



Schedule Scaling of Kubernetes Resources kube-start-stop

Lili Cotic - Software Engineer @ [Weaveworks](https://weaveworks.com)

Twitter: @LiliCotic

Github: lili

kube-start-stop



github.com/lilic/kube-start-stop

Why?



- Automatically start and stop Kubernetes resources in the namespace
- Schedule your dev workload to be down over the weekend
- Autoscaled Kubernetes cluster = save resources
- Not autoscaled Kubernetes cluster = dev cluster doesn't cause havoc
- Similar idea as the AWS Instance Scheduler but for Kubernetes

Design



- Uses *CustomResourceDefinitions* under the hood
- Registers the `Schedule` kind under the `scheduler..io/v1alpha1` apiVersion
- Timezone is UTC
- Scales down deployments based on the selector

How?



```
# Deploy controller
```

```
`kubectl apply -f manifest/deploy/`
```

```
# Apply manifest file with schedule
```

```
`kubectl apply -f schedule.yml`
```

Configure via a Kubernetes manifest file



```
apiVersion: scheduler.io/v1alpha1
kind: Schedule
metadata:
  name: schedule
  Namespace: dev
spec:
  serviceAccountName: kube-start-stop
  schedules:
  - replicas: 0
    selector: my-deployment
    start:
      day: friday
      time:
        hour: 20
        minute: 10
    stop:
      day: monday
      time:
        hour: 8
        minute: 10
```

Future plans



- Add support for local timezones - timezone conversion
- Helm chart support
- And scaling of all resources in the given namespace

Thank you!



github.com/lilic/kube-start-stop

Twitter: @LiliCotic

Github: lilic