

Parallel Rsync (how I believe it's done)

rsync is sooo cool, chances are, if you need to copy some files for whatever reason from one linux machine to another or even from one directory to another, rsync has everything you need. one thing though is terribly missing: parallelism

here is, how i did it when i needed to copy 40 TB of data from one raidset to another while the server was still online serving files to everybody in the company:

Step 1: creat an incremental file list

depending on your needs, there are different options how to do that.

Option 1: rsync --dry-run

one possibly slow option is to do a dry-run of rsync with all your options you want to use and then use the file-list created by the dry-run for your rsync job.

first do the dry run:

```
rsync -aHvx --dry-run --out-format="%n" /source/ /target/ | tee /tmp/rawfilelist
```

use rsync options like you would for a simple rsync run to copy all your files, but add the `--dry-run` `--out-format="%n"` options. the `out-format` option is to make sure you get a simple list of files without the added information about symlinks and hardlinks, that you would get when this option was omitted.

now clean up the resulting file: the problem with the dry-run output is, that you also get directory names before you get the list of the contents of each directory. that's useless if we want to continue later on and run an rsync for each file. so we need to get rid of these directory paths. this will obviously lead to empty directories not being copied, we can fix that later on by running a simple single thread rsync at the end to fix things like exactly that :) so here we go.. let's clean up the filelist (you can do this inplace, but you might just want to use this line and pipe it directly into parallel further down the road)

```
cat /tmp/rawfilelist | sed -e 's/.*\$/\$/' | sed -e 's/sent .* bytes\$/sec$/' | sed -e 's/^total .* (DRY RUN)\$/\$/' | sed -e 's/sending incremental file list\$/\$/' | sed -e '/^$/d' > /tmp/filelist
```

Option 2: using find

after waiting too long for Option 1 to finish on a system that carried tons of backups of other systems, i tried this option:

if you have tons of files and want to skip the lengthy process of producing a file list via rsync, you can

create a list of directories using find and then simply run an rsync per directory. this will give you the full parallelism at the begining but might end with a few ever lasting rsyncs if you don't dig deep enough when doing your initial directory list. still, this might save alot of time.

```
find /source/. -maxdepth 7 -type d | perl -pe 's|^.*?/\./|\1|' > /tmp/filelist
```

with the `-maxdepth` option you can set how deep you want to dive into your directory tree.. the goal is to get directories with a rather small number of files so you don't have to wait too long for the last couple of rsyncs to finish. also note the added `./` at the end of the source path. that's important as we need this to define to which point rsync should be relative. also check out the man page of rsync, i stole the idea from there ;) there is no cleaning needed here, as we really want the directory names to sync directories rather than files.

Step 2: run Rsync with GNU prallel

now it's time to feed our filelist into rsync and run our parallel sync job. in order to parallelize rsync we use the GNU tool `parallel`. it will take a list of files and run a command in parallel with as many processes as are specified by the `-j` option. in the command string, it will replace `{}` with the contents of the respective line. pretty simple :)

```
cat /tmp/filelist | parallel -j 10 rsync -aHvx --relative "/source/.{ }" /target/
```

note how, like in the above mentioned Option 2, we use the `./` separator in the source path to tell rsync where to start with the relative path that it transmits to the client. also make sure you actually use the `--relative` option, otherwise your targets file structure will be very flat :) Also using quotes might be a good idea as it will help you to deal with spaces and other funny characters in your filenames.

Step 3: make sure we didn't miss anything

probably the best feature about rsync is, that it resumes aborted previous jobs nicely and it can be run several times across the same source and target with no harm. so let's use this property to just fix everything we have missed or done wrong by simply running a single thread rsync in the end. now this can take some time, and I know no way around that.

```
rsync -aHvx /source/ /target/
```

From:

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

Permanent link:

http://wiki.psuter.ch/doku.php?id=parallel_rsync&rev=1470678905

Last update: **08.08.2016 19:55**



