

Parallel Rsync (my way)

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

when you copy files with rsync you often see an io performance (by using iotop for example) that is far below what your disks or your network connection are capable of.

1. when you copy a directory containing many small files locally, rsync is slowed down by all the metadata operations it does (copying all the permissions, checking each file for changes by checking file dates etc.)
2. when you copy files across a network, you are slowed down by a single threaded ssh process which can only use one cpu core for encrypting and decrypting data on that connection.

my solution to that: run multiple rsync processes in parallel and leverage the power of several cpu cores in parallel.

here is a bash script function I wrote which i can use in scripts to copy files from A to B through multiple connections.

Use this at your own risk If you are interested in understanding what it does (and i strongly suggest you get interested in that before using this blindly!) you can read through my ideas on how to ideally parallelize rsync below the script.

the bash function

what this function does is as follows: it runs a find across the source directory and gets a list of all files and directories within it. it then extracts a list of directories with a maximum directory depth of \$3 (the 3rd argument to the function). it then queues these directories to sync them and runs \$4 (the fourth argument) rsync processes in parallel to do so using xargs.

once all directories have been synced it runs a single rsync thread as you would to just simply copy the files single threaded. only now we already copied all the files. this step is more some sort of a safety measure to make sure we really copied everything and that all file attributes are correct.

once this step is passed it runs another sanity check and compares md5 sums between all source and target files. this might take very long and is not really necessary but since i programmed this function for an archive script that will copy files to an archive before they are deleted from the source i wanted to be 100% sure everything went okay :)

how many jobs should run in parallel and how many directories deep you want to parallelize your jobs really depends on your sepcific situation. if you have several terabytes of data and you do a complete sync it makes sense to dive deeper into the structure than when you just want to update an already existing copy of the same data, in that case it might be faster to only dive 1 to 2 levels deep into your structure or even not use this script at all, when most of the time is spend by "creating incremental file list". really, read what's behind the script further down to understand how to parametrize it and how to modify it to adjust it to your specific situation

known issues

- when i wrote the script i used excessive escaping to make sure that more complex file names could also be copied. however i do not escape the source directory name or the target directory name. i sadly realized only after writing all this escaping madness, that i could have used the -0 option of xargs and corresponding null options to the other tools so that i could have avoided the need for all this escaping madness in the first place

the code

[prsync.sh](#)

```
#
# Parallel Rsync function 2017 by Pascal Suter @ DALCO AG, Switzerland
#

psync() {
    # $1 = source
    # $2 = destination
    # $3 = dirdepth
    # $4 = numjobs

    source=$1
    destination=$2
    depth=$3
    threads=$4

    # gets directory listing form remote or local using ssh and file
    dirlist(){
        # $1 = path, $2 = maxdepth
        path=$1
        echo "$path" | grep -P "^[^@]*@[^:]*:" > /dev/null
        if [ $? -eq 0 ]; then
            remote=`echo "$path" | awk -F : '{print $1}`
            remotepath=${path:${#remote}+1)}
            ssh $remote "find $remotepath/./ -maxdepth $2 -type d |
perl -pe 's|^.*?/\./|\1|'"
        else
            find $1/./ -maxdepth $2 -type d | perl -pe 's|^.*?/\./|\1|'
        fi
    }

    # get a sorted list of md5sums of all files in a directory (remote
    via ssh or local)
    md5list(){
        # $1 = path
        path=$1
        echo "$path" | grep -P "^[^@]*@[^:]*:" > /dev/null
        if [ $? -eq 0 ]; then
```

```

        remote=`echo "$path" | awk -F : '{print $1}`
        remotepath=${path:${#remote}+1}}
        ssh $remote "cd $remotepath; find -type f -print0 | xargs
-0 -P $threads -n 1 md5sum | sort -k 2"
    else
        cd $path; find -type f -print0 | xargs -0 -P $threads -n 1
md5sum | sort -k 2
    fi
}
# escape wrapper function. will do a double escape if the source is
remote, will do a single escape if source is local
source_escape() {
    echo "$source" | grep -P "[^@]*@[^:]*:" > /dev/null
    if [ $? -eq 0 ];then
        escape | escape
    else
        escape
    fi
}

#magic escape function. it is probably not yet complete but it can
be expanded based on the last "final sync to double check"
#file names that where not or wrongly escaped end up there.
escape() {
    sed -e 's/\\/\\\\/g' -e 's/ /\\ /g' -e 's/\\$/\\\\$/g' -e
's/://\\:/g' -e 's/(\\(/g' -e 's/)/\\)/g' -e 's/"/\\"/g' -e
"s/'/\\\\'/g" -e 's/|/\\|/g'
}

# generate a list of directories to sync
rawfilelist=`dirlist $source $depth`

# separate paths less than DIRDEPTH deep from the others, so that
only the "leaves" get rsynced recursively, the rest is synced without
recursion
i=$(( $depth - 1))
parentlist=`echo "$rawfilelist" | sed -e '/^(.*\\/)\\{'$i'\\}.*/d`
filelist=`echo "$rawfilelist" | sed -e '/^(.*\\/)\\{'$i'\\}.*/!d`

# create target directory:
path=$destination
echo "$path" | grep -P "[^@]*@[^:]*:" > /dev/null
if [ $? -eq 0 ]; then
    remote=`echo "$path" | awk -F : '{print $1}`
    remotepath=${path:${#remote}+1}}
    remotepath=`echo "$remotepath" | escape | escape`
    ssh $remote "mkdir -p $remotepath"
else
    path=`echo "$path" | escape`
    mkdir -p $path

```

```

    fi

    #sync parents first
    echo

    "=====
    ====="
    echo "Sync parents"
    echo

    "=====
    ====="
    echo "$parentlist" | source_escape | xargs -P $threads -I PPP rsync
-aHvx --numeric-ids --relative -f '- PPP/*/' $source/..'PPP'/
$destination/ 2>/tmp/debug
    status=$?
    if [ $status -gt 0 ]; then
        cat /tmp/debug
        rm /tmp/debug
        echo "ERROR ($status): the was an error when syncing the parent
directories, check messages and try again"
        return 1
    fi

    #sync leafs recursively
    echo

    "=====
    ====="
    echo "Sync leafs recursively"
    echo

    "=====
    ====="
    echo "$filelist" | source_escape | xargs -P $threads -I PPP rsync -
aHvx --relative --numeric-ids $source/..'PPP' $destination/
2>/tmp/debug
    if [ $? -gt 0 ]; then
        cat /tmp/debug
        rm /tmp/debug
        echo "ERROR: there was an error when syncing the leaf
directories recursively, check messages and try again"
        return 1
    fi

    #run a single thread rsync across the entire project directory
    #to make sure nothing is left behind.
    echo

    "=====
    ====="
    echo "final sync to double check"
    echo

    "=====
    ====="
    rsync -aHvx --delete --numeric-ids $source/ $destination/

```

```

    if [ $? -gt 0 ]; then
        echo "ERROR: there was a problem during the final rsync, check
message and try again"
        return 1
    fi

    #create an md5 sum of the md5sums of all files of the entire
project directory to compare it to the archive copy
    echo
    "=====
====="
    echo "sanity check"
    echo
    "=====
====="
    diff <( md5list $source ) <( md5list $destination )
    if [ $? -gt 0 ]; then
        echo "ERROR: the copy seems to be different from the source.
check the list of files with different md5sums above. Maybe the files
where modified during the copy process?"
        return 1
    fi

    echo "SUCCESS: the entire directory $project has successfully been
copied."
}

```

Usage you can run this function like so:

```
psync sourceHost:/source/directory target/destination 5 8
```

this will copy the /source/directory to /target/destination and it will dive 5 directory levels deep to parallelize rsyncs. it will run 8 rsync processes in parallel.

caution this is a work in progress.. I am writing down my notes as I go!

caution please be careful with the instructions below and think it through yourself. I will take no responsibility for any data loss as a result of this article.

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:

Before we get startet

one important note right at the beginning: while parallelizing is certainly nice we have to consider, that spinning harddisks don't like concurrent file access. so be prepared to never ever see your harddisks theoretical throughput reached if you copy lots of small files. make sure you don't run too many parallel rsyncs by checking your cpu load with top. if you see the "wa" (waiting) load increase, it

means you have too many processes. On the system I did this all for, first tried with 80 parallel rsyncs using option 2 below and I had a waiting load of about 50% and a throughput of about 20MB/s. I then reduced to 15 parallel rsyncs and the waiting load went down to 25% and the bandwidth went up to over 100MB/s. That is on a RAID set that achieves a raw throughput of over 500MB/s if streaming performance is measured. Just to give you an idea. Besides that you can also use `iostat` to monitor your overall rsync speed.

Step 1: create 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 in place, 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 beginning 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 a lot of time.

```
find /source/. -maxdepth 5 -type d | perl -pe 's|^.*?/\./|\1|' > /tmp/rawfilelist
```

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 ;)

now it's time to clean up the list. we need to move all lines that contain less than the `-maxdepth` number of directories to a separate file list as these directories will need to be synced without recursion. i tried doing this with a loop that went through all lines trying to find the respective lines, but it took way too long for a rawfilelist with more than 300'000 entries, so i tried it with sed inplace and it was incredibly fast!

```
cp /tmp/rawfilelist /tmp/parentslist
cp /tmp/rawfilelist /tmp/filelist
sed -i '/^\(.*\)\{4\}.*$/d' /tmp/parentlist
sed -i '/^\(.*\)\{4\}.*$/!d' /tmp/filelist
```

make sure that the number in the sed regex is your `–maxdepth` number minus 1! now we need to sync the parents without recursion first before continuing to step 2

```
cat /tmp/parentlist | parallel -j 3 'shopt -s dotglob; rsync -aHvx --no-r --relative /tmp/source/{}/ * /tmp/target/'
```

the trick here is to use the `–no-r` option to remove recursion of whatever rsync parameters you have specified before it. also check out the `shopt` command which results in a `*` matching also hidden files like `.htaccess` and so on.

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 :)

note that `parallel` is thorough as far as escaping goes. there are no quotes needed even with funny directory names.

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 --delete /source/ /target/
```

From:

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

Permanent link:

http://wiki.psuter.ch/doku.php?id=parallel_rsync&rev=1524587612

Last update: **24.04.2018 18:33**

