DELL EMC intel

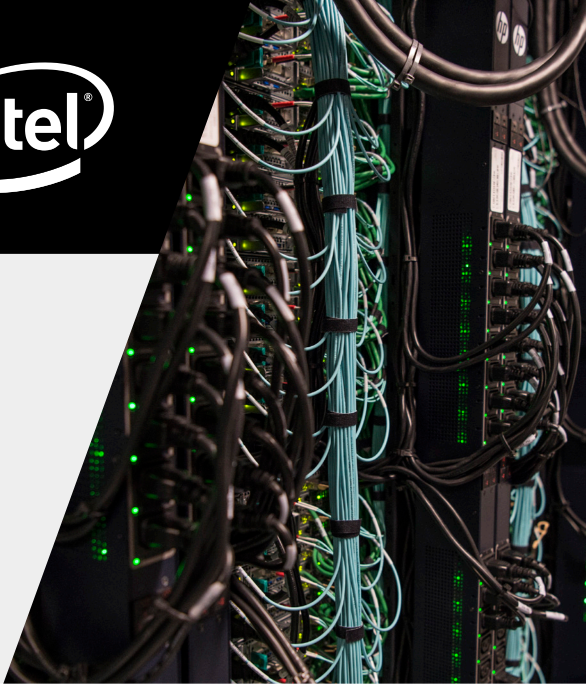

PETAFLOP PERFORMANCE

MIT Lincoln Laboratory Supercomputing Center unveils Intel-based Dell EMC system to accelerate the Nations' research needs in autonomous systems, device physics, and machine learning.

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Academic and Federal Research | United States

Business needs

The MIT Lincoln Laboratory Supercomputing Center (LLSC) needed a next-generation supercomputer to power the work of thousands of scientists and engineers.

Solutions at a glance

- 648 Dell EMC compute notes based on many-core Intel® Xeon Phi™ processors
- Director Class Intel® Omni-Path Architecture network fabric

Business results

- The HPC team from Dell EMC delivered a petascale system in less than a month.
- The LLSC achieved a 6x gain in supercomputing performance

1.03
petaflops
HPL benchmark performance

1.7
petaflops
Peak performance capability

Meeting the computational needs of thousands

MIT Lincoln Laboratory supports research and development aimed at solutions to problems that are critical to the Nation. The research spans diverse fields such as space observations, robotic vehicles, communications, cyber security, machine learning, sensor processing, electronic devices, bioinformatics, and air traffic control. Lincoln Laboratory takes projects from the initial concept stage, through simulation and analysis, to designing and building working prototypes.

The MIT Lincoln Laboratory Supercomputing Center (LLSC) addresses the supercomputing needs of thousands of MIT scientists and engineers by providing interactive, on-demand supercomputing and big data capabilities with zero carbon footprint. The LLSC mission is to address supercomputing needs across all Lincoln Laboratory missions, develop new supercomputing capabilities and technologies, and foster close collaborations with MIT campus supercomputing initiatives.

The supercomputing environment

The LLSC manages a diverse portfolio of supercomputing resources that includes Intel- and AMD-based servers, NVIDIA® GPUs, field-programmable gate arrays (FPGAs) and now Intel® Xeon Phi™ many-core systems. These supercomputing resources are integrated with large parallel files systems from Seagate and DDN, big data file systems, such as the Apache™ Hadoop® platform, and high-speed data networks.

At an application level, the LLSC focuses on providing interactive development environments for rapid prototyping, such as Python, Matlab, GNU Octave, Mathematica, R, and Julia. In addition, the LLSC has had extensive success with merging supercomputing hardware with big data tools, such as map reduce, and high performance databases, such as Apache Accumulo and SciDB.

Targeting petaflop performance

When they set out to build their latest supercomputer, the LLSC team members focused on creating a petascale system, which equates to performance greater than one petaflops, or one quadrillion floating point operations per second.

The technology and vendor selection process led the LLSC to Dell EMC and systems based on Intel® Xeon Phi™ processors and the Intel® Omni-Path Architecture network fabric. The team felt that the Dell EMC solution offered the robust combination of compute performance, network bandwidth and software support needed for the project to meet its long-term objectives of enabling computationally intensive research in the areas of autonomous systems, device physics, and machine learning.

The LLSC's goals for system performance and project timelines were aggressive. At a systems level, the goal was to create petascale supercomputer based on 648 Intel Xeon Phi processor-enabled servers with a dual high-speed OmniPath and 10GB fiber interconnect.
At a project level, the goal was to deliver, deploy and benchmark the system in just one month. The LLSC reached its goals when its supercomputer was online in less than a month and running at 1.032 petaflops using the High Performance Linpack (HPL) Benchmark. The system, which has approximately 40,000 cores, has a peak capacity of approximately 1.7 petaflops.

"Dell EMC has been a great partner of ours for a long time," notes LLSC Systems Engineer Matthew Hubbell. "Dell EMC's high performance computing team is very knowledgeable, and showed a lot of experience. They were extremely responsive. They were critical in our ability to deliver a one-petaflop system in under a month. That was something that could only have been achieved by a very well organized and capable team that came together and provided the resources to really exceed our expectations."

Tackling ever-larger problems

The value of the LLSC's new supercomputer is rooted deeply in the needs of the research scientists and engineers who continually take on bigger and harder problems, according to LLSC Manger Albert Reuther.

"The goal we had for this installation and this project was to reach a petaflop—a real High Performance Linpack petaflop—because many of our applications map reasonably well to HPL," Reuther says. "That, in turn, challenges our users to think big in terms of what they might be able to do with that kind of capability at that scalability."

For the LLSC team and the scientists and engineers it serves, challenges feed on challenges. "When our users are challenged, they start thinking big, and that challenges us back," Reuther says. "And that is exactly the situation we want to be in to do the really hard problems, the MIT hard problems that will have an impact on our national security."

Jeremy Kepner, Laboratory Fellow and head of the LLSC, notes that the capabilities of the supercomputer will have a big impact on the work done by researchers focused on machine learning, device simulation and autonomous systems—three current areas of emphasis for the center.

"We're very excited to push the boundaries of those three areas," Kepner says. "These researchers really need the kind of supercomputing we can provide with the new Dell EMC system with the revolutionary Intel Knights Landing processor. Each processor can do thousands of operations at once, while also being connected together with an incredibly powerful network."

Interactive supercomputing

With the new supercomputer, researchers with the LLSC and MIT now have access to about six times more compute power when compared to their current system, but that doesn't mean that everyone who uses the lab is a supercomputing pro. The LLSC team introduces many new users to parallel computing—and it does so very quickly.

"One of our aims is to make supercomputing easier, more accessible," says LLSC researcher Vijay Gadepally. "We want to get the average user onto a supercomputer in about an hour."

In many cases, new users run their first real parallel computing jobs after a 45-minute consultation with an LLSC team member. Moreover, the supercomputing system is seamlessly integrated into the LLSC's computing fabric, so users can access supercomputing resources directly from their desktop systems.

"We want to enable all researchers," Gadepally says. "At the Lincoln Laboratory Supercomputing Center, that's very core to our mission."

"These researchers really need the kind of supercomputing we can provide with the new Dell EMC system with the revolutionary Intel Knights Landing processor."

Jeremy Kepner
Laboratory Fellow and head of the
MIT Lincoln Laboratory Supercomputing Center

Supercomputing with zero carbon impact

In a unique twist, LLSC provides its supercomputing and big data capabilities with a zero carbon footprint. Its power source from Holyoke Gas & Electric is 99.7 percent carbon free, according to Hubbell. Most of the power consumed by the center comes from a hydroelectric dam that is within a half a mile of the LLSC facility.

"Holyoke Gas & Electric also provides a large portfolio of renewable resources—from wind and solar to other green energy sources," Hubbell says. "All of the power required to run the supercomputing center is derived from these renewable energy sources."

Power sources aside, the center itself operates in an extremely energy-efficient manner. For example, the center relies heavily on outside air to cool its servers, rather than using mechanical cooling. It also relies heavily on dense servers that pack lots of computational power into a small form factor.

Looking ahead

As they look to the future, the LLSC team expects to continue to grow and evolve its interactive, on-demand supercomputing and big data capabilities to advance autonomous systems, device physics, and machine learning. They believe that the combination of the Intel Xeon Phi processors and the Intel Omni-Path network fabric will allow the LLSC to achieve larger simulations and tackle the ever-increasing and important research objectives for thousands of researchers at MIT.

"One of the most exciting things about this new system is really being able to identify where this new many-core Intel architecture and the new interconnects will take the technology, and how we can take advantage of it all," Hubbell says. "It's really a fun and exciting task ahead of us."

And, of course, the LLSC team is already thinking about the work it needs to do to meet the evolving needs of its users. History suggests that the users will continue to come up with larger and larger datasets, bigger and harder problems, that will require bigger and faster supercomputers.

"Our users will probably be satisfied for a few months, and then they will start dreaming up the next big problems," Reuther says. "And we can't wait to find what those next big problems are. We're already listening to our users as to what are the scales of the simulations they want to run, and how we can bring supercomputing to them in an on-demand, interactive way that is as easy for them to use as their desktop and laptop systems."

Intel Inside®. Powerful Solution Outside.

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Learn more about Dell EMC High Performance Computing solutions
Dell.com/HPC

View all customer stories at
Dell.com/CustomerStories

Share this story