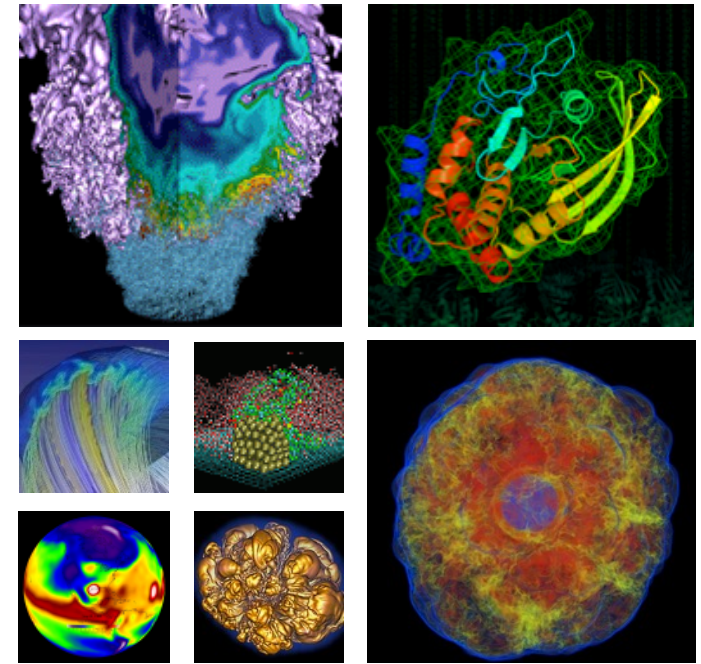


Challenges facing emerging workloads in the age of tiered-everything



Glenn K. Lockwood, Debbie Bard, Quincey Koziol

May 1, 2019

NERSC: the mission HPC facility

for the U.S. Department of Energy Office of Science

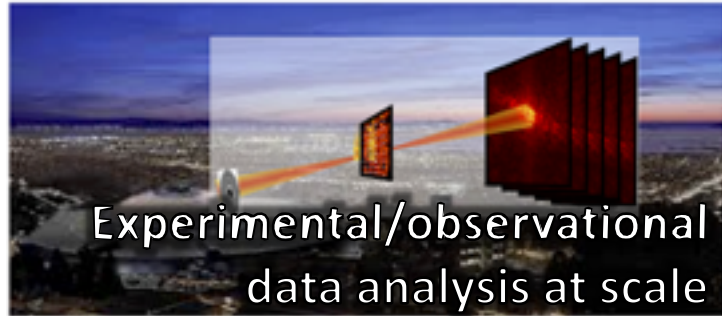
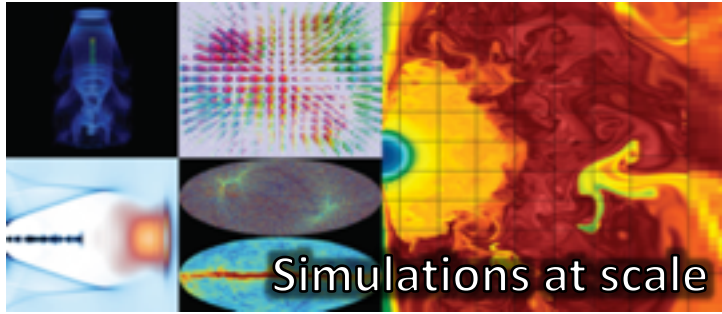
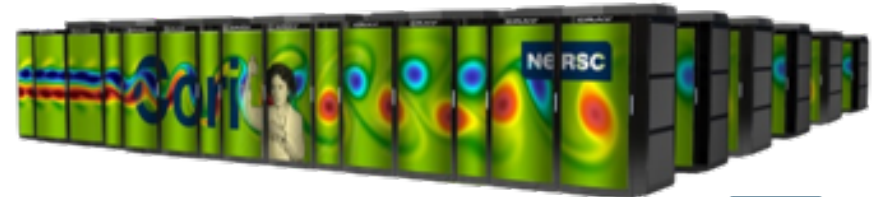
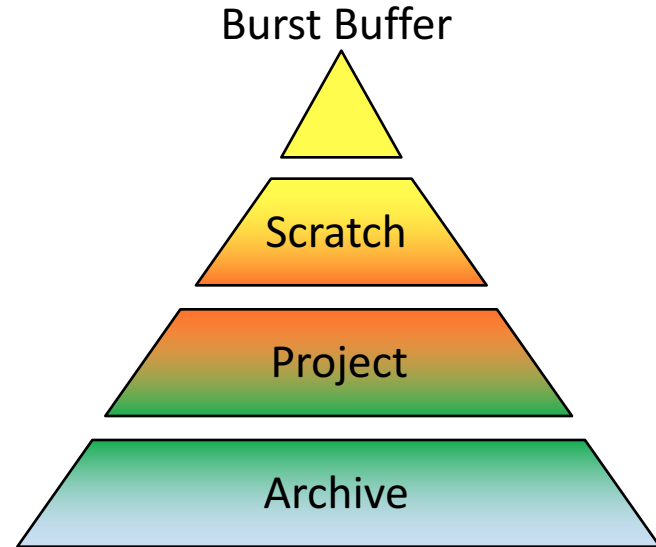


Photo credit: CAMERA

7,000 users, 800 projects, 700 applications
2,000 publications per year



Pains of coupling experiment and HPC



Research Challenges

- **Workflows require manual intervention; effort scales with architectural complexity**
- **Real-time access to HPC requires coscheduling compute, network, and tiers of storage**
- **HPC generally lacks support for searching, publishing, and sharing data**
- **Analytics software does not scale within or across the hierarchy**

Policy Challenges

- **Resilience strategy needed for time-critical experimental analysis**
 - Coordinating maintenances, fault-tolerant pipelines, rolling upgrades, alternative computing/data facilities...
- **Federated identity between wide-area data sources and sinks not yet realized**

Adapted from D. J. Bard et al., *NERSC Superfacility Project*.

Pains of AI at scale



- **Applications are very young, unstable, and rapidly changing**
 - Data management typically last consideration for developers
 - Complex storage hierarchy at odds with rapid innovation
- **AI I/O workloads are extremely demanding compared to traditional HPC**
 - Incumbent parallel file systems are grossly insufficient
 - Burst buffers and node-local NVMe are critical at scale

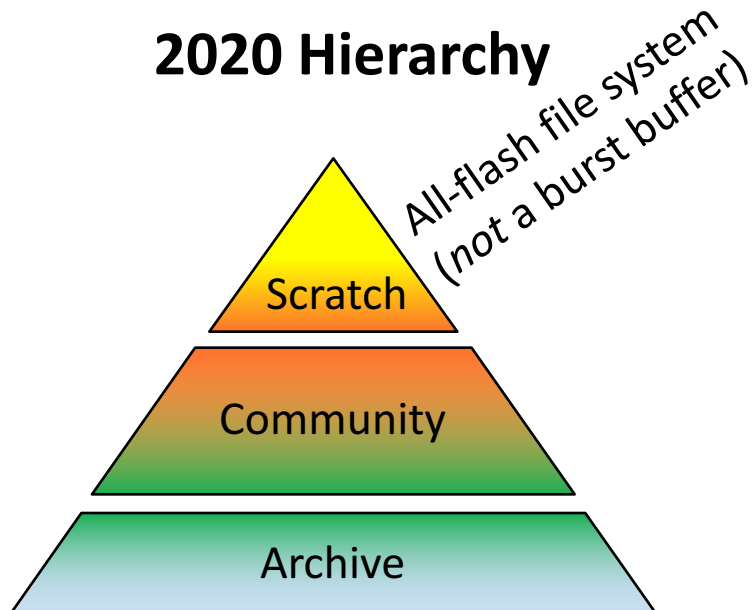
HPC Simulation	Deep Learning
Write-Once, Read-Never	Write-Once Read- <u>Always</u>
Contiguous, large I/O to sequence of files	Random, small I/O to random files (mini-batching)
O(10) of TBs in O(1000) files	O(10) of TBs in O(100,000) files

Adapted from Q. Koziol and Prabhat, *I/O for Deep Learning at Scale*.

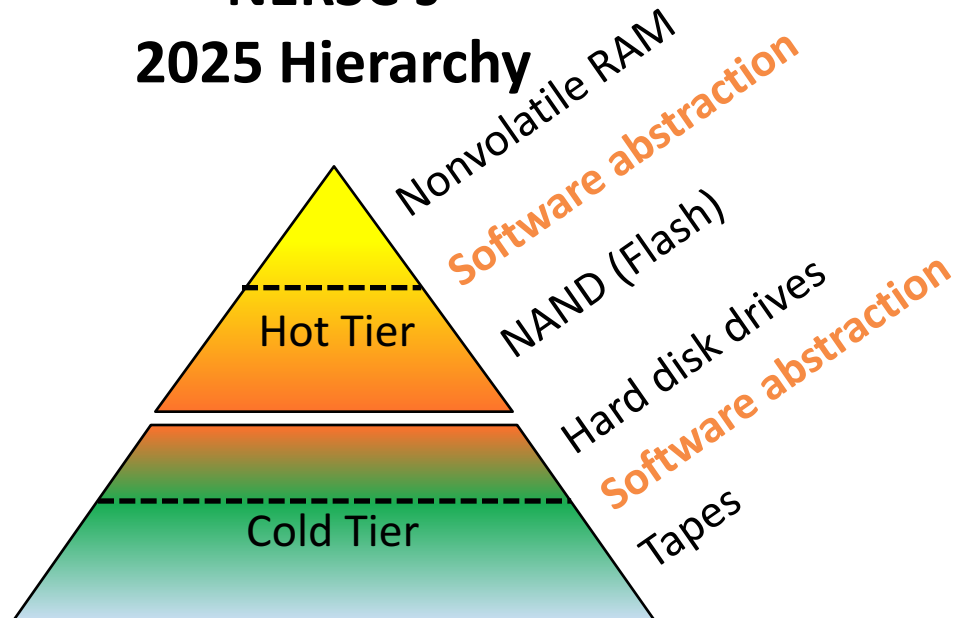
The future requires smarter storage systems



NERSC's 2020 Hierarchy



NERSC's 2025 Hierarchy



More info: G. K. Lockwood *et al.*, "Storage 2020: A Vision for the Future of HPC Storage," Berkeley, CA, 2017.



Thank you!

(and we're hiring!)