

# Migrating Files from HPSS

*Brian Vanderwende*  
*CISL Consulting Services*

December 4, 2019



# The time to migrate important data from HPSS is NOW!

- NCAR's High-Performance Storage System (tape) will go read-only on January 20, 2020 and will reach end-of-life on October 1, 2021.
- This deadline **will not be extended!** The vendor (Oracle) is no longer manufacturing tape drives and spare parts are already in short supply.
- Users have **87 PB** of data stored in HPSS. Except for certain collections, **CISL will not be migrating these data.** The responsibility to preserve needed data falls on you.
- Each HPSS tape drive has limited bandwidth, so it can take quite a bit of time to migrate large data holdings. There is only enough time to migrate about **50%** of existing data before EOL.

# Form a migration plan with your PI

- Work with the project PI to determine which data to save and where it will go. The PI may already have a plan for the files!
- Coordination can prevent unnecessary and/or duplicate transfers.
- Think of the HPSS data drives as highways - the more cars we can keep off of the road, the more likely everyone gets home.
- If the planned destination is Campaign Storage, the PI may need to request an allocation (users can check for CS allocations via **gladequota**).

Visit <https://www2.cisl.ucar.edu/user-support/allocations>  
for allocation requests

# Understand your HPSS data footprint

- HPSS files and directories are associated with projects. While logged into Cheyenne, run the **id** command to list your projects (groups):

```
cheyenne$ id  
uid=8061(nad) gid=1234(ncar) groups=73704(ucbk0099) 7087(cesm0099)
```

- Project files may be in a dedicated directory on HPSS (e.g., /CESM) or in your personal directory. Use the **hsi** interface to list files in these locations using **ls** (inventories can be long, so best to redirect output to files).

```
cheyenne$ hsi ls -lRU /home/username >& ~/hpss-home.txt  
cheyenne$ hsi ls -lRU /USERNAME >& ~/hpss-user.txt
```

# CISL can also provide you with inventories of files you own

For files in all locations on HPSS, you can see the file size, project association, and full path. An example:

```
cheyenne$ cat vanderwb.dat
...
80767844,SCSG0001,/home/vanderwb/wrfrst_d01_2013-09-10_01:00:00
80767844,SCSG0001,/home/vanderwb/wrfrst_d01_2013-09-10_02:00:00
80767844,SCSG0001,/home/vanderwb/wrfrst_d01_2013-09-10_03:00:00
...
```

Submit a Research Computing help ticket via <http://support.ucar.edu> to request an inventory.

# Organize files on tape first

- Moving files to different directories on tape is a metadata operation, so it basically takes no time - **use this fact to organize your data.**

```
cheyenne$ hsi mkdir /PROJ0001/save  
cheyenne$ hsi mv /home/$USER/input_data /PROJ0001/save
```

- Create directories to label project, intention, destination and any other sorting metric useful to your migration.
- Organizing by directory allows you to use recursive **hsi** commands.

```
/home/username/project1/campaign  
/home/username/project1/delete  
/home/username/project2/univ_storage
```

# Migrating files to Campaign Storage using hsi

- Campaign Storage is POSIX-accessible (in the terminal) on data-access nodes, as well as using the hpss and dav partitions in Slurm.
- Use **hsi cget** operation with **-R** (recursive) and **-A** (tape ordering) options to copy files from HPSS to Campaign Storage most efficiently.

```
cheyenne$ ssh data-access.ucar.edu
data-access$ mkdir /glade/campaign/cisl/csg/PROJ0001
data-access$ cd /glade/campaign/cisl/csg/PROJ0001
data-access$ hsi cget -RA /PROJ0001/save
```

# Migrating files to external storage using Globus

- If available on the external platform, Globus provides a fast, robust, traceable method for migrating data off of HPSS.
- First, use **hsi** to copy files from HPSS to your GLADE scratch space.
- Then, use either the Globus web or command line interface to initiate transfers of these data to the external storage.

```
cheyenne$ cd /glade/scratch/$USER/PROJ0001  
cheyenne$ hsi cget -RA /PROJ0001/save
```

Search for the “NCAR GLADE” endpoint on <http://www.globus.org>

# Migrating files to external storage using bbcp/scp/rsync

- If the external site does not have a Globus endpoint, transfer files from scratch using command-line utilities like **bbcp**, **scp**, or **rsync**.
- For small transfers, **scp** is preferred for ease of use.
- For larger transfers, **bbcp** may provide the best performance as it can use multiple transfer streams. It also supports checkpointing of large files.
  - the client must be installed and in user's PATH on both systems
- Use a session program like GNU screen or tmux to run and track long-running transfer commands.

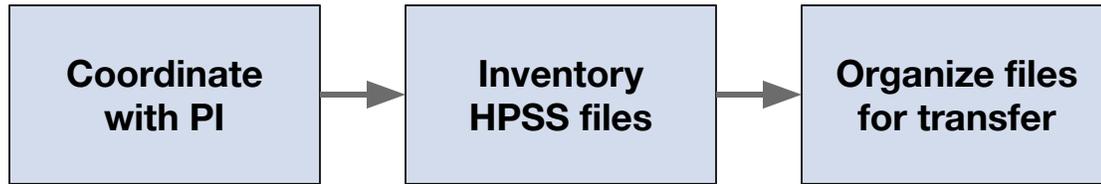
```
cheyenne$ cd /glade/scratch/$USER/PROJ0001
cheyenne$ hsi cget -RA /PROJ0001/save
cheyenne$ scp -r save remote@univ.edu:/projects/PROJ0001
cheyenne$ bbcp -a -r save remote@univ.edu:/projects/PROJ0001
```

# Verifying a successful data migration

- Verify that file inventory and sizes are correct on the destination storage platform using the tools available on that platform (e.g., **ls** queries, Globus transfer logs, verbose modes of bbcp/scp).
- Once you have verified the migration, contact the project PI to get approval to remove files from HPSS.
- Removing files helps CISL track the migration effort and ensures that you don't transfer them again. This step is crucial to the overall process.
  - **Remember to verify first. Data removal is permanent and deleted files are nonrecoverable!**

```
cheyenne$ hsi rm -R /PROJ0001/delete  
cheyenne$ hsi rm -R /PROJ0001/save
```

## Preparation Stages



### Useful commands

#### Preparation

```
hsi ls -lRU /PROJ0001/  
hsi mkdir /PROJ0001/save  
hsi mv /home/$USER/data /PROJ0001/save/data
```

#### Migration

```
hsi cget -RA /PROJ0001/save  
scp -r      save/small \\\  
           remote@univ.edu:/projects/PROJ0001  
bbcp -a -r  save/large  \\\  
           remote@univ.edu:/projects/PROJ0001
```

#### Cleanup

```
hsi rm -R /PROJ0001/save
```

## Verification Stages



Migrate to Campaign Storage

Requires allocation on Campaign Storage

To external using Globus

Requires Globus endpoint on remote system

To external via bbcp, scp, rsync

Requires allocation on remote storage

## Regarding storing new data...

With HPSS going read-only in January 2020, it is critical that you modify your workflows to store data differently moving forward.

- Keep in mind that storage is becoming the bottleneck, so evaluate which data truly need to be preserved and which can be regenerated if needed.
- The current file lifespan on /glade/scratch is 120 days - consider archiving files from scratch **only if necessary**.
- For users with allocations on Campaign Storage, it will serve as the primary cool archive moving forward.
  - Standard terminal access on Casper and data-access nodes
  - Globus access on Cheyenne login and batch nodes via gcert/gci commands (<https://bit.ly/35Sp8PB>)

# Suggested best practices

- Take care to **get** files only a single time in your migration effort - duplicate get operations use valuable tape drive time.
- Record commands and output from those commands into log files.
  - Use verbose modes when available.
- Putting data onto HPSS will slow your migration throughput.
  - You have a five-concurrent-transfer limit on HPSS
- When in doubt, contact CISL for assistance...

# Getting assistance from the CISL Help Desk

<https://www2.cisl.ucar.edu/user-support/getting-help>

- Walk-in: ML 1B Suite 55
- Web: <http://support.ucar.edu>
- Phone: 303-497-2400

Specific questions from today and/or feedback:

- Email: [vanderwb@ucar.edu](mailto:vanderwb@ucar.edu)

