Tricks of the Captains

Adrian Mouat

Chief Scientist

Container Solutions

Tricks of the Captains

A hodgepodge of tips for Docker nirvana compiled from the brains in the Docker Captains program. And me.

Who am I?

Chief Scientist @ Container Solutions

Author of “Using Docker” published by O’Reilly

@adrianmouat

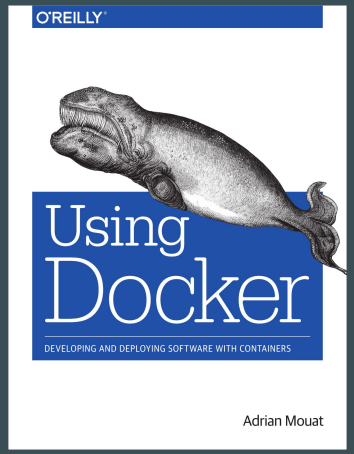

Daily Development

Configure docker ps output

Default output of `docker ps` (or `docker container ls`)

```
$ docker ps
CONTAINER ID      IMAGE               COMMAND                  ...
0f1f72c9aac0      nginx               "nginx -g 'daemon ...
```

Annoyingly takes up far too much screen width, difficult to read.

Configure docker ps output

Solution is to use the `--format` argument

```
$ docker ps --format \
    "table {{.Names}}\\t{{.Image}}\\t{{.Status}}"
NAMES               IMAGE               STATUS
web                 nginx               Up 25 minutes
```

https://docs.docker.com/engine/reference/commandline/ps/#formatting

Configure docker ps output

Make it a permanent default by adding to `config.json`

```
$ cat ~/.docker/config.json
{...
    "psFormat":
    "table {{.ID}}\\t{{.Names}}\\t{{.Image}}\\t{{.Status}}"}
```

https://docs.docker.com/engine/reference/commandline/cli/#configuration-files

Don’t bust the build cache

To keep builds fast, add dependencies before source code in Dockerfiles

```
...
COPY ./ /usr/src/
RUN npm install
...
```

```
...
COPY package.json /usr/src/
RUN npm install
COPY ./ /usr/src/
...
```

File mounting gotcha

Mounting a file as volume might not work as expected

```
$ cat index.html
Moby Rules!
$ docker run -d -p 8000:80 \
  -v $PWD/index.html:/usr/share/nginx/html/index.html nginx
0cdacef2cbaea960f710d90900b23c57550aaf626ccd2752f3a9287b7e51022b
$ curl localhost:8000
Moby Rules!
$ vi index.html
$ cat index.html
Gordon Rules!
$ curl localhost:8000
Moby Rules!
```

File mounting gotcha

Volumes are mounted at the inode level
Text editors save to a new inode

Solutions:

- mount parent directory (`-v $PWD:/usr/share/nginx/html`)
- copy modified file (`cp new.html index.html`)
- overwrite with `>` (`echo "bla" > index.html`)

Cleaning Up

Delete “dangling” images (<none> images)

```
$ docker image prune
WARNING! This will remove all dangling images.
Are you sure you want to continue? [y/N] y
Deleted Images:
deleted:
sha256:708624719836212ccb681d5898a64ebfcc4569f37460537609f6db6…
…
Total reclaimed space: 3.677 GB
```

Cleaning Up

Delete stopped containers

```
$ docker container prune
WARNING! This will remove all stopped containers.
Are you sure you want to continue? [y/N] y
Deleted Containers:
6e5033be3e106d04912fb91b966abc693b77ae47d85946190bdbe73c48112959
…
Total reclaimed space: 304.6 MB
```

Cleaning Up

```
$ docker volume prune
WARNING! This will remove all volumes not used by at least one
container.
…
Total reclaimed space: 3.494 GB

$ docker system prune
WARNING! This will remove:
    - all stopped containers
    - all volumes not used by at least one container
    - all networks not used by at least one container
    - all dangling images
```

Container Lifecycle

Start-up Dependably

Do not require containers to start in sequence

If a container depends on another service:

- It should **wait** for that service
- Do not crash - back off
- Do this in application code
    - or start-up script if you can't

See 12 Fractured Apps by Kelsey Hightower

Shutdown Gracefully

When Docker stops a container, it will

- Send the container a `SIGTERM` signal
- Wait 10s for the container to stop
- Hard kill the container with a `SIGKILL`

Shutdown Gracefully

Proper handling of `SIGTERM` will mean:

- The application gets a chance to “tidy up”
  - close network connections, sockets, handles
  - write data to file or database
  - output to log
- Faster shutdown of the container

Shutdown Gracefully

To ensure your application receives signals either

- Run it as PID 1
  - Use `exec` in any start-up scripts
- Or forward signals to it
  - `tini` can help https://github.com/krallin/tini
- And prefer node to npm for starting node.js apps
  - see Bret Fisher's "Node and Docker Good Defaults"

Use Healthchecks

Used by Docker to determine “health” of container

```
FROM nginx

RUN apt-get update && apt-get install -y curl
HEALTHCHECK --interval=10s --timeout=3s \
  CMD curl -f http://localhost/ || exit 1
```

Use Healthchecks

Used by Docker to determine “health” of container

```
$ docker ps
CONTAINER ID    ...         STATUS
79616fdd4308                Up 3 seconds (health: starting)

$ docker ps
CONTAINER ID    ...         STATUS
79616fdd4308                Up 16 seconds (healthy)
```

Security

Read-only FS

An easy way to improve security

```
$ docker run -d --name n1 --read-only -p 8000:80 \
    --tmpfs /var/run --tmpfs /var/cache/nginx nginx
c1da395bec73ef7933fecb6d8d821140ce203c426c433e5102d25e46cdb66537

$ docker exec n1 /bin/bash -c \
    'echo "HACKED" > /usr/share/nginx/html/index.html'
/bin/bash: /usr/share/nginx/html/index.html: Read-only file system
```

Don’t run as root

Set a **`USER`** in Dockerfiles e.g:

```
FROM debian
RUN groupadd -r mygroup && useradd -r -g mygroup myuser
…
USER myuser
```

Or use the `nobody` user

Don’t run as root

Sometimes need to change user at run-time

sudo works, but has a drawback:

```
$ docker run debian-with-sudo sudo -u nobody ps ax
  PID TTY      STAT   TIME COMMAND
    1 ?        Rs     0:00 sudo -u nobody ps ax
    7 ?        R      0:00 ps ax
```

Don’t run as root

Instead use gosu by Tianon Gravi

```
$ docker run debian-with-gosu gosu nobody ps ax
  PID TTY      STAT   TIME COMMAND
    1 ?        Rs     0:00 ps ax
```

https://github.com/tianon/gosu

Thanks for Listening!

@docker

#dockercon

References

Good Defaults for Node and Docker - Bret Fisher
https://github.com/BretFisher/node-docker-good-defaults

12 Fractured Apps - Kelsey Hightower
https://medium.com/@kelseyhightower/12-fractured-apps-1080c73d481c

Least Privilege Containers - Nathan McCauley and Diogo Monica

gosu - *https://github.com/tianon/gosu*

tini - *https://github.com/krallin/tini*

Thanks to the captains for discussion, particularly Bret, Laura, Marcos and Antonis