Docker File System Isolation

By
Darrin Schmitz
David Huff
Destiny Velasquez

LA-UR-15-25911

Specifications

- HP ProLiant DL380p Gen8 servers
- Head node has 32 cores and 32 GB RAM
- 10 child nodes have 24 cores and 24 GB RAM
- Operating system: CentOS 6.6
- Containers: Docker version 1.6

Abstract Overview

- Our goal
- Technical difficulties
- Overall, we believe Docker is a good security option, even though there are some security risks involved

What is a Container?

- Between a virtual machine and a chroot
- Native hardware utilization
- Able to run different operating systems

Why use Docker?

1. Pre-configures its network bridges
2. Available documentation
3. Portable and recoverable images

Docker Normal Setup

- Docker bridge directly connected to node
- IP forwarding use
- The IP ranges for the containers are 172.17.0.0/20
- Daemon configures iptables

Docker Normal Setup Diagram

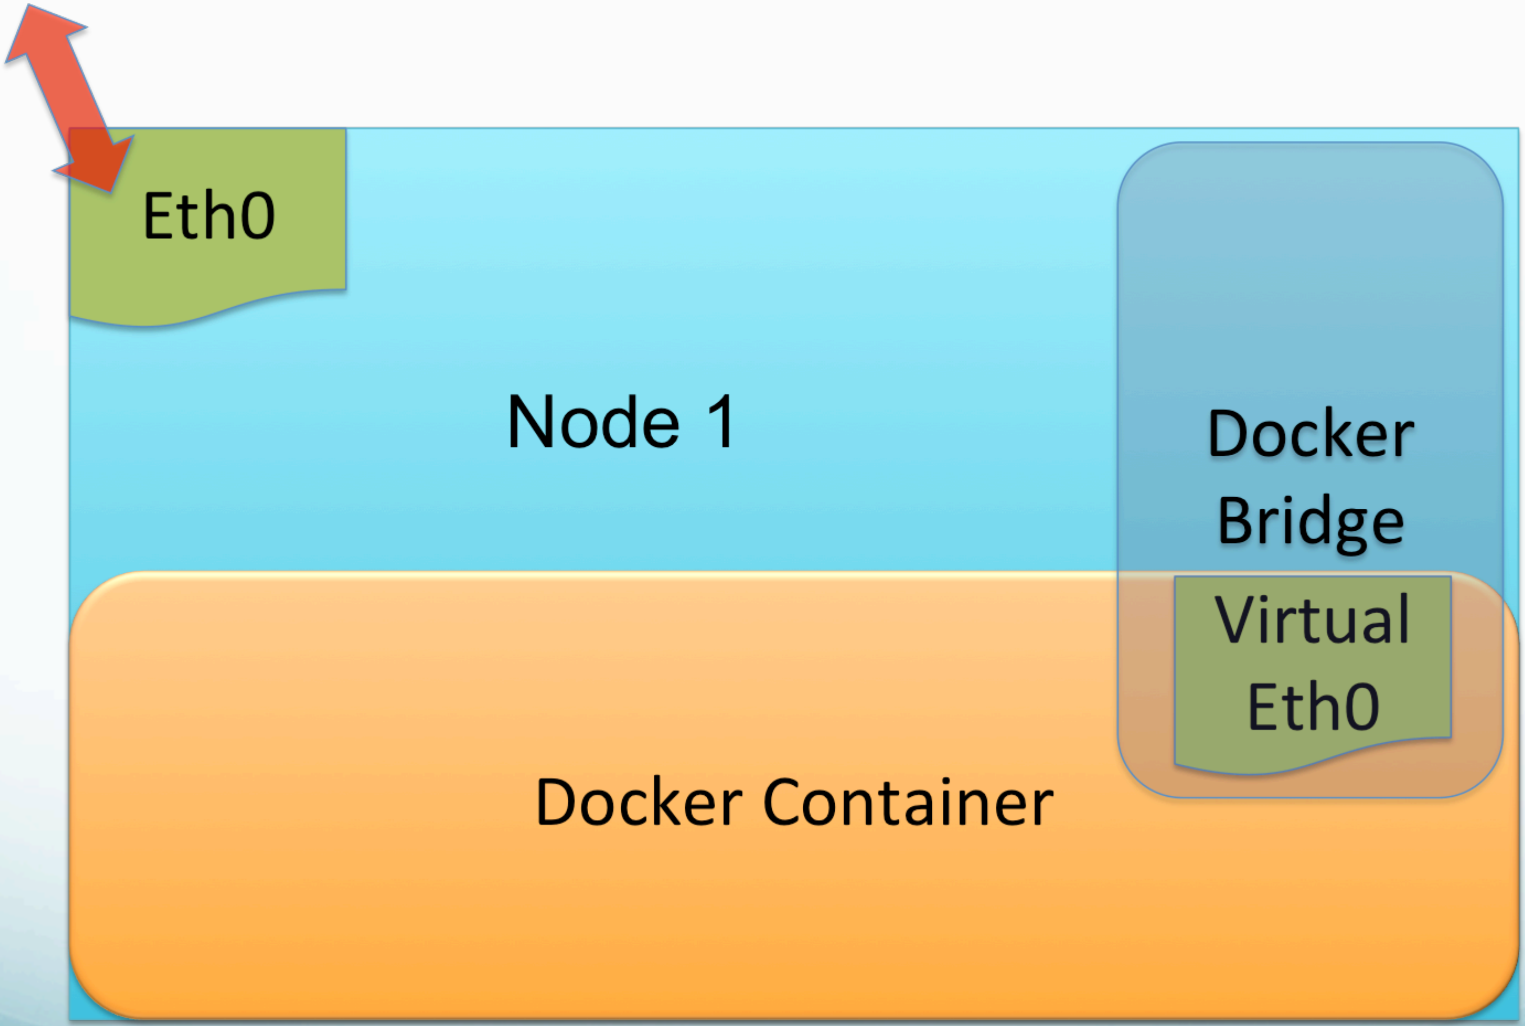

Problems With Default Setup

- Same IP addresses are assigned to different containers on different nodes
- Iptables and bridges are not cleaned up by Docker

Steps to Create a Docker Network With OpenMPI

1. Install Docker
2. Set up the bridge manually
3. Set up SSH-keys
4. Set up OpenMPI

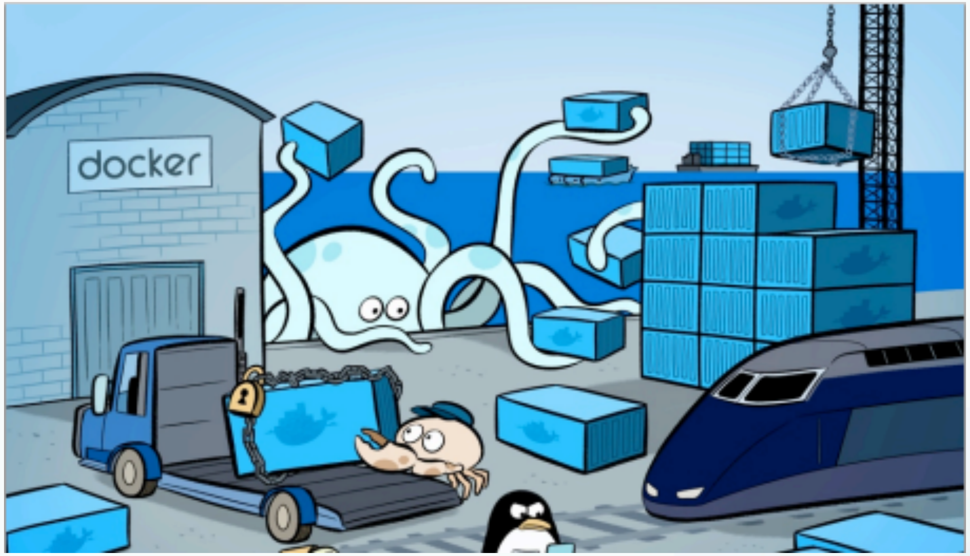

5. Set up the Docker daemon to give out unique IP-addresses

LA-UR-15-25911

https://www.linkedin.com/pulse/docker-containers-kubernetes-smart-ecosystem-solution-yasser-emam

Bridge

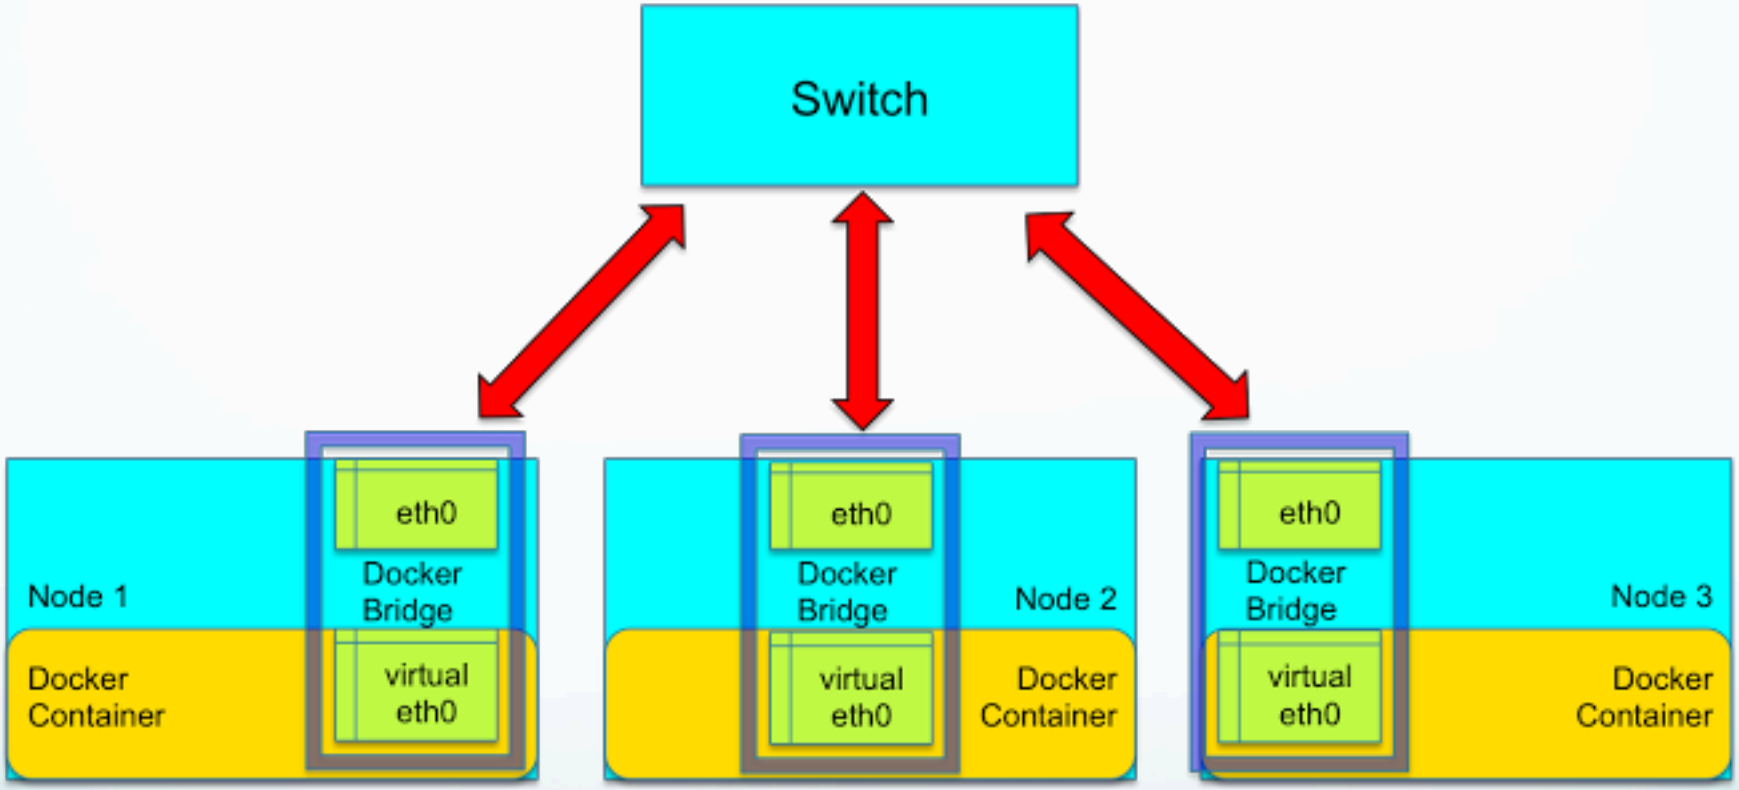
Switch
eth0
Docker
Bridge
Node 1
Docker
Container
virtual
eth0
eth0
Docker
Bridge
Node 2
virtual
eth0
Docker
Container
eth0
Docker
Bridge
Node 3
virtual
eth0
Docker
Container

SSH-Keys & OpenMPI & Mounting

- Generate the SSH-keys and place the public key into the authorized-keys file
- Set up the /etc/openmpi/default-openmpi -hostnames file, and set the path to the OpenMPI libraries
- Mounting is as simple as using Dockers –v flag

Docker Daemon

- The Docker Daemon sets up the bridge
- The IP range for the containers is set up by the daemon
- There is a flag to assign a custom bridge to the daemon

Docker Hub

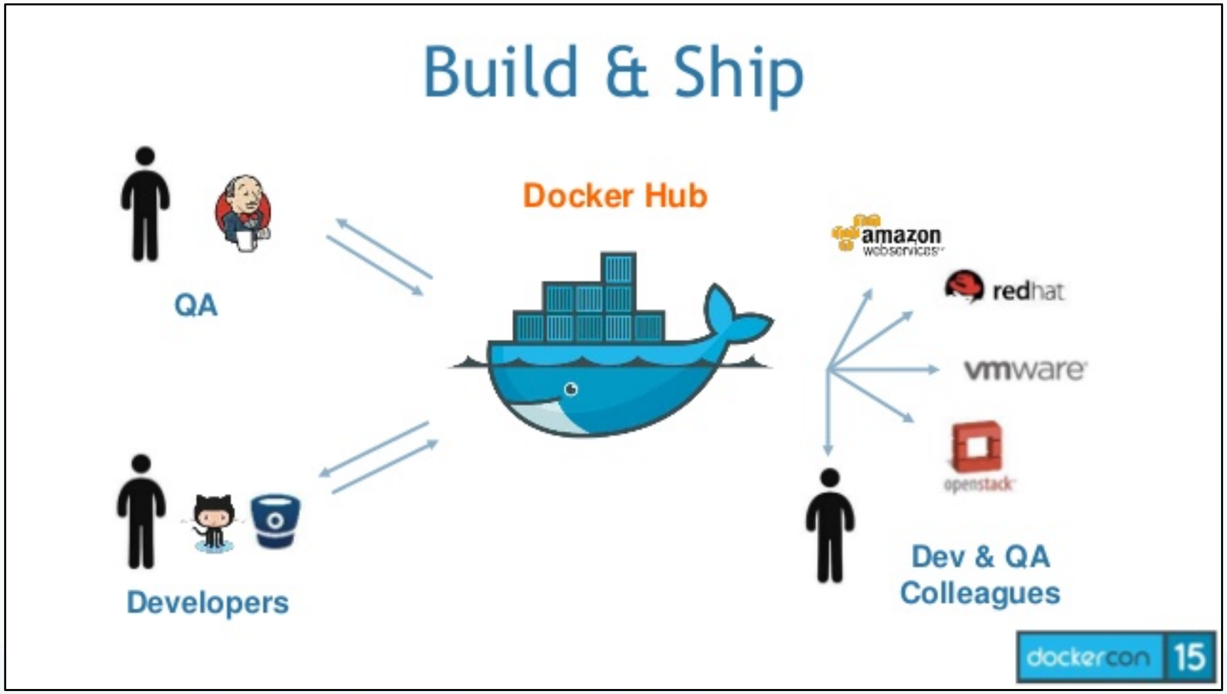

LA-UR-15-25911

http://jenkins-ci.org/content/official-jenkins-lts-docker-image

Problems With Docker

- Docker's bridge needs to connect to the switch directly
- Services do not start at the start of the terminal
- Environment variables are not permanent
- IP-addresses cannot be statically set
- /etc/hosts file is constantly being overwritten

Benchmarks

Write

```
dd if=/dev/urandom of=/Yellow/File bs=1024 count=1024000
dd if=/dev/urandom of=/home/File bs=1024 count=1024000
```

Read

```
dd if=/Yellow/File of=/dev/null bs=1024
dd if=/home/File of=/dev/null bs=1024
```

Benchmark Results

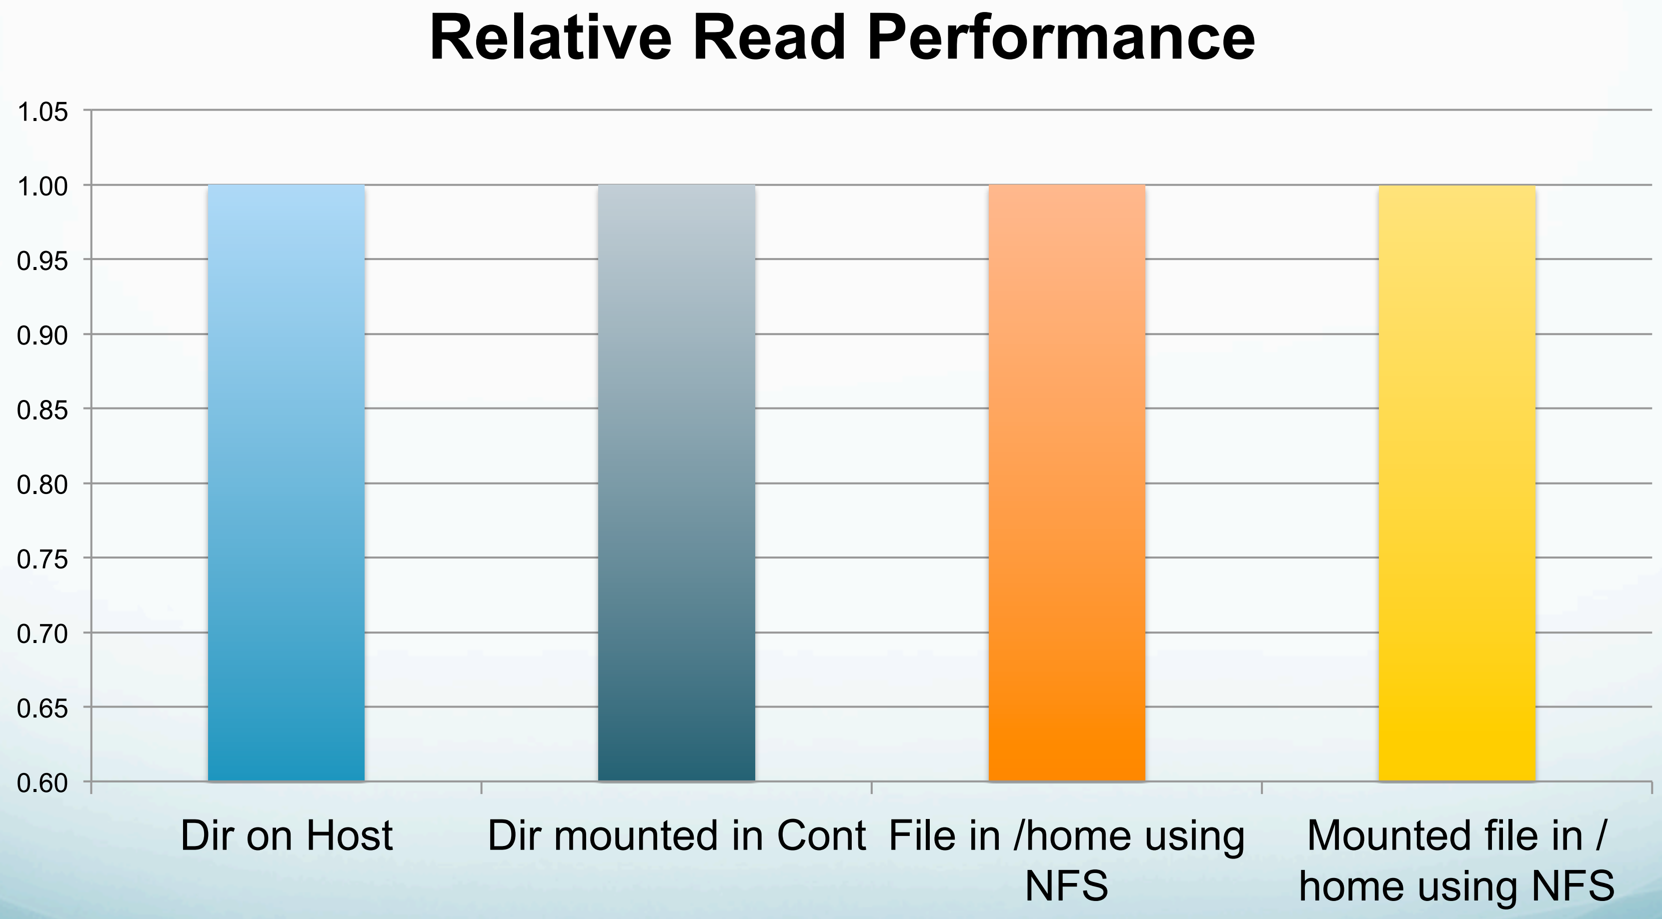

Benchmark Results

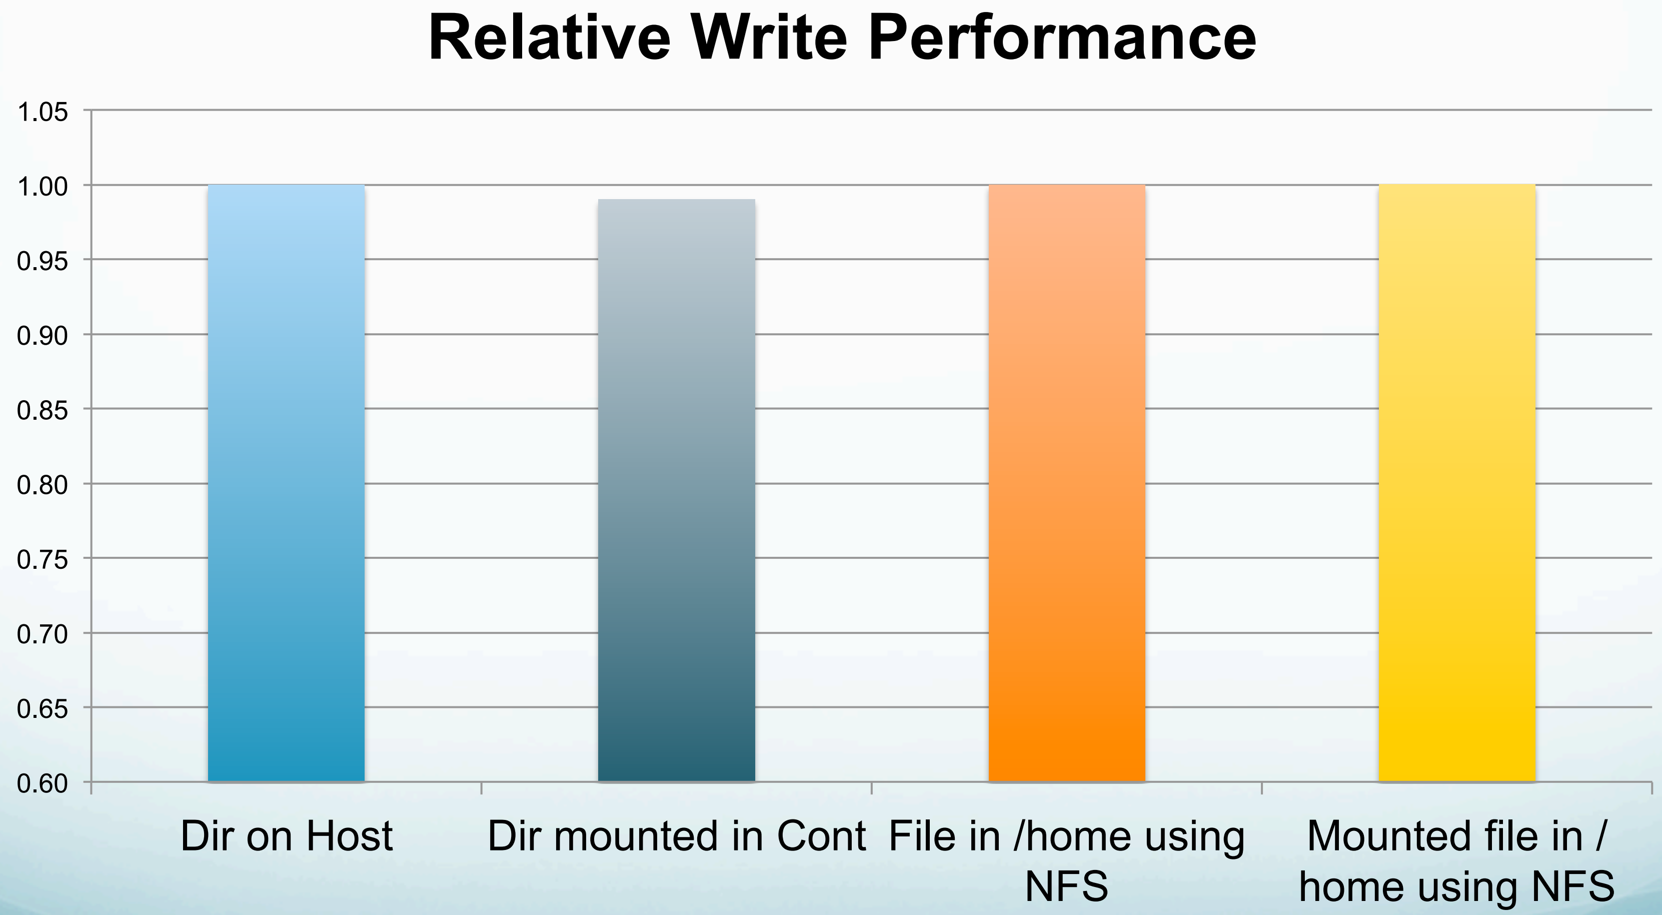

LA-UR-15-25911

CVE’s

- Insecure opening of file-descriptor 1 leading to privilege escalation (CVE-2015-3627)
- Symlink traversal on container respawn allows local privilege escalation (CVE-2015-3629)
- Read/write proc paths allow host modification & information disclosure (CVE-2015-3630)

Security Risks

- The current version of Docker fixes these security holes
- As of the 14th of July, 1.7.1 is compatible with CentOS 6.6
- The isolation provided by Docker is not as robust as the segregation established by hypervisors for virtual machines

Security Recommendations

- Use containers only on unclassified data/file systems
- Containers run with a whitelisted root
- Access control via SSH Keys
- Set up a password between data locations
- Don’t give root to the user
- Set up user account in the container

Future Research

- Write a launch script that works with SLURM/Moab to automatically provision the container environment.

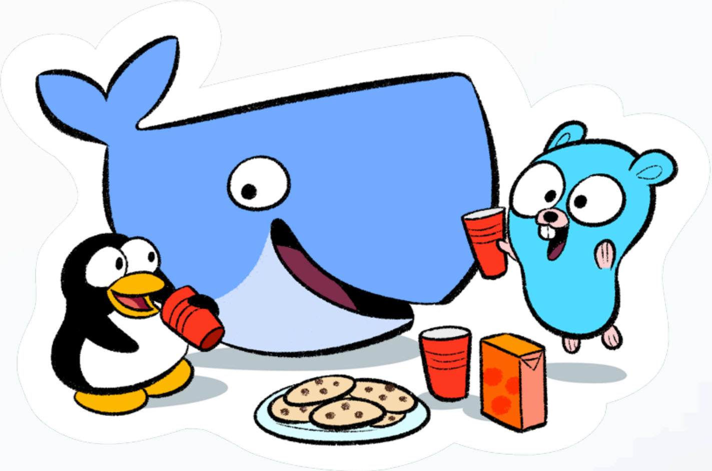

- Investigate bind mounts using Lustre and Panasas.
- Investigate using containers in an SELinux environment.

LA-UR-15-25911

https://docs.docker.com/

Conclusion

- We met the goal of our project by proving Docker is a lightweight security option
- Although there are some security holes to be concerned about, we’ve provided some security recommendations for Docker
- Docker would be a useful option for separating Yellow and Turquoise data

References

1. https://sites.google.com/a/probe.newmexicoconsortium.org/cscnsi-2015-vermilion/
2. https://www.docker.com/
3. https://hub.docker.com/
4. https://nvd.nist.gov/

Questions?

docker