Academic Research | Switzerland

Supporting more sustainable research

EPFL

Leading Swiss technical university EPFL empowers researchers with a state-of-the-art, energy-efficient supercomputer powered by Lenovo ThinkSystem and NVIDIA Accelerated Computing technology.

1 Customer background

Who is EPFL?

EPFL (École Polytechnique Fédérale de Lausanne) is a research institute and technical university in Lausanne, Switzerland, specializing in the natural sciences and engineering. It is one of Europe's most cosmopolitan science and technology institutions, welcoming students, professors, and collaborators of more than 120 nationalities. EPFL has both a Swiss and international vocation and focuses on three missions: education, research, and innovation.

EPFL

2 The challenge

EPFL is home to over 500 laboratories and research groups, each working at the forefront of science and technology—from civil engineering and chemistry to materials science, mathematics, finance, and physics. “We have a goal to better understand our world as we aim to improve it,” begins Gilles Fourestey, Director of Operations at SCITAS.

The university offers a wide range of research facilities across its several campuses, providing cutting-edge technical infrastructure and know-how to help EPFL scientists and students conduct their research and make optimal use of resources. These resources include Scientific IT and Application Support (SCITAS), the central hub for scientific computing at EPFL. SCITAS offers advanced computational resources, high-performance computing (HPC) expertise, and artificial intelligence (AI) and machine learning (ML) solutions.

2 The challenge

Antonio J. Russo, HPC System Manager at SCITAS, comments: “More and more research groups are taking advantage of AI and ML to further their work and so need access to powerful GPU resources. Our existing HPC infrastructure had insufficient GPU resources, and we were also experiencing scalability issues for large-scale simulations. What’s more, we had concerns about the energy efficiency of older hardware. It was time for an upgrade.”

“EPFL research labs produce **huge amounts of data every day**. SCITAS users alone generate 314 TB of data every three months, and we need to keep this data for five to ten years. We looked for new HPC infrastructure to help us process and store all this data efficiently. **Sustainability is a top priority for EPFL**, so we targeted a supercomputer that was as energy efficient as it was powerful.”

Antonio J. Russo

HPC System Manager, SCITAS, EPFL

3 The solution

World-class supercomputer

Following a public tender process, EPFL selected Lenovo to design and deploy a GPU-accelerated HPC cluster based on a Lenovo EveryScale “Best Recipe” framework combining the Lenovo ThinkSystem portfolio and OEM components.

The new cluster, Kuma, has two partitions. The main partition is based on 84 Lenovo ThinkSystem SR675 V3 servers, each equipped with two AMD EPYC 9334 CPUs and four NVIDIA® H100 GPUs. The second partition consists of 21 Lenovo ThinkSystem SR675 V3 servers, each equipped with eight NVIDIA L40S GPUs. The servers are interconnected by a 100Gb/s NVIDIA InfiniBand non-blocking network.

Hardware

Lenovo ThinkSystem SR675 V3
AMD EPYC 9334 CPUs
NVIDIA® H100 GPUs
NVIDIA® L40S GPUs
NVIDIA InfiniBand HDR 100Gb/s
IBM Storage Scale System

Software

Red Hat Enterprise Linux
SLURM

Services

Lenovo EveryScale
Lenovo HPC Services – Cluster Design, Configuration, and Installation

3 The solution

Performance meets energy efficiency

Kuma's scratch storage is based on IBM Storage Scale ESS (Elastic Storage System) with 400 TB usable capacity. As part of the Lenovo EveryScale framework, Lenovo provided cluster design, layout, cable design and labeling, and installation services as well as SLURM support for job scheduling.

The air-cooled cluster occupies 12 racks with Lenovo Rear Door Heat Exchangers (RDHX) in EPFL's energy-efficient data center. The data center is connected to EPFL's new heat-pump-powered heating plant, which draws water from Lake Geneva to cool the facility and then supply heating to the rest of the campus.

Working closely with Lenovo, the SCITAS team was able to get Kuma in production in less than 90 days after the equipment was delivered. "That reflects not only the seamless collaboration between EPFL and Lenovo, but also our efficient DevOps processes at SCITAS, which enabled us to quickly implement a broad suite of research software," says Russo.

4 The results

Based on the latest Lenovo and NVIDIA technology, the Kuma cluster delivers 20x more compute power than EPFL's previous supercomputer, enabling researchers to run larger and more complex simulations.

Delivering an Rpeak performance of 21.66 PFLOPS and an Rmax performance of 11.95 PFLOPS, Kuma is ranked 102nd on the TOP500 list of the world's most powerful supercomputers and 23rd on the Green500 list of world's most energy-efficient HPC systems.[1] This exceptional performance on both fronts makes Kuma one of the top research-computing platforms in terms of performance and carbon emissions.

- **21.66 PFLOPS Rpeak; 11.95 PFLOPS Rmax**
- **105 GPU nodes accelerate AI and ML workloads**
- **102nd TOP500 ranking; 23rd Green500 ranking**

[1] Source: Top500 and Green500 November 2024, https://top500.org/system/180344/

4 The results

Powering groundbreaking research

Designed to meet the ever-growing demands of scientific computing workloads, Kuma at EPFL's SCITAS facility provides researchers with unprecedented access to large-scale HPC GPU resources, enabling breakthroughs in fields ranging from large-scale language models (LLMs), deep learning, and neural network training to protein folding, computational fluid dynamics, nuclear fusion, astrophysics, molecular dynamics, biomedical modeling, and large-scale physics simulations.

Accelerated by a total of 105 NVIDIA H100 and L40S GPU nodes, Kuma has improved accessibility to in-demand GPU resources. "More researchers can use state-of-the-art NVIDIA Accelerated Computing to support AI and ML workloads," says Russo. "What's more, we have optimized the balance between GPU-hour cost and compute power, ensuring a more efficient and cost-effective use of our computational infrastructure."

"The supercomputer's powerful processors will accelerate innovation in applications that already rely heavily on computing capabilities, such as healthcare, climate research, materials science, and other critical fields," adds Michele Ceriotti, director of the scientific steering committee for the platform.

“Thanks to Kuma, EPFL will be able to keep **breaking new ground in sustainable AI and scientific computing**, setting the pace on a global scale.”

Gilles Fourestey

Director of Operations, SCITAS, EPFL

Why Lenovo?

Lenovo's proposal met all of EPFL's technical requirements within budget. "Lenovo's experts actively listened to our expectations and demonstrated a strong understanding of our requirements, offering a well-adapted solution that addressed our specific challenges," recalls Russo.

The SCITAS team was also impressed by Lenovo's EveryScale offering. "Lenovo provided us with a dedicated team to efficiently deploy the hardware in our data center," says Russo. "Their support covered all aspects of the installation, including waste management, cabling, and firmware upgrades. Thanks to their structured approach and expertise, the deployment process was carried out smoothly and within the expected timeframe."

How can research institutes meet growing demand for HPC sustainably?

Working with Lenovo, EPFL deployed a powerful yet energy-efficient new NVIDIA GPU-accelerated supercomputer.

Explore Lenovo HPC Solutions