NERSC Overview

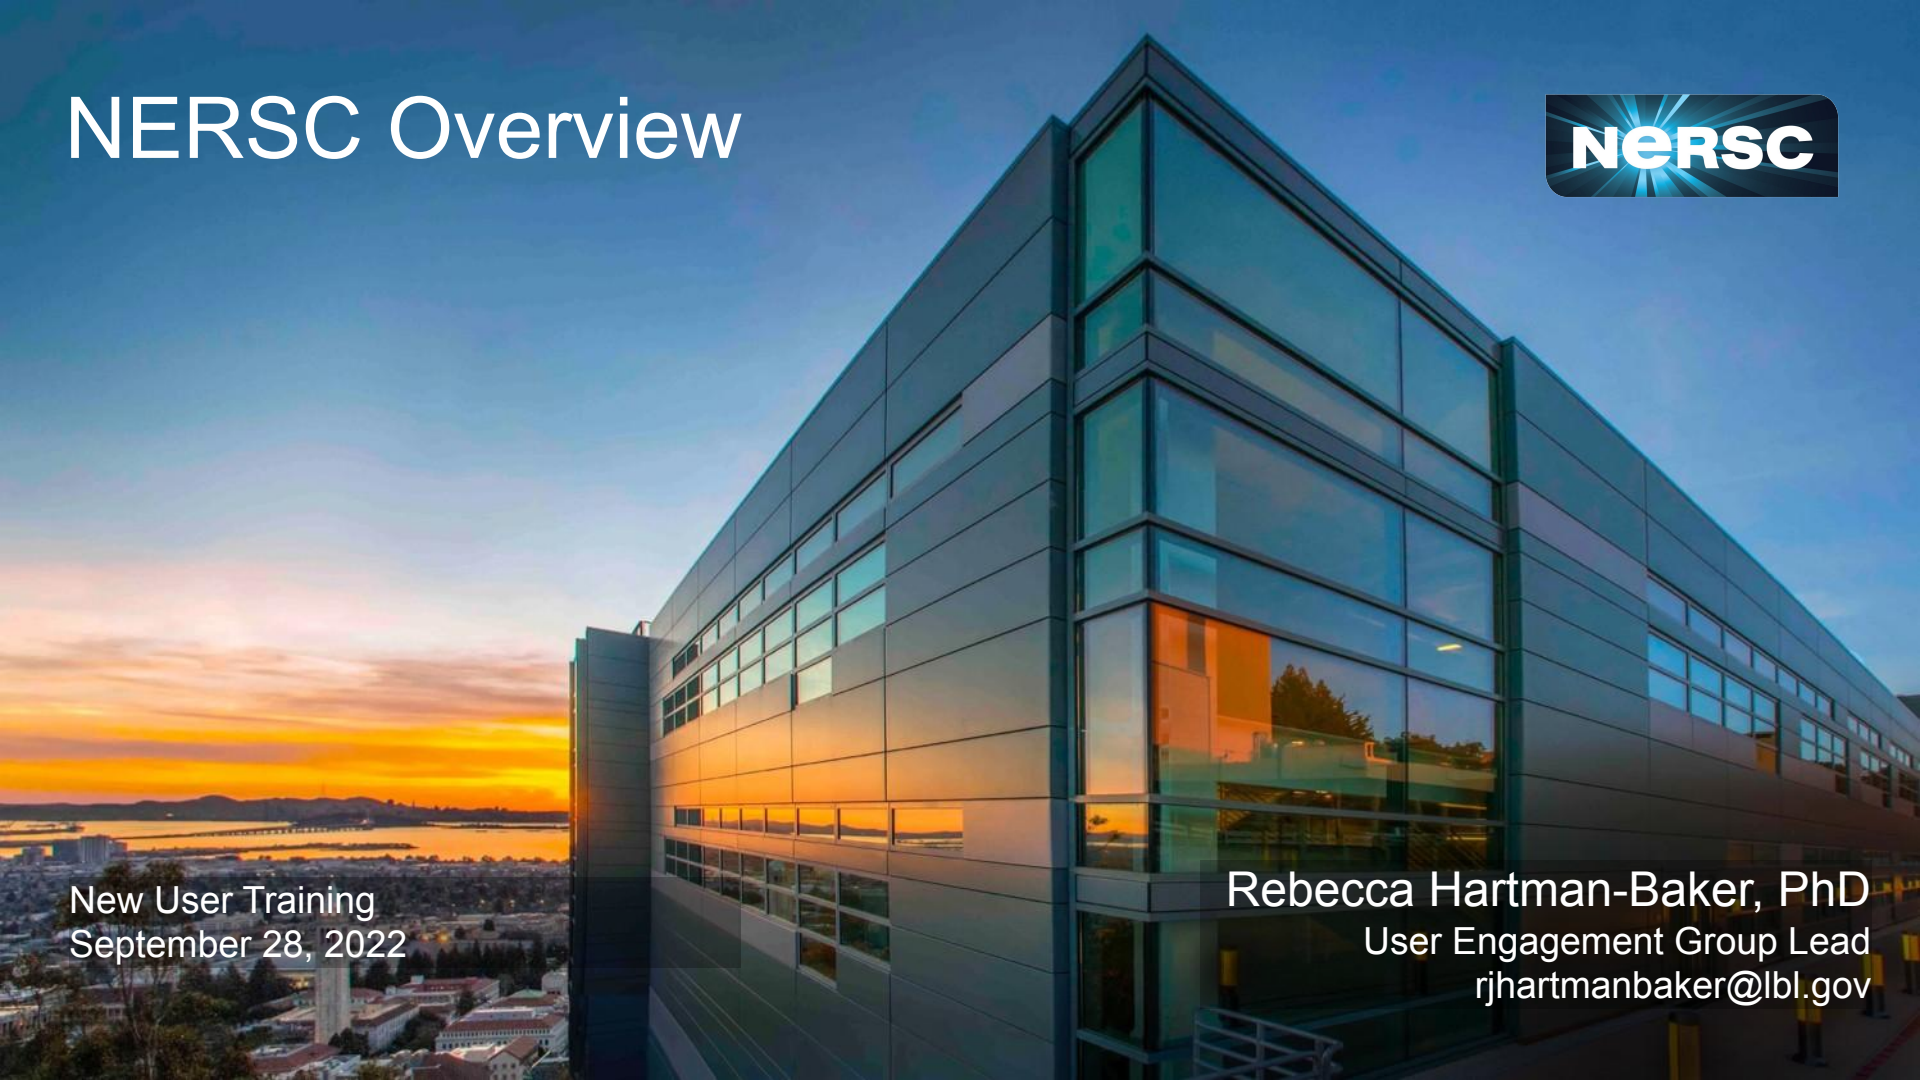

New User Training
September 28, 2022

Rebecca Hartman-Baker, PhD
User Engagement Group Lead
rjhartmanbaker@lbl.gov

Agenda

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

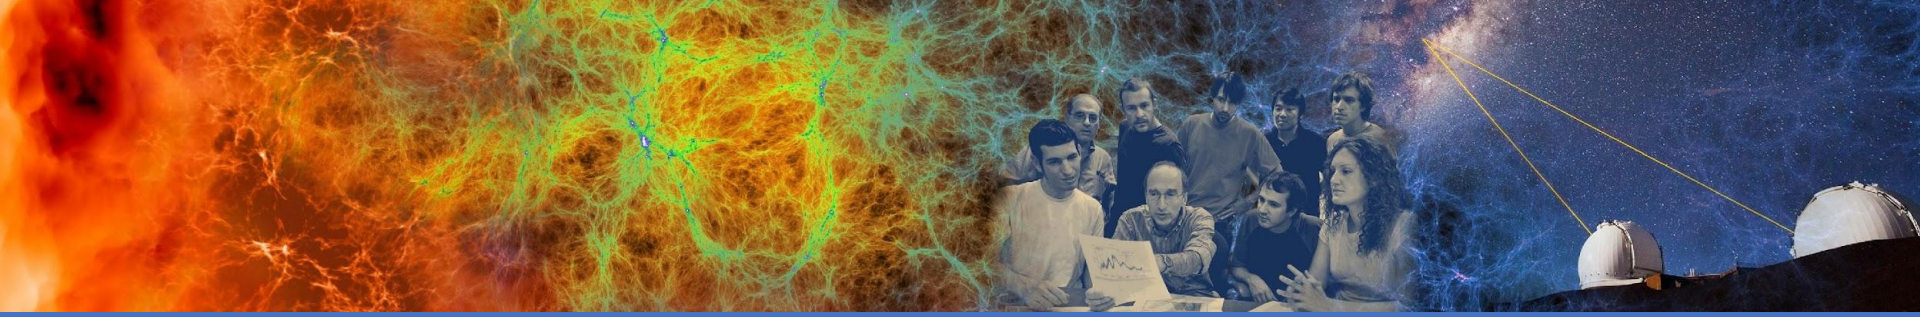

Introduction to NERSC

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
- 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

About NERSC

- Diverse workload:
  - 9000 users, 900 projects
  - 1000 codes, 100s of users daily
- 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

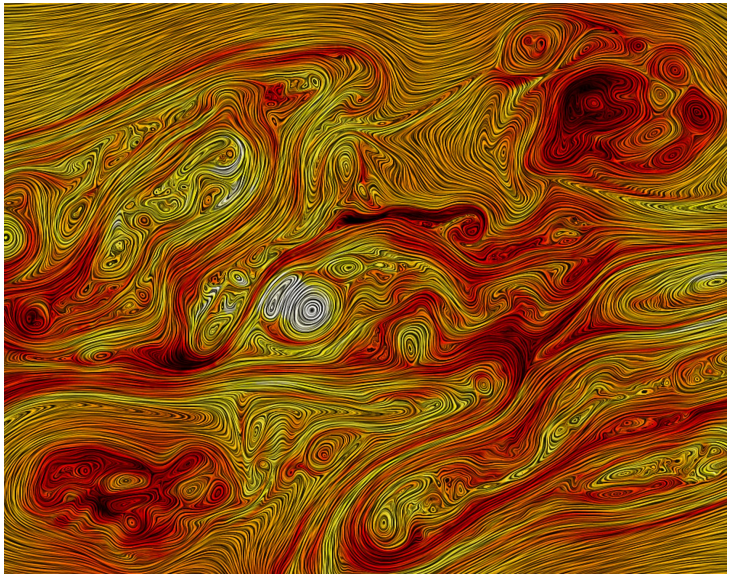

Turbulence in Solar Wind

NERSC Users Produce Groundbreaking Science

NERSC users produce more publications than any other center in the world‡; ~2,500/year

Materials Science

Revealing Reclusive Mechanisms for Solar Cells

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

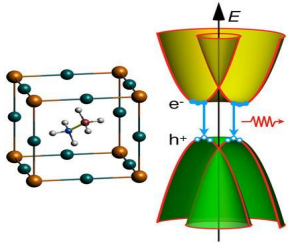

High Energy Physics

Shedding Light on Luminous Blue Variables

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

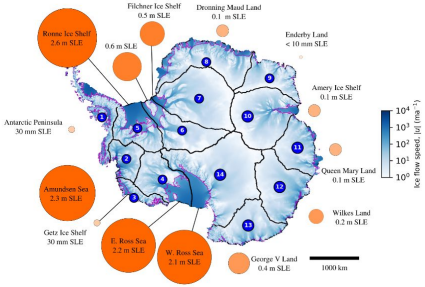

Earth Sciences

Simulations Probe Antarctic Ice Vulnerability

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

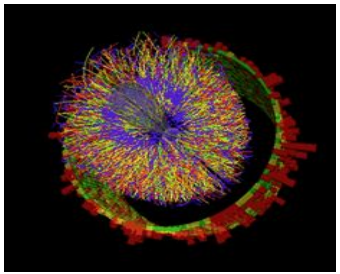

Nuclear Physics

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*

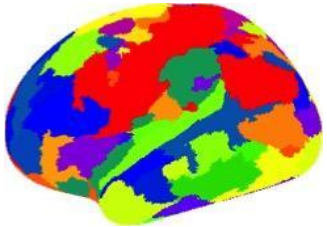

Plasma Physics

Plasma Propulsion Systems for Satellites

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

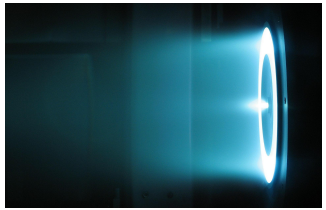

‡ as far as we can tell

DOE View of NERSC Workload

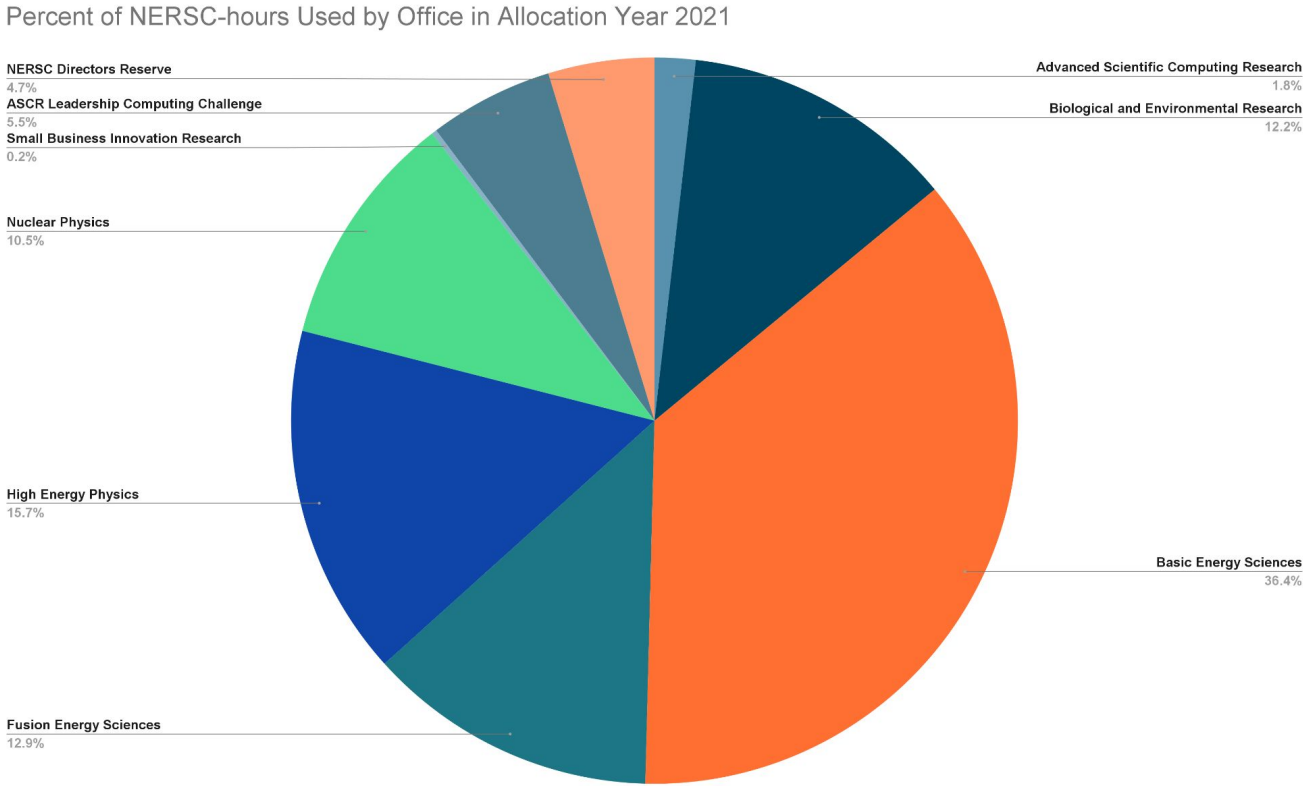

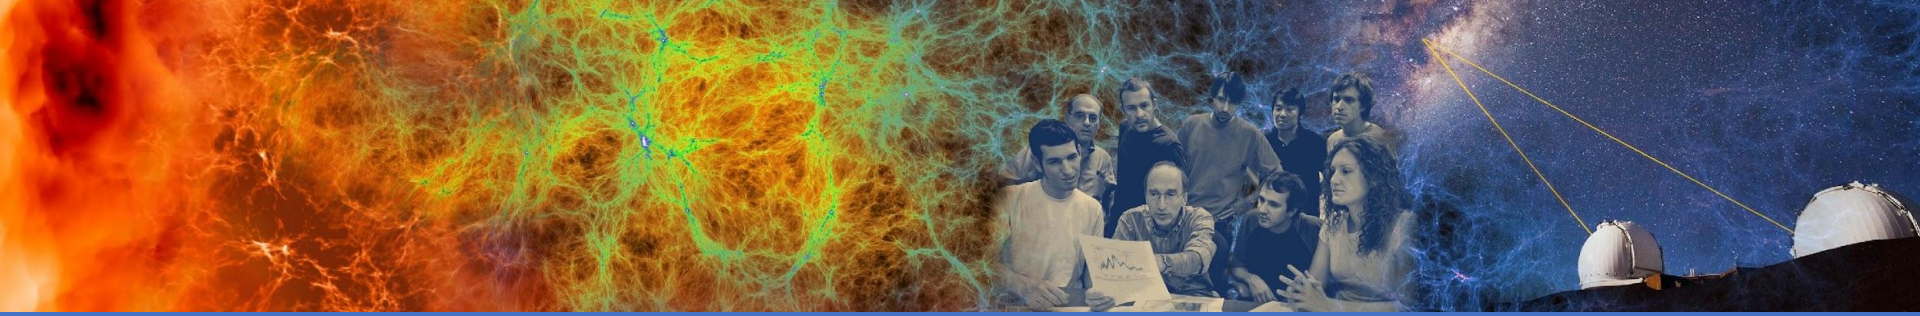

Hardware

NERSC Systems Roadmap

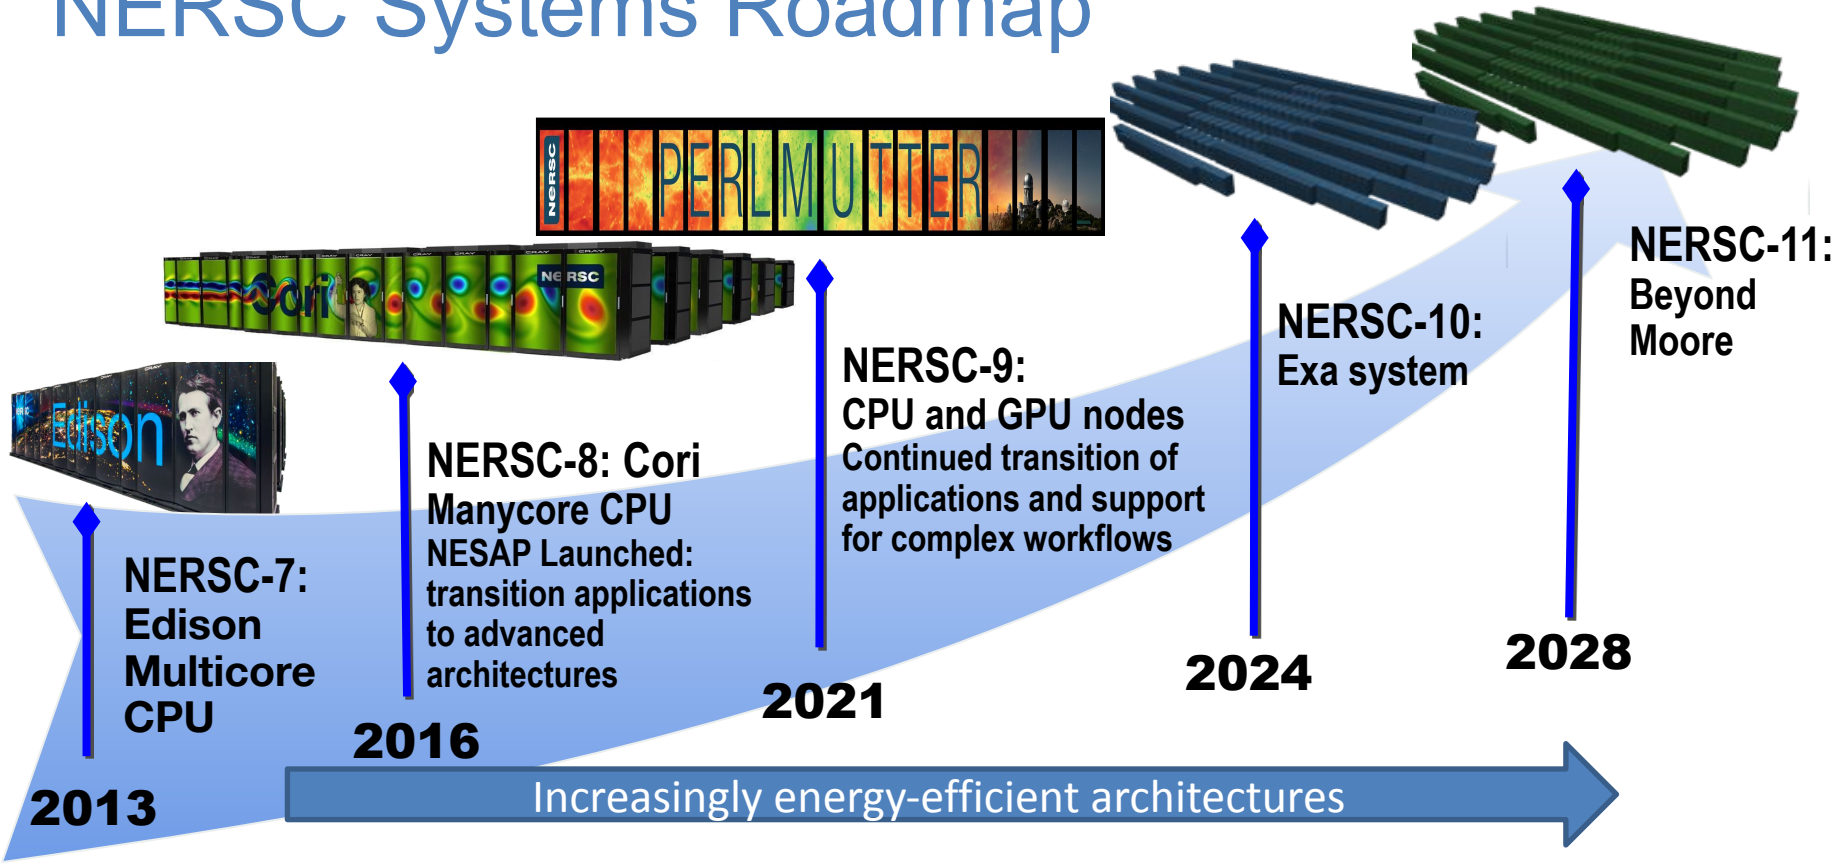

NERSC Systems 2022

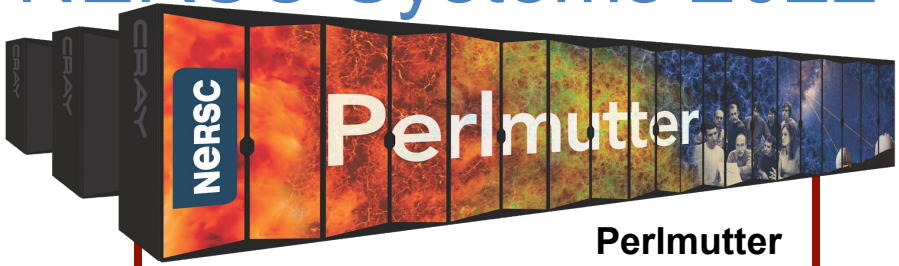

Perlmutter

1,500 NVIDIA A100x4 accelerated nodes
3,000 AMD dual-socket "Milan" CPU nodes
384 TB (CPU) + 240 TB (GPU) memory
HPE Cray Slingshot high speed interconnect
Debuted as World's 5th most powerful system
140 PF Peak

5 TB/s

35 PB Scratch

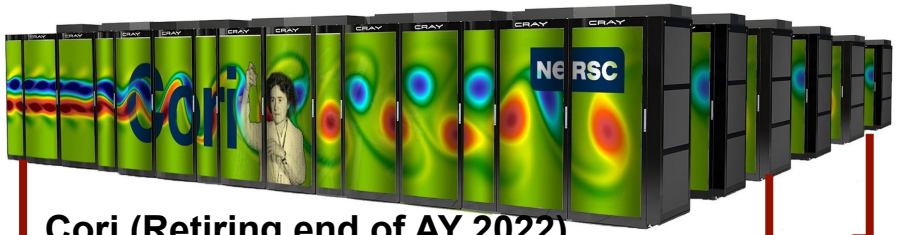

Cori (Retiring end of AY 2022)

9,600 Intel Xeon Phi "KNL" manycore nodes
2,000 Intel Xeon "Haswell" nodes
700,000 processor cores, 1.2 PB memory
Cray XC40 / Aries Dragonfly interconnect
30 PF Peak

1.5 TB/s

2 PB Burst Buffer

700 GB/s

28 PB Scratch

50 GB/s

HPSS Tape Archive ~200 PB

DTNs, Spin, Gateways

ESnet ENERGY SCIENCES NETWORK

2 x 100 Gb/s SDN

Ethernet & IB Fabric

Science Friendly Security
Production Monitoring
Power Efficiency

LAN

100 GB/s

120 PB Common File System

5 GB/s

275 TB /home

BERKELEY LAB Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY | Office of Science

HPC Systems: Cori (Retiring end of AY 2022)

Haswell nodes:

- For throughput
- Queues allow single-core jobs
- Longer walltime limits for smaller jobs
- Long queues

KNL nodes:

- For performance
- Codes should exploit many-core architecture
- Large jobs encouraged; discount for jobs using ≥1024 nodes
- 4x larger than Haswell partition
- Shorter queues
- Flex queue increases throughput & offers substantial discount

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/

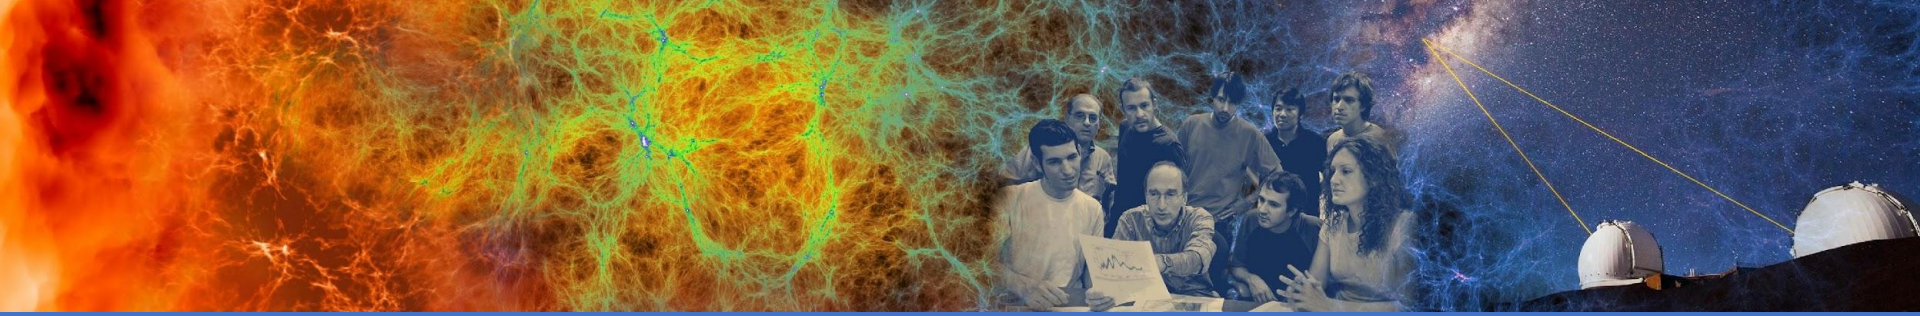

Software

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

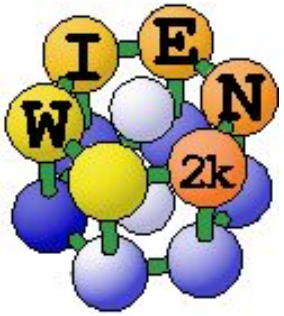

Molpro

QUANTUMESPRESSO

BerkeleyGW

abinit.

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

siesta

NAMD
Scalable Molecular Dynamics

LAMMPS

b-initio
VASP
ackage
imulation
lienna

GAMESS

WANNIER90

CPMD

NWCHEM
HIGH-PERFORMANCE COMPUTATIONAL CHEMISTRY SOFTWARE

NERSC

BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC has a rich data ecosystem!

data transfer and access

data management

data analytics

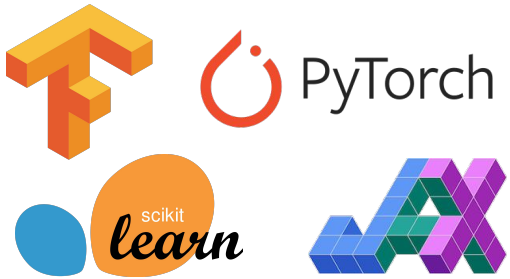

machine learning

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)

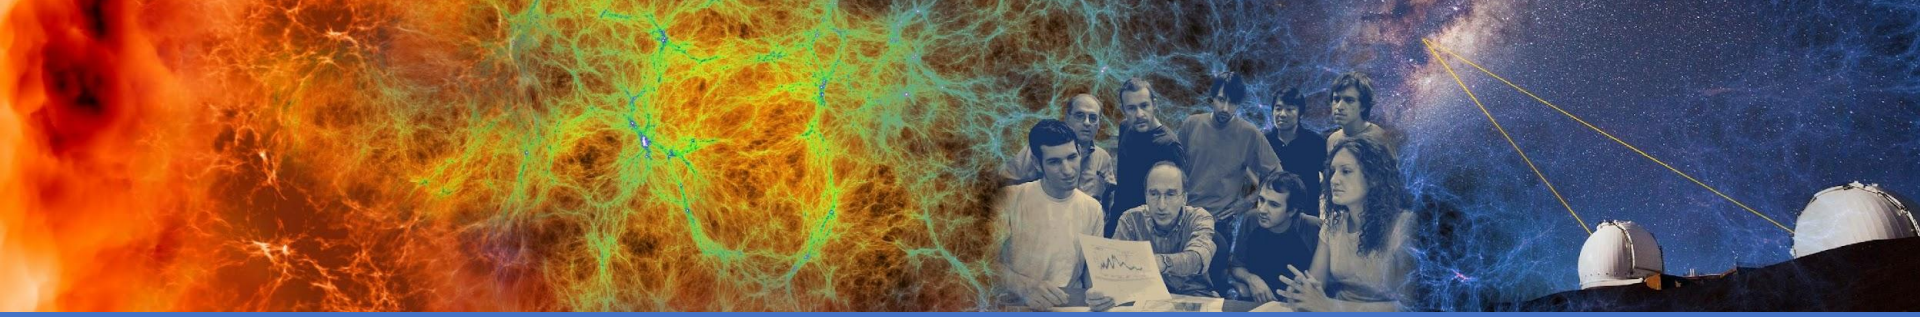

Interacting with NERSC

Interacting with NERSC

- NERSC Consulting & Account Support
  - User Tickets
  - User Appointments
  - User Training
- 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 2021, handled 7,229 tickets from 2,673 unique users

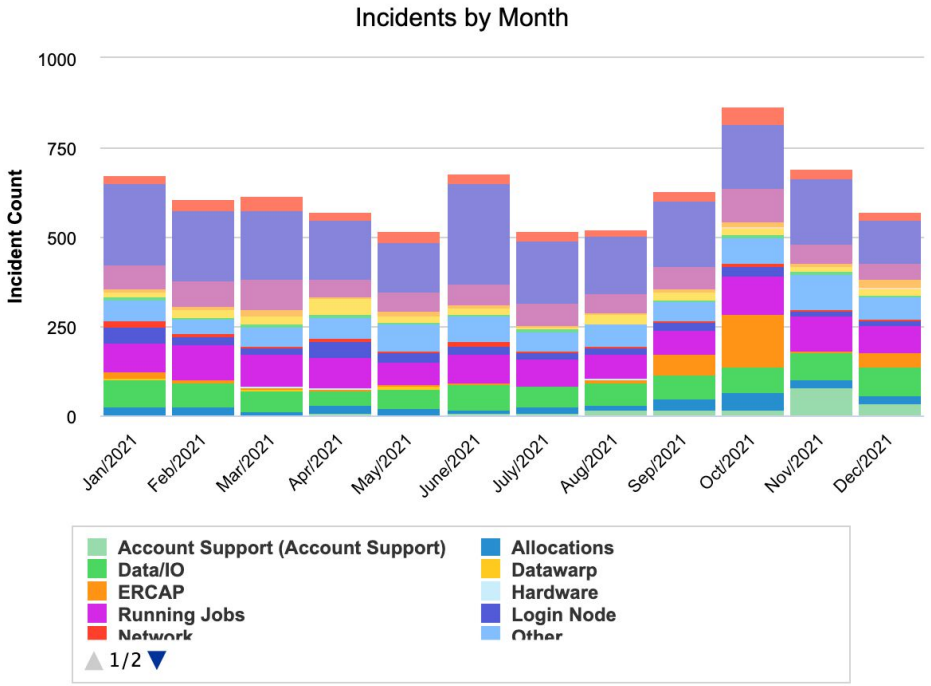

NERSC Consulting: Expectations

- 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
- Upcoming events of interest:
  - GPUs for Science Day (October 25)
  - Data Day (October 26-27)
- 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 online October 12-14, 2022

NERSC Community Engagement

- Special Interest Groups (SIGs) on topics including WRF users, scientific facilities
- Plans for building a cross-disciplinary NERSC users community of practice, expecting to begin in 2023
  - Initial target: graduate students
  - For more information, attend NUG Annual Meeting Friday session

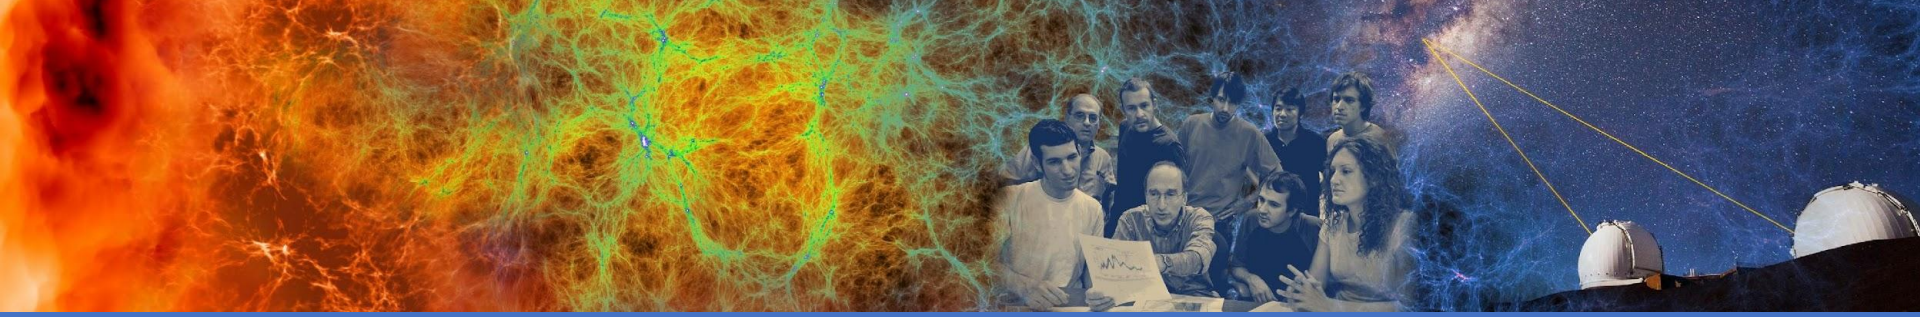

User Responsibilities & Expectations

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