ERKAN YANAR

DOCKER - RETHINKING INFRASTRUCTURE

Erkan Yanar

linsenraum.de

Consulting

Training


openstack™
GALERA CLUSTER
MariaDB
puppet labs
docker
LXC
ANSIBLE
Pacemaker
MySQL®

linsenraum.de


OSDC.de
OPEN SOURCE DATA CENTER CONFERENCE
German Unix User Group
guug
OPEN RHEIN RUHR
Ein Pott voll Software
LINUX TAG
openstack
D A CH e.V.
FRANKFURTER DATENBANKTAGE
2013
committer conf
DevOpsCon
Chemnitzer Linux-Tage
15. und 16. März 2014
iX OpenStack Tag
Die iX Konferenz am 15. April 2015
Eine Veranstaltung von
heinlein
DOAG
Konferenz + Ausstellung


FH FFM
T Systems
HOST EUROPE
trust
INTERNATIONAL
held informatik
DEMOTIX
my xplace
1&1
codership
ch.de
DER Touristik
Reply
ch | schnell
triplesense
hastexo!
Xchanging
linuxhotel
Training & Coworking bei den OpenSource'lern
teuto.net
heinlein

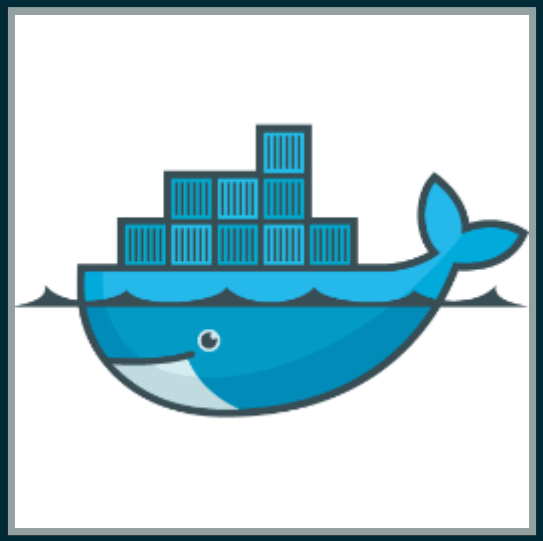

DOCKER HYPE

- Redhat
- Ubuntu
- Google
- Microsoft
- VMWare
- IBM
- OpenStack
- Atomic/Openshift
- Snappy
- Kubernetes
- Azure
- Fargo
- ...
- ...


A magical unicorn! How wonderful!
shoeboxblog.com

AGENDA

- Container vs. Hypervisor
- Kernelfeatures
- Docker
- Architektur Docker
- Architekturen mit Docker
- Docker Ecosystem

DOCKER - DAS WIE

Teil der Containerfamilie:

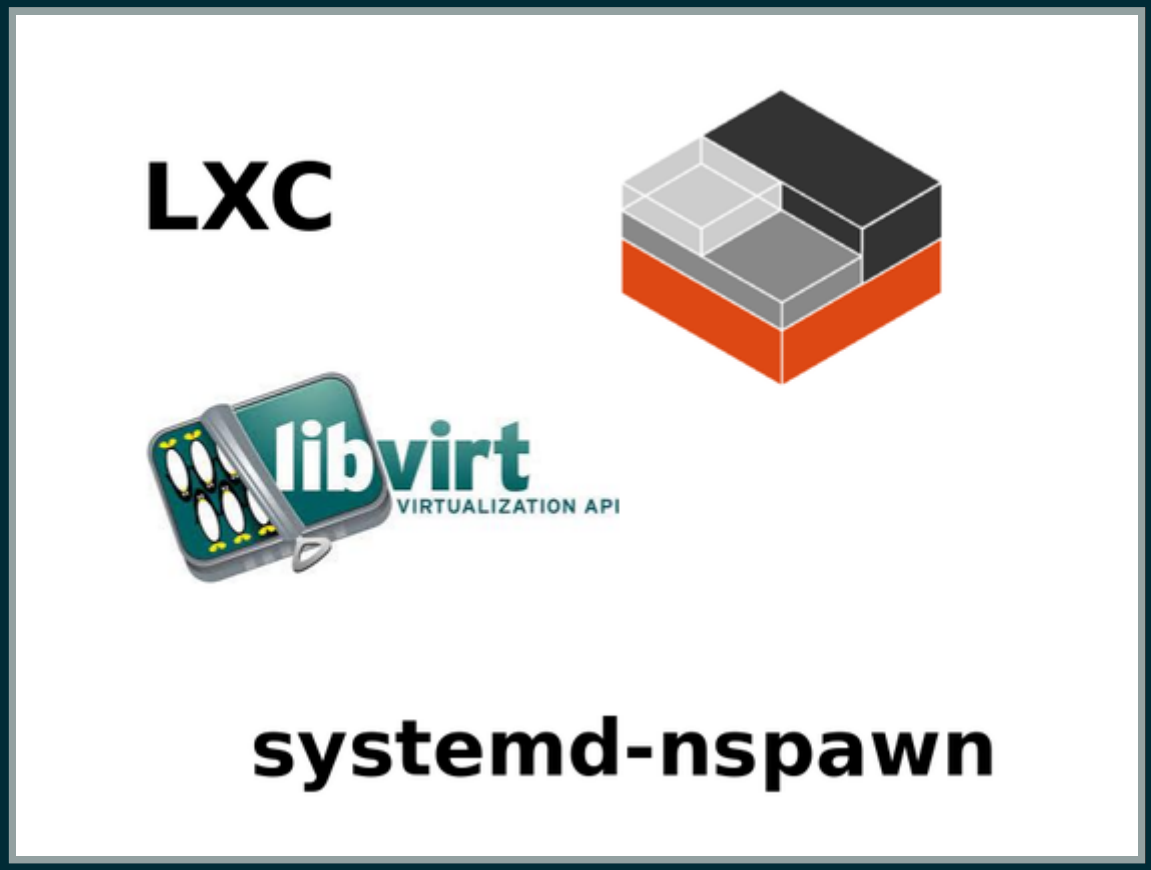

CONTAINER VS. HYPERVISOR

- OS Virtualisierung

vs

- Emulation von Hardware

DOCKER

Auch nur Kerneltechnik

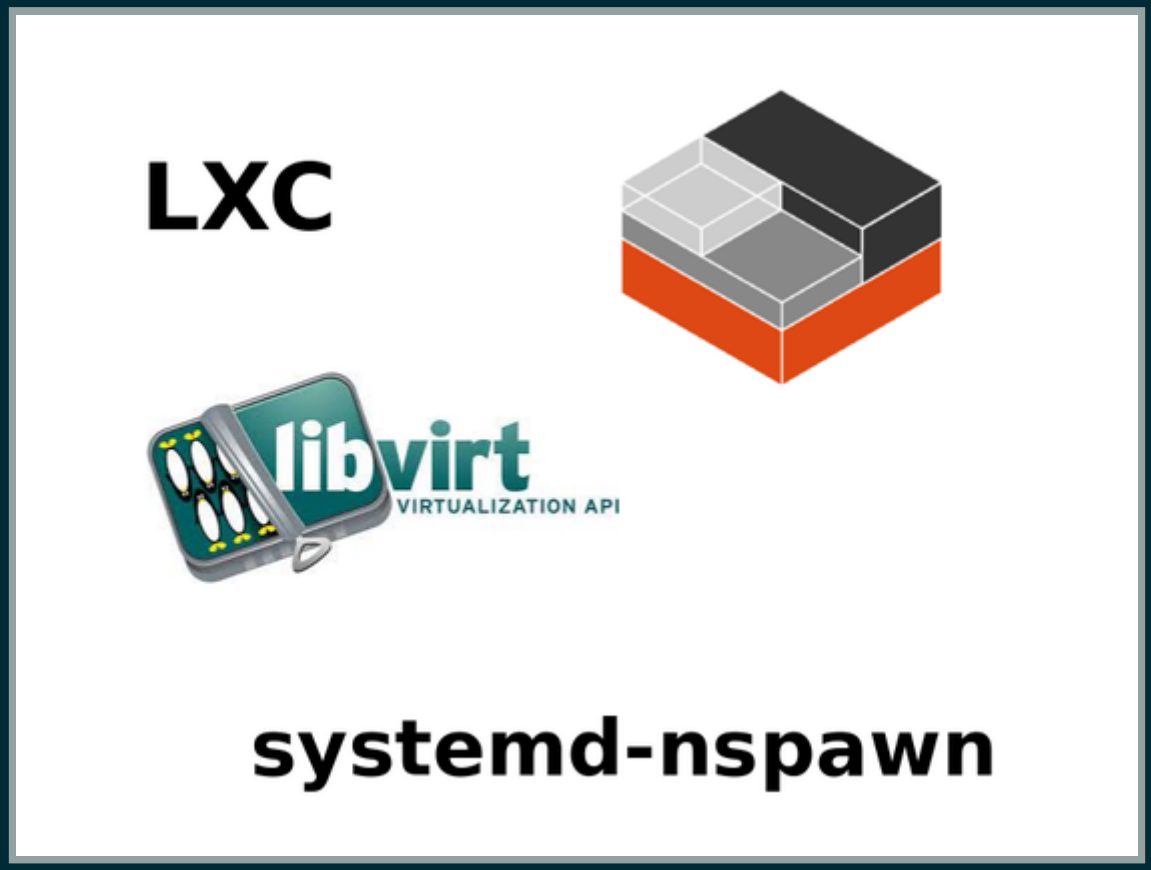

CONTAINERKERNEL

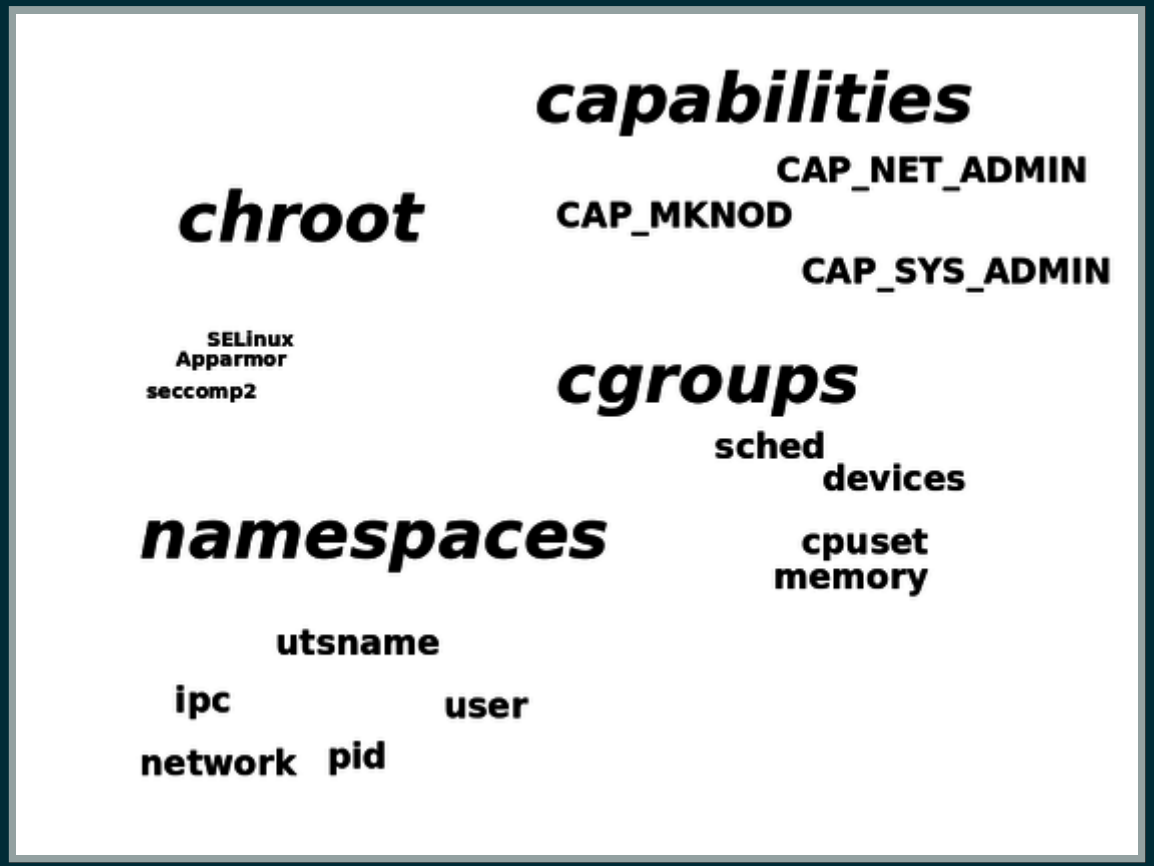

NAMESPACES

```
$ readlink  /proc/self/ns/*
ipc:[4026531839]
mnt:[4026531840]
net:[4026531968]
pid:[4026531836]
user:[4026531837]
uts:[4026531838]
```

Docker implementiert (noch) keinen user-ns.

CAPABILITIES

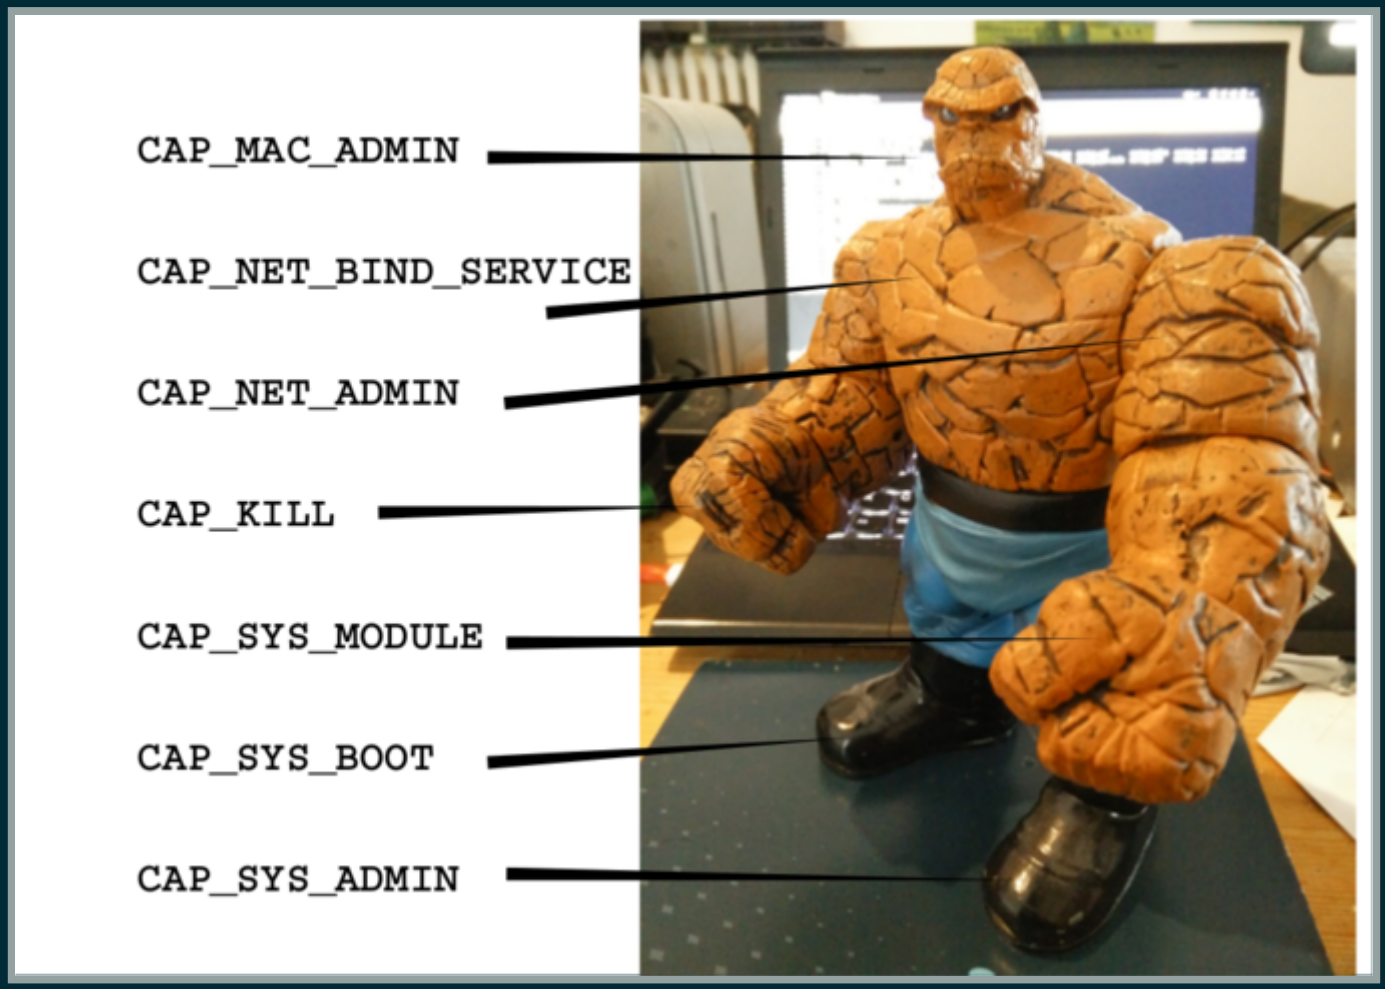

CHROOT

filesystem.png

CGROUPS

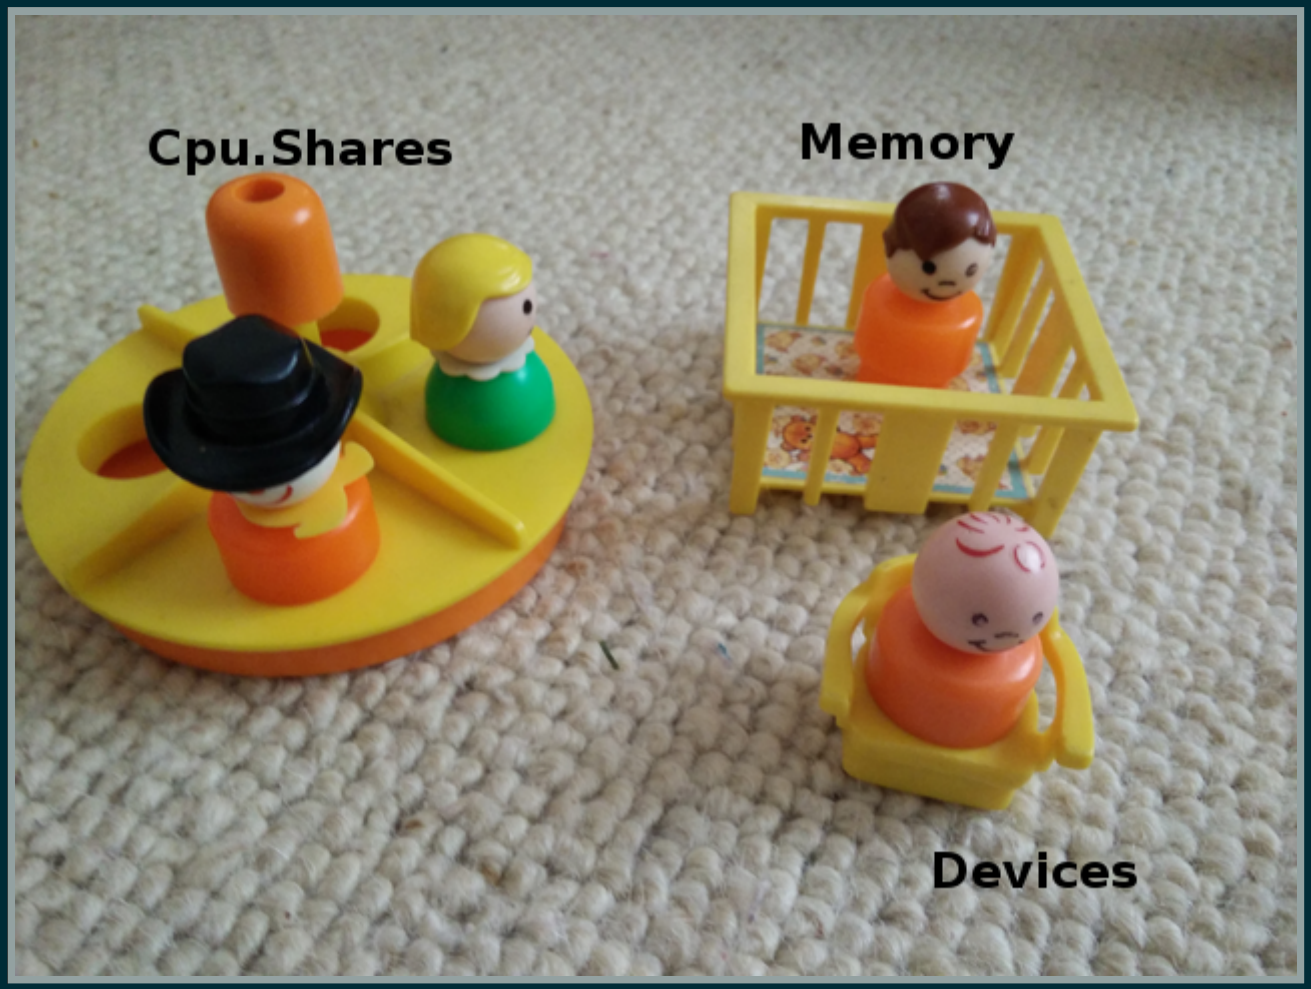

DOCKER

Auch nur Kerneltechnik

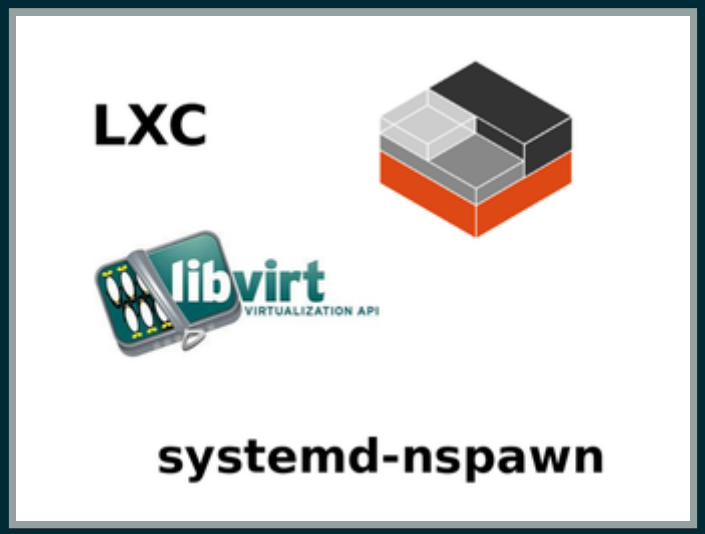

ABER:

DOCKER

Specials:

- Applikationscontainer
- Virtuelles Binary
- Client/Server (REST)
- Portmapping

PORTMAPPING

- NATing der Container
- Ports zum Contaier werden geDNATed

portmapping.png

DOCKER

Specials:

- Transportabele/OS independant Images
- Registry
- Image vs. Container
- Dockerfile

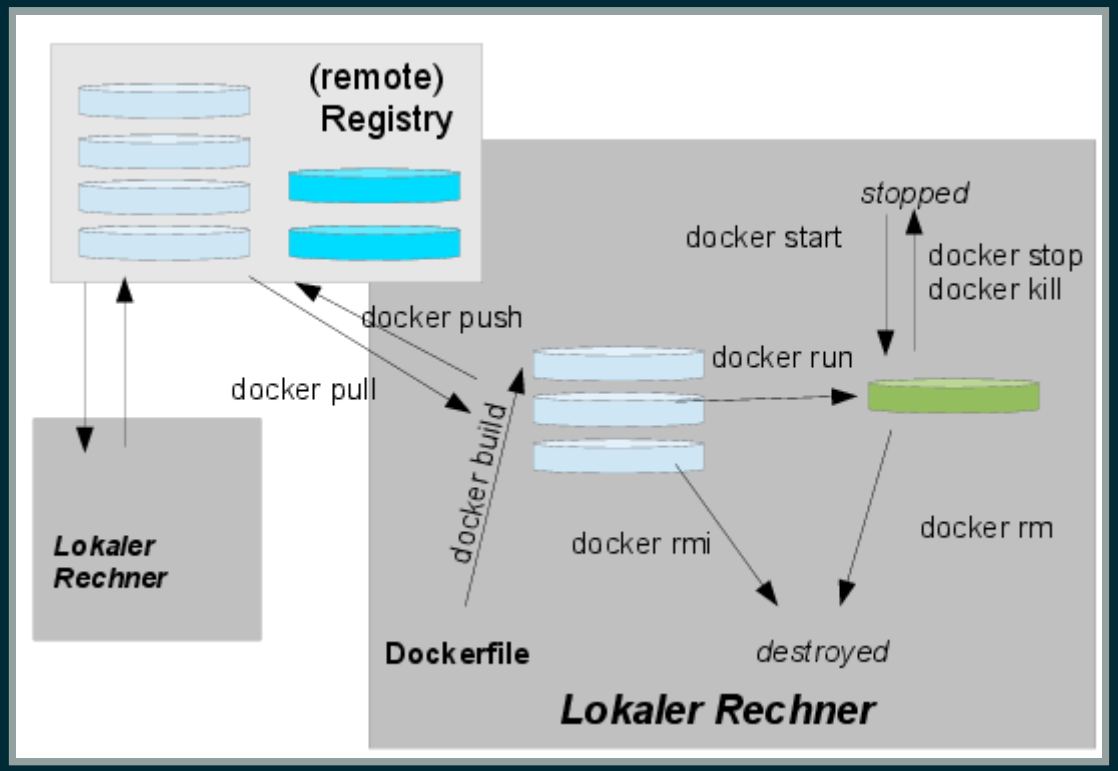
(remote)
Registry
stopped
docker start
docker stop
docker kill
docker push
docker run
docker pull
docker build
Lokaler
Rechner
docker rmi
docker rm
Dockerfile
destroyed
Lokaler Rechner

IMAGES && DOCKERFILE

IMAGE BUILD

```
FROM fedora:heisenbug
MAINTAINER erkan.yanar@linsenraum.de

RUN yum install -y nginx
RUN echo "daemon off;"  >>/etc/nginx/nginx.conf
COPY index.html /usr/share/nginx/html/index.html

EXPOSE 80
ENTRYPOINT ["nginx"]
```

```
docker build -t erkules/nginx .
```

CONTAINER

- Execution Driver (-e)
- Storage Driver (-s)

```
aufs (another unionfs)
devicemapper
btrfs  (B-tree FS)
overlay
```

filesystem.png

DAS WARUM

- Ressourcenschonender
- Build Ship Deploy
- *Works for me* in die Produktion
- Devops

DAS WARUM

- Golden Image
- Microservices
- Immutable Infrastructure

GOLDEN IMAGE

- Vorteile von Golden Image
- Ohne die Nachteile
- Wegen Layerkonzept

MICROSERVICE

IMMUTABLE INFRASTRUCTURE

- Keine Konfiguration zur Laufzeit
- Komponenten (Services) werden ersetzt
- Statless vs. State
- Fehlertolerant
- Versionierte Infrastruktur

HANDSON

```
$
```

DESKTOP/GUI

```
docker run --rm -v /dev/snd:/dev/snd -v /tmp/.X11-unix:/tmp/.X11-unix
      -e DISPLAY=$DISPLAY --privileged erkules/firefox --new-instance
docker run --rm  -v /tmp/.X11-unix:/tmp/.X11-unix -e DISPLAY=$DISPLAY
       erkules/firefox --new-instance
```

SERVER BEISPIELE

100x VM oder 100x Docker?

```
for i in $(seq 1 100); do docker run -d -P --name nginx_$i nginx ; done
docker ps -s
docker ps -q | xargs docker rm -f
```

KOMMUNIKATION

- --link
- --volumes-from

```
mkdir /tmp/SLAC2015
docker run --rm --name mysql -e MYSQL_USER=slac -e MYSQL_PASSWORD=slac
      -e MYSQL_DATABASE=slac -v /tmp/SLAC2015:/var/lib/mysql der_mysql
```

```
docker run -P --rm --name wordi --link mysql:mysql -e WORDPRESS_DB_USER=s
          -e WORDPRESS_DB_PASSWORD=slac -e WORDPRESS_DB_NAME=slac wordpre
```

```
docker run -P --rm --name wordi --volumes-from mysql -e WORDPRESS_DB_USE
          -e WORDPRESS_DB_PASSWORD=slac -e WORDPRESS_DB_NAME=slac wordpre
```

DAS WAS-NOCH

Docker ist doch nur ein virt. Binary!

- Singlehost
- Service Discovery
- Multi-Host
- Cluster
- Orchestrierung
- Netzwerk

multi-hosts.png

DAS WAS-NOCH

Docker ist doch nur ein virt. Binary!

- Docker Swarm
- Docker Network
- Kubernetes
- Consul
- Docker Compose
- CoreOS
- Mesos
- ...

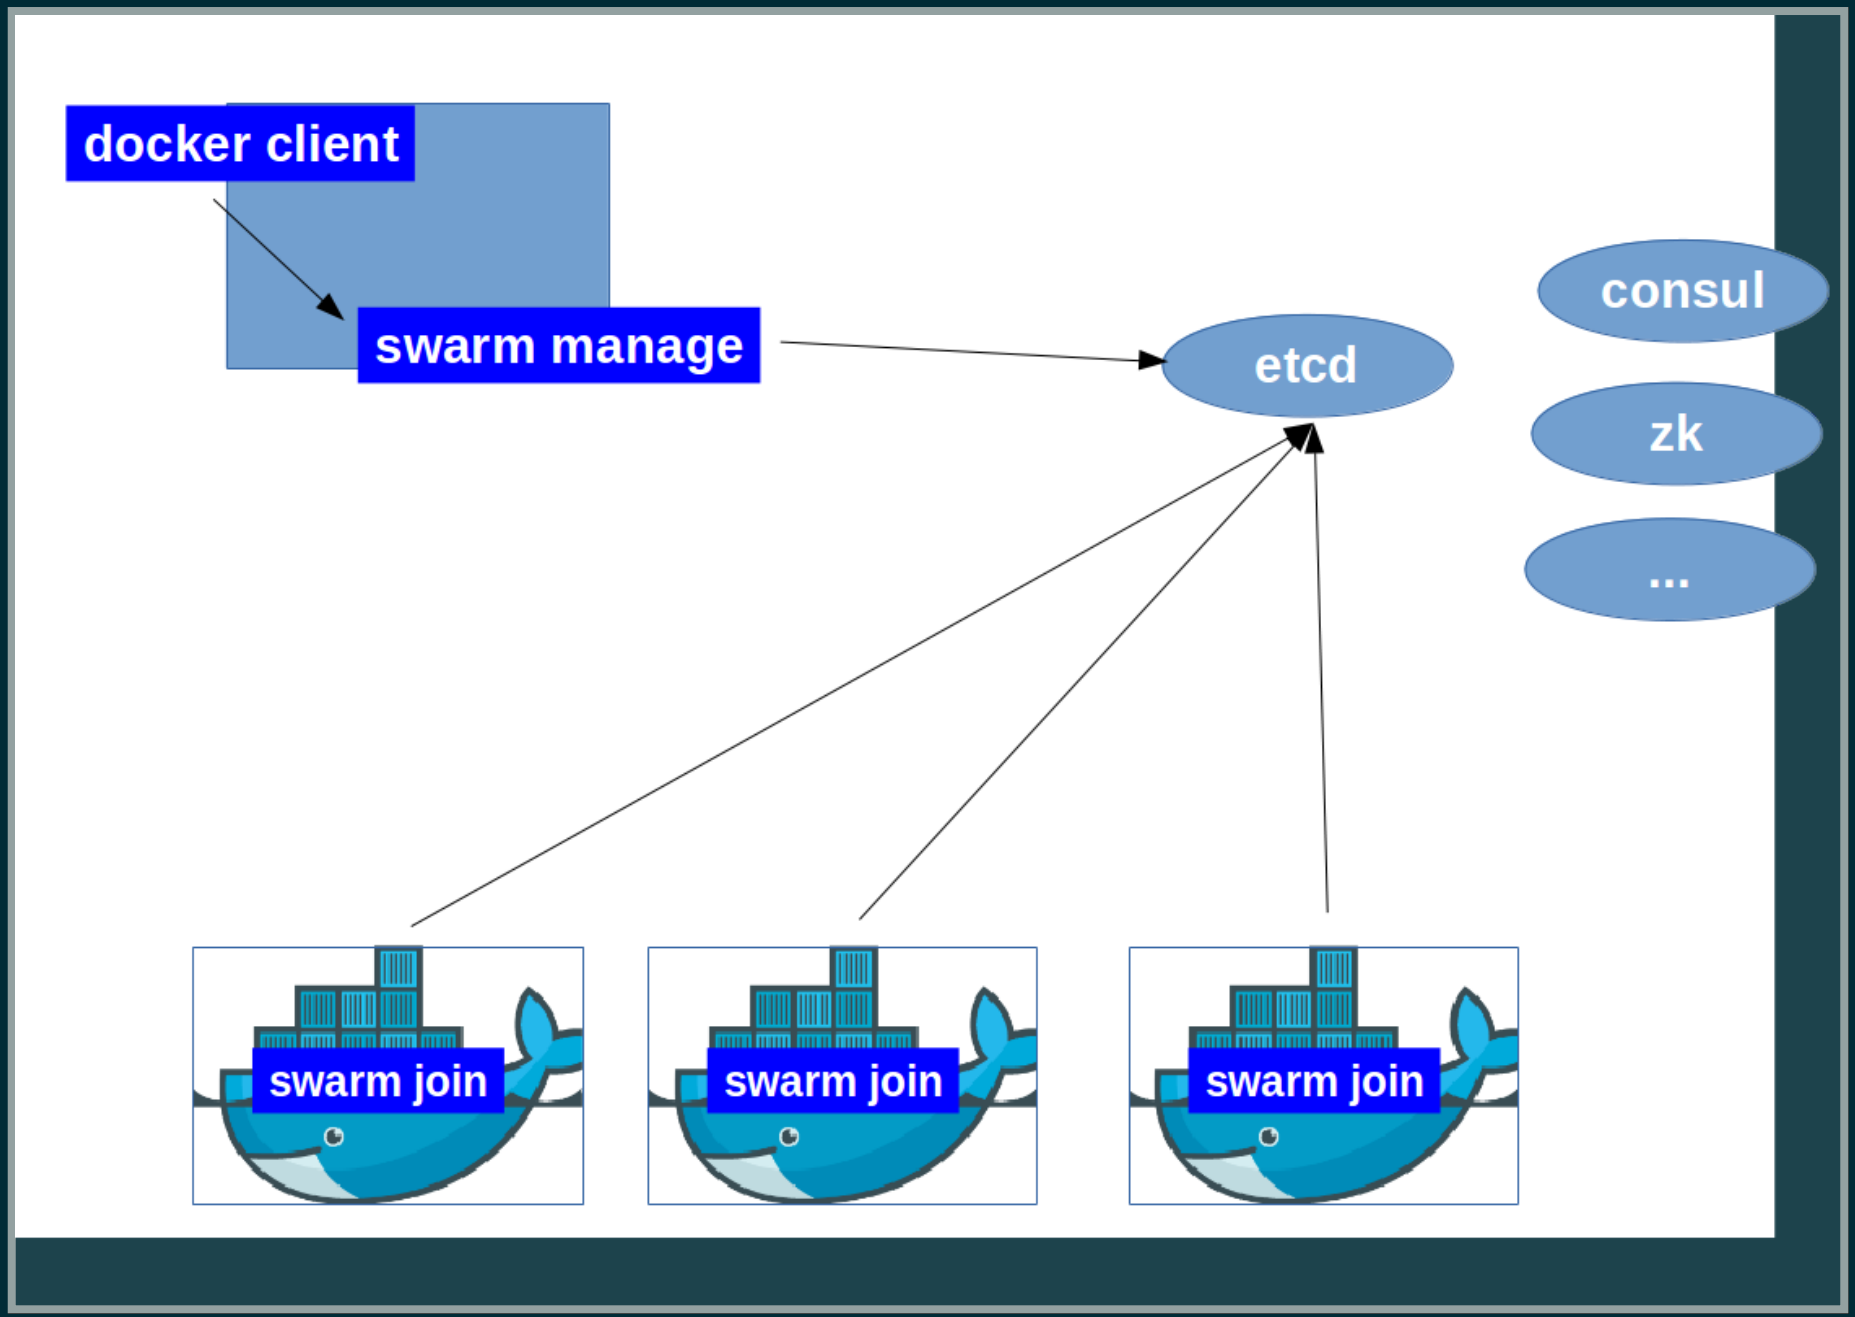
docker client
swarm manage
etcd
consul
zk
...
swarm join
swarm join
swarm join

Daemon:

```
docker -d -H tcp://0.0.0.0:2345
swarm join --addr=10.11.12.8:2345 etcd://10.11.12.8:4001/swarm
```

Manager:

```
swarm manage -H 127.0.0.1:4334 etcd://10.11.12.8:4001/swarm
```

DOCKER COMPOSE

```
web:
  image: nginx
  environment:
    - MYSQL_USER=apli
    - MYSQL_PASS=123
  ports:
    - 80
  links:
    - db:db
datadir:
  image: mysql
  volumes:
    - /srv/mysql:/var/lib/mysql
  net: none
  command: true
db:
  image: mysql
  environment:
```

DOCKER NETWORK

```
socketplane network create web 10.2.0.0/16
socketplane run -n web -itd ubuntu
```


CoreOS

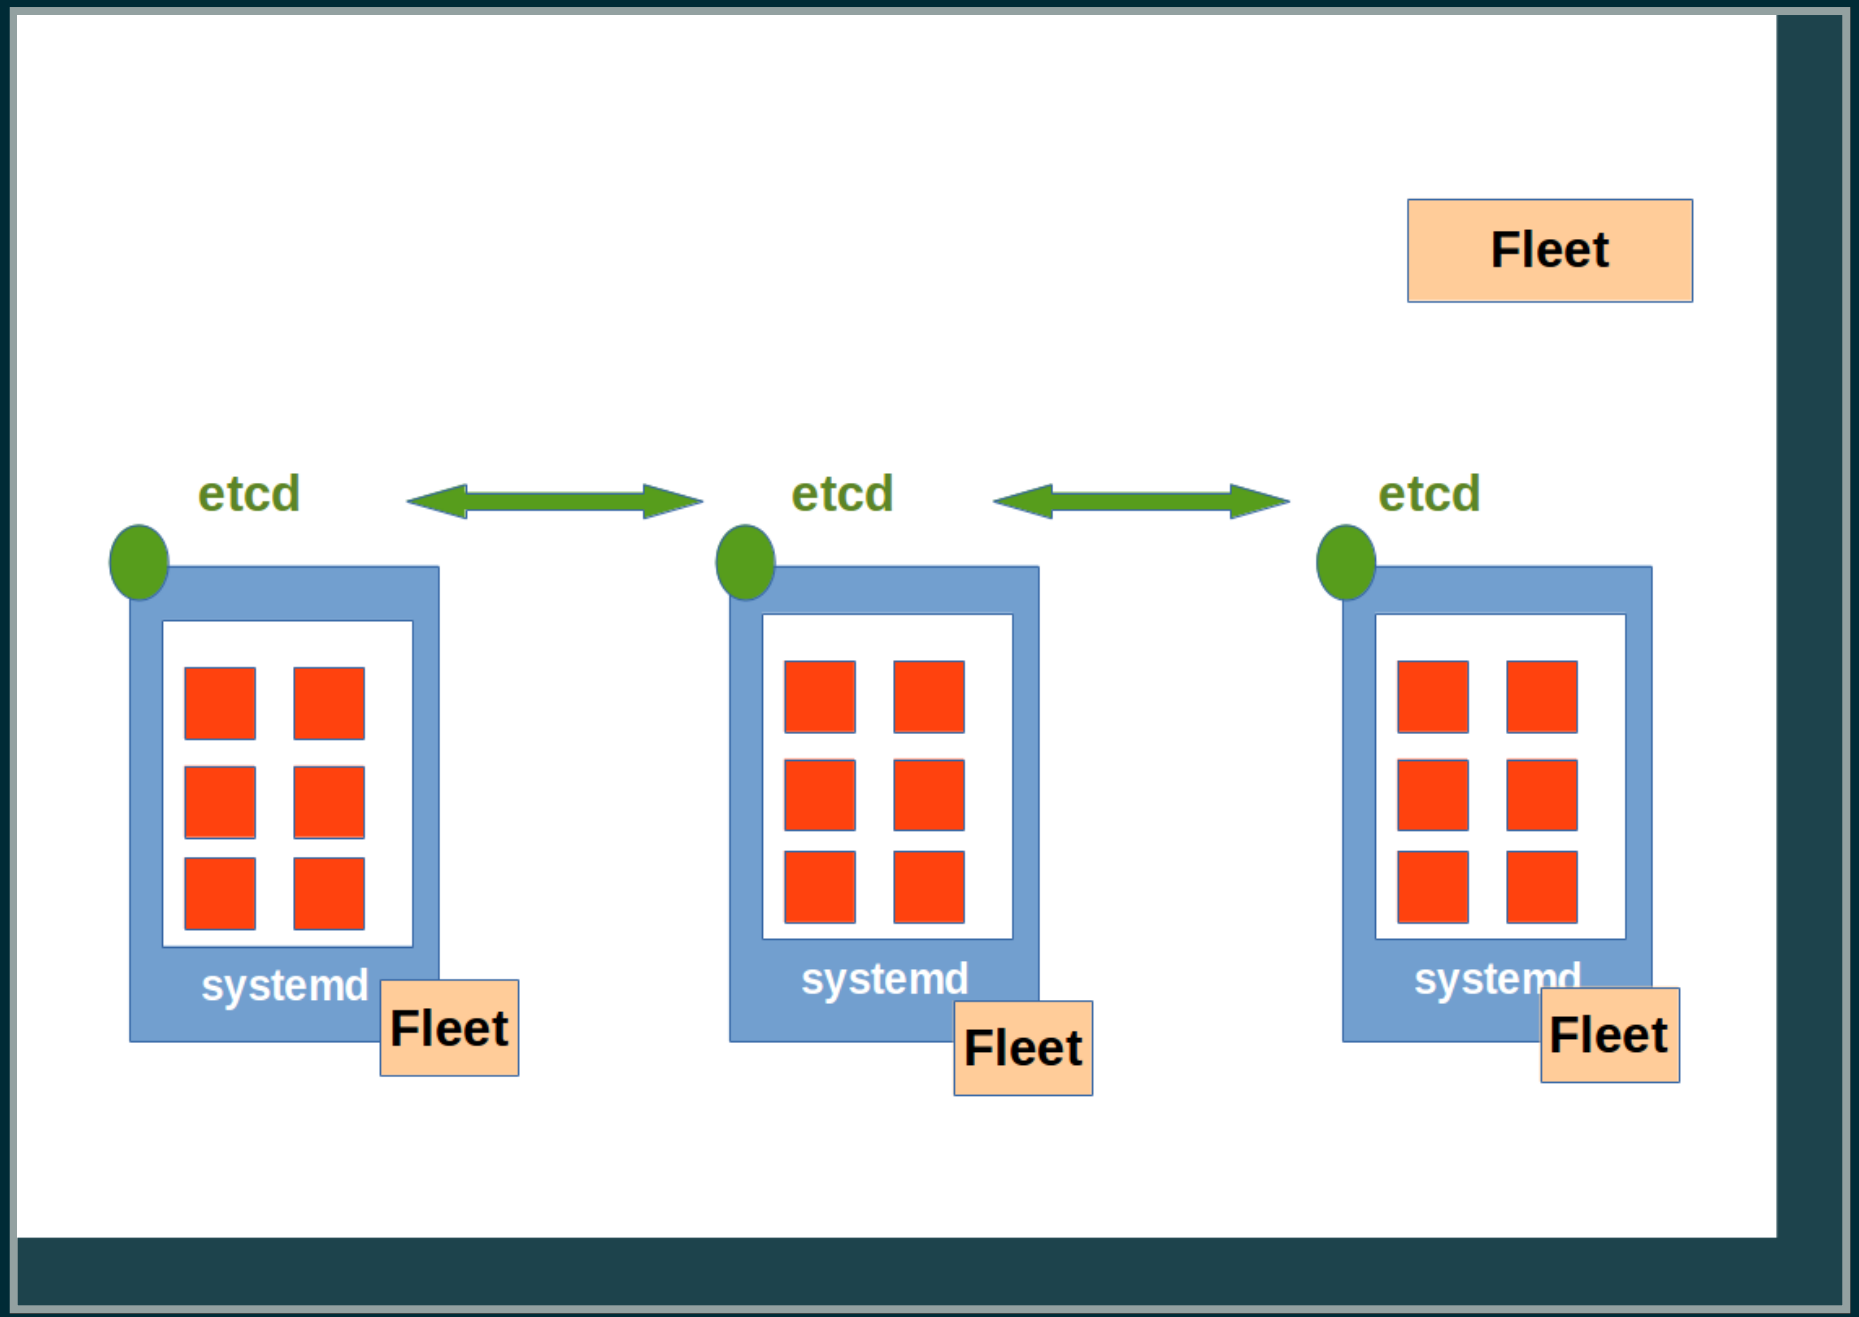
Fleet
etcd
etcd
etcd
systemd
Fleet
systemd
Fleet
systemd
Fleet


kubernetes
by Google

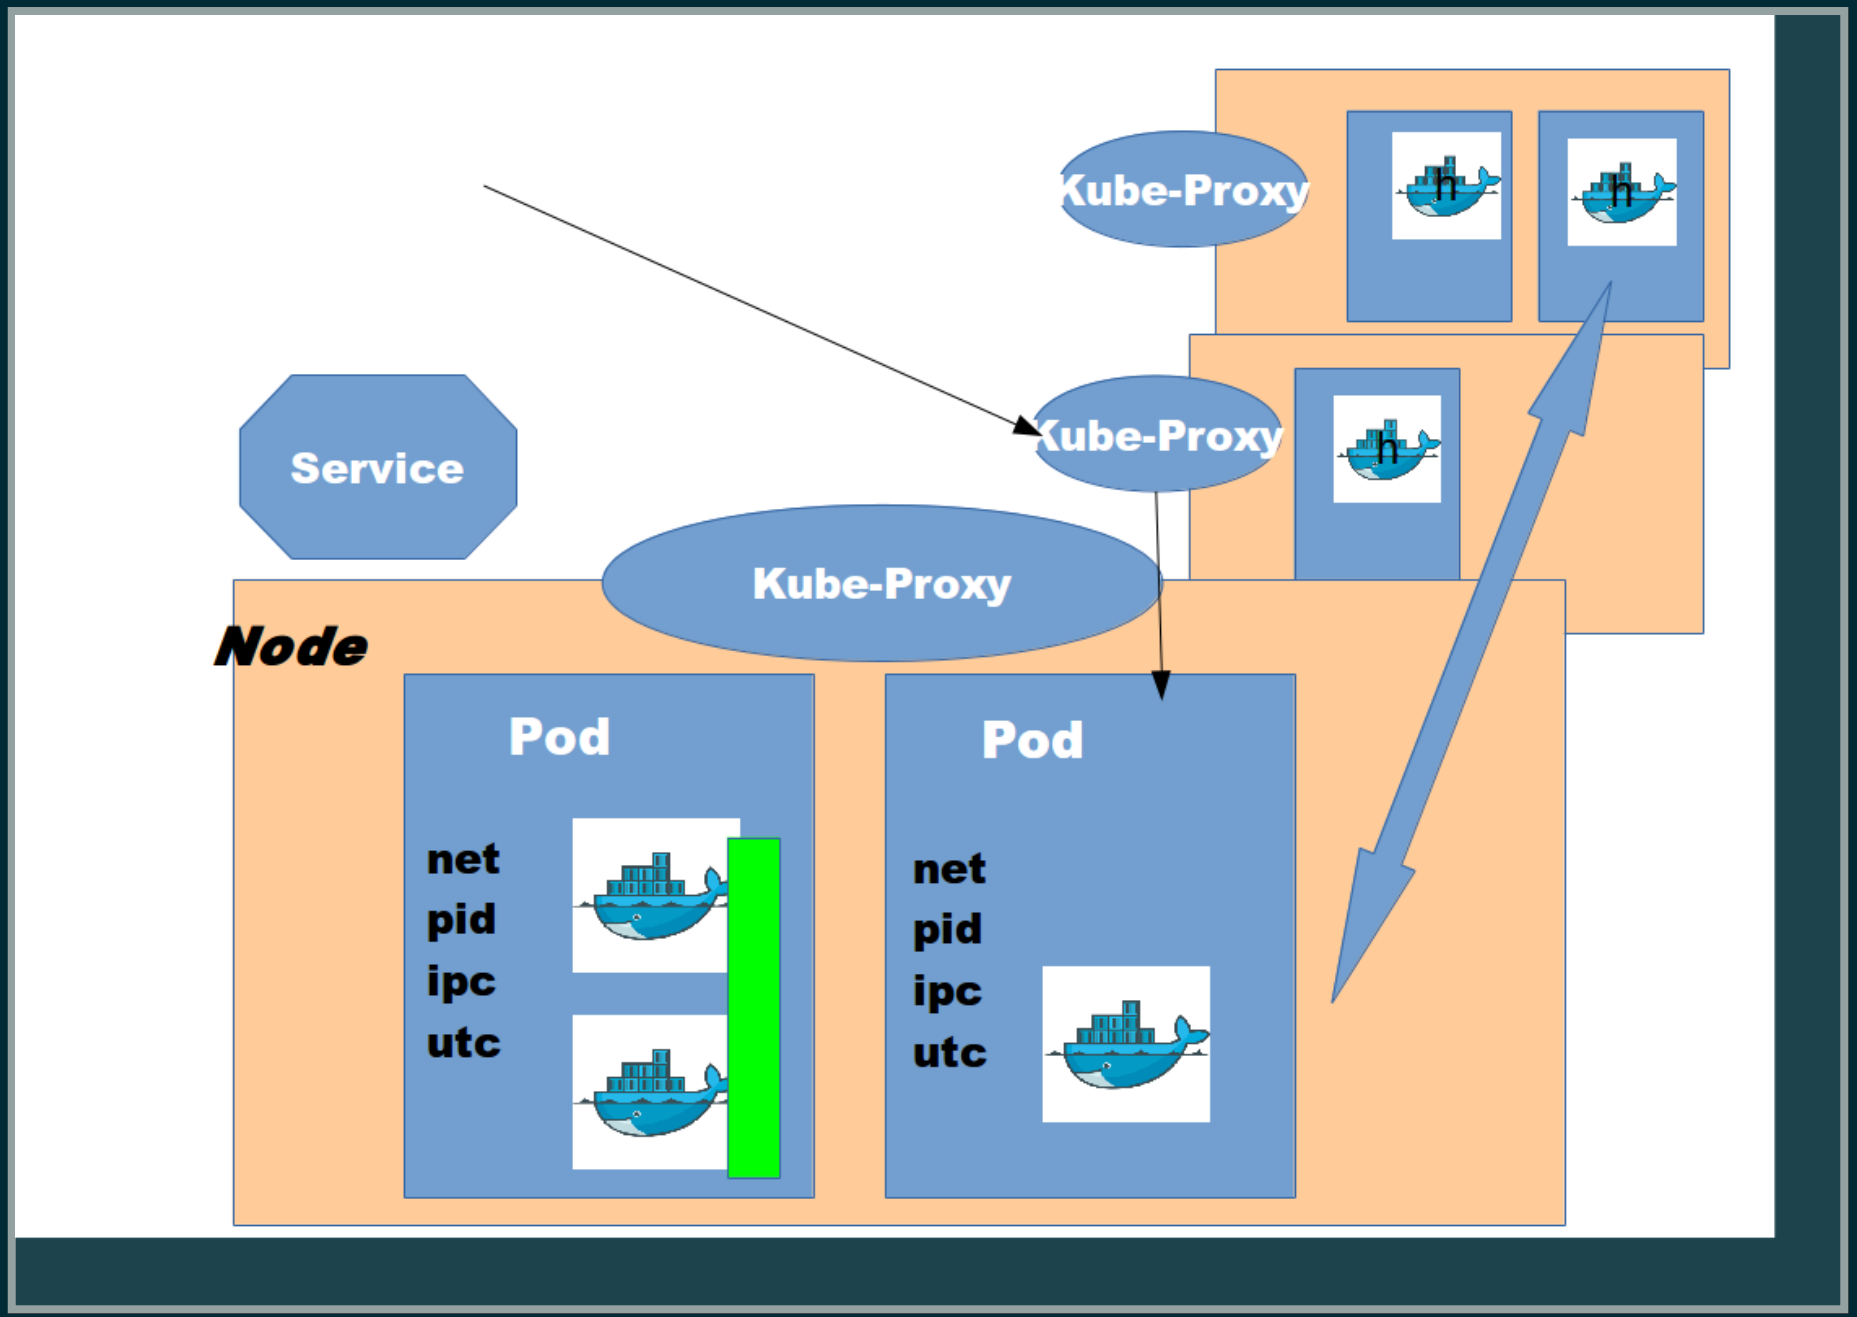
Kube-Proxy
Service
Kube-Proxy
Kube-Proxy
Node
Pod
net
pid
ipc
utc
Pod
net
pid
ipc
utc

ENDE

- Erkan Yanar
- http://linsenraum.de
- Auch xing
- linkedin auch