How to Scale Up Applications

If COVID-19 is Causing a Dramatic Spike for Your Business

Nikita Ivanov

Founder & CTO, GridGain Systems

PMC Member, Apache Ignite

COVID-19: Unique IT/Systems Challenges

- **Up to 10x** increase in demand or system load in a span of few weeks
- Latency **and** throughput are both critical for minimal SLAs

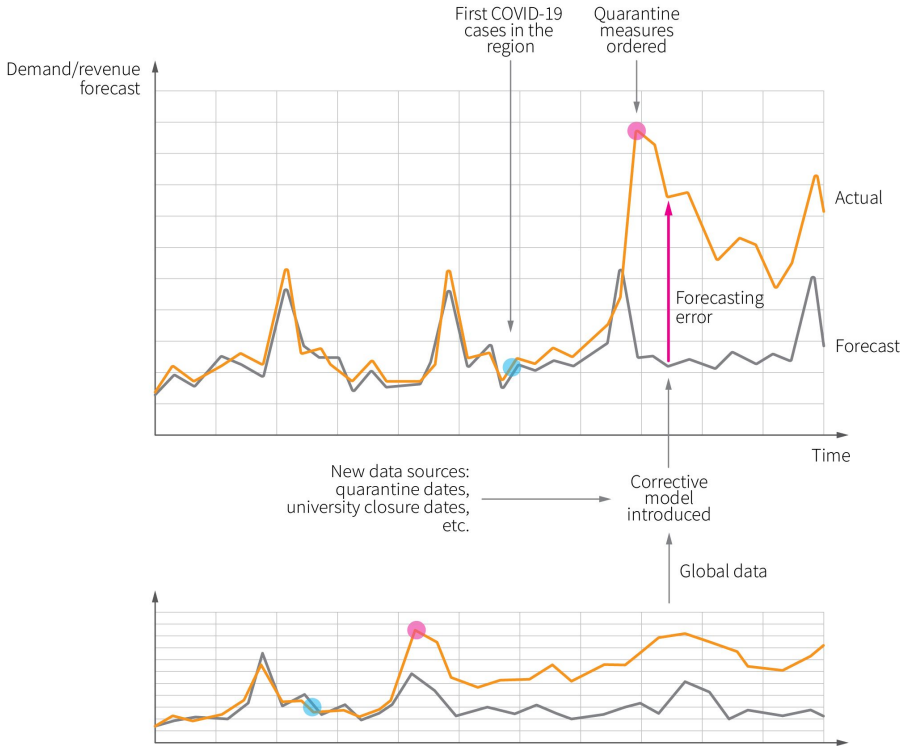

Source: https://blog.griddynamics.com/rapid-response-to-covid-19-supply-chain-and-market-shocks/

GridGain

Short-Term Solutions

- **Apache Ignite:**
  - **In-Memory Computing Platform**
  - **Distributed, Transactional, Write-Through Cache for RDBMS**

- Concentrate today on:
  - **Immediate effect**
  - **Easy-to-implement**
  - **Fit the long-term strategy**
- Long-term solutions remain the same:
  - Heavy cost of re-architecting
  - Strong long term benefits
- Avoid short-term throw-away bandaids

IMC Latency: Memory is Much Faster Than Disk

System Event	Actual Latency	Human Scaled
One CPU cycle	0.4 ns	1 s
Level 1 cache access	0.9 ns	2 s
Level 2 cache access	2.8 ns	7 s
Level 3 cache access	28 ns	1 min
Main memory access (DDR DIMM)	~100 ns	**4 min**
Intel Optane DC persistent memory access	~350 ns	**15 min**
Intel Optane DC SSD I/O	< 10 µs	7 hrs
NVMe SSD I/O	~25 µs	17 hrs
SSD I/O	50-150 µs	**1.5 - 4 days**
Rotational disk I/O	1 – 10ms	**1 – 9 months**
Internet: SF to NY	65 ms	5 years

IMC Throughput: Massively Scalable

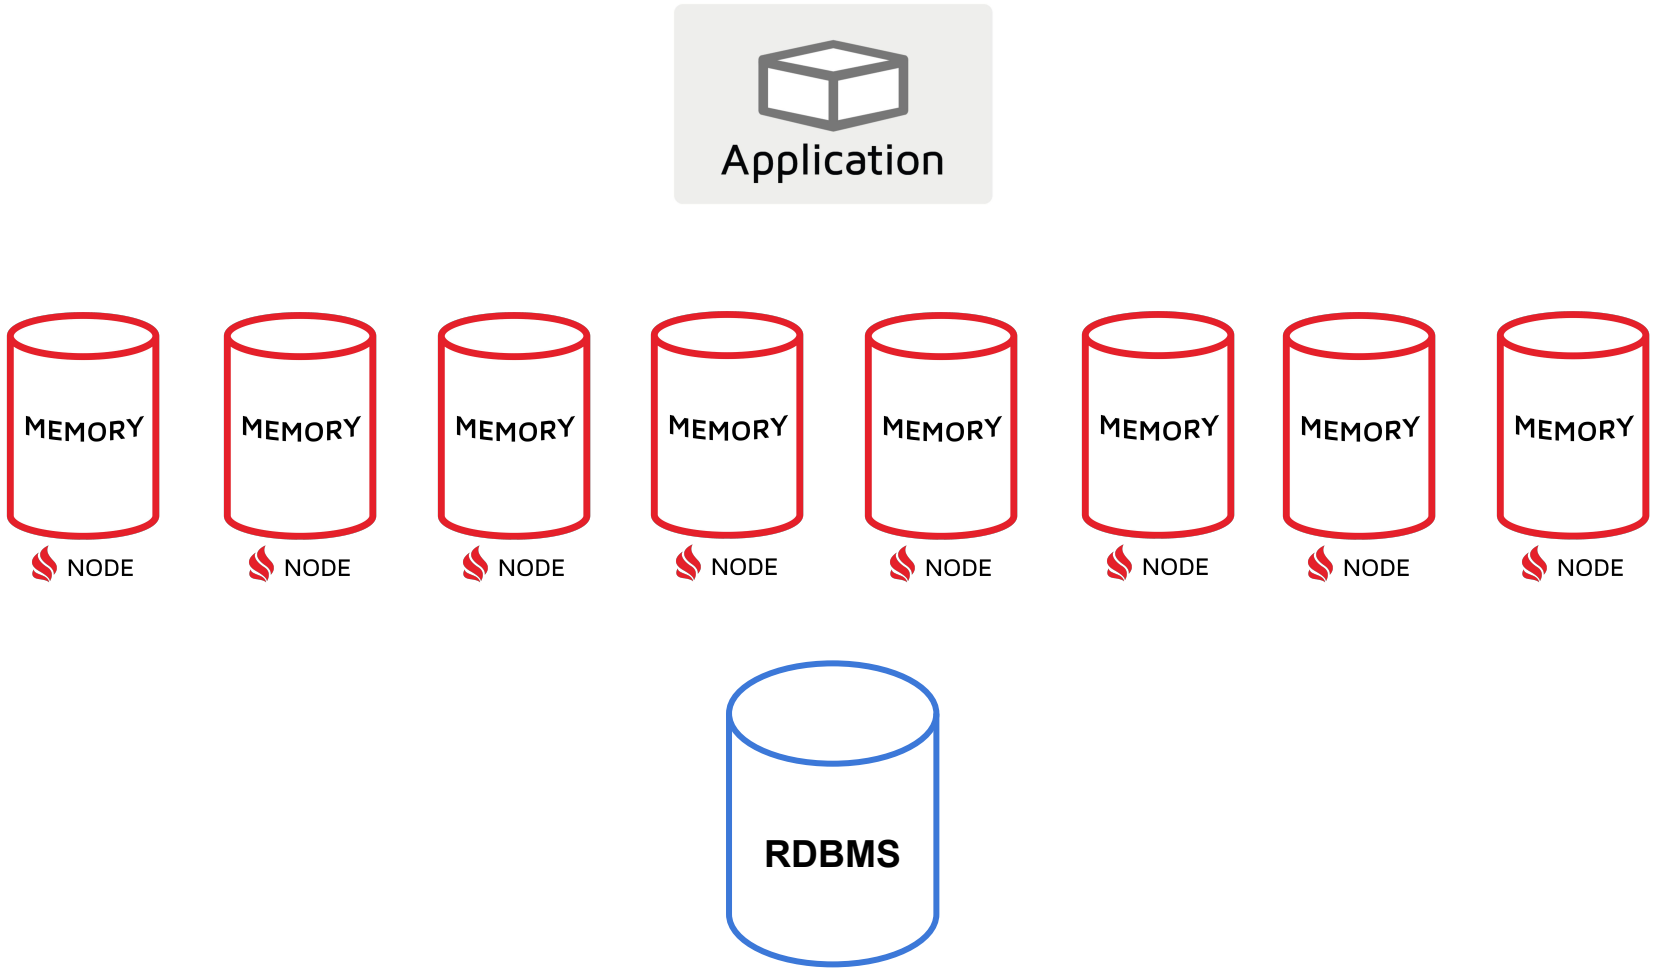

Apache Ignite: Top 5 Apache Projects

A Top 5 Apache Software Foundation Project

Top 5 Dev Mailing Lists	Top 5 User Mailing Lists
1. beam	1. Flink
2. Apache Ignite	2. Lucene
3. kafka	3. Apache Ignite
4. Apache Tomcat	4. cassandra
5. James Enterprise Mail Server	5. kafka

From January 1, 2019 Apache Software Foundation Blog Post:
"Apache in 2018 – By The Digits"

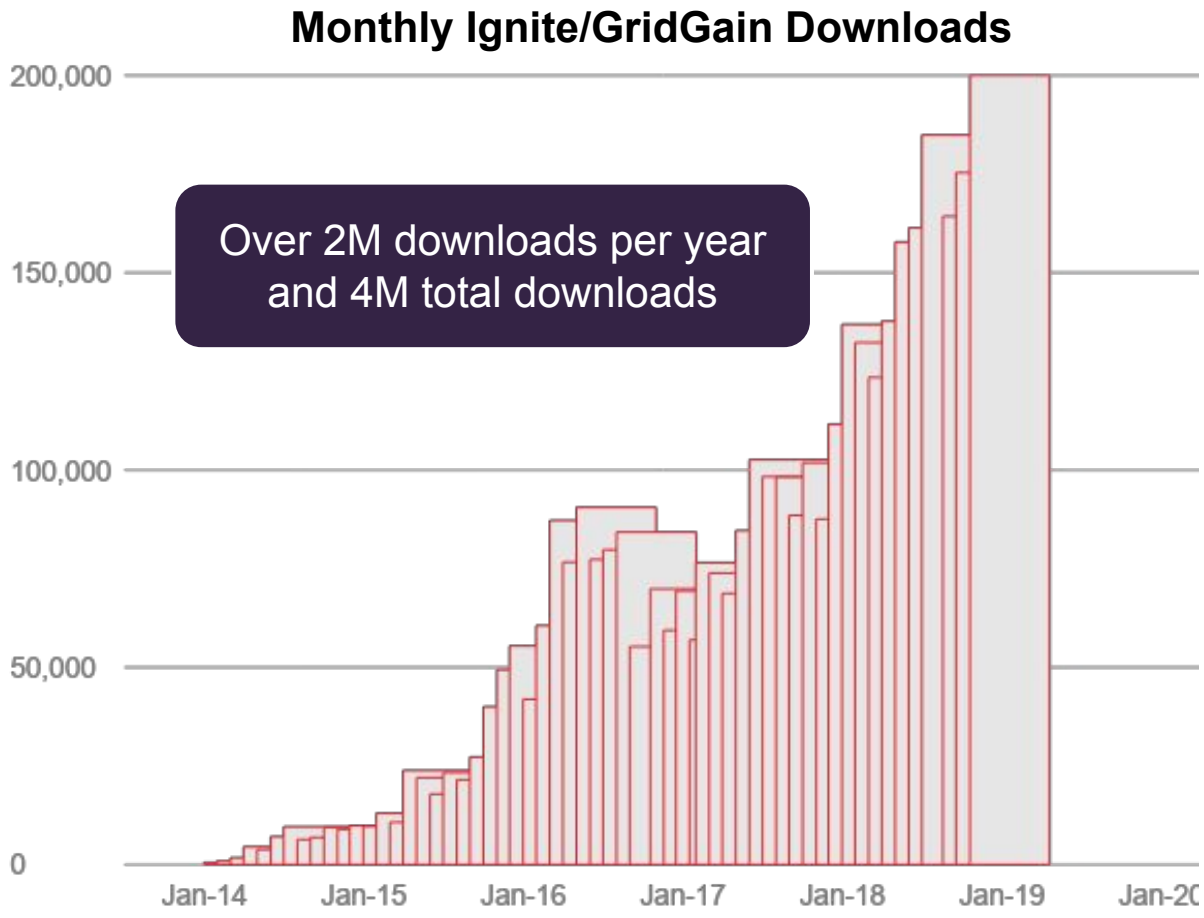

GridGain

Apache Ignite: In-Memory Computing Platform

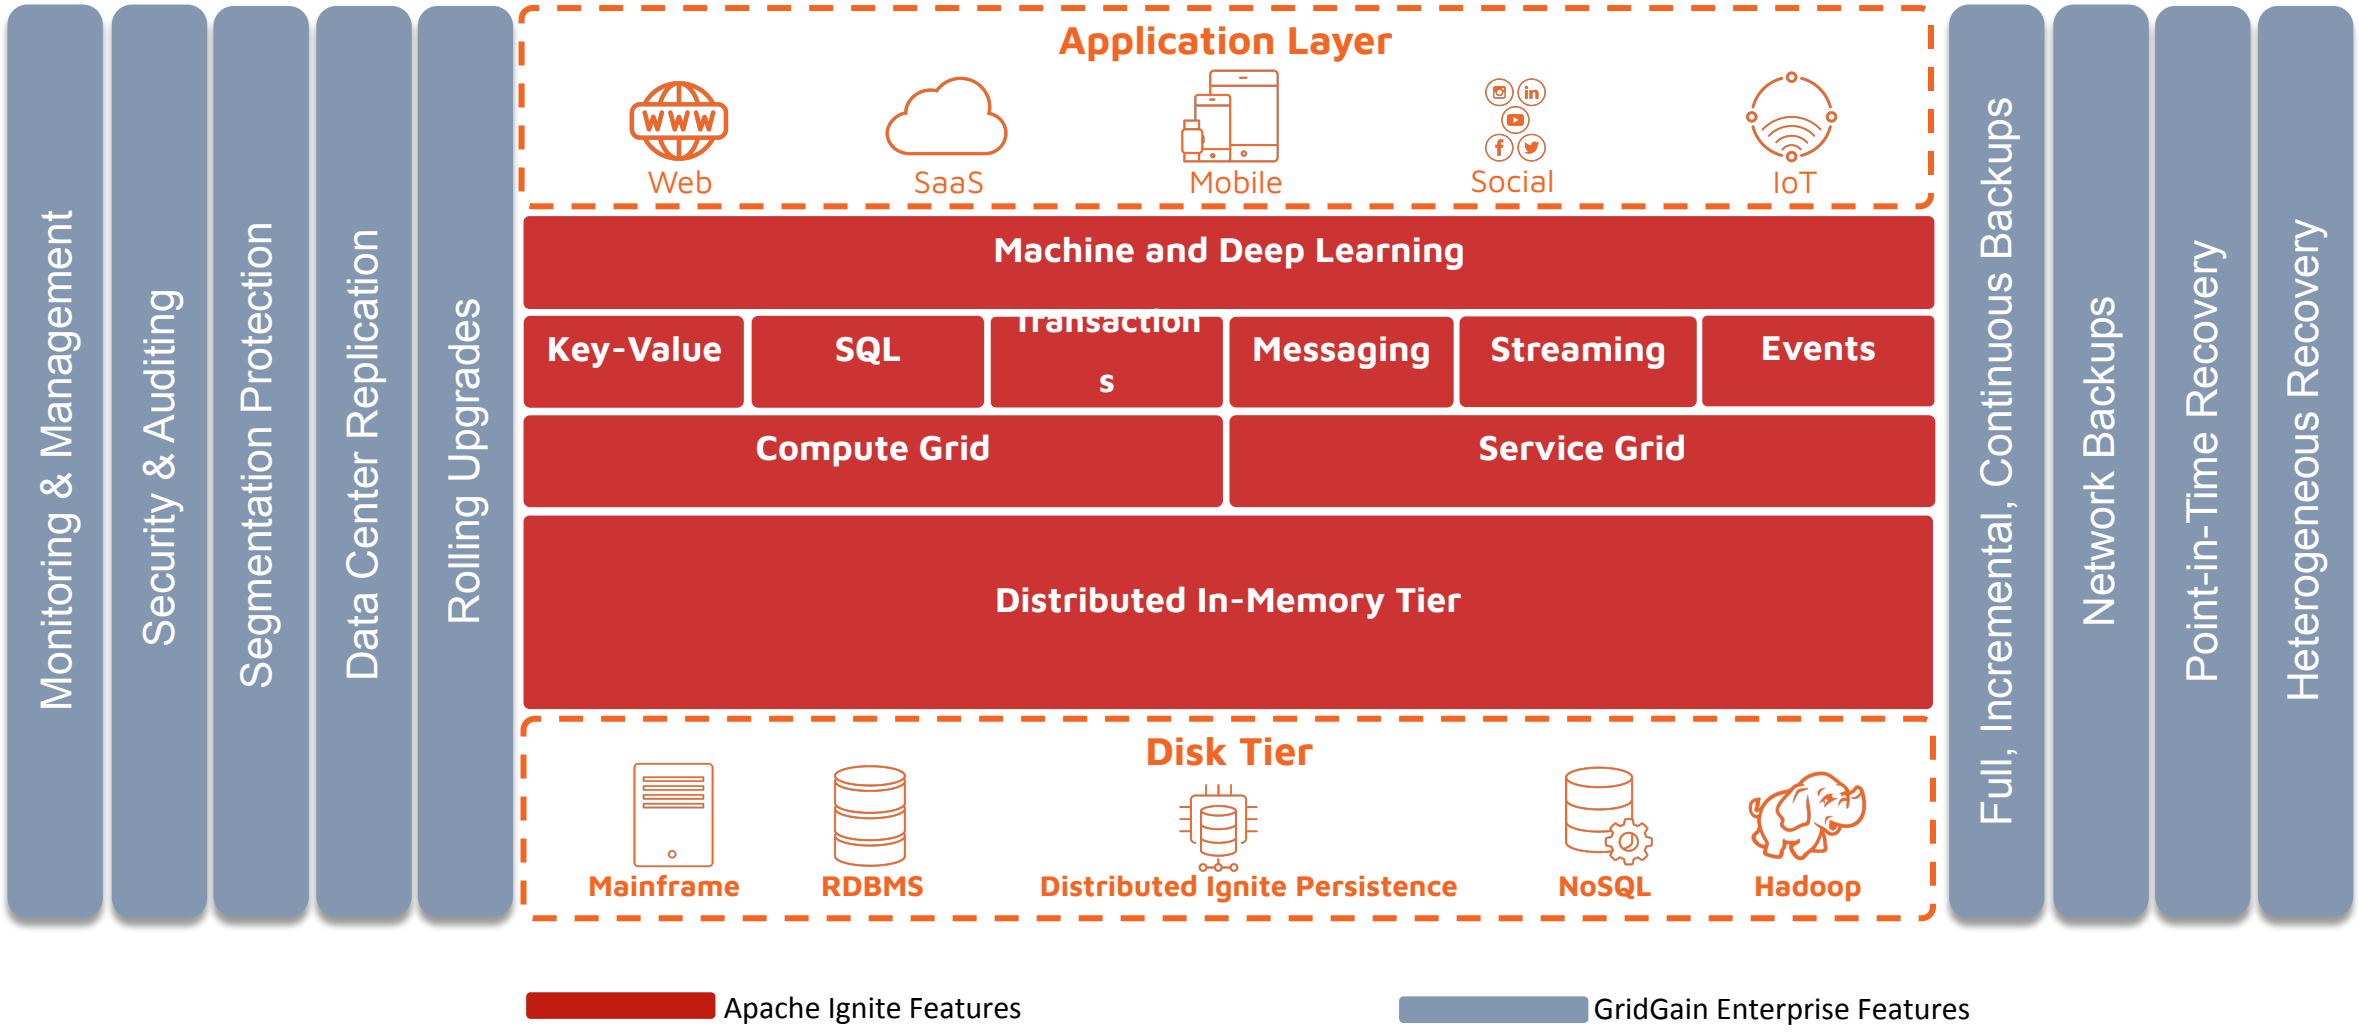

GridGain

Apache Ignite: Distributed Cache or Database

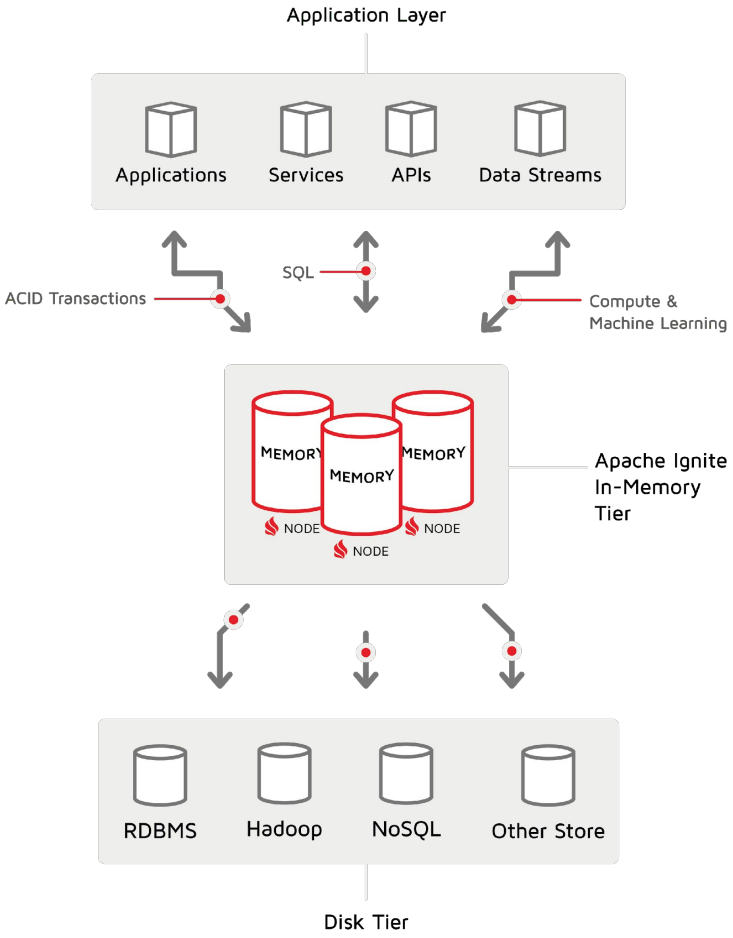

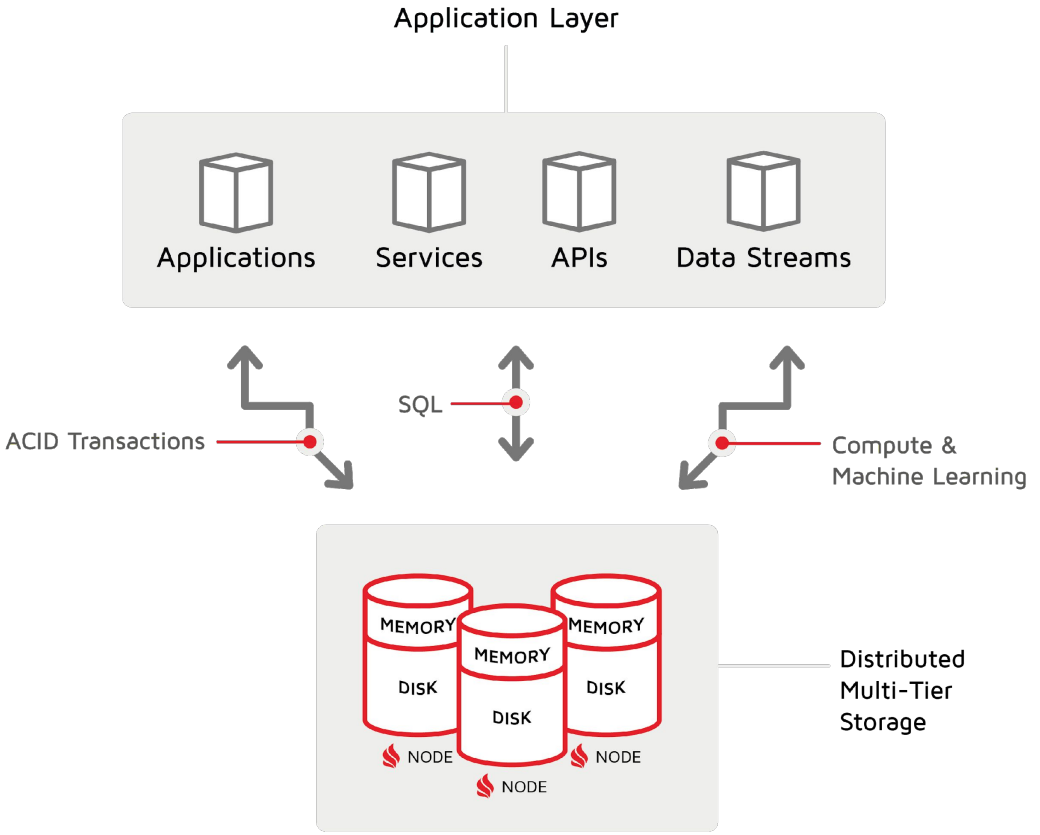

GridGain

Apache Ignite: Distributed Write-Through Cache

- Horizontally Scalable In-Memory Storage
  - ANSI SQL, ACID Transactions, MapReduce
- Automatic Synchronization With RDBMS
  - Write-Through - comparable writes performance
  - Write-Behind - faster writers, eventual consistency
- 10-100x performance increase for reads

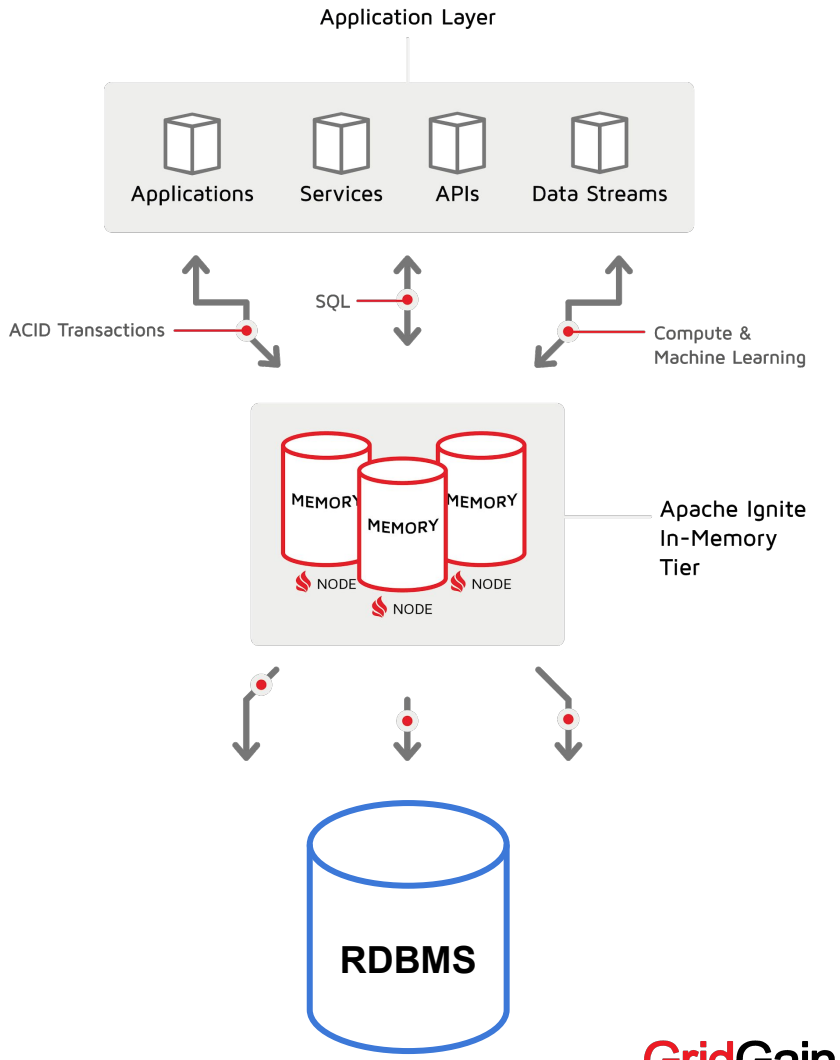

GridGain

Apache Ignite: Default Data Distribution

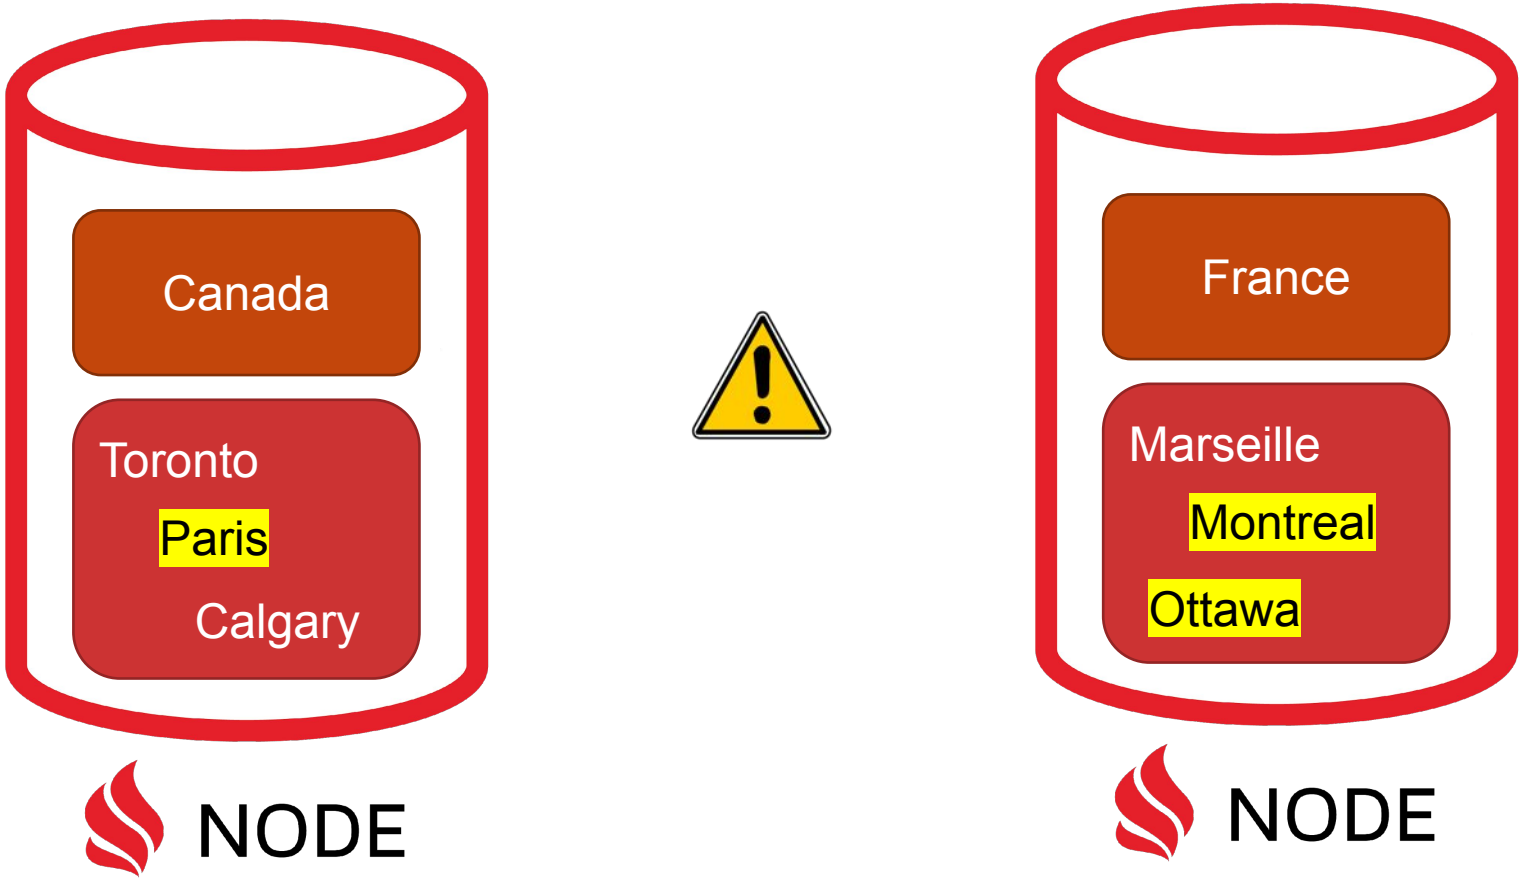

Country Table

City Table

GridGain

Apache Ignite: SQL JOIN With Data Shuffling

Thick Client

1 & 4

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

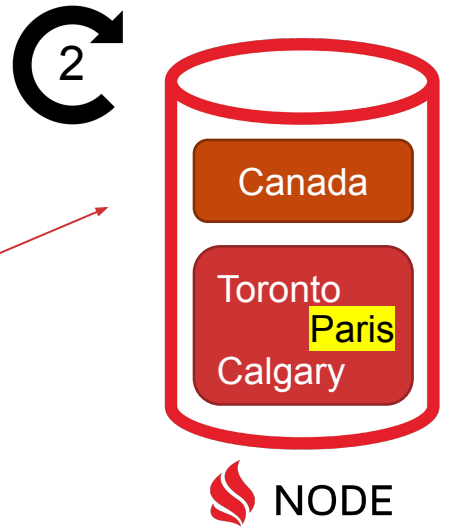

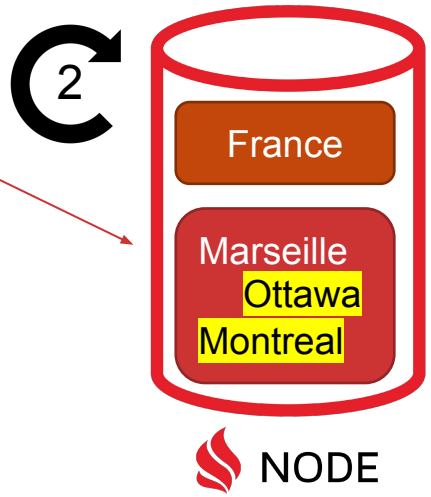

1. Initiating Execution
2. Execution on Servers (map phase)
3. **Data Shuffling**
4. Reduce Phase

GridGain

Apache Ignite: Co-Located Distribution (Affinity Co-Location)

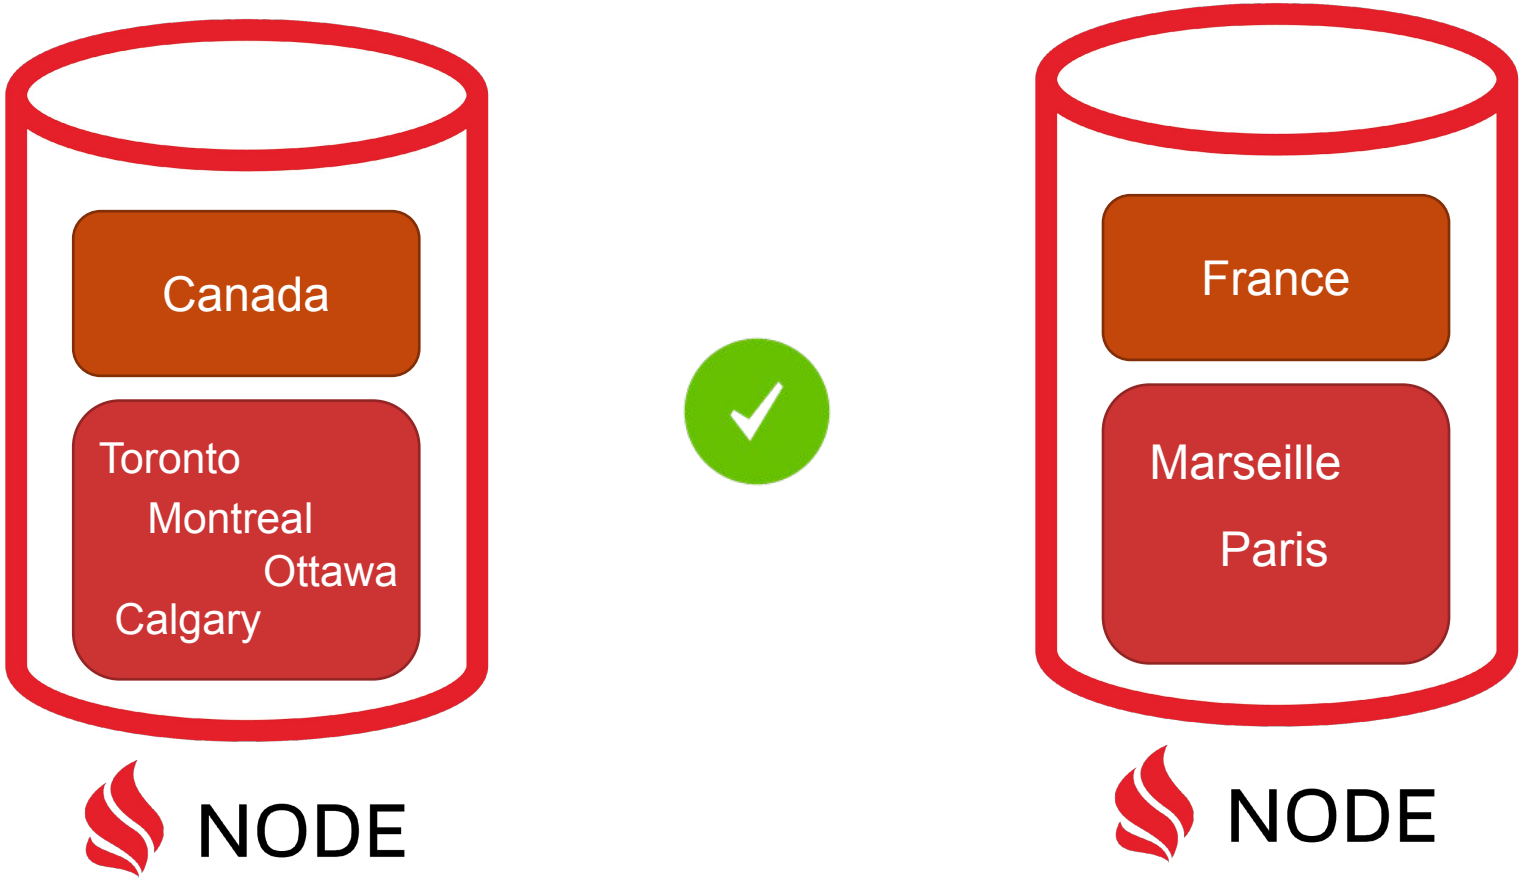

Country Table

City Table

GridGain

Apache Ignite: Configure Affinity Key

```
CREATE TABLE Country (
  Code CHAR(3),
  Name CHAR(52),
  Continent CHAR(50),
  Population INT(11),
  Capital INT(11),
  PRIMARY KEY (Code)
);
```

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

Apache Ignite: Affinity Key to Node Mapping

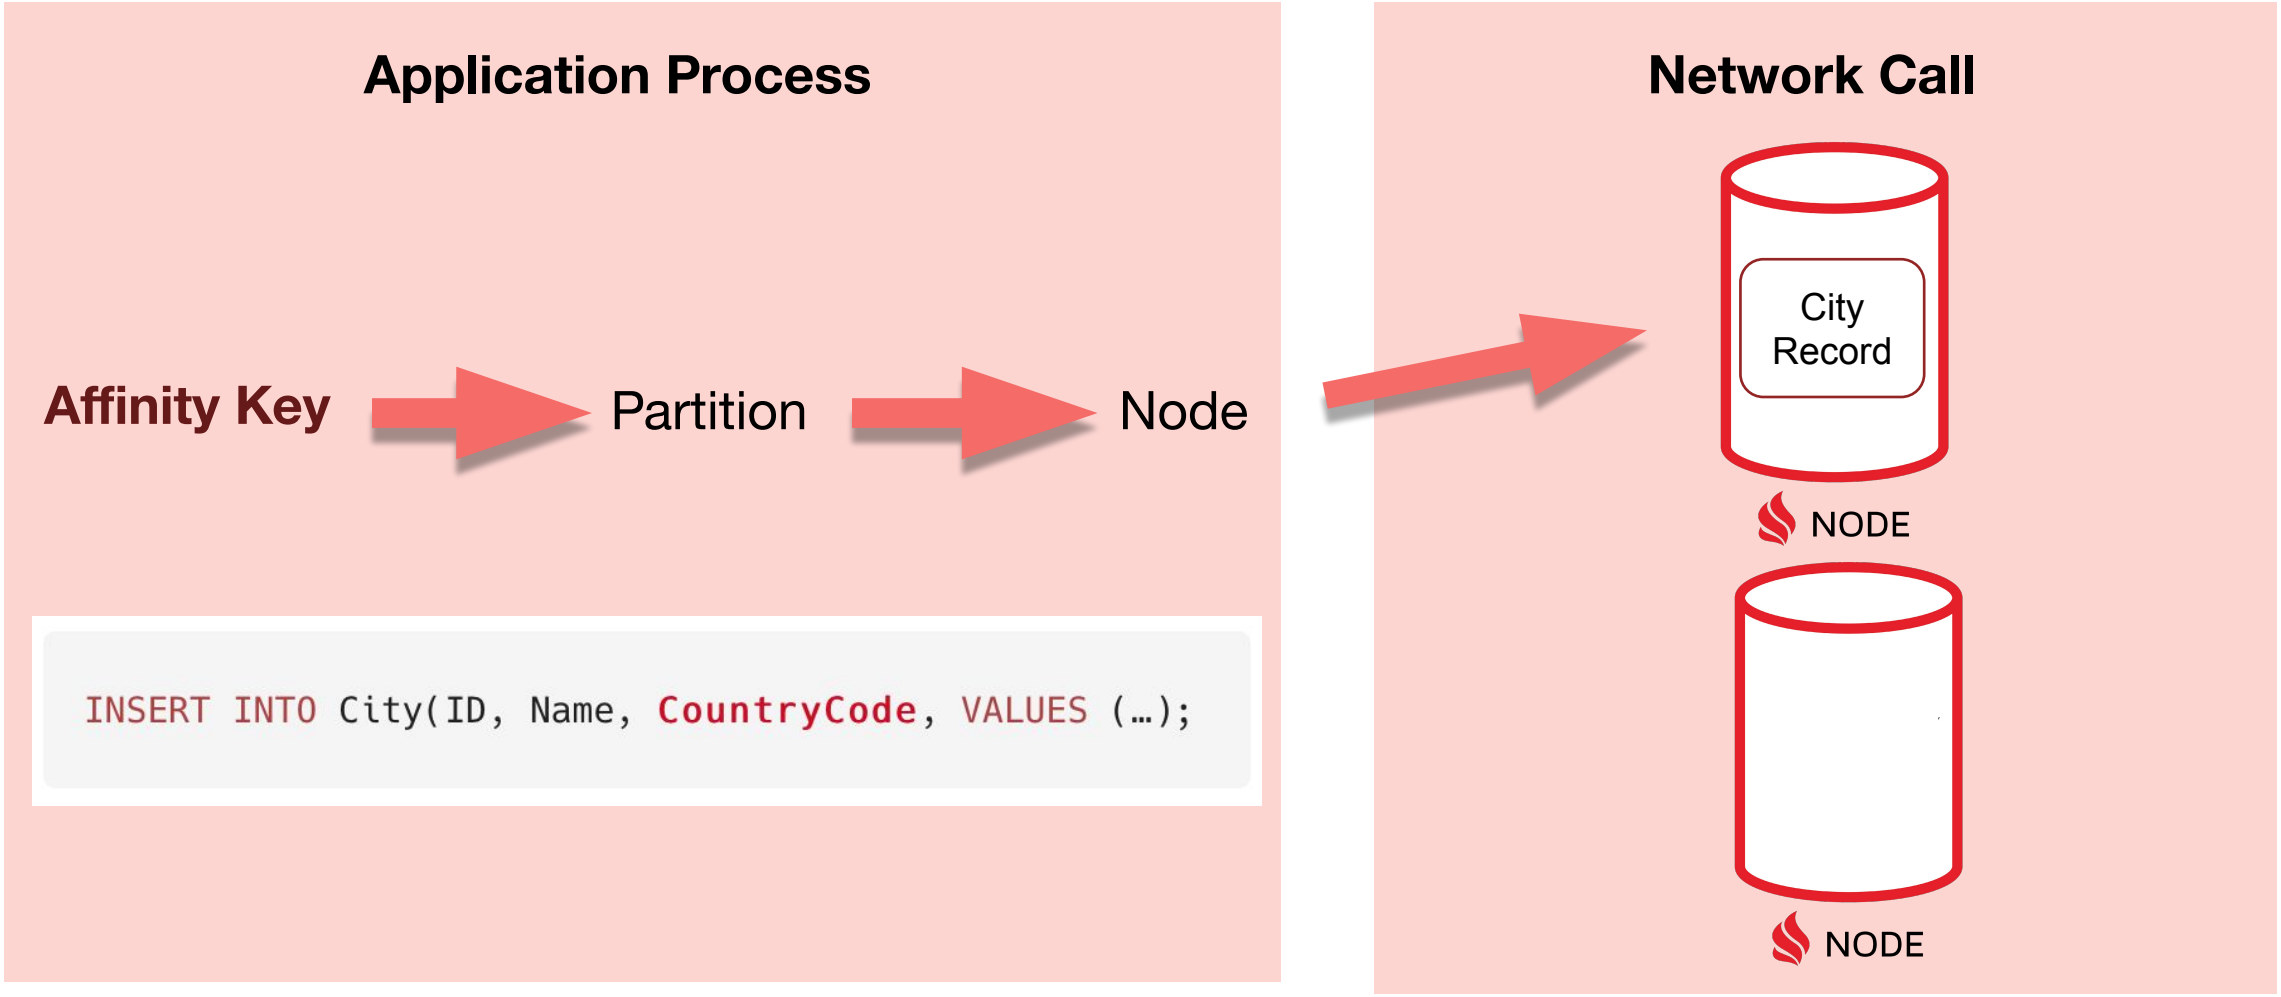

GridGain

Apache Ignite: High-Performance SQL JOIN

Thick Client

1 & 3

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

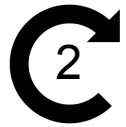

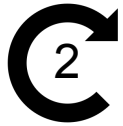

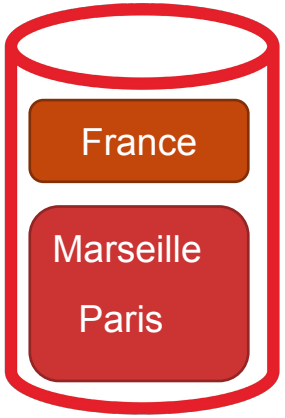

1. Initiating Execution
2. Execution on Servers (map phase)
3. Reduce Phase

GridGain

Apache Ignite: Strong Consistency & Distributed Tx

- Strong or relaxed consistency
  - Transactional and atomic caches/tables
- Distributed ACID Transactions
  - 2-phase commit protocol
- Transactions supported for key-value APIs
  - MVCC for SQL transactions is in experimental mode

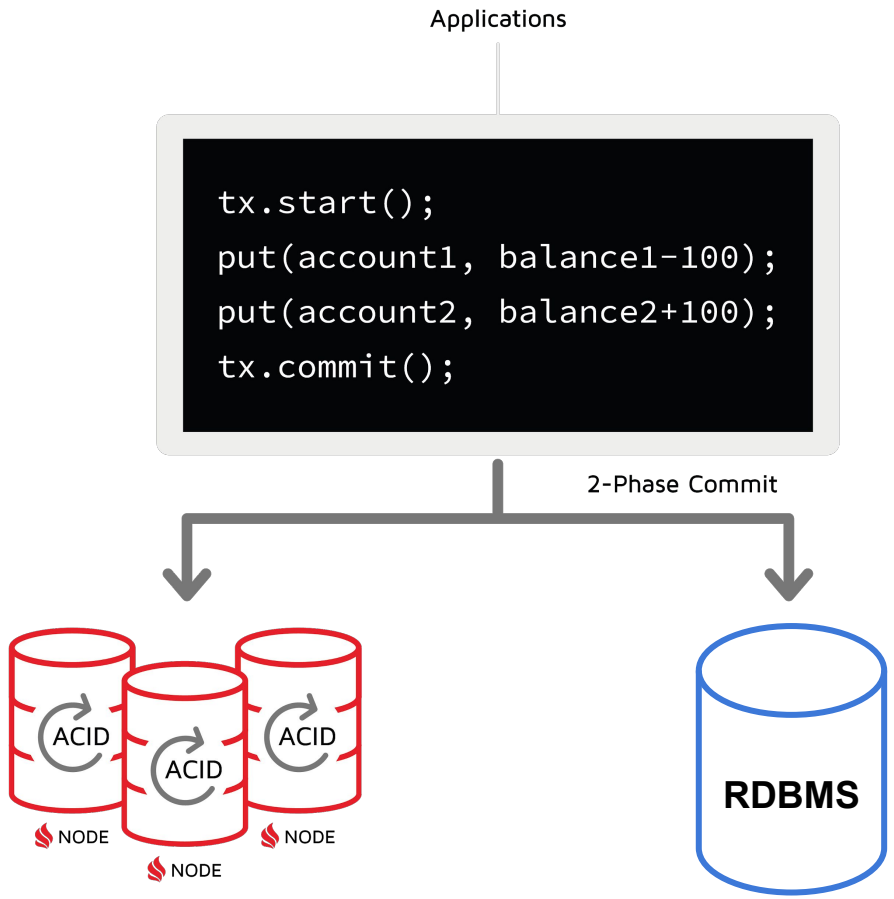

GridGain

GridGain: Enterprise Security

- Open source
  - Data encryption in flight – SSL/TLS
  - Data encryption at rest
- Pluggable auth & auth
  - JAAS, LDAP, Active Directory, Kerberos, etc.
  - Cluster node authentication
  - Fine-grained authorization
- Multi-tenancy
- Comprehensive auditing

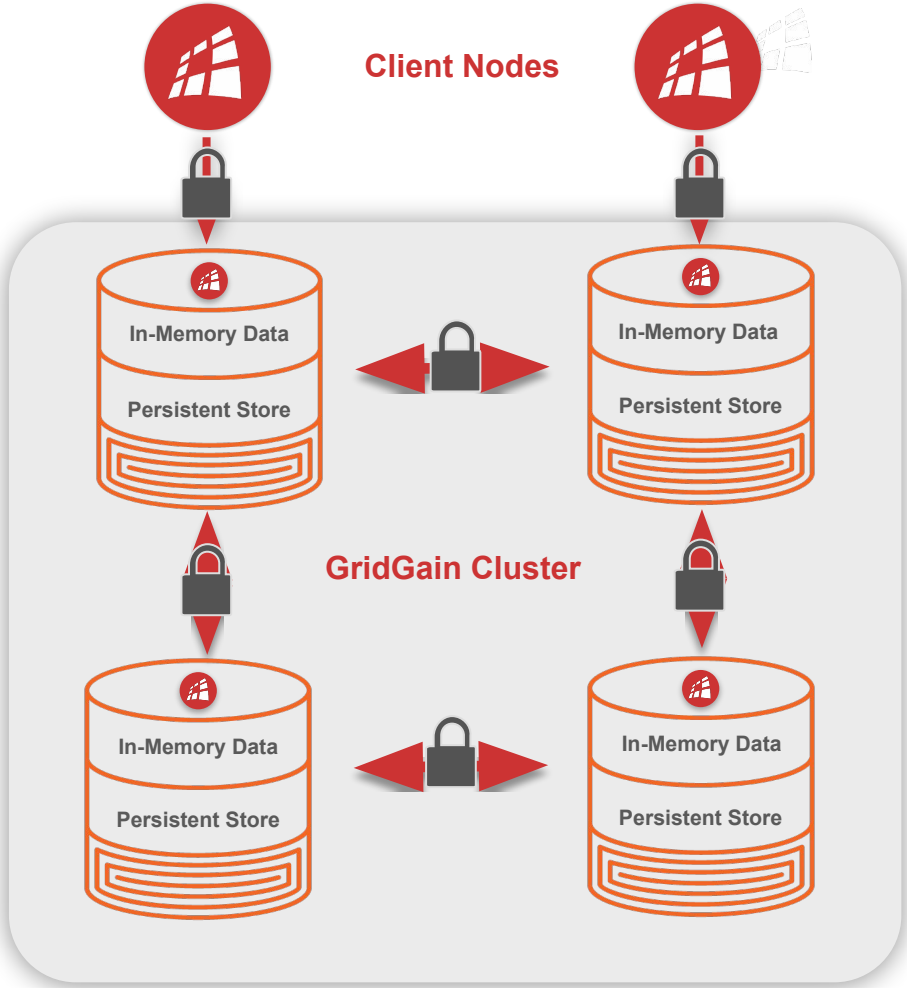

GridGain

GridGain: Multi Datacenter Replication

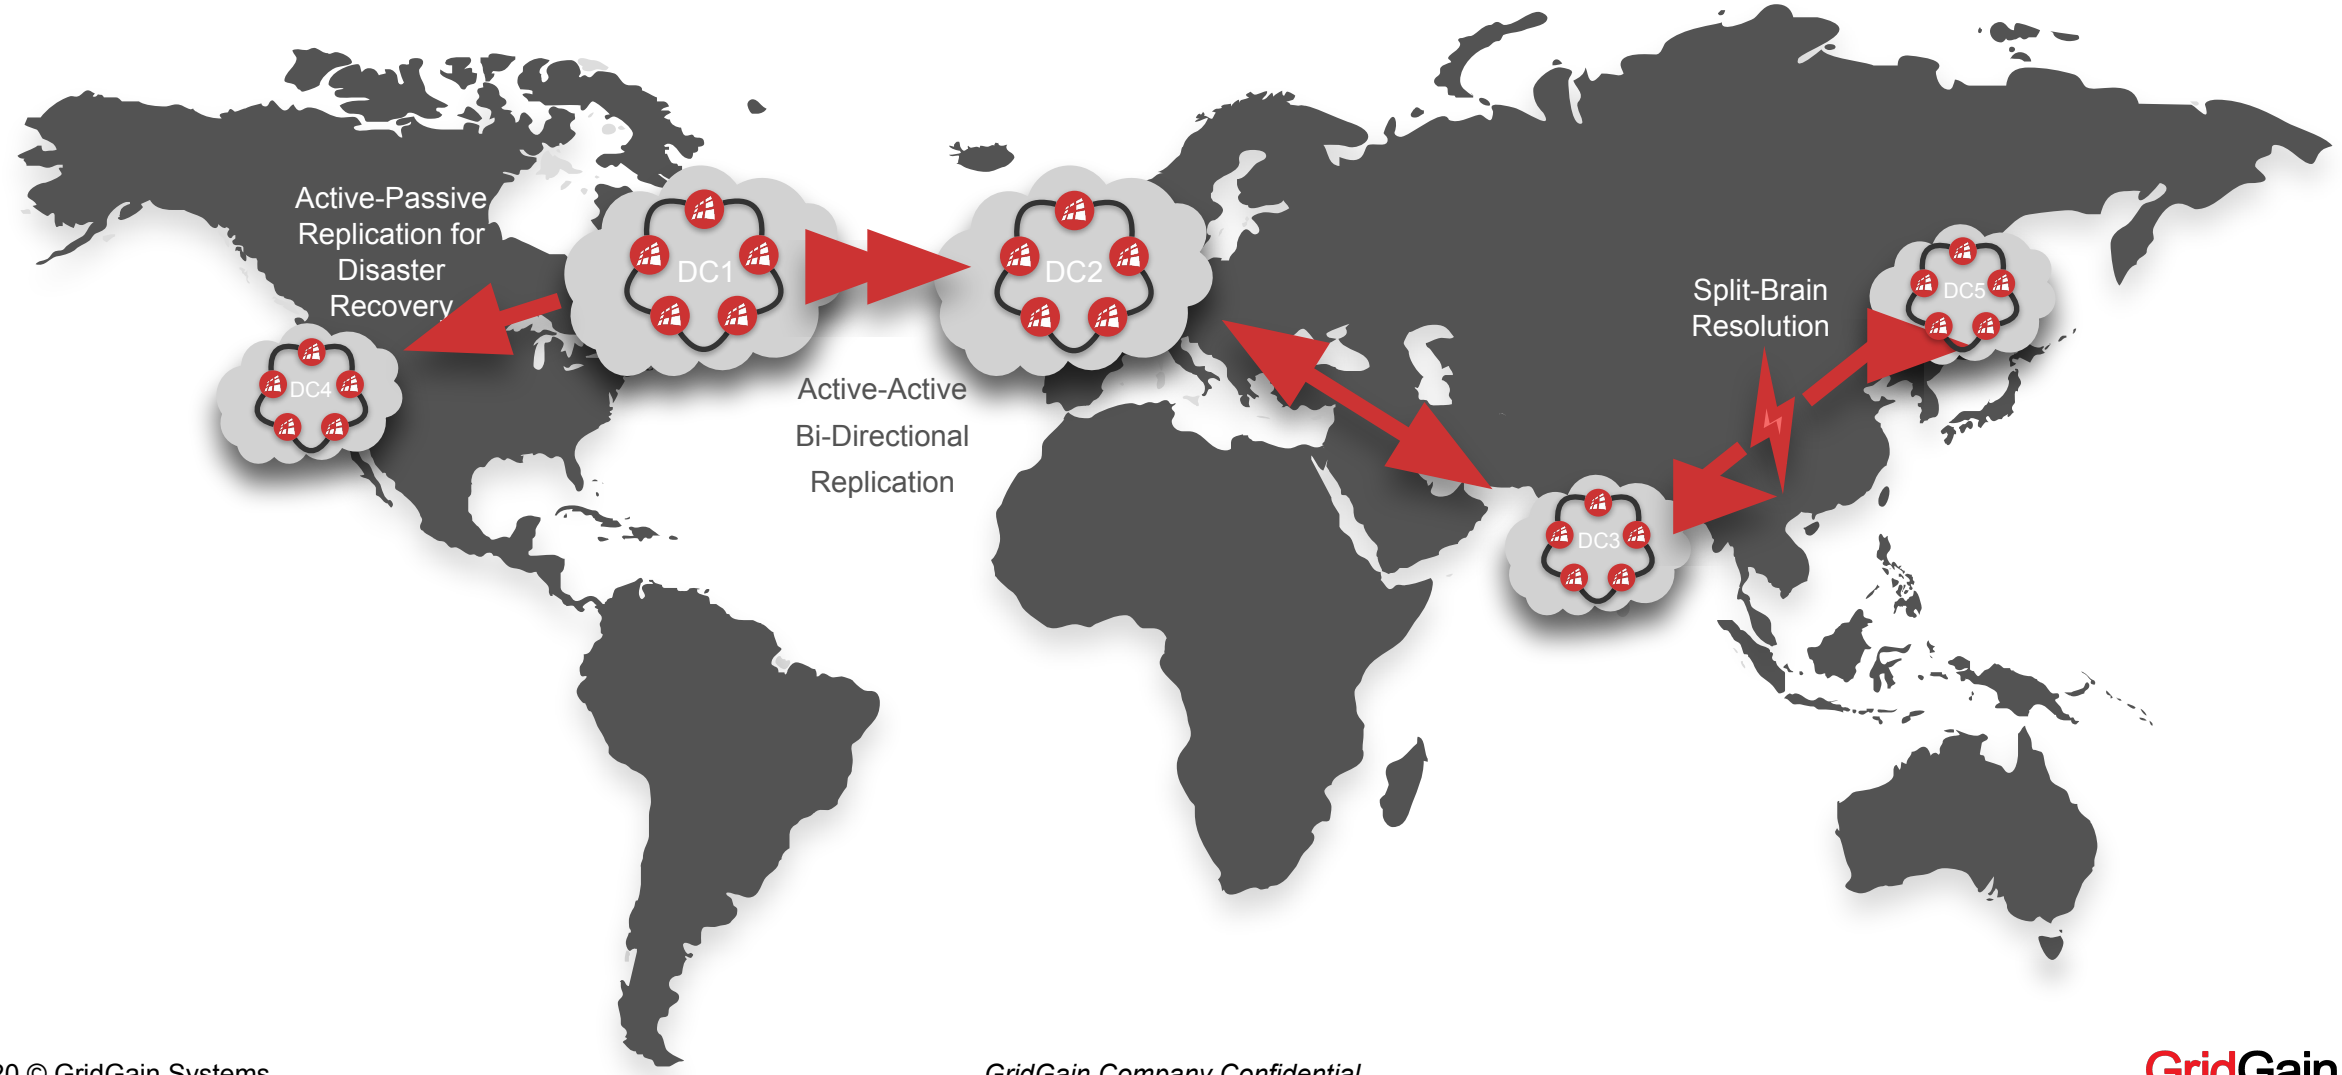

GridGain

Thank you & stay healthy!

ignite.apache.org
www.gridgain.com

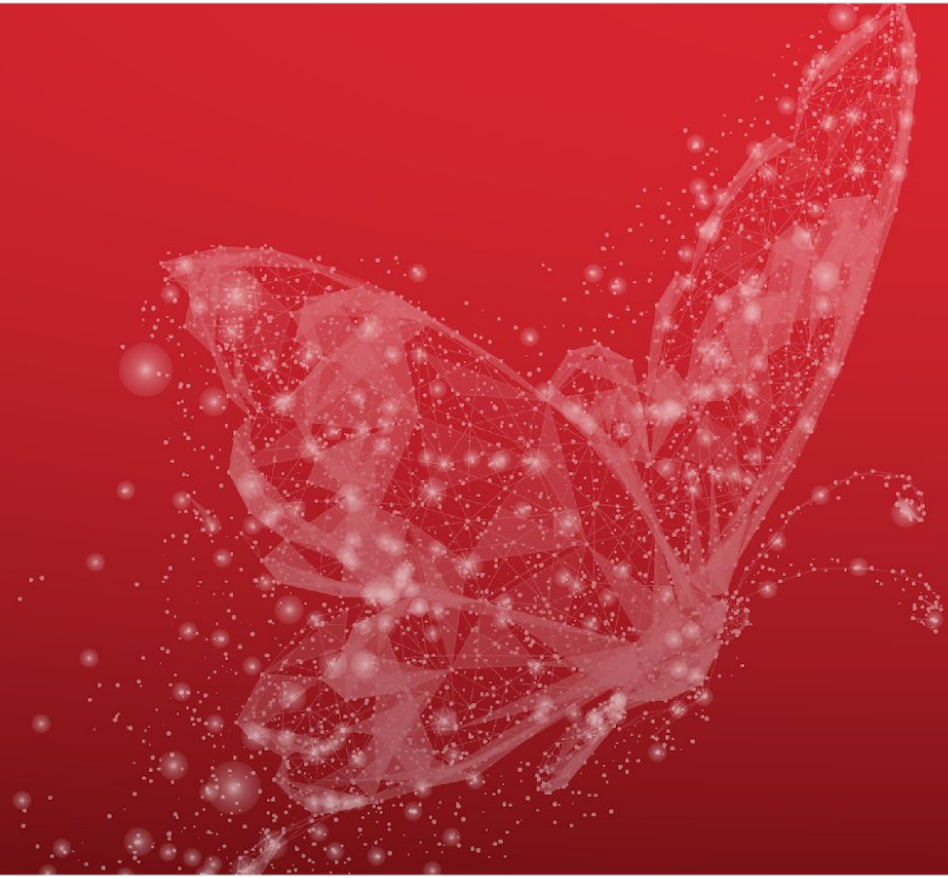

GridGain