Argonne
NATIONAL LABORATORY

ARGONNE'S AURORA EXASCALE COMPUTER

SUSAN COGHLAN
Aurora Technical Lead and ALCF-3 Project Director

29 August 2019
Smoky Mountain Computational Sciences and Engineering Conference

THE BEGINNING

- 2014: the first CORAL RFP was issued by Argonne, Oak Ridge, and Livermore national labs for three next-generation supercomputers to replace Mira, Titan, and Sequoia
  - Two winning proposals were selected, one with IBM/NVIDIA (Summit and Sierra) and one with Intel/Cray
- 2015: the CORAL contract between Argonne and Intel for two systems was awarded
  - Theta, a small Intel KNL based system intended to bridge between ALCF's current many-core IBM BlueGene/Q system, Mira (delivered in 2012) and Aurora
  - Aurora, a 180PF Intel KNH based many-core system intended to replace Mira, scheduled for delivery in 2018
- 2016: Theta was delivered and accepted – well ahead of schedule

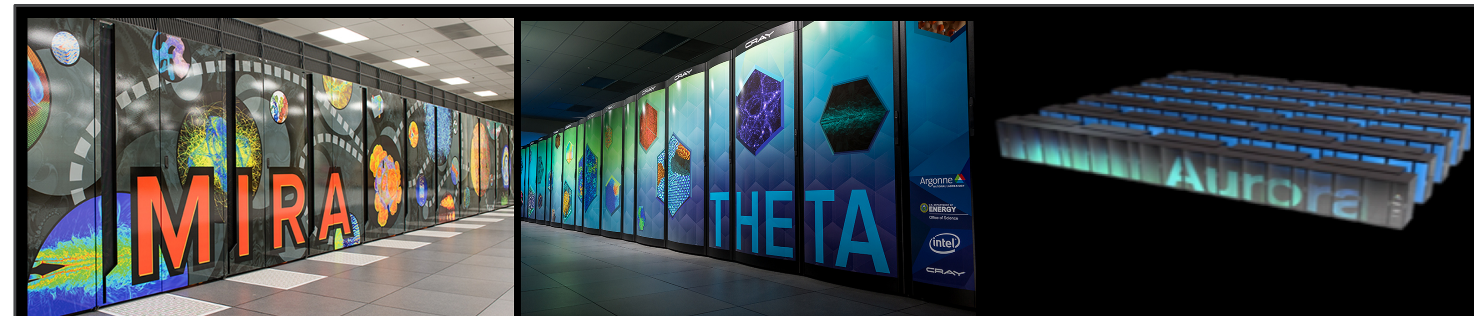

Mira IBM BG/Q - 2012 | **Theta** Intel KNL - 2016 | **Aurora** Intel KNH - coming in 2018

THE CHANGE TO EXASCALE

- 2016: DOE began exploring opportunities to deliver exascale computing earlier than planned
  - DOE revised the target delivery date from 2023 to 2021 based on discussions with vendors and information from an RFI
- 2017: KNH was delayed and Argonne received guidance from DOE to shift from the planned 180PF in 2018 to an exascale system in 2021
- 2018: after **many** reviews, the ALCF-3 project was re-baselined to deliver an exascale system in CY2021
- 2019: after more reviews, contract modifications were completed and the exascale Aurora system was announced
  - Preparations underway in facility improvements, software and tools, and early science

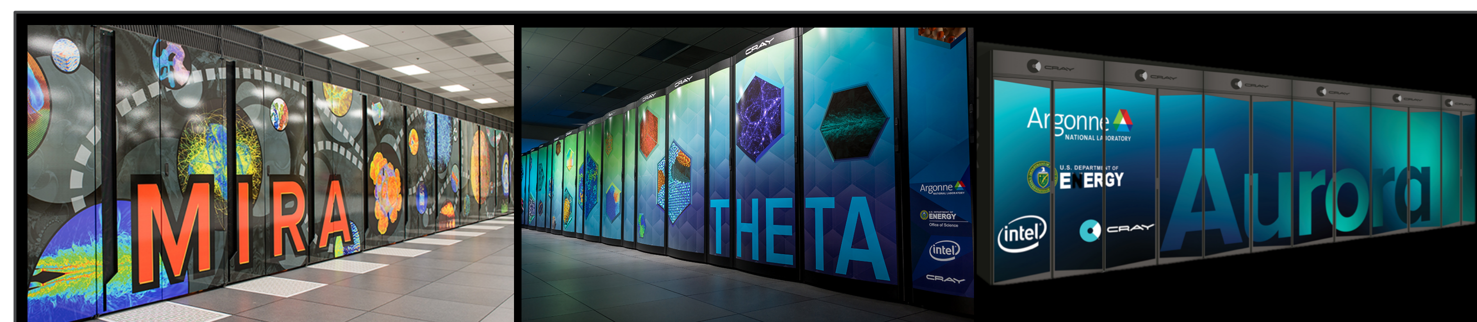

Mira IBM BG/Q - 2012 | **Theta** Intel KNL - 2016 | **Aurora** Intel X^e - coming in 2021

DOE MISSION NEED

- … requires **exascale systems with a 50-100x increase in application performance over today's DOE leadership deployments in the 2021-2023 timeframe**
- Advanced exascale computers needed to model and simulate complex natural phenomena, sophisticated engineering solutions, and to solve a new and emerging class of data science problems
  - Use of **rich analytics and deep learning software**, coupled with simulation software, to derive insights from experimental/observational facilities data
  - Size and complexity of these datasets requires leadership computing resources
  - **Sophisticated data mining algorithms** are needed to steer experiments and simulations in real-time
- DOE leadership computer resources **must support:** statistics, machine learning, deep learning, uncertainty quantification, databases, pattern recognition, image processing, graph analytics, data mining, real time data analysis, and complex and interactive workflows

REQUIREMENTS DRIVING THE DESIGN

- Exascale system delivered in CY2021
- 50X over 20PF Titan/Sequoia for representative applications
  - Aligns with Exascale Computing Project (ECP) application performance goals
- Full support for Simulation, Data, and Learning
  - Includes requirements for optimized Data and Learning frameworks
- Productive user environment for the leadership computing community
  - All the standard stuff
- CORAL RFP requirements
  - Added in requirements to support new targets, in particular, added learning and data application benchmarks
- Within ALCF's budget

Primary driver is to provide the best balanced system (within constraints) for Simulation, Data, and Learning science at the Argonne Leadership Computing Facility (ALCF)

AURORA HIGH-LEVEL CONFIGURATION (PUBLIC)

System Spec	Aurora
Sustained Performance	≥1EF DP
Compute Node	Intel Xeon scalable processor Multiple X^e arch based GP-GPUs
Aggregate System Memory	>10 PB
System Interconnect	Cray Slingshot - 100 GB/s network bandwidth Dragonfly topology with adaptive routing
High-Performance Storage	≥230 PB, ≥25 TB/s (DAOS)
Programming Models	Intel OneAPI, OpenMP, DPC++/SYCL
Software stack	Cray Shasta software stack + Intel enhancements + Data and Learning
Platform	Cray Shasta
# Cabinets	>100

SOFTWARE AND TOOLS (PUBLIC)

Area	Aurora
Compilers	Intel, LLVM, GCC
Programming languages and models	Fortran, C, C++ OpenMP 5.x (Intel, Cray, and possibly LLVM compilers), UPC (Cray), Coarray Fortran (Intel), Data Parallel C++ (Intel and LLVM compilers), OpenSHMEM, Python, MPI
Programming tools	Open|Speedshop, TAU, HPCToolkit, Score-P, Darshan, Intel Trace Analyzer and Collector Intel Vtune, Advisor, and Inspector PAPI, GNU gprof
Debugging and Correctness Tools	Stack Trace Analysis Tool, gdb, Cray Abnormal Termination Processing
Math Libraries	Intel MKL, Intel MKL-DNN, ScaLAPACK
GUI and Viz APIs, I/O Libraries	X11, Motif, QT, NetCDF, Parallel NetCDF, HDF5
Frameworks	TensorFlow, PyTorch, Scikit-learn, Spark Mllib, GraphX, Intel DAAL, Intel MKL-DNN

AURORA EARLY SCIENCE PROGRAM (ESP)

AURORA | EARLY SCIENCE PROGRAM

http://esp.alcf.anl.gov

Applications Readiness

- Prepare applications for Aurora system
  - Architecture
  - *Exascale*
- 5 Simulation, 5 Data, 5 Learning projects
  - Competitively chosen from proposals, based on **Exascale science** calculation and development plan
- 240+ team members, ~2/3 are core developers

- 10 unique traditional simulation applications (compiled C++/C/F90 codes)
- Extensive dependence on ML/DL frameworks
- 10 complex multi-component workflows
  - Includes experimental data
- 3 major Python-only applications

Support

PEOPLE

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

TRAINING

- Training on HW and programming (COE)
- Capturing best practices to share with the community (e.g. Performance Portability Workshop)

COMPUTE RESOURCES

- Current ALCF production systems
- Early next-gen hardware and software
- Test runs on full system pre-acceptance
- 3 months dedicated Early Science access
  - Pre-production (post-acceptance)
  - Large time allocation, access for rest of year

ALCF AURORA ESP SIMULATION PROJECTS

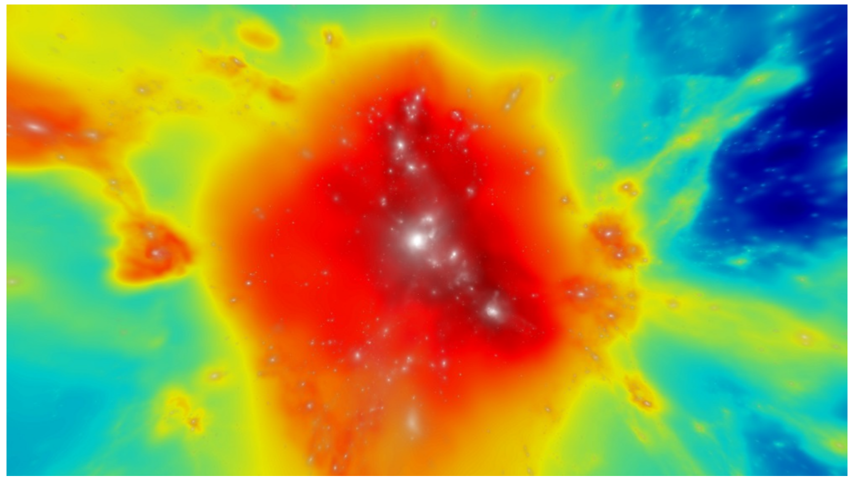
Katrin Heitmann, Argonne National Laboratory

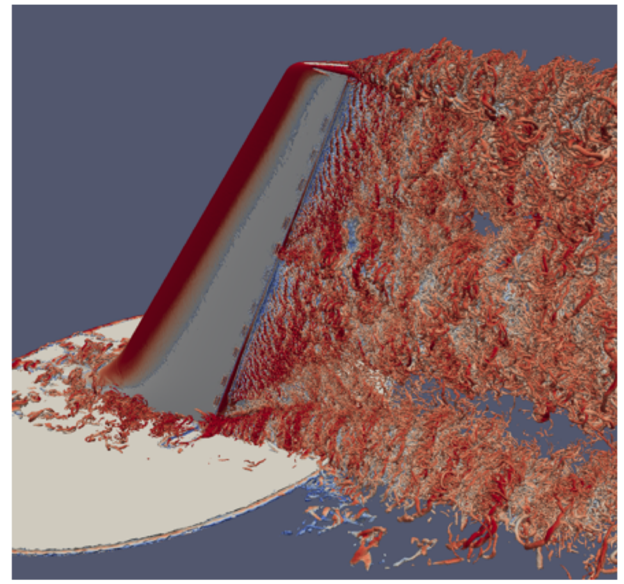
Ken Jansen, U. of Colorado Boulder

Extending Moore's Law computing with Quantum Monte Carlo
Anouar Benali, Argonne National Laboratory
Using QMC simulations, this project aims to advance our knowledge of the HfO2/Si interface necessary to extend Si-CMOS technology beyond Moore's law.

High fidelity simulation of fusion reactor boundary plasmas
C.S. Chang, PPPL
By advancing the understanding and prediction of plasma confinement at the edge, the team's simulations will help guide fusion experiments, such as ITER, and accelerate efforts to achieve fusion energy production.

NWChemEx: Tackling Chemical, Materials & Biochemical Challenges in the Exascale Era
Teresa Windus, Iowa State University and Ames Laboratory
This project will use NWChemEx to address two challenges In the production of advanced biofuels: the development of stress-resistant biomass feedstock and the development of catalytic processes to convert biomass-derived materials into fuels.

Extreme-Scale Cosmological Hydrodynamics
Katrin Heitmann, Argonne National Laboratory
Researchers will perform cosmological hydrodynamics simulations that cover the enormous length scales characteristic of large sky surveys, while at the same time capturing the relevant small-scale physics. This work will help guide and interpret observations from large-scale cosmological surveys.

Extreme Scale Unstructured Adaptive CFD: From Multiphase Flow to Aerodynamic Flow Control
Ken Jansen, University of Colorado Boulder
This project will use unprecedented high-resolution fluid dynamics simulations to model dynamic flow control over airfoil surfaces at realistic flight conditions and to model bubbly flow of coolant in nuclear reactors.

ALCF AURORA ESP DATA PROJECTS

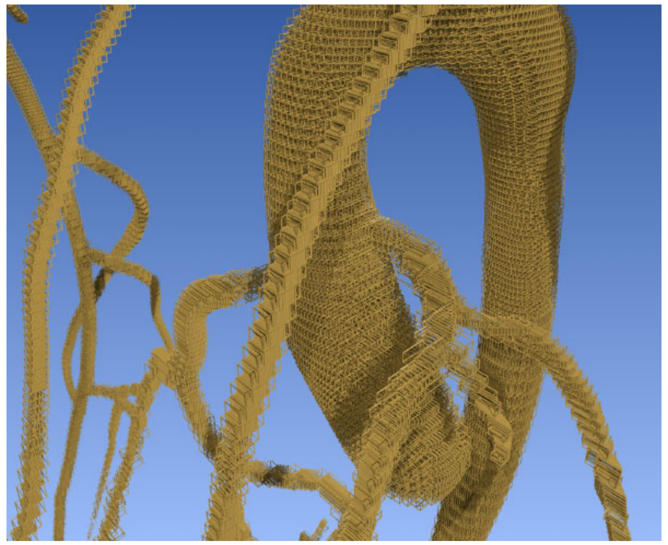

Amanda Randles, Duke University and ORNL

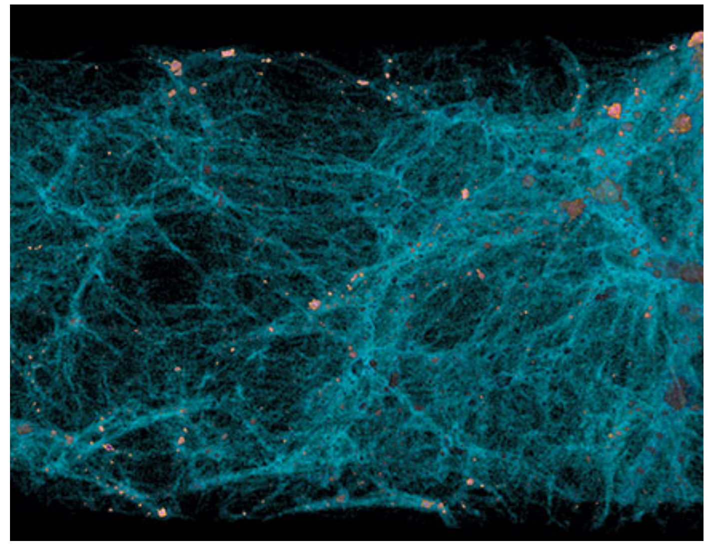

Salman Habib, Argonne

Simulating and Learning in the ATLAS Detector at the Exascale
James Proudfoot, Argonne National Laboratory
This project will develop exascale workflows and algorithms that meet the growing computing, simulation and analysis needs of the ATLAS experiment at CERN's LHC.

Dark Sky Mining
Salman Habib, Argonne National Laboratory
By implementing cutting-edge data-intensive and machine learning techniques, this project will usher in a new era of cosmological inference targeted for the Large Synoptic Survey Telescope (LSST).

Extreme-Scale In-Situ Visualization and Analysis of Fluid-Structure-Interaction Simulations
Amanda Randles, Duke University and Oak Ridge National Laboratory
The research team will develop computational models to provide detailed analysis of the role key biological parameters play in determining tumor cell trajectory in the circulatory system.

Exascale Computational Catalysis
David Bross, Argonne National Laboratory
Researchers will develop software tools to facilitate and significantly speed up the quantitative description of crucial gas-phase and coupled heterogeneous catalyst/gas-phase chemical systems.

Data Analytics and Machine Learning for Exascale Computational Fluid Dynamics
Ken Jansen, University of Colorado Boulder
This project will develop data analytics and machine learning techniques to greatly enhance the value of flow simulations, culminating in the first flight-scale design optimization of active flow control on an aircraft's vertical tail.

ALCF AURORA ESP LEARNING PROJECTS

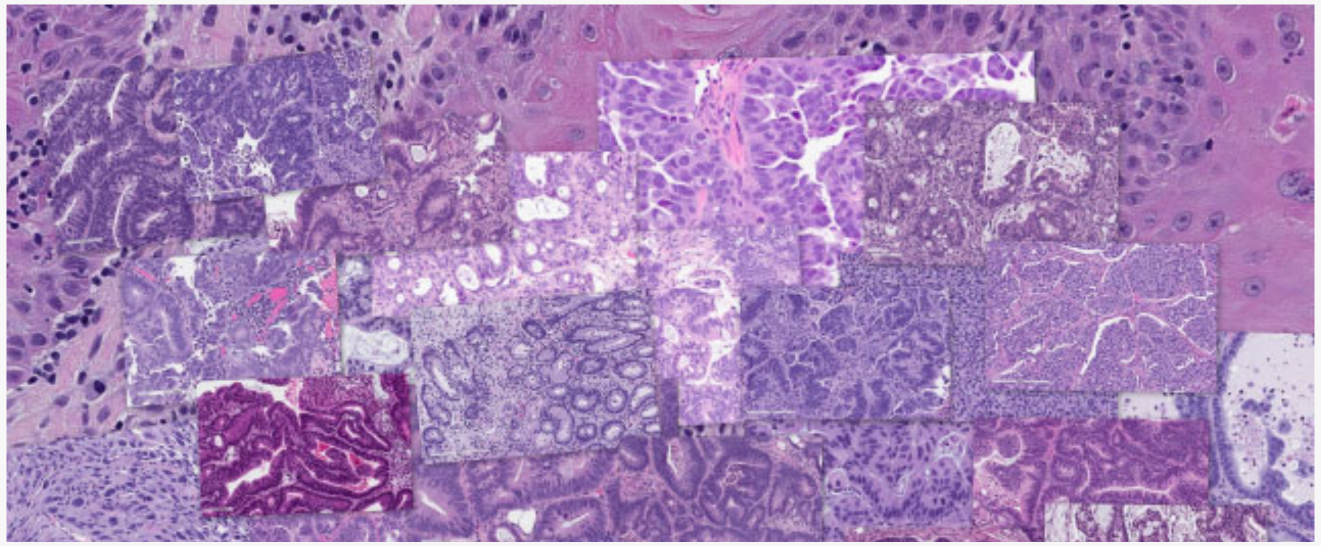

Rick Stevens, Argonne

Virtual Drug Response Prediction
Rick Stevens, Argonne National Laboratory
Utilizing large-scale data frames and a deep learning workflow, researchers will enable billions of virtual drugs to be screened singly and in numerous combinations, while predicting their effects on tumor cells.

Machine Learning for Lattice Quantum Chromodynamics
William Detmold, Massachusetts Institute of Technology
This project couples machine learning and simulations to unravel the mysteries of dark matter while simultaneously providing insights into fundamental particle physics.

Many-Body Perturbation Theory Meets Machine Learning to Discover Singlet Fission Materials
Noa Marom, Carnegie Mellon University
By combining quantum-mechanical simulations with machine learning and data science, this project will harness Aurora's exascale power to revolutionize the computational discovery of new materials for more efficient organic solar cells.

Accelerated Deep Learning Discovery in Fusion Energy Science
William Tang, Princeton Plasma Physics Laboratory
This project will use deep learning and artificial intelligence methods to improve predictive capabilities and mitigate large-scale disruptions in burning plasmas in tokamak systems, such as ITER.

Enabling Connectomics at Exascale to Facilitate Discoveries in Neuroscience
Nicola Ferrier, Argonne National Laboratory
This project will develop a computational pipeline for neuroscience that will extract brain-image-derived mappings of neurons and their connections from electron microscope datasets too large for today's most powerful systems.

AURORA ESP DATA AND LEARNING METHODS

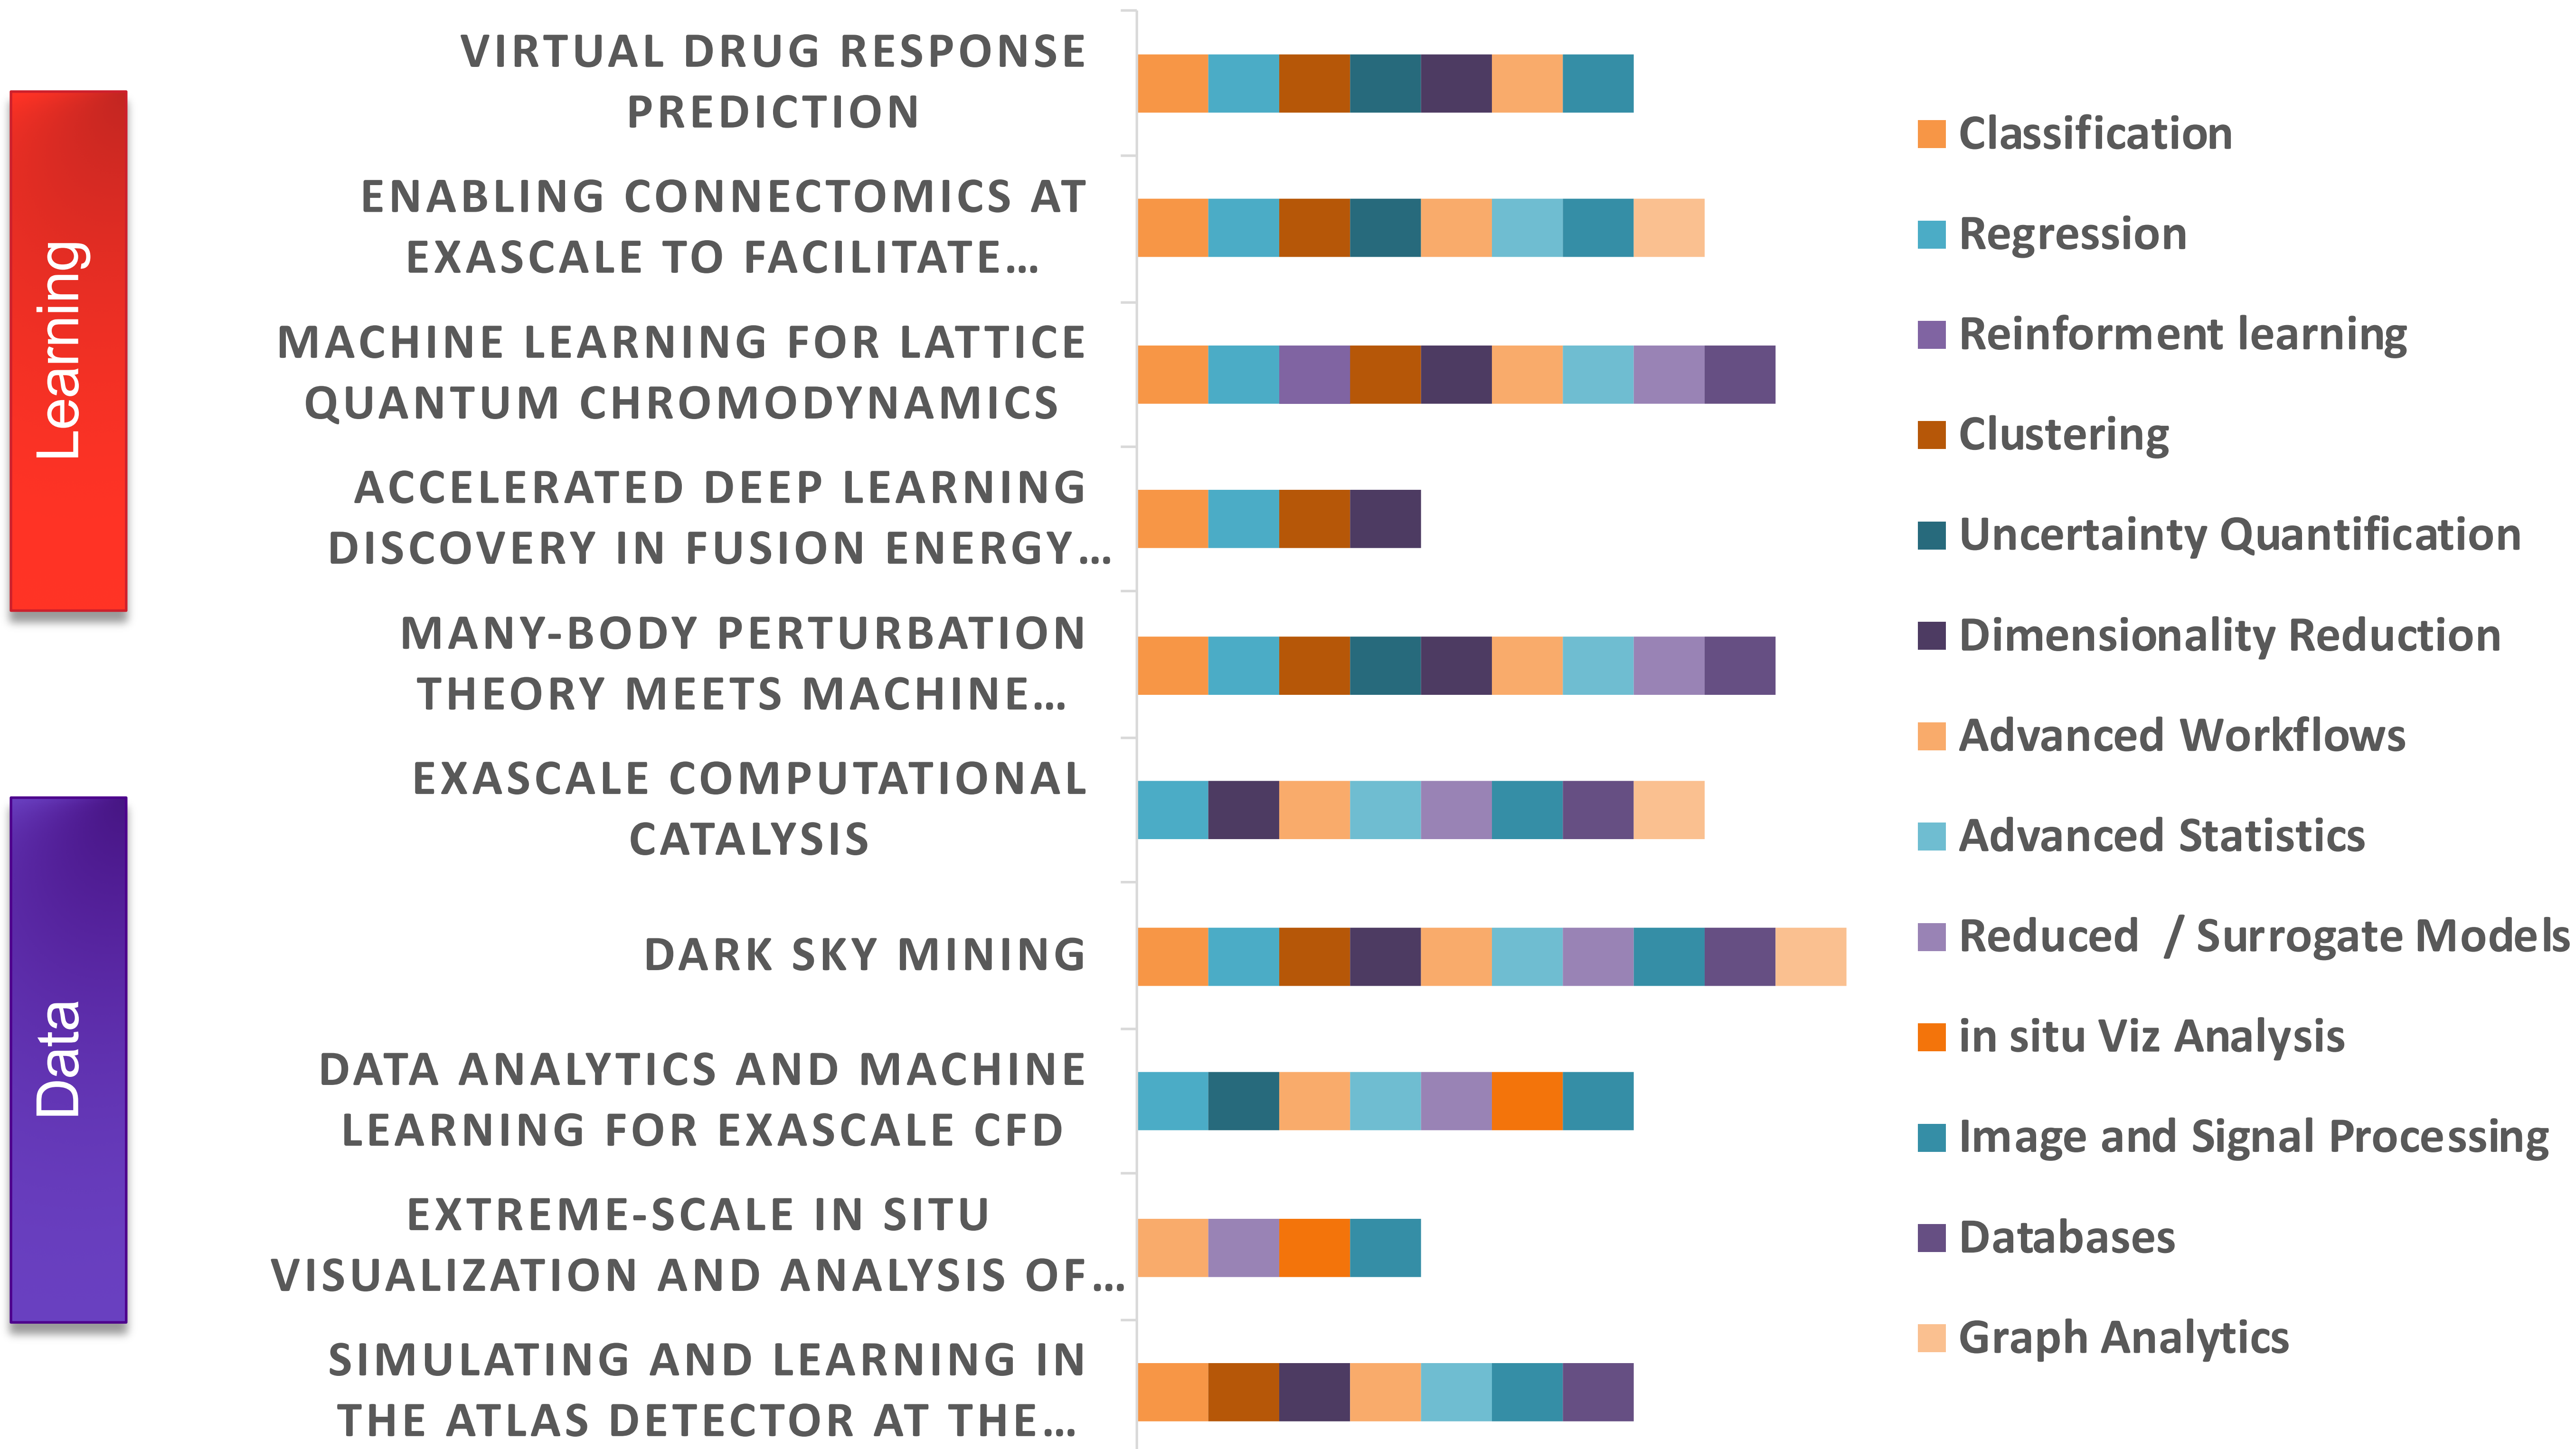

AURORA ESP DATA AND LEARNING METHODS

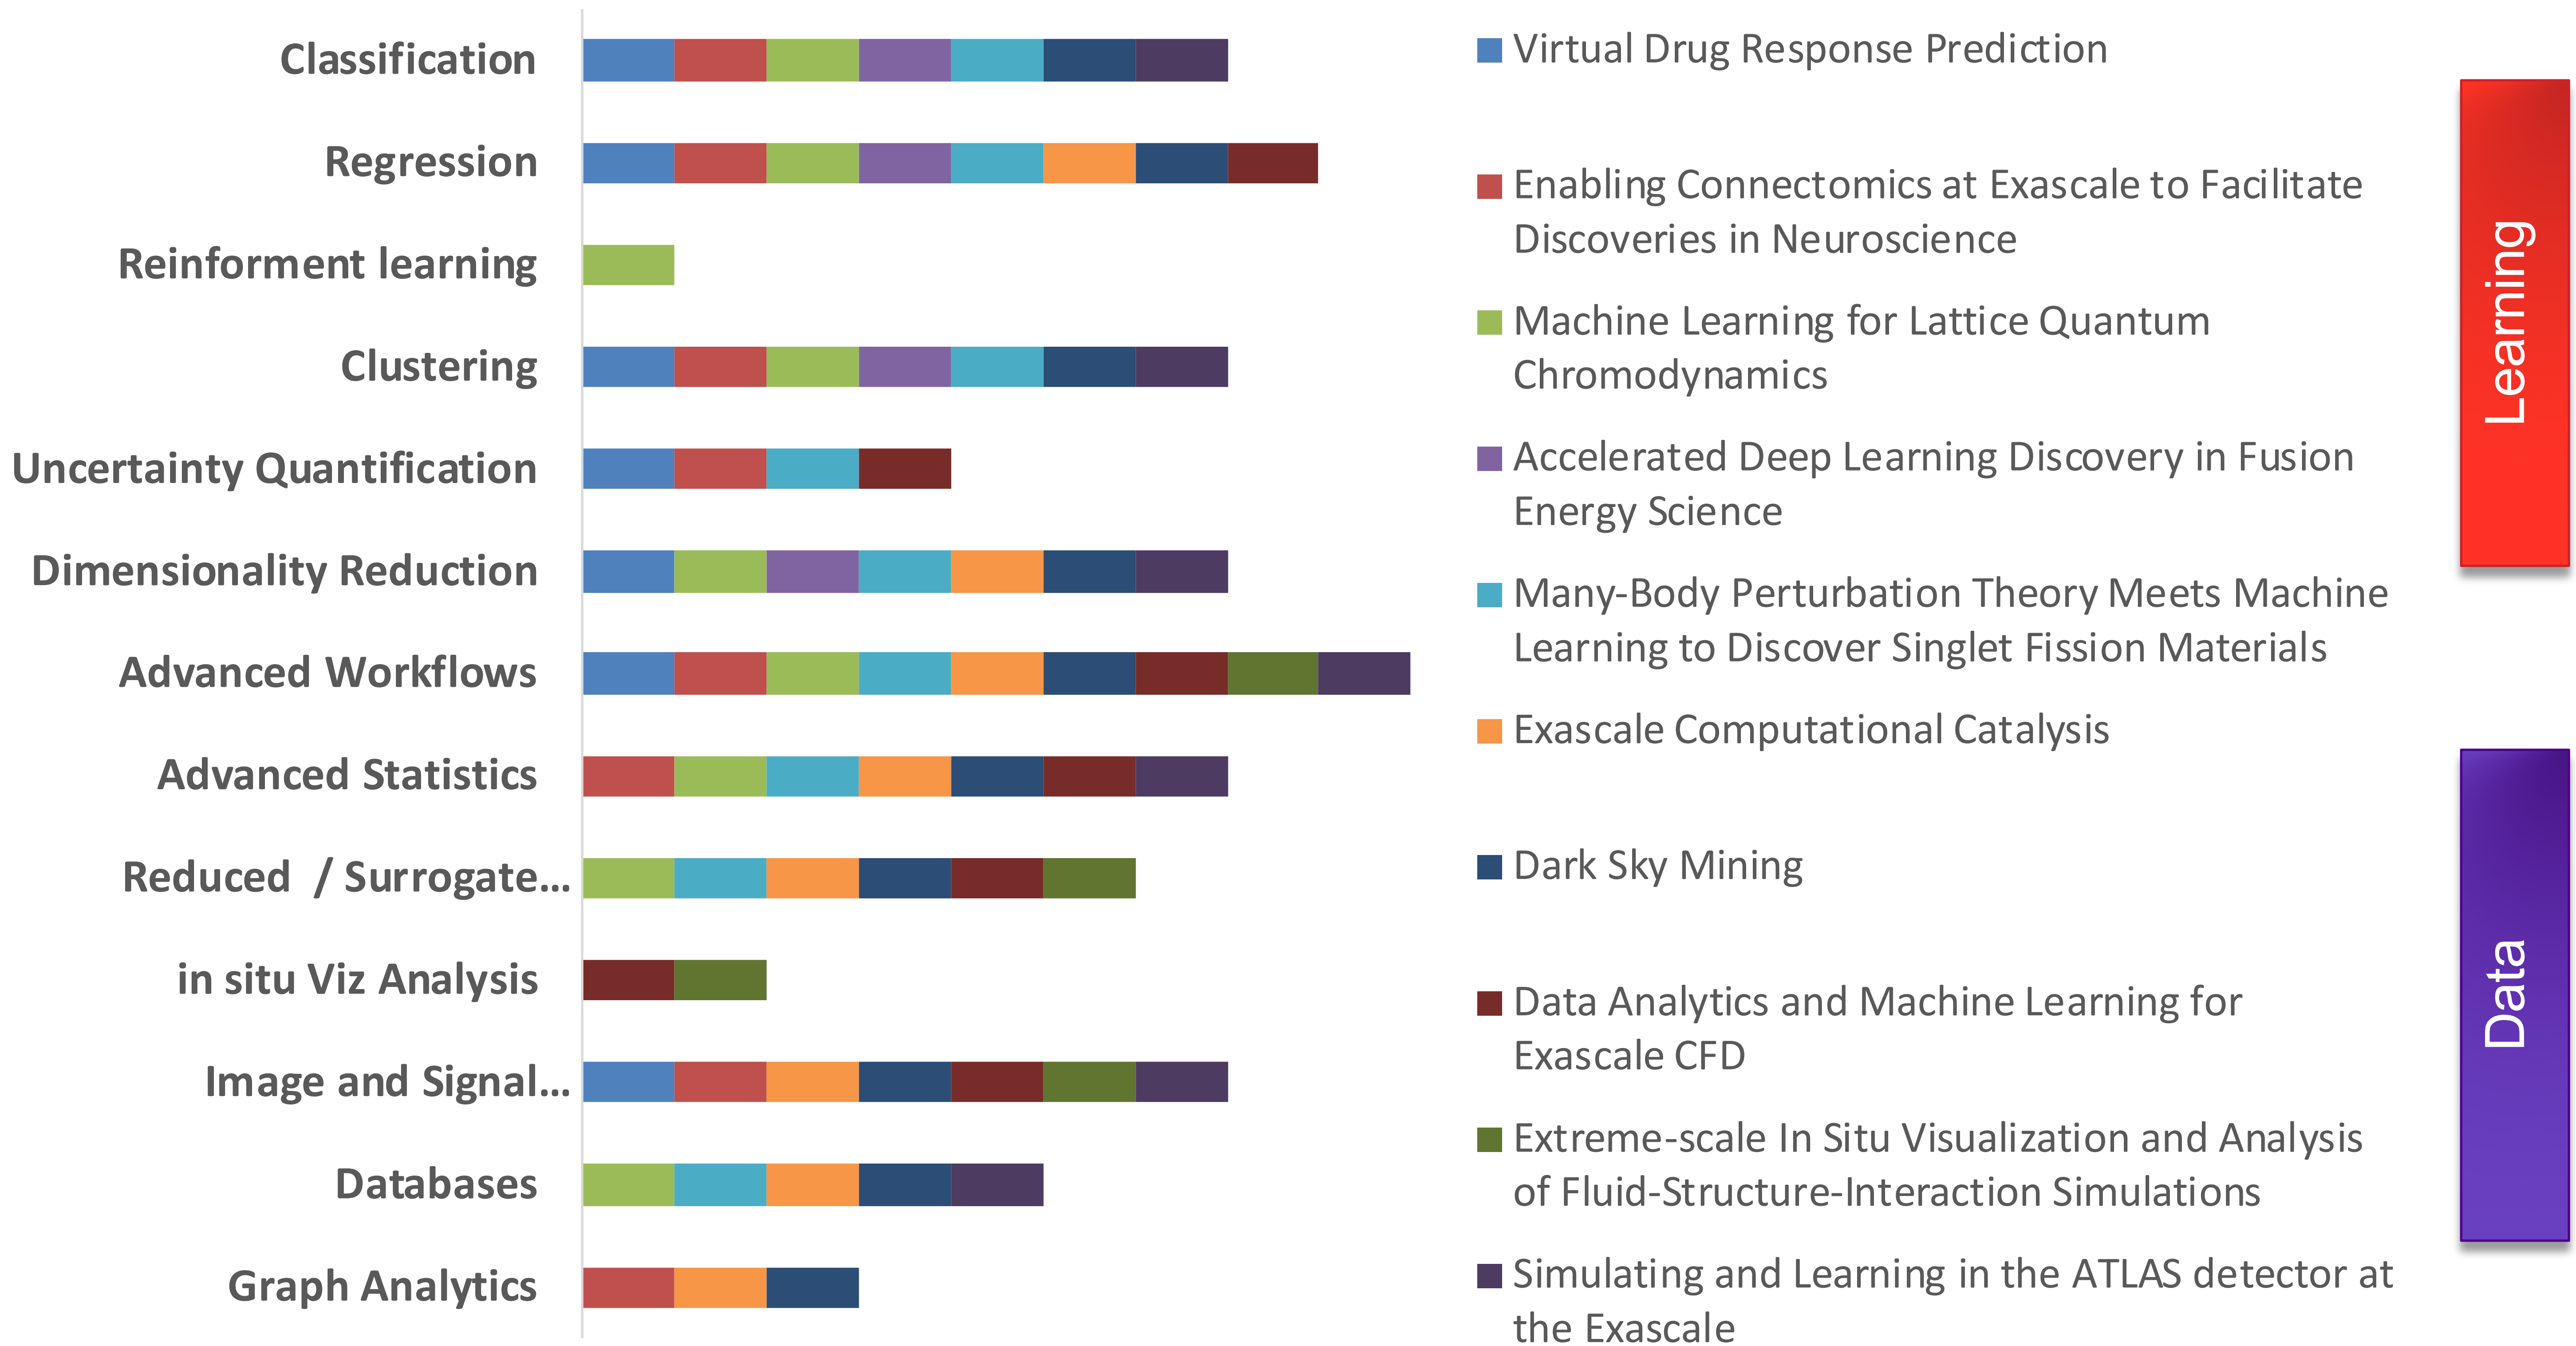

READINESS EFFORTS FOR SIMULATION, DATA, AND LEARNING

15 ALCF-3 Early Science

22 ECP AD Projects

- Flat profile at scale on existing systems
- Flat profile for representative problem
- Characterize kernels of interest
- Use hardware specific advising tools
- Run kernels using representative HW and software SDK

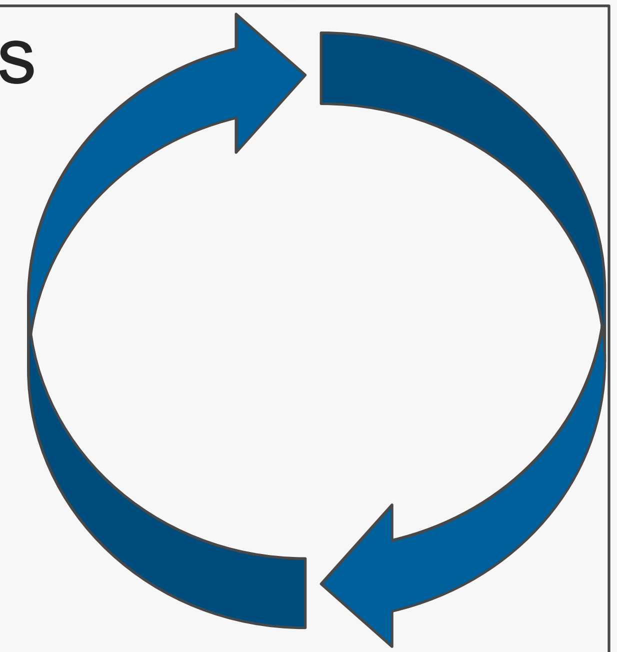

- Prepare workflow technologies
- Optimize libraries, frameworks, and tools
- Harden SW stack

Using: Intel Gen9 GPU + SDK → Early Hardware → Hardware

Feed-back: HW

SW

CURRENT HARDWARE AND SOFTWARE

- Argonne's Joint Laboratory for System Evaluation (JLSE) provides testbeds for Aurora – available today (under the appropriate NDAs)
  - Intel Xeons with Gen 9 Iris Pro Graphics (integrated) (20 nodes)
  - Intel's Aurora software development kit with frequent updates and bug fixes, includes C, C++, Fortran compilers, MKL and DAAL libraries, Vtune, Advisor, SYCL and DPC++, OpenCL, OpenMP 5 (basic)
  - Early generation hardware will be added to JLSE as it becomes available
- Argonne/Intel Outreach
  - Workshops – virtual and in-person (under appropriate NDAs)
  - Hackathons – for targeted ESP projects
  - ESP Training Series – webinars on various topics – upcoming: Machine Learning with TensorFlow, Horovod, and PyTorch on HPC (Sept 9)
  - Sessions at ECP Annual Meeting on Aurora, OpenMP, SYCL, DAOS, and preparing applications for Aurora

SUMMARY

- It's been a long (in computer time) winding road but we are on track and focused on delivering Aurora in 2021
- Aurora will be an Intel Xeon/X^e Exascale system using the Cray Shasta platform
- Aurora blends a user-familiar software stack and programming models with exciting new technologies such as the high-performance DAOS and the new Cray Slingshot interconnect
- A primary driver for Aurora was to provide the best-balanced system for Simulation, Data, and Learning under the given constraints
- Argonne is targeting 5 Simulation, 5 Data, and 5 Learning ESP projects to be ready for Aurora when it arrives, along with the ECP AD projects
- Preparations are well underway in the Facility, Software and Tools, and Early Science

THANK YOU