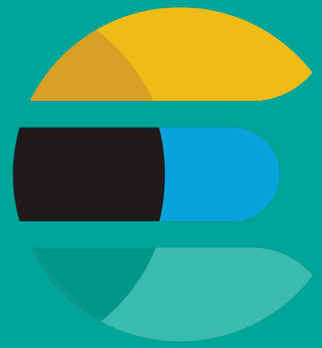

elasticsearch

(R)Evolution

Philipp Krenn

@xeraa

elastic

Revolution

elastic

341 systems in ranking, February 2018

Rank			DBMS	Database Model	Score		
Feb 2018	Jan 2018	Feb 2017			Feb 2018	Jan 2018	Feb 2017
1.	1.	1.	Oracle	Relational DBMS	1303.28	-38.66	-100.55
2.	2.	2.	MySQL	Relational DBMS	1252.47	-47.24	-127.83
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1122.04	-26.03	-81.42
4.	4.	4.	PostgreSQL	Relational DBMS	388.38	+2.19	+34.70
5.	5.	5.	MongoDB	Document store	336.42	+5.47	+0.92
6.			IBM Db2	Relational DBMS	180.97		+2.07
7.			Microsoft Access	Relational DBMS	150.07		-3.32
8.	↑ 9.	↑ 10.	Redis	Key-value store	127.02	+3.88	+12.98
9.	↑ 10.	↑ 11.	Elasticsearch	Search engine	125.32	+2.76	+17.01
10.	↓ 8.	↓ 7.	Cassandra	Wide column store	122.78	-1.10	-11.60
11.	11.	↓ 9.	SQLite	Relational DBMS	117.27	+3.02	+1.96
12.	12.	12.	Teradata	Relational DBMS	72.99	+0.36	-2.60
13.	↑ 15.	↑ 16.	Splunk	Search engine	67.27	+3.27	+11.24
14.	14.	14.	Solr	Search engine	63.87	-0.50	-3.81
15.	↓ 13.	↓ 13.	SAP Adaptive Server	Relational DBMS	63.49	-1.98	-8.25

341 systems in ranking, February 2018

Rank			DBMS	Database Model	Score		
Feb 2018	Jan 2018	Feb 2017			Feb 2018	Jan 2018	Feb 2017
1.	1.	1.	Oracle	Relational DBMS	1303.28	-38.66	-100.55
2.	2.	2.	MySQL	Relational DBMS	1252.47	-47.24	-127.83
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1122.04	-26.03	-81.42
4.	4.	4.	PostgreSQL	Relational DBMS	388.38	+2.19	+34.70
5.	5.	5.	MongoDB	Document store	336.42	+5.47	+0.92
6.	6.	6.	DB2	Relational DBMS	189.97	-0.30	+2.07
7.	7.	↑ 8.	Microsoft Access	Relational DBMS	130.07	+3.37	-3.32
8.	↑ 9.	↑ 10.	Redis	Key-value store	127.02	+3.88	+12.98
9.	↑ 10.	↑ 11.	Elasticsearch	Search engine	125.32	+2.76	+17.01
10.	↓ 8.	↓ 7.	Cassandra	Wide column store	122.78	-1.10	-11.60
11.	11.	↓ 9.	SQLite	Relational DBMS	117.27	+3.02	+1.96
12.	12.	12.	Teradata	Relational DBMS	72.99	+0.36	-2.60
13.	↑ 15.	↑ 16.	Splunk	Search engine	67.27	+3.27	+11.24
14.	14.	14.	Solr	Search engine	63.87	-0.50	-3.81
15.	↓ 13.	↓ 13.	SAP Adaptive Server	Relational DBMS	63.49	-1.98	-8.25

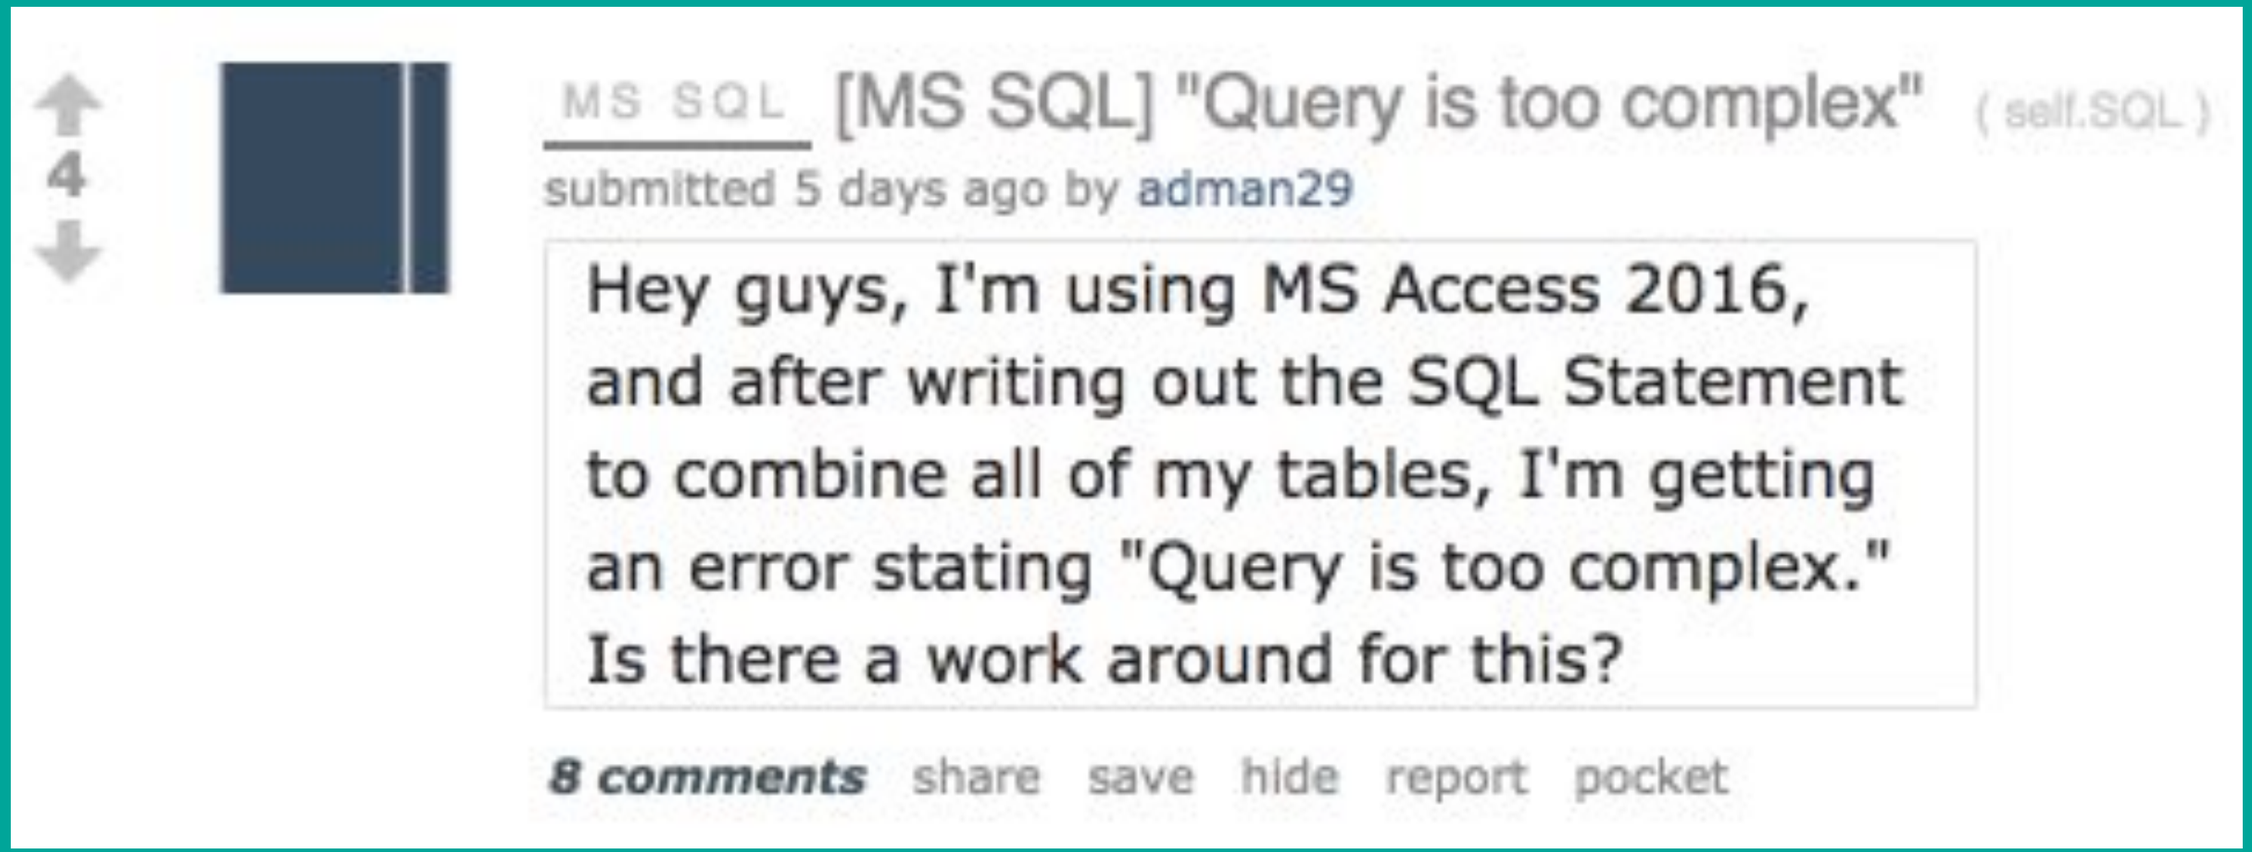

https://www.reddit.com/r/SQL/comments/7i828i/ms_sql_query_is_too_complex/

Who uses Elasticsearch?

elastic

Infrastructure | Developer Advocate

http://thedudeabides.com/articles/the_birth_of_compass

elastic

elasticsearch.

```
$ curl http://localhost:9200
{
  "name": "elasticsearch1",
  "cluster_name": "docker-cluster",
  "cluster_uuid": "sW2BcTZRRCSVVH2sDsWF8A",
  "version": {
    "number": "5.6.7",
    "build_hash": "4669214",
    "build_date": "2018-01-25T21:14:50.776Z",
    "build_snapshot": false,
    "lucene_version": "6.6.1"
  },
  "tagline": "You Know, for Search"
}
```

```
$ curl http://localhost:9200
{
  "name": "elasticsearch1",
  "cluster_name": "docker-cluster",
  "cluster_uuid": "sW2BcTZRRCSVVH2sDsWF8A",
  "version": {
    "number": "5.6.7",
    "build_hash": "4669214",
    "build_date": "2018-01-25T21:14:50.776Z",
    "build_snapshot": false,
    "lucene_version": "6.6.1"
  },
  "tagline": "You Know, for Search"
}
```

logstash

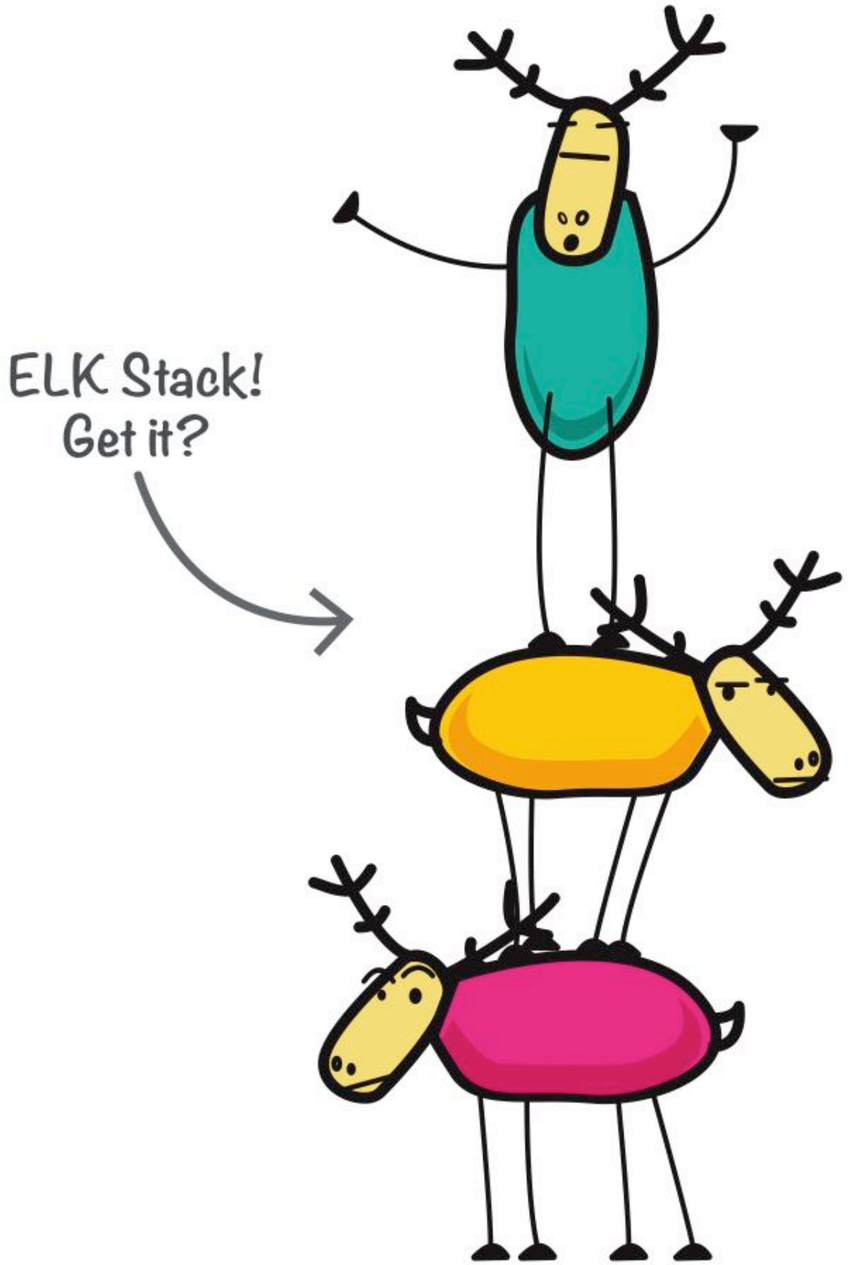

E Elasticsearch

L Logstash

K Kibana


mozilla


slack


BLIZZARD
ENTERTAINMENT

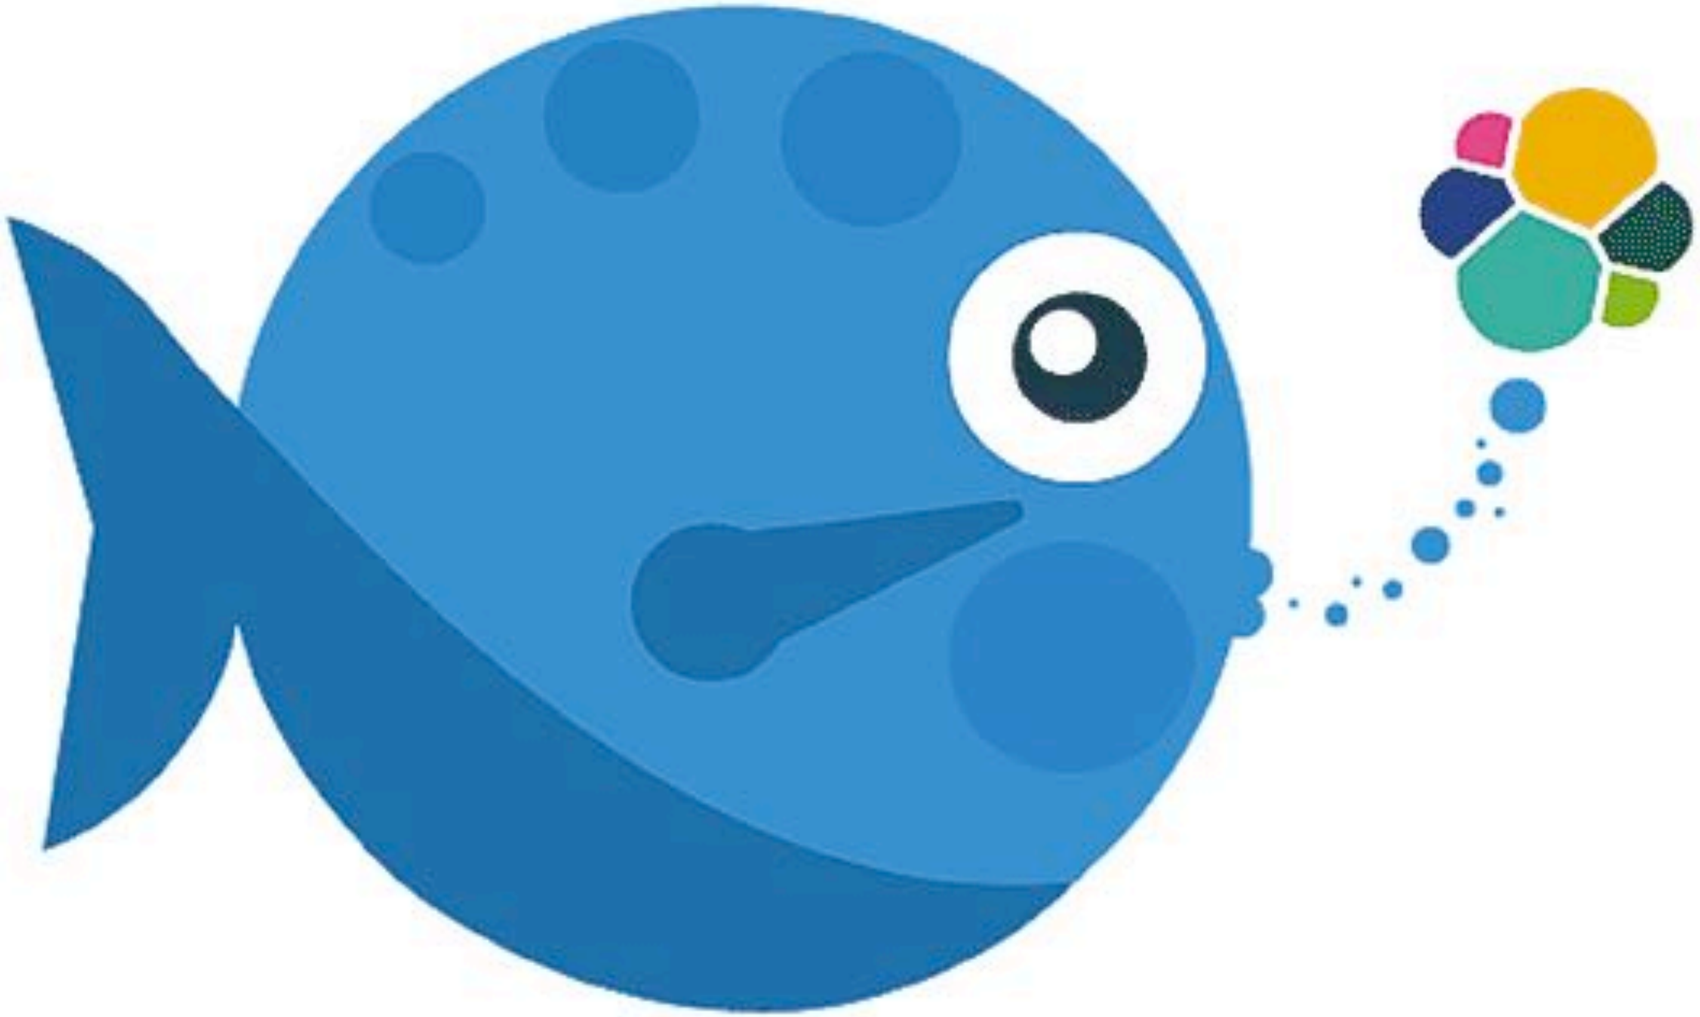

Apparently, I'm an ELKB personality.

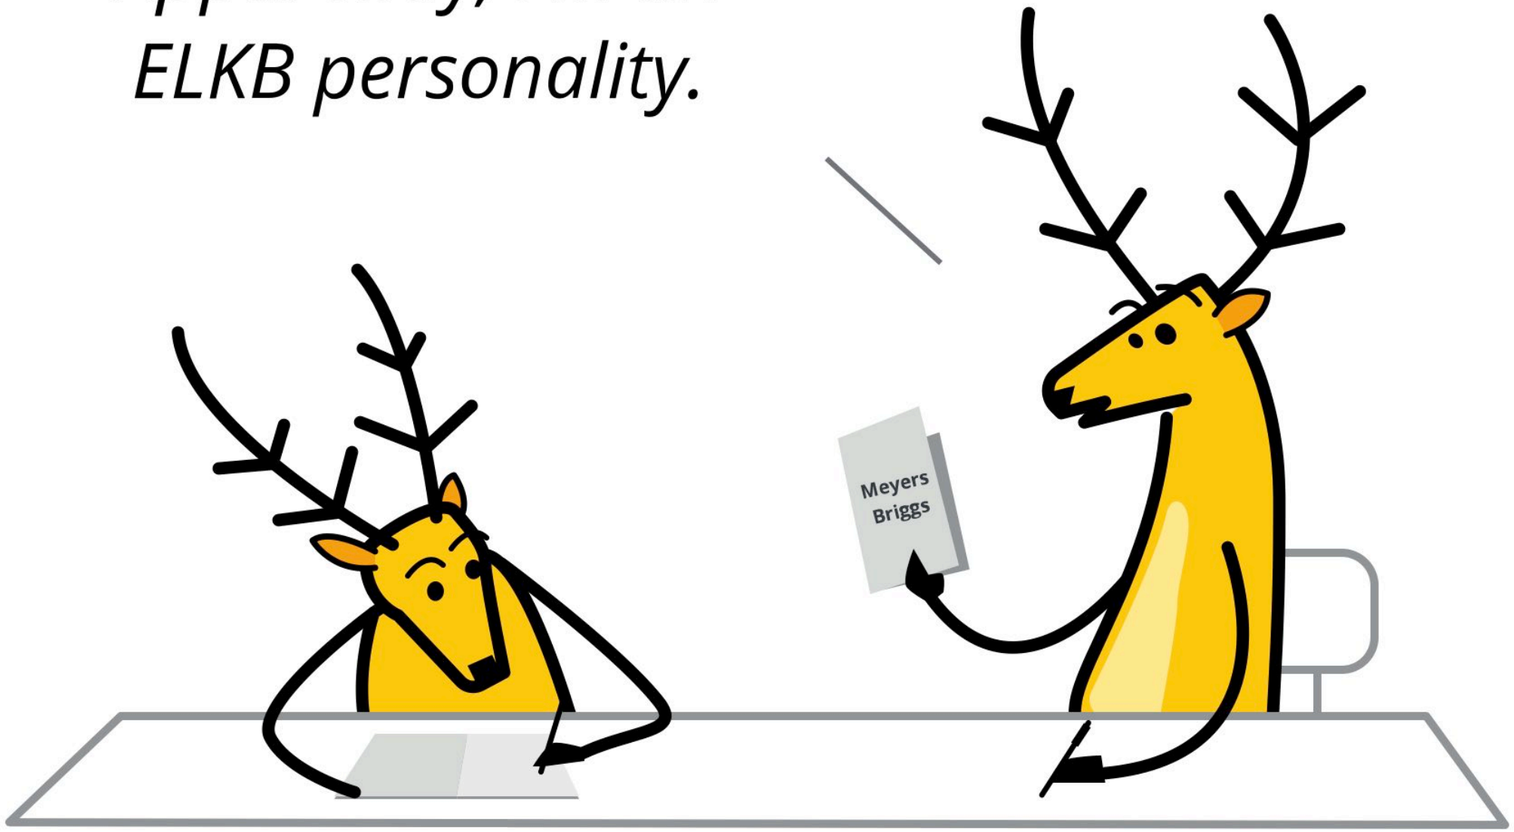

elastic

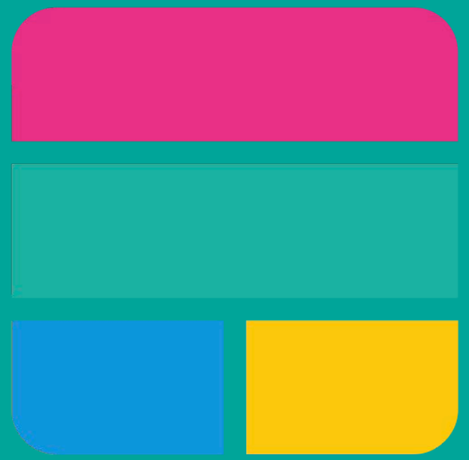

elastic stack

Evolution

Terms

Cluster, Node, Index, Shard, Document


20
30
40

Strictness*

5.0

* Demo

elastic


IT'S BAD

Parameters & configs

Bootstrap checks

Rolling Upgrades*

6.0

* Demo

elastic

Floodstage Watermark*

6.0

* Demo

elastic


Sawyer Rd
NEXT SIGNAL

Low 85%

High 90%

Floodstage 95%

Sequence Numbers*

6.0

* Demo

elastic

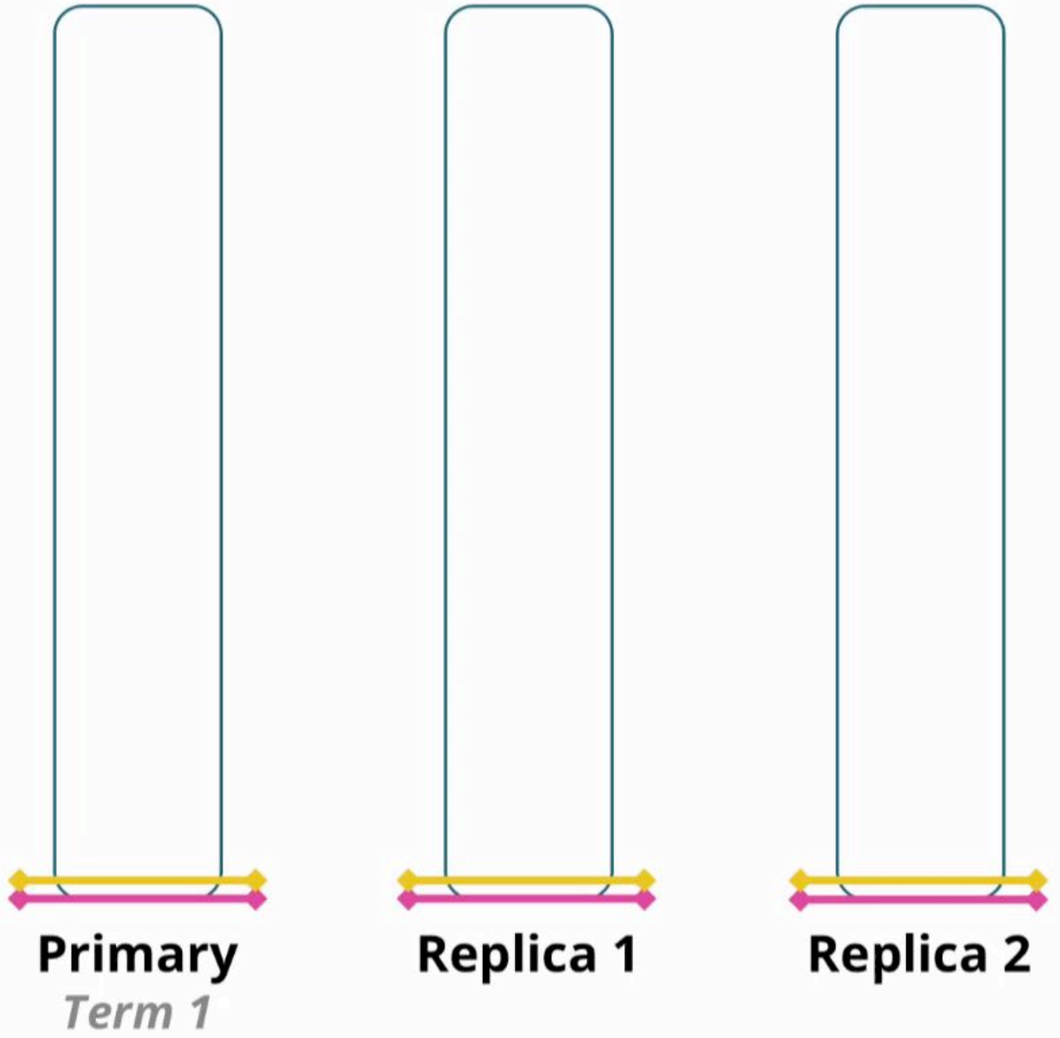
Primary
Term 1
Replica 1
Replica 2
Global checkpoint
Local checkpoint

63 bits ought to be
enough for anyone.

Cross datacenter replication

6.x

~~Types~~*

5.6 to 8.0

* Demo

elastic

Why

Data types

Sparsity

Scoring

How

5.6 opt-in single type

6.x single type

7.x type optional in API

8.x no more types

https://www.elastic.co/guide/en/elasticsearch/reference/6.1/removal-of-types.html#_schedule_for_removal_of_mapping_types

Automatic Queue Resizing

6.0

elastic

Reject and retry instead of long queues

`thread_pool.search.target_response_rate: 2s`

Serving 50 requests/s

Queue size: $2 * 50 = 100$

Adaptive Replica Selection

6.1

elastic

C3: Cutting Tail Latency in Cloud Data Stores via Adaptive Replica Selection

Lalith Suresh, ***Technische Universität Berlin;*** **Marco Canini,** ***Université catholique de Louvain;***
Stefan Schmid, ***Technische Universität Berlin and Telekom Innovation Labs;***
Anja Feldmann, ***Technische Universität Berlin***

https://www.usenix.org/conference/nsdi15/technical-sessions/presentation/suresh

Pick best shard

Exponentially Weighted Moving Average (EWMA)

Piggyback on requests

Test case	Throughput improvement %	50th % change	90th % change	99th % change
1 replica, no load	1.9%	-1.7%	0.5%	1.3%
1 replica, with load	115.8%	129.0%	-62.3%	-57.1%
4 replicas, no load	11.6%	-27.2%	-28.6%	-25.9%
4 replicas, with load	65.8%	-63.5%	-39.8%	-16.9%
1 replica, round robin, no load	6.8%	-7.2%	-16.6%	-28.0%

Shrink & Split*

5.0 6.1

* Demo

elastic

Shrink

Combine shards by a factor

Split

Split into a factor of `number_of_routing_shards`

Benchmarks

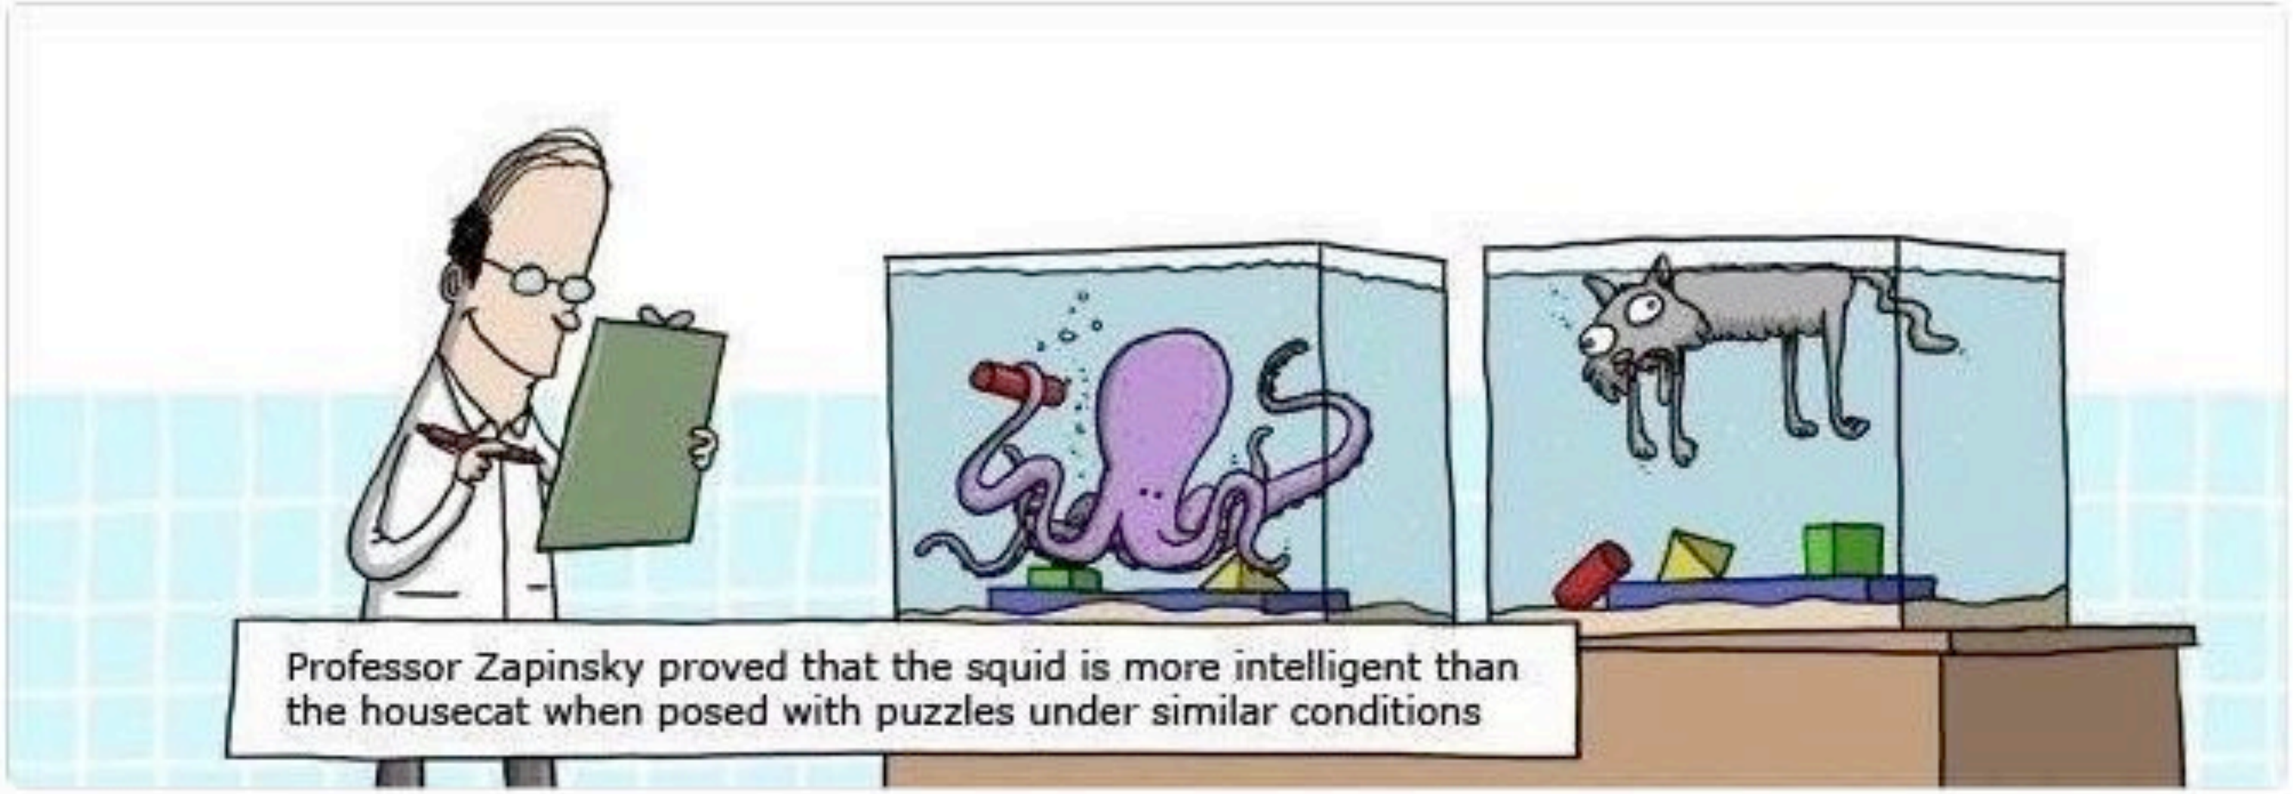
Professor Zapinsky proved that the squid is more intelligent than
the housecat when posed with puzzles under similar conditions

Rally

https://elasticsearch-benchmarks.elastic.co

Conclusion

elastic

Strictness

Rolling Upgrades

Floodstage Watermark

Sequence Numbers

~~Types~~

Automatic Queue Resizing

Adaptive Replica Selection

Shrink & Split

Questions?

Philipp Krenn @xeraa