

Polyglot

Persistence

RED | **VENTURES**

Majid Fatemian

@majidfn





ASSASSIN'S
—CREED—™


FARCRY

TOM CLANCY'S
RAINBOW SIX | SIEGE

YOUR SHAPE




Relational



Non-Relational




Key-Value



Big Data

Application



Relational


Application



Relational




Relational



Non-Relational

Relational



DATA

BIG
DATA

Capture

Search

Storage



Analysis

Query

Privacy

Visualize

Non-Relational



①


Non-Relational


②

Cluster

③

Schemaless








2 x  **XBOX 360**


1 x **WiiU**


1 x  **XBOX ONE**












Application



Relational **Consistency**



ACID

Normalized

Profile

Player ID	Name	Location
1000	John	US
2000	Catherine	CA
3000	Paul	CA
4000	Kylee	US

Activity

Player ID	Map ID	Difficulty	Score
1000	223	E	5
2000	673	H	4.5
1000	451	M	4
2000	980	H	3

Profile

Player ID	Name	Location
1000	John	US
2000	Catherine	CA
3000	Paul	CA
4000	Kylee	US

Activity

Player ID	Map ID	Difficulty	Score
1000	223	E	5
2000	673	H	4.5
1000	451	M	4
2000	980	H	3


Profile

Player ID	Name	Location
1000	John	US
2000	Catherine	CA
3000	Paul	CA
4000	Kylee	US

Activity

Player ID	Map ID	Difficulty	Score
1000	223	E	5
2000	673	H	4.5
1000	451	M	4
2000	980	H	3

Relational **Availability**




Replication


Fault Tolerance



Relational **DevOps**




OPs / DBAs
Tools
Devs










Relational
Scalability

Data Access Layer



Relational
Distribution



```
{  
  player_id: 1234,  
  level: 12,  
  duration: 120,  
  activity: activity  
}
```



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


Relational Schema

**Total
stats**

Players' Profiles

**Categorized
Aggregated
stats**


**Individual Activities
Detailed
stats**




ACID
Transaction
Partitioning
Replication
Tools
OPs



Data Volume
Fixed Schema
Scale Up
Manual Sharding



NoSQL



...

mongoDB



①

OPs

②

Devs

③

Community

Non-Relational Schema-less



player_id: 1234,
level: 75,
duration: 120,
activity: activity




player_id: 1234,
* level: 40,
* duration: 75,
* activity: activity,
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



Non-Relational Map/Reduce



Non-Relational Map/Reduce



Non-Relational Map/Reduce

Map



K,V1



K,V2



K,V3

Non-Relational Map/Reduce

Map



K,V1



K,V2




K,V3

Non-Relational Map/Reduce

Map

Reduce



Result

read / **write**
Disk Space



Non-Relational
Denormalized

Profile

Player ID	Name	Location
1000	John	US
2000	Catherine	CA
3000	Paul	CA
4000	Kylee	US

Activity

Player ID	Map ID	Difficulty	Score
1000	223	E	5
2000	673	H	4.5
1000	451	M	4
2000	980	H	3

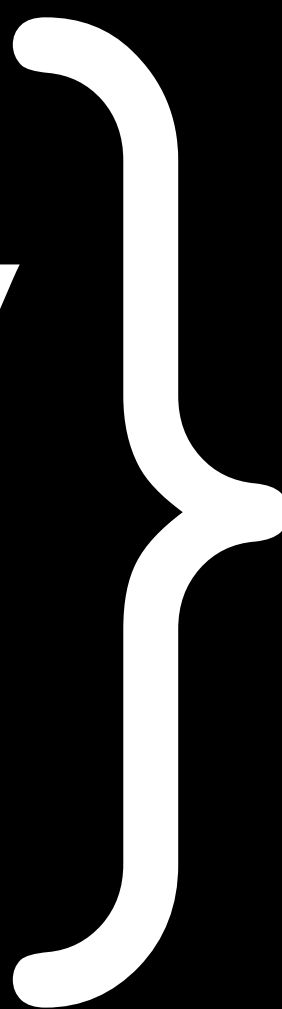


```
{Name: John, game:1, activity: Mission01, Score: 1245}
```


```
{Name: John, game:1, activity: Mission01, Score: 1245}
```

```
{Name: John, game:3, activity: Mission04, score: 5431}
```


Query
Aggregation




Non-Relational
Data Model




Non-Relational **Eventual-** **Consistency**




Non-Relational
Eventual-
Consistency



Non-Relational
Eventual-
Consistency



Non-Relational
Eventual-
Consistency




Non-Relational **Balancer**



+



Non-Relational
Balancer



Non-Relational Balancer



Non-Relational **Resilience**



Non-Relational **Resilience**

DB Level (2.2)

Document Level (3.x)



Non-Relational
Locking


Point-in-time 




Non-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



Relational Data

**Players' Profile
Leaderboards**

Activity Details

**Schema-less
Sharding**









Total Sum

Activity 1
⋮
Activity n
⋮
Activity m







Application



Data Access Layer



The Challenge



Inconsistency

Data Access Layer

2-Phase Commit



Data Access Layer



Pending

2-Phase
Commit

Data Access Layer




Pending

2-Phase
Commit

Data Access Layer

2-Phase Commit



Done

②

Latency

2

Latency
vs. Consistency

3

Operations

③

Operations
Exponential complexity


4

Availability


4

Availability
Plan for failure

Data Access Layer



Data Access Layer



Data Access Layer




5

Data Integrity

5

Data Integrity
Backward compatibility

TEST





Unit

Integration

Functional



Data Migration




Profile

3rd Party



History



```
{  
  profile:  
  {  
    GUID:...,  
    platform:  
    {  
      name: 'xbox',  
      xuid:...  
    }  
  },  
  activity:...,  
  duration:...,  
  datetime:...  
}
```

1

One-Shot Migration

①

One-Shot Migration

4 hours !!


① One-Shot Migration
4 hours !!


② Profile versioning
Live !!





V1






V1






V2







Final Results

Scalability



2011




2012



2013




2014




9x response time

3 *years* uptime



Go Hybrid!



Go Hybrid!
Thoughtfully

Thank
YOU!

 @majidfn

References

NoSQL Distilled

Polyglot Persistence