


Demystifying Docker Networking

Practical guide to black magic

Lorenzo Fontana

March 30, 2017



Lorenzo Fontana

DevOps Expert @Kiratech

Docker Maintainer

<http://fntlnz.wtf>

<https://github.com/fntlnz>

<https://twitter.com/fntlnz>

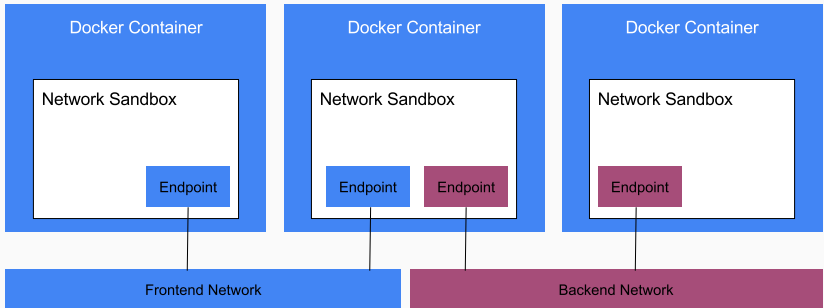
<http://www.kiratech.it>

Container Network Model


CNM: Container Network Model

- Sandbox
- Endpoint
- Network

CNM: Container Network Model (cont'd)




Libnetwork is the Native implementation of CNM



github.com/docker/libnetwork


Network Drivers 101

Null




```
# docker run -it --network=none alpine sh
```


Default Bridge (docker0)




```
# docker run -it alpine sh
```

Custom Bridge (mybridge)



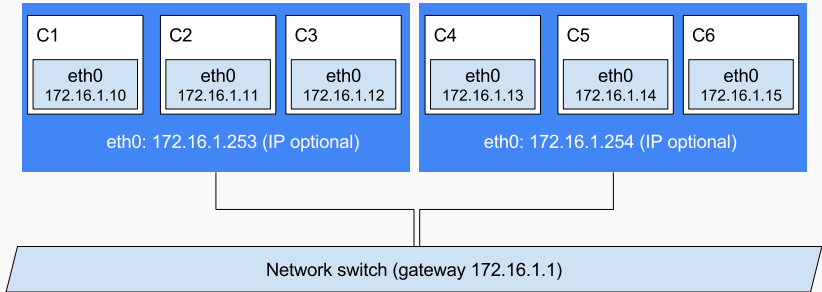
```
# docker network create -d bridge \  
-o com.docker.network.bridge.name=mybridge \  
mybridge  
  
# docker run -it --net mybridge alpine sh
```

Custom Bridge Network IPAM underlay (myunderbr)




```
# docker network create -d bridge --subnet
    =192.168.10.0/24 --gateway=192.168.10.254 \
    --aux-address DefaultGatewayIPv4=192.168.10.1 -o
    com.docker.network.bridge.name=myunderbr
    myunderbr
# brctl addif myunderbr enp5s0
# docker run -it --net myunderbr alpine sh
# docker run -it --net myunderbr --ip 192.168.10.90
    nginx:1.9
# ip a del 192.168.10.254/24 dev myunderbr
```

Macvlan & IPvlan




Macvlan modes: Bridged

Bridged (default): switches packets inside the host



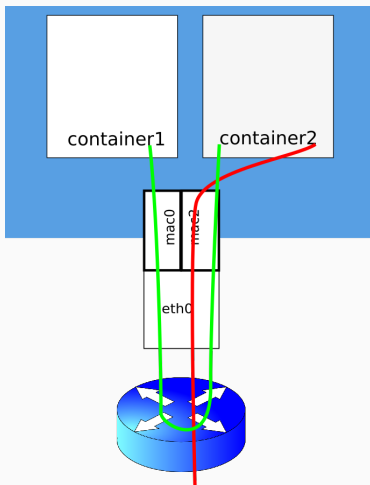
Macvlan modes: Private

Private blocks traffic between two MACVLAN interfaces on the same host




Macvlan modes: VEPA (Virtual Ethernet Port Aggregator)

VEPA requires a downstream switch that supports VEPA 802.1bg that will hairpin traffic back to the host if the destination is on the same host

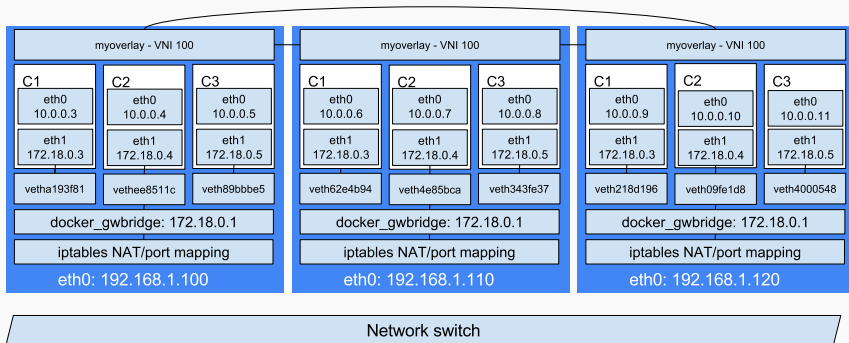


Macvlan modes: Passthru

Passthru is similar to private but relies on an external switch not to hairpin the traffic back to the originating host

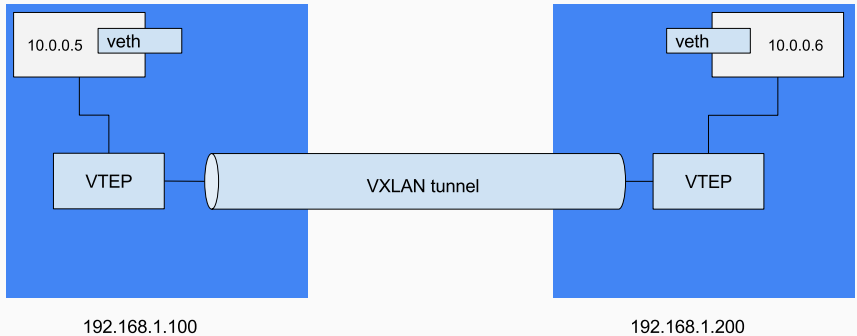


Overlay network (myoverlay)

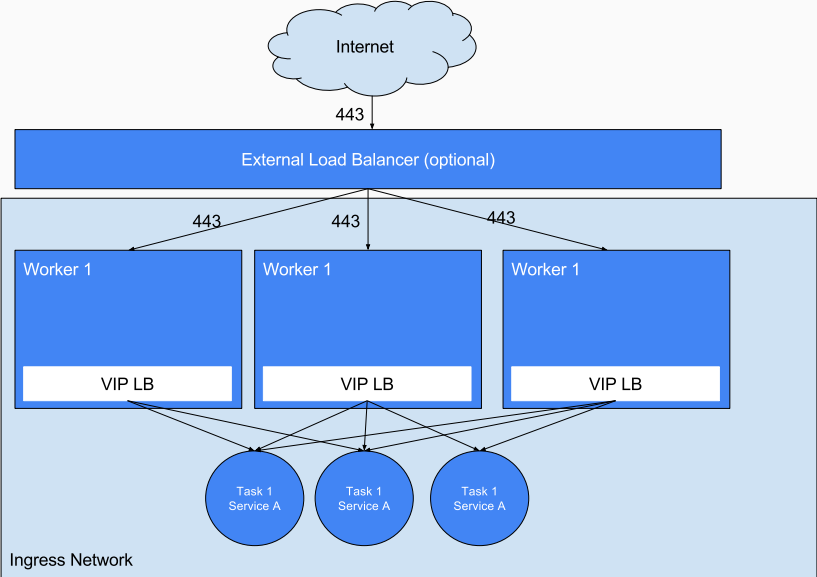


```
# docker network create -d overlay myoverlay
# docker service create --network myoverlay nginx
```

Overlay network (cont'd)



IPVS



Conclusion: Happy promotion!

Kiratech, the company I work for, is **hiring!**
drop me a line at lo@linux.com

Questions?

Thanks for listening!
And thanks to all the organizers!