

PBS equivalent or similar	SLURM directive	Purpose	Example
<code>-l walltime=hhh:mm:ss</code>	<code>--time=hhh:mm:ss</code>	Requests the maximum amount of time that the job can be running. The job will be aborted when this time is reached. If no time is specified, a default value will be assigned (currently 20min).	<code>#SBATCH --time=15:20:00</code> requests 15 hours and 20 minutes
<code>-l nodes=N:ppn=n</code>	<code>--nodes=N</code> <code>--tasks-per-node=n</code>	Requests the number of nodes and the number of tasks/processes per node that will be allocated to the job. By default SLURM will allocate 1 core per task, unless <code>--cpus-per-task</code> is specified.	<code>#SBATCH --nodes=10</code> <code>#SBATCH --tasks-per-node=4</code> requests 10 nodes with 4 cores each.
N/A	<code>--nodes=N</code> <code>--tasks-per-node=n</code> <code>--cpus-per-task=c</code>	Requests the number of nodes, the number of tasks/processes per node, and the number of cores per task that will be allocated to the job. This job will be allowed to start <code>n</code> processes on each node, but each process will be able to run on <code>c</code> different cores (with <code>c</code> threads for example).	<code>#SBATCH --nodes=10</code> <code>#SBATCH --tasks-per-node=4</code> <code>#SBATCH --cpus-per-task=2</code> requests 10 nodes with 4 tasks/processes each. Each process could use 2 cores. (this job will be allocated with 80 cores in total)
<code>nodes=N:ppn=n:gpus=g</code>	<code>--nodes=N</code> <code>--tasks-per-node=n</code> <code>--gres=gpu:g</code>	Requests the number of nodes, the number of cores per node, and the number of GPUs per node that will be allocated to the job.	<code>#SBATCH --nodes=2</code> <code>#SBATCH --gres=gpu:1</code> <code>#SBATCH --tasks-per-node=8</code> requests 2 nodes with 1 GPU and 8 cores each
<code>-l procs=X</code>	<code>--ntasks=X</code>	Requests the number of tasks/processes that will be allocated to the job in an unspecified number of nodes (SLURM could allocate the cores in any number of nodes from 1 to X so that the job starts as soon as possible.	<code>#SBATCH --ntasks=100</code> requests 100 cores that will be distributed across the cluster

<code>-l pmem=m</code>	<code>--mem-per-cpu=m</code>	Requests the maximum amount of memory in megabytes that will be assigned to any core in the job. m must be an integer. It can also be requested in gigabytes adding G at the end. Default value is currently 1940MB.	<code>#SBATCH --mem-per-cpu=4G</code> requests 4 gigabytes per core
<code>-l mem=M</code>	N/A		
N/A	<code>--mem=M</code>	Requests the maximum amount of memory in megabytes that must be allocated per node. M must be an integer. It can also be requested in gigabytes adding G at the end.	<code>#SBATCH --mem=18G</code> requests 18 gigabytes of memory on each node allocated to the job
<code>-N <jobname></code>	<code>--job-name=jobname</code>	Sets the name that your job will have for the scheduler system. This name will be showed, for example, when using qstat or squeue. Default name of the job is the name of the submission script.	<code>#SBATCH --job-name=exp1</code> sets exp1 as the name of the job
<code>-o <outputname></code>	<code>--output=outputname</code>	Specifies the name of the file where SLURM will direct the standard output and standard error of the job. Default name of output log file is <code>slurm-<jobID>.out</code>	<code>#SBATCH --output=results1</code> sets results1 as the name of the output log file
<code>-j oe</code>	N/A	In SLURM, the standard output and error are joined in the same file by default. This can be modified if <code>--error</code> is specified	
N/A	<code>--error=errorname</code>	Specifies the name of the file where SLURM will direct the error output of the job. Default name is <code>slurm-<jobID>.out</code>	<code>#SBATCH --error=errors1</code> sets errors1 as the name of the error log file

<code>PBS_O_WORKDIR</code>	<code>SLURM_SUBMIT_DIR</code>	This environment variable contains the directory from which the job was submitted. By default, a job will start running on this directory, so there is no need anymore to <code>cd \$SLURM_SUBMIT_DIR</code> (or <code>cd \$PBS_O_WORKDIR</code>)	
<code>PBS_JOBID</code>	<code>SLURM_JOB_ID</code>	This environment variable contains the ID of the job allocation	
<code>qsub <my_script></code>	<code>sbatch <my_script></code>	Submits the job described by the batch script <code>my_script</code>	<code>sbatch experiment1</code> submits script <code>experiment1</code>
<code>mpirun <my_mpi_code></code>	<code>srun <my_mpi_code></code>	<code>srun</code> is the native launcher of parallel processes in slurm. We strongly recommend to use <code>srun</code> instead of <code>mpirun</code> or <code>mpiexec</code>. Note that there is no need to specify the number of processes, this info will be obtained from the job request.	<code>#SBATCH --nstacks=25</code> <code>#SBATCH --cpus-per-task=3</code> <code>srun vasp</code> will start 25 vasp processes across the cluster, and allows each process to fork into 3 parallel threads
<code>pbsdsh -v <my_script></code> <code>pbsdsh2 -v <my_script></code>	<code>srun <my_script></code>	<code>srun</code> can also run multiple copies of a script or binary in parallel. Use the environment variable <code>SLURM_PROCID</code> (the equivalent of <code>PBS_VNODENUM</code>) to know the number of process assigned to each copy	<code>#SBATCH --nodes=2</code> <code>#SBATCH --tasks-per-node=8</code> <code>srun exp1</code> will run 16 copies of the script <code>exp1</code> (8 instances on each node). The variable <code>SLURM_PROCID</code> will take values 0,1,2...15
<code>qstat</code>	<code>squeue</code>	Show a list of all jobs in the queues. With <code>squeue</code> the jobs are ordered by priority, not by order of submission as with <code>qstat</code> .	<code>squeue</code>

<code>qstat -f <jobID></code>	<code>scontrol show job <jobID></code> <code>scontrol show job <jobID> -dd</code>	List detailed information of the job with identification number jobID	<code>scontrol show job 321</code> obtains detailed information of job 321
<code>qstat -u <nsid></code>	<code>squeue -u <nsid></code>	List all jobs of user nsid	<code>squeue -u abc123</code> lists all jobs of user abc123 in the queues
<code>showbf</code> <code>showq</code>	N/A	There are no direct equivalence to these commands in SLURM, but rather different options to obtain the same information. Some of these options are described below	
	<code>sprio</code>	Shows the priority of queued jobs and the different factors that compose it.	<code>sprio</code>
	<code>sshare</code>	Shows the usage of an user and its group. It also shows the fairshare value that will be used to calculate the fairshare component of the priority of jobs submitted by the user.	<code>sshare</code>
	<code>sinfo --states=idle</code>	Shows a list of all idle nodes per queue	<code>sinfo --states=idle</code>
	<code>squeue --start</code>	Shows the list of all jobs queued with its estimated start times	<code>squeue --start -u abc123</code> lists all queued jobs of user abc123 with its estimated start times
	<code>sbatch --test-only <my_script></code>	Shows the estimated start time of the job described in my_script without submitting it.	<code>sbatch --test-only experiment1</code> determines the estimated start time for the job described in experiment1 as if it were submitted now. The job is never submitted.