

Operationalizing Artificial Intelligence

Instant Insights to Actions

Rajiv Shah

Director of Solution Architect and
Professional Services

2018





About GigaSpaces

GigaSpaces provides the leading in-memory computing platform for fast data analytics and extreme transaction processing. With GigaSpaces InsightEdge Platform, and its core engine XAP, organizations can overcome the complex obstacles of big data and advanced analytics to deliver real-time business impact. GigaSpaces offices are located in the US, Europe and Asia.

300+

Direct customers

50+ / 500+

Fortune / Organizations

5,000+

Large installations in production (OEM)

25+

ISVs

Leveraging the strength
of combining innovation & mature technology



In-memory Insight Platform
which unifies fast-data analytics,
AI and transactional processing
**Leverages XAP as its core
technology**



In-Memory Computing Platform
for ultra low-latency, high-
throughput transaction
processing



InsightEdge Platform is an open-source in-memory insight platform **unifying fast-data analytics, artificial intelligence and transactional processing**, for instant business insights and actions



**INSTANT
INSIGHTS**



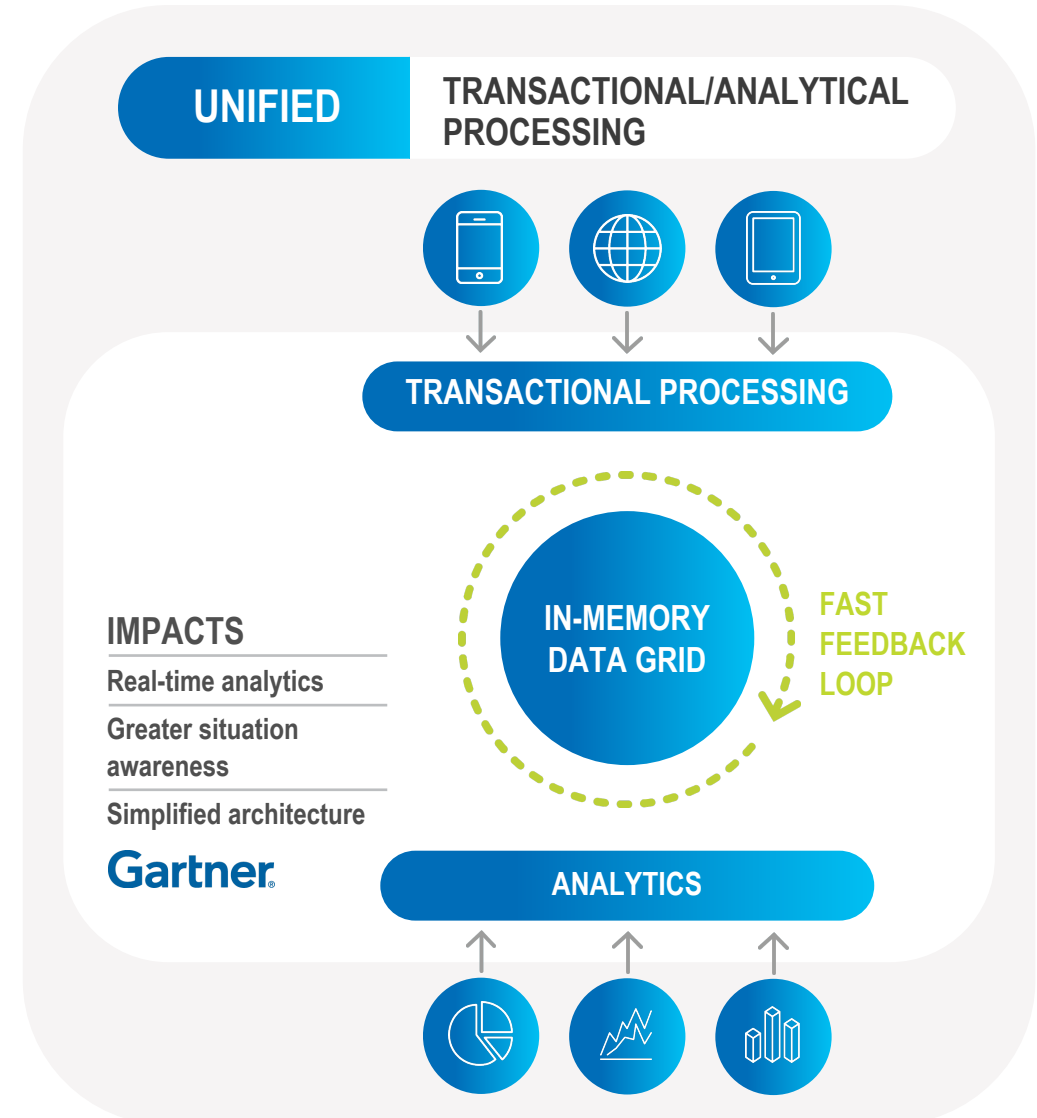
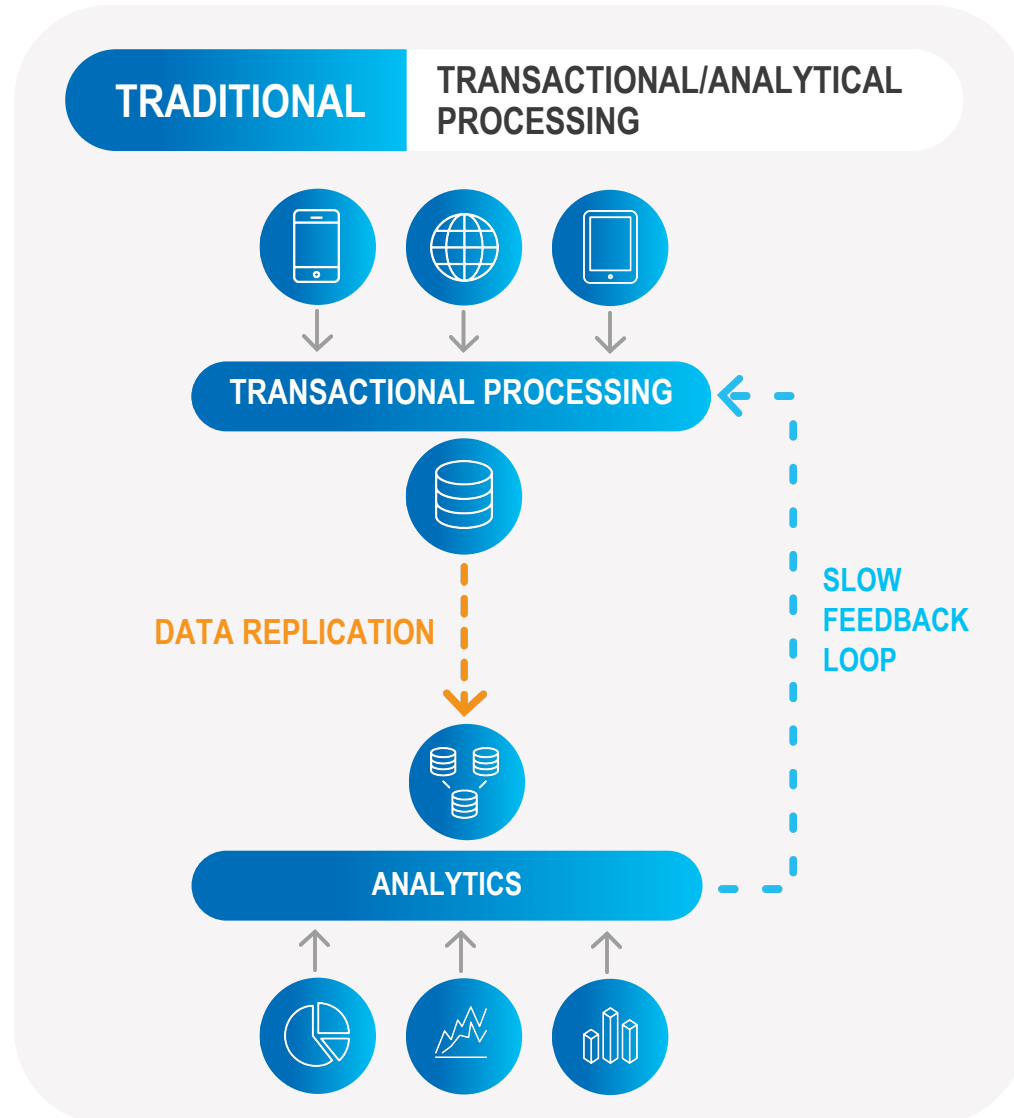
**EXTREME
PERFORMANCE**



**TCO
REDUCTION**

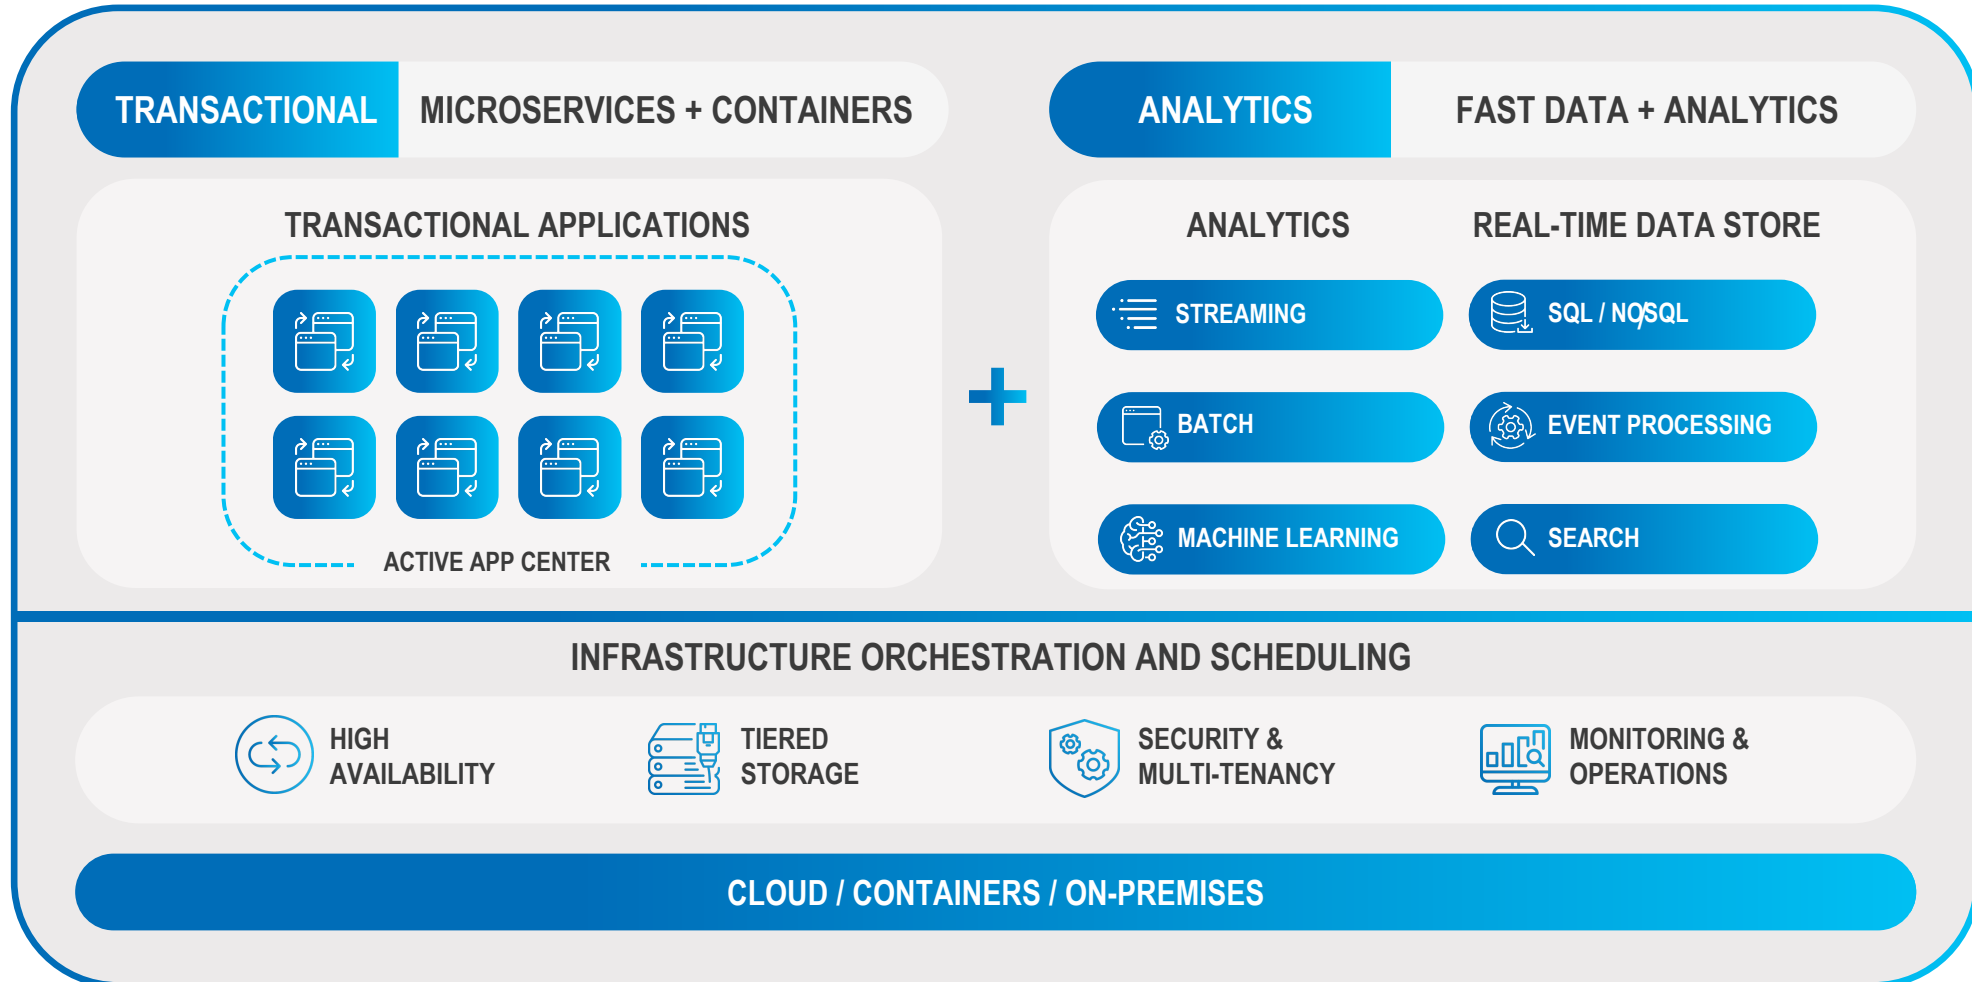


Traditional vs. Unified “Translytical” Processing



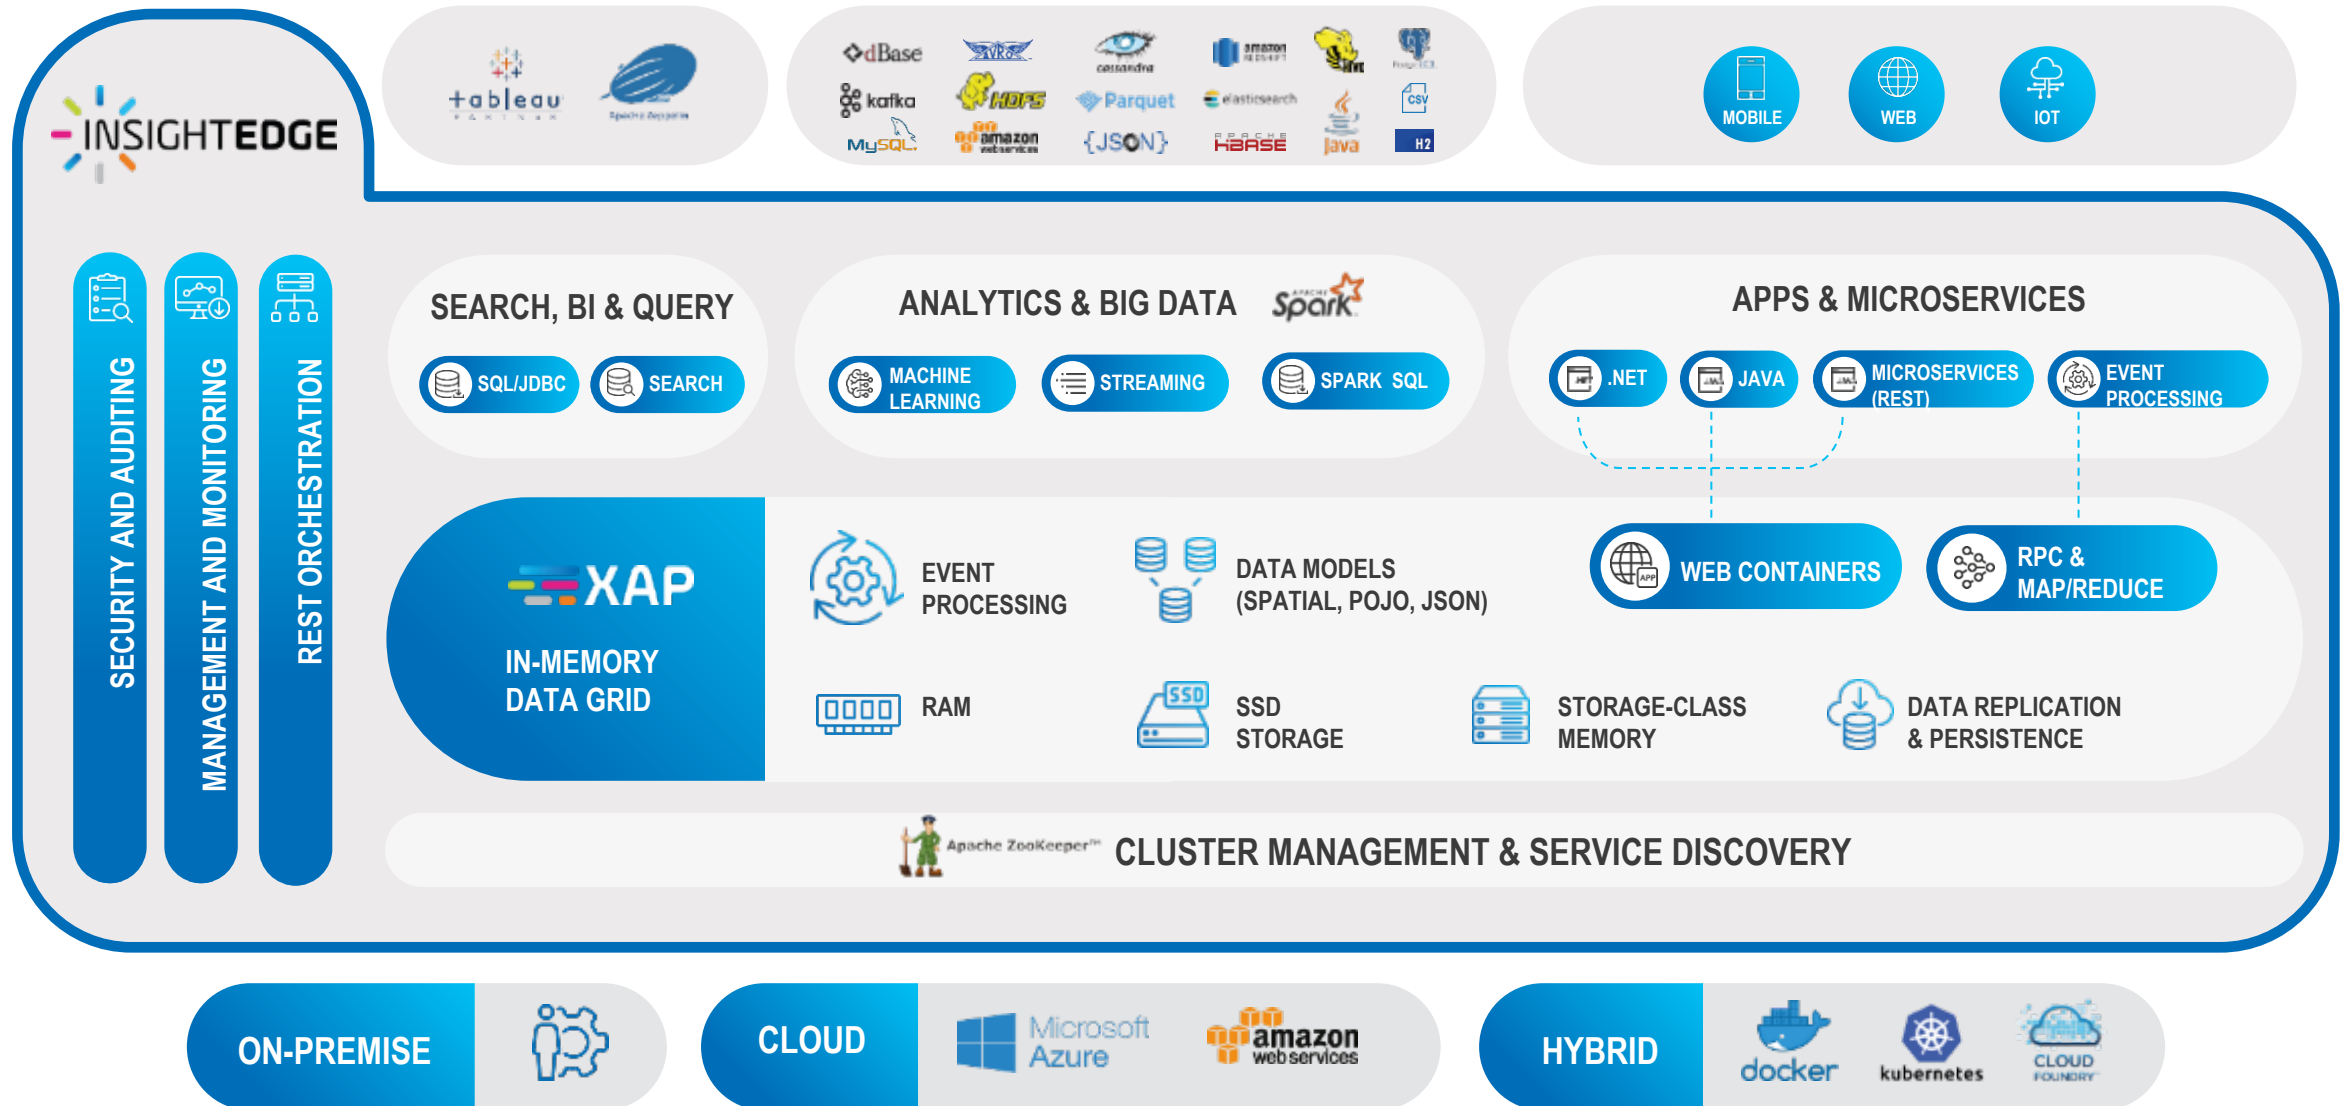


The Digital Core: Unified Fast Data & Microservices



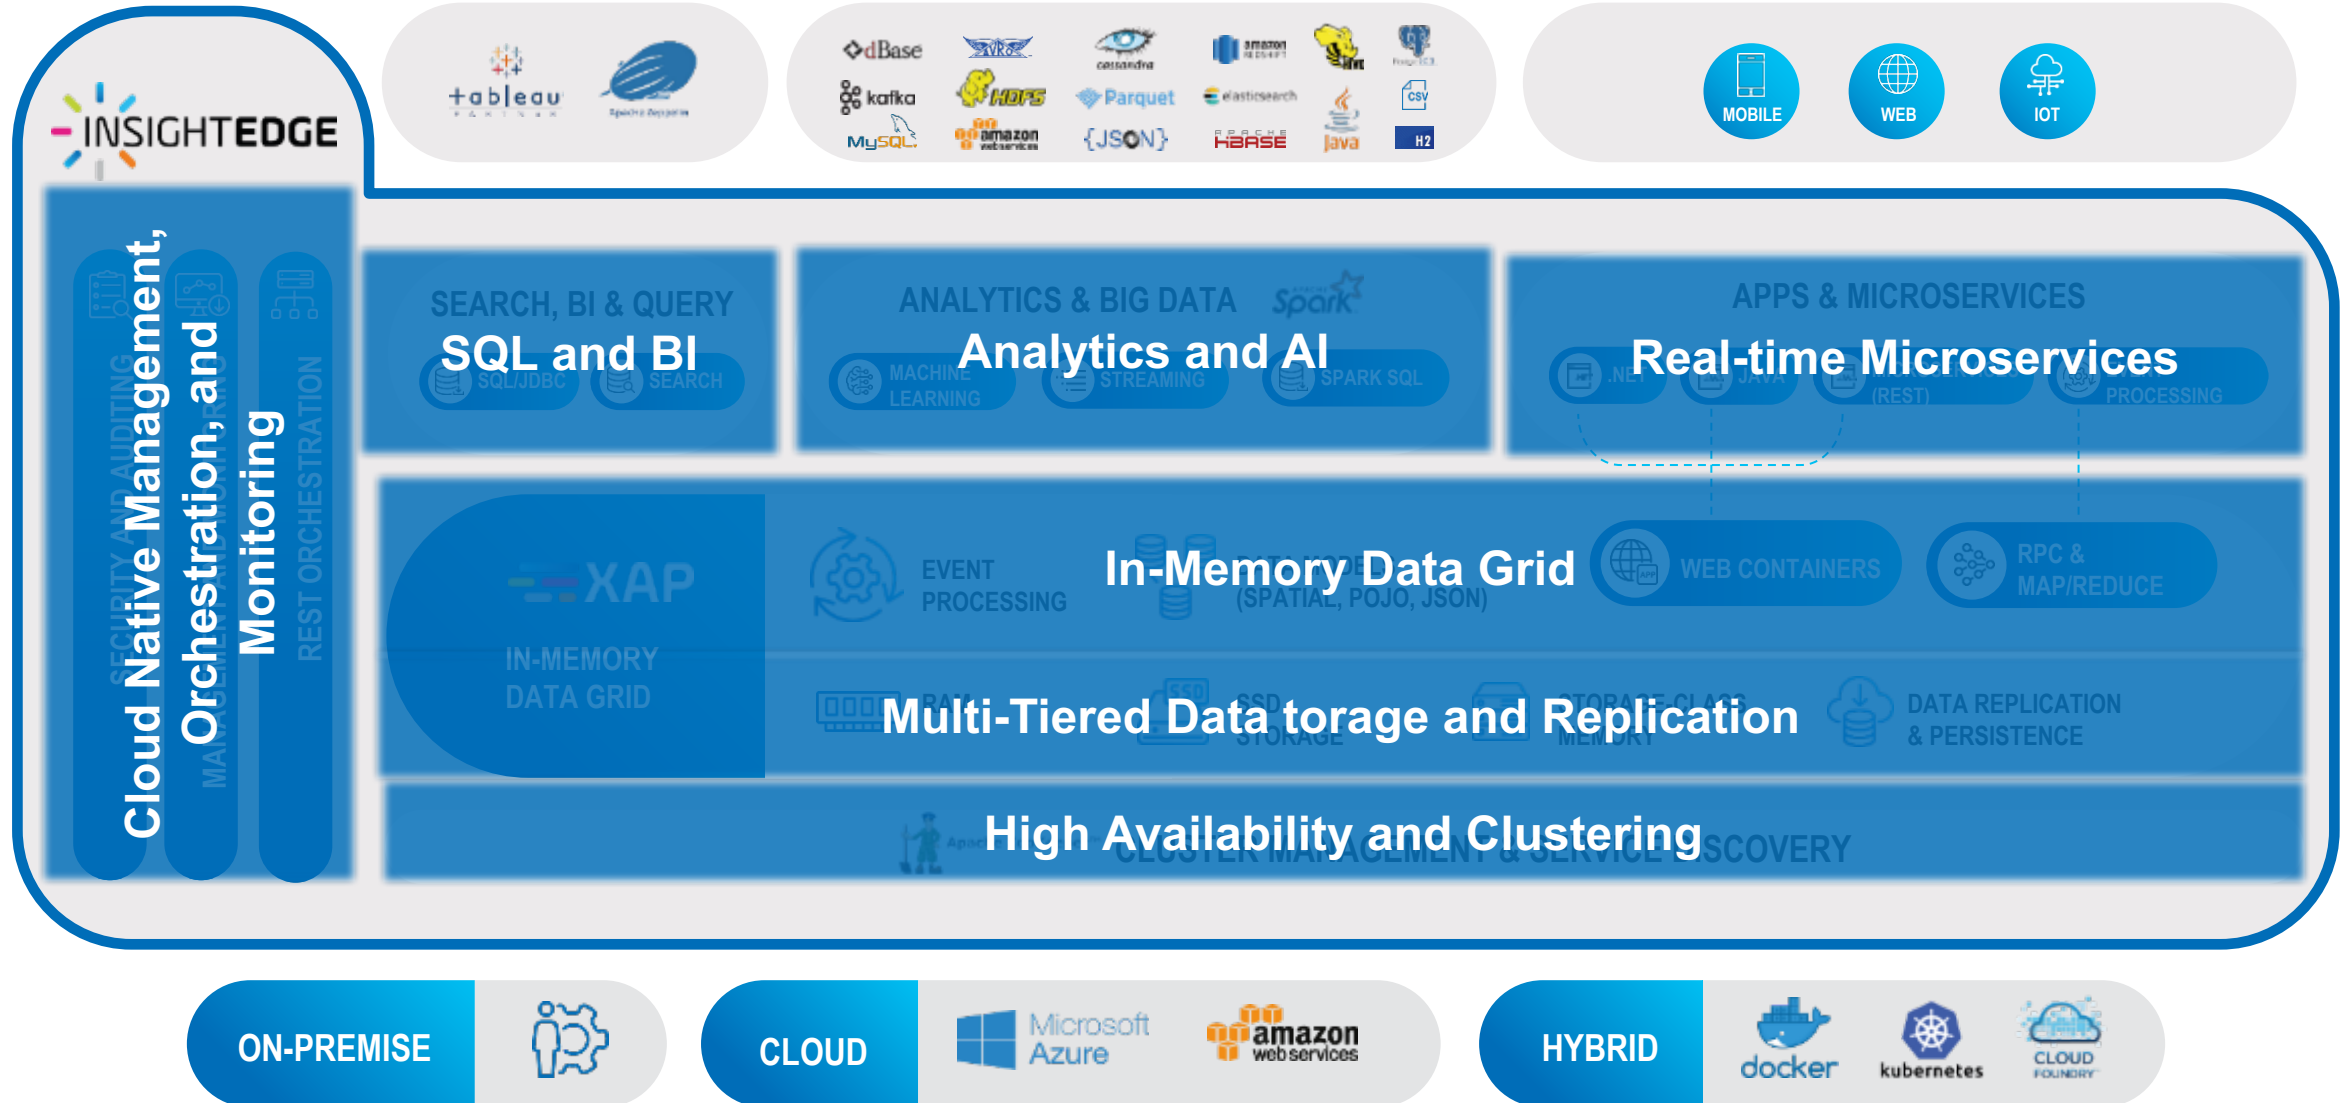


InsightEdge Architecture Overview



InsightEdge

Unifying Fast Data Analytics, AI and Transactional Processing





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

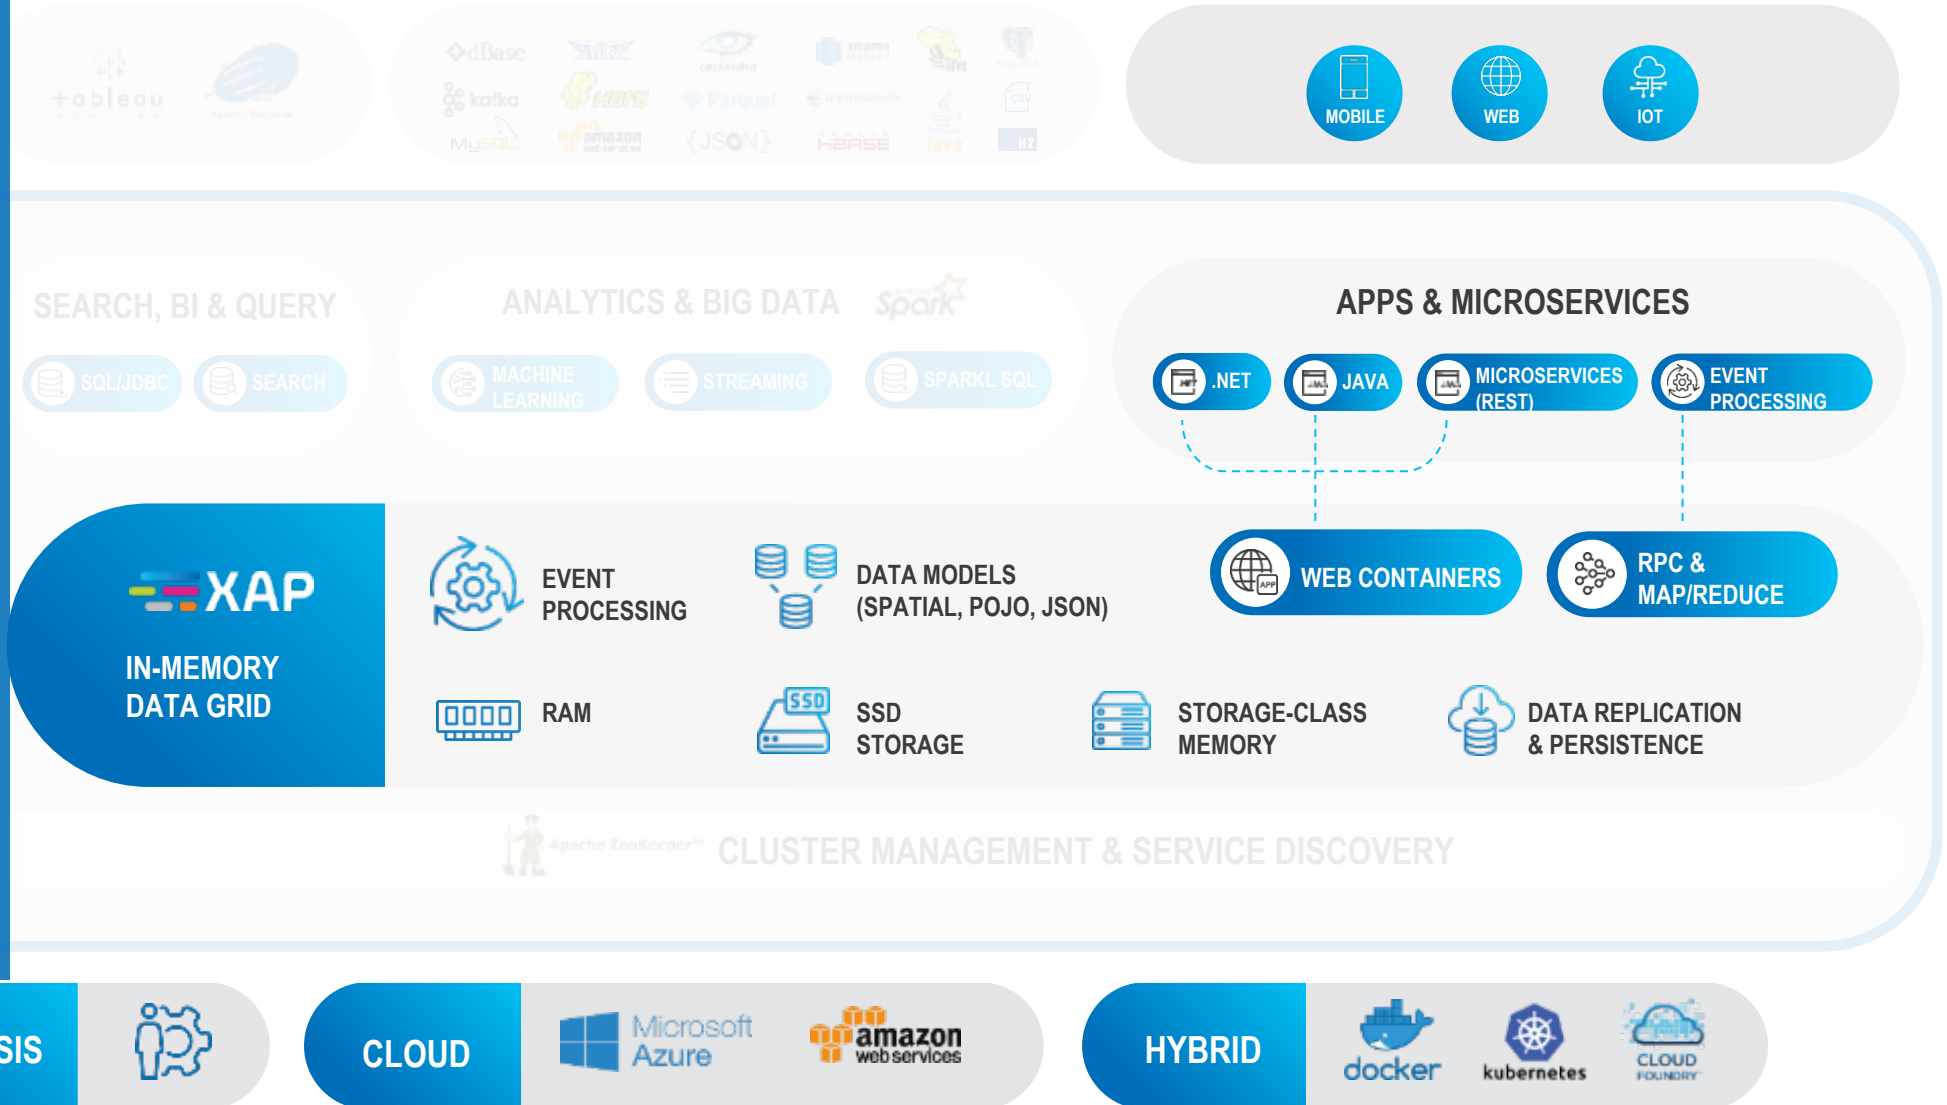
Auto-healing, multi-data center replication, fault tolerance

Fast Indexing Multi-Data Model

POJO, .NET, Document/JSON, GeoSpatial

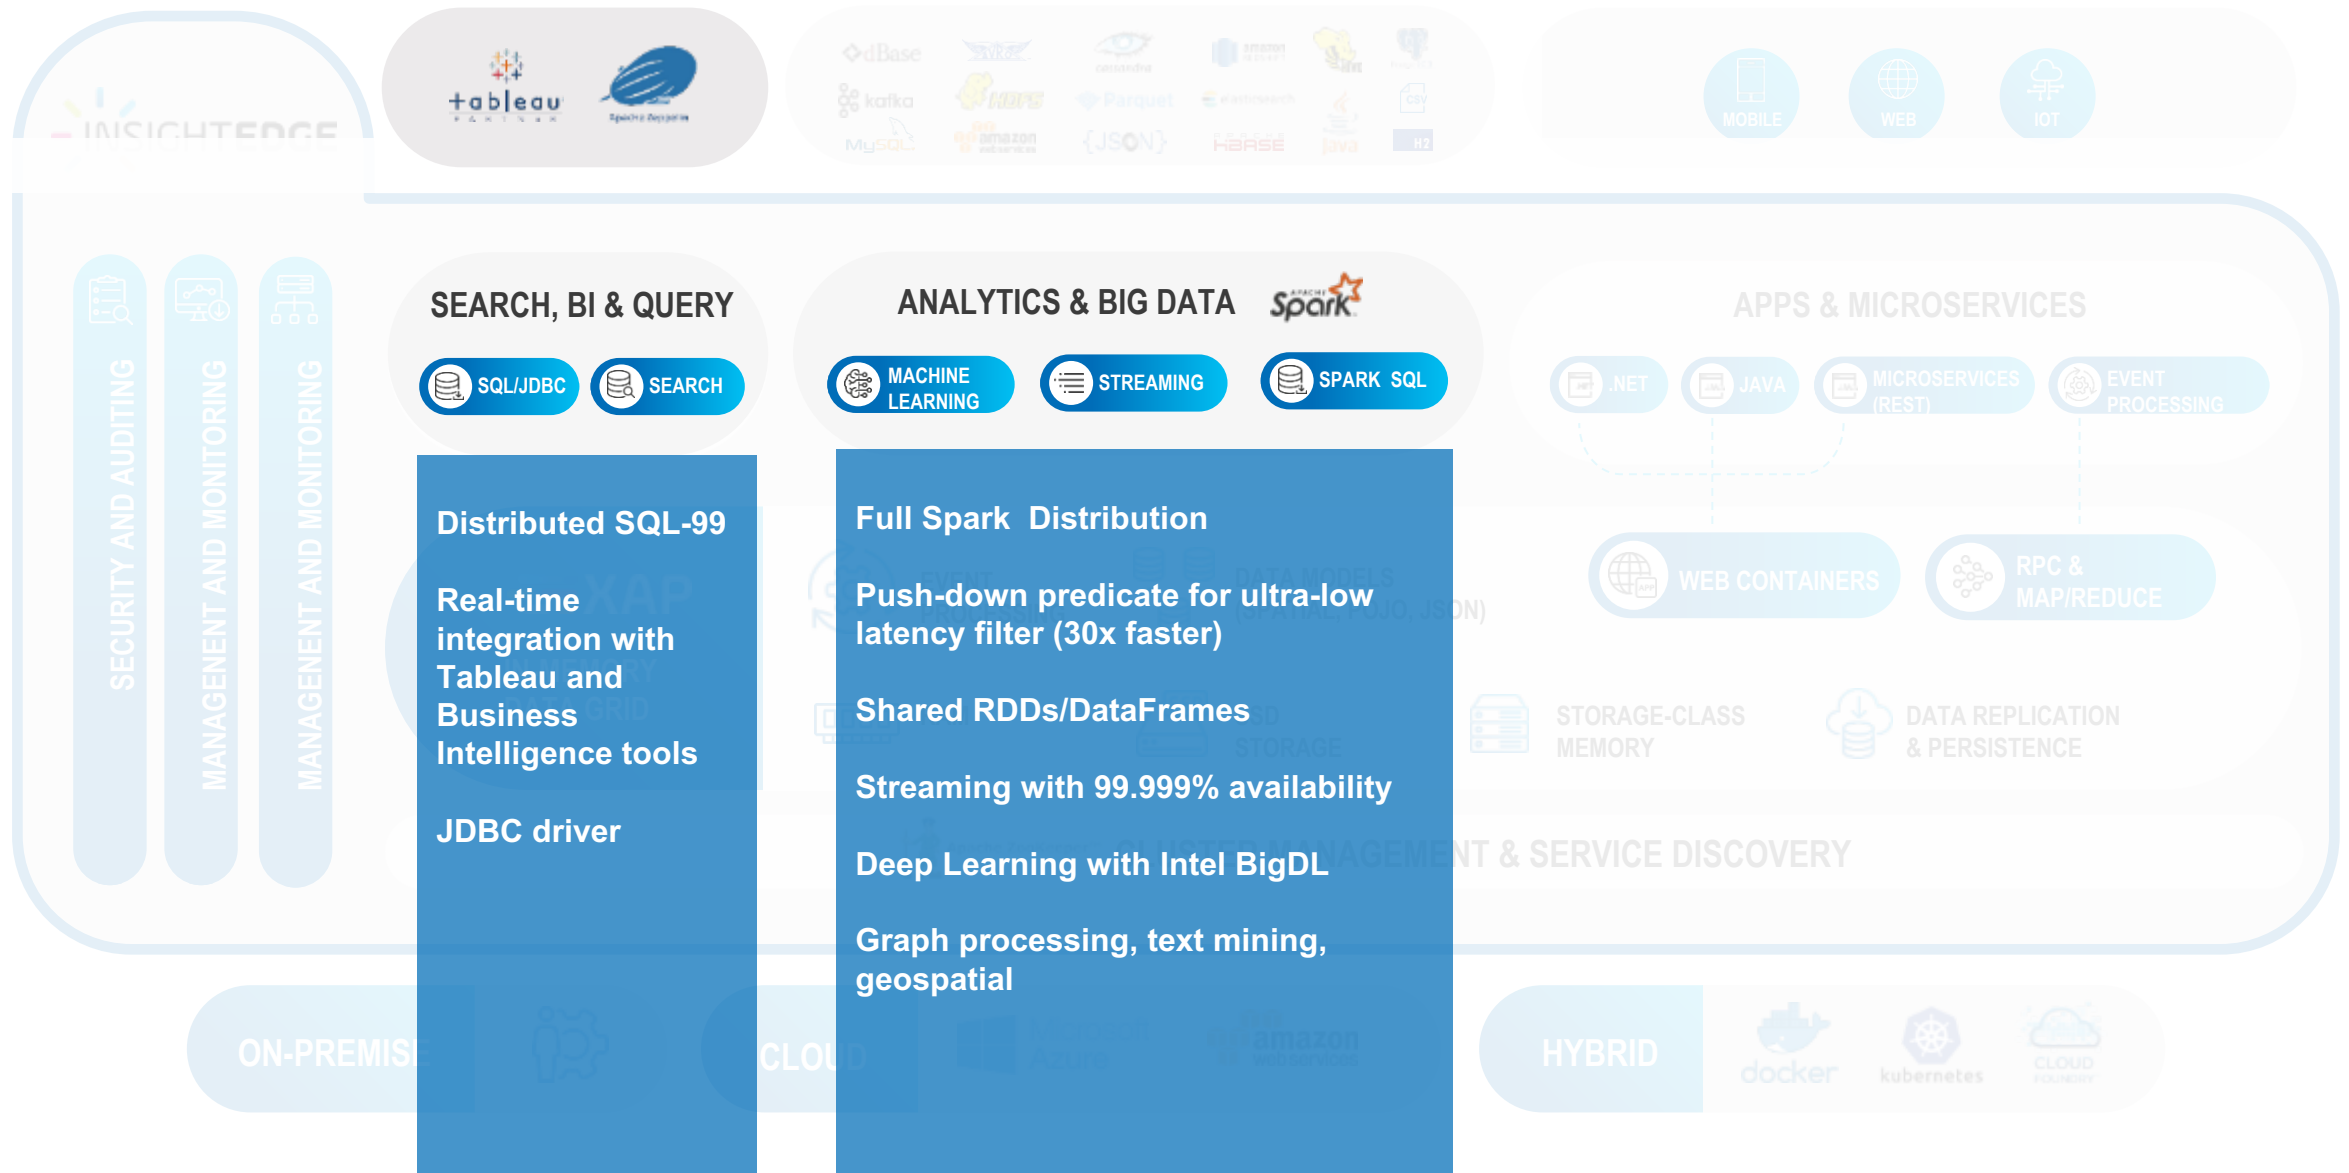
Seamless Integration with Java/Scala ecosystem

Cloud Native





Co-located Analytics and AI with Transactional Processing





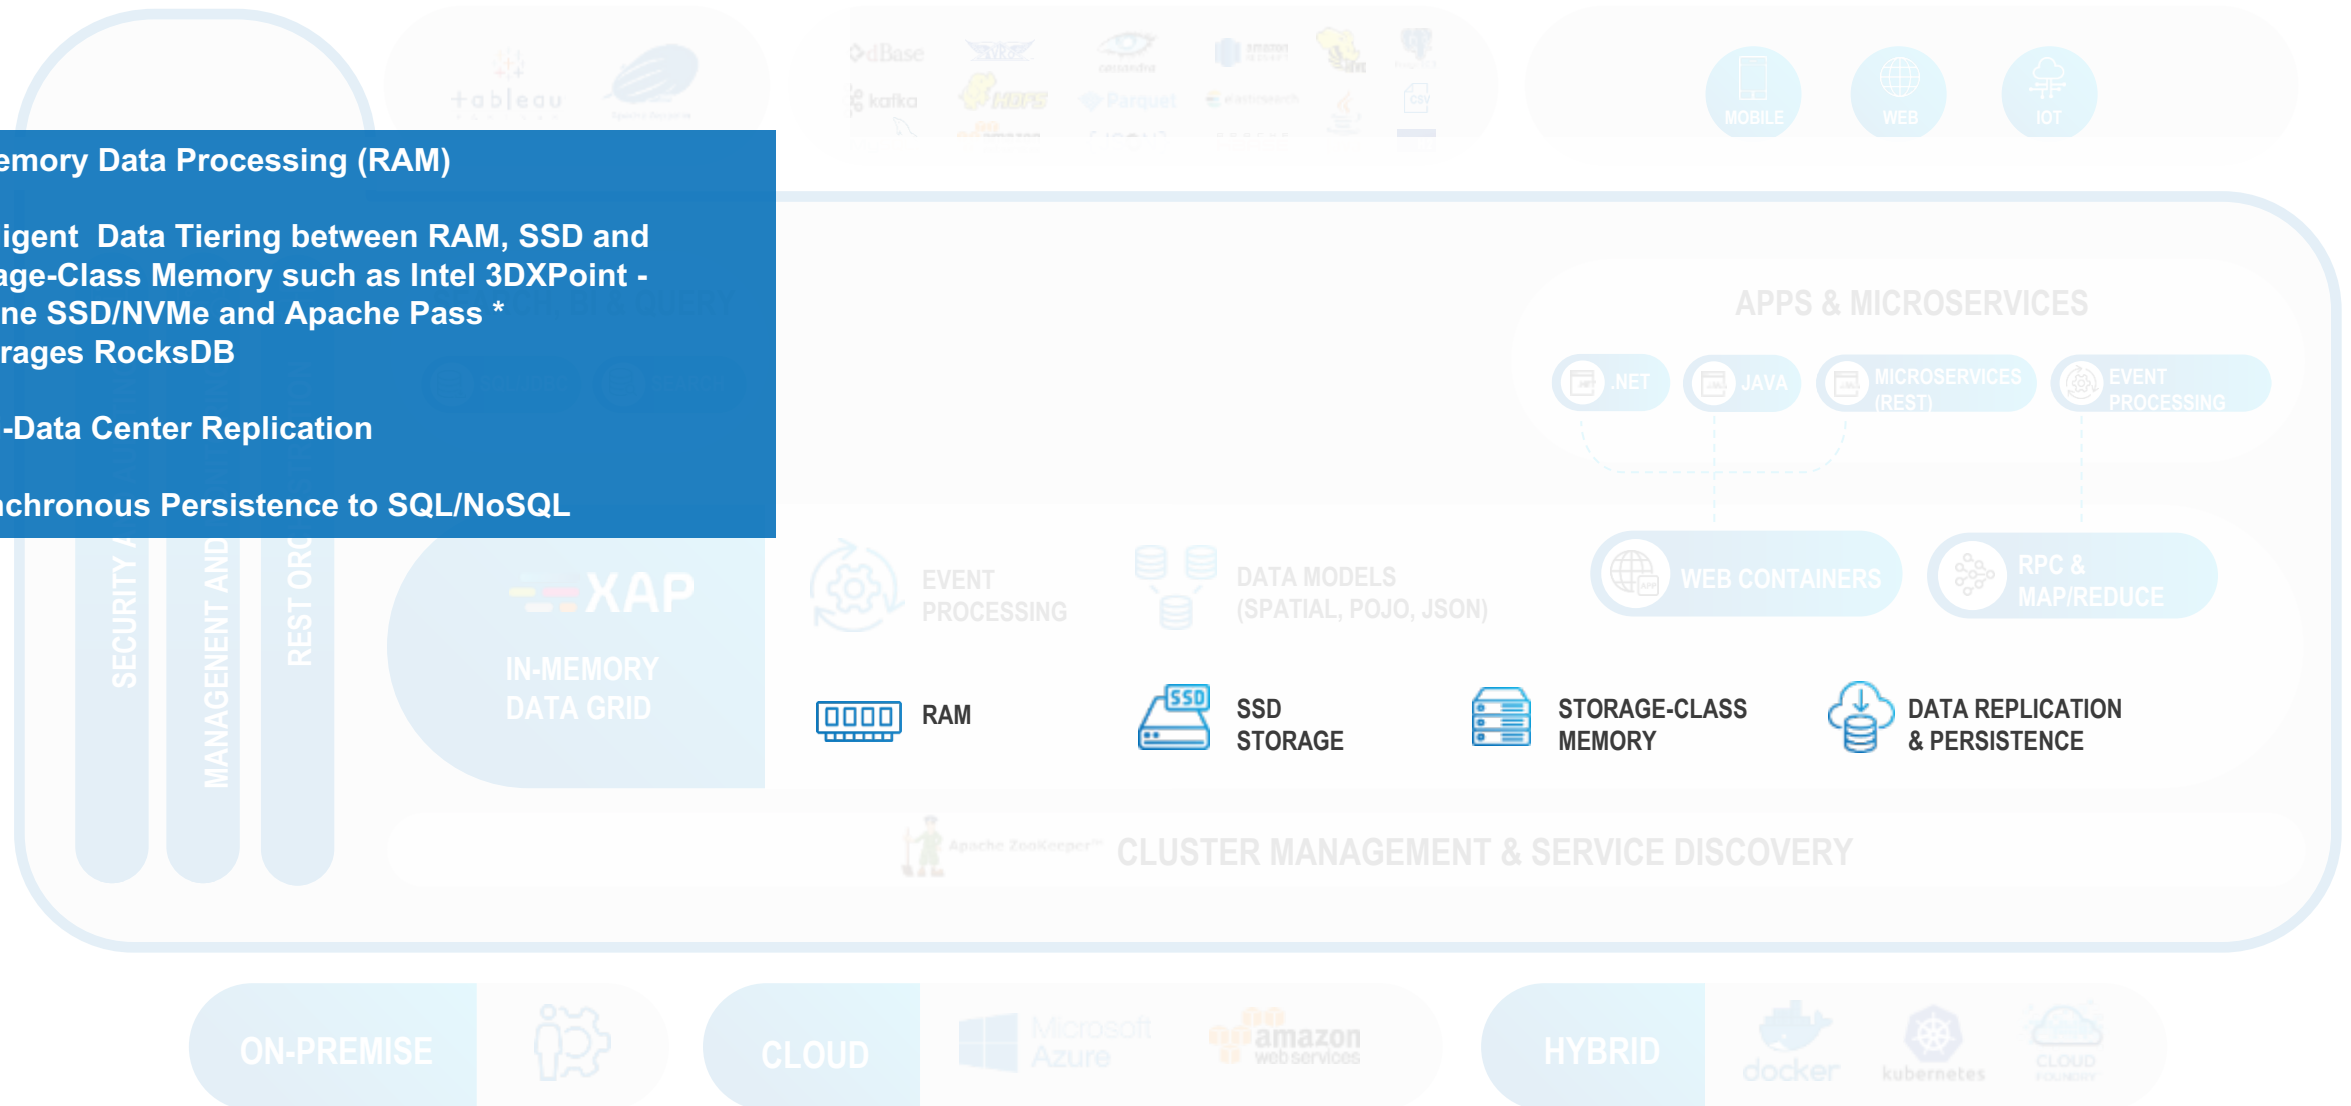
Multi-Tiered Data Storage and Replication

In-Memory Data Processing (RAM)

Intelligent Data Tiering between RAM, SSD and Storage-Class Memory such as Intel 3DXPPoint - Optane SSD/NVMe and Apache Pass *
Leverages RocksDB

Multi-Data Center Replication

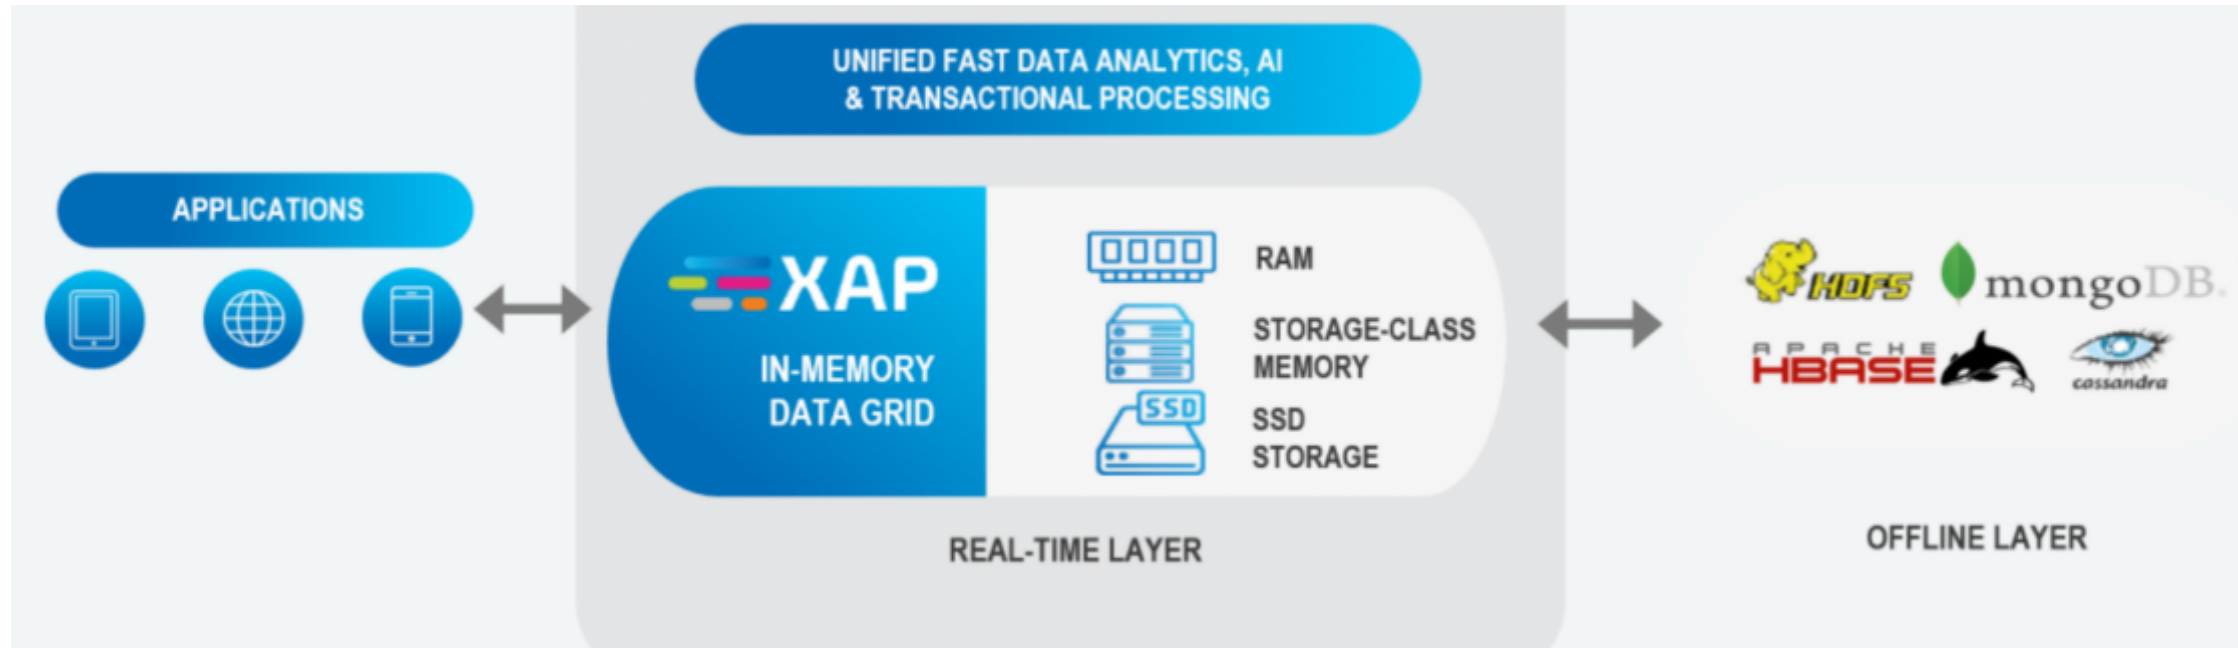
Asynchronous Persistence to SQL/NoSQL



* Apache Pass support in Q4 2018



MemoryXtend Module for Big Data – Optimized TCO



The Right Information in the Right Layer at the Right Time

Higher Performance – Optimized TCO

GigaSpaces InsightEdge Platform with Intel® Technology Accelerates AI Innovations

Latest Benchmark Shows Dramatic Performance Improvements when combining the power of InsightEdge Platform, Intel® Xeon® Scalable processors and Intel® Optane™ Solid State Drives

April 04, 2018 09:00 AM Eastern Daylight Time

NEW YORK--(BUSINESS WIRE)--GigaSpaces, a provider of in-memory computing (IMC) technologies that unify fast-data analytics, AI and real-time applications on a single software platform, today announced the results of a benchmarking study that spotlight improvements of the latest Intel® hardware breakthroughs. The study compared GigaSpaces InsightEdge Platform on Intel® Xeon® Scalable Processors and Intel® Optane™ SSDs against previous-generation Intel products.

Time to Load Data to the GigaSpaces InsightEdge Platform*
(Times shown in seconds. Lower is better.)



Figure 2. The comparison of the time to load data to GigaSpaces InsightEdge Platform* at 36.3 seconds with the Intel® Xeon® processor E5-2680 v4 and the Intel® SSD DC P3700 Series compared 24.3 seconds with the Intel Xeon Platinum 8168 processor and the Intel® Optane™ SSD DC P4800X Series

Time to Read/Query Data with the InsightEdge Platform*
(Times shown in seconds. Lower is better.)


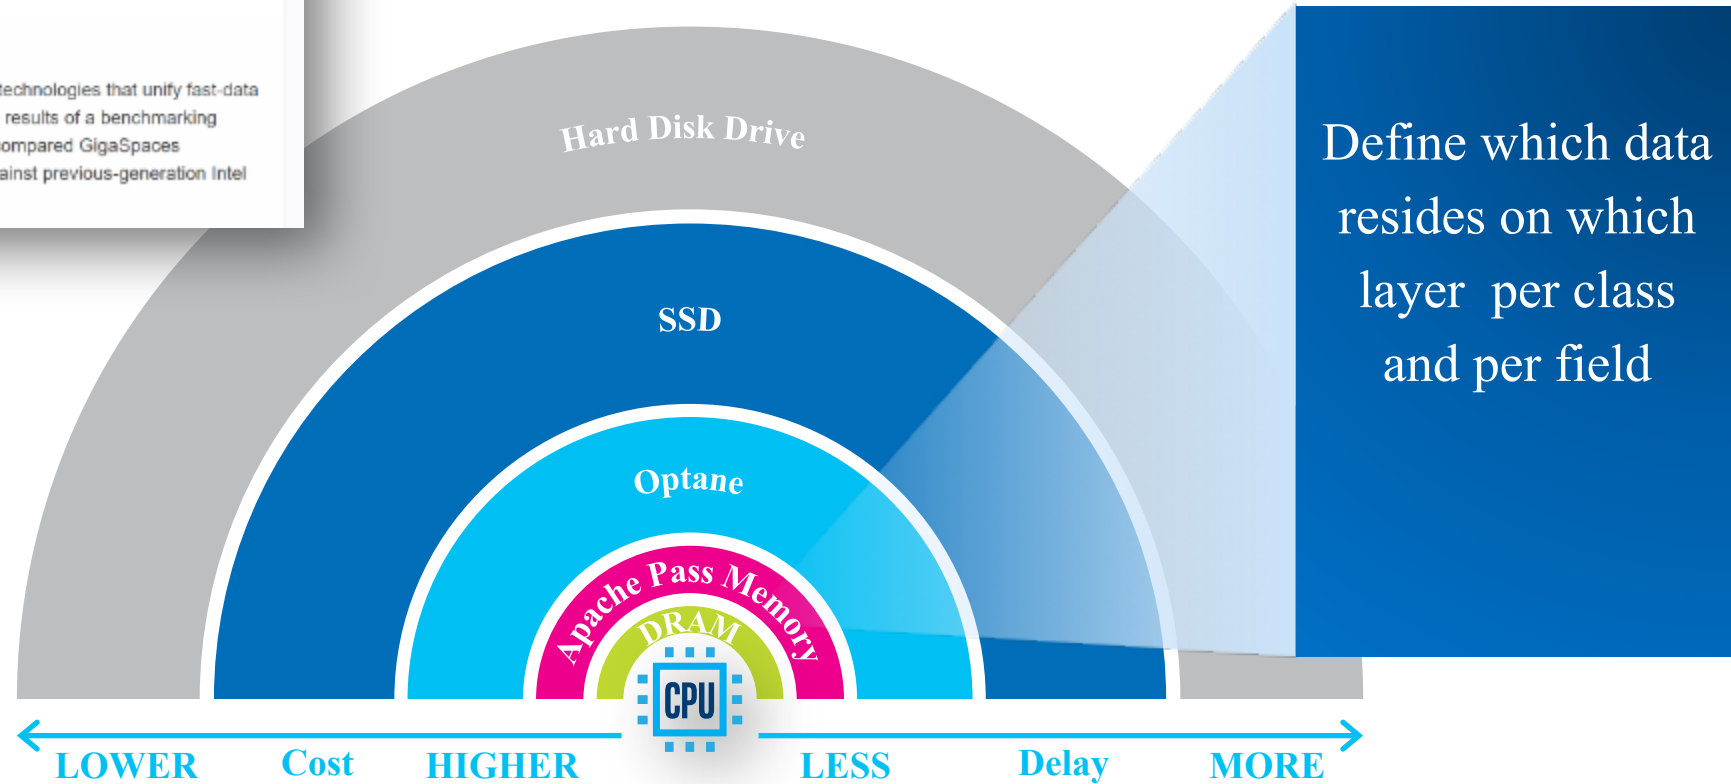


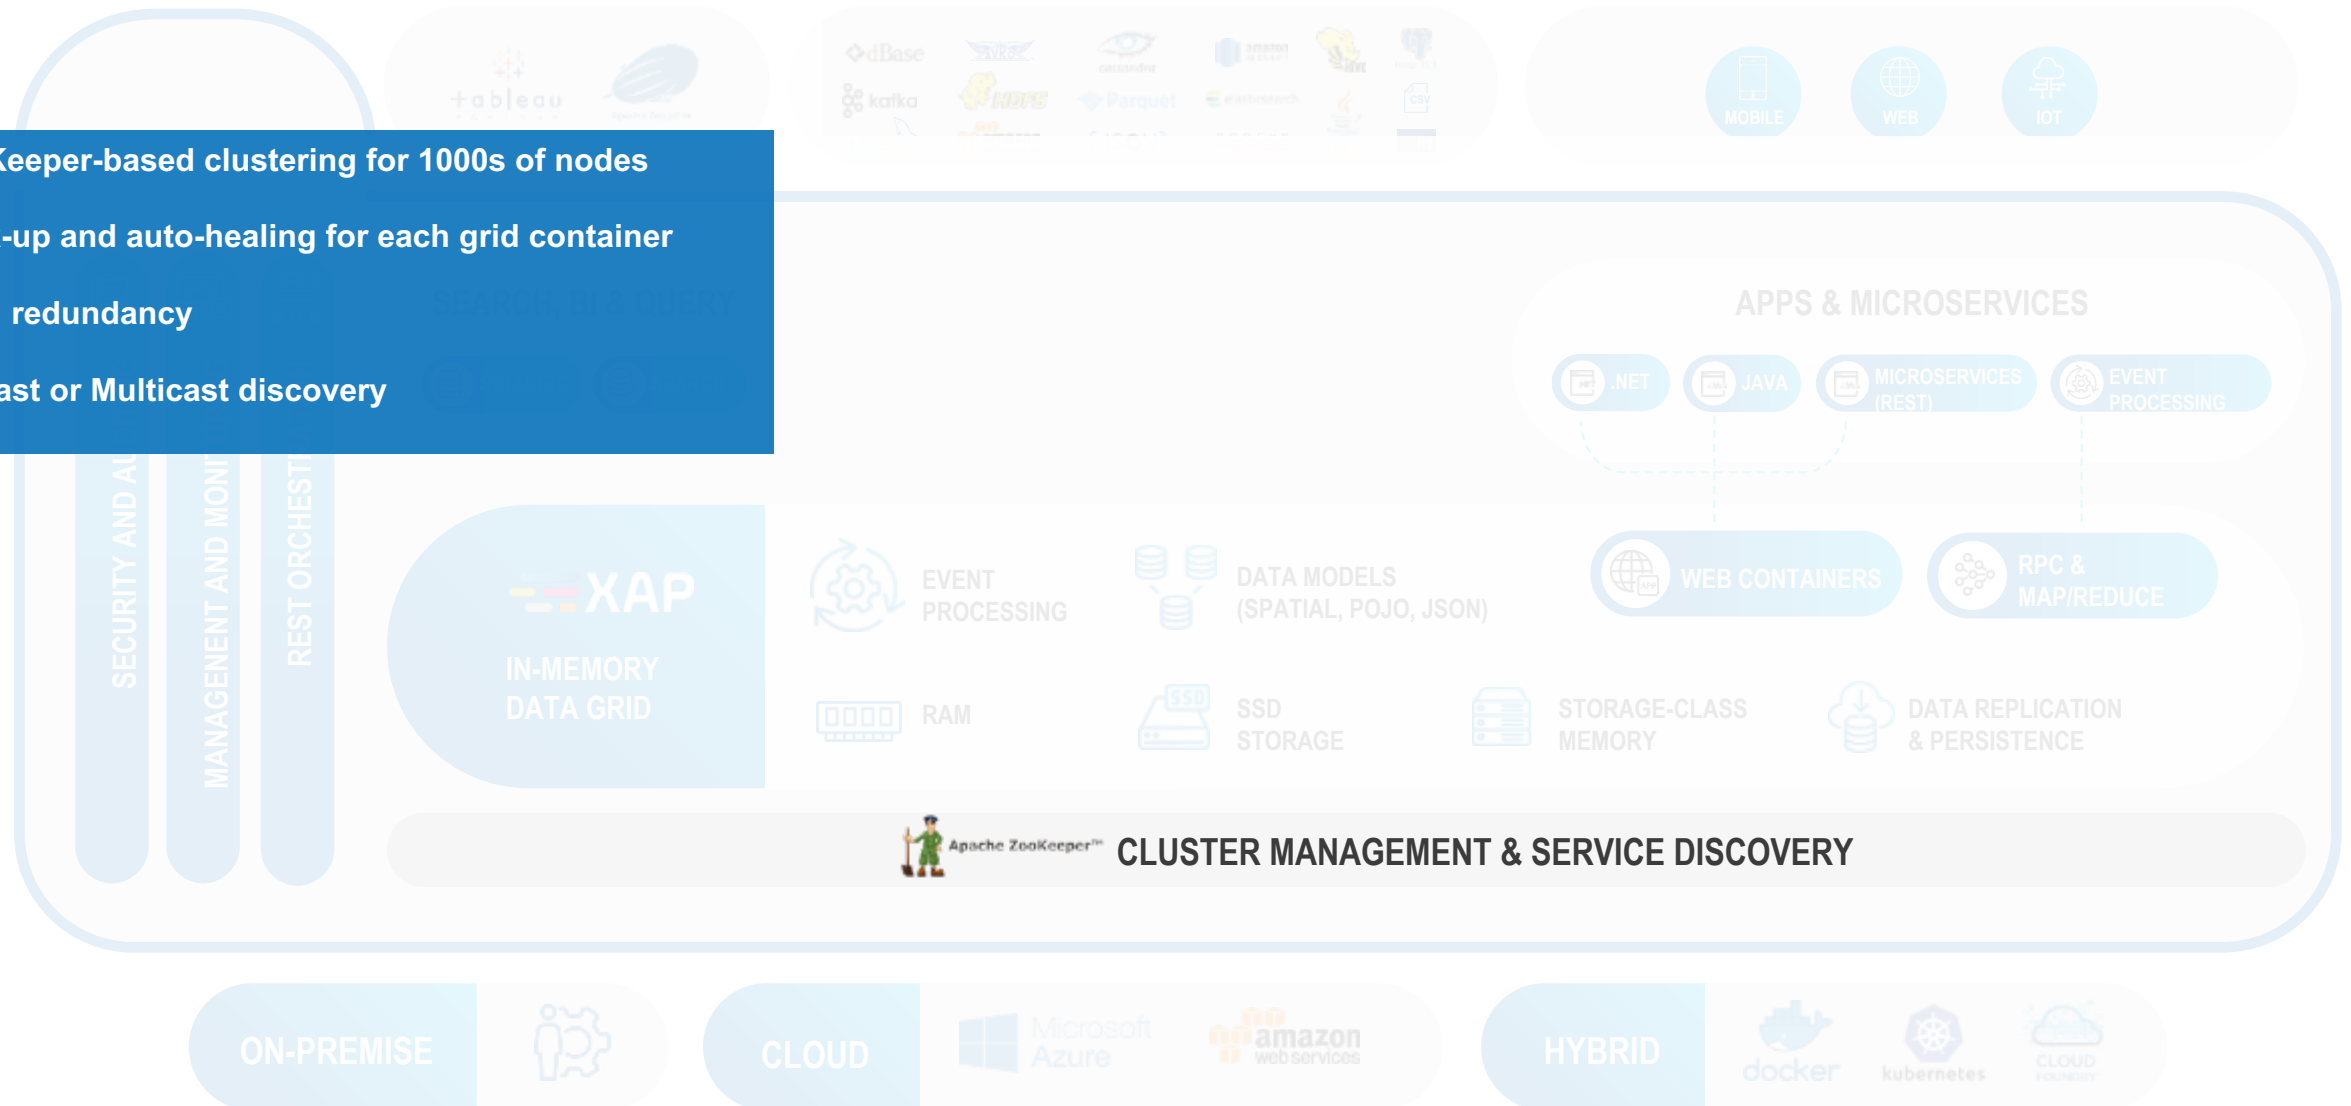
Figure 3. The comparison in time to query data from GigaSpaces InsightEdge Platform* at 9.4 seconds with the Intel® Xeon® processor E5-2680 v4 and the Intel® SSD DC P3700 Series compared to 7.2 seconds with the Intel Xeon Platinum 8168 processor with the Intel® Optane™ SSD DC P4800X Series





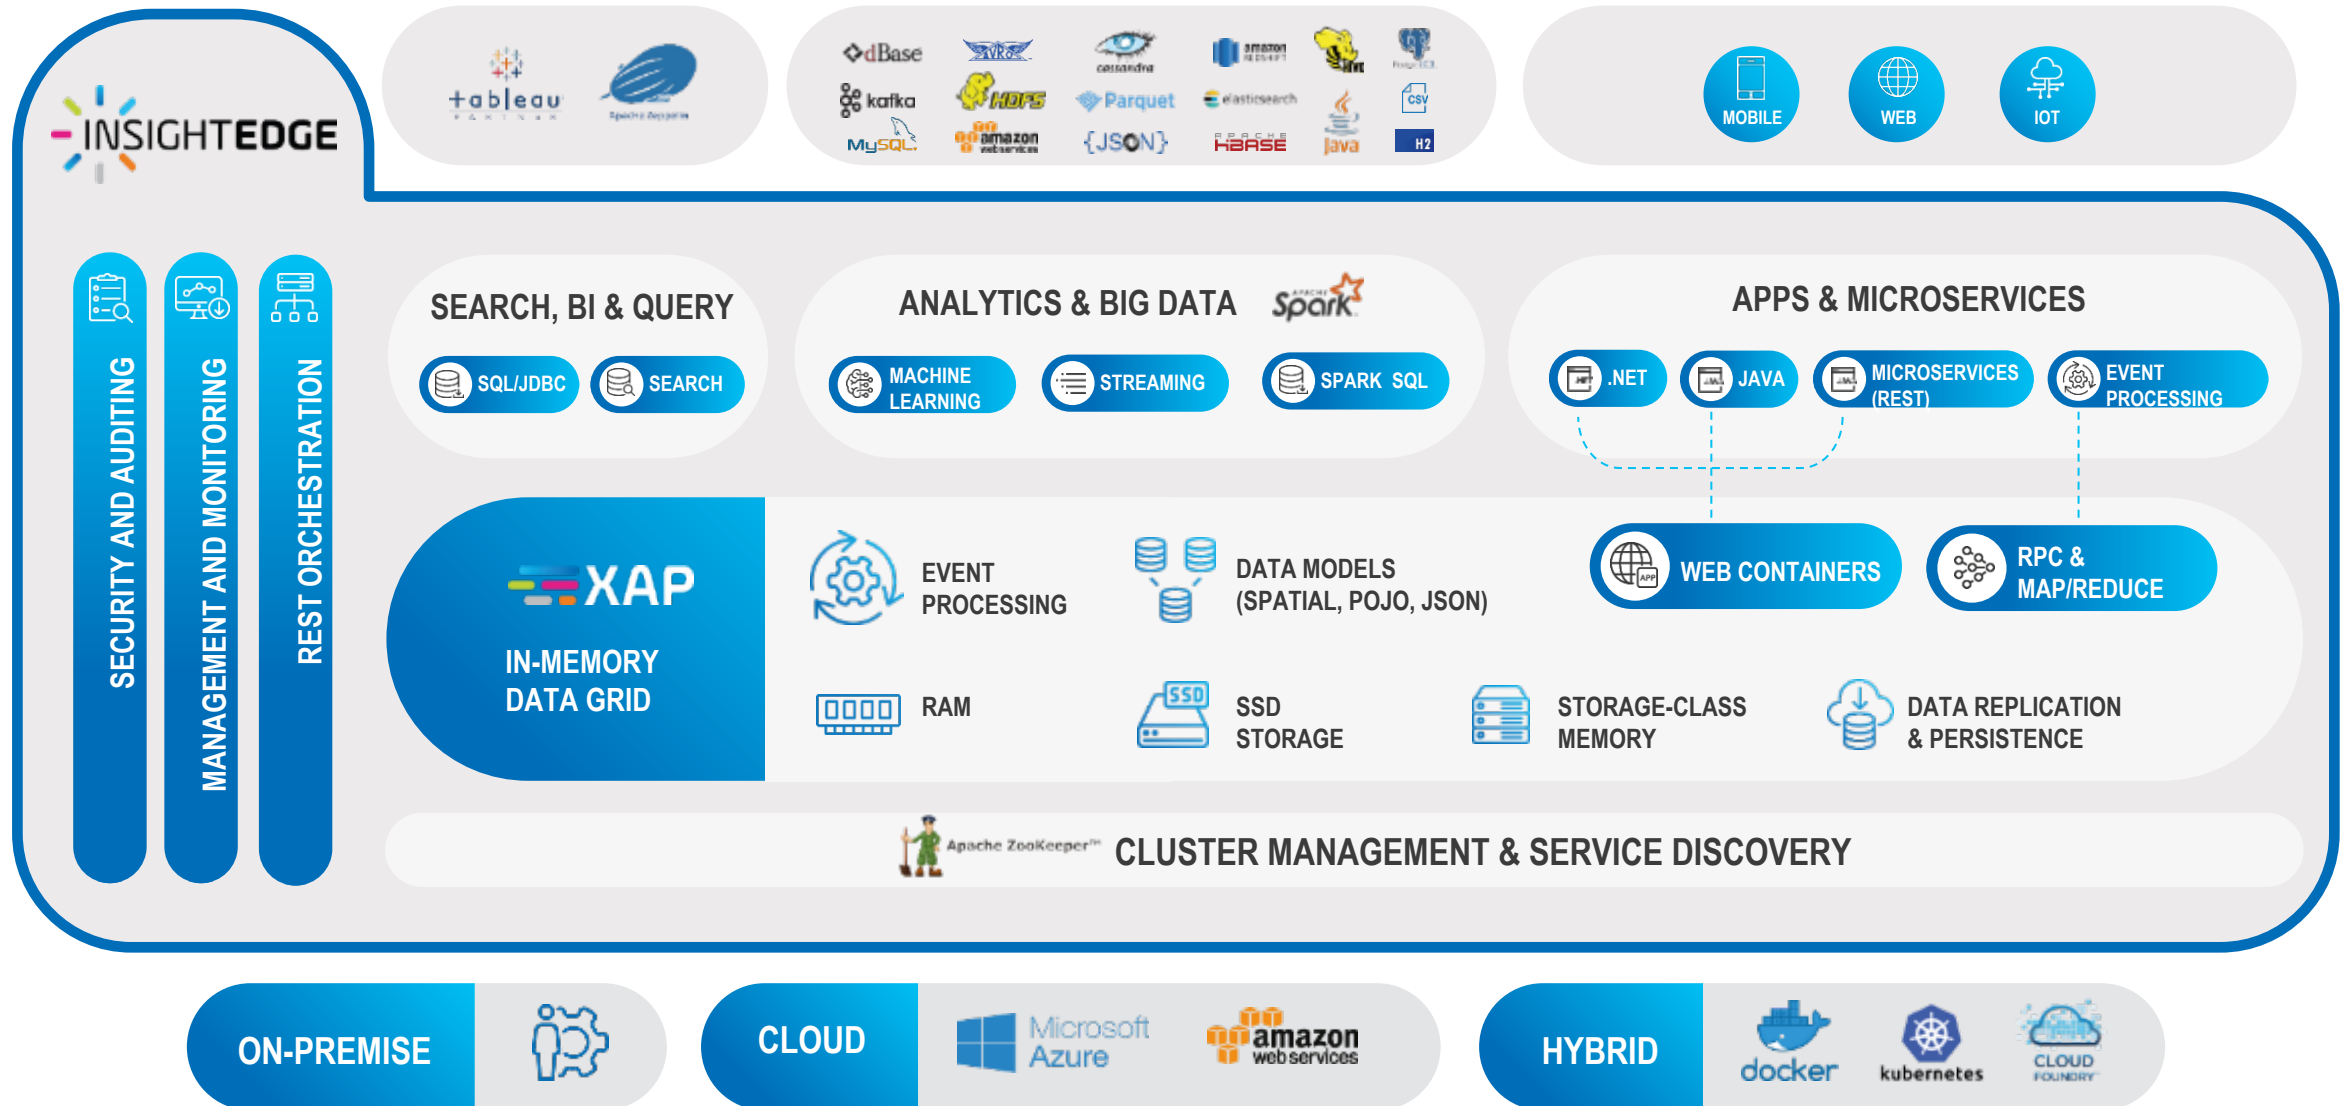
High Availability & Clustering

- ZooKeeper-based clustering for 1000s of nodes
- Back-up and auto-healing for each grid container
- N + 1 redundancy
- Unicast or Multicast discovery



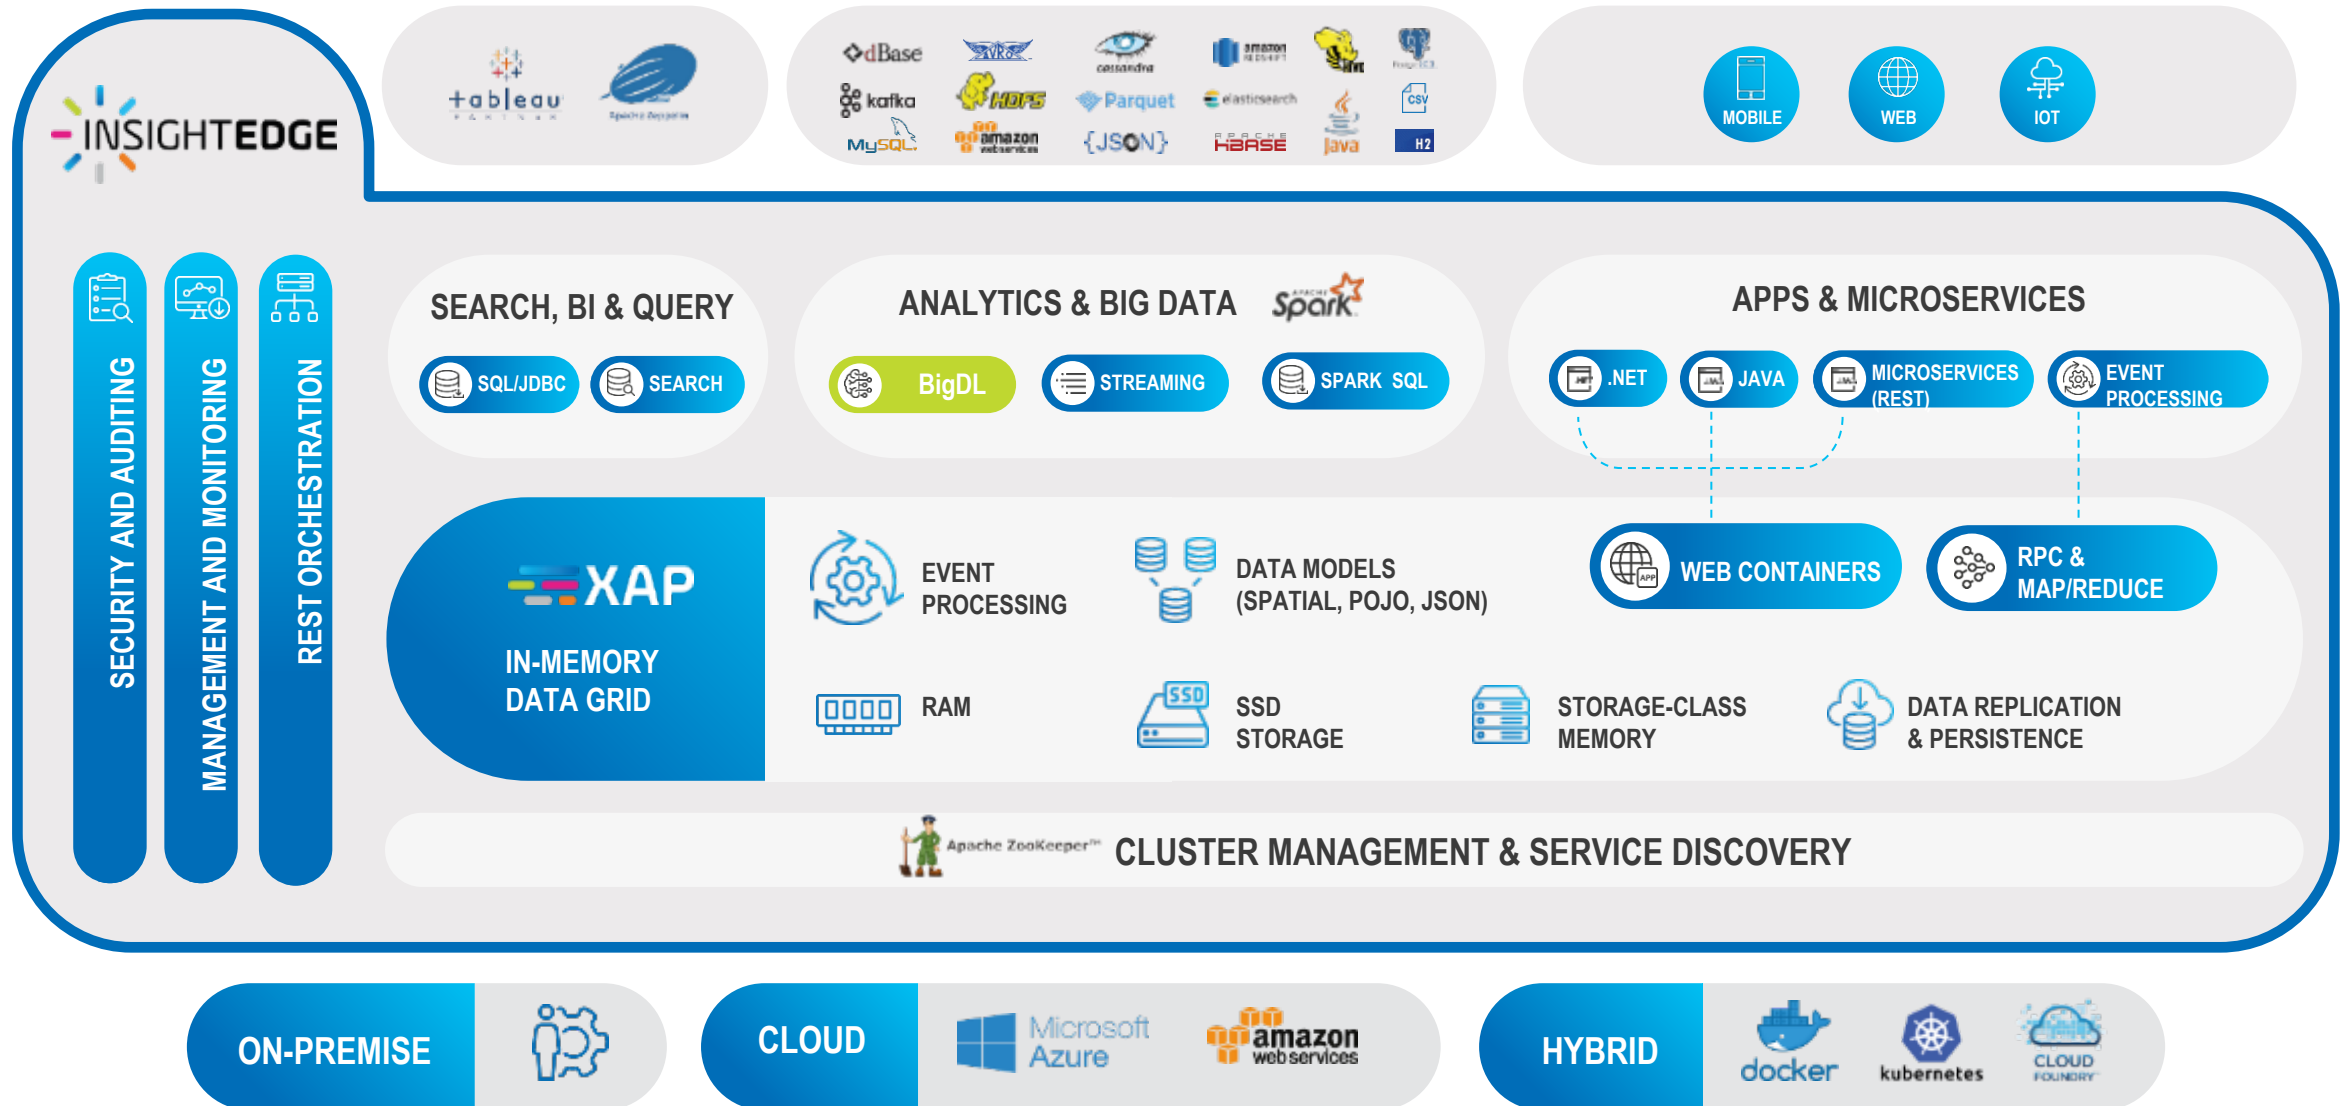


InsightEdge Architecture Overview





InsightEdge Architecture Overview





BigDL

**Distributed deep learning
framework designed for
Apache Spark and Intel Xeon**



**GigaSpaces
InsightEdge Platform**

**In-memory insight platform
unifying fast-data analytics,
artificial intelligence and real-time
applications**

Intel® Xeon® Scalable Processor



Apache Spark

**Open source big data analytics
framework built around speed,
ease of use, and sophisticated
analytics**

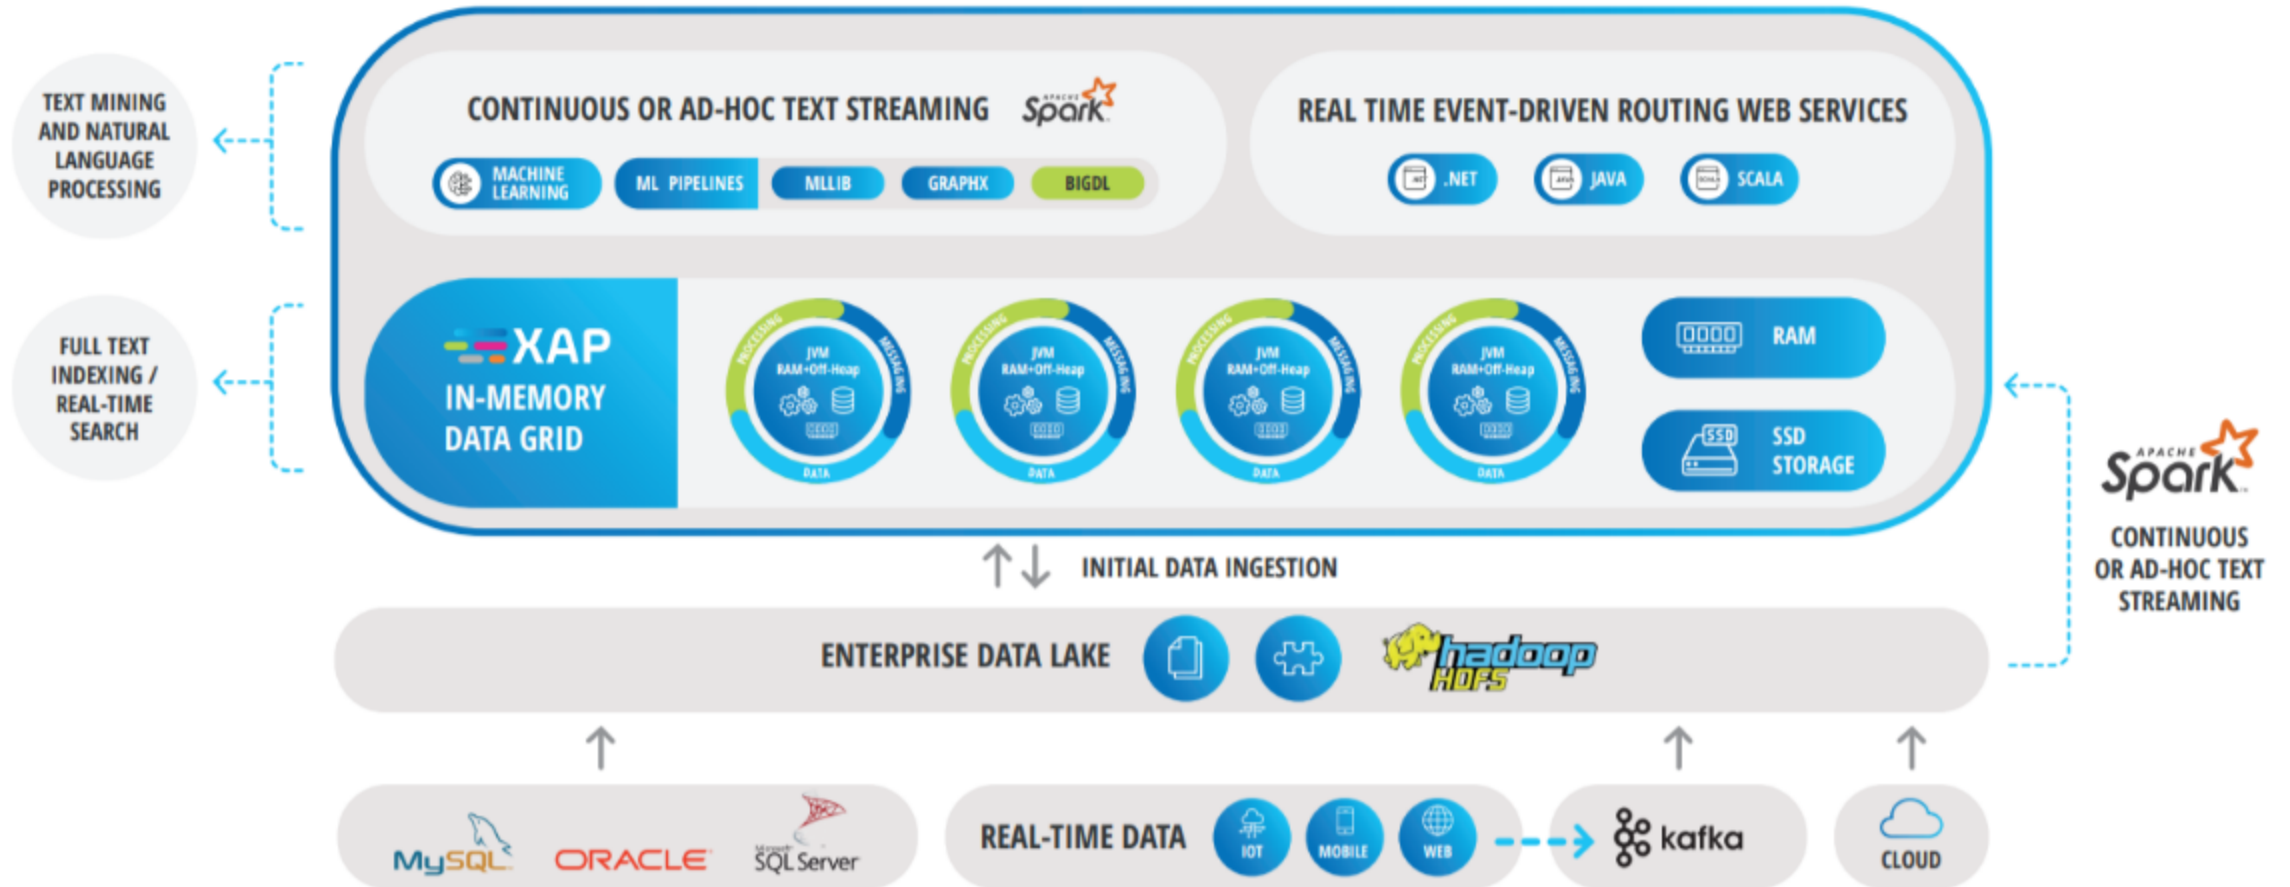


Fast Data Analytics Requires Real-Time Data Processing

- Unifies analytics, AI and real-time applications
- Reduces cluster and component sprawl
- Triggers transactional workflows based on prediction criteria and scoring
- Efficient scale-out computing
- Distributed model training
- Lowers TCO/Decreases Deployment Costs – train and run large-scale deep learning workloads on Intel Xeon Scalable processors without relying on GPUs
- High performance – optimized with Intel® Math Kernel Library (Intel® MKL)



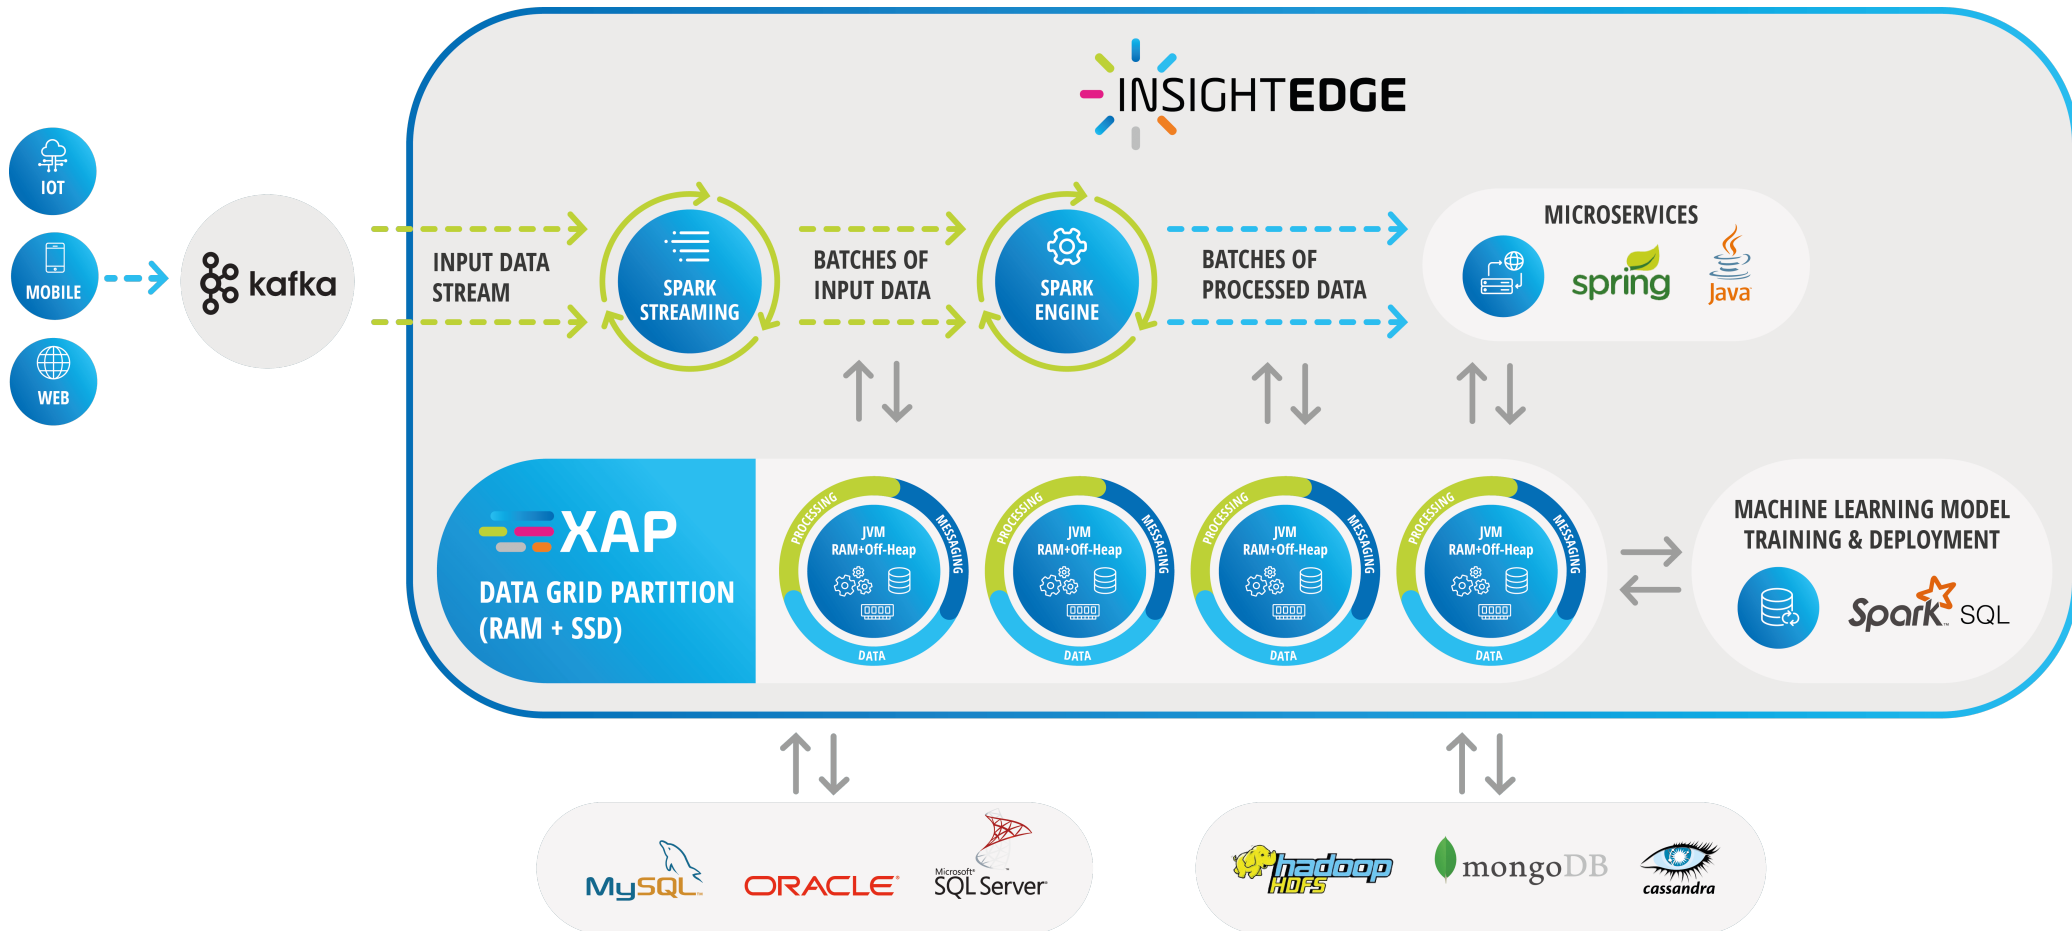
Demo Architecture





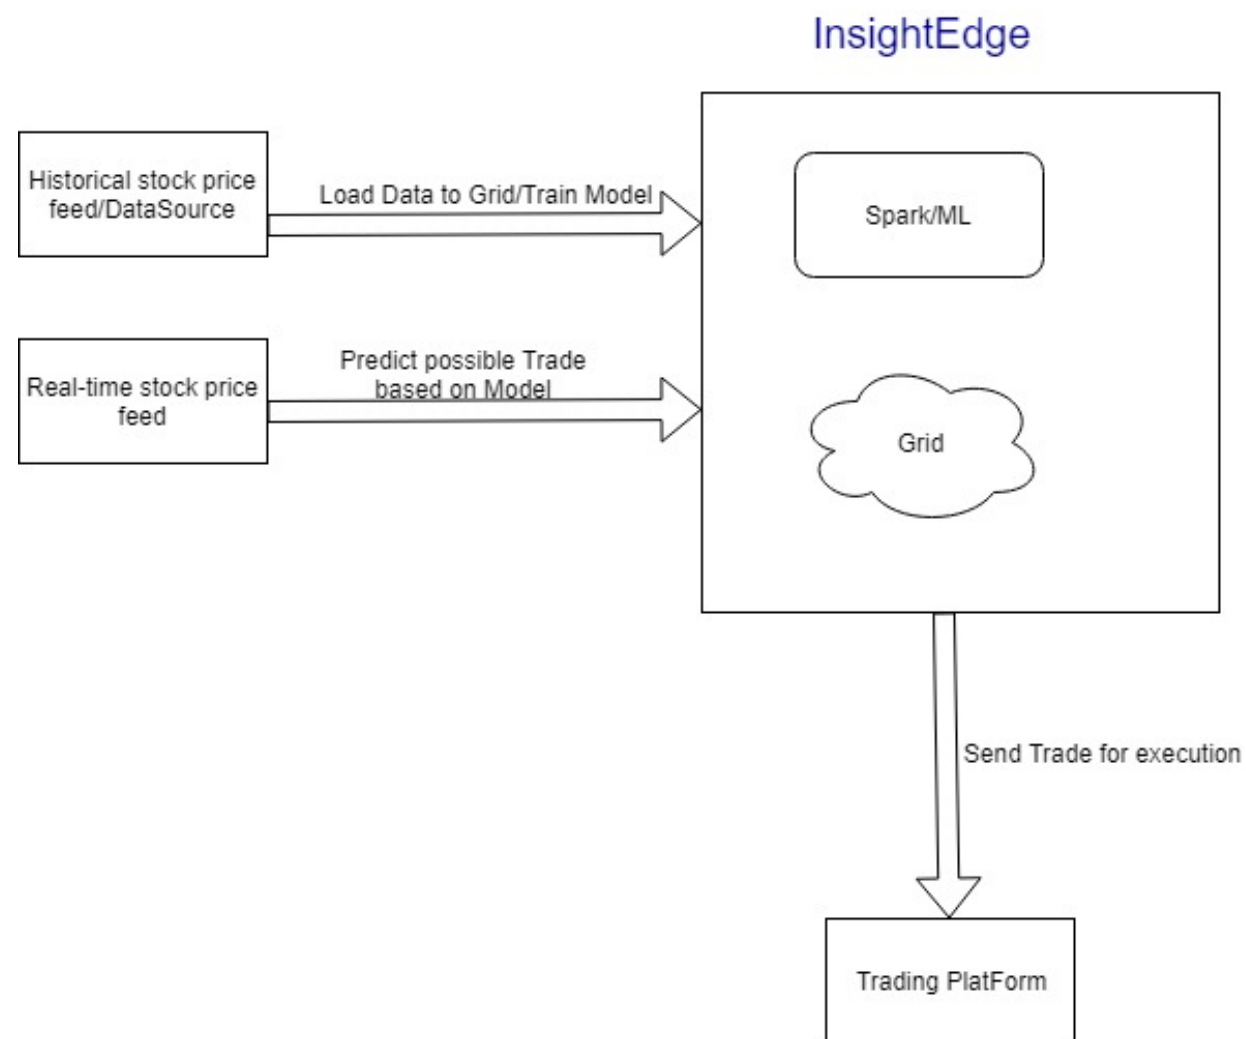
Streaming, Machine Learning and AI

Solution Architecture



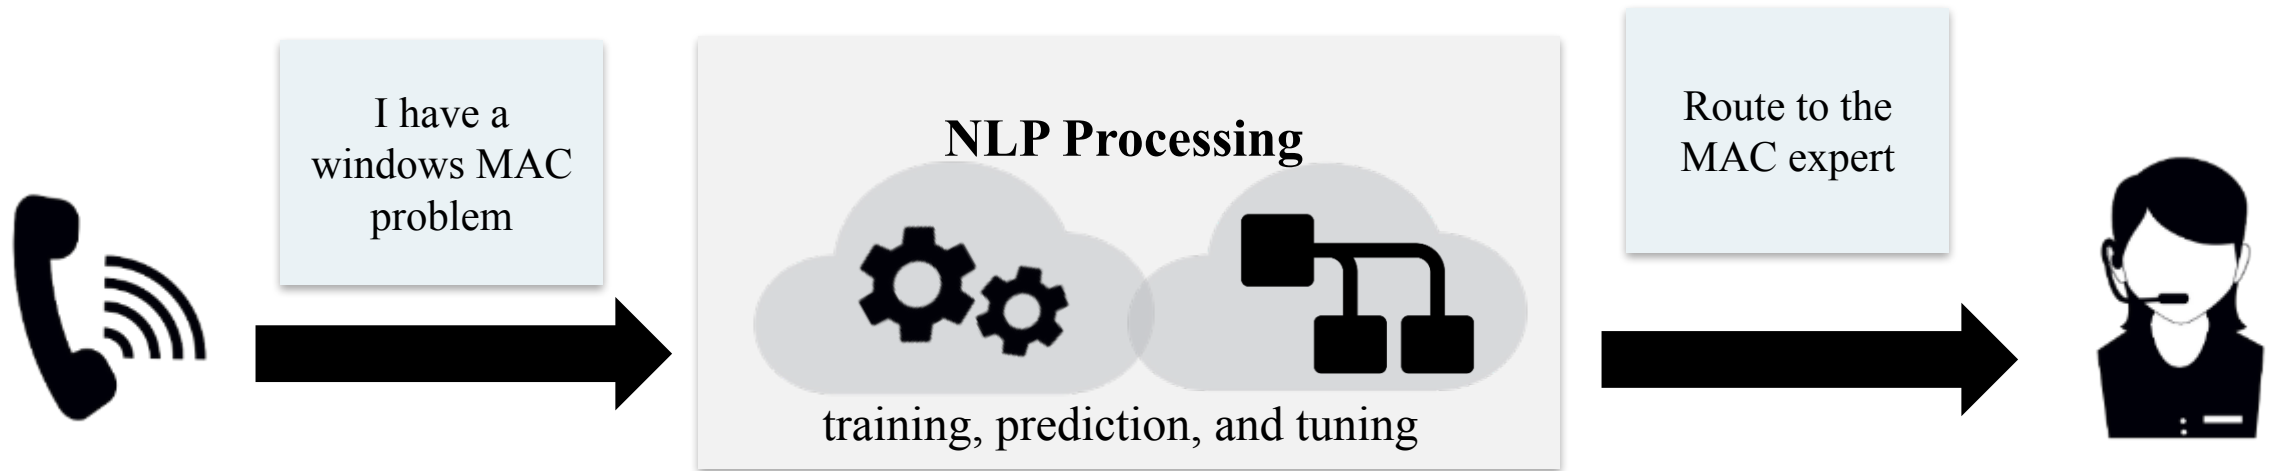


Automated stock trading using InsightEdge and Machine learning





Stop Pressing 0 Or *



User speaks
using web
interface

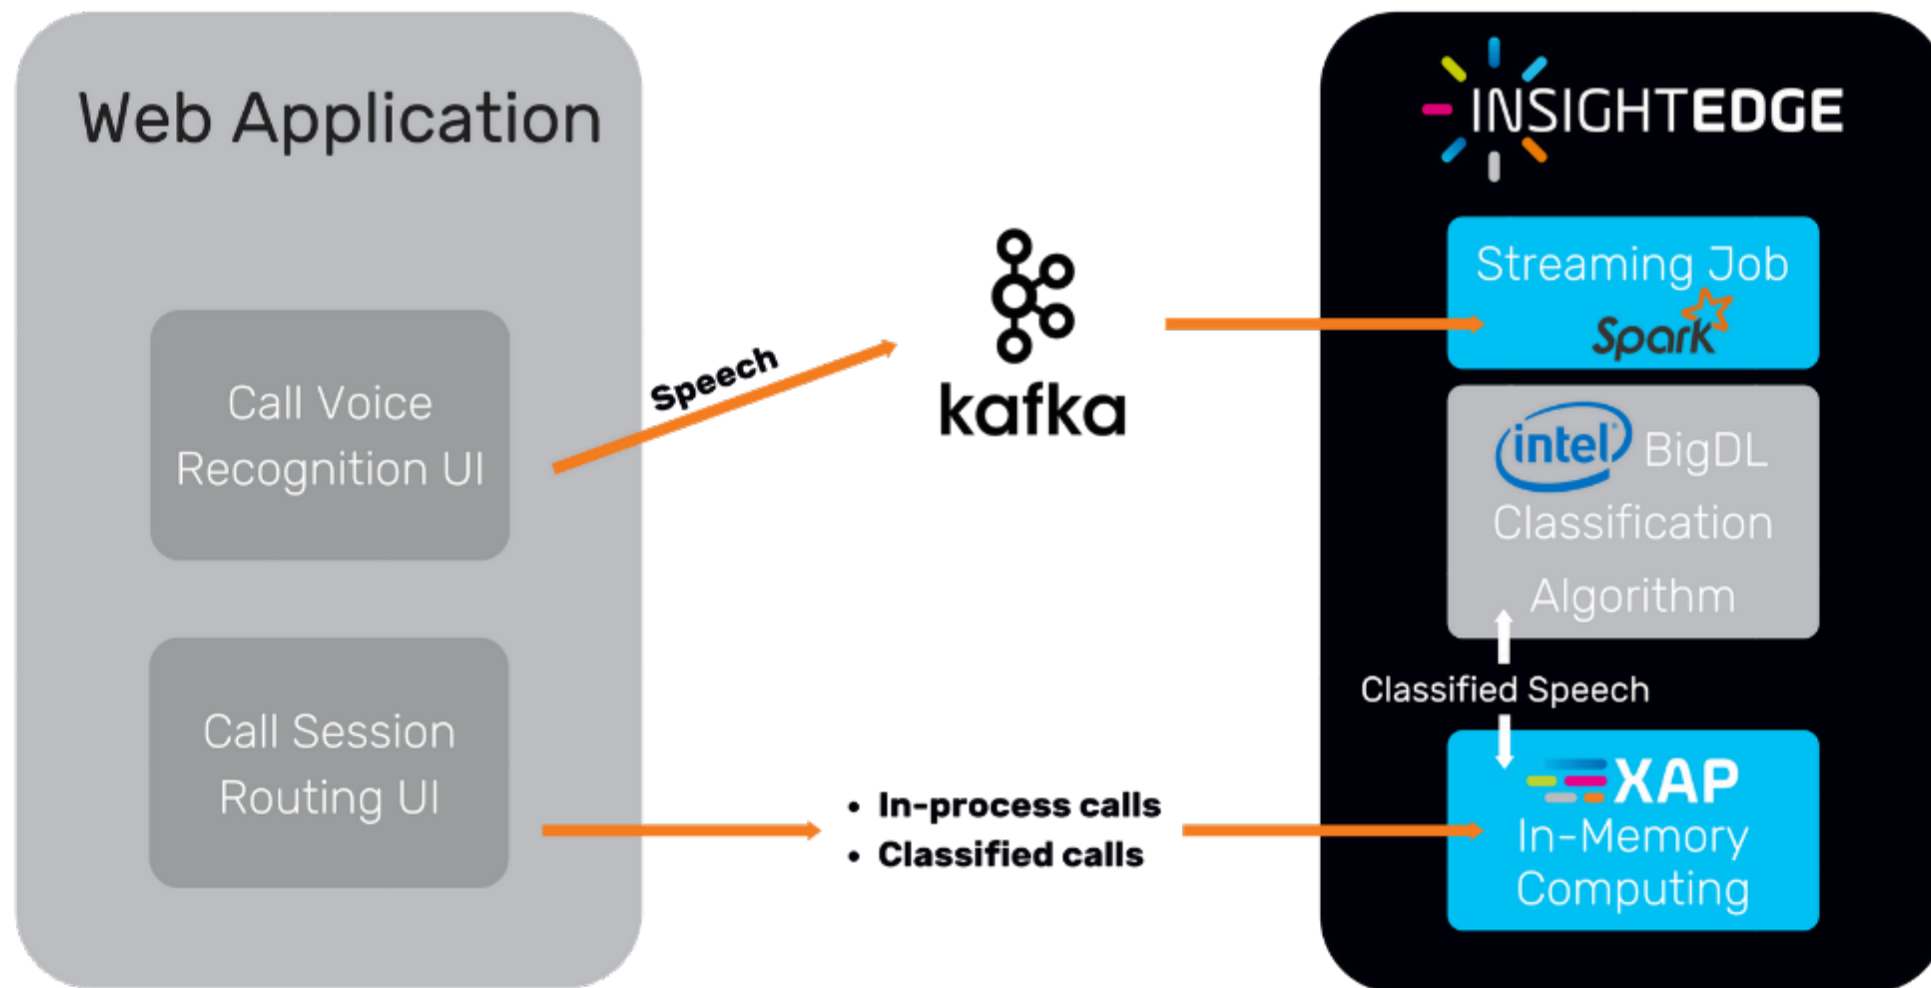
Browser converts
speech to text and
sends to controller

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





Innovate with Confidence



INSTANT INSIGHTS



EXTREME PERFORMANCE



TCO REDUCTION



TOTAL CONFIDENCE





Try it



<http://insightedge.io>

Learn it



<http://software.intel.com/bigdl>

Build it



<http://github.com/InsightEdge>



THANK YOU



GIGASPACE

innovate with confidence