# DOCKER AND PROVENANCE

## WHO ARE YOU AND WHERE THE HELL DID YOU COME FROM?

ADRIAN MOUAT

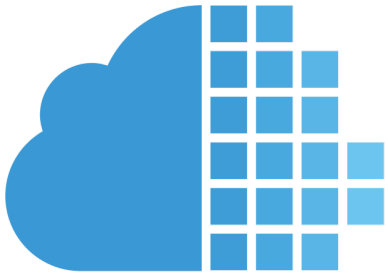

ContainerSolutions

# PROVENANCE

1. origin, source
2. the history of ownership of a valued object or work of art or literature

# AND WHAT'S THAT GOT TO DO WITH DOCKER?

- Want to be sure that downloaded images:
  - Haven't been tampered with
  - Are exactly the same thing the developers tested
- Want to be able to answer our questions!

# OTHERWISE

- Could have been replaced with something malicious
  - or simply corrupted
- May be subtly different to tested version
  - potentially broken
  - breaks promise of Docker!

# HANG ON, WE'VE SOLVED THIS BEFORE

# ENTER SECURE HASHES

- Fingerprint of file or data
- Unique to file (no collisions)
- Changing the file changes the hash
- SHA1, MD5 etc

# USED BY PACKAGE MANAGERS

- Linux package managers
  - Signed package lists with checksums
  - Checksum verified on download before handling file

# DOCKER IMAGE PROVENANCE

- Same idea as package managers
- Image should be "verified" once downloaded

# IMAGE PROVENANCE

```
$ docker pull busybox
busybox:latest: The image you are pulling has been verified
511136ea3c5a: Already exists
df7546f9f060: Already exists
ea13149945cb: Already exists
4986bf8c1536: Already exists
Status: Image is up to date for busybox:latest
```

# CURRENTLY BORKED

- Unpacks images before checking
  - as root
- Only checks manifest file
  - not file contents
- Testing of checksums is broken

https://titanous.com/posts/docker-insecurity

# YOU CANNOT TRUST IMAGES FROM THE HUB!

# WORKAROUND

- Manually download images from secure site
- Test checksum
- Ingest with docker load

https://securityblog.redhat.com/2014/12/18/before-you-initiate-a-docker-pull/

# MORE PROBLEMS

- Dockerfiles download resources
  - apt-get
  - wget/curl
- What if tampered with?
  - in transit or at source
- Or just different/updated?

# SIMPLE SOLUTION

- But requires diligence from Dockerfile author
- Peg versions of apt-get
- Test checksums of other files

# EXAMPLES

- Good
  - WordPress
  - MongoDB
- Bad
  - Python
  - Jenkins

# STILL AN ISSUE

- apt-get will pull in dependencies
  - those can be >=
- Can use tools like aptly
  - http://www.aptly.info/

# THE PRESENT

- Don't use Hub in production!
- Consider RedHat solution
- Dockerfiles
  - Verify downloads
  - Peg versions
- Moving fast has bitten Docker
  - Need to migrate images

# THE FUTURE

- New registry (API and implementation)
- libtrust
- RedHat etc will push for security
- Never achieve perfect security
  - "Reflections on Trusting Trust"
  - http://www.ece.cmu.edu/~ganger/712.fall02/papers/p761-thompson.pdf

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

# LINKS

- TarSum insecurity
  - https://github.com/docker/docker/issues/9719
- Docker Insecurity Blog
  - https://titanous.com/posts/docker-insecurity
- Security advisory on symlinks and path traversal
  - https://groups.google.com/forum/#!msg/docker-user/nFAz-B-n4Bw/0wr3wvLsnUwJ
- LibTrust
  - https://github.com/docker/libtrust
- RedHat advice on pulling images
  - https://securityblog.redhat.com/2014/12/18/before-you-initiate-a-docker-pull/