



[www.GigaSpaces.com](http://www.GigaSpaces.com)

420 Lexington Avenue, Suite 2640 New York, NY 10170

Tel: +1 646 421 2830

Email: [Karen.Krivaa@Gigaspaces.com](mailto:Karen.Krivaa@Gigaspaces.com)

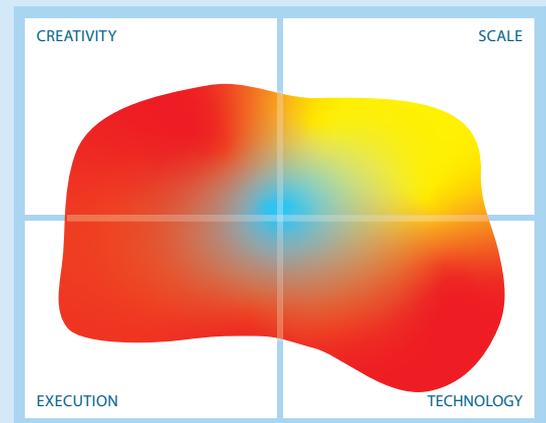
## GigaSpaces InsightEdge

### The company

GigaSpaces was originally founded in Israel in 2000, and currently has headquarters in the United States, with offices in both France and the UK. The company originally targeted financial services but has expanded over the years to support other verticals including transportation, retail, e-Commerce, healthcare, telecom and insurance. Whatever the industry, the majority of users deploy GigaSpaces to accelerate analytical and transactional processing at scale for mission-critical and time sensitive applications.

GigaSpaces has a number of partners of various types. Notably, Avaya, Magic Software, NICE Systems, and Amdocs (amongst others)

“Adoption of GigaSpaces InsightEdge real-time machine learning technology will highly differentiate our services by enabling us to run advanced analytics models on our hot data and instantly predict prices to improve the customer experience.”  
Pricerunner



The image in this Mutable Quadrant is derived from 13 high level metrics, the more the image covers a section the better. Execution metrics relate to the company, Technology to the product, Creativity to both technical and business innovation and Scale covers the potential business and market impact.

GigaSpaces has been working closely with Intel to leverage that company’s new Optane DC Persistent Memory technology which expands capacity while retaining in-memory performance and dramatically lowering hardware footprint.

### What is it?

GigaSpaces initially focused on real-time (microseconds) transaction processing with its XAP product. This is illustrated in **Figure 1** as the “in-memory data grid”. Expanding upon this core technology, is InsightEdge, which allows you to perform real-time analytics thanks to in-memory capabilities, with low latency, high-throughput transaction and stream processing, and the co-location of applications and analytics to act on time-critical data in real-time. InsightEdge is a cloud-native, microservice-based architecture for cloud, on-premises, or hybrid environments; and it supports intelligent, multi-tiered storage across RAM, SSD, storage class memory and persistent memory.

As an in-memory computing data grid, rather than a database, XAP can be implemented on top of, or in front of, more traditional databases, acting as a speed layer for data ingestion supporting millions of operations per second.

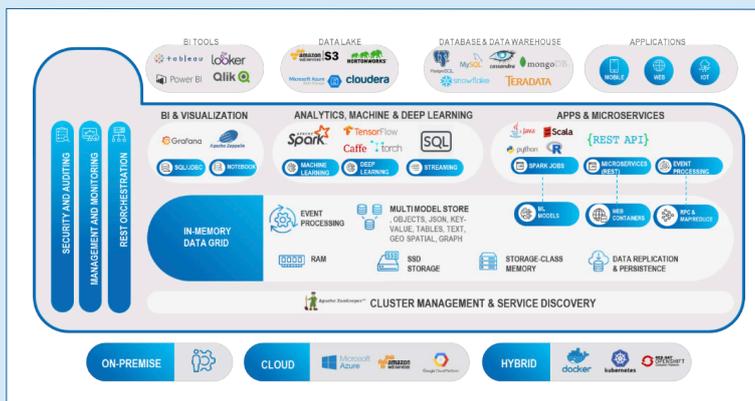


Figure 1 – The “in-memory data grid”

OEM GigaSpaces’ technology. This has resulted in more than 5,000 users of GigaSpaces technologies, either directly or via these partners. Direct users include Barclays, Goldman Sachs, Charles Schwab, Thomson Reuters and many others. HortonWorks/Cloudera, RedHat, Tableau, Informatica and Intel are technology partners.

Architecture	★★★★★	Flexibility	★★★★★½
Concurrent analytics	★★★★★½	In-process analytics	★★★★★½
Criticality	★★★★★	Performance	★★★★★½
Ease of use	★★★★★½	Scale	★★★★★½

### How does it work?

GigaSpaces InsightEdge is available as open source and integrates with big data ecosystems. It supports various data models, including objects, JSON, key-value, tables, text, geo-spatial and graph. The platform is SQL/JDBC compliant and integrates with BI tools such as Tableau, Looker, Power BI and Qlik. Immediate consistency is supported, as is ACID compliance.

For analytics purposes, the architecture of InsightEdge co-locates machine and deep learning frameworks, along with real-time transactional data, and lets you run analytic models on historic

data, reference data and/or data coming in via streams (Kafka and Storm are both supported). To this end, GigaSpaces distinguishes between hot, warm and cold data, as can be seen in **Figure 2**, where hot data is regarded is held in-memory, warm data is persisted in SSD and/or persistent memory and cold data is typically held in Parquet format either on HDFS or in Object

Storage such as Amazon S3 or Azure Blob storage. InsightEdge provides a consistent view across all three of these approaches, supporting analytics via Spark ML and ad hoc query capabilities through Spark SQL.

GigaSpaces also provides an open API to allow machine learning and deep learning, with support for Spark MLlib and Spark GraphX. TensorFlow is also supported as is the loading of Caffe or Torch models. Java, Python, R and Scala are all supported. For ancillary capabilities, Grafana is used to provide operational monitoring, Apache Zeppelin is implemented for query development, there is a RESTful API for management purposes and data operations, and you can leverage popular business intelligence products such as Tableau, Qlik and Looker.

AnalyticsXtreme, which was released recently by GigaSpaces, enables interactive queries and machine learning models to run simultaneously on both real-time mutable streaming data and on historical data that is stored in data lakes based on Hadoop, Amazon S3 or Azure Blob Storage, as well as data warehouses, such as Snowflake, without requiring a separate data load procedure or data duplication.

### Why should you care?

The big advantage of using a data grid is that it is non-disruptive: you can upgrade a legacy (Oracle, Microsoft, IBM and so on) environment to provide both improved performance and scalability at the transactional level, plus adding real-time analytics, without ripping and replacing your existing database systems. Similarly, moving from on-premises to cloud deployments, or changing technology stacks is a relatively simple process. More generally, a data grid can accelerate batch analytics and simplify real-time applications that require analysis of streaming data combined with historical data.

Moreover, you are likely to get better price/performance out of a solution such as InsightEdge compared to expanding your existing database environment, which will often require significant investments in hardware as well as software. In addition, and from a GigaSpaces specific perspective, the company uses a tiered storage architecture, which should provide improved total cost of ownership, along with the sort of mission critical high availability you would expect from an enterprise level offering. On top of this, InsightEdge is richly featured in terms of its support for analytics of all types.

### The Bottom Line

There are multiple ways to implement hybrid environments that support both transactions and analytics. Some of these are more suitable for greenfield deployments, some of them are replacements for traditional environments, and some of them augment those existing systems. GigaSpaces is a leading provider within the last camp, though it is also suitable for greenfield deployments, especially when projects start small.

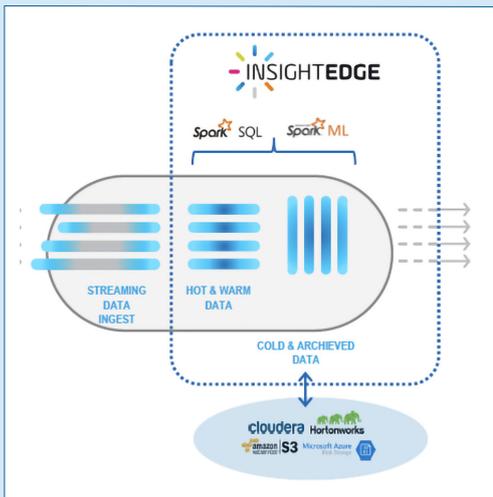


Figure 2 – Differentiating between hot, warm and cold data

[FOR FURTHER INFORMATION AND RESEARCH CLICK HERE](#)