



KubeCon




CloudNativeCon

Europe 2018

Federated Prometheus Monitoring at Scale

LungChih Tung, Software Engineer
Nandhakumar Venkatachalam, Princ Production Engineer



Team



KubeCon




CloudNativeCon

Europe 2018

- Core Infrastructure
 - Infrastructure team powering all Yahoo Media Products

Yahoo Media Products



Journey to Kubernetes



KubeCon





CloudNativeCon

Europe 2018

STAGE colo1	ETCD	Master	Nodes			
----------------	------	--------	-------	--	--	--

CANARY colo1	ETCD	Master	Nodes			
-----------------	------	--------	-------	--	--	--

PROD colo1	ETCD	Master	Nodes			
---------------	------	--------	-------	--	--	--



Monitoring Solutions



KubeCon



CloudNativeCon

Europe 2018

- Inhouse monitoring to get system's health
- Missing insight into kubernetes cluster metrics
- Heapster InFluxDB sink with grafana
- Prometheus 1.x
 - Remote write data into inhouse monitoring system

Heapster with InfluxDB




KubeCon



CloudNativeCon

Europe 2018



Growing Kubernetes Clusters

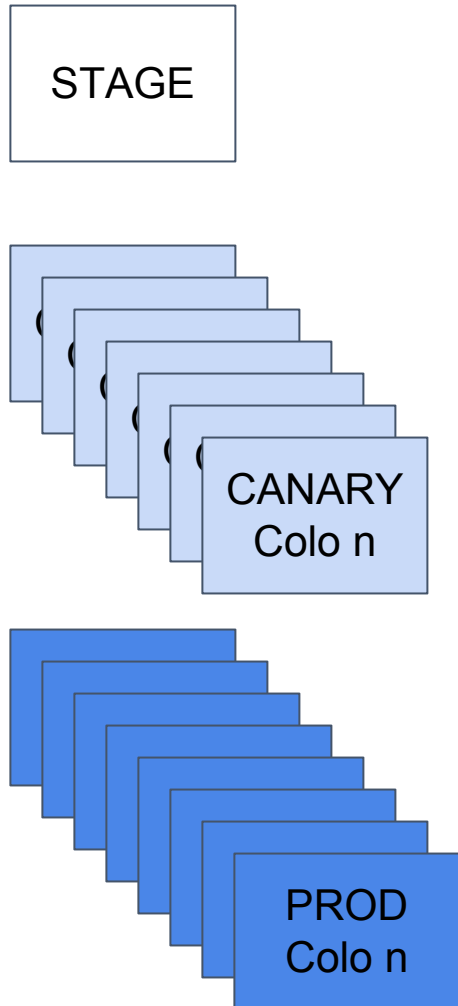


KubeCon



CloudNativeCon

Europe 2018



- Spanning across data centers
- Demand for higher visibility
- Growing monitoring requirement



Our Requirement



KubeCon



CloudNativeCon

Europe 2018

Cluster Health:

- ETCD
- Controller Manager
- Scheduler
- Kubernetes API server
- Kubelet
- Kubelet CAdvisor
- Kube DNS
- Any Add-ons...

Application Health:

- Namespace
- Deployment
- Pod
- Container

Prometheus 2.0 to the Rescue



KubeCon



CloudNativeCon

Europe 2018

- Huge performance improvement specifically in storage
- Simple syntax for aggregation and alerting rules
- Good documentation
- Our focus on pulling every metrics
- Aggregation rules

Metrics Collection



KubeCon



CloudNativeCon

Europe 2018

- Prometheus [example yaml](#)
- Endpoint discovery is simple.
 - annotation `prometheus.io/scrape=true`
- Simple File based discovery for apiservers



KubeCon



CloudNativeCon

Europe 2018

Metrics Collection - API Server

```
- job_name: kubernetes-apiservers
  scrape_interval: 30s
  scrape_timeout: 5s
  metrics_path: /metrics
  scheme: https
  file_sd_configs:
  - files:
    - /etc/prometheus/targets/apiserver-canary1bf1.json
    refresh_interval: 5m
  bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
  tls_config:
    ca_file: /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
    insecure_skip_verify: false
  relabel_configs:
  - separator: ;
    regex: (.*)
    target_label: colo
    replacement: bf1
    action: replace
```

Metrics Proxy



KubeCon



CloudNativeCon

Europe 2018

- ETCD ports are accessible only by master machines
- Scheduler and controller ports bind to 127.0.0.1

Metrics Proxy - ETCD



KubeCon



CloudNativeCon

Europe 2018

```
- job_name: kubernetes-etcd
  params:
    port:
      - "2379"
  scrape_interval: 30s
  scrape_timeout: 5s
  metrics_path: /metrics
  scheme: https
  file_sd_configs:
  - files:
    - /etc/prometheus/targets/etcd-canary1bf1.json
    refresh_interval: 5m
  bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
  tls_config:
    ca_file: /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
    insecure_skip_verify: false
  relabel_configs:
  - separator: ;
    regex: (.*)
    target_label: colo
    replacement: bf1
    action: replace
```

Metrics Proxy - Controller



KubeCon



CloudNativeCon

Europe 2018

```
- job_name: kubernetes-controller
  params:
    port:
      - "10252"
  scrape_interval: 30s
  scrape_timeout: 5s
  metrics_path: /metrics
  scheme: https
  file_sd_configs:
  - files:
    - /etc/prometheus/targets/controller-canary1bf1.json
    refresh_interval: 5m
  bearer_token_file: /var/run/secrets/kubernetes.io/serviceaccount/token
  tls_config:
    ca_file: /var/run/secrets/kubernetes.io/serviceaccount/ca.crt
    insecure_skip_verify: false
  relabel_configs:
  - separator: ;
    regex: (.*)
    target_label: colo
    replacement: bf1
    action: replace
```

Volume of the metrics



KubeCon



CloudNativeCon

Europe 2018




Prometheus Benchmark - 2.0.x




Zoom Out Last 24 hours

Prometheus localhost colo All IP All

Head Time series




Head Active Appenders




Samples Appended/s




Head Chunks



Head Chunks Created



Head Chunks Removed



Federation

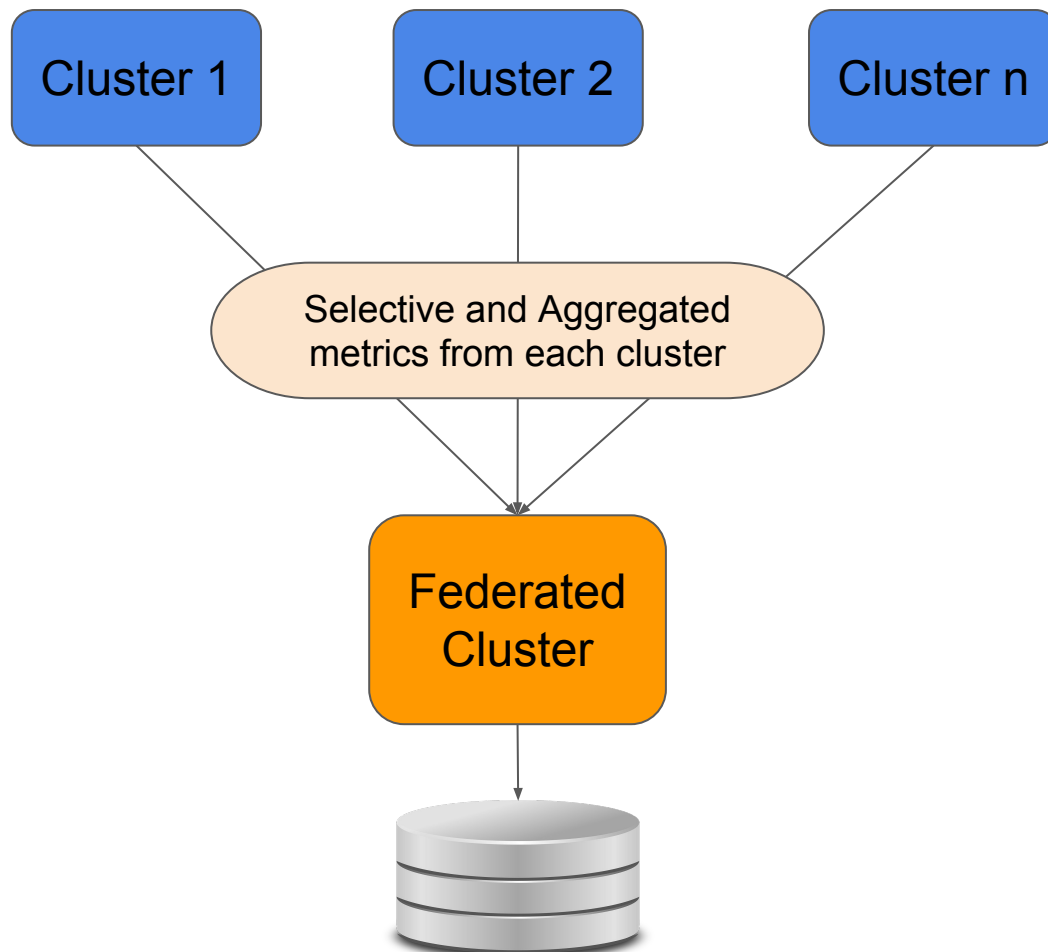


KubeCon



CloudNativeCon

Europe 2018



- Aggregated time series data
- Longer retention period
- Permanent storage
- Unified display of data

Federation Configuration



KubeCon



CloudNativeCon

Europe 2018

- Federate all control plane components
- Selective aggregated metrics

```
- job_name: 'federate'
```

```
honor_labels: true
```

```
metrics_path: '/federate'
```

```
params:
```

```
  'match[]':
```

```
    - '{job="prometheus"}'
```

```
    - '{app="kube-state-metrics"}'
```

```
    - '{job="kubernetes-etcd"}'
```

```
    - '{job="kube-dns"}'
```

```
    - '{job="kubernetes-scheduler"}'
```

```
    - '{job="kubernetes-controller"}'
```

```
    - '{job="kubernetes-nodes"}'
```

```
    - '{job="kubernetes-apiservers"}'
```

```
    - '{__name__=~"cluster.*"}'
```

```
    - '{__name__=~"colo.*"}'
```

```
    - '{kubernetes_name="heapster"}'
```

```
static_configs:
```

```
  - targets:
```


Aggregation Rules



KubeCon



CloudNativeCon

Europe 2018

- Ton of rules from prometheus operator team
- Built CPU and Memory utilization by
 - Colo(cluster) level
 - Namespace/ Deployment/ Pod and Container level

Aggregation - Colo level



KubeCon



CloudNativeCon

Europe 2018

colo level cpu utilization

- record: colo:cpu_percentage:rate

expr: 100 *

CPU usage of all containers per colo

```
sum(label_replace(irate(container_cpu_usage_seconds_total{container_name!="",  
container_name!="POD"}[5m]), "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY  
(colo)
```

/

```
sum(label_replace(container_spec_cpu_shares{container_name!="",  
container_name!="POD"}, "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY  
(colo) * 1000
```

CPU shares allocated to all containers per colo

Aggregation - Namespace level



KubeCon



CloudNativeCon

Europe 2018

namespace level cpu utilization

- record: colo_namespace:cpu_percentage:rate

expr: 100 *

CPU usage of all containers per namespace

```
sum(label_replace(irate(container_cpu_usage_seconds_total{container_name!="",  
container_name!="POD"}[5m]), "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY  
(colo, namespace)
```

/

```
sum(label_replace(container_spec_cpu_shares{container_name!="",  
container_name!="POD"}, "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY (colo,  
namespace) * 1000
```

CPU shares allocated to all containers per namespace

Aggregation - Controller level



KubeCon



CloudNativeCon

Europe 2018

controller level cpu utilization

- record: colo_namespace_controller:cpu_percentage:rate

expr: 100 *

CPU usage of all containers per controller

```
sum(label_replace(irate(container_cpu_usage_seconds_total{container_name!="",  
container_name!="POD"}[5m]), "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY  
(colo, namespace, controller)
```

/

```
sum(label_replace(container_spec_cpu_shares{container_name!="",  
container_name!="POD"}, "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY (colo,  
namespace, controller) * 1000
```

CPU shares allocated to all containers per controller

Aggregation - Pod level



KubeCon



CloudNativeCon

Europe 2018

pod level cpu utilization

- record: colo_namespace_controller_pod:cpu_percentage:rate

expr: 100 *

CPU usage of all containers per pod

```
sum(label_replace(irate(container_cpu_usage_seconds_total{container_name!="",
container_name!="POD"}[5m]), "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY
(colo, namespace, controller, pod_name)
```

/

```
sum(label_replace(container_spec_cpu_shares{container_name!="",
container_name!="POD"}, "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY
(colo, namespace, controller, pod_name) * 1000
```

CPU shares allocated to all containers per pod

Aggregation - Container level



KubeCon



CloudNativeCon

Europe 2018

container level cpu utilization

- record: colo_namespace_controller_pod_container:cpu_percentage:rate

expr: 100 *

CPU usage per container

```
sum(label_replace(irate(container_cpu_usage_seconds_total{container_name!="",
container_name!="POD"}[5m]), "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY
(colo, namespace, controller, pod_name, container_name)
```

/

```
sum(label_replace(container_spec_cpu_shares{container_name!="",
container_name!="POD"}, "controller", "$1", "pod_name", "^(.*)-[a-z0-9]+")) BY
(colo, namespace, controller, pod_name, container_name) * 1000
```

CPU shares allocated to container

Alert Manager



KubeCon



CloudNativeCon

Europe 2018

- Alert manager defines how to handle alerts (email, slack notification, etc.)
 - Grouping alert
 - Silences
 - Inhibition

Alert Rules



KubeCon



CloudNativeCon

Europe 2018

- alert: K8SNodeNotReady

```
expr: kube_node_status_condition{condition="Ready",status="true"} == 0
```

```
for: 1h
```

```
labels:
```

```
  severity: warning
```

```
  colo: bf1
```

```
  environment: production
```

```
annotations:
```

```
description: The Kubelet on {{ $labels.node }} has not checked in with the API,  
  or has set itself to NotReady, for more than an hour
```

```
summary: Node status is NotReady
```


Alert Rules



KubeCon



CloudNativeCon

Europe 2018

- alert: PodCPUPercentage

```
expr: colo_namespace_controller_pod:cpu_percentage:rate{namespace=~"kube-.*"} > 75
```

for: 10m

labels:

```
severity: critical
```

```
colo: bf1
```

```
environment: production
```

annotations:

```
description: 'Pod cpu usage is above 75 for {{ $value }}.
```

```
  Please find out the cause of the spike and if required increase CPU allocation'
```

```
summary: Pod cpu usage is above threshold
```

Alerting on Prometheus



KubeCon



CloudNativeCon

Europe 2018

- Federated prometheus monitors individual prometheus
- A cron job monitors federated prometheus

PrometheusDown (0 active)

```
alert: PrometheusDown
expr: up{job="federate"}
    == 0
for: 15m
labels:
  colo: federation
  environment: production
  severity: critical
annotations:
  description: Prometheus {{ $labels.instance }} are down. Please check the pod status
    and error log by running kubectl. Prometheus uses a lot of memory, so also check
    the memory usage of prometheus in kube dashboard. If prometheus is not schedulale
    due to lack of memory, try to fix a down node or drain a node to make room for
    it.
  summary: Prometheus are down
```



KubeCon



CloudNativeCon

Europe 2018

Dashboards



Cluster View

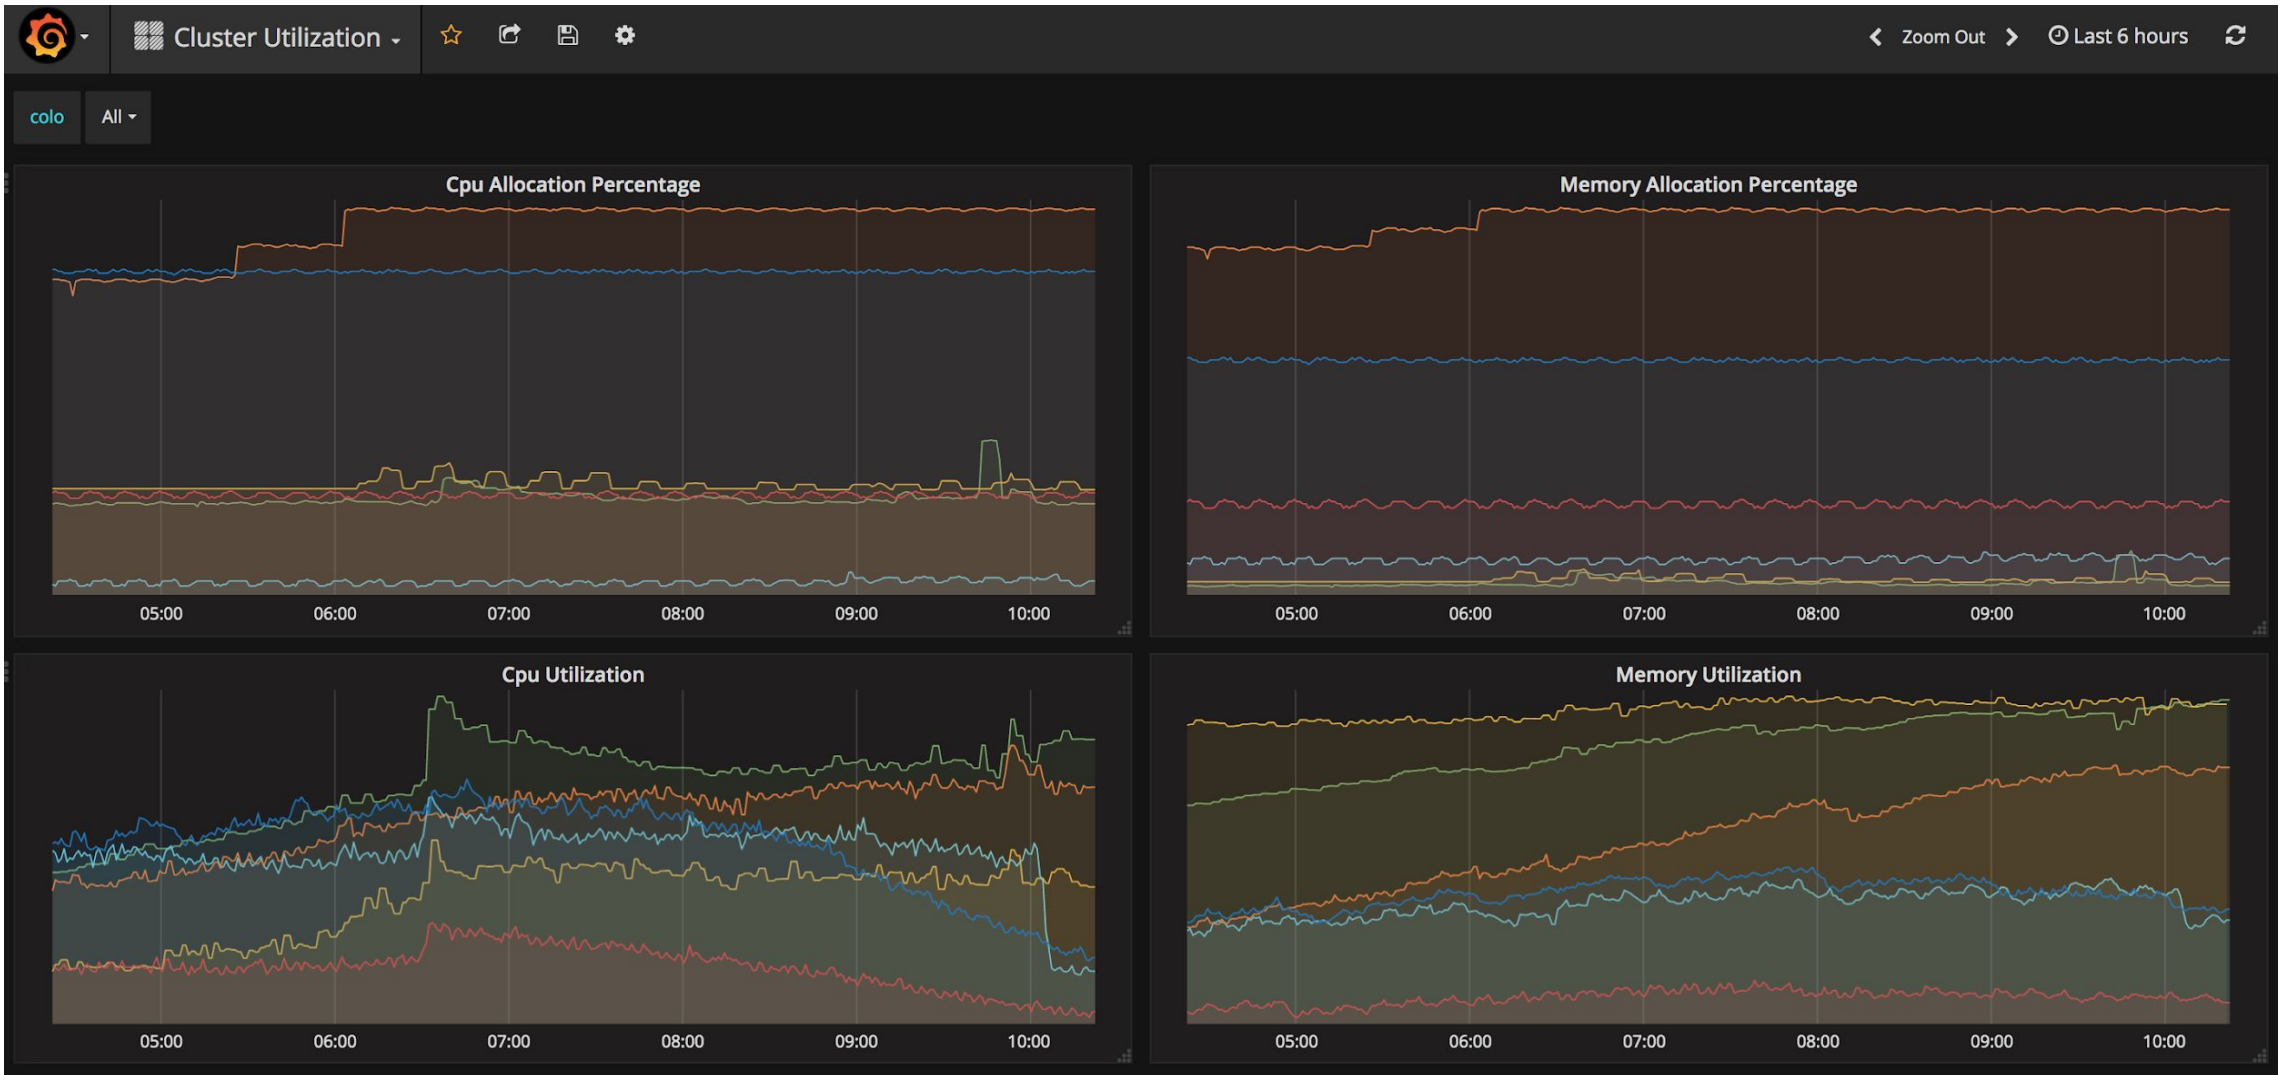


KubeCon



CloudNativeCon

Europe 2018



Namespace

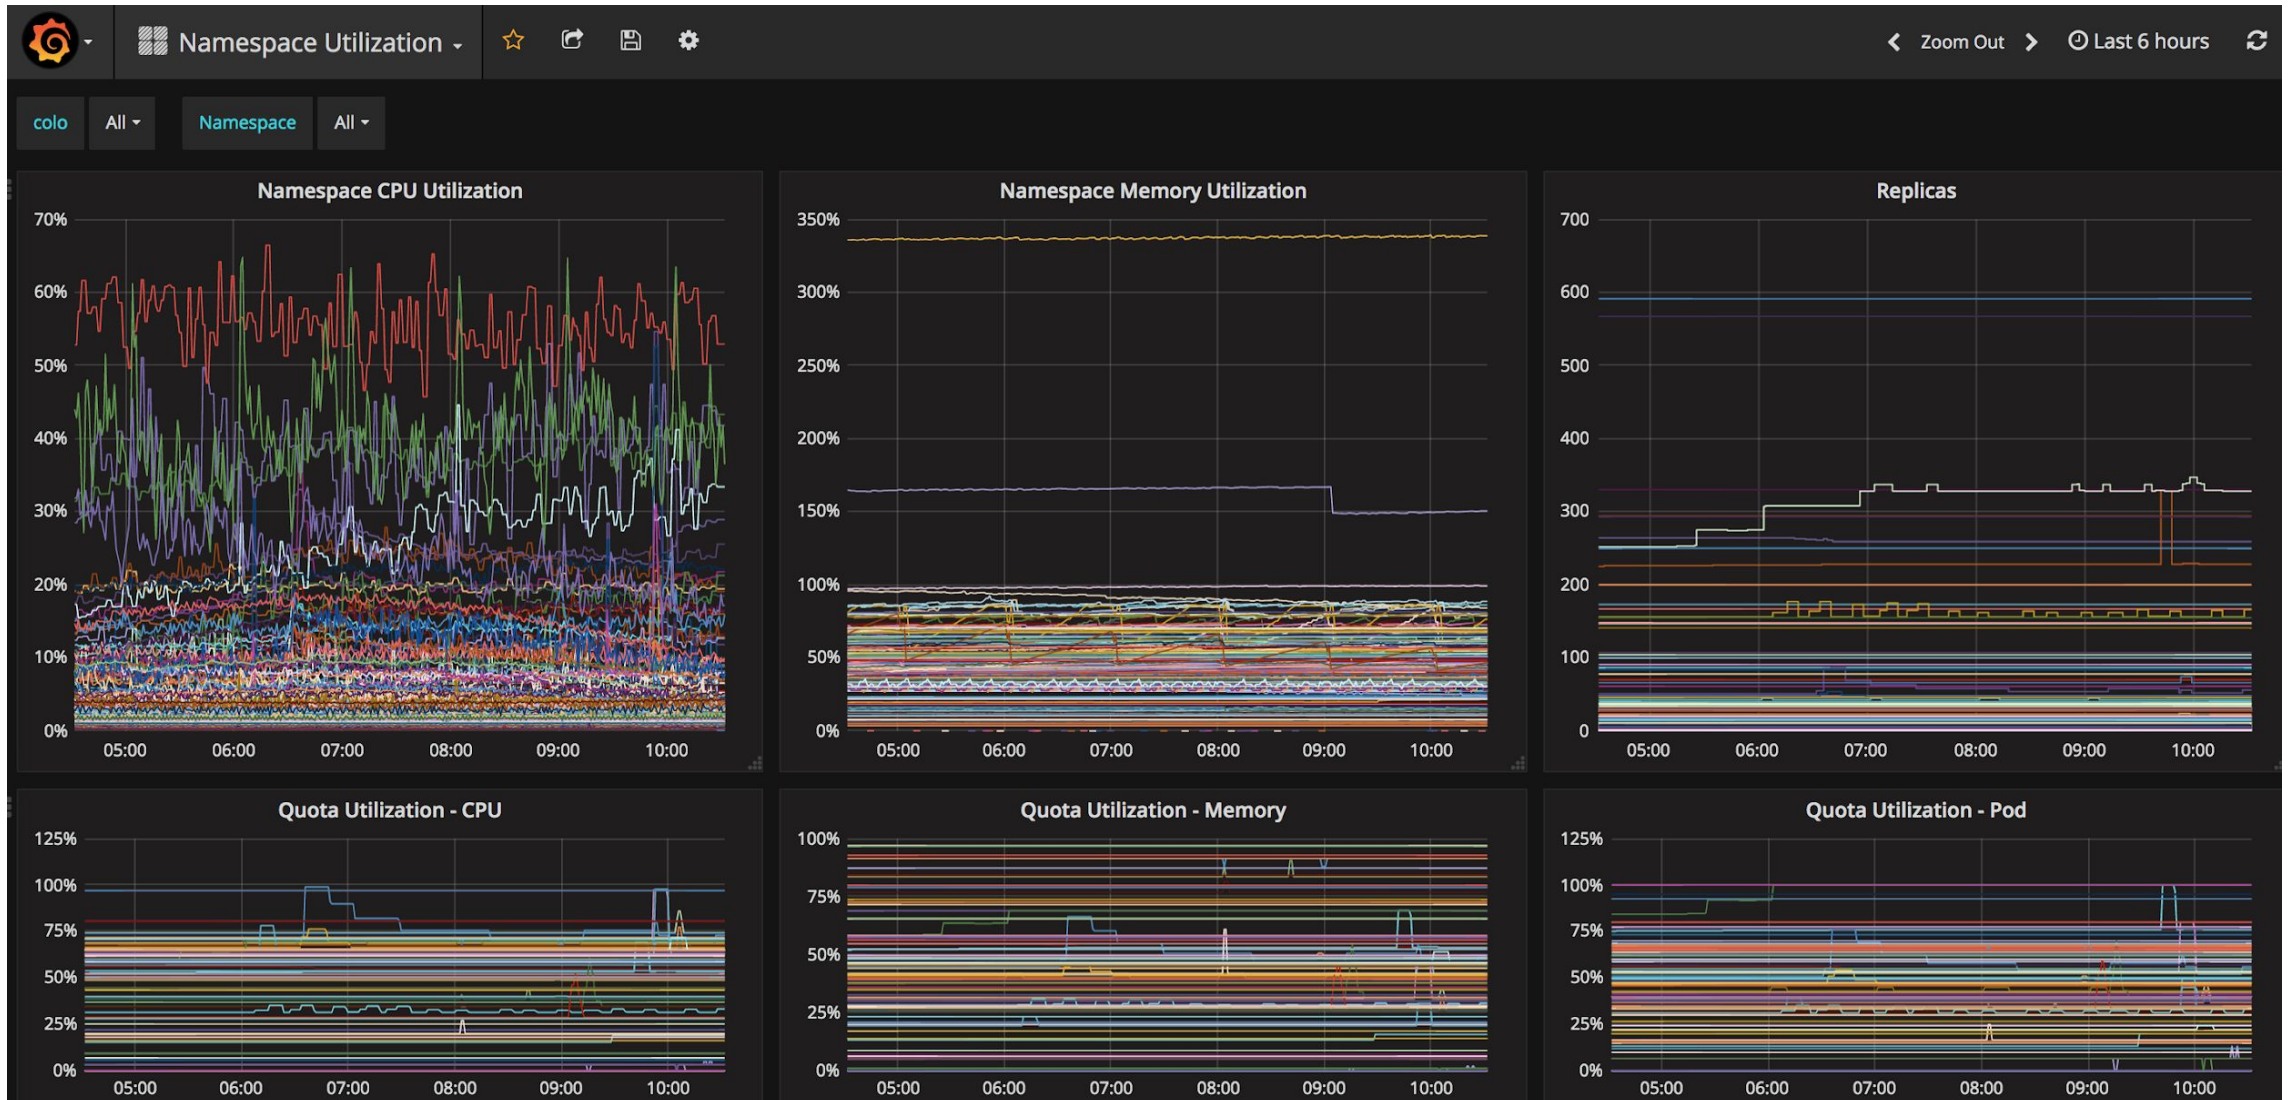


KubeCon



CloudNativeCon

Europe 2018



Deployment

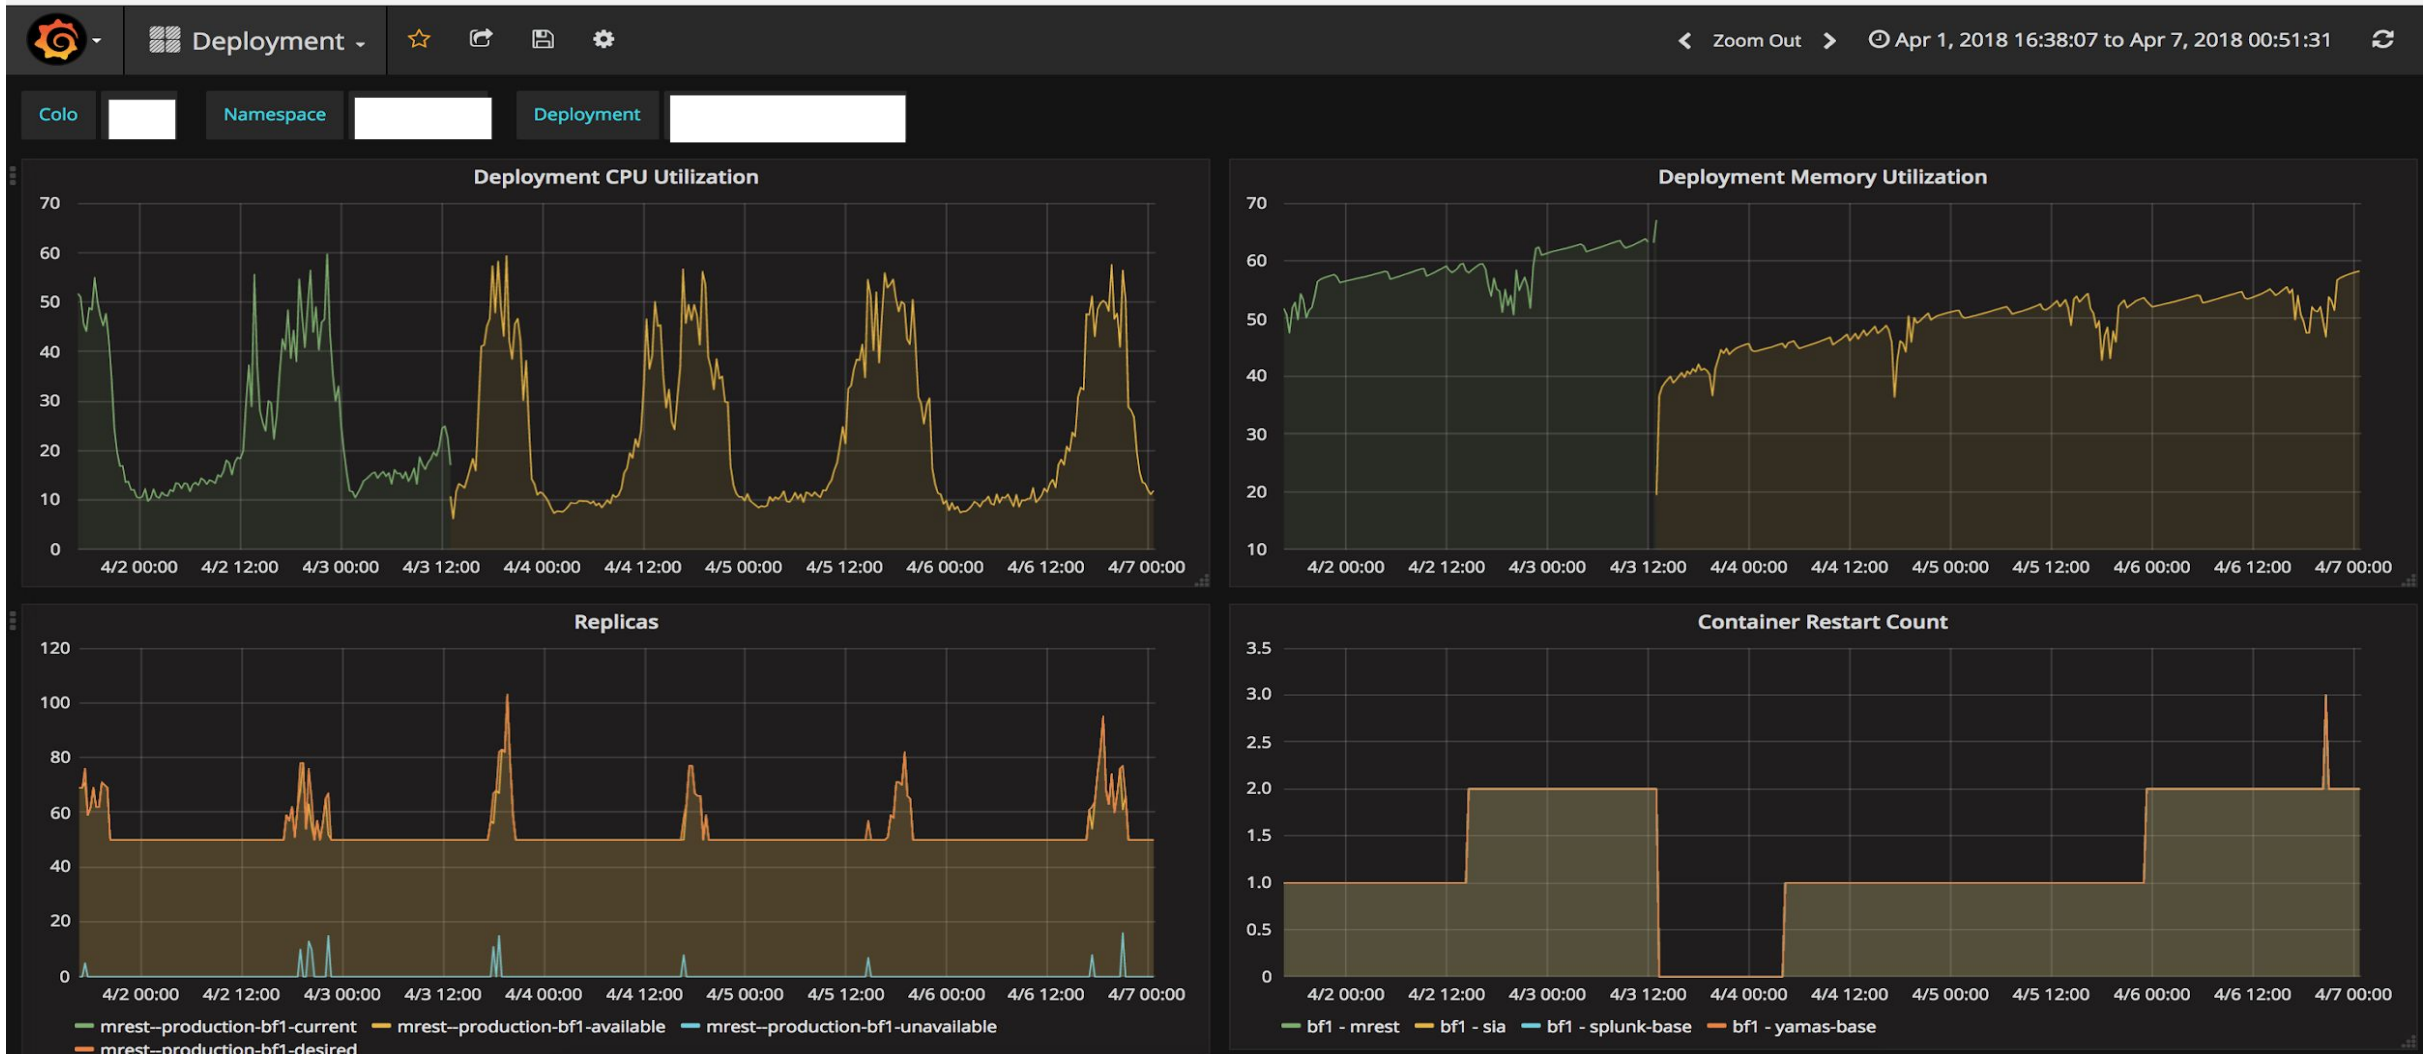


KubeCon



CloudNativeCon

Europe 2018



Controller



KubeCon



CloudNativeCon

Europe 2018



Scheduler

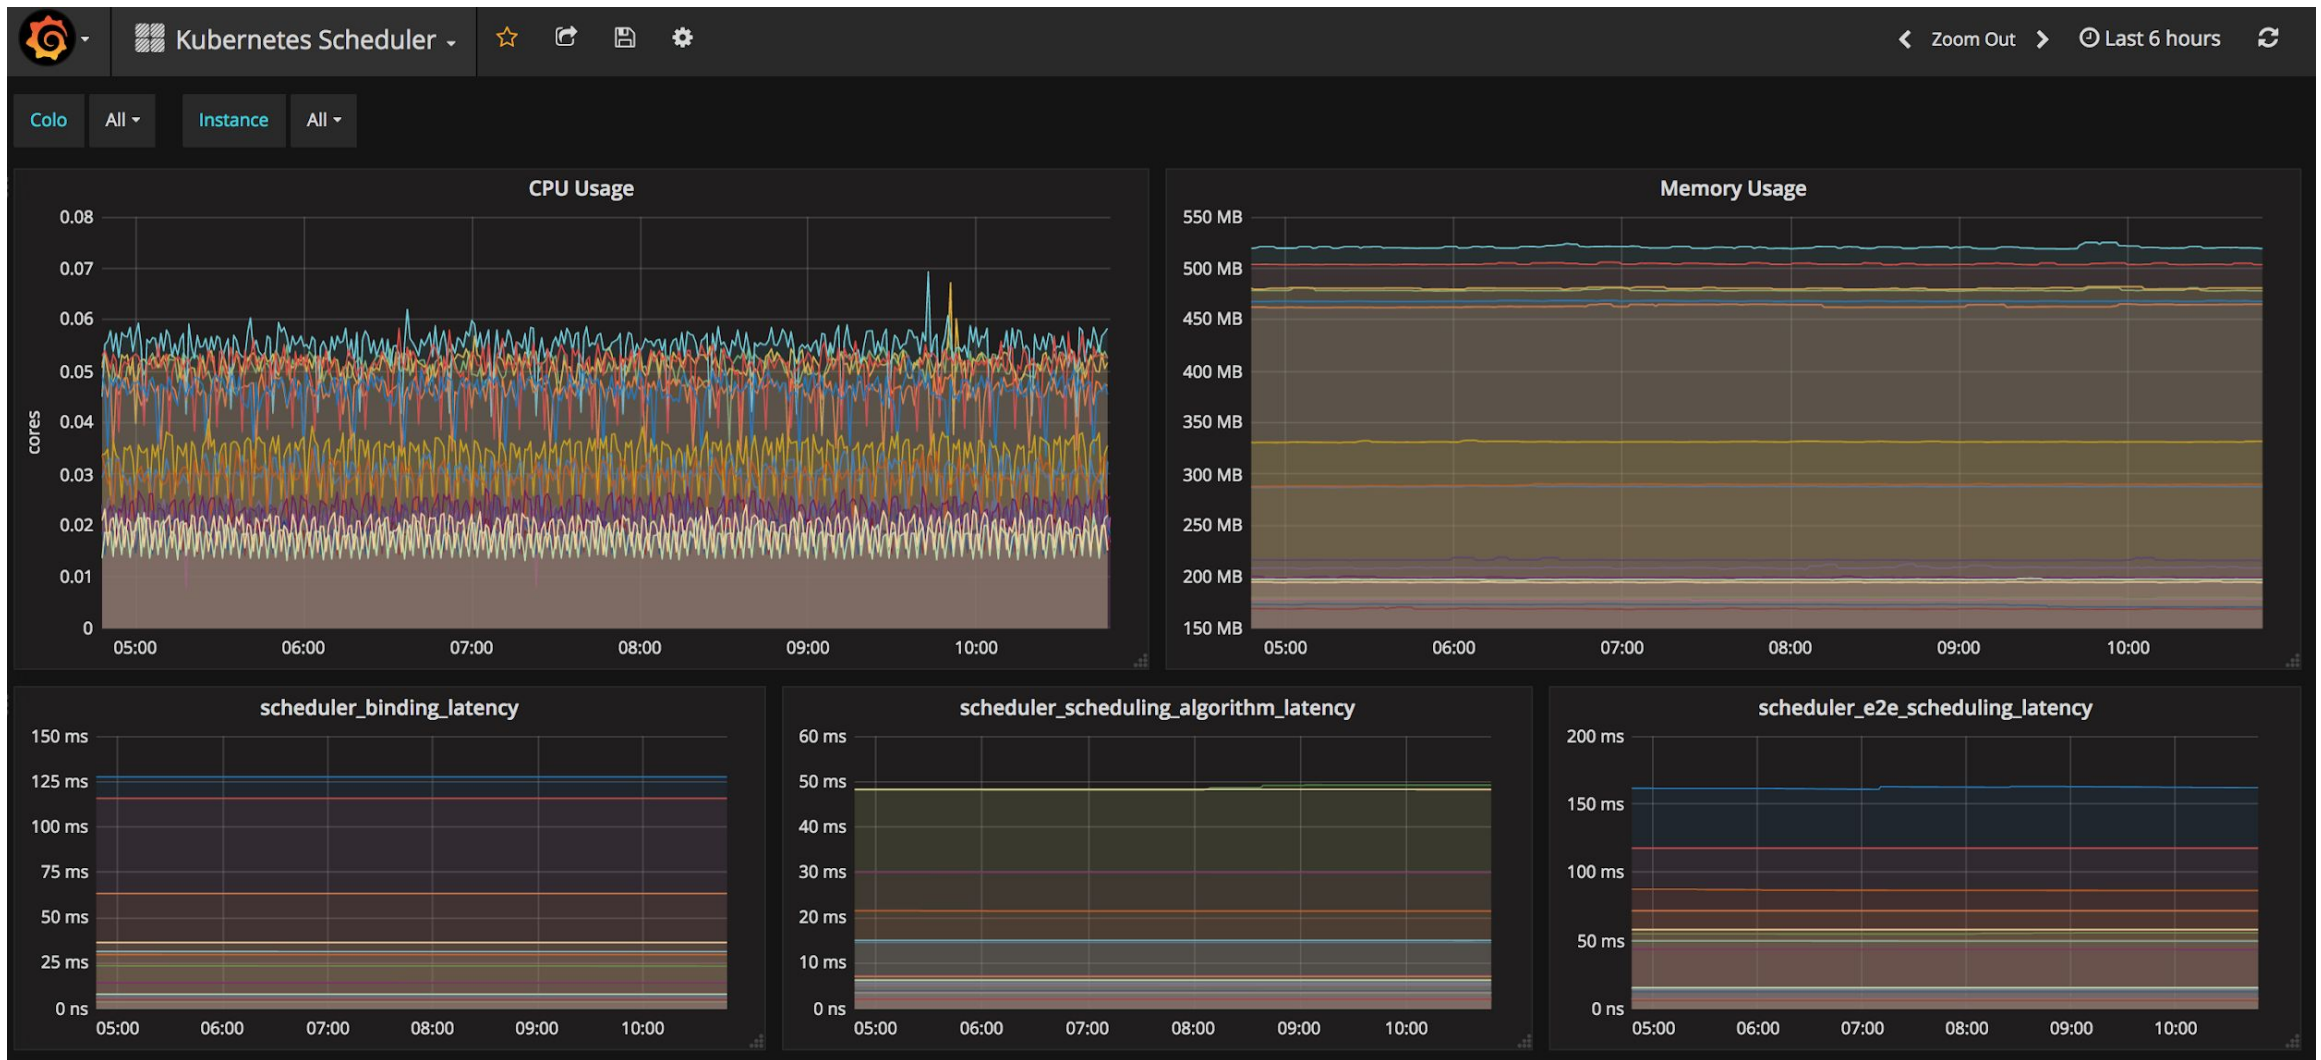


KubeCon



CloudNativeCon

Europe 2018



API server



KubeCon



CloudNativeCon

Europe 2018



Kubelet

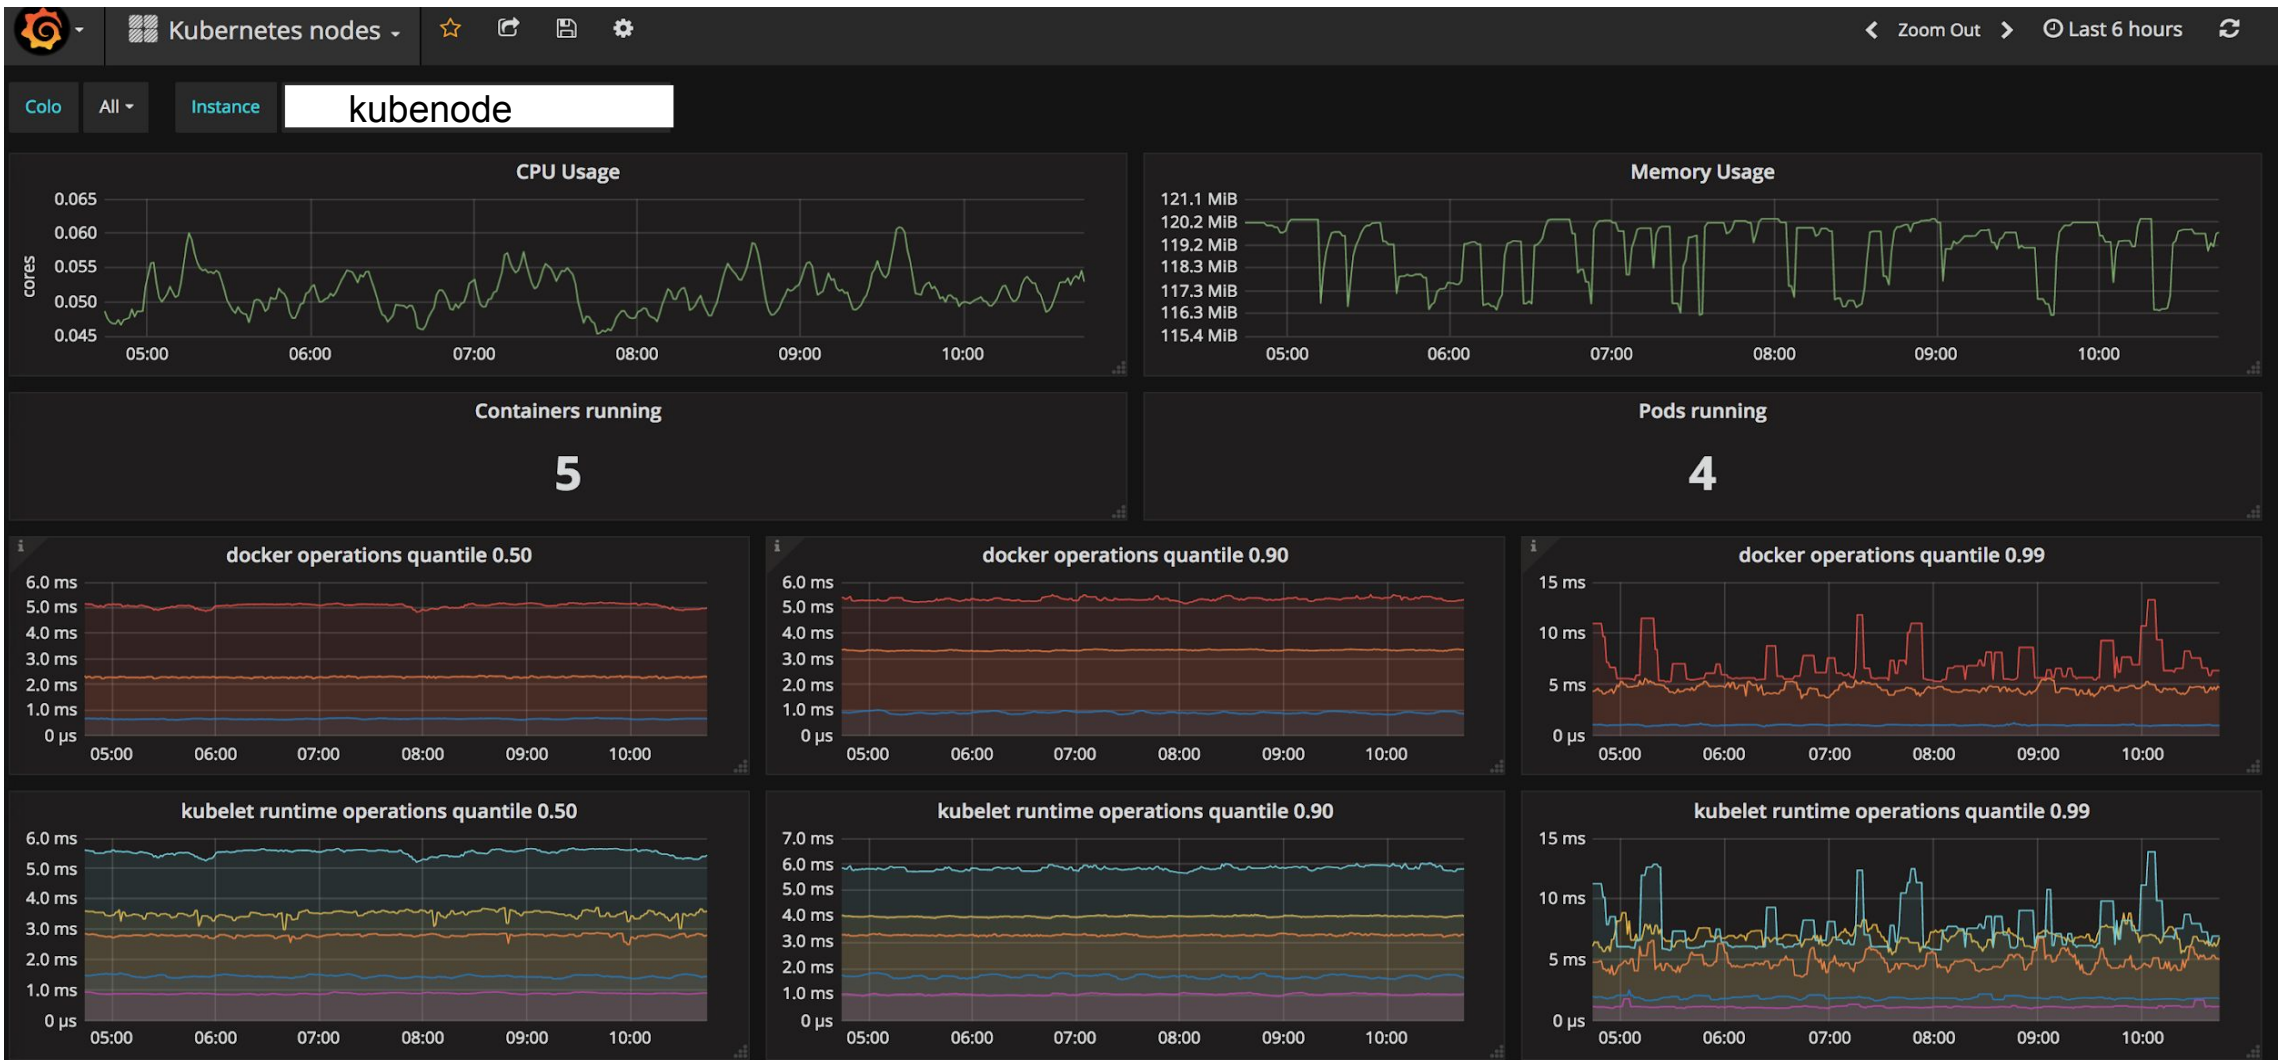


KubeCon



CloudNativeCon

Europe 2018



ETCD

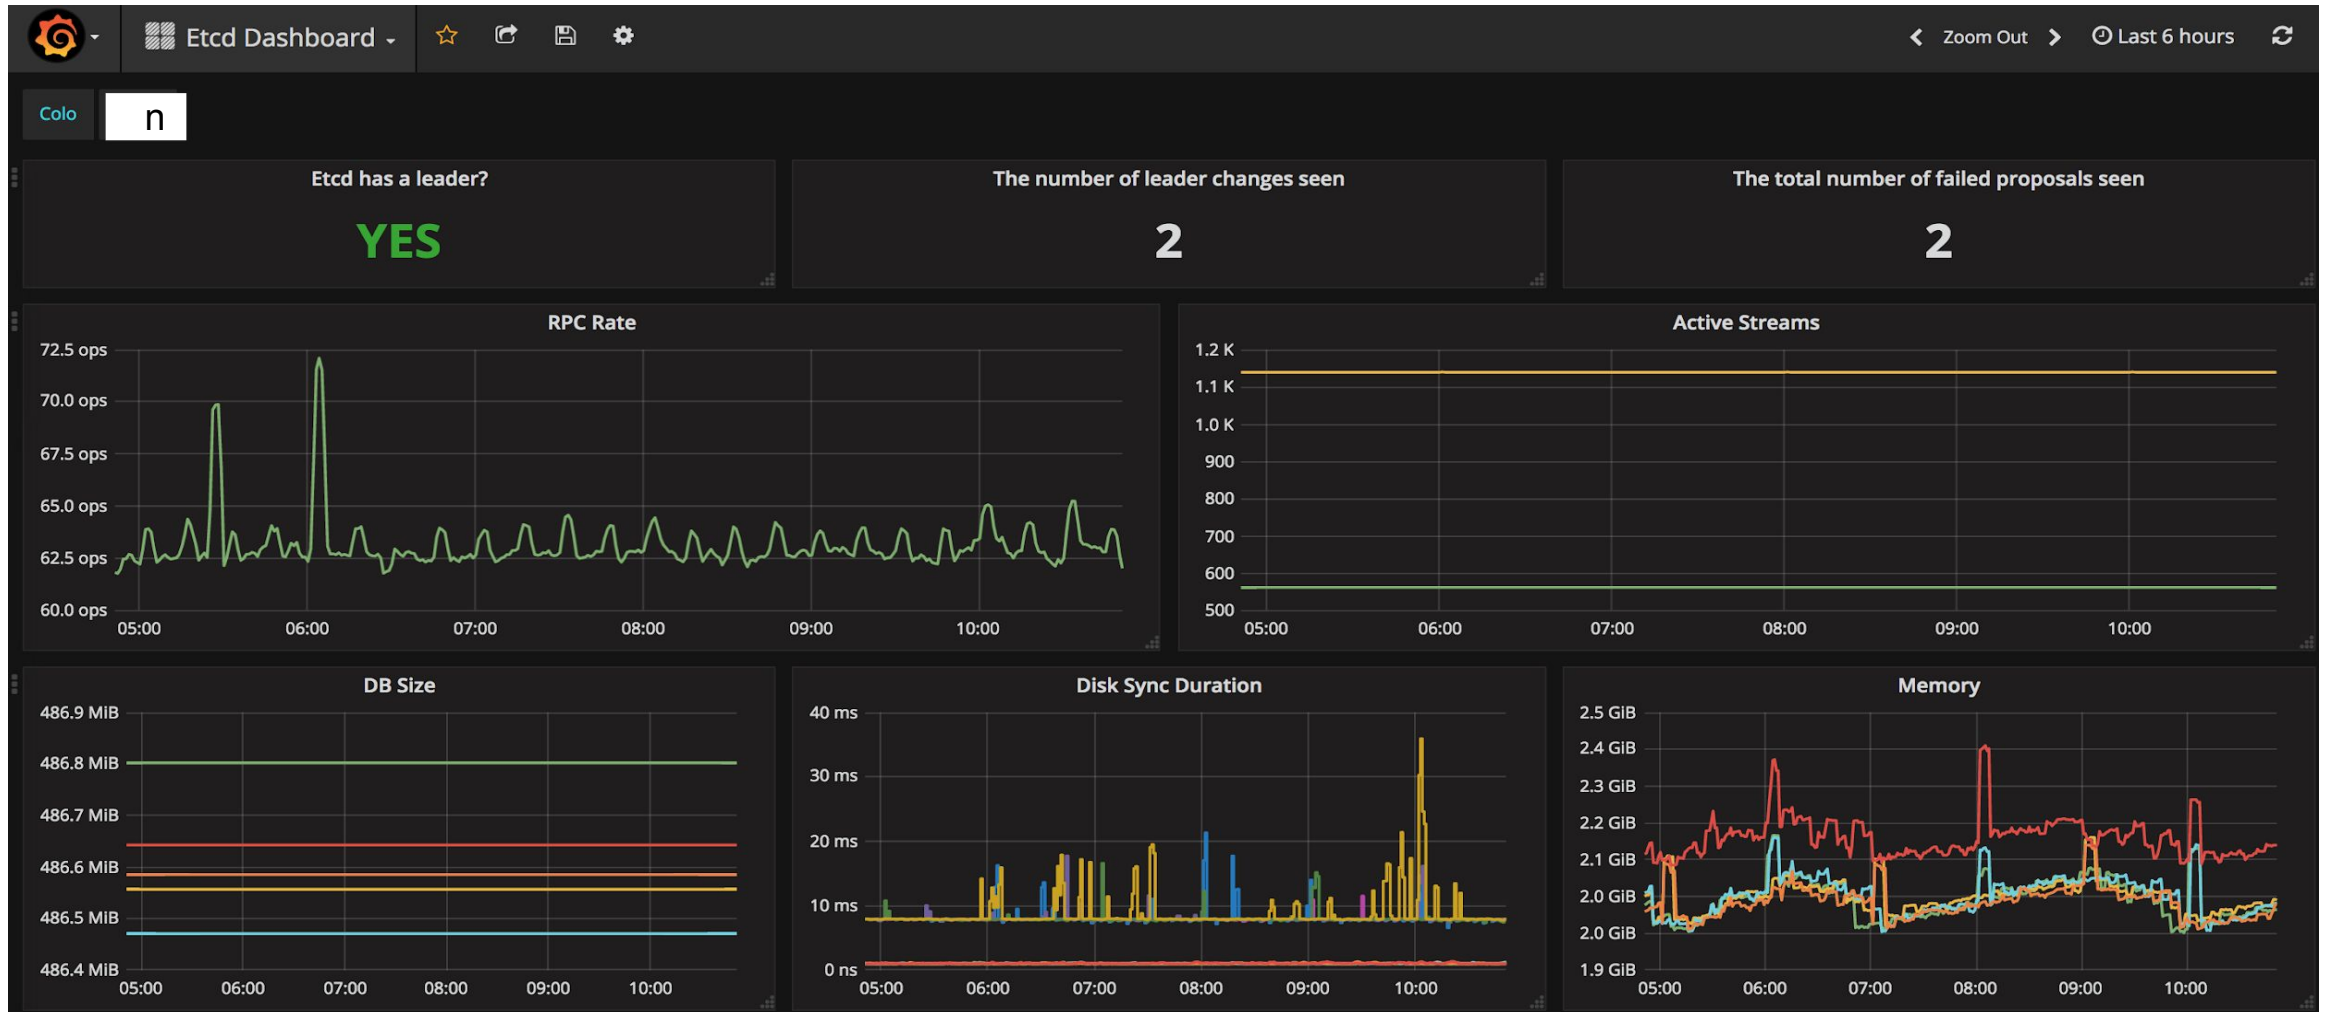


KubeCon



CloudNativeCon

Europe 2018



Package version

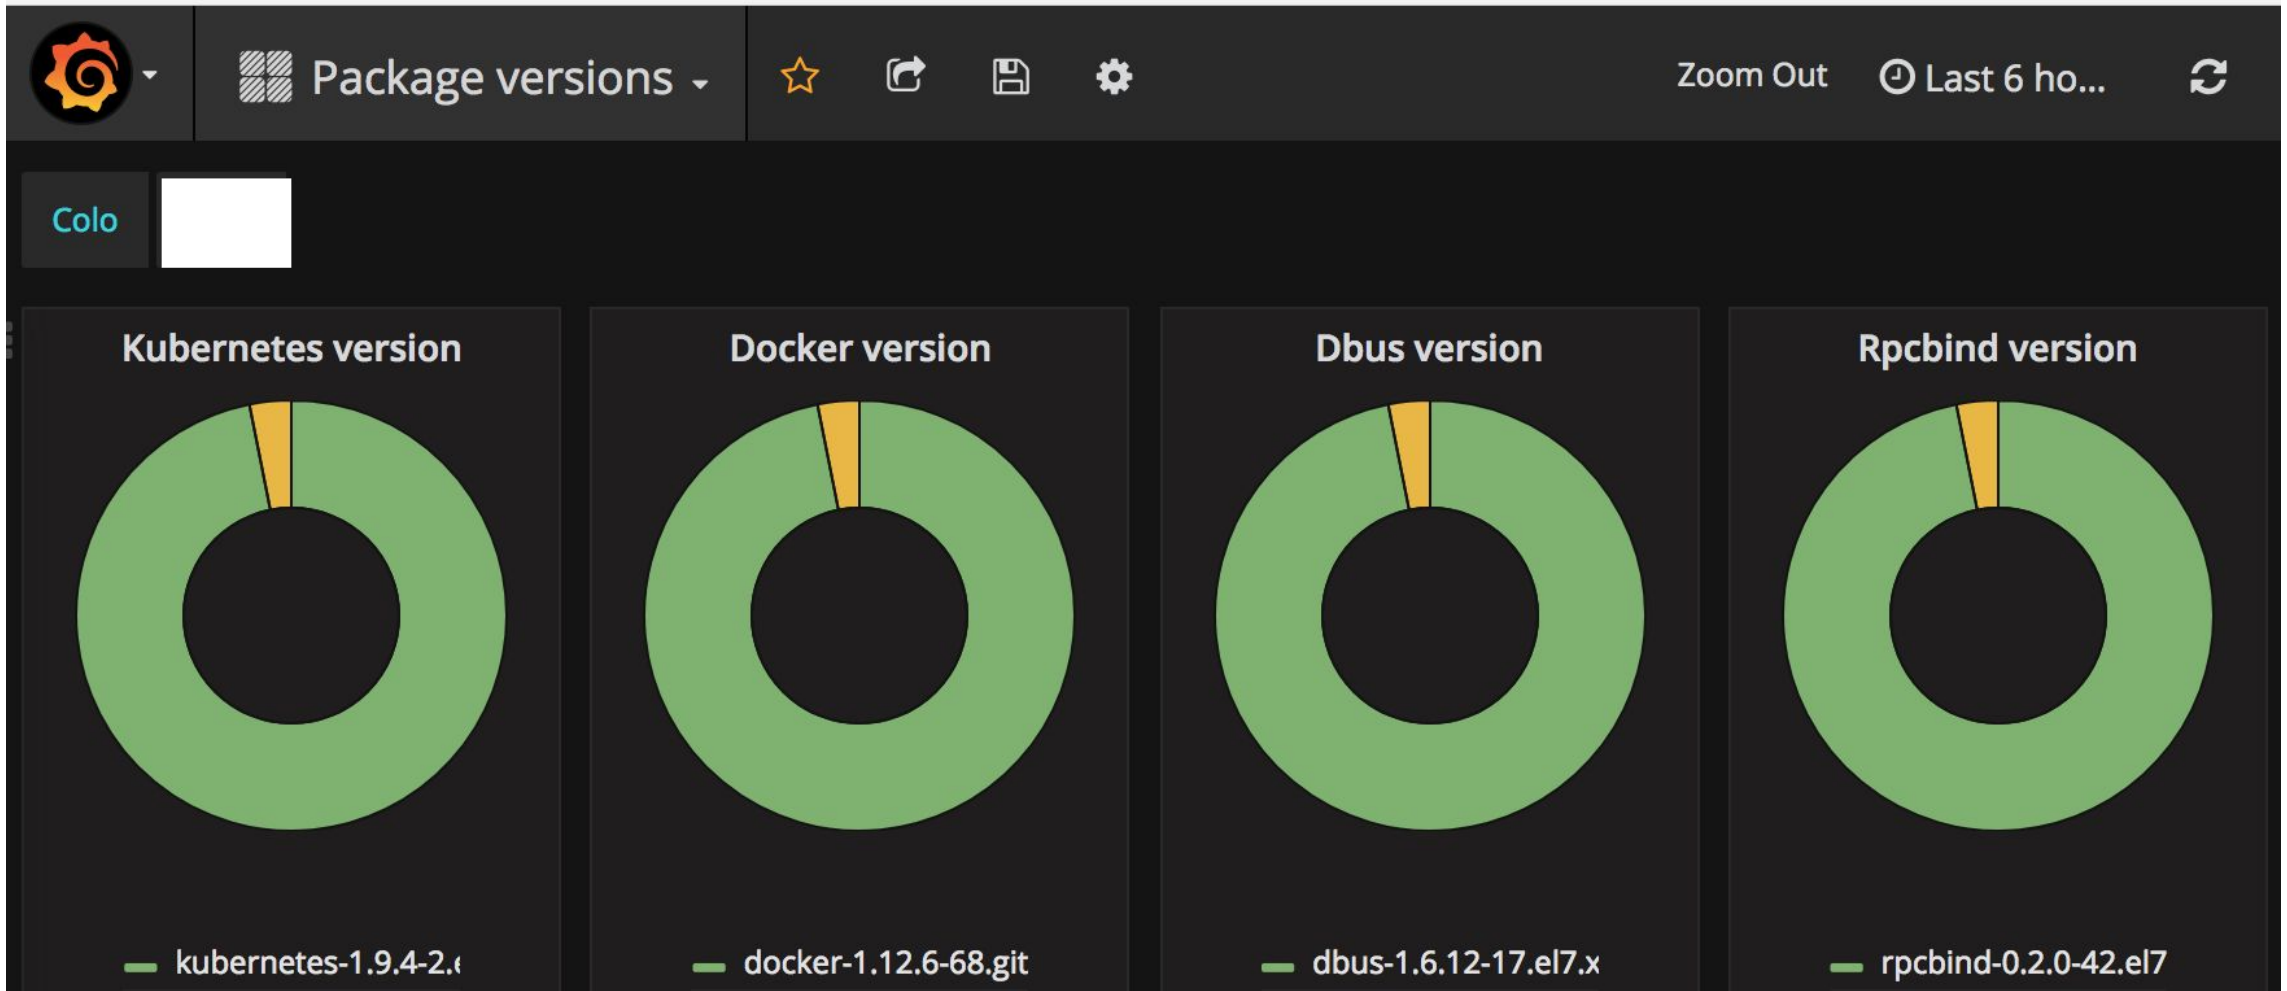


KubeCon



CloudNativeCon

Europe 2018



Thank you



KubeCon



CloudNativeCon

Europe 2018

Q & A