Migrating off of MongoDB to PostgreSQL

Álvaro Hernández Tortosa <aht@8kdata.com>

PgConf.Ru 2017

Who I am

- What we do @8Kdata:
  - ✓Creators of ToroDB.com, NoSQL & SQL database
  - ✓Database R&D, product development
  - ✓Training and consulting in PostgreSQL
  - ✓PostgreSQL Support

Twitter: @ahachete

Linkedin:
http://es.linkedin.com/in/alvarohernandeztortosa/

CEO, 8Kdata.com

Founder, President Spanish Postgres User Group

postgrespana.es ~ 750 members

Agenda

1. MongoDB limitations
2. Pain points when migrating from NoSQL to SQL
   - ✓ Stripe's MoSQL
   - ✓ EDB Mongo FDW
   - ✓ Quasar FDW
   - ✓ Mongo BI Connector, yet another PG FDW ;)
   - ✓ ETL tool
   - ✓ ToroDB Stampede
3. Benchmarks

MongoDB limitations

No ACID is bitter

- Atomic transactions only work within the same document

CART

```
[
...
   {
      user: 567,
      products: [
         {
            id: 47,
            units: 7,
         },
         {
            id: 318,
            units: 2,
         },
      …
]
```

ORDERS

```
[
...
   {
      user: 567,
      orders: [
         ...
         {
            orderId:
24658,
            products: [
               …
             ]
         }
      …
      ]
   }
...
]
```

★This operation is not atomic!!

MongoDB does *not* have consistent reads

David Glasser Follow

langton labs, @meteorjs, etc.

Jun 7, 2016 · 7 min read

MongoDB queries don't always return all matching documents!

When I query a database, I generally expect that it will return all the results that match my query. Recently, I was surprised to find that this isn't always the case for MongoDB. Specifically, if a document is updated while the query is running, MongoDB may not return it from the query—even if it matches both before and after the update! If you use MongoDB, you should be aware of this subtle edge case and make sure your queries don't fall victim to it.

https://blog.meteor.com/mongodb-queries-dont-always-return-all-matching-documents-654b6594a827#.fplxodagr

BI query performance issues

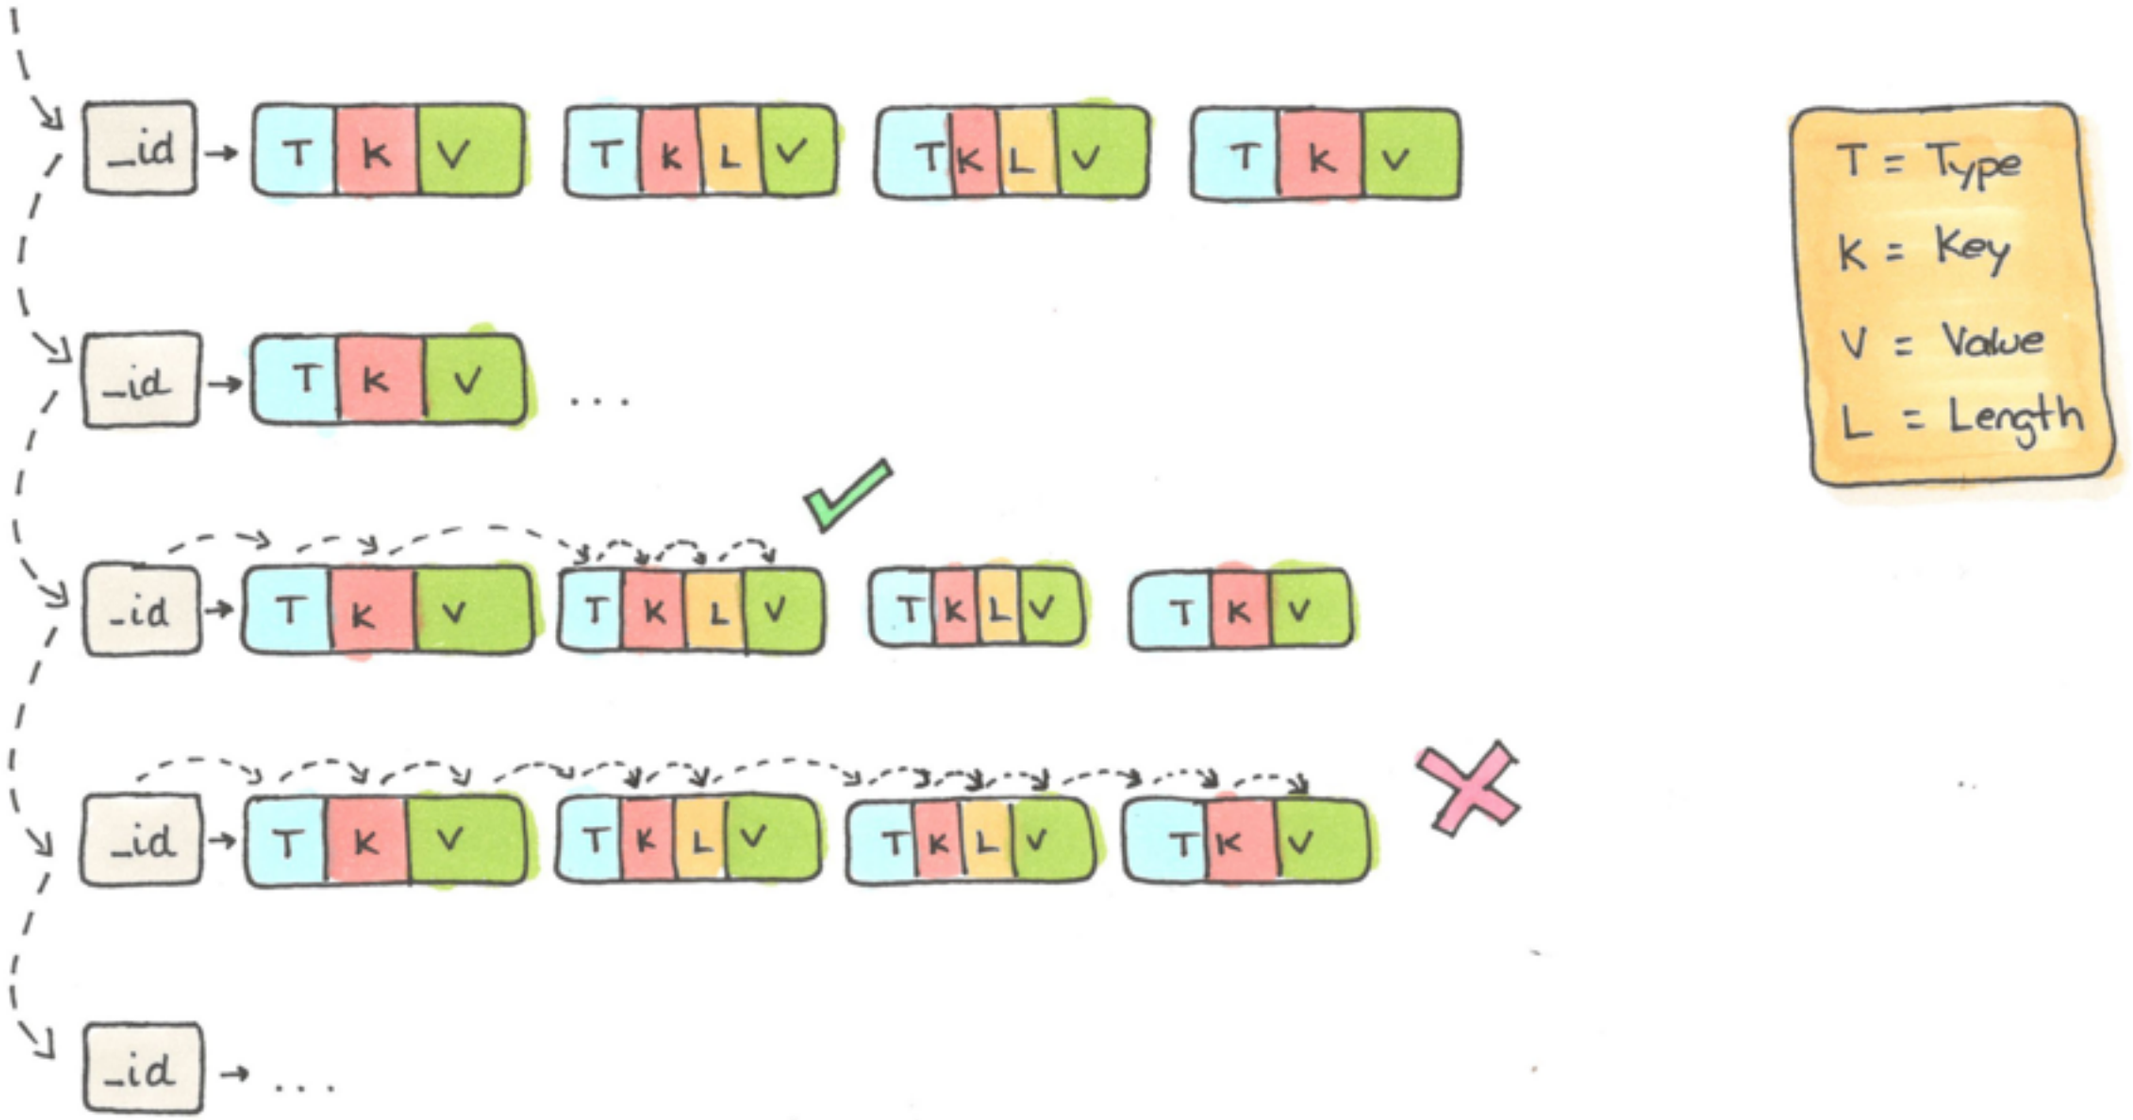

★ MongoDB aggregate query pattern

BI query performance issues

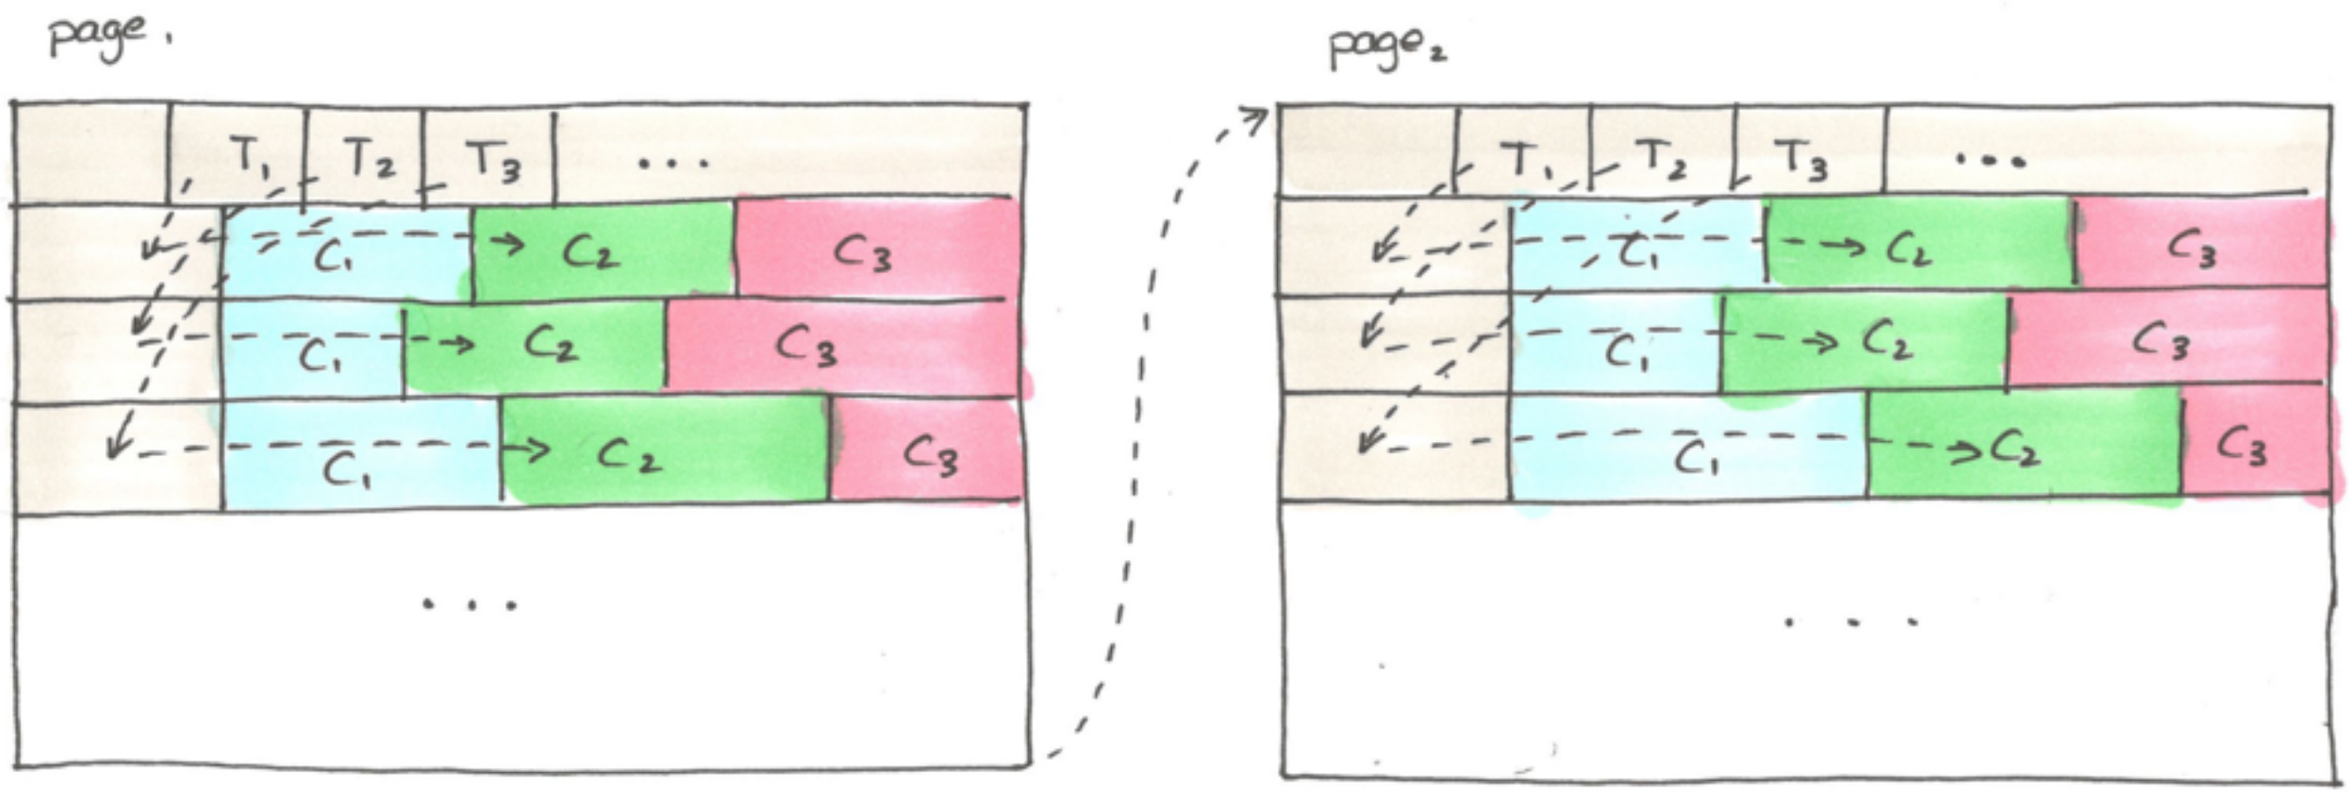

★ PostgreSQL aggregate query pattern

BI query performance issues

What if we use a columnar store?

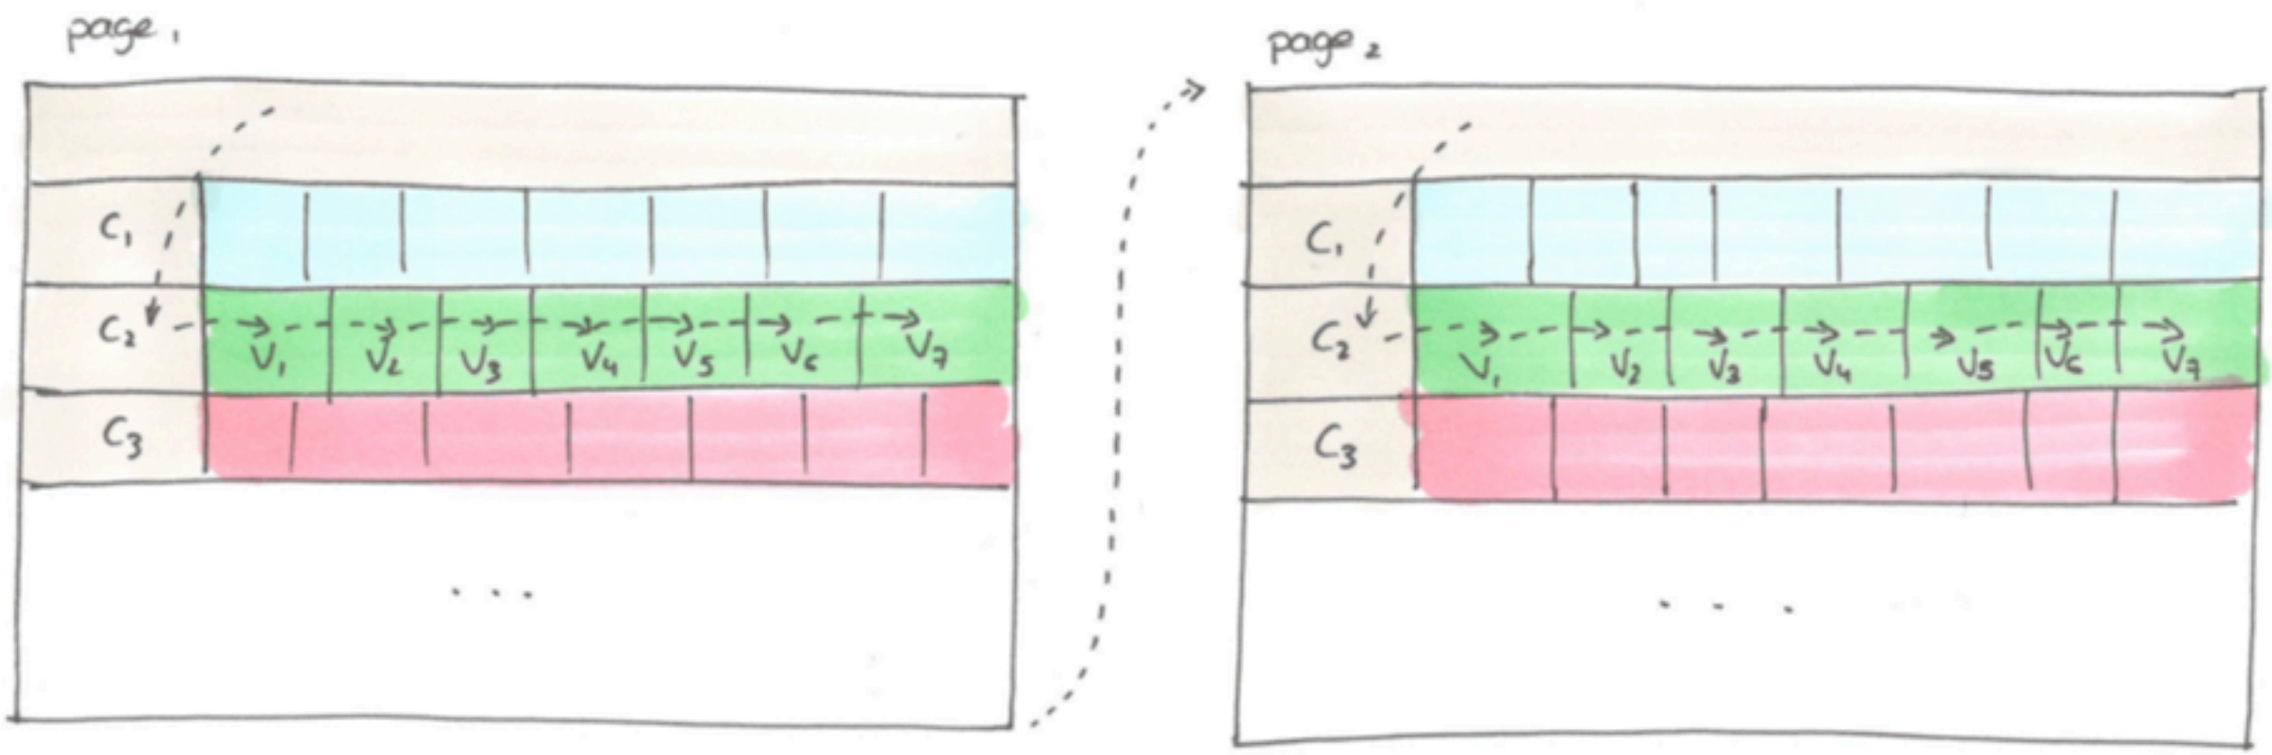

Taking it to the extreme

- We have developed a very difficult video game
- We store in our DB every game a user plays
- We also store every stage a user completes, including the time he had left
- Given that our stages are completed less than 1% of the times…
- How can we list all the players that completed a stage with more than 5 seconds left?

Taking it to the extreme

```
 [
{id:1, player_id:5, max_stage:1, timestamp: …, stagesComplete: {}},
{id:2, player_id:5, max_stage:1, timestamp: …, stagesComplete: {}},
{id:3, player_id:4, max_stage:2, timestamp: …, stagesComplete: {...}},
…
]
```

★MongoDB

Taking it to the extreme

Game				
id	player-id	stage	timestamp	.....
1	5	1	—	....
2	5	1	—	....
3	4	2	—	....
....	....	....	....	....
n	....	....	....	....

Stage-complete			
id	game-id	seconds -left	stage-id
1	3	2	1
2	957	4	1

★RDBMS

Measuring I/O with iotop required to answer a query

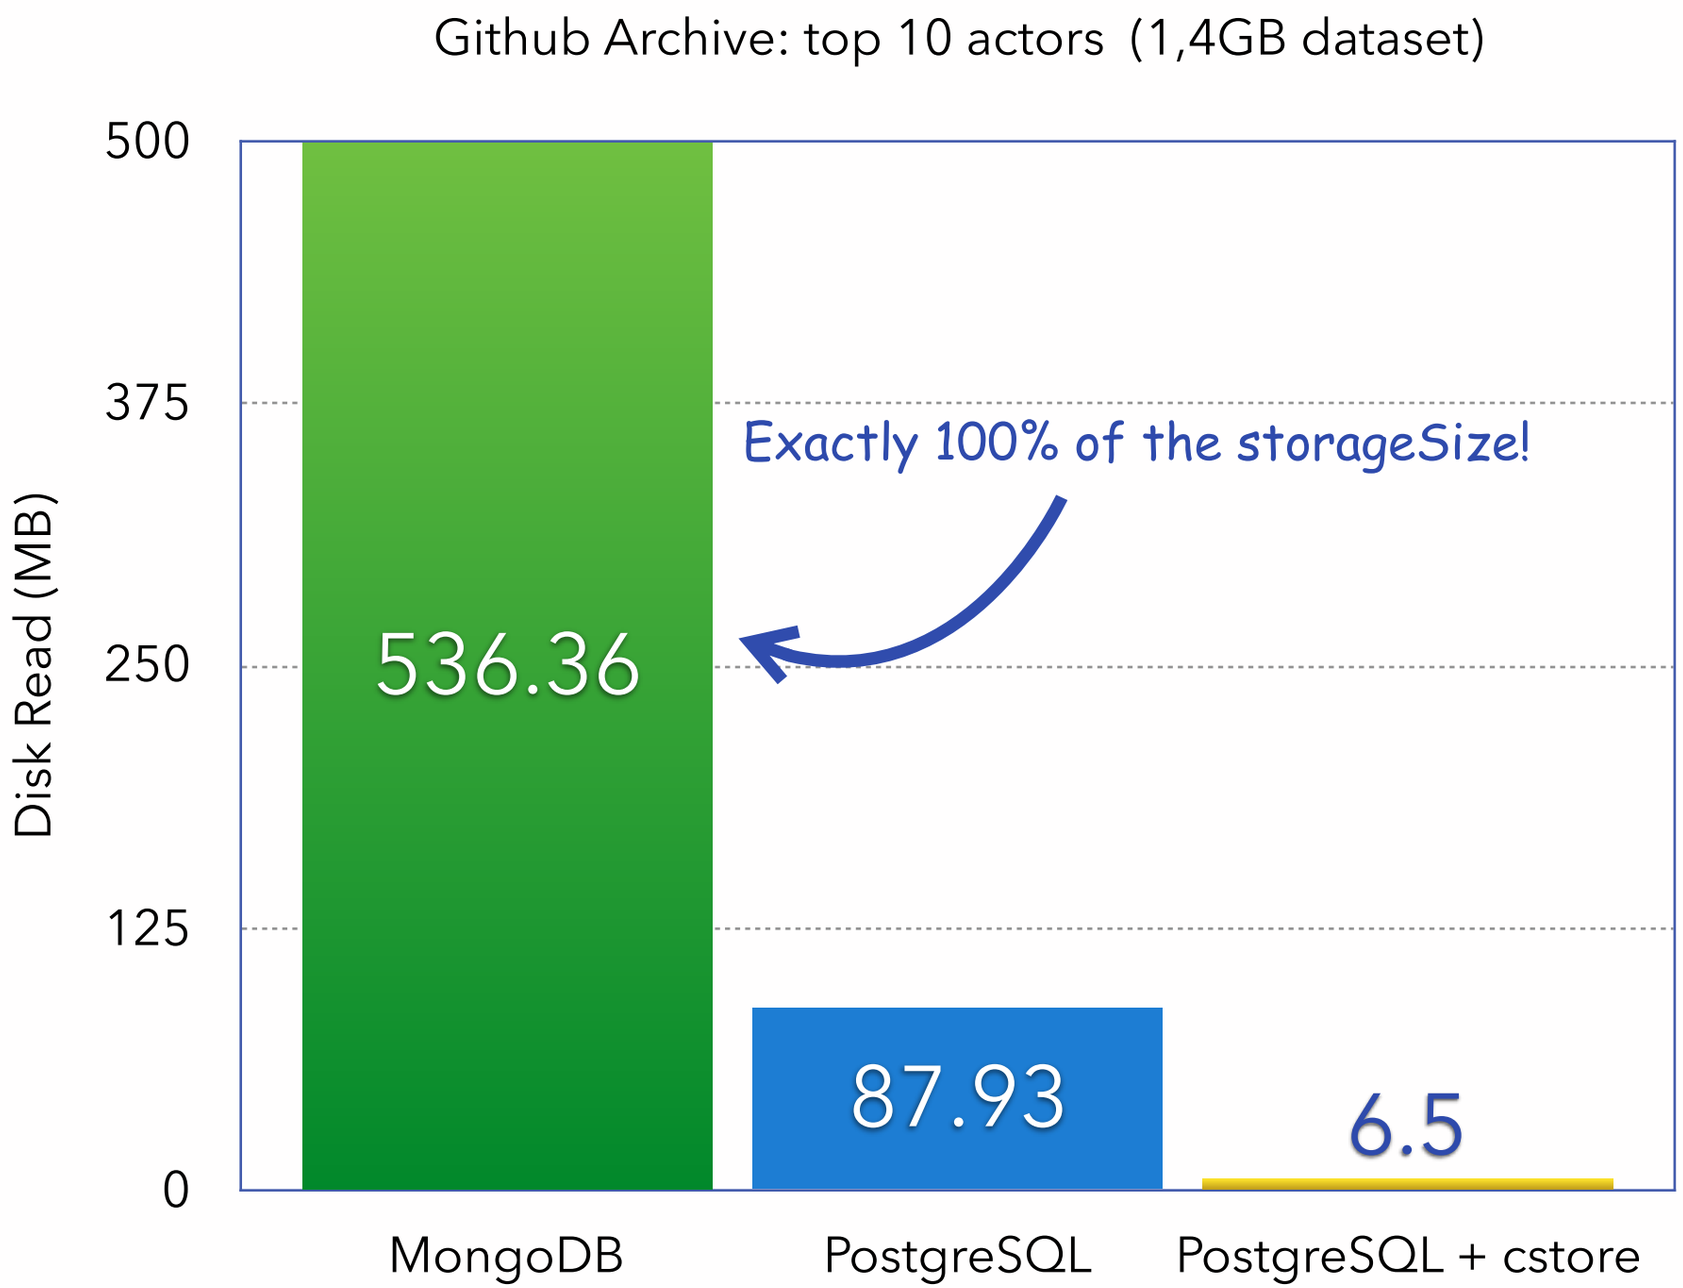

MongoDB storageSize: 536.37 MB

MongoDB size: 1410,35 MB

```
db.githubarchive.aggregate([
    { $group: {
        _id: '$actor.login', events: { $sum: 1 }
    } },
    { $sort: { events: -1 }},
    { $limit: 10 }
])

SELECT count(*), login
FROM actor
GROUP BY login
ORDER BY 1 DESC
LIMIT 10;
```

OK, let's migrate, but... how?

How do we generate a schema?

How do we infer the structure we should apply to the data?

How to react to dynamic changes in the schema?

Available tools in the market

Stripe's MoSQL

https://github.com/stripe/mosql

✓The user specifies the schema in a YAML file

✓It only cares for fields and mappings specified in the schema

✓So it doesn't care for changes in MongoDB schema

✓It creates tables, populates them with MongoDB data and optionally keeps track of Oplog changes and replicates them

But...

Stripe's MoSQL

- It doesn’t infer the schema, you have to tell it!
- It doesn’t support sharding
- MoSQL is not maintained anymore

EDB Mongo FDW

- https://github.com/EnterpriseDB/mongo_fdw

- Based on foreign data wrappers

- Works on MacOS X, Fedora and Ubuntu

But...

EDB Mongo FDW

- Needs a new MongoDB driver for C. Confusing and not self-contained installation.

- The user has to specify the schema:

```
-- load extension first time after install
CREATE EXTENSION mongo_fdw;

-- create server object
CREATE SERVER mongo_server
        FOREIGN DATA WRAPPER mongo_fdw
        OPTIONS (address '127.0.0.1', port '27017');

-- create user mapping
CREATE USER MAPPING FOR postgres
        SERVER mongo_server
        OPTIONS (username 'mongo_user', password 'mongo_pass');

-- create foreign table
CREATE FOREIGN TABLE warehouse(
        _id NAME,
        warehouse_id int,
        warehouse_name text,
        warehouse_created timestamptz)
SERVER mongo_server
        OPTIONS (database 'db', collection 'warehouse');

-- Note: first column of the table must be "_id" of type "NAME".

-- select from table
SELECT * FROM warehouse WHERE warehouse_id = 1;

           _id          | warehouse_id | warehouse_name |     warehouse_created
------------------------+--------------+----------------+---------------------------
53720b1904864dc1f5a571a0|            1 | UPS            | 12-DEC-14 12:12:10 +05:00
```

Quasar FDW

- https://github.com/quasar-analytics/quasar-fdw
- Based on foreign data wrappers
- Developed by SlamData
- Connects to Quasar (https://github.com/quasar-analytics/quasar)

But...

Quasar FDW

- The user has to specify the schema (with extended options compared to EDB Mongo FDW):

```
-- load extension first time after install

CREATE EXTENSION quasar_fdw;

-- load a certain version of extension

CREATE EXTENSION quasar_fdw WITH VERSION 'version';

-- create server object

CREATE SERVER quasar FOREIGN DATA WRAPPER quasar_fdw
       OPTIONS (server 'http://localhost:8080'
                ,path '/local/quasar'
                ,timeout_ms '1000'
                ,use_remote_estimate 'true'
                ,fdw_startup_cost '10'
                ,fdw_tuple_cost '0.01');

-- create foreign table

CREATE FOREIGN TABLE zips(
        city varchar,
        pop integer,
        state char(2),
        loc float[2])
    SERVER quasar
    OPTIONS (table 'zips');


-- create foreign table using column mappings

CREATE FOREIGN TABLE user_comments(
       user_id integer OPTIONS (map 'userId'),
       profile_name varchar OPTIONS (map 'profile.name'),
       age integer OPTIONS (map 'profile.age'),
       title varchar OPTIONS (map 'profile.title'),
       comment_id char(10) OPTIONS (map 'comments[*].id'),
       comment_text varchar OPTIONS (map 'comments[*].text'),
       comment_reply_to_profile integer OPTIONS (map 'comments[*].replyTo[0]'),
       comment_reply_to_comment char(10) OPTIONS (map 'comments[*].replyTo[1]'),
       comment_time date OPTIONS (map 'comments[*]."time"'))
    SERVER quasar
        OPTIONS (table 'user_comments');

-- select from table

SELECT * FROM zips LIMIT 10;

SELECT city, pop FROM zips WHERE pop % 2 = 1 LIMIT 10;

SELECT loc[1] AS lat, loc[2] AS long FROM zips LIMIT 10;
```

MongoDB BI Connector

- https://docs.mongodb.com/bi-connector/master/
- Developed by MongoDB themselves
- Effectively allows to query mongo in a SQL-ish way

But...

MongoDB BI Connector

- How do they do it? You don't know it.
- Harder to detect bugs
- They're part of bigger solutions which may or may not fit you
- If you need Enterprise version, you are paying for the whole solution

Classical ETL tools

- Some ETL tools provide this conversion

  ✓Kettle (Pentaho Data Integration)

  ✓Stitch

- Both have free and paid license modes

  But...

Classical ETL tools

- How do they do it? You don't know it.
- Harder to detect bugs
- They're part of bigger solutions which may or may not fit you

✓If you need Enterprise version, you are paying for the whole solution

ToroDB Stampede

ToroDB: from MongoDB to PostgreSQL

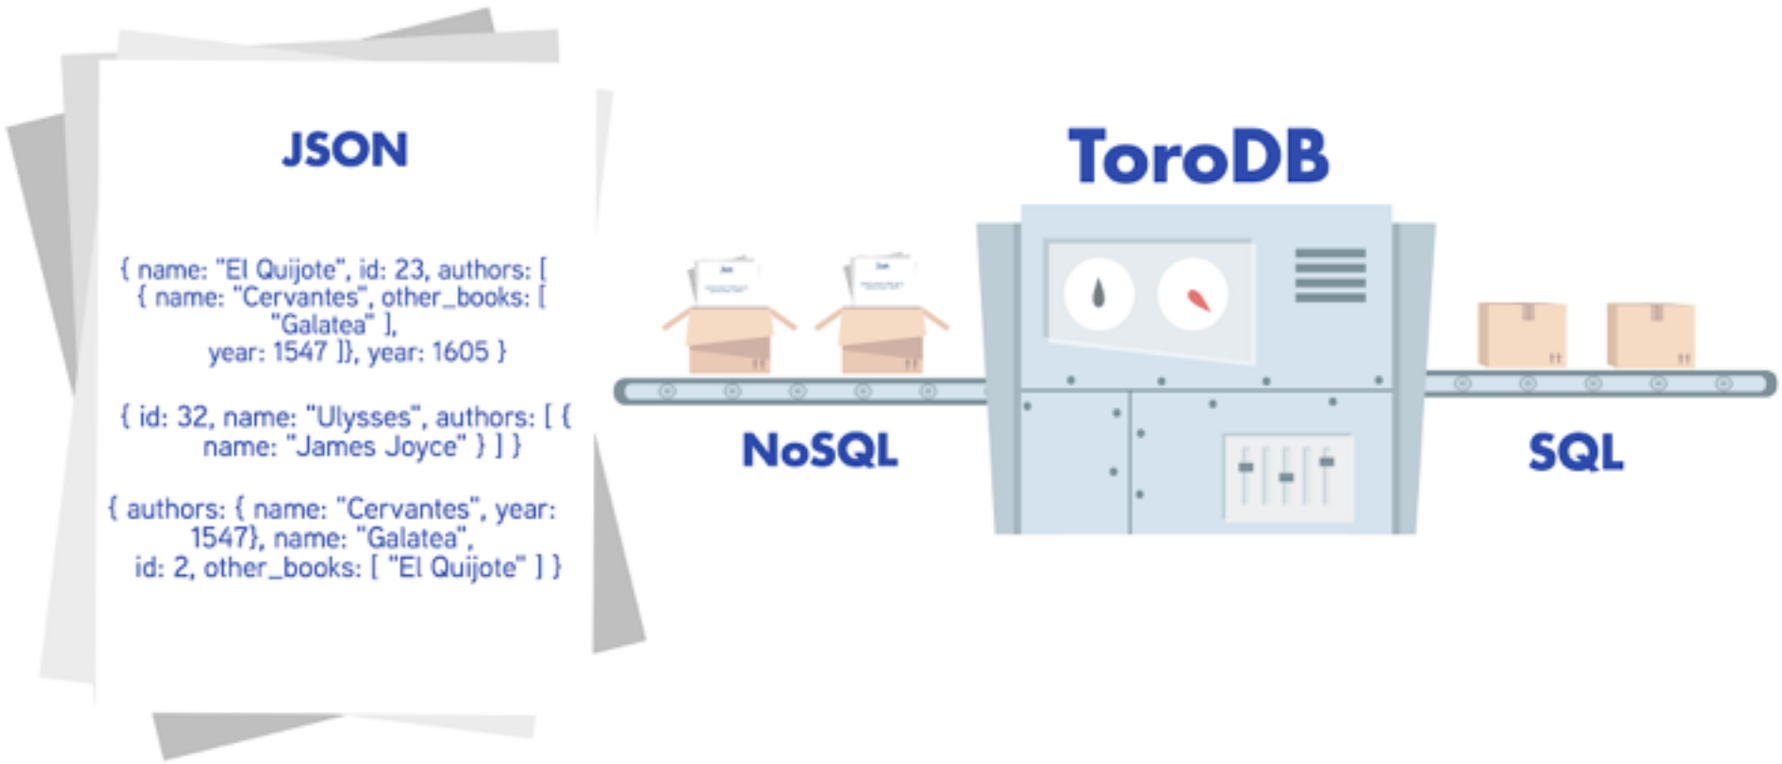

Author	
name	year
Cervantes	1547
James Joyce	

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

How do we generate the schema?

★There's no need to generate it, it's implicit!

```
{
    name: “Alice”,
    address: { street: “Wonderland”, number: 42 },
    hobbies: [ “reading”, “talking to rabbits” ]
}
```

- implicit schema

How do we generate the schema?

- ToroDB analyzes every incoming document and separates metadata (schema) from data (tuples)

- With that schema info, ToroDB creates 1+ tables per MongoDB collection

- ToroDB also creates a RDBMS catalog schema per MongoDB database

How do we generate the schema?

```
SELECT * FROM torodb.people;
```

did	name_s	_id_x
1	Alice	\x01d4a63dc6944f37a61cf3d6

```
SELECT * FROM torodb.people_address;
```

did	name_s	seq	number_d	street_s
1	1		42	Wonderland

```
SELECT * FROM torodb.people_hobbies;
```

did	rid	seq	value_s
1	1	0	“reading”
1	2	1	“talking to rabbits”

```
database: torodb
collection: people
{
  name: “Alice”,
  _id: (implicit val),
  address: {
    street:
      “Wonderland”,
    number: 42},
  },
  hobbies: [
    “reading”,
    “talking to
     rabbits”
  ]
}
```

How to react to dynamic changes in the schema?

```
[
        { name: “Alice” },
        { name: “Harry”, surname: “Potter” }
]
```

★ What happens if a new document in the collection brings new fields?

How to react to dynamic changes in the schema?

Remember: nulls take little or no space in PostgreSQL

did	_id_x	name_s	surname_s
1	\x01d4a63dc6944f37a61cf3d6	Alice	
2	\x21f4f563dc6944f37a61cf436	Harry	Potter

★ToroDB automatically performs:

✓`ALTER TABLE ... ADD COLUMN surname_s text`

How to react to dynamic changes in the schema?

```
[
    { name: “Harry”, surname: “Potter” }
    { name: “Henry”, surname: 8 },
]
```

★ What happens if a new document in the collection brings type conflicts?

How to react to dynamic changes in the schema?

Remember: nulls take no space in PostgreSQL

did	_id_x	name_s	surname_s	surname_i
1	\x01d4a63dc6944f37a61cf3d6	Harry	Potter	
2	\x21f4f563dc6944f37a61cf436	Henry		8

★ Type conflicts are resolved by suffixing column names with the type of the data: "s" for text, "i" for integer and so forth

ToroDB Architecture

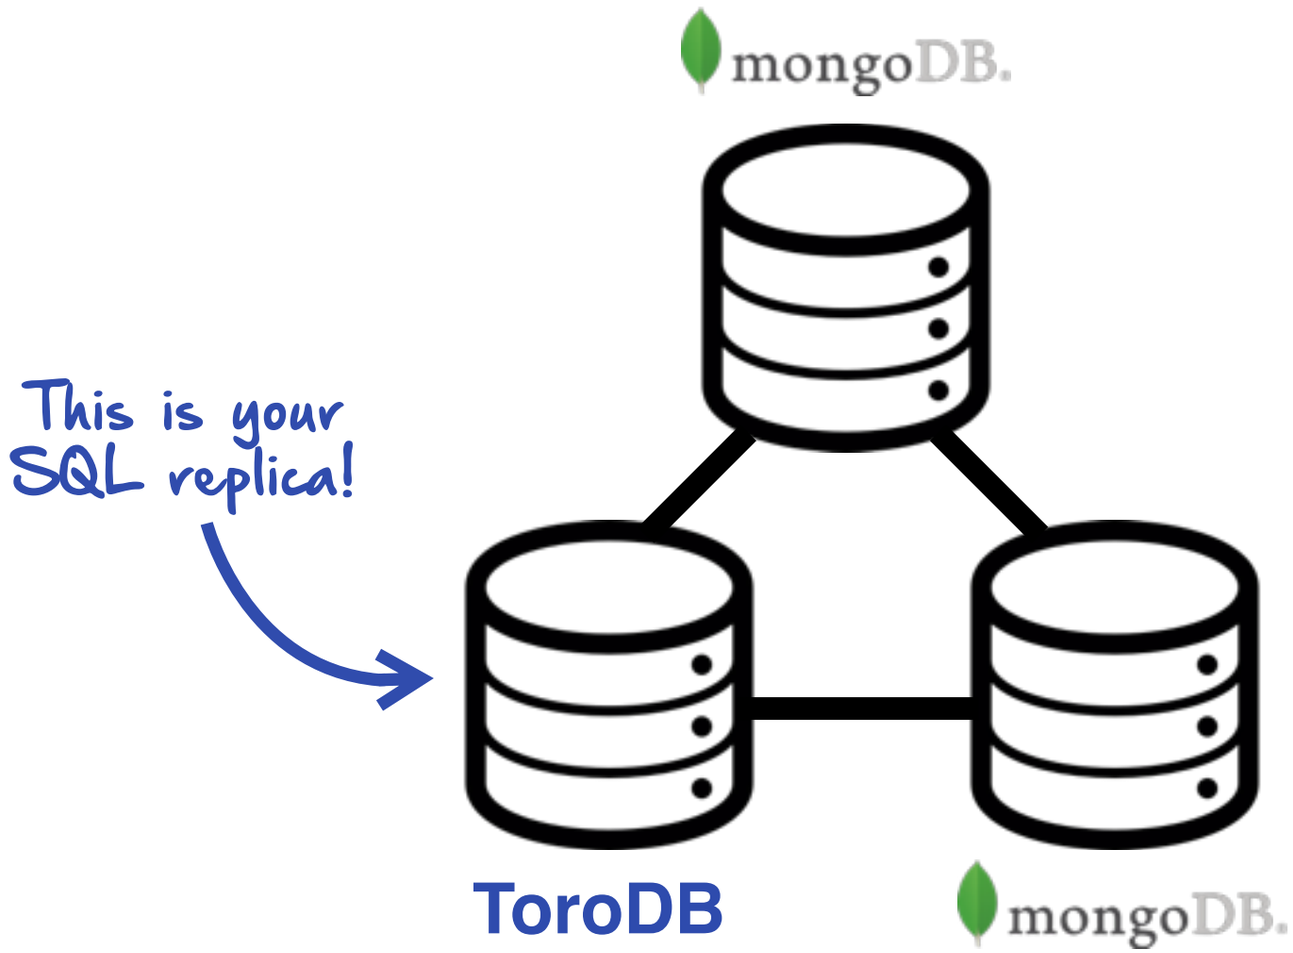

ToroDB works as a secondary node on a MongoDB replica set

ToroDB Architecture

- ToroDB works as a secondary node on a replica set

- This node has two main modes:
  - ✓ First, it imports all MongoDB data into PostgreSQL
  - ✓ Then it stays tuned to changes in MongoDB Oplog and applies them in PostgreSQL

ToroDB Architecture

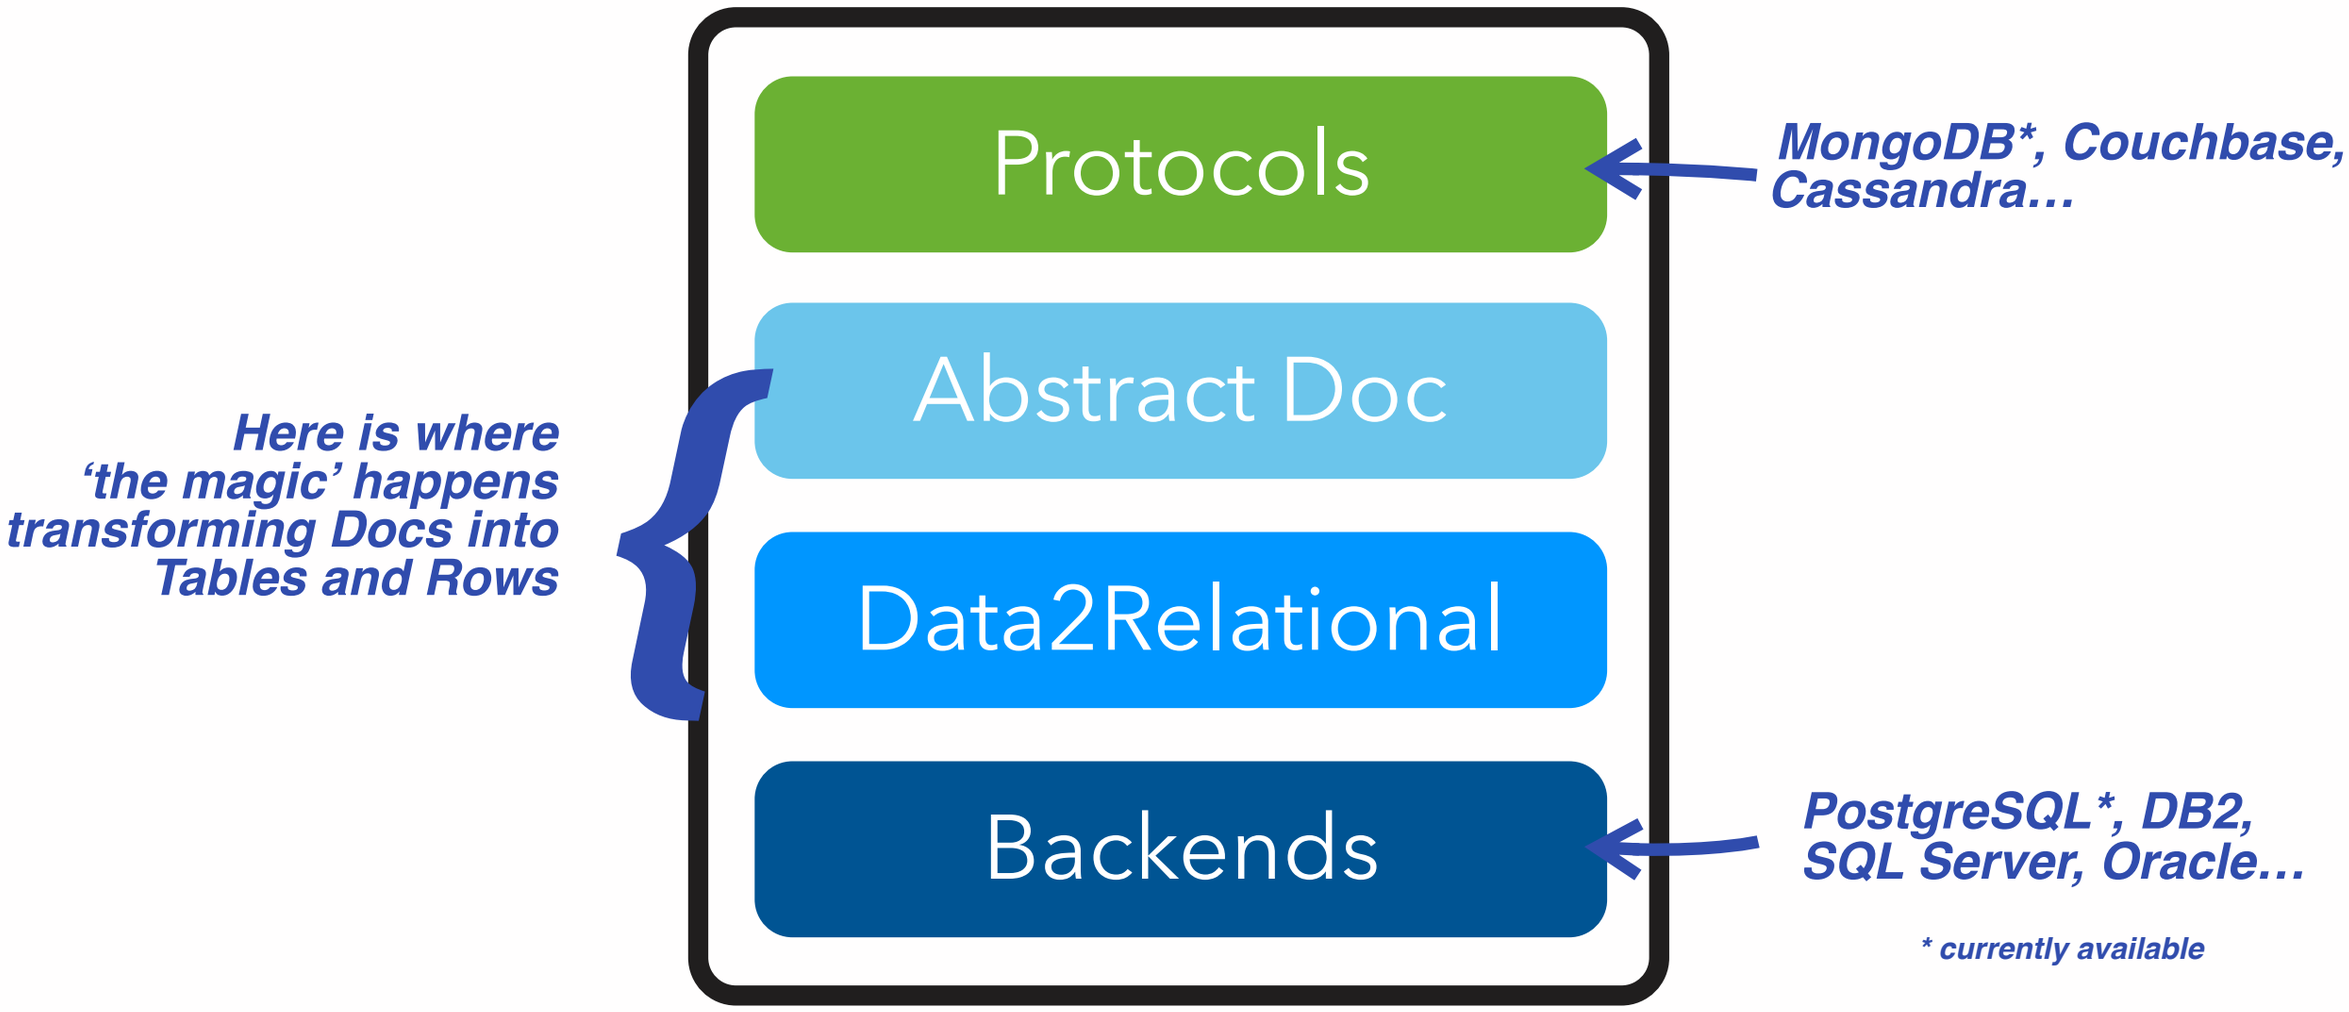

ToroDB offers the interface of any NoSQL database, but persisting data in tables and rows within a SQL database.

ToroDB other use cases

Native SQL BI Connector

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

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

Benchmarks

Building the benchmarks

- Data from Github archive
- Timing three different environments that contain the same data:
  - ✓ MongoDB
  - ✓ MongoDB with 3 shards
  - ✓ PostgreSQL (obtained by migrating MongoDB environment via ToroDB)
- We run 6 queries (PostgreSQL queries being the relational equivalent to MongoDB queries)
  - ✓ Available at https://gist.github.com/ahachete

Github 500GB

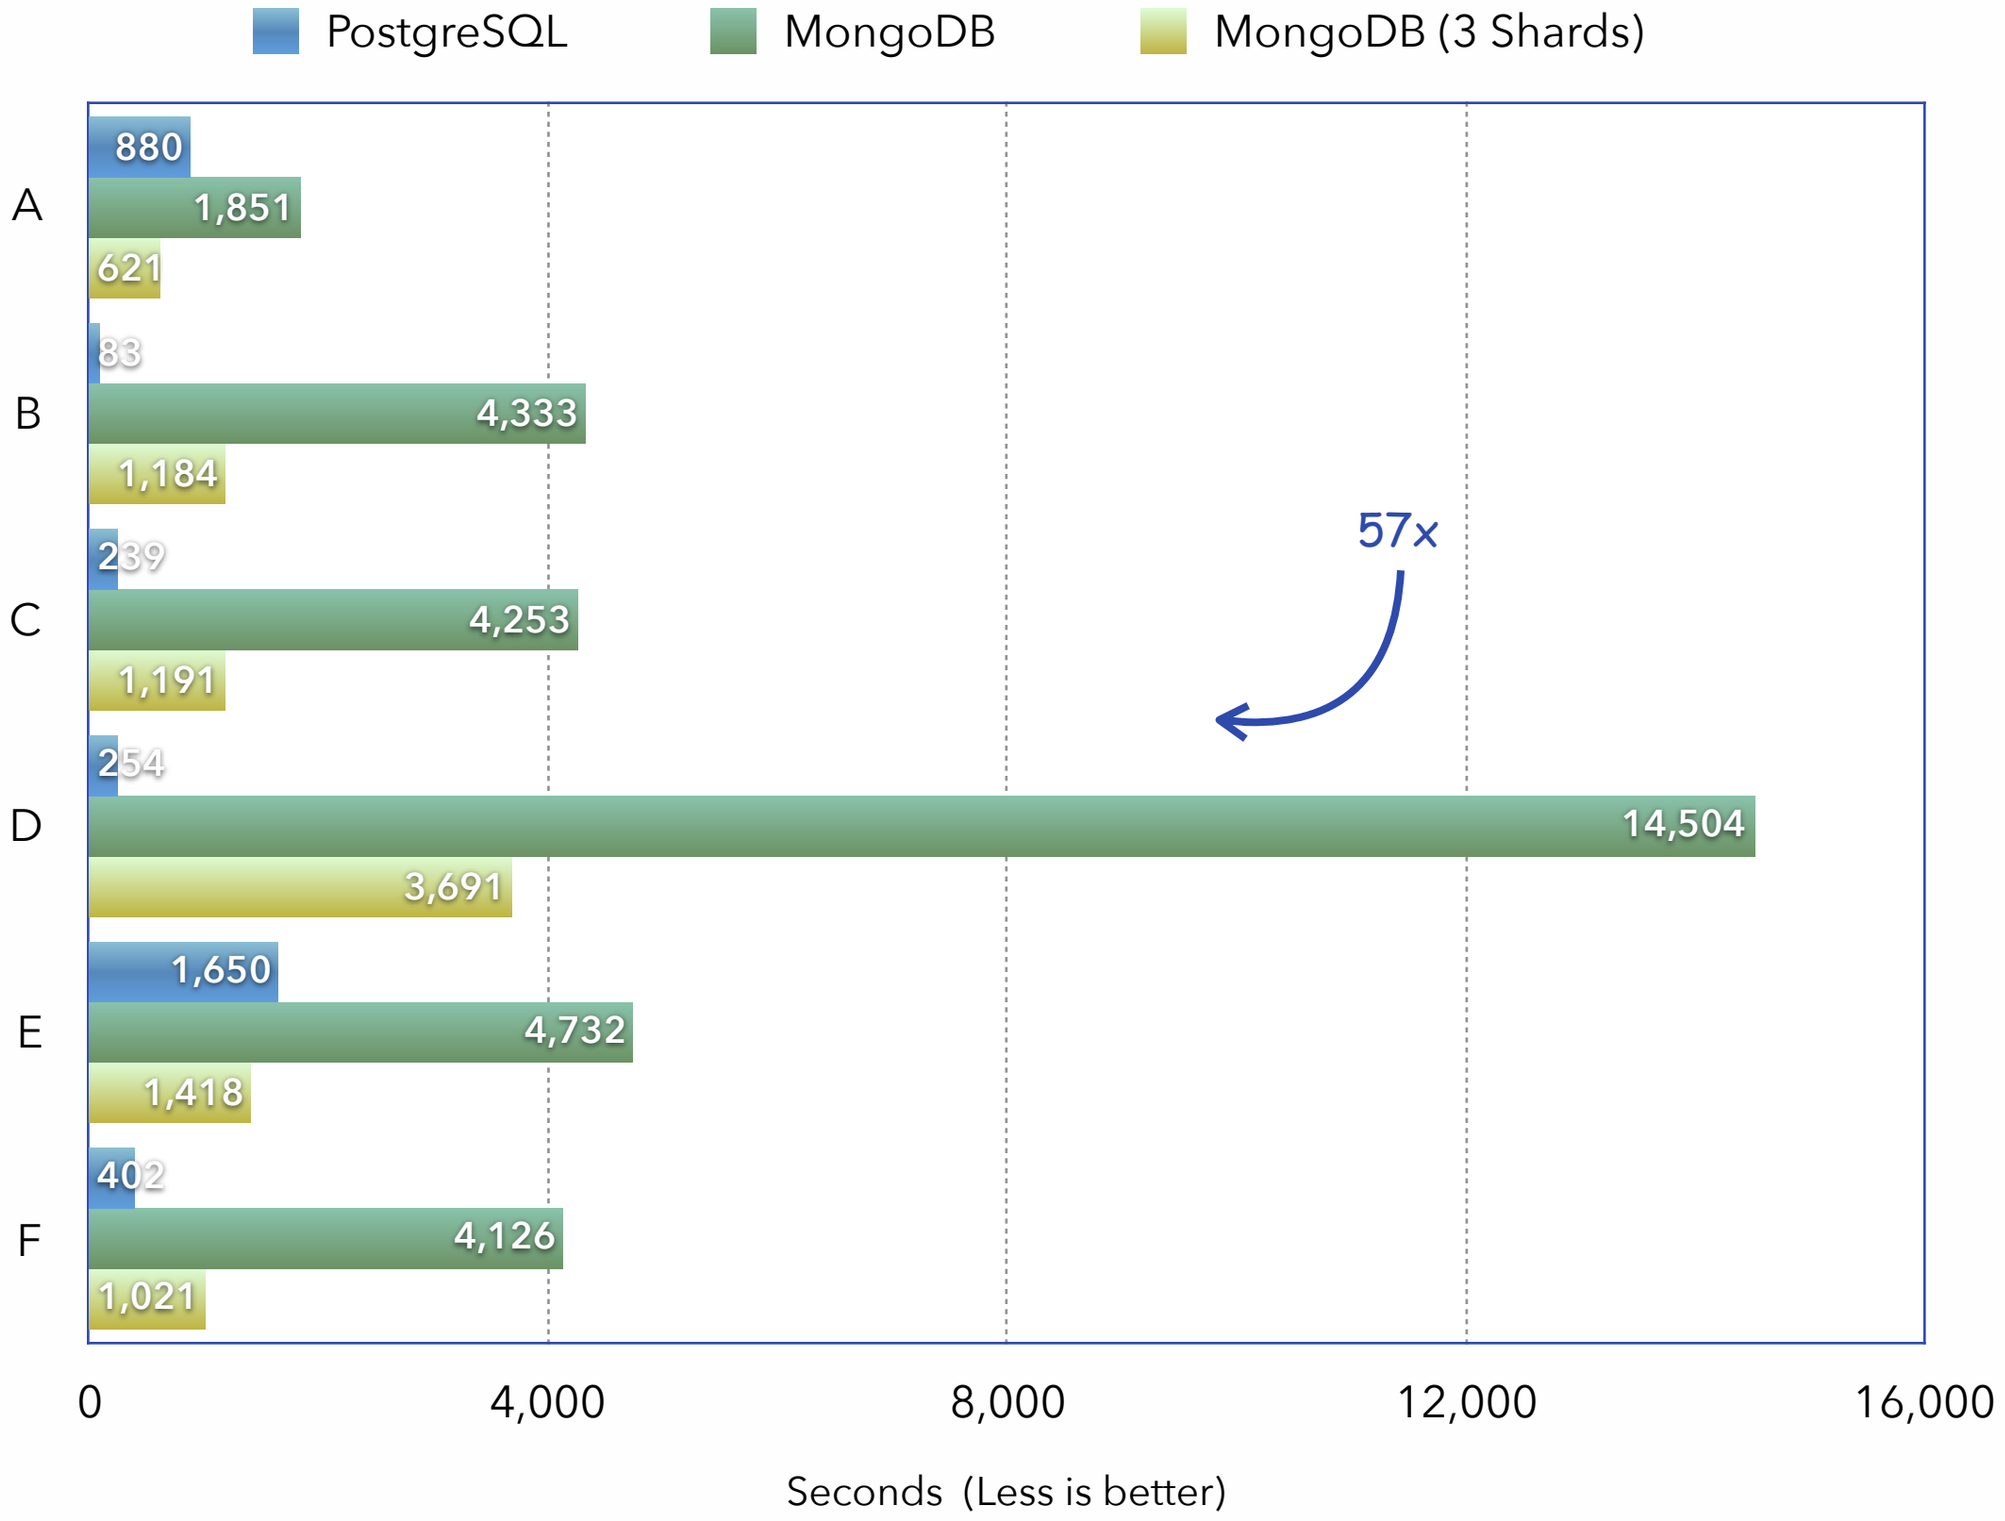

Github 100GB

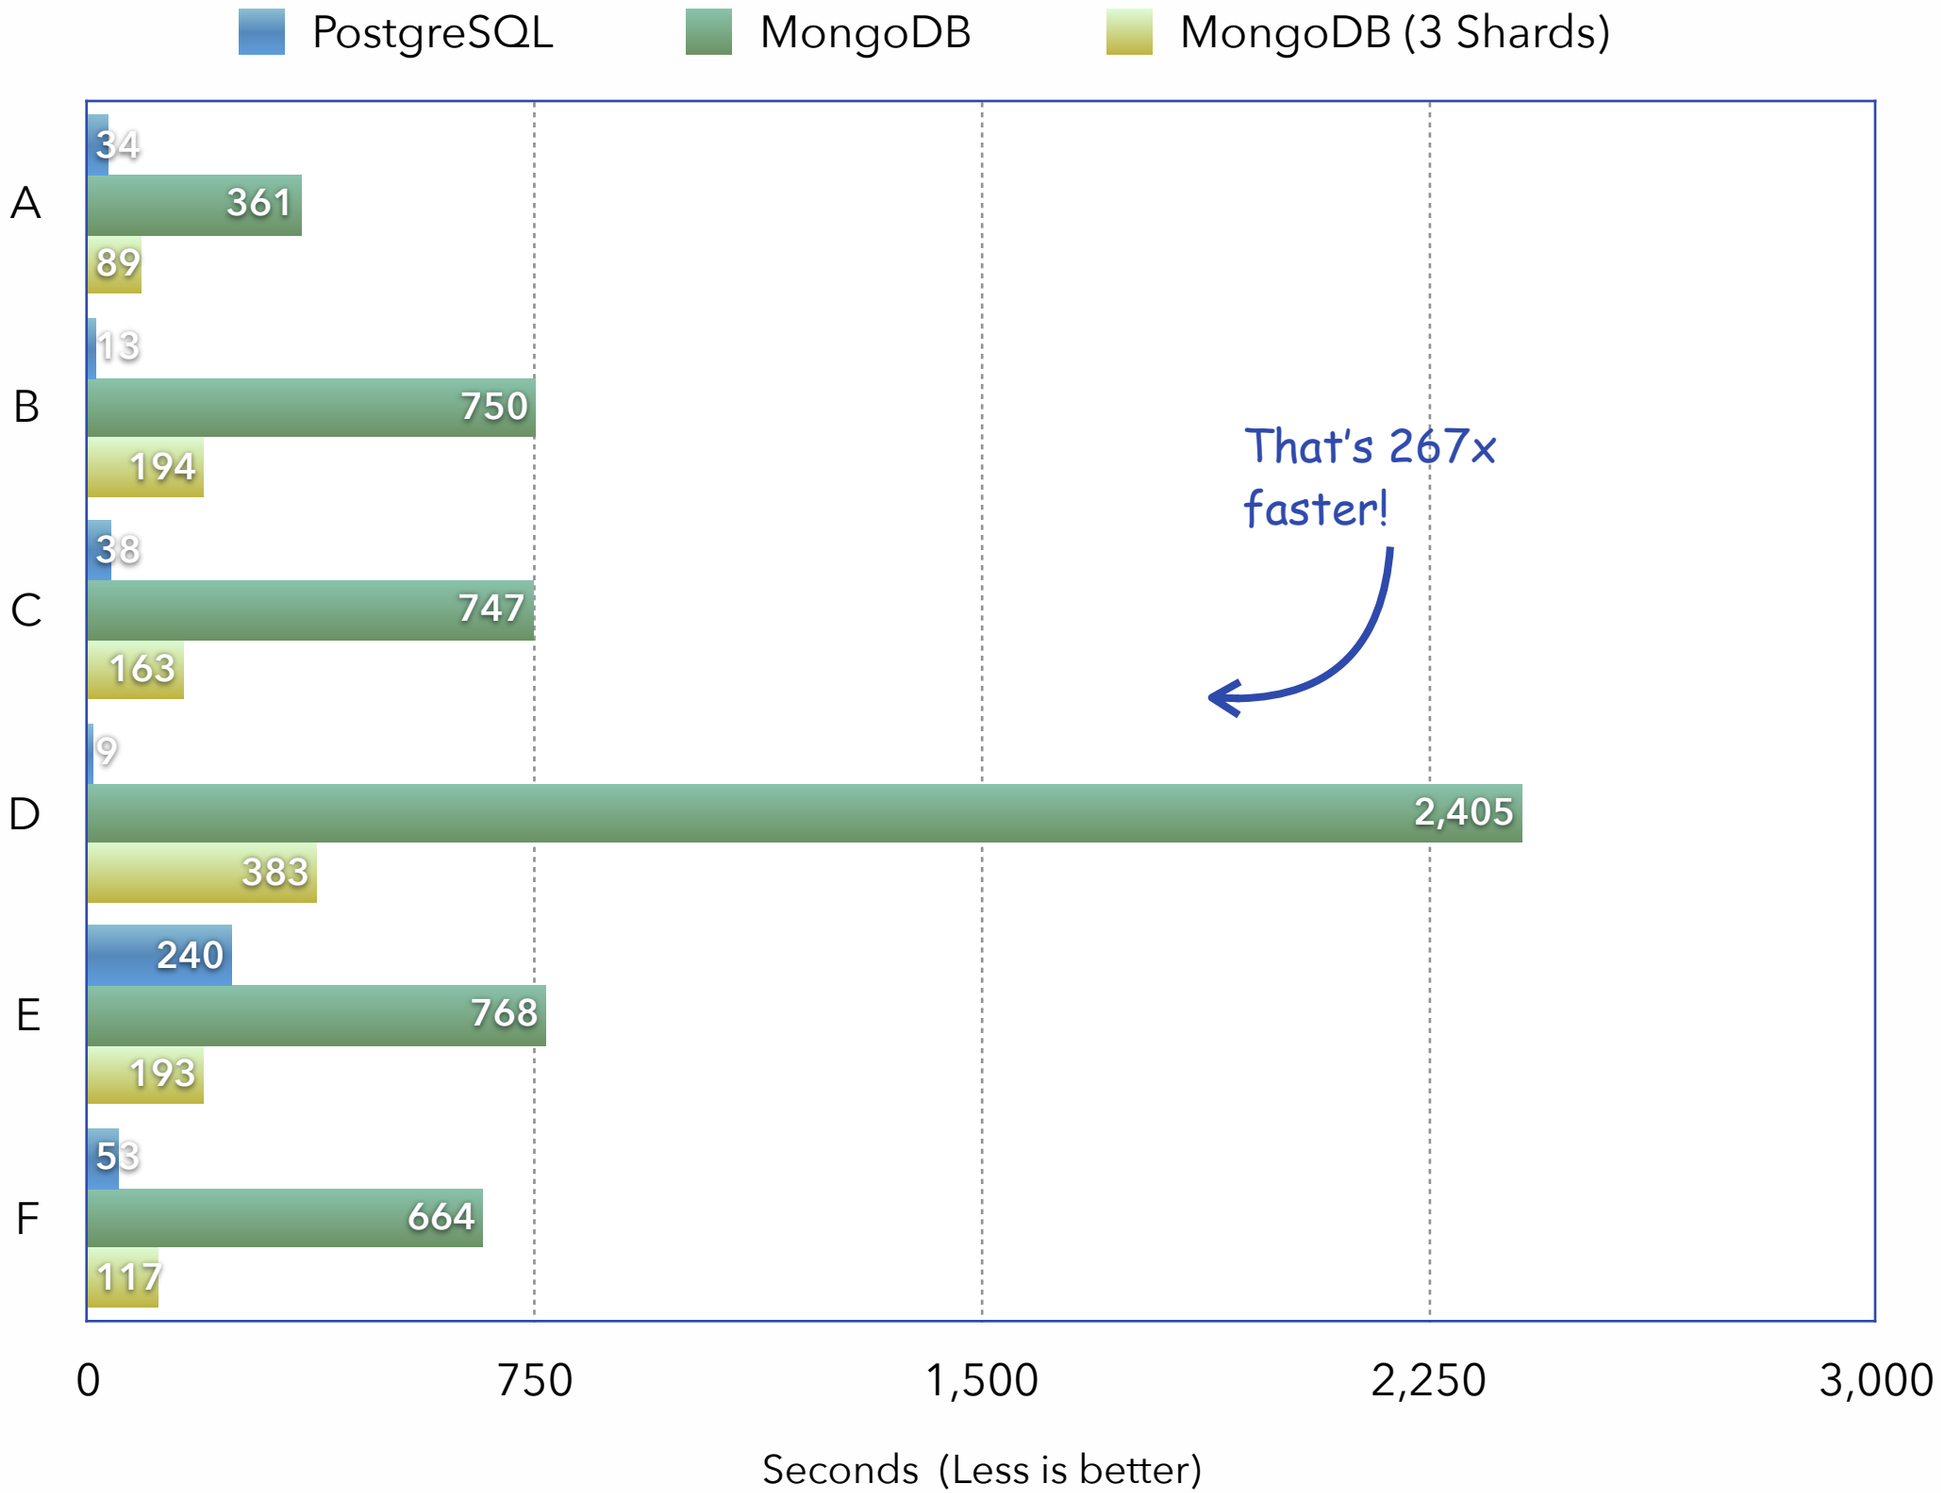

Github 100GB no idx

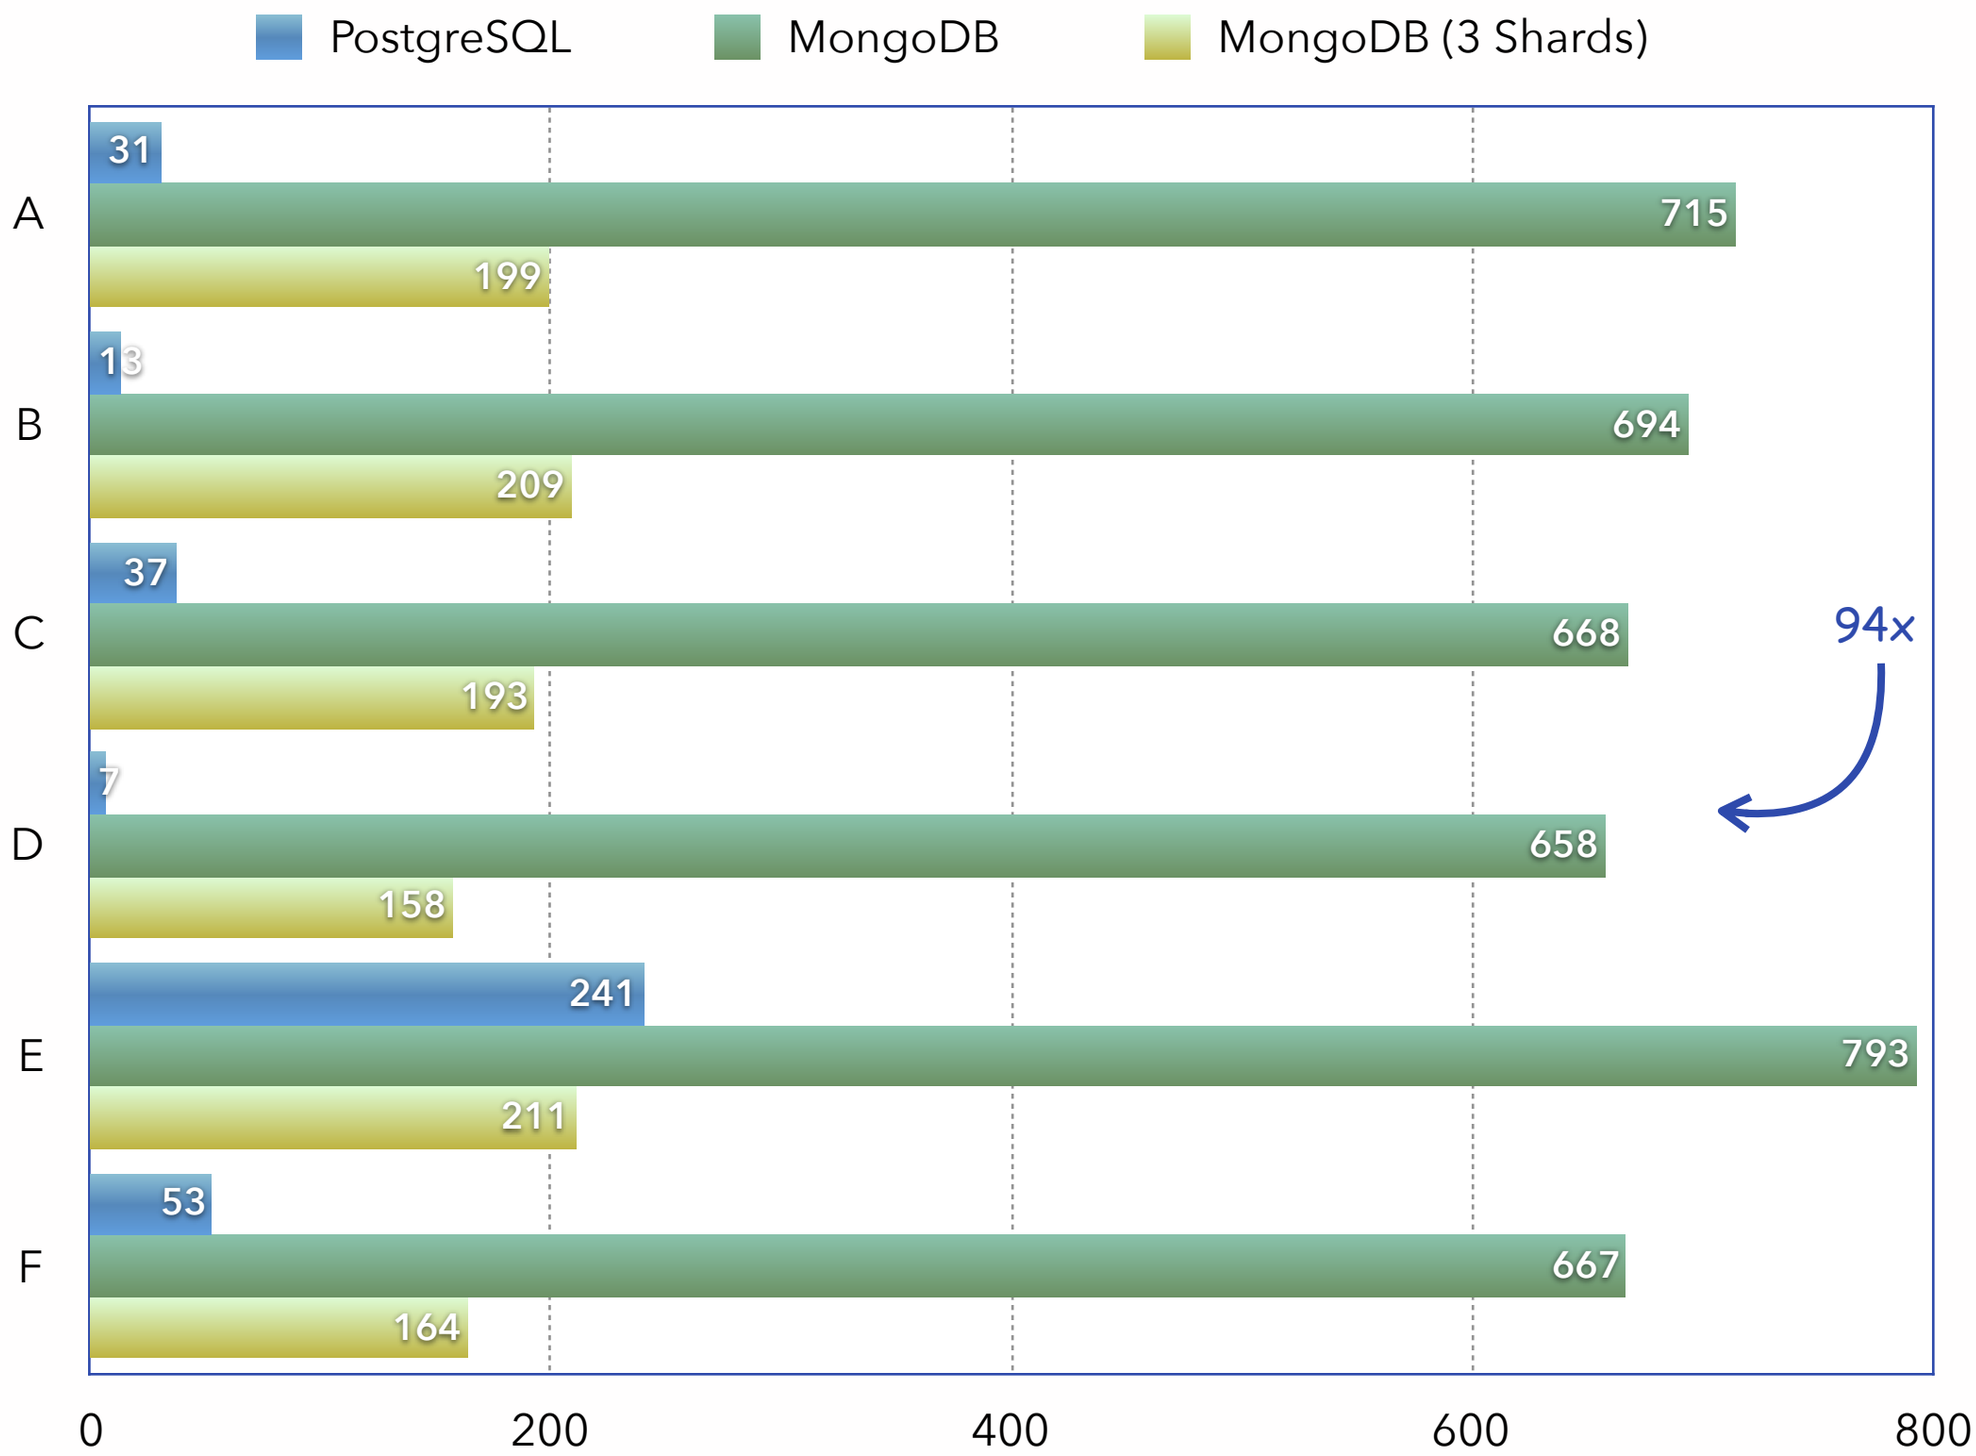

Important because many BI queries don't use indexes

Demo!!!

What's on my MongoDB database?


```
"comments" : "http://delicious.com/url/c666ac39ef072d4eae96a7f062b3b93d", "guidislink" : false, "title_detail
" : { "base" : "http://feeds.delicious.com/v2/rss/recent?min=1&count=100", "type" : "text/plain", "language"
: null, "value" : "Microsoft PressPass - Image Gallery: Microsoft Logos:" }, "link" : "http://www.microsoft.c
om/Presspass/gallery/ms-logos.mspx?FINISH=YES", "source" : {  }, "wfw_commentrss" : "http://feeds.delicious.c
om/v2/rss/url/c666ac39ef072d4eae96a7f062b3b93d", "id" : "http://delicious.com/url/c666ac39ef072d4eae96a7f062b
3b93d#brunokenj", "tags" : [ { "term" : "microsoft", "scheme" : "http://delicious.com/brunokenj/", "label" :
null }, { "term" : "logos", "scheme" : "http://delicious.com/brunokenj/", "label" : null }, { "term" : "image
s", "scheme" : "http://delicious.com/brunokenj/", "label" : null }, { "term" : "logo", "scheme" : "http://del
icious.com/brunokenj/", "label" : null }, { "term" : "design", "scheme" : "http://delicious.com/brunokenj/",
"label" : null } ] }
{ "_id" : ObjectId("58ca54935a9d62483fb2cac3"), "updated" : "Thu, 10 Sep 2009 02:14:33 +0000", "links" : [ {
"href" : "http://www.interactiondesign.com.au/", "type" : "text/html", "rel" : "alternate" } ], "title" : "AC
ID Home", "author" : "amdconnor", "comments" : "http://delicious.com/url/0253b164c08b0c80efaea1b85c66a088", "
guidislink" : false, "title_detail" : { "base" : "http://feeds.delicious.com/v2/rss/recent?min=1&count=100",
"type" : "text/plain", "language" : null, "value" : "ACID Home" }, "link" : "http://www.interactiondesign.com
.au/", "source" : {  }, "wfw_commentrss" : "http://feeds.delicious.com/v2/rss/url/0253b164c08b0c80efaea1b85c6
6a088", "id" : "http://delicious.com/url/0253b164c08b0c80efaea1b85c66a088#amdconnor", "tags" : [ { "term" : "
interactive", "scheme" : "http://delicious.com/amdconnor/", "label" : null } ] }
{ "_id" : ObjectId("58ca54935a9d62483fb2cac4"), "updated" : "Sun, 13 Sep 2009 13:09:23 +0000", "links" : [ {
"href" : "http://www.ibm.com/developerworks/data/library/techarticle/dm-0801sauter/", "type" : "text/html", "
rel" : "alternate" } ], "title" : "The information perspective of SOA design, Part 1: Introduction to the inf
ormation perspective of a Service Oriented Architecture", "author" : "bspaulding", "comments" : "http://delic
ious.com/url/4d194dfd510d774df4507790e0b62f02", "guidislink" : false, "title_detail" : { "base" : "http://fee
ds.delicious.com/v2/rss/recent?min=1&count=100", "type" : "text/plain", "language" : null, "value" : "The inf
ormation perspective of SOA design, Part 1: Introduction to the information perspective of a Service Oriented
 Architecture" }, "link" : "http://www.ibm.com/developerworks/data/library/techarticle/dm-0801sauter/", "sour
ce" : {  }, "wfw_commentrss" : "http://feeds.delicious.com/v2/rss/url/4d194dfd510d774df4507790e0b62f02", "id"
 : "http://delicious.com/url/4d194dfd510d774df4507790e0b62f02#bspaulding", "tags" : [ { "term" : "soa", "sche
me" : "http://delicious.com/bspaulding/", "label" : null } ] }
Type "it" for more
rs1:PRIMARY>
```

Stampede infers and clearly shows you your data

did	rid	seq	rel_s	href_s	type_s
0	0	0	alternate	http://www.linuxforums.org/forum/linux-tutorials-howtos-reference-material/64958-how-install-software-linux.… …html	text/html
1	1	0	alternate	http://progit.org/book/	text/html
2	2	0	alternate	http://www.propagandalarevista.com/	text/html
3	3	0	alternate	http://lifehacker.com/5356882/mobalivecd-updates-and-adds-usb-drive-support	text/html
4	4	0	alternate	http://www.blackholly.com/writingresources.htm	text/html
5	5	0	alternate	http://www.rachaelgibson.co.uk/blog/index.php	text/html
6	6	0	alternate	http://www.buzzfeed.com/scott/nerd-venn-diagram	text/html
7	7	0	alternate	http://online.wsj.com/article/SB125245417031494185.html	text/html
8	8	0	alternate	http://code.google.com/p/cacherl/	text/html
9	9	0	alternate	http://www.mflenses.com/	text/html
10	10	0	alternate	http://www.pspiso.com/showthread.php?t=92302	text/html

did	rid	seq	language_n	type_s	value_s	base_s
0	0	¤	t	text/plain	How To Install Software in Linux - Linux Forums	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
1	1	¤	t	text/plain	Pro Git - Table of Contents	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
2	2	¤	t	text/plain	Propaganda La Revista Quinta Edicion	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
3	3	¤	t	text/plain	MobaLiveCD Updates and Adds USB Drive Support -… … Live CD - Lifehacker	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
4	4	¤	t	text/plain	Holly Black''s Writing Resources	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
5	5	¤	t	text/plain	furcoat	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
6	6	¤	t	text/plain	Nerd Venn Diagram [PIC]	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
7	7	¤	t	text/plain	Finance Overhaul Falters as ''08 Shock Fades - … …WSJ.com	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
8	8	¤	t	text/plain	cacherl - Project Hosting on Google Code	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
9	9	¤	t	text/plain	Manual Focus Forum | Lens, Very, Macro, Quality… …, Photography | Manual Focus Lenses - Manual Fo… …cus Lenses	http://feeds.delicious.com/v2/rss/recent?min=1&count=100
10	10	¤	t	text/plain	*OnEfLyRiDe* over 746 cso''s! last updated 5/14… …/09 - PSPISO	http://feeds.delicious.com/v2/rss/recent?min=1&count=100

Or use SQL tools to show your schema

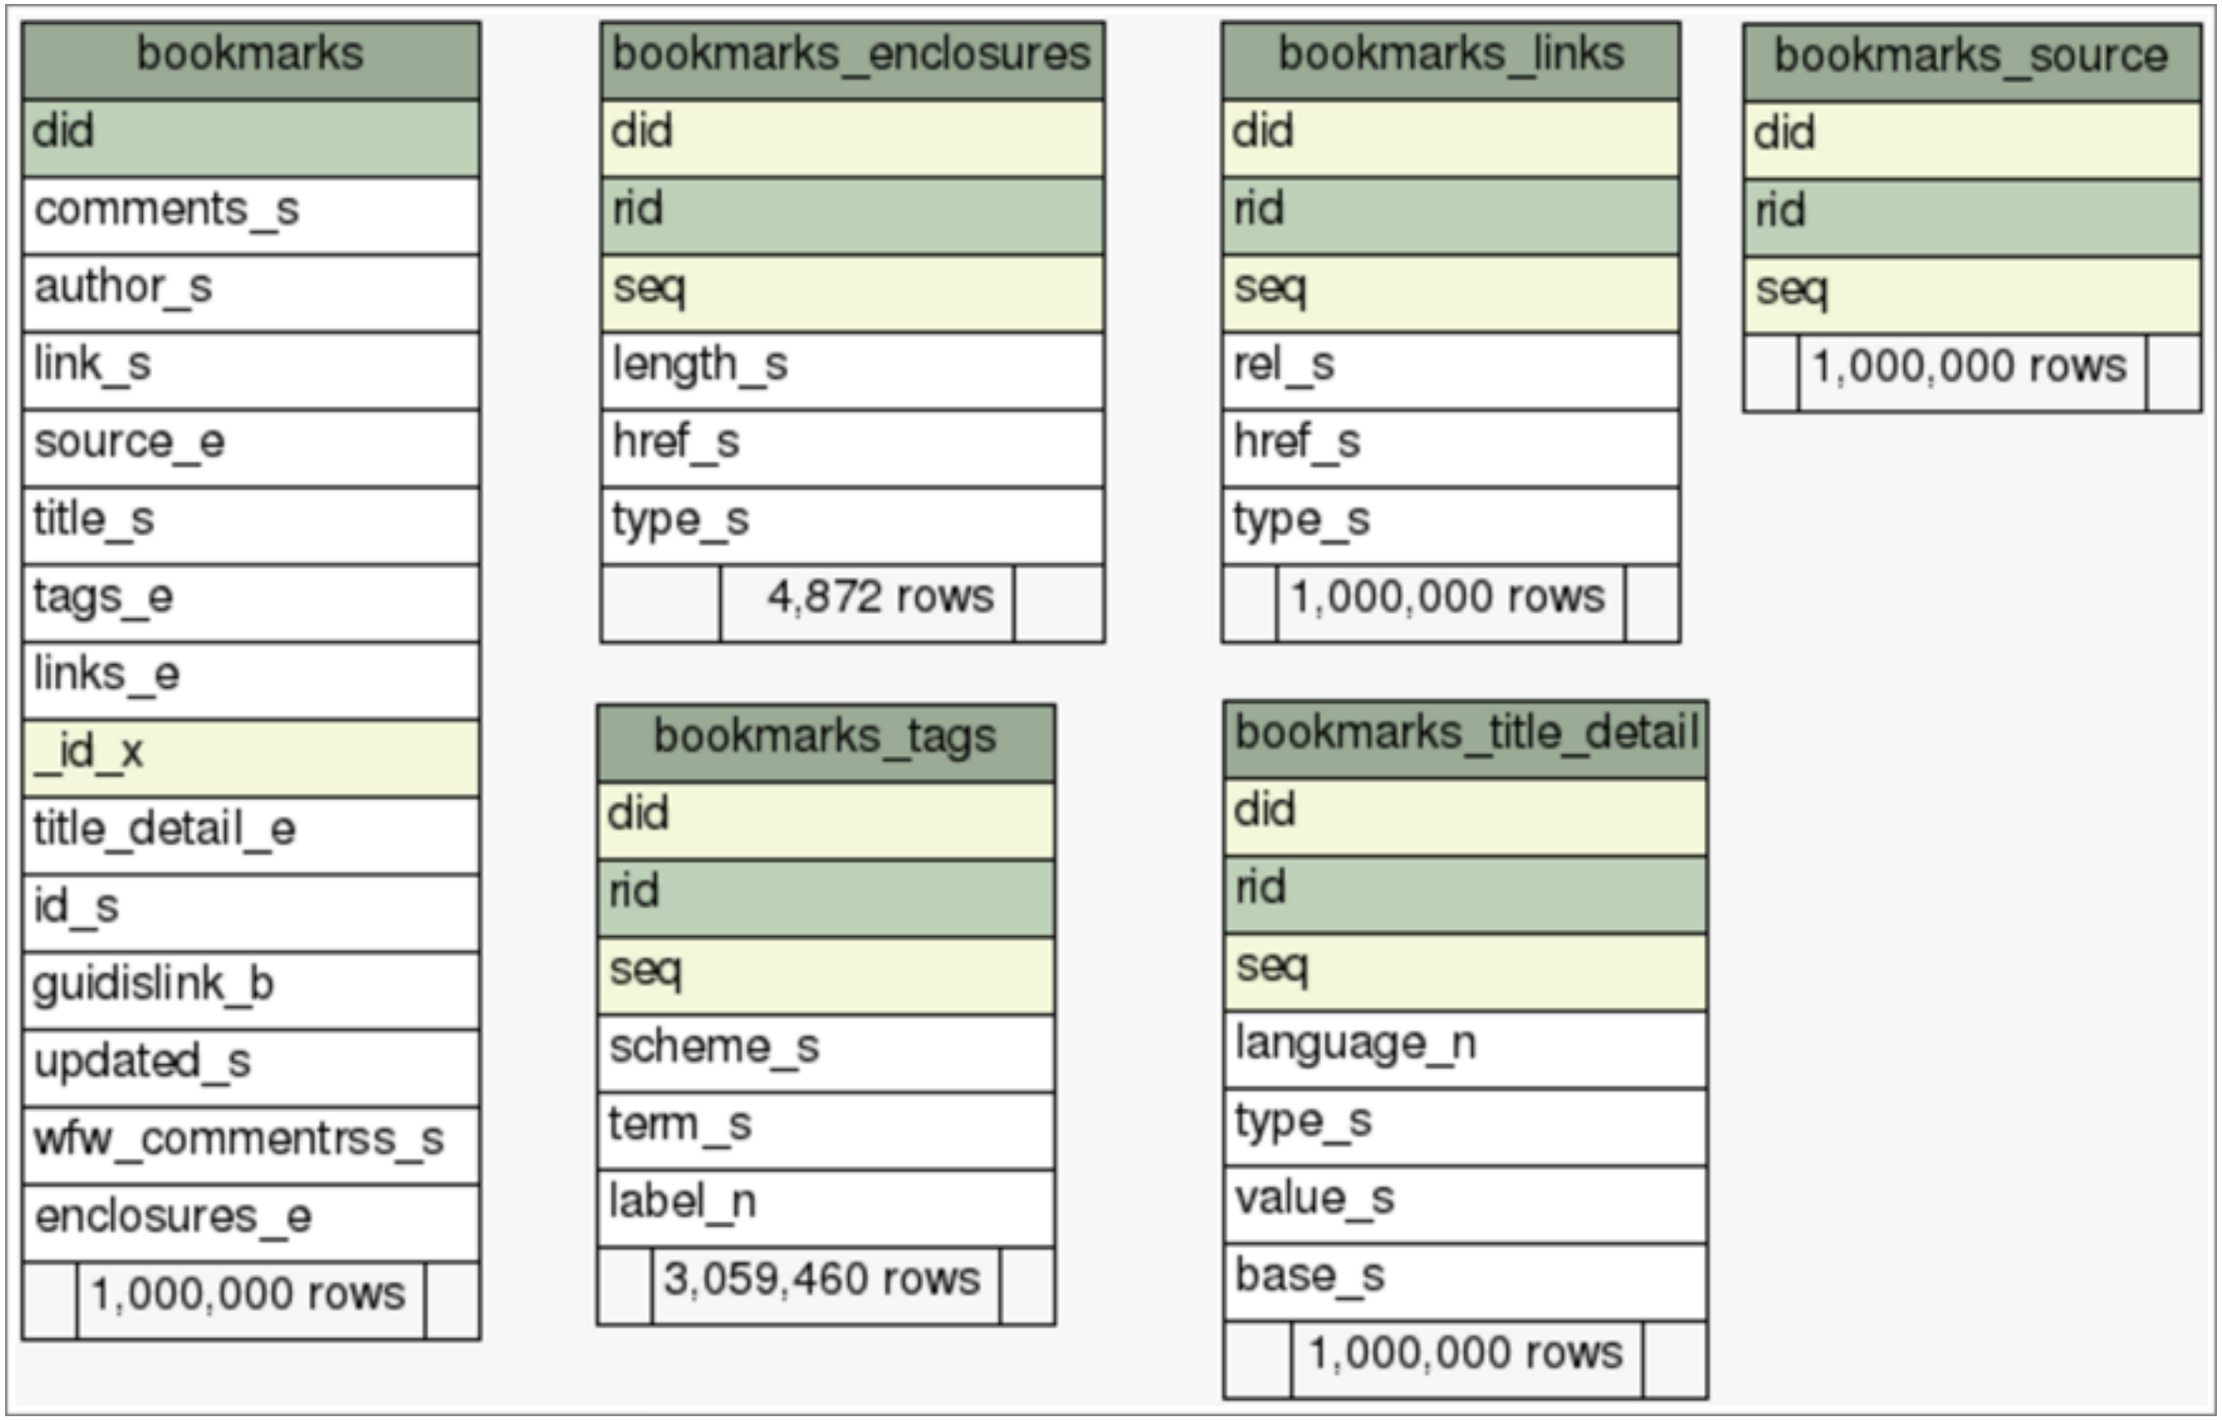

Let's Talk!

Migrating off of MongoDB to PostgreSQL

www.8kdata.com

info@8kdata.com