



Putting finance, applied economics, and information management research on the fast track.

How **Shanghai University of Finance and Economics** makes it quicker and easier for researchers to tackle complex problems with a high-performance computing cluster based on Lenovo ThinkSystem SR670 servers, powered by NVIDIA® GPUs.

Lenovo Infrastructure Solutions
for The Data-Centered

Lenovo

1

Background

Shanghai University of Finance and Economics (SUFU) is a finance- and economics-oriented research university located in Shanghai, China. Founded in 1917, SUFE is the oldest finance university, and has long-standing rankings amongst the country's top finance and economics universities.

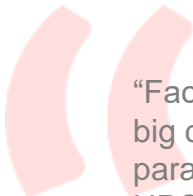
Today, the university is made up of 20 academic schools and departments. Among these is the School of Information Management & Engineering, with three departments and four research centers. Its faculty and students conduct learning and research in a wide array of areas, including artificial intelligence, big data analytics, knowledge engineering, e-commerce, computer science, and quantitative finance.

2

Challenge

Research work at the School of Information Management & Engineering involves increasingly large amounts of simulation and vast volumes of data. Within the School, hundreds of postgraduate and PhD students, as well as faculty members, undertake cross-disciplinary academic research and collaborative research with enterprise and industry.

These research teams require ever-expanding amounts of compute resources to support complex economic simulations, predictive models, and analytics workloads. With the school's existing high-performance computing (HPC) infrastructure straining under the skyrocketing demand, it sought to introduce a new GPU cluster that could easily accommodate current workloads and scale to support future demands.



“Faculty and students teach and conduct research work in various fields, including big data, knowledge engineering, artificial intelligence, quantitative finance, and parallel optimization—all of which require significant GPU resources. Our existing HPC cluster was equipped with 20 NVIDIA P100 GPUs running at full capacity, but it could no longer keep up with our research and development needs.”

Wu Shan

Director of Laboratory Center, School of Information Management and Engineering,
SUFU

Why Lenovo? Stand-out compute performance and density.

When it came to selecting the infrastructure for its new HPC cluster, SUFE decided to make a break with its incumbent hardware provider and form a fresh partnership with Lenovo. The driving factor behind this decision? Lenovo's ability to combine optimal performance with space-saving design.

"Our servers are hosted in the data center at the Education Technology Center at SUFE," explains Wu Shan. "HPC demand has been increasing year on year, from every department. The Education Technology Center had tight restrictions around the amount of space available in the data center. Server density was therefore an important consideration. Lenovo offered high-performance servers with an efficient design. The ThinkSystem SR670 V2 supports up to eight double-wide NVIDIA GPUs per 3U node, while the ThinkSystem SR670 allows up to four double-wide NVIDIA GPUs per 2U node."

The school engaged Lenovo Assessment Services to assess its compute needs and help develop a tailored strategy for HPC. Based on these insights, Lenovo Design Services designed a HPC solution that would meet the school's current and future needs.





“Lenovo packs a lot of computing power into its servers, and that compute density cemented our decision to partner with them.”

Wu Shan

Director of Laboratory Center, School of Information Management and Engineering,
SUFE

Taking HPC higher.

SUFE's School of Information Management and Engineering engaged Lenovo Deployment Services to implement an HPC cluster comprised of 10 Lenovo ThinkSystem SR670 and SR670 V2 servers. The ThinkSystem SR670 servers are equipped with a total of 32 NVIDIA® V100 GPUs, offering 525 TFLOPS of computing power. The ThinkSystem SR670 V2 servers include eight NVIDIA A100 GPUs, with 156 TFLOPS of computing power, accelerated by Tensor Float-32 by up to 2,496 TFLOPS. The NVIDIA GPUs are optimized for AI and HPC workloads, delivering the high performance that the School of Information Management and Engineering needs to tackle increasingly complex research questions.

In addition, the school makes use of IBM Spectrum LSF to simplify administration for its new HPC cluster. The solution supports resource scheduling at the GPU kernel level as well as time and energy consumption monitoring. It also offers advanced workload management capabilities such as CUDA MUS (Multi-Process Service)—helping ensure that users get the most out of their HPC resources.



“With Lenovo and NVIDIA solutions we have established a high-performing, high-density HPC cluster to support valuable research in finance, applied economics, and information management.”

Wu Shan

Director of Laboratory Center, School of Information Management and Engineering, SUFE

3

Results

The Lenovo and NVIDIA HPC cluster has brought a welcome boost to computing power and capacity at the School of Information Management and Engineering. The improvements mean that users can run larger, more complex calculations and complete jobs faster than before. This speeds time-to-answer and helps faculty and students advance important research.


Similarly, with a comprehensive workload management solution now in place, it is now much easier for the IT team to ensure that HPC resources are allocated fairly and efficiently. This helps jobs run at their best performance and in a timely manner, all while easing the administrative burden on IT staff.

As the School of Information Management and Engineering continues to drive valuable research and collaboration between academia, enterprise, and industry, it now has an HPC infrastructure that it can count on well into the future.

Wu Shan concludes: “We are pleased with the performance of the Lenovo and NVIDIA HPC cluster so far, and we know that we can flexibly expand the platform to meet future demands. This will allow us to continue providing strong support for cross-disciplinary research and innovative business development.”



- ✓ Accelerates AI and HPC workloads with powerful GPU-based processing
- ✓ Simplifies workload management and improves resource utilization
- ✓ Shortens time-to-answer for researchers



“The Lenovo and NVIDIA platform opens up new possibilities for our research teams, helping them to tackle more complex and computationally intensive tasks and unlock insights faster.”

Wu Shan

Director of Laboratory Center, School of Information Management and Engineering,
SUFU

What will you do with Lenovo HPC solutions?

The Data-Centered keep vital research moving forward
with Lenovo smarter infrastructure solutions, powered by NVIDIA®.

[Explore Lenovo HPC Solutions](#)



Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo.

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S.
and/or other countries.

Other company, product and service names may be trademarks or service marks of others.

© Lenovo 2022. All rights reserved.