

# Adding Speed and Scale to Existing Applications: **With Apache Ignite**



Denis Magda  
Apache Ignite PMC Chair  
GridGain Product Management

Rob Meyer  
GridGain Outbound  
Product Management

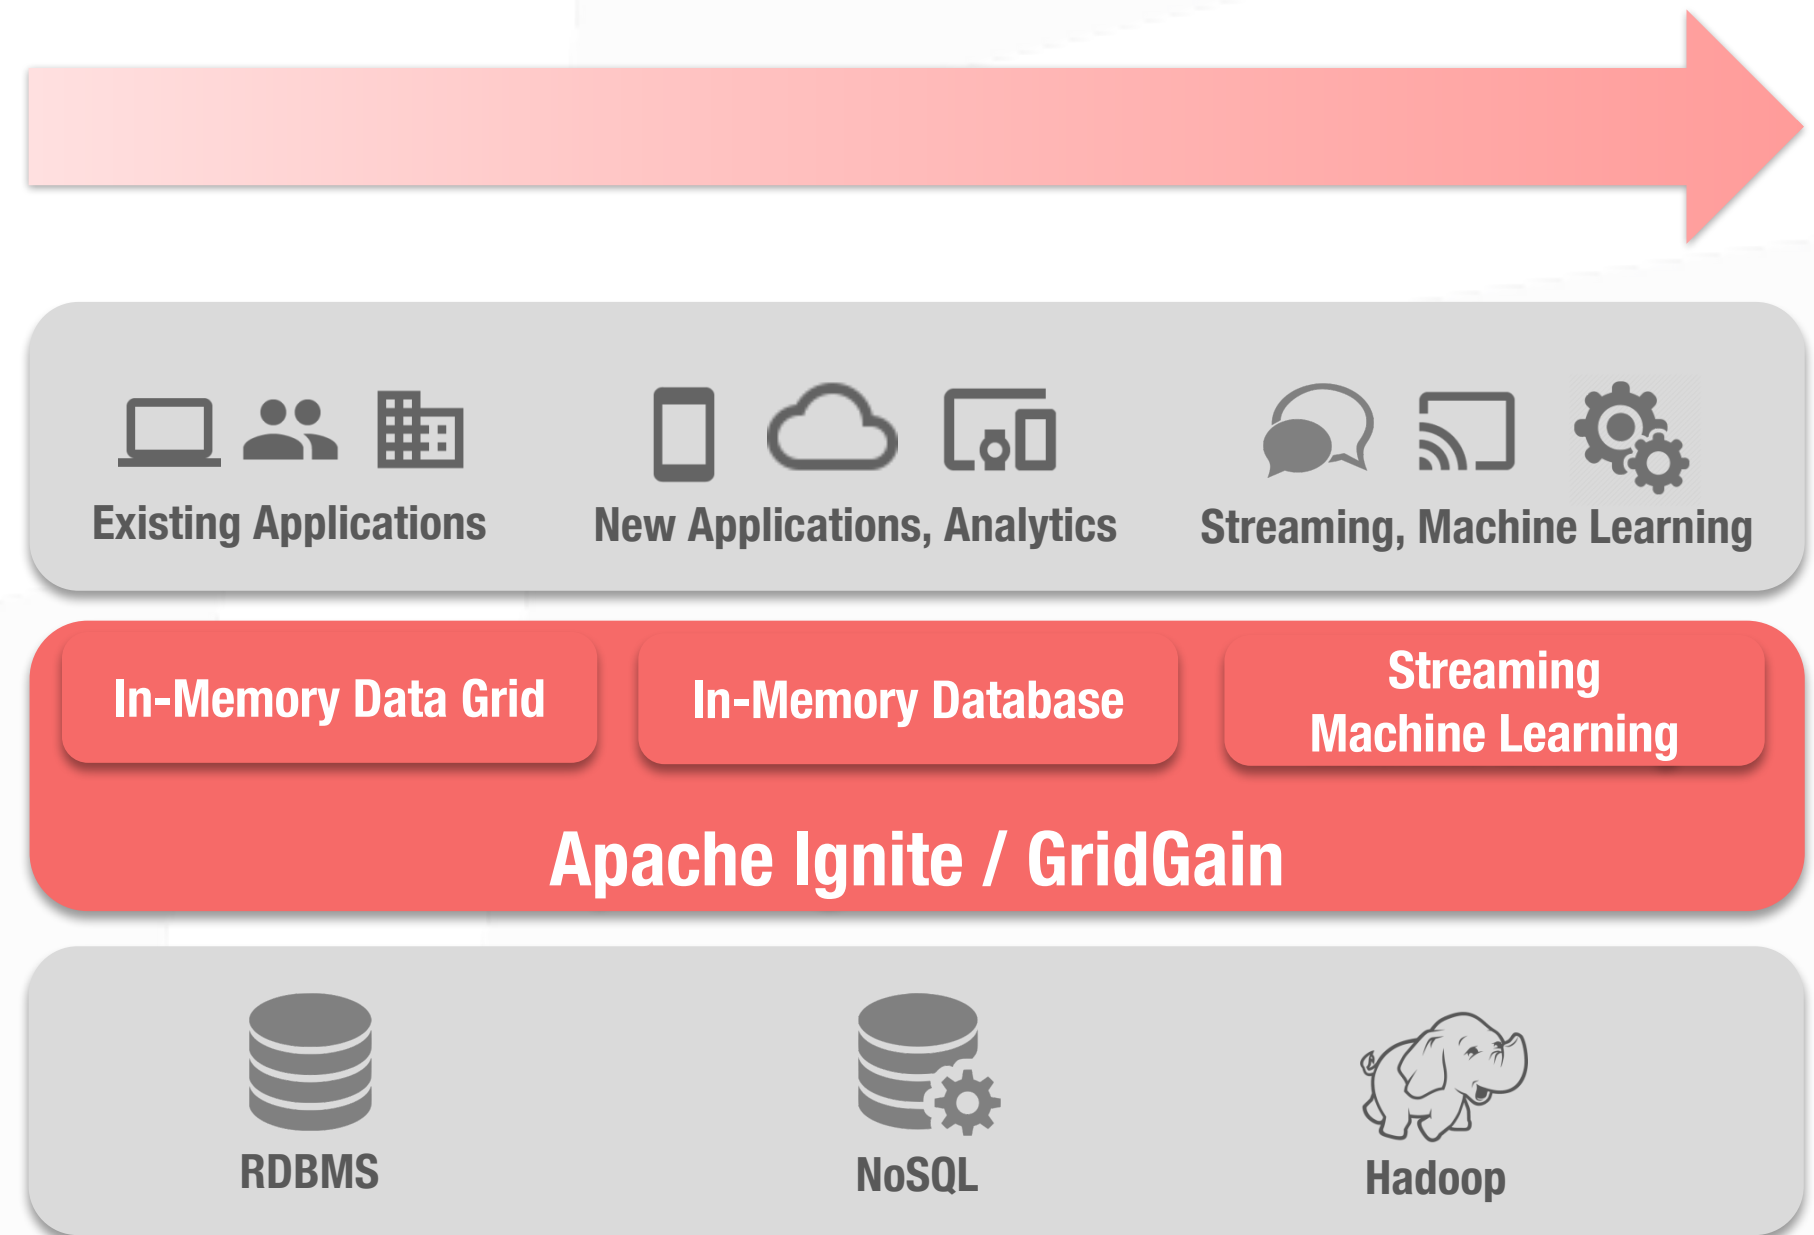
# Agenda

- Reasons to Bring In Ignite?
- Ignite Usage Modes
- Step 1: In-Memory Data Grid
  - Keeping Data in Sync
  - Preserving SQL and Transactions
- Demo
- Step 2: Full Switch to HTAP
- Q & A

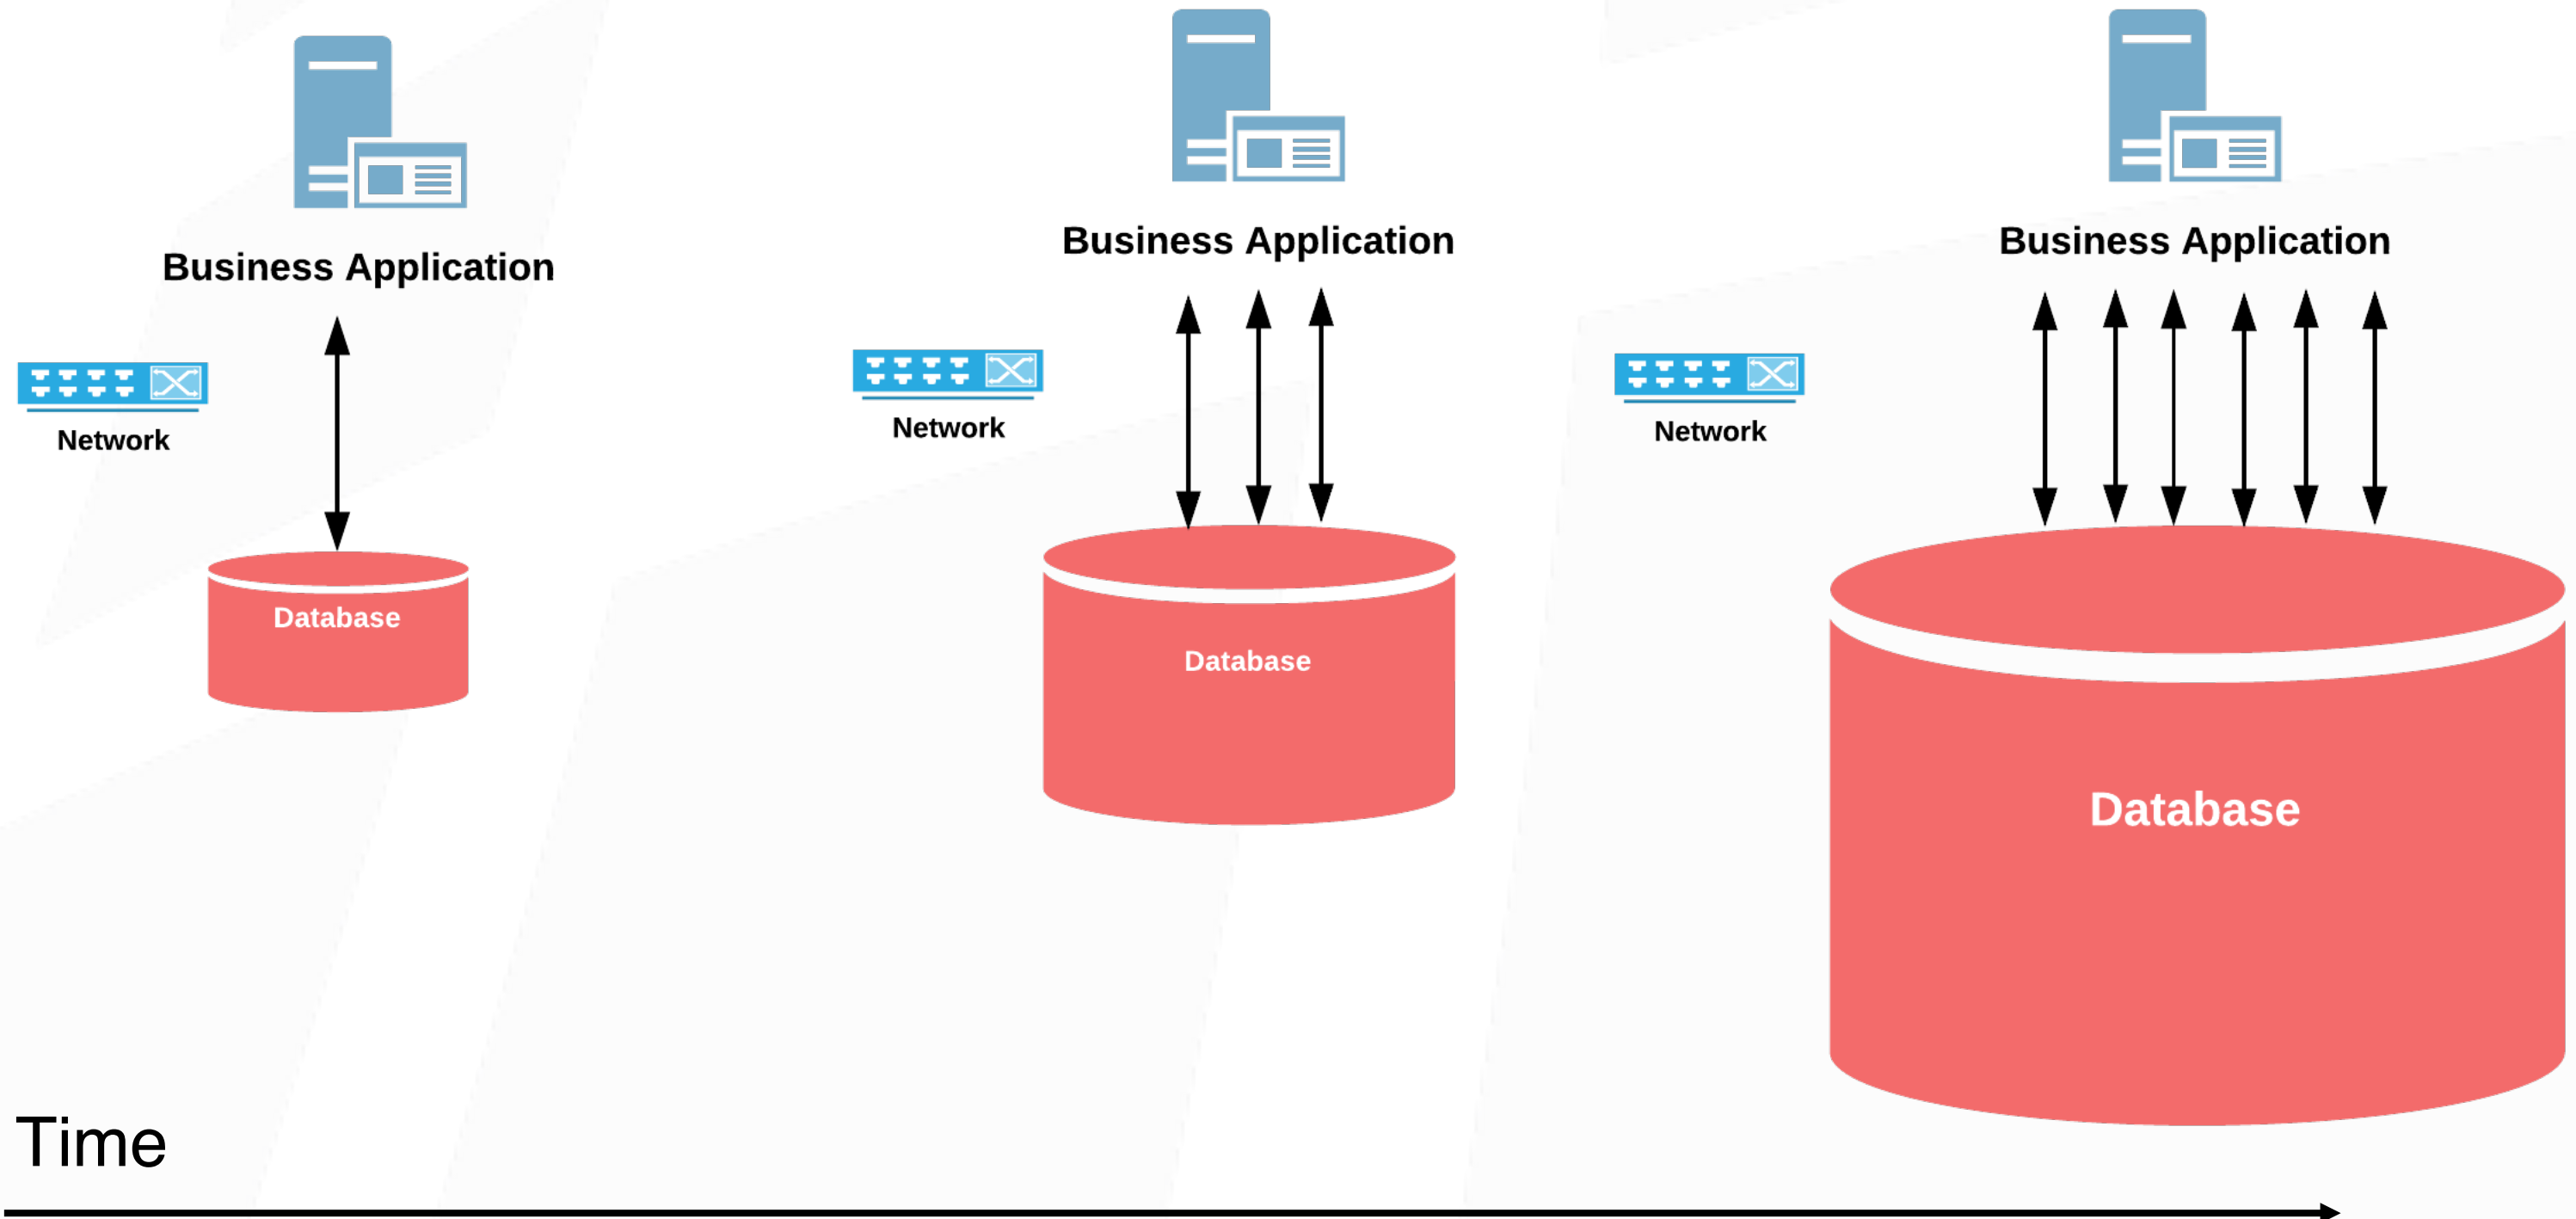
# Reasons to Bring in Ignite?

# Adopting In-Memory Computing – Best Practices

- Adding Speed and Scalability to Existing Applications
  - Web, Mobile (MBaaS)
- Building new applications
  - Data Services for Digital Business
  - Big Data (e.g. Personalization)
  - SaaS and Cloud computing
  - Internet of Things (IoT)
- Building streaming analytics, machine and deep learning
  - Ingestion, computing, analytics
  - Real-time compliance
  - Spark acceleration
  - Continuous model training, automation




# Reasons to Bring in Ignite?



# Scalability and Performance

- Performance Drops
  - Disk I/O
  - Saturated Network
- Scalability Hits a Wall
  - Expensive
  - Limited



# Ignite Usage Modes

# Ignite Usage Modes


- Ignite as In-Memory Store
  - High-Performant In-Memory Computing Platform
- Ignite as In-Memory Data Grid
  - Accelerating Existing Architecture
- Ignite as Memory-Centric Database
  - Modern Greenfield Applications



**In-Memory Store**



**In-Memory Data Grid**




**Memory-Centric Database**



# Step 1: Adding Speed and Scale


# In-Memory Data Grid (IMDG)

- Distributed Cache with Brains
- RAM - primary storage
- Disk – backup copy
- Benefits
  - Boosting and Offloading Reads
  - Horizontal Scalability
  - Database Read/Write-Through
  - Collocated Processing



# Ignite as IMDG

- No rip-and-replace
  - Keep your Database
- Read-Through from Database
  - Key-Value APIs
- Write-Through to Database
  - Key-Value APIs
  - SQL INSERTs, UPDATEs, DELETEs
- ACID Transactions
- Collocated Processing



# Distributed SQL

Cross-platform  
Compatibility

Java

.NET

C++

PHP

REST

DDL & DML  
Support

JDBC

ODBC

SQL

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

## Memory-Centric Storage

DURABLE MEMORY

DURABLE MEMORY

DURABLE MEMORY

Indexes on  
RAM or Disk



Dynamic  
Scaling

ON-DISK

ON-DISK

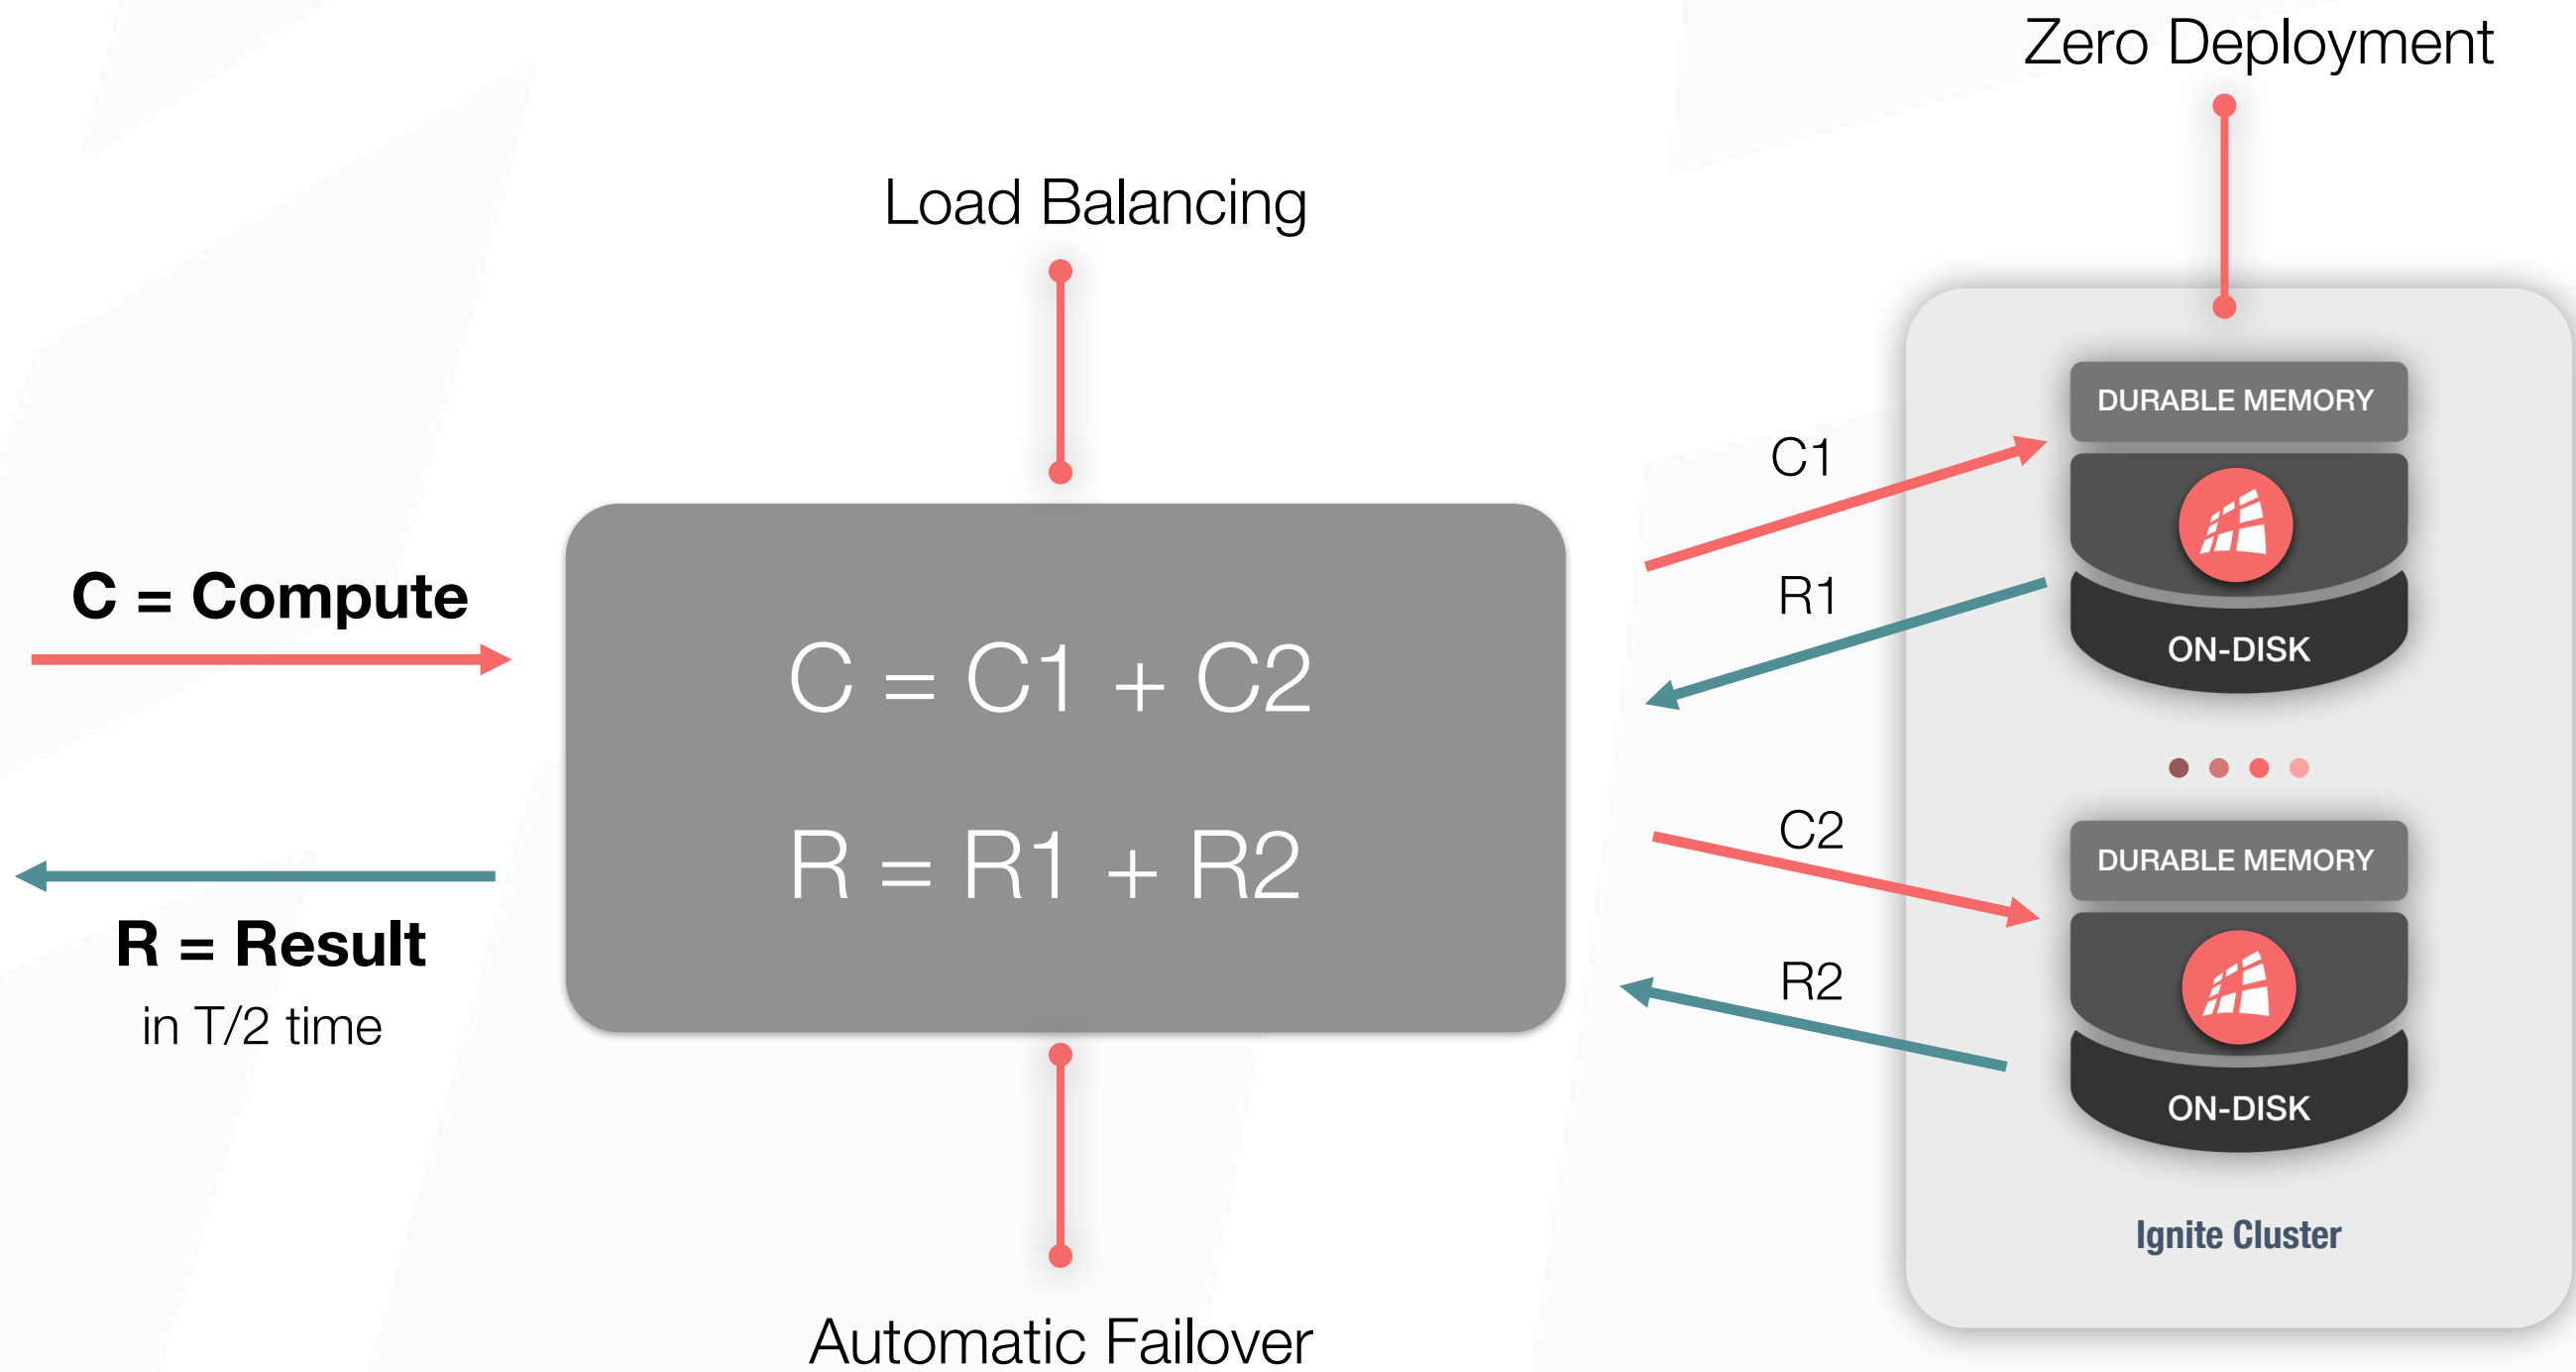
ON-DISK

Server Node

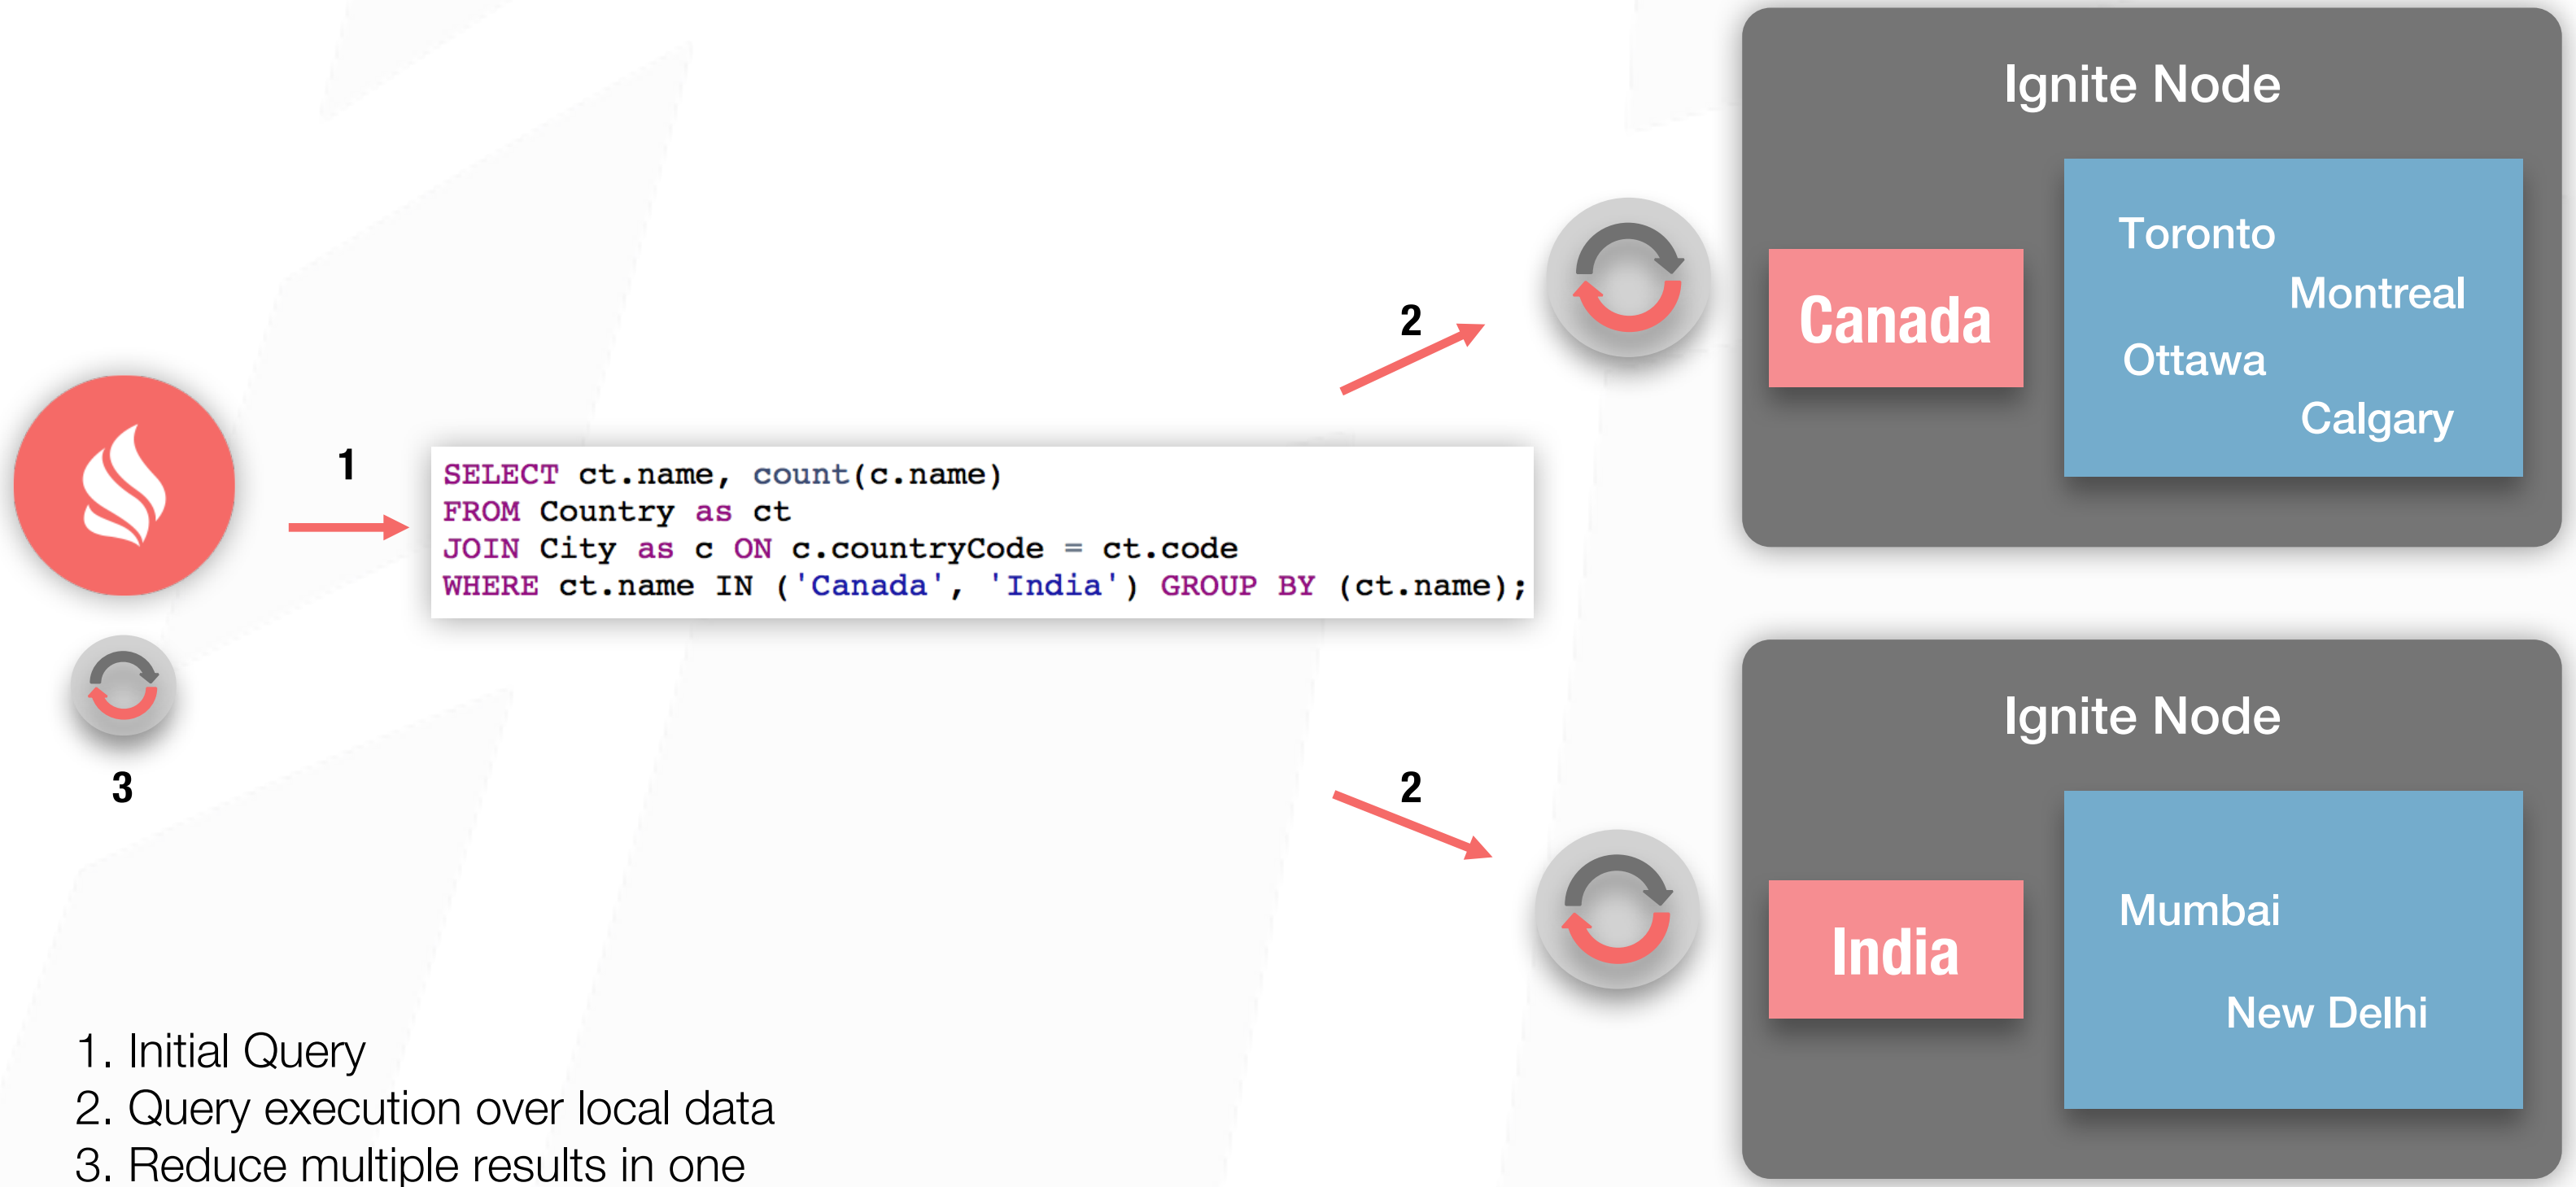
Server Node

Server Node

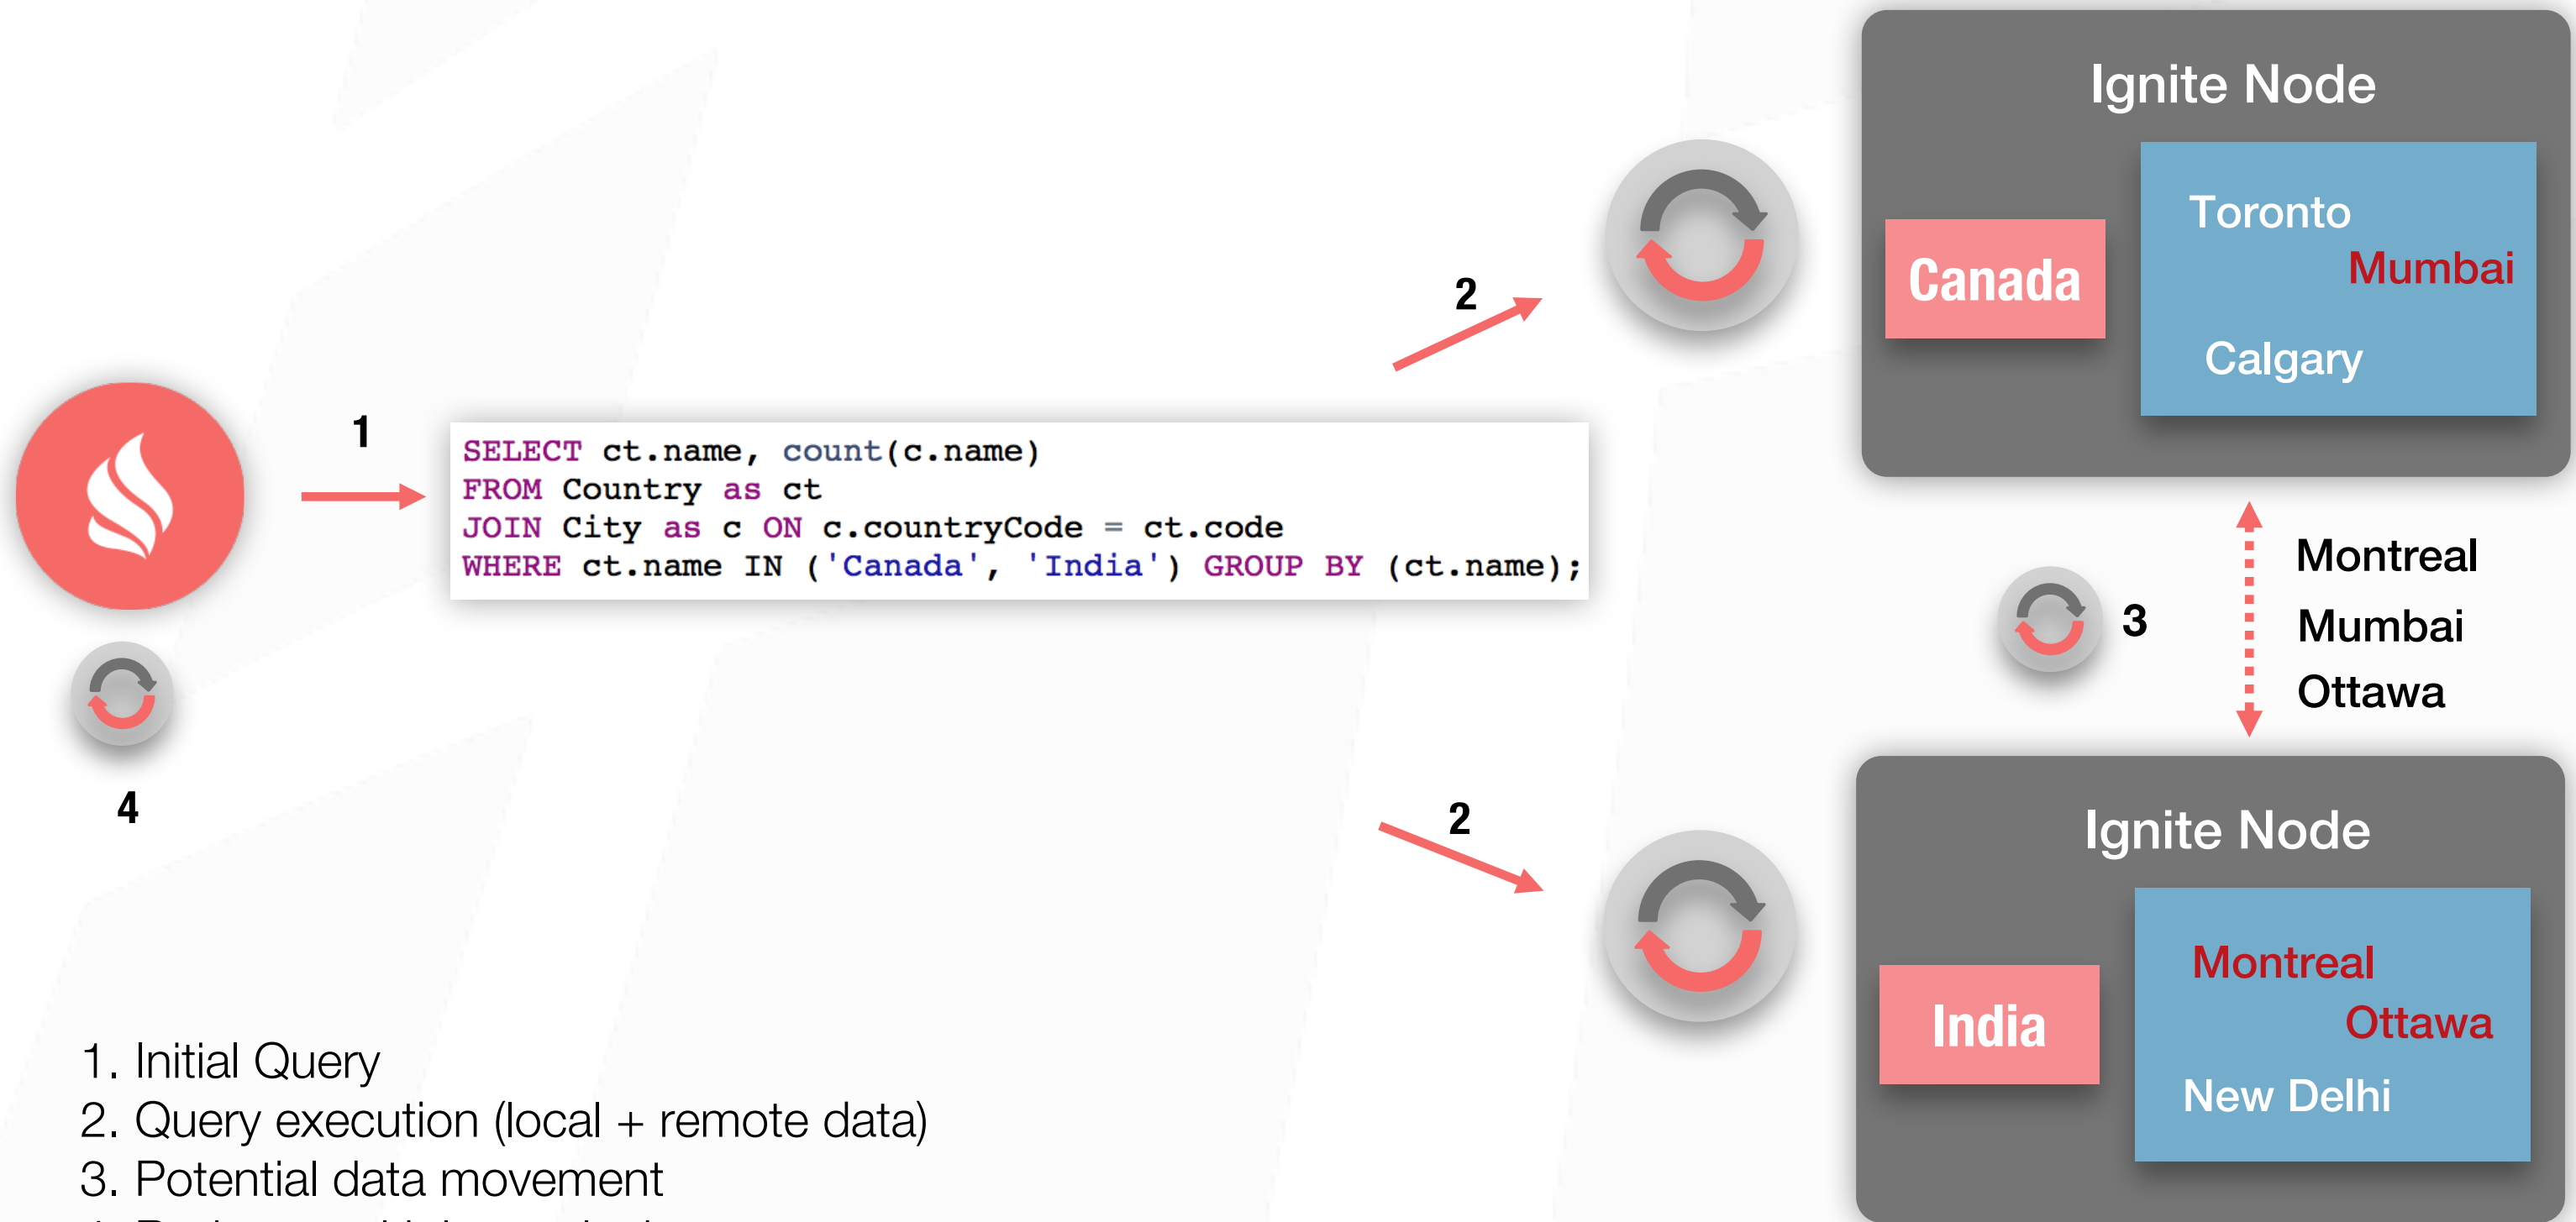
# Compute Grid



# Collocated Joins



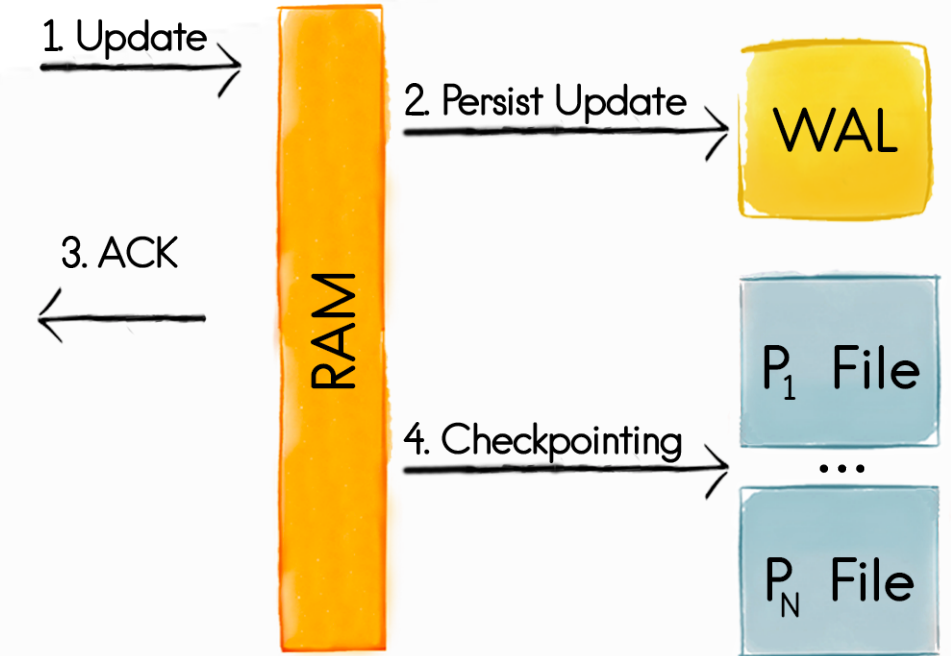
# Non-Collocated Joins



1. Initial Query
2. Query execution (local + remote data)
3. Potential data movement
4. Reduce multiple results in one

# ACID Transactions

- Distributed ACID Transactions
  - Pessimistic/Optimistic
- 2 Phase Commit
  - From RAM to disk
- Deadlock-free Transactions





# Demo

# Step 2: Full Switch to HTAP

## Step 2: Full Switch to HTAP

- Collocated Processing
- Real-Time Streaming
- Machine and Deep Learning
- Memory-Centric Database
  - Boosting Reads/Writes
  - Scaling Beyond RAM capacity
  - Instantaneous Restarts



[Webinar - Part 2: Building HTAP Apps With Ignite](#)

# Questions?

Visit <https://www.gridgain.com/resources/in-memory-computing-resources>

Support | Management Tool | Documentation | Blog | Q Search | [DOWNLOAD >](#)

**GridGain** Technology ▾ Products ▾ Solutions ▾ Customers ▾ **Resources ▾** Company ▾

## In-Memory Computing Resources from GridGain

Home > Resources > In-Memory Computing Resources from GridGain

### GridGain and Apache Ignite Information

GridGain® Systems offers a variety of in-memory computing resources which include information on the GridGain in-memory computing platform and Apache Ignite™. Use the search fields below to identify the best resources for your needs or browse our extensive library of webinars, white papers and more using the navigation to the left.

Resource Type:  Search Term:  Your Role:  [Apply](#)

### Featured In-Memory Computing Resources

#### Choosing the Right In-Memory Computing Technology for Your App

**Webinar**

The need for real-time computing has resulted in the growth of many different in-memory computing technologies including caches, in-memory data grids, in-memory databases, streaming technologies and broader in-memory computing platforms. But what are the best technologies for each type of project? Learn about your options from one of the leading in-memory computing veterans.

[Read More >](#)

Browse

Search

Click Here!

# Thank You!!!

Thank you for joining us. Follow the conversation.

<https://ignite.apache.org>



<https://www.gridgain.com>

@denismagda

#apacheignite  
#gridgain  
#inmemorycomputing

@rdmeyersf