



# Take Telecom to the Next Level with In-Memory Computing

Matt Sarrel  
Industry Consultant

# Agenda

- Introduction
- What is In-Memory Computing?
- GridGain / Apache Ignite Overview
- Survey Results
- Use Cases and Case Studies
- GridGain / Apache Ignite In-depth

# Your Presenter



- Industry Consultant
- Journalist/Reviewer
- Market Analyst
- 30 years in tech (networking, big data, analytics, security)

# Telecom: An industry in Transition

Over 1 billion smartphones sold per year with many using mobile as their primary computing platform.

Customers are increasingly focused on data, especially streaming. Media and video was less than 10% of traffic in 2010 and almost 50% in 2015.

50% decline in mobile voice minutes used between 2005 and 2012

# Telecom is About Data

## Market Shift

- Tiered usage-based pricing for mobile data
- Data sharing for households and businesses

## IoT

- According to Gartner, 26 billion connected devices by 2020
- Consumer and business devices
- 1B M2M connections via mobile networks predicted by 2020 (10% of total mobile connections)

# Network Operators Under Pressure

## Over-the-Top (OTT) Content Providers

- New services challenge operators
- Instant messaging (voice, mobile payments)
- Disrupting the value chain (shifting up to 10% of revenue)

## Adapting to the New Value Chain

- Avoid commoditization
- Focus on digital services
- Leverage traditional competencies

# Adopting a Data Driven Approach in Order to Compete

## Customer Experience Management

- Constantly assessing network performance and customer experience
- Differentiators build a deeper relationship with customers
- Rapidly capture and act on customer feedback

## Service Levels and Personalization

- New generation has never lived in a world without digital
- Expect high-quality, seamless omnichannel interactions
- No carrier loyalty, will switch quickly
- Personalized services and support interactions

## Network Quality

- 4G in-filling
- Backhaul network modernization
- Preparing for 5G
- More sophisticated traffic monitoring and performance management



# Why In-Memory Now?

## • Digital Transformation is Driving Companies Closer to Their Customers

- Customers require real-time interactions
- Companies require real-time insights

## Internet Traffic, Data, and Connected Devices Continue to Grow

- Web-scale applications and massive datasets require in-memory computing to scale out and speed up to keep pace
- The Internet of Things generates huge amounts of data which require real-time analysis for real world uses


## The Cost of RAM Continues to Fall

- In-memory solutions are increasingly cost effective versus disk-based storage for many use cases




# 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

# The Rapidly Growing In-Memory Technology Market

(in \$Billions)



# What is an In-Memory Computing Platform?

## • Multi-Featured Solution

- Includes an in-memory data grid, in-memory database, and streaming analytics

## No Rip and Replace

- Slides in between the existing application and data layers

## Supports OLTP and OLAP Use Cases

- Offers ACID compliant transactions as well as SQL-based analytics support either separately or in combination (HTAP)

## Multi-Platform Integration

- Works with all popular RDBMS, NoSQL and Hadoop databases
- Unified API with support for a wide range of protocols

## Deployable Anywhere

- Can be deployed on premise, in the cloud, or in hybrid environments
- Fully fault-tolerant and load-balanced on node, cluster and data center levels

# The GridGain In-Memory Computing Platform

- *High Performance*
- *Distributed*
- *Memory-Centric*
- *Built on Apache<sup>®</sup> Ignite<sup>™</sup>*

## • Features

- Data Grid

Memory-Centric  
Database

Streaming Analytics

ACID Transactions

SQL, DDL & DML

## Architecture

Advanced Clustering

Key-Value Store

Compute Grid

Service Grid

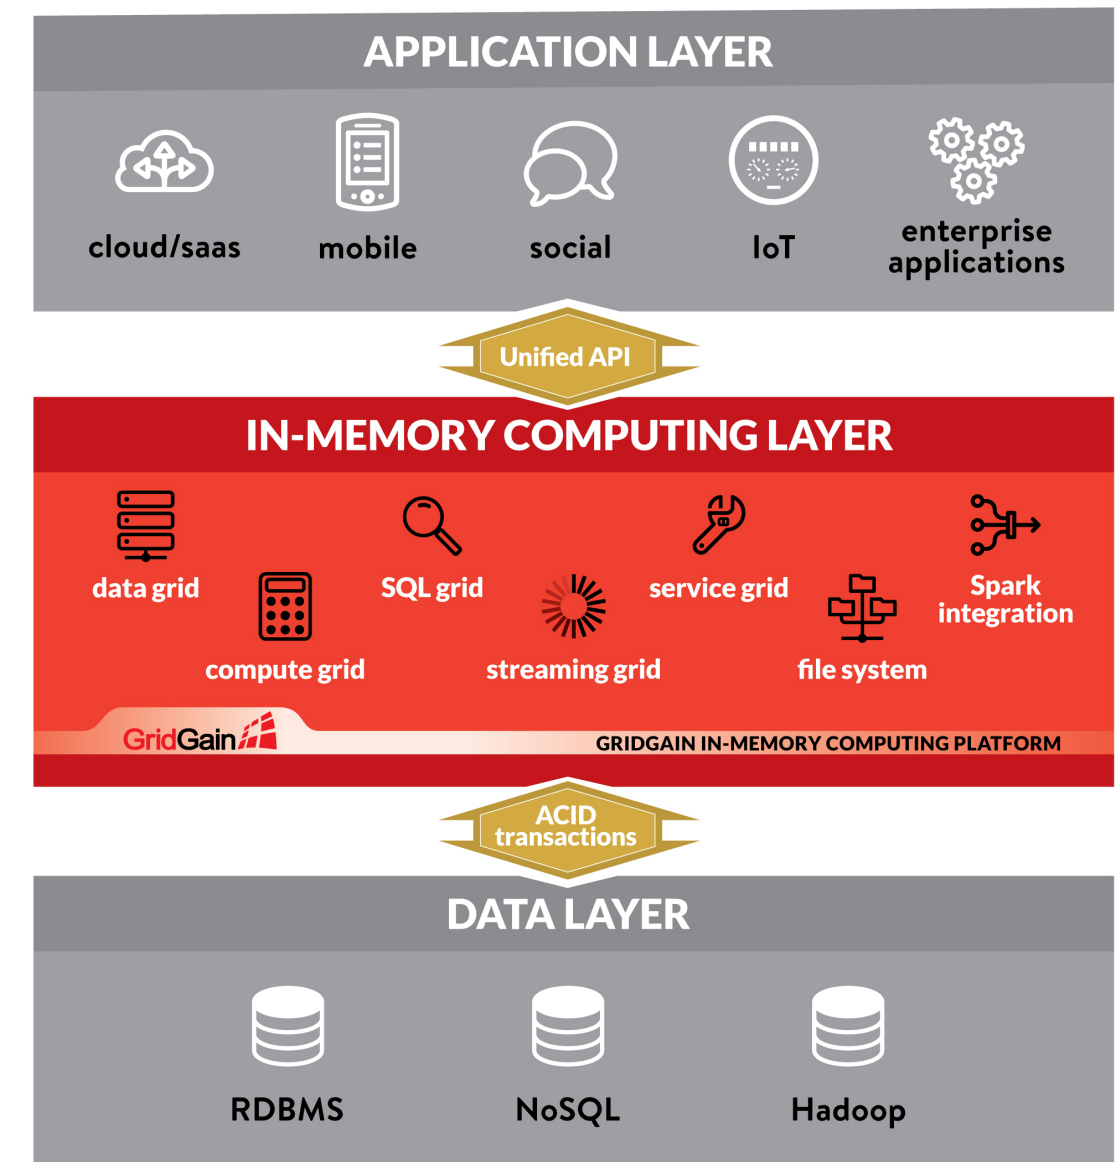
# GridGain is the Leading In-Memory Computing Platform

*(and based on Apache Ignite)*



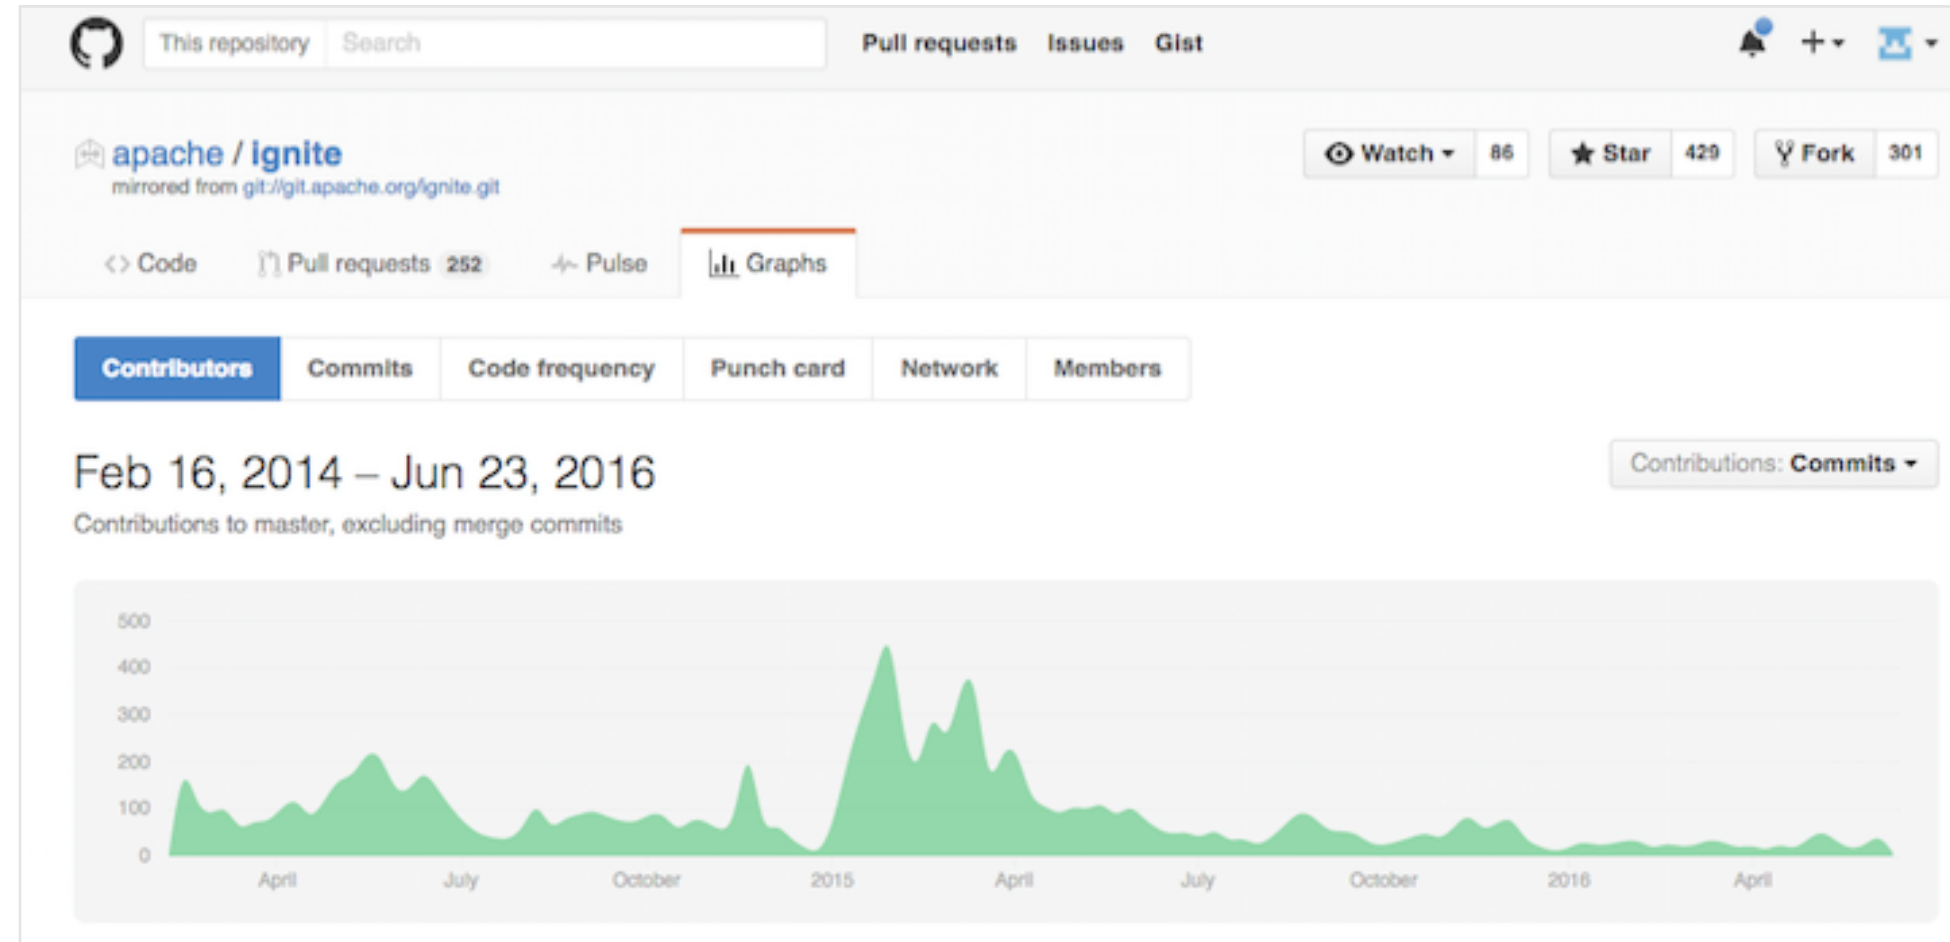
- Delivers Orders-of-Magnitude Improvements in **Speed** and **Scale** for a Variety of Use Cases
  - Big Data analytics
  - SaaS and Cloud computing
  - Mobile and Internet of Things backends
  - Cognitive Computing
  - Streaming and real-time processing
  - Web-scale applications
- Slides Seamlessly Into New or Existing Architectures
  - Between the application and data layers
  - Persistent Store offers transactional SQL database capabilities across a memory-centric, distributed computing environment

## TRANSACTIONAL, ANALYTICAL, OR HYBRID USE CASES




# Apache Ignite Project

- 2007: First version of GridGain
- Oct. 2014: GridGain contributes Ignite to ASF
- Aug. 2015: Ignite is the second fastest project to | graduate after Spark
- Today:
  - 60+ contributors and rapidly growing
  - Huge development momentum - Estimated 192 years of effort since the first commit in February, 2014 [\[Openhub\]](#)
  - Mature codebase: 1M+ lines of code



# Apache Ignite and GridGain Downloads





# Comparison of GridGain Editions and Apache Ignite

| Features                          | Apache Ignite | Professional | Enterprise | Ultimate |
|-----------------------------------|---------------|--------------|------------|----------|
| In-Memory Data Grid               | ○             | ○            | ○          | ○        |
| In-Memory Compute Grid            | ○             | ○            | ○          | ○        |
| Persistent Store                  | ○             | ○            | ○          | ○        |
| In-Memory Streaming               | ○             | ○            | ○          | ○        |
| Distributed SQL with DDL & DML    | ○             | ○            | ○          | ○        |
| In-Memory Service Grid            | ○             | ○            | ○          | ○        |
| Distributed In-Memory File System | ○             | ○            | ○          | ○        |
| In-Memory Hadoop Acceleration     | ○             | ○            | ○          | ○        |
| Advanced Clustering               | ○             | ○            | ○          | ○        |
| Distributed Messaging             | ○             | ○            | ○          | ○        |
| Distributed Events                | ○             | ○            | ○          | ○        |
| Distributed Data Structures       | ○             | ○            | ○          | ○        |
| Portable Objects                  | ○             | ○            | ○          | ○        |

# Comparison of GridGain Editions and Apache Ignite

| Features                        | Apache Ignite | Professional | Enterprise | Ultimate |
|---------------------------------|---------------|--------------|------------|----------|
| Security Updates                |               | ○            | ○          | ○        |
| Maintenance Releases & Patches  |               | ○            | ○          | ○        |
| Management & Monitoring GUI     |               |              | ○          | ○        |
| Enterprise-Grade Security       |               |              | ○          | ○        |
| Network Segmentation Protection |               |              | ○          | ○        |
| Rolling Production Updates      |               |              | ○          | ○        |
| Data Center Replication         |               |              | ○          | ○        |
| Oracle GoldenGate Integration   |               |              | ○          | ○        |
| Cluster Snapshots               |               |              |            | ○        |

# GridGain In-Memory Computing Use Cases

## •Data Grid

Web session clustering

Distributed caching

Scalable SaaS

## Compute Grid

High performance computing

Machine learning

Risk analysis

Grid computing

## Distributed SQL

In-memory SQL

Distributed SQL processing

Real-time analytics

## Streaming

Real-time analytics

Streaming Big Data analysis

Monitoring tools

## Hadoop Acceleration

Faster Big Data insights

Real-time analytics

Batch processing

## Events

Complex event processing (CEP)

Event driven design

# GridGain Customers

## Financial Services



## FinTech



## Software



Siemens PLM Software



## Telecom & Mobile



## IoT



## Pharma & Healthcare



## Logistics & Transportation




## eCommerce & Retail




## AdTech




# Survey Results: What uses were you considering for in-memory computing



# Survey Results: Where do you run GridGain and/or Apache Ignite?




# Survey Results: Which of the following protocols do you use to access your data?






# Survey Results: Which data stores are you/would you likely use with GridGain/Apache Ignite?



# Survey Results: How important are each of the following product features to your organization?



# Case Study:



- Background:
  - Intelligentpipe is a big data software company serving the global telecommunications industry by developing solutions for mobile operators to improve their business and operational processes
- The Challenge
  - Collect and analyze massive amounts of mobile user traffic data in real time
  - Tens of millions of users
  - Consumption of network resources
  - Type of network traffic (voice or data)

# Case Study:

- GridGain Professional Edition used to build a high performance low latency analysis platform
- “GridGain ensures responsiveness regardless of how much information we need to search through.” Sakari Paloviita, CTO, Intelligentpipe
- Collect and analyze multiple terabytes per day

# Case Study:

- Real-time analytics provides fast insight
- Easy integration with existing systems due to GridGain's Unified API and ANSI SQL-99 support
- Linear scaling across deployed server to keep up seamlessly as the business grows

We'll want to use technology GridGain offers so we can focus on our core business ourselves."

- Jari Kuusela, Director of Product Management

# Building a Single Customer View

- Multi-national provider with over 250 million customers across mobile, fixed line, and broadband
- Customer interactions increasingly online
  - Support
  - Billing
  - Self-service
  - 360 degree view

# Building a Single Customer View

- Requirements
  - Accommodate increasing traffic to self-service web portals
  - Real-time data updates to customer account
  - Higher query throughput
  - Future-proof growth in customer base
  - Retention of extensive SQL-based BI code
  - High availability and timely system recovery



# Building a Single Customer View

- POC
  - 5-6 weeks
  - 10 physical nodes (48 core/512GB RAM/512GB disk)
  - Partitioned grid topology
  - Each existing DB stored in it's own cache
  - Updates received via Kafka streaming
  - SQL queries to re-generate and update customer billing and account details
  - SQL queries to respond to user requests via customer web portal

# Building a Single Customer View

- The results
  - Platform provides up to 60,000 updates a second
  - Individual updates complete in average time of 40ms
  - Full customer web view processes within 40ms
  - Full web page render time below 1.4 sec
  - All existing SQL queries and business logic reused with little refactoring

# Case Study:



- Background:
  - Cyber Dust is a platform for text messages: “A safer place to text.”
    - Untraceable
    - Encrypted
    - Disappearing
    - Screenshot blocking
  - Available for Android and iOS
  - Mark Cuban funded
- The Challenge
  - To build a real-time, reliable and highly available server infrastructure to support a mobile messaging platform
  - More than 500K users
  - Millions of messages a day
  - Avoid writing messages to disk

# Case Study:



- GridGain Professional Edition used to build a messaging platform
- Runs completely on Amazon EC2
- All user account data, configurations, and messages held in memory
- Messages deleted without a trace because they were never written to disk
- Extensive use of Unified API

# Case Study:



- “Blast” feature performance: capable of broadcasting disappearing messages to all of a user’s contacts
- Real-world performance of 300,000 messages sent and disappeared in 30 seconds

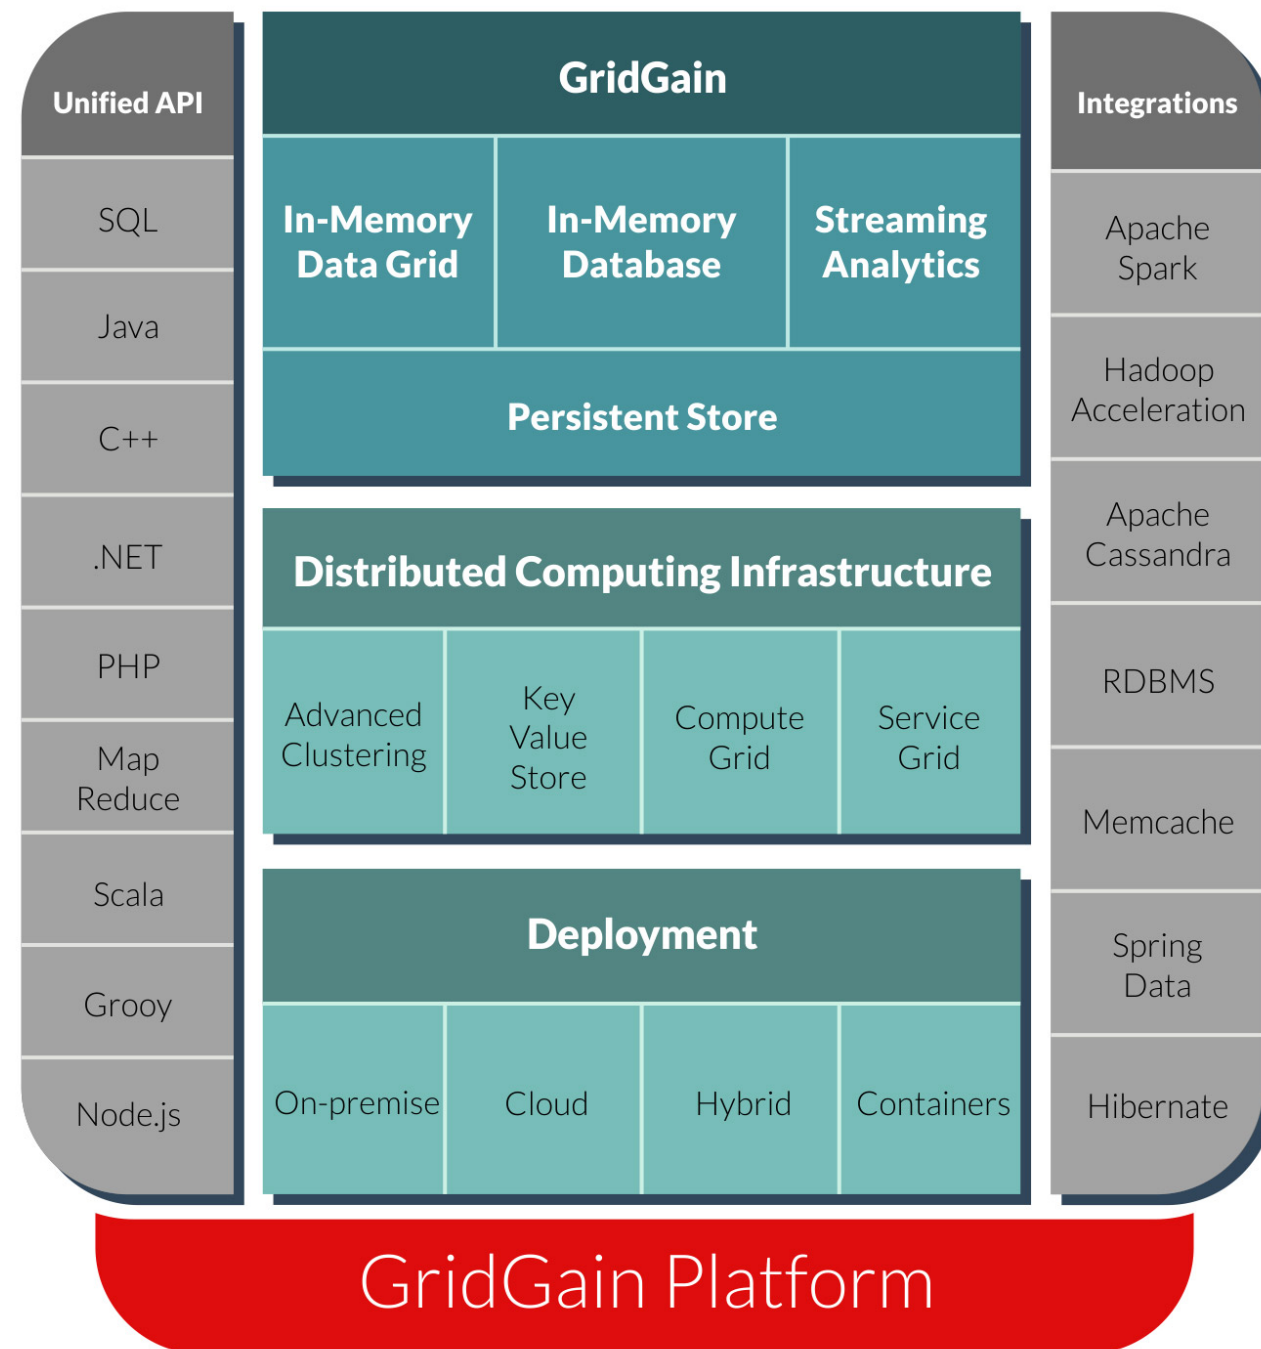
I was pleasantly surprised by  
the GridGain solution and  
performance.

-Igor Shpitalnik, CTO

I keep learning about  
additional capabilities  
GridGain offers. It’s what I  
expected and more.

-Igor Shpitalnik, CTO

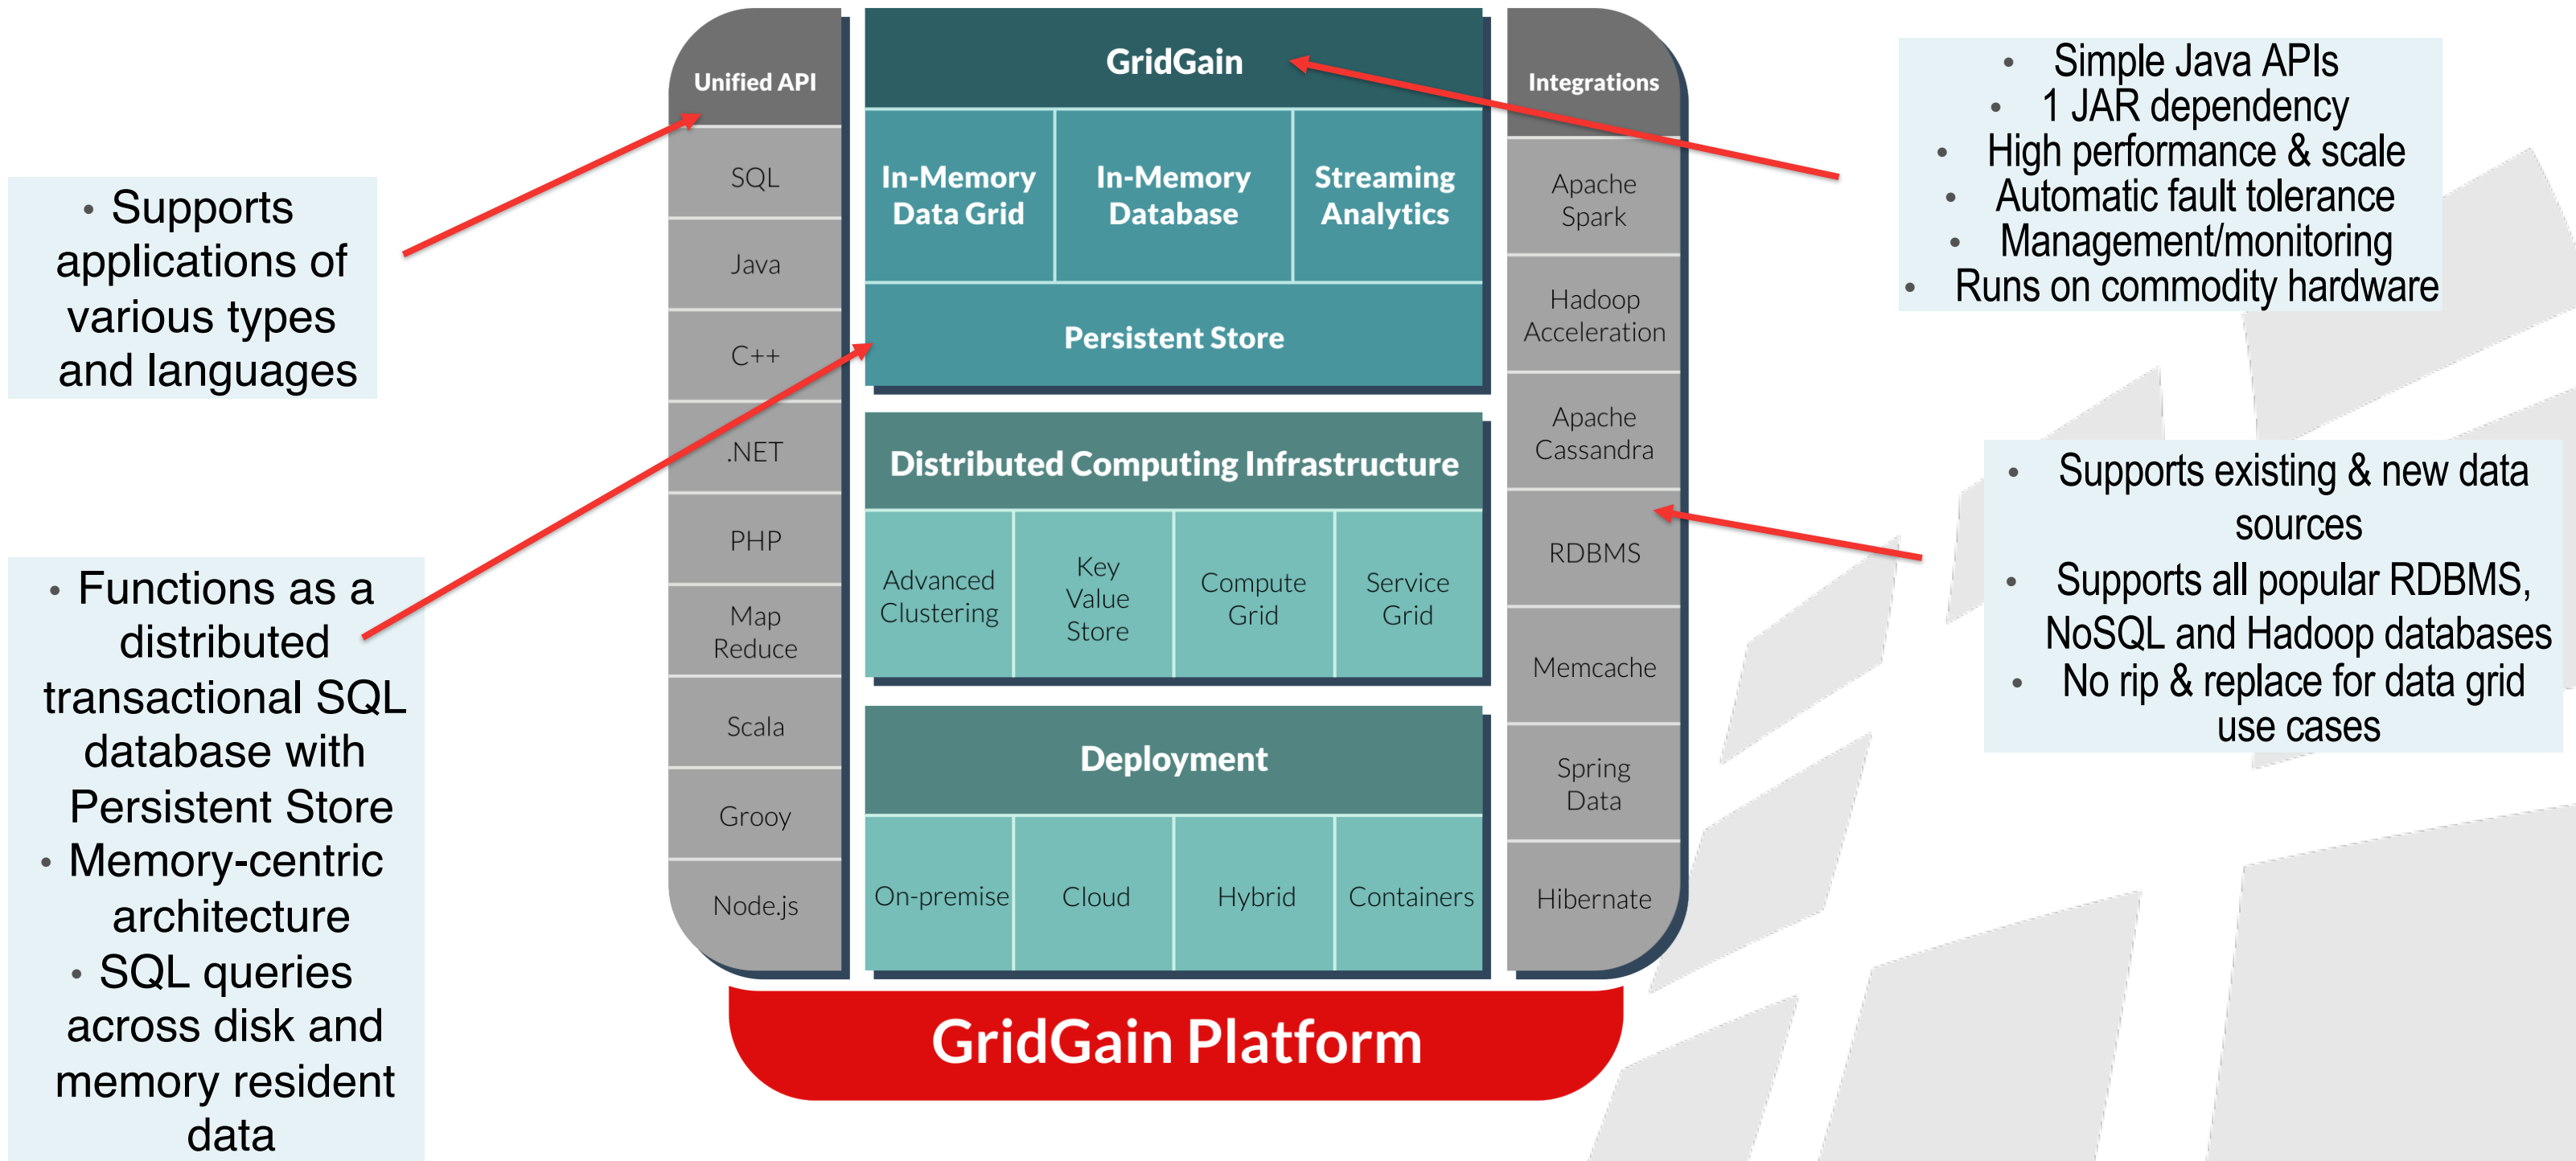
# The GridGain In-Memory Computing Platform



- **Richest functionality – a true in-memory platform**
  - In-memory data grid, in-memory database, streaming analytics
  - Persistent Store for memory-centric computing
  - SQL, ACID transactions, works with any database type, native integrations with many products
- **Built on open source Apache Ignite technology**
  - Ignite is for Fast Data what Hadoop is for Big Data
- **Most strategic, least disruptive approach to in-memory computing speed and scale**
  - No rip-and-replace of existing infrastructure required for IMDG, IMDB or Streaming Analytics
  - Deployable anywhere

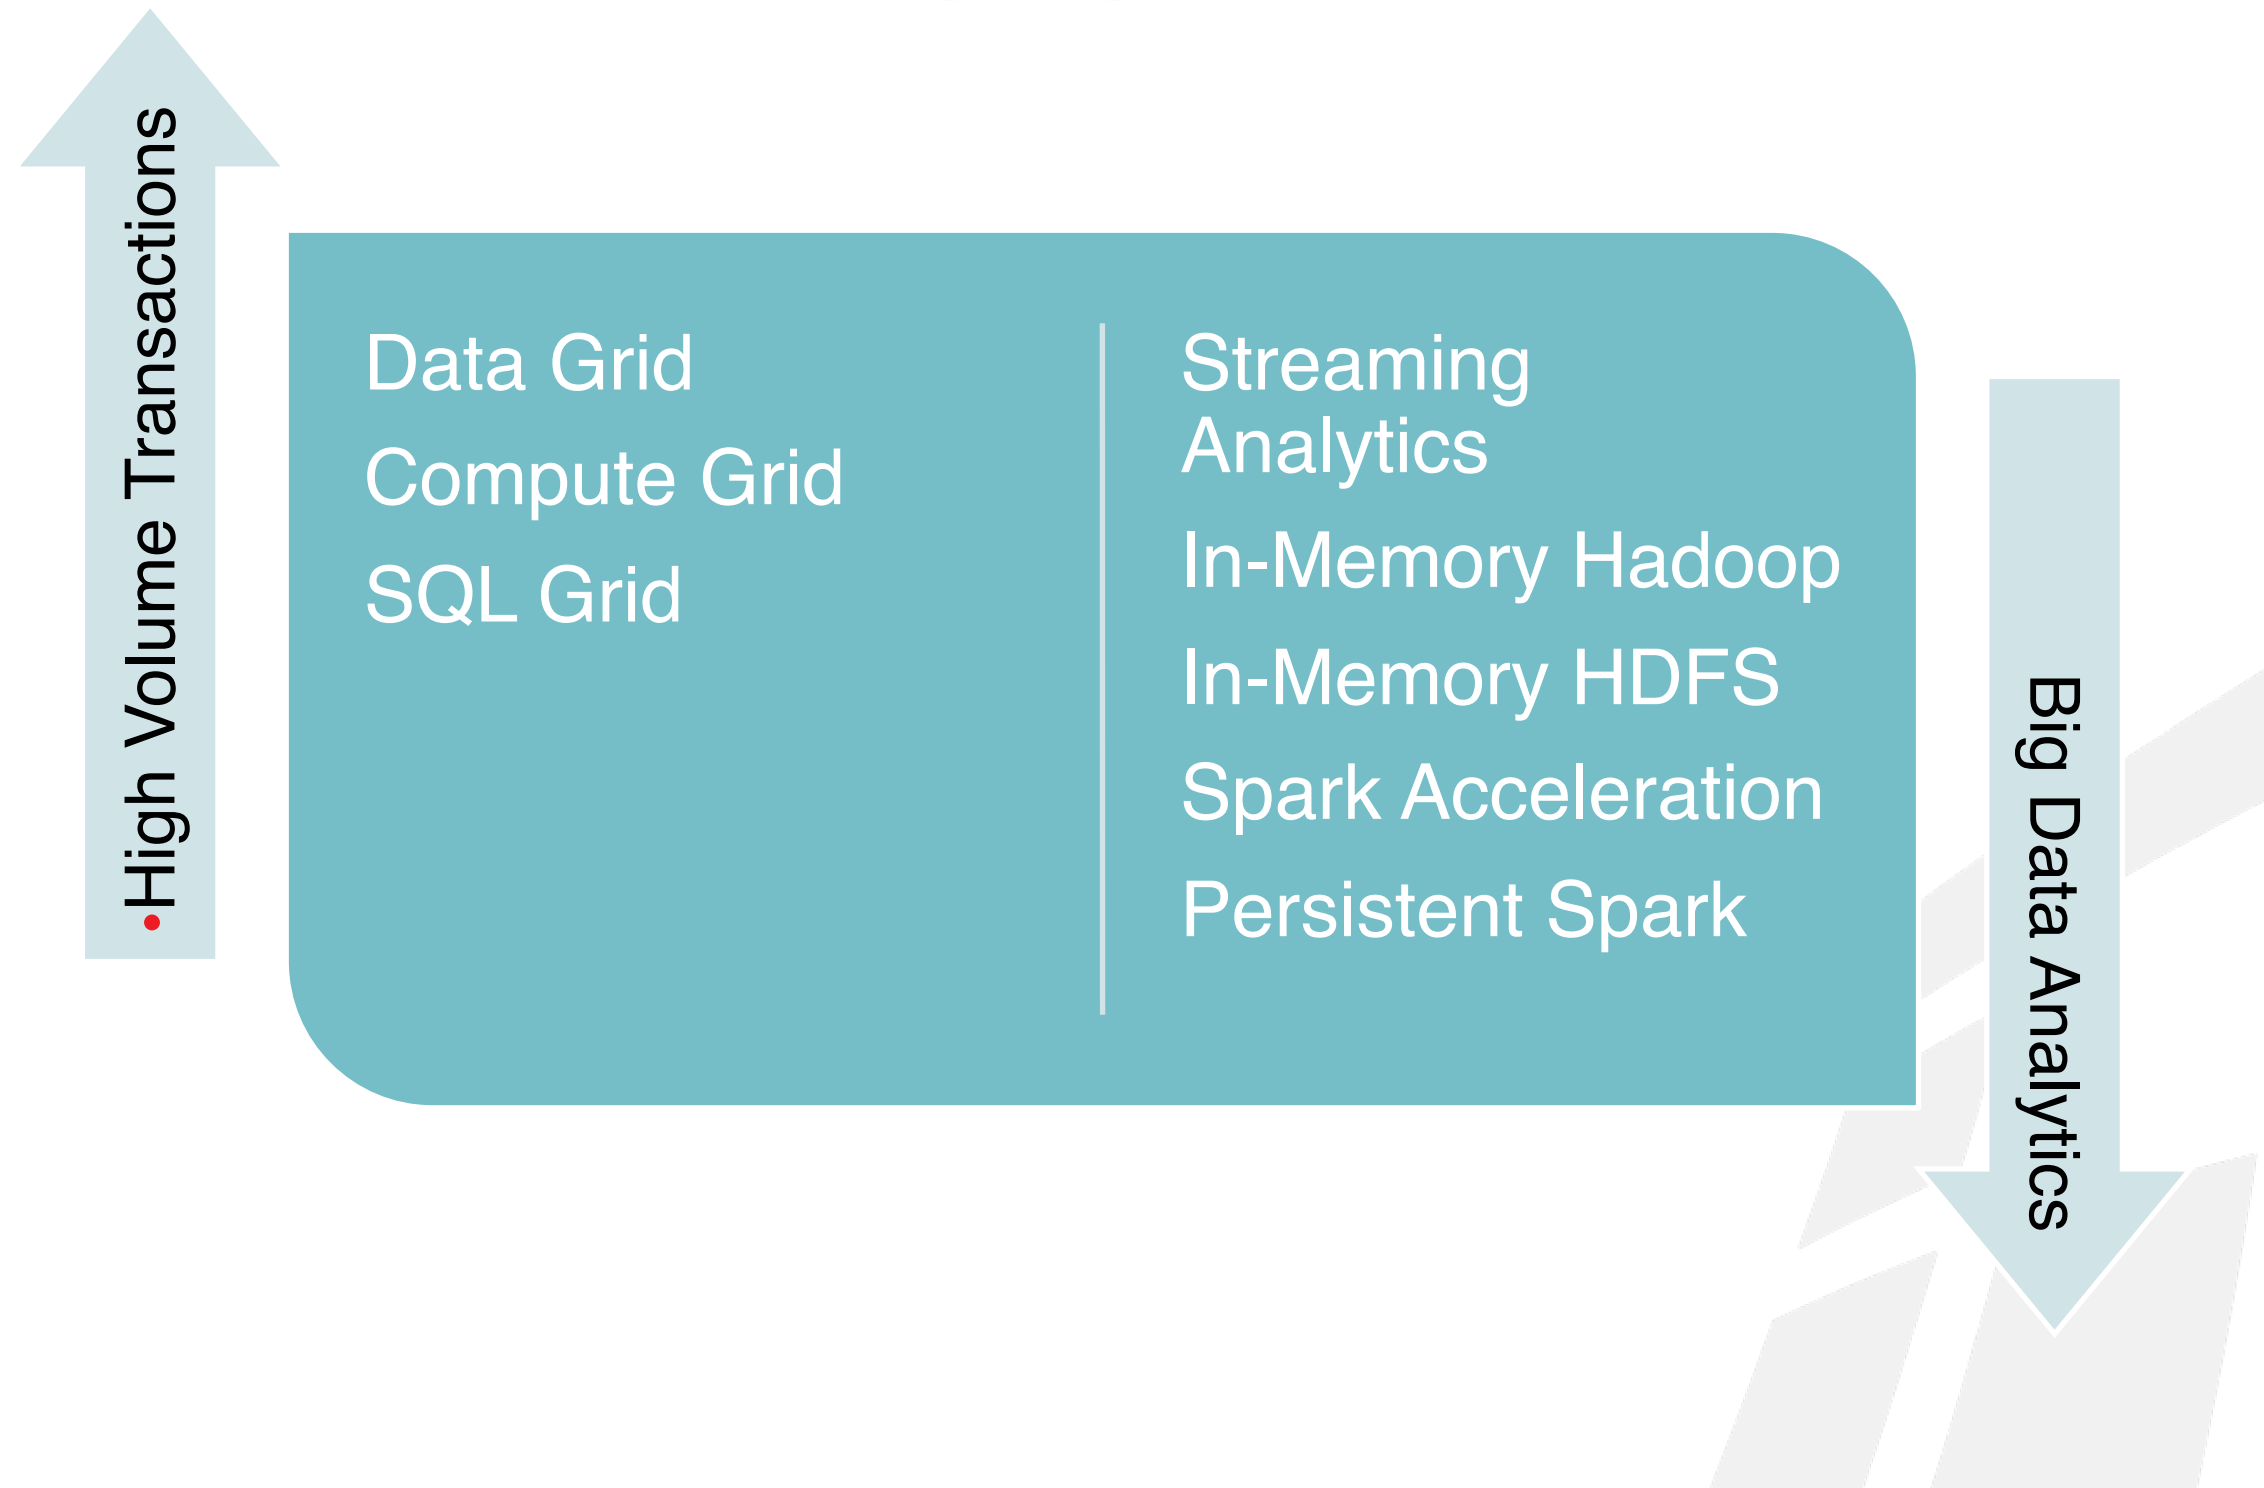


# The GridGain In-Memory Computing Platform



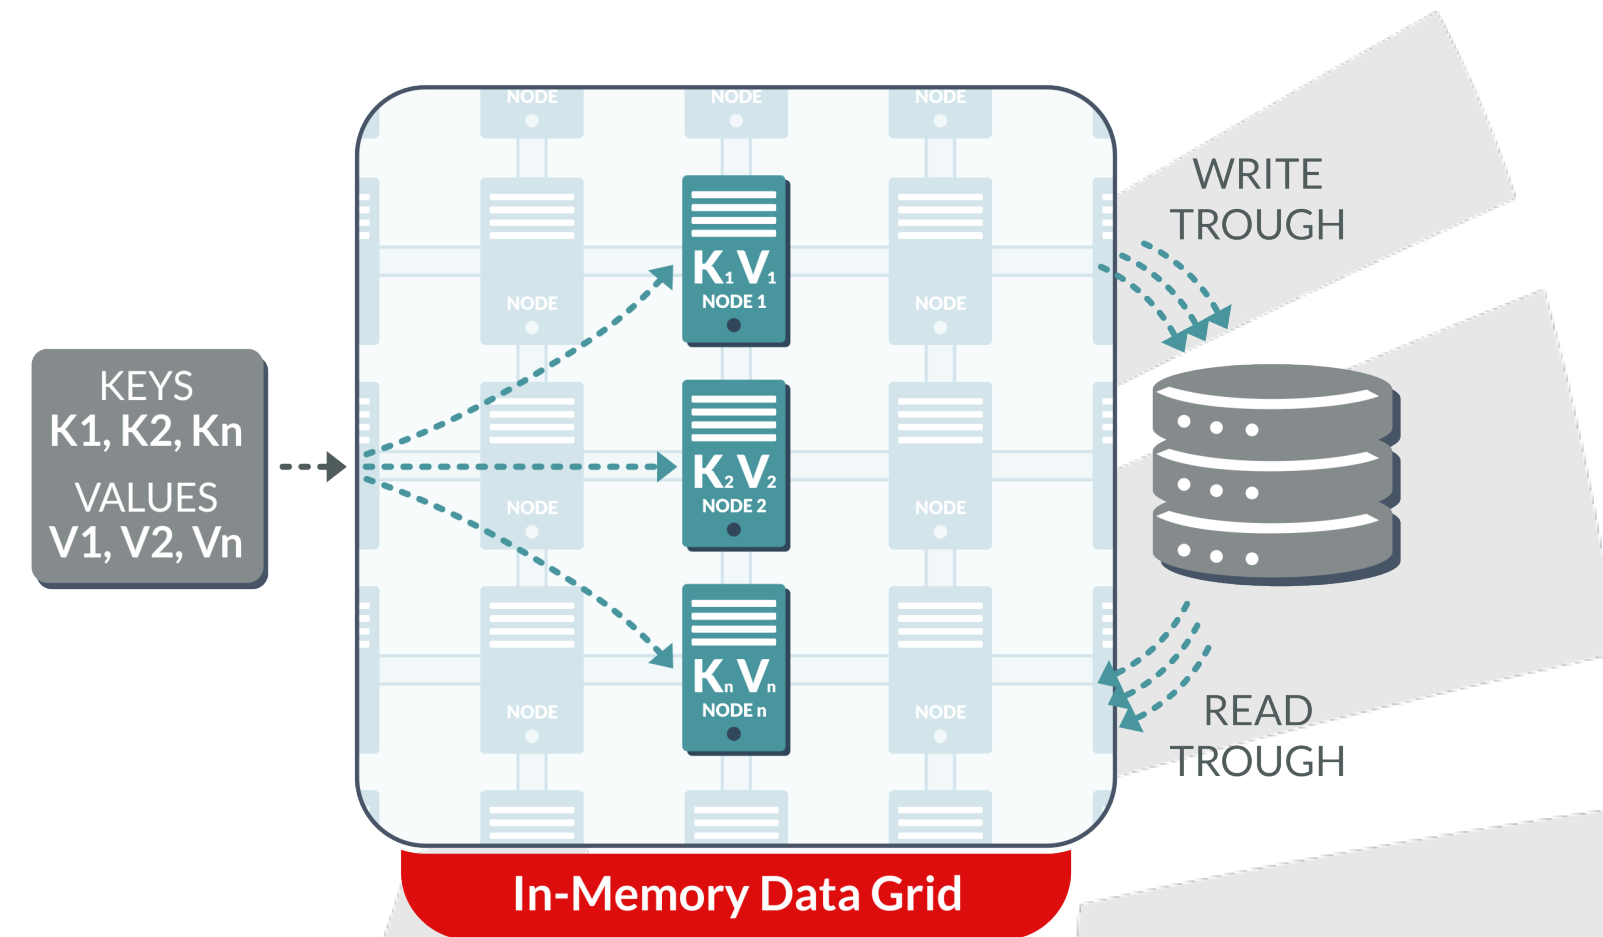


# Transactions & Analytics in a Single Platform



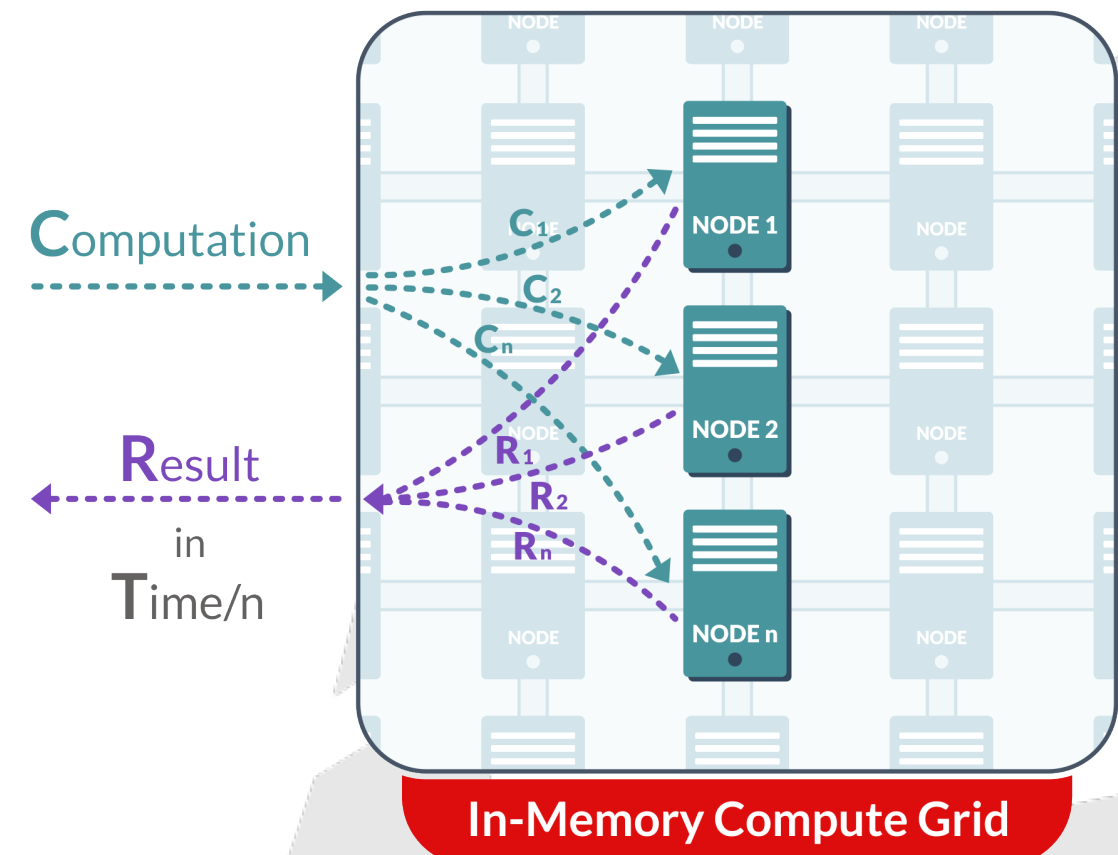
# GridGain In-Memory Data Grid

- *Moves disk-based data from RDBMS, NoSQL or Hadoop databases into RAM*
- **Features:**
  - *Distributed In-Memory Caching*
  - *Lightning Fast Performance*
  - *Elastic Scalability*
  - *Distributed In-Memory Transactions*
  - *Distributed In-Memory Queue and Other Data Structures*
  - *Web Session Clustering*
  - *Hibernate L2 Cache Integration*
  - *Tiered Off-Heap Storage*
  - *Distributed SQL Queries with Distributed Joins*



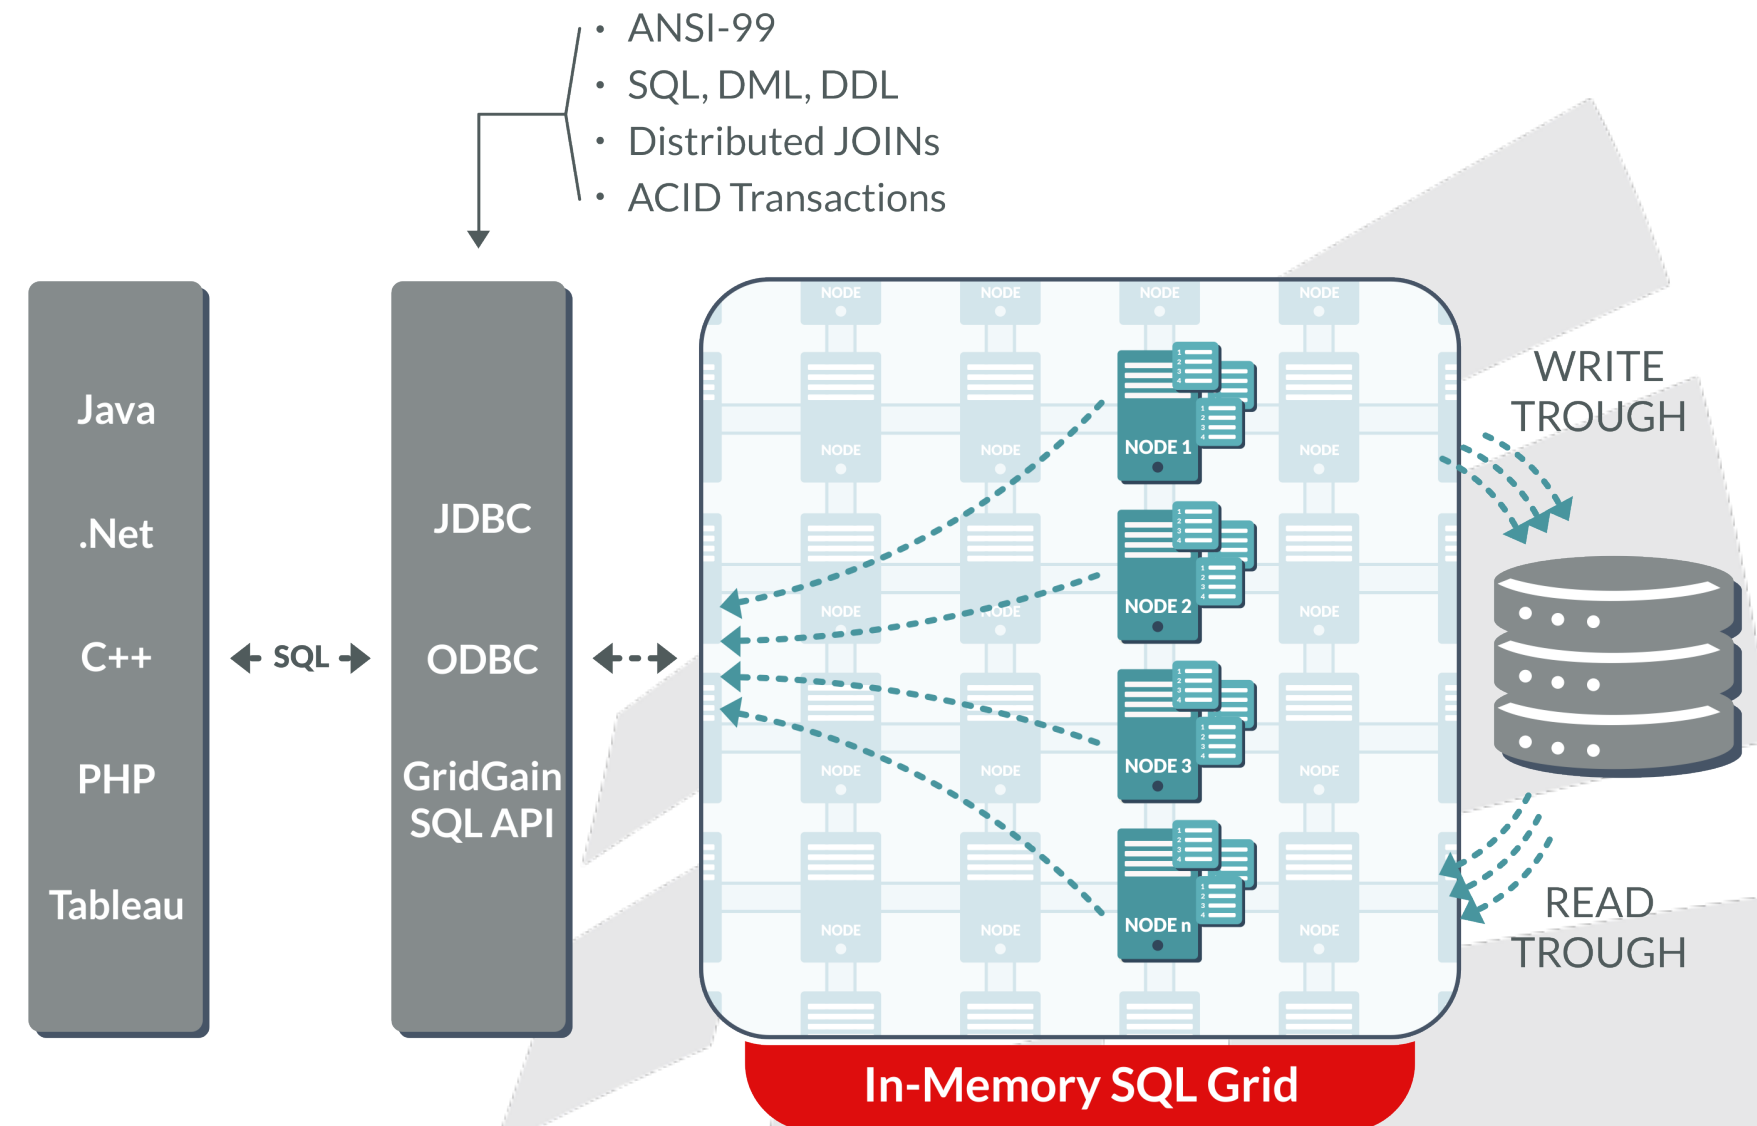
# GridGain In-Memory Compute Grid

- *Enables parallel processing of CPU or otherwise resource intensive tasks*
- **Features:**
  - *Dynamic Clustering*
  - *Fork-Join & MapReduce Processing*
  - *Distributed Closure Execution*
  - *Adaptive Load Balancing*
  - *Automatic Fault Tolerance*
  - *Linear Scalability*
  - *Custom Scheduling*
  - *Checkpointing for Long Running Jobs*
  - *ExecutorService*



# Distributed SQL

- *Horizontally scalable and fault tolerant*
- *ANSI SQL-99 compliant with DDL & DML*
- *Leverage existing SQL applications with in-memory speed and scale*
- **Features:**
  - *Supports SQL and DML commands including SELECT, UPDATE, INSERT, MERGE and DELETE Queries*
  - *DDL for managing caches and SQL schema with commands like CREATE and DROP table*
  - *Distributed SQL*
  - *Geospatial Support*
  - *SQL Communications Through the GridGain ODBC or JDBC APIs Without Custom Coding*
  - *ANSI SQL-99 Compliance*




# GridGain In-Memory Streaming Analytics Engine

- *Allows queries into sliding windows of incoming data*
- **Features:**
  - *Programmatic Querying*
  - *Customizable Event Workflow*
  - *At-Least-Once Guarantee*
  - *Sliding Windows*
  - *Data Indexing*
  - *Distributed Streamer Queries*
  - *Co-Location With In-Memory Data Grid*



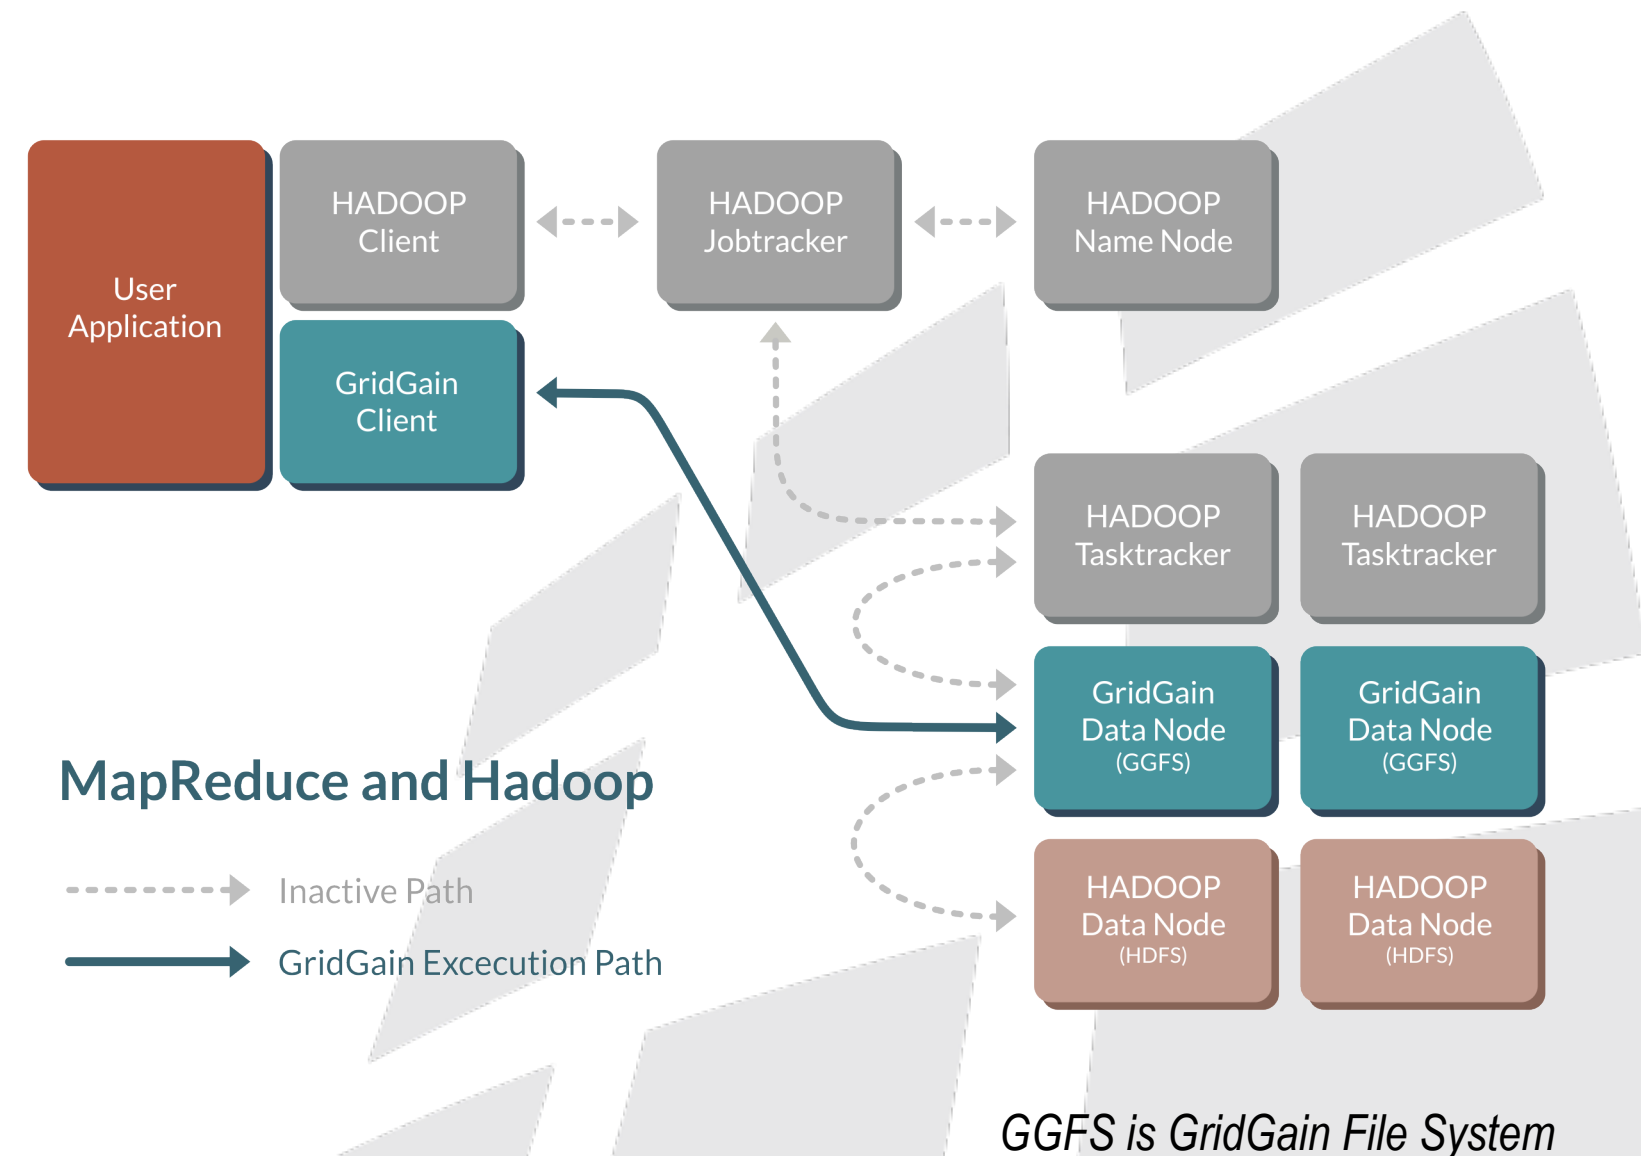
# GridGain Persistent Store

- *Full dataset on disk with some or all of the dataset in memory and ACID transaction and SQL support on full dataset*
- **Features:**
  - *Memory-centric computing*
  - *Distributed, transactional SQL database*
  - *Easily scalable to thousands of nodes*
  - *ACID transaction support*
  - *ANSI SQL-99, DML & DDL across the full dataset*



# GridGain In-Memory Hadoop Acceleration



- *Provides easy to use extensions to disk-based HDFS and traditional MapReduce, delivering up to 10x faster performance*
- **Features:**
  - *100x Faster Performance*
  - *In-Memory MapReduce*
  - *Highly Optimized In-Memory Processing*
  - *Standalone File System*
  - *Optional Caching Layer for HDFS*
  - *Read-Through and Write-Through with HDFS*





# GridGain In-Memory Service Grid

- *Provides control over GridGain and user-defined services on each cluster node and guarantees continuous availability of all deployed services in case of node failures*
- **Features:**
  - *Automatically Deploy Multiple Instances of a Service*
  - *Automatically Deploy a Service as Singleton*
  - *Automatically Deploy Services on Node Start-Up*
  - *Fault Tolerant Deployment*
  - *Un-Deploy Any of the Deployed Services*
  - *Get Service Deployment Topology Information*
  - *Access Remotely Deployed Service via Service Proxy*





# ANY QUESTIONS?

Thank you for joining us. Follow the conversation.

[www.gridgain.com](http://www.gridgain.com)  
[www.gridgain.com/resources/blog](http://www.gridgain.com/resources/blog)



@gridgain  
#gridgain #inmemorycomputing  
info@gridgain.com