



Flux (and a bit of Slurm)

LC User Meeting
March 2025

Ryan Day, LC Resource Management

Prepared by LLNL under Contract DE-AC52-07NA27344.



Documentation

- <https://hpc.llnl.gov/banks-allocations-and-jobs>
- <https://hpc-tutorials.llnl.gov/flux/>
- <https://flux-framework.readthedocs.io/en/latest/>
- <https://github.com/LLNL/mpibind/blob/master/tutorials/flux/module1.md>
- <https://hpc.llnl.gov/documentation/tutorials/using-lc-s-sierra-systems>
- `man flux-COMMAND` , `flux -h` , `flux COMMAND -h`



--exclusive and mpibind

- `--exclusive` is weird in Slurm.
 - `salloc | sbatch | srun (outside of an allocation)` : all of the resources.
 - `srun (inside of an allocation)` : minimal resource set.
 - `srun -N1 -n4 --exclusive -p pdebug my_program` : probably won't do what you want.
 - `salloc -N1 --exclusive -p pdebug srun -n4 my_program` : for 4 tasks that see the whole node.
- Flux defaults to giving a minimal resource set.
 - `--exclusive` : always gives you all of the resources on the node(s).
 - Mpibind works best with all of the resources.
 - `flux run -N1 -n4 my_program` : will only see one GPU on elcap.
 - `flux run -N1 -n4 --exclusive my_program` : will see all 4 GPUs.



Other Flux notes for ElCap and Tuolumne

- ` -b standby ` : jobs automatically killed if doing so would allow highest priority job to run.
- ` flux update JOBID duration=... ` : users can shorten their jobs. LC Staff can extend them.
- ` flux urgency JOBID ... ` : users can hold/release their jobs. LC Staff can expedite them.
- LC Staff can create limited access queues for DATs.
- SPINDLE is on by default. ` -o spindle.level=off ` or ` -o fastload ` to turn it off.
- Mpibind will try to avoid noise reduction / system cores unless you ` unset MPIBIND_RESTRICT ` .
- If using ` --setattr=gpu_mode=... `, also set ` --conf=resource.rediscover=true` .



Documentation

- <https://hpc.llnl.gov/banks-allocations-and-jobs>
- <https://hpc-tutorials.llnl.gov/flux/>
- <https://flux-framework.readthedocs.io/en/latest/>
- <https://github.com/LLNL/mpibind/blob/master/tutorials/flux/module1.md>
- <https://hpc.llnl.gov/documentation/tutorials/using-lc-s-sierra-systems>
- `man flux-COMMAND` , `flux -h` , `flux COMMAND -h`

