

The denormalized query engine design pattern

Simon Willison, DjangoCon US - 16th August 2017

Slides will be here: <http://lanyrd.com/sftkxk>

L A W R E N C E
JOURNAL-WORLD®

YAHOO!

the guardian

Lanyrd.com

Eventbrite

Design patterns

denormalized query engine

- Relational database as single point of truth
- Denormalize all relevant data to a separate search index
- Invest **a lot** of effort synchronizing the two

Relational weaknesses

- They're not great at counting
- You should avoid queries that read more than a few thousand rows
- MySQL can only use one index per query

Search engine strengths

- Horizontal scaling
- Aggregations and counts
- Queries across multiple indexed fields
- Relevance calculations and scoring
- ... and text search too

2005

flickr



Sharding by user

Users
1-
10,000

MySQL

Users
10,000-
20,000

MySQL

Users
20,000-
30,000

MySQL

Users
30,000-
40,000

MySQL

flickr

Explore

Create

Photos, people, or groups



Sign In

Sign Up

Explore

Trending




NEW

More ▾

Tags ▸ racoons


Recently tagged — racoons

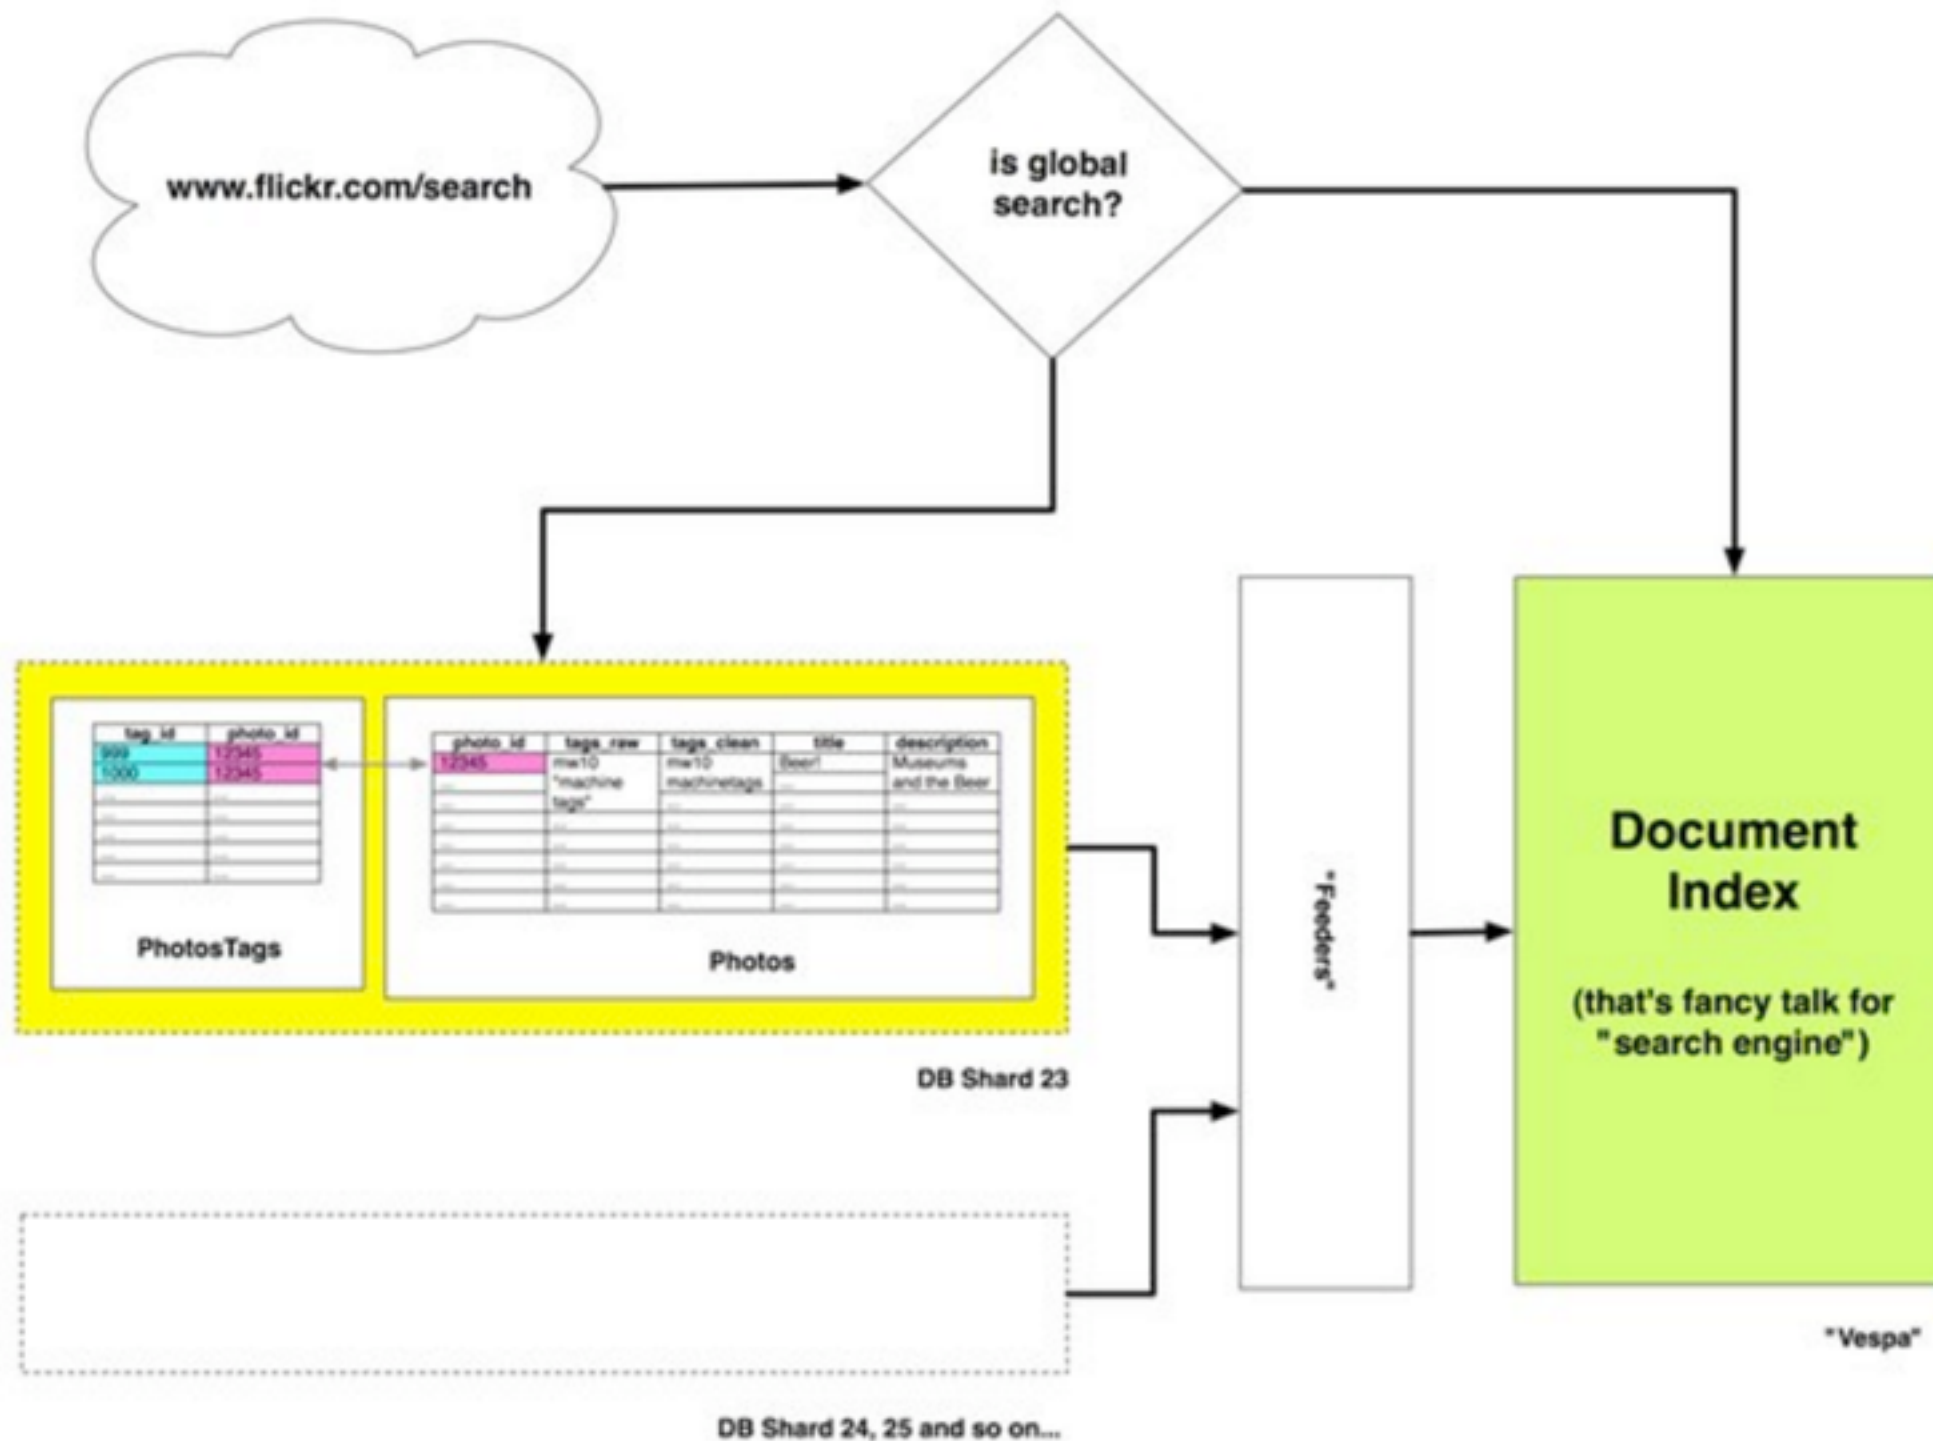
[View all 96](#)



Related groups — racoons

[View all 6](#)





Smart query routing

- Send a user's queries about their own data to the relational database
- Send queries about other users and public data to the search index

Solving a scalability crisis at Lanyrd

[yours](#) 1[all](#) 136[contacts](#) 103[guides](#) 10[topics](#) 38

We've found **136** upcoming events you may be interested in.



Why not try [tracking topics and places](#) that interest you or [some of our guides](#) to find events that are relevant to you?

**Simon Willison**[Your profile page](#)

CONTROL YOUR EMAIL DESTINY

[Edit email preferences](#)

KEEP EXPLORING

[Discover events through curated guides](#)

TODAY

13

DjangoCon US 2017

[United States, Spokane](#)

13th–18th August 2017

[Community](#) [Design](#) [Django](#) [Python](#) [Web Developers](#)

4 contacts attending

From topic [Django](#)[Let others know why you are going](#)**You're speaking**
AT THIS EVENT

14

UX & Digital Design Week 2017

[England, London](#)

14th–20th August 2017

[Agile Product Management](#) [Digital Design](#)
[Digital Product Design](#) [Digital Strategy](#) [LeanUX](#) [UI&UX](#)
[User Experience](#)From topic [User Experience in United Kingdom](#)[Attend](#)[Track](#)

THIS WEEK

17

Build Right: Maker Series with Chris Coyier

[Online conference](#) [United States, Dayton](#)

17th August 2017

1 contact speaking

[Attend](#)[Track](#)

From our blog

[Personalized schedules for any event](#)

Running a multi-track event? Lanyrd makes [event planning](#) easy by allowing event organizers to turn on personalized schedules for all of their attendees, synchronized to their mobile phones and tablets through our iOS, Android and Mobile Web apps.

[Lanyrd acquired by Eventbrite](#)

Lanyrd is becoming part of the Eventbrite family! We're confident that Lanyrd will be better than ever as a result, whether you plan to [sell tickets](#) as an organizer or discover new events as an attendee.

[Help alpha test Lanyrd for Android](#)

We've released a public alpha version of our new native Lanyrd Android app. If you have an Android phone or tablet we'd love to have your help ironing out any bugs.

[Find speakers for your events with Lanyrd's new speaker directory](#)

Today we're pleased to announce our

- Search Solr for events where...
 - attendee_ids matches [giant list of Twitter IDs]
 - start_date is in the future
- Scale Solr horizontally with replication




Elasticsearch

Search & Analyze Data in Real Time

- Open source search engine built on Apache Lucene
- Interface is all JSON over HTTP - easy to use from any language
- Claims to be “real-time” - it’s close enough
- Insanely powerful query language (a JSON DSL)
- Strong focus on analytics in addition to text search
- Elastic **means** elastic: scales horizontally

Faceted search and aggregations

DC Inbox Explorer ([background](#) / [source code](#))



total: 56879

[Message from Rep. Michael Burgess \[R-TX26\] on 2015-06-19](#)

Dr. Burgess weekly address: Tropical Depression Bill

Click here to open this e-mail in its own browser window Click here to open a plain text version of this email Friday, June 19, 2015 WEBSITE | SHARE ON: Dr. Burgess weekly address: Tropical Depression Bill This week, Dr. Burgess thanks the Army ... ([more](#))

[Message from Rep. Frank Guinta \[R-NH1\] on 2015-06-19](#)

SURVEY: What are your thoughts...

News from Congressman Frank Guinta Home About Frank About Our District News Contact Dear Friends, Since the implementation of Obamacare, I have heard from countless Granite Staters regarding their concerns with broken healthcare promises, such ... ([more](#))

[Message from Rep. Emanuel Cleaver \[D-MO5\] on 2015-06-19](#)

EC From DC: Our Veterans Visit Washington

Our Veterans Visit Washington Congressman Cleaver greets Ward Golston, a World War II Veteran You may have heard of the incredible program, Honor Flight. The inaugural Honor Flight took place in May of 2005. Six small planes flew out of Springfield, ... ([more](#))

[Message from Rep. Michael Conaway \[R-TX11\] on 2015-06-19](#)

Conaway Chronicle: Increasing Congressional Oversight of Trade Negotiations, Video Tour of the U.S. Capitol Building, The San Saba Wedding Oak

June 19, 2015 Conaway Chronicle: Increasing Congressional Oversight of Trade Negotiations, Video Tour of the U.S. Capitol Building, The San Saba Wedding Oak _____ Thank you for reading The Conaway Chronicle, a ... ([more](#))

party

Republican 36948

Democrat 19648

Independent 283

role_type

representative 47297

senator 9582

state

CA 5769

TX 4512

FL 3634

VA 2566

PA 2561

```
GET /emails/_search
{
  "query": {
    "match": {
      "Body": "security"
    }
  },
  "aggs": {
    "role_type": {
      "terms": {
        "field": "role_type"
      }
    },
    "party": {
      "terms": {
        "field": "party"
      }
    }
  }
}
```

Relational database sync strategies

- updated/last_touched timestamp
- Indexing queue (redis/kafka)
- Subscribe to database replication stream

last_touched

```
class Conference(models.Model):
    name = models.CharField(max_length=128)
    url = models.URLField(blank=True)
    # ...
    last_touched = models.DateTimeField(
        db_index=True,
        default=datetime.datetime.utcnow,
    )

# Reindex all conferences when associated guide is edited:
guide.conferences.all().update(last_touched=datetime.datetime.utcnow())
```

Indexing code needs to track most recently seen last_touched date time

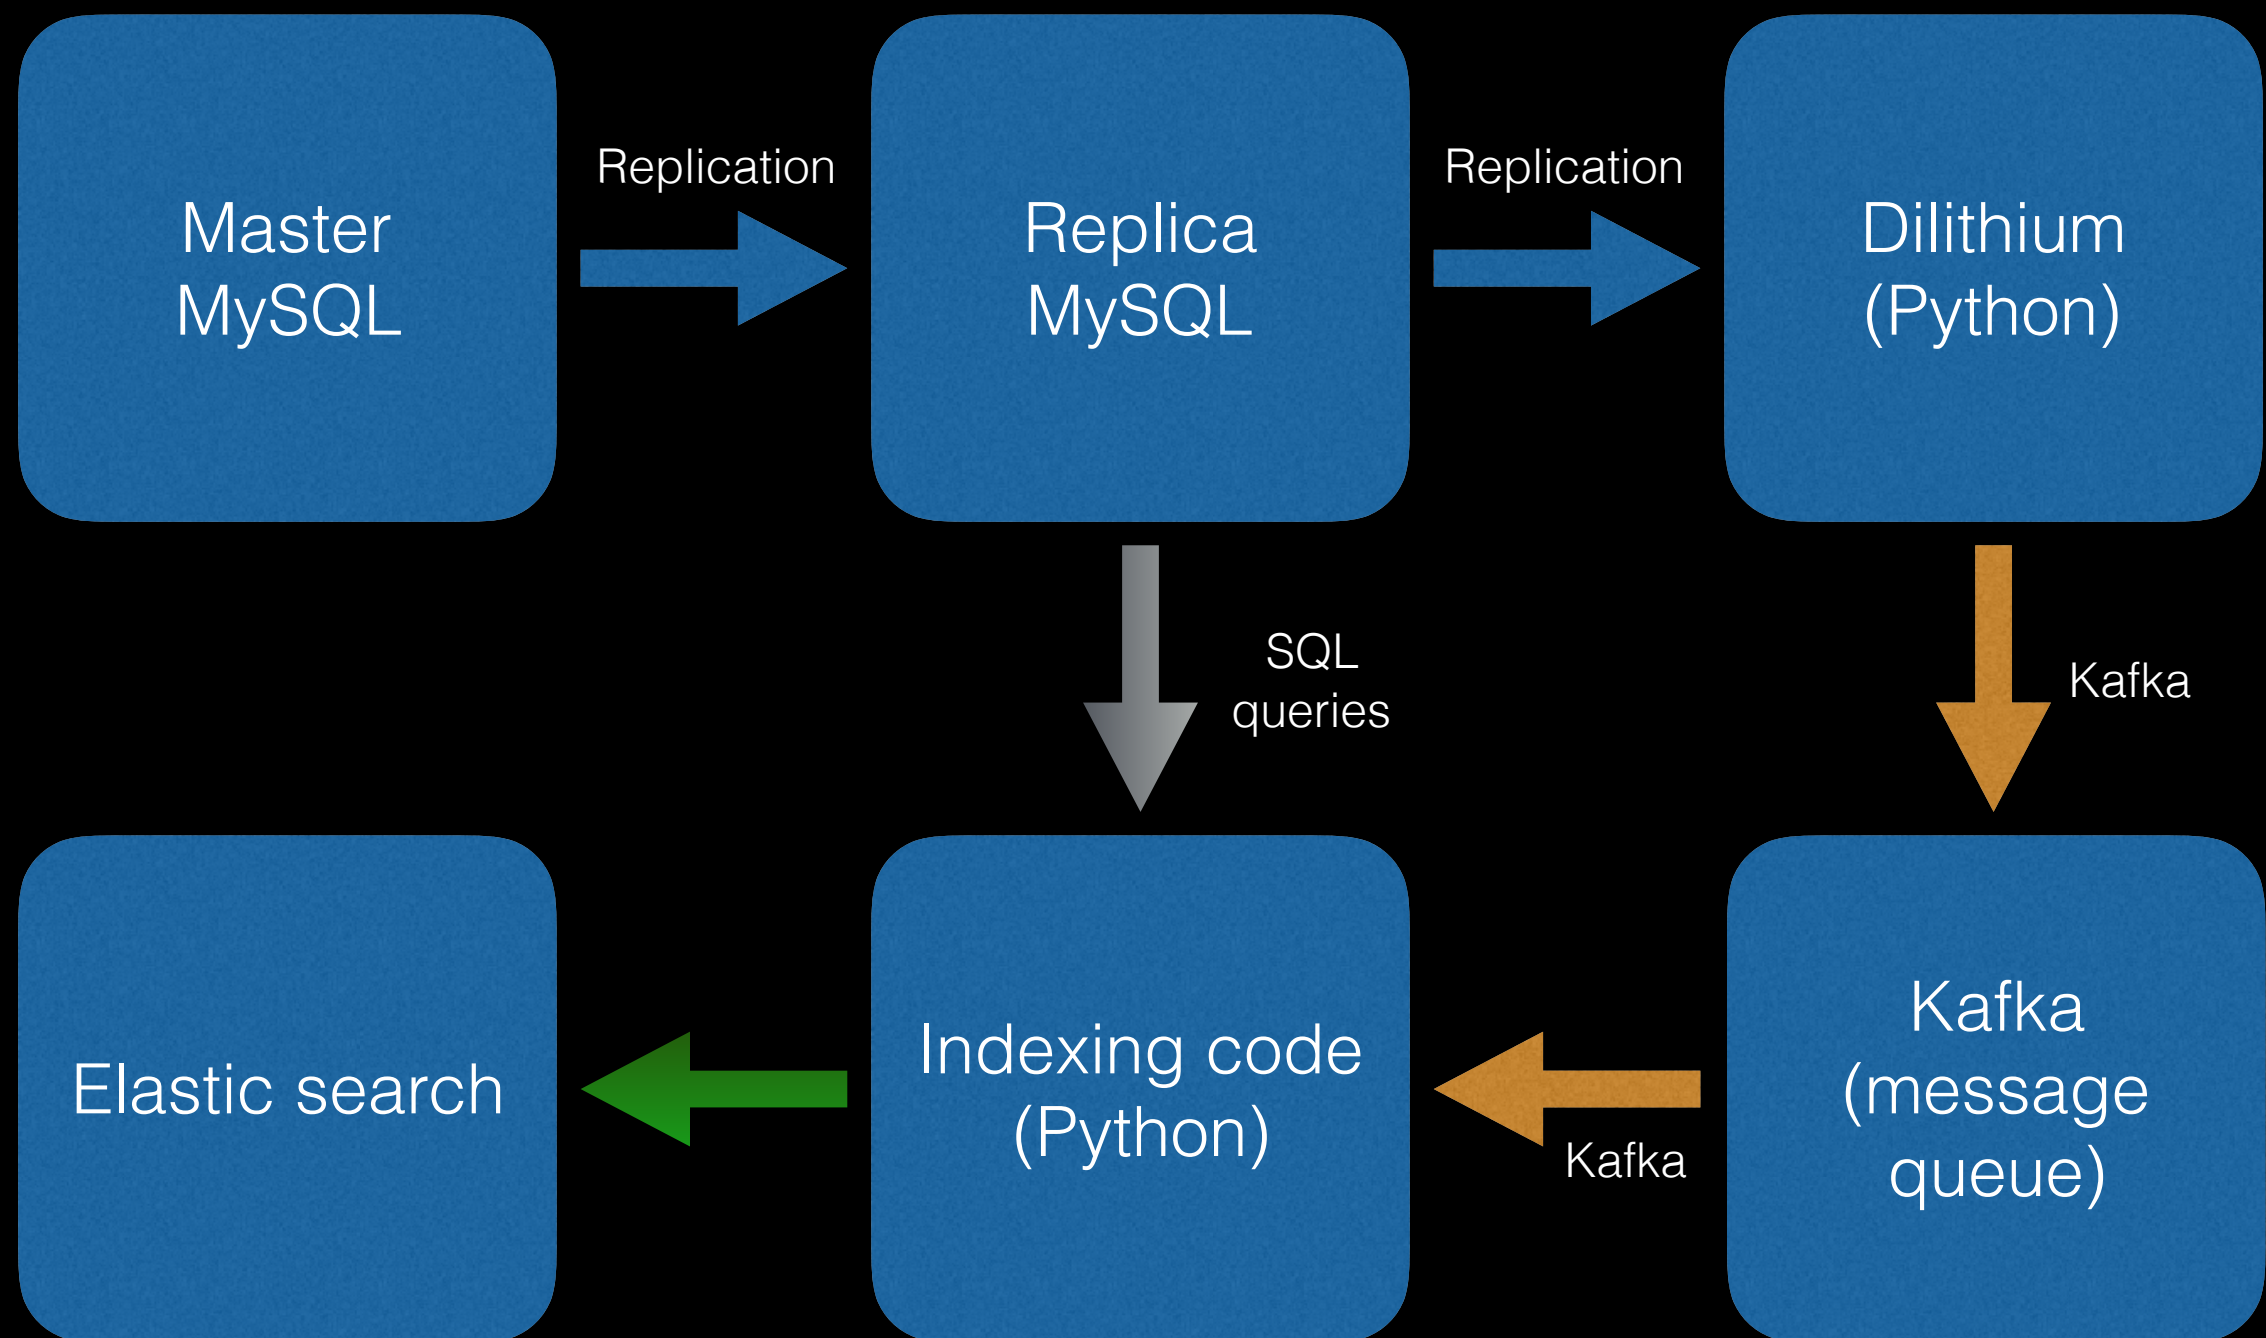
Redis/Kafka queue

- Any time an object needs reindexing, add the type/ID to a queue
- Every few seconds, clear the queue, de-dupe the item IDs, fetch from database and reindex

Replication log

- MySQL supports replication... and you can listen to the replication stream itself
- github.com/noplay/python-mysql-replication

Dilithium



Tips and tricks

Inflate IDs into objects

- Avoid ever showing the user stale data!
- Return row IDs from search...
- ... fetch them from the database for display
- Relational databases are REALLY FAST at primary key lookups (and prefetch_related)

Self-repair on deletes

- If the search index returns an ID that no longer exists in your database...
- Quietly drop it from the output (9 results instead of 10)
- Queue for deletion from the index via celery

Accurate filter trick

```
GET /events/_search
{
  "query": {
    "term": {
      "saved_by_users": "123124"
    }
  }
}
```

Accurate filter trick

```
recent_ids = user.event_saves.filter(  
    # Created within last 5 minutes  
    created__gte=(  
        datetime.utcnow() -  
        timedelta(minutes=5)  
    )  
) .values_list('event_id', flat=True)  
  
# [776, 4124, 3414]
```

Accurate filter trick

```
"query": {  
  "bool": {  
    "should": [  
      {  
        "term": {  
          "saved_by_users": "123124"  
        }  
      },  
      {  
        "ids": {  
          "values": [776, 4124, 3414]  
        }  
      }  
    ]  
  }  
}
```

More use-cases

Recommendations


- Search for events where saved-by matches my-friend-1 or my-friend-2 or my-friend-3 or ...
- Find events similar to my-last-10-saved-events
- Search engines are great at scoring! Boost by in-same-city-as-me, boost more by saved-by-my-friends

Geographic search

- Elasticsearch has great support for geo...
 - Find documents within X radius of point Y
 - Find documents contained in polygon Z
- Combine these with search and filters

Visualizations

DC Inbox Explorer ([background](#) / [source code](#))



total: 56879

[Message from Rep. Michael Burgess \[R-TX26\] on 2015-06-19](#)

Dr. Burgess weekly address: Tropical Depression Bill

Click here to open this e-mail in its own browser window Click here to open a plain text version of this email Friday, June 19, 2015 WEBSITE | SHARE ON: Dr. Burgess weekly address: Tropical Depression Bill This week, Dr. Burgess thanks the Army ... ([more](#))

[Message from Rep. Frank Guinta \[R-NH1\] on 2015-06-19](#)

SURVEY: What are your thoughts...

News from Congressman Frank Guinta Home About Frank About Our District News Contact Dear Friends, Since the implementation of Obamacare, I have heard from countless Granite Staters regarding their concerns with broken healthcare promises, such ... ([more](#))

party

Republican 36948

Democrat 19648

Independent 283

role_type

representative 47297

senator 9582

Real-time map/reduce

- Elasticsearch can be thought of as a real-time map/reduce system
- Like Hadoop, but you can expose it to your users!

in summary...

denormalize to a
query engine!

Elasticsearch is pretty good