

Green computing is super

Massachusetts Institute of Technology

Late this summer, when Lincoln Laboratory scientists and engineers log onto the interactive parallel computing cluster, LLGrid, that answers their high performance computing demands, they will be connecting to a data center situated 90 miles away in Holyoke, Mass., a former textile-manufacturing hub on the Connecticut River. The new Holyoke center offers the technical staff a system six times that of the current LLGrid. The opening of this high-capacity computing facility is the culmination of 10 years of research and planning.

"In 2004, we completed the designs for a data center at Lincoln Lab — LLGrid. But we knew this was only a temporary solution," says Jeremy Kepner, a senior technical staff member in the Computing and Analytics Group, who led the project to develop a computing capability to handle the large datasets and high-fidelity simulations used by researchers across the Laboratory.

David Martinez, currently associate head of the Laboratory's Cyber Security and Information Sciences Division, in 2004 led the Sensor Systems Division in which the Laboratory's advanced computing research was centered. He recalls, "Dr. Kepner had the vision to deploy the most advanced high performance computing to enable Laboratory researchers' rapid prototyping of concepts with national security importance."

Because the size of datasets was exponentially increasing and the complexity of operations performed by computers was growing, Kepner and his team knew the Laboratory would eventually need a facility much bigger than LLGrid's space, a converted lab. "Early on, we did a study on where we should put the next data center," says Kepner. The study considered not only a suitably sized location but also the costs of building and then running a huge computing cluster that consumes megawatts of electrical power around the clock. "Jeremy recognized the need to be close to a power plant to reduce the cost of electricity and to achieve a very environmentally efficient system," says Martinez.

"Holyoke was attractive because the hydroelectric dam provides for less expensive electrical power and land was cheaper," Kepner says. "The cost of electricity in Holyoke is about half the cost of electricity in Lexington." An additional recommendation for this site was that the Holyoke Gas and Electric Department generates about 70 percent of its power from "greener" hydroelectric and solar sources.

The location was right, but there was still the dilemma of a building's pricetag. Kepner knew the construction costs of a brick-and-mortar data center through his participation in a consortium that was seeking computing solutions to support the research of universities and high-tech industries in Massachusetts. He had brought the idea of Holyoke as a site for a computing infrastructure to the consortium that included MIT, Harvard University, the University of Massachusetts, Boston

Green computing is super

Published on Electronic Component News (<http://www.ecnmag.com>)

University, Northeastern University and technology industries. The Massachusetts Green High Performance Computing Center (MGHPCC) in Holyoke, which opened its doors to researchers in late 2012, cost about \$95 million to build.

Kepner proposed a prefabricated alternative to a traditional building. "I heard about Google's containerized computing, essentially putting a supercomputer in a shipping container." Investigation into such a structure led to the decision to purchase two HP PODs, modular data centers that do resemble huge cargo containers. Nicknamed EcoPOD because of energy-efficiencies such as an adaptive cooling system, the data center can be assembled on site in just three months, has 44 racks of space that can accommodate up to 24,000 hard drives, and features security, fire suppression, and monitoring systems. "The EcoPODs' cost is about one-twentieth to one-fiftieth of a building's cost," Kepner says. Moreover, because additional EcoPODs can be easily annexed to create a larger facility when computational needs expand, companies save money and energy by not building and supplying power to a structure larger than their current demand requires.

The resources of the new center are impressive. "We have capacity for 1500 nodes, 0.4 petabytes of memory, and 0.5 petaflops [a petaflop is the ability to perform 1 quadrillion floating point operations per second]," Kepner says. What does all that mean? "With this capability, we can process 1 trillion vertex graphs, do petaflops of simulations, and store petabytes of sensor data. It's like having a million virtual machines."

"MIT Lincoln Laboratory, since its beginnings in 1951, has been in the forefront of computing, starting with the Semi-Automated Ground Environment [SAGE] project. The system deployed at Holyoke continues to be at the vanguard in interactive computing. High-speed connectivity to the system facilitates fast access to massive amounts of data and full access to all of the computational resources available in the system," Martinez says.

After initial checkout and testing, this supercomputing capability will be available to the Lincoln Laboratory research community. The EcoPODs' management and monitoring systems allow for remote operation of the center, although neighboring MGHPCC will provide some support to the facility. Lincoln Laboratory welcomes this hugely enhanced ability to process, analyze, and exploit Big Data, and is proud to be doing it "greenly."

Source URL (retrieved on 05/29/2015 - 2:20pm):

<http://www.ecnmag.com/news/2013/07/green-computing-super>