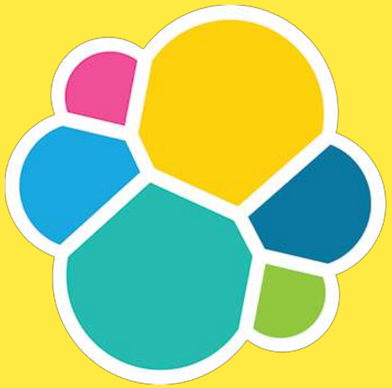

elastic

STACK 7

HIGHLIGHTS

PHILIPP KRENN

@XERAA

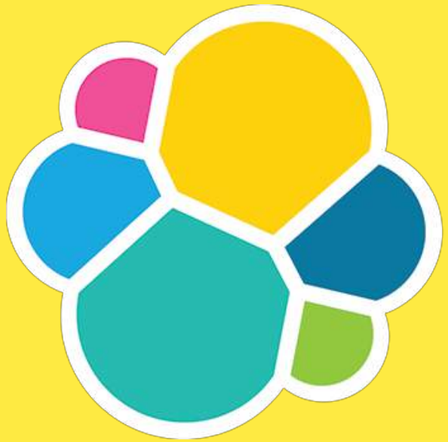

elastic

DEVELOPER

ZEN2

CLUSTER COORDINATION

CLUSTER STATE

MASTER AND QUORUM

`discovery.zen.minimum_master_nodes`

RETHOUGHT AND REBUILT IN 7.0

~~minimum_master_nodes~~

FASTER (SUB-SECOND)

SAFER GROWING AND SHRINKING

BETTER LOGGING AND DEBUGGING

CLUSTER BOOTSTRAPPING

SINGLE HOST: AUTOMATIC

CLUSTER BOOTSTRAPPING

`cluster.initial_master_nodes`

SAFE TO REMOVE NODES SIMPLY BY STOPPING THEM (LESS THAN HALF AT ONCE)

CLUSTER BOOTSTRAPPING

6.7 ROLLING, 6.X RESTART

`minimum_master_nodes` USED FOR BOOTSTRAPPING

RENAMES (OLD DEPRECATED)

OLD NAME ▶️	NEW NAME
`discovery.zen.ping.unicast.hosts` ▶️	`discovery.seed_hosts`
`discovery.zen.hosts_provider` ▶️	`discovery.seed_providers`
`discovery.zen.no_master_block` ▶️	`cluster.no_master_block`

TYPES

... OR NO TYPES

WHY

DATA TYPES

SPARSITY

SCORING

elastic

5.6 OPT-IN SINGLE TYPE
6.x SINGLE TYPE
7.x TYPE OPTIONAL IN API
8.x TYPE REMOVED FROM API

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

IN 7.0, _doc REPRESENTS THE ENDPOINT NAME INSTEAD OF THE DOCUMENT TYPE. THE _doc COMPONENT IS A PERMANENT PART OF THE PATH FOR THE DOCUMENT INDEX, GET, AND DELETE APIS GOING FORWARD, AND WILL NOT BE REMOVED IN 8.0.

OPTIMIZATIONS

`index.search.idle.after: 30s`

IFF DEFAULT `index.refresh_interval`

FASTER RETRIEVAL OF TOP HITS

MAXSCORE ALGORITHM

TOP 10 MATCHES

MAX SCORE "ELASTICSEARCH": 3.0

MAX SCORE "KIBANA": 5.0

10TH BEST SCORE > 3.0: SKIP "ELASTICSEARCH"

BLOCK-MAX WAND

ACTUAL IMPLEMENTATION

IMPLICATIONS

NO NEGATIVE SCORES

TOTAL HIT COUNT

```
"hits" : {
  "total" : {
    "value" : 10000,
    "relation" : "gte"
  },
```

track_total_hits

DEFAULT 10,000

TRADEOFF PRECISION VS SPEED

NEW DEFAULTS

NUMBER OF PRIMARY SHARDS?

1 SHARD PER INDEX

OVERSHARDING

SIMON SAYS

USE A SINGLE SHARD UNTIL IT BLOWS UP

ADAPTIVE REPLICA SELECTION

usenix
THE ADVANCED
COMPUTING SYSTEMS
ASSOCIATION

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

LESS HEAP

FROZEN INDICES

RATIO HEAP : STORAGE

INDEX > FROZEN INDEX > CLOSED INDEX

READ-ONLY

THROTTLED THREAD POOL

1 PARALLEL SEARCH / NODE

100 IN QUEUE

REAL MEMORY CIRCUIT BREAKER

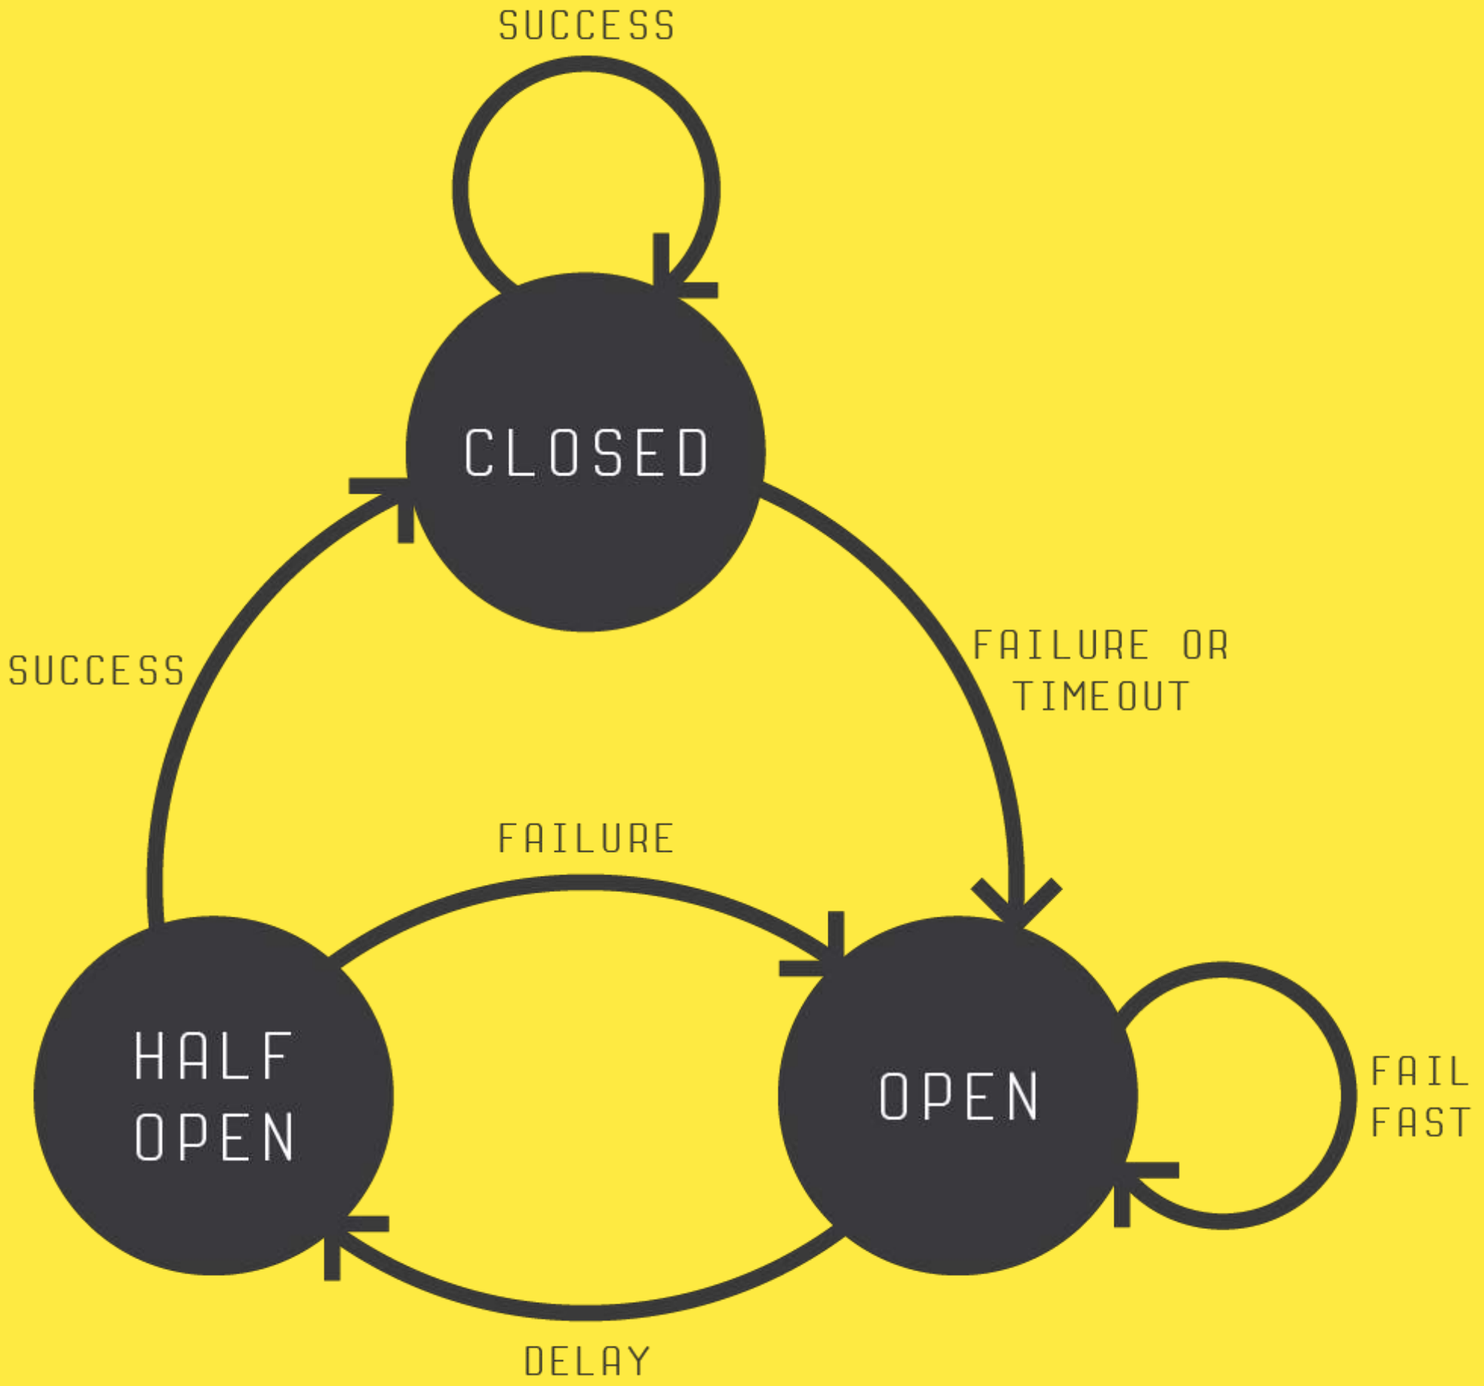
SUCCESS
CLOSED
SUCCESS
FAILURE OR
TIMEOUT
FAILURE
HALF
OPEN
OPEN
FAIL
FAST
DELAY

CIRCUIT BREAKER

GOAL: AVOID `OutOfMemoryError`

UPFRONT ESTIMATION

REAL MEMORY CIRCUIT BREAKER

KNOW ACTUAL USAGE WHEN RESERVING

`MemoryUsage` **FROM** `MemoryPoolMXBean`

elastic

```
{
  'error': {
    'type': 'circuit_breaking_exception',
    'reason': '[parent] Data too large, data for [<http_request>]
               would be [123848638/118.1mb],
               which is larger than the limit of [123273216/117.5mb],
               real usage: [120182112/114.6mb],
               new bytes reserved: [3666526/3.4mb]',
    'bytes_wanted': 123848638,
    'bytes_limit': 123273216,
    'durability': 'TRANSIENT'
  },
  'status': 429
}
```

BETTER RESILIENCY

256MB OF HEAP OR TOO MANY BUCKETS

WATCH OUT FOR BACK OFF / RETRY AND PARTIAL RESULTS

KIBANA

elastic

ELASTIC UI FRAMEWORK

https://github.com/elastic/eui

WARNING WHILE OPEN SOURCE, THE INTENDED CONSUMERS OF THIS REPOSITORY ARE ELASTIC PRODUCTS.

DE-ANGULARIZATION WIP

elastic

CONCLUSION

BREAKING CHANGES

https://www.elastic.co/guide/en/elasticsearch/reference/current/breaking-changes-7.0.html

RECAP

ZEN2, TYPES, OPTIMIZATIONS, NEW DEFAULTS, LESS HEAP, KIBANA

QUESTIONS?

PHILIPP KRENN

@XERAA

elastic