

Nebul Delivers a Powerful AI Cloud That's Sovereign to the EU

Experienced cloud entrepreneur Arnold Juffer, CEO and Founder of Nebul, had predicted the demand for a sovereign, private, and high-performance cloud to power a tidal wave of AI, machine learning and data science. He saw how enterprises were hoarding petabytes of data in his data centers, but without AI and machine learning, they were ill-equipped to capitalize on it. And all that data was growing by 20-50% each year. The result was Nebul. With the goal of building the largest private AI cloud in Europe, Nebul helps companies to unify their data, deploy NVIDIA based Private AI and draw insights from it through powerful and protected infrastructure that's sovereign to the EU.

"Our goal is to be the go-to HPC Cloud Service Provider of Europe. Today, most European enterprises are using US-based hyperscalers for AI, despite government pressure to use providers based in the European Union."
Arnold Juffer, CEO and Founder, Nebul

THE CHALLENGE: LEAPING AHEAD OF THE COMPETITION WITH THE RIGHT TECH STACK

While public AI clouds are great for learning and experimentation, enterprises need private AI clouds to keep corporate data private. This is especially true in Europe, which has some of the world's strictest privacy, sovereignty, and compliance regulations.

Nebul's challenge was finding the right technologies to meet performance and cost goals and ensuring it all works together seamlessly. That's more difficult with AI clouds which require high-performance computing infrastructure that's certified by the GPU vendor. What's more, GPU availability is limited and supply has become a supply chain bottleneck.

Nebul realized they needed to build their AI Cloud quickly to capture the growing number of customers who were adopting AI. "Time to market is critical for our business. A large pharmaceutical company came to us needing immediate access to 32 supercomputers and our cloud wasn't ready. They went with a US-based hyperscaler instead," Juffer shared.

THE SOLUTION: HIGH-PERFORMANCE ELASTIC BLOCK STORAGE FOR AI DATABASES

Nebul started out thinking they didn't need block storage—that they could handle everything with file and object storage—but realized quickly that was not the case. Traditional workloads like virtualization and databases as well as, modern cloud-native workloads like Kubernetes all need elastic block storage with enterprise data services. In particular, they needed fast, scalable block storage for the latency-sensitive vector, real-time, and other NoSQL databases commonly used in AI.



Industry

AI Cloud Service Providers (AlaaS)

Solution

The Lightbits data platform enabled Nebul to build one of Europe's first AI cloud services that's three times more performant at a lower cost than hyperscalers based in the United States.

Business Requirements

- High Performance
- Highly Scalable
- Highly Available
- Cost-Efficient

Business Benefits

- **Performance at Scale:** Performance that rivals local Flash with far greater scalability.
- **Efficiency.** Persistent storage that is disaggregated from compute resources.
- **Availability and Data Protection:** Clustered architecture for availability with built-in 3-way replication.
- **Cost-Efficiency.** Scalable, high performance with efficiency to compete with hyperscalers.

Environment & Workloads

Distributed LLM model training, AI inference and generative AI, vector, real-time, and other AI-oriented databases, and other high-performance computing (HPC) workloads.

Juffer and his team conducted a thorough search for storage solutions, including and beyond the legacy providers. “We considered the usual storage providers, but even the software-defined ones tend to have a ‘closed recipe’ now in terms of infrastructure. We prefer to build with software that’s open, innovative, and especially performant,” Juffer shared.

Nebul tested four storage offerings—two legacy and two more recently introduced solutions—for performance and resilience. Extensive performance benchmarks tested latency, which for their expected workloads was at least as important as throughput. They also tested resource sharing across clients to ensure clients that are doing lots of IOPS are not hogging resources—an important consideration for a cloud service provider. Lightbits emerged as the victor.

Juffer challenged the team to try to break the Lightbits system. As he explained, “If you can break it you know what the limits are. What if half of the power drops or systems start rebooting when you don’t expect it? What if you put the systems at 100% load and then drop a node—how does the system react, how long does it take to rebuild?”

The Nebul team chose Lightbits for its solid performance, comprehensive data services including 3-way replication (a requirement) and QoS, and the company’s commitment to ongoing innovation. Juffer values Lightbits for its product roadmap with new capabilities added at a fast pace, which aligns with Nebul’s own rapid innovation path.

“With Lightbits as part of our platform we can achieve 16 times the performance at half the cost of AI Clouds from American hyperscalers.”

Arnold Juffer, CEO and Founder, Nebul

THE RESULT: A PRIVATE AI CLOUD THAT’S 16X FASTER AND SOVEREIGN TO EUROPE

With Lightbits, Nebul can offer European-based organizations a specialty AI cloud that runs on NVIDIA AI Enterprise with certified tools, framework, and AI apps—and be sure their workloads will get the performance and resilience they need.

When the transactional load increases on an application there’s no noticeable database or application performance drop. “For example, the vector databases used in AI are not always that big, but they’re very latency sensitive. With most storage, if you turn up IOPS, latency goes along with it and database performance drops. Not Lightbits. It behaves well under stress,” Juffer noted. What’s more, with Lightbits as part of their platform, Nebul calculates they can improve LLM model performance by 16 times performance—at half the cost.

Enterprises that adopt next-gen AI like LLMs and Generative AI are 2.6 times more likely to increase revenue by 10% or more. However, they must invest in AI infrastructure to fully reap the benefits. With Nebul, European enterprises can accelerate their AI adoption on a private AI cloud that keeps their data sovereign to the continent.

For more information go to www.lightbitslabs.com.

To get started using Lightbits contact info@lightbitslabs.com