Kubernetes when you do not have cloud niceties

Yiorgos Adamopoulos
VAIX.ai

GRNOG-8 20190612

Infrastructure Overview

VAIX.ai

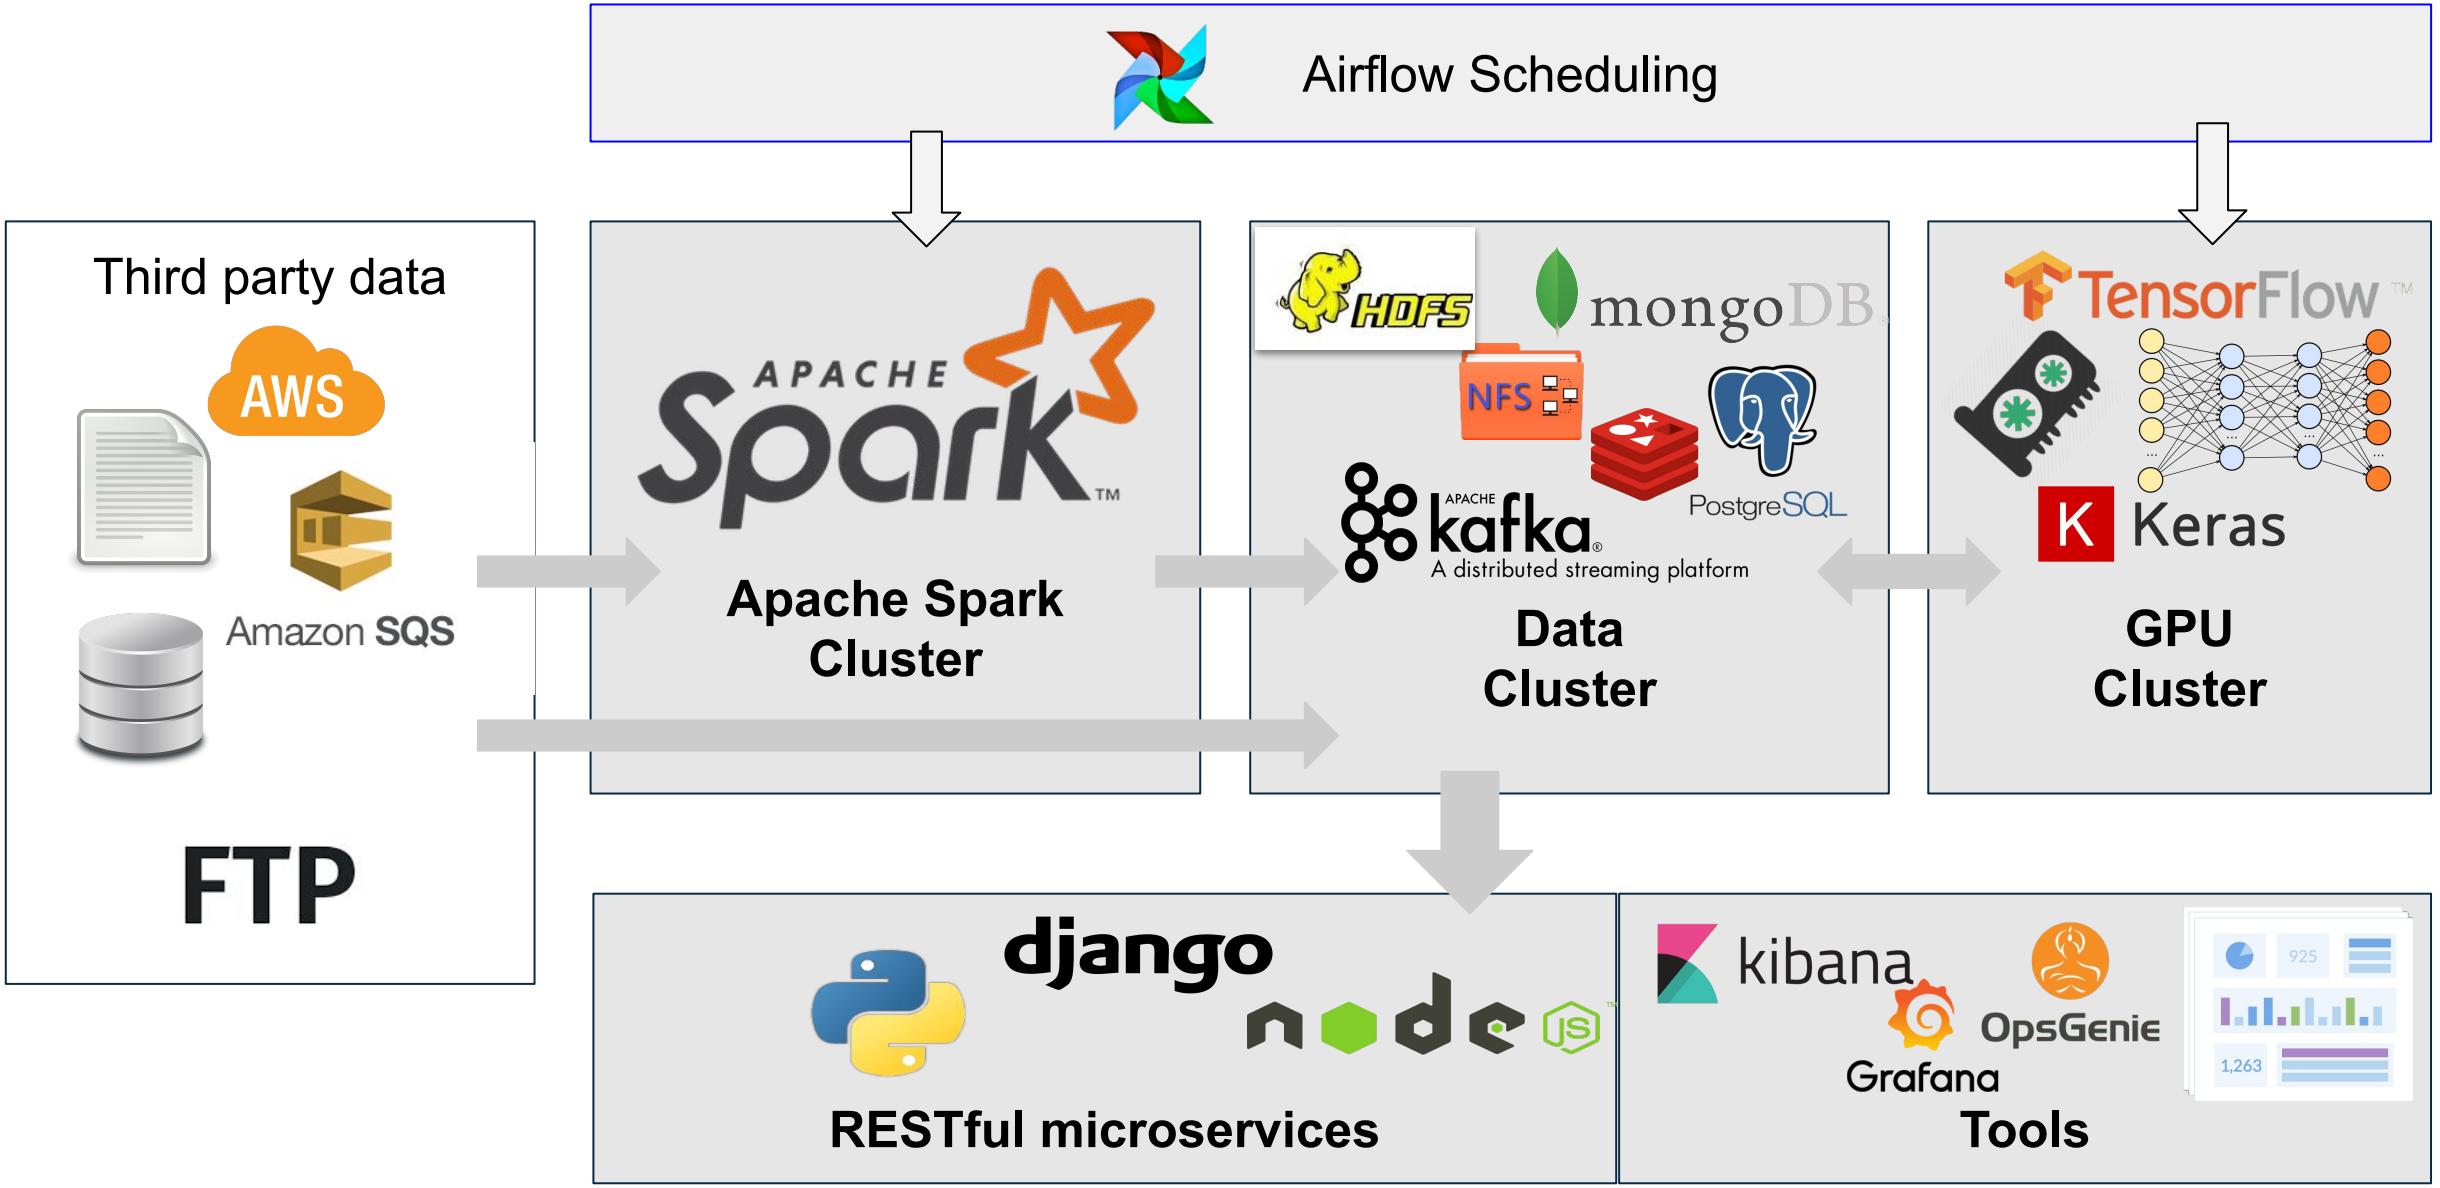

VAIX

Kubernetes is a necessary evil
to hide complexity by ...adding
complexity.

In a world of unlimited funding….

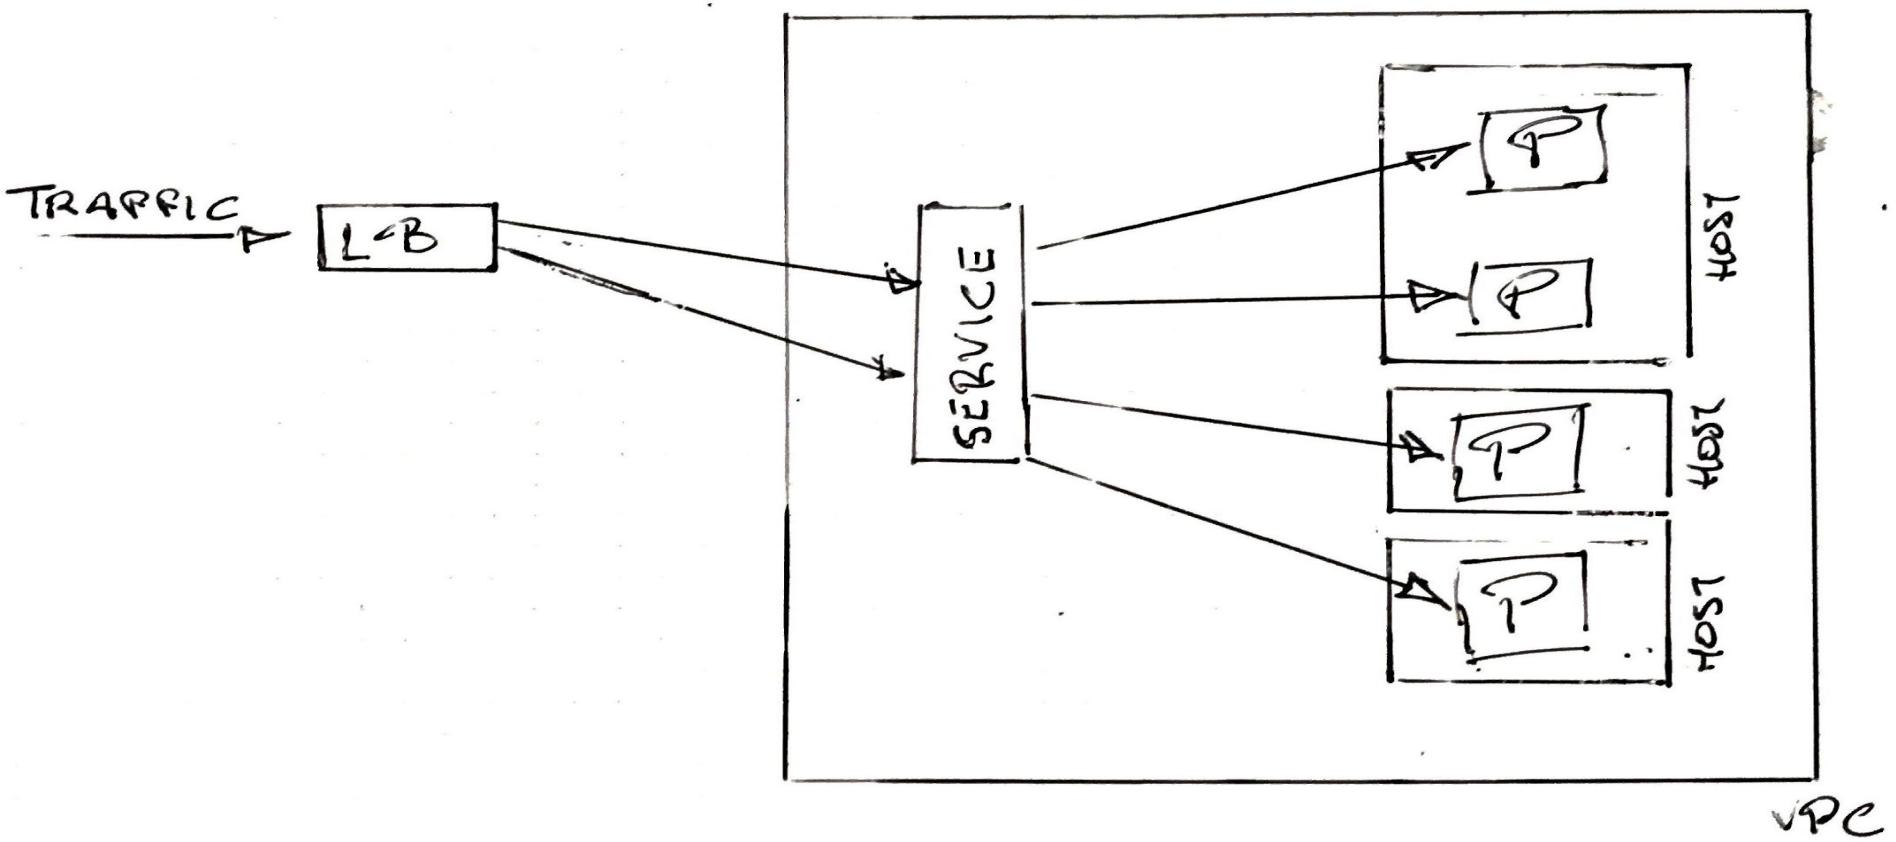

But you cannot have that!

VAIX

Operational constraints:

- No VPC
- No Load Balancer
- No firewall
- No second host interface

So what do you do?

Yes you can run a “LB” in front of NodePorts,

but:

- CVE-2018-1002105
- They are also open to the whole Internet
- Network Policies don’t help

Why not MetalLB?

MetalLB rocks, but:

- The layer-2 option did not work with our provider
- You need to setup BGP between a LB and MetalLB, so no!

So let's build a VPC! Shall we?

N2N VPN

Image source: ntop.org

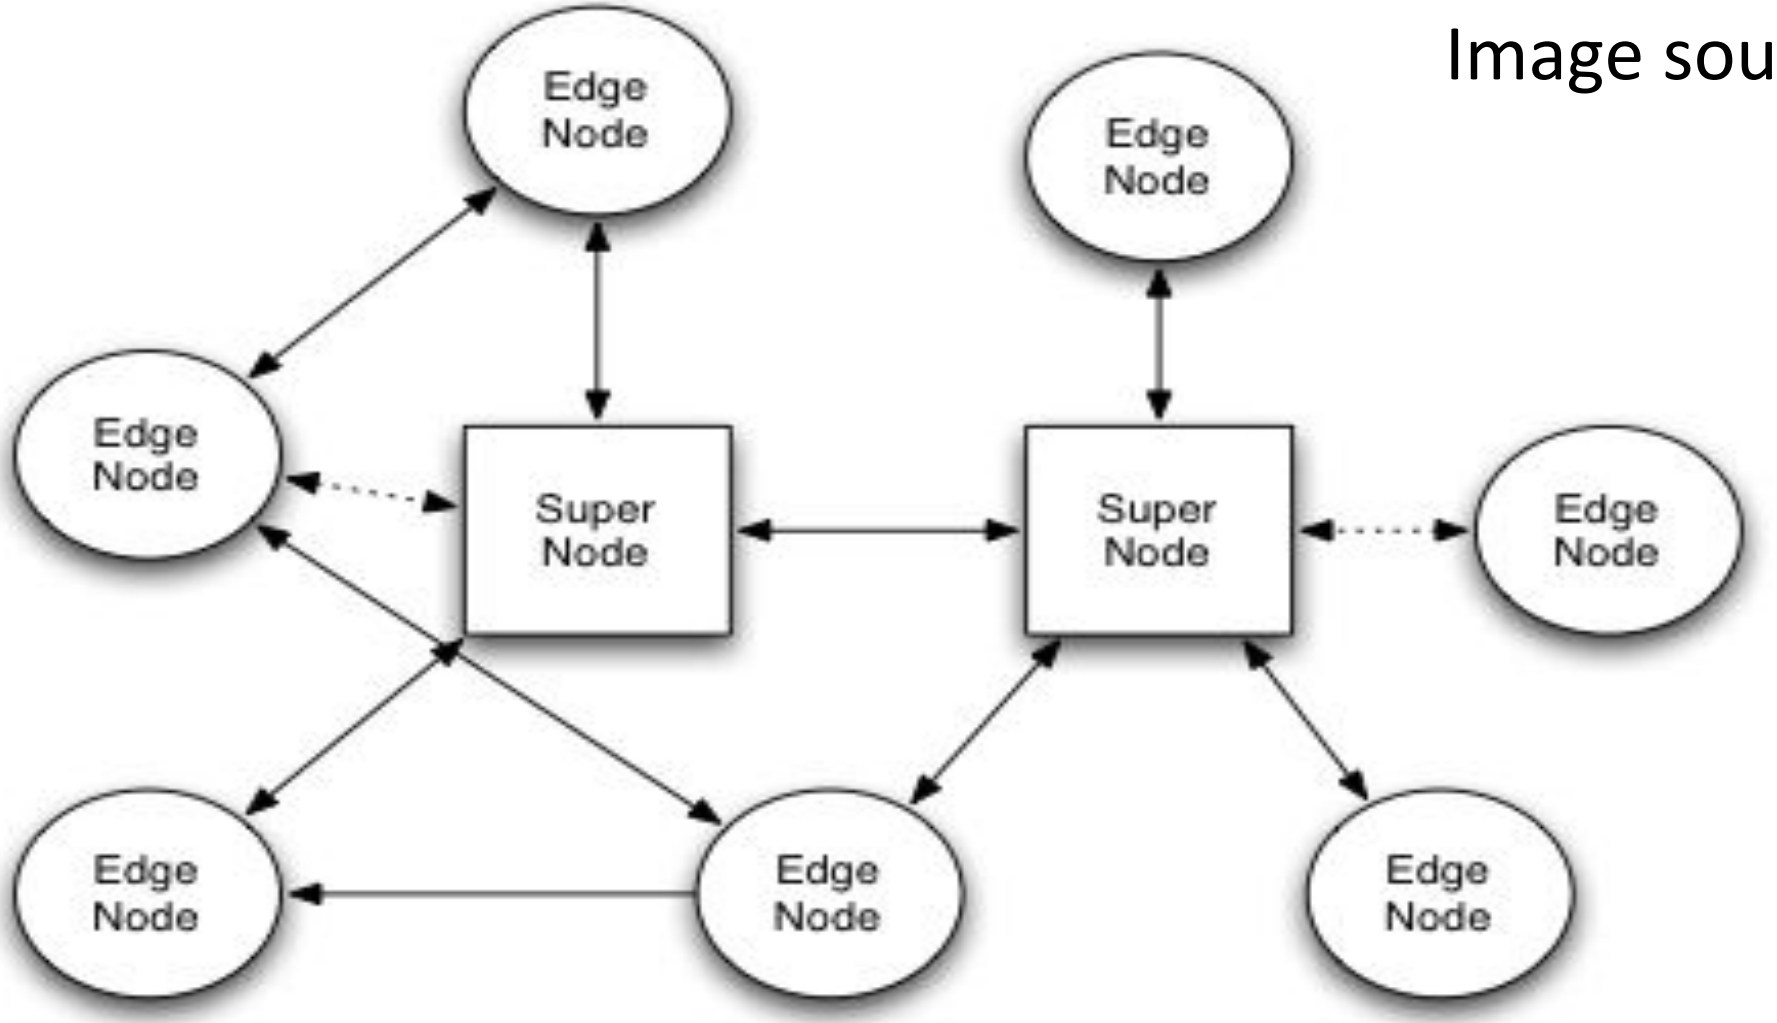

So what do we get with N2N?

- Symmetrically encrypted traffic
- Static IP per node (DHCP also possible)
- Isolated network when using RFC1918 addresses

Now to build the kubernetes cluster

- kubeadm

  The most platform agnostic tool

When it’s not DNS, it is the MTU

- *Everything is a DNS problem*
- Except when it is not
- Calico MTU: 1440
- N2N default MTU 1400 ⟶ 1500

Let’s put NodePorts to N2N only

```
$ kubectl -n kube-system get configmap
kube-proxy -o yaml
:
:
nodePortAddresses: ["172.29.0.0/16"]
:
:
```

GRNOG-8

Let’s check our DNS thingie

```
$ kubectl -n kube-system get configmap
coredns -o yaml
:
:
forward . 1.1.1.1 8.8.8.8 {
        policy sequential
}
:
:
```

What does the kubelet do?

The kubelet does most of the scheduling job on a node. Pain points:

- --address=x.y.z.w
- --node-ip=x.y.z.w

Any pain left?

- You need a lot of book-keeping with NodePorts
- Port 179/tcp still open (calico BIRD) to the world

Look how far we've gone:

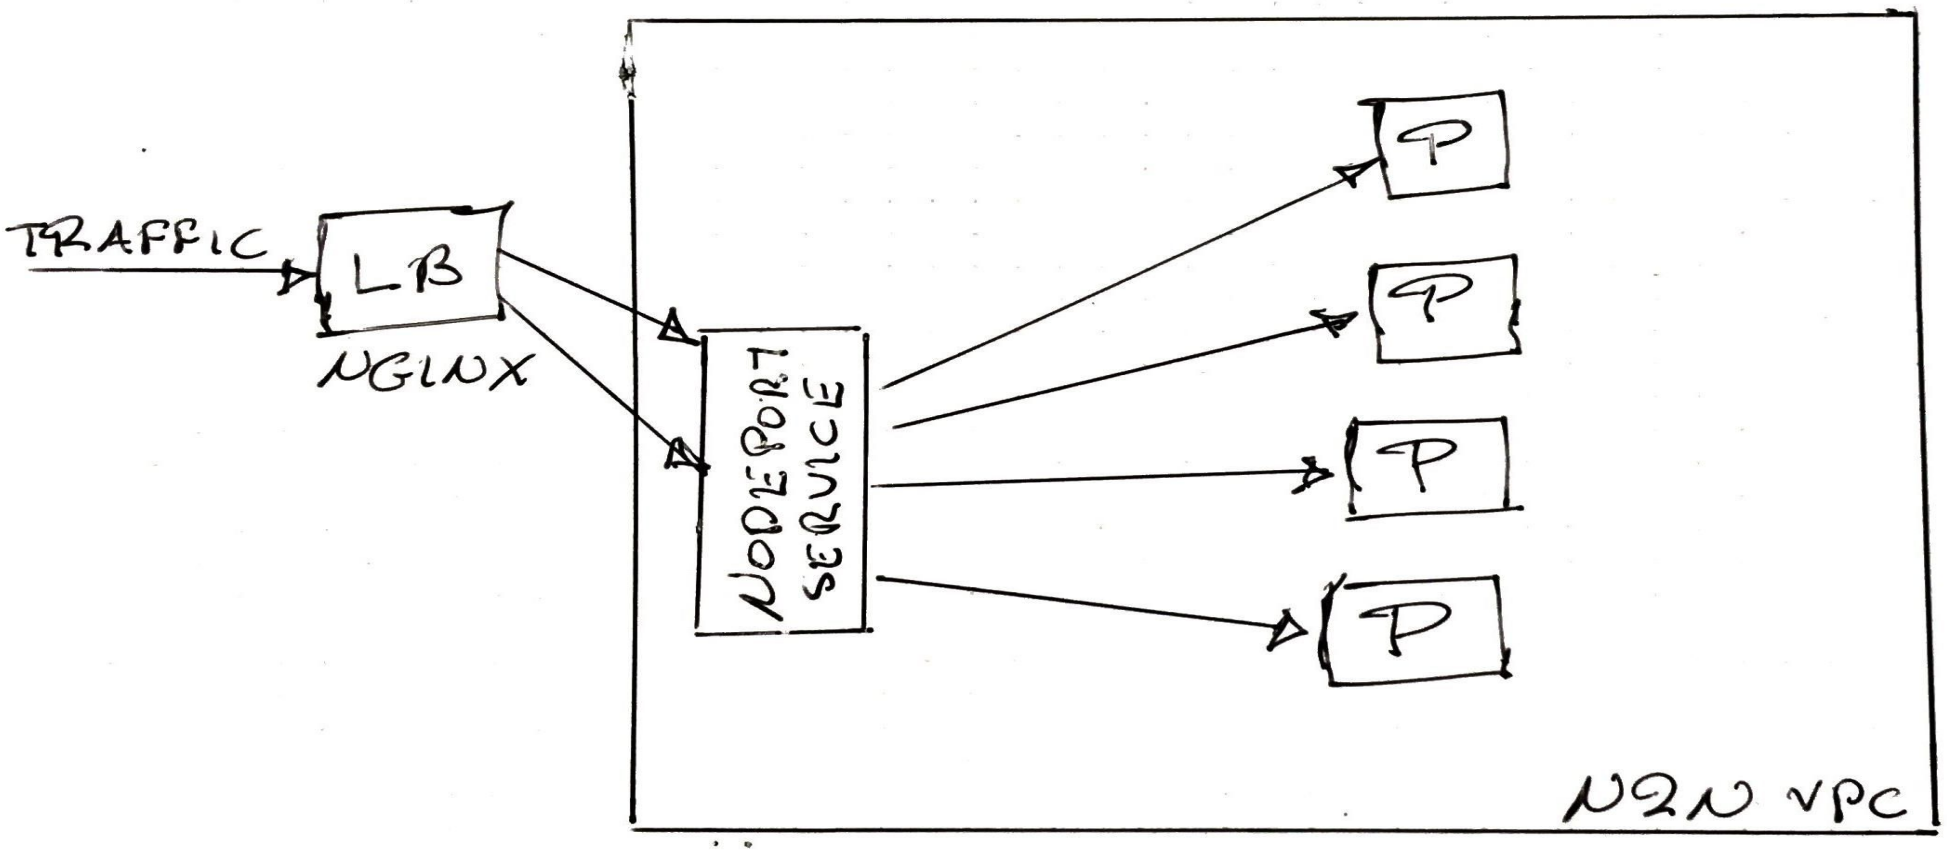

http://gr.linkedin.com/in/yiorgos

Twitter: @hakmem

THANK YOU!