NUG Monthly Meeting

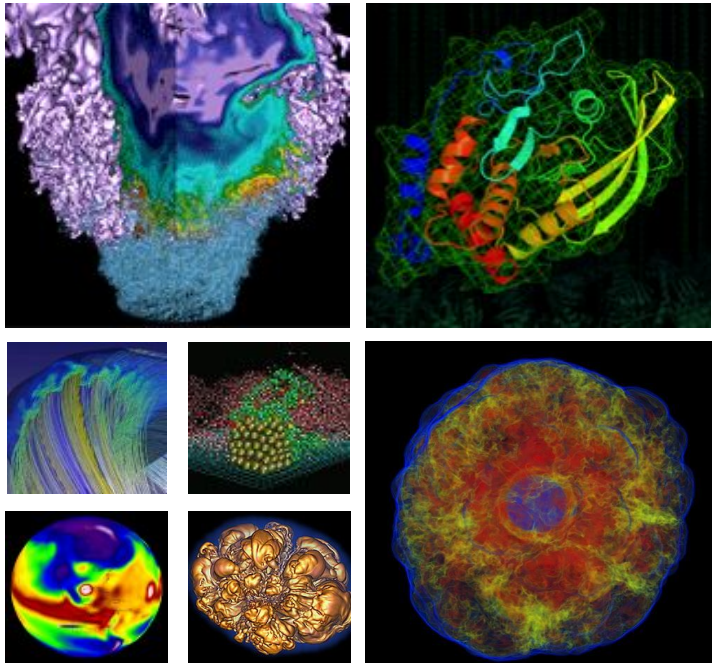

NERSC

16 February, 2023

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

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
  - User Community Survey
  - Announcements/CFPs
  - Topic of the day: **Cori Retirement**. Rebecca Hartman-Baker will provide an overview of NERSC's plans for the retirement of Cori.
  - Coming meetings: topic suggestions/requests?

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/)

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

User Community of Practice

Creating a NERSC User Community of Practice!

A **Community of Practice** requires

- A shared domain of interest
- An actively cultivated & maintained sense of community
- Active practice of the shared domain of interest

Next Step: Community Feedback!

- Help us build a stronger, **more active** community!
- Participate in focus groups to **share your ideas about:**
  - What do you want the NERSC community to look like?
  - What is missing from your NERSC experience?
  - What type of training, programs, and/or events would help you use NERSC resources for your science?

Focus Groups:

- Small (8-10 people) groups to discuss NERSC User Community events and programs
- Held virtually on Zoom for 1 hour max
- Opportunities to discuss ideas with NERSC staff directly!

To make this happen, we need you!

Please take the next 5-7 minutes to complete our survey!*

* you don't have to participate in the focus groups, you can **just** provide us some useful information!!

Announcements and CFPs

See weekly email for these upcoming events:

- Attention Students: NERSC Summer Internships Available!
  - https://www.nersc.gov/research-and-development/internships/
- CfP for AY23 Research in Quantum Information Science on Perlmutter Now Open
  - https://www.nersc.gov/research-and-development/quantum-information-science/quantum-information-science-perlmutter/
- IDEAS-ECP Webinar on "Our Road to Exascale: Particle Accelerator & Laser-Plasma Modeling" March 15
  - http://ideas-productivity.org/events/hpc-best-practices-webinars/
- ECP HPC Workforce Seminar on "Strategies for Inclusive Mentorship" on March 16
  - https://www.exascaleproject.org/event/inclusive-mentorship/
- NERSC is hiring! Several open positions, see https://lbl.referrals.selectminds.com/page/nersc-careers-85

Announcements and CFPs

Migrating from Cori to Perlmutter:

- Training session March 10
  - https://www.nersc.gov/users/training/events/migrating-from-cori-to-perlmutter-training-mar2023/
- Office hours:
  - Thursday, February 23
  - Tuesday, March 7
  - Wednesday, March 15
  - Friday, March 31

Announcements and CFPs

Others?

Cori Retirement

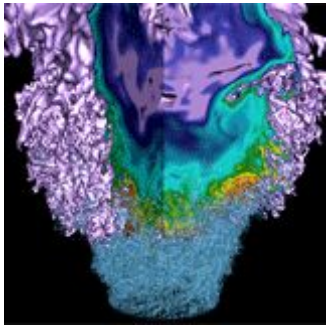

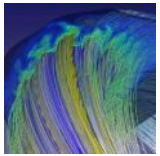
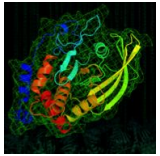
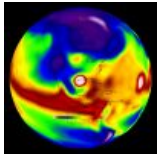
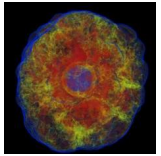
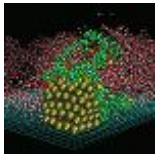

Rebecca Hartman-Baker
NERSC User Engagement

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Cori Retirement

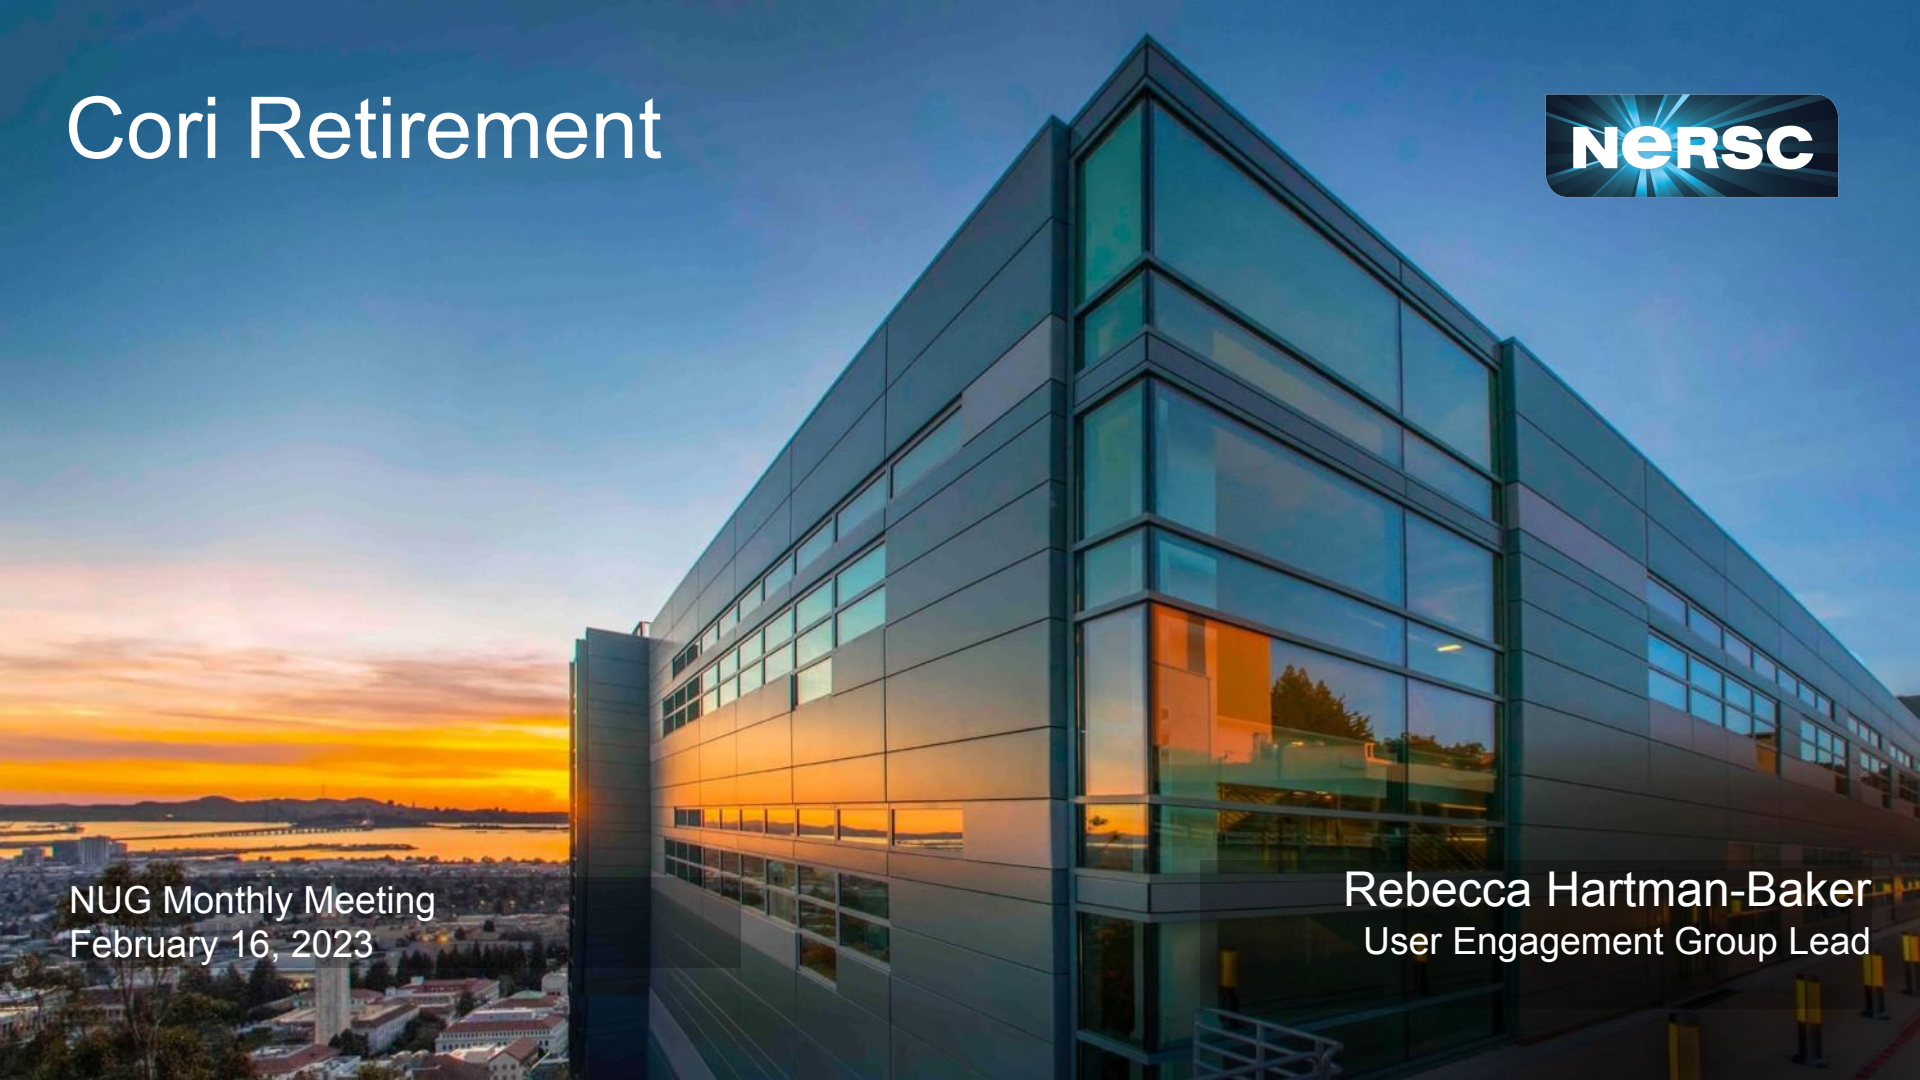

NUG Monthly Meeting
February 16, 2023

Rebecca Hartman-Baker
User Engagement Group Lead

Outline

- Supercomputer lifecycle
- Why retire Cori?
- Cori retirement schedule
- Perlmutter timeline

Supercomputer Lifecycle: Overview

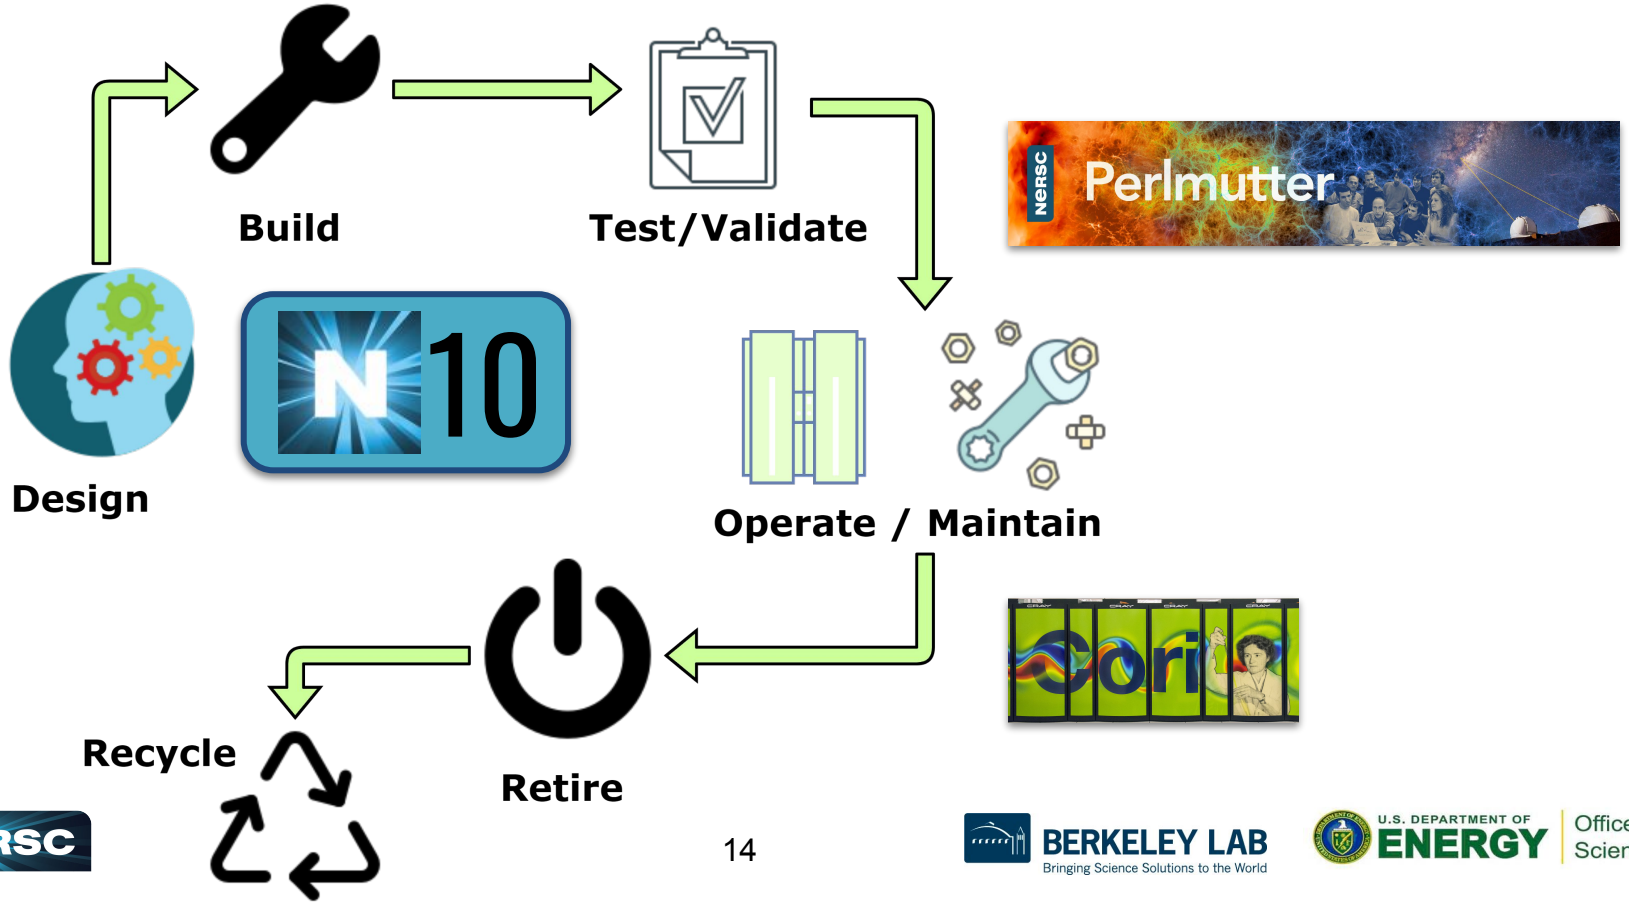

Supercomputer Lifecycle

- Designing the machine is a collaborative process between NERSC & vendor(s)
  - NERSC staff develop requirements for machine, vendors explain what they can provide, & best proposal for price & value is selected
  - Much thought goes into making the best machine with best future technology for users
- Building the machine begins in vendor factory
  - After factory test, machine disassembled & brought to NERSC for reassembly
  - NERSC provides necessary power, water, cooling, etc.

Supercomputer Lifecycle

- Testing begins in factory too
  - Factory test to demonstrate machine is feasible
  - Further testing at NERSC after machine is reassembled
  - Extensive hardware, software, and network testing, which includes letting friendly users on the machine
  - Major tests must be passed before vendor is paid for machine
    - Functionality testing, performance testing, stability testing, reliability testing
    - Includes a 30-day stability test in which machine must remain in service with users on it

Supercomputer Lifecycle

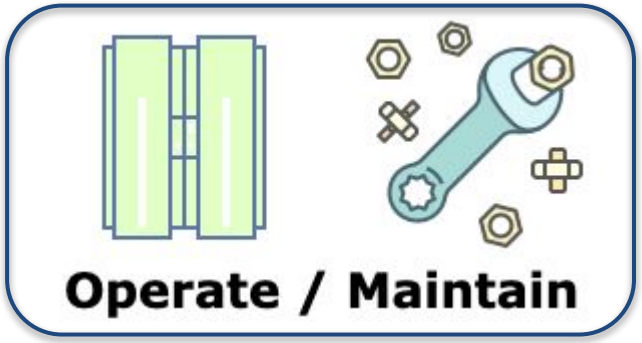

- Operating the machine is a round-the-clock operation
  - At least one NERSC staff member always present onsite
  - Must ensure proper operating conditions
- Regular maintenance is required
  - Vendor staff hired to perform regular physical maintenance & support of the machine, including replacing nodes, network cables, etc.
  - Upgrades to system software to protect from security issues, fix bugs, etc.

Reliability

- A new machine undergoes a “shakeout” period with faulty new hardware that must be discovered & replaced
- Failure rates in first year generally larger than following years
- As machine ages, failure rates rise again
- This is known as the “bathtub curve”

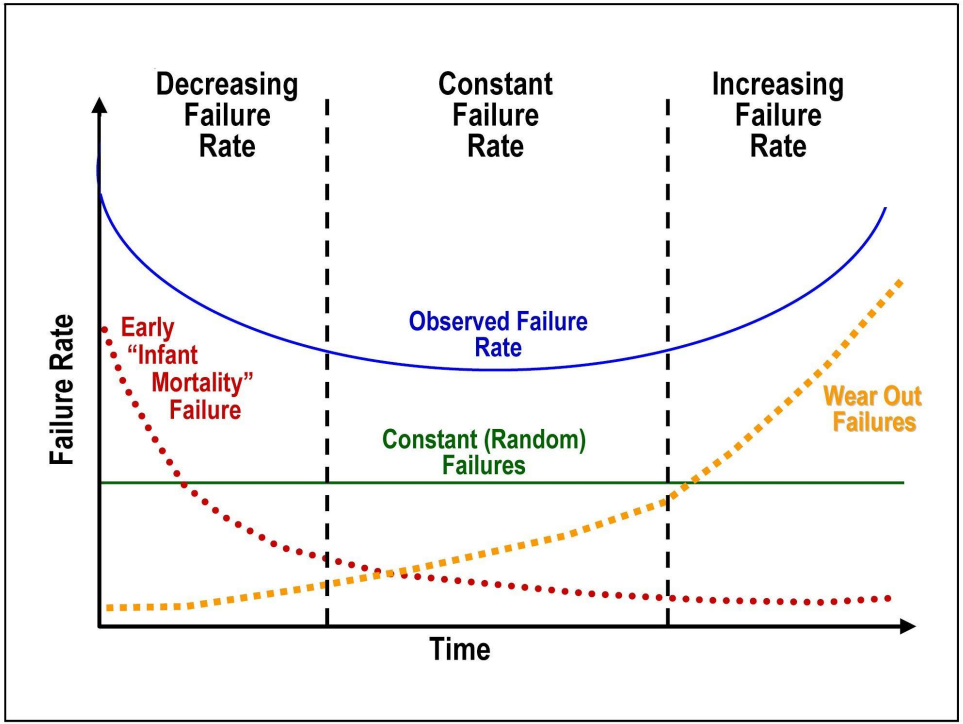

Supercomputer Lifecycle

- NERSC retires machines at the end of their useful life
  - After a certain point, failure rates begin to rise
  - Machine becomes harder to support & less reliable
  - New technologies are more energy efficient and provide more compute power

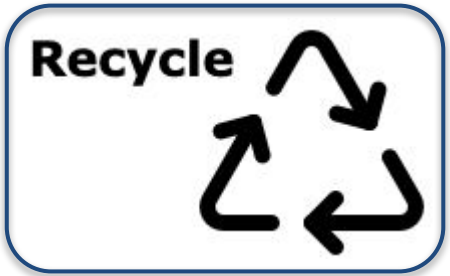

- Machines are returned to vendor to be recycled
  - Some resold, others used for spare parts for similar models still in operation, others have valuable metals or other components removed & recycled

Why Retire Cori?

- Cori has reached the end of its useful lifespan
- Model no longer being produced
  - In addition to processors & memory, cabinet components like fans, electrical parts, & more are no longer made
  - We must rely on remanufactured replacement parts
- Reliability of individual components is going down
  - Observing more frequent failures, especially of electrical components
- Failures are difficult to recover from
  - Of particular concern is the scratch system, for which no spare parts are available
  - Recovery may mean repurposing internally & shrinking the system

Cori Retirement Schedule (Current Plan)

- **March 31:** Auxiliary components will be removed
  - Large memory nodes will be migrated to Perlmutter
  - Cori GPU nodes will be retired
- **End of April:** Cori will be retired (call this date **T**)
  - Access to scratch will continue for 1 week
- **T + 1 week:** Cori powered down
- **T + 1 month:** Cori removal from machine room begins

Perlmutter Timeline

- Perlmutter has not yet completed testing of its final configuration
  - 14 GPU cabinets, 12 CPU cabinets, SS11
  - Final configuration reached in early February 2023
  - Start date unknown, but will be soon
- Cori will not be retired until Perlmutter testing complete
- Components of testing
  - **Functionality:** system provides all required functionality ✅
  - **Performance:** system achieves certain performance level on benchmarks ✅
  - **Stability:** system must remain up during stability testing period
  - **Reliability:** hardware & software failures minimal

Current Perlmutter Status

- Perlmutter is not currently meeting our – or our users' – expectations for reliability
- We understand the importance of a reliable machine to users' scientific progress
- NERSC meets with HPE daily, to address bugs & issues
- HPE experts are at NERSC to focus on resolving the issues users are experiencing
- We are optimistic that this collaboration will improve Perlmutter's reliability

Current Perlmutter Activities

- NERSC & HPE are working together to address
  - Stability of Perlmutter's slingshot network
  - I/O performance on Perlmutter scratch & Community File System
  - Node hardware reliability
- New processes stemming from this collaboration
  - Developed new methodical process for fixing problem nodes before returning to service
  - Configuration changes to CFS & Perlmutter CFS client side to stabilize network communication & performance
  - Rolling out fixes to slingshot network bugs discovered by NERSC

Summary

- In the supercomputer lifecycle, Cori has reached the end
  - No new parts are being manufactured for the machine, making its continued upkeep especially challenging
- We plan to retire Cori at the end of April
- Perlmutter reliability issues are being addressed with high priority by NERSC & our vendor HPE

Coming up

Upcoming topics:

- Suggestions welcome!

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

Nominate a topic at https://forms.gle/WjYx7zV7SAz2CaYz7

Thank You