

Enabling breakthroughs in nuclear fusion

ENEA is at the forefront of clean energy research. To help its scientists make their discoveries faster, ENEA is upgrading its supercomputer architecture to the latest Intel® technologies.

At a glance

- ENEA's central technology division wanted to provide researchers with faster execution times to support research into nuclear fusion.
- ENEA is upgrading to 5th Gen Intel® Xeon® Platinum 8592+ processors, with built-in accelerators. The cluster also has 60 Intel® Data Center GPU Max 1550 accelerators.
- Memory bandwidth constrains some workloads, so the new cluster includes Intel® Xeon® Max processors with high-bandwidth memory (HBM).
- OneAPI enables researchers to run the same application code across the heterogeneous architecture, spanning CPUs, CPUs with HBM, and GPUs.
- ENEA's application benchmark tests showed performance accelerations of up to 8x.^{1,2}
- GPU performance improved by up to 14x and up to 41x.^{1,3}

ENEA is researching nuclear fusion, a potential source of clean energy. To speed up research, ENEA is upgrading its high-performance computing architecture to the latest generation Intel® Xeon® Platinum processors. For workloads constrained by bandwidth, ENEA is adding a cluster of 32 Intel® Xeon® Max processors with high-bandwidth memory. Although the new architecture also includes GPUs, ENEA will be able to use a single code base and the OneAPI Application Programming Interface (API) to run its workloads across all the new clusters.

Challenge

- Accelerate the high-performance computing architecture to support researchers in nuclear fusion better.
- Simplify the coding process for multiple architectures so that researchers do not need to create machine-specific code to use accelerators.

Solution

- ENEA is upgrading to 5th Gen Intel® Xeon® Platinum 8592+ processors, with built-in accelerators including Intel® Advanced Vector Extensions 512 (Intel® AVX-512) and Intel® Advanced Matrix Extensions (Intel® AMX).
- For acceleration, ENEA is deploying 30 additional 5th Gen Intel® Xeon® Platinum 8592+ processors with 60 Intel® Data Center GPU Max 1550 accelerators attached.
- The cluster also includes 32 Intel® Xeon® Max 9480 processors with high-bandwidth memory (HBM) for bandwidth-constrained workloads.
- OneAPI enables researchers to run the same application code across the heterogeneous architecture, spanning CPUs, CPUs with HBM, and GPUs.

Results

- ENEA's application benchmark tests showed performance accelerations of up to 8x with the same number of nodes.^{1,2}
- Comparing the new GPUs with ENEA's existing GPU cluster, performance improved by up to 14x and up to 41x on the new architecture.^{1,3}

Accelerating the quest for nuclear fusion

The impacts of climate change are an urgent global threat. The scientific community is working on several fronts to tackle climate change, including developing clean energy that will be less harmful to the environment. Nuclear fusion has the potential to deliver almost unlimited energy without operational carbon emissions. It's an exciting opportunity, but there is much work to do. So far, nobody has demonstrated a net energy gain from the process.

In Italy, the Divertor Tokamak Test (DTT) facility is now under construction. A tokamak is a machine that captures the energy from nuclear fusion, which happens when hydrogen isotopes are heated to 100 million degrees Celsius—ten times hotter than the Sun's core. Energy is released as atomic nuclei combine to form a heavier, single nucleus.

The DTT is on the outskirts of Rome, at the ENEA Research Center in Frascati. ENEA is Italy's national agency for new technologies, energy, and sustainable economic development. Its researchers working on the DTT and many other projects are supported by the central information and communications technology division, where Giovanni Ponti is head. "We wanted to provide our researchers with better computing facilities and faster execution times so we can support new frontiers in researching nuclear fusion," he said.

Today, ENEA uses its CRESCO 6 cluster as its main computing resource. It was implemented in 2018 with 434 Lenovo compute nodes based on Intel® Xeon® Platinum 8160 processors. CRESCO 6 is complemented by CRESCO 7, a smaller cluster with 144 Lenovo server nodes based on the same processors, implemented in 2023.

In addition to the C++, Fortran, R, and Python languages the researchers used, they also had to learn about programming specifically for GPUs. "We had adopted several GPU technologies that required specific coding," said Ponti. "This was one of the bottlenecks in our research community. We wanted to move the balance of effort to the application and the research, away from hardware-specific coding."

When ENEA secured funding from national, European, and project-specific sources, Ponti and his team set out to find an updated architecture that would give researchers their results faster and be easier to use.

Apart from specifying x86 architecture, ENEA did not lay down hardware requirements. "We provided performance requirements in terms of the petaflops we needed and our power constraints in the data center, and we invited vendors to propose the best solution to us," he says.

Ponti also told vendors which performance indicators he was interested in, a combination of benchmarks from applications used at ENEA and synthetic performance benchmarks.

"The proposal we chose was the most suitable for us overall," said Ponti. "We didn't only take benchmarks and the number of CPUs and GPUs into account, but also the memory bandwidth, architecture, networking, interconnection, and data storage technology."

Solution details

For the CRESCO 8 supercomputing cluster, ENEA chose a solution proposed by Lenovo. Most of the estate comprises 5th Gen Intel® Xeon® Platinum 8592+ processors with 64 cores. These processors have built-in acceleration for workloads including AI, with Intel® Advanced Vector Extensions 512 (Intel® AVX-512), Intel® Data Streaming Accelerator (Intel® DSA), and Intel® Advanced Matrix Extensions (Intel® AMX).

The cluster has 30 additional processors with 60 Intel® Data Center GPU Max 1550 accelerators attached. Many applications across physics, biological and chemical sciences, and AI work out of the box on these accelerators. These include GENE, GROMACS, and CP2K, which ENEA uses and requested benchmarks for.

The built-in and external accelerators will support ENEA as it develops new workloads. "To push the frontiers of scientific applications, we need to adopt accelerators," said Ponti. "It is becoming much more of an everyday requirement for us to use AI, so we want our researchers to have access to accelerators that support their work."

Memory bandwidth constrains some workloads, so the cluster includes Intel® Xeon® Max processors with high-bandwidth memory (HBM). HBM can be used without code changes for workloads of up to 64GB or for caching DDR5 memory. It can also be combined with DDR5 memory for workloads that require large memory capacity. In that case, code changes may be needed to optimize performance.

"Preliminary test results are very interesting and promising for us," said Ponti. "We're excited about testing Intel Xeon Max processors with real workloads."

Ponti will recommend that his researchers adopt OneAPI, a unified programming model. It enables researchers to use a single code base across processors, processors with HBM, and GPUs. In terms of programming languages, OneAPI supports C++, the latest Fortran standards, SYCL, and Python and has enhanced CUDA-to-SYCL code migration capabilities. It includes optimizations for the TensorFlow and PyTorch AI libraries.

As an open and standards-based programming model, OneAPI helps ENEA to future-proof its software, which is important for an organization that has been engaged in research for 25 years and has a lot of legacy code. "We are planning to use oneAPI to better future-proof some legacy code and the new code that takes advantage of the latest processors," said Ponti.

Technical Components of Solution

- **5th Gen Intel® Xeon® Platinum 8592+ processors.** CRESCO 8 uses 1536 of these latest-generation processors, each with 64 processor cores and 320MB L3 cache.
- **5th Gen Intel® Xeon® Platinum 8592+ processors with Intel® Data Center GPU Max 1550 accelerators.** The cluster has 30 additional processors with 60 Intel® GPUs attached for added acceleration.
- **Intel® Xeon® Max 9480 processors.** The cluster includes 32 processors with high-bandwidth memory (HBM) for bandwidth-constrained workloads.

A more sustainable solution

“It is very important for us to support sustainability and energy efficiency at ENEA,” said Ponti. “We want to conduct our research in the most sustainable way we can.”

The new cluster brings Lenovo Neptune, a liquid cooling technology, to ENEA for the first time. Liquid cooling is more efficient than air cooling and helps to reduce energy use while supporting high-performance computing. “We didn’t make it a requirement to have a liquid-cooled machine,” said Ponti, “but with our power and space constraints in the data center, liquid cooling was essential. It helps us to get results for researchers faster because we can push the computational resources harder without exceeding our power or thermal limits.”

Lenovo built the cluster nodes in Hungary, which enabled ENEA to receive the servers sooner.

Intel enables transformation

“Intel has supported Lenovo as much as possible, for which I extend my special thanks,” said Ponti. “Intel has supported us closely, too, by presenting the roadmap and helping us to understand the improvements in the latest technologies. Intel also gives us all the elements needed to address our scientists’ computational needs better.”

Test results

The new cluster will increase ENEA’s computing capacity from 1.01 petaflops on CRESCO 6 to more than 6.5 petaflops on CRESCO 8. That will comprise about 5.4 petaflops in the general-purpose processor partition and 1.18 petaflops in a new accelerated partition with Intel Data Center GPU Max 1550 accelerators.

ENEA carried out testing during its procurement process using several of its application benchmarks (see Table 1), measuring the performance of the latest generation CPU against ENEA’s previous primary architecture. The results showed performance accelerations of up to 8x with the same number of nodes.^{1,2}

CPU Application name (node count)	Time-to-solution CRESCO 6	Time-to-solution CRESCO 8	Improvement observed with the same number of nodes
Quantum Espresso (16)	3186.95"	376.4"	Up to 8x
Gromacs (16)	1659.224 "	264.9"	Up to 6x
Ansys Fluent (4)	379"	87"	Up to 4x
OpenFoam (16)	193.06"	72.2"	Up to 2x
LAMMPS (16)	1160"	255.9"	Up to 4x

Table 1: ENEA’s testing on several benchmarks for applications it uses shows a significant performance increase on the proposed cluster.^{1,2}

In addition, ENEA compared its new cluster architecture with an existing cluster that incorporated GPUs. The existing one was based on IBM POWER9 servers, including NVIDIA V100 GPUs. It was compared with CRESCO 8 nodes configured with two Intel Xeon Platinum processors and four Intel Data Center GPU Max 1550 accelerators. Table 2 shows that performance improved by up to 14x and up to 41x on the new architecture.^{1,3}

Spotlight on ENEA

The ENEA Research Centre in Frascati is one of the major national and international research centers dedicated to studying and developing nuclear fusion, laser sources, and particle accelerators. The Divertor Tokamak Test (DTT) facility is a fusion experiment under construction at the Frascati ENEA Research Center. DTT will play a leading role in fusion research, the main aim being to explore alternative solutions for extracting the heat generated by the fusion process.

<https://www.frascati.enea.it/>

Lessons learned

- Using OneAPI, you can write applications that run across processors and accelerators without writing code specifically for the hardware.
- The Intel® Xeon® Max processor was the first x86-based processor to offer high-bandwidth memory (HBM). It accelerates many HPC workloads without the need for code changes.
- HBM can be used for workloads of up to 64GB or to cache DDR memory without requiring code changes.

Application name (node count)	Time-to-solution Marconi 100 with 4x NVIDIA V100 GPUs	Time-to-solution CRESCO8 Node with 4x Intel® Data Center GPU Max 1550 accelerators	Improvement observed with the same number of nodes
CP2K (8)	26818"	1892.2"	Up to 14x
LAMMPS (8)	5316"	128"	Up to 41x

Table 2: Comparing the previous cluster with the new cluster using Intel® Data Center GPU Max 1550 accelerators.^{1,3}

Ponti said that ENEA understands how benchmarks compare with their real applications, so benchmarks are useful for estimating the likely throughput their applications will achieve.

Learn more

- **Intel® Max series brings breakthrough memory bandwidth and performance to HPC and AI**
- **5th Gen Intel® Xeon® processor**
- **5th Generation Intel® Xeon® Platinum 8592+ processors**
- **Intel® Data Center GPU Max 1550 accelerators**
- **Intel® Xeon® Max 9480 processor**

Find the solution that is right for your organization. Contact your Intel representative or register at [Intel IT Center](#).



Notices & Disclaimers

Performance varies by use, configuration, and other factors. Learn more on the [Performance Index site](#).

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

¹Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

²Tests conducted by F. Iannone (Head of ICT-HPC lab) ENEA in February-March 2023.

CRESCO 6: 434 nodes of 2 socket Intel® Xeon® Platinum 8160 processors (24 cores, 2.10GHz), 192GB RAM, One Intel® Omni-Path Architecture 100Gb/s. Full specs for CRESCO 6: https://www.eneagrid.enea.it/Resources/CRESCO_documents/CRESCO/Sezione6.html#Caratteristiche

CRESCO 8: Testing conducted with pre-production processor codenamed Emerald Rapids (EMR) A0 (64 core, 1.7GHz). This pre-production processor was used to model the performance of the production 5th Gen Intel® Xeon® Platinum 8592+ processor (64 cores, 1.9GHz) deployed in the cluster. 2x EMR A0, 16x32GB DDR5 5600 MHz memory, Mellanox HDR200 networking.

³Tests conducted by F. Iannone (Head of ICT-HPC lab) ENEA in February-March 2023.

Previous architecture: Marconi 100 with IBM POWER9 processors, 4x NVIDIA V100 GPUs

CRESCO 8: 4th Gen Intel® Xeon® Platinum 8480+ processors (56 cores, 20.0GHz) with 4 Intel® Data Center GPU Max 1550 accelerators attached. 512GB of memory at 4800MHz.