

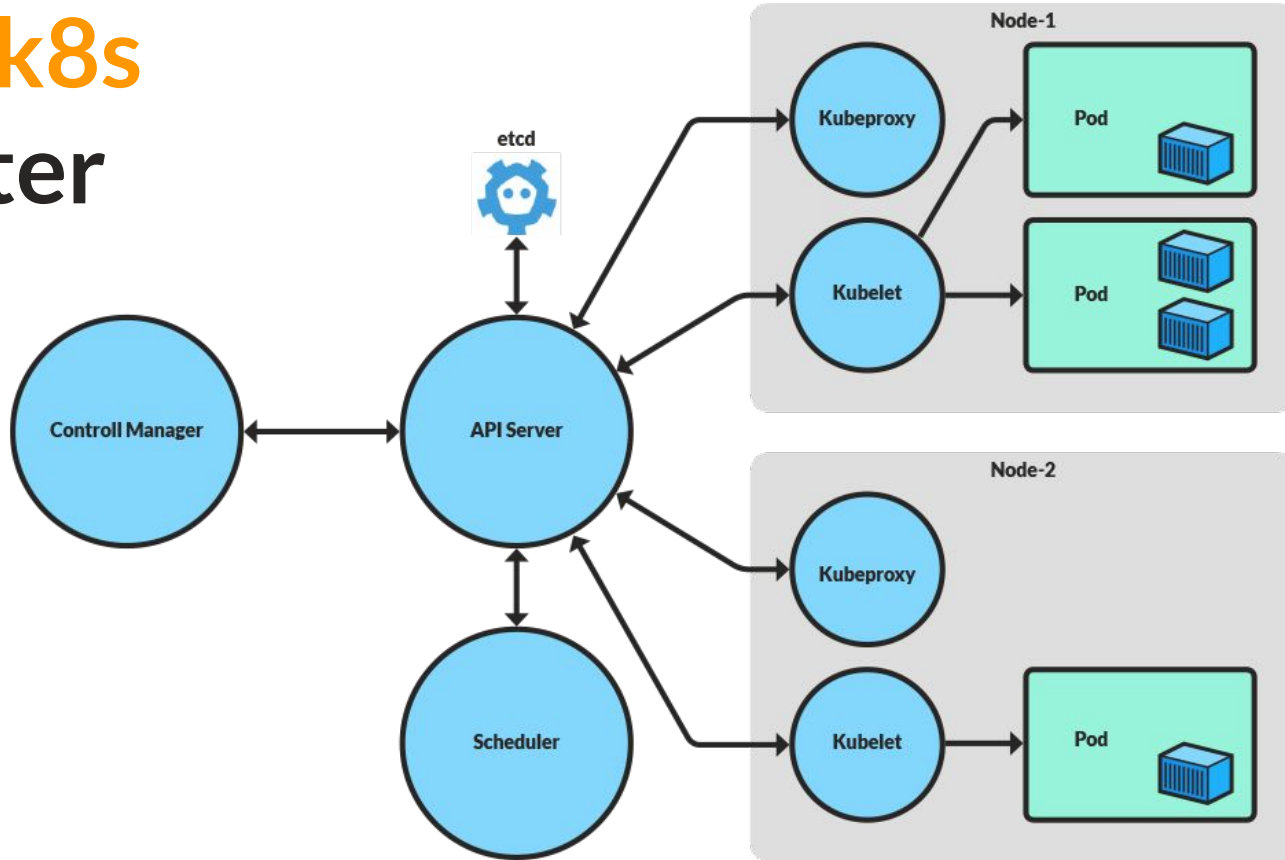
High Availability Kubernetes on Bare-Metal

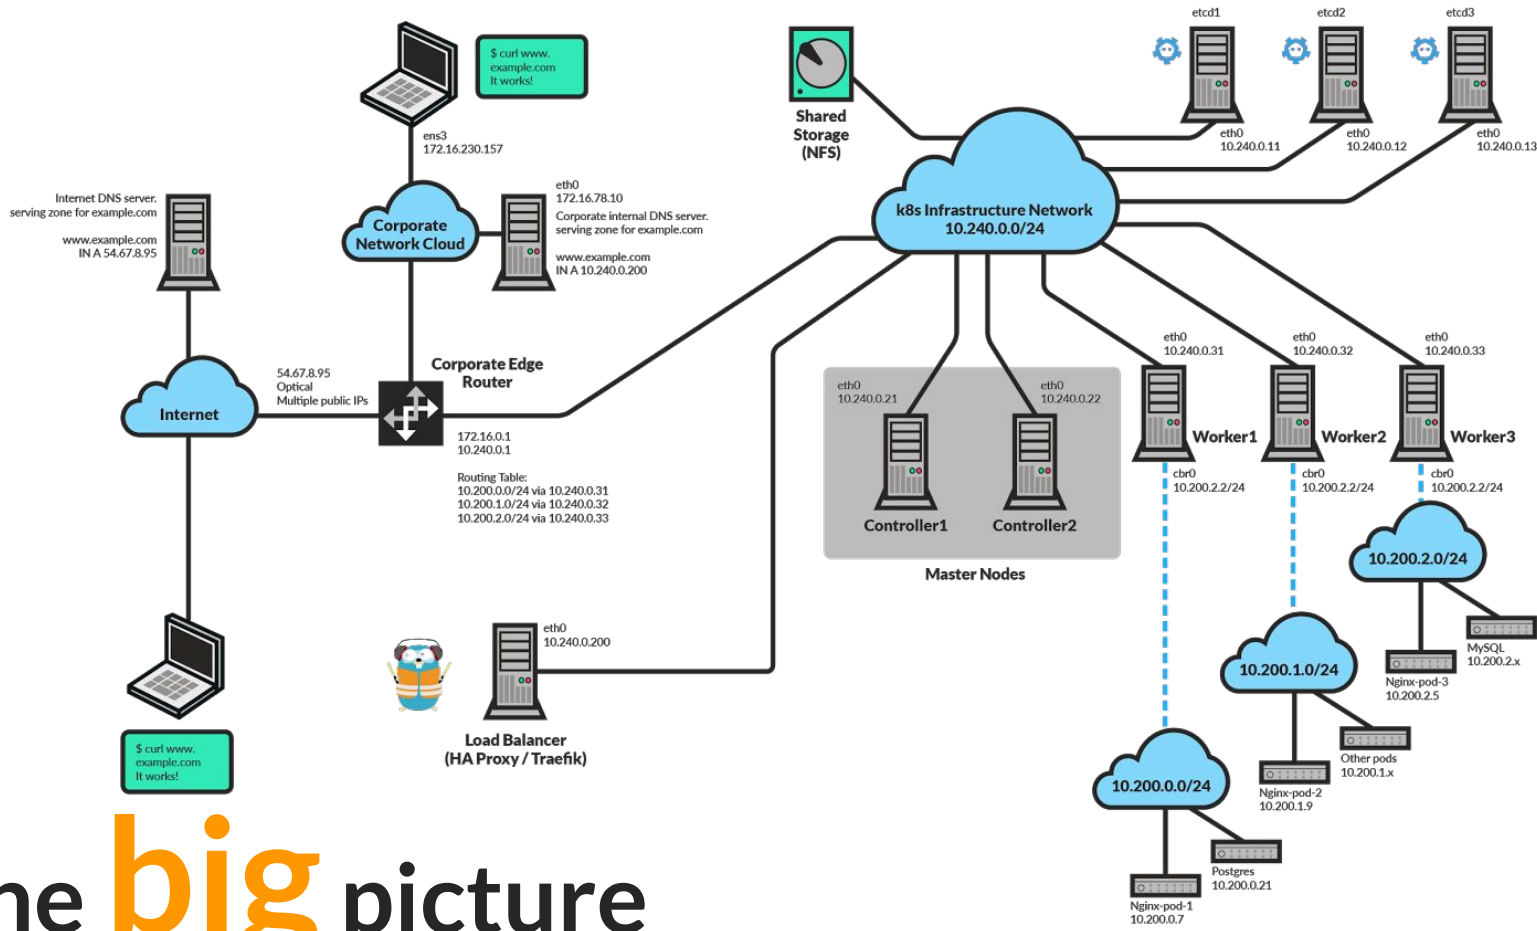
By Muhammad Kamran Azeem & Henrik Høegh



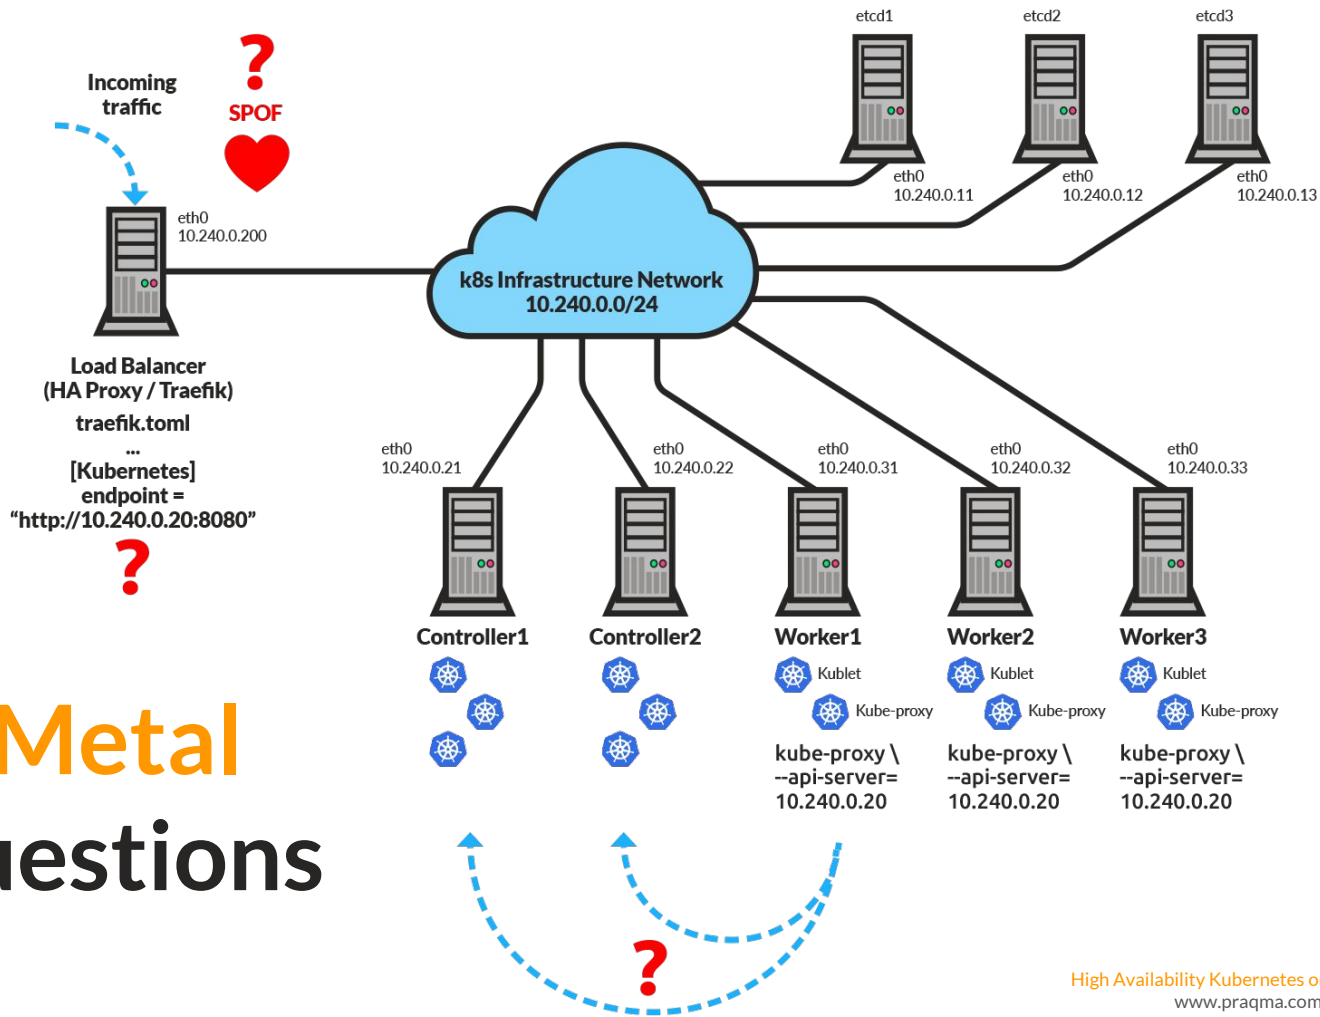
The k8s cluster

The k8s cluster





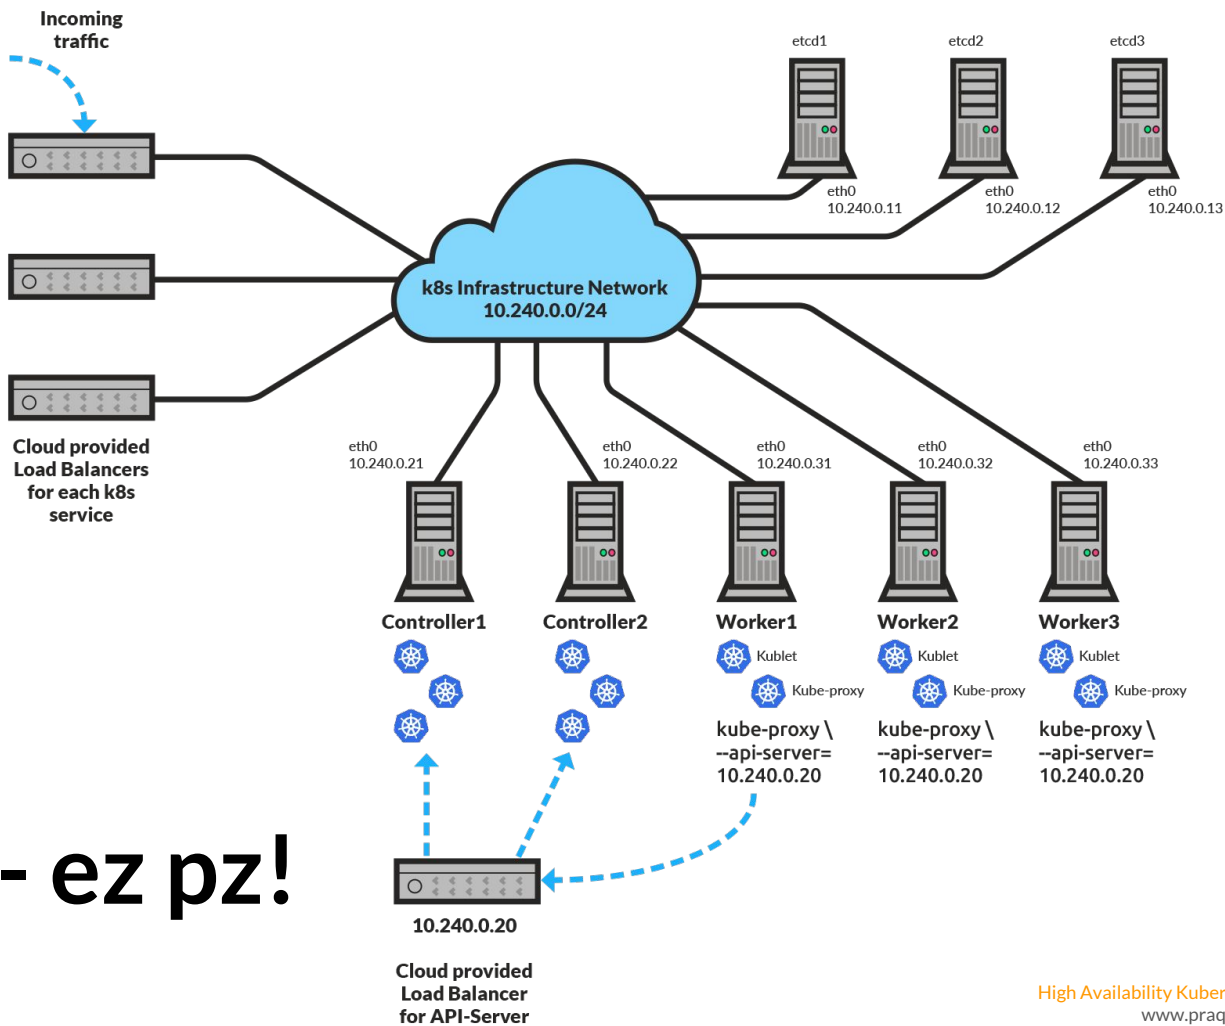
The big picture



Bare-Metal

Big questions

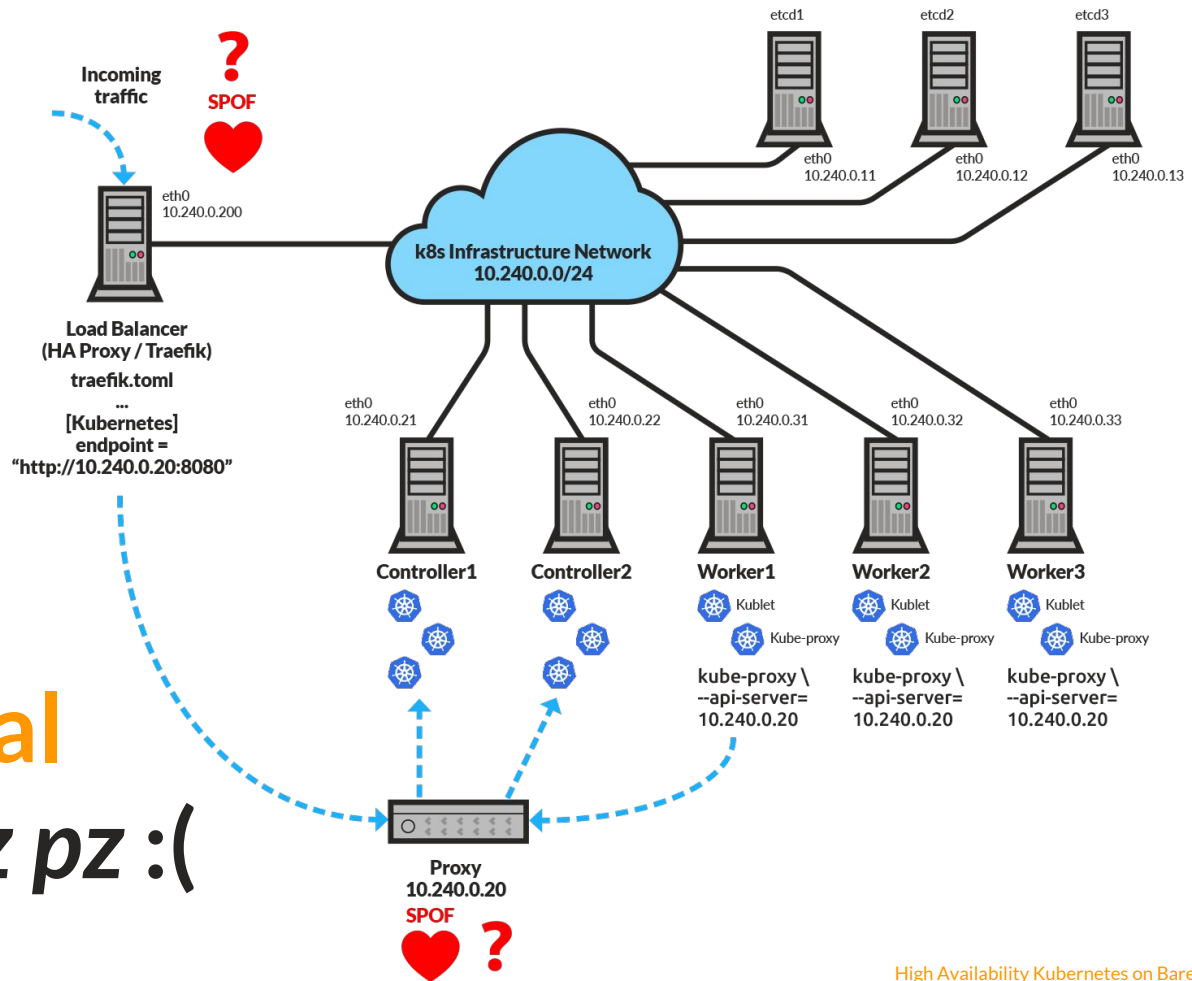
*In the cloud we can simply
use load balancers available
from the cloud providers.*



Cloud - ez pz!



**So - What is the
best way to
introduce HA to
k8s cluster on
bare-metal
infrastructure?**



Bare-metal
...not so ez pz :(

DNS?

- *TTL & Cache issues*
- *DNS becomes Single Point Of Failure*

... no way!

Meet Corosync & PaceMaker

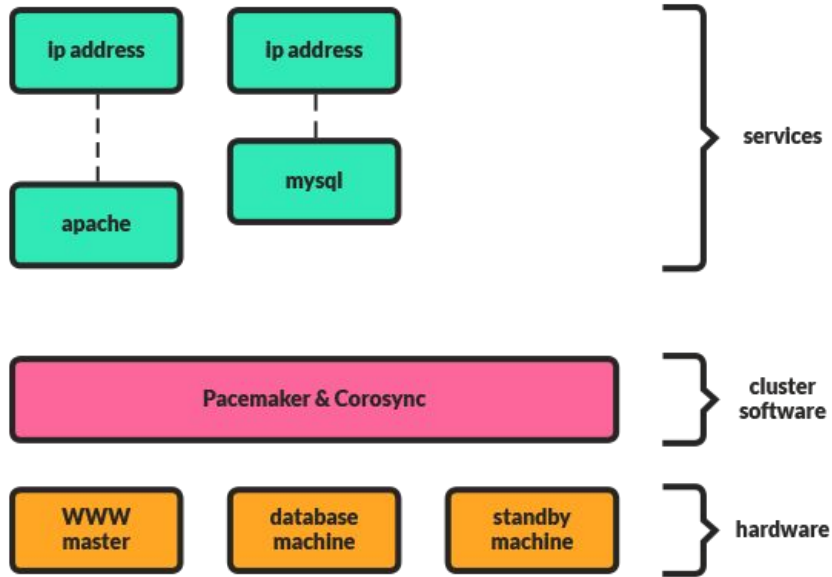
Corosync

The Corosync Cluster Engine



A scalable High Availability
cluster resource manager.

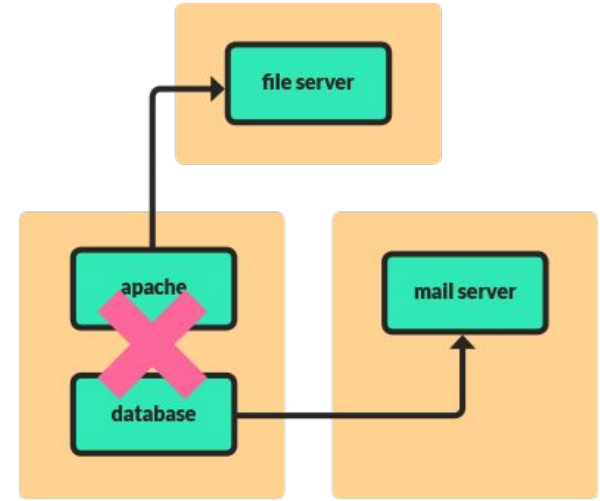
Pacemaker + Corosync - Visual Representation



Deploy

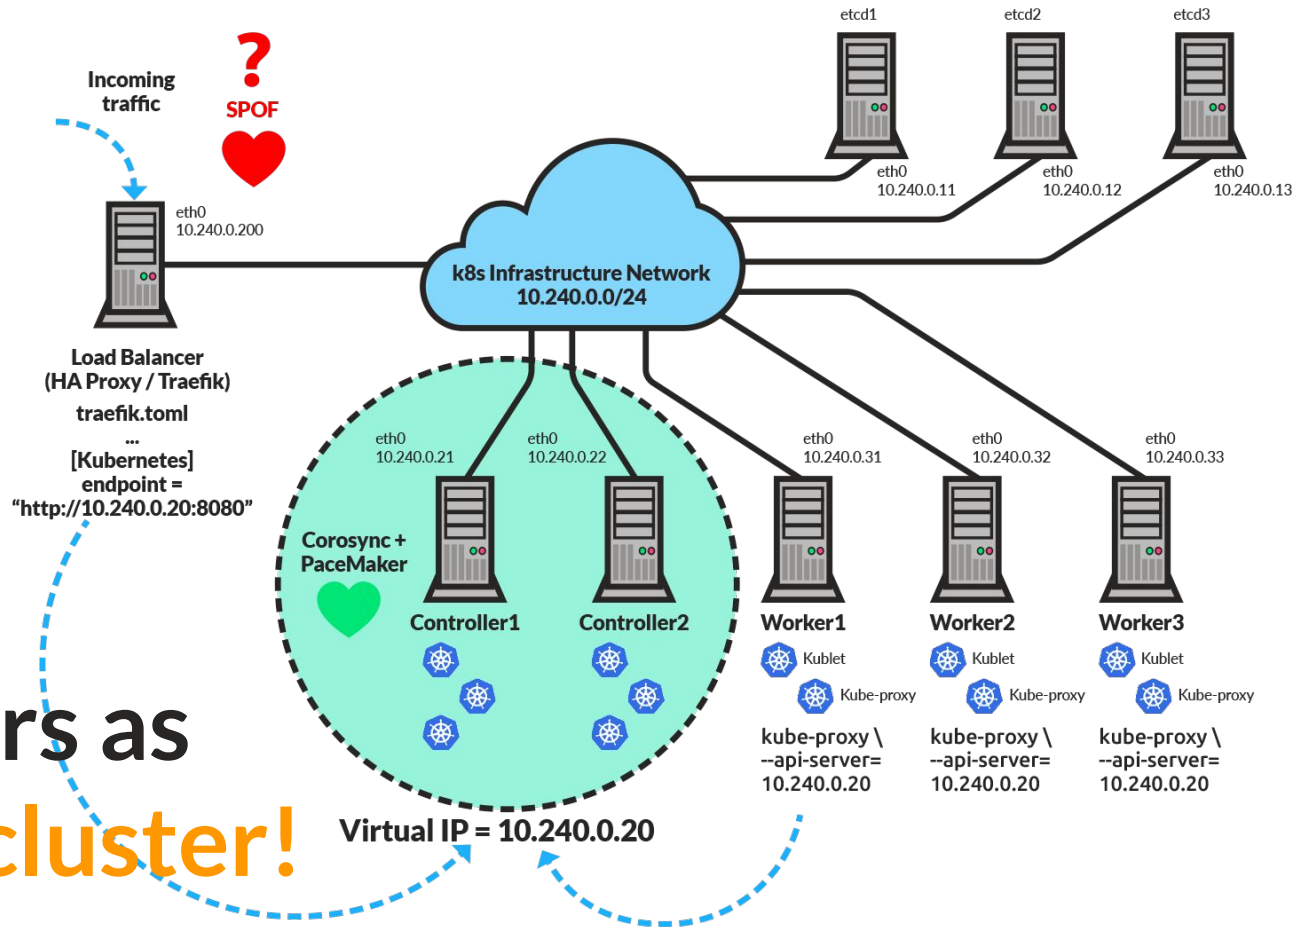


Monitor



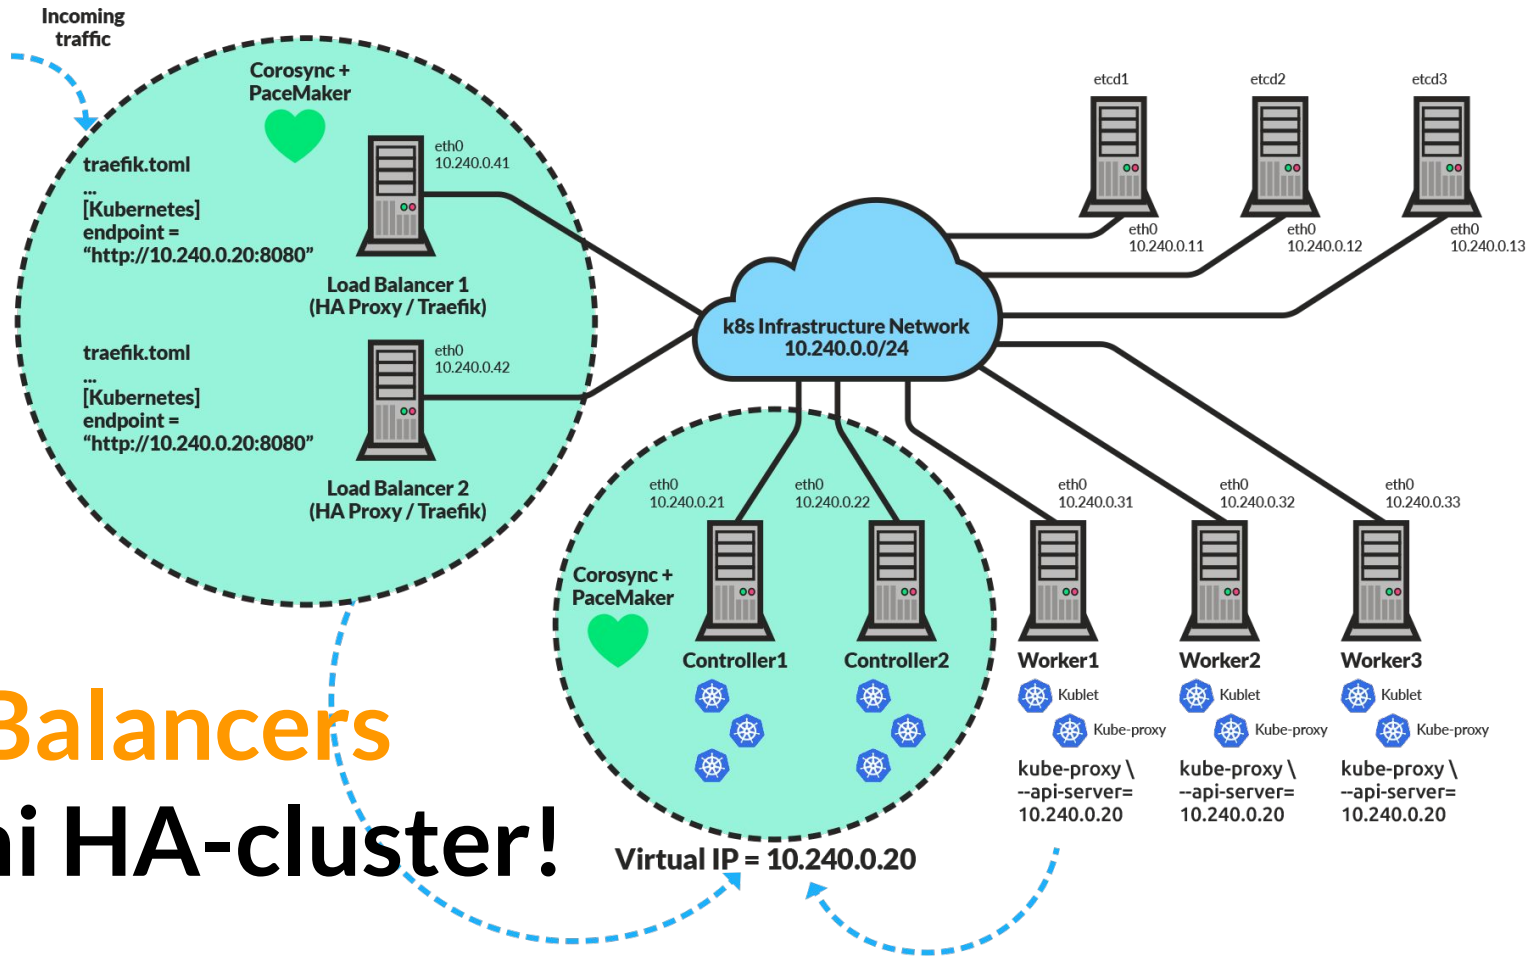
Recover

*So ... Let's configure
controllers as a mini-cluster!
... with heartbeat!*



Controllers as mini HA-cluster!

... *and Load balancers*

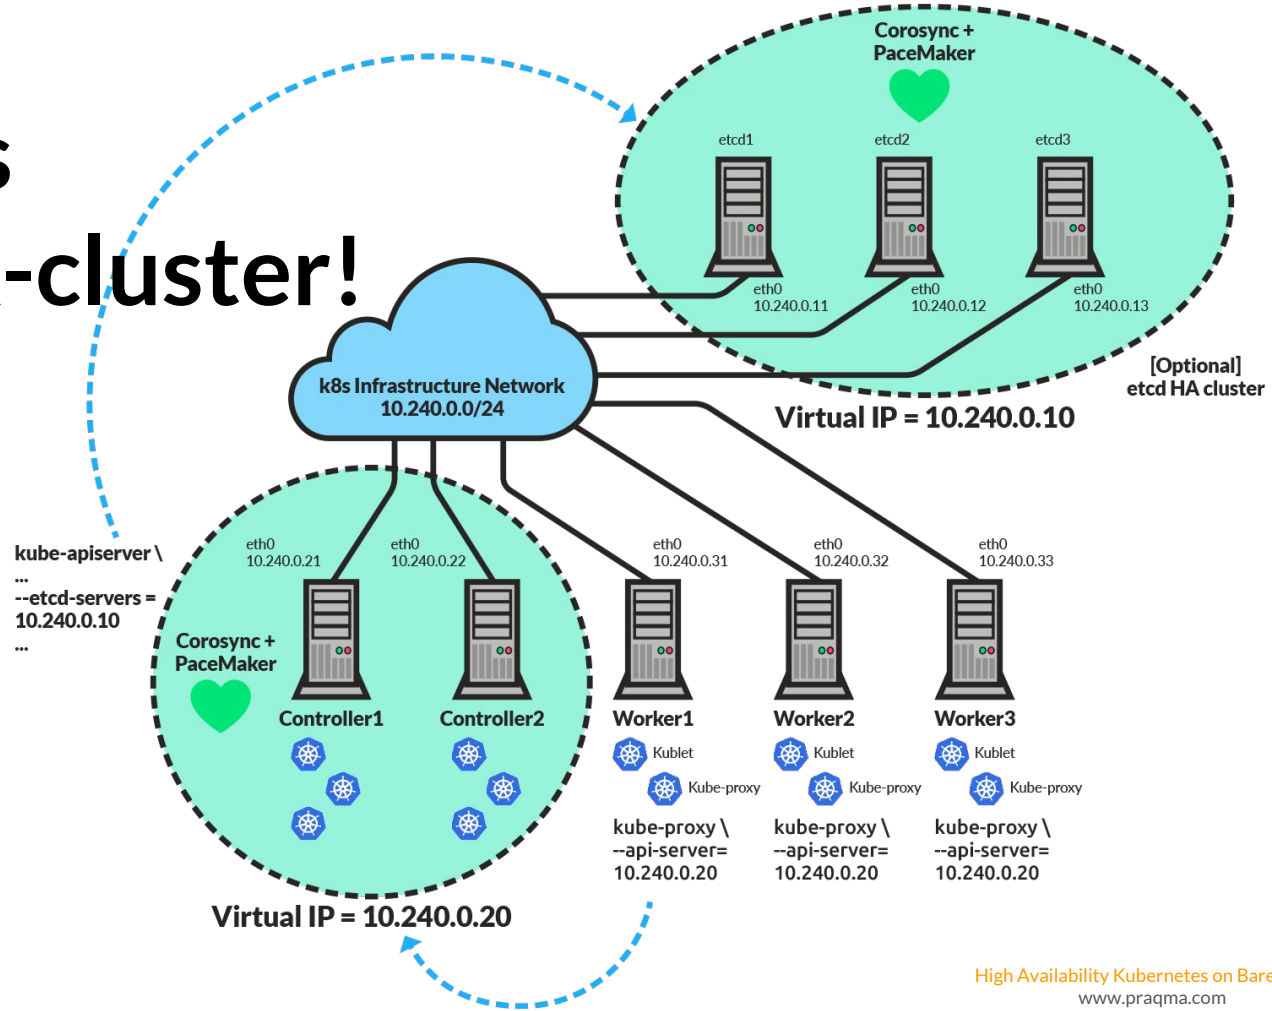


Load Balancers as mini HA-cluster!

... and etcd nodes!

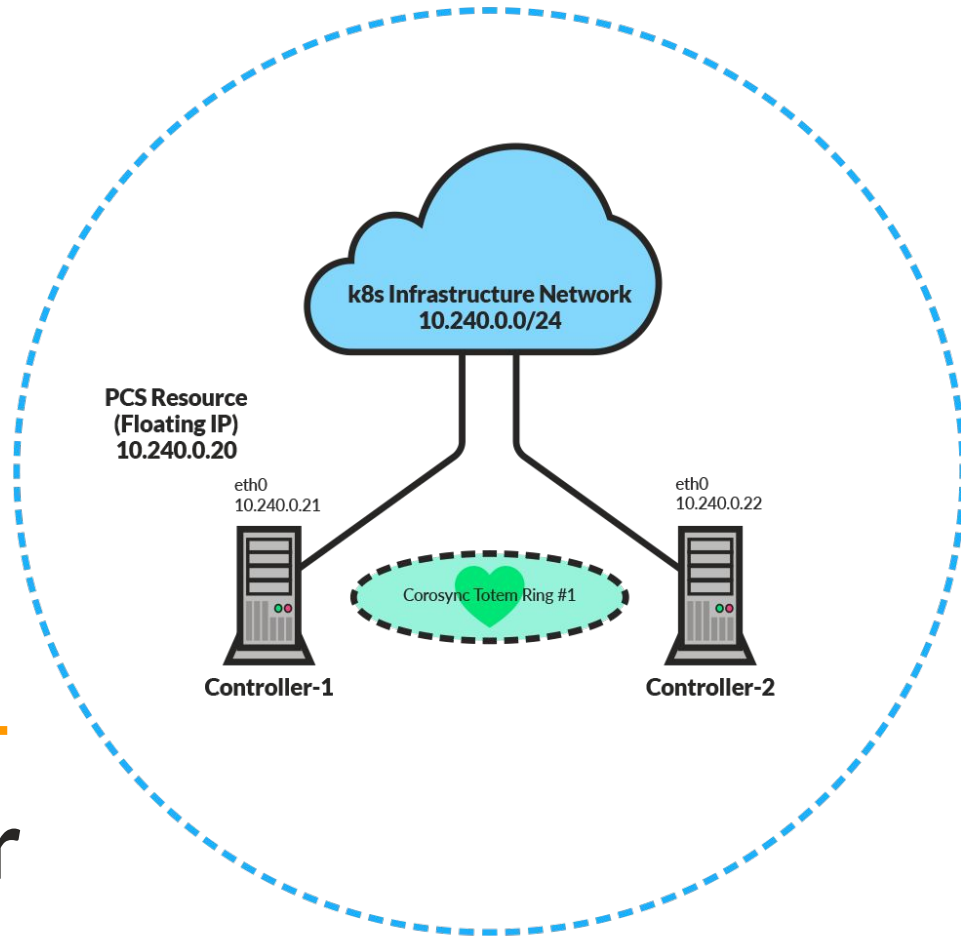
[Optional]...

ETCD as mini HA-cluster!



Some explanation ...

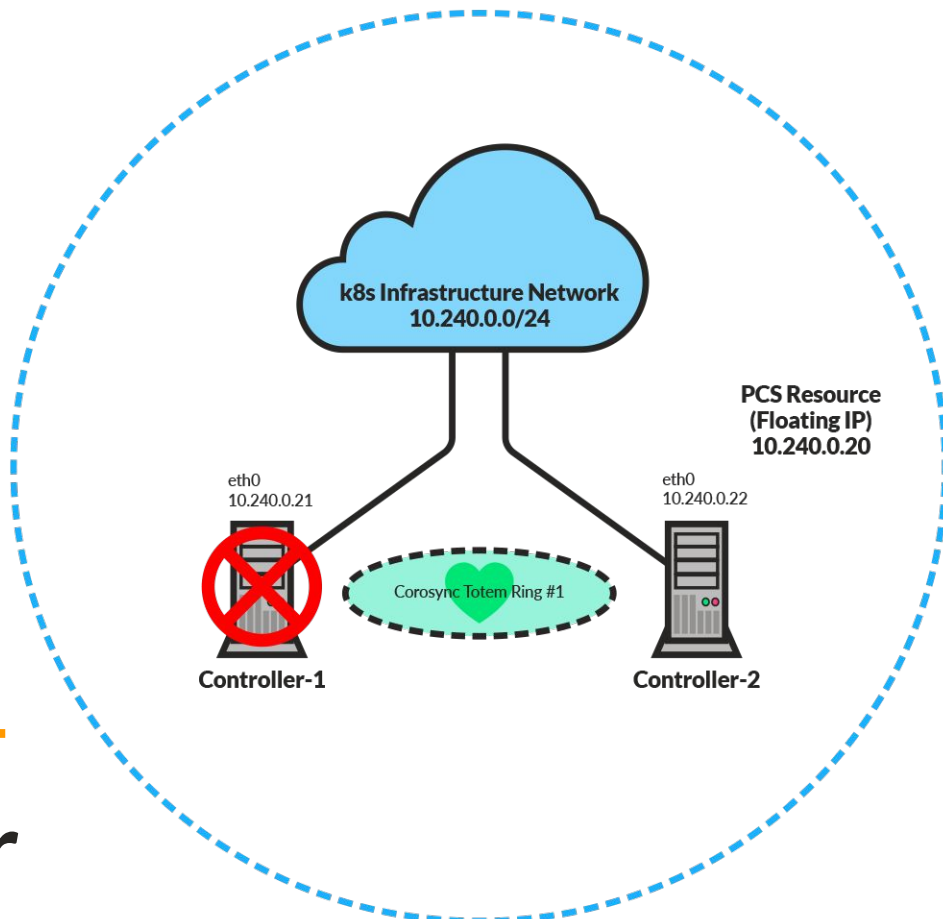
Corosync + PaceMaker



Virtual IP = 10.240.0.20


Node failure
& Recovery

Corosync + PaceMaker



Virtual IP = 10.240.0.20


Multiple
Communication
Rings



Corosync + PaceMaker

Setup

Note: Requires Broadcast and Multicast



Q. Does Amazon VPC support **multicast or **broadcast**?**

No.

So ...

Use VMs on local hypervisor

or

Use bare-metal!

On all (pacemaker) cluster nodes:

```
# yum install corosync pacemaker pcs
```

```
pcs cluster setup --name ControllerHA \  
  controller-1 controller-2 \  
  --force
```

```
pcs cluster start --all
```

```
pcs cluster enable --all
```

```
pcs resource create ControllerVIP \  
  ocf:heartbeat:IPaddr2 ip=10.240.0.20 \  
  cidr_netmask=32 op monitor interval=30s
```

```
[root@controller1 ~]# pcs status
Cluster name: ControllerHA
Stack: corosync
Current DC: controller2.praqma.local - partition with quorum
Last updated: 11:18:22 2017 Last change: 18:48:25 2017 by hacluster via crmd on
controller1.praqma.local
```

```
2 nodes and 1 resource configured
```

```
Online: [ controller1.praqma.local controller2.praqma.local ]
```

```
Full list of resources:
```

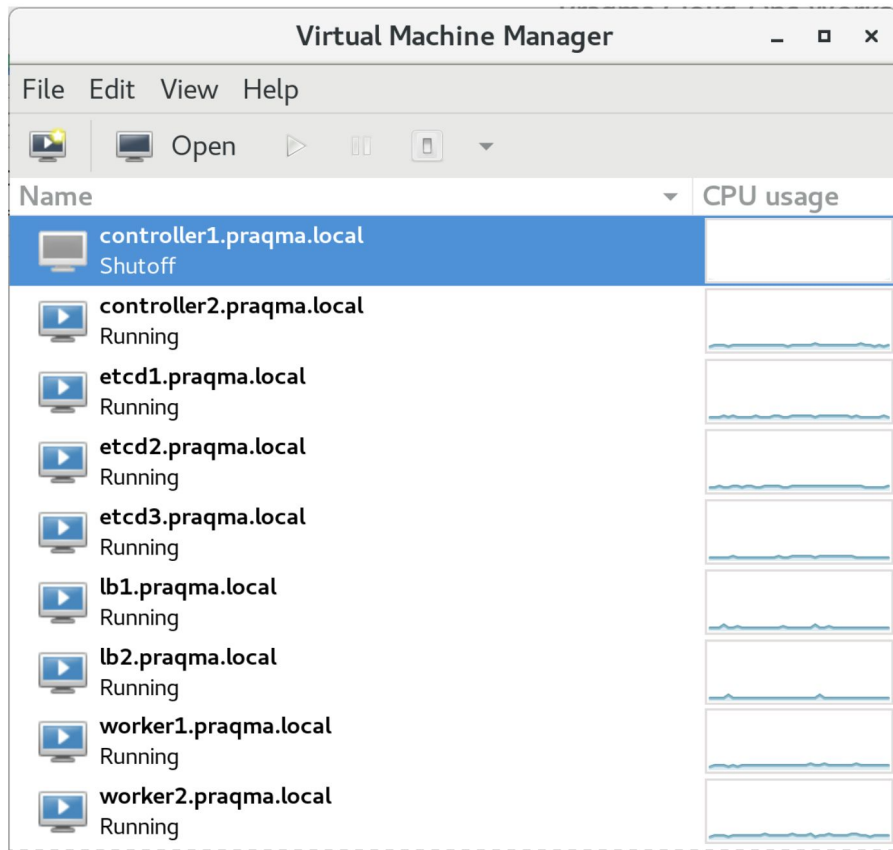
```
ControllerVIP (ocf::heartbeat:IPaddr2): Started controller1.praqma.local
```

```
Daemon Status:
```

```
corosync: active/enabled
pacemaker: active/enabled
pcsd: active/enabled
```

```
[root@controller1 ~]#
```

Demo !!!



- [Link to Kubernetes book \(WIP\)](#)

Thanks!

Questions and comments?

kaz@pragma.net , heh@pragma.net