Drivers for Storage at ALCF

Storage Systems and Input/Output (SSIO) Workshop

Kevin Harms
ALCF

September 19, 2018

Argonne NATIONAL LABORATORY

Computing Resources

Mira IBM BG/Q
49,152 nodes
786,432 cores
768 TiB RAM
Peak flop rate: 10 PF

Theta Cray XC40
4,392 nodes
281,088 cores
892 TiB RAM
Peak flop rate: 11.69 PF

Cetus IBM BG/Q
4,096 nodes
65,536 cores
64 TiB RAM
Peak flop rate: 836 TF

Cooley Cray/NVIDIA
126 nodes
1512 Intel Haswell CPU cores
126 NVIDIA Tesla K80 GPUs
48 TB RAM / 3 TB GPU

Iota Intel/Cray XC40
44 nodes
2,816 cores
8.9 TiB RAM
Peak flop rate: 117 TF

Storage Capability

Disk

- Mira: ~27 PB of GPFS file system capacity with performance of 240 GB/s on the largest file system (19PB).
- Theta: ~18 PB of GPFS/Lustre file system capacity; 9PB is GPFS and 9.2PB is Lustre.

Tape

- The ALCF has three 10,000-slot libraries using LTO 6 tape technology. The LTO tape drives have built-in hardware compression for an effective capacity of 36-60 PB.

Theta – ALCF's newest production system

Features Intel processors and interconnect technology, a new memory architecture, and a Lustre-based parallel filesystem – all integrated by Cray's HPC software stack

Storage Resources

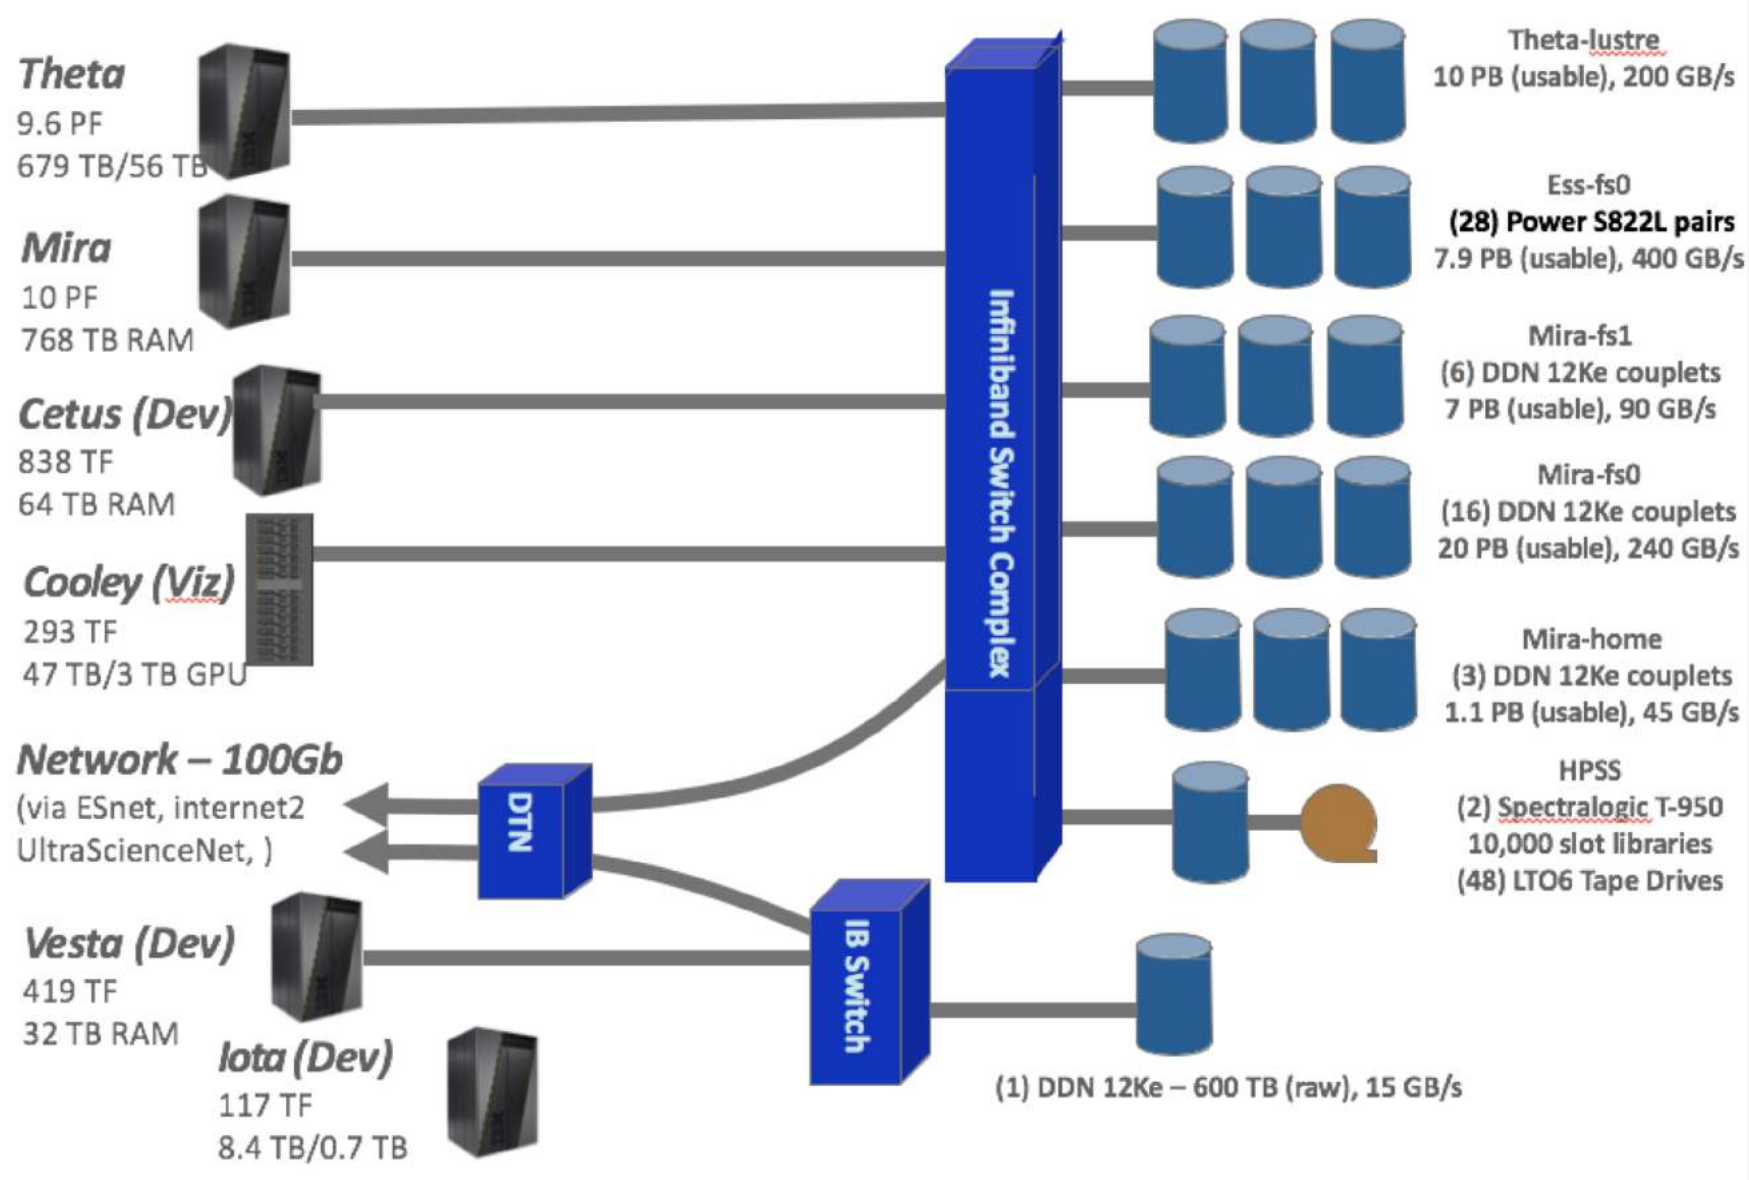

The Future

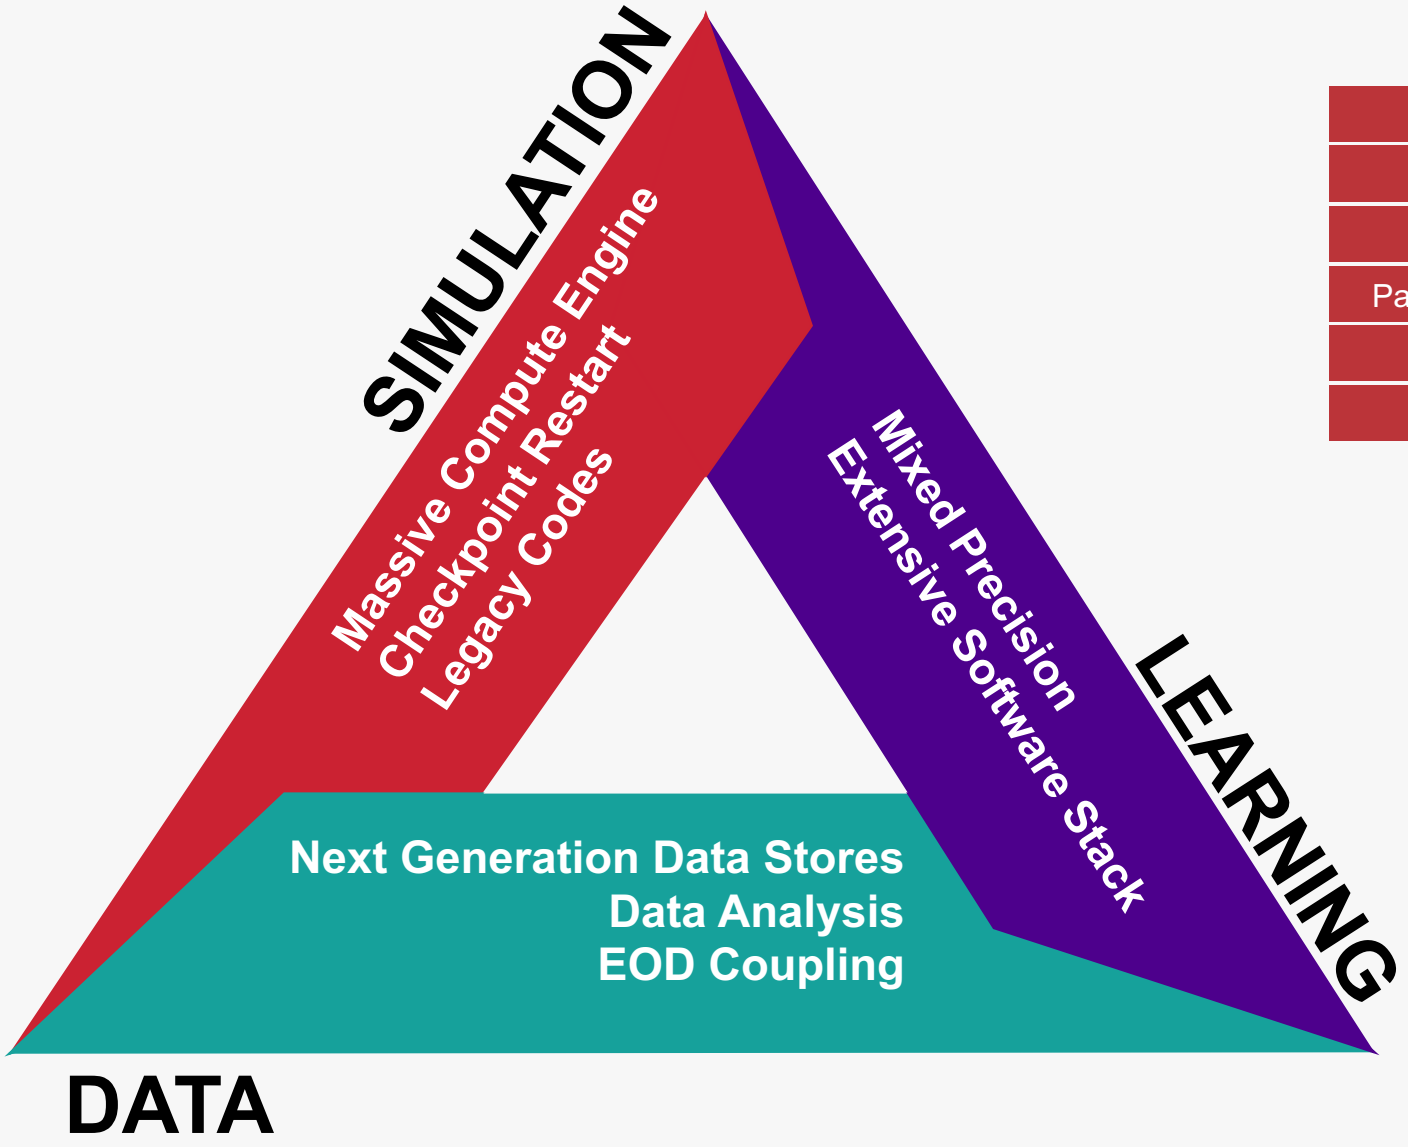

Software Stack

Simulation	Data	Learning
Languages	Languages	Languages
Directives	Big Data Stack	DL Frameworks
Parallel Runtimes	Statistical Libs	Statistical Libs
Solver Libs	Databases	Linear Algebra Libs
I/O Libs	I/O Libs	I/O Libs
	Viz Apps	

Policy Advances

S
D
L

Scheduling
Data Movement/Storage
Security/ID
Award Programs

ALCF Data Science Program

- "Big Data" science that require the scale and performance of leadership computing
- Projects cover a wide variety of application domains that span computational, experimental and observational sciences

https://www.alcf.anl.gov/alcf-data-science-program

- Focus on data science techniques including but not limited to statistics, machine learning, deep learning, Uncertainty Quantification, image processing, graph analytics, complex and interactive workflows

Realistic Simulations of the LSST Survey at Scale

- **PI: Katrin Heitmann, Argonne National Laboratory. 70M core hours**
- **Objectives: Development & execution of end-to-end workflow starting from simulation to the creation of sky maps with realistic galaxies.**
- **Impact: Deliver largest & most detailed synthetic sky maps ever created ready for the first data from LSST.**
- **Approach: Create a virtual survey with images almost indistinguishable from real LSST observations, develop an end-to-end pipeline for LSST data processing and analysis on ALCF supercomputers**

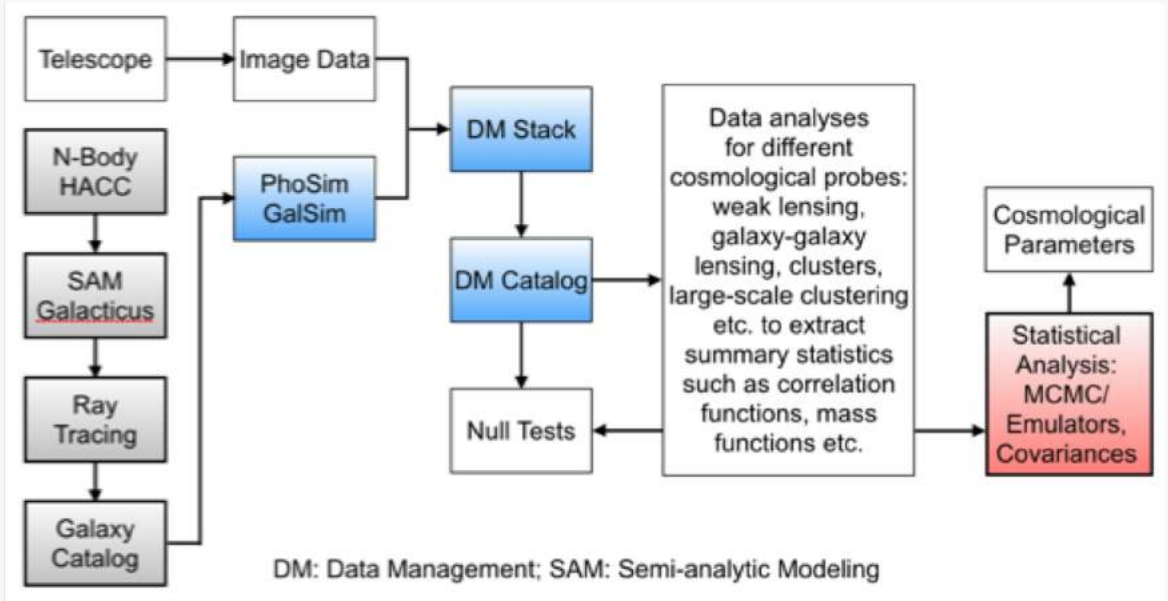

New Services

- Jupyter Lab
  - Mounts main GPFS and Lustre file systems
  - Can launch jobs / interface HPC resources
  - https://jupyter.alcf.anl.gov

- Petrel
  - Data sharing for scientists
  - https://petrel.alcf.anl.gov

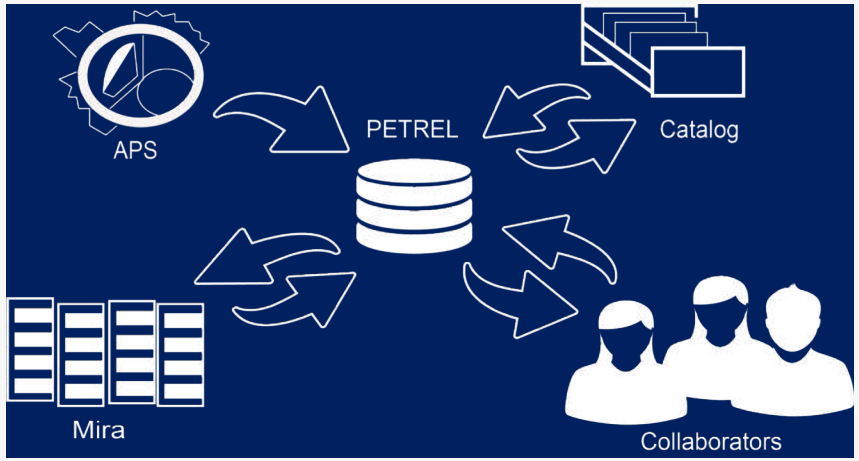

Petrel

- Implementation of the modern research data portal design
- Data management and sharing portal
- Exists outside internal ALCF security domain
- More flexible security model allowing sharing of data outside traditional project boundaries
- Uses Globus with REST APIs to enable transfers

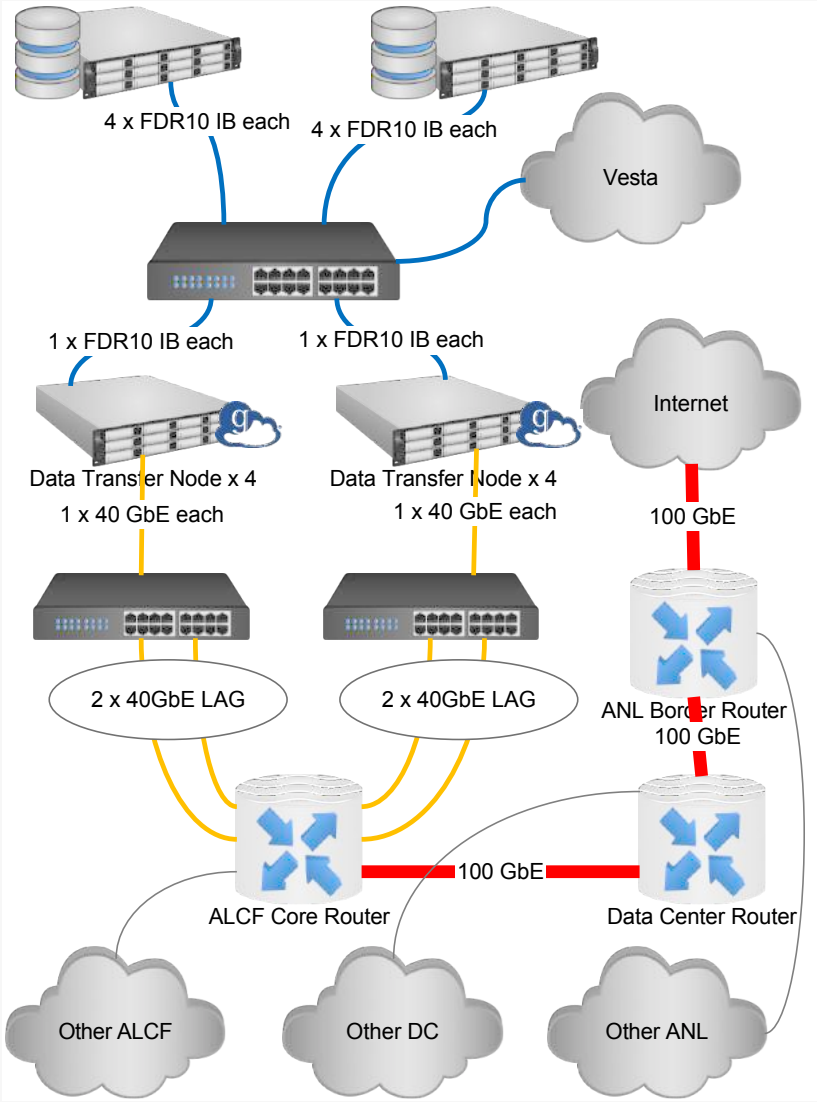

Courtesy Bill Allcock / Ben Allen

New Technologies - Balsam

Balsam is a workflow manager that simplifies the task of running large-scale job campaigns on ALCF resources while minimizing user involvement and improving productivity.

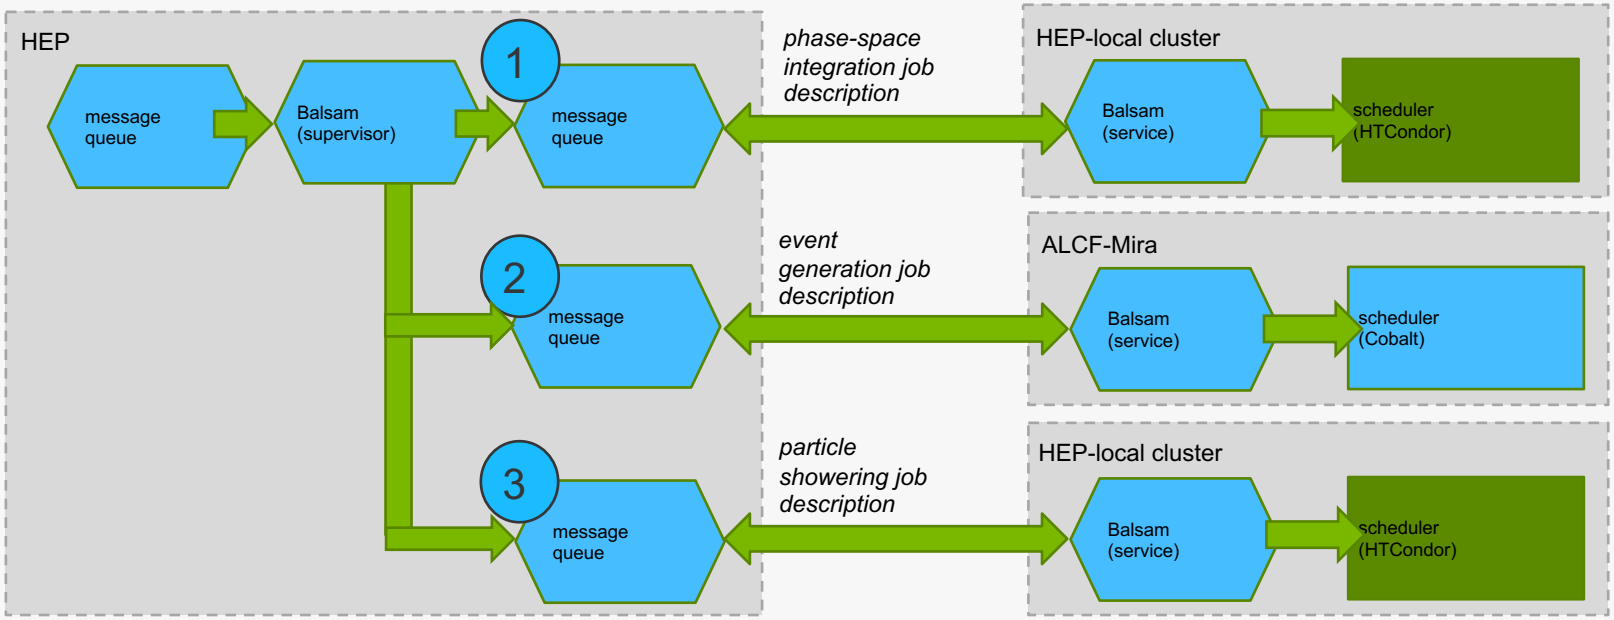

Schematic of ATLAS deployment of Balsam on multiple resources to execute a workflow with alternating serial and parallel stages

https://www.alcf.anl.gov/balsam

New Technologies – Deep Hyper

- Hyperparameter optimization is of paramount importance for deep learning for science. This is expected to be a **key workload on exascale systems**
- Model-based search iteratively refines the model in promising input region by obtaining new outputs at unevaluated input configurations
- General framework:
  - Initialization phase using Random or Latin hypercube sampling
  - Iterative phase wherein a model is fit and then sample using this model

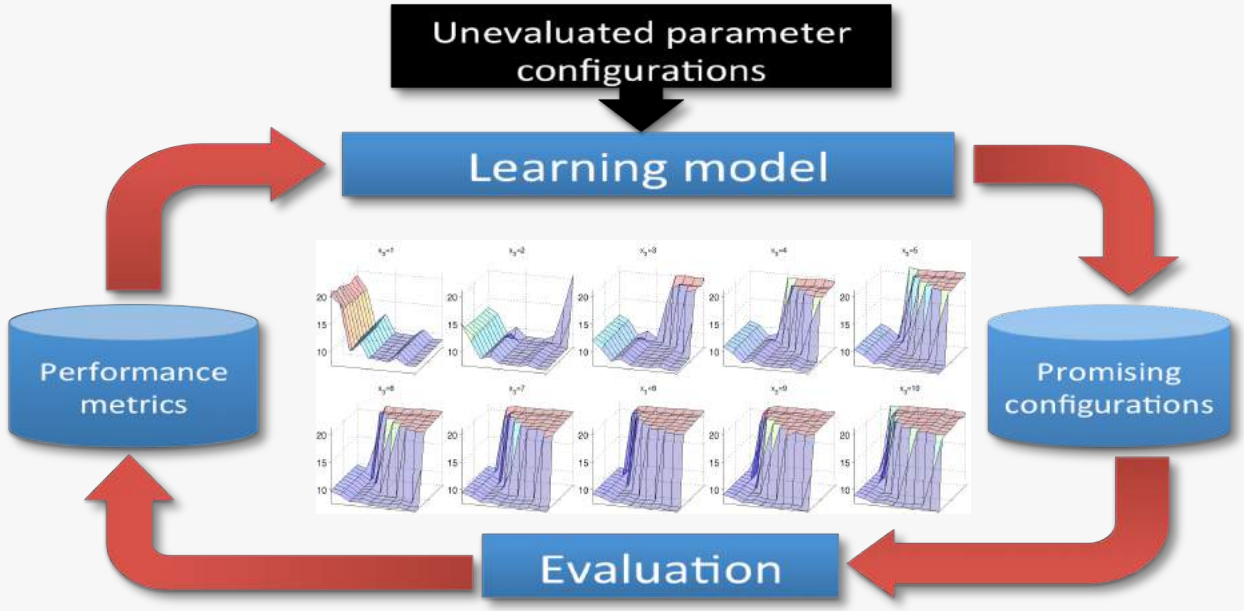

Example Surrogate Model Fitted to Sampled Performance
(iterative refinement improves the learning model)

Object Storage

- Promise of...
  - increased concurrency
  - flexible I/O models
  - designed for new storage hardware

Ceph

- ALCF/JLSE experiment with Ceph for underlying storage stack for Petrel v3 upgrade

DAOS

- ALCF/MCS collaboration with Intel on various underlying technologies within DAOS

Future Storage

- Moving toward global file system
  - Capacity storage for extended project data life time
  - Strong support for legacy interfaces
  - Out-of-band from system acquisitions
- Dedicated storage with systems
  - Performance focused
  - Adequate capacity for month(s)

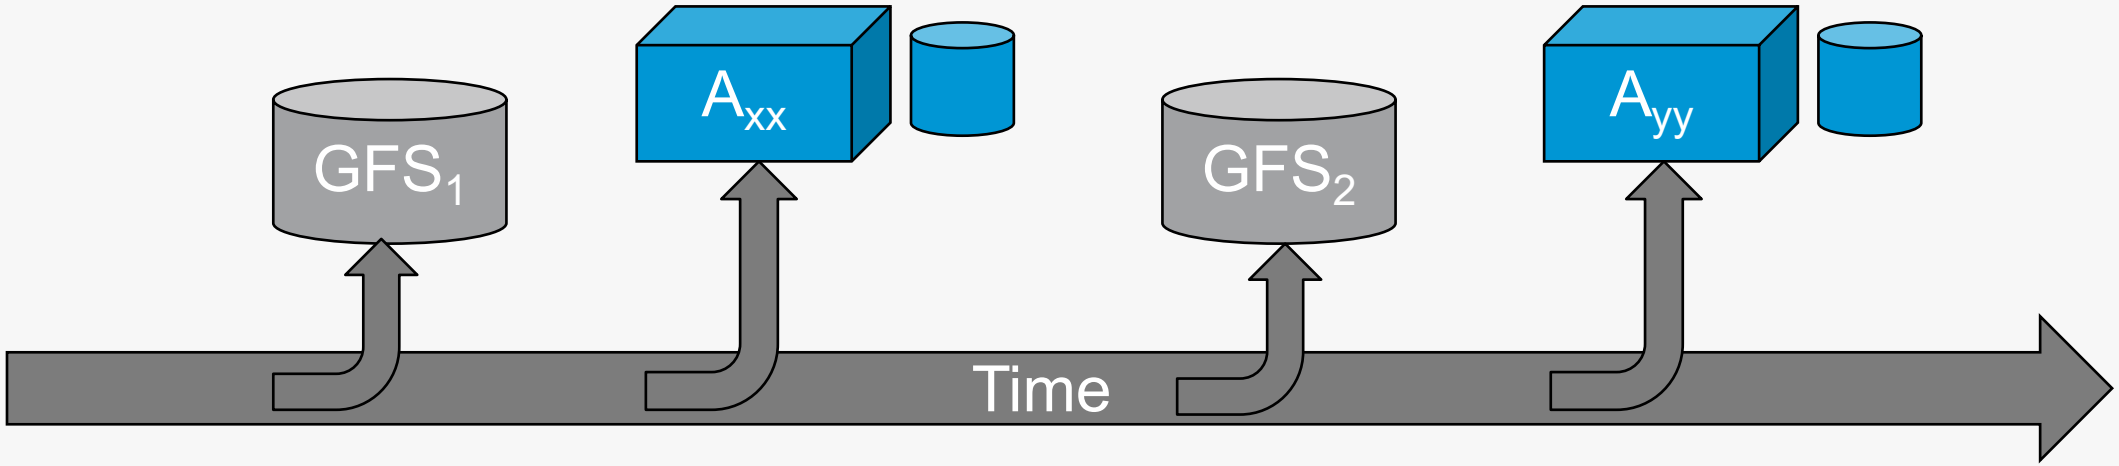

Aurora

Hardware and software optimized for Simulation, Data, and Learning

COMPUTE

- FLOPS *(>1exaFLOPS)*
- Concurrency
- Memory performance
- ML/DL operations

I/O

- Speed
- Capacity
- Flexibility
  - Conventional I/O
  - Database
  - Analytics middleware

Programming Environment

- Optimizing compilers
- Latest OpenMP
- Key Big Data stack components
- Productivity languages
- ML/DL frameworks
- Optimized libraries
  - Math
  - Statistics
  - ML/NN

ALCF Early Science Program (ESP)

Applications Readiness

- Prepare applications for next-gen system:
  - Architecture
  - Scale
- ~Two year lead time

Proposals

- Ambitious targeted science calculation
- Parallel performance
- Development needed
- Team

http://esp.alcf.anl.gov

Support

PEOPLE

- Funded ALCF postdoc
- Catalyst staff member support
- Vendor experts

TRAINING

- Training on HW and programming
- Community workshop to share lessons learned

COMPUTE RESOURCES

- Current ALCF systems
- Early next-gen hardware & simulators
- 3 months dedicated Early Science access
  - Pre-production (post-acceptance)
  - Large time allocation
  - Continued access for rest of year

Extending Moore's Law computing with Quantum Monte Carlo

PI: Anouar Benali (Argonne)

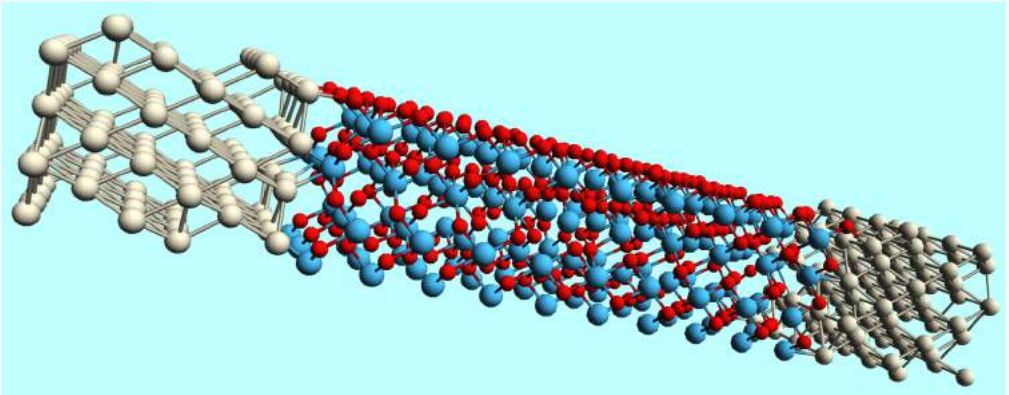

IMPACT:

– Research aims to advance our knowledge of the HfO2/Si interface necessary to extend Si-CMOS technology beyond Moore's law.

Relevance to *Aurora* Hardware/Software:

– The QMCPACK development project will focus on the management of walker associated memory objects on the multi-level memory/storage hierarchy, such as HBM, DDR, NVRAM, SSD and burst buffer.

Materials Science

Relevance to DOE/ANL/ALCF Mission:

– ANL material science and center for nano-materials

Software: QMCPACK

Other Benefits

– Developing the QMCPACK/Quantum Expresso interface.

Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience

PI: Nicola Ferrier (ANL); Collaborators: Princeton Univ.

IMPACT:

– Data and computational pipeline for neuroscience that will extract mappings of neurons and their connections from extremely large electron microscope experiments

Relevance to *Aurora* Early Science Program - Data and Learning (DL):

– Compute-intensive and data-parallel deep learning models, including flood fill networks (RNN) and 3D Unets, together with large and complex data and image processing needs will stress the DL hardware and software and communication fabric
– Data-intensive workflow could benefit from A21 I/O HW and near-real-time coupling of experiments is a novel use case.

Relevance to DOE/ANL/ALCF Mission:

– Crucial to understanding real-time coupling of supercomputing systems with experiments

Other Benefits:

– Data pipeline will benefit APS experiments

Biosciences, Neuroscience

Software: Tomosaic, Tensorflow, Horovod, Flood fill networks, 3D Unets

Simulating and Learning in the ATLAS detector at the Exascale

PI: James Proudfoot (ANL), Other Institutions: LBNL

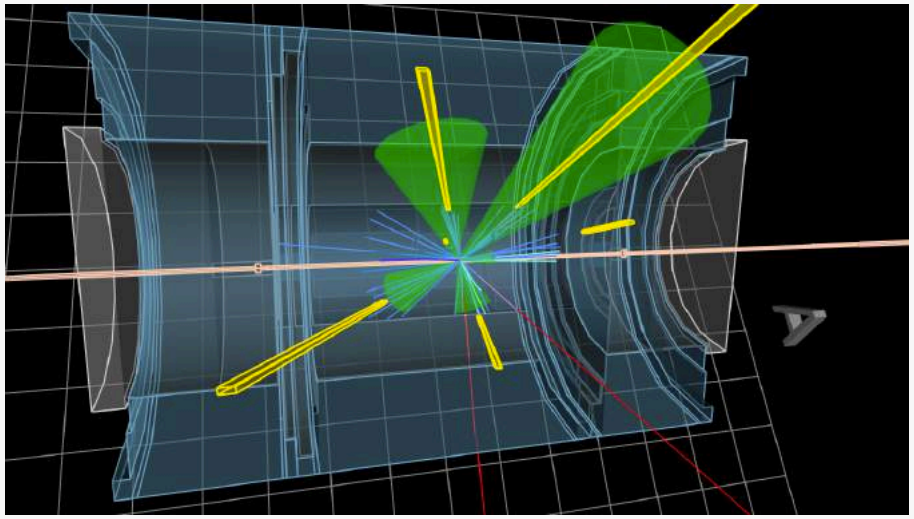

IMPACT:

– Meet the computing and analysis requirements of LHC science enabling the discovery of new particle physics

Relevance to *Aurora* Early Science Program - Data and Learning (DL):

– Diverse end-to-end workflows requirements
– Diverse compute and analysis kernels from relevant ATLAS software stack will stress the DL software and hardware characteristics.
– Machine learning for reconstruction would likely leverage graph convolutions - A novel use-case. ML needs would stress the system characteristics

Relevance to DOE/ANL/ALCF Mission:

– Prepare Aurora to meet the computing needs of HEP experiments
– Crucial to understanding forthcoming experimental data

Other Benefits:

– Deep ALCF staff expertise with ATLAS and good portability exercise

High Energy Physics, ATLAS

Software: AthenaMT, Root, workflows, Containers, Tensorflow

Storage Stats

- 90% of the capacity of the system is located in a few percentage of the files
- Chart shows ~93% of metadata objects, another ~7% are in directories and links

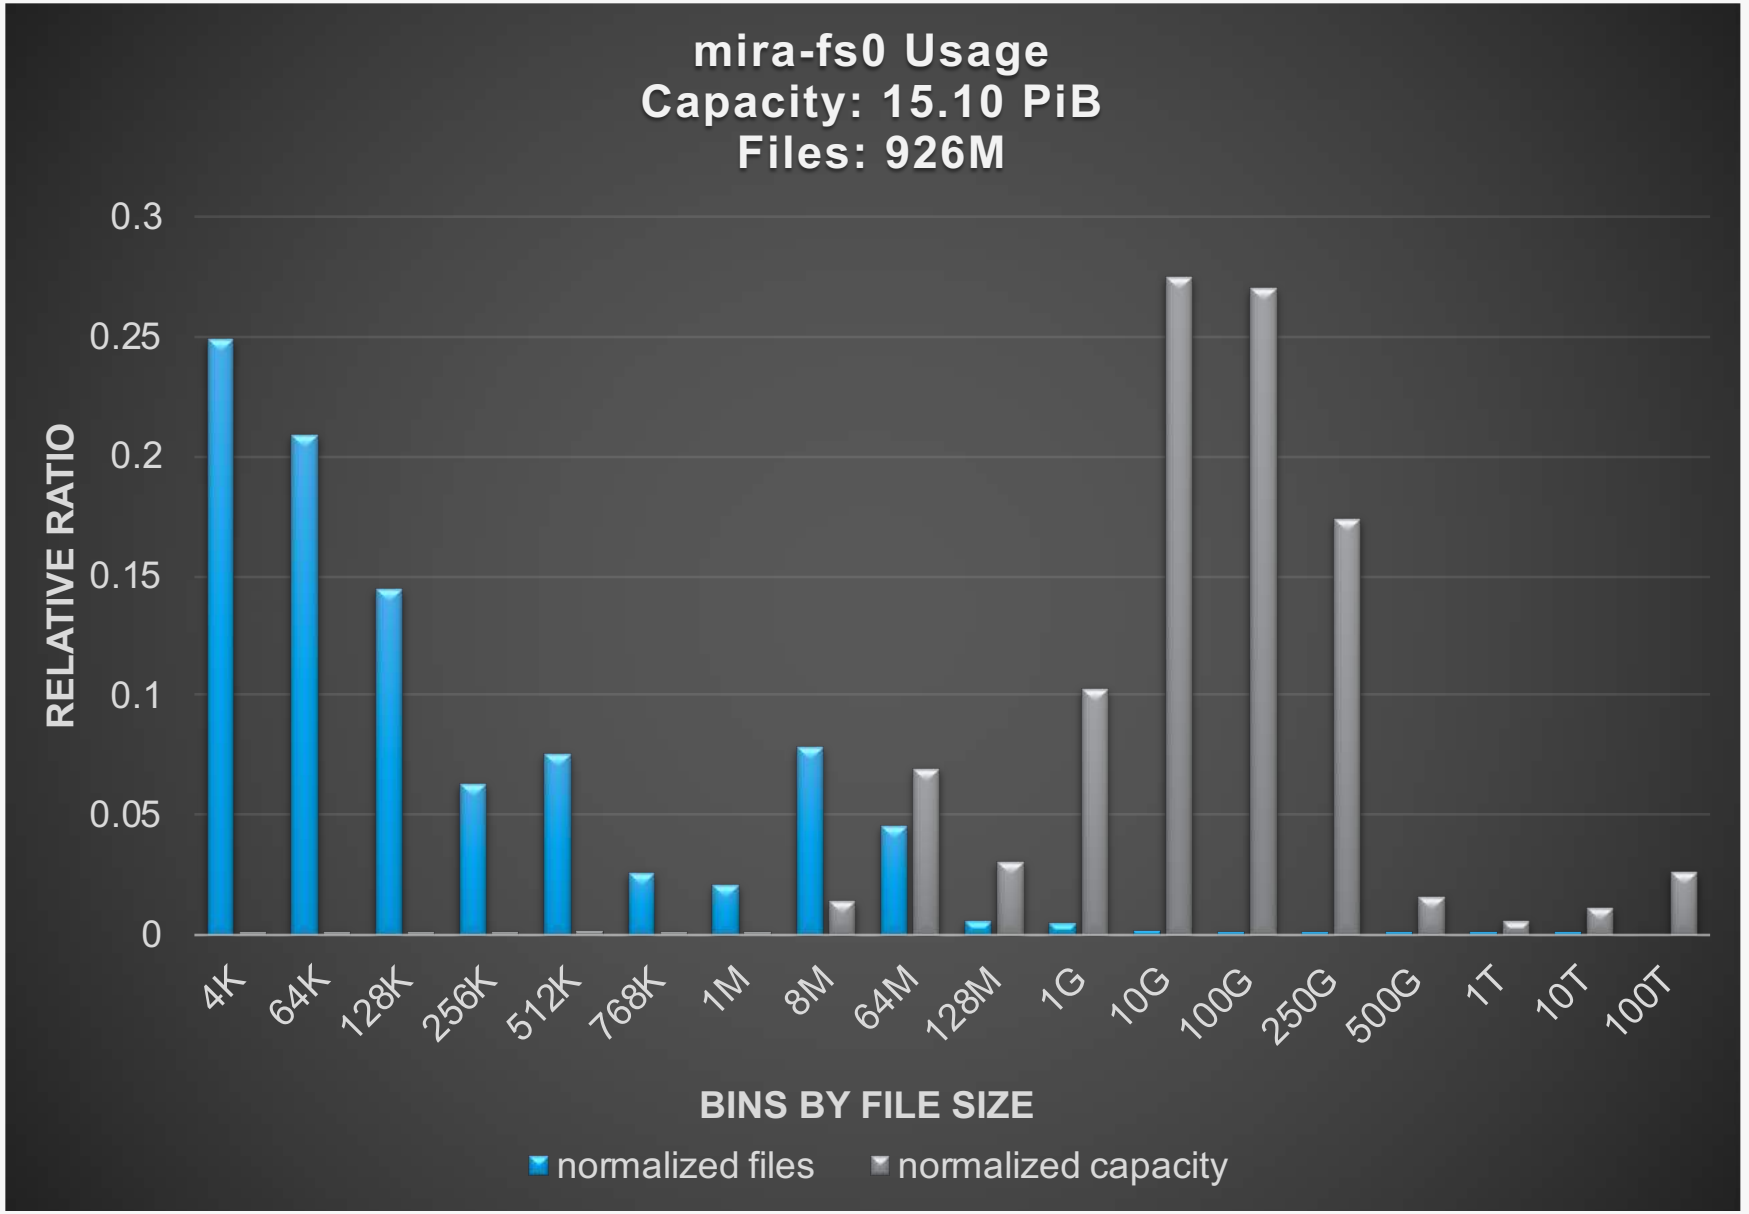

Argonne NATIONAL LABORATORY

Storage Stats - Tape

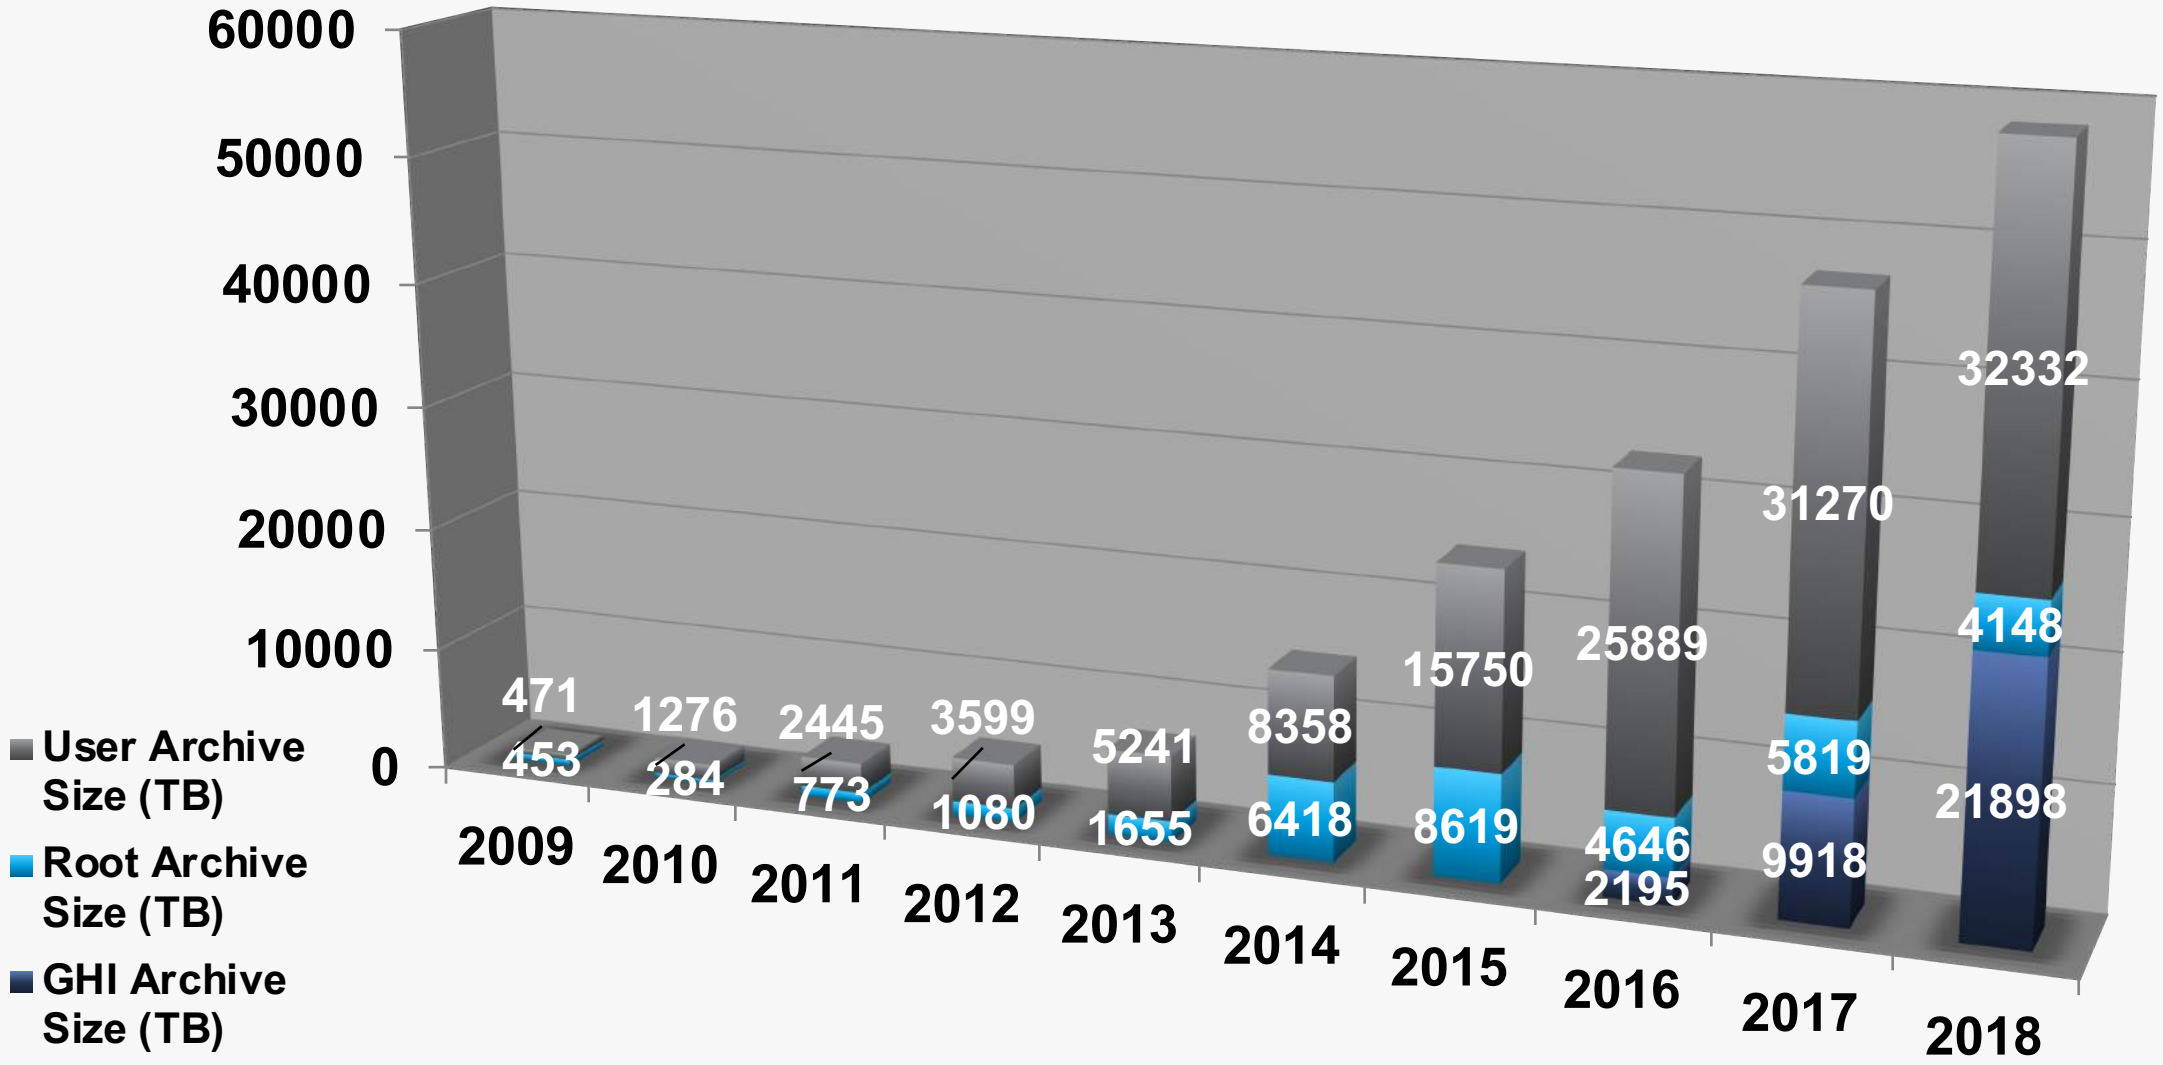

Closing

- Sheer scale of data being generated by simulation and experiment
  - How to annotate and track this data?
  - What is the facilities role?
- Archiving of data for long term (years) reference and reproducibility
  - Fast, easily accessed, tools to enable finding and interacting
  - Is storing data on tape forever a viable solution going forward?
- Longer term - expect science groups are going to be generating massive quantities of data. Even if the scientists are able to effectively reduce the *final* footprint, the intermediate footprint at facilities is going to be large.
- Does the facility/SSIO have a role in research, training and understanding what data really must be preserved

Acknowledgements

Katherine Riley
Elise Jennings
Venkat Viswanath
Tim Williams
Gordon McPheeters
Ben Allen
Jack O'Connell

Thank You!

This research used resources of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357.