

# Exploring OpenSearch deployment in Kubernetes

Cristobal Gimenez Devis

Supervisors: Sokratis Papadopoulos, Luis Pigueiras.

# Table of Contents

**01**

Project definition

**02**

Project motivation

**03**

Project work

# Project definition

- **OpenSearch** is a distributed, search and analytics engine, very popular for log analytics and full-text search
- Offered as a service at CERN since 2016
  - deployed on bare metal machines using **Puppet**
- **Kubernetes** is the most popular container orchestration engine
- **Objective:** explore OpenSearch deployment in **Kubernetes**
  - Create a prototype deployment within the IT Kubernetes service
  - Compare against current puppet deployment



**kubernetes**

# Project motivation

## Flexibility

- Easier to scale, optimized resource management

## Management

- Simplification of operations, easier migration to the cloud

## Community standard

- Increased utilization of community resources



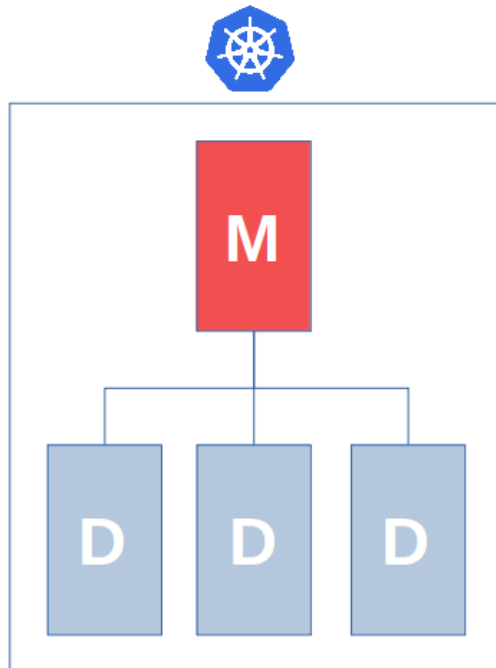
**kubernetes**

# Project work | Kubernetes clusters

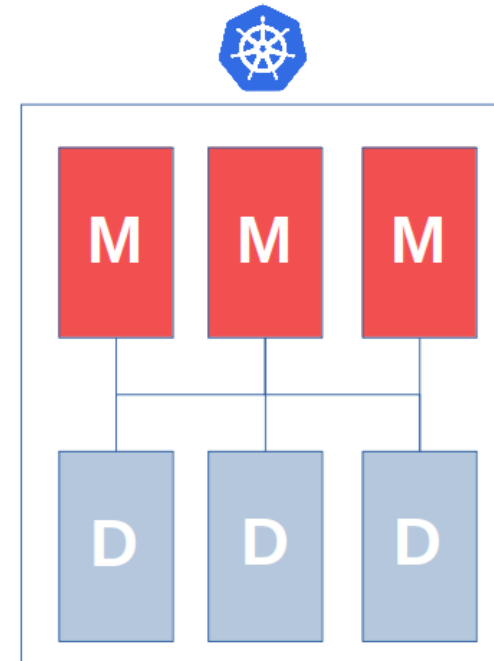


**MAGNUM**  
an OpenStack Community Project

## Single-master cluster



## Multi-master cluster



# Project work | OpenSearch deployment



- Experimented with two ways of deployment: with & without k8s operator

## OpenSearch K8s operator

Offers automation on common operations

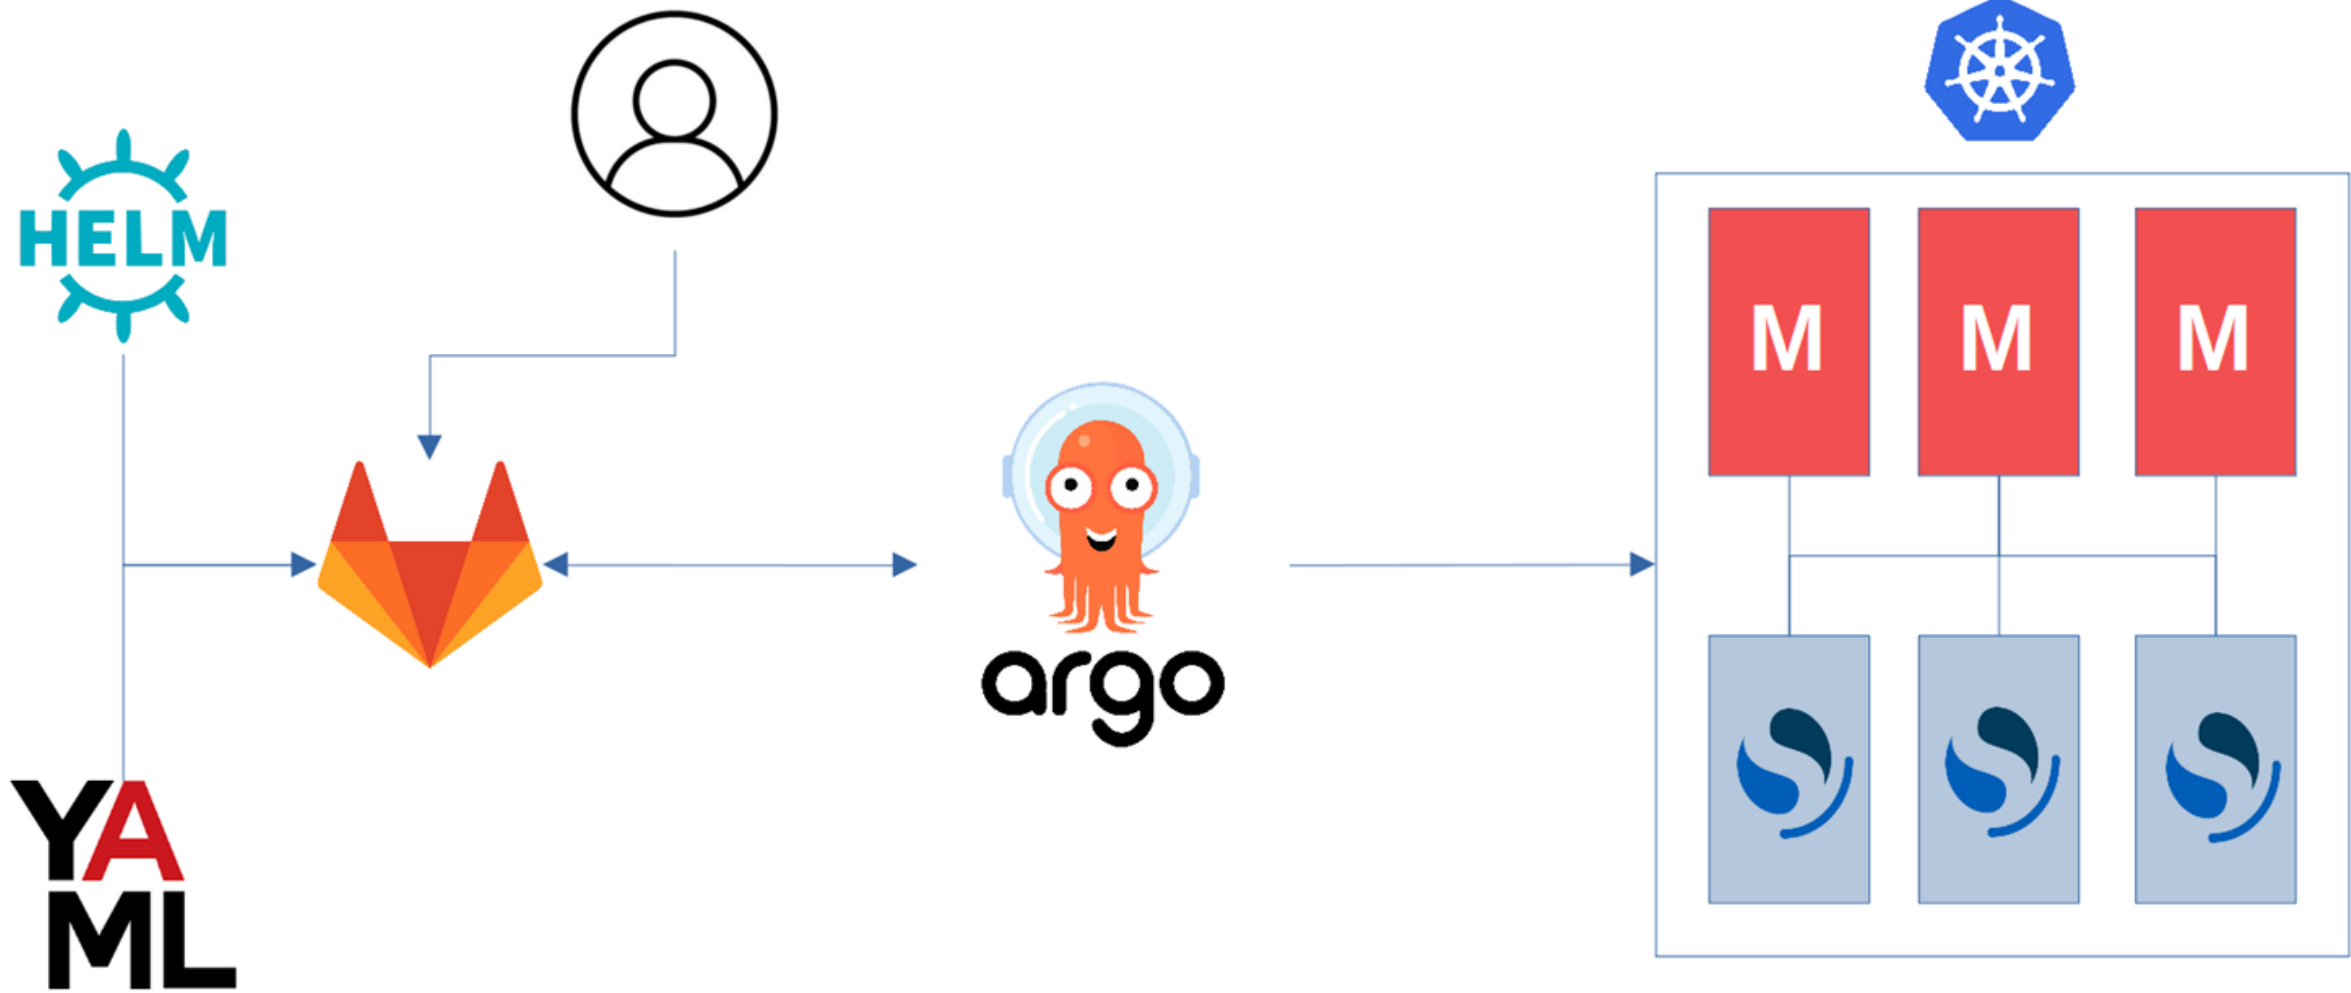
Requires custom work for better GitOps monitoring

## OpenSearch Helm chart

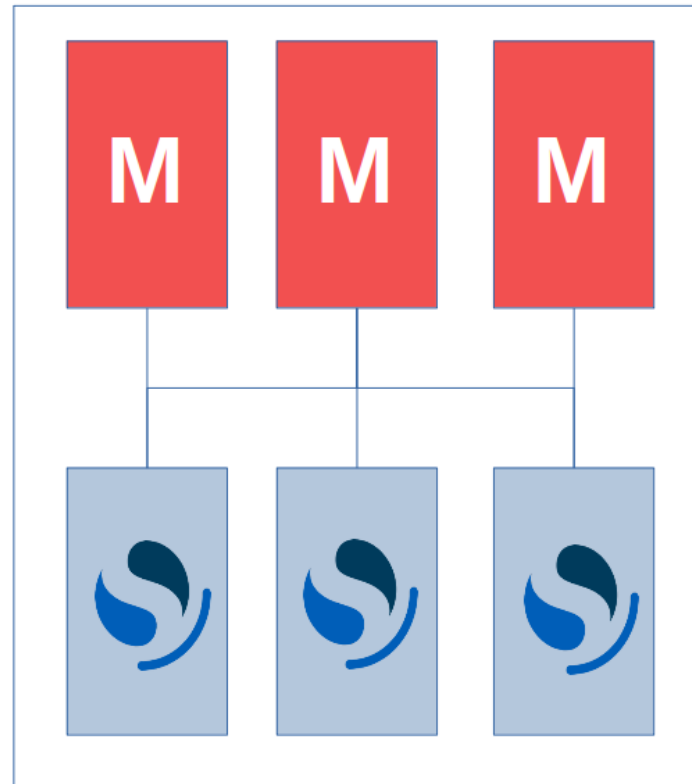
Requires more work to operate

Better default GitOps monitoring

# Project work | GitOps

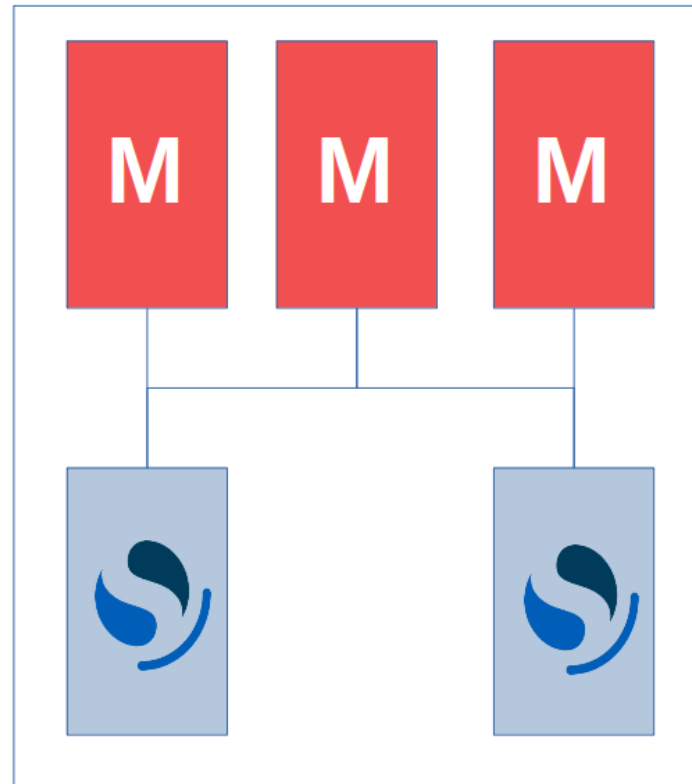


# Project work | Failures

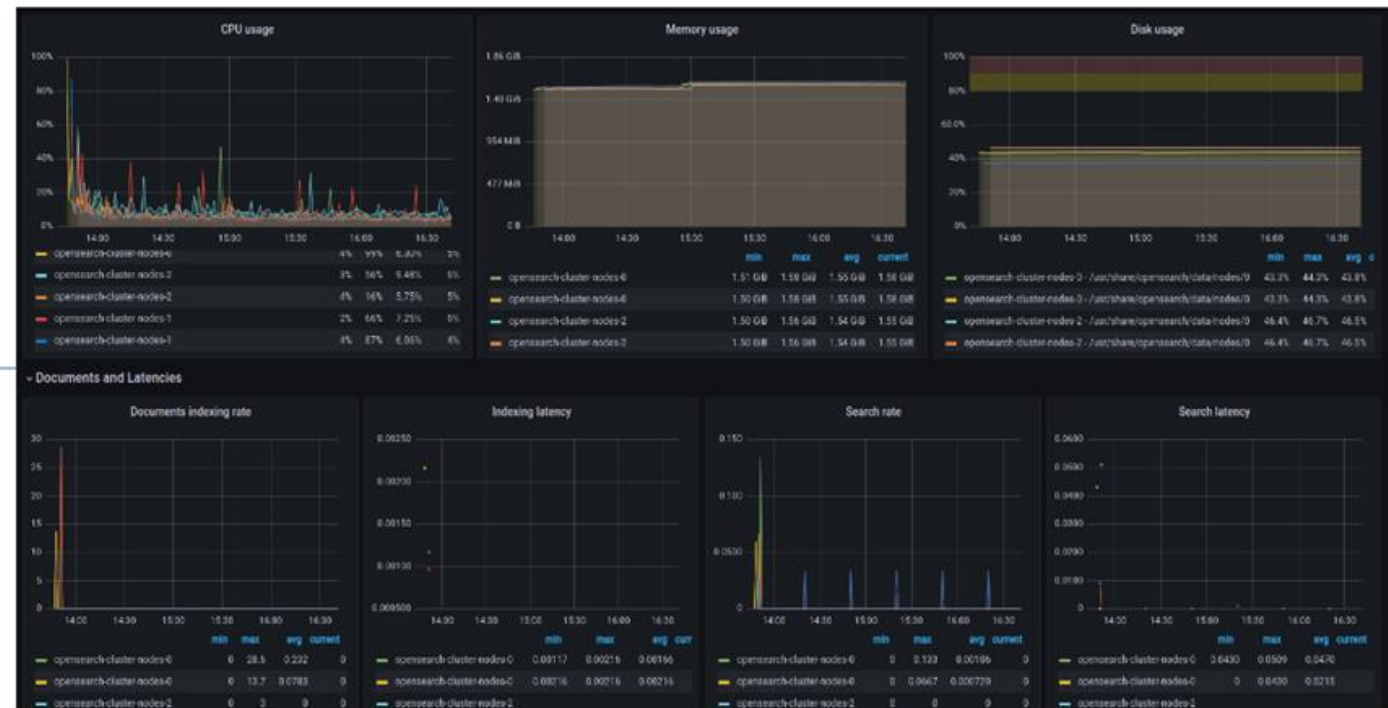






# Project work | Failures



# Project work | Monitoring



CPU usage



Node	min	max	avg	current
opensearch-cluster-nodes-u	4%	99%	6.30%	5%
opensearch-cluster-nodes-2	3%	56%	9.48%	6%
opensearch-cluster-nodes-2	4%	16%	5.75%	5%
opensearch-cluster-nodes-1	2%	66%	7.25%	6%
opensearch-cluster-nodes-1	4%	87%	6.06%	4%

Memory usage



Node	min	max	avg	current
opensearch-cluster-nodes-0	1.51 GiB	1.58 GiB	1.55 GiB	1.58 GiB
opensearch-cluster-nodes-0	1.50 GiB	1.58 GiB	1.55 GiB	1.58 GiB
opensearch-cluster-nodes-2	1.50 GiB	1.56 GiB	1.54 GiB	1.55 GiB
opensearch-cluster-nodes-2	1.50 GiB	1.56 GiB	1.54 GiB	1.55 GiB

Disk usage



Node	min	max	avg	current
opensearch-cluster-nodes-0 - /usr/share/opensearch/data/nodes/0	43.3%	44.3%	43.8%	
opensearch-cluster-nodes-0 - /usr/share/opensearch/data/nodes/0	43.3%	44.3%	43.8%	
opensearch-cluster-nodes-2 - /usr/share/opensearch/data/nodes/0	46.4%	46.7%	46.5%	
opensearch-cluster-nodes-2 - /usr/share/opensearch/data/nodes/0	46.4%	46.7%	46.5%	

Documents and Latencies

Documents indexing rate



Node	min	max	avg	current
opensearch-cluster-nodes-0	0	28.6	0.232	0
opensearch-cluster-nodes-0	0	13.7	0.0783	0
opensearch-cluster-nodes-2	0	0	0	0

Indexing latency



Node	min	max	avg	current
opensearch-cluster-nodes-0	0.00117	0.00216	0.00166	
opensearch-cluster-nodes-0	0.00216	0.00216	0.00216	
opensearch-cluster-nodes-2				

Search rate



Node	min	max	avg	current
opensearch-cluster-nodes-0	0	0.133	0.00106	0
opensearch-cluster-nodes-0	0	0.0667	0.000720	0
opensearch-cluster-nodes-2	0	0	0	0

Search latency




Node	min	max	avg	current
opensearch-cluster-nodes-0	0.0430	0.0509	0.0470	
opensearch-cluster-nodes-0	0	0.0430	0.0215	
opensearch-cluster-nodes-2				




serviceMonitor/default/opensearch-service-monitor/0 (6/6 up) [show less](#)

Endpoint	State	Labels	Last Scrape	Scrape Duration	Error
<a href="https://10.100.213.143:9200/_prometheus/metrics">https://10.100.213.143:9200/_prometheus/metrics</a>	UP	container="opensearch" endpoint="http" instance="10.100.213.143:9200" job="opensearch-service" namespace="default" pod="opensearch-cluster-nodes-1" service="opensearch-service"	31.943s ago	29.700ms	
<a href="https://10.100.213.143:9200/_prometheus/metrics">https://10.100.213.143:9200/_prometheus/metrics</a>	UP	container="opensearch" endpoint="http" instance="10.100.213.143:9200" job="opensearch-service-nodes" namespace="default" pod="opensearch-cluster-nodes-1" service="opensearch-service-nodes"	35.20s ago	27.048ms	
<a href="https://10.100.127.211:9200/_prometheus/metrics">https://10.100.127.211:9200/_prometheus/metrics</a>	UP	container="opensearch" endpoint="http" instance="10.100.127.211:9200" job="opensearch-service" namespace="default" pod="opensearch-cluster-nodes-2" service="opensearch-service"	35.89s ago	28.356ms	
<a href="https://10.100.127.211:9200/_prometheus/metrics">https://10.100.127.211:9200/_prometheus/metrics</a>	UP	container="opensearch" endpoint="http" instance="10.100.127.211:9200" job="opensearch-service-nodes" namespace="default" pod="opensearch-cluster-nodes-2" service="opensearch-service-nodes"	8.213s ago	32.425ms	
<a href="https://10.100.12.15:9200/_prometheus/metrics">https://10.100.12.15:9200/_prometheus/metrics</a>	UP	container="opensearch" endpoint="http" instance="10.100.12.15:9200" job="opensearch-service-nodes" namespace="default" pod="opensearch-cluster-nodes-0" service="opensearch-service-nodes"	5.657s ago	25.709ms	
<a href="https://10.100.12.15:9200/_prometheus/metrics">https://10.100.12.15:9200/_prometheus/metrics</a>	UP	container="opensearch" endpoint="http" instance="10.100.12.15:9200" job="opensearch-service" namespace="default" pod="opensearch-cluster-nodes-0" service="opensearch-service"	28.594s ago	29.732ms	

# Project work | Benchmark



# Project work | Benchmark



# Summary

- **Kubernetes** offers better flexibility compared with the current **Puppet** deployment
- **OpenSearch** clusters using **io2 and io3** show good performance
- **OpenSearch** operator eases common operations (scaling, failures, upgrades)
- **ArgoCD** provides an easy monitored way of configuring **OpenSearch** clusters in K8s
- **K8s** deployment provides easy integration with **Prometheus**, thus aligning with **MONIT** service

**Thank you for  
your attention**