

Basic SLURM

SLURM supports a variety of job submission techniques. By accurately requesting the resources you need, you'll be able to get your work done. This describes how to run both batch and interactive jobs.

Batch Job Submission

Batch jobs involve submitting a file with that contains both your resource requirements and a script of what to run. You can request nodes, cores, memory and GPUs with slurm for a set amount of time. Here is a basic sbatch submit script:

Note: All of the lines that begin with #SBATCH are ready by SLURM for job settings.

Example 1: job running on a single node

```
#!/bin/bash # set the partition where the job will run
#SBATCH --partition=defq
# set the number of nodes
#SBATCH --nodes=1
# set the number of GPU cards to use per node
#SBATCH --gres=gpu:1
#SBATCH --time=20:00:00
# set name of job #SBATCH --job-name=test123
# run the application
srun -n$SLURM_NTASKS <app name>
```

Copy that text into a file and call it imb4.sample.sbatch. To submit it, just run:

```
sbatch imb4.sample.sbatch
```

To see the status of all your jobs, run:

```
squeue
```

Some other useful **squeue** features include:

```
-u for showing the status of all the jobs of a particular user,
e.g. squeue -u bob for user bob;
-l for showing more of the available information;
--start to report the expected start time of pending jobs.
```

To see how much memory your job has used so far, run:

```
sstat --format JobID,NTasks,nodelist,MaxRSS,MaxVMSize,AveRSS,AveVMSize
$JOBNUMBER
```

where \$JOBNUMBER is the number of the job you just submitted. It should return something like this:

```
sstat --format JobID,NTasks,nodelist,MaxRSS,MaxVMSize,AveRSS,AveVMSize
3504          JobID   NTasks          Nodelist        MaxRSS   MaxVMSize
AveRSS   AveVMSize  -----
- -----
12,sh- 2-[1-2]      12735100K  16114480K  12719697K  16110313K
          3504.0          3   sh-1-
```

If your jobs have completed, you can run sacct to job statistics:

```
sacct -format JobID,jobname,NTasks,nodelist,MaxRSS,MaxVMSize,AveRSS,AveVMSize
```

Submitting interactive jobs

To test (or debug) an application using a GPU card, an interactive job can be obtained by using one of the following methods.

The first method is exemplified with the command:

```
salloc --partition=defq --nodes=1
```

This will make a request for 1 node on the default queue.

```
salloc --partition=defq --nodes=1 --gres=gpu
```

that makes a specific request for a GPU card on the **defq** partition and, once the allocation is made, directly lands the user on a node from that partition with a free GPU card. You should adjust the number of gpu cards, as well as the number of nodes and CPUs to your needs.

Finally here is the link for all you want to know about SLURM.

<https://computing.llnl.gov/linux/slurm/documentation.html>