

Nersc User's Group (NUG) Monthly Meeting



18 January, 2024

Today's Pipeline



- Interactive - please participate!
 - [NERSC User Slack](#) (link in chat)
 - **#webinars** channel
- Agenda:
 - Introductions - Charles Lively and Lipi Gupta
 - New Year ICE Breakers - NERSC Edition
 - General NERSC Updates/Announcements
 - Calls for Participation and Upcoming Scheduled NERSC Trainings
 - Topic(s) of the Day:
 - Reminder NERSC Allocation Year Transition Tips
 - Focusing on Power and Energy Efficiency at NERSC (Part 1)
 - Discussion

User Engagement Group - Our People




Lisa Claus



Kevin Gott



Lipi Gupta



Rebecca Hartman-Baker
UEG Group Lead

Alumni:

Tiffany Connors
Zhengji Zhao
Steve Leak
Erik Palmer
Justin Cook
Shahzeb Siddiqui



Helen He




Charles Lively



Kelly Rowland








Kadidia Konate



Woo-Sun Yang

NUG Breakers



NUG Breakers (1)



- What is one way you plan to approach your research differently for 2024?









NUG Breakers (2)



What is one science domain that you're more interested in exploring and learning about in 2024?



Announcements



Announcements - NERSC Updates



- (NEW/UPDATED) Allocation Year 2024 Began Yesterday-January 17, 2024!
 - Reminders to be discussed later
- Save the Date: Celebrate 50 Years of NERSC with Us October 22-24, 2024
 - Celebrations throughout the year and during the Annual NUG Meeting at NERSC
- Changes to Science Gateways using PHP at portal.nersc.gov
- Submit a Science Highlight Today!

Reminder: Please see the Weekly Email for Links!

Announcements - Calls for Participation



- Nominate a Candidate for Rising Stars Workshop (by Jan 28, 2024)
 - held April 30-May 1, 2024 at the Oden Institute in Austin, Texas.
- Apply Today for April SYCL HackFest
 - one-day hacking session with SYCL to be held April 8, 2024
 - Submissions are due February 2, 2024!

Reminder: Please see the Weekly Email for Links!

Announcements - Meetings/Trainings



- New User Training & Updated Best Practices on Perlmutter, February 15-16
 - NERSC is hosting a two half-day online training event for new and existing users
- Join Us for Data Day, February 21-22
 - NERSC is hosting a hybrid two-day event on February 21-22, 2024
- Perlmutter and Data Day Office Hours, February 23
- Argonne Training Program on Extreme-Scale Computing by February 28
- Energy Exascale Earth System Model (E3SM) Tutorial Workshop
 - Applications Now Open for E3SM Tutorial Workshop (due Jan 31 for in-person attendance)
 - First in-person tutorial workshop at NERSC, May 7-10, 2024

Reminder: Please see the Weekly Email for Links!

Announcements - Scheduled Outages









Please see the Weekly Email for Links!

- **HPSS Regent (Backup)**
 - 01/18/24 09:00-13:00 PST, Scheduled Maintenance
 - 01/03/24 09:00-13:00 PST, Scheduled Maintenance
 - Some retrievals may be delayed during tape library maintenance.
- **HPSS Archive (User)**
 - 01/24/24 09:00-13:00 PST, Scheduled Maintenance
 - Possible retrieval delays during tape drive firmware updates.
 - 01/31/24 09:00-13:00 PST, System Degraded
 - Some retrievals may be delayed during tape library maintenance.
- **Spin**
 - 01/19/24 10:00-16:00 PST, Scheduled Maintenance
 - Rancher 2 workloads and the Rancher 2 UI will be unavailable briefly (1-2 min) at least once within the window for upgrades to system software.

<https://www.nersc.gov/live-status/motd/>

Reminders: AY24 Allocation Year Transition



Outline



- On the AY2024 Start Day
- Changes in AY2024
- Discontinued Users

On the AY2024 Start Day, Jan 17 (1)



- **Iris:** Downtime: 7 am - 9:30 am
 - Logout and login to reflect new AY data
- **Perlmutter:** Scheduled maintenance 6 am - 10 pm
 - Login and compute nodes availability info will be provided later
- **Jupyter**
 - Access may be impacted depending on Perlmutter status, such as login and compute nodes availability
- **All other systems and services are up**

These jobs will be deleted on Perlmutter

- Jobs associated with non-continuing projects
- Jobs associated with a continuing project that the user is no longer a member of AY2024
- "overrun" jobs
- Held jobs older than 12 weeks

- **PIs should check/update CPU and GPU allocations allowed for each user in Iris once AY24 starts**
 - Percentages allowed for each user are **inherited** from last year
 - Hours allowed for each user are **NOT inherited** from last year
 - The reason behind this is some users have a huge number inherited previously from before the unit was based on Perlmutter hours
 - Percentage takes precedence over hours if both are set

New Appropriate Use Policy



- NERSC's Appropriate Use Policy is being updated for AY24 to better reflect the way that NERSC resources are used by the NERSC user community.
- All users are required to agree to the updated policy
- More details will be provided later for the process

(New) Updated Software Default in Programming Environment Since Jan 17



Expect most users should not need to re-compile

- CPE upgraded to 23.12
 - Default cudatoolkit module version updated to cudatoolkit/12.2
 - gcc compiler is now provided by the underlying SLES OS
 - Users wishing to access the older programming environment can load the `cpe/23.03` module
- NERSC module updates
 - openmpi: added 5.0.0
 - llvm: added 17.0.6, removed 17.0.2
 - conda: added Miniconda3-py311_23.11.0-2
 - python: updated environments for compatibility with cpe/23.12
- More details:

https://rest.nersc.gov/REST/announcements/message_text.php?id=4890

(New) Default Inode Limit in CFS









- Default inode quota for a project on CFS decreased from 20 M to 10 M.
- No existing user impacts from this change

Discontinued Users



- Users with no active project are “discontinued”, effective AY2024 starts on Jan 17, 2024
- Can login to authorized systems (Perlmutter, HPSS, DTN, etc.) for 60 days until Mar 16, 2024 for data access, but can not run batch jobs
 - Clean up and transfer files back to their home institutions
- Detailed policy at <https://docs.nersc.gov/accounts/policy/#account-deactivation-process>

Focusing on Power and Energy Efficiency at NERSC (Part 1)



Charles Lively III, PhD - UEG

- **DOE's Mission and NERSC**
- **Why is Power and Energy Efficiency Important?**
- **Basics of Power/Energy Consumption**
- **Understanding Power and Energy Consumption in HPC**
- **Ramifications for HPC Users**

Outline-Power and Energy Consumption in HPC

- NERSC and DOE's Mission
- Why is Power and Energy Efficiency Important?
- Basics of Power/Energy Consumption
- Understanding Power and Energy Consumption in HPC
- Ramifications for HPC Users

Outline-Power and Energy Consumption in HPC



- NERSC and the DOE's Mission
- Why is Power and Energy Efficiency Important?
- Basics of Power/Energy Consumption
- Understanding Power and Energy Consumption in HPC
- Ramifications for HPC Users


NERSC: Mission HPC for DOE 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



- **Energy Efficiency is vital to the DOE for several reasons**
 - **Cost Savings**
 - Enables the DOE to save cost without increasing resource usage
 - more resources can be allocated to actual scientific research rather than operational costs.
 - **Enabling Larger and More Complex Simulations**
 - allows for larger and more complex computational tasks to be undertaken.
 - **Advancing Technological Leadership**
 - energy efficiency improvements, this would lead to an unsustainable increase in energy consumption.
 - limited energy infrastructure, energy-efficient computing is crucial.
 - developed for energy efficiency in supercomputing often have broader applications in other areas of technology and industry.
 - **Supporting Environmental Sustainability**
 - Promoting energy efficiency in supercomputing aligns with this goal by reducing the carbon footprint of these large-scale computing facilities.

- **Power and Energy Consumption Affect HPC Centers**
 - **Costs**
 - energy costs can account for up to 50% of the total operational costs for large data centers.
 - electricity cost for running a supercomputer can be several million dollars per year.
 - **Environmental Impact**
 - Data centers globally are estimated to consume about 1-3% of the world's electricity as of 2023, contributing significantly to carbon emissions.
 - **Scalability**
 - energy efficiency improvements, this would lead to an unsustainable increase in energy consumption.
 - limited energy infrastructure, energy-efficient computing is crucial.

- **Definition of Power**

- Power consumption in computing is often estimated using $P = V \times I$, where V is the voltage in volts (V) and I is the current in amperes (A). This formula is crucial for designing and managing energy-efficient computing systems.

- **Definition of Energy**

- The energy consumption (E) in computers can be approximated by $E = P \times t$, where P is the power usage in watts (W) and t is the time in hours (h). This calculation helps in understanding the energy requirements of various computing tasks.
 - This measurement is essential in various fields, from physics to computing, for assessing energy utilization.

- **Limited Increases in Clock Speed for CPUs & GPUs Due to Power Constraints**
 - Reaching the physical limits of increasing clock speeds, power efficiency becomes a crucial factor.
 - highlights the need for innovative ways to enhance performance without significant power increases.
- **More Cores Within Power Budgets**
 - trend towards increasing the number of cores in CPUs and GPUs is influenced by power constraints.
 - adding more cores can lead to higher power consumption, there's a growing emphasis on designing energy-efficient multicore architectures.
- **Redefining Moore's Law in the Context of Power Efficiency**
 - End of Moore's Law in its traditional sense (doubling transistor count every two years) is now coupled with power efficiency challenges.
 - focus is shifting towards improving performance within power limitations, rather than just increasing transistor density.

Power Consumption affects Technology (2)

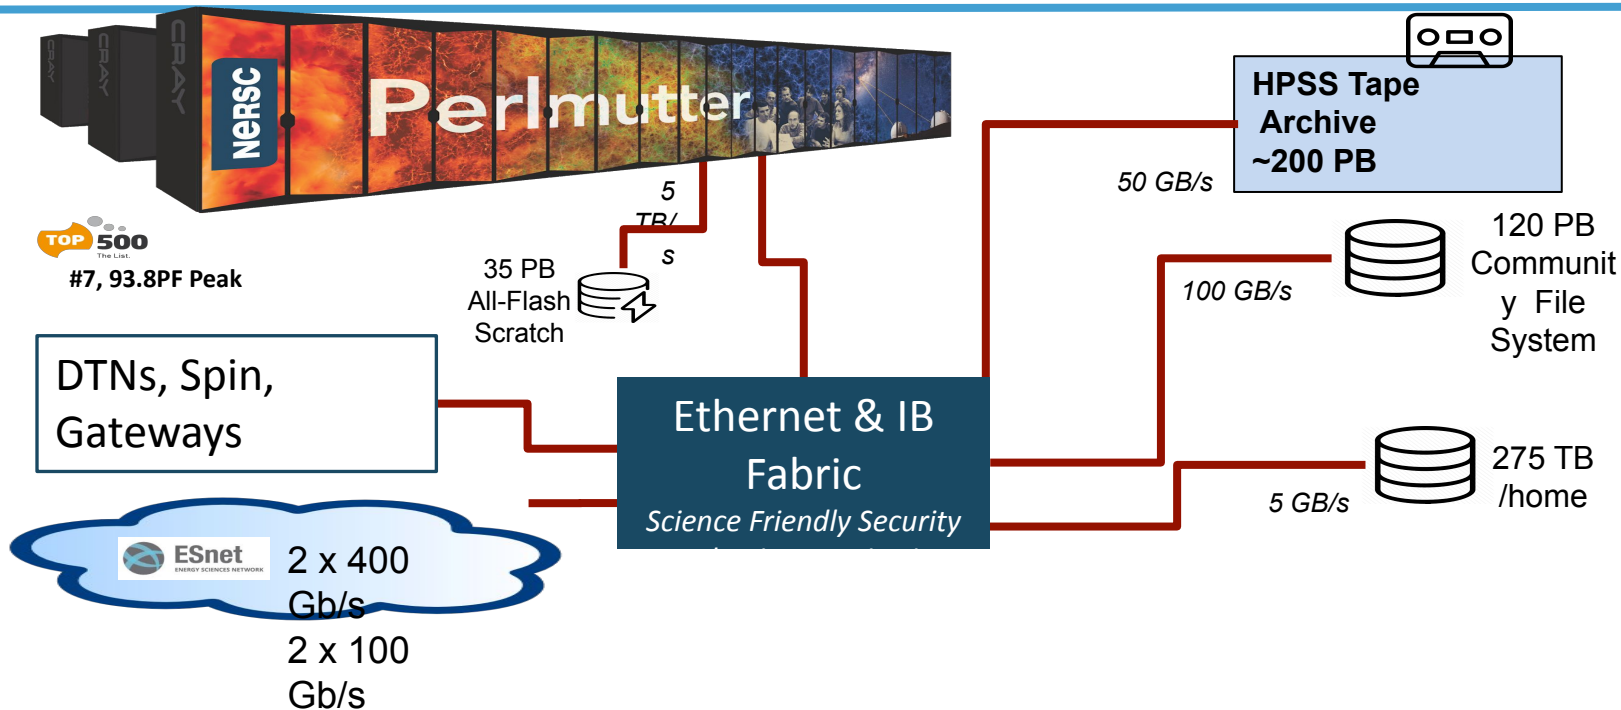


- **Doubling Performance per Socket Within Power Limits**
 - Power-Efficiency can improve performance per socket
 - Industry is looking to boost performance without proportional increases in power consumption.
- **Power Increases as a Driver for Performance Gains**
 - performance gains will be influenced by how effectively power is managed at both the socket and system levels.
 - necessitates advancements in power-efficient computing technologies.
- **Tighter CPU-GPU Integration for Power Efficiency**
 - Trends of tighter CPU-GPU integration,
 - Example: NVIDIA's Grace-Hopper and AMD's MI-300
 - driven by the need to optimize power usage while enhancing computational capabilities.
- **Flash Storage's Rise Aided by Power Efficiency**
 - Increase in flash storage capacity and its encroachment into HDD space is not just due to its speed and capacity advantages but also its power efficiency. As flash storage consumes less power than traditional HDDs, its growth aligns with the trend of reducing power consumption in computing systems.

Power consumption in HPC systems involves hardware performance, software efficiency, and operational demands.


- Components' Energy Requirements:
 - processors, memory, storage, and cooling systems in HPC.

NERSC Center Architecture



Full Perlmutter System Accepted Summer 2024



NERSC



- 1,792 GPU Accelerated Nodes
 - 4 NVIDIA A100 per node
 - 1 AMD “Milan” per node

- 3,072 CPU only nodes
 - 2 AMD “Milan” per node
 - 512 GB memory per node

Perlmutter Acceptance



- Direct liquid cooling
 - NERSC's Power Usage Effectiveness (PUE) is now 1.05-1.08
- HPE Cray "Slingshot 11" network
 - Quadrupled the GPU node bandwidth

Discussion -Power and NERSC Users



- How does the power consumption of HPC resources like Perlmutter influence the efficiency and scheduling of batch jobs?
- In what ways can optimizing power consumption lead to more effective utilization of HPC resources?
- What are the challenges in balancing high-performance computing needs with power limitations?
- How can users adapt their computational methods or algorithms to be more energy-efficient on HPC systems like Perlmutter?
- What tools or metrics are available to NERSC users to monitor and manage the power consumption of their batch jobs?
- How does the power efficiency of HPC systems impact the overall carbon footprint and sustainability goals of research institutions?
-

Coming up



Upcoming topics:

- Power/Energy Consumption Series @ NERSC
- Security @ NERSC
- Community needs/ideas (e.g. new groups/topics, “get to know” <blank>, new docs/training options, career?)
- Other topic suggestions/requests?

We'd love to hear more lightning talks **from NERSC users** about the research you use NERSC for!

Nominate a topic at: <https://forms.gle/WjYx7zV7SAz2CaYz7>

Science Highlights Submission:

<https://docs.google.com/forms/d/e/1FAIpQLScP4bRCtcde43nqUx4Zsz780G9HsXtpecQqIPKvGafDVVKQ/viewform>

Lightning Talks



Highlights



Links to Training Materials and Planning

- Training Upcoming Events Page
 - <https://www.nersc.gov/users/training/events/>
- [Training Past Events Page](#)
 - Can choose “Filter by Year”
- NERSC YouTube Channel
 - <https://youtube.com/c/NERSCTraining-HPC/>
- [Training Events Archive Page](#)
- [Training Materials Page](#)
- [Best Practices for NERSC Training Journal Paper, Slides](#)

We welcome any suggestions, comments, ideas on Training!!

Cheers to a Productive 2024!!!



Thank You