COMPUTING AT ARGONNE

DAVID MARTIN
Manager, Industry Partnerships and Outreach
Argonne Leadership Computing Facility

HPC User Forum
September 10, 2019
Argonne National Laboratory

ARGONNE'S HISTORY OF COMPUTING

The world's first usable amount of electricity from nuclear energy was produced by Experimental Breeder Reactor 1 in Southeastern Idaho and used to light these four light bulbs on December 21, 1951.

AVIDAC, Argonne's first digital computer, began operation in January 1953. It was built by the Physics Division for $250,000. Pictured is pioneer Argonne computer scientist Jean F. Hall.

ARGONNE STRATEGIC PRIORITIES

Hard X-ray sciences	Advanced computing	Universe as our laboratory	Materials and chemistry	Manufacturing science and engineering
			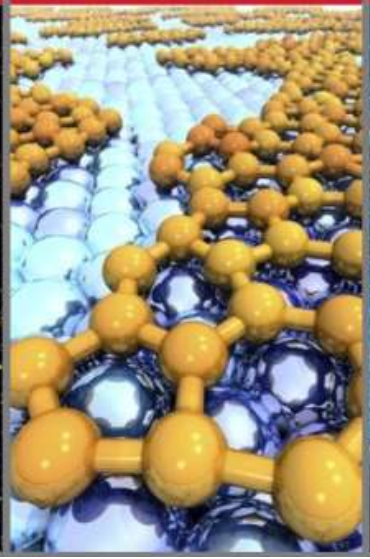	

ARGONNE USER FACILITIES

1. Advanced Photon Source
2. Argonne Leadership Computing Facility
3. Argonne Tandem Linear Accelerator System
4. Atmospheric Radiation Measurement Climate Research Facility
5. Center for Nanoscale Materials

ALCF – A NATIONAL USER FACILITY FOR COMPUTING AND COMPUTATIONAL SCIENCES

- Open to all
  - No restriction on organization, funding source, nationality, or research area
- No charge for open science
  - Cost recovery for proprietary work
- Access through peer-reviewed proposal process
  - Rapid Discretionary access available
- Expert support
  - Dedicated staff help to users utilize unique resources
  - Collaborative work with domain experts

ALCF FOCI

- **Leadership Computing**
  - Maintain a High Performance Computing system designed for the largest, most complex modeling and simulations
  - System in the top 10 of the Top500
- **Expert Support**
  - Computational scientists that are domain scientists and know how to translate problems to computational representations
  - Performance engineers that are experts in porting and optimizing code for massively parallel machines
- **Next Generation Machines and Software**
  - Work with vendors to develop the next generation of HPC
  - Design, procure and install cutting edge computing

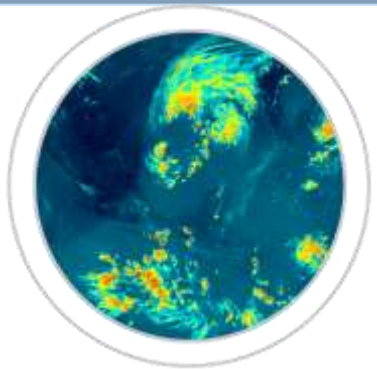

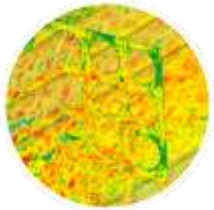

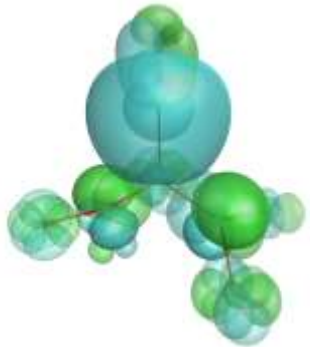

ALCF USER COMMUNITY

- Users are PI-led research teams from academia, national laboratories, and industry
- Represent a wide range of scientific disciplines and a diverse base of users with diverse needs

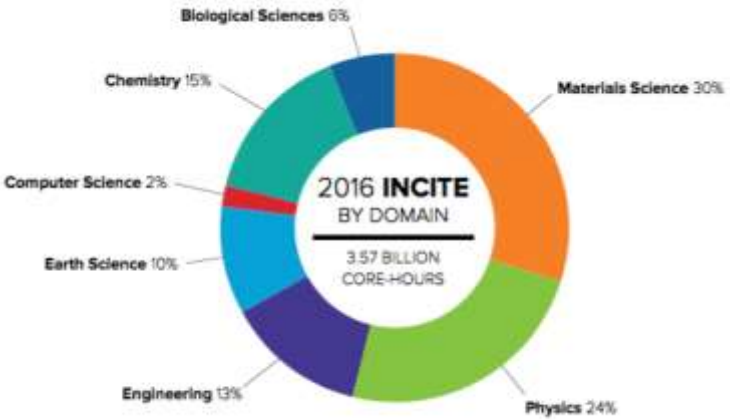

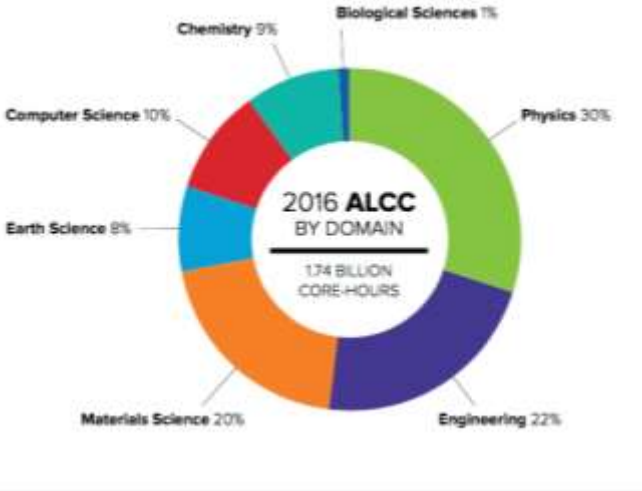

Note: ALCC data is from calendar year 2016.

ALCF SYSTEMS

Mira – IBM BG/Q	Cetus – IBM BG/Q	Vesta – IBM BG/Q	Cooley - Cray/NVIDIA
49,152 nodes 786,432 cores 786 TB RAM Peak flop rate: 10 PF	4,096 nodes 65,536 cores 64 TB RAM Peak flop rate: 836 TF	2,048 nodes 32,768 cores 32 TB RAM Peak flop rate: 419 TF	126 nodes 1512 Intel Haswell CPU cores 126 NVIDIA Tesla K80 GPUs 48 TB RAM / 3 TB GPU

Theta – Intel/Cray

4,392 nodes
281,088 cores
70 TB MCDRAM
679 TB DDR4
562 TB SSD
Peak flop rate: 11.7 PF

Production Storage Capability

HOME: 1.44 PB raw capacity

PROJECT:

- fs0 - 26.88 PB raw, 19 PB usable; 240 GB/s sustained
- fs1 - 10 PB raw, 7 PB usable; 90 GB/s sustained
- fs2 (ESS) - 14 PB raw, 7.6 PB usable; 400 GB/s sustained (not in production yet)

TAPE: 21.25 PB of raw archival storage [17 PB in use]

Argonne NATIONAL LABORATORY

OTHER ARGONNE COMPUTING RESOURCES

- Division Clusters
  - Center for Nanoscale Materials – Carbon
- Transportation Research Computing Center (TRACC)
  - Started as a Department of Transportation resource, but now open to many domains
- Laboratory Computing Resource Center (LCRC)
  - Laboratory-wide computing resource, also open to collaborators

INSTITUTIONAL COMPUTING

LCRC

- Lab-wide resource for mid-range computing and data; LCF on-ramp
- Institutional resources funded by indirect: 1400 nodes, 1.6 petabytes
- Used by all directorates
  - Argonne PIs and their collaborators
  - 350+ projects, 700+ active users

Partnership with Lab Programs

- **Compute/Data hosting & consolidation**
  - 400+ nodes, 1.4 petabytes data
- **Industry Program for SPP, CRADA work**
- **On Demand Resources**
- **Compute and storage strategy w/ EDC**

Bebop – lead production cluster

- **1024 Cray compute nodes, OmniPath**
  - **672 2-socket Broadwell, 128GB**
  - **352 Xeon Phi KNL, 16 GB IPM + 96 GB DRAM**
- **1.5 PF and 46,720 cores**

ALCF EXASCALE SUPERCOMPUTER – AURORA

Intel Aurora supercomputer planned for 2018 shifted to 2021
Scaled up from **180 PF to over 1000 PF**

Support for three “pillars”

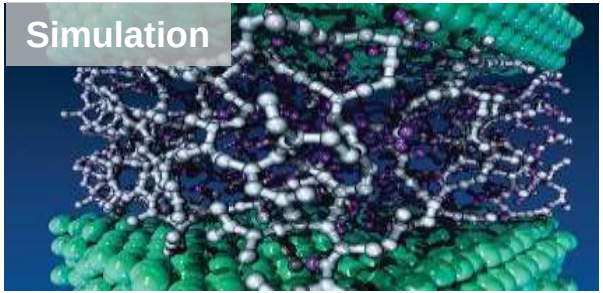

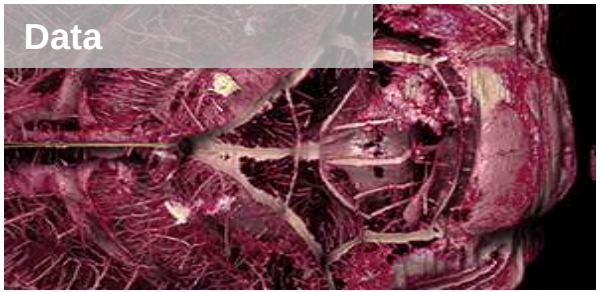

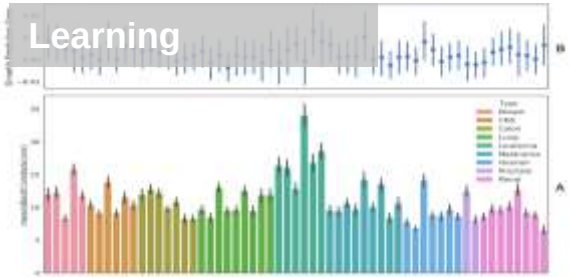

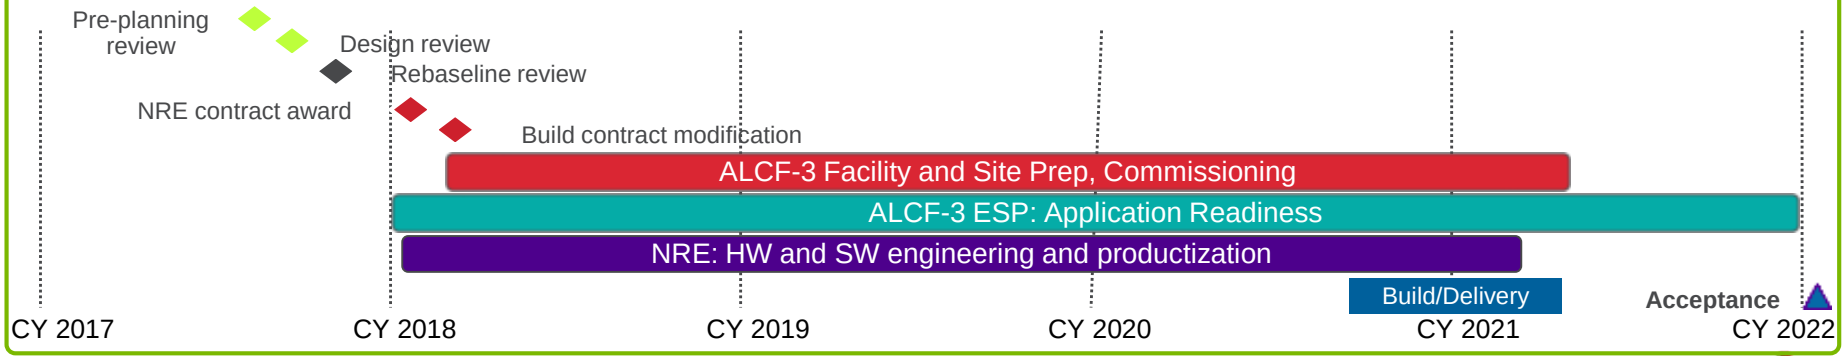

Argonne NATIONAL LABORATORY

AURORA: Argonne's Future Exascale System

Scheduled for delivery in 2021, Aurora is slated to be the US's first exascale supercomputer

- Aurora will feature several technological innovations:
  - A revolutionary I/O system—the Distributed Asynchronous Object Store (DAOS)—to support new types of workloads (230PB, bandwidth >25TB/S)
  - Intel Xeon Scalable Processor + multiple Xe-Architecture based GP-GPUs in each node, >10PB memory
  - Cray Shasta platform with Slingshot interconnect
  - Intel OneAPI, programming techniques in use on current systems will apply directly to Aurora
- Aurora will be highly optimized for simulation, data, and learning applications

The New HPC + AI “Paradigm”

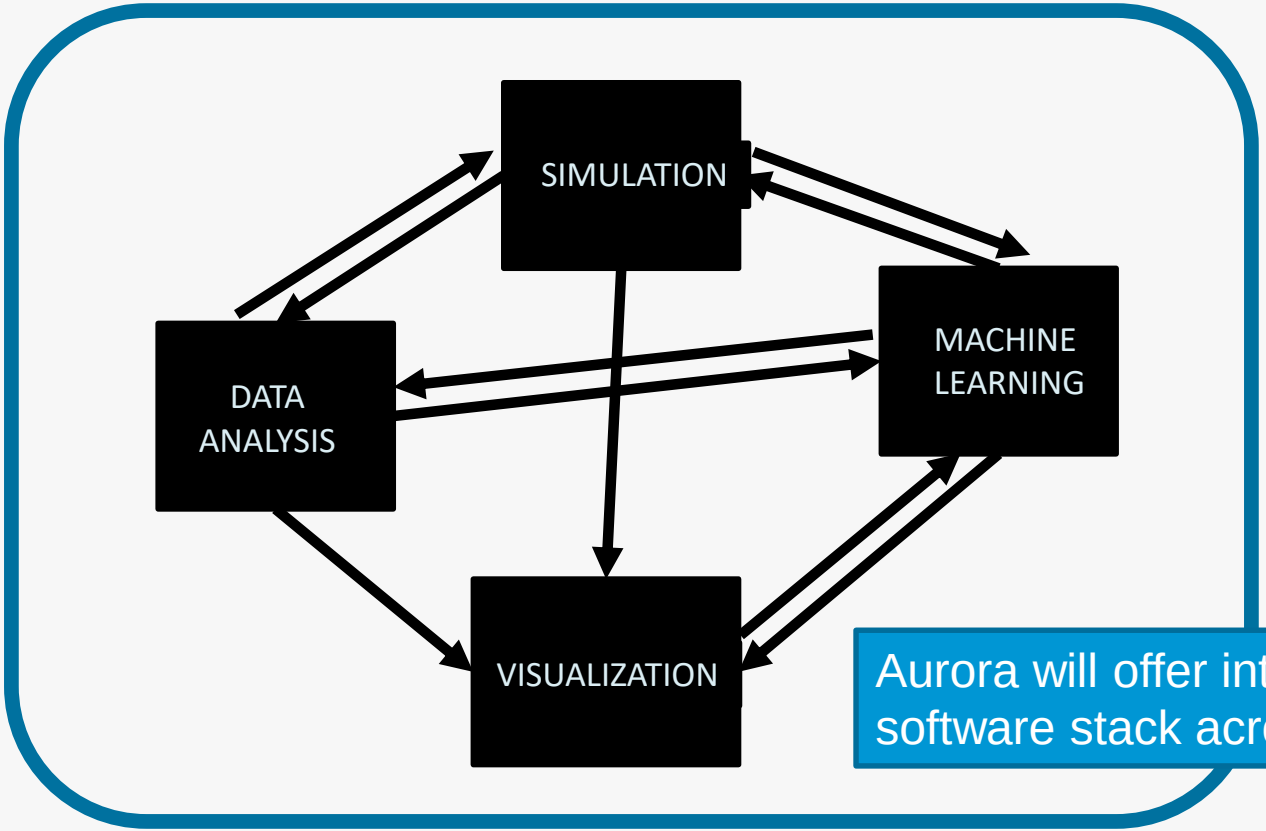

Aurora will offer integrated software stack across all areas

Exascale Programming Realities

- MPI+X
  - Use non-proprietary options
    - OpenMP, OpenCL, SYCL, UPC+Libraries
- Programming Languages
  - C/C++ will come first
  - Fortran will be supported and performant
- Frameworks
  - Primary learning frameworks will be performant
- Data
  - Data Objects
  - Traditional filesystem I/O

LCF Allocation Programs	INCITE 60%	ALCC 20%	Director's Discretionary 20%
Mission	High-risk, high-payoff science that requires LCF-scale resources*	High-risk, high-payoff science aligned with DOE mission	50% Strategic LCF goals 50% ECP
Call	1x/year – Closes June	1x/year – Closes February	Rolling
Duration	1-3 years, yearly renewal	1 year	3m,6m,1 year
Typical Size	**10-15 projects** **1-3M node-hours**	**10-15 projects** **0.5-2M node-hours**	**~100 of projects** **<0.5M node-hours**
Total Hours	**~17.8M Theta**	**~6M node-hours Theta**	**~6M node-hours Theta**
Review Process	Scientific Peer-Review; Computational Readiness	Scientific Peer-Review; Computational Readiness	Strategic impact and feasibility
Managed By	INCITE management committee (ALCF & OLCF)	DOE Office of Science	LCF management
Readiness	**High**	**Medium to High**	**Low to High**
Availability	Open to all scientific researchers and organizations ***Capability > 20% of resource***		

Aurora Early Science Program: Productive Science on Day One

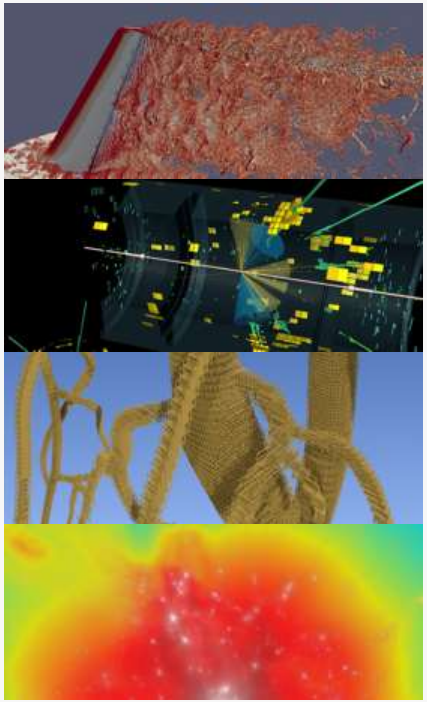

Designed to prepare key applications for the scale and architecture of Aurora

- Supports 15 projects (5 simulation, 5 data, and 5 learning), covering a wide range of scientific areas and numerical methods
- Investigating computational research areas critical to enabling science in the exascale era
  - Mapping and optimizing complex workflows
  - Exploring new machine learning methodologies
  - Stress testing I/O hardware and other emerging technologies
  - Enabling connections to large-scale experimental data sources

ENGAGING WITH ARGONNE

- INCITE and ALCC
- Director’s Discretionary Allocation
- ALCF Early Science Program
- ALCF Data Sciences Program
- Exascale Computing Project
- HPC4Energy Innovations
  - HPC4Manufacturing, HPC4Materials, HPC4Mobility
- CRADA, SPP, proprietary agreements
- Informal Collaborations

COMPUTING EXPERTISE

CELS Divisions Work Together to Solve Critical Problems

- Mathematics and Computer Science
  - Tools and technology for solving critical scientific problems
- Computational Science Division
  - Solving the most challenging scientific problems through advanced modeling and simulation on the most capable computers
- Data Science and Learning
  - Data analysis and artificial intelligence to provide critical insights and accelerate discovery
- Joint Laboratory for System Evaluation
  - Evaluate future hardware and software platforms in partnership with vendors and users

EMERGING AND GROWING EFFORTS

- ABLE - Argonne Biomedical Learning Enclave
  - Establishing the necessary policies, procedures, controls and systems to support analysis of human subject data
  - Moderate enclave that implements NIST 800-53. Eventually compliance with HIPAA and FISMA requirements
- Quantum Computing
  - Architectures, software environments, algorithms, hybrid systems, ecosystems
- AI for Science

Argonne

NATIONAL LABORATORY