

View from NERSC

*Challenges and opportunities for
increasing the impact of DOE
libraries and tools in collaboration
with the facilities*



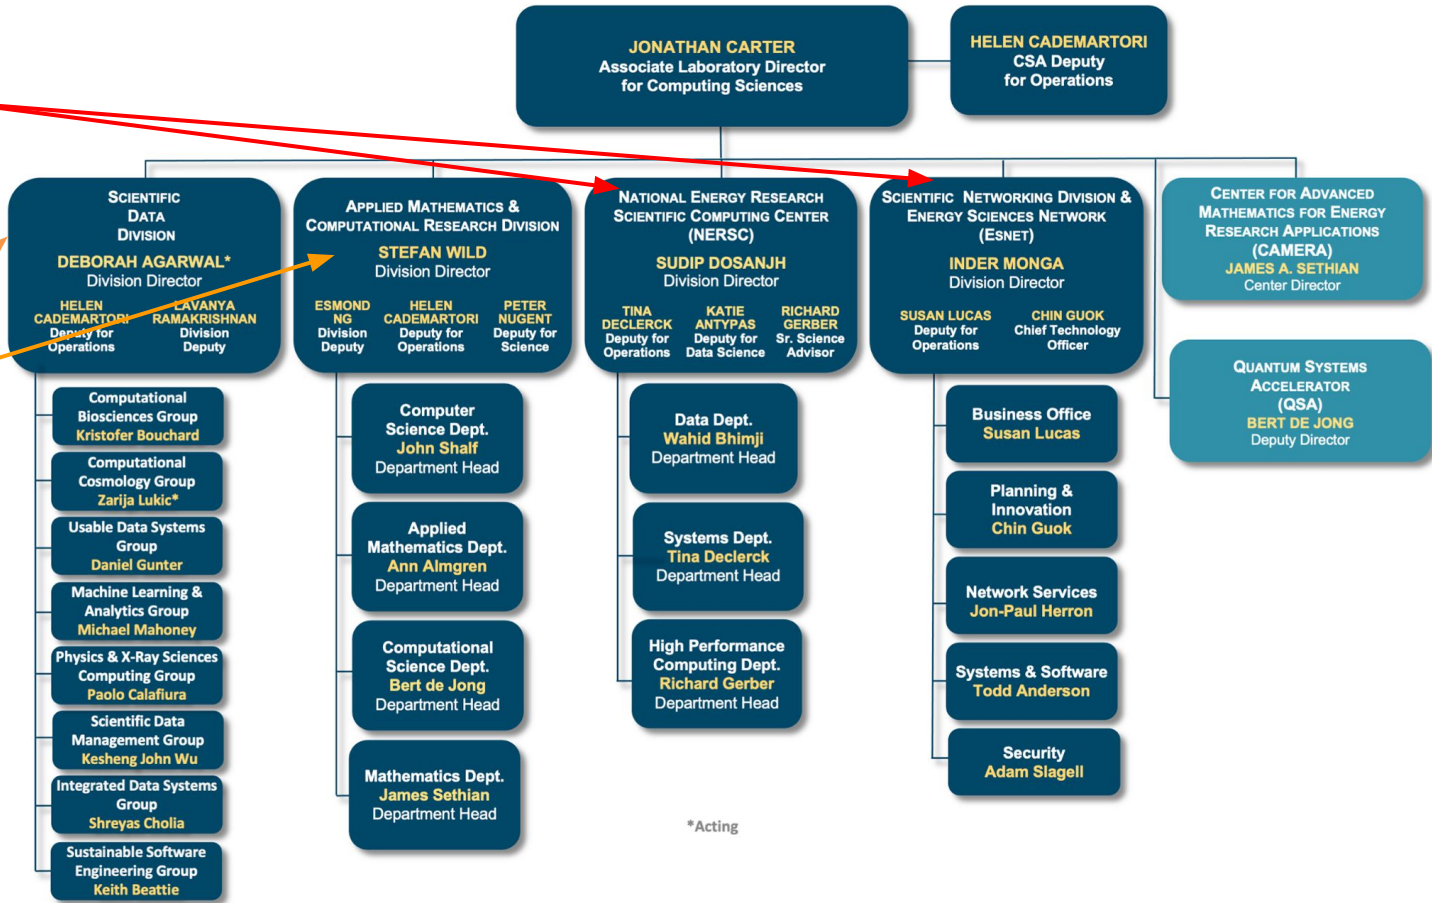
PESO Community Workshop
Argonne National Laboratory

Richard Gerber
HPC Department Head
Senior Science Advisor
June 8, 2023

Berkeley Lab Computing Sciences Area

DOE Office of Science
National User
Facilities

The
Computational
Research
Division (CRD)
split



*Acting




Mission HPC for Office of Science Research




U.S. DEPARTMENT OF
ENERGY

Office of
Science


Largest funder of physical science
research in the U.S.




Bio Energy, Environment




Computing




Materials, Chemistry, Geophysics



Particle Physics, Astrophysics



Nuclear Physics




Fusion Energy, Plasma Physics

9K users, 1K projects, 50 states, 40 countries, universities & national labs



Allocation Programs



Distributed by DOE Office of Science program managers

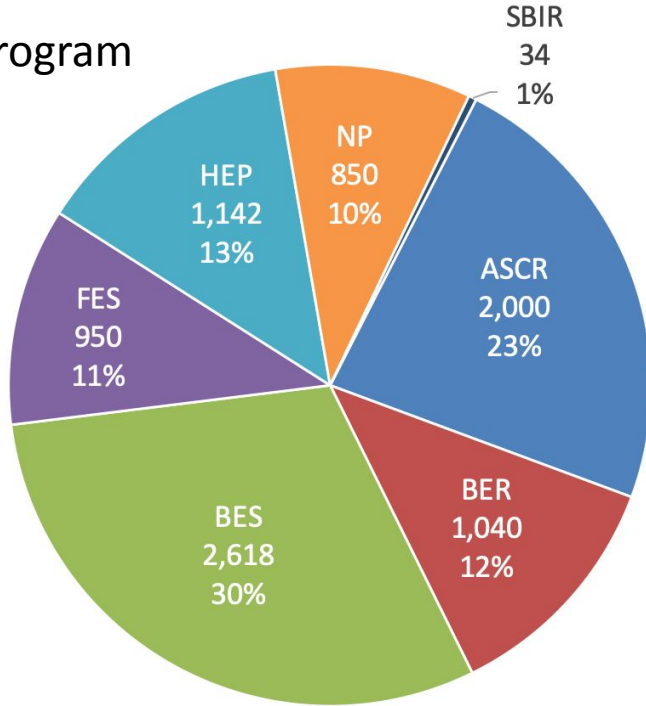
Competitive awards run by DOE Advanced Scientific Computing Research Office

Strategic awards from NERSC

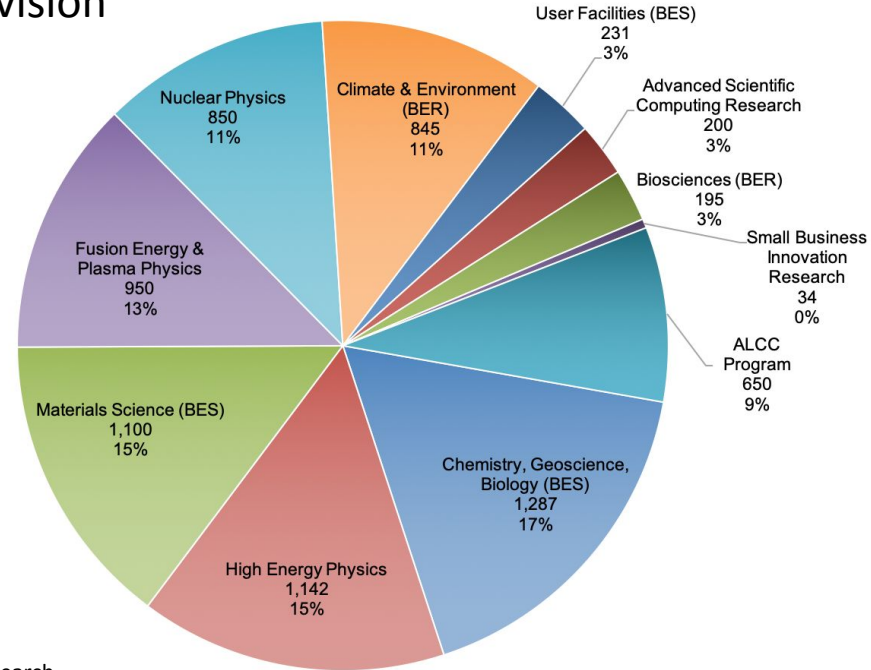
NERSC doesn't pick its users, so we have to meet them where they are and try to guide them towards increased productivity

NERSC Workload

By Program




By Division




ASCR - Advanced Scientific Computing Research
 BER - Biological and Environmental Research
 BES - Basic Energy Sciences (Chemistry, Materials, Geophysics)
 FES - Fusion Energy Sciences (including Plasma Physics)
 HEP - High Energy Physics
 NP - Nuclear Physics
 SBIR - Small Business Innovation Research





A Diverse Scientific Community of Users



Breakdown of Compute Used by DOE Program




~9,000 ANNUAL USERS FROM ~800 Institutions + National Labs



>2,000 Scientific Journal Articles per Year

Perlmutter




- 1,792 GPU Accelerated Nodes
 - 4 NVIDIA A100 per node
 - 7,168 GPUs
 - 1 AMD “Milan” per node
- 35 PB FLASH scratch filesystem
- 3,072 CPU only nodes
 - 2 AMD “Milan” per node
 - 512 GB memory per node
 - 393K cores
- HPE Cray “Slingshot 11” network

Levels of Support and Usage Models

- Our users
 - Develop and use their own codes / Use community codes
 - Small / Medium / Large teams
 - Independent / Need support
 - Low / Medium / High levels of sophistication
- We want to support the full workload
 - But, limited staff
 - Have to consider strategic importance and impact
 - But we also want to help projects and users develop capabilities

Many users run community applications, but we have a very large, long tail that wants to leverage a software infrastructure (2021 data)



NERSC Current User Software Support

- NERSC has a small team that installs and supports a limited number of applications
- Same for software libraries and tools
 - We try to fill the holes left by vendors for oft-requested libraries and tools
 - Have had a lot of success getting vendor support through procurements; future is less promising
- The E4S suite is installed and is being used (Shahzeb!)
- A sustainable and curated software collection has the potential to bring great value to both users and staff
 - Any coordinated effort should also seize the opportunity to create a way to enable a quantitative measure of usage and impact for each package
- We're happy to entertain other enablement models!

We'd be happy to have vendors optimize, install, and support software

Provides clear support model

Reduces staff effort

Vendors have optimization experts for their architecture

But ... may not be agile or as responsive as desired. Versions usually lag.

What Software Do Users Want or Need?

- They evaluate and use what works for them given their optimization targets
 - Cost? Time to get a research result? Capability? Volume of computation? Effort consistent with available personnel resources? Raw performance?
- When choosing software, users' considerations include
 - Functionality and correctness
 - Ease of use
 - Performance
 - Portability and sustainability
 - Documentation and support through software developers or community
 - Ease of Integration into code base/workflow
- Except for the rather obvious packages, it is difficult for facilities to say what users "need" or even "want"
- Many users are open to using new software and methods if they offer advantages
- Would a "standard" HPC user software environment become widely adopted?

What NERSC Considers When Deciding to Support SW

- Level of impact - key enabler of priority science or broad enabler across the board
- Ease of integration
- Forward looking
- Performance - optimization for architecture
- Support model and level of support we have to provide
- Cost
- Impact on system and other jobs
- Security
- Export control and user agreements
- Licenses

Vision of User Support Documentation

<https://docs.nersc.gov/applications/berkeleygw/>