

Advancing research with a liquid-cooled supercomputer.

How **Harvard University Faculty of Arts & Sciences Research Computing** accelerates ground-breaking research with a new high-performance computing cluster based on Lenovo ThinkSystem servers, 2nd Gen Intel® Xeon® Scalable processors, and Lenovo Neptune™ liquid cooling technology.

**Lenovo Infrastructure Solutions
for The Data-Centered**



Lenovo

1

Background

Harvard University is devoted to excellence in teaching, learning, and research, and to developing leaders in many disciplines who make a difference globally. The university, based in Cambridge, Boston, Massachusetts, has an enrollment of over 20,000 degree candidates, including undergraduate, graduate, and professional students.

Harvard's Faculty of Arts & Sciences Research Computing (FASRC) center was established in 2007 with the founding principle of facilitating the advancement of complex research by providing leading-edge computing services.

2

Challenge

FASRC provides researchers with the high-performance computing (HPC) resources they need to process massive data sets, perform complex calculations, and answer important questions in science, engineering, mathematics, medicine and dozens of other disciplines.

To give researchers the very best tools to support their work, and to keep up with growing demand for its services, FASRC refreshes its HPC infrastructure every few years.

FASRC processes more than 290 million jobs a year, with 15,000 jobs running on the cluster at any one time. Researchers need results quickly, so that they can gain new insights, iterate on their experiments, and further their work.



“We aimed to increase our processor count to meet growing demand. We also decided to increase the performance of each individual processor, since 25% of CPU hours are consumed by thousands of single-core computations that are loosely coupled.”

Scott Yockel

University Research Computing Officer, Harvard University



Choosing a new liquid-cooled cluster design.

Wanting to take full advantage of the latest advances in CPU technology with higher wattages, while also enabling more performance per core, FASRC deployed a water-cooled supercomputer cluster from Lenovo.

Yockel comments: “Our previous cluster was air-cooled, so moving to Lenovo Neptune™ liquid cooling technology represented a big change. Liquid cooling supports increased levels of performance much more efficiently, which is crucial to meeting both our current and future computing needs.”

Building a state-of-the-art HPC system.

The new system, named Cannon, is in honor of the pioneering astronomer Annie Jump Cannon, and comprises of 72 Lenovo NeXtScale n1200 enclosures over 12 racks housing 670 Lenovo ThinkSystem SD650 servers with direct-to-node water cooling. Each server is equipped with the Intel® Xeon® Platinum processor family and 192 GB of RAM, giving Cannon a total of 32,160 compute cores. The servers are then clustered together using InfiniBand HDR 100 Gbps fabric. The installation was supported by Lenovo Professional Services and completed on schedule.

The Neptune™ direct-to-node water-cooling technology removes heat from the CPUs, memory, IO, local storage and voltage regulators using a copper-based water loop. This enables FASRC to run the CPUs at a clock rate of 3.5 GHz, compared to their 2.90 GHz base frequency, without any additional air cooling. At 2.076 PetaFLOPs in Linpack performance, Cannon is currently ranked 186 in the TOP500 list of the world's fastest supercomputers.¹

FASRC occupies around 10,000 square feet across three data centers. The primary cluster Cannon is located at the Massachusetts Green High Performance Computing Center (MGHPCC) in Holyoke. Storage and login nodes, virtual machines and specialty computing resources are split between Harvard's Boston and Cambridge campuses, all interconnected through the Northern Crossroads (NOX) network.

The Cannon cluster runs CentOS Linux with Puppet for cluster configuration management and SLURM Workload Manager for job scheduling.

FASRC continues to add Lenovo ThinkSystem servers to the cluster in response to growing demand from researchers. Faculty and research groups can use their own funding to purchase additional nodes, to which they have priority access. "The Lenovo system is easy to expand," confirms Yockel.

¹ Ranked 186 in June 2020 TOP500 (<https://www.top500.org/system/179740/>).

3

Results

With the Cannon cluster, researchers from across Harvard have access to world-class HPC resources via FASRC.

Yockel elaborates: “Our new Cannon cluster delivers four-times greater performance than our previous infrastructure within the same physical footprint, yet it only requires 50% more power.² This is thanks in large part to the direct-to-node water-cooling design, as it enables us to run the Intel® Xeon® Scalable processors at 3.5 GHz for 85% of the time without them overheating. This has considerably increased our processing power, so we can run more jobs faster.”

Today, the Cannon cluster supports thousands of research projects. Significant users include the Center for Brain Science and Center for Astrophysics.

Yockel says: “One example of work currently being done using the Cannon cluster is investigating the relationship between movement and vision in the brain. Researchers use implanted electrodes to measure brain activity in the primary visual cortex as rats moved in an enclosure. The researchers then used that data to create 3D models of the brains and run machine-learning algorithms to study how neurons transmit signals.”

Similarly, the Center for Astrophysics uses the Cannon cluster to process hundreds of terabytes of telescope images to study black holes and uncover new insights. In recent months, FASRC has onboarded a number of new research projects focused on COVID-19, from epidemiologists studying transmission rates to economists examining the financial impact of the pandemic.



4x greater performance than previous cluster with only a 50% increase in power consumption



32,160 compute cores

² 4x performance increase over previous Intel® Xeon® E5 v4 processor family nodes. Data provided by Harvard University Faculty of Arts & Sciences Research Computing.



“Science is all about iteration and repeatability. But iteration is a luxury that is not always possible in the field of university research because you are often working against the clock to meet a deadline. With the increased compute performance and faster processing of the Cannon cluster, our researchers now have the opportunity to try new things, fail, and try again. Allowing failure to be an option makes our researchers more competitive. FASRC is dedicated to furthering research and we are confident that as demand for HPC resources continues to grow, the Lenovo system will support us for years to come.”

Scott Yockel

University Research Computing Officer, Harvard University

What will you do with Lenovo HPC solutions?

The Data-Centered support ground-breaking academic research with
Lenovo smarter infrastructure solutions, powered by Intel®.

[Explore HPC Solutions](#)



Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo.

Intel, the Intel logo, Xeon, and Xeon Inside are trademarks of Intel Corporation or its subsidiaries in the U.S.
and/or other countries.

Other company, product and service names may be trademarks or service marks of others.

© Lenovo 2021. All rights reserved.