Introduction to Perlmutter

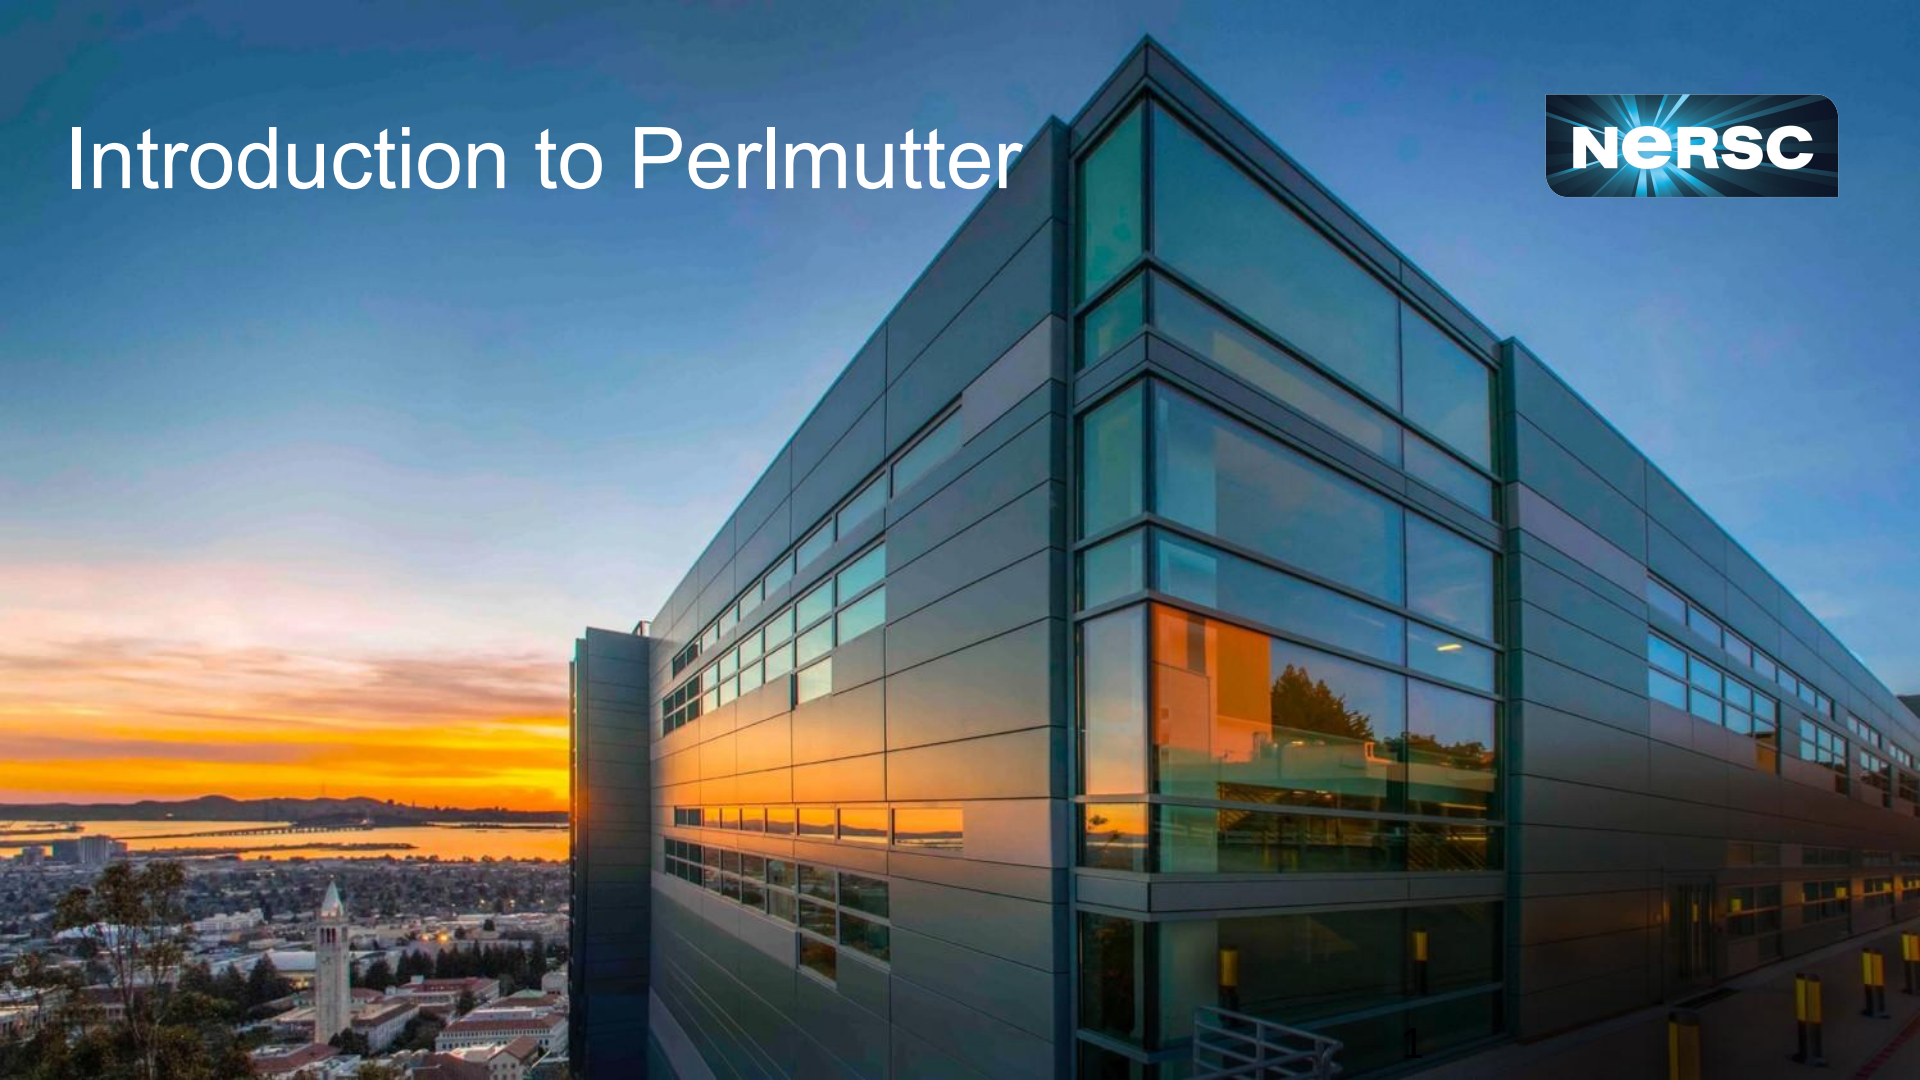

NERSC Systems Roadmap

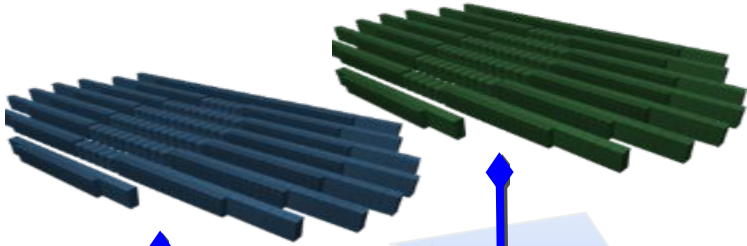

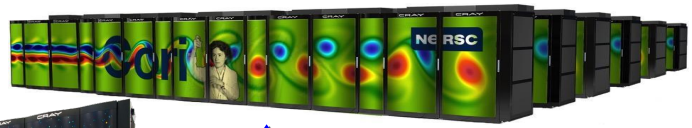

NERSC-7: Edison Multicore CPU

2013

NERSC-8: Cori Manycore CPU
NESAP Launched: transition applications to advanced architectures

2016

NERSC-9: CPU and GPU nodes
Continued transition of applications and support for complex workflows

2021

NERSC-10: Exa system

2025

NERSC-11: Beyond Moore

2029

Perlmutter: A System Optimized for Science

- Cray Shasta System
- GPU-accelerated and CPU-only nodes meet the needs of large scale simulation and data analysis from experimental facilities
- Cray "Slingshot" - High-performance, scalable, low-latency Ethernet-compatible network
- Single-tier All-Flash Lustre based HPC file system

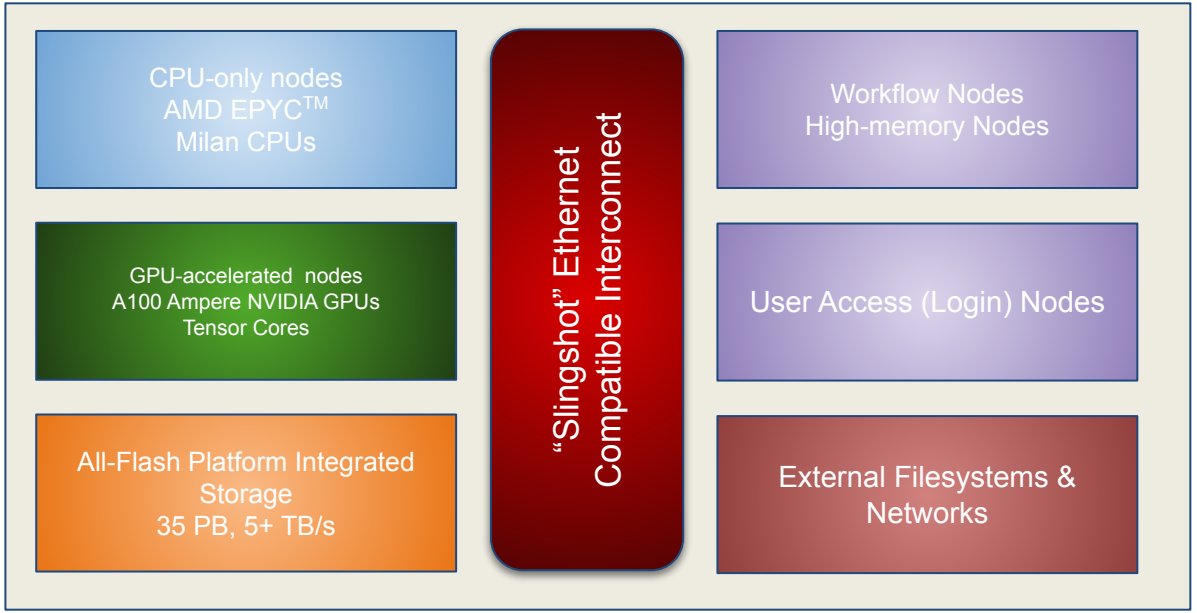

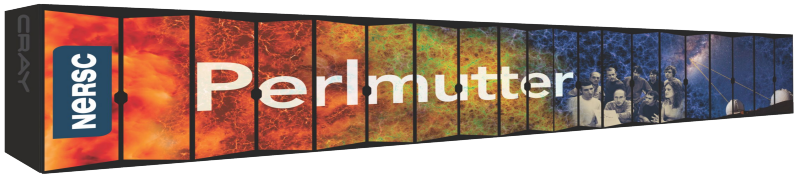

Perlmutter: System Details

- System Management Orchestration using Kubernetes
- GPU enabled Login/Workflow nodes
- Large Memory Nodes with 1 TB memory
- Resilient, High-BW link to the NERSC network and the World

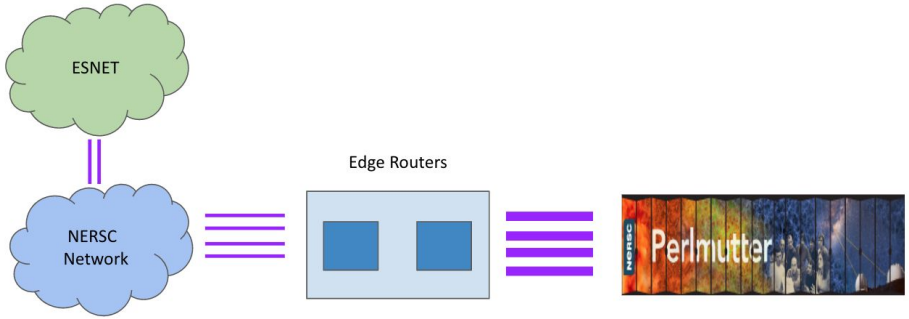

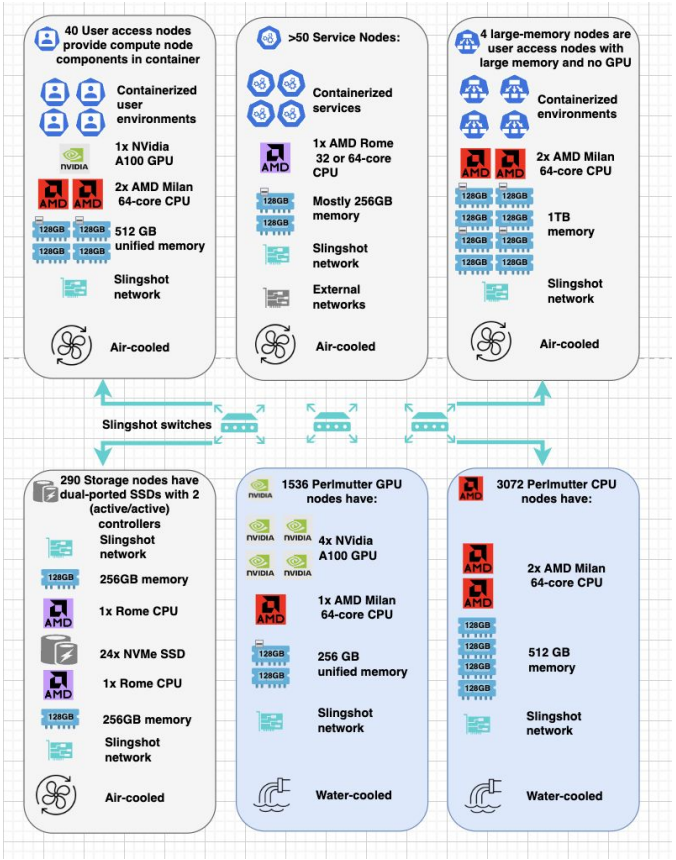

Perlmutter: HW Capabilities

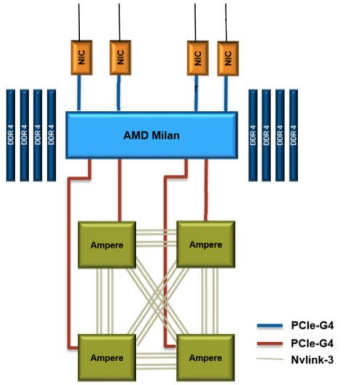

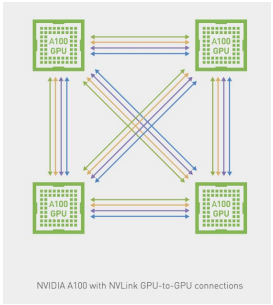

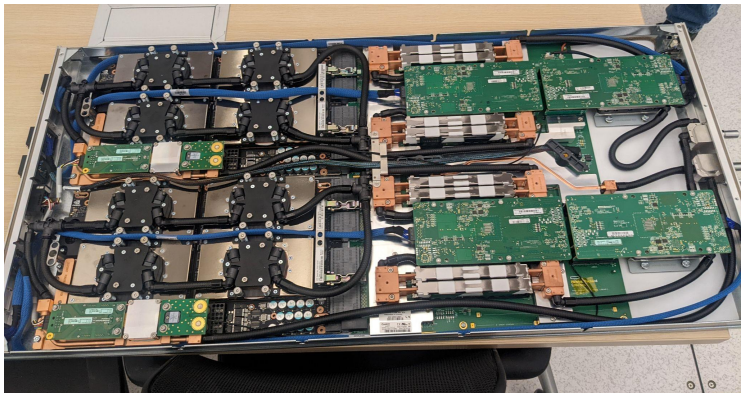

- GPU-Accelerated Nodes
  - 4 Nvidia A100 GPUs per node
    - 40 GB High-BW-Memory per GPU
    - GPUs linked with NVLink-3
  - 1 AMD EPYC 7763 CPU
    - 256 GB DRAM
  - 4 Slingshot-11 (200 Gbps) Network cards per node
- CPU-Only Nodes
  - 2 Socket AMD EPYC 7763 CPU
    - 512 GB DRAM
  - 1 Slingshot-11 (200 Gbps) Network card per node

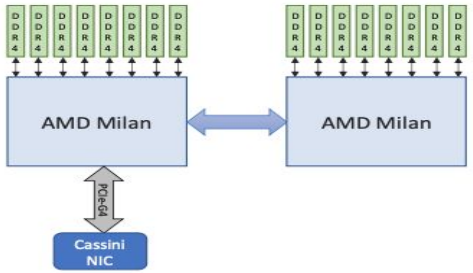

Perlmutter: Shasta Features

- Cloud-managed infrastructure (Kubernetes)
- Service-oriented architecture defines results needed, not operational details
- Services can be tied to specific resources or run within specific classes of nodes
- Vendor value-add software accessible via system-wide application programming interface

Non-compute Nodes (NCN)

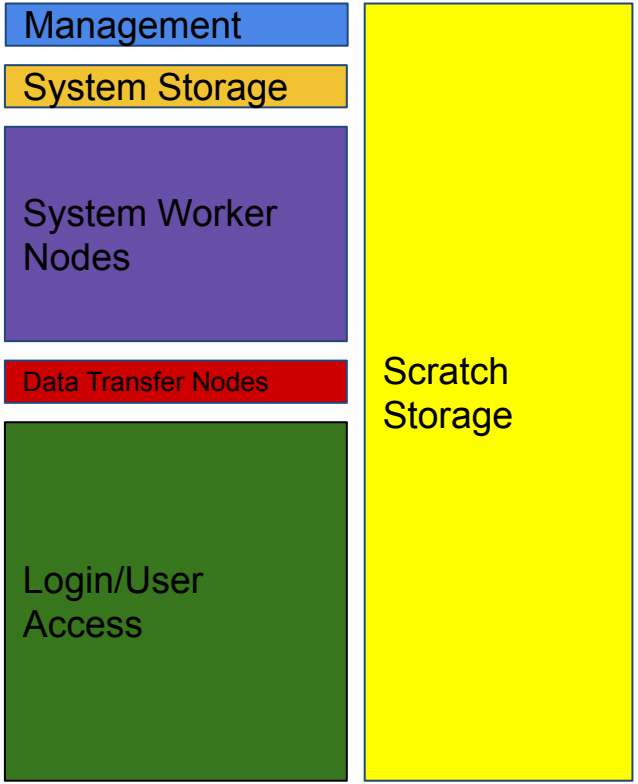

Compute Nodes

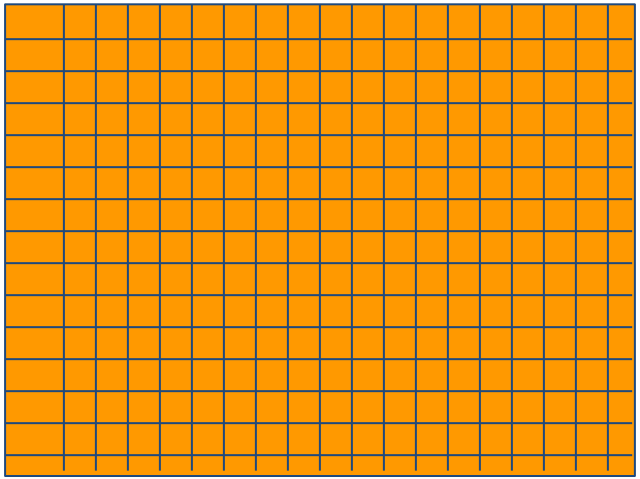

- Enterprise Linux environment (SLES), Vendor modified
- "Bare metal" booted (limited complexity)
- Leverage existing Vendor value add (Programming Environment, Cray Linux)

Perlmutter: Differences from Cori

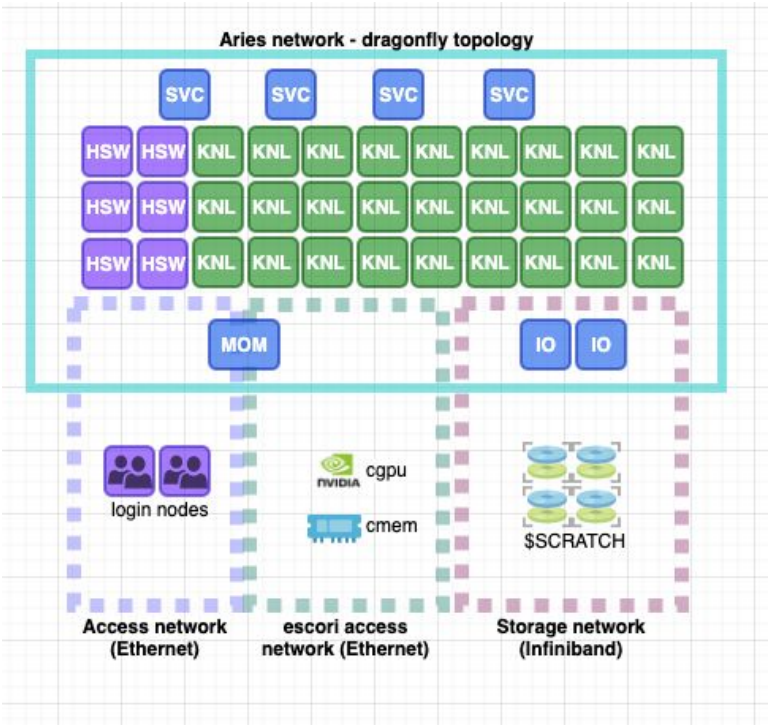

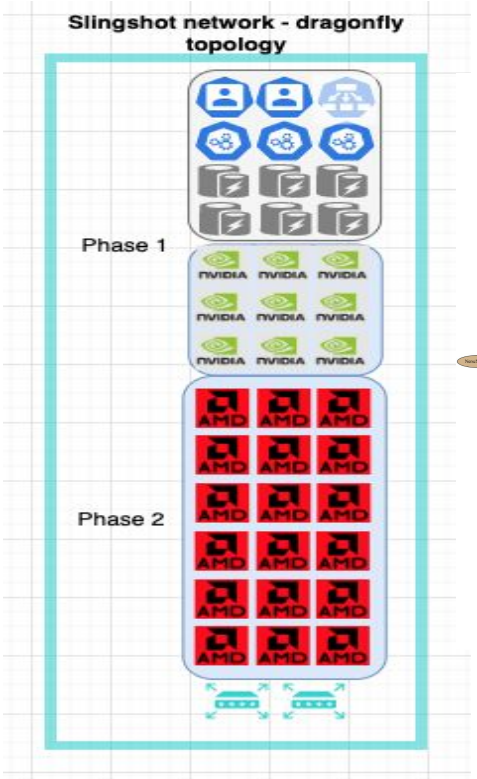

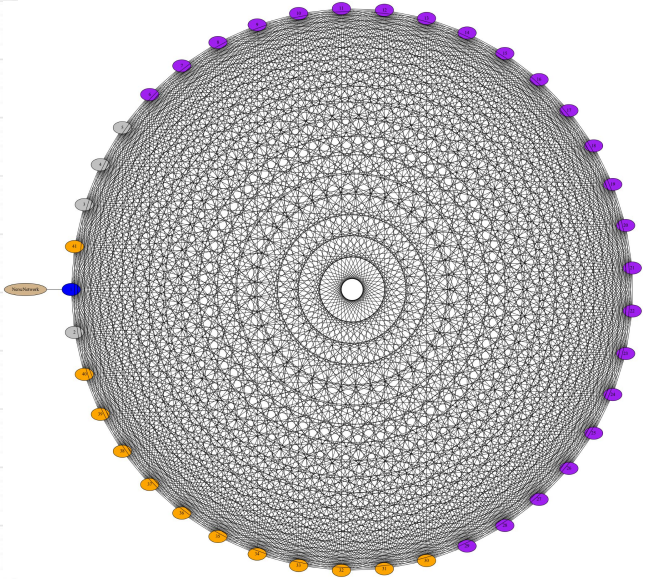

Perlmutter: Software

Programming Environments

	GPU Support	FORTRAN/C /C++	OpenACC 2.x	OpenMP 5.x	CUDA	Kokkos / Raja	Cray MPI
NVIDIA							
CCE							
GNU				(Community Effort)			
LLVM				(Community Effort)			

Vendor Supported

NERSC Supported

Programming models and languages

kokkos

SYCL

OpenMP

OpenACC

Fortran

Community Codes

VASP

LAMMPS

CP2K

WEST!

QUANTUMESPRESSO

Perlmutter: Science

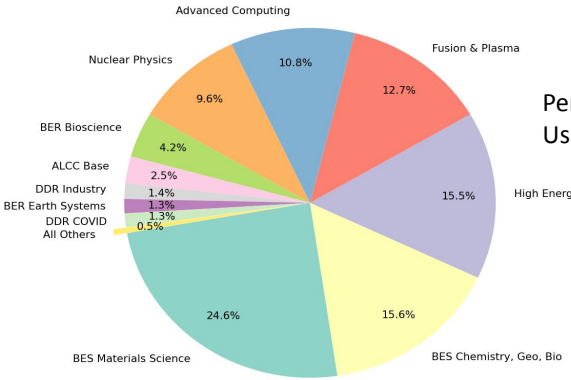

Perlmutter GPU Node Usage 2022

perlmutter Node Hours: 2022-06-01 00:00:00 to 2022-09-01 00:00:00

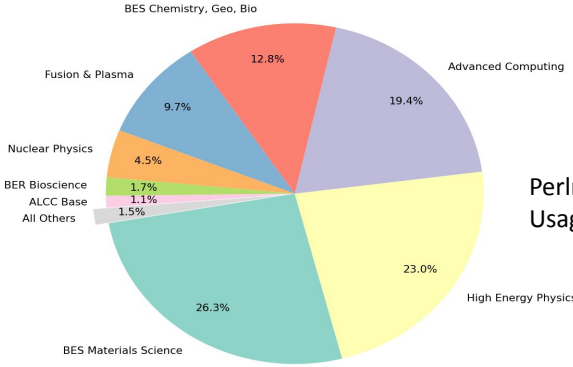

Perlmutter CPU Node Usage 2022

Early Successes in Superfacility

LCLS

- Reconstruct molecular structure form X-Ray scattering data
- Used Perlmutter for live data processing (ie., determining molecular structures during data collection), enabling real-time steering of the experiment

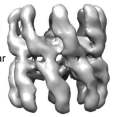

NCEM

- Multi-TB scale electron microscopy image simulation to train NN for materials research
- VASP to calculate XAS spectra to train a ML model for automated assignment of bond valence
- Real-time processing and reduction of 4D-STEM data with distiller.lbl.gov

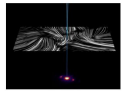

Lux-Zepplin (LZ)

- Dark matter detection experiment used GPUs for ray tracing in detector simulation
- Used Perlmutter to extract limits on dark matter-nucleon interaction for first science results

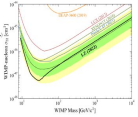

Dark **E**nergy **S**pectroscopic **I**nstrument

- DESI Spectral Extraction is an image processing code implemented in Python.
- **2.5x** improvement in per-node throughput using Perlmutter A100 compared to Cori V100 GPU (x25 compared to Edison).

Exaop Performance for the *Ab-Initio* Molecular Dynamics

Ground-breaking real-world exaop calculation in mixed FP16/32 run on Perlmutter

- The non-orthogonal local submatrix method applied to electronic-structure based molecular dynamics simulations exceeds 1.1 EFLOP/s in FP16/FP32 mixed floating-point arithmetic
- Used 4,400 NVIDIA A100 GPUs on Perlmutter
- The method achieves a sustained fraction of peak performance of about 80%.
- Example calculations are performed for SARS-CoV-2 spike proteins with up to 83 million atoms.

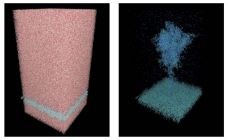

SARS-CoV-2 spike protein in aqueous solution: full cell (left) and without hydrogen and oxygen atoms (right).

Robert Schade, Tobias Kenter, Hossam Elgabarty, Michael Lass, Thomas D. Kühne, Christian Plessl, Paderborn University

arXiv:2205.12182v1 24 May 2022

Deep Learning Reveals How Proteins Interact

Scientific Achievement

- RoseTTAFold neural network software was used to help screen through 8.3 million pairs of yeast proteins
- Identified 1,505 likely to interact
- Built structure models for 106 previously unidentified assemblies and 806 that had not been structurally characterized.

Significance and Impact

- Protein-protein interactions play a key role in biological processes.
- Models of how proteins interact provide insights into function.
- Structures of many complexes are unknown & many interactions not yet identified.
- Advances in evolutionary analysis & deep learning are enabling revolutionary 3D models of how these interactions take place.

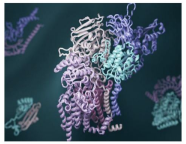

3D computed structural modeling of protein interactions is made possible by deep learning and evolutionary analysis. Credit: University of Washington (Ian Haydon)

*Humphreys, et al., "Computed structures of core eukaryotic protein complexes"; **Science**, 2021 Nov 11, 10.1126/science.abm4805*

Early Successes in Data/Learning

DESC (Dark Energy Science Collaboration)

- Using GPUs with Tensorflow, via Jupyter, for redshift model-fitting
- Distributed TF at scale on GPU w/ NCCL
- Tested up to 2048^3 N-body simulation, distributed or 256 GPUs
- Multiple TB of RAM

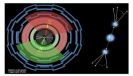

Anomaly Detection, Unfolding & Fast Simulation in Particle Physics

- DL techniques used in searches for fundamental particles at the LHC
- Expanding to more complex models/approaches and higher-fidelity generative networks

Data-driven Atmospheric Modeling

- ML Data-driven prediction of high-resolution atmospheric flow variables
- **2.9x** improvement in throughput using Perlmutter A100 compared to Cori V100 GPUs

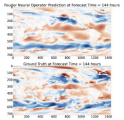

Open Catalyst Project

- Deep learning to accelerate catalyst discovery for reactions that are critical for energy storage and climate change mitigation
- Scaling current models from O(10-100) GPUs to O(1000) GPUs

Perlmutter: Future

Operational Improvements

- Continuous Operations
  - Expect most maintenances & software updates to be done non-disruptively

User Access Improvements

- API-driven interactions
  - RESTful interface to Slurm Workload Manager
  - Start/manage/maintain gitlab runners
  - Data movement operations

User Environment Improvements

- User Access Instance (UAIs)
  - User log-in directly to a dedicated container
  - Provision of standard images with user customized images also possible
  - Capability for long-running Kubernetes managed user services

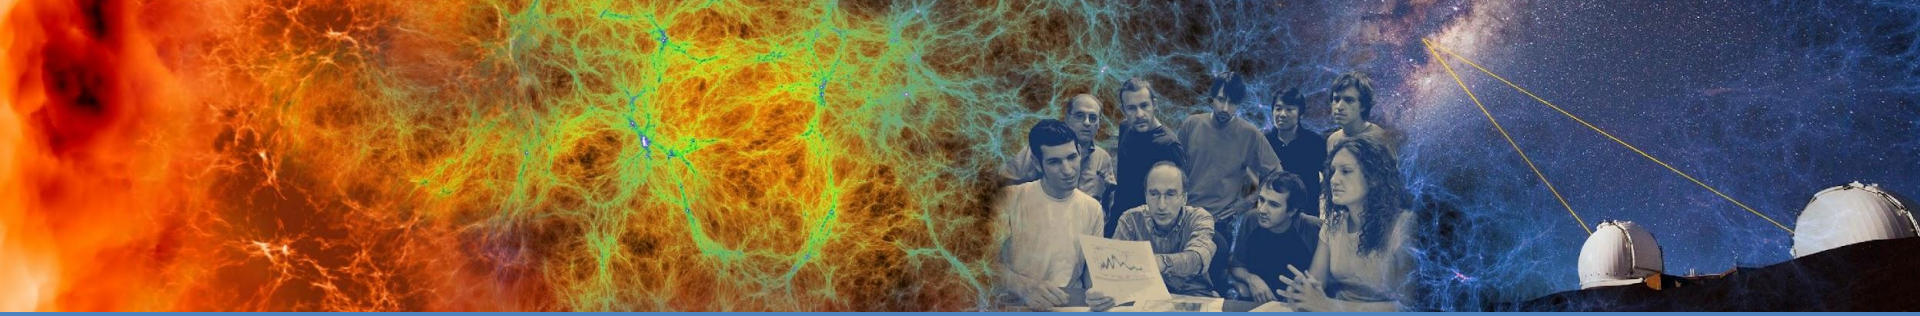

End