

The strategic path to NVIDIA GB200 NVL72 performance

NVIDIA GB200 NVL72 is a monumental leap forward for AI. But accessing its full potential requires more than just hardware; it demands a new class of infrastructure.

Crusoe provides the vertically integrated, purpose-built AI platform designed to unlock the power of NVIDIA GB200 NVL72 – faster, more efficiently, and more sustainably. We eliminate the engineering and operational complexities so you can focus on building the future.



The Crusoe difference: An AI factory designed for builders



Vertically integrated AI factory

We design, build, and operate the entire AI stack, from energy sourcing and data center construction to our AI-native cloud platform. This integration gives us unparalleled control over cost, speed, and reliability.



AI-native performance

Our infrastructure is engineered exclusively for demanding AI workloads. With high-density liquid cooling, ultra-low-latency networking, and 99.98% cluster uptime, we provide a resilient platform for mission-critical AI.



Sustainable by design

We bring AI workloads to abundant, low-cost, and clean energy sources like solar, wind, hydro, and geothermal power. This allows you to scale ambitiously while aligning your compute with the future of the climate.

From complexity to competitive advantage

The challenge: The shift from x86 to Arm

The shift from x86 to the Arm-based NVIDIA Grace™ CPU raises critical questions about code compatibility and migration effort.

The Crusoe advantage:

- **De-risk your transition** with our hybrid deployment capabilities, allowing you to run and test Grace CPU workloads alongside existing x86 instances.
- **Leverage a familiar software environment.** The NVIDIA ecosystem, including CUDA, NVIDIA containers, and the major AI frameworks natively support Arm processors allowing the vast majority of applications to run seamlessly.
- **Gain expert guidance** from our solutions engineers, who specialize in architecting and optimizing complex AI workloads.

The challenge: Liquid cooling reliability

The move to a fully liquid-cooled system can feel like an operational risk, with concerns about reliability and maintenance at scale.

- **Proven at scale** in our Iceland data center, which is powered by 100% renewable geothermal and hydroelectric energy.
- **Unlock superior efficiency** and TCO savings, as our advanced closed-loop systems cut energy use by up to 25% compared to traditional air cooling.
- **Engineered for uptime**, our liquid cooling solutions are a core component of how we deliver industry-leading reliability for the most power-dense AI hardware.

Your infrastructure is your moat

The challenge: Future-proof scalability

Many builders wonder if the GB200 NLV72 performance is excessive for their current needs or too complex to scale incrementally.

The Crusoe advantage:


- **Invest in the future of AI** with an architecture built for the growing complexity of reasoning-based models.
- **Scale seamlessly** from a single rack to a massive cluster on a platform designed for growth without re-architecting.
- **Leverage a unified memory architecture** from day one, which makes our rack-scale systems behave as a single, powerful super-GPU.

The challenge: Maximizing performance

GB200 NVL72 delivers 30x faster real-time large language model (LLM) inference performance, thanks to its architectural innovations.

The Crusoe advantage:


- **Shatter legacy bottlenecks** with the GB200 NVL72 massive unified memory pool. Our platform is architected to maximize this feature for large model training and inference.
- **Maximize parallel training speed** with fifth-generation NVIDIA NVLink™. Our purpose-built infrastructure is designed to leverage its 1.8TB/s of GPU-to-GPU interconnect.
- **Accelerate data pipelines instantly** with the advanced decompression engine. Our tooling and out-of-the-box experience eliminate bottlenecks so you can capitalize on this feature immediately.



Crusoe has been an outstanding partner from the get-go—all of our in-house machine learning models have been trained on Crusoe Cloud. They provide a level of quality of service, responsiveness, and support for early access programs that we couldn't find with any other cloud provider.

Prasanth Veerina, Co-Founder

Pixelcut



Every founder in AI should work with Crusoe. They move as fast as a small startup but provide the robust scaled infrastructure necessary to serve a real-time model to hundreds of thousands of happy users.

Oliver Cameron, CEO

 Odyssey

Ready to unlock the power of GB200 NVL72?

Schedule a trial run with our AI infrastructure experts and see the performance gains firsthand.

Contact us: sales@crusoe.ai

Learn more: crusoe.ai/cloud

Crusoe 