

NERSC Overview




New User Training
June 16, 2020

Rebecca Hartman-Baker
User Engagement Group Leader
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

DOE View of NERSC Workload

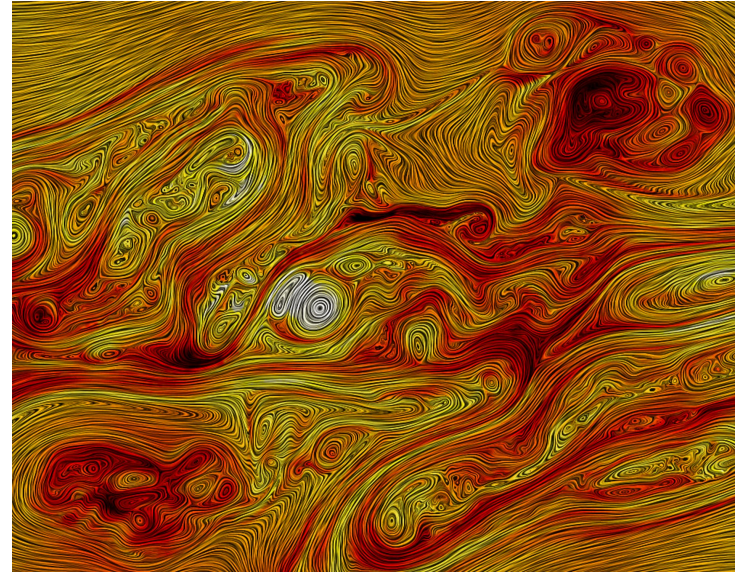
Percent of NERSC-Hours Used By Office in Allocation Year 2019



| | |
|------|--|
| ASCR | Advanced Scientific Computing Research |
| BER | Biological & Environmental Research |
| BES | Basic Energy Sciences |
| FES | Fusion Energy Sciences |
| HEP | High Energy Physics |
| NP | Nuclear Physics |
| SBIR | Small Business Innovation Research |

About NERSC

- Diverse workload:
 - 7000 users, 800 projects
 - 600 codes, 100s of users daily
- Allocations primarily controlled by DOE
 - 80% DOE Annual production awards (ERCAP)
 - O(10K)-O(10M) hour awards
 - Proposal-based, chosen by DOE program managers
 - 10% DOE ASCR Leadership Computing Challenge
 - 10% NERSC reserve




Turbulence in Solar Wind

Over 600 Codes Run at NERSC


- 10 codes make up 50% of workload
- 20 codes make up 66% of workload
- 50 codes make up 84% of workload
- Remaining codes (over 600) make up 16% of workload

Top codes at NERSC, Allocation Year 2018



Focus on Science

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



2018
nature
International weekly journal of science

14 in Nature
31 in Nature Comm.
82 in other journals




11 in
Science

31 in
PNAS

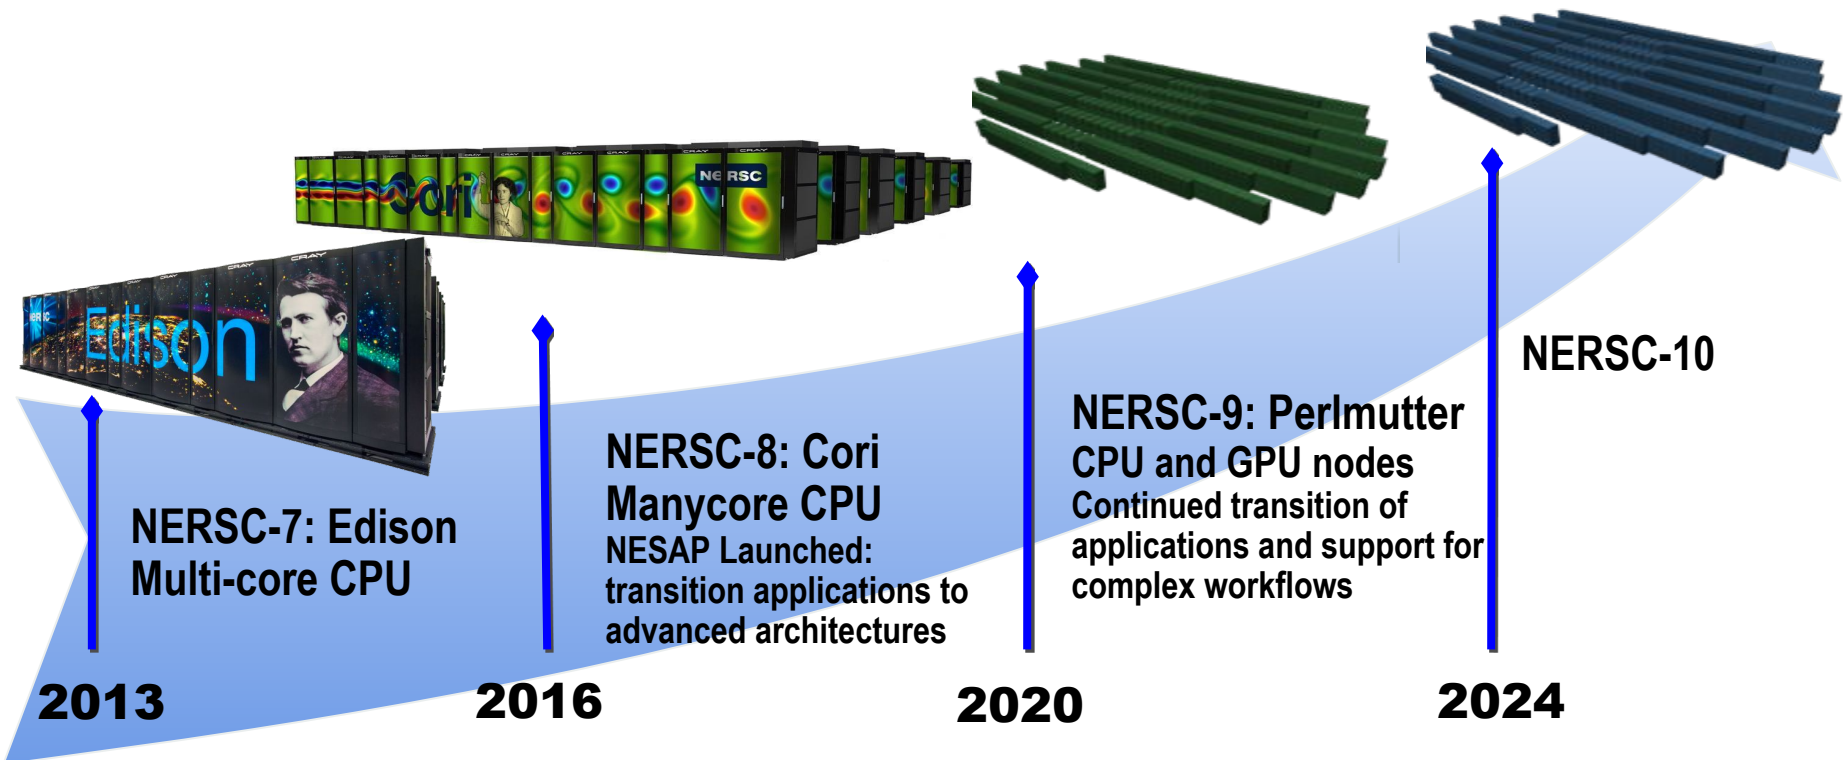


6 Nobel-prize
winning users




NERSC Hardware

NERSC Systems Roadmap



NERSC-9: A System Optimized for Science

- Cray Shasta System providing 3-4x capability of Cori system
- First NERSC system designed to meet needs of both large scale simulation and data analysis from experimental facilities
 - Includes both NVIDIA GPU-accelerated and AMD CPU-only nodes
 - Cray Slingshot high-performance network will support Terabit rate connections to system
 - Optimized data software stack enabling analytics and ML at scale
- All-Flash filesystem for I/O acceleration
- Robust readiness program for simulation, data and learning applications and complex workflows
- Phase 1 delivery in late 2020

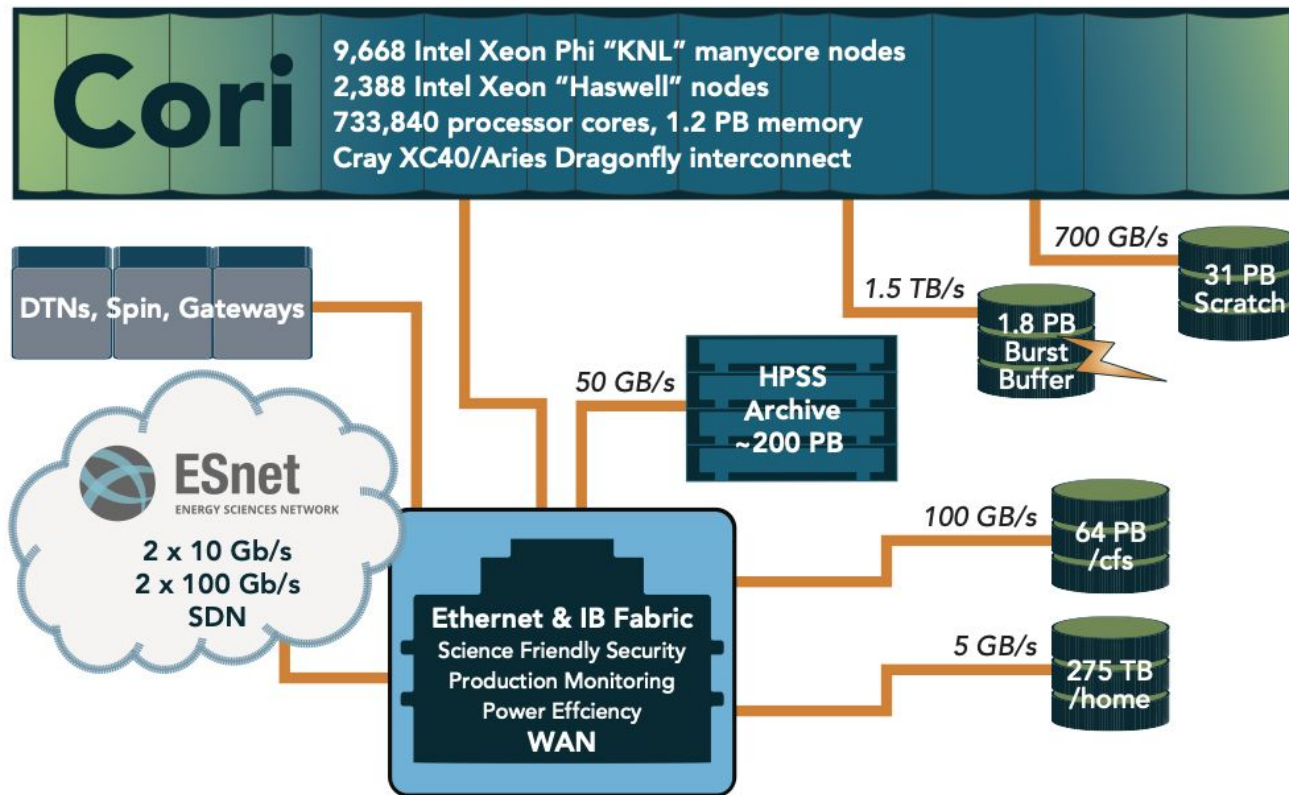


NERSC-9 Named after Saul Perlmutter

- Winner of 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 supercomputers to combine large-scale simulations with experimental data analysis
- Login “saul.nersc.gov”



NERSC Systems Map 2020



HPC Systems: Cori

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

File Systems

- Global File Systems:
 - Home
 - Community (CFS)
- Local File Systems:
 - Scratch
 - Burst Buffer
- 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
- Optimized for read/write operations, NOT storage
- Not backed up
- Purge policy (12 weeks)
- **Perfect for staging data and performing computations**

Burst Buffer

- Temporary per-job storage
- High-performance SSD file system
- Available on Cori only
- **Perfect for getting good performance in I/O-constrained codes**

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



QUANTUM ESPRESSO



BerkeleyGW

abinit.

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



NAMD

Scalable Molecular Dynamics

b-initio
VASP
package
simulation
lenna

GÅMESS

CPMD

LAMMPS

WANNIER90




NWCHEM

HIGH-PERFORMANCE COMPUTATIONAL
CHEMISTRY SOFTWARE



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
- NERSC Operations
- NERSC User Group (NUG)

NERSC Consulting & Account Support Team



NERSC Consulting & Account Support

- The first people you interact with when submitting a ticket or calling
- In 2019, handled 7,825 tickets from 2,709 unique users



| | |
|----------------------------------|-------------------------------|
| Account Support = 1,503 (21.41%) | Running Jobs = 1,175 (16.74%) |
| Software = 1,115 (15.88%) | Login Node = 983 (14%) |
| Data/IO = 831 (11.84%) | Other = 340 (4.84%) |
| Programming = 298 (4.25%) | Allocations = 286 (4.07%) |
| Performance = 157 (2.24%) | Network = 109 (1.55%) |
| Other (more ...) = 223 (3.18%) | |

NERSC Consulting: Expectations

- We will reply to you 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?

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 on August 17, 2020

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
 - Small jobs are great on Cori Haswell, not so good on Cori KNL
- 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!

