



IDLab (Ghent University and the University of Antwerp) Implements Supermicro's Most Powerful Server, Cuts Down AI Research Experiment Time From Nearly Seven Hours Down To 40 Minutes

Supermicro's new 10U GPU server using NVIDIA HGX-2 GPU baseboards helps leading Belgian research team improve calculation times by 10X which leads to faster AI development



INDUSTRY

Higher Education

CHALLENGE

- Find an extremely high-speed GPU server that could run at least 6 different jobs at the same time, and finish most jobs within 40 minutes.
- Increase performance by a factor of 10 just keep up with current demands

SOLUTION

Supermicro's new GPU server contains two NVIDIA HGX-2 boards, each populated with 8 NVIDIA V100 GPUs, to enable extremely fast training of complex AI models.

[IDLab](#) is an imec research lab at Ghent University and the University of Antwerp. The funding for the GPU server was obtained by Ghent University from FWO (Research Foundation – Flanders). Ghent University is a public research university located in Ghent, Belgium. Established in 1817, it is one of the largest Flemish universities, consisting of 44,000 students and 15,000 staff members. Ghent University consistently rates among the top 100 universities in the world and is well-known in biotechnology, medical and agricultural academic areas.

“Supermicro provided us the specs for their new GPU server and the performance capabilities were impressive. We put a unit in our lab and were extremely pleased with its computing power.

—Brecht Vermeulen, PhD, iLab.t at IDLab, Ghent University

Challenge

To keep a leading position in research and education, IDLab is extending its existing [iLab.t testbed infrastructure](#) to support research in three new focus areas—AI Robotics, AIoT, and data mining. IDLab

BENEFITS

- Cut down experiments from nearly seven hours down to 40 minutes.
- Allows data scientists to conduct faster research by running multiple iterations in a short amount of time and converging on the best solution quickly.

ABOUT SUPERMICRO

Supermicro (Nasdaq: SMCI), the leading innovator in high-performance, high-efficiency server technology is a premier provider of advanced Server Building Block Solutions® for Data Center, Cloud Computing, Enterprise IT, Hadoop/Big Data, HPC and Embedded Systems worldwide. Supermicro is committed to protecting the environment through its “We Keep IT Green®” initiative and provides customers with the most energy-efficient, environmentally-friendly solutions available on the market.

ABOUT SERVER STORAGE SOLUTIONS

Server Storage Solutions deliver custom-made server, storage and networking solutions and services to software development firms and service companies all over the world. We bring our clients’ software and services to life and we deliver unparalleled computing power and storage solutions for the most demanding AI and machine learning programs in the world.

already has several Supermicro GPU servers and storage systems in Ghent with traditional PCIe-based solutions in their datacenter. “Our researchers are constantly pushing the limits of our hardware infrastructure,” said Brecht Vermeulen, Ghent University, PhD, head of the research infrastructure iLab.t at IDLab. “As a leader in AI research, it’s critical that we have a solution to meet the growing demands of our team.”

The researchers in IDLab determined they required an extremely high-speed GPU server that could run at least 6 different jobs at the same time, and finish most jobs within 40 minutes. “We conduct several tasks at the same time and it’s imperative that we get timely results,” said Vermeulen. “Our analysis showed that we needed to increase performance by a factor of 10 just to keep up with our current demands.”

Solution

Based on their requirements, IDLab took interest in evaluating the most powerful GPU server on the market. Designed for next-generation AI applications, [Supermicro’s new GPU server](#) contains two NVIDIA HGX-2 boards, each populated with 8 NVIDIA V100 GPUs, to enable extremely fast training of complex AI models. With the convergence of big data analytics and machine learning, and improved machine learning algorithms, deep learning applications require the processing power of multiple GPUs that must communicate efficiently and effectively to train larger and larger AI models.

“Supermicro provided us the specs for their new GPU server and the performance capabilities were impressive. We put a unit in our lab and were extremely pleased with its computing power. It just made sense for us to move forward with Supermicro’s solution,” said Vermeulen.

Besides the hardware design, another critical reason IDLab chose the Supermicro platform is because Supermicro offers high flexibility to choose the appropriate memory, I/O and CPU, and allows users to have their own software. Additionally, IDLab relied on Supermicro partner [Server Storage Solutions](#) for custom-made services, delivering unparalleled technical advice, and fast and accurate on-site support.

“It was certainly advantageous to have the flexibility to use our own software to manage internal jobs,” said Vermeulen. “We have developed our own software that allows our researchers to easily submit deep learning tasks to the system. The Supermicro platform easily adopts to our needs.”

Benefits

IDLab has been using Supermicro’s new GPU server for nearly one year and has completed many AI tasks, at a substantially quicker rate than before. “It’s remarkable how much faster we can now run our tasks—we’ve cut down experiments from nearly seven hours down to 40 minutes. The speed improvements allow our data scientists to conduct faster research by running multiple iterations in a short amount of time and converging on the best solution quickly,” said Vermeulen.

IDLab has already successfully conducted many research assignments. “We recently competed in the 2019 DARPA Spectrum Collaboration Challenge. With the help of Supermicro and our in-house developed GPU Lab software for job scheduling, we learned how to make optimal use of GPUs for optimising the wireless spectrum collaboration using software-defined radios and AI, and our team was able to reach the finals.”



Server Storage Solutions

For more information, visit <https://www.supermicro.com>



© Copyright 2020 Supermicro