



Kris Buytaert
@krisbuytaert

~~Works on My Machine~~ in
my container,
Ops Problem now

Kris Buytaert

- I used to be a Dev,
- Then Became an Op
- CTO and Open Source Consultant @inuits.eu
- Everything is a freaking DNS Problem
- Evangelizing devops
- Organiser of #devopsdays, #cfgmngmtcamp, #loadays, #deliveryconf ?

Enterprise Software Deployment

- “Put this Code Live, here's a tarball” NOW!
- What dependencies ?
- No machines available ?
- What database ?
- Security ?
- High Availability ?
- Scalability ?
- My computer can't install this ?

#devoops in 2009



C(L)AMS

- Culture
- (Lean)
- Automation
- Measurement
- Sharing



Damon Edwards and John Willis

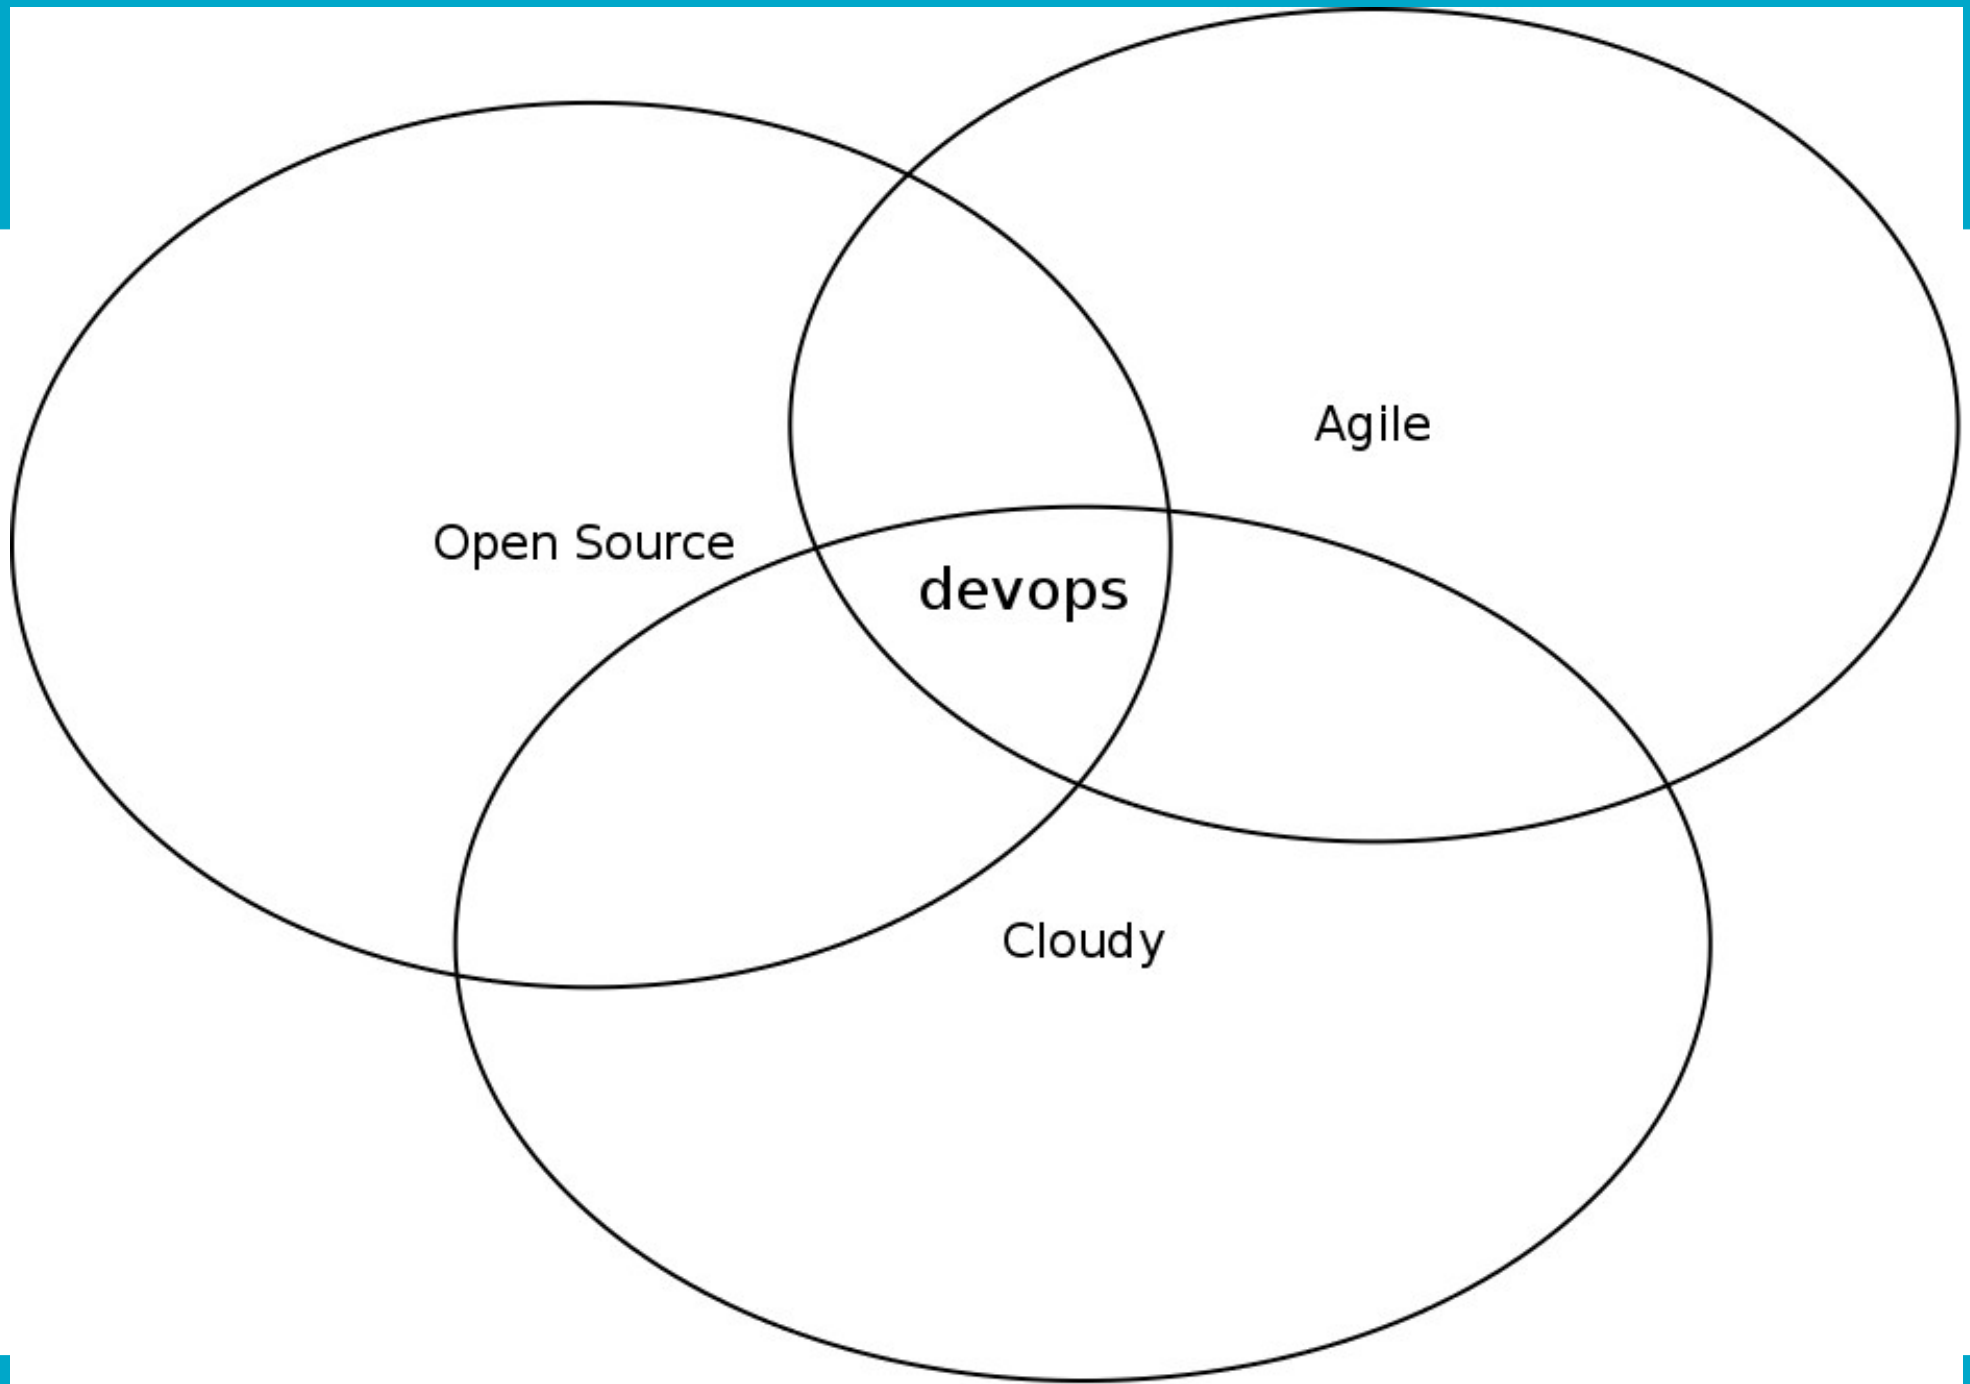
Gene Kim



Dev/Ops

A global movement to improve the quality of software delivery leveraging Open Source experience, started in Gent in 2009





Some History Lessons

openMosix (+/- 2001-2005)

- 2.4 Kernel patch,
- Loadbalance multiple similar processes
- Both short (compilations) and longlived (calculations) processes
- Parts of processes migrated to other nodes
- OMFS allowed identical access to data from all nodes

openMosix Limitations

- shared memory
 - Limited type of applications could migrate
 - Patches available but not successful
- Filesystem access
 - oMFS : unstable => removed

Lessons from openMosix

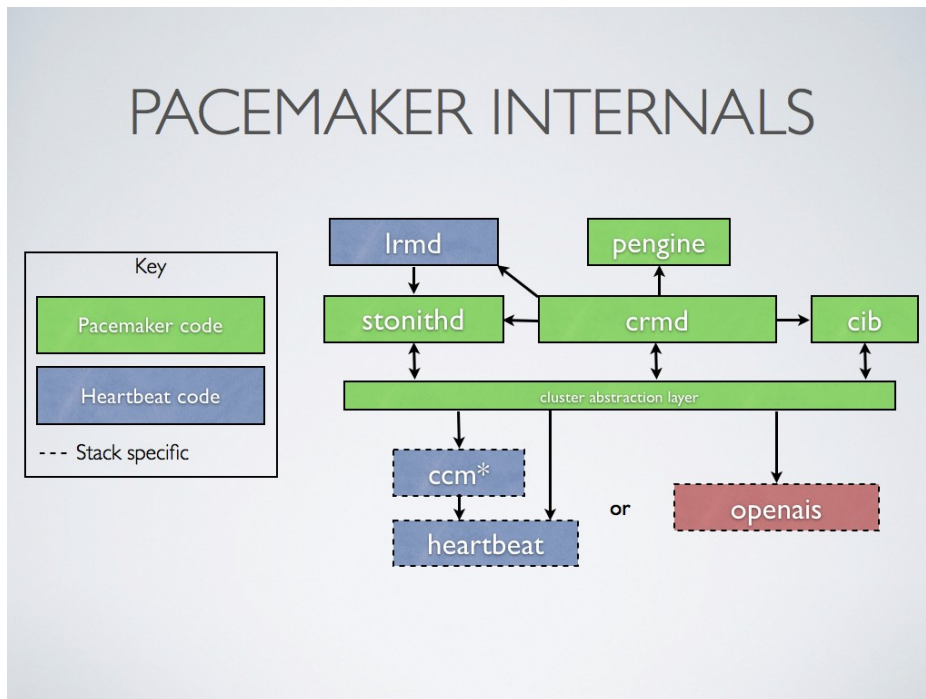
- Applications need to be adapted
- Files need to be accessible (oMFS)

- Huge gap between developers and consumers
=> Limited working use cases

Linux-HA (2005- now)

- Hearbeat, Heartbeat v2, Pacemaker , Corosync
- Define resource (OCF)
- Define constraints, clones, colocation, ...
- Long running services that should not stop

Pacemaker Architecture



- **Stonithd** : The Heartbeat fencing subsystem.
- **Lrmd** : Local Resource Management Daemon. Interacts directly with resource agents (scripts).
- **pengine** Policy Engine. Computes the next state of the cluster based on the current state and the configuration.
- **cib** Cluster Information Base. Contains definitions of all cluster options, nodes, resources, their relationships to one another and current status. Synchronizes updates to all cluster nodes.
- **crmd** Cluster Resource Management Daemon. Largely a message broker for the PEngine and LRM, it also elects a leader to co-ordinate the activities of the cluster.
- **openais** messaging and membership layer.
- **heartbeat** messaging layer, an alternative to OpenAIS.
- **ccm** Short for Consensus Cluster Membership. The Heartbeat membership layer.

Working with developers

- State vs Stateless
 - Most applications have state somewhere
 - Discuss how to share/access state
- Data access
 - Local filesystem is king
 - We needed distributed / shared filesystems
- Health
- Metrics

Isolating Services

- 1 application per VM
 - Easy configuration
 - No accidental changes

Anno 2005 – 2010 people called us crazy

- Multiple VM's with the same functions
(A,B,001,002,256)
- Isolate code from config

Lessons from Linux-HA

- Applications need to be adapted
- Files need to be accessible (shared/distributed FS)
- Monitoring strategy needs change

- Good collaboration between developers and ops folks
=> Many general purpose use cases
- Works better with an Open Source Mindset

Failed (Private) Cloud Projects

- Identical copies of Bare Metal or vm's moved to the cloud
- No config management
- No monitoring
- No resilience
- No API usage



- I want a VM

- Please fill in these 4 forms

- Wait 5 weeks

- Repeat

IT Departments have not adapted,
'Shadow'-IT is winning

This new 'D' hype

- New kid on the block
 - Vagrant-lxc with a nice cli
 - The Ultimate “devops tool”
 - “Unseen” growth
-
- Docker is the new cool thing to do




Docker Docker

Docker Docker

Docker

Docker

Culture vs Tools





Kenny Bastani
@kennybastani

Follow

This Dockerizing legacy apps program is kind of like paying IKEA to give your old furniture a Swedish name



Docker brings containerization to legacy apps | TechCrunch
At the DockerCon conference today in Austin, Docker announced a new service called the Modernize Traditional Applications (MTA) Program that enables cust...
techcrunch.com

RETWEETS 25 LIKES 49



2:54 AM - 24 Apr 2017

4 25 49



Richard Seroter
@rseroter

Follow

Stop pretending that things work the same wherever they run (e.g. cloud). It takes work, says @caseywest. #s1p



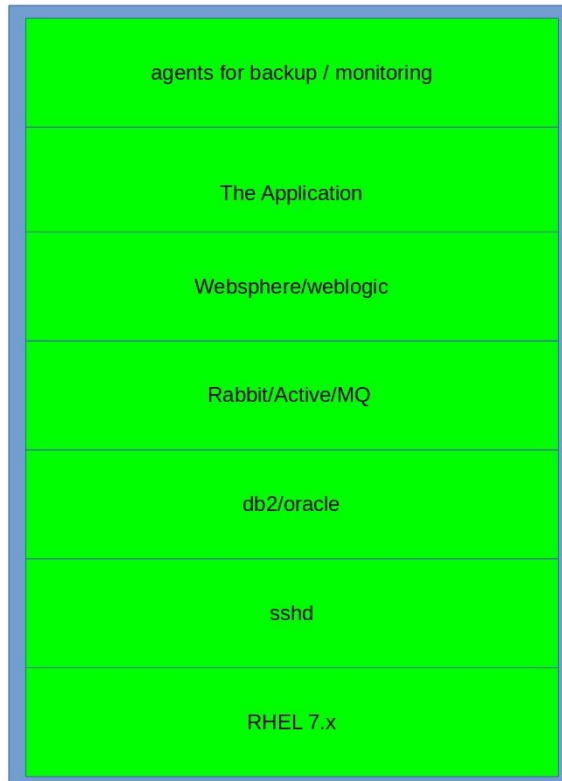
RETWEETS 74 LIKES 68



12:39 AM - 3 Aug 2016


74 68

An Enterprise Container




- No different from a full vm
- Multiple services running in one container
- ssh is the default connection

Image Build by devs, maintained by nobody



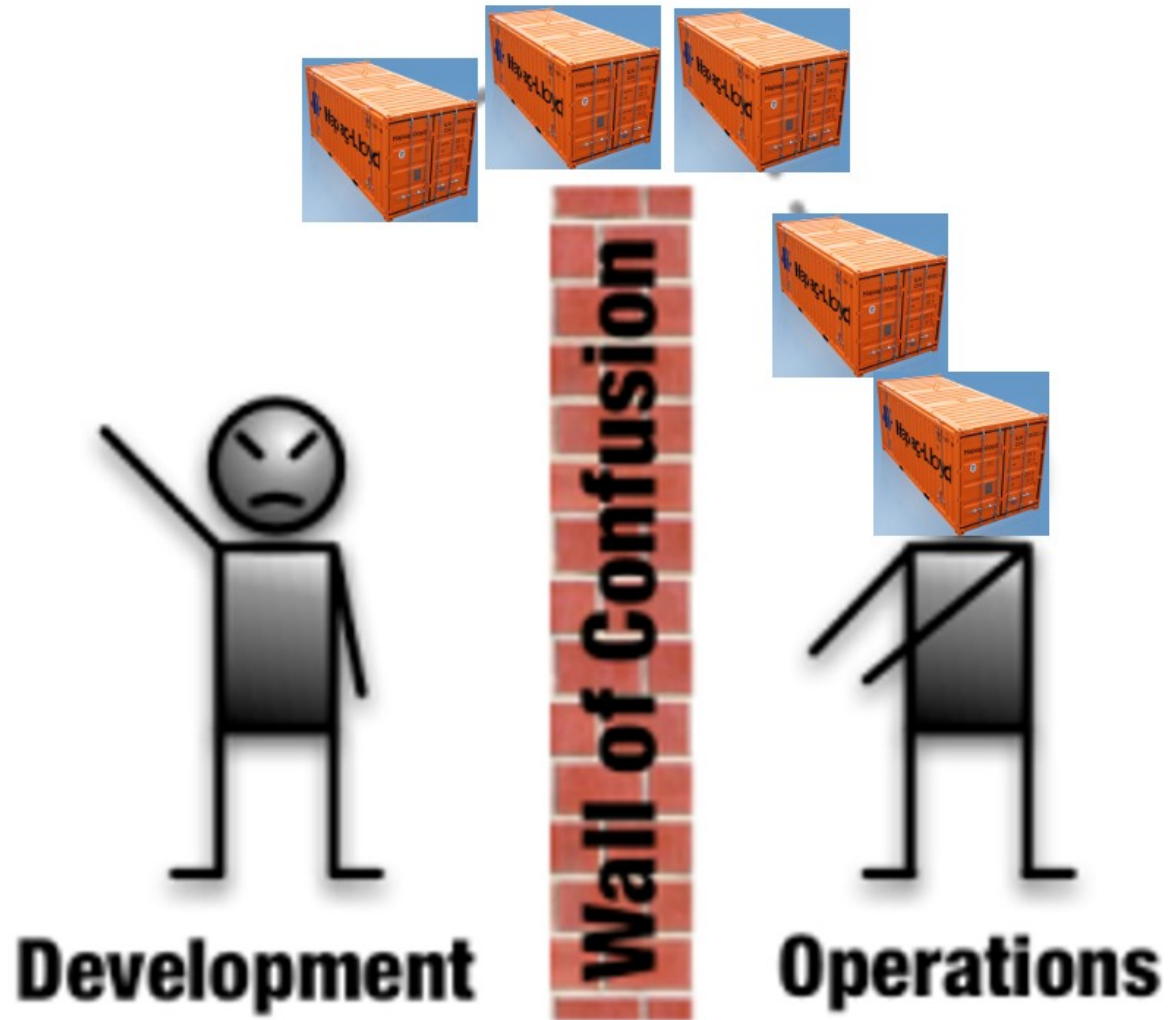
docker



Dev Ooops #container edition

- “Put this Code Live, here's a Docker Image ”
- No machines available ?
- What database ? Where to store the data ?
- Security ? What distro is this even ?
Bad Cows ?
- How do we monitor his ?
- Backups ?
- How did you build this ?





Closing the gaps between dev and ops, **AGAIN !!**

- Where do your containers come from ?
- Who build it ?
- Can you rebuild it ?
- Do you even need a container
- How do you build the hosts that run the containers ?
- Infrastructure as code ++

Kubernetes
Kubernet
Kubernetete
Kubernetete
Docker
Kubernetes
Kube

Container Naive

Home Notifications Messages Search Twitter

Michael Ducey
@mfdll Following

Containers in Dev vs Prod

The diagram is a handwritten note on a whiteboard. At the top, it says 'Containers in:'. Below this, it is divided into two columns: 'DEV' and 'PROD'. Under 'DEV', there is a box labeled 'Docker'. Under 'PROD', there is a list of requirements in boxes: 'Security', 'Code quality', 'Container hosting', 'peer discovery', 'Config changes', 'Supervision', 'monitoring', 'rolling deployment', 'libnetwork', 'Kubernetes/Mesos', and 'Docker'. A red arrow points from the text 'The "learning cliff"' to the gap between the Dev and Prod boxes. The tweet text reads: 'Containers in: DEV PROD The "learning cliff" Docker Security Code quality Container hosting peer discovery Config changes Supervision monitoring rolling deployment libnetwork Kubernetes/Mesos Docker'.

10:27 PM - 10 Feb 2016

1,825 Retweets 1,587 Likes

44 1.8K 1.6K

© 20

10 days into operations



@krisbuytaert

Failed K8s Stories

- Shiny Sales Demos
- Doesn't fit workflow
- Batteries not included
- Multiple attempts at deployment
- Security Ignored

11 days into operations

- Oh I missed 2 other containers
- Security ? What distro is this even ?
Bad Cows ?
- What do you mean “We are a RHEL shop



The New Silo

- Container operations team
- Ignore the old farts
- Who cares about 20 years of experience / security anyhow
- Lets' have some n00bs build a cluster

Dev vs Prod

- 1 local dev machine
- 1 local application
- A cluster
 - ~~Swarm~~
 - Kubernetes
 - ~~Mesos~~
 - Nomad
 - K3s
- Other Applications

Fixing your deployment

- What do you mean we have $n+1$ versions of our code running ($n > 1$)
- My App ain't suitable for that !

Dev vs Prod

- Localhost :8081
- Https
- Dns
- Ingress
- LoadBalancing
- Service Discovery
- Or Vendor Lock In

Dev vs Prod

- Local disk
- No real data
- Database in Container
- Actual customer data
- Distributed Storage
 - 2017 was the year of NFS (again)
- Object Storage
- Backups

Dev vs Prod

- Random image from the internet
- + Some local files
- CI Pipeline
- Image Registry
- Security Scanning

Dev vs Prod

- Works on my machine
- docker pull yolo
- Monitoring
- Metrics
- Security
- Tracing

Containers will not fix your broken culture (and other hard truths)

Clip slide

@bridgetkromhout

Pivotal

@krisbuytaert

inuits
OPEN SOURCE INNOVATORS

The hard truth

- Most organisations have NO usecase for K8s
- Hype is pushing them to use a total misfit between their application architecture and the infrastructure
- Creating even more problems and frustrations
- Large amounts of K8s introductions FAIL
- % of on prem failures is even higher !
- K8s has to be a fit for your architecture

Nothing Wrong with

- docker-compose in a reproducible vm
- Nomad
- No containers at all

Don't fall for

- Cargo Culting
- YoloOps
- Resume Driven development
- Hype Driven development

It's not about the tools
It's about change
It's about the people

Contact

Inuits

Essensteenweg 31
Brasschaat
Belgium
891.514.231

+32 475 961221



Kris Buytaert Kris.Buytaert@inuits.eu

Further Reading

[@krisbuytaert](https://twitter.com/krisbuytaert)

<http://www.krisbuytaert.be/blog/>

<https://inuits.eu/>