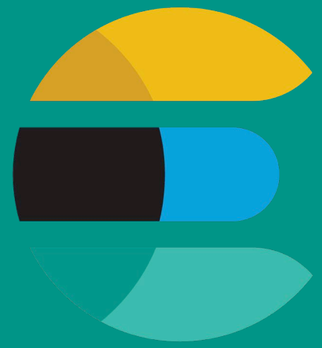

elasticsearch

(R)Evolution

Philipp Krenn

@xeraa

elastic

Revolution

elastic

16 systems in ranking, September 2017

Rank			DBMS	Database Model	Score		
Sep 2017	Aug 2017	Sep 2016			Sep 2017	Aug 2017	Sep 2016
1.	1.	1.	Elasticsearch	Search engine	120.00	+2.35	+23.52
2.	2.	2.	Solr	Search engine	69.91	+2.95	+2.95
3.	3.	3.	Splunk	Search engine	62.57	+1.11	+11.28
4.	4.	4.	MarkLogic	Multi-model	12.10	-0.40	+1.95
5.	5.	5.	Sphinx	Search engine	6.05	-0.10	-1.55
6.	6.	8.	Microsoft Azure Search	Search engine	3.74	+0.14	+1.74
7.	7.	6.	Google Search Appliance	Search engine	2.92	+0.02	-0.00
8.	8.	9.	Algolia	Search engine	2.57	-0.05	+1.26
9.	9.	7.	Amazon CloudSearch	Search engine	2.30	+0.01	-0.13
10.	10.	12.	CrateDB	Multi-model	0.78	-0.07	+0.51
11.	11.	10.	Xapian	Search engine	0.59	+0.02	-0.02
12.	12.	13.	SearchBlox	Search engine	0.27	+0.01	+0.14
13.	13.	15.	Exorbyte	Search engine	0.12	-0.03	+0.08
14.	16.	17.	DBSight	Search engine	0.02	-0.01	+0.02
15.	15.	16.	searchxml	Multi-model	0.01	-0.02	-0.01
16.	14.	11.	Indica	Search engine	0.00	-0.09	-0.34

https://db-engines.com/en/ranking/search+engine

16 systems in ranking, September 2017

Rank			DBMS	Database Model	Score		
Sep 2017	Aug 2017	Sep 2016			Sep 2017	Aug 2017	Sep 2016
1.	1.	1.	Elasticsearch	Search engine	120.00	+2.35	+23.52
2.	2.	2.	Solr	Search engine	69.91	+2.95	+2.95
3.	3.	3.	Splunk	Search engine	62.57	+1.11	+11.28
4.	4.	4.	MarkLogic	Multi-model	12.10	-0.40	+1.95
5.	5.	5.	Sphinx	Search engine	6.05	-0.10	-1.55
6.	6.	↑ 8.	Microsoft Azure Search	Search engine	3.46	+0.17	+1.74
7.	7.	↓ 6.	Google Search Appliance	Search engine	2.92	+0.02	-0.00
8.	8.	↑ 9.	Algolia	Search engine	2.57	-0.05	+1.26
9.	9.	↓ 7.	Amazon CloudSearch	Search engine	2.30	+0.01	-0.13
10.	10.	↑ 12.	CrateDB	Multi-model	0.78	-0.07	+0.51
11.	11.	↓ 10.	Xapian	Search engine	0.59	+0.02	-0.02
12.	12.	↑ 13.	SearchBlox	Search engine	0.27	+0.01	+0.14
13.	13.	↑ 15.	Exorbyte	Search engine	0.12	-0.03	+0.08
14.	↑ 16.	↑ 17.	DBSight	Search engine	0.02	-0.01	+0.02
15.	15.	↑ 16.	searchxml	Multi-model	0.01	-0.02	-0.01
16.	↓ 14.	↓ 11.	Indica	Search engine	0.00	-0.09	-0.34

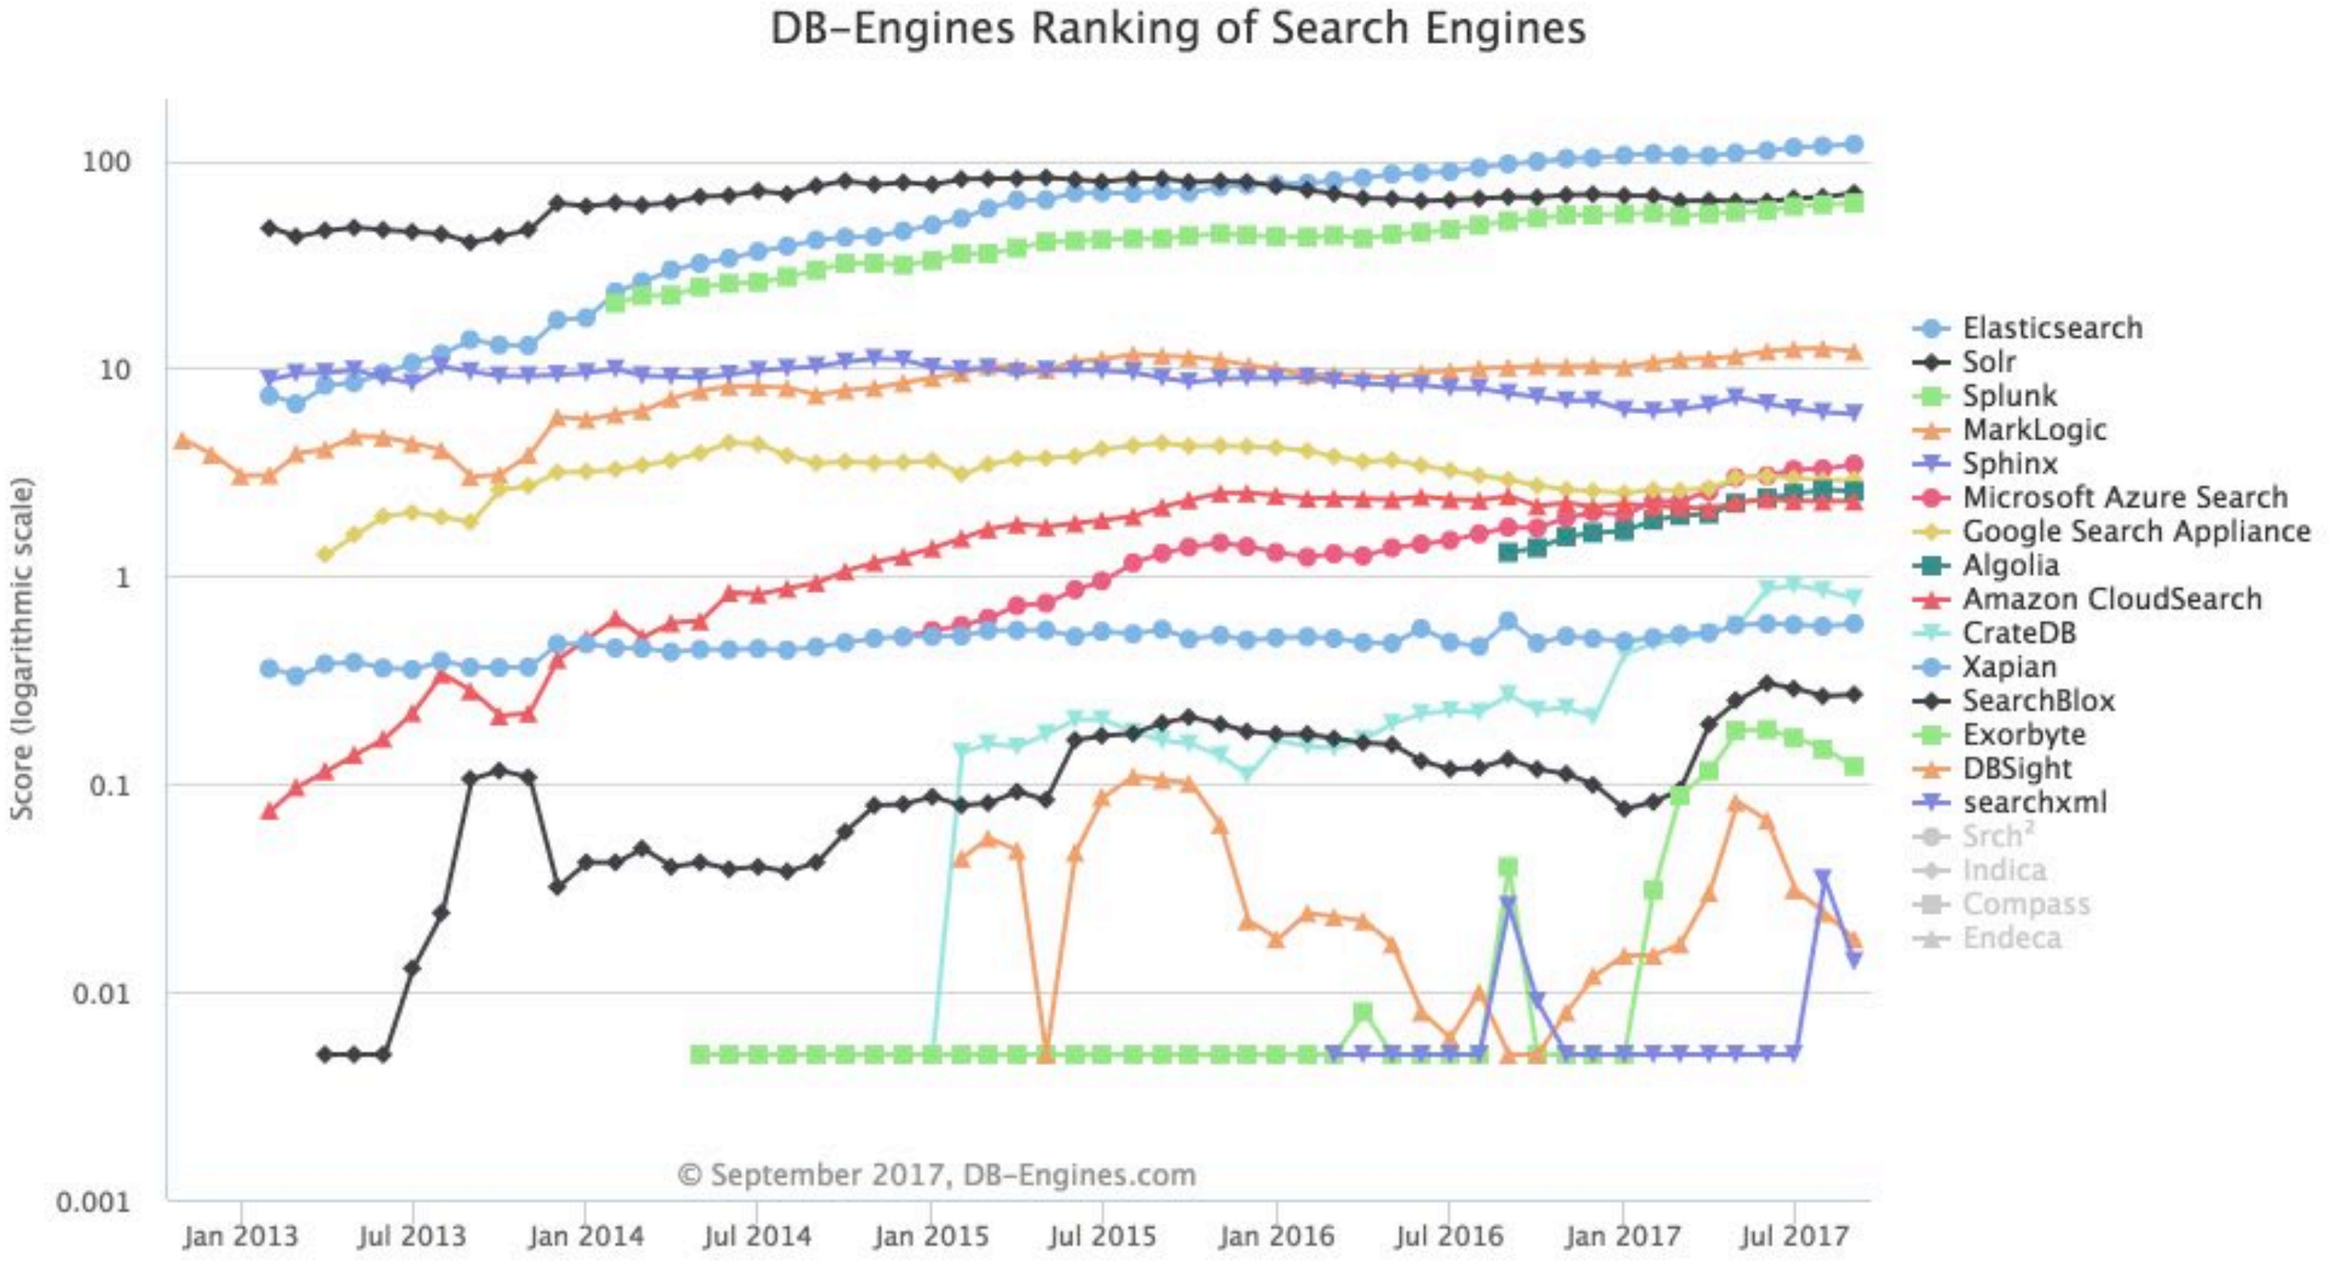
DB-Engines Ranking of Search Engines
Score (logarithmic scale)
100
10
1
0.1
0.01
0.001
Jan 2013
Jul 2013
Jan 2014
Jul 2014
Jan 2015
Jul 2015
Jan 2016
Jul 2016
Jan 2017
Jul 2017
© September 2017, DB-Engines.com
Elasticsearch
Solr
Splunk
MarkLogic
Sphinx
Microsoft Azure Search
Google Search Appliance
Algolia
Amazon CloudSearch
CrateDB
Xapian
SearchBlox
Exorbyte
DBSight
searchxml
Srch²
Indica
Compass
Endeca

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" : "MrmBv92dTQqh-z7pX2OkWw",
  "version" : {
    "number" : "5.6.1",
    "build_hash" : "667b497",
    "build_date" : "2017-09-14T19:22:05.189Z",
    "build_snapshot" : false,
    "lucene_version" : "6.6.1"
  },
  "tagline" : "You Know, for Search"
}
```

```
$ curl -u elastic:changeme http://localhost:9200
{
  "name" : "elasticsearch1",
  "cluster_name" : "docker-cluster",
  "cluster_uuid" : "MrmBv92dTQqh-z7pX2OkWw",
  "version" : {
    "number" : "5.6.1",
    "build_hash" : "667b497",
    "build_date" : "2017-09-14T19:22:05.189Z",
    "build_snapshot" : false,
    "lucene_version" : "6.6.1"
  },
  "tagline" : "You Know, for Search"
}
```

logstash

ELK Stack

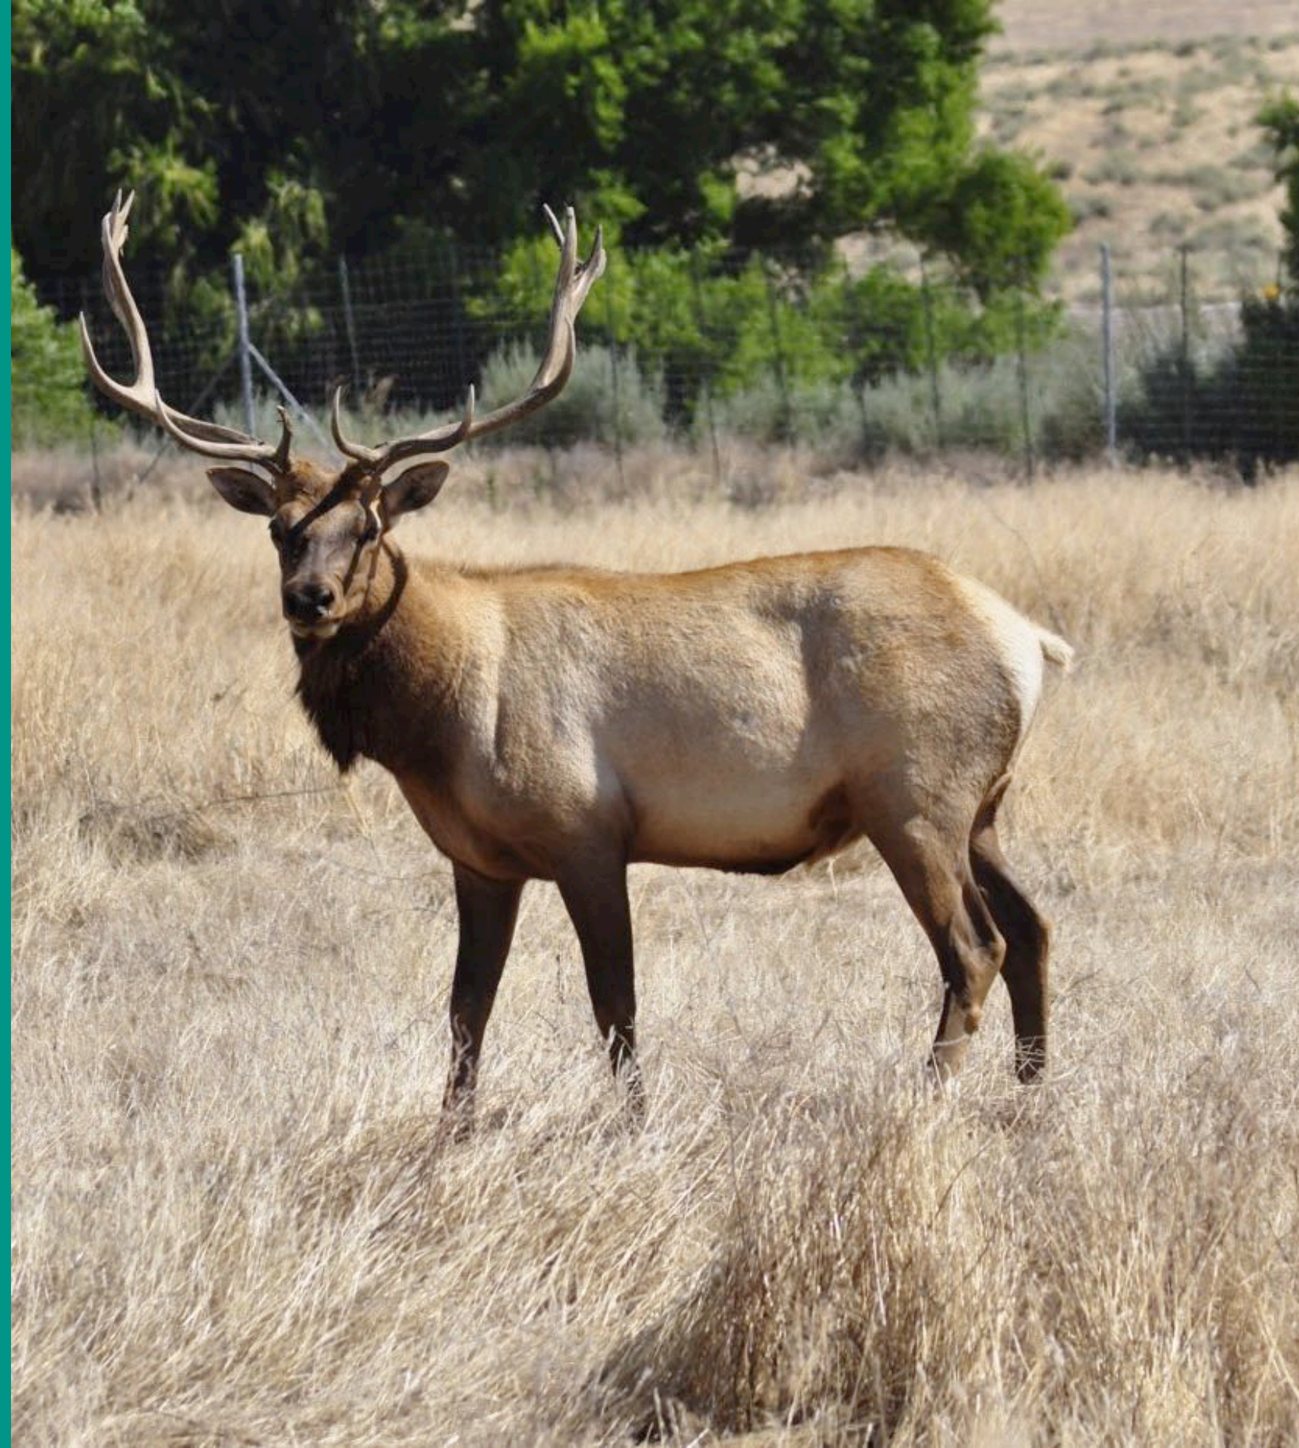

elastic

Goldman Sachs

Sprint

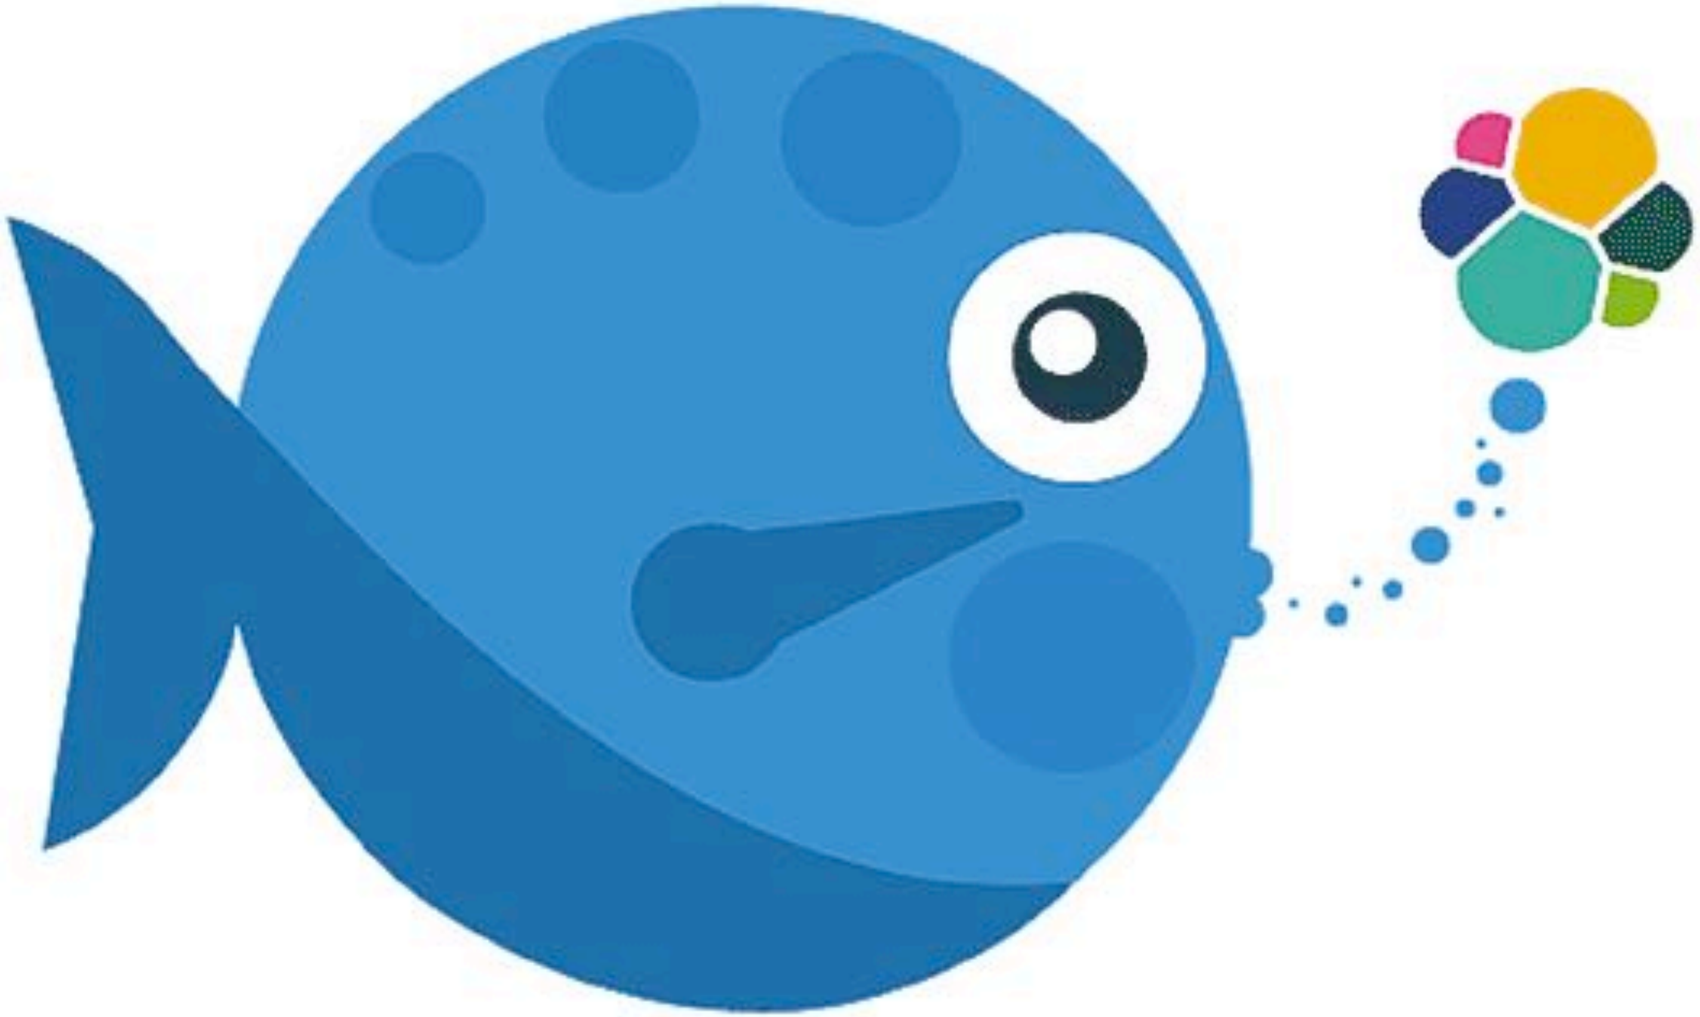

elastic

Elastic Stack

elastic


v5.0

Evolution

elastic

Strictness*

5.0

* Demo

elastic

Parameters & configs

Bootstrap checks

Shard Allocation API*

5.2

* Demo

elastic

Cross Cluster Search

5.3

elastic

Tribe node

https://github.com/elastic/elasticsearch/issues/4708

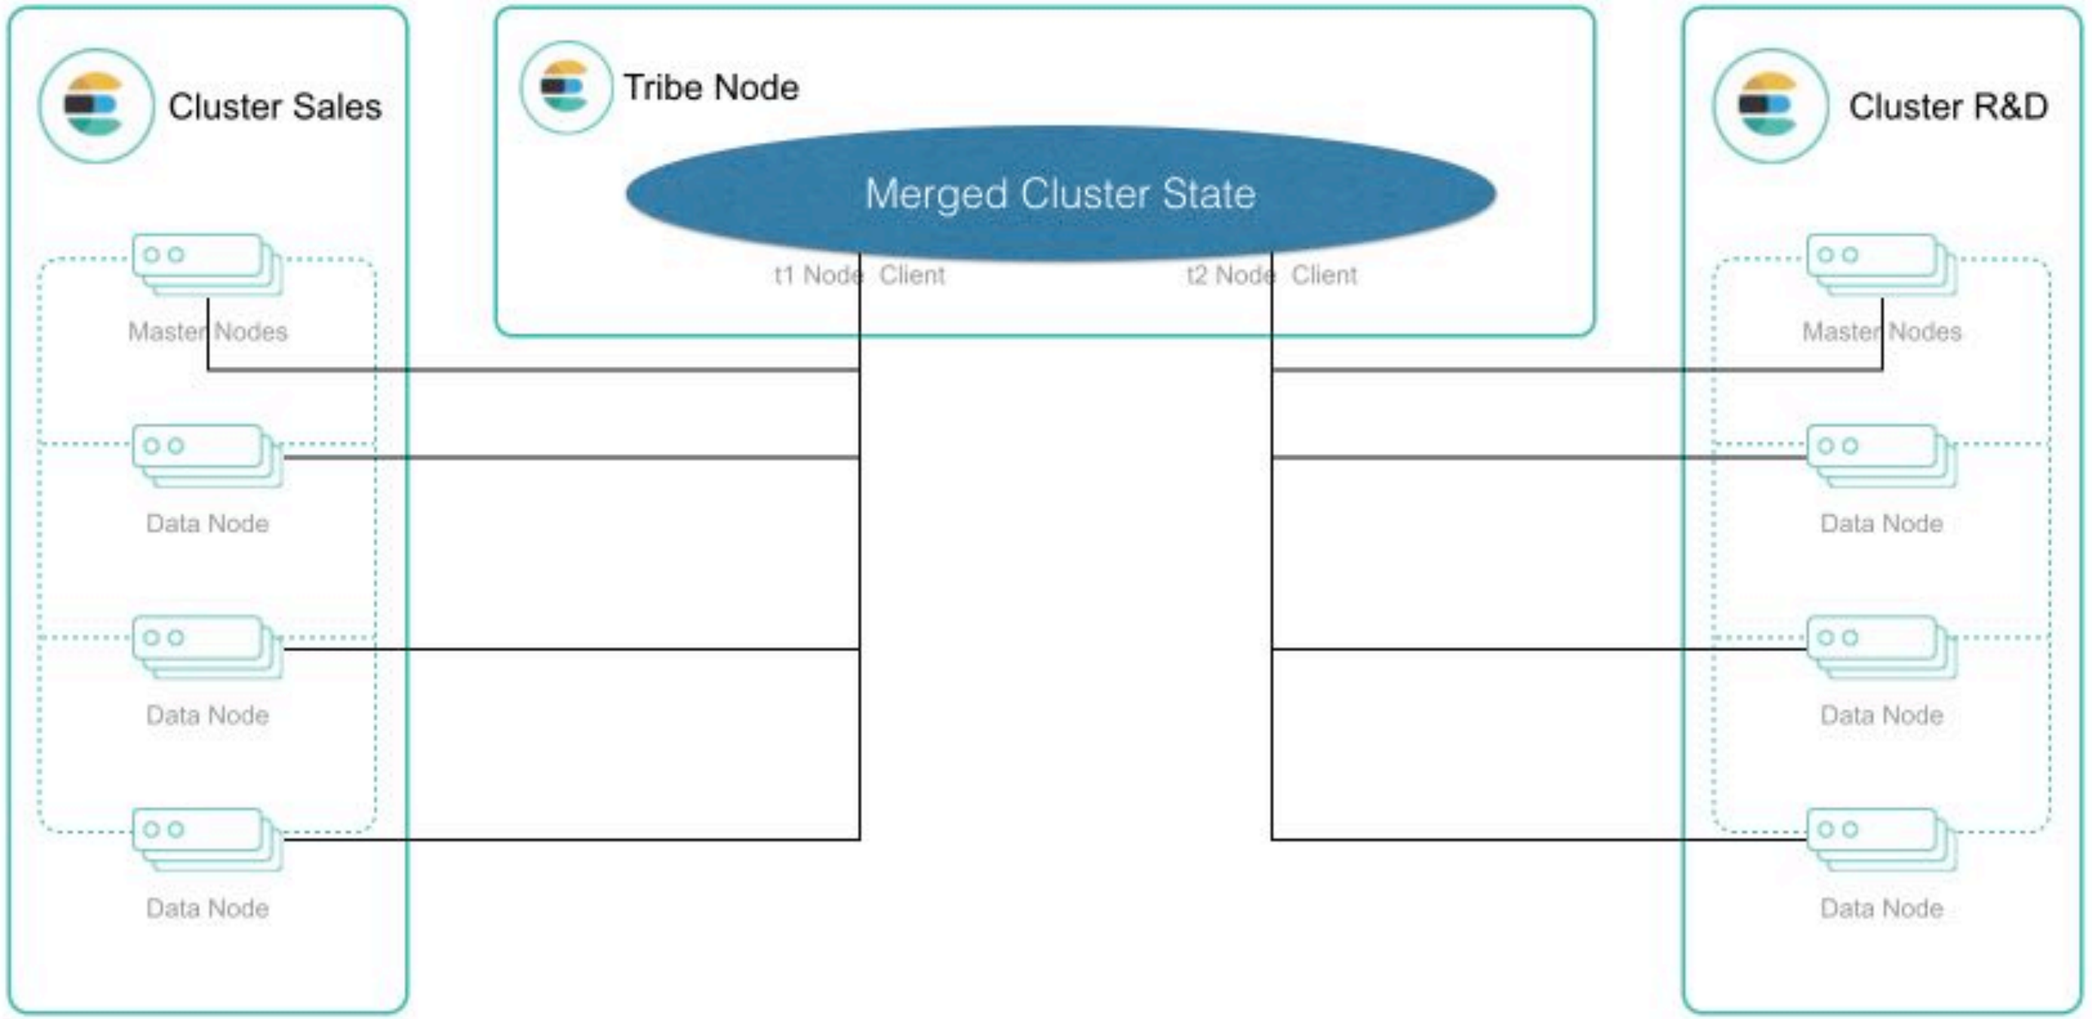
Cluster Sales
Tribe Node
Merged Cluster State
t1 Node
Client
t2 Node
Client
Cluster R&D
Master Nodes
Data Node
Data Node
Data Node
Master Nodes
Data Node
Data Node
Data Node

elastic

Cross cluster search

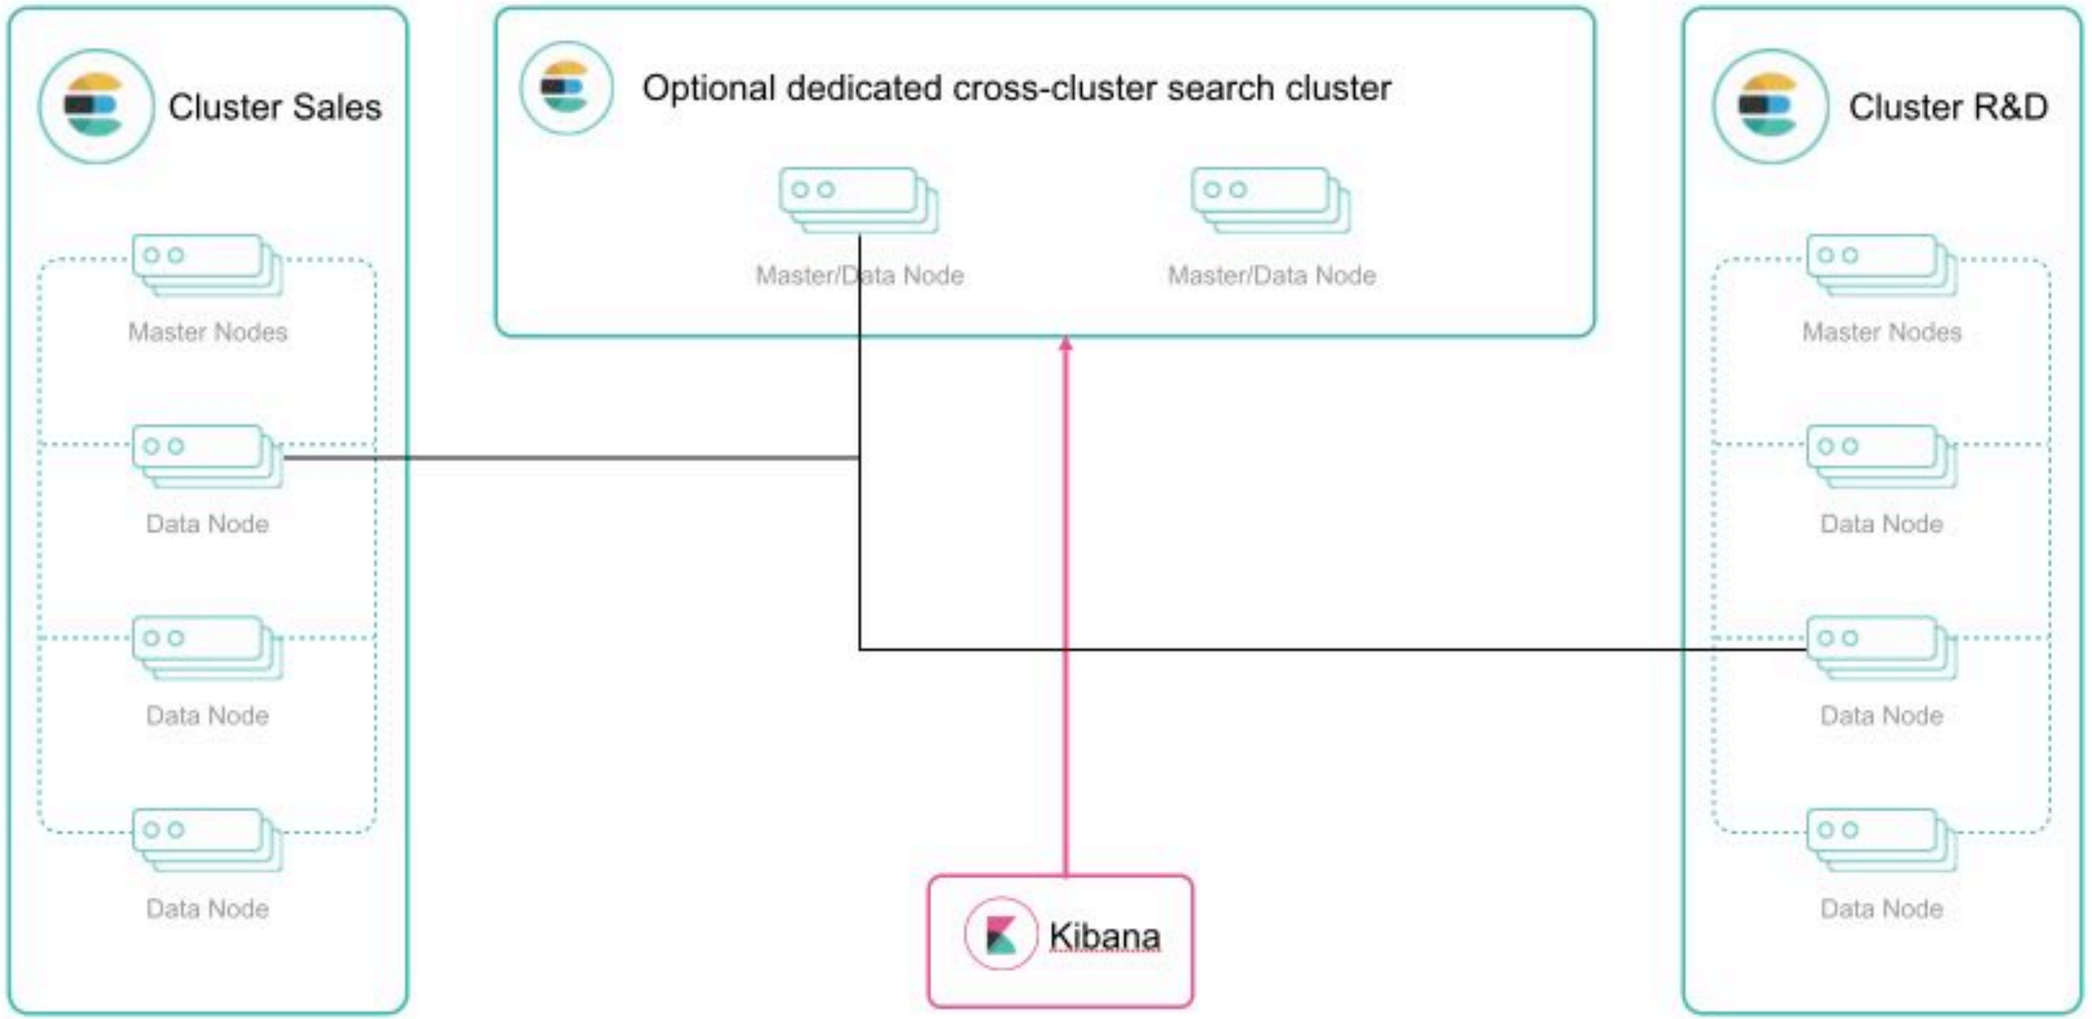
Cluster Sales
Master Nodes
Data Node
Data Node
Data Node
Optional dedicated cross-cluster search cluster
Master/Data Node
Master/Data Node
Kibana
Cluster R&D
Master Nodes
Data Node
Data Node
Data Node

Rolling Upgrades*

6.0

* Demo

elastic

Floodstage Watermark*

6.0

* Demo

elastic

Low 85%

High 90%

Floodstage 95%

Sequence Numbers*

6.0

* Demo

elastic

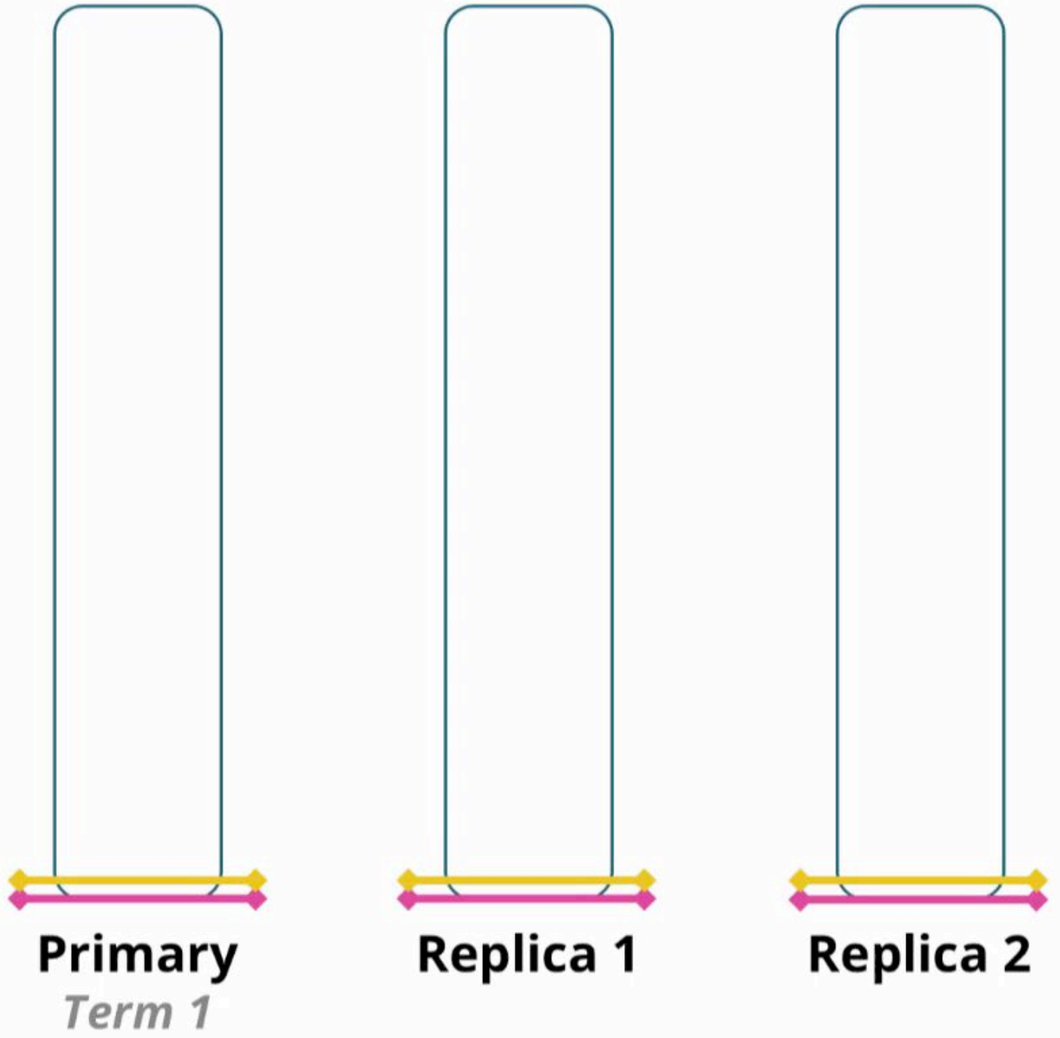
Primary
Term 1
Replica 1
Replica 2
Global checkpoint
Local checkpoint

Cross datacenter replication

6.x

elastic

~~Types~~*

5.6 to 7.0

* Demo

elastic


You lie!
memes.com

Why

Data types

Sparsity

Scoring

How

5.6 opt-in single type

6.x single type

7.x no type

Adaptive Replica Selection

6.1

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%

Off by default

`cluster.routing.use_adaptive_replica_selection`

Benchmarks

∞

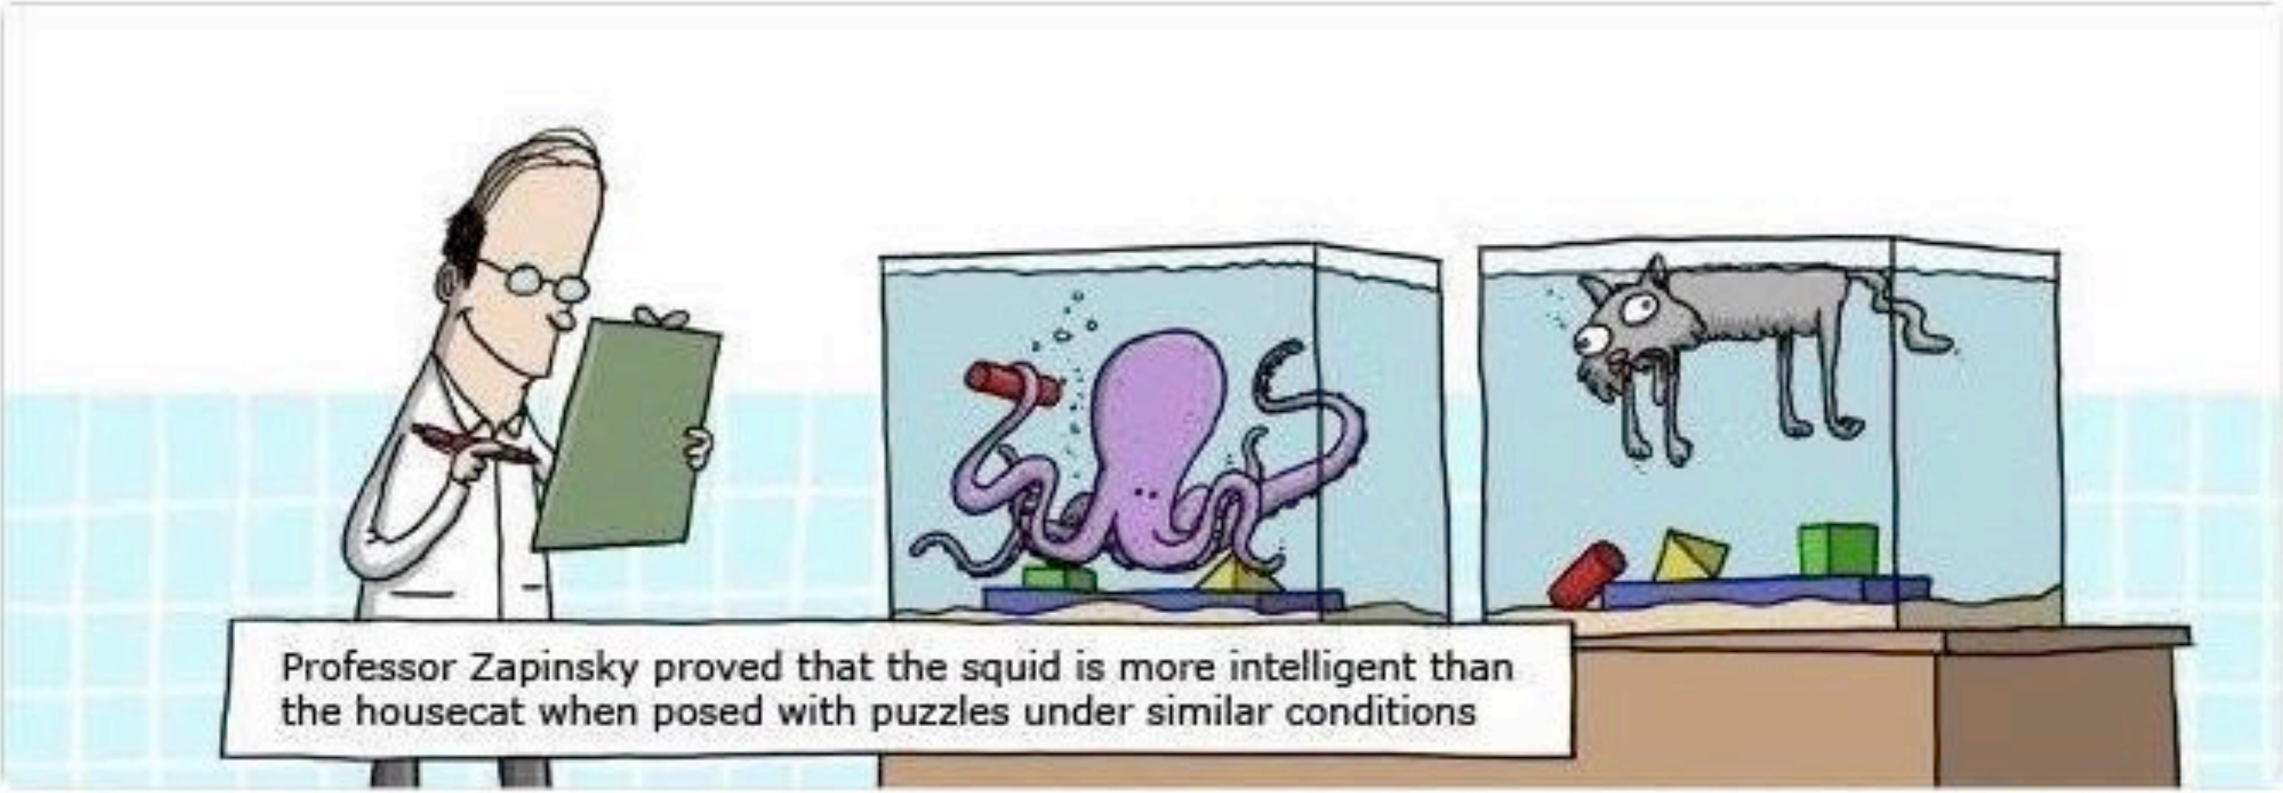
Professor Zapinsky proved that the squid is more intelligent than
the housecat when posed with puzzles under similar conditions

Rally

https://elasticsearch-benchmarks.elastic.co

Conclusion

elastic

Strictness

Shard Allocation API

Cross Cluster Search

Rolling Upgrades

Floodstage Watermark

Sequence Numbers

~~Types~~

Adaptive Replica Selection

Test 6.0 → GitHub / Discuss issue → swag

Questions?

Philipp Krenn

@xeraa