

Translate PBS/Torque to SLURM

How is Slurm different from PBS?

Slurm is different from PBS in a number of important ways, including the commands used to submit and monitor jobs, the syntax used to request resources, and the way environment variables behave.

Some specific ways in which Slurm is different from PBS include:

- Slurm will not allow a job to be submitted if there are insufficient account cycles or no resources matching your request. PBS would queue the job, but it would never run.
- Slurm leverages the kernel's control groups (cgroups) to limit memory and CPU usage for each job and shares those limits across all ssh sessions for the owner of the job. This impacts anyone accessing a compute node running their job. Instead of getting new system resources, you share the same limits as your running job.
- What PBS called queues, Slurm calls partitions.
- Resources are assigned per "task"/process.
- Environmental variables of the submitting process are passed to the job.

How to Monitor Your Job(s):

After submitting your job you can check its status with these new commands:

User commands	PBS/Torque	SLURM
Job submission	qsub [filename]	sbatch [filename]
Job deletion	qdel [job_id]	scancel [job_id]
Job status (by job)	qstat [job_id]	squeue --job [job_id]
Full job status (by job)	qstat -f [job_id]	scontrol show job [job_id]
Job status (by user)	qstat -u [username]	squeue --user=[username]
Get estimated start time of job	showstart <job_id>	squeue -j <job_id> --start

How to Read/Set Slurm Environment Variables:

Like PBS, Slurm sets its own environment variables within your job:

Environment variables	PBS/Torque	SLURM
Job ID	\$PBS_JOBID	\$SLURM_JOBID
Submit Directory	\$PBS_O_WORKDIR	\$SLURM_SUBMIT_DIR
Node List	\$PBS_NODEFILE	\$SLURM_JOB_NODELIST
Hostname job was submitted from	\$PBS_O_HOST	\$SLURM_SUBMIT_HOST

How to Request Resources:

The way that you work with resources in Slurm is similar to PBS but there are some significant differences, including the syntax used to submit requests. See the table below for some of the options most frequently used to request resources.

Job specification	PBS/Torque	SLURM
Script directive	#PBS	#SBATCH
Job Name	-N [name]	--job-name=[name] OR -J [name]
Node Count	-l nodes=[count]	--nodes=[min[-max]] OR -N [min[-max]]
CPU Count	-l ppn=[count]	--ntasks-per-node=[count]
CPUs Per Task		--cpus-per-task=[count]
Memory Size	-l mem=[MB]	--mem=[MB] OR --mem-per-cpu=[MB]
Wall Clock Limit	-l walltime=[hh:mm:ss]	--time=[min] OR --time=[days-hh:mm:ss]
Node Properties	-l nodes=4:ppn=8:[property]	--constraint=[list]
Standard Output File	-o [file_name]	--output=[file_name] OR -o [file_name]
Standard Error File	-e [file_name]	--error=[file_name] OR -e [file_name]
Combine stdout/stderr	-j oe (both to stdout)	(Default if you don't specify --error)
Job Arrays	-t [array_spec]	--array=[array_spec] OR -a [array_spec]

Valid Job States:

Below are the job states you may encounter when monitoring your job(s) in Slurm.

Meaning	Code	State
Job was canceled	CA	Canceled
Job completed	CD	Completed
Job resources being configured	CF	Configuring
Job is completing	CG	Completing
Job terminated with non-zero exit code	F	Failed
Job terminated due to failure of node(s)	NF	Node Fail
Job is waiting for compute node(s)	PD	Pending
Job is running on compute node(s)	R	Running
Job terminated upon reaching its time limit	TO	Timeout