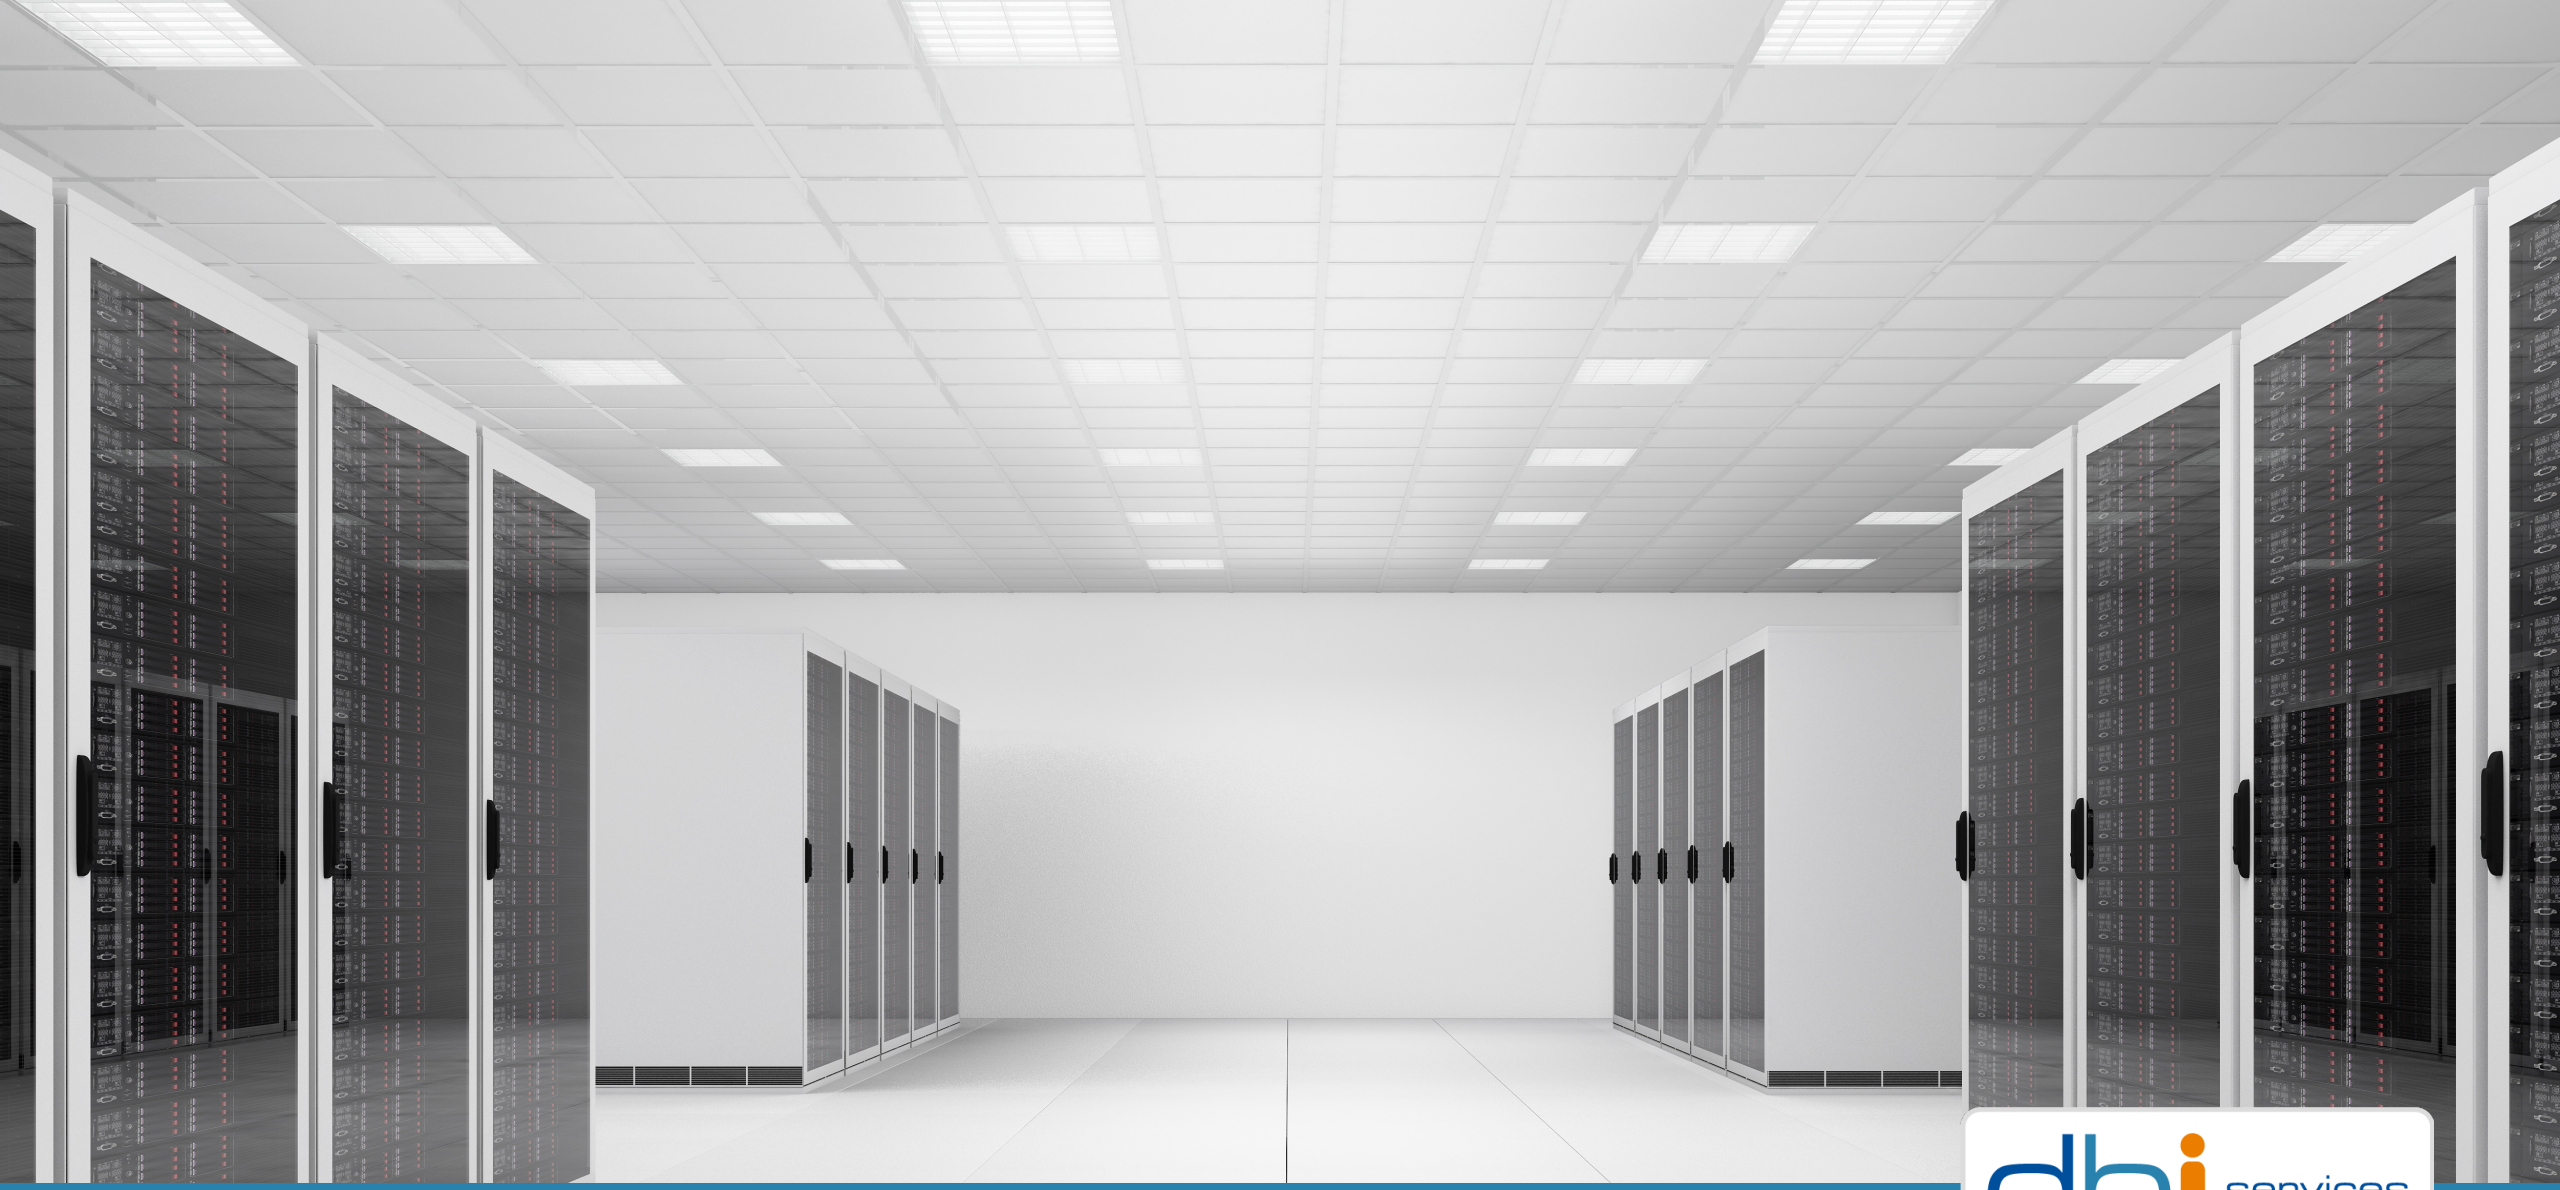

How to migrate data from MongoDB to Postgres with ToroDB

Who we are

Experts At Your Service

> Over 50 specialists in IT infrastructure
> Certified, experienced, passionate

Based In Switzerland

> 100% self-financed Swiss company
> Over CHF8 mio. Turnover

Leading In Infrastructure Services

> More than 150 customers in CH, D & F
> Over 50 SLAs dbi FlexService contracted

dbi services is hiring (career@dbi-services.com)

About me

Mehdi Bada

Consultant

+41 79 928 75 48

mehdi.bada[at]dbi-services.com

Agenda

1. Introduction
2. MongoDB
3. ToroDB
4. Migration: from MongoDB to PostgreSQL
5. Conclusion

Introduction

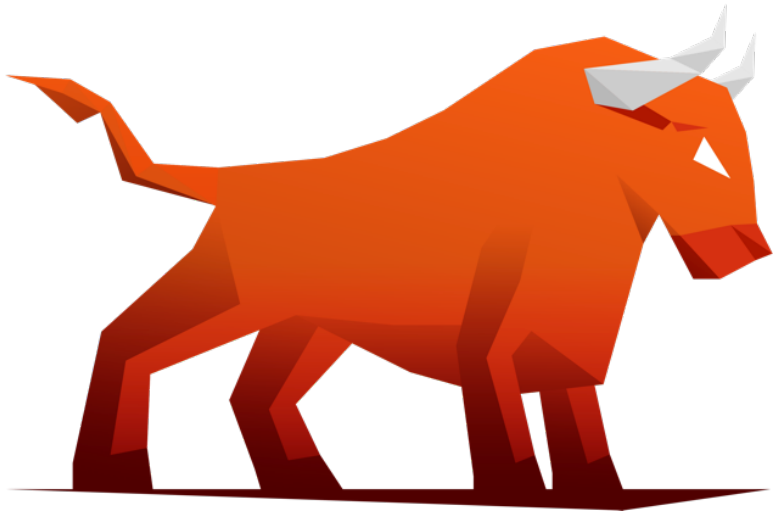

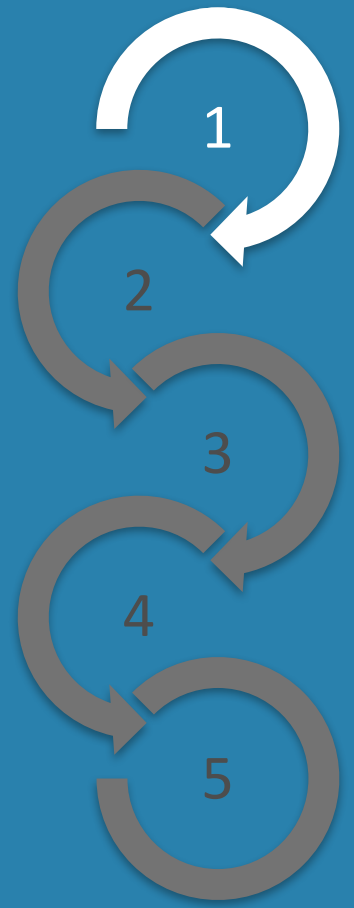

Introduction

MongoDB

- Overview
- Data Model
- High availability
- Horizontal scalability
- Limitations

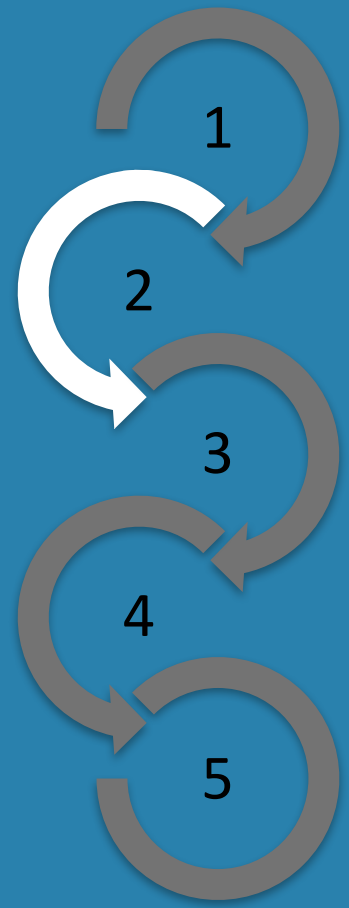

MongoDB
Overview

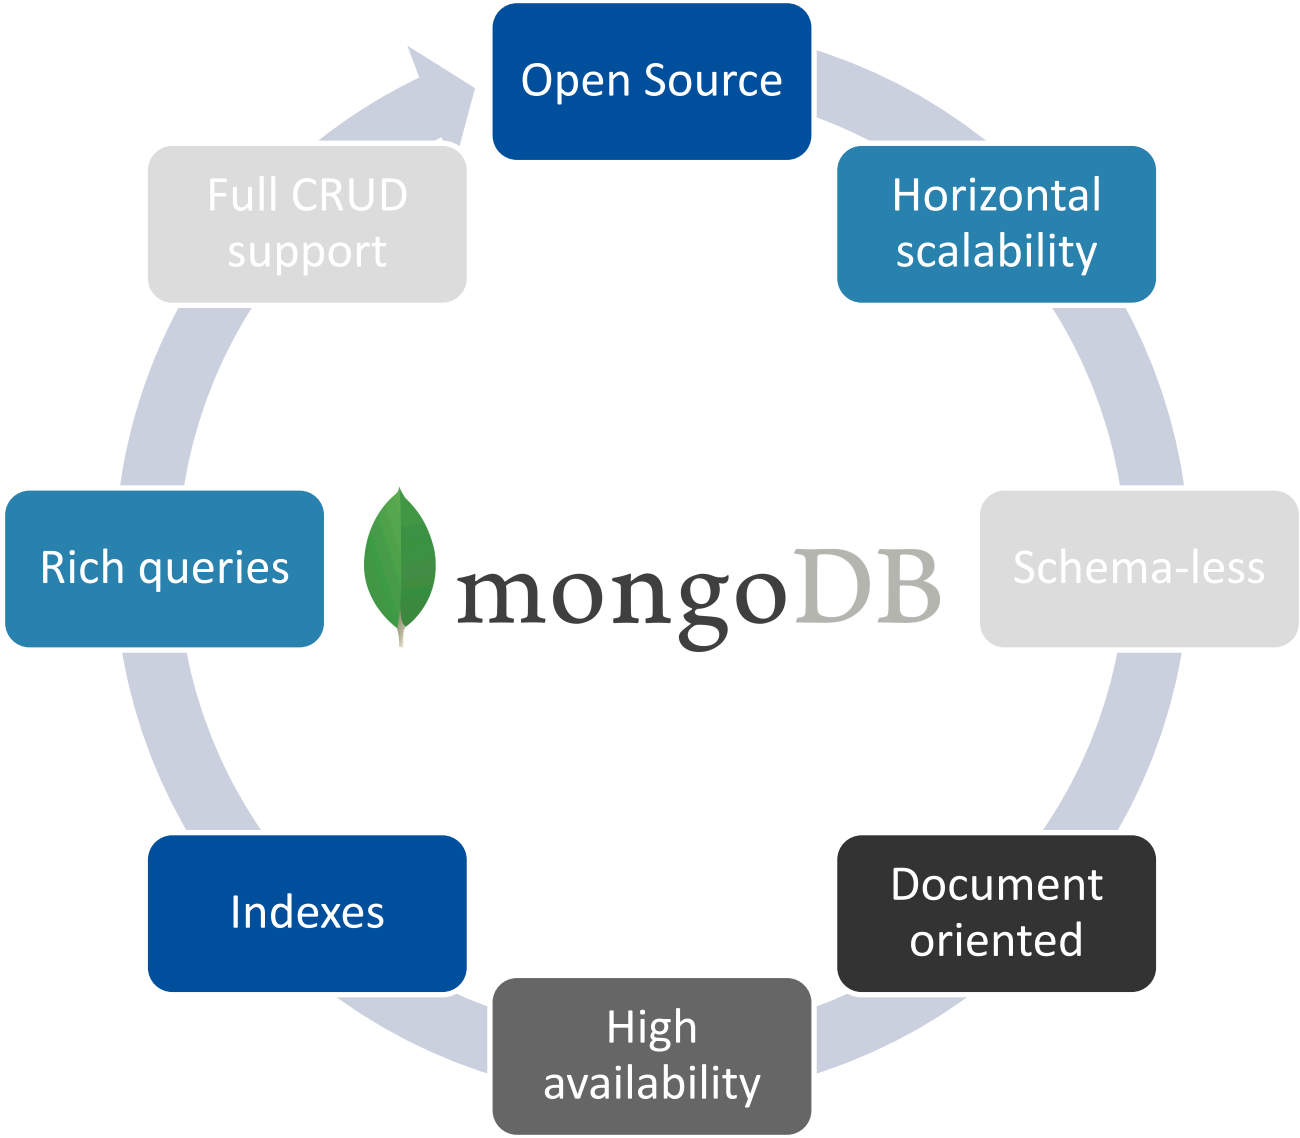

Concept mapping

RDBMS	MongoDB
Tables	Collections
Rows/records	Documents
Queries return rows	Queries return cursor
Join	Embedded document
Partition	Shard

MongoDB
Data Model

Data are stored as documents

MongoDB stores BSON documents (Binary JSON)

Analogous to a database row

Keys and Values

- Key : String
- Values types:
  - String, number, Boolean, null, array, object

```
{
    “_id” : ObjectId(“56e92b9cfdf7bc92bbb3b51f”),
    “first_name” : “Mike”,
    “surname” : “Brody”,
    “city” : ”New-York”,
    “year” : 1987,
}
```

Special key: _id

- Unique identifier
- Object id:

MongoDB
Data Model

Relational

Pers_Id	Surname	First_name	City
0	Mike	Durand	Geneva
1	Pat	Millner	London
2	Ortega	Alvaro	New-York

Car_Id	Model	Year	Pers_Id
0	Ferrari	2013	0
1	Peugeot	2005	0
2	BMW	2016	1

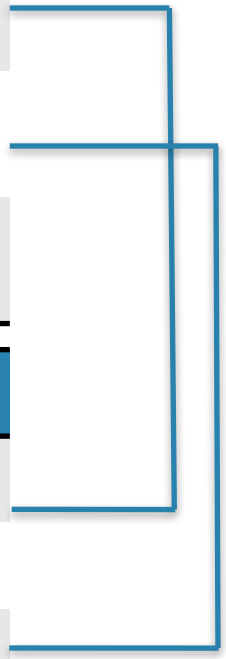

MongoDB schema

```
{
    “first_name” : “Durand”,
    “surname” : “Mike",
    “city” : “Geneva",
    “country” : “Switzerland”,
    “cars”:[
     { “model”: “Ferrari”,
       “year”: 2013
       }
 }
```

MongoDB
High availability

Which mechanism ensure high availability of your data?

> MongoDB Replication

HA is achieved through automatic failover

MongoDB replication allows:

> High availability (HA)

> Disaster Recovery (DR)
  > Data duplication across multiple database servers / storages

> Functional Segregation
  > Topology of replica sets can be used for
    > Backups, Analytics, Reporting, DR, Read operations…

Replication architecture

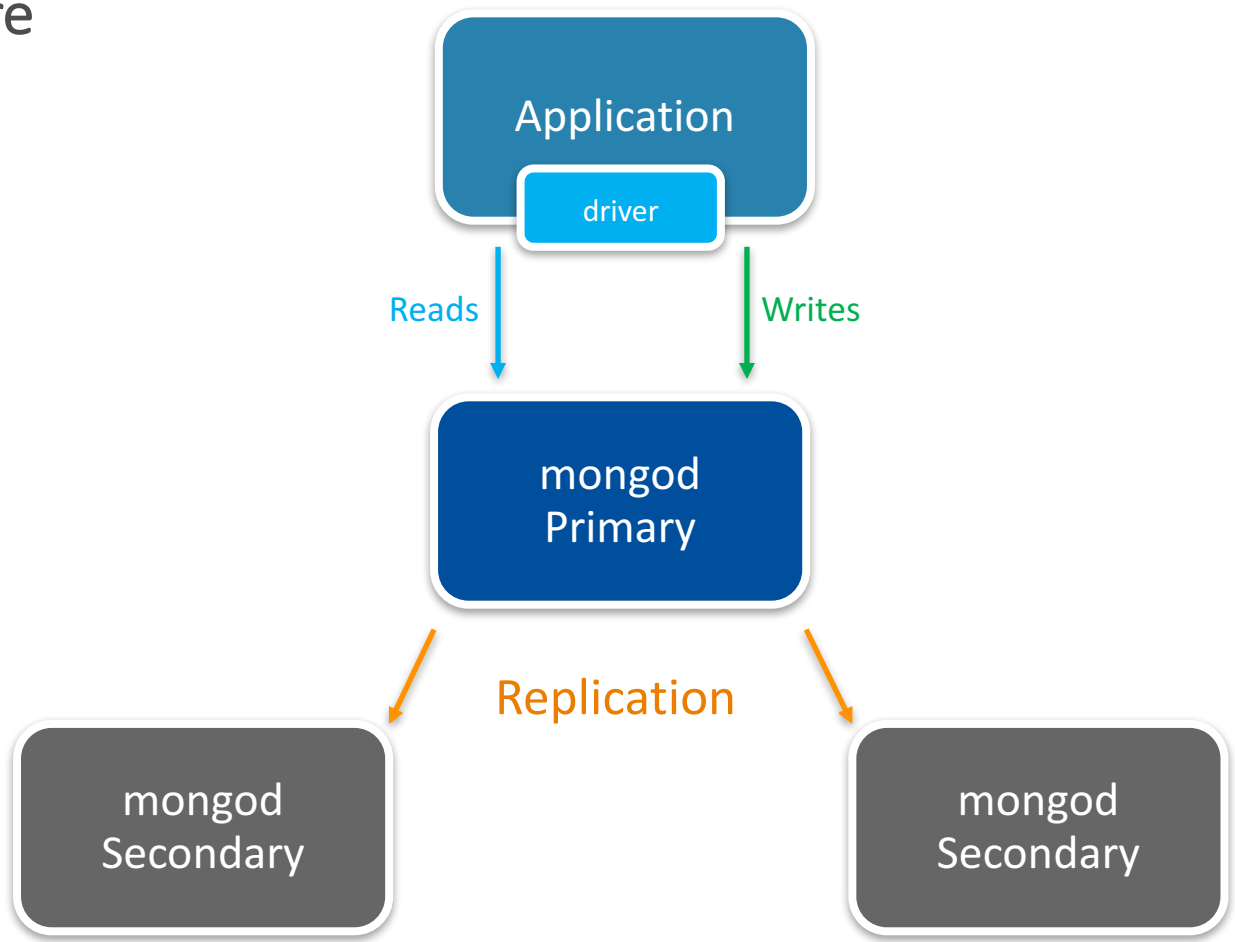

Automatic failover process

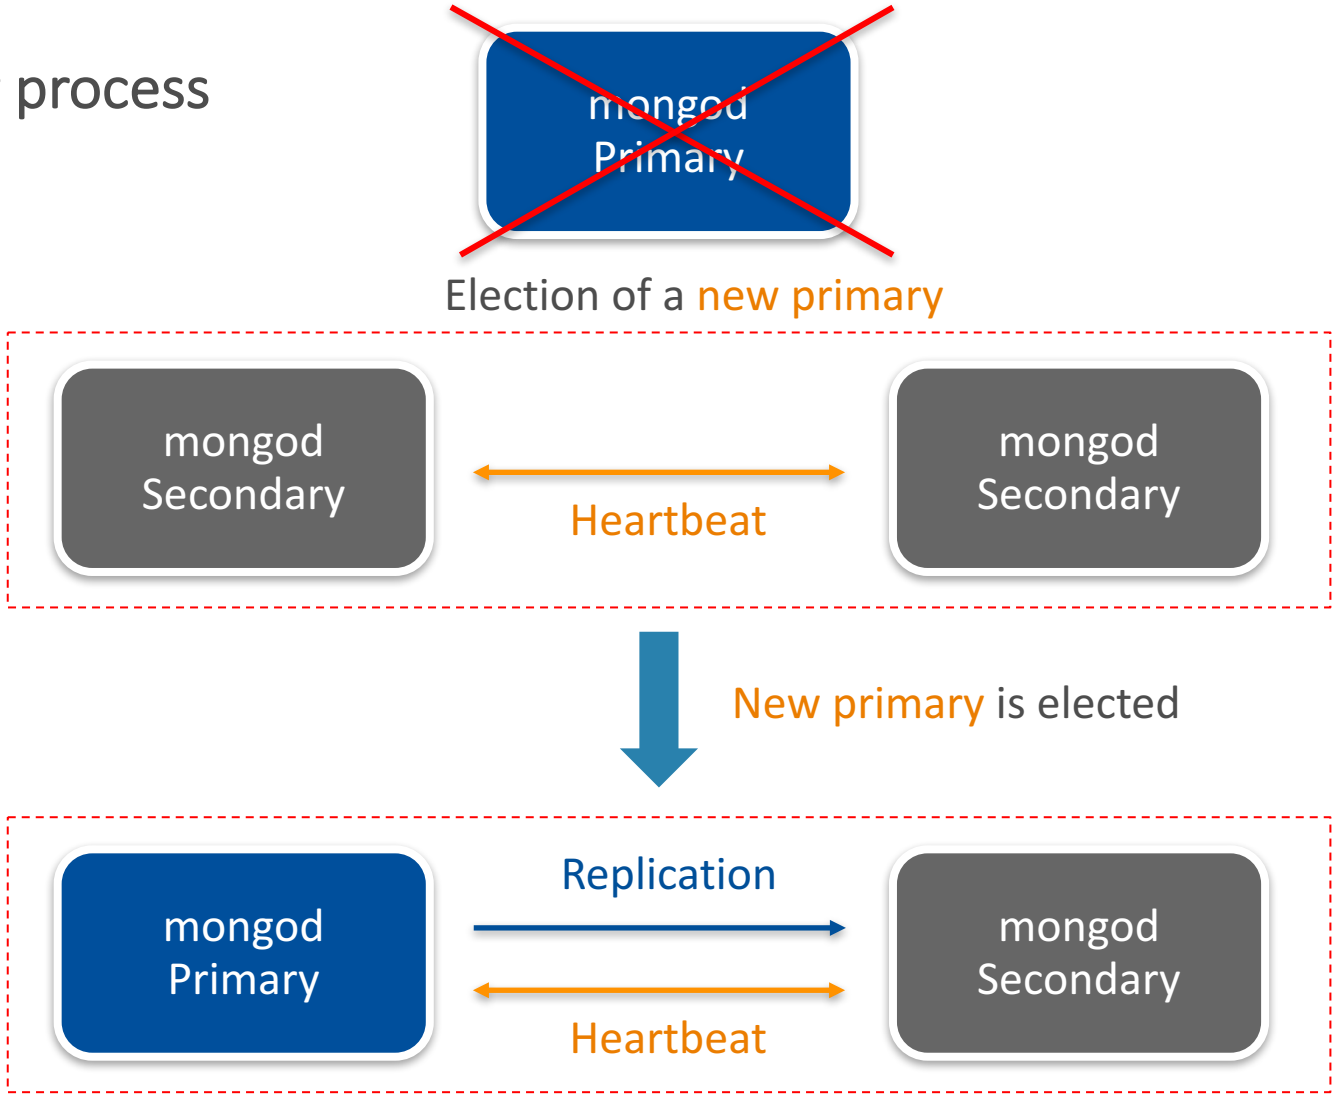

MongoDB
Horizontal scalability

Vertical scalability

Increasing CPU, RAM, I/O

Scaling with MongoDB

> MongoDB Sharding

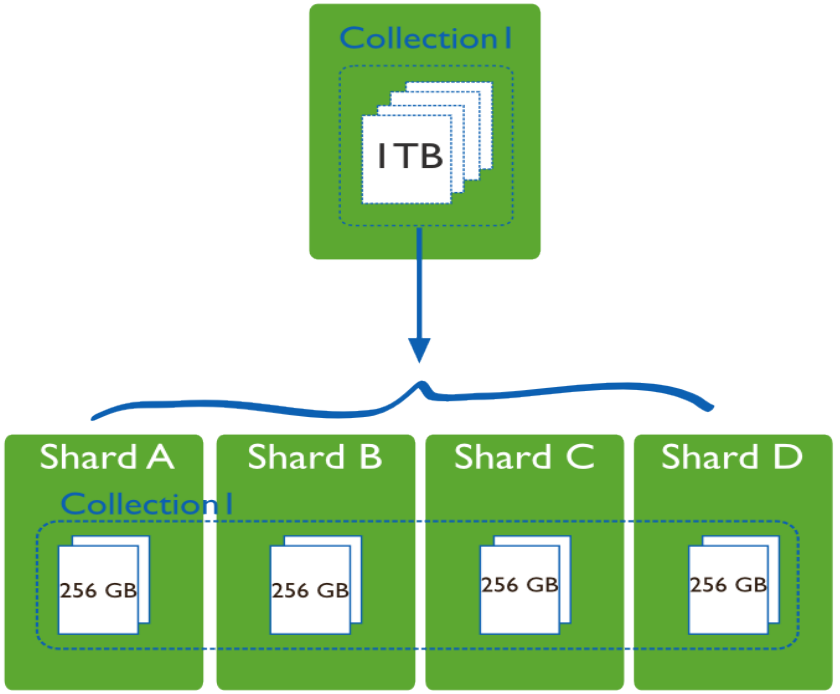

Sharding Architecture

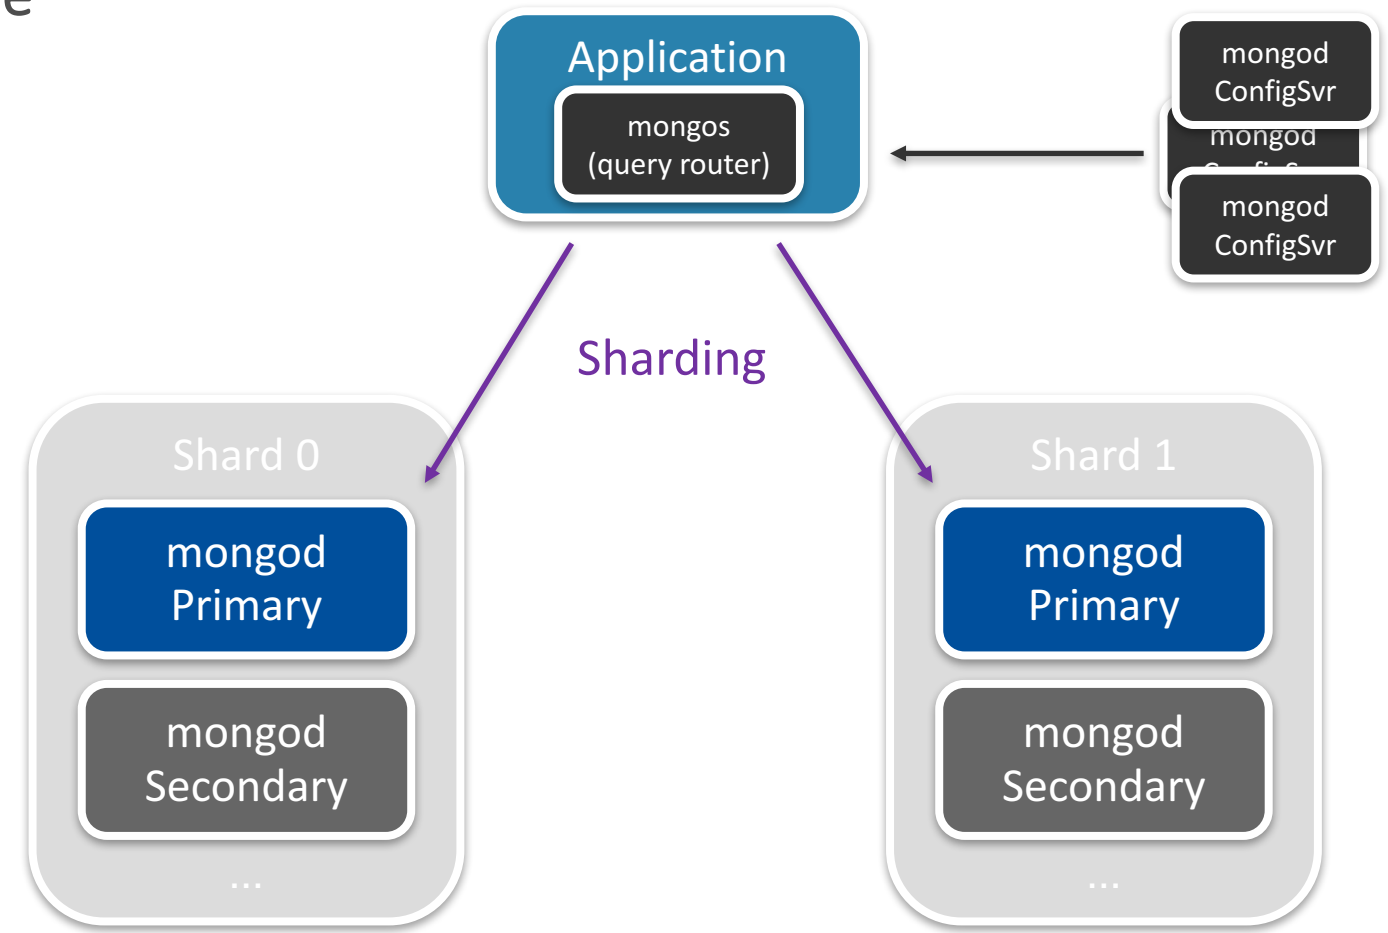

MongoDB
Limitations

No ACID transaction

> Atomic transactions only work within the same document

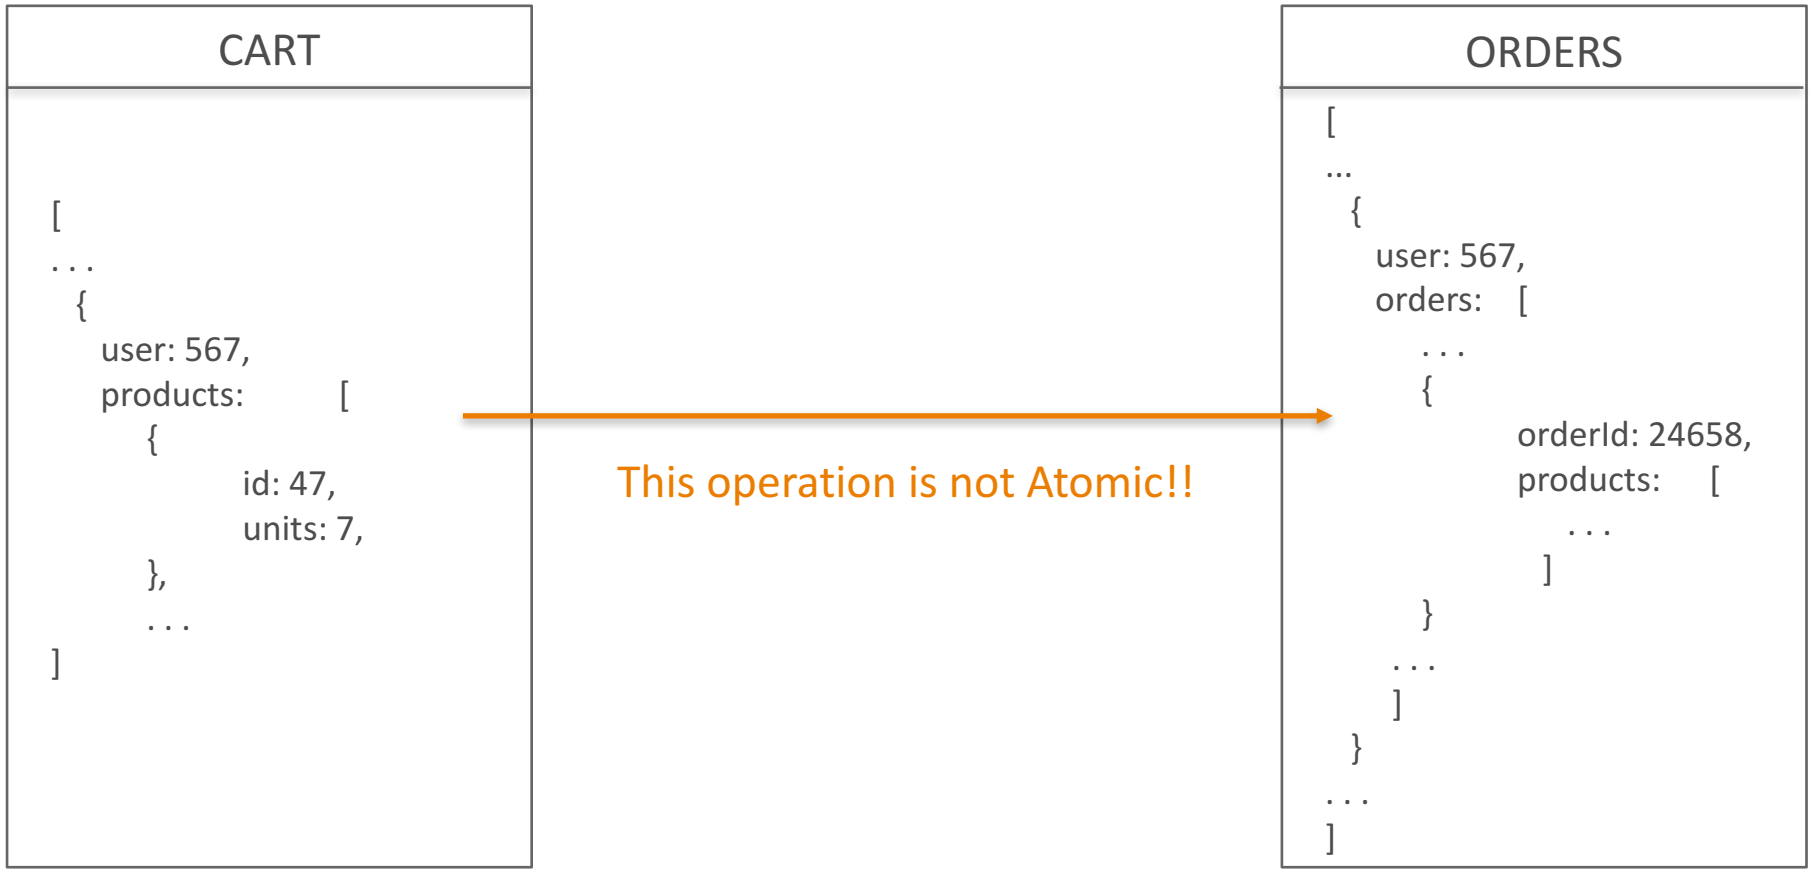

MongoDB
Limitations

MongoDB High availability is not safe!!

> Data loss depending the consistency level you choose

MongoDB consistency levels

> Unacknowledged: Unsafe - 42% of data loss
> Acknowledged: Unsafe
> Journaled: Unsafe
> Fsynced: Unsafe
> Replica Acknowledged: Unsafe
> Only majority is safe

https://aphyr.com/posts/322-jepsen-mongodb-stale-reads

MongoDB
Limitations

BI query performances issues

MongoDB aggregation framework is very slow

> Need to scan multiple documents

> Lots of I/O required to answer the query

Aggregation on a relational design is “100x faster”

Solution for MongoDB BI queries?

> Implement a relational schema!!

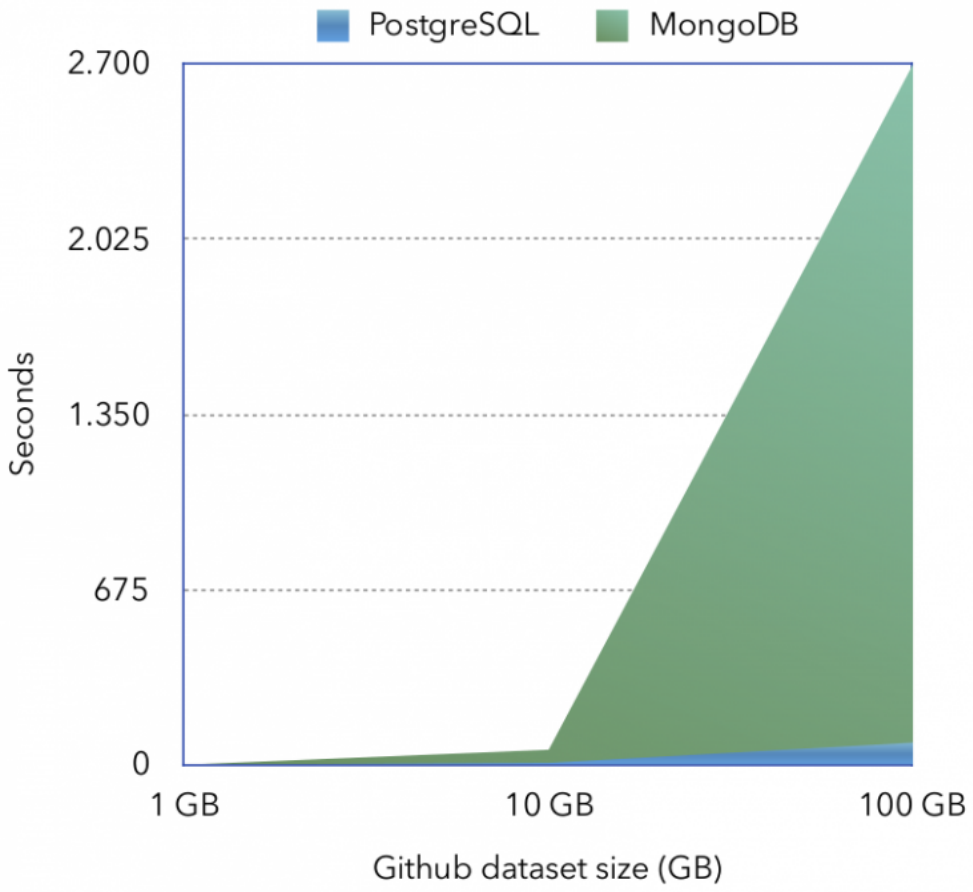

ToroDB

- What is ToroDB?
- How it works?
- Why ToroDB?

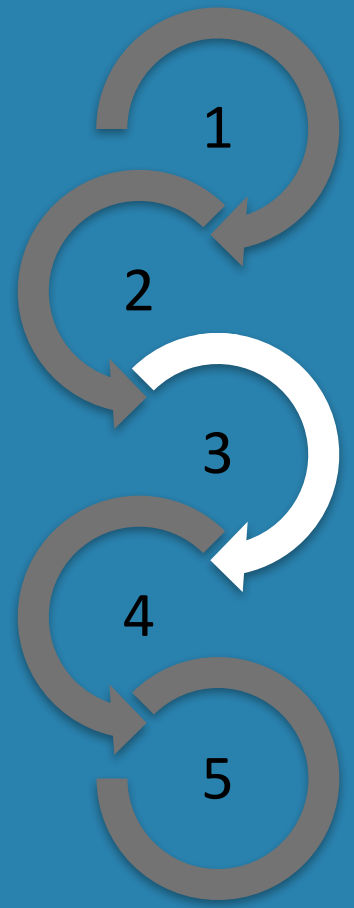

The first database that merges the scalability of a NoSQL with the reliability of SQL

ToroDB
What is ToroDB?

Open source, document-oriented, JSON database that runs on top of PostgreSQL

BI connector for MongoDB

JSON documents are stored relationally
> Significant storage
> I/O savings

MongoDB data is persisting in tables and rows within a SQL database

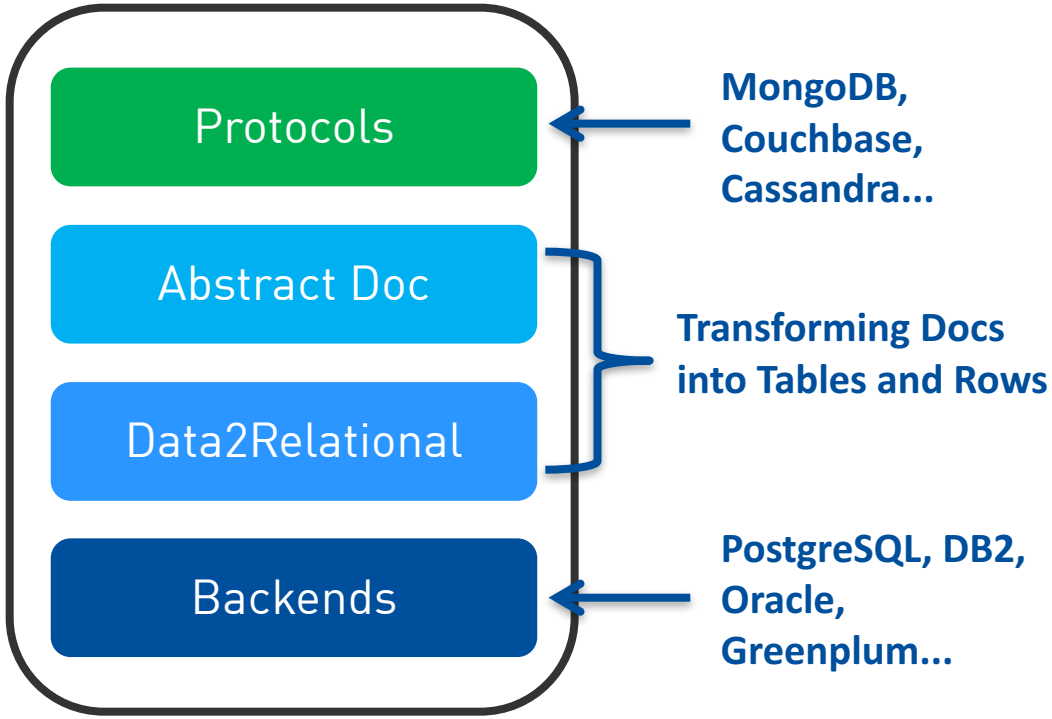

ToroDB
How it works?

ToroDB transforms documents to relational tables

> Data is stored in tables
> ToroDB analyzes every incoming document and separates metadata (schema) form data (tuples)
> 1+ tables per MongoDB collections
> ToroDB creates a RDBMS catalog schema per MongoDB database
> Dynamic and implicit schema generation

Full compatibility with MongoDB

> API programs, clients
> CRUD operations including UPDATE

ToroDB
How it works?

ToroDB can work as a secondary node on a MongoDB replica set

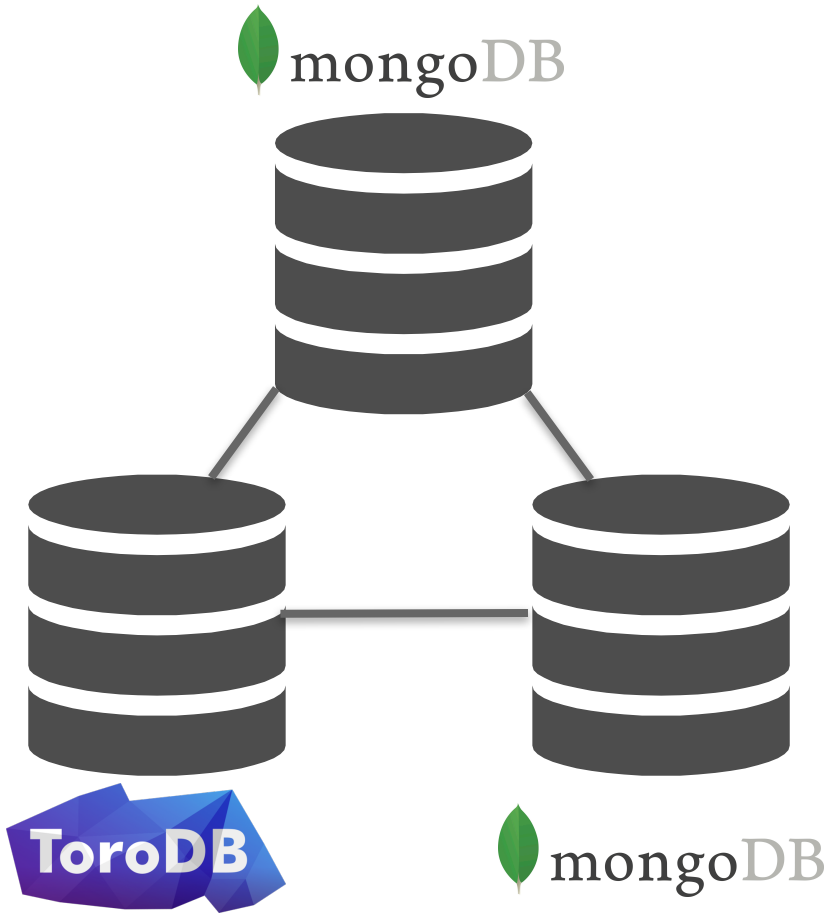

ToroDB
Why ToroDB?

Native SQL BI Connector

Data Integration Platform: SQL and NoSQL apps in the same RDBMS

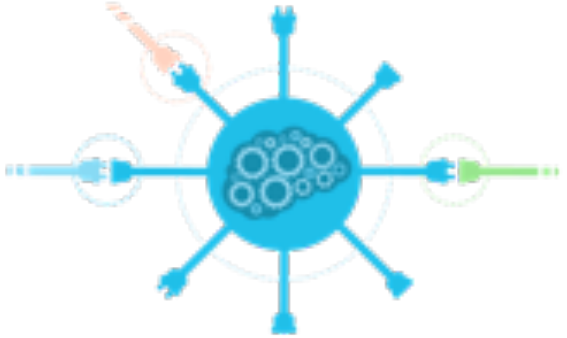

Apps: Write data with Mongo API, query with SQL!

Migration: from MongoDB to PostgreSQL

- Overview
- Prerequisites
- Configuration
- Demo

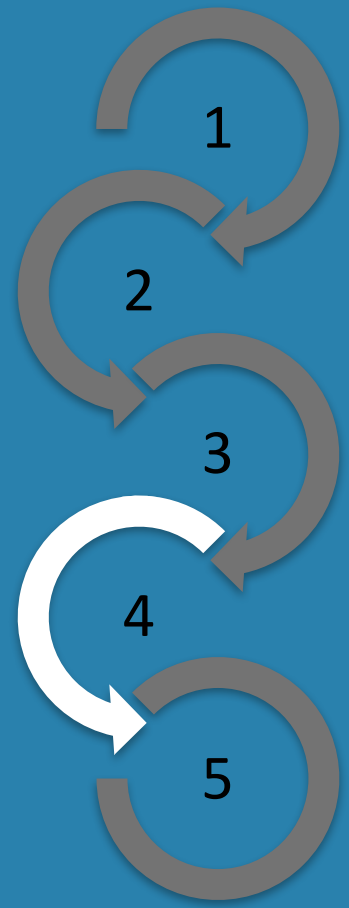

Migration: from MongoDB to PostgreSQL
Overview

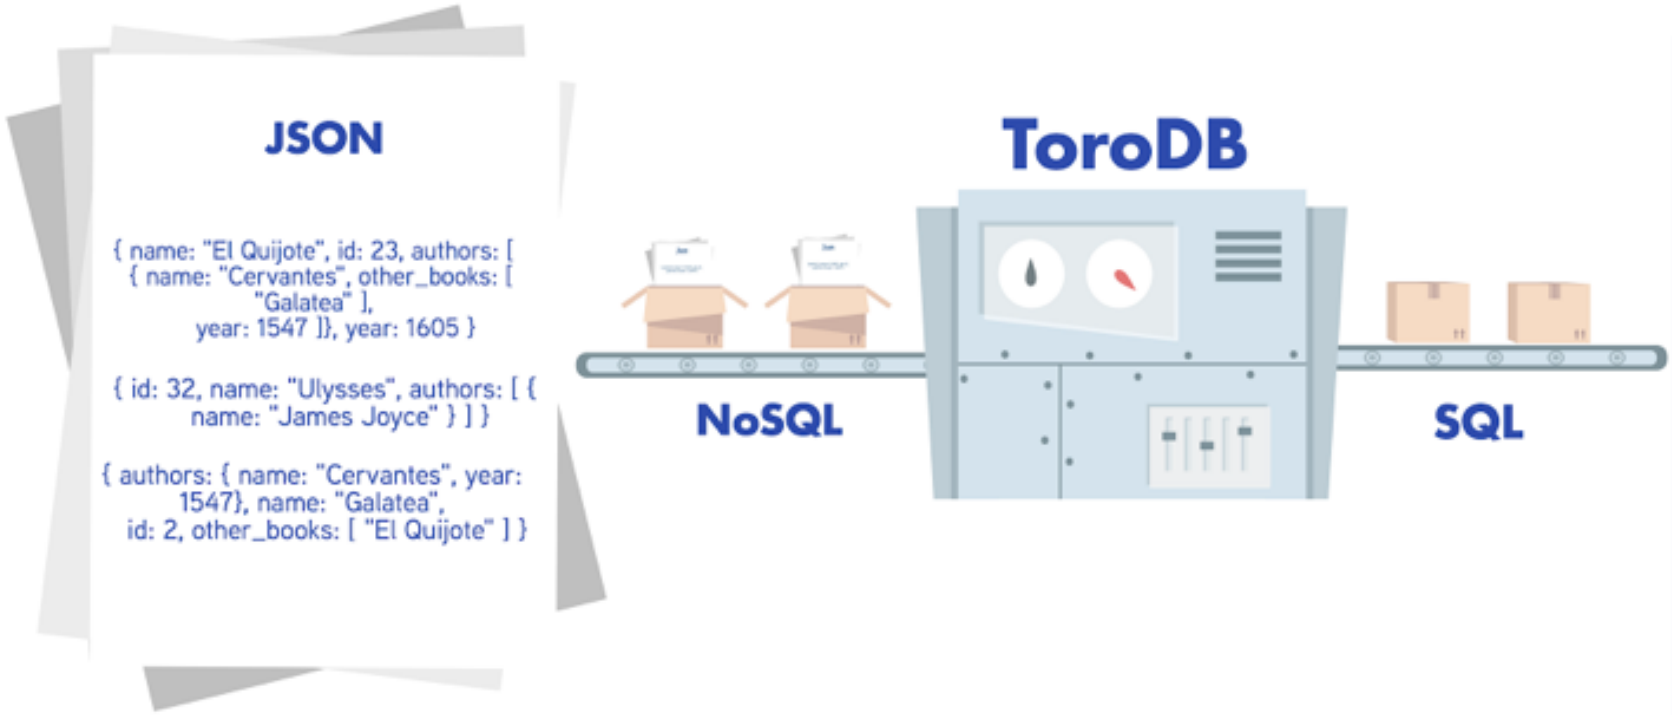

Author

name	year
Cervantes	1547
James Joyce	

Book

id	name	author	year
23	El Quijote	Cervantes	1605
2	Galatea	Cervantes	
32	Ulysses	James Joyce	

Migration: from MongoDB to PostgreSQL
Overview

ToroDB Stampeded uses MongoDB replica set oplog to keep track of the modifications in MongoDB

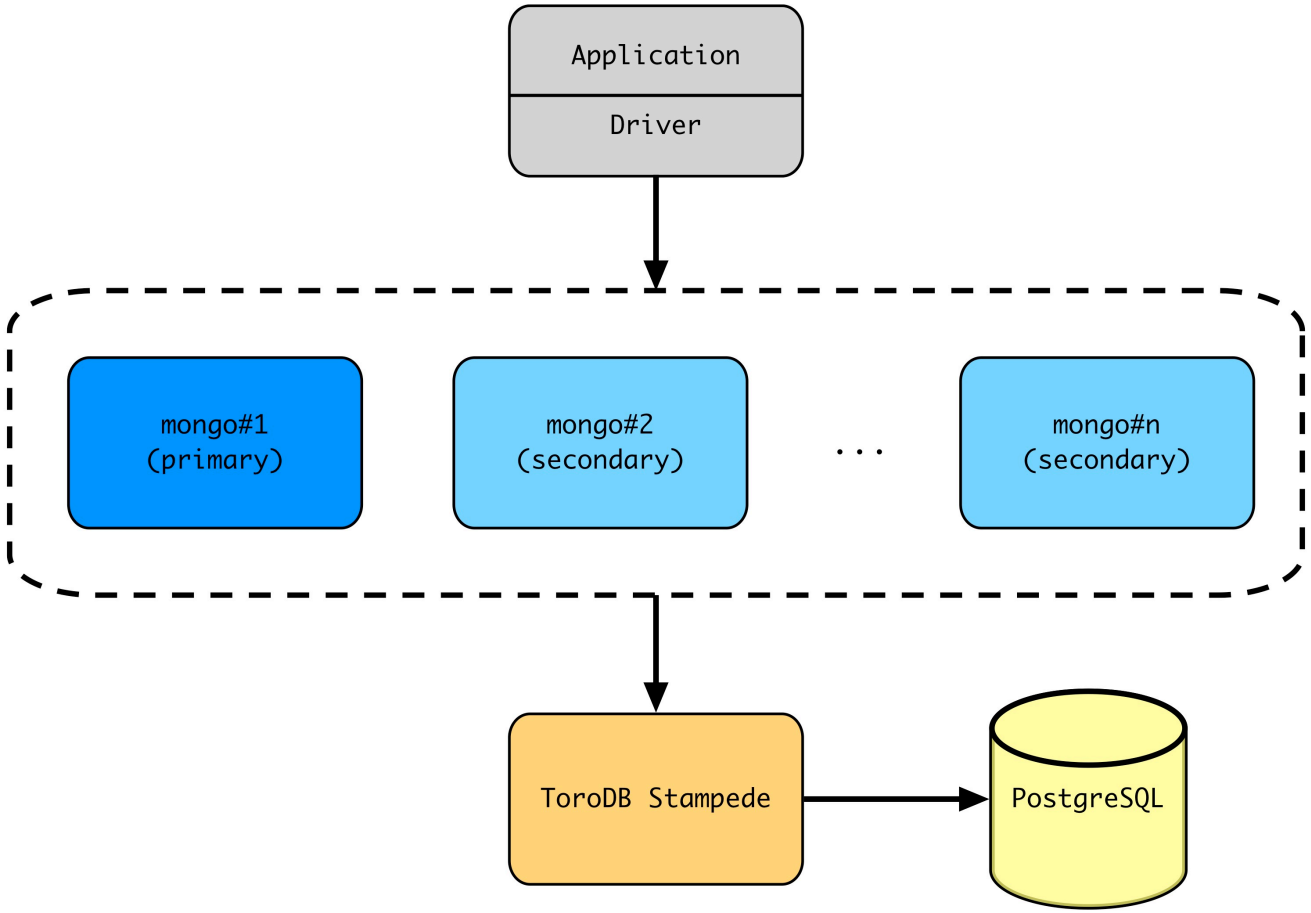

Migration: from MongoDB to PostgreSQL
Overview

During the replication ToroDB Stampeded transforms JSON documents into a relational schema

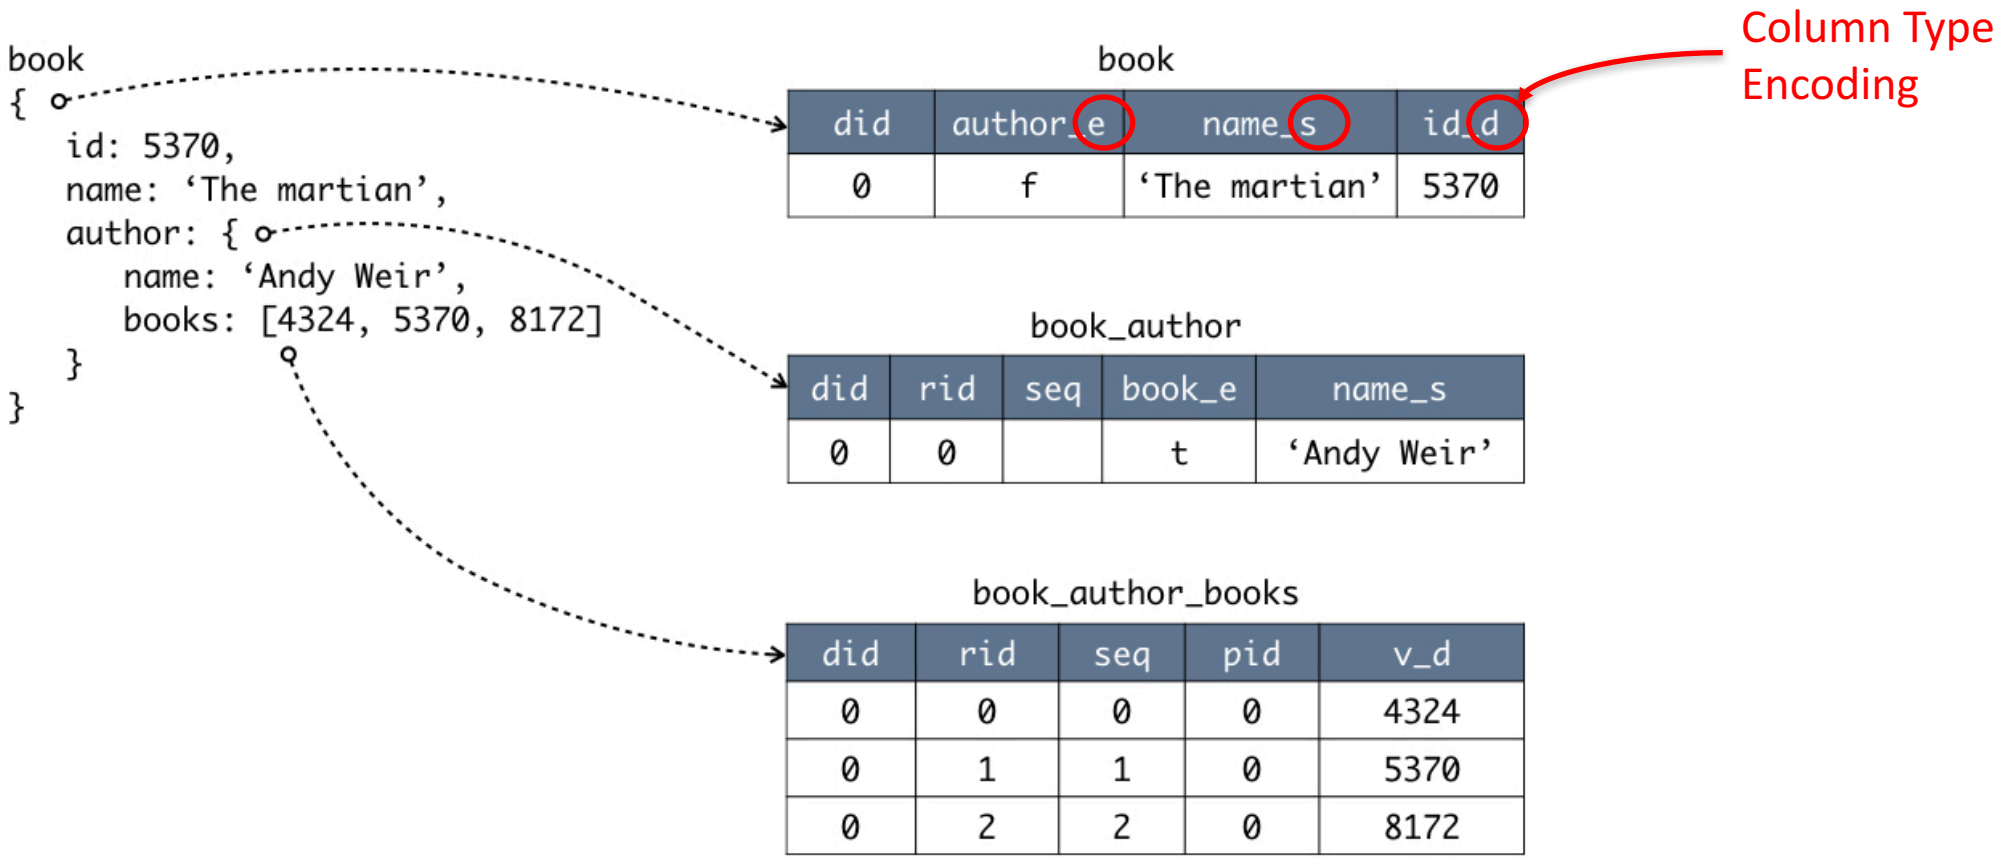

Migration: from MongoDB to PostgreSQL
Prerequisites

Runtime dependencies

Technology	Description
MongoDB	Install and start MongoDB instances with the replication features (dbi services best practices installation)
Replica sets configuration	ToroDB Stampede receives data from a MongoDB replica set. A single-node replica set is sufficient.
PostgreSQL	Install and start a PostgreSQL instance. Create a dedicated user and database. (dbi services best practices)
Java	ToroDB Stampede is written in Java so a Java Runtime Environment (JRE) required to run it. Java 8 is recommended.

Migration: from MongoDB to PostgreSQL
Configuration

Architecture

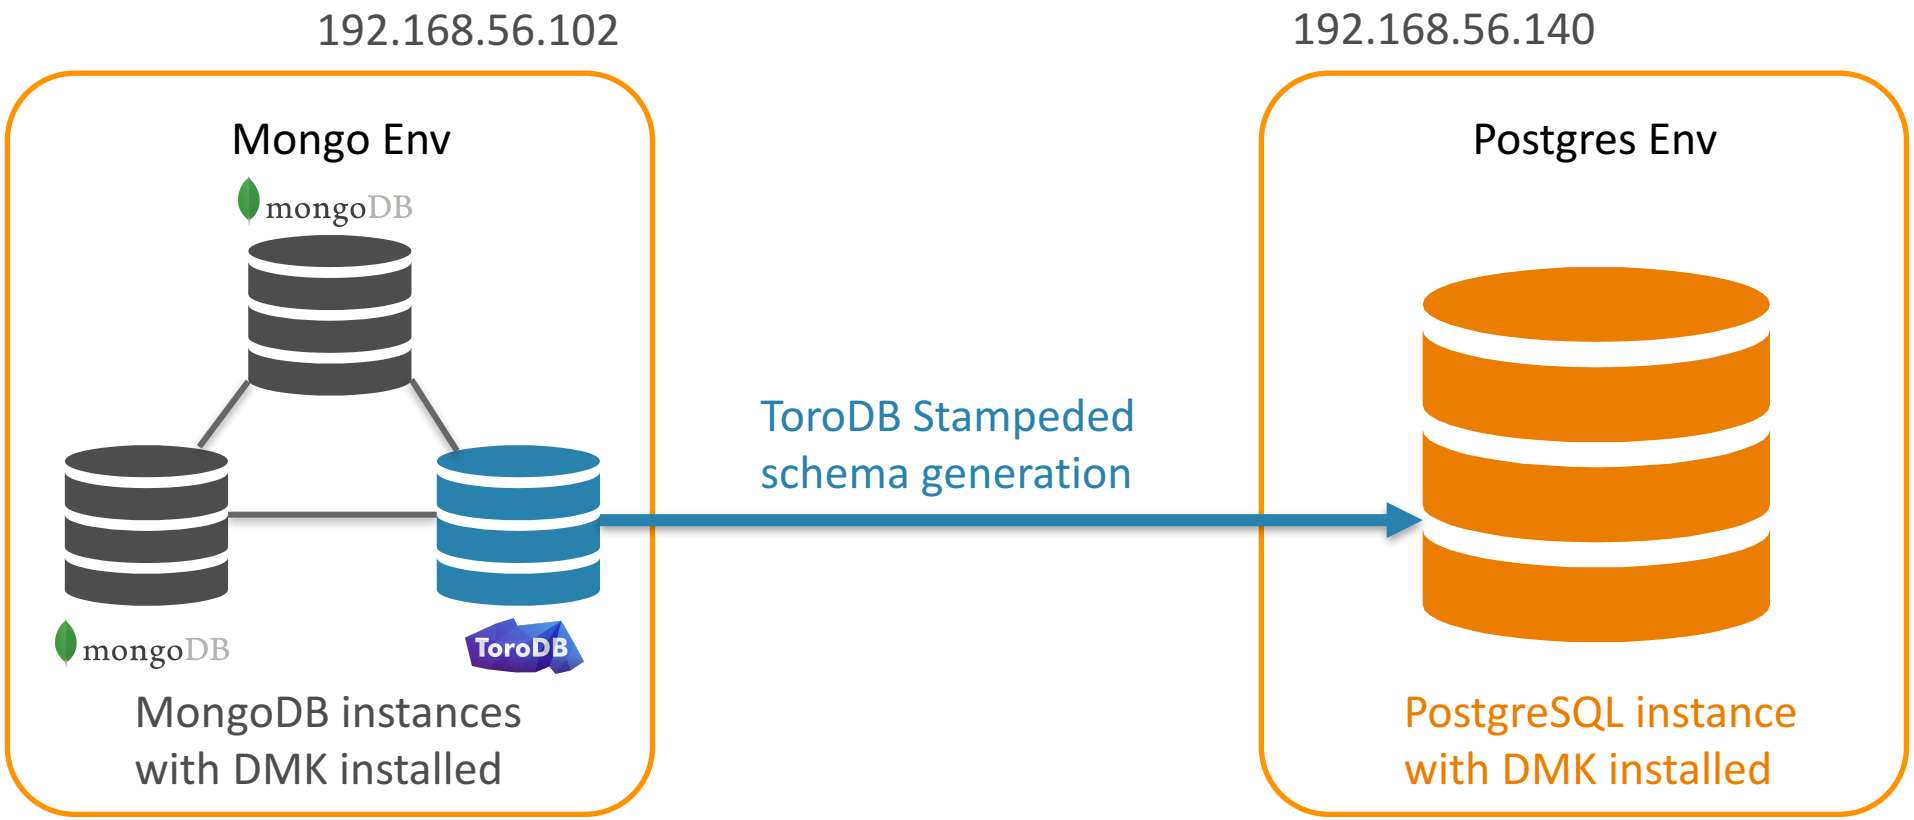

Migration: from MongoDB to PostgreSQL
Configuration

After installation, export $TOROHOME variable

```
mehdi@MacBook-Pro:/u00/app/torodb/ export TOROHOME /u00/app/torodb/torodb-stampede-1.0.0-beta2”

mehdi@MacBook-Pro:/u00/app/torodb/ echo $TOROHOME
/u00/app/torodb/torodb-stampede-1.0.0-beta2
```

Create and adapt the ToroDB configuration file (YAML)

```
mehdi@MacBook-Pro:/u00/app/torodb/ /u00/app/torodb/torodb-stampede-1.0.0-beta2/bin/torodb-stampede -l > ../conf/torodb.yaml

mehdi@MacBook-Pro:/u00/app/torodb/ vi ../conf/torodb.yaml
```

Migration: from MongoDB to PostgreSQL
Configuration

ToroDB Stampeded reads databases credentials from the .toropass file

Create the .toropass file in the home directory

```
mehdi@MacBook-Pro:/u00/app/torodb/ echo ”<host>:<port>:<database>:<user>:<PASSWD>" >
"$HOME/.toropass"

mehdi@MacBook-Pro:/u00/app/torodb/ chmod 0400 ”~/.toropass”
```

Migration: from MongoDB to PostgreSQL
Configuration

Custom configuration for ToroDB Stampeded on startup

> Command line options

> Configuration file

Recommended to use a configuration file

```
mehdi@MacBook-Pro:/u00/app/torodb/ ./torodb-stampede -c myconfiguration.yml
```

Print the current configuration (YAML)

```
mehdi@MacBook-Pro:/u00/app/torodb/ ./torodb-stampede -l
```

Migration: from MongoDB to PostgreSQL
Demo

Steps	Description
1	Java 8 (JRE) Installation and Configuration: http://download.oracle.com/otn-pub/java/jdk/8u131-b11/d54c1d3a095b4ff2b6607d096fa80163/jre-8u131-linux-x64.tar.gz
2	PostgreSQL Installation and Configuration: - Modify, if needed, PostgreSQL instance configuration: /u02/pgdata/PG1/postgresql.conf https://www.torodb.com/stampede/docs/1.0.0-beta2/configuration/postgresql-configuration-tips/ - Create user torodb and database torod (with password) - Test the connection with new user and databases - Adapt pg_hba.conf for new connections
3	MongoDB Installation and Configuration: - Configure and initialize MongoDB replication: replset = torodb - Check replication config. (primary and secondary) - Import data into MongoDB
4	ToroDB Stampede Installation and Configuration: - torodb.yaml configuration file creation - .toropass file creation for creation - Start ToroDB Stampeded

Migration: from MongoDB to PostgreSQL
Demo

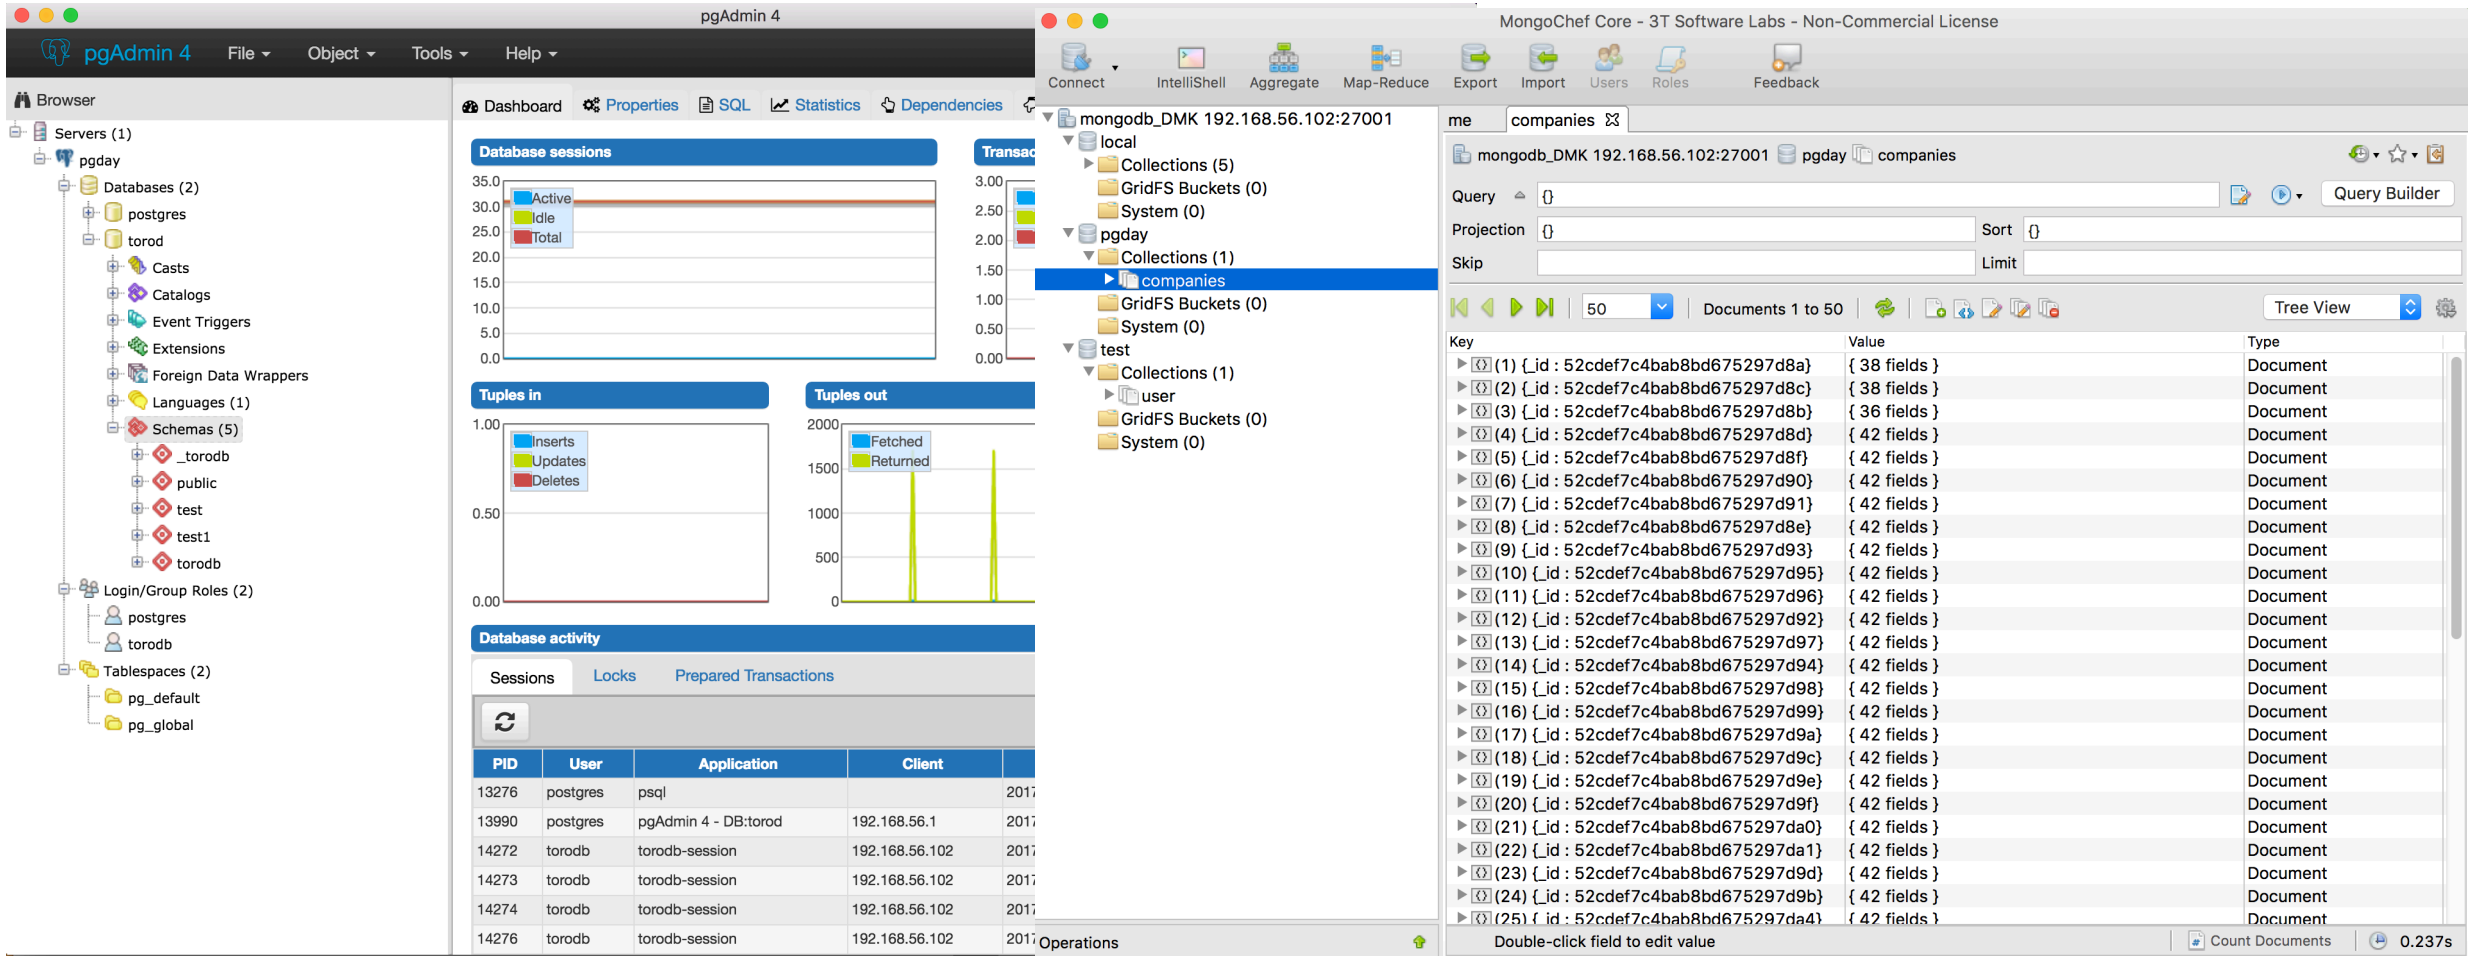

Conclusion

> Advantages vs Drawbacks

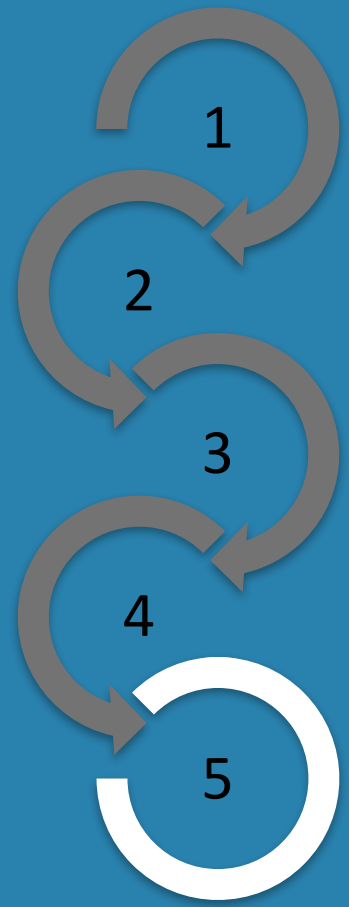

Conclusion
Advantages vs Drawbacks

- Fast and Powerful
- Dynamic schema generation
- Dynamic changes in the schema
- Open source and free!!
- Cross-platform: Linux, Windows

- Only supports PostgreSQL as RDBMS backend
- Need improvements for different schema in the same mongo collection

Any questions?

Please do ask!

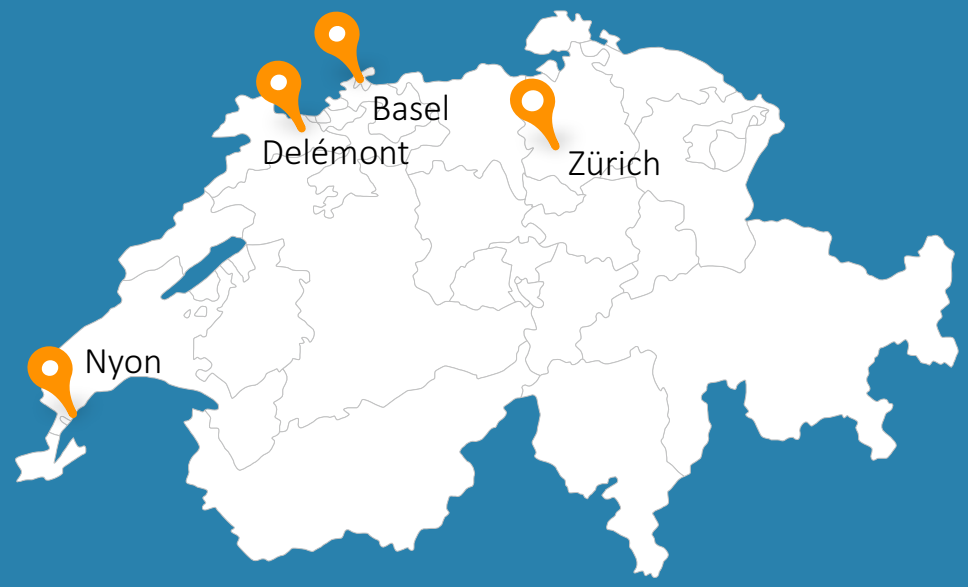

We would love to boost your IT-Infrastructure

How about you?