



Operating Elasticsearch in Kubernetes

[Container] Conf

MIKKEL LARSEN

[@mikkeloscar](https://twitter.com/mikkeloscar)

OLIVER TROSIEN

[@otrosien](https://twitter.com/otrosien)

2019-11-14



WHO ARE WE?



Oliver
Core Search Services
Elasticsearch



Mikkel
Cloud Infrastructure
Kubernetes

SEARCH @ ZALANDO

400k+

Products
per country

~2500

Brands

~1000

Categories

65%

Mobile Traffic

15K

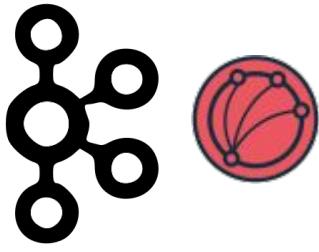
QPS

8K

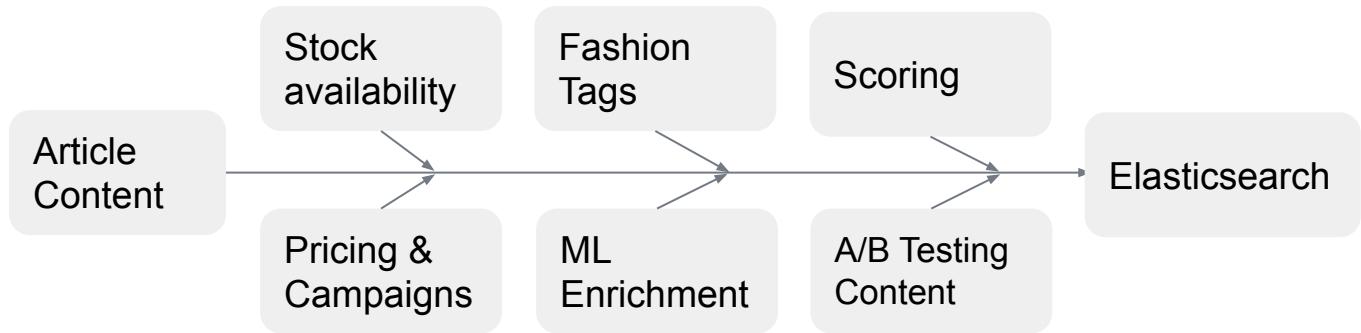
Updates/s



DATA INGESTION



Kafka / Nakadi *



* nakadi.io

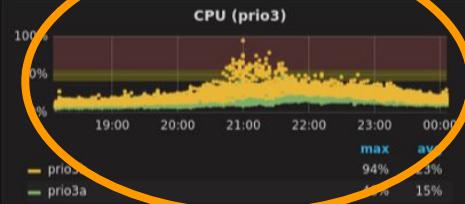
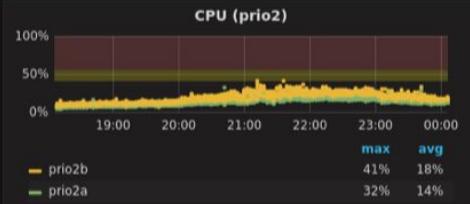
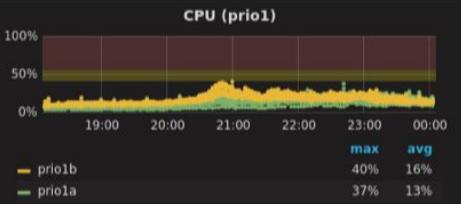
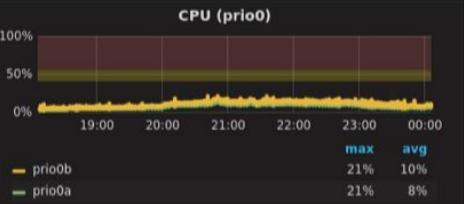
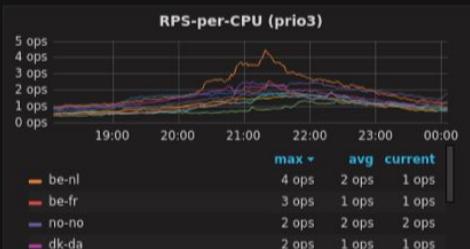
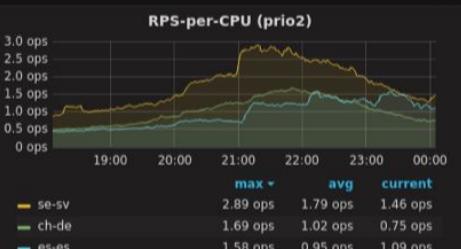
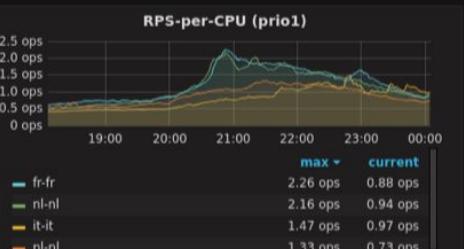
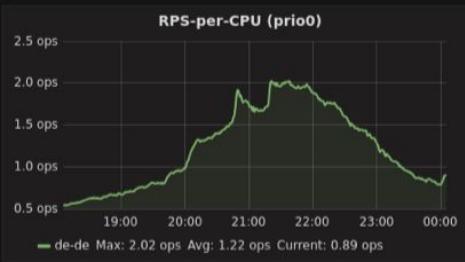
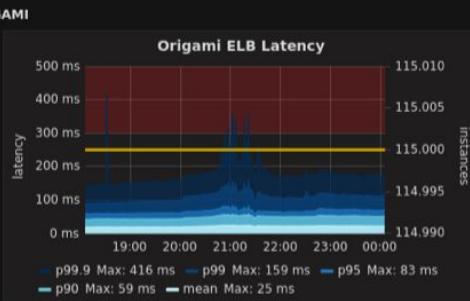
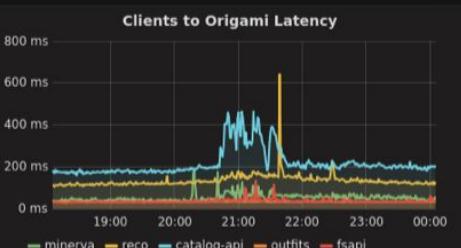
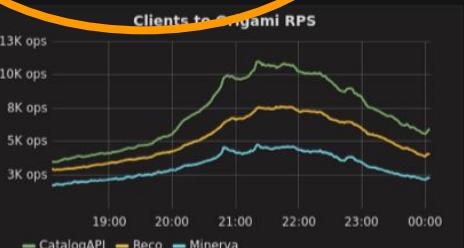


ES Cluster State

ES Active Shards
1391

ES Nodes

ES Indices

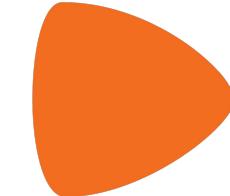


KUBERNETES @ ZALANDO

Default
Deployment
Target

3500~
nodes

Since
Oct 2016



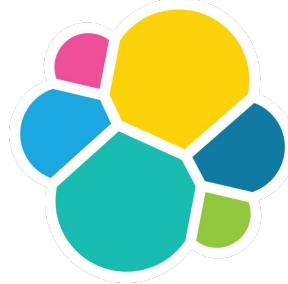
143
clusters

Node
Autoscaling

From v1.4
to v1.14

RUNNING ELASTICSEARCH IN KUBERNETES

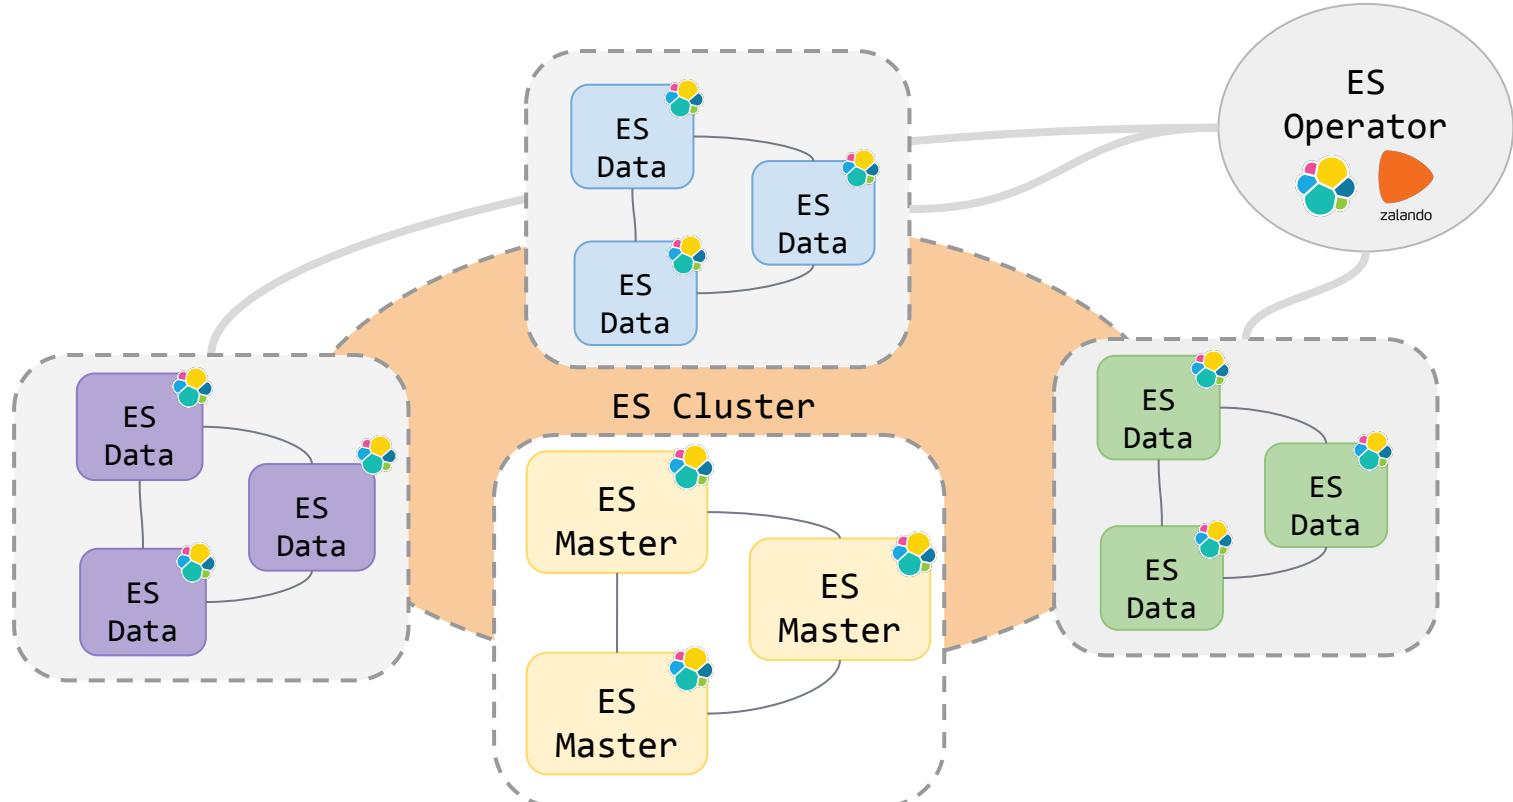
- 1. Safe automatic updates**
(Including Kubernetes cluster updates)
- 2. Advanced auto-scaling for operational stability and cost efficiency**



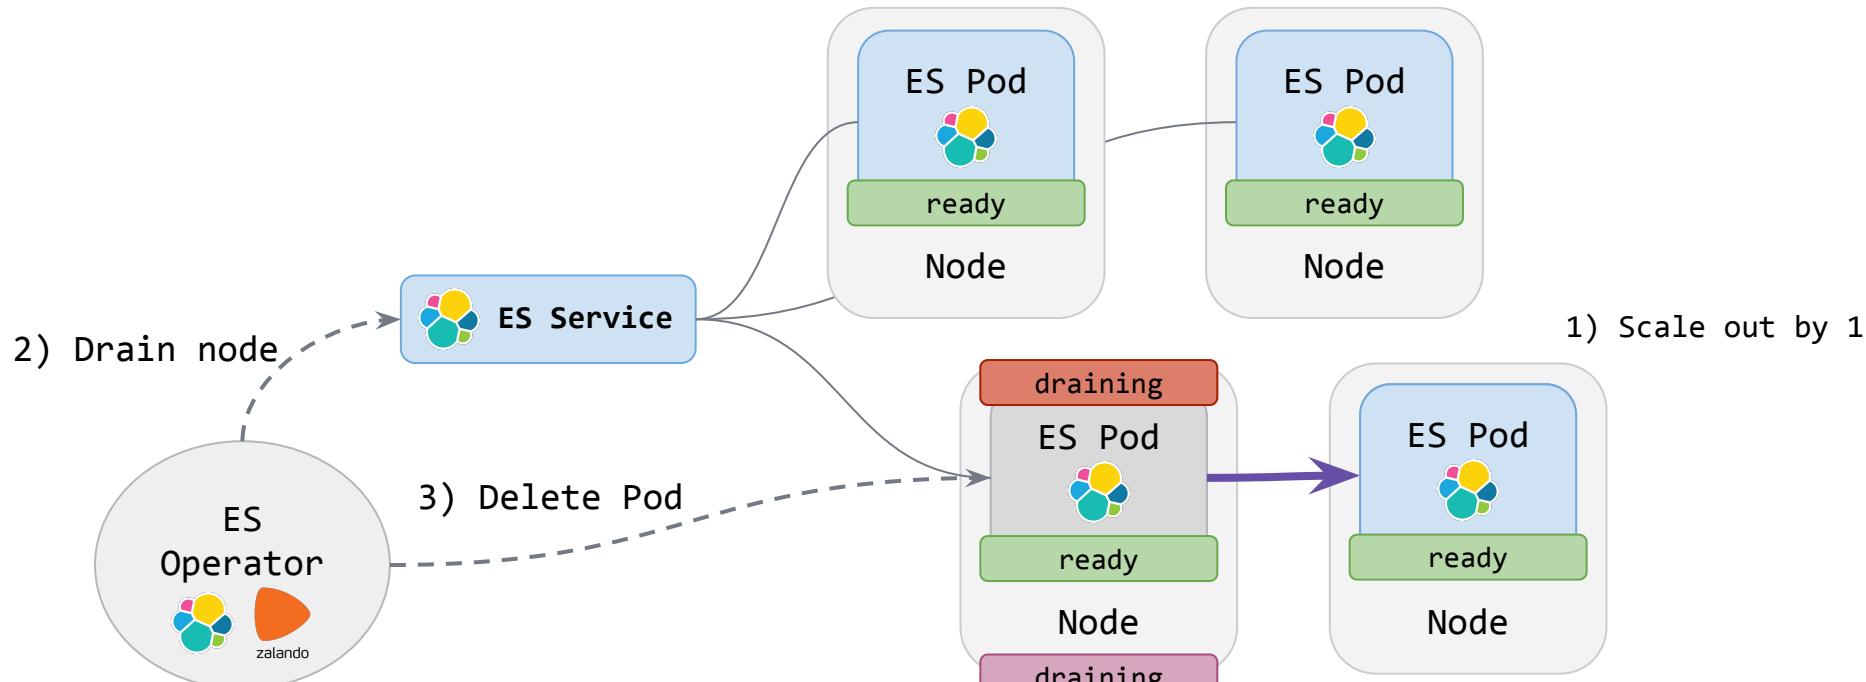
ELASTICSEARCH DATA SETS

```
apiVersion: zalando.org/v1
kind: ElasticsearchDataSet
metadata:
  name: test-cluster
spec:
  scaling:
    {...}
  replicas: 3
  template: # PodTemplate
    {...}
  volumeClaimTemplates:
    {...}
```

ELASTICSEARCH DATA SETS



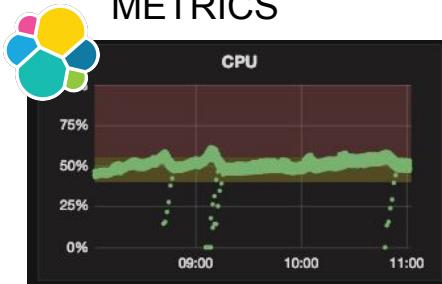
UPDATING ELASTICSEARCH (OPERATOR)



github.com/zalando-incubator/es-operator

SCALING UP ELASTICSEARCH (1)

METRICS

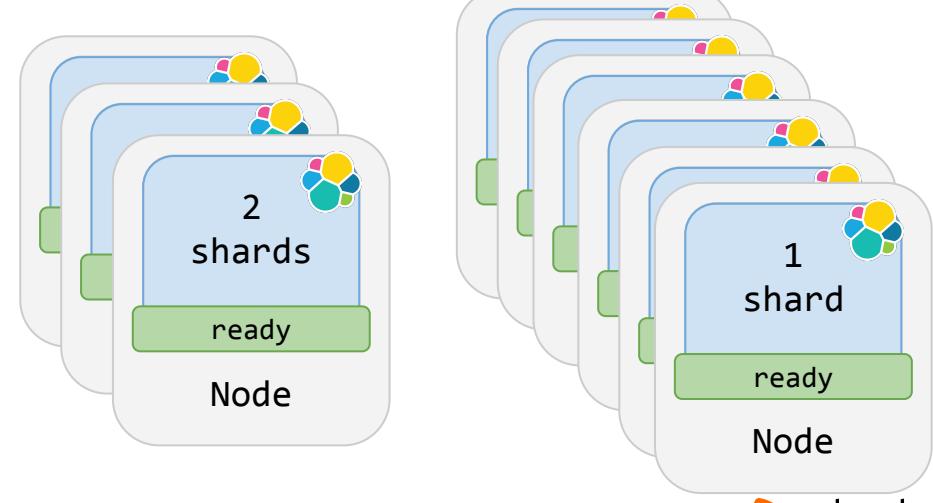
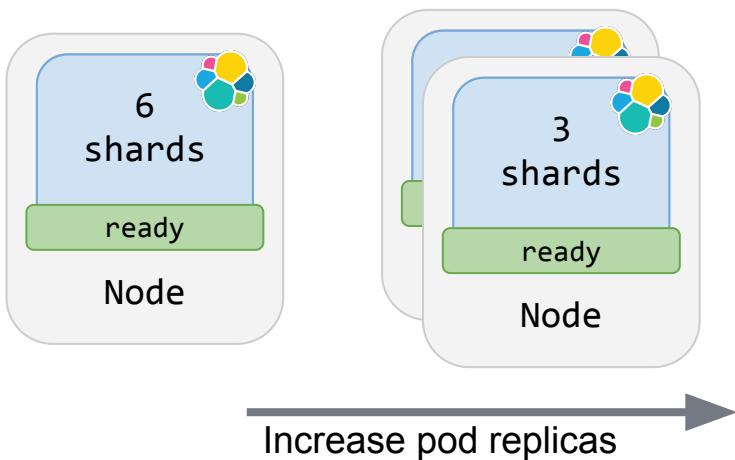


Thresholds

- CPU
- Duration
- Cooldown

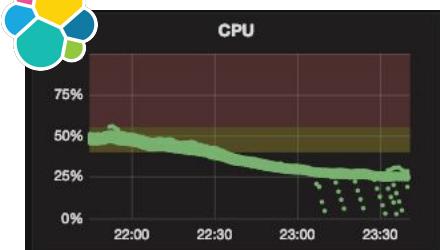
Boundaries

- Max # Pod replicas



SCALING DOWN ELASTICSEARCH

METRICS



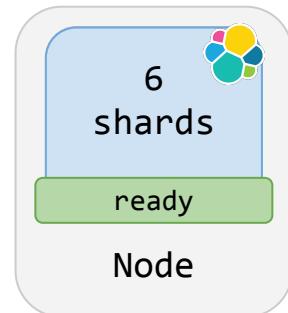
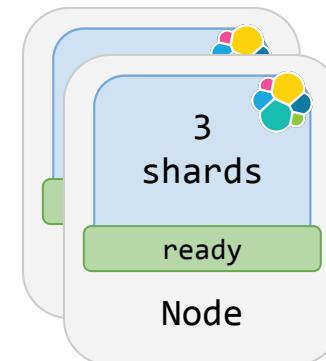
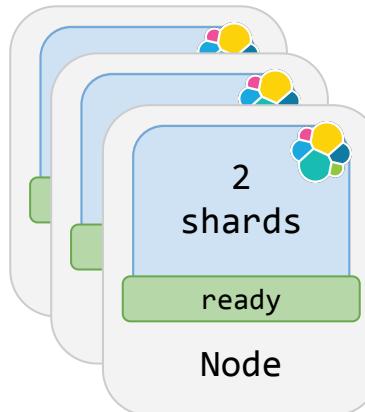
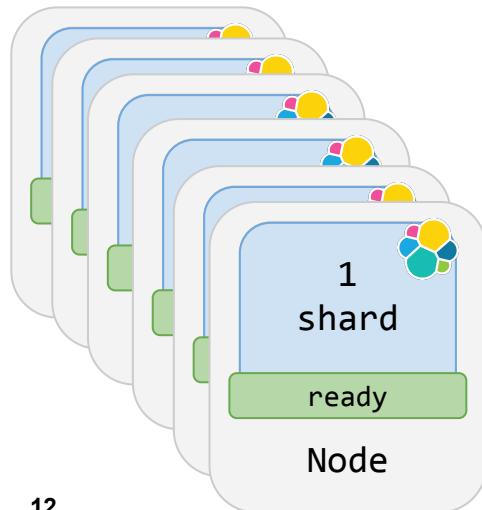
Thresholds

- CPU
- Duration
- Cooldown

Boundaries

- Min # Replica
- Max # Shards per node
- Max disk usage (%)

DON'T OPERATE
WHEN CLUSTER
IS NOT GREEN!



Decrease Pod replicas

SCALING UP ELASTICSEARCH (2)

METRICS

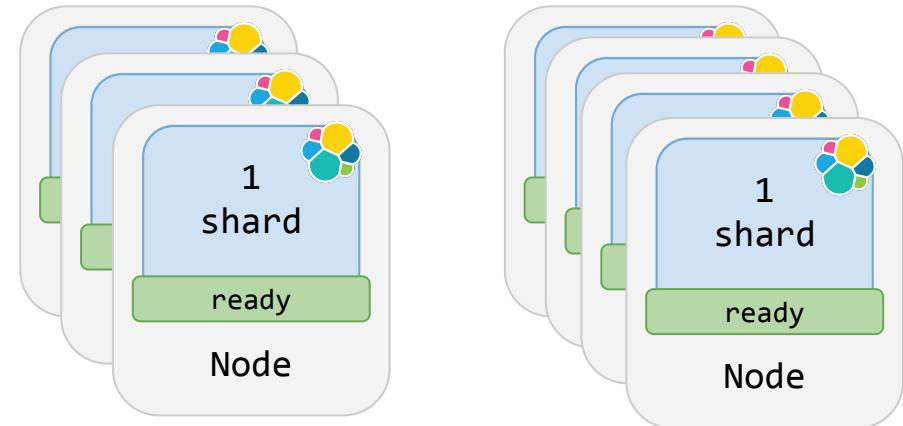
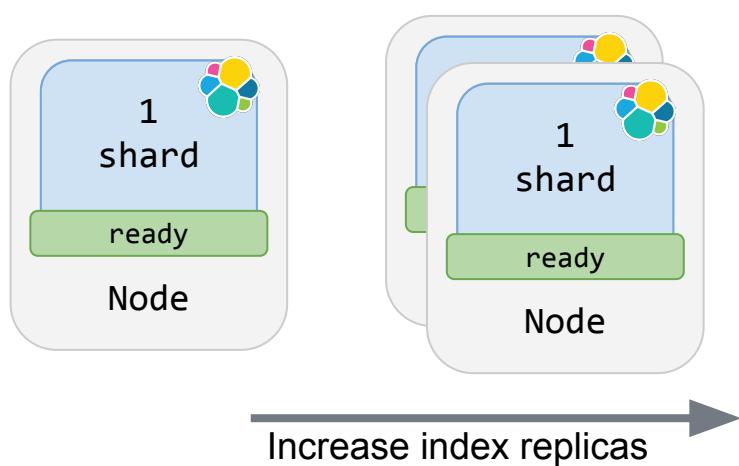


Thresholds

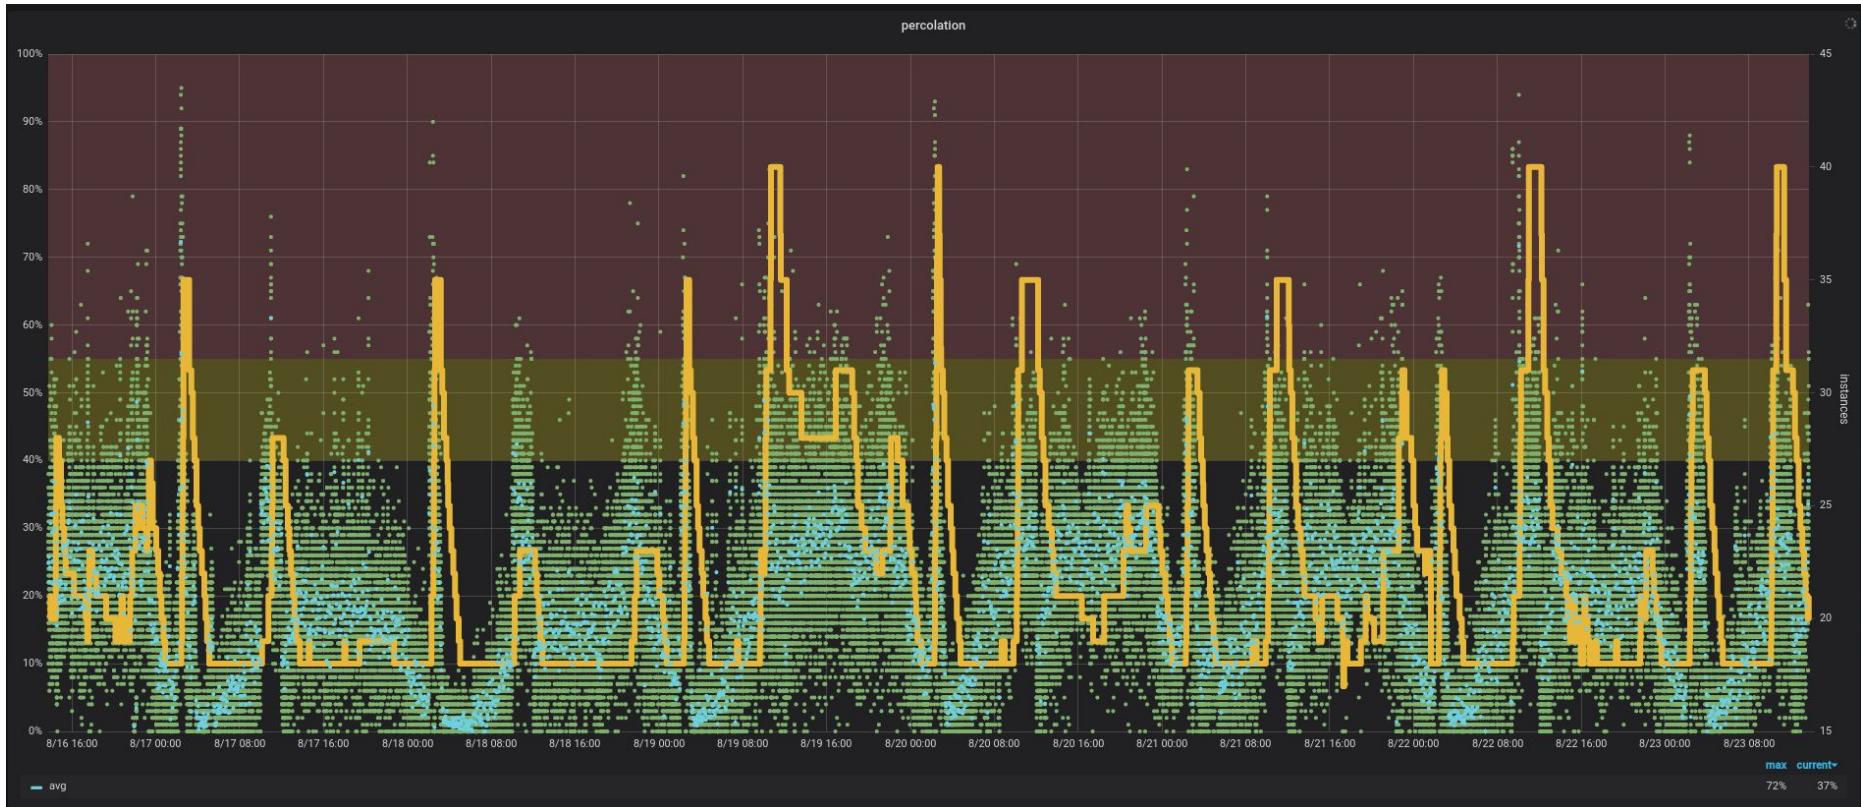
- CPU
- Duration
- Cooldown

Boundaries

- Min # Shards per node
- Max # Pod replicas



AUTO-SCALING STACK (7d)



SAFE SCALING IN PRODUCTION (7d)



HIGHLIGHTS / LOWLIGHTS



- Stacks more sensitive to sudden load spikes (Push Notifications, bot traffic, DDOS)
- Incidents
- Cost reduction
- Improved time-to-rollout of new features / upgrades
- Greatly reduced operational effort
- Helped other teams migrate their ES

DEMO



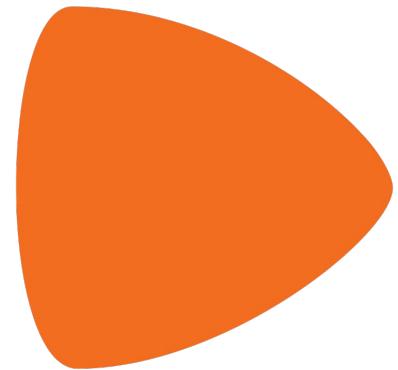
REFERENCES

ES Operator

github.com/zalando-incubator/es-operator

Demo

github.com/otrosien/es-operator-demo





MIKKEL LARSEN

mikkel.larsen@zalando.de

[@mikkeloscar](https://www.twitter.com/mikkeloscar)

OLIVER TROSIEN

oliver.trosien@zalando.de

[@otrosien](https://www.twitter.com/otrosien)

2019-11-14

