



On-Line Trade Execution with XAP In-Memory Computing



March 2015

On- Line Trade Execution

Table of Contents

Introduction.....	03
Main Industry Challenges.....	04
Meeting Trade Execution with In-Memory Computing.....	05
The XAP Advantage.....	07
On-Line Trade Execution Use Cases.....	13
Bottom Line.....	14

Introduction

The financial services industry is undergoing a fundamental transformation from a product-centric to a customer-centric model leveraging cost effective virtualization and low-latency processing technologies. To keep pace with this new business paradigm, large banks around the world are modernizing their IT capabilities by adopting next generation solutions that enable them to achieve cost efficiencies, contain risks, meet regulatory compliance requirements and gain a comprehensive view of their customers.

High frequency trading applications demand higher throughput and lower latency than ever to meet the challenge of processing terabytes of data from various sources, which must be available to multiple parties, making system management increasingly complex. The variety and volume of data is overwhelming traditional classic database solutions and legacy batch-oriented systems that were not designed to meet today's real time analytic requirements. Faster hardware alone will not meet the challenging requirements of real-time data.

This paper will introduce GigaSpaces' XAP In-Memory Computing Platform as the solution for financial services firms to keep pace with the speed of real-time while maintaining transactional integrity running within modern virtualization environments such as private cloud.

Main Industry Challenges

The increasing amount of trading volumes and numbers of markets create bottlenecks for online trading systems. Today's market demands low latency trading for buy-side and sell-side. As a result, trading firms need ultra-fast pricing engines to ensure they find the best venues for trades at any given time, as well as fast-processing order management systems to execute those trades. In a business environment where milliseconds matter, trading applications must be able to carry out transactions with the speed and accuracy customers expect and regulations demand, providing competitive advantage and improving the bottom line.

(Possible sidebar) XAP In-memory data grid (IMDG) combines High Performance, Scalability, Data Protection, Simplicity, Transactional Integrity and Lower Cost to deliver an Open Platform solution that avoids vendor lock in and won't break the bank.

GigaSpaces XAP provides the only complete end-to-end elastic application platform for online trading that meets or exceeds all of your on-line trading requirements - all while leveraging commodity hardware and the latest advances in fast storage (SSD/Flash). With XAP, your team will:

- Accelerate time to market/value
- Lower Trade latency for increased profitability
- Automatically meet peak trading demands (system elasticity)
- Reduce CapEx and OpEx
- Improve business agility and flexibility to speed product innovation
- Easily meet regulatory compliance requirements
- Gain competitive advantage
- Deliver continuous high-availability with ZERO downtime.

To realize new opportunities, financial services firms will have to focus on four key areas:

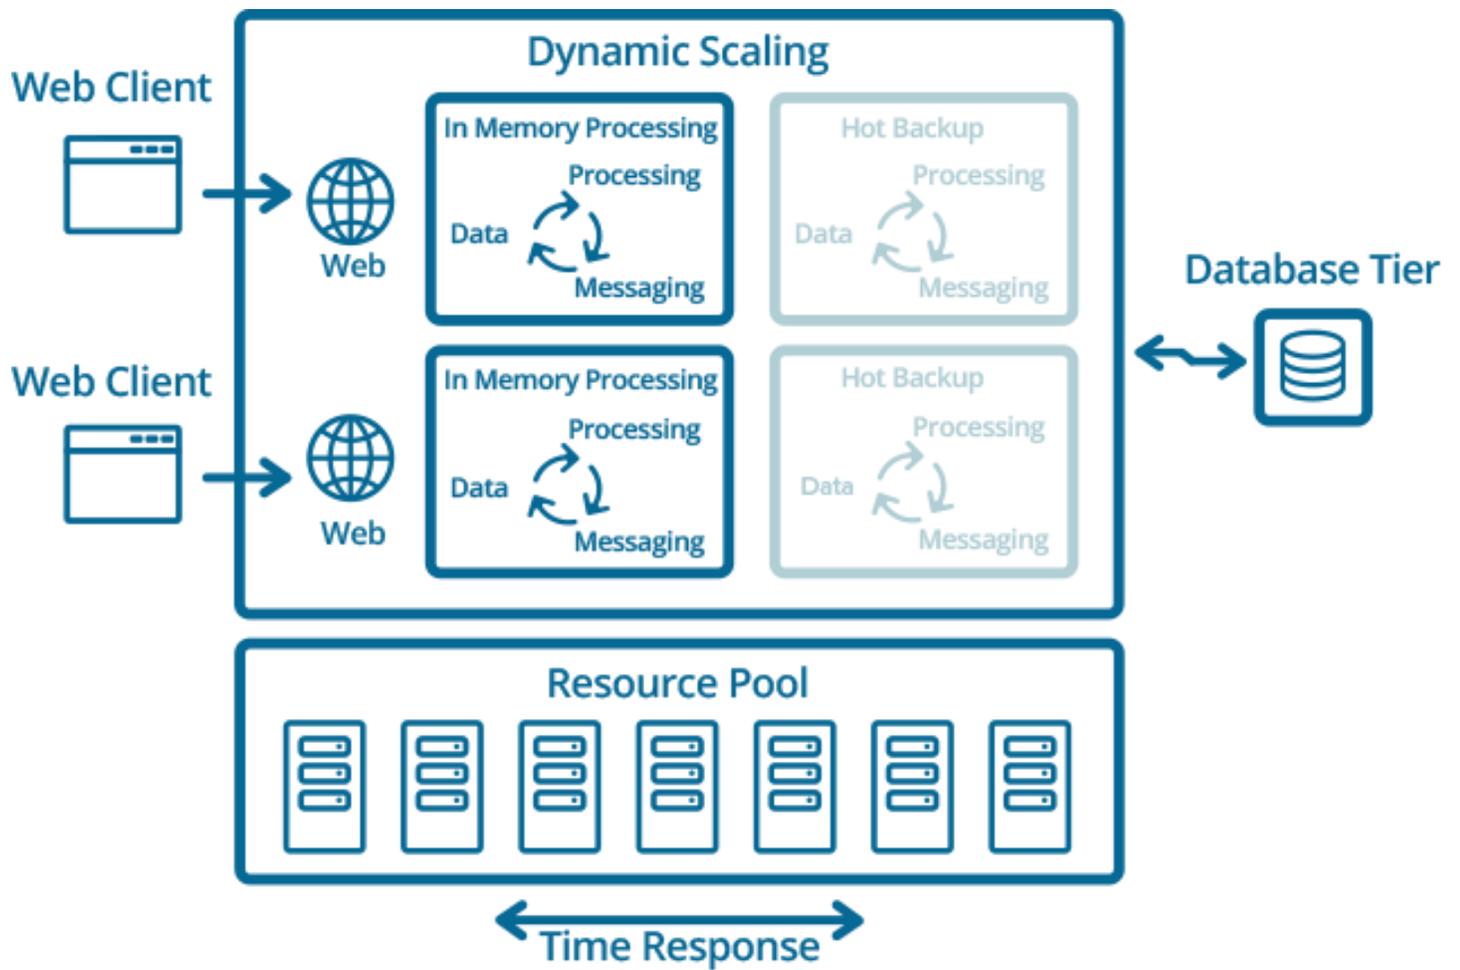
1. **Risk Optimization:** Achieve compliance objectives while mitigating operational risk, fighting crime and optimizing returns
2. **Simplification of operational and technological complexity:** Consolidate disparate sources of customer, product and transactional data. Organizing data as master data, operational data, analytical data accommodating streaming, real-time data and integrating existing data stores

3. **De-duplicate processes, deleverage and pursue componentization of technology:**
Separation of data from applications
4. **Adopt transformational technology solutions that reinforce a simplified, streamlined and agile enterprise:** to balance growth, efficiency and business resiliency in order to support new business innovation.

Meeting Trade Execution Challenges with In-Memory Computing

In-Memory Computing (IMC) is a technology and an architectural style that enables extensive analysis of massive data sets, as well as extreme transaction processing at lightning speeds. These types of architectures rely primarily on RAM as well as CPU power to process the data, which in turn negates the need for rotating disk-stored data access. As posed by Gartner, In-Memory Computing provides transformational opportunities by optimizing the execution of hours-long batch processes into minutes and even seconds. Millions of events can be scanned in a matter of few tens of milliseconds to detect correlations and patterns pointing at emerging opportunities and data insights “as things happen.”

At the root of In-Memory Computing technology is an In-Memory Data Grid (IMDG). This is a cluster of machines that work together to create a resilient shared data fabric for low-latency data access and extreme transaction processing. The XAP IMDG enable an application to run entirely on a single platform with all the tiers (processing, messaging and data) hosted within a single container. The platform provides fast data access by storing all data in-memory, ensuring high availability and scaling mission critical applications automatically and on-demand.



The XAP Advantage



XAP In-Memory Computing is the only platform that enables the development, deployment and management of financial applications running their data, messaging and processing services within a single runtime environment.

XAP delivers true end-to-end scaling with:

- Dynamic provisioning to meet fluctuating loads.
- High performance for latency-sensitive environments.
- Continuous uptime/high availability through automatic failover and self-healing..
- Parallel code execution capability (map-reduce) for improved response time.
- Cost efficiencies through maximum resource utilization, standard development, deployment environments, and advanced management tools.

- Modern grid based architecture allows for infinite scaling leveraging the latest innovations in commodity hardware and flash storage technology.
- Controls and governance to protect the integrity of your data using an enterprise security model.

XAP is built to run real-time, event-driven analytics in memory, co-located with the data, messaging and processing making data immediately available to any process, no matter how complex, with no network hops, vastly speeding up results delivery for better risk management. With built-in multi-tenancy support, the platform enables multiple analytics processes to run on a shared infrastructure while maintaining strict isolation. The result – a more complete solution, providing faster results, at a lower cost, while preserving transactional integrity.

XAP In-Memory Data Grid and Open Source

XAP is a proven, end-to-end in-memory data grid (IMDG) platform ideally suited for analytics requirements. XAP offers a platform approach, which leverages best-of-breed open source tools, including STORM and SPARK for Stream and event processing, data integration tools from Mule and Talend, and Drools for rule execution and real-time decisions.

GigaSpaces is also the creator of Cloudify, an open source automation and orchestration framework integrated with XAP, which can be used with any cloud/virtual environment. While Cloudify is a general-purpose orchestration platform, it can serve many functions, including orchestrating XAP data fabric with other distributed applications such as WebSphere, JBoss, Tomcat and NoSQL databases, such as Mongo or Cassandra. This is accomplished on a bare-metal environment or on a private / public / hybrid cloud environment.

Avoid Vendor Lock-in

XAP runs in the cloud or on premise, easily consolidating disparate sources of customer, product and transactional data from a variety of applications and data formats providing deeper insight in real-time and faster time to market. XAP runs on any OS that supports Java (Windows, UNIX or Linux), providing additional flexibility and portability that allows the decoupling of hardware from the software stack. As a result, firms are able to avoid vertically integrated stack lock-in. In addition, XAP leverages commodity servers and the latest Flash storage technology.

The XAP Price/Performance Edge

Trading applications face an escalating need for performance, while data quantities grow exponentially increasing the need for real-time reporting and analytics.

An independent Total Cost of Ownership (TCO) study conducted by Financial Services Industry experts at Robert Frances Group (RFG) concluded that a **XAP Open Systems Solution is 6x more cost effective than a comparable proprietary Oracle Exadata solution when comparing transaction cost per minute** (Price per core per scan per minute model is derived, in part, from <http://www.tpc.org/tpcc/>)

XAP is Different

For financial transactions, XAP's main differentiator compared to other data grid solutions is XAP's maturity and deep support for reliable, durable and consistent data processing. XAP can ingest a variety of data sources, streaming data, documents, and electronic correspondence (virtually any data source) and process these sources with a real time response.

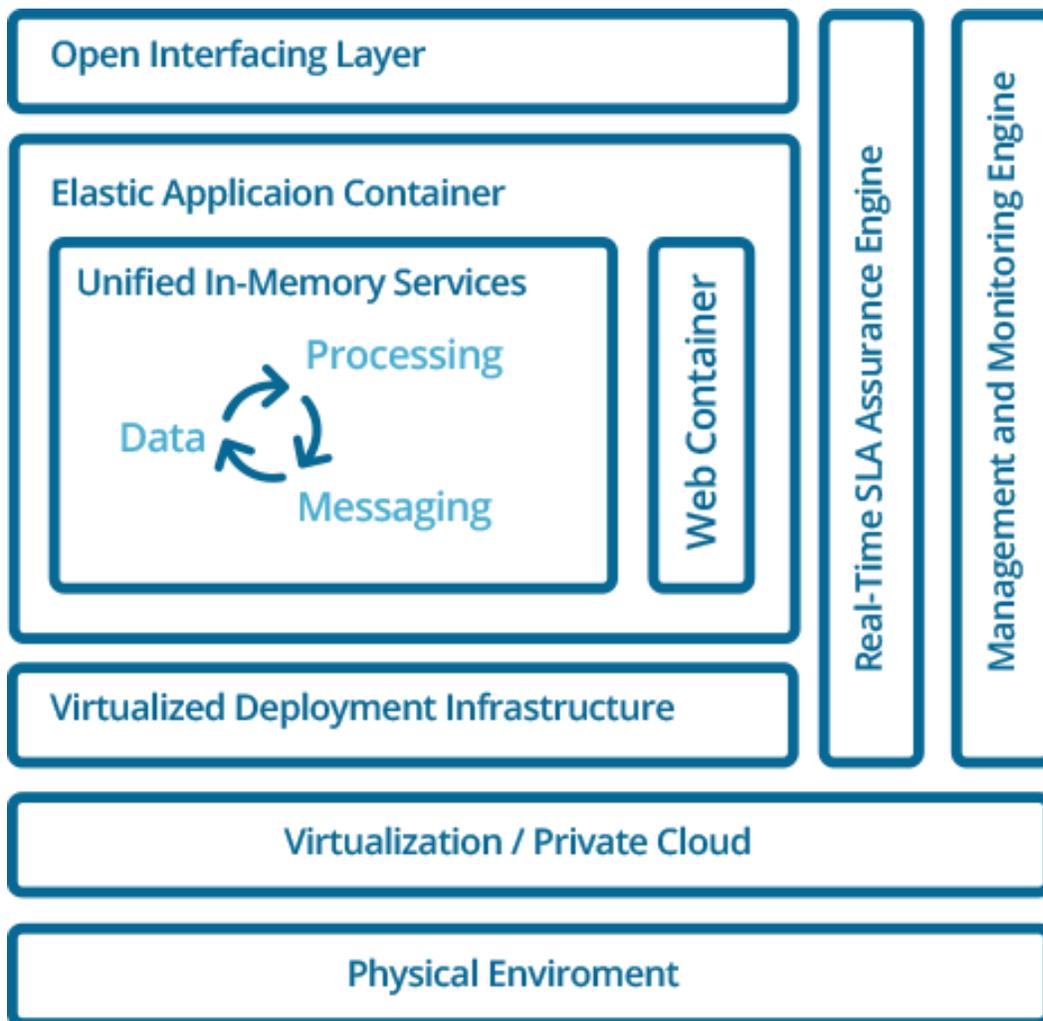
XAP overcomes relational database built-in performance limitations by acting as a single data source or repository, much like a system of record, by replicating and partitioning data while performing "write-behind" to relational or a NoSQL databases preserving entire system consistency.

XAP's Space-Based Architecture (SBA) means each processing unit instance that is acting as the deployment unit with the XAP runtime environment handles only the data preserved within the data grid instance it maintains. In an SBA application model, the data is partitioned in such a way that the services that are triggered as a result will not have to read or write data from other partitions, thus achieving data affinity and in-memory read and write speeds that are far superior to traditional database latencies.

Other major XAP differentiators include:

- **Bidirectional integration** with different data sources, whether to consume data or to push data back out. Integration with real-time processing with Storm, Spark, Hadoop, relational and other databases
- **XAP MemoryXtend** allows XAP IMDG to manage up to 100TB using RAM and local SSD / Central Flash array devices to lower overall grid cost without sacrificing latency requirements.

- XAP leverages Cloudify, an open source automation and orchestration framework developed by GigaSpaces allowing XAP to be deployed seamlessly on any cloud/virtual environment.
- Consistent, flexible transactional behavior and architecture provides eventual consistency or fully ACID, synchronous or asynchronous transactions. Maintain data integrity with 100% transactional data handling through real-time primary with hot-backup replication of in-memory transactions.
- XAP can be used as a Cache or the System of Record interacting transparently with most Big Data, NoSQL solutions and standard SQL databases.
- XAP simultaneously scales the cluster which includes both the data and the business logic in the same grid.
- No network traffic between servers in the grid as XAP divides the data into partitions and the system keeps track of each allocated partition.



XAP Advanced Features

XAP has a number of advanced features that dramatically boost real-time analytic workloads including:

- **Data aggregation** – Similar to relational databases (Sum / Avg / Count / group by, etc.), but much faster execution in a distributed configuration within the data grid scanning millions of items in seconds. Ensures immediate consistency, rich query semantics & data indexing.
- **Stream processing** – Complex data event processing using XAP integration with Storm, Spark, Drools and other frameworks supporting massive data processing in a distributed elastic configuration.
- **Real-time Big Data** – Frontend Hadoop, Cassandra, MongoDB using built-in adaptors and Kafka to propagate real time processing to backend disk based big data systems for further offline analysis. In some cases, XAP can even eliminate the usage of these databases. Optimized for fast data.
- **Replication over WAN** – Scaling across data centers with near real-time replication of massive data streams. Ideal for disaster recovery, cloud bursting and multi-site scenarios.
- **SSD as Data Grid storage medium** – Leveraging both RAM and SSD to manage large amounts of data at a fraction of the cost, reducing hardware footprint.
- **Customer 360** – Ability to consolidate data from various sources, both external and internal and structured or unstructured, into one unified model. XAP supports this via several fundamental features such as a Document / Schema-less model, transparent continuous data load from different databases and the Rest API.
- **Continuous High Availability** – Provided via backups and self-healing features and, sub-second failover capabilities. XAP has no single point of failure.
- **Data Protection** - There is always at least one in-memory backup copy of your data, along with the additional database copy used in case of complete system shutdown to recover the entire IMDG state.
- **Scalability** – Share--nothing architecture, partitioning, automatic rebalancing with smart data routing and elastic provisioning enables XAP to meet fluctuating and variable unexpected loads. Applications achieve linear scalability without significant re-architecting.

- **Front-end mainframe** – Allows the XAP IMDG to pull data from mainframes, or any source for that matter, via public interfaces (SQL or web services for example). In the event of a cache-miss, reading from the IMDG or once the IMDG is started, XAP accesses the mainframe to fetch relevant data seamlessly.
- **Interoperability** – By use of a schema-less data model through documents. Support for multiple languages / protocols (Java, .NET, REST) and Document Native support of popular programming languages (Java,.Net, C++), APIs (Object/SQL, Map/JCache, JMS, JDBC, JPA, etc.), and frameworks (Spring, JEE) enables smooth integration with existing environments & skill sets, essentially eliminating vendor lock-in.
- **Event Processing** – True in-memory computing, achieved by co-locating code and data. Supports multiple execution scenarios, such as distributed task executors and co-located event processors. Reduces moving parts, and supports advanced parallel processing patterns as well as full messaging semantics (p2p or pub-sub).
- **Configuration and Performance Monitoring** – Includes topology & availability monitoring and data grid queries and reports.
- **Easy Setup** – And installation via proven automation tools and an orchestration framework which includes a set of pre-packaged configuration (blueprints) designed to allow a user to automate the provisioning, configuration, setup and management of XAP clusters within your existing data center or in your private cloud.
- **Ease of Management** – Both underlying infrastructure and grid can be managed from a single point of control for fast and safe provisioning, Also includes desktop UI, web UI, admin API for management and monitoring of the cluster.

On-Line Trade Execution Use Cases

High Frequency Trading

GigaSpaces provides extreme performance, high availability, and scalability to meet the extreme demands of high-frequency trading systems. Scaling tier-based architecture is complex – each tier has its own clustering model and H/A scheme, making it difficult to coordinate among them. With constant processing of massive real-time data (orders and ticks), tier-based architecture results in a proliferation of network hops, slowing the system down more and more, until no more increases can be supported without a complete re-architecture.

GigaSpaces changes the paradigm — the entire front office trading process is done in-memory on a single platform. XAP stores orders, prices, and reference data with the OMS and pricing engine codes that process the orders and enable you to execute faster than your competitors. The end-to-end elasticity enables all trading components to scale up as trade volumes peak without a performance impact.

Arbitrage, pair trading, volatility trading, and liquidity detection – all require super-fast processing of market data and high-speed reaction. XAP makes this all possible.

Regulations are reshaping the trading landscape, from the shift of OTC derivatives towards electronic trading to MiFID II's requirements for best execution, transparency and consolidated trade caps. At the same time, liquidity is fragmenting, and clients are demanding new ways to trade.

The largest sell-side firms, for example, still need to offer a broad, global network of exchange connectivity that includes emerging markets, as well as provide direct market access and services around different asset classes. But multiple systems with ad-hoc connections and complex workflows increase costs and hide risks. XAP's simplicity and in-memory capabilities eliminates the connection sprawl and time delays and greatly reduces costs and risk exposures.

To be competitive in today's trading environment, you need to be first to new market opportunities and be able to cost-effectively execute trades across multiple platforms, instruments and geographies. GigaSpaces' solutions enable real-time trading applications by providing a platform with extreme performance, high availability, and scalability capable of handling any trading volume, variety and velocity.

Bottom Line

XAP is the only in-memory data grid (IMDG) solution that allows you to dramatically improve application performance without allocating expensive servers to handle peak loads and future growth. XAP runs on low cost commodity hardware, is compatible with most open source solutions, can run on-premise or in the cloud and is highly available with no data loss concerns. The need for fewer compute resources also means a “greener” data center footprint.

Moreover, XAP TCO is far superior to other solutions whether comparing upfront system investment, cost per transaction (scan) or cost to maintain the system over time. In addition, XAP from GigaSpaces is the only IMDG available today that supports virtually any internal and external data source necessary to provide your business with powerful insights to gain competitive advantage and drive next generation business models by leveraging real-time data.