Scale Your Metrics with Elasticsearch

Philipp Krenn @xeraa

elastic

elasticsearch.

```
$ curl http://localhost:9200
{
  "name" : "elasticsearch-hot",
  "cluster_name" : "metrics-cluster",
  "cluster_uuid" : "O6nHPLLgTrmZEpYli6JW5w",
  "version" : {
    "number" : "6.5.0",
    "build_flavor" : "default",
    "build_type" : "tar",
    "build_hash" : "c53b7d3",
    "build_date" : "2018-11-08T21:28:50.577384Z",
    "build_snapshot" : false,
    "lucene_version" : "7.5.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "tagline" : "You Know, for Search"
}
```

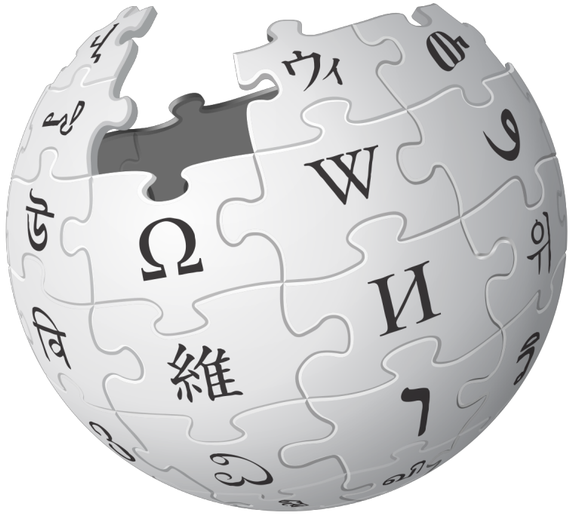

logstash

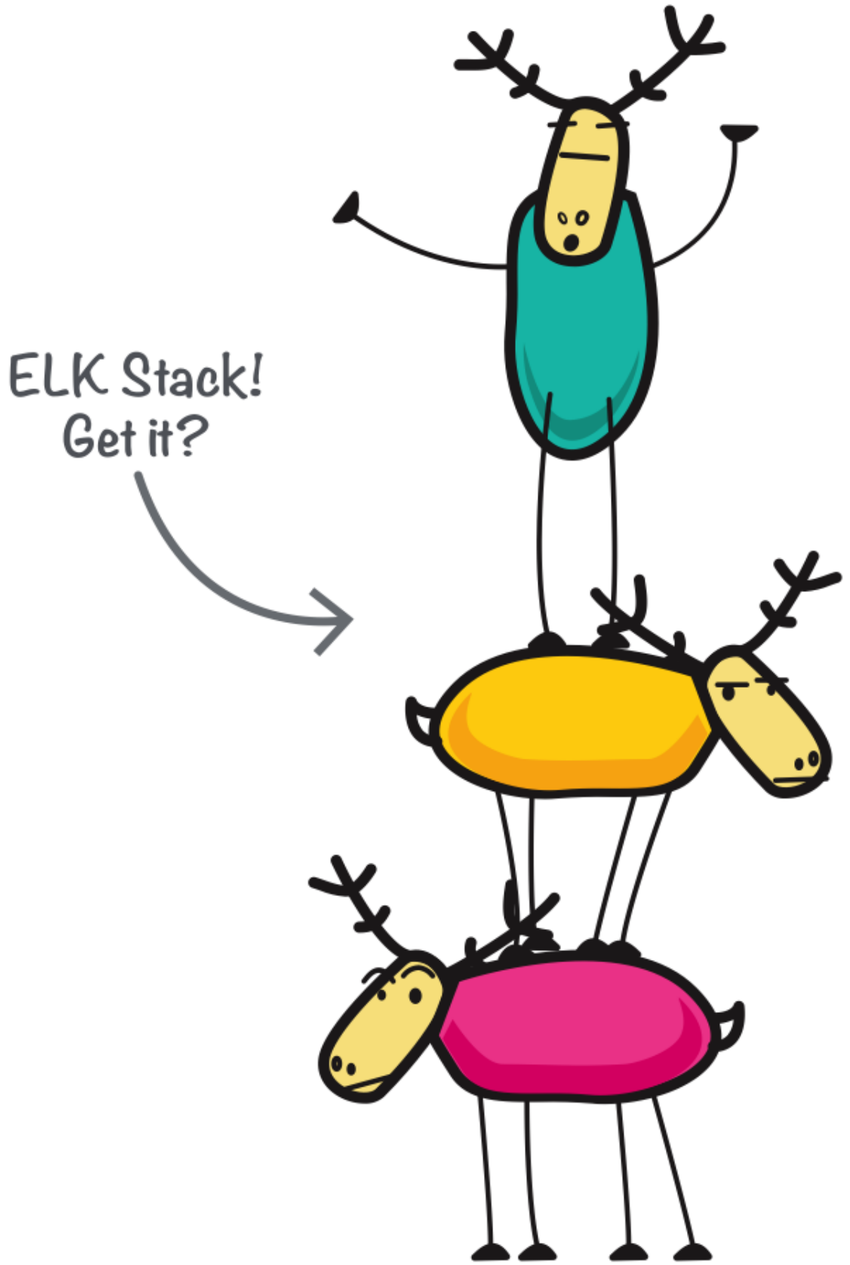

E Elasticsearch

L Logstash

K Kibana


lyft


slack


fitbit

Apparently, I'm an ELKB personality.

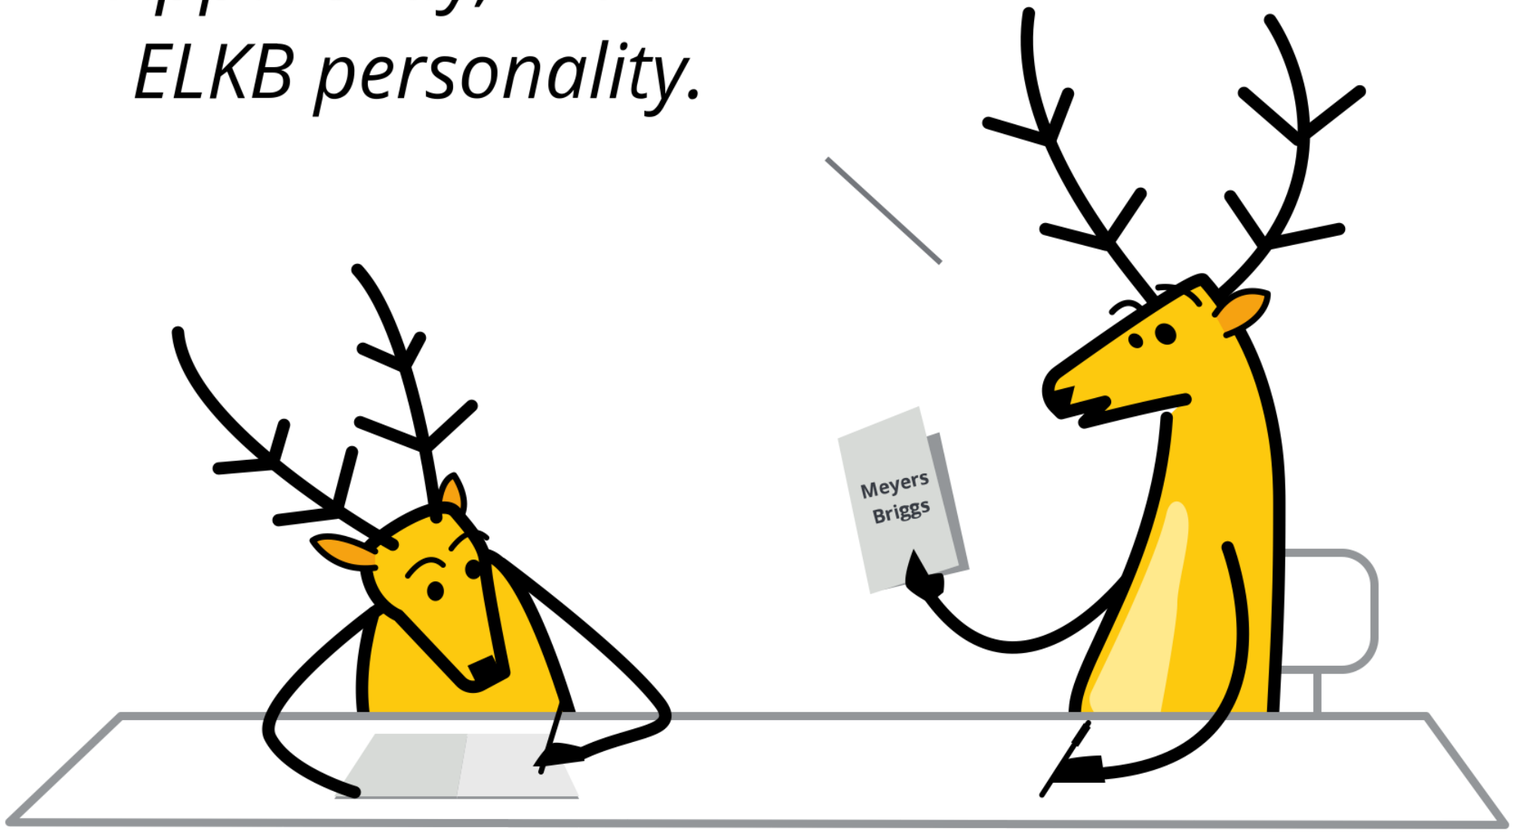

elastic

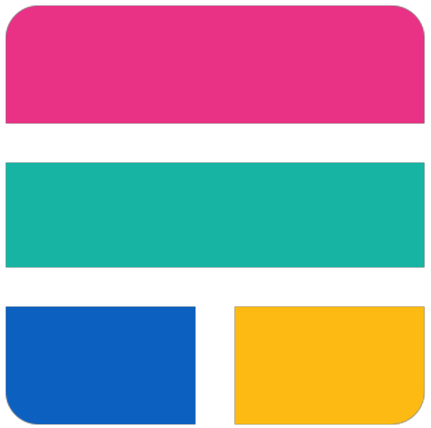

elastic stack


BLIZZARD
ENTERTAINMENT


CAR
2GO


ebay

I'm not going to use a search engine for metrics.

— Too often

elastic

Developer 🥑

Agenda

Building Blocks
Architecture
Demo

Building Blocks

Only accept features that scale.

— **https://github.com/elastic/engineering/blob/master/development_constitution.md**

Horizontal Scaling

Shards
Replication
Writes & Reads

elastic

Cluster, Node, Index, Shard

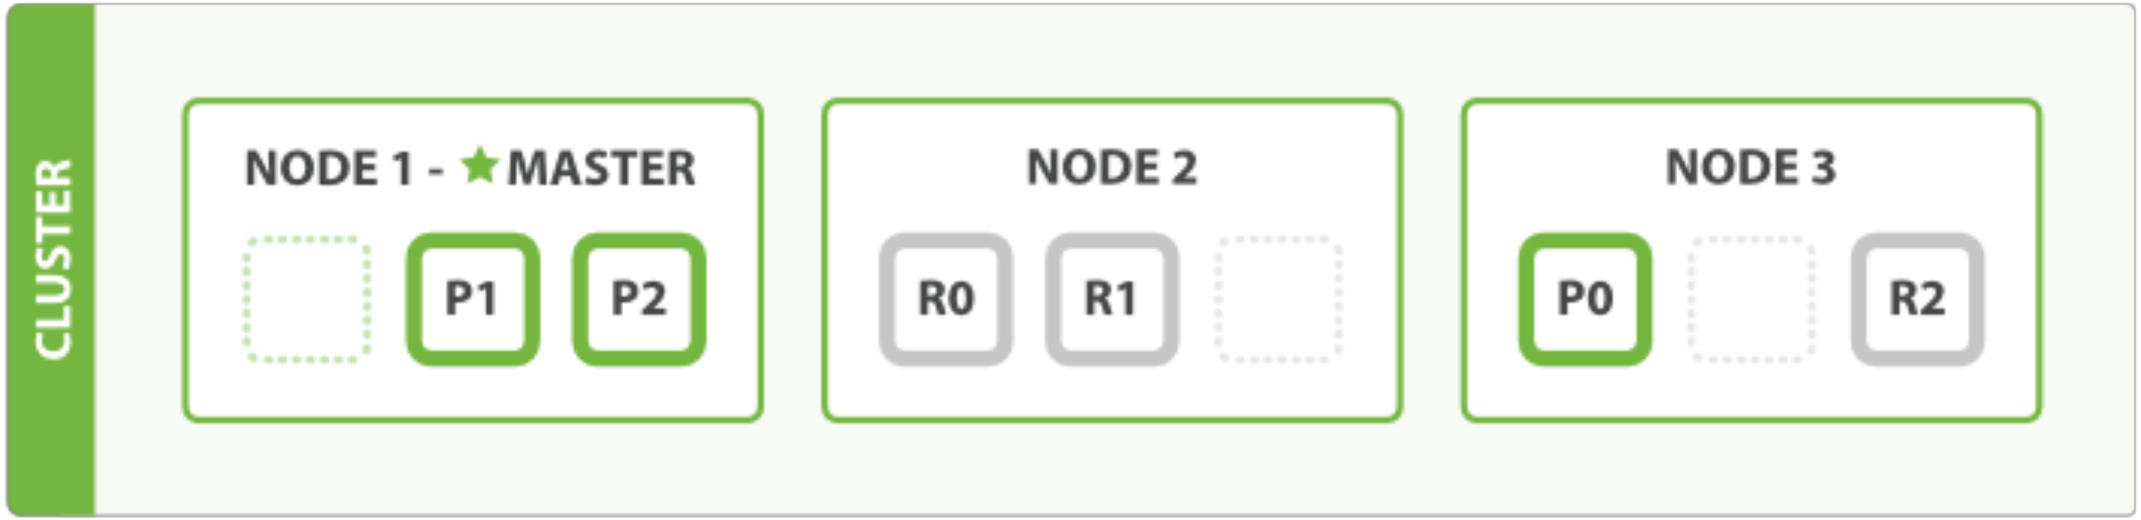

Write

Coordinating Node, Hash, Primary, Replica(s)

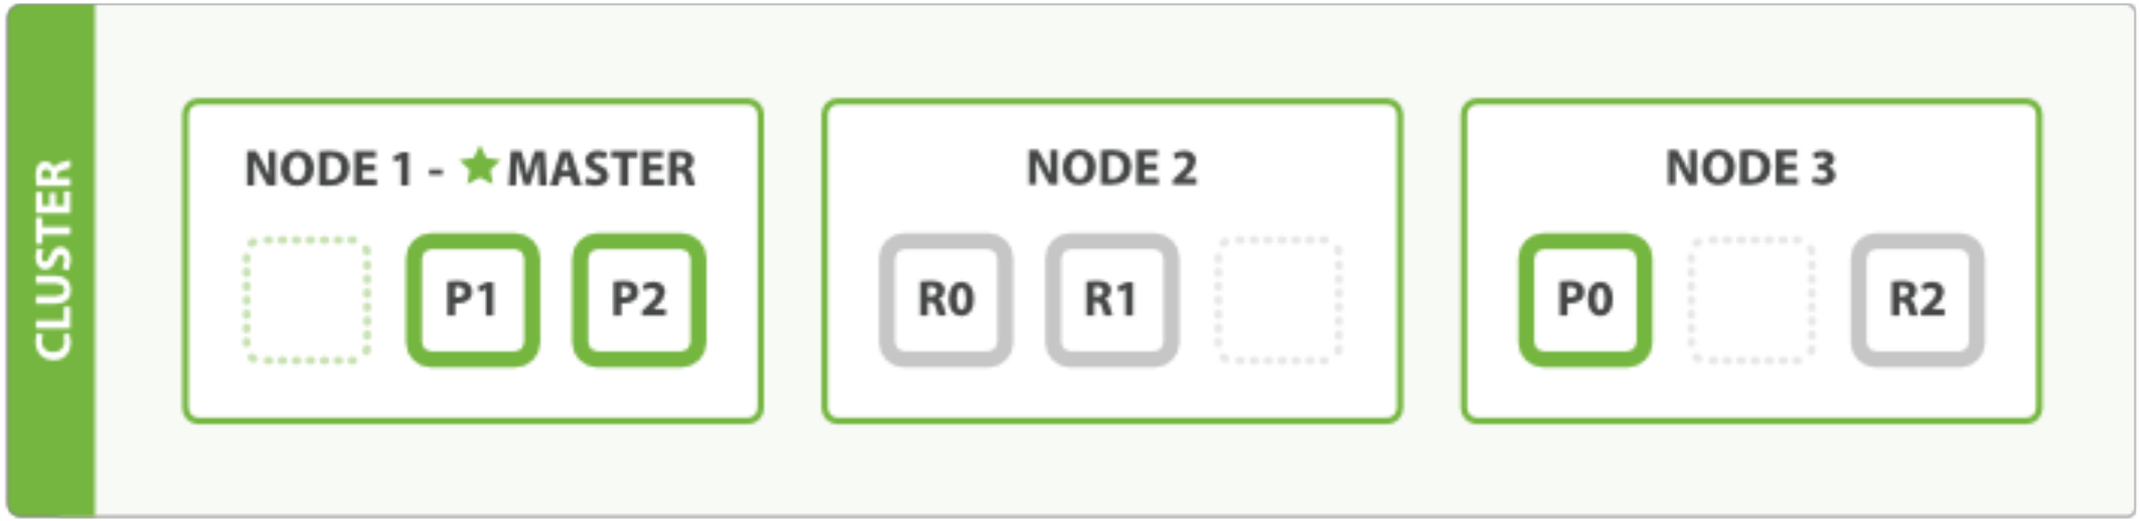

Get

Coordinating Node, Hash, Shard

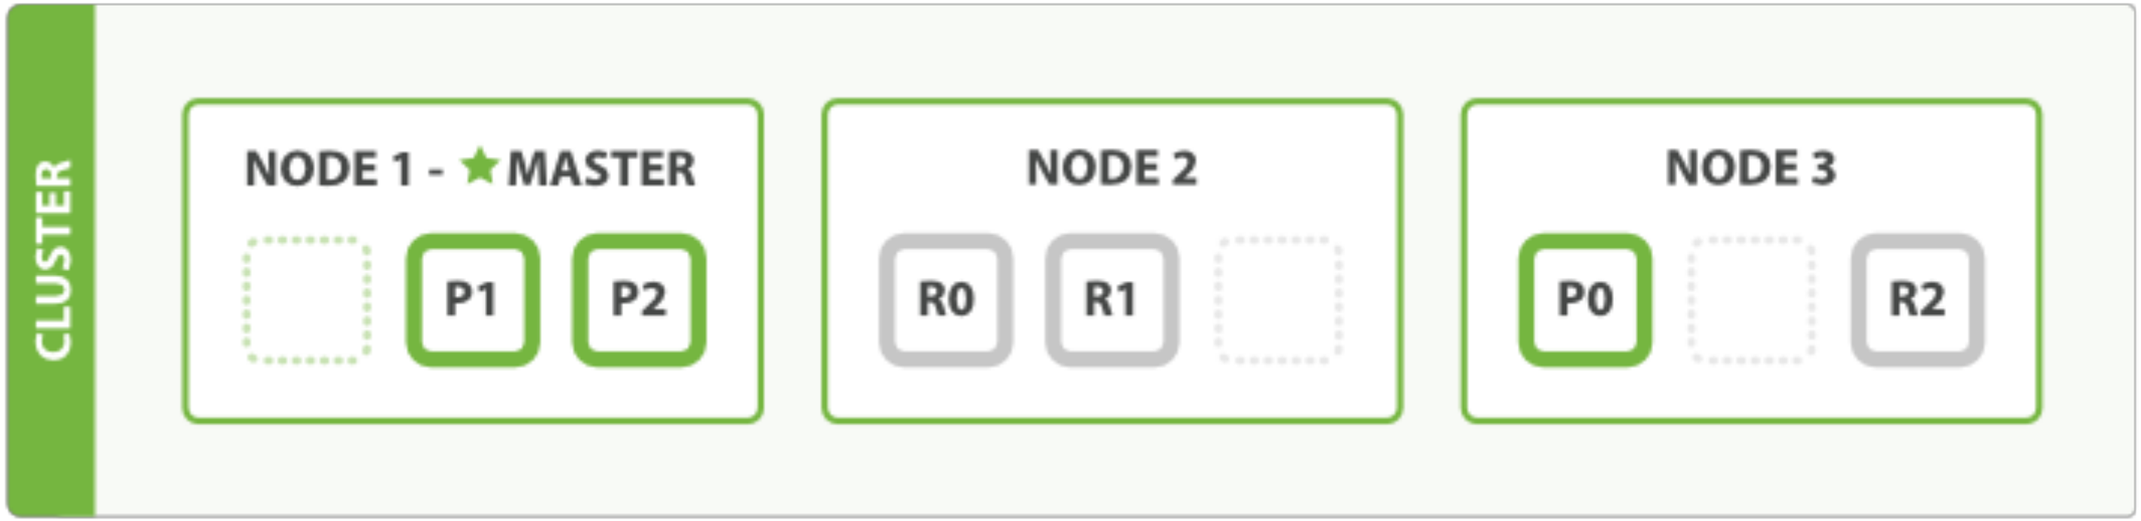

Search

Coordinating Node, Query then Fetch

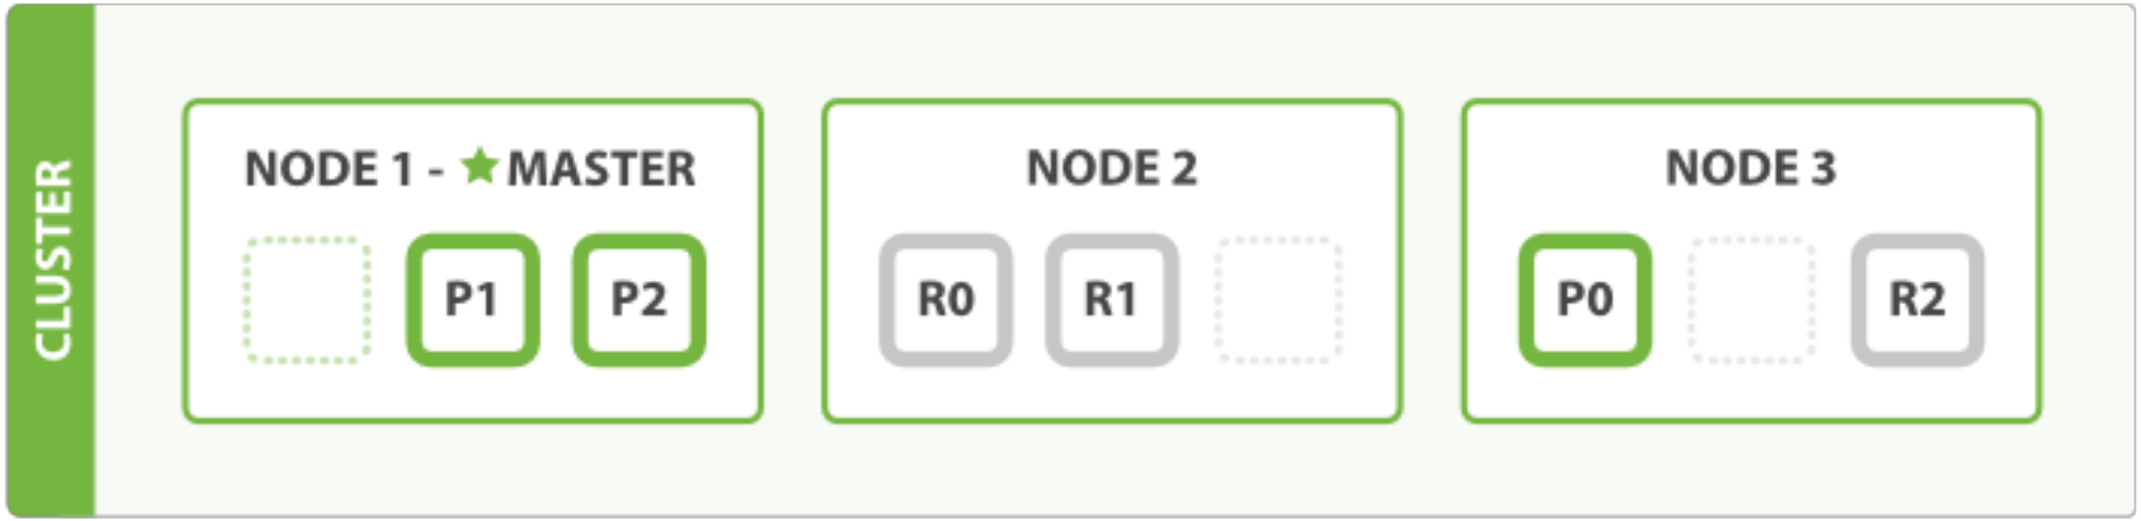

Append-Only Optimization

IDs assigned by coordinating node
Fast add instead of the slow update

https://github.com/elastic/elasticsearch/issues/19813

Lucene Segments

`index.refresh_interval: 1s`

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

Iff default refresh (added in 7.0)

Storage Compression

LZ4 (default)

DEFLATE (`best_compression`)

BKD Trees

Points in Lucene (added in 6.0)

`Integer` (1D 4 byte point) vs legacy `IntField`

elastic

Half & Scaled Floats

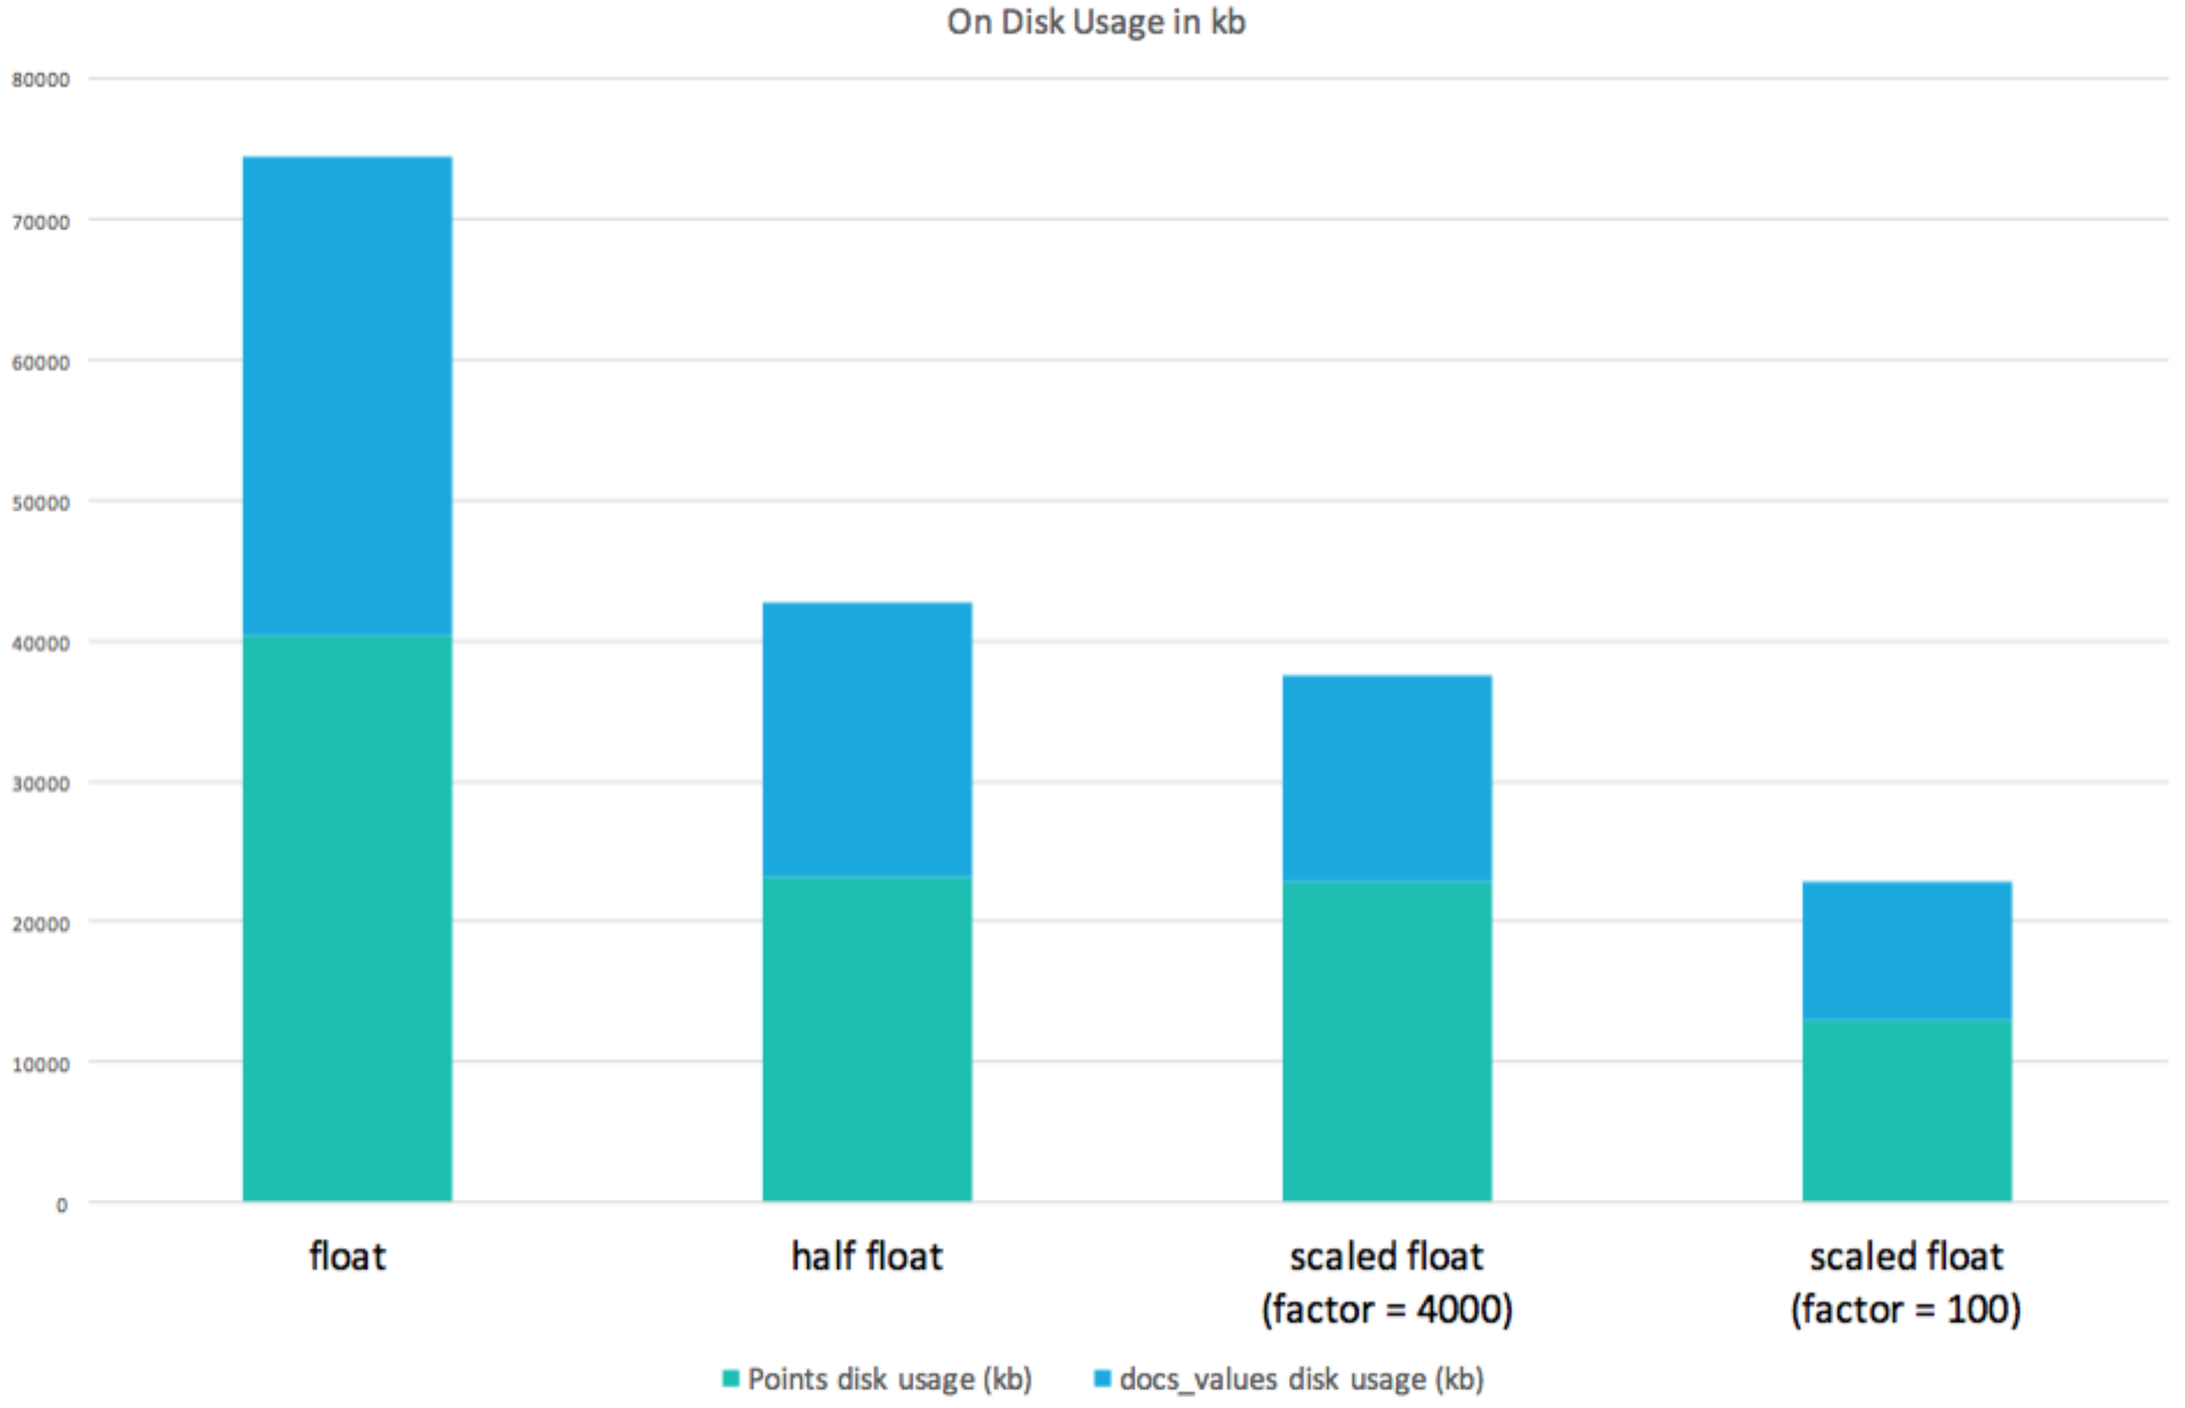
On Disk Usage in kb
80000
70000
60000
50000
40000
30000
20000
10000
0
float
half float
scaled float
(factor = 4000)
scaled float
(factor = 100)
Points disk usage (kb)
docs_values disk usage (kb)

https://github.com/elastic/beats/blob/master/metricbeat/module/system/load/_meta/fields.yml

```yaml
- name: load
  type: group
  description: >
    CPU load averages.
  release: ga
  fields:
    - name: "1"
      type: scaled_float
      scaling_factor: 100
      description: >
        Load average for the last minute.
    - name: "5"
      type: scaled_float
      scaling_factor: 100
      description: >
        Load average for the last 5 minutes.
    ...
```

`_all` Removal

https://www.elastic.co/guide/en/elasticsearch/reference/current/mapping-all-field.html

Doc Values

Replaced Fielddata

https://www.elastic.co/guide/en/elasticsearch/guide/current/_deep_dive_on_doc_values.html

Architecture

Time Based Indices

```
index: "metricbeat-%{[beat.version]}-%{+yyyy.MM.dd}"
```

Rollover Indices

Condition when to switch

think of the bytes

Rollups

Nodes



```
$ bin/elasticsearch
  -Enode.attr.rack=rack1
  -Enode.attr.size=hot


PUT /metricbeat/_settings
{
  "index.routing.allocation.include.size": "hot"
}
```

Frozen Indices

Ratio Heap : Storage

Read-only
No memory

elastic

Frozen Indices

Throttled Thread Pool

1 parallel search / node
100 in queue

Index Lifecycle Management

https://github.com/elastic/curator for snapshots

Features & Order

https://github.com/elastic/elasticsearch/blob/7.0/x-pack/plugin/core/src/main/java/org/elasticsearch/xpack/core/indexlifecycle/TimeseriesLifecycleType.java

```
static final List<String> ORDERED_VALID_HOT_ACTIONS = Arrays.asList(
  SetPriorityAction.NAME, UnfollowAction.NAME, RolloverAction.NAME
);
static final List<String> ORDERED_VALID_WARM_ACTIONS = Arrays.asList(
  SetPriorityAction.NAME, UnfollowAction.NAME, ReadOnlyAction.NAME,
        AllocateAction.NAME, ShrinkAction.NAME, ForceMergeAction.NAME
);
static final List<String> ORDERED_VALID_COLD_ACTIONS = Arrays.asList(
  SetPriorityAction.NAME, UnfollowAction.NAME, AllocateAction.NAME, FreezeAction.NAME
);
static final List<String> ORDERED_VALID_DELETE_ACTIONS = Arrays.asList(
  DeleteAction.NAME
);
```

Demo

Code: https://github.com/xeraa/scale-elasticsearch

Conclusion

Agenda

Building Blocks
Architecture
Demo

Benchmarks

Fair
Reproducible
Close to Production

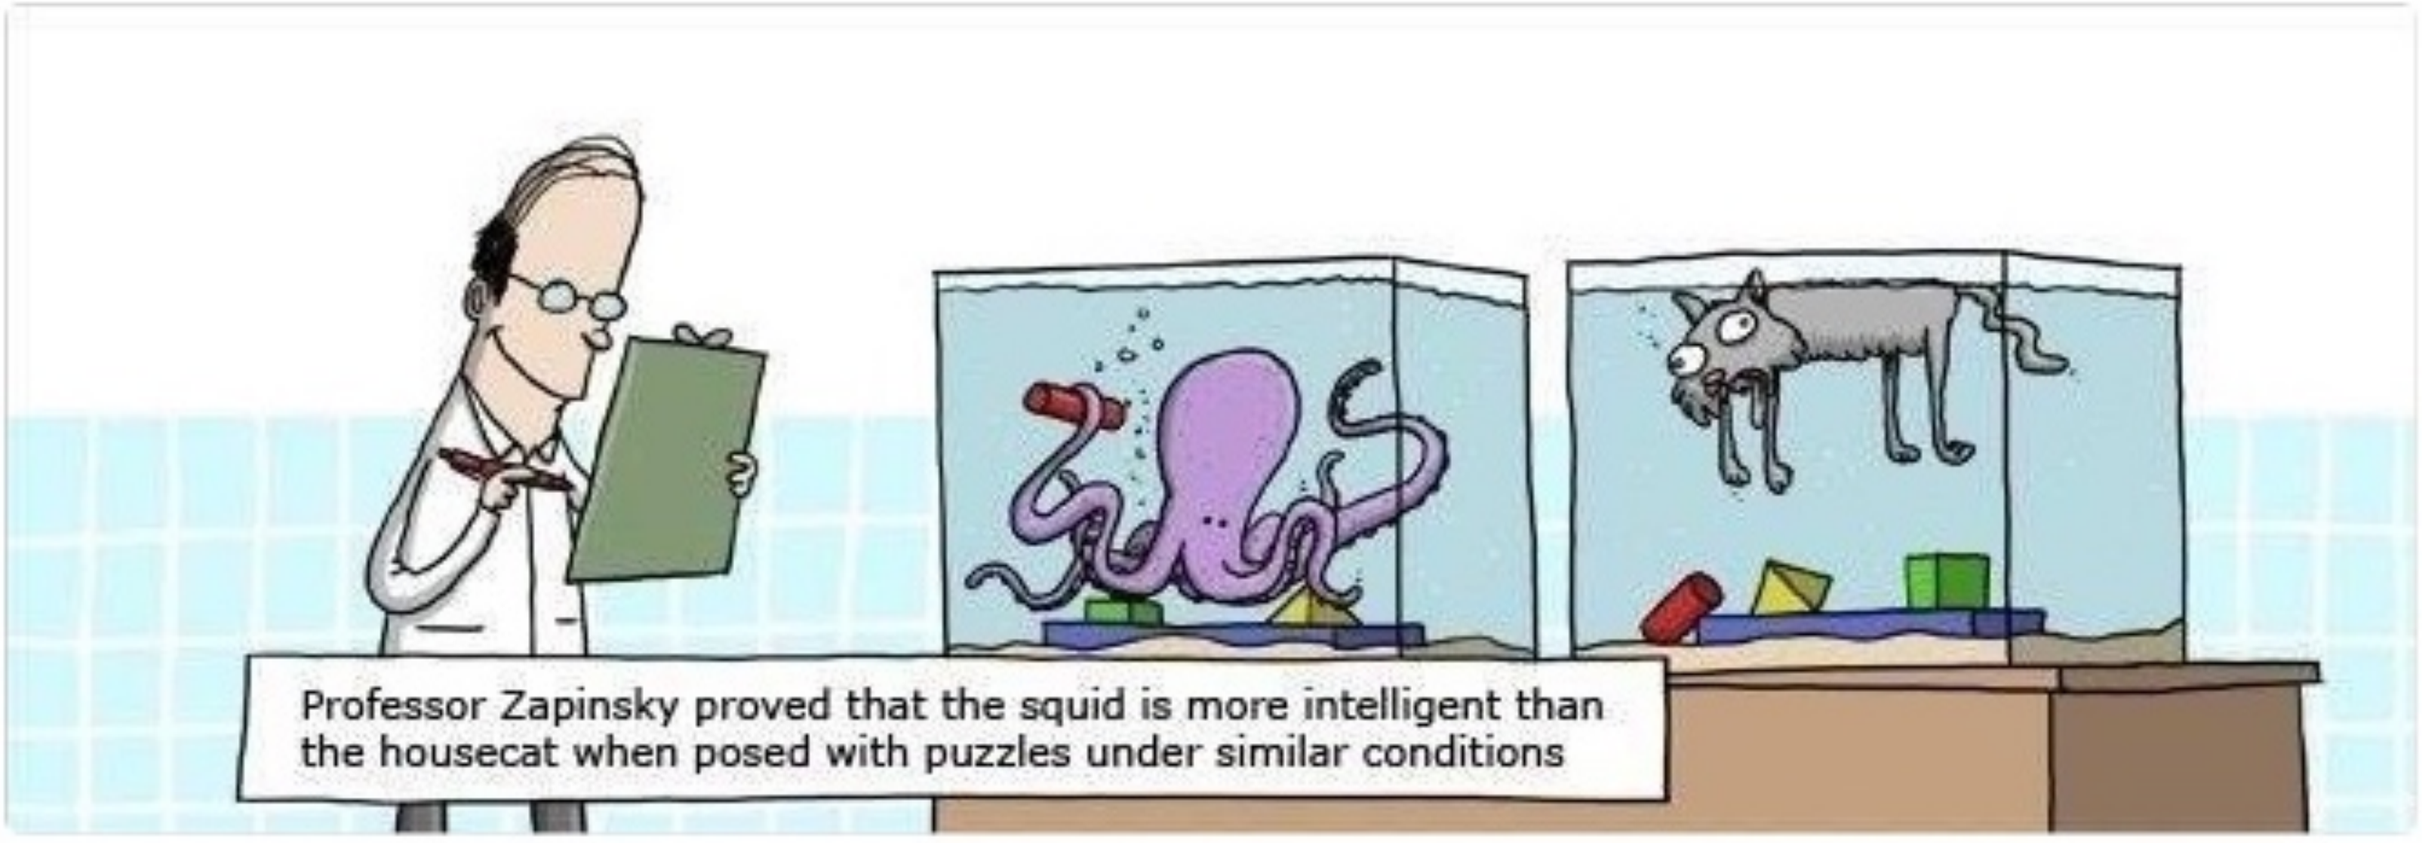
Professor Zapinsky proved that the squid is more intelligent than the housecat when posed with puzzles under similar conditions

Problems

Counters (incremental)
Single values instead of JSON docs

Questions?

Philipp Krenn @xeraa

elastic