Elastic Search as Primary Storage?

NoSQL Roadshow 2013
Jodok Batlogg


@jodok
@lovelysystems

Elasticsearch

cool. bonsai cool

Real-Time

- Hadoop -> Map/Reduce?
- River Implementations

Distributed

- really!
- scale horizontally
- shared nothing

High availability

- replicas
- master node election
- HTTP interface

Multi-Tenancy

- Multiple Indizes
- Aliases
- add/remove indizes on the fly
- query groups of indizes

Full-Text capabilities

- Apache Lucene under the hood
- geo-location
- autocomplete, did-you-mean

Document oriented

- JSON
- all fields indexed by default

Schema free

- JSON documents
- Tweaking with mappings possible

Conflict Management

- Optimistic version control

RESTful API

- (but also native JAVA binding)

So what?

Quora

Search

ElasticSearch MongoDB NoSQL Edit

Why should not I use ElasticSearch as my primary datastore? Edit

I'm just learning about ElasticSearch and I initially thought it was just an indexing/search engine but can now see that it's a full schema-less datastore (as well).

Sorry if this is a naive question but what would be the disadvantage of using ElasticSearch as my primary datastore instead of say MongoDB? Surely the former removes the extra step of having to update indexes from my database all the time? Edit

Comment · Share (1) · Options

Quora (January 2012)

- Security?
- Transactions?
- Durability?
- Maturity of Tools?
- Large Computations?
- Data Availability?

What does a data system do?

Big Data Borat
@BigDataBorat

Following

New data theory promote use of "optimistic persistence" -- define as "hopefully data you lose anyway no was important"

Reply Retweet Favorite More

67 RETWEETS

10 FAVORITES

6:25 PM - 20 Nov 12

ACID

- Atomicity - a transaction is all or nothing
- Consistency - only valid data is written to the database
- Isolation - pretend all transactions are happening serially and the data is correct
- Durability - what you write is what you get

ACID Problems

- Gives you too much, across multiple nodes.
- Down time? Reliable? Multiple Nodes!
- Lots and lots of reads and writes.
- The algorithms don't work in a distributed environment at any acceptable speed.

CAP

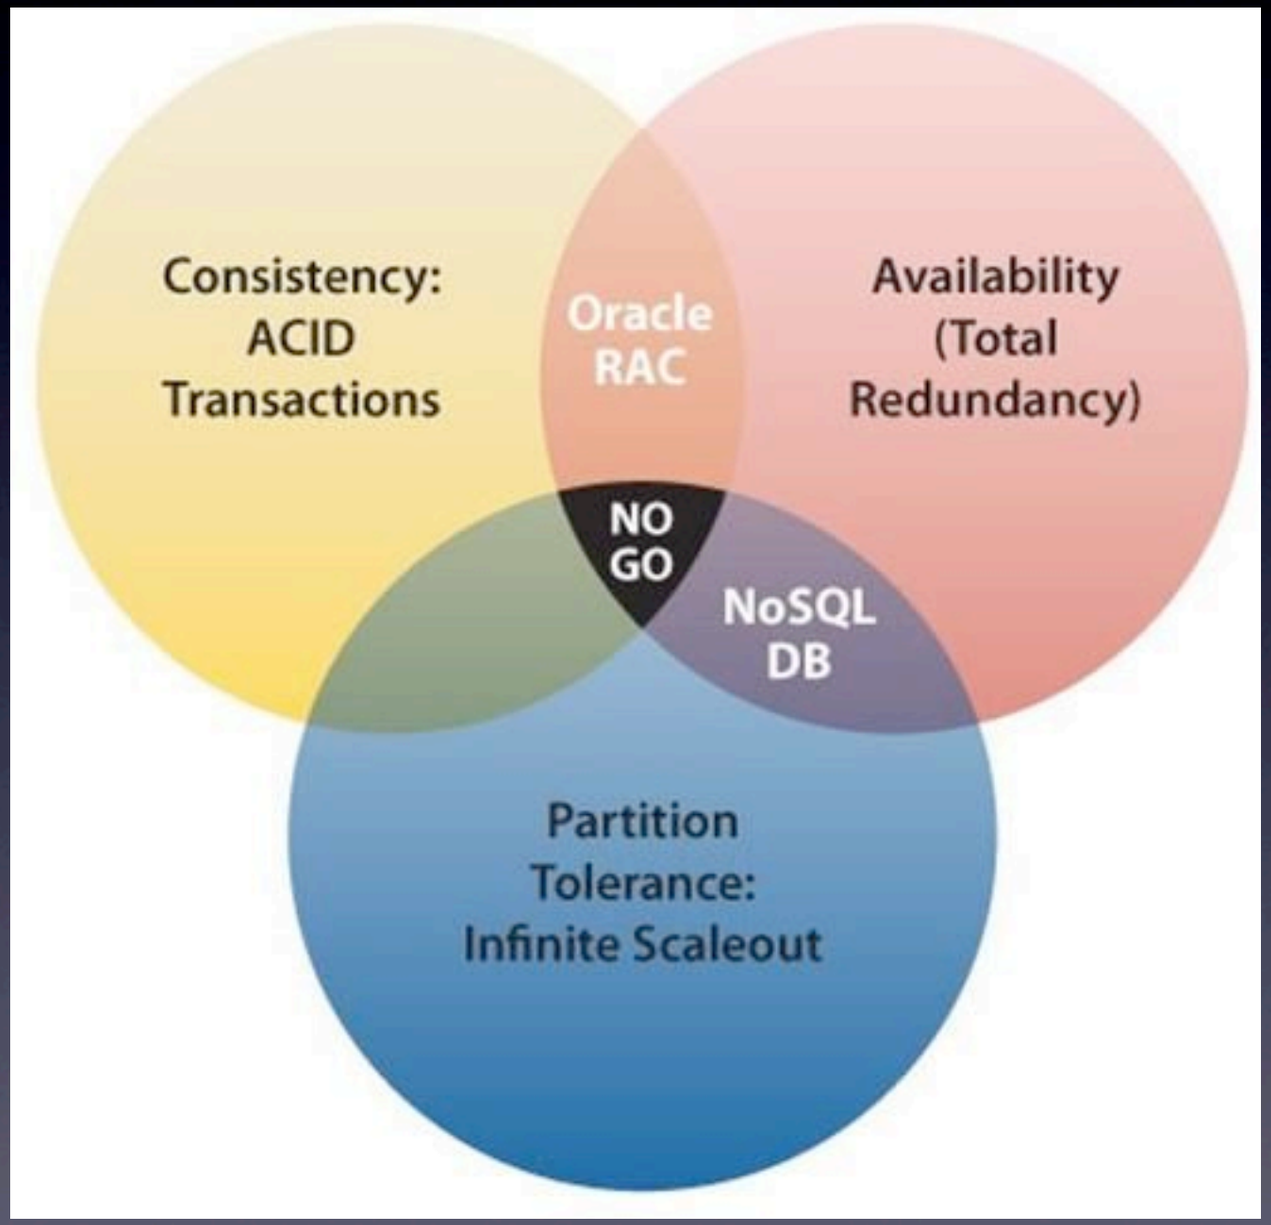

http://bigdatanerd.files.wordpress.com/2011/12/cap-theorem.jpg

BASE

- Basically Available - system seems to work all the time
- Soft State - it doesn't have to be consistent all the time
- Eventually Consistent - becomes consistent at some later time

Everyone who builds big applications builds them on CAP and BASE

Retrieve data that you previously stored

- PUT
- GET

Not really

- Store location information on people
  - How many people live in a particular location
  - Where does Sally live?
  - What are the most populous locations?
- Store pageview information
  - How many pageviews on September 2nd?
  - How many unique visitors over time?
- Store transaction history for bank account
  - How much money does George have?
  - How much money spend people on housing?

Query = Function(All data)

- Sometimes you retrieve what you stored
- Oftentimes you do transformations, aggregations, etc.
- e.g. total number of pageviews to a URL over a range of time

Structured Data?

Schemas

- Hard to change
- Get in the way
- Add development overhead
- Requires annoying information

Nathan Marz:
http://www.slideshare.net/nathanmarz/runaway-complexity-in-big-data-and-a-plan-to-stop-it

Use a schemaless Database!

This is an overreaction.

confuses the poor
implementation of schemas
with the value that schemas
provide

function(data unit)

That says whether data is valid or not.

This is useful.

Value of Schemas

- Structural integrities
- Guarantees on what can be stored and read
- Prevents corruption

Without Schemas

- You'll detect corruption issues at read-time
- Potentially long after the corruption happened
- With little insight into the circumstances of the corruption

Schemas

- Much better to get an exception where the mistake is made, before it corrupts the database
- Saves enormous amounts of time

Potentially long after the corruption happened

Elasticsearch Mappings

A Tweet

```
{
    "tweet" {
        "user" : "kimchy"
        "message" : "This is a tweet!",
        "postDate" : "2009-11-15T14:12:12",
        "priority" : 4,
        "rank" : 12.3
    }
}
```

A Tweet

```
{
  "tweet" : {
    "properties" : {
      "user" : {"type" : "string", "index" : "not_analyzed"},
      "message" : {"type" : "string", "null_value" : "na"},
      "postDate" : {"type" : "date"},
      "priority" : {"type" : "integer"},
      "rank" : {"type" : "float"}
    }
  }
}
```

http://www.elasticsearch.org/guide/reference/mapping/core-types/

Scather/Gather

Elasticsearch Facets

- Terms facet
- Statistical facet
- Histogram facet
- Geo distance

Memory

- relevant field values are loaded into memory
- long vs. short integers

Lovely Systems Facets

- https://github.com/lovelysystems/elasticsearch-ls-plugins
- Uncached Facet
  (Date Histogram for Int/Long Values)
- Distinct Date Histogram Facet
- “Latest” Facet

You know, for search

querying 24 000 000 000 Records in 900ms

@jodok

9962 views

The World's Experts in Search Analytics | Searchmetrics

www.searchmetrics.com/en/

Reader

Pin It bookmark Read Later

searchmetrics

Deutsch | English | Français

Customer Login

Home | SEO Software | Partner | Events | About Us | Blog

Advancement through knowledge.

More SEO! More SEM! More Social!
More competitor analyses!
More knowledge! More success!

Try for free »

Media Coverage | May 22, 2012, 11:41 a.m.
Entrepreneur.com published an interesting article about the Google Penguin update refering to our data. Read more...

Product Notes | May 10, 2012, 12:57 p.m.
Spring Special Offer: Full Access to Searchmetrics Essentials for Two Weeks! Register now - free-of-charge!

Company Notes | April 18, 2012, 10:08 a.m.
Blueclaw Adopts Searchmetrics' Online Search Analytics Service To Underpin Rapid Growth

The World's Experts in Search Analytics

Searchmetrics uses the most advanced engineering to move your company up the search engine ranks. Our integrated suite of search analytics tools accesses the largest and freshest database — 100 million keywords and 75 million domains. The results are unprecedented competitive insights and actionable recommendations for targeting SEO investments with predictable ROI.

more »

Industry Endorsements

Tony Goldstone about Searchmetrics

"Searchmetrics is an invaluable and powerful SEO tool that we find enormously useful in the analysis of websites. It

SOCIAL SPREAD

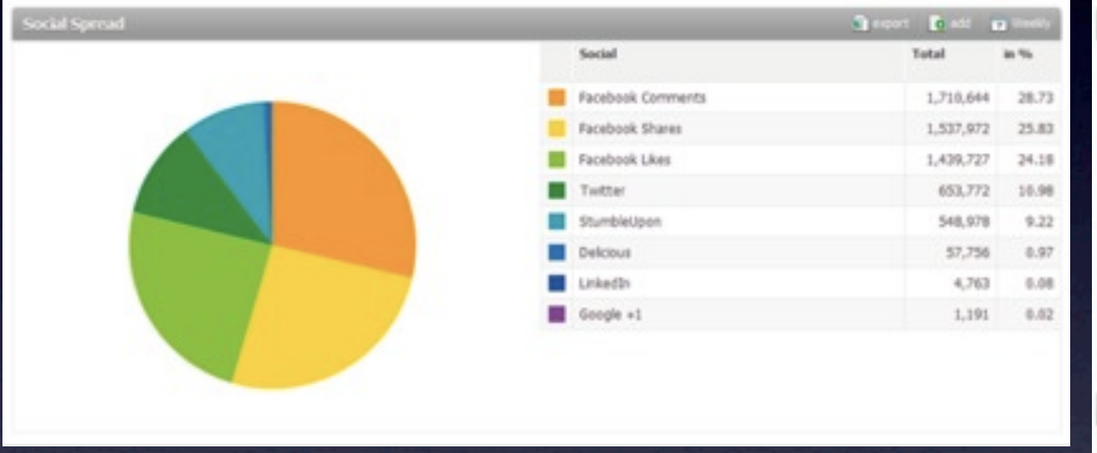

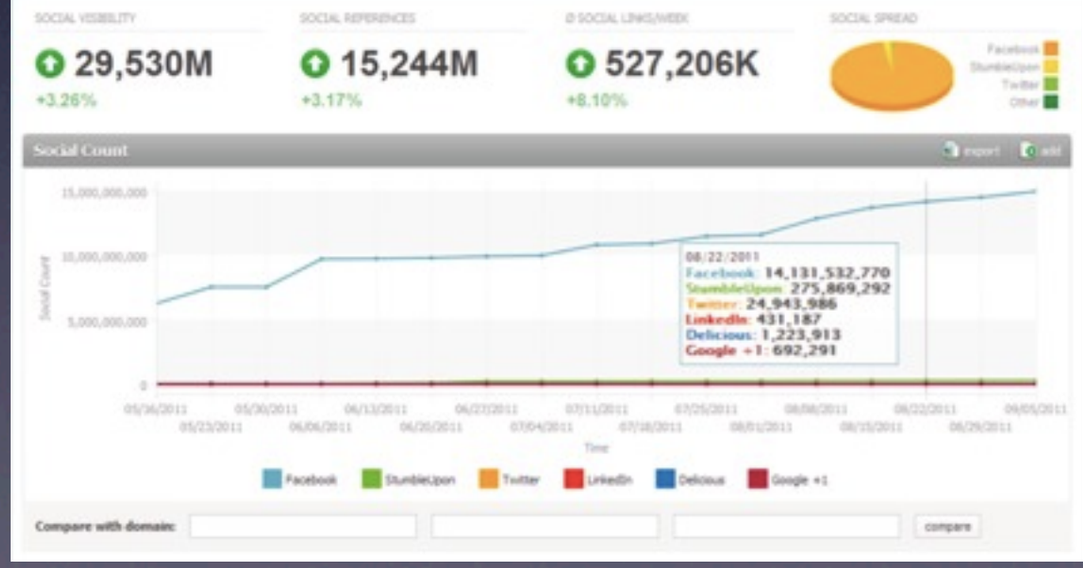

SOCIAL VISIBILITY

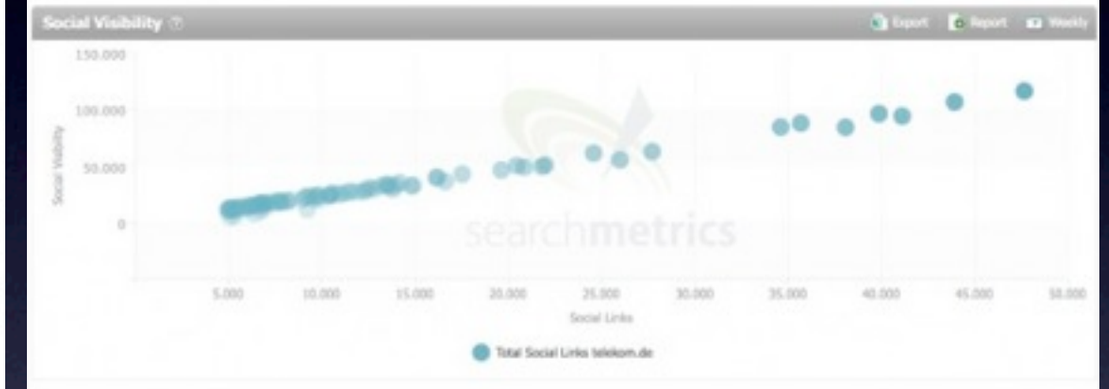

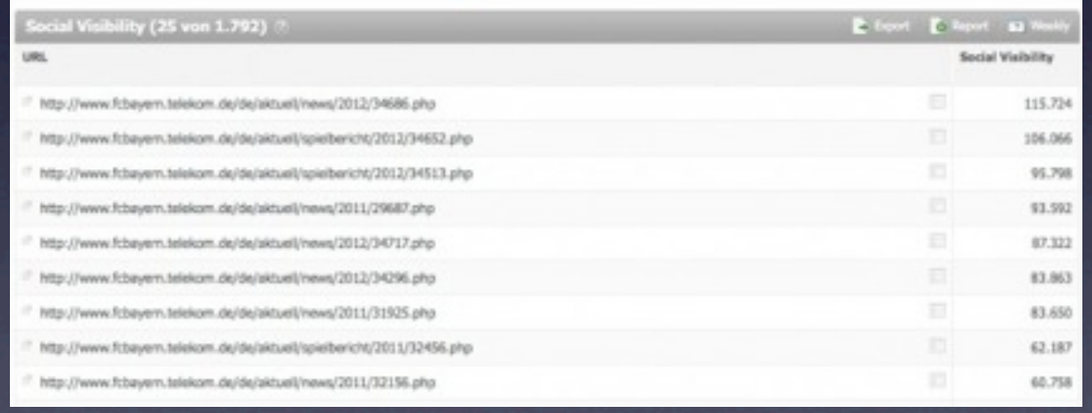

source: www.searchmetrics.com

Tweets per Day

as of September 2010

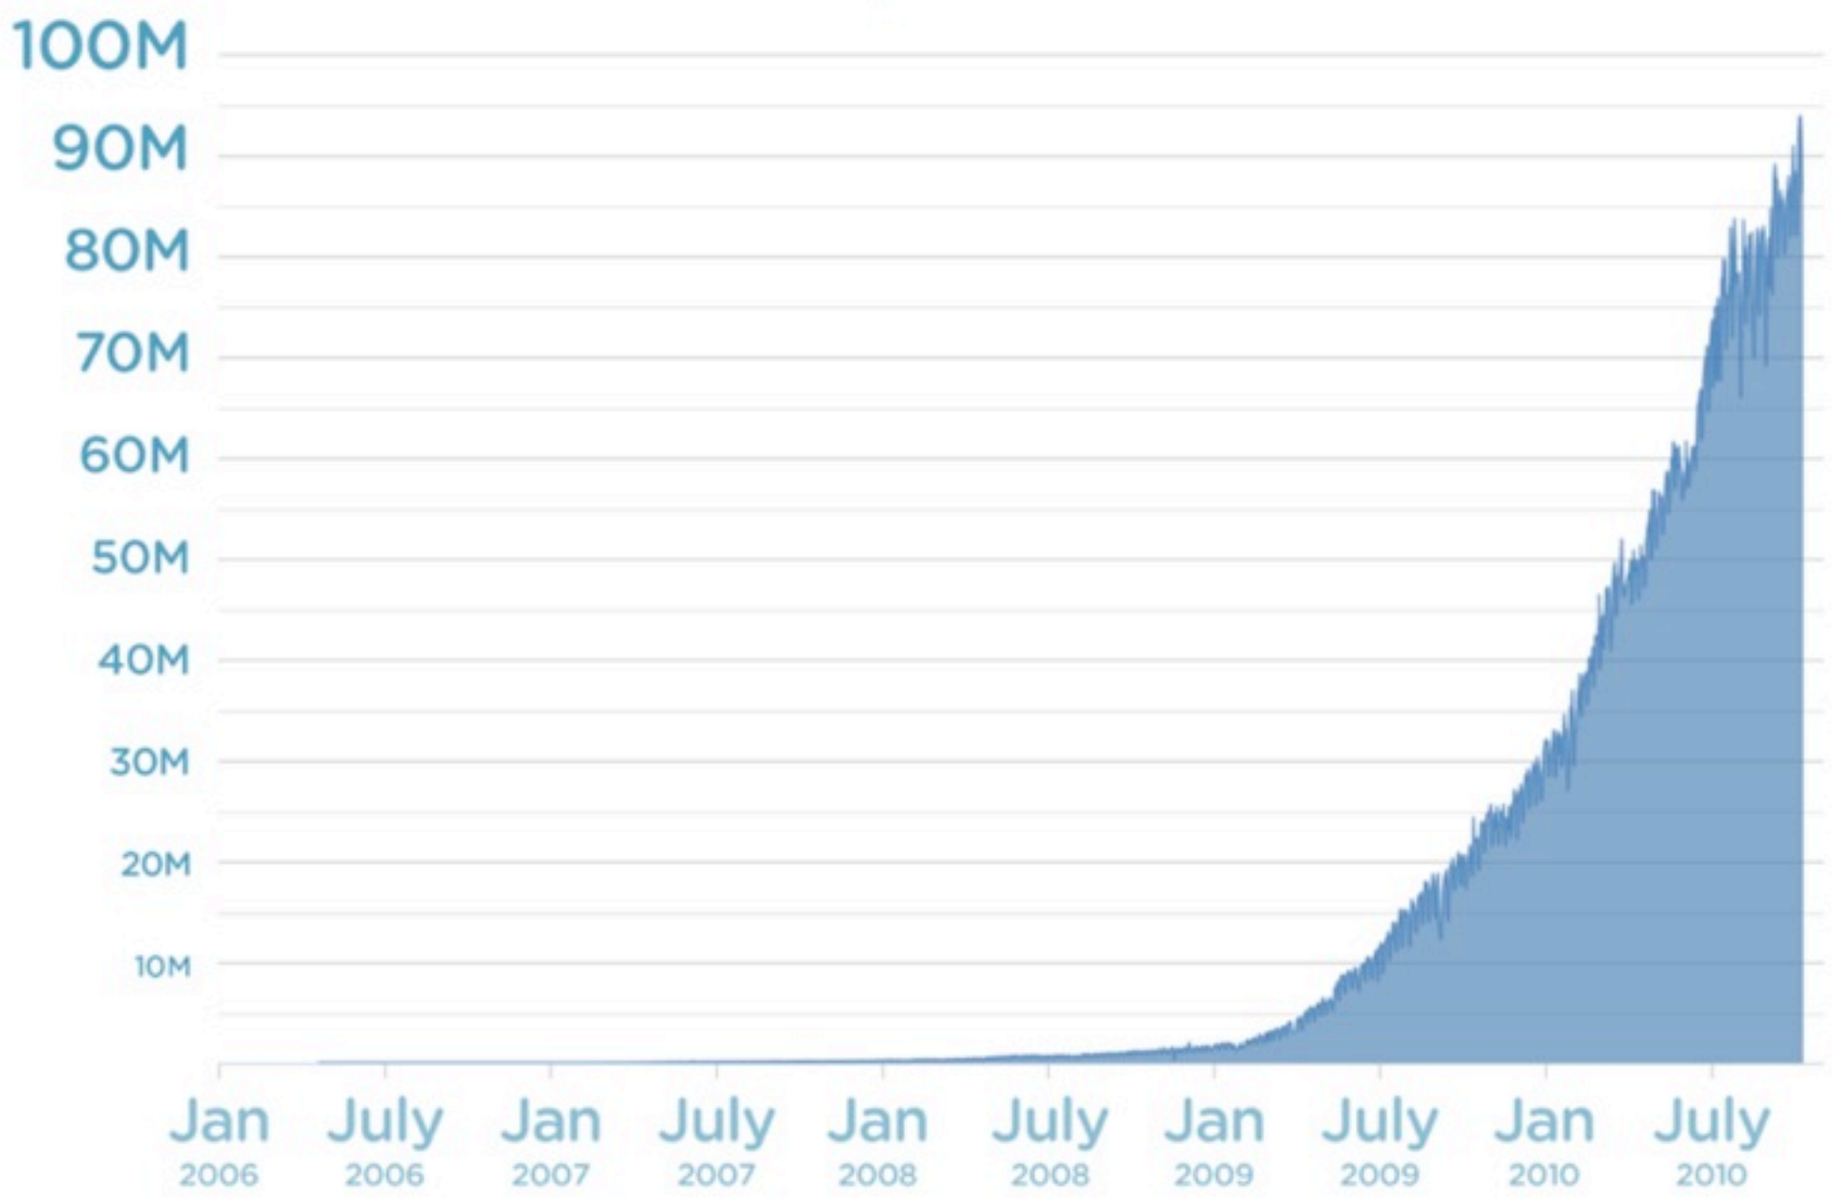

Report: Twitter hits half a billion tweets a day

In London, CEO Dick Costolo tells the audience at IAB Engage that the service now sees 500 million daily tweets, and confirms it's experimenting with a "like" button, says V3.

by Daniel Terdiman | October 26, 2012 5:51 PM PDT

The anatomy of a tweet

http://www.readwriteweb.com/archives/what_a_tweet_can_tell_you.php

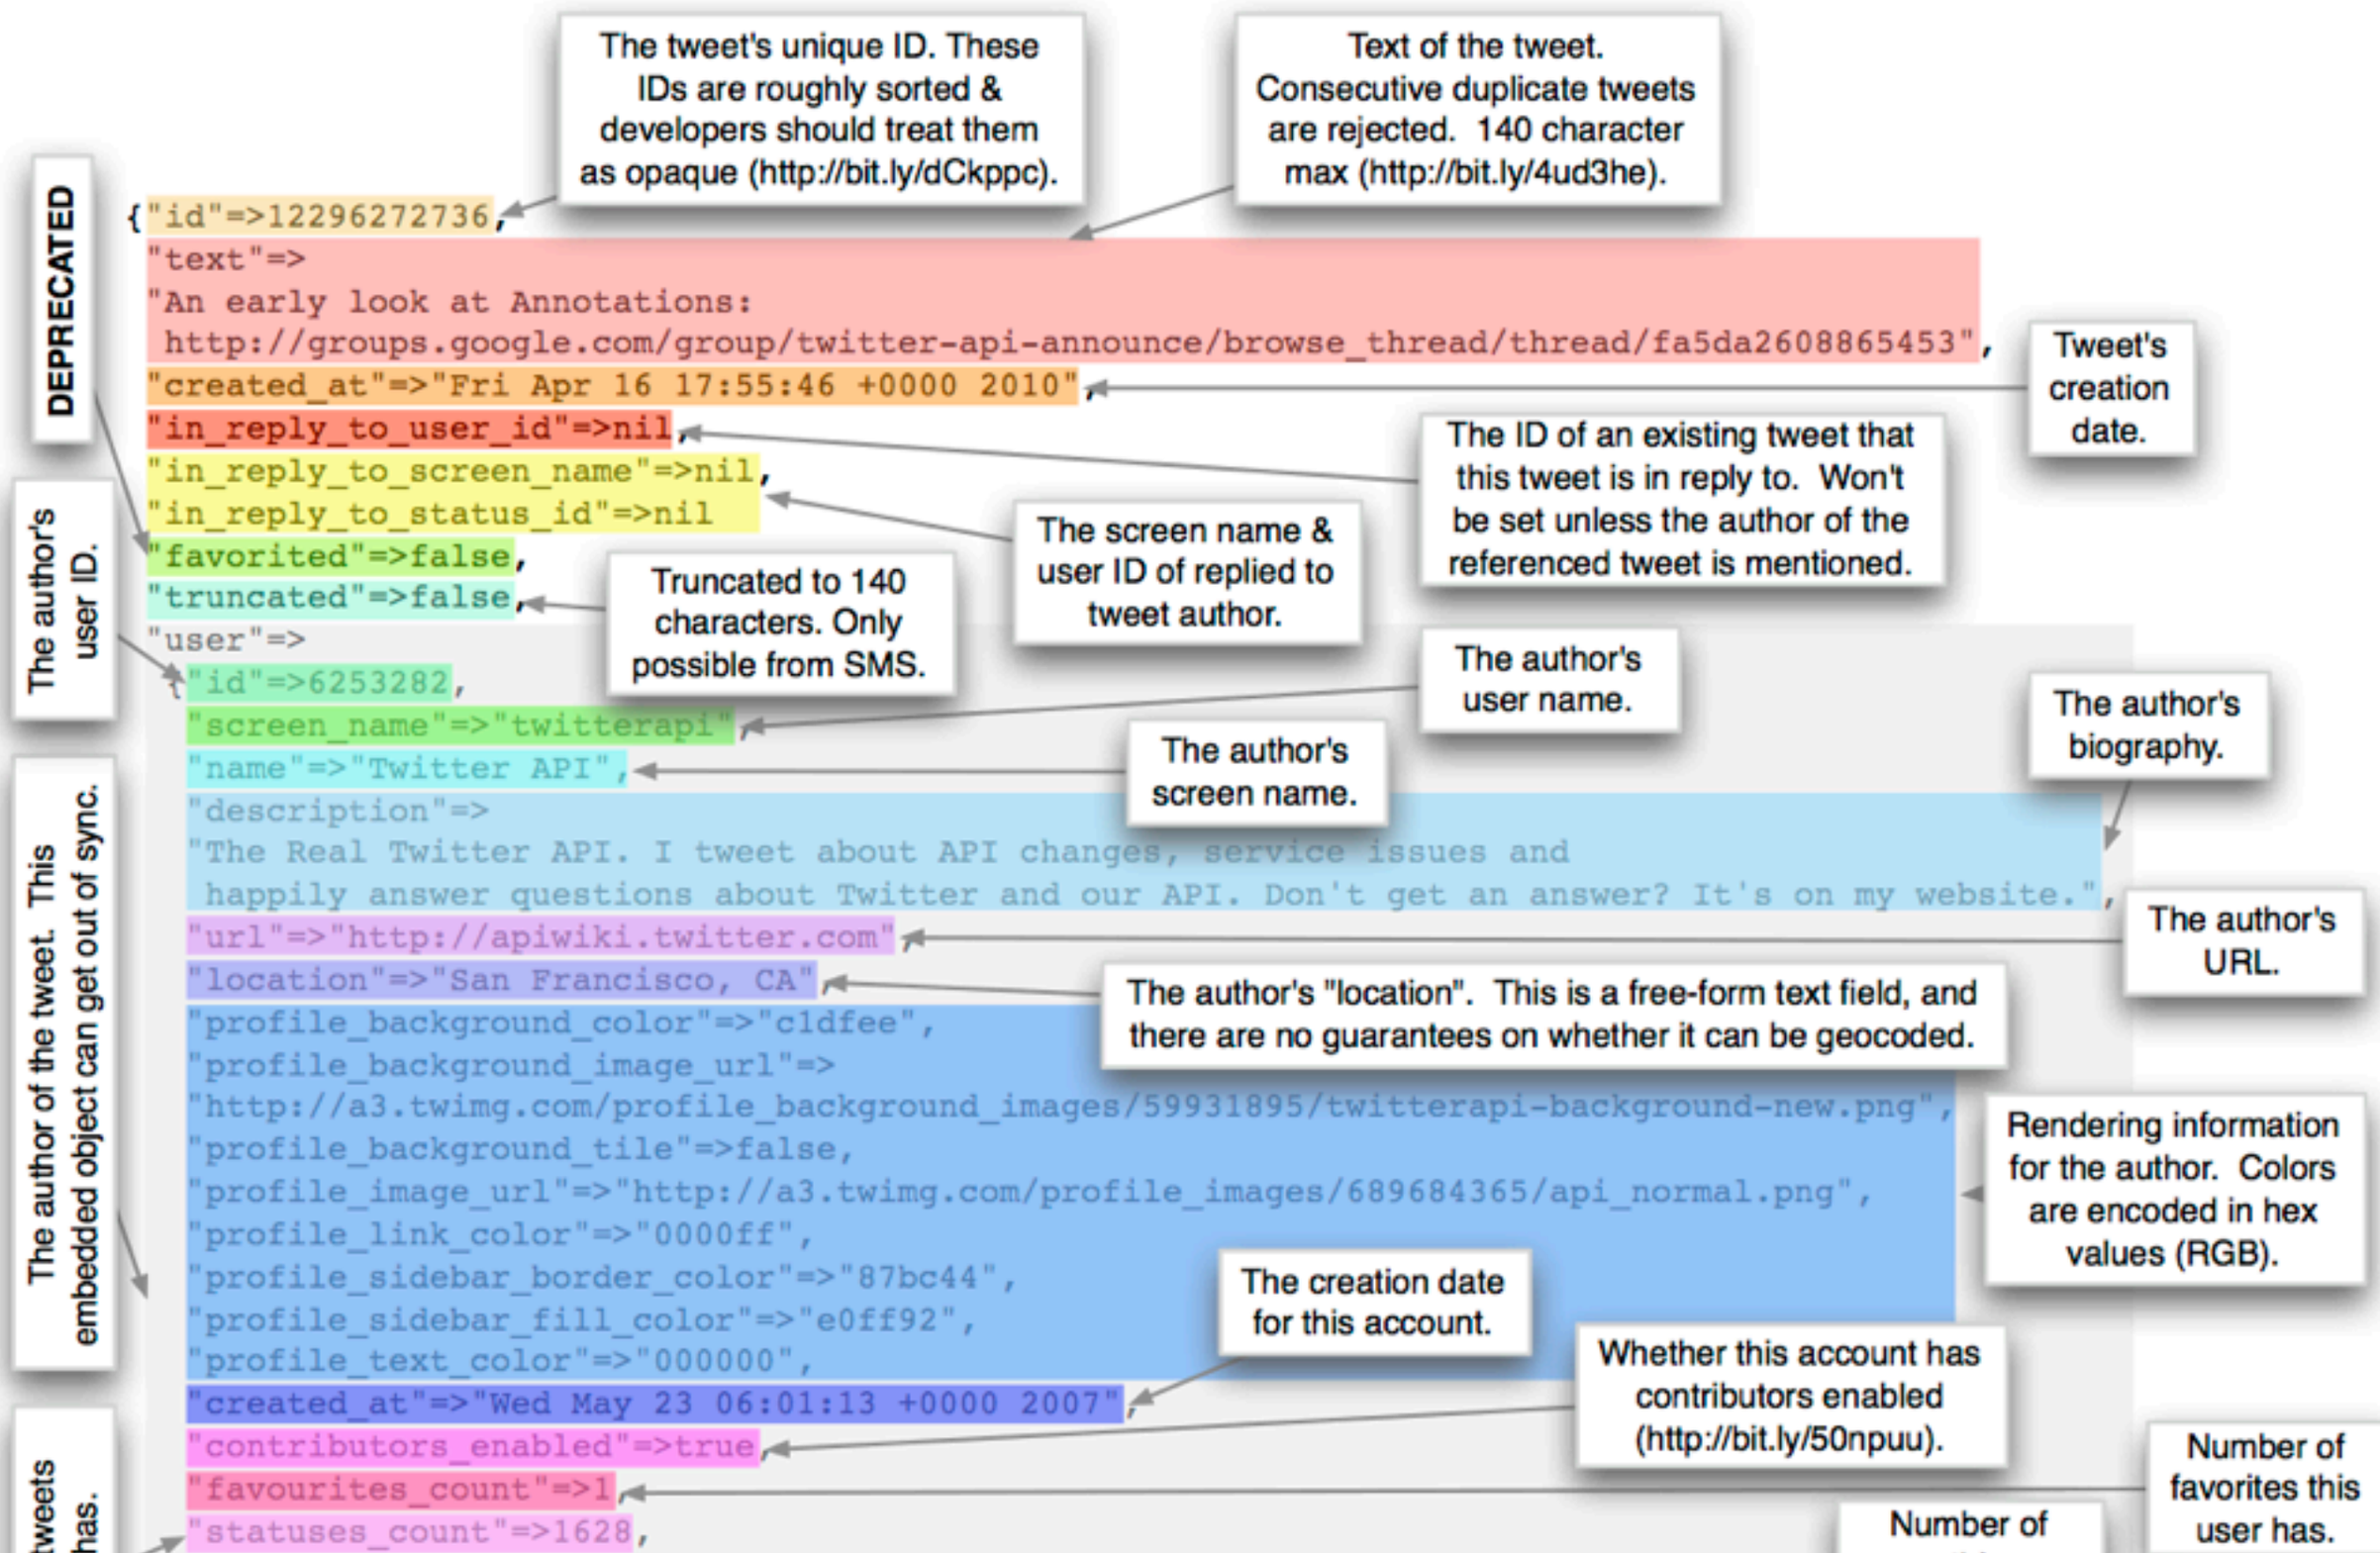

- Map/Reduce to push to Elasticsearch
- via NFS to HDFS storage
- no dedicated nodes

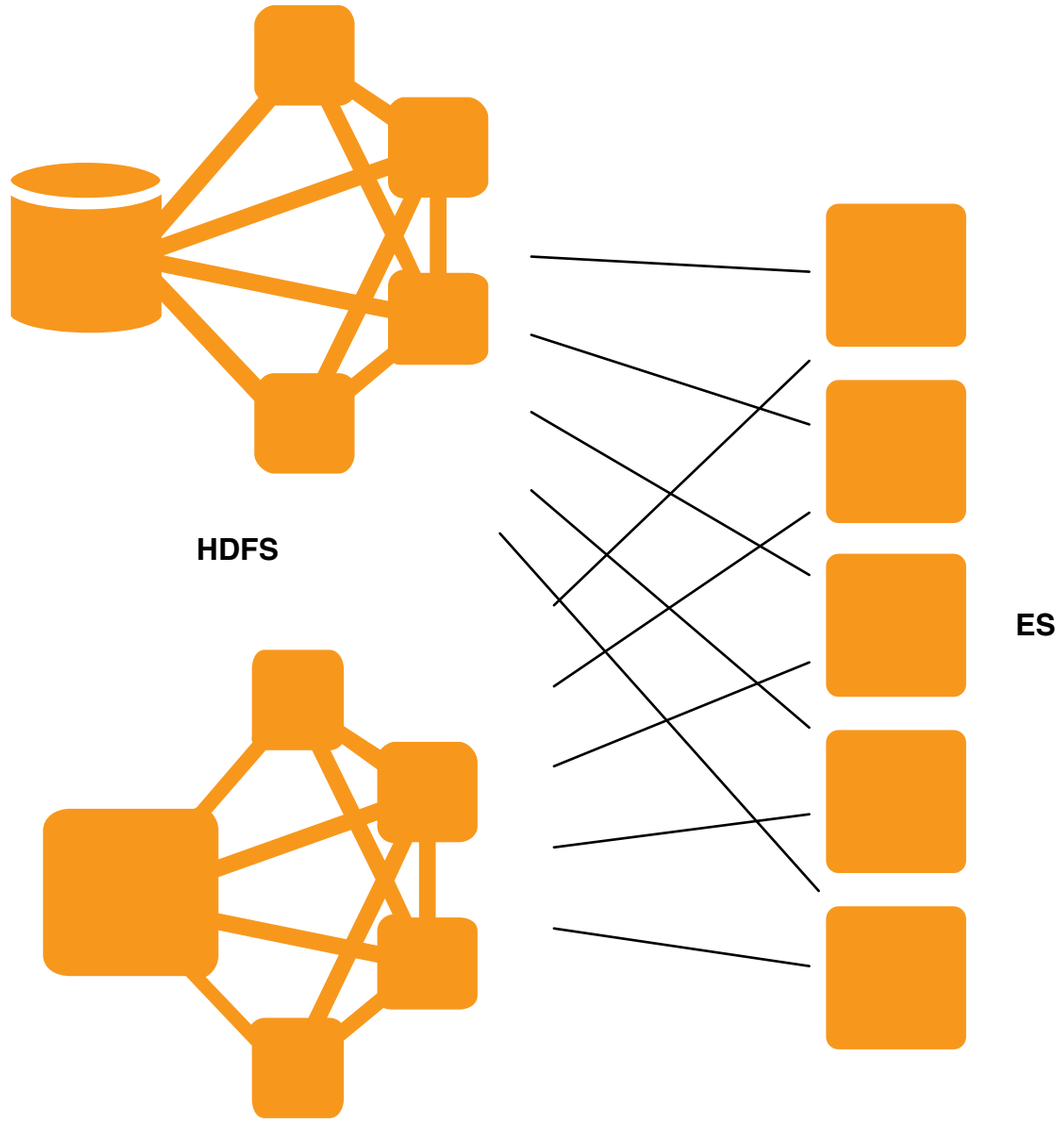

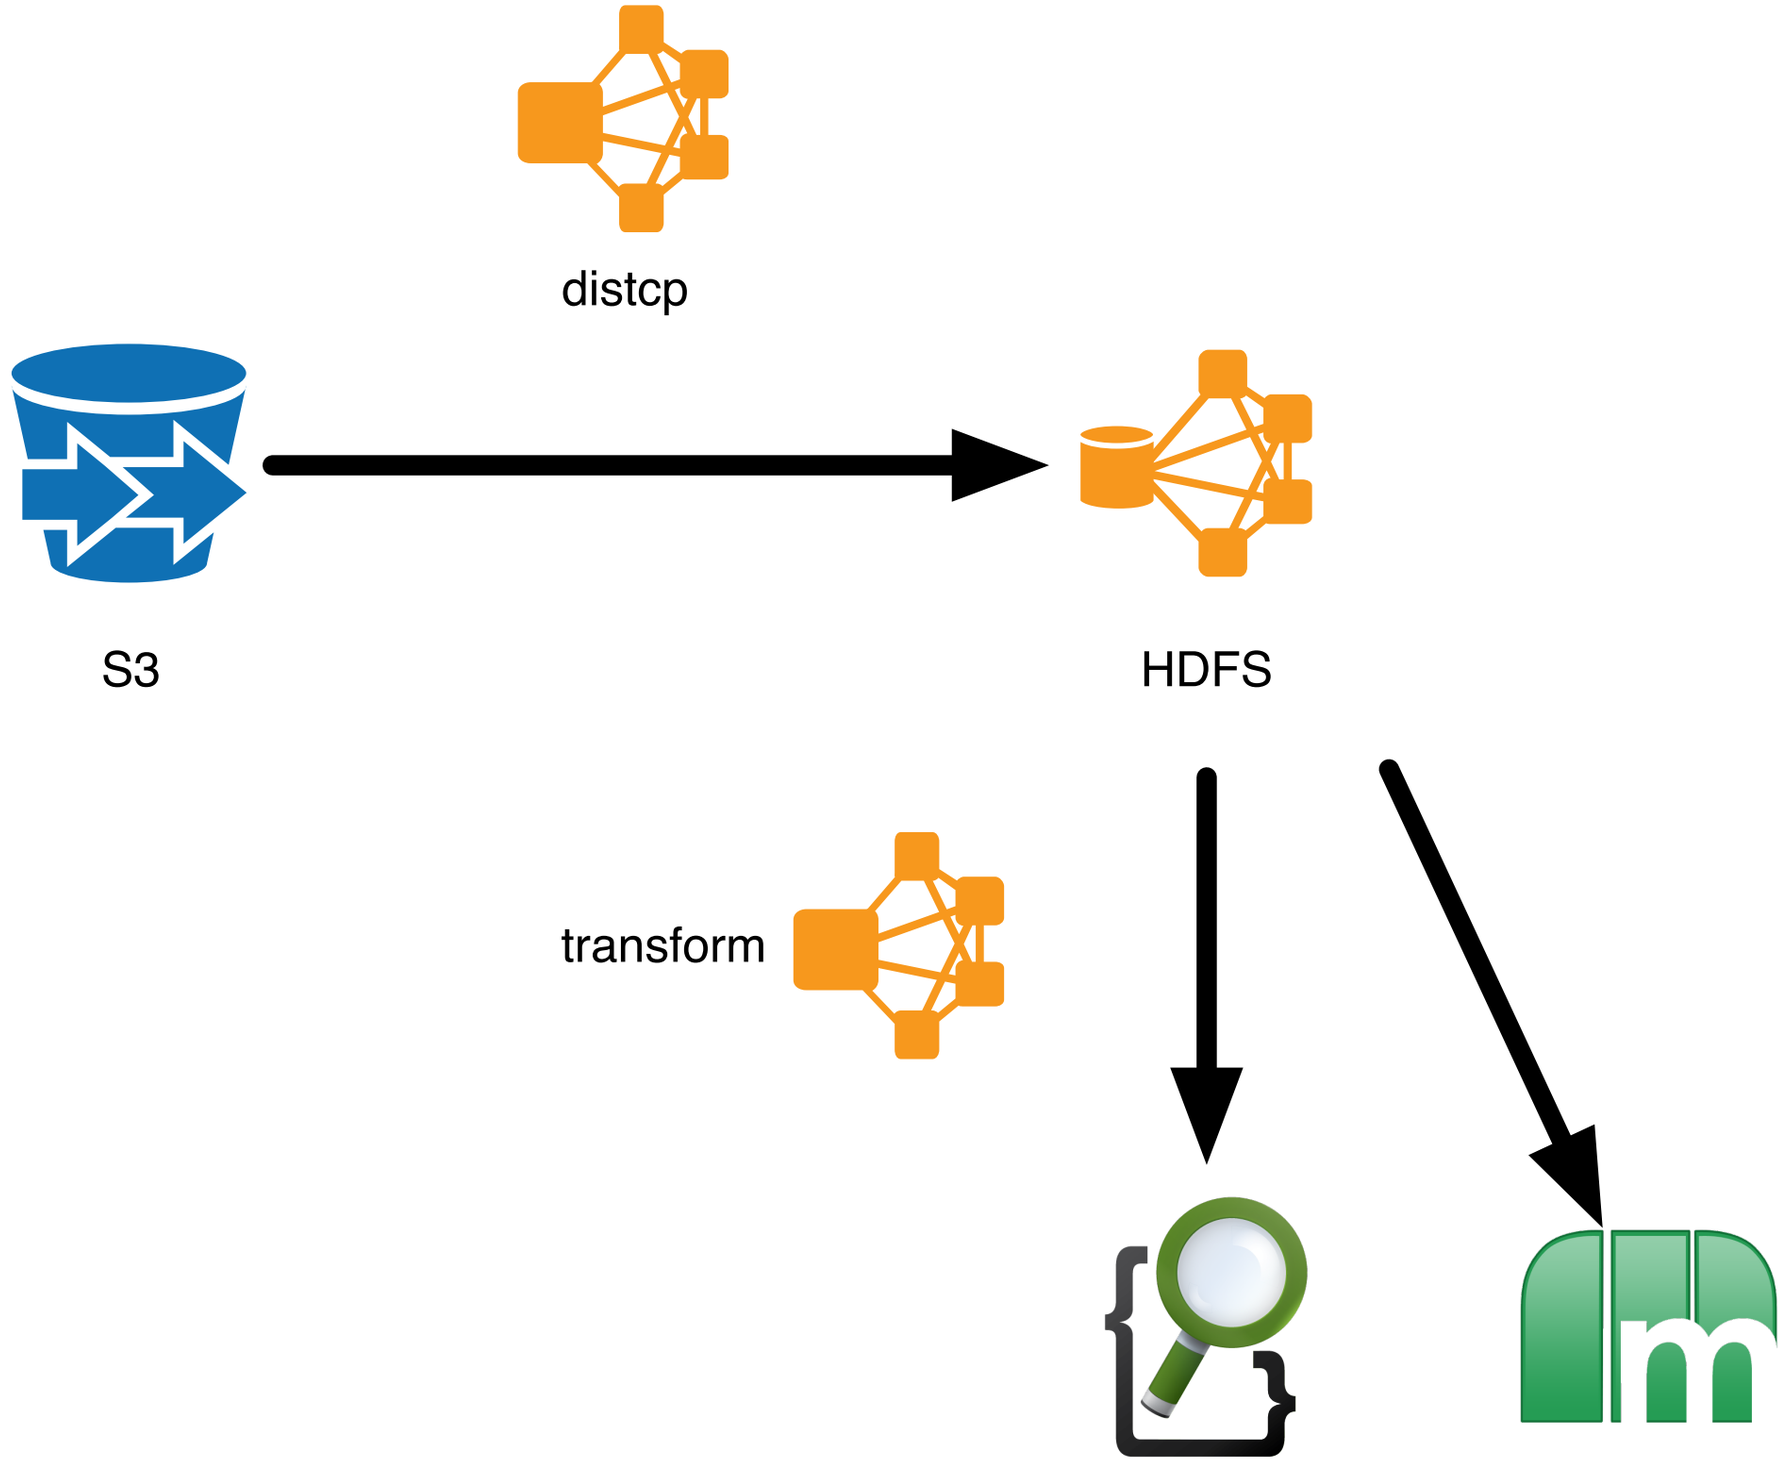

https://github.com/lovelysystems/ls-hive
https://github.com/lovelysystems/ls-thrift-py-hadoop

Memory

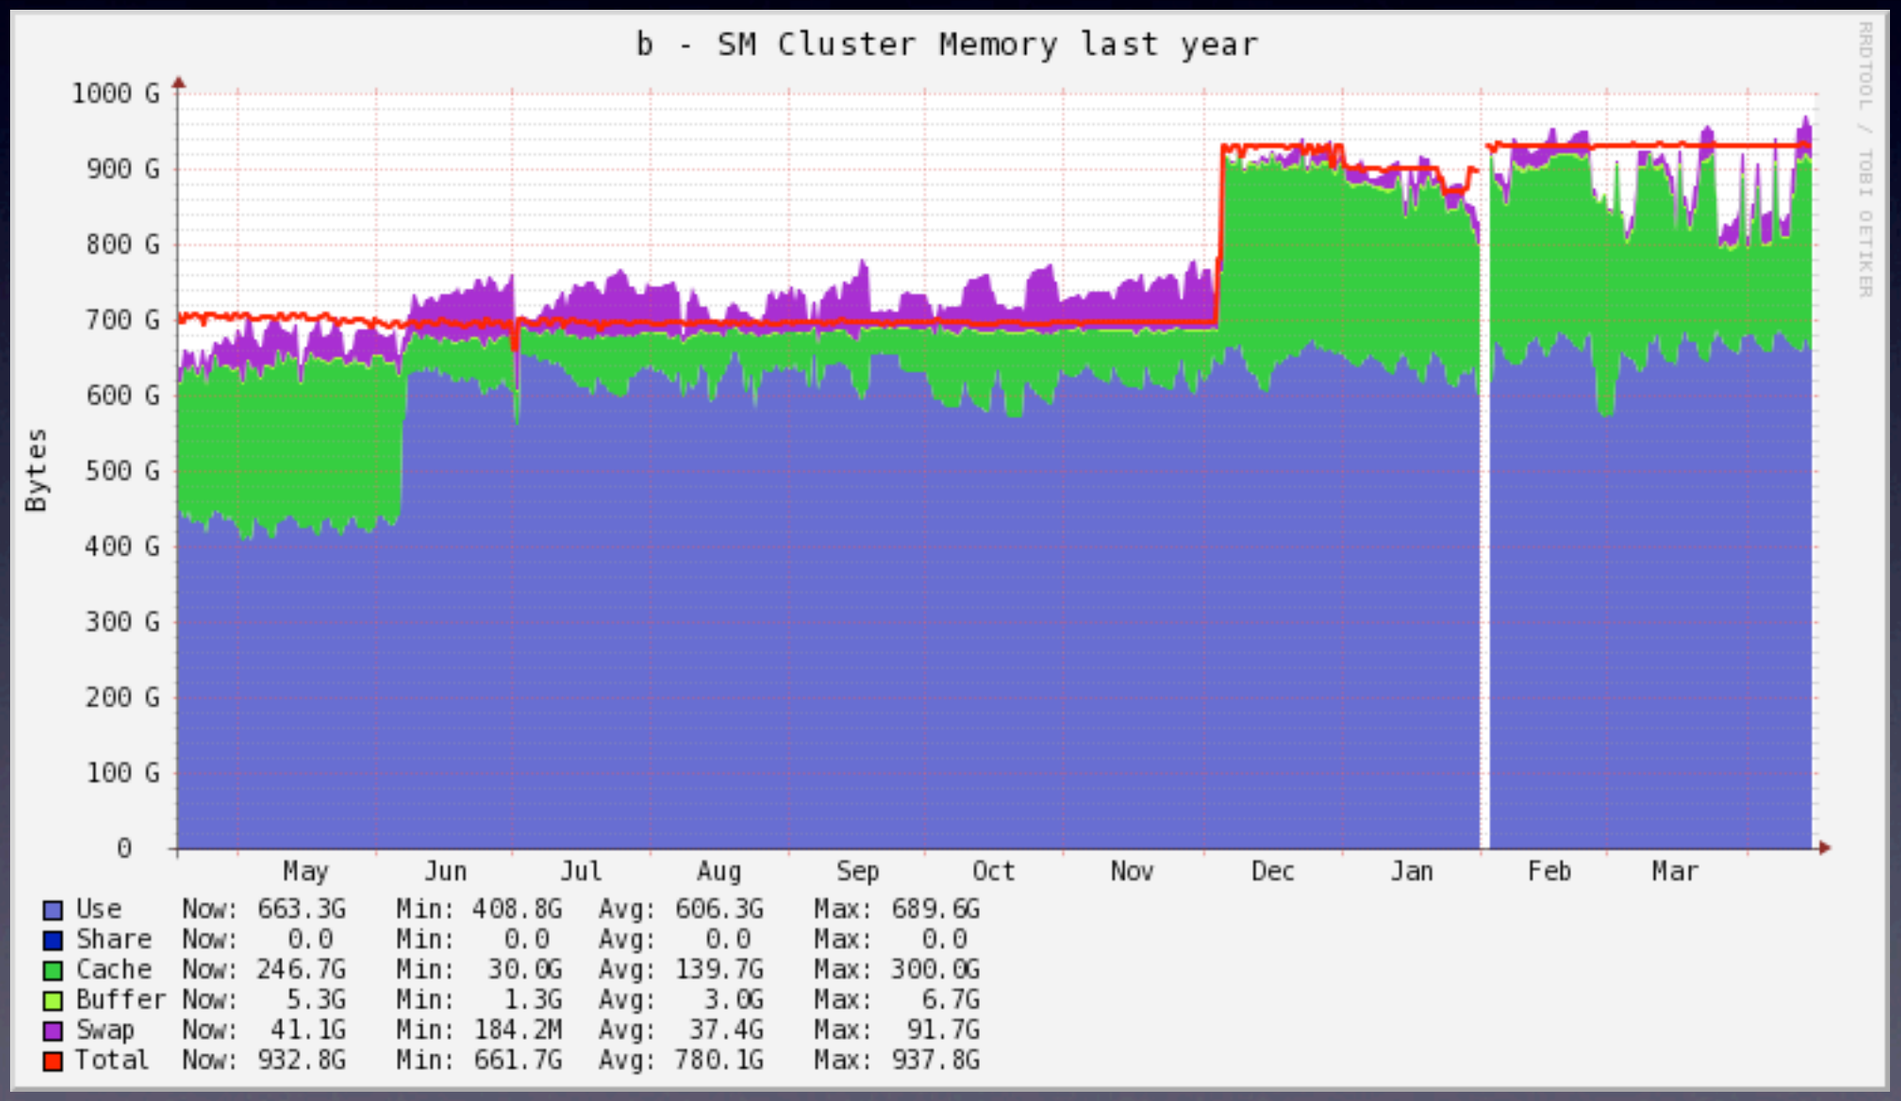
b - SM Cluster Memory last year
1000 G
900 G
800 G
700 G
600 G
500 G
400 G
300 G
200 G
100 G
0
Bytes
May
Jun
Jul
Aug
Sep
Oct
Nov
Dec
Jan
Feb
Mar
Use Now: 663.3G Min: 408.8G Avg: 606.3G Max: 689.6G
Share Now: 0.0 Min: 0.0 Avg: 0.0 Max: 0.0
Cache Now: 246.7G Min: 30.0G Avg: 139.7G Max: 300.0G
Buffer Now: 5.3G Min: 1.3G Avg: 3.0G Max: 6.7G
Swap Now: 41.1G Min: 184.2M Avg: 37.4G Max: 91.7G
Total Now: 932.8G Min: 661.7G Avg: 780.1G Max: 937.8G
RRDTOOL / TOBI OETIKER

Network

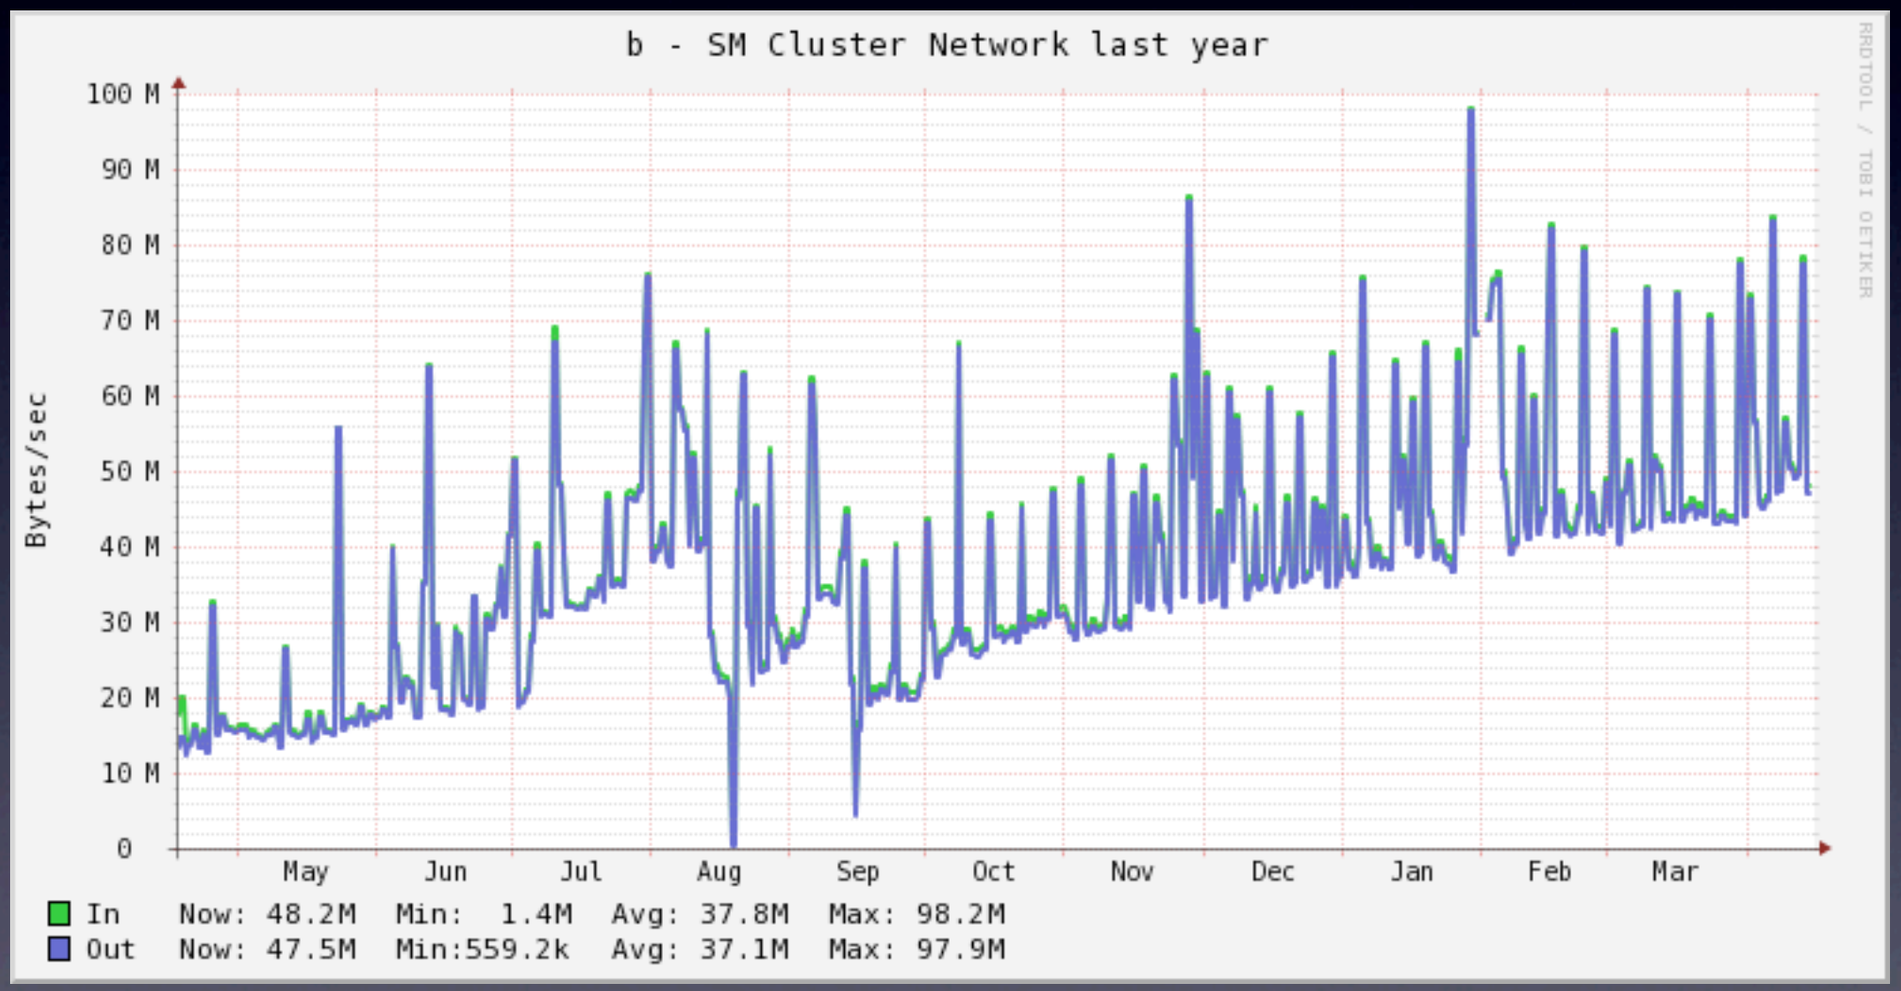

- Daily peaks are 600MByte/s

response time

Overall Average	Slowest Average	Fastest Average
538 ms	804 ms	338 ms

stackoverflow

Questions | Tags | Users | Badges | Unanswered

ElasticSearch vs. MongoDB for Caching User Data

2

1

Up to this point, I have been using MongoDB (Node.js + Mongoose) to save posts which belong to a user, so that I can later retrieve them to display in a stream (just like Facebook, Twitter, etc.)

It recently became necessary to allow the user to deeply search his stream; MongoDB's search was insufficient, so I implemented ElasticSearch on my servers (Amazon EC2 m1.large instances running CentOS, FWIW).

My question: I'm now in a position that I'm duplicating the data between MongoDB (where the user's stream is cached) and ElasticSearch (where it is searched).

Is there any disadvantage to moving my cache ENTIRELY into ElasticSearch, getting rid of the MongoDB all together? It seems a waste to double the storage, and there's no other place that I'm accessing this data (it is only used when presenting/searching the stream of posts).

Specifically, I want to make sure I'm not overlooking anything re: performance. I like the idea of reducing MongoDB as a bottleneck, yet I worry about the memory overhead of ElasticSearch. MongoDB runs on its own server in my cloud setup, whereas ElasticSearch is running on the same instances as node.js. This means I would have **MORE** ElasticSearch servers (the node.js servers are in an auto-scaling array), but they each are not **DEDICATED** servers (unlike MongoDB).

mongodb | elasticsearch

share | improve this question

asked Feb 1 at 18:00

Zane Claes

2,391 ●2 ●6 ●18

The only big obstacle to using ES as a "primary datasource" is that there isn't a good backup mechanism right now.

Gigaset elements

Über elements Blog FAQ Kontakt

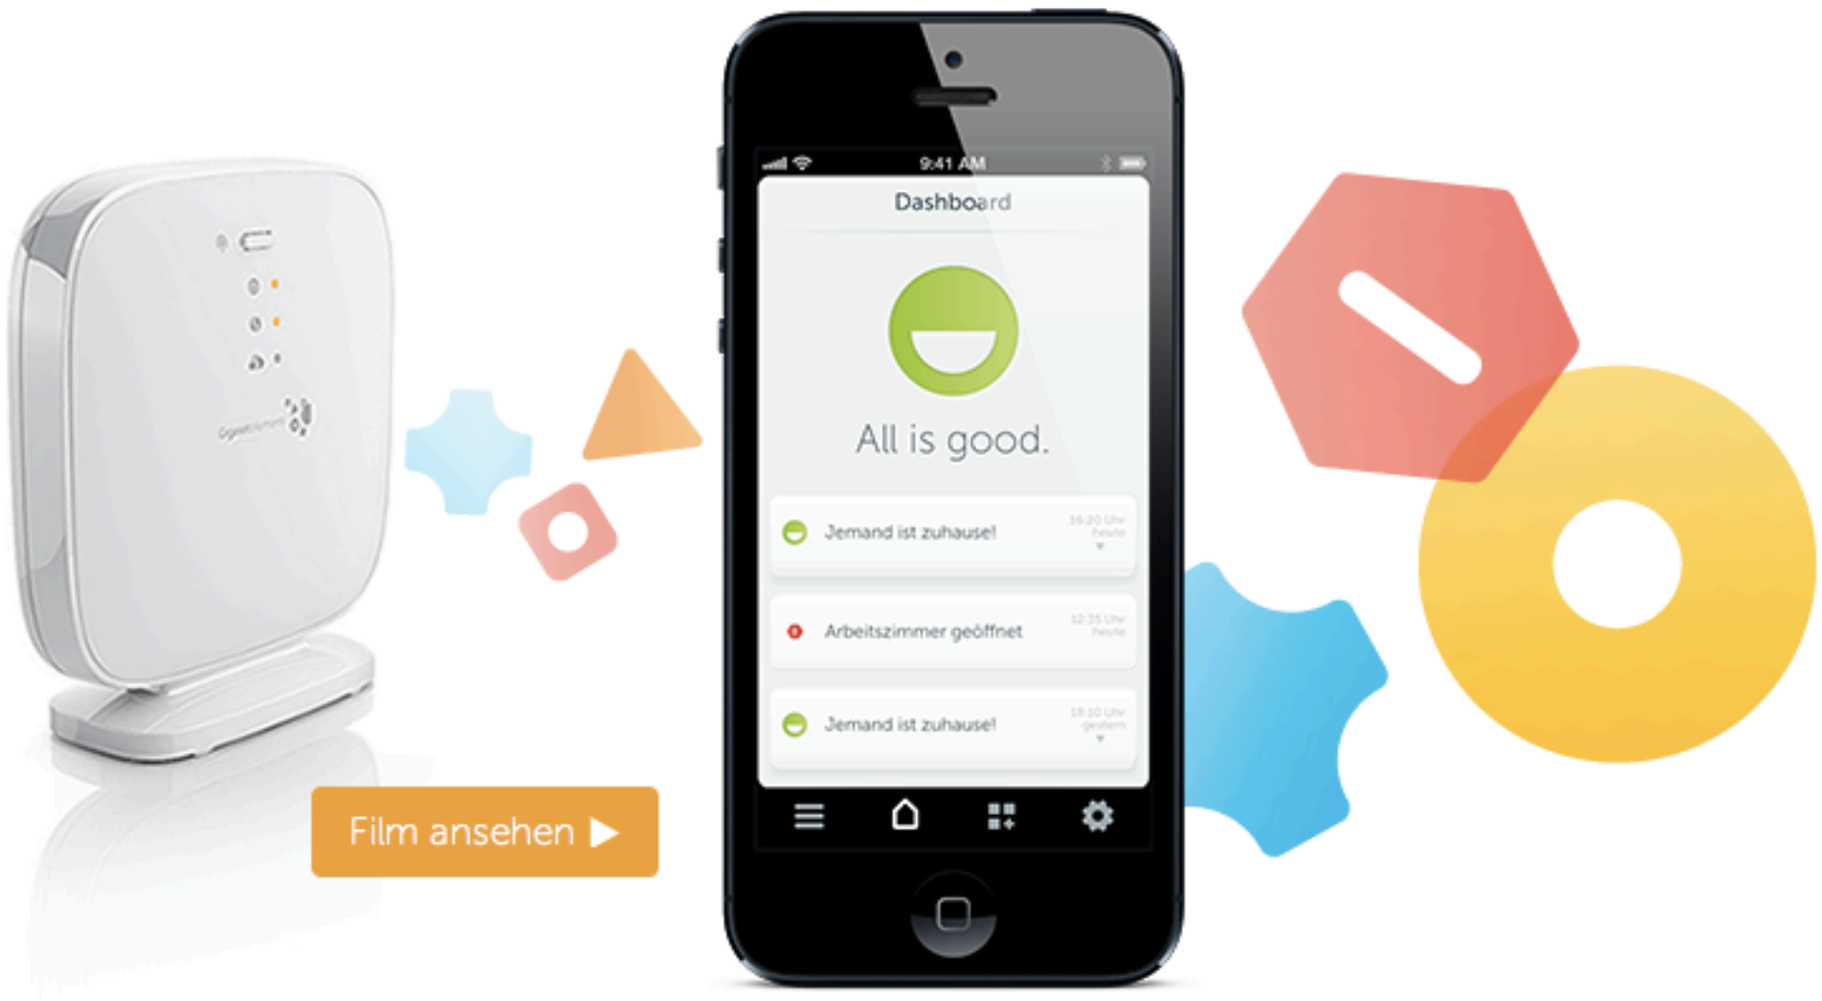

Dein Zuhause und Du werdet ein Team.

Gigaset

Overview

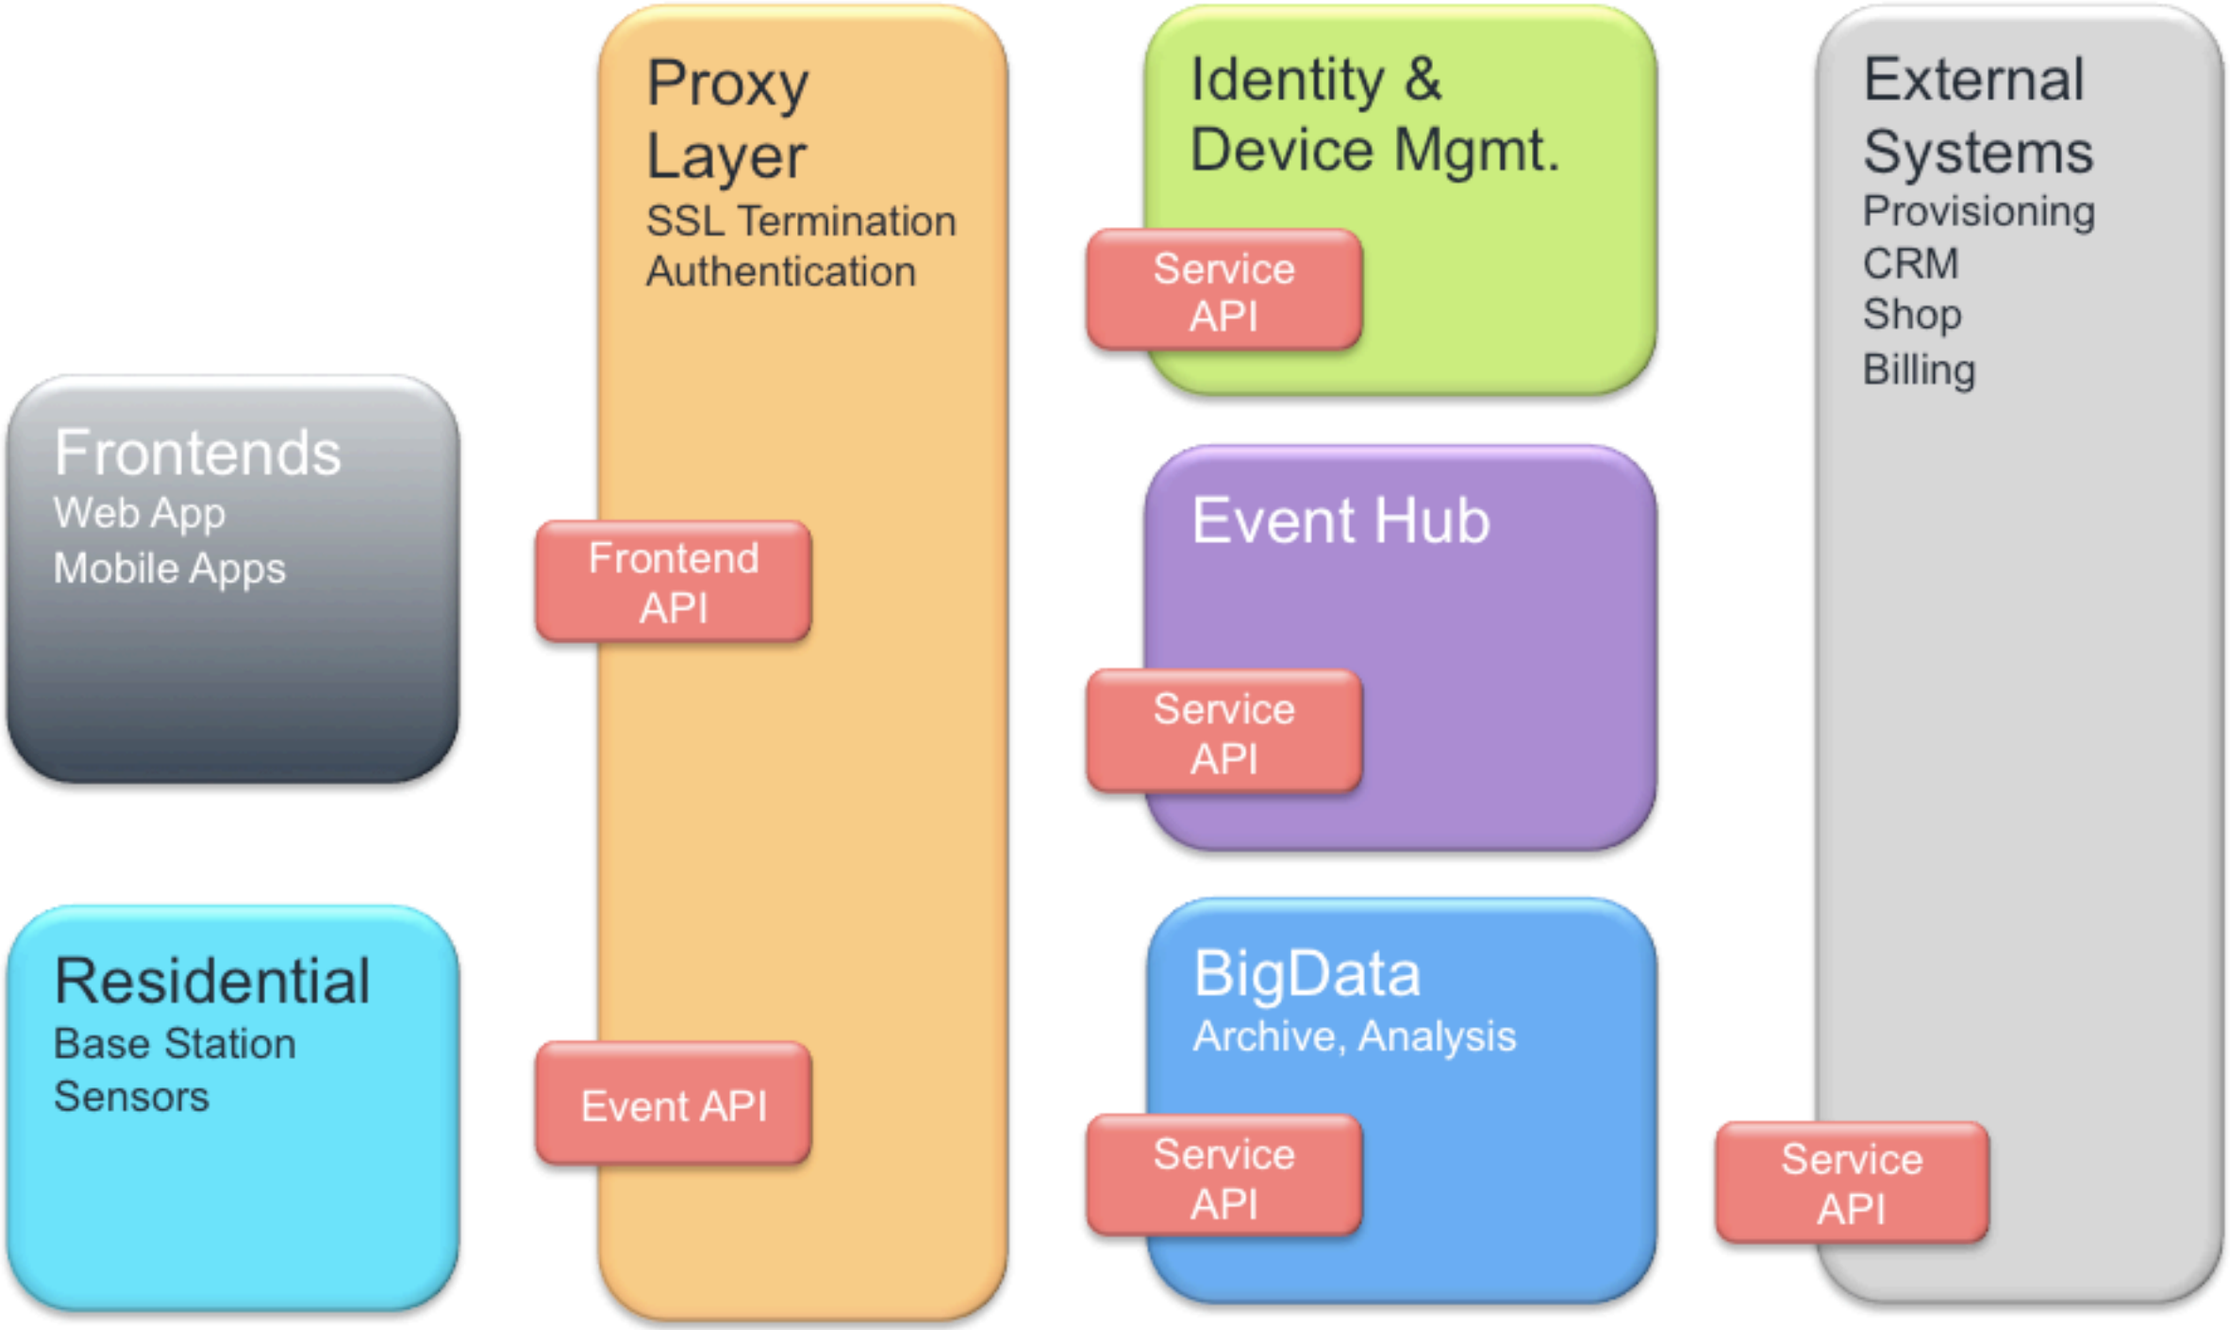

Elasticsearch Settings

- discovery.zen.minimum_master_nodes
- index.mapper.dynamic

Index Settings

- number_of_shards
- number_of_replicas
- replication (async)
- refresh_interval / _optimize / segments
- merge.policy.merge_factor
- routing.allocation.total_shards_per_node

Templates

```
curl -XPUT localhost:9200/_template/template_1 -d '
{
    "template" : "te*",
    "settings" : {
        "number_of_shards" : 1
    },
    "mappings" : {
        "type1" : {
            "_source" : { "enabled" : false }
        }
    }
}
'
```

Time based indexes

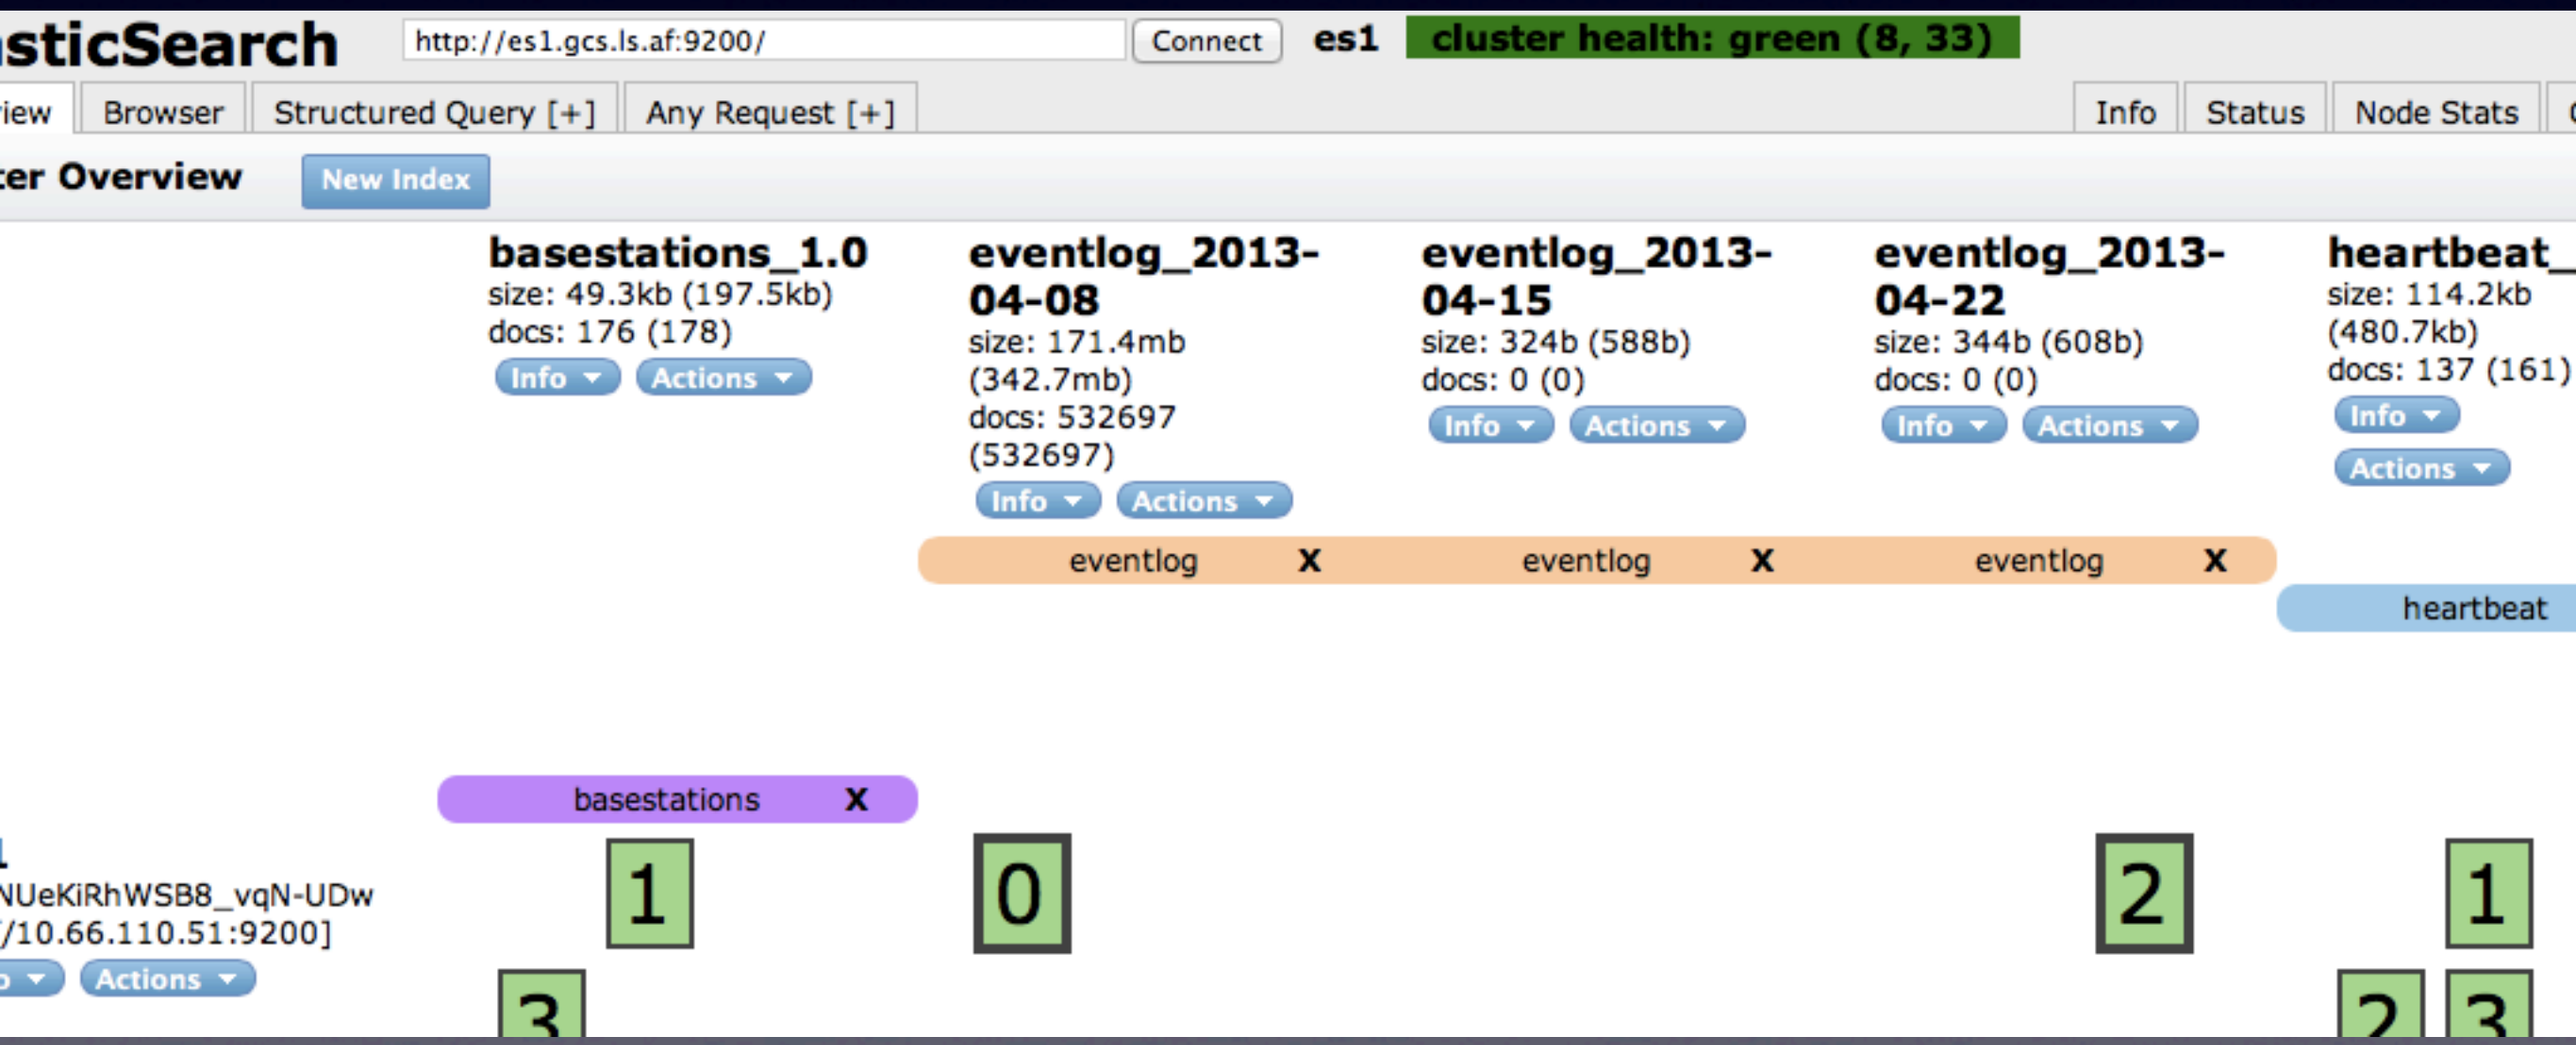

Monitor

```
{
  - fe: {
        status: "GREEN",
        description: "2/2 online "
    },
  - ap: {
        status: "GREEN",
        description: "2/2 online "
    },
  - es: {
        status: "GREEN",
        description: "(8, 130)"
    },
  - sb: {
        status: "RED",
        description: "1/8 online "
    }
}
```

What's next?

- stay tuned, follow @jodok
- will launch "Elastic Search as Primary Storage" and Tools as Package

Lovely Systems

That's thirty minutes away.
I'll be there in ten.

@jodok