

---

# Monitoring Kubernetes with OMD Labs Edition and Prometheus

Michael Kraus - FOSDEM 2017

---

—

# About me

# About me

Michael Kraus



Doing monitoring for 12 years,  
mainly with plain old Nagios,  
open-source only.

Senior Monitoring Consultant  
@ ConSol.

---

---

# Background

# Why

Kubernetes in a  
classical enterprise

Implementation of Kubernetes  
PoC at \$customer:

We have ...

- already running some monitoring instances there.
- but no idea about monitoring Kubernetes.

---

With

Enter Prometheus

Natural choice for kubernetes monitoring:

- Integrated service discovery
- Labels are retained between Kubernetes and Prometheus

---

# How

Where to start

There are excellent tutorials and blog posts available as a starting point, for example by

- [coreos.com/blog/](https://coreos.com/blog/)  
( Fabian Reinartz )
- [robustperception.io/blog/](https://robustperception.io/blog/)  
( Brian Brazil )
- ... many examples on GitHub

---

# Implementation



# Implementation


Prometheus  
kubernetes\_sd

- kubernetes\_sd\_configs
  - role: endpoints
- kubernetes\_sd\_configs
  - role: node
- kubernetes\_sd\_configs
  - role: pod

prometheus-kubernetes.yml  
from prometheus/examples.

---

# kubernetes



# Implementation

Prometheus  
kubernetes\_sd

Metrics:

- apiserver\_\*
- container\_cpu\_\*
- container\_fs\_\*
- deployment\_\*
- etcd\_\*
- kubelet\_\*
- ...

---

# Implementation


node\_exporter

Prometheus exporter for hardware and OS metrics exposed by the kernel.

- DaemonSet
- `prometheus.io/scrape: 'true'`

---

# kubernetes



# Implementation

node\_exporter

Metrics:

- node\_cpu
- node\_disk\_\*
- node\_filesystem\_\*
- node\_netstat\_\*
- node\_vmstat\_\*
- ...

---

# Implementation


kube-state-metrics

“... focused ... on the health of the various objects inside, such as deployments, nodes and pods.”

- `prometheus.io/scrape: 'true'`

---

# kubernetes





# Implementation

kube-state-metrics

Metrics:

- kube\_deployment\_\*
- kube\_node\_\*
- kube\_pod\_\*
- kube\_resource\_quota
- ...

---

# Implementation

Demo environment

Based on minikube:

[github.com/  
kubernetes/minikube](https://github.com/kubernetes/minikube)

Sample config:

[github.com/  
m-kraus/kubernetes-monitoring](https://github.com/m-kraus/kubernetes-monitoring)

---

—

# Demo

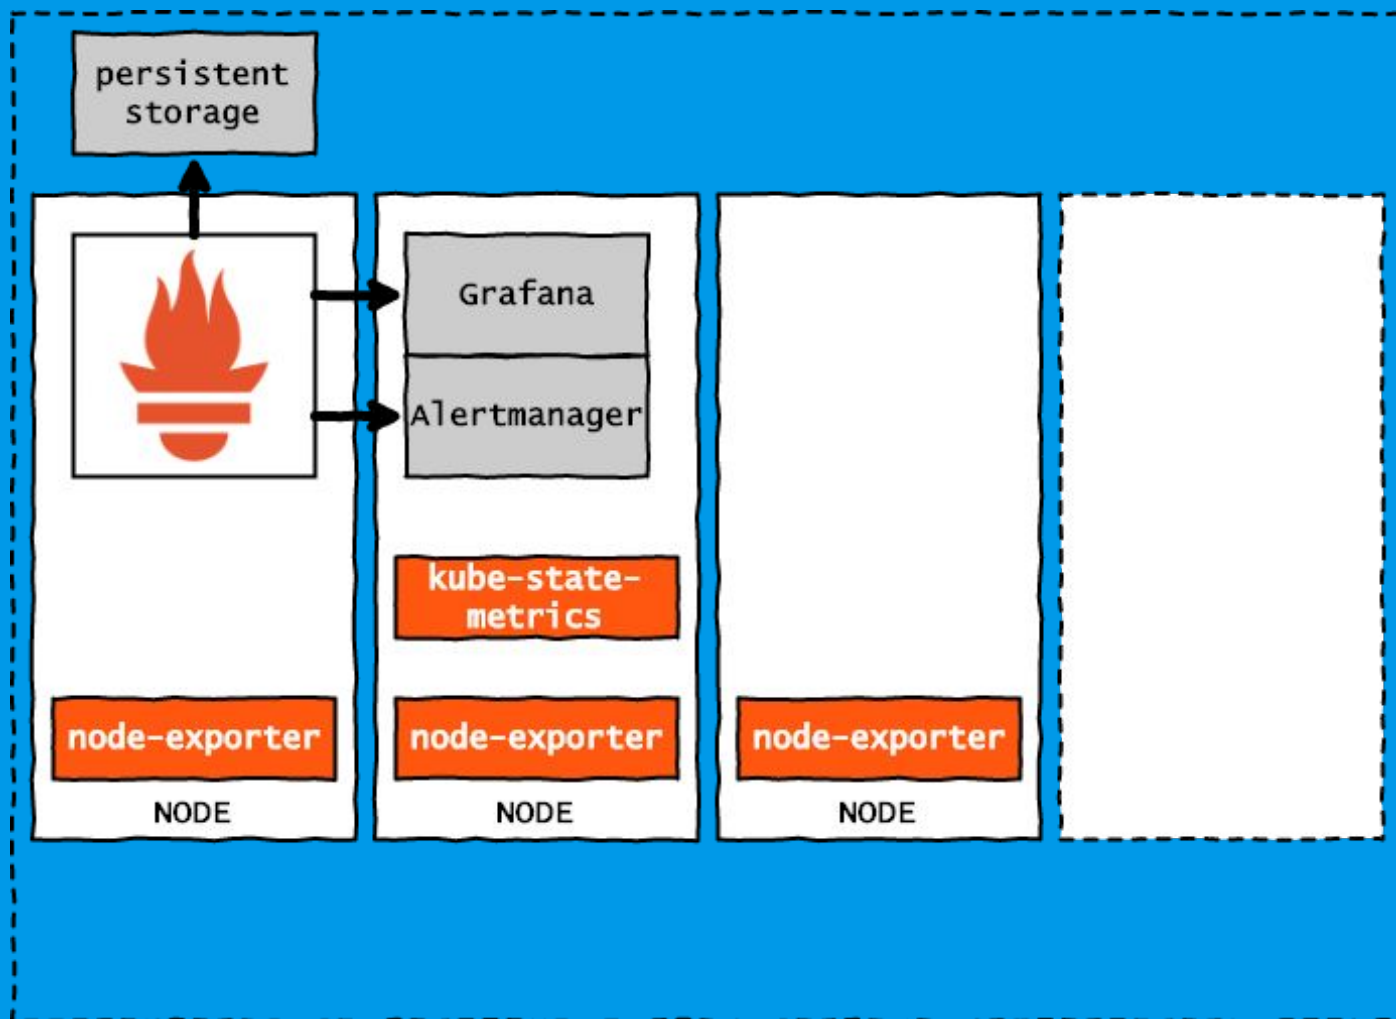
# Implementation

What else?

What we also need:

- persistent storage
- Alertmanager
- Grafana
- Pushgateway
- ...

---



---

**But we have that  
already**

# Classical monitoring

OMD Labs Edition



Monitoring in one package.

- completely open-source
- based on Nagios / Icinga
- bundles “best practices” of many years of experience
- no root required

"Musterlösung" at \$customer  
for monitoring projects:

---

# Classical monitoring

OMD Labs Edition



**Nagios** Icinga1 **Icinga2** Shinken  
Naemon **Thruk** Mod-Gearman  
LMD NagVis **PNP4Nagios**  
Apache MySQL **InfluxDB** Nagflux  
**Prometheus** Dokuwiki  
**Grafana** FreeTDS JMX4Perl  
check\_webinject check\_logfiles  
**Jolokia** check\_mysql\_health  
coshsh check\_mssql\_health rrdcache  
check\_nsc\_web check\_curly  
check\_nwc\_health check\_multi  
check\_oracle\_health **Ansible**



# Classical monitoring

OMD sites and commands



```
omd create <MYSITE>
```

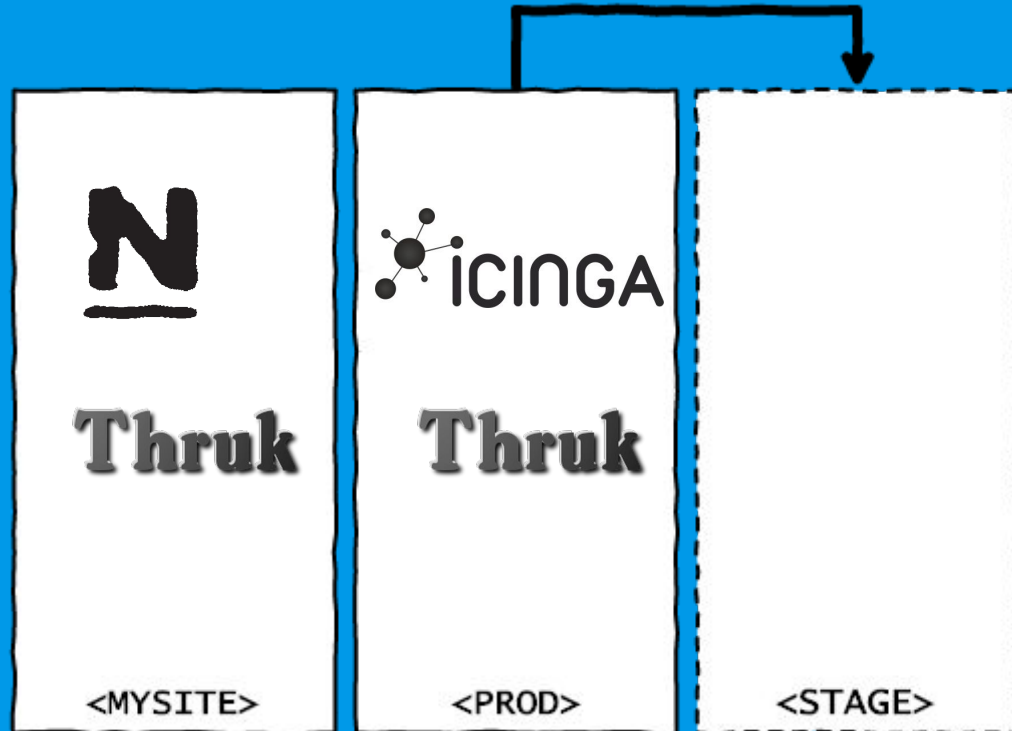
```
omd cp <PROD> <STAGE>
```

```
omd update <STAGE>
```

```
omd version
```

---

```
omd create <MYSITE>  
omd cp <PROD> <STAGE>
```



# Classical monitoring

OMD Labs Edition



<https://labs.consol.de/omd/>

---

# Implementation

Connecting OMD

Why not scrape Kubernetes directly from OMD:

- hard to access pods inside Kubernetes
  - hard to access API from outside Kubernetes
  - API secured via TLS and token only (easily) available from a serviceaccount
-


# Implementation

Connecting OMD


Getting the metrics from  
Kubernetes to OMD:

- federation
    - job\_name: 'kube\_federation'  
metrics\_path: '/federate'  
honor\_labels: true  
params:  
  'match[]':  
  - '{job=~"^kubernetes.+"}'
-

# OMD



# kubernetes



---

# Demo

---

# Issues



# Issues

## Federation

“... Not quite the purpose of federation.”

*Brian Brazil*

*[www.robustperception.io/  
federation-what-is-it-good-for/](http://www.robustperception.io/federation-what-is-it-good-for/)*

- Let's try it anyway ...
-

# Issues

## Securing

"Accessing metrics without authentication is ok for a PoC, but not allowed in production..."

*internal audit*

- How to secure (federated) Prometheus?

---

# Issues

## Integration

"Should Nagios, Alertmanager or both notify?"

"Do we need to define our checks and alerts both, in Nagios and Prometheus?"

- How to route alerts
  - How to ease or centralize configuration
-

# Issues

Long-term storage

“How can we store (some) of our graphs for a longer period of time?”

- InfluxDB ?

---

# Issues

## Coverage

“Our kubernetes cluster died.  
We had no monitoring until it  
up again...”

*operations team*

- external monitoring of crucial components
    - machine health
    - important services
    - important API queries
-

---

**Thanks for watching**