

Experience with Docker and ICAT

Rolf Krahl

ICAT Meeting @ 11th NOBUGS, Copenhagen, October 2016

- Docker is a virtualization environment. Applications are capsuled in *containers*.
- Container are similar to virtual machines, but much less overhead in terms of resources.
- Basic principle: each service in its own container (e.g. Glassfish/ICAT, Database, Apache frontend).
- Container are instances of *images*. An image is a portable file system containing operation system, libraries, and applications.
- Docker maintains a private network for all container at the current host. Forwarding public ports of the host to container is possible.

- Creation of docker images is controlled by a Dockerfile, often using another image as starting point.
- Use file system layers and copy-on-write.
- Directories or individual files from the host can be made accessible to containers by a bind mount (*volume*).

Example: Dockerfile for MariaDB

```
FROM rkrahl/opensuse

RUN zypper --non-interactive install mariadb
RUN mkdir -p /etc/mysql /etc/mysql.d
COPY start-mysql.sh /etc/mysql
RUN chmod 0755 /etc/mysql/start-mysql.sh

CMD ["/etc/mysql/start-mysql.sh "]
VOLUME ["/var/lib/mysql "]
EXPOSE 3306
```

- Cgroups and namespaces are used to isolate containers from the host system.

Example: Show Processes

```
# docker run -d --name test opensuse sleep 120
d147fca80c23ca4cee227a2af2d89205e8b11440cf02df590edfd212d07f6
# docker exec -ti test ps -eH -o pid ,ppid ,stime ,cmd
  PID  PPID  STIME  CMD
    5    0 14:29  ps -eH -o pid ,ppid ,stime ,cmd
    1    0 14:29  sleep 120
# ps -eH -o pid ,ppid ,stime ,cmd
  PID  PPID  STIME  CMD
[... ]
 1278    1 Oct06  /usr/sbin/containerd --listen fd://
14197  1278 16:29  containerd-shim d147fca80c23ca4cee227a2
14209 14197 16:29  sleep 120
```

- Problem: generic docker images versus local configuration. In particular: how to share secrets (e.g. database passwords) between containers?
- One option: keep configuration files in the host system, outside the images. Use bind mounts to make them accessible to the container.
- Consequence: initialization at startup time of the container. Use shell scripts to set the container up.

Hands on MariaDB & ICAT ...

Dockerfile for MariaDB

```
FROM rkrahl/opensuse

RUN zypper --non-interactive install mariadb
RUN mkdir -p /etc/mysql /etc/mysql.d
COPY start-mysql.sh /etc/mysql
RUN chmod 0755 /etc/mysql/start-mysql.sh

CMD ["/etc/mysql/start-mysql.sh "]
VOLUME ["/var/lib/mysql "]
EXPOSE 3306
```

- ICAT can be installed and run dockerized.
- Generic Docker images that contain the software as distributed.
- Configuration is added to the container instance as a volume.
- Setup software and apply configuration at startup time of the container.