

# SGE to SLURM conversion

Sun Grid Engine (SGE) and **SLURM** job scheduler concepts are quite similar. Below is a table of some common SGE commands and their SLURM equivalents.

Some common commands and flags in SGE and SLURM with their respective equivalents:

User Commands	SGE	SLURM
<b>Interactive login</b>	qlogin	<code>srun --pty bash</code> or <code>srun (-p "partition name"--time=4:0:0 --pty bash</code> For a quick dev node, just run "sdev"
<b>Job submission</b>	qsub [script_file]	sbatch [script_file]
<b>Job deletion</b>	qdel [job_id]	scancel [job_id]
<b>Job status by job</b>	qstat -u \* [-j job_id]	squeue [job_id]
<b>Job status by user</b>	qstat [-u user_name]	squeue -u [user_name]
<b>Job hold</b>	qhold [job_id]	scontrol hold [job_id]
<b>Job release</b>	qrls [job_id]	scontrol release [job_id]
<b>Queue list</b>	qconf -sql	squeue
<b>List nodes</b>	qhost	sinfo -N OR scontrol show nodes
<b>Cluster status</b>	qhost -q	sinfo
<b>GUI</b>	qmon	sview
<b>Environmental</b>		
<b>Job ID</b>	\$JOB_ID	\$\$SLURM_JOBID
<b>Submit directory</b>	\$\$SGE_O_WORKDIR	\$\$SLURM_SUBMIT_DIR

<b>Submit host</b>	\$SGE_O_HOST	\$SLURM_SUBMIT_HOST
<b>Node list</b>	\$PE_HOSTFILE	\$SLURM_JOB_NODELIST
<b>Job Array Index</b>	\$SGE_TASK_ID	\$SLURM_ARRAY_TASK_ID
<b>Job Specification</b>		
<b>Script directive</b>	#\$	#SBATCH
<b>queue</b>	-q [queue]	-p [queue]
<b>count of nodes</b>	N/A	-N [min[-max]]
<b>CPU count</b>	-pe [PE] [count]	-n [count]
<b>Wall clock limit</b>	-l h_rt=[seconds]	-t [min] OR -t [days-hh:mm:ss]
<b>Standard out file</b>	-o [file_name]	-o [file_name]
<b>Standard error file</b>	-e [file_name]	e [file_name]
<b>Combine STDOUT &amp; STDERR files</b>	-j yes	(use -o without -e)
<b>Copy environment</b>	-V	--export=[ALL   NONE   variables]
<b>Event notification</b>	-m abe	--mail-type=[events]
<b>send notification email</b>	-M [address]	--mail-user=[address]
<b>Job name</b>	-N [name]	--job-name=[name]
<b>Restart job</b>	-r [yes no]	--requeue OR --no-requeue (NOTE: configurable default)

<b>Set working directory</b>	-wd [directory]	--workdir=[dir_name]
<b>Resource sharing</b>	-l exclusive	--exclusive OR--shared
<b>Memory size</b>	-l mem_free=[memory][K M G]	--mem=[mem][M G T] OR --mem-per-cpu=[mem][M G T]
<b>Charge to an account</b>	-A [account]	--account=[account]
<b>Tasks per node</b>	(Fixed allocation_rule in PE)	--tasks-per-node=[count]
		--cpus-per-task=[count]
<b>Job dependency</b>	-hold_jid [job_id   job_name]	--depend=[state:job_id]
<b>Job project</b>	-P [name]	--wckey=[name]
<b>Job host preference</b>	-q [queue]@[node] OR -q [queue]@@[hostgroup]	--nodelist=[nodes] AND/OR --exclude=[nodes]
<b>Quality of service</b>		--qos=[name]
<b>Job arrays</b>	-t [array_spec]	--array=[array_spec] (Slurm version 2.6+)
<b>Generic Resources</b>	-l [resource]=[value]	--gres=[resource_spec]
<b>Lincenses</b>	-l [license]=[count]	--licenses=[license_spec]
<b>Begin Time</b>	-a [YYMMDDhhmm]	--begin=YYYY-MM-DD[THH:MM[:SS]]

--	--	--

SGE	SLURM
qstat	squeue

qstat -u username qstat -f	squeue -u username squeue -al
qsub qsub -N jobname qsub -m beas qsub -M SUNetID@stanford.edu qsub -l h_rt=24:00:00 qsub -pe dmp4 16 qsub -l mem=4G qsub -P projectname qsub -o filename qsub -e filename qsub -l scratch_free=20G	sbatch sbatch -J jobname sbatch --mail-type=ALL sbatch --mail-user=SUNetID@stanford.edu sbatch -t 24:00:00 sbatch -p node -n 16 sbatch --mem=4000 sbatch -A projectname sbatch -o filename sbatch -e filename sbatch --tmp=20480
# Interactive run, one core	# Interactive run, one core
qcrsh -l h_rt=8:00:00	salloc -t 8:00:00 interactive -p core -n 1 -t 8:00:00
qdel	scancel

### SGE for a single-core application

```
#!/bin/bash
#
#
#$ -N test
#$ -j y
#$ -o test.output
#$ -cwd
#$ -M SUNetID@stanford.edu
#$ -m bea
```

### SLURM for a single-core application

```
#!/bin/bash -l
# NOTE the -l flag!
#
#SBATCH -J test
#SBATCH -o test.output
#SBATCH -e test.output
# Default in slurm
#SBATCH --mail-user SUNetID@stanford.edu
#SBATCH --mail-type=ALL
```

```

# Request 5 hours run time
## -l h_rt=5:0:0
## -P your_project_id_here
#
## -l mem=4G
#
<call your app here>

```

```

# Request 5 hours run time
#SBATCH -t 5:0:0
#SBATCH -A your_project_id_here
#SBATCH --mem=4000
#SBATCH -p normal
# NOTE: if you want more mem consider the hns
node
# it has 1.5 TB RAM use "-p hns"

<call your app here>

```

Comparison of some parallel environments set by sge and slurm

SGE	SLURM
\$JOB_ID	\$SLURM_JOB_ID
\$NSLOTS	\$SLURM_NPROCS