



GridGain In-Memory Data Fabric:

Driving High Frequency Trading and Compliance with In-Memory Computing

ERIC KARPMAN
DIRECTOR, BUSINESS DEVELOPMENT
FINANCIAL SERVICES
ekarpman@gridgain.com

www.gridgain.com



[#gridgain](https://twitter.com/gridgain)



High-Frequency Trading

- Getting Information
- Processing Information
- Execution
- Transaction Cost Analysis and Backtesting

HFT Strategies

- Market Making
- Market-Taker
- Arbitrage
- Pair Trading
- Momentum Trading
- Pinging
- News and Sentiment Trading

HFT Technology

- Dark Fiber Cabling
- Exchange Colocation
- FPGA/GPU
- Hadoop with MapReduce
- Complex Event Processing
- Parallel Processing Clusters
- In-Memory Computing

Evolution of In-Memory Grid Computing

- Move from Disk to 100% In-Memory (RAM)
- Leverage Clustered Memory and Parallel Distributed Processing
- Results: 100x Faster, 10x ROI Improvement
- Making “Big Data” Fast

“In-memory will have an industry impact comparable to web and cloud.”

“RAM is the new disk, and disk is the new tape.”

Gartner

In-Memory Computing Market:

- \$10B in 2019
- CAGR 22%

Gartner

Financial Customer Use Cases

Data Velocity, Data Volume, Data Consistency, Real-Time Performance and Analysis

- **Trading Platforms**
Order Management and Execution Management Systems, algorithmic trading, high volume transactions, ultra low latencies.
- **Risk Management**
Modeling, financial engineering, pricing, hedging, what-if analysis.
- **Financial Analytics**
Real time analysis of trading positions, trending, market data analysis, sentiment analysis, complex event processing, hedging, transaction cost analysis, time series, volatility analysis, Monte Carlo simulations, Black-Scholes, derivatives pricing.
- **Big Data Analytics**
Customer and counter party 360 view, master data management, securities masters, reference data, real-time analysis of P&L, up-to-the-second operational BI.
- **Compliance and Monitoring**
Fraud, AML, KYC, market manipulation and abuse, pre and post trade compliance modeling.
- **Financial SaaS Platforms**
High performance next-generation architectures for Software as a Service Application vendors.



THOMSON REUTERS



Jefferies



APOLLO

Julius Bär



CAMBRIDGE



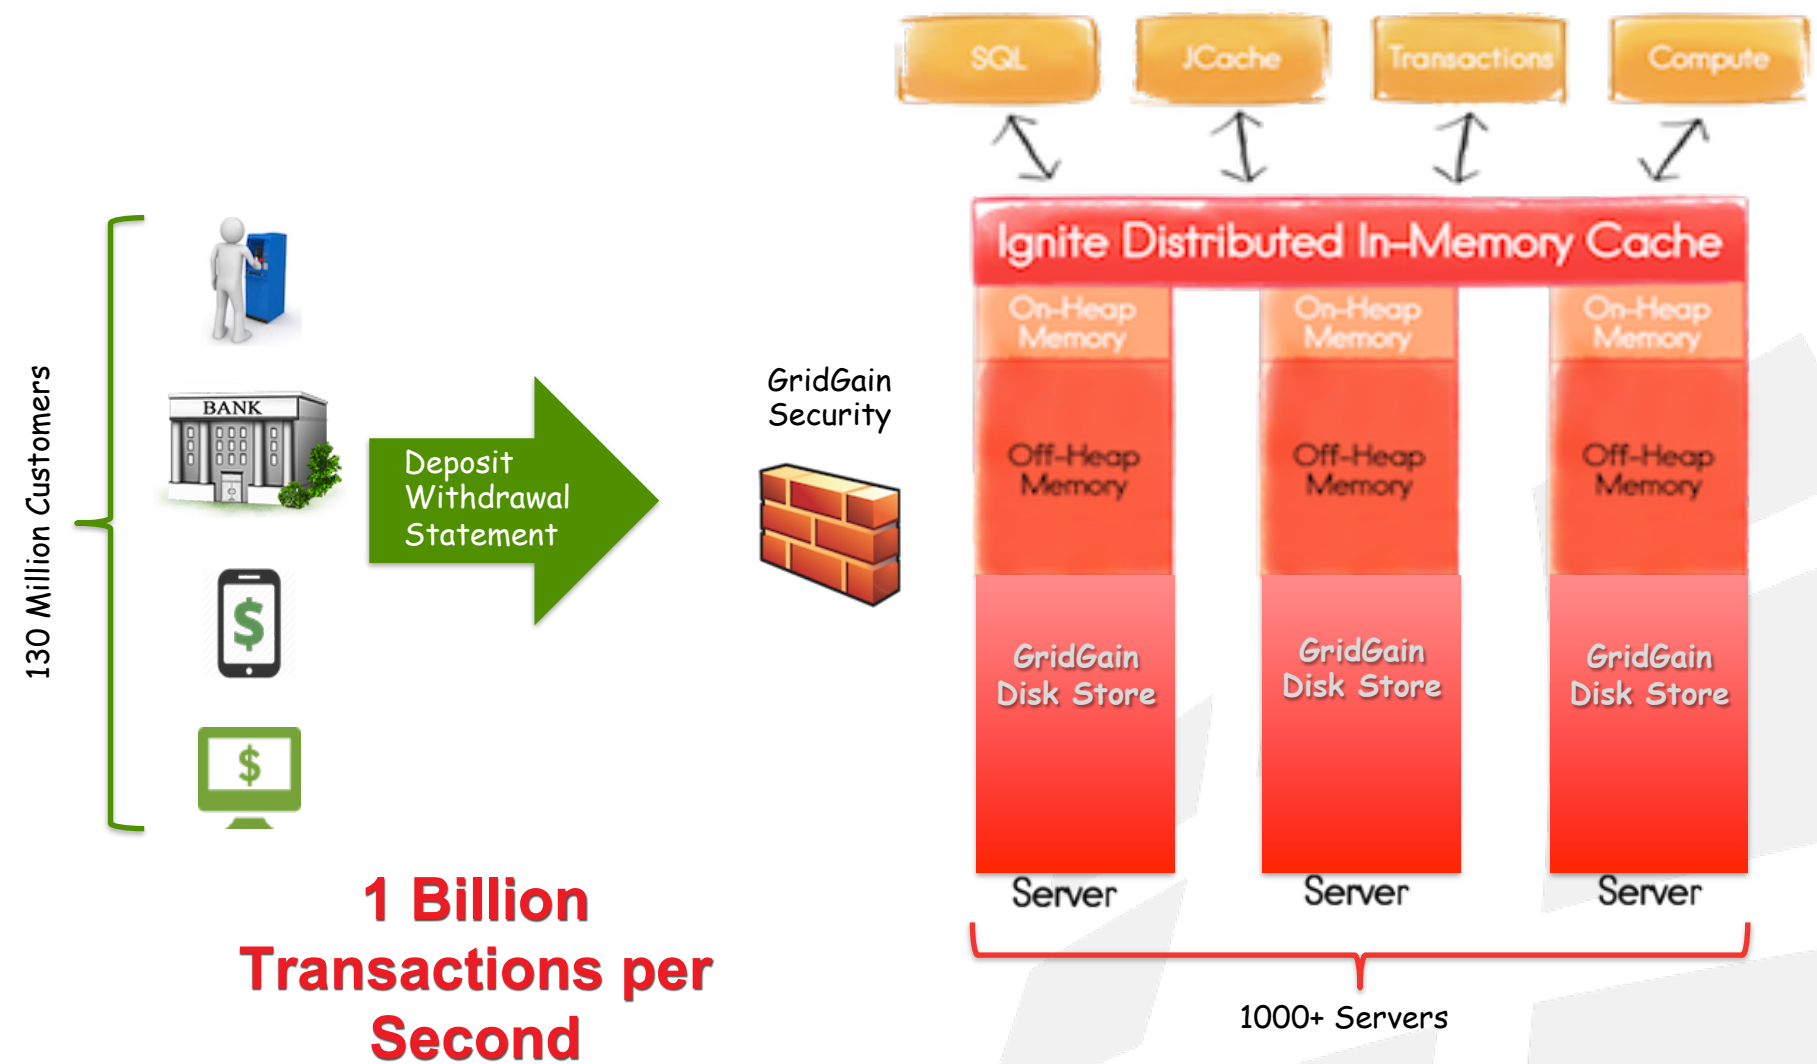
ASSOCIATES



Use Case: **SBERBANK**

Largest bank in Russia and Eastern Europe, and the third largest in Europe

- Sberbank Requirements
 - Migrate to data grid architecture
 - Minimize dependency on Oracle
 - Move to open source
- Why GridGain Won
 - Best performance
 - 10+ competitors evaluated
 - Demonstrated best
 - Fault tolerance & scalability
 - ANSI-99 SQL Support
 - Transactional consistency
 - Strict SLAs
 - Less than 5 min cluster restart
 - Fully Operational from disk & memory



10 Dell R610 blades = \$25K
1 TB Memory

From RBC article – January, 2016



Herman Gref
CEO & Chairman, Sberbank

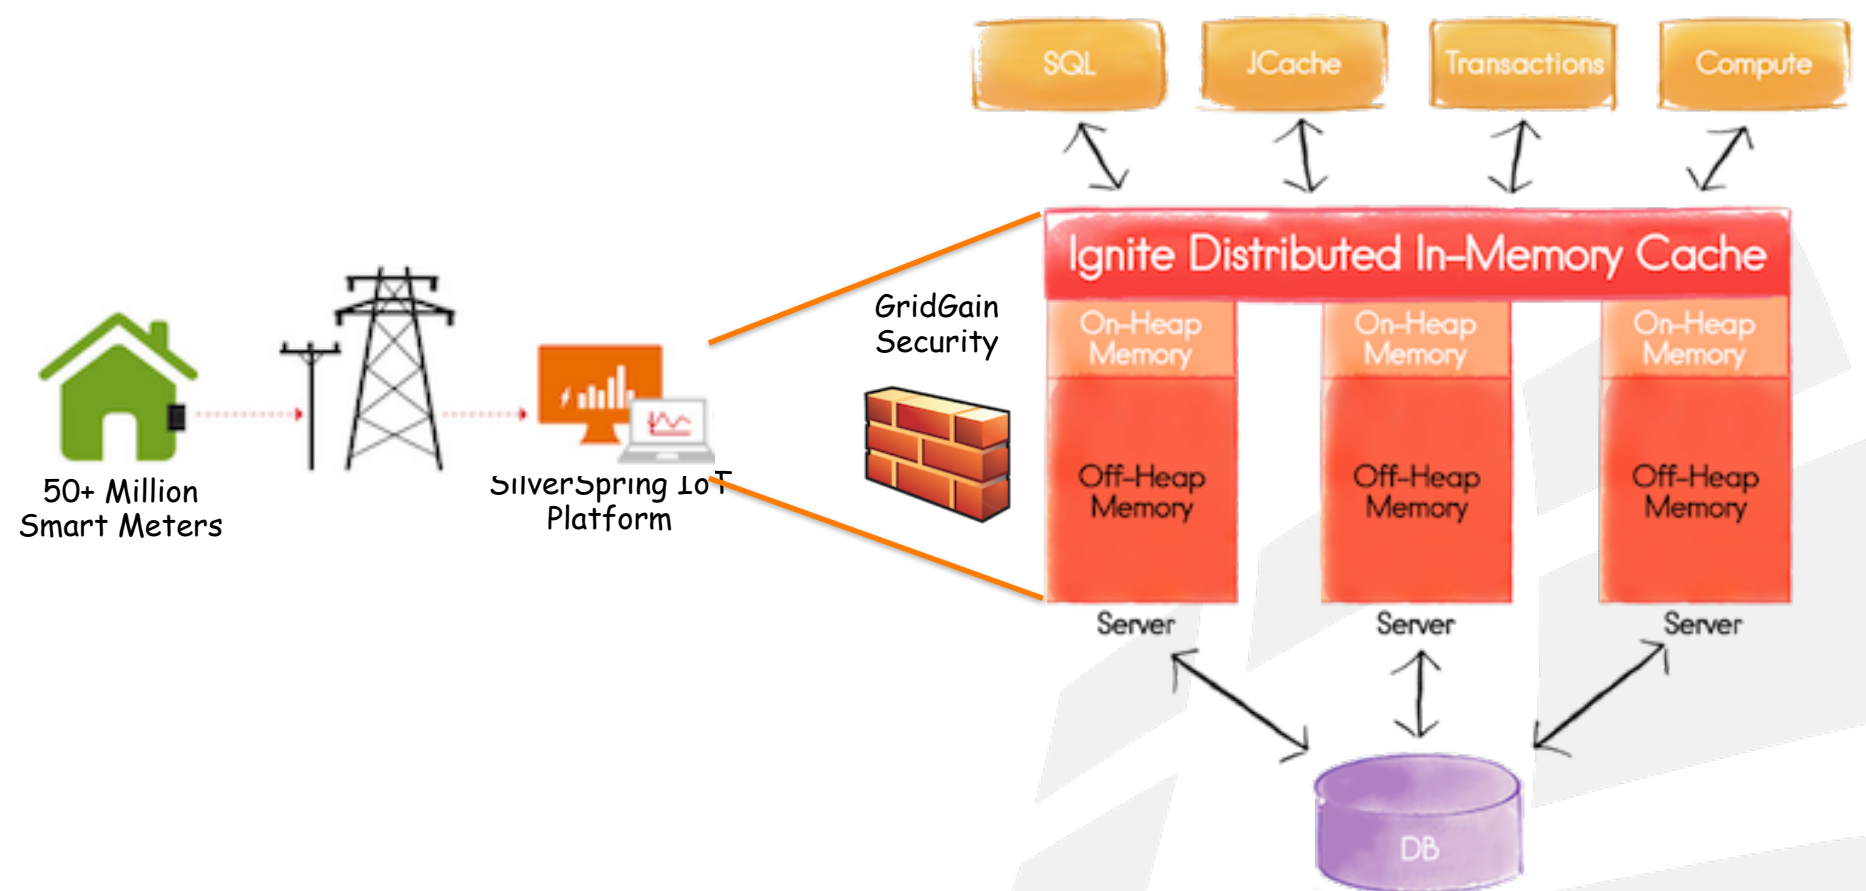
“The new Sberbank IT plan is to create a platform that enables the bank to introduce new products in hours, not weeks. The platform will have virtually unlimited performance and very high reliability. It will be much cheaper and will significantly reduce human interaction during customer transactions. The system will use machine-learning, flexible pricing, and artificial intelligence,” said German Gref, head of Sberbank.

“The new system will use technology from GridGain, which won the tender from Oracle, IBM and others, and turned out to deliver an order of magnitude higher performance than those of the largest companies,” he added.

Use Case:

Smart Metering and Utilities – delivers a comprehensive IOT platform

- SilverSpring Requirements
 - Migrate to in-memory processing
 - Add scalability & elasticity
 - Use open source technologies
- Why GridGain Won
 - Strong compute capabilities
 - Co-located in-memory processing
 - Demonstrated best
 - On-demand elasticity & scalability
 - ANSI-99 SQL Support
 - Transactional consistency




Use Case:




Online consumer electronics store – 11th largest retailer in US

- Newegg Requirements
 - Oracle Coherence replacement
 - Smart in-memory analytics
 - Use open source technologies
- Why GridGain Won
 - Best performance
 - 2x better than incumbent
 - Demonstrated best
 - ANSI-99 SQL Support
 - Resiliency & fault tolerance
 - Transactional consistency




NewEgg Ecommerce Site




Why Now?

Data Growth and Internet Scale Driving Demand



8 zettabytes in 2015 growing to 35 in 2020

Declining DRAM Cost Driving Attractive Economics



Cost drops 30% every 12 months

Gartner names GridGain a 2014 “Cool Vendor” for IMC



“This positions GridGain as one of the few IMC open-source technologies available and the only one... providing such a rich set of functionality.”

Gartner

GridGain named a AlwaysOn Global 250 Top Private Company For 2014 & 2015




GridGain Enterprise and Open Source Strategy



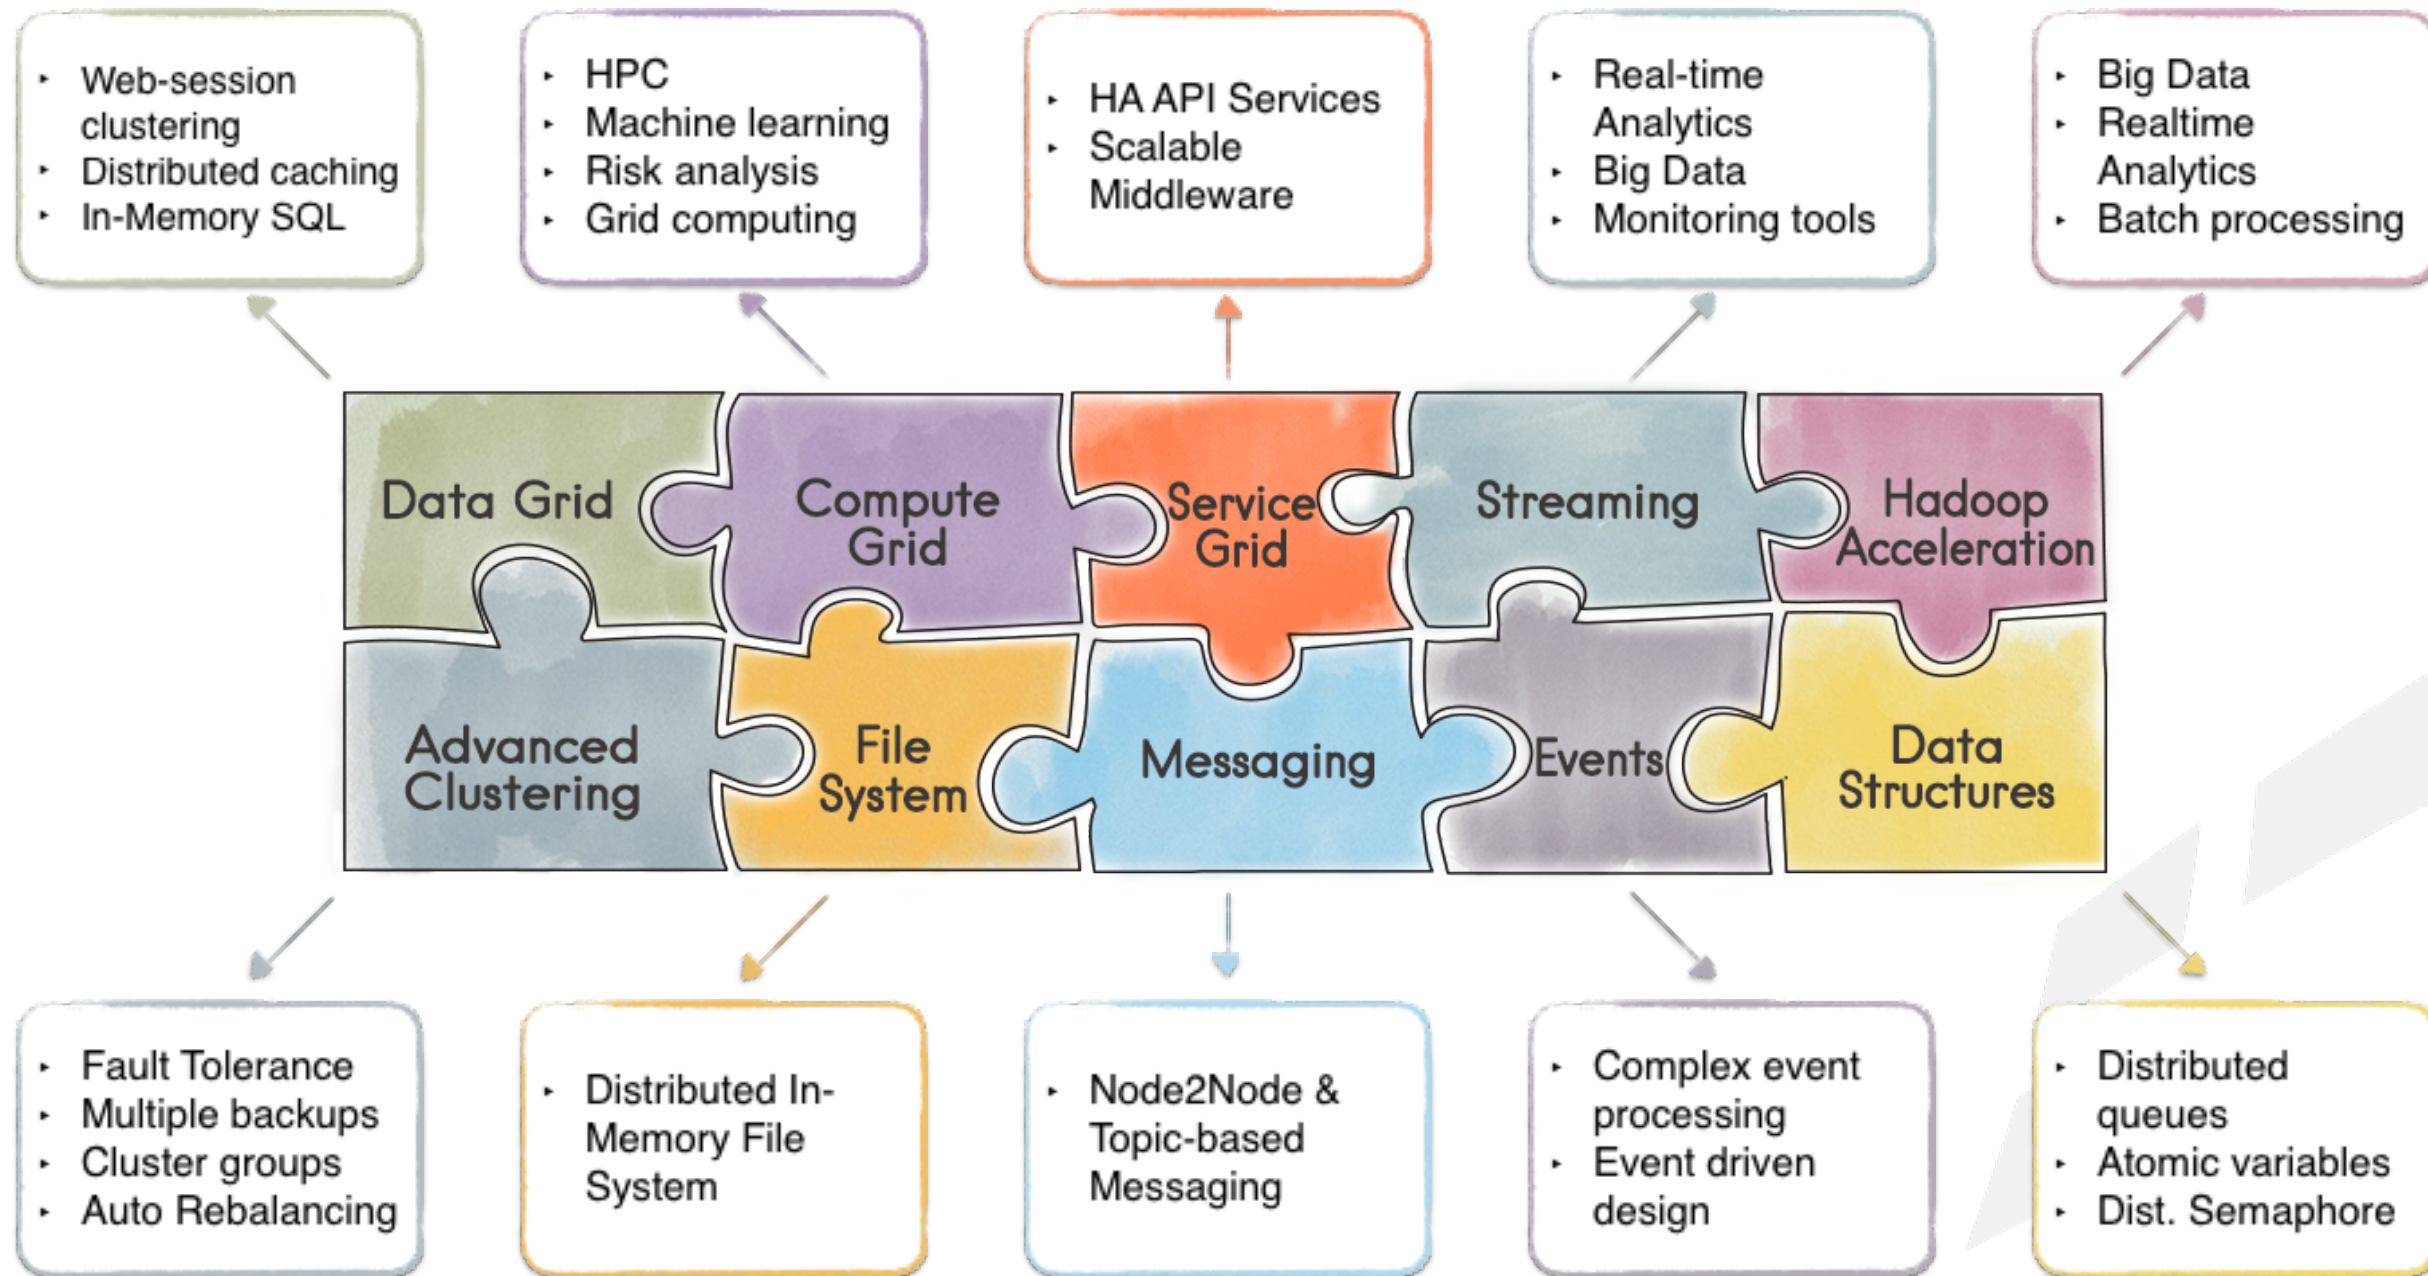
- GridGain Enterprise Edition is based on Apache Ignite
- Open source is intended to provide an easy entry point for learning, testing and non-critical use
- Enterprise Edition customers benefit from many exclusive enterprise-class features along with support and indemnification

What is an In-Memory Data Fabric?




High-performance distributed in-memory platform for computing and transacting on large-scale data sets in near real-time.

In-Memory Data Fabric Use Cases



Flexibility and Enterprise Breadth of In-Memory Data Fabric



In-Memory Data Fabric: Main Benefits

- Performance
 - High Throughput
 - Low Latency
 - Load Balancing
 - Caching
 - In-Memory Indexing
 - Eliminate Java Garbage Collection Pauses
- Scalability
 - Add Cluster Members (cores)
 - Unlimited Vertical Scale
 - Add Memory (RAM)
- High Availability
 - Data Backups
 - Datacenter Replication
 - Automatic Failover
 - Persistence
 - Fault Tolerance
 - Fast Recovery



- Transactions
 - Fully ACID Compliant
 - Optimistic & Pessimistic
 - Data Streaming
- Persistence
 - SQL, NoSQL, Hadoop
 - Tiered Memory: On-Heap ->
- Security
 - Auth & Auth
 - Encryption
 - Tracing & Auditing



GridGain's In-Memory Data Fabric

Enterprise Edition

GridGain Enterprise Subscriptions include:

- > Right to use GridGain Enterprise Edition
- > Bug fixes, patches, updates and upgrades
- > 9x5 or 24x7 Support
- > Training and Consulting Services from GridGain

Features	Apache Ignite	GridGain Enterprise
In-Memory Data Grid	✓	✓
In-Memory Compute Grid	✓	✓
In-Memory Service Grid	✓	✓
In-Memory Streaming	✓	✓
In-Memory Hadoop Acceleration	✓	✓
Distributed In-Memory File System	✓	✓
Advanced Clustering	✓	✓
Distributed Messaging	✓	✓
Distributed Events	✓	✓
Distributed Data Structures	✓	✓
Portable Binary Objects	✓	✓
Management & Monitoring GUI		✓
Enterprise-Grade Security		✓
Network Segmentation Protection		✓
Recoverable Local Store		✓
Rolling Production Updates		✓
Data Center Replication		✓
Integration with Oracle GoldenGate		✓
Basic Support (9x5)	✓	✓
Enterprise Support (9x5 and 24x7)		✓
Security Updates		✓
Maintenance Releases & Patches		✓

Free
w/ optional Paid Support

Annual License
Subscription

Analyst Validation



“This positions GridGain as one of the few IMC open-source technologies available and **the only one... providing such a rich set of functionality.**” (Massimo Pezzini, Gartner)



“IMDGs, in general, suffer from a caching stereotype. **GridGain smashes that stereotype** by offering a general purpose compute grid. **If you need an IMDG and a super-computer all in one then take a look at GridGain Systems.**” (Mike Gualtieri, Forrester)



“The company has strong expertise with regard to in-memory data processing ... The **move to contribute its core in-memory data processing capabilities to the Apache Software Foundation as Apache Ignite could prove to be an inspired one for GridGain** – raising the profile of its core technology and encouraging adopters toward its supported In-Memory Data Fabric product. **The company stands to gain significantly from proposing Apache Ignite** and becoming part of one of the largest families of open source software projects.” (Matt Aslett, 451 Research)



“EMA **considers GridGain a leading innovator** of the type of in-memory processing platforms that helps the transition from traditional processing methodologies to in-memory processing across a wide set of use cases. CIOs, Chief Data Officers (CDOs) and technology architects should regard GridGain as a viable option to add competitive advantage via near-realtime processing requirements now and going forward (John Myers, Enterprise Management Associates)



“Solutions such as GridGain (which is an optimizer layer that sits on top of existing solutions) are **focusing on this market segment...**” (Lynn Langit, GigaOm Research)



“GridGain has now put their key IP out for the world to use as Apache Ignite. **We predict an enthusiastic adoption curve** since GridGain’s approach leverages current developer skills and accelerates existing Hadoop apps without modification, **providing an easy transition into in-memory computing.**” (Mike Matchett, The Taneja Group)



THANK YOU