SUPER POWER APACHE® CASSANDRA™ FOR EXTREME OLTP WORKLOADS

Rachel Pedreschi

Lead Solutions Architect- GridGrain Systems

@rachelpedreschi

Igor Rudyak

Senior Solutions Architect- EPAM

GridGain Enterprise and Open Source Strategy

- GridGain Enterprise is based on Apache Ignite
- Open source is intended to provide an easy entry point for learning, testing and non-critical use
- Enterprise customers benefit from many exclusive enterprise-class features along with support and indemnification

EPAM

ENGINEERING DNA

Our top clients include:

Expedia · Google · UBS

18,000+

Engineers, designers and consultants

CONSTANT GROWTH

4 Continents

25 Countries

Q2 2016 REVENUE

$283.8M

2016 REVENUE GUIDANCE

$1.15B

INDUSTRY FOCUS

26% Financial Services

22% Travel & Consumer

20% Software & Hi-tech

7% Emerging

15% Media & Entertainment

10% Life sciences & Healthcare

SERVICES

Software Engineering & Product/ Platform Development

QA and Test Automation

Managed Services

Infrastructure & Licensing

“These aren’t the unicorns you were looking for.”

Slow Reads?

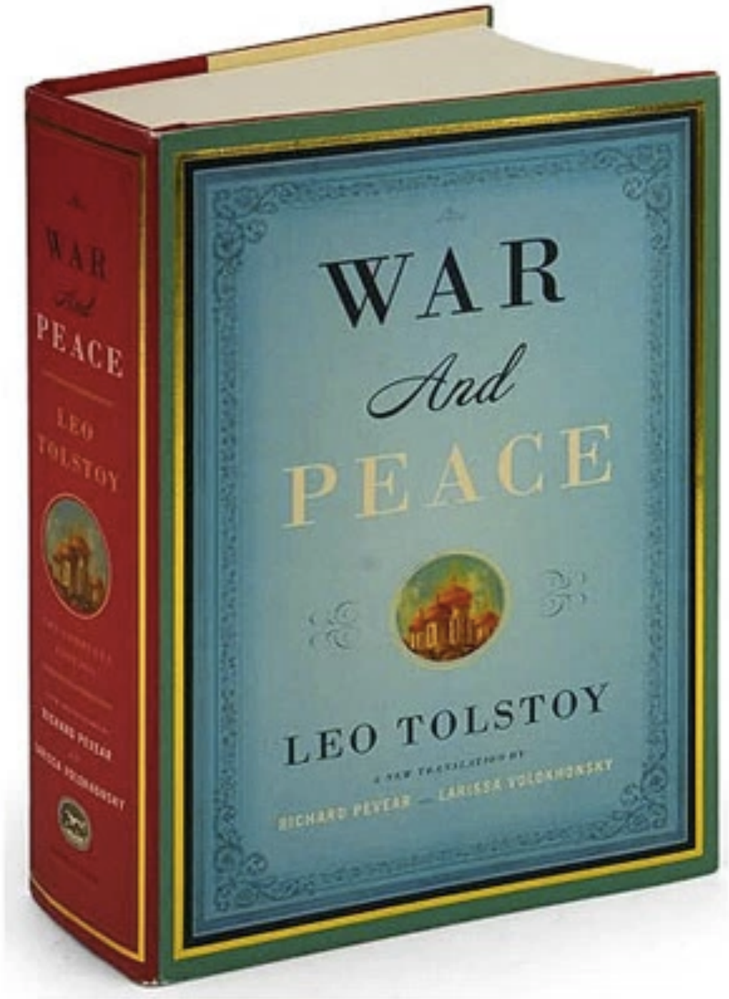

Got OLAP Query?

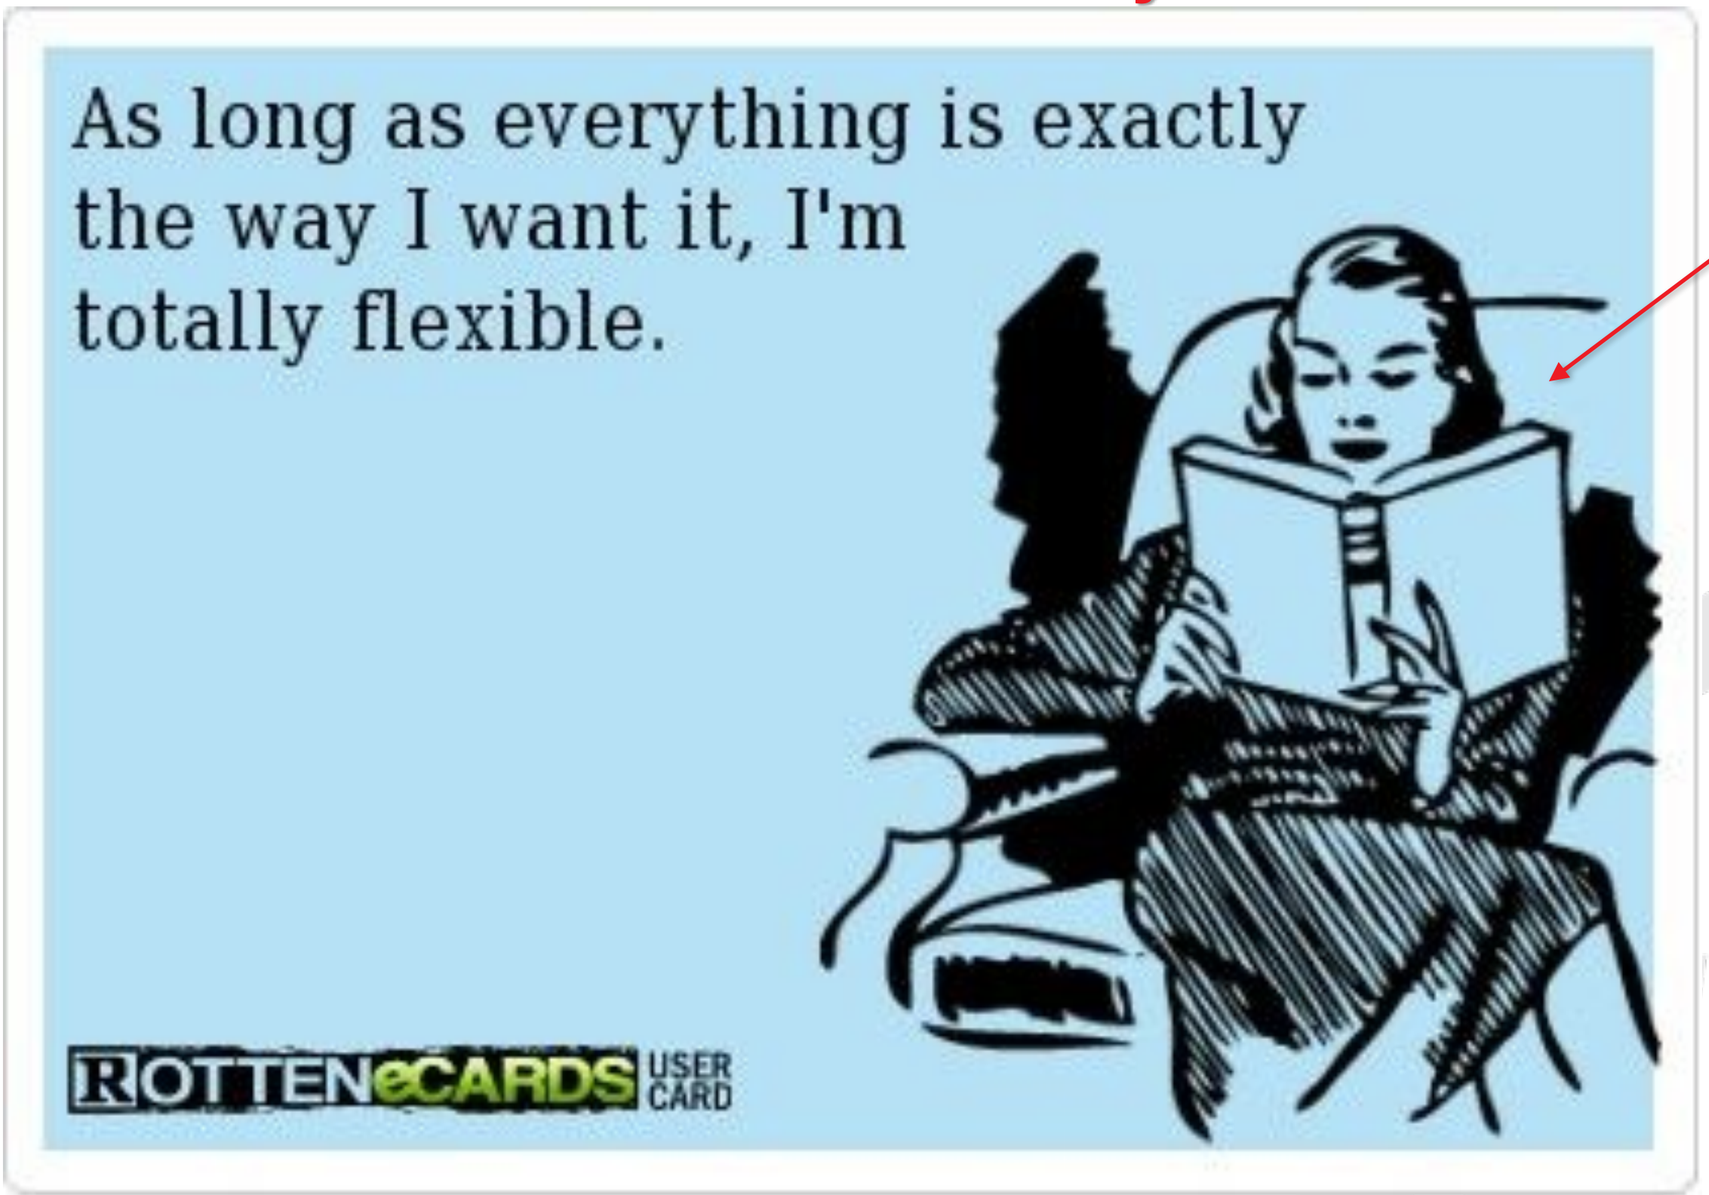

Cassandra

Better Together!

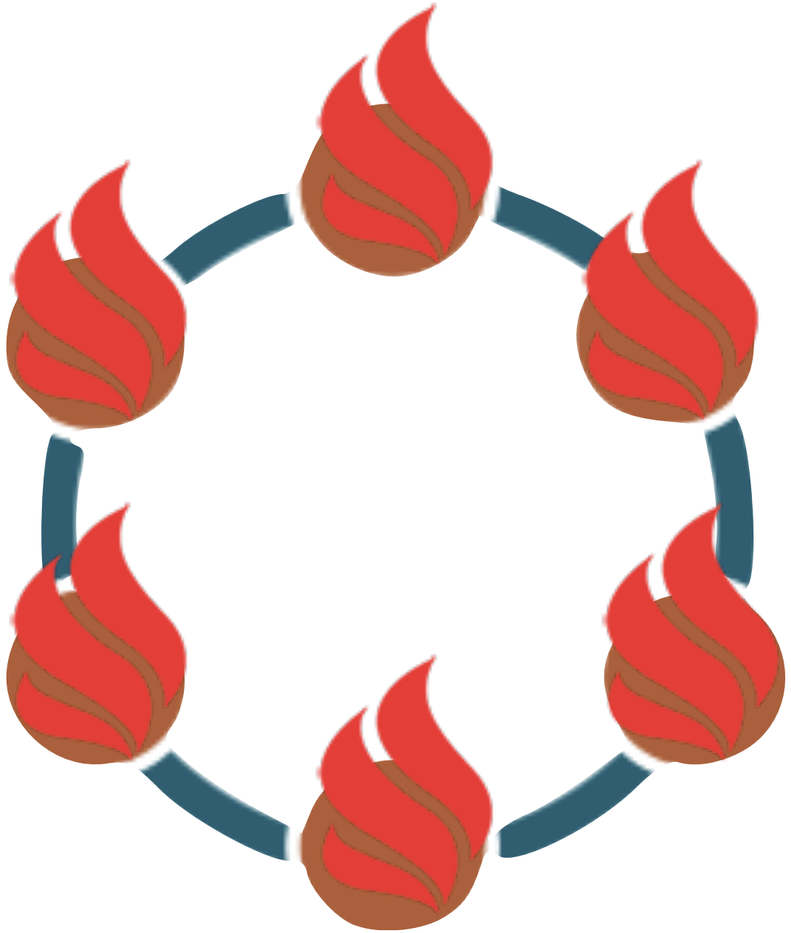

“In-memory will have an industry impact comparable to web and cloud.”

“RAM is the new disk, and disk is the new tape.” Gartner

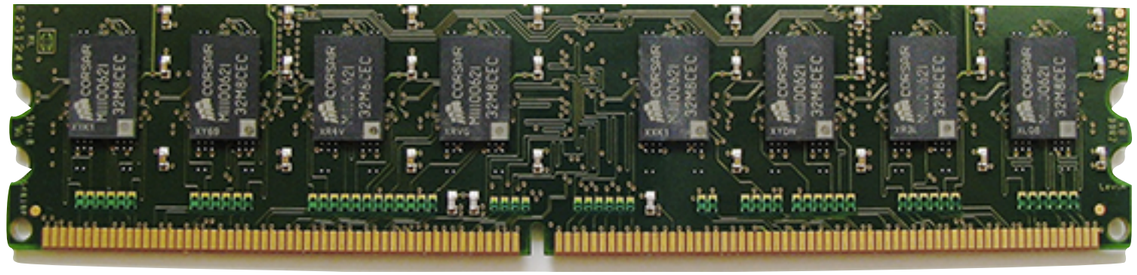

Disk First Architecture

Disk as primary storage, memory for caching

- Access chain: API call <> demarshalling <> OS I/O <> I/O controller <> disk
- Latency: milliseconds

Memory First Architecture

Memory as primary storage, disk as backup

- Access chain: API call <> pointer arithmetic
- Latency: nanoseconds to microseconds

Apache Ignite Project

- **2007:** First version of GridGain
- **Oct. 2014:** GridGain contributes Ignite to ASF
- **Aug. 2015:** Ignite is the second fastest project to graduate after Spark
- **Today:**
  - 60+ contributors and rapidly growing
  - Huge development momentum - Estimated 192 years of effort since the first commit in February, 2014 [Openhub]
  - Mature codebase: 1M+ lines of code

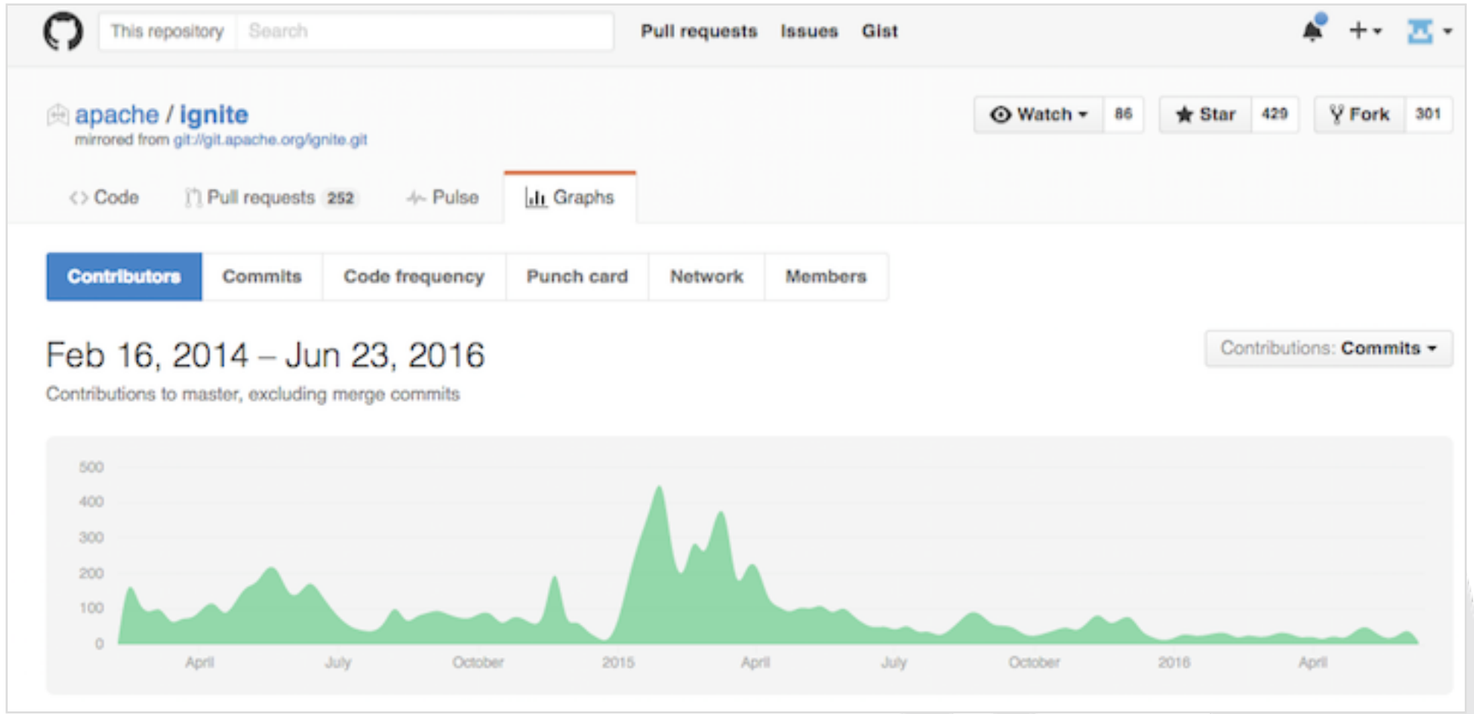

Spark

Application

In-Memory

Scalable

No Rip & Replace

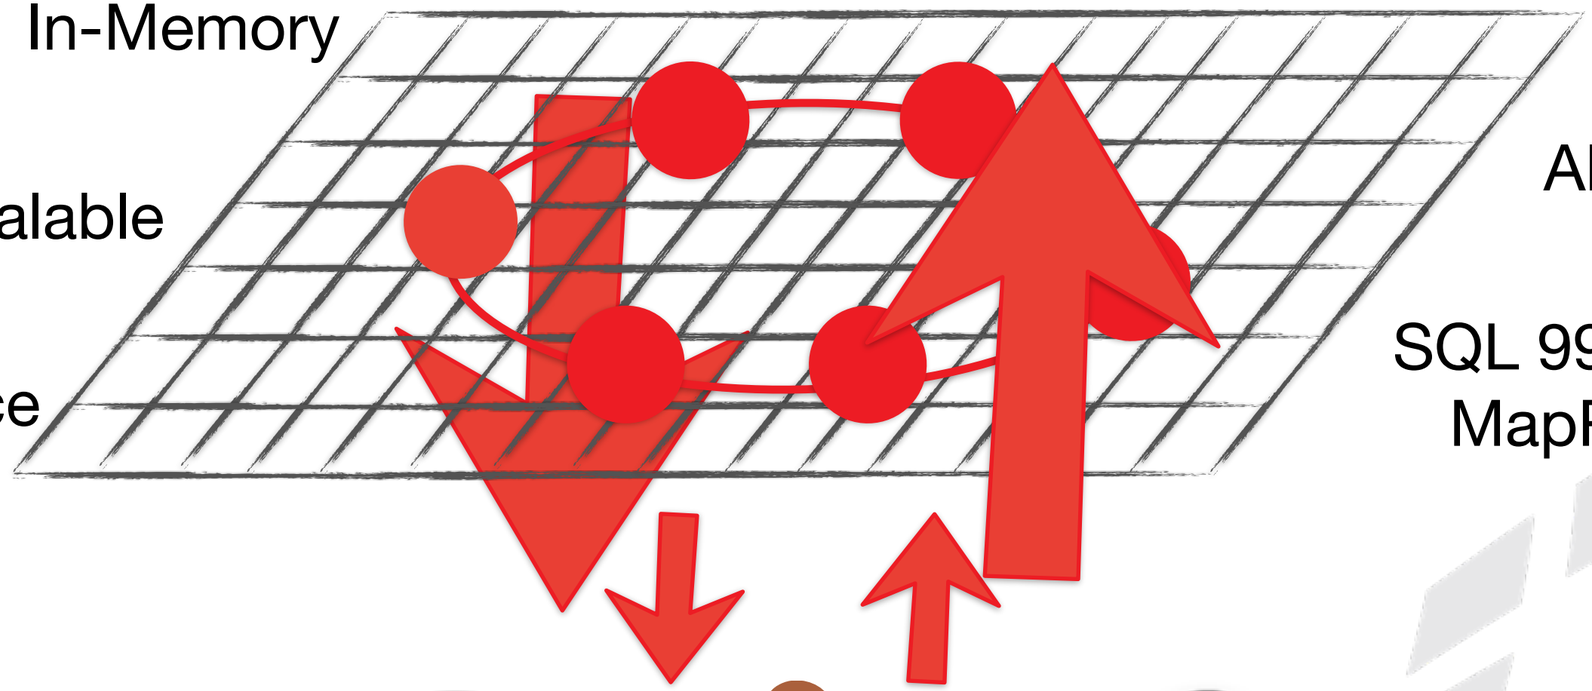

Always Available

SQL 99 / ACID / MapReduce

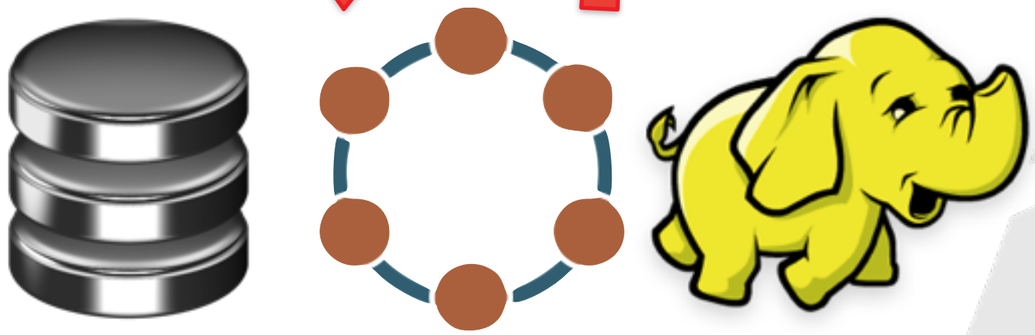

Apache Ignite GridGain

Demo Environment

20 nodes of Cassandra 3.x
3 nodes of Ignite 1.8
3 nodes of test data generation
6500 orders / sec

CPUs: 4 x 2.27 GHz
Memory (RAM): 14.69 GB

Cache being populated by test harness; write through to C*

1: Compare OLTP query between C* and Ignite with no load

```
select * from orders where id = 123456789
```

2: OLAP Queries Ignite / C* under load

```
select p.id as product_id, sum(o.amount) as amount,
sum(o.price) as price
from "product".Product as p, "order".ProductOrder as o
where p.id = o.productId
group by p.id
order by sum(o.amount) desc
limit 10
```

```
Select h.productid, sum(o.amount - h.amount) as amount,
sum(o.price - h.price) as price
from "order_history".ProductOrder h, "order".ProductOrder o
where h.productid = o.productid
group by h.productid
order by 3 desc
limit 10
```

More info

https://ignite.apache.org/

https://github.com/apache/ignite

http://apacheignite.gridgain.org/docs/ignite-with-apache-cassandra

https://issues.apache.org/jira/browse/IGNITE-1371