

Storage Areas & NVMe

Tom Papatheodore

Oak Ridge Leadership Computing Facility

Spock Training – May 20, 2021

ORNL is managed by UT-Battelle, LLC for the US Department of Energy



U.S. DEPARTMENT OF
ENERGY

Available Storage Areas on Spock

NFS Directories – This is where you might want to keep source code, scripts, etc.

NOTE: These directories are read-only from the compute nodes!

`/ccs/home/<userid>`

- Your personal home directory

`/ccs/proj/<projid>`

- Can be accessed by all members of your project

Long-term storage

Backup provided

Not purged

50 GB Quota

GPFS Directories (parallel file system) – This is where you should write data when running on Spock's compute nodes.

`/gpfs/alpine/<projid>/scratch/<userid>`

- Your personal GPFS scratch directory

`/gpfs/alpine/<projid>/proj-shared`

- Can be accessed by all members of your project

`/gpfs/alpine/<projid>/world-shared`

- Can be accessed by all users on the system

Short-term storage

No backup provided

Purged after 90 days

50 TB Quota

Archive Storage (HPSS)

User- and Project-Centric Archive Directories – This is where you should backup important data for long-term storage

NOTE: HPSS is not accessible from Spock, but is accessible from the data-transfer nodes (DTNs) → dtn.ccs.ornl.gov

[/home/<userid>](#)

- User archive directory
- Data from <userid> not specific to any project

[/hpss/prod/<projid>/users/<userid>](#)

- Member archive
- Data specific to <projid> accessible only by <userid>

[/hpss/prod/<projid>/proj-shared](#)

- Project archive
- Data specific to <projid> accessible by all members of <projid>

[/hpss/prod/<projid>/world-shared](#)

- World archive
- Data specific to <projid> accessible by all users of the system

No backup provided

Not purged

100 TB Quota

NVMe (SSD)

```
#!/bin/bash
#SBATCH -A <projid>
#SBATCH -J nvme_test
#SBATCH -o %x-%j.out
#SBATCH -t 00:05:00
#SBATCH -p batch
#SBATCH -N 1
#SBATCH -C nvme

date

# Change directory to user scratch space (GPFS)
cd /gpfs/alpine/<projid>/scratch/<userid>

echo " "
echo "*****ORIGINAL FILE*****"
cat test.txt
echo "*****"

# Move file from GPFS to SSD
mv test.txt /mnt/bb/<userid>

# Edit file from compute node
srun -n1 hostname >> /mnt/bb/<userid>/test.txt

# Move file from SSD back to GPFS
mv /mnt/bb/<userid>/test.txt .

echo " "
echo "*****UPDATED FILE*****"
cat test.txt
echo "*****"
```

Each Spock compute node has [2x] 3.2 TB NVMe

- KIOXIA Model: KCM6XVUL3T20
- 6900 MB/s (read), 4200 MB/s (write)

Users must

- Request access to NVMe (e.g., `#SBATCH -C nvme`)
- Move/copy data to/from NVMe

When `-C nvme` is specified for a job allocation, the SSD will be available at `/mnt/bb/<userid>`

```
$ cat nvme_test-<jobid>.out
Mon May 17 12:28:18 EDT 2021

*****ORIGINAL FILE*****
This is my file. There are many like it but this one is mine.
*****

*****UPDATED FILE*****
This is my file. There are many like it but this one is mine.
spock25
*****
```


Questions?

Summit here



Frontier here

