

APR 2019

Upcoming File Storage Migration

Xiaoge Wang, iCER Research Consultant

As many HPC users have noted, there is ongoing migration of user's home directory and research space to a new file storage system. We advise our users to read this wiki page (<https://wiki.hpcc.msu.edu/display/ITH/ufs18+home+directory+migration>), understand the changes on the new file system, watch the progress of this migration and be prepared to take necessary actions.

Some frequently asked questions and answers are posted for your reference:

Q1: How do I find which server my home or research space is on?

A1: Use the following commands to find the file server where your home/research space is located:

- For home space: `df -h $HOME`
- For research space: `df -h /mnt/research/<your_research_space>`

For example:

```
[wangx147@dev-intel18 ~]$ df -h /mnt/research/common-data
Filesystem                Size  Used Avail Use% Mounted on
ufs-12-b.i:/b12/u/rs/common-data  2.T  1.T  90G  57%
/mnt/research/common-data
```

shows that research space `common-data` is on server `ufs-12-b`.

```
[wangx147@dev-intel18 ~]$ df -h /mnt/home/class0
Filesystem                Size  Used Avail Use% Mounted on
ufs-12-b.i:/b12/u/home/class0  50G  1.0M  50G   1% /mnt/home/class0
```

show that user `class0` has home directory on server `ufs-12-b` as well.

If the output shows the `Filesystem` as `ufs18`, this means the space has already been migrated to the new system.

```
[wangx147@dev-intel18 ~]$ df -h /mnt/research/helpdesk
Filesystem      Size  Used Avail Use% Mounted on
ufs18           3.5P  1.6P  1.9P  45% /mnt/research/helpdesk
```

Q2: What do I need to know or do before my home/research space is migrated?

A2: On the day of scheduled migration, the file servers being migrated will be completely inaccessible to users. All users with home space on those servers will be unable to login. All jobs running on those file servers would be terminated. It is the user's responsibility to manage their jobs to avoid an unnecessary loss of jobs. Users can hold their jobs before the scheduled migration time using `scontrol hold <jobID>` and release the hold using `scontrol release <jobID>`. In this way, your queue time may be saved.

Q3: If I still can't login after migration has been completed, what should I do?

A3: In some exceptional cases, although the migration is completed, user accounts may not be automatically enabled and may thus need to be manually activated. If you notice this problem, please submit a ticket immediately to <https://contact.icer.msu.edu>.

Q4: `df -h $HOME` doesn't show the correct storage usage after migration. How could I find the usage of my disk space?

```
[wangx147@dev-intel18 ~]$ df -h $HOME
Filesystem      Size  Used Avail Use% Mounted on
ufs18           3.5P  1.6P  1.9P  45% /mnt/home/wangx147
```

A4: After migration to `ufs18`, please use the command `quota` or the `powertools` command `myquota` to show the disk usage instead. These two commands provide the same usage information. Additionally, the command `myquota` includes the fileset name.

Example of `myquota` output:

```
[wangx147@dev-intel18 bin]$ myquota

Home Directory:
-----
/mnt/home/wangx147      Quota      Used      Free      Files Quota  Files Used  Files Free
ufs.hpcc.msu.edu/home-023  1024G      567G      457G      314572      186253      128319

Research Groups:
-----
common-data           Space      Space      Space      Files      Files      Files
Quota         Used      Available Quota         Used      Available
-----
helpdesk(rs-011)      10240G      5889G      4351G      31457280    24692116    6765164

Temporary Filesystems:
-----
/mnt/ls15 (legacy scratch)  Inodes Used  Quota      Free
36996         1000000     963004

/mnt/scratch (/mnt/g518)    Space Quota  Space Used  Space Free  Filess Quota  Files Used  Files Free
1G           0G        1G         1048576     498011      550565

/mnt/ffs17  Used      Quota
0By        100.00G
```

The output above shows the user's disk usage. Note that this user's home directory is on `ufs18` with fileset as `ufs.hpcc.msu.edu/home-023`. It also shows that the user has two research group spaces: `helpdesk` on `ufs18` with fileset `ufs.hpcc.msu.edu/rs-011` and

`common-data` which is not yet been migrated (group name is not followed by the fileset name in parenthesis).

Q5: Why might the disk usage shown in `quota` after migration different from the disk usage before migration?

A5: As the noted here (<https://wiki.hpcc.msu.edu/display/ITH/ufs18+home+directory+migration>), the compression algorithm used in the new system is different, so you will likely notice a larger usage on the new file system. In some exceptional cases, the usage *reported* under “Used” may be less than your *actual* usage. This is likely because the space contains some files where the group name or owner name is not consistent with attributes of the space so the `quota` command does not count those files into the “Used” bucket. Please submit a ticket if you find such cases so we can help you fix this issue.

Q6: My program ran well before the migration, but after migration I get a “Disk quota exceeded” message. How do I resolve the problem?

A6: This problem could be due to either the increase in file size or the new quota limit on the number of files implemented upon migration as noted here:

<https://wiki.hpcc.msu.edu/display/ITH/ufs18+home+directory+migration>. The solution is to reduce your usage. Consider the following options for reducing your storage:

- Delete or backup unused files or folders to external spaces if you can,
- Use the `tar` or `zip` command to pack an entire directory into a single archive file and store only the archived file when the directory is not in use. You can later unpack the archived file to your scratch space when you need it.

If you need assistance with resolving disk quota issues, please submit a ticket to <https://contact.icer.msu.edu> or visit our help desk on Mondays and Thursdays from 1-2p in BPS 1440.