Cassandra at ebay

Ahbaid Gaffoor, Sr Manager
Feng Qu, Sr MTS

Agenda

- NoSQL Database Strengths & Cautions
- NoSQL Usages at eBay
- Cassandra Use Case Study
- Conclusions

About me

Sr MTS @ eBay Inc.

- Worked on Oracle since 1990s at Osage, DoubleClick, Yahoo and Intuit
- Joined eBay on 2011 to work on company wide NoSQL (Cassandra, MongoDB and Couchbase) initiatives exclusively
- Named 2014 and 2015 DataStax Cassandra MVP
- Speaker at 2013, 2014 and 2015 Cassandra Summit
- Speaker at EDW 2016

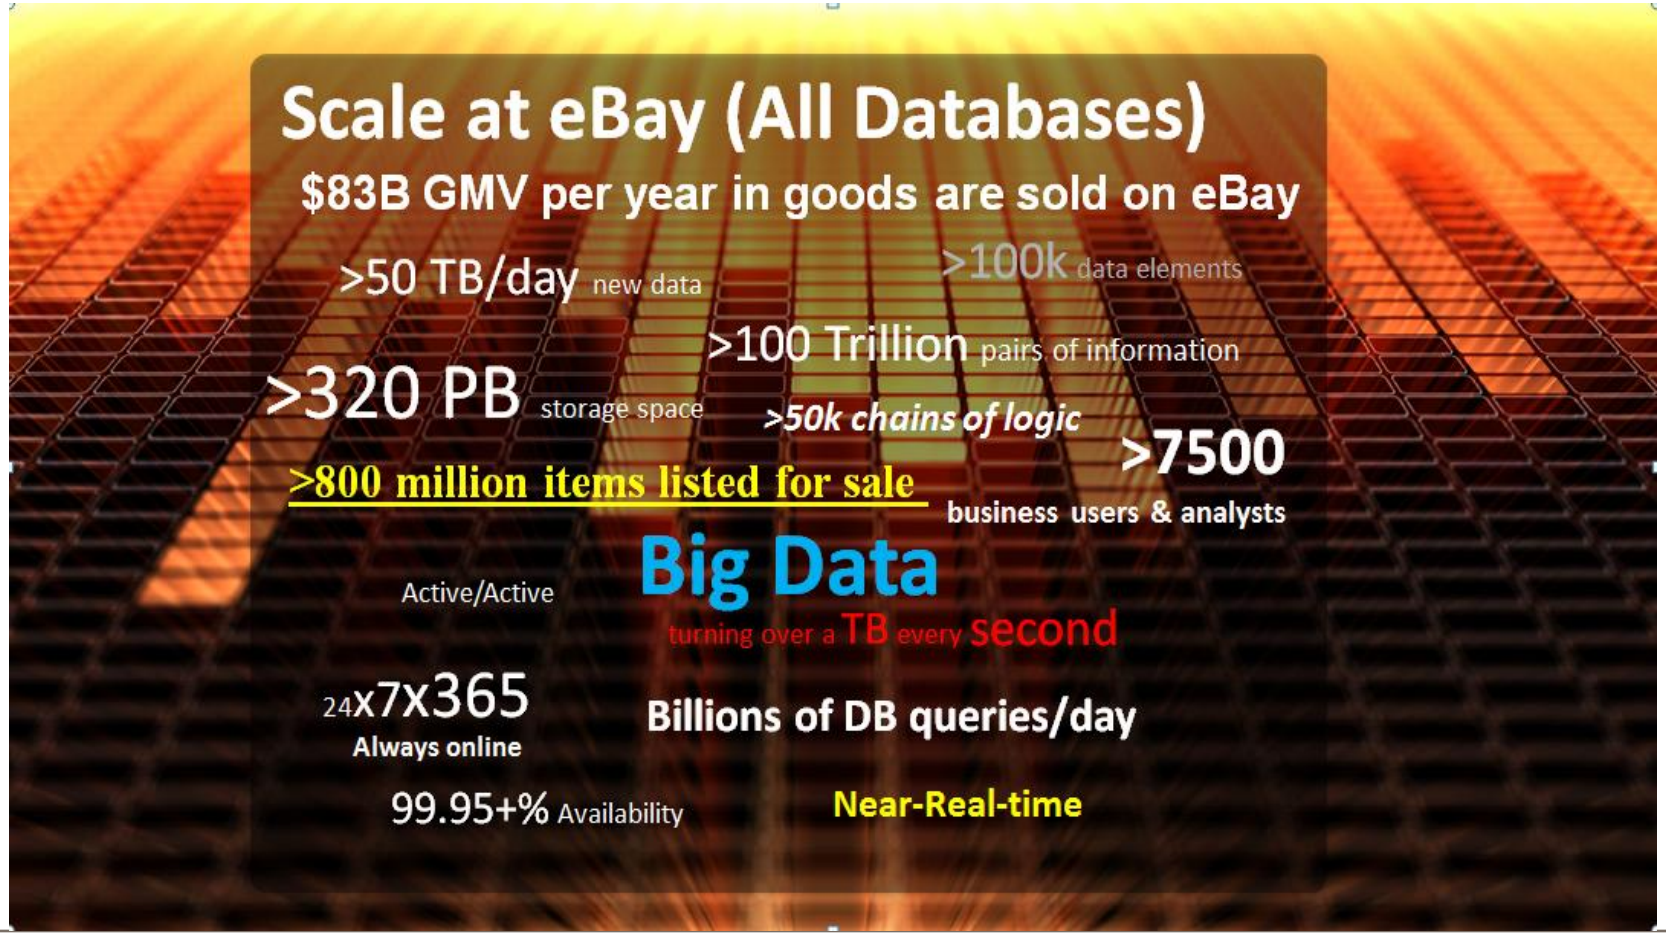
Scale at eBay (All Databases)
$83B GMV per year in goods are sold on eBay
>50 TB/day new data
>100k data elements
>100 Trillion pairs of information
>320 PB storage space
>50k chains of logic
>800 million items listed for sale
>7500
business users & analysts
Active/Active
Big Data
turning over a TB every second
24x7x365
Always online
Billions of DB queries/day
99.95+% Availability
Near-Real-time

Why NoSQL?

- Challenges of traditional RDBMS
  - Performance penalty to maintain ACID features
  - Lack of native sharding and replication features
  - Lack of linear scalability
  - Cost of software/hardware
  - Higher cost of commit
- Time to think about NoSQL?

Gartner Magic Quadrant for Operational Databases

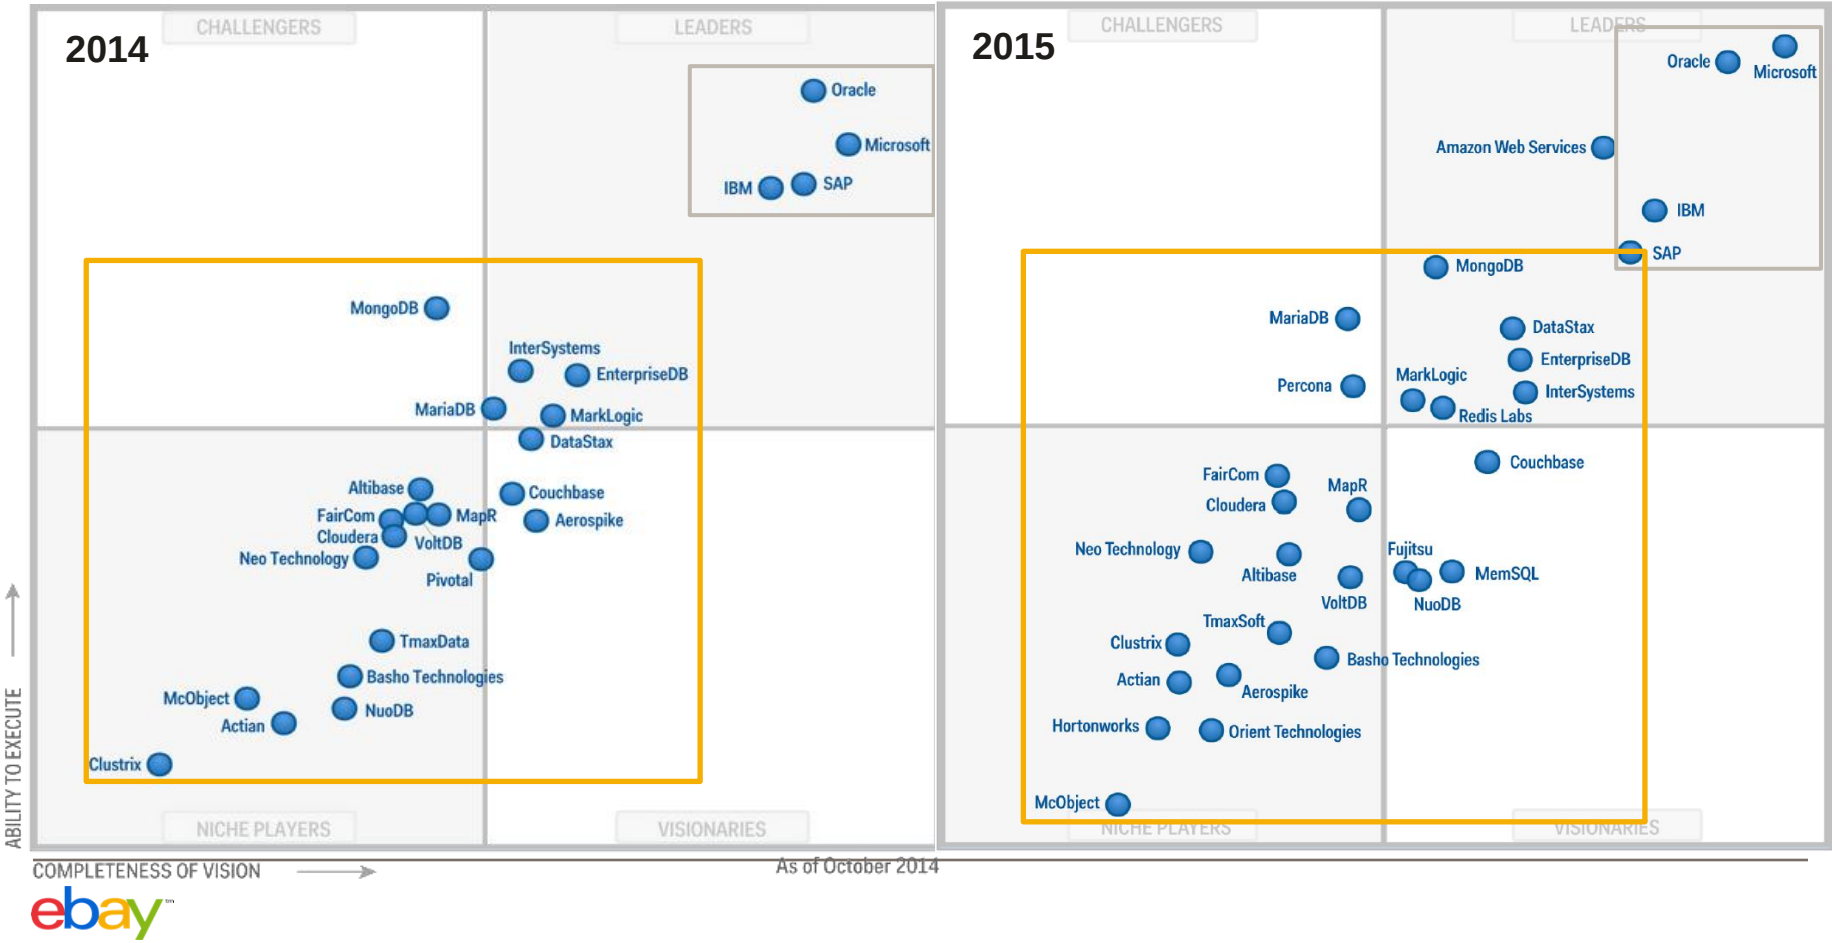

ebay

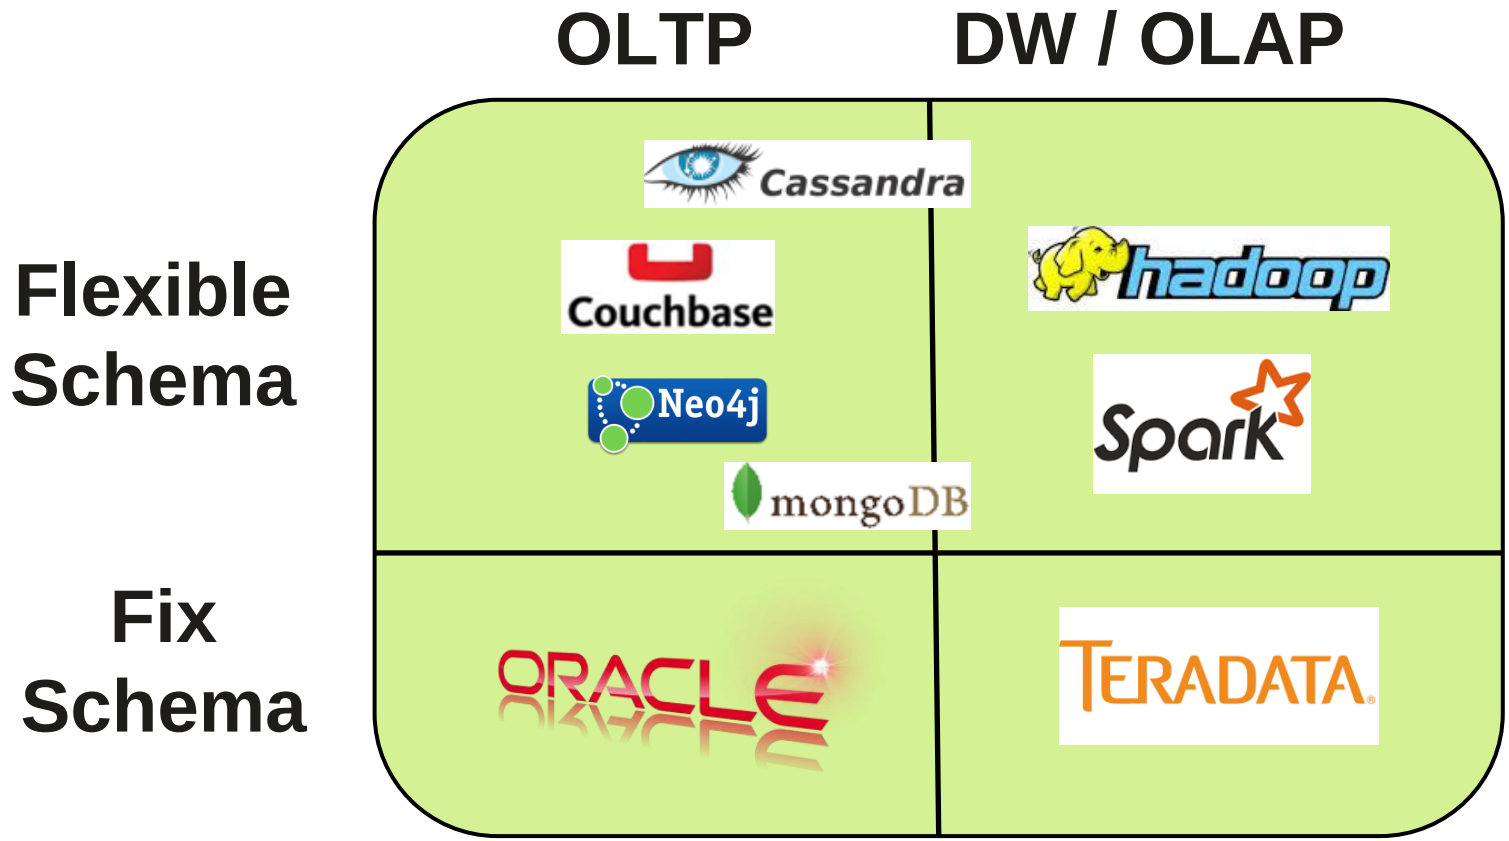
OLTP
DW / OLAP
Flexible Schema
Fix Schema
Cassandra
Couchbase
Neo4j
mongoDB
hadoop
Spark
ORACLE
TERADATA

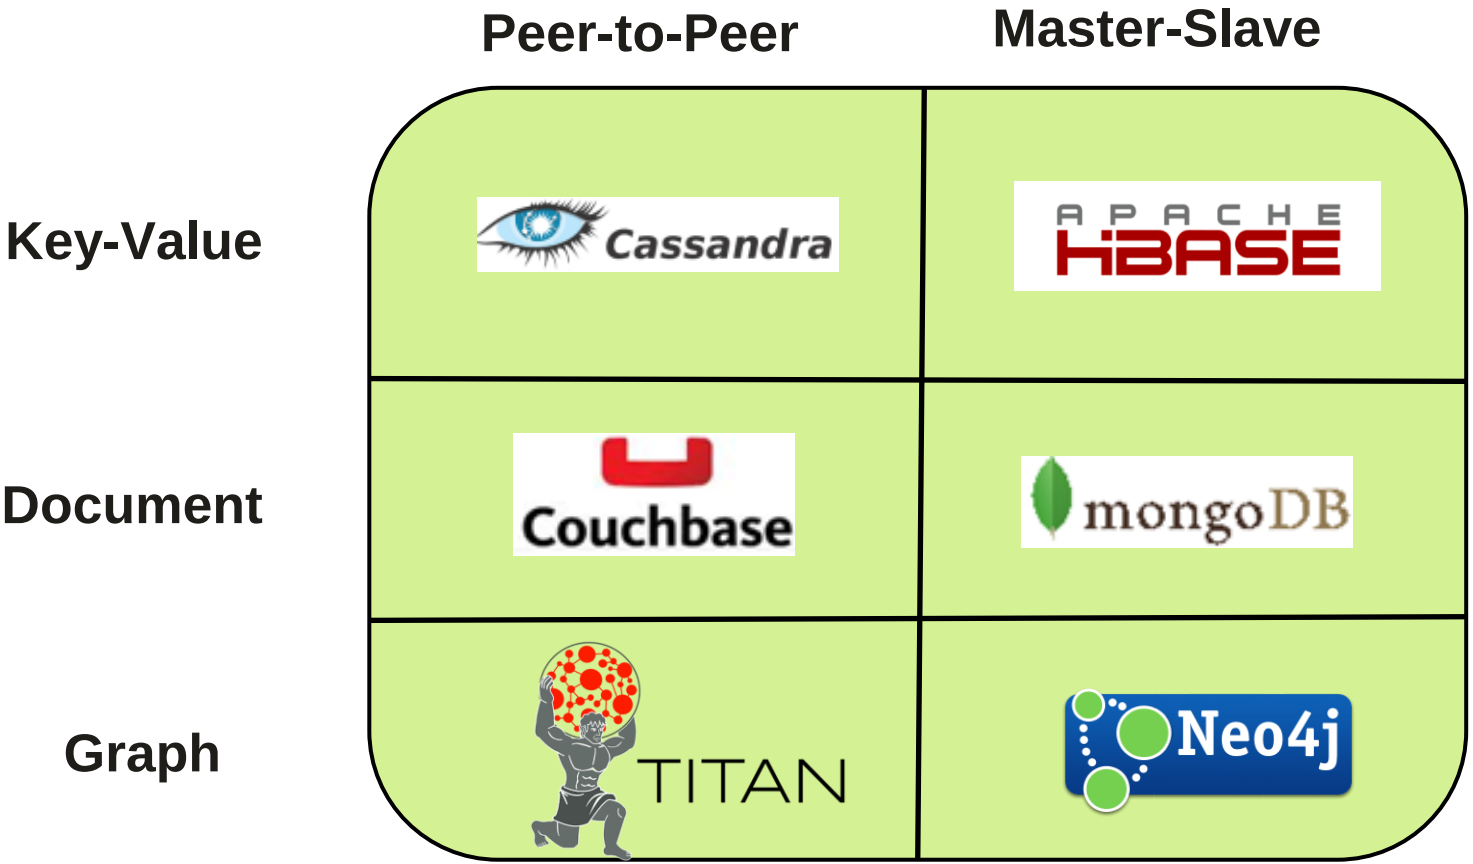
Peer-to-Peer
Master-Slave
Key-Value
Document
Graph
Cassandra
APACHE HBASE
Couchbase
mongoDB
TITAN
Neo4j

It’s all about Data!

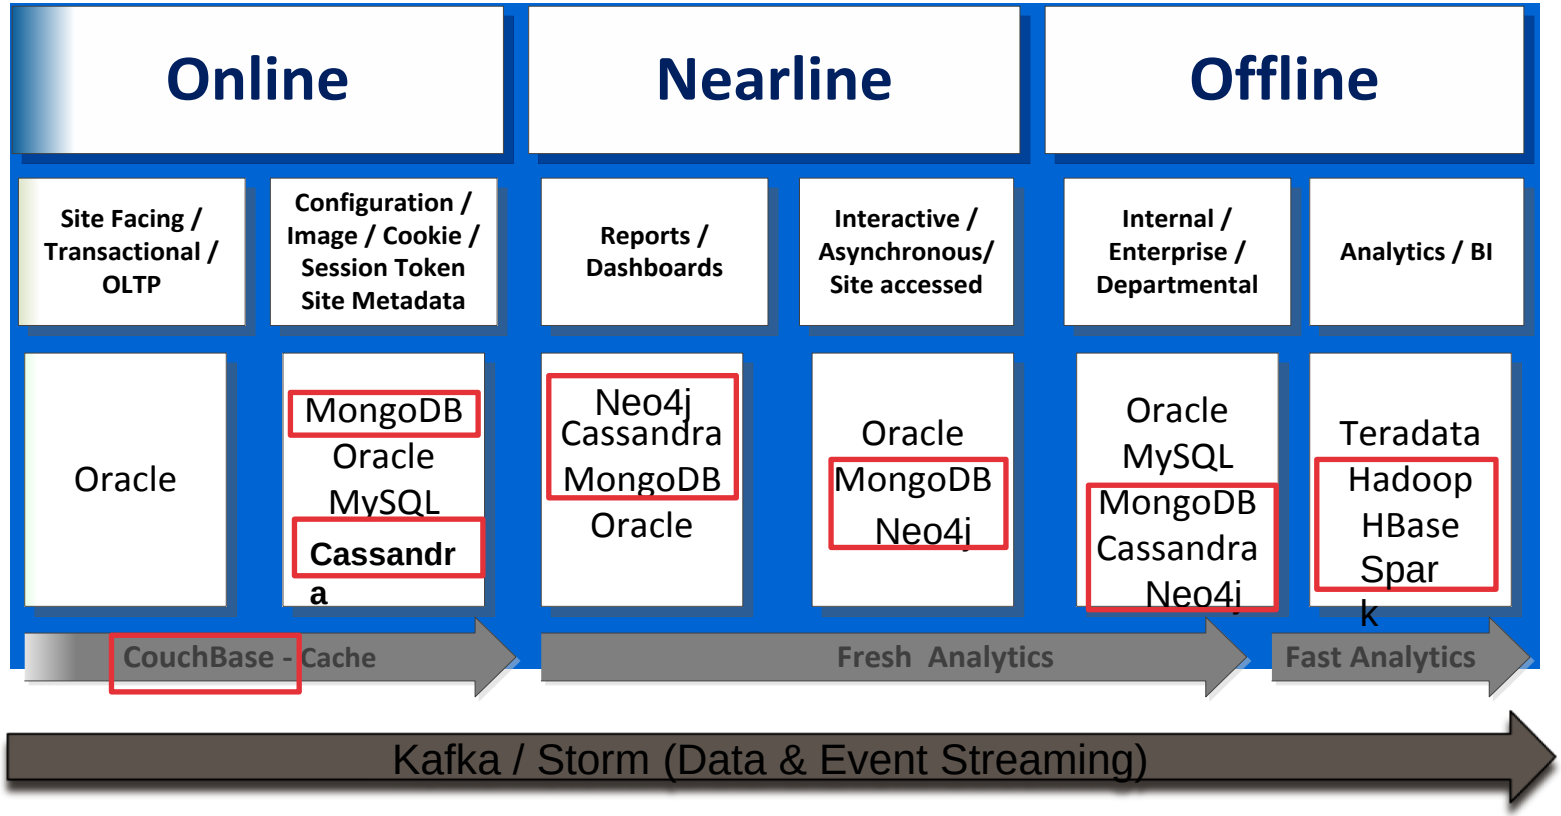

NoSQL Journey @ eBay

- MongoDB since 2010
- Cassandra since 2011
- Couchbase since 2014
- Thousands of NoSQL nodes in multi data centers
  - Bare metals or VM clusters
  - Dedicated or multi-tendency
- Multi petabytes of data
- Tens of billions database calls per day

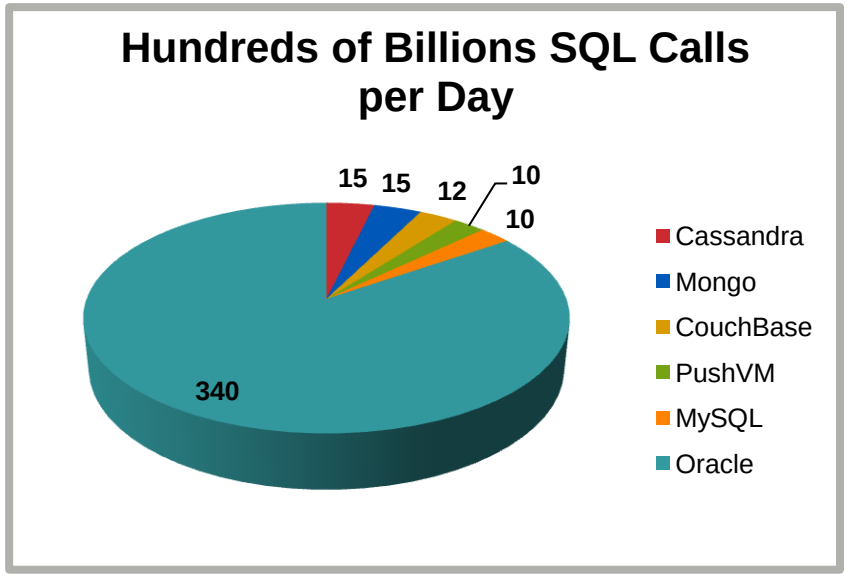

NoSQL Database Strengths and Cautions

- Geo distributed replication & sharding
- Location aware low latency query performance
- Workload & access pattern optimized
- Linear scalability with reduced disruption to business
- Supports semi- or un-structured data
- Flexible schema provides significant increase in Dev agility

- Lack strict ACID compliant transaction
- Lack strong data model control & governance
- Not suitable for ad-hoc workload & random access pattern
- Requires change of mindset, ecosystem and infrastructure
- Rapidly changing technology & competitive landscape
- Requires Dev expertise in nuances of distributed systems

MongoDB

- Dev friendly rich JSON document model
- Secondary index enables mixed access patterns
- High business value (semi-) structure data
- Balanced scale-out reads & writes (with optional sharding)
- Straightforward admin effort

- Short write interruption during primary re-election
- Not suitable for nanosecond latency writing
- Potentially high TCO for large scale sharded cluster
- Lack resource isolation

Cassandra

- Peer-to-peer without SPOF (Single Point of Failure)
- Active-active cross Datacenter
- High read & very high write performance
- Absolute linear scalability

- Inefficient secondary index (pre-V3)
- Not suitable for mixed user query & access patterns
- High compaction overhead for frequent random deletes
- Require JVM tuning to mitigate GC pauses
- Lack resource isolation
- Slow cluster rebalancing

Couchbase

- Memcached compatible persistent document store
- Peer-to-peer architecture
- High read & write performance
- Active-active cluster replication
- Strong local cluster RW consistency
- Resource isolation

- Short write interruption during node failover
- Inefficient secondary index (pre-V4)
- Counter intuitive cross DC write conflict resolution
- Slow cluster rebalancing
- Slow warm-up

Typical Cassandra Use Cases

- Write Intensive: metrics collection
  - Collecting metrics from tens of thousands devices periodically
- Read Intensive: home page feeds
  - Recommendation backend to generate dynamic taste graph
- Mixed workload: personalization, classification
  - Data is loaded from data warehouse periodically in bulk and from user events stream consistently
  - Data is retrieved in real time when user visits eBay site

Use Case Study – Next Gen Classification(NGC)

- Event based classification framework service used by critical eBay business workflows like selling, buying & search on listing content, pricing, attribute …
- It provides flexible & easy-to-use editorial with on-demand deployment on new or revised classification items
- What's special about NGC?
  - Tens of billions of transactions per day and workload continues to grow year over year
  - Very tight reads & write performance
  - Significant GMV impact
  - Global operation

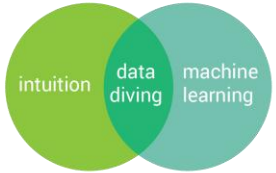

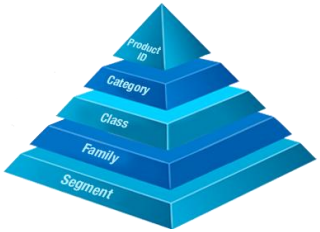

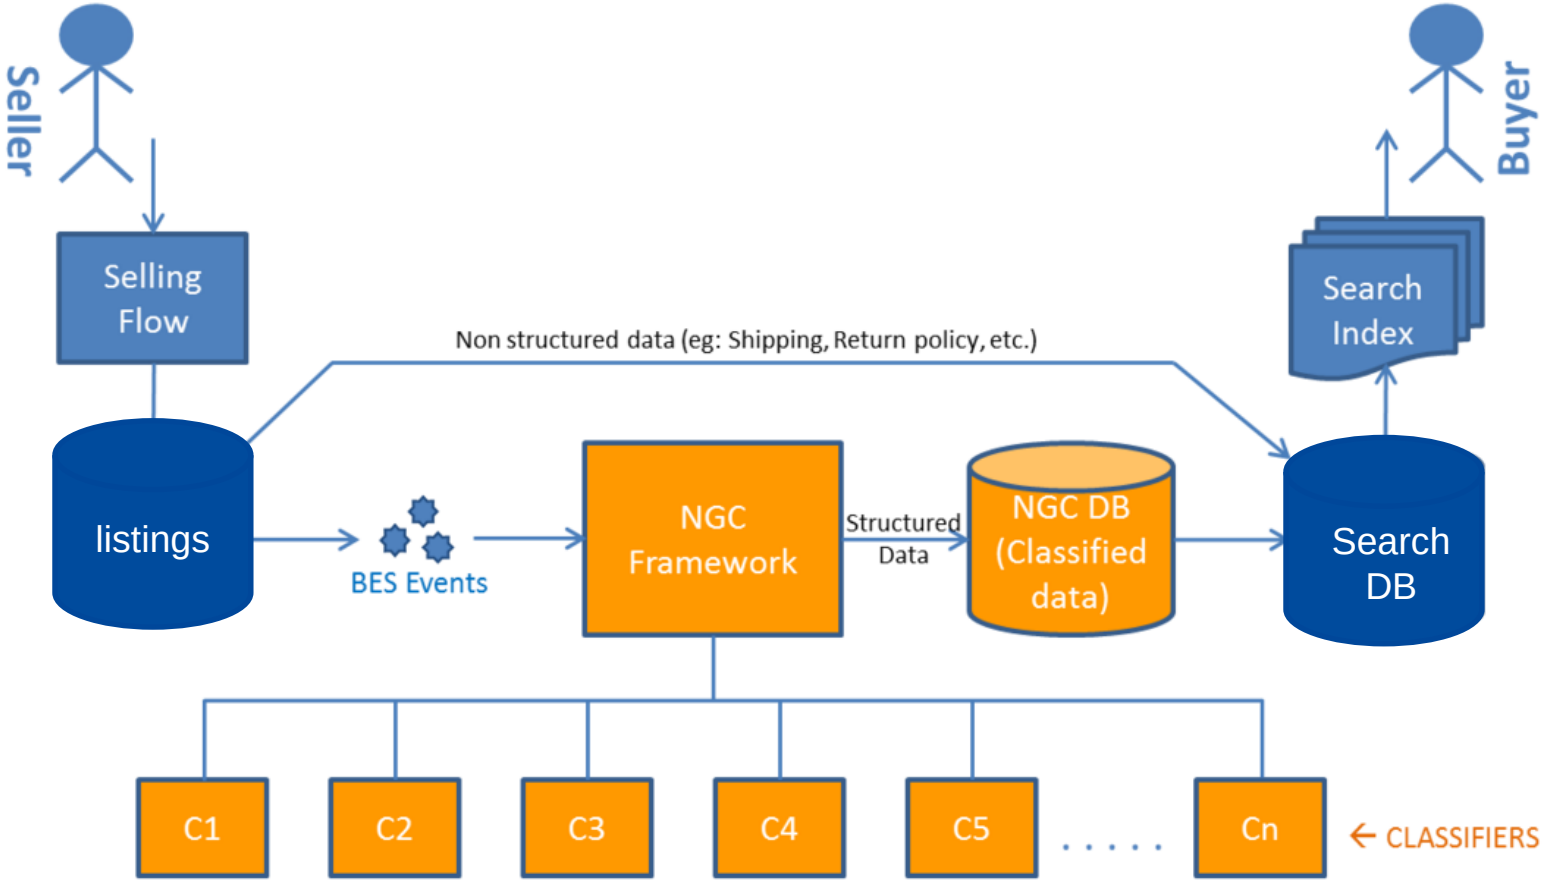
Seller
Buyer
Selling Flow
Search Index
Non structured data (eg: Shipping, Return policy, etc.)
listings
BES Events
NGC Framework
Structured Data
NGC DB (Classified data)
Search DB
C1
C2
C3
C4
C5
. . . . .
Cn
← CLASSIFIERS

Operation Best Practices

- Thorough use case qualification is critical!
  - Availability, DR, SLA, data model, access pattern …
- Set proper stakeholder mindset & expectation
  - BO/PM, DA, PD, DBA, DevOps …
- Understand infrastructure lifecycle management capability
- Benchmark testing and capacity planning
- In doubt of stated capacity numbers, over provision to handle unplanned capacity increase
- Although some NoSQL are easier to manage than others, they all need SOPs! Operation automation comes to rescue.
- Open sourced NoSQL requires additional management skills
- Ongoing performance tuning
- Scale out vs. scale up

Conclusions

- Compared RDBMS with 4 leading NoSQL databases
- Different databases for different jobs
- NoSQL close to end of early maturity cycle
- New NoSQL product feature trend
  - Pluggable storage engines, such as WiredTiger and ForestDB
  - SQL like query capability, such as CQL and N1SQL
  - Always available distributed ACID Transactions
  - Data model governance & security
- NewSQL
  - Is this where SQL meets NoSQL? Time will tell.


Q&A