



CUSTOMER SUCCESS STORY

Verda and Supermicro Accelerate the Future of Sovereign AI with Vertically Integrated AI Cloud Infrastructure

Verda and Supermicro's deployment in Europe runs on the latest NVIDIA technologies and is powered by 100% Renewable Energy

SUMMARY

Pioneering Sovereign AI in Europe: Verda's Full-Stack AI Cloud Meets Supermicro's Latest NVIDIA Blackwell Systems

Verda is a leading European AI cloud provider, chosen by pioneering AI teams worldwide for its Full-Stack AI Cloud reinforced by its internal AI Lab. Its data centers are located in the Nordics, allowing the company to operate as a sovereign, neutral enabler of AI technologies, comply with regional regulations, and are the first in Europe to utilize 100% renewable energy sources. Verda is trusted by AI-native innovators, enterprises, research institutions, and public-sector organizations across Europe, the US, and Asia.

By partnering with Supermicro, Verda was among the very first cloud providers to offer NVIDIA HGX™ B200, HGX B300, RTX PRO™ 6000 Blackwell Server Edition, and GB300 NVL72 systems in the EU locations. To date, the availability of NVIDIA Blackwell and Blackwell Ultra systems remains scarce in the EU regions, as deploying the latest generation requires deep technical expertise across both hardware installations and software ecosystems. Verda is one of the very few companies in the EU to have the required expertise to deploy and operate next-gen GPU infrastructure at this scale.

INDUSTRY

Cloud Service Provider

CHALLENGES

Meeting Explosive Demand for Frontier Level AI Models– The Need for NVIDIA Blackwell on a Scale Never Before Seen in Europe

Verda's customers push the boundary of what's possible with AI through frontier R&D, large-scale foundation training, and cost-efficient inference at unprecedented scales. These workloads are GPU-intensive and require the most advanced systems available to date, namely NVIDIA's Blackwell and Blackwell Ultra. Verda needed the latest generation systems with the most powerful GPUs to deliver the compute intensity and performance their customers required. A deep knowledge of the systems, global co-engineering, and data center optimization expertise was required as well.

SOLUTION

Rapid Deployment of Liquid-Cooled Supermicro NVIDIA Blackwell, Blackwell Ultra, and GB300 NVL72 Systems Across Europe

Verda chose Supermicro for delivery speed, state-of-the-art liquid-cooling engineering, and a co-engineering partnership at rack scale, the only way to bring GB300 NVL72 into production on a credible timeline. The combination of the latest generation Supermicro systems, with the most performant CPUs and GPUs, created an unbeatable performance across all systems acquired by Verda:

- Supermicro NVIDIA GB300 NVL72 Rack
 - 36 Grace CPUs
 - 72 NVIDIA B300 GPUS
 - System Memory 17TB DDR5
 - GPU Memory 21TB HBM3e
 - Liquid Cooled
- Supermicro NVIDIA HGX B300 8-GPU system rack
- Supermicro NVIDIA HGX B200 8-GPU system rack
- Supermicro NVIDIA RTX PRO 6000 Blackwell Server Edition system rack

BENEFITS

Accelerating European AI Innovation with Superior Performance, Regulatory Compliance, and Unmatched Sustainability

With Supermicro servers, Verda met the ever-increasing customer demand for NVIDIA Blackwell and Blackwell Ultra servers, deployed in EU locations, and adhered to local regulatory requirements, such as GDPR, AI Act, and sector-specific rules. The servers unlock the performance required by Verda's customers for numerous use cases, such as large-scale mixture-of-experts training and serving, agentic inference on long contexts, and real-time media generation (including video world models), across multiple verticals, such as Financial Services, Media & Entertainment, and Robotics.

Verda is not only committed to its customers, the company is also focused providing a positive impact to the local community. In addition to utilizing 100% renewable energy sources, Verda has partnered with Helsinki utility companies to transfer the heat emitted by the data center cooling system to heat 5,000 to 6,000 homes, with peak potential of up to 15,000 homes.

Customers also benefit from flexible access to these cutting-edge servers with Verda's



“Verda's mission is to provide our customers with the fastest access to the latest hardware available – without compromising on service quality. Being one of the earliest adopters of Blackwell Ultra in 2025, we continue to expand our capacity to meet the market demand through 2026.”

– **Jorge Santos**
COO, Verda



extended service offering:

- **Full-Stack AI Cloud** – A robust application layer supporting AI development across the full model lifecycle, from experimentation to training and inference, at any scale. Paired with security and compliance, the Verda Cloud Platform offers a complete ecosystem with a web console, API, SDK, Terraform, and partner integrations.
- **Instant Clusters** – 16x-128x NVIDIA HGX B200 GPUs with NVIDIA Quantum-2 InfiniBand available via self-service and fully provisioned in 20-30 minutes. Verda is one of the only providers to offer NVIDIA GPU clusters with InfiniBand interconnect via the self-service and pay-as-you-go model.
- **GPU instances** – 1x, 2x, 4x, or 8x virtual machines available with NVIDIA HGX B300 and NVIDIA HGX B200 GPUs, both with NVLink™, and NVIDIA RTX PRO 6000 Blackwell Server Edition.
- **Rack-scale Clusters** – Large-scale GPU clusters built to customers' specifications, including multi-rack configurations with Supermicro NVIDIA GB300 NVL72.
- **Serverless GPUs** – Managed infrastructure for AI inference and batch processing, with autoscaling from multi-GPU deployments to zero when idle.
- **Confidential Computing** – Confidential virtual machines (CVM) that run inside a hardware-enforced Trusted Execution Environment (TEE), encrypting both CPU and GPU memory and cryptographically isolating customer workloads in memory, in transit, and at rest

With Supermicro, Verda can deliver compute when customers need it, at the scale they need it.

“*Supermicro’s Blackwell and Blackwell Ultra systems have been instrumental in helping us scale high-performance AI inference efficiently. Combined with Verda’s flexible cloud capacity and enterprise-grade compliance standards, this infrastructure enables us to deliver reliable, optimized AI services at scale.*”

– **Nikola Borisov**
CEO, DeepInfra

SUPERMICRO

Supermicro is a global leader in high performance, green computing server technology and innovation. We provide our global customers with application-optimized servers and workstations customized with blade, storage, and GPU solutions. Our products offer proven reliability, superior design, and one of the industry's broadest array of product configurations, to fit all computational need.

www.supermicro.com

VERDA

Verda is a full-stack AI cloud platform, purpose-built for GPU workloads from bare metal infrastructure through managed inference endpoints. Running on 100% renewable energy in European data centers and certified to ISO 27001, Verda gives AI teams worldwide the performance, cost efficiency, and infrastructure independence they need to build, train, and deploy at any scale. Verda serves customers across 50+ countries, from AI-native startups to global enterprises.

www.verda.com



©Super Micro Computer, Inc. Specifications subject to change without notice. All other brands and names are property of their respective owners. All logos, brand names, campaign statements and product images contained herein are copyrighted and may not be reprinted and/or reproduced, in whole or in part, without express written permission by Supermicro Corporate Marketing.