


NUG Monthly Meeting



18 February, 2021

Today's plan



- Interactive - please participate!
 - Raise hand or just speak up
 - [NERSC User Slack](#) (link in chat), **#webinars** channel

- Agenda:
 - Win-of-the-month
 - Today-I-learned
 - Announcements/CFPs
 - Topic of the day: **Making the most of Slurm at NERSC**
 - Coming meetings: topic suggestions/requests?
 - Last month's numbers

Win of the month



Show off an achievement, or shout out someone else's achievement, e.g.:

- Had a paper accepted
- Solved a bug
- A scientific achievement (maybe candidate for Science highlight, or **High Impact Scientific Achievement award**)
- An **Innovative Use of High Performance Computing** (also a candidate for an award) (<https://www.nersc.gov/science/nersc-hpc-achievement-awards/>)

Note that early-career nominations for these awards are due this week! Please let us know of award-worthy work from you or your colleagues

Tell us what you did, and what was the key insight?

Today I learned



What surprised you that might benefit other users to hear about?
(and might help NERSC identify documentation improvements!)

Eg:

- Something you got stuck on, hit a dead end, or turned out to be wrong about
 - Give others the benefit of your experience!
 - Opportunity to improve NERSC documentation
- A tip for using NERSC
- Something you learned that might benefit other NERSC users

"If we knew what it was we were doing, it would not be called research, would it?" -
Einstein

NERSC Early Career HPC Achievement Awards Nominations Due Tomorrow!

- **High Impact Scientific Achievement:** recognizing work that has or is expected to have an exceptional impact on scientific understanding, engineering design for scientific facilities, or a broad societal problem.
- **Innovative Use of High-Performance Computing:** recognizing researchers who have used NERSC's resources in innovative ways to solve a significant problem or have provided a new methodology with the potential to have a large scientific impact. Examples might include application of HPC to a new scientific field or combining computing, data, networking, and edge services to do something entirely new in a domain where HPC is already established.







Eligibility: Research that used NERSC resources during allocation years 2019 and/or 2020 by any NERSC user who received their degree during or after 2015

Announcements and CFPs



Others?

Topic of the day: Making the most of Slurm at NERSC



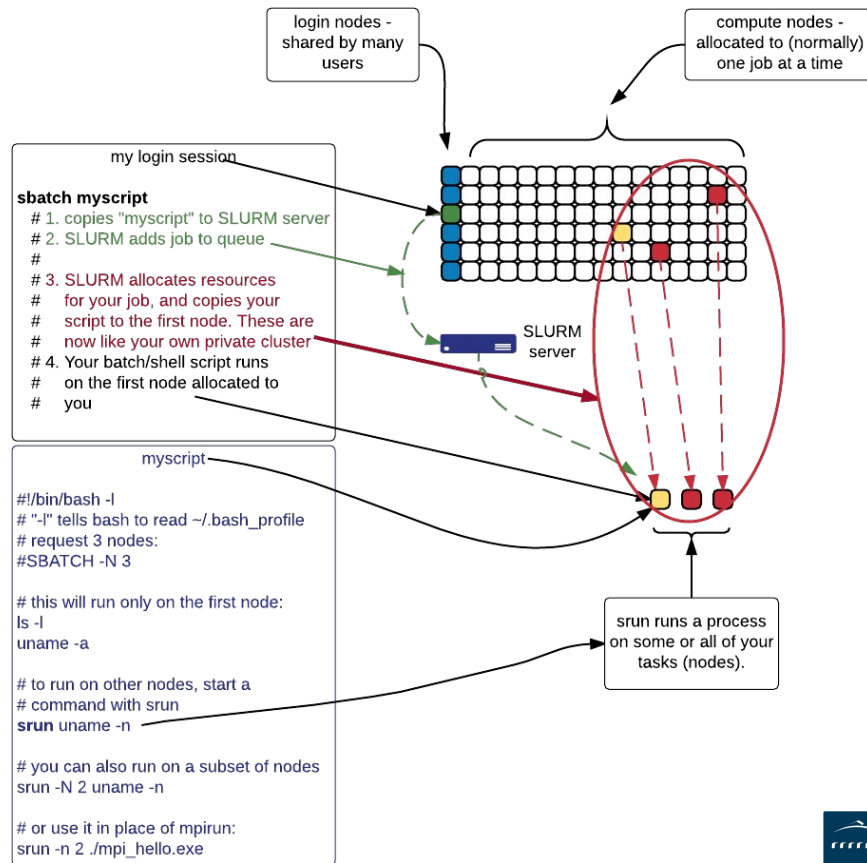
Shahzeb Siddiqui
NERSC User Engagement Group

Refresher: How jobs work



- **Start on login node**
 - shared by many users
 - **not for computational work**
- **Access compute nodes with `sbatch` or `salloc`**
 - A number of nodes are allocated *exclusively** to your job for a period of time
- **Your batch script:**
 - Gets copied to queue (so editing now has no effect)
 - Has directives for SLURM, and shell commands to perform on first compute node
- **Access your other allocated nodes with `srun`**
- **`stdout`, `stderr` will be saved to file**
 - (when running in batch mode)

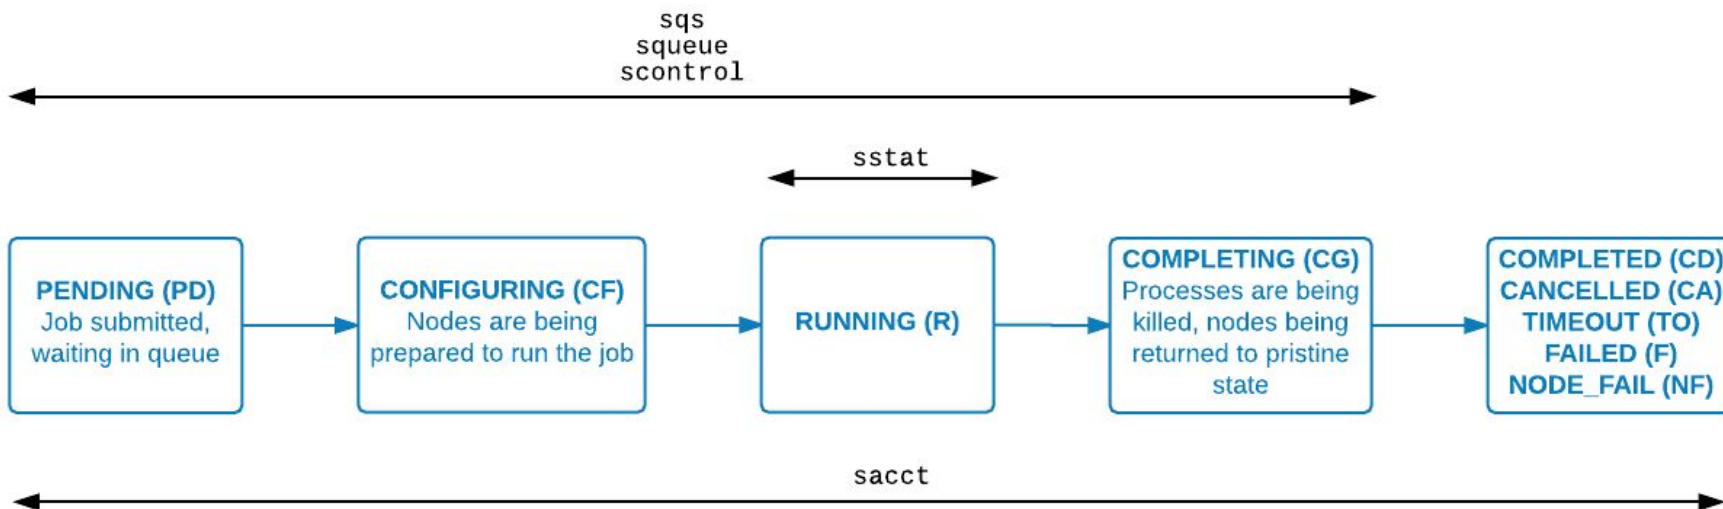
* normally exclusively. But there are eg "shared" queues (you get part of a node)



Watching your job



Job states you might see:



How scheduling works




Backfill pass (first part of queue)

- Slurm will look for first available slot and schedule job there

Instant pass (rest of queue)

- "Can this job be started *right now* without delaying a scheduled job?"

Outcome: mostly first-come-first-served, but short (and small) jobs can often jump the queue



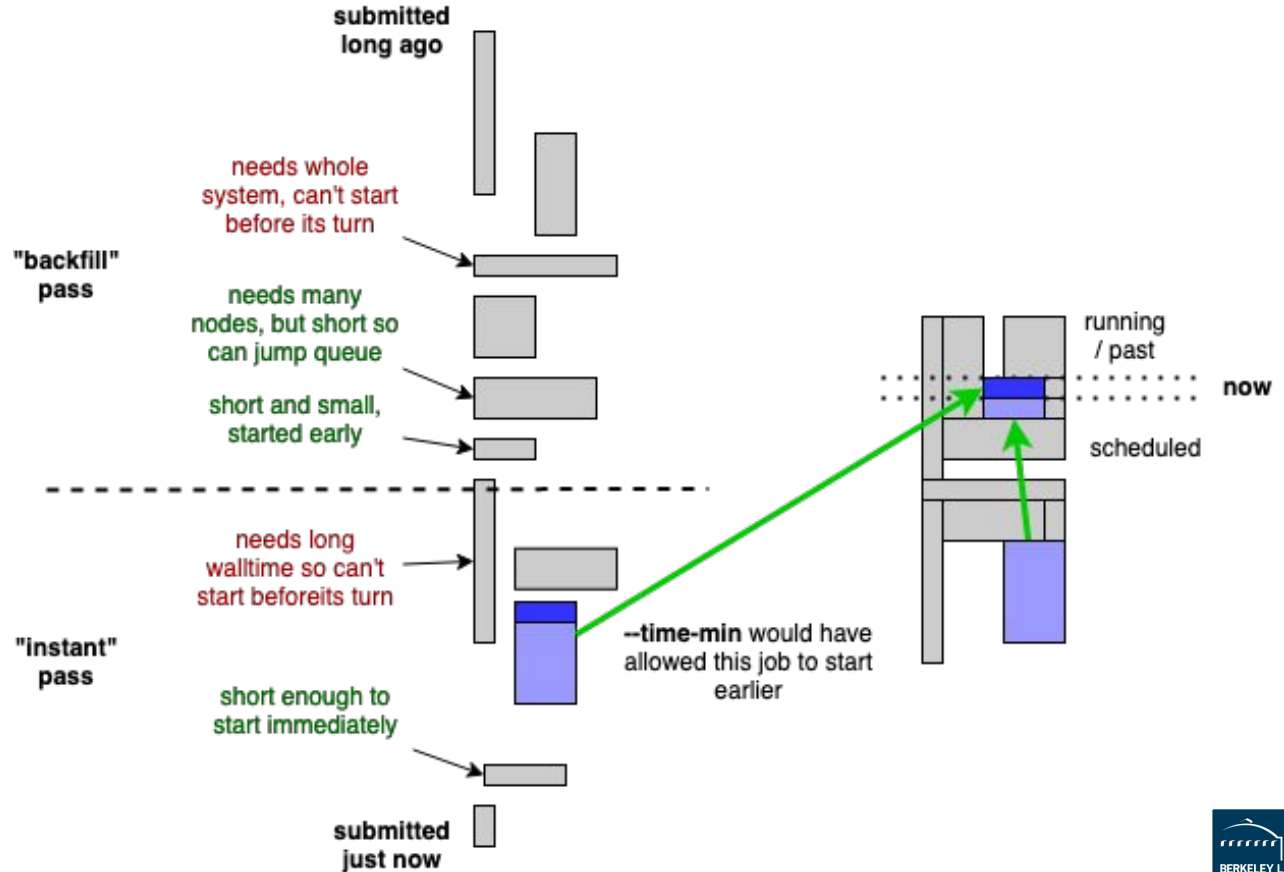
Use --time-min to adjust job to fit gap



You **don't** need to guess the best wallclock time!

(if your job is amenable to stop-and-continue)

Flex queue adds 75% discount, but you can do this in regular queue too, and still get reduces queue wait



Available queues



Queues	Cost	Typical time in queue	Workload
-q regular	\$\$\$	Hours to days	For most work
-q premium	\$\$\$\$\$	Minutes to hours	Limited access, emergency-only
-q debug	\$\$\$	Minutes to hours	Small, short jobs
-q interactive	\$\$\$	5-minute or abort	Interactive debugging (use salloc)
-q low	\$	days to weeks	Low-urgency work
-q flex	\$	fills gaps	Workloads that can handle flexible walltimes. Requires 2 hrs min time (--time-min) and only available for KNL partition. Ideal for jobs that support checkpoint/restart
-q shared	\$\$\$	hours	Serial/small tasks. Node is shared with multiple users. Packing several into a regular job is often better. Haswell only .
-q overrun	-	days to weeks	Out-of-time projects can run jobs in unfilled schedule gaps
-q realtime_*		near-immediate	Special purpose only

Available queues



Most jobs should use -q regular

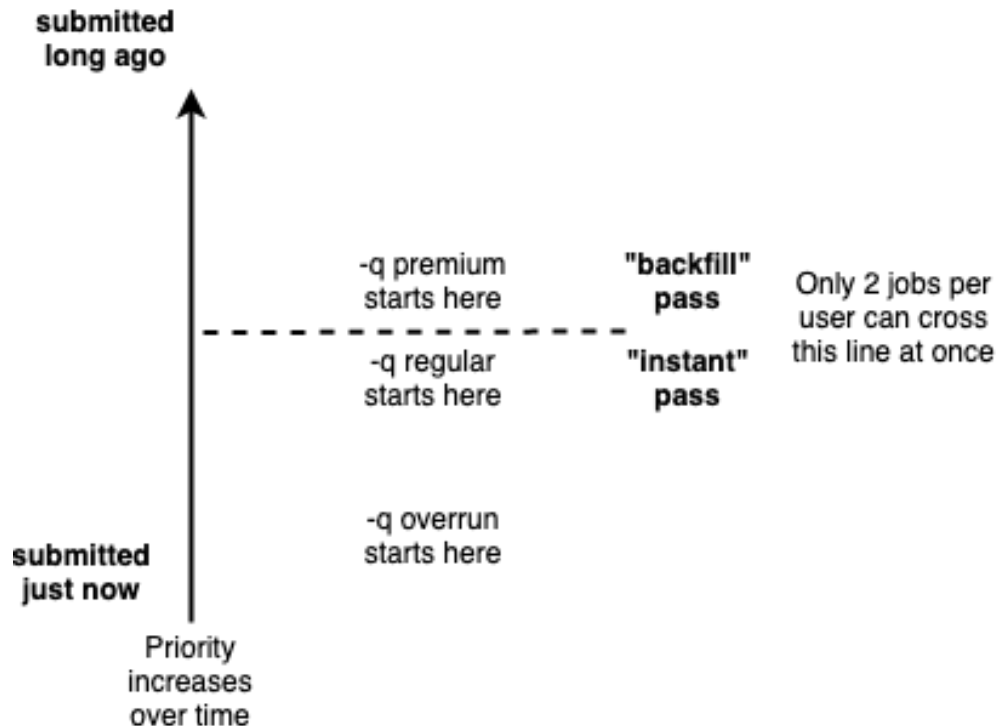
-q premium is for scientific emergencies ("I need to re-run this simulation in response to peer review recommendation, before this deadline")

Flex queue: **75% discount** (KNL-only)

- Job must have --time-min no more than 2 hours
- (Aim is to improve Cori utilization by filling schedule gaps)

Large jobs: **50% discount** (KNL-only)

- Regular queue only
- 1024+ nodes, KNL only



Summary of tips



- Queues are long!
- Jobs that can fit in schedule gaps will start sooner!
 - Short-timelimit gaps are common
 - Long-timelimit gaps are rare
- --min-time gives you the longest walltime that will fit in the soonest-available gap
 - And using this in -q flex gives a 75% discount
- Large KNL jobs get a 50% discount
- -q overrun allows out-of-time project to continue work

Best practices:

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

Coming up



Topic requests/suggestions?

- Eg:


We'd love to hear some lightning talks **from NERSC users** about the research you use NERSC for!

Last month's numbers - January



Scheduled and overall availability:	Scheduled	Overall
Cori	99.8%	94.3%
HPSS	100%	100%
CFS	100%	100%

Cori:



Last month's numbers - December



Cori Utilization: 95.3%

Large jobs: 41.3%

New Tickets: 673

Closed Tickets: 544

Backlog at 1 Feb: 620



Thank You



U.S. DEPARTMENT OF
ENERGY

Office of
Science

