



In-Memory Performance
Durability of Disk





Apache Ignite and Apache Spark

Where Fast Data Meets the IoT



Akmal Chaudhri
GridGain Systems

Agenda

- IoT Demands to Software
- IoT Software Stack
 - Device OS/RTOS
 - Data Collection and Enrichment
 - NewSQL Database
 - 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

NewSQL Database

Data Collection and Enrichment

Device OS/Real-Time OS

Apache IoT Software Stack

Application APIs



NewSQL Database



Data Collection and Enrichment




Device OS/Real-Time OS



Apache MyNewt



Data Collection and Enrichment



Apache Ignite Database, Caching and Processing Platform



Financial
Services



Telco



Travel &
Logistics



E-Commerce



Pharma &
Healthcare



IoT

SQL

Key/Value

Transactions

Compute

Services

Streaming

ML

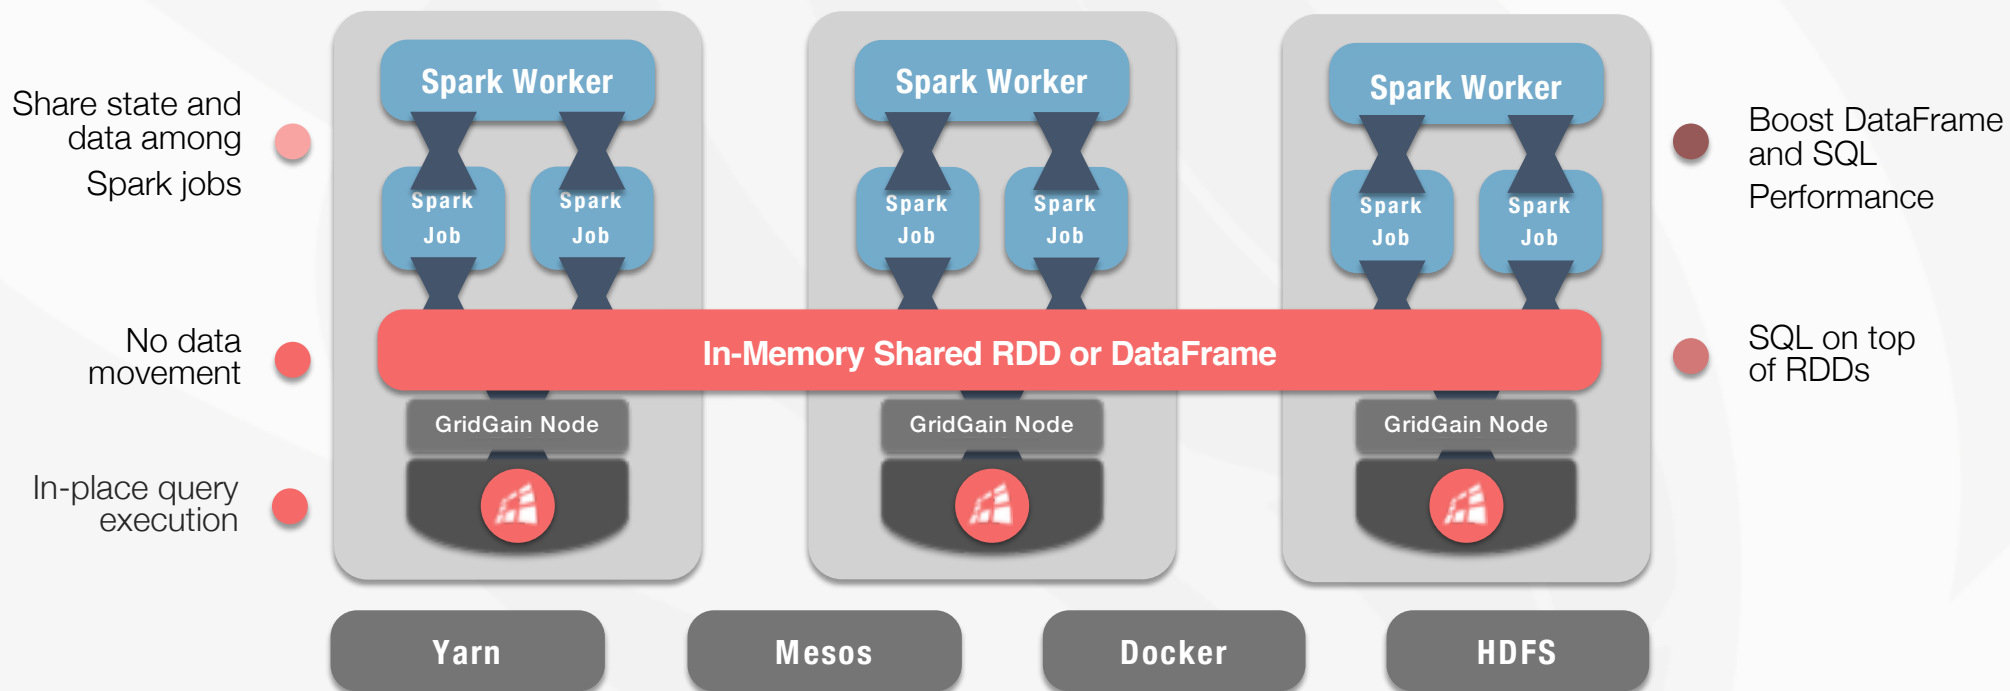
Memory-Centric Storage

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

Third-Party Persistence
(RDBMS, HDFS, NoSQL)

Ignite and Spark Integration

Spark Application



SQL Queries Execution Flow



```
SELECT ct.name, count(c.name)
FROM Country as ct
JOIN City as c ON c.countryCode = ct.code
WHERE ct.name IN ('Canada', 'India') GROUP BY (ct.name);
```



Ignite Node

Canada

- Toronto
- Montreal
- Ottawa
- Calgary



Ignite Node

India

- Mumbai
- New Delhi



3

1. Initial Query
2. Query execution over local data
3. Reduce multiple results in one

Comparing Ignite and Spark



- Distributed memory-centric **database**
- Fully fledged compute platform: SQL, transactions, key-value, collocated processing, ML/DL
- OLAP and OLTP



- Ingests data from HDFS or another storage
- Streaming and compute engine
- Inclined towards OLAP and focused on MR payloads

Ignite and Spark Together

Ignite is a memory-centric store for Spark

- No data movement from Ignite to Spark
 - In-place query execution
- Boost DataFrame and SQL performance
- Share state and data among Spark jobs
- Faster data and streaming analytics



DEMO



Any Questions?

Thank you for joining us. Follow the conversation.
<http://ignite.apache.org>



#apacheignite