Container

jnlin

Outline

- Virtualization: machine level
- Virtualization: OS level
- FreeBSD jail
- Docker

Virtualization – machine level (1)

- Hardware virtualization
  - Emulate CPU, RAM, HDD, Network Interface …
- Host OS and Guest OS
  - Isolated between each guest OS
- Hypervisor
  - QEMU, VirtualBox, VMWare…
- EC2, GCE (Google Computed Engine)

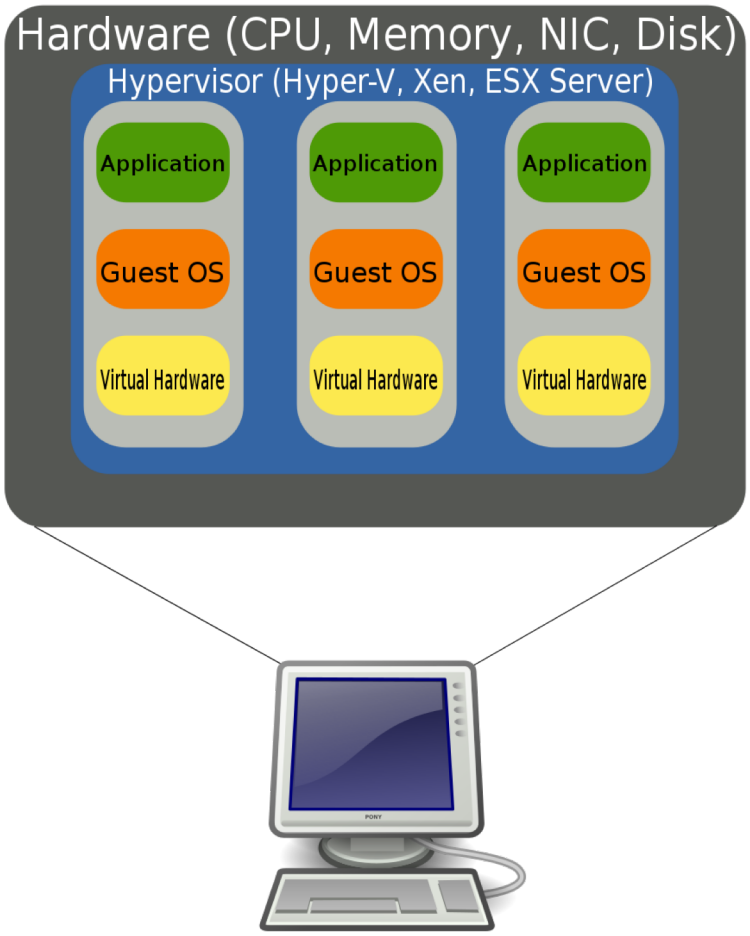

Virtualization – machine level (2)

- Intel VT-x and AMD-V
  - An extension of CPU instructions to improve performance of virtualization
  - Include CPU and I/O virtualization
- Pros
  - Better use of IT resources
  - Elasticity
- Cons
  - Performance

Virtualization – OS level (1)

- Multiple isolated user space instances
  - Share the same kernel
  - Must use the same operating system as the host one
- Like chroot but more powerful
- Containers can be live migrated (without restart service)

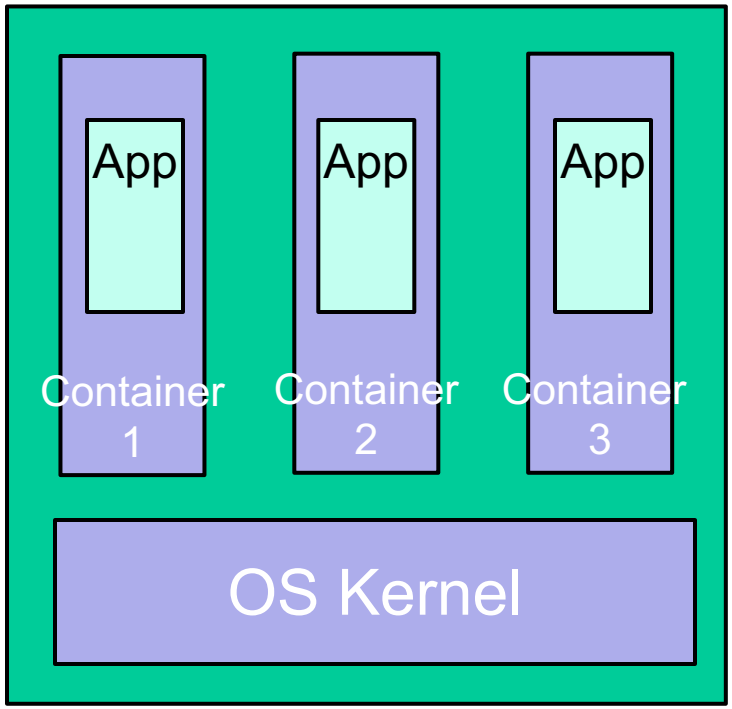

Virtualization – OS level (2)

- Technologies
  - Docker – works on Windows, Mac OS X, and Linux as host OS
  - Jails – works on FreeBSD
  - LXC, OpenVZ
- Pros
  - Better performance
  - Better security
  - Use the same environment in development and production
- Cons
  - Complex to understand (by software developers, if they did not take the SysAdm course ;-)

FreeBSD jail (1)

- jail(8)
- Preparation (from base image)
  - % mkdir -p /home/jails/firstjail
  - % export DESTDIR=/home/jails/firstjail
  - % export DESTRELEASE=12.0-RELEASE
  - % export DESTARCH=`uname -m`
  - % export SOURCEURL=http://ftp.freebsd.org/pub/FreeBSD/releases/$DESTARCH/$DESTRELEASE/
  - % fetch $SOURCEURL/$base.txz
  - % tar -xf base.txz -C $DESTDIR

FreeBSD jail (2)

- Start jail while booting
  - /etc/jail.conf
    - www {
    - host.hostname = www.example.org; # Hostname
    - ip4.addr = 192.168.0.10; # IP address of the jail
    - path ="/usr/jail/www"; # Path to the jail
    - devfs_ruleset = "www_ruleset"; # devfs ruleset
    - mount.devfs; # Mount devfs inside the jail
    - exec.start = "/bin/sh /etc/rc"; # Start command
    - exec.stop = "/bin/sh /etc/rc.shutdown"; # Stop command
    - }
  - jail_enable="YES" (in /etc/rc.conf)

FreeBSD jail (3)

- jls – list all jails
  - JID IP Address Hostname Path
  - 3 192.168.0.10 www /usr/jail/www
- jexec – execute commands in a jail
  - jexec 3 ps -auxww

Docker

- Most popular OS level virtualization technology in the 2010s
- Open sourced (Apache License 2.0), Developed by Docker Inc.
- Use different interfaces to access virtualization features of Linux kernel
- Infrastructure as Code

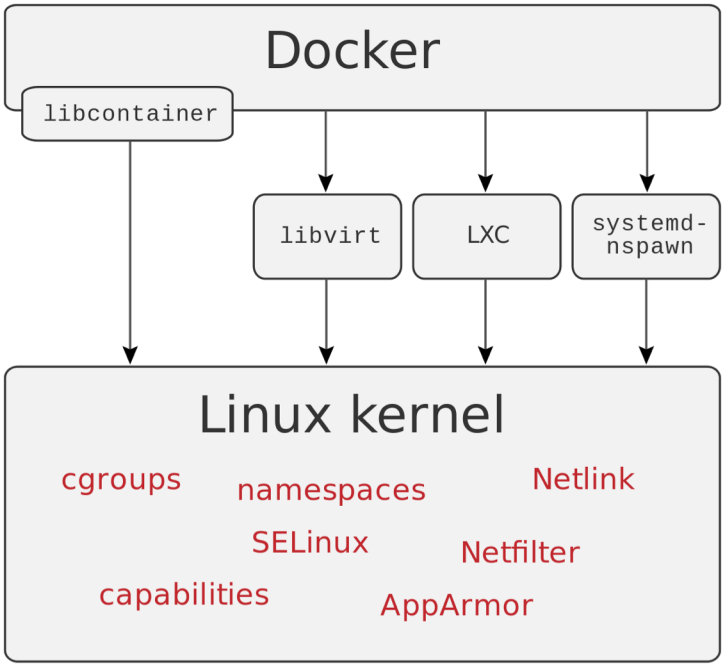

Docker - Dockerfile (1)

- Reuse pre-built images
- Automate the process of building environment
- Example
  - FROM alpine
  - RUN apk update && apk install curl
  - COPY myapp /app/myapp
  - CMD /app/myapp

Docker - Image Layer (1)

- A writeable layer on top of a bunch of read-only layers
- Each RUN has its own commit
  - FROM alpine
  - RUN apk add curl
  - RUN https://xxx.yyy.zzz/data.tgz
  - RUN rm data.tgz
- % docker history <image>
- Keep image as small as possible

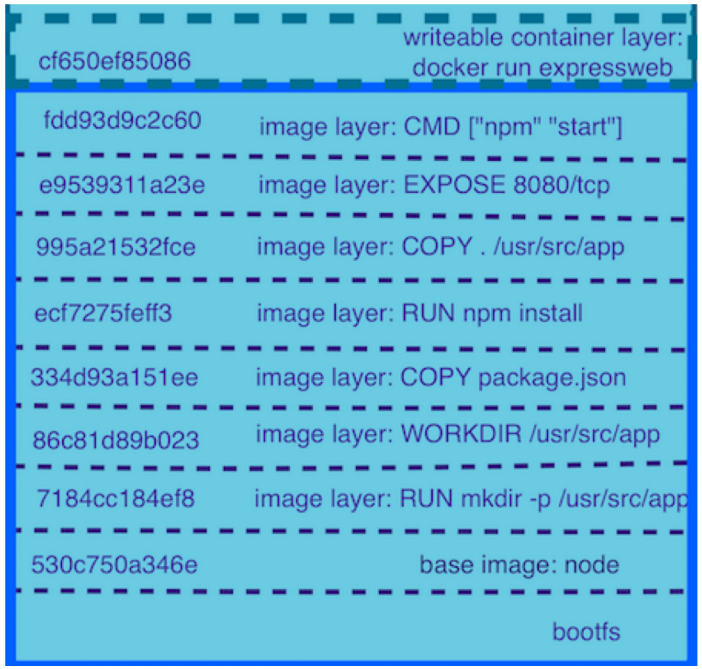

Docker - Image Layer (2)

- overlayfs
  - combining numerous directories into one directory that looks like it contains the content from all the them.

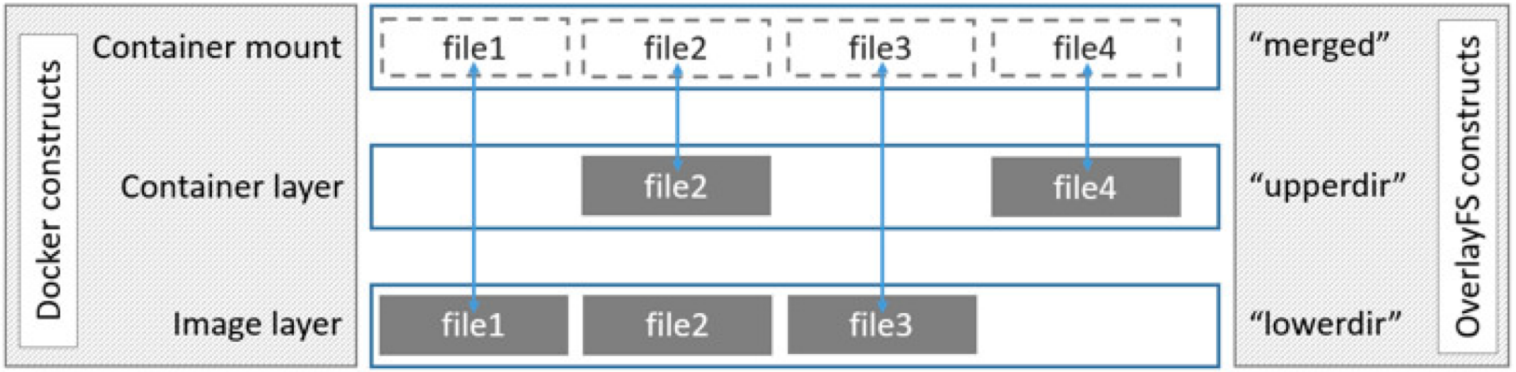

- You also can use ZFS to implement the same feature

Docker - Docker (Command Line)

- docker pull
  - Pull an image from public or private repository
- docker build
  - Build image from Dockerfile
- docker run
  - Start a docker instance
- docker kill
  - Stop a docker instance
- docker rm
  - Remove resource used by a docker instance
- docker ps
  - Show all docker instances (running or stopped)
- docker push
  - Push the image to docker repository

Docker

- Pros
  - IaC simplify the operating effort
    - Works on my ~~computer~~ server
- Cons
  - Security
  - Scalability