Case Study

High Performance Computing
4th Gen Intel® Xeon® Processors

Using High Performance Computing to Make Stars in a Jar

Conesus, the Laboratory for Laser Energetics' new supercomputer at the University of Rochester, uses 4th Gen Intel® Xeon® processors to better understand inertial confinement fusion for national security and for use as a future energy source.

Solution Summary

- LLE scientists research inertial confinement fusion.
- Conesus supercomputer supports target design and study of materials under extreme conditions.
- 384 Dell EMC PowerEdge C6220 nodes
- 56-core Intel® Xeon® Platinum 8480+ processors

intel XEON

DELLTechnologies

Executive Summary

The Laboratory for Laser Energetics (LLE) at the University of Rochester is a unique resource for the United States. It is one of the few facilities in the world where scientists are studying and experimenting with inertial confinement fusion (ICF) to harvest energy from the same process that power stars—nuclear fusion. Nuclear fusion promises to provide clean energy from a nearly inexhaustible source for powering the future, but getting there has required decades of research, computing, and experimentation to safely ignite and confine nuclear 'fuel'.

The LLE has a cooperative agreement with the Department of Energy's (DOE) National Nuclear Security Administration (NNSA), making it a national resource in addition to being a university laboratory for the study of high-density-energy physics and ICF.

Inertial confinement fusion is done with sophisticated equipment that includes powerful lasers. But the path to ICF starts with supercomputers to model the materials, the lasers, and the experiments themselves. The LLE has hosted many supercomputers over the decades since its inception. About every five to seven years, the LLE reaches a threshold where computational demand outgrows the capacity of existing resources. That happened recently, triggering the need for a new supercomputer.

They acquired a system built by Dell Technologies with 4th Gen Intel® Xeon® processors, which will allow them not only to simulate experiments with more precision, but also explore using machine learning and artificial intelligence (AI) to garner insight into how to harvest power from nuclear energy.

Conesus is built on 4th Gen Intel Xeon processors with built-in Intel® Accelerator Engines. University of Rochester photo / J. Adam Fenster.

Challenge

The LLE and its sister lab, the National Ignition Facility at Lawrence Livermore National Laboratory (LLNL), are both ICF labs. For the first time, on December 5, 2022, scientists at the NIF, achieved controlled fusion creating a net positive return from their experiment: they generated more energy than the laser energy used in the experiment. That's a significant step.
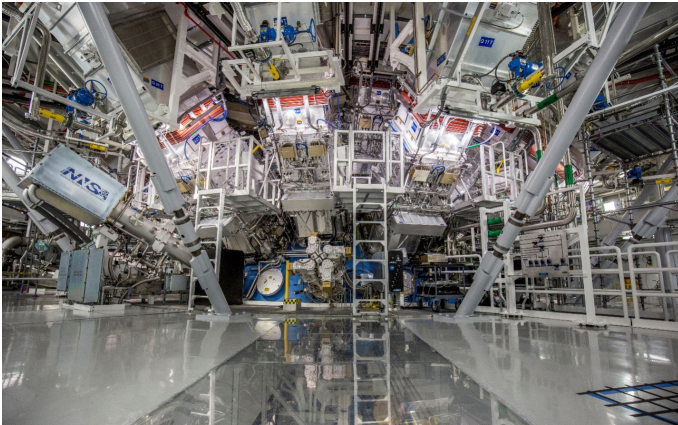
The target chamber of LLNL's National Ignition Facility, where 192 laser beams delivered more than 2 million joules of ultraviolet energy to a tiny fuel pellet to create fusion ignition on Dec. 5, 2022.

LLE and NIF use different methods to compress the nuclear target to pressures that cause a fusion reaction. The NIF uses indirect-drive, while the LLE uses direct-drive. Both use powerful lasers. Direct-drive directly bombards the nuclear target with laser energy to set off the necessary reactions, while indirect-drive uses lasers to bombard an interim medium surrounding the target. Reactions within the interim medium create x-rays that compress and ignite the target.

The LLE has two very powerful lasers called OMEGA and OMEGA EP, which scientists at the lab designed with the help of supercomputers.

"Approximately ten times a day, our lasers are used to make a star in a jar," William Scullin HPC Lead at LLE's computing facility said.

Designing the lasers and simulating the experiments is computationally demanding.

"A lot of our compute cycles go into simulations of experiments," Scullin explained. "We've got 1D, 2D, and 3D modeling capabilities for modeling inertial confinement fusion. We do simulations of materials and plasmas at extremes of temperatures and pressures. High power lasers are not a commercially available component. So, in house we designed a lot of our own optics and laser systems, which can include materials modeling for the development of things like liquid crystals coatings. Additionally, increasingly, there's a lot of statistical work to be done."

For example, according to Scullin with the increase in statistical analysis needed, computational scientists are exploring

how they can use machine learning to see what can be discovered from older data and available data. They needed new computing resources to help make such discoveries possible. Plus, the LLE is getting bigger.

"The Laboratory is growing," Scullin added. "We're physically expanding. We've got new partnerships with faculty on campus and new partnerships in the wider community. We're also coming up on the renewal of our cooperative agreement with the NNSA. We reached the point where the availability of resources meant our users had to wait in long queues to do research. It all justified the purchase of a new, large cluster."

In 2022, with computational resources being strained, they began the acquisition process of a new HPC system.

Solution

"Having a cooperative agreement with NNSA means we benefit from the CTS2 supercomputer design work of the Tri-Labs—Sandia National Laboratory, Los Alamos National Laboratory, and Lawrence Livermore National Laboratory," Scullin said. "CTS2 defines configurations of efficient and cost-effective computing systems for the kinds of nuclear problems the Tri-Labs solve."

Under the CTS2 program, the LLE acquired the Conesus supercomputer, named after one of the finger lakes in the region around Rochester, New York. Conesus is built on 4th Gen Intel Xeon processors with built-in Intel® Accelerator Engines.

"We have always been a CPU shop," Scullin stated. "The majority of our integrated modeling codes use the finite volume method. So, things like memory bandwidth become incredibly important to us very quickly. Likewise, a lot of our production codes are written in Fortran. Intel compilers have always delivered excellent performance for Fortran."

Built by Dell Technologies, Conesus comprises 384 PowerEdge C6220 nodes with Intel Xeon Platinum 8480+ processors containing 56 cores in each of the two sockets. New Dell PowerEdge servers with the Intel Xeon Platinum 8480+ processors can support up to eight DIMMs per CPU and up to 4800 MTS. The CPU's architecture delivers up to 50 percent higher memory bandwidth (4800MTS (1DPC)/4400MTS(2DPC) versus the previous generation processor. The 43,008-core machine recently ranked 311 on the June 2023 Top500 list at 2.59 petaFLOPS. The new supercomputer also ranked 77 on the June 2023 Green500 list.

Built-in Intel® Accelerator Engines offer speedups to many key workloads in HPC. Scullin expects scientists will take advantage of Intel® Advanced Vector Extensions 512 (Intel® AVX-512), which provides significant acceleration for floating point calculations. Security is also a key requirement at LLE.

"We don't do classified work here," Scullin added, "but we're sensitive to security needs around our work. So, another important aspect of the architecture was to make sure we built an environment that would allow for export control. Everything follows NIST SP 800-171, which includes security controls, such as encryption at rest, encryption at boot, and encryption in flight."

Many hardware-enabled Intel® Security technologies, including Intel® Crypto Acceleration, help support these types of NIST requirements. LLE work could also benefit from other accelerator engines, such as Intel® Data Streaming Accelerator for analyzing large data sets.

The LLE is the largest NNSA university-based research program in the nation, making them "the perfect venue for the new system," says Chris Deeney, director of the LLE. "Conesus will provide unique educational and research opportunities for students and faculty in the Rochester region and across the nation."

The system is supported by multiple groups at the DOE and NNSA, as well as state funding from the New York State Energy Research and Development Authority (NYSERDA) and Empire State Development.

"We could not have made this leap forward without the support of NNSA and their willingness to let us partner with the National Laboratories," said Deeney. "Our NY State funding enables LLE to make strategic investments and the new storage system to exploit the power of the new computer will be another great example."

Result

"Everything with Conesus comes back to fusion and the needs of the LLE," Scullin explained. "We've been resource constrained, and users have been waiting in queues to get work done. Having the extra capacity will allow more work to get done sooner. We expect to see significant improvements in bandwidth due to DDR5 and big payoffs in terms of efficiency and throughput. It will also allow users to scale."

Scullin says scientists will now have computing resources to make more runs, collect more data, and perform higher resolution studies, including using machine learning on bigger datasets. One researcher had estimated that one of his projects taking 30 days to a week to run on earlier LLE HPC systems could be completed in a matter of a few days on Conesus.

While many of LLE's scientific codes are written in Fortran, Conesus with Intel technologies and Intel® software will offer performance optimizations in other frameworks.

"Many young scientists have begun using Jupyter notebooks for analysis," Scullin commented. "so, we're looking at things like Python for workflows. The optimization work Intel has done with the Intel® Distribution for Python should directly translate into better abilities to do machine learning types of problems."

Several early science projects are planned for Conesus. The researchers are working from projects they have done on other machines to see what Conesus can do and do at scale. These runs include testing statistical modeling of the OMEGA laser systems' cryogenic implosions; simulating alpha particle stopping and burning plasmas; and studies of liquid crystals that produce a large response with a very high degree of thermal stability.

Conesus will go into production this summer.

The Commodity Technology Systems 2 (CTS-2) program is a National Nuclear Security Administration supercomputer procurement program for the Tri-Labs (Sandia National Laboratory, Los Alamos National Laboratory, and Lawrence Livermore National Laboratory). Since 2007, commodity system procurement programs have allowed the Tri-Labs to obtain cost-effective HPC resources built on a common platform using commodity components for robust capacity computing.

Previous programs were the Tri-Lab Capacity Clusters 1 (TLCC1, 2007-2010), TLCC2 (2011-2015), and CTS-1 (2016-2021). CTS-2 procurements will run from 2022 through 2025.

CTS-2 machines are built with 4th Gen Intel Xeon processors. They provide a common computing environment when used with the Tri-Labs software stack (TOSS) and Tri-Labs Common Environment (TCE). Built by Dell EMC, CTS-2 machines comprise the following technologies:

- Dell C6620 compute nodes
- Dell 760 login/management/gateway servers
- 4th Gen Intel Xeon processors
- Cornelis Networks Omni-Path or Mellanox InfiniBand fabric
- CoolIT direct-to-chip liquid cooling
- GPU Options

Several CTS-2 machines were deployed in 2022, while others continue to arrive in 2023, including Conesus at the LLE.

Solution Summary

The LLE is a resource to study inertial confinement fusion to harvest power from nuclear fusion. High Performance Computing at LLE allows scientists to simulate materials, build and test their lasers, and simulate high-density-energy physics experiments for ICF. With researchers waiting in long queues for computing time on existing resources, the Laboratory acquired a new system, called Conesus, under the NNSA CTS2 program. Conesus is built on 4th Gen Intel Xeon processors with built-in Intel Accelerator engines. While the new 2.59 petaFLOPS system will go into production later this year, it has already ranked in the Top500 and the Green500. Conesus will allow researchers to continue their ICF work with more capacity and explore large data sets with machine learning to help extend their insights into ICF.

Where to Get More Information

Learn more about Laboratory for Laser Energetics.

Explore the capabilities of the 4th Gen Intel Xeon processors with built-in Intel® Accelerator Engines.

Solution Ingredients

- LLE scientists research inertial confinement fusion

- Conesus supercomputer supports simulations and materials design

- 384 Dell EMC PowerEdge C6220 nodes

- 56-core Intel® Xeon® Platinum 8480+ processors

intel
xeon