

Polyglot

Persistence



Majid Fatemian
@majidfn
Ubisoft Montreal


Relational



DATA

**BIG
DATA**

NoSQL



1

Non-Relational

2

Cluster

3

Schemaless

1. Document

MongoDB, CouchBase

2. Key-Value

Memcache, Redis


3. Graph

Neo4J, OrientDB

4. Column

Cassandra, HBase

Availability



Consistency

Partition Tolerance

Consistency **A**vailability



MySQL

Postgres

Availability **P**artition



Cassandra

Riak

Consistency **P**artition



MongoDB



CouchBase

YOUR SHAPE

Fitness Evolved



YOUR SHAPE

Fitness Evolved



 XBOX 360

 XBOX ONE

Wii U

YOUR SHAPE


Fitness Evolved



 **XBOX** 360™

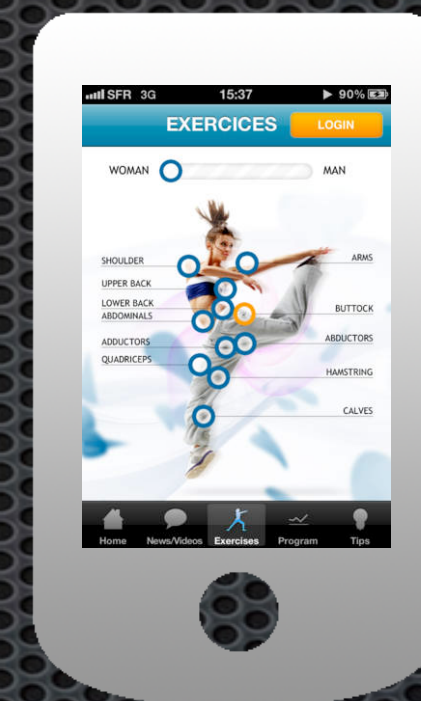
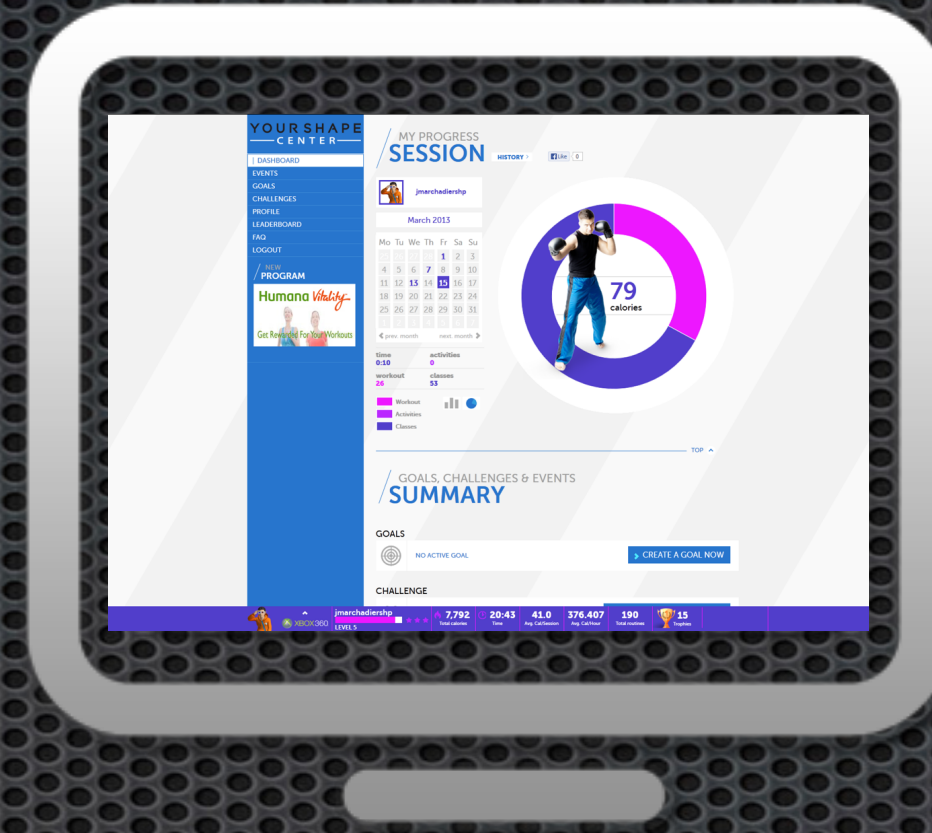
 **XBOX** ONE™

Wii 


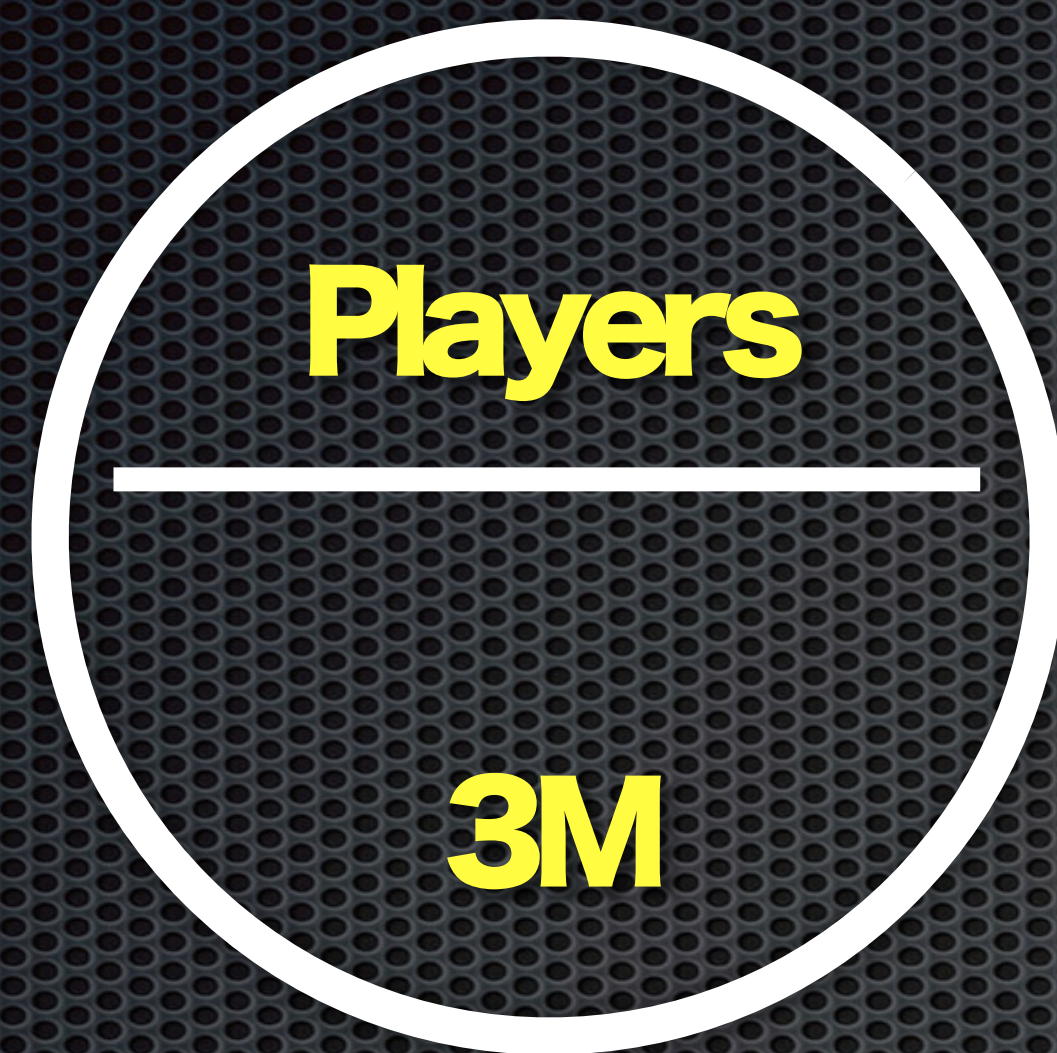


YOUR SHAPE

Fitness Evolved




STATS






2012



2014

2011



2013





2012




2014


2011




2013







Application



Norm

Relational
Consistency

{ ACID

Trans

Relational
Availability

{ Replication
Table Partitioning

OPS /

{ Tools

Devs

Relational
DevOps




2012



2014

2011



2013



Aggregated



Profile
Community


Detailed



History
Leaderboards

Single Point Reliability

Relational



Relational
Scalability



```
{  
  player_id: 1234,  
  calories: 75,  
  duration: 120,  
  activity: boxing  
}
```




```
{  
  player_id: 1234,  
  * calories: 40,  
  * duration: 75,  
  * activity: jumping_ropes,  
  game_id: 4001,  
  platform: WiiU,  
  difficulty_level: medium,  
  repetition: 178,  
  score: 450,  
  endurance: 3  
}
```

Relational Schema

Player	Game	Cals	Duration	Activity	Difficulty	Score	Reps.
1000	1	100	80	Box	-	-	-
2000	2	90	100	Jump	H	4.5	180
1000	2	80	120	Pushup	M	4	130
2000	1	120	110	Situps	-	-	-

Player	Game	Cals	Duration	Activity	Difficulty	Score	Reps.
1000	1	100	80	Box	-	-	-
2000	2	90	100	Jump	H	4.5	180
1000	2	80	120	Pushup	M	4	130
2000	1	120	110	Situps	-	-	-



Relational
Distribution




ACID
Transaction
Partitioning
Replication
Tools
OPs



Data Volume
Fixed Schema
Scale Up
Manual Sharding



NoSQL



mongoDB



1

OPs


2

DEVs

3

Community

Non-Relational Schemaless




player_id: 1234,
calories: 75,
duration: 120,
activity: boxing

player_id: 1234,
* calories: 40,
* duration: 75,
* activity: jumps,
game_id: 4001,
platform: WiiU,
difficulty_level: 2,
repetition: 178,
score: 450,
endurance: 3

Non-Relational Scalability



Non-Relational Scalability



Non-Relational Sharding



Non-Relational Sharding



Non-Relational Sharding

- DB Managed
- Balanced Load

- Sharding key
- Bouncing
- Decreased resilience

Non-Relational **Map/Reduce**



Query

Non-Relational
Map/Reduce



Non-Relational Map/Reduce

Map



Non-Relational Map/Reduce

Map

Reduce




rewrite / Denormalized
Disk Space

Non-Relational

Data Model

Query


Non-Relational




Non-Relational
Eventual-
Consistency

X2

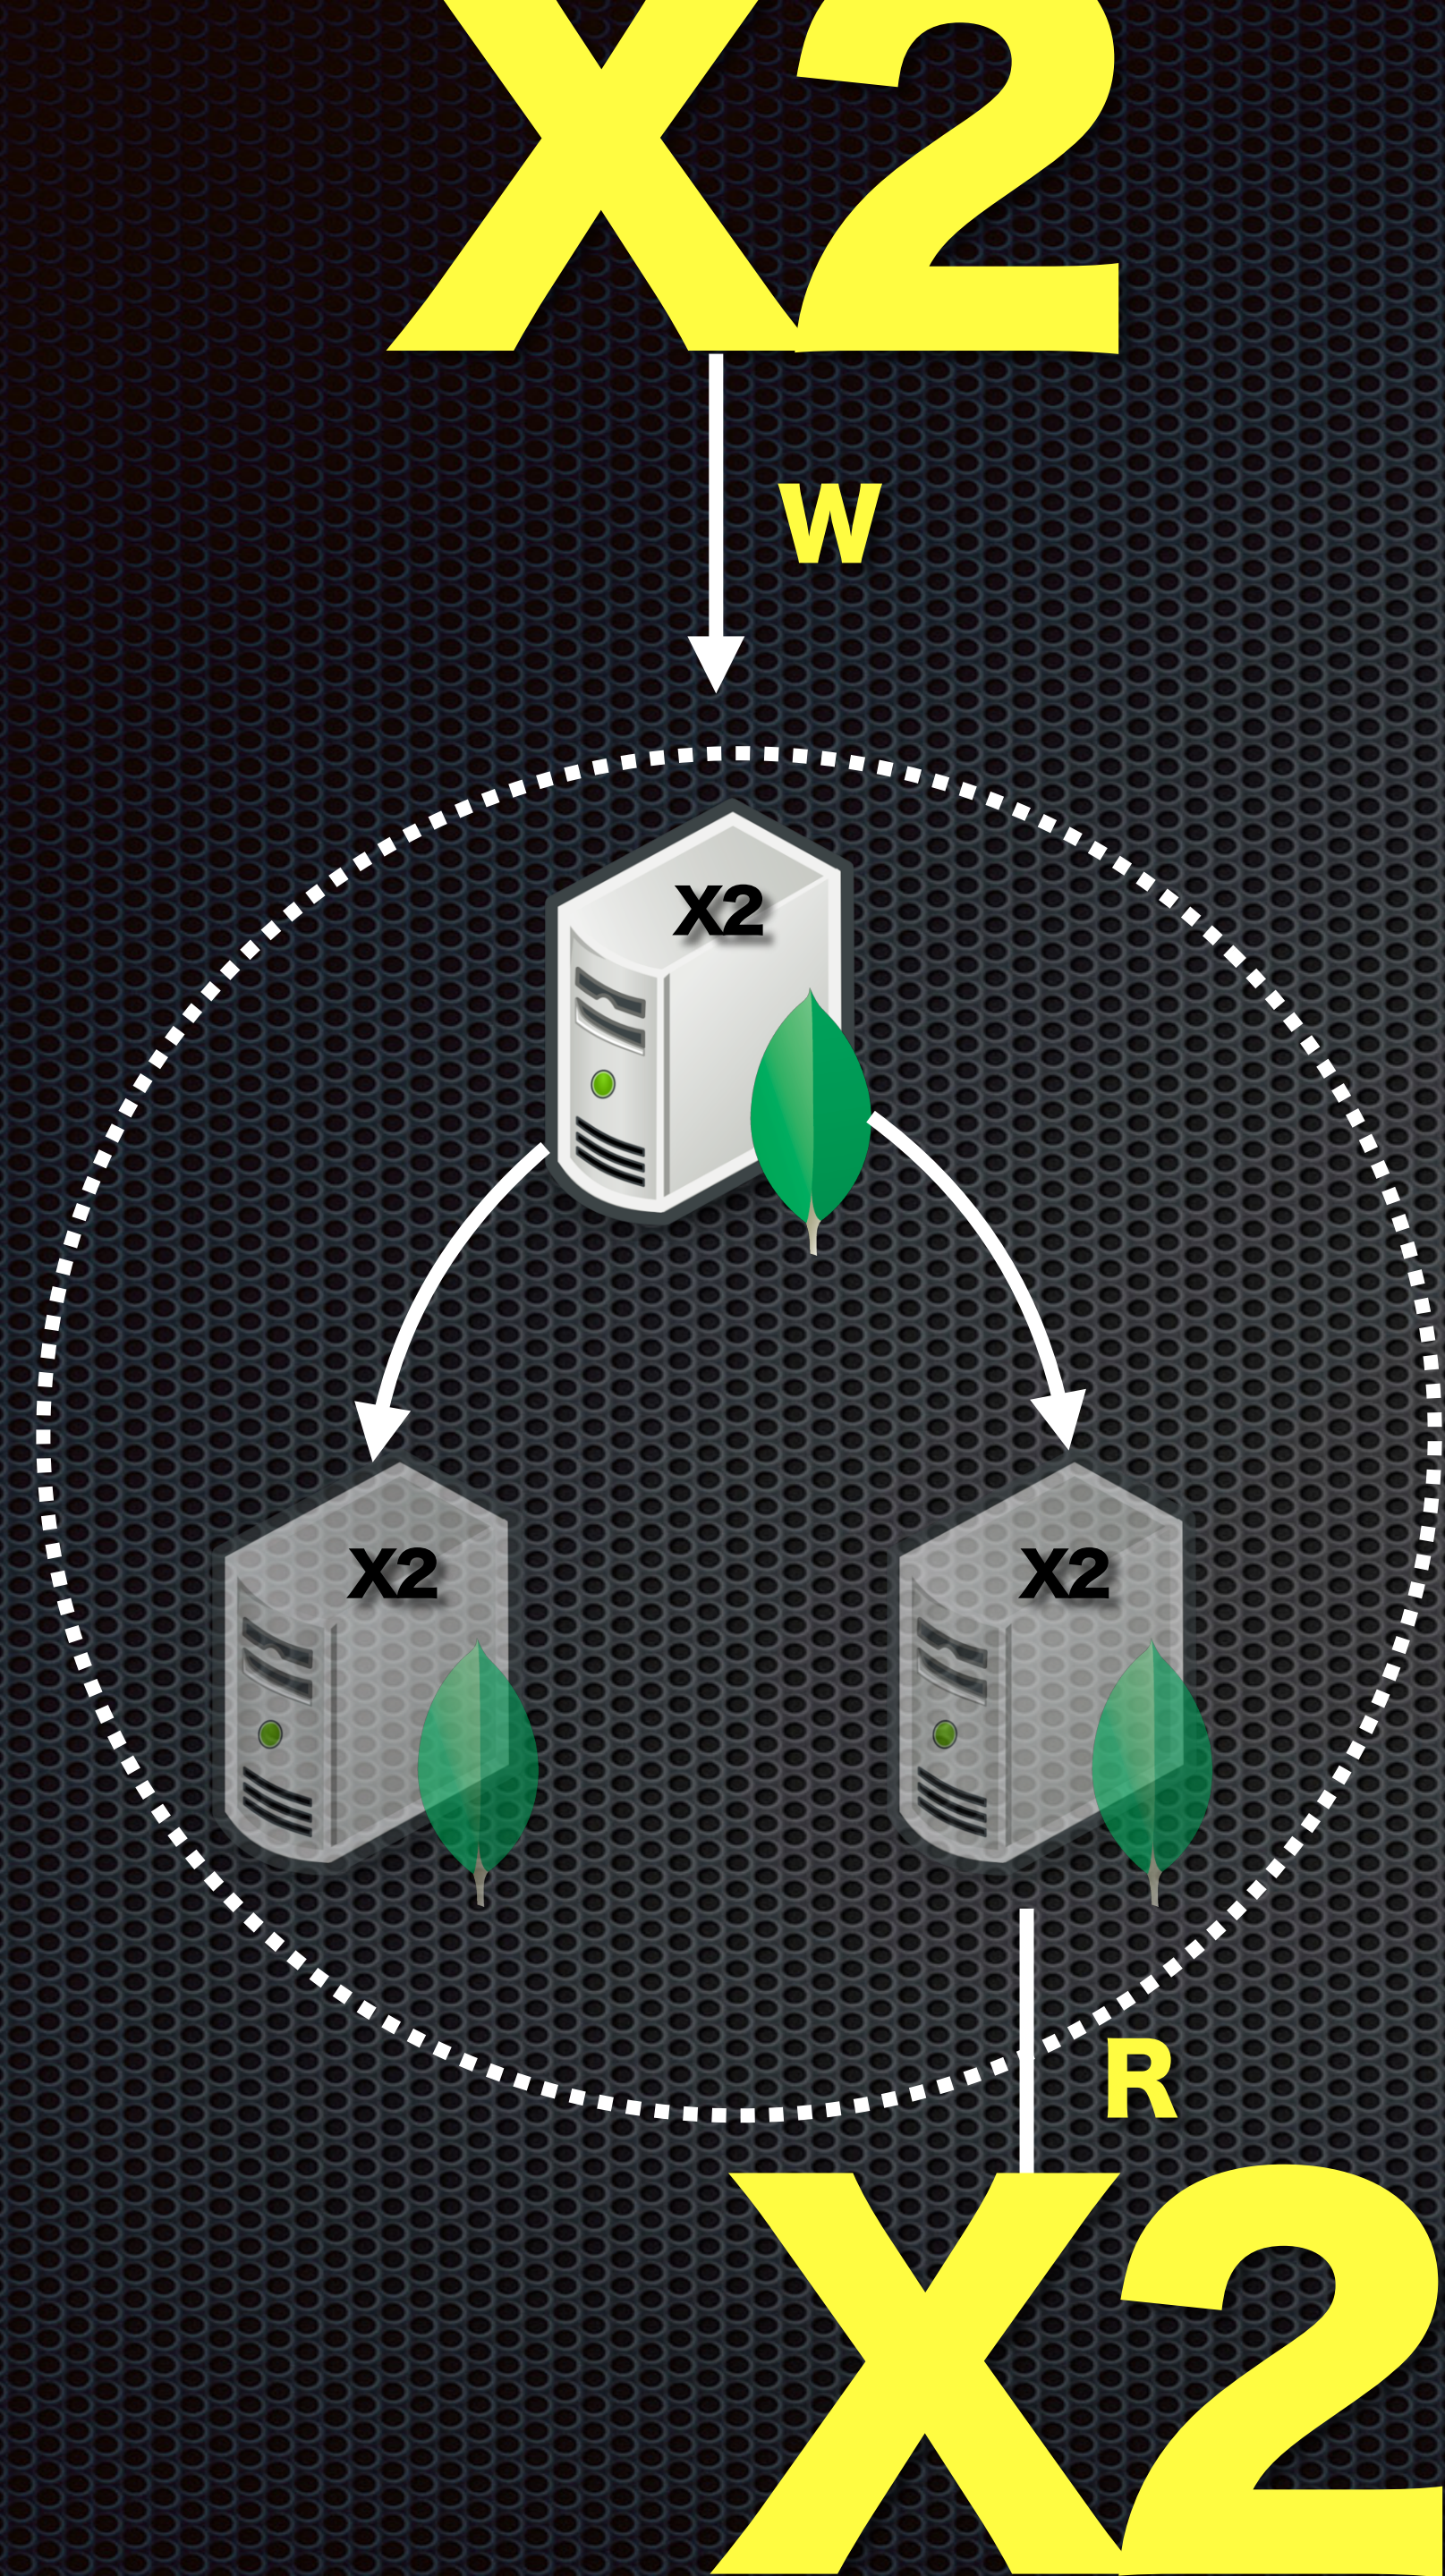
W



Non-Relational
Eventual-
Consistency



Non-Relational
Eventual-
Consistency



Non-Relational
Eventual-
Consistency

Document Level
Locking

(3.0)

Point-in-time Relational Backup

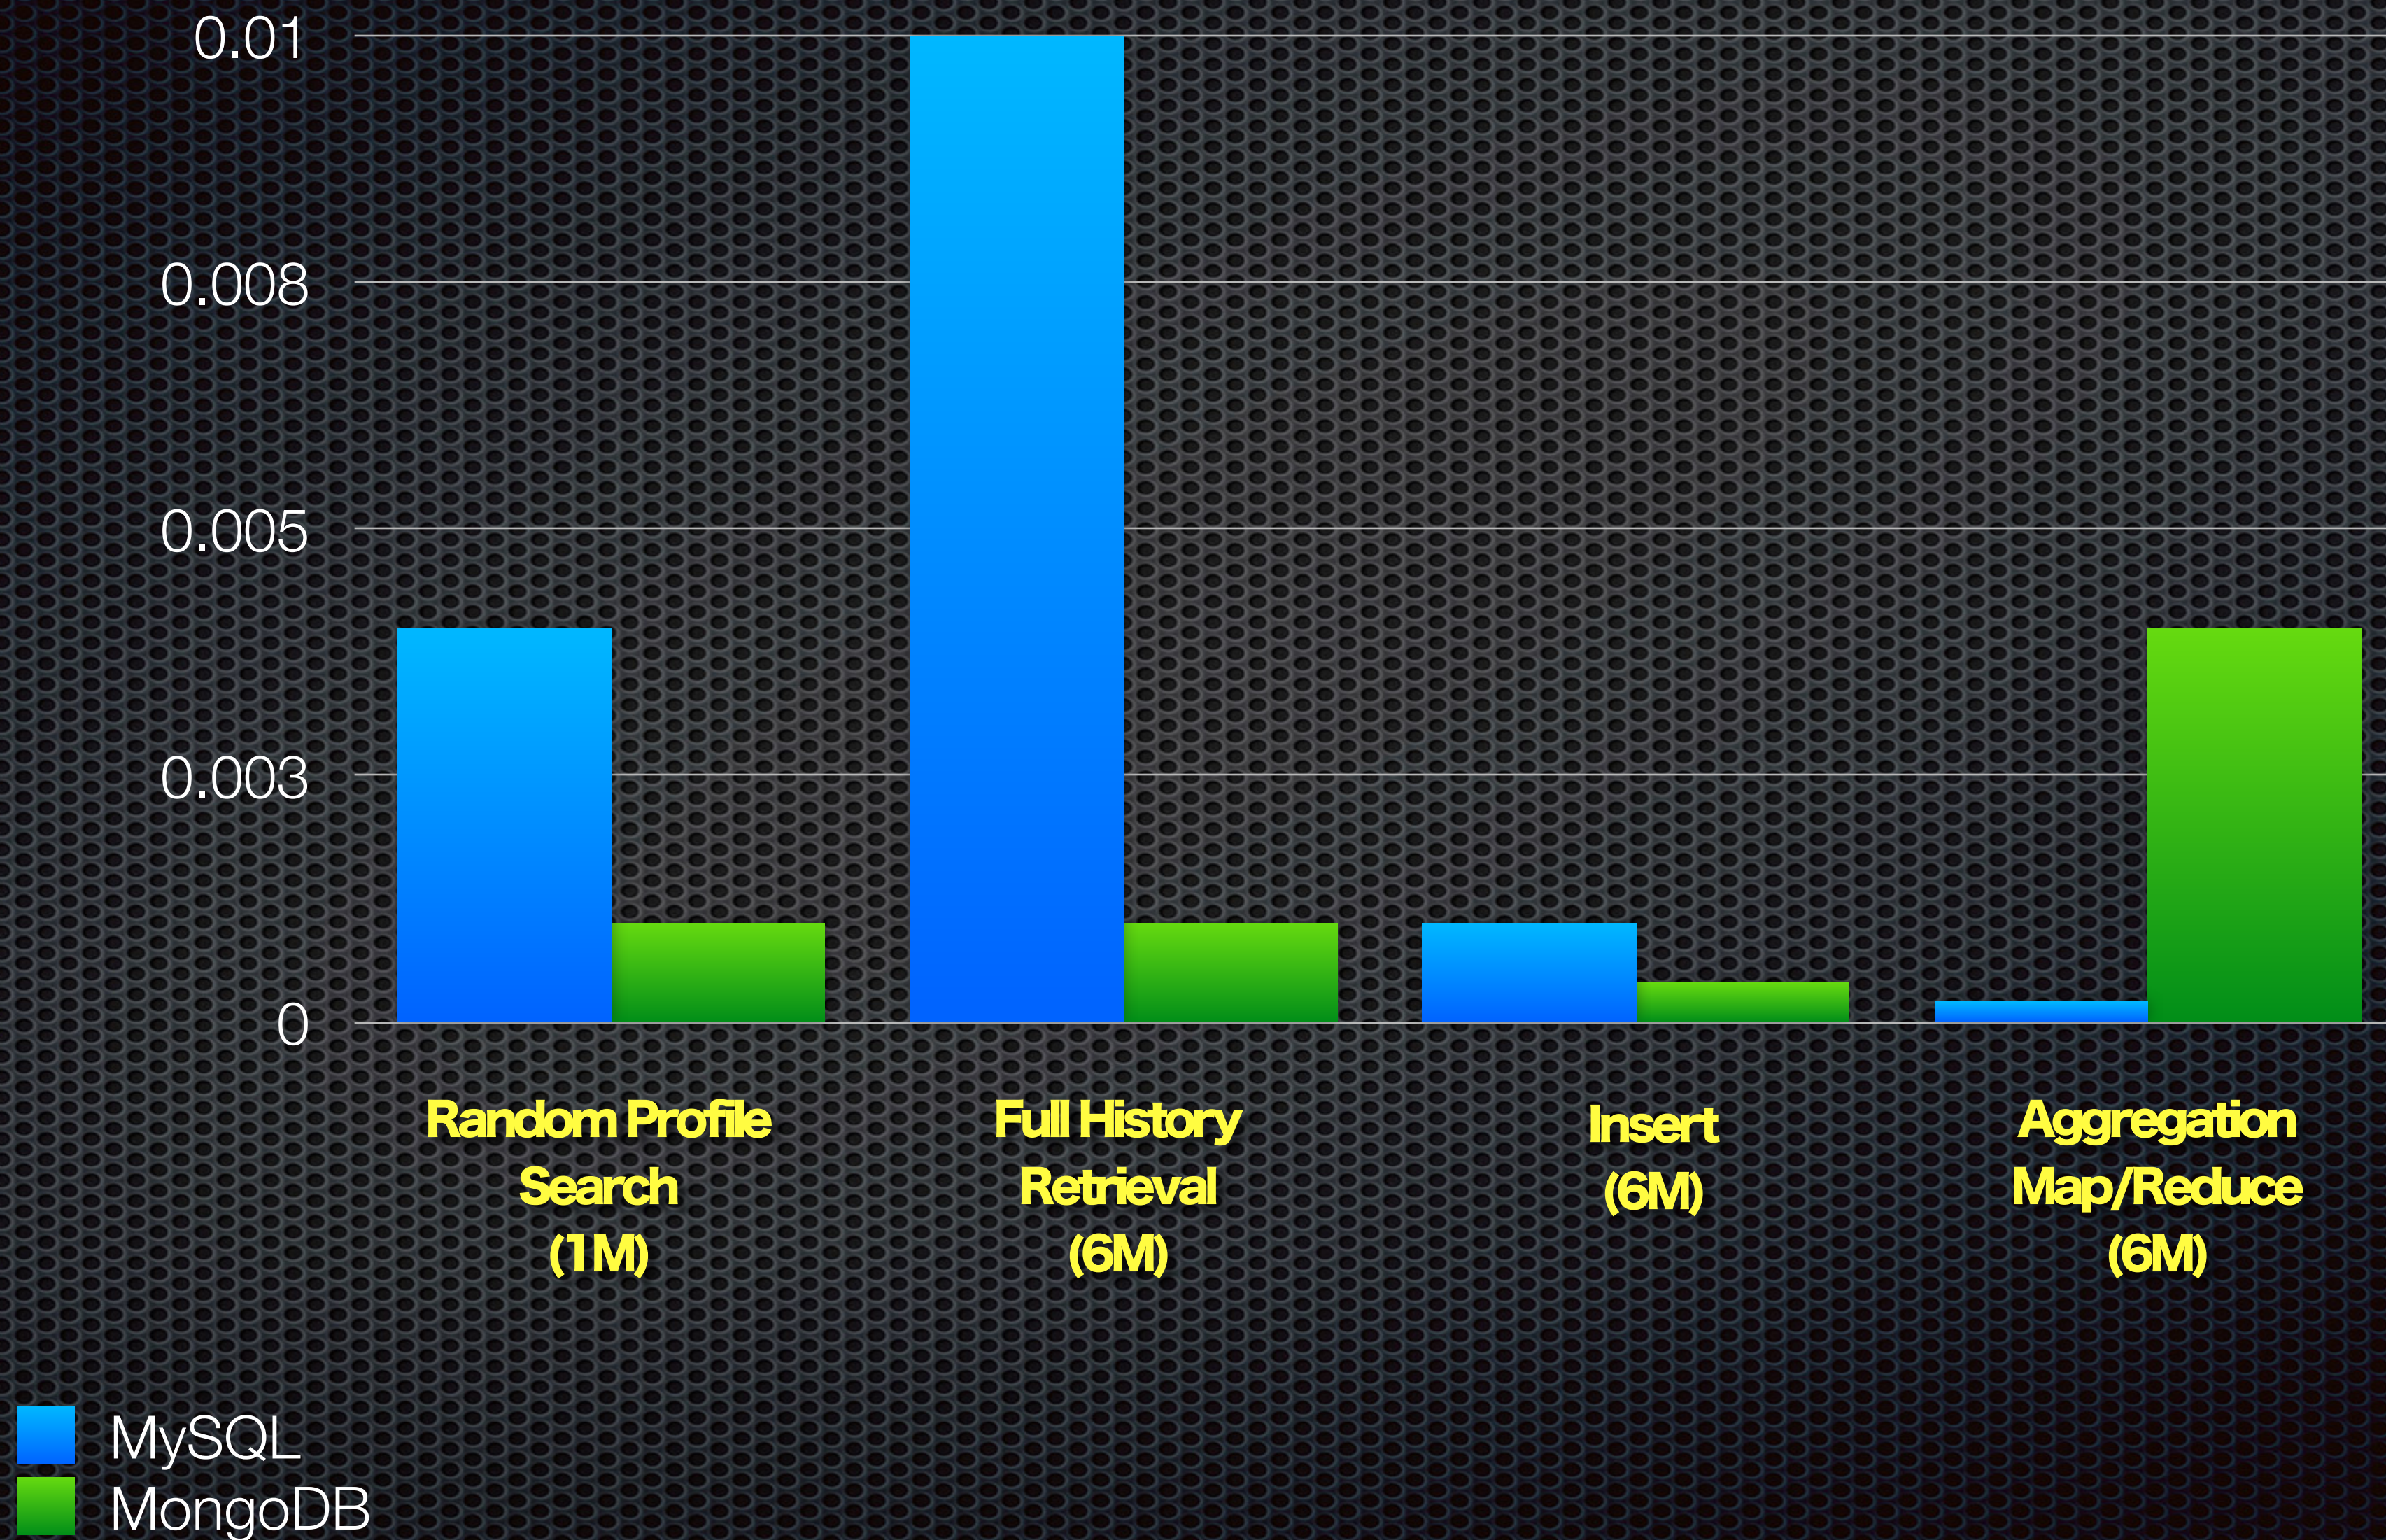




Schema-Less
Aggregated Data
Large Scale Data
Sharding
Map/Reduce
Memory Storage
Journaling



Denormalized data
Disk Space
Expertise
Complex Querying
Eventual Consistency
Resource Usage
DB-Level Locking



Polyglot Persistence




Profile (Total Sum)

Leaderboards

...

Activity Details
Activity Aggregations





Data Access Layer




The Challenge




Inconsistency

Data Access Layer

2-Phase Commit



Data Access Layer



**2-Phase
Commit**

Pending

Data Access Layer




**2-Phase
Commit**

Pending

Data Access Layer

**2-Phase
Commit**



Done

2

Consistency
vs.
Latency

3

Operations

4

Availability

Data Access Layer



Data Access Layer



Data Access Layer



5

Data Integrity

TEST

**{ Unit
Integration
Function }**

DATA Migration



Profile

3rd Party

Activity



Profile



3rd Party

Activity



1

One-Shot Migration

2

Profile Versioning

4
Hours!!!





RAW




RW



V2


Final Results

2022020314




9x response time

3 years uptime



Go Hybrid, wisely!



Go Hybrid, wisely!



Schema-Less
Large Scale Data
Sharding
Map/Reduce

Denormalized data
Eventual Consistency
Evolving



ACID
Transaction
Replication
Tools, OPs

Data Volume
Fixed Schema
Scaling

Thank
YOU!



@majidfn

<http://joind.in>

/13728