

## Gridengine to Slurm conversion

### User Commands

**Interactive Login**  
with X11 forwarding

### Job submission

specific PE/partition  
wildcard PE/partition  
partial wildcard PE/partition  
run on 2 nodes (1 process / node)  
run on 4 nodes with 2 processes / node  
resource request  
resource flags  
negative resource request  
Example: Avoid cluster 17  
run time limit  
memory limit  
request GPUs  
request GPU with type  
request GPU with GPU memory  
request GPU with compute capability

### Job deletion

### Job status by job id

### Job status by user

### Job hold

### Job release

### Queue list

### List nodes

### Cluster status

### Expand nodelist

### Gridengine

qrsh  
qrsh

qsub *jobscript*  
qsub -pe *pname nslots jobscript*  
qsub -pe \\* 10 *jobscript*  
qsub -pe omp\\* 10 *jobscript*  
qsub -pe omp\\*\_1 2 *jobscript*  
qsub -pe omp\\*\_2 8 *jobscript*  
qsub -l *feature[=value]* ...  
qsub -l *feature*  
qsub -l *feature='!value'* ...  
qsub -l cluster='!17' ...  
qsub -l h\_rt=*seconds* ...  
*not implemented, just a reservation*  
qsub -l gpu=*gpus* ...  
qsub -l gpu=1,gputype=*gputype* ...  
qsub -l gpu=1,gpu\_mem=40g ...  
qsub -l gpu=1,ccc=7.0 ...  
qdel *jobid*  
qstat -j *jobid*  
qstat -u *user*  
qhold *jobid*  
qrls *jobid*  
qstat  
qhost  
qhost -q

### Slurm

srun --pty bash  
srun -x11 -pty bash  
*or*  
salloc sh -c 'ssh -Y \$SLURM\_NODELIST' <sup>1</sup>  
sbatch *jobscript*  
sbatch -p *pname jobscript* <sup>2</sup>  
sbatch -p \\* *jobscript*  
sbatch -p omp\\* -N 10 *jobscript*  
sbatch -p omp\\* -N 2 *jobscript*  
sbatch -p omp\\* -N 4 -c 2 *jobscript*  
sbatch -C *feature[:value]* ...  
sbatch -C *feature* ...  
*impossible, use other positive feature*  
sbatch -C ib:mlnx ....  
sbatch --time=*minutes* .... <sup>3</sup>  
sbatch --mem=size[M|G] <sup>4</sup>  
sbatch --gres=gpu:*gpus* ...  
sbatch --gres=gpu:*gputype:gpus* ...  
sbatch --gres=gpu:1 --gres=gpu\_mem:8000M  
sbatch --gres=gpu:1,ccc:70 ... <sup>5</sup>  
scancel *jobid*  
squeue -j *jobid*  
squeue -u *user*  
scontrol hold *jobid*  
scontrol release *jobid*  
squeue  
sinfo -N -o "%10N %6c %10m %G" <sup>6</sup>  
sinfo  
scontrol show hostnames 'node[782-784]'

- 1 The two methods are subtly different, the 2nd uses ssh to forward the X11 connection. ssh into a node works only if there is a job or allocation on that node, and the ssh session will be attached to that job
- 2 The number of nodes and slots defaults to 1
- 3 Unless specified, the default run time limit is 12 hours.
- 4 The default unit is megabytes
- 5 The values are multiplied by 10 for slurm (it can't handle floating point values here)
- 6 Feel free to adjust the formatting and/or create an alias. Check the `sinfo` man page for details on the output formats.

## Batch Scripts:

### Gridengine:

```
#!/bin/bash --login
#$ -N gpu
#$ -o gpu.$HOSTNAME.$JOB_ID.out
#$ -j y
#$ -cwd
#$ -pe mp 2

#$ -l gpu=1

module load cuda/12.0
nvidia-smi

#end
```

### Slurm:

```
#!/bin/bash --login
#SBATCH --job-name=gpu
#SBATCH --output=gpu.%N.%A.out

#SBATCH --ntasks=1
#SBATCH --cpus-per-task=2
#SBATCH --gres=gpu:1

module load cuda/12.0
nvidia-smi

# end
```

For more options and details, please check the man pages for *sbatch*.