Perlmutter
and the Next Procurement - NERSC-10

HPC User Forum 2023
Tucson AZ

Nick Wright
Chief Architect
& Advanced Technologies Group Lead

7th Sept 2023

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.

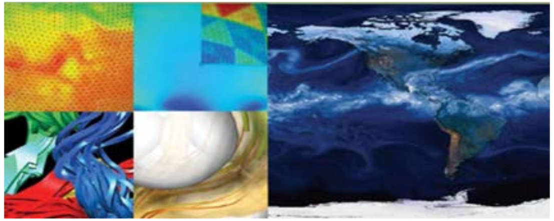

Biological and Environmental Research

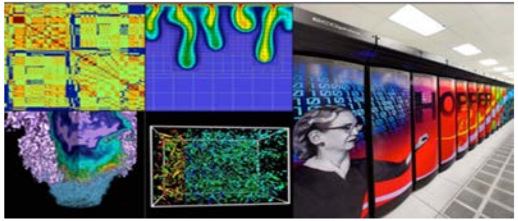

Computing

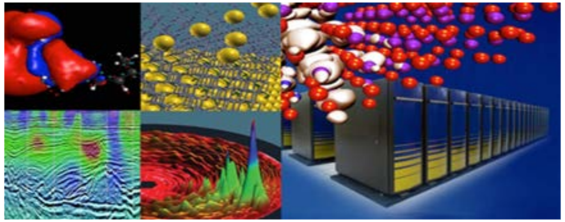

Basic Energy Sciences

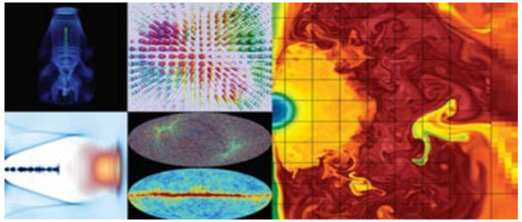

High Energy Physics

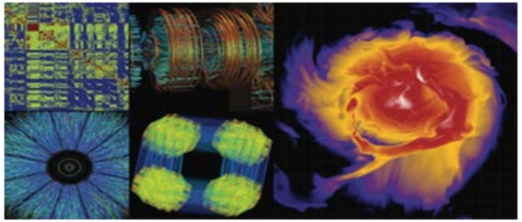

Nuclear Physics

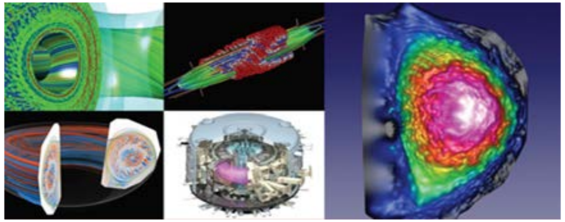

Fusion Energy, Plasma Physics

Nobel-Prize Winning Users

2013 Chemistry

for the development of multiscale models for complex chemical systems

Martin Karplus

2011 Physics

for the discovery of the accelerating expansion of the Universe through observations of distant supernovae

Saul Perlmutter

2006 Physics

for the discovery of the blackbody form and anisotropy of the cosmic microwave background radiation

George Smoot

2007 Peace

for their efforts to build up and disseminate greater knowledge about man-made climate change

Warren Washington

2017 Chemistry

for developing cryo-electron microscopy for the high-resolution structure determination of biomolecules in solution

Joachim Frank

2015 Physics

for the discovery of neutrino oscillations, which shows that neutrinos have mass

SNO Collaboration

NERSC by the Numbers

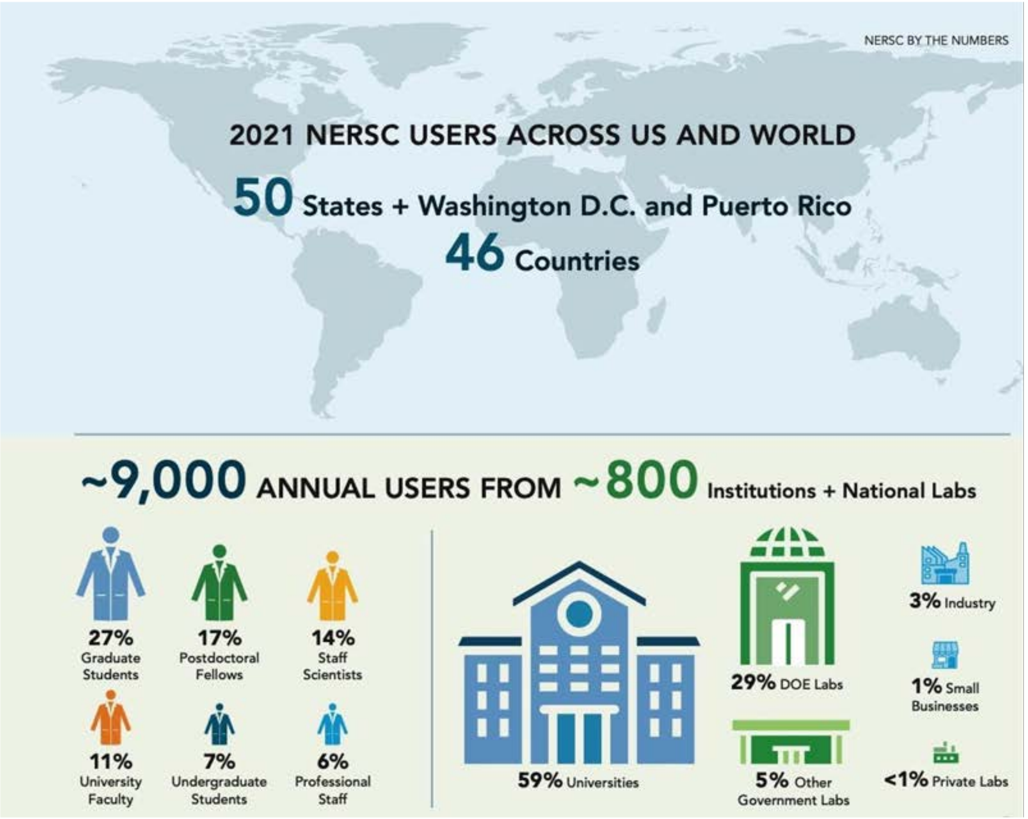

NERSC has been acknowledged in 5,829 refereed scientific publications & high profile journals since 2020

- Nature [32]
- Nature Communications [116]
- Proceedings of the National Academy of Sciences [55]
- Science [21]
- Nature family of journals [232]
- Monthly Notices of the Royal Astronomical Society [248]
- Physical Review B : Condensed Matter and Materials Physics [206]
- Physical Review D : Particles, Fields, Gravitation, and Cosmology [200]

We Accelerate Scientific Discovery for Thousands of Office of Science Users with 3 Advanced Capability Thrusts

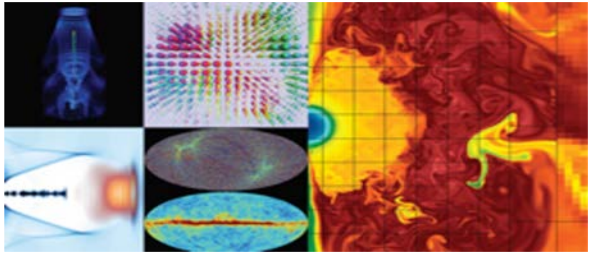

Large-scale applications for simulation, modeling and data analysis

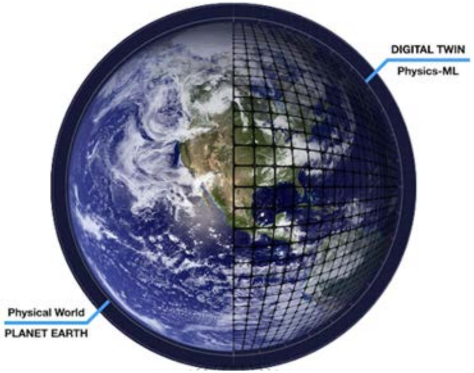

Complex experimental and AI driven workflows

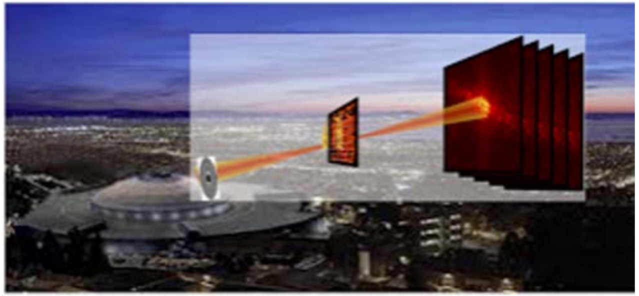

Time-sensitive and interactive computing

The NERSC workload is diverse with growing emphasis on integrated research workflows

NERSC Ecosystem

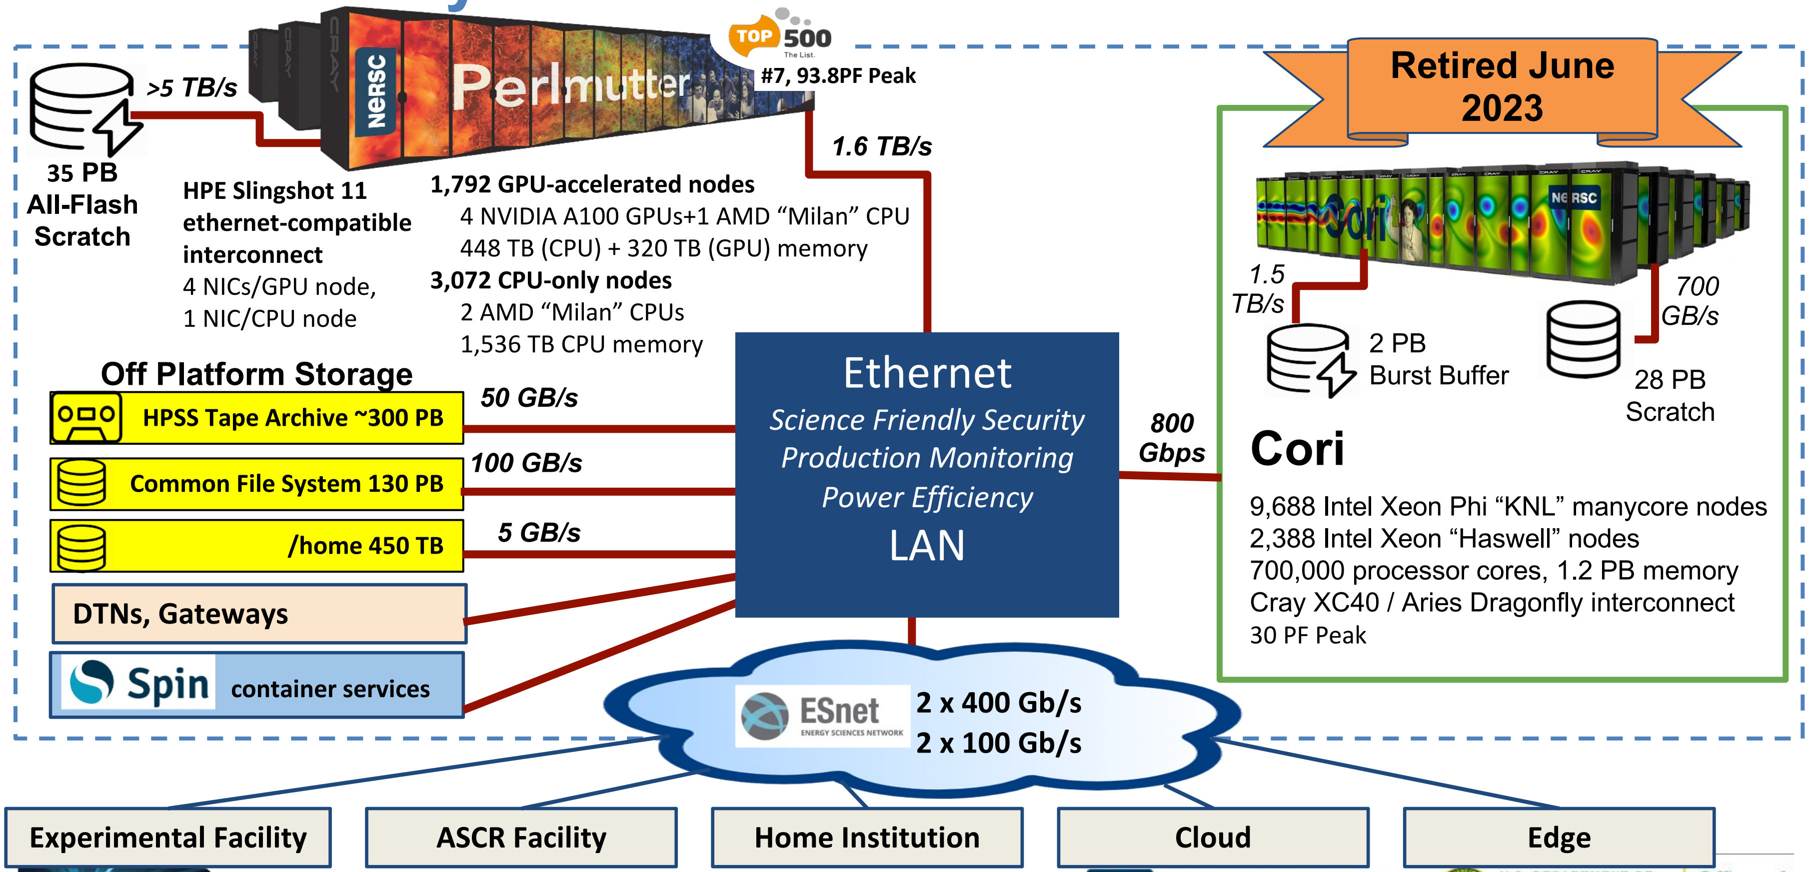

NERSC Systems Roadmap

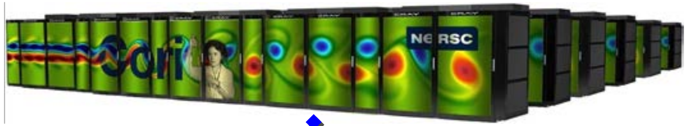

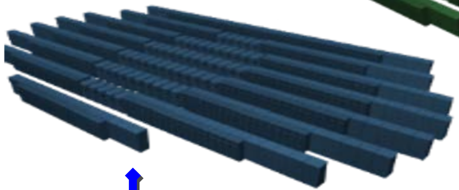

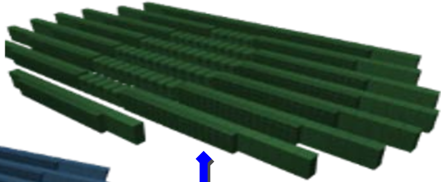

2013

NERSC-7: Edison
Multicore CPU

2016

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

2020

NERSC-9: Perlmutter
CPU and GPU nodes
NESAP Expanded Simulation, Learning & Data: Continued transition of applications and support for complex workflows

2026

NERSC-10:
Exa system
NESAP Workflows: Accelerating end-to-end workflows with technology integration

2030+

NERSC-11:
Beyond Moore

NERSC Systems Roadmap

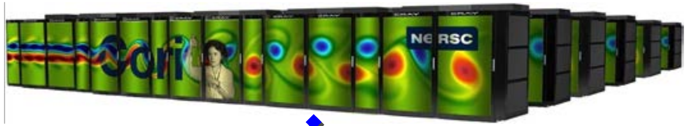

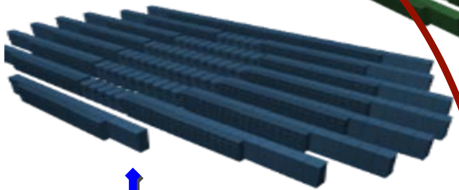

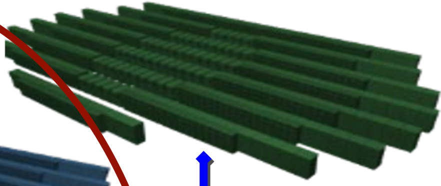

2013
NERSC-7: Edison
Multicore CPU

2016
NERSC-8: Cori
Manycore CPU
NESAP Launched: transition applications to advanced architectures

2021
NERSC-9: Perlmutter
CPU and GPU nodes
NESAP Expanded
Simulation, Learning & Data: Continued transition of applications and support for complex workflows

2026
NERSC-10:
Exa system
NESAP Workflows: Accelerating end-to-end workflows with technology integration

2030+
NERSC-11:
Beyond Moore

Forming NERSC-10 Strategy

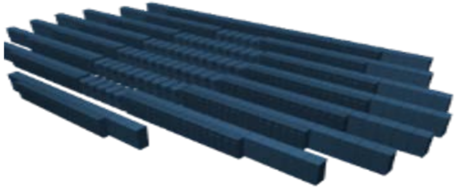

- Target 4QCY2026 Delivery
  - 10x Perlmutter on applications
- Can we use the same strategy as we did Perlmutter – (~2017 for 2020 delivery) ?
- Examine trends in
  - Technology
  - Market & Vendor Landscape
  - User Community

Technology Trends

- No more increases in clock speed for CPUs & GPUs
  - More & more cores
- End of Moore's Law
  - Performance per socket may continue to double through
- Increases in performance will primarily be obtained through power increases
  - At the socket & the system level
- Tighter & Tighter CPU-GPU integration
  - Grace-Hopper from NVIDIA
  - MI-300 from AMD
- Flash Storage will continue to increase in capacity and eat into HDD space

Software Technology Trends

Chip-vendor provides user software toolchains

Service-oriented architectures and microservices enable resilience and extreme scale for workflows

- Containerized services (Docker, LXC)
- "serverless" computing (Lambda)

Software-defined/programmable infrastructure

- Software-defined networking (SDN, SD-MPLS, EVPN)
- Software-defined storage (SPDK)

AI for operations and resource management

- Anomaly detection, cybersecurity
- Energy efficiency and automated controls
- Complex scheduling

OPENSHIFT

FSx

splunk>

DeepMind

Vendor Landscape Has Changed Dramatically

APEX (N9) [11/2015]		NERSC-10 [11/2022]
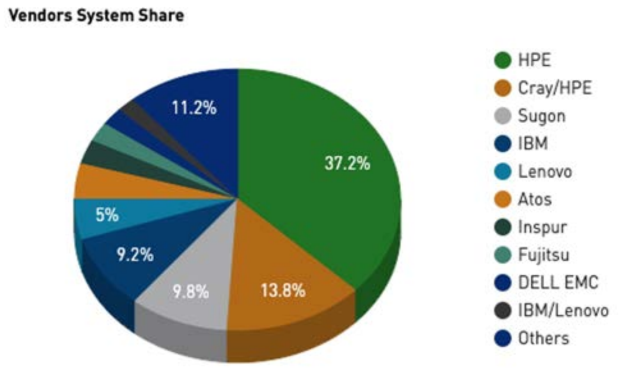 **199 Top 500 Systems in the US**	IBM? HPE/Cray market share decreased Growing industry need for HPC Many large systems are not showing up on the Top500 (AI & cloud)	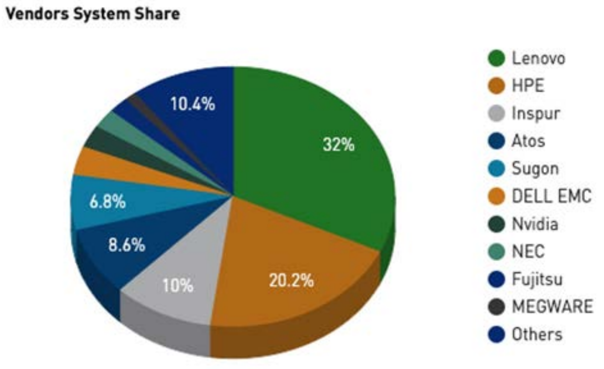 **127 Top 500 Systems in the US**

Engage with numerous vendors to reinvigorate and redefine the landscape of technology providers and integrators

What About Deep Learning?

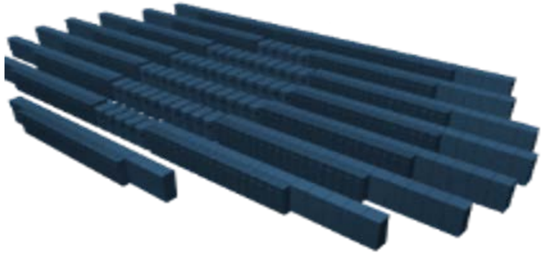

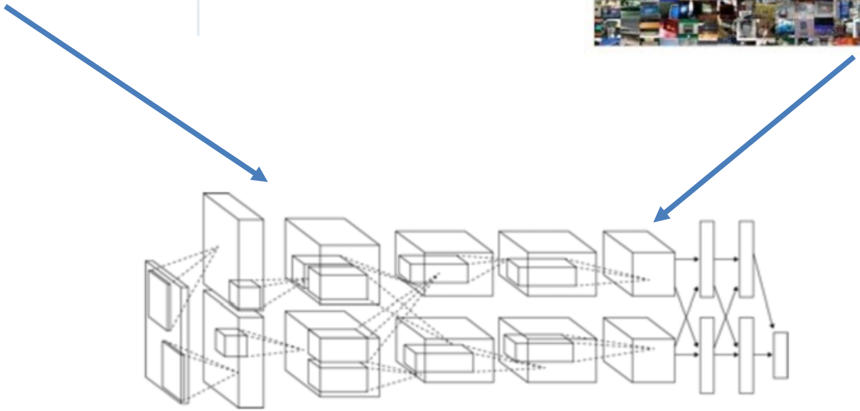

AI Training Uses Supercomputers

SIGN IN / UP — The Register

AI + ML — 7

Google boffins pull back more of the curtain hiding TPU v4 secrets

And Nvidia's having none of Big G's claims of superiority

Thomas Claburn — Thu 6 Apr 2023 // 01:34 UTC

Google on Wednesday revealed more details of its fourth-generation Tensor Processing Unit chip (TPU v4), claiming that its silicon is faster and uses less power than Nvidia's A100 Tensor Core GPU.

TPU v4 "is 1.2x–1.7x faster and uses 1.3x–1.9x less power than the Nvidia A100," said researchers from Google and UC Berkeley in a paper published ahead of a June presentation at the International Symposium on Computer Architecture. Our pals over at *The Next Platform* previously dived into the TPU v4's architecture, here based on earlier material released about the chips.

After Google's reveal this week, Nvidia coincidentally published a blog post in which founder and CEO Jensen Huang noted that the A100 debuted three years ago and that Nv's more recent H100 (Hopper) GPUs deliver 4x more performance than A100 based on MLPerf 3.0 benchmarks.

Google's TPU v4 also entered service three years ago, in 2020, and has since been refined. The Google/UC Berkley authors explain that they chose not to measure TPU v4 against the more recent H100 (announced in 2022) because Google prefers to write papers about technologies after they have been deployed and used to run production apps.

SIGN IN / UP — The Register

HPC — 7

Tesla hedges Dojo supercomputer bet with 10K Nvidia H100 GPU cluster

Keeping full self-driving dream on the road just needs more graphics chips?

Tobias Mann — Wed 30 Aug 2023 // 10:29 UTC

Tesla still dreams of fueling its motors with actual full self-driving (FSD) capabilities, and it's blowing piles of cash on AI infrastructure to reach that milestone.

The American EV manufacturer's latest investment is in a 10,000 GPU compute cluster, revealed in a xeet by Tesla AI Engineer Tim Zaman over the weekend. The system, which came online Monday, will help crunch the data collected by its vehicles and accelerate development of the FSD functionality we've heard so much about. The automaker declined to comment further.

Tesla has been teasing fully autonomous driving capabilities since 2016. So far what's been delivered is essentially super-cruise-control: a driver assistance system that is not truly self-driving and requires a human to keep their hands on the wheel.

CEO Elon Musk has no problem throwing money at his goal of achieving FSD. Last month Tesla revealed it would invest $1 billion to build out its Dojo supercomputer between now and the end of 2024 to speed the development of its autonomous driving software.

AI Training Uses Supercomputers

The System

- Each system consists of 64 Google racks, deployed in 8 groups of 8
  - 4096 interconnected chips sharing 256TiB of HBM memory
  - Total compute >1 ExaFLOP
  - Each group of 8 racks gets a CDU (Coolant Distribution Unit)
- Dozens of systems deployed [Sundar, Google I/O]
  - Up to 8 superpod systems in a single cluster!

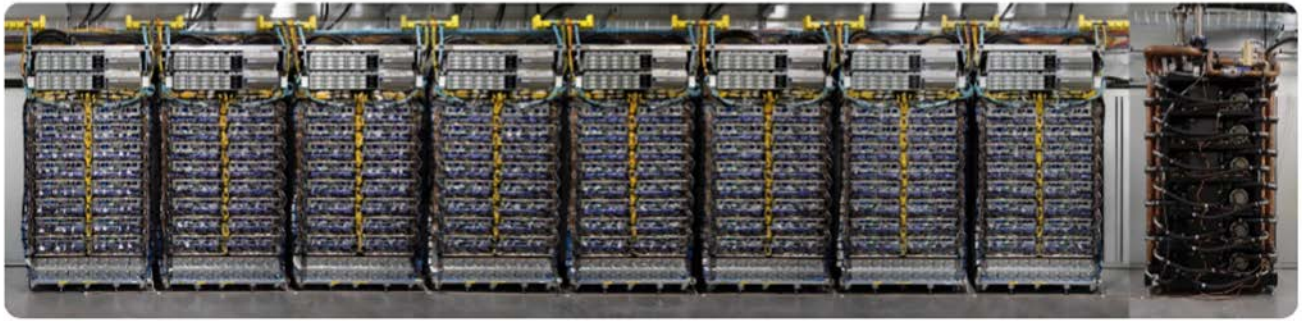

8x TPU Racks | 1x CDU

Google 10

A Machine Learning Supercomputer With An Optically Reconfigurable Interconnect and Embeddings Support
Norman Jouppi & Andy Swing, Google
Hotchips 2023

AI Training Uses Supercomputers

The Fiber

- Each Superpod has enough fiber to encircle the state of Rhode Island!
- Over 16,000 individual connections
- Major focus on deployability and serviceability

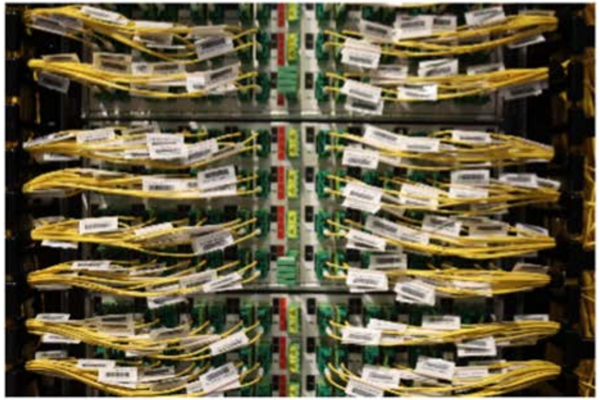

Google 13

A Machine Learning Supercomputer With An Optically Reconfigurable Interconnect and Embeddings Support
Norman Jouppi & Andy Swing, Google
Hotchips 2023

AI Training Uses Supercomputers

Scalability

- Goal was to create a highly scalable balanced system
- Hence TPUs connected by high BW to distributed shared memory
- We have ~linear speedups up to 3072 chips on internal workloads except for DLRMs

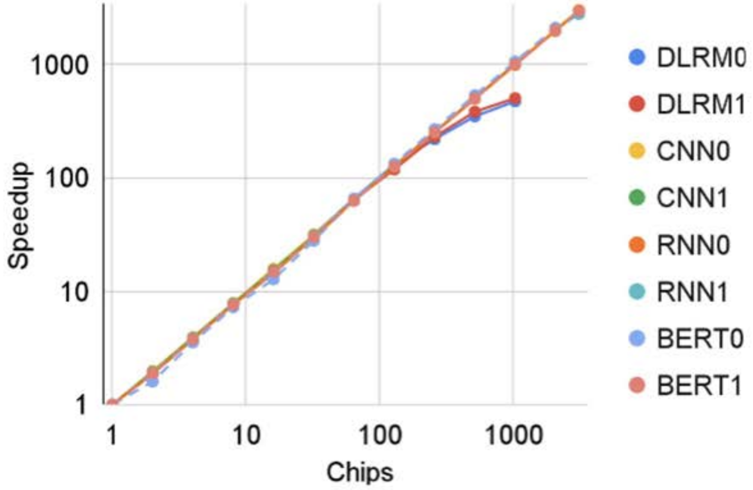

20

A Machine Learning Supercomputer With An Optically Reconfigurable Interconnect and Embeddings Support
Norman Jouppi & Andy Swing, Google
Hotchips 2023

AI – Training and Inference vs HPC

AI

- ~$300B market by 2026
- Smaller breadth of applications
  - GEMM dominant operation
  - FP16 and lower precision
  - Well understood persistent communication patterns
  - Complicated software stack
  - Storage ?

HPC

- ~$10B market by 2026
- Large breadth of applications
  - Often limited by memory bandwidth (not GEMM)
  - Need FP64
  - (Usually) not communication limited
  - Very complicated software stack
  - Parallel I/O at scale
- No real HPC analogue of Inference

- Both need liquid-cooled, serviceable and scalable deployments

Impact of AI/machine Learning on HPC?

- Liquid-cooling technologies will become commoditized
- Routine deployment of supercomputing scale resources *should* facilitate better, more robust solutions
  - Today capabilities of Cloud-based Deep Learning Supercomputer resources are equal to (or greater than) .gov ones
- HPC centers will need to focus on where they can add unique value

User Community - HPC Facility Workload Balance is Evolving

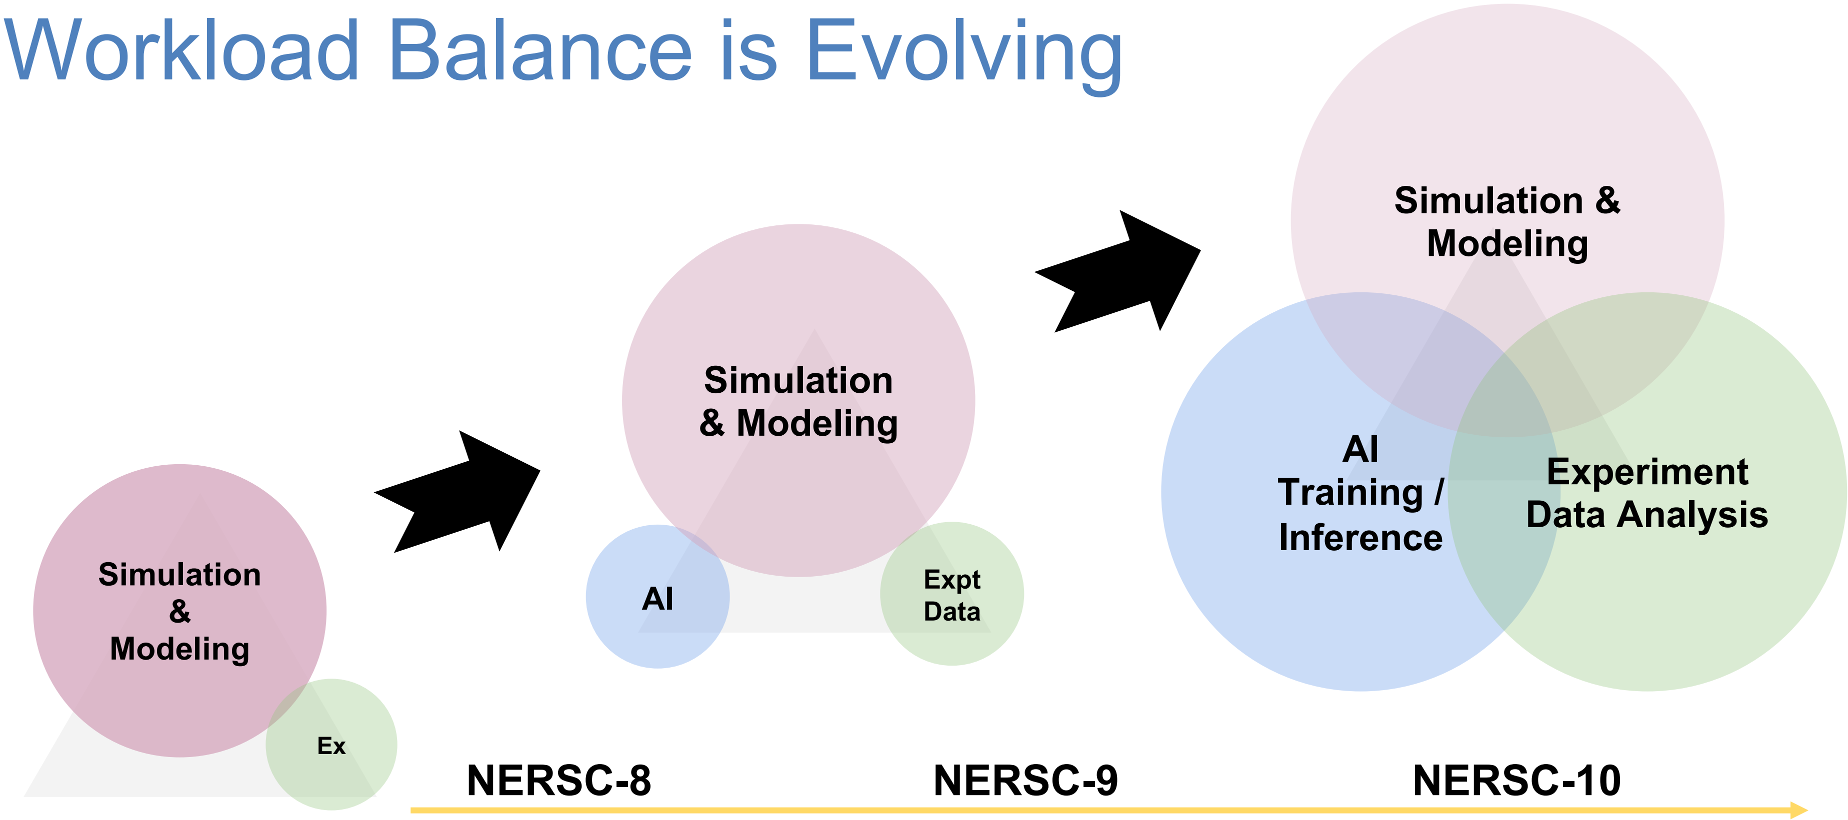

N10 User Requirements

Users require support for new paradigms for data analysis with **real-time interactive feedback between experiments and simulations**.

Users need the ability to search, analyze, reuse, and combine data from different sources into **large scale simulations and AI models.**

NERSC-10 Mission Need Statement:
*The NERSC-10 system will **accelerate end-to-end** DOE SC **workflows** and enable new modes of scientific discovery through the integration of experiment, data analysis, and simulation.*

BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY | Office of Science

What is an HPC Workflow?

Workflows are interconnected computational and dataflow tasks with data products. They have task coupling (control flow) and/or data movement between tasks (data flow).

High performance computing (HPC) workflows interconnect computational and data manipulation steps across one/some/all of:

- **High performance simulation and modelling**
- **High performance AI workflows**
- **High performance data analytics**

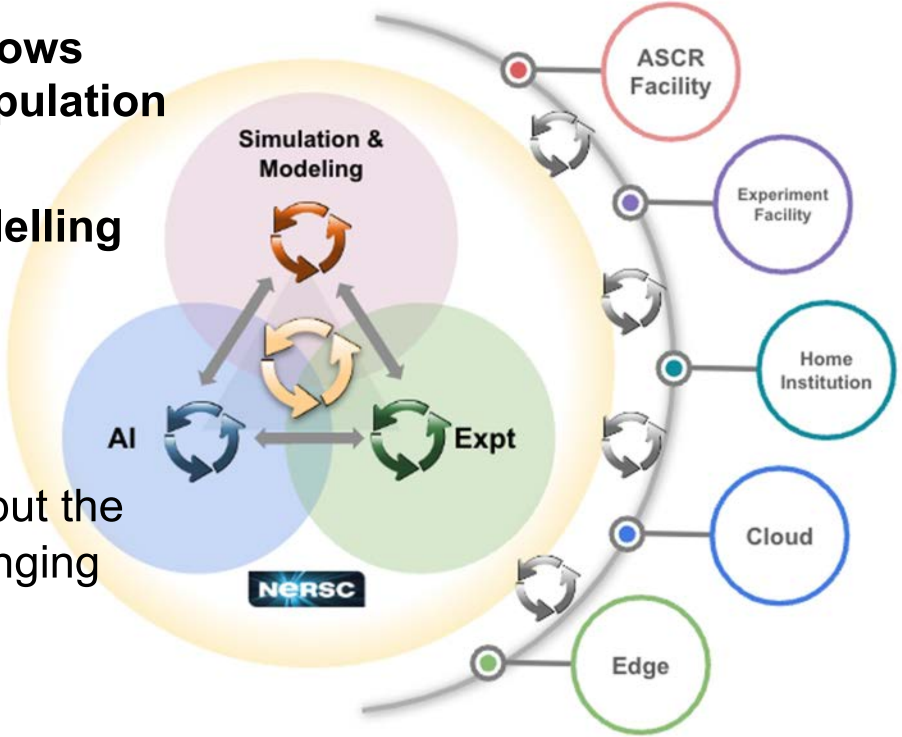

We've been running workflows for decades - but the complexity and timeliness of workflows is changing which motivates a new approach with N10.

Example of Cross-facility Workflow: Fusion Experiment

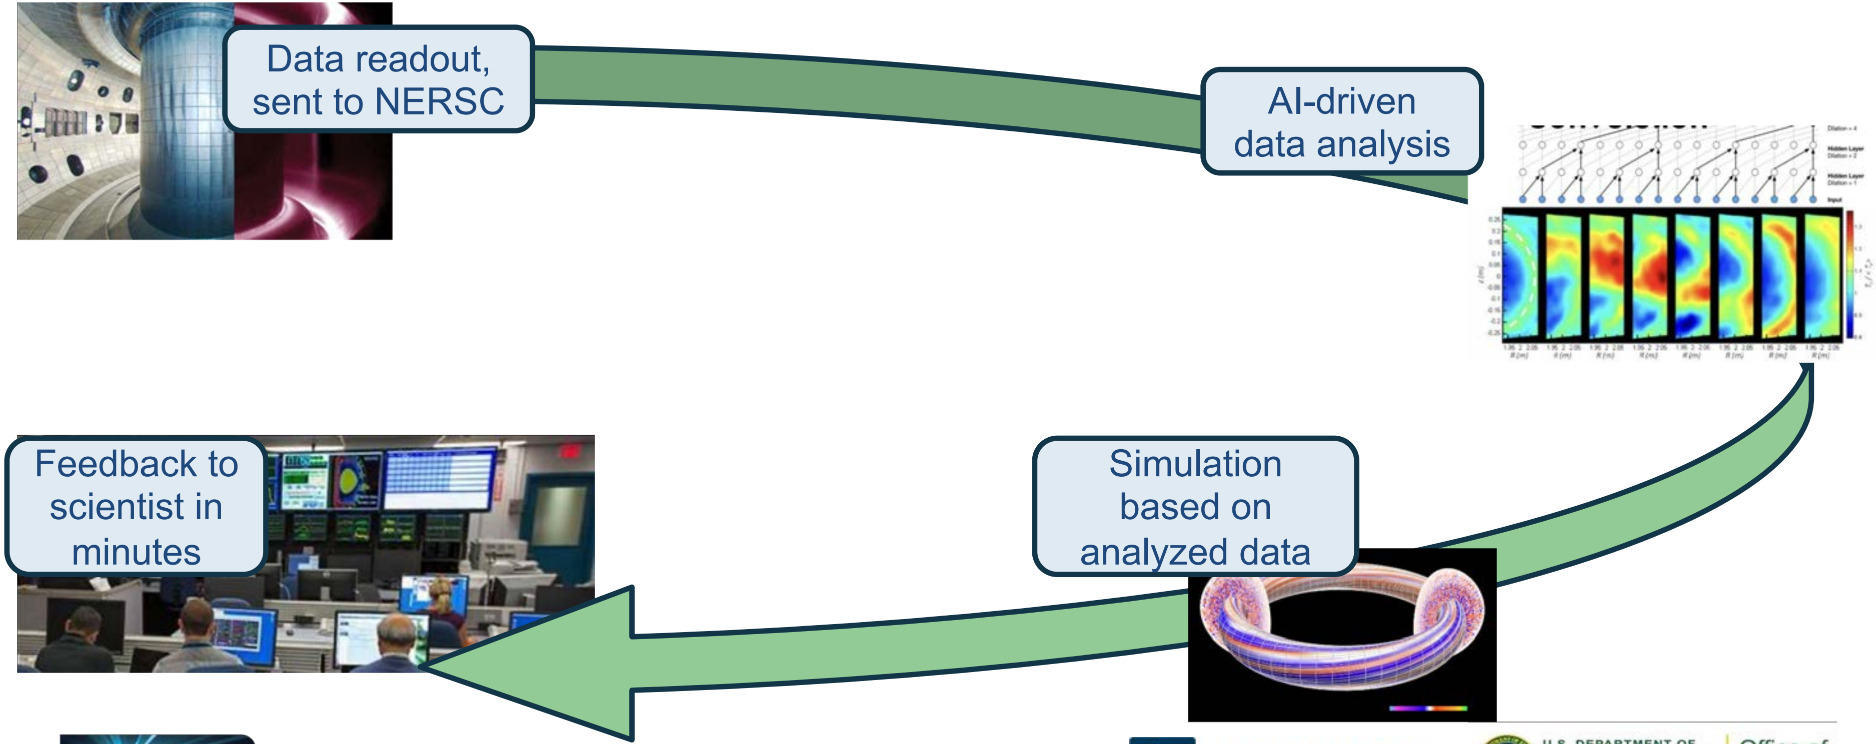

Example of Cross-facility Workflow: Fusion Experiment

Time-sensitive workflow requires **QSS** for deterministic performance and **network QOS** for guaranteed response in O(min)

Data reac
sent to NERSC

AI-driven
data analysis

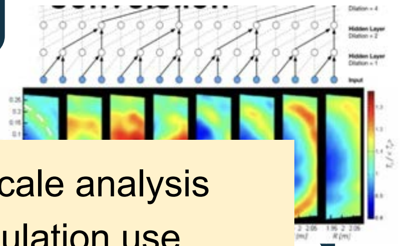

Data movement and compute progress tracked using **APIs** by automated workflow orchestrator and databases on **WENs**

Large-scale analysis and simulation use **containerized** apps and **accelerated** nodes.

Simulation
based on
analyzed data

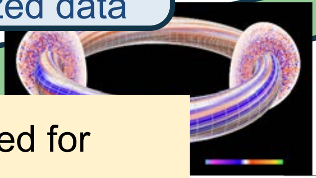

Feedback to
scientist in
minutes

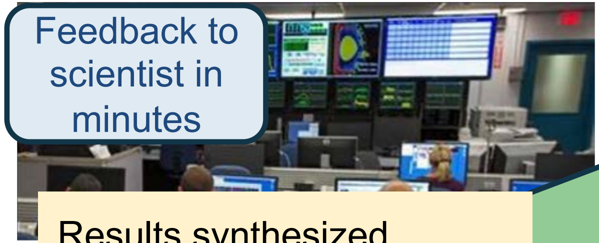

Results synthesized, displayed and shared via **Jupyter** and **python** ready for the next shot

Portable workflows designed for resiliency, possibly running on other resources if NERSC is unavailable

We identified 6 workflows archetypes to help define our vision for N10

1. High-performance simulation & modeling workflow	large-scale multi-physics applications with checkpoint/restart, data post-processing, visualization
2. High-performance AI (HPAI) workflow	data integration-intensive science patterns such as training, inference, hyperparameter optimization
3. Cross-facility workflow: Rapid data analysis and real time steering	time-sensitive science patterns such as superfacility, edge, and hybrid cloud
4. Hybrid HPC-HPAI-HPDA workflow	long-term campaign science patterns, AI-in-the-loop, AI-around-the-loop
5. Scientific data lifecycle workflow: Interactive, data-analytics and viz	data integration-intensive science patterns such as Jupyter, scientific databases, VSCode
6. External event-triggered and API-driven workflow	time-sensitive science patterns such as function-as-a-service, microservices

We identified 6 workflows archetypes to help define our vision for N10

1. High-performance simulation & modeling wo	large-scale multi-physics applications with … , visualization
2. High-perfor	…ns such as … mization
3. Cross-facili… analysis and r	…superfacility,
4. Hybrid HPC	…l-in-the-loop, AI-around-the-loop
5. Scientific data… Interactive, data-analytics and viz	…terns such as Jupyter, scientific databases, VSCode
6. External event-triggered and API-driven workflow	time-sensitive science patterns such as function-as-a-service, microservices

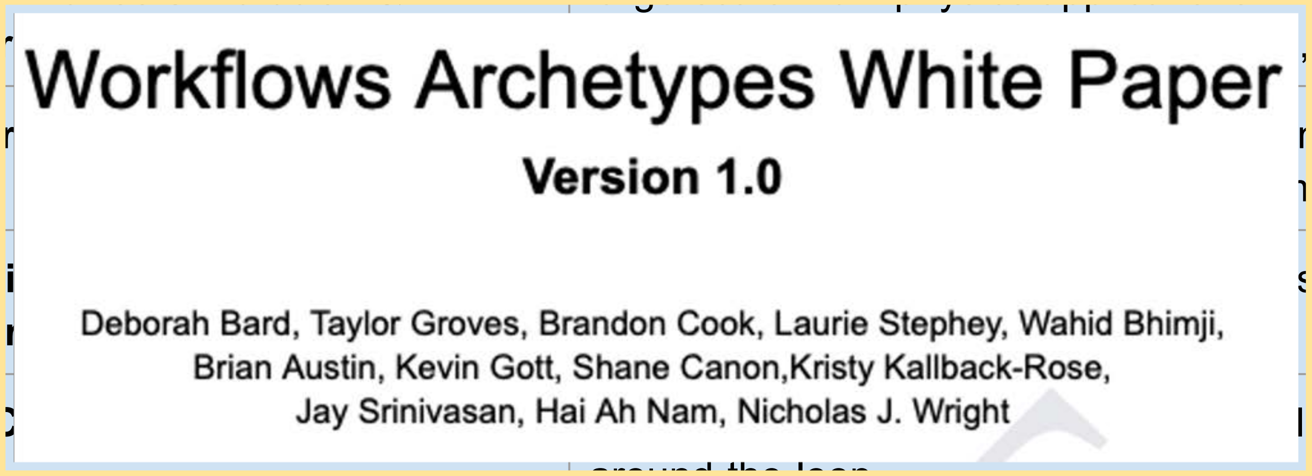

Workflows Archetypes White Paper

Version 1.0

Deborah Bard, Taylor Groves, Brandon Cook, Laurie Stephey, Wahid Bhimji, Brian Austin, Kevin Gott, Shane Canon,Kristy Kallback-Rose, Jay Srinivasan, Hai Ah Nam, Nicholas J. Wright

search for “NERSC workflows white paper”

HPC Workflows Drive Advanced Technology Capabilities

	Cloud native/ containers	QoS storage system (QSS)	End -to-end API	Network/ scheduling QoS	IRI/ Multi-site workflows	Smart networkin g	Prog. Env	Workflow Enablement Nodes (WEN, fka Spin)
1.Simulation & modeling		X	X			X	X	
2.AI	X	X	X	X	X	X	X	X
3.Cross-facility	X	X	X	X	X	X		X
4.Hybrid HPC-HPAI-HPDA	X	X	X	X	X	X	X	X
5.Scientific data lifecycle	X	X	X	X			X	X
6.Event-triggered & API-driven	X	X	X	X		X	X	X

HPC Workflows Drive Advanced Technology Capabilities

	Cloud native/ containers	QoS storage system (QSS)	End -to-end API	Network/ scheduling QoS	IRI/ Multi-site workflows	Smart networkin g	Prog. Env	Workflow Enablement Nodes (WEN, fka Spin)
1.Simulation & modeling		X	X			X	X	
2.AI	X	X	X	X	X	X	X	X
3.Cross-facility	X	X	X	X	X	X		X
4.Hybrid HPC-HPAI-HPDA	X	X	X	X	X	X	X	X
5.Scientific data lifecycle	X	X	X	X			X	X
6.Event-triggered & API-driven	X	X	X	X		X	X	X

Pink: cannot be done today
Orange: can be done only with extraordinary effort
Green: can be done today in limited way

NERSC-10 Architecture: Designed to Support Complex Simulation and Data Analysis Workflows at High Performance

- **Quality of Service** – computation, storage and networking designed to emphasize response-time plus throughput/utilization.
- **Seamlessness** – tight integration of system components to enable high performance across workflow steps.
- **Portability** – Modular workflow execution across heterogeneous HPC, edge and cloud.
- **Programmability** – APIs to manage data, execute distributed code, and interact with system resources.
- **Orchestration** – coordinate resource management across different resource domains.
- **Security** – authentication, authorization and auditing (e.g., identify proofing, access/privacy control, records of transactions).

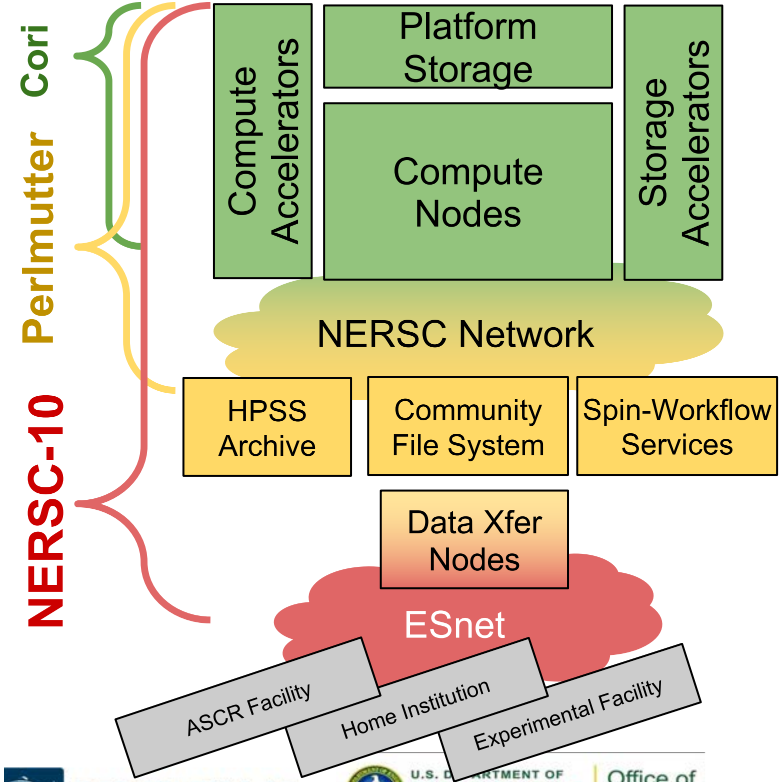

Resulting NERSC-10 Strategy

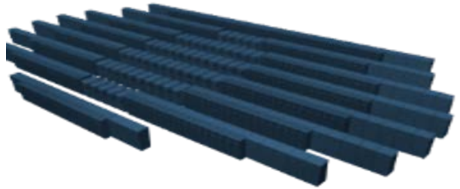

- Allow vendors who have not responded to DOE leadership-class RFP before to participate
  - Reduce number of requirements (90 TR-1, 23 TR-2, 15 TR-3)
  - No mandatory requirements - request DOE Independent Review Board (IRB)
  - Extensive, inclusive market survey – include cloud/AI vendors
- Early release of technical requirements draft – (available April 2023)
  - Frequent and often communication
- Do not prescribe a solution - describe problem
  - Partially necessitated by advanced timeline
  - Provide conduit for vendor discussion and eventual collaboration
- Enable Complex Workflows
  - Co-design software through user engagement
- Focus remains on maximizing science within existing constraints
  - Peak FLOPS will not appear in RFP

Innovation in software is key to enabling complex workflows

New capabilities: FaaS/serverless, specialized HW, AI deployment, data lifecycle, quantum…

Support usage of both ssh and Jupyter

Meet federal security requirements

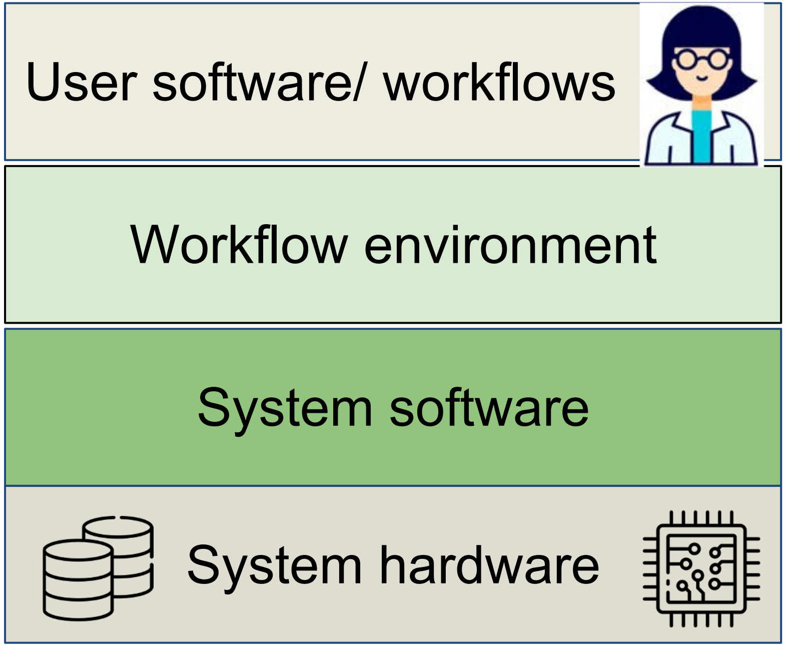

RESTful user-facing APIs support automation

System-side APIs for workflow observability, administration and reconfigurability

Containerize the user environment

Summary

- HPC is at an inflection point
  - ~~Zettascale~~
    - End of Moore's Law
  - Deep Learning training is routinely performed using supercomputers today
- N10 will deliver 10x Perlmutter performance on HPC workflows
- The N10 RFP is expected next year, system delivery in 2026

Thanks!