

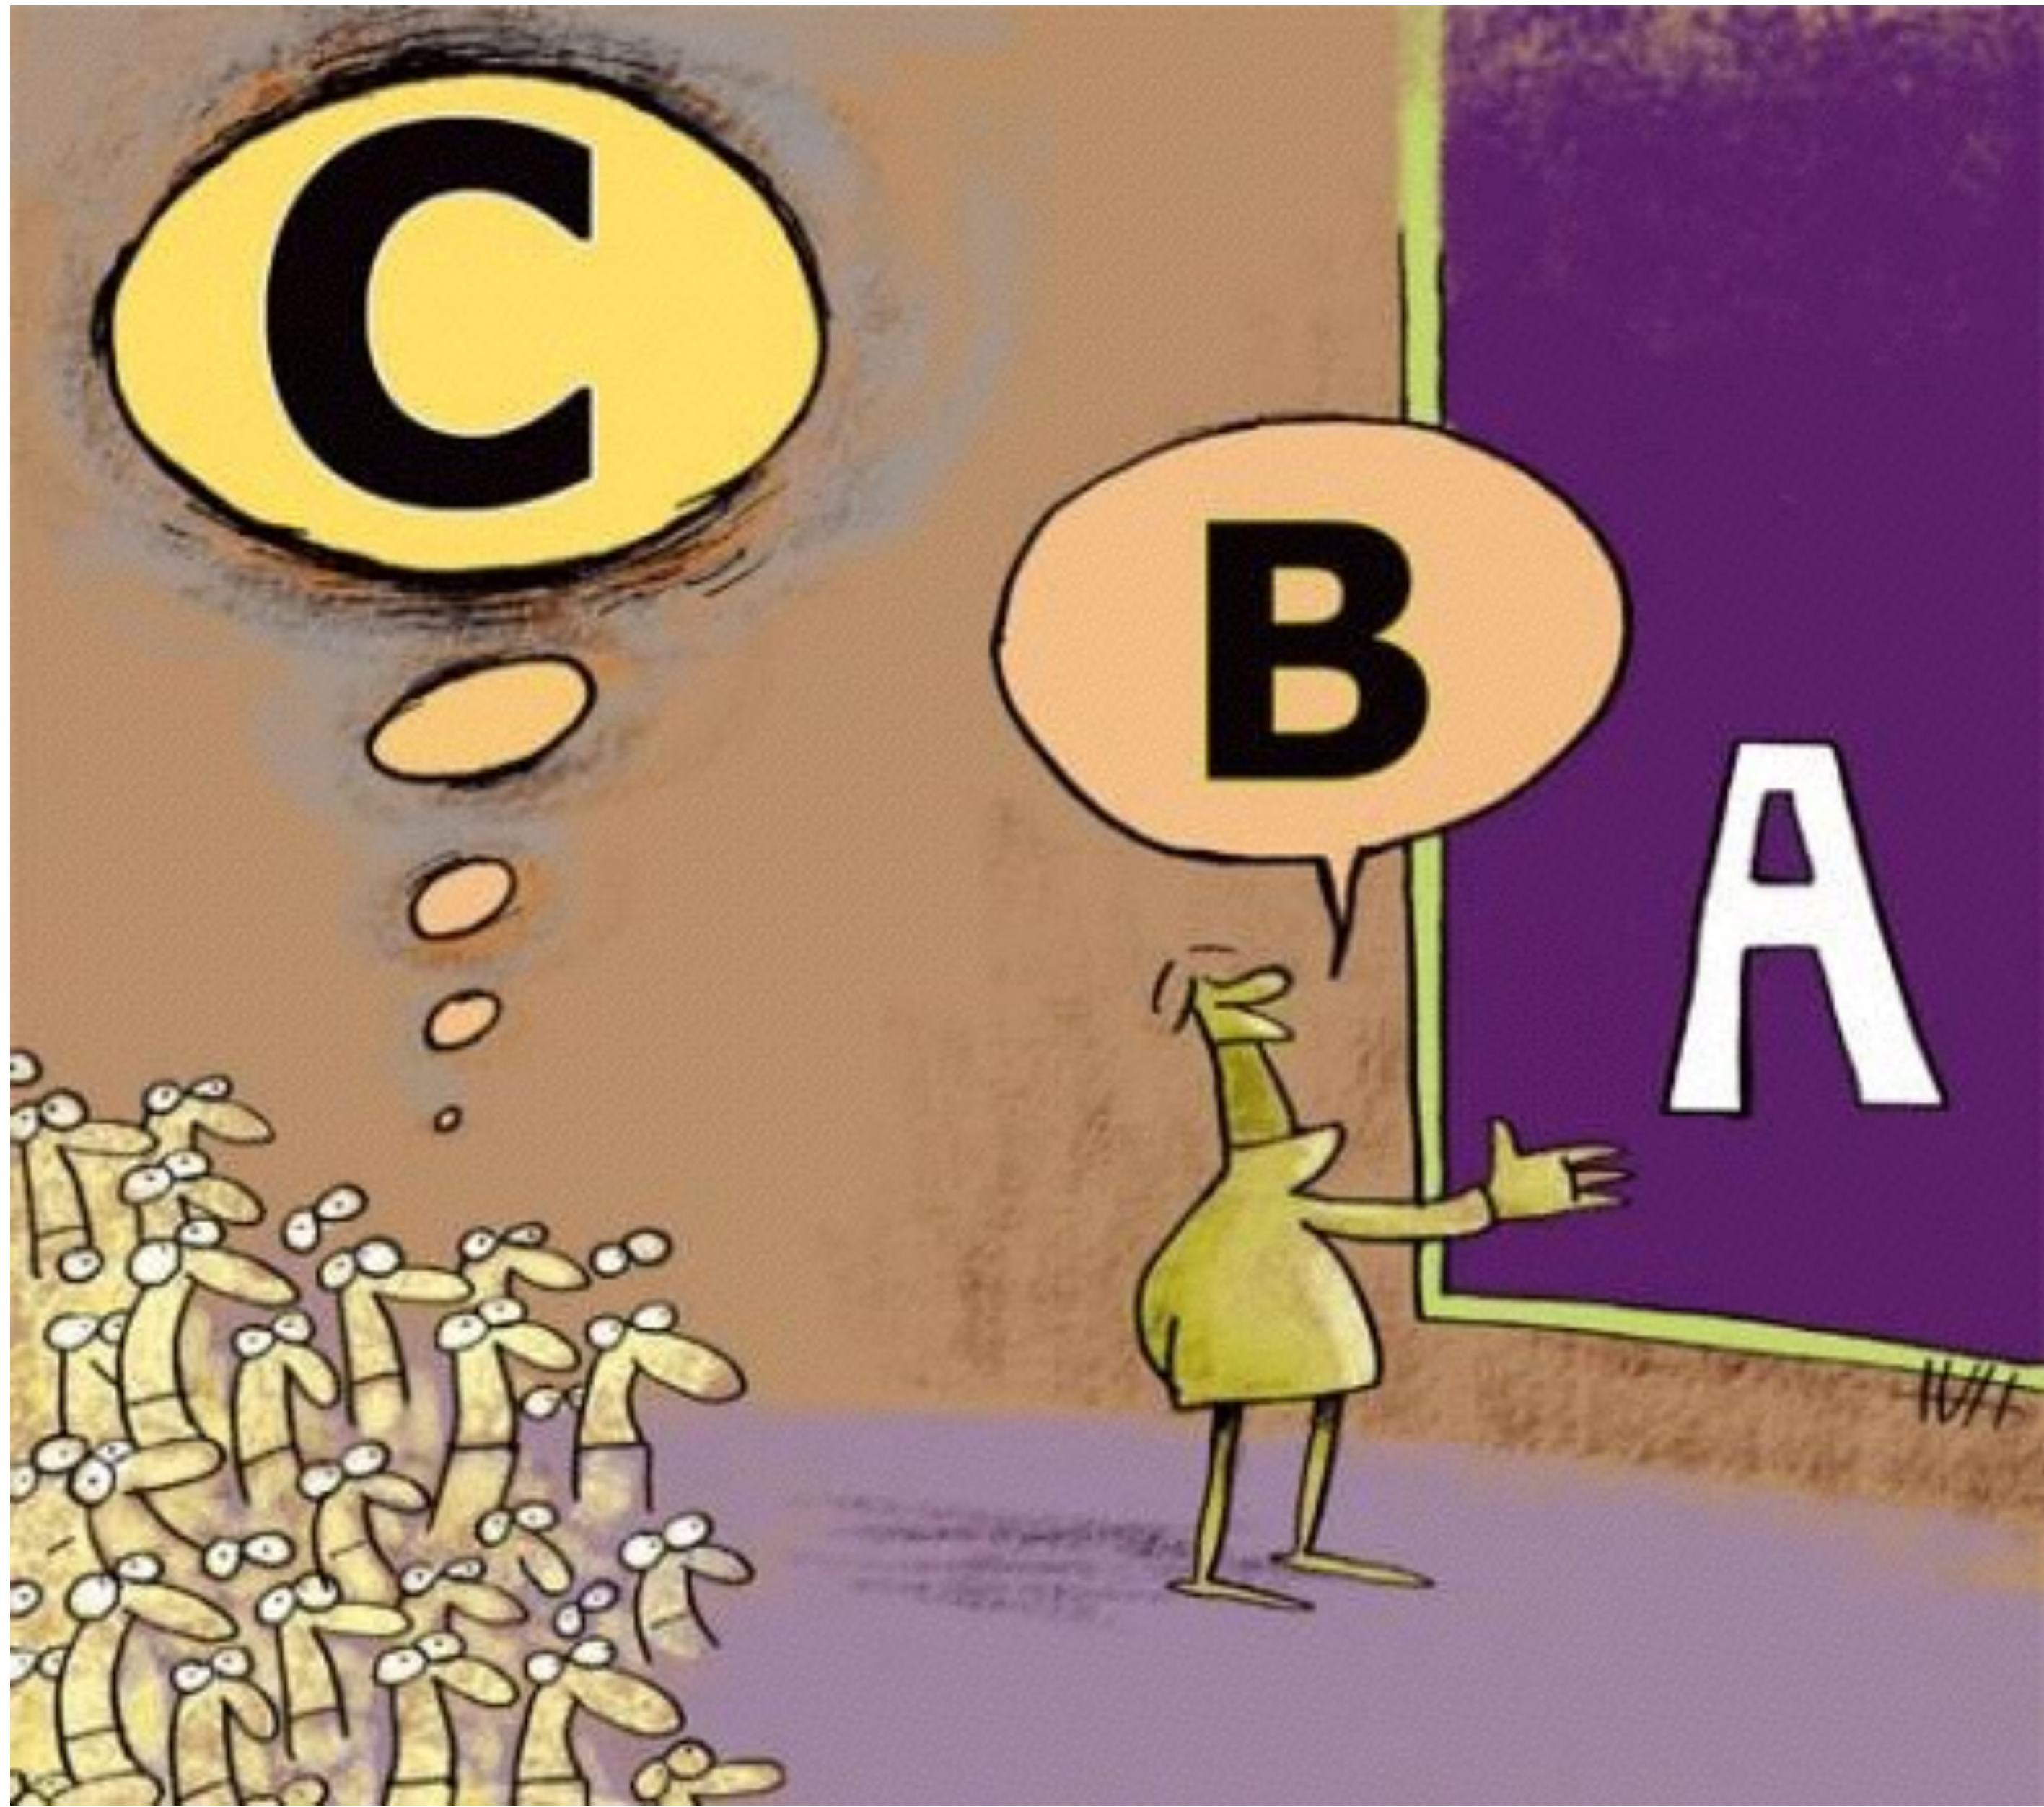






“Business group showing diversity in a meeting”





Z+6N&

hb:40:JFH88>

g

..H\$88Z||EUVZ

M\$:d:
Pff+a
U

UFZH|:}tgr~

\$H+2
L L ?

HfNX:Y6i50+
]A&05:++&

XZS4&f^208G

hTg7y08&TF7
YFEt[,49DD2v

@C=ttt
MXttt
Ott888tt0

z
s##zJKJd&y
|

q:,@00xx

v:

H96FELN81.

bom

/Oig0-(yjo
OvwLOZdh#u

w:|1k4

vY|F|F|

<64/NMFDW|

pocb6e^axM

vK^

oLFHOXPB

H9/X8/vR<

-OttMtt

v^PEP^vLVA^
70-88PMLN^

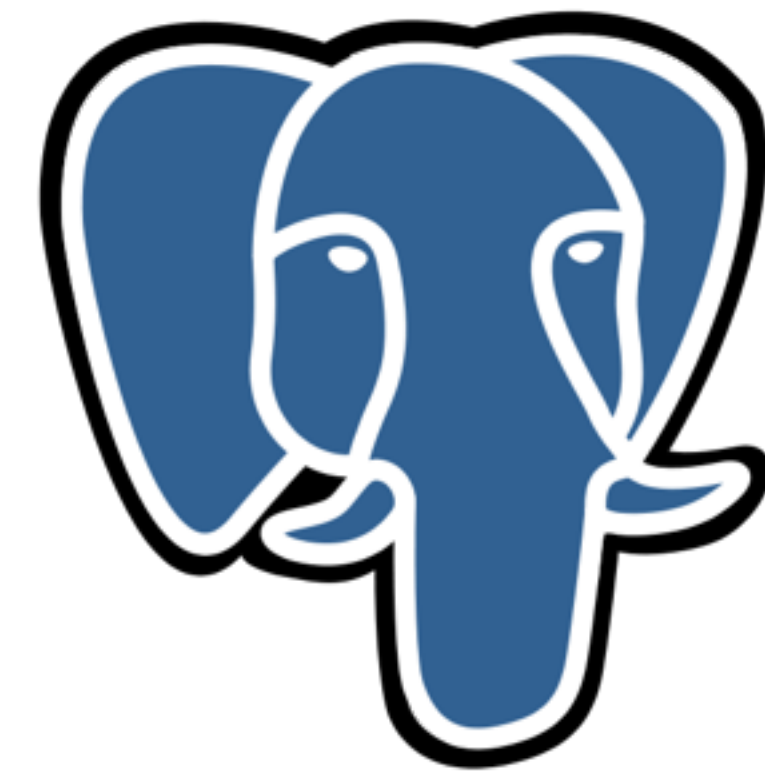
vO

00HCOX.DF<g

F







PostgreSQL











Mono culture = clean

Diversity = messy

Embracing Diversity in Databases



Frank @Lyaruu

CTO Dexels

Technology Hipster

Amsterdam



How can we reduce the cost of
diversity in software?

Software BS-Bingo

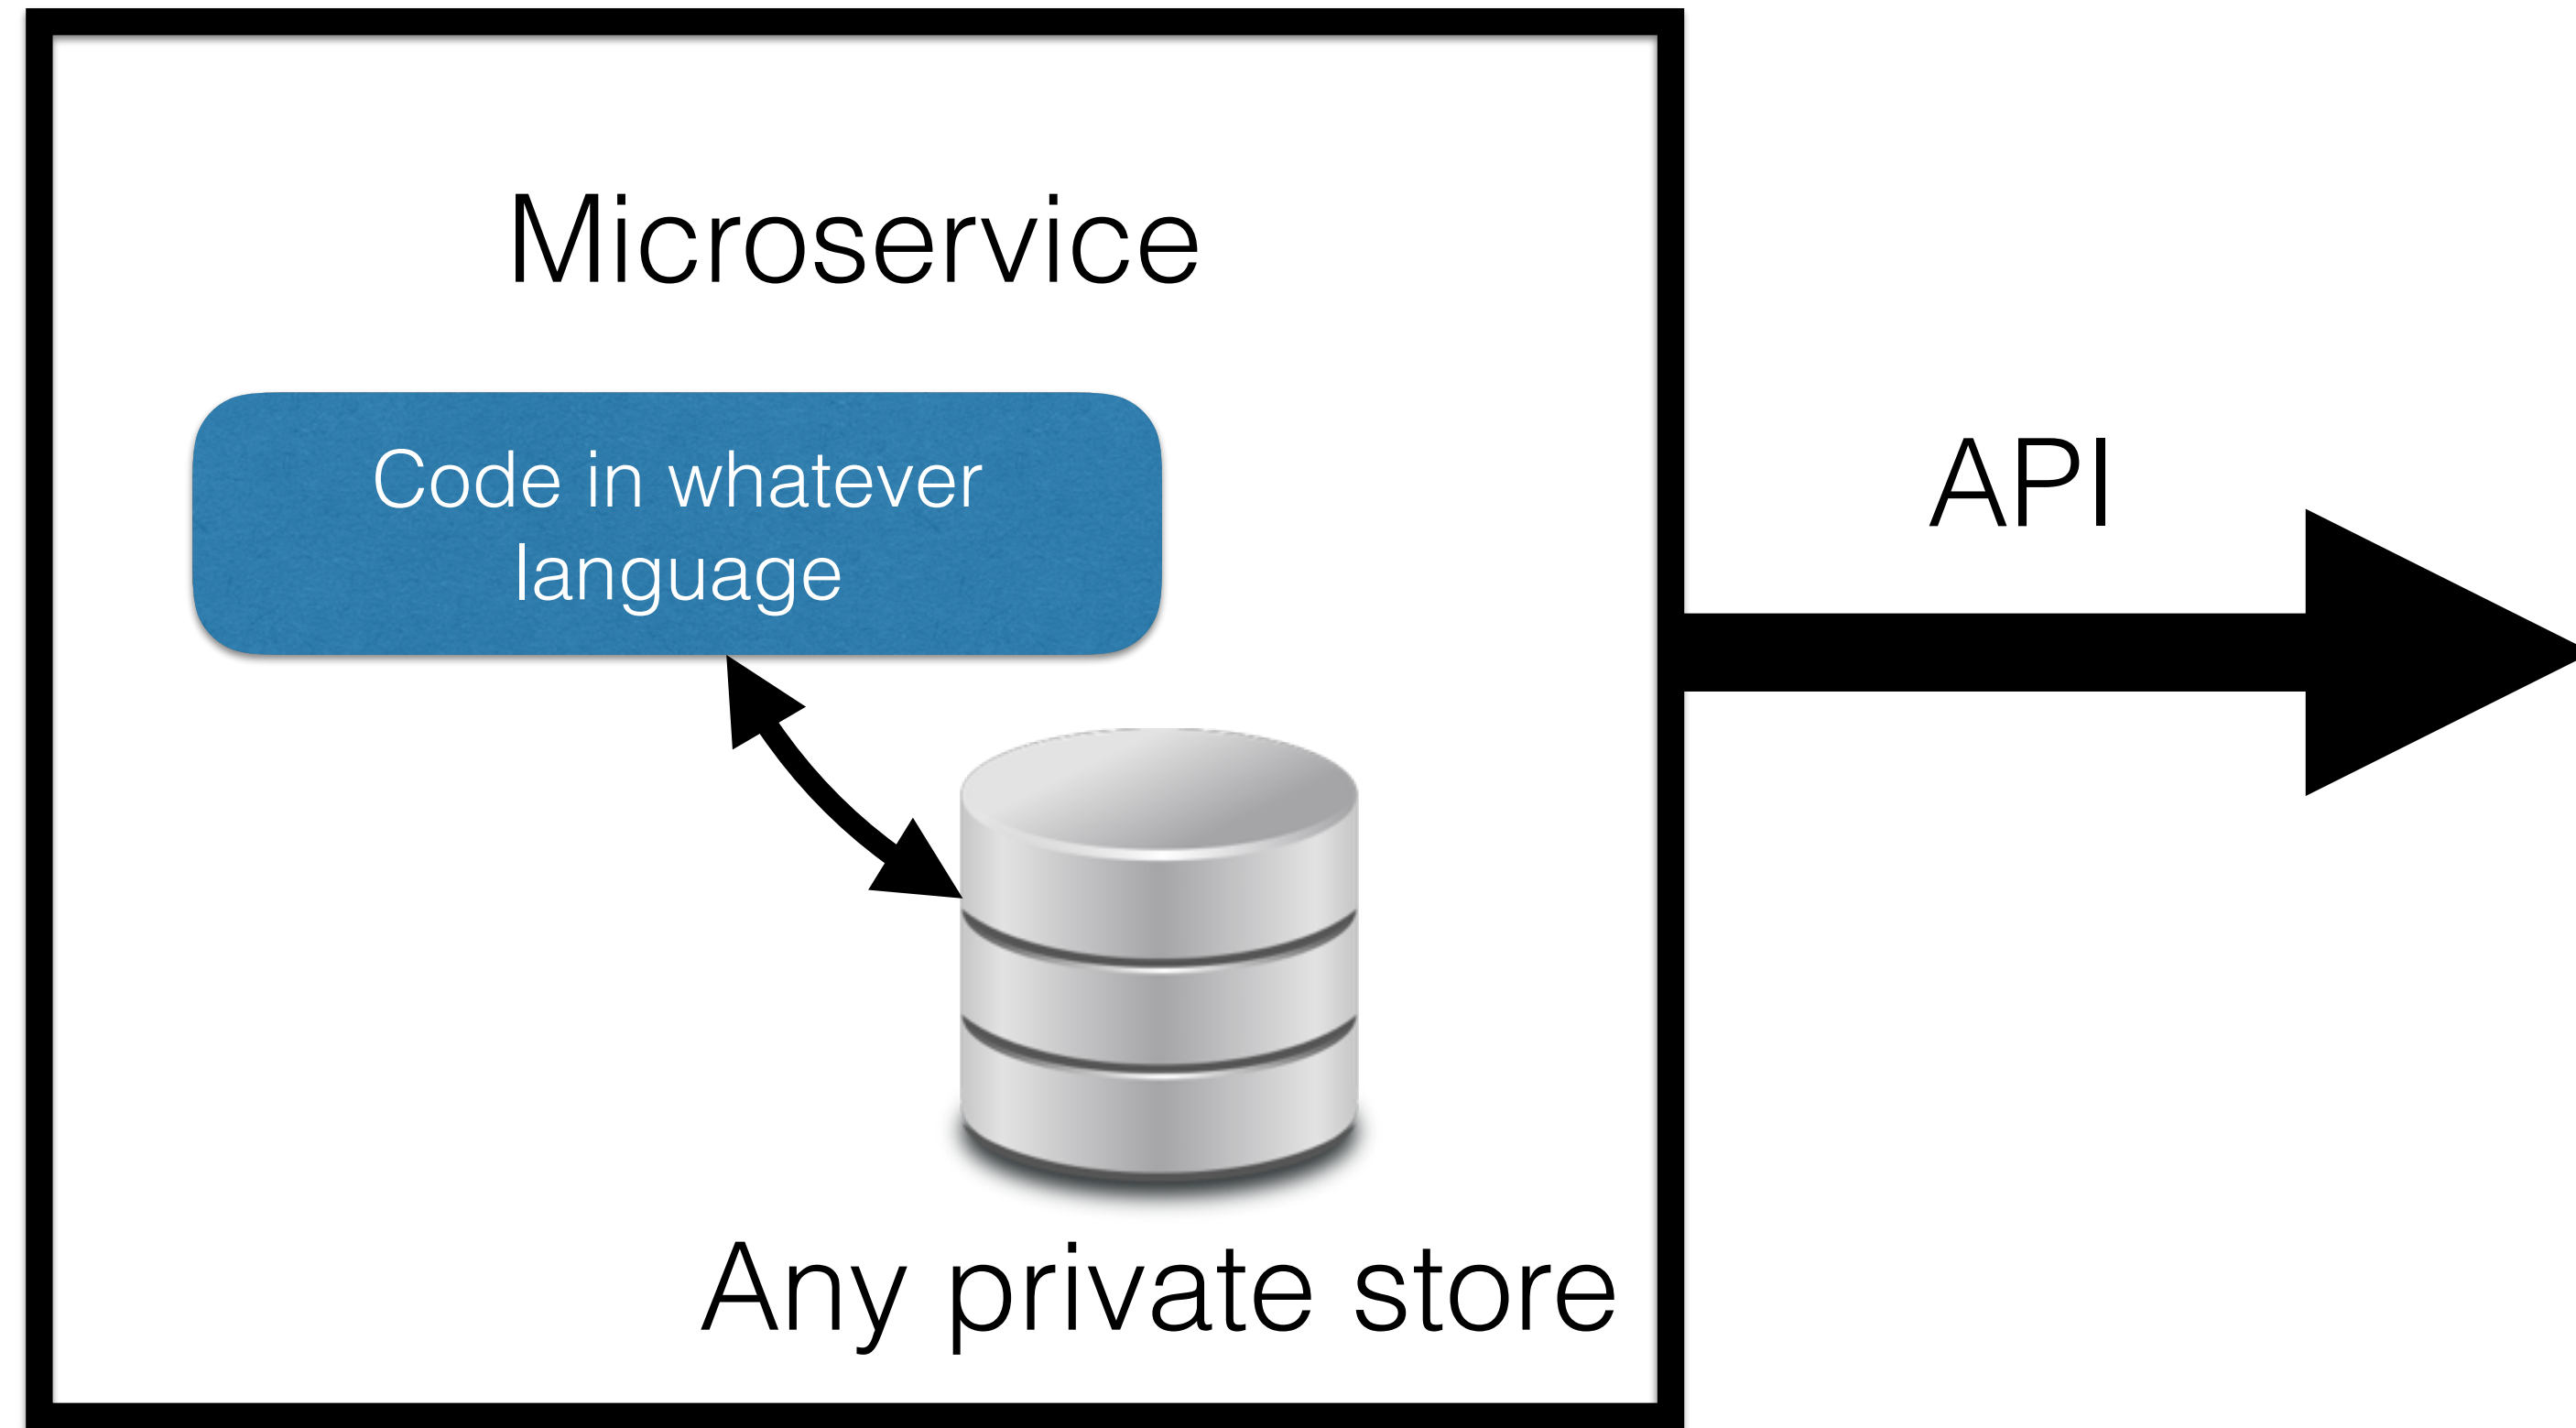
kubernetes	docker	polyg
docker	micro service	streaming
docker	cloud native	reactive



API



How can we reduce the cost of
diversity in ~~software~~?
databases?



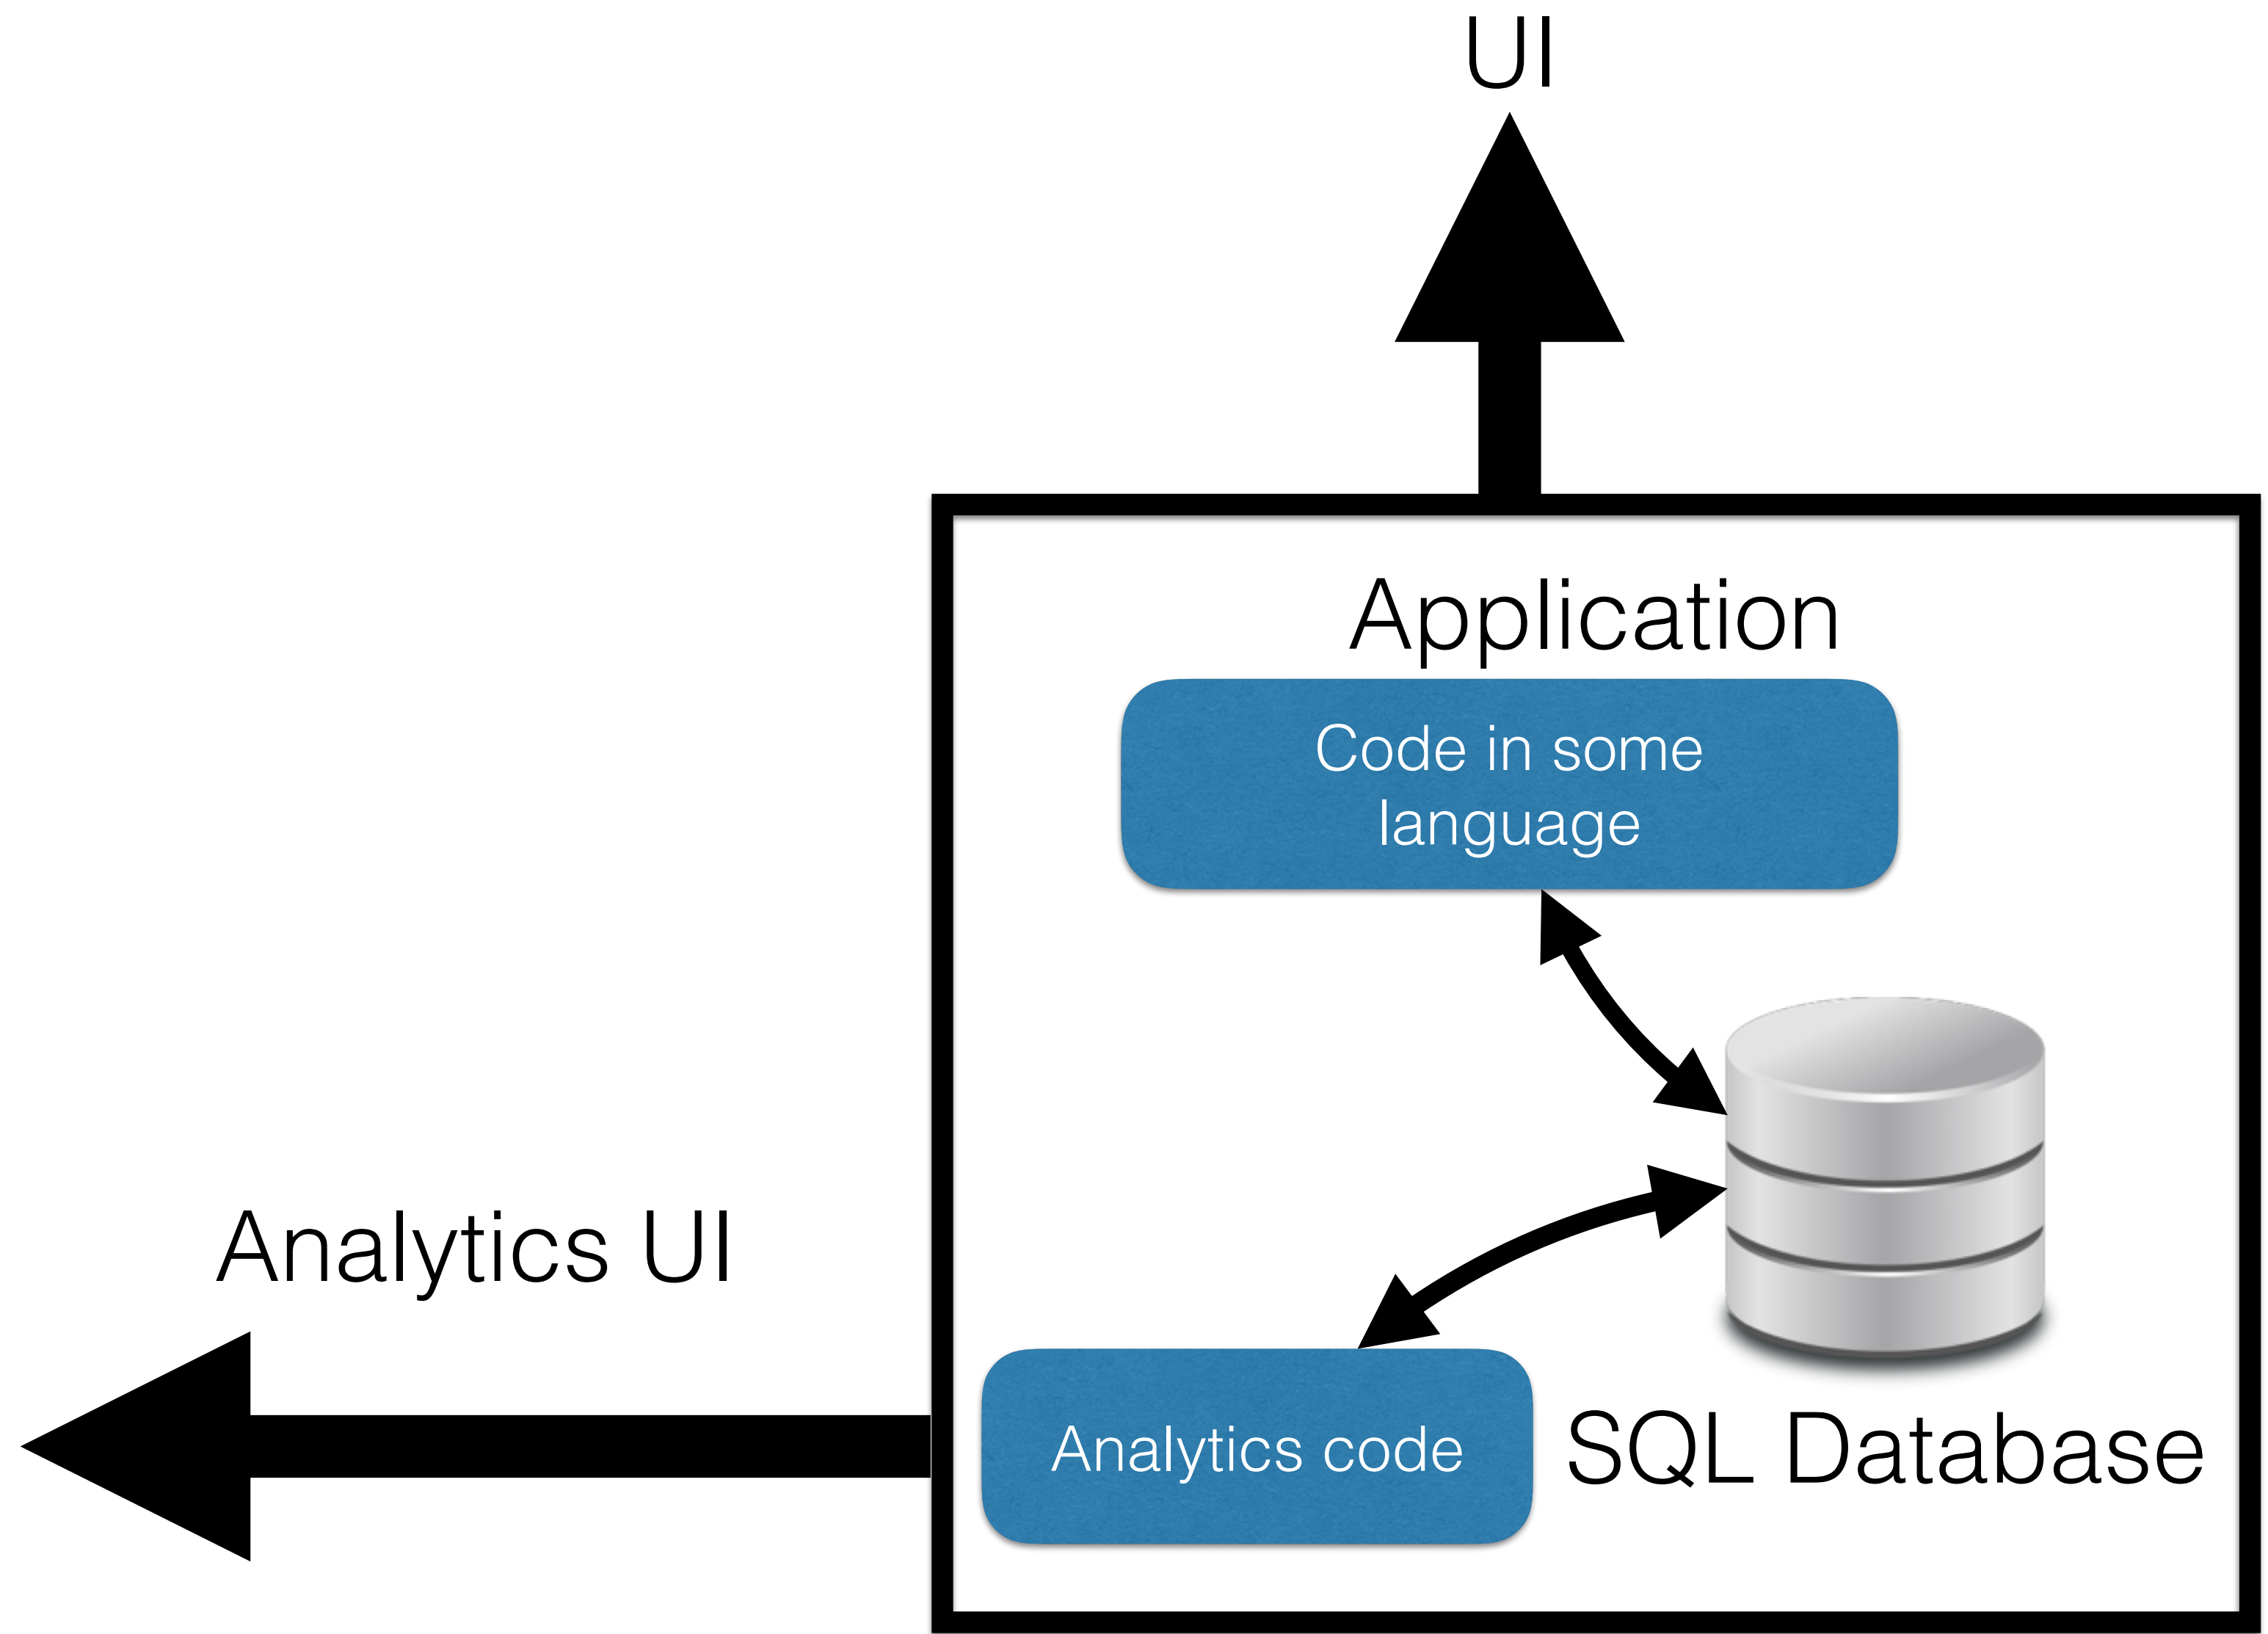


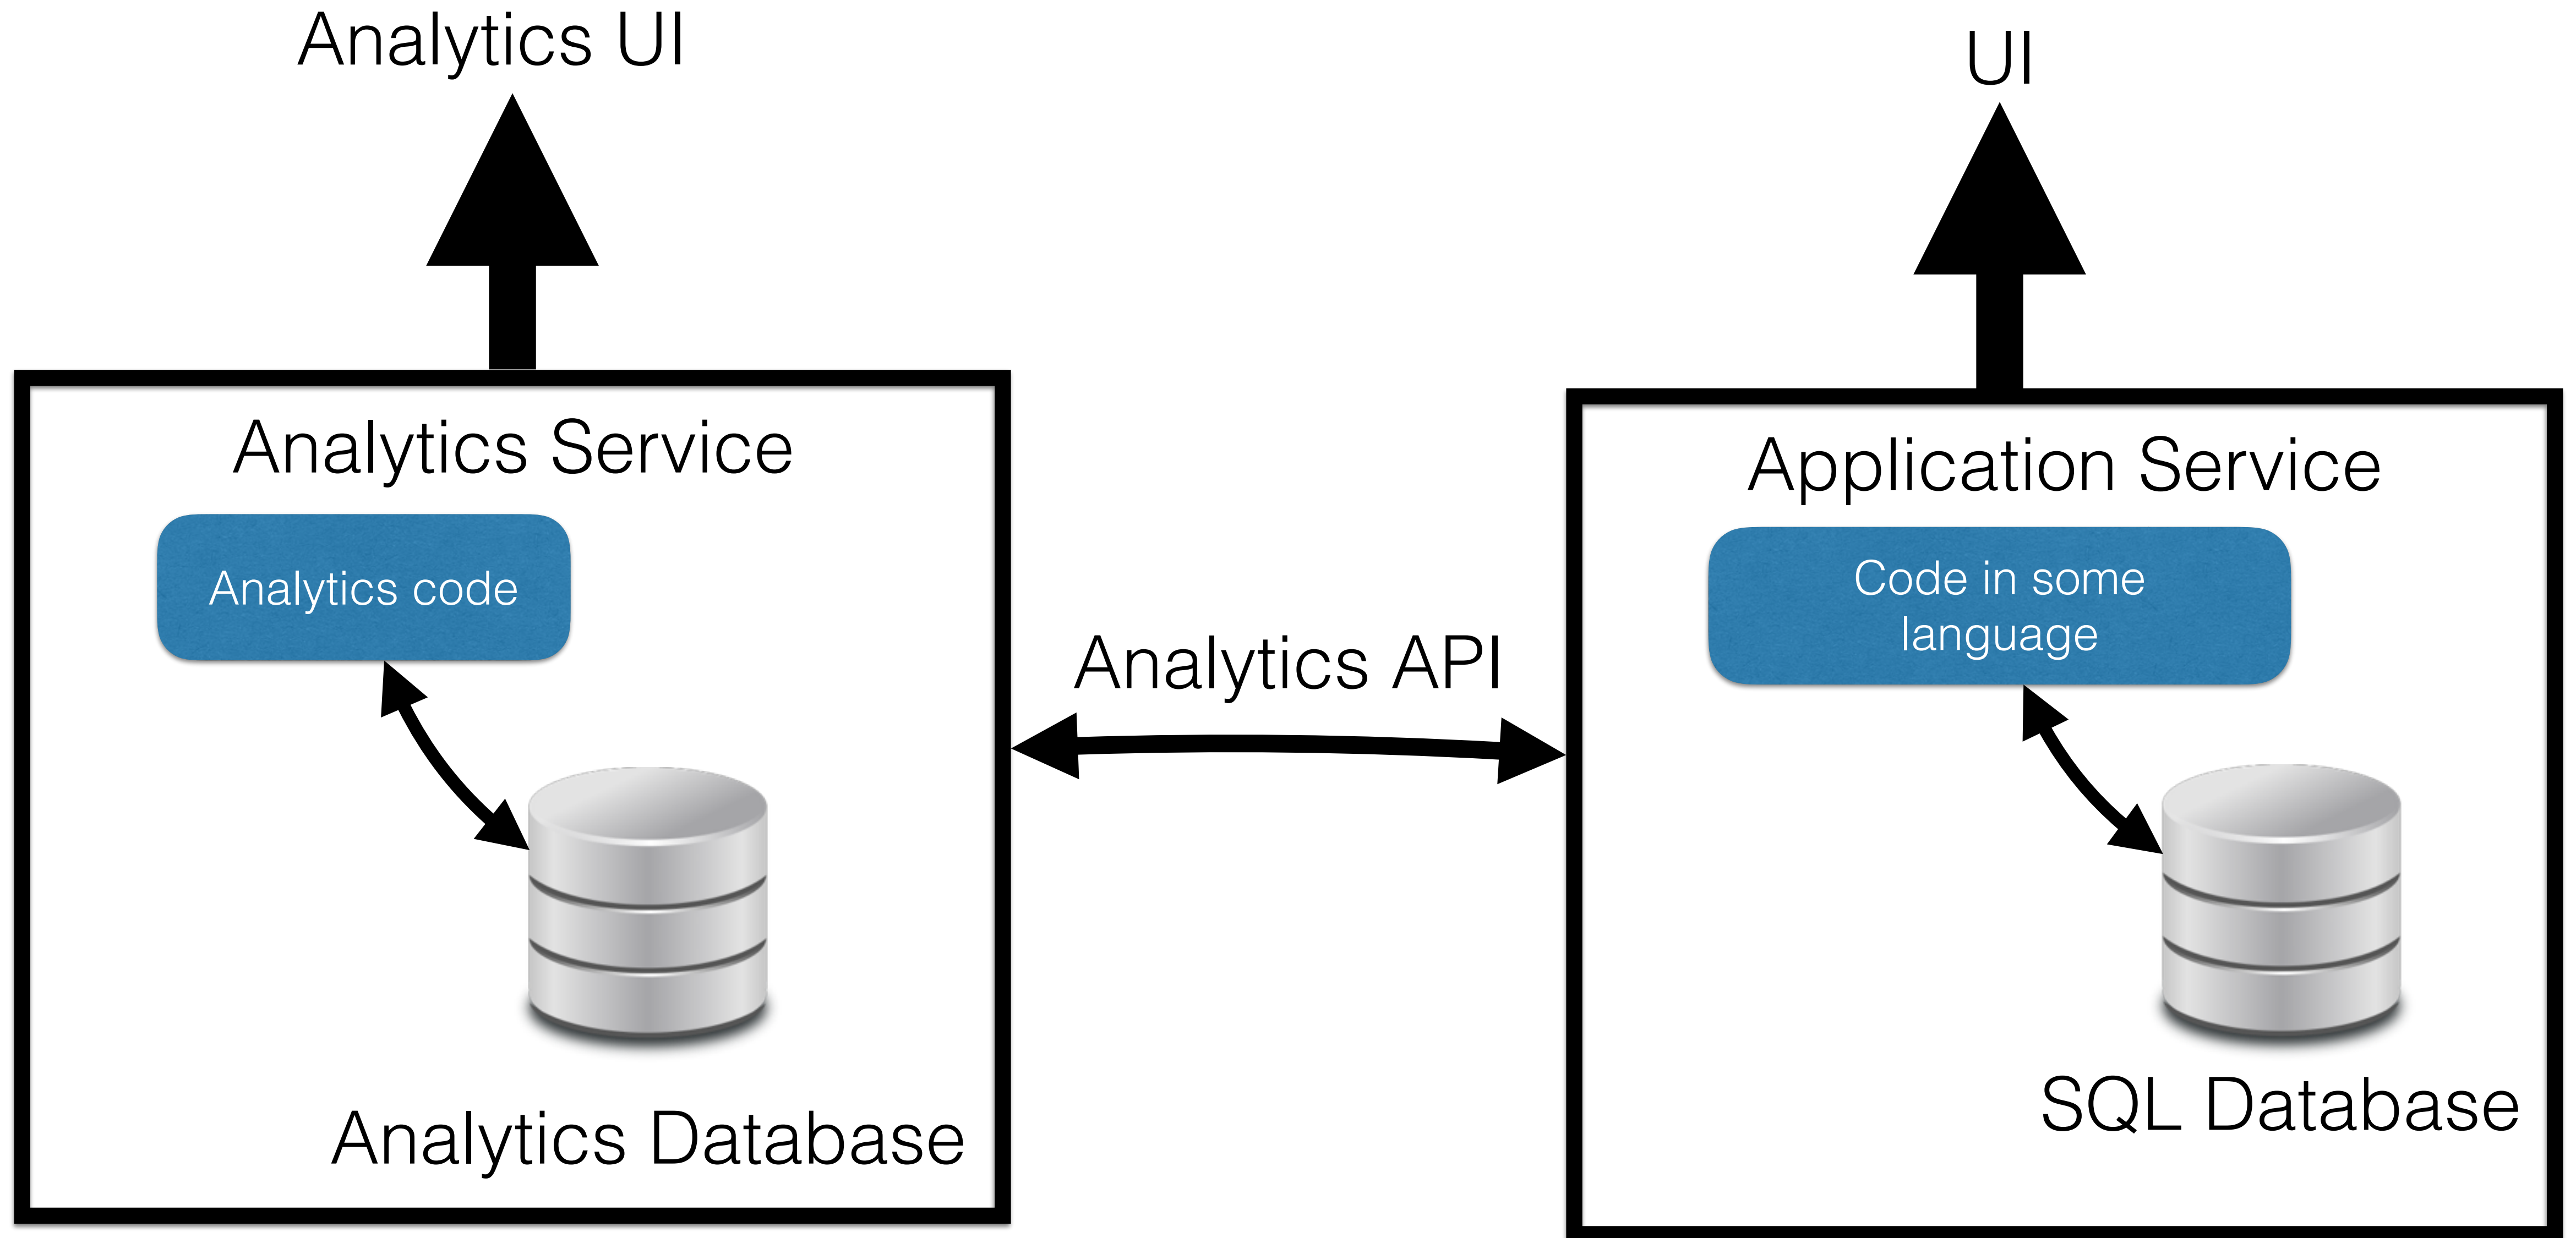
MISSION ACCOMPLISHED

NOT EXACTLY

Different parts need the same
data

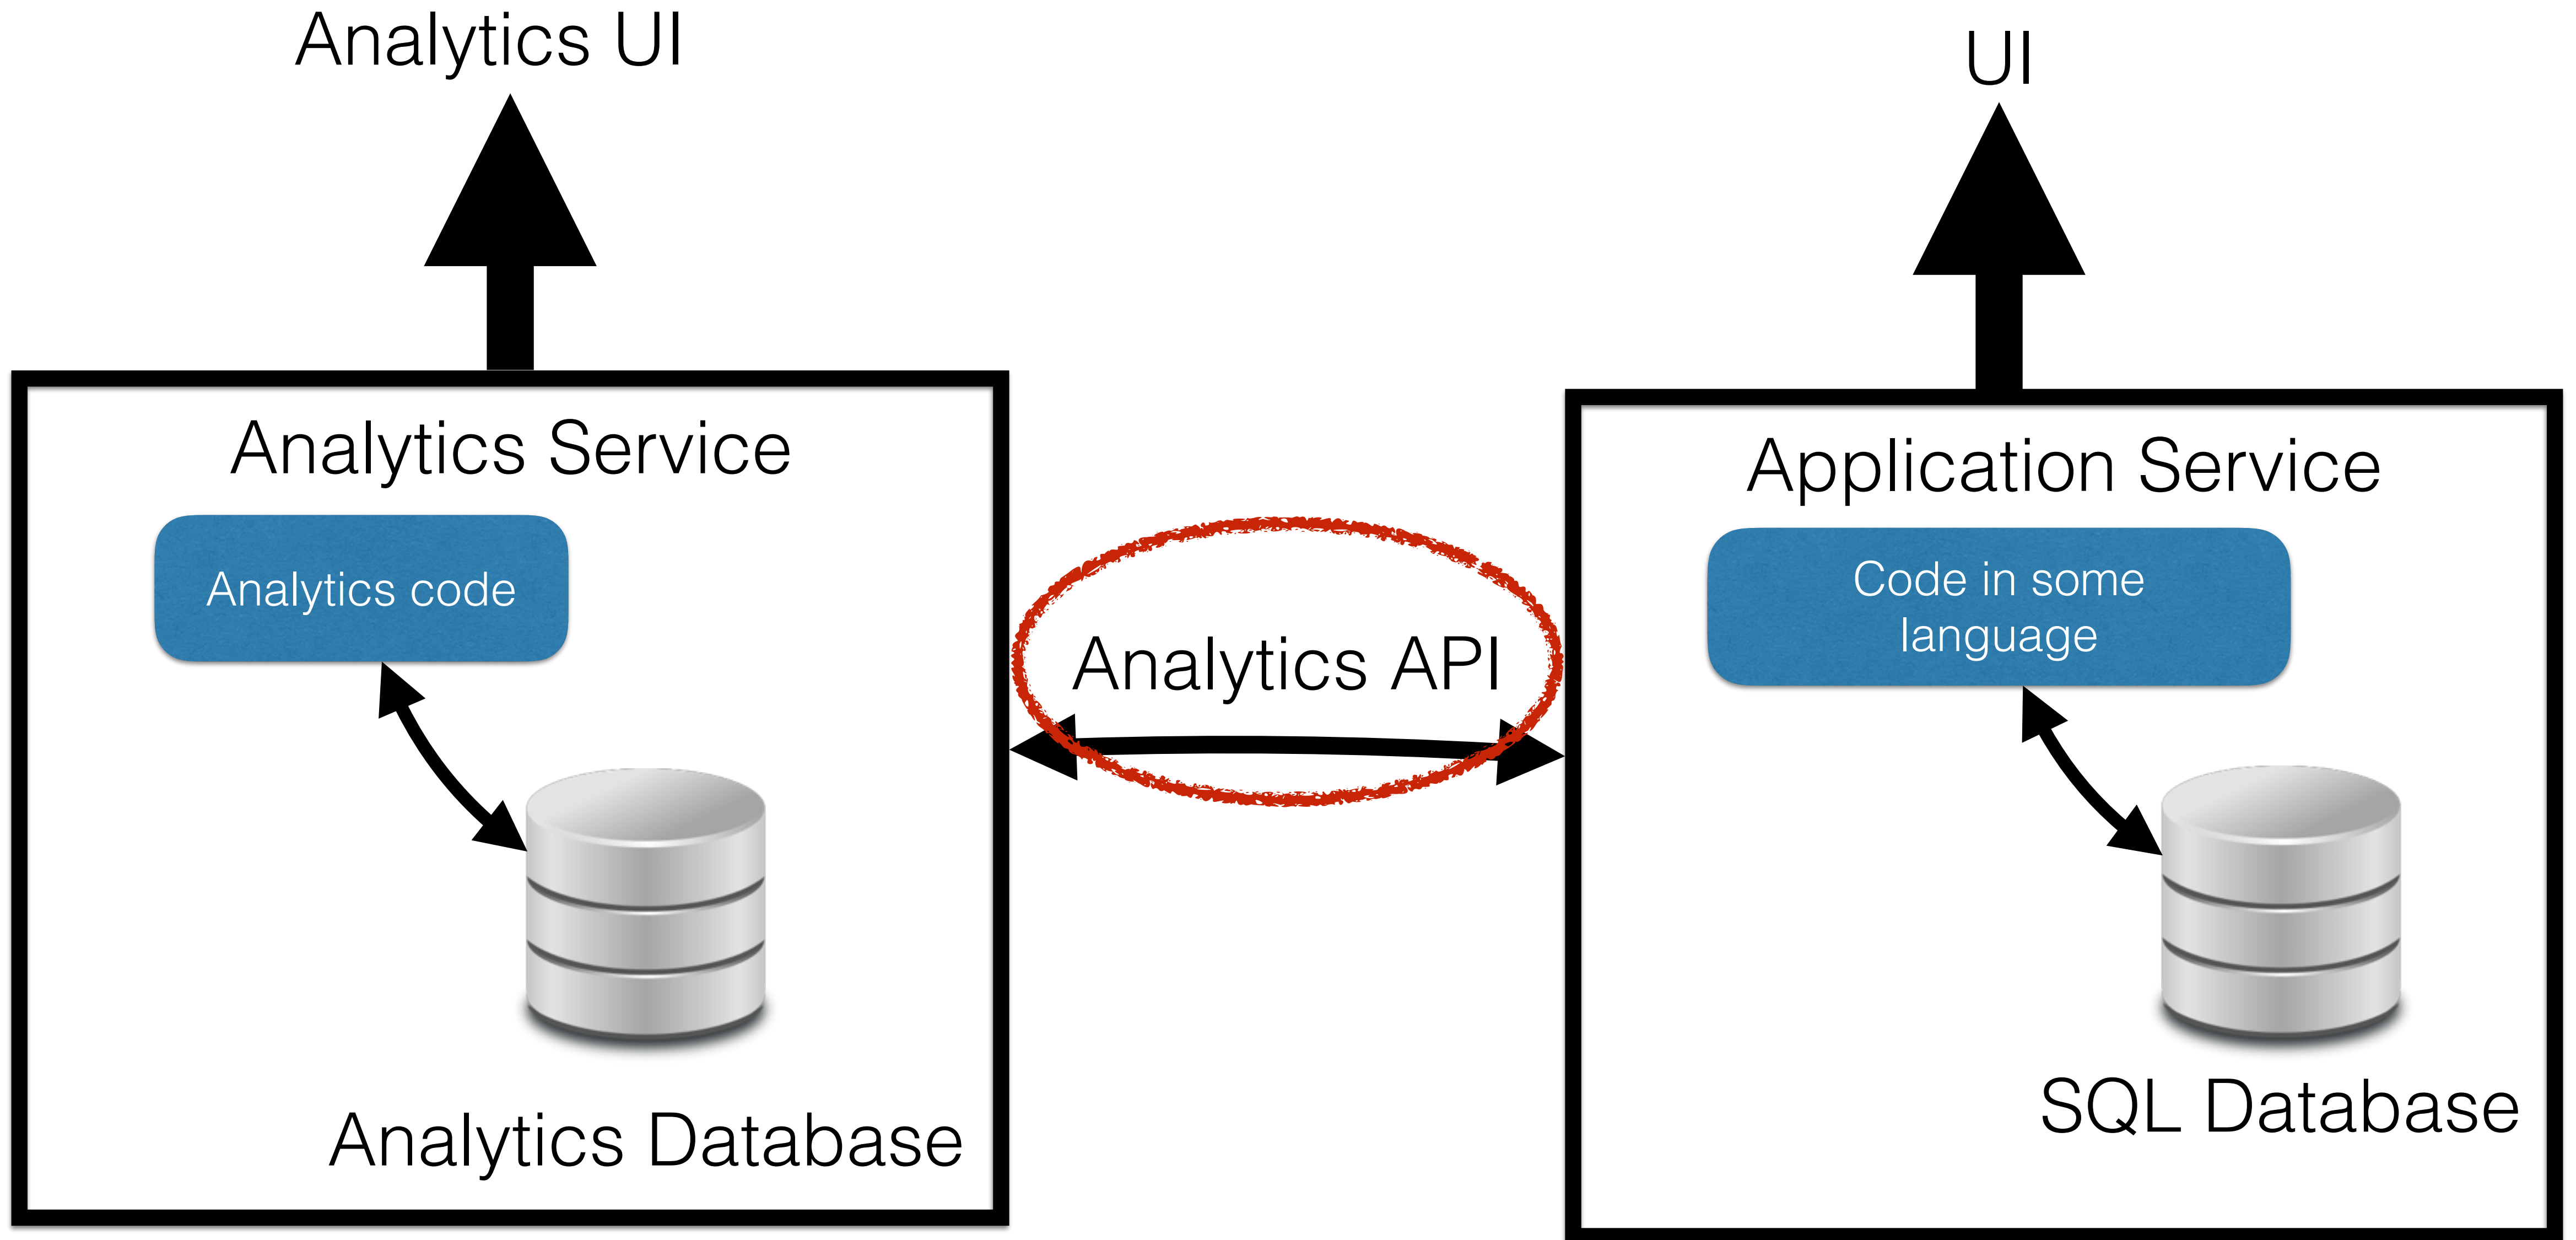
... but in a different way





That was easy!







What does that even mean?

How would that work?

GET /give-me-everything
?

GET /get-all-personids

+

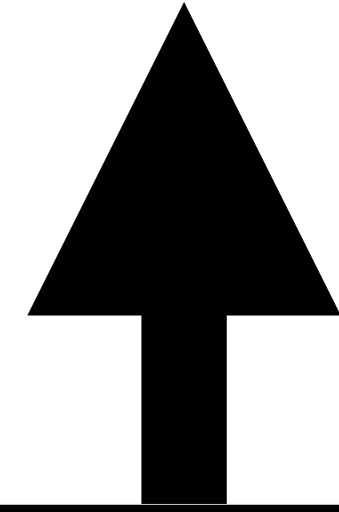
GET /person?id=123

?

GET /data?query="SELECT AVG(age) FROM PERSON"

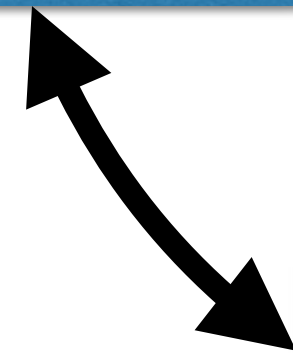


Analytics UI



Analytics Service

Analytics code

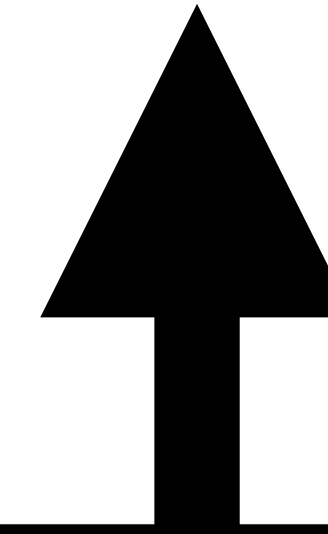


SQL Database

Magic* Replication Tool

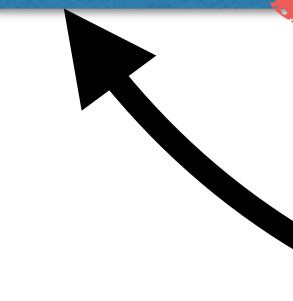


UI



Application Service

Code in some language



SQL Database

*Expensive

Issues

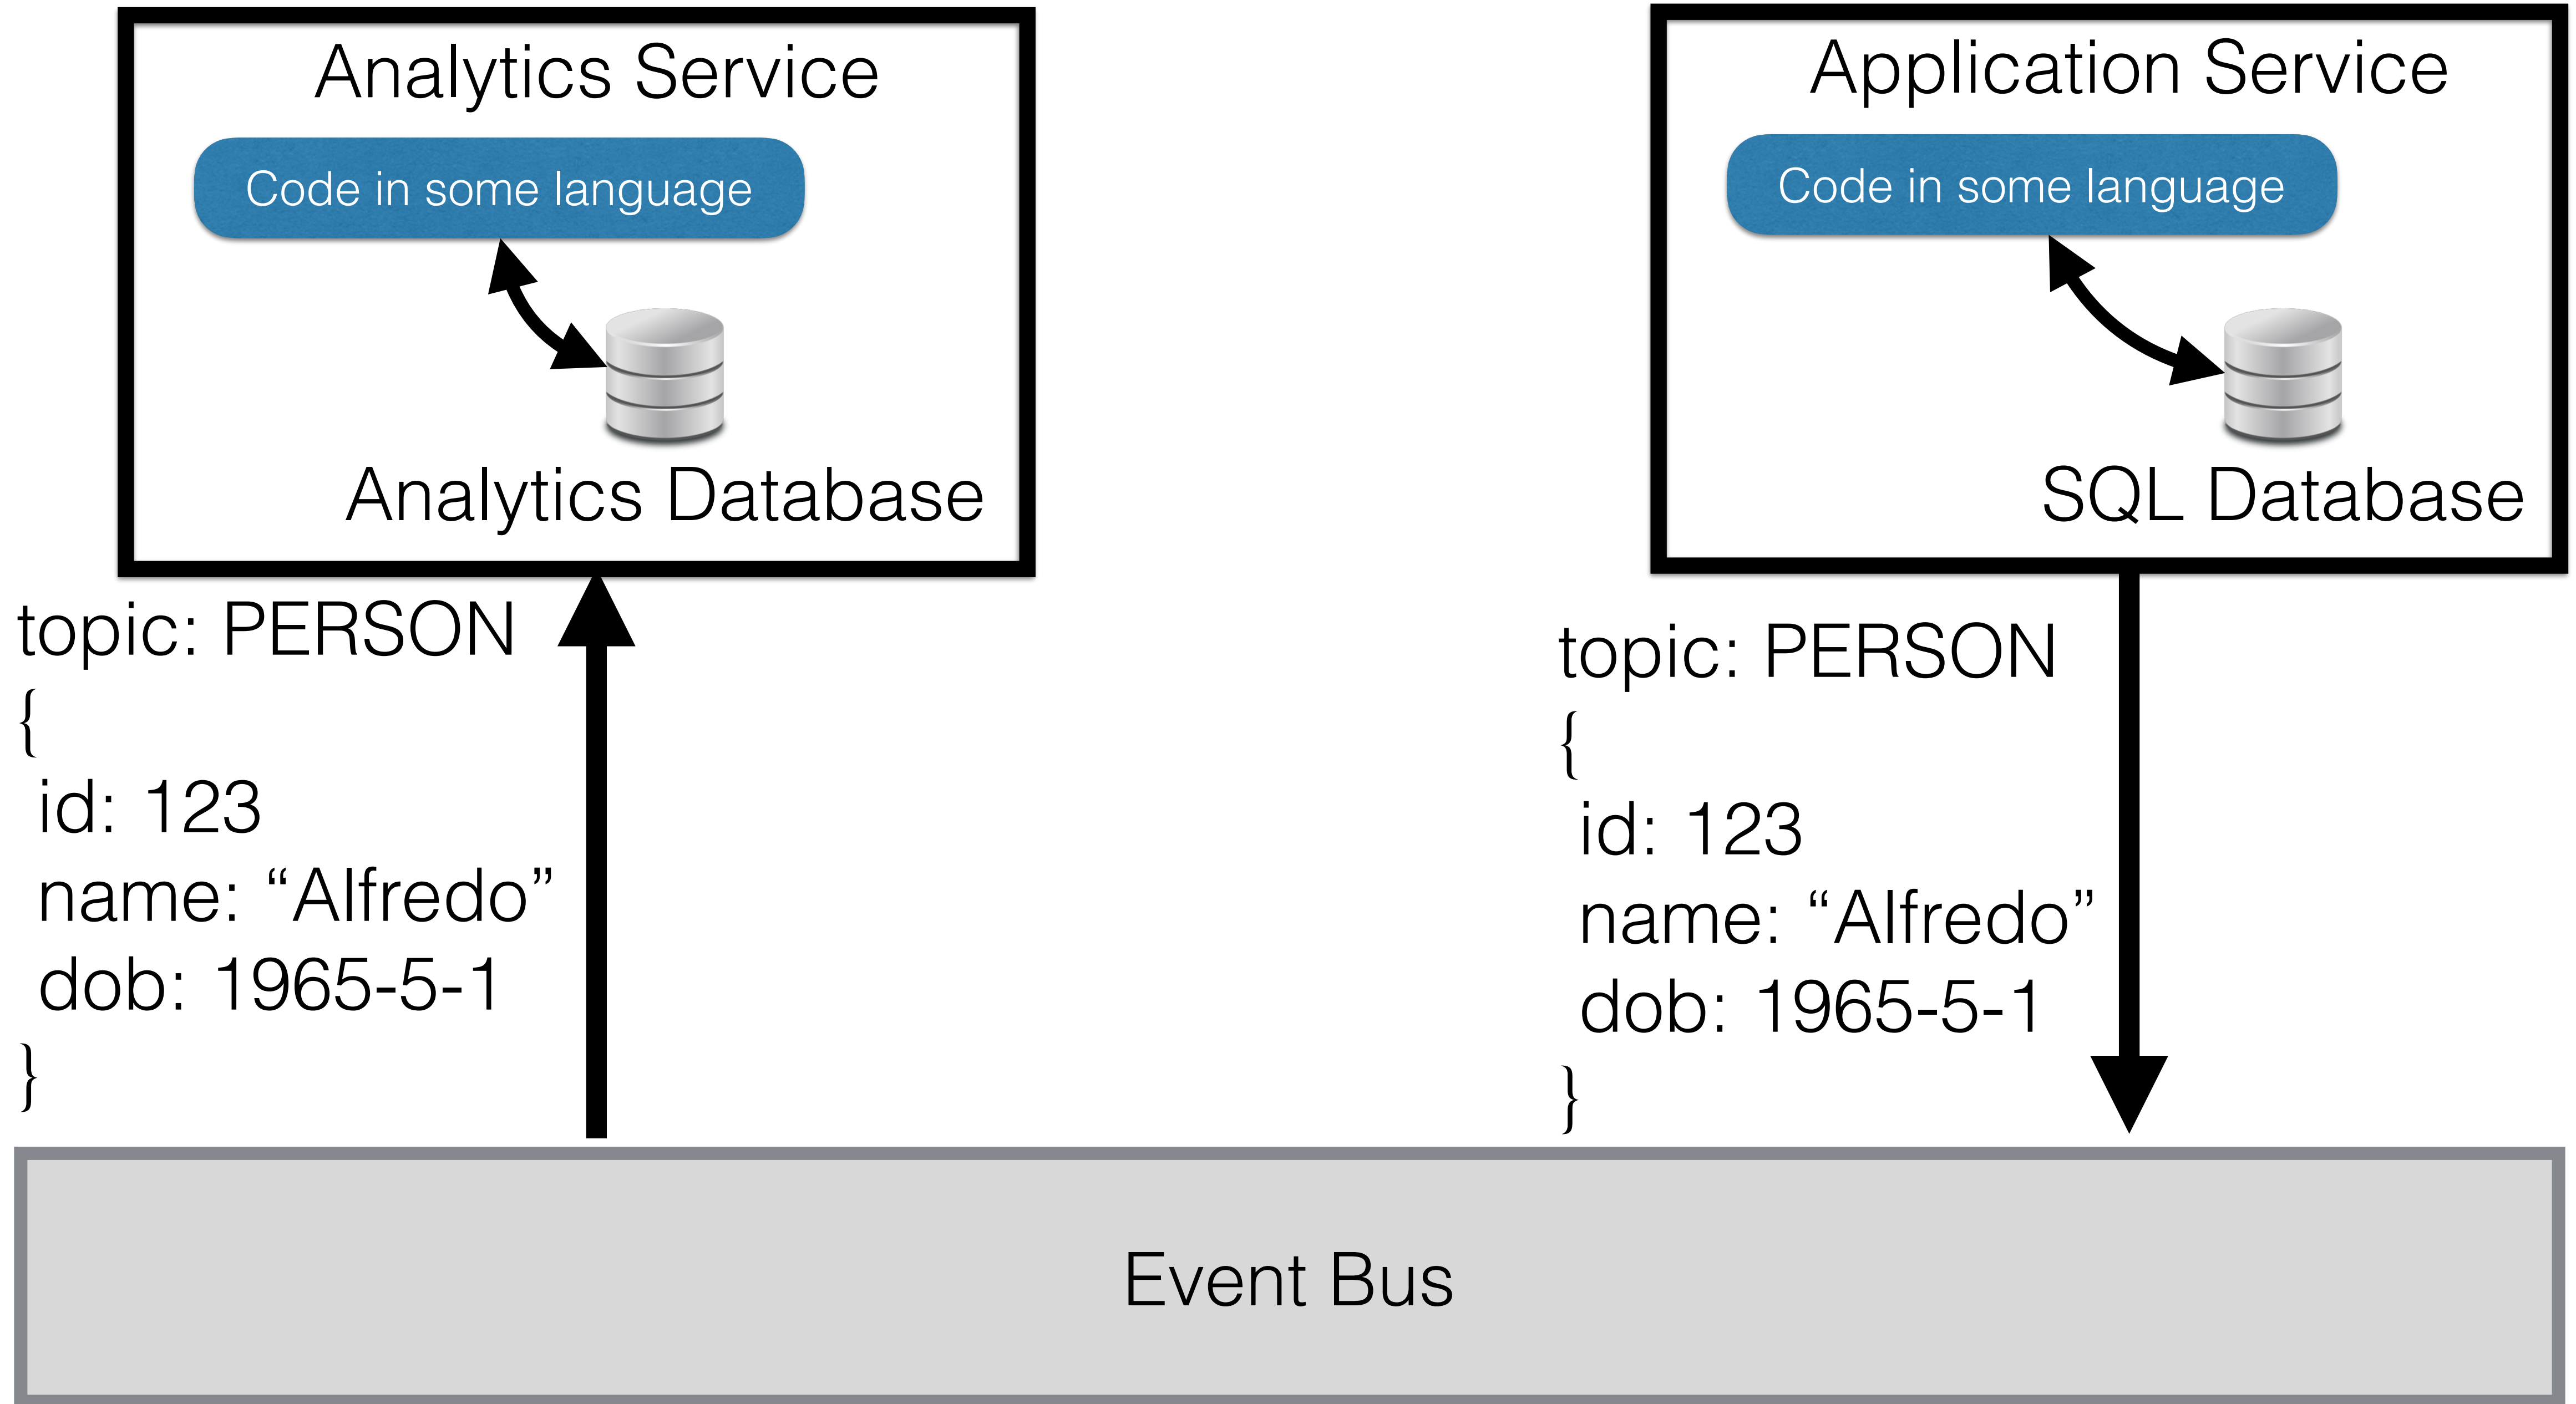
- Data stores are no longer private. We have a tight coupling between databases. It's not a micro service
- We can't use any interesting databases :-)
- Difficult to scale out to many services

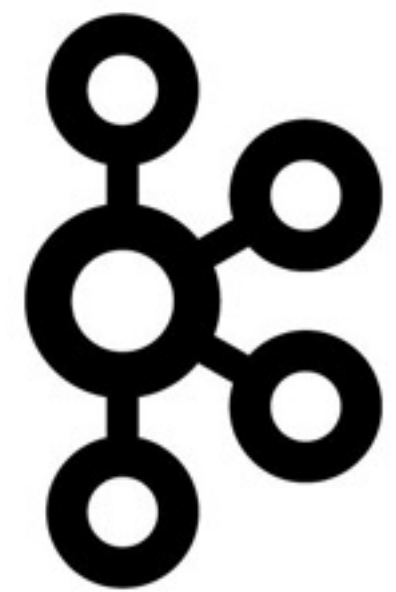
Event Driven Microservices

Not your father's micro services

Event Driven Microservices

- Services push events instead of a request/response model
- Usually backed by a publish/subscribe bus





Kafka

Kafka

- Persistent pub/sub message bus
- High throughput
- Subscribers can consume at their own speed
- Subscribers can request a 'rewind' and re-consume a topic
- Has some tricks to keep the data volume down
- Having both fast and slow consumers is not a problem

sport.link

- Service provider for professional and amateur team sports in the Netherlands
- 10+ years old
- Managing personal data, planning competitions, assigning officials, supplying data feeds

sport.link

- 1M+ players
- 4K+ clubs
- 40K matches a week
- Spikey but predictable load

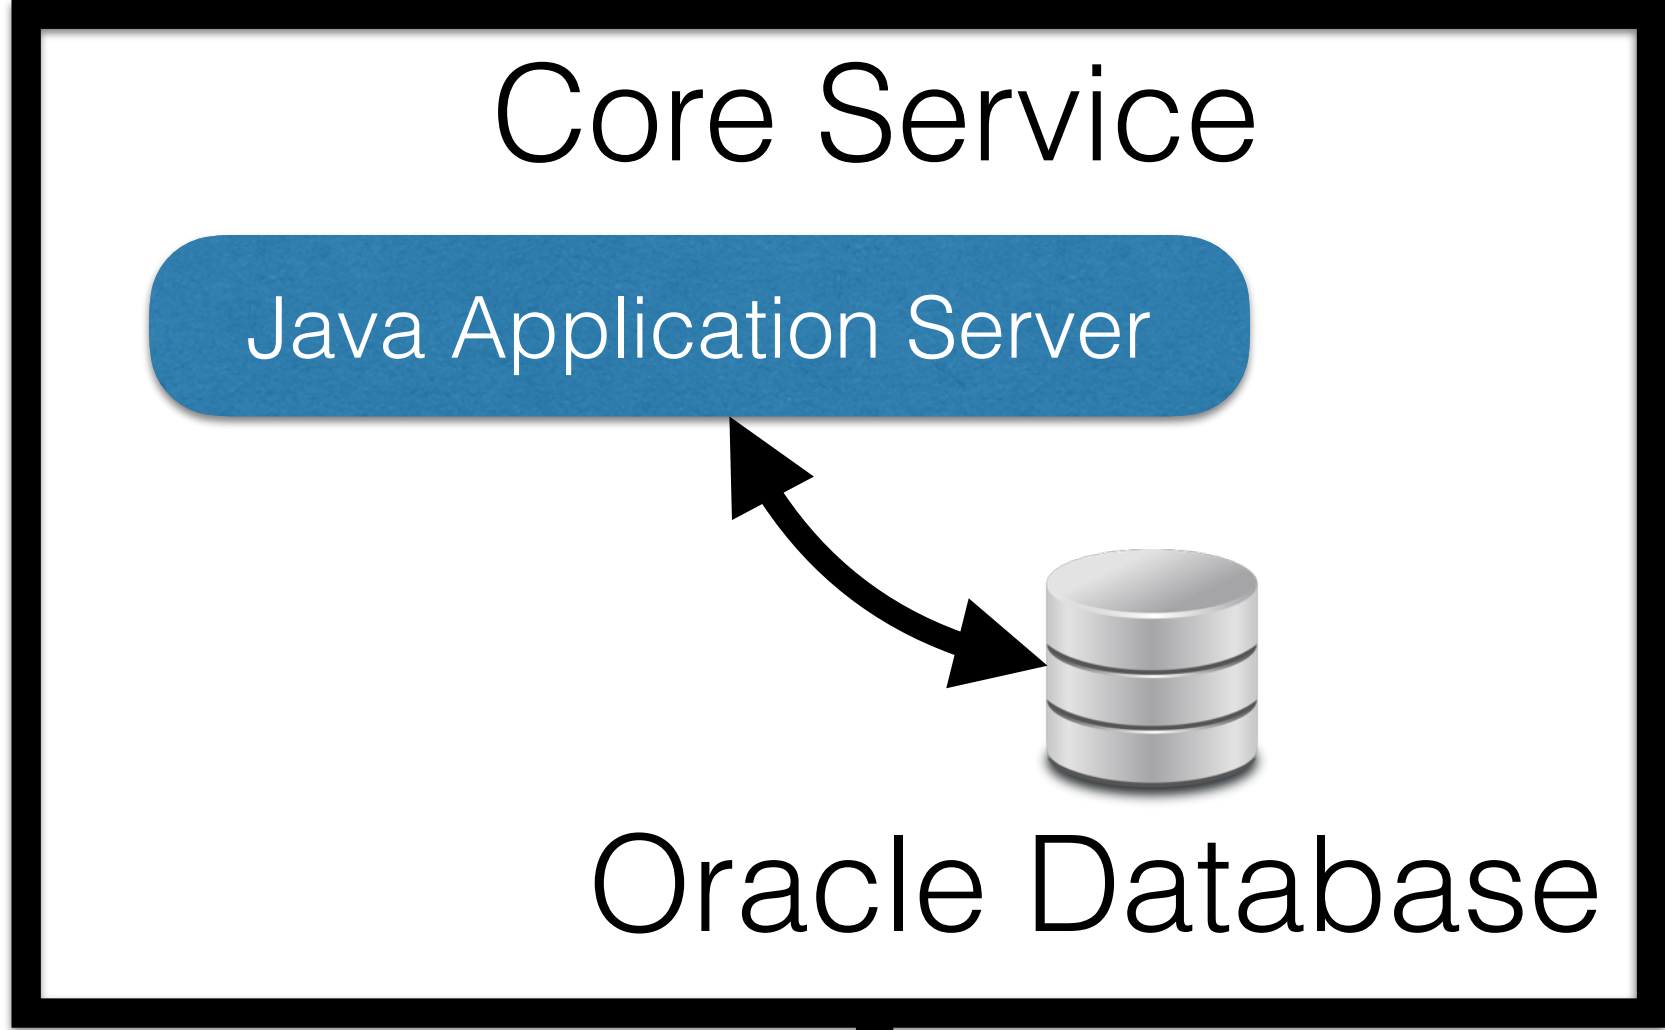
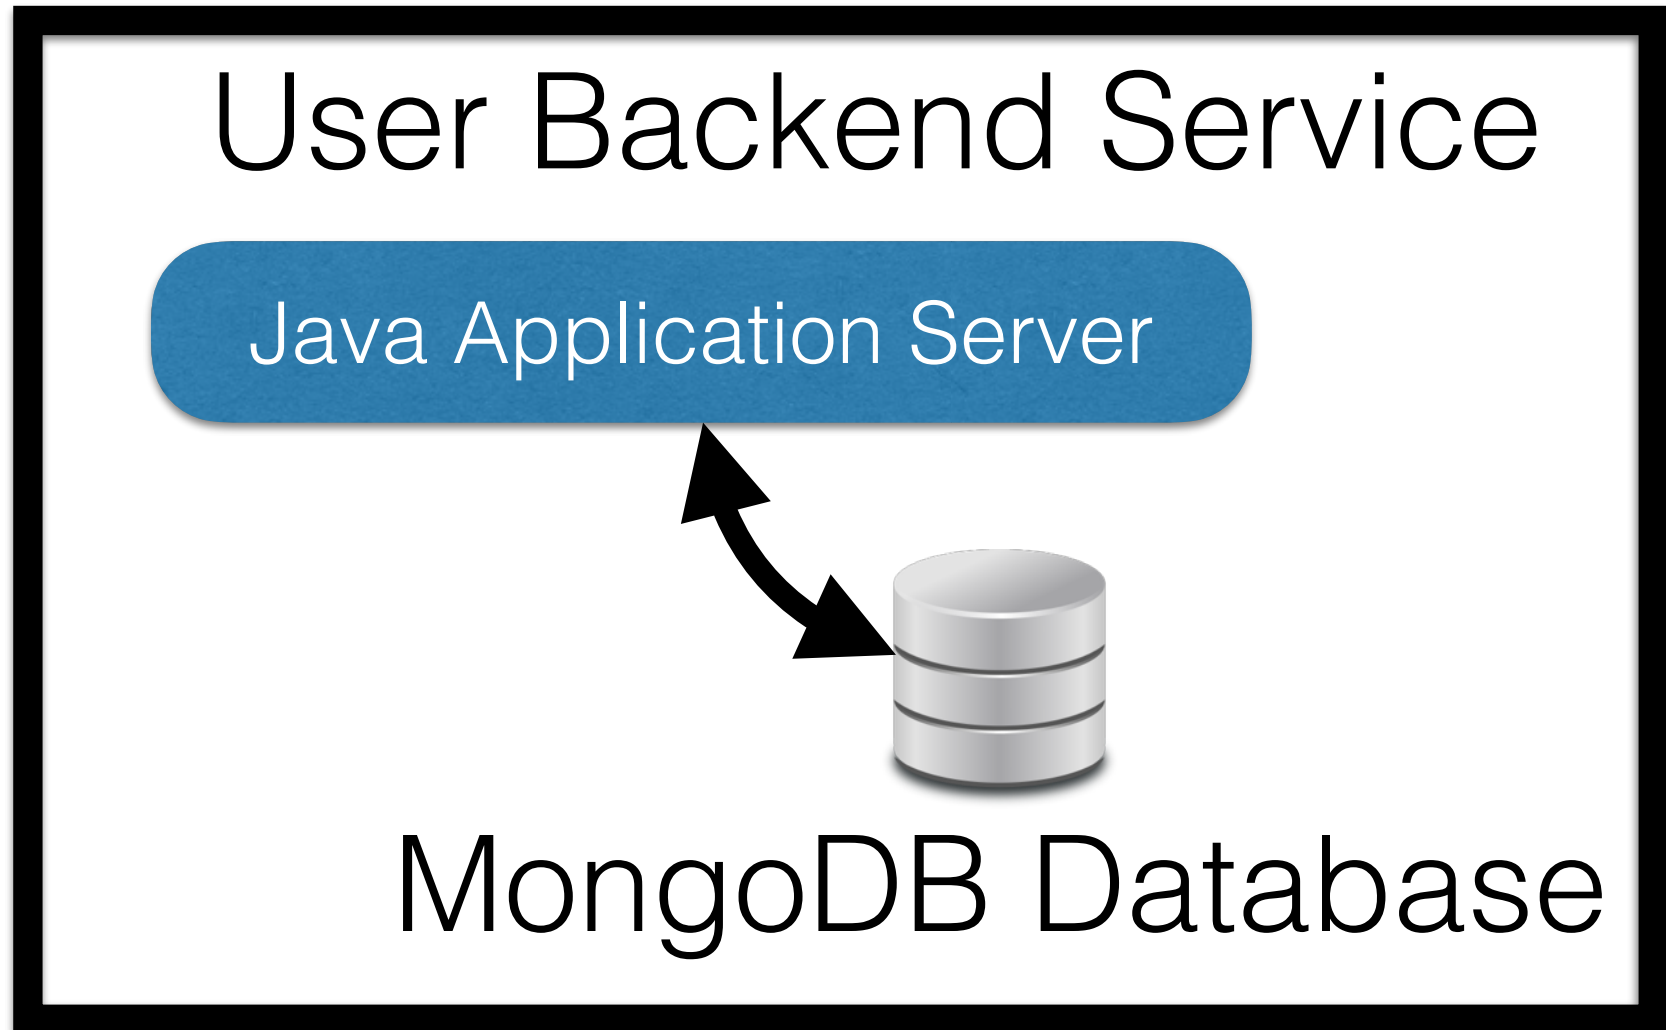
Technology stack

sport.link

- Oracle database
- Cluster of Java based application servers
- Diverse set of clients

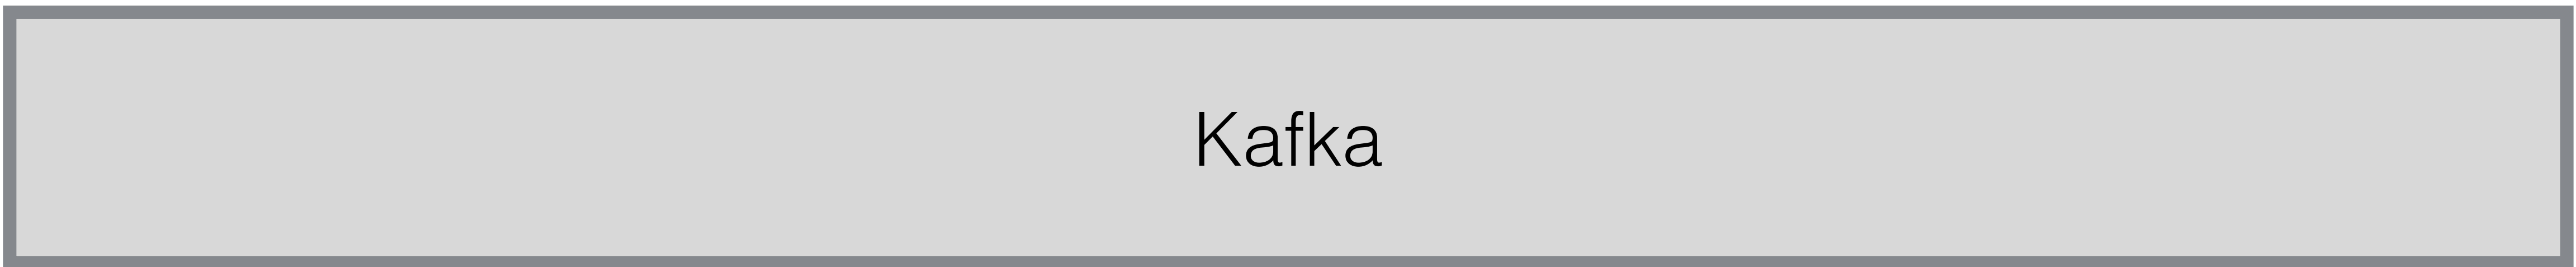
Challenge

- Move to a player centric model instead of a club centric
- Order of magnitude more users and load
- Moving away from Oracle is not feasible in the short term
- Scaling Oracle is just too expensive



```
topic: PERSON
{
  id: 123
  name: "Alfredo"
  dob: 1965-5-1
}
```

```
topic: PERSON
{
  id: 123
  name: "Alfredo"
  dob: 1965-5-1
}
```





Example

Telephoneld	PersonId	Type	Telephone
1	1	Mobile	12345
2	1	Email	alfredo@aol.com
3	1	Twitter	@alfredo

PersonId	Name	Birthdate
1	Alfredo	1990-01-01
2	Ben	1991-01-02

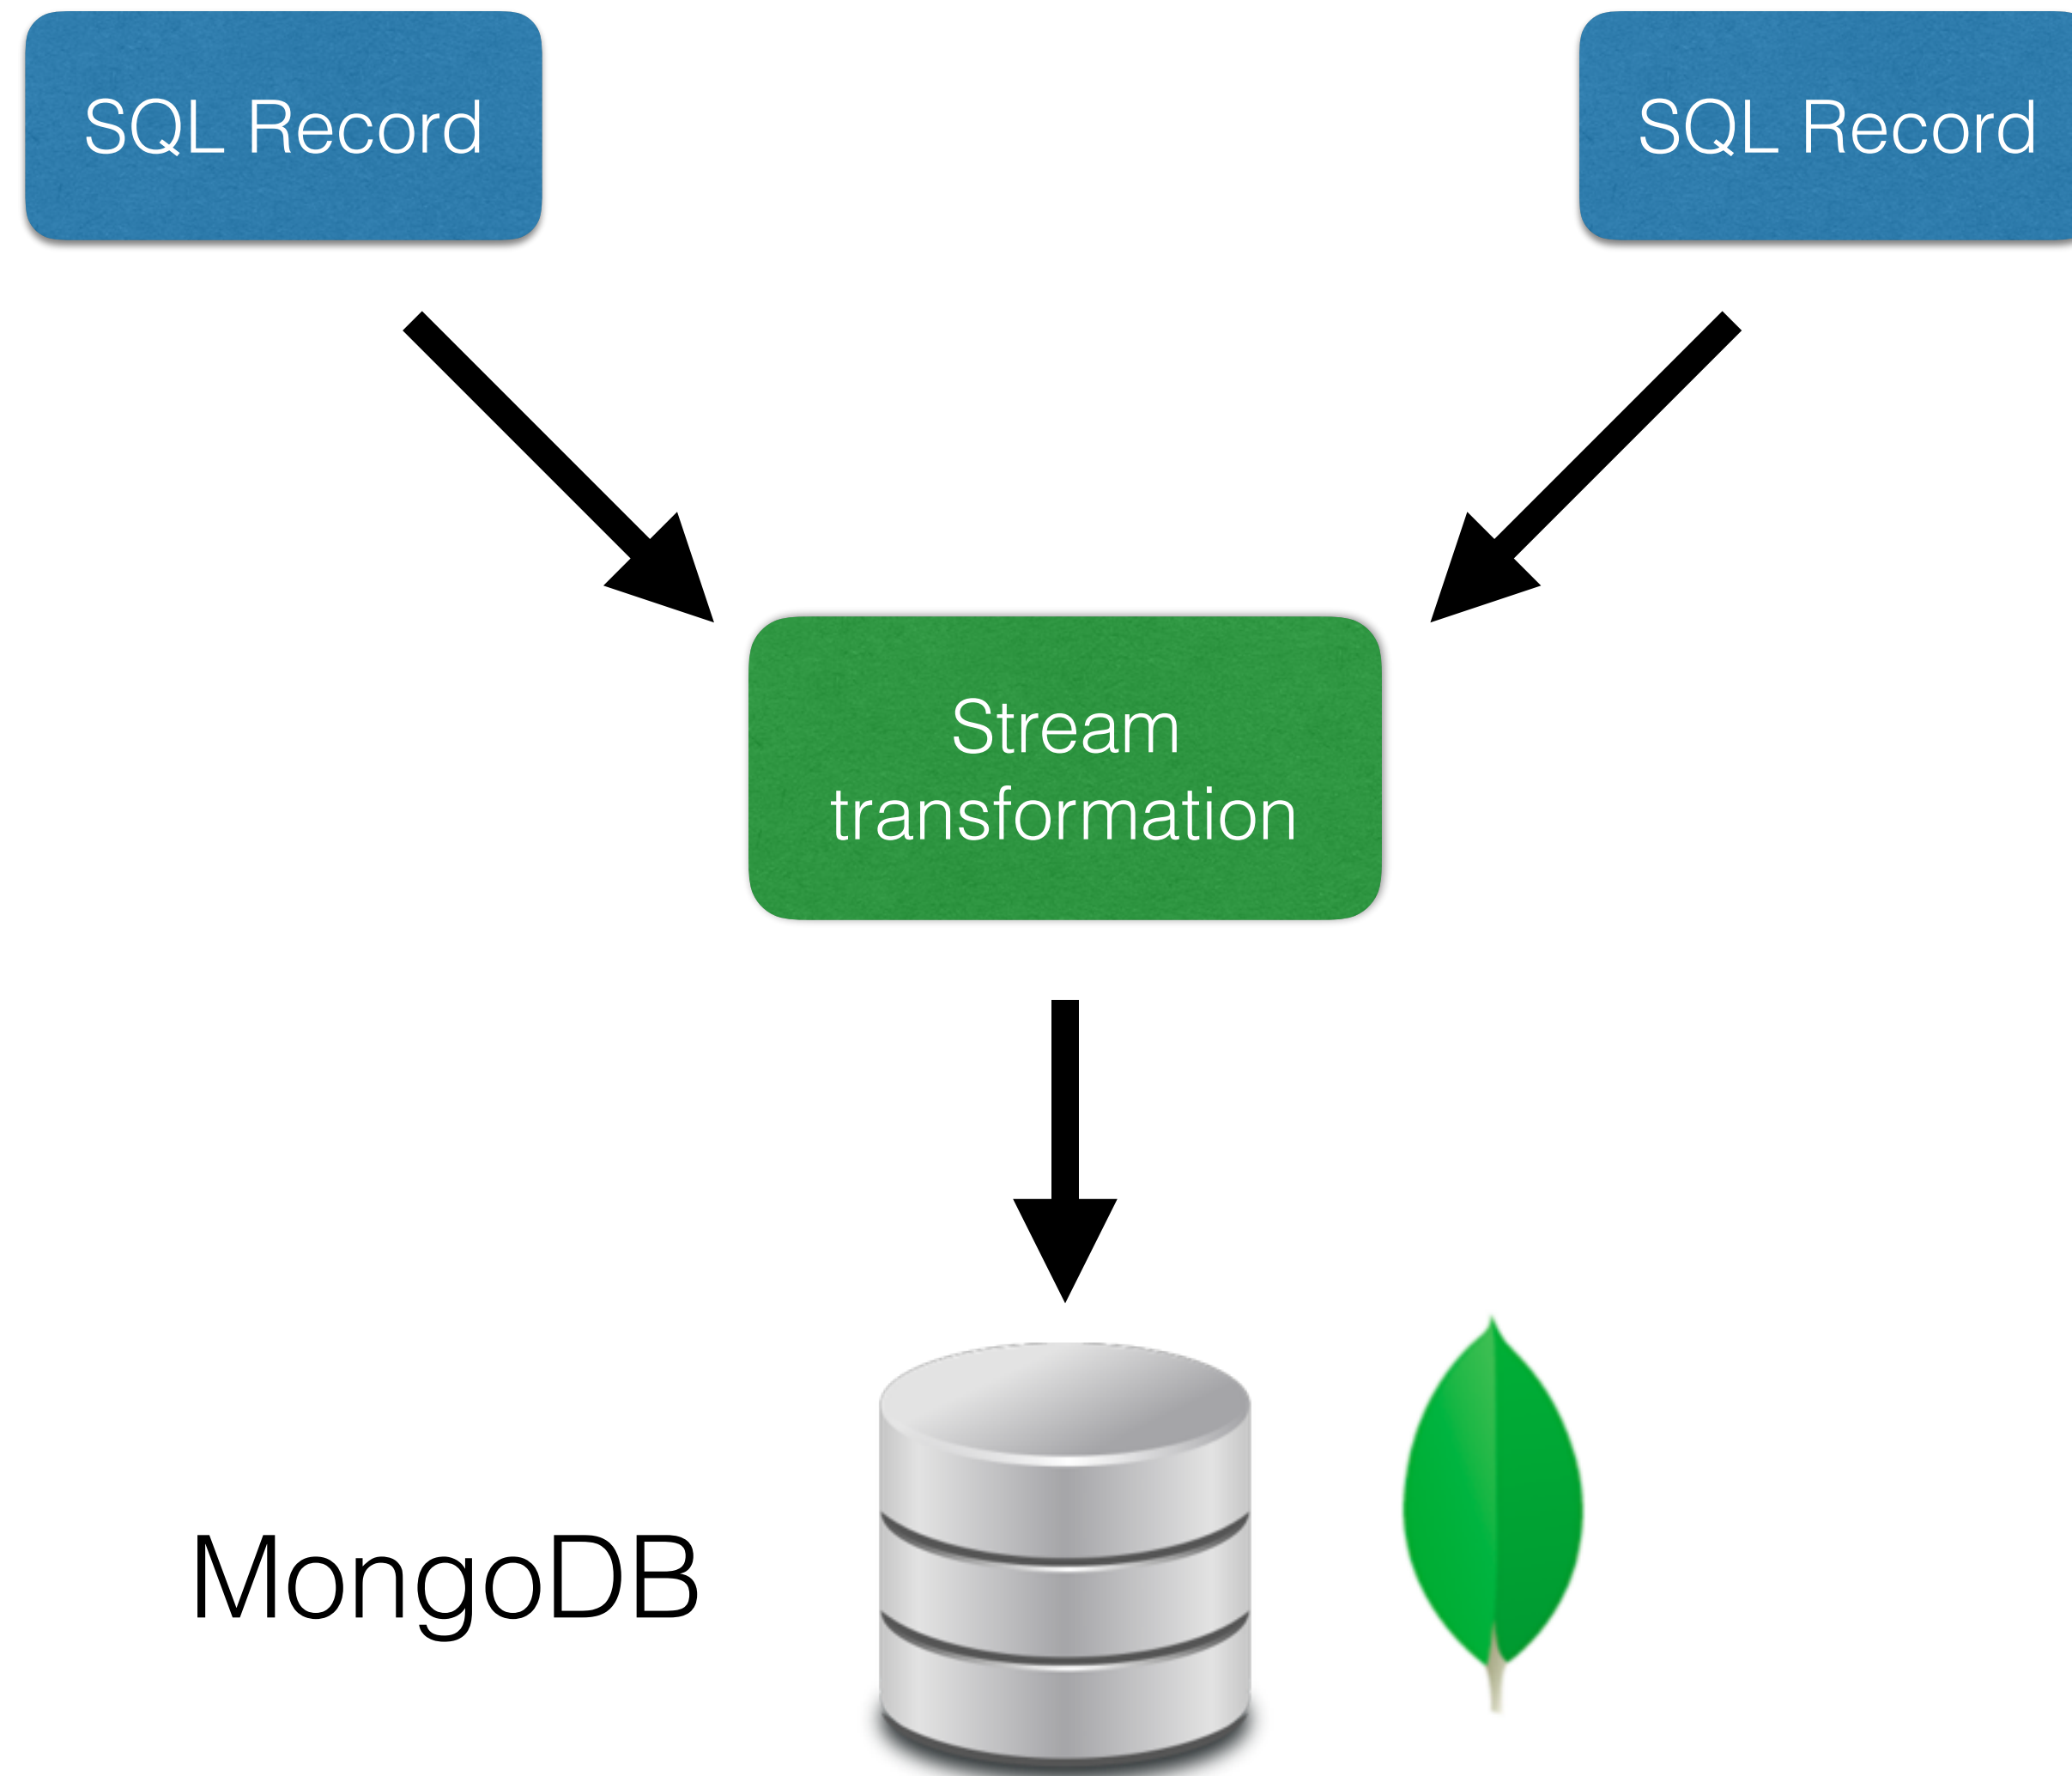
```
SELECT * FROM Communication C WHERE PersonId = 1
```

```
SELECT * FROM Person WHERE PersonId=1
```

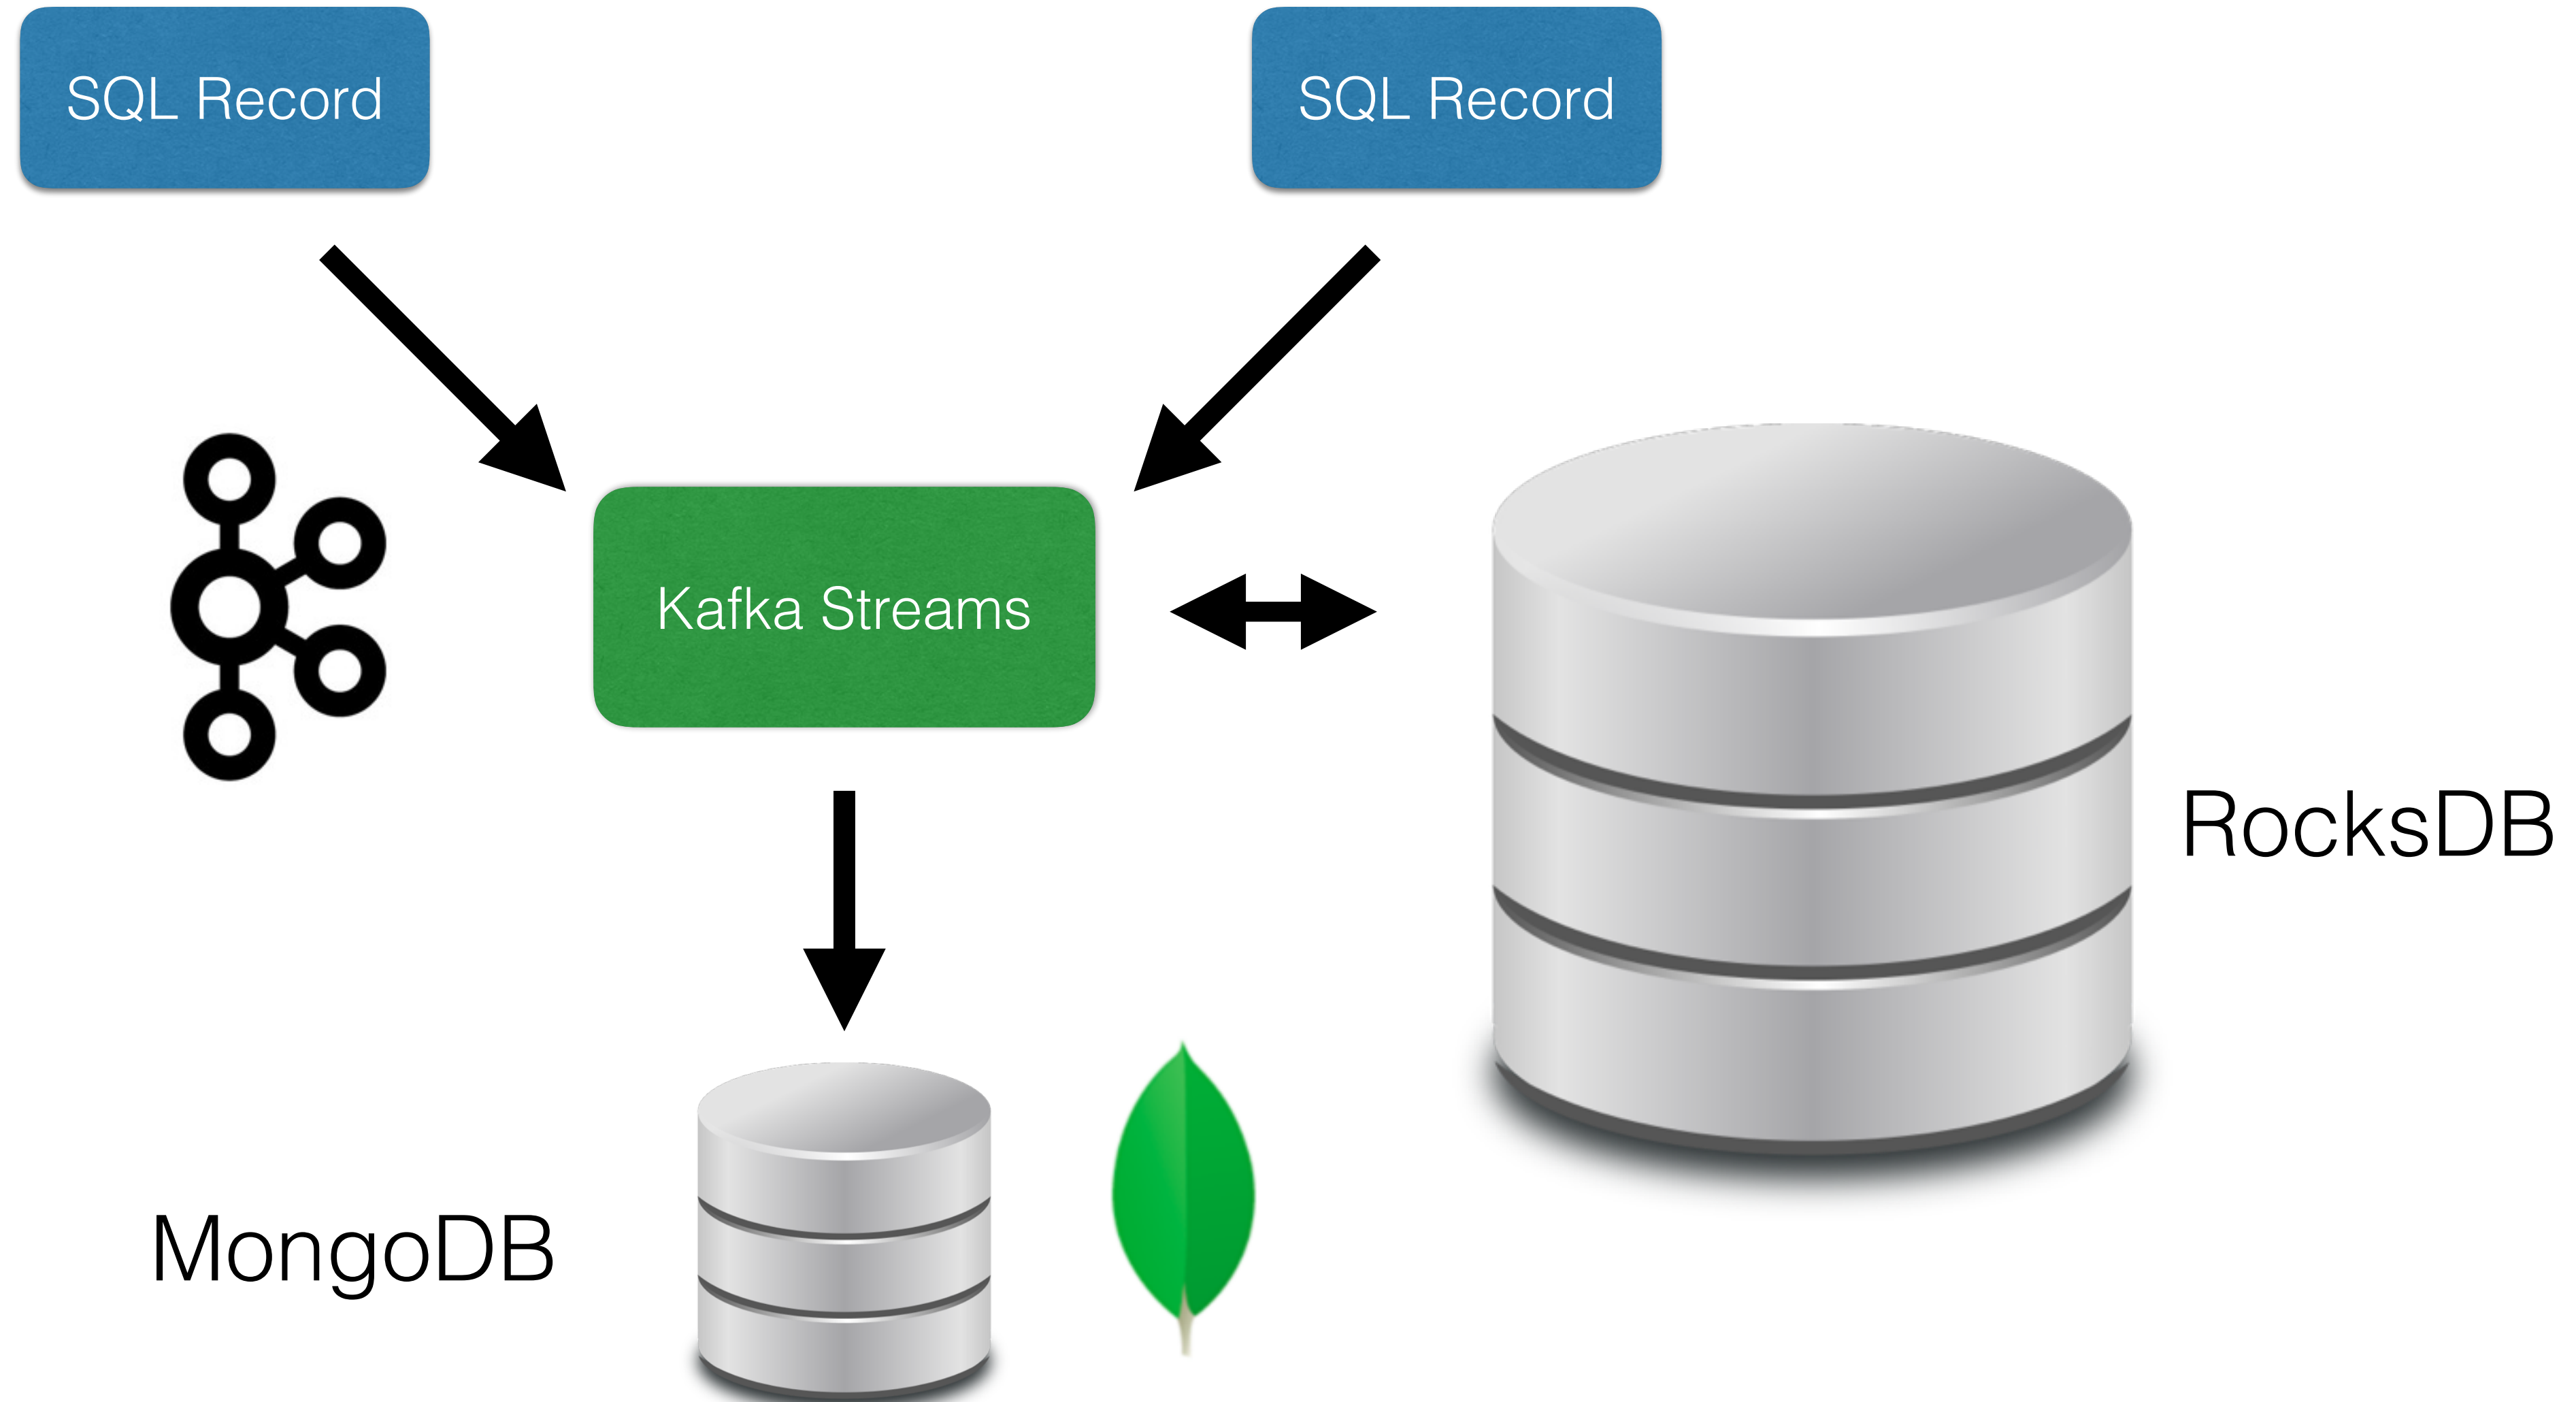
MongoDB

```
{  
  "_id":1,  
  "Name":"Alfredo",  
  "DOB":"1990-1-1",  
  "Communication":{  
    "Mobile":"12345",  
    "Email":"alfredo@aol.com",  
    "Twitter":"@alfredo"  
  }  
}
```

Stream Processing



Stream Processing



Kafka Streams at Scale

- \pm 500M rows of SQL data
- \pm 50 joins
- 500 topics
- 400 Gb of Kafka Data
- 300 Gb of RocksDb data
- Building a complete replica from scratch takes many hours
- After that <100 ms latency for changes

Development cycle

- Developing and testing is hard for stateful code
- Starting a new 'generation' is costly
- Contaminated data might show up

Conclusions

- Went into production early June
- Generally behaves well (aside from some glitches)
- Kafka Streams is in a lot better state than a few months ago



NOW THE FUN BEGINS

Elasticsearch

- Add unstructured search to our application
- Reduces load on our source databases
- Users expect google-like interfaces

Neo4J

- Graph Database
- Some analytics are much easier to express in terms of graphs

Firestore Realtime

- Real-time database
- 'Backend As a Service'
- Essentially one big JSON document
- Very easy to use client libraries for web and mobile
- Safe to develop

Caches

- We can use our streaming engine to update caches

Push to clients

- We can push data to clients in real time

Things I ignored

- Change Capture (“Change Data Capture: The Magic Wand We Forgot”)
- Eventual consistency
- Kafka compaction / partitioning



DATABASES ARE
PEOPLE



made with ♥ @ Recite.com



Thank you
Questions?*

* Any question that is not: "Why don't you use Postgres? Postgres can do anything"