

Getting Started with Apache® Ignite™

as a Distributed Database



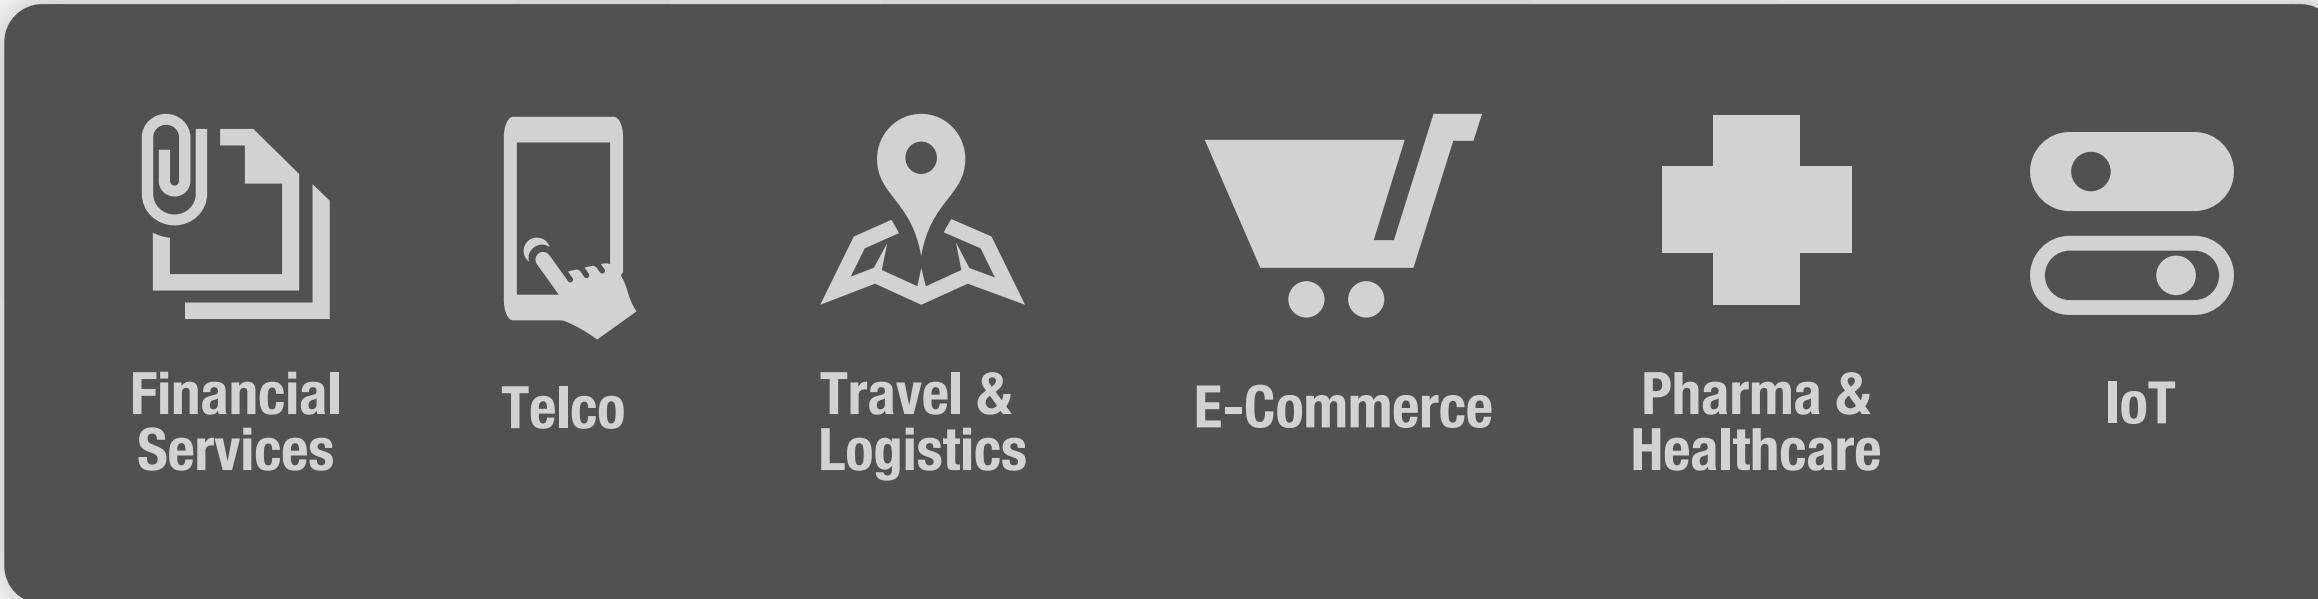
VALENTIN KULICHENKO

Lead Architect
GridGain Systems, Inc.

Agenda

- Apache Ignite as a Distributed Database
 - Connectivity
 - Data Definition Language
 - Data Manipulation Language
 - Collocated and Non-collocated Joins
- Under the Hood
 - Durable Memory Architecture
 - Native Persistence
- Data Storage Configuration
 - Data Regions
 - Data Eviction and Page Replacement
 - Performance Tips
- Demo
- Q&A

What is Apache Ignite?



SQL

Key/Value

Transactions

Compute

Services

Streaming

ML

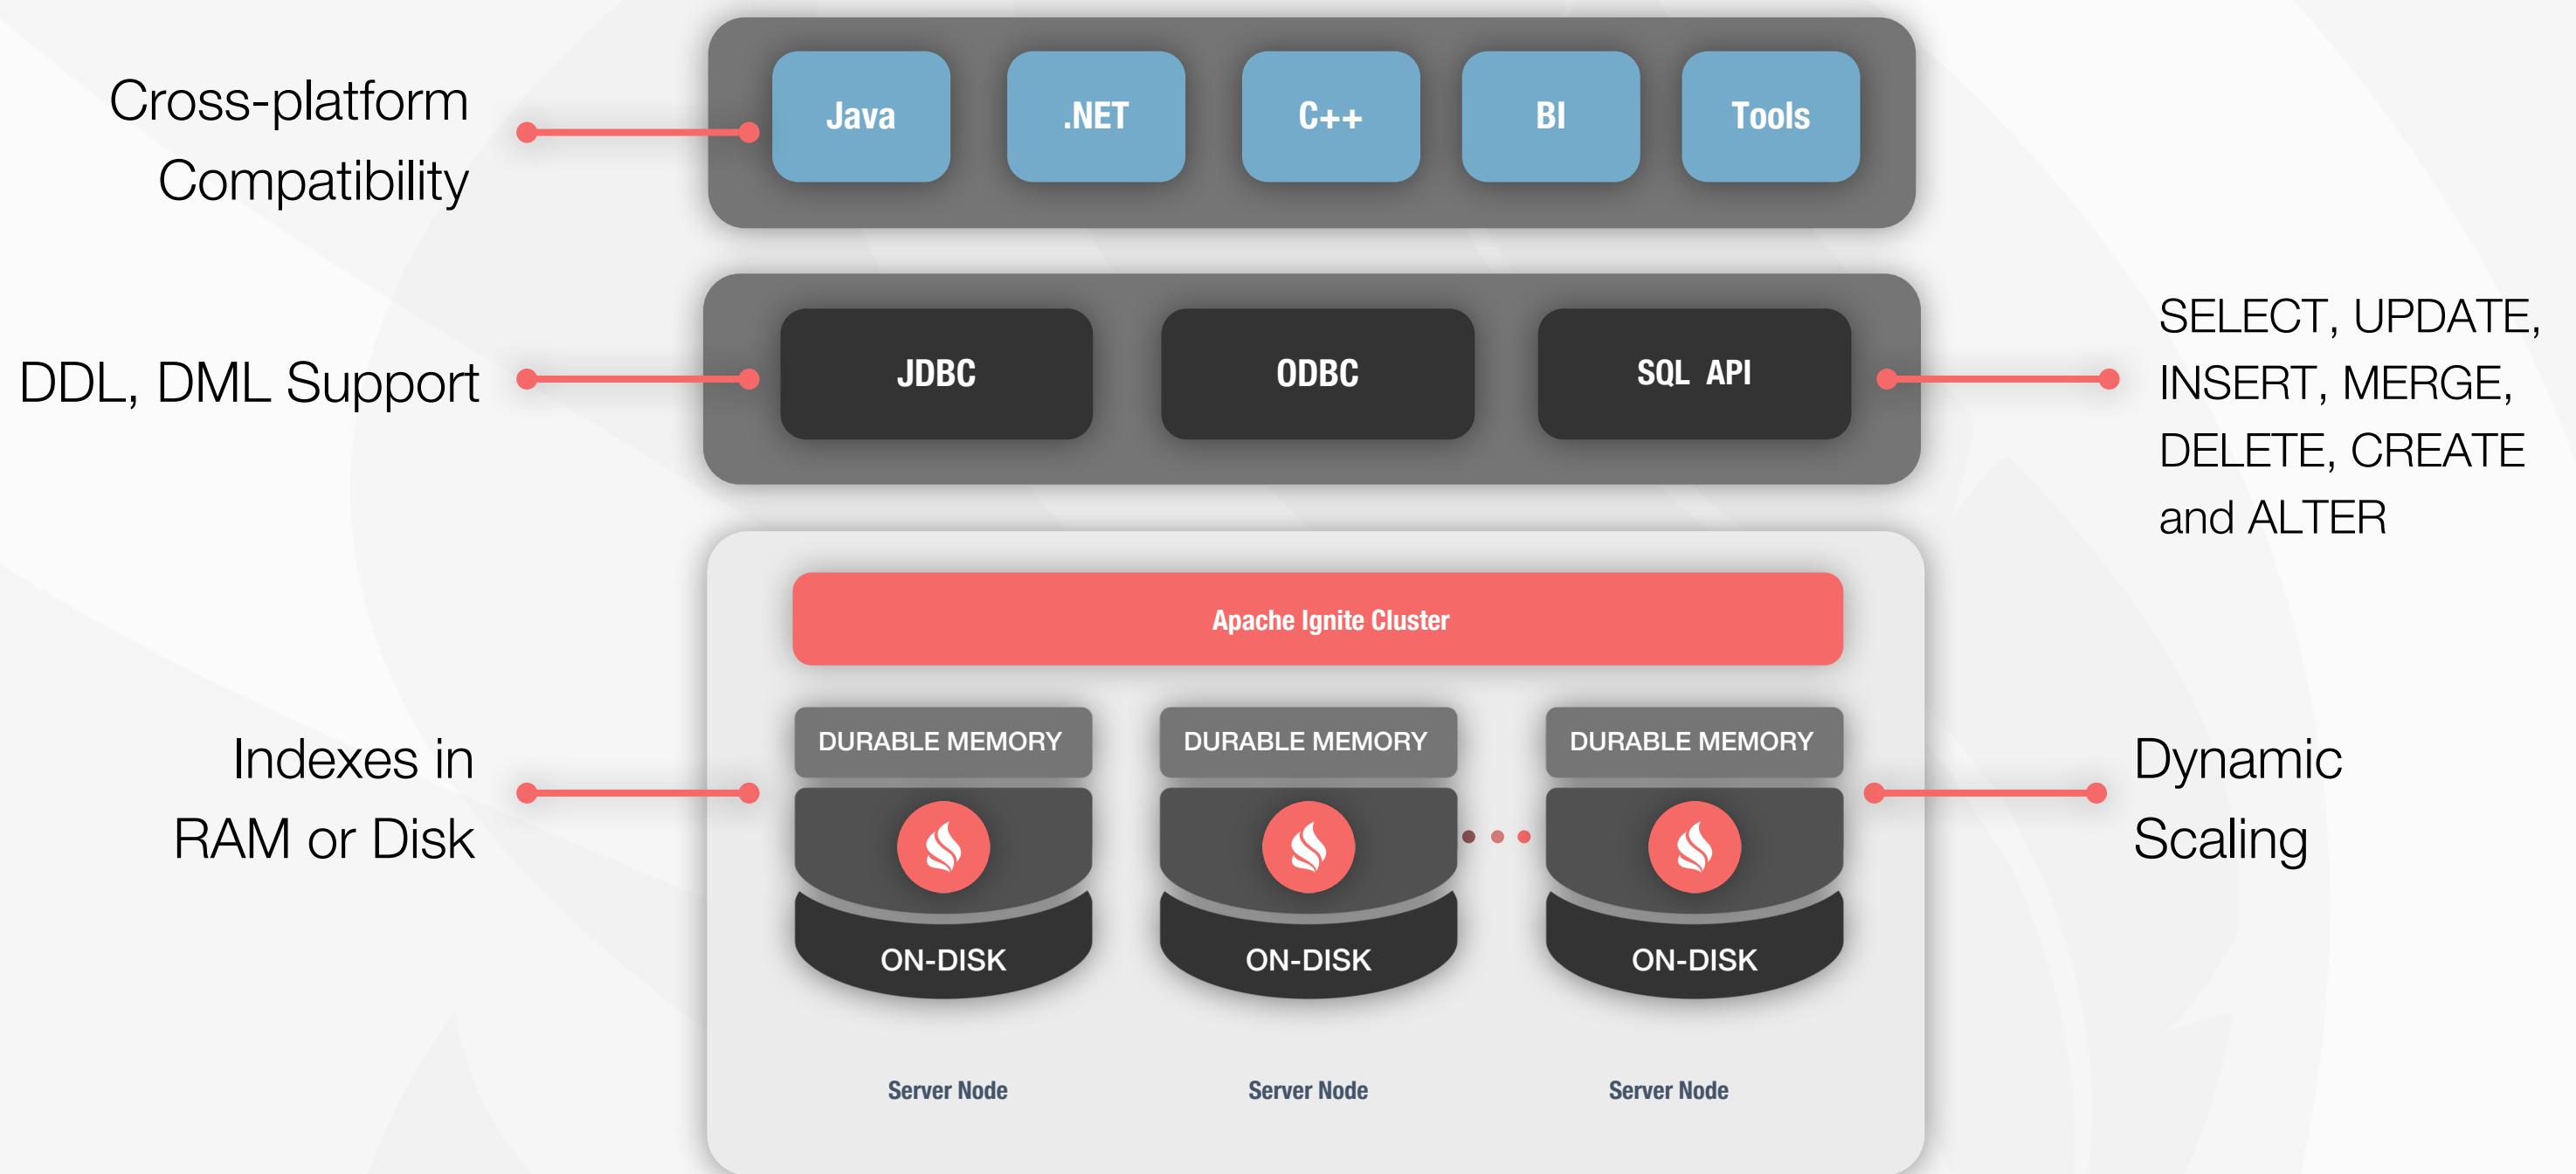
Memory-Centric Storage

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

Third-Party Persistence
(RDBMS, HDFS, NoSQL)

Ignite as a Distributed Database

Distributed SQL



Connectivity

- JDBC
- ODBC
- REST
- Java, .NET and C++ APIs

```
// Register JDBC driver.  
Class.forName("org.apache.ignite.IgniteJdbcThinDriver");  
  
// Open the JDBC connection.  
Connection conn = DriverManager.getConnection("jdbc:ignite:thin://192.168.0.50");  
  
. ./sqlline.sh --color=true --verbose=true -u jdbc:ignite:thin://127.0.0.1/
```

Data Definition Language

- CREATE/DROP TABLE
- CREATE/DROP INDEX
- ALTER TABLE
- Changes Durability
 - Ignite Native Persistence

```
CREATE TABLE `city` (
    `ID` INT(11),
    `Name` CHAR(35),
    `CountryCode` CHAR(3),
    `District` CHAR(20),
    `Population` INT(11),
    PRIMARY KEY (`ID`, `CountryCode`)
) WITH "template=partitioned, backups=1, affinityKey=CountryCode";
```

<https://apacheignite-sql.readme.io/docs/ddl>

Data Manipulation Language

- ANSI-99 specification
- Fault-tolerant and consistent
- INSERT, UPDATE, DELETE
- SELECT
 - JOINs
 - Subqueries

```
SELECT country.name, city.name, MAX(city.population) as max_pop
FROM country JOIN city ON city.countrycode = country.code
WHERE country.code IN ('USA', 'RUS', 'CHN')
GROUP BY country.name, city.name ORDER BY max_pop DESC LIMIT 3;
```

<https://apacheignite-sql.readme.io/docs/dml>

Collocated Joins



1

```
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);
```

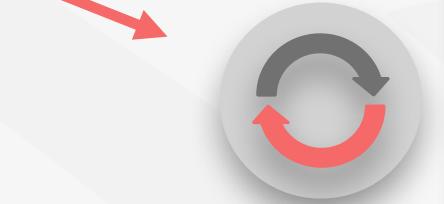
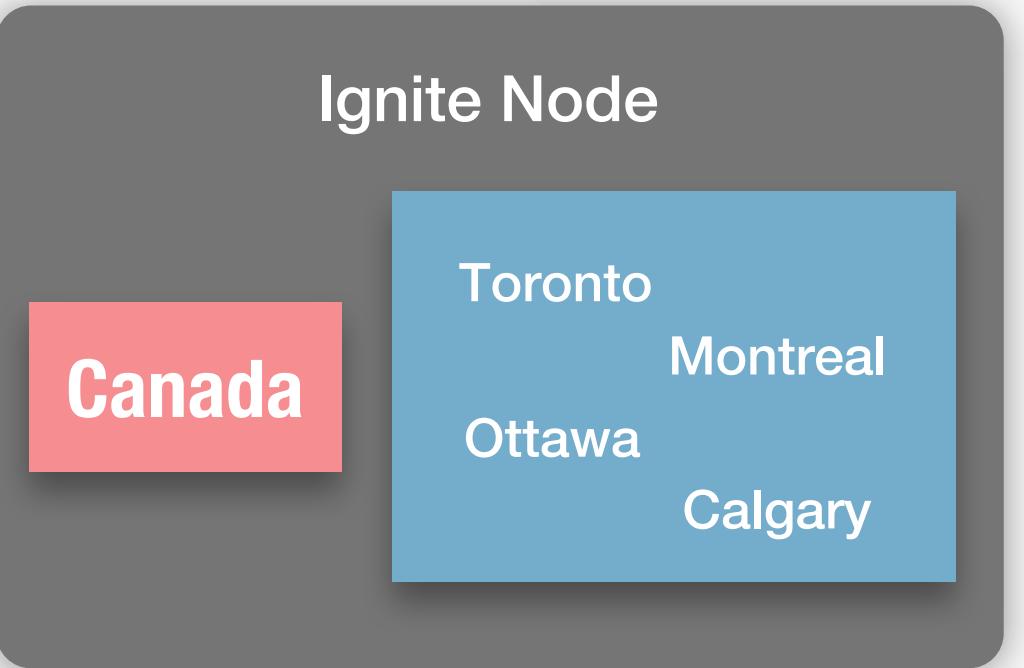


3

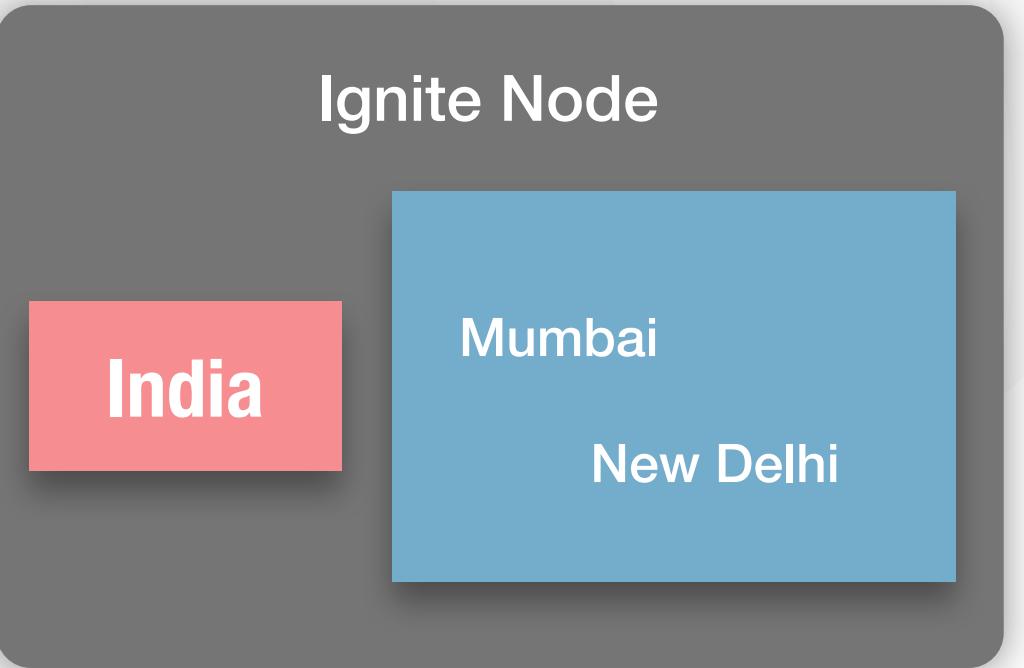
1. Initial Query
2. Query execution over local data
3. Reduce multiple results in one



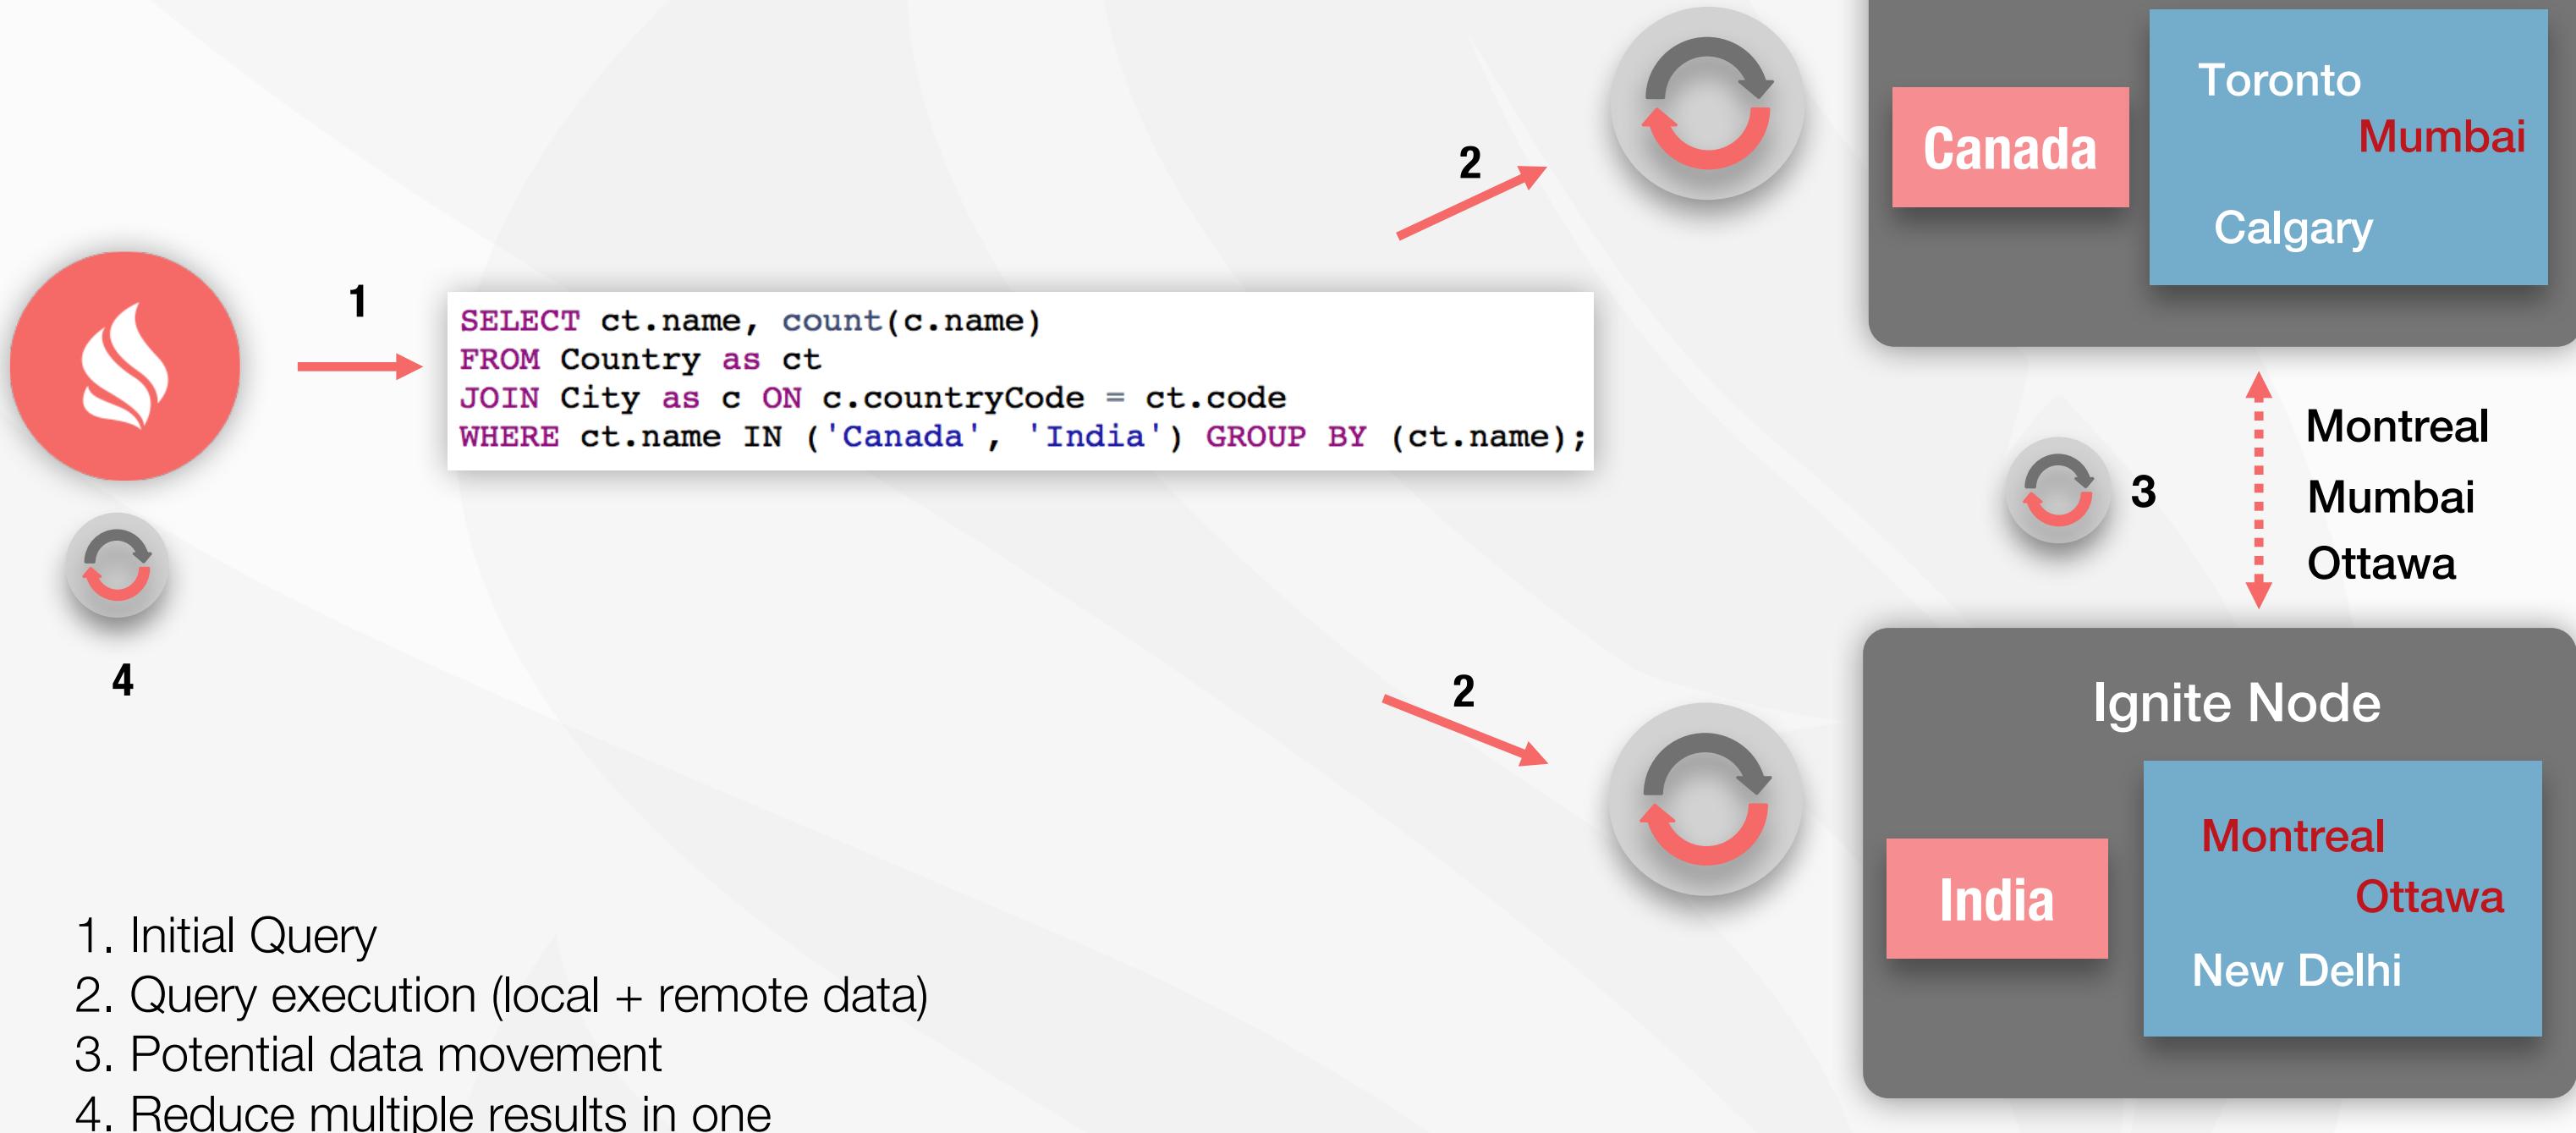
2



2

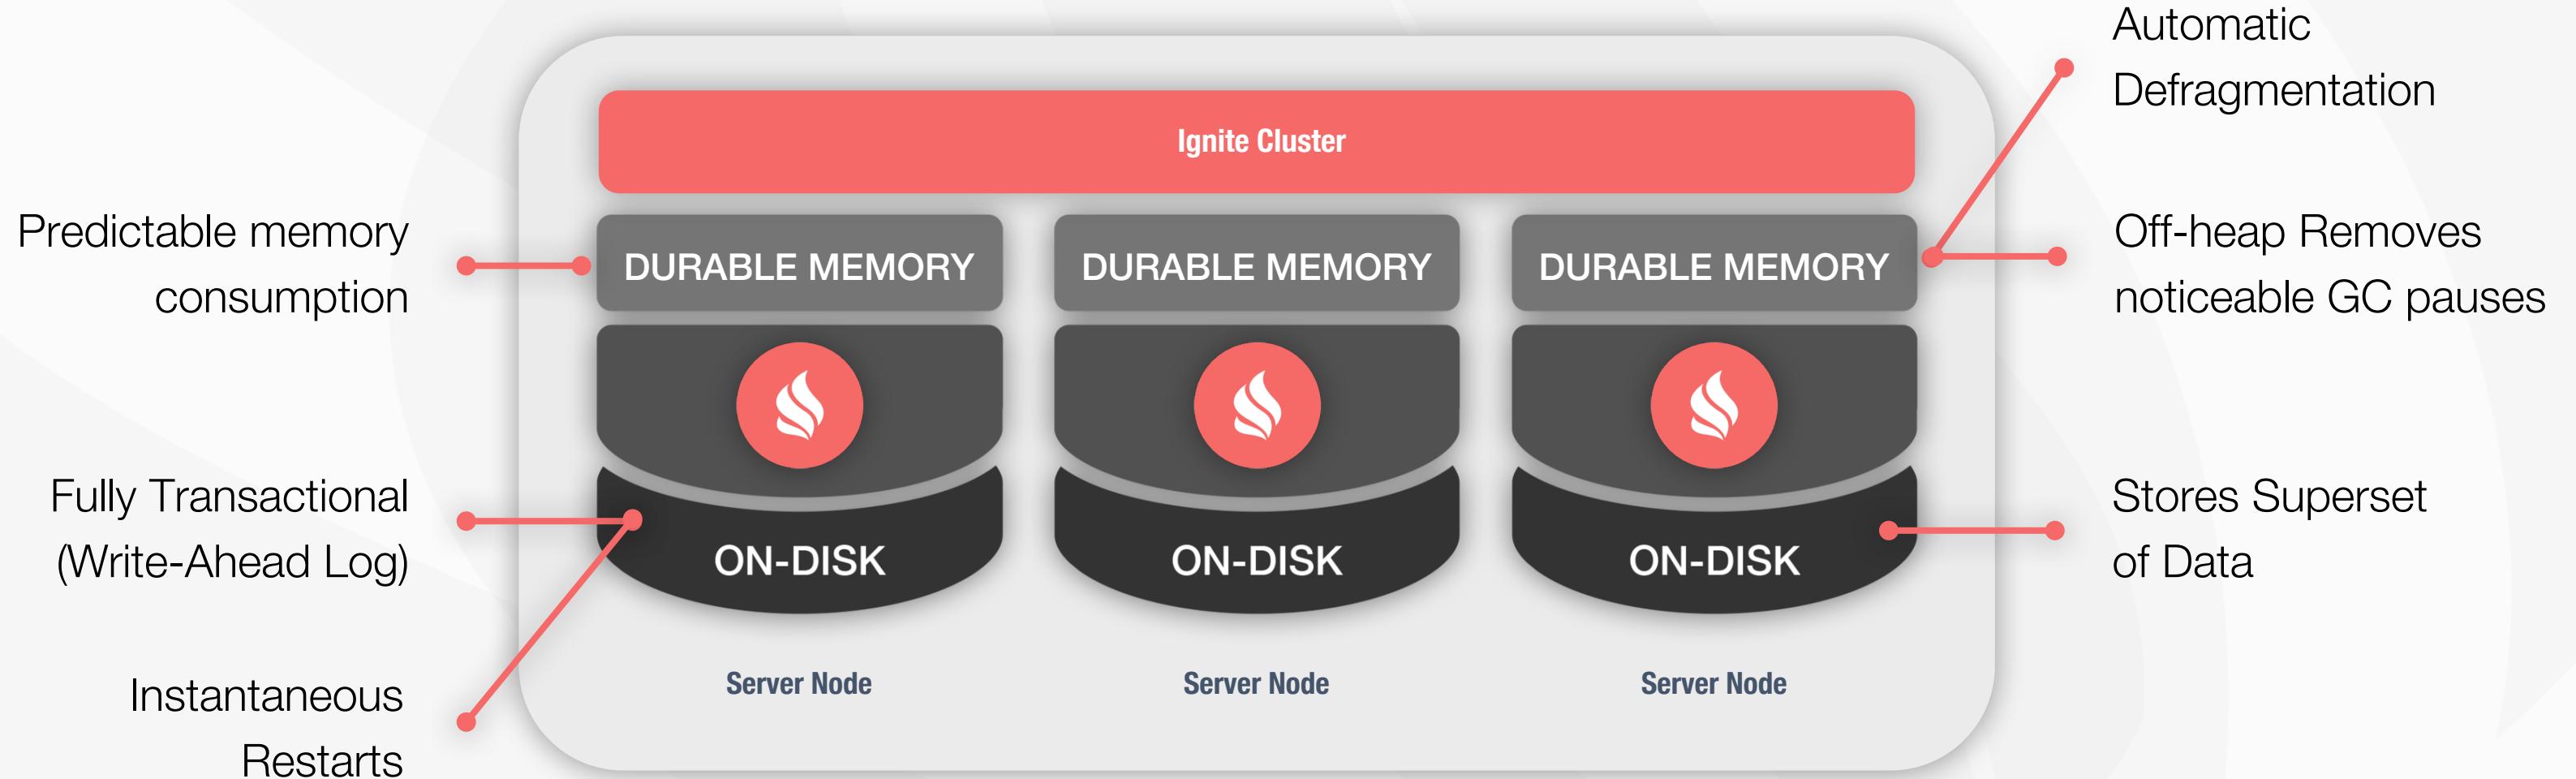


Non-Collocated Joins

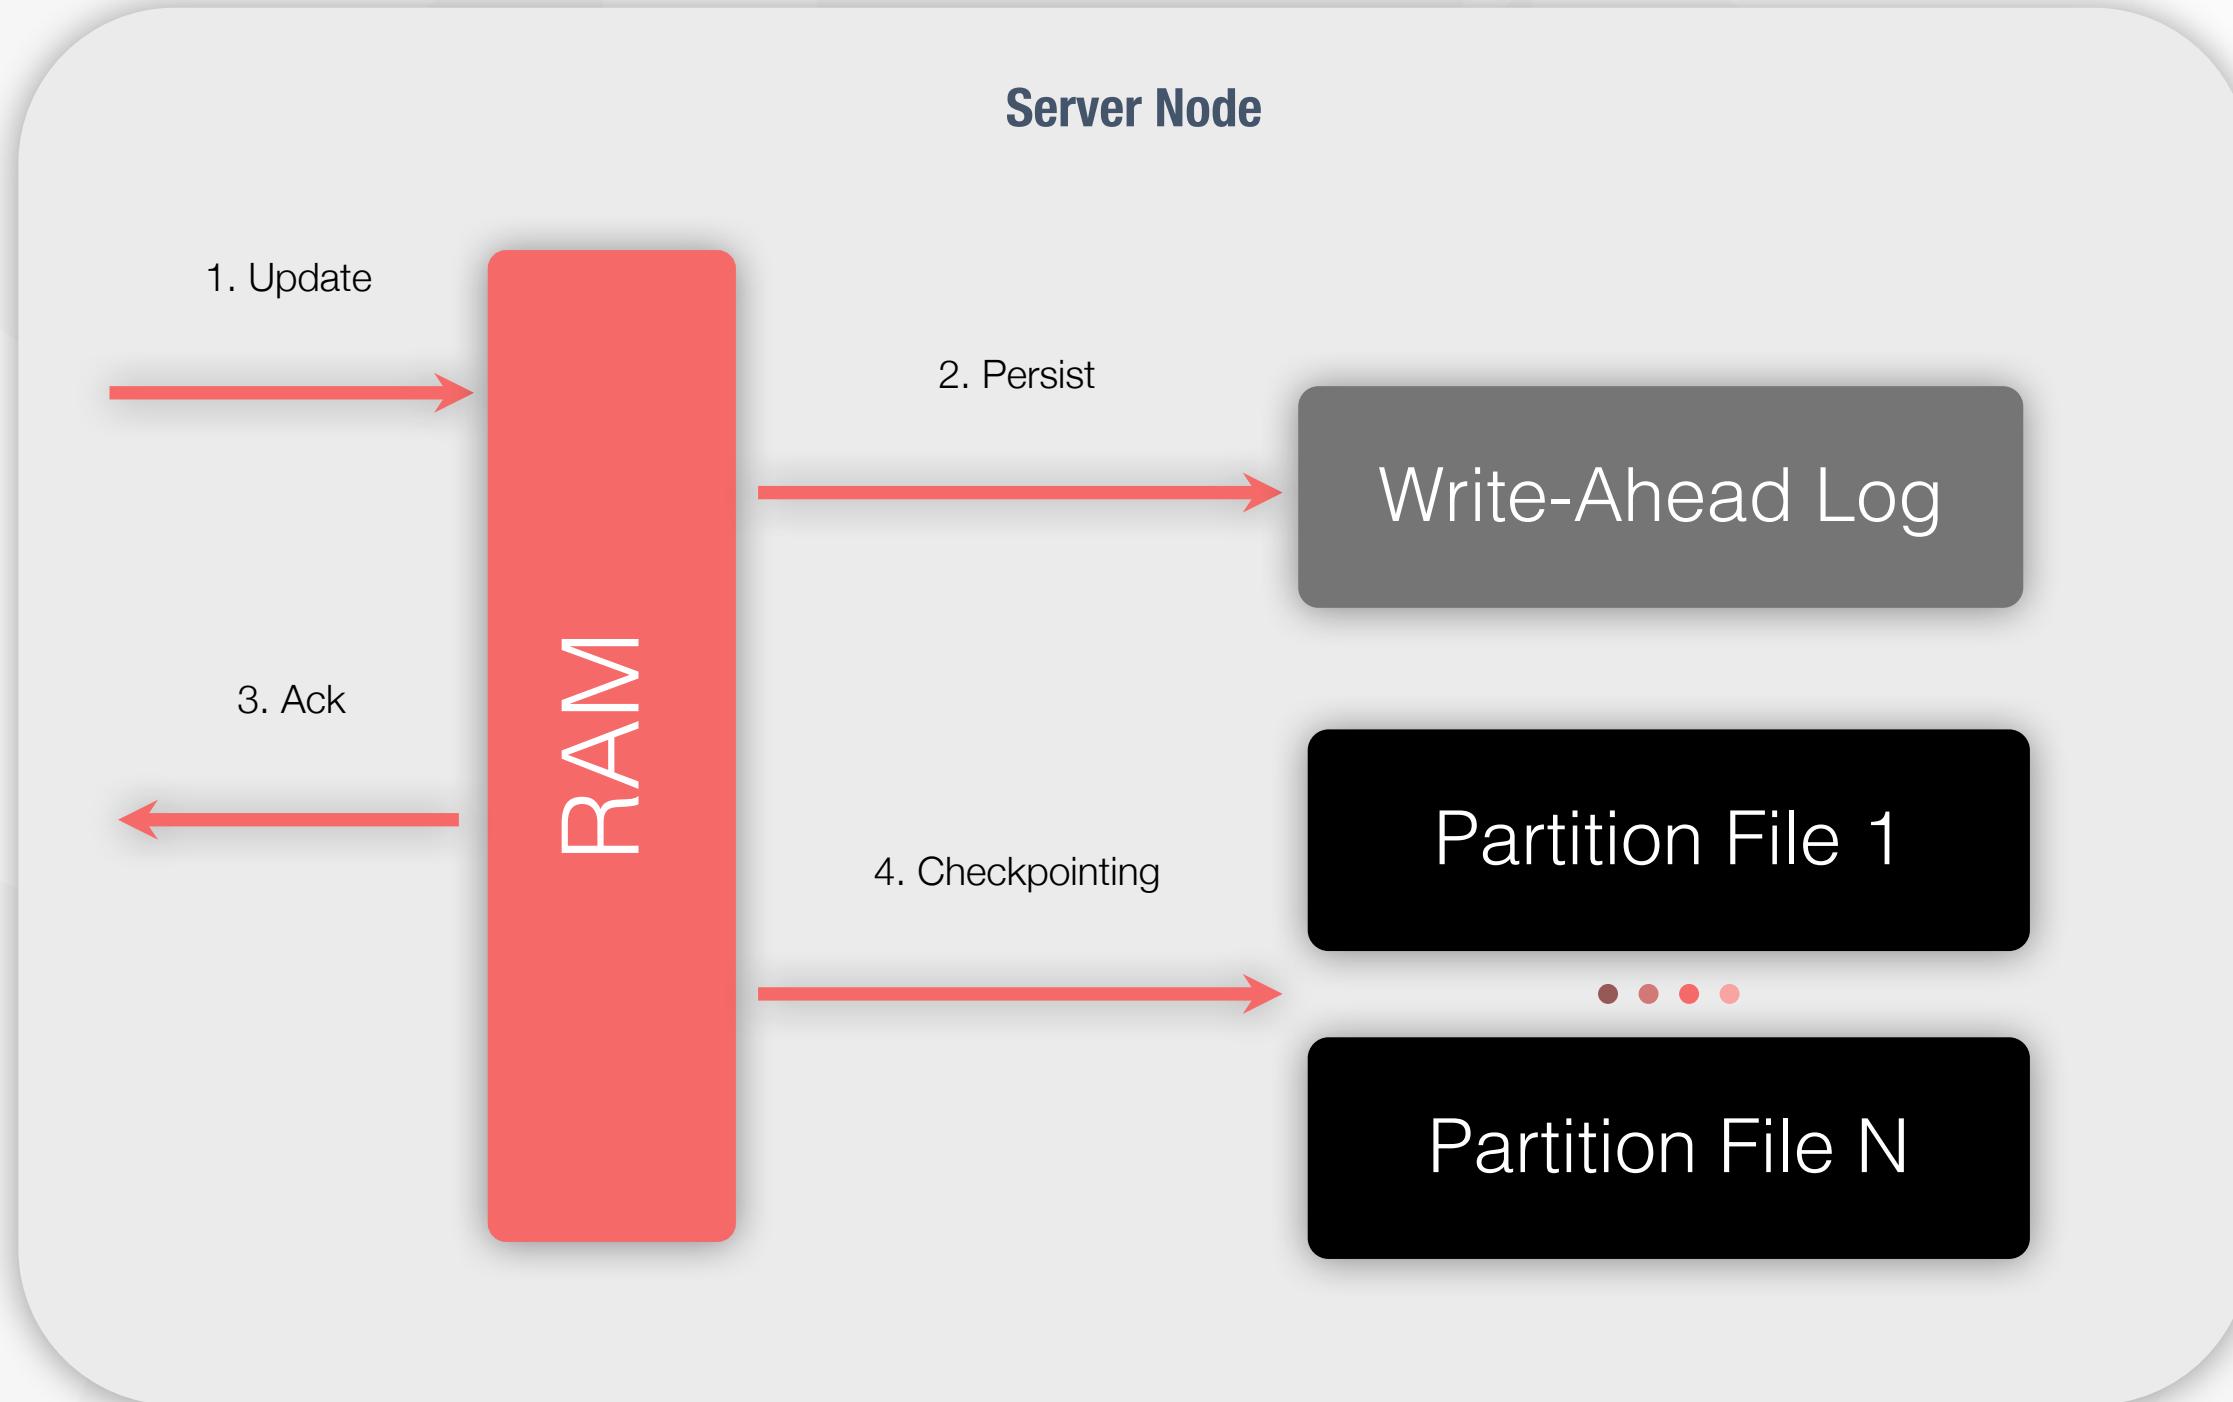


Under the Hood

Durable Memory



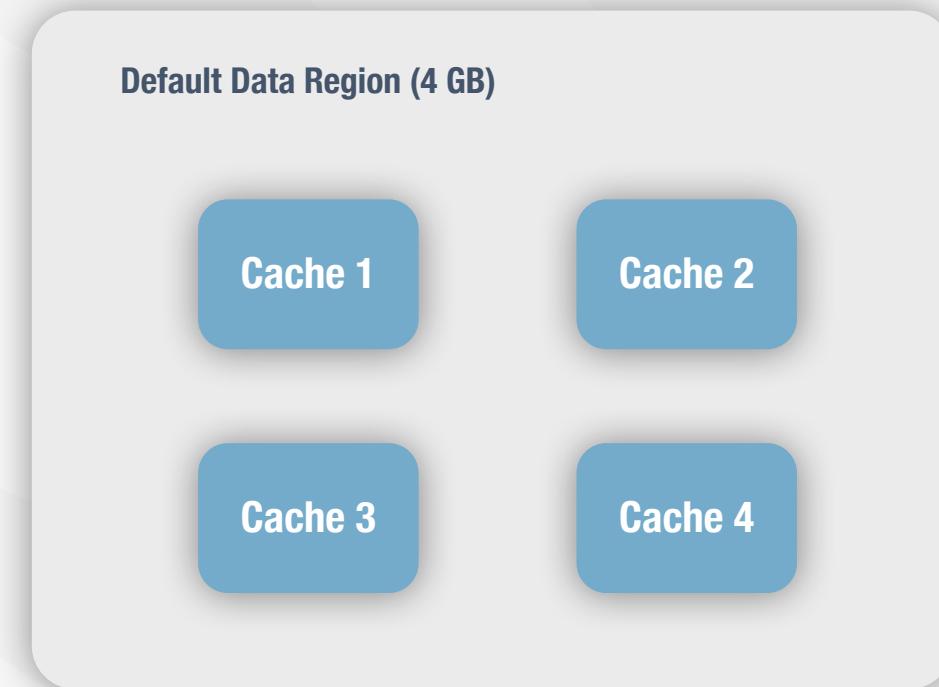
Ignite Native Persistence



Data Storage Configuration

Single Data Region

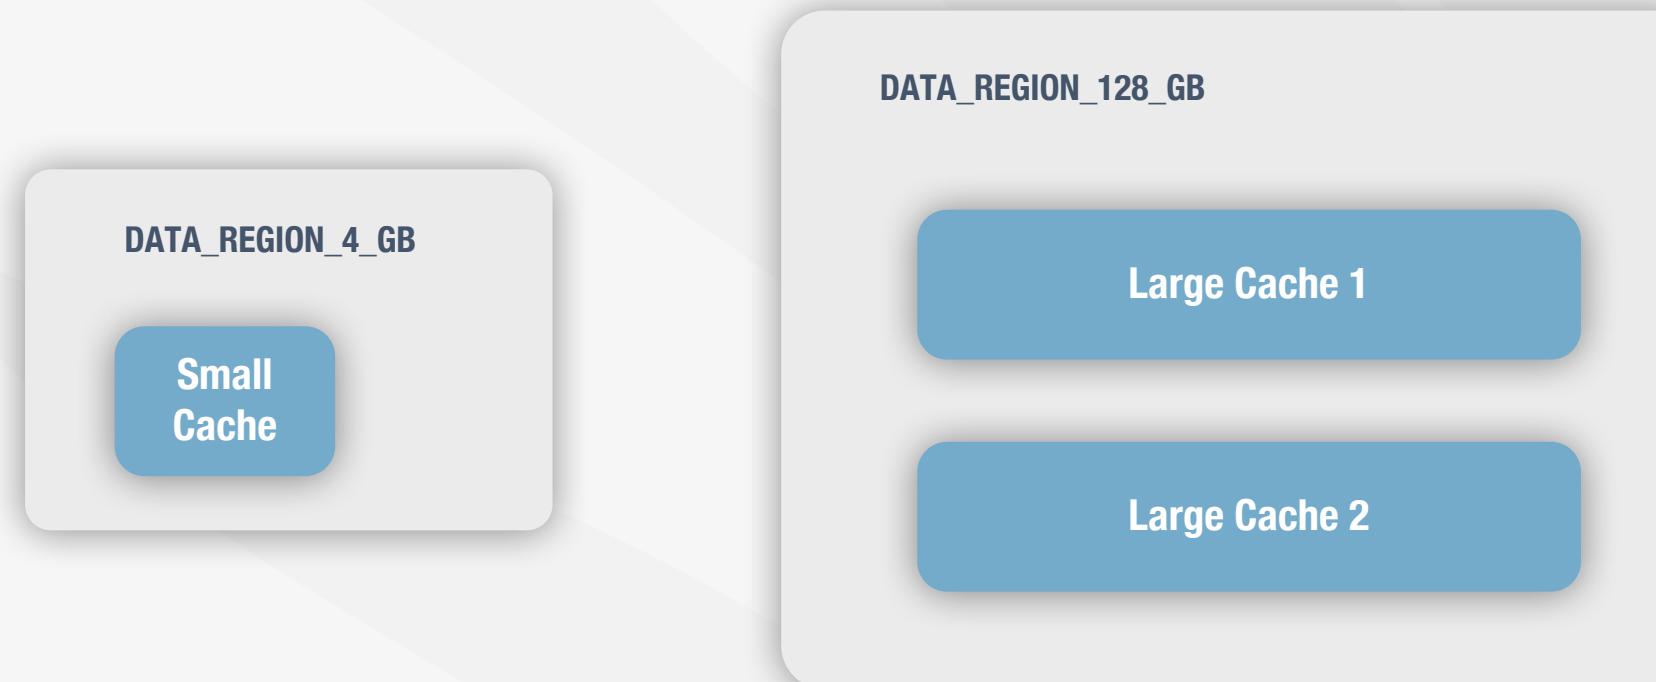
```
<property name="dataStorageConfiguration">
    <bean class="org.apache.ignite.configuration.DataStorageConfiguration">
        <property name="defaultDataRegionConfiguration">
            <bean class="org.apache.ignite.configuration.DataRegionConfiguration">
                <property name="maxSize" value="#{4L * 1024 * 1024 * 1024}"/> <!-- 4 GB. -->
            </bean>
        </property>
    </bean>
</property>
```



Multiple Data Regions

```
<property name="dataStorageConfiguration">
    <bean class="org.apache.ignite.configuration.DataStorageConfiguration">
        <property name="dataRegionConfigurations">
            <list>
                <bean class="org.apache.ignite.configuration.DataRegionConfiguration">
                    <property name="name" value="DATA_REGION_4_GB"/>
                    <property name="maxSize" value="#{4L * 1024 * 1024 * 1024}" /> <!-- 4 GB. -->
                </bean>

                <bean class="org.apache.ignite.configuration.DataRegionConfiguration">
                    <property name="name" value="DATA_REGION_128_GB"/>
                    <property name="maxSize" value="#{128L * 1024 * 1024 * 1024}" /> <!-- 128 GB. -->
                </bean>
            </list>
        </property>
    </bean>
</property>
```



Enabling Data Eviction

```
<property name="dataStorageConfiguration">
    <bean class="org.apache.ignite.configuration.DataStorageConfiguration">
        <property name="defaultDataRegionConfiguration">
            <bean class="org.apache.ignite.configuration.DataRegionConfiguration">
                <property name="maxSize" value="#{4L * 1024 * 1024 * 1024}"/> <!-- 4 GB. -->

                <property name="pageEvictionMode" value="RANDOM_LRU"/>
                <property name="evictionThreshold" value="0.9"/>
            </bean>
        </property>
    </bean>
</property>
```

- Eviction Modes:
 - DISABLED
 - RANDOM_LRU
 - RANDOM_2_LRU

Enabling Native Persistence

```
<property name="dataStorageConfiguration">
    <bean class="org.apache.ignite.configuration.DataStorageConfiguration">
        <property name="defaultDataRegionConfiguration">
            <bean class="org.apache.ignite.configuration.DataRegionConfiguration">
                <property name="maxSize" value="#{4L * 1024 * 1024 * 1024}"/> <!-- 4 GB. -->

                <property name="persistenceEnabled" value="true"/>
            </bean>
        </property>

        <property name="walMode" value="LOG_ONLY"/>
        <property name="checkpointFrequency" value="60000"/> <!-- 60 sec. -->
    </bean>
</property>
```

- WAL Modes:
 - DEFAULT (full sync)
 - LOG_ONLY
 - BACKGROUND
 - NONE

Performance Tips

- Leave memory space for heap and OS
- Use production ready SSDs
- Have separate physical disk for WAL
- Advanced data storage configuration:
 - Change page size
 - Enable write throttling
 - Increase checkpoint buffer size

<https://apacheignite.readme.io/docs/durable-memory-tuning>

Demo

Resources

- GridGain in-memory computing resources
 - Learning: <https://www.gridgain.com/resources/in-memory-computing-resources>
 - In-Memory Computing Summit: <https://www.imcsummit.org/>
 - Local Events: <https://www.gridgain.com/company/news/events>
 - Customer Stories: <https://www.gridgain.com/customers/customers-and-industries>
- Getting started with Apache Ignite:
 - <https://ignite.apache.org/>
 - Getting Started doc: <https://apacheignite.readme.io/docs/getting-started>
 - GridGain Web Console (free): <https://www.gridgain.com/resources/manage>
 - Apache Ignite User Forum: <http://apache-ignite-users.70518.x6.nabble.com/>
 - Stack Overflow: <https://stackoverflow.com/questions/tagged/ignite>
- Getting started with GridGain:
 - <https://www.gridgain.com/>
 - <https://www.gridgain.com/resources/download>

HAPPY VALENTIN'S DAY!



VALENTIN KULICHENKO

Lead Architect
GridGain Systems, Inc.