

Ready for Aurora



TIM WILLIAMS

Deputy Director

Computational Science Division (CPS)

Argonne National Laboratory

26 March 2025

Abstract

The webinar will highlight the wide variety of applications that are ready to run on Aurora's PVC GPU-accelerated architecture. In addition to representing many different scientific domains, the project workflows include many numerical methods, AI models for training and inference loads, and a variety of programming approaches for the PVC GPUs.

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

Aurora



NOW IN PRODUCTION!

Int
Ma

#3 in Top500 w/9234 nodes – HPL 1.012 exaFLOPS

4th

#1 in Top500 w/9500 nodes – **HPL-MxP** 11.6 exaFLOPS

Se

Band

Platform

HPE Cray-Ex

Dragonfly topology with adaptive routing
Cassini NIC, 200 Gb/s (25 GB/s), 8 per node
Network Switch:

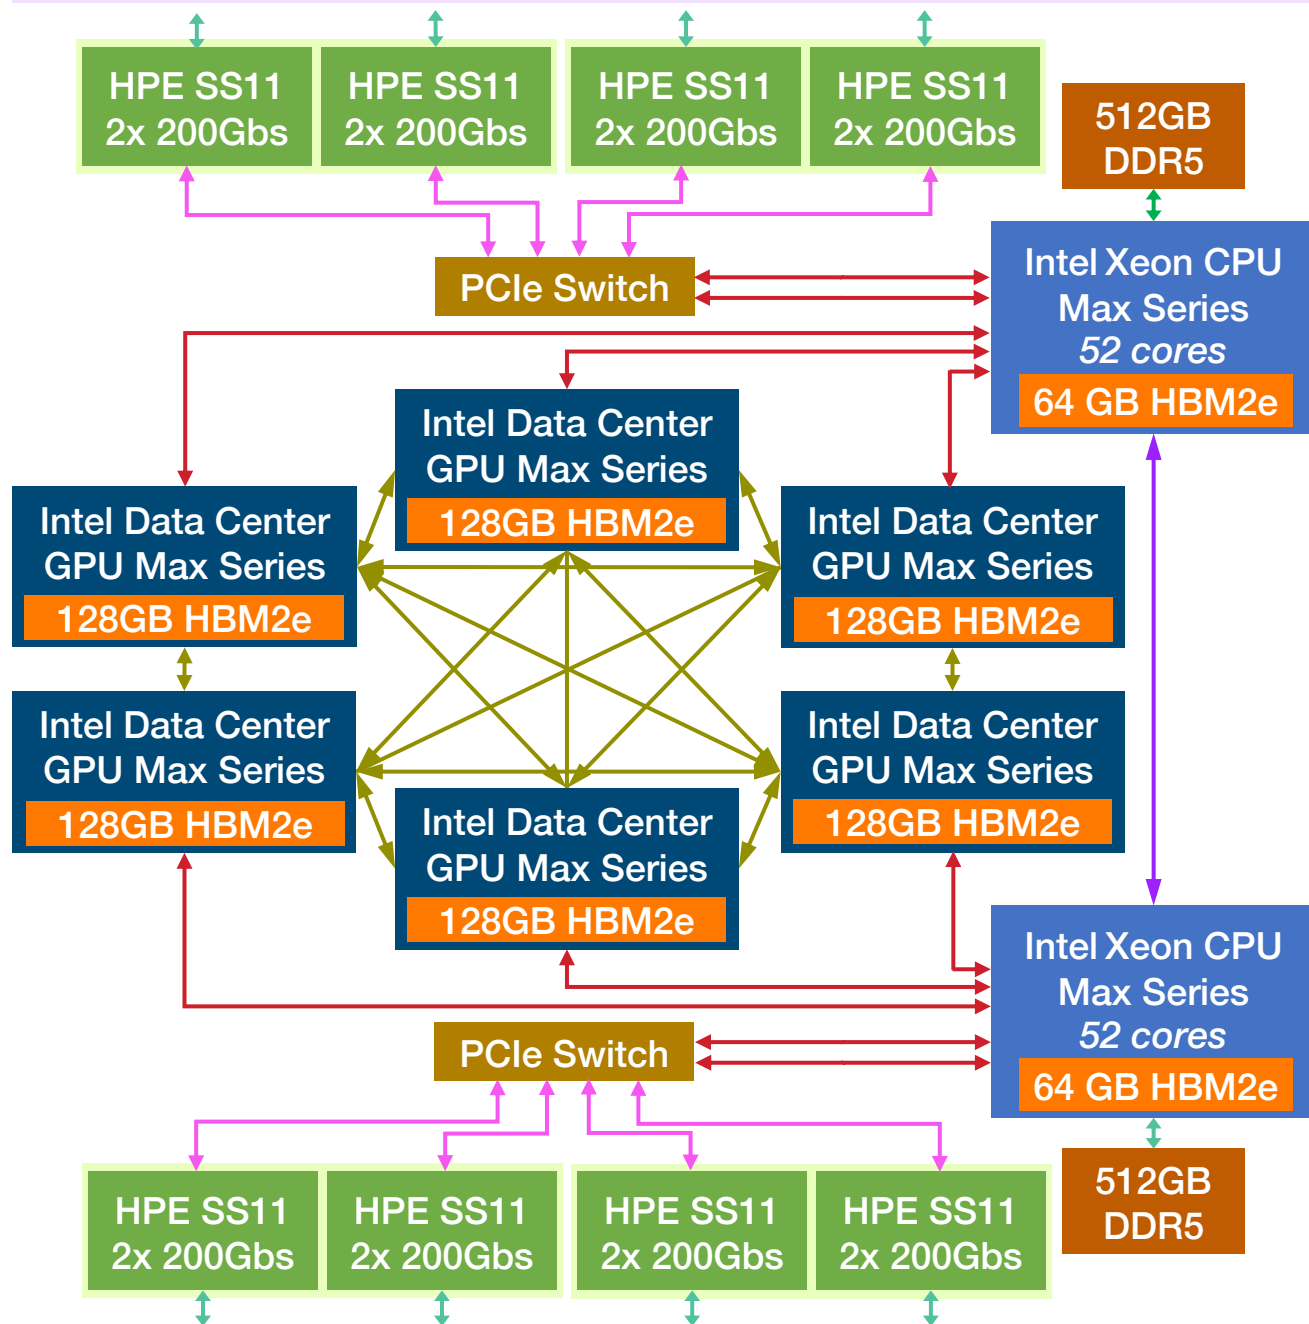
25.6 Tb/s per switch (64 200 Gb/s ports)
Links with 25 GB/s per direction

0.69 PB/s Peak Bisection BW

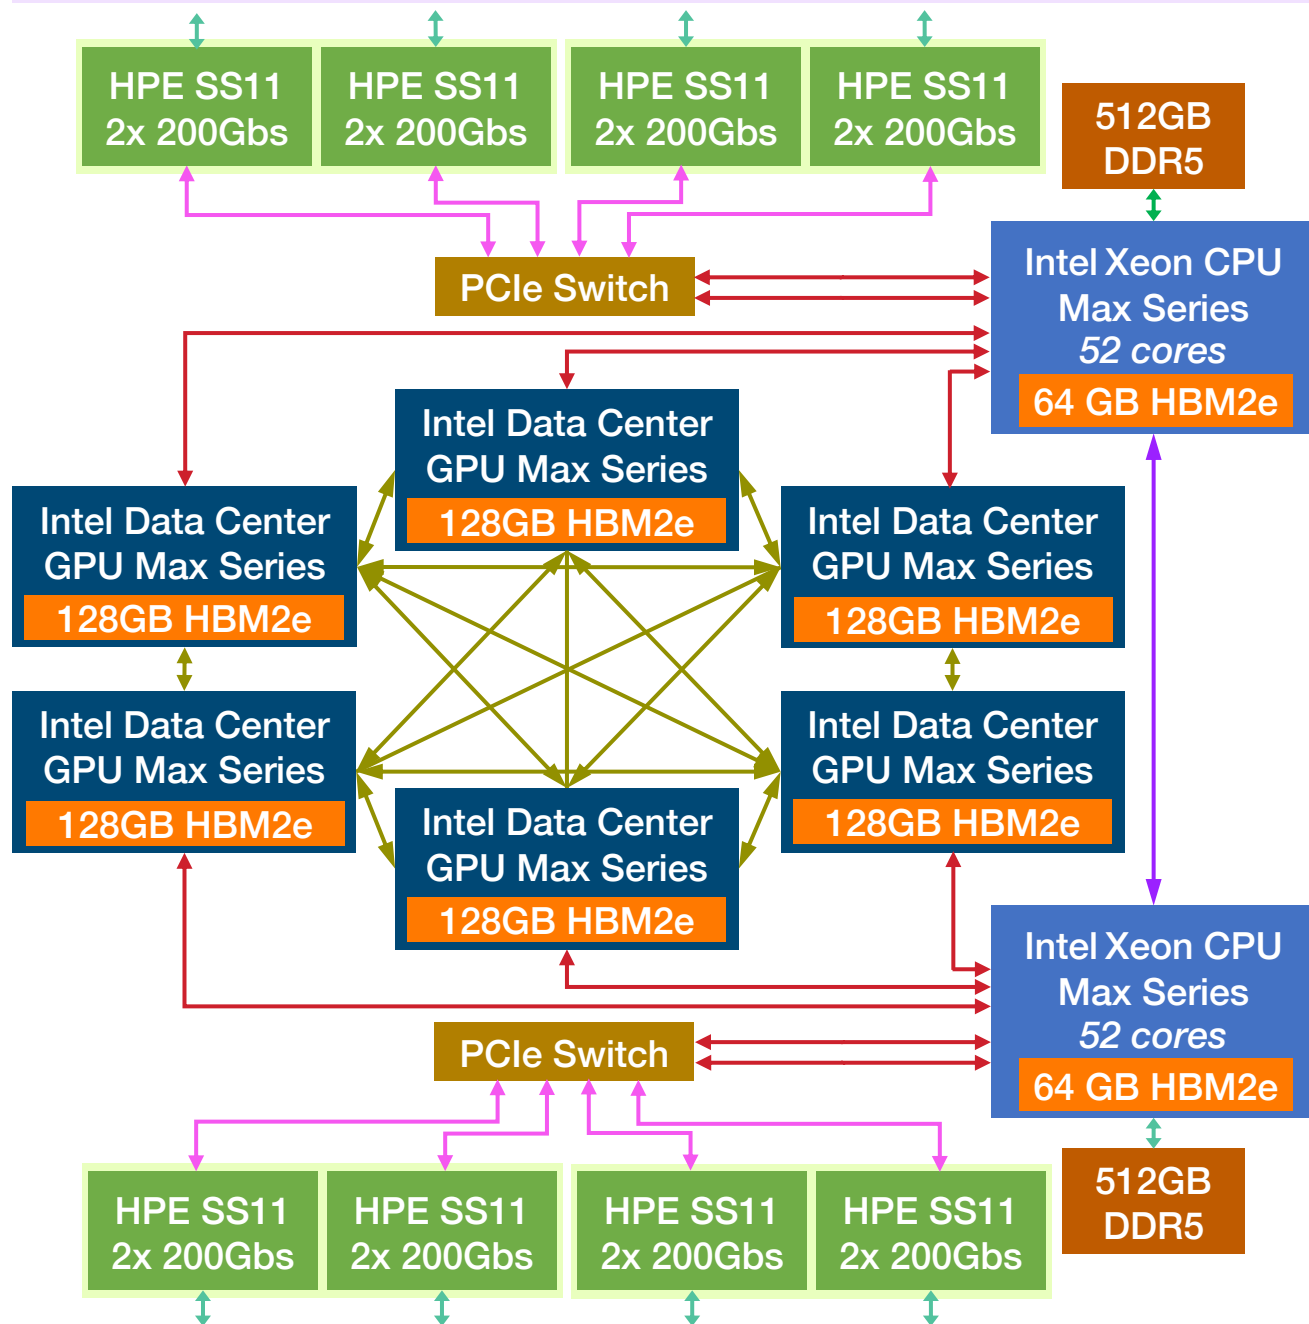
Storage

230PB DAOS Capacity
31 TB/s DAOS Bandwidth

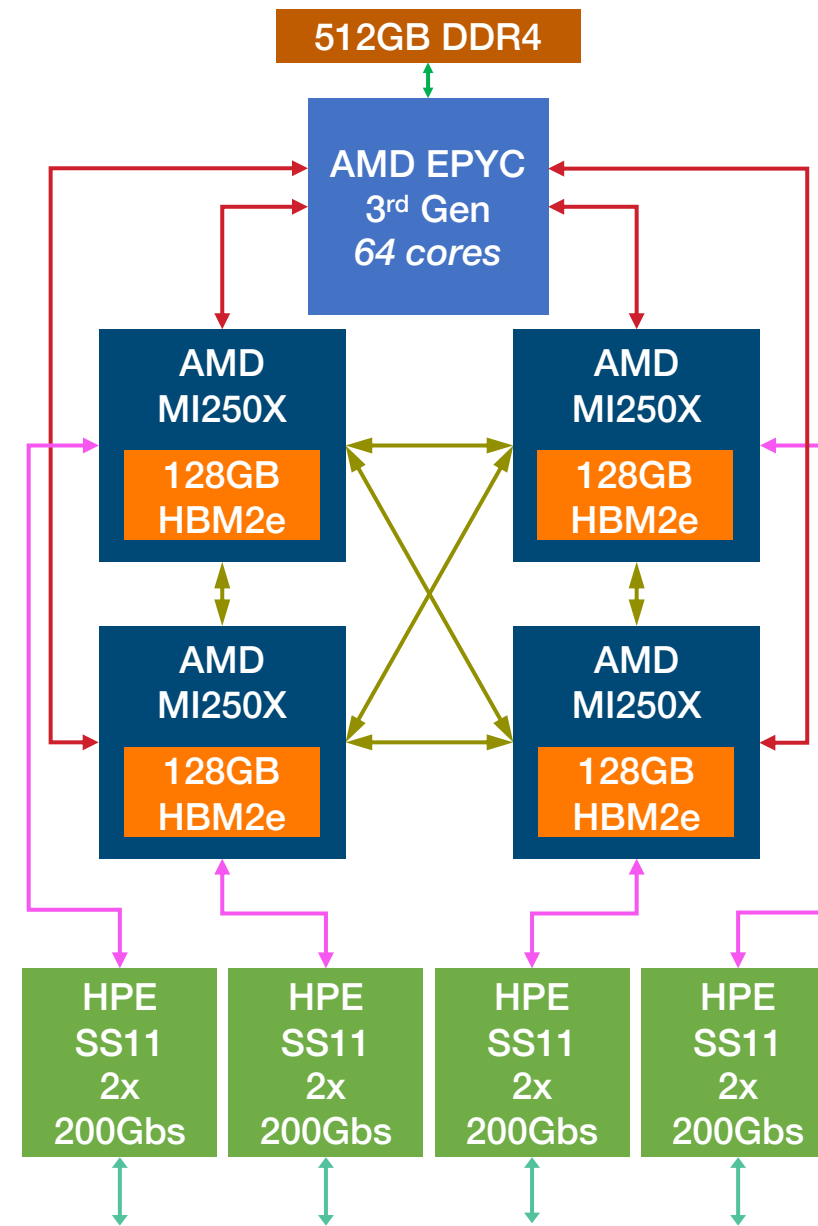
Aurora Node



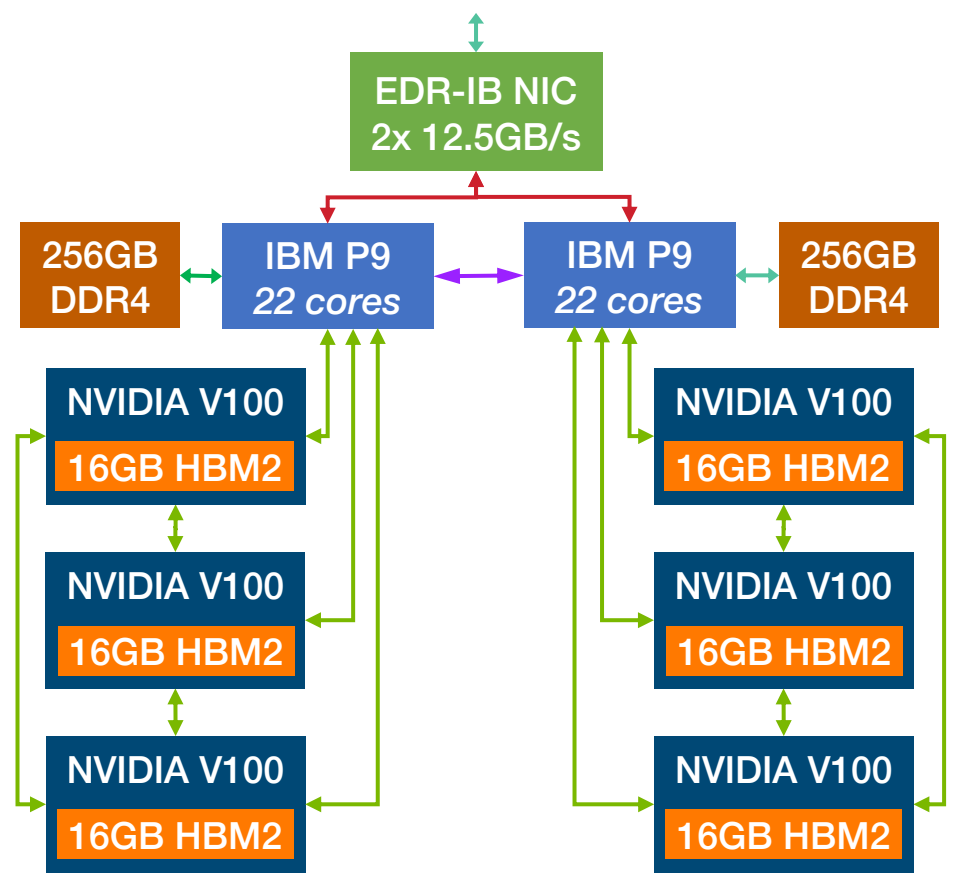
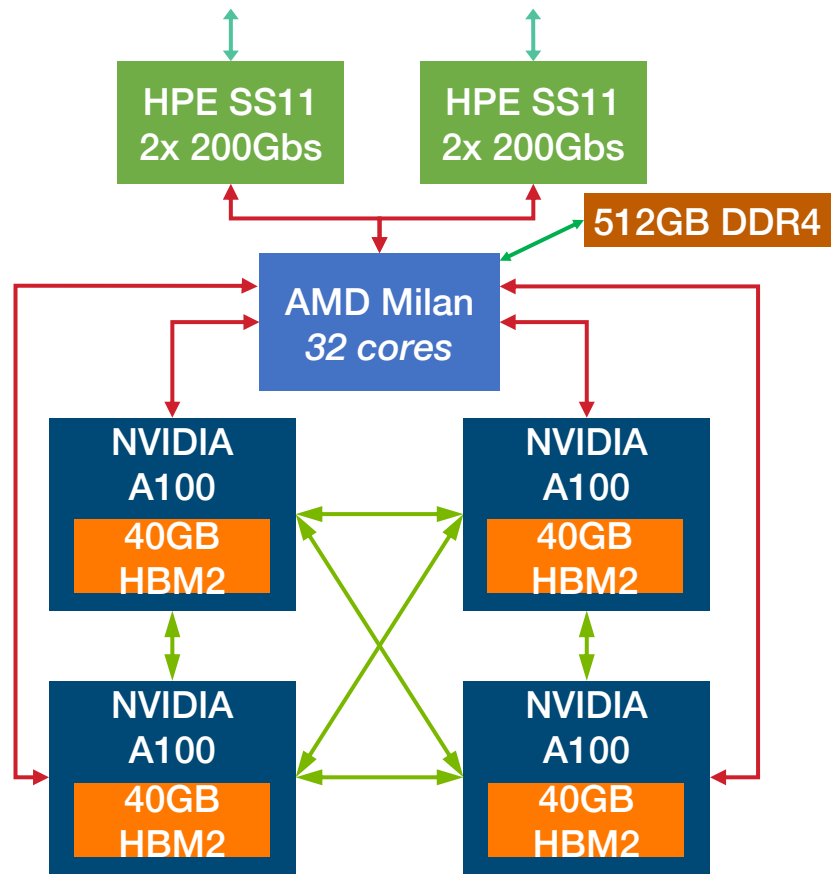
Aurora Node



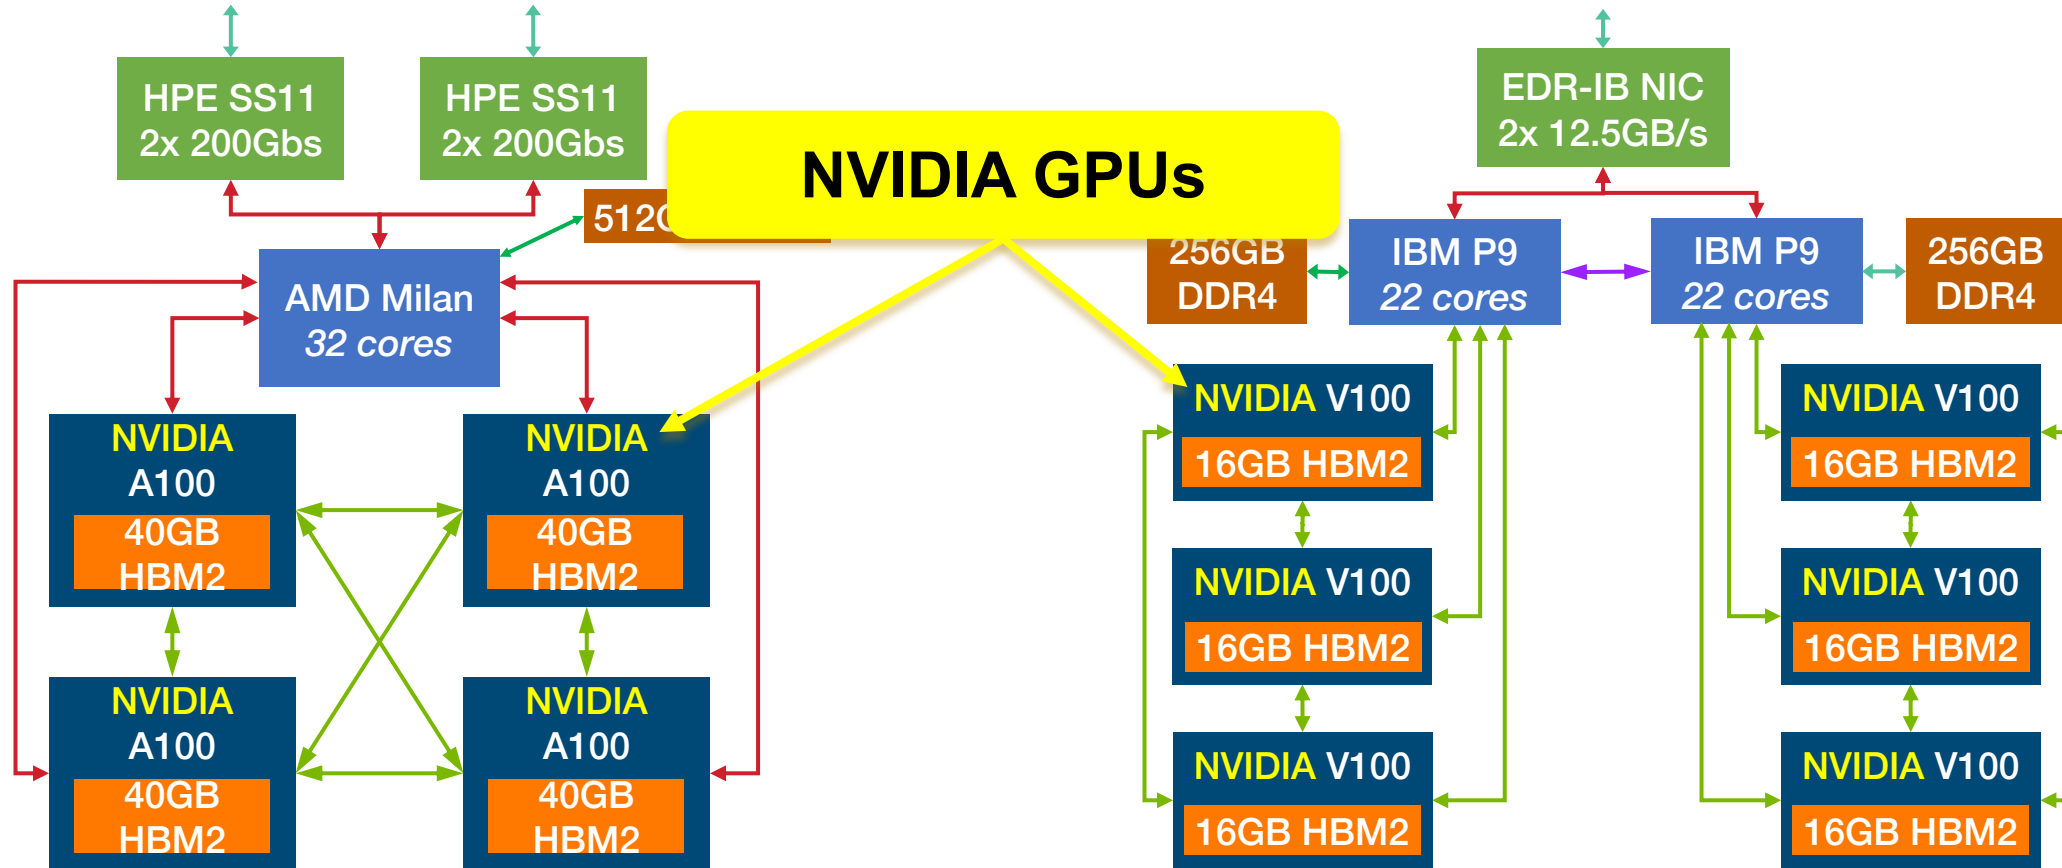
Frontier Node



Pre-Exascale Systems







Pre-Exascale Systems



Aurora Applications/Software Readiness Efforts


- 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



Aurora Early Science Program


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

Simulation



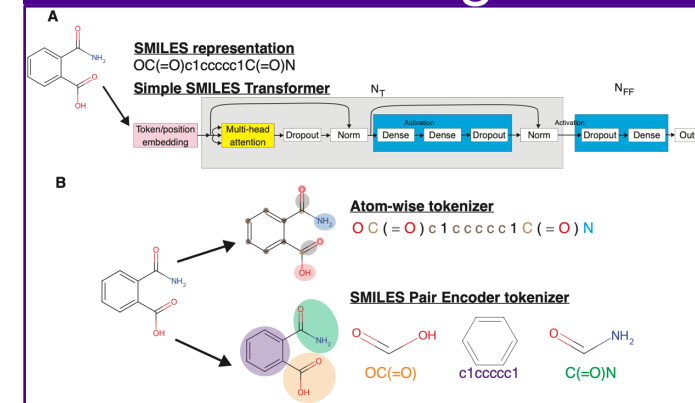
- Reactive, quantum, and classical MD for materials and biophysics
- ITER tungsten impurity ions kinetics—PIC
- Cosmological P³M + SPH hydrodynamics
- CCSD(T) quantum chemistry
- Quantum Monte Carlo
- FEM CFD, CFD+combustion

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 P³M simulation data
- Chemistry kinetics database
- Tokamak experiment data

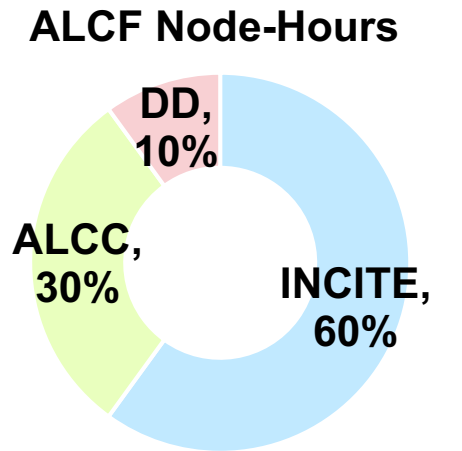
Learning



- 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

First-Year Science Projects on Aurora

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



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



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.

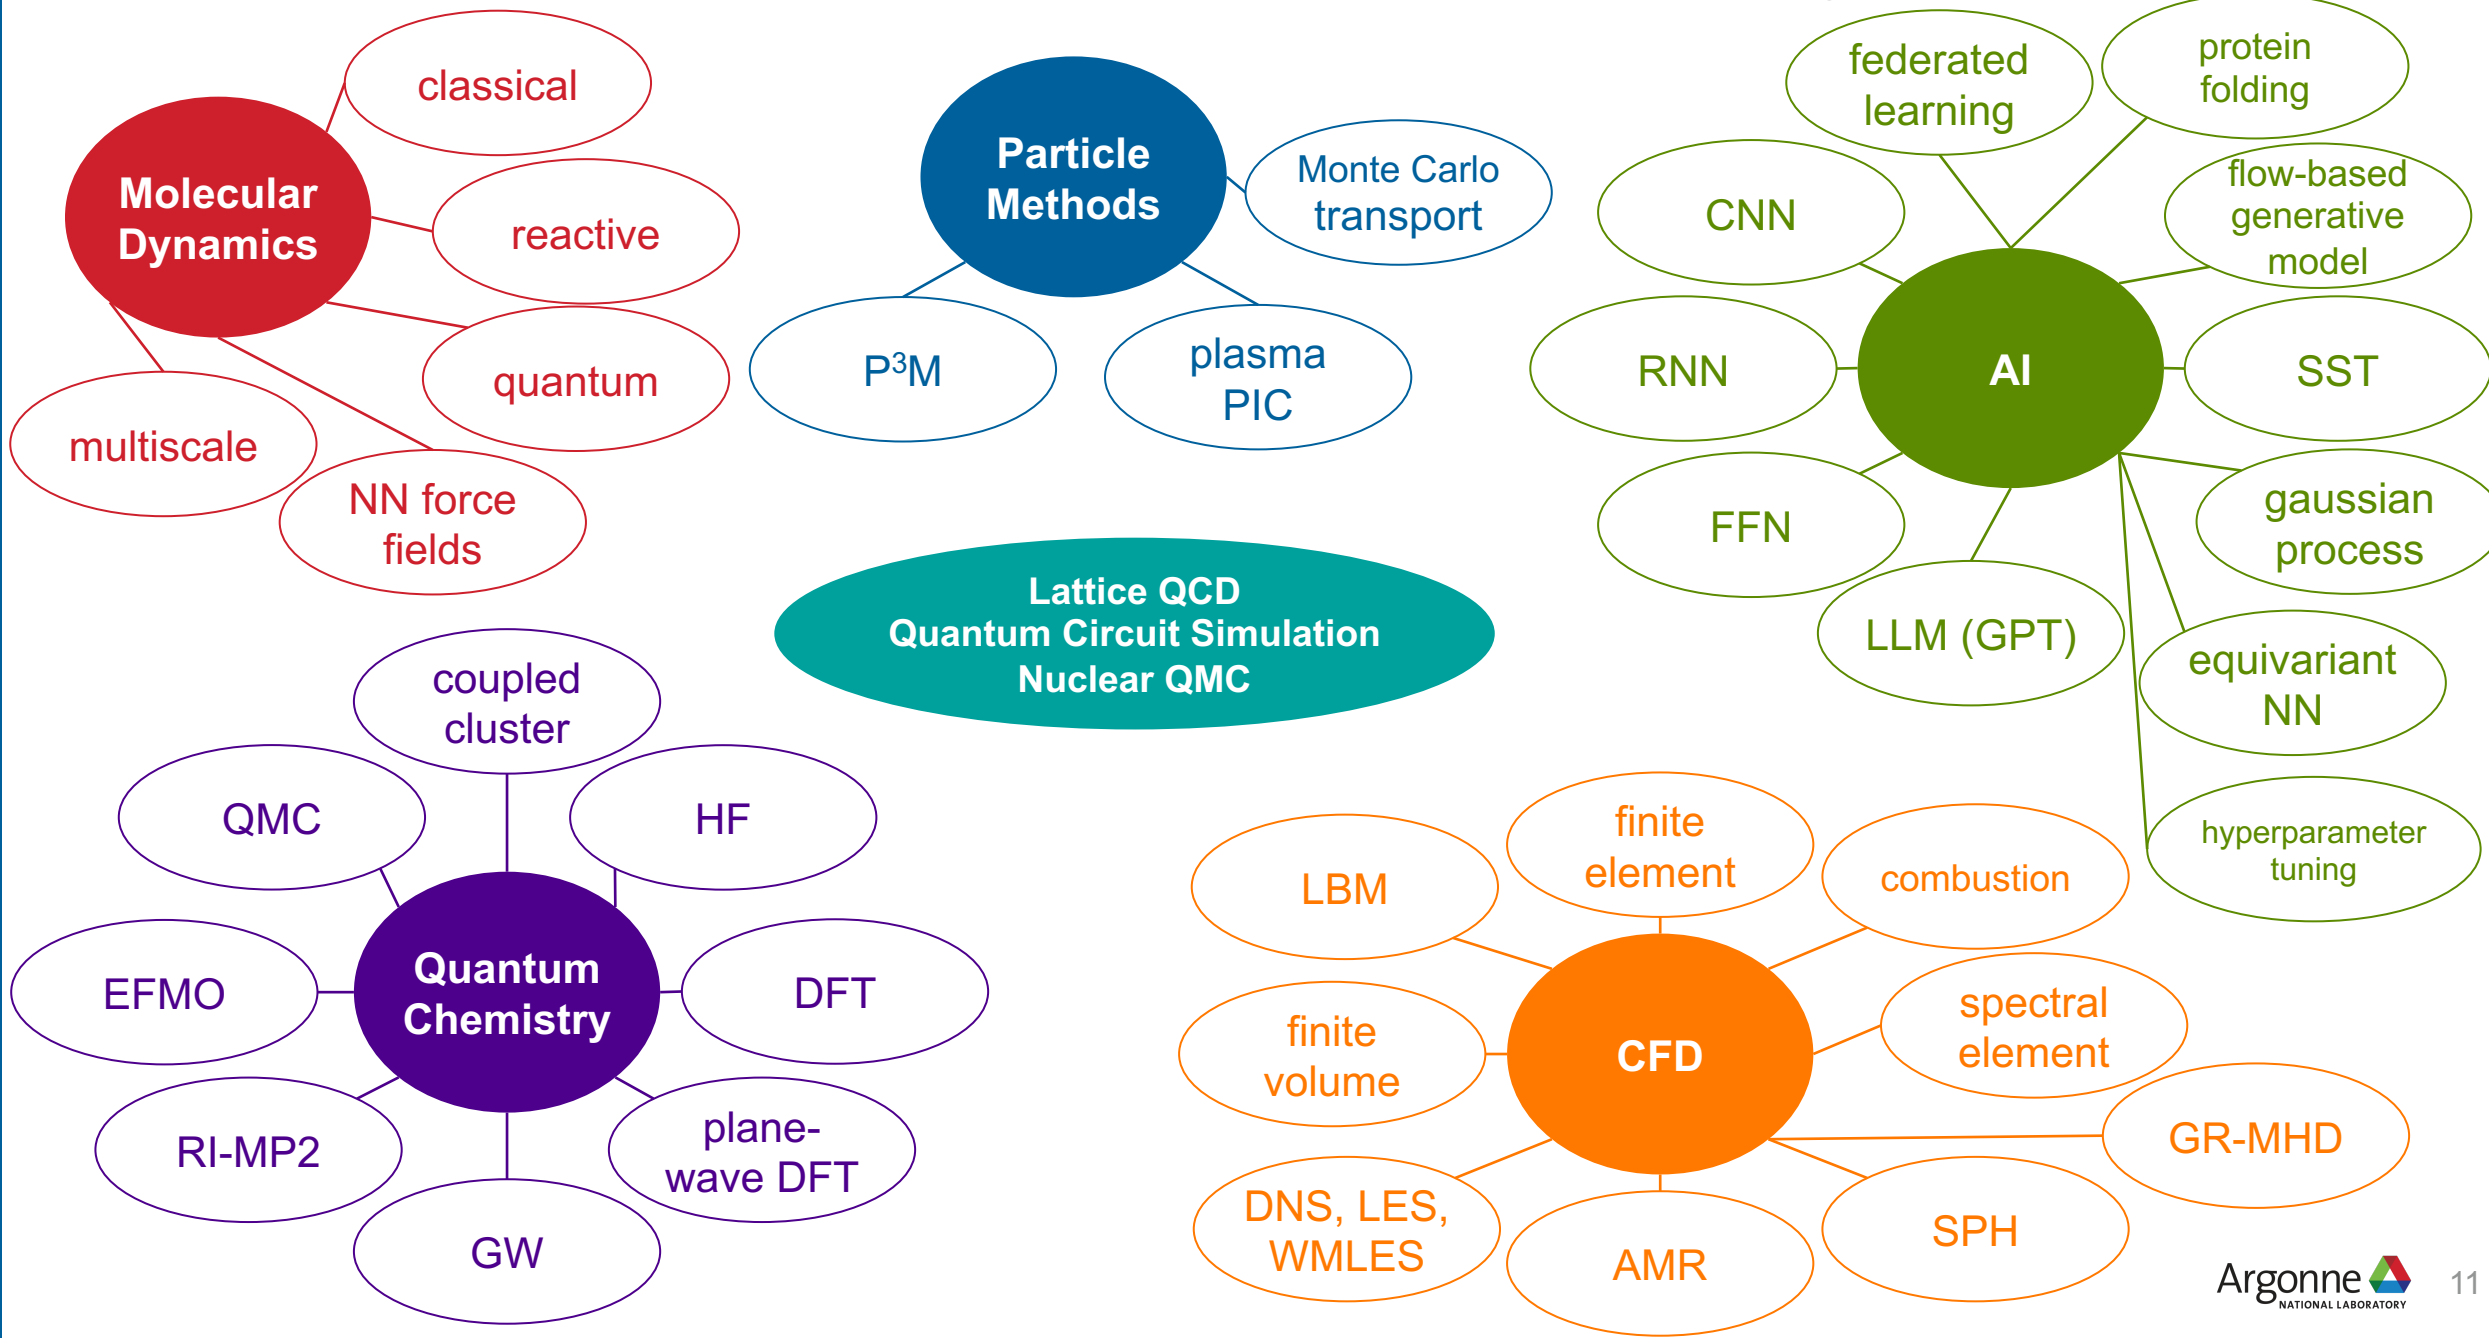
- ALCC (ASCR Leadership Computing challenge)
 - 13 projects
 - 5 million node-hours



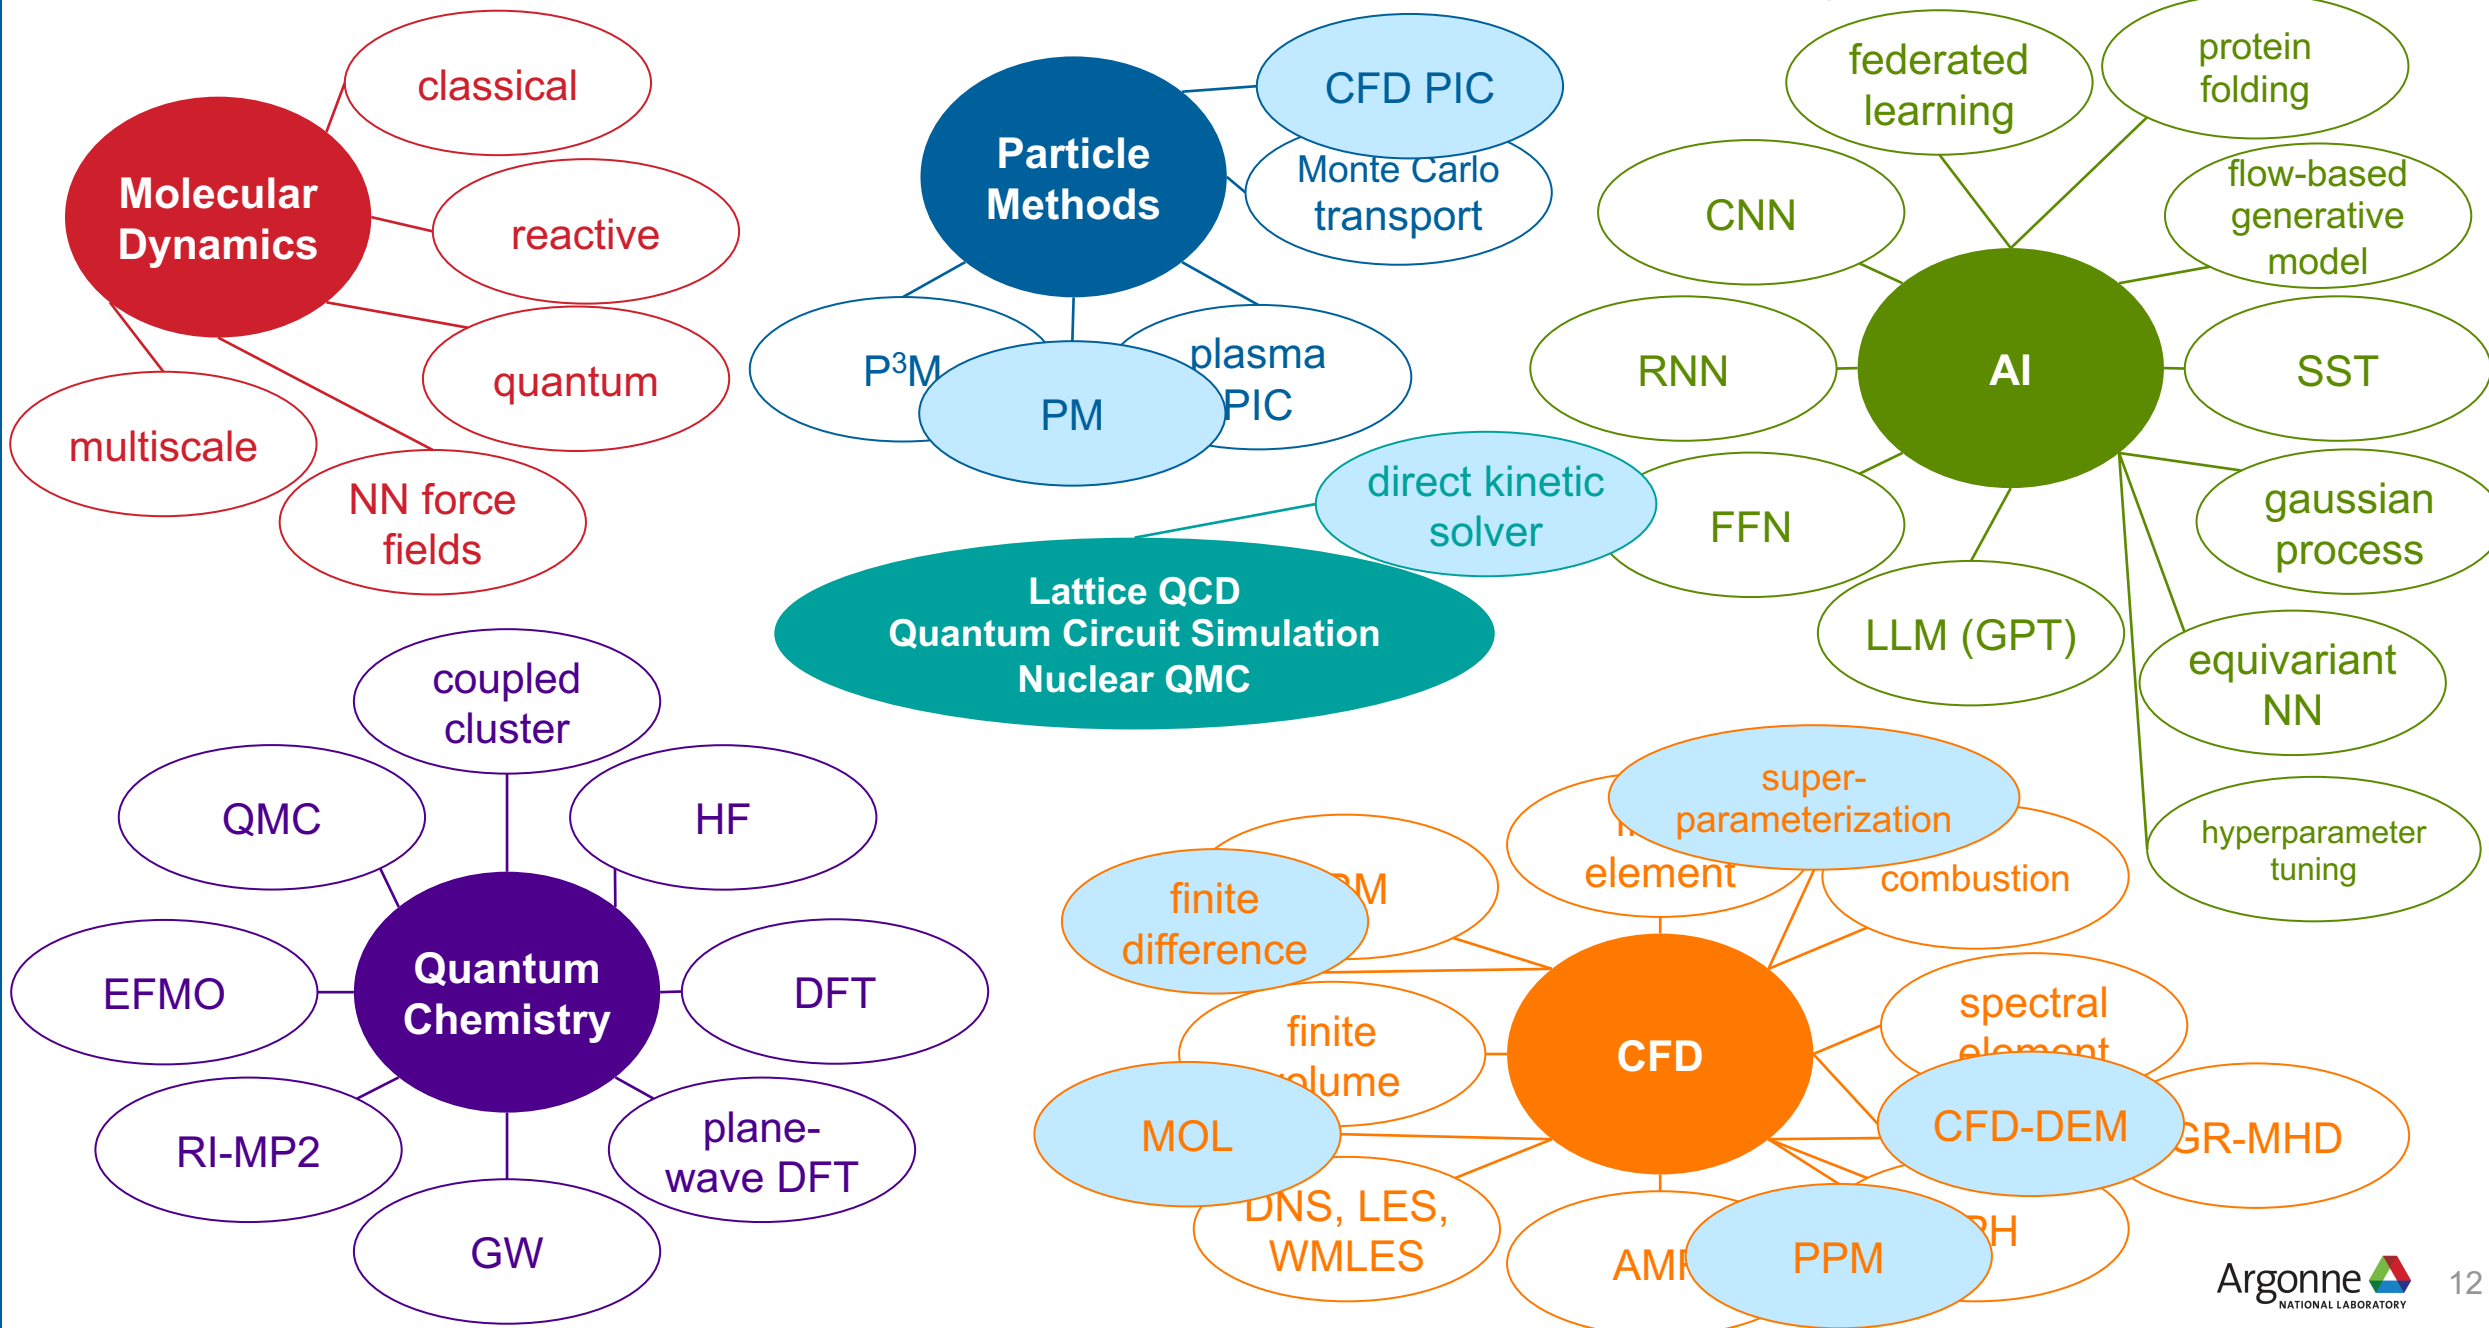
U.S. DEPARTMENT OF
ENERGY

Office of
Science

Computational Science Methods in Y1 Aurora Projects



Computational Science Methods in Y1 Aurora Projects + ECP



Science Domains in Y1 Aurora Projects

Aerodynamics

Combustion

Thermal hydraulics for fusion & fission

Turbulent Rayleigh-Benard convection

Buoyancy-driven turbulence

Wind energy

Fusion radiation transport

Fusion plasma kinetics

Tokamak control

Neuroscience foundation model

Science foundation model

Federated learning

Quantum computing

AI assistant for simulation

Core collapse supernovae

Binary neutron star mergers

Core collapse supernovae

Black hole accretion

Neutron star remnants

Functional quantum materials

Fusion device materials

Multiscale materials modeling

Quantum materials

Materials science

QCD

Cosmology

Nuclear structure and reactions

Collider data analysis

MOF design

Carbon capture

Transition metal chemistry

Heterogeneous catalysis

Heterogeneous reactions

Biofuels

Membrane transport proteins

Predictive molecular epidemiology

Blood flow & transport

Cancer drug design

Heteropolymer design

Brain connectomics

Pathogenesis

Climate modeling

Runoff inundation

Science Domains in Y1 Aurora Projects + ECP

Aerodynamics

Combustion

Thermal hydraulics for fusion & fission

Turbulent Rayleigh-Benard convection

Buoyancy-driven turbulence

Wind energy

Fusion radiation transport

Fusion plasma kinetics

Tokamak control

Neuroscience foundation model

Science foundation model

Federated learning

Quantum computing

AI assistant for simulation

Core collapse supernovae

Binary neutron star mergers

Core collapse supernovae

Black hole accretion

Neutron star remnants

Functional quantum materials

Fusion device materials

Multiscale materials modeling

Quantum materials

Materials science

QCD

Cosmology

Nuclear structure and reactions

Collider data analysis

MOF design

Carbon capture

Transition metal chemistry

Heterogeneous catalysis

Heterogeneous reactions

Biofuels

Membrane transport proteins

Predictive molecular epidemiology

Blood flow & transport

Cancer drug design

Heteropolymer design

Brain connectomics

Pathogenesis

Climate modeling

Runoff inundation


Seismic wave propagation

Aurora Y1 and ECP Applications by Name

INCITE, ALCC, ESP, ECP Applications With PVC Offload			
AMR-Wind	DiffLinker	HARVEY	QUDA
Allegro	E3SM-MMF	HemeLB_GPU	Quantum Espresso
AlphaFold2	ExaChem	LAMMPS	RXMD-NN
AthenaK	FLARE	MFDn	Rosetta
AuroraGPT	FUN3D	MFIX-Exa	SCREAM
BerkeleyGW	FastCaloSim	MILC	SR-APPFL
CANDLE/Uno	Flood Fill Network	MadGraph	SST
CEED-PHASTA	Fornax	Masala	SwiFT
CHROMATIX	FusionDL	NAMD	TRITON
CP2K	GAMESS	NWChemEx	Thornado/Flash-X
CPS	GENE	NekRS	UIntah
Chroma	GTC	OpenMC	Uno
Contegra	GURU	OpenMM	XGC
DCMESH	Grid	PHASTA	
DarkSkyML	H-AMR	PWDFT	
DataDrivenCFD	HACC	QMCPACK	

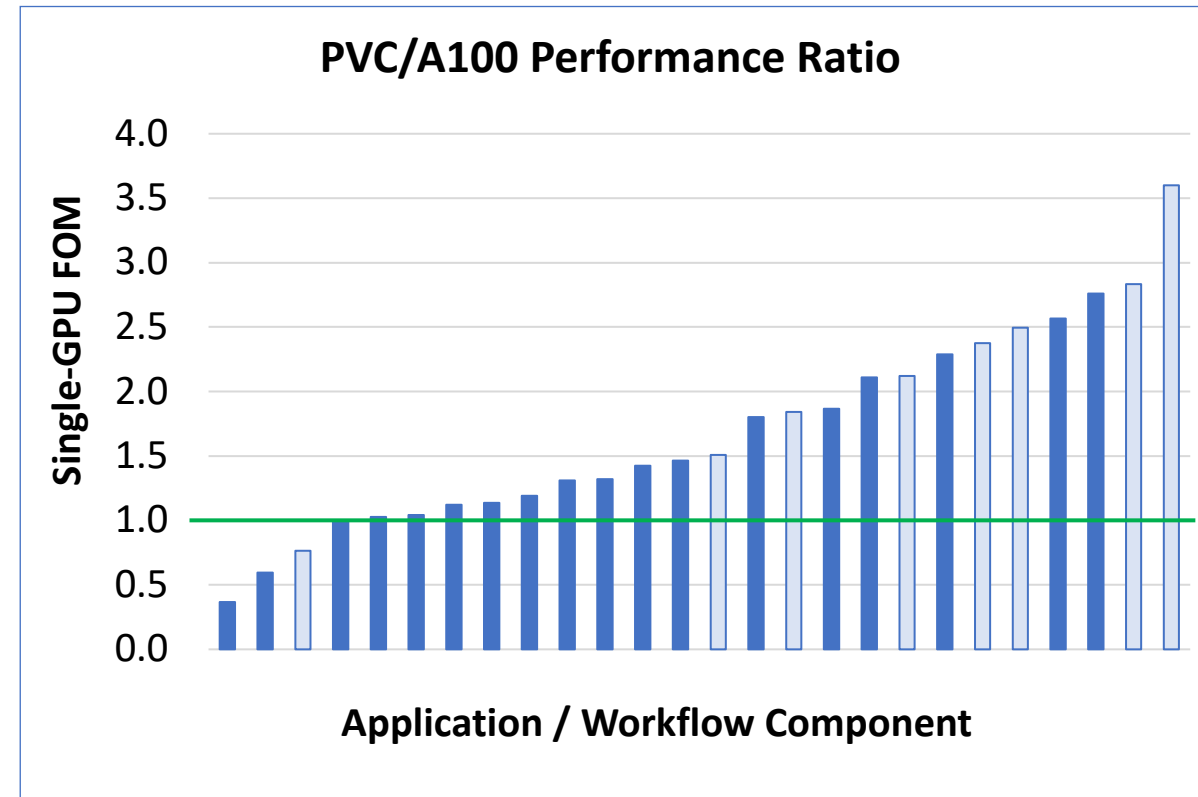
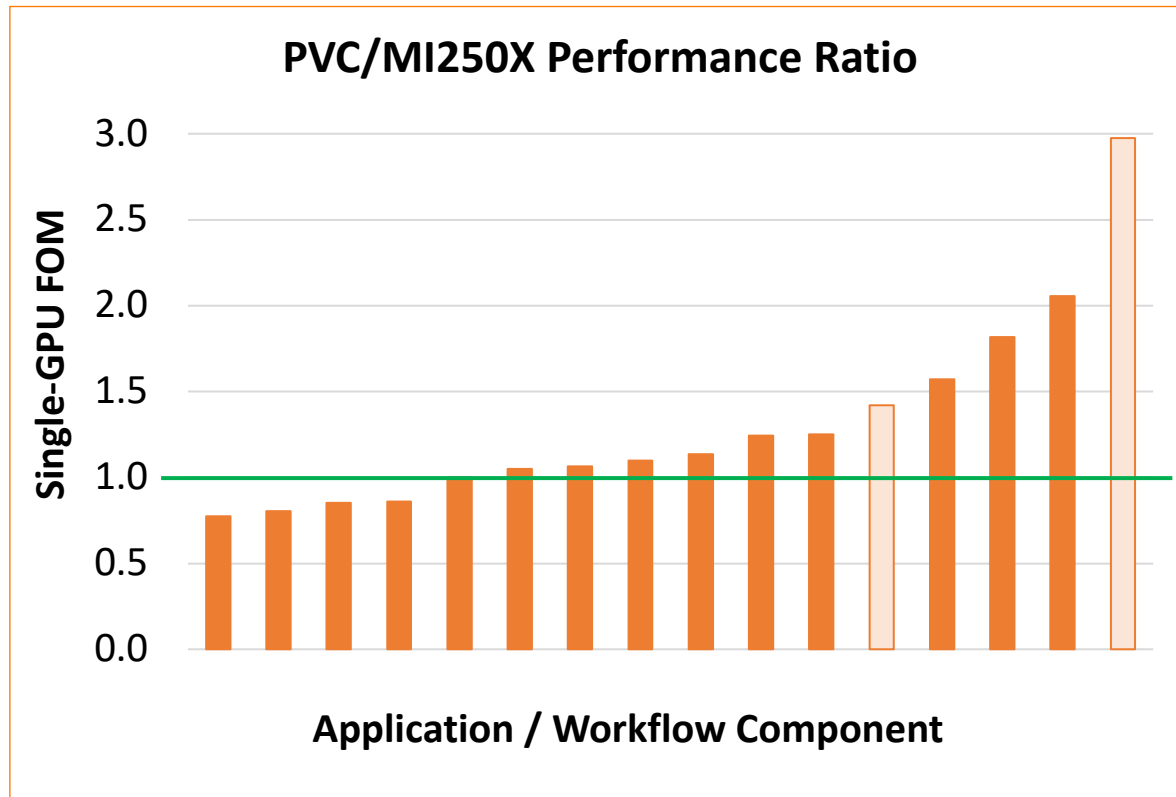
Aurora Performance and Expectations

- 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+external review: 16/19 projects met goal in Fall 2024



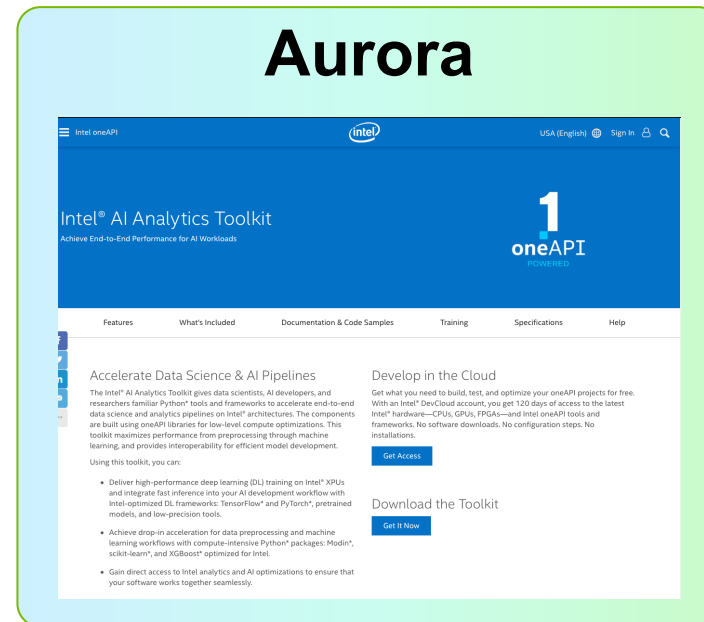
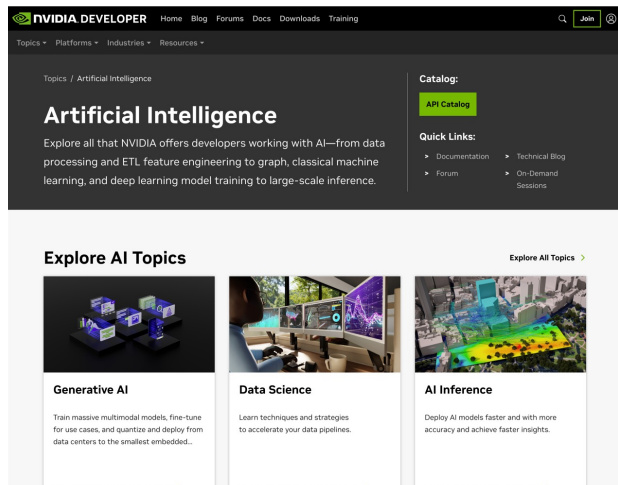
Aurora Performance and Expectations

- 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)

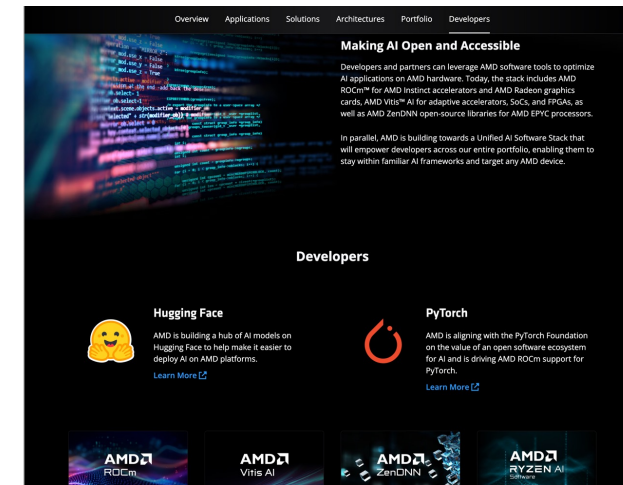


Exascale Programming Models – Data & Learning

- 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

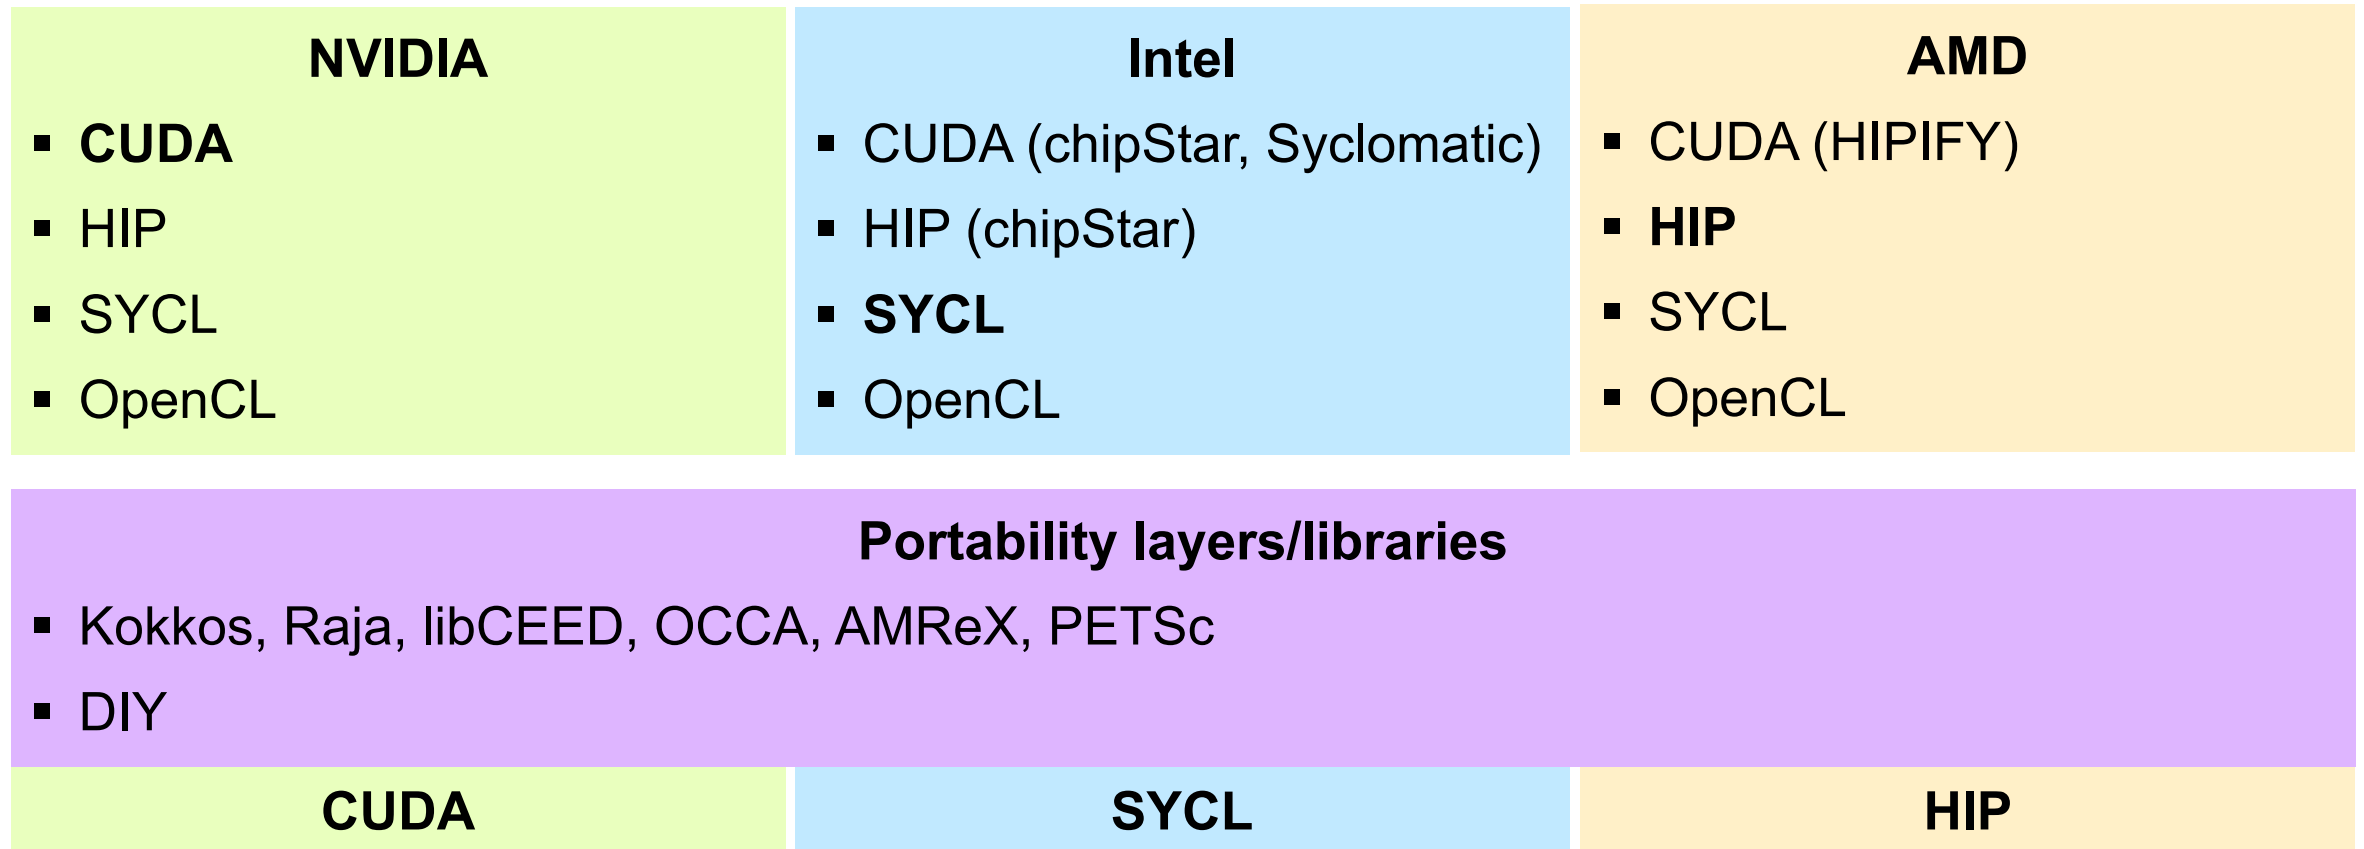


https://docs.alcf.anl.gov/aurora/_data-science/python/



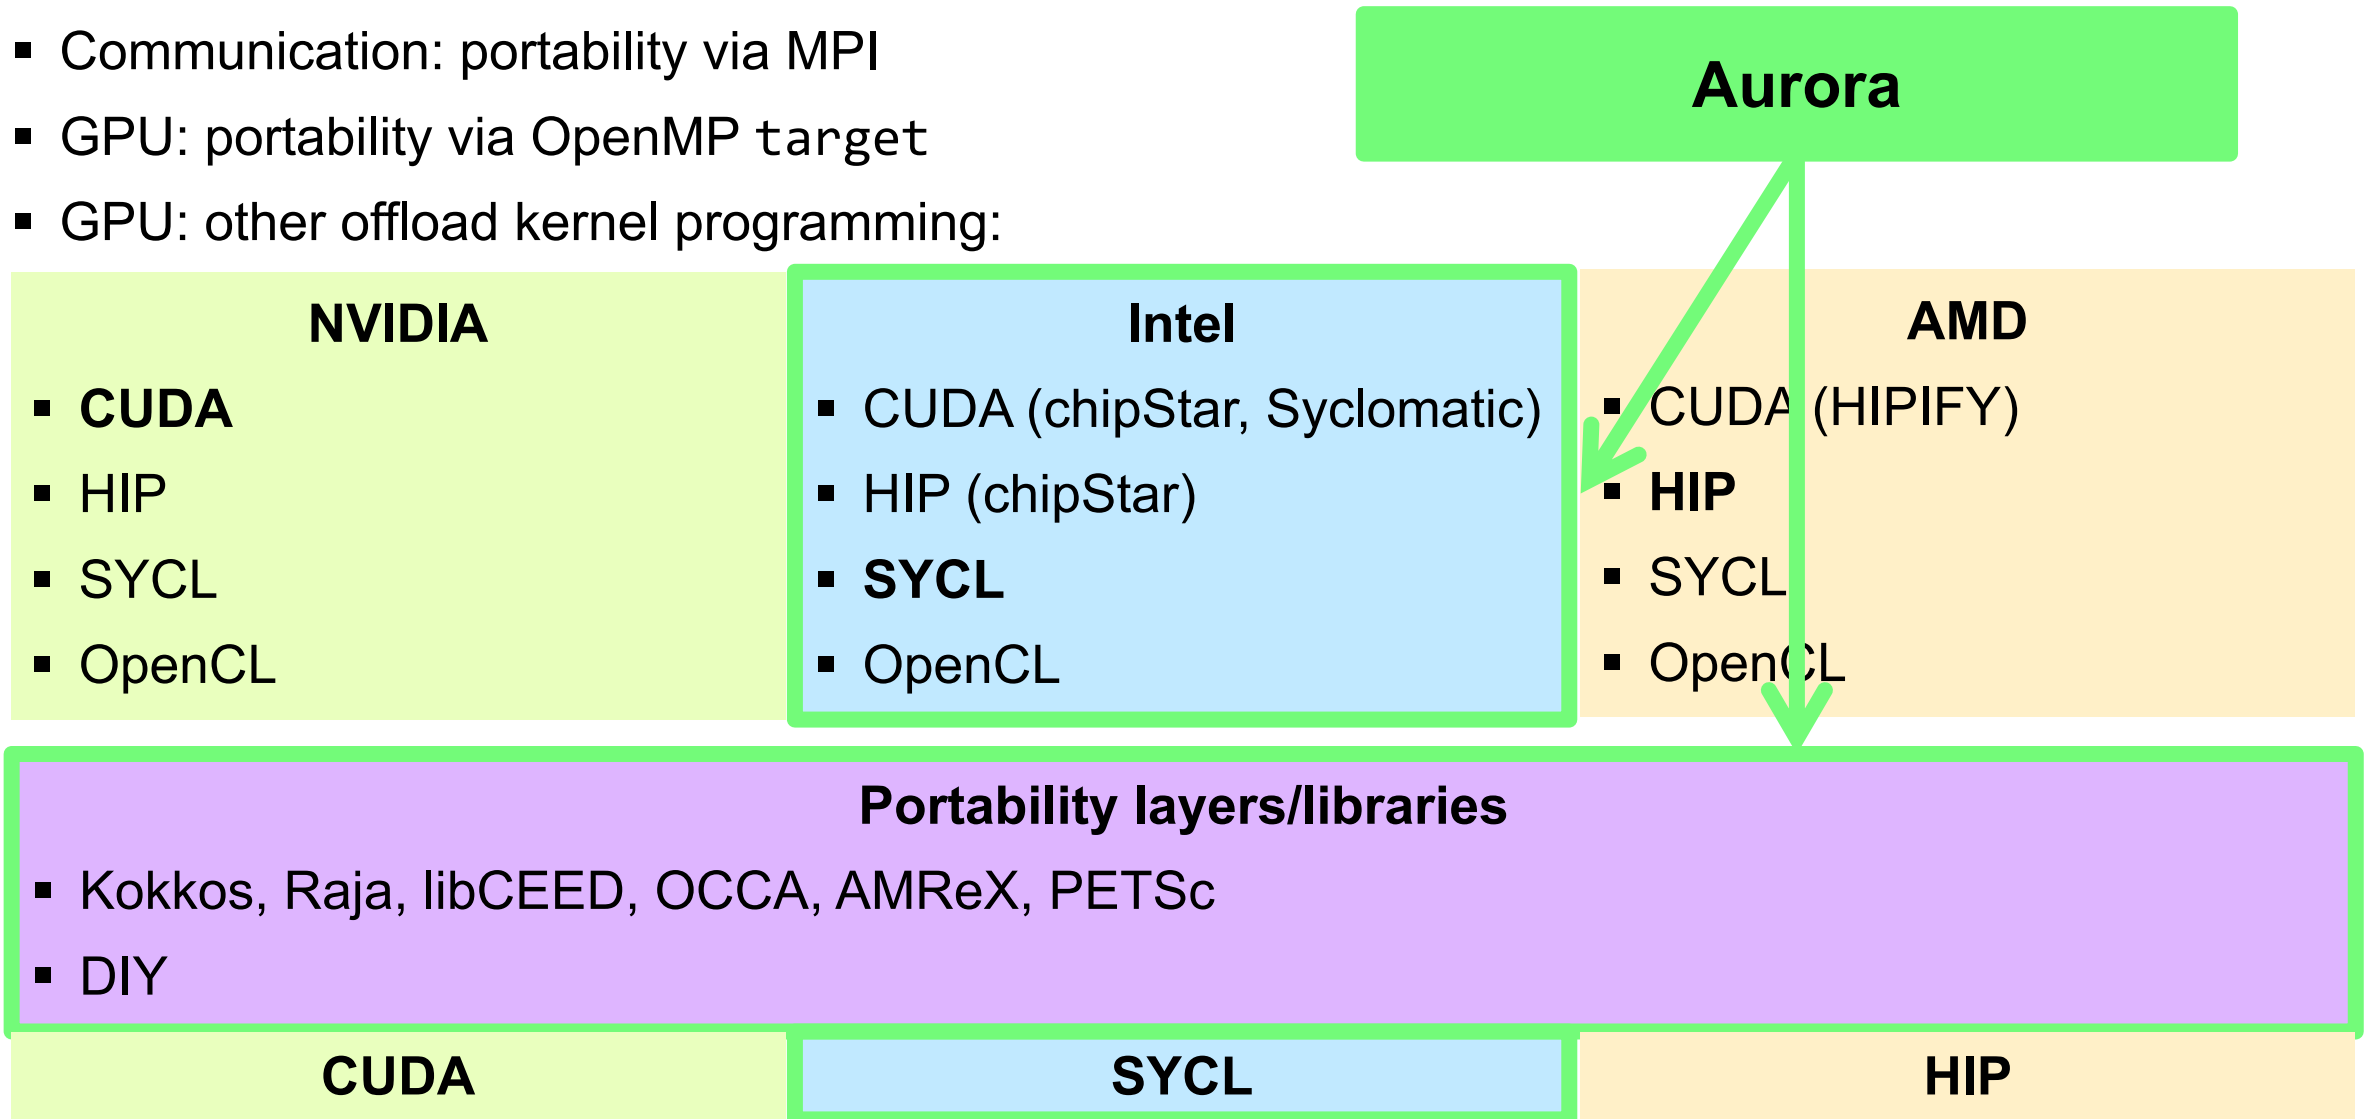
Exascale Programming Models – Simulation

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

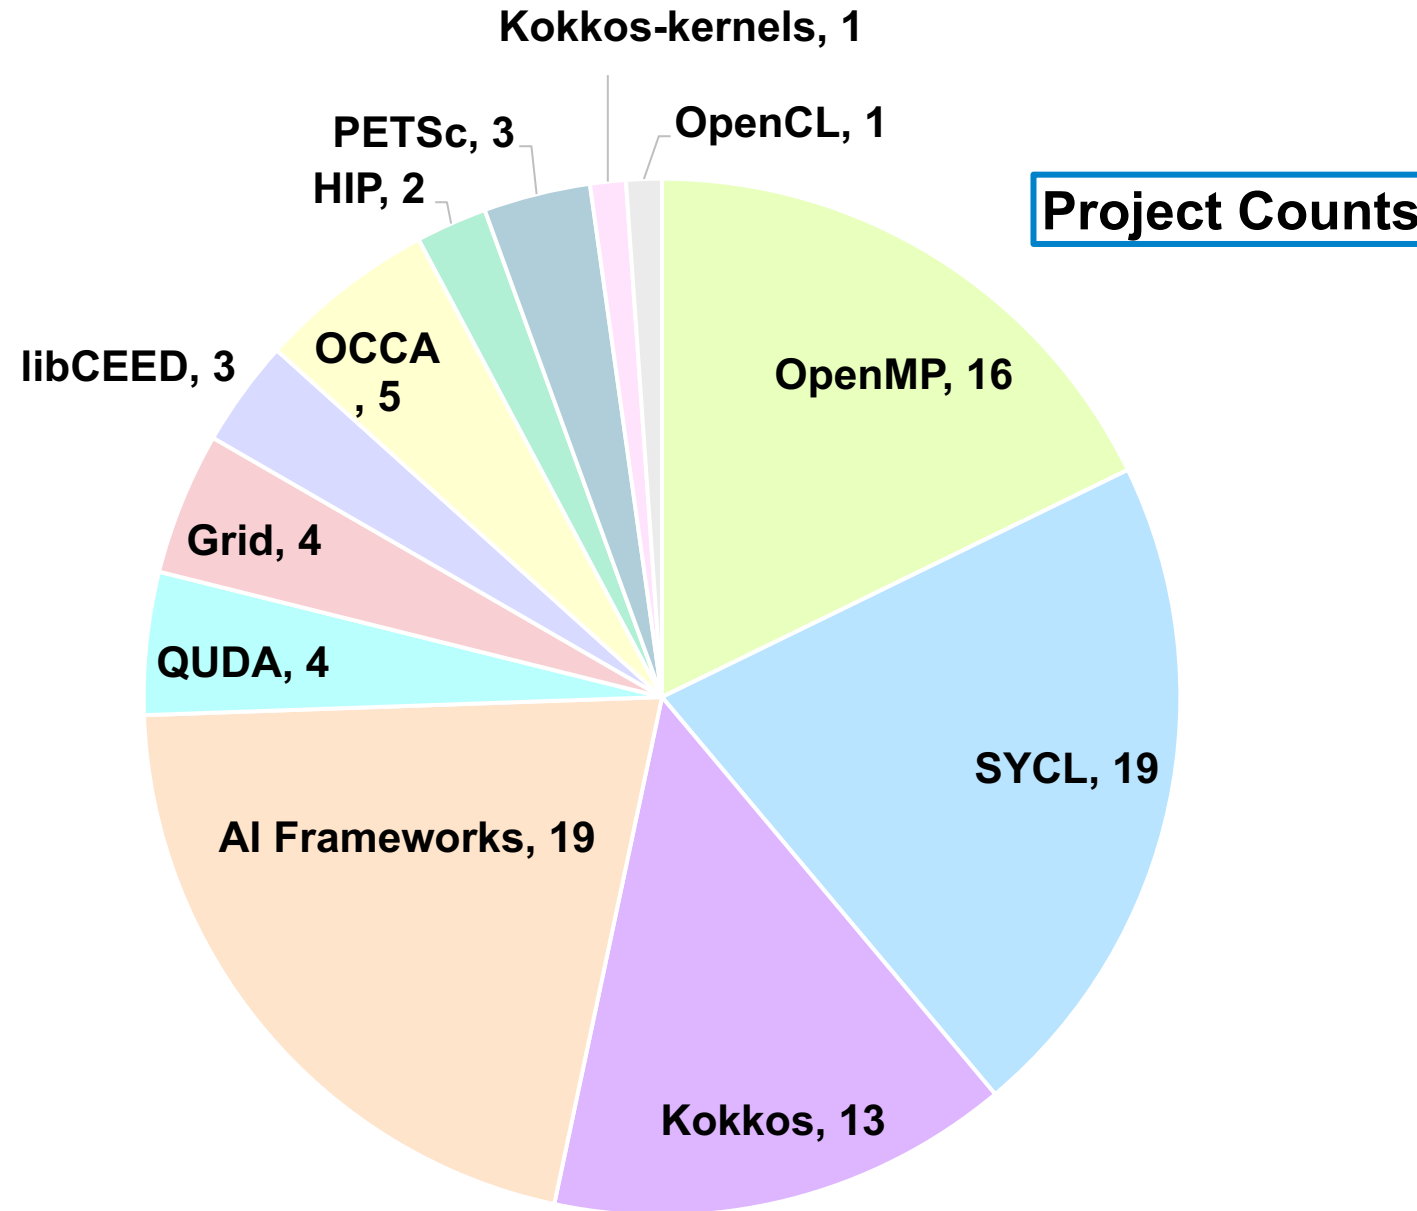


Exascale Programming Models – Simulation

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



Programming Model Choices by Y1 Aurora Projects



Programming Model Choices by Y1 Aurora Projects + ECP

