

VO-Box in a container

Maxim Storetvedt

Western Norway University of Applied Sciences

July 20 2017

Background

- ▶ Recurring issues with several VO-Box VMs at CERN
 - ▶ Persistent network freezes
 - ▶ Debugging attempts have failed so far

Background

- ▶ Recurring issues with several VO-Box VMs at CERN
 - ▶ Persistent network freezes
 - ▶ Debugging attempts have failed so far
- ▶ Possible workaround: **containers**
 - ▶ May provide better reliability
 - ▶ Added performance
 - ▶ Simplified deployment

Singularity(?)

- ▶ New tool for launching containers
 - ▶ Rapidly gaining popularity within WLCG
- ▶ Simple and lightweight
- ▶ Great for isolation, but...
- ▶ Too simple?
 - ▶ Mostly a namespace wrapper
 - ▶ No individual container networking

In essence, what we want is something more of a "VM replacement"

Docker

- ▶ Well established container platform
- ▶ Mature codebase
- ▶ Well documented
- ▶ Can provide a more "VM oriented" approach to containers
- ▶ Newer Docker versions come with a built-in networking stack
 - ▶ Can provide containers with individual IP/MAC addresses
 - ▶ Provided through MACVLAN (available on Docker 1.12+)
 - ▶ Indistinguishable from conventional computers on a network
 - ▶ Straightforward to set up(!)

Creating a networked container

Create a new container running CentOS:

```
$ docker run -it -v /cvmfs:/cvmfs -h myalienvo.cern.ch  
--name=myvocontainer centos /bin/bash
```

Create a new MACVLAN bridge named "docknet":

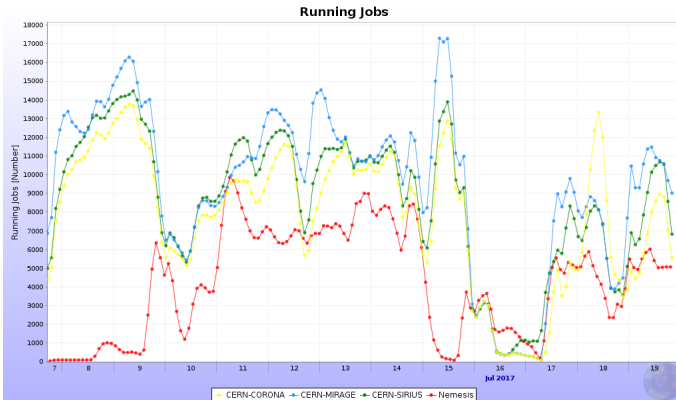
```
$ docker network create -d macvlan --subnet=  
137.138.47.192/26 --gateway=137.138.47.193 --ipv6  
--subnet 2001:1458:201:b50e::/64 --gateway  
2001:1458:201:b50e::1 -o parent=eth0 docknet
```

Connect to container to bridge:

```
$ docker network connect docknet myvocontainer  
--ip=137.138.47.210 --ip6=2001:1458:201:b50e::100:c
```

Current status

- ▶ One VO-Box test-site is currently active
 - ▶ "Nemesis" - Set up by **Maarten Litmaath** (thanks!)
- ▶ Shows promising results in terms of stability:



Issues and limitations

- ▶ Requires Linux 3.10+ and Docker 1.12+
 - ▶ Final version of Docker for CentOS 6 was 1.7
 - ▶ Thus, CentOS 7 is required
- ▶ Host Connectivity
 - ▶ The host and its containers can not reach/ping each other
 - ▶ Specific issue to macvlan bridging
 - ▶ Workarounds are available
- ▶ Kernel Access
 - ▶ Containers have limited access privileges by default
 - ▶ Several tools and services may fail to launch
 - ▶ Most networking tools are affected
 - ▶ Torque
 - ▶ More access can be given
 - ▶ But give too much, and you risk breaking isolation

Outlook

- ▶ Stability over time...?
- ▶ Deployment and configuration
 - ▶ The current site was manually configured from a "clean" CentOS 6 image.
 - ▶ Need preconfigured / prepackaged images for simplicity
 - ▶ Included packages
 - ▶ Required configuration steps
 - ▶ Distribution
 - ▶ Maintenance
 - ▶ Default access privileges

Questions?