

k8s autoscale

- (Horizontal Pod Autoscaler, HPA)
- HPA CPU 3 CPU 5 CPU
- HPA (Cluster Autoscaler, CA)

yml

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: web-php
spec:
  replicas: 1
  selector:
    matchLabels:
      run: web-php
  template:
    metadata:
      labels:
        run: web-php
    spec:
      containers:
        - image: maho/web-php:0.2
          name: web-php
          resources:
            requests:
              cpu: 200m
---
apiVersion: v1
kind: Service
metadata:
  name: web-php
spec:
  type: NodePort
  selector:
    run: web-php
  ports:
    - port: 80
      protocol: TCP

```

nodePort: 31446

Command

```
$ kubectl autoscale deployment web-php --cpu-percent=50 --min=1 --max=10
horizontalpodautoscaler.autoscaling/web-php autoscaled
```

kubectl describe node []	가 HPA
kubectl autoscale [] []	HPA . kubectl autoscale help 가
kubectl get hpa	HPA . get describe

Docs

- [kubectl_delete](#)

- [Kubernetes](#)

From:
<https://jace.link/> - **Various Ways**

Permanent link:
<https://jace.link/open/k8s-autoscale>

Last update: **2021/10/07 02:22**

