NITSolutions
Werkt voor je

Kubernetes for Network Developers

Kubernetes is cool. And with BGP makes it way cooler.

By Egbert Verhage
AS216195

kubectl explain events

- Every developer or DevOps Engineer wants to use Kubernetes
- For most engineers Kubernetes is still a blackbox
- Kubernetes gets deployed without much thought, so the developer can just get on with their work.

- But firstm a quick recap what is Kubernetes?

`kubectl explain` give api details of api resource

kubectl get service Kubernetes

> Kubernetes is a platform for managing containerized workloads.

Main Kubernetes components:

- **API Server**: The api service
- **Controller Manager**: The state manager
- **Scheduler**: The pod scheduler
- **Kubelet**: The node manager
- **etcd**: The database

`kubectl get` list a resource, this case the Kubernetes service in default namespace

kubectl get pod -n api-gateway

The pod! Your OCI container(s) that runs on a node.

Pod Properties:

- Has its own IP address (IPv4/IPv6)
- Can have one or more containers, with the same network
- Pods are disposable; expect them to come and go
- But, want to change something? Recreate it!

Want to update or scale your pod? Use a `Deployment` / `StatefulSet` *for pod deployment.*

`kubectl get pod -n api-gateway` list all the `pods` in `api-gateway` namespace

kubectl create service webserver --tcp=8080:8080

Everytime the pod is created, it gets a new IP address; A service always points towards pods to solve this.

A service is an internal Kubernetes resource for pods to address other services on the network.

It gets a static IP address and DNS entry via CoreDNS.

But there is more, you can create different types of services:

- ClusterIP
- LoadBalancer
- NodePort
- ExternalName

`kubectl create` can create new resources. This create a service for webservice on port 8080:8080

kubectl create namespace nlnog-day-2026

Small thing about security: **namespaces**

Namespaces in Kubernetes are a way to isolate resources.

- They groups Pods, Services, Deployments, NetworkPolicies
- Can create resources with same name in different **namespaces**
- It is by default **not** a network boundary.
- RBAC 🔒. *(But not for today)*

`kubectl create namespace nlnog-day-2026`: Create a new namespace `nlnog-day-2026` for resources of this presentation

kubeadm init

Yes! We are creating a new Kubernetes cluster...

Wait, before you create a new cluster we need to configure pod-cidr and service-cidr:

Pod CIDR; IP addresses assigned to Pods (IPv4 Default: 10.244.0.0/16)

- Usually managed by your CNI
- We can announce them over BGP

Service CIDR; Virtual IP addresses for Kubernetes Services (IPv4 Default: 10.96.0.0/12)

- Virtual addresses; These can be shared between different clusters
- Only used on Kubernetes nodes / pods

`kubeadm` is tool create and manage a Kubernetes cluster

helm install flannel

Now the cluster is running, but some pods won't start.

One important piece is still missing: pod networking. That's where the CNI comes in.

> **CNI** stands for Container Network Interface

The CNI connects the nodes and pods to same network. This can be via VXLAN, Direct Routing or WireGuard.

First, lets setup flannel as CNI and inspect it. *(Demo time)*

`helm` is deployment tool for Kubernetes

asciinema play deploy-flannel.cast

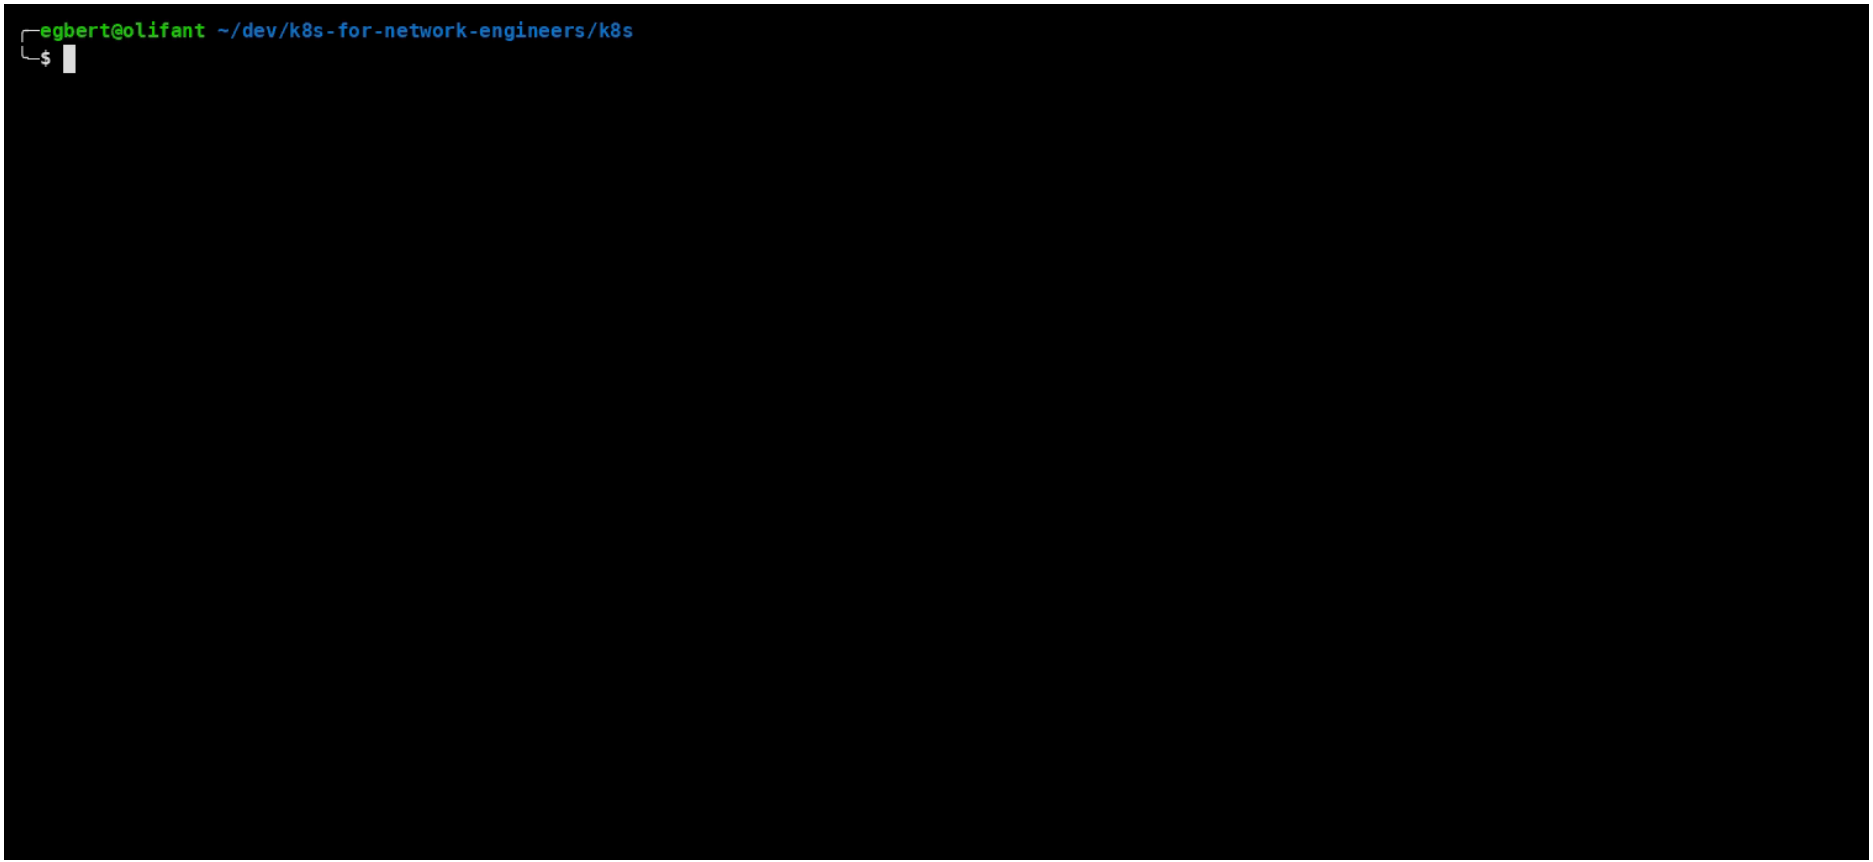

kubectl apply -f ip-pool.yaml

Yes! Your Kubernetes cluster is running, now for the BGP stuff.

Now we want to do BGP stuff. By default Kubernetes does not support BGP, but maybe your CNI does.

For now lets show an example (demo) with MetalLB.

> **MetalLB** is load-balancer implementation for a *bare metal* Kubernetes clusters.

`kubectl apply -f [file]` is easy way to apply yaml resource directly. You can use `-` is for stdin

asciinema play metallb.cast

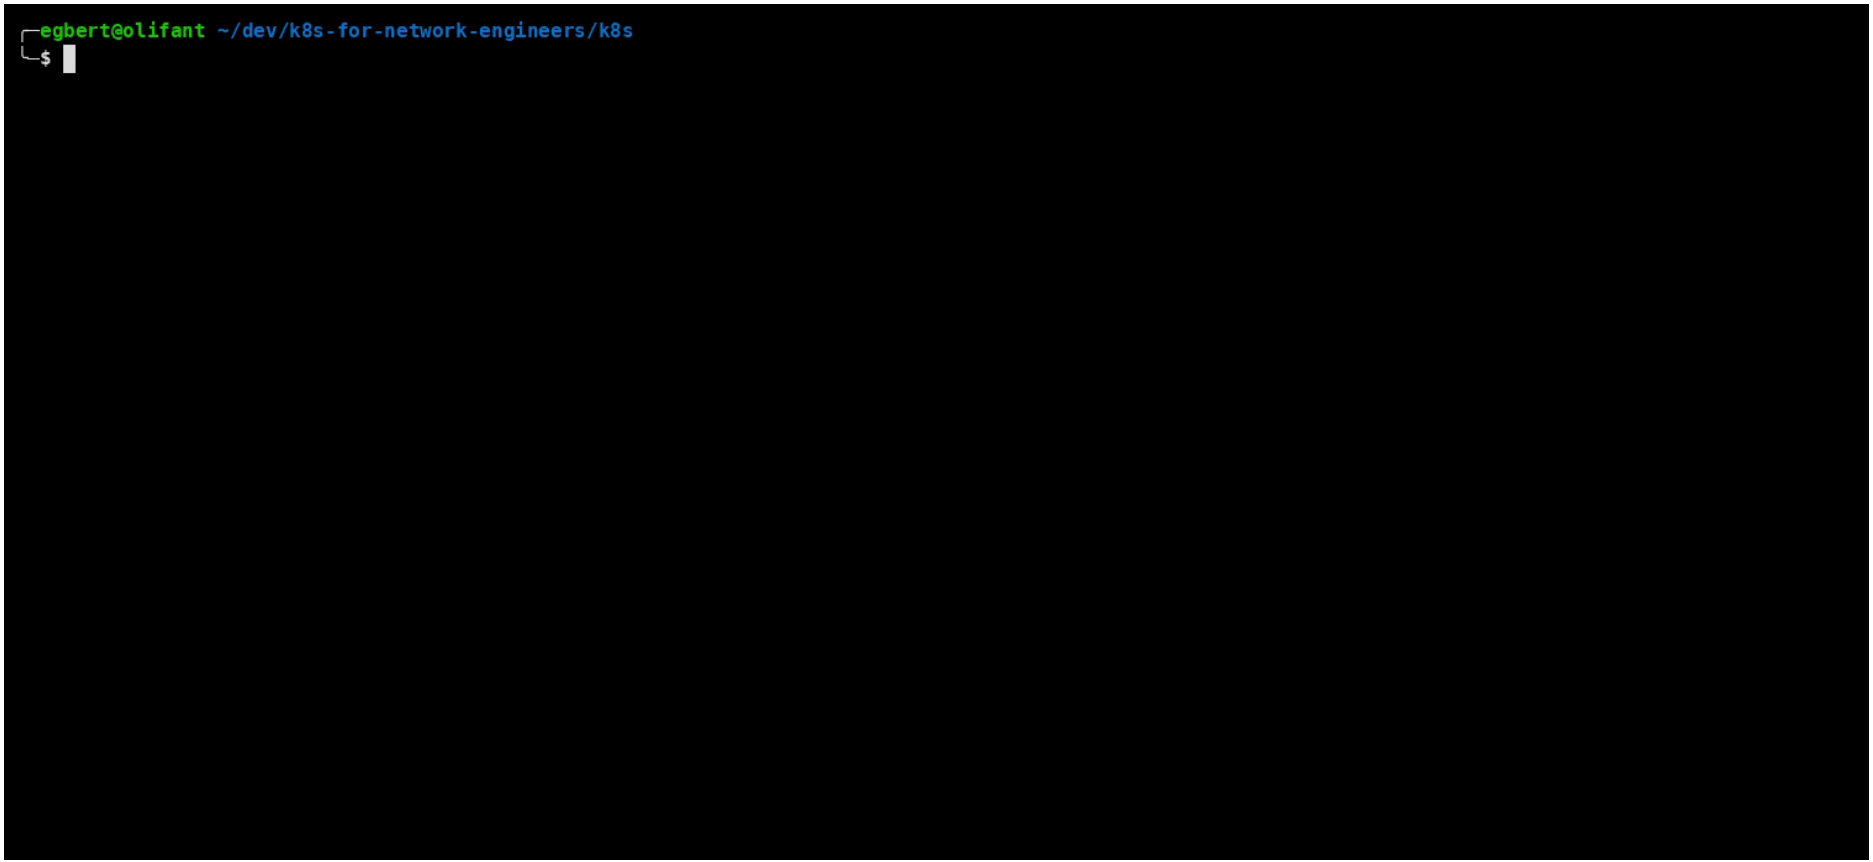

kubectl explain ciliumbgpadvertisements.cilium.io

Now we have a simple cluster running with BGP for a LoadBalancer Service.

This is just the start, next up:

- Now I swap the CNI for Cilium and configure cilium without Source NAT (masquerade) from the node
- Announce the pod pools over BGP
- Announce LoadBalancer over BGP

`ciliumbgpadvertisements.cilium.io` is custom CRD's of cilium to configure bgp advertisements

asciinema play cilium.cast

ip addr show dev net0

But wait, there is more!

Besides connecting pods internaly in kubernetes, we add extra interfaces to a pod; Via:

`multus`

`multus` is an intermediate (meta-plugin) for your CNI and can add extra interfaces to your pod. For example `macvlan` to can connect a pod directly to different network via extra interface.

`ip`; the linux network tool

asciinema play multus.cast

kubectl get vmi

And way way more.

The way more networking thing you can do. For example:

- `KubeVirt` for KVM VM's. Integrates great with Multus
- `Rook` ; Ceph in Kubernetes.
- `SR-IOV CNI` ; Attach system interface directly in pod.
- `Whereabouts` ; IPAM for Multus

vmi stands for `VirtualMachineInstance`, so vm's in kubernetes

kubectl version

And special thanks to:

- NLnog - For having me
- Mark Bertels - Checking my work and Language support
- Tjerk Jan Vonk - Checking my work and support
- Dion Groeneveld - Supporting me

`kubectl version`: show version of the client and server

NITSolutions
Werkt voor je

Questions?

Egbert Verhage
egbert@eggie.nl

AS216195

linkedin.com/in/eggiecode

`kubectl tail`: Awesome plugin to show logs of multple pod logs