Microsoft Azure

kinetica

Partner Brief

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

Complete Native Visualization and Geospatial Capabilities
- Real-time geoprocessing
- A fully GPU-accelerated distributed rendering pipeline
- Kinetica Reveal, an extensible and flexible visualization framework, enables interactive, real-time data exploration
- Simply drag and drop data tables to slice and dice data and create on-the-fly analytics dashboards
- Visualize billions of points in seconds

With Microsoft Azure, you can easily leverage Kinetica's GPU acceleration in the hybrid cloud for advanced analytical processing

As streaming analytics, Machine Learning, Deep Learning, and accelerated analytics led by an explosion of data and data types become more ubiquitous, higher performance is needed to accelerate GPU compute workloads. The Microsoft Azure N-Series with GPU Compute is ideal for anyone in high performance computing who wants to enable powerful visual simulations and engineering scenarios. Microsoft Azure virtual machines provide dynamic rendering for GPU-intensive work and compute scenarios for deep learning.

With Microsoft's N-Series of Azure Virtual Machines with GPU capabilities, you can:

- Access data in mere milliseconds vs. 10s of seconds with normal queries
- Render large volumes location-based IoT on the fly in the cloud
- Power decisions and apps with real-time insights

Microsoft Azure N-Series with GPU Compute

Azure is the only public cloud that provides Cognitive APIs, Bots, Machine Learning, and Blockchain as a Service (BaaS) capabilities for developers and data scientists. With Microsoft Azure N-Series Virtual Machines, you can build intelligent solutions at scale by pairing these capabilities with Kinetica's GPU-accelerated analytics database to run deep learning models, HPC simulations, visualizations, real-time data analytics, and many more GPU-accelerated tasks—all on one platform.

The Azure N-Series makes GPU computing accessible for data scientists and developers alike to solve the world's hardest problems. Azure NC-Series VMs are powered by NVIDIA Tesla K80 GPUs. They provide the fastest computational GPU power available in the public cloud. The Tesla K80, with its 4,992 CUDA cores in a dual-GPU design, can deliver up to 2.91 Teraflops of double-precision and up to 8.93 Teraflops of single-precision performance.

NC Instances: Pricing and Availability

Kinetica is available on NC24 Instances for Microsoft Azure. This solution provides customers and developers access to industry-leading accelerated computing and visualization experiences.

Instance	Cores	RAM	GPU Dies
NC24	24	224 GB	4x K80

Standard Managed Disk Type and Size

Managed disks handle storage for you behind the scenes, and offer simple and scalable VM deployment. Kinetica can leverage the following managed disk type for persistent storage:

Disk Type	S30
Disk Size	1024 GB (1 TB)

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