



# Building Consistent and Highly Available Distributed Systems with **Apache Ignite™** and **GridGain**

**Valentin Kulichenko**

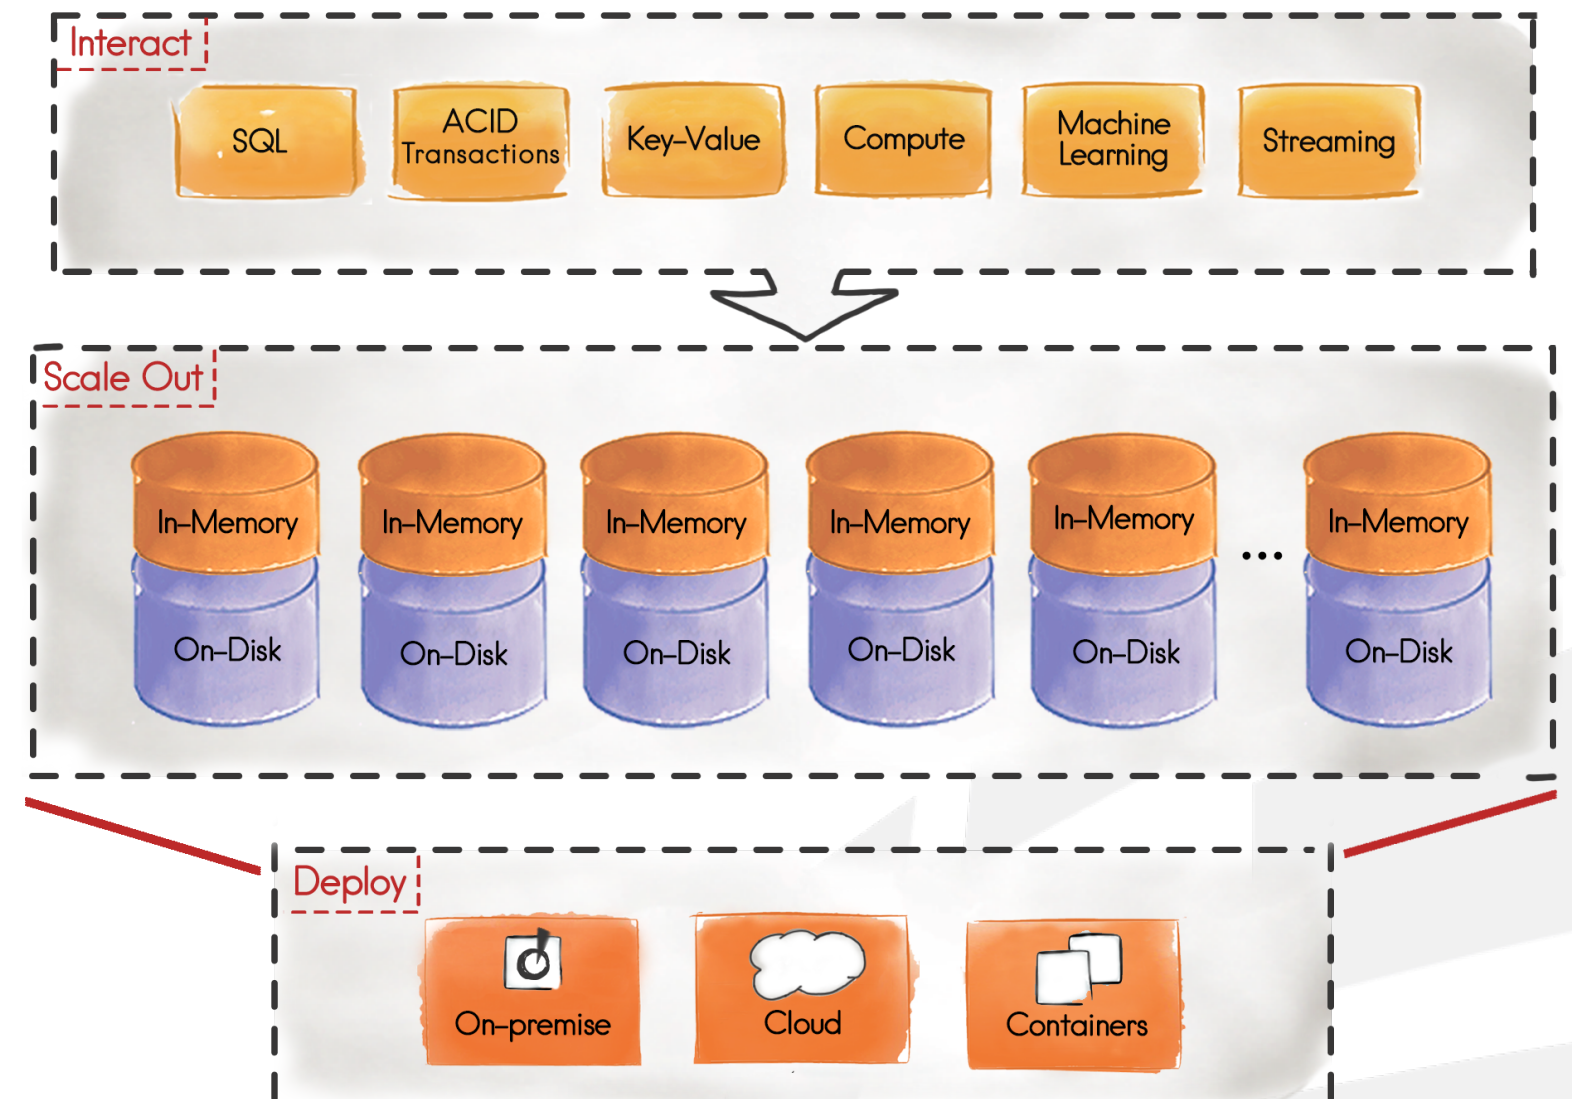
GridGain Lead Architect  
Apache Ignite PMC

# Agenda

- What is Ignite and GridGain
- Ignite and CAP Theorem
- High Availability in Ignite
  - Cache Backups
  - Persistence
  - Data Center Replication
  - Data Snapshots (+ Demo)
- Q&A

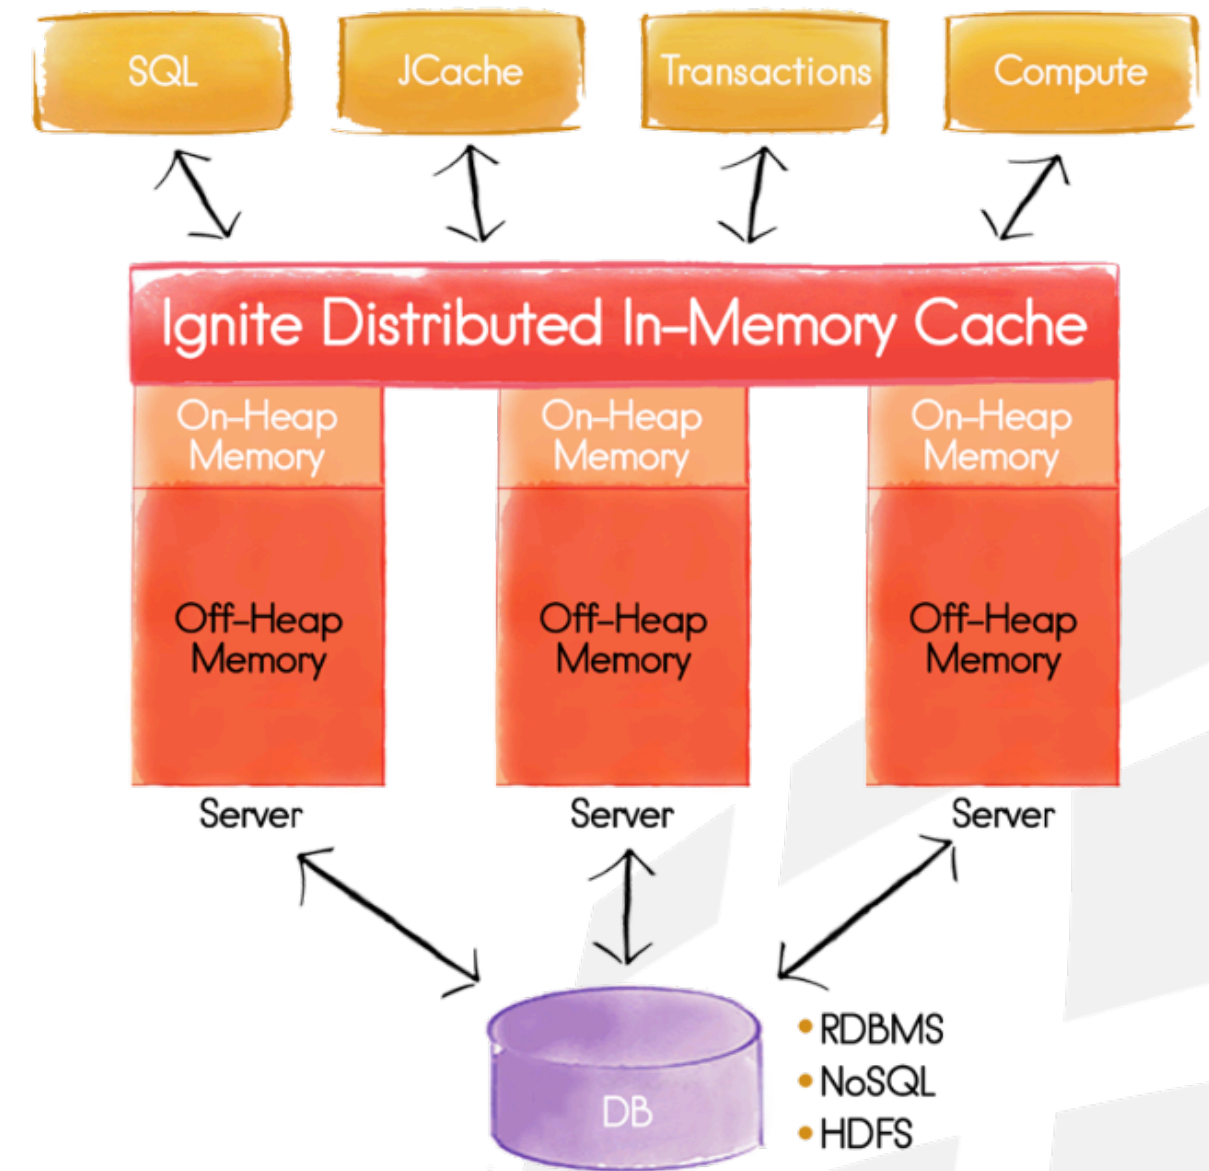
# What is Ignite?

- A **memory-centric** data platform
- that is **strongly consistent**
- and **highly available**
- with powerful **SQL**,
- **key-value** and **processing APIs**




# Key-Value Data Grid

- In-Memory Key-Value Store
- ACID Compliant
- Collocated Processing
- Persistence
  - Native Ignite persistence
  - Pluggable 3<sup>rd</sup> party persistence
- Usage
  - Database Caching
  - Web Session Clustering



# Distributed SQL Database

- Both In-Memory and On-Disk
- Full SQL Support
  - SELECT, UPDATE, DELETE, INSERT
- Data Definition Language
  - Configure schema and caches





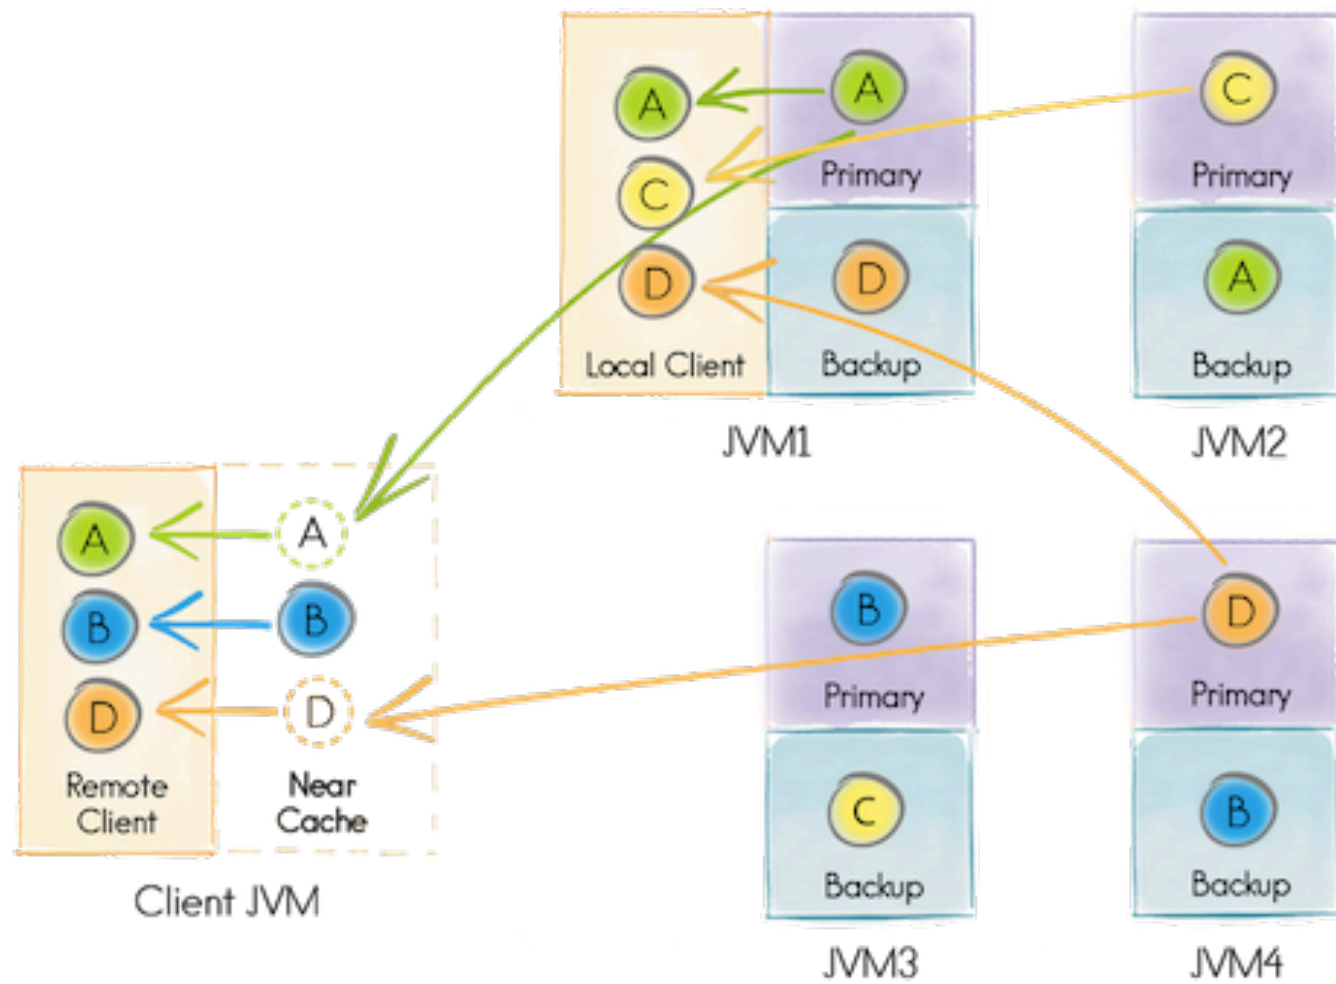
# Ignite and CAP Theorem

- CAP
  - Consistency (C)
  - Availability (A)
  - Network Partition Tolerance (P)
  - Most of distributed systems are CP or AP
  - Impossible: CAP
- Ignite
  - Strongly CP
  - But still highly available?

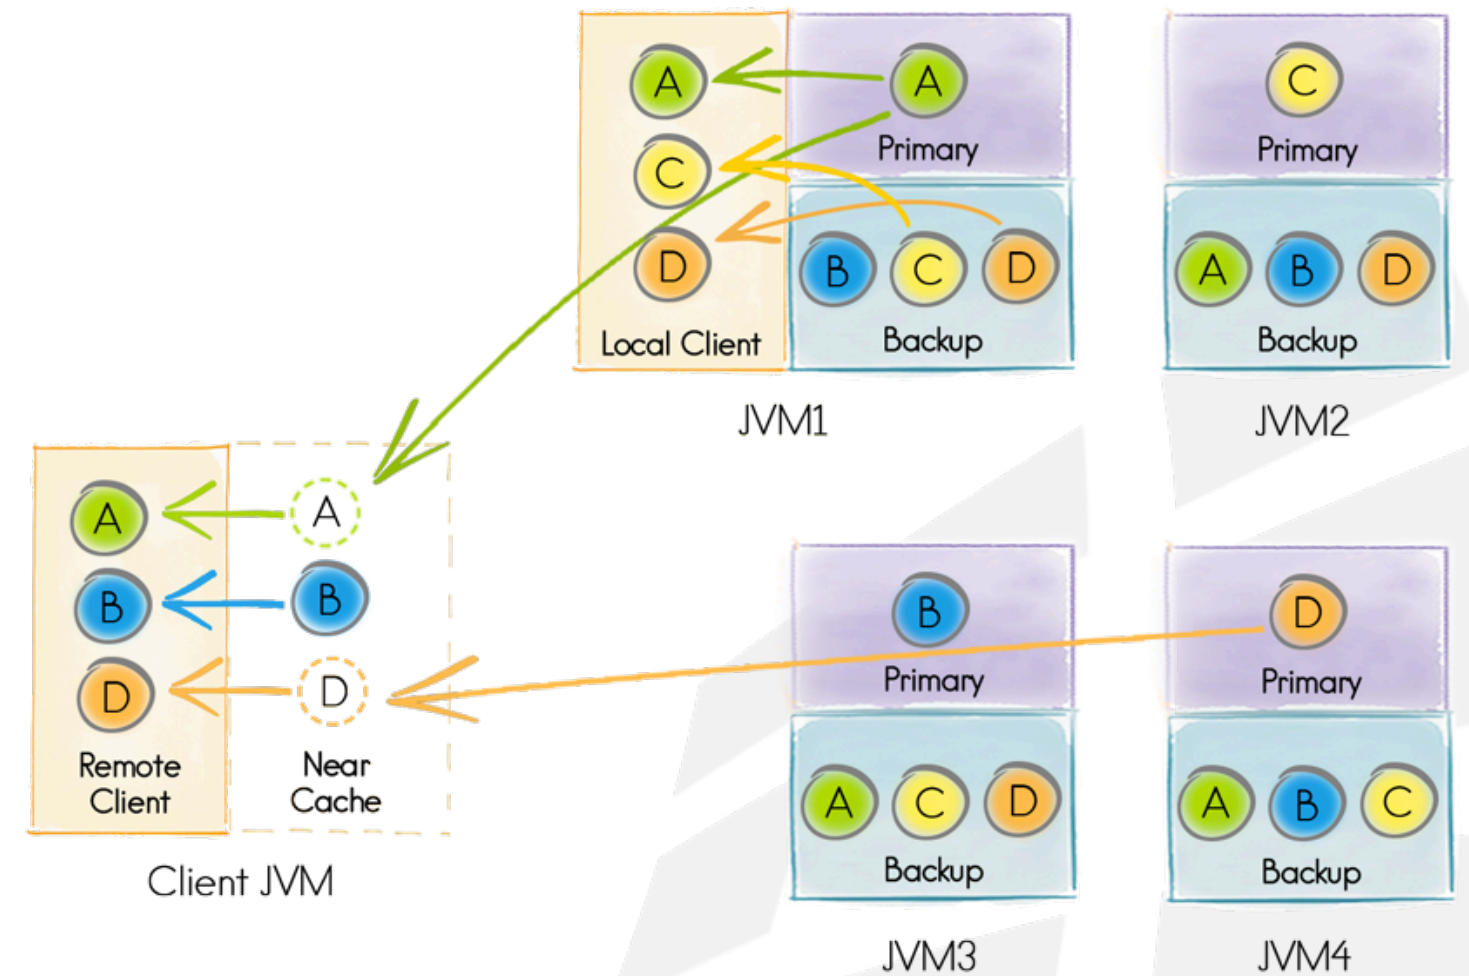
C  
A?  
P

# Cache Backups: Partition vs. Replicated

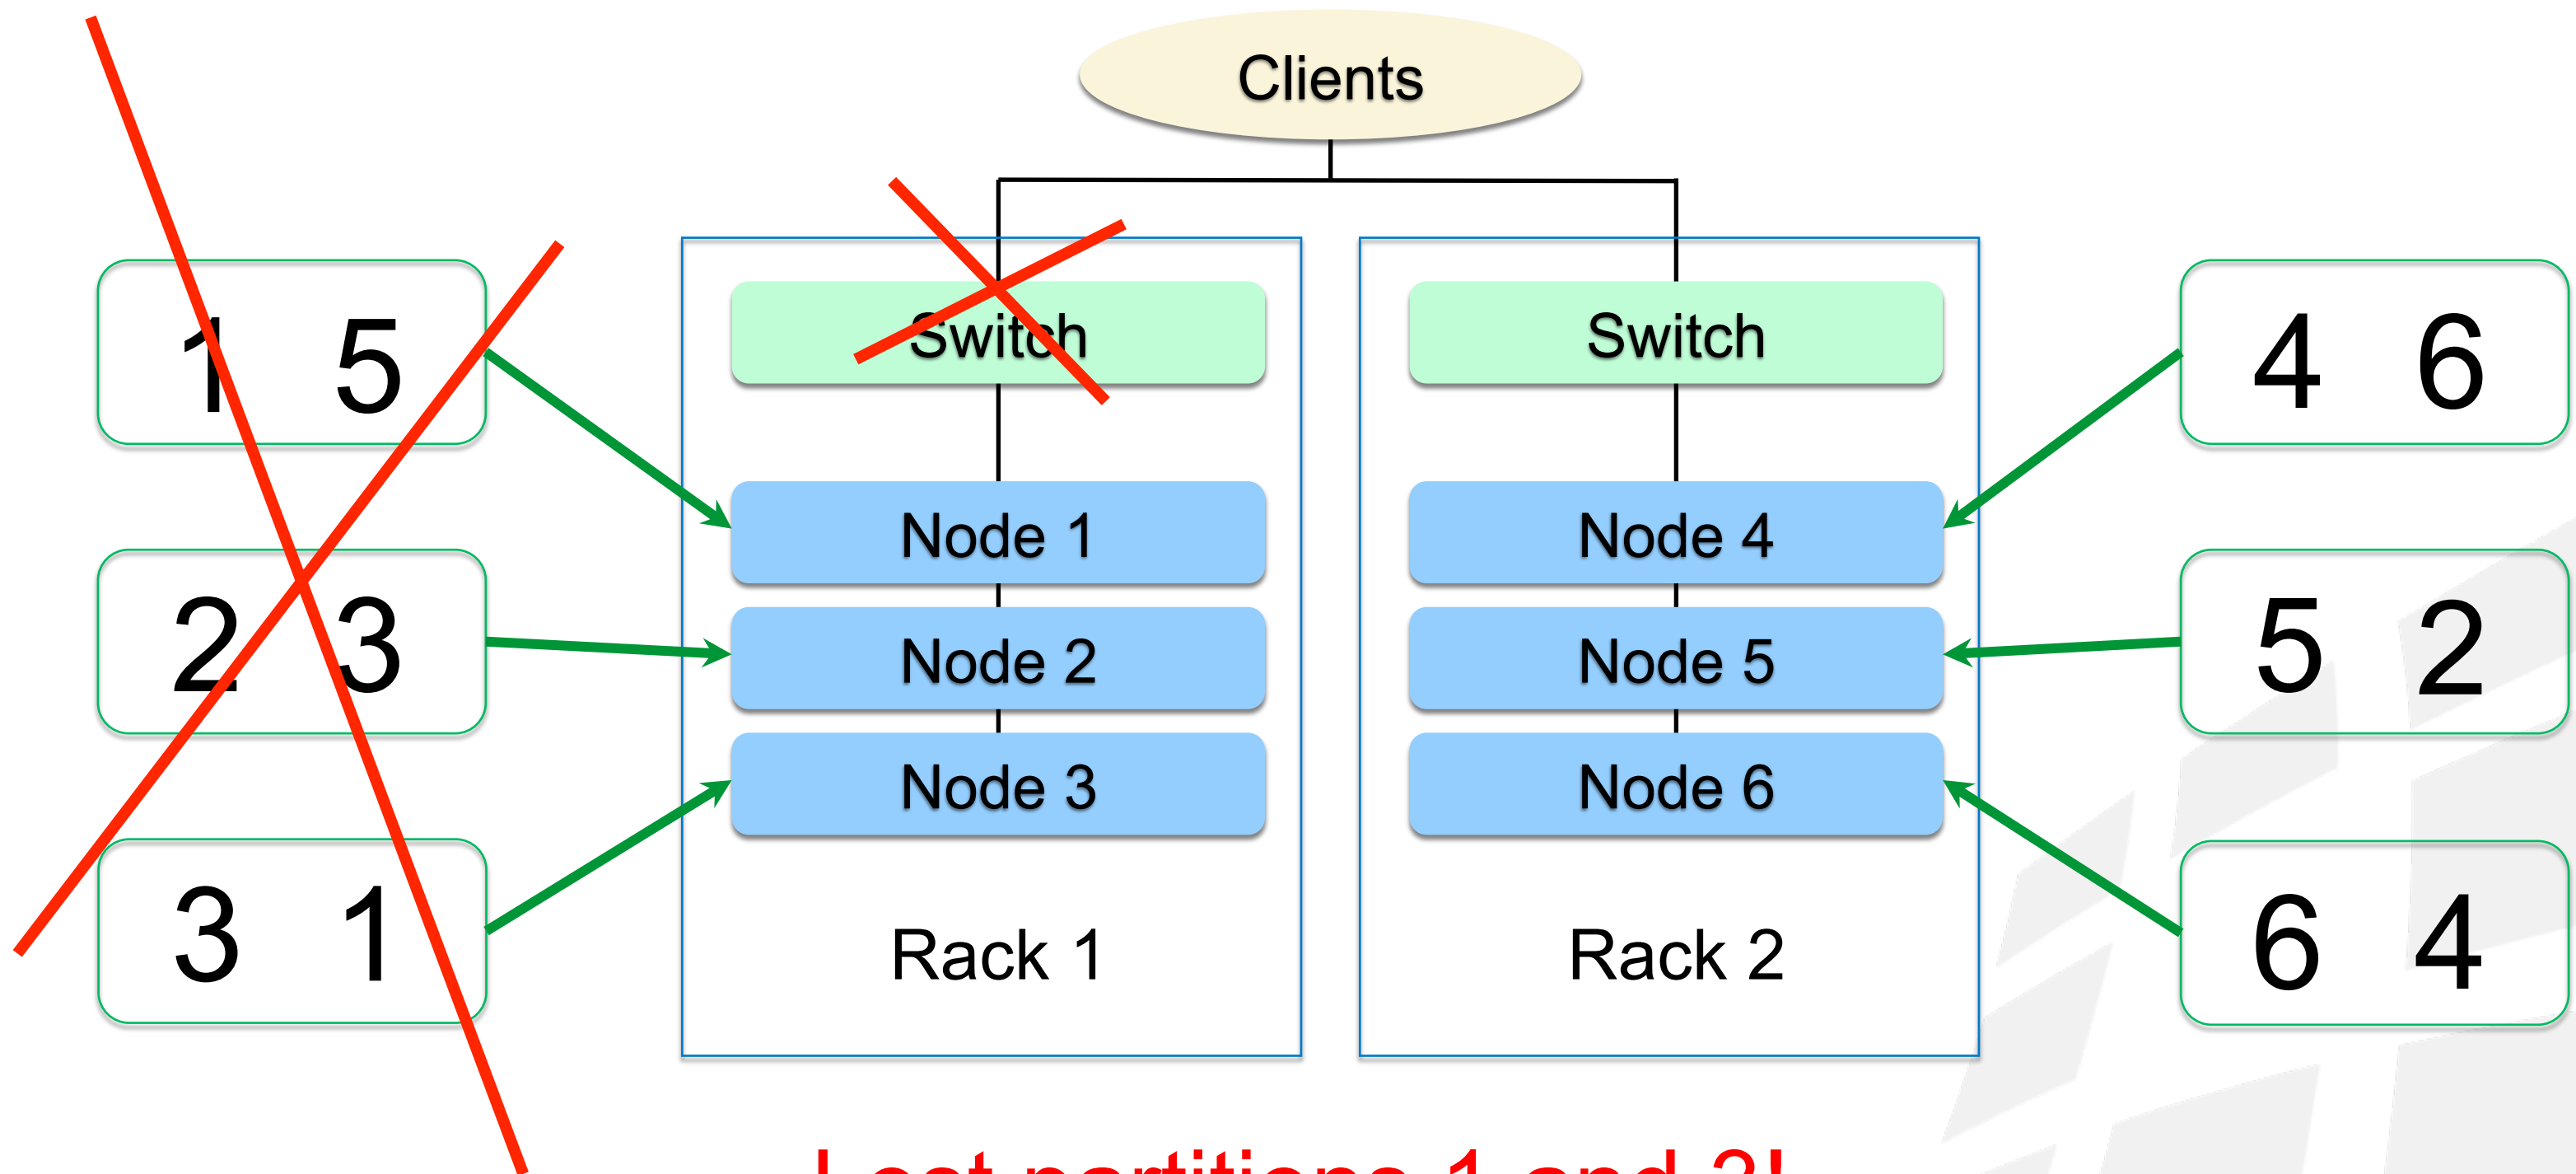
Partitioned Cache



Replicated Cache



# Cache Backups and Network Segmentation



Lost partitions 1 and 3!



# Cache Backups and Network Segmentation: Solution

- Assign attribute to each node:

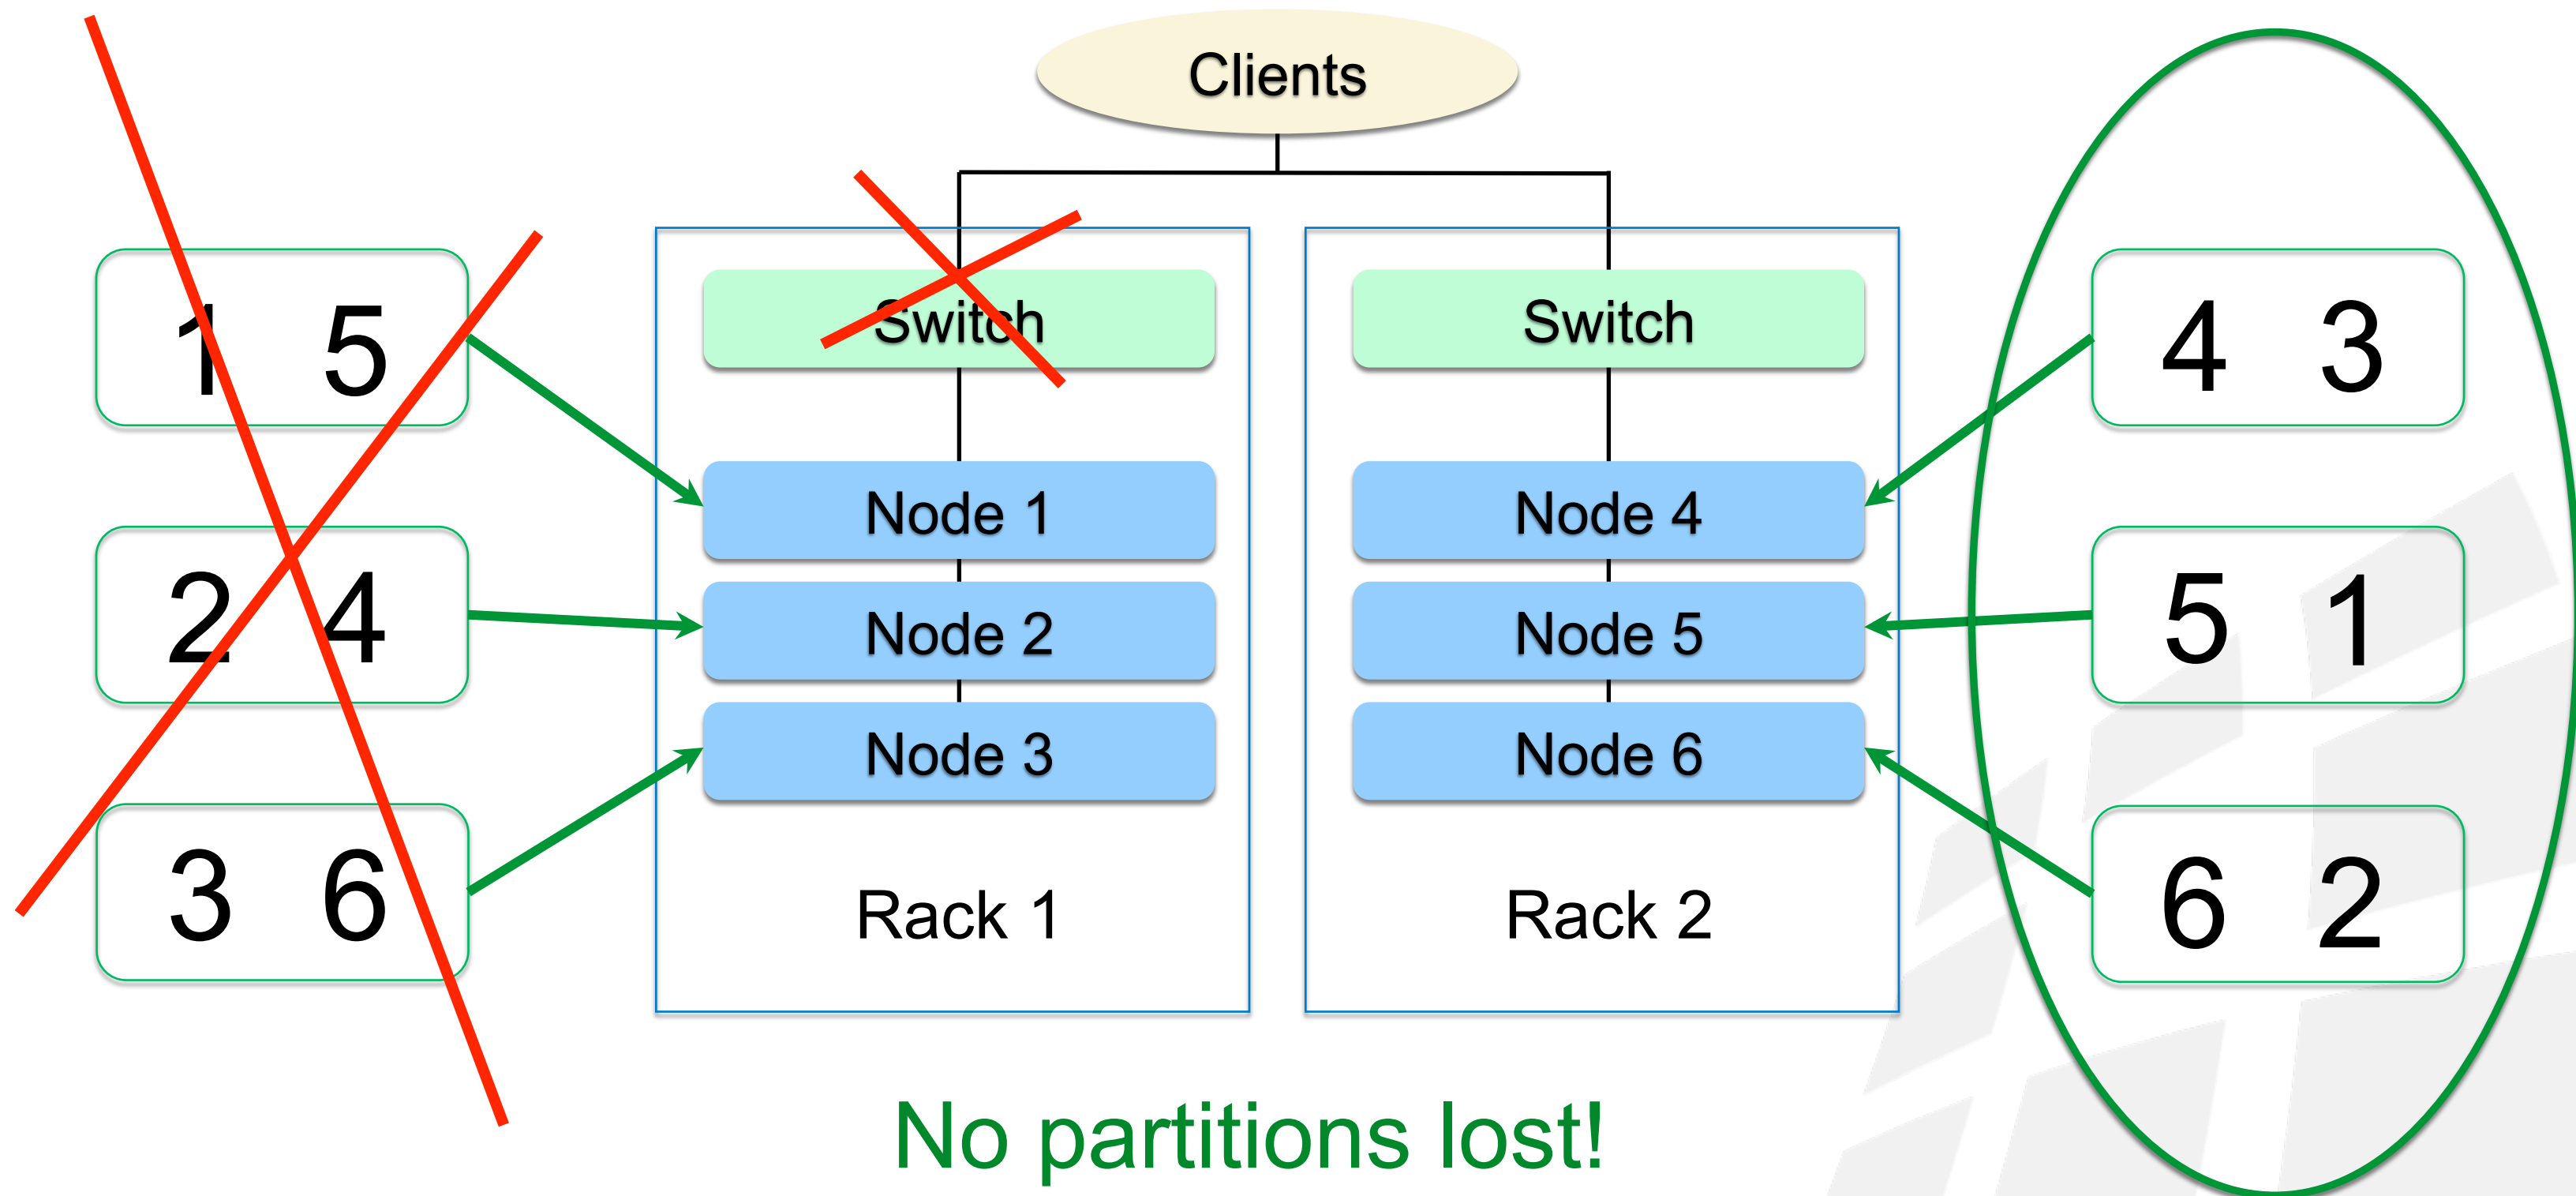
```
<property name="userAttributes">
  <map>
    <entry key="rack" value="rack-1"/>
  </map>
</property>
```

- Implement backup filter:

```
public class RackFilter implements IgniteBiPredicate<ClusterNode, List<ClusterNode>> {
    @Override public boolean apply(ClusterNode candidate, List<ClusterNode> assigned) {
        String candidateRack = candidate.attribute("rack");
        String primaryRack = assigned.get(0).attribute("rack");

        return !Objects.equals(candidateRack, primaryRack);
    }
}
```

# Cache Backups and Network Segmentation: Solved!



# Persistence

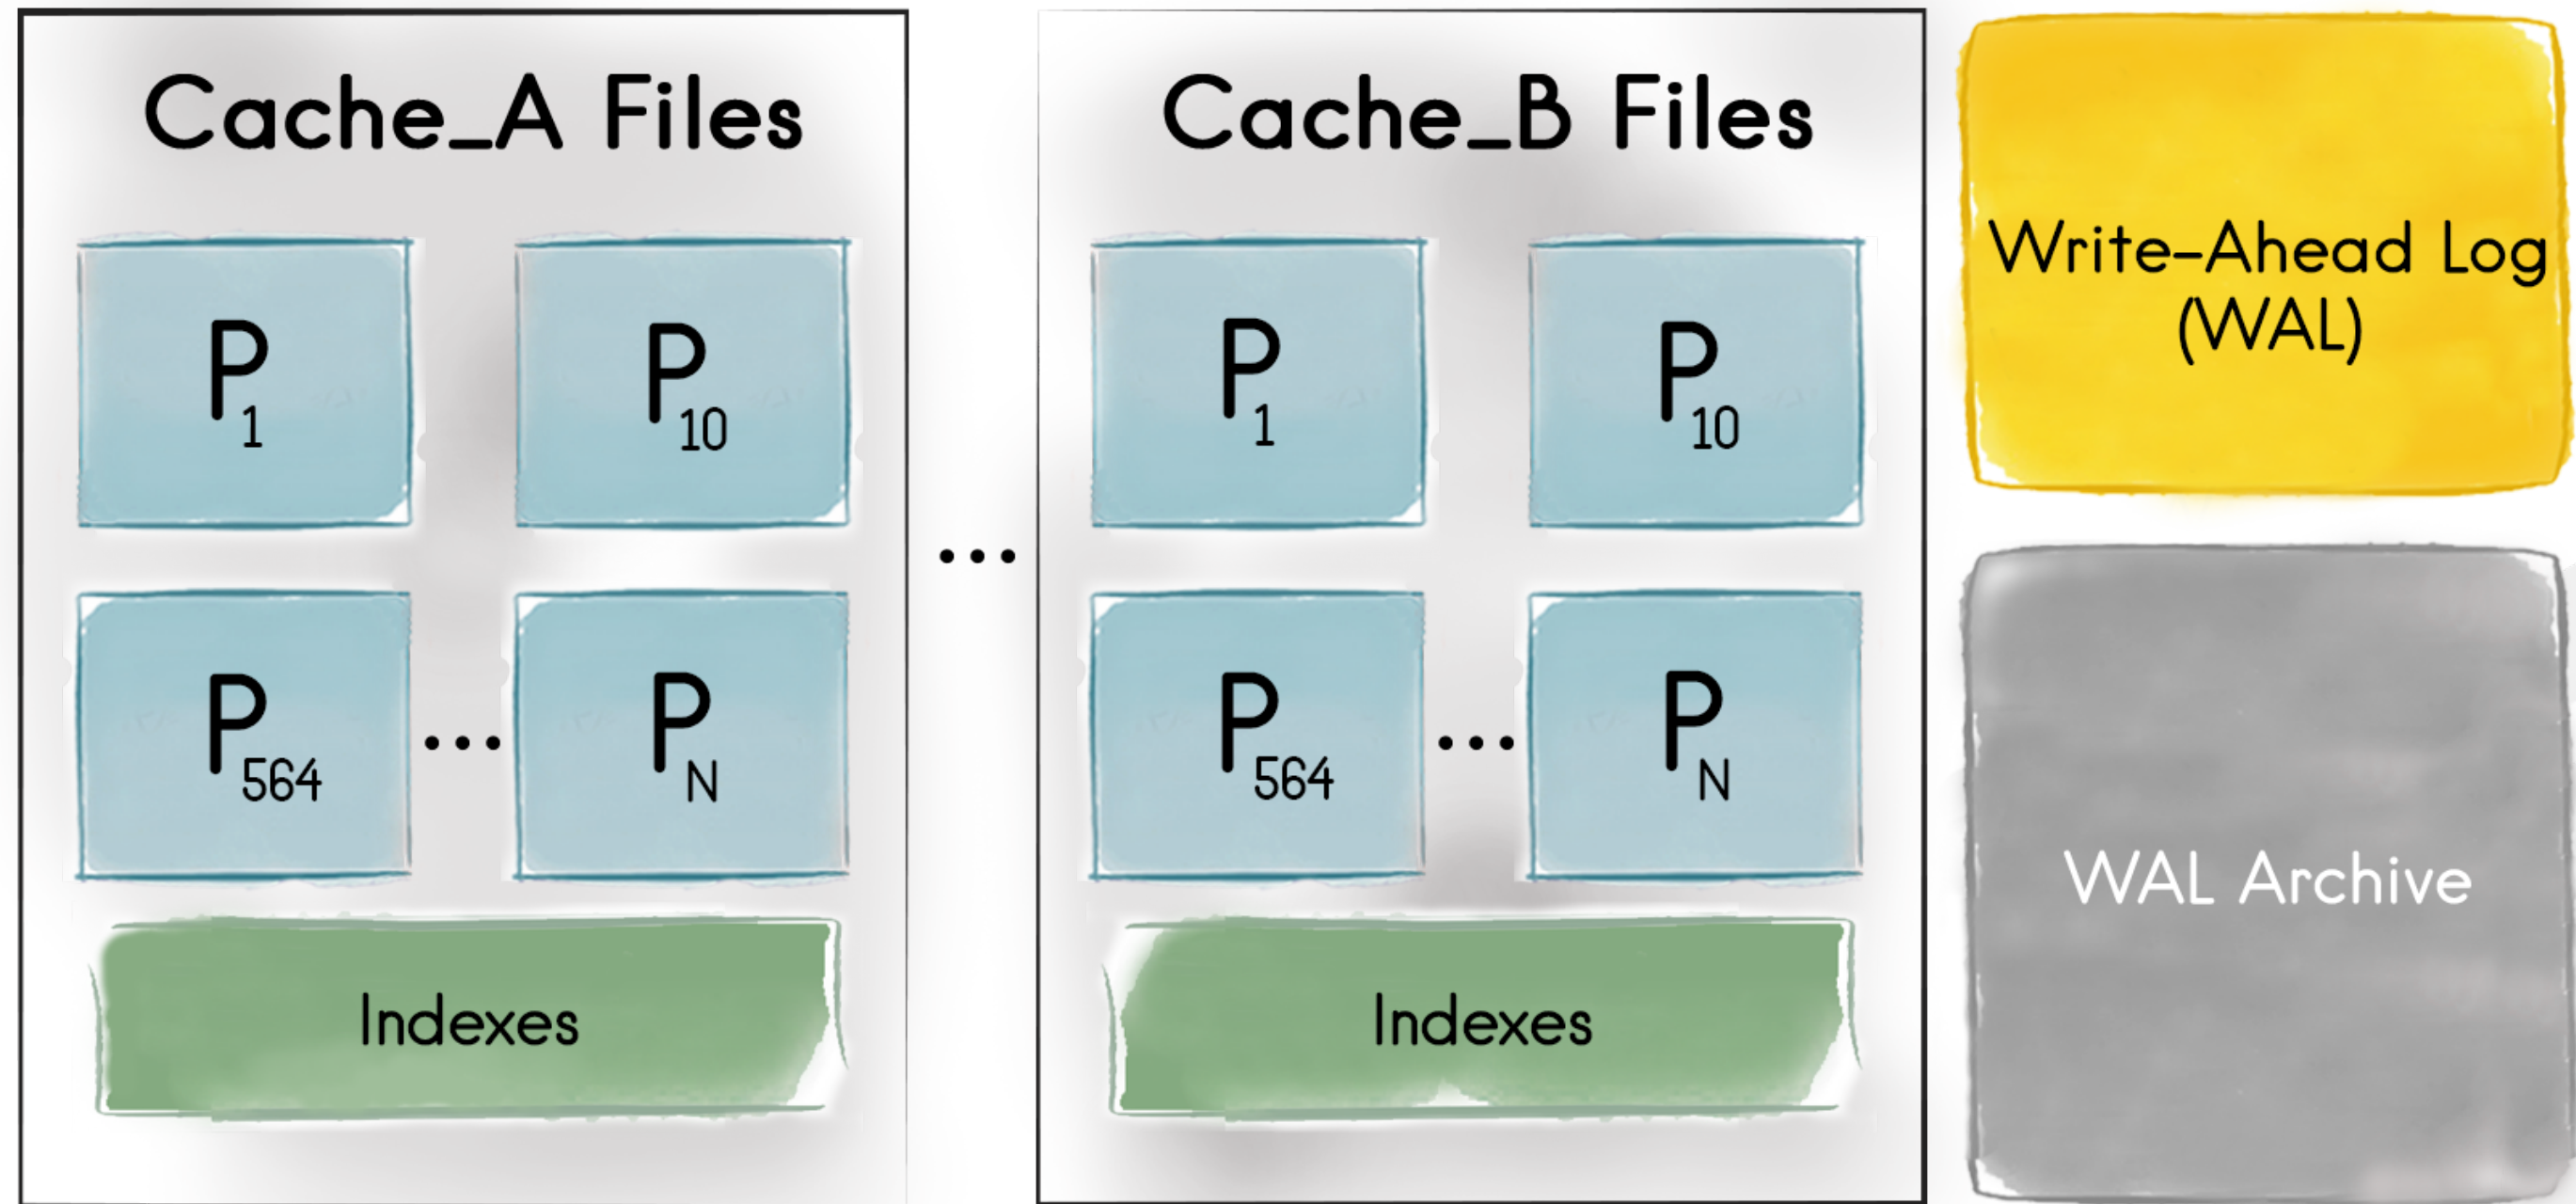
## Ignite Persistent Store

- Additional Disk Tier
- SQL & key-value over both memory and disk
- Memory warmup never required
- Distributed disk storage
- Reads/writes are distributed

## 3<sup>rd</sup> Party Store

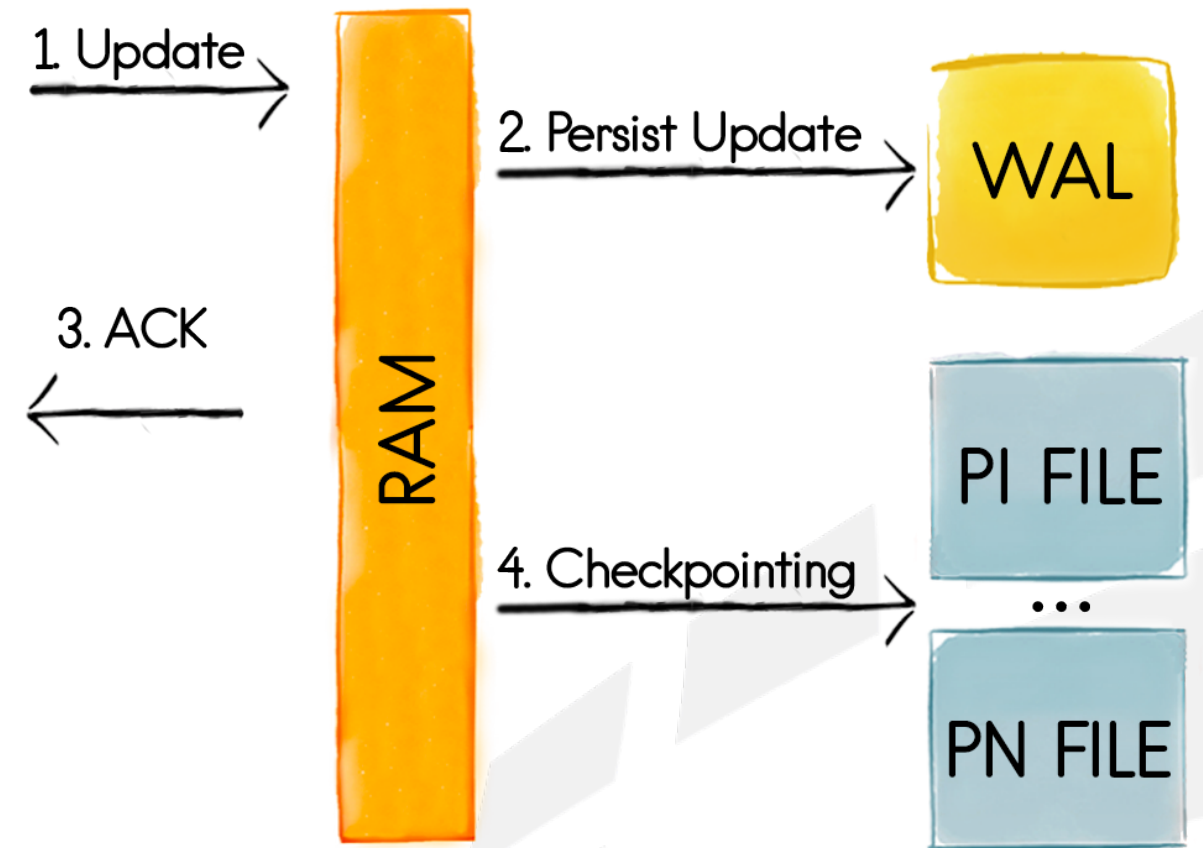
- RDBMS, MongoDB, etc.
- Key-value over both memory and disk
- Memory warmup might be required (SQL usage)
- Single point of failure (RDBMS)
- Bottleneck for writes (RDBMS)

# Ignite Persistence Store: Per-Node Architecture



# Ignite Persistence Store: Consistency and Durability


- Write-Ahead Log (WAL)
  - Append only file
  - Propagates updates to disk
  - Provides recovery mechanism
- Checkpointing
  - Triggered periodically
  - Copies actual data from RAM to disk
  - Reduces WAL size





# Data Center Replication

- Multiple (up to 32) Data Centers
- Active-Active & Active-Passive
- Smart Conflict Resolution
- Durable Persistent Queues
- Automatic Throttling



*GridGain Enterprise Edition Only!*

# Data Snapshots

- Full snapshots
  - Full state
  - May take long time
- Incremental snapshots
  - Partial state
  - Only delta since last full snapshot
- Scheduled snapshots
- Restore on different clusters
- Fully managed

The screenshot displays the GridGain web interface for managing snapshots. The top navigation bar includes 'Configure', 'Queries', 'Monitoring', 'Start Demo', 'Cluster 6CD0FF49', and 'Alex Smith'. The 'Snapshots' section has buttons for 'Set Schedule' and 'Create Incremental Snapshot'. Below this is a table for 'My schedule' with columns: Name, Command, Schedule, Type, Last Run, Next Run, and Enable. The table lists three items: 'Snapshot' (create), 'Cleanup' (delete), and 'Move to archive' (move), all with 'Enable' toggle switches turned on. Below the schedule table is a section for 'All' snapshots, showing '1 item selected' and a period from '19/06/2017 00:00 to 20/06/2017 00:00'. The table has columns: ID, Start Time, Duration, Mode, Type, and Caches. It lists four snapshots, with the first one selected (checkbox checked).

Name	Command	Schedule	Type	Last Run	Next Run	Enable
Snapshot	create					<input checked="" type="checkbox"/>
Cleanup	delete					<input checked="" type="checkbox"/>
Move to archive	move					<input checked="" type="checkbox"/>

ID	Start Time	Duration	Mode	Type	Caches
<input checked="" type="checkbox"/> 1 497 863 437 767 INC 2	19/06/2017 16:10		console	FULL	5
<input type="checkbox"/> 1 497 863 443 613	19/06/2017 16:10		wizard	INC	5
<input type="checkbox"/> 1 497 863 448 438	19/06/2017 16:10		console	INC	5
<input type="checkbox"/> 1 497 863 768 026 INC 1	19/06/2017 16:16		console	FULL	1

*GridGain Ultimate Edition Only!*



# ANY QUESTIONS?

Thank you for joining us. Follow the conversation.

<http://ignite.apache.org>



#apacheignite