

Question and Answer

There were not many questions this month, so we discussed various things more generally.

Backups

The particular issue was how to ensure that old backups did not contain older copies of a particular file, in case it was restored and thought to be up to date. The file was updated by someone else and updates sent regularly, so only the latest version would be valid. One solution would be to place the file in an area that was not backed-up at all, and if there was a problem, then a request to the originator would restore the correct version.

If it was placed in a separate area, then it could also be found in the “usual”, or convenient, place by inserting a link to the correct file at the appropriate place in the file system. This is easy in a Unix system.

Ubuntu backup

Be aware that the Ubuntu backup system will restore only to the same machine as it was backed up from. This can cause difficulties if the reason for trying to do a restore is that the machine died and the files will have to be restored to a new one.

The new Mint backup system should solve those issues.

Timed searches in Google

Often when doing searches with Google, you may find that a lot of the results are to older documents that are now out of date. There is a way to restrict the dates to a certain time frame.

After the first search is done, Click on Search Tools just below the search query field. This brings up some other menus, one of which is Any Time. Clicking that gives a menu from which you can select a recent time, or define the timeframe you are interested in.

Photographs and Dates

When you take a photograph with your camera, it knows the date and time it was taken. This time is used when it creates a file in the camera. It also inserts both date and time into the metadata, which is the area of the file containing things like exposure, flash, and other information relating to the actual photo itself.

When you copy the file out of a camera under Linux, the programs will ensure that the time of the files where they are stored on your machine also are given the same time stamps.

However, when you copy those files, the system may change the timestamps because new files are created. It depends on how the copying is performed.

File browsers are generally good at doing the right thing here and set the times of any copies to be same as the files being copied. However, other programs may not be so good and may change the times to when you made the copy. So loading up a picture in a viewer to look at it, and then saving a new copy will probably reset the time to when this new version is made.

But the metadata will still contain the original dates. How you can view the metadata depends on the viewer in use, but something like Image → Properties, etc. will usually find it.

On Linux, there are command line programs to manipulate the metadata to insert extra comments, keywords, and copyright information, if that is of interest. Useful packages in this regard are *imagemagick*, *libimage-exiftool-perl*, *libjpeg-progs* and *netpbm*.

Clouds

We briefly talked about cloud storage. If you want to use one of these, then you need to think about where the data is held and who might have access to it, especially in light of recent news about the US government and the way they are involving commercial organisations. This leads to issues of the jurisdiction under which the data is held, and whether it conforms to your requirements.

Another question to ask is whether the data is encrypted and who holds the keys.

The way Firefox manages its Sync process seems to be the model we are looking for. They create a key on your machine, and do not upload it. This key is used to encrypt all the data in your sync'ed storage, and Mozilla cannot decrypt it. Of course, if you lose the key, you've lost the data.

But the data stored is only your bookmarks, add-ons, preferences, history and standard tabs for Firefox, so not usable for normal files, but the model of how it is done looks good preserving your privacy.

There is a suite of programs called git-annex that can synchronise with remote serves, but currently it is still too early to use by everybody – it is still work in progress, but worth keeping an eye on.

Everything comes down in the end to: Whom do you trust?