Akamai Customer Story

Deepsolver Unified Global AI Inference

AI poker solver provider unified global infrastructure while improving latency, scalability, and cost predictability

- Gained access to fast GPUs globally
- Reduced inference latency
- Significantly reduced cloud costs

Delivering real-time AI insights for professional poker players

Deepsolver built an AI-native poker analysis platform that helps professional players and gaming companies generate strategic insights from complex gameplay data in seconds instead of hours. As the Poland-based SaaS company scaled globally, its fragmented cloud environment created operational complexity, unpredictable costs, and regional limitations. Those challenges threatened performance and growth. By consolidating infrastructure onto Akamai Inference Cloud, Deepsolver deployed production-grade GPU inference infrastructure across Europe, North America, and Asia. Doing so helped it maintain low latency, predictable economics, and stability for global customers.

Replacing slow poker analysis with real-time AI inference

Traditional game theory optimal (GTO) poker solvers require enormous compute resources and long processing times. Players often need servers with one terabyte of memory and high-end CPUs — hardware that has become increasingly difficult and expensive to acquire because of soaring memory costs.

"We built an advanced AI poker solver and the infrastructure for professional poker players," explained Tomasz Cybulski, CEO of Deepsolver. "The challenge was giving them analytics and insights instead of simply showing raw data."

Deepsolver's neural network-based platform dramatically accelerated poker analysis workflows for both B2C and B2B customers. Professional players can upload session histories and receive strategic insights within seconds, while gaming operators and developers can integrate Deepsolver's API into their own products.

"With traditional GTO solvers, computations take anywhere from minutes to hours," he continued. "On our platform, it takes seconds because we run inference on GPU servers."

That performance fundamentally changed how players use poker analytics. "In the past, you might schedule calculations overnight and review them the next day," Cybulski said. "Now you can run aggregate reports within minutes. It completely changed how players use these tools."

deepsolver

Location
Lublin, Poland
deepsolver.com

Industry
Software and SaaS

Solution
Cloud Computing

Product
Akamai Inference Cloud

Akamai Cloud

Managing growth across fragmented cloud environments

As Deepsolver expanded globally, its infrastructure became increasingly fragmented across multiple cloud providers, including OVHcloud, Google Cloud Platform (GCP), and Scaleway.

"We tried nearly every cloud provider," Cybulski said. "AWS, Azure, DigitalOcean, you name it."

The company initially selected providers based on specific infrastructure needs. OVHcloud handled GPU workloads, GCP powered CPU-based services, and Scaleway supported storage and database functions. But over time, that fragmented architecture created operational inefficiencies.

"The biggest issue was regional availability," Cybulski explained. "OVH worked well in Europe, but we didn't have access to the United States or Tokyo. We also wanted to standardize on one provider because using three or four clouds simultaneously creates cost complexity."

Deepsolver's growing enterprise business amplified those challenges. The company deploys dedicated infrastructure for each B2B customer in the geographic region closest to its users, making regional cloud coverage a direct driver of product performance.

"When we sell our API to clients, we build dedicated infrastructure in the region where they want it," Cybulski said. "Regional coverage became extremely important."

The company also wanted to eliminate unpredictable hyperscaler pricing models. "Costs on other cloud providers were sometimes hidden," Cybulski said.

Scaling production GPU inference without performance trade-offs

After evaluating multiple cloud providers, Deepsolver selected Akamai Inference Cloud because of its GPU performance, regional coverage, and operational simplicity.

According to Tomasz Cybulski, Akamai offered the combination of predictable economics and global infrastructure the company needed to scale its AI-native platform efficiently. "Akamai offers everything a business needs and is cost effective compared to hyperscalers," Cybulski said.

Before committing fully, Deepsolver benchmarked Akamai GPU infrastructure against competing hyperscaler offerings. "We tested performance against H100s and other NVIDIA GPUs," Cybulski said. "The RTX 6000 PRO™ Blackwell gave us about 10% better performance while costing significantly less."

The company also found the RTX 4000 Ada GPUs particularly effective for cost-efficient solving workloads tied to large-scale API libraries.

Most importantly, the migration introduced no production performance regressions. "At one point, we ran 150 GPU servers on Akamai," Cybulski recalled. "We simply had no issues."

> **Akamai's predictable costs, excellent regional coverage, and GPU performance are perfect for our AI workloads.**
>
> **Tomasz Cybulski**
> **CEO, Deepsolver**

Consolidating global infrastructure onto Akamai Inference Cloud

Rather than attempt a large-scale cutover, Deepsolver deliberately migrated in phases. "We never intended to do a 'big bang' migration," he explained. "We took a phased approach to make sure everything worked as intended — and it did."

After first migrating GPU inference workloads to Akamai, Deepsolver standardized nearly its entire platform stack on Akamai cloud computing services, including:

- **GPU:** Dedicated GPU infrastructure powered by NVIDIA RTX PRO 6000 Blackwell Server Edition® and NVIDIA RTX 4000 Ada GPUs for AI inference workloads.
- **CPU:** Dedicated CPU infrastructure supporting application services and backend workloads
- **LKE: Managed Kubernetes Engine:** Managed Kubernetes service powering four production clusters.
- **Managed Databases:** PostgreSQL instances supporting core application data services.
- **Akamai Object Storage:** Scalable object storage supporting datasets and application assets.
- **Akamai NodeBalancers:** Load balancing services supporting distributed application delivery.

Today, Deepsolver operates production infrastructure across three core Akamai regions:

- Frankfurt for European users
- Chicago for North American users
- Tokyo for Asia-Pacific users

"We basically run everything on Akamai now," Cybulski said. "Plus, we can deploy everything inside one data center for each client deployment, giving us the best latency possible. Akamai's latency is very low compared to other cloud providers," Cybulski explained.

Supporting global growth with predictable infrastructure economics

Standardizing on Akamai Inference Cloud — and accessing highly responsive support — also simplified Deepsolver's operational model. "Akamai saves my time and energy, and that's very important when running a business." Cybulski said. "I don't have to constantly think about resource usage or hidden costs."

The company gained a far more predictable operating model while improving scalability for future AI growth. "Don't just go for the cheapest provider," Cybulski advised other AI-native SaaS companies. "You can end up spending more because of complexity and developer overhead."

Today, Deepsolver continues expanding its inference infrastructure and API footprint globally while building new AI capabilities on top of its neural network engine. "We plan to keep developing our engine to provide better speed, accuracy, and solver extensions," Cybulski said.

The company also plans to continue scaling dedicated customer infrastructure globally on Akamai Inference Cloud. "We are significantly growing on Akamai and look forward to a long-term relationship," Cybulski concluded.

deepsolver

Deepsolver is a tech start-up imagined, designed, and founded by professional poker players, operating as a fully remote organization with an office in the heart of Lublin, Poland. We deliver educational SaaS solutions for poker professionals all around the world. Our next-generation poker solver calculates Nash equilibrium strategy for professional poker players and coaches them with optimal play strategy, aids them with postgame analysis, and facilitates the learning process so they can increase their skill, gain an edge over the opposition, and achieve their full potential.

| Published 06/26