Supercomputer debuts as world's fastest, breaking exascale barrier

May 30 2022

Frontier has arrived, and ORNL is preparing for science on Day One. Credit: Carlos Jones/ORNL, U.S. Dept. of Energy

The Frontier supercomputer at the Department of Energy's Oak Ridge National Laboratory earned the top ranking today as the world's fastest on the 59th TOP500 list, with 1.1 exaflops of performance. The system is the first to achieve an unprecedented level of computing performance known as exascale, a threshold of a quintillion calculations per second.

Frontier features a theoretical peak performance of 2 exaflops, or two quintillion calculations per second, making it ten times more powerful

than ORNL's Summit system. The system leverages ORNL's extensive expertise in accelerated computing and will enable scientists to develop critically needed technologies for the country's energy, economic and [national security](#), helping researchers address problems of national importance that were impossible to solve just five years ago.

"Frontier is ushering in a new era of exascale computing to solve the world's biggest scientific challenges," ORNL Director Thomas Zacharia said. "This milestone offers just a preview of Frontier's unmatched capability as a tool for scientific discovery. It is the result of more than a decade of collaboration among the national laboratories, academia and [private industry](#), including DOE's Exascale Computing Project, which is deploying the applications, software technologies, hardware and integration necessary to ensure impact at the exascale."

Rankings were announced at the [International Supercomputing Conference 2022](#) in Hamburg, Germany, which gathers leaders from around the world in the field of [high-performance computing](#), or HPC. Frontier's speeds surpassed those of any other supercomputer in the world, including ORNL's Summit, which is also housed at ORNL's Oak Ridge Leadership Computing Facility, a DOE Office of Science user facility.

Frontier, a HPE Cray EX supercomputer, also claimed the number one spot on the Green500 list, which rates energy use and efficiency by commercially available supercomputing systems, with 62.68 gigaflops performance per watt. Frontier rounded out the twice-yearly rankings with the top spot in a newer category, mixed-precision computing, that rates performance in formats commonly used for artificial intelligence, with a performance of 6.88 exaflops.

The work to deliver, install and test Frontier began during the COVID-19 pandemic, as shutdowns around the world strained

international supply chains. More than 100 members of a public-private team worked around the clock, from sourcing millions of components to ensuring deliveries of system parts on deadline to carefully installing and testing 74 HPE Cray EX supercomputer cabinets, which include more than 9,400 AMD-powered nodes and 90 miles of networking cables.

"When researchers gain access to the fully operational Frontier system later this year, it will mark the culmination of work that began over three years ago involving hundreds of talented people across the Department of Energy and our industry partners at HPE and AMD," ORNL Associate Lab Director for computing and computational sciences Jeff Nichols said. "Scientists and engineers from around the world will put these extraordinary computing speeds to work to solve some of the most challenging questions of our era, and many will begin their exploration on Day One."

Frontier's overall performance of 1.1 exaflops translates to more than one quintillion floating point operations per second, or flops, as measured by the [High-Performance Linpack Benchmark test](#). Each flop represents a possible calculation, such as addition, subtraction, multiplication or division.

Frontier's early performance on the Linpack benchmark amounts to more than seven times that of Summit at 148.6 petaflops. Summit continues as an impressive, highly ranked workhorse machine for open science, listed at number four on the TOP500.

Frontier's mixed-precision computing performance clocked in at roughly 6.88 exaflops, or more than 6.8 quintillion flops per second, as measured by the [High-Performance Linpack-Accelerator Introspection](#), or HPL-AI, test. The HPL-AI test measures calculation speeds in the computing formats typically used by the machine-learning methods that drive advances in artificial intelligence.

Detailed simulations relied on by traditional HPC users to model such phenomena as cancer cells, supernovas, the coronavirus or the atomic structure of elements require 64-bit precision, a computationally demanding form of computing accuracy. Machine-learning algorithms typically require much less precision—sometimes as little as 32-, 24- or 16-bit accuracy—and can take advantage of special hardware in the graphic processing units, or GPUs, relied on by machines like Frontier to reach even faster speeds.

ORNL and its partners continue to execute the bring-up of Frontier on schedule. Next steps include continued testing and validation of the system, which remains on track for final acceptance and early science access later in 2022 and open for full science at the beginning of 2023.

Facts about Frontier

The Frontier supercomputer's exascale performance is enabled by some of the world's most advanced pieces of technology from HPE and AMD:

- Frontier has 74 HPE Cray EX supercomputer cabinets, which are purpose-built to support next-generation supercomputing performance and scale, once open for early science access.
- Each node contains one optimized EPYC processor and four AMD Instinct accelerators, for a total of more than 9,400 CPUs and more than 37,000 GPUs in the entire system. These nodes provide developers with easier capabilities to program their applications, due to the coherency enabled by the EPYC processors and Instinct accelerators.
- HPE Slingshot, the world's only high-performance Ethernet fabric designed for next-generation HPC and AI solutions, including larger, data-intensive workloads, to address demands for higher speed and congestion control for applications to run smoothly and boost performance.

- An I/O subsystem from HPE that will come online this year to support Frontier and the OLCF. The I/O subsystem features an in-system storage layer and Orion, a Lustre-based enhanced center-wide file system that is also the world's largest and fastest single parallel file system, based on the Cray ClusterStor E1000 storage system. The in-system storage layer will employ compute-node local storage devices connected via PCIe Gen4 links to provide peak read speeds of more than 75 terabytes per second, peak write speeds of more than 35 terabytes per second, and more than 15 billion random-read input/output operations per second. The Orion center-wide file system will provide around 700 petabytes of storage capacity and peak write speeds of 5 terabytes per second.
- As a next-generation supercomputing system and the world's fastest for open science, Frontier is also energy-efficient, due to its liquid-cooled capabilities. This cooling system promotes a quieter datacenter by removing the need for a noisier, air-cooled system.

Provided by Oak Ridge National Laboratory