Quantitative methods for procuring supercomputers at NERSC – what can we learn from instrumentation?

Nicholas J. Wright
NERSC Chief Architect
& Advanced Technologies Group Lead

Instrumentation Colloquium
24 Feb 2021

Abstract

The National Energy Research Scientific Computing (NERSC) center is the production High Performance Computing (HPC) center for the Office of Science in the US Dept of Energy. NERSC purchases and deploys HPC infrastructure to enable its more than 7,000 users to perform basic research across a wide range of disciplines. In this talk, I will describe the process by which NERSC purchases supercomputers, with a focus upon the Perlmutter machine which will be delivered to NERSC/LBNL in 2021. I will describe our research efforts to instrument our current HPC resources that are targeted at gaining a deeper understanding of the ways in which NERSC is used today. I will also describe current technology trends and how they might impact the upcoming NERSC-10 & NERSC-11 procurements.

NERSC is the mission High Performance Computing facility for the DOE SC

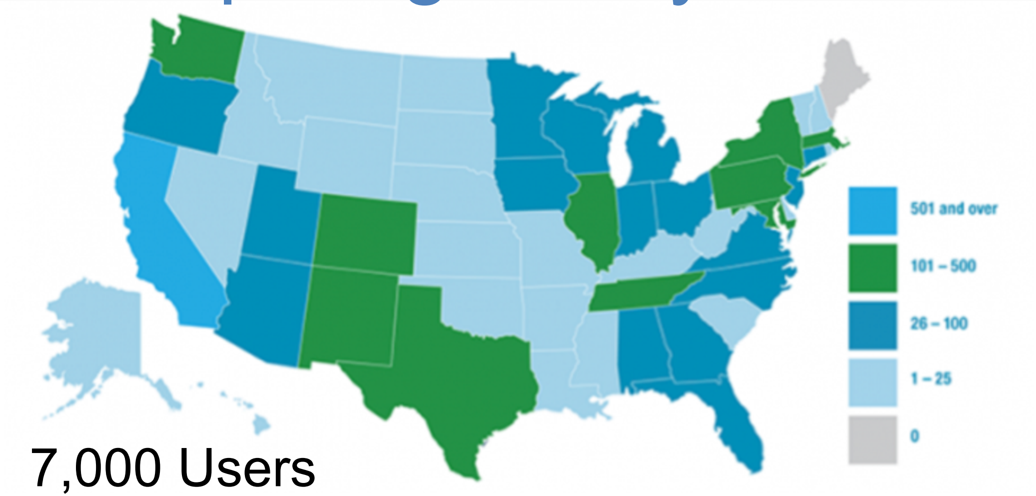

7,000 Users
800 Projects
700 Codes
2000 NERSC citations per year

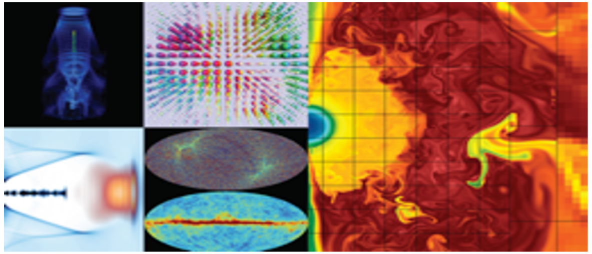

Simulations at scale

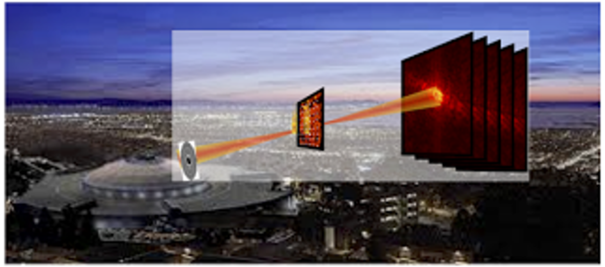

Data analysis support for DOE's experimental and observational facilities

Photo Credit: CAMERA

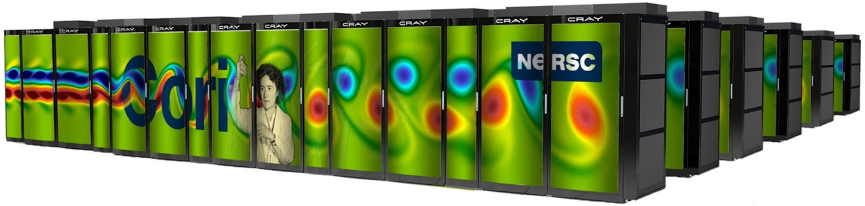

NERSC has a dual mission to advance science and the state-of-the-art in supercomputing

- We collaborate with computer companies years before a system's delivery to deploy advanced systems with new capabilities at large scale

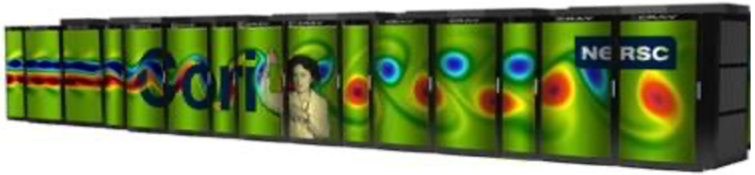

- We provide a highly customized software and programming environment for science applications

- We are tightly coupled with the workflows of DOE's experimental and observational facilities – ingesting tens of terabytes of data each day

- Our staff provide advanced application and system performance expertise to users

NERSC's infrastructure for science

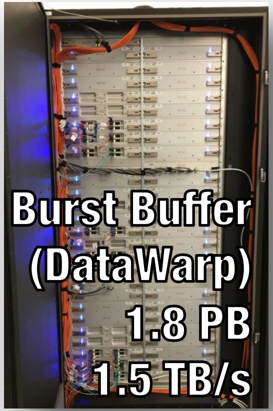

NERSC Systems Roadmap

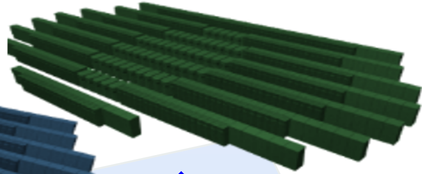

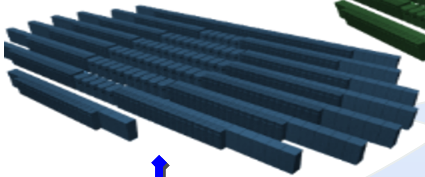

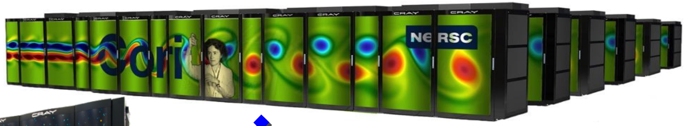

NERSC-11:
Beyond Moore

2030

NERSC-10:
Exa system

2025

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

2020

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

2016

NERSC-7:
Edison
Multicore CPU

2013

Increasing need for energy-efficient architectures

What is a Supercomputer?

- “A computer with a high level of performance as compared to a general-purpose computer” – Wikipedia
- Scientific instrument
    - Tool for doing science with

What is a Supercomputer?

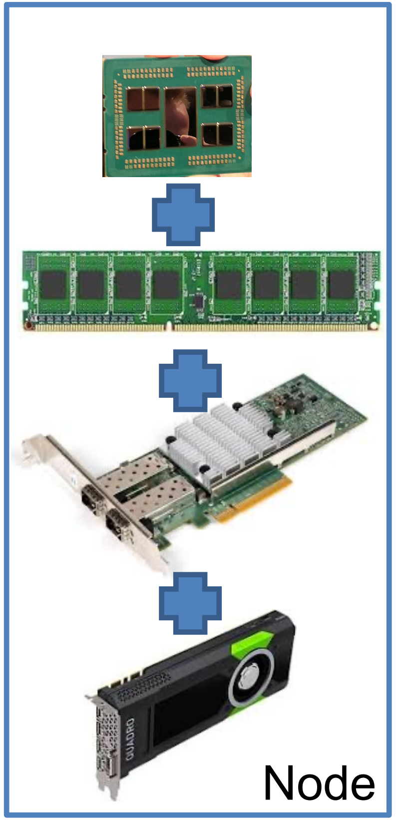

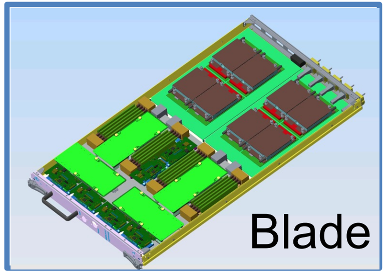

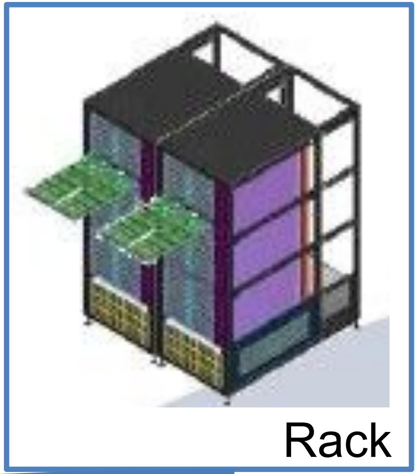

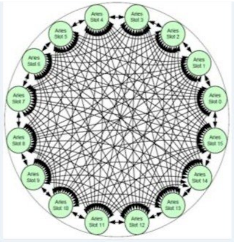

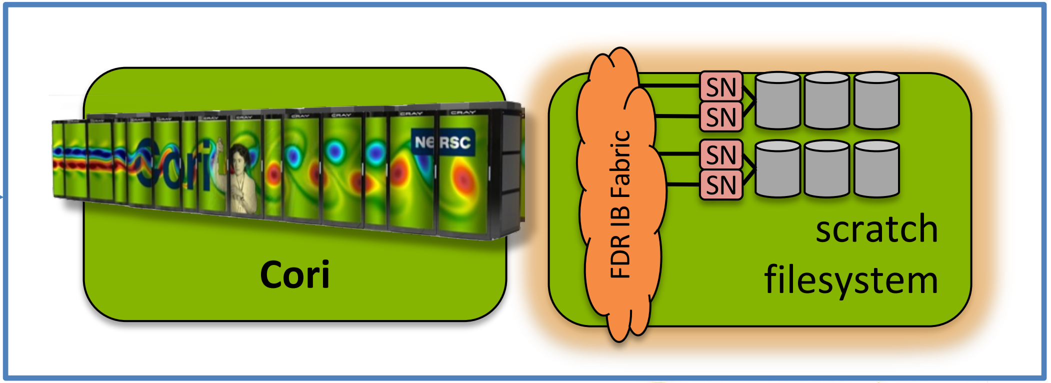

What is a Supercomputer?

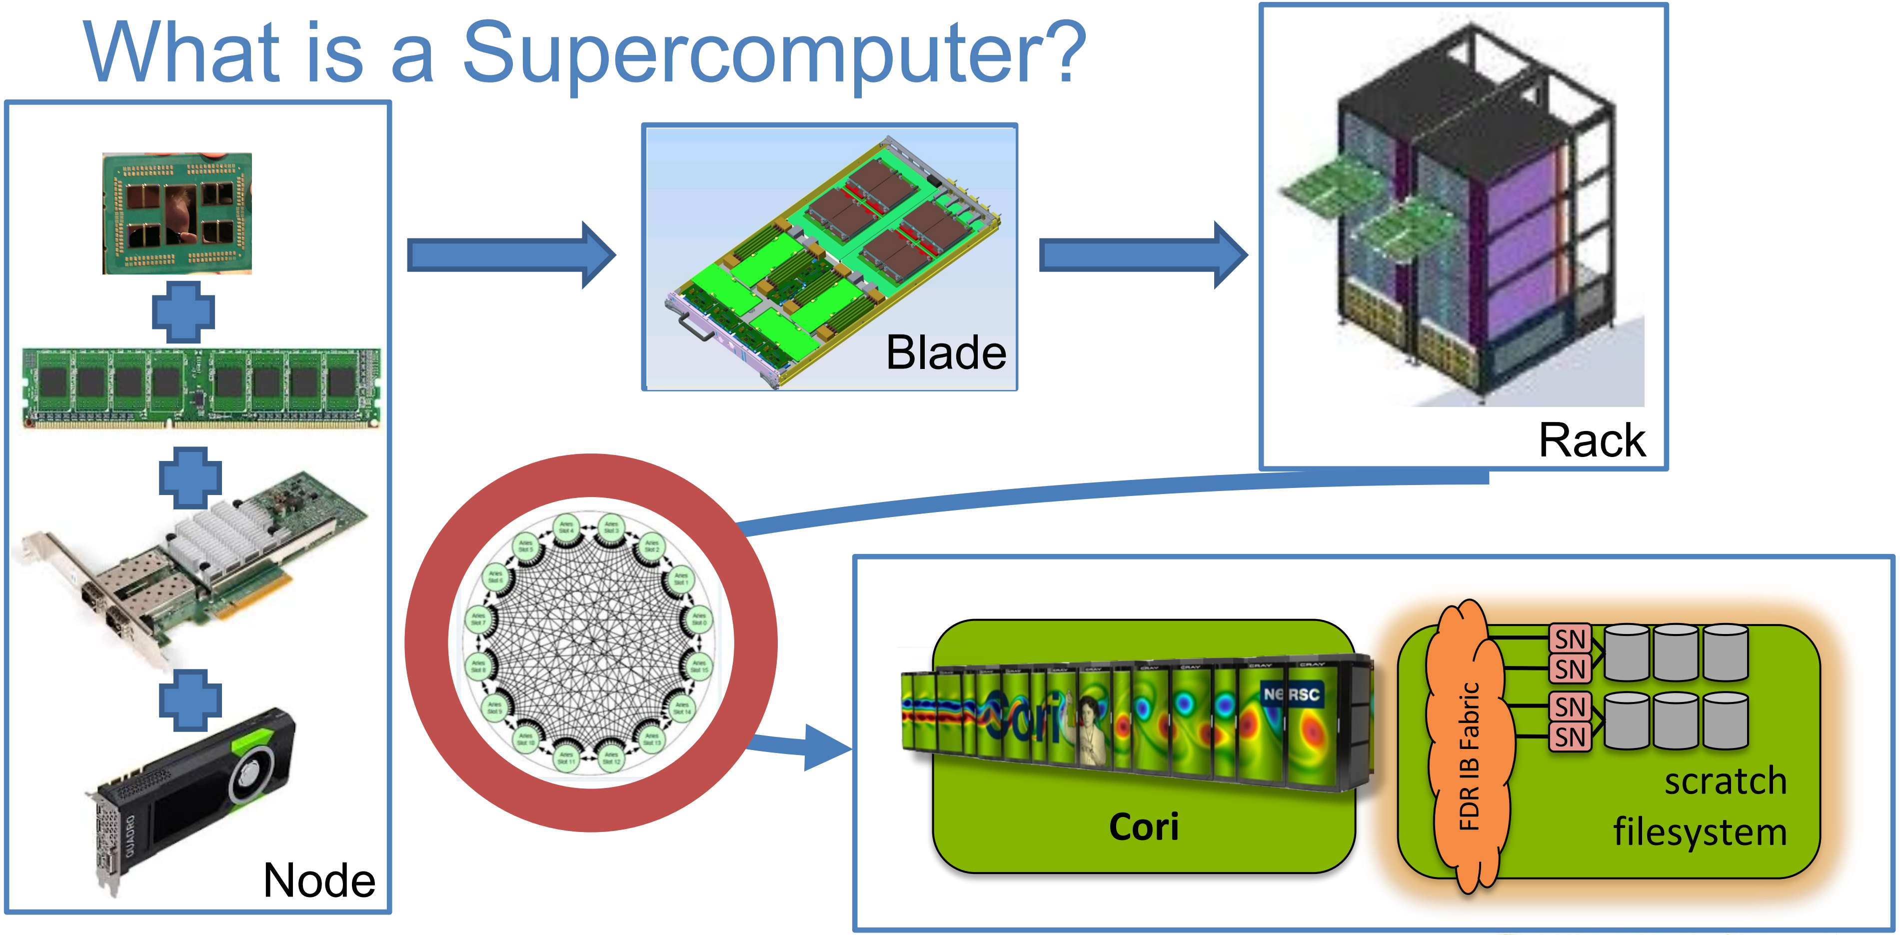

The NERSC-8 System: Cori

- Cray system w/ 9,688 Intel Knights Landing (KNL) nodes
  - Self hosted processor, manycore processor
  - 192 GB DDR, 16 GB MCDRAM
- Data Intensive Science Support
  - 2,388 Haswell processor cabinets to support data intensive apps
  - Burst Buffer to accelerate data intensive applications ~1.5TB/sec, 1.5 PB
  - 28 PB of disk, >700 GB/sec I/O bandwidth

NERSC's approach to strategic planning ~2016

User Requirements

Workload Analysis

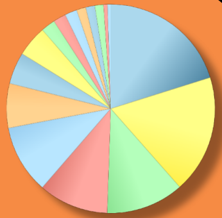

Technology Trends

NERSC Strategy

Technology Trends – 2016 looking to 2020

- GPUs emerging as completive solution for scientific computing
  - How many NERSC users can use them?
- Next generation HPC networks would be available from Cray & Mellanox
- Flash storage emerging as economically viable alternative to Harddrives for HPC
  - Can we afford enough capacity?

NERSC workload is extremely diverse, but not evenly divided.

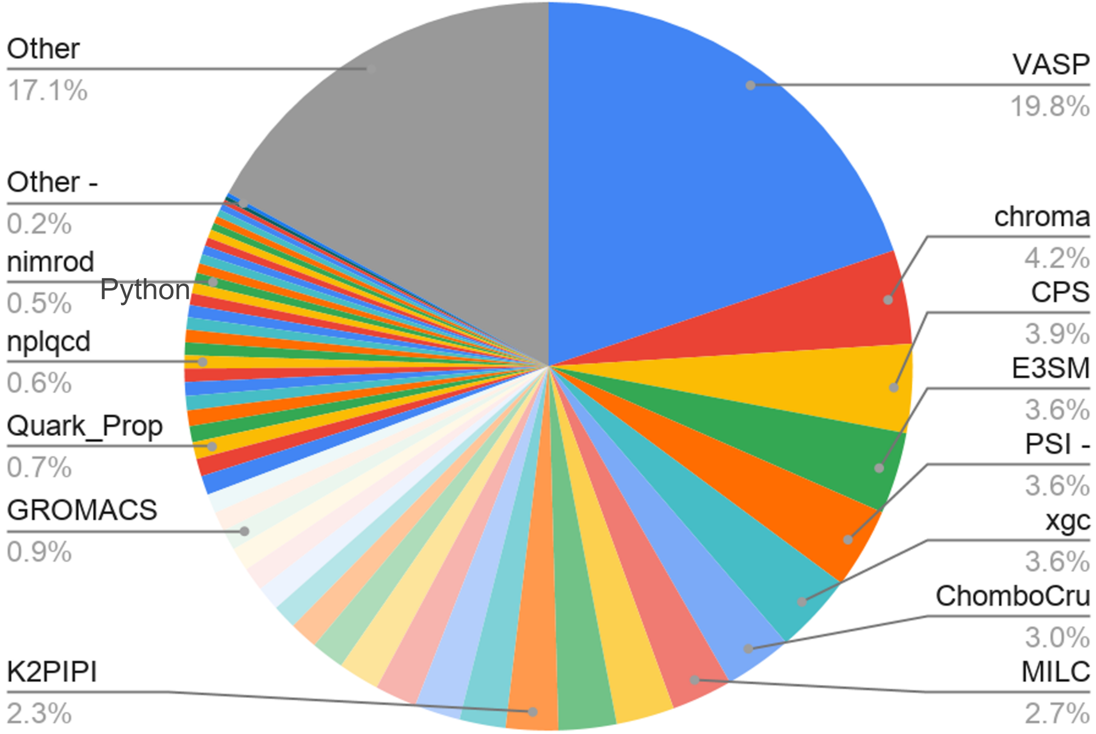

- **10 codes make up 50% of workload.**
- **20 codes make up 66% of workload.**
- **50 codes make up 84% of workload.**
- **Remaining codes (over 600) make up 16% of workload.**

Should the strategy target each user equally or each CPU hour ?

Much of the NERSC workload already runs well on GPUs

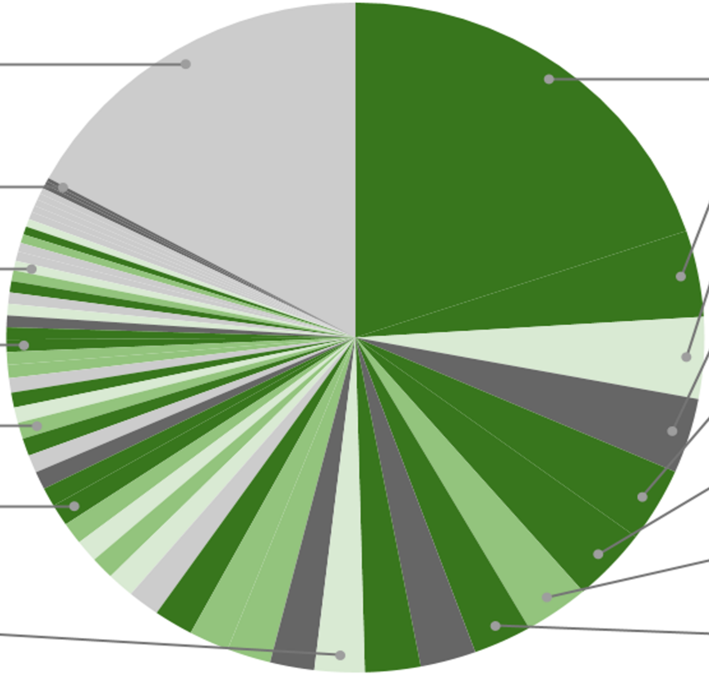

GPU Status & Description	Fraction
Enabled: Most features are ported and performant	43%
Kernels: Ports of some kernels have been documented.	8%
Proxy: Kernels in related codes have been ported	14%
Unlikely: A GPU port would require major effort.	10%
Unknown: GPU readiness cannot be assessed at this time.	25%

Hetero system design & price sensitivity: Budget for GPUs increases as GPU price drops

Chart explores an isocost design space

- **Vary the budget allocated to GPUs**
- **Assume GPU enabled applications have performance advantage = 10x per node, 3 of 8 apps are still CPU only.**
- **Examine GPU/CPU node cost ratio**

GPU / CPU $ per node	SSI increase vs. CPU-Only (@ budget %)	
8:1	None	No justification for GPUs
6:1	1.05x @ 25%	Slight justification for up to 50% of budget on GPUs
4:1	1.23x @ 50%	GPUs cost effective up to full system budget, but optimum at 50%

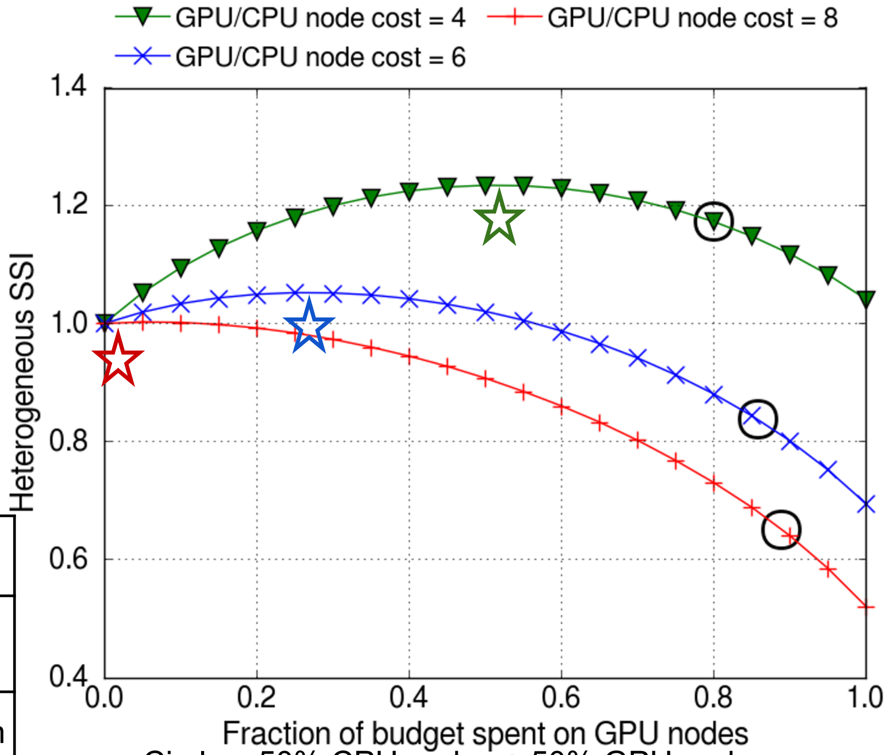

Circles: 50% CPU nodes + 50% GPU nodes
Stars: Optimal system configuration.

B. Austin, C. Daley, D. Doerfler, J. Deslippe, B. Cook, B. Friesen, T. Kurth, C. Yang, N. J. Wright, "A Metric for Evaluating Supercomputer Performance in the Era of Extreme Heterogeneity", *9th IEEE International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS18)*, November 12, 2018,

All-flash filesystem *should* be affordable in 2020

- Performance was *not* considered—we assumed:
  - performance doesn't matter with all-NVMe
  - capacity will be most scarce resource
- With a fixed budget, how much do we spend on…
  - OST capacity?
  - NVMe endurance?
  - MDT capacity for inodes and Lustre DOM?
- Gamble on price of commodities in +2 years
  - Share risk: fix price for NAND in contract
  - Allow renegotiation if 2020 price is > ± 5%

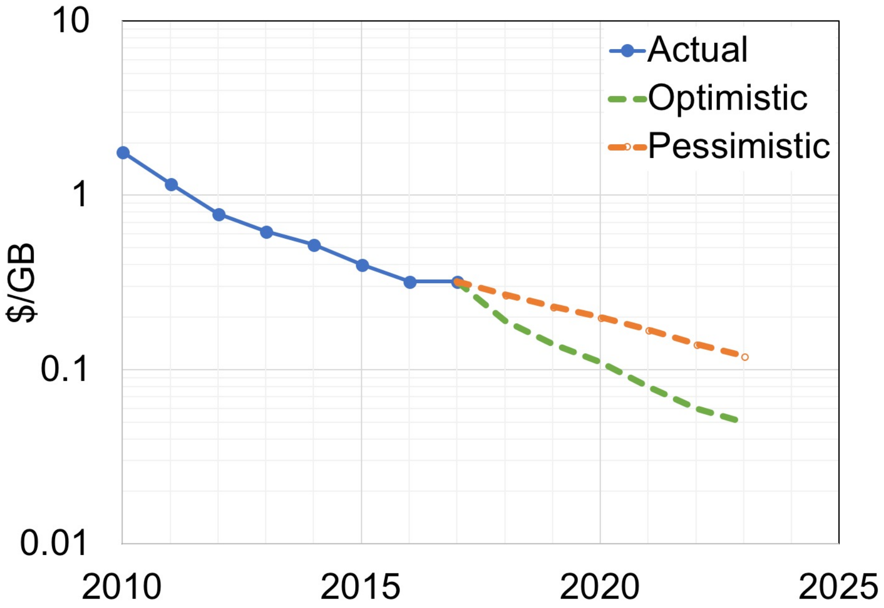

Perlmutter: A System Optimized for Science

- 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, >6x Cori’s bandwidth
- Dedicated login and high memory nodes to support complex workflows
- Delivery in early FY21

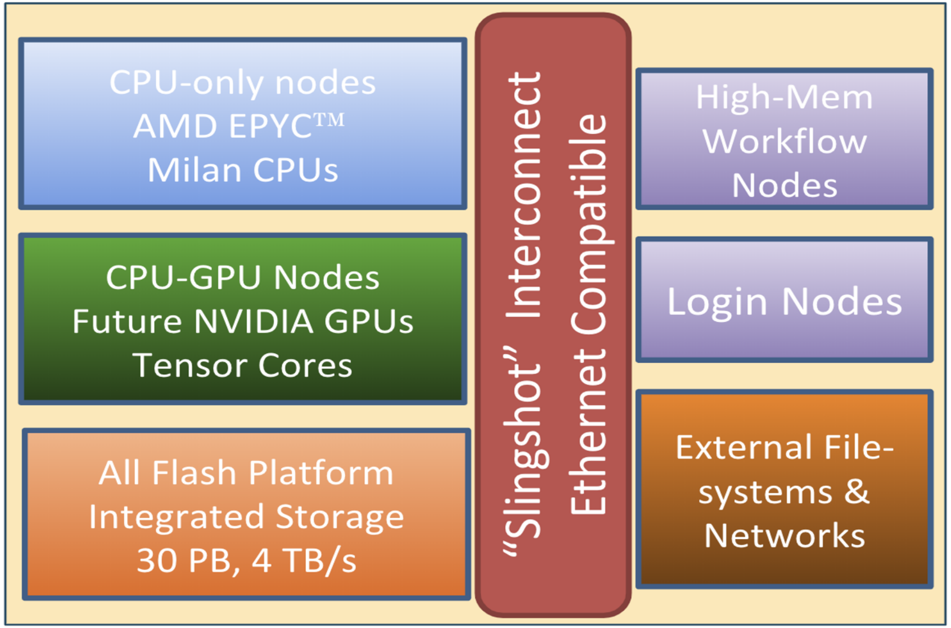

NERSC Systems Roadmap

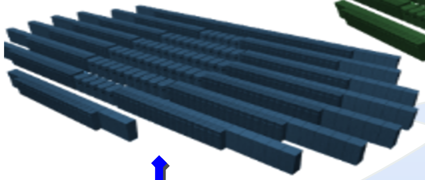

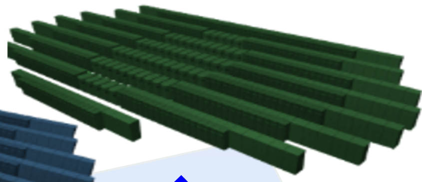

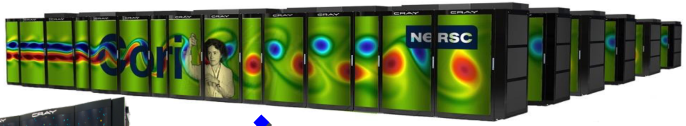

NERSC-7: Edison Multicore CPU — **2013**

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 — **2020**

NERSC-10: Exa system — **2025**

NERSC-11: Beyond Moore — **2030**

Increasing need for energy-efficient architectures

NERSC's approach to strategic planning ~2020

User Requirements

Workload Analysis

Technology Trends

NERSC Strategy

Users require more than an increase in compute hours

Requirements reviews and a multitude of workshops indicate that users need a ***significant*** increase in computational hours -- but ***also require*** a more ***integrated ecosystem*** that support new paradigms for ***data analysis, movement, management*** and ***resilience*** of scientific ***workflows***

NERSC

HEP: CMB-S4 Data Analysis and Workflow

Telescopes at the South Pole and Atacama Desert

FABRIC

ESnet
ENERGY SCIENCES NETWORK

N10 will be primary computing and archive center

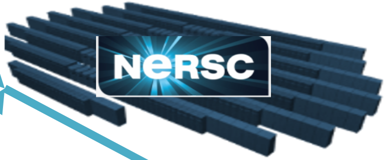

Open Science Grid

Argonne NATIONAL LABORATORY

Leadership Computing Facility

XSEDE
Extreme Science and Engineering Discovery Environment

(secondary data centers)

- Live monitoring, alerts
- Timestream simulation
- Timestream reduction to maps
- Data publication

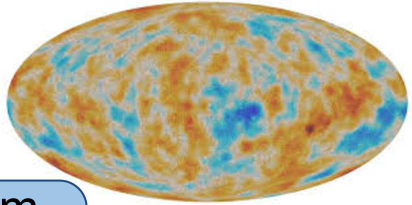

Map of Cosmic Microwave Background (CMB)

- Multi-continent workflow
- Requires ~exadays compute, O(100)PB per year
- Integration of simulation and data analysis for a large collaboration

NERSC

NERSC's approach to strategic planning ~2020

User Requirements

Workload Analysis

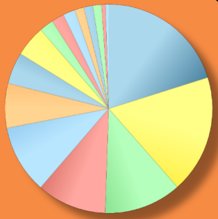

Technology Trends

HGST a Western Digital brand
DDN STORAGE
SAMSUNG
Micron
intel
ORACLE
IBM
SEAGATE
CRAY

NERSC Strategy

End of Moore's Law ?

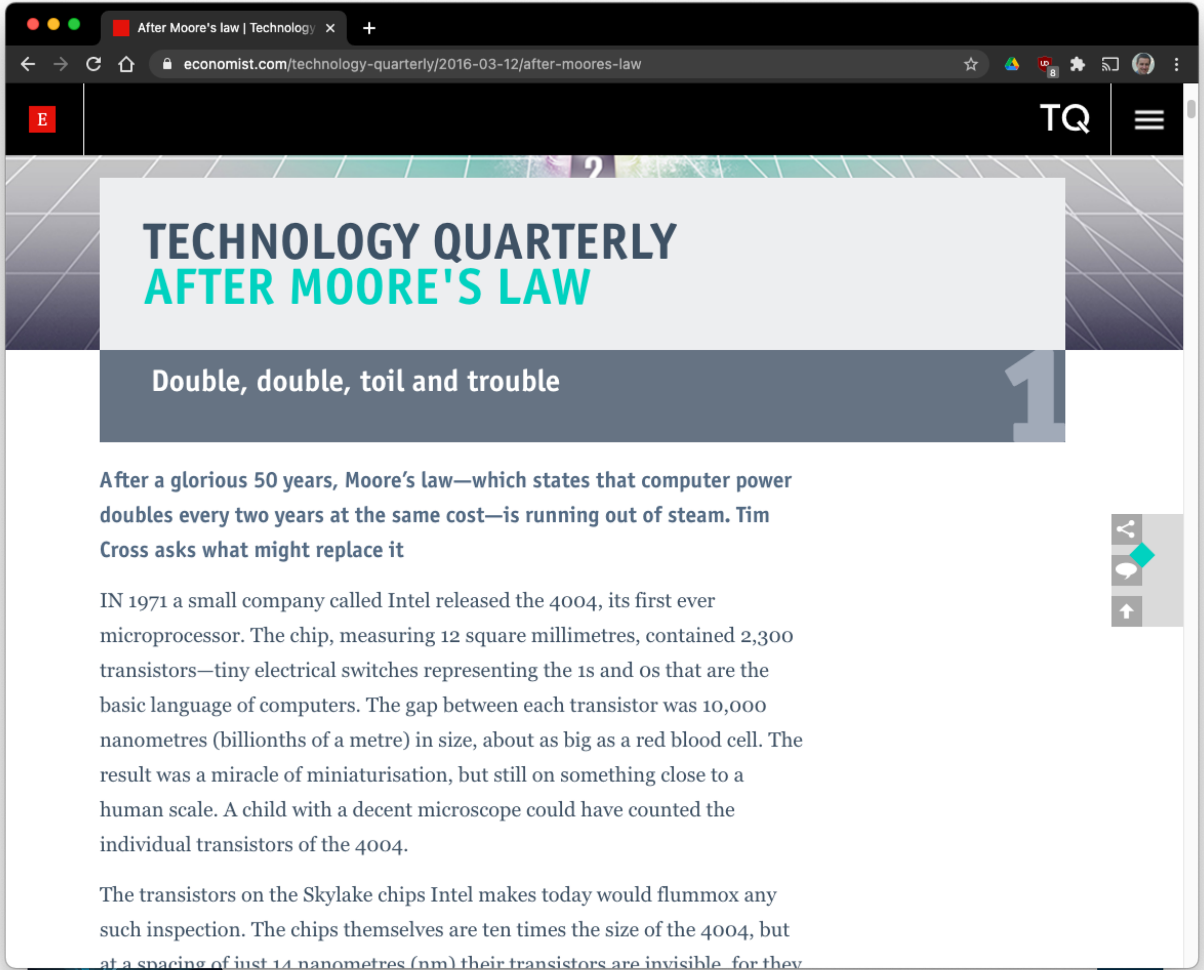
After Moore's law | Technology
economist.com/technology-quarterly/2016-03-12/after-moores-law

TQ

TECHNOLOGY QUARTERLY
AFTER MOORE'S LAW

Double, double, toil and trouble

After a glorious 50 years, Moore's law—which states that computer power doubles every two years at the same cost—is running out of steam. Tim Cross asks what might replace it

IN 1971 a small company called Intel released the 4004, its first ever microprocessor. The chip, measuring 12 square millimetres, contained 2,300 transistors—tiny electrical switches representing the 1s and 0s that are the basic language of computers. The gap between each transistor was 10,000 nanometres (billionths of a metre) in size, about as big as a red blood cell. The result was a miracle of miniaturisation, but still on something close to a human scale. A child with a decent microscope could have counted the individual transistors of the 4004.

The transistors on the Skylake chips Intel makes today would flummox any such inspection. The chips themselves are ten times the size of the 4004, but at a spacing of just 14 nanometres (nm) their transistors are invisible, for they

End of Moore's Law ?


DOI:10.1145/3282307

Innovations like domain-specific hardware, enhanced security, open instruction sets, and agile chip development will lead the way.

BY JOHN L. HENNESSY AND DAVID A. PATTERSON

A New Golden Age for Computer Architecture

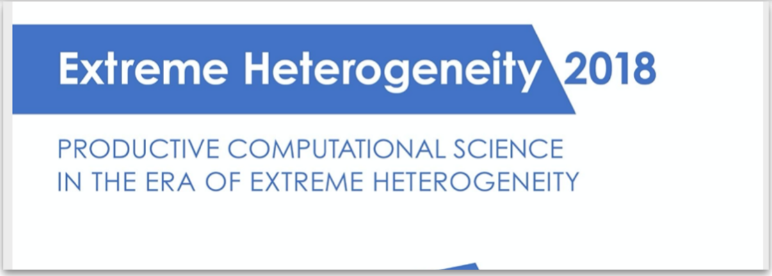
Extreme Heterogeneity 2018

PRODUCTIVE COMPUTATIONAL SCIENCE IN THE ERA OF EXTREME HETEROGENEITY

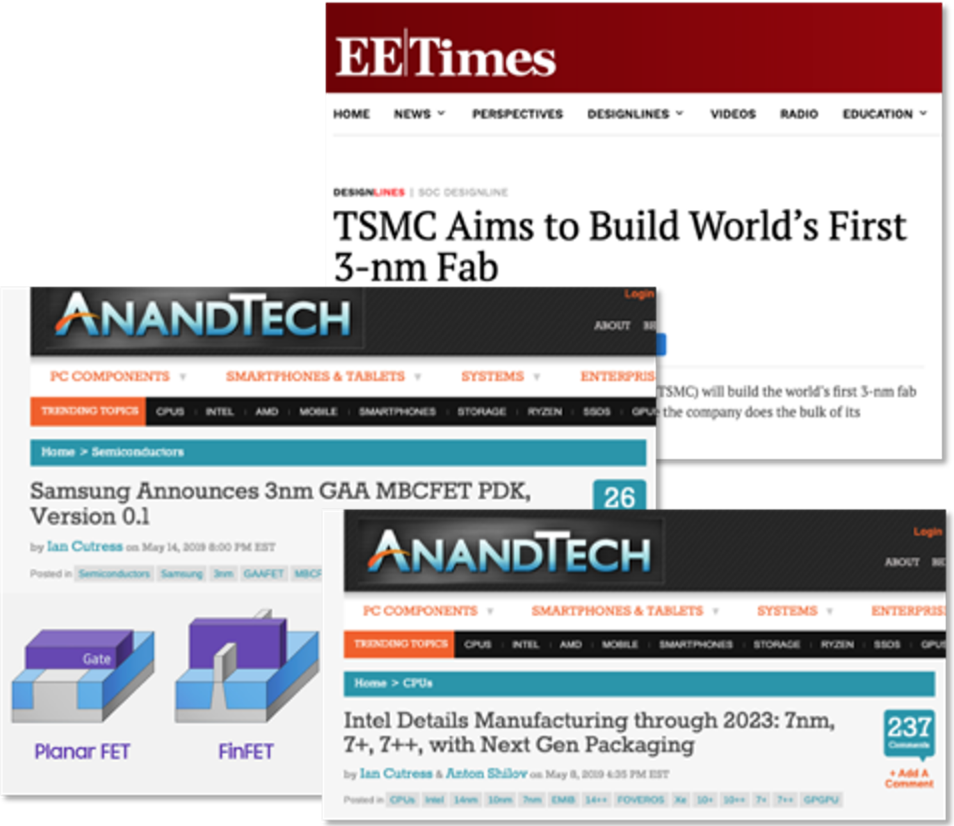
EETimes

HOME NEWS PERSPECTIVES DESIGNLINES VIDEOS RADIO EDUCATION

DESIGNLINES | SOC DESIGNLINE

TSMC Aims to Build World's First 3-nm Fab

TSMC) will build the world's first 3-nm fab
the company does the bulk of its

ANANDTECH

PC COMPONENTS SMARTPHONES & TABLETS SYSTEMS ENTERPRIS

TRENDING TOPICS CPUS INTEL AMD MOBILE SMARTPHONES STORAGE RYZEN SSDS

Home > Semiconductors

Samsung Announces 3nm GAA MBCFET PDK, Version 0.1

26

by Ian Cutress on May 14, 2019 8:00 PM EST

ANANDTECH

PC COMPONENTS SMARTPHONES & TABLETS SYSTEMS ENTERPRISE

TRENDING TOPICS CPUS INTEL AMD MOBILE SMARTPHONES STORAGE RYZEN SSDS

Home > CPUs

Intel Details Manufacturing through 2023: 7nm, 7+, 7++, with Next Gen Packaging

237

by Ian Cutress & Anton Shilov on May 8, 2019 4:35 PM EST

End of Moore's Law ?

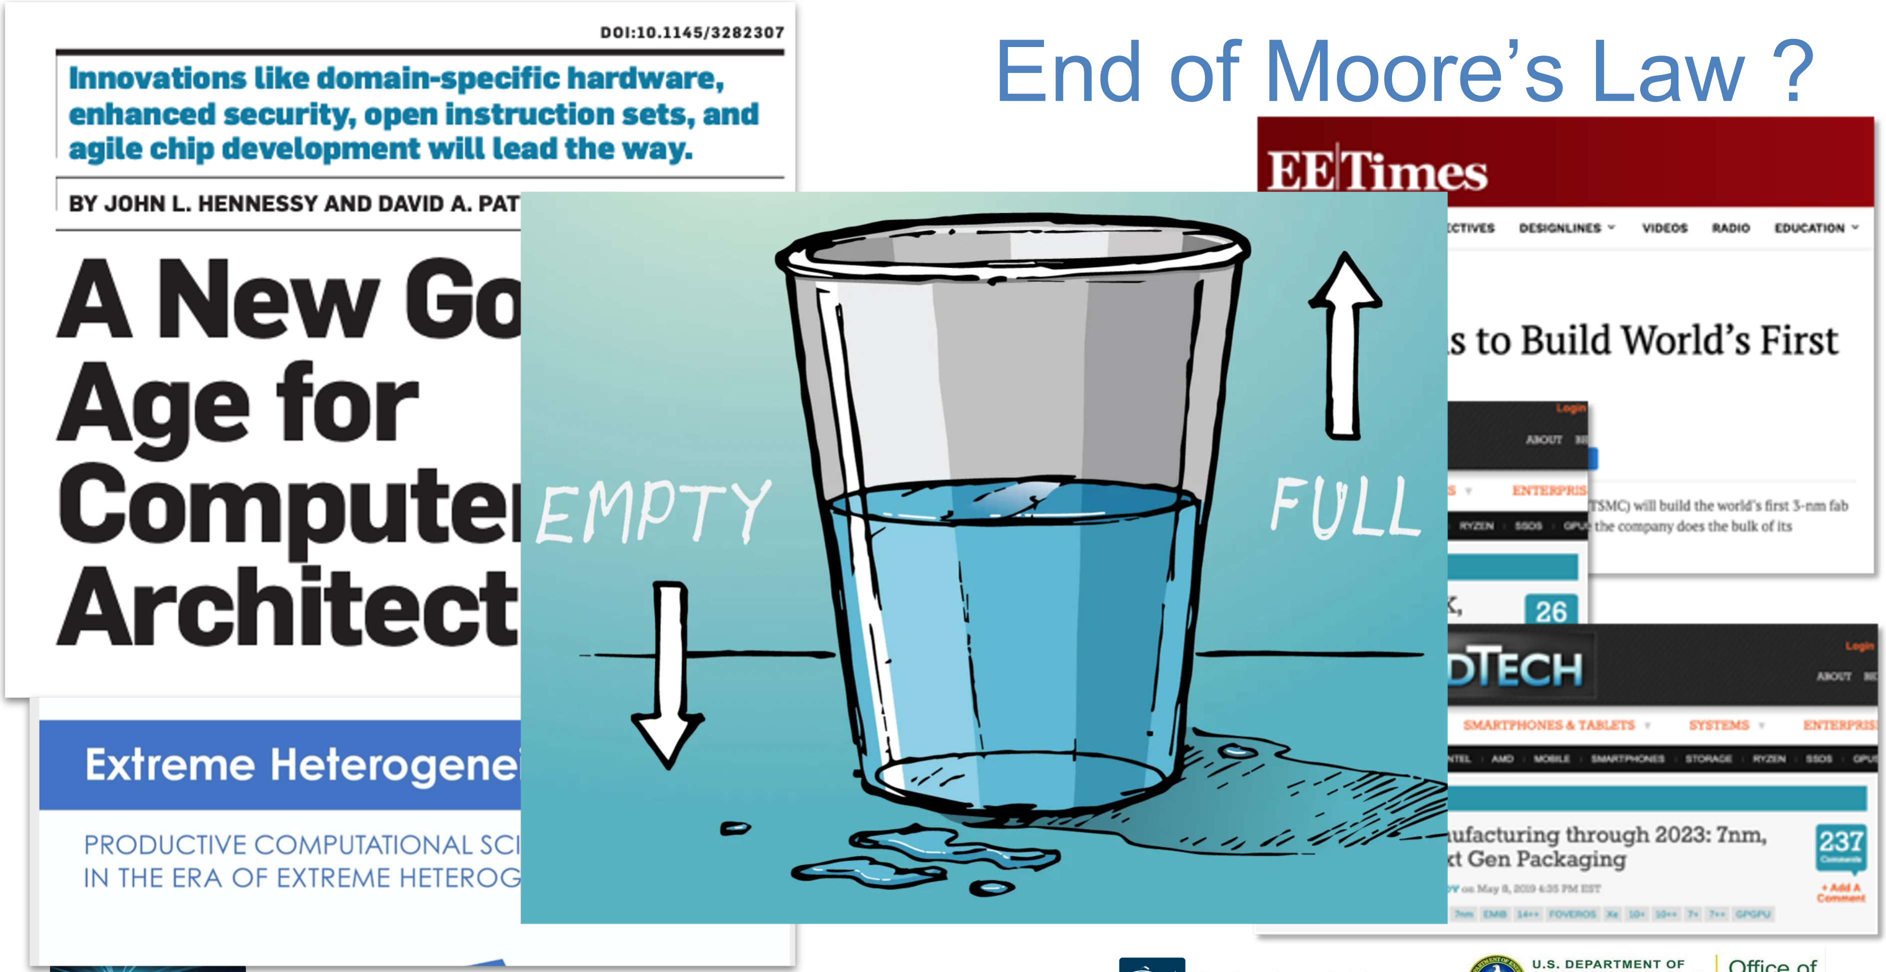

BERKELEY LAB
Bringing Science Solutions to the World
U.S. DEPARTMENT OF ENERGY
Office of Science

Hardware Technology Trends

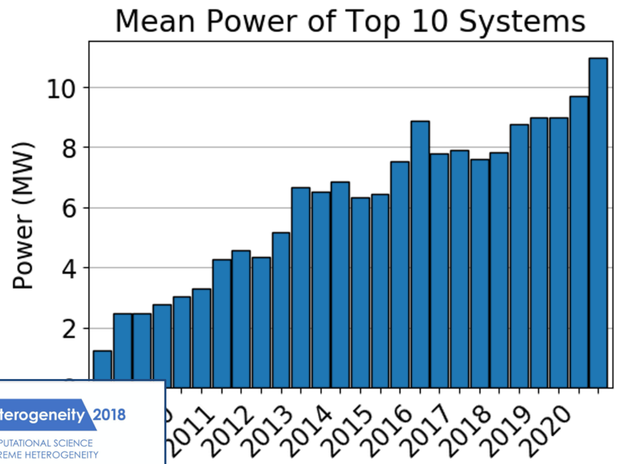

Moore's law is slowing down

- Flops/$ continues to increase
- Flops/W also increasing
  (more performance = more power)

Extreme heterogeneity is emerging

- Compute in network, computational storage
- Specialized AI accelerators
- FPGAs, dataflow, non-von Neumann

Reconfigurable computing

- disaggregated storage, memory, compute
- software-defined storage, networking, computing

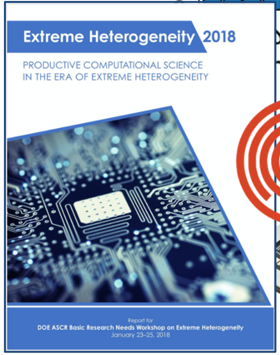

cerebras

CXL Compute Express Link

Software Technology Trends

Service-oriented architectures, microservices enable resilience and extreme scale

- Containerized services (Docker, Kubernetes)
- Lambda functions and "serverless" computing

Software-defined/programmable infrastructure

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

AI for operations and resource management

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

OPENSHIFT

FSx

splunk>

DeepMind

NERSC-10 – summary so far

- Technology trends - Extreme Heterogeneity is coming!
    - Many more potential computational elements will be available
        - Which ones will work for NERSC?
- User Requirements- Workflows are emerging as a new usage modalities
    - How can we adapt our workload analysis to respond?

NERSC's approach to strategic planning

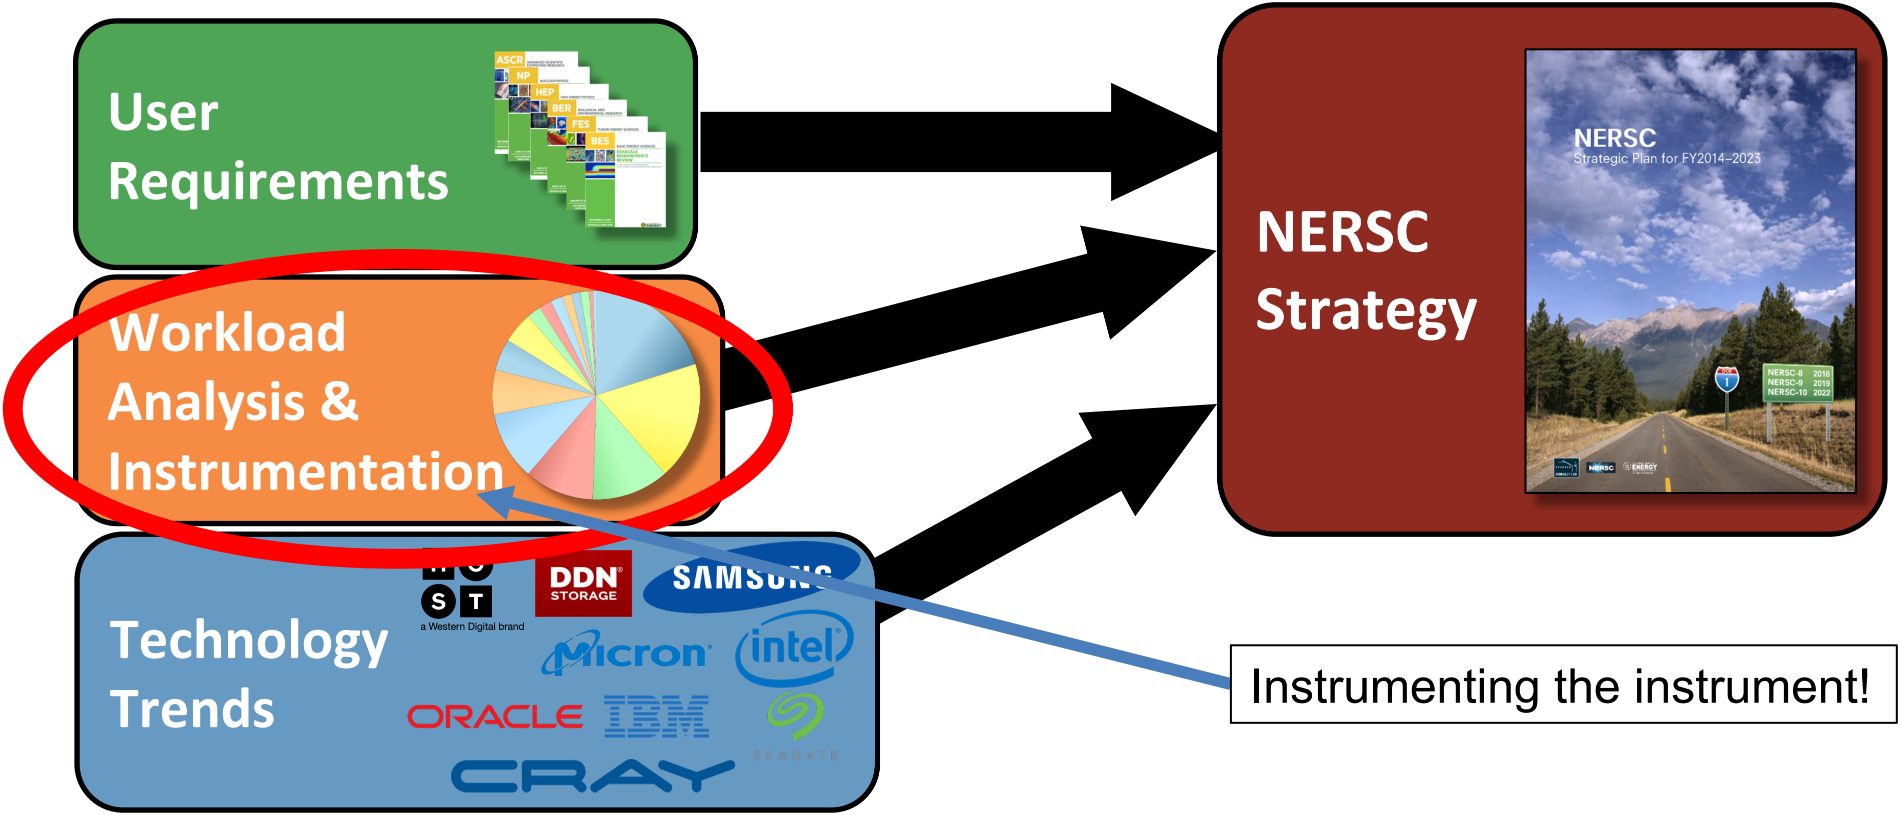

Instrumenting a Supercomputer

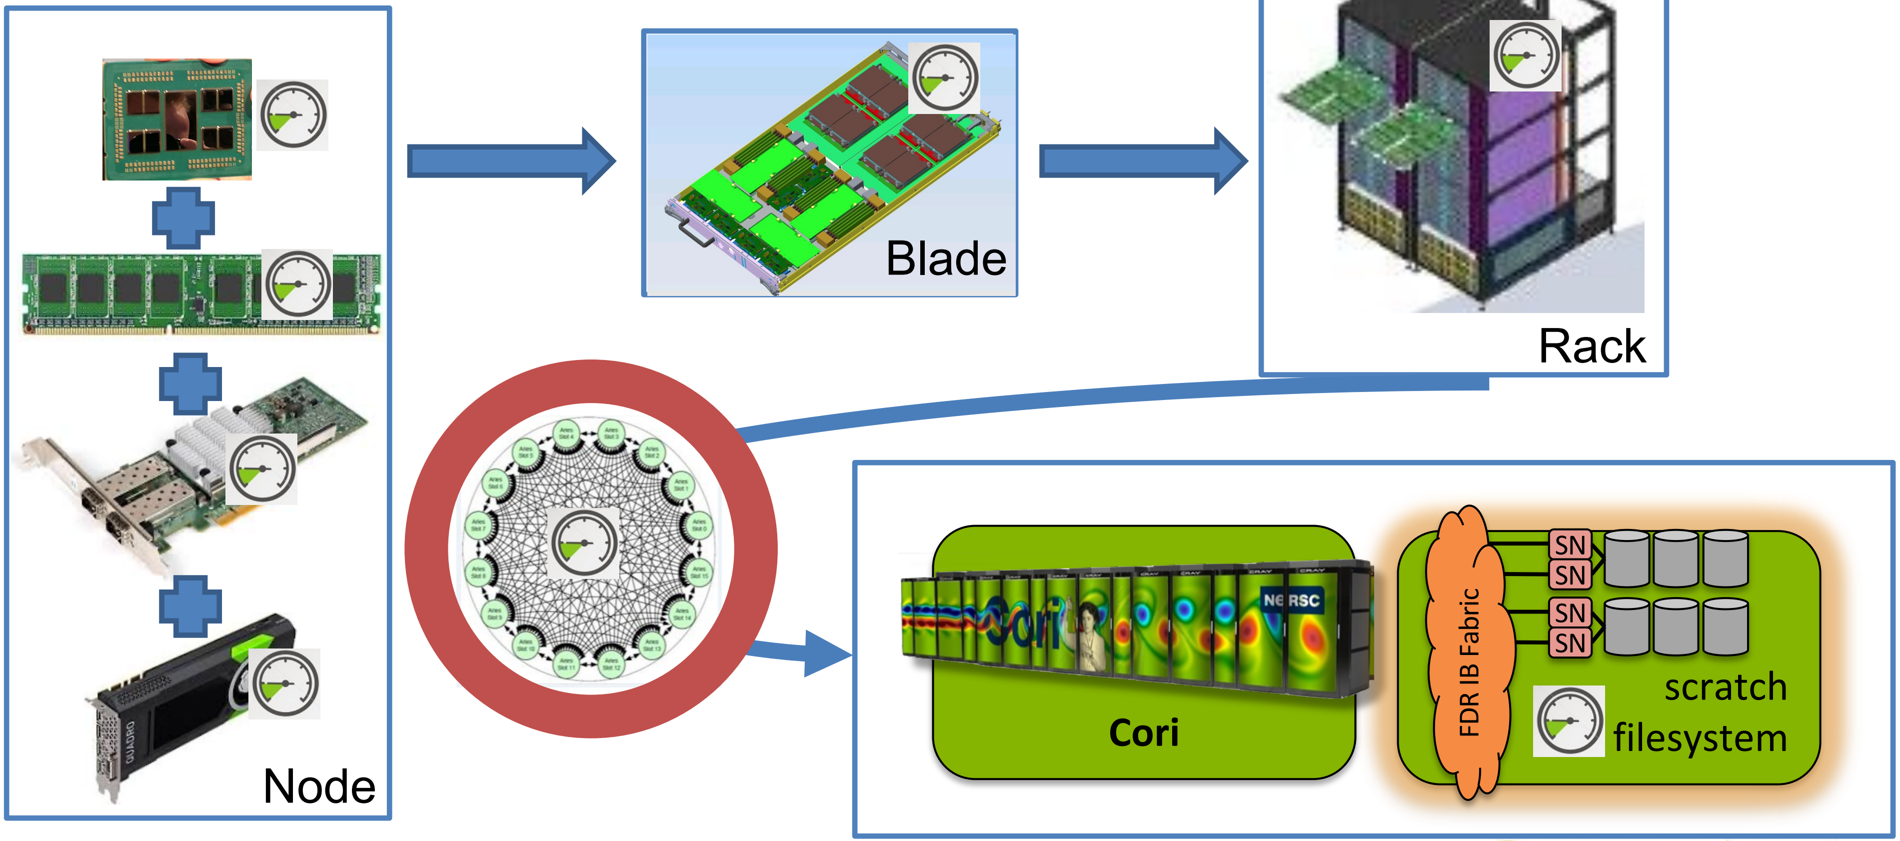

Scientific computing is more than compute!

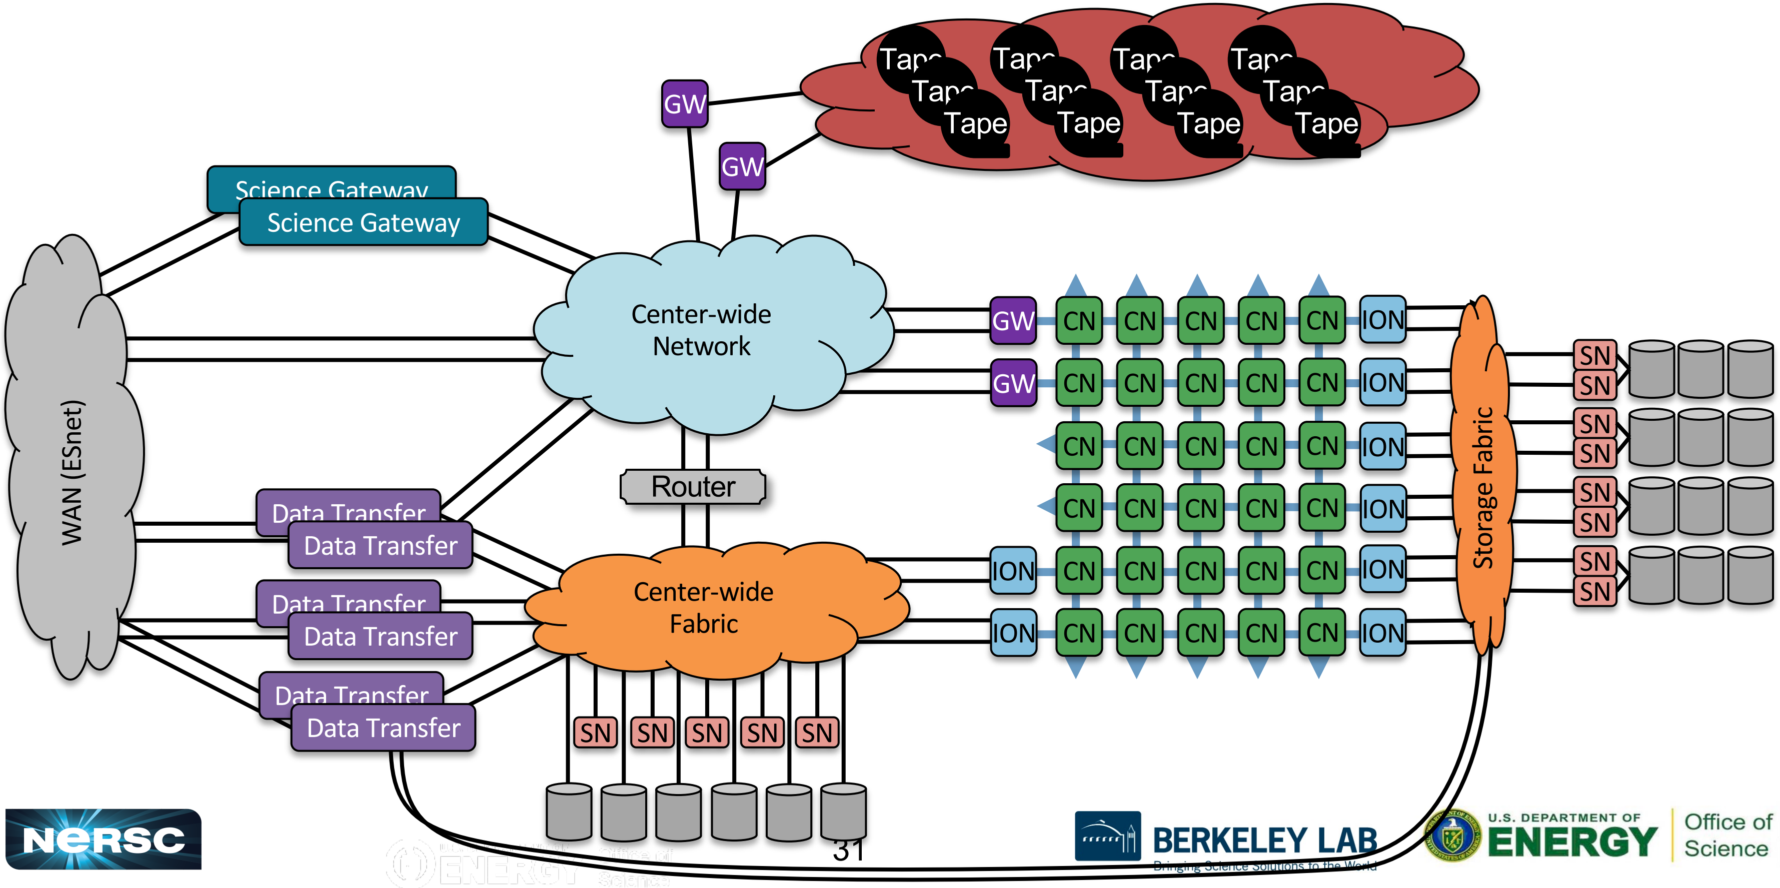

Scientific computing is more than compute!

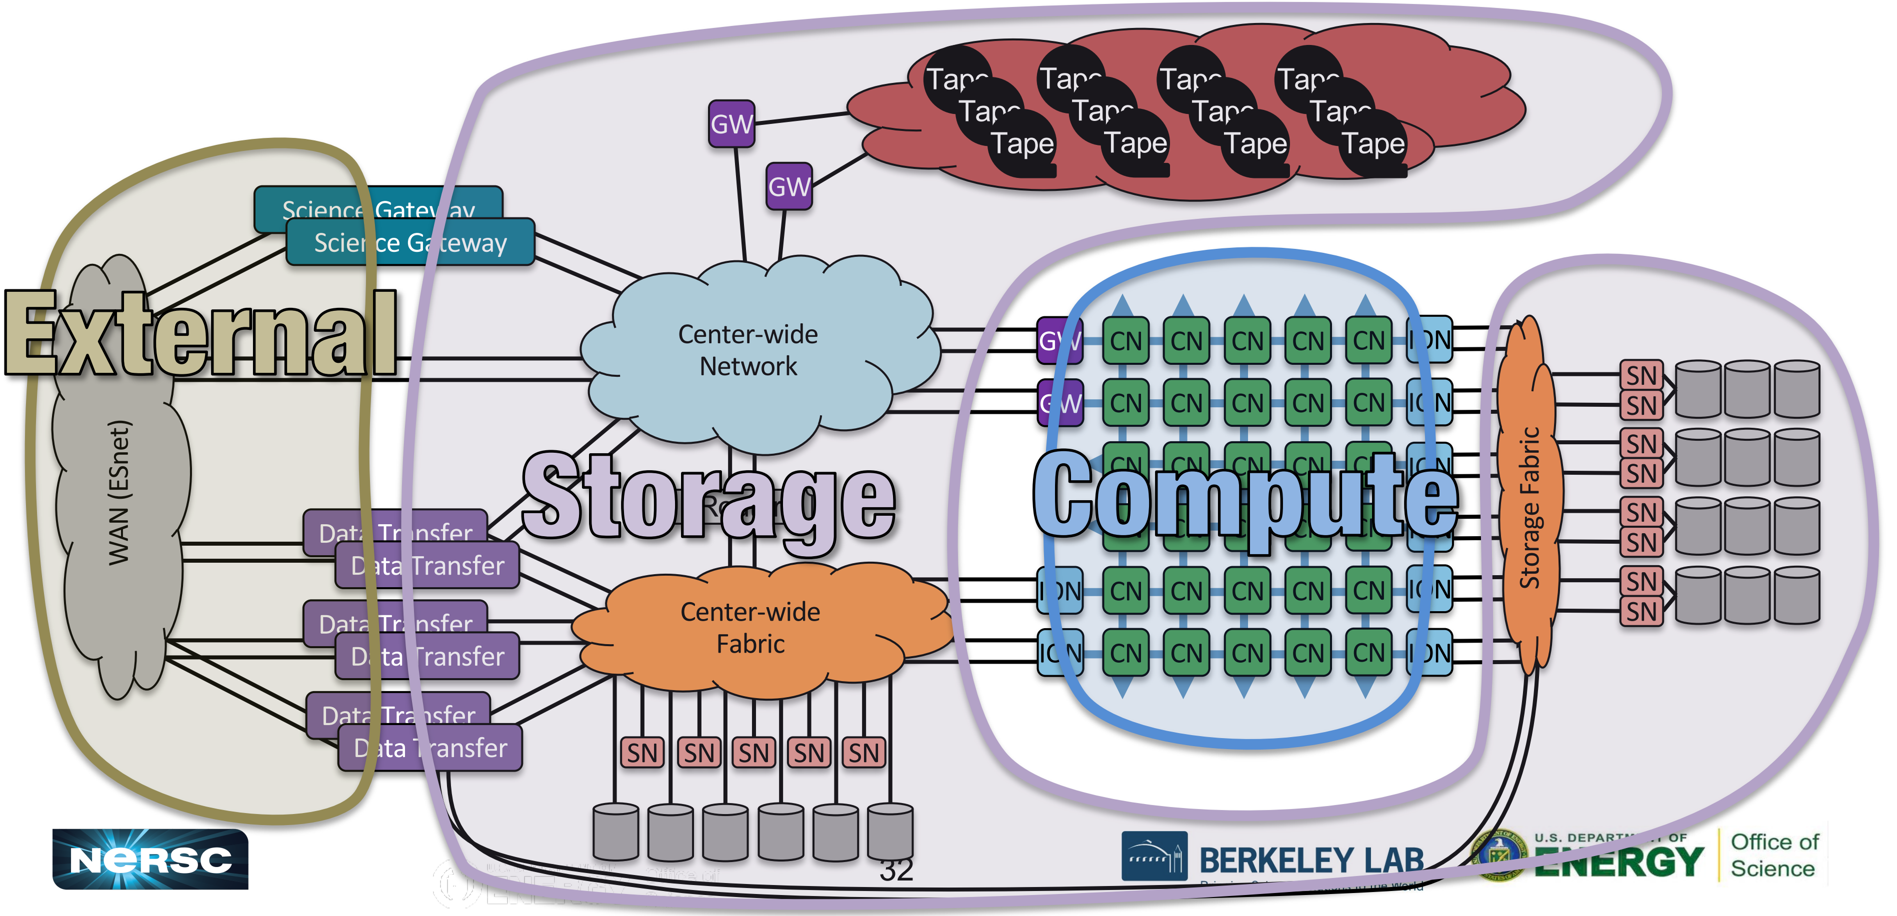

Scientific computing is more than compute!

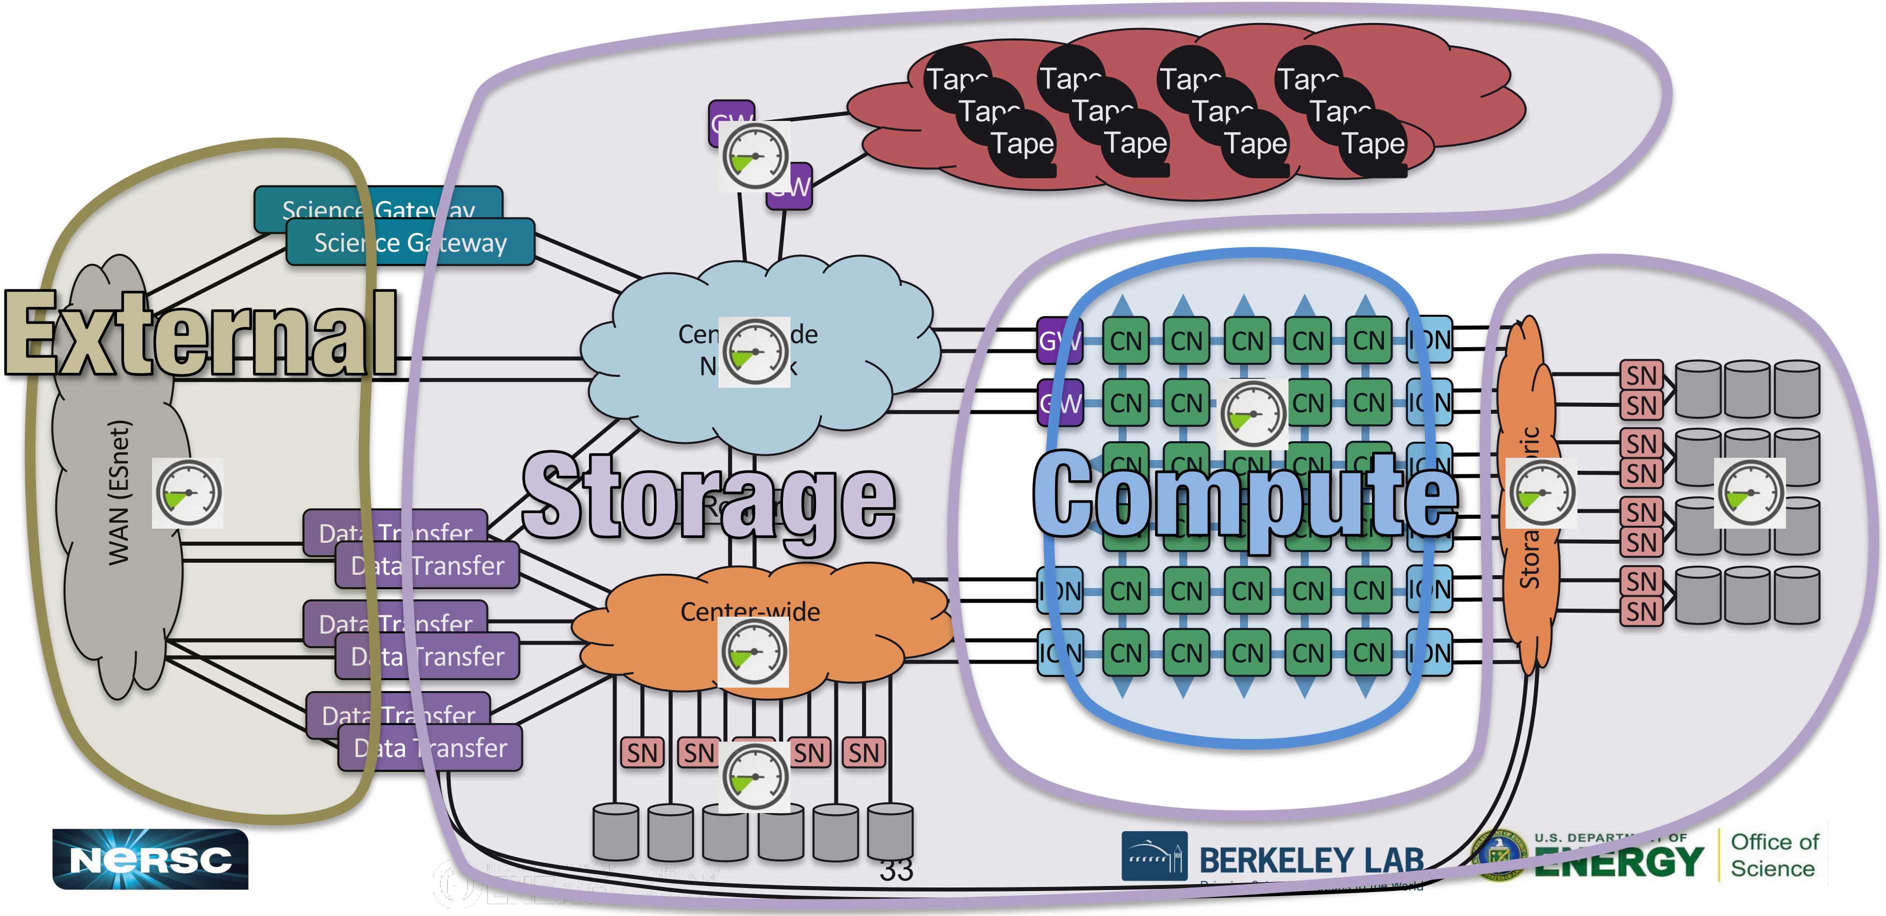

What is possible with this approach?

May 1 – August 1, 2019
- 194 million transfers
- 78.6 PiB data moved

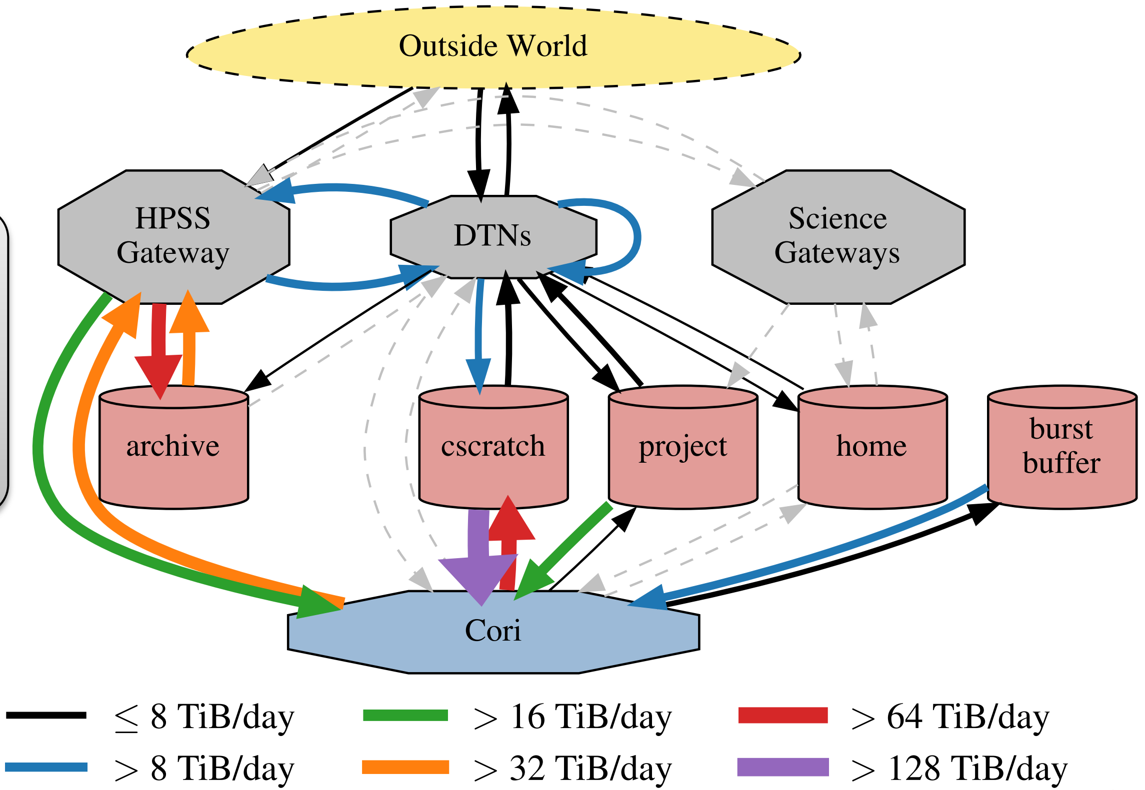

Understanding data motion in the modern hpc data center GK Lockwood, S Snyder, S Byna, P Carns, NJ Wright
2019 IEEE/ACM Fourth International Parallel Data Systems Workshop (PDSW), 74-83

Few users result in the most transfers

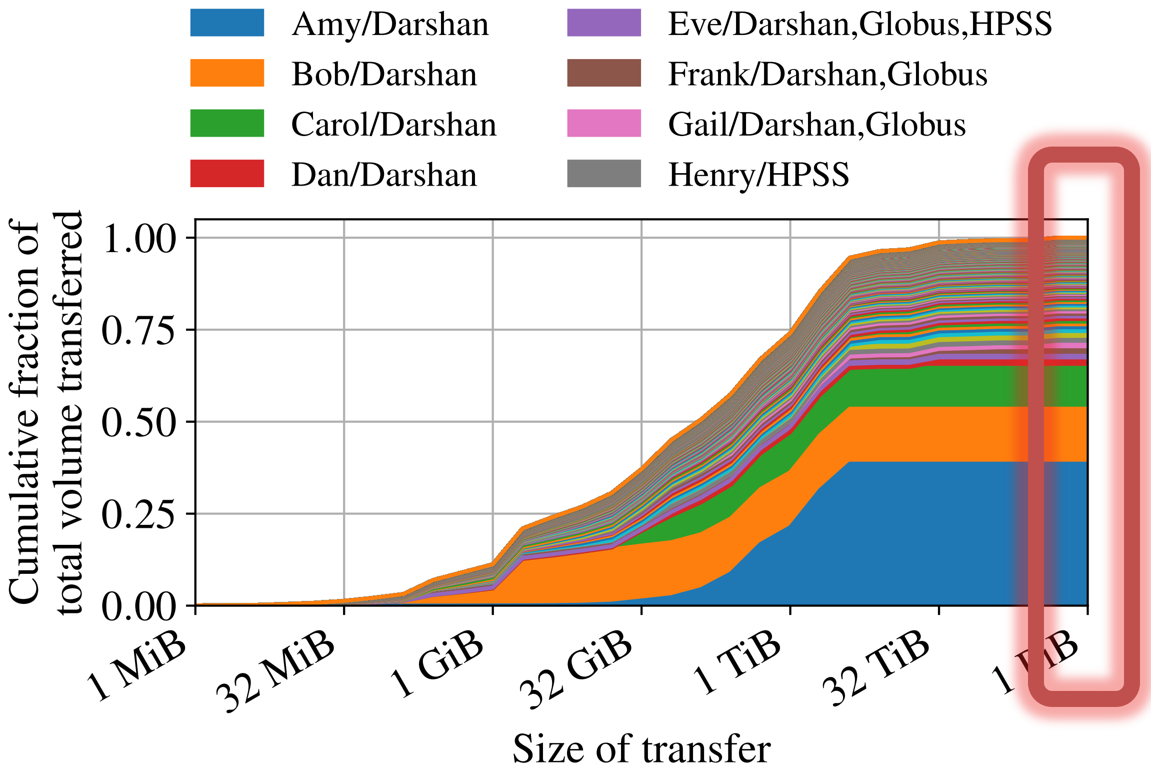

- 1,562 unique users
- Top 4 users = 66% of volume transferred
- Users 5-8 = 5.8%
  - All used multiple transfer vectors
  - Henry is a storage-only user

Understanding data motion in the modern hpc data center GK Lockwood, S Snyder, S Byna, P Carns, NJ Wright
2019 IEEE/ACM Fourth International Parallel Data Systems Workshop (PDSW), 74-83

Total Knowledge of I/O (TOKIO) Framework

Transforms monitoring data from across the data center into answers to answer "why is my I/O slow?"

TOKIO

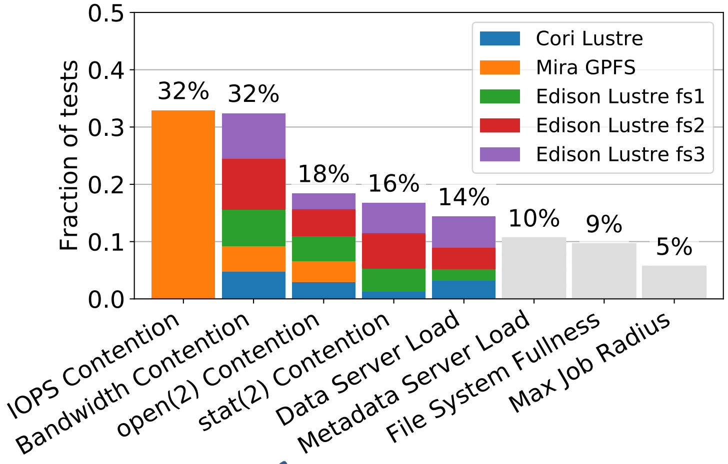

Lockwood, G.K. et al. 2018. A Year in the Life of a Parallel File System. *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis* (Piscataway, NJ, USA, Nov. 2018), 931–943.

https://www.nersc.gov/research-and-development/tokio/
https://www.github.com/nersc/pytokio

NERSC

BERKELEY LAB
Bringing Science Solutions to the World

U.S. DEPARTMENT OF ENERGY
Office of Science

Supporting Workflows at NERSC

Scientific Achievement

- **Towards understanding large scientific workflows running at NERSC, we develop two methods that identify temporal connections and data dependencies in user jobs analyzed three months of log data for Cori.**

Significance and Impact

- **Our work has helped identified key workflow execution and I/O patterns. Our analysis shows that a) Sequence+Parallel pattern is dominant in time-window workflows b) Single-Job and Sequence are common patterns (>95%) for data-dependent workflows, and c) Single-Job workflows may not effectively use all the CPUs. Our results give us new insights into I/O patterns of HPC workloads, showing that a) workflows with Single-Job and Sequence patterns predominantly read/write less than one GB of data, and b) parallelism seems correlated to the amount of data read in Parallel workflows.**

- ## Research Details

Using batch queue, I/O logs to understand workflows on NERSC

Identified workflow patterns using two methods and opportunities to improve resource utilization in workflows. Interacted with user groups to understand data and workflow journeys.

Using Gridftp log analyses:

Many workflows are transferring in and out over 100 GB of data. Workflows retain the transferred data at the NERSC file systems for over a week. Data is moved across the different filesystems at NERSC during entire workflow lifecycle

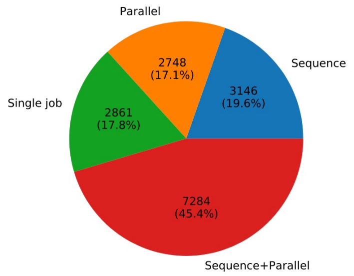

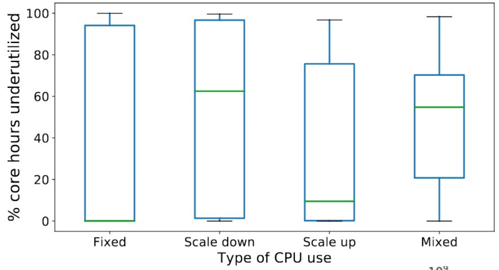

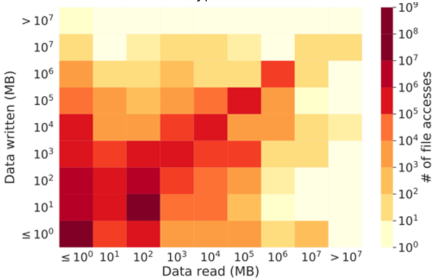

Devarshi Ghoshal, Brian Austin, Deborah Bard, Christopher Daley, Glenn Lockwood,Nicholas J. Wright, Lavanya Ramakrishnan, Characterizing Scientific Workflows on HPC Systems using Logs, In Submission.

What’s next ?

- Add more sources of data
- More sophisticated analysis techniques
    - Combine multiple data sources
    - Machine Learning?
- Enable programmability & automation

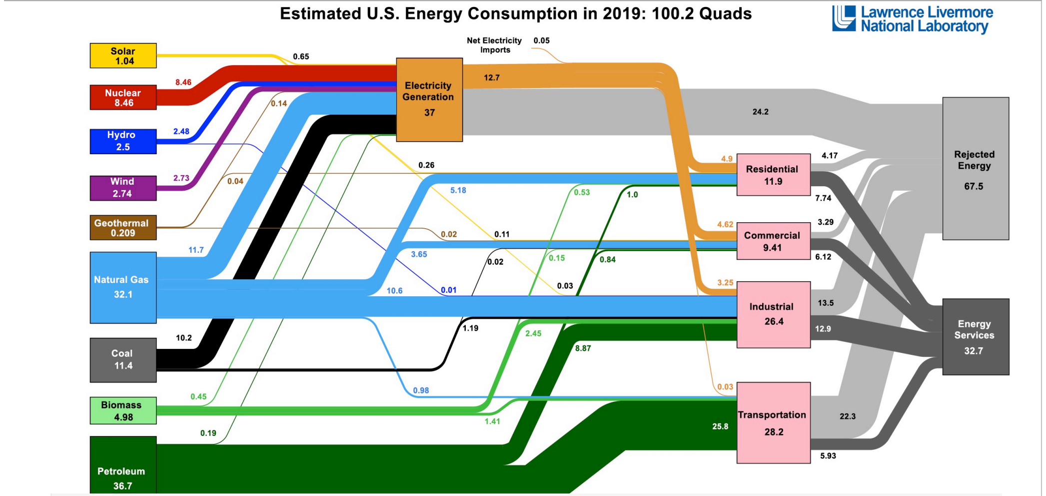

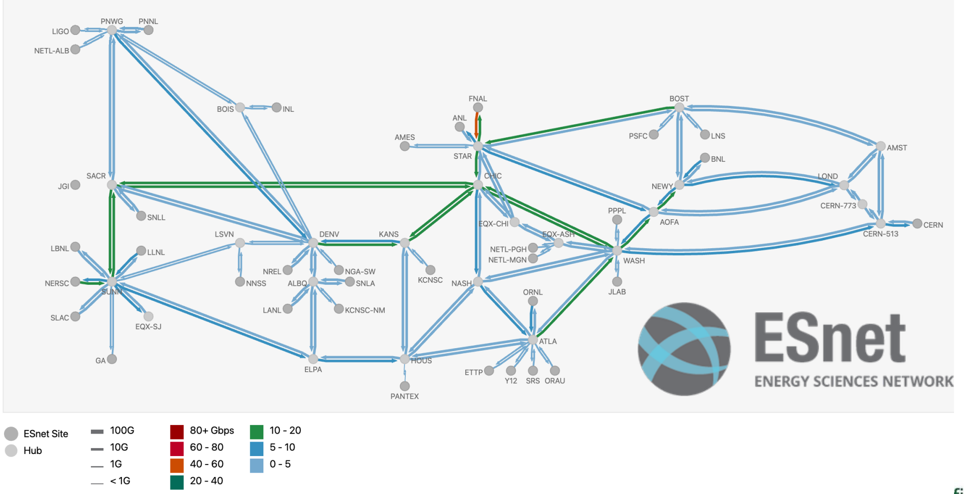

NERSC-10: Architecture Optimized for Workflows

Extreme heterogeneity offers unique optimization opportunities for workflows

- **Dense/low-precision math**: GPUs, AI accelerators
- **Low-latency processing**: CPUs
- **I/O acceleration**: Smart NICs/SSDs
- **Extreme IOPS/bandwidth**: Nonvolatile memory
- **Extreme availability/resilience**: Object stores

Complexity and heterogeneity managed using complementary technologies

- **Programmable infrastructure**: avoid downfalls of one-size-fits-all, monolithic architecture
- **AI and automation**: reduces complexity

NERSC-10 will be heterogeneous and dynamically composable to deliver on-demand, resilient workflow acceleration across the data center

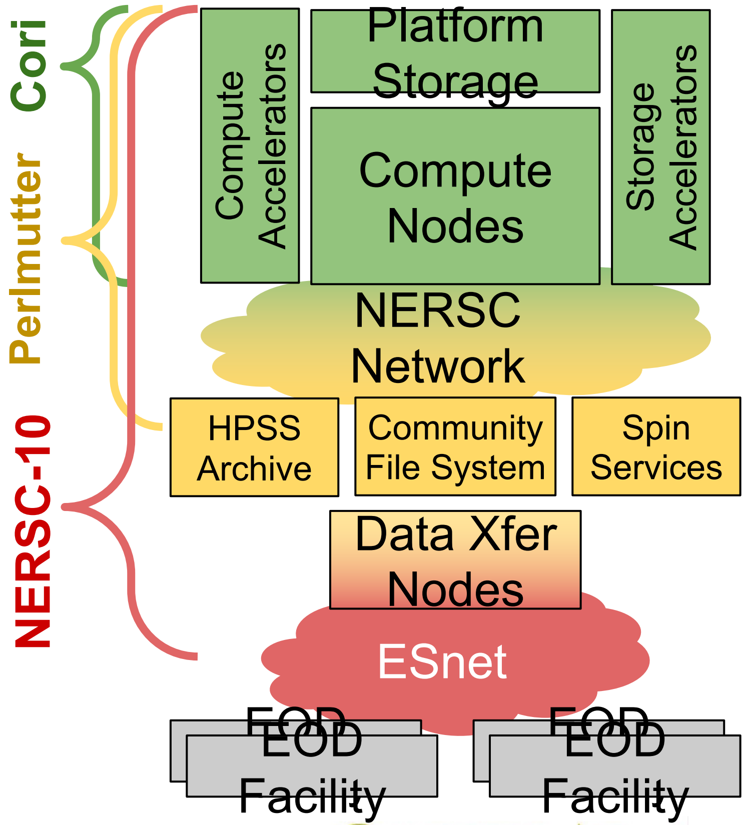

NERSC-10: Programmable data center in practice

NERSC-10 will be programmable to optimize for each workflow

1. User requests hardware resources, connections between them, and data placement
2. System schedules CPU, accelerators, storage, networking, and data movement
3. Same resources are later reconfigured to adapt to new requirements

NERSC-10 will achieve this by embracing technology trends

- Disaggregated, software-defined infrastructure to connect heterogeneous components
- AI and automation to manage
  - complexity of scheduling and operations
  - data movement between reconfigurations
  - complexity for users - sensible defaults

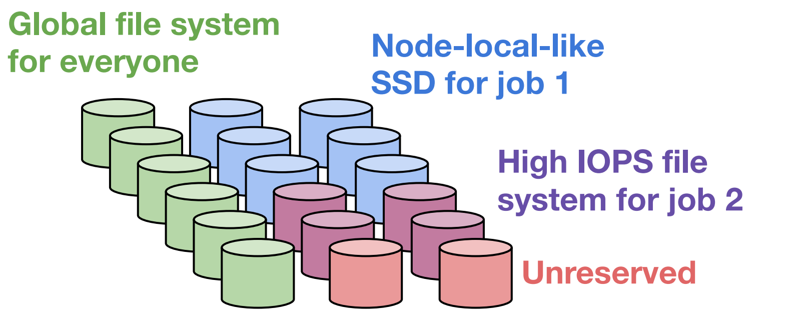

Later that day...

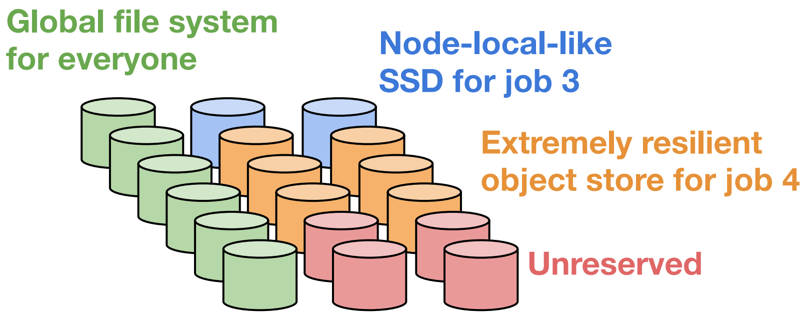

Summary

- NERSC's workload is evolving and complex workflows are emerging as the primary usage modality
  - more users from experimental facilities
  - increase in AI usage for simulation and data analysis
- NERSC-10 will accelerate these complex workflows by enabling users to program the data center by holistically allocating heterogeneous resources
- Instrumenting the NERSC data center will help us navigate the design space and enable automation
  - Can *you* help ? We are hiring !

Thank you !

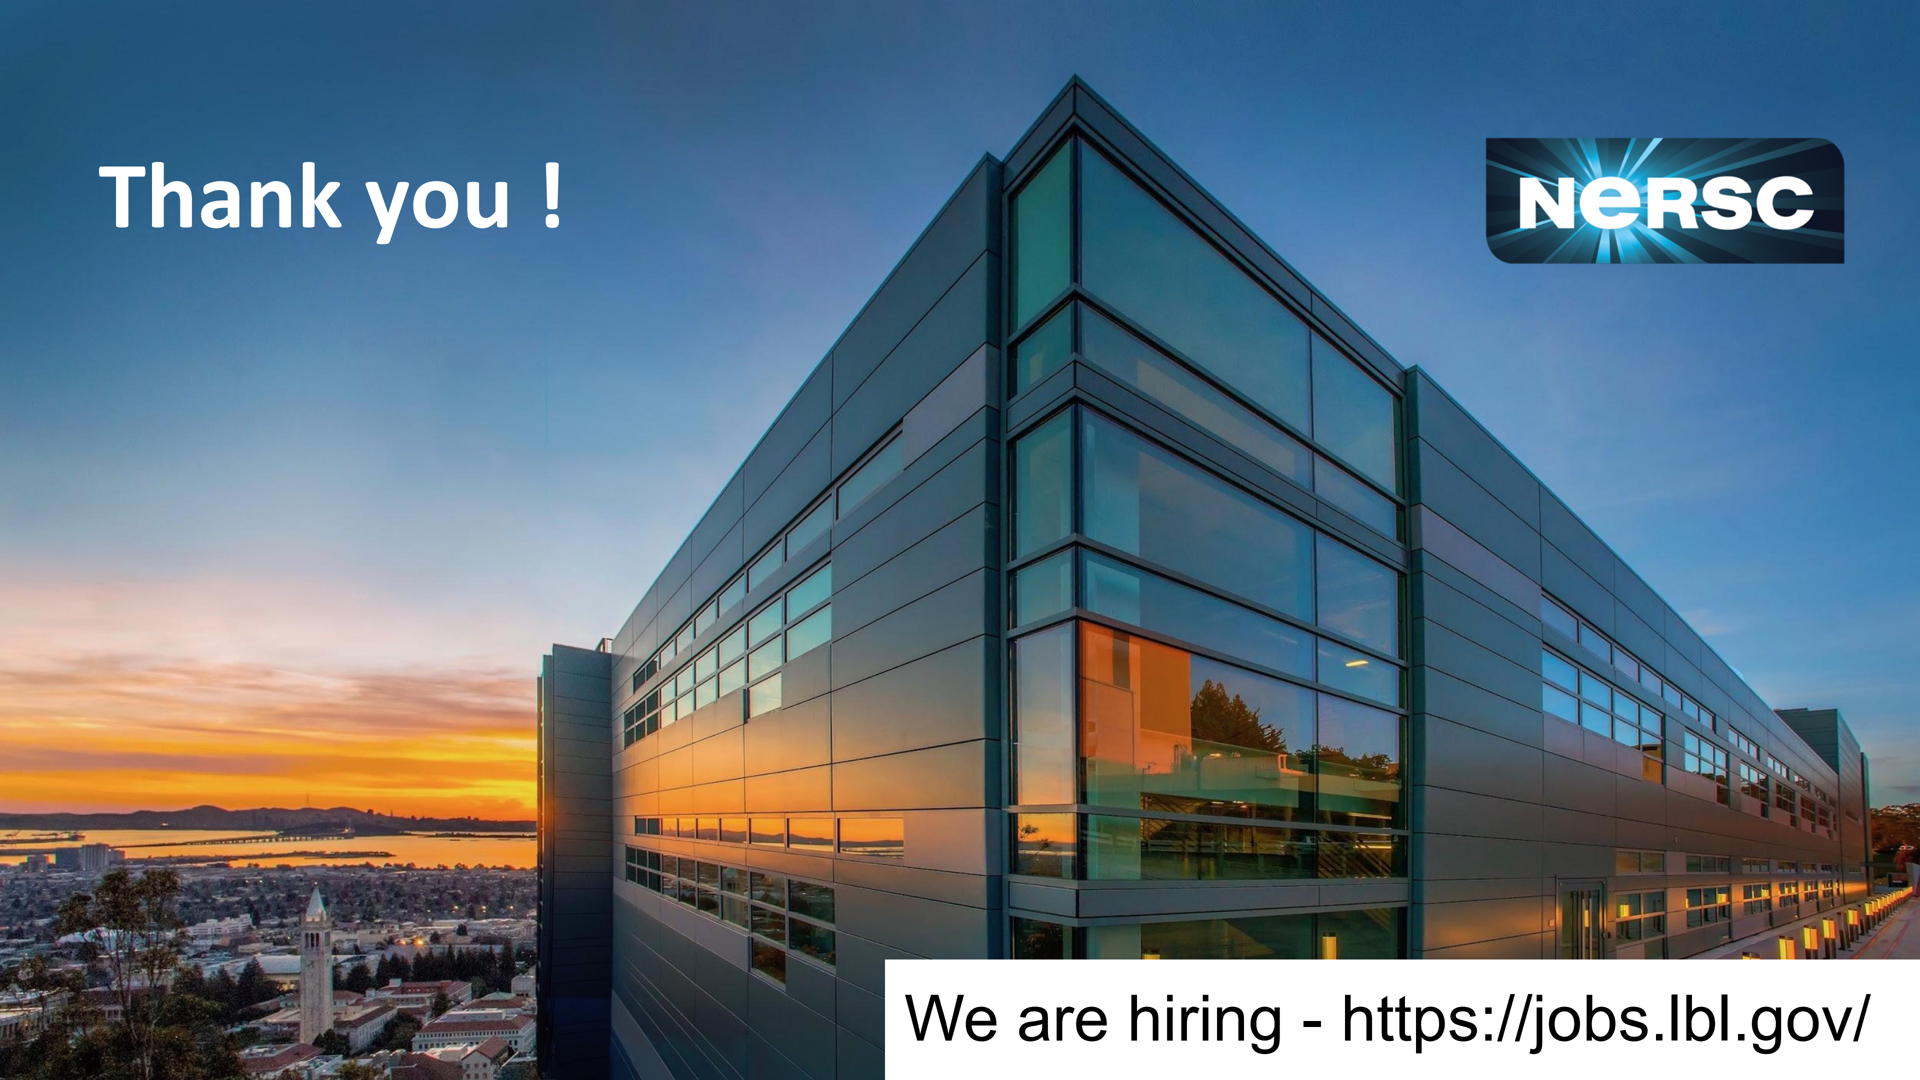

We are hiring - https://jobs.lbl.gov/