

SIG Instrumentation: Deep Dive

Frederic Branczyk @fredbrancz

Piotr Szcześniak @piosz



Recap

From SIG Instrumentation Intro



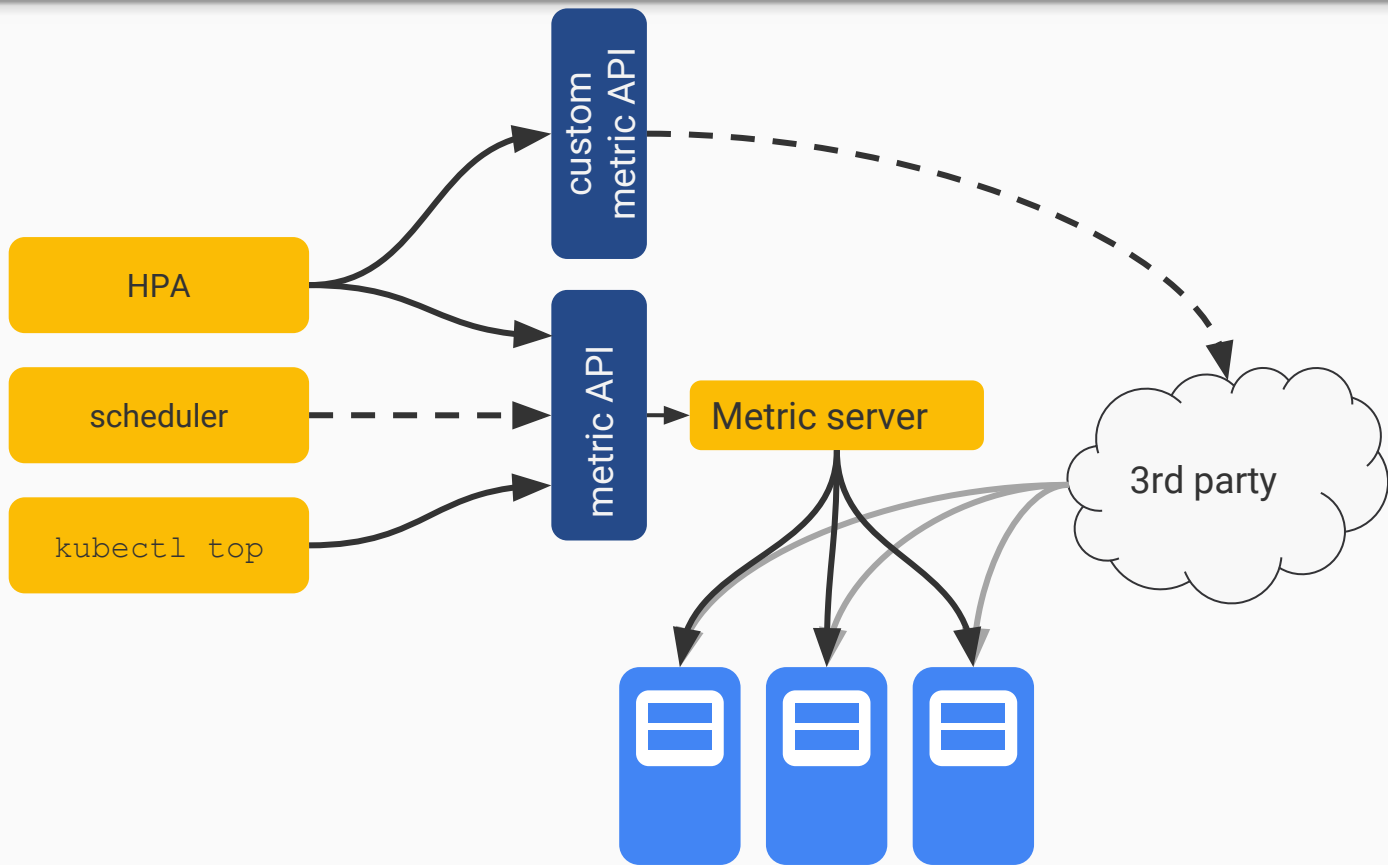
Recap

- `kubectl top`
- `kubectl logs`
- `kubectl get events`
- Prometheus format

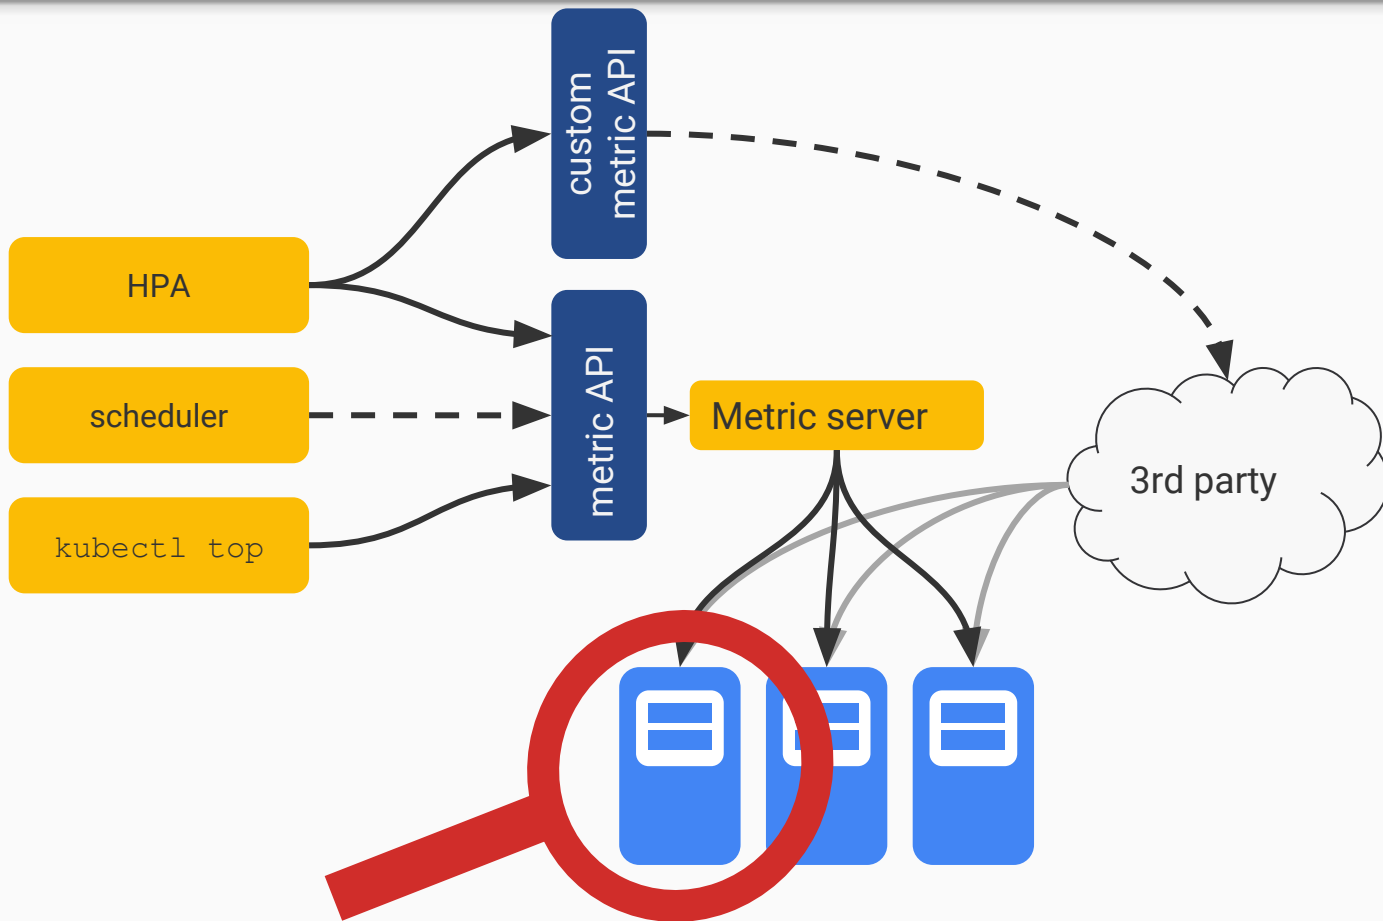
Architecture overview



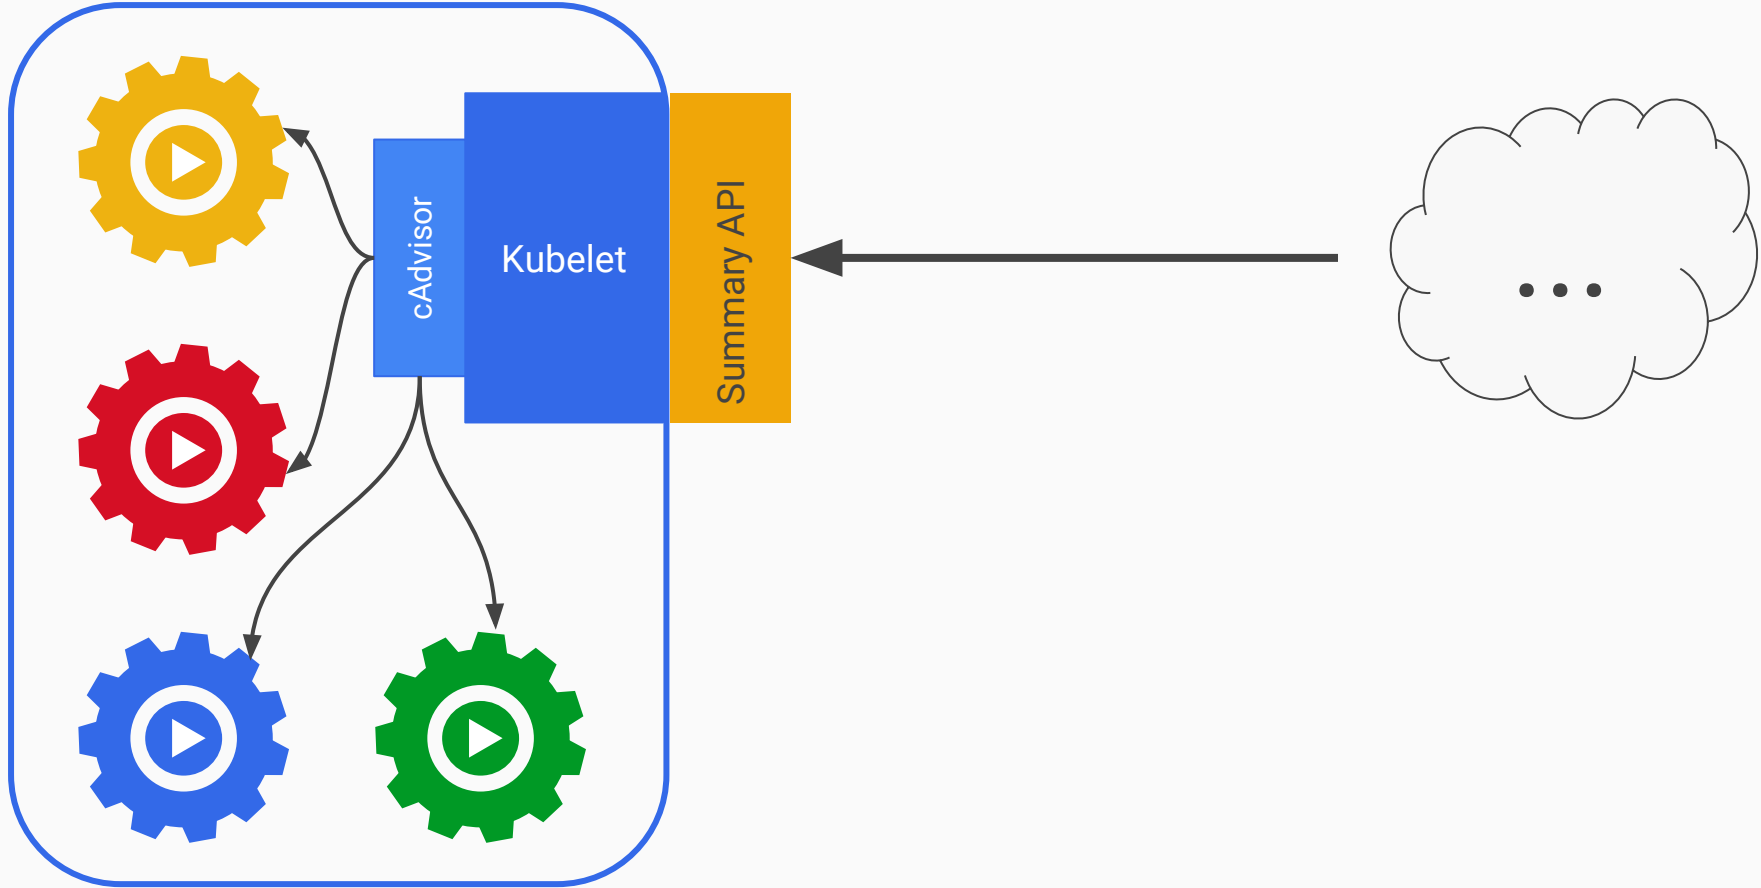
Monitoring architecture overview



Monitoring architecture overview



Monitoring architecture overview - node



APIs



Metrics related APIs

- Metrics API
- Custom Metrics API
- External Metrics API

Metrics related APIs

- Metrics API -> default implementation
- Custom Metrics API
- External Metrics API

Metrics related APIs

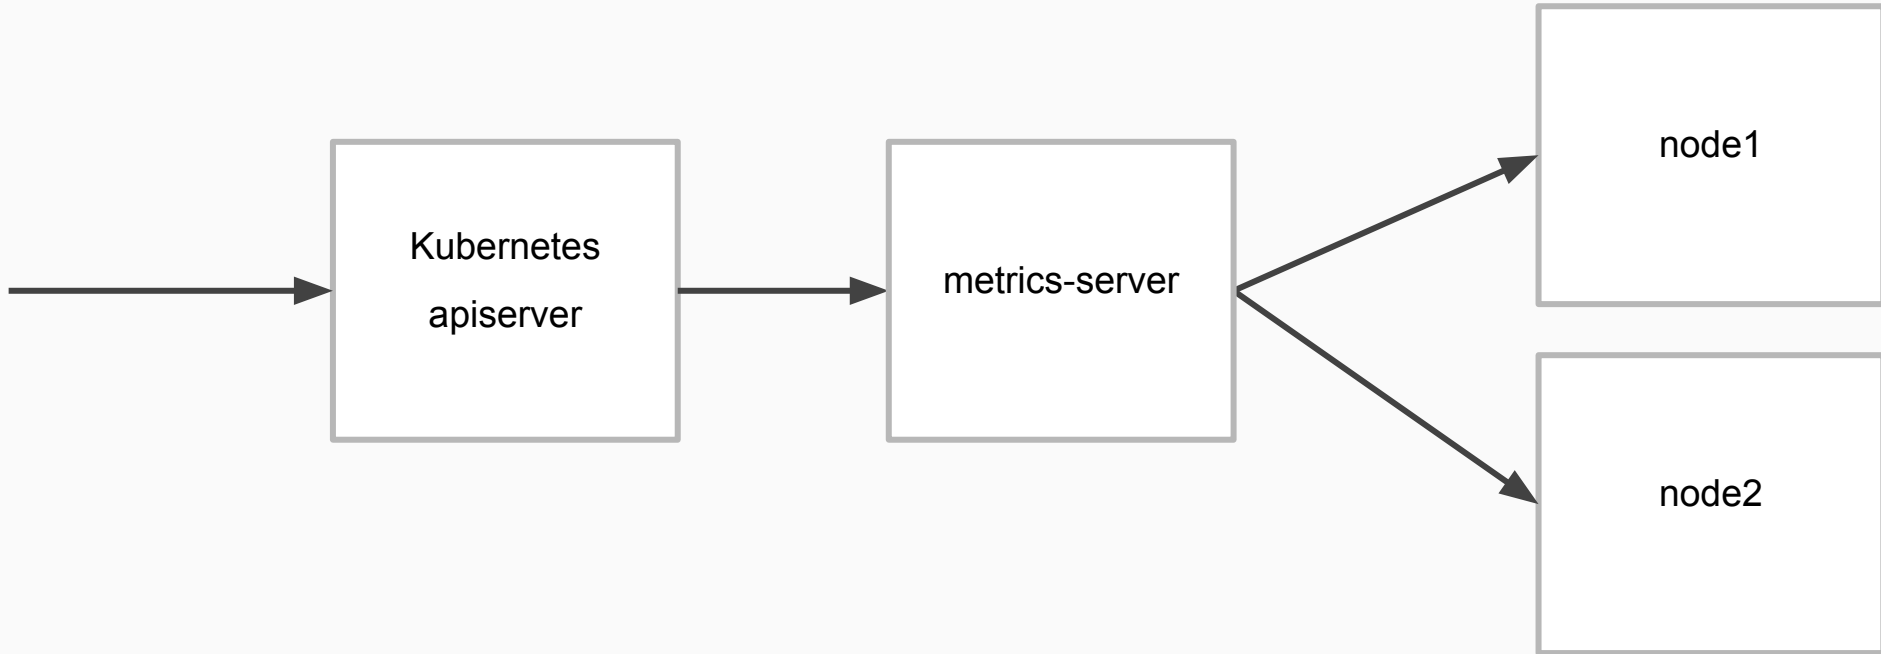
- Metrics API -> default implementation
- Custom Metrics API -> provider specific
- External Metrics API

Metrics Server internals

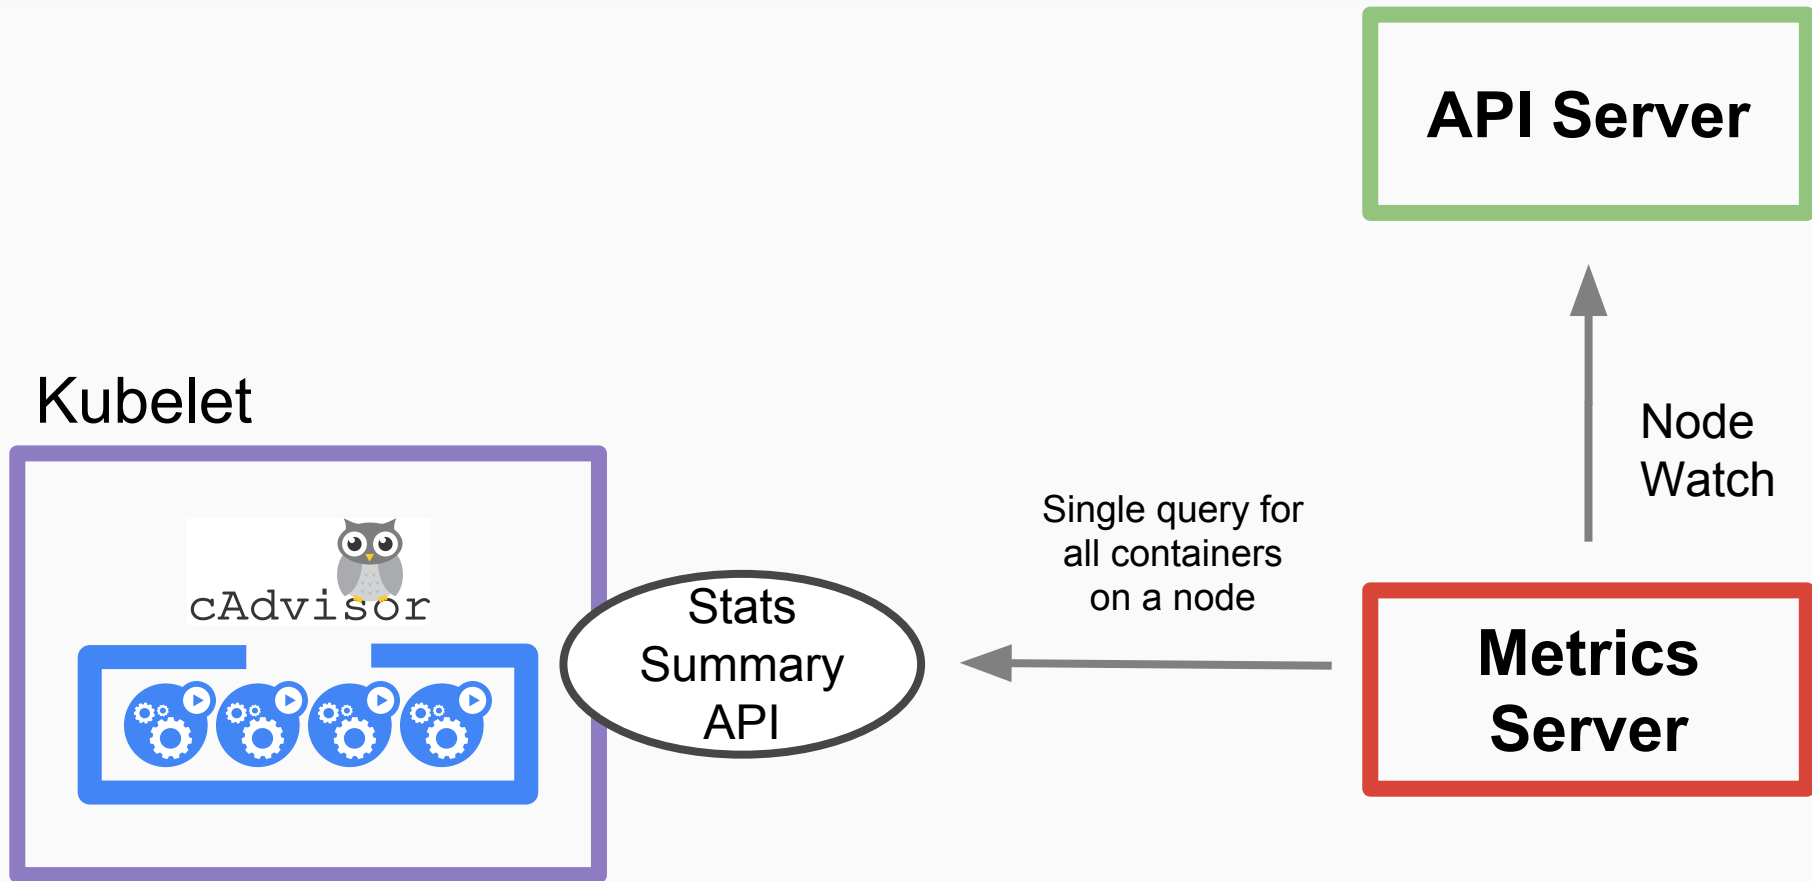


Metrics Server internals

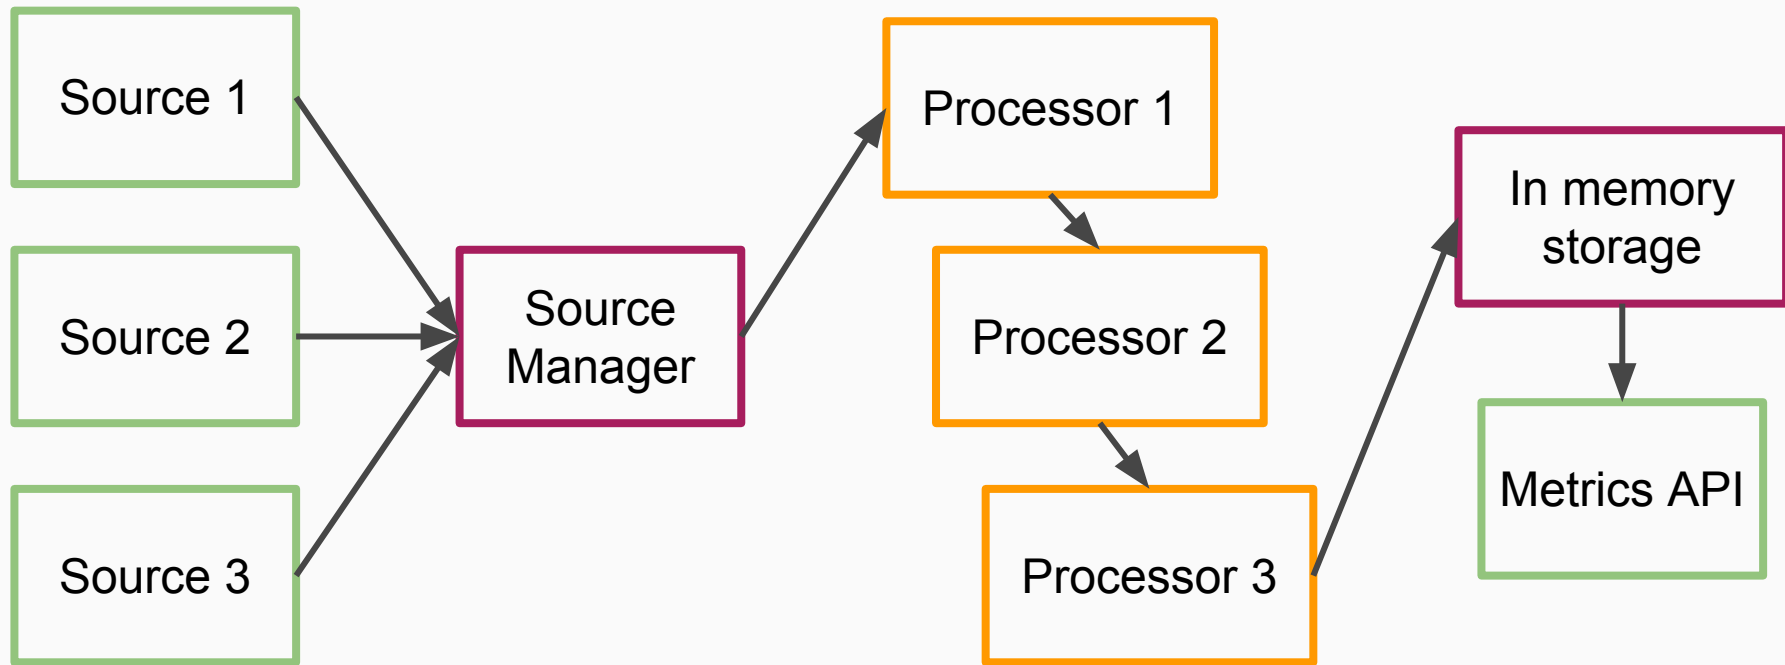
- Aggregated API



Metrics Server internals



Metrics Server internals



CM API adapters



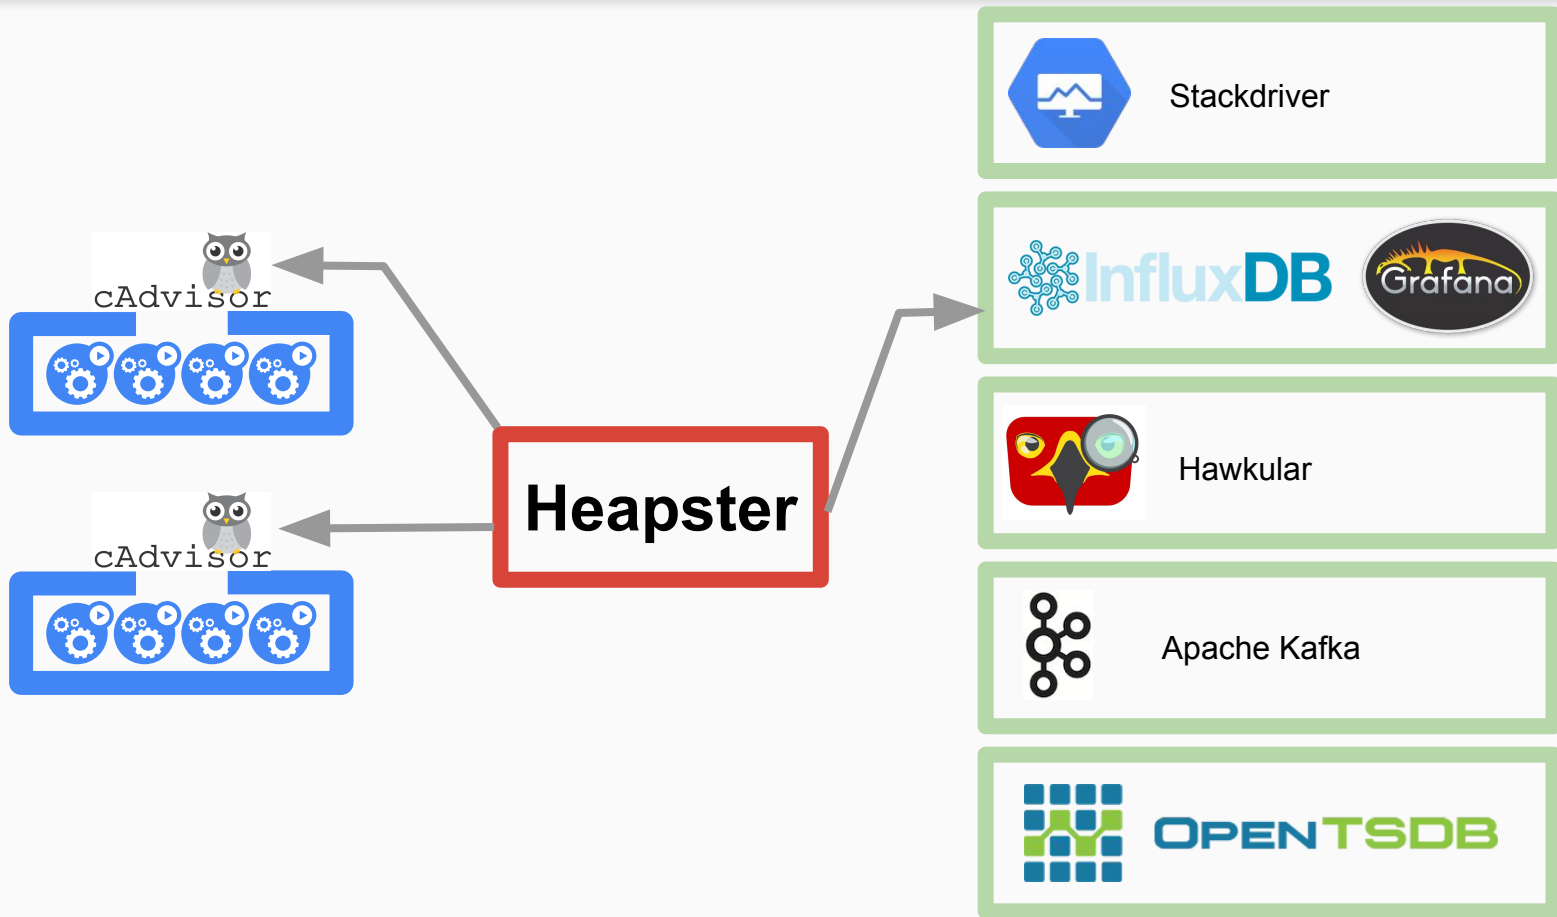
CM API adapters

- Prometheus
- Stackdriver
- Datadog?
- ...

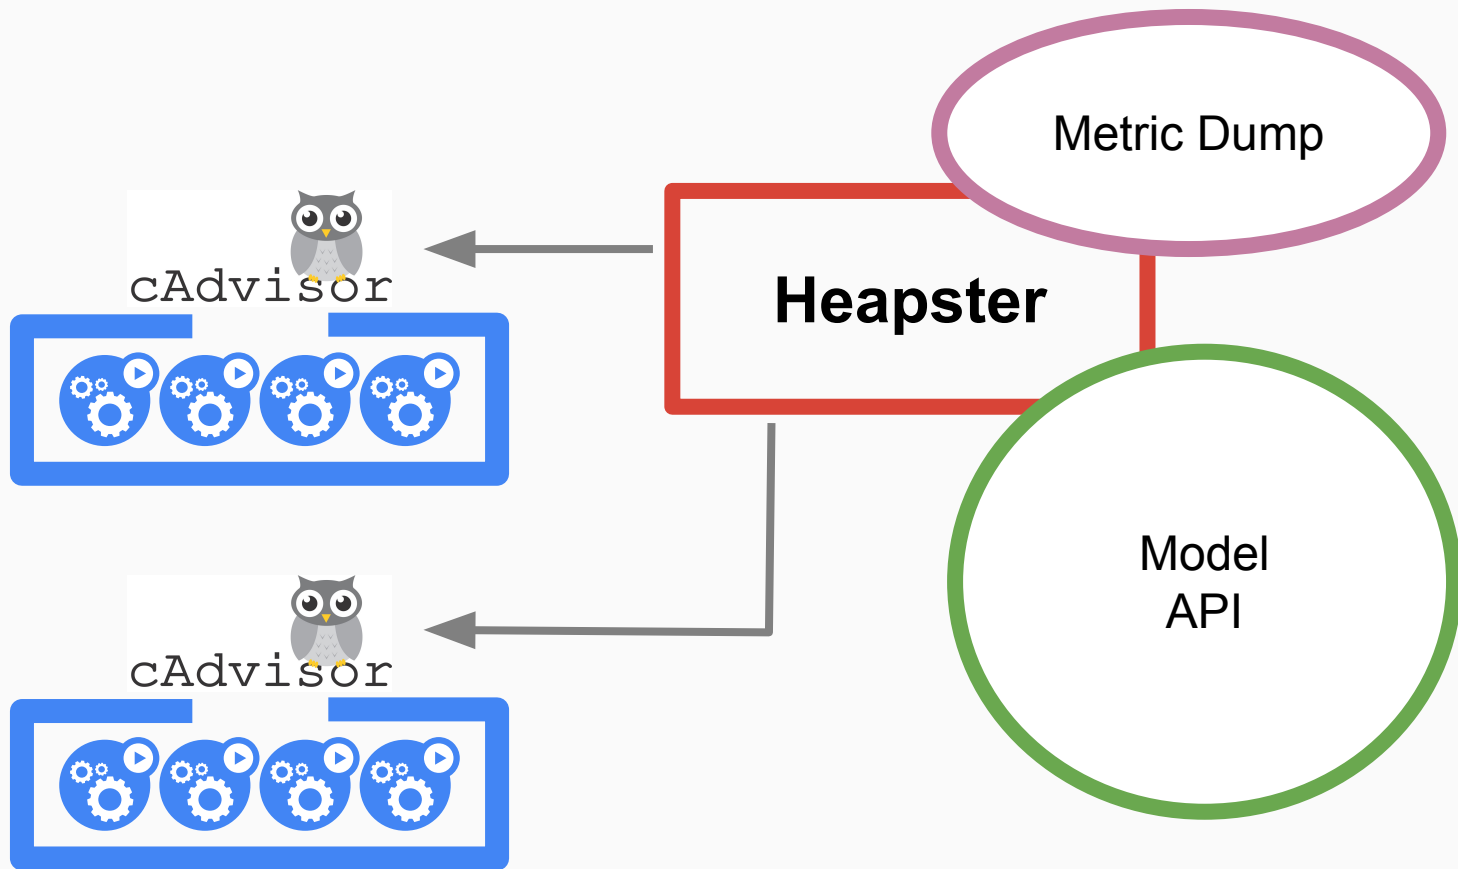
Heapster



Heapster internals



Heapster internals



Heapster problems

- hard to maintain
- non Kubernetes quality APIs

Heapster problems

- Maintenance
- non Kubernetes quality APIs

Solution: Deprecate

kube-state-metrics

History and internals



~1.5 years ago

We asked ourselves:
What's missing in
Kubernetes monitoring?

Debugging with kubectl = meh

→ Cluster state

Research ...

Enter kube-state-metrics

```
apiVersion: apps/v1beta1
kind: Deployment
metadata:
  name: app1
spec:
  replicas: 3
status:
  replicas: 2
```



```
kube_deployment_spec_replicas 3
kube_deployment_status_replicas 2
```

2 metrics

Untouched for half a year

Sig-instrumentation was founded

→ picked up maintainership

Fast forward

Interest, contributions,
involved companies
→ Road to v1.0

Load tests results

1Gb memory

0.3 cores CPU

1000 nodes

v1.0 released in August 2017

~ 531 Commits

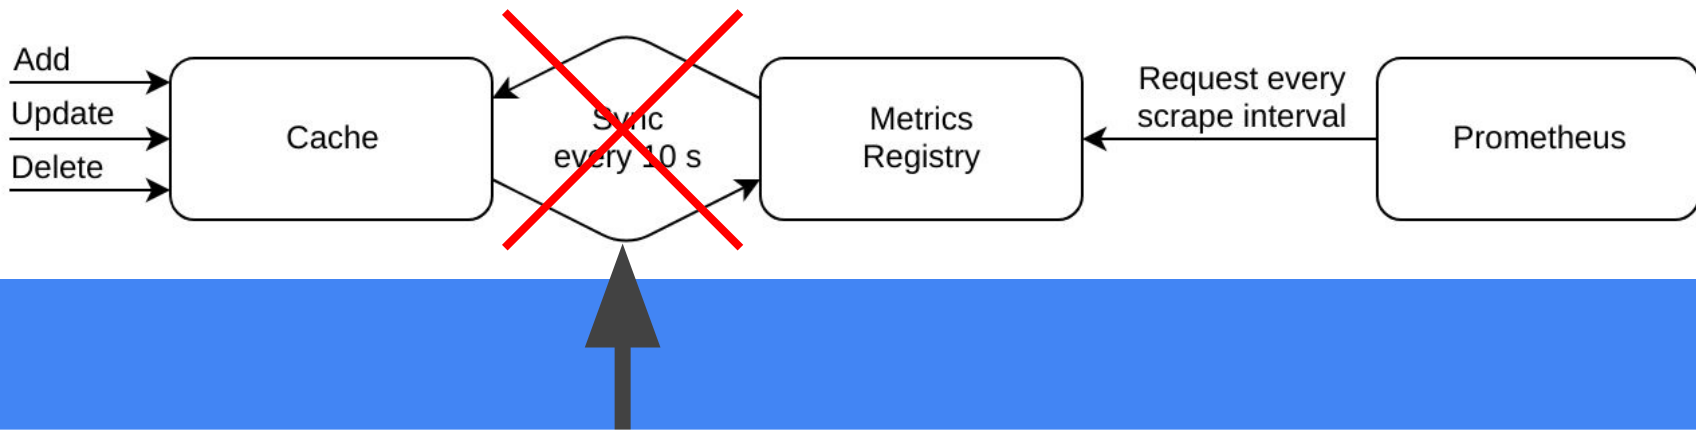
~ 145 Metrics

~ 72 contributors

* stats from 27th of April 2018


Nice, but deep dive!

Evolution of architecture




```
deploymentReplicas.WithLabels("my-namespace", "my-deployment-1").Set(7)
```

Today's architecture



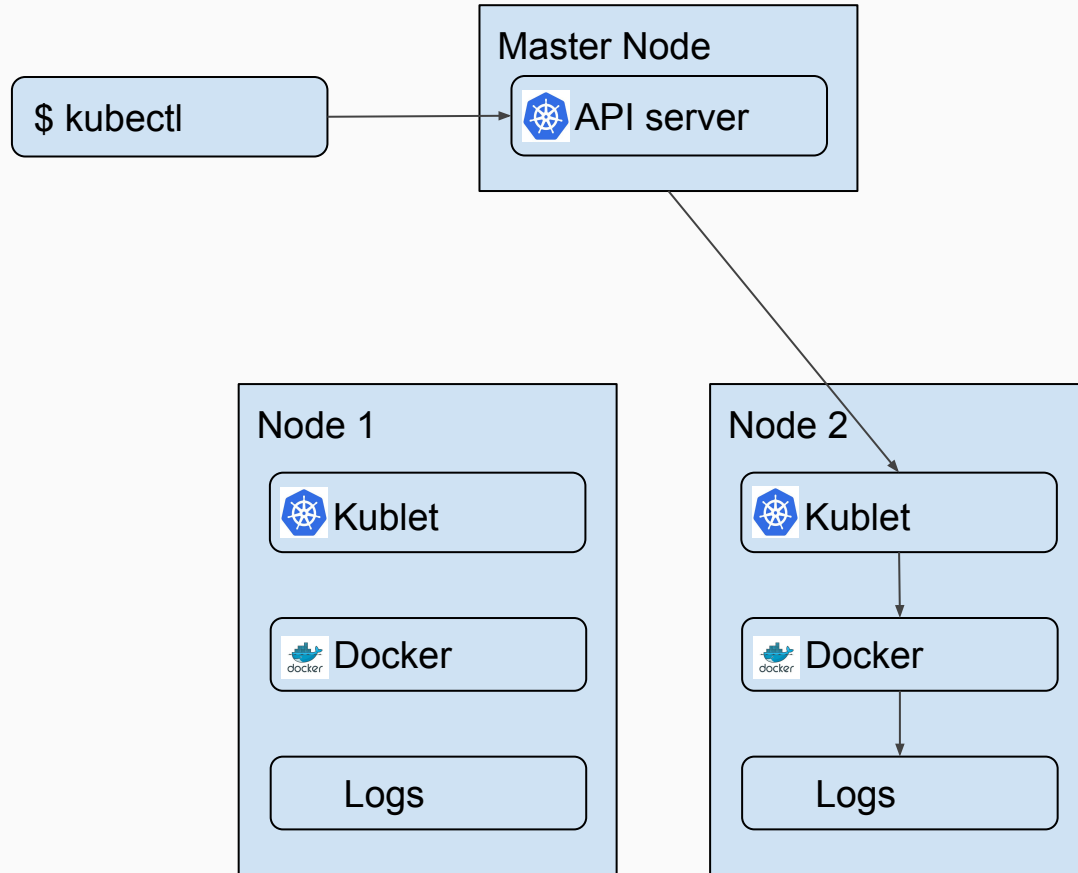
Future architecture



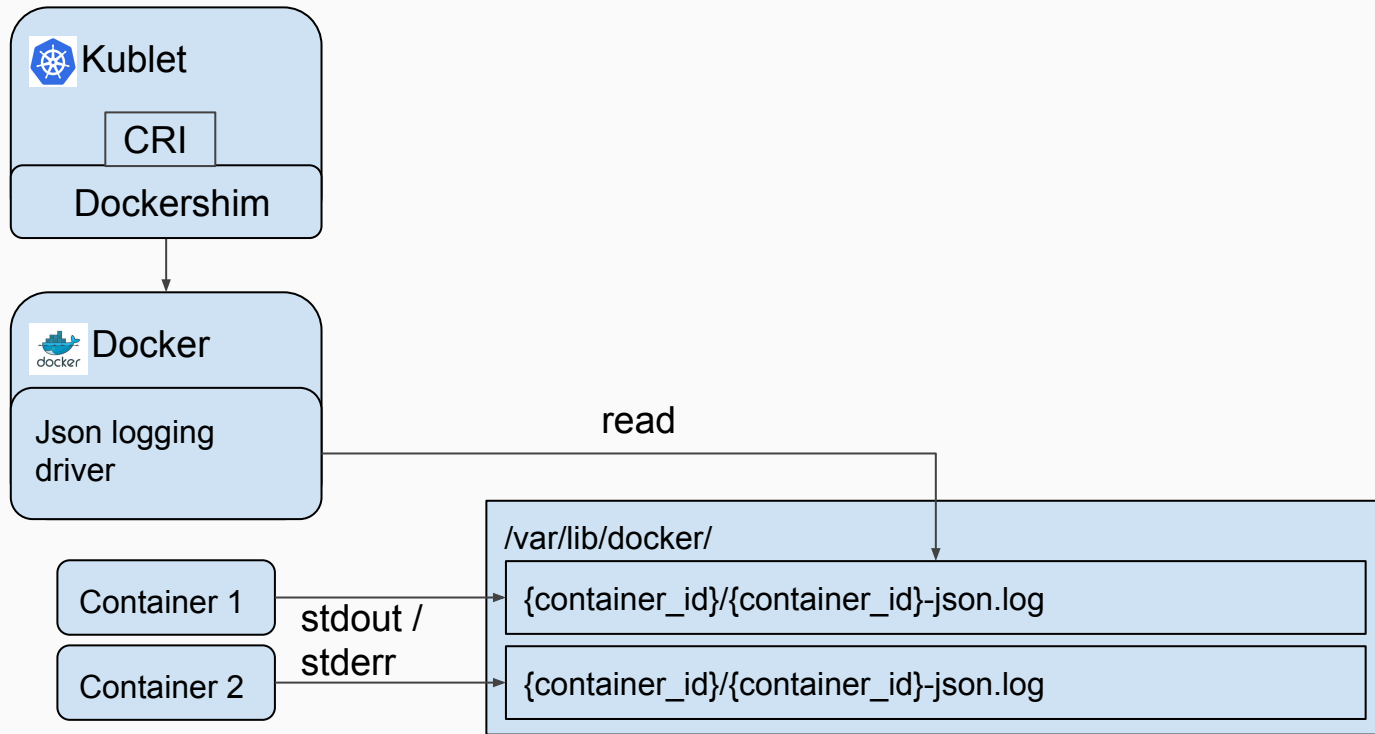
Additional optimizations possible

Logging architecture

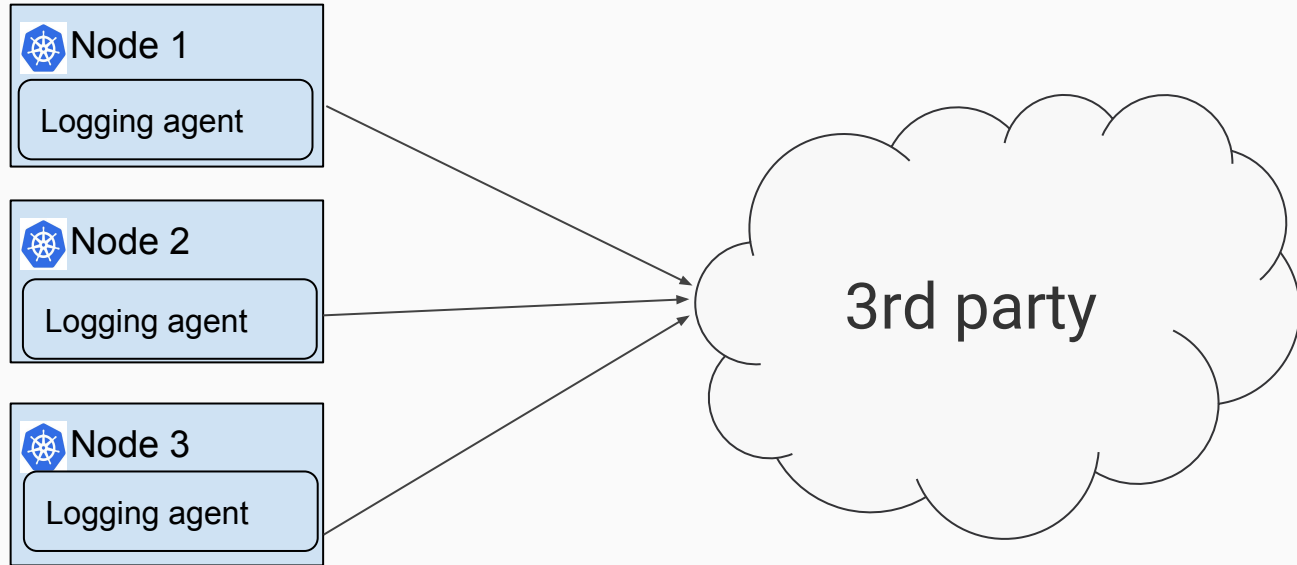
Logging architecture



Handling logs on node



3rd party logging solution



Thank you!

Questions?