

Need Fast Data but tired of Hadoop's cluster sprawl?

Accelerate time-to-value for your big data initiatives with GigaSpaces InsightEdge, a high performance and enterprise-ready OLTP, OLAP and Streaming analytics solution based on Apache Spark.

➔ Real-time Big Data without Hadoop's Complexity

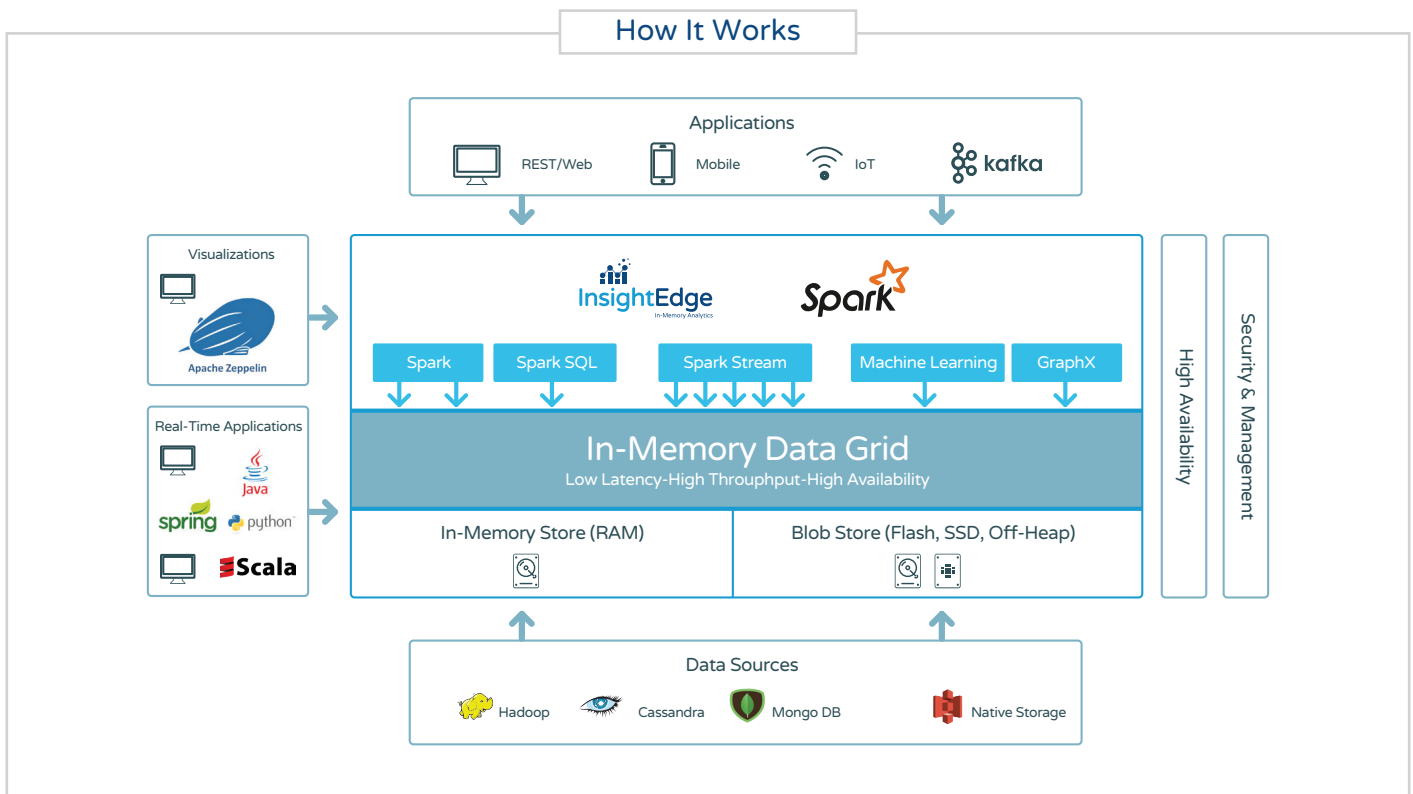
- Process huge volumes of data for real-time analytics at in-memory speeds.
- Eliminate the complexity of collecting, processing, analyzing and visualizing big data without massive cost investment in Hadoop.
- Reduce your Hadoop infrastructure by 3x leveraging an in-memory data virtualization layer.

➔ OLTP, OLAP and Streaming in One Unified Solution

- Eliminate the need to move data from your transactional systems by running analytics directly against the data source.
- In-memory technology allows you to run transactional workloads side-by-side with analytics queries.
- Ingest data directly from relational databases to analyze in a real-time OLAP or streaming analytics fashion.

➔ Optimized Spark Applications

- Native integration between Spark and in-memory data grid to avoid the need to manage yet another data store.
- Improve Spark job performance by minimizing network traffic serialization costs and JVM tuning bottlenecks.
- action lifecycle.





Airline Use Case: Forecasting Live Data Mart

Application


- Flight data yield management: optimization of seat inventory by forecasting against 300 million records

Challenges

- Spark + Hadoop – extensive shuffling between memory and HDFS leads to high latency
- Inability to re-process all data for what-if analysis

Benefits

- Real-time optimization/forecasting against 2.2 million flight objects
- Eliminate Hadoop's latency by using native SSD/Flash as a storage medium for Spark
- Real-time query interface against in-memory storage for result data



Transportation Use Case: Spark Streaming

Application

- Predictive maintenance of equipment, field data ingestion and stream processing

Challenges

- Process streaming data at scale and query from a live data mart
- Providing correlation and statefulness (the IoT streaming problem) for many small, low volume streams

Benefits

- Simplified big data pipeline
- High performance stream processing with HA
- Create triggers to your backend that direct the output to a workflow (M2M) or a human analyst (live dashboards)

Scale your Spark workloads to process terabytes of data 30X faster

