



NERSC

NERSC Overview



Rebecca Hartman-Baker

User Engagement Group Leader

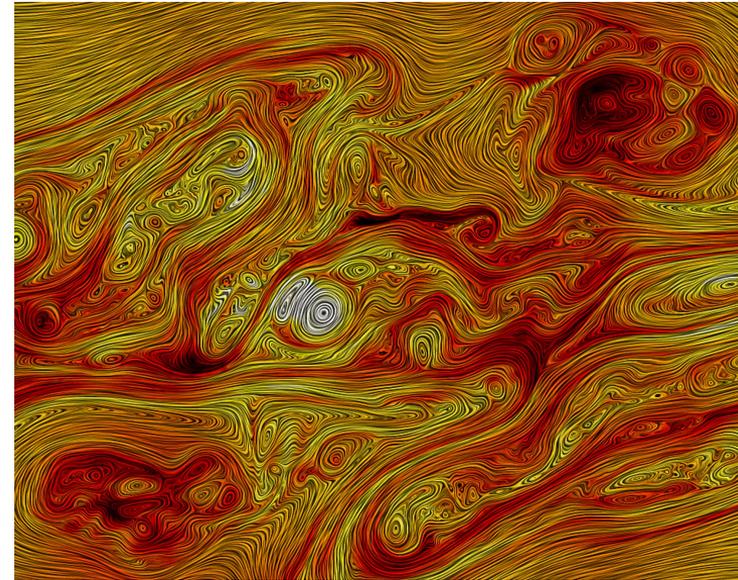
NERSC New User Training

January 25, 2019

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

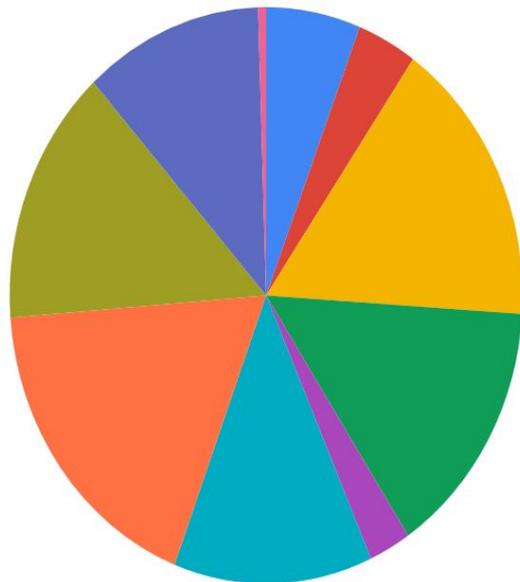
- National Energy Research Scientific Computing Center
 - Established 1974, first unclassified supercomputer center
 - Original mission: to enable computational science as complement to magnetically controlled plasma experiment
- Today's mission: ***Accelerate scientific discovery at the DOE Office of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

- 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

Initial Allocation of Hours by DOE Office of Science Program Office 2017



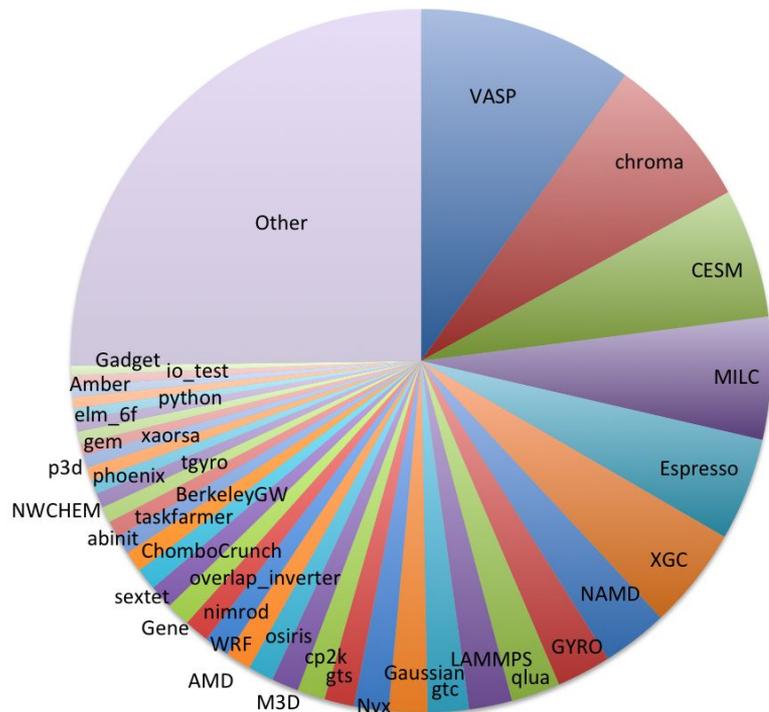
- ASCR
- BER Biosciences
- BER Climate/Env
- BES Chemistry/Geo
- BES Facilities
- BES Materials
- FES
- HEP
- NP
- SBIR

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

Over 600 Codes Run at NERSC



NERSC 2015 Code Usage



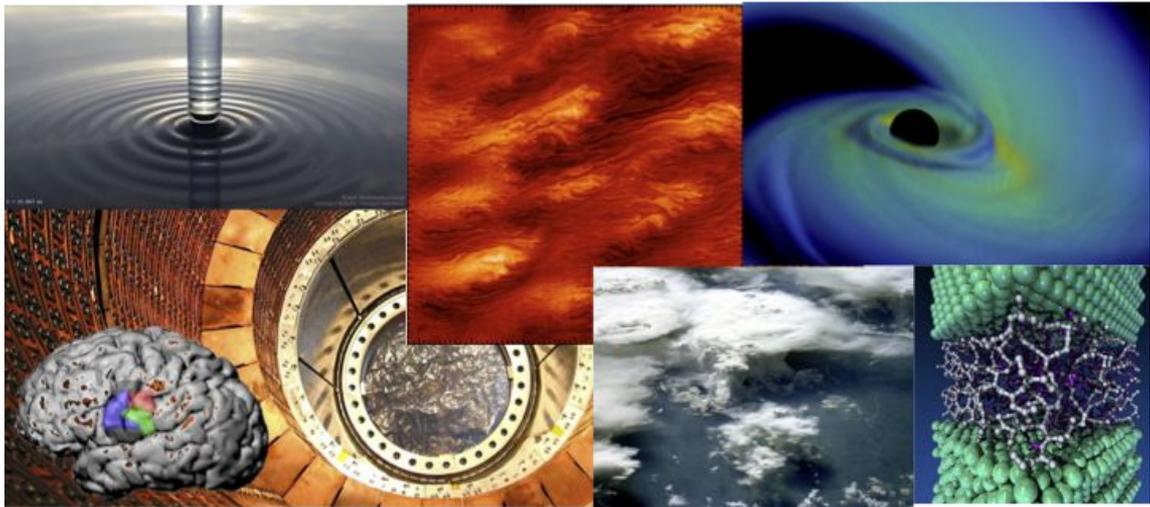
- Ten codes make up 50% of NERSC workload
- 25 codes make up 66% of NERSC workload

Focus on Science



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

1,300 citations via Web of Science in 2017 (underestimate!)



2017



8 in Nature
40 in Nature Comm.
118 in 12 journals



9 in Science

15 in PNAS



6 Nobel-prize winning users



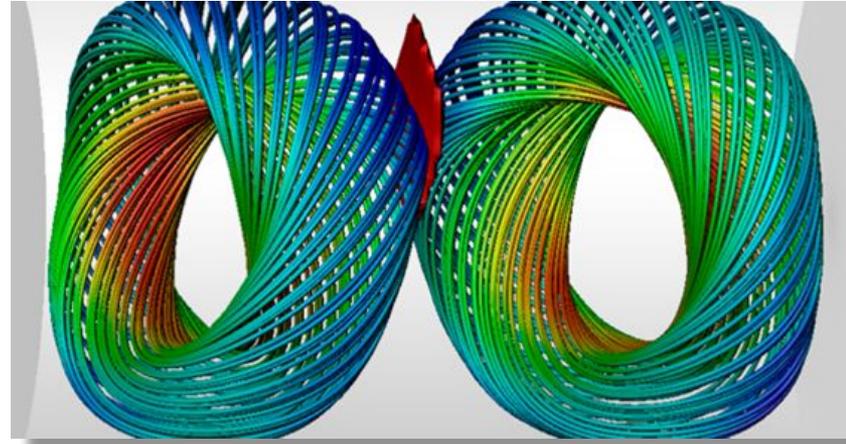
* as far as we can tell



Help NERSC Help You!

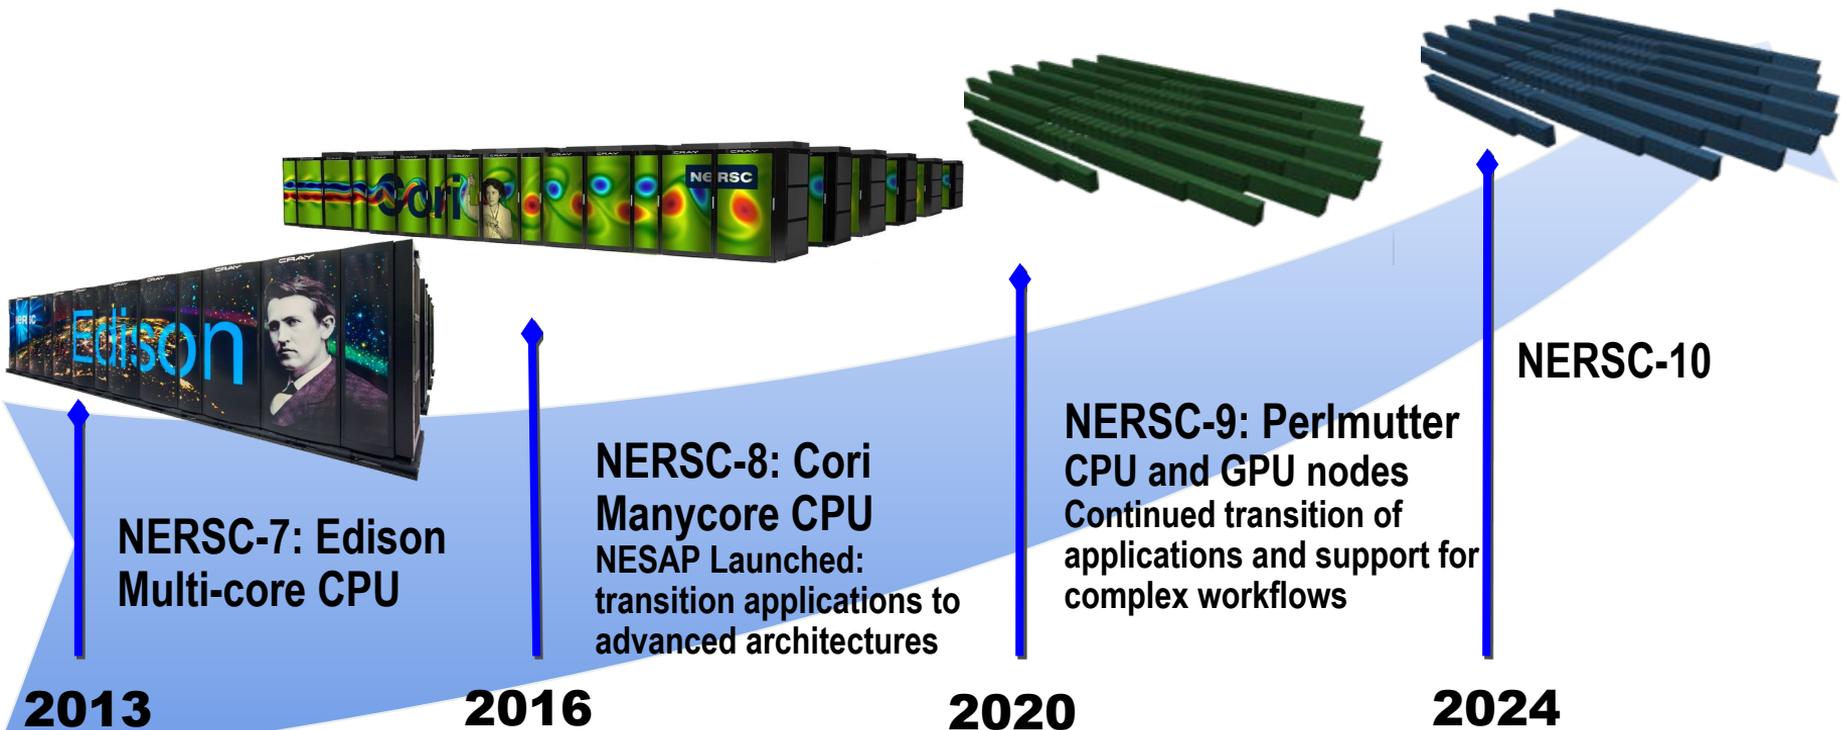


- Be sure to acknowledge NERSC in publications!
 - This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.
 - Available at:
<https://www.nersc.gov/users/accounts/user-accounts/acknowledge-nersc/>
- Science highlights sent to DOE each quarter
 - Please send us links to your publications!



***Magnetic field lines from HiFi
simulations of two spheromaks.
NERSC repo m1255
Image courtesy of Vyacheslav Lukin
(NRL)***

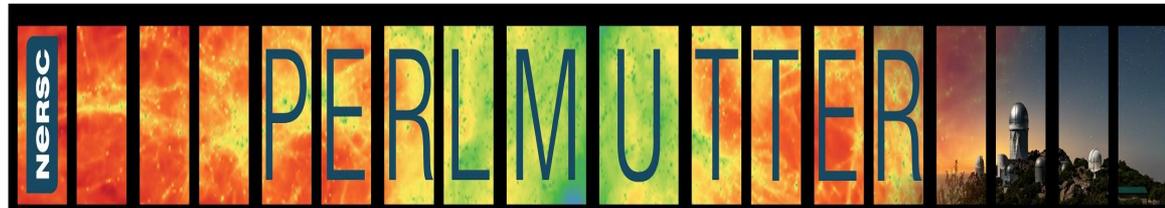
NERSC Systems Roadmap



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**
- **Delivery in late 2020**

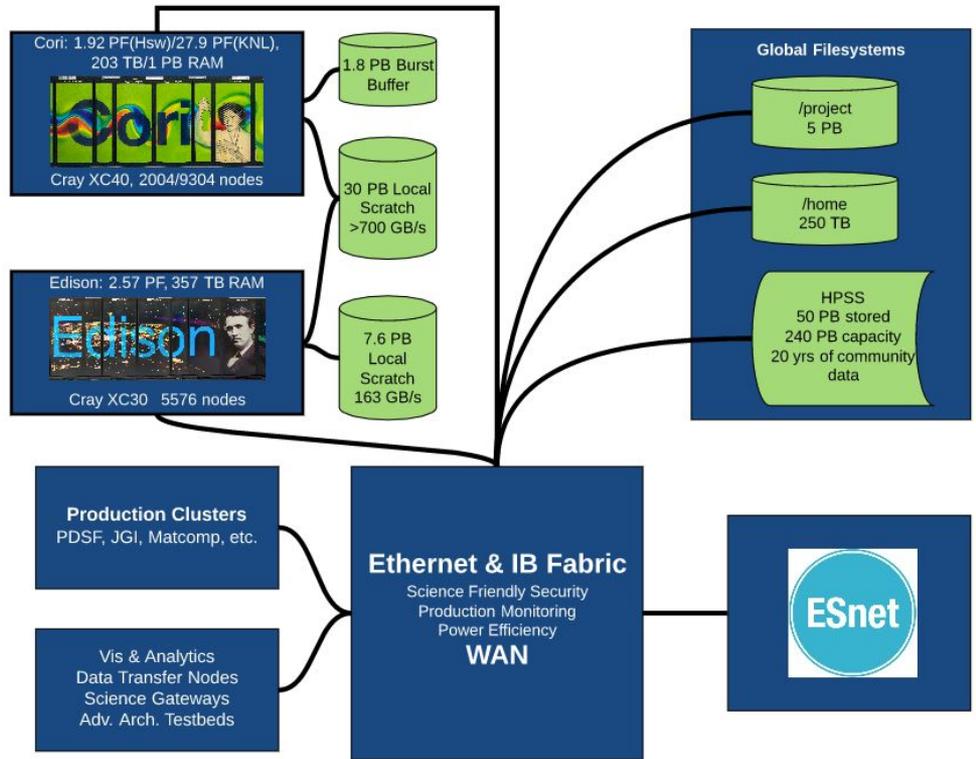


NERSC-9 will be named after Saul Perlmutter

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



NERSC Systems Map



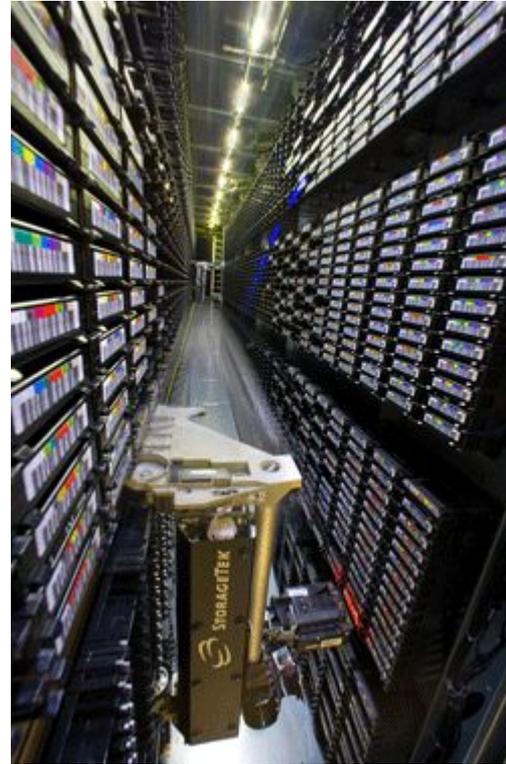
Edison:

- Large, stable machine
- Shorter queues
- Lower charge factor
- Retiring March 31, 2019

Cori:

- 2 types of nodes: Haswell and KNL
- *Haswell nodes*: for throughput
 - Queues allow single-core jobs
 - Longer walltime limits for smaller jobs
- *KNL nodes*: for performance
 - Codes should exploit many-core architecture
 - Large jobs encouraged

- Global Filesystems:
 - Home
 - Project
- Local Filesystems:
 - Scratch
 - Burst Buffer
- Long-term Storage System:
 - HPSS



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**

Project

- 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**

Scratch

- Large, temporary storage
- Local on Edison, Cori scratch also mounted on Edison
- 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**

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 Filesystems (1)

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



Using NERSC Filesystems (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 filesystem



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

Using NERSC Filesystems (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!)



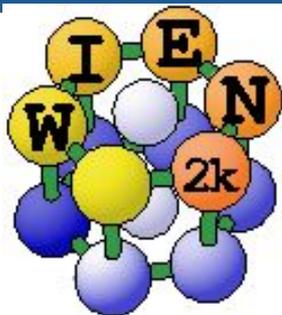
Queen cakes cooling on a wire rack by James Petts

<https://www.flickr.com/photos/14730981@N08/13475333725/>

- 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 Applications

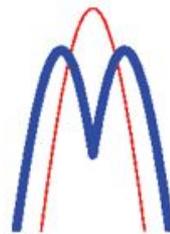
NERSC



BerkeleyGW



QUANTUM ESPRESSO



Molpro

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

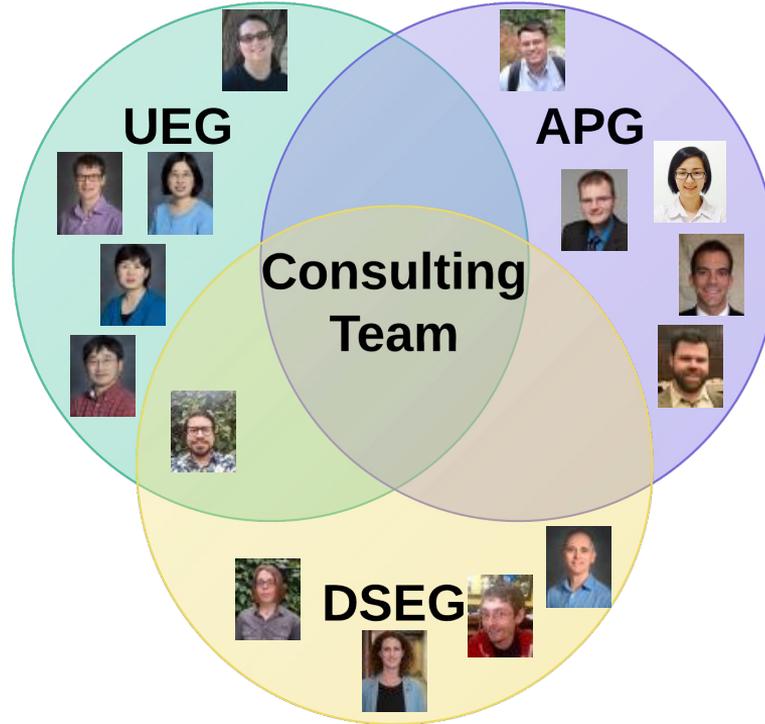
WANNIER90



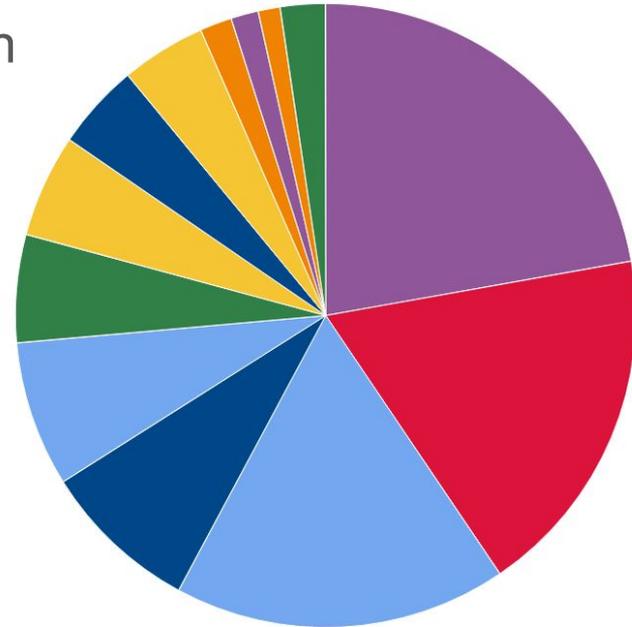
NWCHEM

HIGH-PERFORMANCE COMPUTATIONAL
CHEMISTRY SOFTWARE

- NERSC Consulting
- NERSC Operations
- NERSC User Group (NUG)



- The first people you interact with when submitting a ticket or calling
- In 2018, handled 7194 tickets from 2350 unique users



Account Support = 1,598 (22.21%) Running Jobs = 1,317 (18.31%) Software = 1,248 (17.35%) Data/IO = 589 (8.19%) Other = 543 (7.55%)
ERCAP = 402 (5.59%) Allocations = 385 (5.35%) Login Node = 320 (4.45%) Programming = 318 (4.42%) Performance = 120 (1.67%)
Data/IO = 103 (1.43%) Network = 82 (1.14%) Other (more ...) = 169 (2.35%)

- 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

- 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?

- Operations staff are on site 24/7/365 to supervise operation of the machine room
- Operations staff answer phone and forward to consultants during business hours if necessary
- Operations know the health of the machines and can help users with some tasks (password reset, killing jobs, changes to running reservation, etc.)

- 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)

- 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!

