NIMBIX

kinetica

Solution 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

Nimbix and Kinetica accelerate deep learning and analytics capabilities on the cloud.

Nimbix provides on-demand and scalable compute resources that enable organizations to run large-scale High Performance Computing (HPC) workloads in the cloud. With the Nimbix cloud platform for high performance computing and Kinetica, you can explore the full potential in your data and deliver truly real-time insights. The Nimbix Cloud is an award-winning high performance computing (HPC) platform for enterprises and end users who demand performance and ease of use. It brings an unprecedented level of performance and ease of use in running high-performance applications and workflows in the cloud. Nimbix is designed to run applications on accelerated hardware without the overhead of virtualization, so you benefit from faster access and superior hardware performance.

Enabling the highest performance for researchers, data scientists, and engineers looking to leverage optimized workflows.

Through an easy to use portal or a powerful processing API, the Nimbix Cloud runs compute intensive applications on the industry's first true supercomputing cloud. Unlike other public clouds that rely on virtualization for machine provisioning, the Nimbix Cloud runs your applications on bare-metal machines, which translates to better performance at the lowest cost, getting you answers faster than any other solution alternative.

JARVICE is what makes the Nimbix Cloud unique. It's the platform that delivers true HPC in the cloud. Run your enterprise HPC applications in a browser or hand JARVICE a job via its API. Nimbix and the JARVICE platform can scale to meet almost any demand for your large data sets.

Kinetica on the Nimbix Cloud harnesses the power of parallel GPUs to deliver real-time analytics, while also providing tremendous ease-of-use via the Nimbix Cloud environment. Virtualizing GPUs in the cloud gives you unprecedented computing capabilities to help you solve some of your biggest data analytics challenges.

The NVIDIA GPU Advantage

NVIDIA GPUs enable Kinetica to perform brute-force queries on large datasets by leveraging the parallel processing nature of GPUs with their thousands of cores per device, versus 18 to 32 cores on a typical CPU. The outcome is remarkable performance increases, and tangible savings on hardware. On internal benchmarks, NVIDIA GPUs help Kinetica to deliver 100x faster analytic performance than other CPU-based in-memory databases.

Recommended hardware for use with Kinetica

POWER8 in the Nimbix Cloud

Nimbix uses IBM Power Systems S822LC (Minsky) for high performance computing powered by the 128-thread POWER8 CPU. The machines are uniquely suited to GPU-accelerated, high performance computing applications such as Kinetica and are fine-tuned for deep learning, artificial intelligence, and cognitive computing.

Capabilities

- 128 threads, 4 x NVIDIA Tesla P100 GPUs, and 128 GB of RAM
- NVIDIA NVLink for multi-GPU utilization
- PushToCompute leverages familiar continuous integration and continuous deployment workflows for deploying applications to POWER8 architecture machines
- PushToCompute build service for compiling POWER8 applications and deploying application images directly to JARVICE
- Turn-key workflows on JARVICE transparently leverage the compute capabilities of IBM POWER without architecture-specific expertise

The hardware is equipped with their standard supercomputing network fabric, FDR Infiniband, and Pascal-based NVIDIA Tesla P100s equipped with NVIDIA NVLink. The IBM and NVIDIA deployment on Nimbix is the first of its kind in the public cloud and is one most advanced computing platforms available on demand.

Kinetica runs on IBM Power 8 with four NVIDIA P100s equipped with NVLink to power Kinetica's real-time GPU data processing. You can get up and running with Kinetica in just a few minutes at **https://power.jarvice.com**.

kinetica

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