GridGain

IN-MEMORY COMPUTING

In-Memory Computing Principles
and Technology Overview

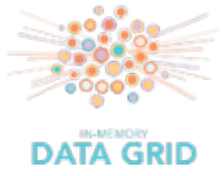

Agenda

- Overview / Why In Memory
- Use Cases
- Concepts & Approaches
- In-Memory Computing Platform 6.1

In-Memory Data Grid
  In-Memory HPC
  In-Memory Streaming
In-Memory Hadoop Accelerator

GridGain

Why In-Memory Computing?

- Cloud/SaaS apps, Mobile Computing back-ends, Internet of Things, Big Data analytics, Social Networks – all need to be done in-memory to reach Internet scale

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

Gartner

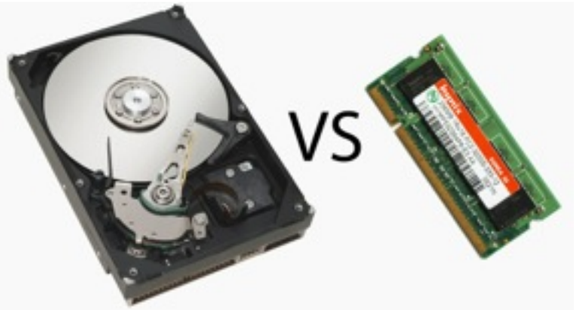

RAM is 3,000 times faster than spinning disks. By moving data from disk to RAM and employing modern in-memory data grid technology, things get fast. Really, really fast.

GridGain

"In-memory computing will have a long term, disruptive impact by radically changing users' expectations, application design principles, products' architectures and vendors' strategies."

Gartner

In-memory computing is the future of computing… it offers a massive potential not only in TCO reduction but across all four value dimensions: performance, process innovation, simplification and flexibility.

Deloitte.

GridGain

Egham, UK, April 3, 2013

View All Press Releases

Gartner Says In-Memory Computing Is Racing Towards Mainstream Adoption

In-Memory Data Grid Market to Reach $1 Billion by 2016

Analysts to Examine the Opportunities and Challenges of IMC at the Gartner Application Architecture, Development & Integration Summit 2013, May 16-17, in London

The rapid maturation of application infrastructure technologies and a continued dramatic decline in the cost of semiconductor technologies are paving the way for mainstream use of in-memory computing (IMC). Gartner, Inc. said that although the in-memory data grid (IMDG)* market, a key IMC segment, is small, it is likely to grow fast and to reach $1 billion by 2016.

"The relentless declines in DRAM and NAND flash memory prices, the advent of solid-state drive technology and the maturation of specific software platforms have enabled IMC to become more affordable and impactful for IT organizations," said Massimo Pezzini, vice president and Gartner Fellow.

Until recently, only the most technologically savvy organizations — for example, in vertical markets like financial trading, telecommunications, military and defense, online entertainment and logistics — could cope with the high cost and complexity of adopting IMC. But IMC technology is now more affordable and more proven. "Organizations that do not consider adopting in-memory application infrastructure technologies risk being out-innovated by competitors that are early mainstream users of these capabilities," said Mr. Pezzini.

GridGain

Egham, UK, April 3, 2013

View All Press Releases

Gartner Says In-Memory Computing Is Racing Towards Mainstream Adoption

In-Memory Data Grid Market to Reach $1 Billion by 2016

Analysts to Examine the Opportunities and Challenges of IMC at the Gartner Application Architecture, Development & Integration Summit 2013, May 16-17, in London

decline in the cost of ting (IMC). Gartner, small, it is likely to

e drive technology and le and impactful for IT

Until recently, only the most technologically savvy organizations — for example, in vertical markets like financial trading, telecommunications, military and defense, online entertainment and logistics — could cope with the high cost and complexity of adopting IMC. But IMC technology is now more affordable and more proven. "Organizations that do not consider adopting in-memory application infrastructure technologies risk being out-innovated by competitors that are early mainstream users of these capabilities," said Mr. Pezzini.

"Organizations that do not consider adopting in-memory application infrastructure technologies risk being out-innovated by competitors that are early mainstream users of these capabilities"

GridGain

In-Memory Computing: Why Now?

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:
- $13.23B in 2018
- 2013-2018 CAGR 43%

MARKETSANDMARKETS

Data Growth

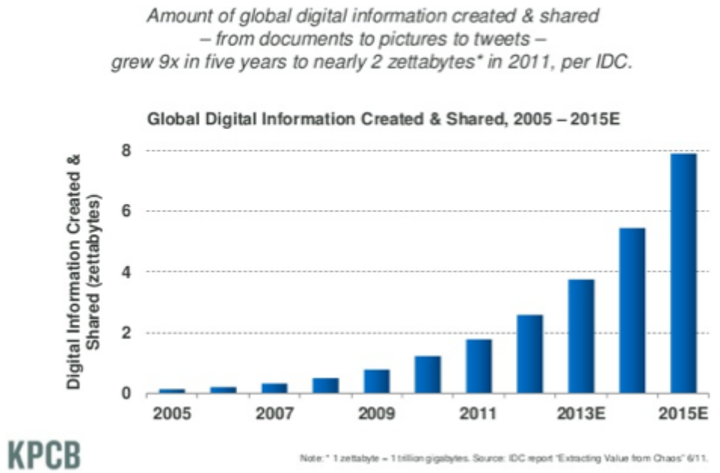

Less than 2 zetabytes in 2011, 8 in 2015

DRAM Cost, $

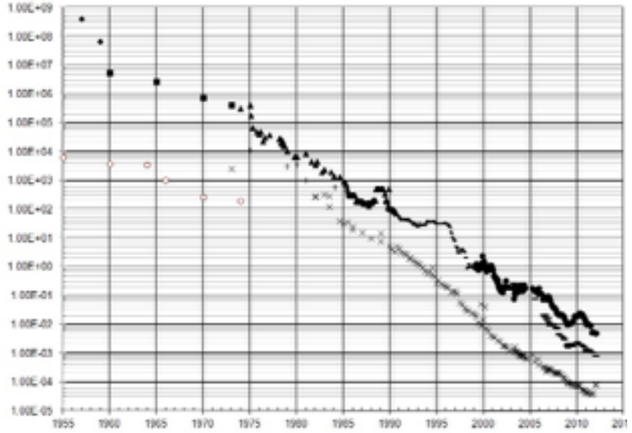

Cost drops 30% every 12 months

BigData Technologies Planned

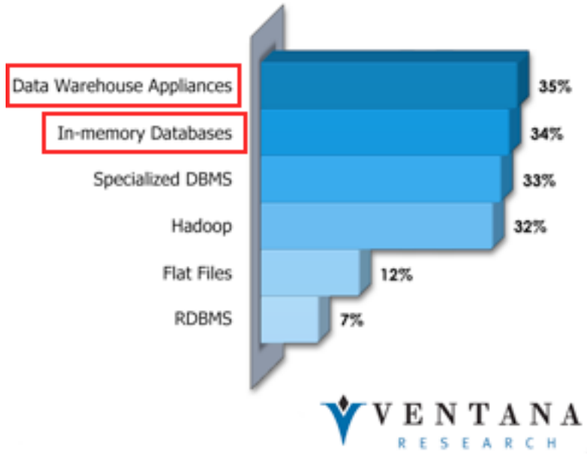

34% will use in-memory technology

GridGain

Top 3 Reasons for In-Memory Computing

1. Performance
2. Performance
3. Performance

GridGain

Top 3 Reasons for In-Memory Computing

1. Performance
2. Scalability
3. Future-proofing

GridGain

How In-Memory Computing Works: The Basic Idea

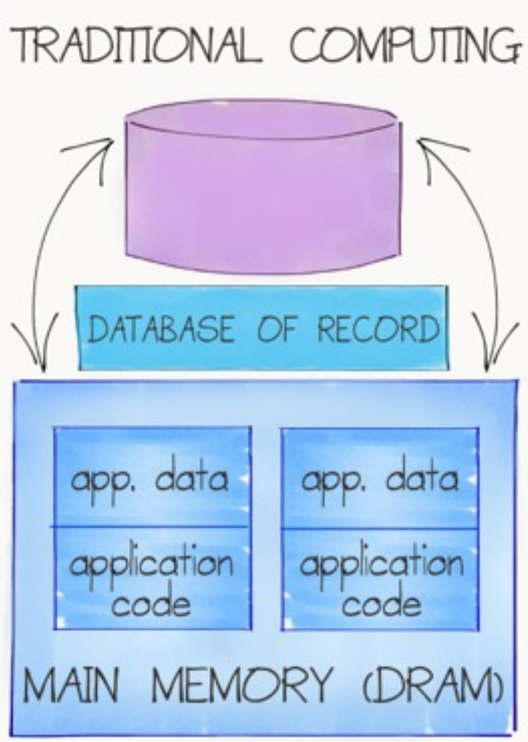

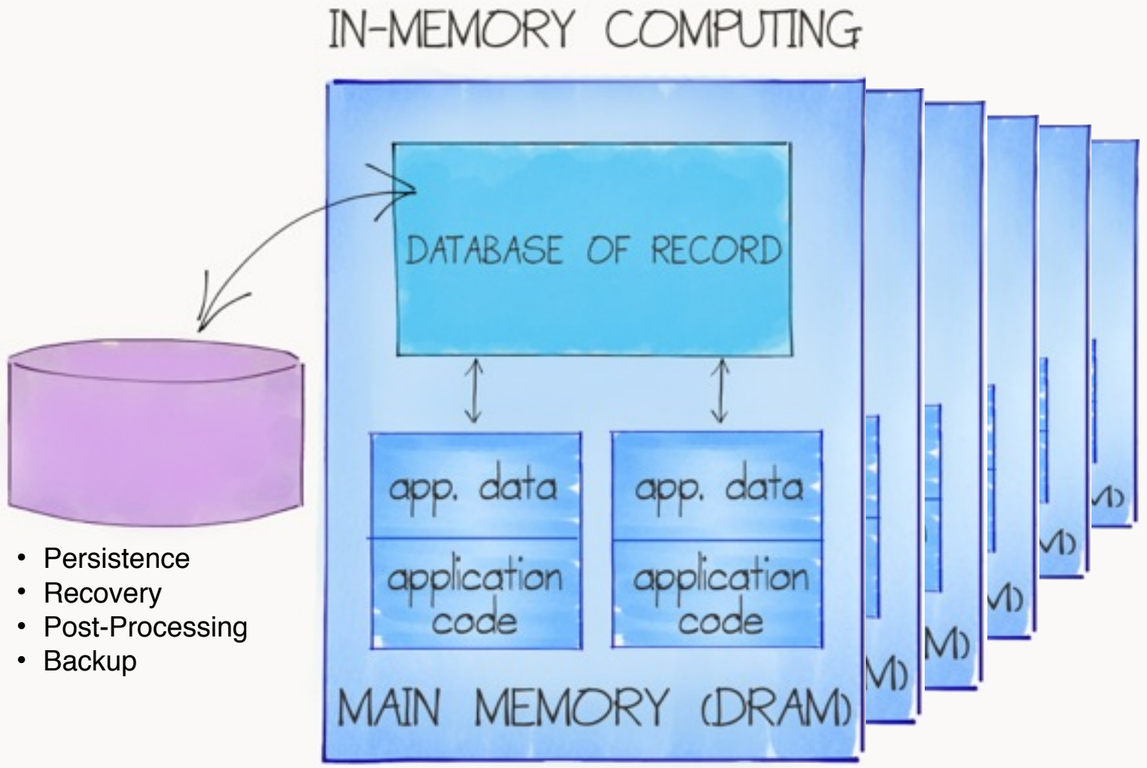

- Persistence
- Recovery
- Post-Processing
- Backup

GridGain

In-Memory Technology: Use Cases

Data Velocity, Data Volume, Real-Time Performance

Automated Trading Systems

Real time analysis of trading positions & market risk. High volume transactions, ultra low latencies.

Financial Services

Fraud Detection, Risk Analysis, Insurance rating and modeling.

Online & Mobile Advertising

Real time decisions, geo-targeting & retail traffic information.

Big Data Analytics

Customer 360 view, real-time analysis of KPIs, up-to-the-second operational BI.

Online Gaming & Mobile Back-ends

Real-time back-ends for mobile and massively parallel games and mobile applications.

Bioinformatics & Sciences

High performance genome data matching. Environmental simulation.

GridGain

The Evolution of Architectures

Traditional Architecture

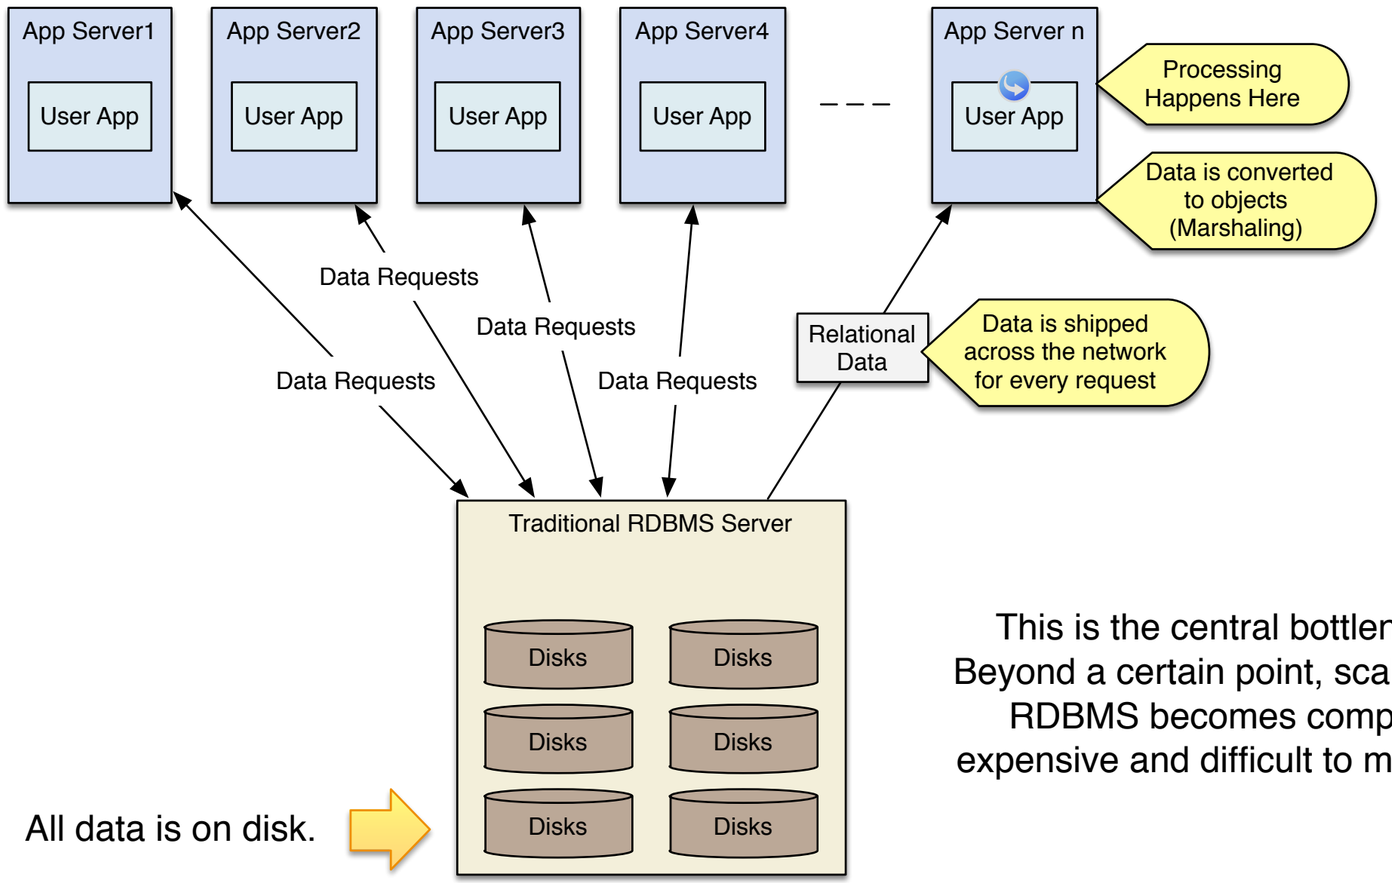

This is the central bottleneck. Beyond a certain point, scaling the RDBMS becomes complex, expensive and difficult to manage.

GridGain

Horizontally Scale the RDBMS

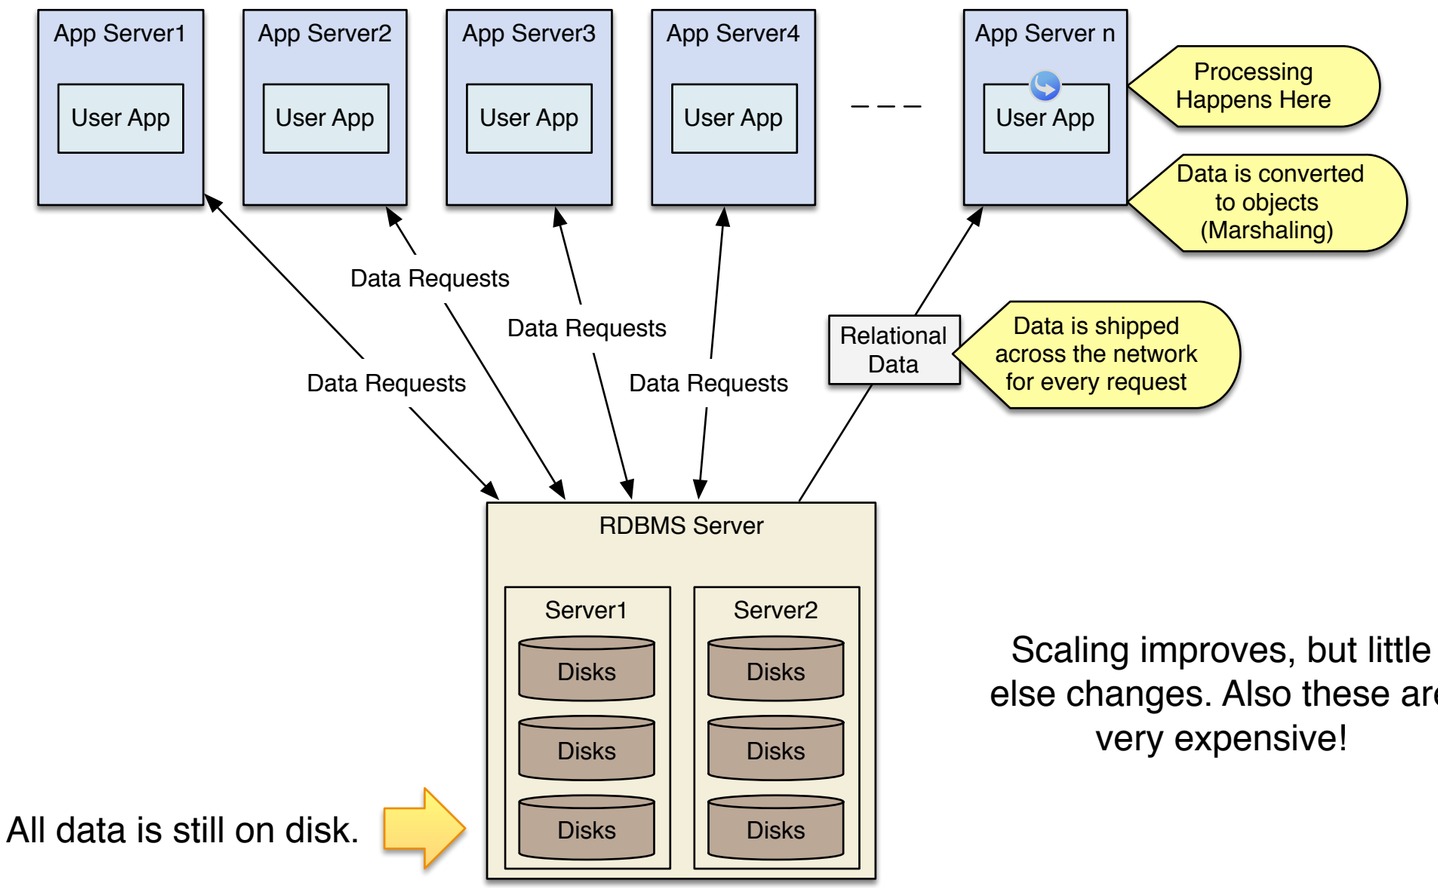

All data is still on disk.

Scaling improves, but little else changes. Also these are very expensive!

GridGain

IMDG: Distributed Caching

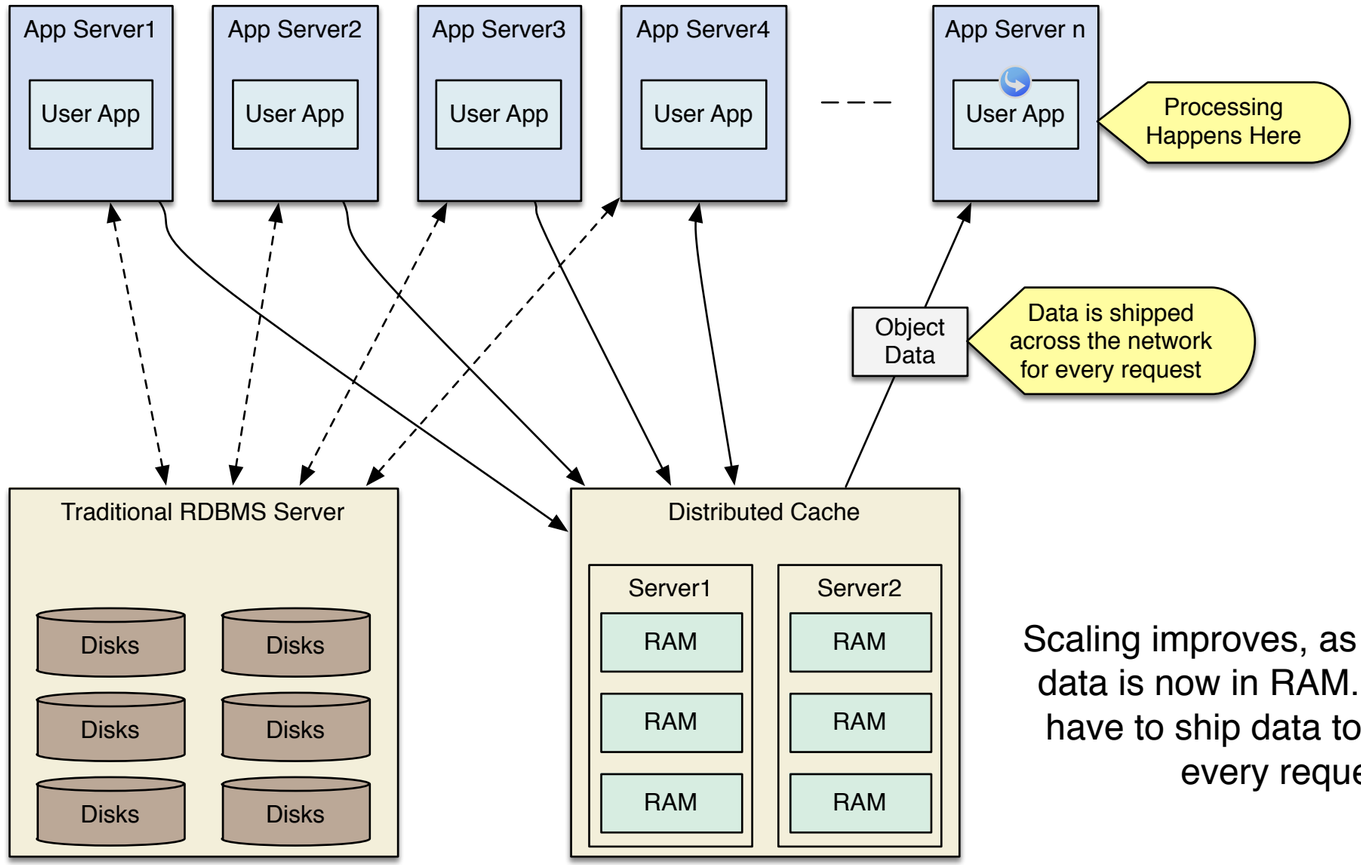

Scaling improves, as some (or all) data is now in RAM. But, we still have to ship data to the app for every request.

GridGain

GridGain: IMDG + Grid Compute

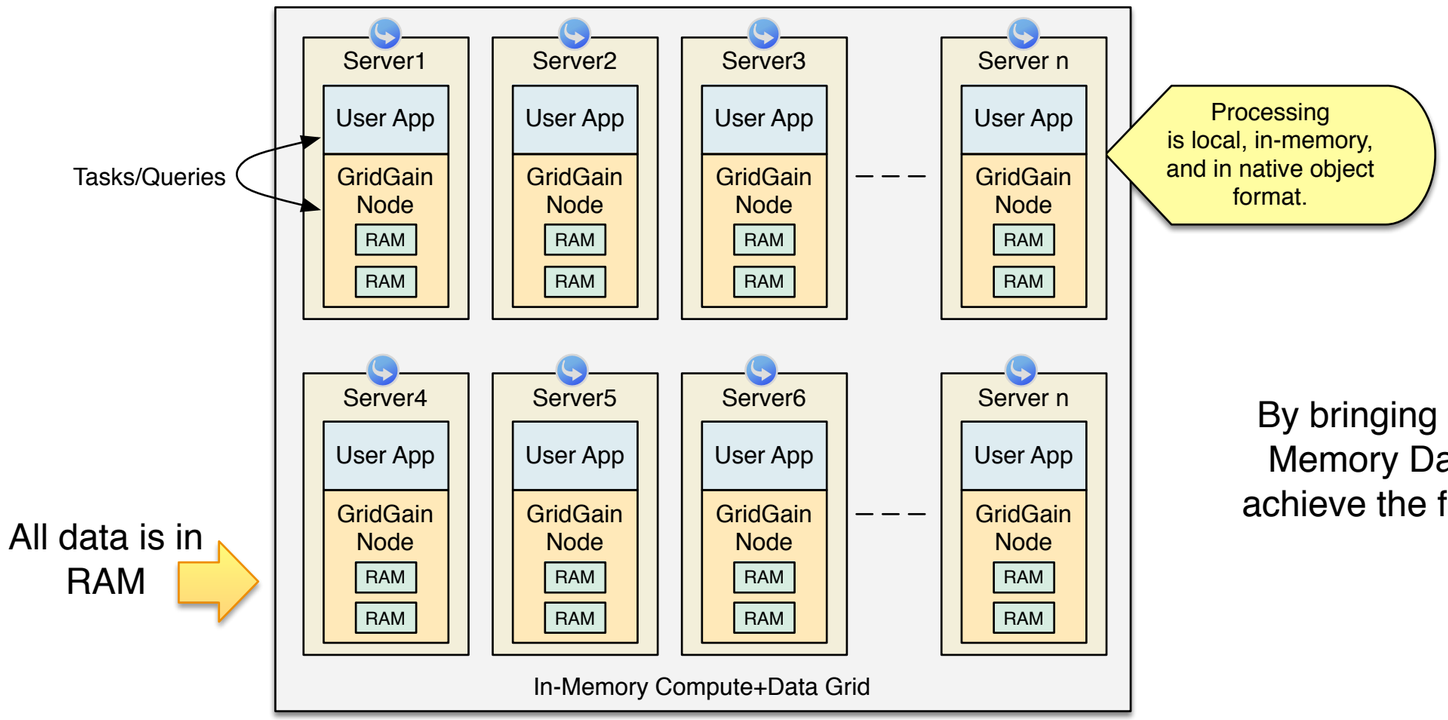

By bringing computation to the In-Memory Data Grid, you can now achieve the fastest possible results.

GridGain

In-Memory Data Grids: Considerations

1. Performance
2. Scaling:
   Vertically, Horizontally,
   and without downtime
3. Consistency
4. Persistence
5. Transactions
6. Search/Query

GridGain

Introduction to GridGain

GridGain: Company Facts

Facts:

- Founded in 2007

  Founders:
  Nikita Ivanov
  Dmitriy Setrakyan

- Locations:

  HQ: Foster City, CA
  R&D: Silicon Valley
  Saint Petersburg, Russia

Technology:

- 5+ years in production
- Starts every ~10 sec. worldwide
- Over 18M starts globally

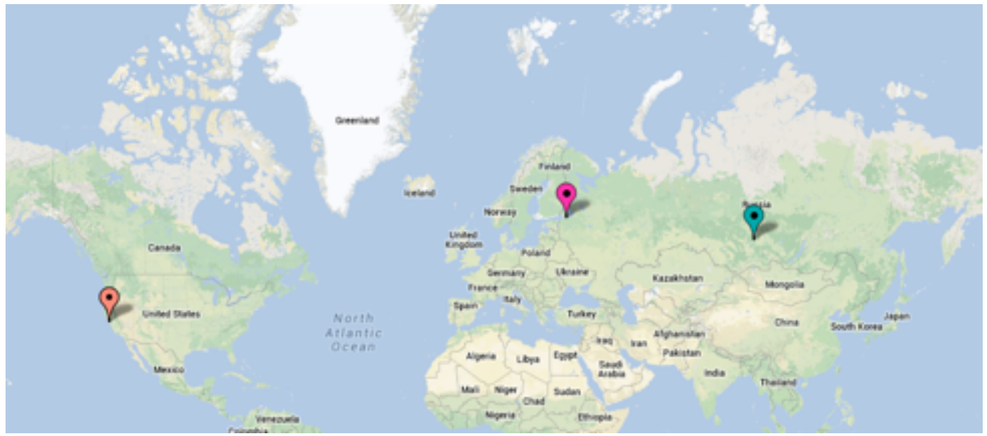

GridGain

Customer Use Cases

Data Velocity, Data Volume, Real-Time Performance

> Automated Trading Systems

Real time analysis of trading positions & market risk.
High volume transactions, ultra low latencies.

> Financial Services

Fraud Detection, Risk Analysis, Insurance rating and modeling.

> Online & Mobile Advertising

Real time decisions, geo-targeting & retail traffic information.

> Big Data Analytics

Real time analysis of inventory & purchasing
Operational up-to-the-second BI.

> Online Gaming

Real-time back-ends for mobile and massively parallel games.

> Bioinformatics

High performance genome data matching.

GridGain

Use Case: SBERBANK

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

- Open tender won by GridGain
  - Goal: Real-time risk and leverage reporting on their global financial trading portfolio
  - Performed a detailed evaluation and software assurance test
  - Delivered best performance, scale and high availability

1 Billion
Transactions per Second

---

10 Dell R610 servers
1 TB Memory

< $45K

GridGain

In-Memory Computing Platform

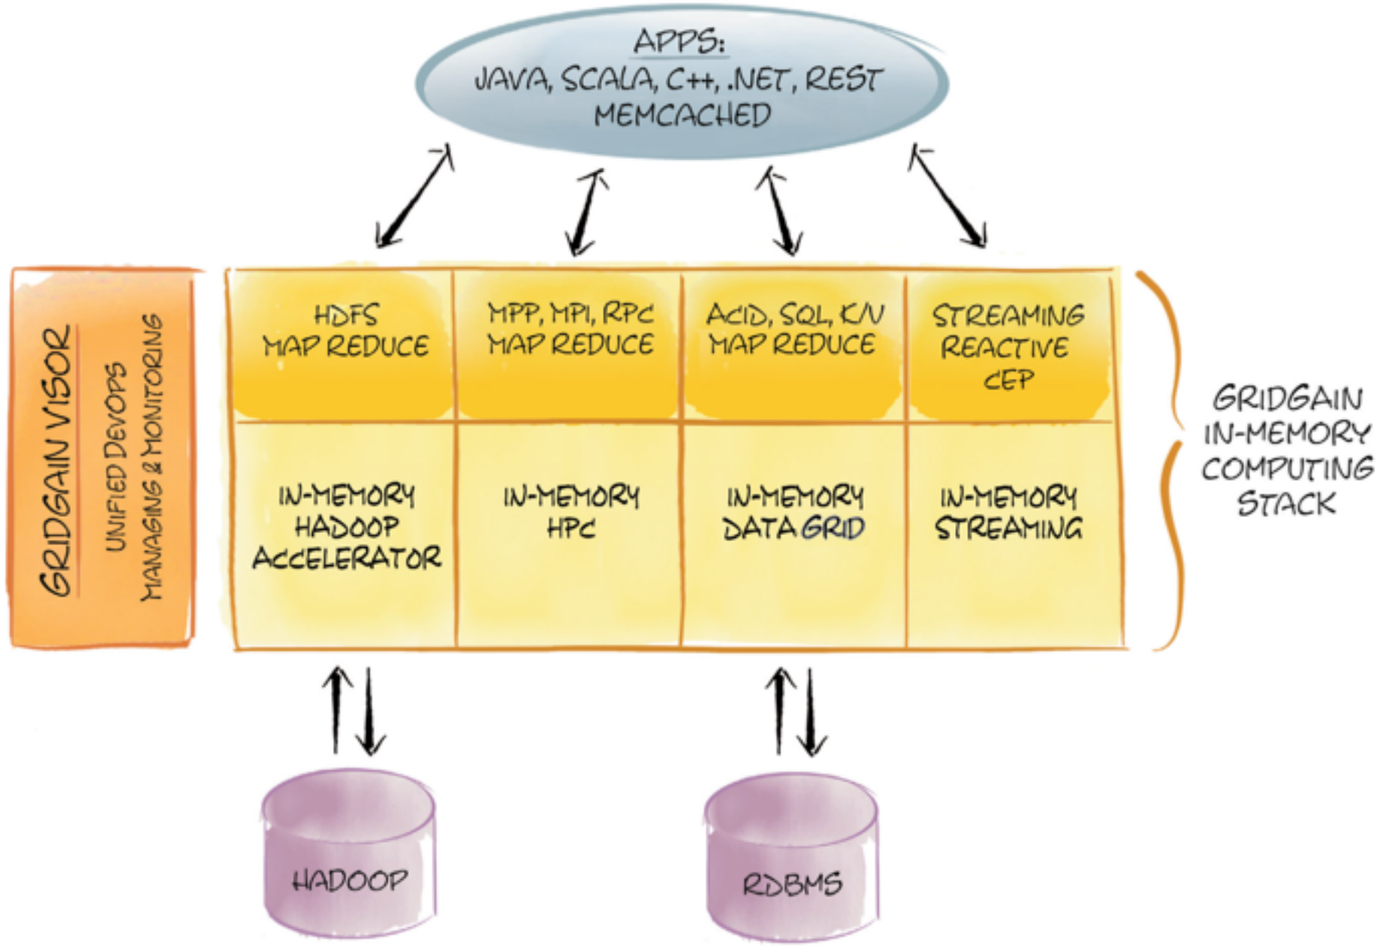

GridGain

In-Memory HPC: Key Features (all editions)

- Direct API for MapReduce
- Support for MPP, MPI, RPC
- Zero Deployment
- Cron-like Task Scheduling
- State Checkpoints
- Distributed Task Sessions
- AOP support
- Early and Late Load Balancing
- Advance Fault Tolerance
- Full Cluster Management
- Advanced Features
  - Redundant mapping support
  - Partial asynchronous reduction
  - Distributed continuations
- Pluggable SPI design
  - 14 pluggable subsystems
  - Discovery, communication, deployment, security…
- Colocation of Compute + Data
  - 100% integration with IMDG

GridGain

In-Memory HPC (all editions)

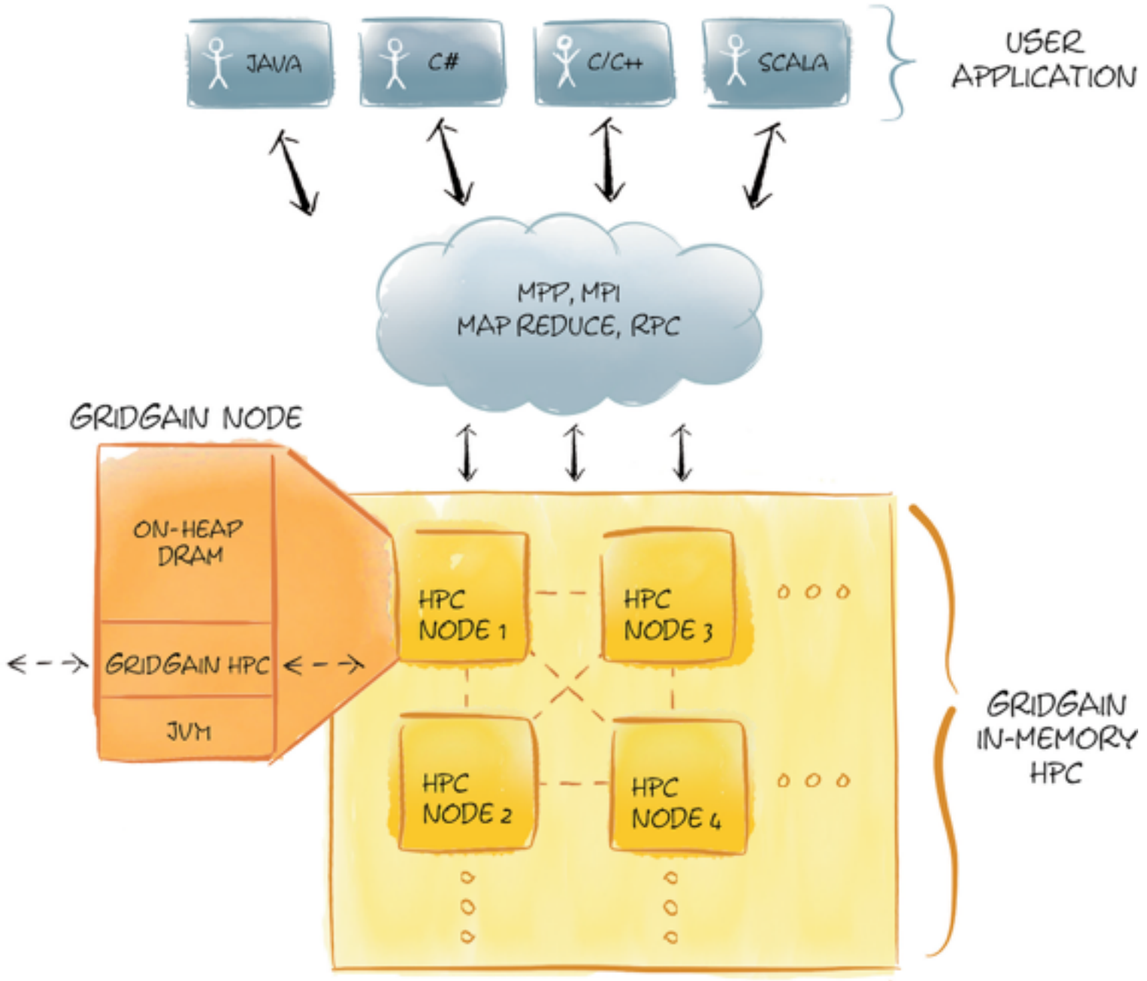

GridGain

In-Memory HPC: MPP & MapReduce (all editions)

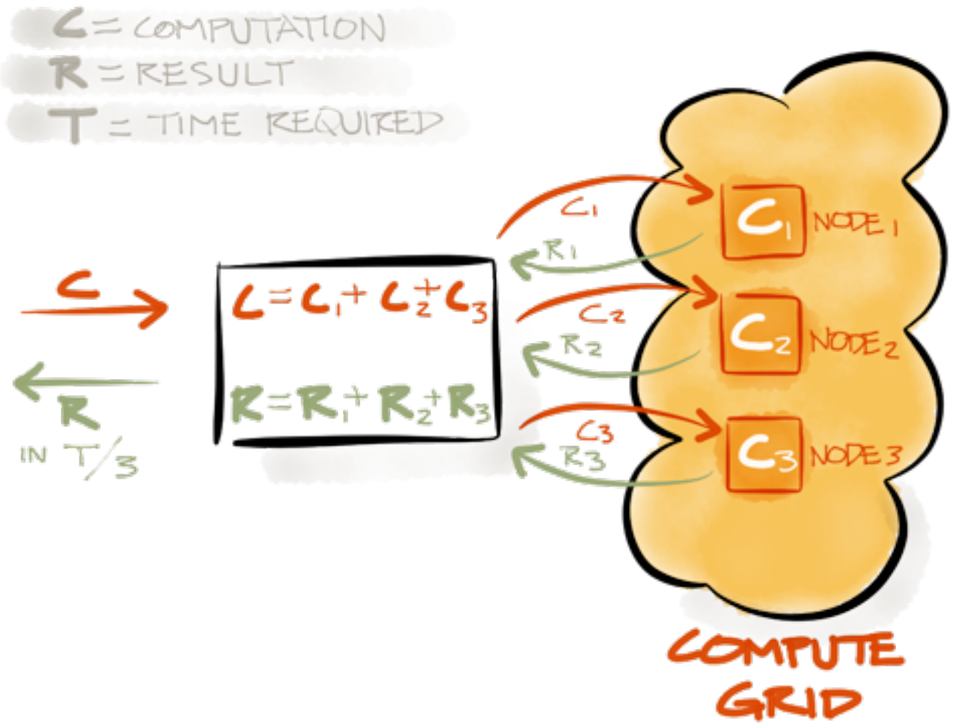

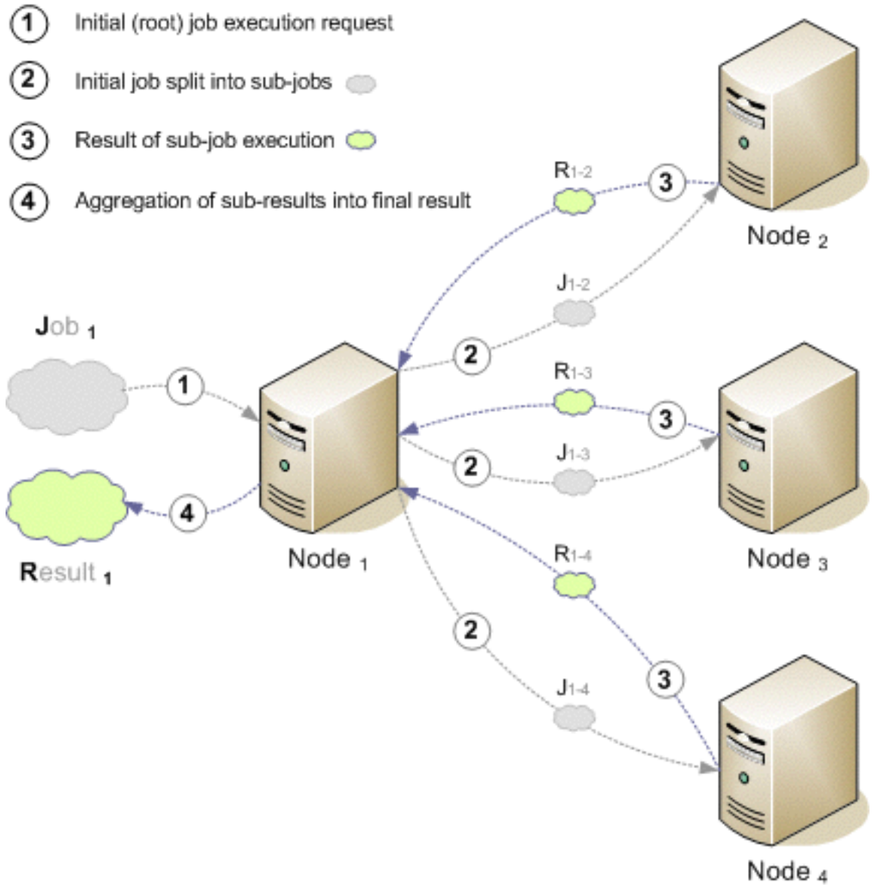

GridGain

In-Memory HPC

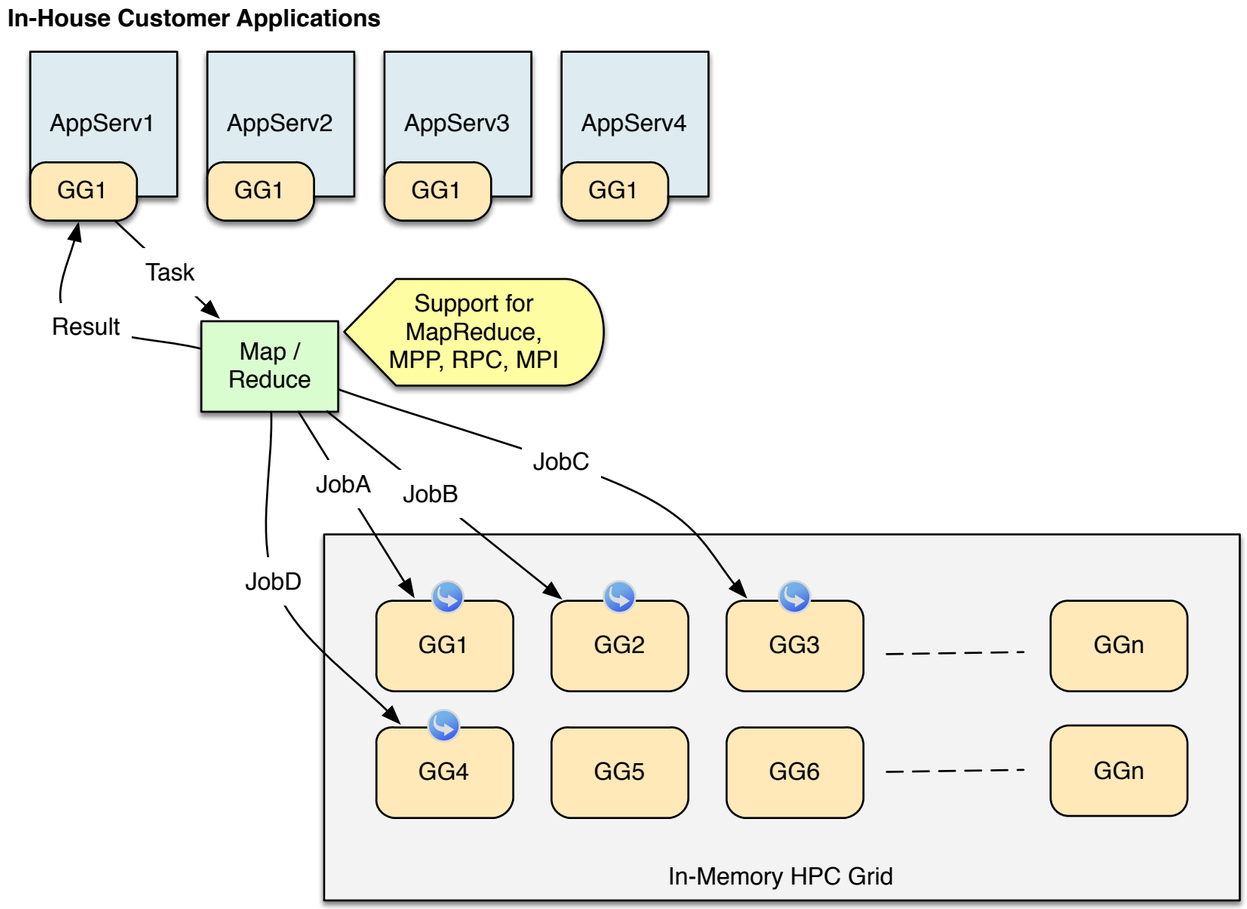

GridGain

In-Memory HPC: Auto-Failover

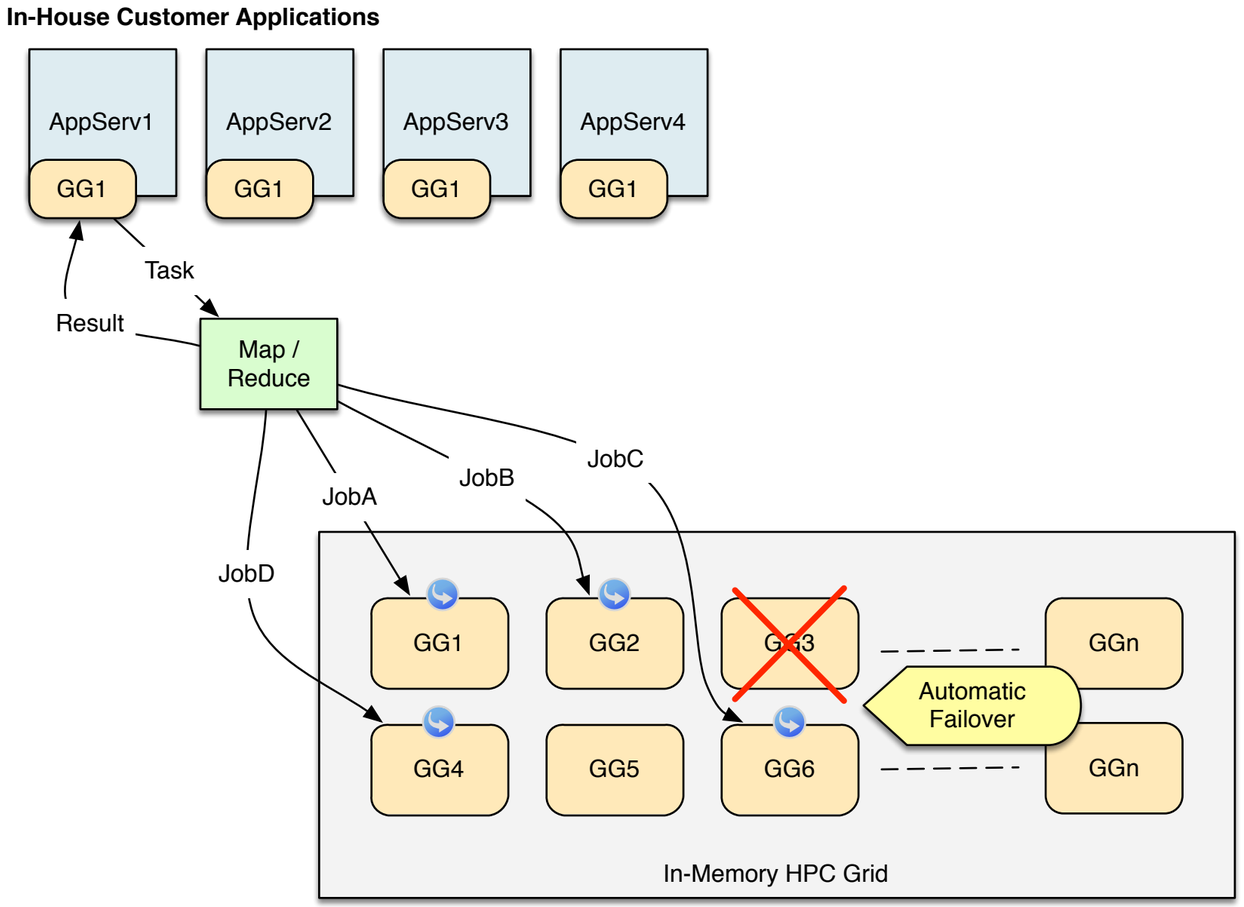

GridGain

Deployment Options

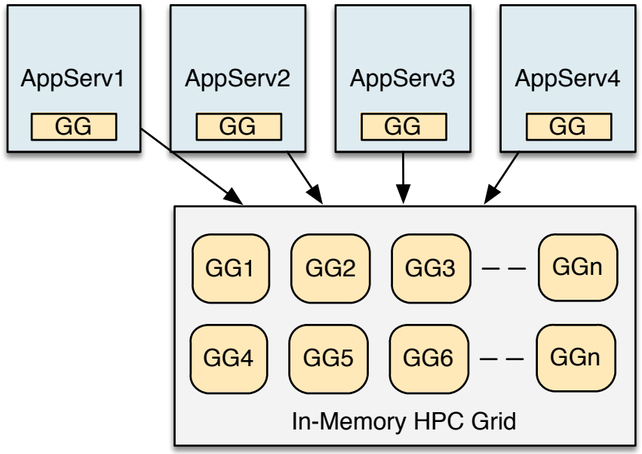

Pseudo Client/Server

Co-Deployed Grid
(Internal Deployment)

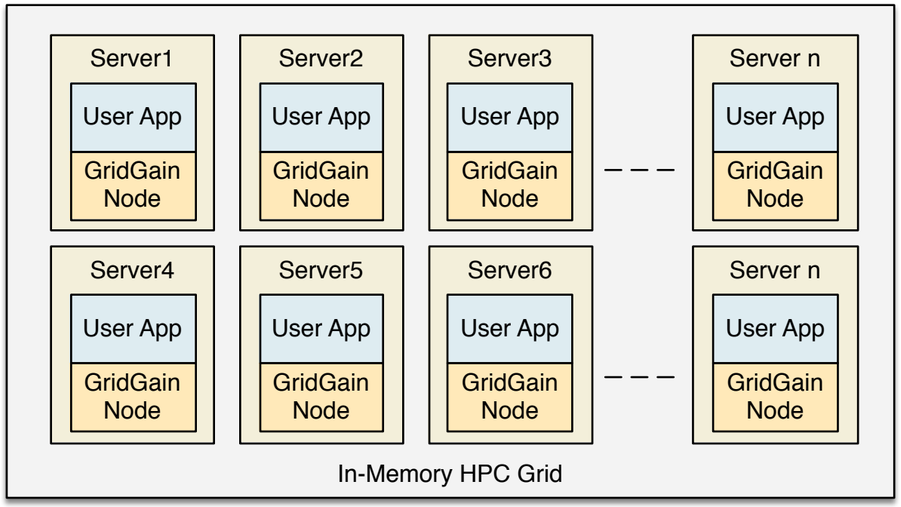

GridGain

In-Memory Computing Platform

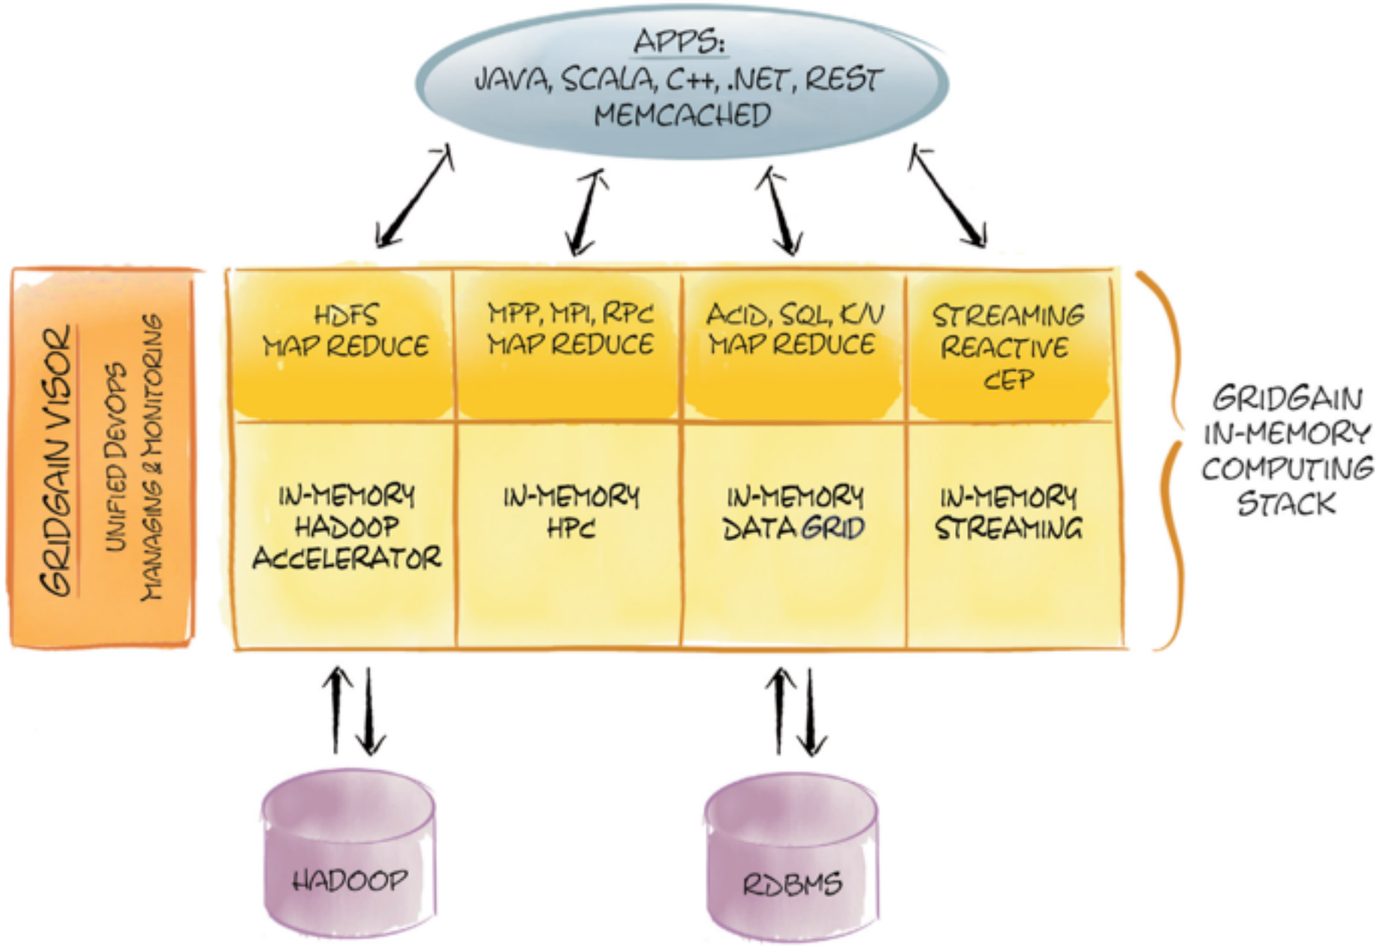

GridGain

In-Memory Data Grid: Key Features

- Distributed KV store of objects
- TBs of data, of any type
- Vertical & Horizontal Scale
  - Use Heap and Off-Heap Memory
  - Dynamic topology management
  - Tiered storage model
- High-Availability
  - Active replicas, automatic failover
- Data Consistency
  - ACID distributed transactions
- SQL queries and JDBC driver
  - Advanced indexing
- Colocation of Compute + Data
  - Affinity Routing
- REST, Memcached, SQL, MapReduce
- Synchronous and Asynchronous ops
- Local, Replicated, and Partitioned
- Pluggable Implementations
  - Expiration, overflow, indexing
- Advanced Features
  - Preloading, HyperLocking, datacenter replication

GridGain

In-Memory Data Grid

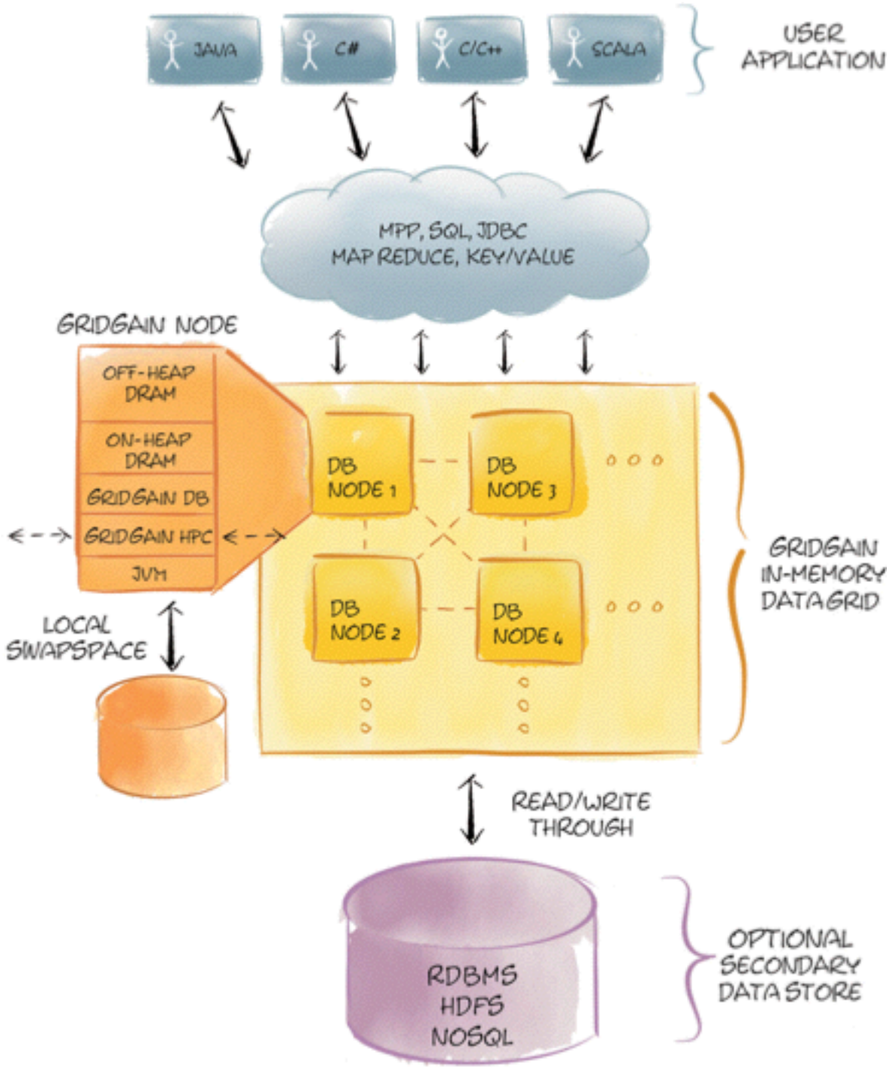

GridGain

In-Memory Data Grid

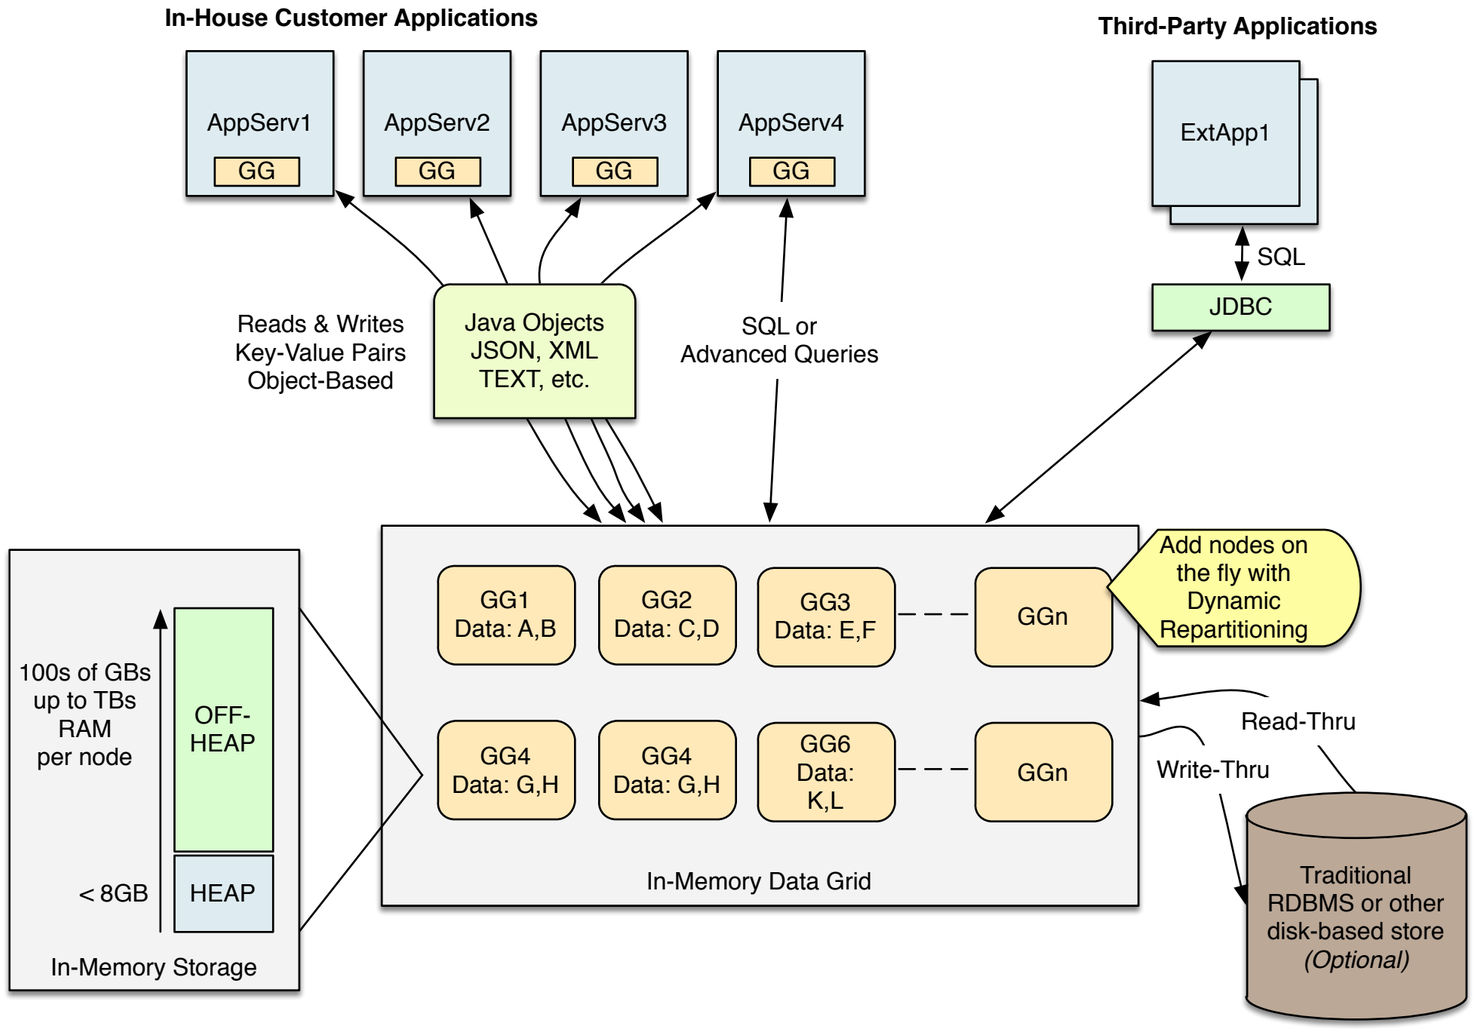

GridGain

In-Memory Data Grid: Compute + Data

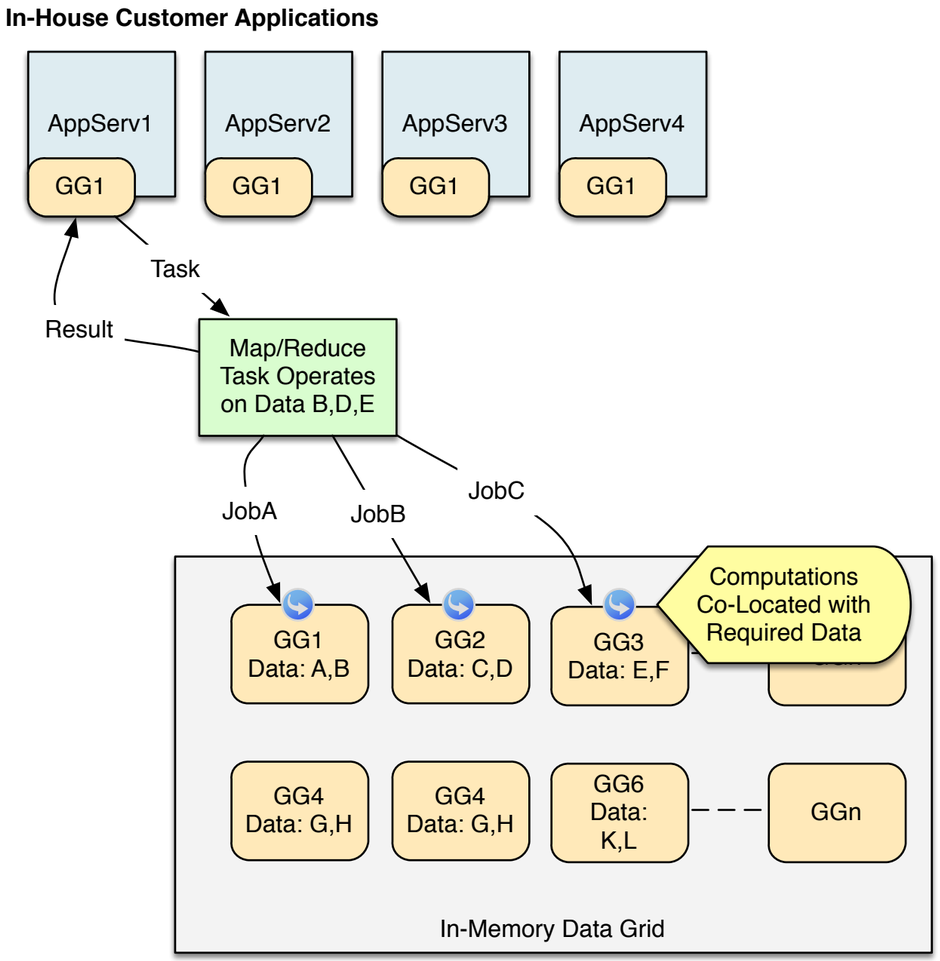

GridGain

In-Memory Data Grid: Affinity Co-Location

Other solutions

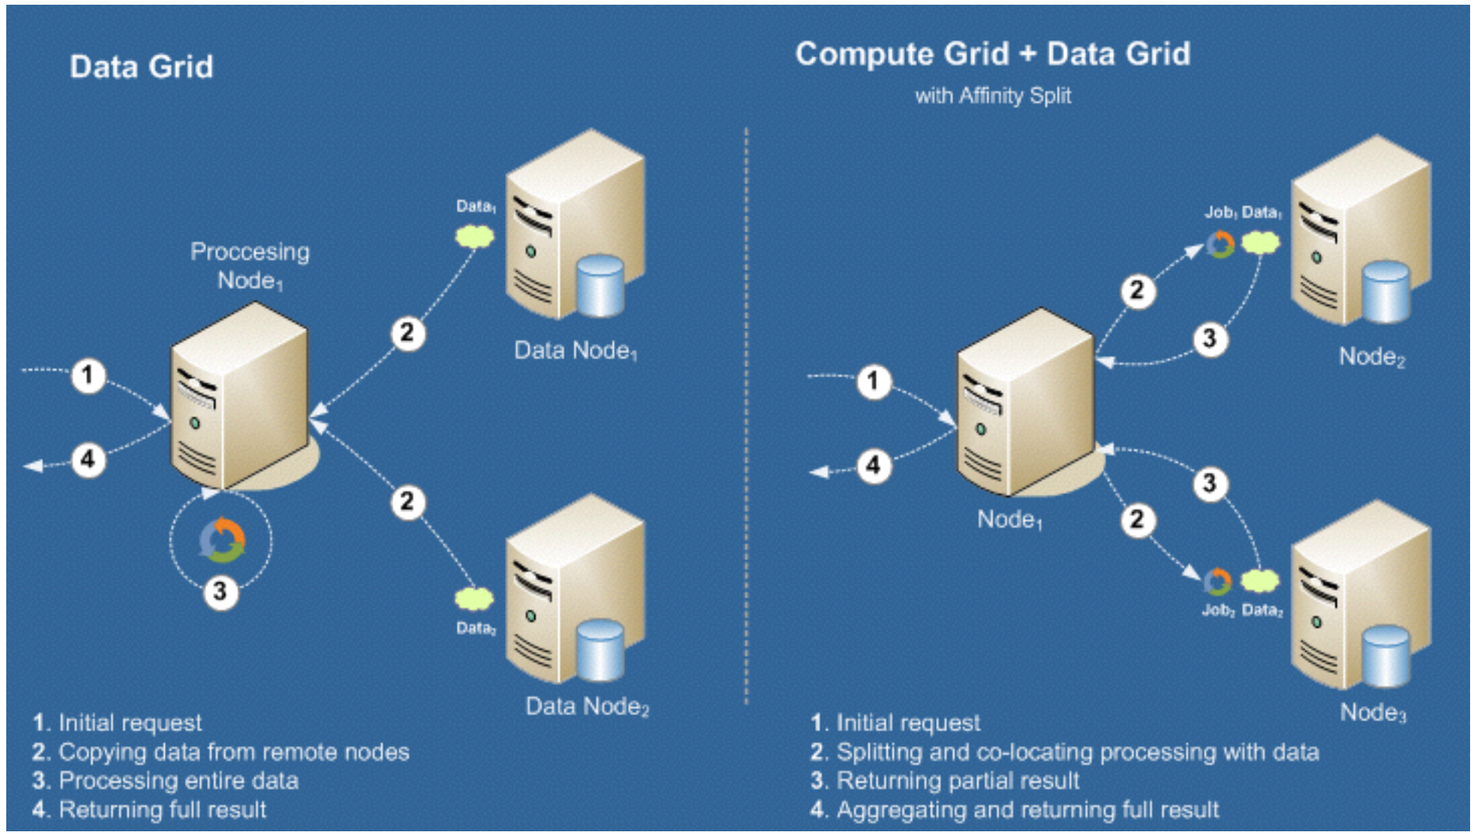

GridGain IMDG

GridGain

In-Memory Computing Platform

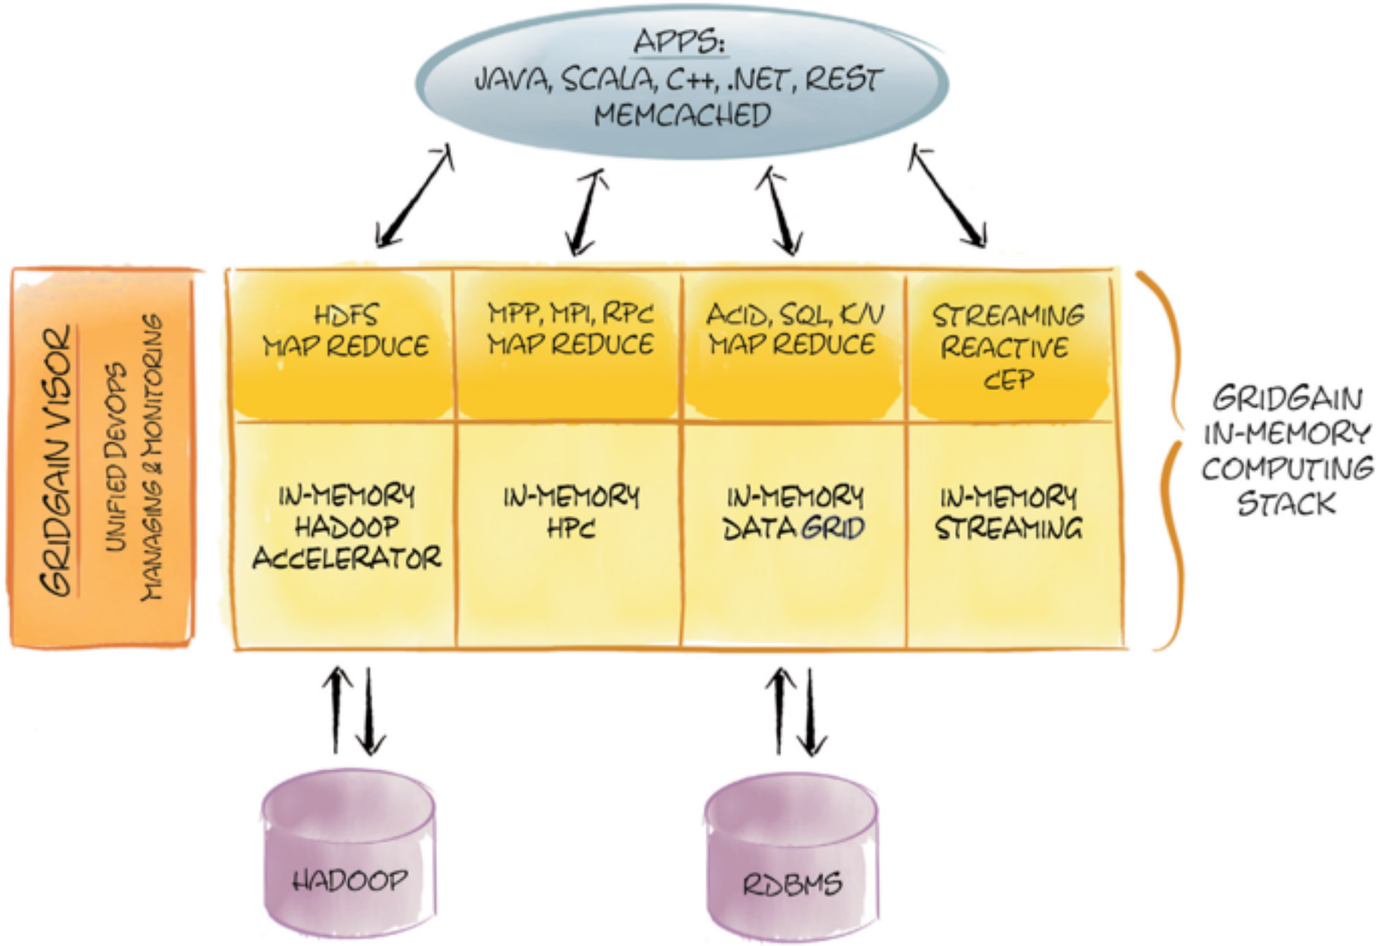

GridGain

In-Memory Streaming: Key Features

- Streaming Data Ingest
- Branching Pipelines
- Pluggable Routing
- Real-Time Analysis
- Configurable Data Windows
- CEP / Continuous Query
- Streaming Data Indexing
- 1000+ Nodes Scalable
- Full GridGain Integration
- Optimized for high-velocity data

GridGain

In-Memory Streaming

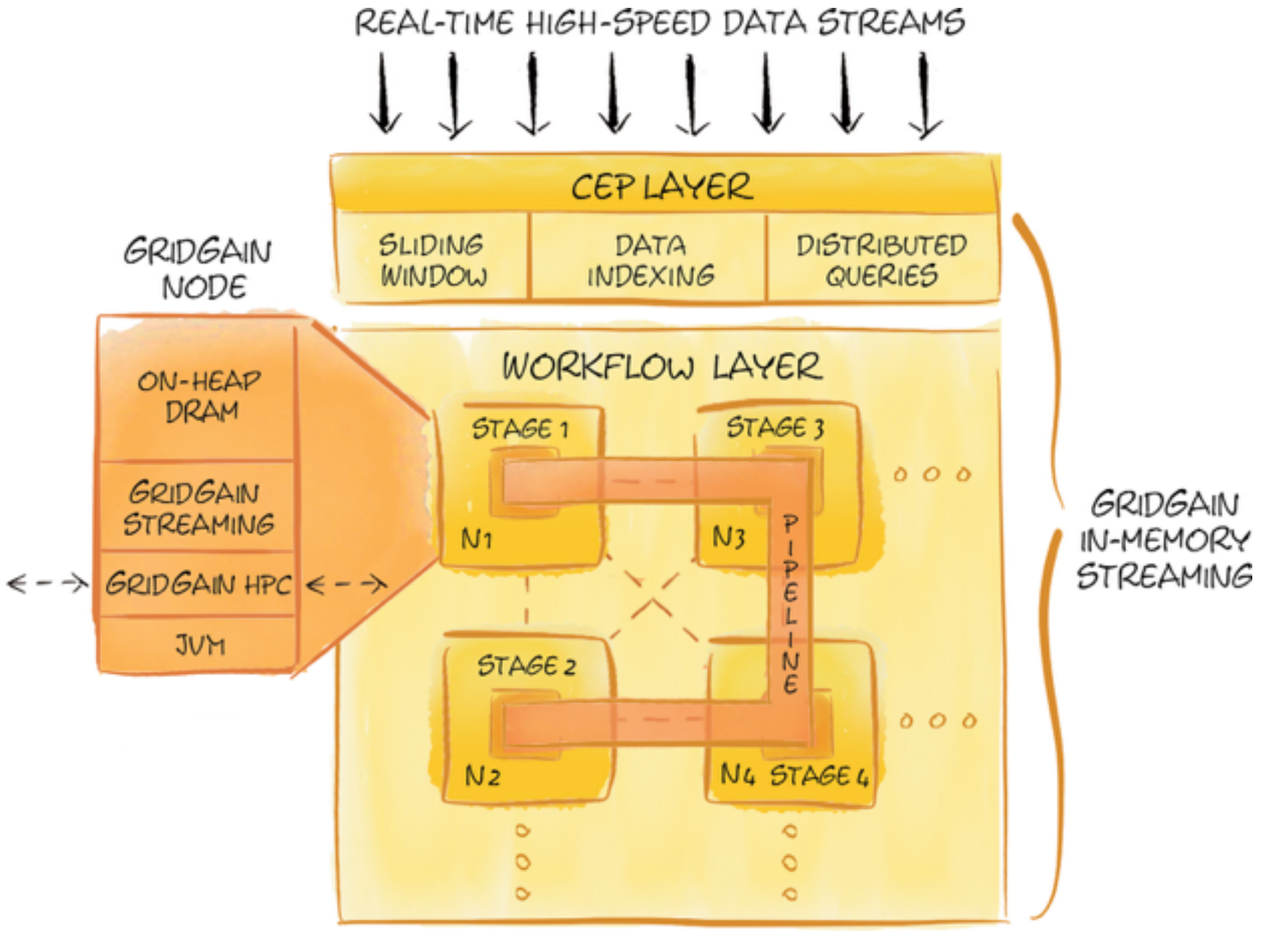

GridGain

In-Memory Streaming

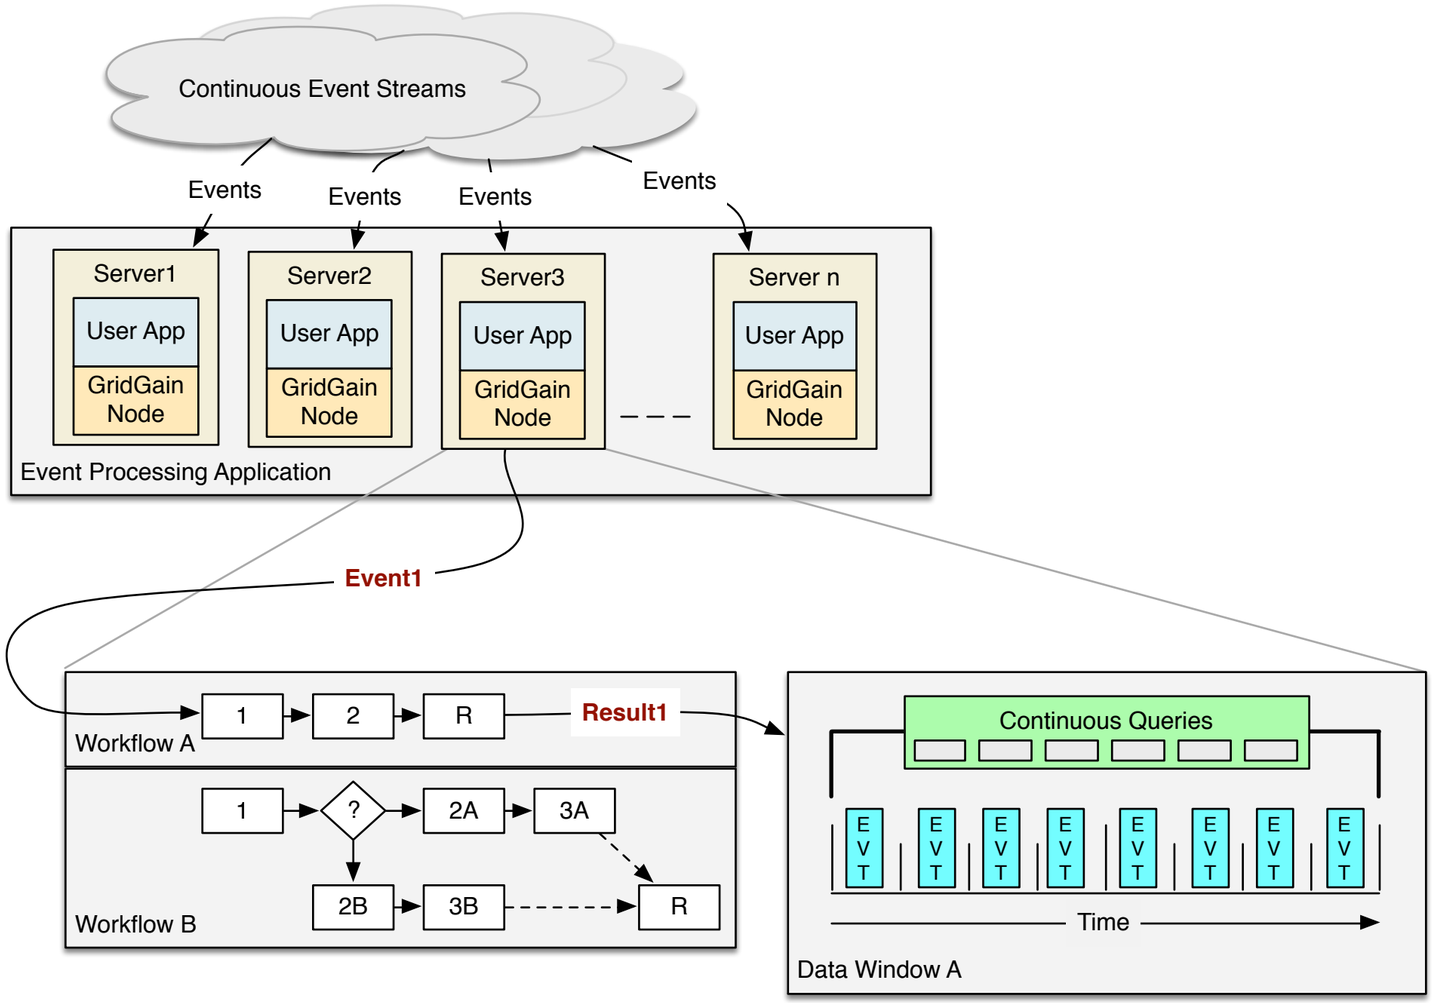

GridGain

In-Memory Computing Platform

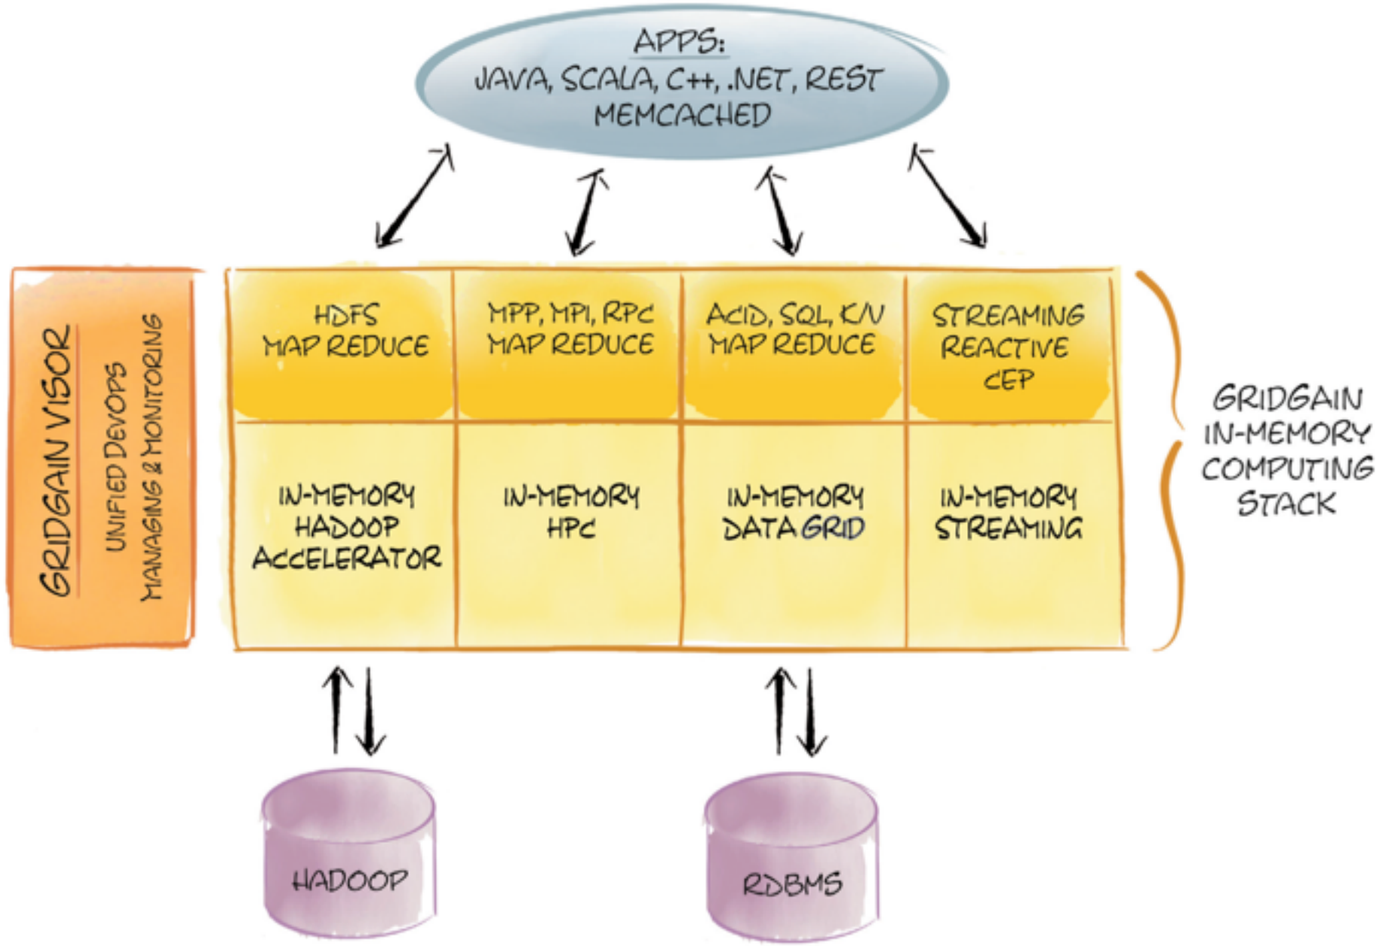

GridGain

In-Memory Hadoop Accelerator: Key Features

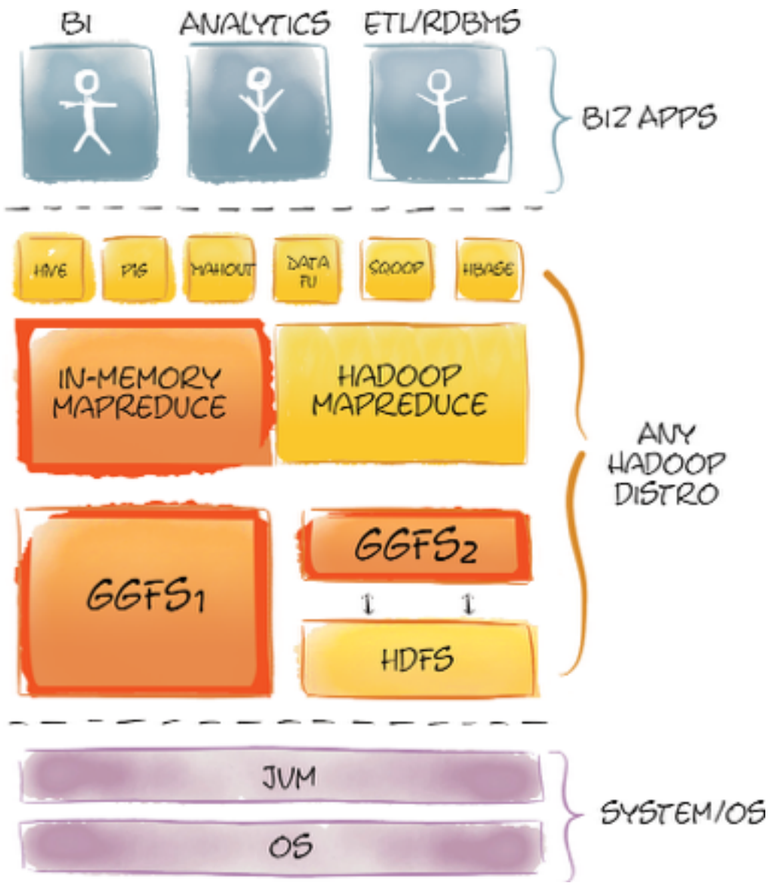

- **Boost HDFS Performance**
  Dual-mode In-Memory File System (GGFS)
  100% HDFS compatible
- **Eliminate Hadoop MapReduce Overhead**
  In-Memory MapReduce
  In-process execution & data co-location
  Record-based vs. K/V
- **Any Hadoop 1.x & 2.x Distributions**
- **Plug & Play: Works with existing Hadoop applications**
- **No ETL Required**
  Read-through, write-through from and to HDFS
- **Speed Up Any MapReduce Jobs**
- **GUI Management**

GridGain

In-Memory Hadoop Accelerator: GridGain File System

Benchmark	GGFS ms	HDFS ms	Boost %
File Scan	27	667	2,470%
File Create	96	961	1,001%
File Random Access	413	2,931	710%
File Delete	185	1,234	667%

GridGain

In-Memory Hadoop Accelerator: MapReduce

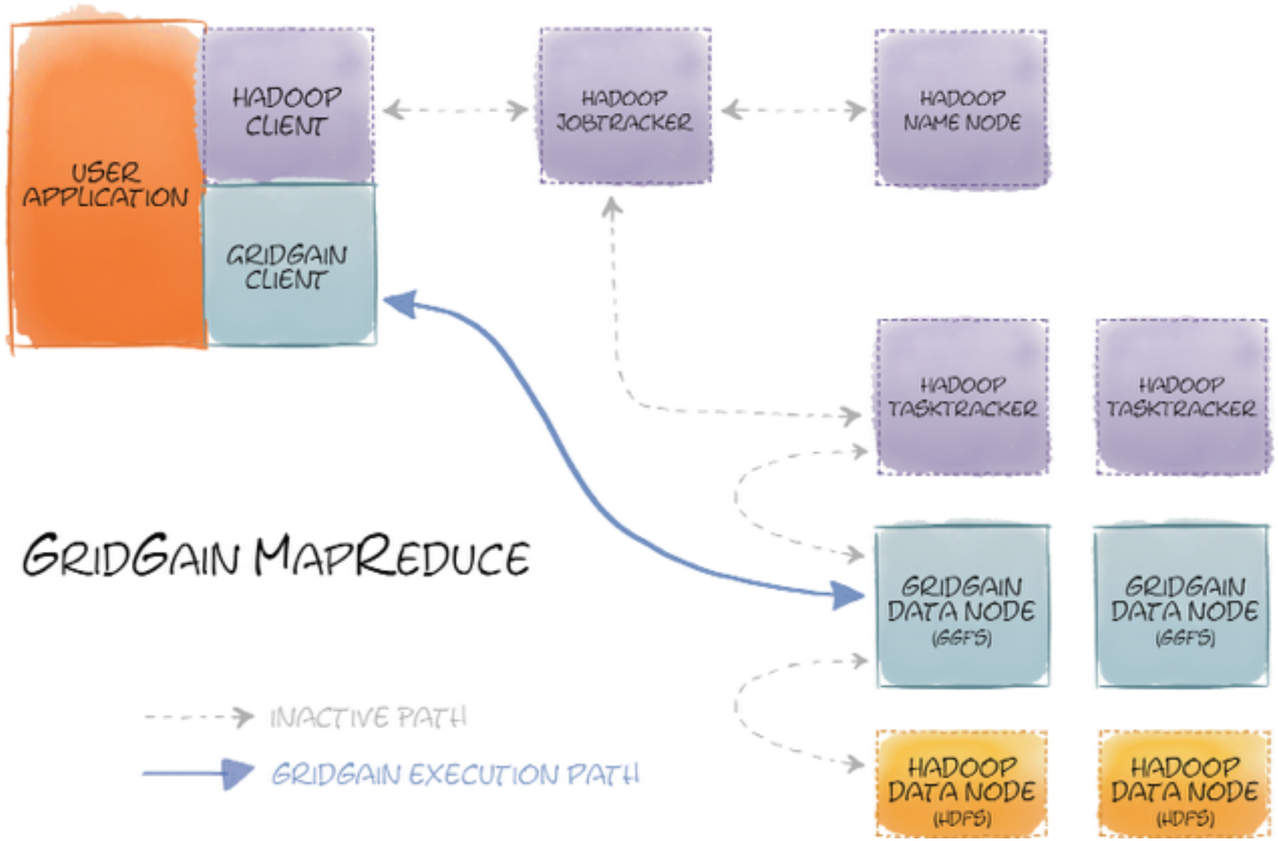

GridGain

GridGain Visor: Unified DevOps

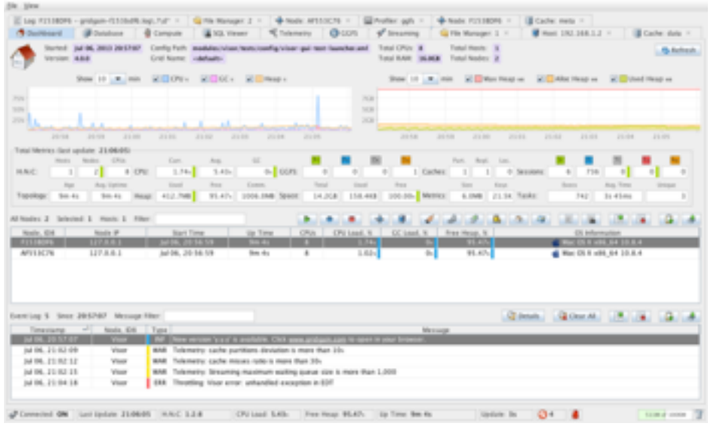

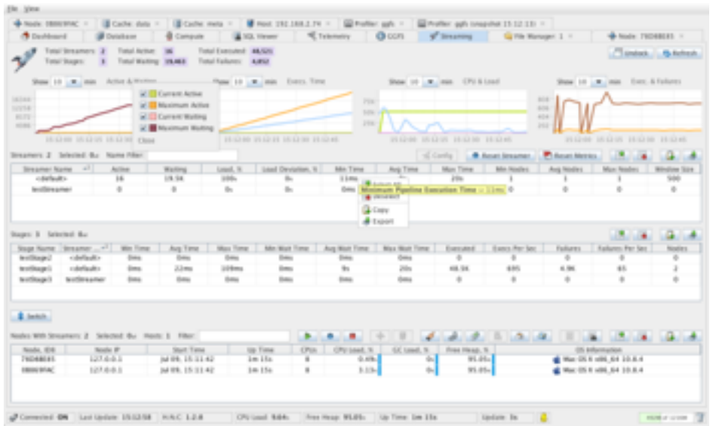

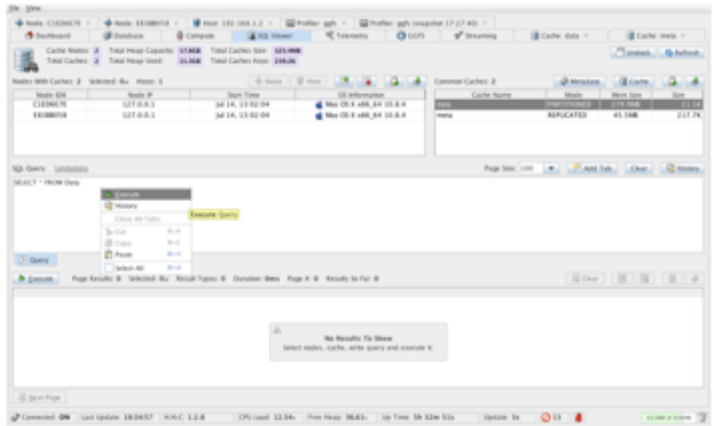

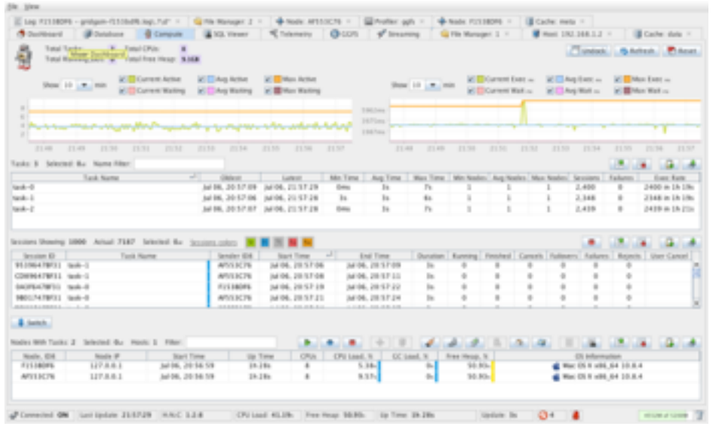

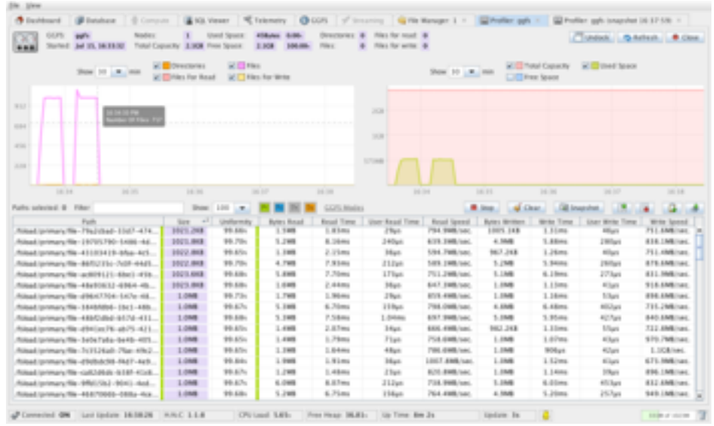

GridGain

Enterprise Edition: Exclusive Features

- **Management & Monitoring**
  GridGain Visor for GUI-based DevOps
- **Local Restartable Store**
  fast recovery during planned outages or DR
- **Data Center Replication**
  Multi-datacenter WAN support
- **Network Segmentation Protection**
  Configurable fault-tolerance for network interruptions
- **Security Features**
  Client Authentication and related SPIs
- **Rolling Production Updates**
  Perform software upgrades without downtime
- **Support & Maintenance**
  Support incidents, ticket access, upgrades, patches
- **Training & Consulting**
  Technical training and customized consulting services
- **Deploy with Confidence**
  Indemnification for Enterprise Customers

GridGain

Enterprise & Open Source Comparison Chart

Features	Open Source	Enterprise Edition
In-Memory Data Grid	√	√
In-Memory Streaming	√	√
In-Memory HPC	√	√
In-Memory Accelerator for Hadoop	√	√
Local Restartable Store		√
Security Features & Updates		√
Network Segmentation Protection		√
Rolling Production Updates		√
Management & Monitoring Visor		√
Data Center Replication		√
9x5 and 24x7 Support		√

In summary, GridGain Enterprise Subscriptions will include the following during the term of an Enterprise Subscription:

- Right to Use the Enterprise Edition of the GridGain product.
- Bug fixes, patches, updates and upgrades to the latest version of the product.
- 9x5 or 24x7 Support for the product.
- Ability to procure Training and Consulting Services from GridGain.
- Confidence and protection, not provided under Open Source licensing, that only a commercial vendor can provide, such as Indemnification.

GridGain

GridGain

IN-MEMORY COMPUTING

GridGain Systems
www.gridgain.com

1065 East Hillsdale Blvd, Suite 230
Foster City, CA 94404