

Using Moab on BlueGene/L

General information on running jobs on BlueGene/L (BGL) can be found online at:

https://asc.llnl.gov/computing_resources/bluegenel/basics/

Moab and SLURM 1.3 have been installed on the BlueGene machines. There are only a few differences in submitting jobs to Moab to run on BGL and bgldev compared to the other LC machines.

1. Use the `msub -l nodes=<value>` option to specify the number of BlueGene compute nodes (c-nodes).
Example: `msub -l nodes=1024 ...`
2. In the nodes specification, “k” is recognized as shorthand for 1024.
Example: `msub -l nodes=2k`
3. The job’s requested nodes must exactly match an available block in the (default or specified) SLURM partition. To see block and partition information, invoke `smap -Db -c`. Use the `msub -q <partition>` option to target your job to a specific SLURM partition.
4. On bgldev, you can request a portion of the mid-plane. You should specify a node count supported by the hardware (e.g., 128, 512, or 1024 for bgldev).
5. Use the new `msub --slurm` option to pass BGL-specific options to SLURM. The `--slurm` option **must** be the last `msub` option on the command line. All options after that are passed directly to SLURM’s `sbatch` command.
Example: `msub <program> -l nodes=512 --slurm --no-rotate`
6. Where a pdebug node partition is available on a BGL machine, we have configured the pdebug partition to disallow job submissions directly to SLURM using the `sbatch` command. This allows us to schedule the pdebug nodes based on fair-share instead of simple FIFO. This breaks with the policy in place on other LC machines with pdebug partitions.
7. The `psub` wrapper exists on BGL and bgldev, but its use is deprecated. The following options are not supported:
 - a. The `psub -ln k` shorthand notation (to represent 1024 compute nodes) is not supported. Either specify the node count in numeric form only or use the `msub` command described in item 2 above.
 - b. The `psub -bgl <attributes>` option is not supported. See item 5 above for the recommended alternative.