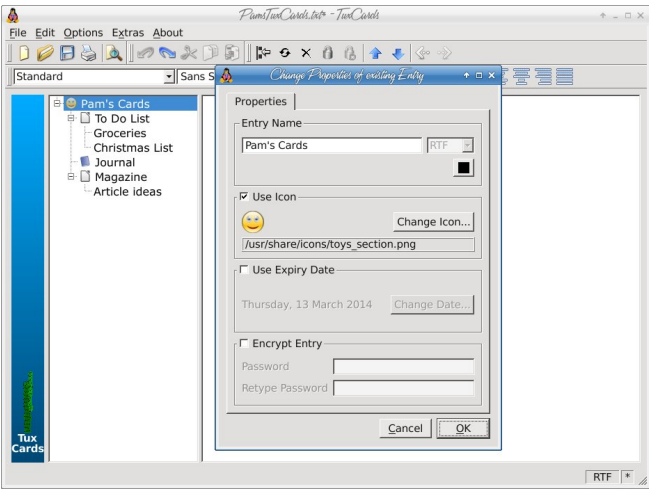
Tux Cards is similar to KeepNote. You will see a similar layout, but only two sections instead of three. You'll also see a small sidebar which says Tux Cards in it, but it has no other use as far as I can see (well, you can add a bit to it, but just for fun).

In the left section will be the word “root”, which is the starting card. You can leave it that way, or you can right-click on the word and choose **Change Properties**. A window will appear which will let you change the name of the card. you can also assign an icon, and even password-protect that card (next page, top left).

Each card you want can be created by right-clicking the main title and choosing **Add Entry**. You can also click the **Insert** icon in the toolbar.



Handy Utilities To Organize Your Life, Part One



You can add child cards under other cards, as well. I added a card named Magazine, then, right-clicking on it, I added a child card named Article Ideas. Like KeepNote, you can insert the current date and/or time, and even an image.

You can add as many cards as you want, and even create a new file with other information.

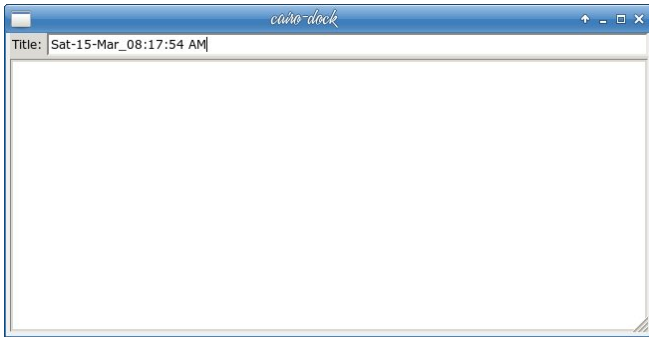
As for the sidebar, it is solely for fun. You can change the colors and the location of the text. Another fun option is the sidebar cactus. In the **Options** window, it can be configured in the **SideBar** tab. You can enable the cactus, and while you use TuxCards, random cactus plants will grow in the sidebar.

TuxCards is a little different, but it seems to be a good note organizer. The notes can be exported to HTML.

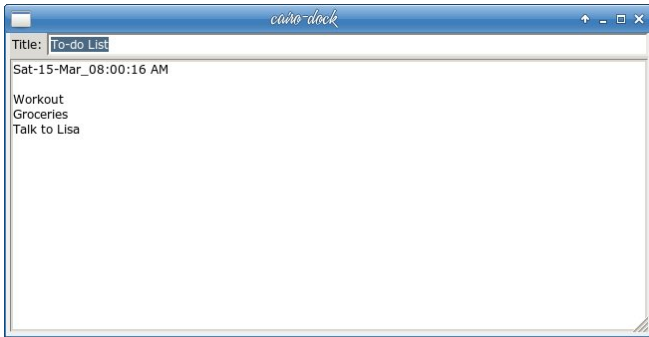
Note-Taking Applet for Cairo Dock

I noticed this one in Synaptic. The manual says you can “control your Gnote or Tomboy notes directly from Cairo Dock.” When you add it to Cairo Dock, you will see a sticky note icon with the number zero

in the middle. Middle-click the icon to add a note (or right-click and choose Add a note). A window will appear with the current date and time in the title.



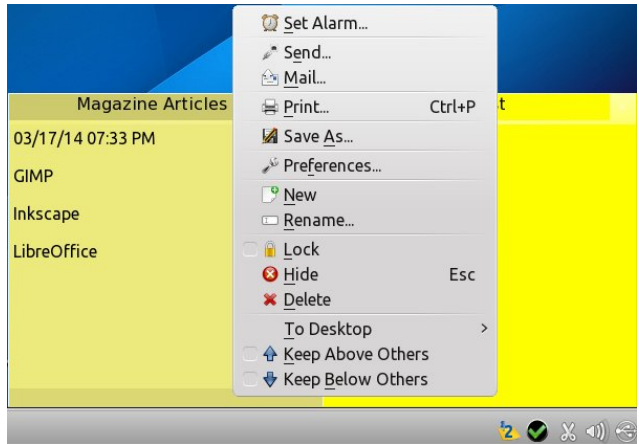
You can rename the note. I have been moving the date & time to the edit area and entering a meaningful title.



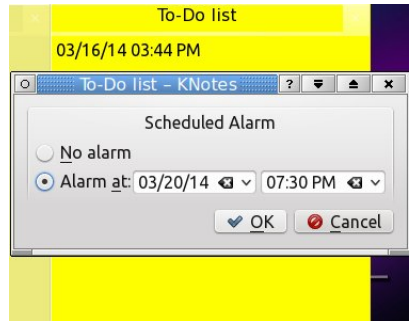
When you want to open a note, hover your mouse over the icon on the dock. Icons representing your notes will appear next to the dock, and from there you can click on the one you want to open. If you want to delete a note, right-click on it and choose **Delete This Note**.

KNotes

KDE has several programs that work well. One of the sticky notes programs is KNotes. The package kde-pim must be installed to get KNotes. When you open KNotes, your first note will appear. The current date and time will be in the title bar, but you can right-click it or the bottom bar and choose **Rename**.



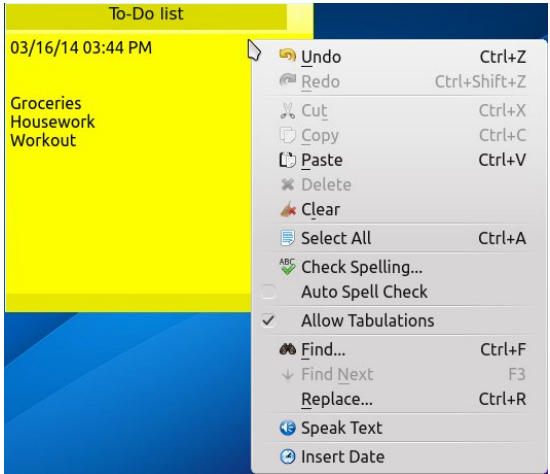
This menu also has the Preferences, which allow you to choose the color of your note and text. Notice that you can print your note, lock, hide or even delete your note. This note program also has the feature of letting you set an alarm on this note, so you can have it “pop up” later to remind you of something in your note.



If you right-click in the body of the note, you will get another menu which gives you some additional



options including cut, copy and paste, spell check, speak text, and insert date.

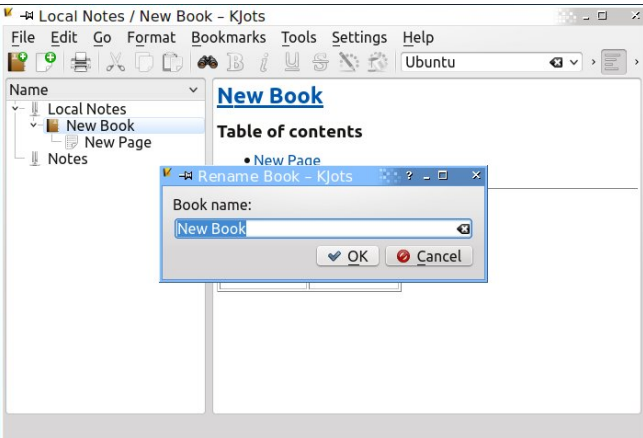


For KDE users, this is a nice sticky note program. You can print your notes, or email your notes from here, and can set it up so you can send notes to someone on your network.

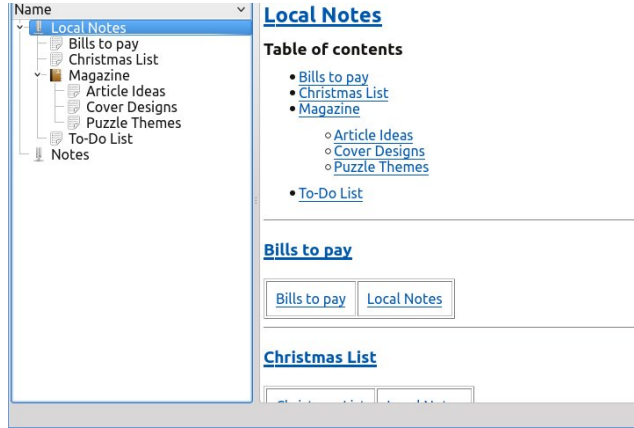
KJots

KJots is a little different, but still in the file explorer format. When you open it for the first time, it imports Akonadi information, if there is any (I don't use Akonadi), and adds pages to the tree. It will then show two sections to start: Notes and Local Notes. However, it may show two additional pages named `akonadi_akonotes_resource_0` and `akonadi_akonotes_resource_1`.

You may then start adding books and pages. My to-do list is only one page long right now, so it can be added as a page. My magazine notes, however, have several different items I prefer to keep separate, so I added a new book, then added pages to it. When you click on **New Book**, KJots adds it with a New Page already in it. You can right-click on each and choose rename (center, top).

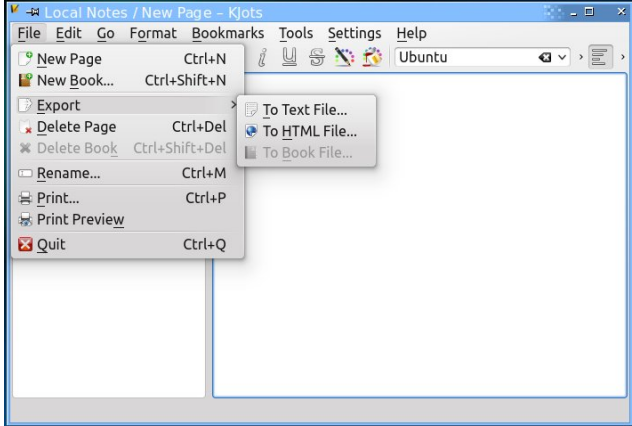


After you organize your notes into books and pages, you can click at the top of the tree and see an outline of sorts on the right. It is organized in the same manner, but it has links to your notes or even back up to the main part of the tree. This allows you to jump to different notes, if you wish. Scrolling down the right side of the window will show you all your notes, but if you just click on the links at top, you will be taken to that note farther down.



You can export your file as well. Going to **File > Export** will allow you to export everything as a text file, an HTML file, or a book file. The book file will look like a book with the section names appearing as chapter names. The HTML file opens in your web

browser, and looks very similar to the right side of the KJots window. It is handy if you want to send an outline to a co-worker. Click on the heading for that outline and choose export as HTML. KJots will export that section only, so I could send my Magazine outline to parnote without sending him my Christmas List or To-Do List. There is also an option to print your notes.

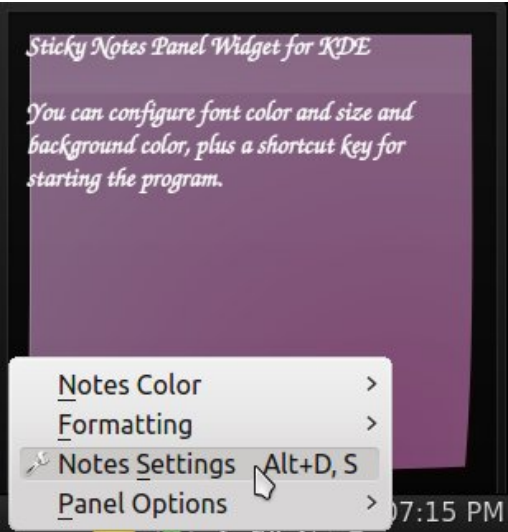


As with many of the others, KJots has many good points. I noticed, however, that the Preferences window only has a way to change the frequency the file is saved. I didn't find a way to change the default text formatting other than to change it in every window. Of course, it is using the system default, so you have to go back to KDE's Configure My Desktop and change the system default if you want it different.

Notes Panel Widget for KDE

This one was unexpected. You add this just as you add other widgets in KDE. It also uses the system defaults unless you change them in Preferences. The strange thing is that you can only have one note. You can apparently make it as large as you want, but you can't open a new one without adding the widget to the panel again and having two widget icons.





Screenshot Showcase



Posted by Just17, March 15, 2014, running KDE.

If you use KDE and need more than one sticky note to organize, you would probably be better off with KNotes or KJot.

I'm sure there are more notes apps in Linux, so if you are using a different one than one of these, please PM me in the forum and tell me what it is. Many people just use simple text files, which they update regularly. However you organize your thoughts, there is probably an app for it that you will love!



Looking for an old article?
Can't find what you want? Try the
PCLinuxOS Magazine's
searchable index!

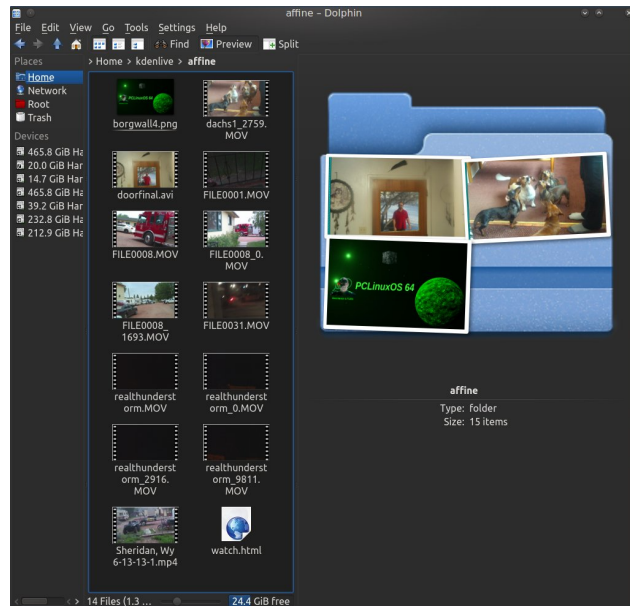
The **PCLinuxOS** magazine



KDenLive, Part Five

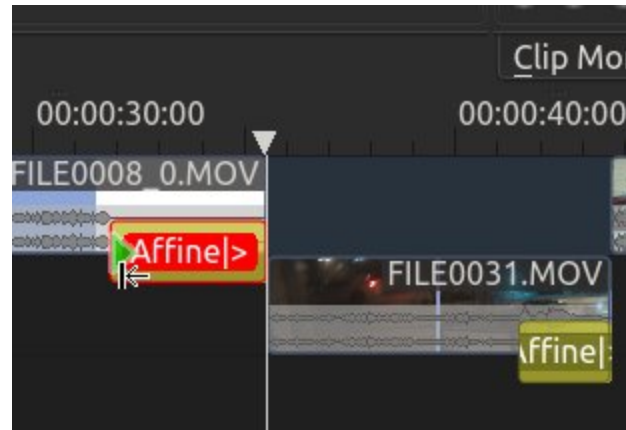
by loudog

Hello again, fellow PCLinuxOS users. In this installment, we will cover some interesting effects that are used often in the video editing community. I had initially planned on a different effect to cover, but I could not, for the life of me, find my camera tripod, so we will do these. You will recognize them easily, having seen them used in many movies and YOUTube videos. So, lets fire up that KDenLive and get started. The first thing we want to do is collect all the video clips together in one folder. Start a new project and import the clips into the project tree.

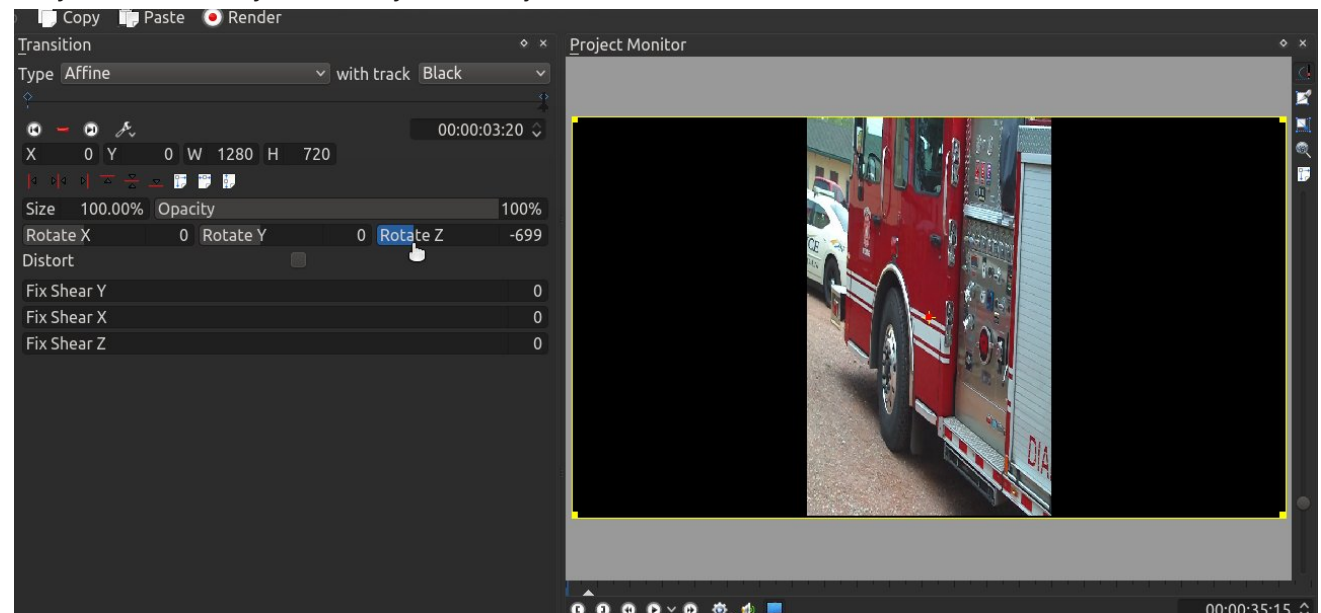


Since this will be the partial Affine transition tutorial, I have named the folder Affine, and this folder is placed in the KDenLive main directory. We will select our beginning video from the project tree, and drop it into the project timeline after our intro clip. Now add

the transition "affine" at the ending of the clip. Stretch the effect out a little.

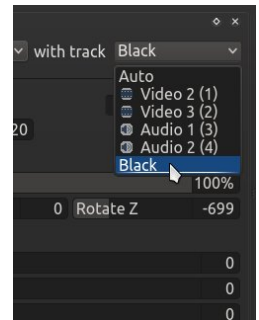


When the transition is selected, you can view the controls in the transition window. If you have been following along in the previous articles, this will be fairly familiar territory now. Many of the adjustments




are similar to the other effects, effect timeline, keyframe selection, center image, left/right, etc. The zoom works the same, but next we come to something new, called the axis control sliders. As you can see in the next image, I have set an ending keyframe and have started to adjust the Z axis. The Z axis turns the image/video from left to right or right to left depending on which way you slide the slider, so you can end up with a mirror image/video. Notice the background color, black.

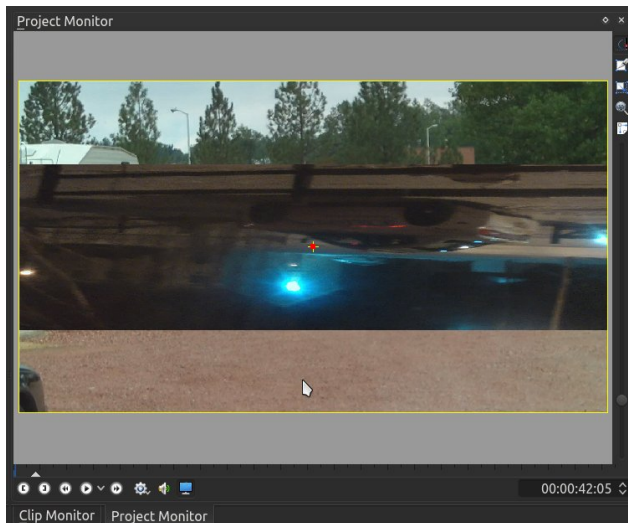
Now, at the top right of the transition window is a dropdown menu that is used for the background. We have chosen the black selection for the background as you can see in the project monitor with the firetruck partially turned.



We will review the transition in the monitor and make our adjustments until we are happy with the transition. If you want a longer/slower transition, just stretch out the effect a little more. In the next clip, I will do a standard fade from black effect. Let's take a closer look at this effect in the effect stack.

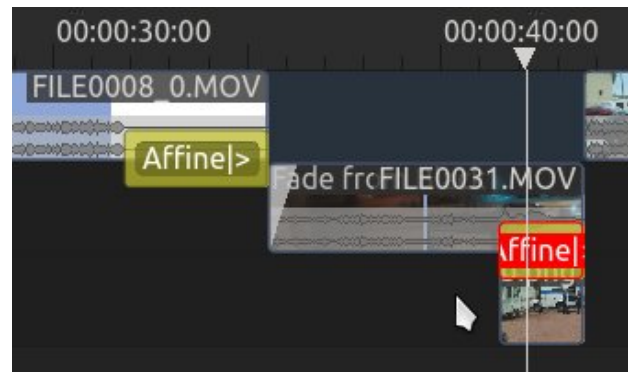


As you can see, here is where we set the duration time of the fade from/to black. I like my fades to be a little longer than default so I adjust the slider accordingly. At the end of this next track, I will use the affine transition again except a little differently this time. As you look at the image, notice that I have the video partially turned on the Y axis (the vertical control). This will flip the image/video upside down. Also notice that the background is not black this time, but an image.

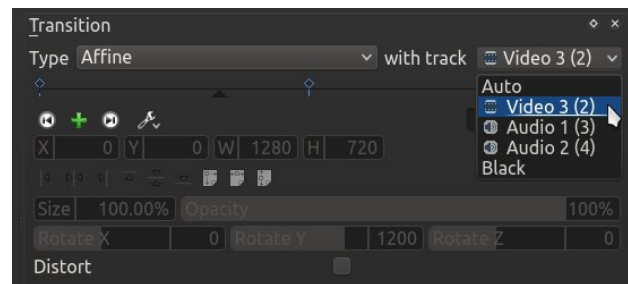


As the first video is rotating (police cruiser), the background image is a frame I selected from the next video, giving a little preview of what the next part of the video will contain. To do this, first extract

a frame from the video clip and appropriately name it. I named this one amb, short for ambulance. Import the image into the project tree and place it into the third video track in the project timeline. Align it and stretch it out so it matches the transition "affine" length. Examine this next screenshot closely. You will see all three video tracks have something in them.



Look at the background selection I have chosen. You have to select the proper configuration for this effect to become enabled. Video track "3" will be the background of video track "(2)" instead of a black background. You can manipulate this transition many, many ways.

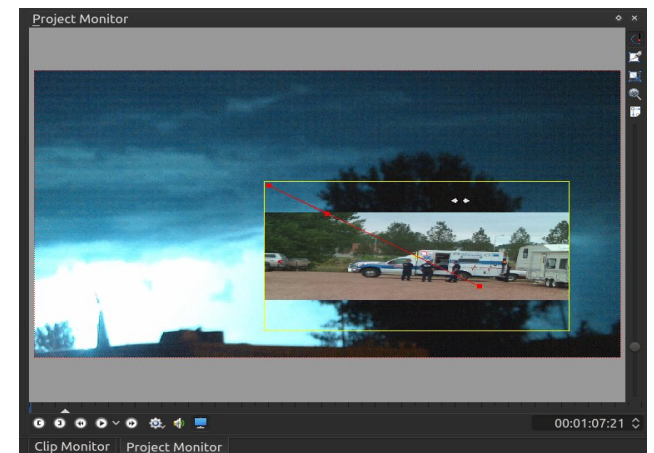


In the next clip I will use the affine in a different way. I will overlap the videos, add the affine transition on the first video, select video 2 (background) (1) (foreground) from the drop down menu and start adding my keyframes. Ready, set, here we go.



First keyframe: 3 second duration, zoom out to 80%, slide the X axis (spins the video in a circle) all the way over (left or right it doesn't matter). You will notice the image in the monitor will show with the familiar yellow outline for purposes of positioning it, because you have zoomed out. Decide where you want it and move it there. Many of the controls can be used much in the same manner as the pan and zoom effect we covered in earlier installments. I will leave it in the default corner for this keyframe.


Second keyframe: 5 second duration, zoom out to 40%, move the Y axis slider all the way over to the right (or left) and move the image to the bottom right corner of the monitor. As the first video (foreground) rotates to the lower left it also shrinks (zooming out) giving just enough room to see the lightning flash in the background of the video that is playing in the background.



Third keyframe: 5 second duration, zoom out to 5%, X axis slider all the way over and position the image at the top right. As the thunder storm rages, the ambulance flips this way and that and fades away. At the end of this next clip, I will use the default standard dissolve effect. One little tip on dissolve. If you look at the effect in the effect stack window you will notice a checkbox option to reverse the effect. Sometimes the effect looks better reversed than default, so review it both ways to see what will work best for your personal preference.

may view the finished video [here](#), but because of some coarse and strong language, it may not be suitable for younger ears. This is all we will cover in this installment of KDenLive, but will finish the affine transition tools next month, so don't you worry.

Getting a little more in depth with this transition in upcoming articles will be fun. Until then, stay under cover from the software slavery storm and remember to watch over Tux during the hailstorms. Until next month, enjoy spring!




For the next transition, to the massive hail storm video I took, I will again extract an image to preview in the background of what's coming up in the next segment, but instead of just rotating the foreground image, I will also zoom all the way out over 5 seconds so the preview is noticed quickly. The possibilities are fantastic with this transition, and I'll bet I haven't come close to all the combinations.



This is a video I am putting together for the family over on the west coast. This installment is on summer weather. I plan on making a few of these so the family can get a little taste of midwest living. You

Screenshot Showcase



Posted by Meemaw, March 16, 2014, running Xfce.

Back To School With Free Linux Intro Course

by Paul Arnote (parnote)

One of the fastest growing segments of the IT world is in Linux. IT managers around the globe are desperately seeking Linux professionals. According to a press release from The Linux Foundation, nine out of 10 surveyed hiring managers are seeking to hire Linux professionals within the next six months. Who knows – maybe some PCLinuxOS users could land gainful employment as a Linux professional, having cut their Linux teeth on our favorite operating system.



On March 6, 2014, [The Linux Foundation](#) announced, in collaboration with the free online education site [edX](#), the development of a series of Massive Open Online Courses (MOOC) that enable users to learn Linux. The first of these courses is

called [Introduction to Linux](#), and the \$2,400 (U.S. dollars) course is being offered free to anyone and everyone, anywhere and everywhere. Within the first 24 hours, more than 2,500 people had enrolled in the course, illustrating the desire for such courses.

 **LINUX FOUNDATION**

From The Linux Foundation [press release](#):

“Our mission is to advance Linux and that includes ensuring we have a talent pool of Linux professionals,” said Jim Zemlin, executive director at The Linux Foundation. “To widen that talent pool and give more people access to the opportunities in the Linux community and IT industry, we are making our training program more accessible to users worldwide. We are partnering with the leader in MOOCs to enable this access.”

“Linux is the world’s largest collaborative development project, and people everywhere want to understand how it works and how to tap into that massive community to advance their careers and general understanding of how technology today is built,” said Anant Agarwal, edX president. “Like edX, Linux is a non-profit, open source initiative, and we are pleased to partner with The Linux Foundation to make their basic learning materials accessible to more people around the globe.”

Additionally, The Linux Foundation is partnering with AT Computing, Enea, OlinData and SolutionWare to produce an Authorized Training Partner program. This new program, coupled with the new MOOCs, will increase access for Linux professionals to high quality, technically advanced, community based training materials. The courses from the Authorized

Training Partner program will be fee based, but should allow greater access to the training available from Linux Foundation Authorized Instructors.



edX was founded in 2012 as a non-profit learning platform by Harvard and MIT. Since then over 31 universities have partnered with edX, providing access for more than 2,000,000 people to its courses. All of the courses on the edX learning site are free of charge.



The PCLinuxOS Magazine

Created with Scribus



More Screenshot Showcase



Posted by daniel, March 8, 2014, running LXDE.



Posted by ferry_th, March 23, 2014, running KDE.



Posted by Texstar, March 7, 2014, running e18.



Posted by francesco bat, March 16, 2014, running KDE.