USING DOCKER SAFELY

ADRIAN MOUAT

NLUUG 28 MAY 2015

ContainerSolutions

LOT OF NEGATIVE COMMENTS ON DOCKER SECURITY

- "Containers Don't Contain"
  - Daniel Walsh, RedHat
  - https://opensource.com/business/14/7/docker-security-selinux
- "... total systemic failure of all logic related to image security"
  - Jonathan Rudenberg, Flynn.io
  - https://titanous.com/posts/docker-insecurity
- "... gives the apps root access"
  - Alex Larrson, RedHat
  - https://news.ycombinator.com/item?id=9086751

SO CAN CONTAINERS BE USED SECURELY?

YES!

OVERVIEW

- THINGS TO WORRY ABOUT!
- PRIMARY DEFENCES
- TIPS AND TECHNIQUES

KERNEL ATTACKS

DENIAL OF SERVICE

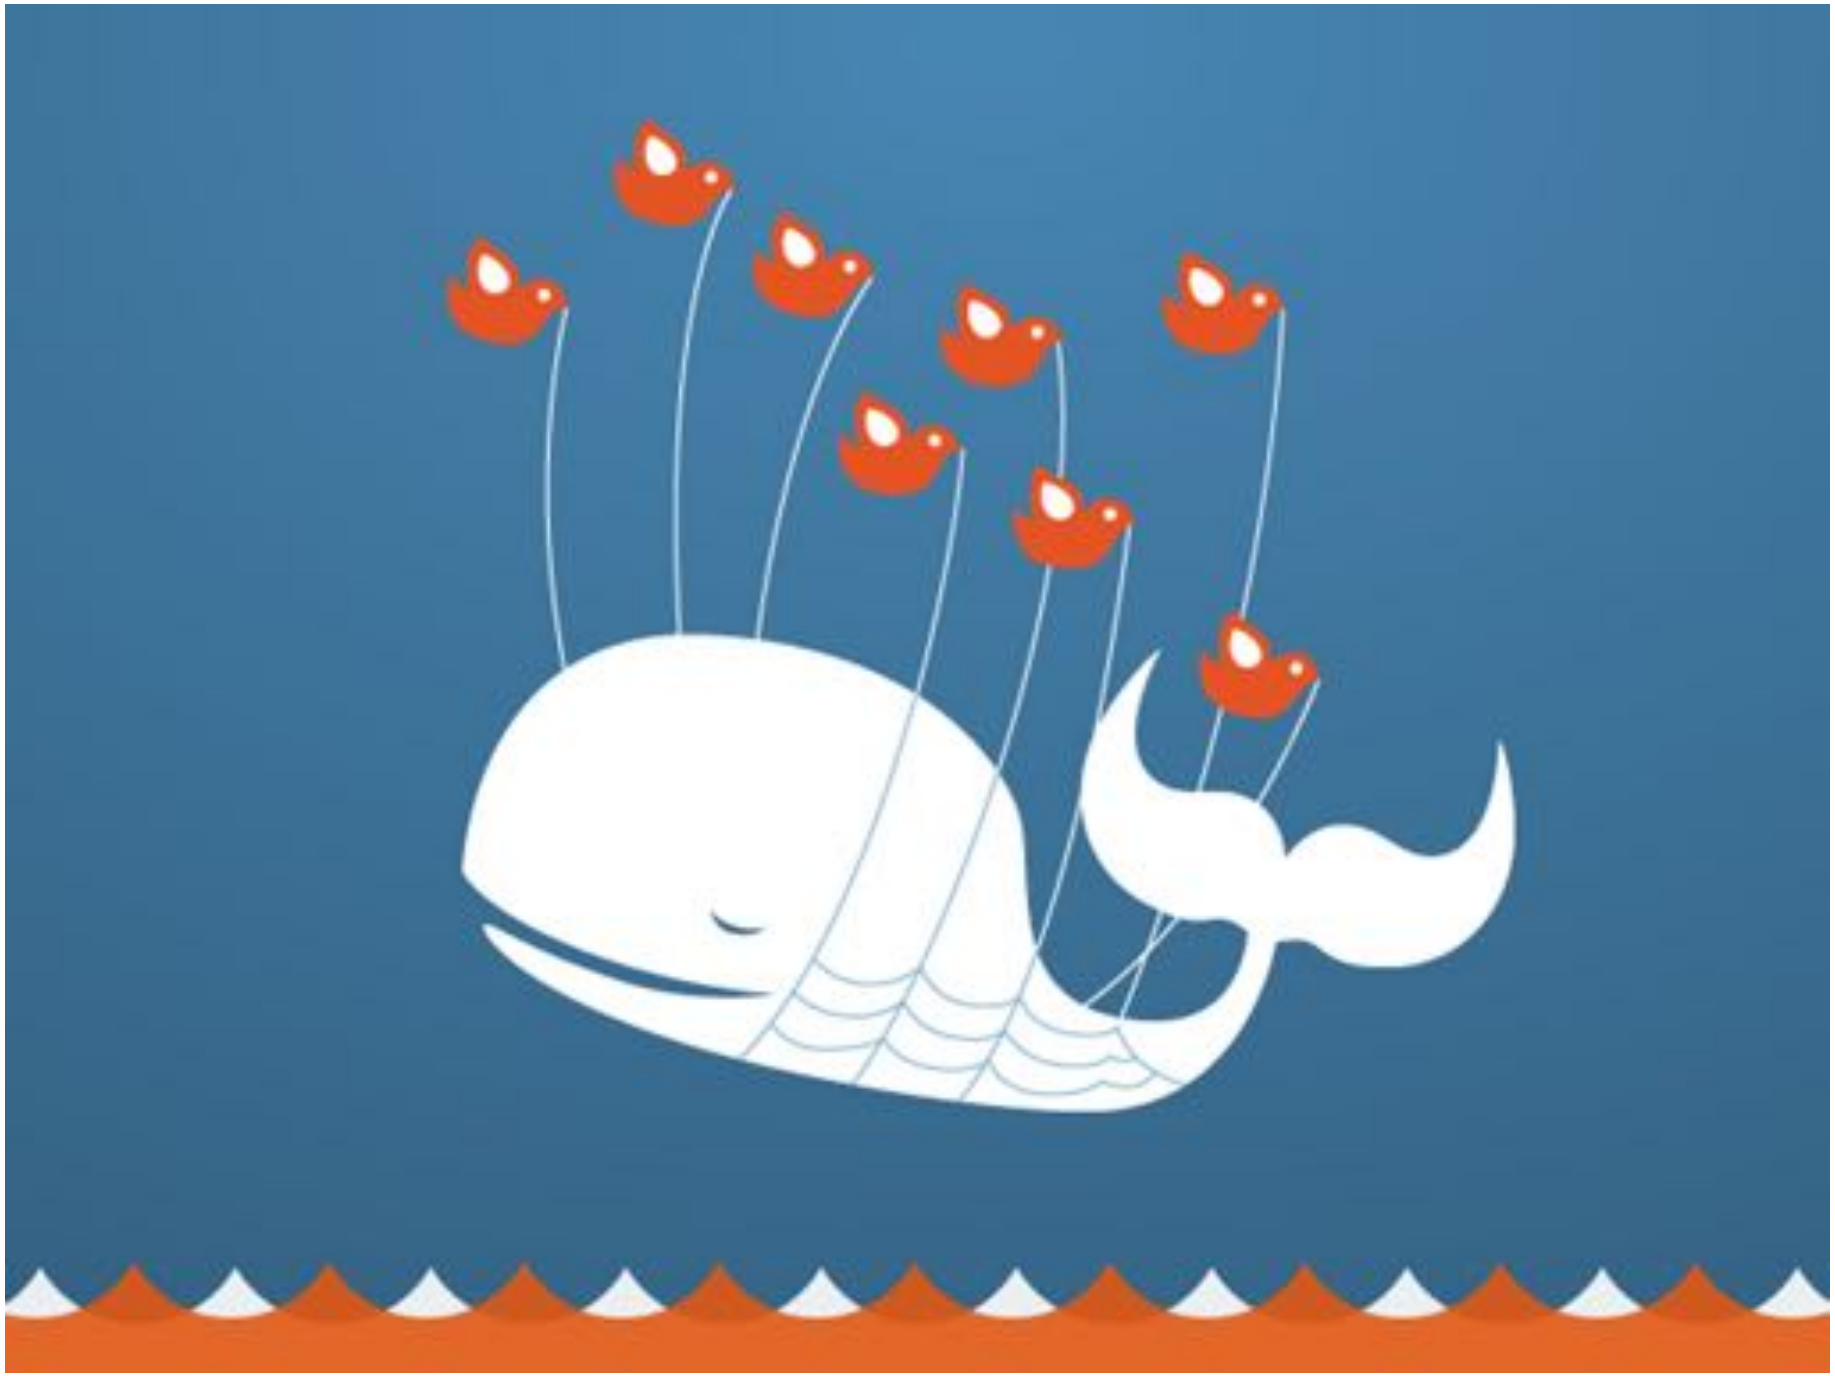

CONTAINER BREAKOUTS

POISONED IMAGES

SNIFFING SECRETS

THINK "DEFENCE IN DEPTH"

MULTIPLE LINES OF DEFENCE

- CONTAINERS
- VMS
- ENCRYPTION
- MONITORING
- AUDITING
- ...

VIRTUAL MACHINES

- Use VMs to segregate groups of containers

DOCKER PRIVILEGES

==

ROOT PRIVILEGES

- BE CAREFUL WHO YOU GIVE ACCESS!
- SECURE REMOTE API

USERS ARE NOT NAMESPACED

- Root in container is root on host

SET A USER

- Create a user in your Dockerfile
- Change to the user via USER or su/sudo/gosu

```
RUN groupadd -r user && useradd -r -g user user
USER user
```

SET CONTAINER FS TO READ-ONLY

```
$ docker run --read-only debian touch x
touch: cannot touch 'x': Read-only file system
```

SET VOLUMES TO READ-ONLY

```
$ docker run -v $(pwd)/secrets:/secrets:ro \
         debian touch /secrets/x
touch: cannot touch '/secrets/x': Read-only file system
```

DROP CAPABILITIES

```
$ docker run --cap-drop SETUID --cap-drop SETGID myimage
$ docker run --cap-drop ALL --cap-add ...
```

FINER GRAINED LIMITING

SELINUX

- By NSA!
- Policy based
- MAC not DAC
- File access, sockets, interfaces
- Also AppArmor

SET CPUSHARES

```
$ docker run -d myimage
$ docker run -d -c 512 myimage
$ docker run -d -c 512 myimage
```

SET MEMORY LIMITS

```
$ docker run -m 512m myimage
```

TURN OFF INTER-CONTAINER COMMUNICATION

```
$ docker -d --icc=false
```

NOW CONTAINERS CAN'T ATTACK EACH OTHER

PEACE :)

BUT A BIT USELESS

ALLOW LINKED CONTAINERS TO COMMUNICATE

```
$ docker -d --icc=false --iptables
```

BEWARE BUGS

- Dependent on Kernel Parameters
  - /proc/sys/net/bridge/bridge-nf-call-iptables
  - /proc/sys/net/bridge/bridge-nf-call-ip6tables
  - https://github.com/docker/docker/pull/11405
- Drop Rule Placement
  - https://github.com/docker/docker/pull/11526

VERIFY IMAGES

- Only use automated builds, check Dockerfile
- Build yourself
- Pull by digest

```
$ docker pull debian@sha256:0ecb2ad60
```

DEFANG SETUID/SETGID BINARIES

- Applications probably don't need them
- So don't run them in production

TO FIND THEM

```
$ docker run debian \
  find / -perm +6000 -type f -exec ls -ld {} \; 2> /dev/null
-rwsr-xr-x 1 root root 10248 Apr 15 00:02 /usr/lib/pt_chown
-rwxr-sr-x 1 root shadow 62272 Nov 20  2014 /usr/bin/chage
-rwsr-xr-x 1 root root 75376 Nov 20  2014 /usr/bin/gpasswd
-rwsr-xr-x 1 root root 53616 Nov 20  2014 /usr/bin/chfn
...
```

TO DEFANG THEM

```
FROM debian:wheezy
RUN find / -perm +6000 -type f -exec chmod a-s {} \; \
    || true
```

RESULT

```
$ docker build -t defanged-debian .
...
Successfully built 526744cf1bc1
$ docker run --rm defanged-debian \
  find / -perm +6000 -type f -exec ls -ld {} \; \
  2> /dev/null | wc -l
0
$
```

SHARING SECRETS

BAKE IT INTO THE IMAGE


NO.

ENVIRONMENT VARIABLES

```
$ docker run -e API_TOKEN=MY_SECRET myimage
```

- Suggested by 12 factor apps
- Can be seen too many places
  - linked containers, inspect
- Can't be deleted

MOUNTED VOLUMES OR DATA VOLUME CONTAINERS

```
$ docker run -v /secretdir/keyfile:/keyfile:ro myimage
$ docker run --volumes-from my-secret-container myimage
```

- Works, but icky
- Files can get checked in by accident

KEY-VALUE STORE

- etcd (plus crypt)
  - https://github.com/coreos/etcd
  - https://github.com/xordataexchange/crypt
- vault
  - https://hashicorp.com/blog/vault.html
- keywhiz
  - https://github.com/square/keywhiz/
- Can control leases, store encrypted
- Still requires some sort of authentication token

CONCLUSION

- Many aspects to container security
- Get it wrong and you hand over the keys to your host
- Get it right and you have defence in depth
  - More secure than VMs alone

- Chief Scientist @ Container Solutions
- http://www.container-solutions.com
- Writing "Using Docker" for O'Reilly
- @adrianmouat

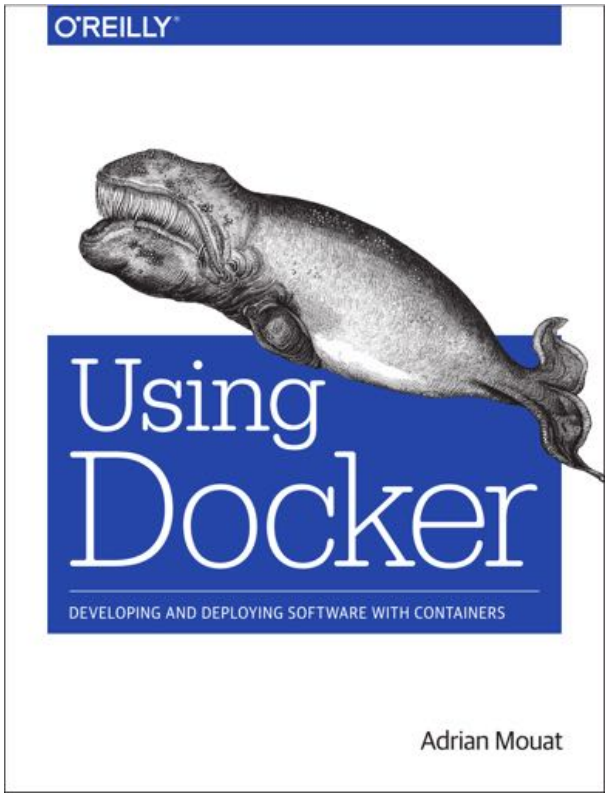
O'REILLY
Using
Docker
DEVELOPING AND DEPLOYING SOFTWARE WITH CONTAINERS
Adrian Mouat