GIGASPACES
innovate with confidence

Simpler, Smarter and Faster Insights Big Data Analytics Processing on streaming, hot and historical data

Rajiv Shah
Director of Solution Architect and Professional Services

2019

About GigaSpaces

We deliver the fastest big data analytics processing platform to run your analytics & machine learning in production, at scale

INSIGHT**EDGE**

300+
Direct customers

50+ / 500+
Fortune / Organizations

5,000+
Large installations in production (OEM)

25+
ISVs

GigaSpaces Select Customers

OEMs / ISVs / Partners

AnalyticsXtreme:
Accelerating Your Data Lake by 100X for Real-time Analytics

Your data is immediately searchable, queryable, and available for analytics

- Single logical view for hot, warm and cold data

- Hot data resides on in-memory data grid and historical data on HDFS/Object Store

- Hot data is **mutable** and historical data is **immutable** (parquet)
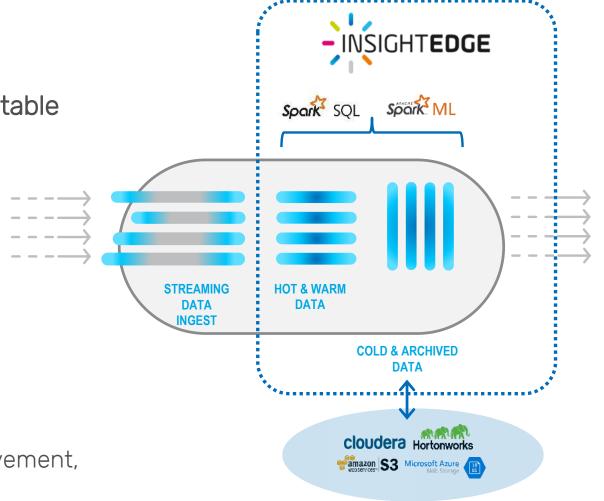
Fast Access

- Fast access to frequently used historical data

Access any data through a unified layer
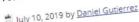
- Analytics (Spark ML)
- Query (Spark SQL)

Automatic lifecycle management

- Automatically handles the underlying data movement, optimization and deletion

INSIGHTEDGE

Spark SQL   Spark ML

STREAMING DATA INGEST   HOT & WARM DATA   COLD & ARCHIVED DATA

cloudera   Hortonworks
amazon web services S3   Microsoft Azure Blob Storage

GigaSpaces Coverage

The insideBIGDATA IMPACT 50 List for Q3 2019

July 10, 2019 by Daniel Gutierrez Leave a Comment

tweet | share | share | email

The team here at insideBIGDATA is deeply entrenched in following the big data ecosystem of companies from around the globe. We're in close contact with most of the firms making waves in the technology areas of big data, data science, machine learning, AI and deep learning. Our in-box is filled each day with new announcements, commentaries, and insights about what's driving the success of our industry so we're in a unique position to publish our quarterly IMPACT 50 list of the most important movers and shakers in our industry. These companies have

IMPACT 50
THE INDUSTRY'S MOST IMPACTFUL COMPANIES
inside BIGDATA

#37 Gigaspaces – In-memory Computing Platform +2

---

451 Research®

REPORT REPRINT

GigaSpaces adds 'data lake accelerator' to its InsightEdge in-memory computing platform

MARCH 12 2019

By Csilla Zsigri, Matt Aslett

GigaSpaces has a long track record of providing in-memory data-processing for data grid and cache use cases. With the latest enhancements to its InsightEdge Platform, it aims to further simplify and accelerate big-data processing for 'instant' business insight.

O'REILLY® Artificial Intelligence Conference presented with intel AI

PUT AI TO WORK
April 15-18, 2019
New York, NY

Operationalizing real-time ML and DL with GigaSpaces, Intel Analytics Zoo, and Optane DC Persistent Memory

Add to Your Schedule
Add Comment or Questic

Yoav Einav (GigaSpaces), Vin Costello (GigaSpaces)
:30pm–2:40pm Tuesday, April 16, 2019
Intel® AI Builders Showcase
Location: Grand Ballroom West

Strata DATA CONFERENCE PRESENTED BY O'REILLY + CLOUDERA

MAKE DATA WORK
29 April–2 May 2019
London, UK

See passes & pricing

How NLP is helping a European financial institution enhance customer experience

Add to Your Schedule
Add Comment or Question

Yoav Einav (GigaSpaces)
Case studies
Location: Capital Suite 13

20th banking technology AWARDS

THE GOLDEN BRIDGE AWARDS
GRAND
BEST 2019
BUSINESS WORLD

NETWORK PRODUCTS GUIDE
IT World Awards 2019 GOLD
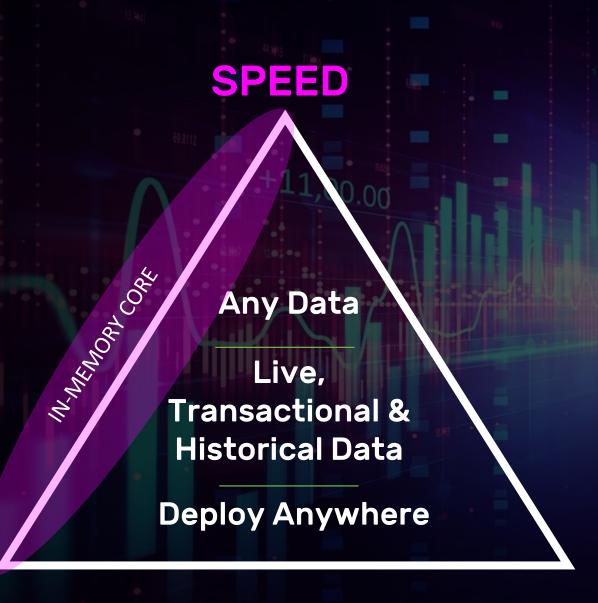
Data Analytics: Undeniable Value to your Business

Dynamic Pricing
Helps grow sales by **30% annually**

Predictive Maintenance
Reduces maintenance costs by up to **75%** per mile (transportation example)

Optimized Operations
Saves **$100sK** in annual savings (banking example)

Personalized Recommendation
Increases conversions by **up to 20X** for brick & mortar stores via location-based promotions

Risk Analysis
Reduces loan losses **by 10-30%**

Fraud Analytics
Reduces losses by **3 to 5%** in mature environments and by **over 30%** in evolving contexts

Call Center Automation
Increases efficiency by **over 90%**

The Velocity of Business

"To prevent fraud, anomaly detection needs to happen against 500,000 txn/sec in less than 200 milliseconds"

"A typical e-commerce website will experience 40% bounce if it loads in more than 3 seconds, including personalization offers"

"A call center receives 450,000 calls/day, each call needs to be routed in less than 60 milliseconds"

FINANCIAL SERVICES

ECOMMERCE

TELCO

Use Cases Spanning Industries Benefit from Near Real-time AI Decision Support Systems Built on GigaSpaces

FINANCIAL SERVICES
- Fraud
- Credit risk scoring
- Customer 360
- Customer churn

INSURANCE
- Usage based insurance
- Customer 360
- Customer churn
- Claims management

RETAIL ECOMMERCE
- Personal recommendations
- Intelligent inventory mgmt.
- Customer 360
- Locations-based promotions

TRANSPORTATION
- Predictive maintenance
- Fleet management
- Customer 360

INDUSTRIAL IOT
- Inventory planning
- Customer 360
- Predictive maintenance

MEDIA/ TELCO
- Customer 360 (incl. churn)
- Intelligent call center routing
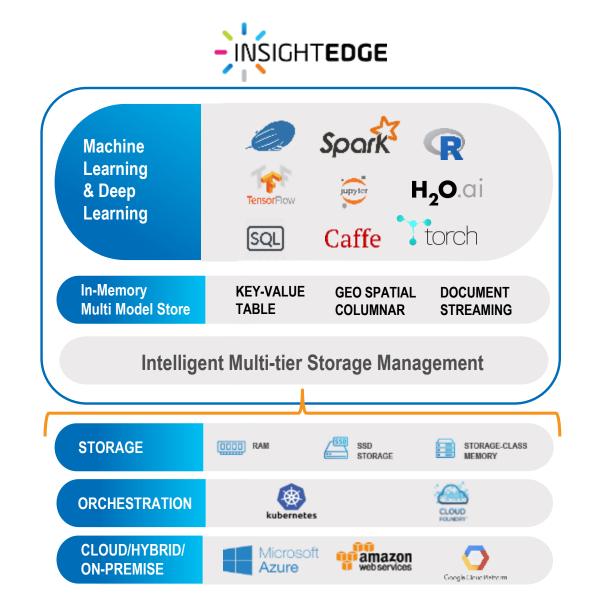- Data Center Infrastructure Monitoring (DCIM)
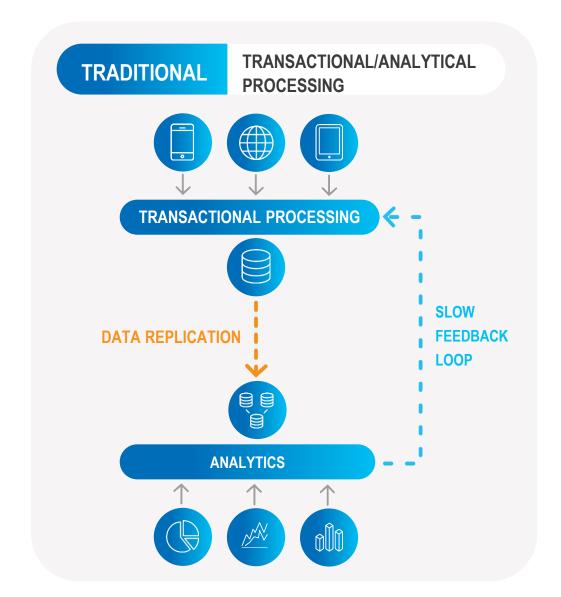- Predictive maintenance

InsightEdge: Unifying Real-Time Analytics, AI and Transactional Processing in One Platform
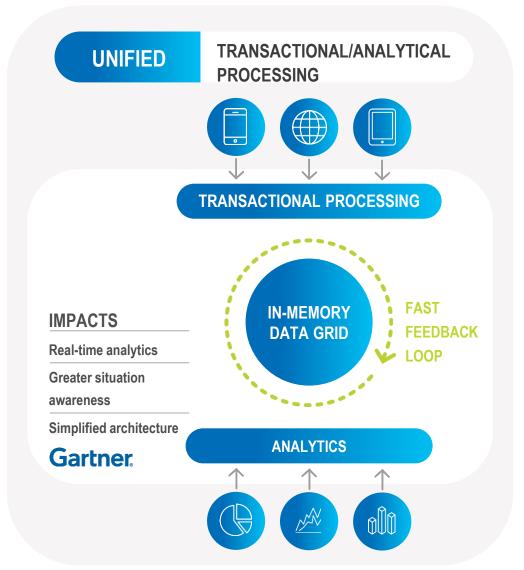
- Rich ML & DL support
- Extreme performance
- Fully Transactional
- ACID Compliance
- Enterprise-grade (Security, High Availability)
- Co-located Apps and Services
- Seamless integration with Big Data ecosystem
  - Data sources (Kafka/Nifi/Talend/etc.)
  - Data lakes (S3/Hadoop/etc.)
  - BI tools (Tableau/Looker/etc.)

Traditional vs. Unified "Translytical" Processing
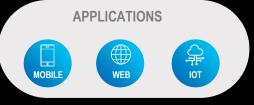

TRADITIONAL TRANSACTIONAL/ANALYTICAL PROCESSING

TRANSACTIONAL PROCESSING

DATA REPLICATION

SLOW FEEDBACK LOOP

ANALYTICS

UNIFIED TRANSACTIONAL/ANALYTICAL PROCESSING

TRANSACTIONAL PROCESSING

IN-MEMORY DATA GRID

FAST FEEDBACK LOOP

IMPACTS

Real-time analytics

Greater situation awareness

Simplified architecture

Gartner.

ANALYTICS

UNIFYING Analytics and Transactional Processing at SCALE & SPEED

INSIGHTEDGE

BI TOOLS
- tableau
- looker
- Power BI
- Qlik Q

DATA LAKE
- amazon web services | S3
- cloudera
- Microsoft Azure Blob Storage

DATABASE & DATA WAREHOUSE
- PostgreSQL
- MySQL
- cassandra
- mongoDB
- snowflake
- TERADATA

APPLICATIONS
- MOBILE
- WEB
- IOT

SECURITY AND AUDITING

MANAGEMENT AND MONITORING

REST ORCHESTRATION

BI & VISUALIZATION
- Grafana
- Apache Zeppelin
- SQL/JDBC
- NOTEBOOK

ANALYTICS, MACHINE & DEEP LEARNING
- Apache Spark
- TensorFlow
- Caffe
- torch
- SQL
- MACHINE LEARNING
- DEEP LEARNING
- STREAMING

APPS & MICROSERVICES
- Java
- Scala
- python
- R
- {REST API}
- SPARK JOBS
- MICROSERVICES (REST)
- EVENT PROCESSING

CORE

IN-MEMORY DATA GRID

MULTI MODEL STORE
OBJECTS, JSON, KEY-VALUE, TABLES, TEXT, GEO SPATIAL, GRAPH

EVENT PROCESSING

PERSISTENT MEMORY

CDC Engine
CR8

MAP/REDUCE

RPC & MAP/REDUCE

MemoryXtend - MULTI-TIERED STORAGE
- RAM
- SSD
- DATA LAKE

WAN GW - MULTI SITE REPLICATION
WAN GATEWAY

Apache ZooKeeper™ CLUSTER MANAGEMENT & SERVICE DISCOVERY

ON-PREMISE

CLOUD
- Microsoft Azure
- amazon web services
- Google Cloud Platform

HYBRID
- docker
- kubernetes
- RED HAT OPENSHIFT Container Platform
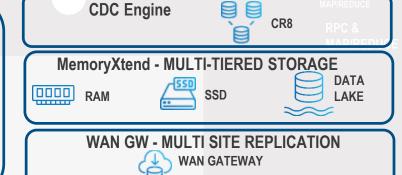
Ultra-low latency and high throughput transactional processing IMDG

Partitioned In-Memory Grid
Shared-nothing, linear scalability, elastic capacity

Co-Location of Data and Business Logic
Co-located ops, event-driven, fast indexing
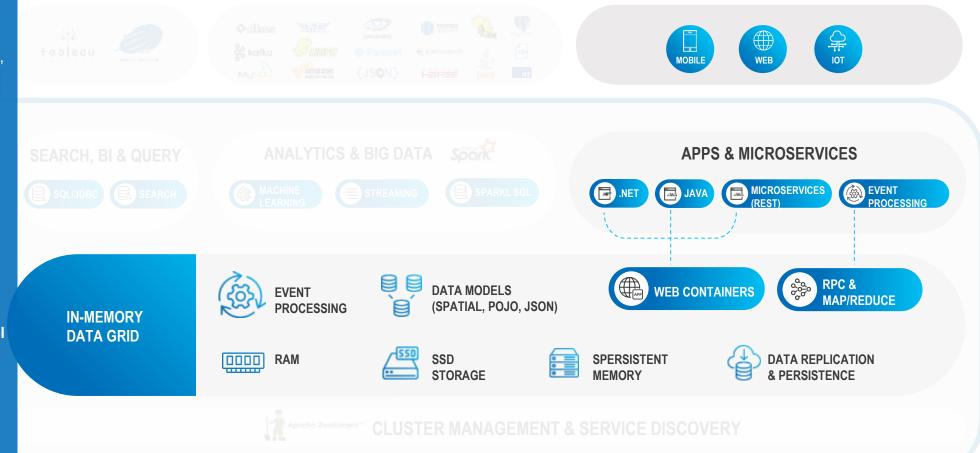
Event-Driven Processing and Map/Reduce

No Downtime
Auto-healing, multi-data center replication, fault tolerance

Fast Indexing Multi-Data Model
POJO, .NET, Document/JSON, Geospatial, Time-series

Seamless Integration with Java/Scala ecosystem

Cloud, Kubernetes, Docker Native

MOBILE
WEB
IOT

SEARCH, BI & QUERY

SQL/JDBC   SEARCH

ANALYTICS & BIG DATA   Spark

MACHINE LEARNING   STREAMING   SPARKL SQL

APPS & MICROSERVICES

.NET   JAVA   MICROSERVICES (REST)   EVENT PROCESSING

IN-MEMORY DATA GRID

EVENT PROCESSING

DATA MODELS (SPATIAL, POJO, JSON)

WEB CONTAINERS

RPC & MAP/REDUCE

RAM

SSD STORAGE

SPERSISTENT MEMORY

DATA REPLICATION & PERSISTENCE

CLUSTER MANAGEMENT & SERVICE DISCOVERY

ON-PREMISE

CLOUD   Microsoft Azure   amazon web services

HYBRID   docker   kubernetes   CLOUD FOUNDRY

Co-located Analytics and AI with Transactional Processing

SEARCH, BI & QUERY

SQL/JDBC   SEARCH

Distributed SQL-99

Real-time integration with Tableau and Business Intelligence tools

JDBC driver

ANALYTICS & BIG DATA — Spark

MACHINE LEARNING   STREAMING   SPARK SQL

Spark for ML and leading DL frameworks

Push-down predicate for ultra-low latency filter (30x faster)

Shared RDDs/DataFrames

Streaming with 99.999% availability

Deep Learning with Intel BigDL

Graph processing, text mining, geospatial

APPS & MICROSERVICES

.NET   JAVA   MICROSERVICES (REST)   EVENT PROCESSING

WEB CONTAINERS   RPC & MAP/REDUCE

MOBILE   WEB   IOT

SECURITY AND AUDITING   MANAGEMENT AND MONITORING   MANAGEMENT AND MONITORING

ON-PREMISE   CLOUD   HYBRID

Benchmark (in IOPS)

Random Reads
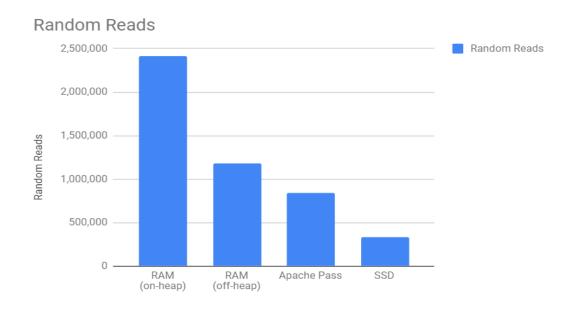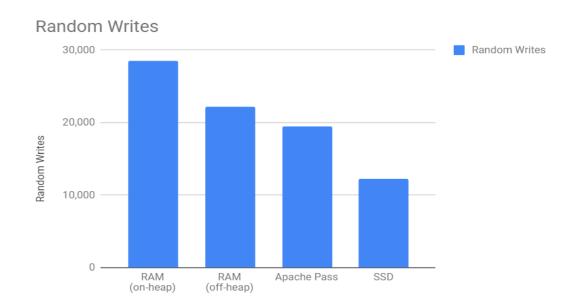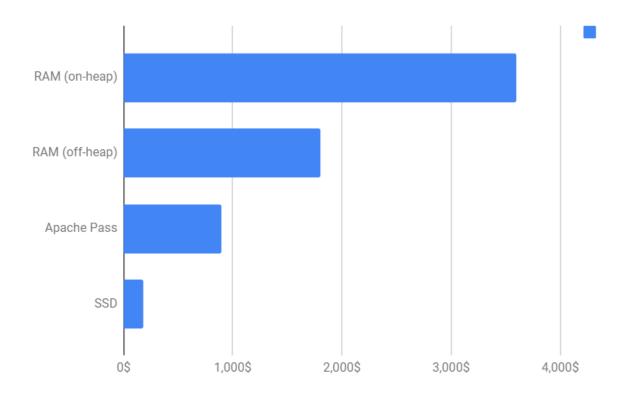
- Persistent Memory +249% than SSD
- RAM (off-heap) +350% than SSD

Random Writes

- Persistent Memory +159% than SSD
- RAM (off-heap) +180% than SSD

Costs Analysis
for 5GB usable data

- CAPEX reduction of up to 50% with RAM off-heap vs. on-heap
- CAPEX reduction of up to 75% with AEP vs. RAM on-heap
- OPEX reduction by X10

Tiered Storage Architecture
Higher Performance – Optimized TCO

10X less expensive
than only RAM maintaining
in-memory performance

- Volatile Memory
- Load/Store Instructions
- Cache Line Granularity

- Non-Volatile Storage
- Load/Store Instructions
- Cache Line Granularity

- Non-Volatile Storage
- I/O Commands
- Bock Granularity

Cost ($/GiB)

CPU Registers ~0.1ns

CPU Caches (L1, L2, L3, L4) 1-10ns

DDR DRAM ~80-100ns

Persistent Memory <1us*

NAND SSD 10-100us

Hard Disk Drives (HDD) ~10ms

Tape ~100ms

Latency

Capacity

(*) See vendor specifications

Figure 2: Memory-Storage Hierarchy with Persistent Memory Tier
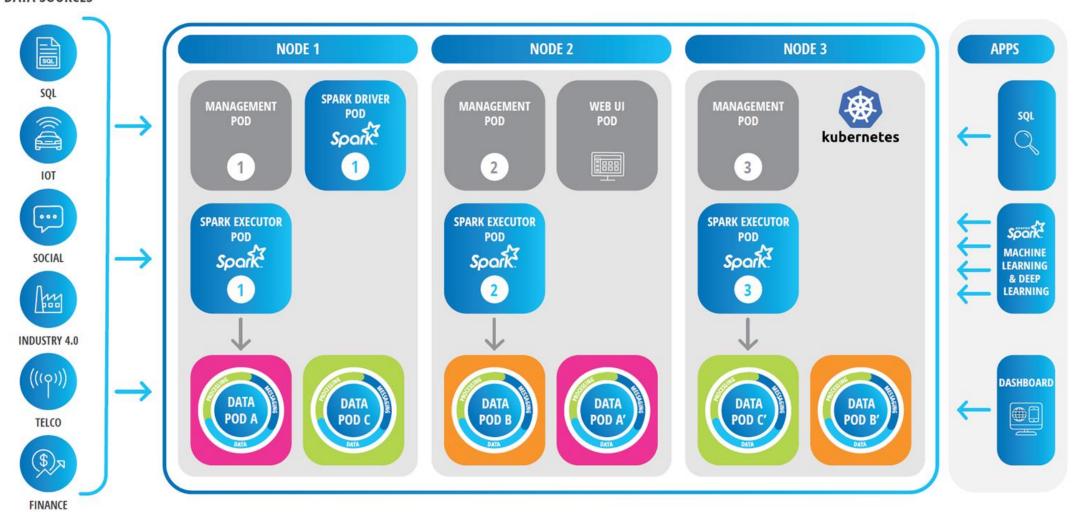
Define which data resides on which layer per class and per field

Kubernetes and Docker

LAMBDA ARCHITECTURE IS COMPLICATED
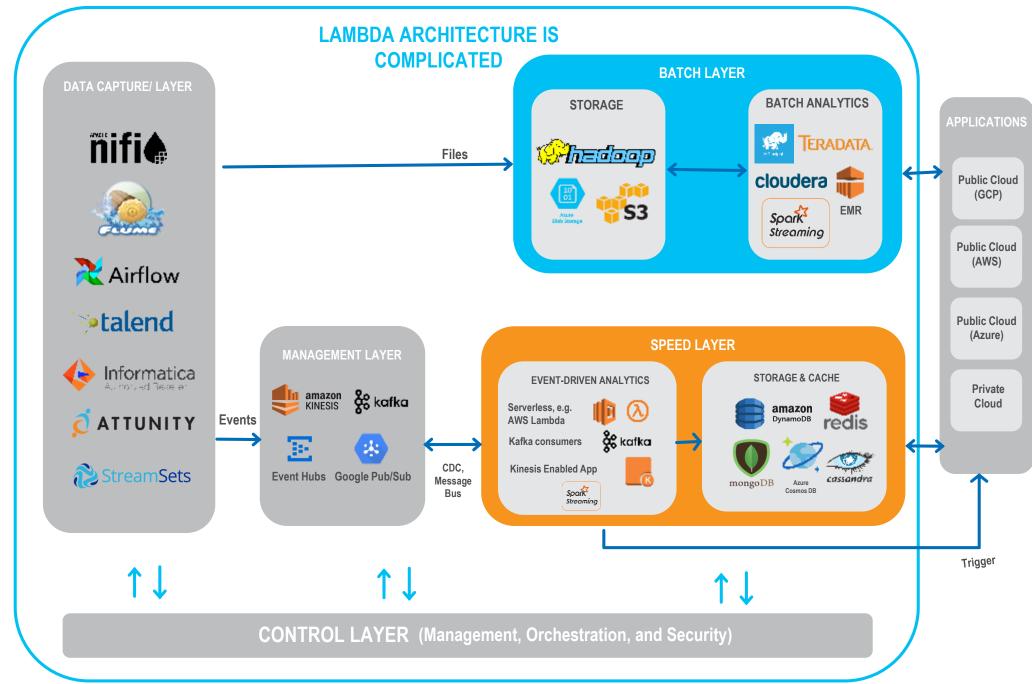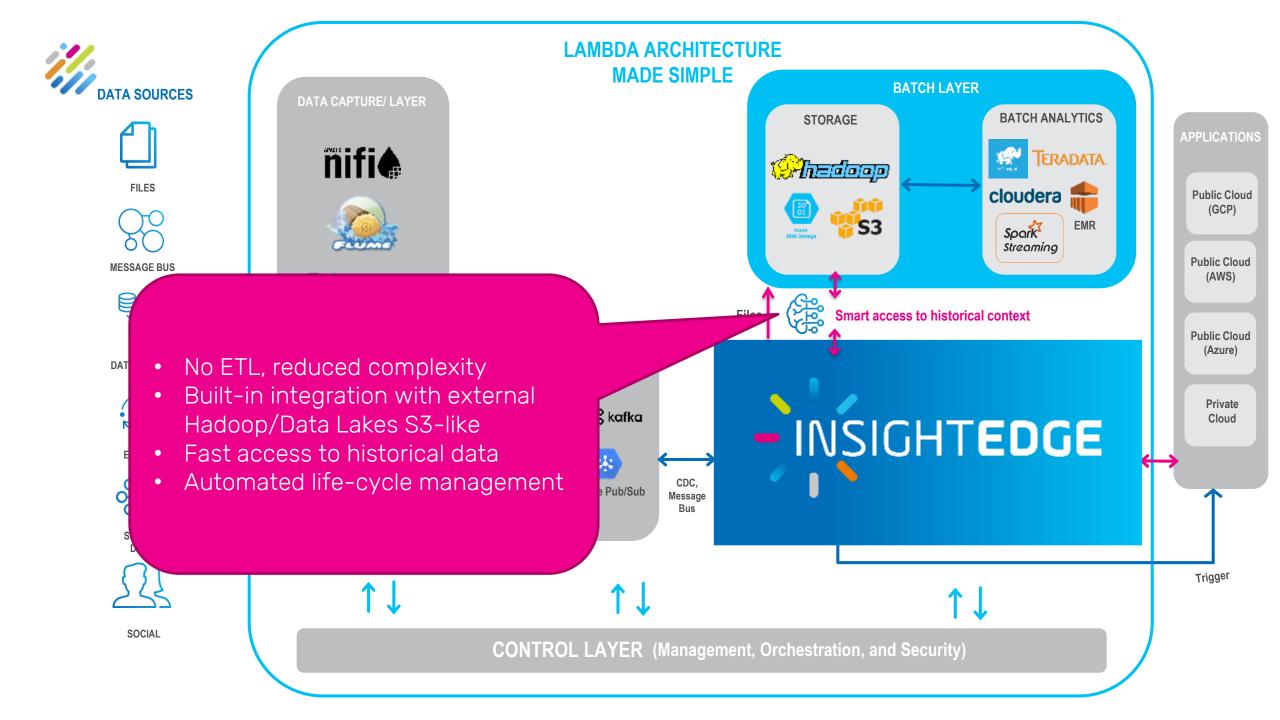
Leverage leading BI Platforms

Tableau

Looker

Qlik

Power BI

YCSB

Workload A: Update heavy
This workload has a mix of 50/50 reads and writes.
Workload B: Read mostly
This workload has a 95/5 reads/write mix.
Workload C: Read only
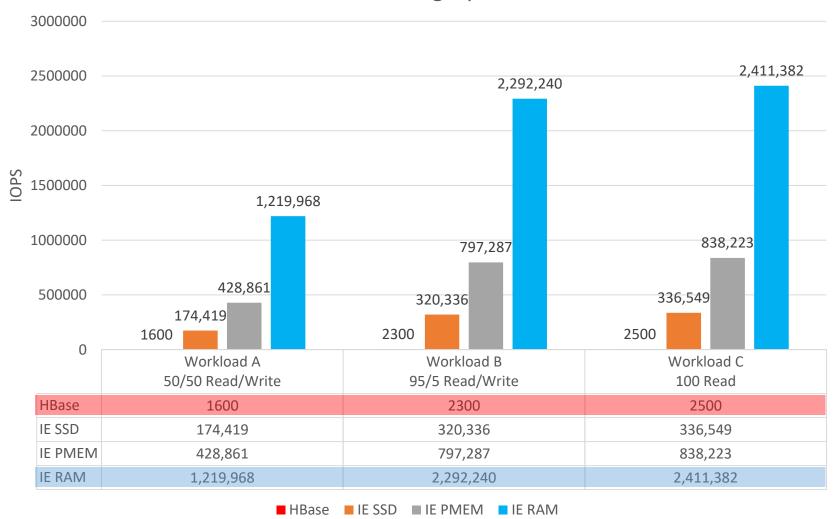This workload is 100% read.
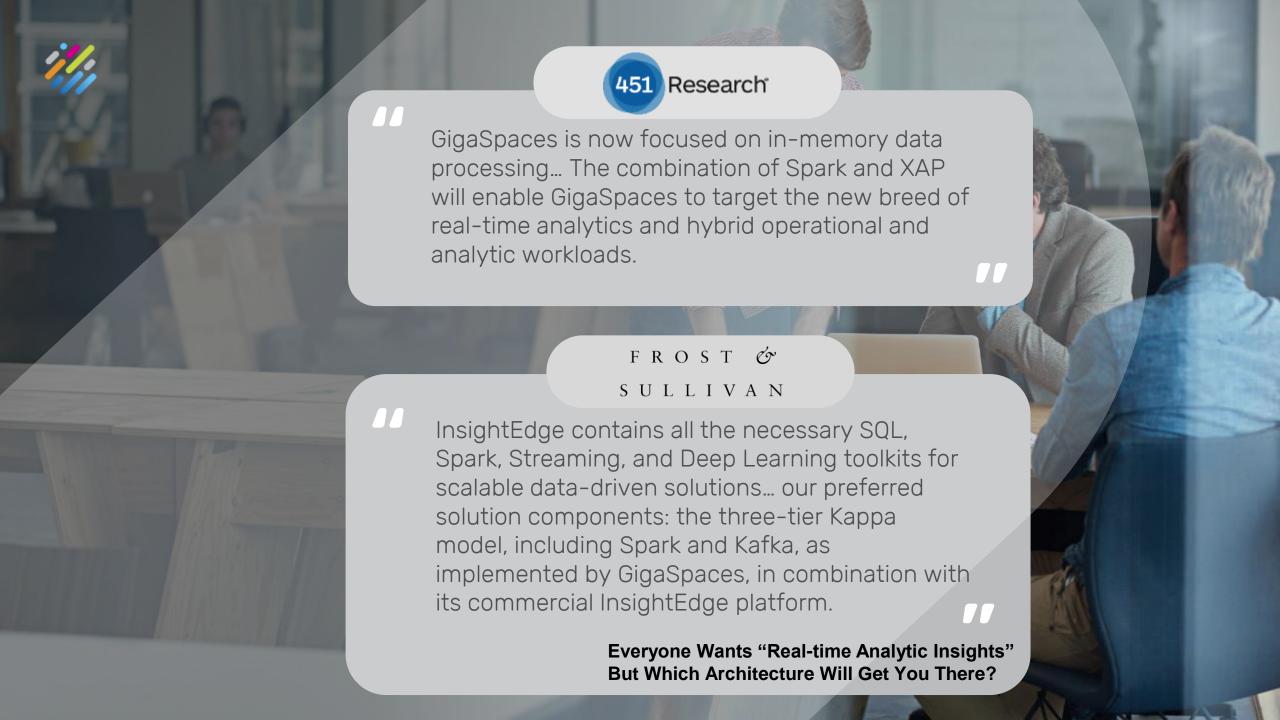
Per Node
Replication Factor: 2
Record size:1KB
RAM: 32GB
CPU: 16 cores
Disk: 1.2TB SSD

NoSQL vs. GigaSpaces

	Workload A 50/50 Read/Write	Workload B 95/5 Read/Write	Workload C 100 Read
HBase	1600	2300	2500
IE SSD	174,419	320,336	336,549
IE PMEM	428,861	797,287	838,223
IE RAM	1,219,968	2,292,240	2,411,382

■ HBase ■ IE SSD ■ IE PMEM ■ IE RAM

X100 Faster

451 Research

"GigaSpaces is now focused on in-memory data processing... The combination of Spark and XAP will enable GigaSpaces to target the new breed of real-time analytics and hybrid operational and analytic workloads."

FROST & SULLIVAN

"InsightEdge contains all the necessary SQL, Spark, Streaming, and Deep Learning toolkits for scalable data-driven solutions... our preferred solution components: the three-tier Kappa model, including Spark and Kafka, as implemented by GigaSpaces, in combination with its commercial InsightEdge platform.

Everyone Wants "Real-time Analytic Insights" But Which Architecture Will Get You There?

CASE STUDY: **Fast Global Fabric for Risk, Trading and Market Data**

RISK

SOCIETE GENERALE

BUSINESS CHALLENGE:

- Prior to executing a trade, a credit check needs to run and guarantee that the counterparty is not exceeding their limit

TECHNICAL CHALLENGE:
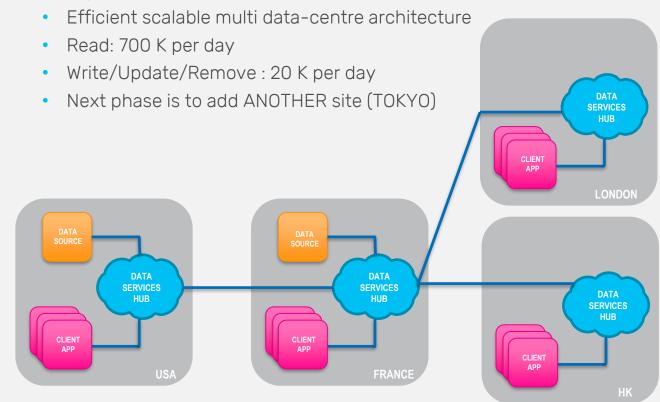
- Complete control over all eTrading platforms
- Regulatory enforcement set by RISK rules on all users trades on a daily basis
- Regulation analysis and checks
- Client onboarding
- Traversal framework
- Referential data for other apps

IMPLEMENTATION:

- All reservations, limits and client data is stored in the GigaSpaces in-memory platform
- All the requests are executed via the platform
- GigaSpaces is used in front of the database to speed up data access
- A worldwide deployment is done (Paris, NY and London) with GS asynchronous replication between each site to populate the data in NY and London

RESULTS:

- Three sites with 99.999 HA, replicated WW (Paris, London, New-York and Hong Kong)
- Reduced cluster and component sprawl
- Real-time risk analysis and credit checks complying with regulations
- Efficient scalable multi data-centre architecture
- Read: 700 K per day
- Write/Update/Remove : 20 K per day
- Next phase is to add ANOTHER site (TOKYO)

CASE STUDY: DYNAMIC PRICING & OPTIMIZATION

BUSINESS CHALLENGE:
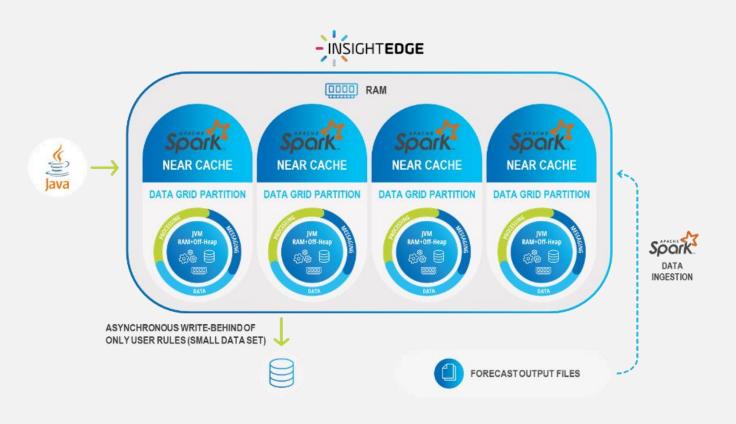
- Demand forecasting and price optimization in real-time based on threshold changes

TECHNICAL CHALLENGE:

- Ingest ~ billion of records in minutes
- Ability to query data from multiple geographies in real-time at low latency
- Ability to update with low latency multiple locations to adjust forecast and influence
- Cloud nativeness

RESULTS:

- Agility: Reduced forecasting ingestion from 3 hours to 8 minutes
- Live interactive querying and analytics through Spark SQL < 150ms latency

CASE STUDY: BOOKING AND FLIGHT AVAILABILITY

BUSINESS CHALLENGE:
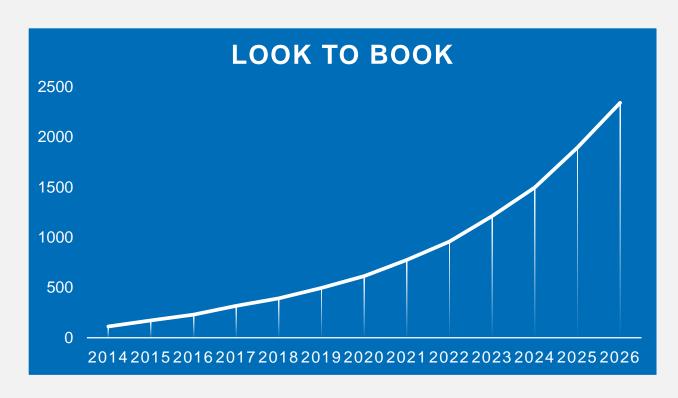
- Flight availability forecasting real-time based on various factors: date, city pair, #seats requested, marketing class, Point of Sale (PoS), quota limits, traffic restrictions, etc.

TECHNICAL CHALLENGE:

- Various internal systems (Reservation, Shopping, eCommerce Systems)
- Open API for external systems: Airlines, Global Distribution Systems (GDSs) and BOTs (automated searching).
- Auto scaling and sub-sec latency
- Multi tenancy (small/med/large airlines)

RESULTS:

- Querying and analytics response time < 50ms latency
- High Performance with up to 200K transaction/sec
- Scaling Near Linear (X100)
- Increase throughput by X& and reduce network overhead by 10%

LOOK TO BOOK

2500	
2000	
1500	
1000	
500	
0	

2014 2015 2016 2017 2018 2019 2020 2021 2022 2023 2024 2025 2026

Ratio of Bookings per Availability Requests increases by 100

CASE STUDY: PriceRunner
Compares Prices for Millions of Offers in Milliseconds

BUSINESS CHALLENGE:

- PriceRunner receives prices from 18,000 different merchants and has 4.4 million unique visitors per month, needed to ensure real-time comparisons for their customers at high peak periods such as the night before Black Friday where traffic increases between 10-20 times the normal traffic.

TECHNICAL CHALLENGE:

- Support scalability requirements at peaks without compromising performance
- No downtime
- Real-time analytics on transactional data
- Event-driven applications powering integrated applications
- Microservices architecture for rapid development and deployment

18,000 different merchants

200 million prices updates

1 Billion requests a month

5-8 millisecond performance

"Innovation is a key tenant of our strategy, and 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."

Roger Forsberg, CTO
PriceRunner

CASE STUDY: Real-Time Pricing Engine

BUSINESS CHALLENGE:
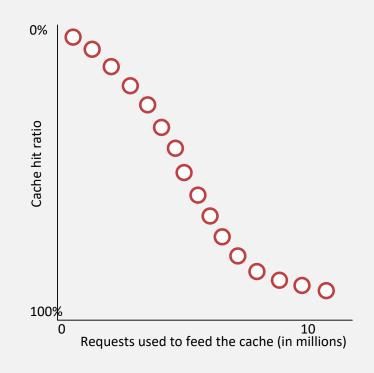
- Dynamic pricing engine based on CO2 tax regulations for B2B and B2C

- Many car configurations are unique, but all parts are not significant for CO2 calculation

TECHNICAL CHALLENGE:

- The current pricing engine workload is around 60 to 80 calculations/s, expected to increase to 2000

- Pricing calculations are obsolete after 24 hours.

- Each CO2 returned value must be exact

- All requests (both internal and external) must be equally treated

RESULTS:

- Querying and analytics response time < 100ms latency

- Reduce infrastructure footprint by a factor of X4-6

- Scaling up by X20

Cache hit ratio vs. Requests used to feed the cache (in millions): 0% at 0, trending to 100% at 10+

Pricing Requests increases by 20x

USE CASE:

SMART AGENT ASSISTANCE

BUSINESS CHALLENGE:

- Enhance customer experience with quicker First Call Resolution
- Reduce Average Handle Time for optimized efficiency

TECHNICAL CHALLENGE:

- Ingestion of millions of CRM cases and data from other repositories into a unified analytics platform
- Leveraging ML models in real time
- Continuous model training

DATABANK

Ticket ID   #54367

Customer Name
#54367

Type
Enterprise

Support Level
Bronze

Last Contact Date
20.12.18

GIGASPACES

| DATA SOURCES | CUSTOMER | CUSTOMER TICKET |

95.32% — Credit Limit exceeded #56409
93.05% — Authentication required #33487
86.16% — Beneficiary account unknown #180762
77.98% — Beneficiary account dormant #180762
71.53% — Intermediary bank changes #60975

International Payment declined

Case Description
International payment to supplier declined
Read more

Case Resolution
Check that credit limit is not exceeded
Read more

Case Description
Check here to email instructions to customer
Email

Reducing mean time to resolution by 5-10X
Average time of 50ms to search and find similar cases

CASE STUDY: Fraud and Money Laundering Detection in Real-time

BUSINESS CHALLENGE:

- Detecting fraud on mobile payment applications in real-time
- Detecting the deposit of the same check in multiple accounts at different banks in real-time
- User experience: application availability 24×7
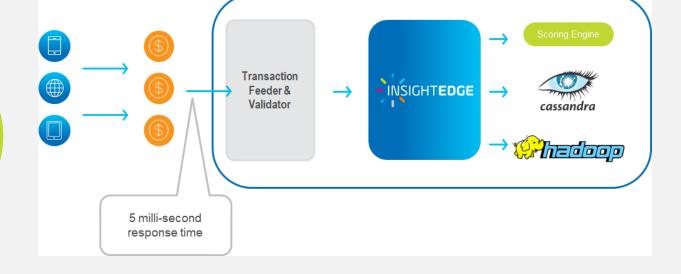- TCO reduction: reduce dependency on expensive RDBMS (Oracle)

TECHNICAL CHALLENGE:

- IMC Platform to ingest 4 TB of data daily
- Fully consistent transactional In-Memory Map-Reduce
- Millisecond response
- Analyze and validate against a large dataset of live (multiple TB) in memory and archived data (to Cassandra NoSQL and Hadoop)

RESULTS:

- Sub-second response for accurate fraud detection to stop the transaction
- TCO Reduction: RAM and SSD for runtime data compared to Oracle DB or SAN
- Fault-tolerant, highly available, scaling on demand

Ingest **4 TB** daily
Handle **1.5M** events per second

Transaction Feeder & Validator

INSIGHTEDGE

Scoring Engine

cassandra

hadoop

5 milli-second response time

CASE STUDY: Instant Payments for real-time transactions and high reliability to enhance the overall customer experience

BUSINESS CHALLENGE:
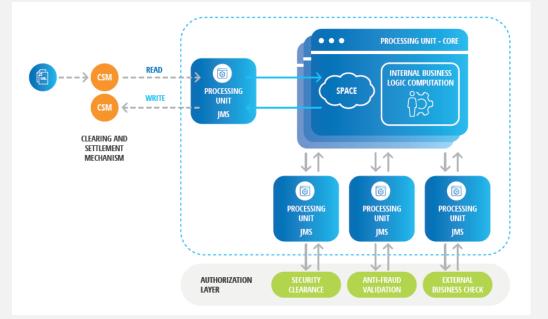
- Enable and accelerate instant payment solutions and meet regulatory requirements on a global scale
- Automatically track purchases and other server-to-server communication in real time
- Store payment transactions, order information and other sales internally

TECHNICAL CHALLENGE:

- Ability to handle added data volumes 15k payment/sec receipts introduced by management of new SEPA European payment regulation
- Assure no-downtime for mission critical service

RESULTS:

- Running low-latency payment and business logic calculations
- **No downtime** assured
- Real-time analytics and Machine Learning - preventing fraud and adherence to regulations
- Design to deployment in just a few months leveraging microservices architecture
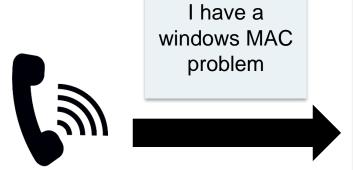
Payment transaction in 500 milliseconds
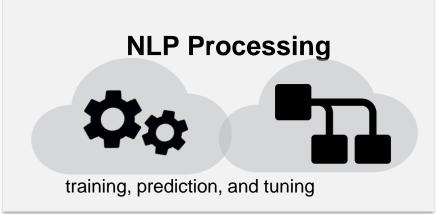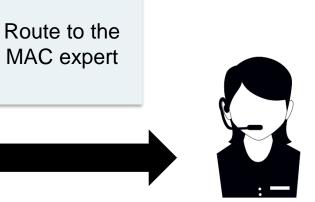
End-to-end validation in seconds

Stop Pressing 0 Or *

Automatic routing to the right agent for the perfect personalized experience

I have a windows MAC problem

NLP Processing

training, prediction, and tuning

Route to the MAC expert

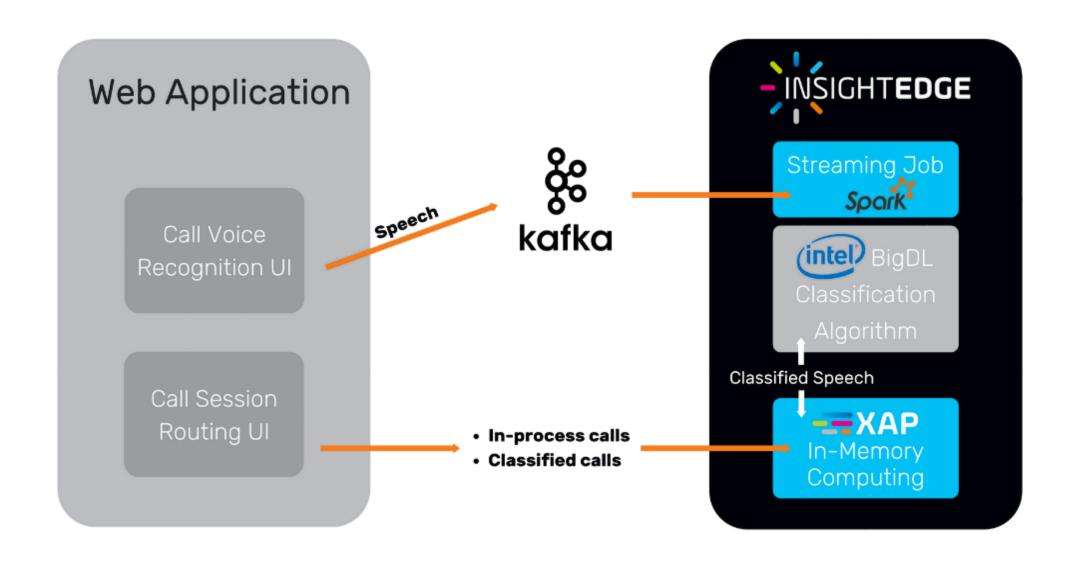| User speaks using web interface | Browser converts speech to text and sends to controller | Controller writes data to InsightEdge and to Kafka topic | Spark job listens on Kafka topic and using BigDL model, creates prediction | BiGDL writes Prediction to InsightEdge data grid | InsightEdge event processor listens for Prediction data and routes call session |

Operationalizing AI Example – Automatic Call Routing

CASE STUDY: PriceRunner
Compares Prices for Millions of Offers in Milliseconds

eCommerce

BUSINESS CHALLENGE:

- PriceRunner receives prices from 18,000 different merchants and has 4.4 million unique visitors per month, needed to ensure real-time comparisons for their customers at high peak periods such as the night before Black Friday where traffic increases between 10-20 times the normal traffic.
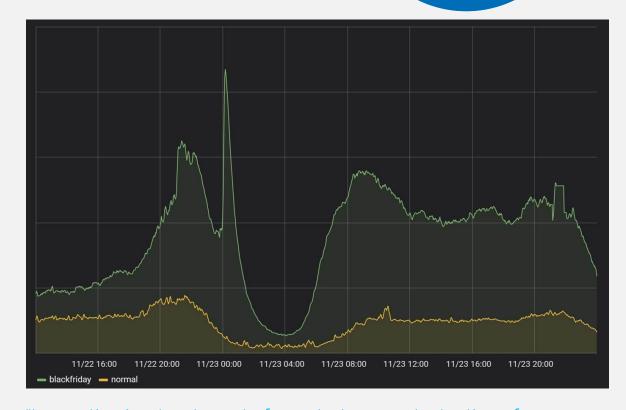
TECHNICAL CHALLENGE:

- Support scalability requirements at peaks without compromising performance
- No downtime
- Real-time analytics on transactional data
- Event-driven applications powering integrated applications
- Microservices architecture for rapid development and deployment

18,000 different merchants

200 million prices updates

1 Billion requests a month

5-8 millisecond performance

"Innovation is a key tenant of our strategy, and 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."

Roger Forsberg, CTO
PriceRunner

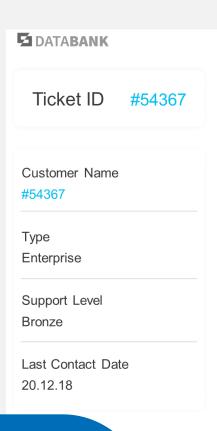# USE CASE:

SMART AGENT ASSISTANCE

BUSINESS CHALLENGE:

- Enhance customer experience with quicker First Call Resolution
- Reduce Average Handle Time for optimized efficiency

TECHNICAL CHALLENGE:

- Ingestion of millions of CRM cases and data from other repositories into a unified analytics platform
- Leveraging ML models in real time
- Continuous model training

Reducing mean time to resolution by 5-10X
Average time of 50ms to search and find similar cases

INSTANT INSIGHTS TO ACTION

EXTREME PERFORMANCE

TCO OPTIMIZATION

MISSION CRITICAL AVAILABILITY

<1 sec from data to insight to action

millions of IOPS

10X less expensive than only RAM with In-memory performance

No Downtime at leading enterprise customers for YEARS And still counting

WHY GIGASPACES?

- Real-time insights
- Boost your performance
- Simplify your architecture
- Lower TCO / Enhance ROI

Enterprise Grade System of Record

Optimized Data Replication:
Field-proven, reliable, high performance replication mechanism to replicate data between peer nodes in the data grid

Locking Support:
RDBMS locking and transaction isolation for robust and hassle-free data access

Data Partitioning:
Transparent content-based data partitioning to evenly and intelligently distribute data across your cluster

Multi-Site Deployment: Replicate and share data between multiple, geographically-distributed, active clusters for global activity

Transaction Support:
Full transaction support, including local, distributed and XA transactions

Network Segmentation Protection:
Ensure data remains consistent in case of network segmentations of all types

Write Behind:
Asynchronous and reliable propagation of data to any external data source

Security:
Role-based authentication for data and operations, Support for Kerberos, Spring, TLS and more

Advanced Querying & Indexing

Querying:
Sophisticated query engine with support for SQL and example queries

Indexing:
Predefined and add-hoc Property indexing for fast data access

Projection API
Customize the query's result set by defining which fields should be returned

Geospatial:
Enhance your data model with shapes and use spatial operations to find matches

Change API:
Update data by specifying only the required change instead of the entire updated object

Full Text Search:
Go beyond plain text with regular expressions, fuzzy search, proximity matching and more

Aggregations
Sum, Avg, Min, Max, GroupBy and more, or even your own user-defined aggregations

SQL Functions:
Abs, Round, Length, Upper, Lower and more, or even your own user-defined functions

Data Model Flexibility & Interoperability

Native:
Highly optimized, POJO driven API which exposes all the unique capabilities of the platform

JPA:
Support data grid access using the standard JPA API for seamlessly scaling your JEE data access layer

Document:
Completely schema-free data API that supports upgrading the application's data model on the fly

Key-Value:
Simple and intuitive Map-based interface for simple caching scenarios

.Net:
Native C# interface that enables any .NET application to access the data grid

REST API:
Standard REST endpoint provides access to the data grid from any app, Platform and programming language

Cross Language Access:
support for heterogeneous environments, with seamless interoperability among them all

Messaging & Event Features

Publish/Subscribe Messaging:
Propagation of any event that takes place in the data grid to listeners using the publish/subscribe paradigm

Point-to-Point Messaging:
Support for implementation of complex workflows and triggering of processing logic across the data grid

Content Based Routing:
Routing of events to relevant cluster members based on their content

Durable Notifications
Fully durable pub/sub messaging for data consistency and reliability

FIFO Groups
Ensure in-order and exclusive processing of events belonging to the same group, while parallelizing across groups

Workflow Support:
Implement complex workflows using event propagation and sophisticated event filtering

43

Collocation of Data and Business Logic

Spring on Steroids:
Deployment, provisioning and proactive management of any spring application, with or without a data grid

Master-Worker Support:
Intuitive and highly scalable master-worker implementation for distributing computation-intensive tasks

Code and Data Collocation:
Deployment of business logic and data as a single coherent unit for optimized performance

Dynamic Code Execution:
Dynamic code shipment and map/reduce- like execution across the grid for optimized processing and data access

Robust Remoting Support:
Built on top of the data grid to provide fault tolerance, service auto discovery, cluster wide invocations and more

Security:
Customizable security policy to control who can run dynamic code on the grid

Grid Health Transparency & Monitoring

UI Based Management
Web-based dashboard app for easy monitoring & management of deployed app. Enhanced data grid console for cluster wide queries or single Space instance queries.

REST Admin API:
Comprehensive and intuitive API for monitoring and controlling every aspect of your cluster and application

Event Tracking:
Trace Cluster Events as they happen for improved visibility & easier troubleshooting
(available through both admin API and UI)

Single Click Deployment:
Support for distribution, provisioning and management of application deployments across any number of hosts

Alerts:
Out-of-the-box identification & notification of risky situations (e.g., above-normal CPU utilization or data replication failure)

Application Dependencies:
Deploy modules as an application ensuring order of deployment

Event Containers Monitoring:
Trace embedded and remote event containers

Client side Cache Monitoring:
Discover client-side cache and views connected to your spaces

Extensible Metrics Framework:
Measure both space and user-defined metrics, integrated with any tool (InfluxDB and Grafana out of the box)

WAN Replication Monitoring:
Discover client-side cache and views connected to your spaces