

The Kinetica Advantage

Unmatched Performance

- Ingest streaming data—billions of records per minute—and get “up to the moment” analytics
- Realize 100x performance improvement on queries compared to CPU-based in-memory solutions
- Holds 100s of TB of data in-memory for extremely low-latency analytics

Advanced Analytics with In-Database Processing

- User-defined functions (UDFs) enable compute as well as data-processing, within the database
- Machine learning/AI libraries such as TensorFlow, BIDMach, Caffe, and Torch can run in-database alongside, and converged with, BI workloads.

Simplicity

- No typical tuning or indexing required; ask and answer any question in real time
- Connect with common BI tools like Tableau, Kibana and Caravel
- A converged, unified suite; not multiple disparate components

Predictably Scalable

- Easily scale up or out
- Data written to Kinetica is automatically routed to parallel connections across the cluster
- OLAP queries are executed using fully distributed GPU-accelerated processing across the cluster

Easy APIs and Integration

- Open source integration components include Apache NiFi, Spark and Spark Streaming, Storm, Kafka and Hadoop
- Kinetica's APIs are fully supported in REST, Java, Python, C++, Javascript and Node.js
- ODBC and JDBC drivers integrate with industry-standard BI and SQL tools

Kinetica and Supermicro work together to dramatically accelerate your computationally-intensive applications and innovations

Supermicro and Kinetica have partnered to provide a GPU-optimized range of servers that can run Kinetica. Those working in diverse sectors such as finance, logistics, telecommunications, energy, and retail can dramatically accelerate application performance with minimal investment in development with this supercomputing solution, which is comprised of Supermicro's servers with NVIDIA GPUs and Kinetica's GPU-accelerated analytics database.

Supermicro has servers that support Kinetica's GPU database and NVIDIA GPUs across a wide range of high-performance servers. Supermicro's computing platforms achieve higher parallel processing capability with NVIDIA's newest Tesla Volta GPU workhorse, delivering the highest quality with extreme optimization for the most computationally-intensive applications.

Supermicro's servers offer superior energy-efficient performance for compute-intensive data analytics, machine learning, and engineering applications. Supermicro's SuperComputing servers offer several key advantages:

- Up to 960 TFLOPs FP16 performance in a single SuperServer
- Unparalleled system NVIDIA Tesla density (up to 4 Tesla V100 in 1U)
- Highest scalability with NVIDIA GPUDirect RDMA capable servers
- Industry-leading Titanium Level (96%+) power efficiency
- Optimized GPU peer-to-peer with single-root architecture
- Performance up to 1/10th the cost and 1/20th the power and consumption of traditional CPU-only servers
- The most comprehensive V100 SXM portfolio in the market

Supermicro offers a wide range of X10 GPU servers that allow customers to accelerate their applications and innovations to address the most complex, real-world problems:

X11 GPU Server Portfolio

Ratio: GPU:CPU	Tower	Rack	Deep Learning
GPU OPTIMIZED	 7049GP-TRT 4:2 (4U)	 4029GP-TRT 8:2 (4U)	 4029GP-TRT 8:2 (4U)
		 2029GP-TR 6:2 (2U)	 1029GP-TR 3:2 (1U)
		 1029GQ-TRT 4:2 (1U)	 1029GQ-TVRT 4:2 (1U)
		 1019GP-TT/5019GP-TT 2:1 (1U)	 4029GP-TRT2 10:2 (4U)

X11 GPUs for Deep Learning

Supermicro offers best-in-class technology designed for supercompute-level performance for artificial intelligence and deep learning applications. It's an ideal system for running Kinetica with deep learning frameworks like Caffe, Torch, and TensorFlow.

X11 for GPU DEEP LEARNING TRAINING for AI

4029GP-TRT2

Flexibility



- 4U Chassis
- DP SKYLAKE CPU w/ 3UPI
- 24 DDR4 DIMMs
- 24 2.5" HS HDD bays
- 12 x16 PCIe 3.0 slots
- 4 (2+2) 2000W Titanium PWS



- 10 Double-Wide GPUs

1029GQ-TVRT

Scalability



- 1U Chassis
- DP SKYLAKE CPU w/ 3UPI
- 12 DDR4 DIMMs
- 2 HS HDD bays
- 4 x16 PCIe 3.0 slots
- 2 2000W Titanium PWS



- 4 Tesla V100 SXM2

4029GP-TVRT

HyperScale



- 4U Chassis
- DP SKYLAKE CPU w/ 3UPI
- 24 DDR4 DIMMs
- 16 HS HDD bays (w/ NVMe)
- 6 x16 PCIe 3.0 slots
- 4 (2+2) 2000W Titanium PWS



- 8 Tesla V100 SXM2



For more information on Kinetica and GPU-accelerated databases, visit kinetica.com

Kinetica and the Kinetica logo are trademarks of Kinetica and its subsidiaries in the United States and other countries. Other marks and brands may be claimed as the property of others. The product plans, specifications, and descriptions herein are provided for information only and subject to change without notice, and are provided without warranty of any kind, express or implied. Copyright © 2017 Kinetica