

CernVM-FS for
Docker image distribution
in Cloud Foundry

George Lestaris, Pivotal

Credentials

- **2012-2014** DN=/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=glestari/CN=739XXX/CN=George Lestaris
 - CernVM & Geant4 & cern.ch/alice20
- **2015** - `int pivot_root(const char *new_root, const char *put_old);`

Pivotal

- Founded in 2013
- Educating the industry in
 - Agile
 - eXtreme Programming
 - Pair programming

Pivotal

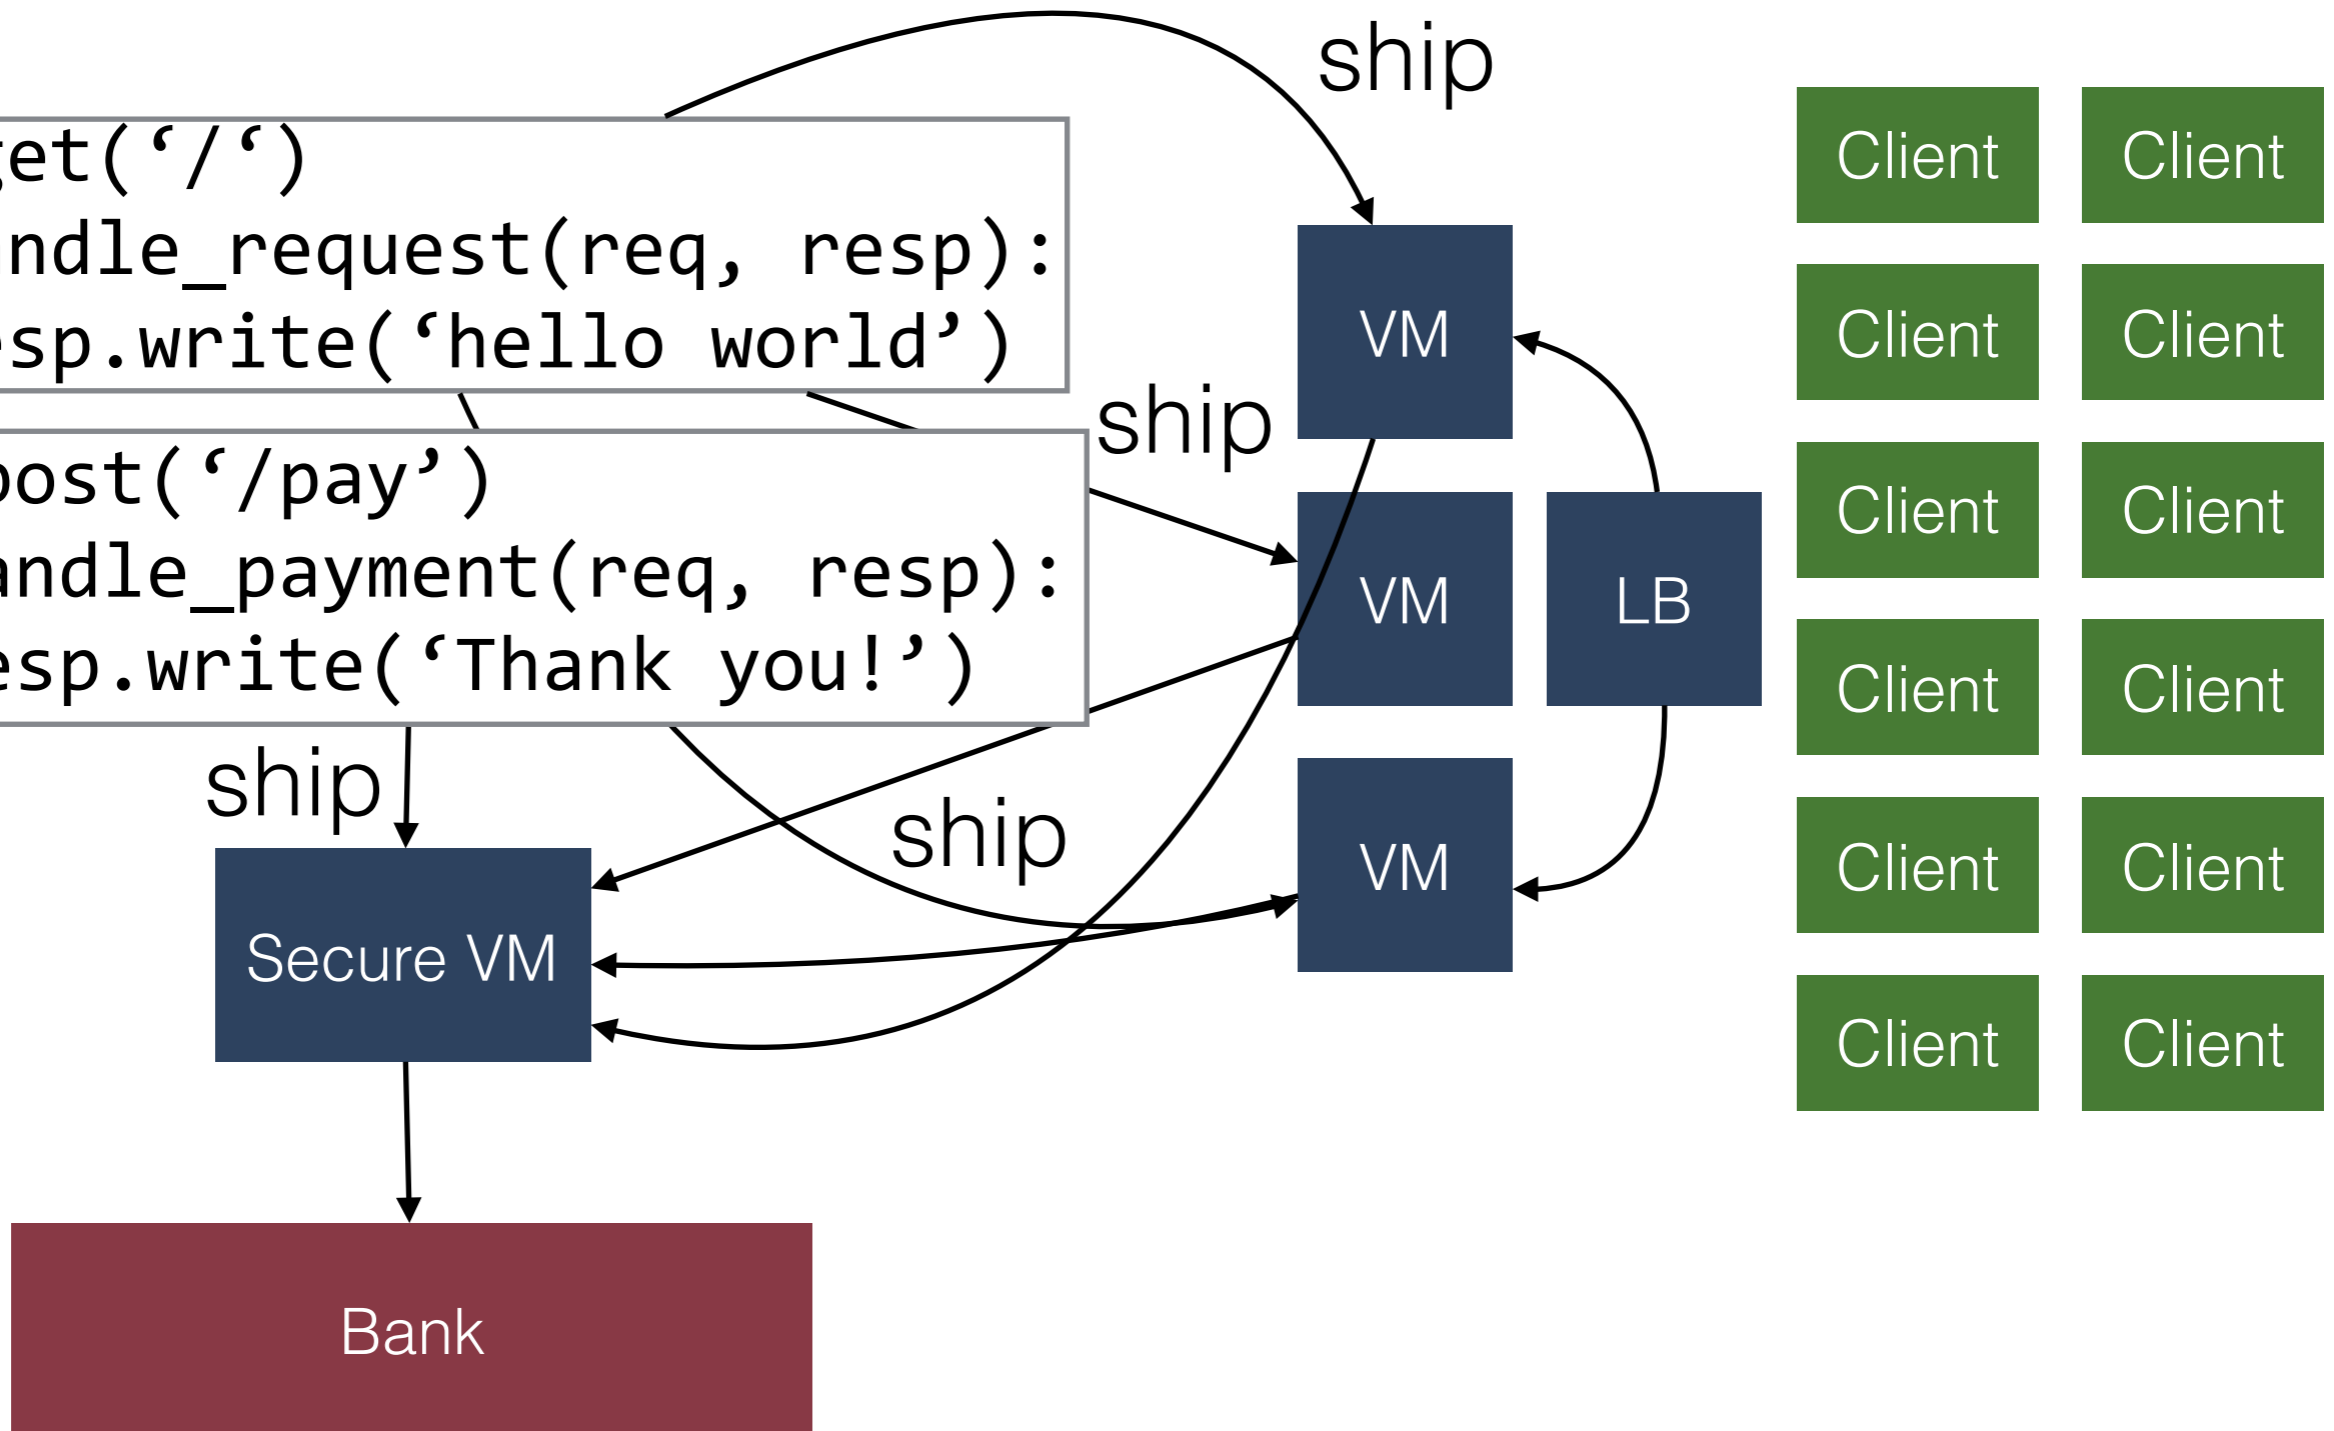


the-story-of-every-
startup.io

the-story-of-every-startup.io

```
@app.get('/')  
def handle_request(req, resp):  
    resp.write('hello world')
```

```
@app.post('/pay')  
def handle_payment(req, resp):  
    resp.write('Thank you!')
```



the-story-of-every-startup.io

- I could go on for ever
- But in this short journey the teams learn:
 - operations are hard
 - infrastructure is unreliable
 - deploying fast is crucial: time to market
 - spending more time coding
 - reproducibility is important

Enterprise software

- [the-story-of-every-startup.io](#) happens a lot in enterprises
 - bureaucracy, slow and complex processes, large organizations

DOCKER: DELIVERING AGILITY ONE CONTAINER AT A TIME

🕒 JULY 12, 2013 👤 MIKE KAVIS 💬 2 COMMENTS

<https://www.virtualizationpractice.com/docker-delivering-agility-one-container-at-a-time-22328>

Cloud Foundry in enterprise

NOV 30, 2011 @ 01:58 PM **52,966** VIEWS

Now Every Company Is A Software Company

<http://www.forbes.com/sites/teconomy/2011/11/30/now-every-company-is-a-software-company>



Nomad



kubernetes



Apache
MESOS™



but which one?



kubernetes



Apache
MESOS™

Startups

- Quick feedback
- Time-to-market
- Less time deploying
- More time coding

Enterprise

- Support
- Common strategy within the organization
- Reduce or remove bottlenecks
- Availability



CLOUD FOUNDRY

Cloud Foundry

- Platform-as-a-Service
- Originally built in VMware in 2011
 - was using containers (!)
 - quickly got open sourced
- Moved to Pivotal in 2013
- End-to-end system

Cloud Foundry: End-to-end system

- **Scalability:** Deploys your application and scales it to multiple instances seamlessly
- **Orchestration:** multiple workers - orchestration of app instances
- **Load balancing:** distributes load to all the app instances
- **Isolation:** did I mention containers?
- Logs
- Metrics
- Data services

Demo time

Cloud Foundry haiku

Here is my source code
Run it on the cloud for me
I do not care how

Haiku (俳句) (plural haiku) is a very short form of Japanese poetry.

12 factor apps

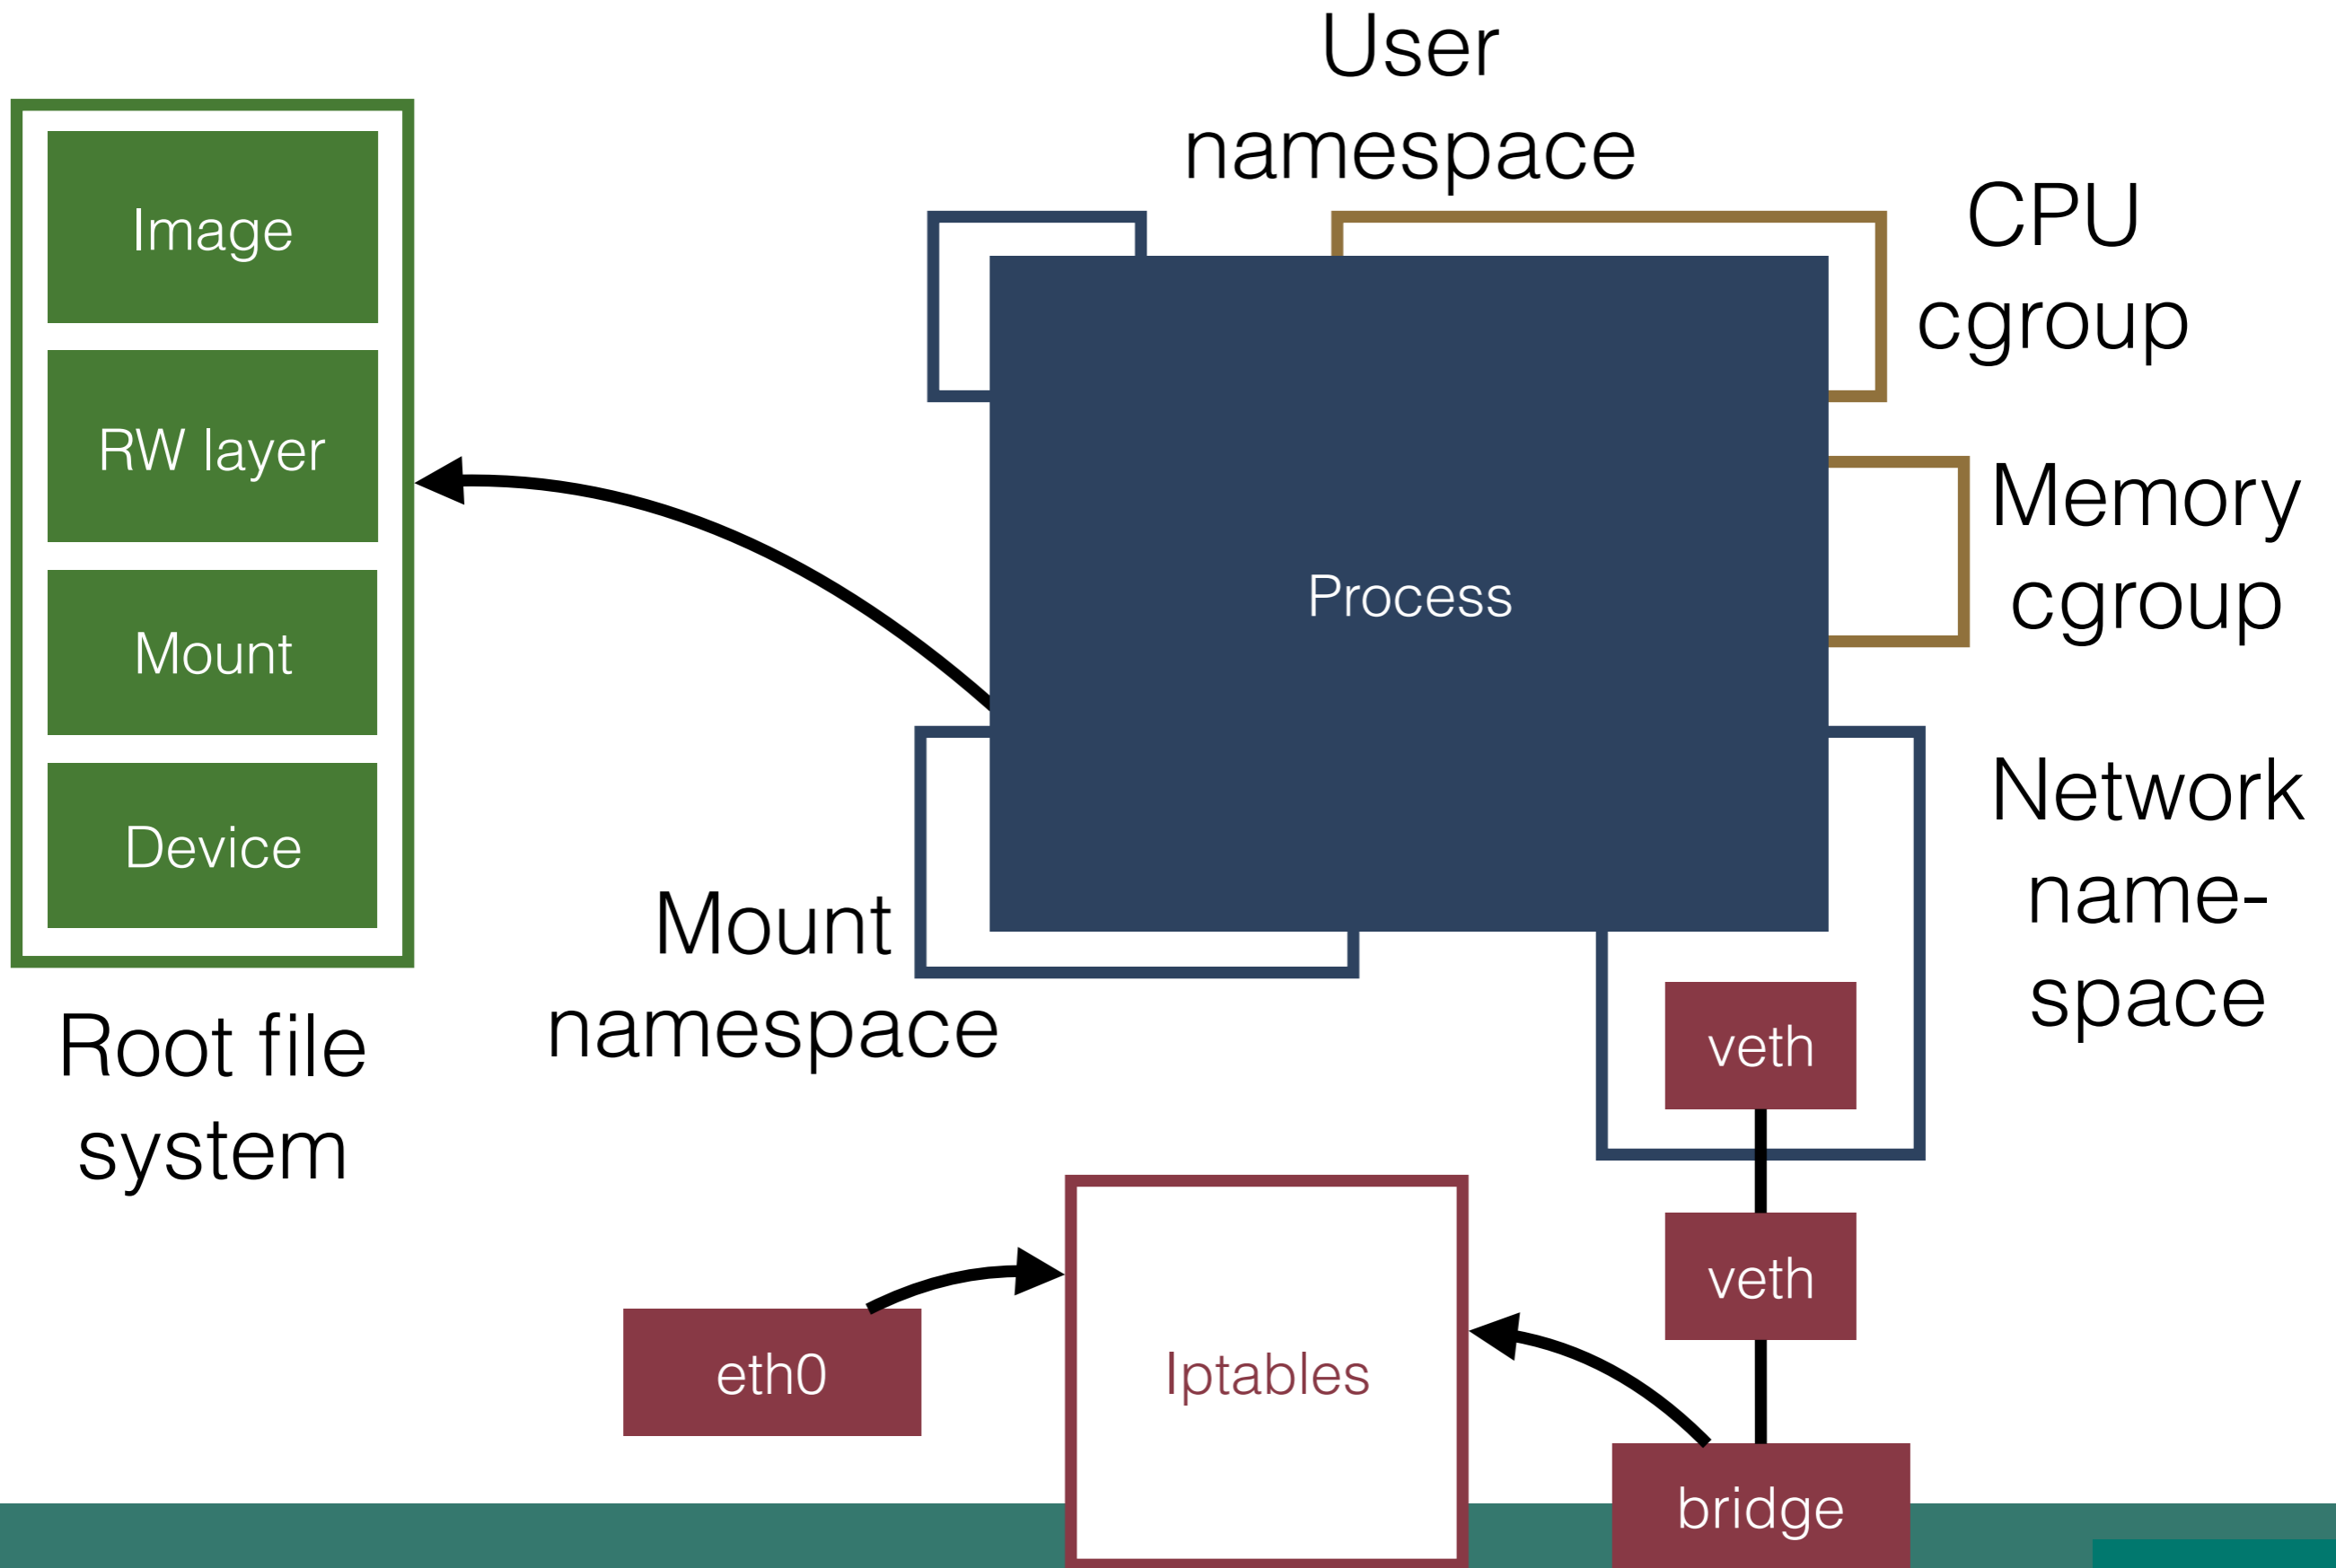
- Containers are disposable
- Configuration via environment
- Export services via port binding
- 8 more factors :)

Containers



What are containers?

1. Docker container
2. Way to isolate processes and share resources
3. Containers that wrap your shipped application
4. Linux perspective - combination of: namespaces, cgroups, capabilities, rlimits, AppArmor, SELinux, iptables, virtual ethernet pairs, network bridges, overlay filesystems, etc



Docker images in Cloud Foundry

- Last year Diego and Garden were shipped in Cloud Foundry
 - Diego: container orchestrator (Kubernetes)
 - Garden: container runtime (Docker)
 - why not using Docker? - talk about that later

Demo

So what?

- No vendor lock-in
- Can containerize your applications and swap between CF, Mesos, Kubernetes etc

Potential use cases of CernVM-FS in CF

- Docker image distribution/caching
- Cloud Foundry root file systems management
- Provide build packs or stacks

Prototype

- `docker-image-export` exports Docker images in a directory: using it to push Docker images to a stratum-0
- **BOSH:** multi-cloud infrastructure
 - using it to deploy CF
 - built a `cvmfs-client-boshrelease`
- `garden-runc` the new container runtime of CF
 - makes runC containers

Demo

What is the right
interface?

Open container initiative

- Linux foundation collaborative project
- Aims to provide standards around container formats and runtime
- Founded in 2015
- Open technical governance



runC

- Reference implementation of the OCI specification
- Originally it was donated by Docker (libcontainer)
- In use by Docker (v1.11)
- Will soon be used by Cloud Foundry




image-spec

- OCI's attempt to standardize the container **image format**
- Originally, image format was out of scope
 - the OCI runtime specification expects simply a directory
- Currently
 - **Layer based**
 - Distribution is out of scope

Open governance vs open source

- Open governance models (Linux foundation):
 - OCI: runC
 - Cloud native computing foundation: Kubernetes
 - Cloud foundry foundation
- Open source software:
 - Docker (uses runC)
 - OpenShift (uses Docker and Kubernetes)

Cloud Foundry foundation

- Part of the **Linux foundation**
- Governs the open source CF
- Teams are made of engineers and PMs from **foundation members** (Pivotal, IBM, EMC, SAP,...)
- Members usually provide their own CF products: Pivotal Cloud Foundry, IBM Bluemix, SAP HANA