

## Accelerating climate change modeling

By adopting processors with high-bandwidth memory, the Euro-Mediterranean Centre on Climate Change (CMCC) aims to get faster results for its climate research, increase accuracy, and serve more users.

### At a glance

- CMCC in Italy was looking for a new supercomputer architecture that would enable it to improve the speed and accuracy of its climate predictions.
- The new cluster is based on Intel® Xeon® Max processors. These are the only x86 CPUs with high-bandwidth memory (HBM).
- Intel ran tests using CMCC's Nucleus for European Modelling of the Ocean (NEMO) workloads.
- The Intel Xeon Max 9480 processor with HBM delivered 3.36x and 3.47x the performance of the 3<sup>rd</sup> Gen Intel® Xeon® Platinum 8360Y processors used in CMCC's newer existing cluster.<sup>1</sup>

CMCC needed a new supercomputing cluster to keep up with the demands of its research into climate change. Intel ran performance tests of CMCC's ocean modeling workloads and demonstrated that the Intel® Xeon® Max 9480 processor with high-bandwidth memory (HBM) would deliver up to 3.47x the performance of CMCC's old Zeus cluster.<sup>1</sup> CMCC is now in the process of deploying a new cluster called Cassandra based on this processor.

### Challenge

- Speed up time to results for complex climate modeling workloads, including Nucleus for European Modelling of the Ocean (NEMO).
- Implement an architecture that improves performance for memory-bound applications.
- Improve modeling resolution up to 7km.
- Cut electricity costs in the data center, which reached €600,000 in 2023 for the data center where the old cluster Zeus was located.

### Solution

- CMCC is implementing a new cluster called Cassandra, based on Intel® Xeon® Max 9480 processors with high-bandwidth memory (HBM).
- The Lenovo servers use direct liquid cooling, which is more energy-efficient than the air cooling used for Zeus.

### Results

- Running single-node tests using CMCC's NEMO workloads, the Intel® Xeon® Max 9480 processor with high-bandwidth memory (HBM) delivered 3.36x and 3.47x the performance of CMCC's newest existing cluster, Juno.<sup>1</sup>
- The Intel Xeon Max 9480 processor with HBM was 2.84x faster than the Juno architecture on a 32-node test of the larger NEMO GLOB16 workload.<sup>1</sup>



## Predicting climate change

“Global warming will increase the likelihood of extreme weather events, droughts, and floods,” said Giovanni Aloisio, strategic advisor at the Euro-Mediterranean Centre on Climate Change (CMCC) and former director of the CMCC Supercomputing Center. “We are studying how to project the temperature increase in 50 to 100 years to inform our governments.”

The simulations produced by CMCC, along with work from other international climate science institutions, are used by the Intergovernmental Panel on Climate Change (IPCC). It publishes reports every two years with advice on mitigation and adaptation measures governments should consider.

CMCC has eleven research divisions, which are organized into three institutes: the Institute for Earth System Prediction (IESP), the Institute for Climate Resilience (ICR), and the European Institute on Economics and the Environment (EIEE).

The researchers at these institutes share the supercomputing center. Today, they use two generations of supercomputers (see Table 1): Zeus (established in 2019) and Juno (2022). The plan was for Zeus to be decommissioned when Juno was available. Still, there were huge challenges in migrating models and their huge volumes of data between the machines and their locations.

As a result, Zeus is still being used by 275 researchers. “Modeling the climate system is a complex problem,” said Aloisio. “There are many interacting processes to model and a great range of timescales and geographic scales to analyze. It requires sophisticated mathematics and computational resources. We also need to manage large volumes of data produced by the simulation.”

The 1-year Mediterranean simulation, for example, produces 288 GB of output.

CMCC’s research work includes using coupled models, which combine models from different domains. For example, the terrestrial vegetation model could be coupled with the atmospheric model, which is coupled with the sea ice and oceanic model.

One of the applications CMCC runs is Nucleus for European Modelling of the Ocean, better known as NEMO. It has three engines for modeling ocean thermodynamics, sea ice, and biogeochemical processes.

“NEMO, like many climate models, is an example of memory-bound code,” said Francesca Mele, who leads the NEMO group at the CMCC Advanced Computing Division. “That means it spends most of its execution time accessing memory to do the computation.”

	Zeus	Juno	Cassandra
<b>Year established</b>	September 2019	October 2022	November 2024
<b>Number of nodes</b>	348 (dual-processor nodes)	170 (dual-processor nodes)	180 (dual-processor nodes)
<b>Processor type</b>	Intel® Xeon® Gold 6154 processors	Intel® Xeon® Platinum 8360Y processors	Intel® Xeon® Max 9480 processors
<b>Cores per processor</b>	18	36	56
<b>Processor speed</b>	3.0 GHz	2.4 GHz	1.9 GHz
<b>Total core count</b>	12,528	12,240	20,160
<b>Memory per node</b>	96 GB	512 GB	1024 GB

Table 1: CMCC’s existing supercomputing architectures of Zeus and Juno, compared with the new Cassandra architecture

## Case Study | Accelerating climate change modeling

The time it takes for models to run varies from minutes to days, depending on the configuration of the simulation. Increasing the memory bandwidth would allow CMCC to get results faster and enable more research to be conducted using the supercomputing resources.

Additionally, CMCC wanted to be able to improve the accuracy of its modeling. “We need to go to a higher resolution,” said Aloisio. “If we use a resolution of 1 degree of the Earth, that’s about 100km. We want to increase the resolution to 1/16 of a degree, which is about 7km. When we increase the resolution, the complexity of the model increases, and we need more computation resources.”

The time to results for several models is highly dependent on the number of cores available for processing, so increasing the core count will shorten the time to results. CMCC was also seeing increasing demand with the number of users growing and the number of core hours they needed for their applications also increasing.

CMCC was looking for a new supercomputer architecture that would enable it to improve the speed and accuracy of its climate predictions.

### Solution details

CMCC is implementing a new cluster, Cassandra (see Table 1), based on Intel® Xeon® Max processors. These are the only x86 CPUs with high-bandwidth memory (HBM). HBM can be used without code changes for workloads of up to 64GB or for caching DDR5 memory. HBM 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.

“HBM will enable us to increase memory use in our models,” said Aloisio. “We will start by using the cache mode without any code changes, and we will look at testing the HBM-only mode, which also works without code changes.”

The new cluster, which will replace Zeus, uses processors with a higher core count of 56, delivering 20,160 cores across the cluster. The memory per node will increase from 96 GB on Zeus to 1024 GB on Cassandra. “Having more memory per node, well balanced with the CPU speed, enable us to increase the resolution of our models,” said Mele.

The Juno cluster has a limited number of GPUs, which are continuously at full utilization. “We’re investigating the use of machine learning for our predictive model,” said Aloisio. “In the future, we expect to more fully exploit the accelerators built into the Intel Xeon processors, such as Intel® Advanced Vector Extensions 512 and Intel® Advanced Matrix Extensions.”

The new cluster includes an investment in 13.6 petabytes of usable storage capacity to accommodate the data generated by the simulations.

Using a new 10 gigabit internet connection, CMCC can now move data quickly from the old data center where Zeus is to the data center where Juno is, and Cassandra will be. The transfer rate has increased from about 5 TB per day to 55 TB per day. Zeus will be decommissioned, and the old data center will be closed when Cassandra is operational. Zeus’s users will run their models on Cassandra. Juno will continue to operate.

### Improving sustainability

The new cluster is based on Lenovo liquid-cooled servers, which are more energy-efficient than air-cooled servers. These are manufactured in Hungary, which reduces transportation emissions and the delivery time.

“Considering the mission of CMCC, sustainability is hugely important to us,” said Osvaldo Marra, director of the CMCC Supercomputing Center. “In 2023, our electricity costs for the Zeus data center were €600,000, and €450,000 for the new data center where Juno is based. Cassandra will be water-cooled, so we will save money in our operations compared to Zeus.”

Aloisio added: “Having the system manufactured in Europe is also important for us. We were impressed when we visited the facility in Budapest.”

#### Technical Components of Solution

- **Intel® Xeon® Max 9480 processor.** The new Cassandra cluster comprises 360 processors with high-bandwidth memory (HBM) for bandwidth-constrained workloads.
- **Lenovo ThinkSystem SD650 V3 Neptune DWC servers.** The Lenovo servers use direct water cooling to enable CMCC to achieve high performance with a lower cooling cost.

## Intel is a strategic ally

Intel worked closely with CMCC on running tests using CMCC’s own models. “Intel was very helpful in benchmarking,” said Aloisio. “The team was able to confirm that NEMO is memory bound and demonstrate the results of using the Intel® Xeon® Max 9480 processor with DDR5 memory. In the past, we’ve only asked for standard benchmarks, but this time, we needed to have benchmarks based on NEMO, a real use case. It gave us a real indication of the performance we could achieve using the processor.”

CMCC is also working with Intel to explore quantum computing for climate modeling, and Aloisio is in the process of setting up a quantum computing group at CMCC.

“We don’t want to see Intel as just a vendor, but more as a partner for developing our performance,” said Aloisio. “It’s a strategic partnership for us.”

## Test results

Intel ran tests using NEMO workloads. These compared the performance of the 3rd Gen Intel® Xeon® Platinum 8360Y processors used in the newer existing cluster, Juno, with the Intel Xeon Max 9480 processor with high-bandwidth memory (HBM), and the 4th Gen Intel® Xeon® Platinum 8480+ processor.

While the 4th Gen Intel Xeon Platinum 8480+ processor delivered 1.61x and 1.77x the performance of the Juno architecture, the Intel Xeon Max 9480 processor with

HBM delivered 3.36x and 3.47x in the single-node tests conducted (see Figure 1). These tests used two generic NEMO workloads.

We also tested the larger NEMO GLOB16 model on a 32-node cluster, comparing the same processors. The Intel Xeon Max 9480 processor with HBM was 2.84x faster than the Juno architecture.<sup>1</sup>

CMCC selected the Intel Xeon Max 9480 processor with HBM for Cassandra.

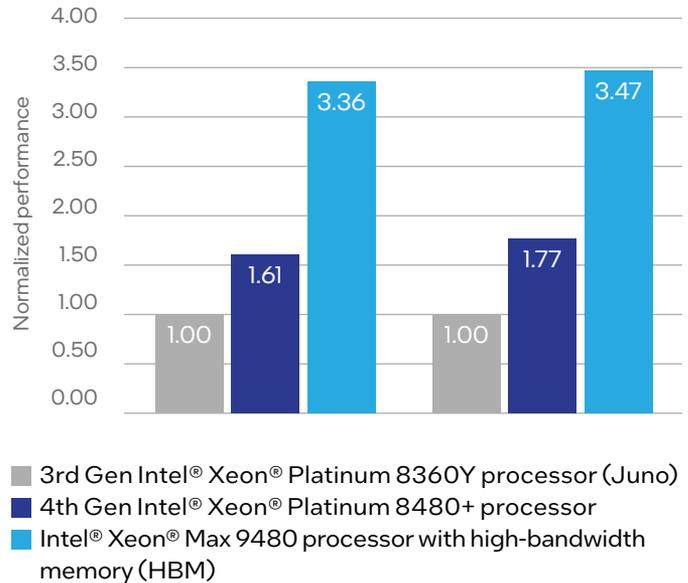


Figure 1: Comparing the single-node performance of three Intel® Xeon® processors for CMCC’s NEMO v4.2 workloads.

	3rd Gen Intel® Xeon® Platinum 8360Y processor (Juno) with DDR4 memory	4th Gen Intel® Xeon® Platinum 8480+ processor with DDR5 memory	Intel® Xeon® Max 9480 processor with high-bandwidth memory (HBM) (cache mode)
Memory footprint	~3 TB	~5 TB	~4.1 TB
Average time per timestep	1.9335 seconds	1.2423 seconds	0.6808 seconds
Normalized performance	1x	1.56x	2.84x

Table 2: Test results for NEMO GLOB16 on a 32-node cluster.<sup>1</sup>

## Lessons learned

- The Intel® Xeon® Max processor is ideal for applications where the memory bandwidth is a bottleneck, thanks to its use of high-bandwidth memory (HBM).
- HBM can be used to cache DDR5 memory without requiring code changes, enabling the acceleration of memory-bound workloads such as NEMO.
- The AI accelerators built into Intel® Xeon® Scalable processors can speed up AI applications. Intel has incorporated optimizations in popular frameworks, including TensorFlow and PyTorch, to take advantage of them.

## Spotlight on CMCC

The Euro-Mediterranean Centre on Climate Change (CMCC) is a non-profit international research center that collaborates with experienced scientists, economists, and technicians to provide full analyses of climate impacts on socio-economic systems. It is Italy's national focal point for the Intergovernmental Panel on Climate Change (IPCC). CMCC's consortium members and institutional partners include six universities, the National Institute of Geophysics and Volcanology, and the American non-profit Resources for the Future.

CMCC's mission is to investigate and model our climate system and its interactions with society to provide reliable, rigorous, and timely scientific results. These, in turn, aim to stimulate sustainable growth, protect the environment, and develop science-driven adaptation and mitigation policies for the changing climate.

<https://www.cmcc.it/>

### Learn more

- [Intel® Max series brings breakthrough memory bandwidth and performance to HPC and AI](#)
- [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.

### <sup>1</sup>Test configurations:

• 3<sup>rd</sup> Gen Intel® Xeon® Platinum 8360Y processor (Juno), 36 cores @ 2.40 GHz. Memory 256GB DDR4 3200 MHz. Interconnect: Intel® OmniPath 100Gbps. Storage: Lustre.

• 4<sup>th</sup> Gen Intel® Xeon® Platinum 8480+ processor, 56 cores @ 2.00 GHz. Memory 512GB DDR5 4800 MHz. Interconnect: Mellanox HDR InfiniBand 200Gbps. Storage: Lustre.

• Intel® Xeon® Max 9480 processor with High-Bandwidth Memory (HBM), 56 cores @ 1.90 GHz. Memory 512GB DDR5 4800 MHz. Interconnect: Mellanox HDR InfiniBand 200Gbps. Storage: Lustre.

OS: Rocky Linux 8.7. Kernel: 4.18.0-372.32.1.el8\_6.crt2.x86\_64. BIOS: SE5C7411.86B.9409.D04.2212261349. OFED stack: mlnx-5.7-1.0.2.0. Compilers and MPI: Intel® oneAPI 2023.0. NEMO v4.2. Input: GLOB16. 5760x3962x98 grid points, 2.5h forecast starting at 2020-0101, dt=200 seconds.

Testing by Intel, February 2023.