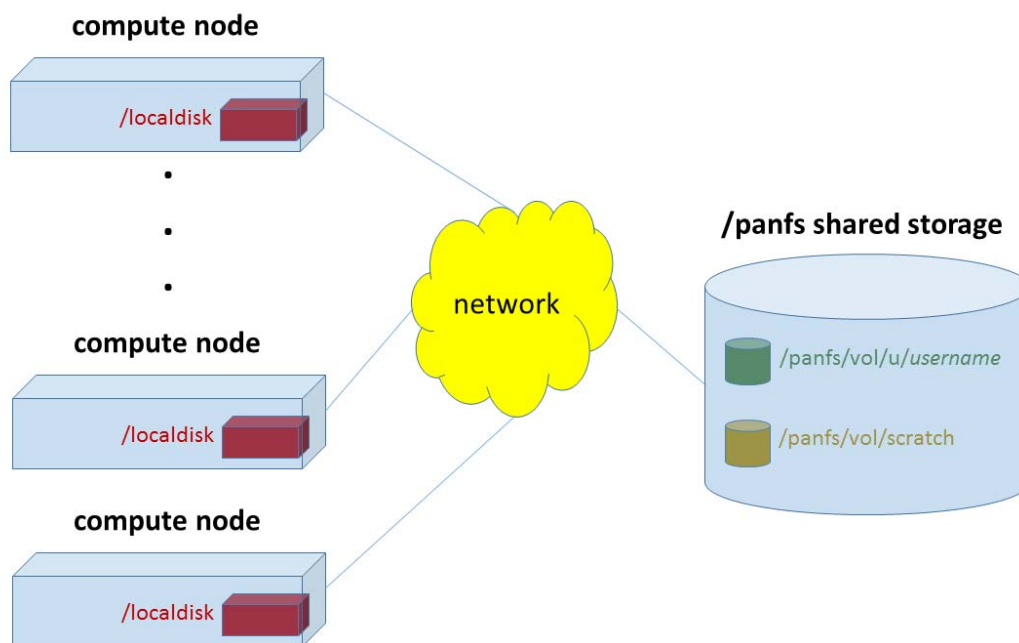


Disk Space Selection for your PBS Jobs on Raad

We would like to inform and/or remind users that there are three ways they can use disk resources for their jobs on raad:

- (a) use the space within your home directory for the needs of you running jobs
- (b) use the `/panfs/vol/scratch` space for job disk needs for the duration of the job and copy back results to your home directory at the end of the job
- (c) use the `/localdisk` space on the compute node itself for job disk needs for the duration of the job and copy back results to your home directory at the end of the job



The diagram above illustrates the locations of these three disk areas in the context of the cluster's architecture. Because the raad storage system (`/panfs`) has nearly reached the limits of its capacity, we discourage users from employing options *a* and *b* above. Not only is option *c* viable for the vast majority of existing jobs, it will provide better disk performance than the other two options. The only users who will likely be unable to benefit from option *c* are those who run multi-node MPI jobs that do I/O to file(s) shared by the MPI tasks, or those who produce very large output files (>80GB).

Most people have been using option *a* as it appears to be the least complicated. However, to employ options *b* or *c*, you must first include the following directive in your PBS job file...

```
#PBS -W sandbox=PRIVATE
```

...then inside your job, you have access to two different disk spaces:

1) By doing a “cd \$PBS_JOBDIR” in the job, you change directory to a temporary location under /panfs/vol/scratch that is created ONLY for the duration of your job. In using this space, it will be your responsibility to copy any files needed for your execution into this space at the beginning of the job script and copying back any results file from this space at the end of the job script. This temporary directory and its contents will be deleted automatically by PBS once the job concludes. Remember, we are recommending that users *do not* use this method unless they absolutely need to.

2) By doing a “cd \$TMPDIR” in the job, you change directory to a temporary location under /localdisk on the node where the job is running. This directory is also created ONLY for the duration of your job. In using this space, it will be your responsibility to copy any files needed for your execution into this space at the beginning of the job script and copying back any results file from this space at the end of the job script. This temporary directory and its contents will be deleted automatically by PBS once the job concludes.

The difference between 1 and 2 is that \$PBS_JOBDIR resides on the Panasas shared storage system whereas \$TMPDIR resides on the local SSD disk of the compute node. The SSD has limited space (< 80GB) for user needs while the Panasas storage can give you much more disk space. On the other hand, SSD will almost always be much faster.

Sample Job File to Use /localdisk

Assume you have made a sub-directory within your home directory called mySubmitDir. All the files you need to launch your job have been placed in this folder:

```
[fachaud74@raad mySubmitDir]$ ls
myJobFile.job  myInputFile  simulationConfigFile.cfg
[fachaud74@raad mySubmitDir]$
```

Your job file (myJobFile.job) reads like this:

```
#PBS -S /bin/bash
#PBS -N myTestJob
#PBS -l ncpus=16,walltime=168:00:00,mem=48gb
#PBS -j oe
#PBS -W sandbox=PRIVATE

cd $TMPDIR
cp -Rp $PBS_O_WORKDIR/* .

module load softwareOfInterest
softwarePackage.exe -I myInputFile > myResultsFile
cp myResultsFile $PBS_O_WORKDIR
```

You submit your job with “qsub myJobFile.job” and when the job finishes, you should see something similar to the following in your job folder:

```
[fachaud74@raad mySubmitDir]$ ls  
myJobFile.job myInputFile simulationConfigFile.cfg myResultsFile  
myTestJob.o6063450  
[fachaud74@raad mySubmitDir]$
```