Special Analysis

A Broad International Roll Out of Exascale Systems on the Horizon

Bob Sorensen, Steve Conway, and Earl Joseph
September 2020

HYPERION RESEARCH OPINION

Hyperion Research projects that between 2021 and 2026, worldwide installations of leading-edge exascale and near-exascale systems will total 28-38 new systems, worth an estimated $10-$15 billion, with China, the EU, and the United States each fielding 7 -10 systems in the six year interval. Being an exascale trailblazer can be expensive, with systems on the near horizon costing upwards of $500 million or more each, and only the most ambitious governments can commit to participating in these early rounds of exascale progress. However, as HPC price/performance moderates, a long-verified phenomenon in the HPC space, additional nations will also be looking to provide their domestic R&D base with access to exascale-class HPCs either developed domestically or procured from foreign sources.

CURRENT SITUATION

Hyperion Research projects that between 2021 and 2026, worldwide installations of leading-edge exascale and near-exascale systems will total 28-38 new systems, worth an estimated $10-$15 billion. As seen in Table 1, the new systems will be distributed across the major HPC user nations, with China, the EU, and the United States each fielding 7 -10 systems over the next six years. Japan will likely field around four new systems, while tier-two HPC nations, including India, Singapore and South Korea, could add to the list with 3-4 exascale-class system installations among them, albeit in the 2024 and later time period. Hyperion Research does not foresee see any significant negative impacts to these forecasts as a result of the economic downturn due to the covid-19 pandemic.

TABLE 1

Exascale and Near-Exascale Systems (2021 to 2026)

Year Accepted	China	EU & UK	Japan	US	Other Countries*	Total Systems	Total Value
2021	1 or 2 near-exascale systems ~$400M each	1 pre-exascale system ~$150M	1 exascale system Fugaku ~$1B	--	--	3-4	$1.6B-$2.0B
2022	1 or 2 exascale systems ~$350M - $400M each	2 pre-exascale systems ~$150M each	1 near-exascale	2 exascale systems	--	6-7	$2.0B-$2.4B

TABLE 1

Exascale and Near-Exascale Systems (2021 to 2026)

Year Accepted	China	EU & UK	Japan	US	Other Countries*	Total Systems	Total Value
			system ~$200M	~$550M each			
2023	1 or 2 exascale system ~$350M - $400M each	1 or 2 exascale systems ~$375M	1 near-exascale system ~$150M	1 exascale system ~600M	--	4-6	$1.5B-$2.3B
2024	1 exascale system ~$350M - $400M each	1 or 2 exascale systems ~$375M	?	1 or 2 exascale systems ~$500M each	1 exascale system ~$250M	4-6	$1.5B-$2.4B
2025	2 exascale systems ~$350M - $400M each	1 exascale systems ~$375M	1 exascale system ~$150M	1 or 2 exascale systems ~$500M each	1 exascale system ~$200M	6-7	$1.9B-$2.5B
2026	1 or 2 exascale systems ~$350M - $400M each	1 or 2 exascale systems ~$375M	?	2 exascale systems ~$500M each	1 or 2 exascale systems ~$200M	5-8	$2.0B-$3.0B
Total	**7-11**	**7-10**	**4+**	**7-9**	**3-4**	**28-38**	**$10B-$15B**

* Includes Australia, Canada, India, Israel, Russia, Saudi Arabia, Singapore, and South Korea

Source: Hyperion Research, September 2020

National Exascale Plans

The demand for exascale-class systems is being addressed by a number of government-funded exascale development programs including those currently underway in the United States, China, Japan, the EU, and the UK, each with its own goals, budgets, and vision for realizing exascale-class performance.

For the US, the three earliest exascale systems will be installed in three US Department of Energy laboratories: the Aurora System at Argonne National Laboratory, the Frontier System at Oak Ridge National Laboratory, and the El Capitan at the Lawrence Livermore National Laboratory. All three machines, each with a price tag on the order of $500 million, will be built on the HPE/Cray Shasta architecture but with a mix of CPU-GPU configurations. Recent information indicates, however, that

shipment of the Aurora system, targeted to be the first US exascale machine, will likely be delayed due to technical difficulties by Intel in delivering their GPU on schedule.

For China, different exascale systems are currently under development by three different Chinese HPC developers based on technology selected for prototype development in 2017. When announced, the plan was for one or more of the prototypes to be selected for full-up implementation as an exascale system. Hyperion Research expects that at least one of each type will be selected for full exascale operation.

Each Chinese developer went with a different architectural emphasis, and many expect one of these three Chinese HPCs will be announced as the world's first exascale-class system:

- The National University of Defense Technology (NUDT) is working on the Tianhe-3, the latest implementation of its world-class Tianhe HPC line. The Tianhe-3 is based on an NUDT-designed CPU, possibly an Arm-based Phytium Xiaomi, coupled in a one-to-one ratio with the Matrix-300 DSP, an NUDT 96-core accelerator with a peak performance of nearly 10 teraflops. Each blade of the Tianhe-3 will likely be equipped with eight DSPs paired with eight CPUs, providing 96 teraflops per blade with 128 blades per cabinet.
- The National Research Center of Parallel Computer (NRCPC) is developing the Sunway exascale system that likely will use an upgraded processor based on the ShenWei 26010 chip used in the NRCPC-designed system, the Sunway TaihuLight HPC, installed at the National Supercomputing Center in Wuxi, China in June 2016. That system was the fastest HPC in the world at that time. The yet to be named system will be a CPU-only design with no special accelerators.
- Sugon is developing a heterogenous HPC architecture consisting of 2 CPUs and 2 DCU accelerators per node. The Hygon CPU likely is a licensed clone of the US-based AMD Gen 1 EYPC processors, and the DPU is an internal Sugon design. The Sugon machine is using a unique board-level liquid immersion system that likely will support high levels of board level power and cooling capabilities. Sugon is the only commercial Chinese entity involved in the development of next generation Chinese exascale systems.

Japan has already demonstrated its world-class HPC capability in its first exascale class system, the Fugaku machine at Riken, Japan's largest and most comprehensive research organization for basic and applied science. The Fugaku system recently became the fastest HPC in the world, based on the June 2020 Top500 HPC list, with a High Performance Linpack result of 415.5 PFlops. The Fugaku system is the first near-exascale class system and the first supercomputer to cost around $1 billion dollars.

Although not strictly an exaflops system, the overall design features of the Fugaku ensures that the system will likely outperform many other exascale systems with higher peak values on real-world application for at least the next few years. Key features of the system include the use of a specially designed Fujitsu A64 ARM8.2A processors consisting of 48/52 compute cores, each with GPU-like vector intension and in-die network hardware that provides ~400 Gbps bandwidth across the world-class 6-D Tofu-D network. Hyperion Research expects that smaller, less powerful, and less expensive versions of this system will be successfully installed by Fuitsu at a number of leading Japanese, and perhaps foreign research facilities in the next few years.

In the EU, a key exascale development program, EuroHPC, was stood up in 2018 to develop EU-wide HPC development capabilities and use. Membership includes 33 participating States and the EU, represented by the Commission. Three sites have already been selected for 150-200 Pflops systems,

with installations to start in 4Q2020, at Kajaani, Finland; Barcelona, Spain; and Bologna, Italy. Total investment is expected to be 650 million euros, with 50% funding coming from the EU and the rest from consortium members.

Longer-term EU plans call for the acquisition of two exascale systems in the 2022-2023 time frame, with at least one to use European technology: specifically, using a processor developed under the European Processor Initiative. Hyperion Research projects there could be two additional, non-EuroHPC exascale procurements in Germany in 2024 and 2025, while EU plans may include two additional exascale systems in the 2023-2026 time frame at yet to be determined sites.

The UK has not joined EuroHPC and will likely not be eligible to fully take part in EuroHPC projects or access calls when Horizon 2020 ends this year. There is a UK-only exascale initiative in the works to address UK HPC modeling and simulation requirements as well as newer AI-based deep learning applications for both scientific and industrial users. Any such system would likely be targeted for operation by Q42022/Q12023 with a price tag possibly exceeding $500 million.

Finally, Hyperion Research expects that as exascale performance becomes more common and affordable over time, a number of tier-two HPC nations will either develop or purchase exascale class systems to support their overall S&T-related scientific, national security, and economic competitiveness agendas. Nations with the technical wherewithal include Australia, Canada, India, Israel, Russia, Saudi Arabia, South Korea, and Singapore.

FUTURE OUTLOOK

Exascale development is happening around the world as developers seek to build indigenous HPC ecosystems and skills meet the requirements for ever faster HPC performance and address an increasing range of HPC workloads. Being an exascale trailblazer can be expensive, with systems on the near horizon costing $500 million or more each. To that end, only the most ambitious governments can commit to participating in the early rounds of exascale progress. However, as HPC price/performance moderates, a long-verified phenomenon in the HPC space, additional nations will soon also be looking to provide their domestic R&D base access to exascale-class HPCs either developed domestically or procured from foreign sources.

Hyperion Research expects that building second-generation 10-100 exaflops systems could potentially cost $1 billion or more (including both the R&D and system acquisitions costs), requiring considerable up-front return on investment analysis from both a national security and economic competitiveness perspective. These costs can be mitigated through carefully planned exascale programs that seek not only to enable important advances in leading-edge scientific and engineering use cases, but also spur development in innovative hardware, software, and related application skills that can have broad applicability for a wide class of less costly and more accessible HPCs, and in some cases advance the state-of-the-art in IT markets beyond HPC.