Oak Ridge Leadership Computing Facility: Summit and Beyond

Justin L. Whitt

OLCF-4 Deputy Project Director,
Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory

March 2017

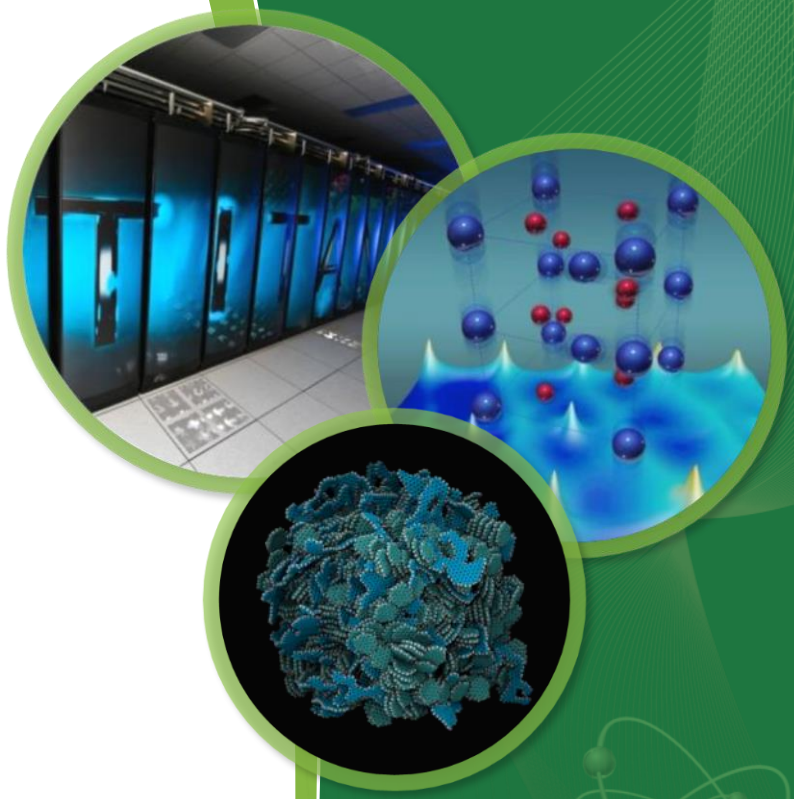

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Summit will replace Titan as the OLCF's leadership supercomputer in 2019

- Many fewer nodes
- Much more powerful nodes
- Much more memory per node and total system memory
- Faster interconnect
- Much higher bandwidth between CPUs and GPUs
- Much larger and faster file system

Feature	Titan	Summit
Application Performance	Baseline	5-10x Titan
Number of Nodes	18,688	~4,600
Node performance	1.4 TF	> 40 TF
Memory per Node	38GB DDR3 + 6GB GDDR5	512 GB DDR4 + HBM
NV memory per Node	0	800 GB
Total System Memory	710 TB	>6 PB DDR4 + HBM + Non-volatile
System Interconnect (node injection bandwidth)	Gemini (6.4 GB/s)	Dual Rail EDR-IB (23 GB/s)
Interconnect Topology	3D Torus	Non-blocking Fat Tree
Processors	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	32 PB, 1 TB/s, Lustre®	250 PB, 2.5 TB/s, GPFS™
Peak power consumption	9 MW	13 MW

Summit Project Scope

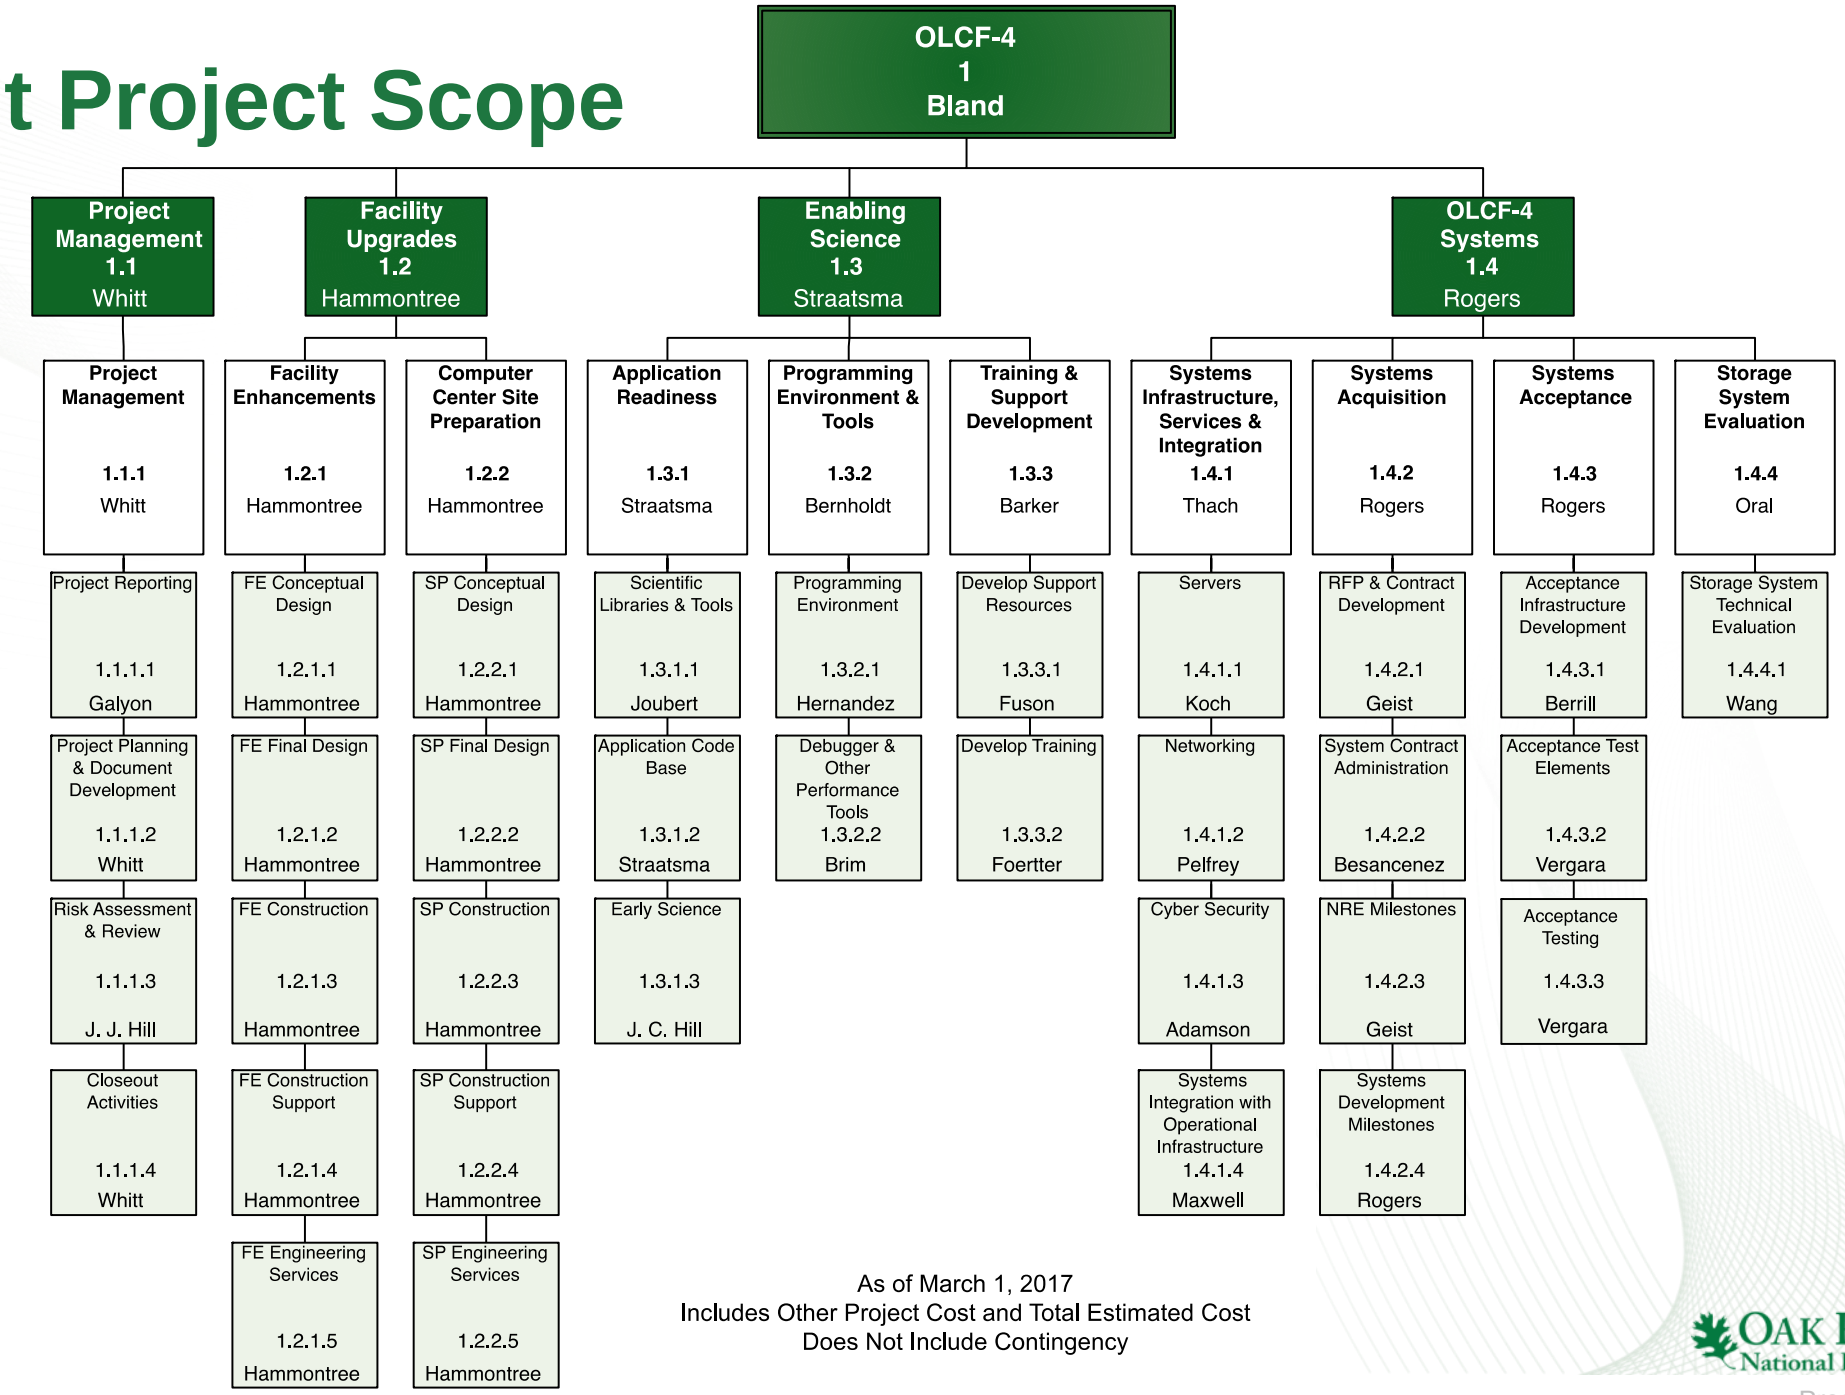

As of March 1, 2017
Includes Other Project Cost and Total Estimated Cost
Does Not Include Contingency

Facility Enhancements

- 13 MW power
- 20 MW of cooling capacity
- Preparation of the bare room for the computers
- Electrical distribution
- Cooling Water distribution
- Fire protection
- Controls systems

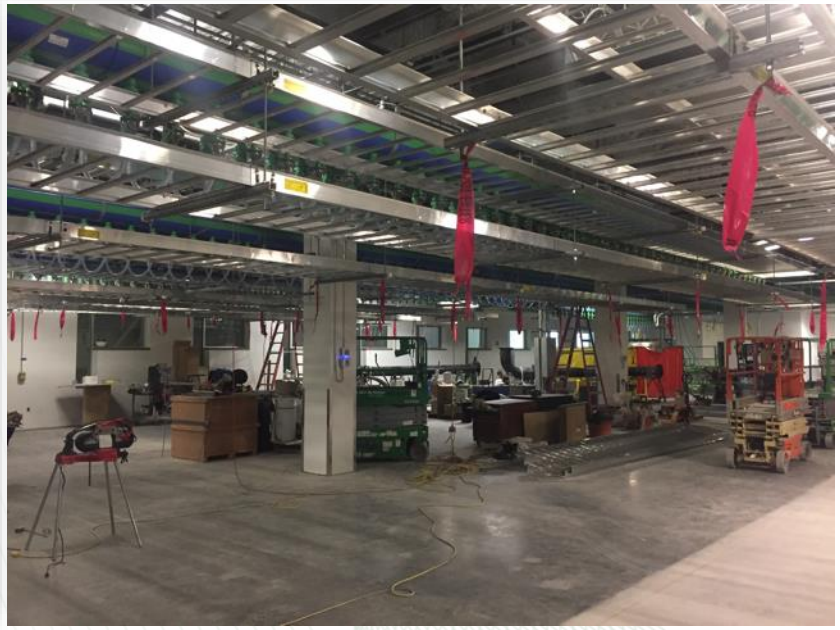

Application Readiness for Summit

The **Center of Accelerated Application Readiness (CAAR)** remains the OLCF's forward facing program to facilitate application readiness on evolving architectures

- ***Build on the experience of a successful application readiness program for OLCF-3 (Titan)***
- ***Thirteen CAAR projects were selected after a call for proposals***
  - Partnership: Application Developers,
    OLCF Scientific Computing staff
    Vendor Center of Excellence

- ***Resources available to CAAR projects***
  - Dedicated collaboration with OLCF Scientific Computing staff
  - Support and consultation from other OLCF staff and vendor Center of Excellence
  - Access to early test systems
  - Eight associated postdoctoral fellow in CSEEN program associated with CAAR projects
  - Allocations to available compute resources at OLCF, ALCF and NERSC in ALCC program
  - Early Science allocations

CAAR Applications

Domain	Application	Methods	PI	Institution	Related to INCITE/ALCC	Related to SciDAC
Astrophysics	FLASH	Grid, AMR	Bronson Messer	ORNL	Zingale	SciDAC II
Chemistry	DIRAC	Particle, LA	Lucas Visscher	VUA	Dixon	
Climate Science	ACME (N)	Unstr Mesh	David Bader	LLNL	Taylor	SciDAC III
Engineering	RAPTOR	Kokkos	Joseph Oefelein	SNL	Oefelein	SciDAC II
Materials Science	QMCPACK	MC	Paul Kent	ORNL	Kent, Ceperley	SciDAC III
Nuclear Physics	NUCCOR	Particle	Gaute Hagen	ORNL	Vary	SciDAC III
Plasma Physics	XGC (N)	PIC, PETSc	CS Chang	PPPL	Chang	SciDAC III
Seismic Science	SPECFEM	Unstr Mesh	Jeroen Tromp	Princeton	Tromp	
Astrophysics	HACC(N,A)	Grid	Salman Habib	ANL	Habib	SciDAC III
Biophysics	NAMD (N)	Particle	Klaus Schulten	UIUC	Klein, Schulten, Tajkhorshid	SciDAC II
Chemistry	NWCHEM (N)	Particle, LA	Karol Kowalski	PNNL	Dixon, Sumpter	SciDAC III
Chemistry	LSDALTON	Particle, LA	Poul Jørgensen	Aarhus	Jørgensen	
Plasma Physics	GTC (N)	PIC	Zhihong Lin	UCI	Lin	SciDAC III

N: NERSC application; A: ALCF application

CAAR: Architecture and Performance Portability

ALCF, NERSC and OLCF Joint Activities and Resources

- ALCF, NERSC and OLCF participated in each other's proposal reviews
- ALCC Award to support NESAP, CAAR and ESP
- Common applications teams in NESAP, CAAR and ESP will collaborate
- Leveraging training activities at NERSC, OLCF and ALCF
- SC15 workshop "Portability Among HPC Architectures for Scientific Applications" on Sunday, November 15 was chaired by Tim Williams (ALCF), Katie Antypas (NERSC) and Tjerk Straatsma (OLCF)
- All three ASCR facilities have representation on the standards bodies for programming models that facilitate portability (OpenACC and OpenMP)
- We are working with vendors to provide programming environments and tools that enable portability, as part of CORAL and Trinity procurements
- Organized portability workshops
- Portability Research Project shared between OLCF, ALCF, NERSC and their CoE's

Oakland, September 24-25, 2014

Oak Ridge, January 27-29, 2015

Synergy between Application Readiness Programs

NESAP at NERSC - *NERSC Exascale Science Application Program*

- Call for Proposals – June 2014
- 20+26 Projects selected
- Partner with Application Readiness Team and Intel/Cray
- 8 Postdoctoral Fellows

CAAR at OLCF - *Center for Accelerated Application Readiness*

- Call for Proposals – November 2014
- 13 Projects selected
- Partner with Scientific Computing group and IBM/NVIDIA Center of Excellence
- 8 Postdoctoral Associates

ESP at ALCF - *Early Science Program*

- Call for Proposals – May 2015
- 6 Projects selected in first round
- Partner with Catalyst group and Intel/Cray Center of Excellence
- Postdoctoral Appointee per project

Summit Early Evaluation System

Each IBM S822LC node has:

- 2x IBM POWER8 CPUs
  - 32x 8GB DDR4 memory (256 GB)
  - 10 cores per POWER8, each core with 8 HW threads
- 4x NVIDIA Tesla P100 GPUs
  - NVLink 1.0 connects GPUs at 80 GB/s
  - 16 GB HBM2 memory per GPU
- 2x Mellanox EDR InfiniBand
- 800 GB NVMe storage

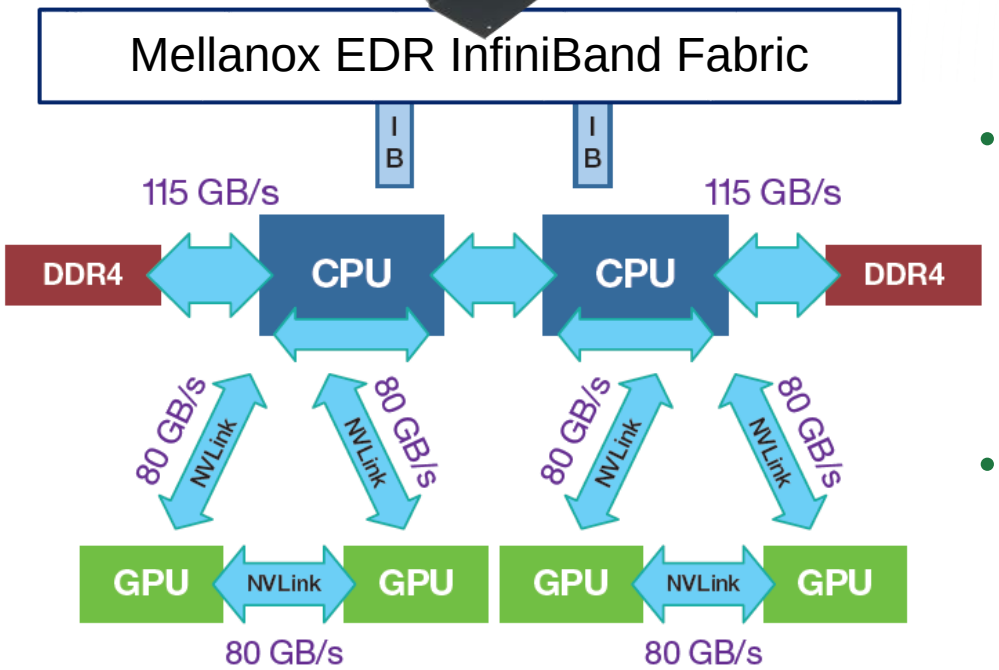

Summit EA System:

- Three racks, each with 18 nodes
- One rack of login and support servers
- Nodes connected in a full fat-tree via EDR InfiniBand
- Liquid cooled w/ heat exchanger rack
- We will get an additional rack to add to Summit EA for Exascale Computing Project testing, giving us a 54 node system
- One additional 18-node rack is for system software testing

Information and drawing from IBM Power System S822LC for High Performance Computing Data Sheet

Spider 3 @ OLCF

Spider 3 is a center-wide single namespace POSIX file system to serve all OLCF resources, eliminating data islands and enabling seamless data sharing between resources

- Built on IBM's Elastic Storage Server and uses Spectrum Scale (formerly known as GPFS) parallel filesystem technology utilizing GPFS Native RAID with 8+2 redundancy
- Provides a usable capacity of 250 PB
- Performs at an aggregate sequential peak read/write bandwidth of 2.5 TB/s
- Performs at an aggregate random peak read/write bandwidth of 2.2 TB/s
- Provides rich metadata performance; single directory parallel create rate of 50,000/s
- Provides rich interactive performance; @32 KiB I/O 2.6 million IOPs
- Disk-based, with tens of thousands of disks
- Connected to OLCF's SION 3 SAN with IB EDR
- Will also serve as the Summit Burst Buffer sink and source on the end-to-end I/O path

OLCF Programming Environment and Tools Focus Areas

Programming Environment

- Directive-based programming
  - **OpenACC** – accelerator offload, unified memory support, error handling
  - **OpenMP** – threading and tasks with accelerator offload under intensive development, memory hierarchy support, tools API, task reductions.
  - **SPEC High-Performance Group** – benchmark suites to drive performance and correctness
- Runtime
  - **MPI** – resilience, collectives, scalability

Tools

- Co-design of hardware and software to ensure maximum capabilities for tools on Summit system
  - **CPU, GPU, memory system, network**
  - **CORAL vendors, labs, tool developers**
- Target tools
  - HPCToolkit, Open|Speedshop, TAU, Valgrind, PAPI, and DynInst
  - Allinea DDT, MAP; Score-P/VAMPIR

- Evaluation of CORAL NRE products, including compilers, tools, and infrastructure
- Support for the Center for Accelerated Applications Readiness (Summit) applications
- Support for current OLCF users (Titan)

Timeline for Summit

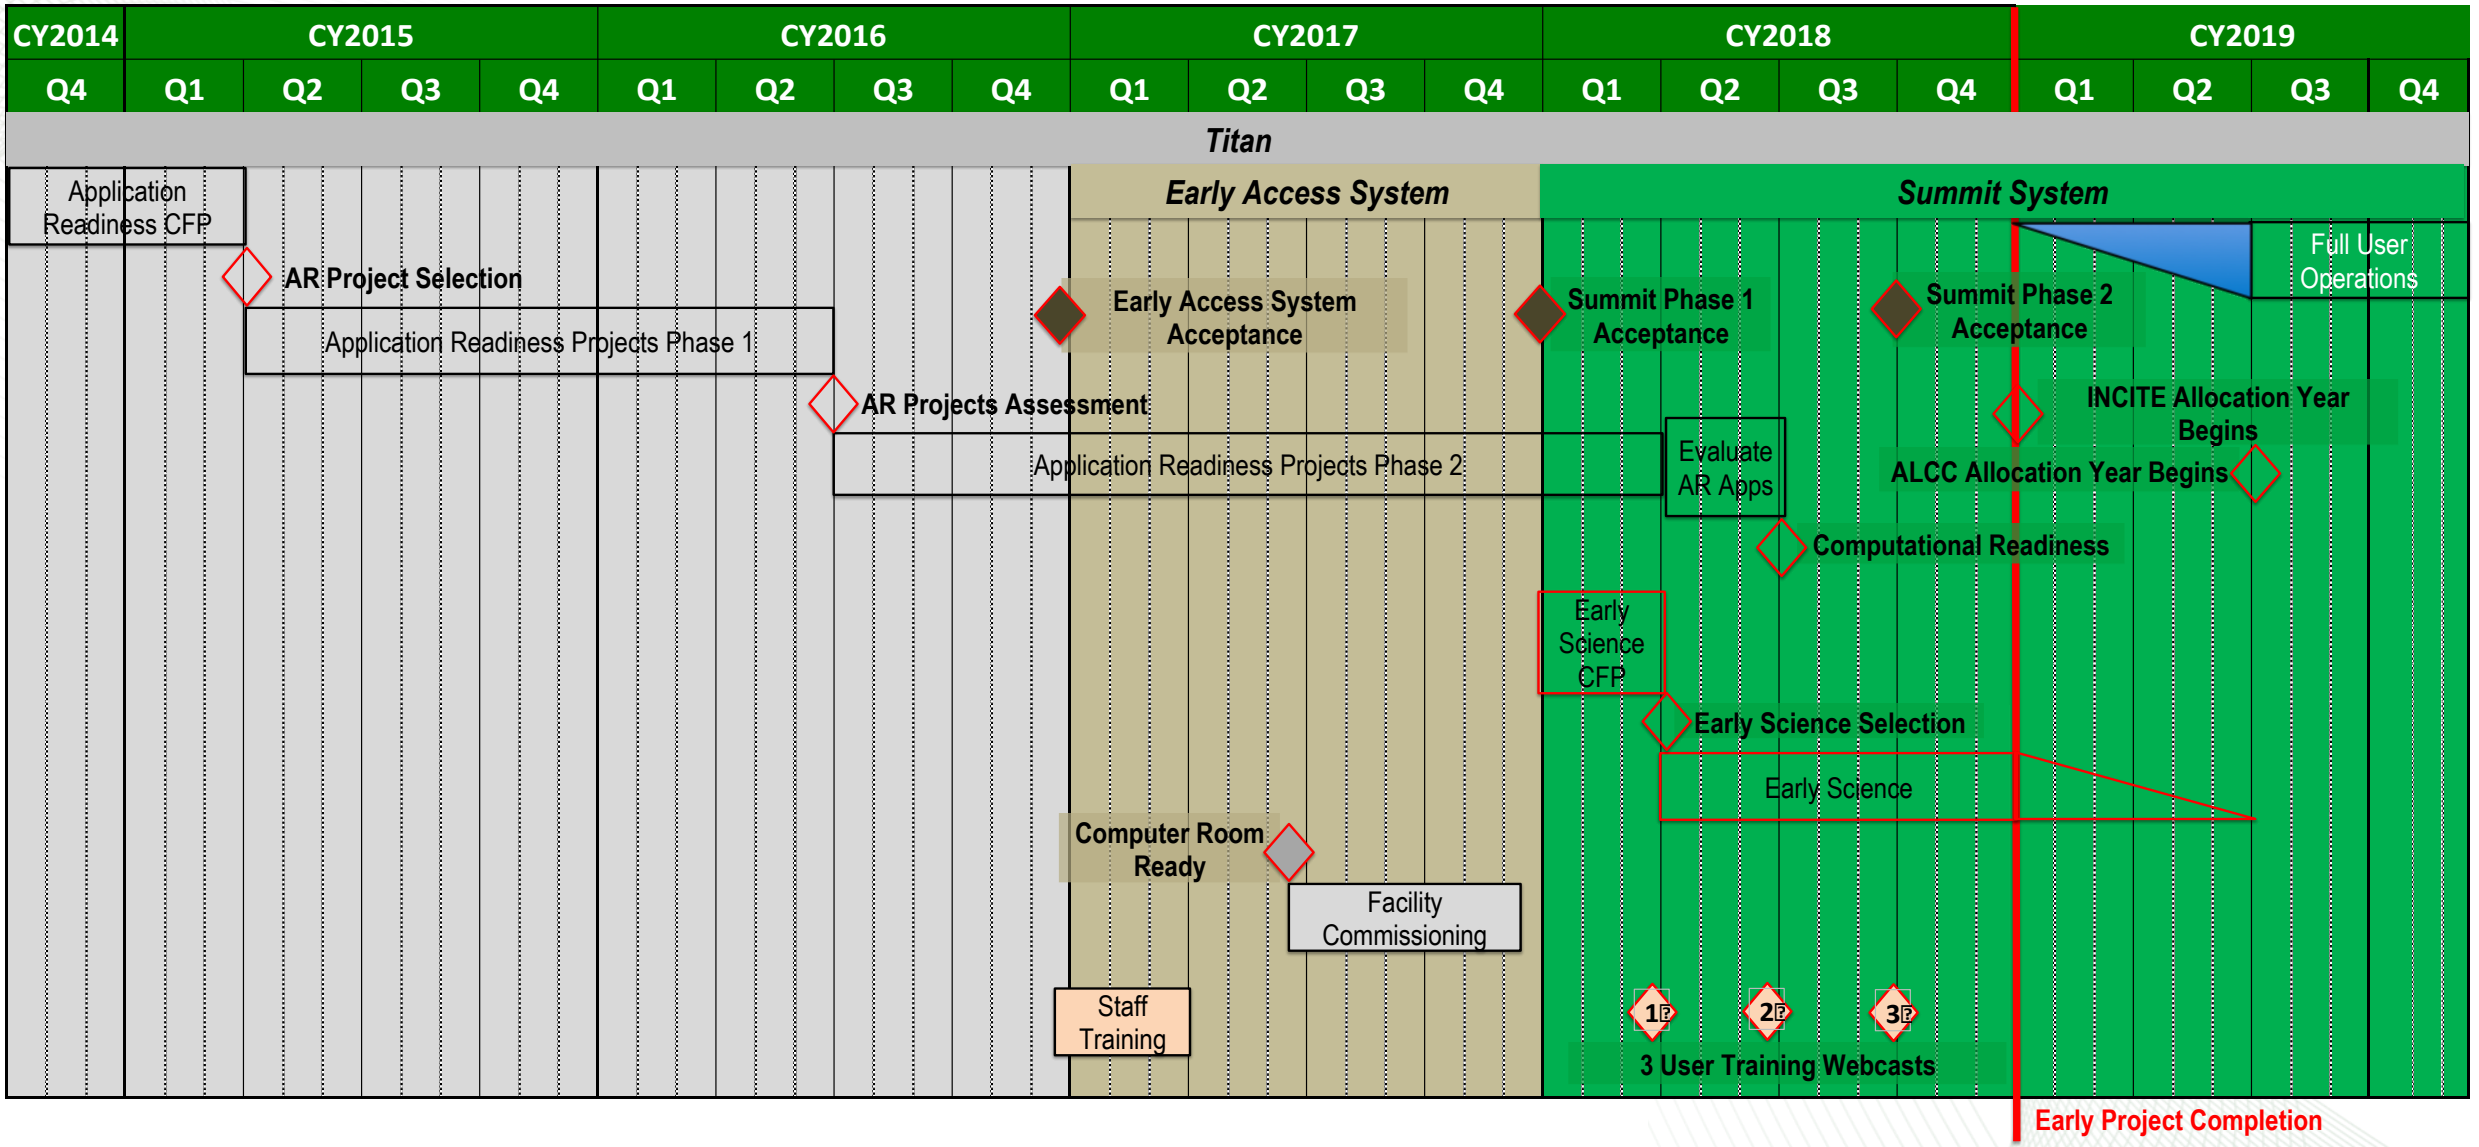

NVIDIA DGX-1

Specifications

- 8 Tesla P100 GPUs
- GPU Mem: 16 GB per GPU
- System Mem: 512 GB
- Storage: 4x 2 TB SSDs
- Out of the box libraries (theano, caffe, cuDNN, cuBLAS, etc)
- Integrated into the CADES environment

Preparation for Summit

- **DGX-1 architecture and use of NVLink similar to Summit-dev**
  - Ease the preparation for DNN training at larger scales

DGX-1

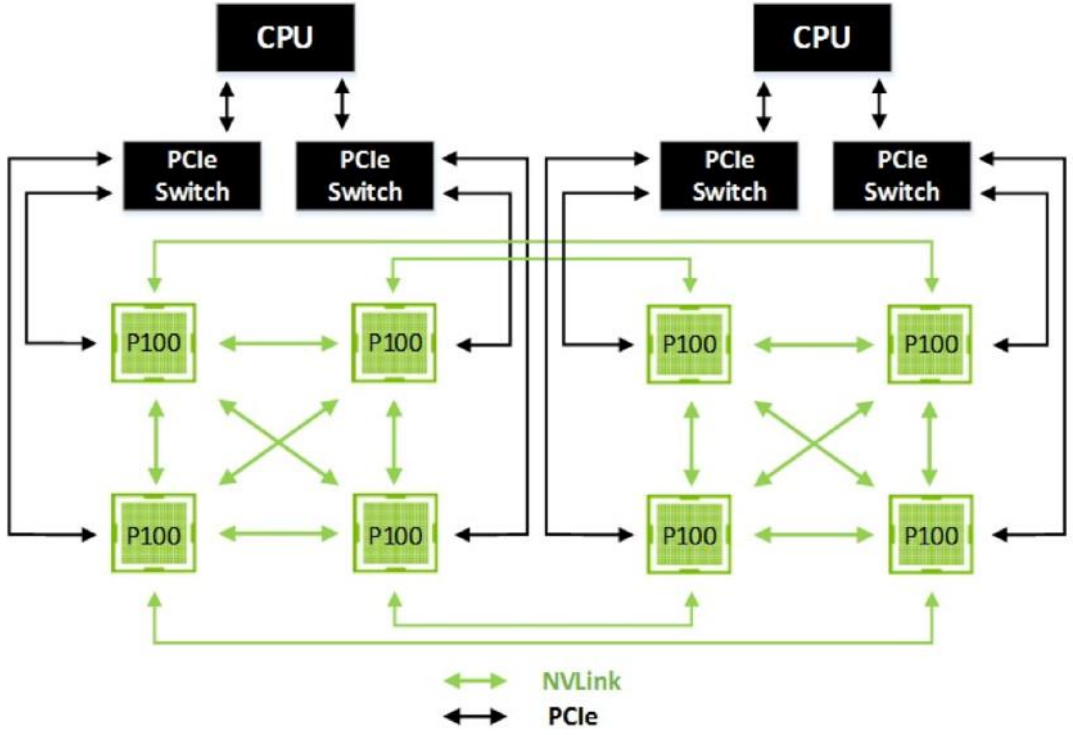

Summit-dev

(IBM's S822LC)

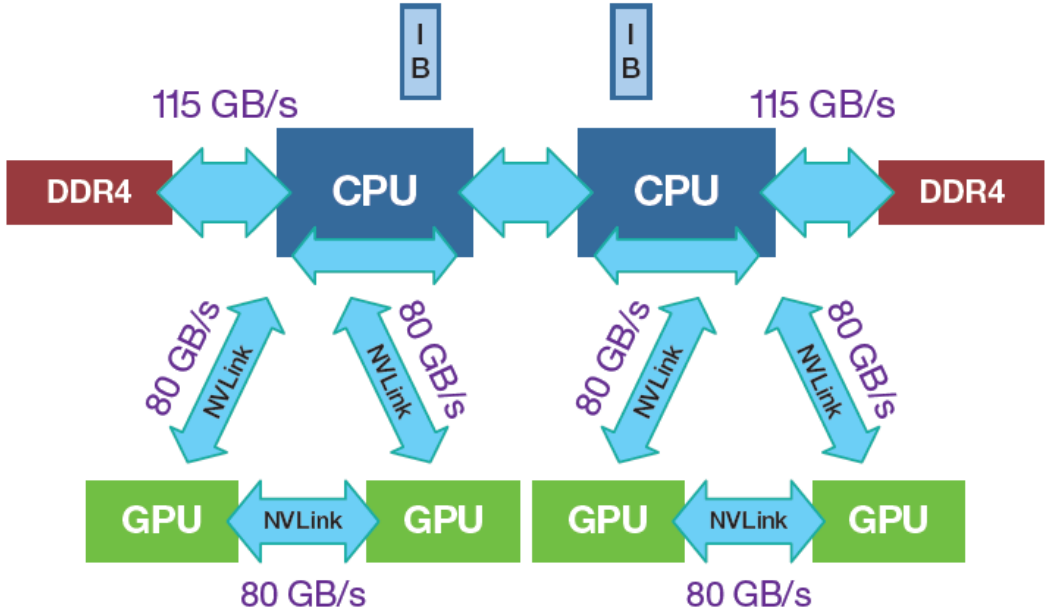

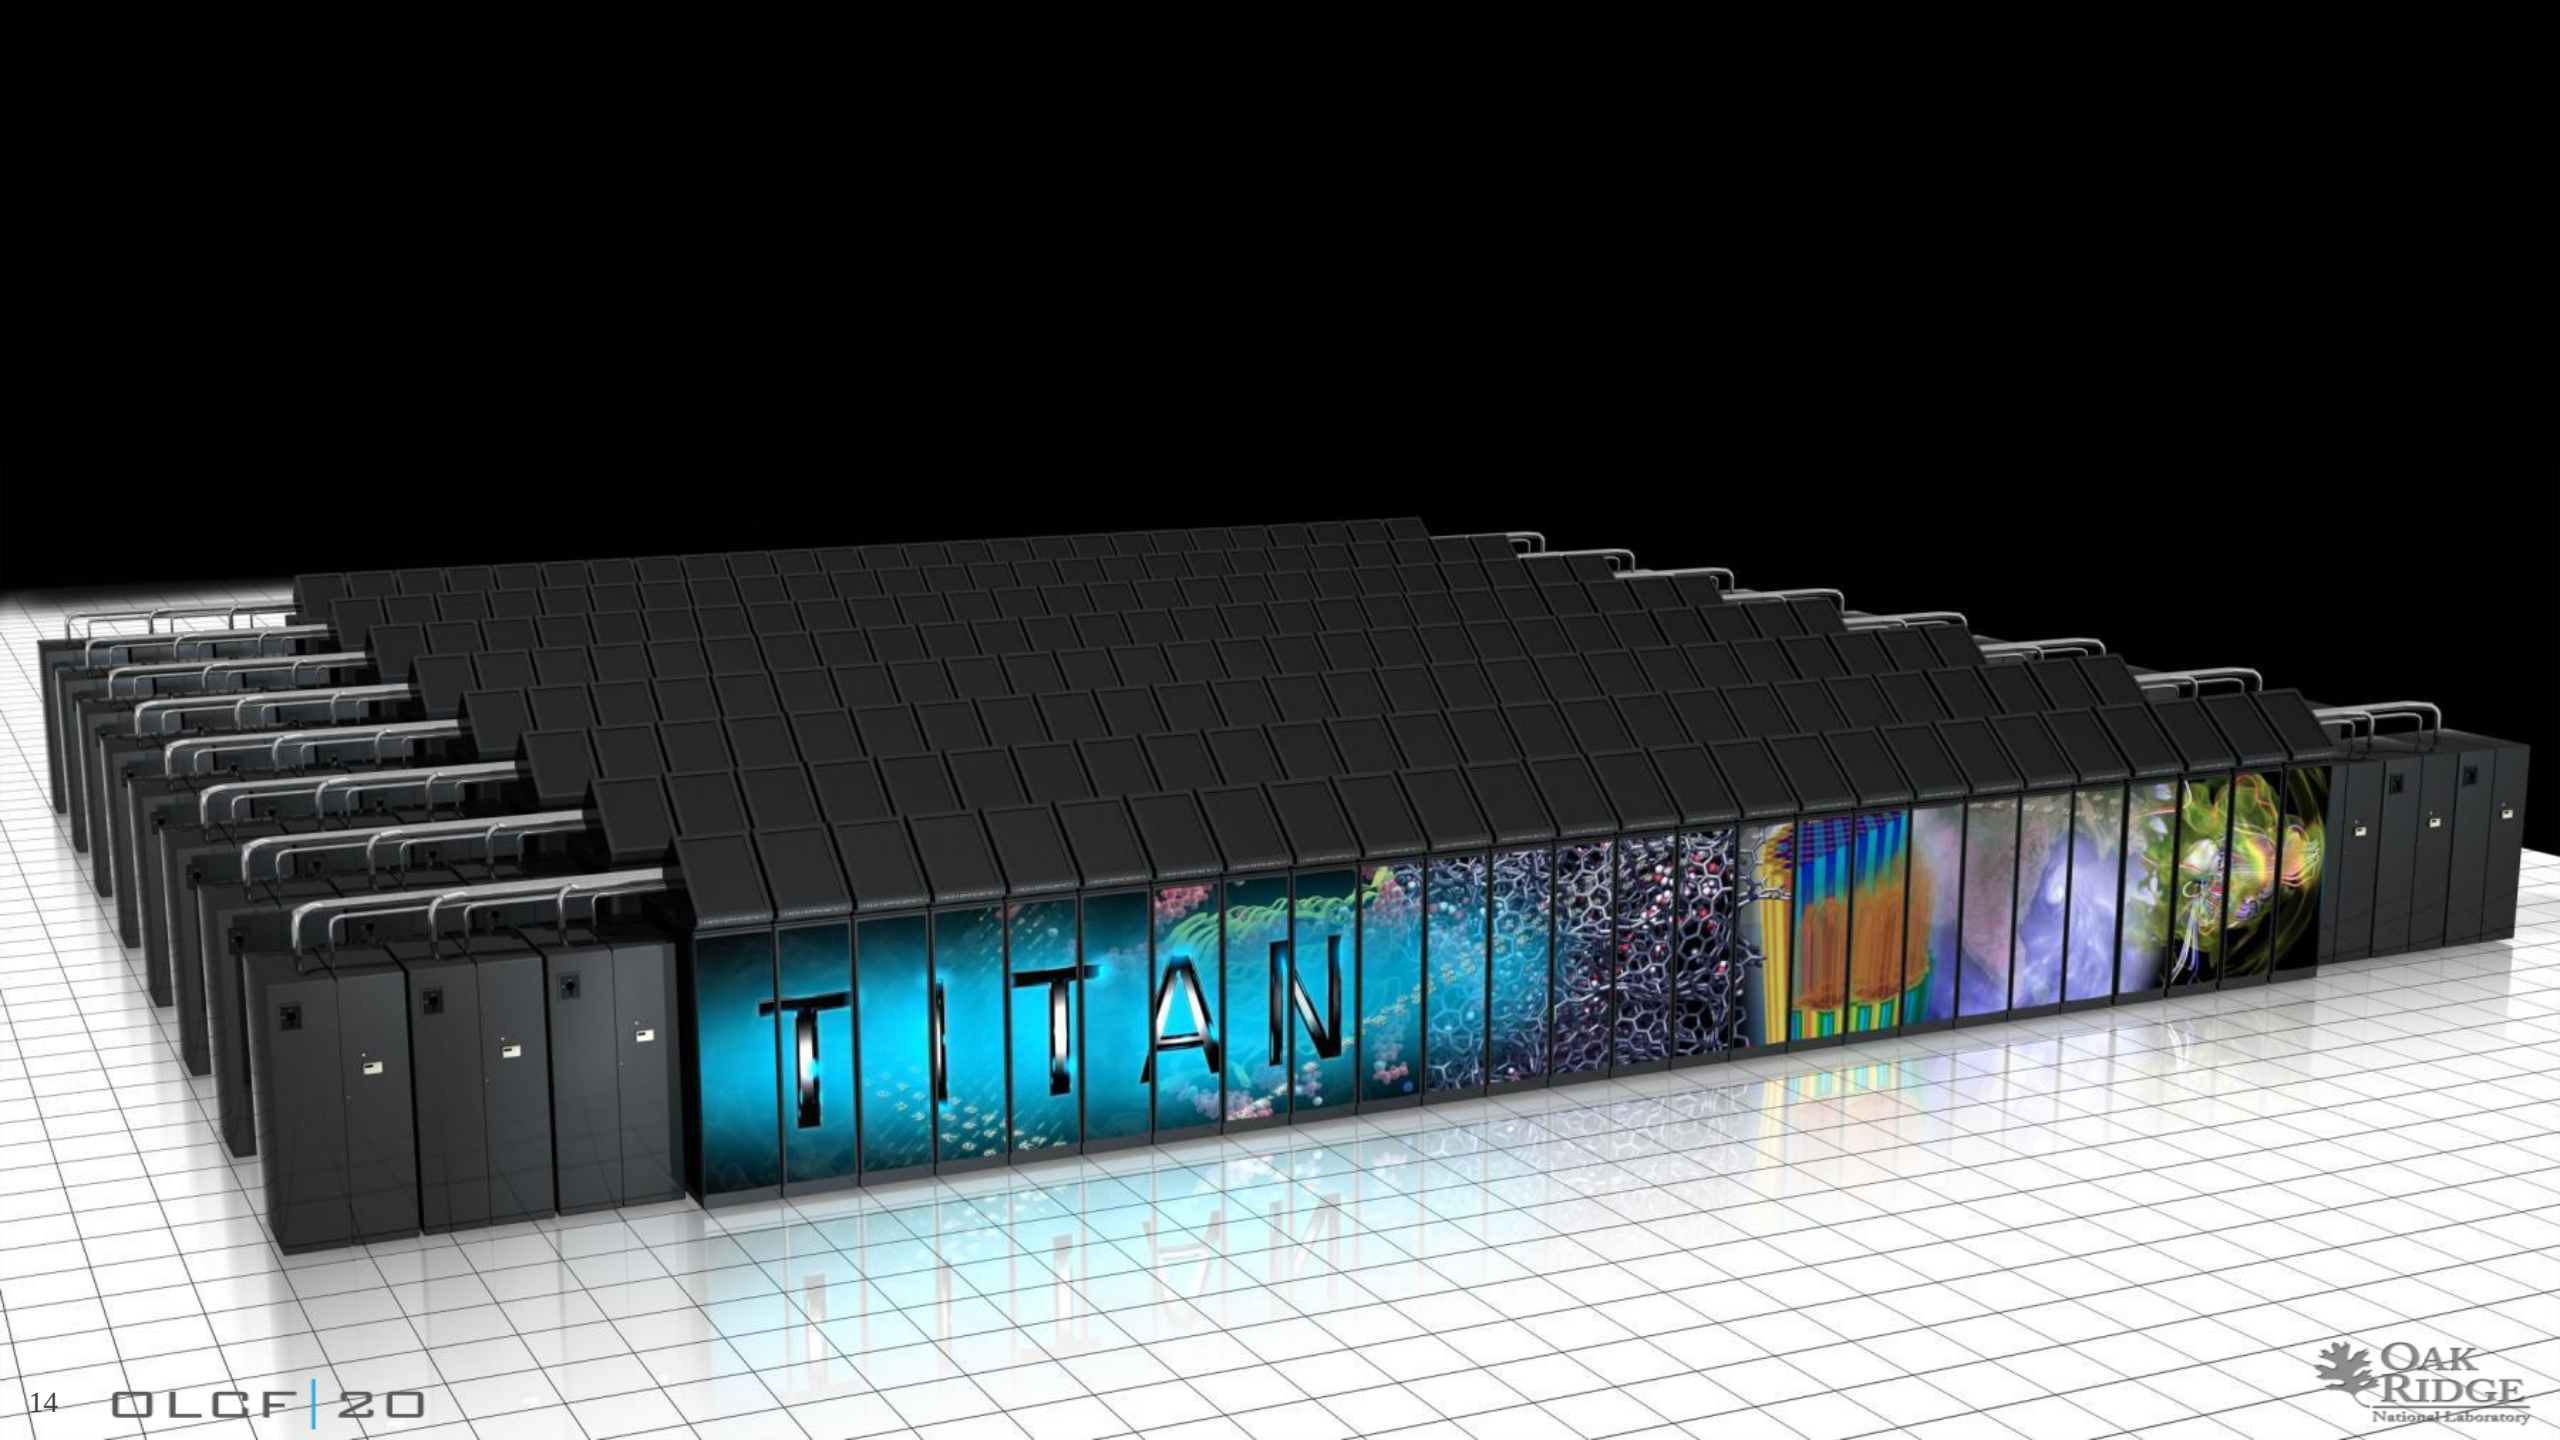
TITAN