

# Transferring data from ATNF telescopes

## LBA Recorders

These instructions are for data transfer from Parkes (pam-store-ext, pkvsi1-ext, pkvsi2-ext), ATCA (cavsi1-ext, cavsi2-ext, caxcube-ext), Mopra (caxcube-ext, mpvsi1-ext, mpvsi2-ext) and ASKAP (cira10, akxcube-ext).

Requirements:

- ssh access to venice.atnf.csiro.au (ATNF Unix account)
- vlbi password
- access to lba-1.it.csiro.au or vlbi-data

ssh authentication key setup:

In order to use globus GridFTP to push data to lba-1 (or vlbi-data or Pawsey) via your own user account, it helps to have ssh keys set up. On each recorder host that you need to transfer data from (iff each has its own independent home directory), use `ssh-keygen` to generate your own ssh keys with a passphrase that you can remember. `ssh-keygen` will prompt for a filename to save to; use a new unique file in the `~/.ssh` area, e.g. "myname\_rsa". Append the public key to your `authorized_keys` on the destination machine with e.g. `ssh-copy-id -i ~/.ssh/myname_rsa.pub user@cass-01-per.it.csiro.au`. After this you should be able to ssh from the recorder host to `cass-01-per` without needing to enter your password.

To use your keys on the recorder host type

```
ssh-agent bash
ssh-add ~/.ssh/myname_rsa
```

`ssh-add` will ask for your passphrase. After this you should be able to log into `cass-01-per` without being prompted for your passphrase. Doing this within screen will save you having to repeat this every time you log in (man `screen` for info on screen).

Note that the alias `gcopy` described below is defined on execution of the `.bashrc`, so you may want to use a bash shell (rather than the default `csh`) even if you don't use `ssh-agent`.

To save typing your password every time you log into a recorder host from venice, you may also want to append your own ssh public key from your ATNF user account to the vlbi account's `ssh authorized_keys` on the recorder hosts.

After each run (including mini-runs):

- Using the [Current Block Schedule](#) on the wiki or [web page](#), check experiment names and telescopes which took part.
- Update the [google docs spreadsheet](#) population from the `recmon` database. Hover the mouse over the A1 cell to view instructions on how to do this. Note that the `updateDisks` script is strict about the formatting of the date you enter. If you stuff it up, however, you can go back to the previous revision of the spreadsheet.
- Using the above google docs spreadsheet (or [recmon disk label view](#)) work out which

experiments were recorded where - note some experiments may be recorded on multiple locations.

- For each station, and for each disk array, transfer data to lba-1 (see below). Only transfer from one of At, Pa, Mp at a time (and from one machine).
- ASKAP can be run in parallel with an east coast telescope.

#### How to transfer data:

- ssh to venice (using your own ident) then to the data host (one of the above machines) as vlbi to push the data over. It's best to run transfers within screen - see above.
- See the notes on how to use [globus-url-copy](#) to transfer directories to lba-1.it.csiro.au (or hpc-data.pawsey.org.au) - can create an alias to save time entering all the options (see gcopy on most VLBI machines). You can also use parallelisation, e.g. `-p 2`, or `udt -udt` (not both!) to speed up transfers. The environment variable `$GLOBUS_LOCATION` should be set. (Add to `.mybashrc` and source this from `.bashrc`, if not already there.)
- Data at the Pawsey centre takes the structure:  
`/scratch/pawsey0169/transfers/{year}/{experiment}/{antenna}`
- Data at lba-1 take the structure:  
`/datasets/data_cass/data/{year}/{experiment}/{antenna}/`
- Data at vlbi-data take the structure: `/data/lba/{year}/{experiment}/{antenna}`
- Note that for dual-DAS recording, transfer DAS1 and DAS2 data into `{experiment}/{antenna}/{DASn}` directories.
- You can run `globus-url-copy` a second time (be sure to include the `-sync` option) to check everything transferred OK.
- Where multiple experiments are to be transferred from one machine (perhaps from multiple raids) you can avoid having to start multiple transfers by using soft links to have all the data appear in one source directory:

```
cd ~/transfers
ln -s /path/to/data
ln -s /path/to/other/data
gcopy file://$PWD/ sshftp://user@remote//remotedir/
```

where `gcopy` is a bash function: `globus-url-copy -cd -r -sync -sync-level 1 -restart -vb -c $@`

- Once the whole lot is transferred, the data can be moved from `{antenna}/{experiment}` to the correct `{experiment}/{antenna}` directories. **Note** the following is specific to Pawsey, but could be also made available on cass-01-per: To help with this rearrange, Chris has a script `~cphillips/bin/movedir.sh` which works with alias 'movedir' (to be run from within the experiment directory where the data are currently located). To use it, you can add `alias movedir='~cphillips/bin/run_movedir.sh'` to your own `~/.profile` or similar, and add Chris' `~/bin` directory to your path (add the line `export PATH=/home/cphillips/bin:$PATH` to your `.bashrc`).
- Use the script `filesun.sh` (`~cphillips/bin/filesun.sh` at Pawsey, `~big040/bin/filesun.sh` on cass-01-per) on both the copy and the original data at the observatories to compare the two sets of data (check that total data volume, #files, smallest and largest files and first and last are the same). Once this comparison is made, use alias "movedir" on the original data (at the observatories; the alias is already set in the vlbi account) to move it to a directory called "checked".
- Update the google docs spreadsheet to indicate data transfer has occurred.

## Medusa Pkstore (FlexBuff) recorder

Medusa records data to a FlexBuff recorder (pkstore).

```
ssh vlbi@pkstore.atnf.csiro.au
```

(via venice if necessary, address may also be pkstore-ext depending where you connect from).

Data are stored on a VBS (virtual) file system which you need to access with the vbs command set and/or jive5ab:

[jive5ab](#)

[vbs](#)

To get a file listing:

```
vbs_ls -lrth vt031a*
```

To transfer, make sure jive5ab is running at both source and destination, then use m5copy:

```
m5copy --resume -p 40040 vbs://localhost/vt031a*  
file://lba-1.it.csiro.au:40100/datasets/data_cass/data/2021/vt031a/md/
```

There is also an alias jcopy on pkstore which sets `--resume` and runs under a re-try loop in case the connection falls over. This is very convenient for large unattended transfers.

## Mopra Mark6

Used mostly for the OCTAD backend.

```
ssh oper@kasi-mark6.atnf.csiro.au
```

data are typically found under `/mnt/mark6g*` and appear as a more-or-less normal file system (YMMV). Use `du --apparent -size` to check file sizes.

Use jive5ab to transfer data similar to procedure for Medusa (pkstore) above.

## RadioAstron data transfers

All data recorded by ATNF telescopes for RadioAstron experiments needs to go to cass-01-per, usually in the directory `/OSM/PER/CASS_LBA/data/{year}/Radioastron/{experiment}`. On pawsey, the script for moving RadioAstron data to the correct subdirectory structure can be accessed using `alias movera=' . ~/cphillips/bin/run_movera.sh'` - this could be copied to cass-01-per.

The data are transferred to ASC for correlation, and need to be converted to mark5b then copied using the tsunami transfer protocol. ASC requests a copy of all Radioastron data, however some experiments are correlated at JIVE or Bonn and can also be transferred there directly (see below).

You will need to use the following programs/scripts. Make an alias or add e.g. `/home/phi196/bin` to your path.

```
/home/phi196/bin/lba2mk5b  
/home/phi196/bin/tsunamid  
/home/big040/bin/domd5.pl
```

The following assumes these are in your path.

To convert LBADR data to mark5b change directory to where the data are and run `lba2mk5b`, e.g.

```
> cd transfers/2016/radioastron/rk12o1/ATCA  
> lba2mk5b *.lba
```

If you wish to keep all the mark5b data in a subdirectory add the option `-o mark5b`. You must create the directory by hand first.

ACS insist on md5sums calculated on all files. `domd5.pl` will do this for you. It will do multiple directories recursively, only runs `md5sum` on files with `.m5b` extension and checks if the md5 sum has already been calculated for a particular file. This means you can redun `domd5.pl` if more data is added or the script is interrupted. Note however if a file is updated, you will have to edit the `dir.md5` file by hand to remove the file entry (or remove the file to redo the whole directory).

To md5sum a number of directories at the same time run:

```
> domd5.pl rg14e rgs14d rgs14e
```

If you do not pass any directories on the command line it does not do anything.

Once the mark5b files and created and `md5sum` run, you need transfer the data using Tsunami. Multiple directories can be transferred at the same time. Run something like

```
> tsunamid --port=40070 rg14e/*/*.m* rg12o1/*/*.m*
```

The exact command will depend on the directory structure used and number of experiments/stations transferred at once. Note the above should make the `.m5b` data files and `.md5` md5sum files available.

Finally send an email to ACS staff (currently Dmitry) to let them know which experiments/stations you have made available and note the Tsunami port used (40070 above, you can choose anything between 40000 and 40100) and which machine (e.g. `cass-01-per.it.csiro.au`). Request Dmitry lets you know when transfer is completed. It can 24hrs or more for the transfer to start, so run `tsunami` from a VNC or screen session.

## Other transfers

The below assumes firewall settings allow traffic on the relevant IP addresses/ports. Can test with `traceroute` (uses UDP by default), e.g. `traceroute -p 40070 {IP address} (-T option for TCP)`

## To Bonn:

Use either `jive5ab` (see below) or `tsunami`. (`jive5ab` needs to be running at both ends.) Coordinate with Gabriele Bruni re. Radioastron data transfers to Bonn.

To run the `tsunami` client, log in to the remote machine with `ssh evlbi@{IP address of machine to send data to}` (NB requires that `ssh` keys are set up). Then, e.g. if 350Mbps has been reserved for the transfer: (`dir` lists the available files, just to make sure the server is working)

```
> cd /data{directory where data are to go}
> tsunami set port 40070 set udpport 40070 rate 350m connect cass-01-per.it.csiro.au
> dir
> get *
```

## To JIVE (or anywhere else, with `jive5ab`):

Use `jive5ab`. This should always be running already on the JIVE end. Coordinate with Martin/Bob/e-Bob. If not running already on `cass-01-per` (check with `lsof -i`), start in `screen` with `jive5ab` (the port number here is irrelevant as we only ever read from the “local” disk). Then run `m5copy` with something like:

```
m5copy -udt -p 40001 file://$PWD/*.m5b vbs://{IP address of flexbuff at JIVE}:40000/
```

See `m5copy -h` for more options/usage. For file to file transfers, use something like:

```
m5copy -udt -p 40001 --resume
file:///OSM/PER/CASS_LBA/data/2018/{path/to/data}/* file://{remote IP}/path/to/data_directory/
```

NB wildcards are supported for SRC file URIs only if `m5copy` is running on the local machine (remote wildcards are not supported).

## Tuning Kernel for High Speed Transfers

We normally use the following settings in `/etc/sysctl.conf` to get better performance in both TCP and UDT transfers.

```
# increase TCP max buffer size
net.core.rmem_max = 201326592
net.core.wmem_max = 33554432
# increase Linux autotuning TCP buffer limits
# min, default, and max number of bytes to use
net.ipv4.tcp_rmem = 4096 87380 33554432
net.ipv4.tcp_wmem = 4096 87380 33554432
#net.ipv4.udp_rmem_min = 262144 #Traditional value (see
M recommendation below)
```

```
#net.ipv4.udp_wmem_min = 262144 #Traditional value (see  
M recommendation below)  
net.ipv4.udp_mem = "536870912 805306368 1073741824" #Marjolein  
recommendation for Monsterbuff @16Gbps  
net.core.netdev_max_backlog = 8192
```

To check current settings:

```
sysctl -a
```

To set parameters on the fly:

```
sysctl -w variable=value
```

To load a new sysctl.conf after editing:

```
sysctl -p /etc/sysctl.conf
```

From:

<https://www.atnf.csiro.au/vlbi/dokuwiki/> - **ATNF VLBI Wiki**

Permanent link:

<https://www.atnf.csiro.au/vlbi/dokuwiki/doku.php/lbaops/datatransfers?rev=1653460885>

Last update: **2022/05/25 16:41**

