

elasticsearch

(R)Evolution

Philipp Krenn

@xeraa





elastic

Developer 🥑

Questions: <https://sli.do/xeraa>

Answers: <https://twitter.com/xeraa>



Revolution

343 systems in ranking, June 2018

Rank			DBMS	Database Model	Score		
Jun 2018	May 2018	Jun 2017			Jun 2018	May 2018	Jun 2017
1.	1.	1.	Oracle +	Relational DBMS	1311.25	+20.84	-40.51
2.	2.	2.	MySQL +	Relational DBMS	1233.69	+10.35	-111.62
3.	3.	3.	Microsoft SQL Server +	Relational DBMS	1087.73	+1.89	-111.23
4.	4.	4.	PostgreSQL +	Relational DBMS	410.67	+9.77	+42.13
5.	5.	5.	MongoDB +	Document store	343.79	+1.67	+8.79
6.	5.	5.	DB2 +	Relational DBMS	85.51	+0.00	-1.86
7.	7.	↑ 9.	Redis +	Key-value store	136.30	+0.95	+17.42
8.	↑ 9.	↑ 11.	Elasticsearch +	Search engine	131.04	+0.60	+19.48
9.	↓ 8.	↓ 7.	Microsoft Access	Relational DBMS	130.99	-2.12	+4.44
10.	10.	↓ 8.	Cassandra +	Wide column store	119.21	+1.38	-4.91
11.	11.	↓ 10.	SQLite +	Relational DBMS	114.26	-1.19	-2.44
12.	12.	12.	Teradata	Relational DBMS	75.77	+1.36	-1.55
13.	↑ 14.	↑ 18.	MariaDB +	Relational DBMS	65.85	+0.85	+12.95
14.	↓ 13.	↑ 16.	Splunk	Search engine	65.78	+0.68	+8.26
15.	15.	↓ 14.	Solr	Search engine	62.06	+0.55	-1.55

<https://db-engines.com/en/ranking>

343 systems in ranking, June 2018

Rank			DBMS	Database Model	Score		
Jun 2018	May 2018	Jun 2017			Jun 2018	May 2018	Jun 2017
1.	1.	1.	Oracle	Relational DBMS	1311.25	+20.84	-40.51
2.	2.	2.	MySQL	Relational DBMS	1233.69	+10.35	-111.62
3.	3.	3.	Microsoft SQL Server	Relational DBMS	1087.73	+1.89	-111.23
4.	4.	4.	PostgreSQL	Relational DBMS	410.67	+9.77	+42.13
5.	5.	5.	MongoDB	Document store	343.79	+1.67	+8.79
6.	6.	6.	DB2	Relational DBMS	185.64	+0.03	-1.86
7.	7.	9.	Redis	Key-value store	136.30	+0.95	+17.42
8.	9.	11.	Elasticsearch	Search engine	131.04	+0.60	+19.48
9.	8.	7.	Microsoft Access	Relational DBMS	130.99	-2.12	+4.44
10.	10.	8.	Cassandra	Wide column store	119.21	+1.38	-4.91
11.	11.	10.	SQLite	Relational DBMS	114.26	-1.19	-2.44
12.	12.	12.	Teradata	Relational DBMS	75.77	+1.36	-1.55
13.	14.	18.	MariaDB	Relational DBMS	65.85	+0.85	+12.95
14.	13.	16.	Splunk	Search engine	65.78	+0.68	+8.26
15.	15.	14.	Solr	Search engine	62.06	+0.55	-1.55

↑ 9.

Redis 

11.

Elasticsearch 

↓ 7.

Microsoft Access

↓ 8.

Cassandra 

10

SQLite 

Evolution

Strictness*

5.0

* Demo

A black and white close-up portrait of a man with short hair, wearing dark sunglasses and a black leather jacket with silver studs. He has a slight, enigmatic smile. The background is a plain, light-colored wall.

IT'S BAD

Bootstrap Checks

<https://www.elastic.co/guide/en/elasticsearch/reference/current/bootstrap-checks.html>

Bootstrap Checks

Heap size check

File descriptor check

Memory lock check

Maximum number of threads check

Maximum size virtual memory
check

Max file size check

Maximum map count check

Client JVM check

Use serial collector check

System call filter check

OnError and OnOutOfMemoryError
checks

Early-access check

G1GC check

All permission check

Parameters & Configs

Rolling Upgrades*

6.0

* Demo



Floodstage Watermark*

6.0

* Demo



Low 85%

High 90%

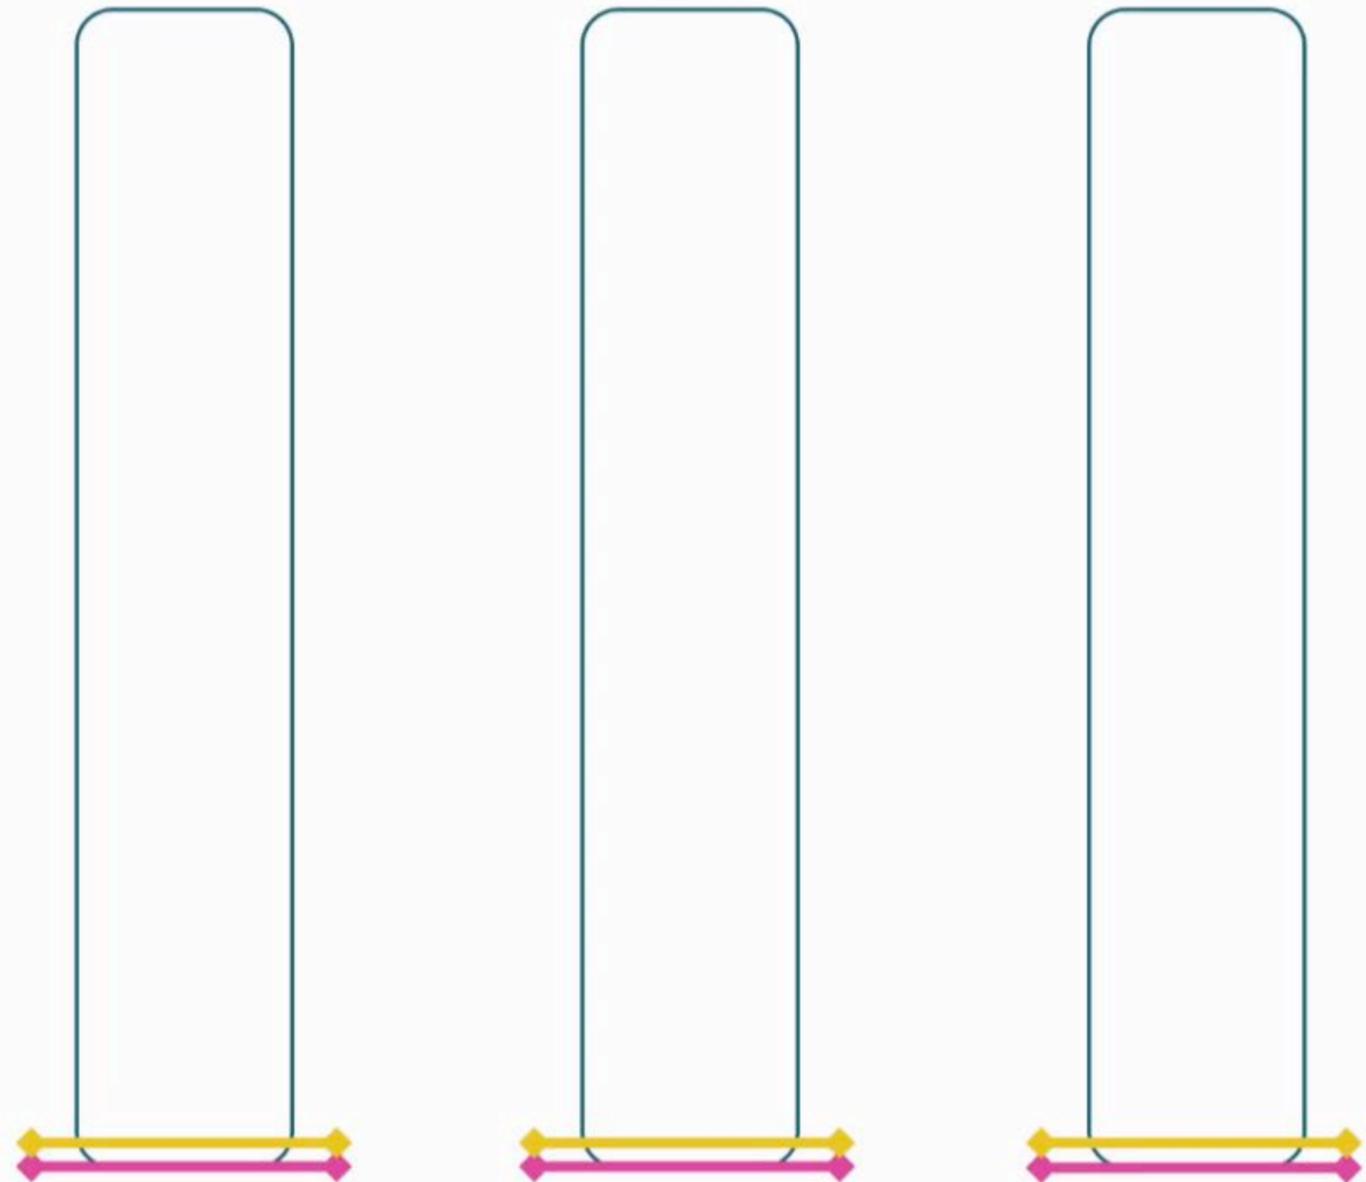
Floodstage 95%

Sequence Numbers*

6.0

* Demo





Primary
Term 1

Replica 1

Replica 2

Global checkpoint
Local checkpoint



63 bits ought to be
enough for anyone.

Tradeoff

```
index.translog.retention.size: 512MB
```

```
index.translog.retention.age: 12h
```

Cross Datacenter Replication

6.x or 7.x

~~Types~~*

5.6 to 8.0

* Demo



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/current/removal-of-types.html#_schedule_for_removal_of_mapping_types

Automatic Queue Resizing

6.0

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 (enabled by default in 7.0)

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 responses to coordinating node

Shrink & Split*

5.0

6.1

* Demo



Shrink

Combine shards by a factor



Split

Split into a factor of `number_of_routing_shards`

Not required in 7.0+

Open Code

6.3





 Public Repo

elasticsearch

< code >

License: **Apache 2.0**

 Private Repo

x-pack-elasticsearch

< code >

License: **Elastic EULA**

 Public Repo

elasticsearch

< **oss code** >
Apache 2.0

< **x-pack code** >
Elastic EULA

Community License: **Apache 2.0 + EULA**

Graph Reporting
Machine Learning
Grok Debugger Canvas
SQL Monitoring CSV

X-PACK

PAID
FREE

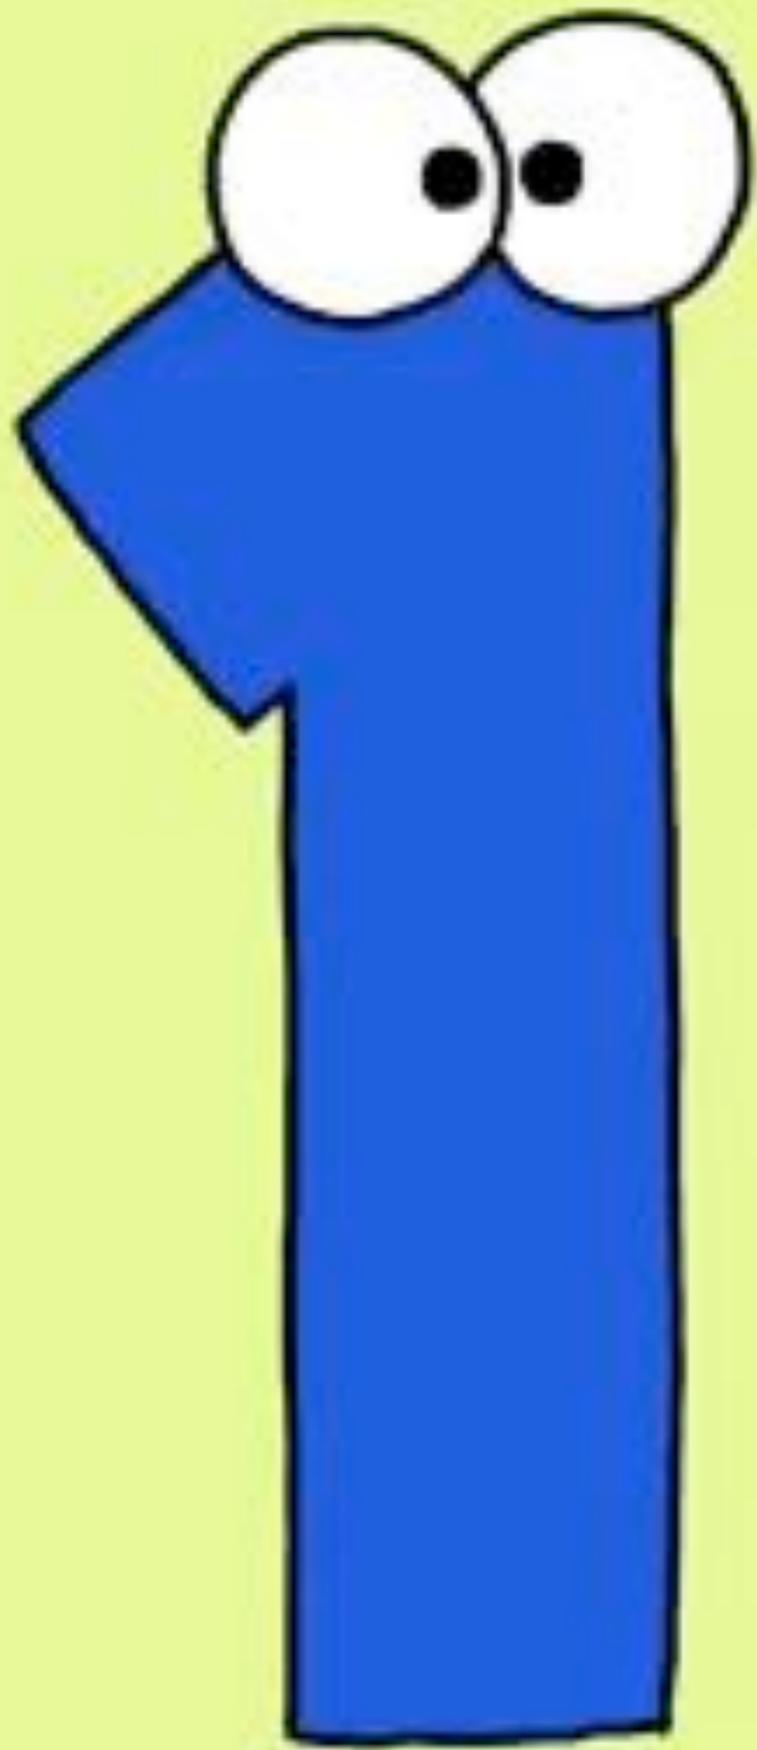
OPEN

Enrichment Ingestions Distributed
Search Engine **Data Collection**
Agents Dashboard Visualizations
Visualizations Ingestions Enrichment
Edge Collection Distributed
Search Engine Data Collection

OSS

Shards

7.0



Default: 1 Shard per Index

Oversharding

Simon Says

Use a single shard until it
blows up



Beats: Rollover

Planning phase: size based, depends on ILM

JDK 11?

7.0

Can users upgrade?
Version for Java clients?

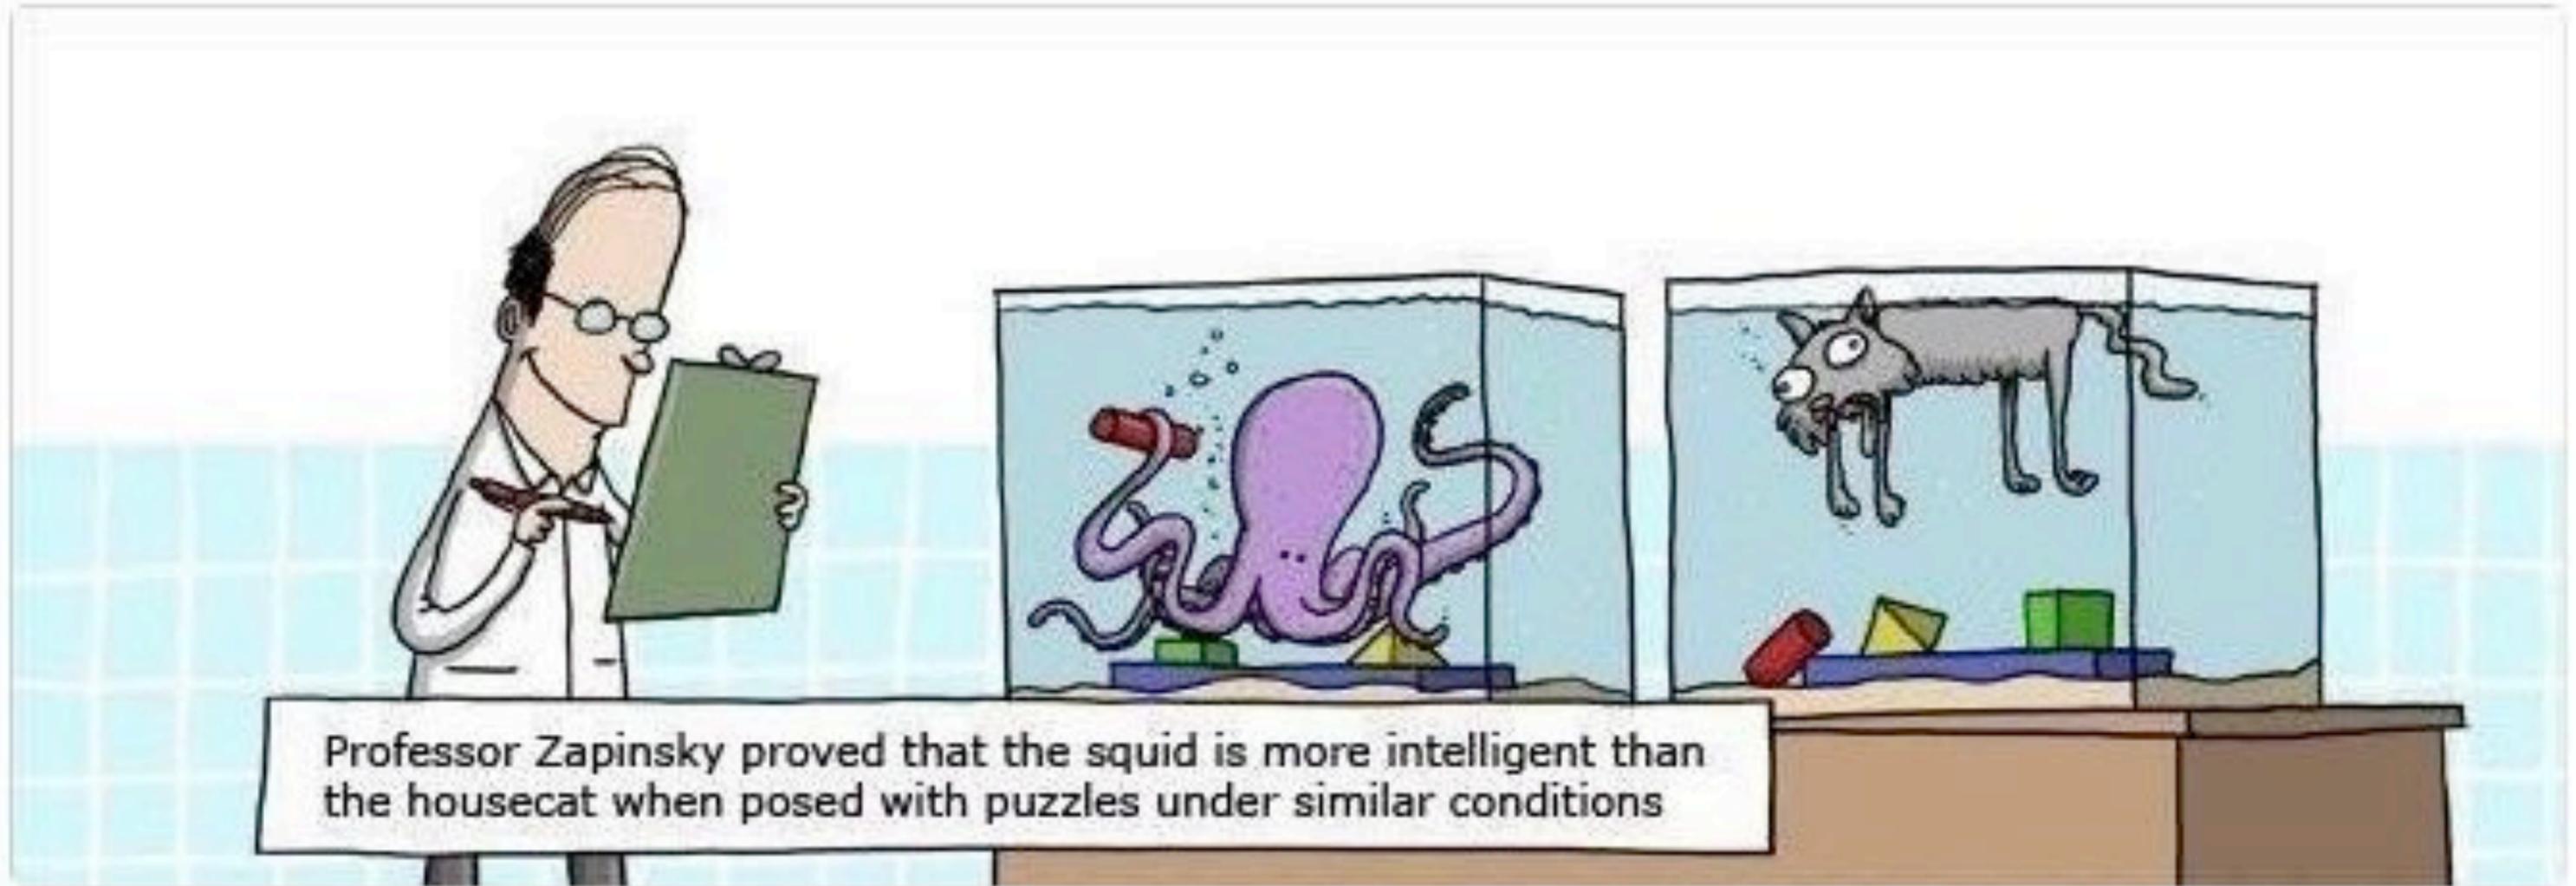
Low-level REST: independent (JDK 7)

High-level REST: JDK 11

Transport: deprecated

Benchmarks







Rally

<https://elasticsearch-benchmarks.elastic.co>

Conclusion

Don't be Shay to
change



Strictness

Rolling Upgrades

Floodstage Watermark

Sequence Numbers

~~Types~~

Automatic Queue Resizing & Adaptive Replica Selection

Shrink & Split

Open Code

Default Shard Number

JDK 11

meetup

<https://www.meetup.com/Elasticsearch-Berlin/>

Thursday, June 14, 19:00, Springer Nature



Questions?

Philipp Krenn

@xeraa