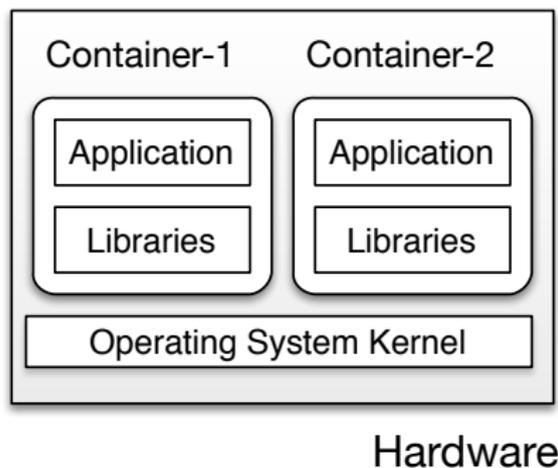
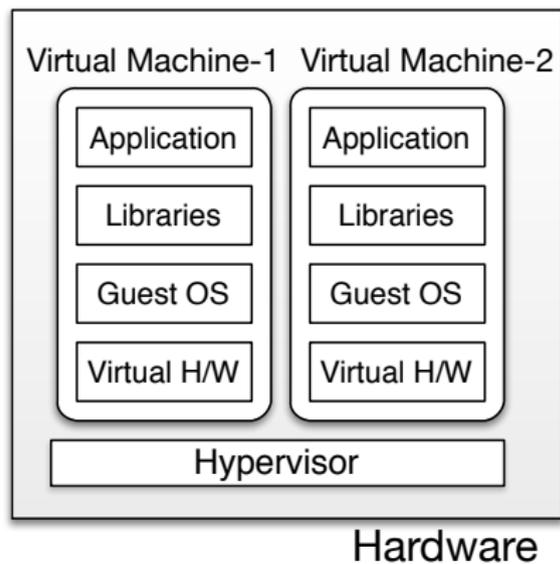


OS Virtualization and Containers

Hardware vs. OS Virtualization



OS Virtualization

- Virtualize the OS for each application
- *Not* the hardware
- Allow applications to run in isolation from each other
- If this seems similar to the process abstraction, that's because it is
- Resource sharing is a core OS primitive anyway.
- However, some services are shared:
 - 1 File system, libraries, config files
 - 2 Users, uids
 - 3 Pid tree (pid's are unique)
 - 4 Network interfaces, rules (iptables), ports
 - 5 Time
 - 6 Procfs, sysfs

Dual View of Containers

Containers can be viewed as either:

- 1 Lightweight virtual machines
- 2 Process groups with better isolation

Containers as OS virtualization analogue of VMs:

- Applications/processes run inside containers
 - Isolated from each other using sandboxing techniques
 - Applications should not affect or monitor apps in other containers
- Containers run on a “virtualized OS interface”
- Need no additional interposition
- No/negligible performance degradation

Essence of OS Virtualization

- Virtualize the shared OS services
- Ideally: applications not affected by others in any way
- Implemented through namespaces
- Resource isolation. Fine-grained control of CPU, memory, I/O resources consumed
- OS must try to provide isolation anyway
- OS virt is thus the natural progression:
 - Early UNIX: File system isolation with chroot
 - FreeBSD: Jails [2000] (chroot+containing the omnipotent root)
 - Solaris: Zones [2005] (refined, more isolated Jails)
 - Linux: [Virtuzzo-2000, LXC-2008, ...]

OS Virtualization Desiderata

- Security
- Isolation
- Virtualization
 - HW devices, network IP, hostname, ...
- Granularity
 - Containers can be arbitrarily sized
 - No dedicated CPUs required
- Transparency. No porting required. Exact same ABI/API as running on bare-metal.

OS Virtualization Challenges

- Lots of OS → App interfaces to isolate
- Security challenges abound
- Needs careful understanding of the OS
- Replace process-id with (process-id, container-id) throughout the kernel
- Update virtualization and isolation with each new OS feature
- In contrast, HW, and HW virt, are relatively stable
- What should be the limitations on the “root” user in a container?

Namespaces

- Pid
- Network : nic, iptables, routing tables
- UTS: hostname
- Mount : Mount different file-system trees
- User: User-ids
- IPC: POSIX shmем, etc

Control groups

- CPU, memory, blkio
- Control via sysfs
- Resource limiting, prioritization, accounting
- Control (freeze, checkpoint processes)

Linux Containers (LXC)

- Create “containers” using namespaces, cgroups, seccomp (security policies)
- Similar control abstraction as VMs:
- Containers have names, “FS images”, resource allocation (CPU, mem,..)
- Operations: create, start, shutdown, pause, migrate

Docker

- LXC + layered file system + image repository
- Copy-on-write file system allows images to be composed layer-wise:
 - 1 Base layer: Debian
 - 2 Layer 2: Essentials (Emacs)
 - 3 Layer 3: Apache web server
- Common use-case: CI/CD
- Continuous Integration and Deployment
- Create docker container in dev environment, “push” into production

OS Virtualization Benefits

- “Lightweight” virtualization. No performance overhead
- Negligible resource overhead (mem, cpu)
- VMs need resources to run guest OS
- Near-instant startup
 - Guest OS boot can take tens of seconds
 - Hypervisor optimizations can reduce it to ~ 100 ms [ClearContainers]
- Dynamic resource management
- Change cpu, mem, IO resources at run-time
- Less wasted resources
- Free resources inside a VM are considered “allocated”

OS Virtualization Drawbacks

- Linux containers not secure due to large surface area
- Resource isolation may be weak
- Leads to performance interference
- Two CPU intensive containers
- One container runs fork-bomb
- Isolation is hard
- Many common, shared resources
- /procfs . Useful for system monitoring
- Exposed to container or not?

Other Virtualization Options

Library virtualization

- Implement glibc on Windows (Cygwin, mingw)
- Implement windows APIs on Linux (Wine)

System-call virtualization

- Run applications from different OS
- Intercept syscalls made by application and reimplement them
- Linux apps on FreeBSD, SmartOS (Solaris)
- Linux apps on Windows

Docker Commands

- `docker inspect image-name` to see image layer storage
- `docker save` to tar file
- `commit` to create a new image

Dockerfiles

```
FROM ubuntu:18.04  
COPY . /app  
RUN make /app  
CMD python /app/app.py
```

```
docker build
```