

Welcome and NERSC Overview




New User Training and Updated
Best Practices on Perlmutter
February 15, 2024

Helen He, Rebecca Hartman-Baker, Charles Lively
User Engagement Group

Agenda

- Introduction to NERSC
- Hardware
- Software
- Interacting with NERSC
- User Responsibilities & Expectations



Introduction to NERSC

About NERSC

- National Energy Research Scientific Computing Center
 - Established 1974, first unclassified supercomputer center
 - Original mission: to enable computational science as complement to magnetically controlled plasma experiment
 - 2024: NERSC 50th anniversary!
- Today's mission: ***Accelerate scientific discovery at the DOE Office of Science through High-Performance Computing and Extreme Data Analysis***
- NERSC is a national user facility
- NERSC is part of Berkeley Lab



NERSC: Mission HPC for DOE Office of Science Research




Office of
Science


Largest funder of physical
science research in U.S.




Bio Energy, Environment




Computing




Materials, Chemistry,
Geophysics



Particle Physics,
Astrophysics



Nuclear Physics



Fusion Energy,
Plasma Physics

Allocations primarily
controlled by DOE

- 80% DOE Annual production awards (ERCAP)
 - O(100)-O(10000) hour awards
 - Proposal-based, chosen by DOE program managers
- 10% DOE ASCR Leadership Computing Challenge
- 10% NERSC reserve

NERSC by the Numbers

NERSC USERS ACROSS US AND WORLD

50
States,
Washington D.C.
& Puerto Rico

53
Countries

~10,000 Annual Users from **~800** Institutions + National Labs



32%
Graduate
Students



19%
Postdoctoral
Fellows



15%
Staff
Scientists



13%
University
Faculty



8%
Undergraduate
Students



5%
Professional
Staff



60%
Universities



29%
DOE Labs



5%
Other
Government Labs



4%
Industry




1%
Small
Businesses



<1%
Private Labs

2,500 Referred Publications per Year




- ~1000 codes
- hundreds of users log in each day

NERSC Users Produce Groundbreaking Science

Materials Science

Revealing Reclusive Mechanisms for Solar Cells

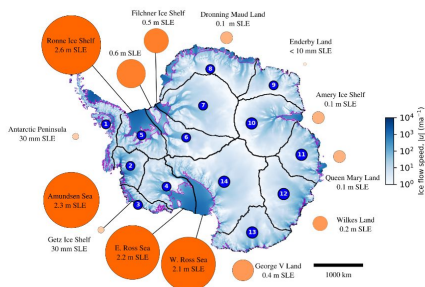
NERSC PI: C. Van de Walle, UC Santa Barbara, *ACS Energy Letters*



High Energy Physics

Shedding Light on Luminous Blue Variables


NERSC PI: Yan-Fei Jiang, UC Santa Barbara. *Nature*



Earth Sciences

Simulations Probe Antarctic Ice Vulnerability

NERSC PIs: D. Martin, Berkeley Lab; E. Ng, Berkeley Lab; S. Price, LANL. *Geophysical Research Letters*




Nuclear Physics

Enabling Science Discovery for STAR

NERSC PI: J. Porter, Berkeley Lab. *J. Phys.: Conference Series*

Advanced Computing


Scalable Machine Learning in HPC
NERSC PI: L. Oliker, Berkeley Lab, *21st International Conference on AI and Statistics*



Plasma Physics

Plasma Propulsion Systems for Satellites

NERSC PI: I. Kaganovich, Princeton Plasma Physics Lab, *Physics of Plasmas*



Nobel-Prize Winning Users



for the development of multiscale models for complex chemical systems

2013 Chemistry

Martin
Karplus



for the discovery of the accelerating expansion of the Universe through observations of distant supernovae

2011 Physics

Saul Perlmutter



for the discovery of the blackbody form and anisotropy of the cosmic microwave background radiation

2006 Physics

George Smoot



for their efforts to build up and disseminate greater knowledge about man-made climate change

2007 Peace

Warren Washington



for developing cryo-electron microscopy for the high-resolution structure determination of biomolecules in solution

2017 Chemistry

Joachim Frank




for the discovery of neutrino oscillations, which shows that neutrinos have mass

2015 Physics

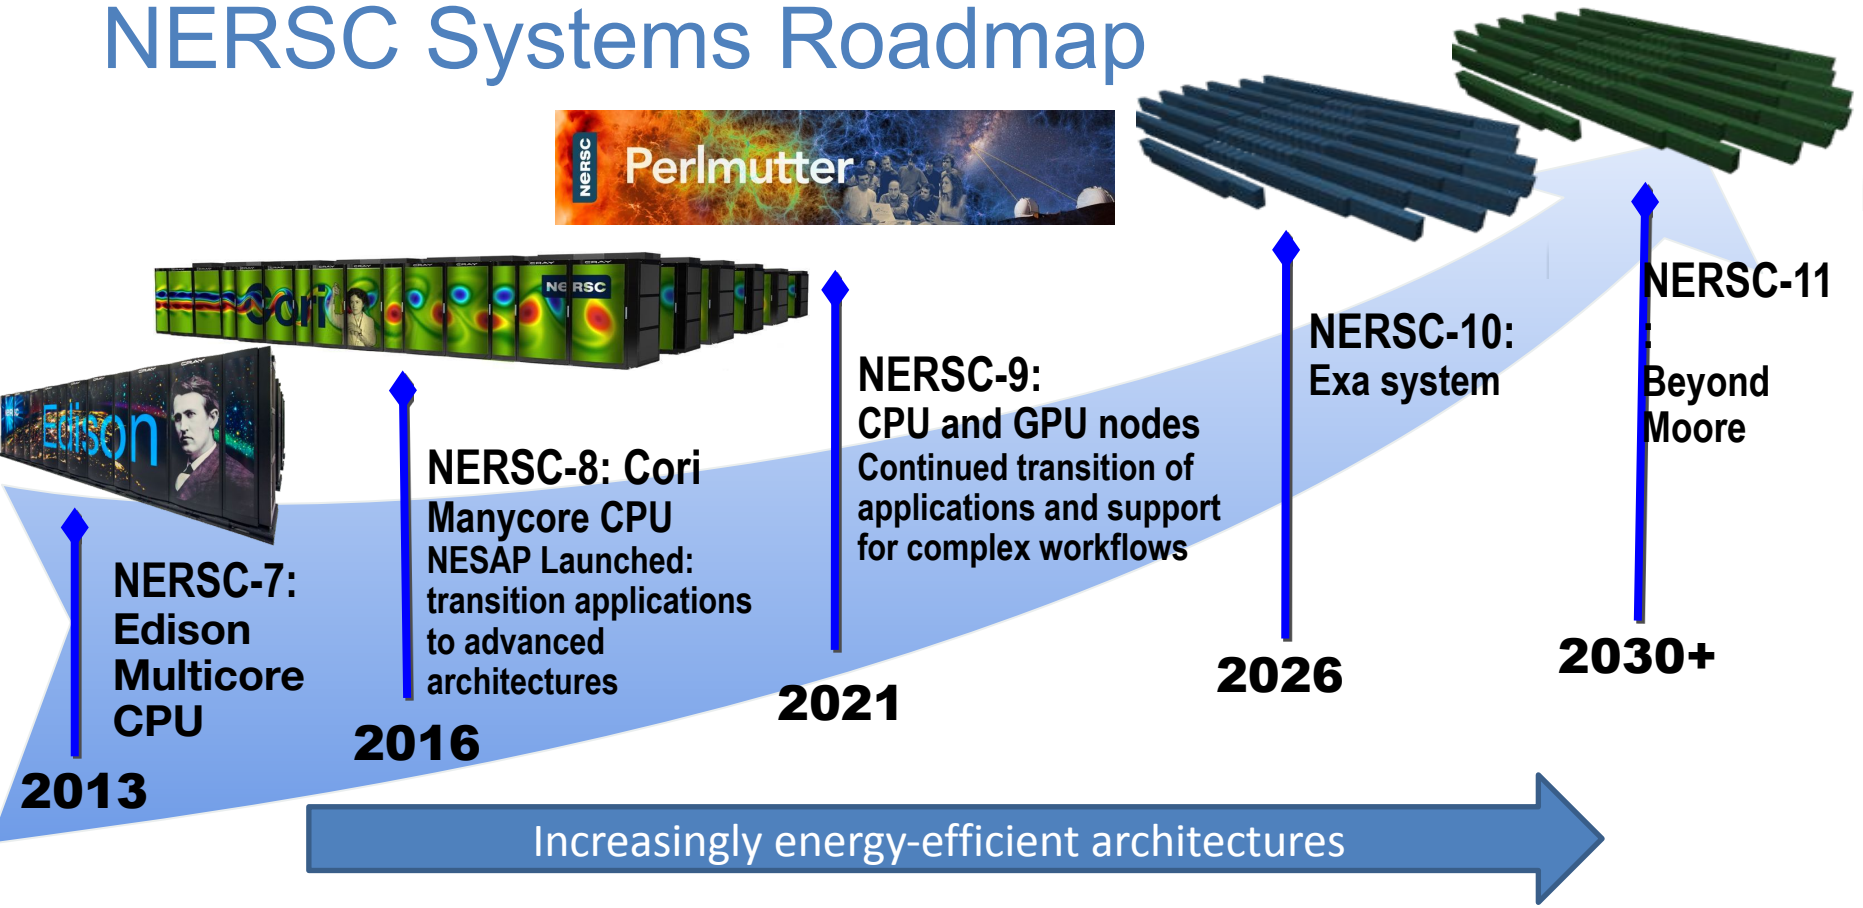
SNO Collaboration





Hardware

NERSC Systems Roadmap




NERSC-9 is named after Saul Perlmutter

- Shared 2011 Nobel Prize in Physics for discovery of the accelerating expansion of the universe.
- Supernova Cosmology Project, lead by Perlmutter, was a pioneer in using NERSC supercomputers combine large scale simulations with experimental data analysis
- Login “saul.nersc.gov”



Perlmutter: Optimized for Science



- HPE Cray System with 3-4x capability of Cori
- GPU-accelerated and CPU-only nodes
- HPE Cray Slingshot high-performance network
- All-Flash filesystem
- Application readiness program (NESAP)

Phase I: Arrived in 2021

- **1,536 GPU-accelerated nodes**
- 1 AMD “Milan” CPU + 4 NVIDIA A100 GPUs per node
- 256 GB CPU memory and 40 GB GPU high BW memory

Phase II Addition: Arrived in 2022

- **3,072 CPU only nodes**
- 2 AMD “Milan” CPUs per node
- 512 GB memory per node

HPC Systems: Perlmutter

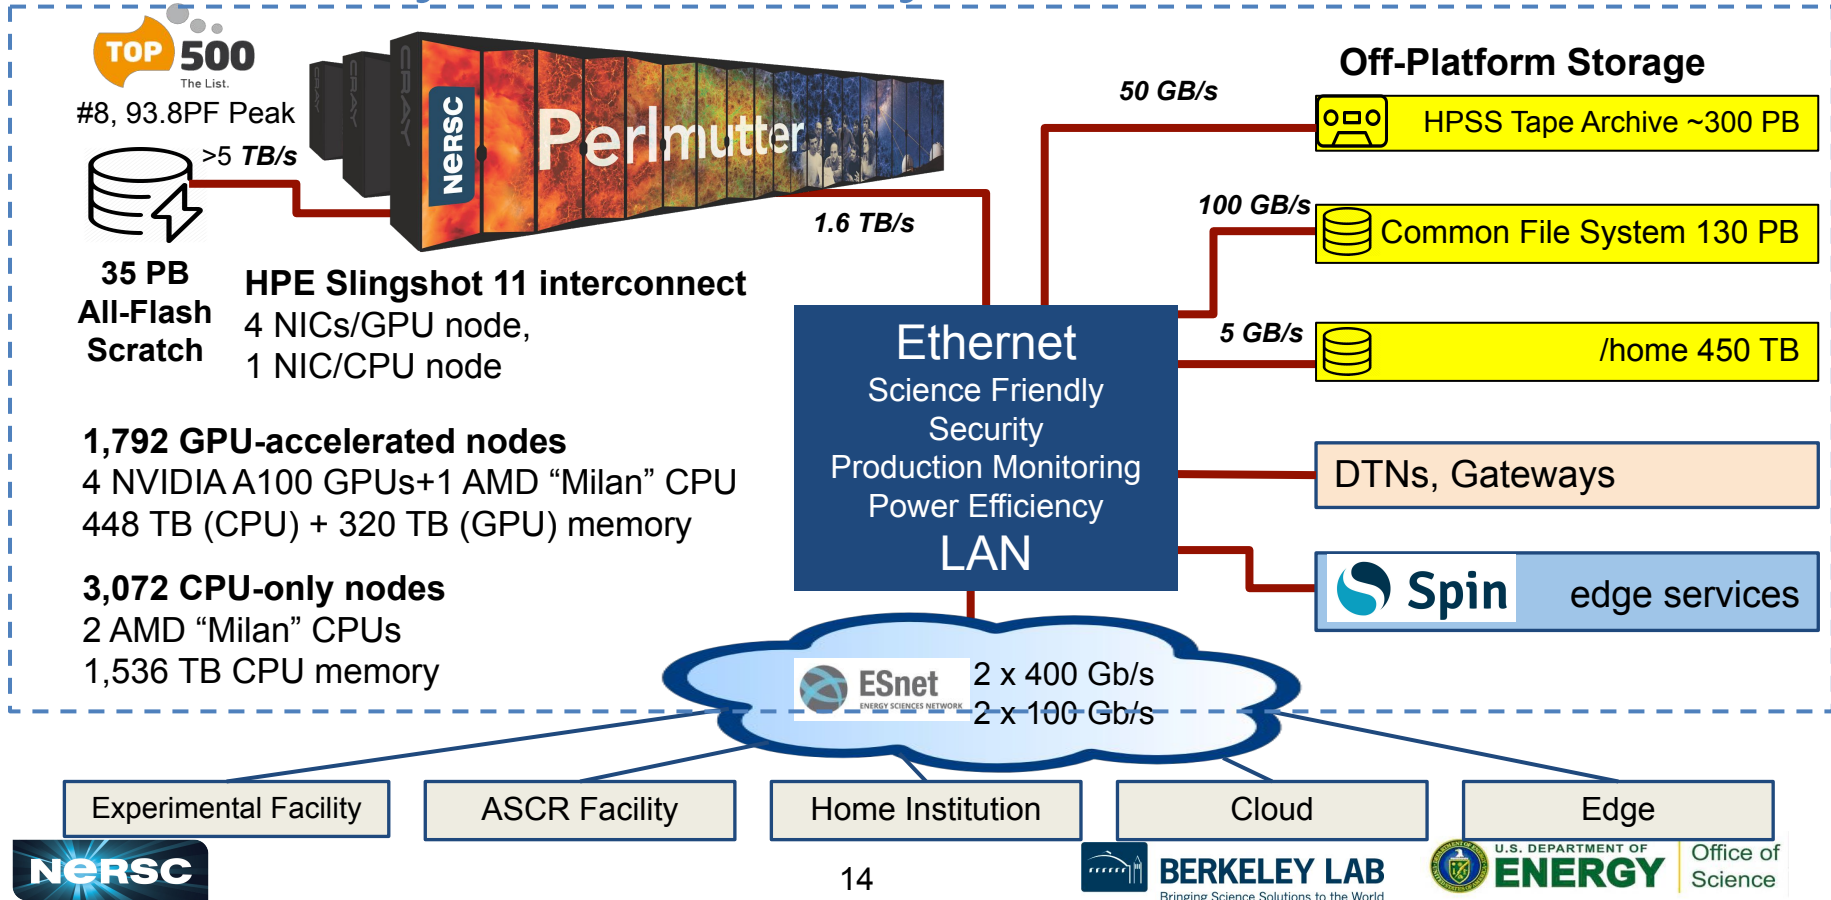
GPU nodes:

- Immense compute power from GPUs
- Large jobs using many GPUs encouraged
- Great for codes that can exploit GPU compute power

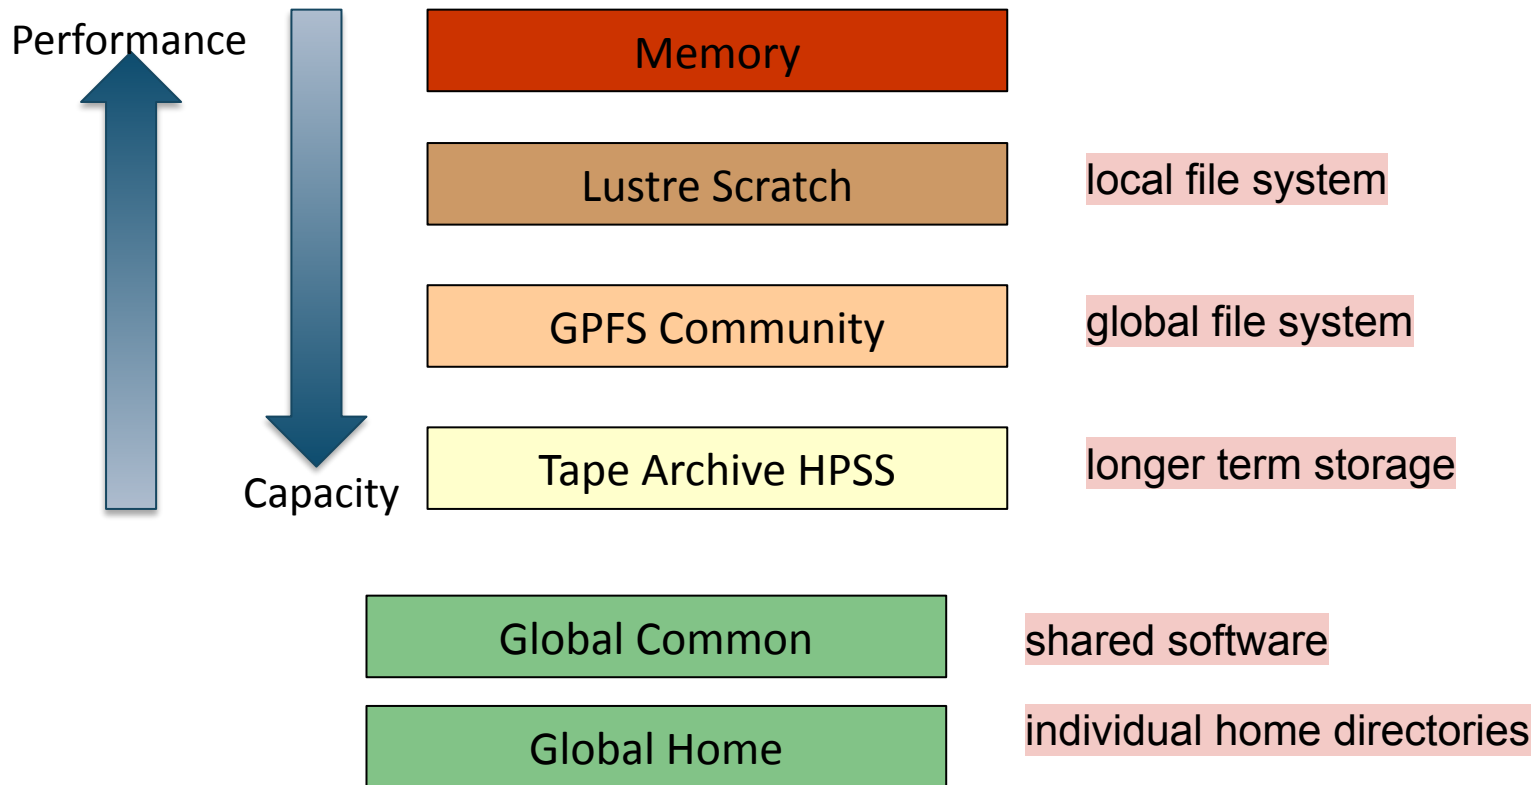
CPU nodes:

- Powerful CPUs (but only 10% of GPU compute power)
- Equivalent in compute power to all of Cori (previous system)
- More like a traditional cluster
- Great for throughput jobs

NERSC Systems Ecosystem



Simplified NERSC File Systems



Perlmutter File Systems

Global Home

- Permanent, relatively small storage
- NOT tuned to perform well for parallel jobs
- Snapshot backups
- **Perfect for storing data such as source codes, shell scripts**
- **cd \$HOME**

Global Community File System (CFS)

- Permanent, larger storage
- Medium performance for parallel jobs
- Snapshot backups
- **Perfect for sharing data within research group**
- **cd \$CFS**

Local Scratch

- Large, temporary storage
- Optimized for read/write operations, NOT storage
- Not backed up
- Purge policy (8 weeks)
- **Perfect for staging data and performing computations**
- **cd \$SCRATCH**

Long-Term Storage System

HPSS

- High-Performance Storage System
- Archival storage of infrequently accessed data
- Hierarchical storage:
 - Data first ingested onto high-performance disk arrays
 - Migrated to large enterprise tape subsystem for long-term retention

Using NERSC File Systems (1)

- Analogy:
 - Computing = baking
 - Input = baking ingredients
 - Output = cake
- NERSC is gigantic shared kitchen space with all the latest kitchen gadgets
 - Computers = ovens
 - Home, CFS = pantry, fridge
 - HPSS = freezer
 - Scratch = kitchen counter



Using NERSC File Systems (2)

- When baking, stage ingredients from pantry and fridge (plus maybe rarely used ingredients from freezer) onto kitchen counter
 - Likewise, stage data and executable onto scratch file system




New Mexico. Mrs. Fidel Romero proudly exhibits her canned food, 1946 US National Archives NWDNS-33-S-12785

Using NERSC File Systems (3)

- After baking, clean up after yourself!
- It's okay to let cake cool on kitchen counter, but need to leave space clean for next user
 - After a while, we will clean up if you don't, but not like you would want
 - We will throw all your materials in the trash (even your cake!)



Queen cakes cooling on a wire rack by James Petts
<https://www.flickr.com/photos/14730981@N08/13475333725/>





Software

Software

- Cray supercomputers OS is a version of Linux
- Compilers are provided on machines
- NERSC strategically provides software to users
- Libraries: many libraries are provided by vendor, still others provided by NERSC
- Applications: NERSC compiles and supports many software packages for our users
- Extreme-scale Scientific Software Stack (E4S) at NERSC is a curated software stack contains hundreds of packages delivered via the SPACK package manager

Chemistry & Materials Science Applications



Molpro



QUANTUM ESPRESSO



BerkeleyGW

abinit.

- *More than 13.5 million lines of source code Compiled, Optimized, and Tested*



NAMD

Scalable Molecular Dynamics



GÅMESS

LAMMPS

WANNIER90

CPMD



NWCHEM

HIGH-PERFORMANCE COMPUTATIONAL CHEMISTRY SOFTWARE

NERSC has a rich data ecosystem!



globus online



jupyter



data transfer and access



mongoDB®




MySQL™

data management



julia



data analytics



PyTorch



scikit learn



machine learning



ParaView
Parallel Visualization Application

visualization




SHIFTER



Spin

containers



GNUparallel



Parsl



papermill

workflows




FireWorks

Perlmutter Supports Multiple Compilers and Every GPU Programming Model

	Fortran/ C/C++	CUDA	OpenACC 2.x	OpenMP 5.x	CUDA Fortran	Kokkos / Raja	MPI	HIP	DPC++ / SYCL
NVIDIA	Vendor Supported	Vendor Supported	Vendor Supported	Vendor Supported	Vendor Supported	Vendor Supported	Vendor Supported	NERSC Supported	NERSC Supported
CCE	Vendor Supported	NERSC Supported	Vendor Supported	Vendor Supported	NERSC Supported	Vendor Supported	Vendor Supported	NERSC Supported	NERSC Supported
GNU	Vendor Supported	Vendor Supported	Vendor Supported	Vendor Supported	NERSC Supported	Vendor Supported	Vendor Supported	NERSC Supported	NERSC Supported
LLVM	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported
Intel	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported	NERSC Supported

Vendor
Supported

NERSC
Supported



Interacting with NERSC



BERKELEY LAB





U.S. DEPARTMENT OF
ENERGY

Office of
Science

Interacting with NERSC

- NERSC User Engagement Group
 - User Engagement, outreach
 - User Training and Documentation
- NERSC Consulting & Account Support
 - User Tickets
 - User Appointments
- NERSC Operations
- NERSC User Group (NUG)


Consulting & Account Support Team



NERSC Consulting & Account Support

- The first people you interact with when submitting a ticket or calling
- In 2023, we handled 8,260 tickets from 3,296 unique users

2023 Incidents by Month



Login Node	Network
Other	Performance
Programming	Scratch/Slurm Hanging
Security	Software
2/3 ▼	

(empty)	Account Support (Account Support)
Allocations	Data/IO
Datawarp	ERCAP
Hardware	Running Jobs
1/3 ▼	

Account Support (support)	System Outage
Consulting Triage	Workstation Support
3/3 ▼	



NERSC Consulting: Expectations

- Our first response will be within four business-hours
- We will help you resolve your problem, and keep you apprised of progress
- We will attempt to accommodate user needs that don't fit within our operating structure
- We welcome user feedback and constructive criticism

NERSC Consulting: Tips & Tricks

- Help us help you!
- Provide specifics:
 - What is the problem?
 - What machine? Which node? Which file path?
 - When did it happen?
 - What modules were loaded?
 - How did you try to fix or work around it?
 - How can I reproduce the problem?

- Tips for filing a good ticket:

<https://docs.nersc.gov/getting-started/#how-to-file-a-good-ticket>

NERSC User Appointments

- In 2018, we began offering “office hours”
 - Open Zoom meeting which users could join to get help with a particular topic, e.g., MFA, KNL Optimization, ERCAP, etc.
 - Shortcoming: long periods with no participants, then many jump on simultaneously
- Appointments: more efficient use of everyone’s time
- 30-minute appointments offered on a variety of topics:
 - GPU basics, Optimization, File Systems, Using GPUs in Python, Containers, NERSC 101, Checkpoint/Restart jobs with MANA, Spin, Appentra Codee
- Schedule an appointment: nersc.as.me

NERSC User Training

- NERSC provides a robust training program for users of all skill levels, interests, and personas
 - All trainings are recorded, professionally captioned, & posted to [NERSC YouTube channel](#)
 - Slides posted to training event webpage
- For more information on upcoming and past events, see <https://www.nersc.gov/users/training/events/>
- Collection of [Categorized Training Materials](#)
- [Training Events Archive](#)

A few Upcoming Training Events

<https://www.nersc.gov/users/training/events/>

- Data Day, Feb 21-22
- Office Hour, Feb 23
- Performance Portability Series, next event Feb 26
- Forge, Mar 13
- AMReX, Mar 14
- Fortran, Mar 26-27



NERSC Data Day, Feb 21-22 2024 »

February 21, 2024


NERSC is hosting Data Day, a hybrid two day event on February 21-22, 2024. Join us for exciting talks and demonstrations designed to showcase the latest and greatest data-focused tools for scientific computing on Perlmutter.

[Read More »](#)

Perlmutter and Data Day Office Hours, February 2024 »

February 23, 2024

Users are welcome to bring questions and their own applications to get one-on-one help with using Perlmutter's GPU nodes on Friday February 23, 2024 for Office Hours from 9am - 12pm PACIFIC/UTC-7 [Read More »](#)



Performance Portability Training Series: 2023 - 2024 »

February 26, 2024

This performance portability training series, offered by OLCF, NERSC, and ALCF, features training sessions on various performance portable programming solutions to help ease developer transitions between current and emerging high-performance computing (HPC) systems.

[Read More »](#)

Performance Portability for Next-Generation Heterogeneous Systems, Feb 26, 2024 »

February 26, 2024

This performance portability overview session, presented by Tom Deakins of University of Bristol is part of the Performance Portability training series. [Read More »](#)



NERSC Operations

- Operations staff are on site 24/7/365 to supervise operation of the machine room
- Operations know the health of the machines and can help users with some tasks (killing jobs, changes to running reservation, etc.)
- [NERSC MOTD](#) (message of the day, live status)
- Please avoid contacting Operations except in urgent cases

NERSC User Group (NUG)

- Community of NERSC users
- Source of advice and feedback for NERSC (we listen!)
- Executive Committee: 3 representatives from each office + 3 members-at-large
- Monthly teleconferences hosted by NERSC (usually 3rd Thursday of the month, 11 am to noon)
- NUG Slack: <https://www.nersc.gov/users/NUG/nersc-users-slack/>
- Special Interest Groups (SIG): experimental facilities, Fortran, WRF, and more forming (such as graduate students)
- Join us NUG Annual Meeting (Oct 22-24, 2024) with exciting 50th anniversary of NERSC related events!

User Responsibilities & Expectations

- Be kind to your neighbor users
 - Don't abuse the shared resources!
- Use your allocation smartly
 - Pick the right resource for your job and your data
- Back your stuff up
 - Especially from scratch, which has a purge policy
- Acknowledge NERSC in your papers
 - Acknowledge us so we can stay in business!
- Pay attention to security
 - Don't share your account with others!

Thank You and
Welcome to
NERSC!

