

Apache Ignite 101: Key Deployment Strategies for Database Acceleration

Valentin Kulichenko

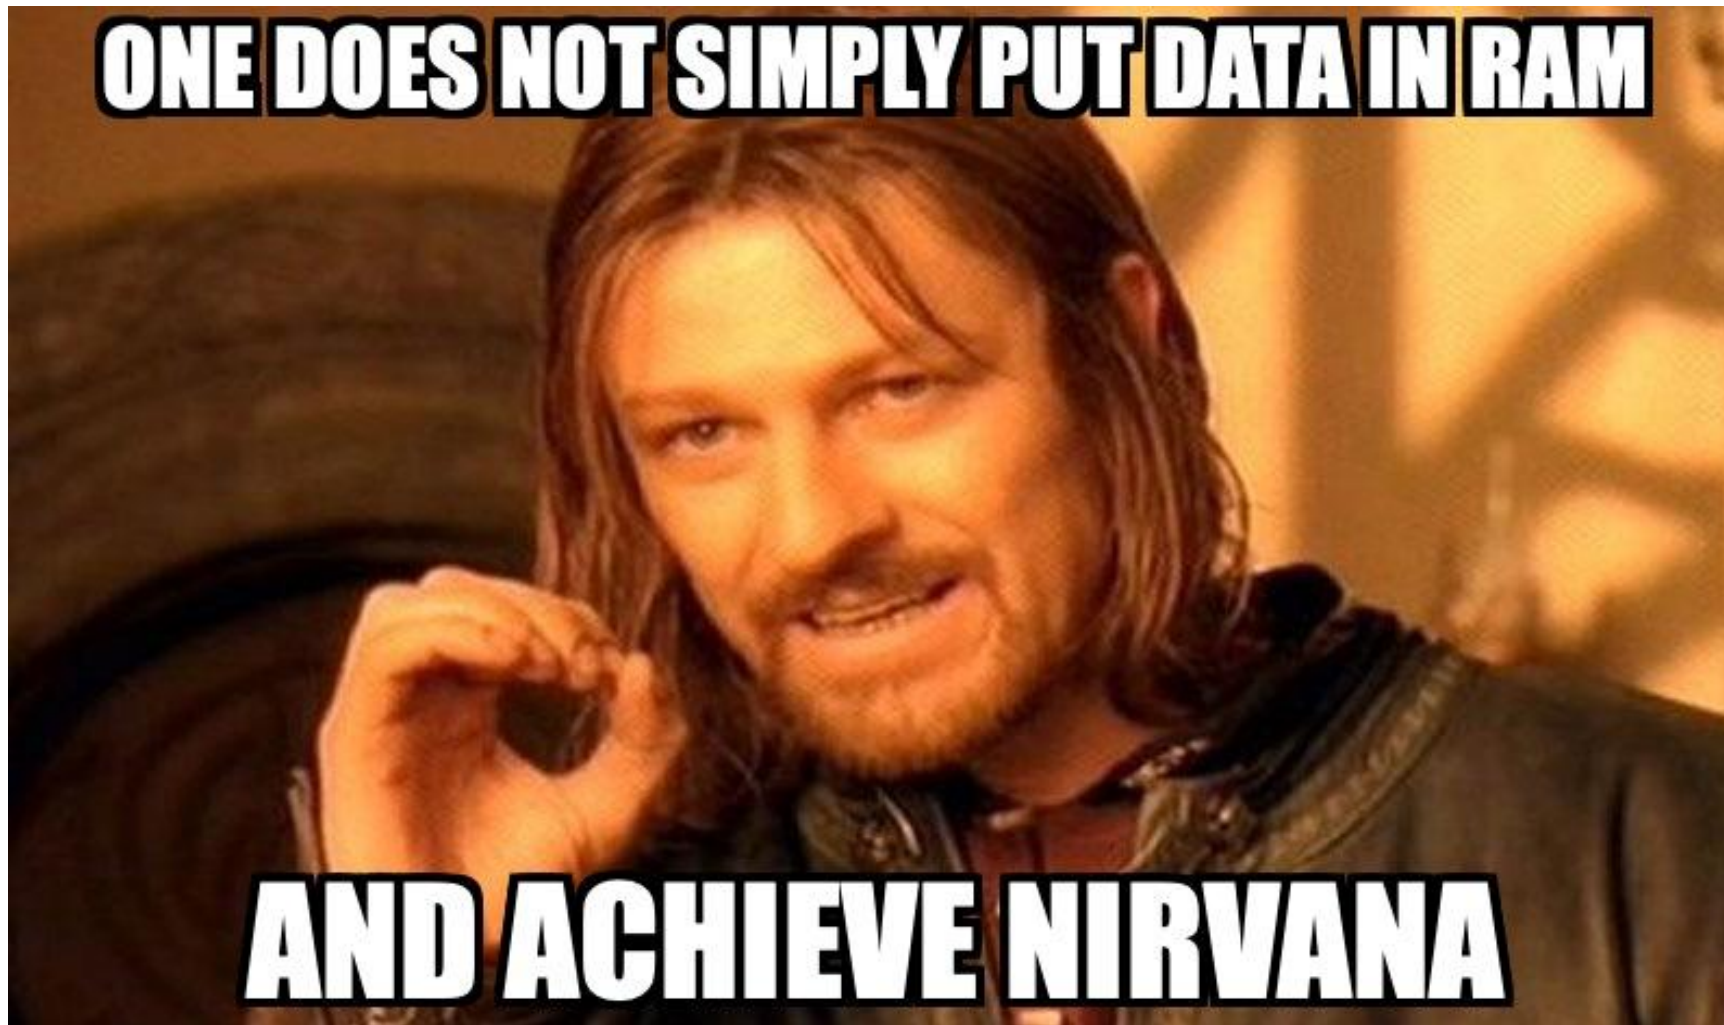
Aug 5, 2020

Memory is Much ... Much Faster Than Disk

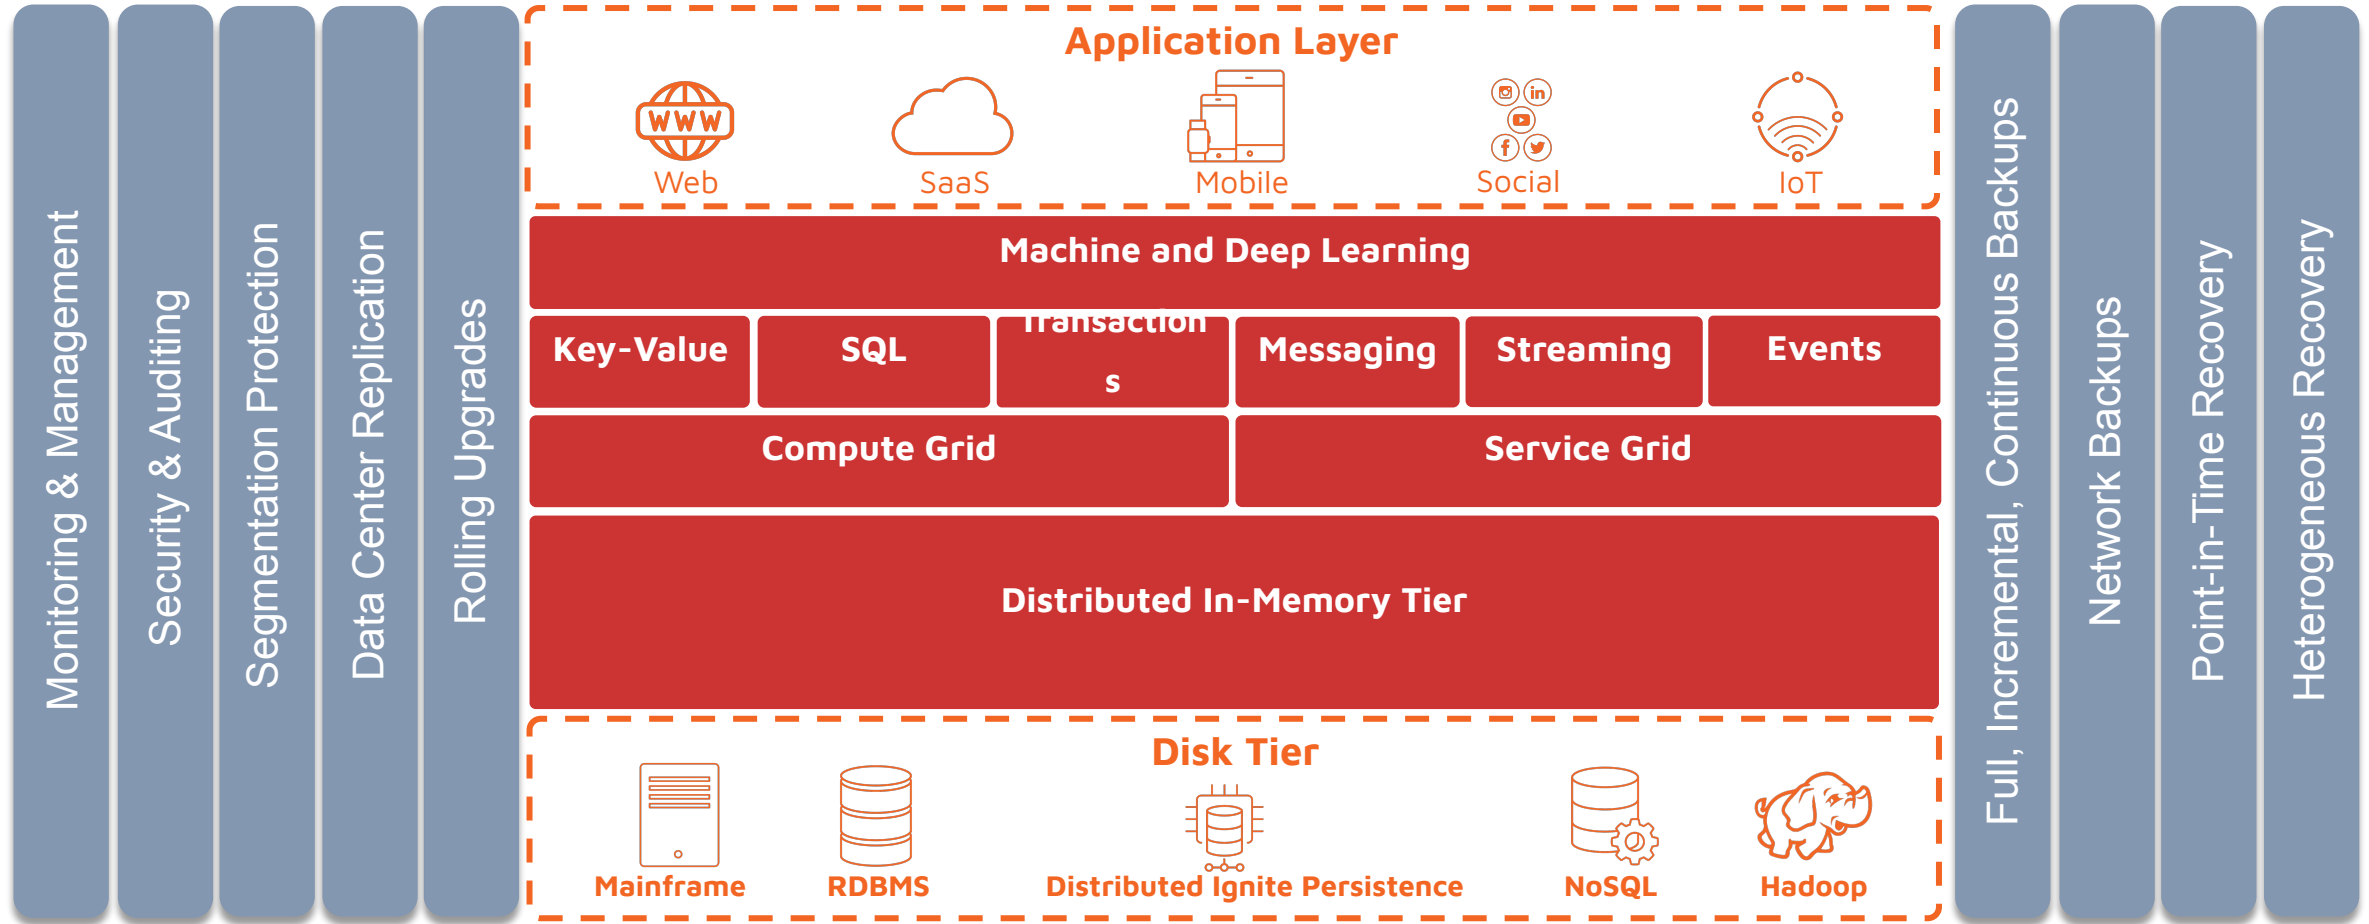


System Event	Actual Latency	Computer Latency at a Human Scale
One CPU cycle	0.4 ns	1 s
Level 1 cache access	0.9 ns	2 s
Level 2 cache access	2.8 ns	7 s
Level 3 cache access	28 ns	1 min
Main memory access (DDR DIMM)	~100 ns	4 min
Intel Optane DC persistent memory access	~350 ns	15 min
Intel Optane DC SSD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

Dirty Secret of In-Memory Systems



Apache Ignite In-Memory Computing Platform



Apache Ignite Features

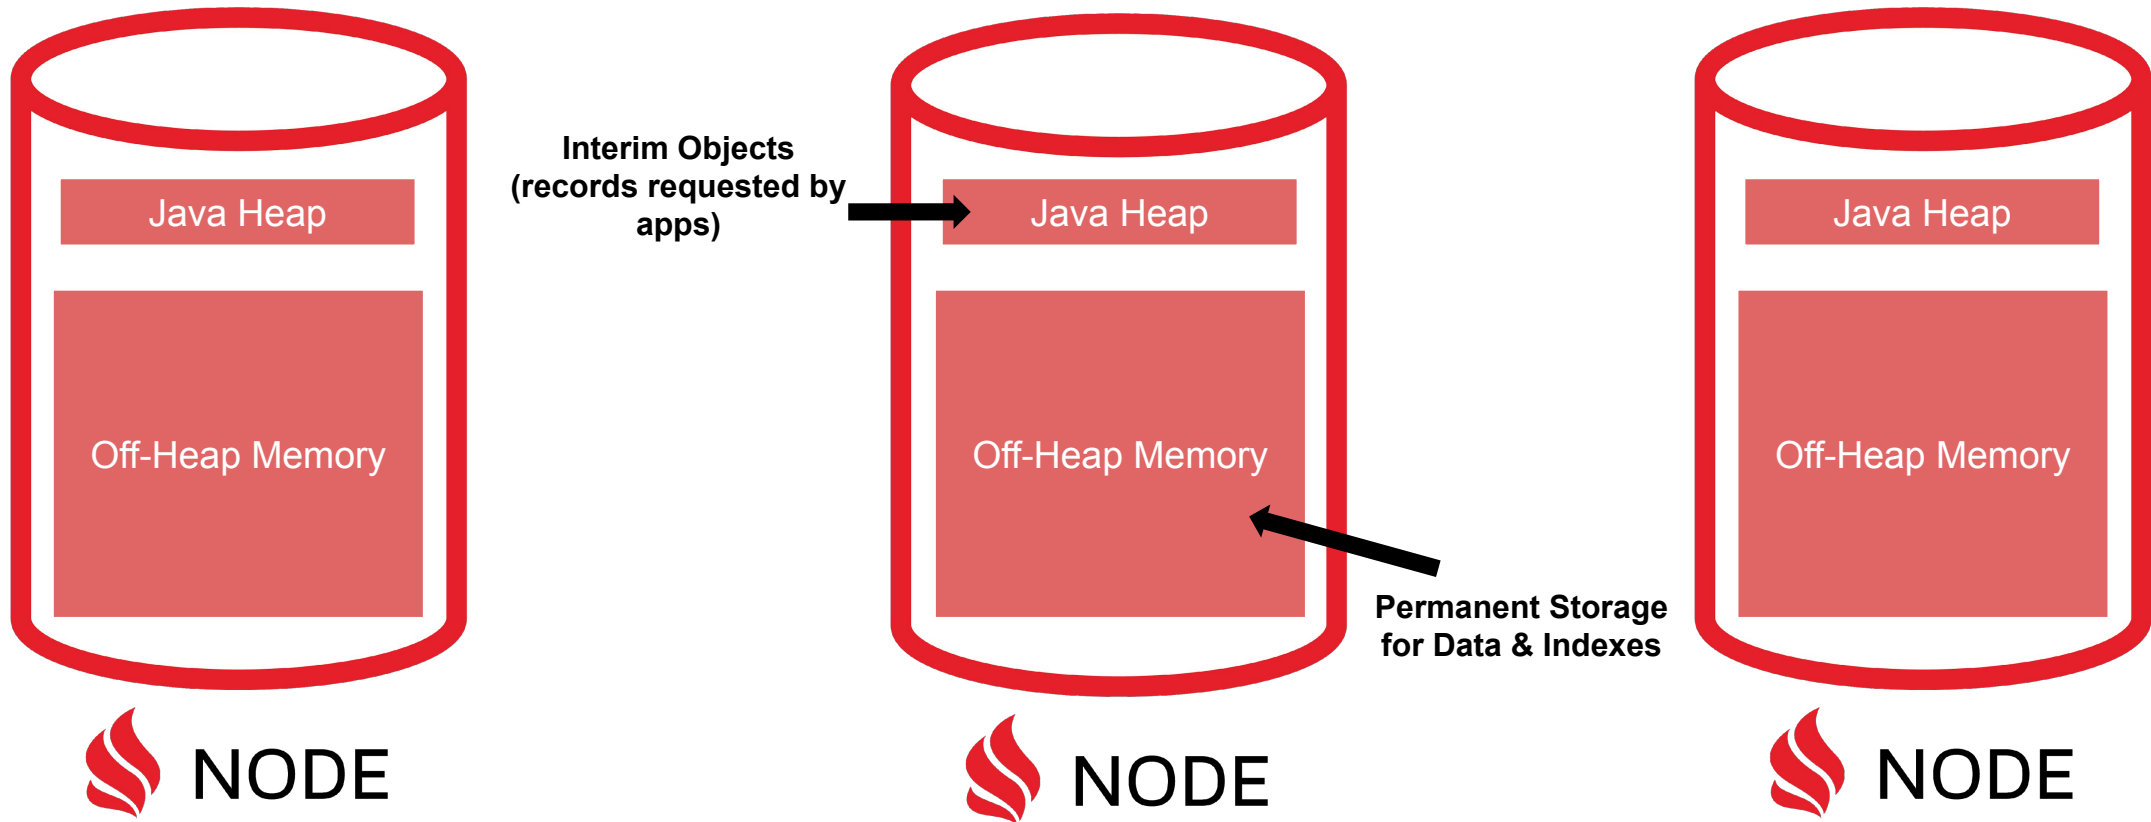
GridGain Enterprise Features

Multi-Tier Architecture Advantages



Mode	Description	Major Advantage
In-Memory	100% data in the In-Memory Store (only)	Maximum performance possible (data is never written to disk)
In-Memory + 3rd Party DB	Data in the In-Memory Data Store as a caching layer (aka. in-memory data grid) 3 rd Party DB (RDBMS, NoSQL, etc) used for persistence	Horizontal scalability Faster reads and writes
In-Memory + Persistent Store	The whole data set is stored both in memory and on disk	Survives cluster failures
100% on Disk + In-Memory Cache	100% of data is in GridGain Persistent Store and a subset is in memory	Unlimited data scale beyond RAM capacity

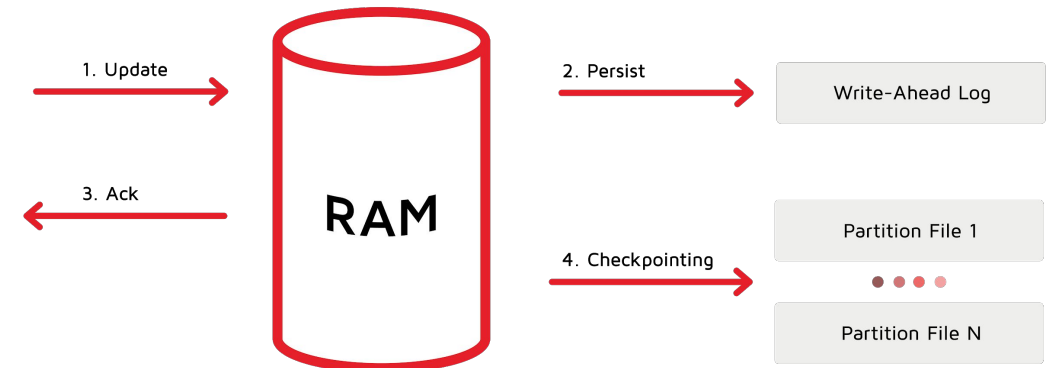
Ignite Memory Tier



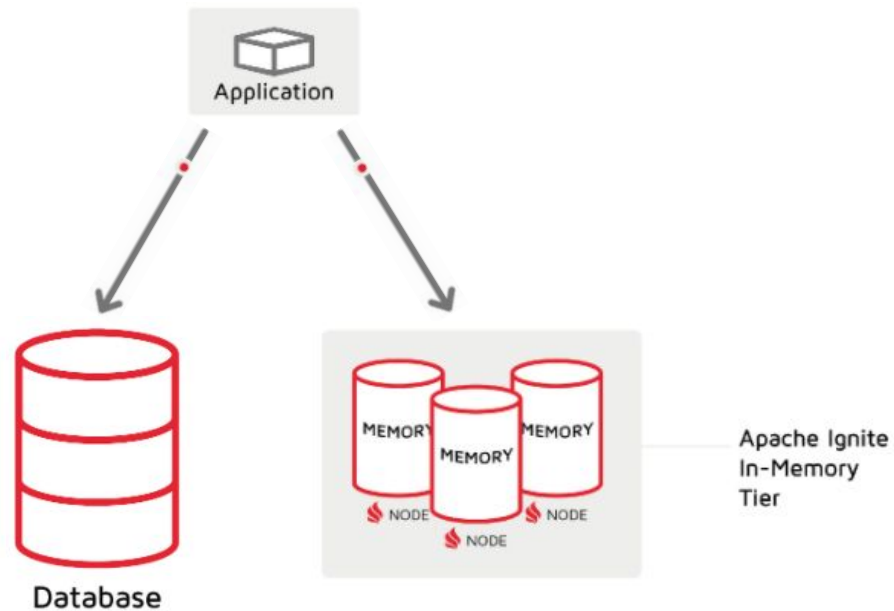
Ignite Native Persistence



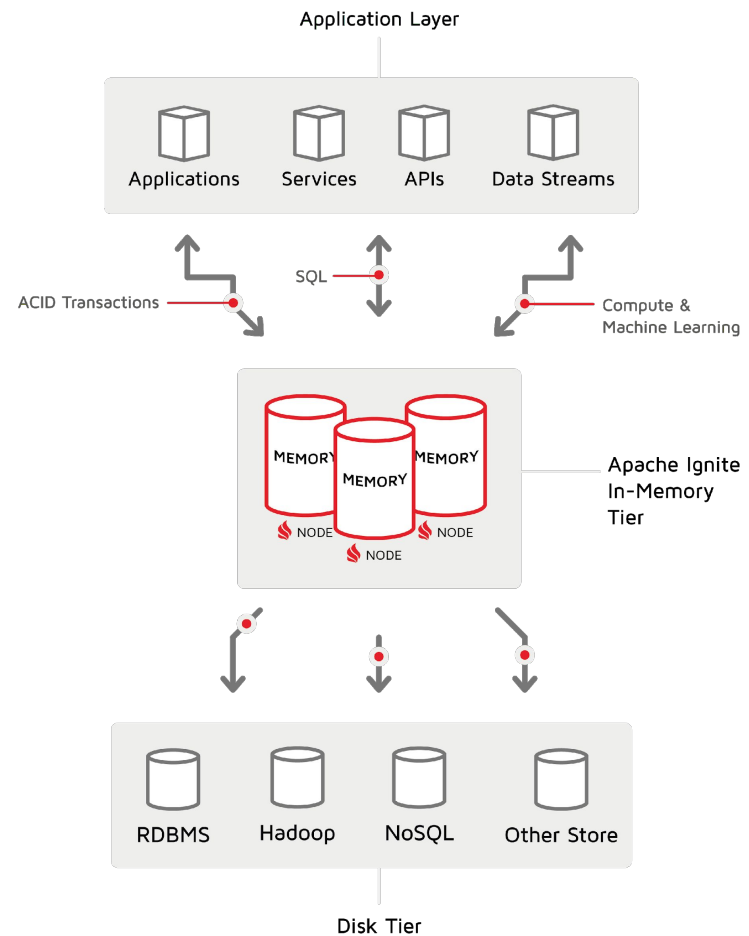
- Distributed Persistence Tier
 - Fully transactional and consistent
 - No need to cache 100% of data in RAM
 - No need to warm-up RAM on restarts
- Performance vs. Cost Tradeoff
 - Cache more for fastest performance
 - Cache less to reduce infrastructure costs



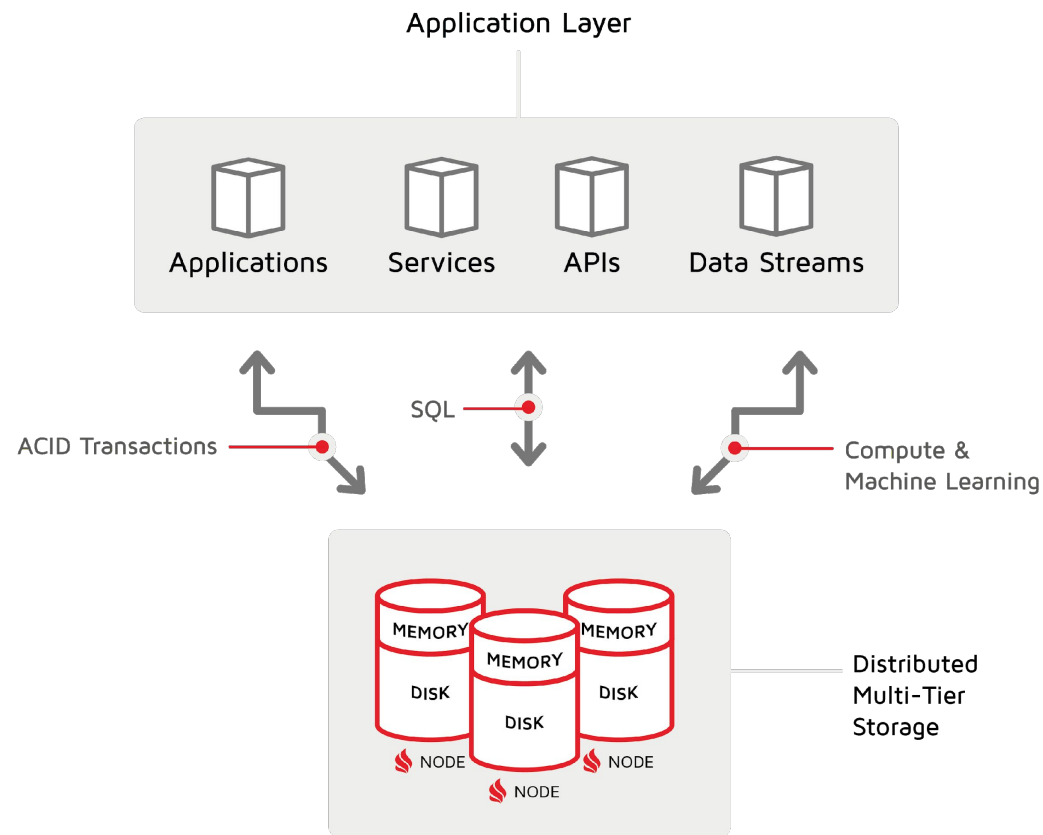
Apache Ignite as a Cache



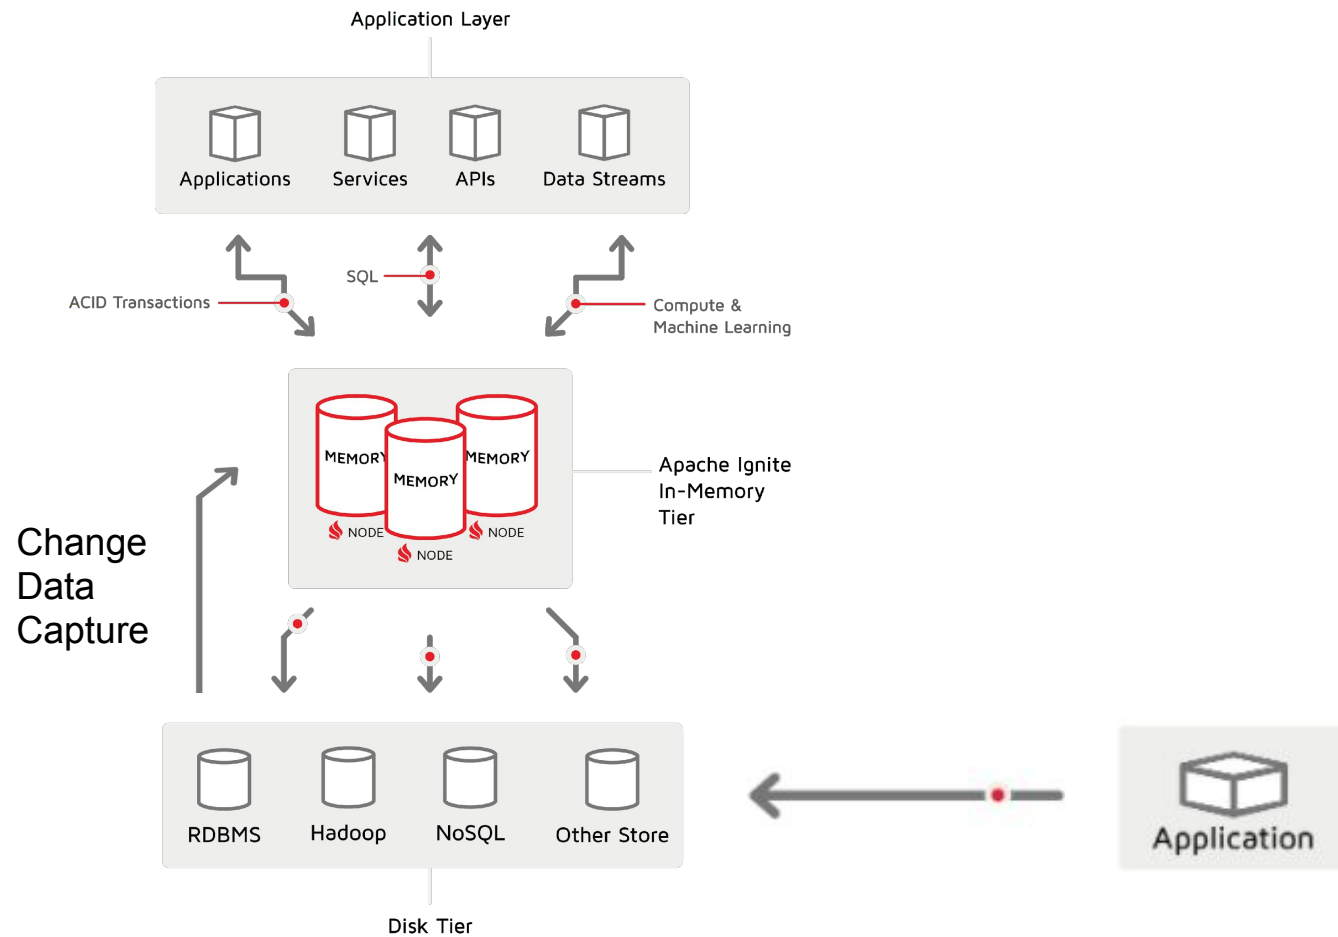
Apache Ignite as a Data Grid



Apache Ignite as a Database



Change Data Capture

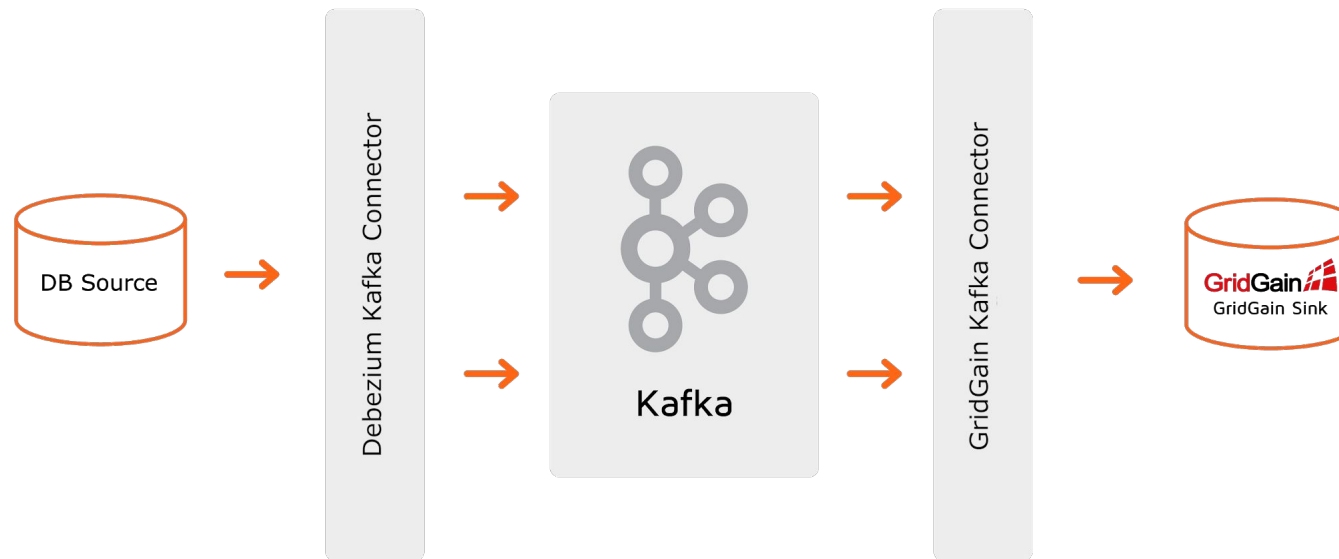


Change Data Capture Options



- Database triggers
- Oracle GoldenGate
- Streaming (e.g. Debezium+Kafka)

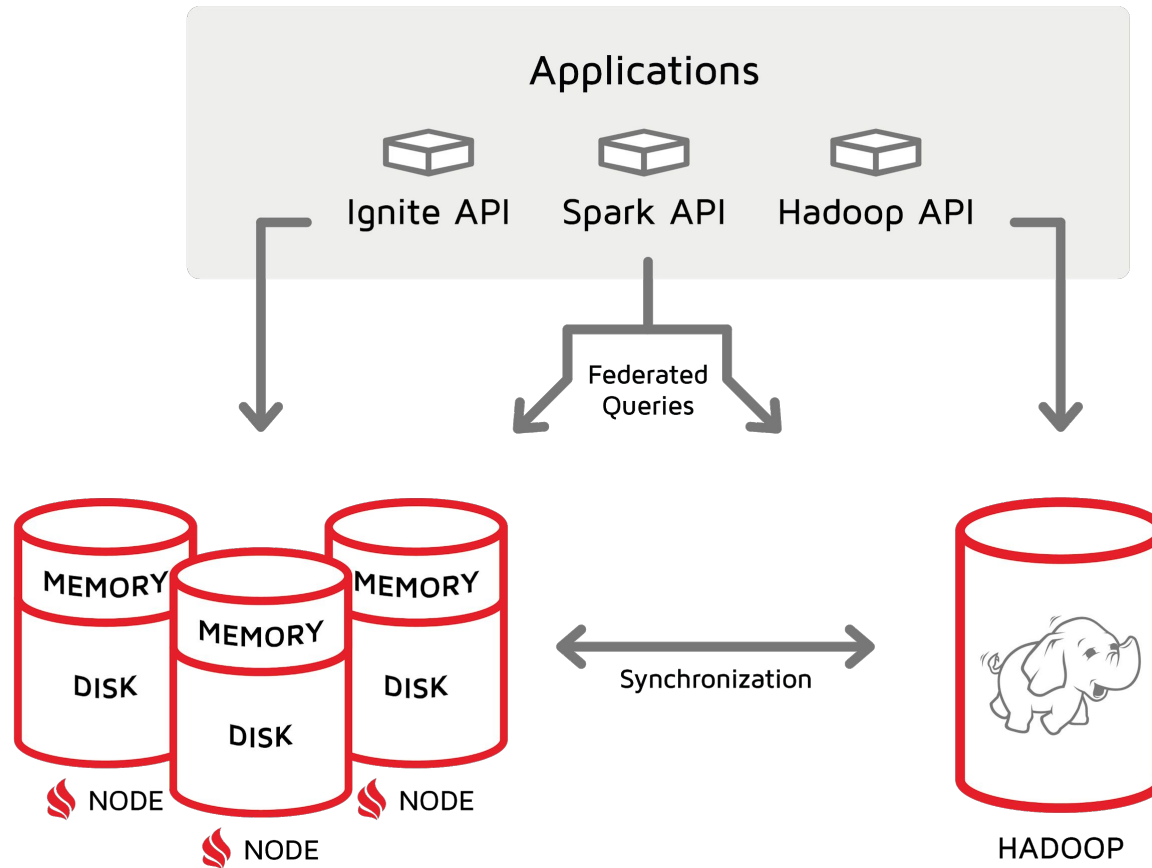
CDC with Debezium and Kafka



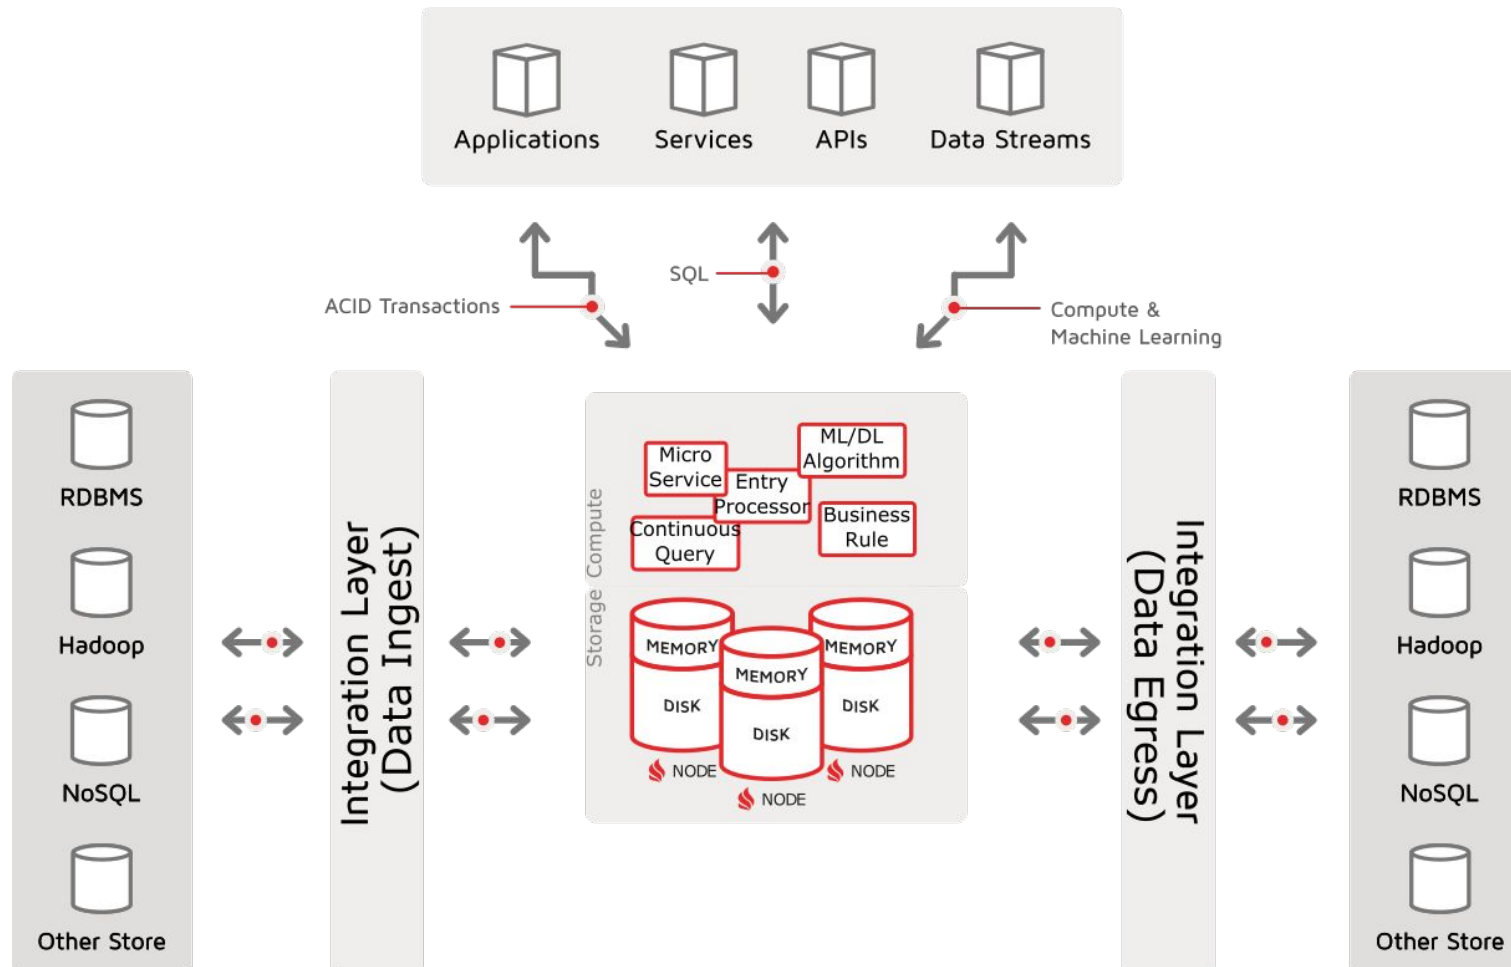
Blog and Demo by Evgenii Zhuravlev:

<https://www.gridgain.com/resources/blog/change-data-capture-between-mysql-and-gridgain-debezium>

Hadoop Acceleration



Digital Integration Hub



Stay connected with Apache Ignite users & experts



meetup.com/Apache-Ignite-Virtual-Meetup/

Join Apache Ignite Community



Discuss <https://ignite.apache.org/community/resources.html#mail-lists>

Check demos <https://github.com/GridGain-Demos/>

Connect <https://twitter.com/apacheignite>

Join events <https://ignite.apache.org/events.html>

Contribute <https://ignite.apache.org/community/contribute.html>

Questions?

