

ACE3P on Carver@NERSC Cheat Sheet, Feb 24 2011

Carver is a supercomputer at NERSC (National Energy Research Center) in Berkeley, CA. It has a scheduler to run computational jobs based on a queue. The user has to use a job submission script to run a job. Once a job is submitted, the scheduler will decide when and where the job will be executed.

- You first need to get a NERSC account
(send your information to ace3p@slac.stanford.edu :
First and Last names, nationality, phone, address, and email)
- General overview at <http://www.nersc.gov/nusers/systems/carver/>
- Or go directly to
http://www.nersc.gov/nusers/systems/carver/running_jobs/
- Transfer mesh and input file(s) from your desktop to your directory on Carver's high-performance file system:

```
scp -pr <problem_dir>  
carver.nersc.gov:/global/scratch/sd/<user>/
```

(Note, that since we want access to the high-performance file system, we need to run our ACE3P simulations from within the /scratch directory)

- Login on Carver:
`ssh carver.nersc.gov`
- Go to the above directory:
`cd /global/scratch/sd/<user>/<problem_dir>/`
- Note that you can find the ACE3P executables for Carver in the directory
`~candel/.community/carver/`

How to run a Job on Carver (Similar to Franklin, but not quite the same) http://www.nersc.gov/nusers/systems/carver/running_jobs/

- The following is a sample job submission script. You can find it at
~candel/.community/carver/run.pbs

```
#PBS -S /bin/bash
#PBS -N ACE3P
#PBS -q regular          <- which queue to use (magellan is best)
#PBS -A m349             <- which account to charge (this is for magellan)
#PBS -l nodes=1:ppn=4    <- requested nodes and CPUs per node (ppn=1,2,4,8)
#PBS -l pvmem=5GB        <- memory per MPI process: set to 20GB*1/ppn
#PBS -l walltime=00:30:00 <- requested time (don't need to use it all)
#PBS -j eo               <- combine standard output + error
#PBS -V                 <- transfer environment variables
cd $PBS_O_WORKDIR       <- go to the right place on compute nodes
module unload pgi openmpi <- unload default PGI environment
module load gcc/4.4.2 openmpi-gnu <- load GNU environment

mpirun -np 4 ~candel/.community/carver/omega3p input.o3p
```

Carver has 8 CPUs and 24 or 48 GB of memory per compute node. For large problems, you can force using only the nodes with 48 GB of memory with the “bigmem” and qualifier, see also

http://www.nersc.gov/nusers/systems/carver/running_jobs/memory.php

```
#PBS -l nodes=4:ppn=2:bigmem
#PBS -l pvmem=22GB      ### use 44GB*1/ppn for bigmem nodes
```

Important:

#PBS -l nodes=<NNodes> says how many nodes (and thus memory!) you reserve

Then, for a given number of nodes, specify the number of processes to start:

```
mpirun -np <NNodes>*8  starts 8 processes per node, each has around 3 / 6 GB
mpirun -np <NNodes>*4  starts 4 processes per node, each has around 6 / 12 GB
mpirun -np <NNodes>*2  starts 2 processes per node, each has around 12 / 24 GB
mpirun -np <NNodes>    starts 1 process per node, it has all the 24 / 48 GB for itself!
```

For Omega3P/S3P and large problems, it may be better to start fewer than all the <NNodes>*8 processes, as memory scalability is hard and the O/S and support

libraries also occupy memory. For all other codes, it is usually OK to use all CPUs. Remember to use Solver: CG for T3P, to reduce memory requirements.

- Submit the job

```
qsub run.pbs
```

- Look at the status of the batch jobs, and get the Job ID

```
qstat -u <username>
```

- Watch the job's printout as it runs

```
tail -f <Job ID>.cvrsvc09-ib.ER
```

- Cancel the job

```
qdel <Job ID>    (e.g., qdel 228409)
```

- Change the request time (before the job starts)

```
qalter -l walltime=HH:MM:SS <Job ID>
```

- Copy the results back to your computer (from your computer)

```
scp -pr carver.nersc.gov:/<directory> .
```

- Clean up on Carver:

1) Move the results back into your project directory

```
mv <directory> /project/projectdirs/gc2/<user>/
```

or, even better

2) upload it into HPSS (high performance storage system) for archival

please consult <http://www.nersc.gov/nusers/systems/hpss/>