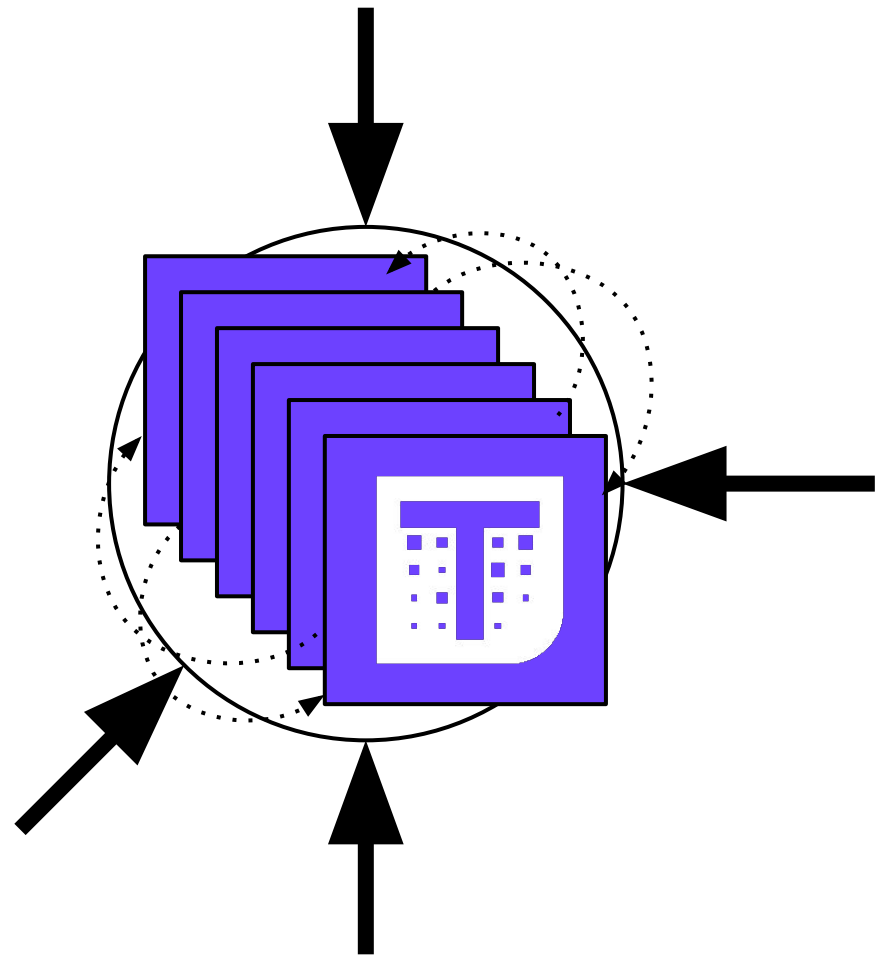


***Turn It Up to a Million***  
*Ingesting Millions of Metrics*  
*With Thanos Receive*

Lucas Servén Marín





Lucas Servén Marín



@squat



Thanos



# Related KubeCon Talks



## **Thanos: Cheap, Simple, and Scalable Prometheus**

- Giedrius Statkevičius, Adform & Matthias Loibl, Red Hat
- Wednesday, August 19
- 14:30 - 15:05 UTC

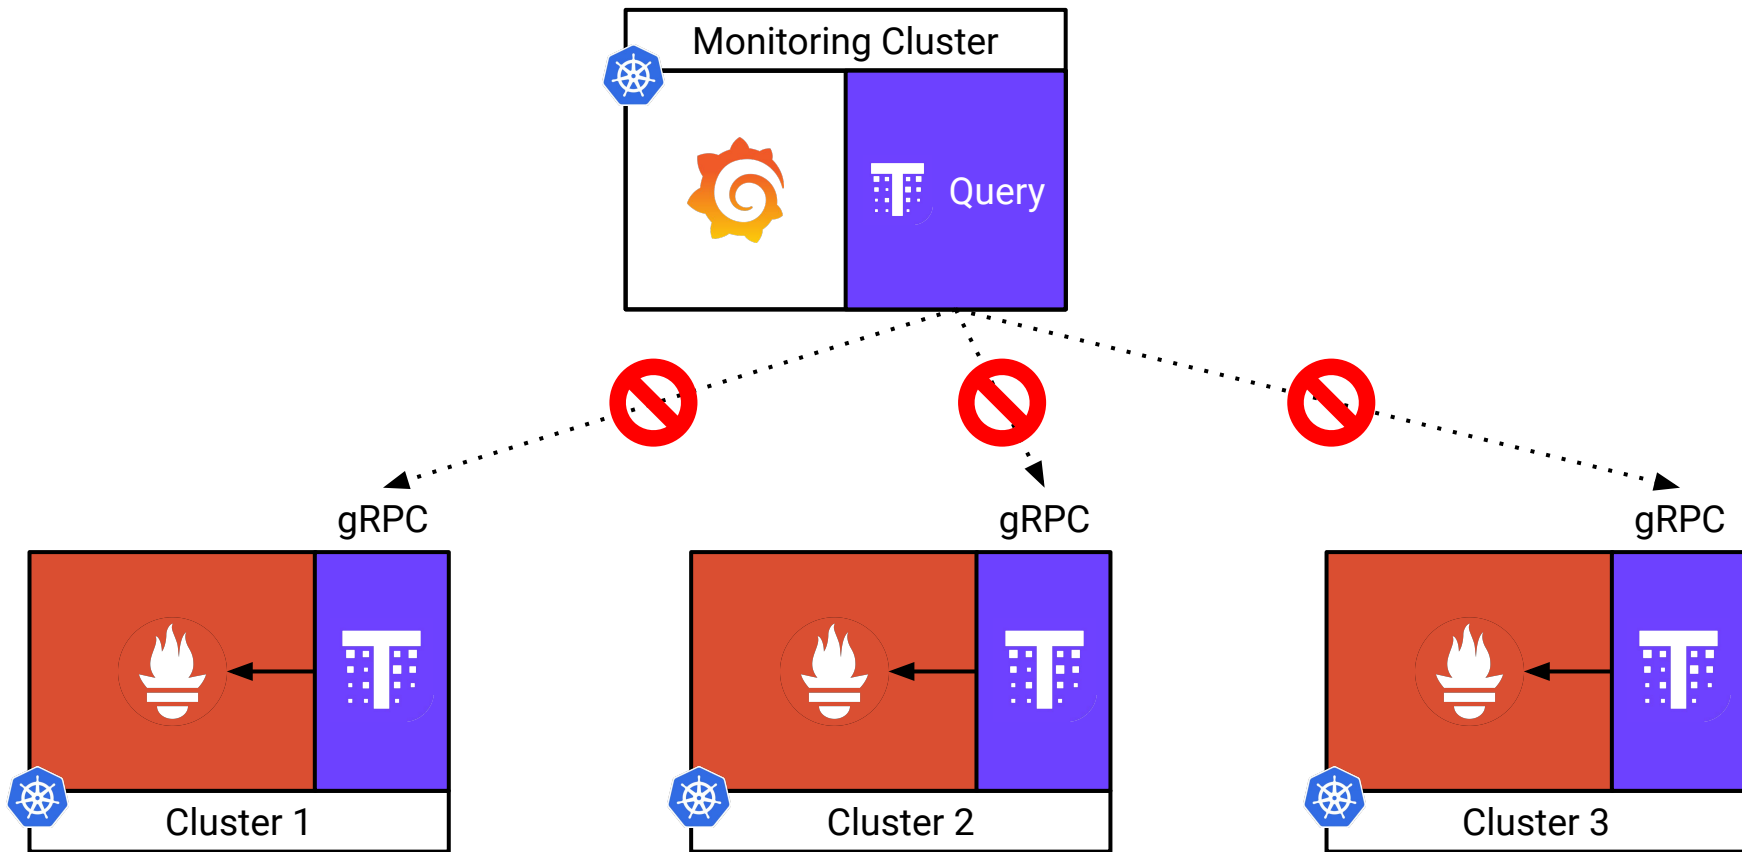


## **Scaling Prometheus: How We Got Some Thanos Into Cortex**

- Thor Hansen, HashiCorp & Marco Pracucci, Grafana Labs
- Tuesday, August 18
- 14:45 - 15:20 UTC




# Global View






# Remote Write API



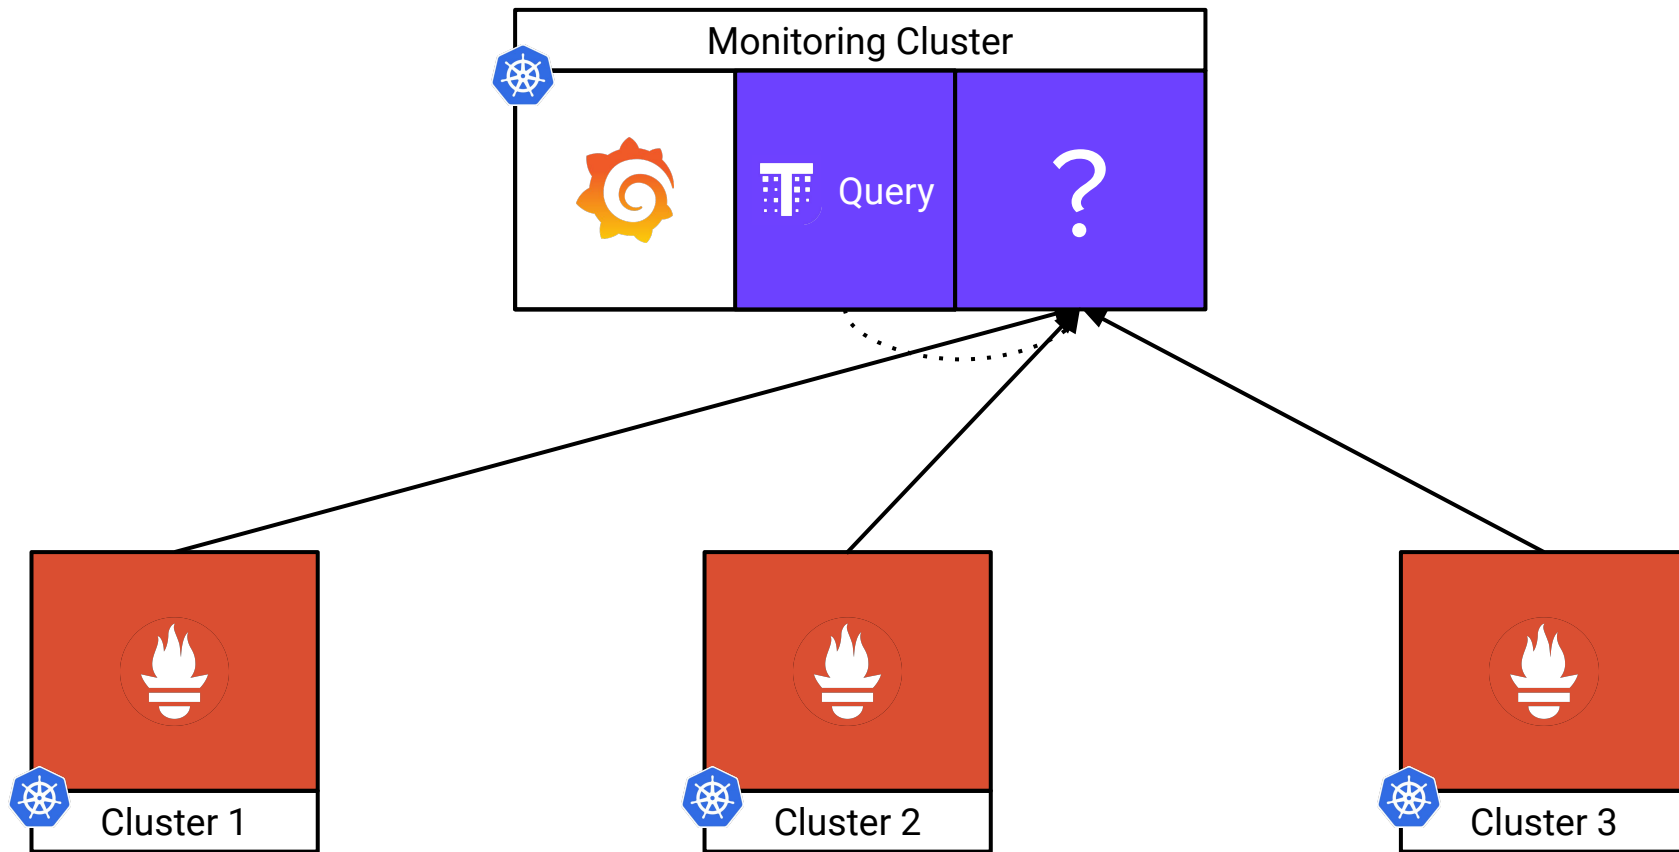


# Push Metrics






# Push Metrics





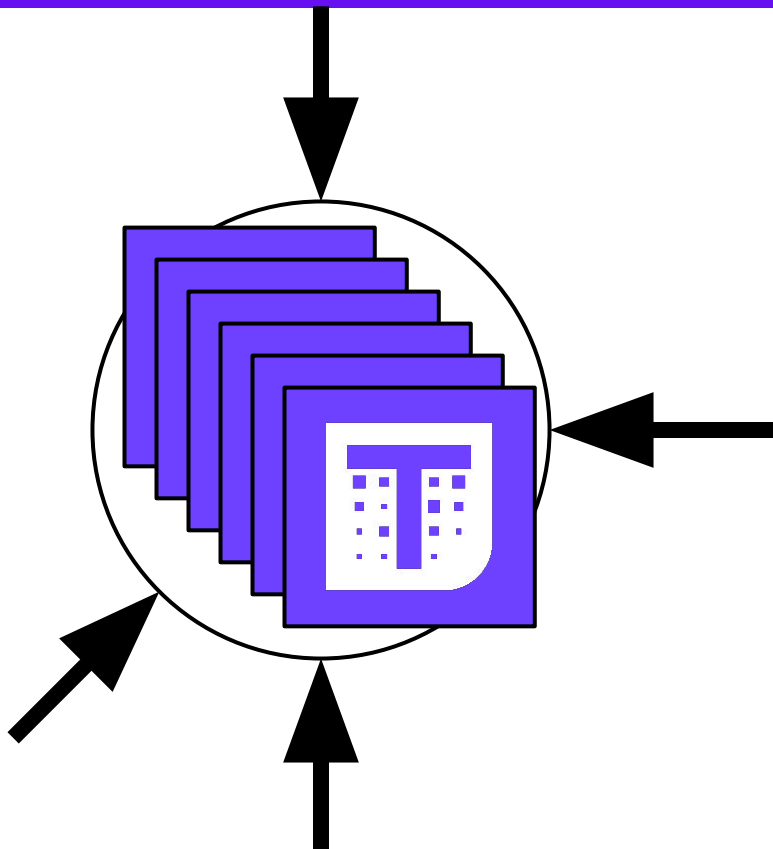


# Scalability






# Hash Ring








# Hash Ring



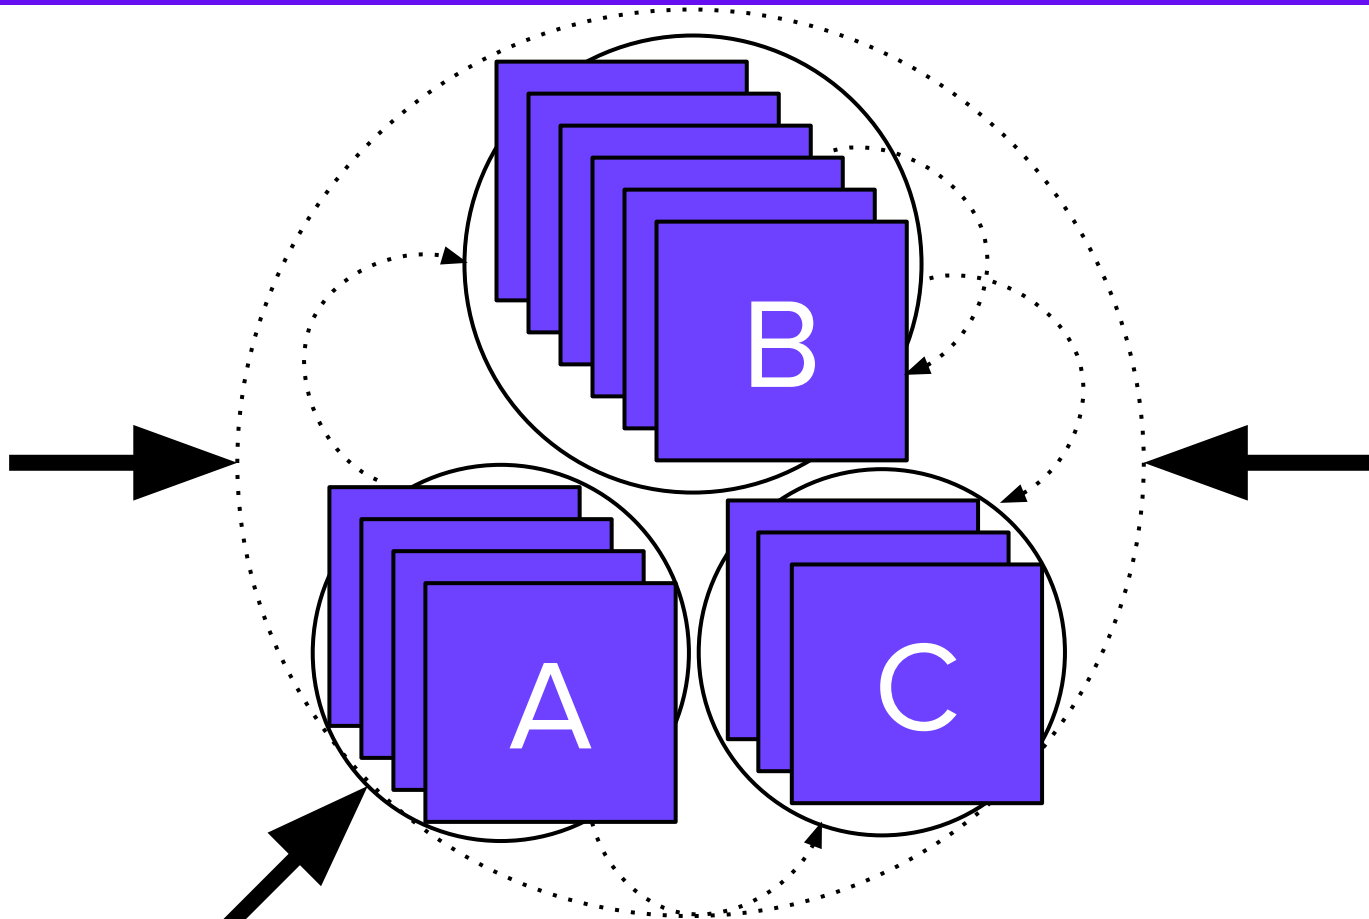


# Multitenant



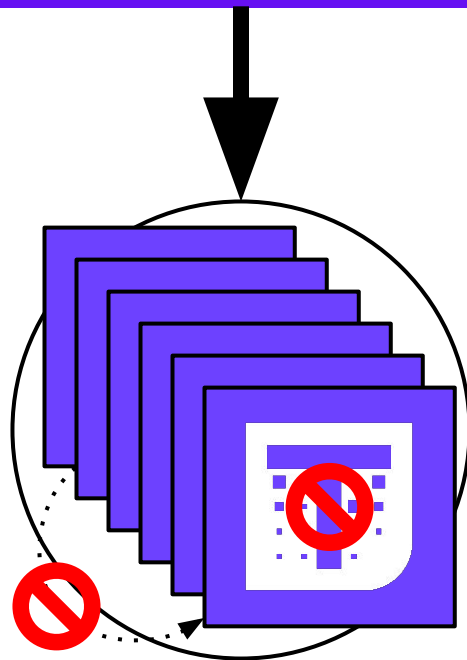


# Multitenant



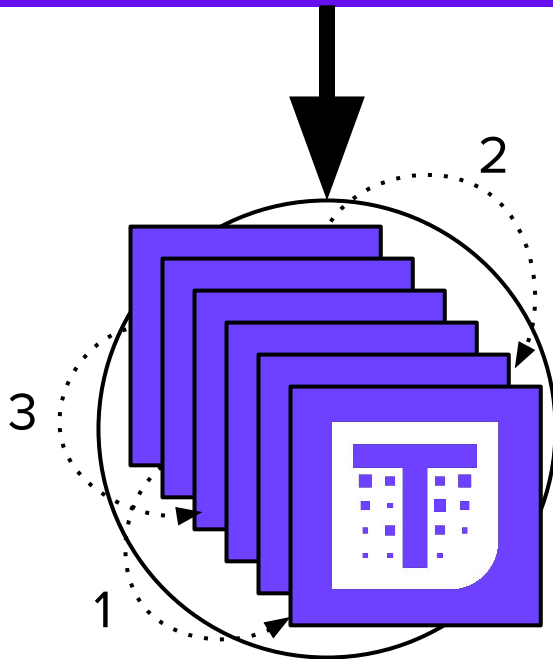


# Availability






# Replication





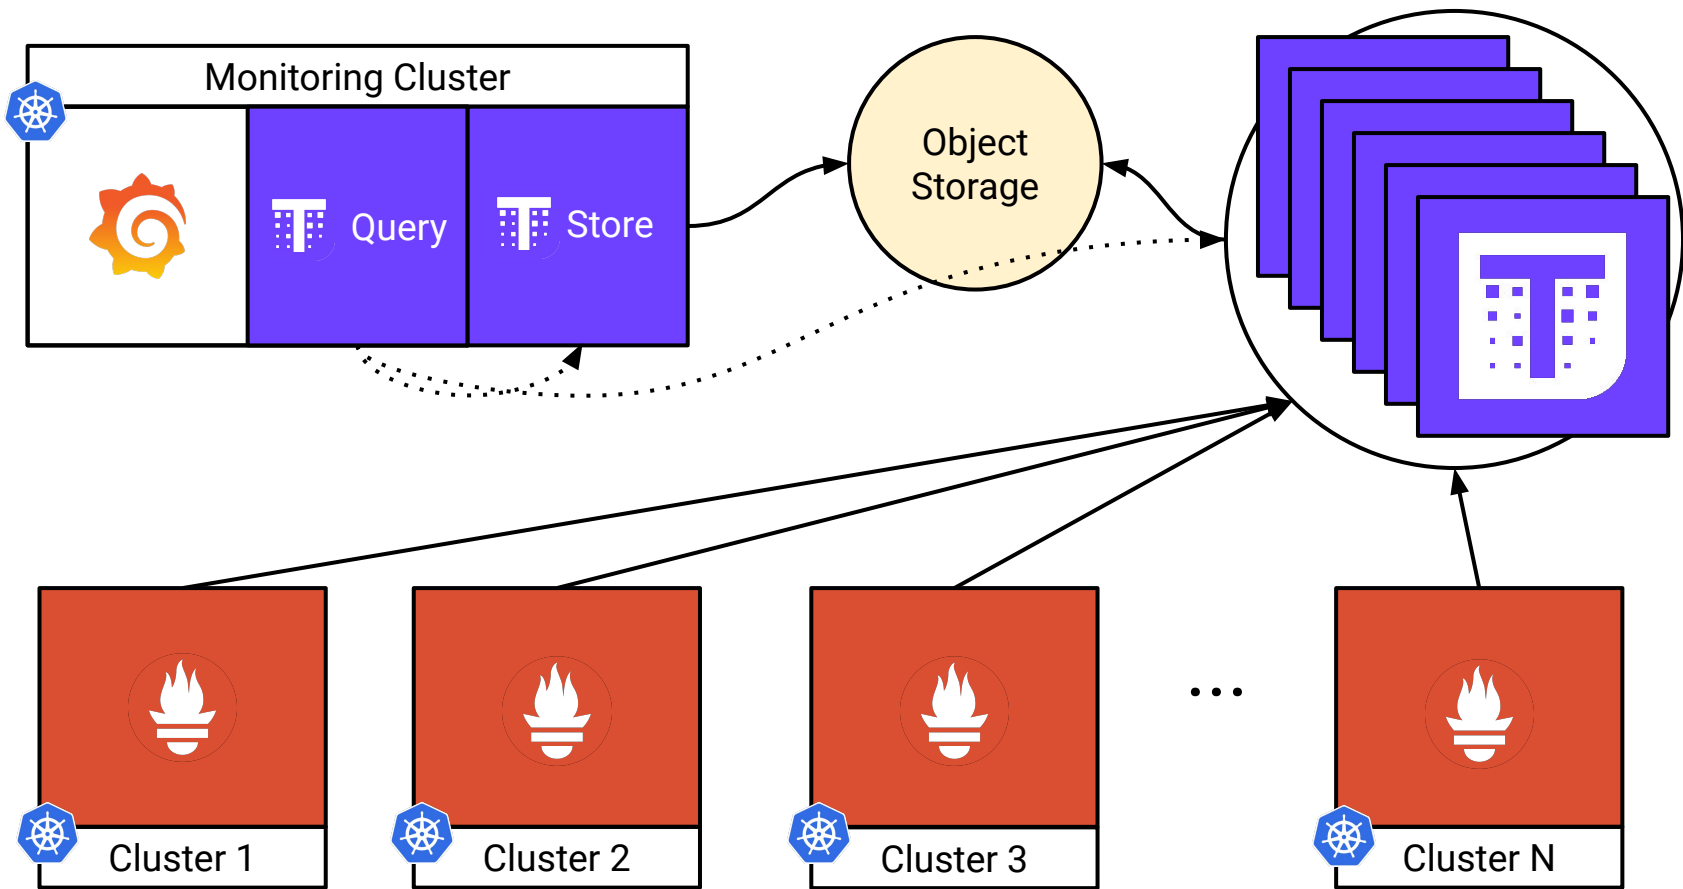
# Scalability







# Scalability





Running



+





# hashrings.json

```
[
  {
    "hashring": "default",
    "tenants": [
      "ci",
      "telemetry"
    ],
    "endpoints": [
      "127.0.0.1:1337",
      "127.0.0.1:1338",
      "127.0.0.1:1339"
    ]
  }
]
```

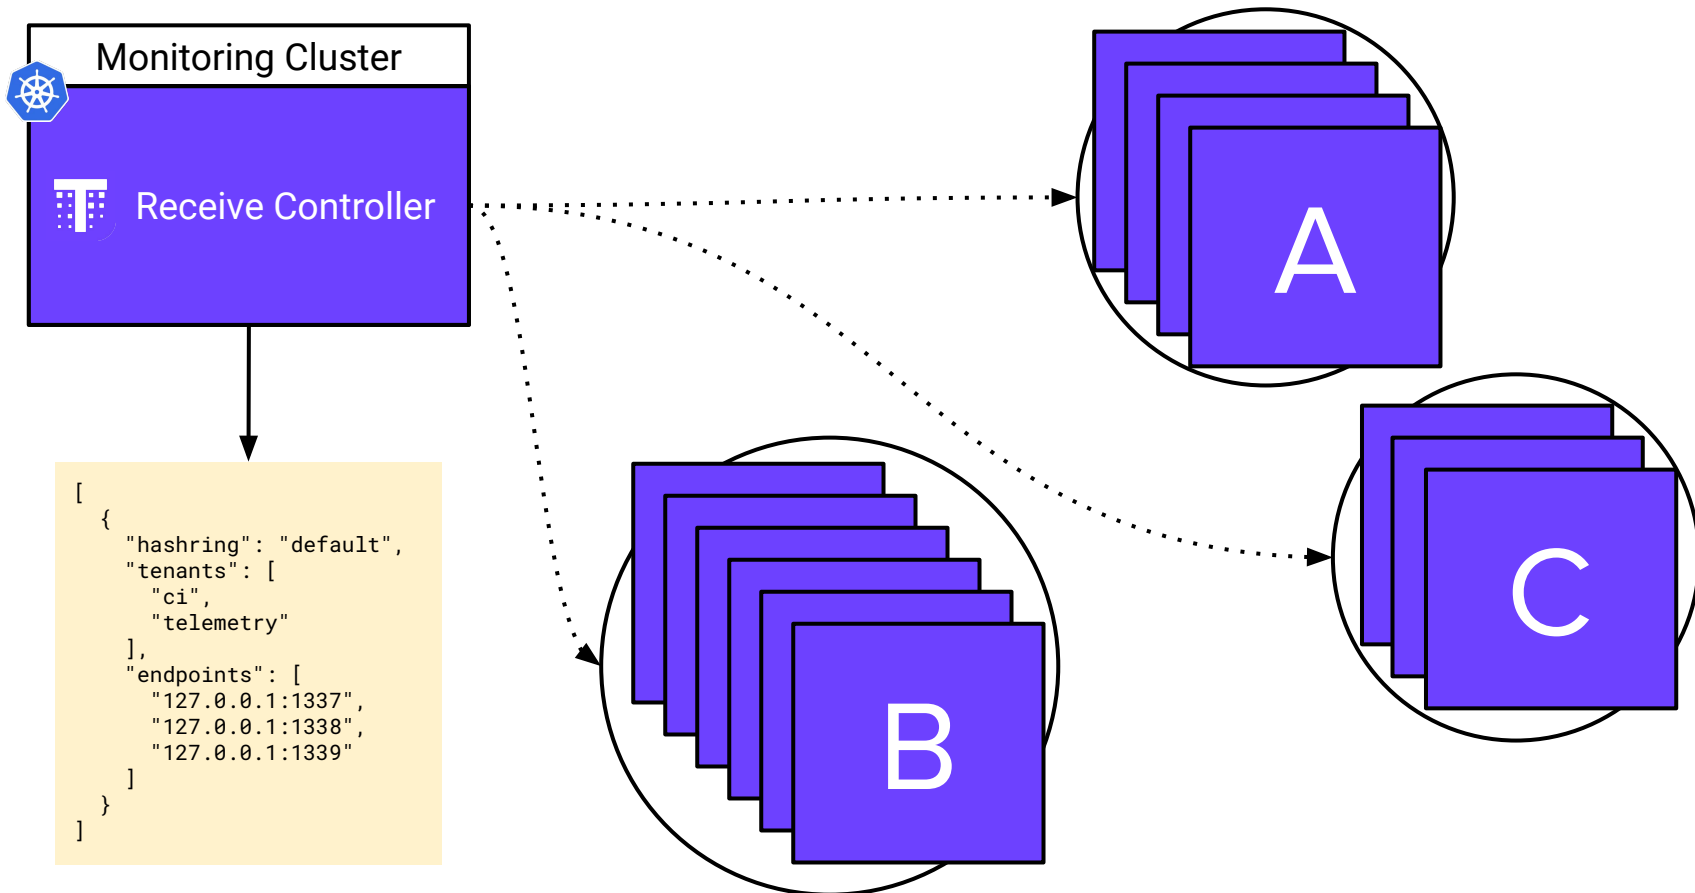


# Thanos Receive Controller

[github.com/observatorium/thanos-receive-controller](https://github.com/observatorium/thanos-receive-controller)




# Thanos Receive Controller





# Deployment #1





# Deployment #1

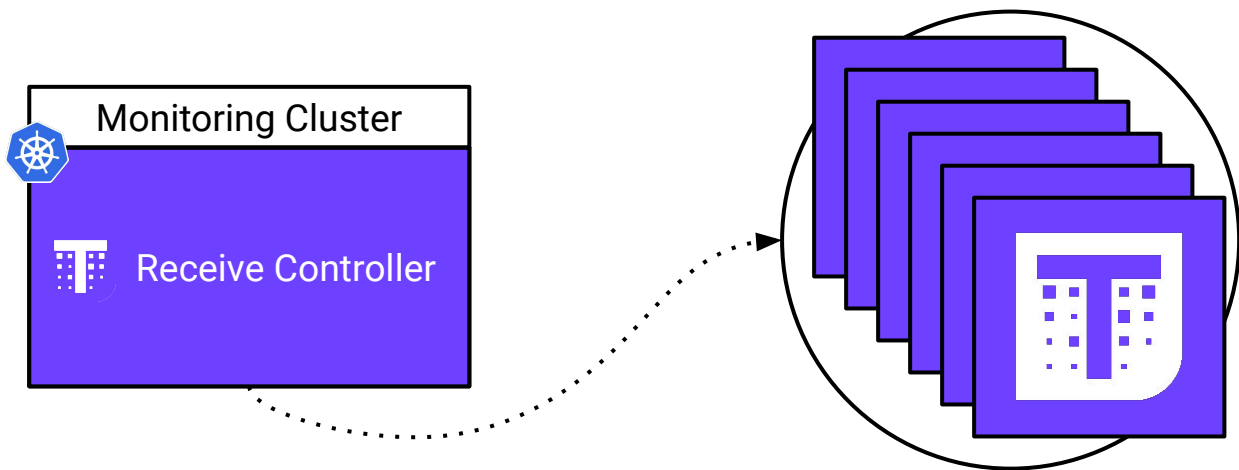


Monitoring Cluster

```
kubectl scale sts thanos-receive --replicas=20
```



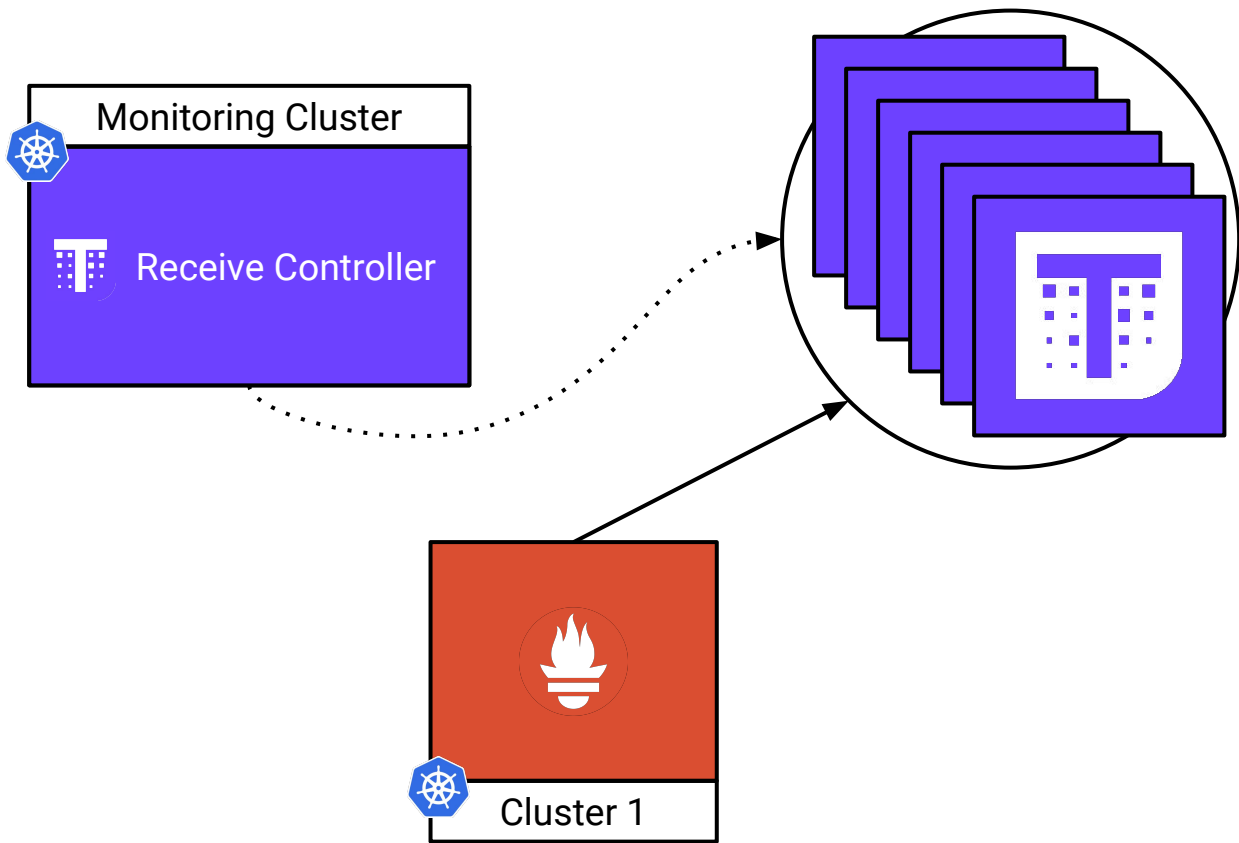
# Deployment #1







# Deployment #1

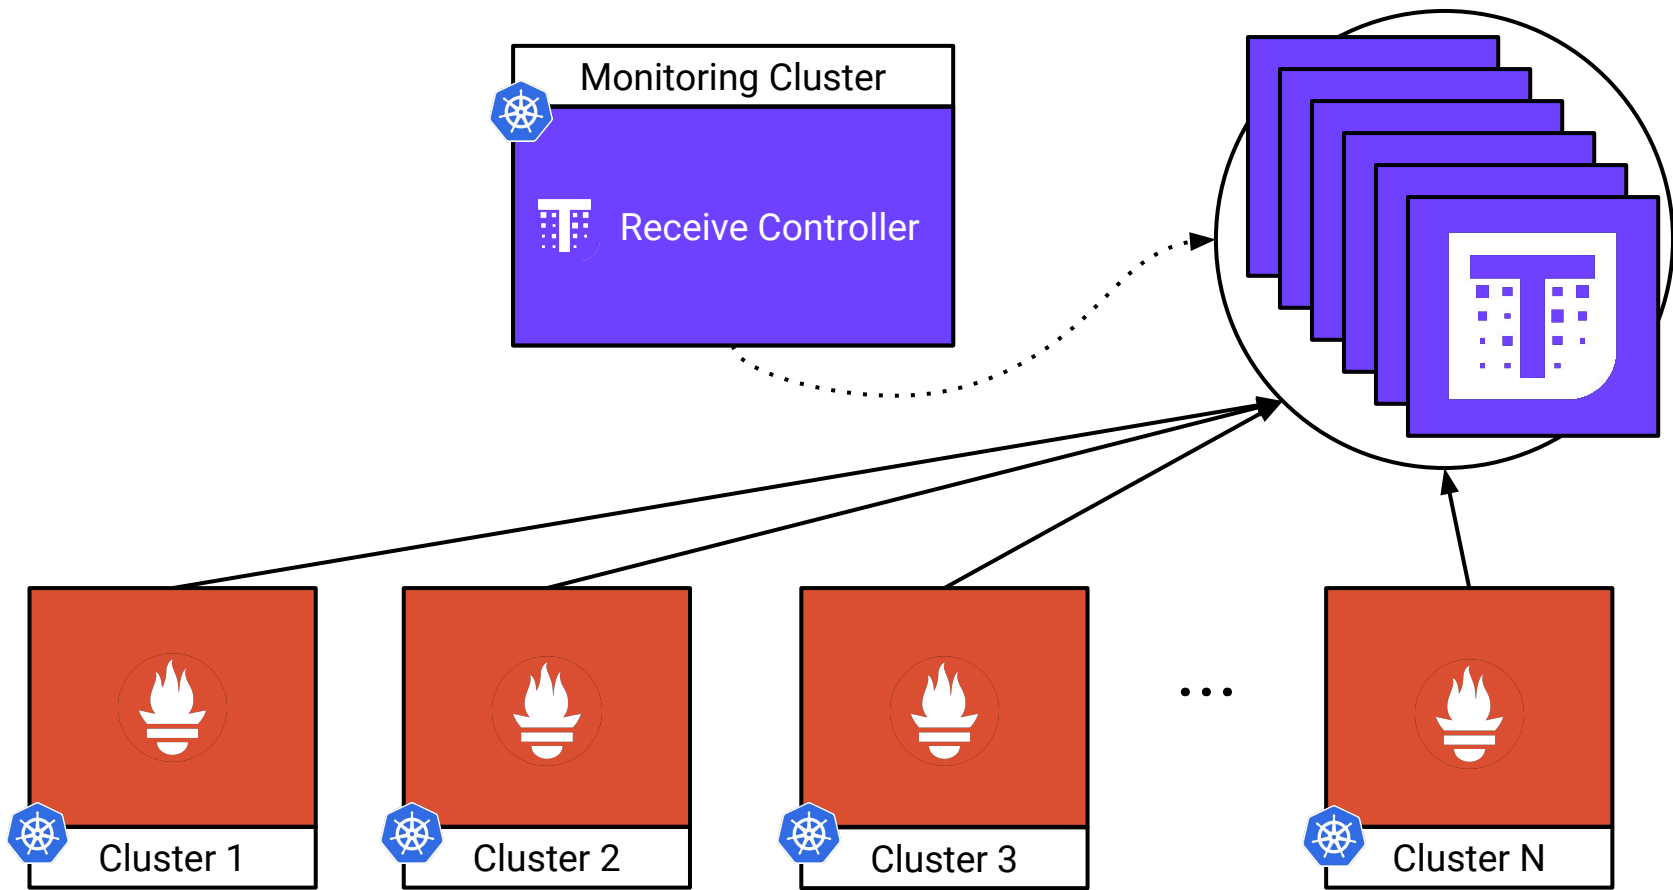


Demo

[github.com/squat/kubeconeu2020](https://github.com/squat/kubeconeu2020)



# Tiny But Mighty





# Active Time Series

```
sum(prometheus_tsdb_head_series{job="observatorium-thanos-receive-default"})
```


Load time: 275ms  
Resolution: 4838s  
Total time series: 1

Execute

- insert metric at cursor - ▾

Graph Console

- 2w + ◀ Until ▶ Res. (s)  stacked





# Samples/Second

```
sum(rate(prometheus_tsdb_head_samples_appended_total{job="observatorium-thanos-receive-default"}[5m]))
```


Load time: 169ms  
Resolution: 9676s  
Total time series: 1

Execute

- insert metric at cursor -

Graph Console

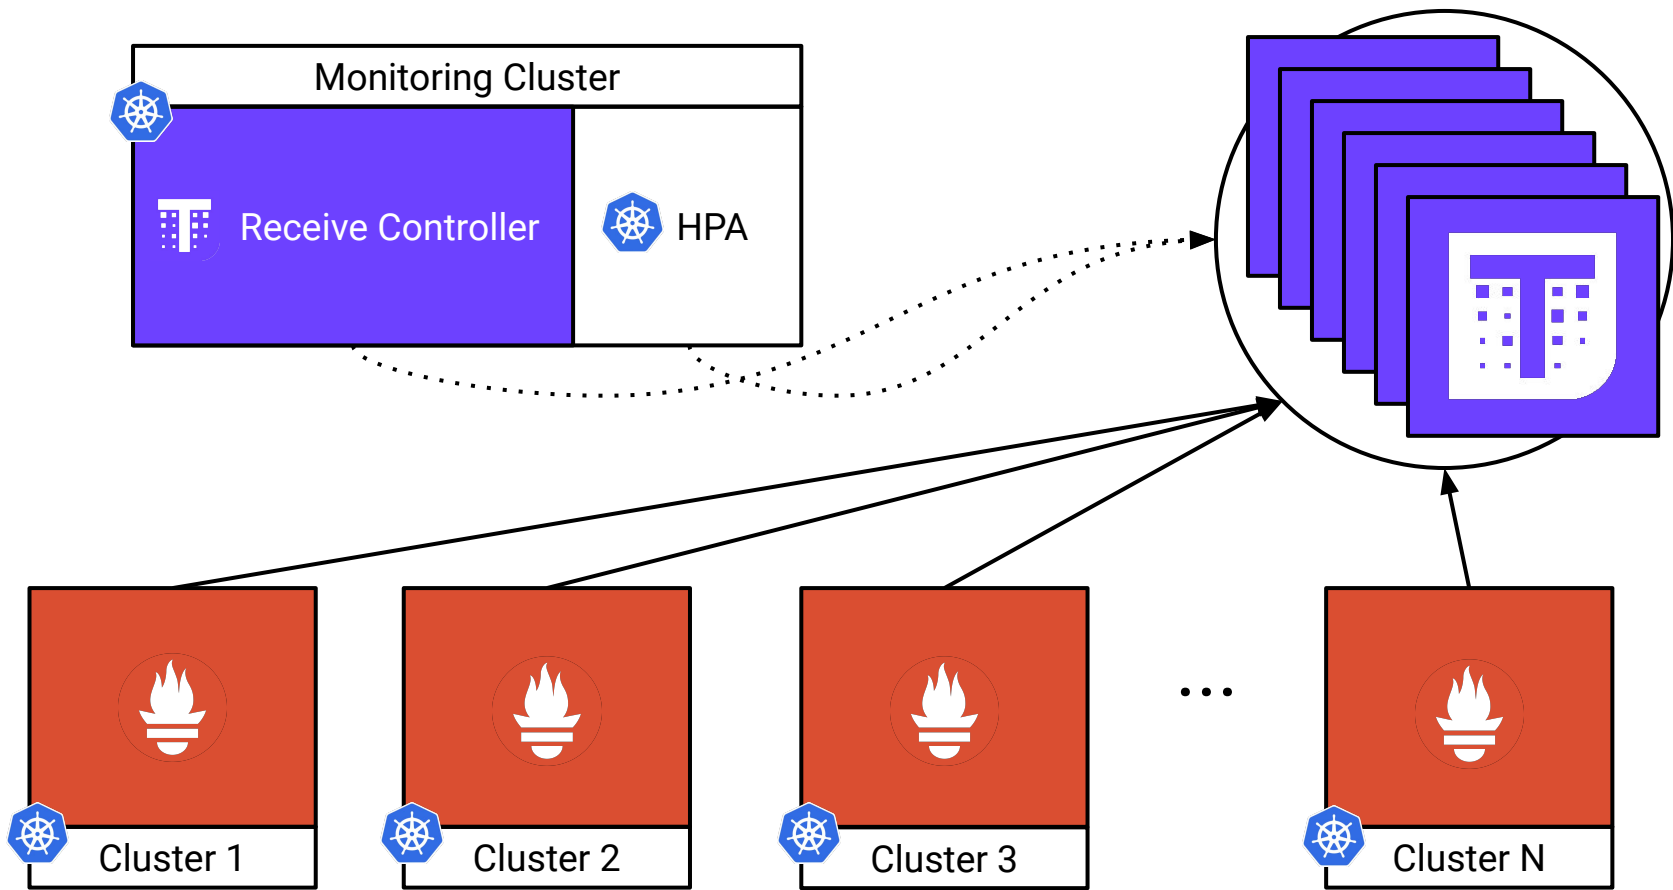
- 4w + << Until >> Res. (s)  stacked



✓ 0



# Autoscaling



Demo

[github.com/squat/kubeconeu2020](https://github.com/squat/kubeconeu2020)




# Autoscaling

WHAT HAPPENED???





# Unstable Equilibrium





# Problems

`spec.replicas != #pods`



# Problems

```
[
  {
    "hashring": "default",
    "tenants": [
      "ci",
      "telemetry"
    ],
    "endpoints": [
      "127.0.0.1:1337",
      "127.0.0.1:1338",
      "127.0.0.1:1339"
    ]
  }
]
```




```
[
  {
    "hashring": "default",
    "tenants": [
      "ci",
      "telemetry"
    ],
    "endpoints": [
      "127.0.0.1:1337",
      "127.0.0.1:1338",
      "127.0.0.1:1339",
      "127.0.0.1:1340",
      "127.0.0.1:1341",
      "127.0.0.1:1342",
      "127.0.0.1:1343"
    ]
  }
]
```

slow + split-brain + downtime



# Problems



Request Multiplication



# Solutions

```
[
  {
    "hashring": "default",
    "tenants": [
      "ci",
      "telemetry"
    ],
    "endpoints": [
      "127.0.0.1:1337",
      "127.0.0.1:1338",
      "127.0.0.1:1339"
    ]
  }
]
```



```
[
  {
    "hashring": "default",
    "tenants": [
      "ci",
      "telemetry"
    ],
    "endpoints": [
      "127.0.0.1:1337",
      "127.0.0.1:1338",
      "127.0.0.1:1339",
      "127.0.0.1:1340",
      "127.0.0.1:1341",
      "127.0.0.1:1342",
      "127.0.0.1:1343"
    ]
  }
]
```

fast + no downtime



Thank You!

*<https://thanos.io>*



1. Is the Thanos receiver ready for production?
2. What causes more trouble in the hashing, time series, samples, large batches, small batches?
3. Is it better to have a large hashing with many small replicas or a small hashing with many large replicas?
4. What is the best metric to use for scaling?
5. How is the Thanos receiver different from Cortex?