OAK RIDGE National Laboratory

Opening Remarks

OLCF User Meeting

Georgia (Gina) Tourassi, Director

National Center for Computational Sciences

June 22, 2021

ORNL is managed by UT-Battelle LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

Virtual Meeting Etiquette

Identify Yourself

If you join as a call-in user and do not enter your name, the hosts will ask you to identify yourself.

Turn off your video

In order to preserve bandwidth, we request that non-presenters please turn off your video.

Remain on Mute

Everyone should remain on mute unless you are presenting, asking or responding to a question.

NCCS Org Chart

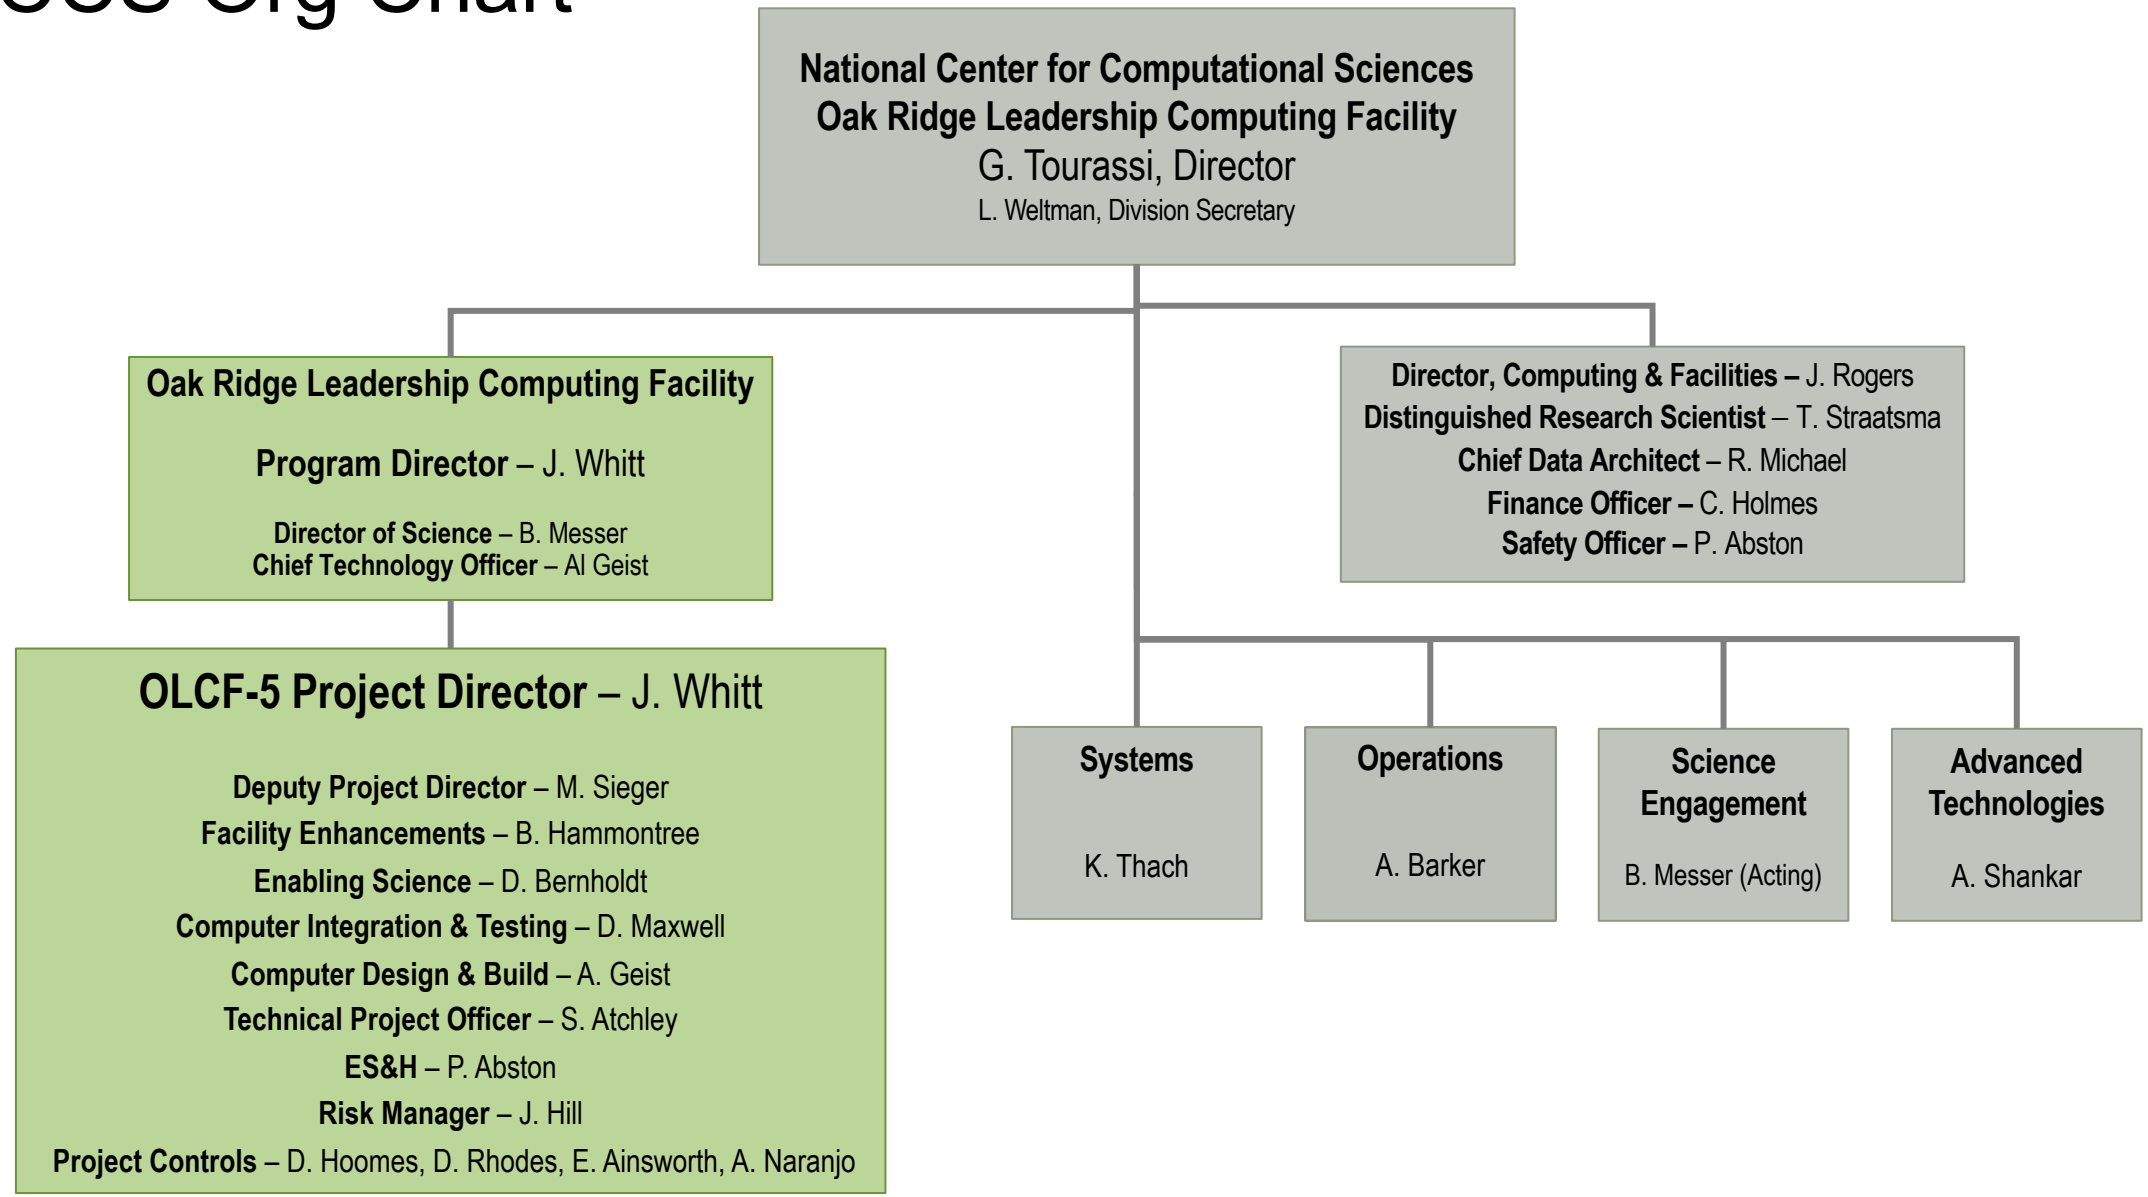

DOE Leadership Computing Facility

- **The Leadership Computing Facility** is a collaborative, multi-lab, DOE/SC initiative ranked top domestic priority; 2 centers/2 architectures to address diverse and growing computational needs of the scientific community
- **Mission:** Provide an ecosystem, including partnering opportunities, that enables unsurpassed capability computing opportunities and the associated science and engineering breakthroughs
- Administer and support two highly competitive user allocation programs (INCITE, ALCC)
  - Innovative and Novel Computational Impact on Theory and Experiment (INCITE)
  - ASCR Leadership Computing Challenge (ALCC)
  - Computational allocations typically 100 times greater than routinely available for university, laboratory, and industrial scientific and engineering environments

The OLCF has Successfully Delivered Six Systems Since 2004

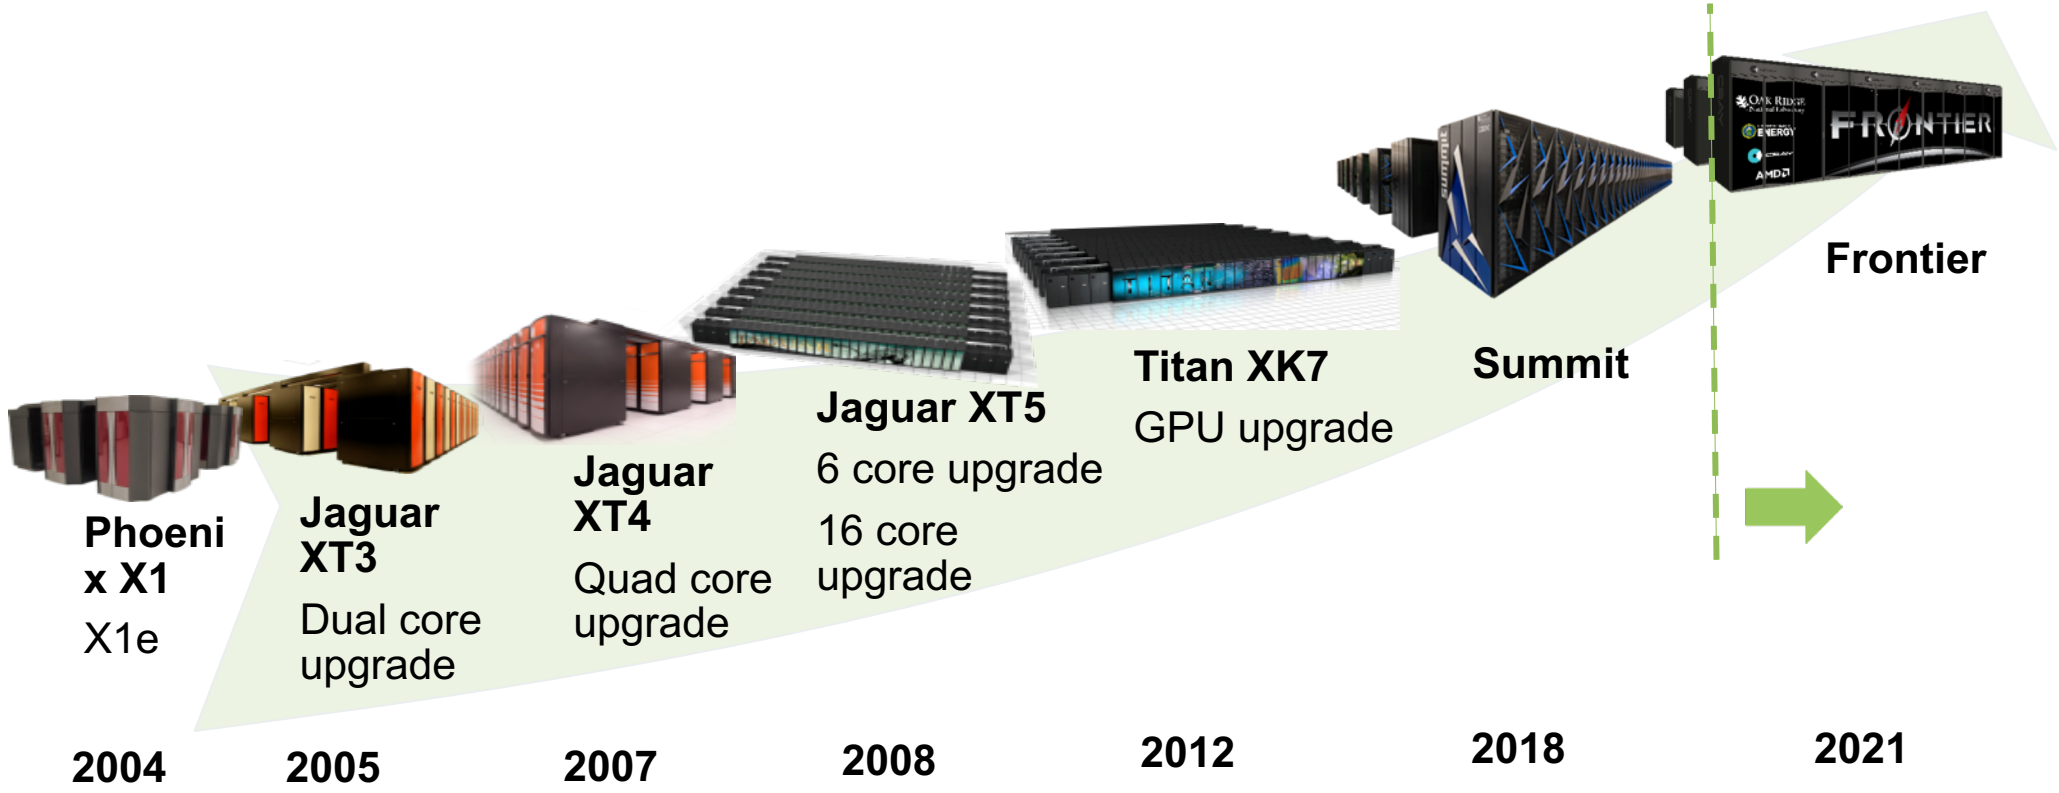

- All projects delivered on scope, schedule, and budget
- Frontier will be system number seven and will provide an increased capability of over 80,000x
- Large part of success has been strong user partnerships to scale & refactor codes/methods
- Partnering has been essential to delivering science in a rapidly changing computational environment

Summit IBM AC922

Specifications and Features

- Processor: IBM Power9™ (2/node)
- GPUs: 27,756 NVIDIA Volta V100s (6/node)
- Nodes: 4,734
- Node Performance: 42TF
- Memory/node: 512GB DDR4 + 96GB HBM2
- NV Memory/node: 1.6TB
- Total System Memory: >10PB DDR4 + HBM + Non-volatile
- Interconnect Topology: Mellanox EDR 100G InfiniBand, Non-blocking Fat Tree
- Peak Power Consumption: 13MW

Uninterrupted operations at the OLCF in 2020 were critical to the fight against COVID-19 and allowed DOE-SC research to continue through the pandemic

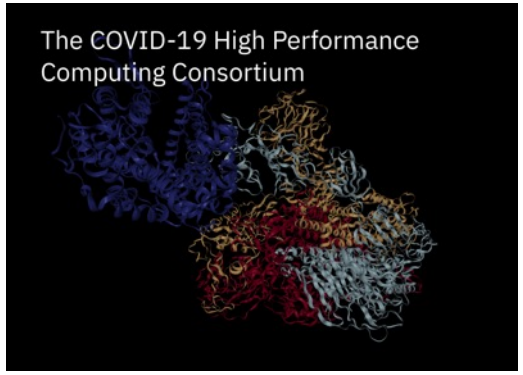

- 2M hours provided to 21 COVID-19 research projects
- Nearly 40M total hours of compute time provided
- 99% availability to ASCR user programs
- Facility continuously operated and staffed

COVID-19 Funding: Summit Node Architecture Upgrade

Adding 54 "larger memory" nodes for apps needing larger memory

- Add 54 nodes, each with double the HBM and 4x the DDR4 and NVMe
- Allows jobs that need larger on-node memory to run on up to 54 nodes
- Applications that benefit
  - Computational Chemistry
  - AI/DL for medical image analysis

Feature	Current Summit Nodes	Large Memory Nodes
Peak $FLOPS_{64}$	200 PF	203 PF
Number of Nodes	4,608	54
Node performance	43 TF	43 TF
Memory per Node	512 GB DDR4 + 96 GB HBM2	2048 GB DDR4 + 192 GB HBM2
NV memory per Node	1.6 TB	6.4 TB
Total System Memory	2.8 PB + 7.4 PB NVM	2.9 PB + 7.7 PB NVM
System Interconnect	Dual Rail EDR-IB (25 GB/s)	Dual Rail EDR-IB (25 GB/s)
Interconnect Topology	Non-blocking Fat Tree	Non-blocking Fat Tree
Bi-Section Bandwidth	115.2 TB/s	115.2 TB/s
Processors on node	2 IBM POWER9™ 6 NVIDIA Volta™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	250 PB, 2.5 TB/s, GPFS™	250 PB, 2.5 TB/s, GPFS™

OLCF Systems Enable Breakthrough Science

Gordon Bell Prize: Machine Learning for MD Simulation

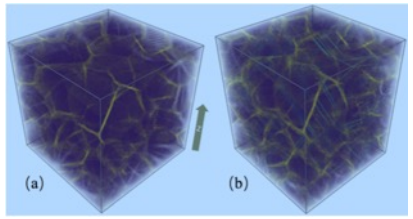

Weile Jia
UC Berkeley/LBNL

A LBNL-led team successfully tested the DeePMD-kit software package on Summit, simulating a system of 127.4 million atoms. This software package bridges classical MD and AIMD to produce complex simulations that are both large and accurate—a first. The authors contend that their protocol achieves the first efficient MD simulation of 100 million atoms with ab initio accuracy.

Jia, et al. 2020. *International Journal of High Performance Computing Applications, SC20.*

Super-conductivity

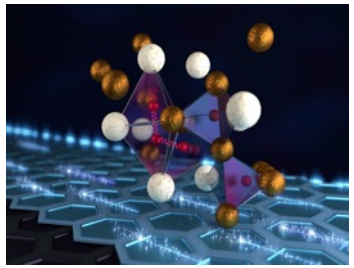

Timmy Ramirez-Cuesta
ORNL

Using neutron scattering at ORNL, researchers studied samples of zirconium vanadium hydride at atmospheric pressure and at temperatures from -450 degrees Fahrenheit (5 K) to as high as -10 degrees Fahrenheit (250 K)—much higher than the temperatures where superconductivity is expected to occur in these conditions. Their findings detail the first observations of such small hydrogen-hydrogen atomic distances in the metal hydride.

Borgschulte, et al. 2020. *Proceedings of the National Academy of Sciences.* **117**.

Turbulent Flow

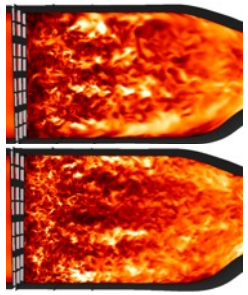

Michael Osusky
General Electric

A GE team used the Summit supercomputer at ORNL to complete first-of-their-kind 3D flow simulations that are providing breakthrough insights about these fluid flows and revealing how to better cool engine parts for more efficient and durable jet engines. The team performed the computationally intensive large eddy simulations (LES) with unprecedented speed and exceptional detail that more closely matched results of actual engine tests.

Quantum Materials

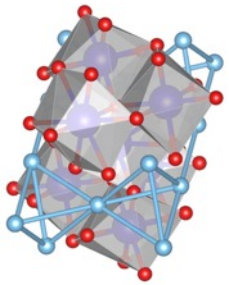

Alan Tennant
ORNL

An ORNL-led team is using artificial intelligence for the first time to find patterns in neutron scattering data that can lead to an understanding of the physics inside quantum materials or complex magnetic materials. The team trained an artificial neural network to successfully interpret data from a neutron scattering experiment performed at ORNL's Spallation Neutron Source.

Samarakoon et al., 2020. *Nature Communications.* **11.**

Earth Systems Modeling

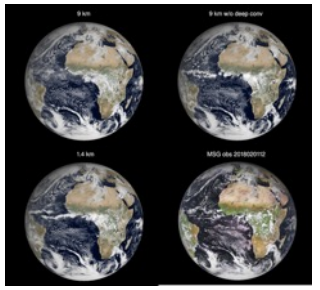

Nils Wedi
ECMWF

Using Summit, a team of researchers from ECMWF and ORNL achieved a computational first: a global simulation of the Earth's atmosphere at a 1-square-kilometer average grid-spacing for a full 4-month season. The milestone marks a big improvement in resolution for the "European Model," which currently operates at 9-kilometer grid-spacing for routine weather forecast operations.

Wedi, et al. 2019. *Journal of Advances in Modeling Earth Systems*. **12**.

All from 2020 work conducted on Titan and Summit.

OLCF Systems Enable Breakthrough Science: COVID-19

Gordon Bell Special Prize: AI Workflows

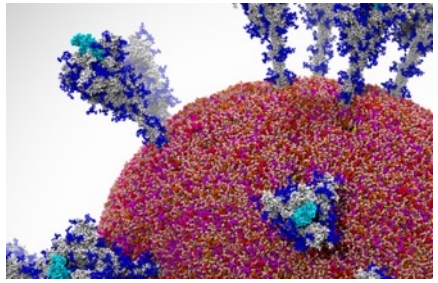

Rommie Amaro
UC San Diego

A UC San Diego and Argonne-led team built a first-of-its-kind workflow based on AI and ran it on the Summit to simulate the SARS-CoV-2 virus's spike protein in numerous environments, including within the virus's viral envelope comprising 305 million atoms—the most comprehensive simulation of the virus performed to date.

Casalino, et al. 2020. *International Journal of High Performance Computing Applications, SC20.*

Treatment

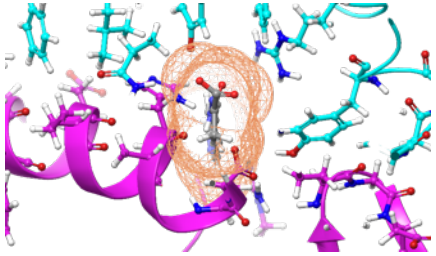

Jeremy Smith
ORNL

An ORNL and UTK-led team used Summit to identify small-molecule drug compounds that might warrant further study in the fight against the SARS-CoV-2 coronavirus, which is responsible for the COVID-19 disease outbreak. They performed simulations of more than 8,000 compounds to screen for those that are most likely to bind to the main "spike" protein of the coronavirus, rendering it unable to infect host cells, and found 77 compounds of interest that might have value further studies of the virus.

Smith, et al. 2020. *ChemRxiv*.

Response Pathways

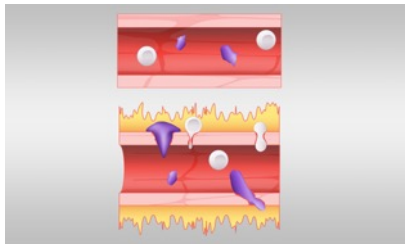

Dan Jacobson
ORNL

An ORNL-led team used Summit to analyze genes from cells in the lung fluid of COVID-19 patients compared with control patients and found that genes related to one of the systems responsible for lowering blood pressure—the bradykinin system—appear to be excessively "turned on" in the lung fluid cells of those with the virus, information that could inform new drug targets. They needed the power of Summit to run 2.5 billion correlation calculations to understand the normal regulatory circuits and relationships for the genes of interest.

Garvin, et al. 2020. *eLife. 9.*

All from 2020 work conducted on Summit.

Frontier System Overview

System Performance

- Peak performance of >1.5 exaflops
- Anticipated to debut in 2021 and available to users in 2022.

The system includes

- 1 HPC and AI Optimized AMD EPYC CPU and 4 Purpose Built AMD Instinct™ GPUs per node
- Approximately 10 PB of combined high bandwidth and DDR memory
- Cray four-port Slingshot network 100 GB/s

On-Node interconnect

- AMD Infinity Fabric
- Coherent memory across the node

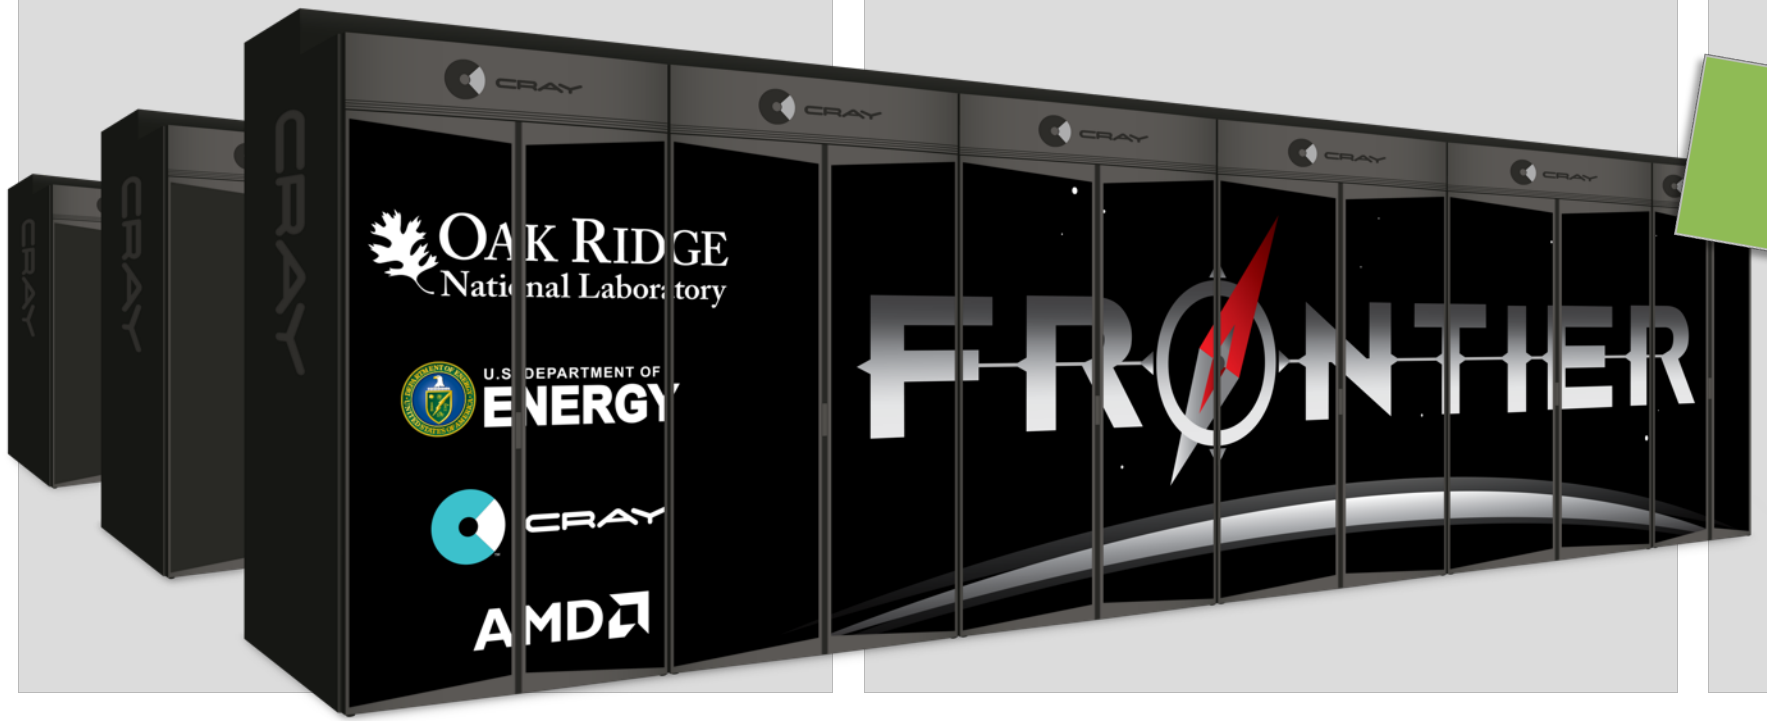

COMING IN 2021!

Frontier: Data Center Prep Progress to Date

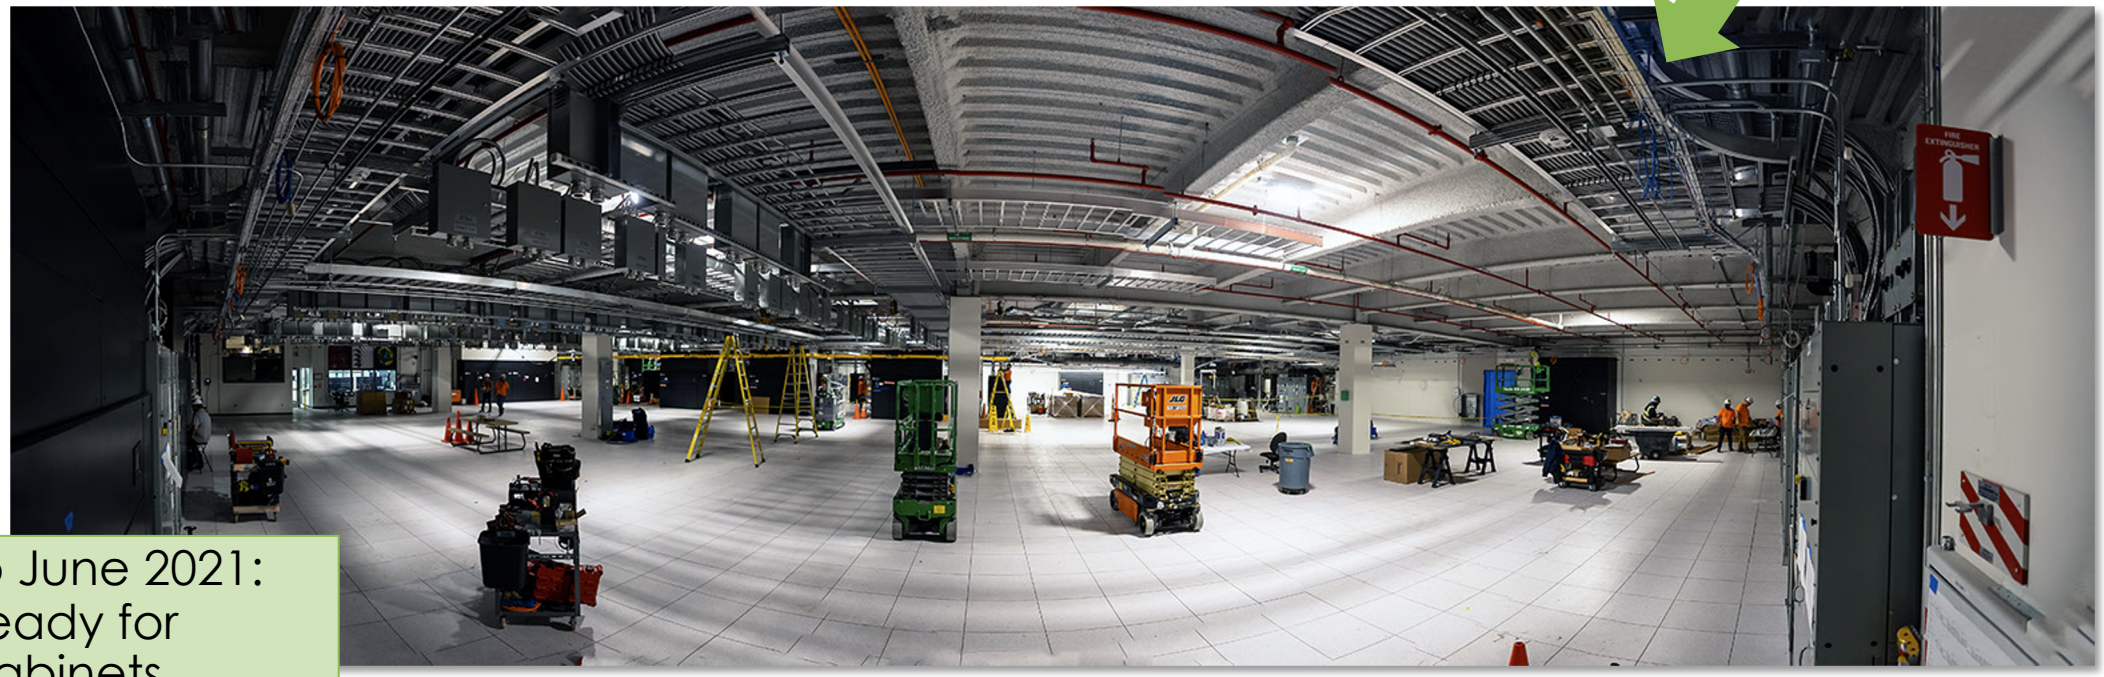

To June 2021:
Ready for
Cabinets

Power and cooling are in place for the Frontier computer

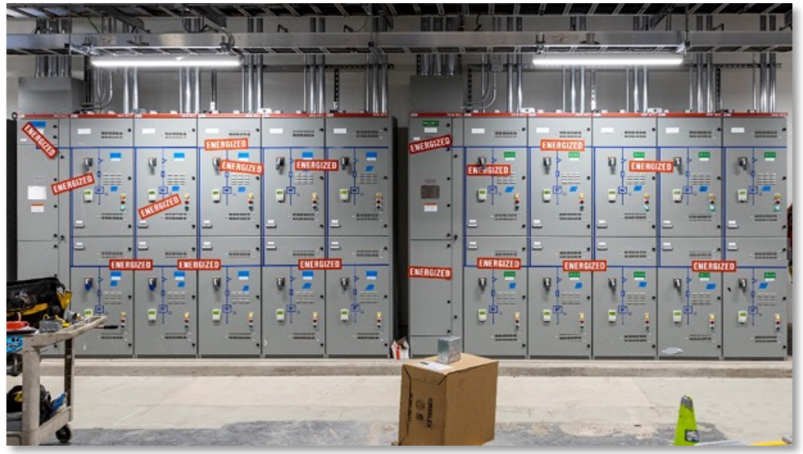

Agenda

Time	Title	Speaker
11:00 AM – 11:15 AM	Welcome & Introduction	Gina Tourassi Director, NCCS
11:15 AM – 12:00 PM	Summit Updates: • System Updates • Jupyter Notebooks on Summit • Training • Virtual Facility Tour	Session Leader: Thomas Papatheodore, OLCF System Acceptance and User Environment
12:00 PM – 1:00 PM	Table-Top Breakout Discussions Day #1	Session Leader: Chris Fuson, OLCF User Assistance
1:30 PM – 3:00 PM	User-Led Talks: • Roberto Car • Amanda Randles • Jens Glaser	Session Leader: Jens Glaser, OLCF Advanced Computing for Chemistry and Materials

+ Full agenda here: https://www.olcf.ornl.gov/calendar/2021-olcf-user-meeting/

Questions?