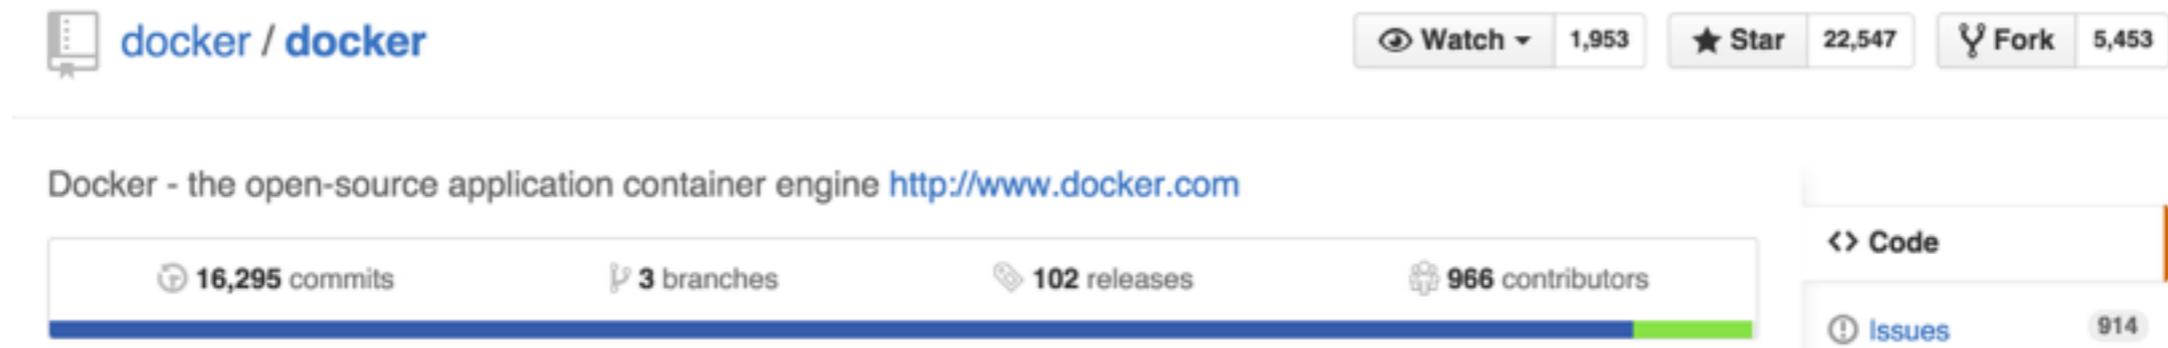


# Docker Recipes for Java Developers

Arun Gupta, @arungupta

# What is Docker?

- Open source project and company



The screenshot shows the GitHub repository page for Docker. At the top left, it says "docker / docker". To the right, there are buttons for "Watch" (1,953), "Star" (22,547), and "Fork" (5,453). Below this, the repository description is "Docker - the open-source application container engine" with a link to "http://www.docker.com". A progress bar shows statistics: 16,295 commits, 3 branches, 102 releases, and 966 contributors. On the right side, there are buttons for "Code" and "Issues" (914).

- Used to create containers for software applications
- Package Once Deploy Anywhere (PODA)

## Docker Contributors by employers/hackers

 [gistfile1.txt](#)

```
1 Top changeset contributors by employer
2 (Unknown) 4520 (47.9%)
3 "Docker" 3821 (40.5%)
4 "Red Hat" 685 (7.3%)
5 "IBM" 232 (2.5%)
6 "Google" 119 (1.3%)
7 "Cisco" 49 (0.5%)
8 "Amadeus" 4 (0.0%)
9 "VMWare" 2 (0.0%)
10 "CoreOS" 1 (0.0%)
11
```

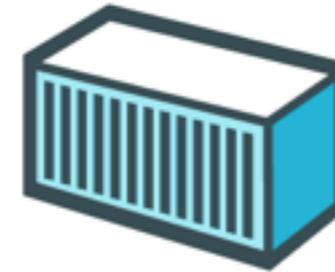
<https://gist.github.com/arun-gupta/7d5a373099ff831d7213>

<http://www.infoworld.com/article/2925484/application-virtualization/look-whos-helping-build-docker-besides-docker-itself.html>



## Build

Develop an app using Docker containers with any language and any toolchain.



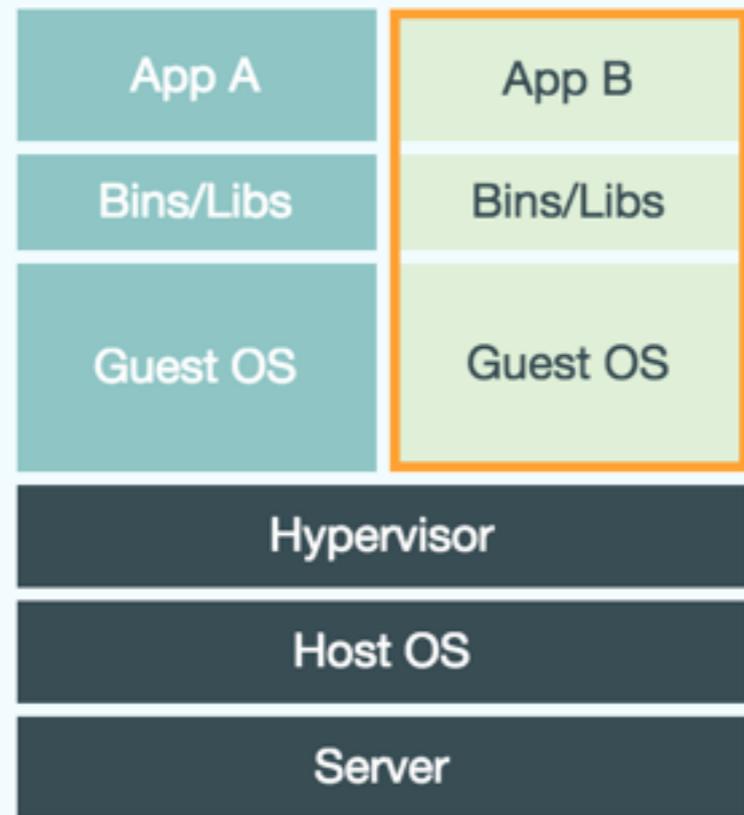
## Ship

Ship the “Dockerized” app and dependencies anywhere - to QA, teammates, or the cloud - without breaking anything.



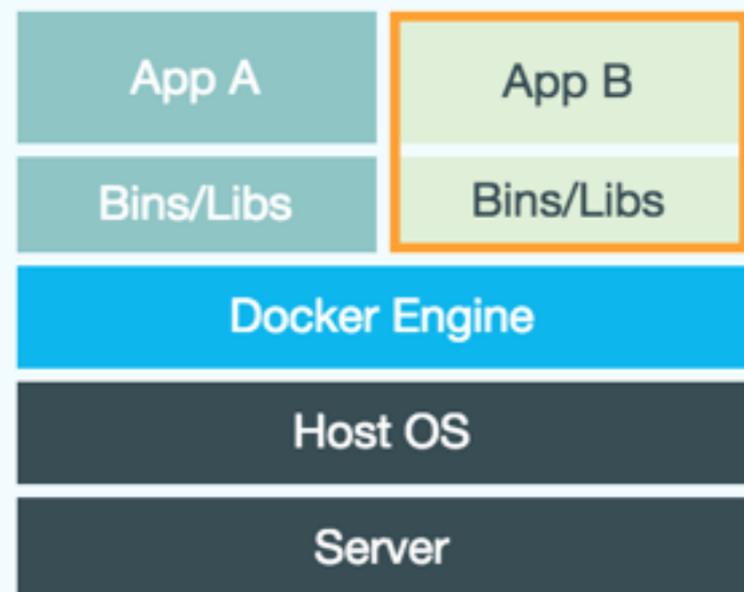
## Run

Scale to 1000s of nodes, move between data centers and clouds, update with zero downtime and more.



## Virtual Machines

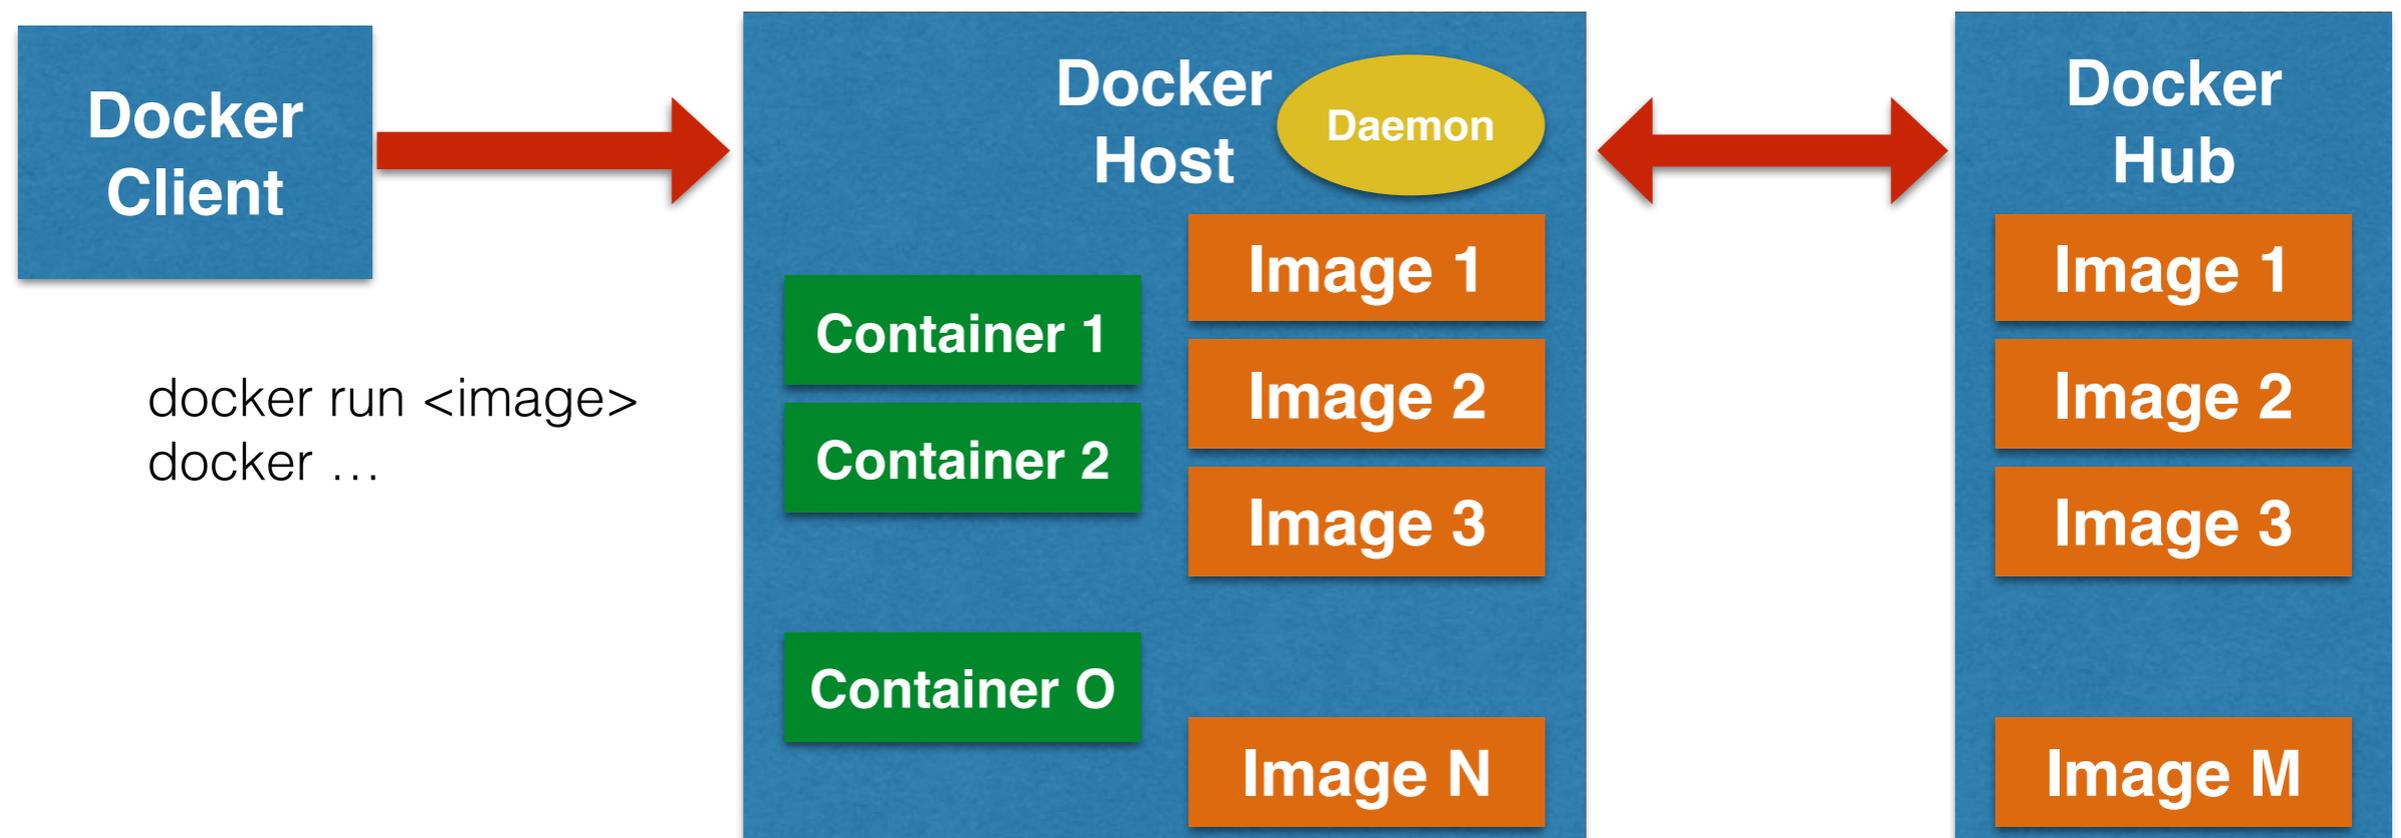
Each virtualized application includes not only the application - which may be only 10s of MB - and the necessary binaries and libraries, but also an entire guest operating system - which may weigh 10s of GB.



## Docker

The Docker Engine container comprises just the application and its dependencies. It runs as an isolated process in userspace on the host operating system, sharing the kernel with other containers. Thus, it enjoys the resource isolation and allocation benefits of VMs but is much more portable and efficient.

# Docker Workflow

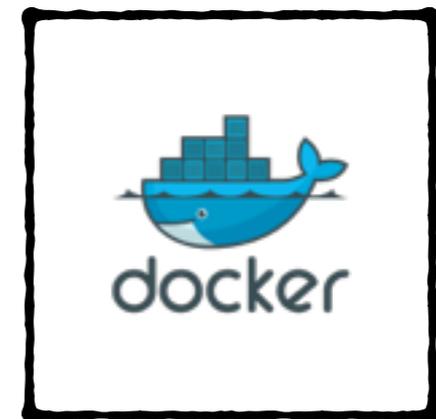


# Docker Machine

- Create Docker Host on computer or cloud provider

```
docker-machine create --driver=virtualbox  
myhost
```

- Configure Docker client to talk to host
- Create and pull images
- Start, stop, restart containers
- Upgrade Docker
- Not recommended for production yet



# Docker Compose

- Defining and running multi-container applications
- Configuration defined in a single file
- Great for dev, staging, and CI
- Not recommended for production yet

# docker-compose.yml

**mysql**

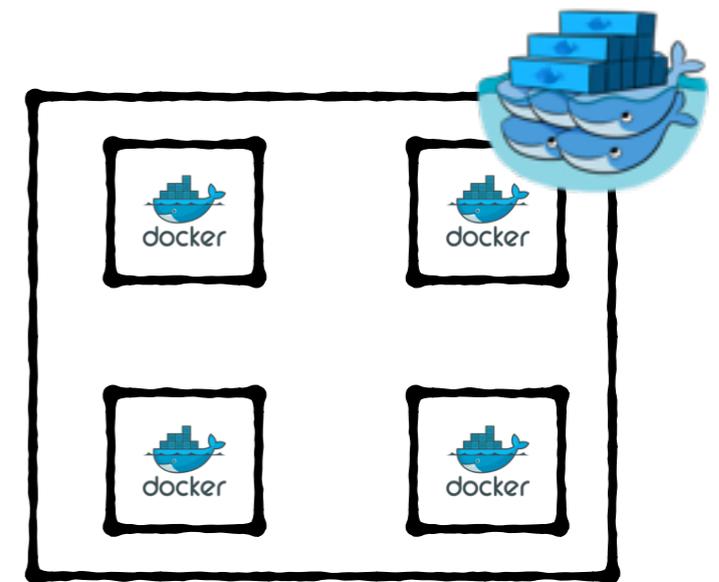
```
image: mysql
environment:
  MYSQL_DATABASE: sample
  MYSQL_USER: mysql
  MYSQL_PASSWORD: mysql
  MYSQL_ROOT_PASSWORD: supersecret
```

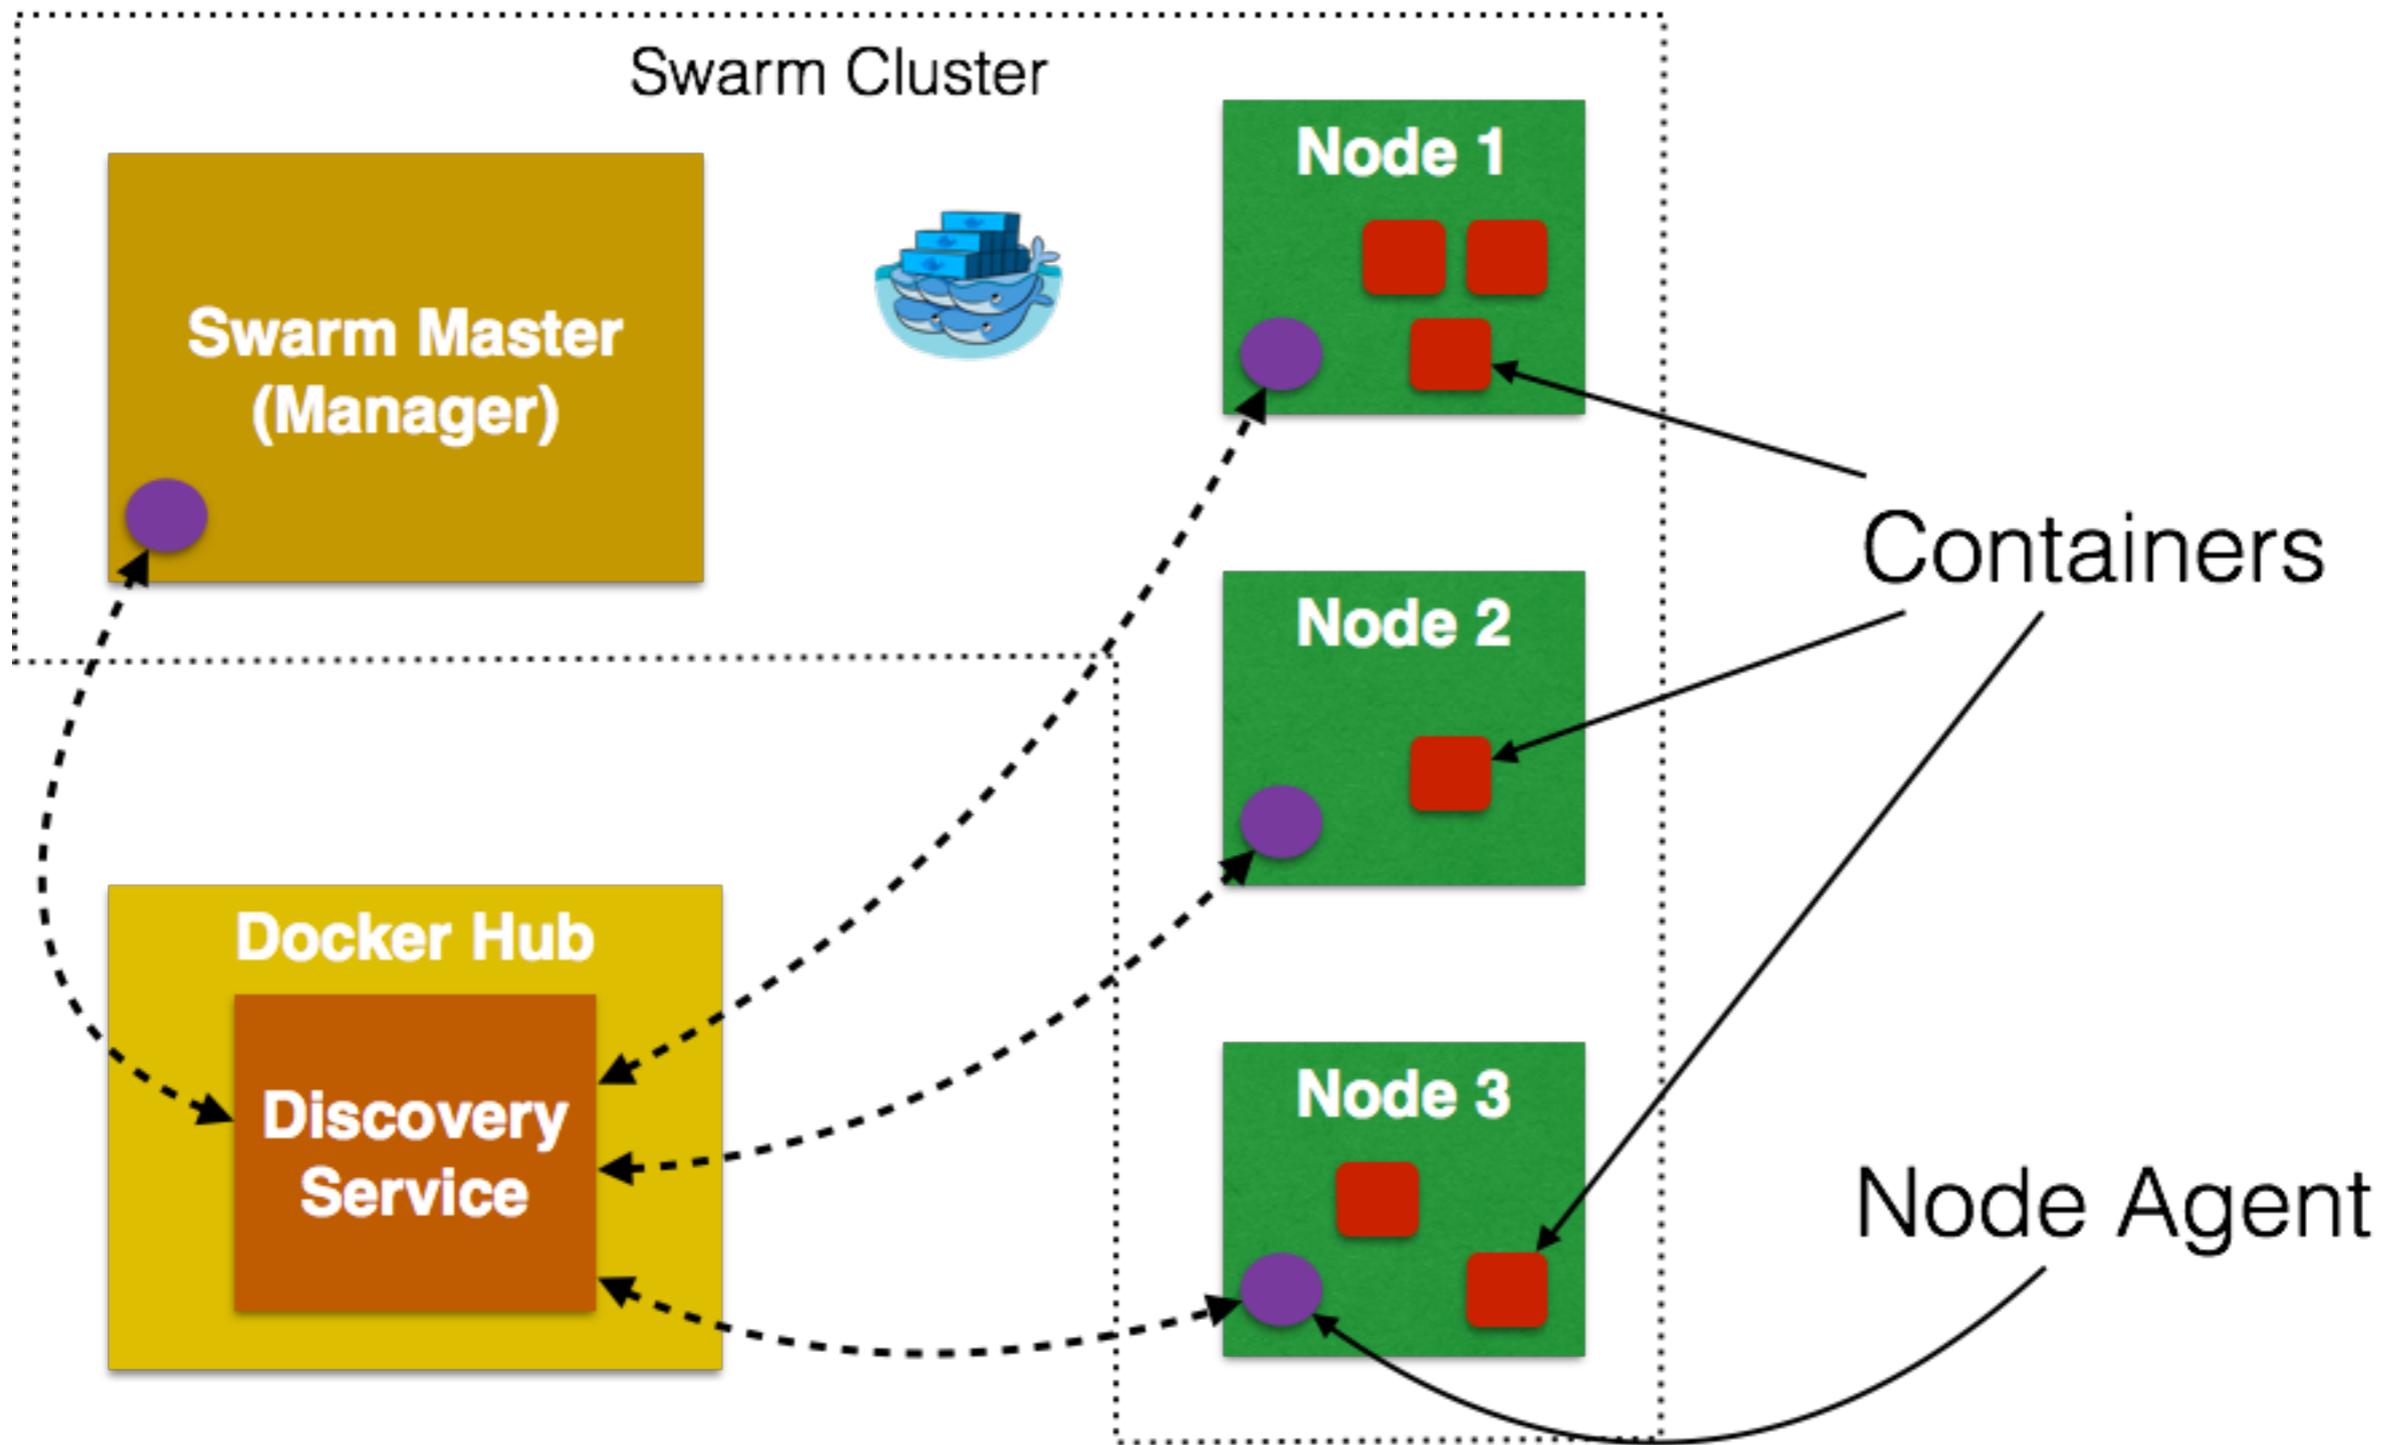
**mywildfly**

```
image: arungupta/wildfly-mysql-javaee7
links:
  - mysql:db
```

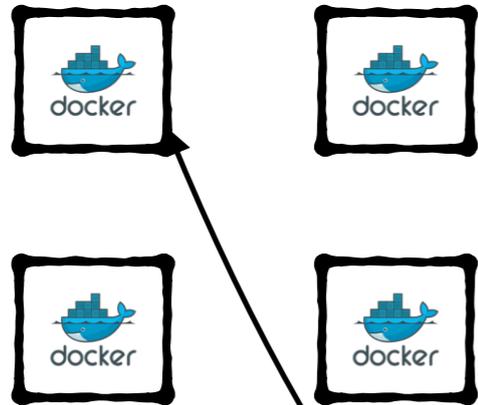
# Docker Swarm

- Native clustering for Docker
- Provides a unified interface to a pool of Docker hosts
- Fully integrated with Machine
- Serves the standard Docker API
- Partially integrated with Compose
- Not recommended for production yet

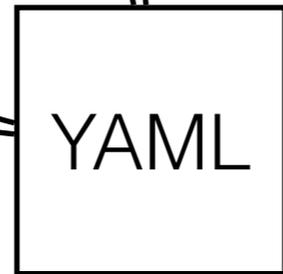
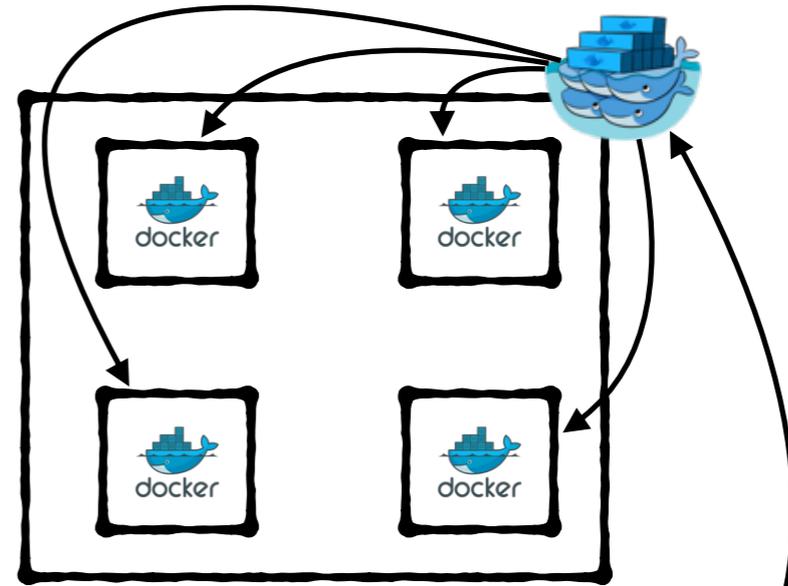




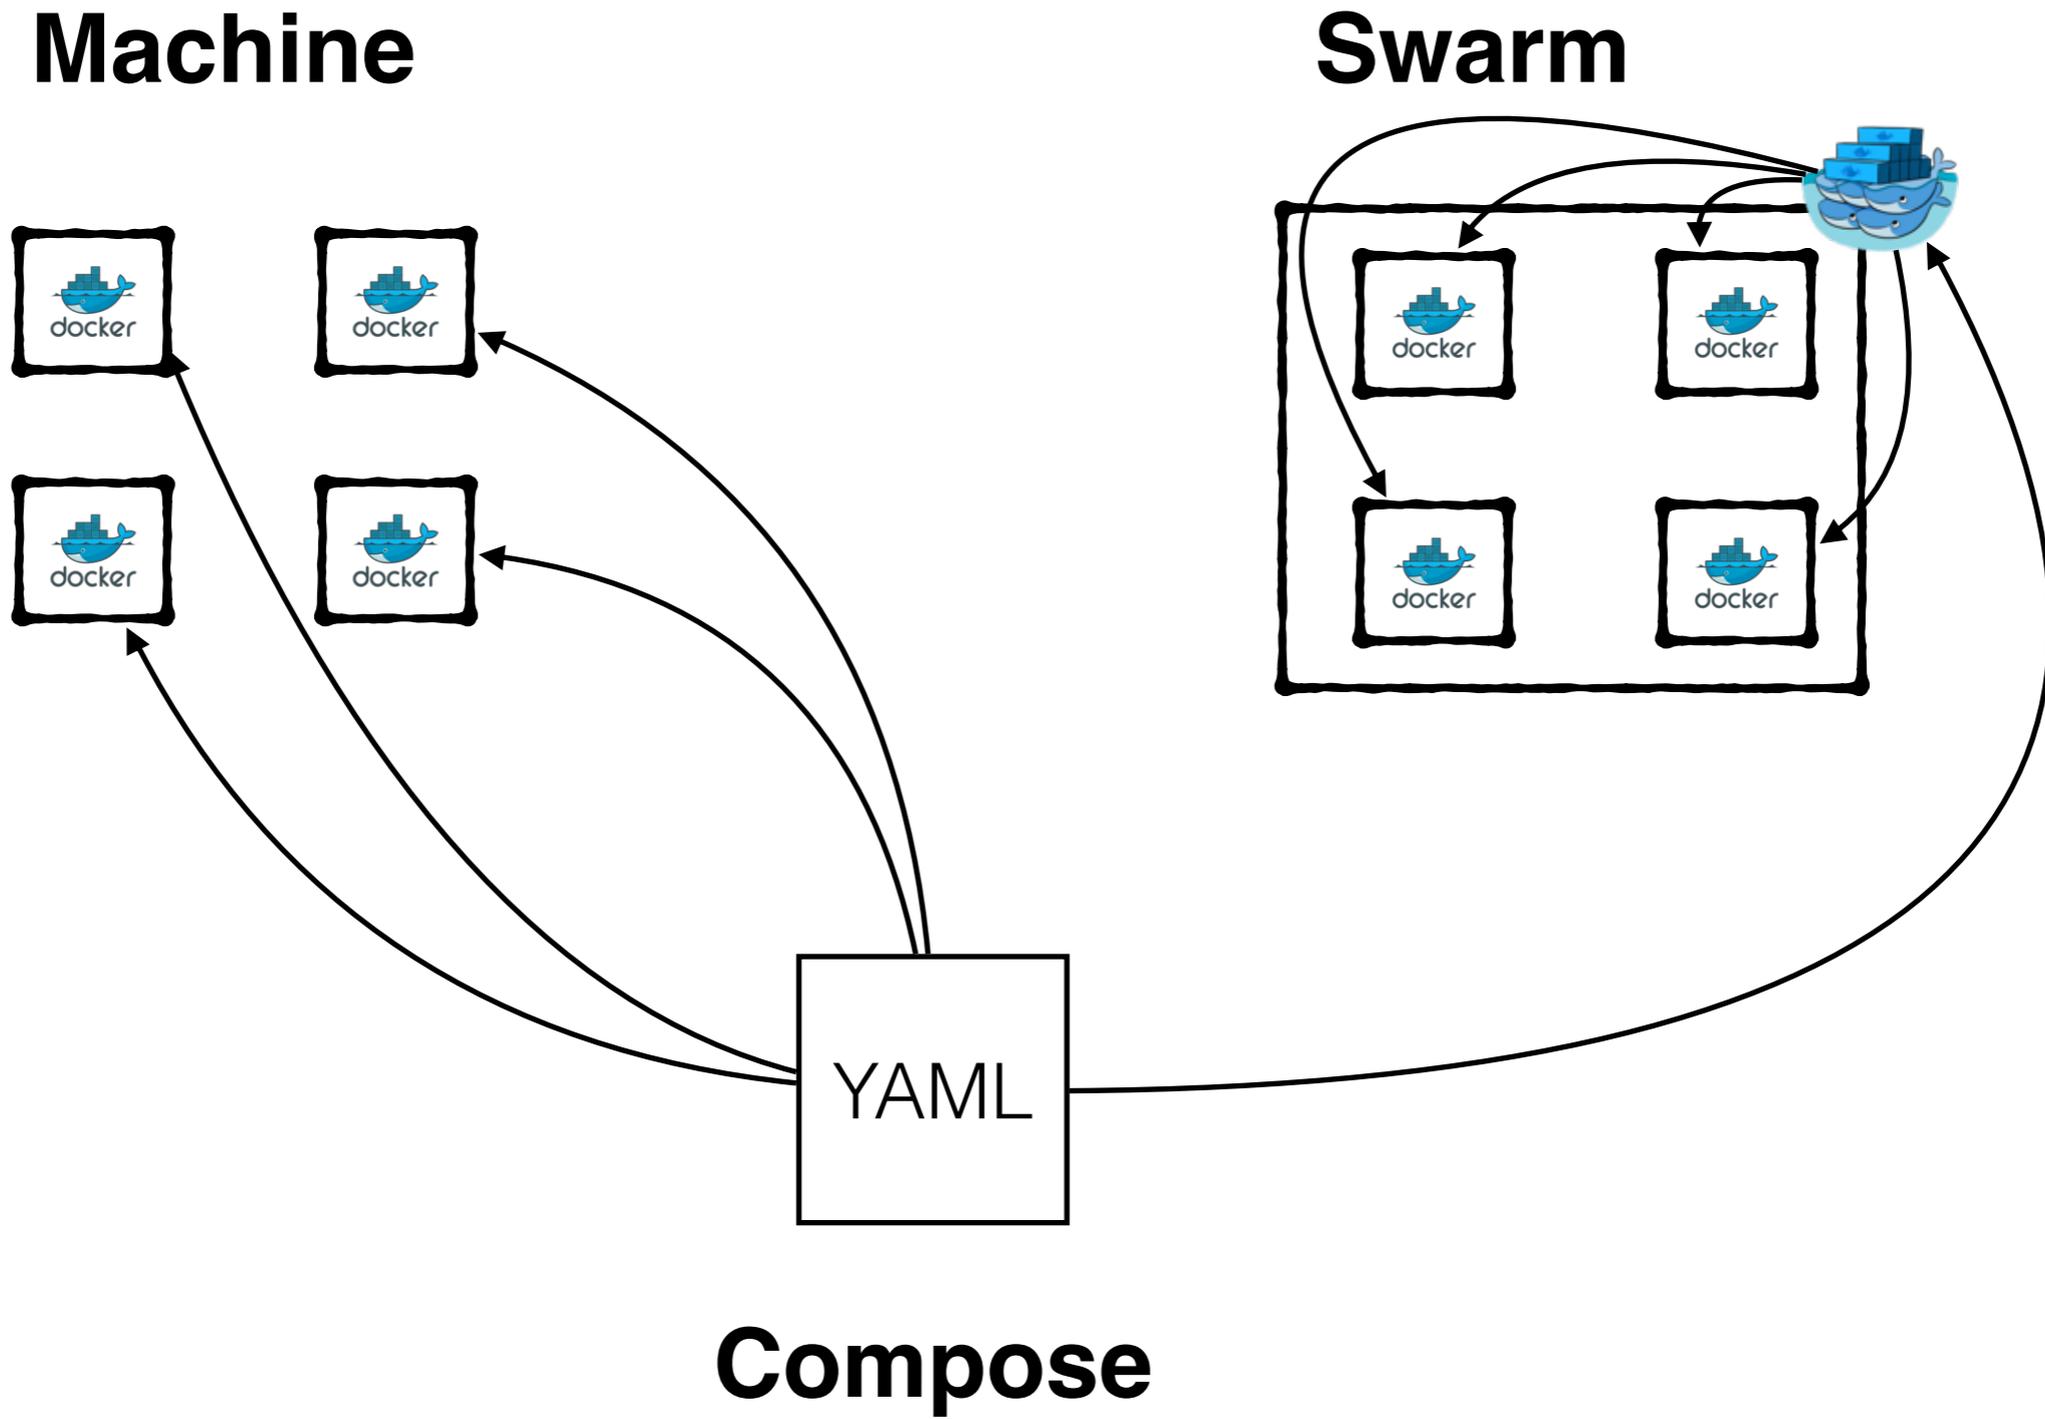
# Machine



# Swarm



# Compose



# Advantages of Containers

- Immutability
- Reproducibility
- Isolation
- Faster deployments
- Portability - “it works on my machine”
- Snapshotting
- Security sandbox
- Limit resource usage
- Simplified dependency
- Sharing

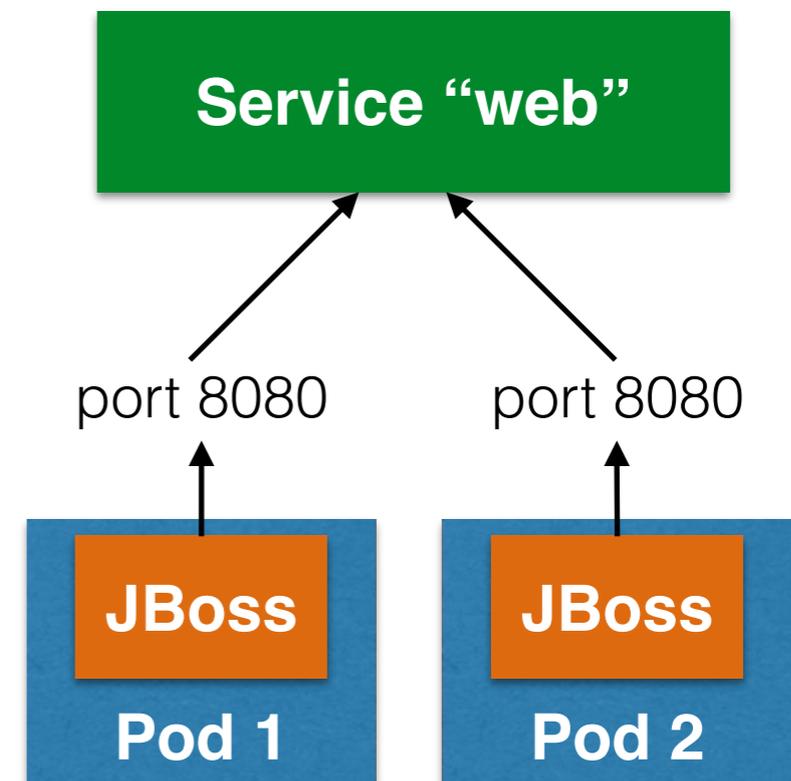
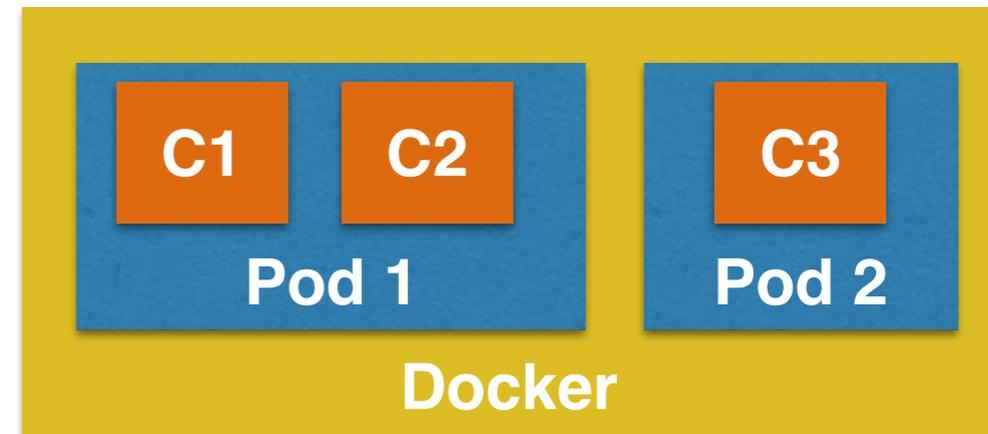
# Kubernetes

- Open source orchestration system for Docker containers
- Provide declarative primitives for the “desired state”
  - Self-healing
  - Auto-restarting
  - Schedule across hosts
  - Replicating

# Concepts



- **Pods:** collocated group of Docker containers that share an IP and storage volume
- **Service:** Single, stable name for a set of pods, also acts as LB
- **Replication Controller:** manages the lifecycle of pods and ensures specified number are running
- **Label:** used to organize and select group of objects



## Collect statistics about who wrote kubernetes

gitdm (git data mine) is a tool written by Jonathan Corbet at LWN which he uses to do his 'who wrote the kernel' articles. I spent a couple of minutes to run it against kube. Some interesting results...

...

### Developers with the most changesets

Brendan Burns	643 (10.3%)
Daniel Smith	485 (7.8%)
Clayton Coleman	453 (7.3%)
Tim Hockin	312 (5.0%)
derekwaynecarr	209 (3.4%)

### Developers with the most changed lines

Brendan Burns	455888 (18.0%)
Nan Monnand Deng	379610 (15.0%)
Yifan Gu	219191 (8.7%)
Patrick	169265 (6.7%)
nikhiljindal	124915 (4.9%)

### Top changeset contributors by employer

"Google"	3562 (58.3%)
"Red Hat"	1252 (20.5%)
(Unknown)	899 (14.7%)
"CoreOS"	144 (2.4%)
"FathomDB"	142 (2.3%)

