

Quickstart for Using NAS GPU Resources During GPU Hackathon

This quickstart guide assumes that you already have a NAS account and can ssh from your local desktop to the Pleiades front-ends (PFEs). If you experience difficulty to ssh to the PFEs, please send an email ASAP to support@nas.nasa.gov to request assistance.

There are 5 nodes (each with 36 CPU cores, 4 V100 GPU cards and about 360 GB CPU memory) reserved from 7 AM Sept 28 to 7 PM Oct 7 (through a special GID n2008 and PBS queue 'hackathon') for use by mentors and team members of this hackathon. Here are a few guidelines to follow when using the reserved resources:

- Use the GID n2008 with the option -W group_list=n2008

Please first check if you are a member of n2008 with either of these two commands:

```
pfe% groups your_NAS_username  
or  
pfe% grep n2008 /etc/group
```

If not, send an email to henry.jin@nasa.gov and provide your team name and username. If confirmed, add the -W group_list option in the qsub command line or in PBS script:

```
qsub -W group_list=n2008  
or  
#PBS -W group_list=n2008
```

- Use the PBS queue name 'hackathon' served by the PBS server pbspl4:

```
qsub -q hackathon@pbspl4  
or  
#PBS -q hackathon@pbspl4
```

- Each team is entitled to 2 GPU cards, M-F from 7 AM to 7 PM PDT. The 2 GPU cards can be requested by a single team member in one PBS job or split into two PBS jobs each requesting only 1 GPU card. There will be no automatic enforcement. However, please cooperate with this guideline so that each team gets to use the reserved resources during the time specified above. Outside of the specified time, more resources can be requested. However, if you experience difficulty in getting any resources to use, please ask the organizer to coordinate.
- Sample PBS script to request 1 GPU card

Note that the ncpus and mem amount can be adjusted. If you want to use 2 GPU cards in one PBS job, replace ngpus=1 with ngpus=2.

This sample script assumes that you want to use bash. If your default shell is csh, and you want to use csh for your PBS job, then remove the first line (`#!/bin/bash`) and the line that sources `global.profile`.

```
##### Sample PBS script #####
#!/bin/bash
#PBS -q hackathon@pbspl4
#PBS -W group_list=n2008
# The -l in the following three lines is l as in list
#PBS -l select=1:ncpus=9:ngpus=1:mem=90GB:model=sky_gpu
#PBS -l walltime=12:00:00
#PBS -l place=free:shared

# source this file so that you can use the module commands
source /usr/local/lib/global.profile
module use -a /nasa/modulefiles/testing
module load comp-pgi/20.4
module load cuda/11.0

# put what you want to do next
...
# End of Sample PBS script
```

- Use interactive PBS session

If you want to use the resources in an interactive environment, then create an interactive PBS session through the `qsub -I` (I as in I am) option:

```
pfe% qsub -I -q hackathon@pbspl4 -W group_list=n2008
-lselect=1:ncpus=9:ngpus=1:mem=90GB:model=sky_gpu,walltime=12:00:00,place=free:shared
```

- Check your job status with `qstat` and delete job with `qdel`

```
pfe% qstat -nu your_NAS_username @pbspl4
pfe% qdel jobid.pbspl4
```

You can also login from a pfe to the PBS server `pbspl4` directly, and run the `qsub`, `qstat` and `qdel` commands without `@pbspl4`.

```
pbspl4% qsub -I -q hackathon -l ...
pbspl4% qstat -a
pbspl4% qstat -nu your_NAS_username
pbspl4% qdel jobid
```

Ref: https://www.nas.nasa.gov/hecc/support/kb/using-gpu-nodes_298.html