

Local K8S deployment with Microk8s



kubernetes

Table of contents

1	Microk8s installation	3
1.1	Snap installation	3
1.2	Microk8s installation	3
1.3	MetalLB installation	3
1.4	Deploy Metrics Server	4
2	Configuration of the cluster	4
2.1	MetalLB configuration.....	4
2.2	Metrics server configuration	5
3	Microk8s usage.....	7

1 Microk8s installation

The easiest and fastest way to create a local cluster is using microk8s. It's possible to make containers, push them, and deploy them directly in the laptop.

1.1 Snap installation

Snap is needed to install microk8s.

```
$ sudo apt update
```

```
$ sudo apt install snapd
```

1.2 Microk8s installation

The easiest and fastest way to create a local cluster is using microk8s. It's possible to make containers, push them, and deploy them directly in the laptop. One line installation:

```
$ sudo snap install microk8s --classic
```

After a few seconds, microk8s is installed. To check if kubernetes is running:

```
$ microk8s.kubectl get all --all-namespaces
```

We will see the following:

NAMESPACE	NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE
default	service/kubernetes	ClusterIP	10.152.183.1	<none>	443/TCP	15m

1.3 MetalLB installation

We need a load balancer in our microk8s cluster in order to assign floating IPs to the pods.

```
$ kubectl apply -f
```

```
https://raw.githubusercontent.com/google/metallb/v0.7.3/manifests/metallb.yaml
```

This will deploy the required pods:

```
$ microk8s.kubectl get all --all-namespaces
```

We can see the following pods running:

NAMESPACE	NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)	AGE			
default	service/kubernetes	ClusterIP	10.152.183.1	<none>	443/TCP	8m16s			
NAMESPACE	NAME	TYPE	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE	NODE SELECTOR	AGE
metallb-system	daemonset.apps/speaker		1	1	0	1	0	<none>	5s
NAMESPACE	NAME	TYPE	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE	AGE	
metallb-system	deployment.apps/controller		0/1	1	0	1	0	5s	
NAMESPACE	NAME	TYPE	DESIRED	CURRENT	READY	AGE			
metallb-system	replicaset.apps/controller-7cc9c87cfb		1	1	0	5s			


```
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: my-ip-space
      protocol: layer2
      addresses:
      - 10.0.2.1-10.0.2.21
```

After creating the file, we need to apply the configuration to the load balancer:

```
$ microk8s.kubectl apply -f metallb-conf.yaml
```

2.2 Metrics server configuration

The configuration of the metrics server is very easy, we just only need to add some lines (the ones inside the red box) in the metrics server deployment file. This file is in *deploy/1.8+/metrics-server-deployment.yaml*.

```
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: metrics-server
  namespace: kube-system
---
apiVersion: extensions/v1beta1
kind: Deployment
metadata:
  name: metrics-server
  namespace: kube-system
  labels:
    k8s-app: metrics-server
spec:
  selector:
    matchLabels:
      k8s-app: metrics-server
  template:
    metadata:
      name: metrics-server
      labels:
        k8s-app: metrics-server
    spec:
      serviceAccountName: metrics-server
      volumes:
        # mount in tmp so we can safely use from-scratch images and/or
        read-only containers
        - name: tmp-dir
          emptyDir: {}
      containers:
        - name: metrics-server
          image: k8s.gcr.io/metrics-server-amd64:v0.3.1
          imagePullPolicy: Always
          volumeMounts:
            - name: tmp-dir
              mountPath: /tmp
          command:
            - /metrics-server
            - --kubelet-insecure-tls
            - --kubelet-preferred-address-types=InternalIP
```

After saving the file, we need to apply the configuration:

```
$ microk8s.kubectl apply -f deploy/1.8+/metrics-server-deployment.yaml
```

At this point we should have a complete functional kubernetes cluster running inside our laptop.

3 Microk8s usage

The commands in Microk8s are the same as in kubectl but we need to add the *microk8s.* prefix to all of them. Here are some examples:

See the pods deployed on the cluster:

```
$ microk8s.kubectl get po
```

Check the services of the cluster:

```
$ microk8s.kubectl get svc
```

It is possible to deploy in microk8s an existing network service from a different cluster. We can copy and paste the deployments and services *yaml* files in a folder in our laptop. For example, we have a folder called *media-pilot-deployments/*, it is very to deploy the whole network service with this command:

```
$ microk8s.kubectl apply -f media-pilot-deployments/
```

This command will create the deployments and services, after a few seconds we will have all the pods running in Microk8s.