NERSC: Welcome to Scientific Discovery through Computation with High-Performance Computing

Rebecca Hartman-Baker, PhD
User Engagement Group Lead
Charles Lively III, PhD and Lipi Gupta, PhD
Science Engagement Engineer(s)

New User Training Fall 2023
September 7, 2023

The Plot

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

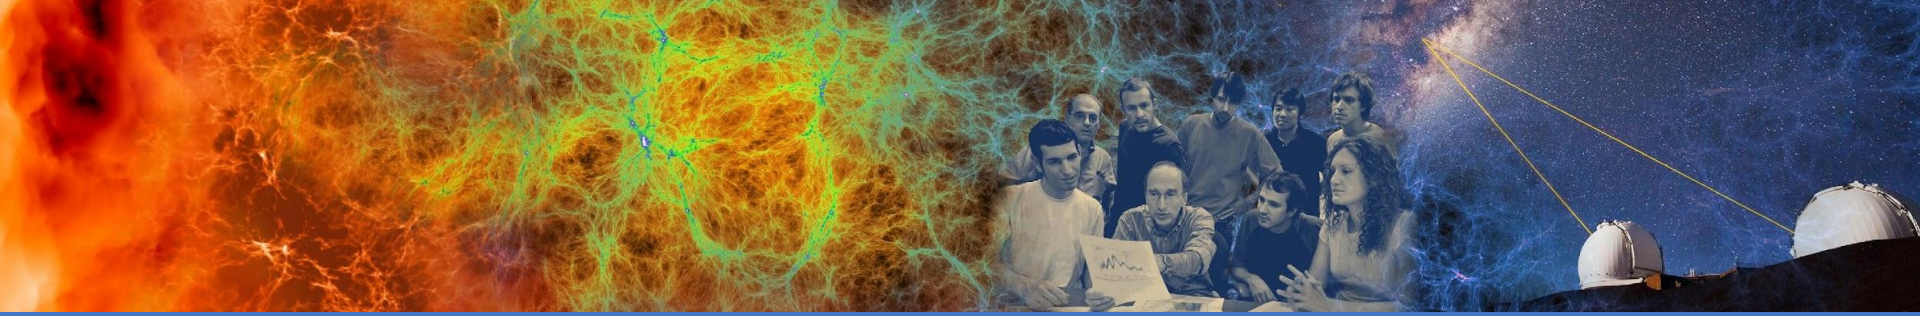

Introduction to NERSC

National Energy Research Scientific Computing Center

- NERSC is a national supercomputer center funded by the U.S. Department of Energy Office of Science (SC)
  - Supports SC research mission
  - Part of Berkeley Lab
- If you are a researcher with funding from SC who needs computational resources at scale, you can use NERSC
  - Other researchers can apply if research is in the SC mission
- NERSC supports 9,000 users, 1,000 projects
  - From all 50 states + international; 65% from universities
  - Hundreds of users log on each day

NERSC is the Production HPC & Data Facility for DOE Office of Science Research

U.S. DEPARTMENT OF ENERGY | Office of Science

Largest funder of physical science research in U.S.

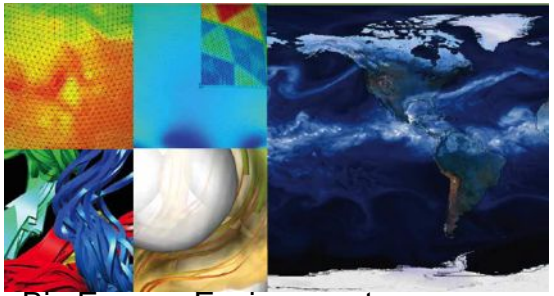

Bio Energy, Environment

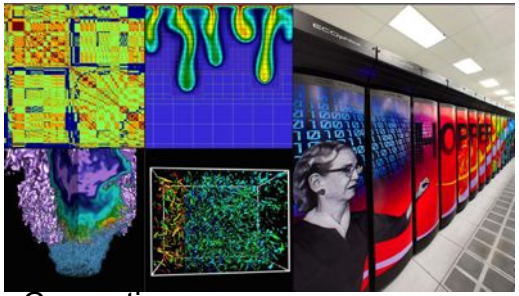

Computing

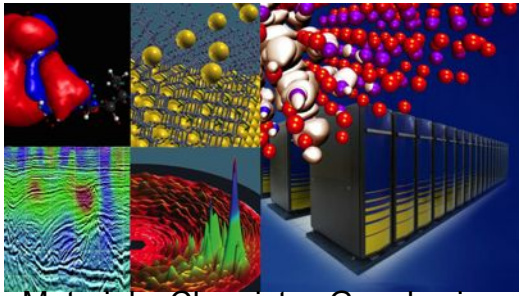

Materials, Chemistry, Geophysics

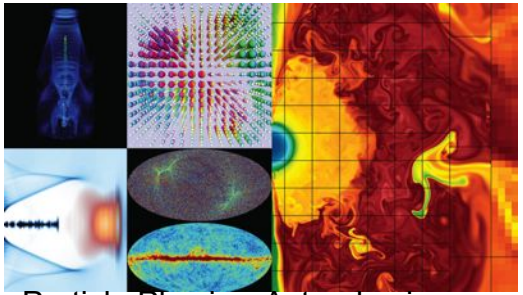

Particle Physics, Astrophysics

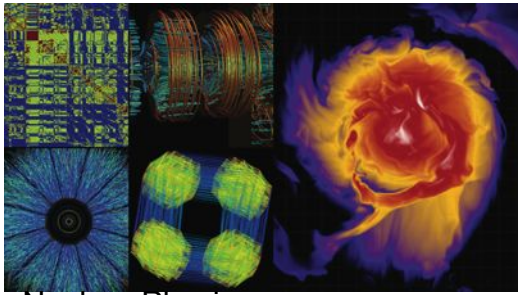

Nuclear Physics

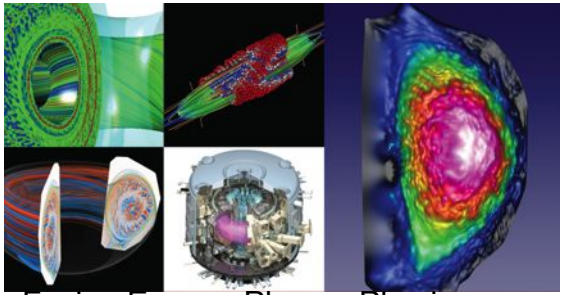

Fusion Energy, Plasma Physics

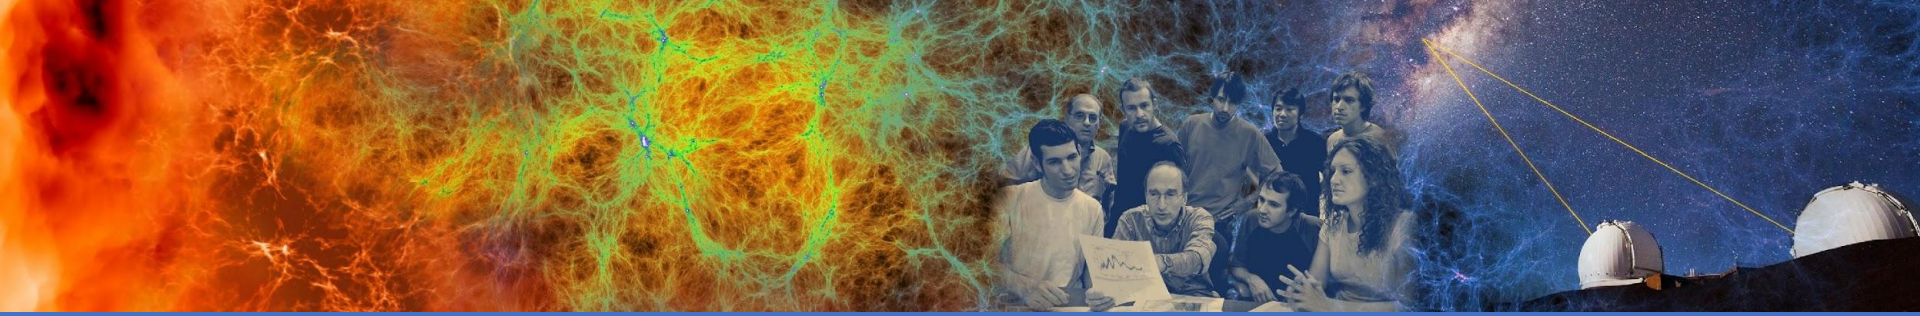

NERSC + Science = Discovery

About the NERSC Community

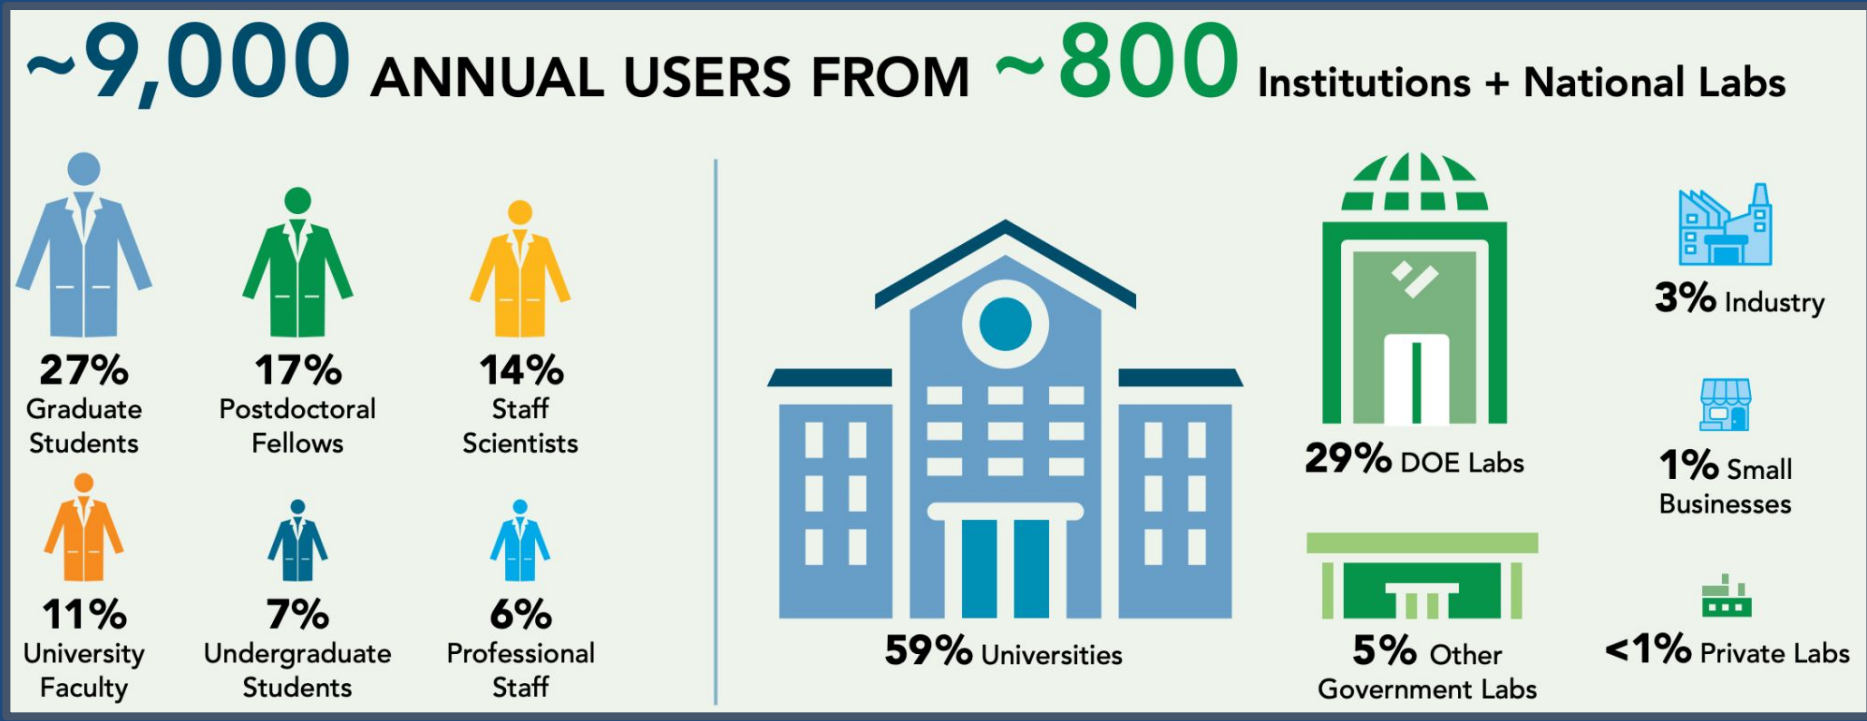

NERSC's Users Produce Groundbreaking Science

Materials Science

Revealing Reclusive Mechanisms for Solar Cells

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

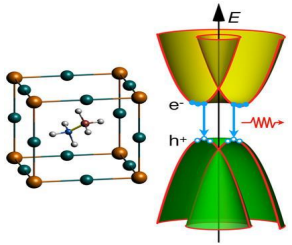

High Energy Physics

Shedding Light on Luminous Blue Variables

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

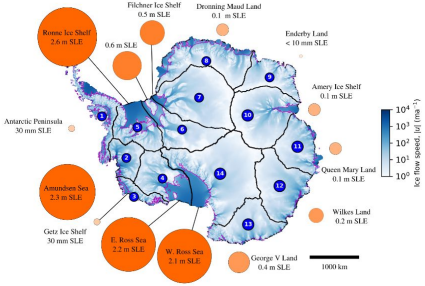

Earth Sciences

Simulations Probe Antarctic Ice Vulnerability

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

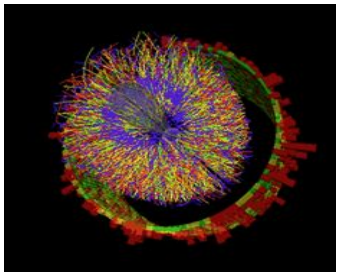

Nuclear Physics

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*

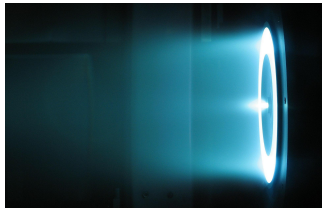

2,500 Refereed Publications per Year

Nobel-Prize Winning Users

2013 Chemistry

for the development of multiscale models for complex chemical systems

Martin Karplus

2011 Physics

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

Saul Perlmutter

2006 Physics

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

George Smoot

2007 Peace

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

Warren Washington

2017 Chemistry

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

Joachim Frank

2015 Physics

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

SNO Collaboration

NERSC

BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY | Office of Science

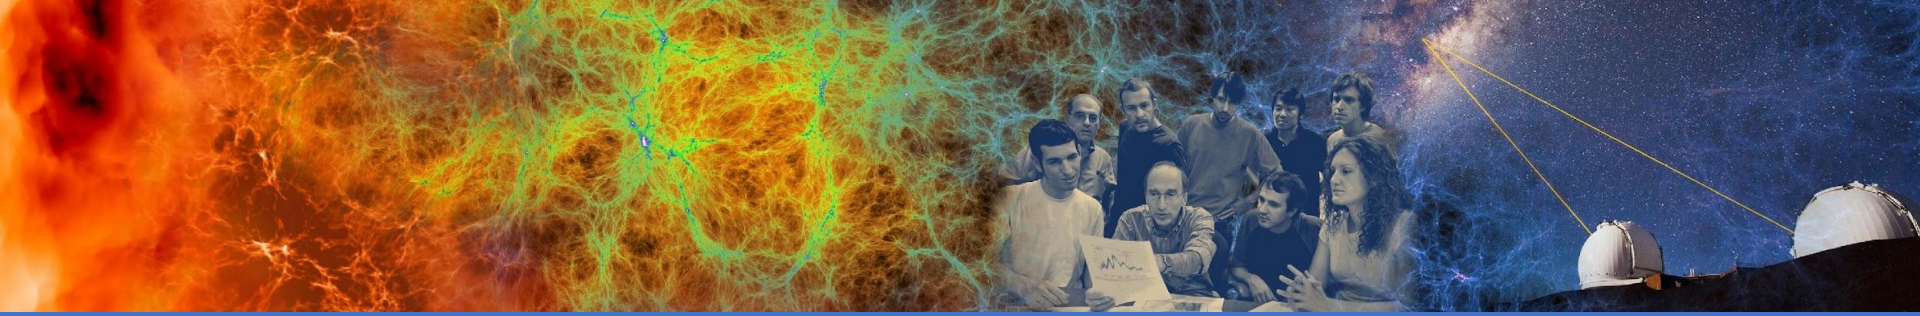

Hardware

NERSC Systems Roadmap

NERSC-7:
Edison
Multicore
CPU

2013

NERSC-8: Cori
Manycore CPU
NESAP Launched:
transition applications
to advanced
architectures

2016

NERSC-9:
CPU and GPU nodes
Continued transition of
applications and support
for complex workflows

2021

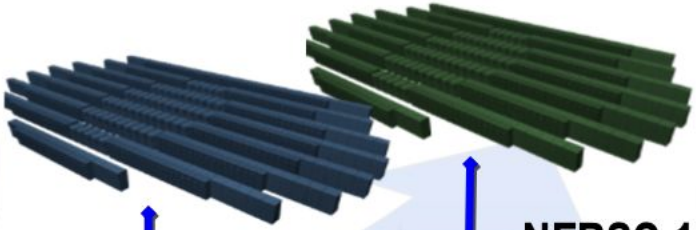

NERSC-10:
Exa system

2026

NERSC-11:
Beyond
Moore

????

Increasingly energy-efficient architectures

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”

A True Supercomputer…but…

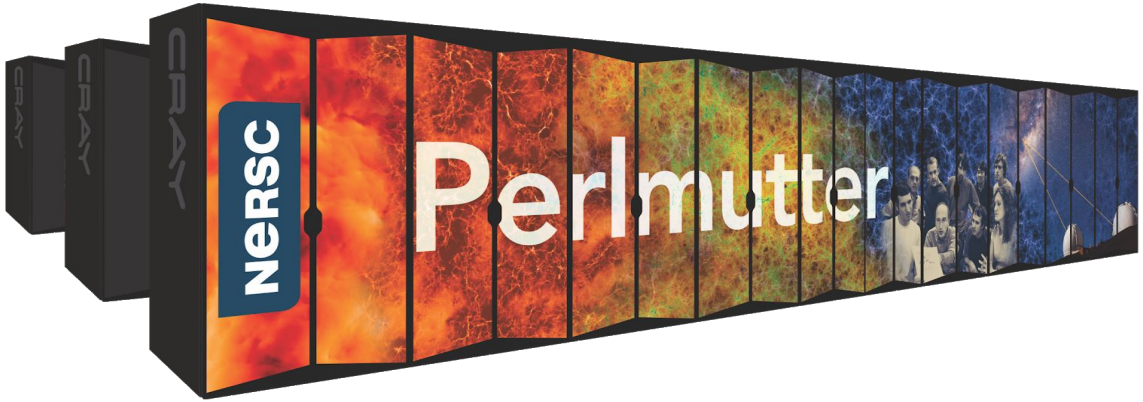

vs.

… not so different from a super high-end desktop computer.

Or rather, a lot of super high-end desktop computers.

Perlmutter (left) has ~13,300 nodes (~ high-end desktop computers)

Over 760,000 compute cores

Perlmutter: Optimized for Science

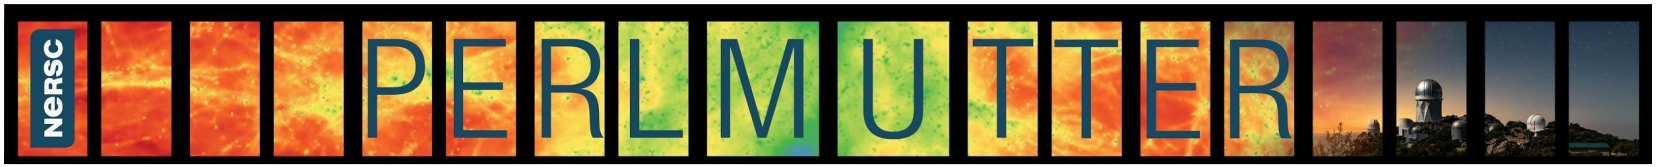

- HPE Cray System with 3-4x capability of Cori
- GPU-accelerated and CPU-only nodes
- HPE Cray Slingshot high-performance network
- All-Flash filesystem
- Application readiness program (NESAP)

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
- 35 PB FLASH scratch file system
- User access and system management nodes

Phase II Addition: Arrived in 2022

- 3,072 CPU only nodes
- 2 AMD “Milan” CPUs per node
- 512 GB memory per node
- Upgraded high speed network
- CPU partition exceeds performance of entire Cori system

NERSC

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 (Haswell + KNL)
- More like a traditional cluster
- Great for throughput jobs

File Systems

- Global File Systems:
  - Home
  - Community (CFS)
- Local File Systems:
  - Scratch
- Long-term Storage System:
  - HPSS

Global File Systems

Home

- Permanent, relatively small storage
- Mounted on all platforms
- NOT tuned to perform well for parallel jobs
- Quota cannot be changed
- Snapshot backups (7-day history)
- **Perfect for storing data such as source code, shell scripts**

Community File System (CFS)

- Permanent, larger storage
- Mounted on all platforms
- Medium performance for parallel jobs
- Quota can be changed
- Snapshot backups (7-day history)
- **Perfect for sharing data within research group**

Local File Systems

Scratch

- Large, temporary storage
- Local to machine
- Optimized for read/write operations, NOT storage
- Not backed up
- Purge policy (12 weeks)
- **Perfect for staging data and performing computations**

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
- (For more info please see later presentations)

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/

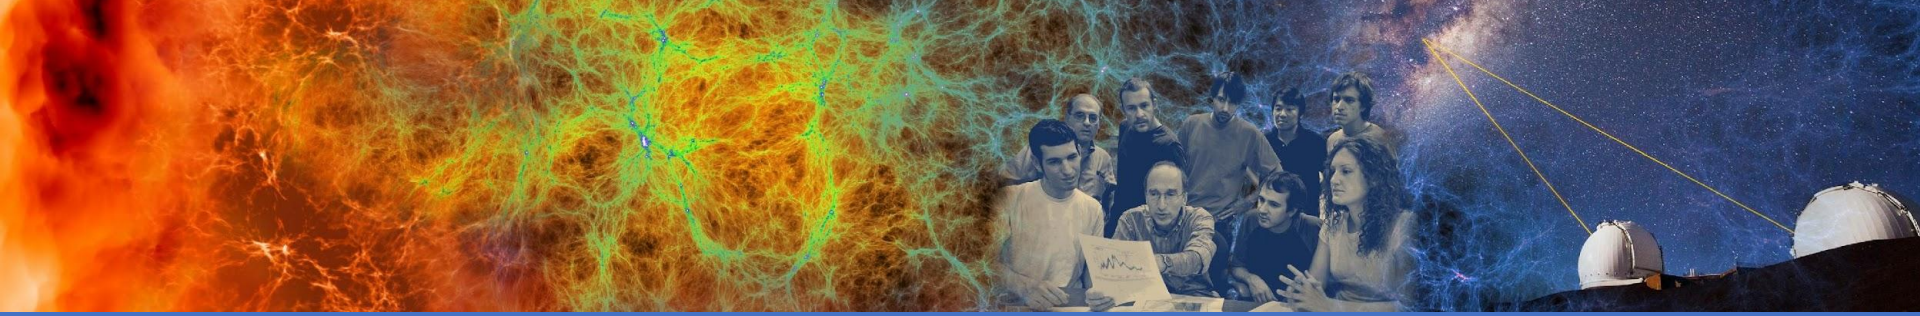

Software

Software

- Cray supercomputers OS is a version of Linux
- Compilers are provided on machines
- Libraries: many libraries are provided by vendor, still others provided by NERSC
- Applications: NERSC compiles and supports many software packages for our users
- (For more details, please see later presentations!)

Chemistry & Materials Science Applications

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

LAMMPS

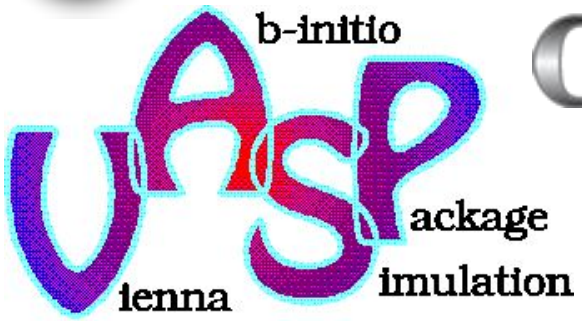

NERSC has a rich data ecosystem!

data transfer and access

data management

data analytics

machine learning

visualization

containers

workflows

Software: Policy

- Software version defaults consistent for allocation year
  - Same Cray programming environment software will be available all year, with exceptions for security issues or major OS upgrades
- Software at NERSC classified into 4 support levels
  - **Priority:** provided by NERSC, high priority, NERSC performs functionality & performance testing regularly
  - **Provided:** provided by NERSC, moderate priority, NERSC performs functionality testing regularly
  - **Minimal:** not generally provided by NERSC, low priority, NERSC performs no testing
  - **Restricted:** not allowed on NERSC resources (e.g., export controlled software, Gaussian)

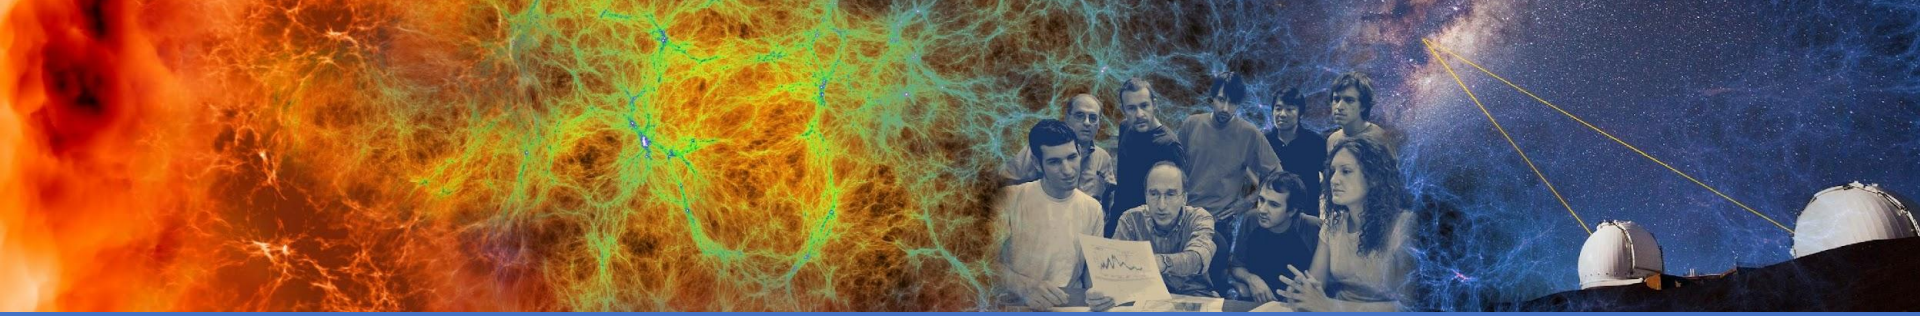

Interacting with NERSC

Interacting with NERSC

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

Our People

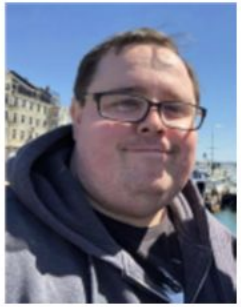

Justin Cook

Kevin Gott

Lipi Gupta

Rebecca Hartman-Baker

Helen He

Kadidia Konate

Charles Lively

Erik Palmer

Kelly Rowland

Shahzeb Siddiqui

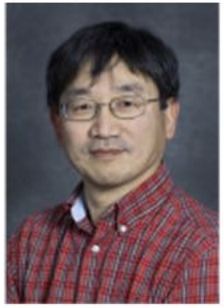

Woo-Sun Yang

Alumni:

Tiffany Connors
Zhengji Zhao
Steve Leak

UEG Mission

The User Engagement Group engages with the NERSC user community to increase user productivity via **advocacy, support, training,** and the **provisioning of usable computing environments.**

Consulting & Account Support Team

NERSC Consulting & Account Support

- The first people you interact with when submitting a ticket or calling
- In 2022, we handled 6,079 tickets from 2,664 unique users

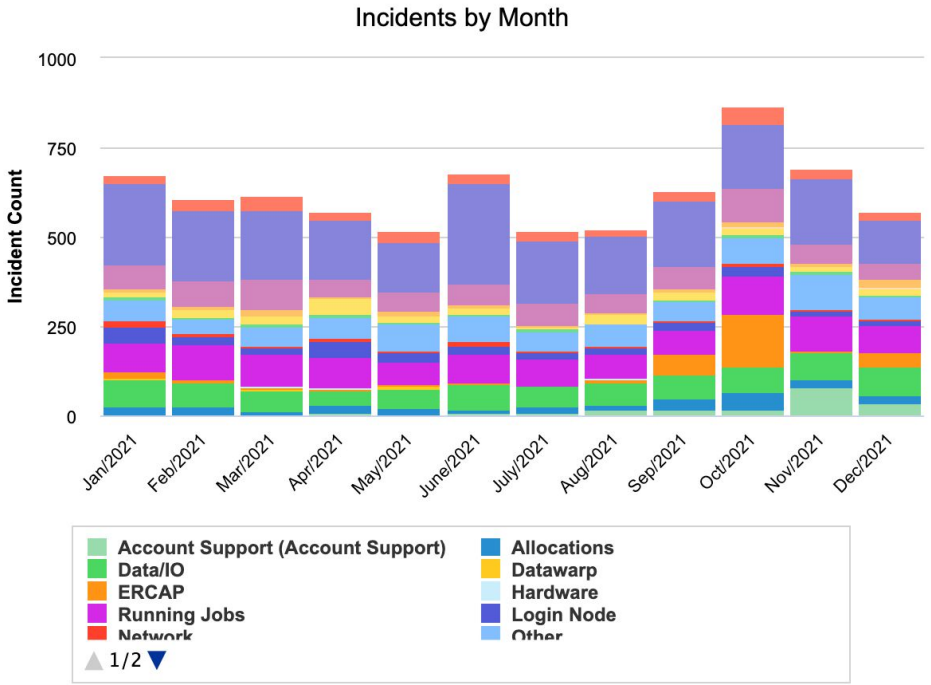

NERSC Consulting: Expectations

- 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?
  - When did it happen?
  - What modules were loaded?
  - How did you try to fix or work around it?
- 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, KNL 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 and interests
  - All trainings are recorded, professionally captioned, & posted to NERSC YouTube channel
  - Slides posted to training event webpage
- For more information on current/upcoming events, see https://www.nersc.gov/users/training/events/

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.)
- 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: join at https://www.nersc.gov/users/NUG/nersc-users-slack/ (login required)
- Join us for the NUG Annual Meeting (in-person or online) September 26-28, 2023

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!
NERSC

Questions?