Technology | United States

Delivering AI computing capacity at speed

Corvex

To provide AI product and model builders with rapid access to the latest GPU technology, Corvex deployed an innovative liquid-to-air cooled GPU cluster built on Lenovo ThinkSystem servers with NVIDIA's accelerated computing and cutting-edge networking. Incorporating advanced liquid cooling technology, Corvex commissioned the infrastructure in just two weeks, bypassing the industry's typical lengthy data center upgrade timelines.

1 Customer background

Who is Corvex?

Corvex is on a mission to help AI pioneers move at the speed of innovation by providing high-performance AI computing capacity when they need it. With its AI Factory program, GPU-as-a-Service (GPUaaS), and upcoming Token Factory offerings, Corvex delivers the infrastructure that AI product and model builders require to train, fine-tune, and run inference without waiting months for traditional data center builds.

CORVEX

2 The challenge

AI product and model builders face a critical bottleneck: they need access to the latest GPU technology now, but traditional data center upgrade cycles take 7-12 months or longer. This delay can mean the difference between leading the market and falling behind competitors.

“The customers we work with are building the next generation of AI products and models,” says Seth Demsey, Co-Founder and Co-CEO of Corvex. “They can't wait a year for capacity. When they're ready to scale, they need infrastructure operational in weeks, not months. Speed to deployment is everything.”

Corvex’s customers consistently request immediate access to the newest hardware for their advanced workloads. To continue supporting these AI builders with cutting-edge capacity on their timeline, Corvex set out to rapidly deploy the latest generation of GPUs, including the deployment of the advanced liquid cooling technology necessary to keep the cluster running at optimum performance.

“The customers we work with need the newest hardware for their advanced AI workloads, **and they need it fast**. As GPU thermal densities continue to rise, liquid cooling delivers performance and reliability advantages. The question was: could we deploy this advanced cooling technology quickly? **With Lenovo, we proved you can have both cutting-edge liquid cooling and rapid deployment**—we had this cluster of NVIDIA B200 GPUs operational in two weeks.”

Seth Demsey

Co-Founder and Co-CEO, Corvex

3 The solution

Deploying the latest and greatest GPUs with advanced cooling—at speed

Corvex knew that its long-term technology partner, Lenovo, would be essential for achieving rapid deployment of the latest GPU technology with advanced liquid cooling.

Working closely with Lenovo, Corvex designed and deployed a cluster consisting of liquid-cooled Lenovo ThinkSystem SR780a V3 servers, accelerated by the latest-generation NVIDIA® HGX B200 GPUs. The cluster leverages NVIDIA Quantum-2 QM9700 InfiniBand switches for GPU traffic, and NVIDIA Spectrum SN5600 switches for the storage and SN2201 switches for management networks.

Hardware

Lenovo ThinkSystem SR780a V3 accelerated by NVIDIA® HGX B200 GPUs
Lenovo Neptune® Liquid Cooling
NVIDIA Quantum-2 QM9700 InfiniBand switches
NVIDIA Spectrum™ SN5600 and SN2201 Ethernet switches

Services

Lenovo Installation Services

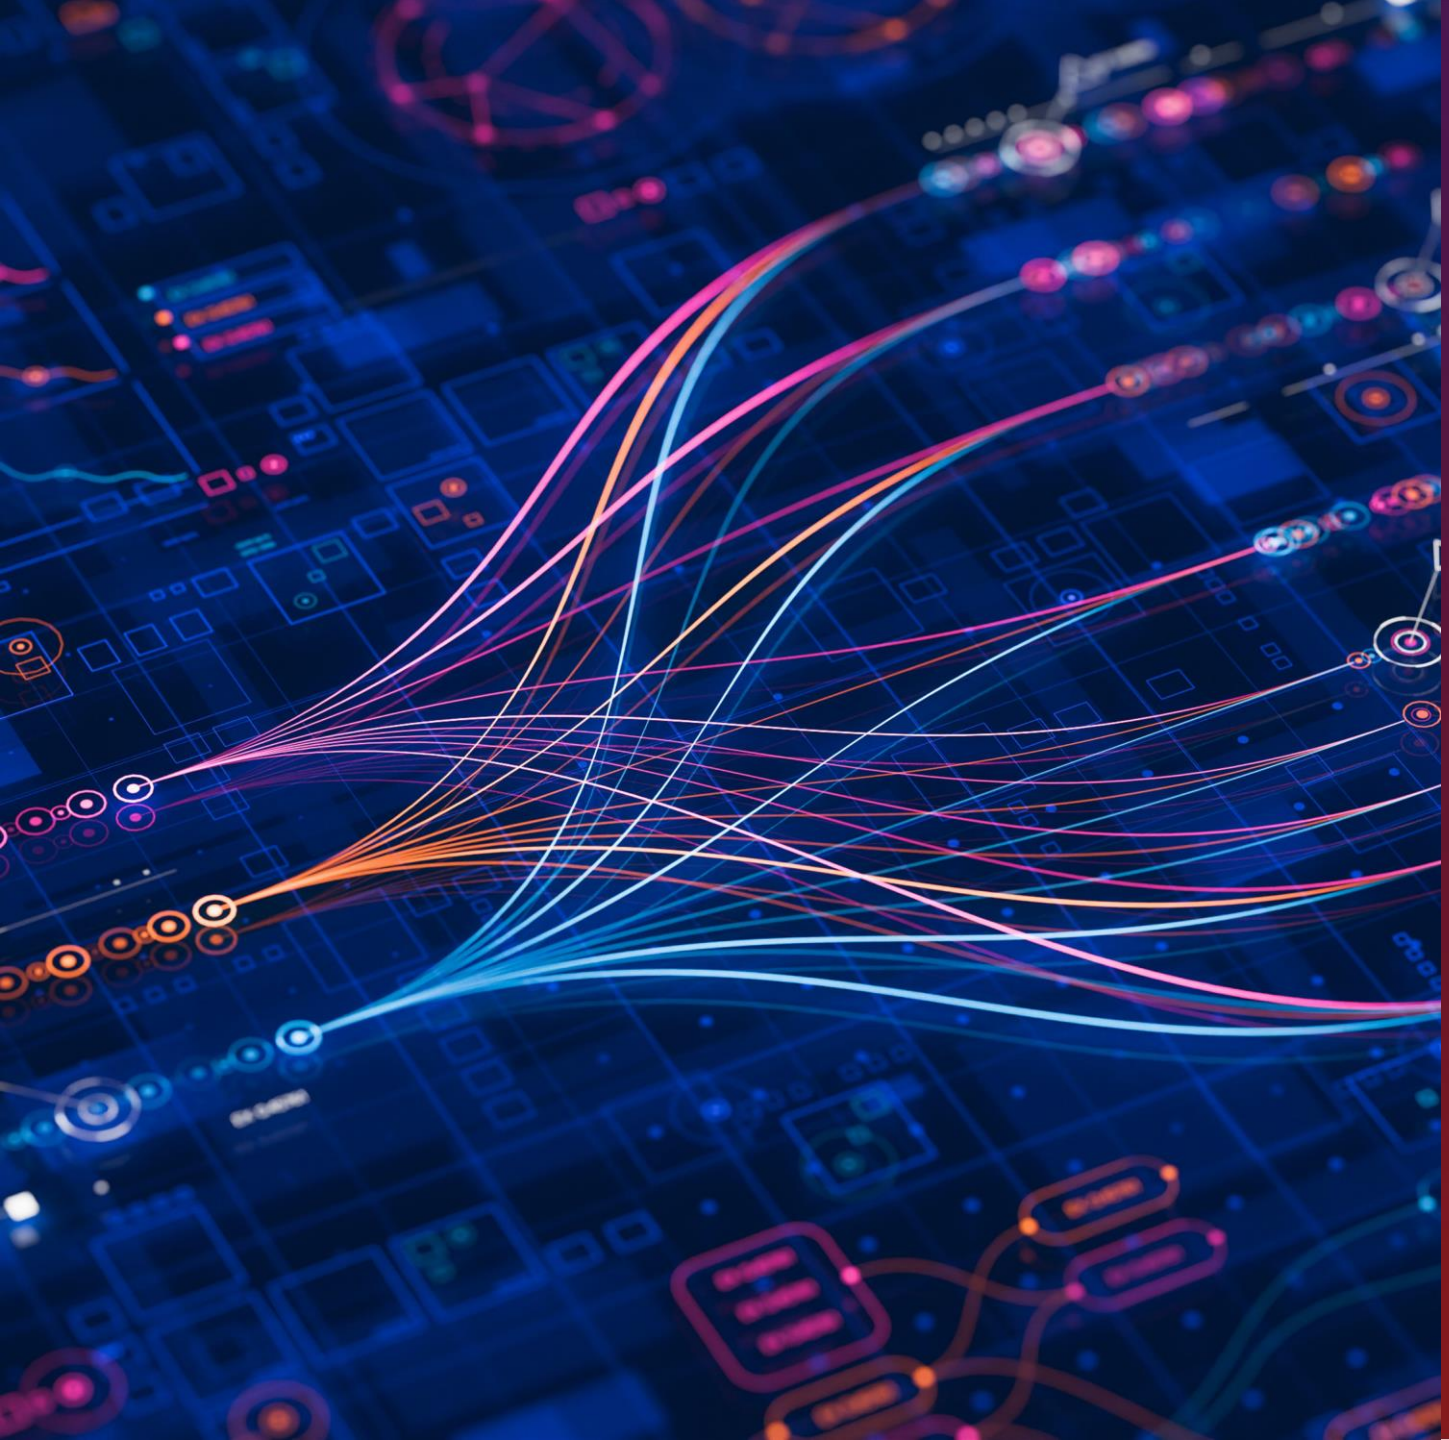

Integrated with Lenovo Neptune® Liquid Cooling, the cluster delivers best-in-class thermal performance to accelerate generative AI and large language model (LLM) workloads.

Corvex commissioned the cluster just two weeks after completion of the hardware load-in.

3 The solution

Retrofitting to unlock performance

Traditional liquid cooling deployments can add many months to data center projects, requiring extensive retrofitting and infrastructure changes. Corvex needed a solution that delivered the performance benefits of liquid cooling without sacrificing deployment speed.

Supported by Lenovo, the company installed Lenovo Neptune Liquid to Air Modules in the servers and integrated them with its existing air conditioning system. Liquid circulates through each rack in the cluster, absorbing heat and evacuating it into the air-cooled data center. This innovative approach enabled rapid installation of advanced liquid cooling technology with the latest, most heat-dense GPU systems without major structural changes to the data center.

“What’s remarkable is that even with the inclusion of advanced liquid cooling technology, we maintained our speed advantage,” confirms Demsey. “Lenovo’s liquid-to-air approach gave us both the thermal performance we needed for the NVIDIA B200 GPUs and the rapid deployment our customers expect.”

“**The deployment process was remarkably fast**, including the liquid cooling. The Lenovo team was with us at every step of the way, from the design to the electrical installation, as well as providing training for our employees. Our customers had access to the latest liquid-cooled NVIDIA HGX B200 GPUs **on their timeline**, not on a traditional data center timeline.”

Seth Demsey

Co-Founder and Co-CEO, Corvex

4 The results

With the new liquid-cooled Lenovo and NVIDIA cluster operational in weeks rather than months, Corvex delivered on its core promise: latest-generation computing capacity with advanced cooling technology, deployed on their timeline.

“AI product and model builders need to move fast, but they also need the best performance. This deployment meant our customers could start running larger, more complex models with superior thermal performance,” says Demsey. “We're delivering that capacity now.”

By deploying Lenovo Neptune Liquid Cooling technology and NVIDIA accelerated computing rapidly, Corvex can continue rolling out the latest, most thermally-demanding GPUs on customer timelines—even as heat densities rise—maintaining speed-to-market advantage for its customers.

- Rapid delivery of cutting-edge AI computing capacity
- Energy-efficient liquid-to-air cooling
- Competitive price-performance ratio

“**Our goal is to be the fastest path to AI Factory capacity with the best technology**. AI builders shouldn't have to choose between cutting-edge infrastructure and rapid deployment. They need both. Our new liquid-cooled Lenovo and NVIDIA cluster proves what’s possible when **innovation meets execution**. Two weeks from hardware arrival to fully commissioned and operational. Our customers are already scaling their most advanced workloads on these liquid-cooled NVIDIA HGX B200 GPUs, with all the thermal performance benefits.”

Seth Demsey

Co-Founder and Co-CEO, Corvex

Why Lenovo and NVIDIA?

As an NVIDIA Cloud Partner, Corvex sought a server solution that aligns with the NVIDIA Cloud Reference Architecture for Enterprise AI, and Lenovo ThinkSystem SR780a V3 ticked all the boxes.

“Lenovo’s reputation preceded them,” says Demsey. “Ultimately we found that Lenovo’s systems architects could deliver exactly what we needed, and Lenovo is always first in the conversation with us as we look to continually expand our data center.”

How can AI builders get liquid-cooled GPUs fast?

Corvex deployed an advanced liquid-cooled Lenovo and NVIDIA GPU cluster, commissioning the latest NVIDIA B200 GPUs in just two weeks after hardware delivery.

Explore Lenovo ThinkSystem Solutions