

Perlmutter Access, Allocations, & User Environment



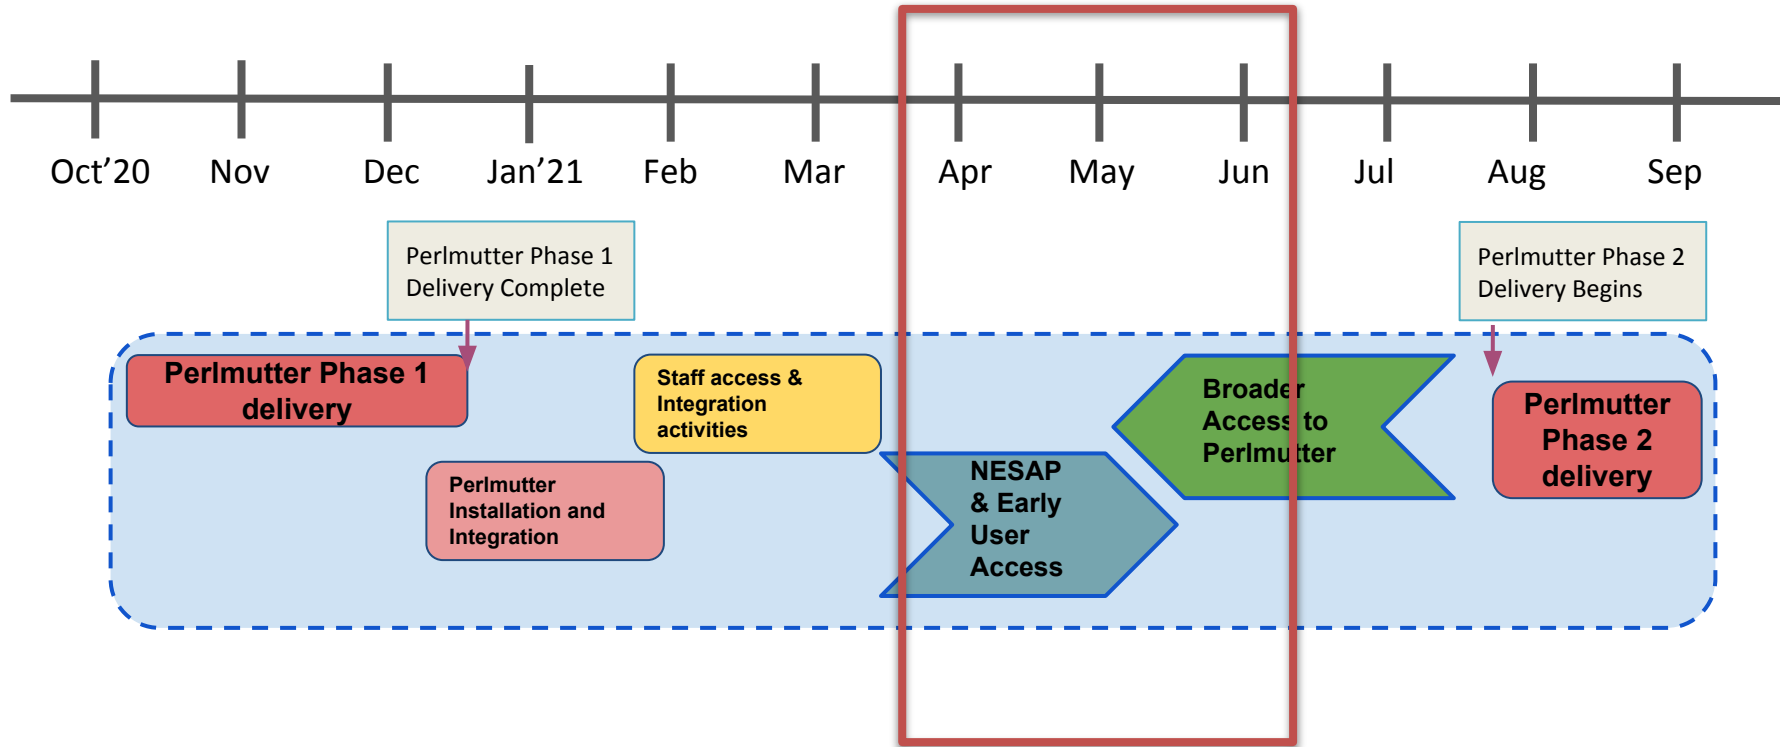
2020 NERSC User Group
Annual Meeting

Rebecca Hartman-Baker
User Engagement Group Lead
August 17, 2020

Outline

- Access & Allocations
 - Timeline review
 - Access & Allocations
- User Environment
 - Conceptual Architecture
 - Logging in
 - Programming environment
 - Software
 - Running jobs
- Learning Opportunities
 - Training
 - Office Hours
 - Hackathons

Perlmutter Phased Timeline



User Access to Perlmutter Phase I (GPU Nodes)

~ April 2021:

- All users given accounts to Phase I for code development and small-scale testing
- Priority access for large-scale testing and project milestones
 - NESAP teams
 - Exascale Computing Project teams

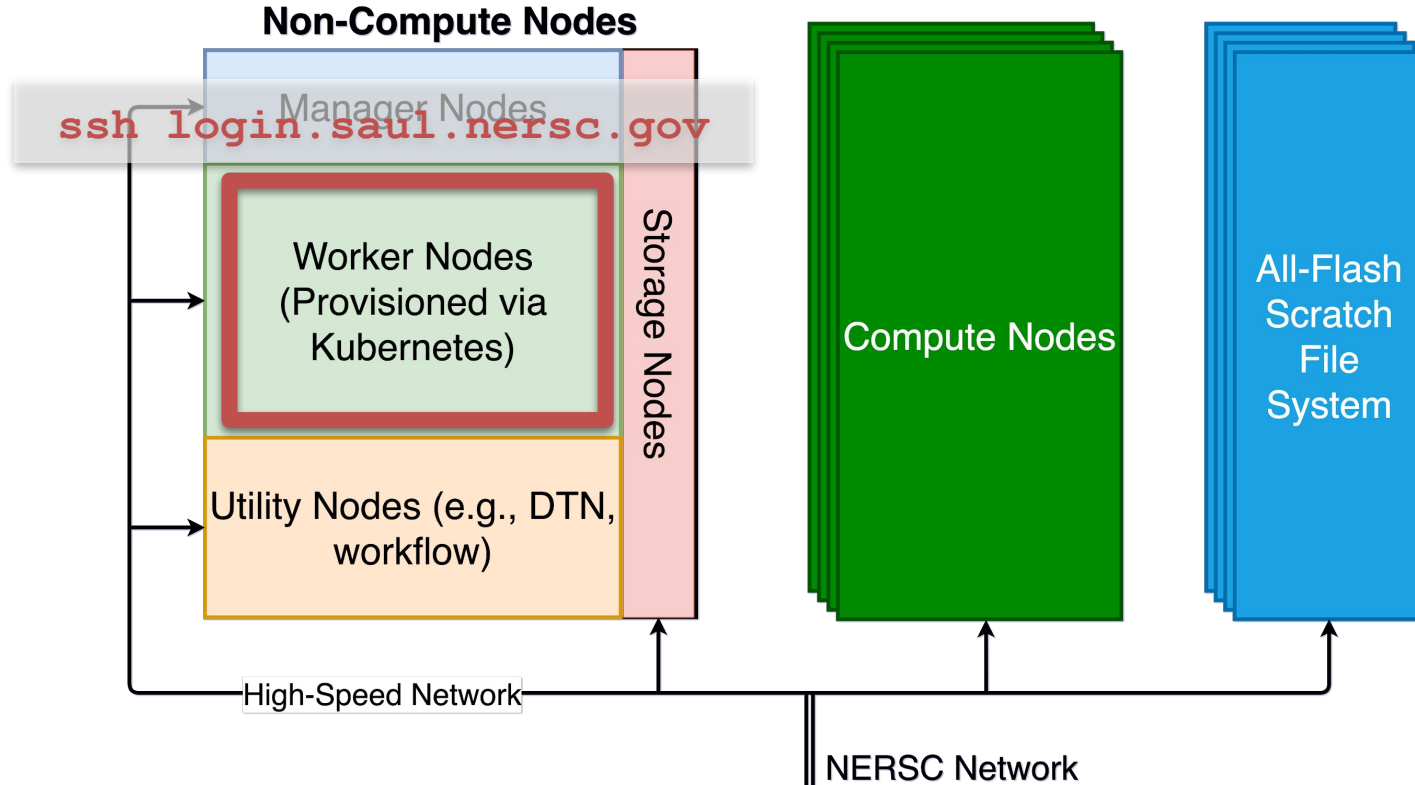
~ May 2021:

- Large-scale scientific computing access for all GPU-capable projects
- GPU-readiness evaluation form required
- Key GPU-enabled community apps will be available

Allocations & Charging

- In Allocation Year (AY) 2021:
 - All time will be “free”
 - However, no commitment for hours available
 - *Expect that there will be outages for testing & stabilization*
- ERCAP requests & charging will begin in AY 2022
- Allocation and Charging Units
 - **GPU Node Hours** for Perlmutter GPU accelerated nodes
 - **CPU Node Hours** for Perlmutter CPU, Cori Haswell and KNL nodes
 - **GPU** & **CPU** hours will not be interchangeable!
 - CPU Node Hour charges will incorporate performance scale factors for 3 types of nodes

Perlmutter Architecture: Conceptual Overview



The Login Experience

- **ssh login.saul.nersc.gov**
- Each user within a container orchestrated by Kubernetes
 - Spin up resources as needed
 - Insulation from other users' behavior
- Jupyter also available for login



Login Node Environment

- Will be largely familiar to users
- Differences:
 - LMOD instead of modules (also on Cori in AY21*)
 - Hierarchical based modules, should be easier to find & load modules
 - Different programming environments
 - No Intel

*probably



Programming Environment

	GPU Support	Fortran/ C/C++	OpenACC 2.x	OpenMP 5.x	CUDA	Kokkos / Raja	Cray MPI
PGI							
CCE							
GNU							
LLVM							

Vendor
Supported

NERSC
Supported

Software

- NERSC and/or Vendor will provide many libraries for users
 - Available through LMOD modules (similar to current environment, with a few advantages)
- Users will be able to compile software not provided
 - User Spack instance, to draw upon pre-existing recipes
 - Help from NERSC consultants to install in user or project space
- New: Cray Minerva Data & Analytics Software Stack

Software: Minerva Data & Analytics Stack

Library	Vendor Supported?	GPU Enabled?
Python (Anaconda)		
Spark		
R		
TensorFlow		
Keras		
Caffe		
PyTorch		

Running Jobs

- Slurm will continue to be used for job scheduling
- Additional flag needed for requesting GPU resources
- Queue policies TBD, but should resemble Cori



Learning Opportunities: Training

- NERSC will hold training sessions for diverse interests and levels of experience
 - How to use the system
 - How to compile codes
 - Performance optimization
 - NVIDIA A100 Architecture deep dive
 - Development and tools
 - Machine Learning
 - Chemistry / materials science applications
 - and more!

Learning Opportunities: Office Hours

- Over past 2 years, we have developed & deployed office hours concept modeled after University office hours
 - Past topics included ERCAP, MFA, KNL
- We will continue to offer opportunities for one-on-one help with targeted topics
 - Replace open office hours with appointment scheduling
- We have some ideas for topics, but also will seek your feedback
 - Please look for survey in a few weeks

Learning Opportunities: Hackathons


- Hackathons come in a variety of forms, but generally:
 - Pair code teams with experienced mentors.
 - Give an opportunity learn new profiling techniques and tools.
 - Identify, explore and directly fix problems with your codes.
 - Research and learn about new coding strategies and methods.
 - Develop contacts for future collaboration, code development and support.
 - Virtual-only formats are being adopted, tested and expanded on now.

Types of Hackathons

	GPU Hackathons	Center of Excellence (COE) Hackathons	NERSC Hackathons
Host	NVIDIA, OpenACC organization & a local host.	NERSC / HPE	NERSC
Format	Intense, 4 day events with preparation & follow-up.	B.C.: 6 weeks of intense preparation with HPE/NVIDIA, 1 week of full-time hacking. Now: ~12 weeks of lower intensity work with HPE/NVIDIA Collaborative events	Variety of formats, but often 1-2 day events with speakers to start followed by hands-on with your code.
Frequency	Many events year round	Quarterly	periodically, thematic
More Info	https://www.gpuhackathons.org	See next slide!	NERSC Weekly Email & NERSC website

Upcoming COE Hackathons

- Seeking 4 NESAP teams to commit to ~12 weeks of collaboration with HPE/NVIDIA on GPU performance optimization/tuning of their codes
- All NESAP teams encouraged to apply
 - Teams who have already participated in hackathons are welcome to apply again
 - Tier 2 NESAP teams also encouraged to apply
- Please contact Brian Friesen <bfriesen@lbl.gov> if you are interested in participating



Any Questions?

Thank You!

