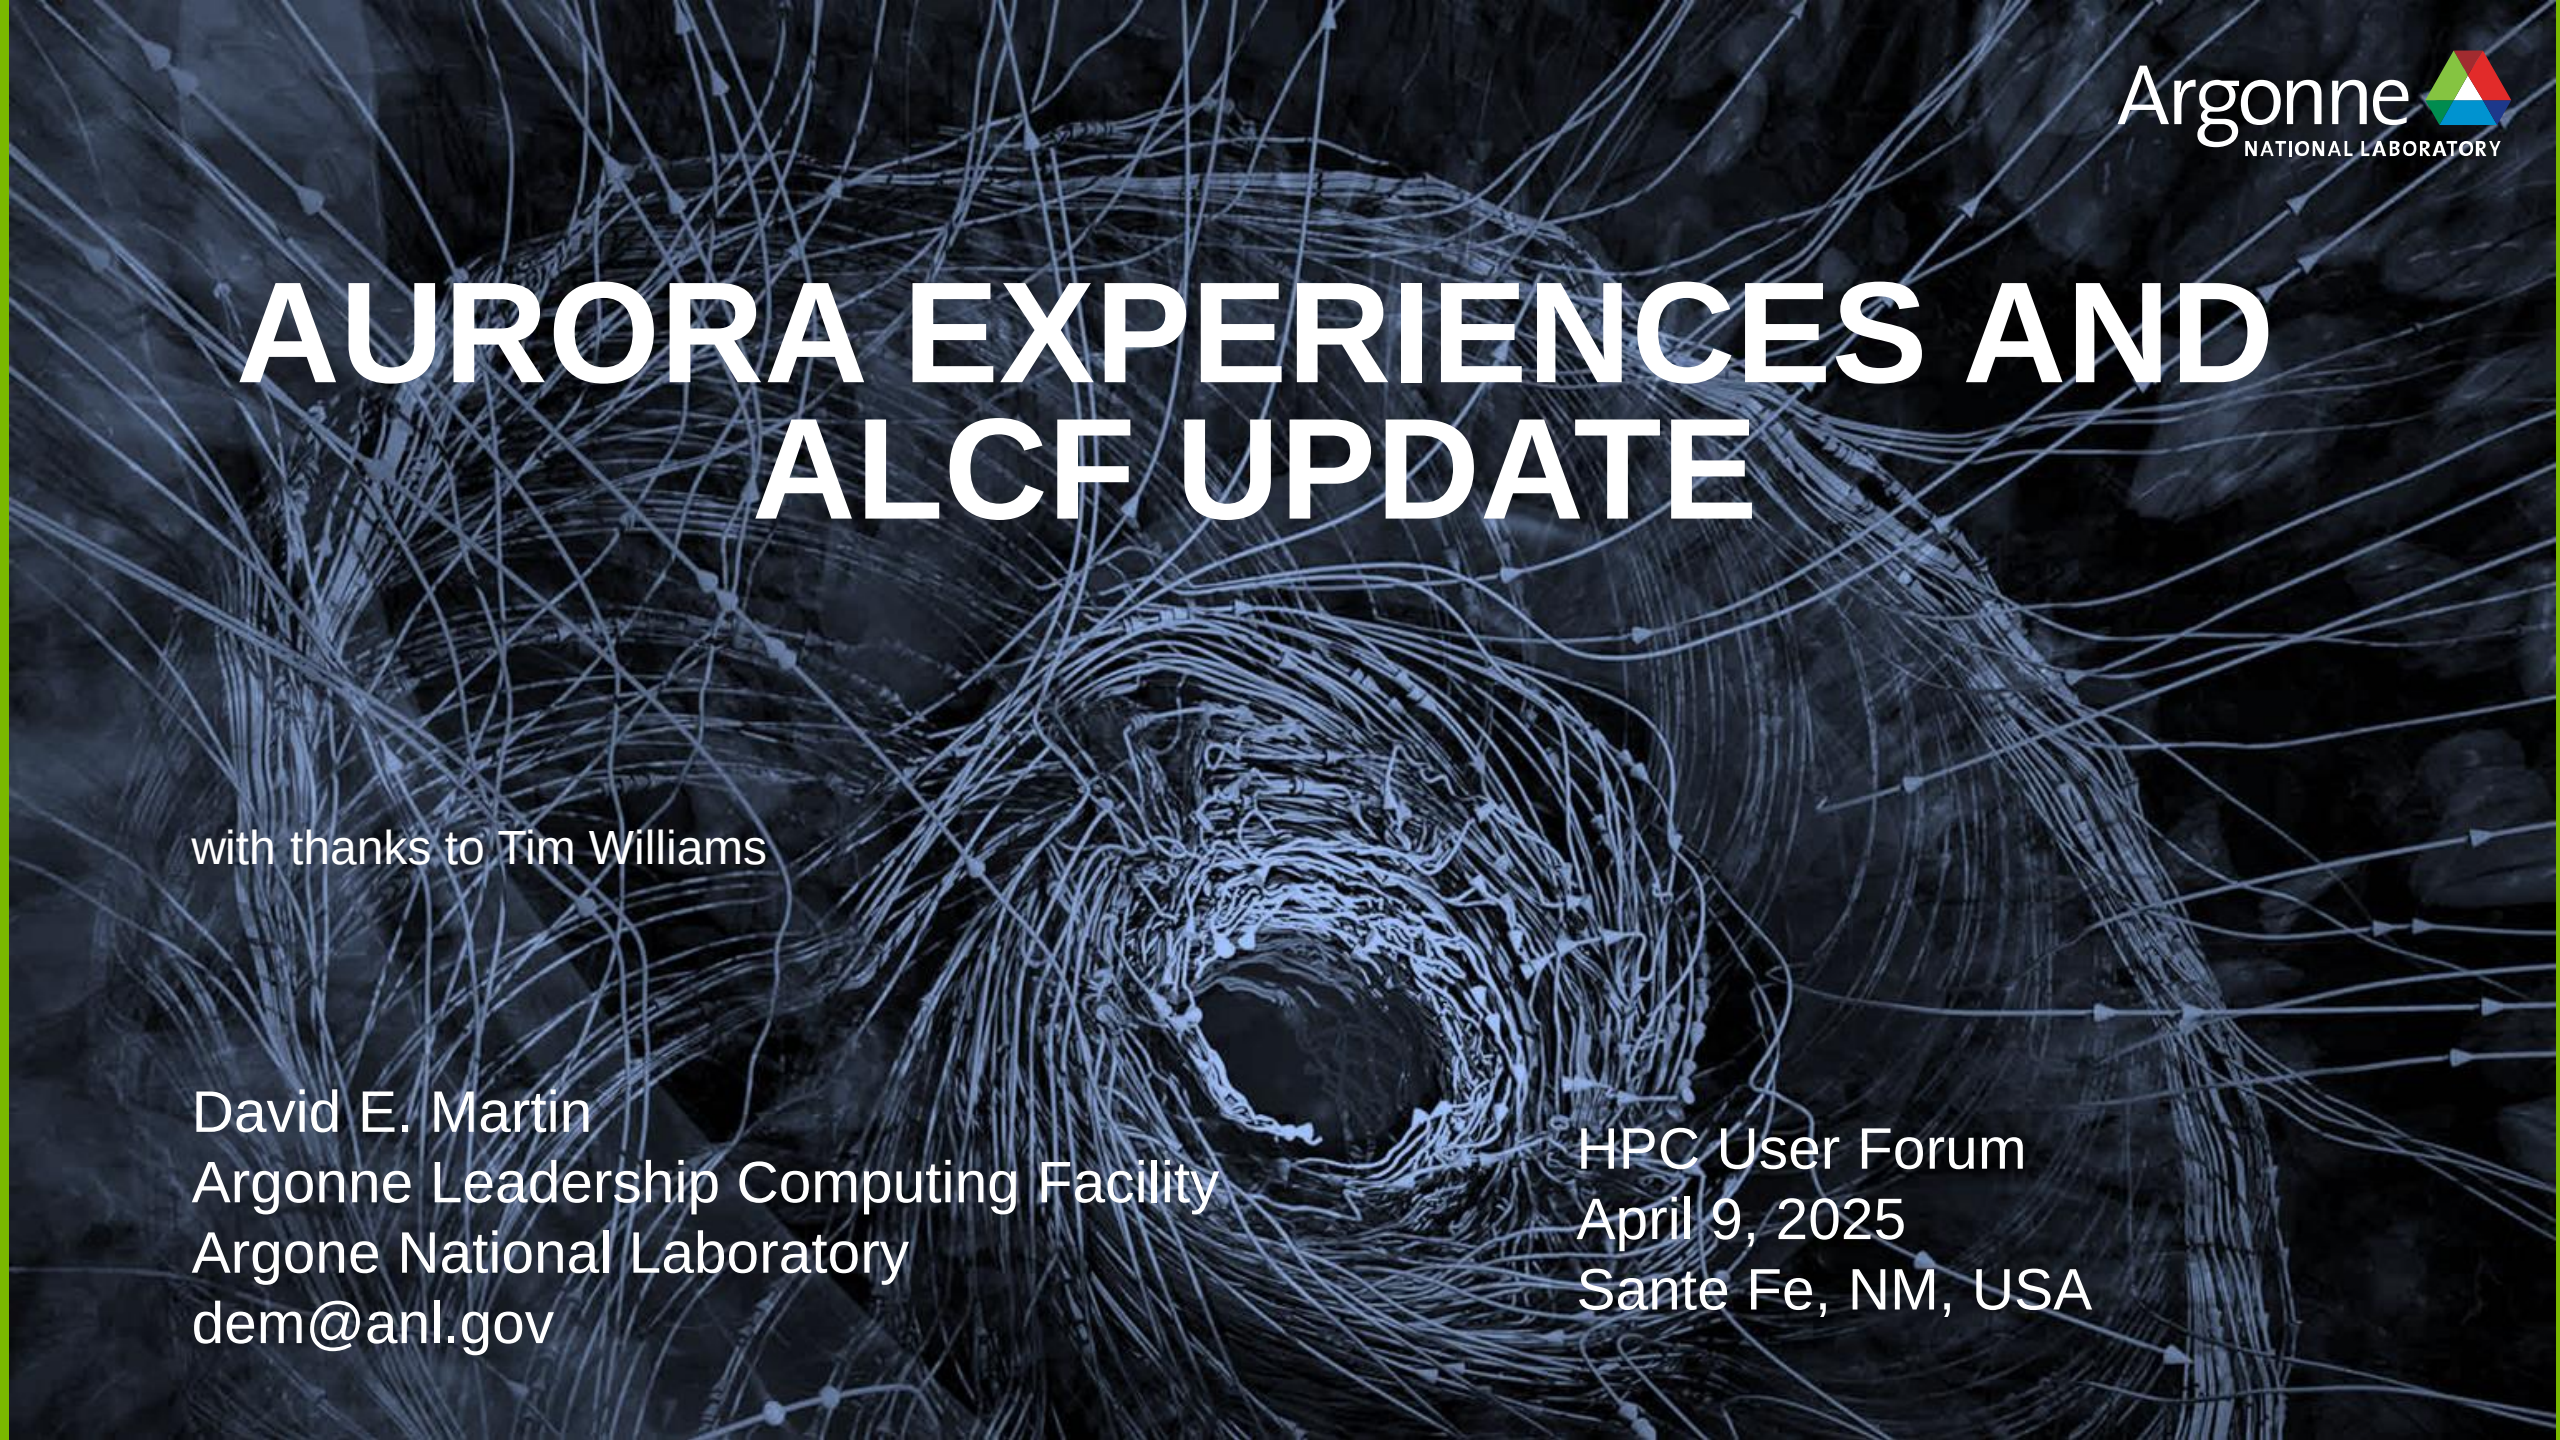

AURORA EXPERIENCES AND ALCF UPDATE

with thanks to Tim Williams

David E. Martin
Argonne Leadership Computing Facility
Argone National Laboratory
dem@anl.gov

HPC User Forum
April 9, 2025
Sante Fe, NM, USA

DOE SC
Advanced Scientific Computing Research User Facilities

The Advanced Scientific Computing Research (ASCR) program leads the nation and the world in supercomputing, high-end computational science, and advanced networking for science.

ALCF and OLCF make up the DOE Leadership Computing Facility

Argonne Leadership Computing Facility (ALCF)

Oak Ridge Leadership Computing Facility (OLCF)

National Energy Research Scientific Computing Center (NERSC)

Energy Sciences Network (ESnet)

High Performance Data Facility (HPDF)

DOE Leadership Computing Facility

- Established in 2004 as a collaborative, multi-lab initiative funded by DOE's ***Advanced Scientific Computing Research*** program
- Operates as **one facility** with two centers, at Argonne and at Oak Ridge National Laboratory
- Deploys and operates at least two advanced architectures that are **10-100 times more powerful** than systems typically available for open scientific research
- **Fully dedicated** to open science to address the ever-growing needs of the scientific community

COMMON RULES FOR NATIONAL SCIENTIFIC USER FACILITIES

- Open to all
  - No restriction on organization, funding source, nationality, or research area
- Access through peer-reviewed proposal process
  - Project must enable breakthrough science
  - Rapid discretionary access available
- Two ways to “pay”
  - Publish significant scientific results
  - Pay cost recovery to keep *everything* proprietary
- 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 for the largest, most complex modeling and simulations
- Maintain storage, networking and software

EXPERT SUPPORT

- Computational scientists that are domain scientists who translate problems to computational representations
- Performance engineers port and optimize 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

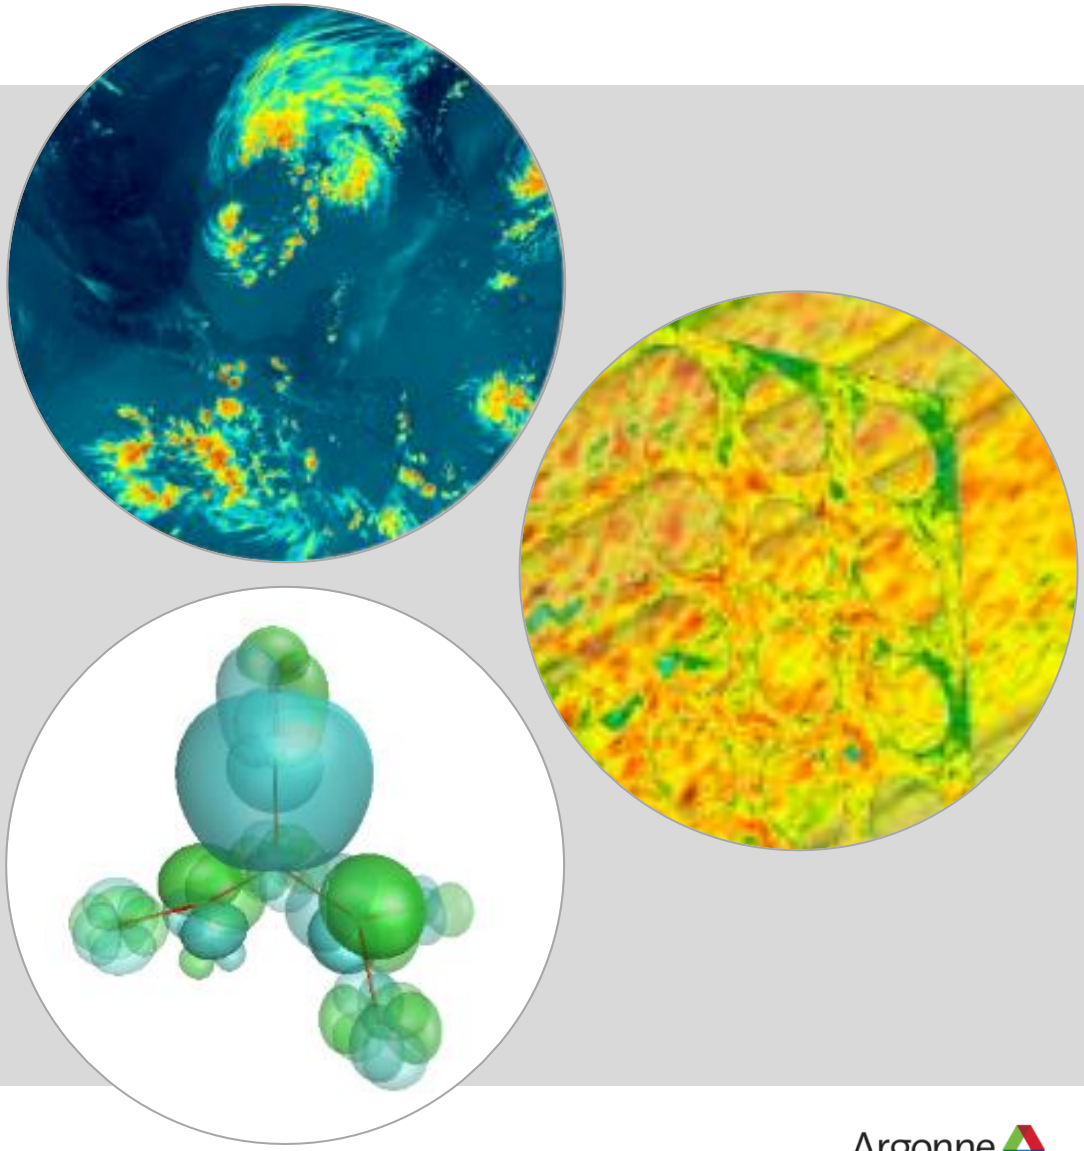

ALCF Allocation Programs

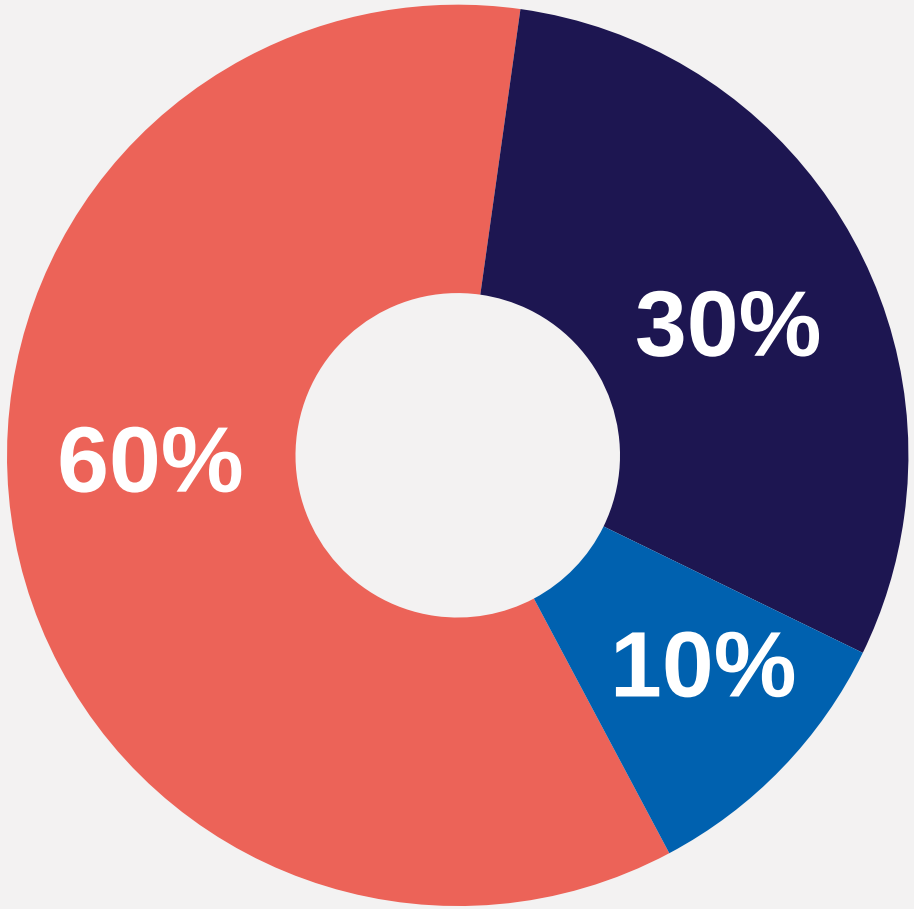

INCITE: Innovative and Novel Computational Impact on Theory and Experiment

- Yearly call with computational readiness and peer reviews
- Open to all domains and user communities

ALCC: ASCR Leadership Computing Challenge

- Yearly call with peer reviews
- Focused on DOE priority

DD: Director's Discretionary Program

- Rapid allocations for project prep and immediate needs
  - Early Science Program (ESP)
  - Exascale Computing Project (ECP)
  - ALCF Data Science Program (ADSP)
  - Proprietary Projects

ALCF SYSTEMS EVOLUTION

5.7 TF

IBM BG/L

2004

557 TF

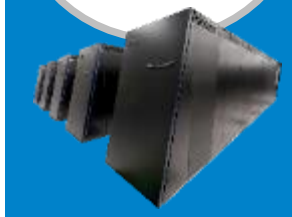

Intrepid
IBM BG/P

2007

10 PF

Mira
IBM BG/Q

2012

11.7 PF

Theta
Intel-Cray XC40

2017

15.6 PF

NVIDIA
DGX A100

2020

35-45 PF

Polaris & Crux
HPE-AMD-NVIDIA

2021

1+ EF

Aurora
Intel-HPE

2024

US Exascale Computers for Science and Engineering

- Frontier
  - AMD CPUs, AMD GPU accelerators
  - Deployed at OLCF (Oak Ridge Leadership Computing Facility)

in production

- Aurora
  - Intel CPUs, Intel GPU accelerators
  - Deployed at ALCF (Argonne Leadership Computing Facility)

in production

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

Open to researchers from academia, government laboratories, and industry, the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program is the major means by which the scientific community gains access to the nation's fastest supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering.

- El Capitan
  - AMD MI300a APUs – truly unified memory space
  - Deployed at Lawrence Livermore National Laboratory
  - For use in US DOE NNSA-ASC (National Nuclear Security Administration – Advanced Simulation and Computing)

operational

US Exascale Computers for Science and Engineering

- Frontier

  #2 in Top500 – HPL 1.353 exaFLOPS
  #2 in Top500 – HPL-MxP 11.4 exaFLOPS

in production

- Aurora

  #3 in Top500 w/9234 nodes – HPL 1.012 exaFLOPS
  #1 in Top500 w/9500 nodes – HPL-MxP 11.6 exaFLOPS

in production

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

Open to researchers from academia, government laboratories, and industry, the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program is the major means by which the scientific community gains access to the nation's fastest supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering.

- El Capitan

  #1 in Top500 – HPL 1.742 exaFLOPS

operational

Aurora Node

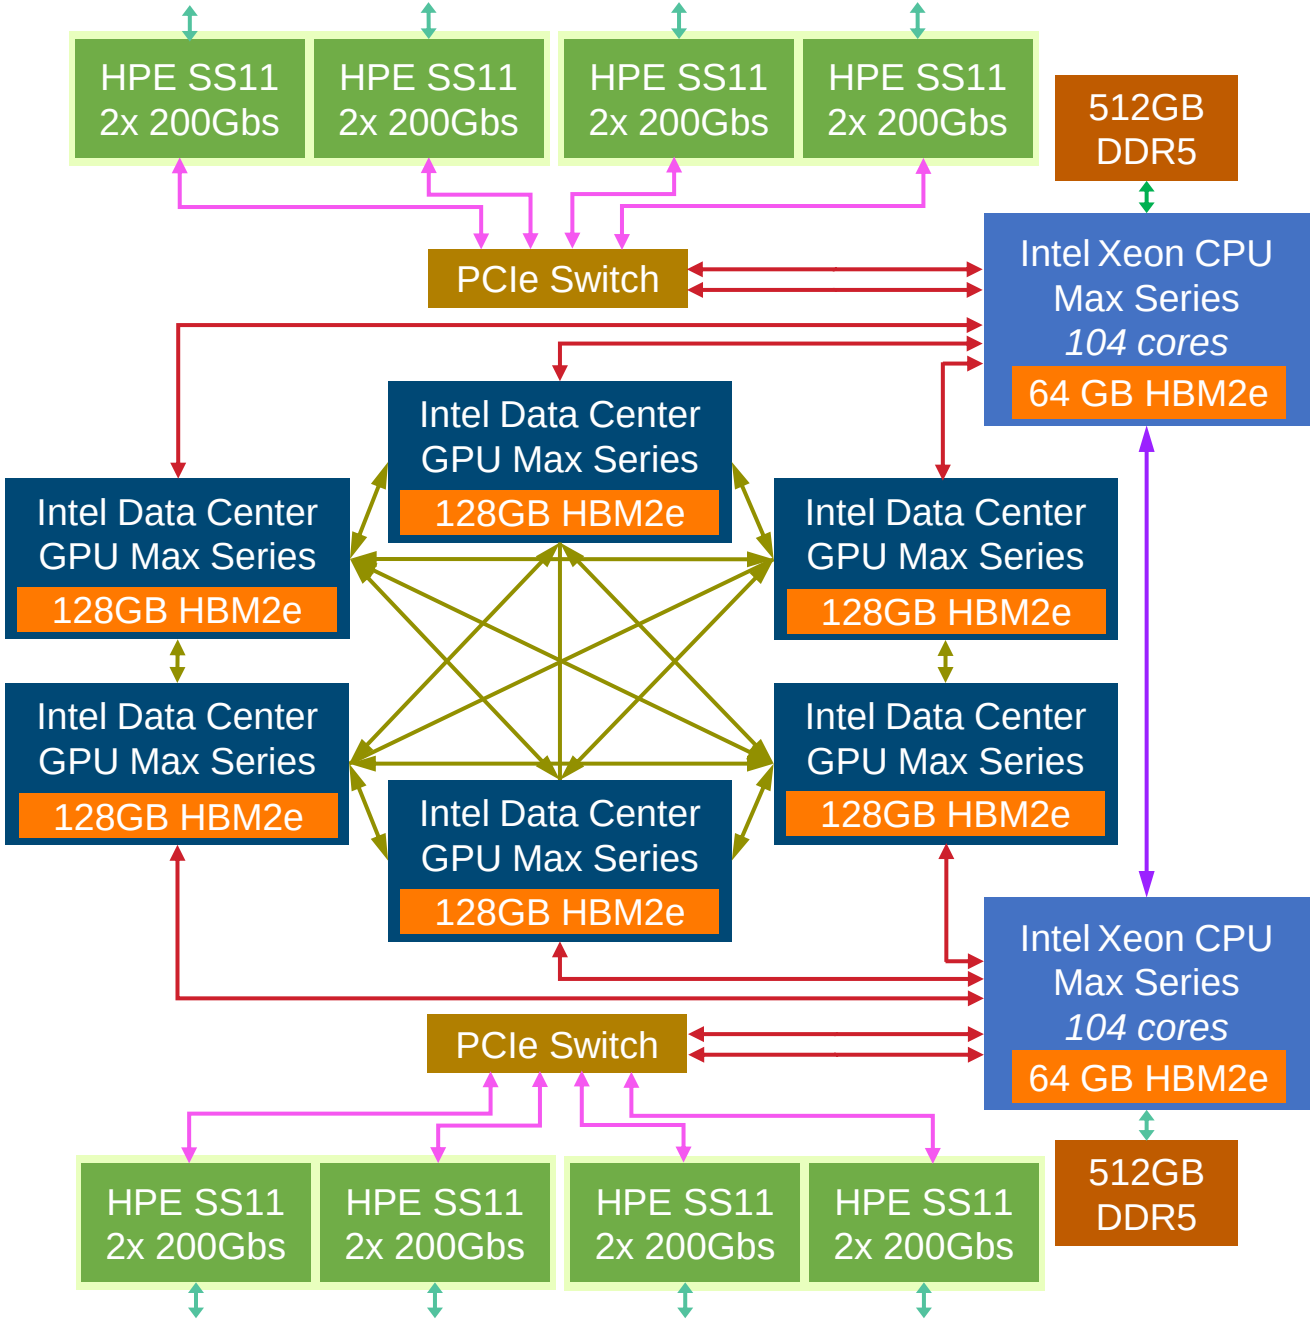

Frontier Node

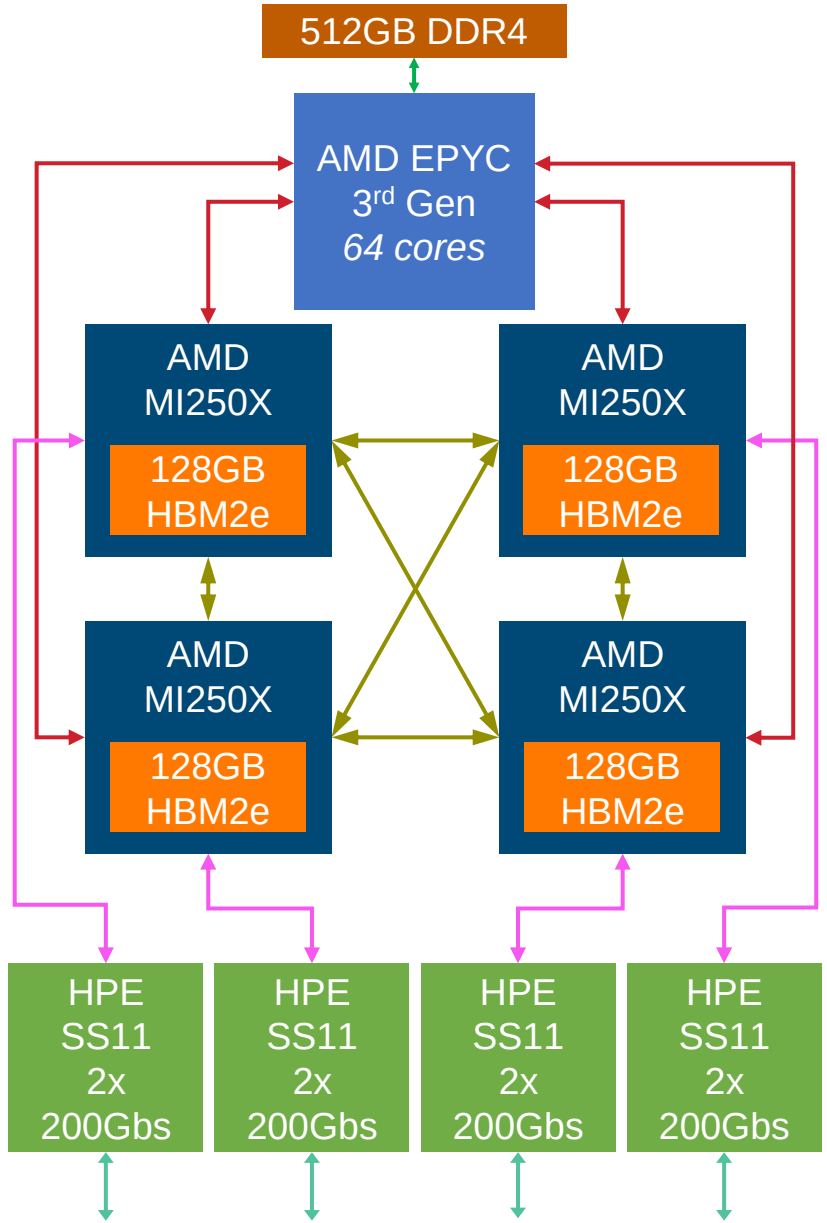

Aurora

Intel® Data Center GPU Max Series

4th Gen Intel XEON Max Series CPU *with High Bandwidth Memory*

Platform
HPE Cray-Ex

Racks - 166
Nodes - 10,624
CPUs - 21,248
GPUs – 63,744

Interconnect
HPE Slingshot 11
Dragonfly topology with adaptive routing
Cassini NIC, 200 Gb/s (25 GB/s), 8 per node
Network Switch:
25.6 Tb/s per switch (64 200 Gb/s ports)
Links with 25 GB/s per direction

Peak FP64 Performance
≧ 2 exaFLOPS

Memory
10.9PiB of DDR @ 5.95 PB/s
1.36PiB of CPU HBM @ 30.5 PB/s
8.16PiB of GPU HBM @ 208.9 PB/s

Network
2.12 PB/s Peak Injection BW
0.69 PB/s Peak Bisection BW

Storage
230PB DAOS Capacity
31 TB/s DAOS Bandwidth

Frontier

AMD® Instinct MI250X GPU

AMD® EPYC "Trento" CPU

Platform
HPE Cray-EX

Racks – 74
Nodes – 9,408
CPUs – 9,408
GPUs – 37,632

Interconnect
HPE Slingshot 11
Dragonfly topology with adaptive routing
Cassini NIC, 200 Gb/s (25 GB/s), 4 per node

Peak FP64 Performance
2.0 exaFLOPS

Memory
4.6PiB of DDR @ 1.93 PB/s
4.6PiB of GPU HBM @ 123.0 PB/s

Network
0.94 PB/s Peak Injection BW

Storage
695PB Lustre Capacity (11PB flash tier)
9.4 TB/s Lustre Bandwidth
36 PB on-node storage 75/38 TB/s R/W BW

Pre-Exascale Systems

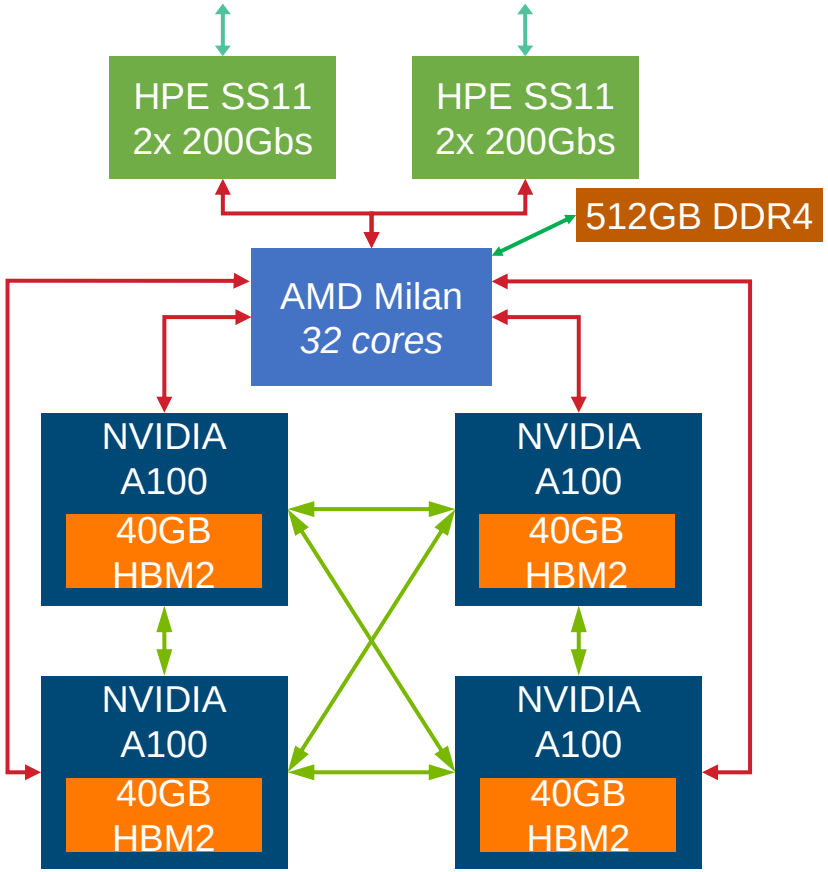

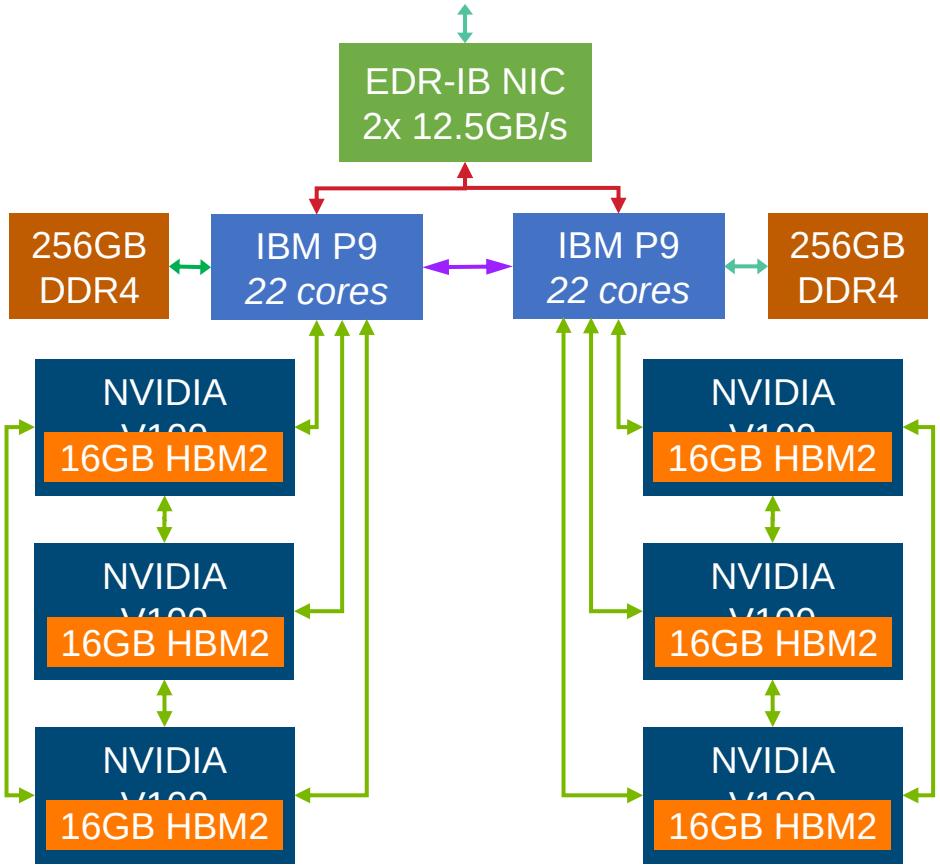

Pre-Exascale Systems

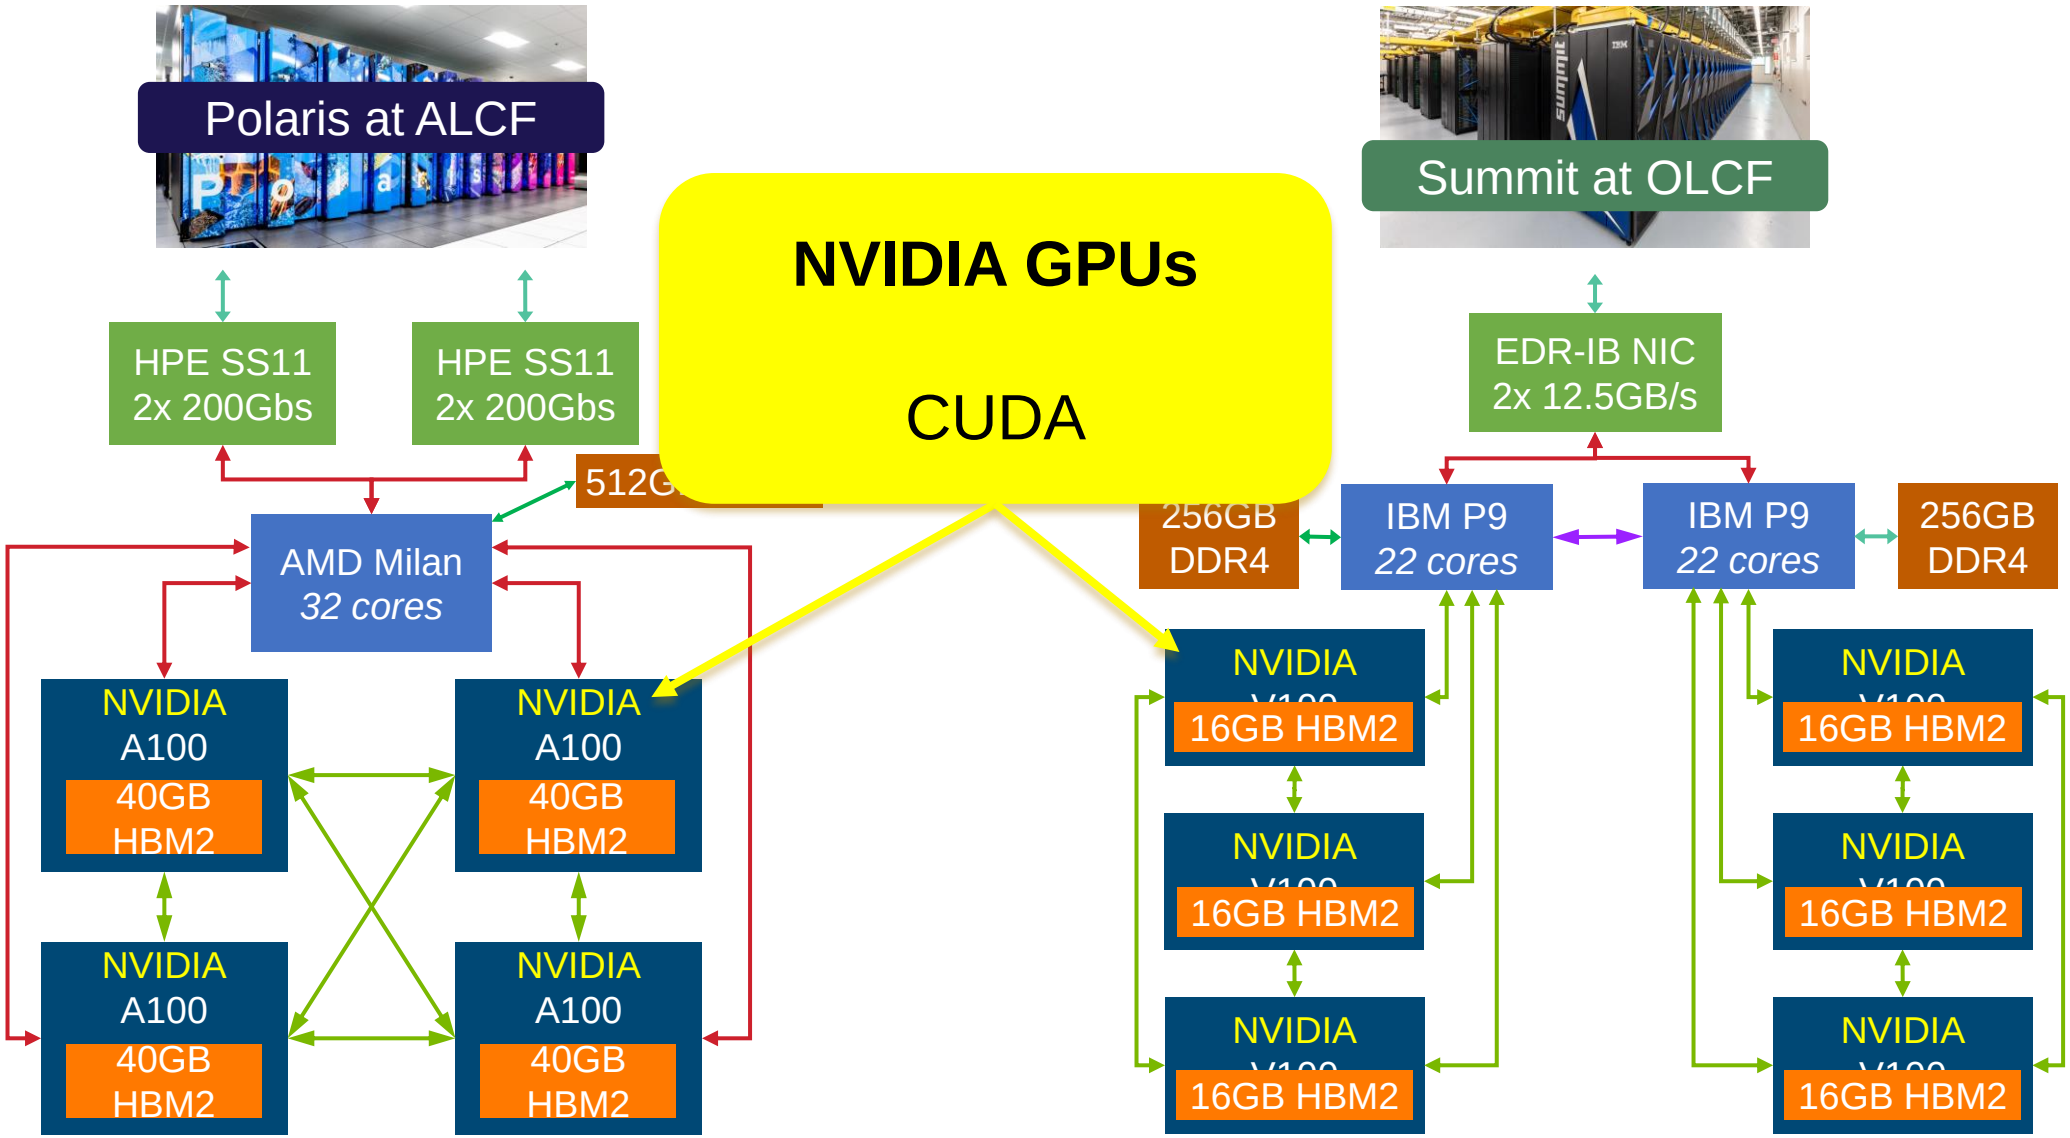

Exascale Applications/Software Readiness

- Exascale Computing Project (ECP) 2016-2023
  - 7 years, $1.8B, 2800 people
    - 15 (of 21) Applications Development projects engaged with ALCF toward Aurora
    - ALCF POC for each, via ECP Hardware & Integration funding
- Aurora Early Science Program (ESP)
  - 19 project teams with over 200 multi-institutional team members
    - 16 ALCF postdocs
    - 19 ALCF Catalysts/POCs (10% effort for ~7 years)
- Intel-Argonne Center of Excellence
    - 5 Intel staff
- Frontier Center for Advanced Application Readiness (CAAR)
  - 8 projects + 12 ECP AD projects

First-Year Science Projects on Aurora

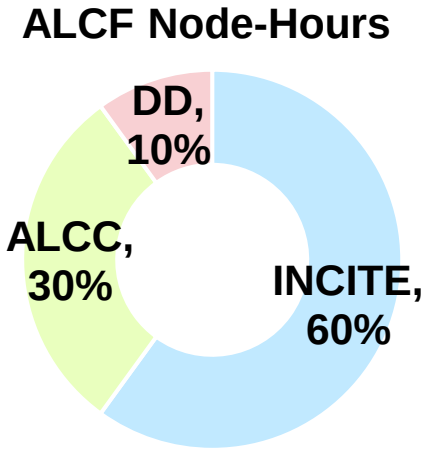

- Aurora Early Science Program (ESP)
  - 19 projects
  - 19.8 million node-hours

AURORA EARLY SCIENCE PROGRAM

- INCITE
  - 30 projects (+ 4 porting/optimization projects)
  - 22.6 million node-hours

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

Open to researchers from academia, government laboratories, and industry, the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program is the major means by which the scientific community gains access to the nation's fastest supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering.

- ALCC (ASCR Leadership Computing challenge)
  - 13 projects
  - 5 million node-hours

U.S. DEPARTMENT OF ENERGY | Office of Science

Aurora Early Science Program

Science campaigns start 1/27 with workflows incorporating the 3 pillars

Simulation

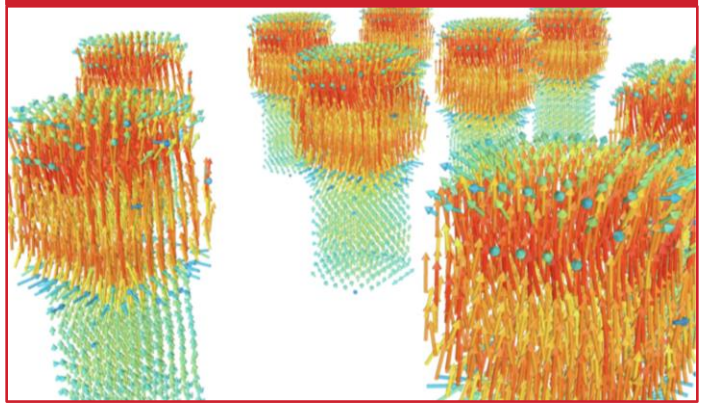

- Reactive, quantum, and classical MD for materials and biophysics
- ITER tungsten impurity ions kinetics—PIC
- Cosmological P^3M + SPH hydrodynamics
- CCSD(T) quantum chemistry
- Quantum Monte Carlo
- FEM CFD, CFD+combustion

Data

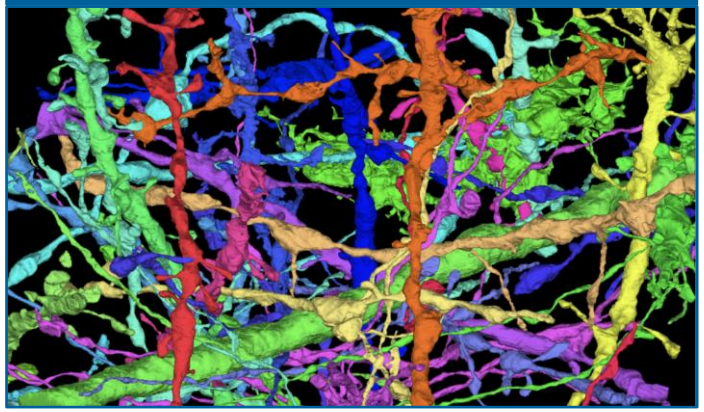

- High-throughput electron microscopy data
- In situ data from blood flow and aerodynamics simulation
- Atlas detector event Monte Carlo analysis
- Cosmological sky survey and P^3M simulation data
- Chemistry kinetics database
- Tokamak experiment data

Learning

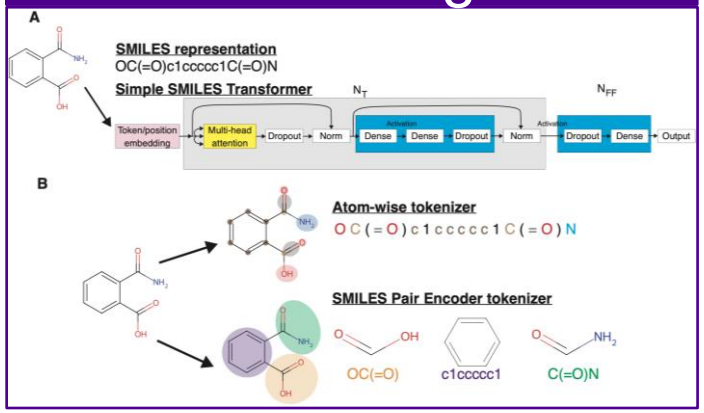

- Simple SMILES transformer for drug design
- Generative model for lattice QCD
- Material band gap predictor
- Tokamak disruption predictor
- Flood fill model for connectomics
- MLP turbulence closure model

Computational Science Methods in Y1 Aurora Projects + ECP

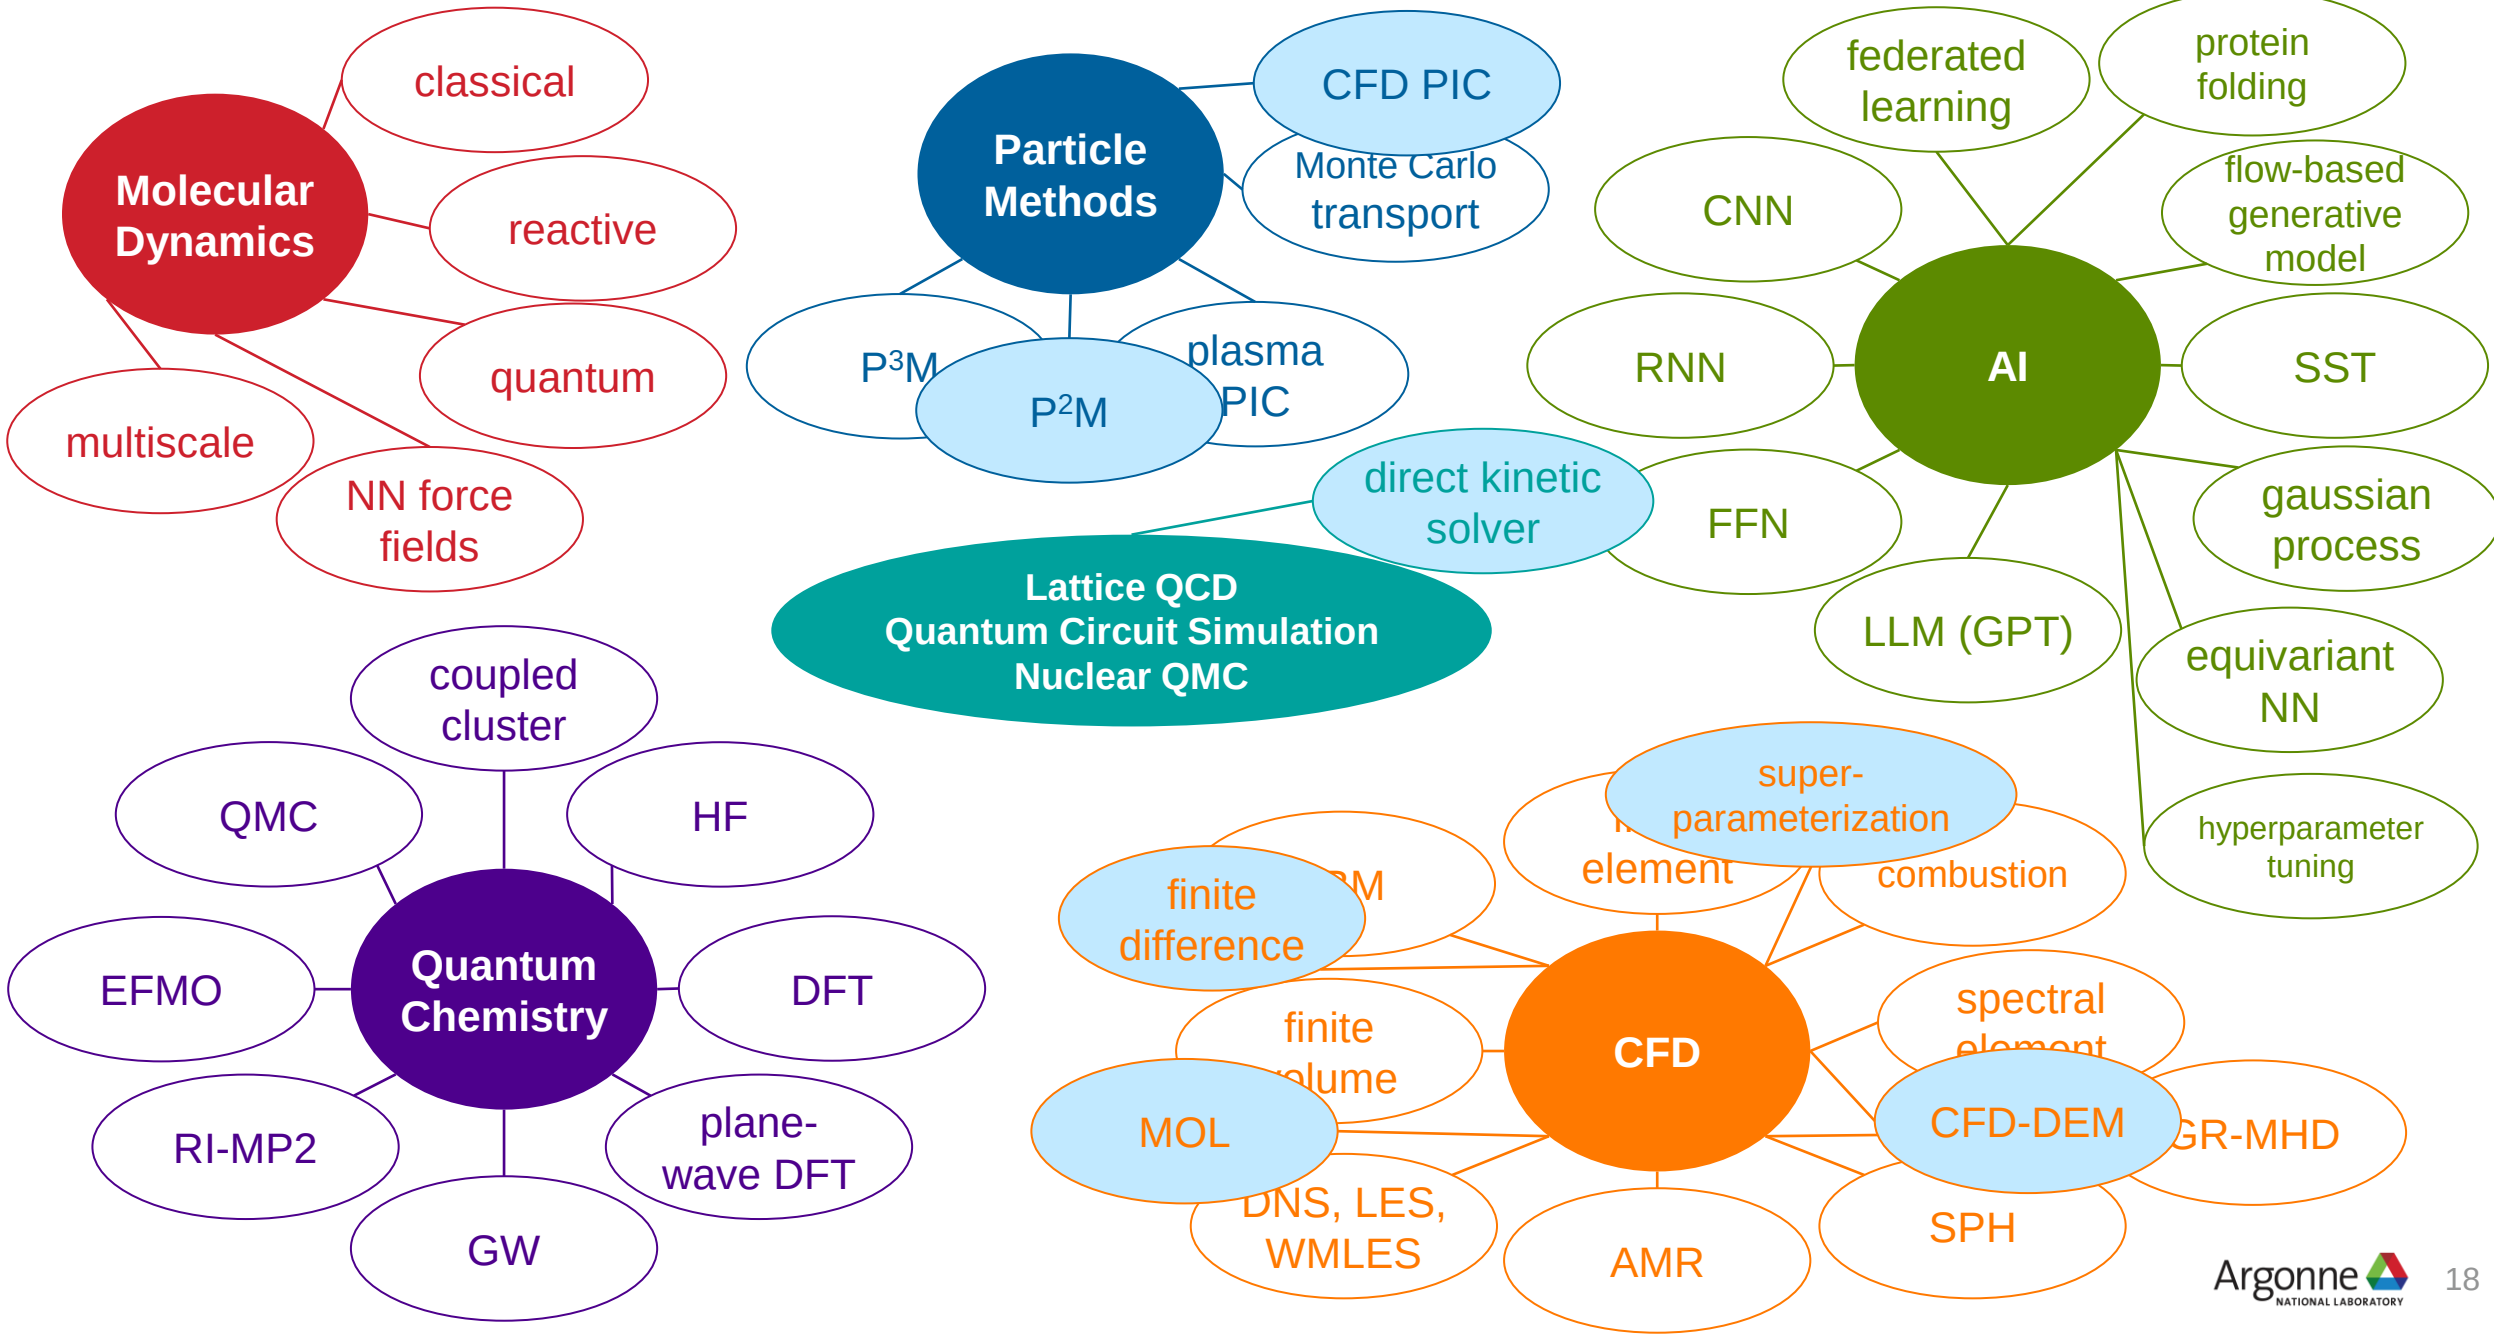

Science Domains in Y1 Aurora Projects + ECP

Aerodynamics
Combustion
Thermal hydraulics for fusion & fission
Turbulent Rayleigh-Benard convection
Buoyancy-driven turbulence
Wind energy

Core collapse supernovae
Binary neutron star mergers
Core collapse supernovae
Black hole accretion
Neutron star remnants

MOF design
Carbon capture
Transition metal chemistry
Heterogeneous catalysis
Heterogeneous reactions
Biofuels

Membrane transport proteins
Predictive molecular epdimiology
Blood flow & transport
Cancer drug design
Heteropolymer design
Brain connectomics
Pathogenesis

Fusion radiation transport
Fusion plasma kinetics
Tokamak control

Functional quantum materials
Fusion device materials
Multiscale materials modeling
Quantum materials
Materials science

Neuroscience foundation model
Science foundation model
Federated learning
Quantum computing
AI assistant for simulation

QCD
Cosmology
Nuclear structure and reactions
Collider data analysis

Climate modeling
Runoff inundation
Seismic wave propagation

Exascale Programming Models – Data & Learning

- AI / data science community expects portability via Python ecosystem
  - Python – Numba, NumPy, etc.
  - Deep learning frameworks
    - PyTorch, TensorFlow, JAX, DeepSpeed, Horovod
- Hardware vendors optimize performance (NVIDIA, Intel, AMD)
  - GPU offload
  - communication for distributed training & inference

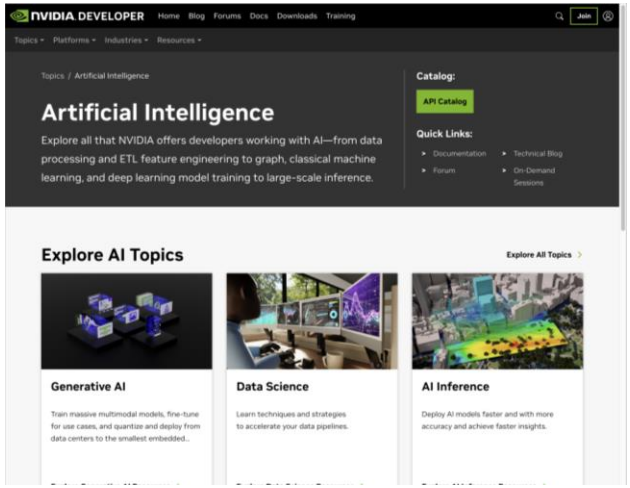

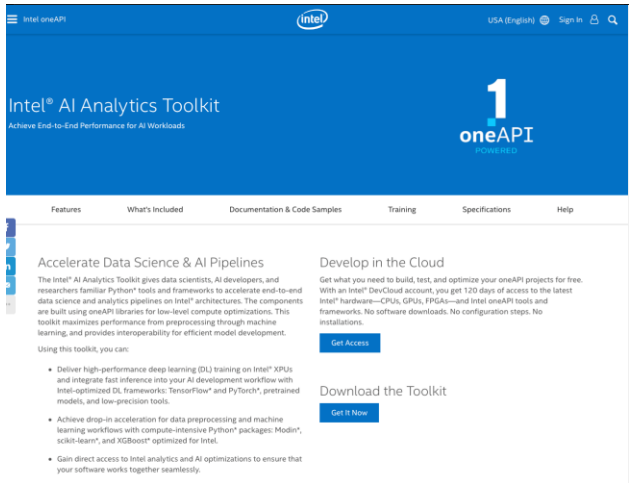

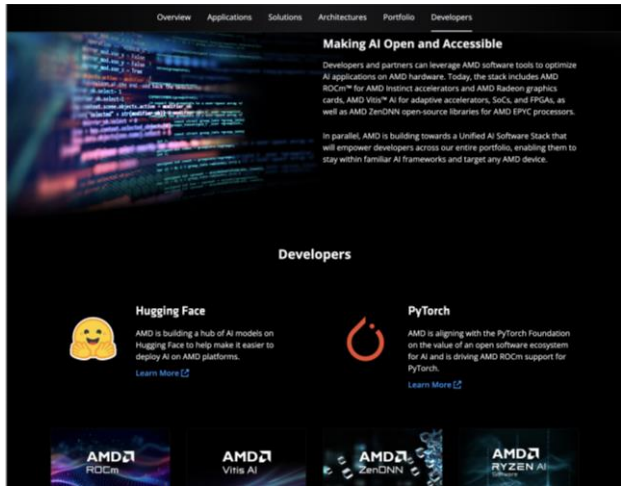

Exascale Programming Models – Simulation

- CPU: portability via OpenMP threads and SIMD
- Communication: portability via MPI
- GPU: portability via OpenMP `target`
- GPU: other offload kernel programming:

NVIDIA	Intel	AMD
CUDA	CUDA (chipStar, Syclomatic)	CUDA (HIPIFY)
HIP	HIP (chipStar)	**HIP**
SYCL	**SYCL**	SYCL
OpenCL	OpenCL	OpenCL

Portability layers/libraries

- Kokkos, Raja, libCEED, OCCA, AMReX, PETSc
- DIY

CUDA	SYCL	HIP

Exascale Programming Models – Simulation

- CPU: portability via OpenMP threads and SIMD
- Communication: portability via MPI
- GPU: portability
- GPU: other offlo

Starting from CPU-only code?

- *Congratulations!* You get to do the hard part.
- Refactor/rewrite code kernels to run efficiently on GPU
  - drastically limit branching
  - coalesced memory accesses
    - modify data structures
  - fuse/break kernels
  - CPU vectorization ≠ GPU SIMT/SIMD
  - *...and more!*
- Use any GPU kernel programming approach*
- Portability layers (Kokkos...) span CPU and GPU

NVIDI

- **CUDA**
- HIP
- SYCL
- OpenCL

AMD

PIFY)

- Kokkos, Raja, li
- DIY

CUDA | **SYCL** | **HIP**

Programming Model Choices by Y1 Aurora Projects + ECP

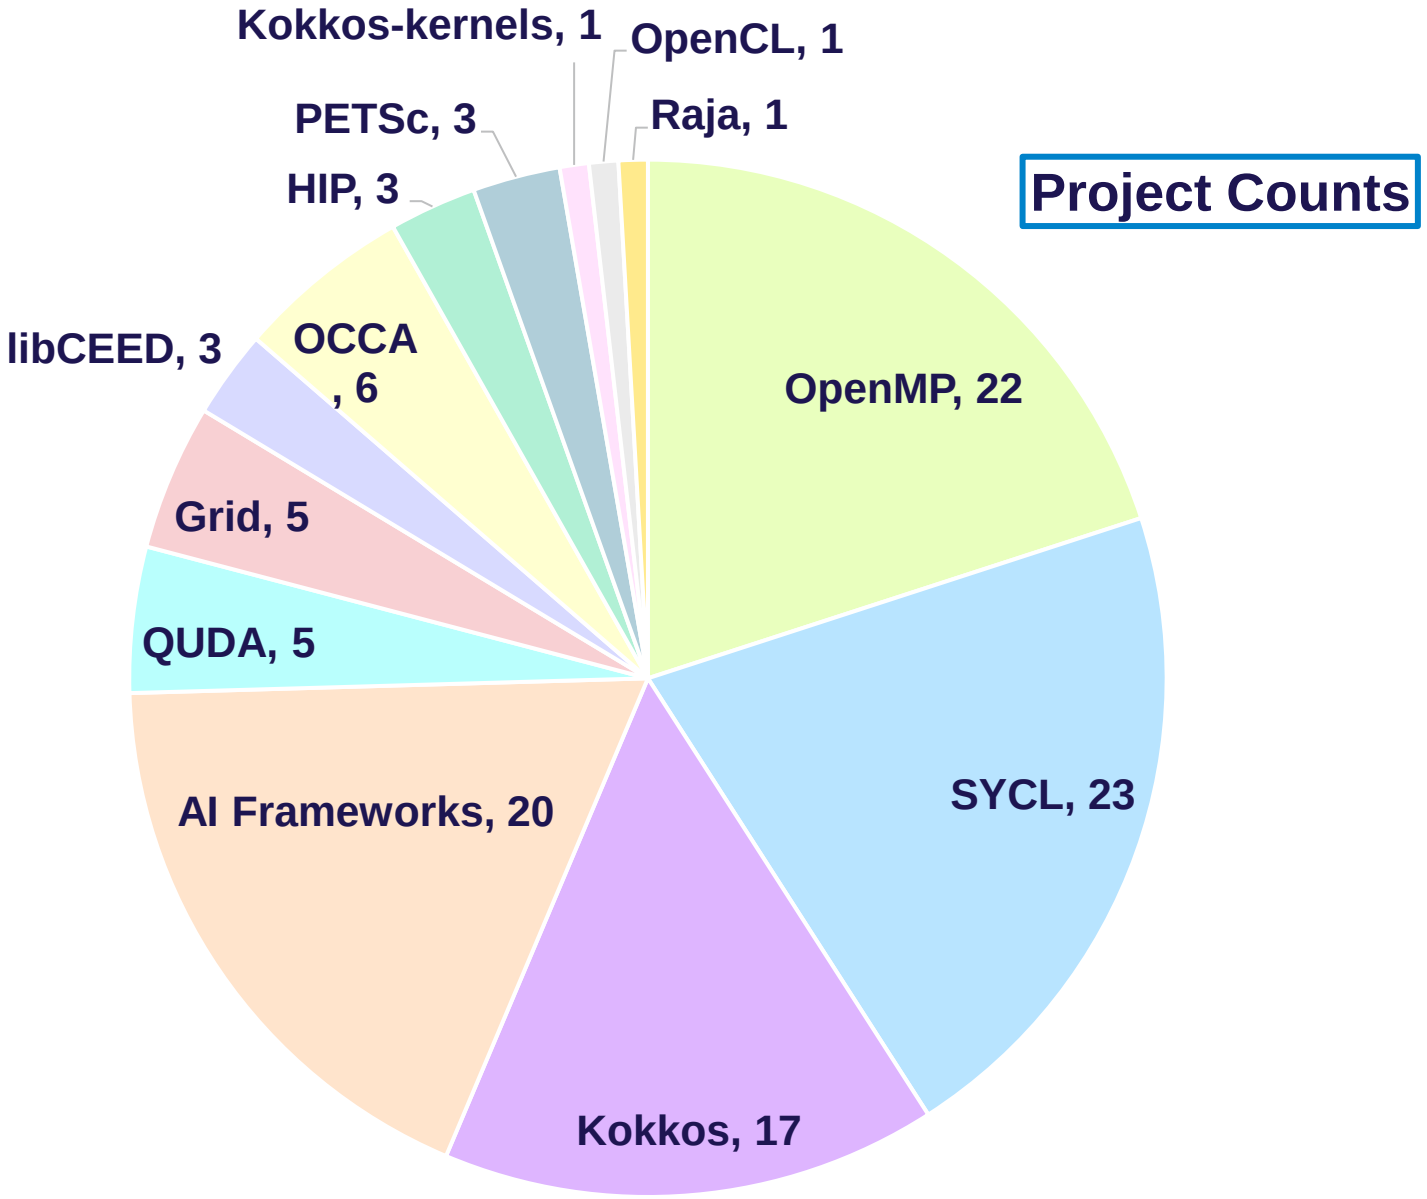

Aurora Performance and Expectations

- DOE Exascale Computing Project (ECP)
  - Achieve 50X improvement in figure-of-merit (FOM) w.r.t. pre-ECP baseline (2017 systems)
  - Aurora: 3 projects measured this:
    - EXAALT – 89X on 1024 nodes
    - ExaSMR – 84X on 512 nodes
    - ExaSky – 277X on 4096 nodes

- Aurora Early Science Program (ESP)
  - Demonstrate INCITE-level computational readiness
    1. Scale efficiently to 20% of full system
    2. Make use of GPUs
    3. Ready to run science in 1-3 months
  - Internal review: 16/19 projects met goal in Fall 2024

Aurora Performance and Expectations

- Argonne Applications Working Group performance tracked single-GPU performance figures of merit across >25 ECP and ESP applications for several years
  - Expectations
    - Comparable performance to AMD MI-250X (Frontier GPU)
    - Better performance than NVIDIA A100 (Polaris GPU)

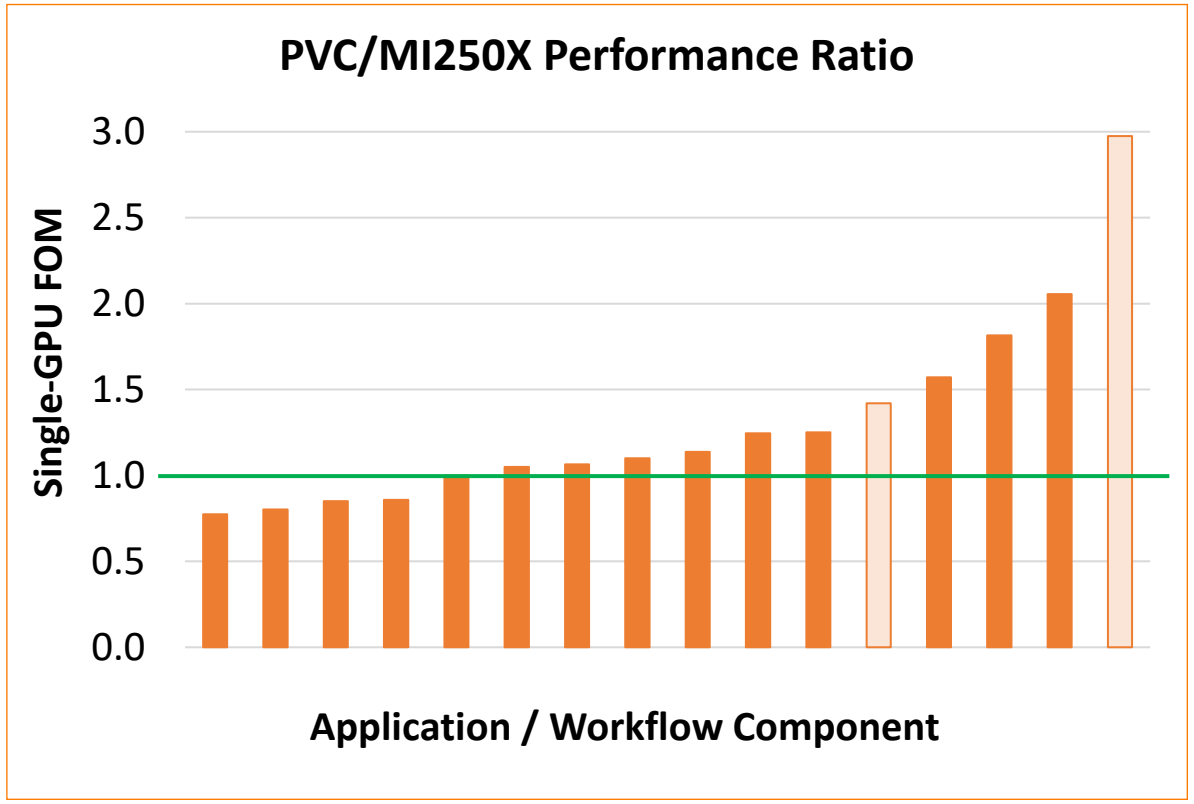

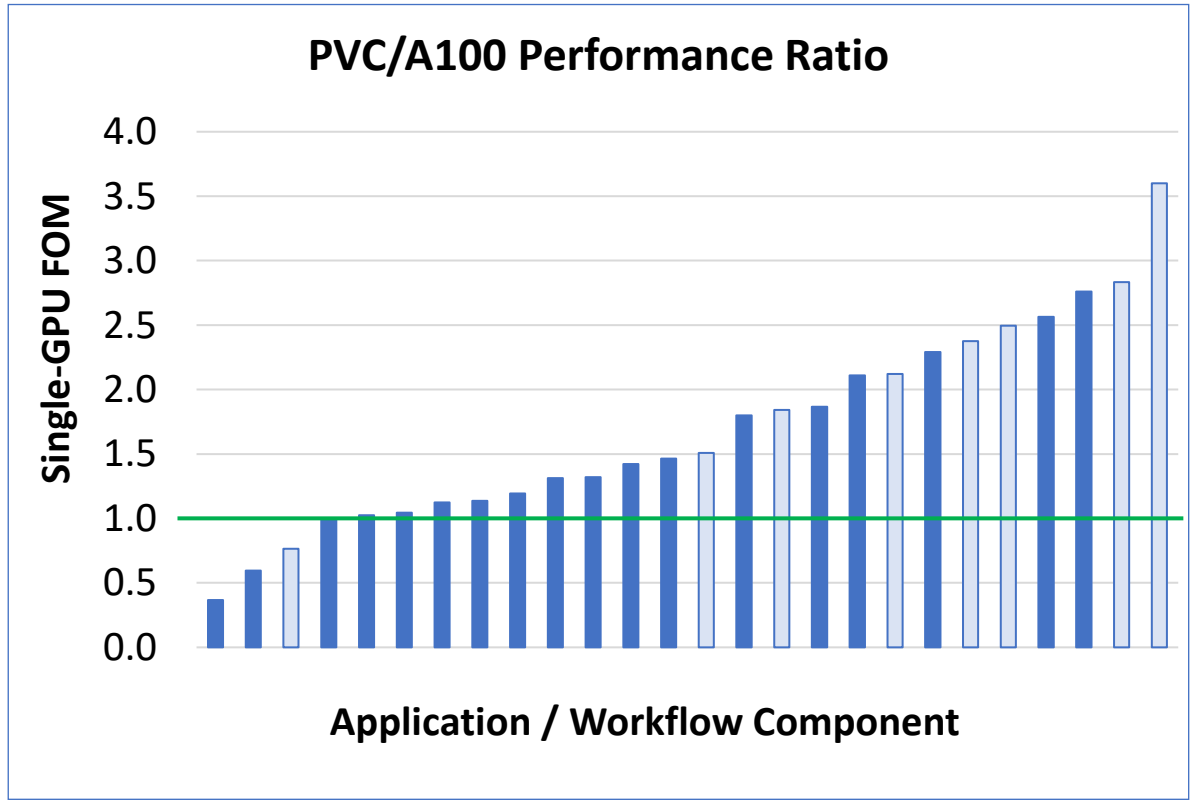

ALCF AI Testbed

Advancing science with HPC and AI

- ALCF AI pathfinding effort provides insights on cutting-edge AI technology and how it improves science outcomes
- Evaluates the usability and performance of machine learning-based applications running on these accelerators
  - a deep learning accelerator, reconfigurable dataflow units, intelligent processing unit- (IPU) based systems
- Ongoing work is guiding the facility toward a future marked by extreme heterogeneity in the compute: CPUs, GPUs, AI, and other accelerators

AI testbeds include:

- SambaNova DataScale
- GraphCore MK1
- Groq
- Cerebras CS-2
- Habana Gaudi

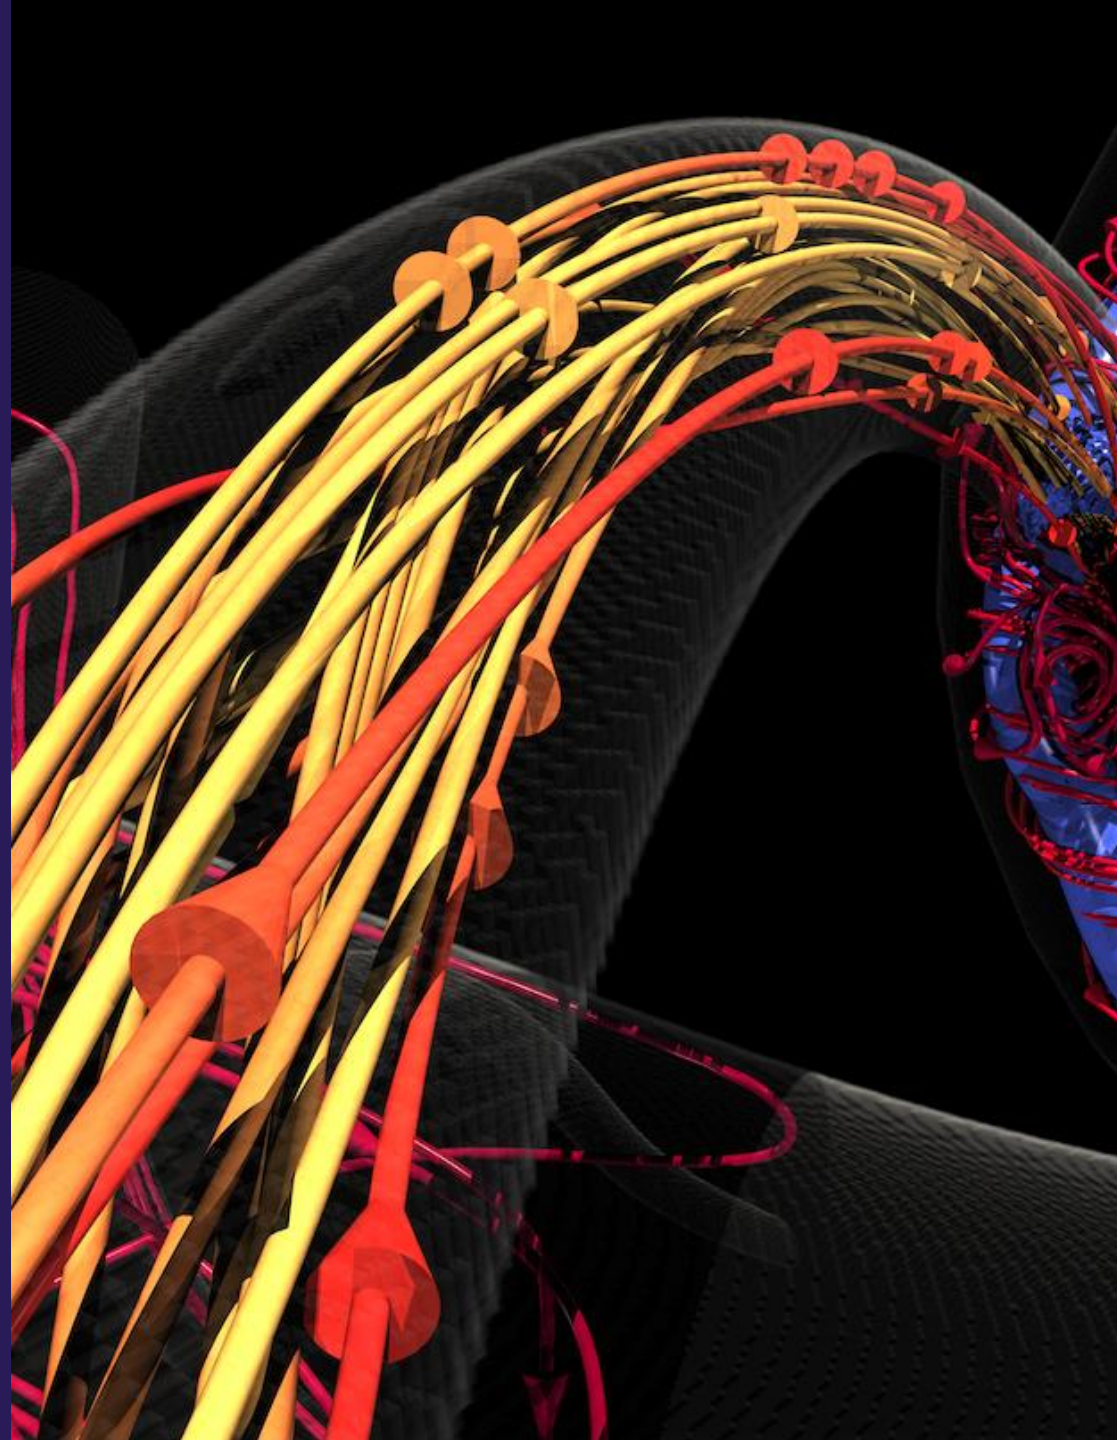

ALCF AI Testbed

Cutting-edge AI accelerators for science

GroqRack (Available for Allocation Requests)

GroqRack Inference

System Size: 72 Accelerators (9 nodes x 8 Accelerators per node)

Compute Units per Accelerator: 5120 vector ALUs

Performance of a single accelerator (TFlops): >188 (FP16) >750 (INT8)

Software Stack Support: GroqWare SDK, ONNX

Interconnect: RealScale TM

Cerebras CS-2 (Available for Allocation Requests)

Cerebras CS-2 Wafer-Scale Cluster WSE-2

System Size: 2 Nodes (each with a Wafer scale engine) including Memory-X and Swarm-X

Compute Units per Accelerator: 850,000 Cores

Performance of a single accelerator (TFlops): >5780 (FP16)

Software Stack Support: Cerebras SDK, Tensorflow, Pytorch

Interconnect: Ethernet-based

SambaNova Dataflow (Available for Allocation Requests)

SambaNova DataScale SN30

System Size: 64 Accelerators (8 nodes and 8 accelerators per node)

Compute Units per Accelerator: 1280 Programmable compute units

Performance of a single accelerator (TFlops): >660 (BF16)

Software Stack Support: SambaFlow, Pytorch

Interconnect: Ethernet-based

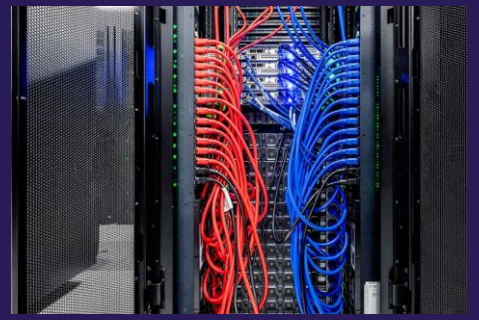

Graphcore Bow Pod64 (Available for Allocation Requests)

Graphcore Intelligent Processing Unit (IPU)

System Size: 64 Accelerators (4 nodes x 16 Accelerators per node)

Compute Units per Accelerator: 1472 independent processing units

Performance of a single accelerator (TFlops): >250 (FP16)

Software Stack Support: PopArt, Tensorflow, Pytorch, ONNX

Interconnect: IPU Link

Habana Gaudi-1

Habana Gaudi Tensor Processing Cores

System Size: 16 Accelerators (2 nodes x 8 Accelerators per node)

Compute Units per Accelerator: 8 TPC + GEMM engine

Performance of a single accelerator (TFlops): >150 (FP16)

Software Stack Support: Synapse AI, TensorFlow and PyTorch

Interconnect: Ethernet-based

Recent ALCF AI Testbed Updates

ALCF AI Testbed Systems are in production and available for allocations to the research community

https://www.alcf.anl.gov/science/directors-discretionary-allocation-program

SambaNova upgraded to latest 2nd generation SN30 accelerators and scaled to 8 nodes with 64 AI accelerators

SambaNova SN30

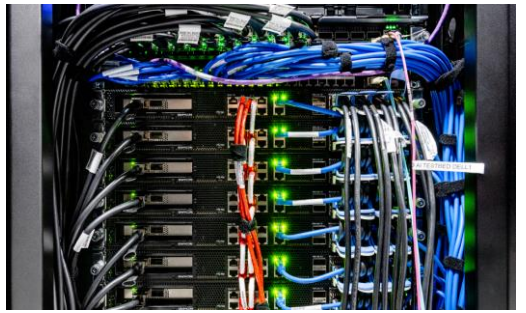

Graphcore upgraded to latest Bow generation accelerators and scaled to a Pod-64 configuration with 64 accelerators

Graphcore BowPod64

Cerebras CS-2 upgraded to an appliance mode to include Memory-X and Swarm-X technologies to enable larger models and scaled to two CS-2 engines

Cerebras CS-2

Groq system has been upgraded to a GroqRack with nine nodes, each consisting of eight GroqChip Tensor streaming processors accelerators

GroqRack

https://nairrpilot.org

Preparing Users for Exascale & AI

- ALCF hosts hundreds of researchers for a multitude of training opportunities for GPU-accelerated systems and AI.
  - **All-hands Workshops** are held for a multiple days for project members
  - **Hackathons** are project deep-dives hosted by ALCF experts and industry leaders
  - **Bootcamps** provide intensive and immersive programs designed to jumpstart the ability of users to run jobs on our systems
  - **Webinars and On-Demand Videos** are open to all researches and available online

Training Activities
ALCF conducted over 30 training activities, reaching more than 500 participants and 25 teams.

TRILLION PARAMETER CONSORTIUM

Global consortium of scientists to address the challenges of building large-scale artificial intelligence (AI) systems and advancing trustworthy and reliable AI for scientific discovery

GOALS

- **Build an open community of researchers** interested in creating state-of-the-art large-scale generative AI models aimed broadly at advancing progress on scientific and engineering problems by sharing methods, approaches, tools, insights, and workflows.
- **Incubate, launch, and coordinate projects** voluntarily to avoid duplication of effort and to maximize the impact of the projects in the broader AI and scientific community.
- **Create a global network of resources and expertise** to facilitate the next generation of AI and bring together researchers interested in developing and using large-scale AI for science and engineering.

TPC WORK AREAS

- **Identifying and preparing high-quality training data**, with teams organized around the unique complexities of various scientific domains and data sources.
- **Designing and evaluating** model architectures, performance, training, and downstream applications.
- **Developing crosscutting and foundational capabilities** such as innovations in model evaluation strategies with respect to bias, trustworthiness, and goal alignment, among others.

TPC will not hold intellectual property, code or money

AuroraGPT

- Large Language Model for Science
- Interactivity of ChatGPT
- Training from science publications and data (graphs, datasets, etc.)
  - Goal of one trillion parameters
  - Training on Aurora, optimized for Intel PVC
- Looking into multi-modal queries and results
- Argonne-led with external collaborators

AI for Science - Concept

Scientific & Engineering Datasets

Mathematics
Biology
Materials
Chemistry
Particle Physics
Nuclear Physics
Computer Science
Climate
Medicine
Cosmology
Fusion Energy
Accelerators
Reactors
Energy Systems
Manufacturing

Text and Code Corpora

General Text
Media
News
Humanities
History
Law
Digital Libraries
OSTI Archive
Scientific Journals
arXiv
Code repositories
Data.gov
PubMed
Agency Archives

DOE and NNSA Exascale Systems
FASST Common AI Software Frameworks
FASST Responsible AI Techniques

Training

Open Science Foundation Model

Training

National Security Foundation Model

Integrated Research Infrastructure
Online Experimental Facilities
Strategic Partnerships

Tuned and Adapted Downstream Models

Exemplar DOE Mission Tasks

Scientific Discovery
Digital Twins
Inverse Design
Code Optimization
Accelerated Simulations
Autonomous Experiments
Secure Data Infrastructure
Co-Design

From Rick Stevens, ANL

U.S. DEPARTMENT OF ENERGY

Integrating Different Types of HPC

ALCF is driving a new paradigm in high performance computing for science and engineering

Modeling & Simulation

Used to study things that are too big, too small, or too dangerous to study in a laboratory setting.

Data Science

Researchers can glean insights from very large datasets produced by experimental, simulation, or observational methods.

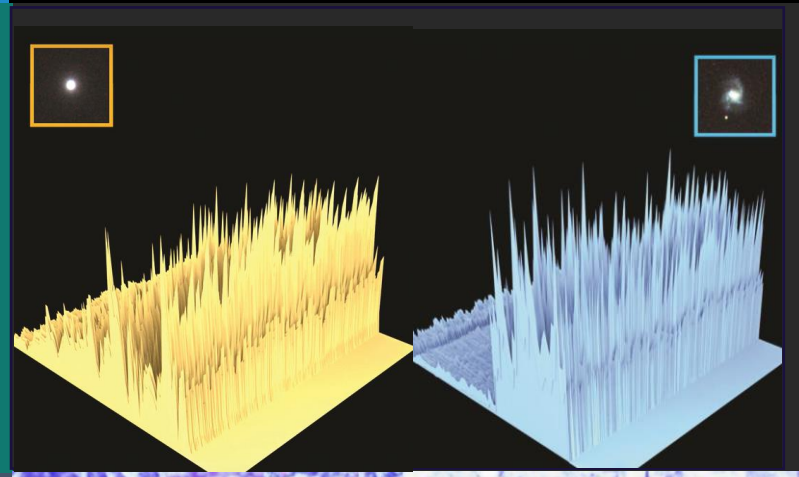

Machine Learning

A type of artificial intelligence that trains computers to discover hidden patterns in data to make novel predictions without being explicitly programmed.

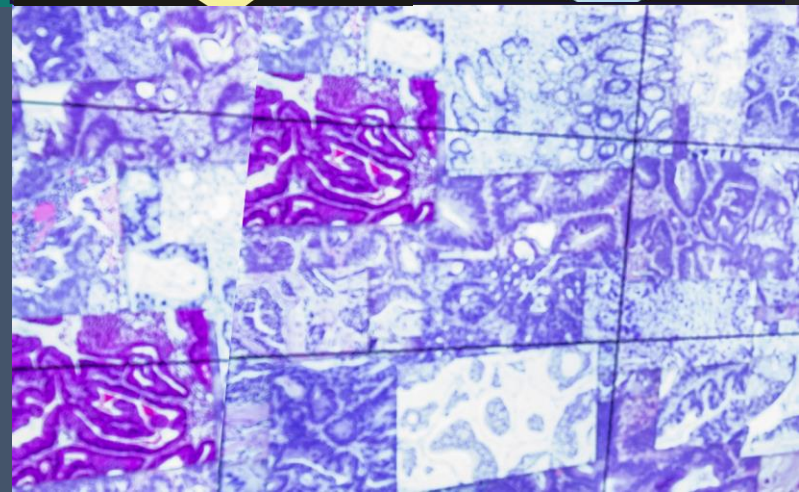

APPROACH TO INDUSTRY COLLABORATIONS

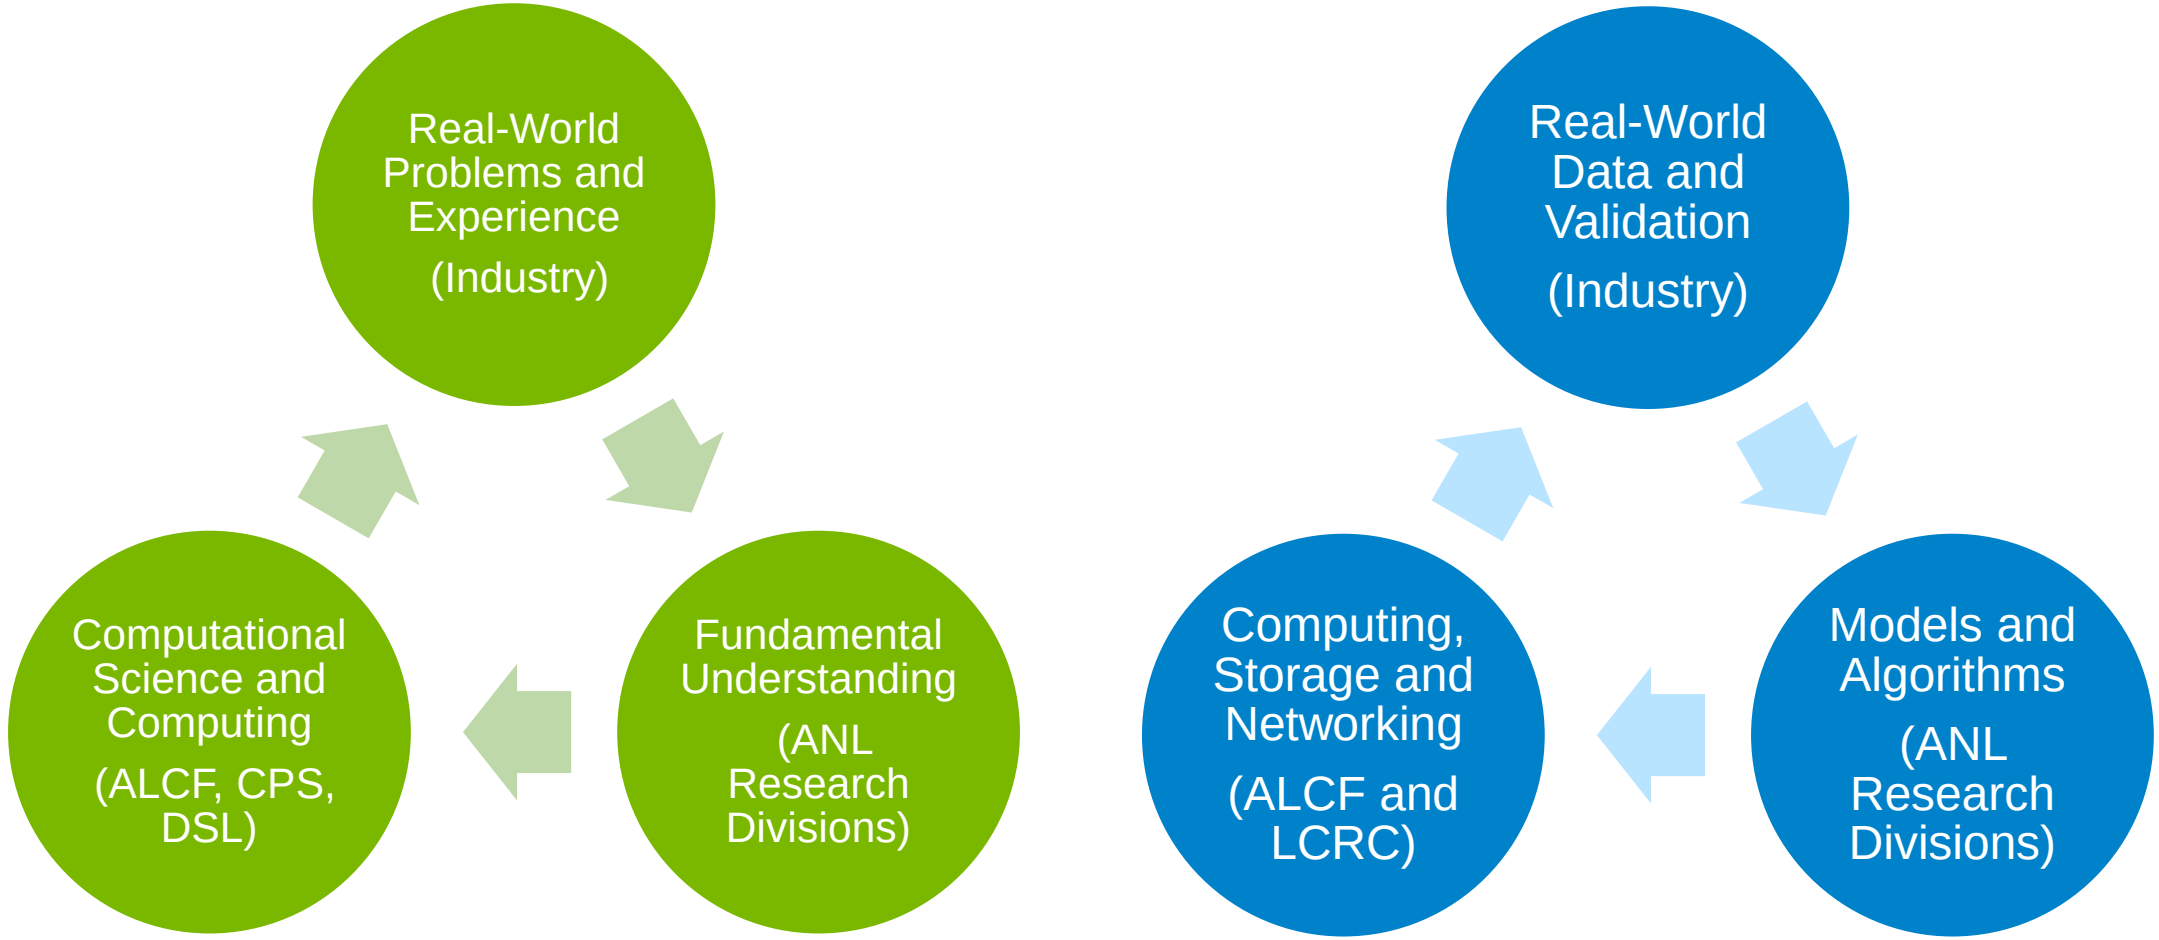

SWEET SPOT FOR COLLABORATIONS

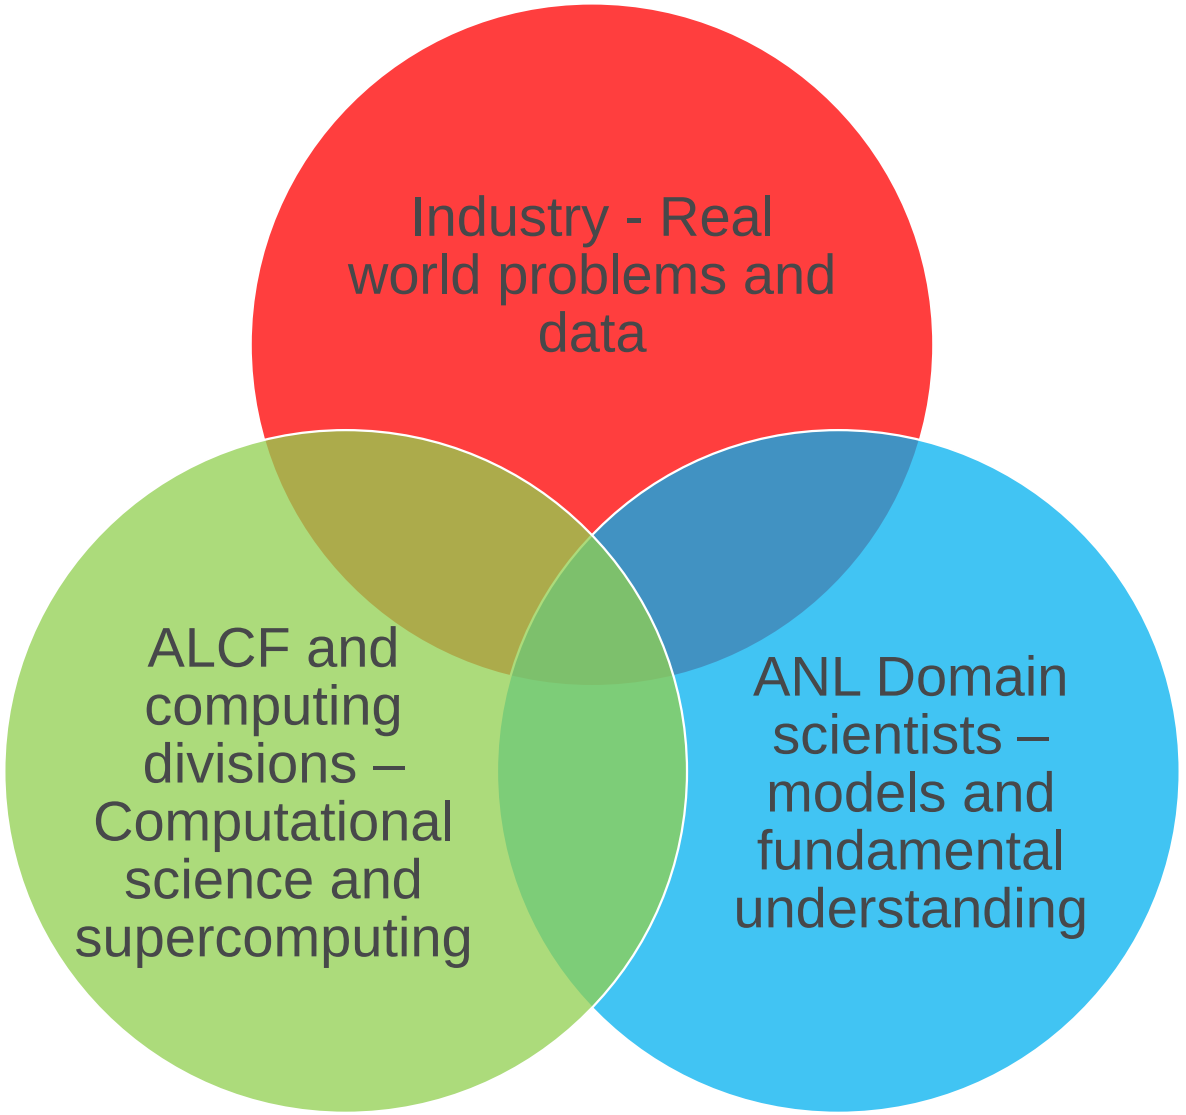


Argonne
NATIONAL LABORATORY