

Best Practices for Jobs



NERSC New User Training
Sept 08, 2023

Justin Cook
User Engagement Group

CPU vs GPU

- Account (--account / -A)
 - -A <account> for CPU
 - -A <account_g> for GPU
- Constraints (--constraint / -C)
 - -C cpu
 - -C gpu
- Modules
 - Load `cpu`
 - Load `gpu` (default, CUDA aware MPI)

Always specify your account

Ensure the appropriate allocation is charged

- Default is set in iris, not always the one you want

```
#SBATCH --account=<NERSC Account>
```

Do not run production jobs in your home directory

Run in \$SCRATCH

- Large, temporary storage
- Optimized for reads and writes

Install Software to Global Common

- /global/common/software
- Optimized for reads
- Mounted read-only to compute nodes via DVS (more later)

Set job time limits

Due to backfill scheduling, short and variable-length jobs generally start quickly resulting in much better job throughput

```
#SBATCH --time-min=<lower_bound>
```

```
#SBATCH --time=<upper_bound>
```

Set job filesystem licenses

A batch job will not start if any of the specified file systems are unavailable due to maintenance or an outage

```
#SBATCH --licenses=scratch,cfs
```

Specify the logical core per task

Each CPU core has two hardware threads

A CPU-only compute node has a total of 128 physical cores , or 256 logical CPUs total.

The GPU compute node has a total of 64 physical cores, or 128 logical CPUs total.

Why double? CPU-only compute nodes have two processors, GPU compute nodes have one.

More: <https://docs.nersc.gov/jobs/affinity/#process-and-thread-affinity>

Specify the logical core per task (CPU)

Logical cores per task (-c, --cpus-per-task)

- CPU: $2 \times \lfloor 128 / \text{tasks per node} \rfloor$
- GPU: $2 \times \lfloor 64 / \text{tasks per node} \rfloor$

Example: 5 MPI tasks per GPU node

- $2 \times \lfloor 64 / 5 \rfloor = 2 \times 12 = -c 24$

What is it for a CPU node?

- -c 50

And there's more

<https://docs.nersc.gov/jobs/best-practices/>

- Some are more advanced than others
- Some require knowing more about your code
- Some require knowing more about the system

<https://www.nersc.gov/users/training/events/2023/migrating-from-cori-to-perlmutter-training-march2023/>

Thank you and
welcome to
NERSC!



Justin Cook
User Engagement
September 8, 2023