

Kubernetes 101

Oleg Fiksel | Sean Mitchell | Michael Siebertz

Email: oleg@fiksel.info | sean@srm.im | mail@michaelsiebertz.de

Matrix: [@oleg:fiksel.info](https://matrix.org/join/!oleg:fiksel.info) | [@sean:fiksel.info](https://matrix.org/join/!sean:fiksel.info) | [@captain.usan:matrix.org](https://matrix.org/join/!captain.usan:matrix.org)

2018-11-04 OpenRheinRuhr 2018

MOTIVATION

Our motivation to do this talk

MOTIVATION

Our motivation to do this talk

Provide an overview of container architecture on the example of Docker and Kubernetes.

AGENDA

AGENDA

▶ Docker Fundamentals

AGENDA

- ▶ Docker Fundamentals
- ▶ **Kubernetes from 10.000 feet**

AGENDA

- ▶ Docker Fundamentals
- ▶ Kubernetes from 10.000 feet
- ▶ **Live Demo**

OLEG FIKSEL (DOCKER FUNDAMENTALS)

¹Accenture Interactive

²DevSecOps Workshop

OLEG FIKSEL (DOCKER FUNDAMENTALS)

About me:

- ▶ DevOps Engineer @ Accenture Interactive ¹
(DevOps Chapter)
- ▶ **Main topics**
 - ▶ Big fan of CI (especially GitLab) and a Gentoo user
 - ▶ Learning Golang and like to automate everything
 - ▶ Continuous Security (DevSecOps) ²

¹Accenture Interactive

²DevSecOps Workshop

SEAN MITCHELL (KUBERNETES FROM 10.000 FEET)

SEAN MITCHELL (KUBERNETES FROM 10.000 FEET)

About me:

- ▶ Site Reliability Engineer (SRE) @ Grandcentrix ¹
- ▶ **Main topics**
 - ▶ Support developers building and deploying projects
 - ▶ Ensure customers' cloud infrastructure runs
 - ▶ Home Automation geek

¹Grandcentrix

MICHAEL SIEBERTZ (LIVE DEMO)

MICHAEL SIEBERTZ (LIVE DEMO)

About me:

- ▶ DevOps Engineer @ HRS ¹
- ▶ **Main topics**
 - ▶ Linux-User for long time
 - ▶ Trying to automate as much as possible
 - ▶ Part of Central DevOps Team

¹HRS

Container fundamentals

A BIT OF HISTORY

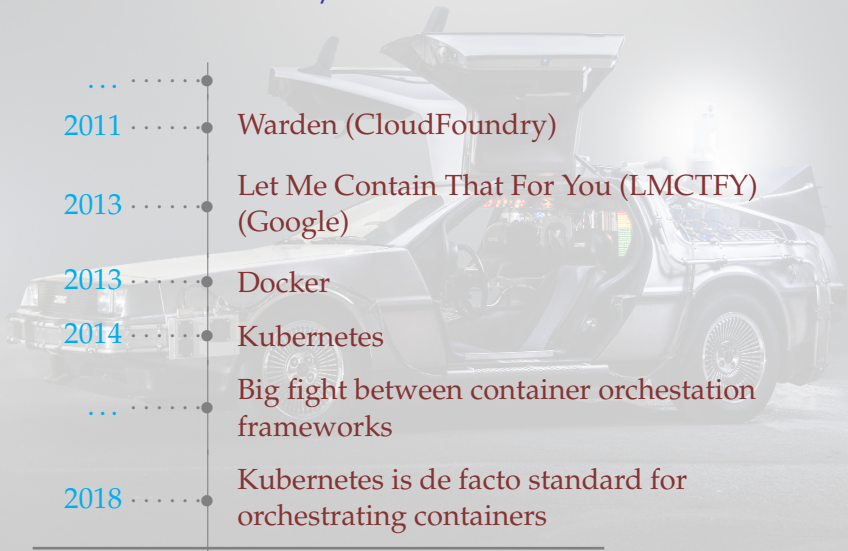


A BIT OF HISTORY 1/2



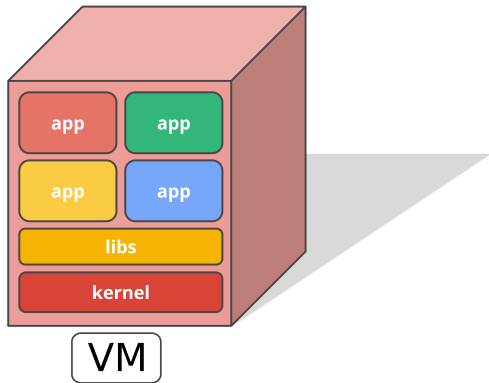
- 1979 ● chroot
- 2000 ● FreeBSD Jails
- 2001 ● Linux VServer
- 2004 ● Solaris Containers
- 2005 ● OpenVZ
- 2006 ● Process Containers (Google)
- 2008 ● LXC
- ●

A BIT OF HISTORY 2/2




VMs AND CONTAINERS

VMs AND CONTAINERS



VMs AND CONTAINERS



VMs AND CONTAINERS

VMs AND CONTAINERS

▶ Storage

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.
 - ▶ **Container: No emulation, just separation**

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.
 - ▶ Container: No emulation, just separation
- ▶ **Networking**

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.
 - ▶ Container: No emulation, just separation
- ▶ Networking
 - ▶ VM: NIC emulation

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.
 - ▶ Container: No emulation, just separation
- ▶ Networking
 - ▶ VM: NIC emulation
 - ▶ **Container: Uses network stack of the host**

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.
 - ▶ Container: No emulation, just separation
- ▶ Networking
 - ▶ VM: NIC emulation
 - ▶ Container: Uses network stack of the host
- ▶ **Runtime**

VMs AND CONTAINERS

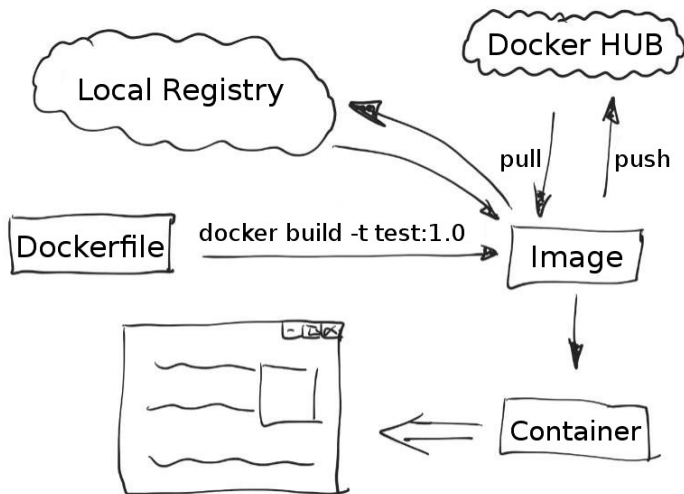
- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.
 - ▶ Container: No emulation, just separation
- ▶ Networking
 - ▶ VM: NIC emulation
 - ▶ Container: Uses network stack of the host
- ▶ Runtime
 - ▶ VM: stateful

VMs AND CONTAINERS

- ▶ Storage
 - ▶ VM: Binary Image
 - ▶ Container: Layered plain FS
- ▶ Emulation
 - ▶ VM: BIOS, Hardware, etc.
 - ▶ Container: No emulation, just separation
- ▶ Networking
 - ▶ VM: NIC emulation
 - ▶ Container: Uses network stack of the host
- ▶ Runtime
 - ▶ VM: stateful
 - ▶ **Container: stateless**


CONTAINER LIFECYCLE - OVERVIEW

CONTAINER LIFECYCLE - OVERVIEW



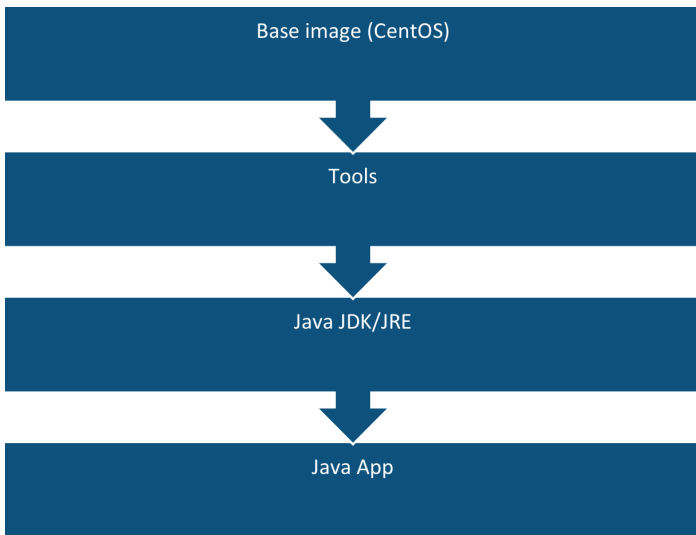
CONTAINER LIFECYCLE - IMAGE VERSIONS

CONTAINER LIFECYCLE - IMAGE VERSIONS




LAYERED FILE SYSTEM

LAYERED FILE SYSTEM



LAYERED FILE SYSTEM




LAYERED FILE SYSTEM - IMAGE DEPENDENCIES


LAYERED FILE SYSTEM - IMAGE DEPENDENCIES



LAYERED FILE SYSTEM - IMAGE DEPENDENCIES



LAYERED FILE SYSTEM - IMAGE DEPENDENCIES



LAYERED FILE SYSTEM - BACKENDS

LAYERED FILE SYSTEM - BACKENDS

- ▶ aufs (legacy)
- ▶ overlayfs (legacy)
- ▶ devicemapper
- ▶ overlayfs2
- ▶ btrfs

PERSISTENT STORAGE (VOLUMES)

PERSISTENT STORAGE (VOLUMES)

- ▶ Containers start without state

PERSISTENT STORAGE (VOLUMES)

- ▶ Containers start without state
- ▶ State is provided via mounts

PERSISTENT STORAGE (VOLUMES)

- ▶ Containers start without state
- ▶ State is provided via mounts
 - ▶ **Binds (local FS)**

PERSISTENT STORAGE (VOLUMES)

- ▶ Containers start without state
- ▶ State is provided via mounts
 - ▶ Binds (local FS)
 - ▶ **Docker volumes**

PERSISTENT STORAGE (VOLUMES)

- ▶ Containers start without state
- ▶ State is provided via mounts
 - ▶ Binds (local FS)
 - ▶ Docker volumes
 - ▶ **Docker Storage plugins**

SEPARATION AND RESOURCE LIMITING

SEPARATION AND RESOURCE LIMITING

▶ cgroups

SEPARATION AND RESOURCE LIMITING

- ▶ cgroups
 - ▶ Linux kernel feature

SEPARATION AND RESOURCE LIMITING

- ▶ cgroups
 - ▶ Linux kernel feature
 - ▶ Limiting: CPU, memory, disk I/O, network

SEPARATION AND RESOURCE LIMITING

- ▶ cgroups
 - ▶ Linux kernel feature
 - ▶ Limiting: CPU, memory, disk I/O, network
- ▶ namespaces

SEPARATION AND RESOURCE LIMITING

- ▶ cgroups
 - ▶ Linux kernel feature
 - ▶ Limiting: CPU, memory, disk I/O, network
- ▶ namespaces
 - ▶ Linux kernel feature

SEPARATION AND RESOURCE LIMITING

- ▶ cgroups
 - ▶ Linux kernel feature
 - ▶ Limiting: CPU, memory, disk I/O, network
- ▶ namespaces
 - ▶ Linux kernel feature
 - ▶ process isolation
Each container "thinks" it's the only one in the system
(*ps uax*)

SUMMARY

SUMMARY

- ▶ The is no question "Use container or not?".
The question is "Where to start?"

SUMMARY

- ▶ The is no question "Use container or not?".
The question is "Where to start?"
- ▶ Docker is the standard de facto (for now) for running containers

SUMMARY

- ▶ The is no question "Use container or not?".
The question is "Where to start?"
- ▶ Docker is the standard de facto (for now) for running containers
- ▶ Docker alone is not enough to run containers big time
- ▶ **Next up: Container Orchestration (Kubernetes)**

Q & A

Thanks!

Oleg Fiksel | Sean Mitchell | Michael Siebertz

Email: oleg@fiksel.info | sean@srm.im | mail@michaelsiebertz.de

Matrix: @oleg:fiksel.info | @sean:fiksel.info | @captain.vsan:matrix.org