For safe, long-term data storage, transfer your files to the Lou mass storage system, which allows you to retrieve your stored files quickly and securely whenever you need them.

The Lou mass storage system is an Intel Xeon Gold 6154 "Skylake"-based cluster combined with a Spectra Logic tape storage archive. Lou uses a parallel Data Migration Facility (DMF) system to provide high speed and bandwidth for data transfers between Lou's disks and tapes.

Connecting to Lou

Lou's four hosts, known as Lou front-end systems (LFEs), are designated as lfe[*5-8*].

You can connect automatically to the LFE that has the lowest load by using the "lou" or "lfe" hostnames in your ssh command line: ssh lfe or ssh lou. This behavior is similar to the load balancing process that occurs on the PFEs when you run ssh pfe.

Transferring Files Between /nobackup Filesystems and Lou

The Lustre (/nobackup) filesystems are mounted on the LFEs. To transfer files between your /nobackup filesystem and your Lou home filesystem, use the local file transfer commands shiftc, cp, or mcp. For example:

lfe% cp /nobackup/*your_username*/filename /u/*your_username*

You can also create a tar file that contains the data in one of your /nobackup subdirectories and then store the tar file in the LFE home filesystem. For example:

lfe% cd /nobackup/*your_username*
lfe% tar cf /u/*your_username*/*mydir*.tar *mydir*

Transferring Files from Your Local System to Lou

Recommended: Use the Shift tool (shiftc), which accepts both the alias and specific hostnames: lou, lfe, or lfe[*5-8*].

For more information about transferring files to and from your local system, see Remote File Transfer Commands.

Data Migration Between Disk and Tapes

In addition to 7.6 petabytes of disk space, Lou has 51 LTO-8 tape drives. Each LTO-8 tape holds 12 TB of uncompressed data, for a total storage capacity of approximately 1040 petabytes, or one exabyte (with normal 35% compression). Data migration (from disk to tape) and unmigration (from tape to disk) are managed by the Data Migration Facility (DMF).

Data stored on Lou's home filesystems (on disk) are automatically migrated to tape. Two copies of your data are written to tape media in silos located in separate buildings. When it is necessary to make room for more data, some files that have been written to tape may have their data "released" from disk. This means that the file is still visible on the filesystem, but the data must be retrieved from tape before the file can be used.

When a file completes writing to tape on Lou, its ctime attribute is updated to signal that it is available to be backed up by system software. This should not have an affect on any file transfer applications such as rsync.

If you need to retrieve data that is on tape, be sure to unmigrate the data from tape to your home filesystem on Lou before transferring it to other systems.

TIP: If you use the Shift tool (shiftc) for file transfers, it will automatically ensure that files on Lou are online before the transfer.

If you are not using Shift, you can use the following DMF commands to retrieve your files from tape:

```
$ dmls -al file1 file2 . . .    # show the status of your files.
$ dmget file1 file2 . . . &     # retrieve your file from tape.
```

At this point, you can start your transfer and the files will transfer as they come online.

WARNING: Do not use the /nobackup filesystems for long-term data storage. As the names suggest, these filesystems are not backed up, so any files that are removed cannot be restored. You should store essential data on more permanent storage devices, such as Lou.

For more tips on how to use the Lou storage systems more effectively, see:

- Portable File Names and Sizes
- Dealing with Slow File Retrieval

Post-Processing Your Data

To perform post-processing tasks on Lou data, use the Lou data analysis nodes, which provide dedicated PBS resources for that purpose.

Post-processing is not permitted on the LFEs. To enforce this, a system monitoring tool called `query_wms` runs on the LFEs and kills any user process that uses more than 1 GB of memory.

Quota Limits On Lou

There are no disk quota limits on your Lou home filesystem. However, there *are* limits on the number of files (inodes):

- 250,000 inode soft limit (14-day grace period)
- 300,000 inode hard limit

To check your quota status and usage on your home filesystem, log into Lou (lfe) and run the `quota -ls` command as follows:

```
lfe% quota -ls
Disk quotas for user username (uid xxxx):
     Filesystem  blocks   quota   limit   grace   files   quota   limit   grace
/dev/cxvm/sfa2_s2n
               87422M       0       0           59381     250k    300k
```

See Quota Policy on Disk Space and Files for more information about HECC system quotas.

The Lou data analysis nodes (LDANs) provide dedicated PBS resources to perform post-processing tasks on your Lou mass storage data. You can also use them for performing pre-processing or post-processing tasks on active data residing in the home or /nobackup filesystems of NAS compute resources. The LDANs can be accessed only through PBS jobs in which resources are dedicated to the job owner.

The LDANs contain Intel Xeon Gold 6154 "Skylake" processors, equipped with more memory than the Electra Skylake nodes, as follows:

- Each ldan[*11-12*] has 36 cores and 768 GB of memory
- Each ldan[*13-14*] has 36 cores and 1.5 TB of memory

The Lou, Lustre, and home filesystems of NAS compute resources are mounted on the LDANs. You can access them in a PBS job running on an LDAN by using the following paths:

Lou:
    /u/*username* or /lou/lax where X is one of 1-6
Lustre:
    /nobackup/*username*
Home:
    /homeX/u/*username* where X is one of 1,3,4,5,6,7

Software Modules

All of the software packages that you use on the front ends or compute nodes of any NAS x86_64 compute resouces are also available on the LDANs. However, both the Lou front-end (LFE) nodes and LDANs use startup files from your Lou home directory rather than your home directory of the NAS compute systems.

Therefore, if your PBS jobs rely on loading specific software modules or environment variables in the system startup files (such as .cshrc), in order to run them on the LDANs you must modify your startup files on Lou to load the modules and set appropriate environment variables.

Running PBS Jobs on LDANs

The PBS server for the LDANs is pbspl1. To use the LDANs, submit your jobs to the ldan queue. Each job can use only *one* LDAN for up to *three days*, and each user can have a maximum of *two jobs* running simultaneously.

Before You Begin: If you want to process archive data that has been migrated offline to tape, use the Data Migration Facility command dmget to migrate the data back to disk before submitting jobs to the LDANs.

Submitting Your PBS Job

You can submit interactive PBS jobs to the LDANs from either the LFEs or the PFEs. You can submit PBS job scripts from either your Lustre filesystem (/nobackup/*username)* or your Lou home filesystem (/u/*username)*.

WARNING: Do not submit job scripts from the home filesystem of the NAS compute resources (/homeX/username) because PBS error and output files cannot be copied back there from the LDANs.

Use the :mem=*xxx*GB attribute to specify how much memory you need for your job. This will instruct PBS to allocate an appropriate LDAN.

Note: Keep in mind that the amount of memory on a node that is available for your job is slightly less than the total physical memory, because the system kernel can use up to 4 GB of memory in each node.

Sample PBS Script

```
#PBS -lselect=1:ncpus=36:mem=750GB
#PBS -lwalltime=2:00:00
#PBS -q ldan

### Remember to load necessary modules
### or set environment variables that you need

module load xxx

### Note that the default directory a PBS job
### is in when it starts is the Lou home filesystem.
### Change directory to the appropriate directory
### for your pre- or post- processing work.

cd $PBS_O_WORKDIR

### If your executable for pre- or post- processing
```

is located under your Pleiades home filesystem,
use the pathname /homeX/*username*/bin/....
do not use the pathname /u/*username*/bin/...

/homeX/*username*/bin/*your_processing_code* < input > output

Running Graphics Applications on LDANs

If you want to run graphics applications that consume enough memory to require one of the larger-memory LDANs, submit an interactive PBS job from one of the LFEs that specifies mem=1040GB. For example:

```
lfe% qsub -I -q ldan -lselect=1:mem=1040GB,walltime=1:00:00
```

This will provide you with an LDAN for the duration of the specified wall time. Then, from any other terminal window, you can use SSH to connect directly to the LDAN assigned to your job, and run your graphics application interactively. If you have SSH passthrough set up, you can even SSH directly in to the LDAN from your workstation by adding the LDANs as allowed SSH passthrough targets in your .ssh/config file.

To run graphics applications directly from a PBS session submitted from an LFE, you need to export your DISPLAY environment by adding the -X option to the qsub command, as follows:

```
lfe% qsub -I -X -q ldan -lselect=1:mem=750GB,walltime=1:00:00
```

If you are a remote user, you can also explore using Virtual Network Computing (VNC) to connect to NAS systems.