

# Let's Embrace Green CloudNative Engineering



# Henrik Rexed

---



## CNCF Ambassador & Cloud Native Advocate

- 15+ years of Performance engineering


- Owner of



- Producer of

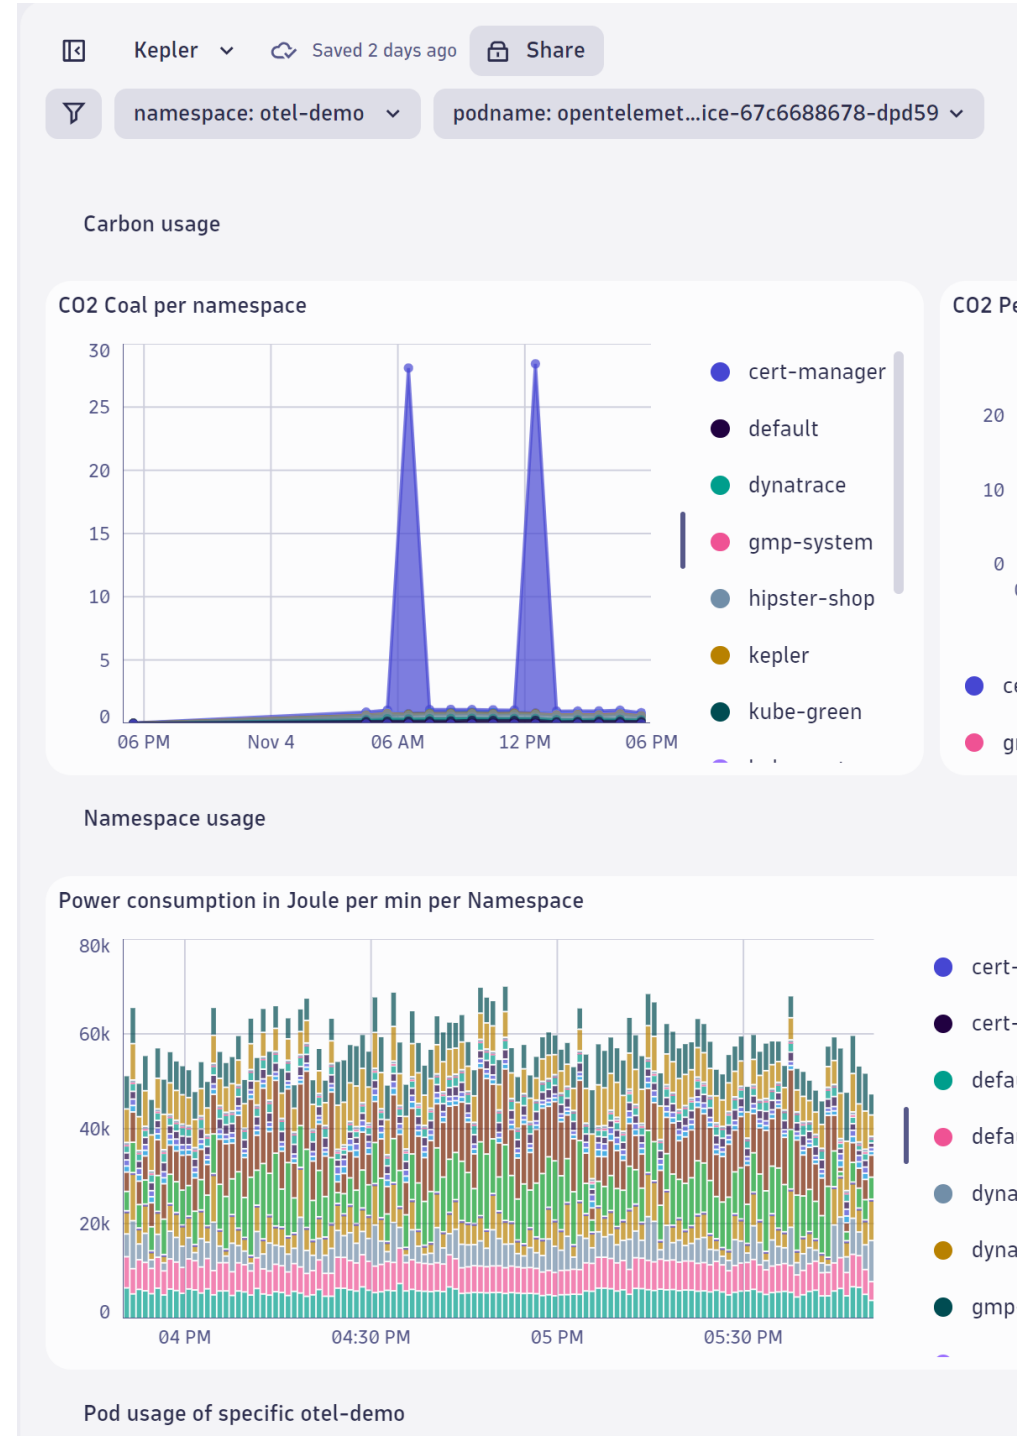


# Kaptain Green




# If you stay with me you will ...

- Learn about the **importance** of reducing the energy footprint of **your k8s workload**
- The **impact of not optimized** resource allocation
- The various way of saving energy
- Finaly we will present :
  - Kepler
  - KubeGreen






We need to save our planet



# Everything started with

## Carbon footprint

- Datacenter
- road & transportation
- building
- petroleum refining
- other energy
- land use change & forestry
- fuel production
- waste
- electricity & heat
- industry
- non road transportation



Number shared from Kaptain Green

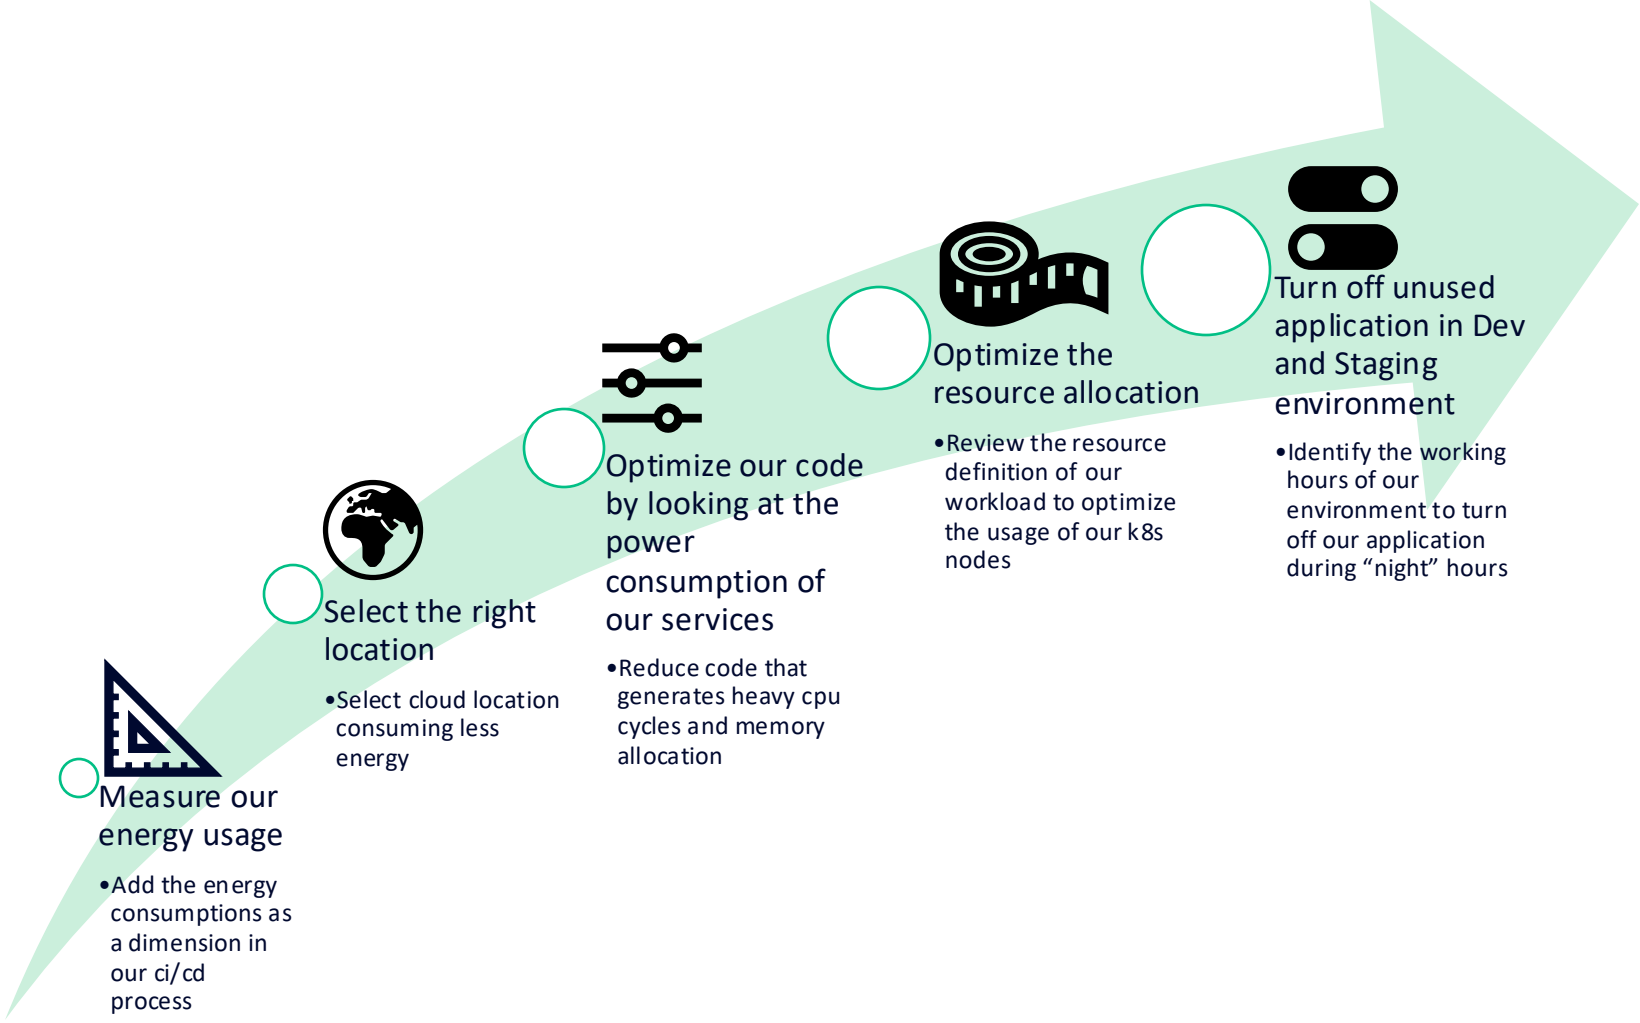


## But the usage of datacenter kept increasing

- Every industry relies on datacenter to manage :
  - software managing sales...and more
  - The usage of AI is increasing
  - The explosion of data required for any AI applications
- The energy required for our datacenter is increasing by 1-2% per year.
- The energy usage of our datacenter should reach 15 by 2040



# How can we avoid the disaster?



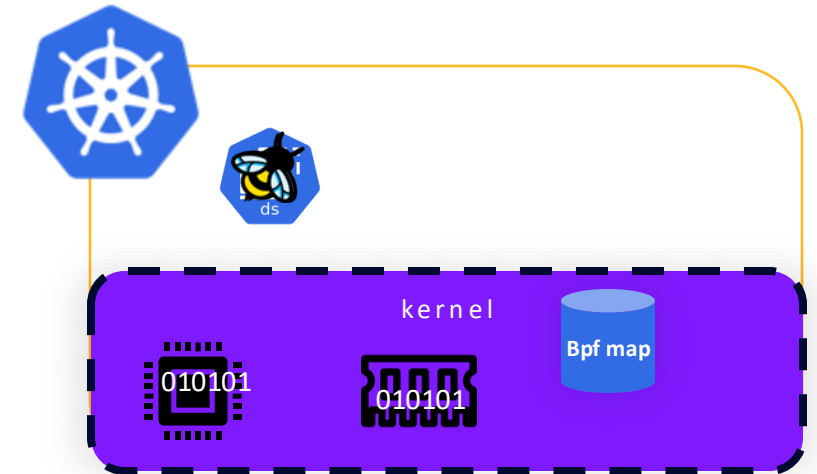


How can we measure ?

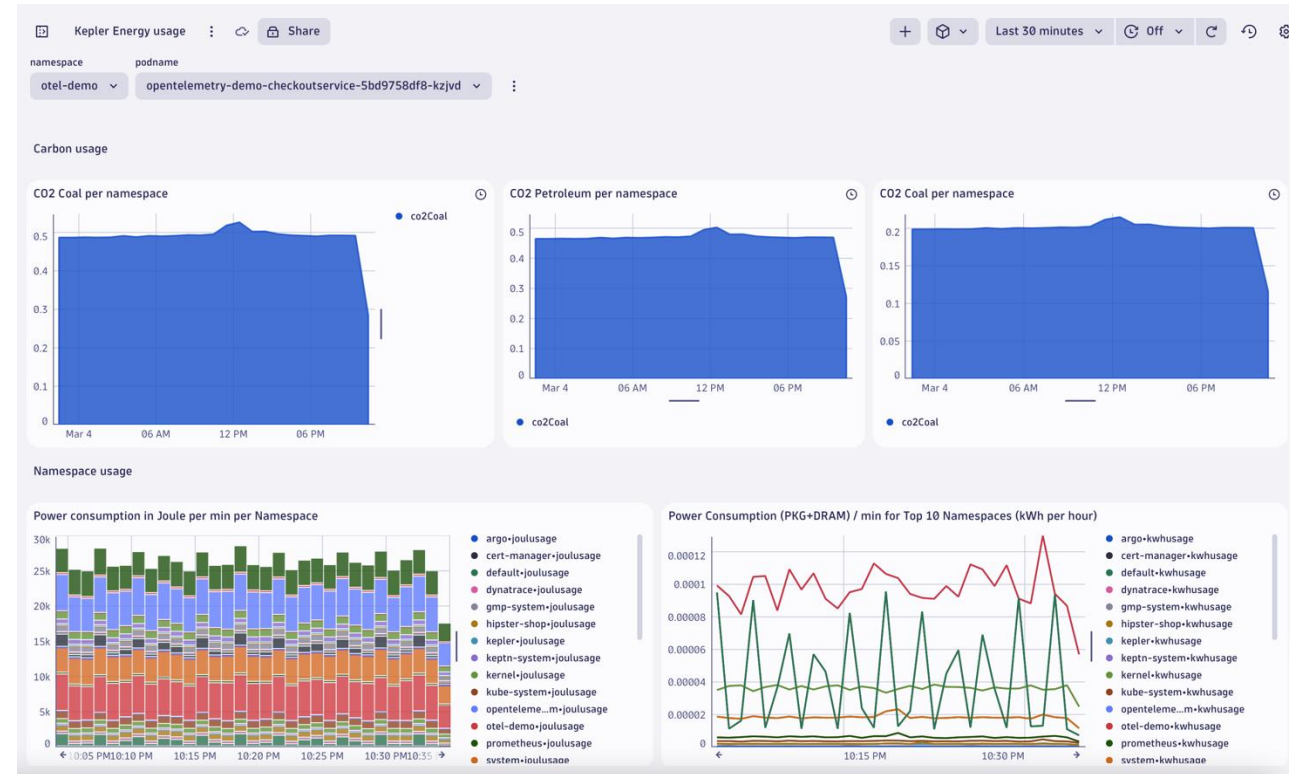


# Kepler

- Kepler utilize ebpf programs to estimate the power consumption of our workload
- Kepler will estimate the power usage by measuring :
  - the number of cpu cycles required for a process
  - The memory allocation requested by a process
- It adds the k8s metadata to the generated Prometheus metrics

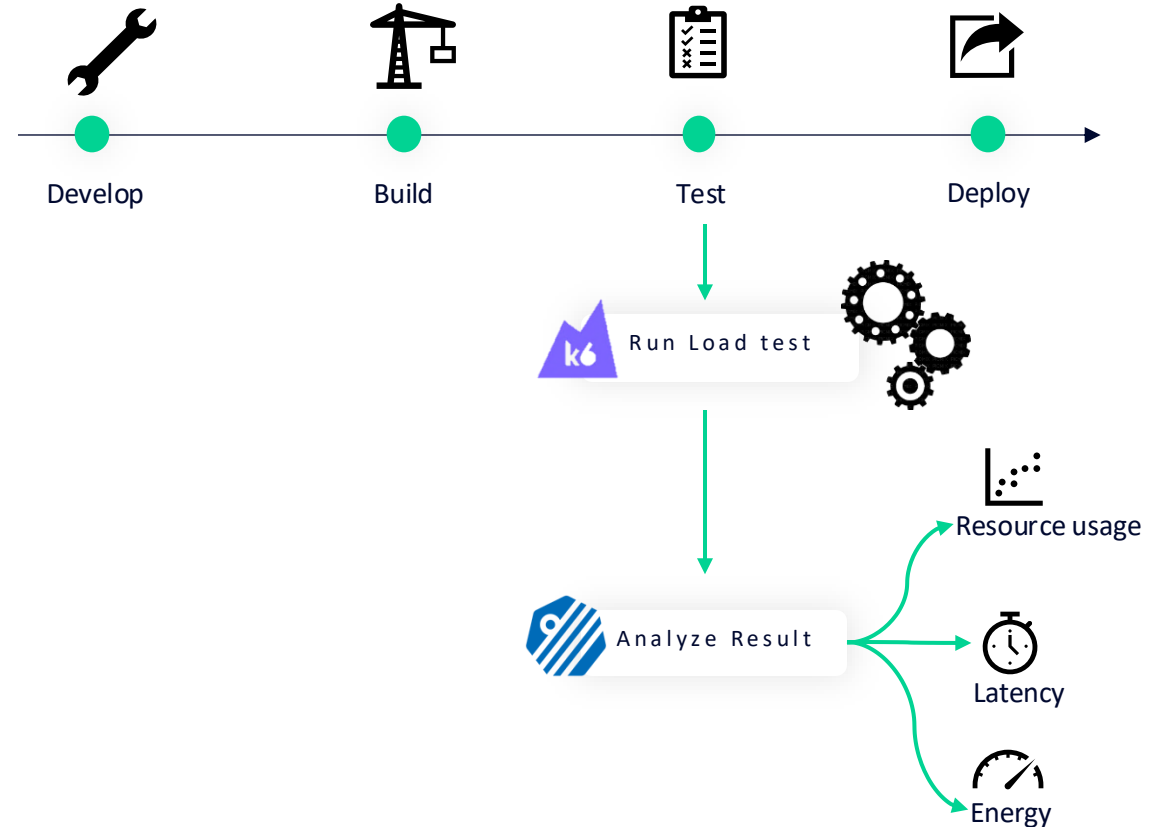


# Tools required for our journey



# Include Energy Consumption into your pipeline

- Energy usage should be considered to be one of our key kpi to evaluate a successful pipeline
- You can use any CI/CD system of your choice, but need a solution helping you to evaluate based on metrics
- Keptn project provides k8s CRD to build your evaluation based on external metrics.





# Keptn Analyse Defintion




# SERVICE LEVEL INDICATOR

METRIC



# SERVICE LEVEL OBJECTIVE

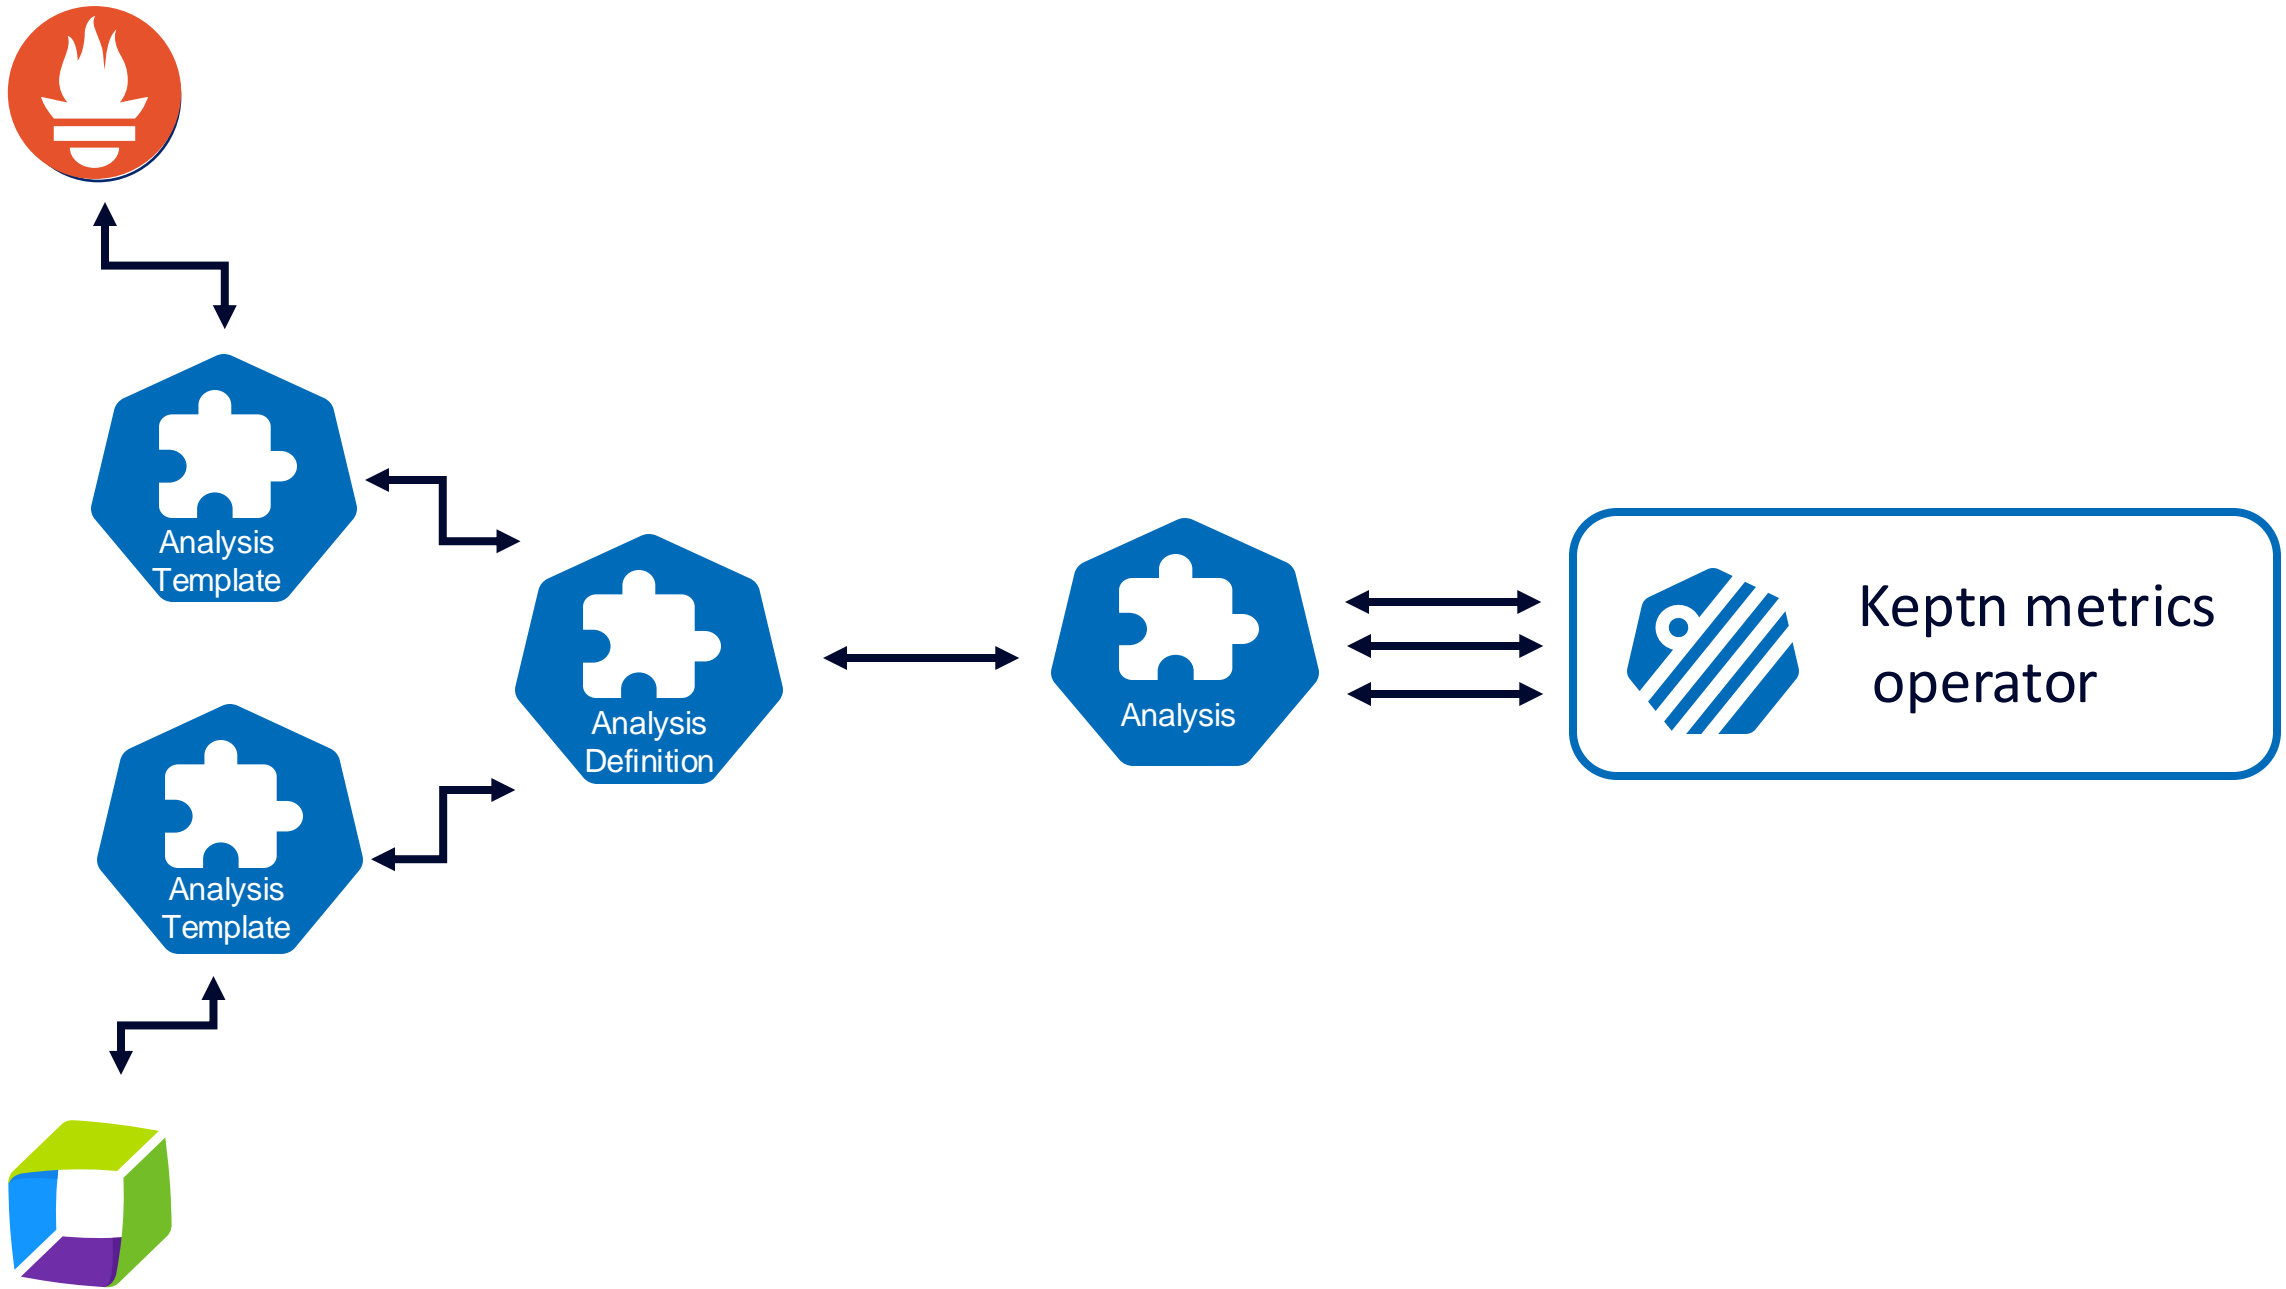
GOAL



# SERVICE LEVEL AGREEMENT

CONTRACT





```
apiVersion: metrics.keptn.sh/v1alpha3
```

```
kind: AnalysisValueTemplate
```

```
metadata:
```

```
  name: value-sample
```

```
spec:
```

```
  provider:
```

```
    name: my-provider
```

```
  query:
```

```
'(sum(kepler_container_joules_total{pod  
_name=~"{{.pod}}.*"})/sum(kepler_conta  
iner_joules_total{container_namespace=  
"{{.ns}}"})))*100'
```

# SERVICE LEVEL INDICATOR

METRIC





apiVersion: metrics.keptn.sh/v1alpha3

kind: AnalysisDefinition

metadata:

name: definition-sample

spec:

objectives:

- analysisValueTemplateRef:

name: value-sample

target:

failure:

lessThan:

fixedValue: 15

weight: 2

keyObjective: false


totalScore:

passPercentage: 75

warningPercentage: 50

# SERVICE LEVEL OBJECTIVE

GOAL



apiVersion: metrics.keptn.sh/v1alpha3

kind: Analysis

metadata:

name: analysis-sample

spec:

timeframe:

recent: 5m

args:

ns: „otel-demo„

pod:“opentelemetry-demo-checkoutservice“

analysisDefinition:

name: definition-sample

status:

RESULTS GO HERE

# SERVICE LEVEL AGREEMENT

CONTRACT




# Demo



# Our Technical Environment



**Kepker**



**Keptn**



**Argo**



**prometheus**



**hipstershop**



**Istio**

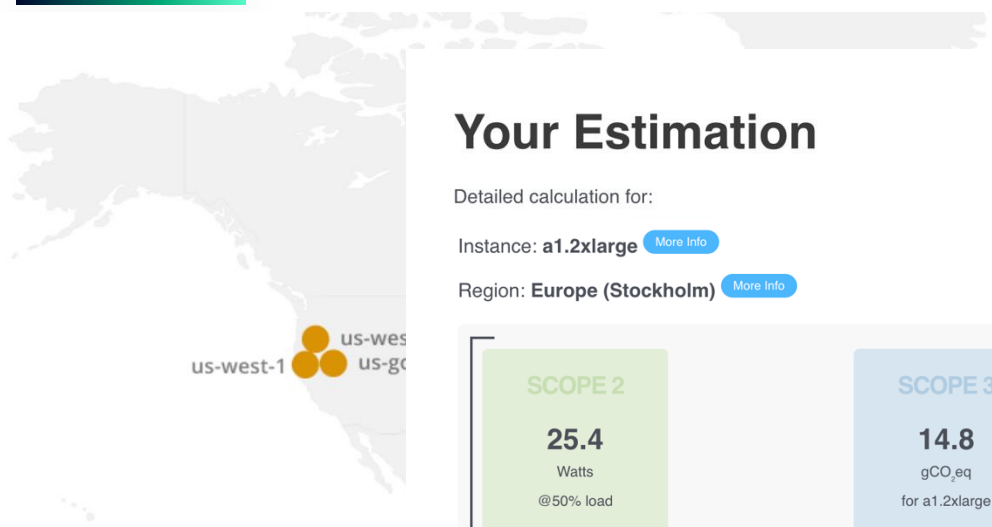




# Move your workload To sweden



# All locations are not equal...




## Your Estimation


Detailed calculation for:

Instance: **a1.2xlarge** [More Info](#)

Region: **Europe (Stockholm)** [More Info](#)




# How to optimize our resource allocations?




# Resize your Workload Size

- Node resource allocation is based on our workload definition



- So, we need reduce our Request/Limit

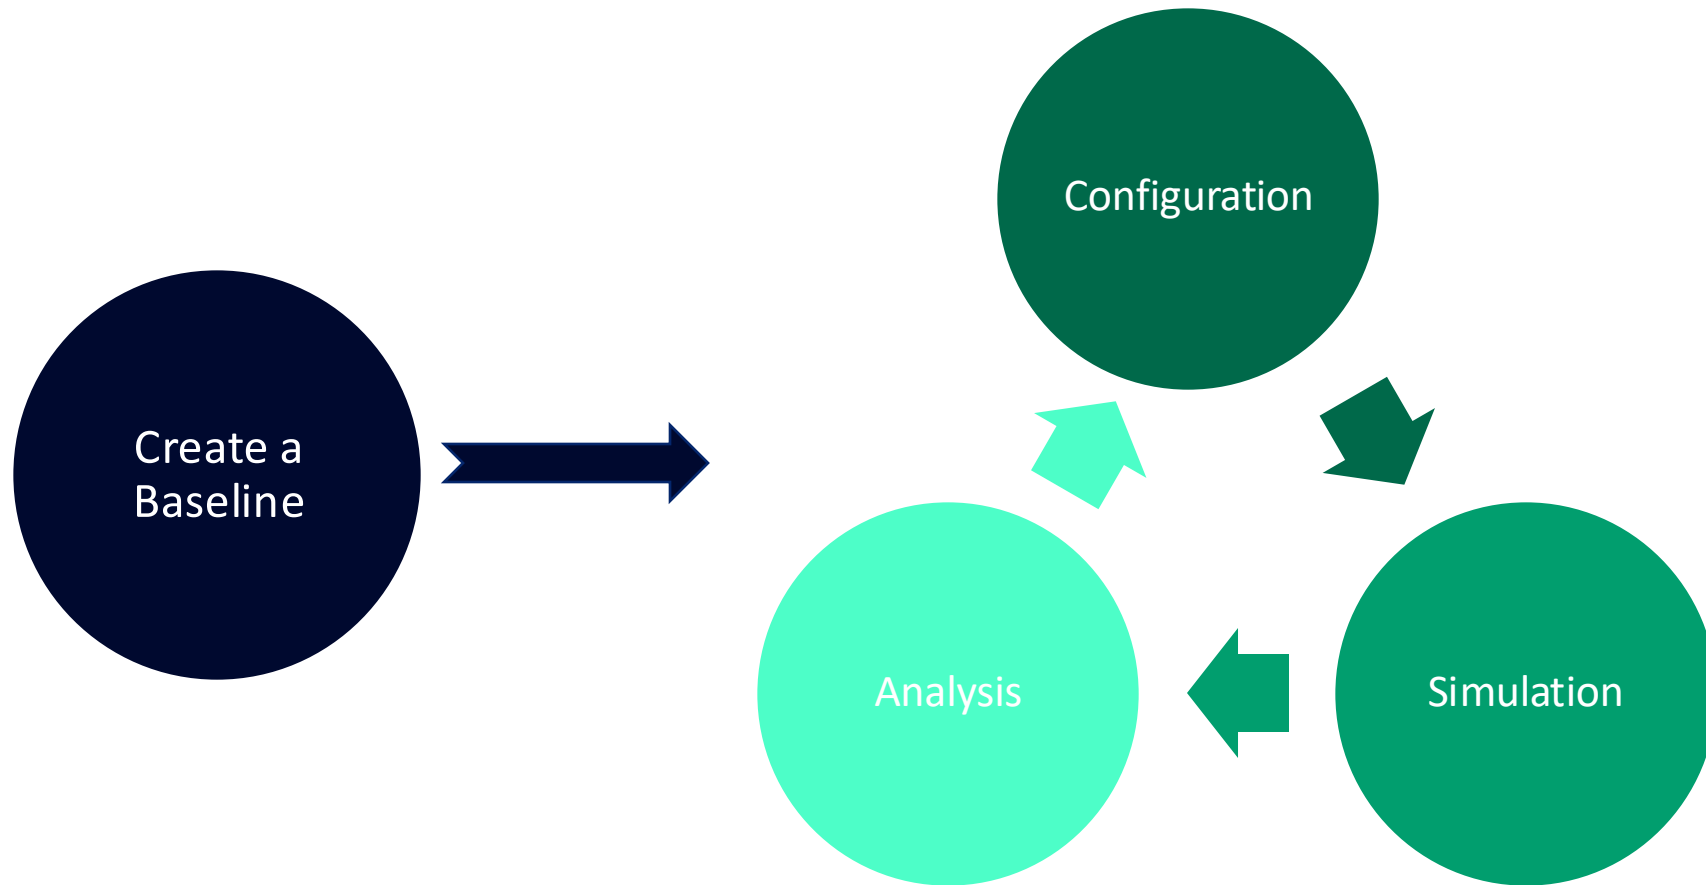
- Look at the usage vs request
- Look at the usage vs limit




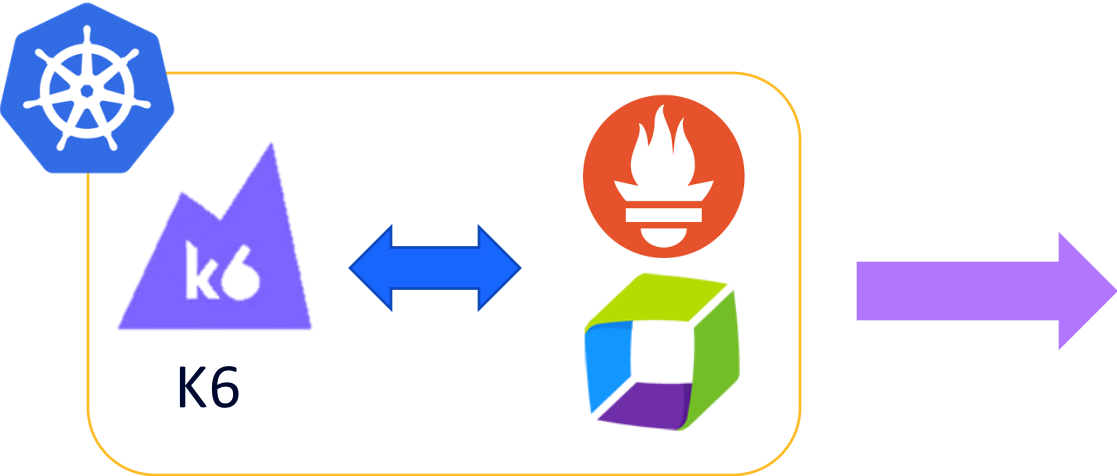


# Optimize : What is the approach?

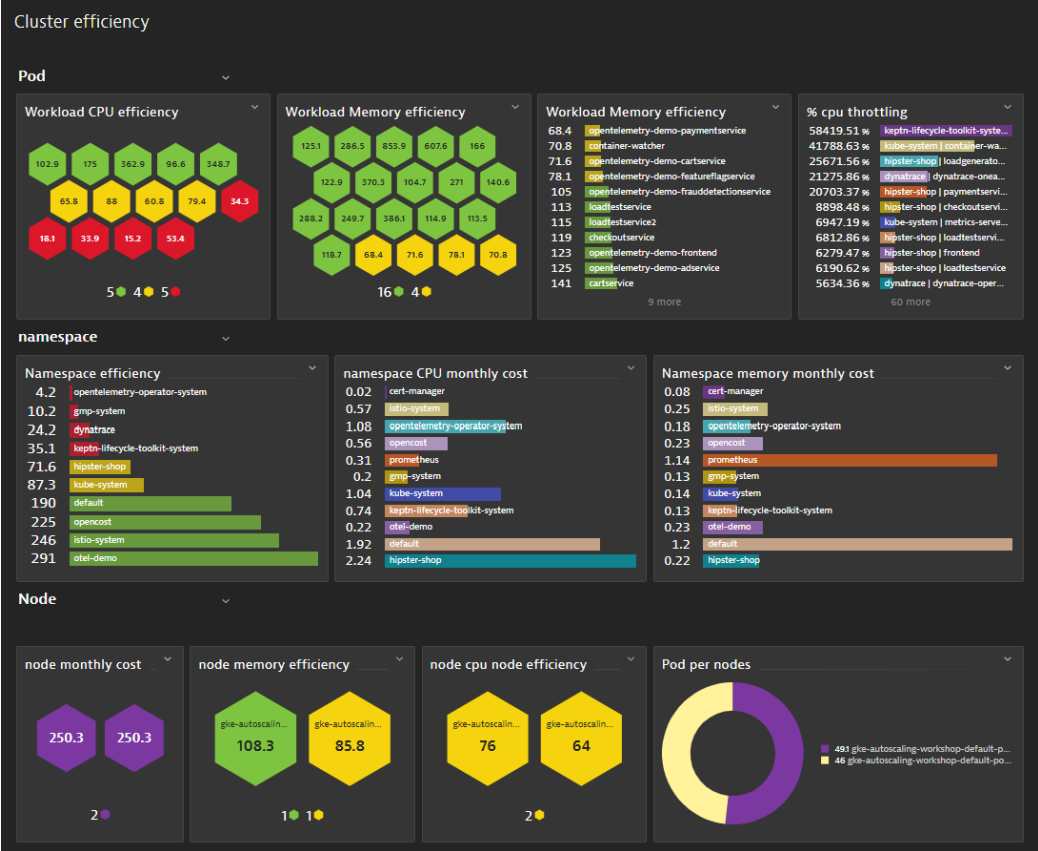
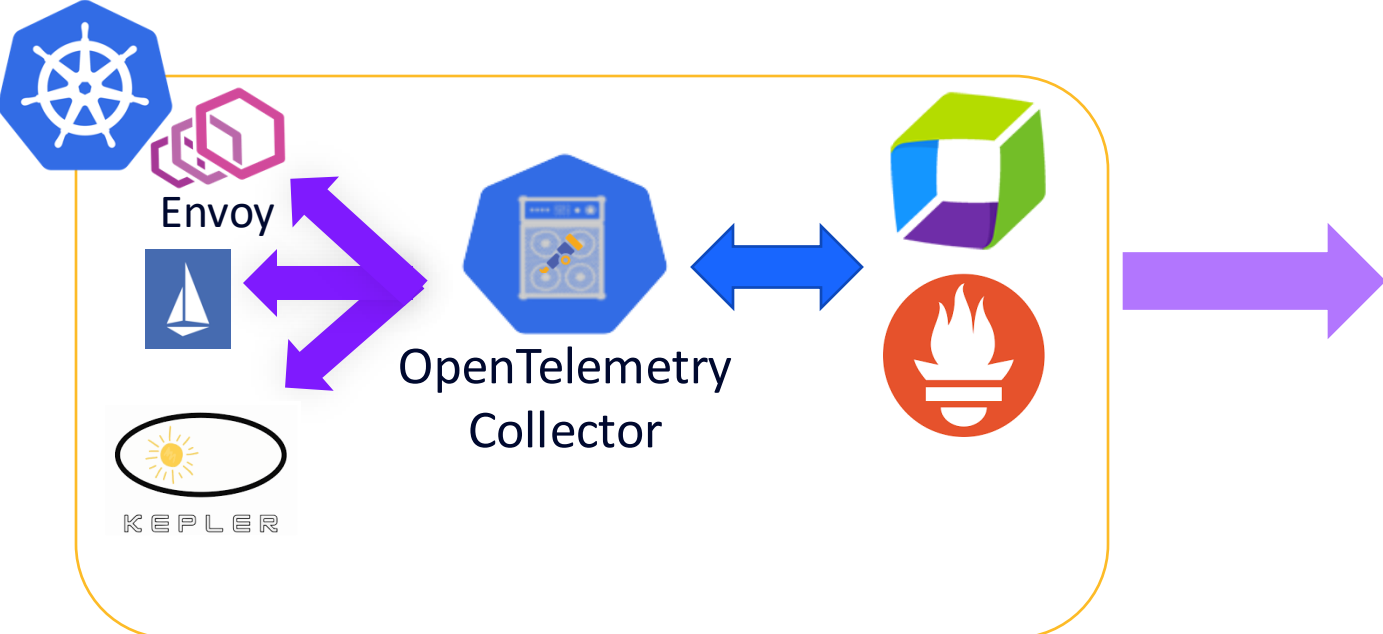
---



# K6 – The Dynatrace integration



# Kepler to report the energy usage of the workload



# How to remove unuseful workload



# Our environments


---



Dev




Testing



Production

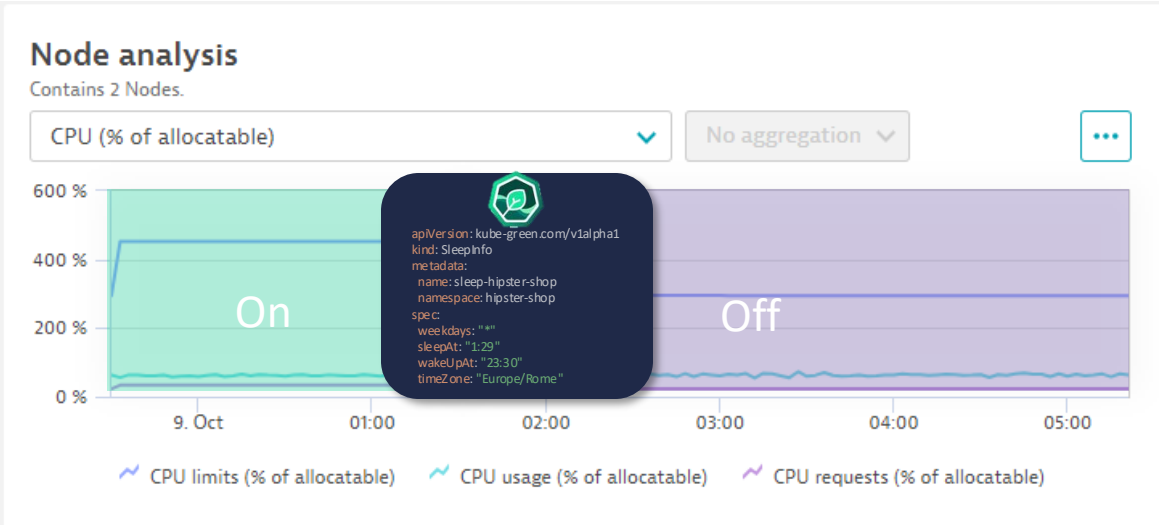

# Testing Environment

- There are always several hours where our non production environment are not handling any traffic.
- We should remove the workload when they are not required.
- The project KubeGreen is designed for that





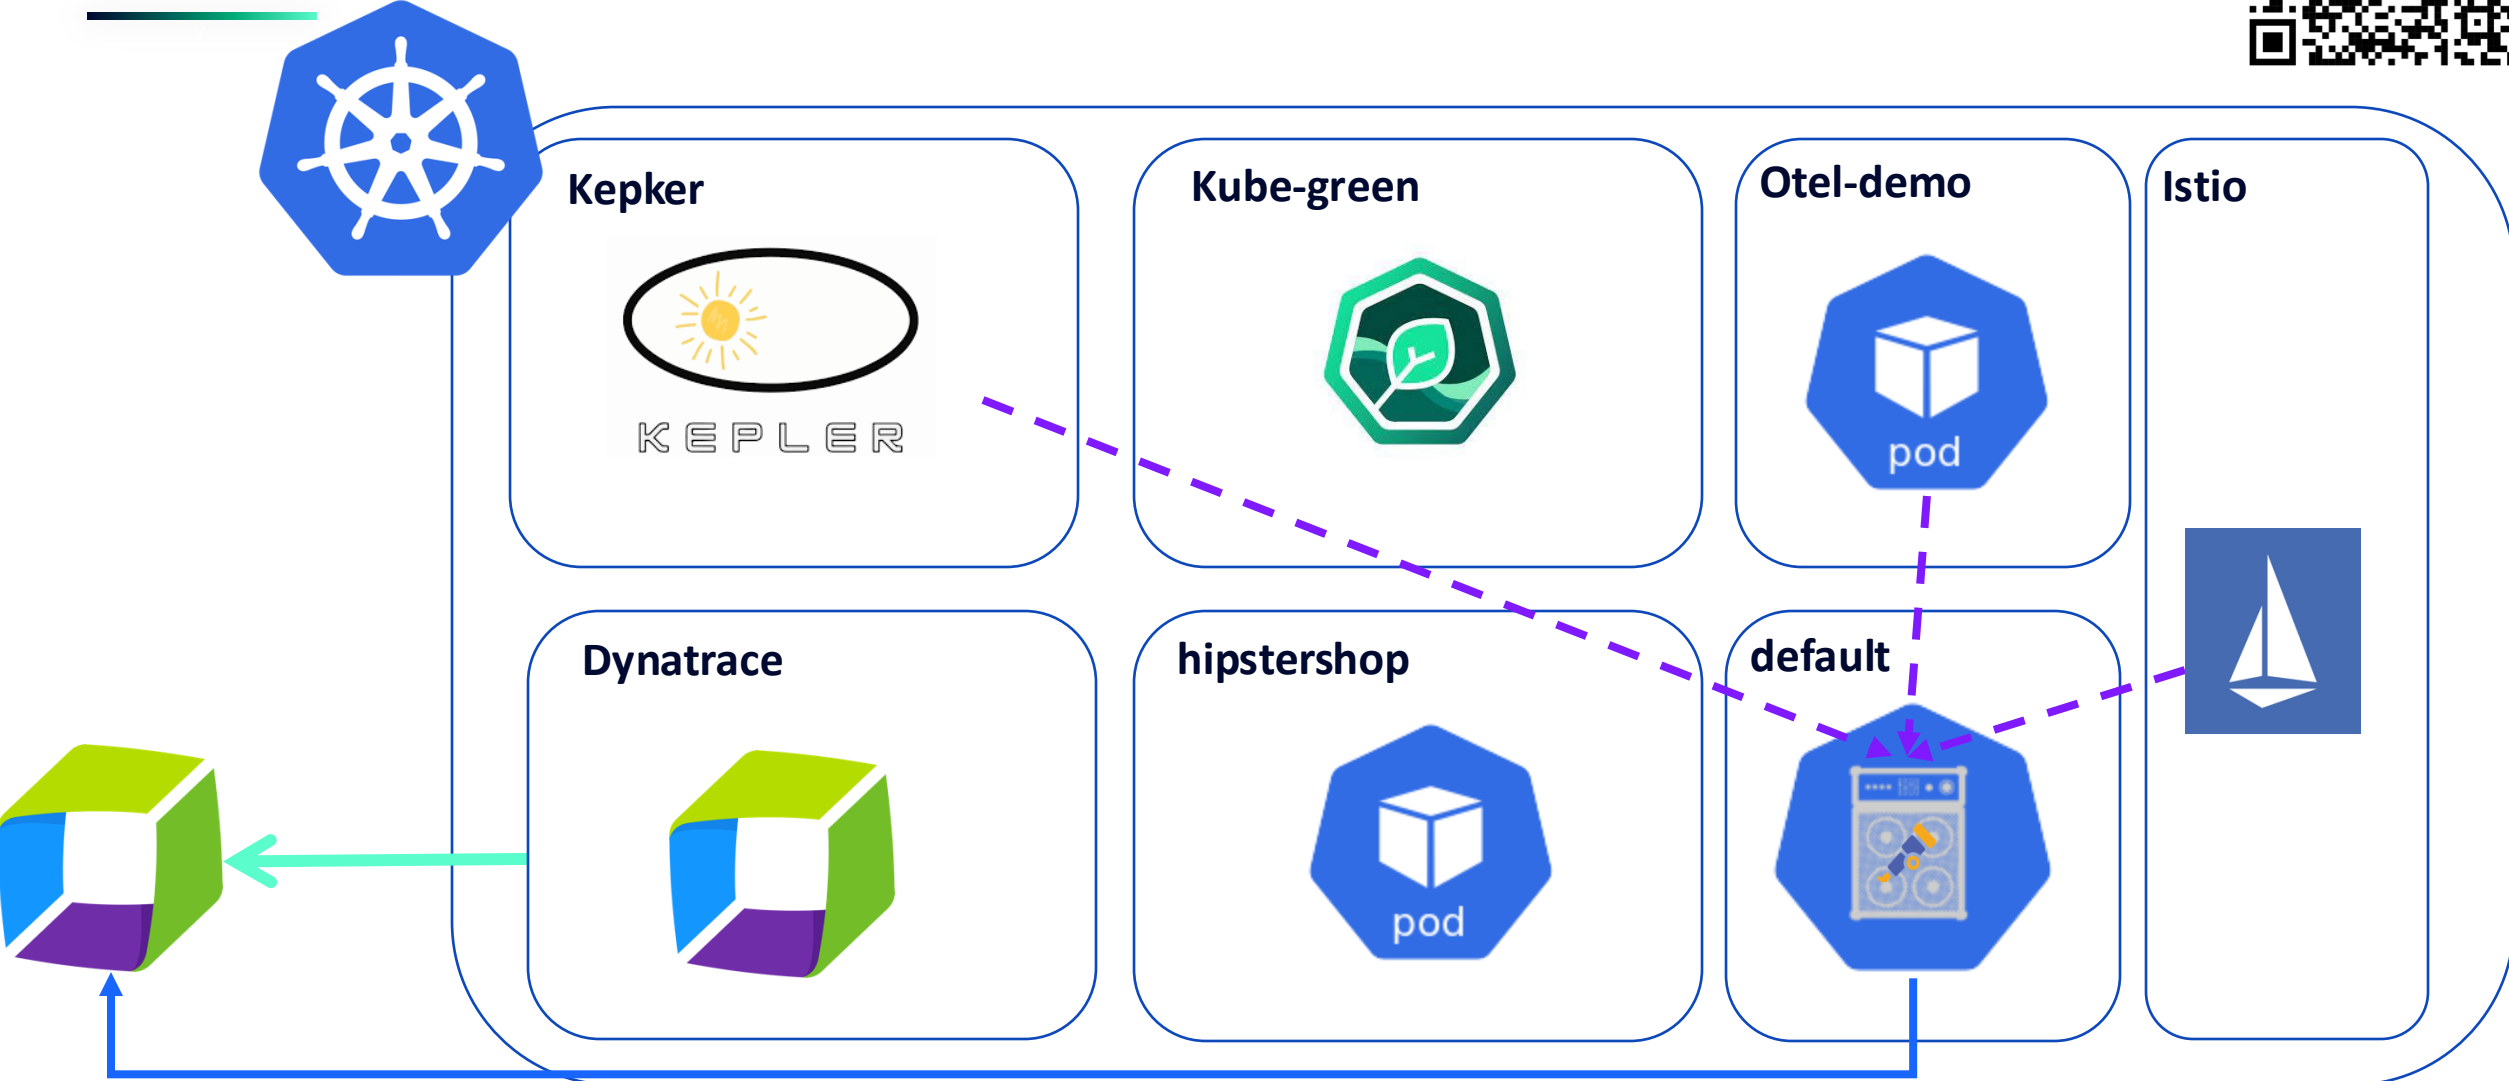
# KubeGreen



# Demo



# Our Technical Environment



# Take Away

---

## Measure

- Deploy Kepler to report energy usage

## CI/CD

- Include energy consumption in your CI/CD pipeline
- Create QualityQuates to automate the feedback loop

## Optimize

- Use Profiling solutions to identify code consuming CPU
- Guide your optimization with Kepler metrics

## Sleep

- Use Kubegreen to remove un-useful workload
- Make your application sleep in Testing environment

# Is it observable

- Looking for educational content on Observability , Checkout the YouTube Channel :

## Is It Observable





**Thank you**