

# Encrypted Backups to the Cloud

## What I have

Here is my situation: I have a server at home where I store all my private data. I also have an extra harddrive in the server where I keep local backups of my storage with daily snapshots for about half a year. However, I currently have zero protection against data loss in case of a total loss of my server be it by theft, fire, water (the server is in the basement) or whatever other reason there might be.

## What I want

I want to synchronize the latest backup I keep with a remote server. There are a few challenges though: I want to only transmit changes, as all data currently accumulates to about 2TB, doing full backups is not practical. The remote backup should be encrypted in a way that the operator of the remote server has no access to my file (meaning, the files need to be encrypted before they are sent to the other server). In order for incremental backups to work, I need **file level encryption** which means, i can't encrypt an entire harddisk or a single image that then contains all my files, as this image would need to be copied in full upon each file change which is the same as a 2TB full backup in my case.

## Options I found

### pure Rsync

that's not an option, but i mention it here to just point out why: while rsync uses an encrypted channel for file transfer, it cannot encrypt the files themselves. so the transfer is secure, nobody between me and the remote server can see my files, but on the remote server they are stored in "plain text" other than that, rsync would be ideal.

### Duplicity

[Duplicity](#) is another popular candidate. actually i wanted to use that first, but realized that it creates archives of changes which are basically "added" to the initial full backup. this means that the backup will grow over time and never get smaller again as files can't be deleted. also restores will take longer and longer, the older the initial full backup gets. in order to circumvent this problem, duplicity needs to make a full backup every now and then which in my case is not an option.

### EncFS

[EncFS](#) would be the perfect solution for me. encfs is FUSE based and does encrypt every file automatically when stored to an EncFS mount and it then saves it to a directory somewhere on your

harddrive. even better, EncFS supports a `--reverse` option which does the opposite: you can mount an unencrypted directory to a mountpoint and in there you will see the encrypted version of that directory. once you have this, it is simply a matter of running a regular rsync on that encrypted reverse-mount and you're all set!

The Problem with EncFS is, that it is very old (started somewhere around 2003-ish). unfortunately it does not comply with modern best practices for cryptography and therefore did not get a good review when it was looked at by a payed security consultant in 2014. at the same time, the development of the project had prettymuch stalled. there are still some fixes and maintenance updates every now and then, but the main developer seems to have no time and motivation to keep this project alive and most of the security concerns from that review where not fixed until today. With all that, EncFS is not a viable option anymore as it is not particularlyly secure and its future is unknown, so starting a new setup for a backup I want to be able to access in the future as well is not such a good idea.

## CryFS

In 2015 a student of KIT (Karlsruhe Insitute of Technology in Germany) started a new Project called [CryFS](#) as his master thesis basically based on the story of EncFS. So CryFS does what EncFS does but it does it a lot more secure. CryFS even hides your directory structure and file size, so if you are afraid of someone profiling your file structure to then compare it to known profiles (i.e. song albums in mp3 format, windows installation cd's etc.) to prove you have stolen something somewhere or to know that you posses certain dataset and therefore are worth capturing this is a good option for you. there was also a master thesis reviewing the security of CryFS which seems to be pretty good. CryFS also supports very long filenames which is a plus over EncFS and it seems to be pretty alive for now.

There are two downsides I found while researching CryFS: there is no reverse-mode and it is pretty expensive as far as cpu load and required storage space goes. the whole hiding of file sizes and directory structures seems to demand its toll.

## rclone

[rclone](#) was invented basically to do exactly what I want. it can sync a local directory to a number of cloud services and it does that even without a separate client from those cloud services. It also has a crypt function that allows to encrypt your files before sending them to the cloud and it supports prettymuch all major cloud storage providers. In my case the cloud storage i will be using is a VPS linux server, so the recently added support for FTP and the sftp support that preexisted are what would work perfectly fine. There are no reviews regarding its crypto mechanism but they seem to use proven systems that seem to be safe and uncracked so far. however, i'm by no means an expert and can only say what i have read somewhere on the internet.. so don't regard this as a qualified oppinion and research it yourself.

the problem with rclone is: it is very much focused on basically replacing the standard sync clients from dropbox and the likes and not so much on providing a complete backup for a linux filesystem which is what i am looking for. it has unfortunately no support what so ever for any linux attribuets, not even permissions and owner information. it would save a copy of all my data and in case my house burns down i would certainly be happy with my backup even without permissions and owner information as i would clearly have other things to worry about and with only about two users on the server this would really not be such a big issue. However, I skipped this option as finding something

that would also work for corporate servers in the future was more important to me than Dropbox etc. interoperability.

there is a possible [workaround to store file permissions](#) which could be scripted, but since I want to backup both my data and also the server installation, i would really like to have some more posix features kept in the backup than simply permissions.

## gocryptfs

[gocryptfs](#) is another answer to the nearing death of EncFS. The Austrian Guy who started and maintains the project seems to have also used EncFS before and if I understood correctly he has even discussed on how to continue EncFS and how to go about building a version 2.0 of it to solve those security issues. His suggestion was to start from scratch and build a replacement and that's just what the aim of gocryptfs is. It is written in Go, hence the name.

gocryptfs seems to be under active development but it is already considered stable by the developer. the webpage features a nice [comparison with other solutions](#) which also compares it to CryFS. both CryFS and gocryptfs are about the same age. Gocryptfs has recently been [Audited](#) and the developers seem to take the audit serious and seem to work on fixing the more important issues. there is a [issue on github](#) that tracks the progress of solving the findings. From what i understand from the audit, gocryptfs is safe to use as long as the attacker has no read-access to the mounted (unencrypted) filesystem and as long as one is not concerned about receiving tampered data back from the backup. In my case that is sufficient as I only want to protect my backup against being read by someone else. Hopefully i will never have to read my backup again anyway, so the tampering part is really not that big of a concern and it would be very uninteresting for an attacker to get access to my server this way.. there are easier more efficient ways I suppose.. and still, the information on my server would probably not be worth all that effort as my private holiday photos aren't really that amazing anyway ;)

## The solution I am aiming for

I decided to go with gocryptfs in reverse mode and rsync to backup my files to a remote server. As i can put a physical harddrive into the remote backup server to which the VPS will then have exclusive access, i will first run the backup locally from disk to disk and then introduce ssh as a transport tunnel to do remote incrementals in the future.

## the setup

### Prepare the Target

first i need to mount the harddisk i am going to install to the remote server locally on my home-server.. i have a bunch of 2TB disks laying around, so i will simply create a RAID 5 across 3 of them to get 4TB of usable storage. i will later add another 2TB disk to convert this into a RAID 6 but at home i only have 3 slots available in my server, so 3 disks it is ;)

```
mdadm --create --verbose /dev/md3 --level=5 --raid-devices=3 /dev/sd{h..j}
```

now lvm on top of that so i could expand my volume with additional disks/raidsets etc. in the future without redoing my initial 2TB sync: `pvccreate /dev/md3 vgcreate psuter /dev/md3 vgdisplay lvcreate -n psoffsite -l 953800 psuter` the number behind -l is simply the free PE number from `vgdisplay` in order to create a logical volume covering the full size available.

now let's put a filesystem on top of that.. i'll use xfs:

```
mkfs.xfs /dev/psuter/psoffsite
```

now mount it temporarily to do the initial backup.

```
mkdir /mnt/offsiteBackup  
mount /dev/psuter/psoffsite /mnt/offsiteBackup/
```

## gocryptfs installation

download a static binary from the [releases](#) page

From:  
<http://wiki.psuter.ch/> - **pswiki**

Permanent link:  
[http://wiki.psuter.ch/doku.php?id=encrypted\\_backups\\_to\\_the\\_cloud&rev=1502620854](http://wiki.psuter.ch/doku.php?id=encrypted_backups_to_the_cloud&rev=1502620854)

Last update: **13.08.2017 12:40**

