

The AI Era is Here

How Companies Like Kinetica are Powering the AI Era with GPUs

Kinetica is a database designed for real-time data and Generative AI, capable of executing complex analytics on massive datasets with unparalleled speed. Its powerful engine provides accelerated functions for spatial, time series, graph, OLAP, and vector analytics. Trusted by large organizations such as Citibank, the US Army, and Verizon, Kinetica ensures that these entities can perform rapid, sophisticated analyses to maintain their competitive edge.

Kinetica's primary users are enterprises and government agencies that build high-performance analytic and GenAI applications on real-time data, such as IoT sensor readings, vector embeddings, and stock market feeds. These organizations rely on Kinetica to deliver fast, accurate insights that support informed decision-making in high-stakes environments. By offering a platform that integrates seamlessly with existing systems, Kinetica empowers its users to fully leverage their data, driving innovation and achieving operational excellence.

"Kinetica is the world's only fully-distributed GPU accelerated database for generative AI and real-time applications. Our customers rely on us to execute complex analytics on real-time and batch data." said Nima Negahban, CEO of Kinetica.

"Generative AI applications need a retrieval engine that can execute ad-hoc and often complex analytics on large amounts of real-time data. Kinetica's raw analytic power makes this possible no matter the scale." explained Negahban.

Nima Negahban,
CEO, Kinetica

"Generative AI applications need a retrieval engine that can execute ad-hoc and often complex analytics on large amounts of real-time data. Kinetica's raw analytic power makes this possible no matter the scale."

kinetica

Industry

AI Technology

About Kinetica

Kinetica is a real-time, GPU-accelerated database that delivers lightning-fast analytics and insights. It uses the power of GPUs and SIMD-CPU, for advanced data processing, Generative AI, OLAP, time-series analysis, and more.

kinetica.com

Finding the right cloud platform

When evaluating cloud platforms, Kinetica prioritized finding a solution that provided ease of use, robust support for modern hardware such as the NVIDIA GH200 Grace Hopper™ architecture, and cost-effectiveness. The platform needed to offer flexibility, including comprehensive support for virtual machines (VMs) and managed Kubernetes environments. These criteria were essential to ensure that Kinetica could deliver its accelerated analytics capabilities efficiently and adapt to the diverse needs of its enterprise and government clients.

Enter Vultr

Kinetica first connected with the Vultr team while attending the NVIDIA GTC 2024 conference. After evaluating Vultr's capabilities, it was clear that Kinetica had found a reliable cloud partner.

Kinetica now runs NVIDIA GH200 Grace Hopper™ Superchip nodes on Vultr Cloud GPUs.

"We use Vultr to run benchmarks including TPC-DS and vector similarity search to demonstrate how Kinetica and modern GPU hardware can deliver exceptional performance at reduced cost," said Phil Darringer, VP of Product, Kinetica.

"We are also considering supporting Vultr as a cost-effective option in our fully managed SaaS environment, Kinetica Cloud. Vultr Kubernetes with GPU support would allow us to provide a highly reliable and cost-effective option for our cloud customers," said Darringer.

With Vultr, Kinetica delivered the same performance as Databricks at half the cost

Kinetica has extensive experience using hyperscale cloud providers, but they chose Vultr because of its simple, attractive pricing. And given that other providers have yet to offer GH200s, Kinetica was also drawn to the access to the latest generation of GPU technology.

The results speak for themselves. Kinetica estimates that they are spending about half as much for the nodes on Vultr as they would be with equivalently powered EC2 nodes. They're also able to achieve the same performance as platforms like Databricks – at half the cost.

"Kinetica is a high-end analytical platform that executes complex analytics on modern hardware. Vultr made it possible for us to achieve exceptional performance on cutting-edge GPU technology, all while significantly reducing our costs. The combination of Vultr's competitive pricing and access to the latest hardware allowed us to maintain our performance standards at a market leading price point, a benefit that sets Vultr apart from other providers," said Darringer.

"The combination of Vultr's competitive pricing and access to the latest hardware allowed us to maintain our performance standards at a market leading price point, a benefit that sets Vultr apart from other providers."

Phil Darringer, VP of Product, Kinetica

Kinetica makes AI innovation possible, with a little help from Vultr

Kinetica is proving that there is no shortage of potential use cases across a multitude of industries that can benefit from ground-breaking AI technology. With Vultr's help deploying NVIDIA GH200 Grace Hopper Superchips, Kinetica can process and execute vast amounts of data and complex analytics.

"With the power of Vultr and NVIDIA's latest GPUs, Kinetica delivers a cost-effective solution for today's real-time generative AI applications," said Negahban.

Learn more about how
Vultr powers organizations
of every industry

To get started with your
own Vultr success story,
contact us at vultr.com/sales.

vultr.com/solutions

