More on Backups and Scheduling thereof

by Andy Pepperdine

This paper is a supplement the previous one of 22 Sept 2011, but concentrating only on the simplest and usual case.

What to backup?

For the detail in this paper, the assumption is that only the home directory will be backed up. If you have any other data area that you wish to back up, then an appropriate change to the descriptions should be obvious.

On a standard Linux system all your private documents and other data will be accessible under the home directory, which is the directory /home/<*user*>. I will the the notation <*user*> to indicate the name of the user.

What format to backup to?

A lot of simple backup tools will compress the data as they back it up. In this age of cheap memory, there is little point in doing so, because the typical private user will have too little data to make it necessary. A significant advantage of taking a straight copy is that if you realise that you have deleted a file in error, the backup copy will have the same directory structure and can be accessed with all the normal tools, without having to decompress or in some way unravel.

Similarly, you may or not wish to encrypt it, depending on where you will store the backup. If you want to take copies onto a portable drive, then encryption may be worth considering if you have sensitive data that should no be lost.

How frequently to do them?

This depends entirely on what you want the backup for. If you think you might lose some file by deleting in error, or editing the wrong file, etc, then a frequent snapshot of your working area would be suitable. That way you can always arrange to recover the latst, or a defined version, from some date in the past.

However, if you want to have a complete accurate copy of what you now have (for example before a significant system upgrade) then taking a copy of your home directory onto an external drive would be appropriate.

Where to back to?

If you back up to the same local disk, the setup for taking backups is easy and recovery is always available. However, that means that if your disk crashes or becomes unusable, then you lose all the backups as well.

On the other hand, an external disk has always to be present, plugged in, and accessible, if you want regular snapshots of your working areas.

A combination of regular local snapshots that are for immediate recovery of deleted files, coupled with a manual backup once a week externally may suit most cases. The external one can be protected more thoroughly, and would be needed only when there is a serious crash. In that case, a complete backup of all the data would be required, not just a file or two.

How to make regular snapshots

A graphical interface to a thorough backup tool is backintime, which is available in the Ubuntu repositories. Mint users can choose the gnome version. Search for it in Synaptic.

First, make a suitable target directory where all the snapshots will be saved. For instance, you could place then in parallel with the home directories, by creating a directory with the terminal command:

```
    sudo mkdir /home/backup /home/backup/<user>
    sudo chown <user>:<user> /home/backup/<user>
```

The backup tool will create quite a complicated directory structure under the target directory, but the naming of the directories is obvious and simple.

When backintime is first invoked, it will show a simple dialog where you can specify where the snapshots are to be placed. There is one default "project" which identifies the type of snapshot to be taken. Under the General tab, leave the mode as Local.

Then browse to the directory you just created, for example, /home/backup/<*user*>, to set Wherer to save the snapshots.

Initially you can leave the other settings under the General tab as the defaults.

Next, under the Include tab, click the Add Folder and browse to your home directory.

The other tabs have sensible defaults and can be left alone. Later, when you are familiar with it, you can adjust some of those things as you wish.

Close the dialog and it will initialise its data areas.

Alternative arrangements

Other positions for the snapshot directory are possible, For instance, you could define the snapshot directory under your own home directory, and then ensure that the name of that directory is specified in the Exclude tab (or else, it might loop and keep snapshotting itself!).

Taking a snapshot manually

The leftmost icon will take a snapshot immediately, and is probably a good idea to do, to make sure it will work, and provides a baseline from which changes will be recorded when later ones are taken.

For the first one, this will probably take a significant amount of time. My first snapshot onto an external device of my photos took about 2 ½ hours. Later ones will be much faster.

Making them on external devices.

A manual snapshot can be taken on a different device if you first ensure the device is accessible, and then create a new project under backintime, which you can then invoke manually, or regularly if required. If it is not, then backintime will complain until it is plugged in.

NOTE: the filesystem you use to contain the backups must be able to support links, so a pre-formatted disk with VFAT will not work. They claim that NTFS will work, but I always reformat my disks, so cannot check.

Experiments with encrypted directories appear to work, so I suspect the documentation I've seen is out of date, as it would be difficult setting up hard links for such a case.

Scheduling regular snapshots

In the Settings dialog (the icon showing sliders) you can set up a regular regime for backups. This will be useful if you want to have snapshots always available to recover old files. But note that you would have to ensure any external device used to store them must be plugged in.

Alternatives for a manually controlled copy

When you want to make an external copy of all of your files, then an alternative manual method would be more useful – one which will not require the existence of the backintime program to be present when recovering the total data, such as after a major upgrade, or disk crash, for example.

The best command for that task is the command `rsync`. It will be present on all Linux systems. It will only copy those files that have been changed since the last time it was used under the same circumstances.

You will need the name of the external drive, as seen by the system. To find that, after you've plugged the drive in, open a terminal and execute the command

```
mount
```

Then look for the likely line in the output. It will start with directory name /dev/sdb1 (or something similar – it depends on your hardware and what else is plugged in), and after that you will see a name of the type /media/<*user*>/<*diskname*> where <*user*> is your user name, and <*diskname*> is the name of the partition as known by the system. That name is the one you need.

Next, create a suitable directory on the disk to contain all of your data (which here, I will assume is just your home directory), suppose it is given the name homecopy.

The following command will then make a copy of your home directory under that directory on the hard drive:

```
    rsync -av /home/<user>/   /media/<user>/<diskname>/homecopy
```

NOTE: the final / character on the first named directory is essential if you to copy the contents into the right place. Miss it out and extra level of directory will be added.

The first time will be long as it copies everything. Later runs wil copy only changed files.

References

Documentation on backintime can be found at: http://backintime.le-web.org/documentation/ although it is not very useful; I suspect it is out of date and not been updated for several years.