



Es-operator

Building an Operator From the Bottom Up



KubeCon



CloudNativeCon

Europe 2019

MIKKEL LARSEN

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

2019-05-21



\$ whoami



Mikkel Larsen

Software Engineer

Cloud Infrastructure @ Zalando SE

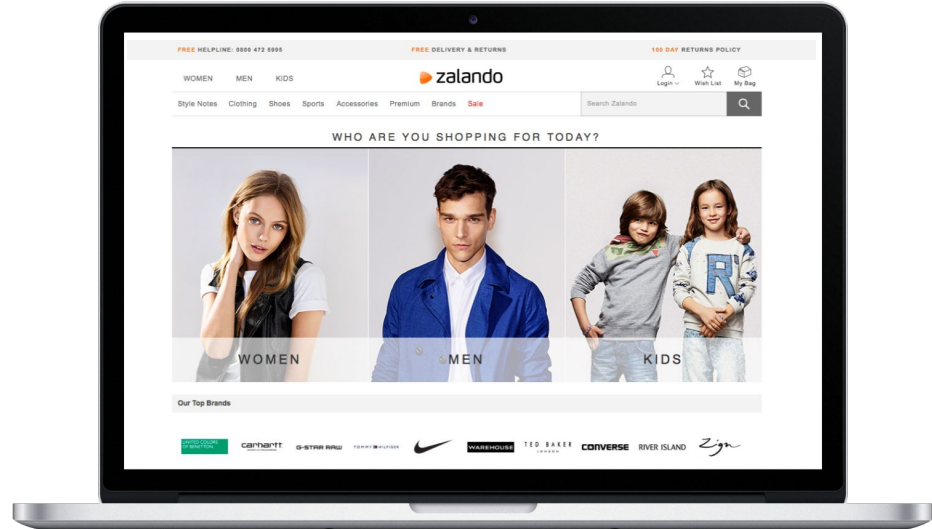



[@mikkeloscar](https://twitter.com/mikkeloscar)



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

“EUROPE’S LEADING ONLINE FASHION PLATFORM”



WE BRING FASHION TO PEOPLE IN 17 COUNTRIES

17 markets


7 fulfillment centers

26 million active customers

5.4 billion € revenue 2018

250 million visits per month

15,000 employees in Europe




KUBERNETES @ ZALANDO

Default
Deployment
Target

1400~
nodes

Since
Oct 2016



zalando

~125
clusters

Node
Autoscaling

From v1.4
to v1.13

SEARCH @ ZALANDO

300k+
Products
per country

~2000
Brands







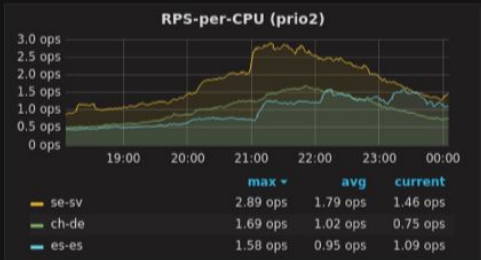



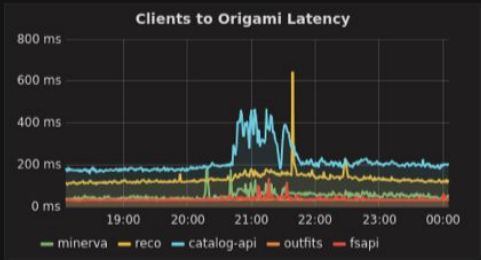
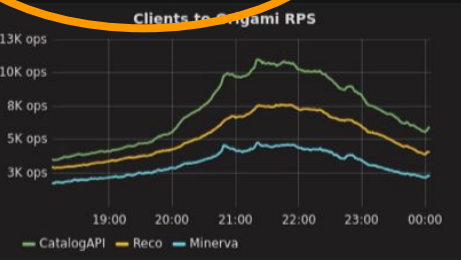


~700
Categories



45%
Mobile Traffic

12K
QPS

8K
Updates/s



RUNNING ELASTICSEARCH IN KUBERNETES

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



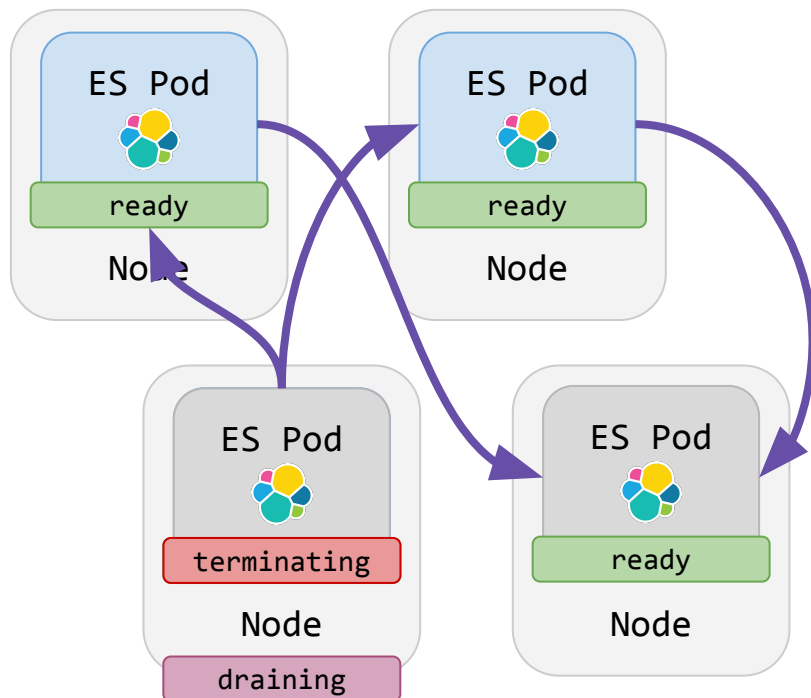
UPDATING ELASTICSEARCH (STATEFULSET)

1) PreStop Hook ([bash script](#))

- Exclude node in ES
- Wait for node to drain (up to 1h)
- Data is moved to existing nodes

2) PostStart Hook ([bash script](#))

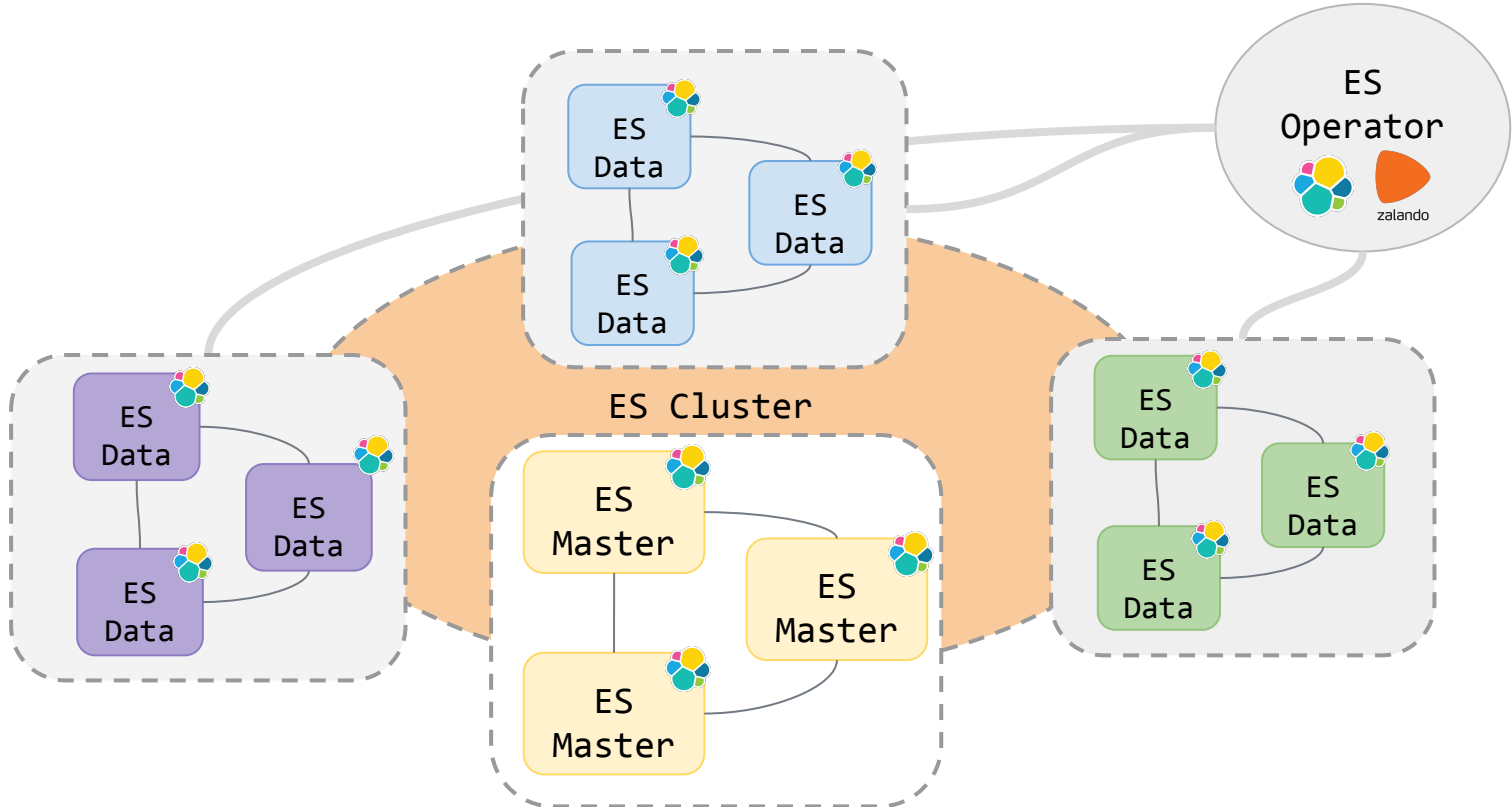
- Remove all excludes
- Let ES rebalance from existing nodes



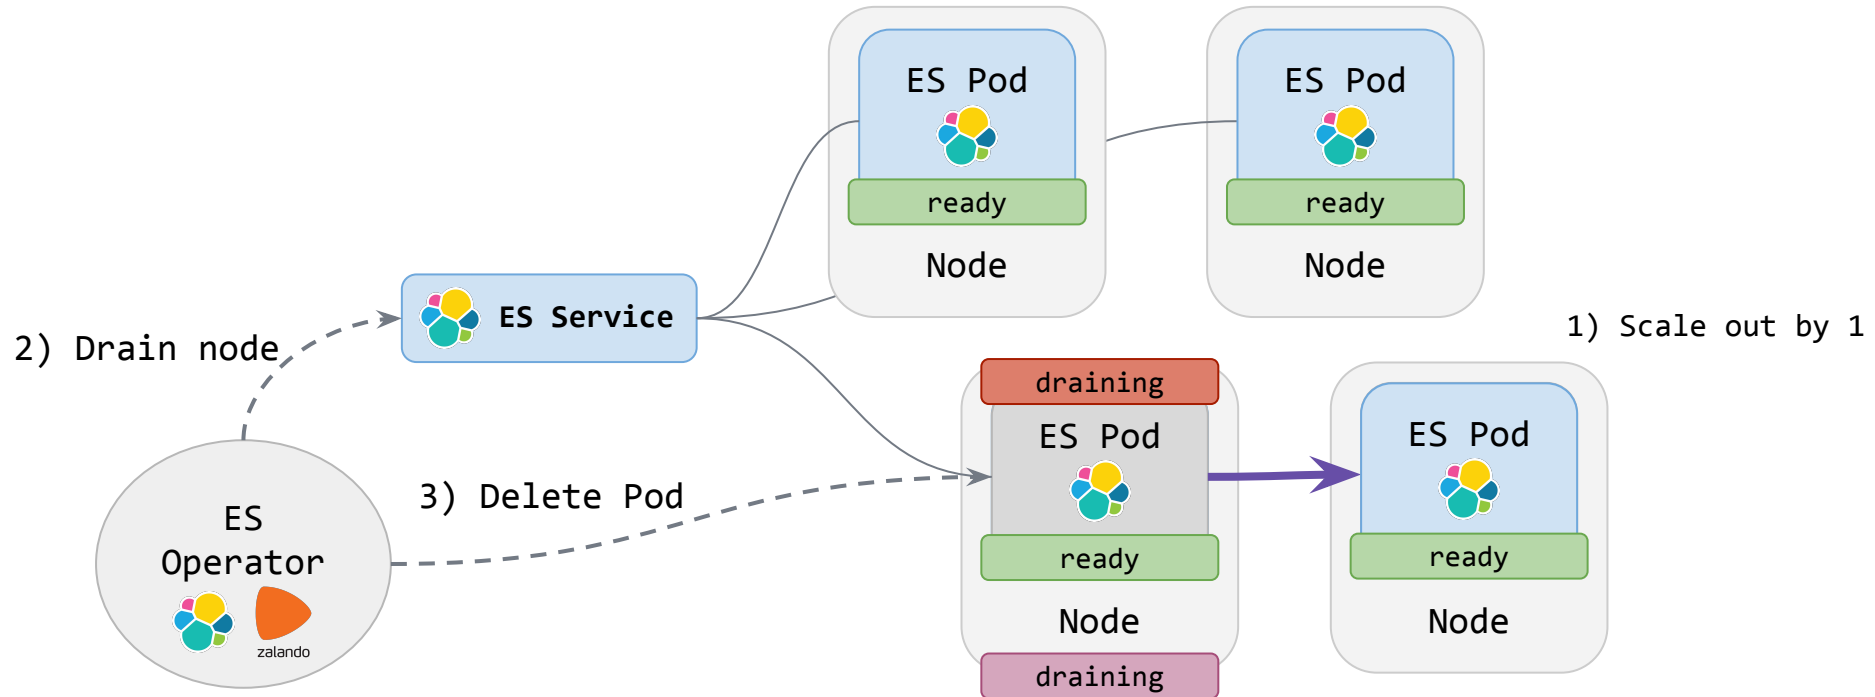
ELASTICSEARCH DATA SETS

```
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

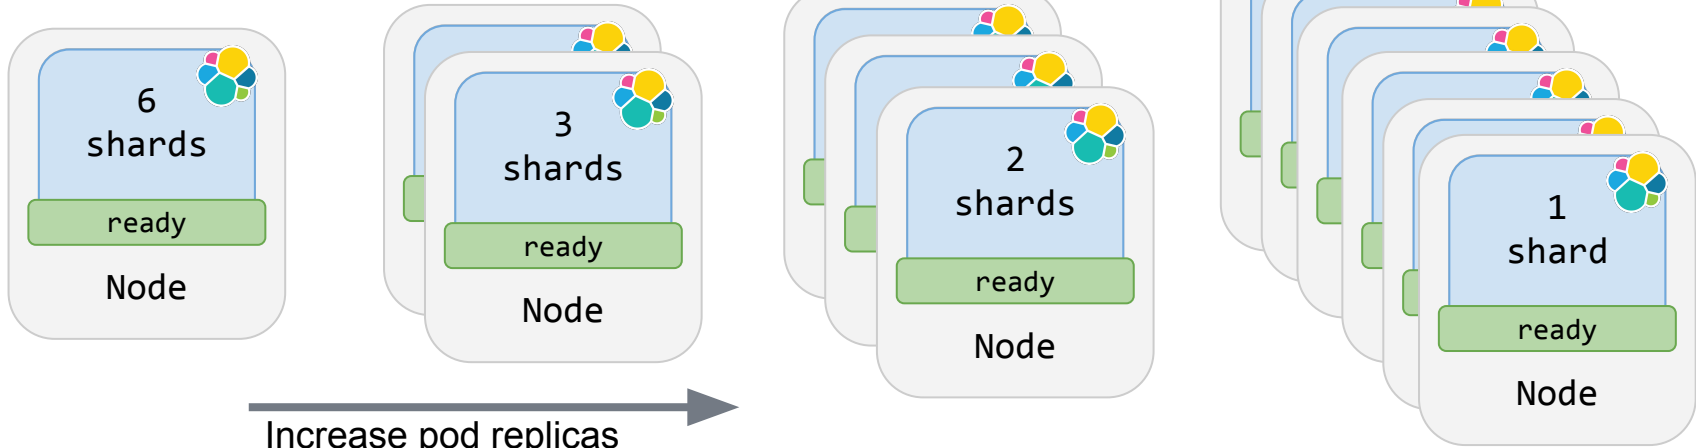


Thresholds

- CPU
- Duration
- Cooldown

Boundaries


- Max # Pod replicas



SCALING DOWN ELASTICSEARCH

DON'T OPERATE WHEN CLUSTER IS NOT GREEN!

METRICS







Thresholds

- CPU
- Duration
- Cooldown

Boundaries


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



Decrease Pod replicas →

SCALING UP ELASTICSEARCH (2)

METRICS

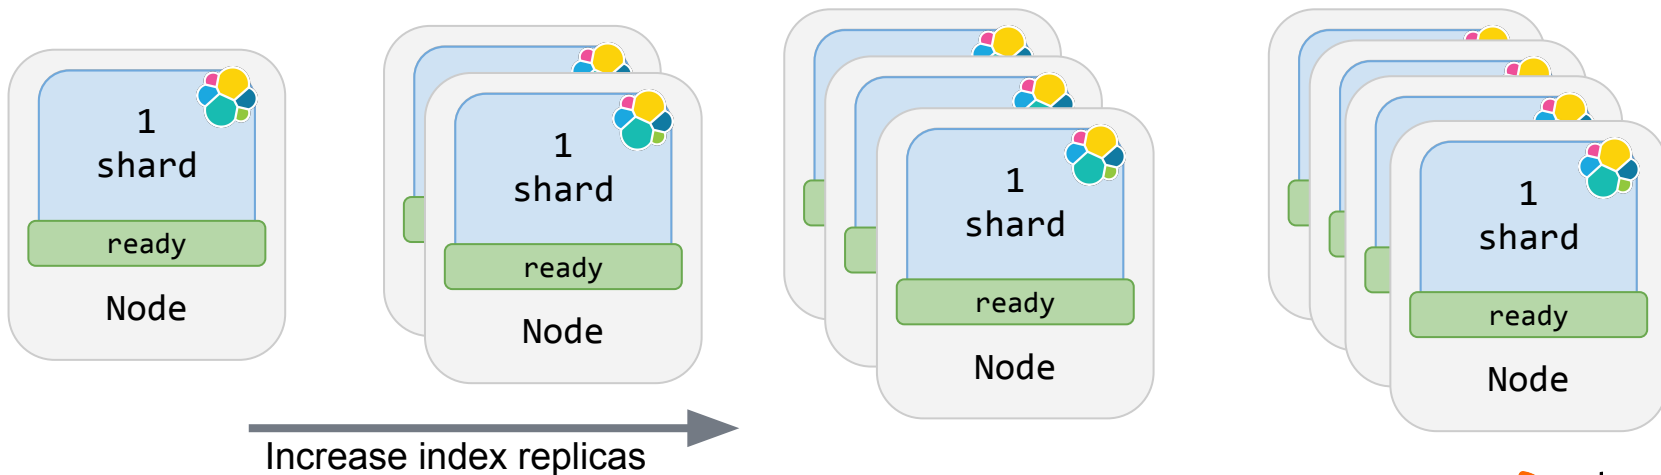


Thresholds

- CPU
- Duration
- Cooldown

Boundaries

- Min # Shards per node
- Max # Pod replicas



OPEN SOURCE

Elasticsearch Operator

github.com/zalando-incubator/es-operator

Kubernetes on AWS

github.com/zalando-incubator/kubernetes-on-aws

Kubernetes Operator Pythonic Framework (Kopf)

github.com/zalando-incubator/kopf



zalando



¡GRACIAS!



MIKKEL LARSEN

mikkel.larsen@zalando.de



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



2019-05-21