Update on Frontier

June 2022 OLCF User Conference Call

Bronson Messer
Director of Science
Leadership Computing Facility
Oak Ridge National Laboratory

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

Oak Ridge National Laboratory's Journey from Petascale to Exascale

Mission: Providing world-class computational resources and specialized services for the most computationally intensive global challenges	Vision: Deliver transforming discoveries in energy technologies, materials, biology, environment, health, etc.

Steady progress of 10x per generation

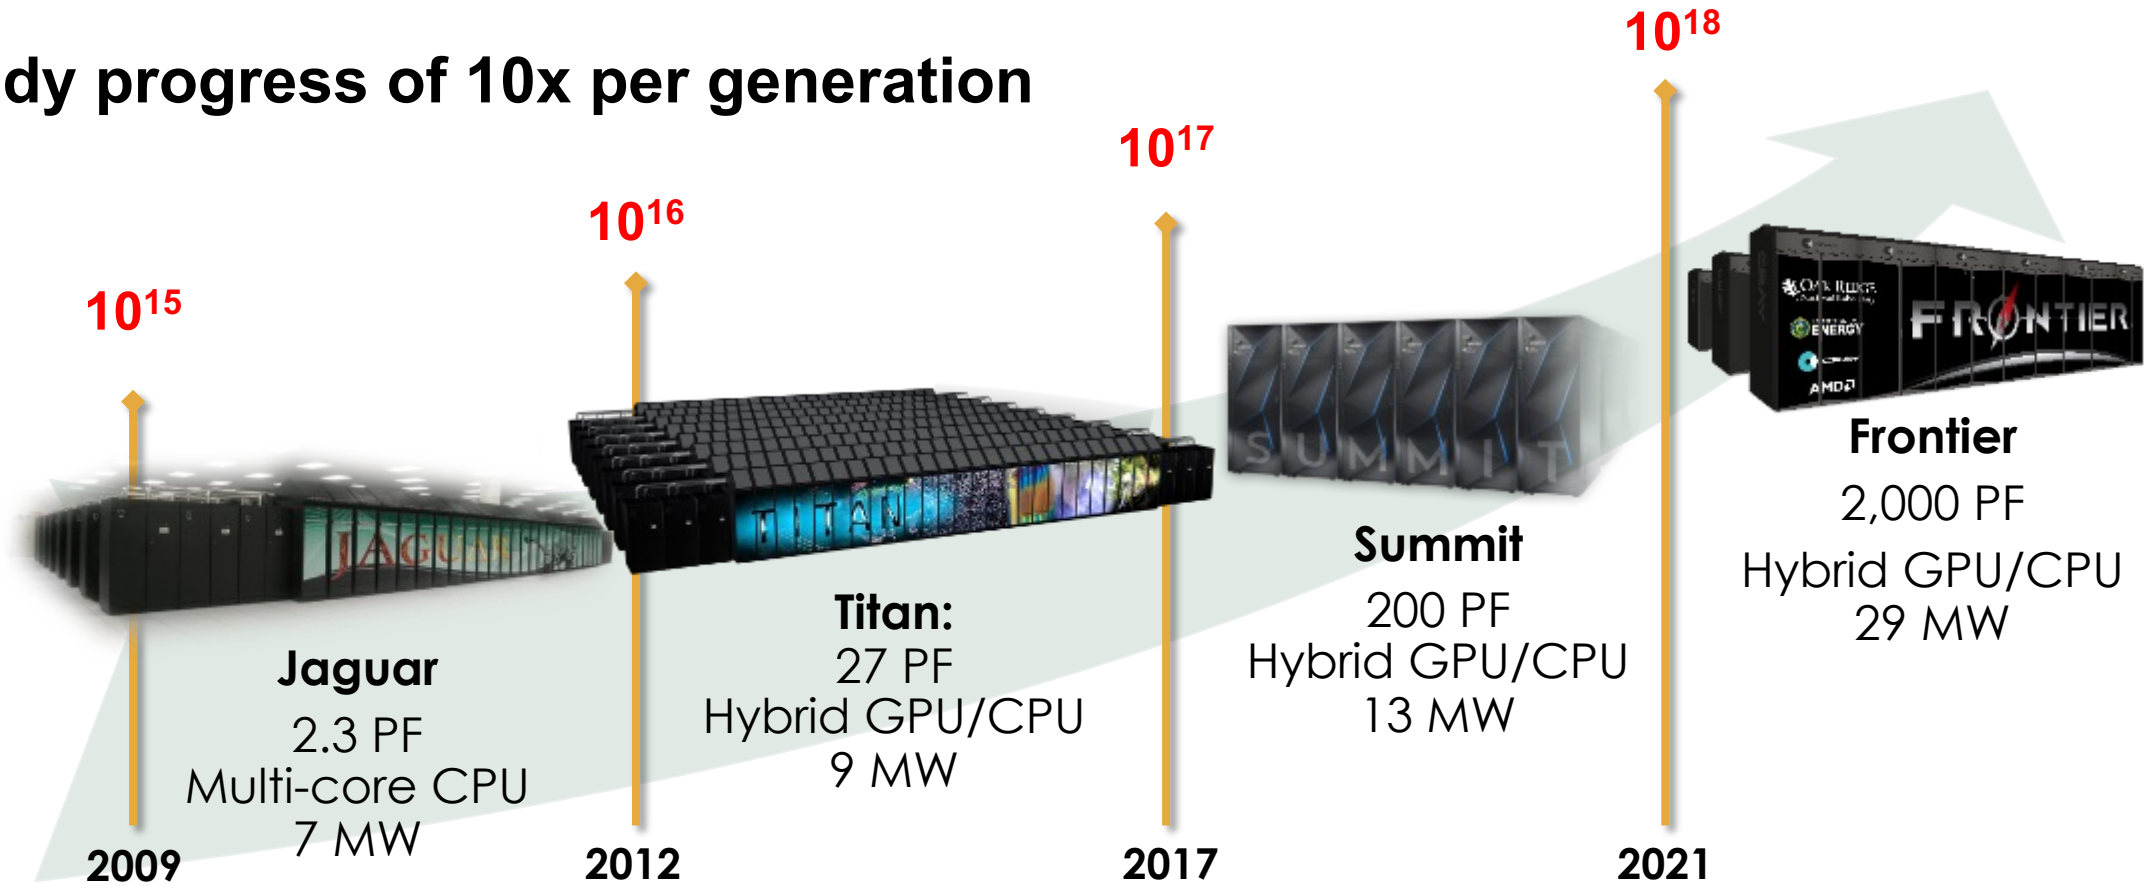

Frontier Overview

Extraordinary Engineering

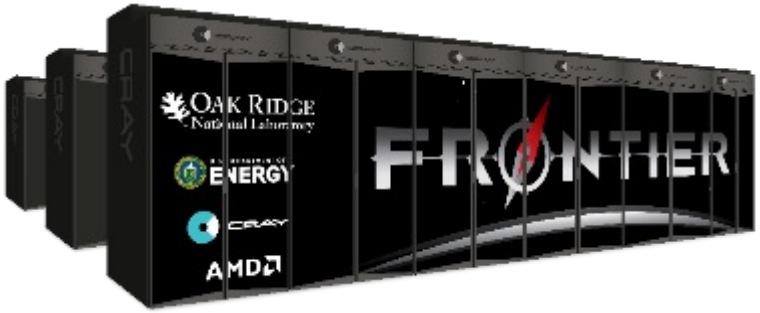

System

- 2 EF Peak DP FLOPS
- 74 compute racks
- 29 MW Power Consumption
- 9,408 nodes
- 9.2 PB memory
  (4.6 PB HBM, 4.6 PB DDR4)
- Cray Slingshot network with dragonfly topology
- 37 PB Node Local Storage
- 716 PB Center-wide storage
- 4000 ft^2 foot print

Built by HPE

Olympus rack

- 128 AMD nodes
- 8,000 lbs
- Supports 400 KW

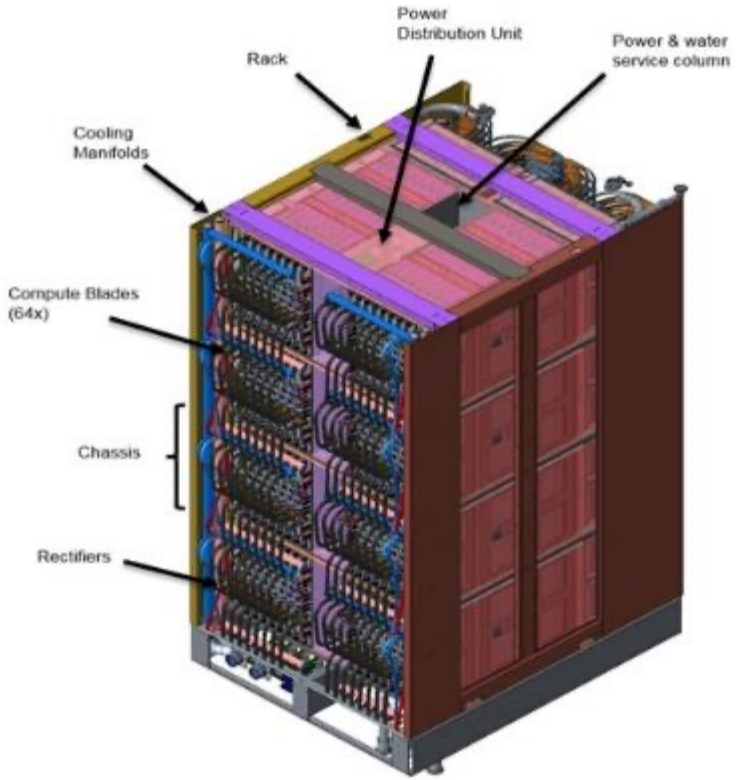

Powered by AMD

AMD node

- 1 AMD “Trento” CPU
- 4 AMD MI250X GPUs
- 512 GiB DDR4 memory on CPU
- 512 GiB HBM2e total per node
  (128 GiB HBM per GPU)
- Coherent memory across the node
- 4 TB NVM
- GPUs & CPU fully connected with AMD Infinity Fabric
- 4 Cassini NICs, 100 GB/s network BW

Compute blade

- 2 AMD nodes

All water cooled, even DIMMS and NICs

Energy-efficient computing – Frontier achieves 14.5 MW per EF

Since 2009 the biggest concern with reaching Exascale has been energy consumption

- **ORNL pioneered GPU use in supercomputing** beginning in 2012 with Titan thru today with Frontier. Significant part of energy efficiency improvements.
- **ASCR [Fast, Design, Path] Forward vendor investments** in energy efficiency (2012-2020) further reduced the power consumption of computing chips (CPUs and GPUs)..
- **200x reduction in energy per FLOPS** from Jaguar to Frontier at ORNL
- ORNL achieves additional energy savings from using warm water cooling in Frontier (32 C).
  ORNL Data Center PUE= 1.03

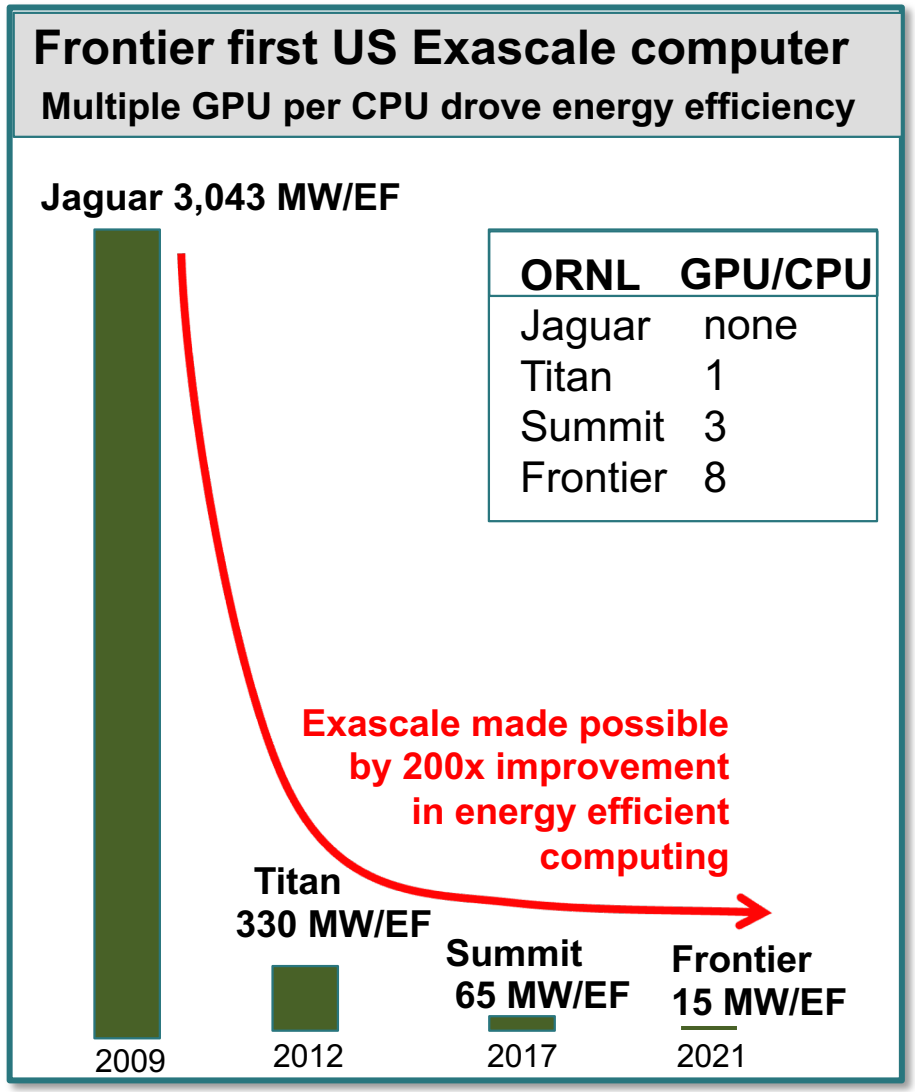

Frontier multi-tier storage system

	Capacity	Performance	
Multi-tier I/O Subsystem		**Read**	**Write**
	37 PB Node Local Storage	65.9 TB/s	62.1 TB/s
		11 Billion IOPS	
	11 PB Performance tier	9.4 TB/s	9.4 TB/s
	695 PB Capacity tier	5.2 TB/s	4.4 TB/s
	10 PB Metadata	2M Transactions per sec	

Two 2TB SSD NVM per node Local Storage (Flash)

Gazelle SSD Storage board (Performance Tier and Metadata)

Moose HDD Storage board (Capacity Tier)

During Frontier build -- the chip shortage hit in earnest!

When HPE began ordering parts, suppliers said the lead time on orders was increasing an additional 6-12 months.

60 Million parts needed for Frontier

685 Different part numbers used in Frontier

167 Frontier part numbers affected by the chip shortage
(more than 2 million parts from dozens of suppliers worldwide)

12 Part numbers blocked building the first compute cabinet

15 Part numbers shortage for AMD building all the MI200 cards for Frontier

It wasn't exotic parts like CPUs or GPUs, rather parts needed by everyone – in cars, TVs, electronics, such as voltage regulators, oscillators, power modules, etc.

Last Cabinet of Frontier Delivered to ORNL October 18th

Thanks to Heroic Efforts of the HPE and AMD teams

Last cabinet being rolled into place.
(Each cabinet weighs 8,000 lbs.)

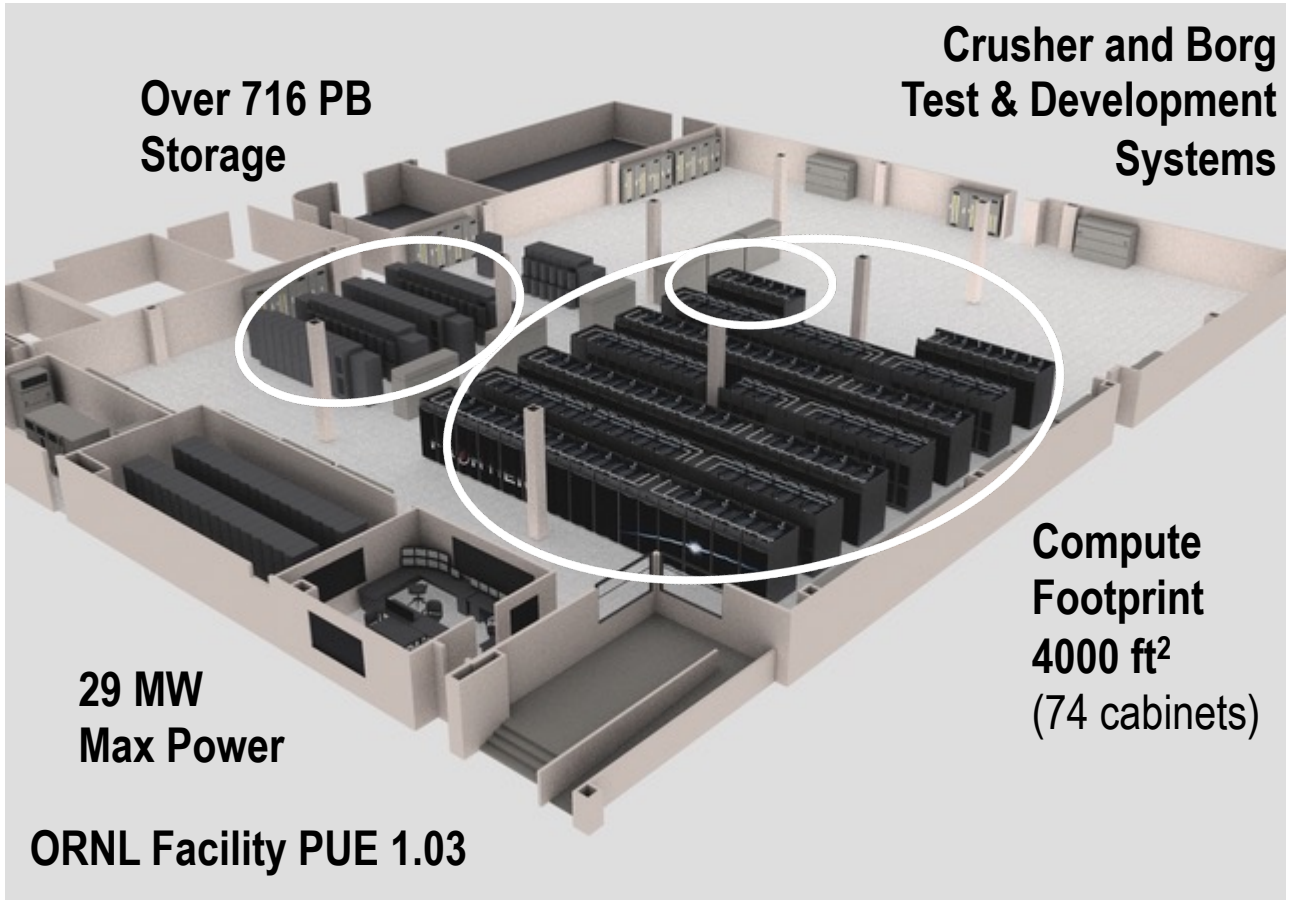

After the cabinets arrived they had to be connected. There are 81,000 cables between all the Frontier nodes

CAAR

The **Center for Accelerated Application Readiness (CAAR)** is the primary OLCF program to achieve and demonstrate application readiness

- ***Build on the experience from the successful CAAR programs for OLCF-3 (Titan) and OLCF-4 (Summit)***
- ***CAAR project resources***
  - Dedicated collaboration with OLCF staff
  - Support and consultation from other project personnel, particularly from the Programming Environment and Tools area, and the vendor Center of Excellence
  - OLCF Postdoctoral fellows (both during application readiness and early science)
  - Allocations to available compute resources (Summit, early access systems)

Current Access Plan

- Frontier will be in production starting January, 2023 (INCITE)
- ECP, CAAR, and a handful of new ALCC awardees have access to the Frontier test & development system (Crusher)

- Director's Discretionary allocations will be available on Frontier after the machine is formally accepted
- Current target for full acceptance is September, 2022

Questions?