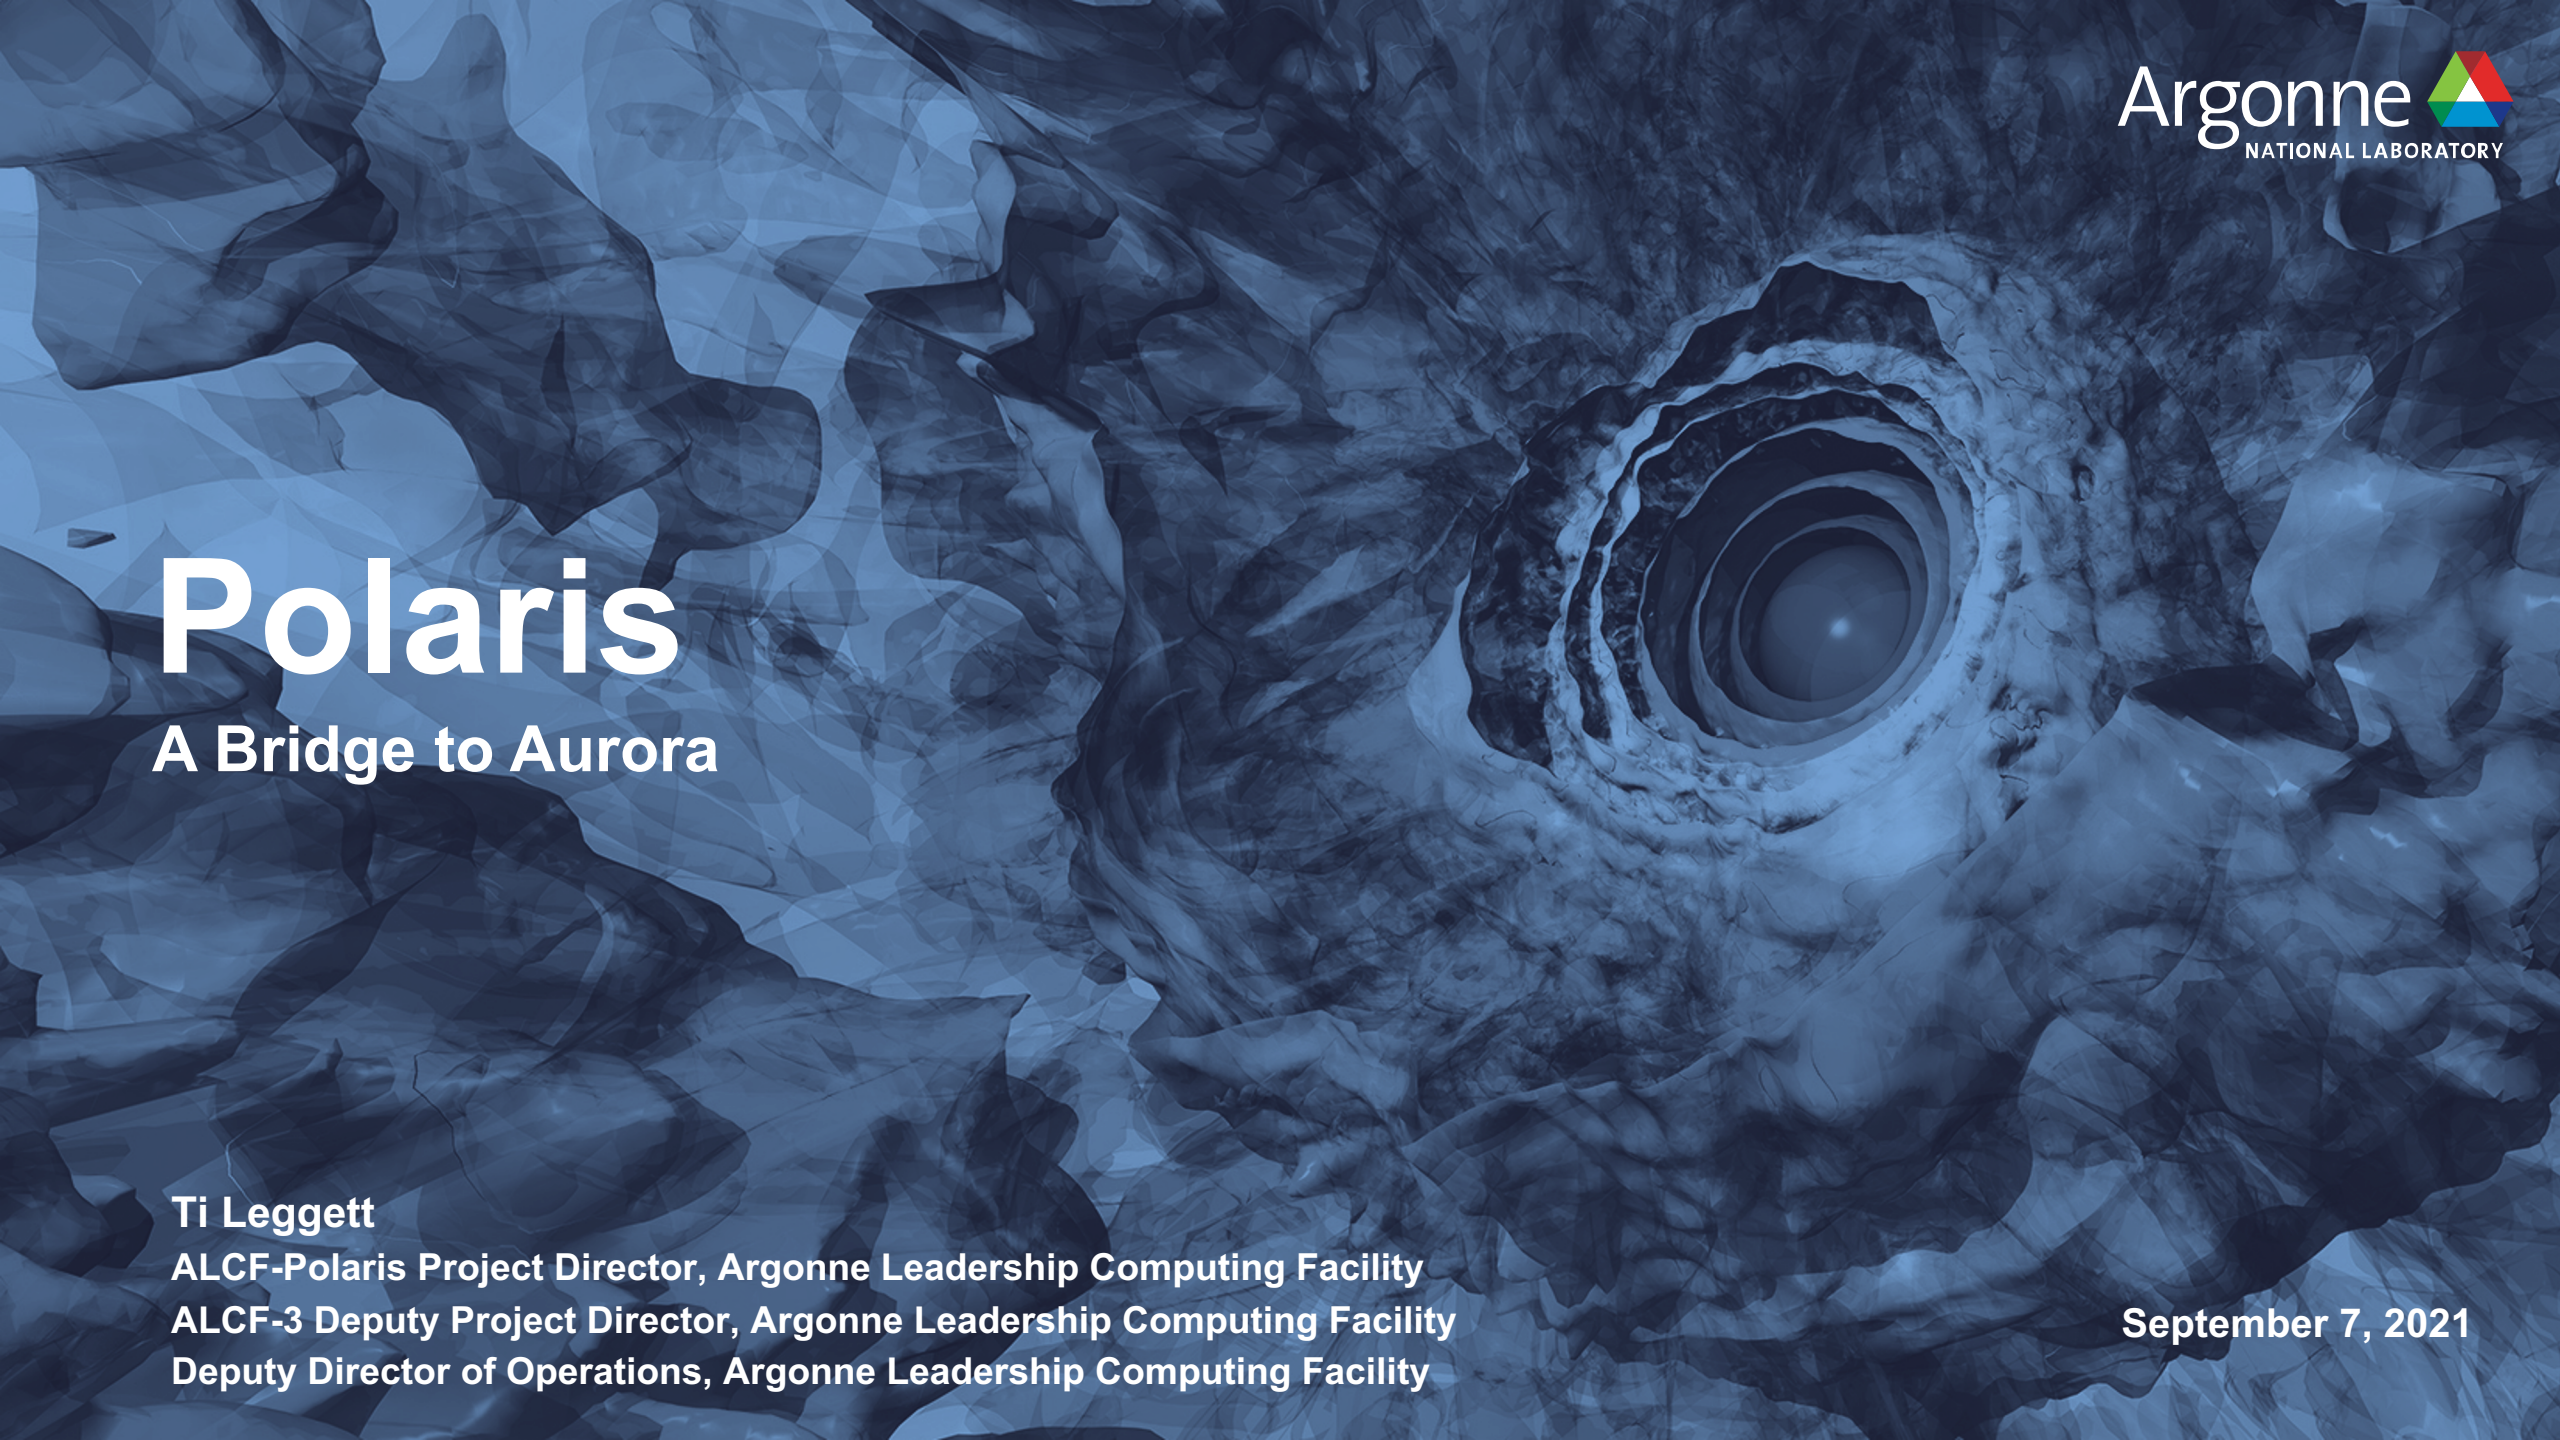

Polaris

A Bridge to Aurora

Ti Leggett
ALCF-Polaris Project Director, Argonne Leadership Computing Facility
ALCF-3 Deputy Project Director, Argonne Leadership Computing Facility
Deputy Director of Operations, Argonne Leadership Computing Facility

September 7, 2021

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

A World-Class Computing Resource Provider

The ALCF enables breakthroughs in science and engineering by providing supercomputing resources and expertise to the research community.

- Users pursue scientific challenges
- In-house experts to help maximize results
- Resources fully dedicated to open science

ALCF Systems Evolution

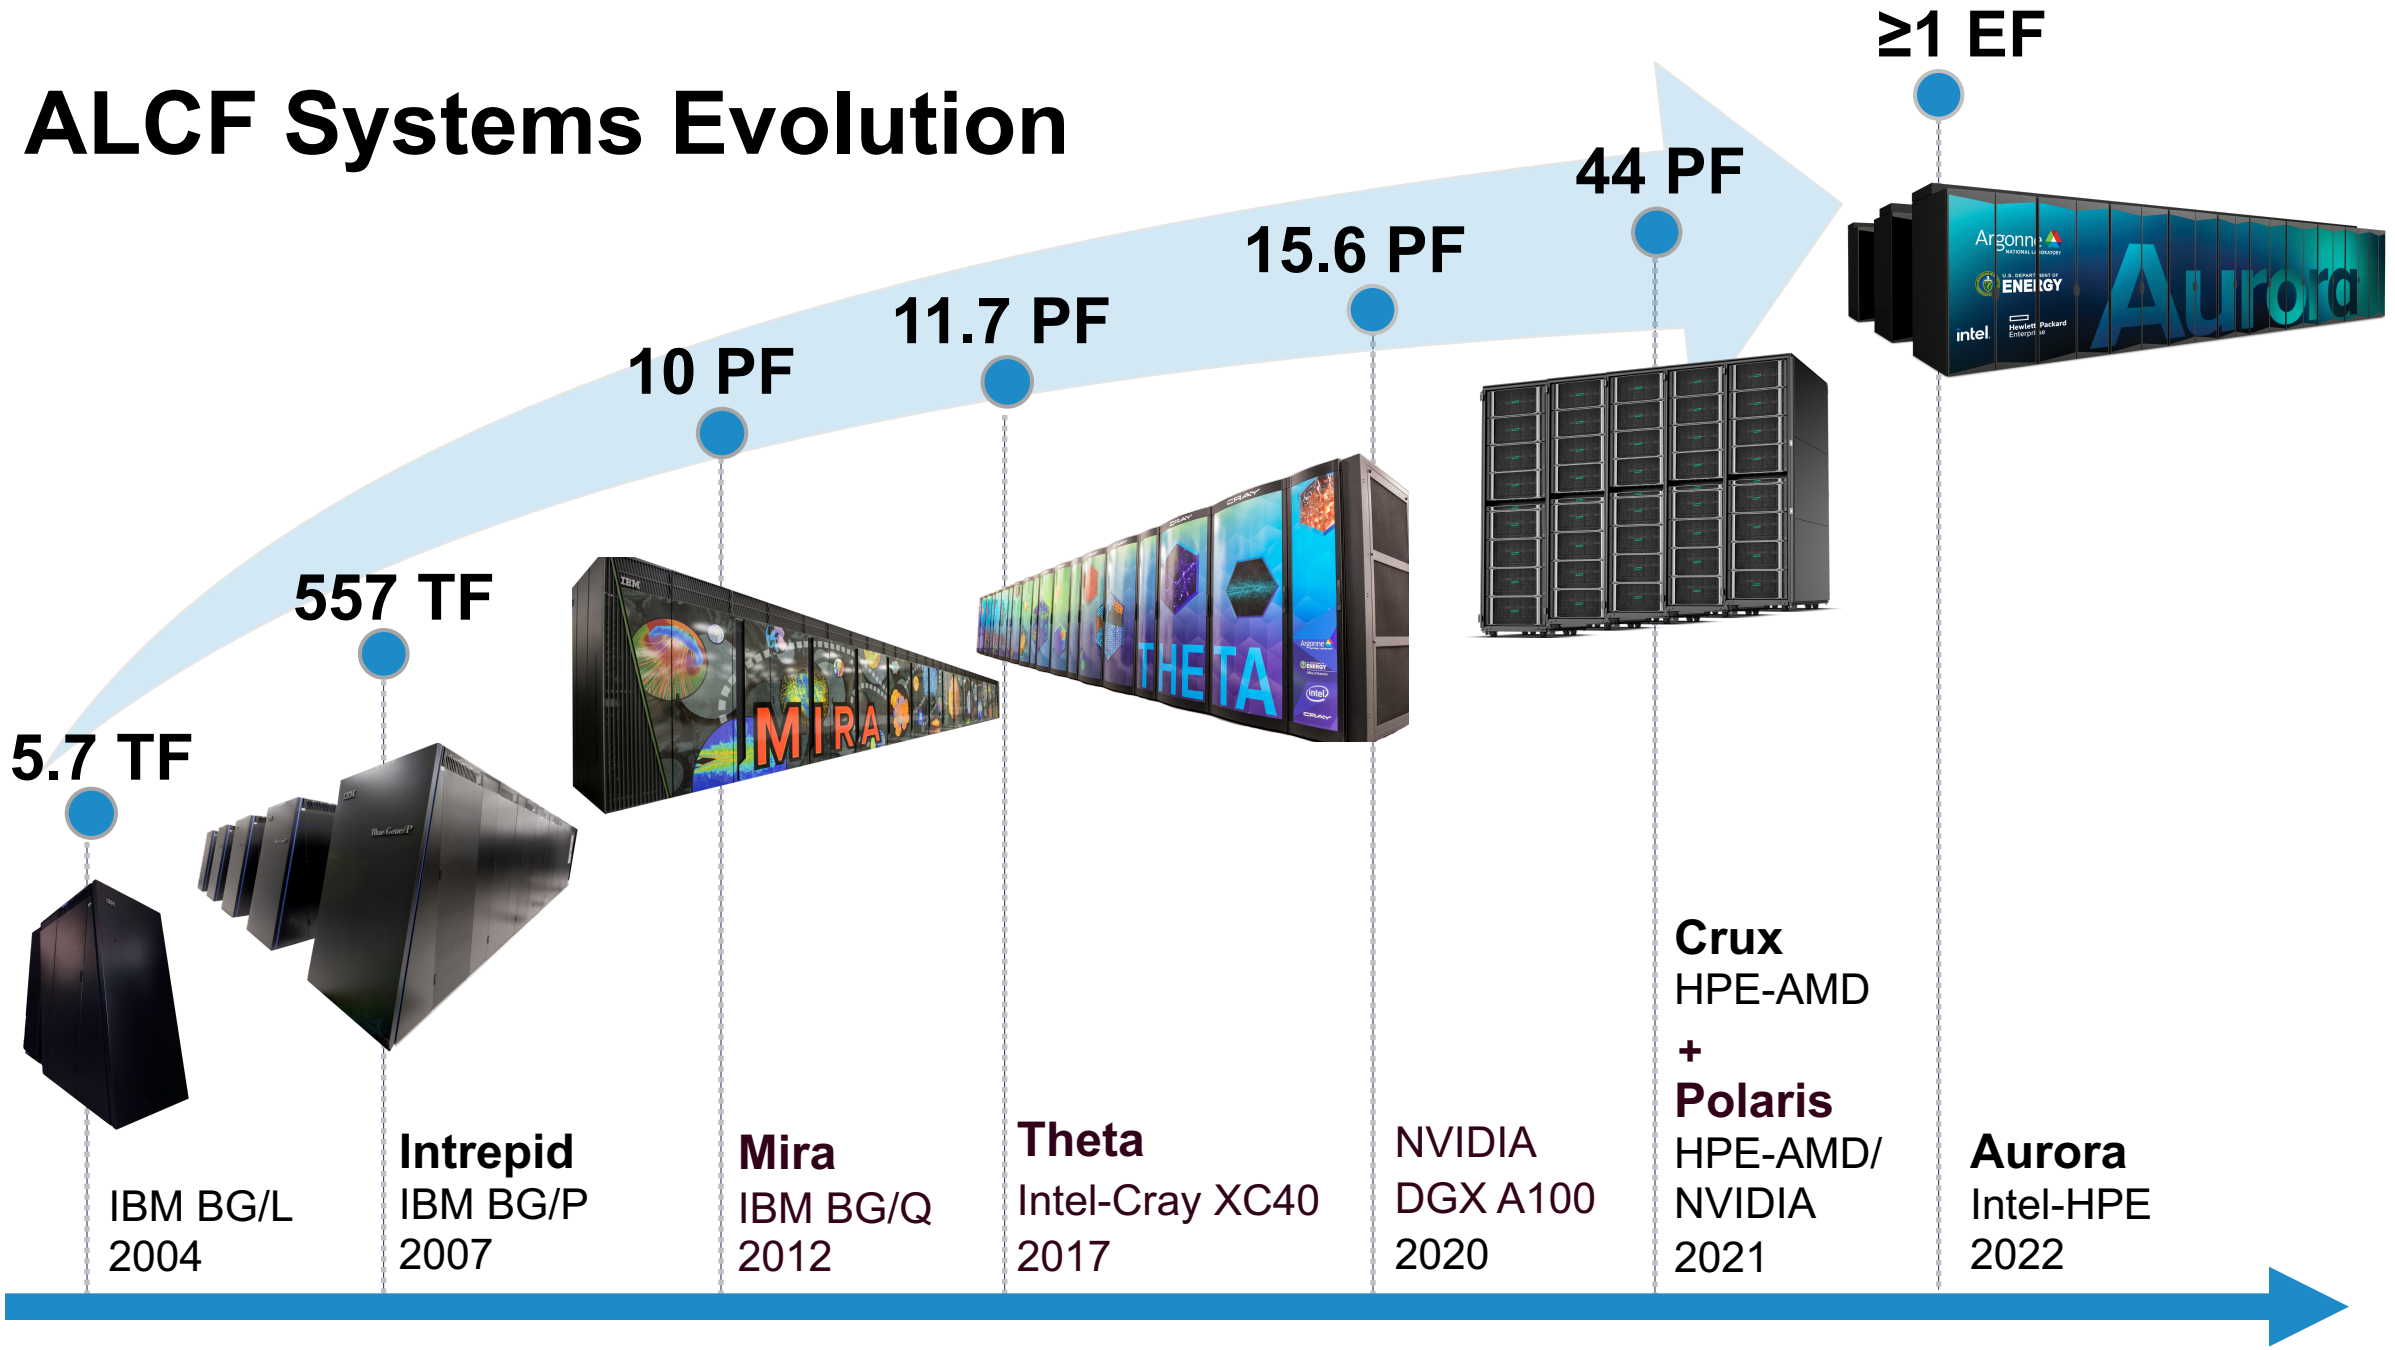

Aurora

Argonne's upcoming exascale supercomputer will leverage several technological innovations to support machine learning and data science workloads alongside traditional modeling and simulation runs.

SUSTAINED PERFORMANCE
≥1 Exaflop DP

Xe ARCHITECTURE-BASED GPU
Ponte Vecchio

INTEL XEON SCALABLE PROCESSOR
Sapphire Rapids

PLATFORM
HPE Cray EX

Compute Node
2 Intel Xeon scalable "Sapphire Rapids" processors; 6 Xe arch-based GPUs; Unified Memory Architecture; 8 fabric endpoints; RAMBO

GPU Architecture
Xe arch-based "Ponte Vecchio" GPU; Tile-based chiplets, HBM stack, Foveros 3D integration, 7nm

CPU-GPU Interconnect
CPU-GPU: PCIe
GPU-GPU: Xe Link

System Interconnect
HPE Slingshot 11; Dragonfly topology with adaptive routing

Network Switch
25.6 Tb/s per switch, from 64–200 Gb/s ports (25 GB/s per direction)

High-Performance Storage
≥230 PB, ≥25 TB/s (DAOS)

Programming Models
Intel oneAPI, MPI, OpenMP, C/C++, Fortran, SYCL/DPC++

Node Performance
>130 TF

System Size
>9,000 nodes

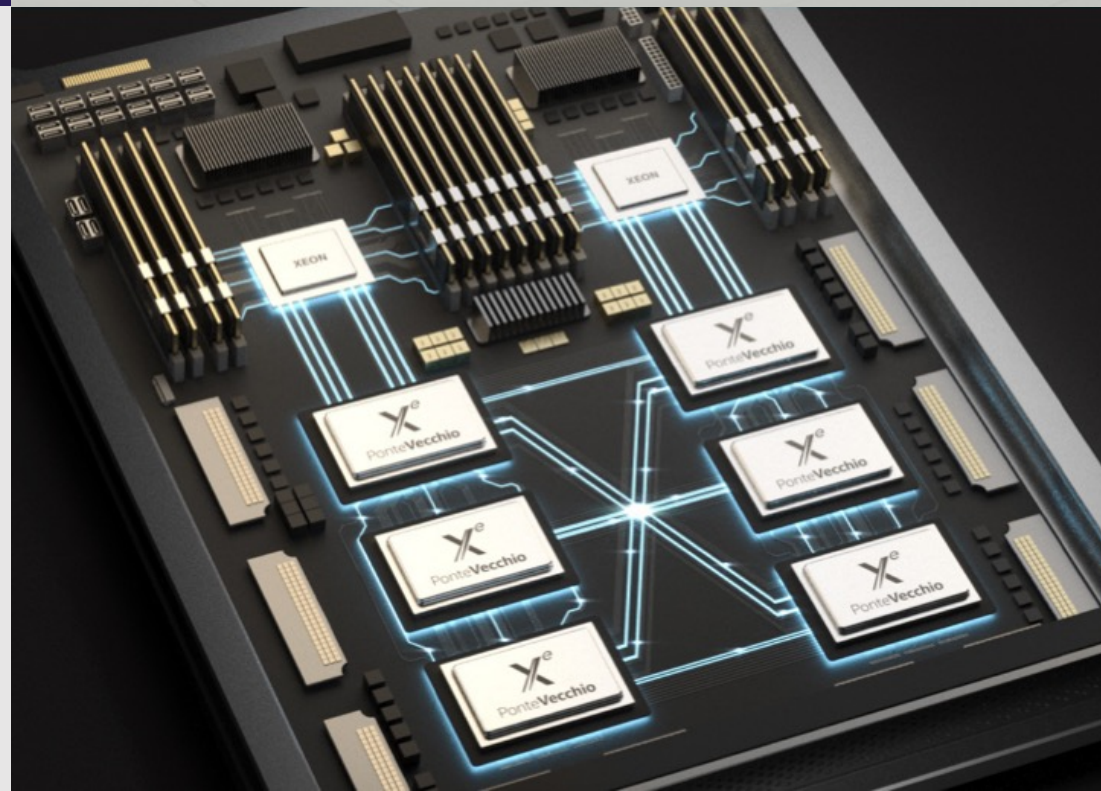

Preparing Users for Exascale

Early Science Program (ESP)

- ALCF conducts ESP to ensure the facility's next-generation systems are ready for science on day one

- Provides research teams with critical pre-production computing time and resources
  - prepares applications for the architecture and scale of a new supercomputer
  - solidifies libraries and infrastructure for other
  - production applications to run on the system

Polaris Overview

Polaris

Polaris will provide a platform utilizing several of the Aurora technologies and similar architectures to provide ALCF staff and users a platform for early scaling and testing purposes.

PEAK PERFORMANCE

44 Petaflop DP

NVIDIA GPU

A100

AMD EPYC PROCESSOR

Milan 7532

PLATFORM

HPE Apollo Gen10+

Compute Node
1 AMD EPYC 7532* processor; 4 NVIDIA A100 GPUs; Unified Memory Architecture; 2 fabric endpoints; 2 NVMe SSDs

GPU Architecture
NVIDIA A100 GPU; HBM stack

CPU-GPU Interconnect
CPU-GPU: PCIe
GPU-GPU: NVLink

System Interconnect
HPE Slingshot 10*; Dragonfly topology with adaptive routing

Network Switch
25.6 Tb/s per switch, from 64–200 Gb/sports (25 GB/s per direction)

Programming Models
CUDA, MPI, OpenMP, C/C++, Fortran, DPC++

Node Performance
78 TF

System Size
560 nodes

*Initial technology to be upgraded later

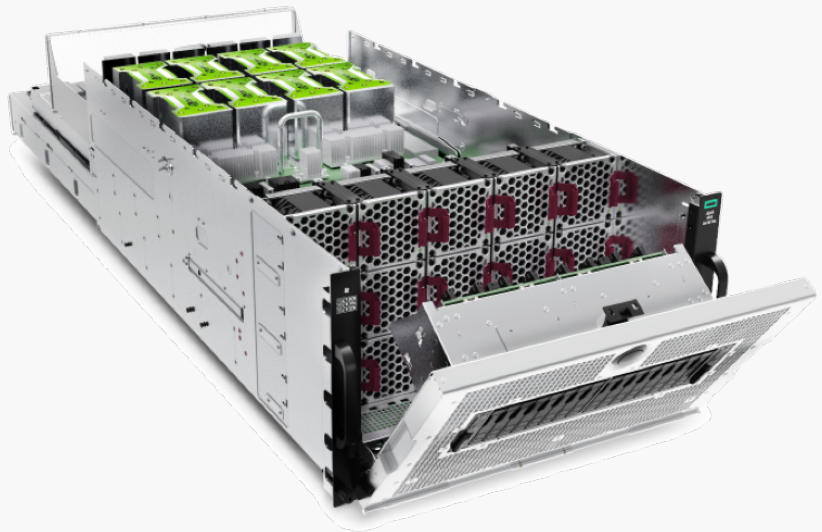

Polaris Single Node Configuration

# of AMD EPYC 7532 CPUs	1
# of NVIDIA A100 GPUs	4
Total HBM2 Memory	160 GB
HBM2 Memory BW per GPU	1.6 TB/s
Total DDR4 Memory	512 GB
DDR4 Memory BW	204.8 GB/s
# OF NVMe SSDs	2
Total NVMe SSD Capacity	3.2 TB
# of Cassini NICs	2
Total Injection BW (w/ Cassini)	50 GB/s
PCIe Gen4 BW	64 GB/s
NVLink BW	600 GB/s
Total GPU DP Tensor Core Flops	78 TF

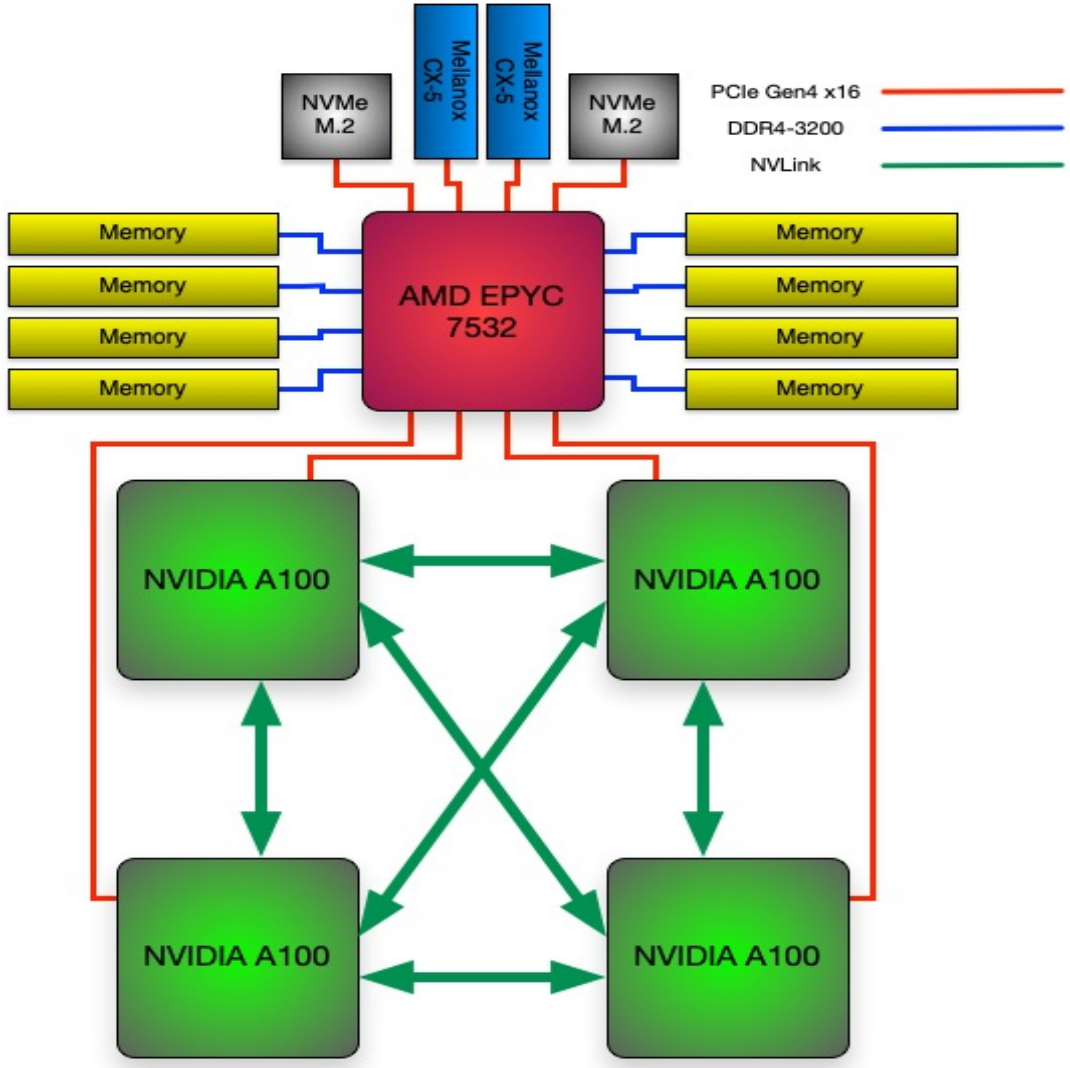

Single AMD EPYC “ROME” 7532 CPU Specs

Base Frequency	2.4 GHz
Max Boost Clk	3.3 GHz
# of Zen2 Cores	32
# of Threads	64
Total DDR4 Memory	512 GB
# of Memory Channels	8
DDR4 Memory BW	204.8 GB/s
Total Shared L3 Cache	256 MB
L2 Cache per Core	512 KB
L1 Cache per Core	32 KB
PCIe Gen 4	128 lanes (8 ports)
PCIe Gen4 BW	64 GB/s
TDP	200 W

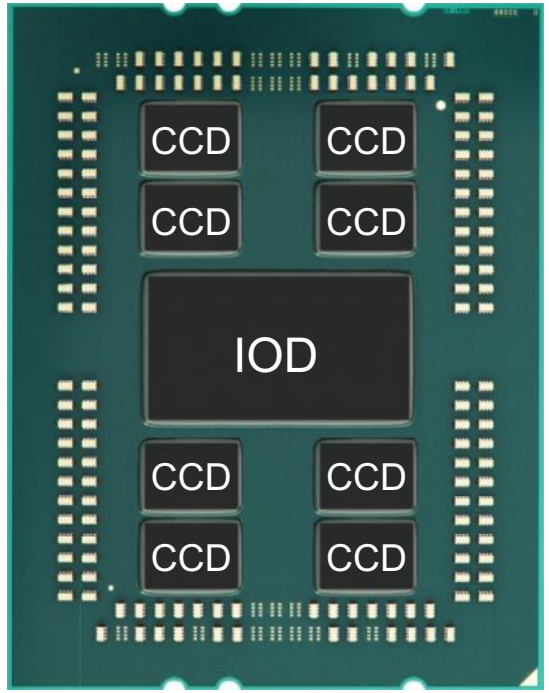

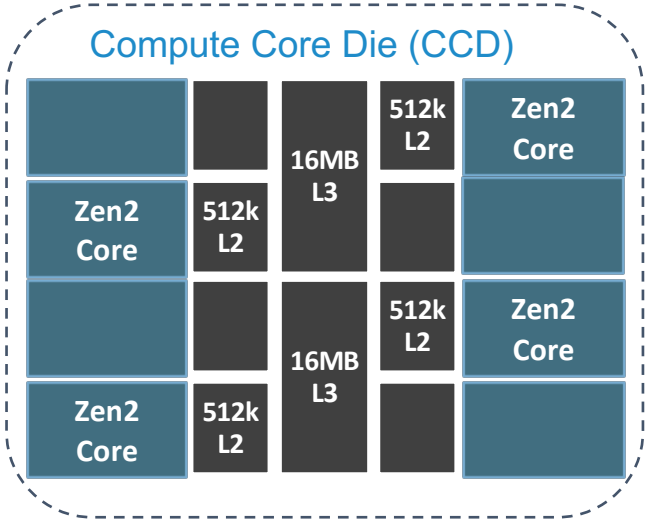

Single AMD EPYC “MILAN” 7543P CPU Specs

Base Frequency	2.8 GHz
Max Boost Clk	3.7 GHz
# of Zen2 Cores	32
# of Threads	64
Total DDR4 Memory	512 GB
# of Memory Channels	8
DDR4 Memory BW	204.8 GB/s
Total Shared L3 Cache	256 MB
L2 Cache per Core	512 KB
L1 Cache per Core	32 KB
PCIe Gen 4	128 lanes (8 ports)
PCIe Gen4 BW	64 GB/s
TDP	225 W

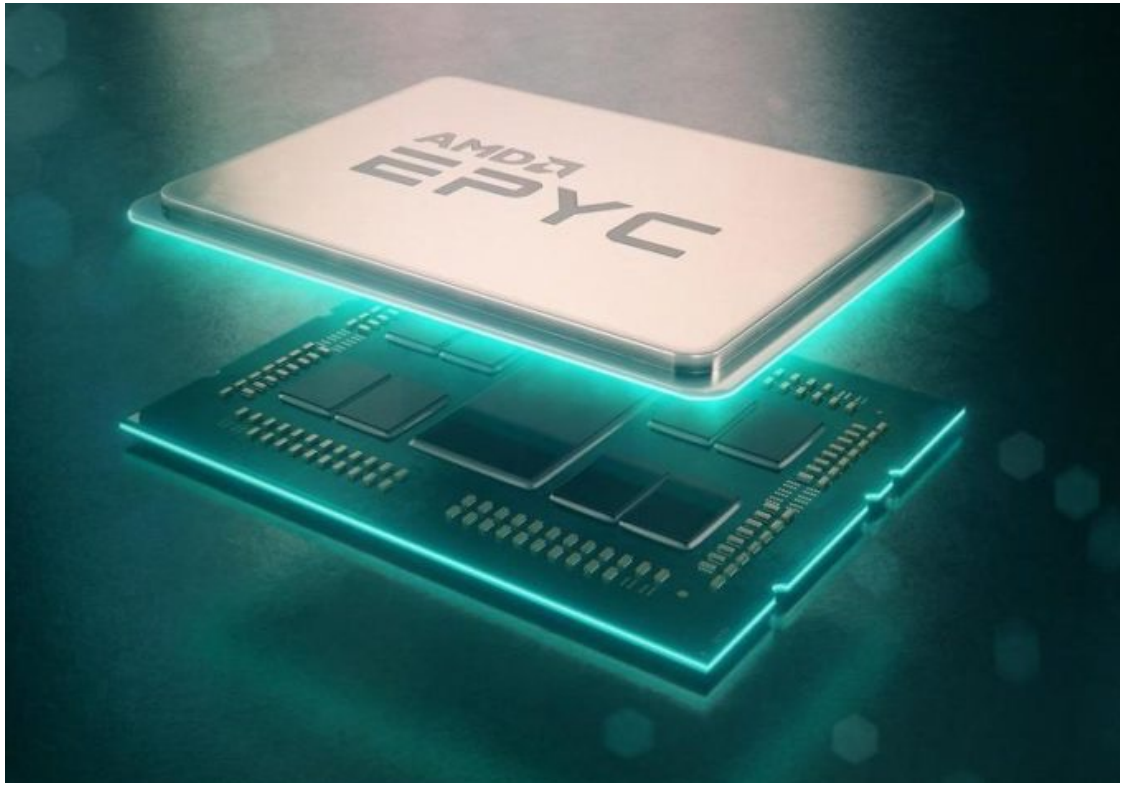

NVIDIA HGX A100 Specs

	A100 PCIe	HGX
FP64	9.7 TF	38.8 TF
FP64 Tensor Core	19.5 TF	78 TF
FP32	19.5 TF	78 TF
BF16 Tensor Core	312 TF	1.3 PF
FP16 Tensor Core	312 TF	1.3 PF
INT8 Tensor Core	624 TOPS	2496 TOPS
GPU Memory	40 GB HBM2	160 GB HBM2
GPU Memory BW	1.6 TB/s	6.4 TB/s
Interconnect	PCIe Gen4 64 GB/s	NVLink 600 GB/s
Max TDP Power	250W	400W

Ampere 7nm

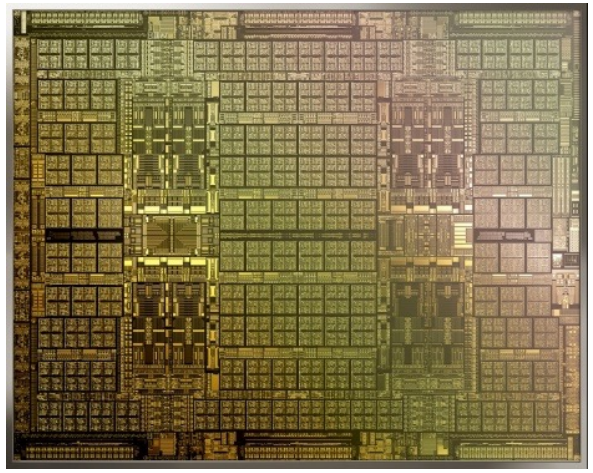

A100 PCIe

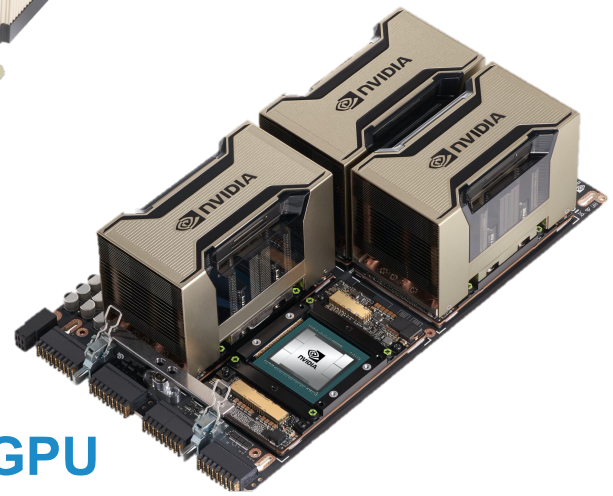

HGX A100 4-GPU

Slingshot Interconnect

Rosetta Switch

- Multiple QoS levels
- Aggressive adaptive routing
- Advanced congestion control
- Very low average and tail latency
- High performance multicast and reduction

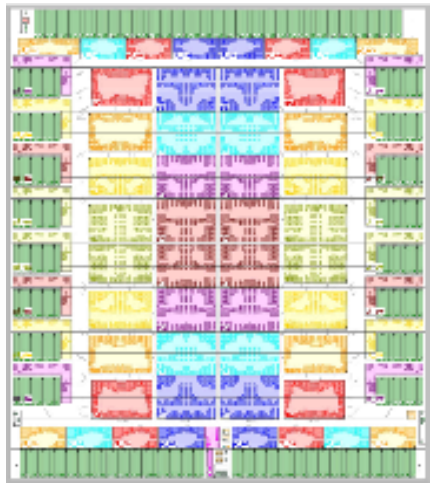

64 ports x 200 Gbps

SS-10 (100Gb)
Injection: ~14 TB/s
Bisection: ~24 TB/s

SS-11 (200Gb)
Injection: ~28 TB/s
Bisection: ~24 TB/s

Mellanox ConnectX NIC

Slingshot 10

- HPE Cray MPI stack
- Ethernet functionality
- RDMA offload

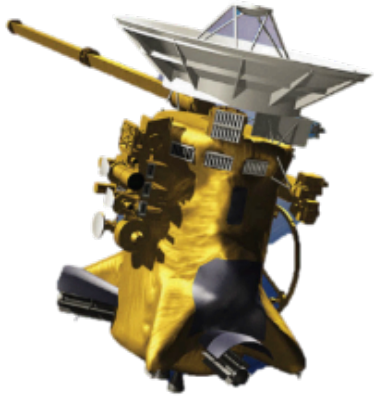

Cassini NIC

Slingshot 11

- MPI hardware tag matching
- MPI progress engine
- One-sided operations
- Collectives
- 2X injection bandwidth

Slingshot Configuration

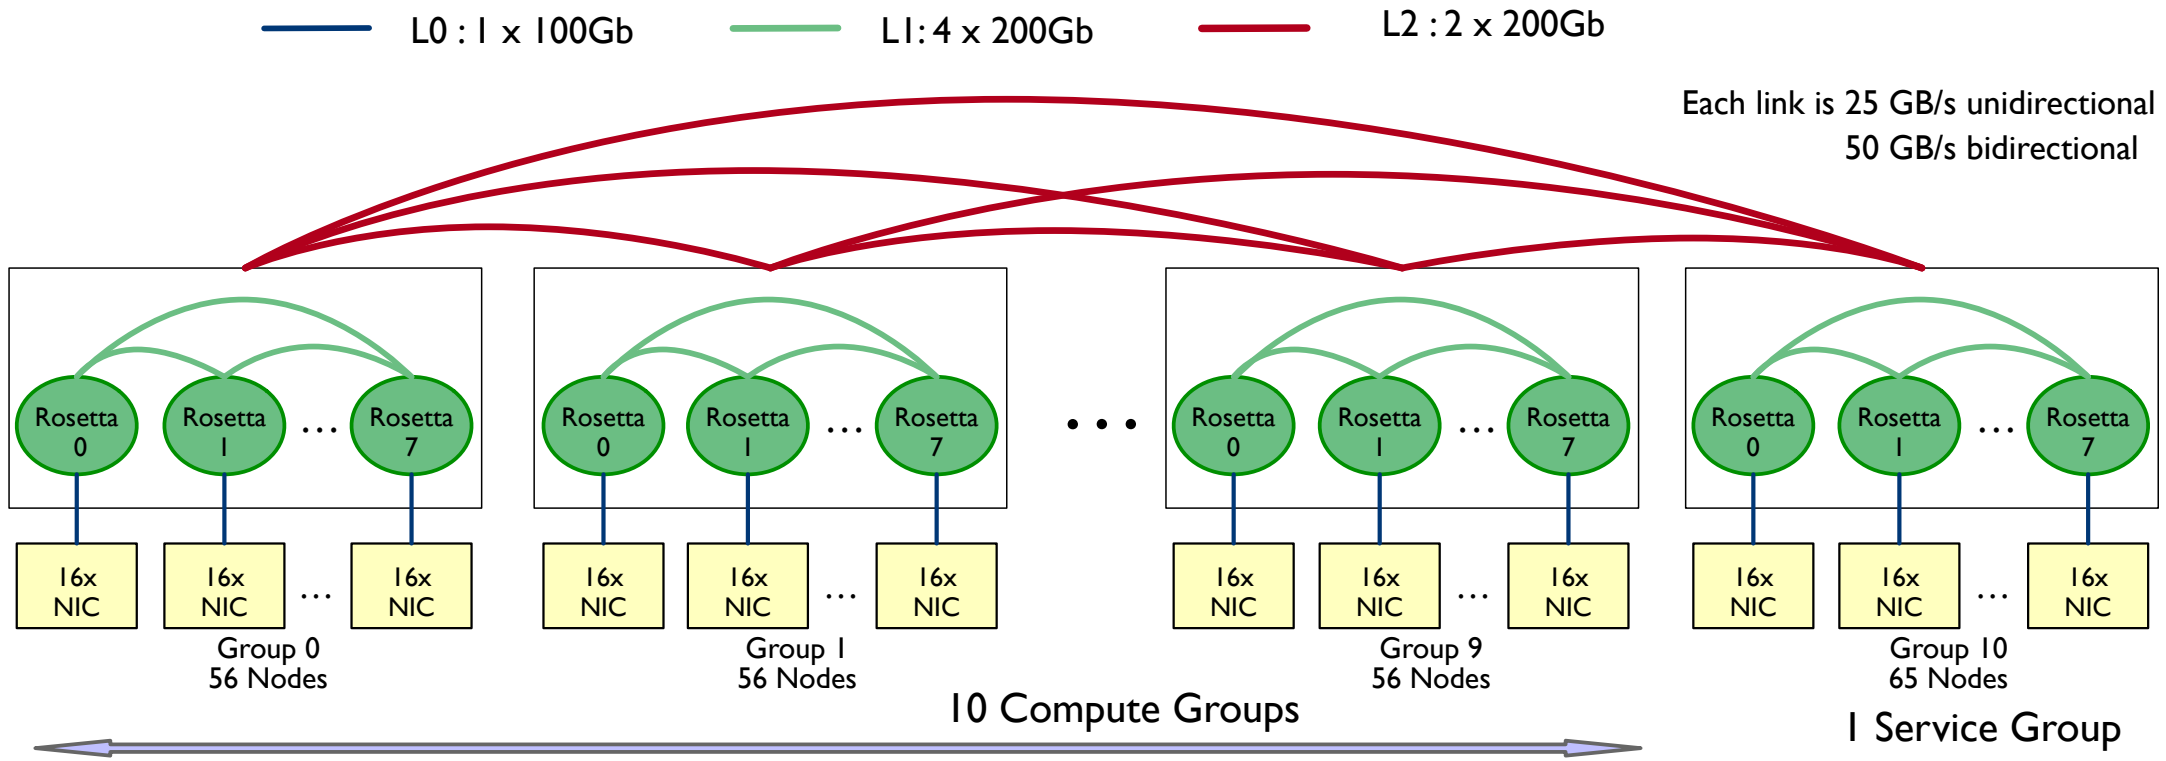

- 11 Total dragonfly groups, 10 compute groups and 1 non-compute group
- 2 links/arc between each group
- 4 links/arc within each group (between switches of a group)
- 1 link from each NIC (100Gb with SS10, 200Gb when upgraded to SS11)

Storage

Polaris will be connected to existing ALCF storage resources

- Grand – Global/Center-wide file system providing main project storage
  - 100 PB @ 650 GB/s
  - Accessed via Lustre LNET routers using Polaris gateway nodes
  - Gateway nodes can provide >1 GB/s
- Eagle – Community file system providing project storage that can be shared externally via Globus sharing
  - 100 PB @ 650 GB/s
  - Accessed via Lustre LNET routers using Polaris gateway nodes
- Home – shared home file system for convenience not for performance or bulk storage

Polaris Software Overview

- Provides an excellent platform for preparing application codes for Aurora
  - All programming models available on Aurora can be tested
  - Features HPE Cray Programming Environment (PE)
  - Built with HPE HPCM system software

- Provides excellent capabilities in simulation, data and learning using NVIDIA's existing HPC SDK

- Support for HPE Cray MPI and MPICH via *libfabric* using Slingshot provider
  - Initial SS10 feature support
  - Later, full SS11 feature support for testing all MPI features available on Aurora

Programming Environment

- HPE Cray PE for Polaris
  - HPE Cray MPI support for PGI offload to A100 for Multi-NIC and Multi-GPU support
  - Full Rome and Milan support
- NVIDIA HPC SDK will provide primary support for programming A100
- SYCL/Data Parallel C++ provided via
  - CodePlay computecpp compiler with NVIDIA support
  - LLVM via Intel DPC++ branch which supports offload to NVIDIA GPUs as well as Intel GPUs

Programming Models

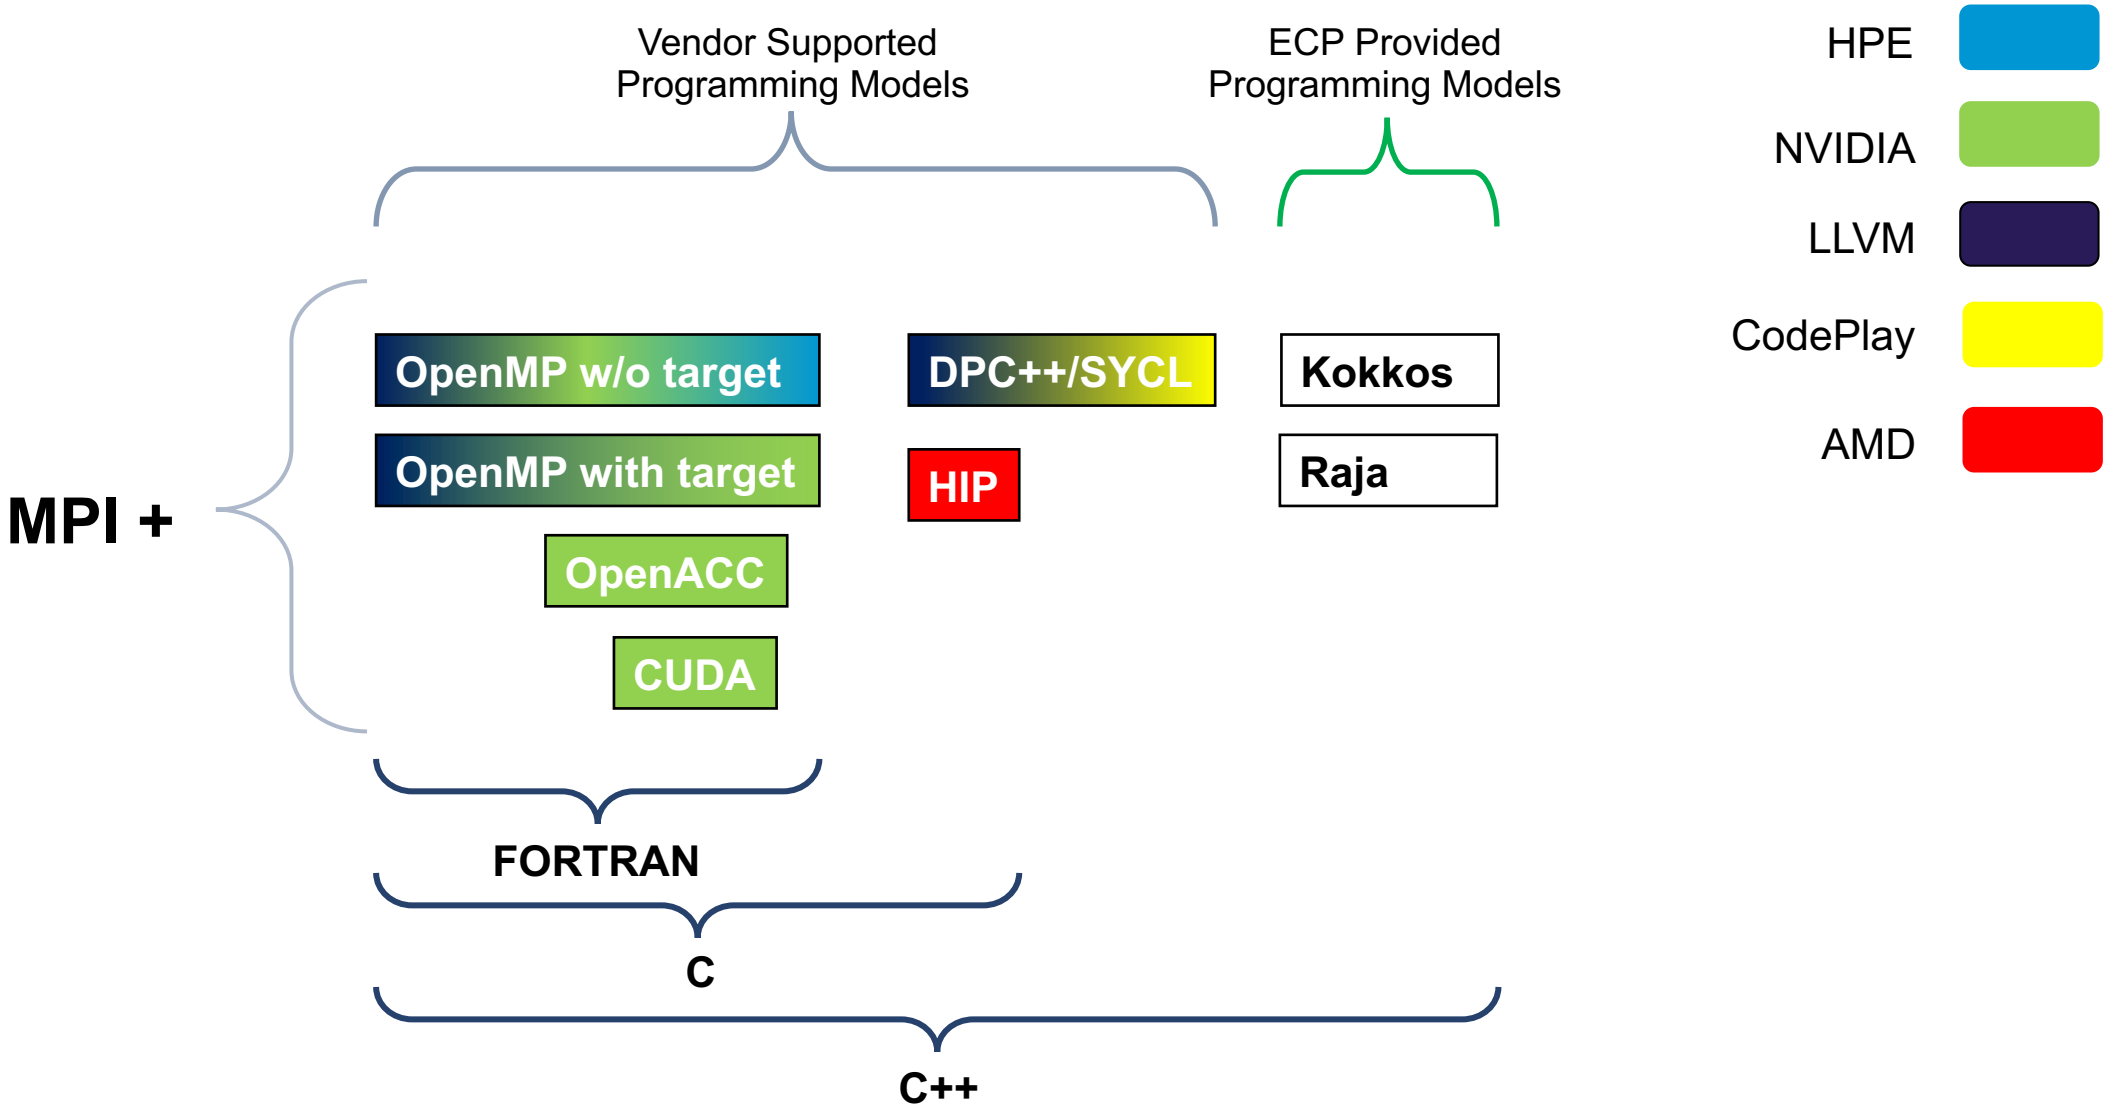

Data Analytics And Learning

- Deep Learning
  - Tensorflow and PyTorch are optimized with CuDNN.
    - Other DL frameworks available and optimized by NVIDIA libraries.
  - Distributed data-parallel deep learning, using libraries such as Horovod, optimized to leverage the NCCL communication library from NVIDIA
- Data Analytics
  - NVIDIA RAPIDS library for data science and analytics
  - Python environment and Numba for GPU using the CUDA backend
- Cray/HPE AI Stack
  - Hyper-parameter optimization (HPO) library
  - Jupyter notebooks
  - SmartSim library for in-situ coupling of simulations with analytics and learning
- Spark BigData Stack with modules optimized for GPU
- Container environment based on Singularity

HPE Performance Cluster Manager (HPCM)

HPCM delivers an integrated system management solution for Linux®-based HPC clusters. HPCM provides complete provisioning, management, and monitoring for clusters scaling up to Exascale levels.

- Currently deployed on many HPE clusters today
- Supports integration with a wide variety of existing schedulers and monitoring tools
- Support integration with Slingshot through the external Slingshot fabric manager

Bridge to Aurora

Component	Polaris	Aurora
System Software	**HPCM**	**HPCM**
Programming Models	**MPI, OpenMP, DPC++, Kokkos, RAJA, HIP,** CUDA, OpenACC	**MPI, OpenMP, DPC++, Kokkos, RAJA, HIP**
Tools	**PAT, gdb, ATP,** NVIDIA Nsight, cuda-gdb	**PAT, gdb, ATP,** Intel VTune
MPI	**HPE Cray MPI, MPICH**	**HPE Cray MPI, MPICH,** Intel MPI
Multi-GPU	1 CPU : 4 GPU	2 CPU : 6 GPU
Data and Learning	**DL frameworks, Cray AI stack, Python/Numba, Spark, Containers,** Rapids	**DL frameworks, Cray AI stack, Python/Numba, Spark, Containers,** oneDAL
Math Libraries	cu* from CUDA	oneAPI

Acknowledgement

This publication used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357.