

Pleiades Lustre Filesystems

Category: Lustre on Pleiades

Summary: The Lustre filesystems on Pleiades are called "nobackup." As the name suggests, these filesystems are for temporary use, and are not backed up. Lustre can handle many large files, but you cannot store those files on Pleiades; if you want to save them, move them to Lou.

Pleiades has several Lustre filesystems (`/nobackupp[1-6]`) that provide a total of about 6.795 petabytes of storage and serve thousands of cores. These filesystems are managed under Lustre software version 1.8.6.

Lustre filesystem configurations are summarized at the end of this article.

WARNING: As the names suggest, these filesystems are not backed up, so any files that are removed *cannot* be restored. Essential data should be stored on Lou[1-2] or on other, more permanent storage.

Which /nobackup Should I Use?

Once you are granted an account on Pleiades, you will be assigned to use one of the Lustre filesystems. Find out which Lustre filesystem you have been assigned to by typing the following:

```
pfel% ls -l /nobackup/your_username
lrwxrwxrwx 1 root root 19 Feb 23 2010 /nobackup/username -> /nobackupp2/username
```

In the above example, the symlink from `/nobackup` to `/nobackupp2` shows that the user's assigned nobackup system is `/nobackupp2`.

Default Quota and Policy on /nobackup

Disk space and inodes quotas are enforced on the `/nobackup` filesystems. The default soft and hard quota limits for inodes are 75,000 and 100,000, respectively. Those for the disk space are 500 gigabytes and 1 terabyte, respectively. To check your disk space and inodes usage and quota on your `/nobackup`, use the `lfs` command and type the following:

```
%lfs quota -u username /nobackup/username
Disk quotas for user username (uid nnnn):
  Filesystem  kbytes      quota   limit   grace   files   quota   limit   grace
/nobackup/username 1234  530000000 1100000000  -      567   75000 100000  -
```

The NAS quota policy states that if you exceed the soft quota, an email will be sent that lists your current usage and remaining grace period. It is expected that users will occasionally exceed their soft limit, as needed; however after 14 days, users who are still over their soft limit will have their batch queue access to Pleiades disabled.

If you anticipate having a long-term need for higher quota limits, please send a justification via email to support@nas.nasa.gov. This will be reviewed by the HECC Deputy Project Manager for approval.

For more information, see also, [Quota Policy on Disk Space and Files](#).

NOTE: If you reach the hard limit while your job is running, the job will die prematurely without providing useful messages in the PBS output/error files. A Lustre error with code **-122** in the system log file indicates that you are over your quota.

In addition, when a Lustre filesystem is full, the jobs writing to it will hang. A Lustre error with code **-28** in the system log file indicates that the filesystem is full. The NAS Control Room staff normally will send out emails to those using the most space, asking them to clean up their files.

Lustre File Systems Configurations

In the table below, /nobackupp[1-6] are abbreviated as nbp[1-6]. P=Petabytes; T=Terabytes

Pleiades Lustre Configurations						
Filesystem	nbp1	nbp2	nbp3	nbp4	nbp5	nbp6
# of MDSEs	1	1	1	1	1	1
# of MDTs	1	1	1	1	1	1
size of MDTs	0.9T	0.9T	0.6T	0.6T	0.8T	0.9T
# of usable inodes on MDTs	~256x10 ⁶	~256x10 ⁶	~173x10 ⁶	~173x10 ⁶	~512x10 ⁶	~256x10 ⁶
# of OSSes	8	8	8	8	8	8
# of OSTs	120	120	60	60	120	120
size/OST	15T	15T	7.1T	7.1T	15T	7.1T
Total Space	1.7P	1.7P	424T	424T	1.7P	847T
Default Stripe Size	4M	4M	4M	4M	4M	4M
Default Stripe Count	1	1	1	1	1	1

NOTE: After January 13, 2011, directories without an explicit stripe count and/or stripe size adopted the new stripe count of 1 and stripe size of 4MB. However, old files in that directory retain their old default values. New files that you create in these directories will adopt the new default values.

Each Pleiades Lustre filesystem is shared among many users. To get good I/O performance for your applications and avoid impeding the I/O operations of other users, read the related articles listed below.

Article ID: 225

Last updated: 24 Jul, 2012

Computing at NAS -> Best Practices -> Lustre on Pleiades -> Pleiades Lustre Filesystems
<http://www.nas.nasa.gov/hecc/support/kb/entry/225/?ajax=1>