



**University of Stuttgart**  
Institute of Software Technology  
Reliable Software Systems

# Performance Engineering for Microservices

## Research Challenges and Directions

**André van Hoorn**

(and 7 more)

# The Authors



Robert Heinrich  
KIT  
Germany



Lucy E. Lwakatare  
U Oulu  
Finland

André van Hoorn  
U Stuttgart  
Germany



Claus Pahl  
Free U Bozen-Bolzano  
Italy



Holger Knoche  
Kiel University  
Germany



Stefan Schulte  
TU Wien  
Austria

Fei Li  
Siemens AG  
Austria



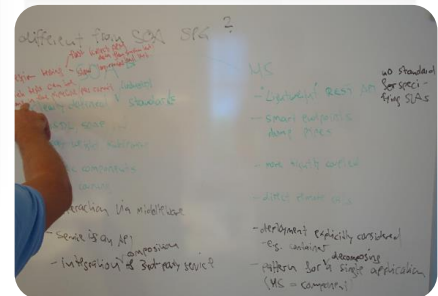
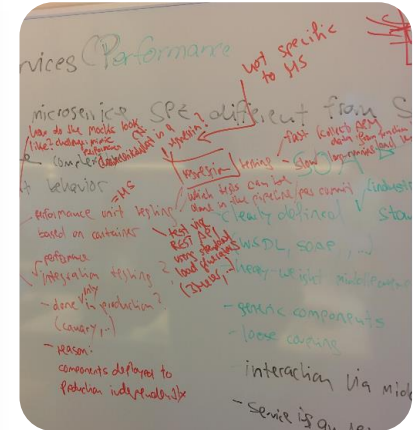
Johannes Wettinger  
U Stuttgart  
Germany




# Result of Break-Out-Group @ GI Dagstuhl Seminar 16394

## Software Performance Engineering in the DevOps World

<http://www.dagstuhl.de/16394>





# Flow of Discussion – Here and There



# DevOps and Microservices

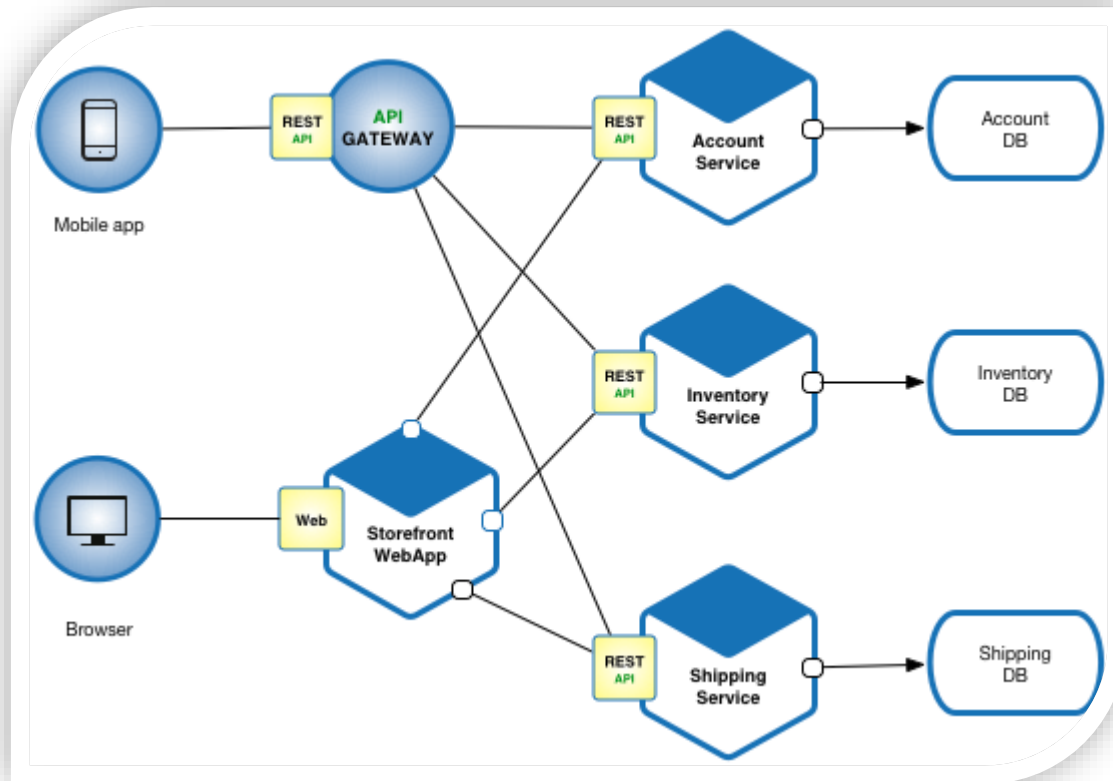
# DevOps and Continuous Deployment



*“DevOps is a set of practices intended to reduce the time between committing a change to a system and the change being placed into normal production, while ensuring high quality.”* – Bass et al., 2015

- Practiced in different magnitudes
  - cars.com => 700 deployments/Year (~ 2 deployments/Day)
  - flickr => 10+ deployments/Day
  - amazon.com => Every 11.6 seconds (~ 7500 deployments/Day)
- Beneficial properties to reduce the time to production
  - No or very few coordination with other teams is required
  - Multiple versions of the same service can be run in the production environment simultaneously
  - Roll back/forward changes made to a running service in the event of errors

# Adequate Architectural Style: Microservices




<http://microservices.io/patterns/microservices.html>

# Common Platform: Container-based Virtualization

- **Docker**

- Open platform for building, shipping and running distributed applications



`$ docker run account`

- **Kubernetes**

- Container cluster manager
- Orchestrates deployment of containers
- + much more:

Load-balancing, auto-scaling, fault-management, monitoring, service discovery, ...


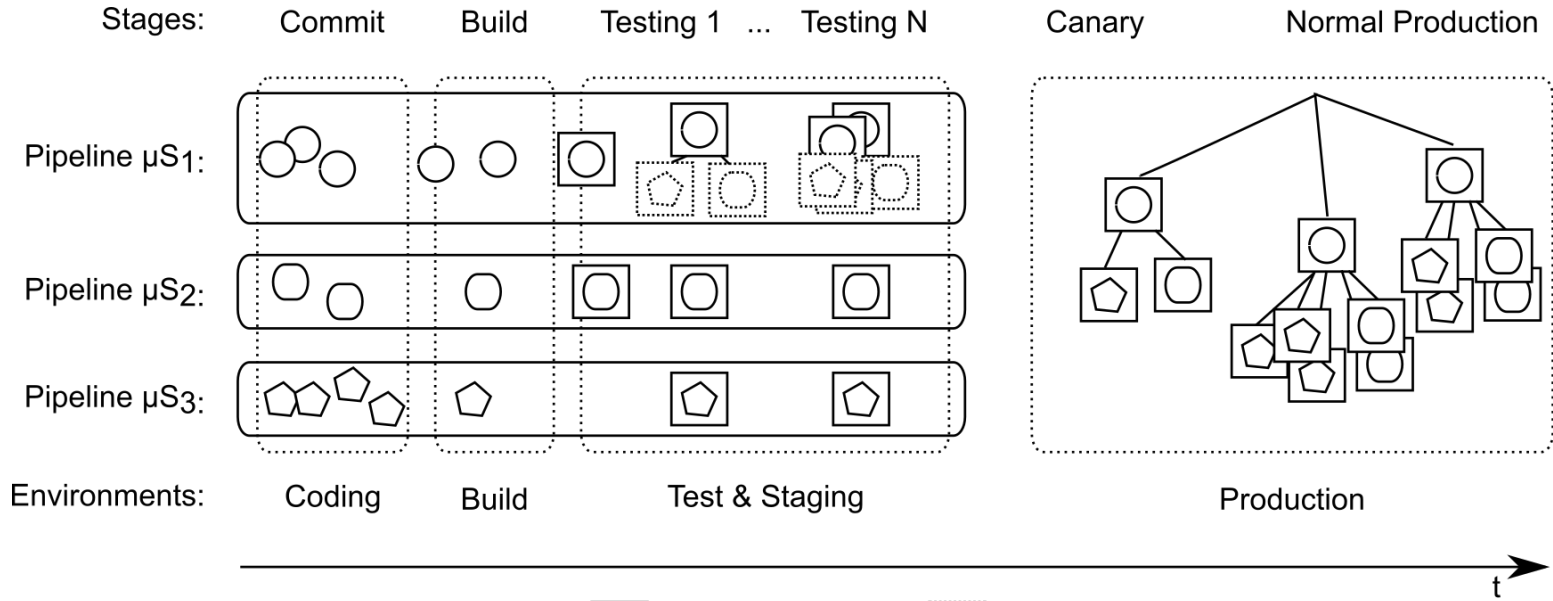


**kubernetes**

`$ kubectl scale rc account --replicas=3`  
`$ kubectl --rolling-upgrade inventory inventory:5.3`

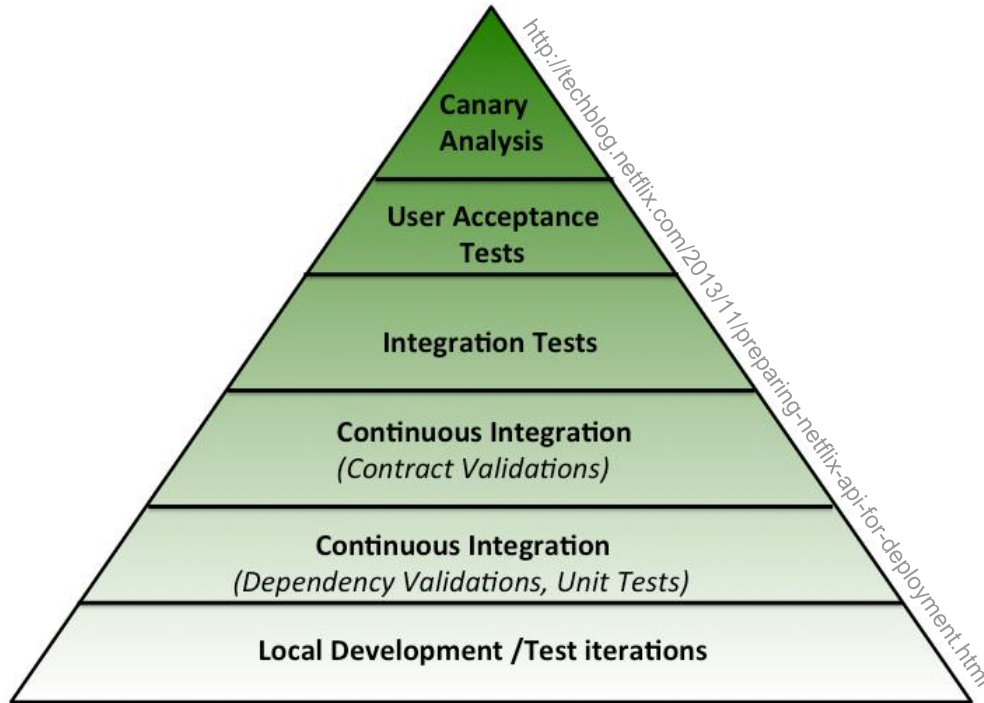


# Microservices in the CD Pipeline



**Testing**

# Common Testing Practices for DevOps/Microservices

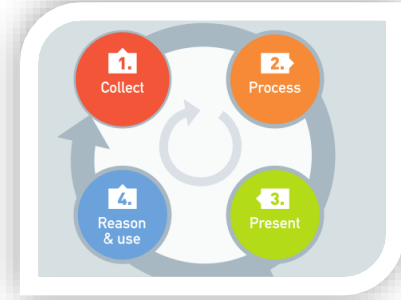


# Performance Testing Microservices

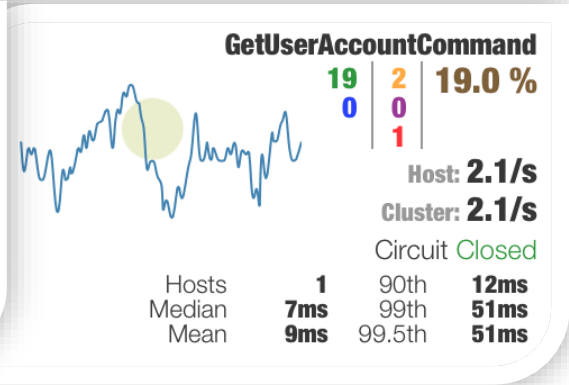
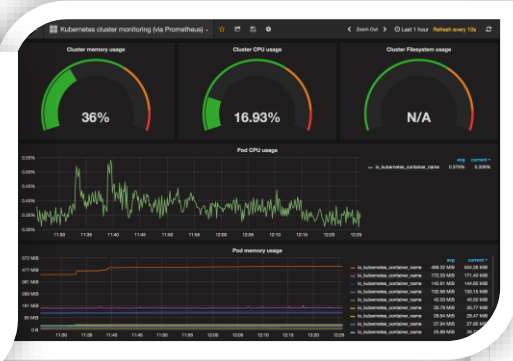
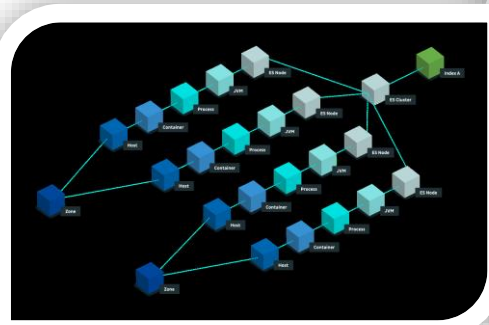
- Conflict with CD practices:  
extensive performance tests vs. pace of stages
- Concerns
  - Which tests are run for each commit?
  - Which tests are run for a consolidated set of commits?
  - Which tests to run in which situations/environments?
- Research directions
  - Splitting test suites by different scopes + (dynamic) test case selection
    - Decide which tests should be performed
  - Innovative online performance testing strategies (canary, ...)
  - Use of input from monitoring data, operational failures, and the actual application status
    - Extraction and refinement of tests

# Monitoring

# Monitoring Microservices



- Monitoring is a first-class property of microservices/DevOps
- Extensive monitoring facilities available
  - Application, infrastructure, and system level:
    - mature APM tools
    - monitoring capabilities in frameworks (Hystrix, ...)
  - Container level: built-in monitoring (Docker, Kubernetes, ...)



<https://www.insta.com/product/>

<https://docs.giantswarm.io/>

<https://github.com/Netflix/Hystrix>

# Monitoring Microservices – Challenges

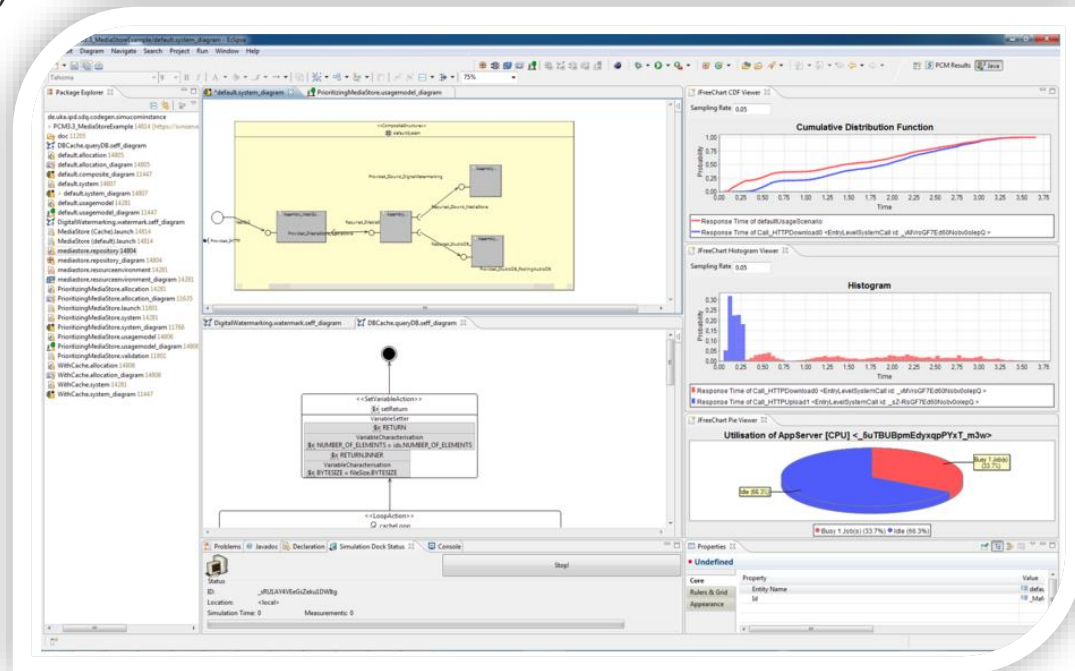
1. Instrumentation challenge: polyglot technology stacks
2. Classic metrics meaningful? New metrics? For instance,
  1. Short-lived (stateless) containers: life time, capacity/elasticity
  2. State of resilience mechanisms (e.g., circuit breakers)
3. Accurate and precise anomaly detection and diagnosis under frequent change

# Modeling and Prediction



# Performance Modeling of Microservice Architectures

Performance modeling (and prediction) gained considerable attraction and maturity in SPE over the past two decades (UML SPT/MARTE, LQN, Palladio, ...)



# Why Don't We Simply Use the Existing Ones?

## 1. Traditional use case – capacity planning – less relevant (auto-scaling)


- But: resilience and design of runtime adaptation strategies

## 2. New abstractions needed to adequately capture the recent advances in deployment technology (container orchestration, scale, ...)

- More light-weight techniques possible?

## 3. Model creation is challenging

1. Application architecture can be extracted (e.g., from execution traces)
2. How to learn the characteristics of the execution platform?



# Conclusions

# Summary of Observations and Challenges

- Observations

- Microservices are an emerging architectural style enabling the efficient use of cloud technologies and DevOps practices
- So far, performance engineering for microservices has not gained attraction in the relevant communities
- Existing performance engineering techniques—focusing on testing, monitoring, and modeling—cannot simply be re-used

- Particular (activity-specific) challenges include

1. strategies for efficient performance regression testing
2. performance monitoring under continuous software change
3. appropriate performance modeling concepts for shifted use cases

# Future Directions – Combining Testing, Monitoring, Modeling

- The use of **performance models to prioritize tests cases**, including the decision whether and when (pipeline vs. production) tests are executed
- The use of **monitoring data to create and refine performance tests**, including the creation of representative usage profiles and performance-aware service mockups
- The use of **tests and resulting monitoring data to automatically create performance models**, including the combination of data from different environments
- The use of **models to guide the diagnosis** of performance problems
- **Learning and assessing deployment strategies** based on monitoring data

# Advertisement / Call for Collaborations

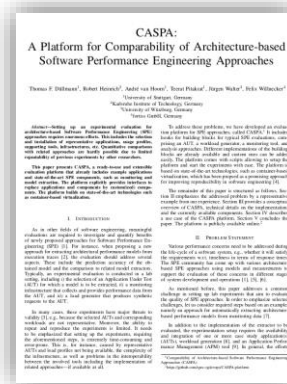
- SPEC RG DevOps Performance



**Mission:** foster and facilitate research in combining **measurement-based** application performance management (APM) and **model-based** software performance engineering (SPE) activities for **business-critical** application systems.

<http://research.spec.org/devopswg>

- **CASPA:**  
**A Platform for Comparability of Architecture-based Software Performance Engineering Approaches**  
*Düllmann et al., ICSA '17*



<https://github.com/spec-rgdevops/CASPA-platform>