A Quick Update on Exascale Systems

November 2021

www.HyperionResearch.com
www.hpcuserforum.com

Bob Sorensen

Near-Term US Exascale Plans

Three systems over two years with budget of ~ $1.8 billion

- **Frontier: DOE Office of Science: Oak Ridge National Laboratory**
  - Currently being installed: >9,000 nodes, ~29MW, >1.5EF
  - Cray Slingshot with AMD EPYC CPU and AMD Radeon Instinct GPUs
  - Delivery in late 2021, full access in 3Q 2022
- **Aurora: DOE Office of Science, Argonne National Laboratory**
  - >9,000 nodes, **60MW**, > 1EF DP sustained
  - Delivery in late 2022, acceptance in 2023 (delayed at least 12 months)
  - Cray Shasta architecture with Intel Xeons and Intel Xe GPU
  - 08/21: Polaris testbed system (44PF DP and 1.5EF AI)
    - Stand up in 2Q2022
- **El Capitan: DOE NNSA's LLNL**
  - Early access RZNevada (100 of TFs) delivered in 02/2021
  - Full system delivery in late 2022, with full production, late 2023
  - Cray Shasta architecture AMD EPYC processors, next generation Radeon Instinct GPUs

Mid-Size US Exascale Plan

Crossroads on the horizon: LANL

- **Procurements by the Alliance for Computing at Extreme Scale (ACES) partnership between Los Alamos National Laboratory and Sandia National Laboratories**
- **$105 million contract awarded to Hewlett Packard Enterprise (HPE) to deliver Crossroads, a next-generation HPE Cray EX supercomputer to be sited at Los Alamos for 2022 operation**
- **Crossroads will be among the first to receive a supercomputer equipped with NVIDIA's Grace CPUs**

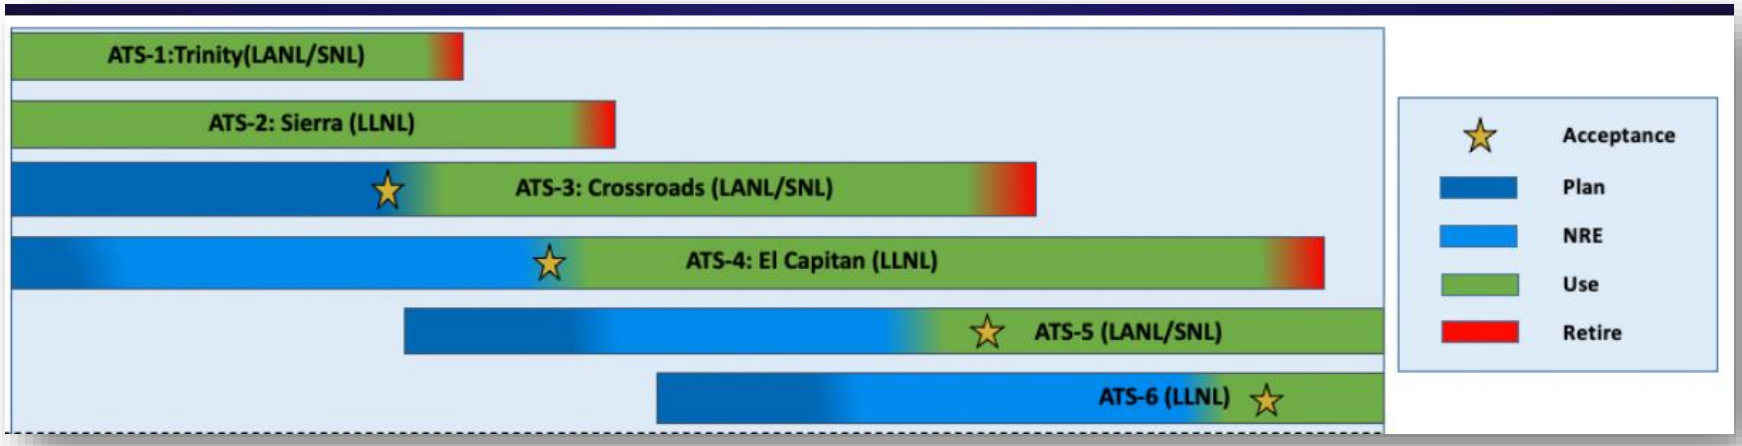

Source: Gary Grider LANL, DoE, LA-UR-21-22315

The Cycle Continues

Looking ahead started already

Post-Exascale Computing for the National Nuclear Security Administration

As requested in section 3172 of the FY 2021 National Defense Authorization Act, an ad hoc committee of the National Academies of Sciences, Engineering, and Medicine will conduct a consensus study "reviewing the future of computing beyond exascale computing to meet national security needs at the National Nuclear Security Administration." (Exascale refers to a computer that performs near or above 10^18 floating point operations per second.)

The study will review:
(1) NNSA's computing needs over the next 20 years that exascale computing will not support;
(2) Future computing technologies for meeting those needs including quantum computing and other novel hardware, computer architecture, and software;
(3) The likely trajectory of promising hardware and software technologies and obstacles to their development and their deployment by NNSA; and
(4) The ability of the U.S. industrial base, including personnel and microelectronics capabilities, to meet NNSA's needs.

Dates		Amounts	
Date Signed *(mm/dd/yyyy)* :	05/28/2021	Action Obligation:	$1,077,278.00
Period of Performance Start Date *(mm/dd/yyyy)* :	06/01/2021	Base And Exercised Options Value:	$1,077,278.00
Completion Date *(mm/dd/yyyy)* :	05/31/2023	Base and All Options Value (Total Contract Value):	$1,077,278.00
Est. Ultimate Completion Date *(mm/dd/yyyy)* :	05/31/2023	Fee Paid for Use of IDV:	$0.00
Solicitation Date *(mm/dd/yyyy)* :	04/12/2021		

Source: National Academies, 2021

Japan's Exascale System

Riken's Fugaku #1: June 2020 - June 2021 Top 500 list

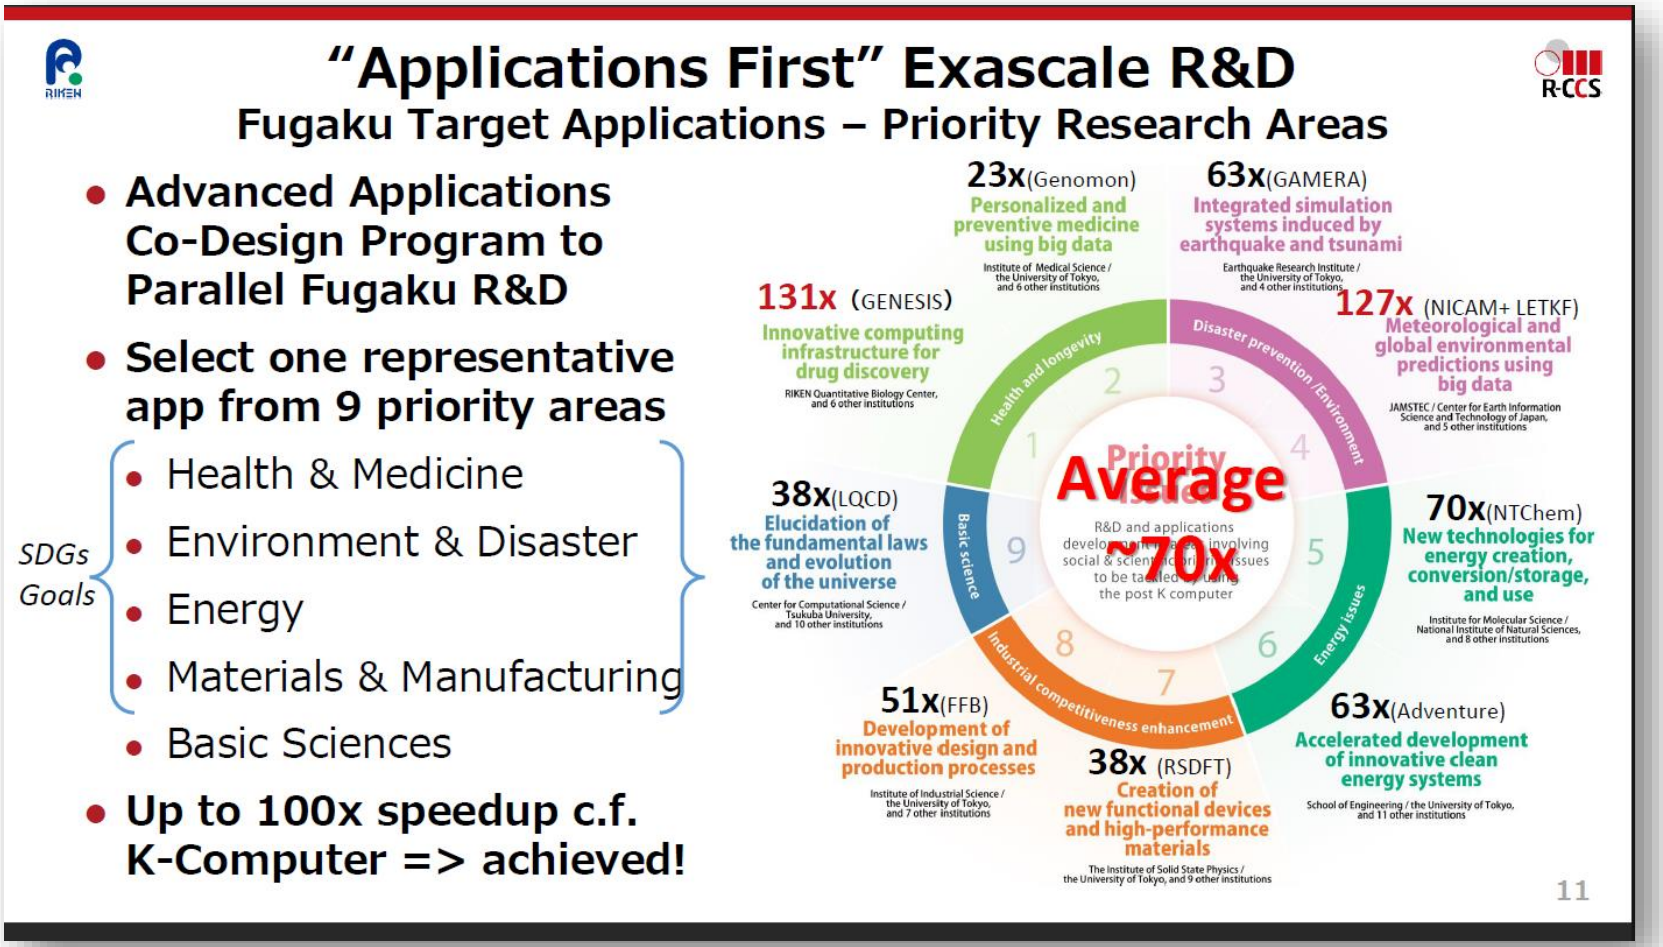

Source: Riken, 2021

China Exascale Status

Two (and first) exascale systems are already operating, the 3rd is delayed

- **Sunway Pro OceanLight**
  - Completed March 2021, ~1.3 EFlops Rpeak, ~1.05 EFlops Rmax
  - Full Linpack run
  - 35 MW
  - ShenWei post-Alpha CPU
  - Est. 96 cabinets X 1024 SW39010 (390 core) cpus -> 38 million cores
  - National Supercomputing Center in Wuxi
- **NSCC-Tianjin Tianhe-3**
  - Dual-chip FeiTeng ARM and matrix accelerator node
  - Full completion this month
  - Est ~ 1.7 EFlops Rpeak, 1.3 EFlops Rmax
  - No Linpack results currently
- **NSCC-Shenzhen Phase 2 - Sugon 2-Flops at Guangdong Province**
  - Scheduled for 2022—delayed with possible platform change
  - Hygon processors (low confidence), may go AMD Zen4
- **Skipped June Top 500 listing? Maybe again?**

EU HPC Plans

EuroHPC program: confusion in the making

- **Chartered to develop EU-wide HPC development program**
  - 33 participating States + EU
  - Operational duration: November 2018-2026
- **Three sites recently selected for 150-200 Pflops systems**
  - Kajaani Finland, Barcelona Spain, and Bologna Italy
  - Total investment: 650 million Euros
    - 50% EU
    - 50% Consortium
- **Systems are owned by EuroHPC Joint Undertaking**
- **Trouble afoot :**
  - July 2021: EuroHPC JU canceled the public procurement process for the MareNostrum 5 supercomputer
    - The voting result did not achieve the needed majority to reach an agreement to adopt the selected tender
    - Officials expect that it will be reopened in autumn, leading to a selection early in 2022

EU HPC Plans (continued)

Exascale plans going forward

- **EU plan calls for acquisition of two exascale systems in the 2021-2024 timeframe**
  - At least one to use European technology: specifically using an EPI-developed processor
  - Additional procurements in Germany in 2024, 2025
  - EU may include 2 additional ES systems in 2023-2026
- **Post Exascale System around 2027**
  - Plans call for integration and deployment of the first hybrid HPC/quantum infrastructure in Europe

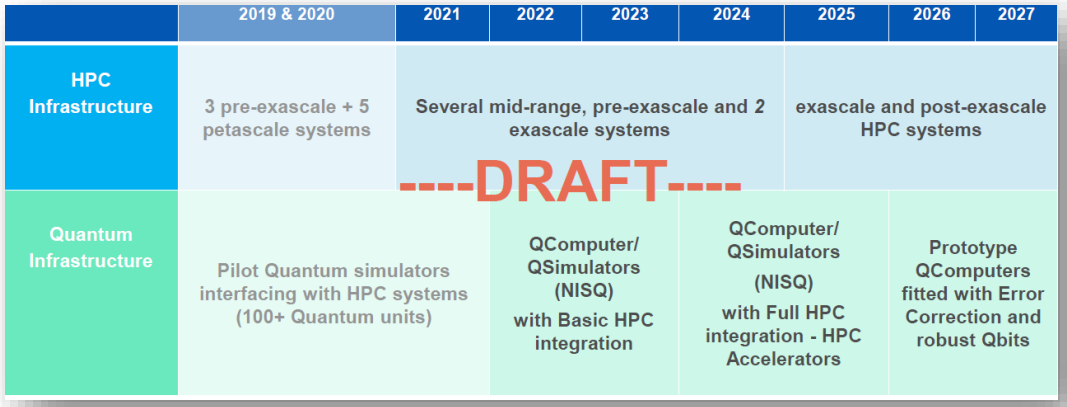

Source: Leonardo Flores Añover, Senior Expert

DG CNECT, HPC & Quantum Technology Unit - European Commission 2021

UK EU Plans

UK looking at 2024

- **The UK, which will not likely be eligible to fully take part in EuroHPC projects or access calls when Horizon 2020, has plans for a domestic exascale system**
  - Exascale project requirements include support for both traditional modeling and simulation as well as AI/Deep Learning
  - System targeted for both scientific community and industrial users
- **Exascale rollout schedule**
  - Procurement during 2022
  - Assembly and installation 2023
  - Final changes to hosting environment 2023
  - Planned service opening April 2024
- **System will be hosted at Advanced Computing Facility of EPCC, formerly the Edinburgh Parallel Computing Centre, a supercomputing centre based at the University of Edinburgh**

HYPERION RESEARCH

QUESTIONS?

Questions or comments are welcome.

Please contact us at:

info@hyperionres.com