



Metric Storage for Capacity Management of Kubernetes/OpenShift Clusters

Ulrike Klusik, <https://www.consol.de>

Motivation

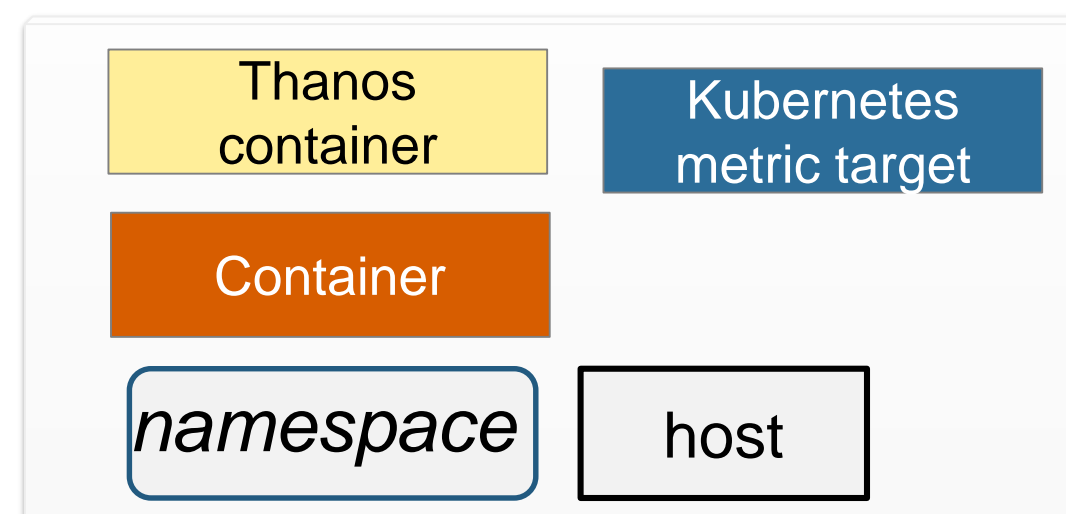
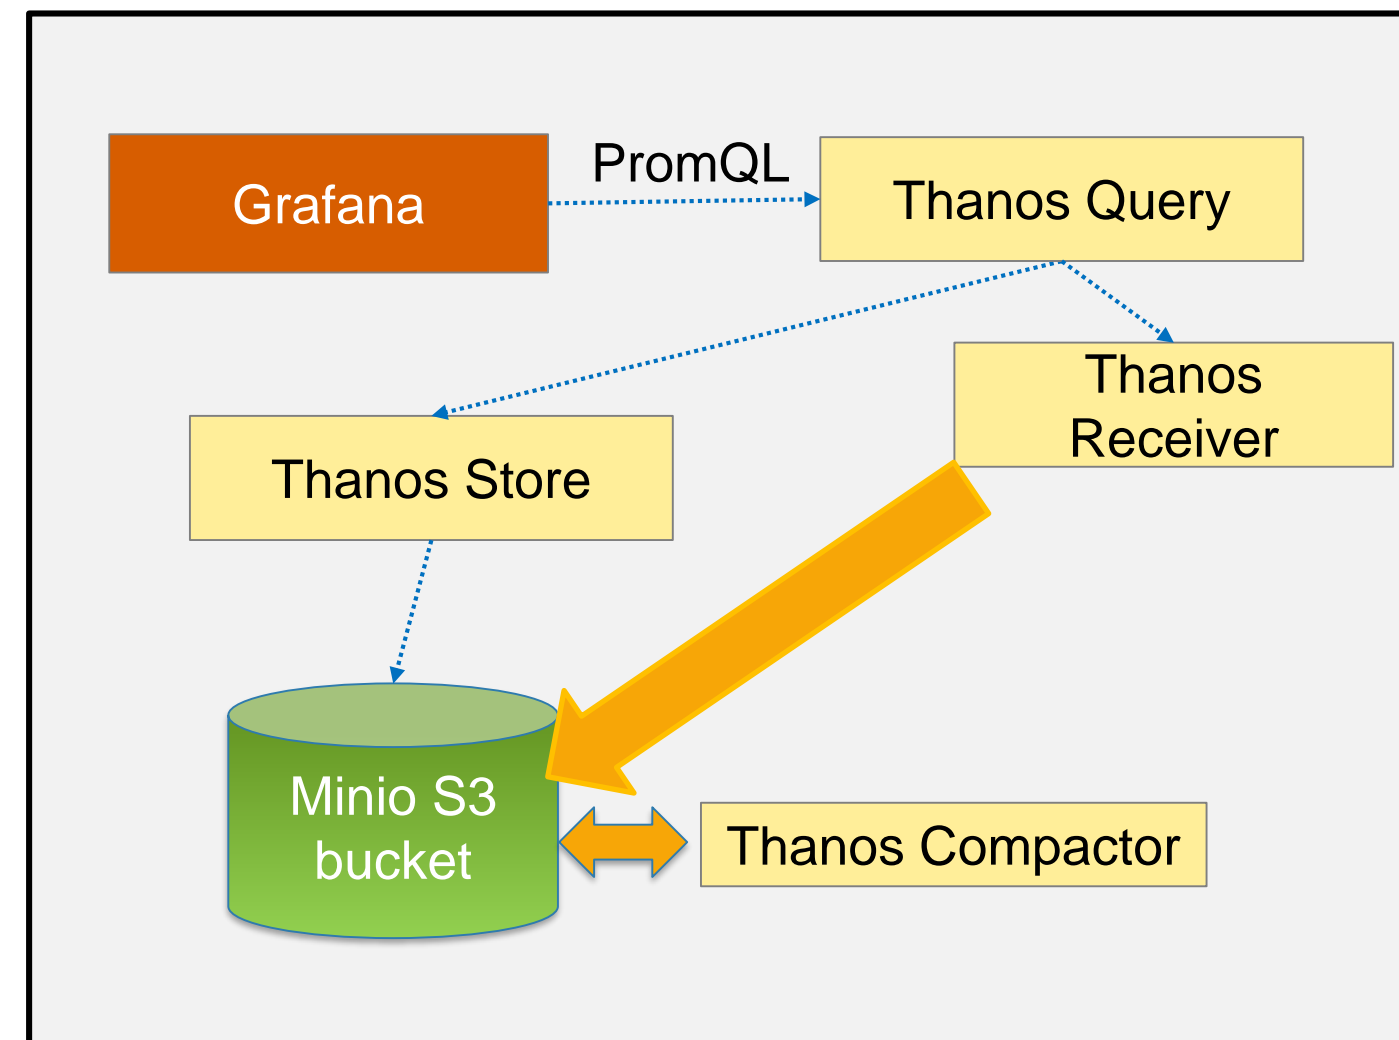
- Capacity Management needs aggregated node capacity, quota and resource usage metrics for several month and years.
- Prometheus has only one retention for all metrics
- Conflict:
 - Detailed metrics for post mortem analysis (days or upto a month)
 - Highly aggregated data for capacity management. E.g. one value of CPU Quota per per cluster and application_type (infrastructure vs. application)Metrics wanted for years.



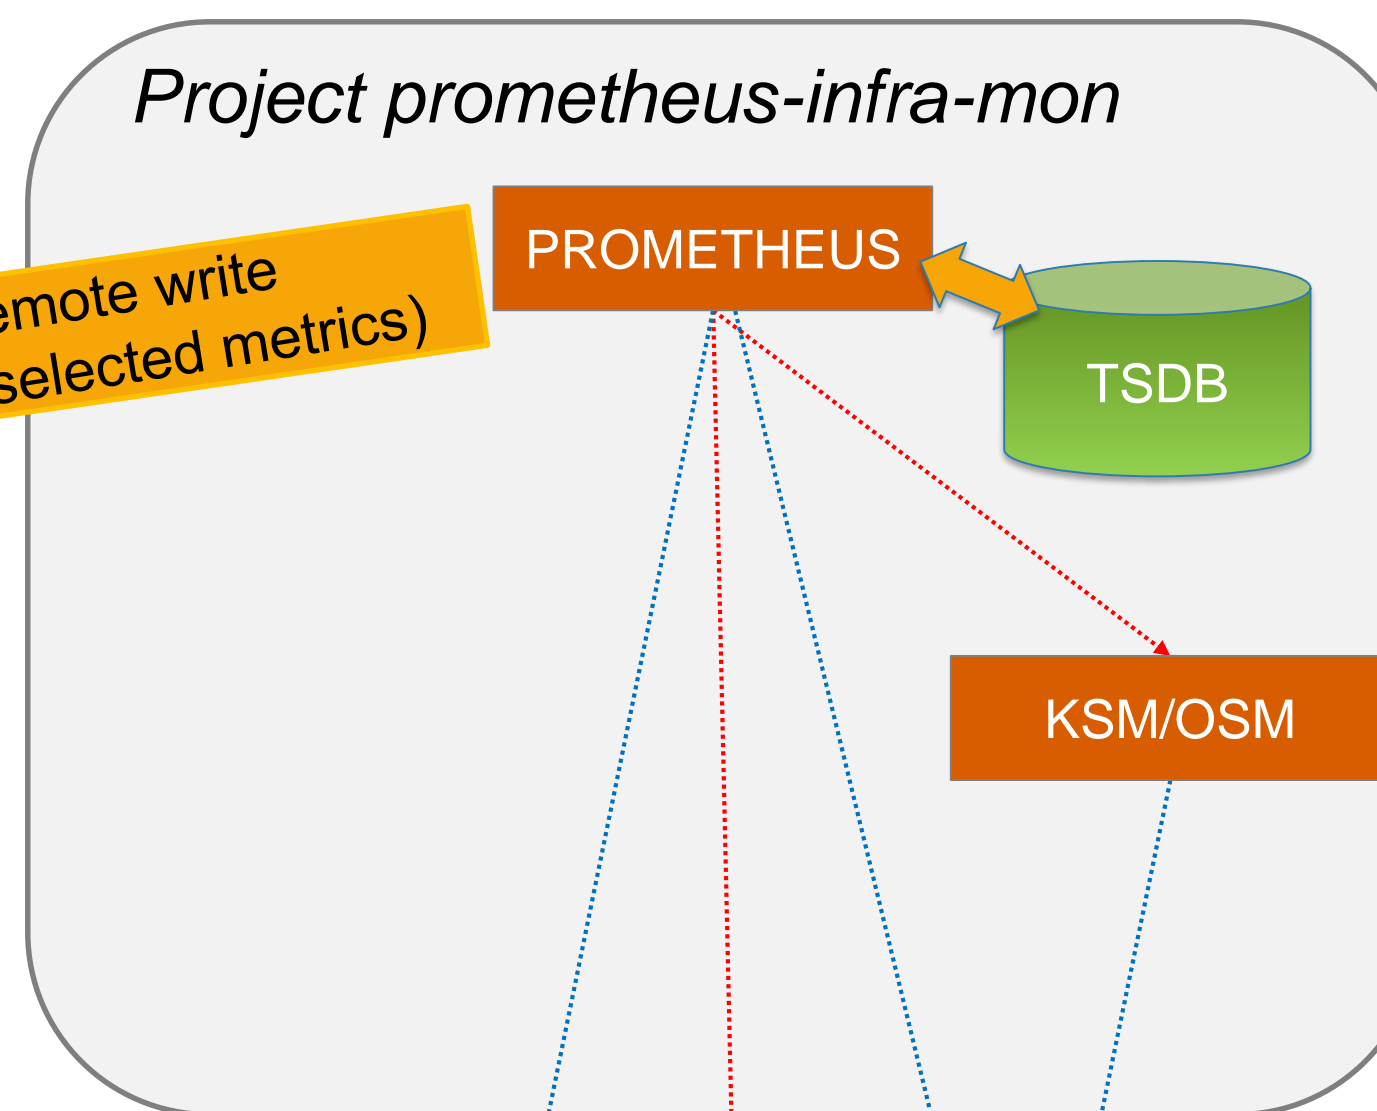
- Prometheus Remote write feature: selected metrics can be stored in an external time series database with longer retention policies.
- Thanos (<https://thanos.io/>) provides a remote write target and PromQL interface

Metric Collection and Storage Architecture

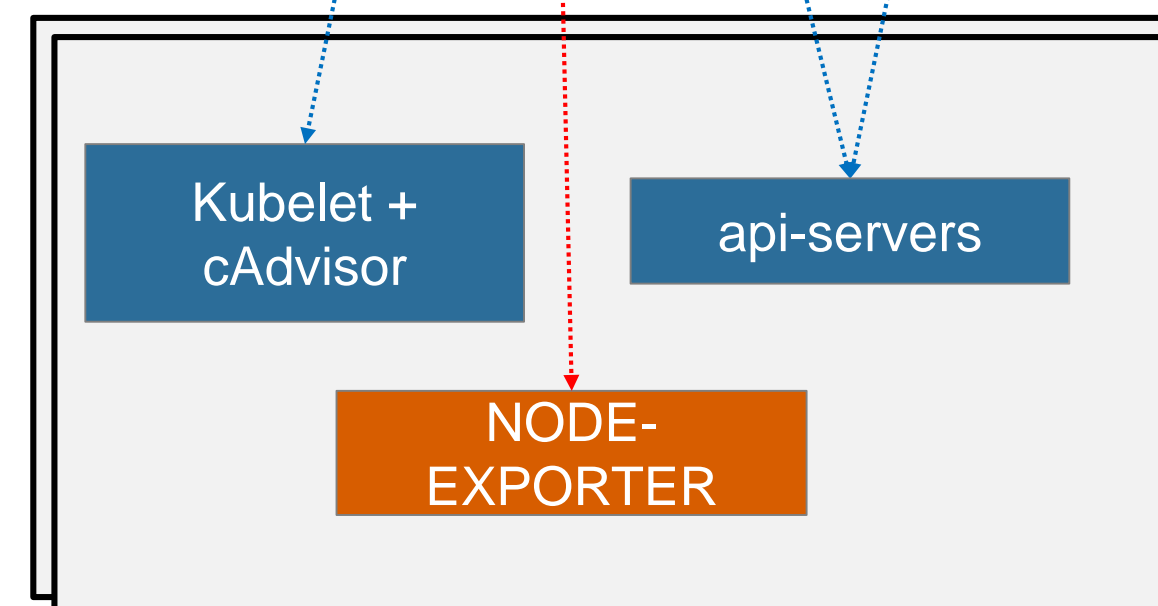
central storage and visualization



Kubernetes/OpenShift Cluster

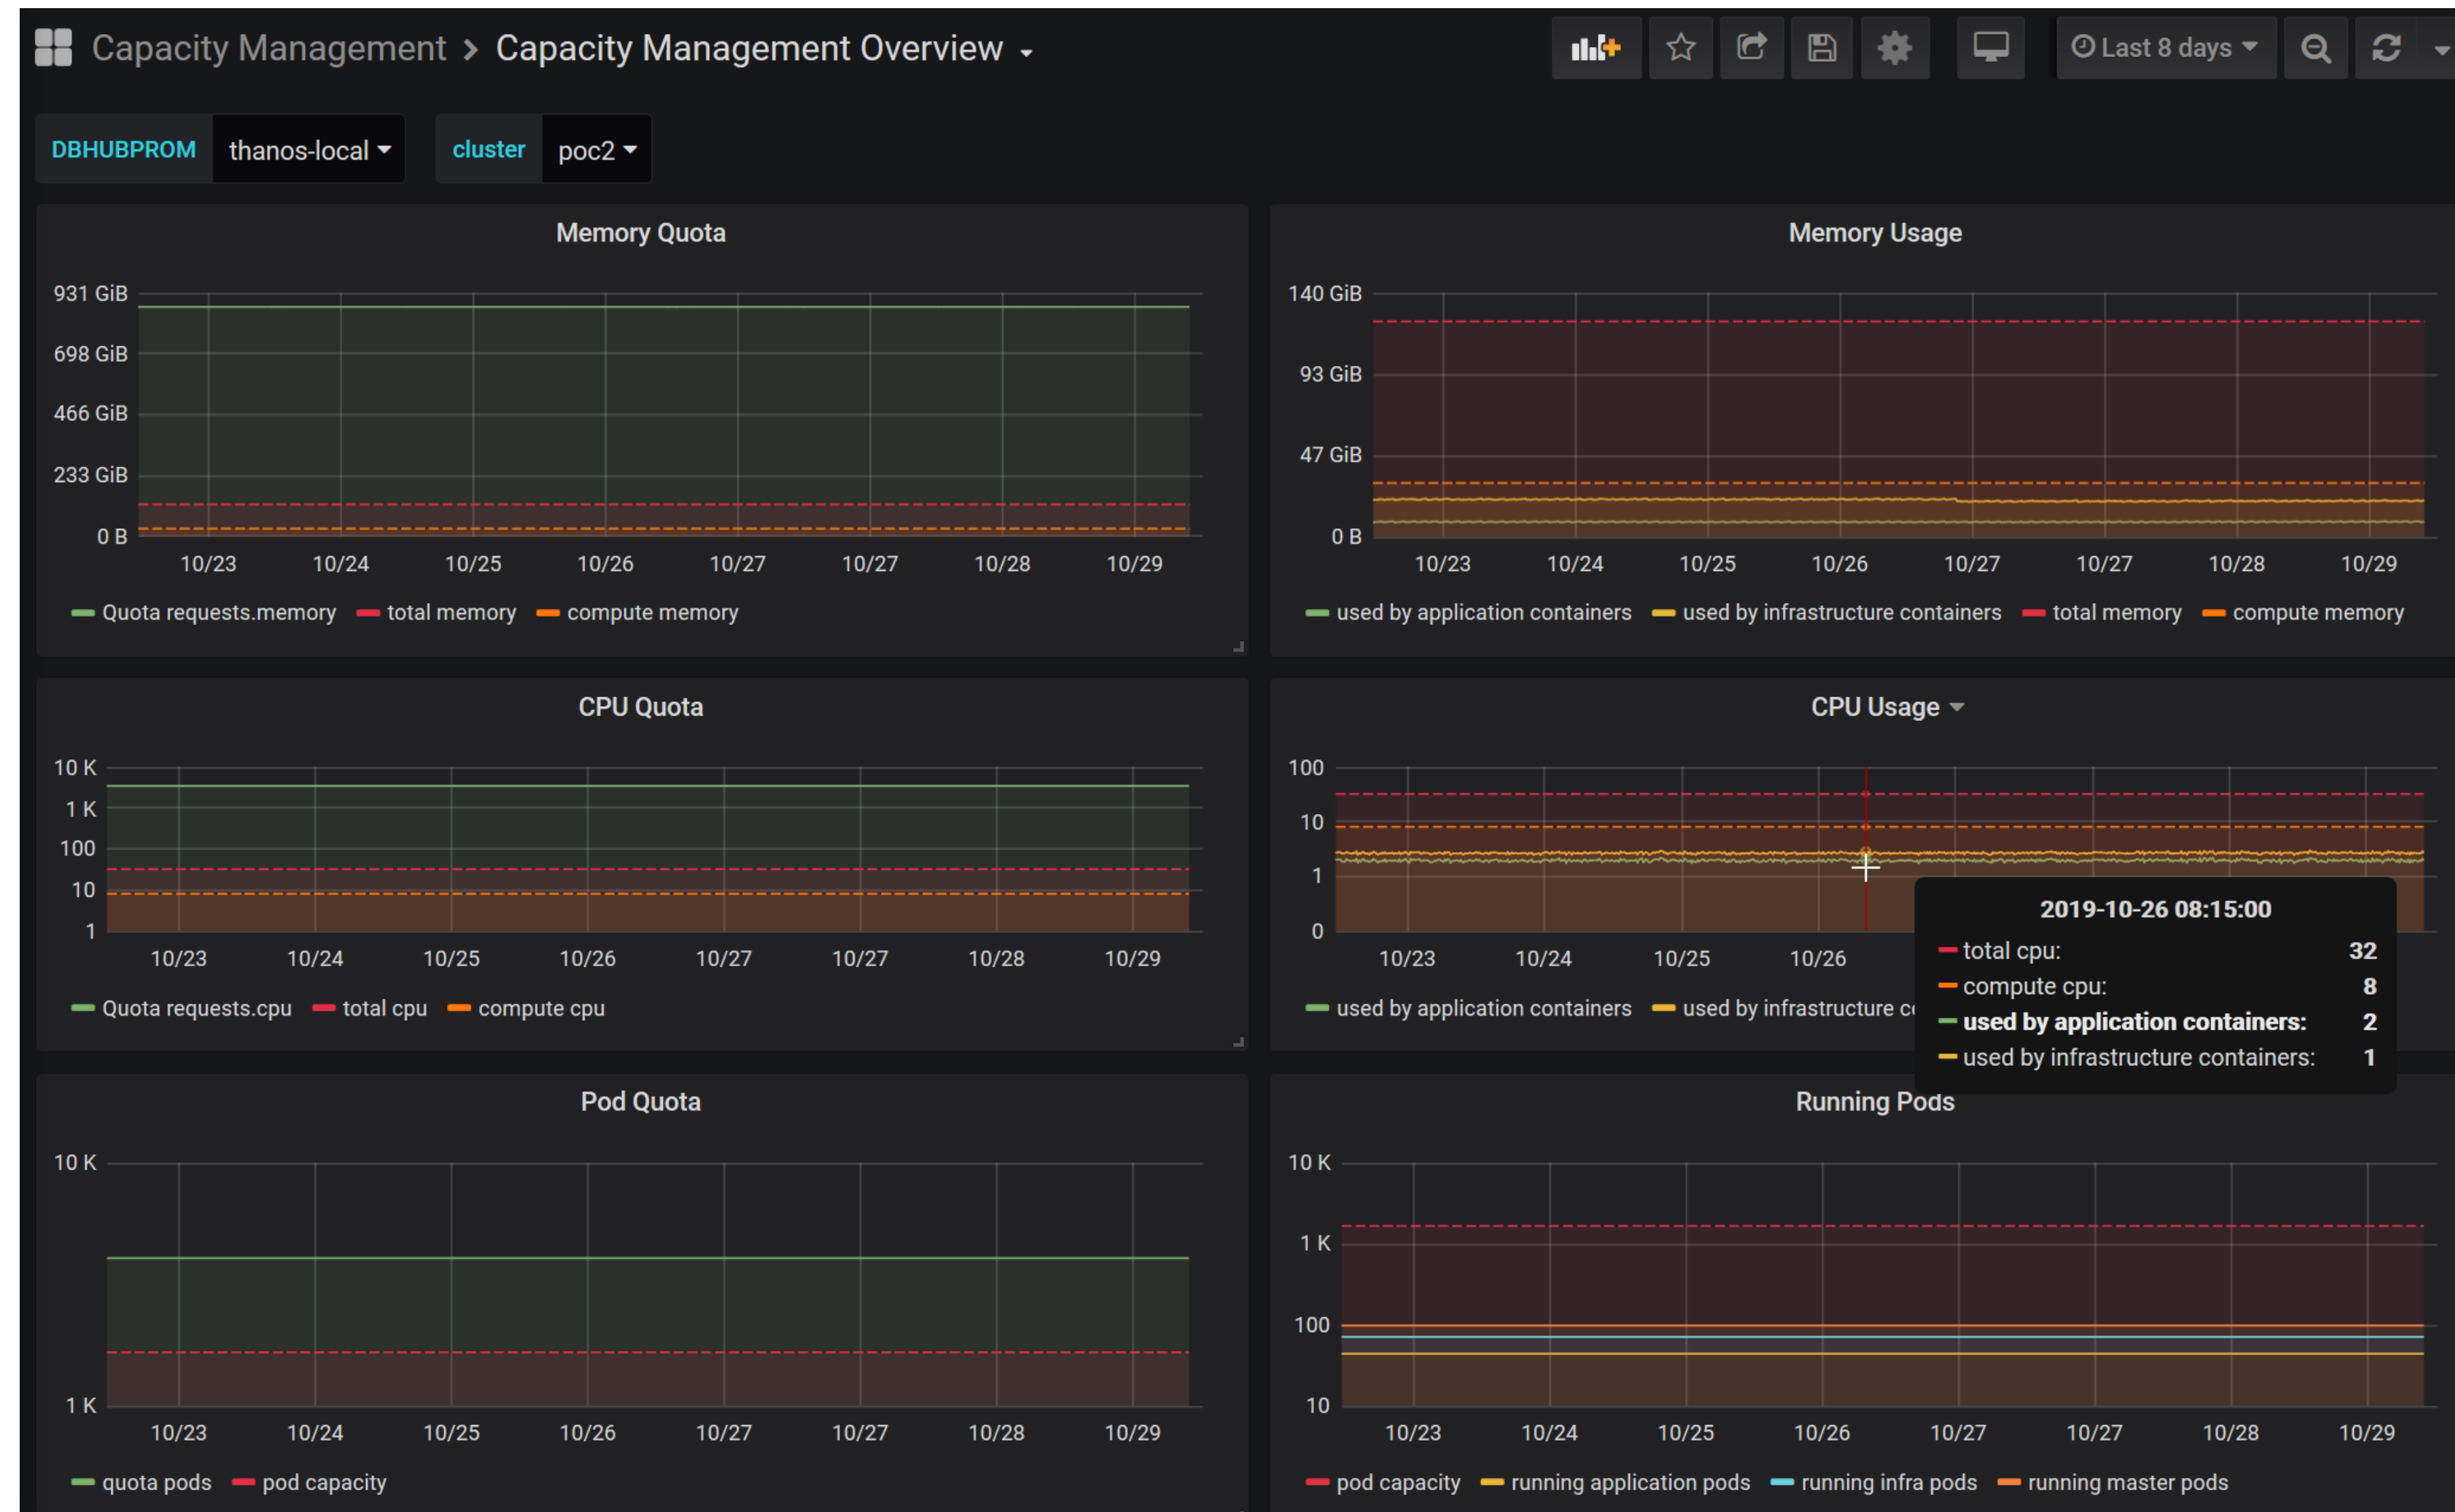


Nodes



- KSM (Kube state metrics)/OSM (OpenShift state metrics): for running Pods per Namespace, ResourceQuotas and ClusterResourceQuotas
- Node-Exporter for operating system metrics of resource usage

Dashboard: Capacity Management Overview



Computing Effective Quotas per Namespace

```
- record: namespace:kube_resourcequota:effective
  expr: min by(namespace,namespace_type,resource_base,type)
(label_replace({__name__=~"^(kube_resourcequota)$",namespace=~".+",
resource=~"(requests.)?(memory|cpu)",resource!~"limit.*",type="hard"},"resource_base","requests.$2",
"resource","(requests.)?(.+)"))
- record: namespace:kube_resourcequota:effective
  expr: max by(namespace,namespace_type,resource_base,type)
(label_replace({__name__=~"^(kube_resourcequota)$",namespace=~".+",
resource=~"(requests.)?(memory|cpu)",resource!~"limit.*",type="used"},"resource_base","requests.$2",
"resource","(requests.)?(.+)"))
...
- record: namespace:kube_resourcequota:effective
  expr: min by(namespace,namespace_type,resource_base,type)
(label_replace({__name__=~"^(kube_resourcequota)$",namespace=~".+",resource=~"limits.+","type="hard"}
,"resource_base","$1","resource","(.+)"))
- record: namespace:kube_resourcequota:effective
  expr: max by(namespace,namespace_type,resource_base,type)
(label_replace({__name__=~"^(kube_resourcequota)$",namespace=~".+",resource=~"limits.+","type="used"}
,"resource_base","$1","resource","(.+)"))
```



Thank you!



ConSol

Consulting & Solutions Software GmbH

St.-Cajetan-Straße 43

D-81669 München

Tel.: +49-89-45841-100

info@consol.de

www.consol.de

Twitter: @consol_de