PBS to Slurm Migration

Quick reference for users migrating from the old PBS/Torque system to Slurm.

Command Equivalents

PBS/Torque	Slurm	Description
`qsub job.sh`	`sbatch job.sh`	Submit a job
`qsub -I`	`srun --pty bash`	Interactive session
`qstat`	`squeue`	View jobs
`qstat -u username`	`squeue -u username`	Your jobs
`qdel JOBID`	`scancel JOBID`	Cancel a job
`pbsnodes`	`sinfo`	Node status
`qstat -f JOBID`	`scontrol show job JOBID`	Job details

Directive Equivalents

PBS/Torque	Slurm	Description
`#PBS -N name`	`#SBATCH --job-name=name`	Job name
`#PBS -q queue`	`#SBATCH --partition=partition`	Queue/partition
`#PBS -l nodes=1:ppn=8`	`#SBATCH --nodes=1 --cpus-per-task=8`	Nodes and cores
`#PBS -l mem=4gb`	`#SBATCH --mem=4G`	Memory
`#PBS -l walltime=02:00:00`	`#SBATCH --time=02:00:00`	Wall time
`#PBS -o output.log`	`#SBATCH --output=output.log`	Output file
`#PBS -e error.log`	`#SBATCH --error=error.log`	Error file
`#PBS -M email`	`#SBATCH --mail-user=email`	Email
`#PBS -m abe`	`#SBATCH --mail-type=BEGIN,END,FAIL`	Mail events
`#PBS -V`	`#SBATCH --export=ALL`	Export environment
`#PBS -t 1-10`	`#SBATCH --array=1-10`	Job array

Environment Variables

PBS/Torque	Slurm	Description
`$PBS_JOBID`	`$SLURM_JOB_ID`	Job ID
`$PBS_JOBNAME`	`$SLURM_JOB_NAME`	Job name
`$PBS_NODEFILE`	`$SLURM_JOB_NODELIST`	Node list
`$PBS_ARRAYID`	`$SLURM_ARRAY_TASK_ID`	Array task ID
`$PBS_NP`	`$SLURM_NTASKS`	Number of tasks
`$PBS_O_WORKDIR`	`$SLURM_SUBMIT_DIR`	Submission directory

Key Differences

- Slurm requires `--account` and `--qos` — run `check_my_partitions` to find yours
- Slurm uses `--cpus-per-task` instead of `ppn`
- Memory in Slurm is per-CPU (`--mem-per-cpu`) or total (`--mem`)
- Slurm does not automatically change to the submission directory — add `cd $SLURM_SUBMIT_DIR` to your script if needed

---

Created 2026-06-14 08:41:00 UTC by levk
Updated 2026-09-07 15:16:46 UTC by levk