

docker: Linux containers for everyone

Sébastien Binet

LAL/IN2P3

2013-12-18

- What ?
- Why ?
- How ?
- (When ?)

Docker: what is it ?

- <http://www.docker.io/>
- an open source project to pack, ship and run any application as a lightweight container

High level description

- kind of like a **lightweight** VM
- runs in its own process space
- has its own network interface
- can run stuff as `root`

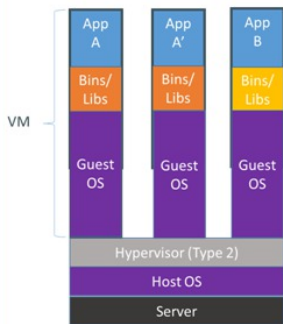
Low level description

- `chroot` on steroids
- container = isolated process(es)
- share kernel with host
- no device emulation


Docker: why ?

- same use cases than for VMs
- **speed**: boots in *ms*
- **footprint**: 100-1000 containers on a single machine/laptop. small disk requirements

Containers vs. VMs



Containers are isolated, but share OS and, where appropriate, bins/libraries



Efficiency: *almost* no overhead

- processes are isolated but run straight on the host
- CPU performance = **native** performance
- memory performance = a few % shaved off for (optional) accounting
- network performance = small overhead

Efficiency: storage friendly

- unioning filesystems
- snapshotting filesystems
- copy-on-write

- provisioning takes a few milliseconds
- ... and a few kilobytes
- creating a new container/base-image takes a few seconds

Docker: how ?

- Linux Containers (LXC)
- Control Groups and Namespaces
- AUFS (overlying filesystem)
- Client-Server with an `http/REST` API

LXC

- Let's you run a Linux system within another Linux system
- A container is a group of processes on a Linux box, put together is an isolated environment
- From the inside, it looks like a VM. From the outside, it looks like normal processes
- *"chroot on steroids"*

Control Group & Namespaces

Linux kernel feature to limit, account and isolate resource usage:

- CPU
- Memory
- disk I/O

AUFS

- File system that implements union mount

```
$ mount -t aufs -o br=/files1:/files2 none /files
```

- Supports copy-on-write (COW)


```
$ mount -t aufs -o br=/tmp=rw:/bin=ro none /files
```

Go

Docker is actually a Go binary scripting LXC and AUFS



docker



Requirements (`docker < 0.7`)

- Linux Kernel 3.8 or above
- AUFS
- LXC
- 64-bit

Requirements (`docker ≥ 0.7`)

- Linux Kernel 3.8 or above
- Devicemapper
- LXC
- 64-bit

More on: <http://docs.docker.io/en/latest/installation>

Hello World

- get a base container (ubuntu, centos, ...)

```
$ docker pull ubuntu
```

- list images already pulled in:

```
$ docker images
```

ubuntu	12.04	8dbd9e392a96	5 months ago	131.5 MB (virtual 131.5 MB)
ubuntu	latest	8dbd9e392a96	5 months ago	131.5 MB (virtual 131.5 MB)
ubuntu	precise	8dbd9e392a96	5 months ago	131.5 MB (virtual 131.5 MB)
ubuntu	12.10	b750fe79269d	6 months ago	24.65 kB (virtual 180.1 MB)
ubuntu	quantal	b750fe79269d	6 months ago	24.65 kB (virtual 180.1 MB)

- run an executable inside a container

```
$ docker run ubuntu:12.10 echo ``hello world``
```

<http://www.docker.io/gettingstarted/>

Detached mode

- run a container in detached mode:

```
$ docker run -d ubuntu sh -c \  
  while true; do echo "hello"; sleep 1; done;
```

- get the container id:

```
$ docker ps
```

ID	IMAGE	COMMAND	CREATED	STATUS	PORTS
78c88e279f26	ubuntu:12.04	/bin/sh -c while tru	14 seconds ago	Up 11 seconds	

- attach to the container

```
$ docker attach 78c88e279f26
```

- start/stop/restart a container

```
$ docker stop 78c88e279f26
```

Public index

- pull an apache container from the index:

```
$ docker search apache  
$ docker pull creack/apache2
```

- run the image and check the ports

```
$ docker run -d creack/apache2  
$ docker ps
```

ID	IMAGE	COMMAND	CREATED	STATUS	PORTS
369602483ae9	creack/apache2:latest	/usr/sbin/apache2ctl	4 seconds ago	Up 1 seconds	49153->80, 49154->443

Also available from the browser:

<https://index.docker.io/>

- **run docker interactively:**

```
$ docker run -i -t ubuntu bash
root@bf72b1a06e6c:/# apt-get update
Reading package lists... Done
```

```
root@bf72b1a06e6c:/# apt-get install memcached
[...]
root@bf72b1a06e6c:/# exit
```

- **commit the resulting container**

```
$ docker commit `docker ps -q -l` sbinet/memcached
ab59e4b14266
```

- **run the image**

```
$ docker run -d -p 11211 -u daemon sbinet/memcached memcached
ab59e4b14266
```

Docker: creating a customized container - Dockerfile

```
## build mana
from sbinet/slc:6.5
maintainer binet@cern.ch

## install compiler and libs
run yum -y update && yum -y install file which gcc gcc-c++

## install hwaf
RUN (cd /usr && \
    curl -L http://cern.ch/hwaf/downloads/tar/hwaf-latest.tar.gz \
    | tar -zxf -)

## fetch mana sources
RUN mkdir -p /build/mana
RUN (cd /build/mana && \
    curl -L http://cern.ch/mana-fwk/downloads/src/mana-latest-next.t
    | tar --strip-components 1 -zxf -)

## build
RUN (cd /build/mana && ./scripts/build-release)
```

- build the container

```
$ docker build -t=sbinet/mana-x86\_64-slc6-gcc44-opt \  
path/to/dir/holding-Dockerfile
```

- run the container (and test the build)

```
$ docker run -d sbinet/mana-x86\_64-slc6-gcc44-opt \  
mana-20131023
```

- bind mounts

```
$ docker run -d sbinet/mana-x86\_64-slc6-gcc44-opt \  
-v /host/build/results:/build/mana \  
mana-20131023
```

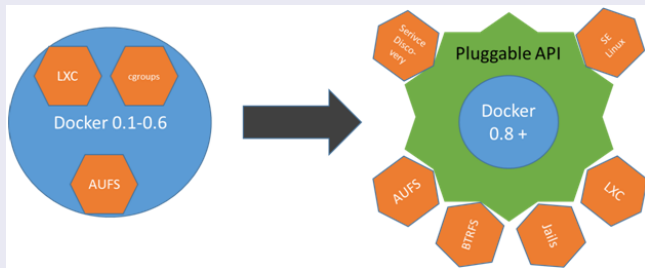
- copy files from container to host

```
$ docker cp mana-20131023:/build/mana /host/build/results
```

<https://github.com/atlas-org/docker-containers>

Docker: when ? where ?

- OpenStack-Havana has support for running docker containers
 - so beginning 2014: be able to run those on `openstack.cern.ch`
-
- `docker-0.7` removes AUFS requirements
 - packaged for Fedora-20/19
 - available on RHEL/SLC/CentOS-6.5
 - ▶ `yum install docker-io`



- easily distribute dev-environments
- easily provision build and dev-environments
- provision efficient performance-wise environments (production)

