

CASE STUDY

High Performance Computing (HPC)
2nd Generation Intel® Xeon® Scalable Processor
Intel® Optane™ DC Persistent Memory



Texas Advanced Computing Center Installs Frontera for Massive Scale Computing

448,448-core cluster of 2nd Generation Intel® Xeon® Platinum processors will help solve the world's biggest challenges

Frontera Highlights:

- 8,008 dual-socket Dell PowerEdge* C6420 servers with 2nd Generation Intel® Xeon® Scalable Processors (448,448 cores total)
- Peak performance of 38.7 petaFLOPS¹
- 50 nodes with Intel® Optane™ DC Persistent Memory
- #5 most powerful supercomputer in the world, and the fastest at any university



Executive Summary

The [Texas Advanced Computing Center \(TACC\)](#) continuously re-invents supercomputing at larger and larger scale to enable breakthrough research and deliver the resources that scientists need. Frontera, a 38.75 petaFLOPS cluster, that earned the #5 ranking on the June 2019 Top500 list,¹ is its latest supercomputing system comprising nearly a half-million cores of 2nd Generation Intel® Xeon® Scalable processors inside Dell EMC PowerEdge* servers.

Challenge

The Texas Advanced Computing Center (TACC) is a world-renowned facility for supercomputing, enabling new discoveries across a range of disciplines in science and industry.

“Our mission here at the Texas Advanced Computing Center,” said TACC’s Executive Director, Dr. Dan Stanzione, “is to provide groundbreaking new computing capabilities to enable new kinds of scientific discoveries, and new kinds of engineering research.”

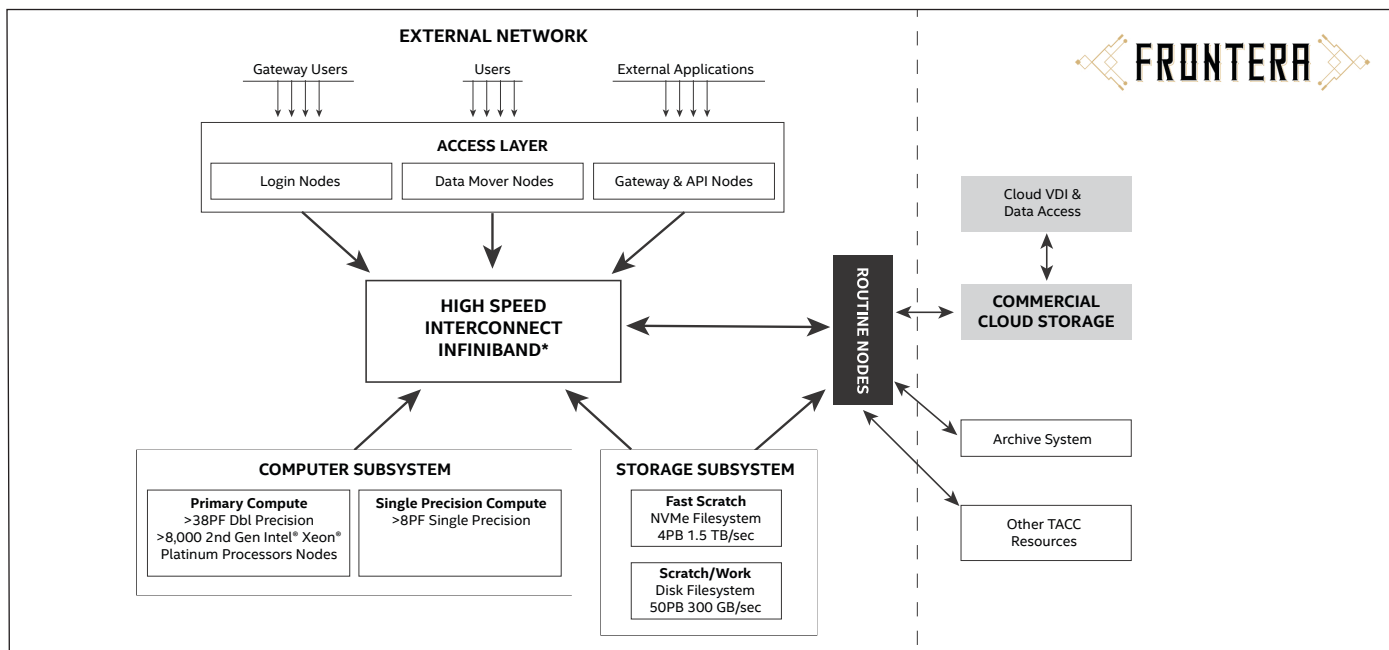
Deployed in 2017, TACC’s Stampede2 supercomputer incorporated the latest Intel® Xeon® Scalable Processors inside Dell EMC PowerEdge servers and including Intel® Omni-Path architecture fabric. Designed as a capability machine, Stampede2 will support three to four thousand projects over its lifetime. But, every few years, TACC looks at the kinds of problems that researchers are tackling and what types of architecture will offer the best support for that science. Some of those problems address the ‘grand challenges’ of our time and require computing on a massive scale.

“We’re looking at control problems around fusion reactors,” commented Stanzione as he offered an example of the kinds of massive scale research that will require new levels of supercomputing performance. “We’re looking at mantle convection as a whole Earth problem, where you see single simulations across the entire planet.”

Such a scale of problems requires a different scale of supercomputer than Stampede2.

Solution

Frontera is TACC’s newest supercomputer, supported by a \$60 million award from the U.S. National Science Foundation. It contains a large main system that will deliver peak performance of 38.71 petaFLOPS, according to Stanzione. The main system is built on the 2nd Gen Intel® Xeon® Platinum Processor with 8,008 dual-socket nodes of 56 cores per node, interconnected by InfiniBand* Architecture at 100 Gbps. Its 448,448 cores give TACC more computing capacity and memory capacity than the center has had in the past.



Frontera hardware and software system overview.

By selecting Intel's latest server processor, Frontera offers:

- A higher clock rate than previous systems, delivering higher single-thread performance,
- More processor cores to run more threads at the same time, and
- More memory bandwidth that can feed data to all those cores.

“Frontera will address a narrower mission than Stampede2,” explained Stanzione. “Instead of supporting thousands of projects, we’ll have a few hundred that have an extraordinary computational need and massive scale of computation. It’ll solve the very biggest sort of grand challenge projects in the scientific ecosystem. We’ll be running calculations at a speed and at a scale that we’ve never been able to do before.”

Frontera will also support new technologies previously unavailable, including Intel® Deep Learning Boost (Intel® DL Boost) targeted for artificial intelligence workloads. These new technologies will help TACC supercomputer designers understand better which of these are useful to researchers, so the technologies can be integrated into the next next-generation TACC machine slated for 2025. One such technology is Intel® Optane™ DC Persistent Memory.

“Intel Optane DC persistent memory,” commented Stanzione, “has several unique characteristics for us that offer advantages over traditional memory and advantages over traditional storage. There are many potential interesting use cases, such as very, very large memory nodes—multiple terabytes per node—or simple fault tolerance. When a server fails, we can keep the state of memory and allow the computation to keep running, versus having to restart it across the whole 8,008 nodes that make up the machine.”

Result

Grand challenge problems need massive computing capacity.

“It’s going to be a remarkably productive system,” said Stanzione. “We think, in terms of real science throughput, we’ll get three or four times the performance of its predecessor.”

Beyond the Standard Model

With the discovery of the Higgs boson using the Large Hadron Collider (LHC) at CERN in Geneva, Switzerland, the final piece of the Standard Model of Physics was put in place. Now, scientists around the world are looking Beyond the Standard Model to gain a finer sense of what makes up high-energy particle physics. The LHC, with one of its detectors called ATLAS (A Toroidal LHC ApparatuS), will again be at the center of their research. CERN plans on increasing the number of LHC collisions by a factor of ten in the coming years.

The LHC requires enormous amounts of computing capacity to interpret its collisions. CERN scientists have run workloads on Stampede2. Now that Frontera is operational, CERN will have a much larger system to use to understand what is happening at these subatomic scales.

“We simulate the detector response to a given physics model,” said Robert Gardner, a research professor in the Enrico Fermi Institute at the University of Chicago, who co-leads the distributed computing facility group for the U.S. ATLAS collaboration.

“When we’re doing the analysis on the actual data, we may plot some distributions such as the particle mass, transverse momentum, or the ‘missing energy’ in the collision. And you get the number of candidates that we have for the raw data coming off the detector. Then we compare those to different kinds of models and see if we can match up the distributions. This provides clues to what might be actually happening during the collisions.”

From Nuclear Fission to Fusion Power

Another area involving global scientific collaboration is innovating new resources for supplying the world’s power needs. From more efficient wind generation to battery research and hydrogen mining from water, science is trying to find clean alternatives to fossil fuels.

Nuclear fusion—the merging of nuclei to release massive amounts of energy, like Earth’s Sun does—is considered the holy grail of energy production, without the drawbacks

of today's fission reactors. In France, such a reactor—the International Thermonuclear Experimental Reactor (ITER)—is being built by a consortium of seven governments. Scheduled for a 2025 completion date, it is designed to produce 20 to 25 times more power than it uses.

An urgent problem for designers is to be able to accurately and reliably predict—and avoid—large-scale disruptions. But for years, scientists have struggled to match physics models and simulations with the dynamics in a real reactor.

“If you try to use conventional theoretical methods, buttressed by high performance computing, you still aren't going to be able to make predictions,” said William Tang, principal research physicist at the Princeton Plasma Physics Laboratory—the U.S. DOE National Lab for fusion studies. “You needed the impact of big data analytics that can deal with a lot of data that's relevant to disruptions.”

Tang and his team have turned to Artificial Intelligence to help solve the problem. The team developed the Fusion Recurrent Neural Net (FRNN) Code, deploying deep learning for better predictions. Their code can predict disruption events with 90+ percent accuracy more than 30 milliseconds ahead of the disruption trigger event. Tang will take advantage of Frontera's new resources for deep learning to further his research with the FRNN code and develop a control system that can avoid disruptions in ITER.

Computation for World Problems

Other challenges requiring massive computing scale include using precision agriculture and genomics to feed the world's growing population and innovating cleaner coal combustion, which is still a leading source of energy.

“We need systems like Frontera to answer the big questions of our time, such as the sustainability of the environment and renewable energy,” said Professor Gardner. “We have to continue to work on frontier science and everything that comes after it, and we can't do that without computation.”

Solution Summary

Frontera was built to support a new, much larger scale of scientific computing than TACC previously was able to. Built on 2nd Generation Intel Xeon Platinum Processors inside Dell EMC PowerEdge servers, with nearly half a million cores, Frontera will deliver a peak performance of 38.7 petaFLOPS, according to TACC's Executive Director Dan Stanzione.

The new supercomputer will also allow scientists to test new technologies, including Intel Optane DC Persistent Memory, to assess how the supercomputing center might implement these technologies on their next next-generation supercomputer.

Where to Get More Information

Find out more about 2nd Generation Intel® Xeon® Scalable Processor family at <https://www.intel.com/content/www/us/en/products/processors/xeon/scalable.html>.

Learn more about Intel® Optane™ DC Persistent Memory at <https://www.intel.com/content/www/us/en/architecture-and-technology/optane-dc-persistent-memory.html>.

Find out more about Frontera at <https://www.tacc.utexas.edu/systems/frontera>.

Learn more about Dell EMC PowerEdge Servers at dell.com/hpc.

Solution Ingredients

- 8,008 Dell EMC PowerEdge C6420 compute nodes, consisting of 2nd Generation Intel® Xeon® Platinum processors, 56 cores per node
- Intel® Optane™ DC Persistent Memory



A view between two rows of Frontera servers in the TACC Data Center.



¹ Testing conducted by TACC for July 2019 TOP500 rating. See <https://www.top500.org/system/179607>.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at <http://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-architecture-fabric-overview.html>.

Intel, the Intel logo, Xeon, Optane and Phi are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

0819/RJ/KR/PD



Please Recycle

340938-002