WHAT'S REALLY THE DIFFERENCE

Between a VM and a Container?

YOUR CLOUDS.
OUR EXPERTISE.

WELCOME

ONE SIMPLE IDEA

CHANGED EVERYTHING.

1873


COLGATE'S
ANTISEPTIC
DENTAL POWDER
COLGATE & CO.
COLGATE'S
DENTAL
POWDER
COLGATE & CO.
ENDORSED BY EMINENT DENTISTS

1896

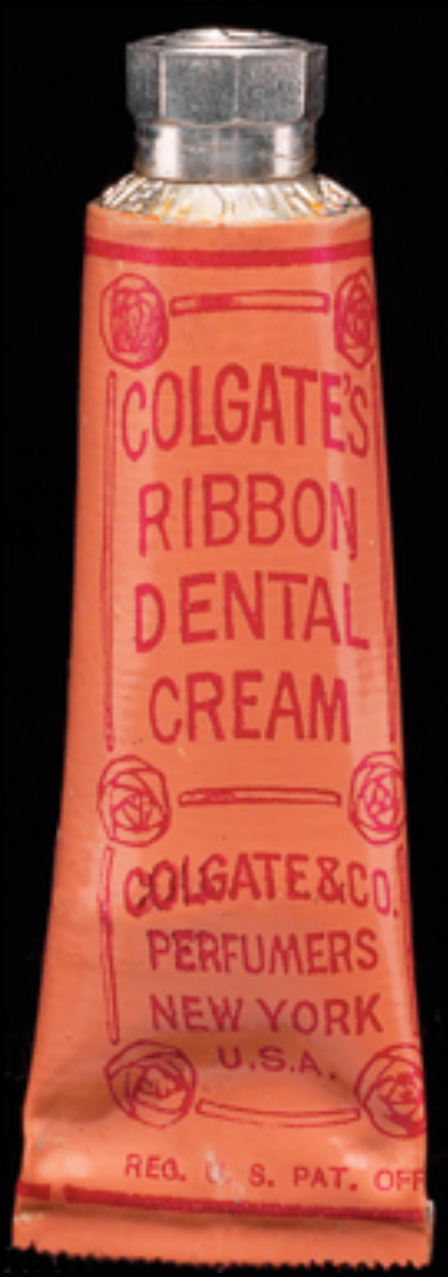
COLGATE'S
RIBBON
DENTAL
CREAM
COLGATE&CO.
PERFUMERS
NEW YORK
U.S.A.

“

WE COULDN'T IMPROVE THE PRODUCT SO WE IMPROVED THE TUBE.

”

- Colgate, 1908

rackspace®


1962
Colgate Research Center

1978

THE LAB ASSISTANT


Questions or comments?
Call toll-free 1-800-468-6502
Dist. by: COLGATE-PALMOLIVE COMPANY
New York, NY 10022 www.colgate.com
Colgate®
Fluoride Toothpaste
Cavity Protection
Strengthens Teeth with Active Fluoride
Great Regular Flavor
ADA ACCEPTED American Dental Association
NET WT 4.6 OZ (130 g)
Colgate®
Fluoride Toothpaste
Cavity Protection
Strengthens Teeth with Active Fluoride
Great Regular Flavor
ADA ACCEPTED American Dental Association
NET WT 4.6 OZ (130 g)

OMG, SALES DOUBLED!

YOUR CLOUDS.
OUR EXPERTISE.

WELL, ACTUALLY…

TOOTHPASTE TUBE THEORY

1) PRESSURE BUILT UP IN A FINITE BOUNDED SYSTEM NEEDS TO BE RELEASED SOMEWHERE OR THE SYSTEM WILL BREAK.

2) THERE ARE DIMINISHING RETURNS TO SQUEEZING THE TUBE AFTER A CERTAIN POINT.

IDEA

ADRIAN OTTO

Distinguished Architect, Rackspace
Founder, OpenStack Containers Team
Founder and PTL, OpenStack Magnum
Organizer, Docker Los Angeles

THE DIFFERENCE

1. EFFICIENCY
2. PERFORMANCE
3. SECURITY

HISTORY OF VIRTUALIZATION

- 1960's IBM S/360 Mainframes are the 800# Gorilla
  - Single user system designed for batch jobs
- 1963 MIT Project MAC ($2M grant from DARPA)
  - MAC = Multiple Access Computing: Multics
  - Vendor Choice == GE (Commercial interest in time sharing computer)
  - Whoops! IBM panicked! Created CP-40 for Bell Labs, CP-67.
    - Virtual Machines on the CP-67 using "CP (Control Program)" in 1967!
- 1987 Insignia Solutions "SoftPC"
- 1997 Apple (Connectrix) "VirtualPC"
- 1999 VMWare "VMWare Workstation"

APPLICATION VIRTUALIZATION

- 1990 Sun Microsystems “Stealth”
  - Address C/C++ Portability problems
  - Renamed Oak -> Webrunner -> Java (1995)
- 1996 Sun Microsystems “Java”
  - Java Development Kit (JDK)
  - Java Runtime Environment (JRE)
    - Java Virtual Machine (JVM)

OPEN SOURCE VIRTUALIZATION

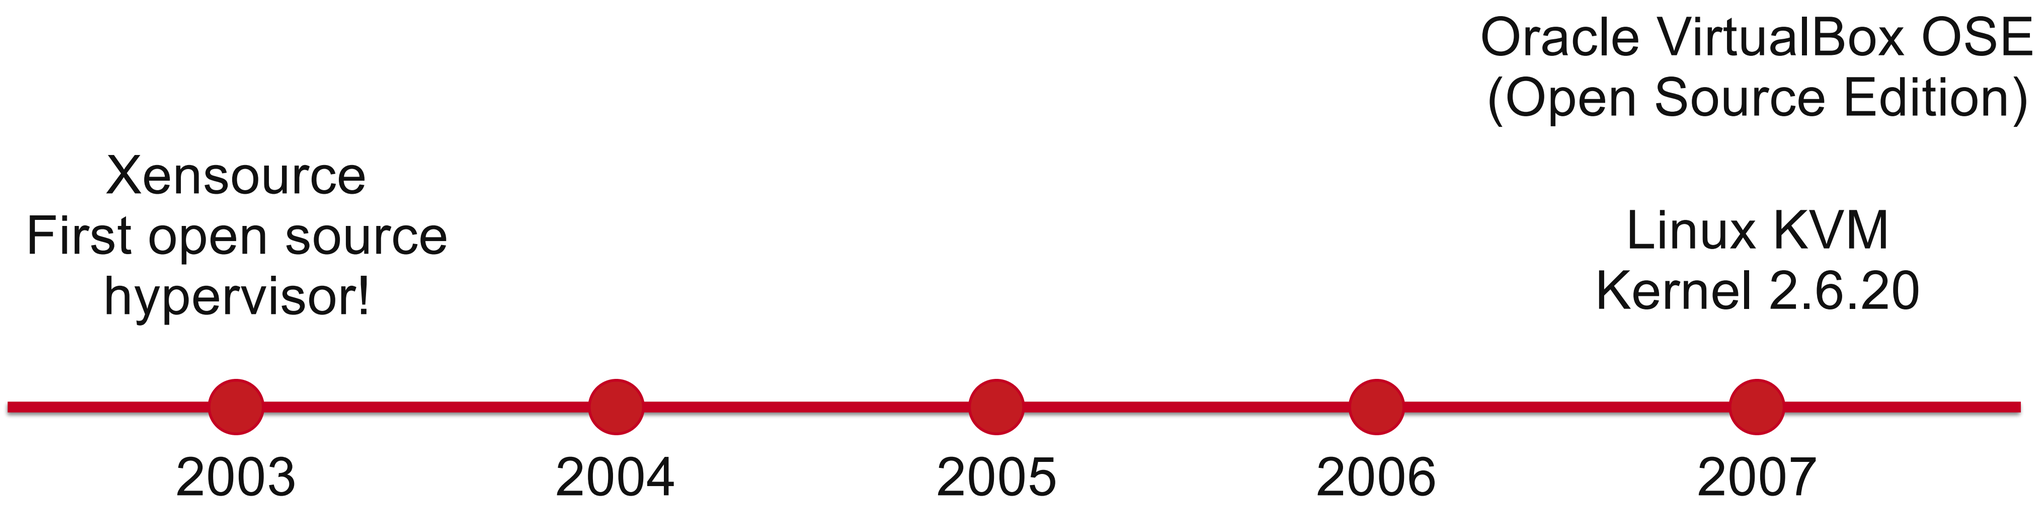

HISTORY OF CONTAINERS (1/2)

- 1979 UNIX chroot (added to BSD in 1982)
- 2000 FreeBSD Jails (filesystems, users, networks)
- 2001 Linux VServer (VPS Solution)
- 2005 OpenVZ (filesystems, users/groups, process tree, networks, devices, IPC)
- 2006 Process Containers (Linux Kernel 2.6.24, limit CPU, mem, disk, network IO)
- 2008 Control Groups (cgroups added to Linux Kernel)
- 2008 LXC (LinuX Containers, CLI and language bindings for 6 languages)
- 2011 Warden, CloudFoundry
- 2013 LMCTFY, Google

HISTORY OF CONTAINERS (2/2)

DotCloud becomes Docker, Inc.

CoreOS introduces Rocket

Microsoft Windows Containers

2013 2014 2015 2016 2017

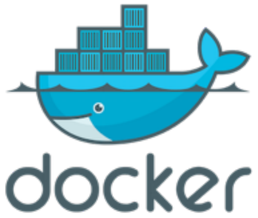

EVERYTHING CHANGED IN 2013

2013

DOCKER IMAGE

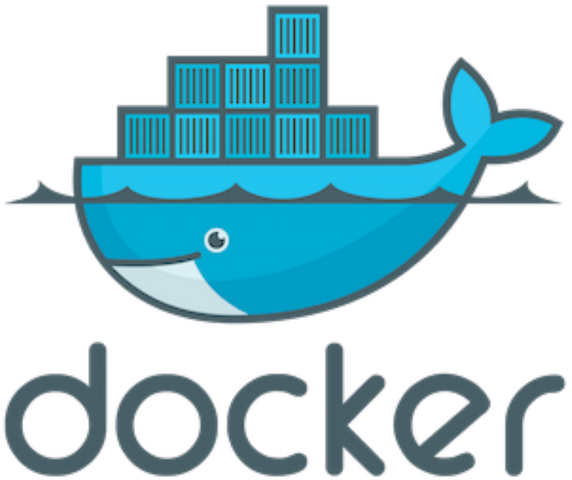

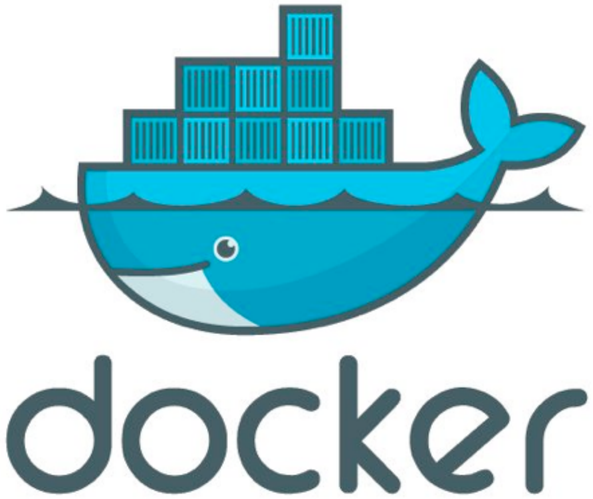

LINUX CGROUPS

- Kernel Feature
- Groups of processes
- Control resource allocations
  - CPU
  - Memory
  - Disk
  - I/O
- May be nested

LINUX KERNEL NAMESPACES

- Kernel Feature
- Restrict your view of the system
- Mounts (CLONE_NEWNS)
- UTS (CLONE_NEWUTS)
  - uname() output
- IPC (CLONE_NEWIPC)
- PID (CLONE_NEWPID)
- Networks (CLONE_NEWNET)
- User (CLONE_NEWUSER)
  - See also: privileged/unprivileged modes
- May be nested

DOCKER CONTAINER IMAGE

- NOT A FILESYSTEM
- NOT A VHD
- Basically a tar file
- Has a hierarchy
- Arbitrary depth
- Layered filesystem
  - Top layer can be writable
- Fits into the Docker Registry

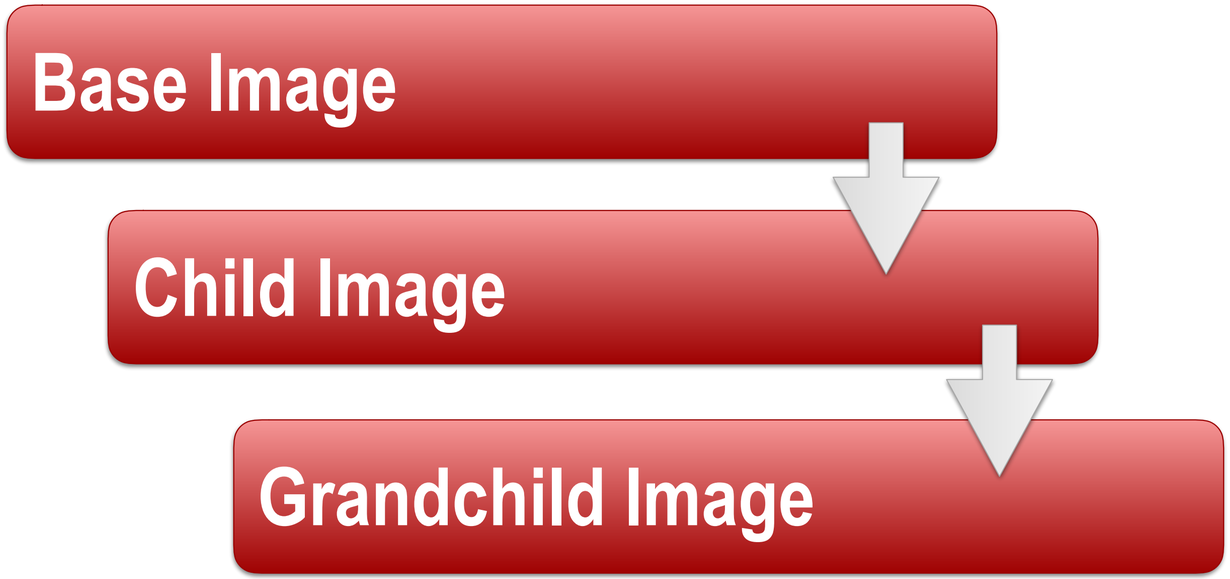

DOCKER REGISTRY

- Git Repo Semantics
  - Pull
  - Push
  - Commit
- Hierarchy

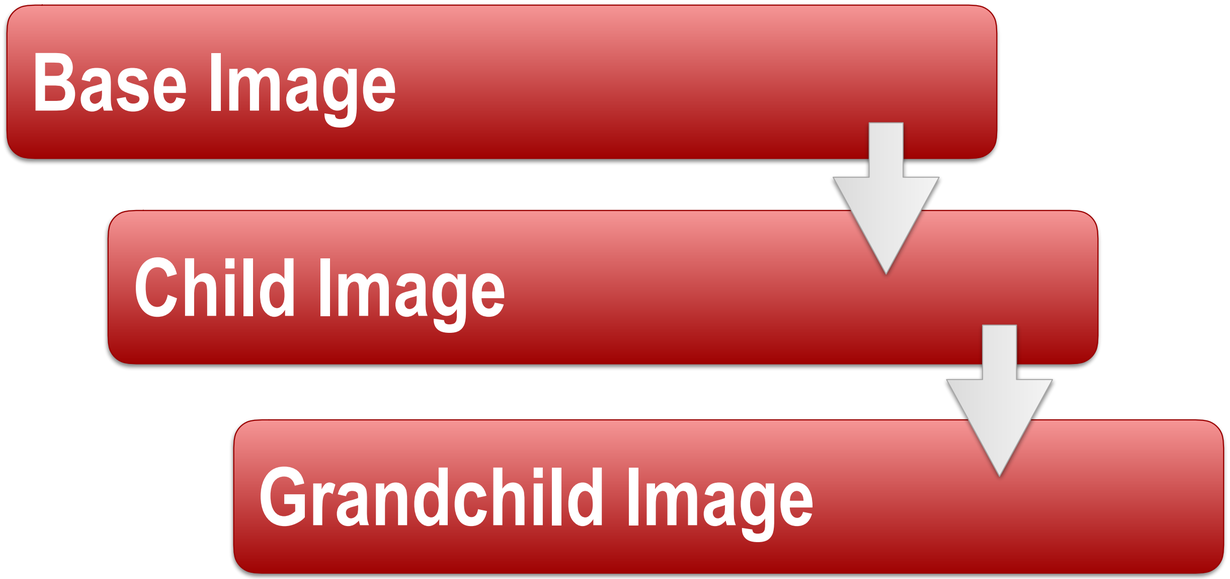

CONTAINER

- Combines several things
  - Linux Cgroups
  - Kernel Namespaces
  - Docker Image
  - Has a lifecycle

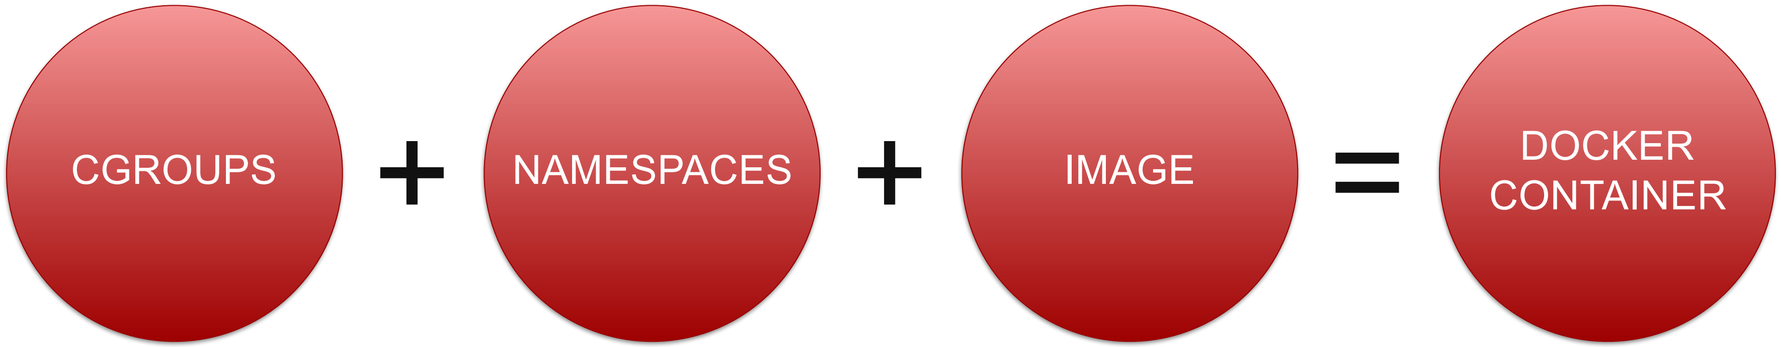

DOCKERFILE

- Like a Makefile (shell script with keywords)
- Extends from a Base Image
- Results in a new Docker Image
- Imperative, not Declarative

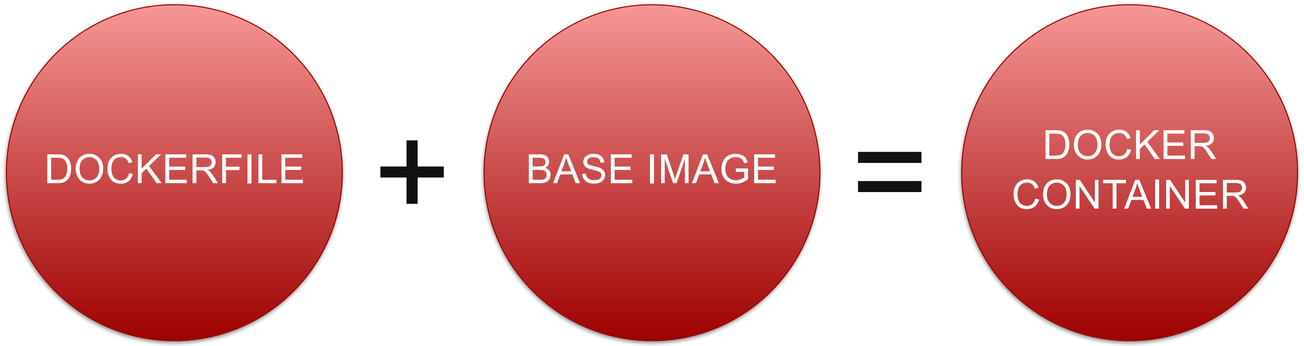

DOCKERFILE EXAMPLE

```
FROM centos:centos6
MAINTAINER Adrian Otto <aotto@aotto.com>
RUN yum -y install httpd
EXPOSE 80
ADD start.sh /start.sh
CMD /start.sh
```

```
$ docker build -t webserver .
```

DOCKERFILE EXAMPLE

```
FROM webserver
MAINTAINER Adrian Otto <aotto@aotto.com>
RUN yum -y install mysql-server php
EXPOSE 80
ADD start.sh /start.sh
CMD /start.sh
```

```
$ docker build -t lampstack .
```

THE DIFFERENCE

1. EFFICIENCY
2. PERFORMANCE
3. SECURITY

THE DIFFERENCE

1 EFFICIENCY

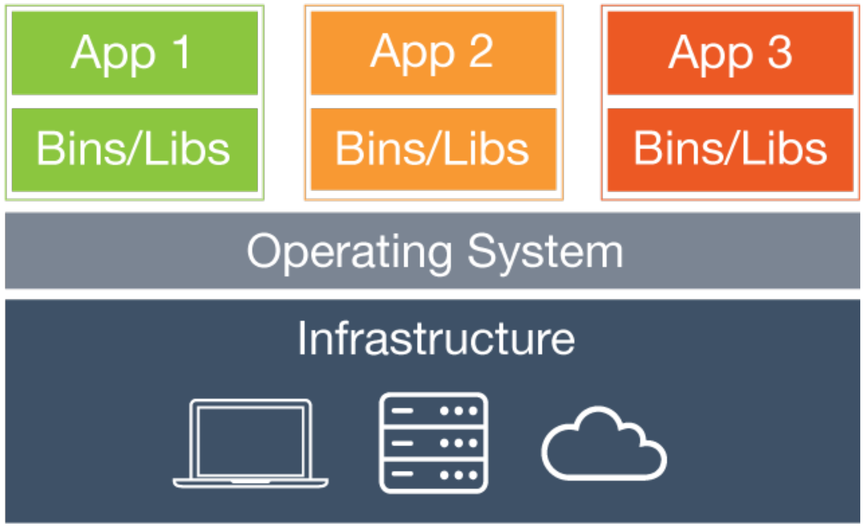

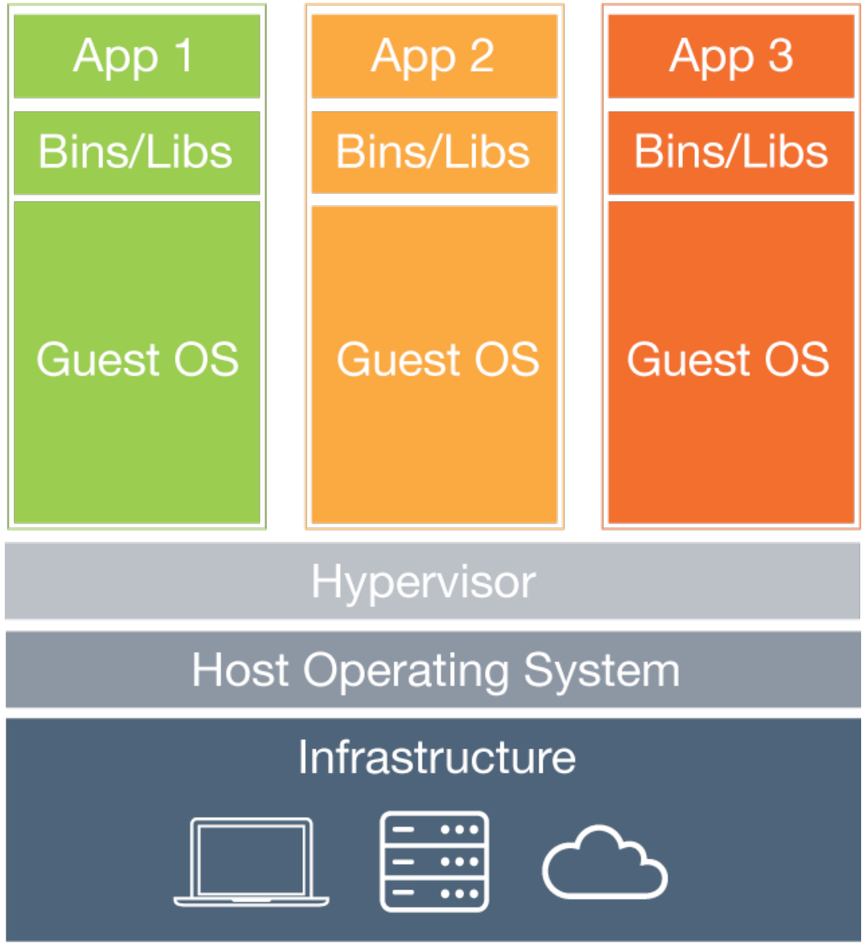

THE DIFFERENCE

2 PERFORMANCE

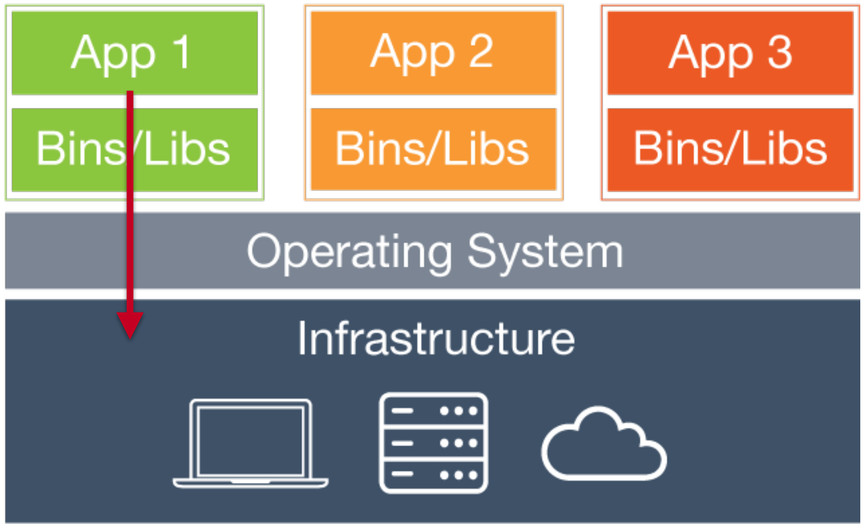

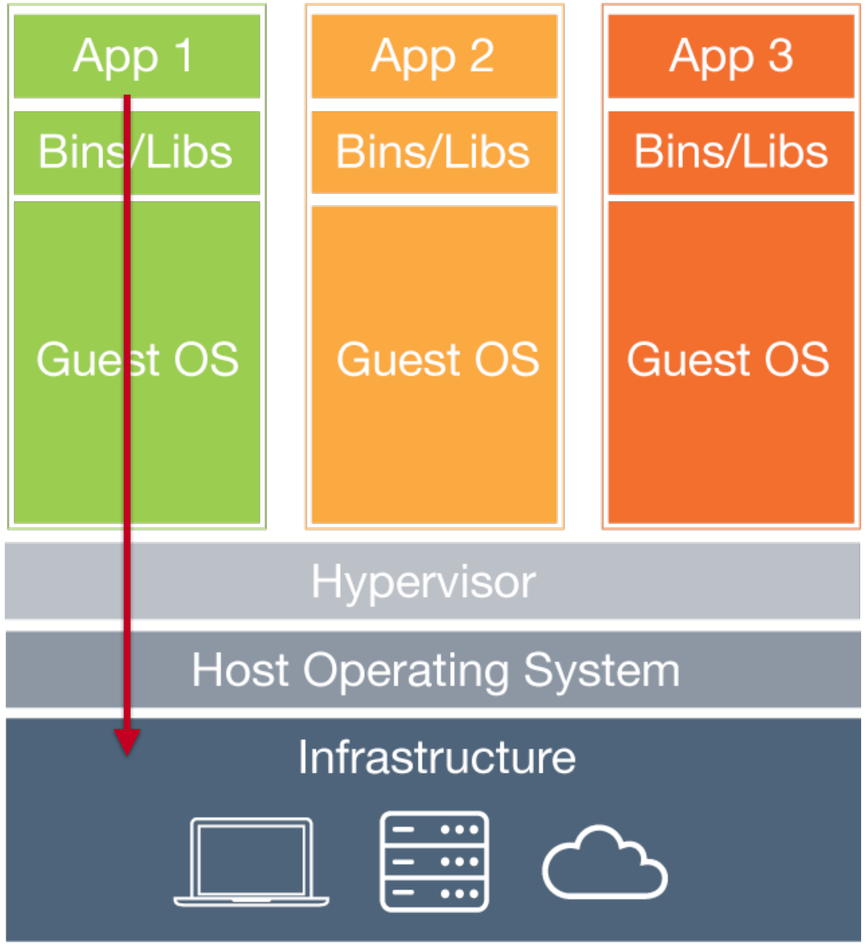

THE DIFFERENCE

3 SECURITY

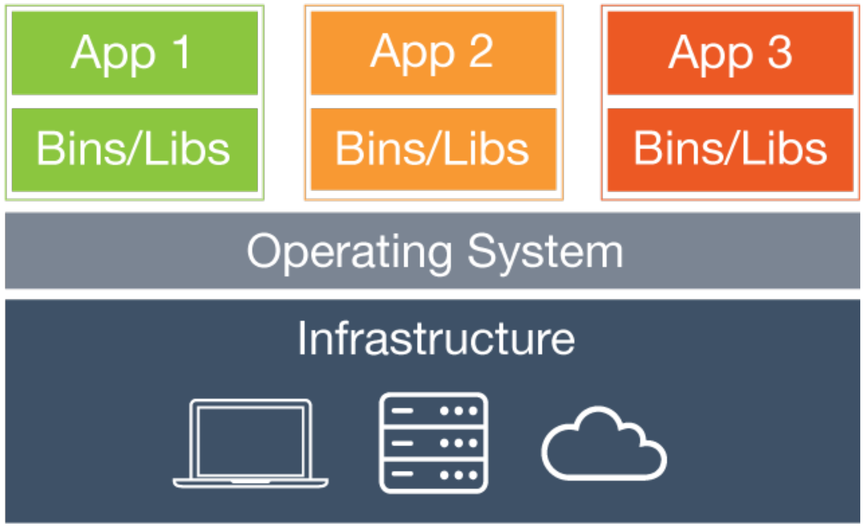

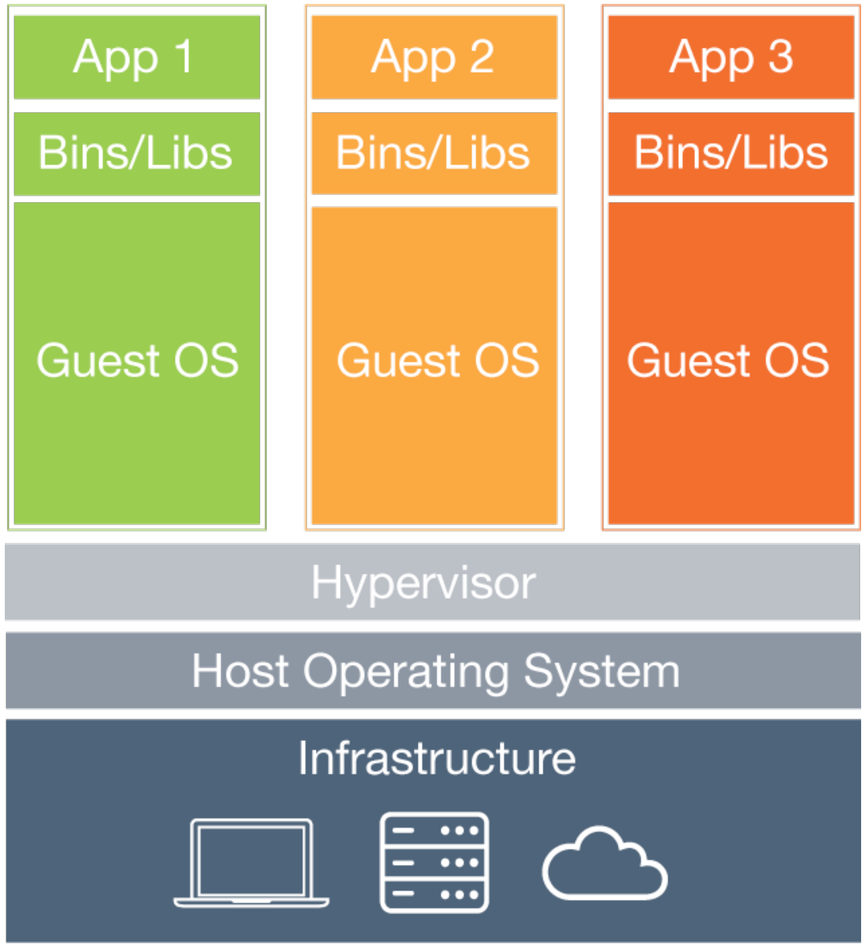

CASTILLO DE SAN MARCOS

VIRTUALIZATION MAPPINGS

Physical	Virtual
System	Partition
Logical Processor	Virtual Processor
Advanced Programmable Interrupt Controller (APIC)	Virtual APIC + Synthetic Interrupt Controller (SynIC)
Physical Address = System mPhysical Address (SPA)	Guest Physical Address (GPA)

LINUX SYSCALL INTERFACE

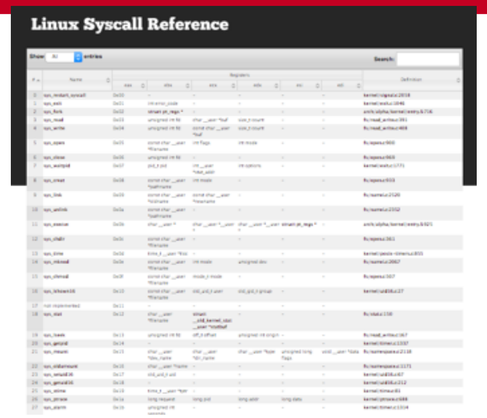

397 CALLS IN KERNEL 3.19

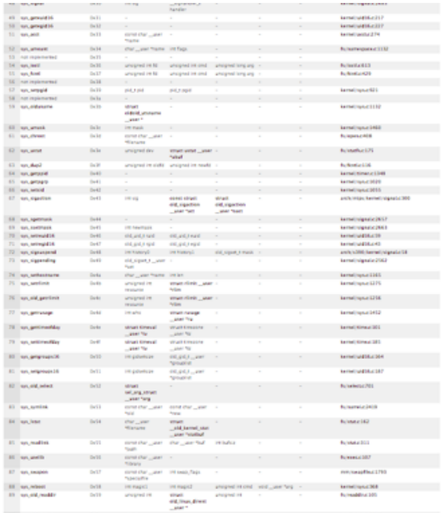

THE DIFFERENCE

3 SECURITY

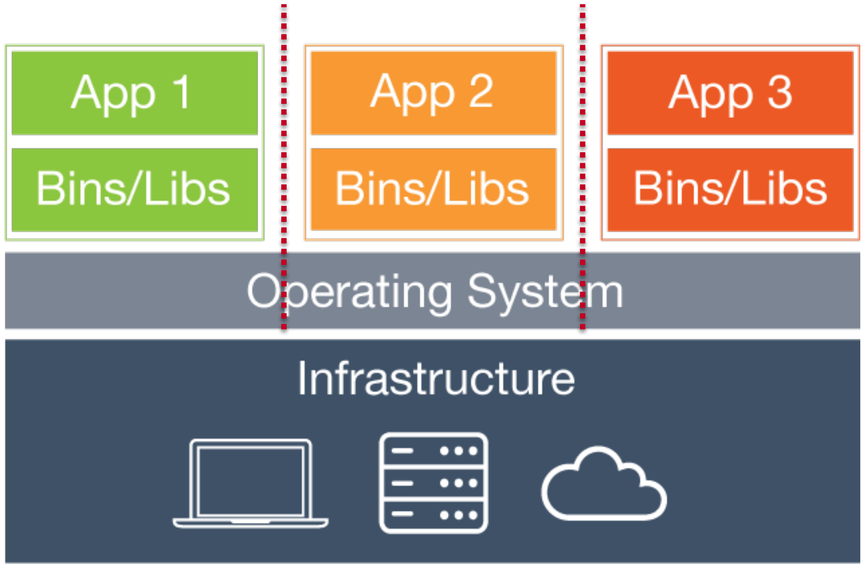

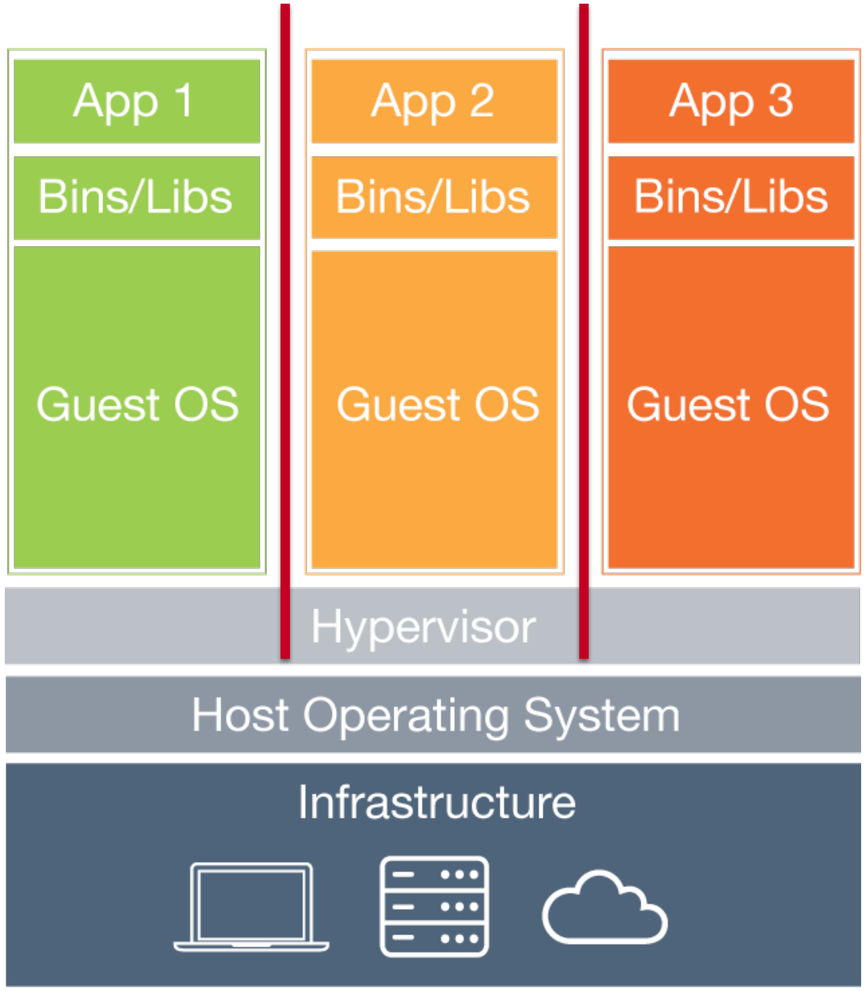

CONTAINTER ISOLATION TECHNIQUES

- SELinux / App
- Secure Compu
- Container Nes
- Docker Auth P
- User Namespa
- Encrypted File
- Address Space
- Hardware Sec

THE DIFFERENCE

1. EFFICIENCY
2. PERFORMANCE
3. SECURITY

rackspace | YOUR CLOUDS. OUR EXPERTISE.

THANK YOU

ONE FANATICAL PLACE | SAN ANTONIO, TX 78218

US SALES: 1-800-961-2888 | US SUPPORT: 1-800-961-4454 | WWW.RACKSPACE.COM


rackspace®
YOUR CLOUDS.
OUR EXPERTISE.