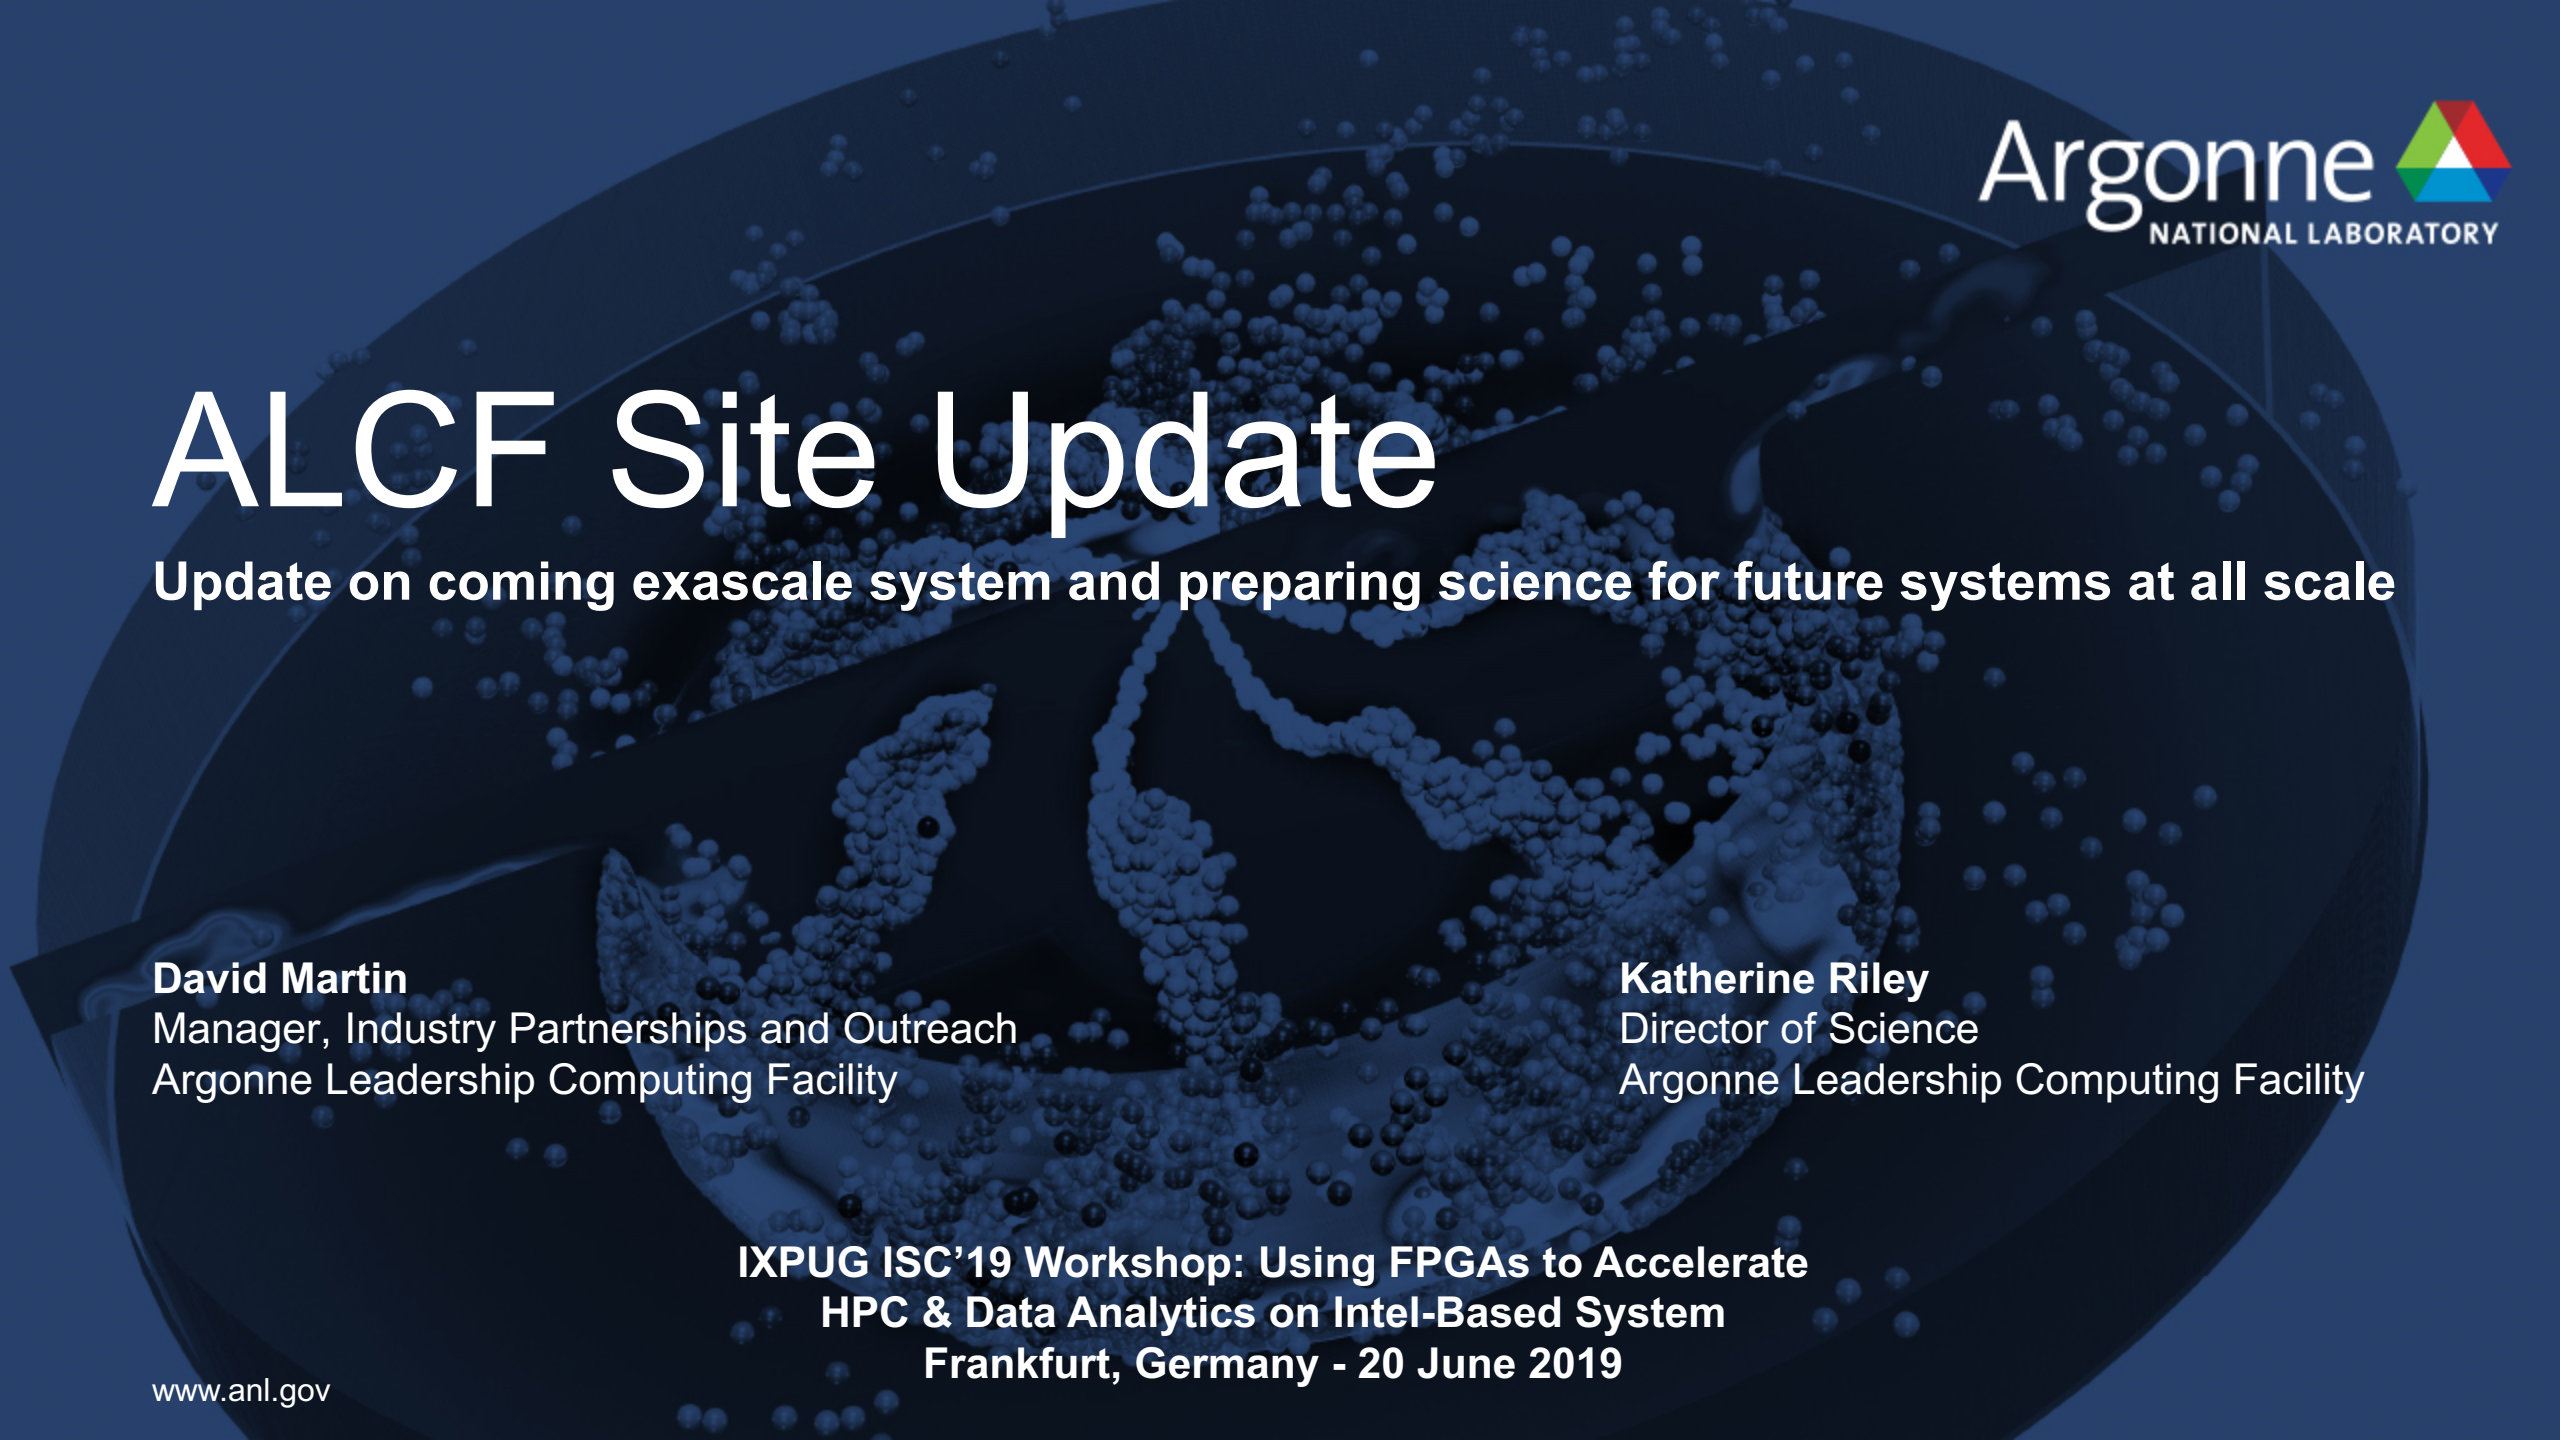

ALCF Site Update

Update on coming exascale system and preparing science for future systems at all scale

David Martin
Manager, Industry Partnerships and Outreach
Argonne Leadership Computing Facility

Katherine Riley
Director of Science
Argonne Leadership Computing Facility

IXPUG ISC'19 Workshop: Using FPGAs to Accelerate HPC & Data Analytics on Intel-Based System
Frankfurt, Germany - 20 June 2019

DOE Office of Science Advanced Scientific Computing Research Facilities

Providing the Facilities – High-End and Leadership Computing

National Energy Research Scientific Computing Center (NERSC)
Lawrence Berkeley National Laboratory

– Delivers high-end capacity computing to entire DOE SC research community / **Over 6,000 users and 800 projects**

Leadership Computing Centers at Argonne National Laboratory (ALCF) and Oak Ridge National Laboratory (OLCF)

– Delivers highest computational capability
  - Open to national and international researchers, including industry
  - Not constrained by existing DOE or Office of Science funding or topic areas
    – Approximately 1,000 users and 50-60 large projects at each center, each year

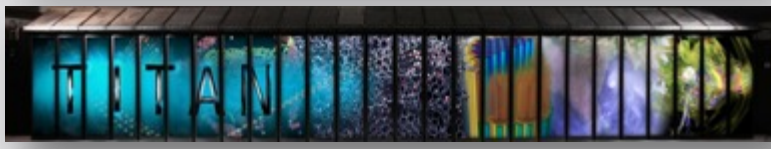

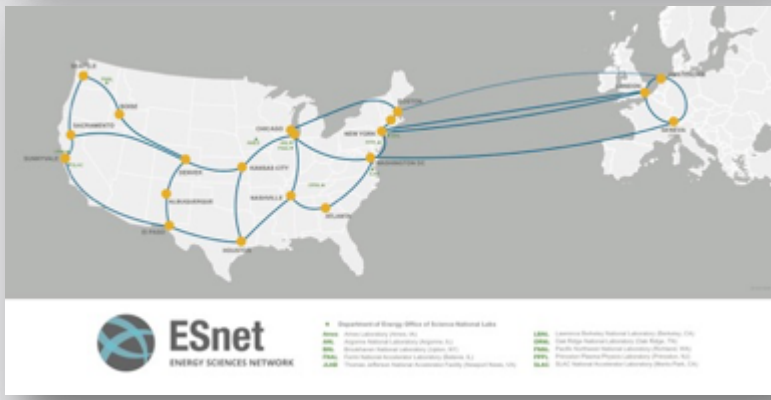

Linking it all together – Energy Sciences Network (ESnet)

DOE Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

Pre-Exascale Systems | **First US Exascale Systems**

2012	2016	2018	2020	2021-2023

ORNL
Cray/AMD/NVIDIA

ORNL
IBM/NVIDIA

ORNL
Cray/AMD

ANL
IBM BG/Q

ANL
Cray/Intel KNL

ANL
Intel/Cray

LBNL
Cray/Intel Xeon/KNL

LBNL
Cray/AMD/NVIDIA

LLNL
IBM BG/Q

LANL/SNL
Cray/Intel Xeon/KNL

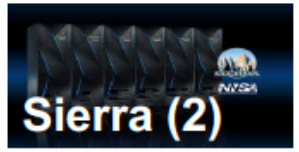

LLNL
IBM/NVIDIA

LANL/SNL
TBD

LLNL
TBD

ECP EXASCALE COMPUTING PROJECT

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"

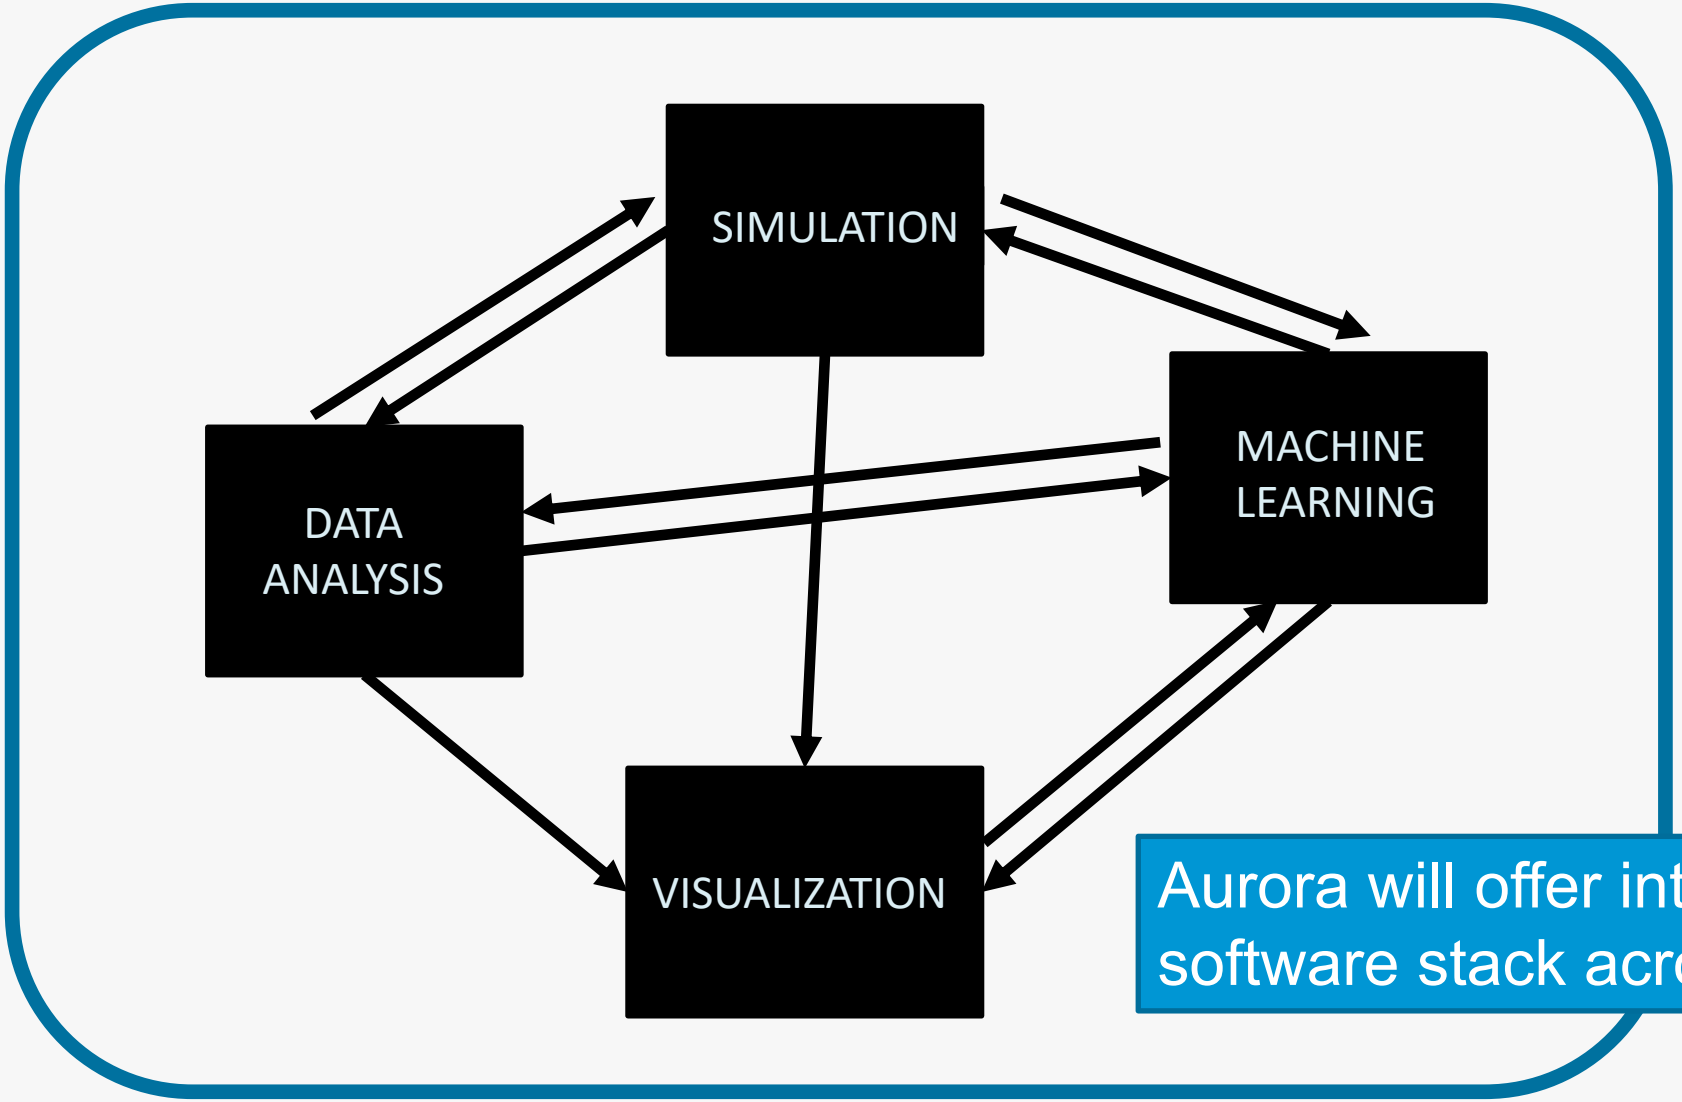

Aurora will offer integrated software stack across all areas

Aurora Early Science Program: Productive Science on Day One

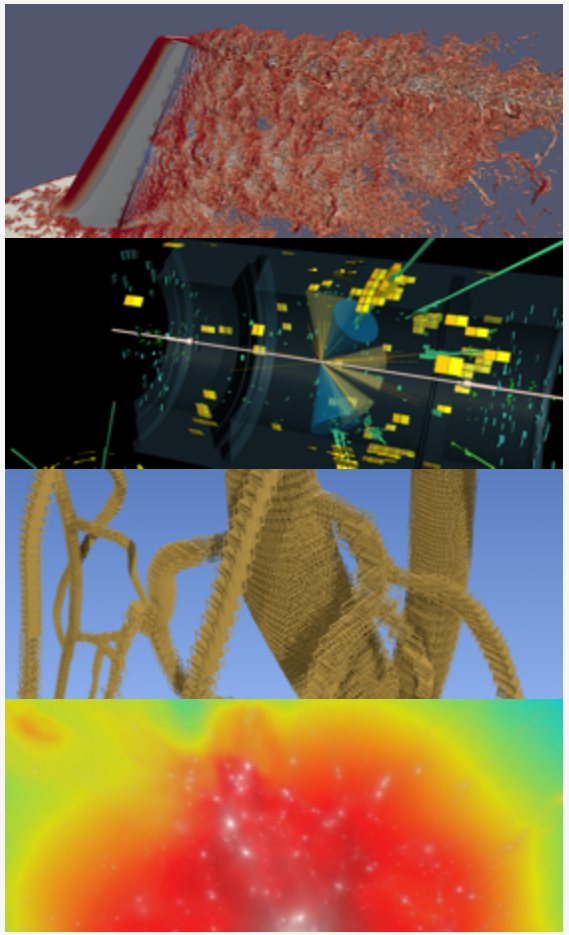

Designed to prepare key applications for the scale and architecture of Aurora

- 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

Programming Model Status for Exascale

MPI+OpenMP with Fortran/C/C++ is plan of record
Standard Learning Frameworks

Zoo of new developments in programming models
- Likely to coalesce
- Critical mass is important
- Performance, abstractions/survivability, interoperability

- On node: ECP projects are experimenting with new approaches that aim at device portability: KOKKOS, RAJA, OpenACC, OpenCL, Swift, etc
- Internode: Few projects are looking beyond MPI+X
  - Legion, UPC++, Global Arrays
  - *Most challenging path*

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

Conclusions

- Integration of data, learning, and simulation needs
- Accelerators are the new norm – not just GPUs
- Programming techniques in use on current systems will apply directly to Aurora
- Good software engineering will provide flexibility for future systems