CASE STUDY

High Performance Computing (HPC)
Intel® Omni-Path Architecture

MIT Supercomputer, TX-Green, Enables New Breakthroughs in Interactive Supercomputing

40,000+ core Intel® Xeon Phi™ processor-based Dell EMC cluster delivers insights in machine learning training

MIT Lincoln Lab Supercomputing Center's New TX-Green Pioneer Machine Learning Cluster:

- 648 nodes of bootable Intel® Xeon Phi™ processor-based servers
- Interconnected with Dell Networking H-Series Chassis Switch based on Intel® Omni-Path Architecture
- Designed for training phase of machine learning
- Relies on non-carbon energy sources—solar, wind, hydro, and nuclear

Executive Summary

Massachusetts Institute of Technology (MIT) Lincoln Lab Supercomputing Center (LLSC) was developed to enhance computing power and accessibility for over 1,000 researchers across the Institute. Its mission is to address supercomputing needs, develop new supercomputing capabilities and technologies, and to collaborate with MIT campus supercomputing initiatives. Their newest system, called TX-Green Pioneer, is an interactive, on-demand parallel computing cluster that enables Laboratory researchers to augment the power of desktop systems. LLSC collaborates closely with the MIT Beaver Works Center for Engaging Supercomputing.

Challenge

MIT Lincoln Lab has used a methodology they call "interactive supercomputing" for decades, augmenting researcher's desktops and workstations with large parallel supercomputing clusters. Recently, the LLSC has evolved into the MIT SuperCloud—a fusion of multiple ecosystems comprising supercomputing, enterprise computing, big data, and traditional databases—a coherent, unified platform for rapid prototyping across all four systems.

"Our approach to interactive supercomputing," said Jeremy Kepner, laboratory fellow and head of the LLSC, "makes fast turnaround possible not only for rapid prototyping, but for all projects designed for parallel processing on new HPC platform architectures." MIT is known for its quick prototyping capabilities, enabled by supercomputing for analysis and simulation, artificial intelligence (AI), and big data.

Engineers, scientists, faculty, and students use LLSC's resources in a wide range of fields, including space observation, robotic vehicles, communications, cyber security, sensor processing, electronic devices, bioinformatics, and air traffic control. Three areas of extreme interest that demand more computational capabilities include autonomous systems, device physics, and machine learning.

Data analysis is a major focus for most of the Lab's users, according to Kepner. "Our long-term familiarity with Big Data allows us to apply the principles of interactive supercomputing to data collection and analysis, designing sensors, and developing algorithms. The right combinations of mathematics and processors is beneficial to our users," he added.

So, when MIT LLSC was designing their next supercomputer, they looked for a solution for high-performance data analytics, and specifically for the learning phase of Deep Learning/Artificial Intelligence. "Early in 2016, we talked to MIT about the Intel® Xeon Phi™ processor, which is a sweet spot for machine learning, and about Intel® Omni-Path Architecture," said Gretchen Stewart, Intel's representative for MIT's Lincoln Lab. "We felt we could easily meet their performance numbers with an Intel-based solution that would land them on the Top500." Other Intel® Xeon® processor E5 v4-based HPE systems already existed for

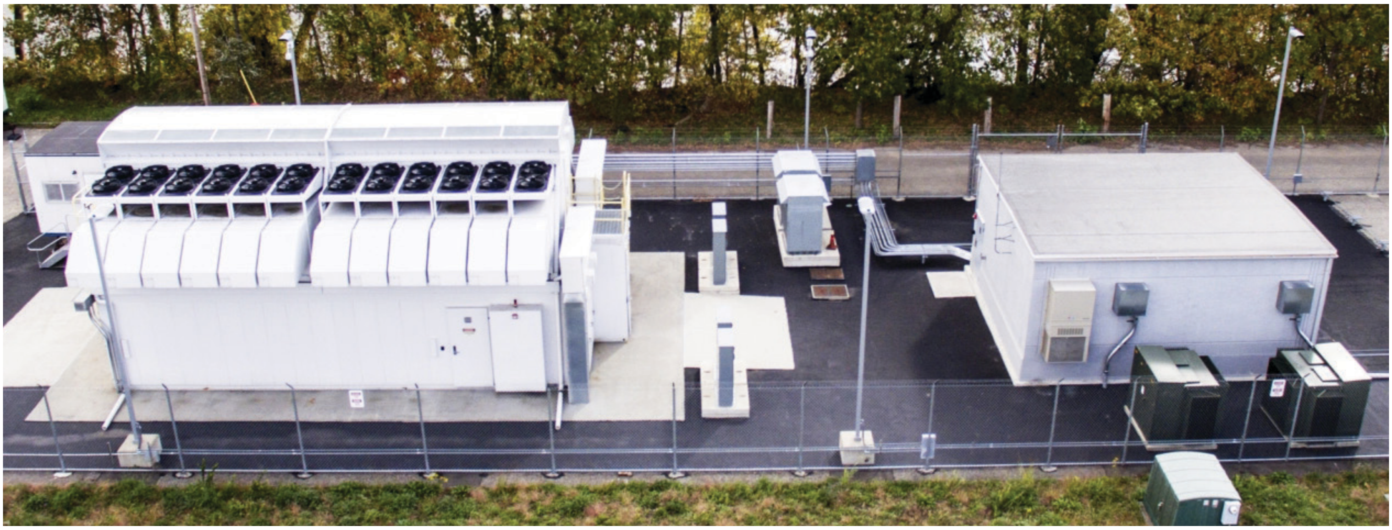

This EcoPOD was assembled on site in three months, providing space for 44 racks that accommodated up to 24,000 hard drives, and featured security, fire suppression, and monitoring systems. It offers energy-efficiencies, such as an adaptive cooling system, to help keep operational costs down.

running the inference phase of Machine Learning. The new cluster would be installed alongside these servers at MIT.

Solution

LLSC selected Dell EMC to install a 648 node HPC system built on bootable Intel Xeon Phi processors. The cluster included Intel Omni-Path Architecture as the high-speed fabric between compute nodes. The 1+ petaFLOPS system contains 41,472 cores.

LLSC's new TX-Green Pioneer system, one of the largest of its kind on the US East Coast, was installed in record time. "Dell EMC has been a great partner in enabling us to dramatically increase the capabilities of our supercomputing center," commented Kepner. "The Dell HPC and Intel teams were responsive and able to deliver, install, and benchmark our petaFLOPS-scale system in less than a month. This was a great example of a well-coordinated and dedicated organization that was able to allocate the appropriate resources to exceed customer expectations."

"In the world of supercomputers, you don't order, get up and running, and get certifiable test results in such a short time. It's unheard of," added Gretchen Stewart.

Space was at a premium at the Lab. And the costs to acquire, build, and manage a data center for the cluster would be high. Instead, Kepner proposed a prefabricated alternative, which led to the purchase of two HPE Performance Optimized Data Centers, or PODs—modular data centers that resemble huge cargo containers. Nicknamed EcoPOD, the PODs offered energy-efficiencies, such as an adaptive cooling system, to help keep operational costs down. Using the PODs, the data center was assembled on site in three months, provided space for 44 racks that accommodated up to 24,000 hard drives, and featured security, fire suppression, and monitoring systems. According to Kepner, the cost is about one-twentieth to one-fiftieth of a building-based data center.

Notably, the TX-Green Pioneer lives up to its name. It runs on hydro-, wind-, solar-, and nuclear-generated energy—it is 93% carbon free.[1] "Because of our focus on interactive supercomputing for high-performance data analysis and relatively low carbon footprint, the center is unlike other supercomputing centers," said Albert Reuther, manager of LLSC.

Result

TX-Green Pioneer has been named one of the most powerful supercomputers in New England. Since installation, Lincoln Laboratory has used TX-Green for several important projects:

- Develop, prototype, and transition next-generation signal and image processing algorithms.
- Create a novel method to quickly analyze hyperscale network traffic.[2]
- Enable deployment of massively parallel computations for machine learning on 40,000 cores in seconds.[3]
- Develop a method to interactively launch 16,000 Windows* instances onto 16,00 cores in five minutes.[4]
- Design, generate, and validate extreme-scale power-law graphs in seconds using Kronecker products.[5]

"This is the kind of computation for which signal processing, image processing, machine learning, and physical simulation are a natural fit," commented Reuther.

"The LLSC vision is to enable the brilliant scientists and engineers at Lincoln Laboratory to analyze and process enormous amounts of information with complex algorithms," said Kepner. "Our new system is specifically focused on enabling new research in machine learning, advanced physical devices, and autonomous systems."

"We've been working on connecting different types of databases—SQL, NoSQL, NewSQL—all on the same system, all for the same purpose," explained Dr. Vijay Gadepally, LLSC Senior Scientist. In one of his projects, Gadepally and colleagues are analyzing patient data from ten-years of publicly available intensive care patient information and vital signs. They are developing algorithms that are capable of quickly analyzing thousands of patients records from both historical and real-time sources to help doctors and medical personnel get a better understanding of what interventions were used on patients that may help in guiding diagnosis, prognosis and treatment plans for patient care. "We were able to complete the analysis from two to ten times faster. That's something we couldn't do before," concluded Gadepally.[6]

SOLUTION SUMMARY

MIT's LLSC new supercomputer was built for machine learning and high performance data analytics. The system comprises over 40 thousand cores of bootable Intel Xeon Phi processors, interconnected by the Intel OPA fabric. Part of the MIT SuperCloud, the system has been used since installation for important breakthroughs in network analytics, fast deployment of massively parallel computations in machine learning in seconds, large numbers of non-virtualized Windows environment launches in minutes, and the design, generation, and validation of extreme-scale power-law graphs in seconds.

Where to Get More Information

Learn more about MIT Lincoln Lab, visit https://www.ll.mit.edu.

Learn more about Intel Omni-Path Architecture at https://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-architecture-fabric-overview.html.

Learn more about Intel Xeon Phi processors at https://www.intel.com/content/www/us/en/products/processors/xeon-phi.html.

Learn more about Dell EMC HPC at dellemc.com/hpc.

Solution Ingredients

648 nodes of Intel Xeon Phi processor 7250 (bootable)

Intel Omni-Path Architecture Host Fabric Interface 100 Series

Dell Networking H-Series (H9124-OPF) 24 Slot DCS based on Intel Omni-Path Architecture

[1] http://news.mit.edu/2016/lincoln-laboratory-establishes-supercomputing-center-0511
[2] https://arxiv.org/abs/1808.08353
[3] https://arxiv.org/abs/1807.07814
[4] https://arxiv.org/abs/1808.04345
[5] https://arxiv.org/abs/1803.01281
[6] https://arxiv.org/pdf/1609.07548.pdf

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

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 computer system 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 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.

© Intel Corporation

0819SY/RJMJ/PDF ♻ Please Recycle 341084-001US