



Accelerating the development of game-changing image recognition solutions.

How **DPControl** used a high-performance Lenovo ThinkSystem SR670 server with cutting-edge NVIDIA® A100 Tensor Core GPUs to train its neural networks for machine vision 28x faster, enabling it to bring innovative edge camera solutions to market ahead of the competition.

Lenovo Infrastructure Solutions
for The Data-Centered

Lenovo

1

Background

Founded in 2002 and headquartered in Salerno, Italy, DPControl (DPC) is a specialist in machine vision, neural networks, artificial intelligence, and smart cameras. Based on its own in-house research and development, DPC offers custom, high-performance hardware and software systems, imaging and multi-modality sensors for a wide range of use cases, including industrial internet of things, factory automation, robotics and drones, security and biometrics, and more.

DPC's offerings hinge on its ability to extract relevant data from large volumes of still images and video quickly and cost-effectively. To enhance its product portfolio, the company is continually researching new image-processing techniques to refine its capabilities and enable innovative, value-added services.

2


Challenge

Moving machine vision workloads to the edge creates valuable opportunities to develop new solutions for a wide range of industries. To capitalize on the new market, DPC aims to enable small, lightweight devices to perform demanding machine vision use cases.

Mario Vigliar, Principal Chief Executive Officer at DPCControl, takes up the story: “Edge cameras have untold potential in industries such as retail—but while there are many new opportunities, there are also new challenges. Unlike traditional image processing where data is sent to a high-performance server for analysis, all the data processing must be completed using the hardware onboard the camera. To make this possible, we need to train neural networks that can deliver rapid and accurate results using minimal compute resources.”

Training these neural networks has increased the complexity of DPC’s deep learning models 100 fold—driving up consumption of CPU and GPU resources. Through its participation in NVIDIA Inception, a program to help startups evolve faster through access to cutting-edge GPU technology, DPC was already using advanced NVIDIA® A100 Tensor Core GPUs to train its neural networks. However, the compute performance of its existing workstation-based infrastructure was becoming a significant bottleneck.

“The faster we can train our neural networks, the sooner we can go to market with edge camera solutions—but our workstation environment was no longer capable of delivering on our requirements,” Vigliar confirms.



“We manage training datasets of around 2 TB, which we were keen to keep on-premises. As a result, we looked for a cost-effective infrastructure platform that could deliver on our demanding performance requirements.”

Mario Vigliar


Principal Chief Executive Officer, DPControl

Why Lenovo? High performance in a cost-effective package.

After evaluating several data center solutions, DPC selected the high-performance Lenovo ThinkSystem SR670 server with low-latency NVIDIA ConnectX 2x100GbE network adapter as its new on-premises infrastructure for training neural networks. The ThinkSystem SR670 server is NVIDIA-certified, which means that it has been validated to provide the optimal performance and scalability for a wide range of accelerated workloads when configured with NVIDIA GPUs and networking.

“By combining our four NVIDIA A100 Tensor Core GPUs into a Lenovo ThinkSystem SR670, we can complete all our neural network training workloads on one server,” explains Vigliar. “The Lenovo solution doesn’t just increase the speed of data processing, it also greatly simplifies our network topology, which makes management far easier.”

With experts from Lenovo, DPC determined the optimal configuration of storage resources to avoid bottlenecks between its CPUs and GPUs. Vigliar comments: “Lenovo helped us find the right balance between cost and performance, and size our system appropriately in terms of storage capacity, power consumption, and physical footprint.”



“NVIDIA A100 Tensor Core GPUs offer incredible performance, but our previous infrastructure didn’t allow us to feed data to the GPUs fast enough to truly take advantage of it. Thanks to the Lenovo ThinkSystem SR670 server, we’ve solved that challenge.”

Mario Vigliar

Principal Chief Executive Officer, DPControl

Fine-tuning the solution.

DPC deployed the new Lenovo ThinkSystem SR670 solution, working with Lenovo to configure and tune the environment.

“Compatibility with our NVIDIA GPUs was an important consideration during the evaluation stage of the project, because at the time the NVIDIA A100 Tensor Core was brand new,” says Vigliar. “Lenovo was one of the first vendors to offer support for the NVIDIA GPU, and helped us with the tuning, maintenance, and monitoring during the implementation. I have expertise in GPUs from my academic background, which means we didn’t need any training—but we greatly appreciated that Lenovo could offer that to us if we wanted it.”



“Lenovo hardware has a great reputation for value, reliability and cost-effectiveness, and our experience with the Lenovo ThinkSystem SR670 solution shows that it’s well-earned. For demanding workloads, power consumption can have a big impact on total cost of ownership—and we’re very pleased with the efficiency of the solution.”

Mario Vigliar

Principal Chief Executive Officer, DPCControl

3

Results

Since moving from its workstation environment to a single Lenovo ThinkSystem SR670 server, DPC has slashed the time taken to complete its neural network training workloads.


“We’ve shrunk our training models from seven days to just six hours since moving to the new Lenovo and NVIDIA platform—28x faster than before,” says Vigliar. “The new solution is a game-changer. Because we can complete our training workloads in hours rather than days, we can carry out more training cycles in the same amount of time—enabling us to branch out into new research areas, evolve the most promising systems, and accelerate innovation.”

Crucially, the new capabilities allow DPC to bring edge camera solutions to market faster, helping it to sharpen its competitive edge.

“The Lenovo and NVIDIA platform has definitely improved our time-to-market—something that would not have been possible without the extra muscle of the Lenovo ThinkSystem SR670 server,” elaborates Vigliar. “We’re now sprinting ahead with an edge camera solution for the retail industry, which represents a big business growth opportunity for the company. We’re so pleased with the new solutions that we’ve already recommended Lenovo to other companies in our group that perform similar deep-learning workloads—and we’re planning to deploy additional Lenovo servers ourselves in the coming year.”



- ✓ **28x faster training for deep learning models**
- ✓ **Cuts time-to-market for edge camera solutions**
- ✓ **Delivers cost-effective performance**



“Lenovo and NVIDIA solutions are already empowering us to carry out innovative work that simply wasn’t possible before. For example, we’re providing the University of Turin with access to our server to experiment with our data remotely—helping us support the next generation of deep-learning research.”

Mario Vigliar

Principal Chief Executive Officer, DPControl

What will you do with Lenovo Analytics & AI solutions?

The Data-Centered accelerate training for deep learning models with Lenovo smarter infrastructure solutions, powered by NVIDIA®.

[Explore Lenovo Analytics & AI 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.