

Kubernetes Jobs and the Sidecar Problem

Lightning Talk

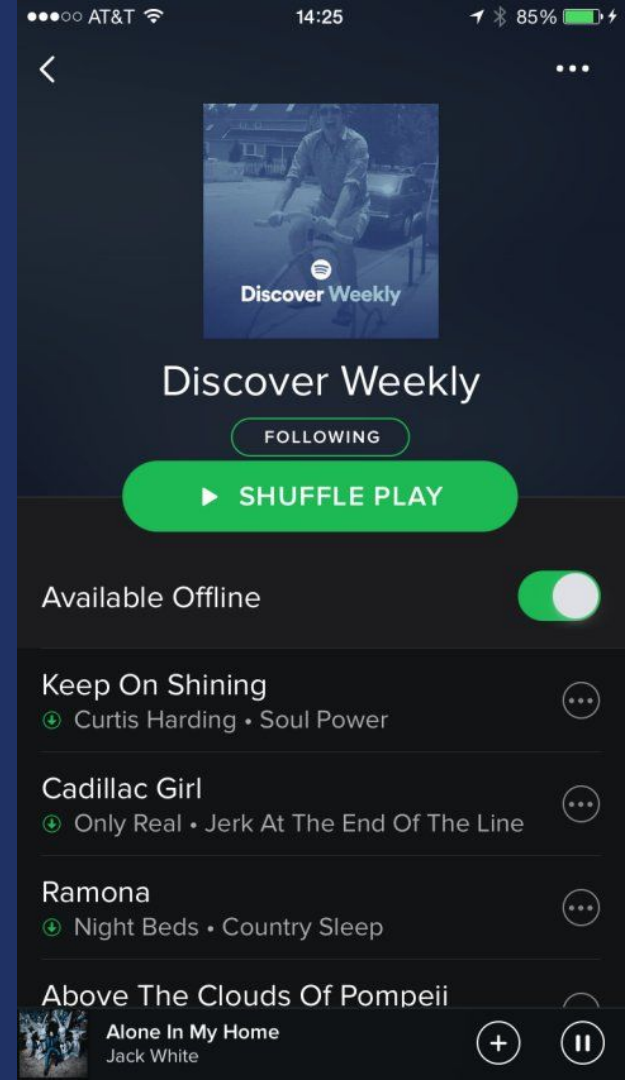
James Wen - jameswen@spotify.com
[@rochesterinnyc](https://twitter.com/rochesterinnyc)

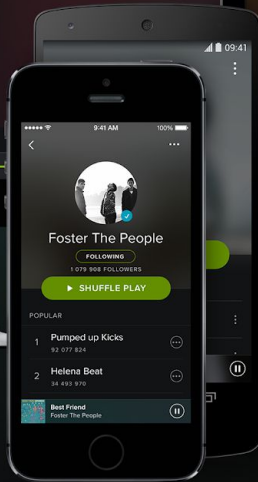
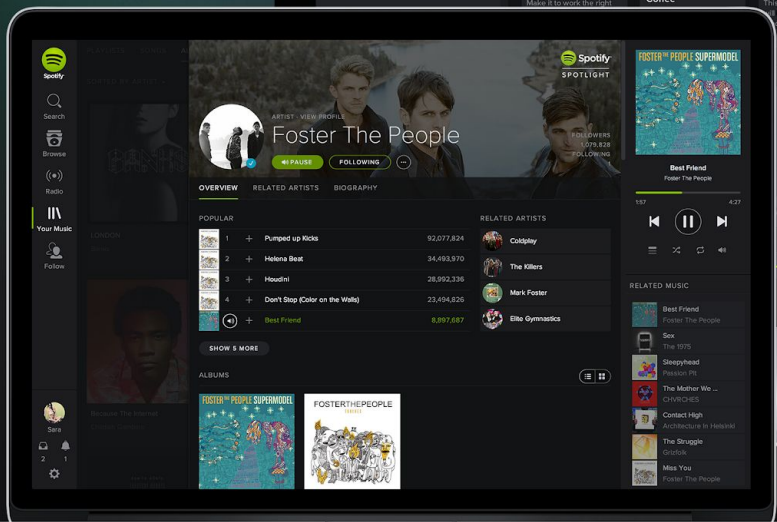
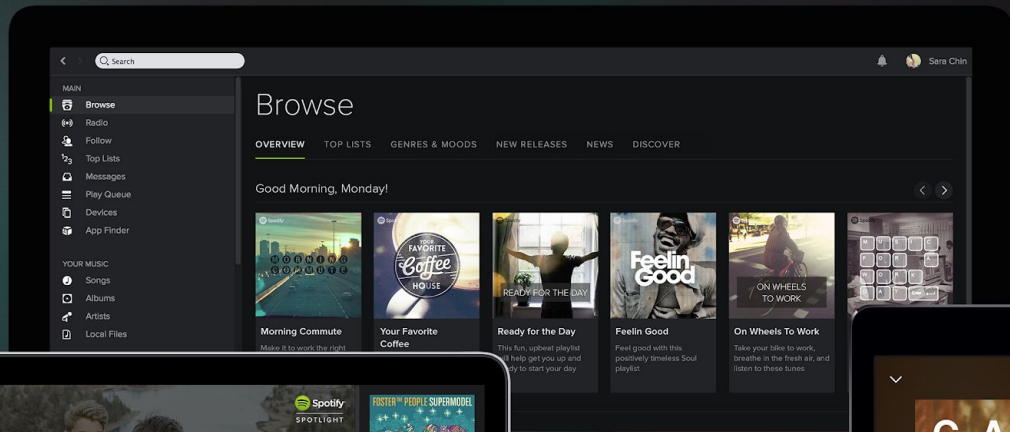
KubeCon + CloudNativeCon Europe 2019
May 20th, 2019



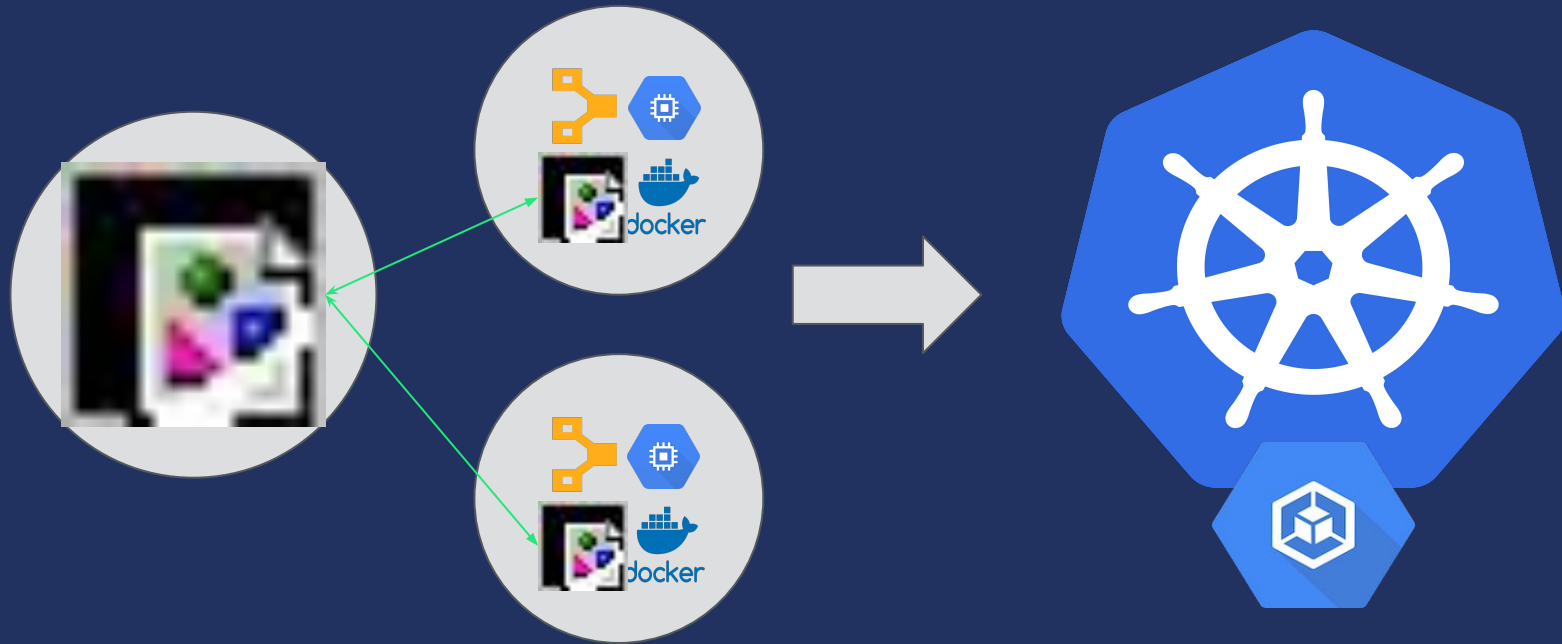
About Me

- Senior Site Reliability Engineer
- Runtime + Kubernetes Migration
- Lots of automation → CronJobs

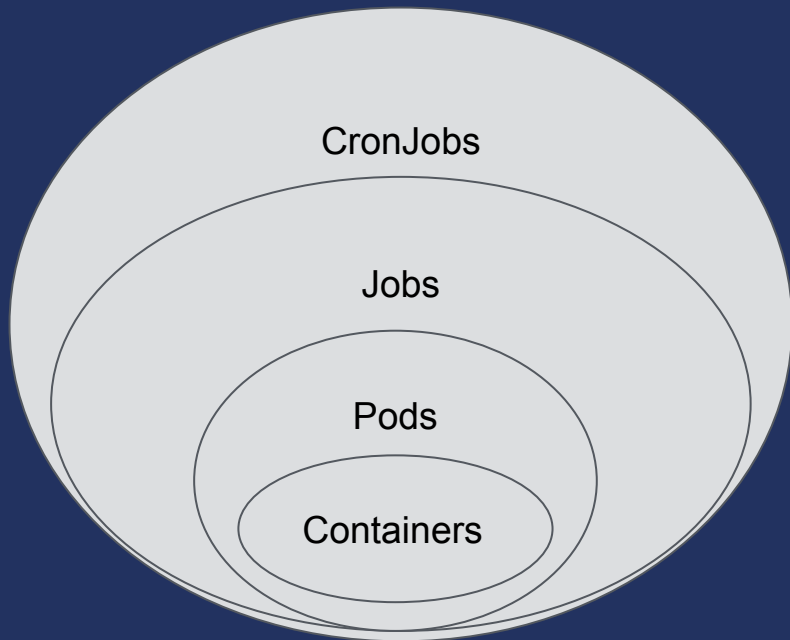




Kubernetes at Spotify



Jobs



```
apiVersion: batch/v1
kind: Job
metadata:
  name: cleanup-cluster
spec:
  template:
    spec:
      containers:
      - name: cleanup-cluster
        image: alpine
        restartPolicy: Never
      backoffLimit: 4
```

History of Jobs in Kubernetes

- September 2015: Jobs go Beta in 1.1
- March 2016: Jobs go GA in 1.2
- May 2016: ScheduledJobs created in v1.3.0-alpha.4
- May 2016: Issue for "Better support for sidecar containers in batch jobs" created

Jobs + Sidecars



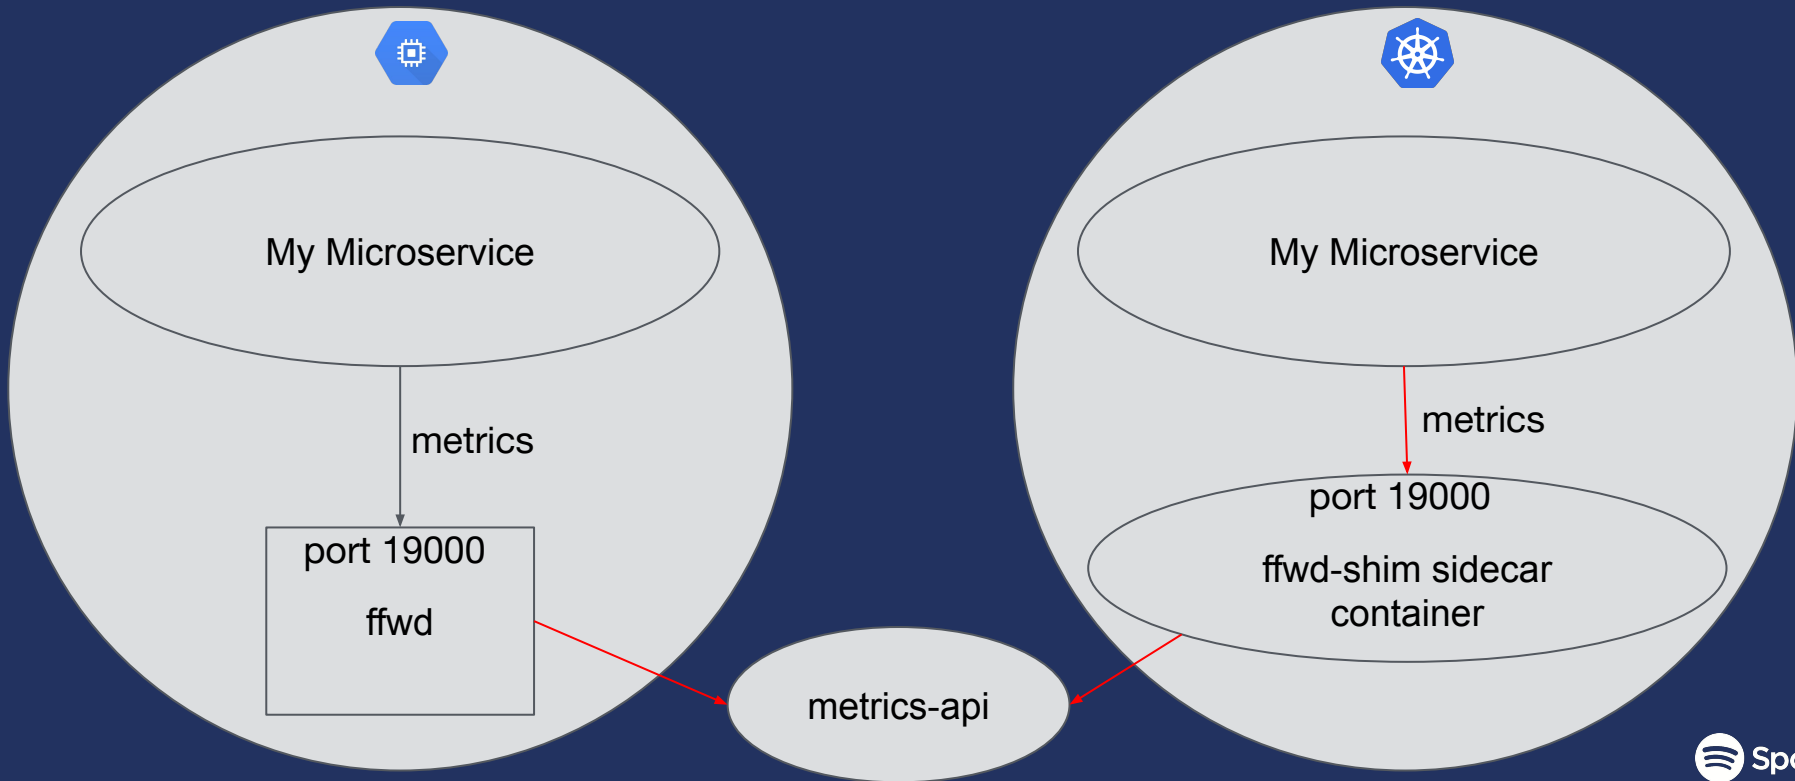
Jobs + Sidecar Problem



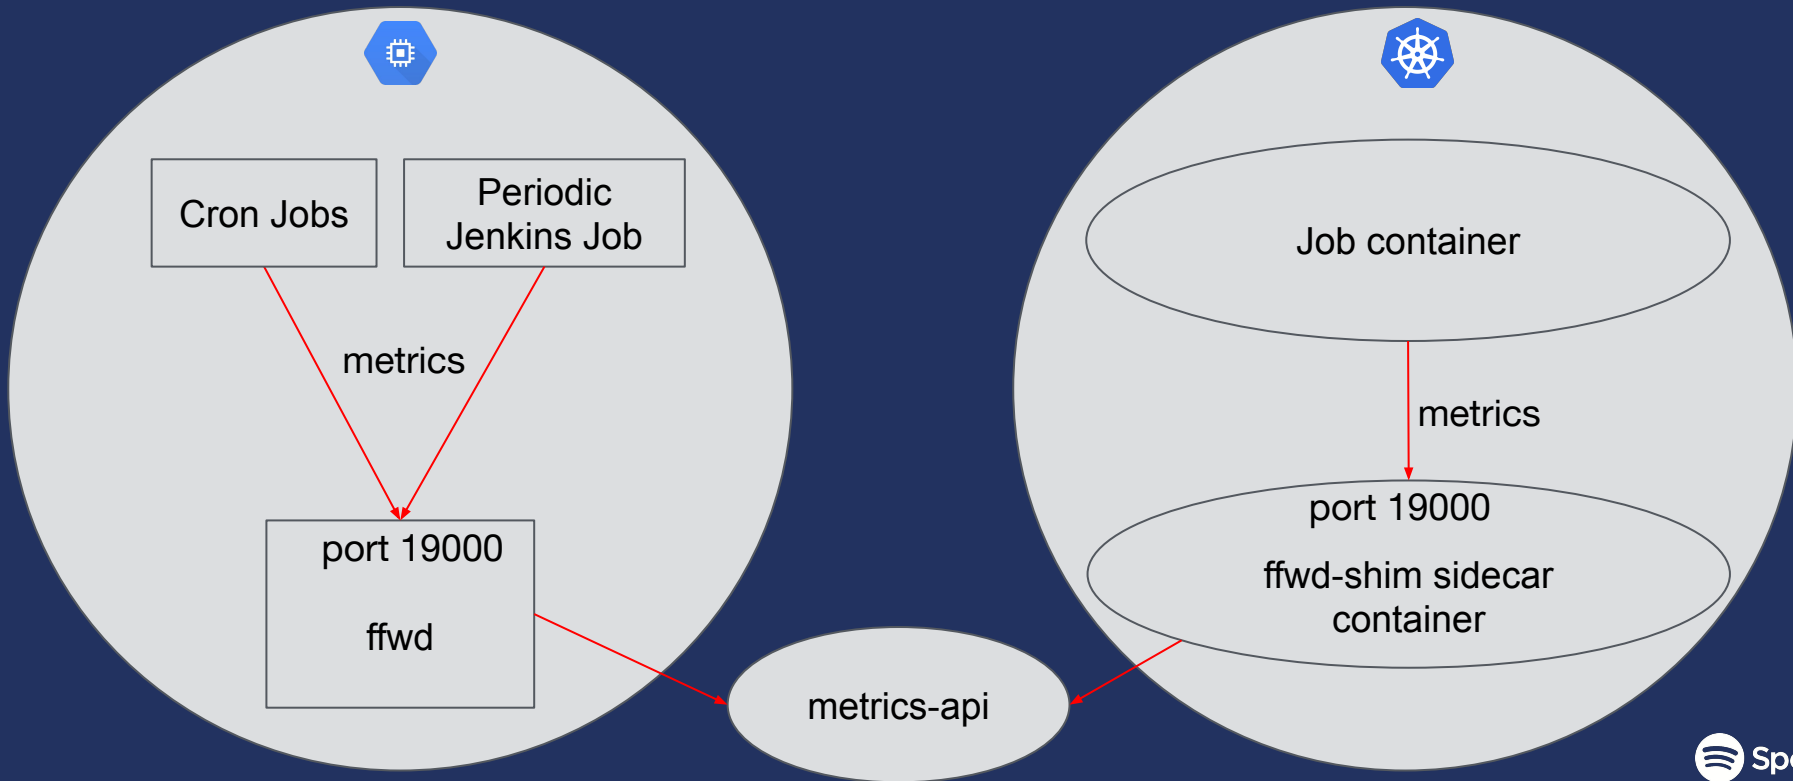
Jobs + Sidecar Problem



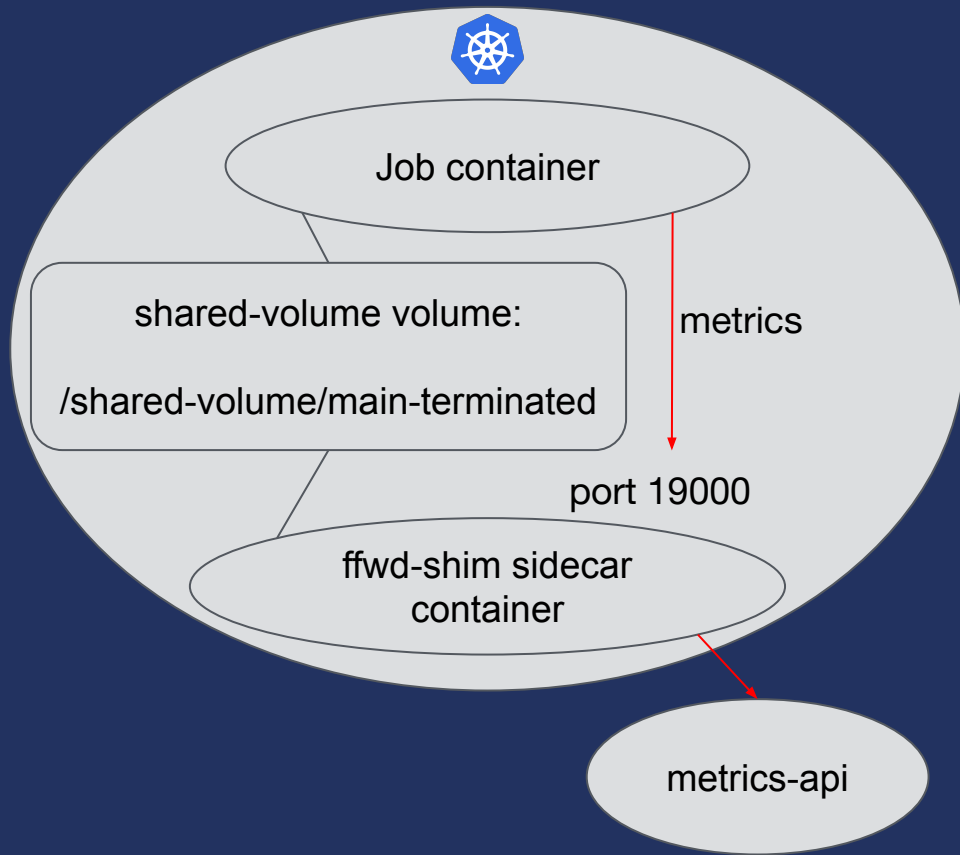
Sidecars at Spotify



Sidecar Problem at Spotify



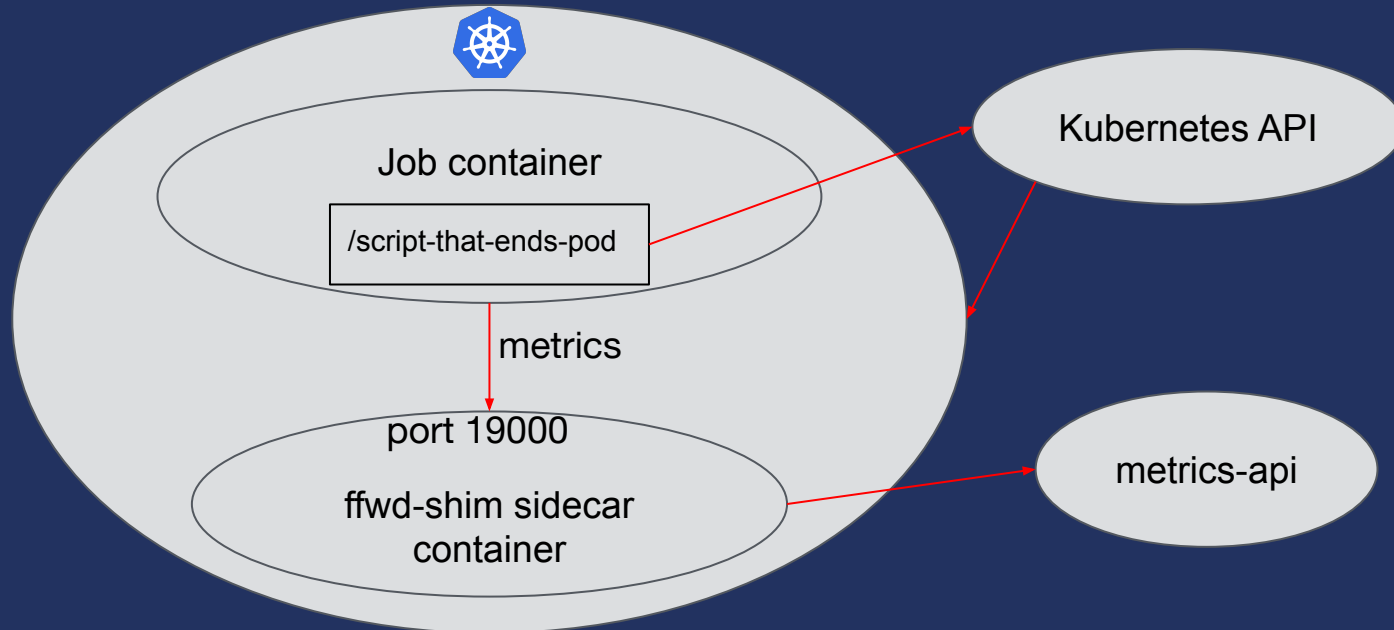
Solution: Shared Volume Communication



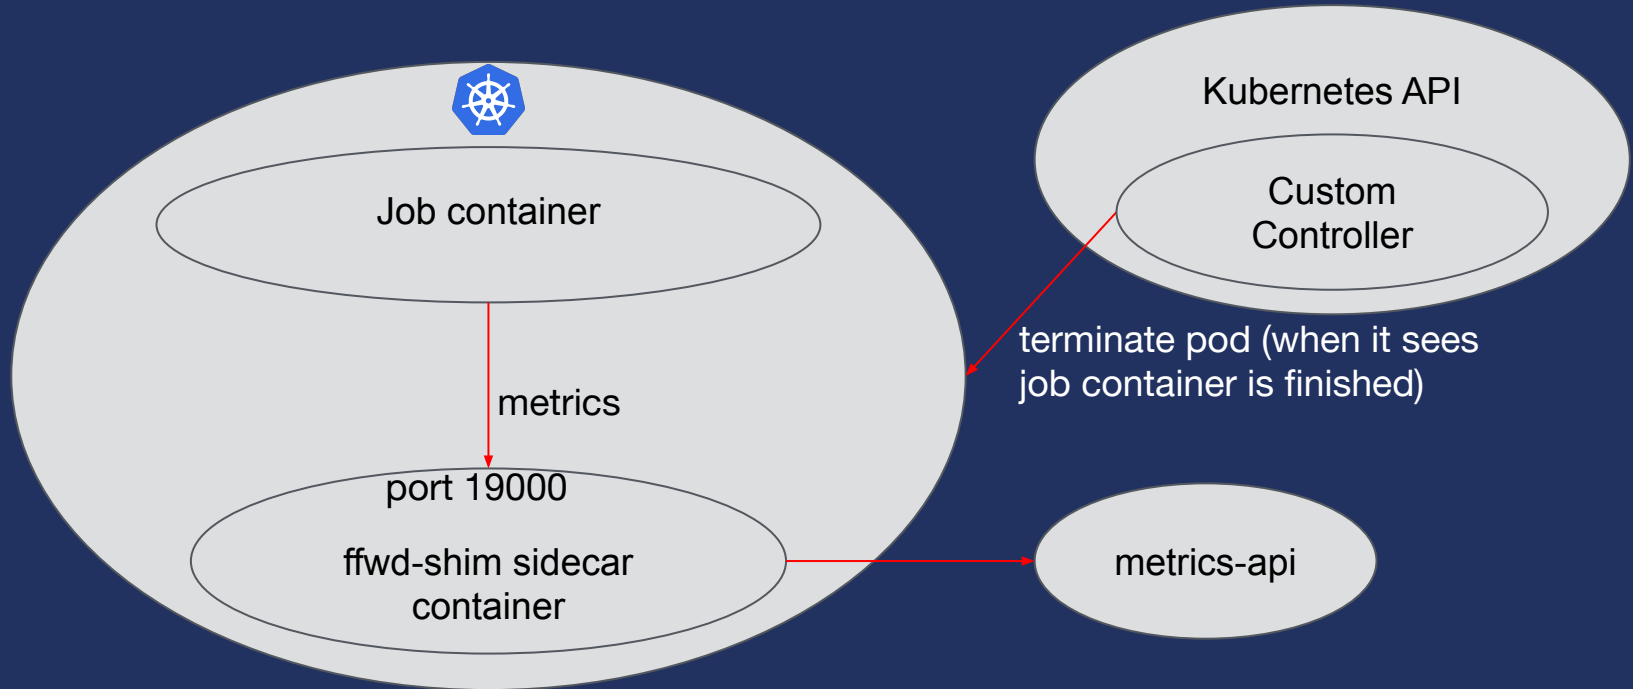
```
volumes:  
  - name: shared-volume  
    emptyDir: {}  
containers:  
  - name: main-job-container  
    image: gcr.io/my-project/my-job  
    command: ["/bin/bash", "-c"]  
    args:  
      - |  
        trap "sleep 5; touch /tmp/pod/main-terminated" EXIT  
          /usr/src/app/run-job-script.sh  
    volumeMounts:  
      - name: shared-volume  
        mountPath: /tmp/pod  
  - name: ffwd-java-shim  
    image: ffwd-java-shim  
    command: ["/bin/bash", "-c"]  
    args:  
      - |  
        /run-ffwd-java-shim.sh &  
        CHILD_PID=$!  
        (while true; do if [[ -f "/tmp/pod/main-terminated" ]];\  
          then kill $CHILD_PID; fi; sleep 1; done) &  
        wait $CHILD_PID  
        if [[ -f "/tmp/pod/main-terminated" ]]; then exit 0; fi  
    volumeMounts:  
      - name: shared-volume  
        mountPath: /tmp/pod  
        readOnly: true
```

Solution: PreStop Termination

```
containers:  
- name: main-job-container  
  image: gcr.io/my-project/my-job  
  lifecycle:  
    preStop:  
      exec:  
        command: ["bash", "-c", "/script-that-ends-pod"]
```



Solution: Custom Controller



Call to Action

- Finally a KEP out: [Joseph Irving's Blog Post on the process](https://github.com/kubernetes/enhancements/blob/master/keps/sig-apps/sidecarcontainers.md)
<https://github.com/kubernetes/enhancements/blob/master/keps/sig-apps/sidecarcontainers.md>
- Jobs have been out for ~4.5 years
- The issue for sidecar containers has been out for 3 years
- Let's get some more movement on the KEP!



Thanks for listening!

Join the Band!

spotifyjobs.com

James Wen - jameswen@spotify.com

[@rochesterinnyc](https://twitter.com/rochesterinnyc)

KubeCon + CloudNativeCon Europe 2019

May 20th, 2019