

A High-Schooler's Guide to Kubernetes Network Observability

Drew Ripberger



Who am I?



Virtual

North America 2020

- Graduated high school in 2020
- Ohio State University computer science
- Science fair: it's not all baking soda volcanoes
- Raft brought me to etcd which introduced me to Kubernetes

Enter Nirmata



Virtual

- Started looking of part time summer work in the winter
- Honestly wasn't expecting to get a job in software
- Applied to Nirmata on a whim because they listed Golang

The Project: *kube-netc*



Virtual

North America 2020

- An open source, eBPF based Kubernetes network monitor deployed as a DaemonSet
- Collect node level networking statistics using eBPF
- Expose these stats using a Prometheus endpoint
- Deploy our monitor as a pod on each node in the cluster
- What was the goal of kube-netc?

github.com/nirmata/kube-netc

The Goal of *kube-netc*



Virtual

- All encompassing networking statistics at the TCP layer
- Easily interpretable statistics to be visualized or process upstream
- Simplicity

Map



KubeCon



CloudNativeCon

North America 2020

Virtual

1. Monitoring using eBPF
2. Tracking our stats with Prometheus
3. DaemonSets with Kubernetes
4. kube-netc: tying it all together
5. Demo



Prometheus



kubernetes

eBPF



Virtual

North America 2020



eBPF Crash course



North America 2020

- eBPF: extended Berkeley Packet Filter
- Allows running custom code in the Linux kernel
- At its most basic level, it's a bytecode
- Sandboxed environment for interaction with the Linux kernel
- Implications in networking, observability, security and much more

eBPF Crash course



Virtual

North America 2020

- How do we implement eBPF?
- Unless you want to write eBPF byte code, probably a C subset BCC or the Go bindings: github.com/iovisor/bcc and github.com/iovisor/gobpf
- BCC lets you write eBPF byte code as a subset of C using LLVM as a backend
- For a beginner this is a start!

Open-Source eBPF Projects



Virtual

North America 2020

Some of the amazing projects being developed

- `bcc` (github.com/iovisor/bcc)
- `bpftrace`
- Cilium
 - Hubble
- Falco
- `kubectl-trace`

Want to learn more?



Virtual

- ebpf.io
- Beyond the Buzzword: BPF's Unexpected Role in Kubernetes
 - Andrew Randall & Alban Crequy, Kinvolk
 - Thursday November 19, 2020 5:40pm - 6:15pm EST

Implementing Prometheus



Virtual

North America 2020



Prometheus

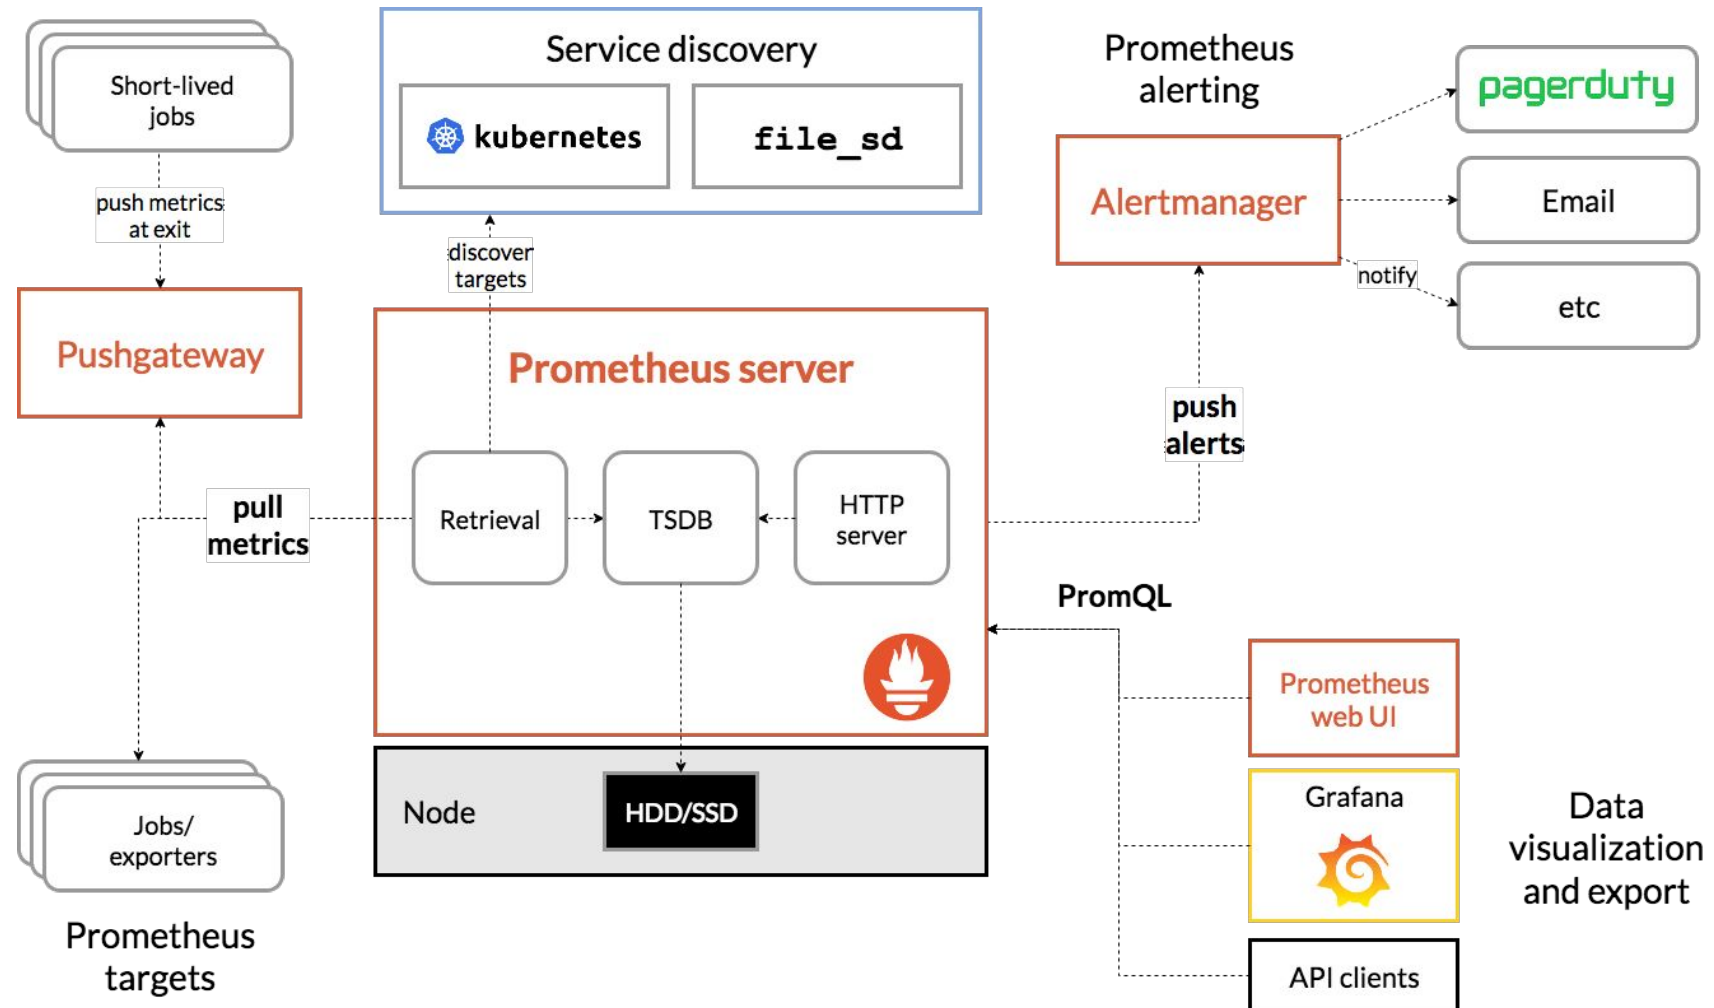
Basics of Prometheus



North America 2020

- Pull vs. Push: applications query exporters to pull data
- Handle time series data
- Easily monitor your services
- Great opportunities for visualization

The Landscape



Grafana Visualization



KubeCon



CloudNativeCon

North America 2020

Virtual

- Easy method to quickly visualize stats
- Customizable dashboards to display your data
- PromQL queries for quick processing



<https://prometheus.io/docs/visualization/grafana/>

Kubernetes




Virtual

North America 2020



kubernetes

What is Kubernetes?



What is Kubernetes?



KubeCon



CloudNativeCon

North America 2020

Virtual




Illustration courtesy of Reagan Whipkey

The DaemonSet



Virtual

North America 2020

```
apiVersion: apps/v1
kind: DaemonSet # specify that this is a DaemonSet
metadata:
  name: fluentd-elasticsearch
  namespace: kube-system
  labels:
    k8s-app: fluentd-logging
spec:
  template:
    ...
    spec:
      containers:
        - name: fluentd-elasticsearch # this container gets deployed on each node as a pod
          image: quay.io/fluentd_elasticsearch/fluentd:v2.5.2
      ...
```

github.com/nirmata/kube-netc

- Basic Install Instructions
- Grafana dashboard
- Great start for new contributors!
- bit.ly/netcblogpost

README.md

kube-netc: A Kubernetes eBPF network monitor

build passing go report A+

kube-netc (pronounced *kube-net-see*) is a Kubernetes network monitor built using eBPF

Getting Started

To test the current capabilities of **kube-netc**, this guide will walk you through viewing the network statistics of your nodes.

Install kube-netc

First, install the daemon set using the `install.yaml`:

```
kubectl apply -f https://github.com/nirmata/kube-netc/raw/master/config/install.yaml
```

View results

This will start the **kube-netc** DaemonSet on your cluster and setup the required roles. Then, we get the name of the kube-netc pod:

```
kubectl get pods | grep kube-netc
```

For example, my **kube-netc** pod is:

```
kube-netc-j56cx
```

In a new terminal, we port-forward the port of our pod so we can access it with `curl` outside the cluster with:

Great Starter Issues!



North America 2020

☐ 5 Open ✓ 10 Closed

Author ▼

Label ▼

Projects ▼

Milestones ▼

Assignee ▼

Sort ▼

☐ create helm chart that includes kube-netc + Prometheus + Grafana enhancement

#20 opened on Jul 30 by drewrip



☐ additional metrics for key statistics enhancement

#19 opened on Jul 28 by drewrip

☐ narrow metrics by namespace enhancement

#18 opened on Jul 27 by drewrip

☐ Metrics incorrectly being reported for pods that node have any ports bug

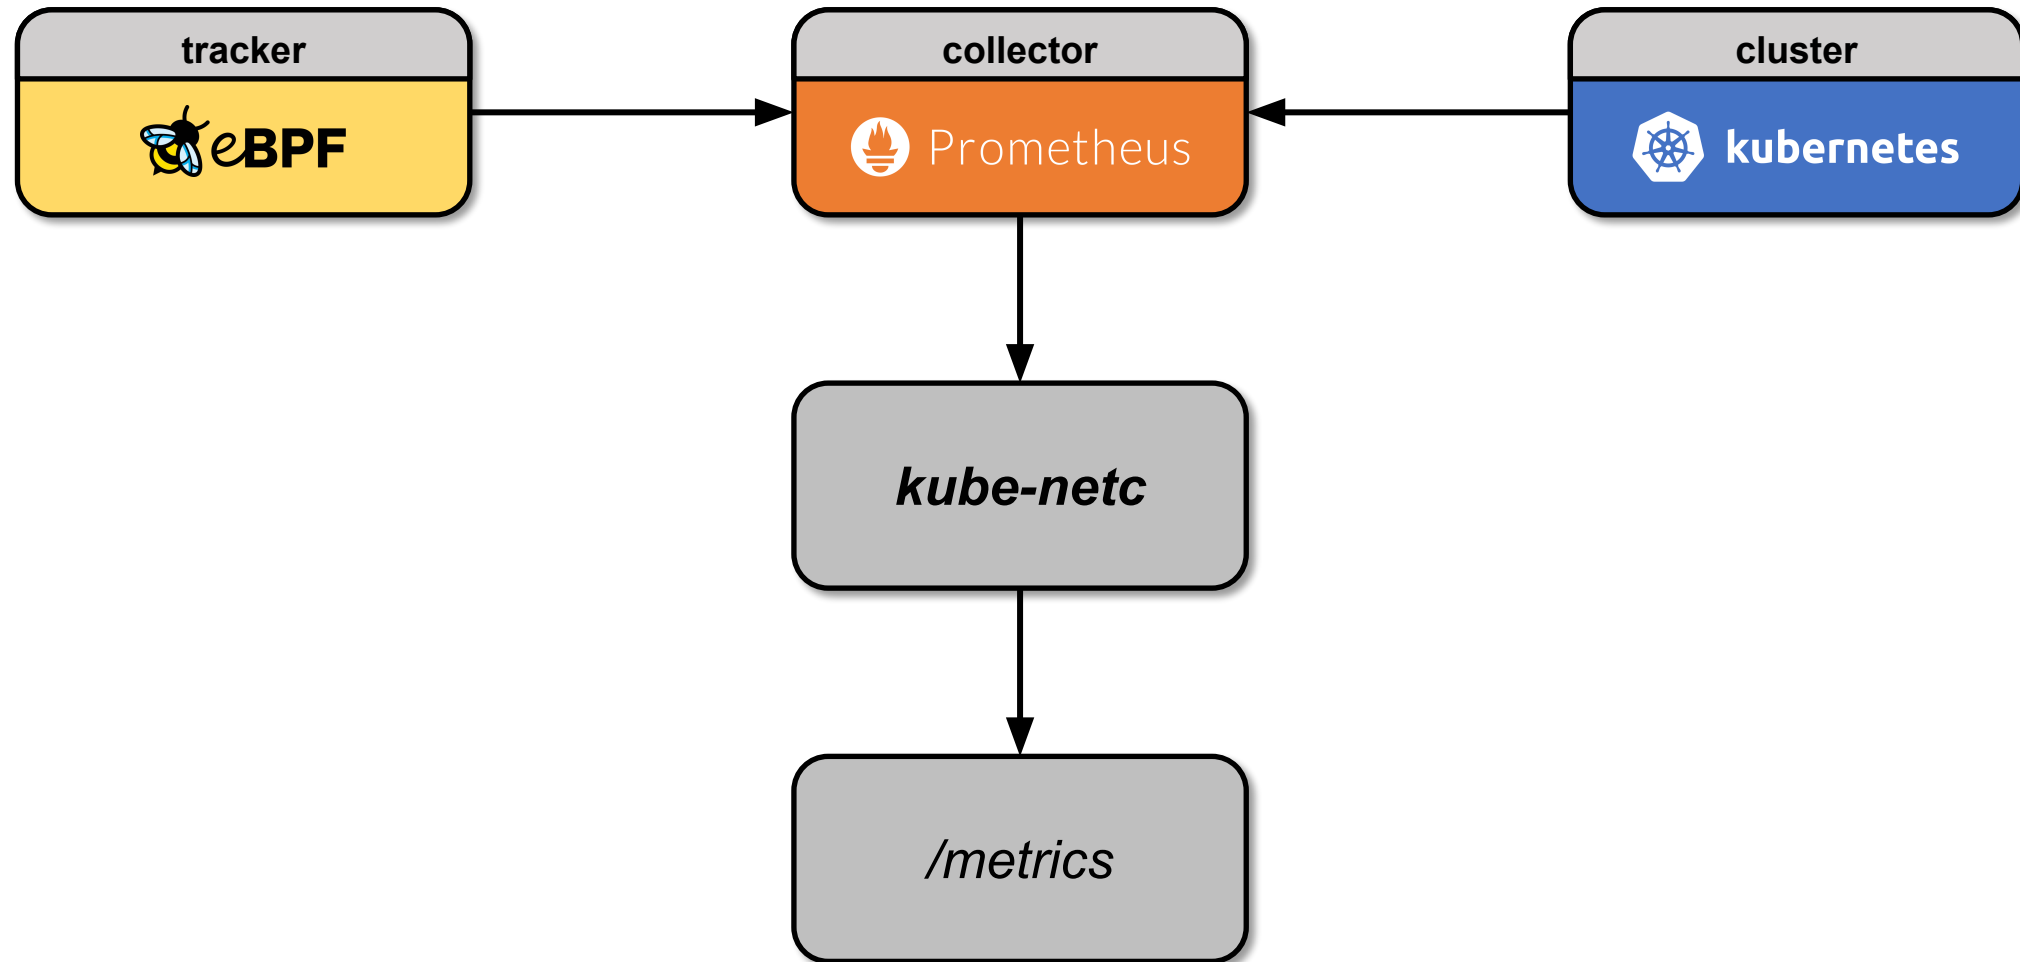
#13 opened on Jun 21 by patelrit

☐ kube-netc restarts several times bug

#7 opened on Jun 14 by JimBugwadia ➔ Release 0.0.2

6

Map of the Components



Now to demo eBPF and some of kube-netc!

Key Takeaways



- OSS and CNCF projects are for everyone
- Learn by doing
- When learning, leverage open source examples!

Thanks!



Virtual

North America 2020

Me

- drew.ripberger@gmail.com
- github.com/drewrip



- nirmata.com
- github.com/nirmata
- Twitter: [@NirmataCloud](https://twitter.com/NirmataCloud)

