amazon web services™

kinetica

Partner Brief

The Kinetica Advantage

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

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
- Visualize billions of points in seconds

Get real-time actionable intelligence on large, complex, and streaming data sets with AWS and Kinetica

Some of the world's most challenging analytics problems can be solved in real time by using Kinetica's GPU-accelerated database on Amazon EC2 P2 instances.

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. Kinetica is certified to run on Amazon EC2 P2 instances, which are powerful, scalable instances that provide GPU-based parallel compute capabilities. P2 instances are designed for general-purpose GPU compute applications. The AWS and Kinetica solution is ideally suited for accelerated visual analytics, machine learning, high performance databases, computational fluid dynamics, computational finance, seismic analysis, molecular modeling, genomics, rendering, and other server-side GPU compute workloads.

With Kinetica and Amazon EC2 P2 instances, you can query and visualize billions of rows of data in mere milliseconds. With up to 16 NVIDIA Tesla K80 GPUs, P2 instances are the most powerful GPU instances available in the cloud. You can easily increase or decrease capacity within minutes, not hours or days. You can commission one, hundreds or even thousands of server instances simultaneously. Your application can automatically scale itself up and down depending on its needs since it's controlled with web service APIs.

P2 instances allow you to build and deploy compute-intensive applications using the CUDA parallel computing platform or the OpenCL framework without up-front capital investments. To offer the best performance for these high-performance computing applications, the largest P2 instance offers 16 NVIDIA GPUs with a combined 192 Gigabytes (GB) of video memory, 40,000 parallel processing cores, 70 teraflops of single precision floating point performance, and over 23 teraflops of double precision floating point performance.

Accelerated Computing Instances – P2

P2 instances are intended for general-purpose GPU compute applications.

Features:
- High Frequency Intel Xeon E5-2686v4 (Broadwell) Processors
- High-performance NVIDIA K80 GPUs, each with 2,496 parallel processing cores and 12GiB of GPU memory
- Supports GPUDirect™ (peer-to-peer GPU communication)
- Provides Enhanced Networking using the Amazon EC2 Elastic Network Adaptor with up to 20Gbps of aggregate network bandwidth within a Placement Group
- EBS-optimized by default at no additional cost

(Continued on back)

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.

P2 Instance Details

	GPUs	vCPUs	RAM (GIB)	Network Bandwidth
p2.xlarge	1	4	61	High
p2.8xlarge	8	32	488	10 Gbps
p2.16xlarge	16	64	732	20 Gbps

kinetica

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