Data-Intensive Science Requirements at Leadership Computing Facilities

Jack C. Wells
Director of Science
Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory

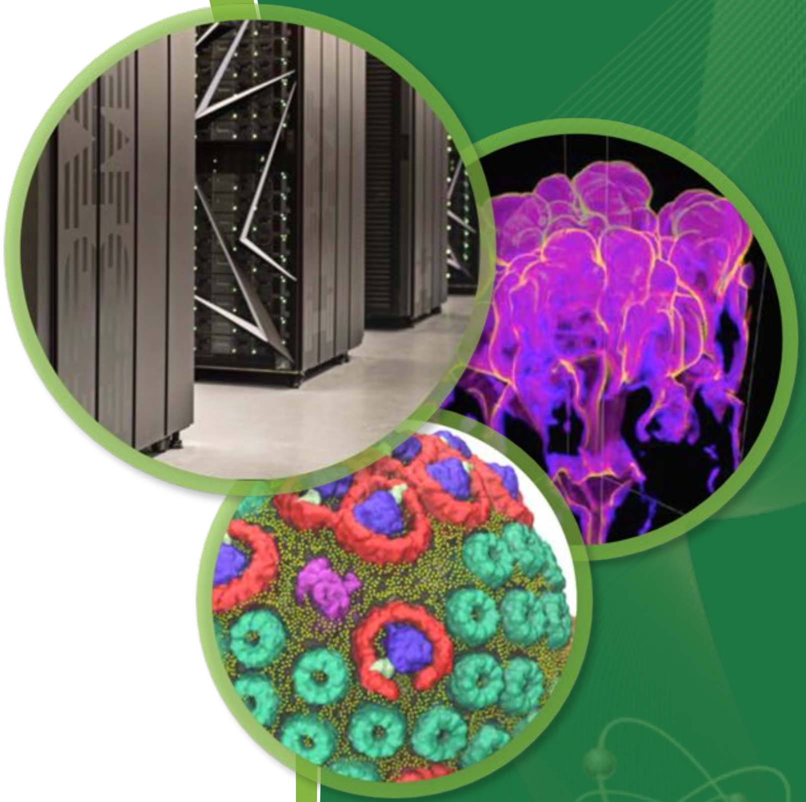

Storage Systems and Input/Output (SSIO) Workshop

19 September 2018

Gaithersburg, MD, USA

ORNL is managed by UT-Battelle
for the US Department of Energy

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

2013 OLCF User Requirements Survey – Memory Bandwidth is Priority Requirement

- 76% of users said there is still a moderate to large amount of parallelism to extract in their code, but…
- 85% of respondents rated the difficulty level of expressing that parallelism as moderate or difficult – underlining the need for more support, training, and assistance in this area
- ***For the first time, FLOPS was NOT the #1 ranked requirement***
- Local memory capacity was not a driver for most users, perhaps in recognition of cost trends
- Data sharing, long-term archival storage of 100s of PB, data analytics, and faster WAN connections also showed up as important requirements

https://www.olcf.ornl.gov/wp-content/uploads/2013/01/OLCF_Requirements_TM_2013_Final.pdf

Hardware feature	Ranking
Memory Bandwidth	**4.4**
Flops per Node	4.0
Interconnect Bandwidth	**3.9**
Archival Storage Capacity	**3.8**
Interconnect Latency	3.7
Disk Bandwidth	**3.7**
WAN Network Bandwidth	3.7
Memory Latency	**3.5**
Local Storage Capacity	**3.5**
Memory Capacity	**3.2**
Mean Time to Interrupt	3.0
Disk Latency	2.9

Rankings from OLCF users
1=not important, 5=very important

What is CORAL?
The program through which ORNL, ANL, and LLNL are procuring supercomputers.

- Several DOE labs have strong supercomputing programs and facilities.
- DOE created CORAL (the Collaboration of Oak Ridge, Argonne, and Livermore) to jointly procure these systems, and in so doing, align strategy and resources across the DOE enterprise.
- Collaboration grouping of DOE labs was done based on common acquisition timings. Collaboration is a win-win for all parties.

U.S. DEPARTMENT OF ENERGY

Oak Ridge National Laboratory

Lawrence Livermore National Laboratory

Argonne NATIONAL LABORATORY

Paving The Road to Exascale Performance

ORNL Summit System Overview

System Performance

- Peak of 200 Petaflops (FP_{64}) for modeling & simulation
- Peak of 3.3 ExaOps (FP_{16}) for data analytics and artificial intelligence

The system includes

- 4,608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM file system transferring data at 2.5 TB/s

Each node has

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory
  (96 GB HBM2 + 512 GB DDR4)
- 1.6 TB of NV memory

Oak Ridge National Laboratory | Leadership Computing Facility

Supercomputer Specialization vs ORNL Summit

- As supercomputers got larger and larger, many expected them to be more specialized and limited to a reduced number of applications that can exploit their growing scale.
- We also predicted that power consumption would be a dominant issue as we approached exascale.
- Summit's architecture seems to have stumbled into a sweet spot. It is not showing power consumption growth and has broad capability across:
  - Traditional HPC modeling and simulation
  - High performance data analytics
  - Artificial Intelligence

Slide ACK: Al Geist (ORNL)

Summit Excels Across Simulation, Analytics, and AI

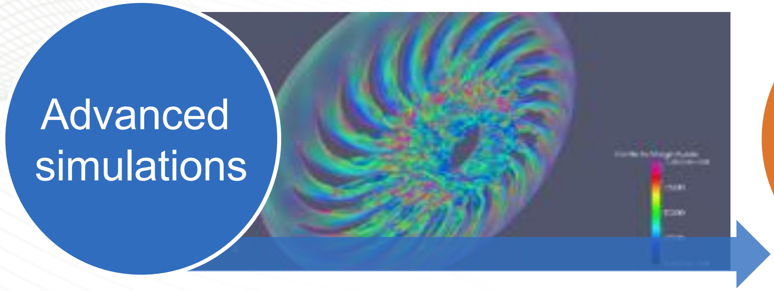

- Data analytics – CoMet bioinformatics application for comparative genomics. Used to find sets of genes that are related to a trait or disease in a population. Exploits cuBLAS and Volta tensor cores to solve this problem 5 orders of magnitude faster than previous state-of-art code.
  - **Has achieved 2.36 ExaOps** mixed precision (FP_{16}-FP_{32}) on Summit
- Deep Learning – global climate simulations use a half-precision version of the DeepLabv3+ neural network to learn to detecting extreme weather patterns in the global climate output
  - **Has achieved a sustained throughput of 1.0 ExaOps (FP_{16})** on Summit
- Nonlinear dynamic low-order unstructured finite-element solver accelerated using mixed precision (FP_{16} thru FP_{64}) and AI generated preconditioner. Answer in FP_{64}
  - **Has achieved 25.3x speedup** on Japan earthquake – city structures simulation
- **Many Early Science codes are reporting >10x speedup on Summit vs Titan**

Summit compared to Titan

- Many fewer nodes
- Much more powerful nodes
- Much more memory per node and total system memory
- Faster interconnect
- Much higher bandwidth between CPUs and GPUs
- Much larger and faster file system

Feature	Titan	Summit
Peak Flops	27 PF	200 PF
Application Performance	Baseline	5-10x Titan
Number of Nodes	18,688	~4,600
Node performance	1.4 TF	> 40 TF
Memory per Node	32 GB DDR3 + 6 GB GDDR5	512 GB DDR4 + 96 GB HBM
NV memory per Node	0	1600 GB
Total System Memory	710 TB (600 TB DDR3 + 110 TB GDDR5))	10 PB (2.3 PB DDR4 + 0.4 PB HBM + 7.4 PB NVRAM)
System Interconnect (node injection bandwidth)	Gemini (6.4 GB/s)	Dual Rail EDR-IB (23 GB/s)
Interconnect Topology	3D Torus	Non-blocking Fat Tree
Processors per node	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	32 PB, 1 TB/s, Lustre®	250 PB, 2.5 TB/s, GPFS™
Peak power consumption	9 MW	13 MW

Summit Node Schematic: Complex Memory Architecture

- Coherent memory across entire node
- NVLink v2 fully interconnects three GPUs and one CPU on each side of node
- PCIe Gen 4 connects NVM and NIC
- Single shared NIC with dual EDR ports

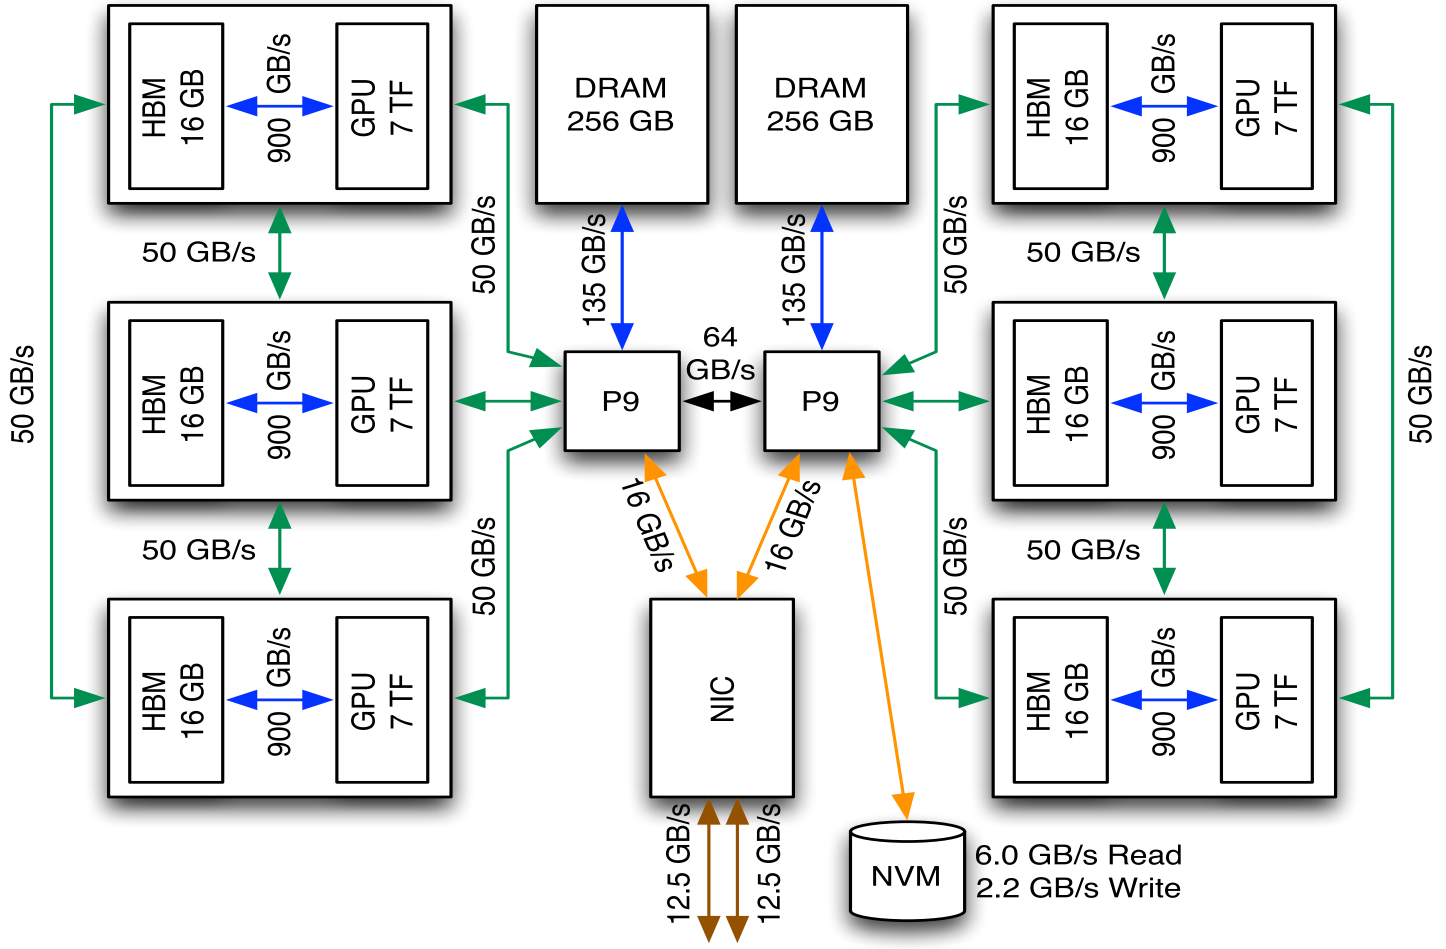

TF	42 TF (6x7 TF)
HBM	96 GB (6x16 GB)
DRAM	512 GB (2x16x16 GB)
NET	25 GB/s (2x12.5 GB/s)
MMsg/s	83

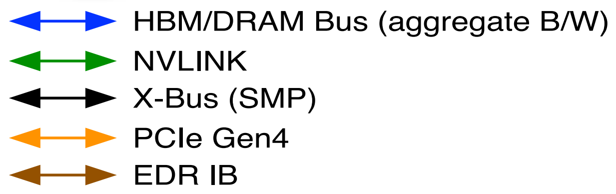

HBM & DRAM speeds are aggregate (Read+Write).
All other speeds (X-Bus, NVLink, PCIe, IB) are bi-directional.

Experimental and Observational Science Data is Exploding

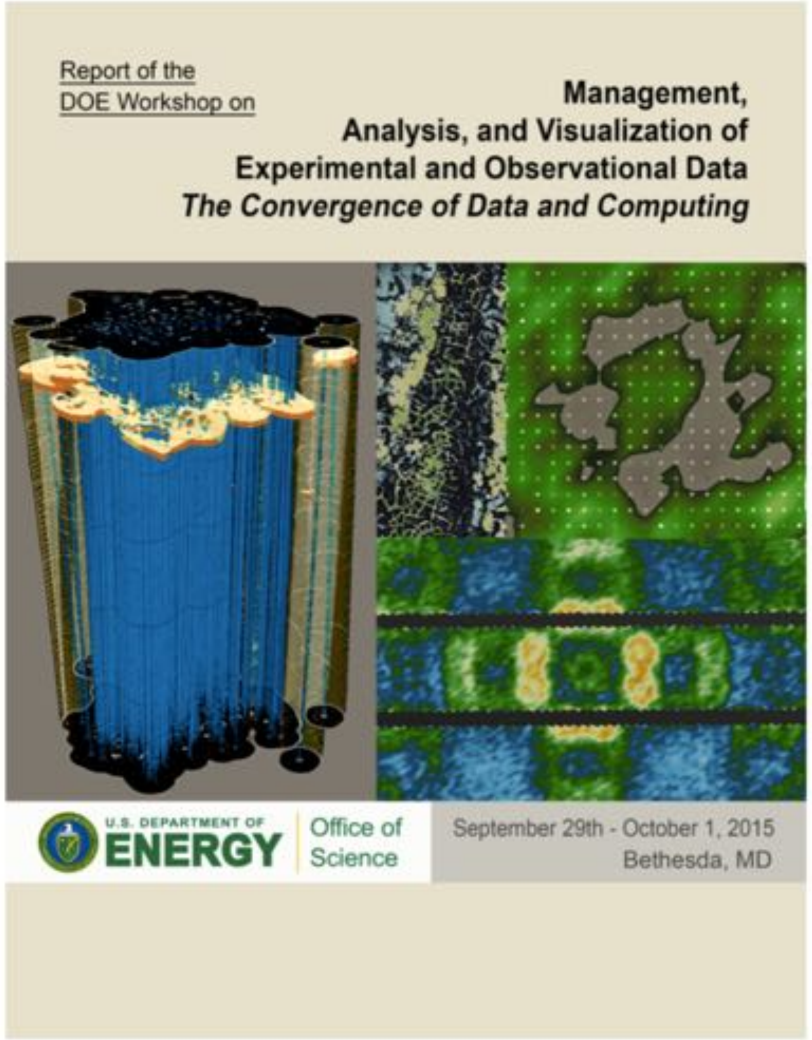

DOE Workshop: Data Management, Analysis and Visualization for Experimental and Observational Data (EOD) Workshop, (2015)

Program Leader: Lucy Nowell, DOE/ASCR

https://extremescaleresearch.labworks.org/events/data-management-visualization-and-analysis-experimental-and-observational-data-eod-workshop

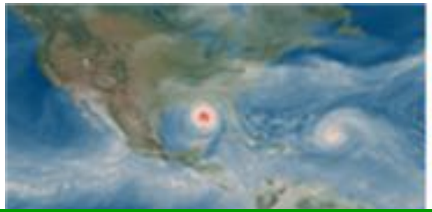
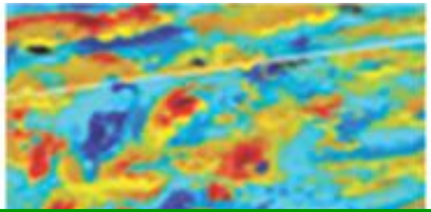
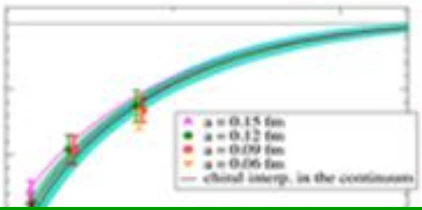
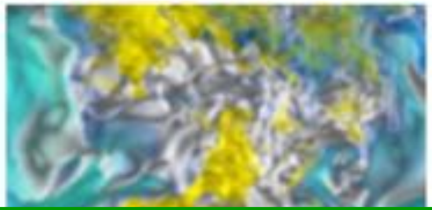


DOE-ASCR Exascale Requirements Reviews

ASCR facilities conducted six exascale requirements reviews in partnership with DOE Science Programs

- Goals included:
  - Identify mission science objectives that require advanced scientific computing, storage and networking in exascale timeframe
  - Determine future requirements for a computing ecosystem including data, software, libraries/tools, etc.

Schedule

June 10–12, 2015	**HEP**
November 3–5, 2015	**BES**
January 27–29, 2016	**FES**
March 29–31, 2016	**BER**
June 15–17, 2016	**NP**
Sept 27–29, 2016	**ASCR**
March 9–10, 2017	**XCut**

All 7 workshop reports are available online: http://exascaleage.org/

Data: Common Themes Across Offices

Across all the exascale requirements reviews a number of common themes emerged as particularly challenging areas. Let's have a discussion about defining the best categories.

Large-scale data storage and analysis

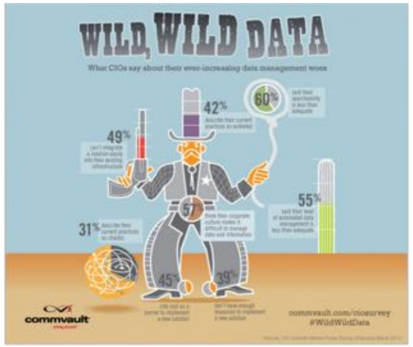

Data management, archiving, and curation

Input/Output (I/O) performance

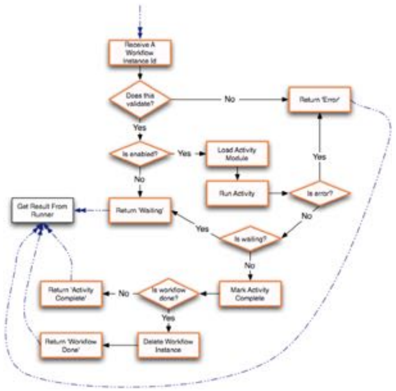

Experimental and simulation workflows

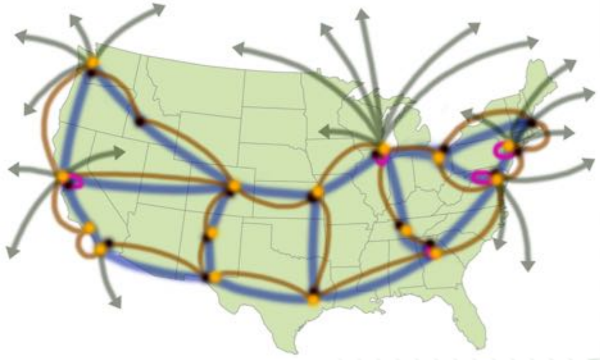

Remote access, sharing, and data transfer

Emerging Science Activities:
Selected Machine Learning Projects on Titan: 2016-2017

Program	PI	PI Employer	Project Name	Allocation (Titan core-hrs)
ALCC	Robert Patton	ORNL	Discovering Optimal Deep Learning and Neuromorphic Network Structures using Evolutionary Approaches on High Performance Computers	75,000,000
ALCC	Gabriel Perdue	FNAL	Large scale deep neural network optimization for neutrino physics	58,000,000
ALCC	Gregory Laskowski	GE	High-Fidelity Simulations of Gas Turbine Stages for Model Development using Machine Learning	30,000,000
ALCC	Efthimions Kaxiras	Harvard U.	High-Throughput Screening and Machine Learning for Predicting Catalyst Structure and Designing Effective Catalysts	17,500,000
ALCC	Georgia Tourassi	ORNL	CANDLE Treatment Strategy Challenge for Deep Learning Enabled Cancer Surveillance	10,000,000
DD	Abhinav Vishnu	PNNL	Machine Learning on Extreme Scale GPU systems	3,500,000
DD	J. Travis Johnston	ORNL	Surrogate Based Modeling for Deep Learning Hyper-parameter Optimization	3,500,000
DD	Robert Patton	ORNL	Scalable Deep Learning Systems for Exascale Data Analysis	6,500,000
DD	William M. Tang	PPPL	Big Data Machine Learning for Fusion Energy Applications	3,000,000
DD	Catherine Schuman	ORNL	Scalable Neuromorphic Simulators: High and Low Level	5,000,000
DD	Boram Yoon	LANL	Artificial Intelligence for Collider Physics	2,000,000
DD	Jean-Roch Vlimant	Caltech	HEP DeepLearning	2,000,000
DD	Arvind Ramanathan	ORNL	ECP Cancer Distributed Learning Environment	1,500,000
DD	John Cavazos	U. Delaware	Large-Scale Distributed and Deep Learning of Structured Graph Data for Real-Time Program Analysis	1,000,000
DD	Abhinav Vishnu	PNNL	Machine Learning on Extreme Scale GPU systems	1,000,000
DD	Gabriel Perdue	FNAL	MACHINE Learning for MINERvA	1,000,000
		TOTAL		**220,500,000**

Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program for 2019

- Access to the most capable, most productive, fastest open science supercomputers in the nation
- Call for proposals submission window:
  - Apr 16 – Jun 22, 2018
- Applicable to a broad array of science, engineering, and computer science domains
- Proposals must be:
  - High-impact, computationally and/or data intensive campaigns
  - Must take advantage of unique HPC architectures
  - Research that cannot be performed anywhere else.
- For more information visit http://www.doeleadershipcomputing.org/

U.S. DEPARTMENT OF ENERGY | Office of Science

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

Argonne NATIONAL LABORATORY

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Question 1: Trends or challenges coming from OLCF science that mandate research in storage and I/O (1)

- Strong demand for LCF resources from quality, meritorious R&D projects
  - Demand is broad across diverse areas of science and engineering, not narrow or focused
  - Tremendous benefit gained by provisioning general-use data and commute capabilities.
- NSCI theme of “convergence of compute-intensive and data-intensive systems” is rapidly advancing
  - Modeling and simulation remains the LCF’s largest market segment
  - Data analytics, machine learning/AI are rapidly growing LCF segments in our user programs
  - Many teams are integrating all “segments” through workload management systems
- Architectural complexity is increasing across the memory hierarchy
  - A capable storage system and scalable I/O are fundamental to each of these NSCI goals
  - LCFs have complex memory arrangements along with a complex storage layer.
  - Accelerated-node supercomputing – a lot of work, but mission outcomes are being achieved

Question 1: Trends or challenges coming from OLCF science that mandate research in storage and I/O (2)

- Strong demand for LCF resources from Experimental and Observational (EOD) R&D projects
  - Tremendous diversity in workloads and workflows
    - Broad application base & diverse requirements for time criticality
    - Need to use familiar EOD tools AND have the ability to “surge” capability or large ensembles into LCF data centers.
- Trends/Opportunity for the reuse of science reference data generated at LCFs.
- Long LCF procurement cycle “bumps up against” R&D’s “5-10 year horizon” – challenge for R&D – Facilities integration and planning.
  - For example, CORAL (I) proposal due date was February 18, 2014!
  - Summit is being deployed and accepted now (in 2018)
  - ECP Project can help with this.

Question 2: Key issues related to the OLCF's perspective on science requirements (1)

- Priority requirements for SSIO design remain: (1) Reliability, (2) Utility/ease-of-use, (3) Capacity, (4) Performance
- OLCF has a preference for a large, center-wide parallel file system
  - Significant science productivity from our center-wide Spider file system over Jaguar and Titan eras. Our preference is to continue with such a center-wide design concept.
- Since demand for mod/sim is strong, we do not anticipate a "phase change" in our SSIO requirements. Based on experience:
  - ~10% of our files use ~90% of the space
  - ~10% of the workloads account for 90% of the streaming bandwidth
  - ~10% of the workloads account for 90% of the metadata operations
- That being said, we see an increase in large streaming reads, and random small writes (less than 32KB)

Question 2: Key issues related to the OLCF's perspective on science requirements (2)

- Traditional mod/sim application segment will continue to play a major role in defining SSIO requirements
  - But data analytics and AI/ML are definitely hot topics right now
  - Most significant new requirement is the inclusion of large-streaming reads, smaller writes, and heavy metadata usage
- LCF storage system performance requires increased balance wrt writes and reads:
  - Current HPC filesystems are write-optimized, for machine learning/data analytics purposes we need a read-optimized file system
  - Today's HPC filesystems get good performance for reads on contiguous data, whereas ML needs random access of portions of data (tiny in relation to the whole, potentially very large)
  - ML/AI training sets do and will exceed memory capacity. Solutions will require, e.g., streaming from disk-based storage solutions, or accelerated with storage that can handle random access (e.g., flash).

Question 2: Key issues related to the OLCF's perspective on science requirements (3)

- Federated, distributed computing architectures are required to advance our mission
  - Many significant “proof-of-concept” demonstrations have been achieved, e.g., HPC integration into global-scale grid-computing campaigns, integration of EOD facility data workloads into HPC workflows, global-scale science-data-streaming demos
  - But significant effort is need to deliver on the promise of this vision
- Users need a coherent view of SSIO resources

Summary, Comments, & Conclusions

- Over the past 5 years, Titan has delivered the DOE Leadership Computing Mission at ORNL: delivering science and constraining energy consumption.
- ORNL is advancing hybrid-accelerated supercomputing based on success of Titan; Summit is our next step taking place this year.
- DOE Office of Science is advancing an integrated vision for exascale computing ecosystem, including data-intensive and distributed applications dependent on networks and storage, e.g., experimental and observational data.
- Summit's high-bandwidth, GPU-accelerated architecture should be very effective for data analytics and machine learning.

Acknowledgements

- DOE Office of Science, Advanced Scientific Computing Research
  - Oak Ridge Leadership Computing Program, Manager: Christine Chalk
- Contributions from:
  - ORNL: K. Thach, S. Vazhkudai, S. Klasky, B. Maccabe
  - ORNL: W. Joubert, D. Jacobson, R. Kannan, R. Patton
  - LBNL: Thorsten Kurth, Prabhat, Kjiersten Fagnan
- Collaboration with:
  - OLCF's Users
  - Computer Vendor Partners: Cray, NVIDIA, AMD, IBM, Mellanox, others
  - CORAL Partners: ALCF, LLNL
  - DOE SciDAC Projects & ASCR R&D Programs
  - DOE ECP project