



# Novicell Workshop: DevOps and Docker

April 2016

*Henrik Bærbak Christensen*  
Associate Professor  
Computer Science  
University of Aarhus



## The speaker...

### Associated Professor since 2003

- Dept. of computer science / University of Aarhus
- Interests: Software architecture & Teaching
- Leader for SWK part-time education at AU

### Industrial experience

- Architect *and* developer for a product suite of meteorological systems for Danish airports.
- Collaborations with Danish companies: Danfoss, Systematic A/S, B&O, Jyske Bank, TDC, Mjølner, KMD, and many others.
- Recently software intensive projects: Net4Care, EcoSense, CloudArch



# ... have SW in Production ☺

The image shows two overlapping screenshots. The background is the RabbitMQ Management interface, displaying an overview of a cluster named 'rabbit@karibu'. It includes charts for 'Queued messages' and 'Message rates', and a table of nodes. The foreground is a Nagios monitoring dashboard showing the 'Service Status Details For Host Group 'production-servers''. The Nagios interface includes sections for 'Current Network Status', 'New Status Totals', and 'Service Status Totals'. A blue box with the text 'EcoSense Karibu' is overlaid on the bottom right of the screenshots.

CS@AU  
Henrik Bærbak Christensen

# Characteristics

Designed for 24/7/365 operations

– No service windows at all

- Last week we extended disk space from 12 TB -> 15 TB without service interruption
- All security updates and reboots – without service interruption
- All services can be updated and restarted – without service interruption [MQ is only exception!]

– Services to be added at run-time

- New dataformats, new apps, new producers – without service interruption

– Performant and Scalable

- ~65 kB/s – and about 1% CPU load
- More power? Add machines!

But cannot survive:  
power cuts on all machines ☹



## Agenda: Take Away Points

Imhotep Vidensformidling inden for Softwareudvikling

### DevOps:

- Agility in development as well as in production
- *Full stack development*: Teams do the full stack

### MicroServices:

- Decentral data and governance, *products* not projects
- Design for failure

### Docker:

- *Infrastructure as code*
- Lightweight virtualization: *containers ship anywhere*

### Software Architecture!

- *Technology fix will not help if architecture is wrong*

Henrik Bærbak Christensen

5



## Literature

Imhotep Vidensformidling inden for Softwareudvikling

- [Anderson, 2015] Charles Anderson (2015) *Docker*, IEEE Software, May/June 2015. ([Online](#))
- [Bahga and Madisetti, 2014] A. Bahga and V. Madisetti (2014). *Cloud Computing - A Hands-On Approach* Published by >A. Bahga and V. Madisetti
- [Bernstein, 2014] David Bernstein (2014) *Cloud: From LXC to Docker to Kubernetes*, IEEE Cloud Computing, September, 2014. ([Online](#))
- [Buschmann, 2011] Frank Buschmann, 2011 *Tests: The Architect's Best Friend* In IEEE Software, Maj/June 2011 ([Online](#))
- [Fink, 2014] John Fink (2014) *Docker: a Software as a Service, Operating System-Level Virtualization Framework*, code4lib journal, Issue 25, 2014. ([Online](#))
- [Fowler, 2006] Martin Fowler, 2006 *Continuous Integration* ([Online](#))
- [Lewis et al., 2014] James Lewis and Martin Fowler (2014) *Microservices*. ([Online](#))
- [Nygard, 2007] Michael T. Nygard (2007). *Release It! Design and Deploy Production-Ready Software* Pragmatic Bookshelf
- [Rosenblum et al., 2005] Rosenblum, M., Garfinkel, T. (2005) *Virtual Machine Monitors: Current Technology and Future Trends* IEEE Computer, May 2005 (vol. 38 no. 5) ([Online](#))
- [Smith et al., 2005] Smith, J. E., Nair, R. (2005) *The Architecture of Virtual Machines*. IEEE Computer, May 2005 (vol. 38 no. 5). ([Online](#))
- [Christensen, 2010] Henrik Bærbak Christensen *Flexible, Reliable Software -- Using Patterns and Agile Development* CRC Press, 2010

Henrik Bærbak Christensen

6



# DevOps

## The Problem and a Solution

Henrik Bærbak Christensen

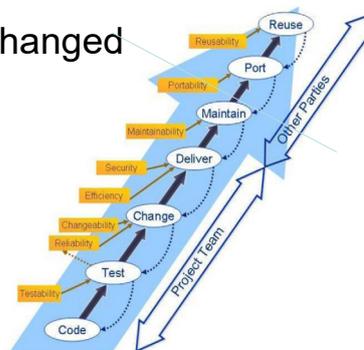
7



# DevOps ?

I am a software developer! I develop software, and then someone else puts it into production – Makes sense for ‘Word’ and ‘Excel’

However, ‘software’ has changed



Henrik Bærbak Christensen

8

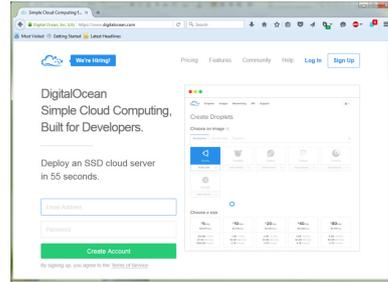


# Software Anno 2016

Imhotep Vidensformidling inden for Softwareudvikling

Software no longer comes on a CD. It is already live on the internet

 Windows Small Business Server 2011



## Computing as Utility

Henrik Bærbak Christensen

9



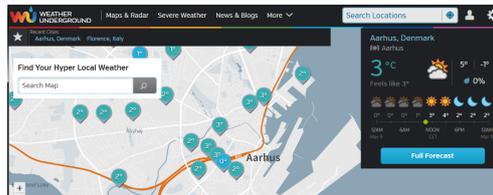
# Software Anno 2016

Imhotep Vidensformidling inden for Softwareudvikling

I buy the capacity I need, not the machinery to execute on...

\$5/mo	\$10/mo	\$20/mo	\$40/mo	\$80/mo
12MB Memory	1GB Memory	<b>Most Popular Plan</b> 5GB Memory	4GB Memory	8GB Memory
1 Core Processor	1 Core Processor	2 Core Processor	2 Core Processor	4 Core Processor
20GB SSD Disk	30GB SSD Disk	60GB SSD Disk	60GB SSD Disk	80GB SSD Disk
1TB Transfer	2TB Transfer	4TB Transfer	4TB Transfer	8TB Transfer
SIGN UP	SIGN UP	<b>SIGN UP</b>	SIGN UP	SIGN UP

And why develop services, if they are already there?



## Computing as Utility

Henrik Bærbak Christensen

10



## Software Anno 2016



### New Years eve:

- 60 bookings per second

### Multiple data centers

- 1000s of servers

### Millions of connected devices

Example: Uber

### Services

- Fare estimation
- Tracking car movement
- Select nearest car
- Fare splits
- Supply positioning
  - Heatmaps
- Fraud detection
- Dynamic pricing
- Backend processing
  - Optimization

*All made using open-source software!*



## But How to Develop?

### From 2015 Master of IT thesis

- A typical development team includes UI designers/developers, backend developers, and database experts. [...] and many other companies split the development process into several components, such as UI, backend, database and mobile applications.

Why does this make sense?

UI designers

Application Server developers

Database designers



## Bottlenecks

Crossing the *boundaries of environments* for the full development-testing-deployment cycle is costly...

- Reconfigurations
- Setup
- Changed architecture
- Changed HW

Even worse

- Different people
- *Someone else's problem*



## DevOps

### Charles Anderson: Docker

The DevOps movement, for example, emerged from one of the classic stumbling blocks in a lot of organizations. Developers build code and applications and ship them to the operations people, only to discover that the code and applications don't run in production. This is the classic "it works on my machine; it's operations' problem now."





## The problems

- Boundaries between *people*
  - *Is it the fault of the UI guys, or the DB guys*
- Boundaries between *environments*
  - *But it works on my machine ???*



## DevOps Movement

### DevOps

From Wikipedia, the free encyclopedia

**DevOps** (a clipped compound of "development" and "operations") is a culture, movement or practice that emphasizes the collaboration and communication of both **software developers** and other **information-technology (IT)** professionals while automating the process of software delivery and infrastructure changes.<sup>[1][2]</sup> It aims at establishing a culture and environment where building, **testing**, and releasing software, can happen rapidly, frequently, and more reliably.<sup>[3][4][5]</sup>



Agile software development has broken down some of the silos between requirements analysis, testing and development. Deployment, operations and maintenance are other activities which have suffered a similar separation from the rest of the software development process. The DevOps movement is aimed at removing these silos and encouraging collaboration between development and operations.

DevOps has become possible largely due to a combination of new operations tools and established agile engineering practices [1], but these are not enough to realize the benefits of DevOps. **Even with the best tools, DevOps is just another buzzword if you don't have the right culture.**

The primary characteristic of DevOps culture is **increased collaboration** between the roles of development and operations. There are some important cultural shifts, within teams and at an organizational level, that support this collaboration.

Virtualization, cloud computing, automated configuration management



<http://martinfowler.com/bliki/DevOpsCulture.html>

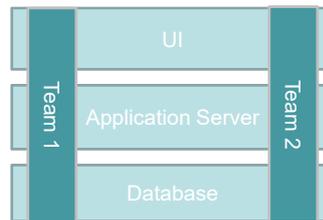


## Full stack developer

- Small teams do *everything*: db design, code development, UI, config, testing, deploy, monitor, production

## Uber development

- 08-12 | develop it
- 12-16 | deploy it
- 16-08 | monitor it
  - (pager at bed table!)



## The new type of code on the block



```

1 regservice:
2   build: .
3   ports:
4     - 4567:4567
5   links:
6     - mongodb:mongo
7
8 mongodb:
9   image: mongo:latest
10 |

```

- **Coding application logic:** The programming of logic for the core business functionality under the assumption that all services work correctly. Example: Developing application server code to access an external inventory service and retrieve the price and stock availability of an item.
- **Coding quality attribute logic:** The programming of logic for architectural quality attributes [4] (non-functional requirements) like availability or performance, such as to handle the situations where one or several services fail, are slow, or produce unexpected results. Example: Developing code that implement graceful failure modes in case the inventory service is too slow to respond due to a peak load or simply unavailable due to network failure.
- **Coding infrastructure logic:** The programming of logic for the deployment of services. Traditionally handled by manual procedures (installing, configuring, and linking services), but in face of large-scale deployments, this too must be coded. Example: Developing scripts that start the application server, inventory service and associated database, initialize them, and connect them correctly—i.e. create a *staging environment*.

Teaching DevOps and Cloud Computing using a Cognitive Apprenticeship and Story-Telling Approach

Henrik Bærbak Christensen  
Computer Science Aalborg University  
4000 Aalborg Øst, Denmark  
hb@cs.aau.dk



# MicroServices



# Micro Services

## The agony of being an old man...

- Same ideas, new packaging, pops up every 10 years
  - Modular programming, object-orientation, component-based, service-oriented,... and now *micro services*
- All the same idea, dating to David Parnas
  - *Information hiding* in modular programming
  
- Gamma et al.
  - ❶ *program to an interface*
  - ❷ *favor object composition*
  - ❸ *encapsulate what varies*



## What varies?

Imhotep Vidensformidling inden for Softwareudvikling

The idea is the same, only the packing and wiring varies

- Modula 2 module = compilation unit, static link
- Class = compilation unit, static link
- Components = deployment/static unit, static link
- SOA = deployment/run-time unit, dynamically linked
- MicroService = deployment/run-time unit, dynamically linked

Henrik Bærbak Christensen

21



## What is New?

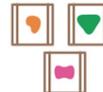
Imhotep Vidensformidling inden for Softwareudvikling

The *real* new concept here, IMO, is the *self-containedness* and *autonomy* of a MS

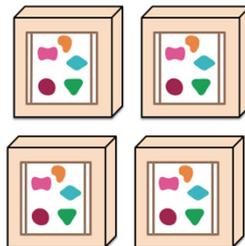
A monolithic application puts all its functionality into a single process...



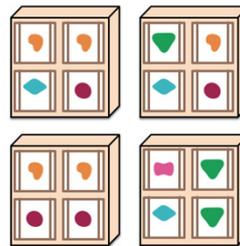
A microservices architecture puts each element of functionality into a separate service...



... and scales by replicating the monolith on multiple servers



... and scales by distributing these services across servers, replicating as needed.



<http://martinfowler.com/articles/microservices.html>

Henrik Bærbak Christensen

22



# MS = Business Capability

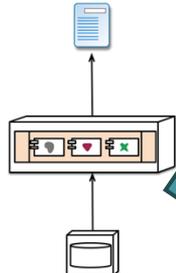
Imhotep Vidensformidling inden for Softwareudvikling

UI specialists

middleware specialists

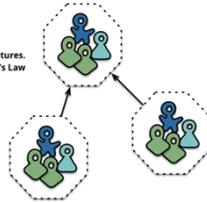
DBAs

Siloed functional teams...

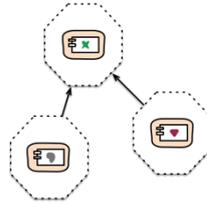


... lead to siloed application architectures. Because Conway's Law

Ownership: you build it, you run it!  
Products, not projects



Cross-functional teams...



... organised around capabilities Because Conway's Law

Henrik Bærbak Christensen

23



# Decentralizing...

Imhotep Vidensformidling inden for Softwareudvikling

## Governance

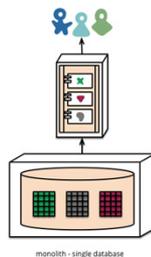
- Choice of platform and technology by the team
  - *Pick the right tool for the job*

## Data management

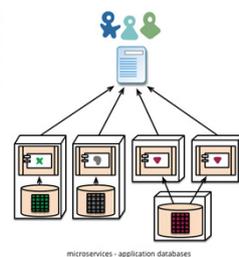
- MS = Products own their data

## Communication

- Typical
  - Network
  - RESTish
    - Lightweight
    - Extensible
    - Fault tolerant



monolith - single database



microservices - application databases

Henrik Bærbak Christensen

24



## Design for Failure !

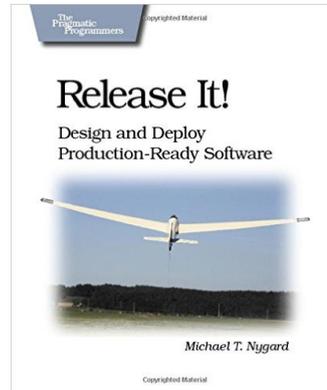
Get that book! →

*Safe failure modes*

- *Failing services*
- *Slow responses*
- *Cascading failures*

*Redundancy*

*Horizontal scaling*



CAP theorem

- *Consistency* (all nodes see the same data at the same time)
- *Availability* (a guarantee that every request receives a response about whether it succeeded or failed)
- *Partition tolerance* (the system continues to operate despite arbitrary partitioning due to network failures)



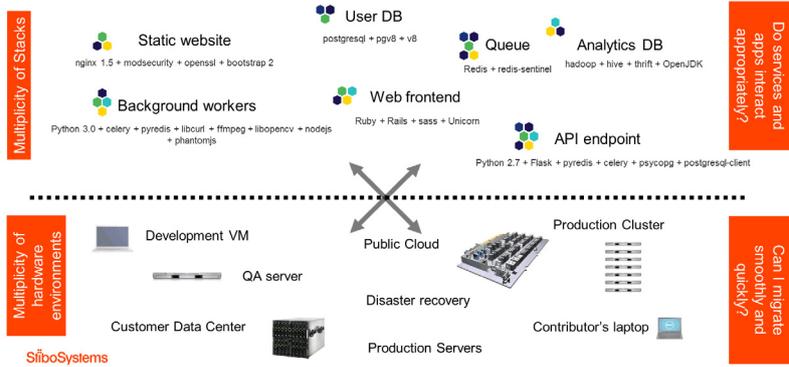
## Containers and Docker



# The Problem Again

Imhotep Vidensformidling inden for Softwareudvikling

## Crossing boundaries, that is, *moving code*



Source: Torben Haagh, Sibosystems

Henrik Bærbak Christensen

27



# Was Solved in 1960'sies

Imhotep Vidensformidling inden for Softwareudvikling



Henrik Bærbak Christensen

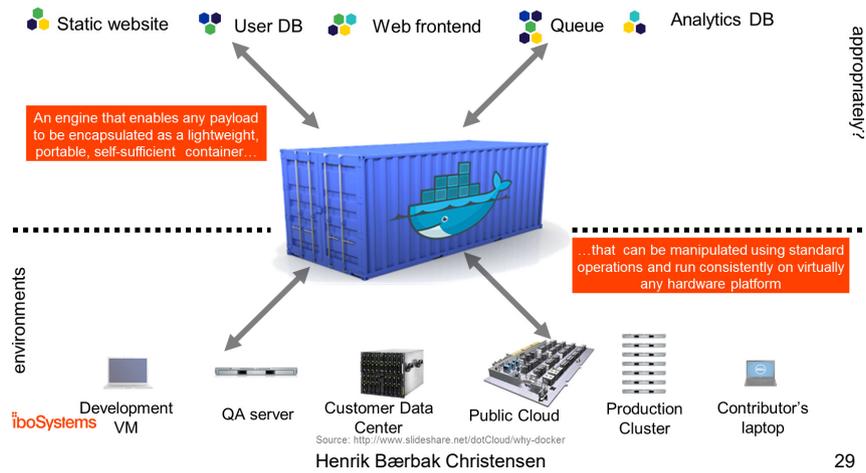
28



# Docker = Container

Imhotep Vidensformidling inden for Softwareudvikling

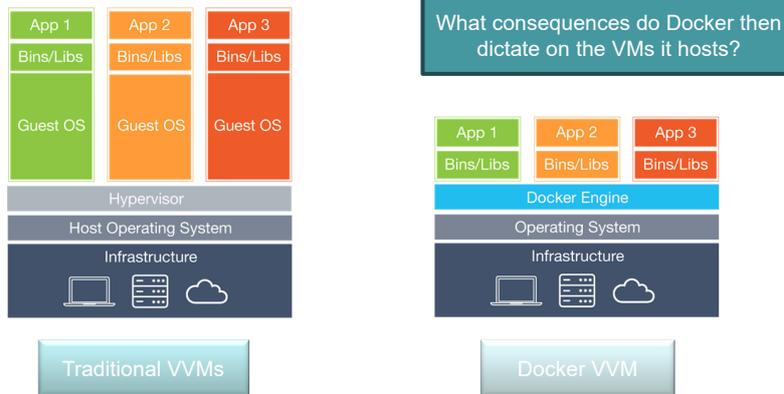
## Docker is a shipping container system for code



# Light-Weight virtualization?

Imhotep Vidensformidling inden for Softwareudvikling

## Moving the virtualization boundary





## Docker Engine

### Core concepts

- Image            Deployment Unit – a Virtual Machine
  - Naming: **ownername/imagename:tag**
- Container        Executing process(es) in an Image
  - Naming: **'trusty\_thar'** or give it your own name

### Lifecycle *classic* for building *your* image

- container = instantiate(image)
  - `docker run -d --name c1 baerbak/dockerdemo:v1`
- modify container
  - Install software, change files, add stuff, ...
- commit container → image2
  - `docker commit c1 baerbak/dockerdemo:v2`



## Tech Glimpse

### Onion file system: *Copy-on-Write*

- Every operation basically creates a new *file layer*
  - *Changing 'hans.txt' in layer N creates a (modified) copy of 'hans.txt' in layer N+1*

### Base images = 'prebaked file system'

- All layers up-till N forms an Image

### I.e. `henrikbaerbak/cloudarch:base_v2`

- Ubuntu 14.04 LTS server base image
- Java, Ant, Ivy, Git, ... are all layered on top



# Docker Engine

Imhotep Vidensformidling inden for Softwareudvikling

But but – *infrastructure code* ???  
Lifecycle *scripted* for building *your* image

## Dockerfile

- *Infrastructure code* at the image level
- Defines the contents of the image
- docker build -t (name)
  - creates the image from local 'Dockerfile'

```

1 FROM henrikbaerbak/cloudarch:base_v1
2
3 # CaveReg docker file. Requires a running mongodb container
4 # named 'mongo' running.
5
6 MAINTAINER Henrik Baerbak Christensen <hbc@cs.au.dk>
7
8 # Create the /root/regservice
9
10 RUN mkdir /root/regservice
11 WORKDIR /root/regservice
12
13 # Copy local stuff there
14 COPY src/ /root/regservice/src/
15 COPY test/ /root/regservice/test/
16 COPY resource/ /root/regservice/resource/
17
18 COPY build.xml /root/regservice/build.xml
19 COPY ivy.xml /root/regservice/ivy.xml
20
21 # Setup the entry point
22 COPY entry-point.sh /root/regservice/entry-point.sh
23 RUN chmod 755 /root/regservice/entry-point.sh
24
25 # Default port exposed, standard spark-java port
26 EXPOSE 4567
27
28 # Entrypoint, you need to supply the 'web' as the ant target to run
29 ENTRYPOINT ["./root/regservice/entry-point.sh"]
30 CMD ["web"]

```

Henrik Bærbak Christensen

33



# Docker Hub

Imhotep Vidensformidling inden for Softwareudvikling

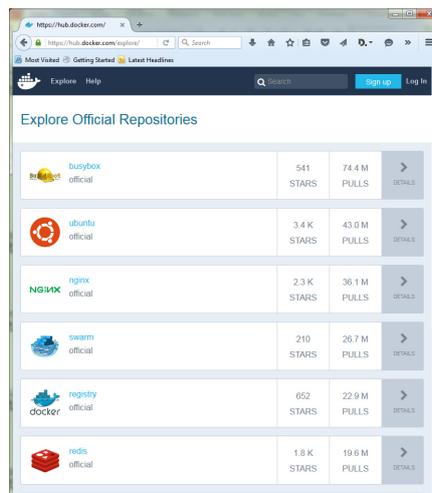
Similar to Maven Repo

Repo of open-source images

- docker pull (name)

Free hosting

- Register and
- docker push (name)



Henrik Bærbak Christensen

34





## Docker in Action

From my  
Cloud Computing Course



## SkyCave

Massive Multiplayer Online (MMO)  
Social Networking  
Domain: 'Colossal Cave' – reimagined

### Characteristics

- Subscription-based
- Single Sign-on
- NoSQL
- REST services

```
Colossal Cave Adventure • Score: 36 • Turns: 3
your surroundings. Typing "inventory" tells you what you're
carrying. "Get", "drop" and "throw" helps you interact with
objects. Part of the game is trying out different commands and
seeing what happens. Type "help" at any time for game
instructions.

Would you like more instructions?

> no

You are standing at the end of a road before a small brick
building. Around you is a forest. A small stream flows out
of the building and down a gully.

> e

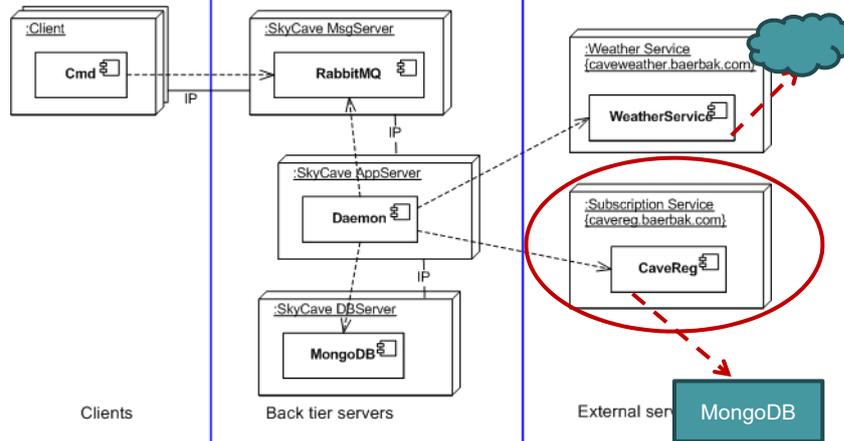
You are inside a building, a well lit room for a large spring.
There are some keys on the ground here.
There is a shiny brass lamp nearby.
There is tasty food here.
There is a bottle of water here.
What's next?
```



# Deployment Viewpoint

Imhotep Vidensformidling inden for Softwareudvikling

## The central nodes in SkyCave / Three Tier + MS



CS@AU

Henrik Bærbak Christensen

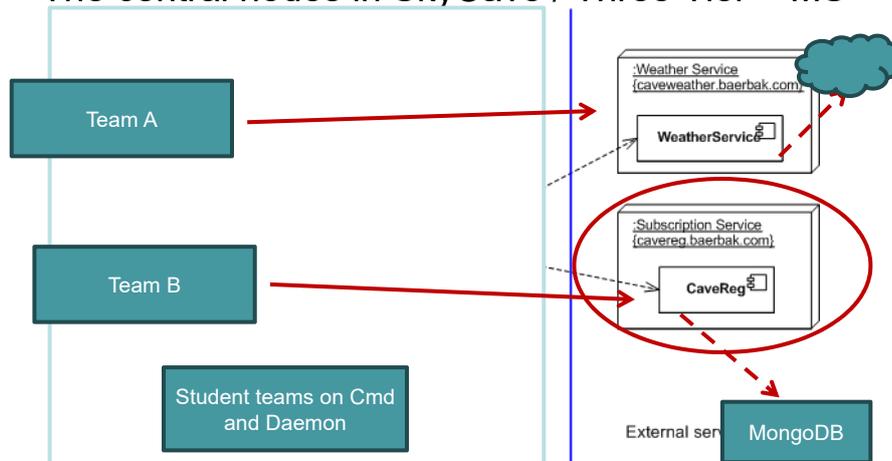
39



# Governance

Imhotep Vidensformidling inden for Softwareudvikling

## The central nodes in SkyCave / Three Tier + MS



CS@AU

Henrik Bærbak Christensen

40



# MicroService: Subscription

Imhotep Vidensformidling inden for Softwareudvikling

The Subscription service is a web site, backed up by a MongoDB database, which allows:

- Creating a new subscriber
- Authenticating through REST call

Thus, to run, a MongoDB container must be live first!

Henrik Bærbak Christensen

41



# Example: Docker Compose

Imhotep Vidensformidling inden for Softwareudvikling

Multiple compose files

- Layering
  - Base: Core stuff for staging environment
  - Prod: Additional stuff for real production running

For staging

- docker-compose up
  - Read only docker-compose.yml

For production

- docker-compose -f docker-compose.yml -f docker-compose.prod.yml up -d

```

1 regservice:
2   build: .
3   ports:
4     - 4567:4567
5   links:
6     - mongodb:mongo
7
8 mongodb:
9   image: mongo:latest
10 |
11
12 mongodb:
13   volumes:
14     - /mongo_datadir:/data/db
15

```

Henrik Bærbak Christensen

42

# Example

Imhotep Vidensformidling inden for Softwareudvikling

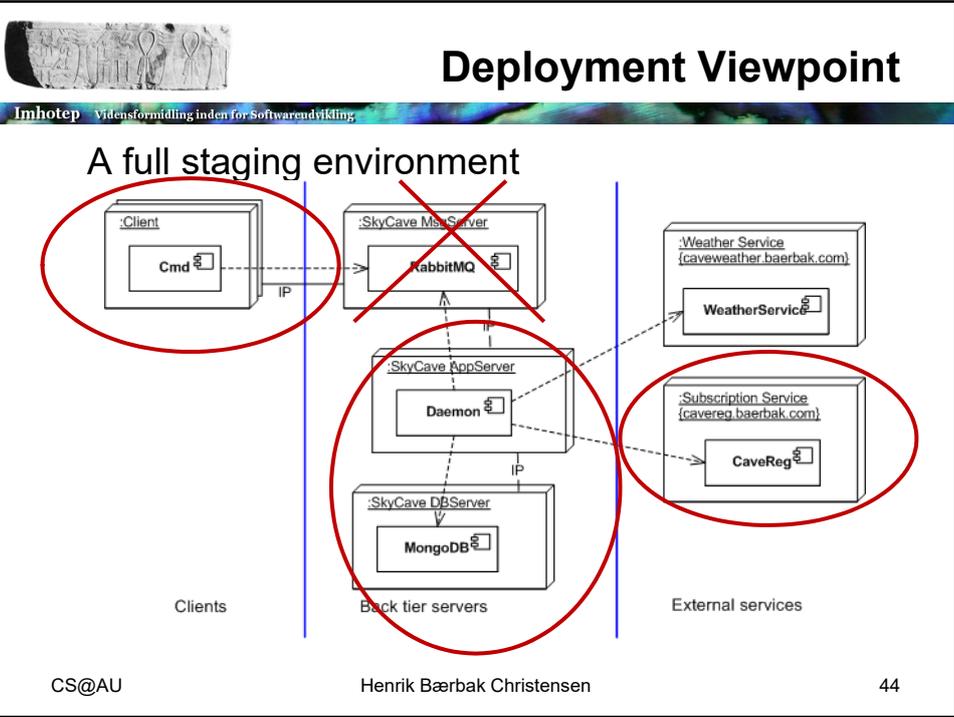
The screenshot shows a web browser window with the URL localhost:4567/home. The page displays "Welcome, mikkel!" and two buttons: "Edit Registration" and "Logout". To the right, a terminal window shows the following logs:

```

sasp@SaspDev: ~/dev/dockerdemo/registration
[... logs about MongoDB connections and application startup ...]
> use cave; findOne()
> use cave
switched to db cave
> show collections
SubscriptionRecord
> db.SubscriptionRecord.findOne()
{
  "_id": ObjectId("5e11fb0e400b1161fa4d3f"),
  "className": "cloud.cave.web.models.SubscriptionRecord",
  "loginName": "7777777",
  "playerName": "mikkel",
  "passwordHash": "$2a$10$YGBCuWnErJyomdss6u/OKR0K0JTFa09JfF/WS.Yfsokt7UMbx6",
  "region": "NODENSE",
  "groupName": "novicell",
  "dateCreated": ISODate("2016-03-10T07:18:15.144Z")
}

```

Henrik Bærbak Christensen 43







## Discussion

### Benefits:

- *Infrastructure as code*
  - Fast, reproducible execution of deployments
  - Excellent documentation of deployment view

### Liabilities:

- Docker is a moving target
  - E15: Linking was the central composition tool; now deprecated!
  - Documentation erosion!
- Layered file system: Take care with security!
  - No, you cannot delete the credential file again...



**... And software architecture**



## Discussion

MicroServices and Docker are **patterns** and **technology**, they do not by themselves solve the deep architectural issues of 'hard to maintain monoliths'

- *Loose coupling* and *High cohesion* are challenging no matter what technology you use!
- Software Architecture =
  - Making the *right* decisions on
    - Splitting the behaviour in the proper units
    - Providing adequate connections between the units
    - Controlling quality attributes
      - » Especially availability, modifiability, performance in MS



## Discussion

If a MS architecture

- Have bad boundaries between units
  - Require all units' behaviour and the networks between them to operation correctly
- ... you still have a Monolith that is very hard to maintain and grow...



## Discussion

### Design for failure, means

- Introduce proper techniques to handle cascading failures at the integration points
  - Detect failure
- Consider how graceful degradation works
  - How to continue operation once the failure occurs?

### Example:

- Player wants to log into SkyCave, but the subscription server is not responding
  - Reject login? (Loss of interest, bad reputation)
  - Accept limited login? ("Karl, I know him, let him in...")



## Summary



## Agenda: Take Away Points

### DevOps:

- Agility in development as well as in production
- *Full stack development*: Teams do the full stack

### MicroServices:

- Decentral data and governance, *products* not projects
- Design for failure

### Docker:

- *Infrastructure as code*
- Lightweight virtualization: *containers ship anywhere*

### Software Architecture!

- *Technology fix will not help if architecture is wrong*



**Thanks...**

Questions?