



Apache Ignite and Apache Spark

Where Fast Data Meets the IoT



Denis Magda
Ignite PMC Chair
GridGain PM

Agenda


- IoT Demands to Software
- IoT Software Stack
 - Device OS/RTOS
 - Data Collection and Enrichment
 - HTAP Platform
 - Application APIs
- Demo

IoT Demands to Software

Real-time Processing

Analytics (BI, ML)

Simple Scalability



SQL, Geo-Spatial

High-Availability

IoT Software Stack

Application APIs

HTAP Platform

Data Collection and Enrichment

Device OS/Real-Time OS

Apache IoT Software Stack

Application APIs



HTAP Platform



Data Collection and Enrichment



kafka



Device OS/Real-Time OS



Apache MyNewt

Open Source RTOS
Cortex M, MIPS




Bluetooth, Wifi,
TCP/IP



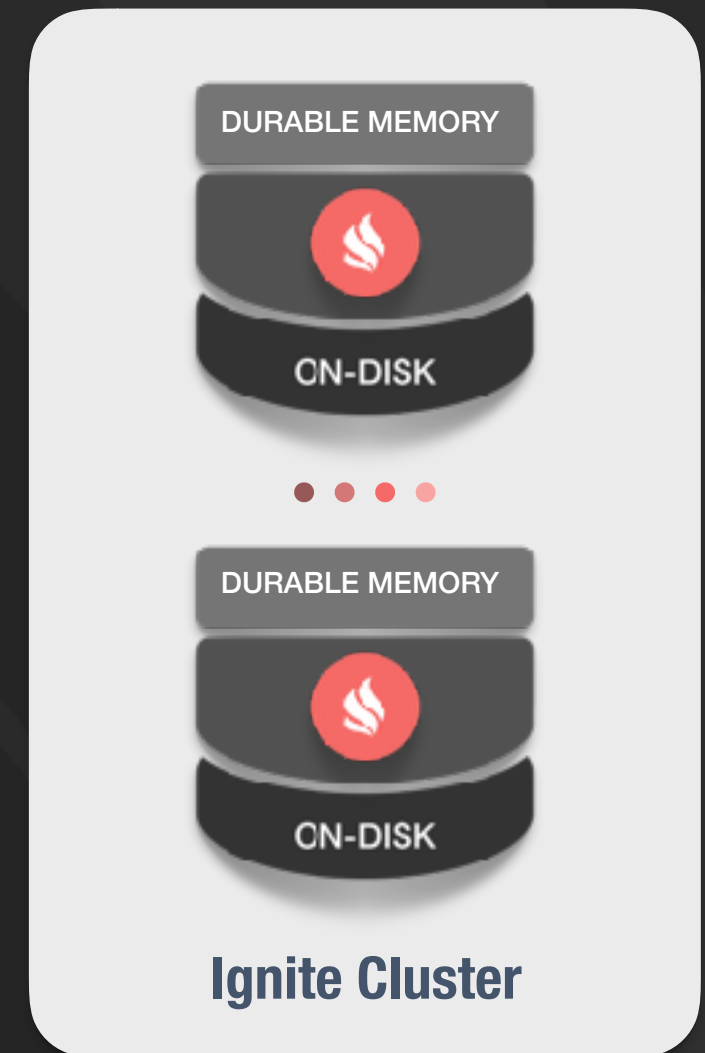



Secure Bootloader



Remote Firmware
Upgrade



Data Collection and Enrichment



Apache Ignite In-Memory Computing Platform

Applications



Financial
Services




Telco




Travel &
Logistics



E-Commerce



Pharma &
Healthcare



IoT

SQL

Key/Value

Transactions

Compute

Services

Streaming

ML


Memory-Centric Storage

Ignite Native Persistence
(Flash, SSD, Intel 3D XPoint)

Third-Party Persistence
(RDBMS, HDFS, NoSQL)


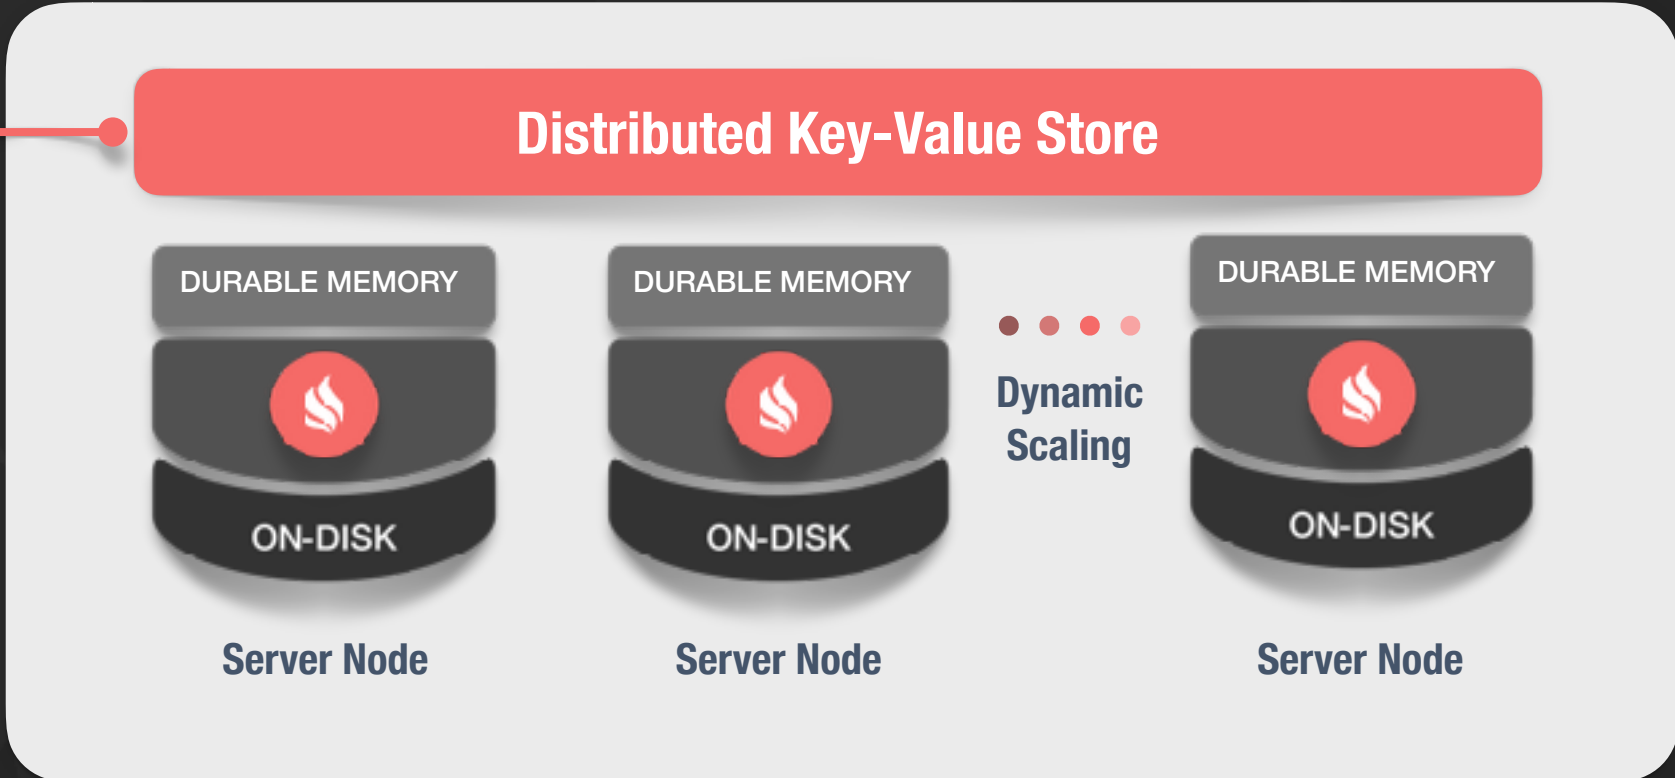
Data Grid

JCache & SQL



ACID Transaction

Distributed partitioned hash map



Distributed SQL

Cross-platform
Compatibility

Java

.NET

C++

BI

Tools

DDL, DML Support

JDBC

ODBC

SQL API

SELECT, UPDATE,
INSERT, MERGE,
DELETE, CREATE
and ALTER

Indexes in
RAM or Disk

Apache Ignite Cluster

DURABLE MEMORY

DURABLE MEMORY

DURABLE MEMORY

ON-DISK

ON-DISK

ON-DISK

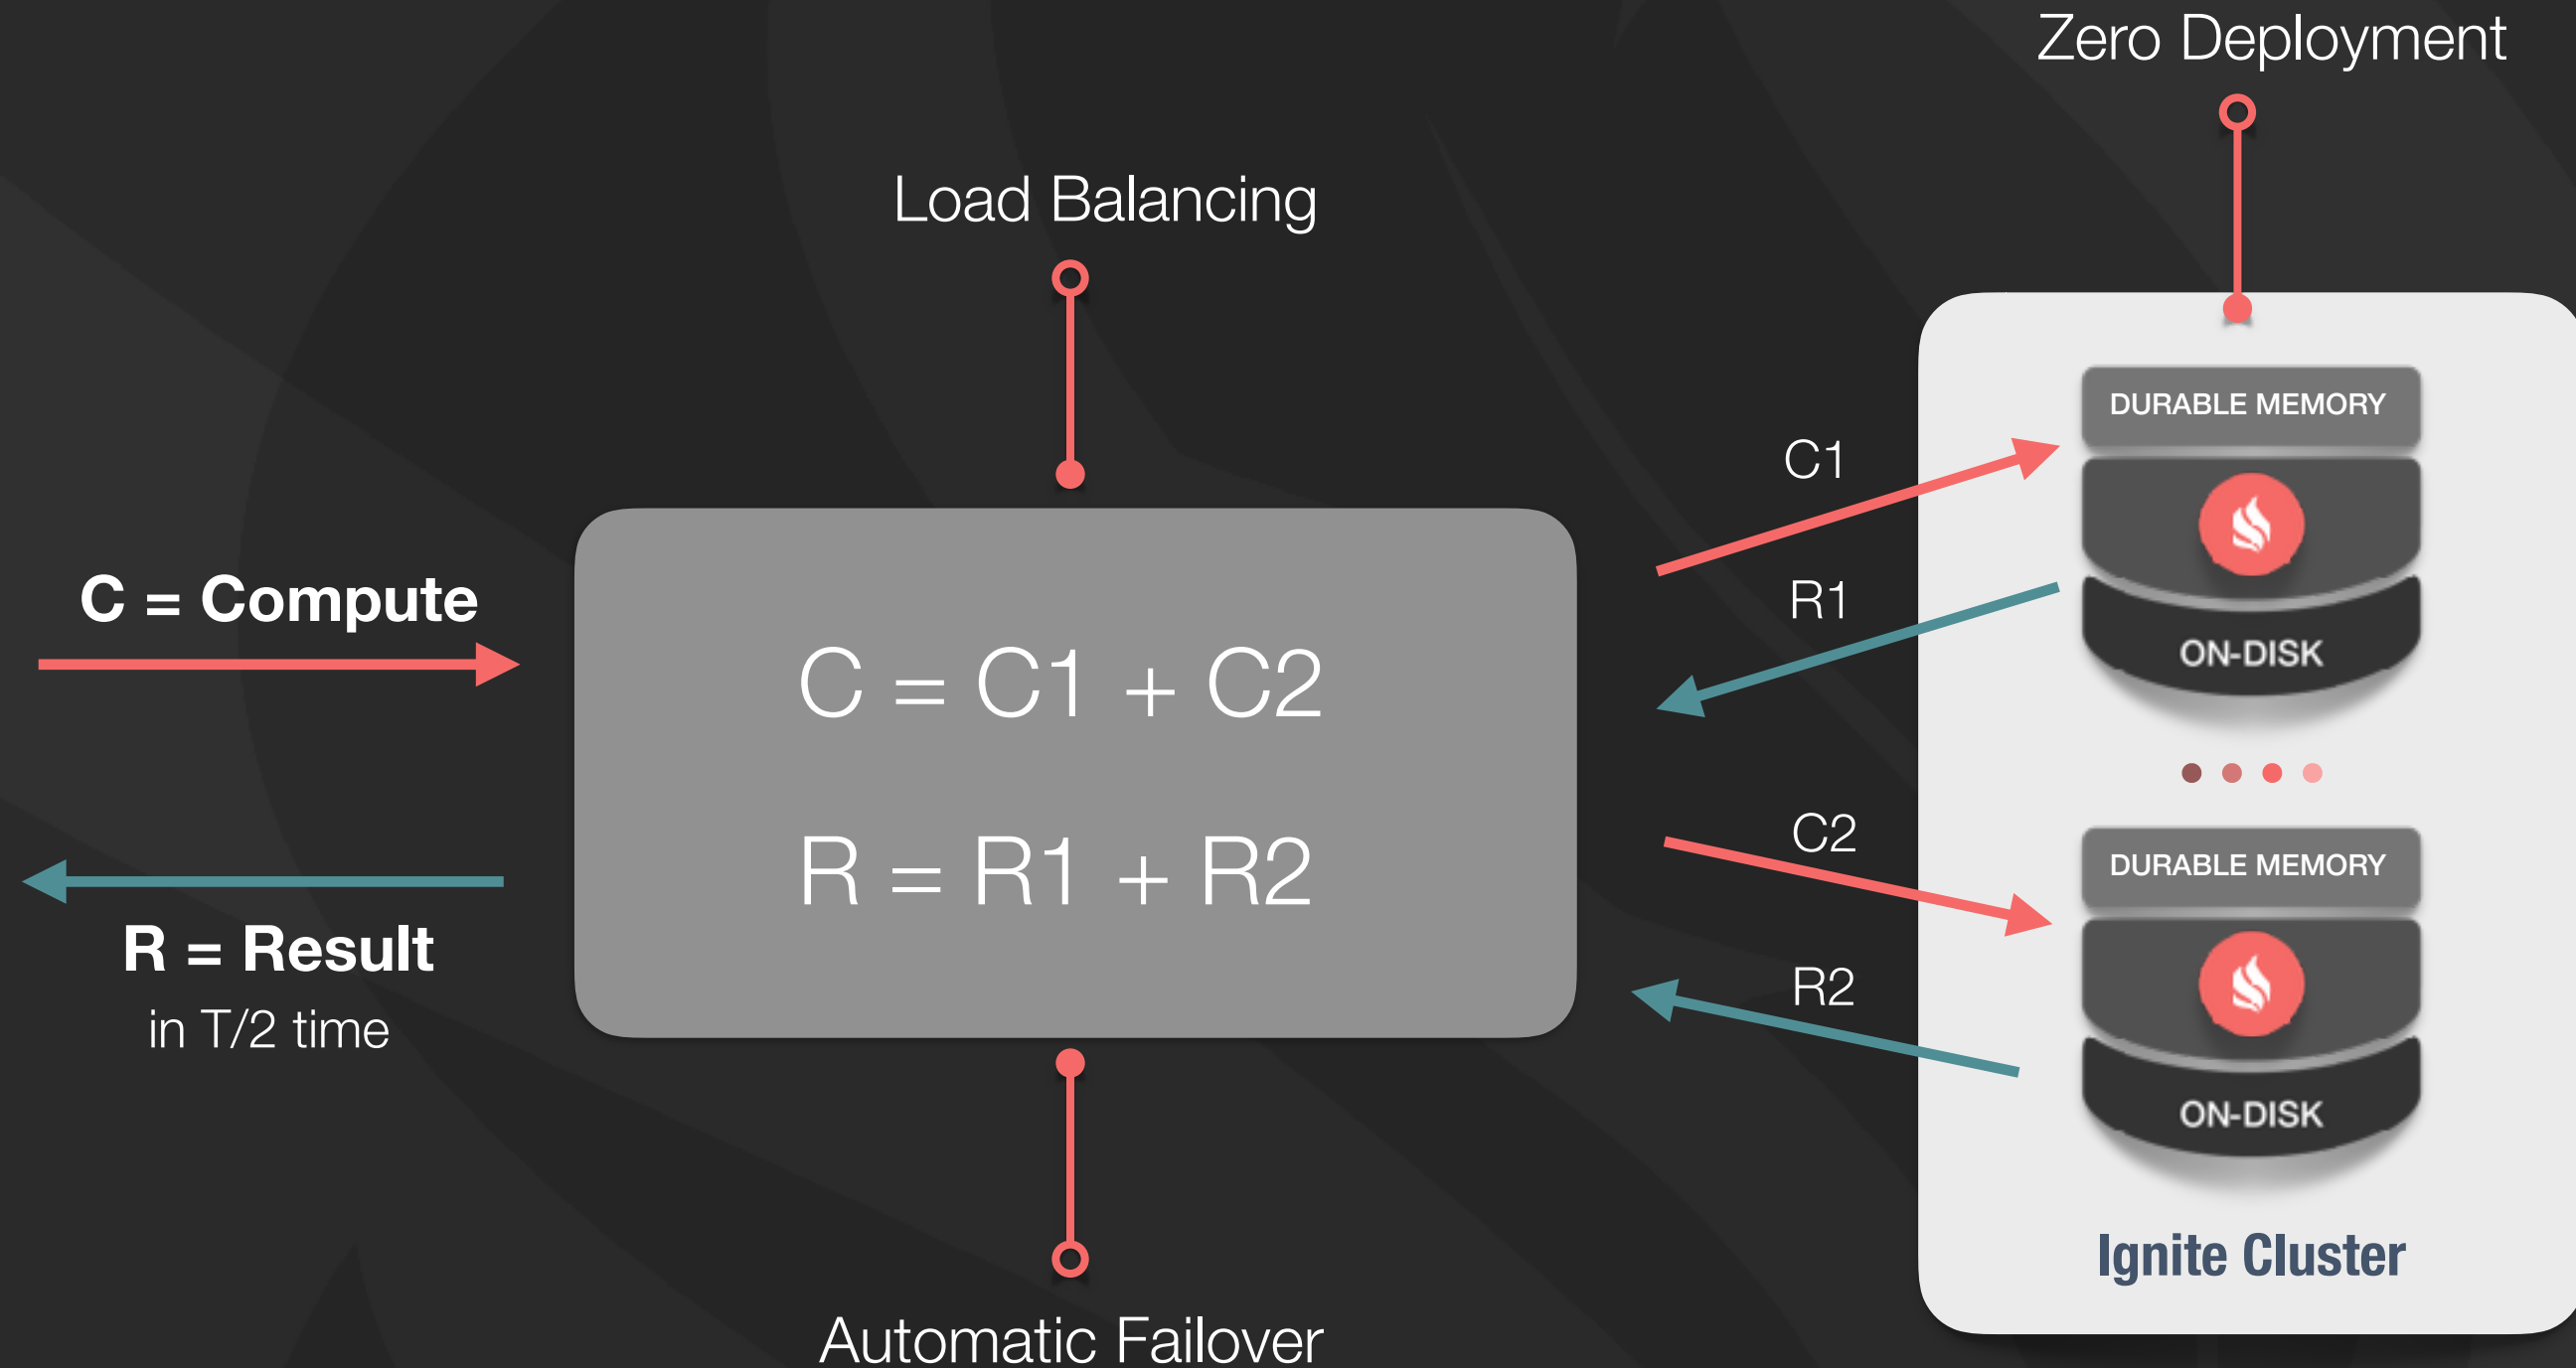
Server Node

Server Node

Server Node

Dynamic
Scaling

Compute Grid



Ignite and Spark Integration

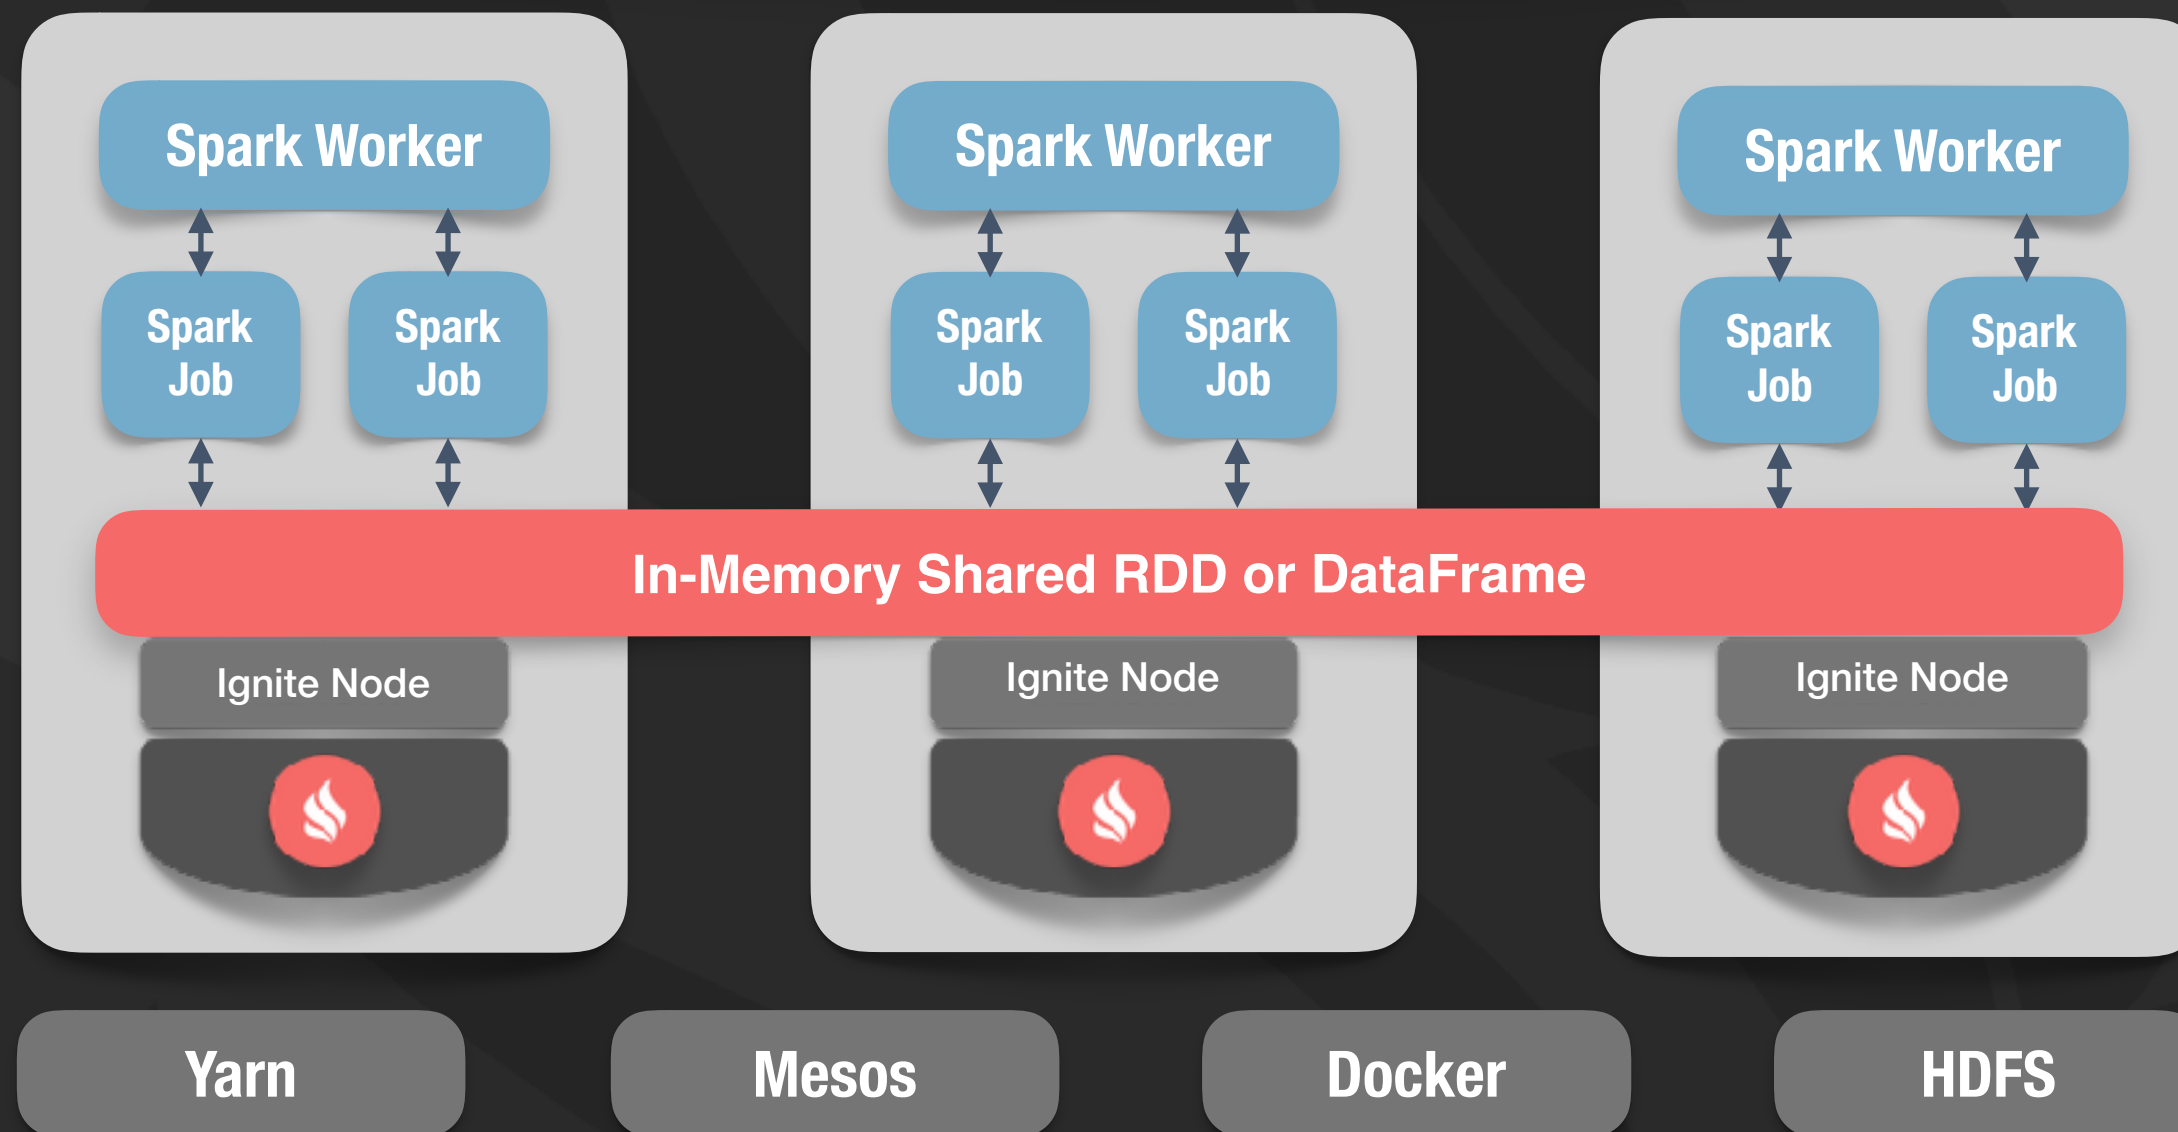
Spark Application

Share RDD
across jobs on
the host

Share RDD
Globally

In-Memory
Indexes

SQL on top of
RDDs




Machine Learning Grid

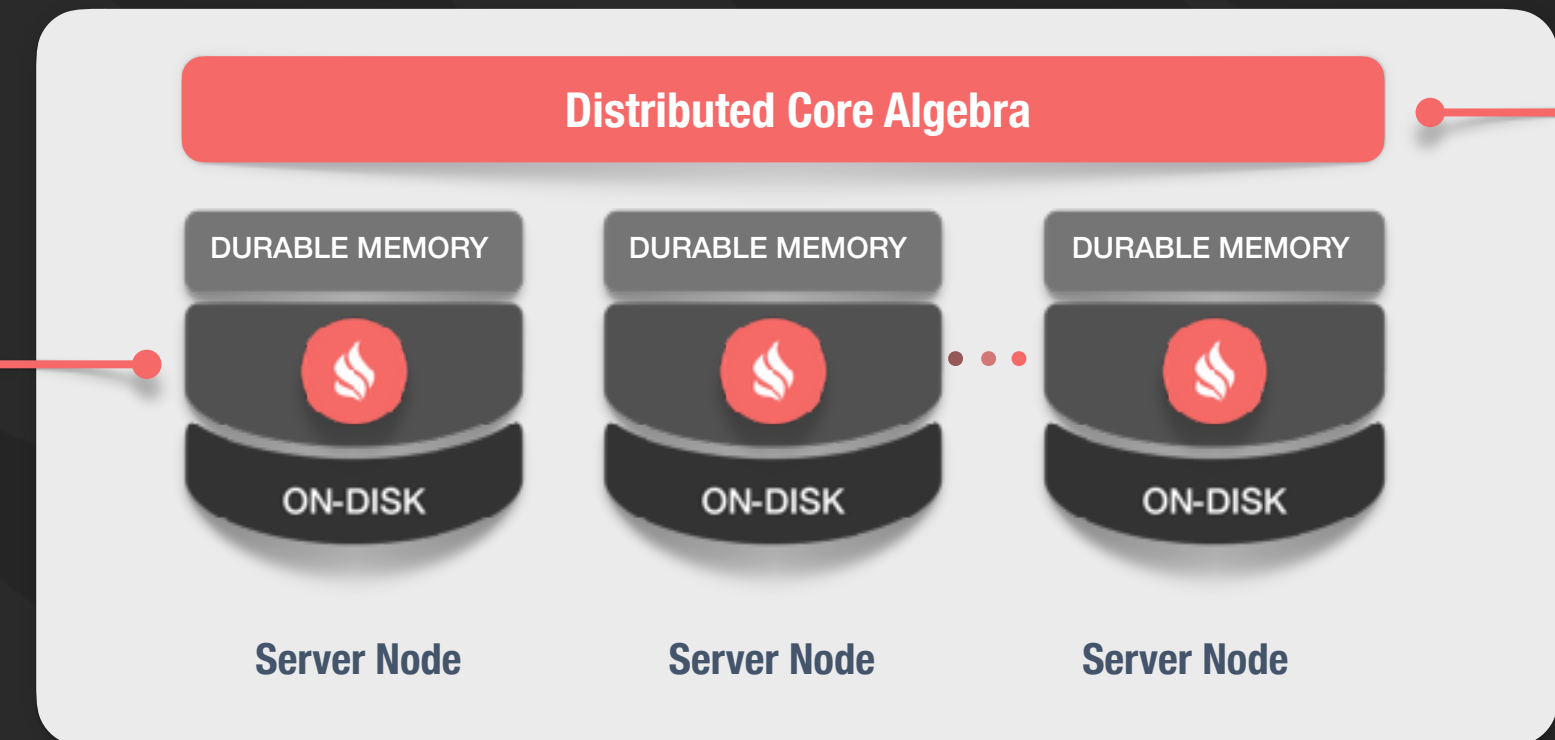


Multi-Language Support

Distributed Algorithms



Collocated Processing



Dense and Sparse Algebra

Apache Ignite Users

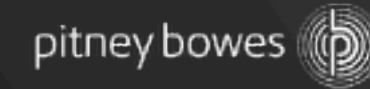
Financial Services



Software




Logistics & Travel



E-commerce




FinTech




Telco



Pharma & Healthcare



IoT



Adtech



Large IoT Provider - Smart Metering and Utilities

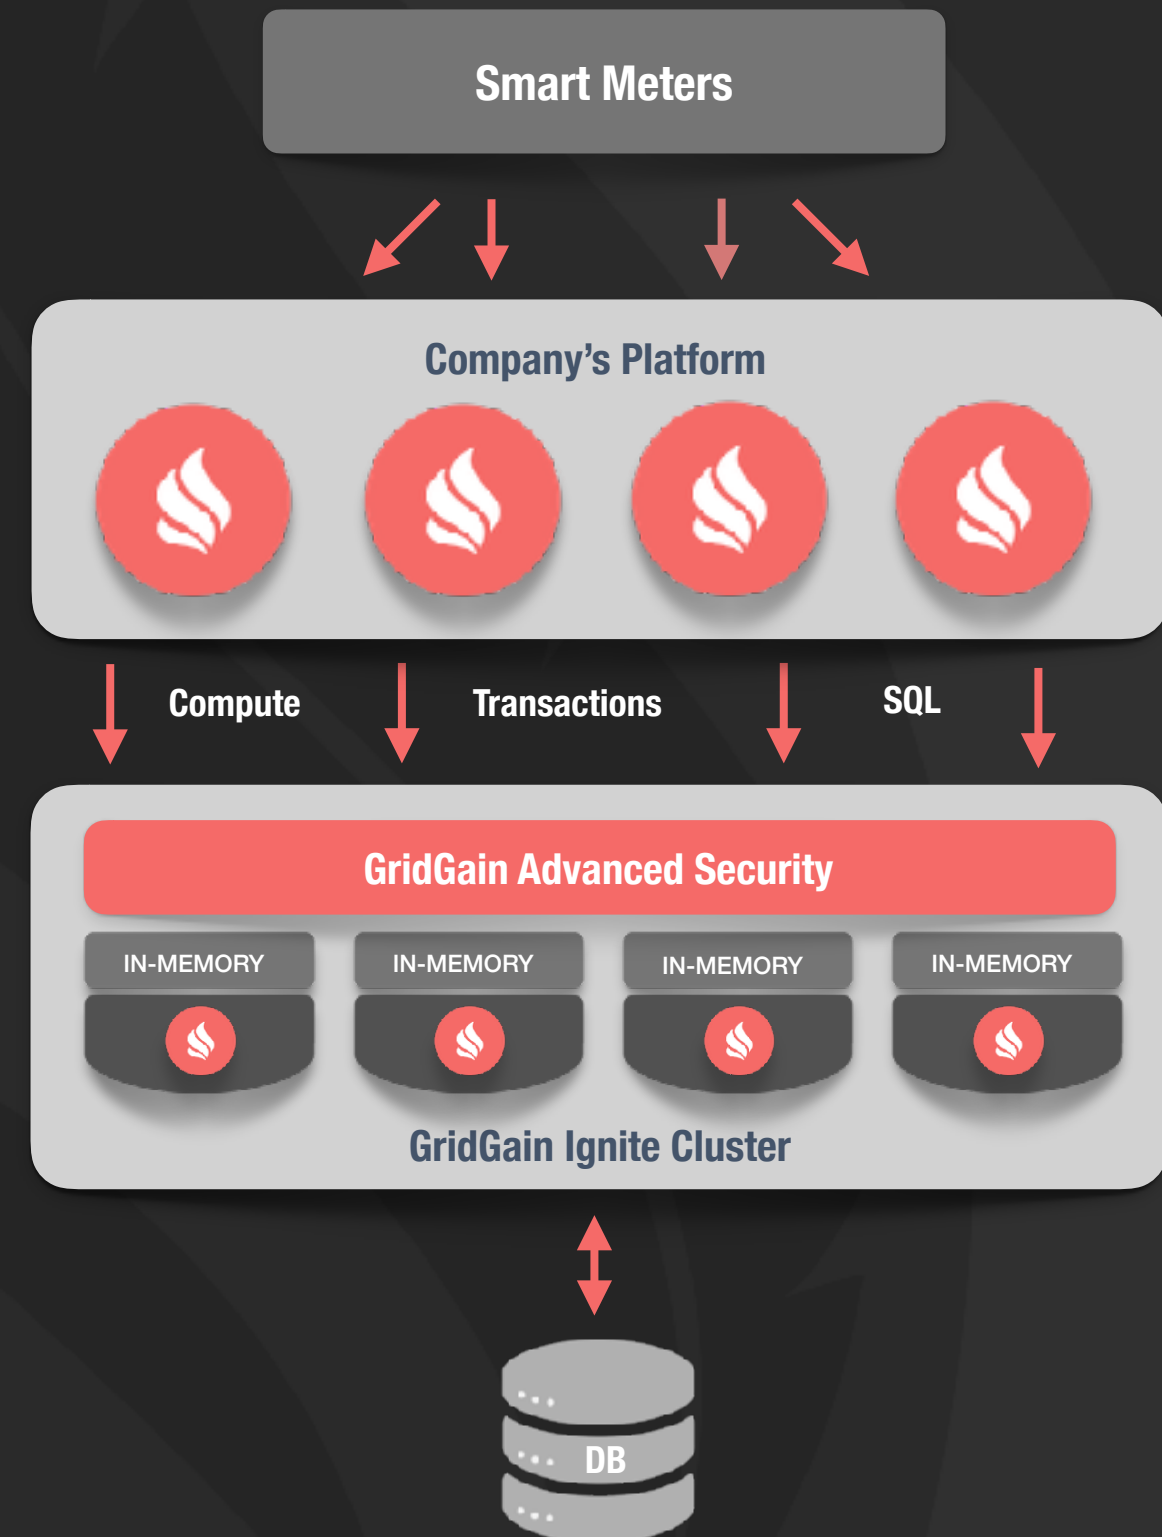
The company develops IoT solutions that transmit energy consumption data between meters, consumers and utilities in real time.

Problem

- Could not meet latency and throughput SLAs
- Missing scalability and elasticity

GridGain Solution

- 50 millions meters stream the data back in real-time
- Collocated in-memory processing
- Advanced security and multi-tenancy



Demo



Any Questions?

Thank you for joining us. Follow the conversation.

<http://ignite.apache.org>



#apacheignite
#denismagda