Performance testing

Tools and Techniques for Software Testing - Barbara Russo

SwSE - Software and Systems Engineering group

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

Modern (online) systems may underperform as they are overloaded

unibz

Performance testing

- Performance testing is the **process** of determining the **speed, responsiveness and stability** of a computer, network, **software program** or device **under a workload**
- Performance testing can involve quantitative tests done in a lab, or occur in the production environment in limited scenarios

wikipedia

Open question: how to test in in-production environments?

Example of performance testing: Load Testing

Tools and Techniques for Software Testing - Barbara Russo

SwSE - Software and Systems Engineering group

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

“Load testing is the process of assessing the behavior of a system under load in order to detect load-related problems”

Jiang et al., 2015

Non-functional testing

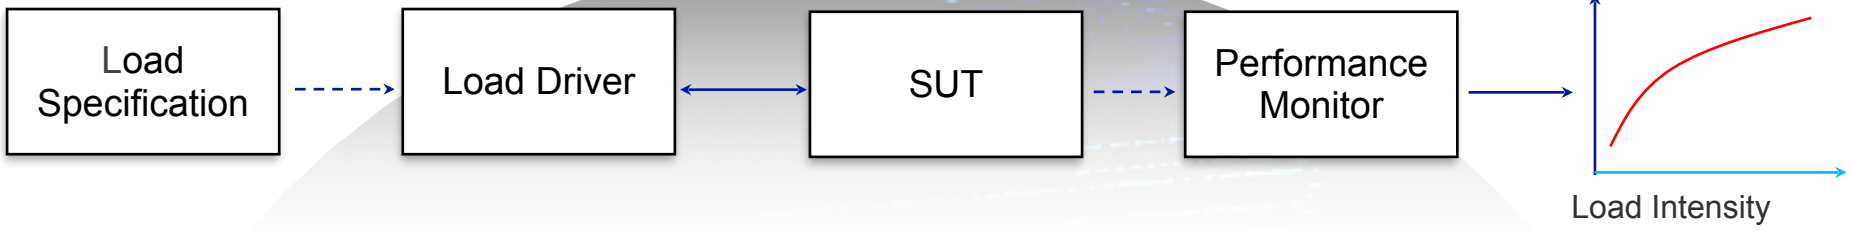

What is load?

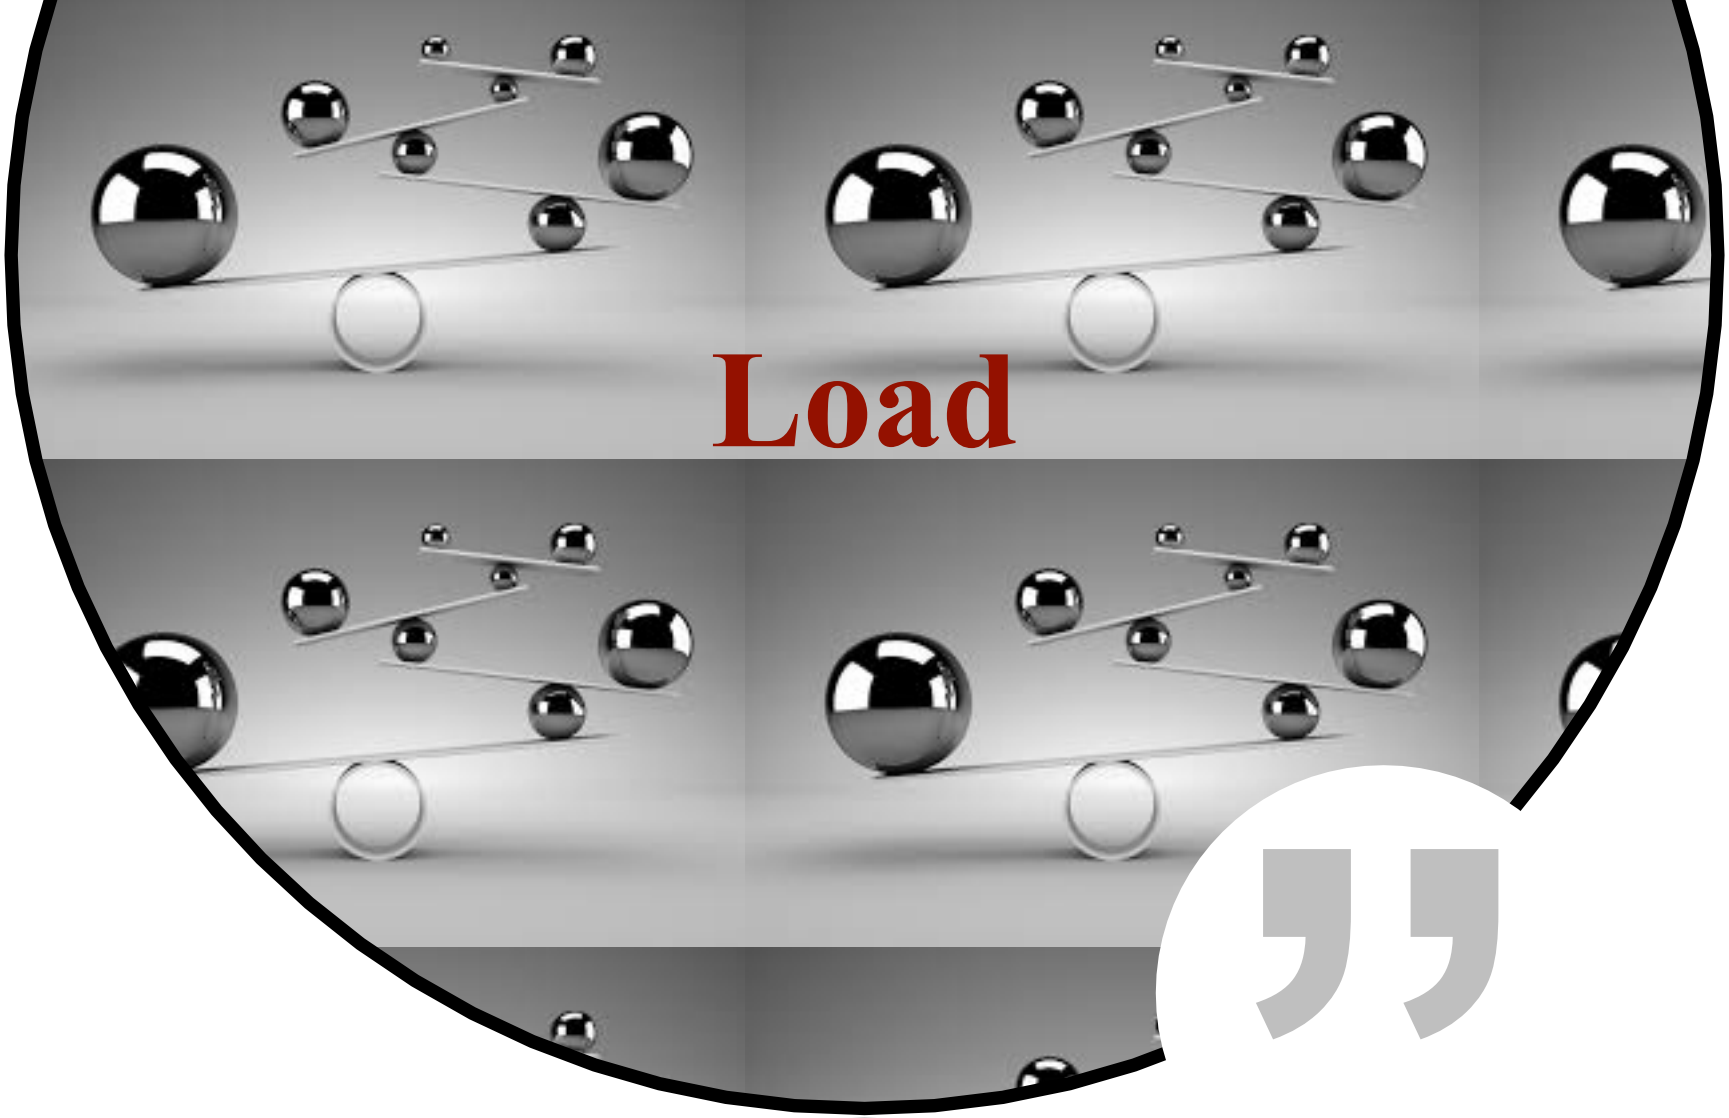

Load

Amount of computational work being performed by a software system

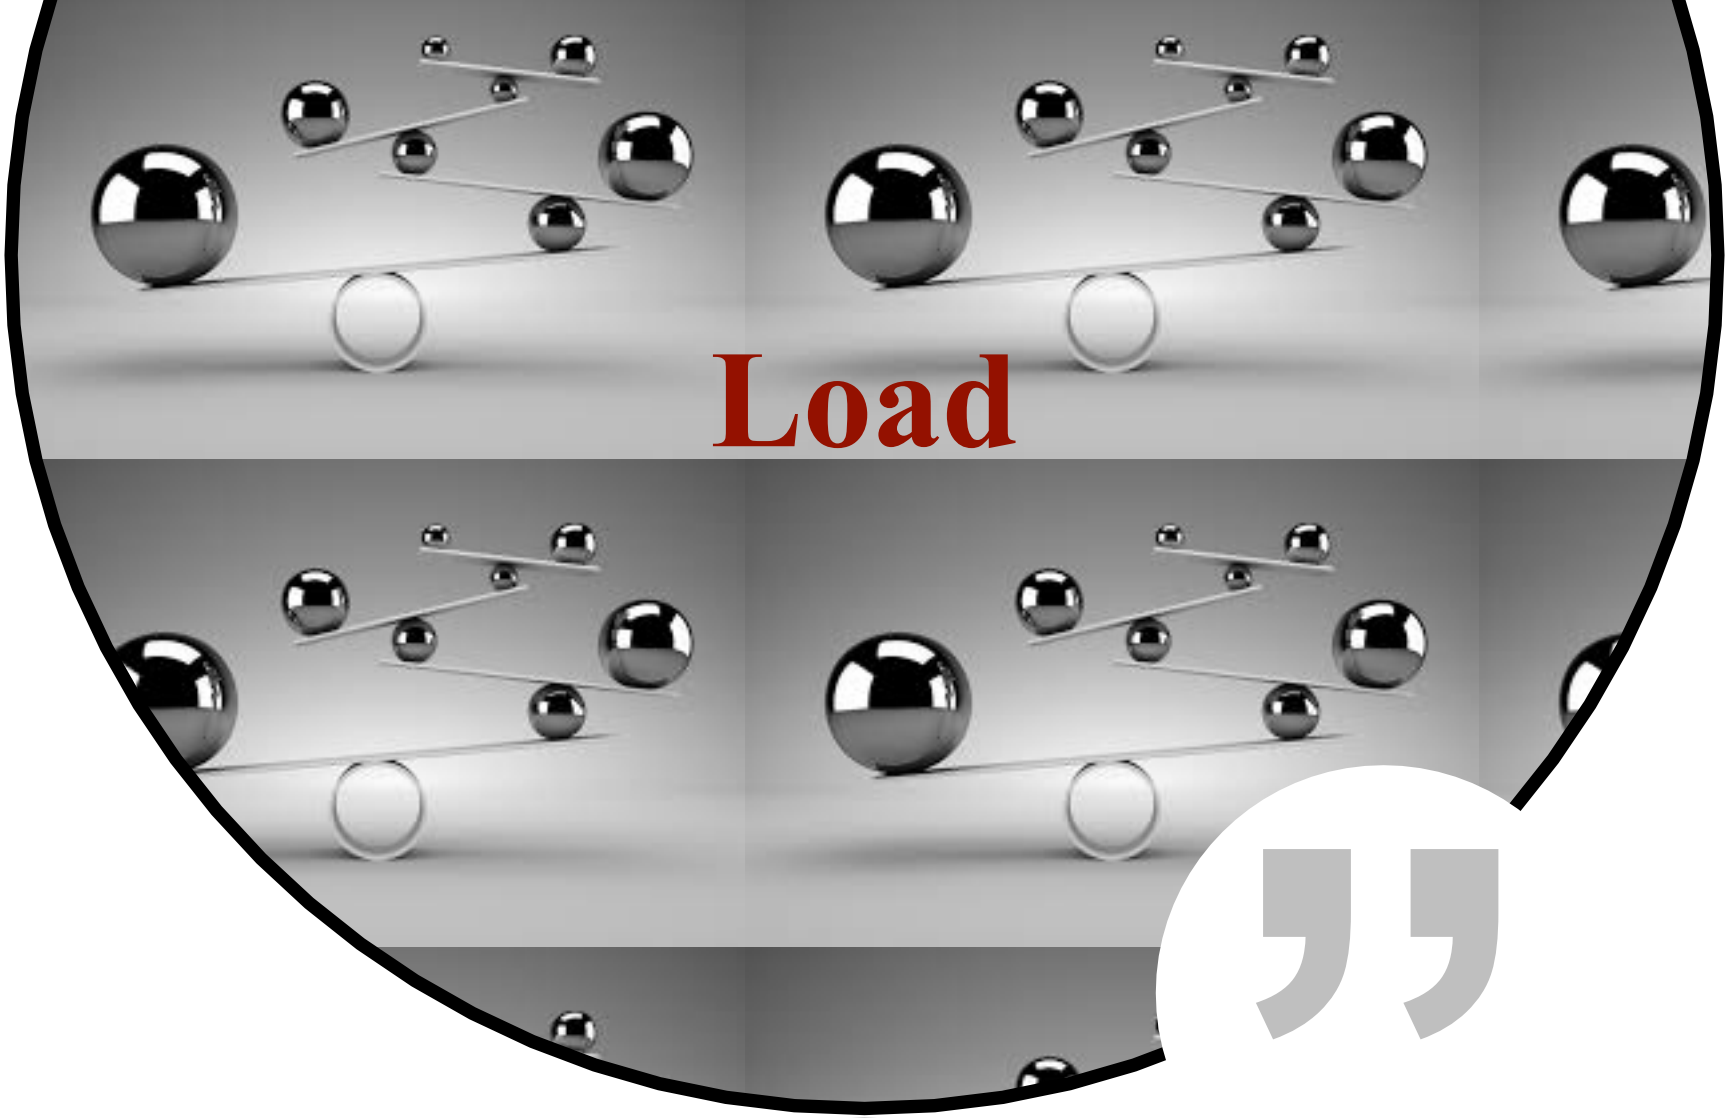

Amount of concurrent users

The Classic Load Testing Approach

… and Classic Problems

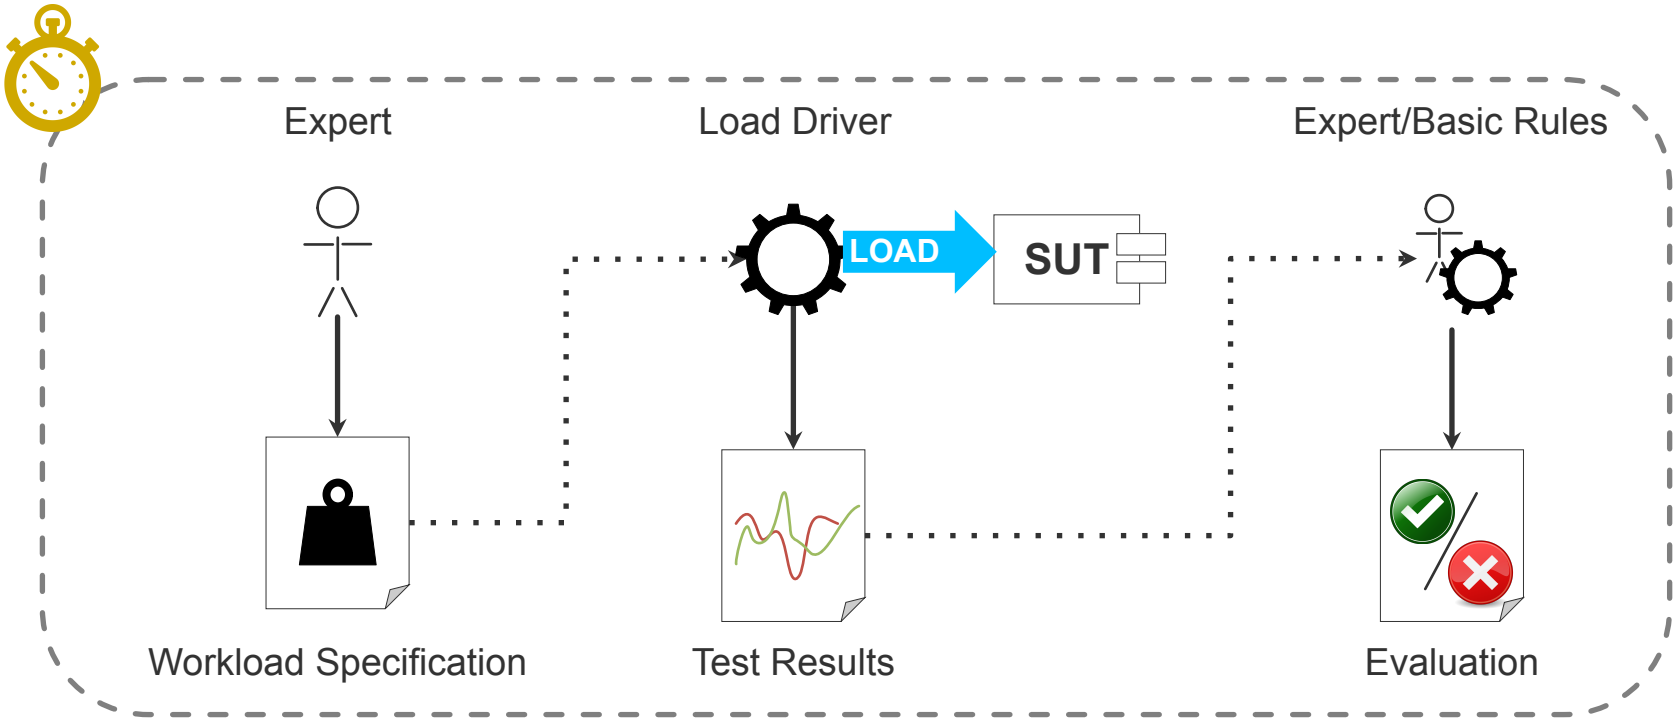

The Classic Load Testing Approach

… and Classic Problems

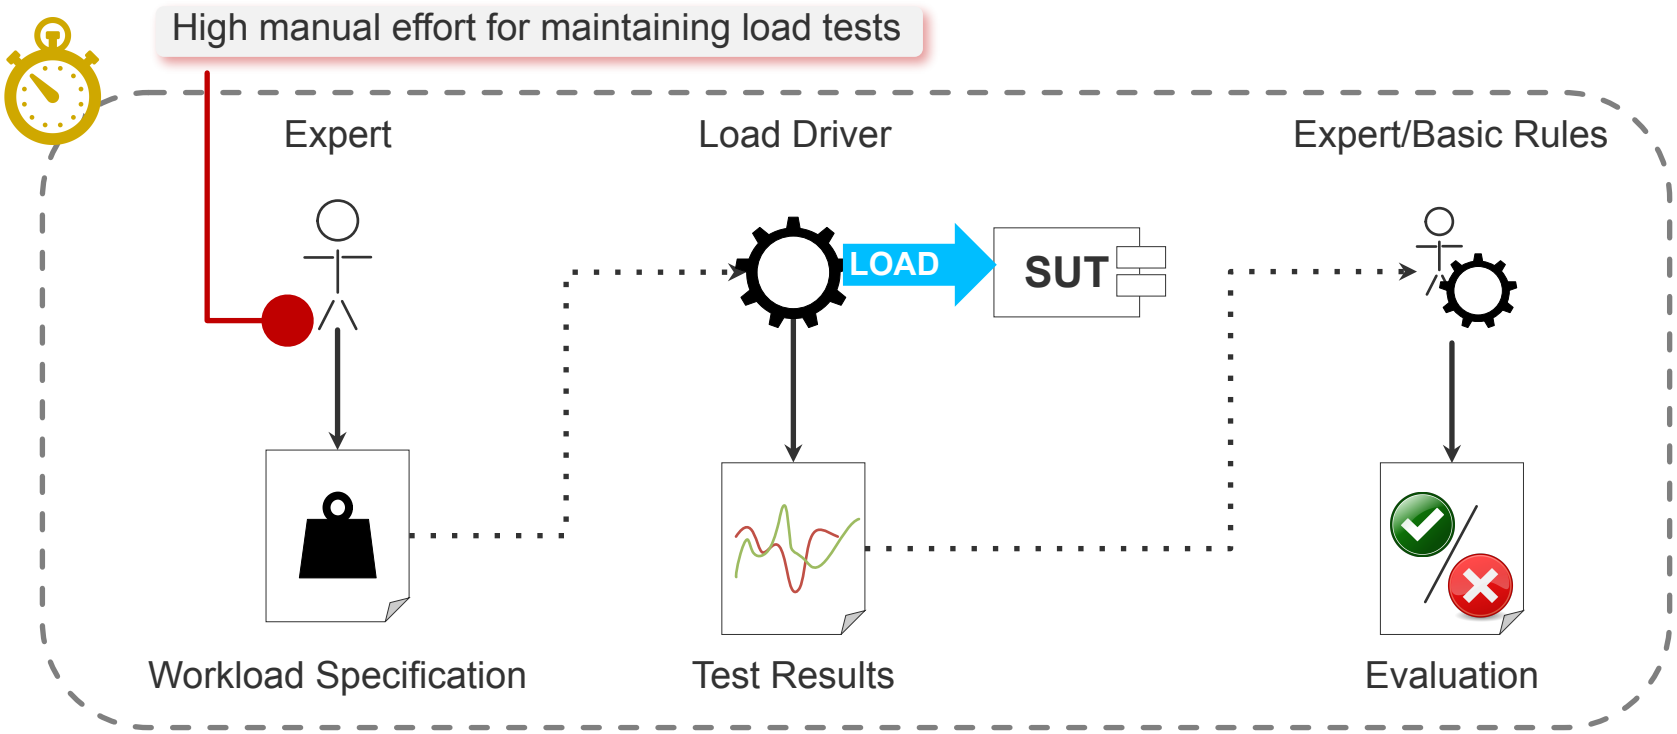

The Classic Load Testing Approach

… and Classic Problems

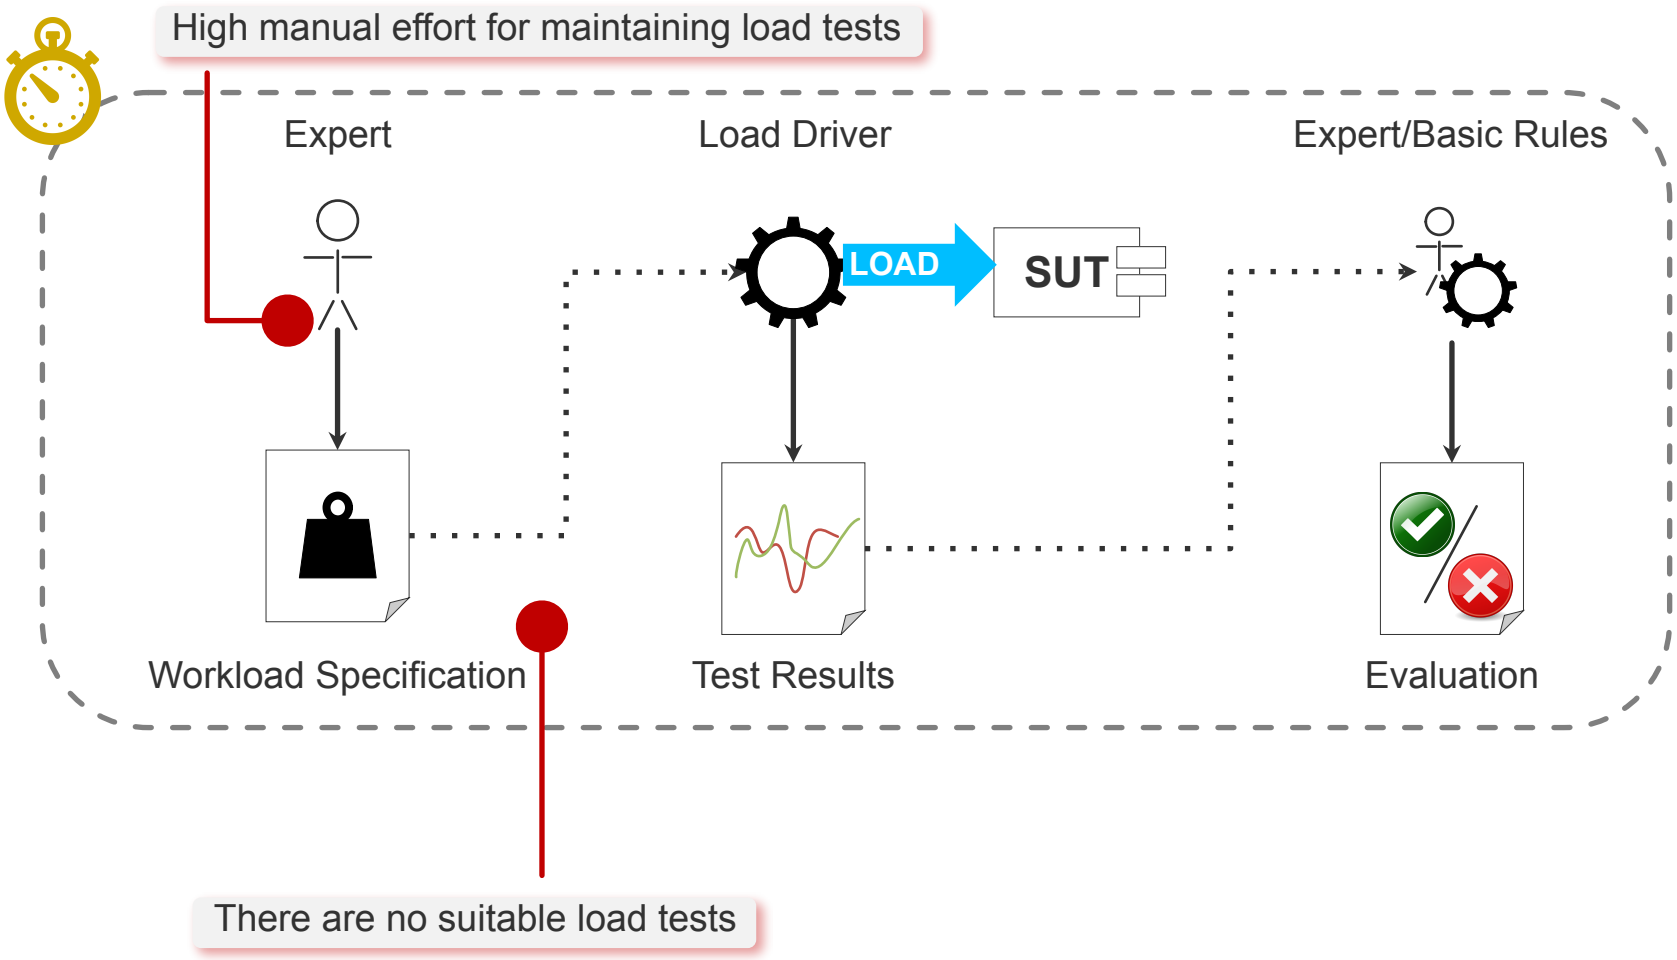

The Classic Load Testing Approach

… and Classic Problems

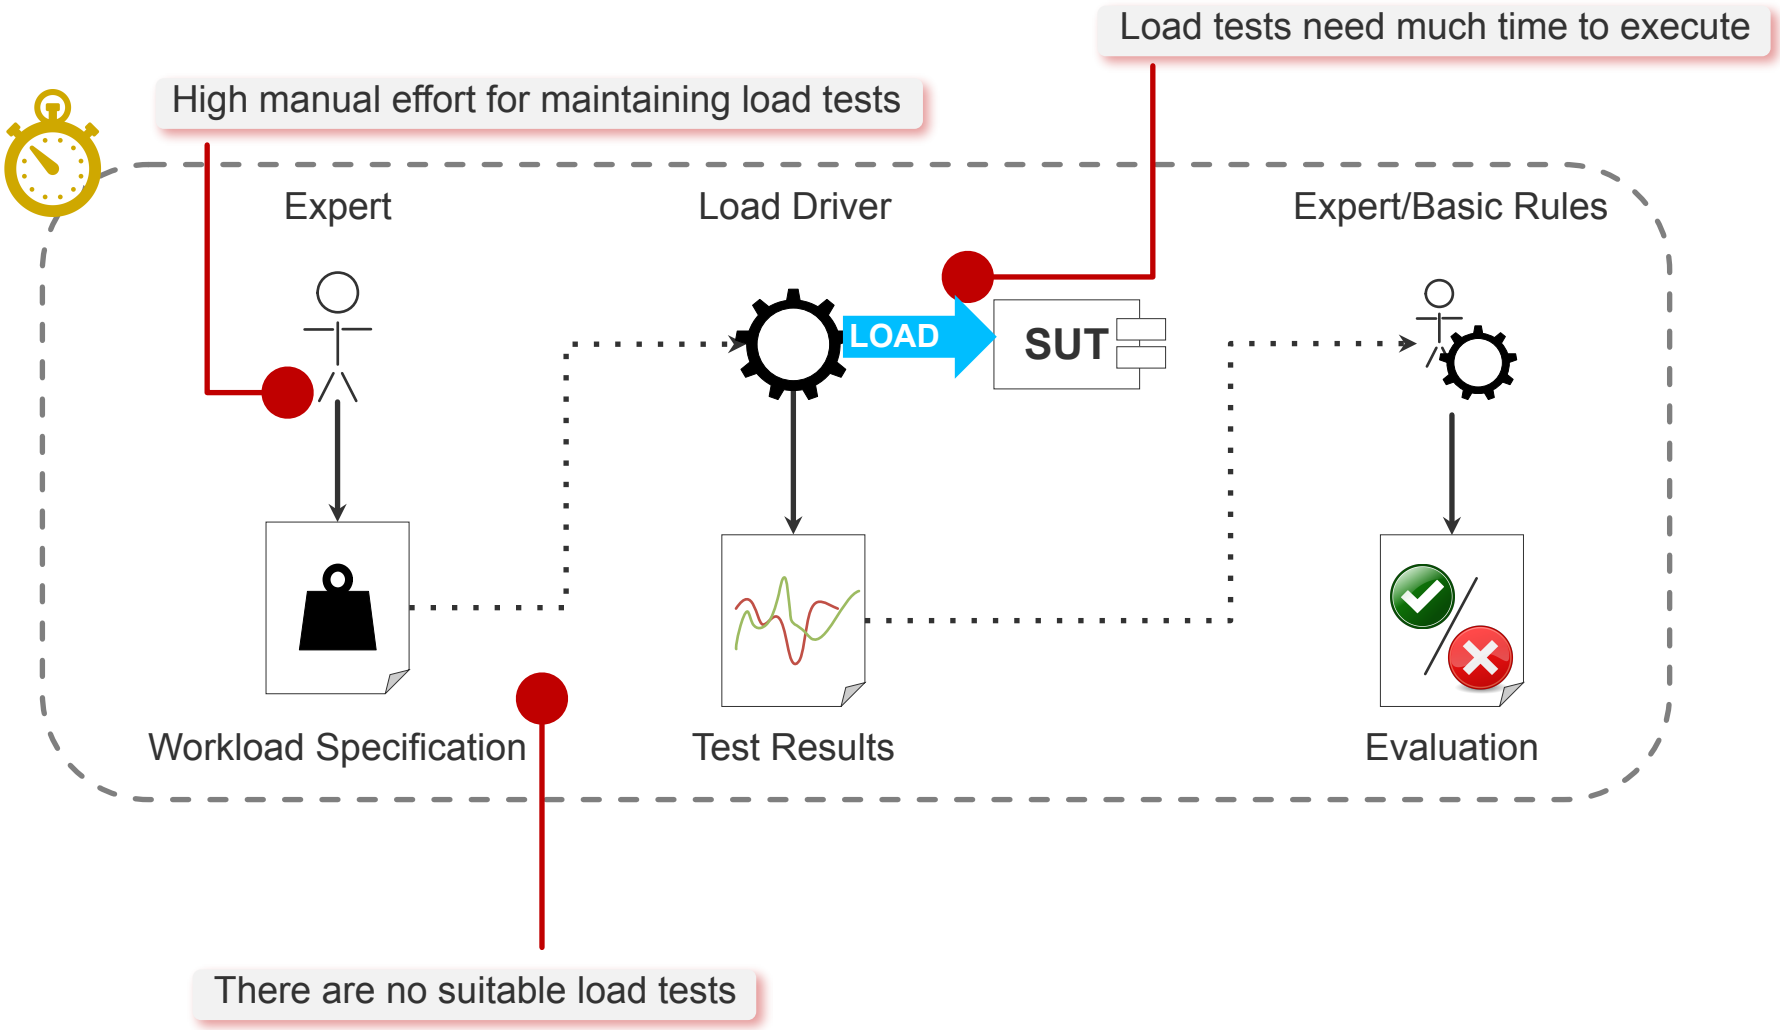

The Classic Load Testing Approach

… and Classic Problems

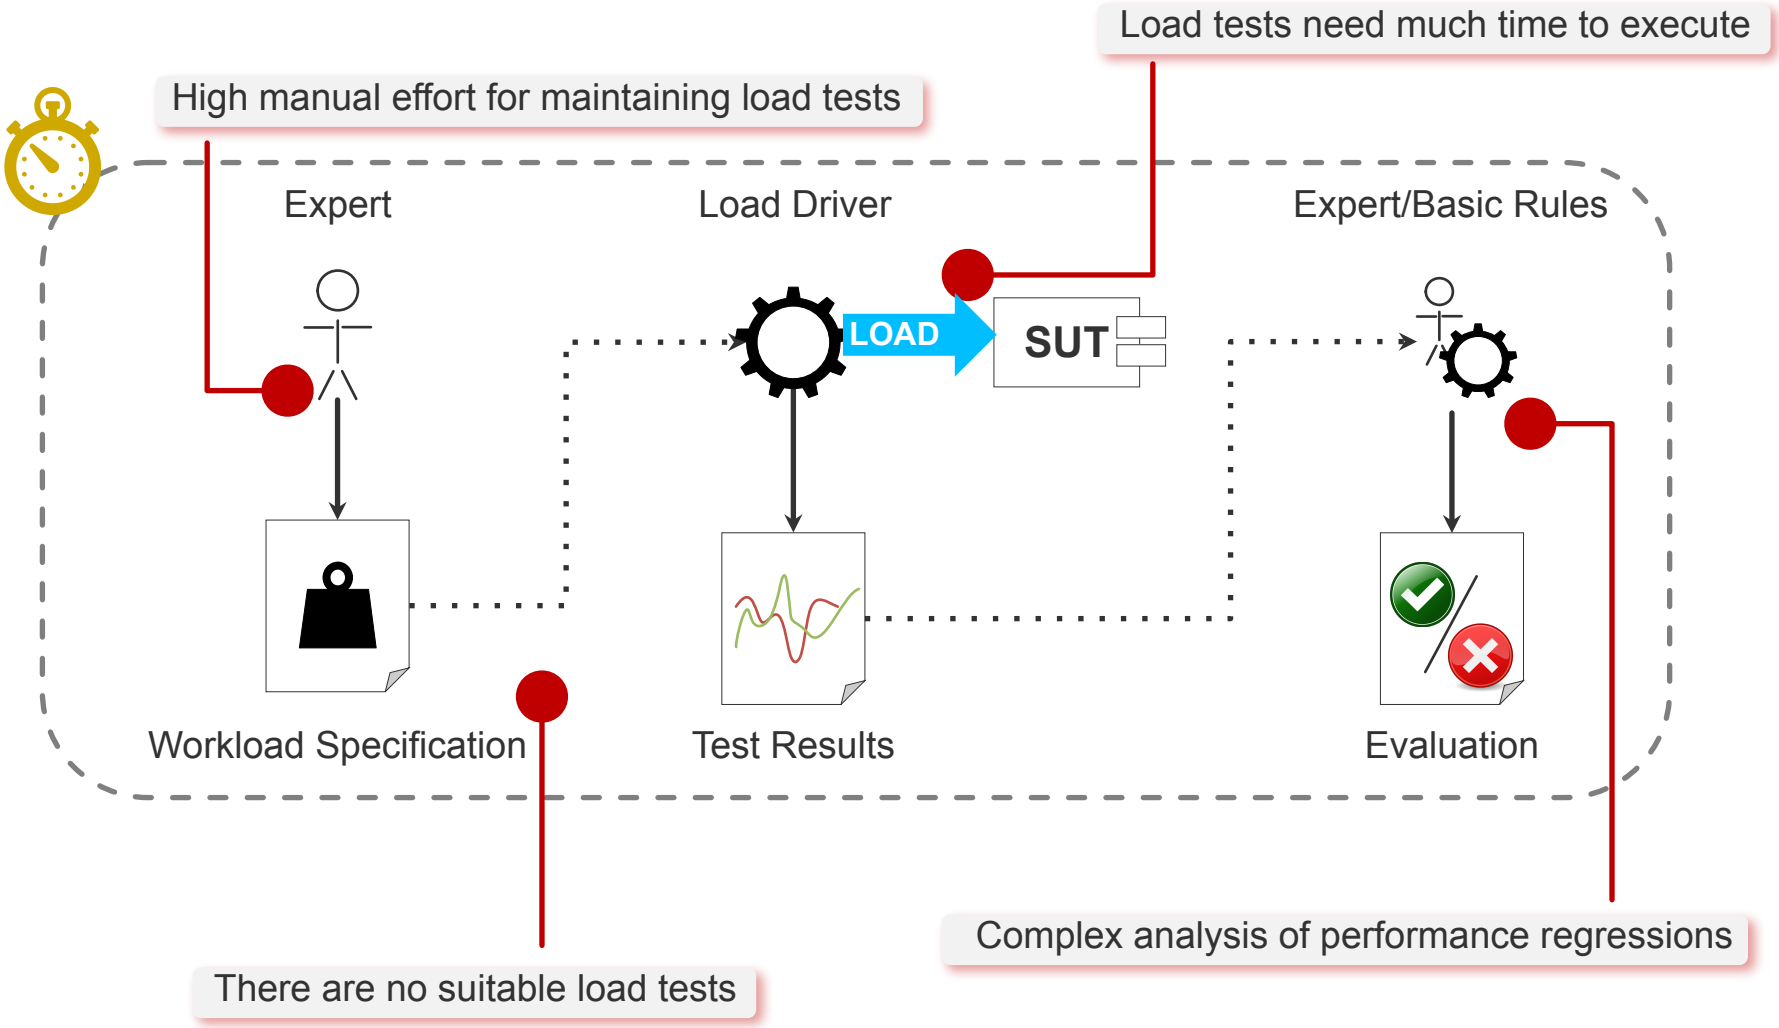

Load Testing in Continuous Delivery Pipelines

… How Problems Get Worse

Load tests need much time to execute

High manual effort for maintaining load tests

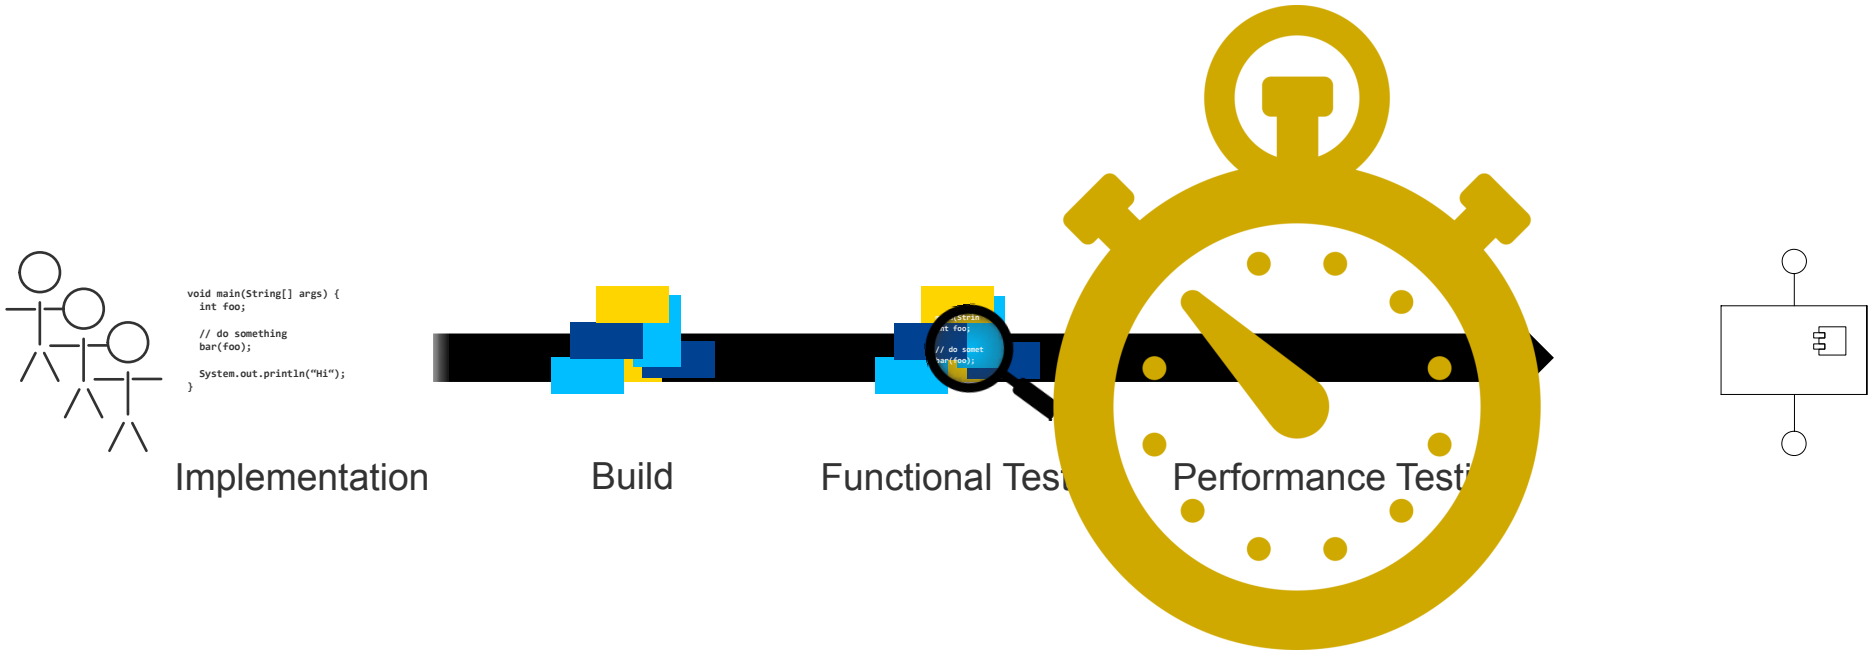

Complex analysis of performance regressions

There are no suitable load tests

Load Testing in Continuous Delivery Pipelines

… How Problems Get Worse

Load tests need much time to execute

High manual effort for maintaining load tests

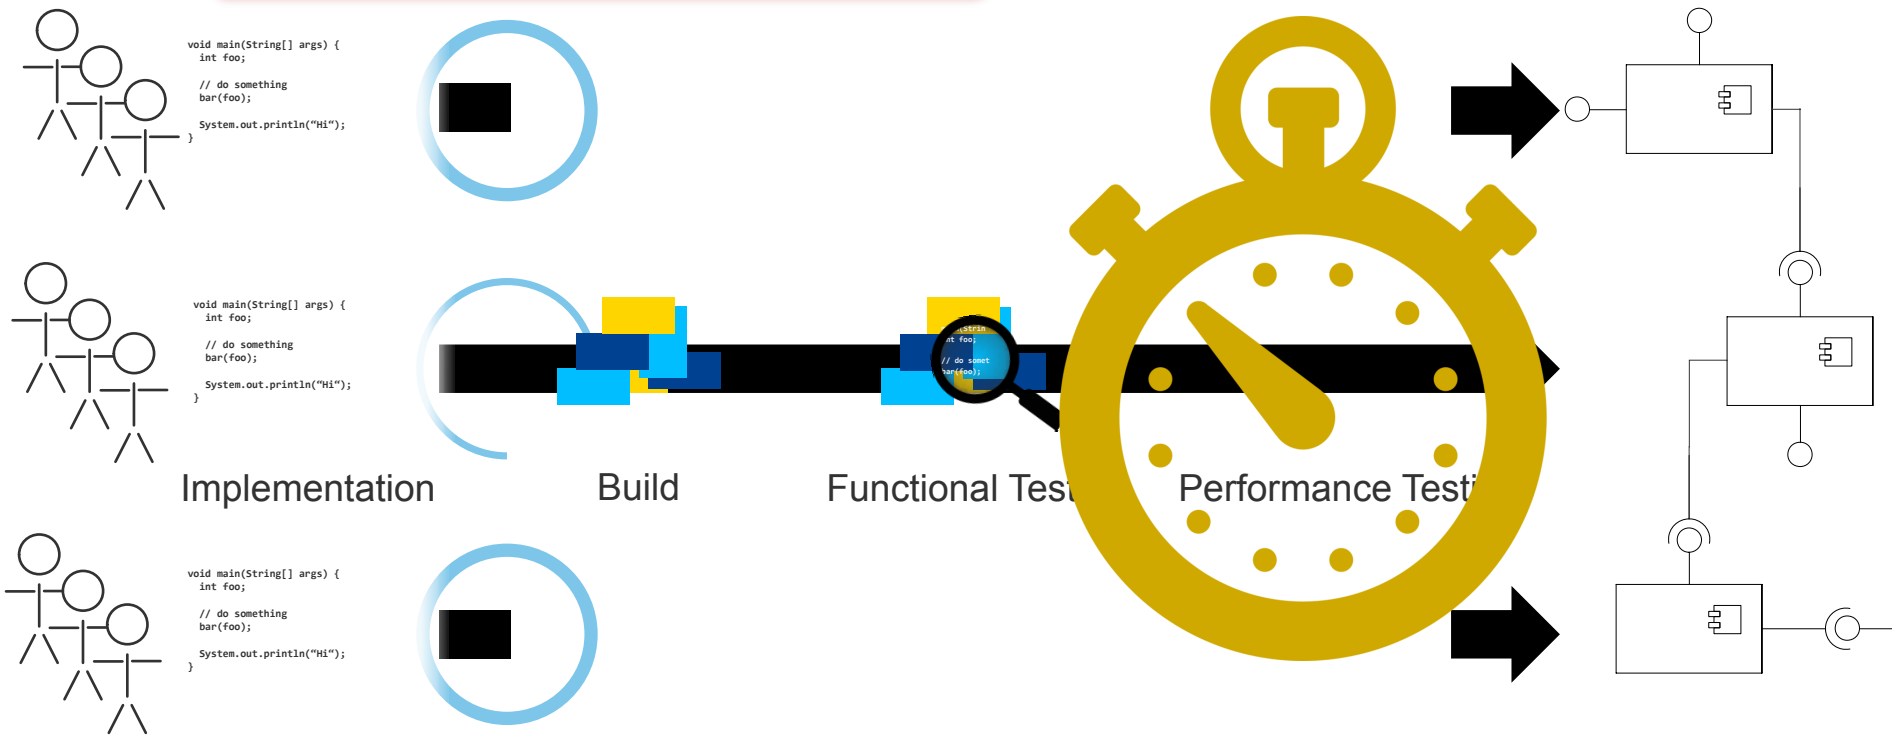

Complex analysis of performance regressions

There are no suitable load tests

Load Testing in Continuous Delivery Pipelines

… How Problems Get Worse

Load tests need much time to execute

vs.

Fast & frequent releases

High manual effort for maintaining load tests

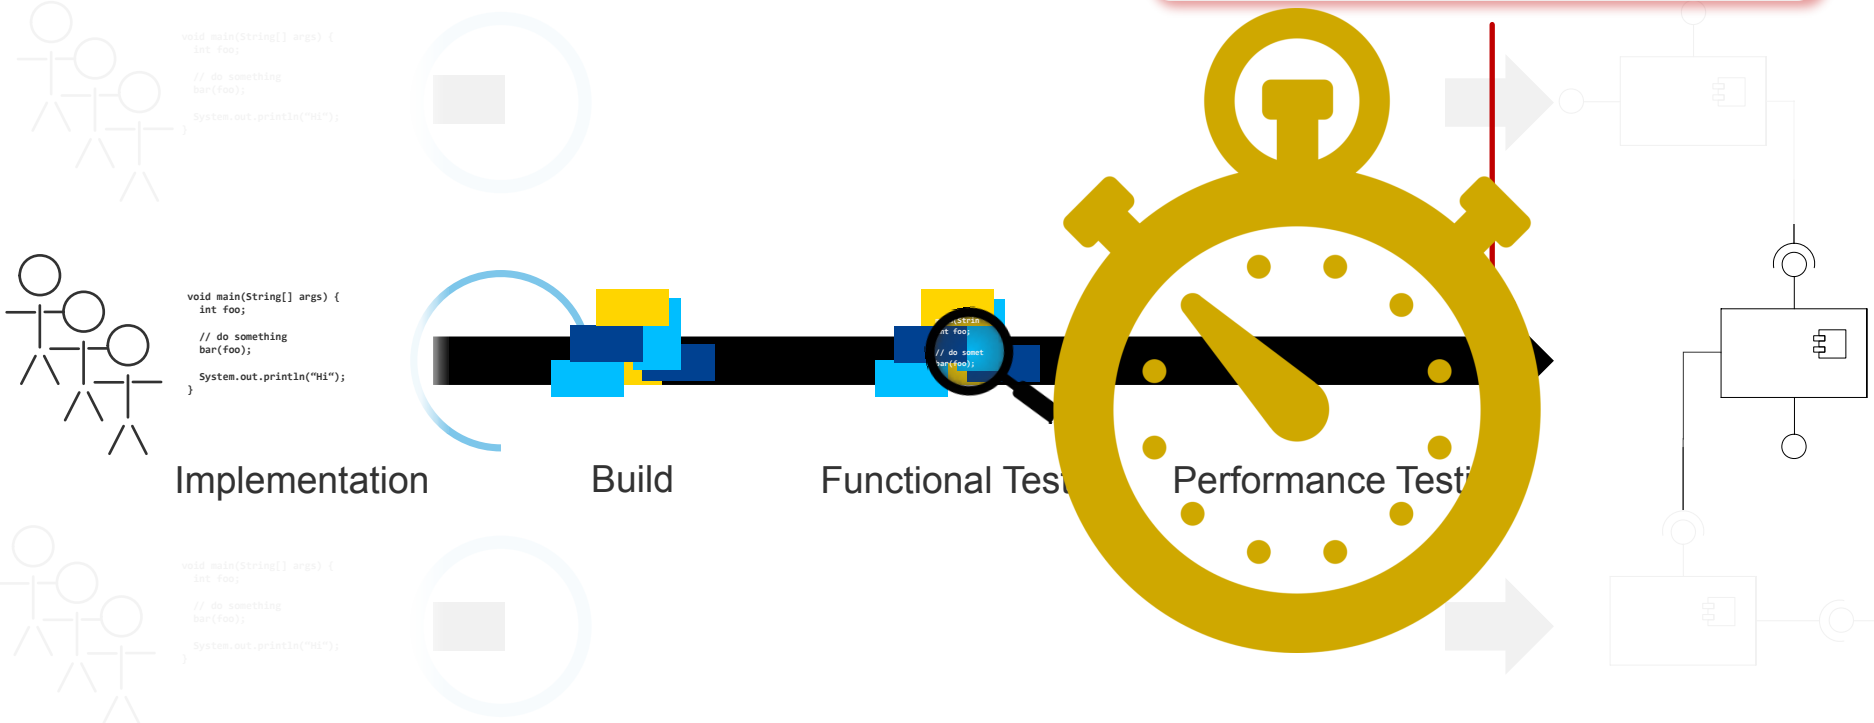

Complex analysis of performance regressions

There are no suitable load tests

Load Testing in Continuous Delivery Pipelines

… How Problems Get Worse

High manual effort for maintaining load tests

vs.

Pipeline automation

Load tests need much time to execute

vs.

Fast & frequent releases

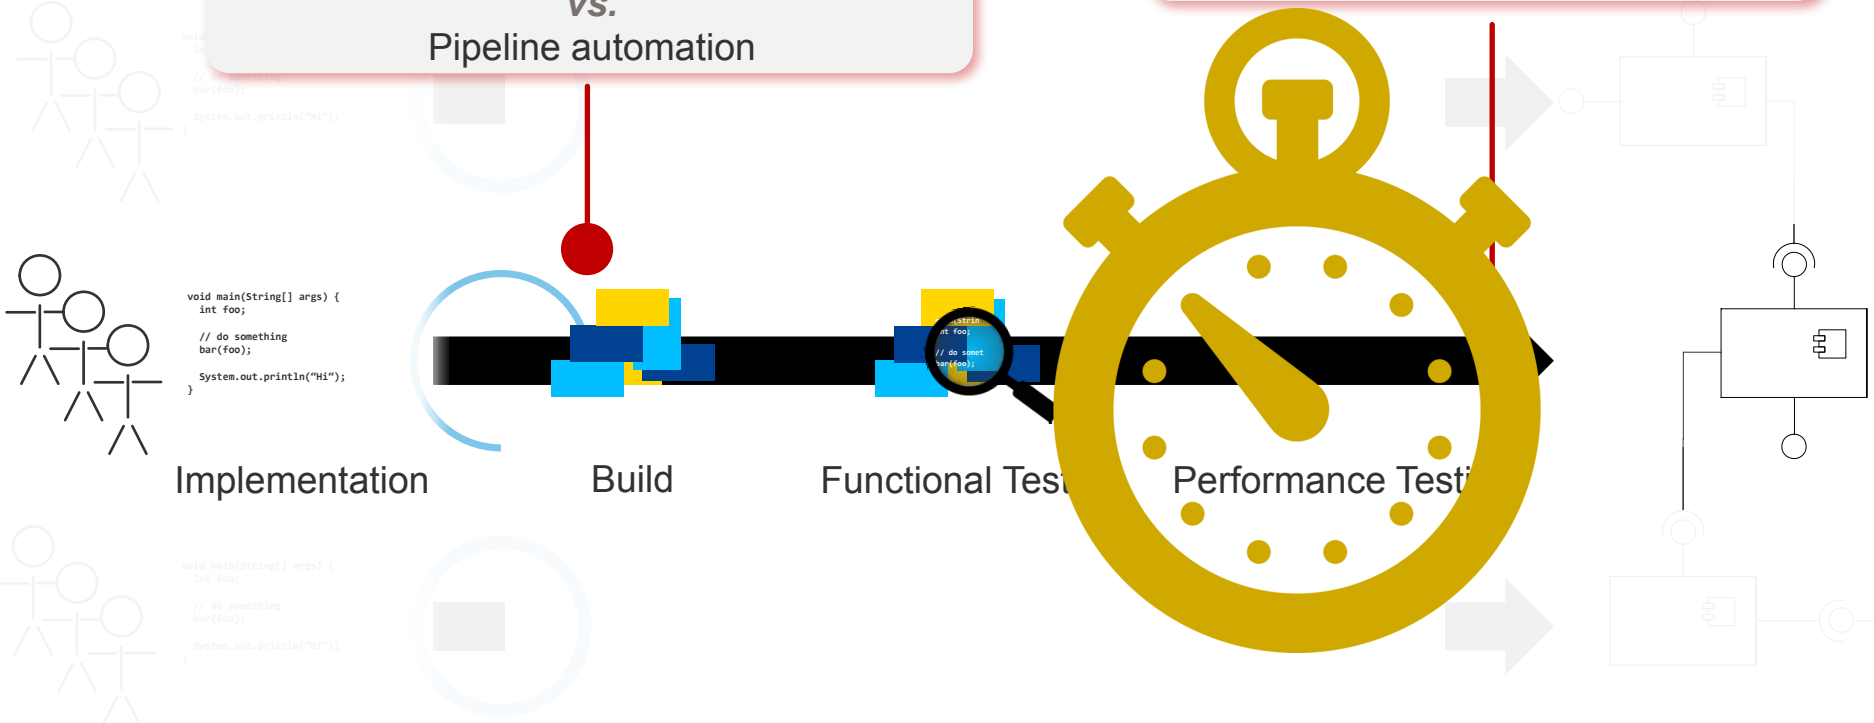

There are no suitable load tests

Complex analysis of performance regressions

Load Testing in Continuous Delivery Pipelines

… How Problems Get Worse

High manual effort for maintaining load tests

vs.

Pipeline automation

Load tests need much time to execute

vs.

Fast & frequent releases

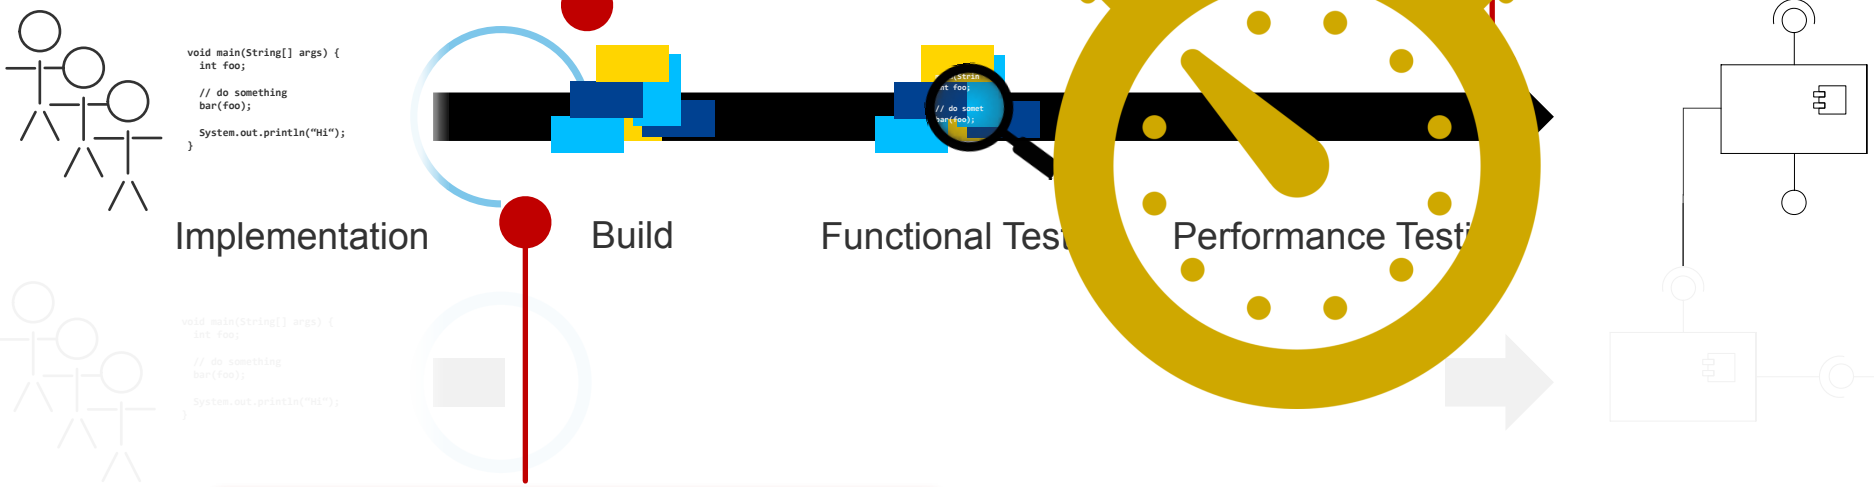

Service-focus requires multiple tests

vs.

There are no suitable load tests

Complex analysis of performance regressions

Load Testing in Continuous Delivery Pipelines

… How Problems Get Worse

High manual effort for maintaining load tests

vs.

Pipeline automation

Load tests need much time to execute

vs.

Fast & frequent releases

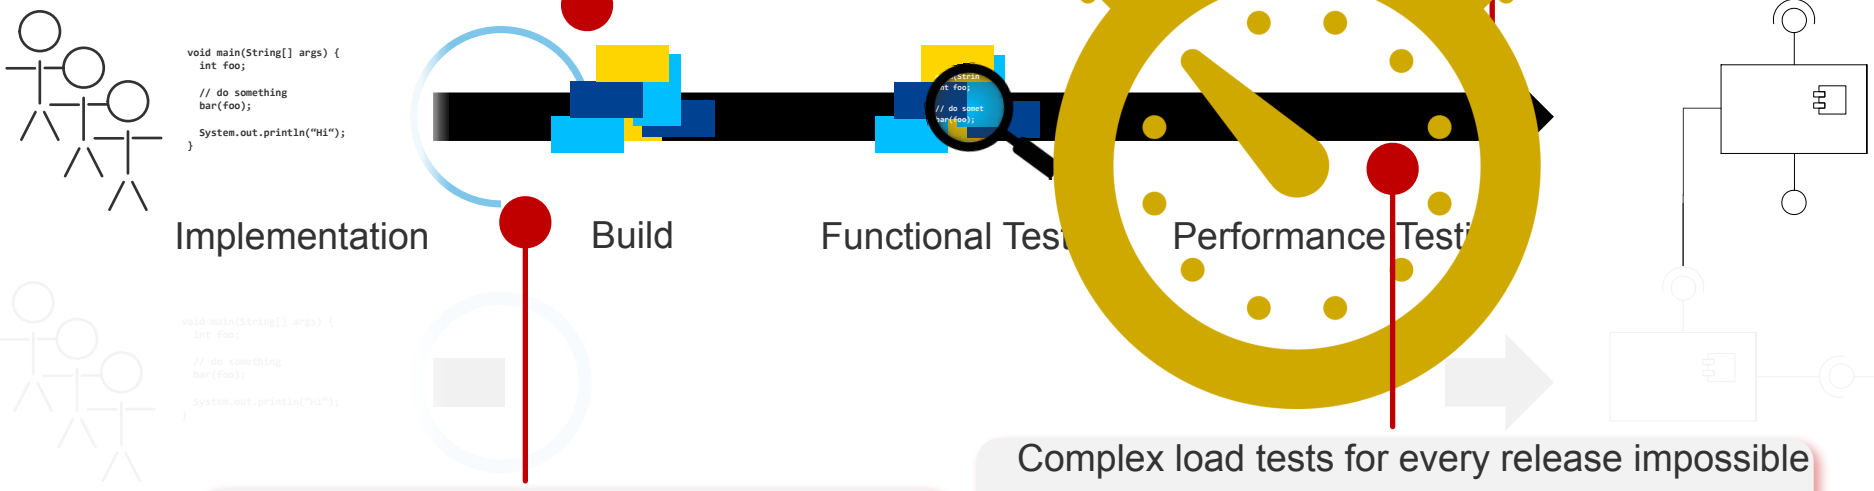

Service-focus requires multiple tests

vs.

There are no suitable load tests

Complex load tests for every release impossible

vs.

Complex analysis of performance regressions

Workload specification

- The workload specification Model consists of:
- An *Application Model*, specifying **allowed sequences of service invocations** and SUT-specific details for generating valid requests
- A set of *Behavior Models*, each providing a probabilistic representation of user sessions in terms of **invoked services and think times** between subsequent invocations as **Markov Chains**
- A *Behavior Mix*, specified as **probabilities (frequencies) for the individual Behavior Models to occur** during workload generation
- A *Workload Intensity* that includes a function which specifies the (possibly varying) **number of concurrent users** during the workload generation execution

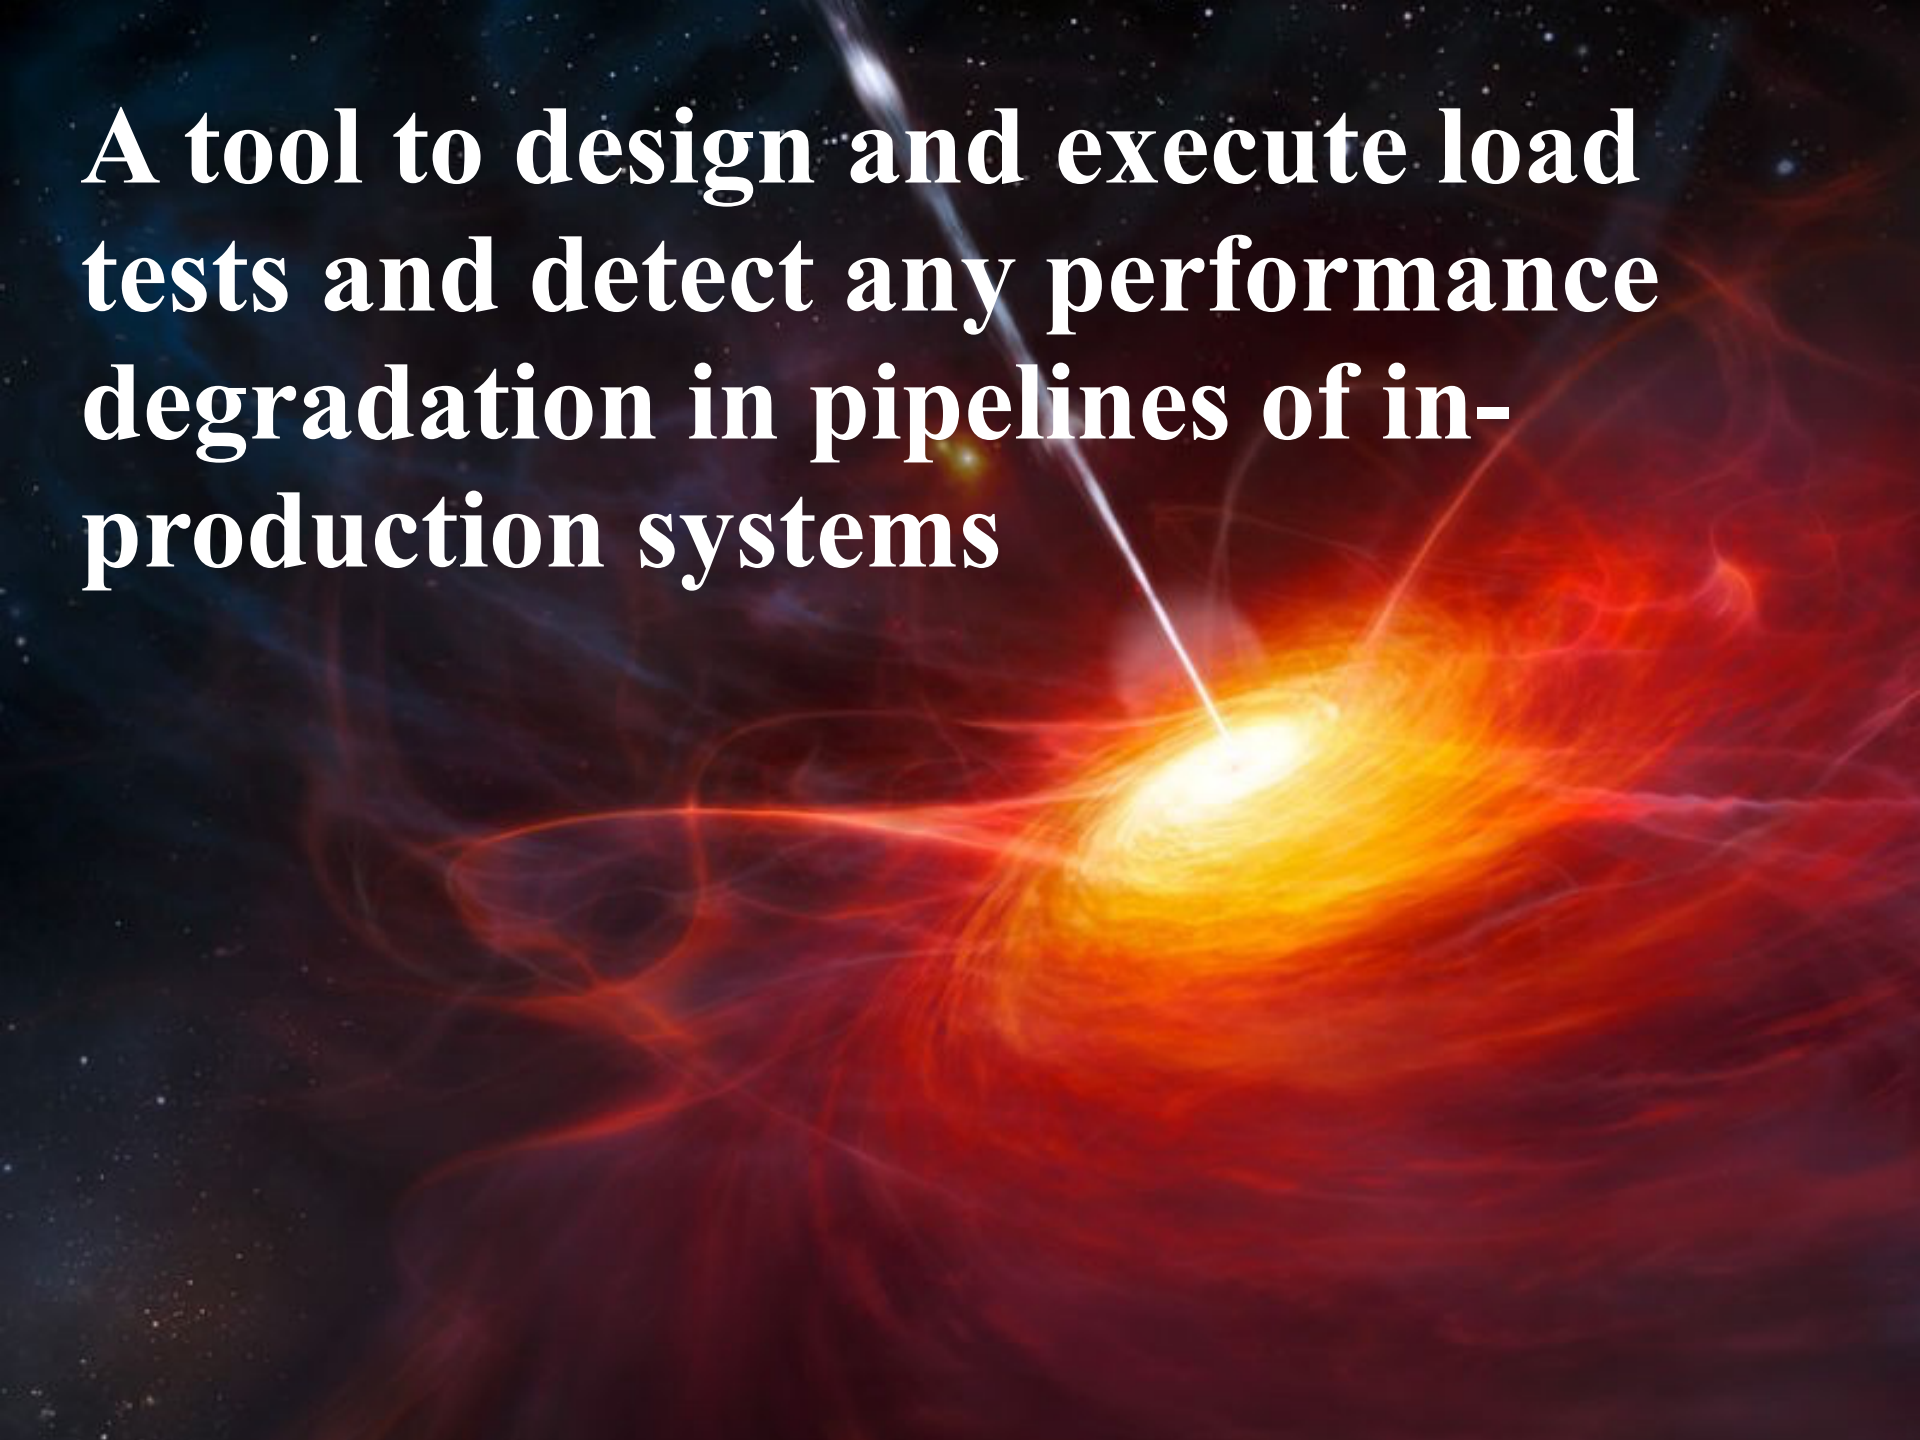
A tool to design and execute load tests and detect any performance degradation in pipelines of in-production systems

Monitoring data

- **Loads frequencies**, **Request logs, traces, and response times** of the service interfaces
- The data is enriched by various **contextual information**, e.g., marketing campaigns, public holidays, or sports events

ContinuITy

Load Test Generation & Execution

ContinuITy

Load Test Generation & Execution

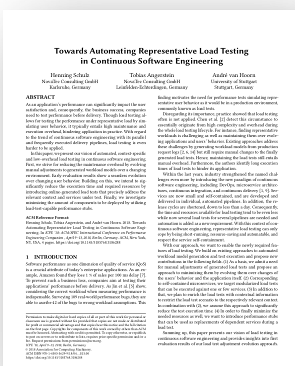
Towards Automating Representative Load Testing in Continuous Software Engineering

Henning Schulz, NovaTec Consulting GmbH, Karlsruhe, Germany
Tobias Angerstein, NovaTec Consulting GmbH, Leinfelden-Echterdingen, Germany
André van Hoorn, University of Stuttgart, Stuttgart, Germany

Henning Schulz, Tobias Angerstein, and André van Hoorn:
Towards automating representative load testing in continuous software engineering.

In Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE 2018) Companion (7th International Workshop on Load Testing and Benchmarking of Software Systems, LTB 2018), pages 123–126. ACM, 2018

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

ContinuITy

Load Test Generation & Execution

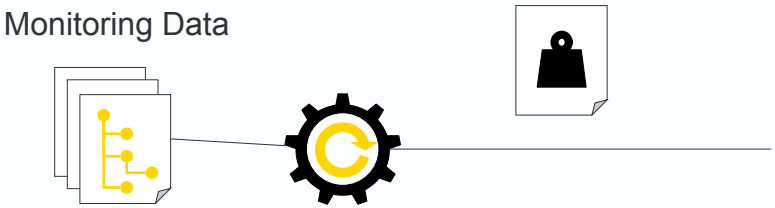

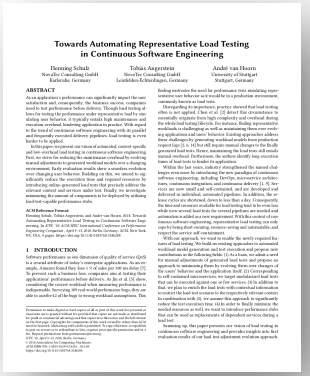

Towards Automating Representative Load Testing in Continuous Software Engineering

Henning Schulz, Tobias Angerstein, and André van Hoorn: *Towards automating representative load testing in continuous software engineering*.

In Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE 2018) Companion (7th International Workshop on Load Testing and Benchmarking of Software Systems, LTB 2018), pages 123–126. ACM, 2018

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

ContinuITy

Load Test Generation & Execution

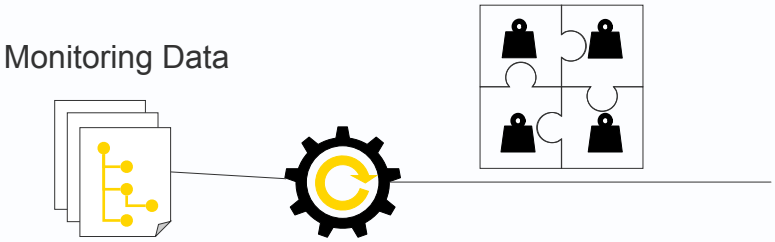

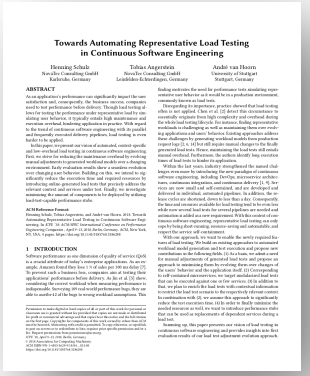

Towards Automating Representative Load Testing in Continuous Software Engineering

Henning Schulz, Tobias Angerstein, and André van Hoorn: *Towards automating representative load testing in continuous software engineering*.

In Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE 2018) Companion (7th International Workshop on Load Testing and Benchmarking of Software Systems, LTB 2018), pages 123–126. ACM, 2018

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

ContinuITy

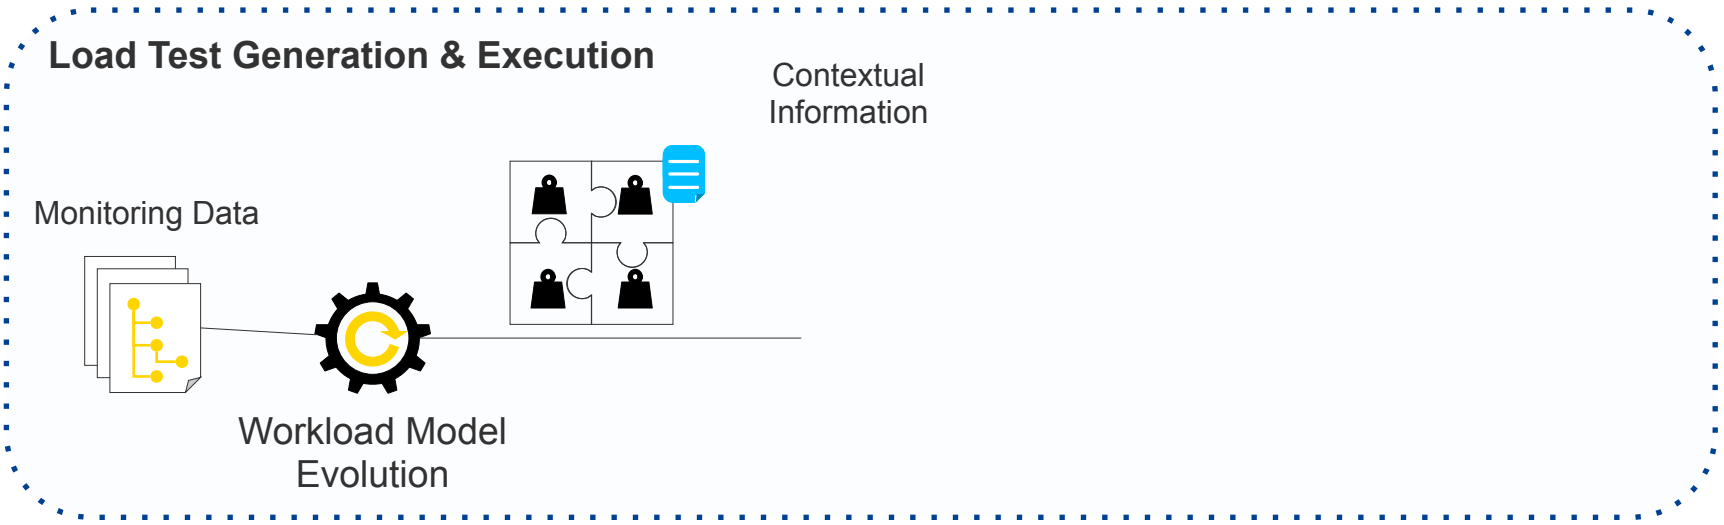

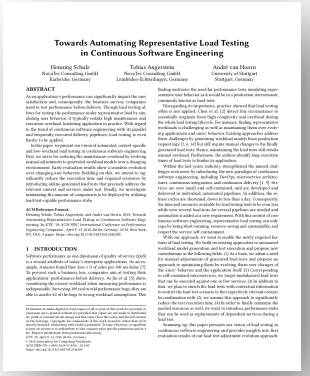

Towards Automating Representative Load Testing in Continuous Software Engineering

Henning Schulz, Tobias Angerstein, and André van Hoorn: *Towards automating representative load testing in continuous software engineering*.

In Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE 2018) Companion (7th International Workshop on Load Testing and Benchmarking of Software Systems, LTB 2018), pages 123–126. ACM, 2018

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

ContinuITy

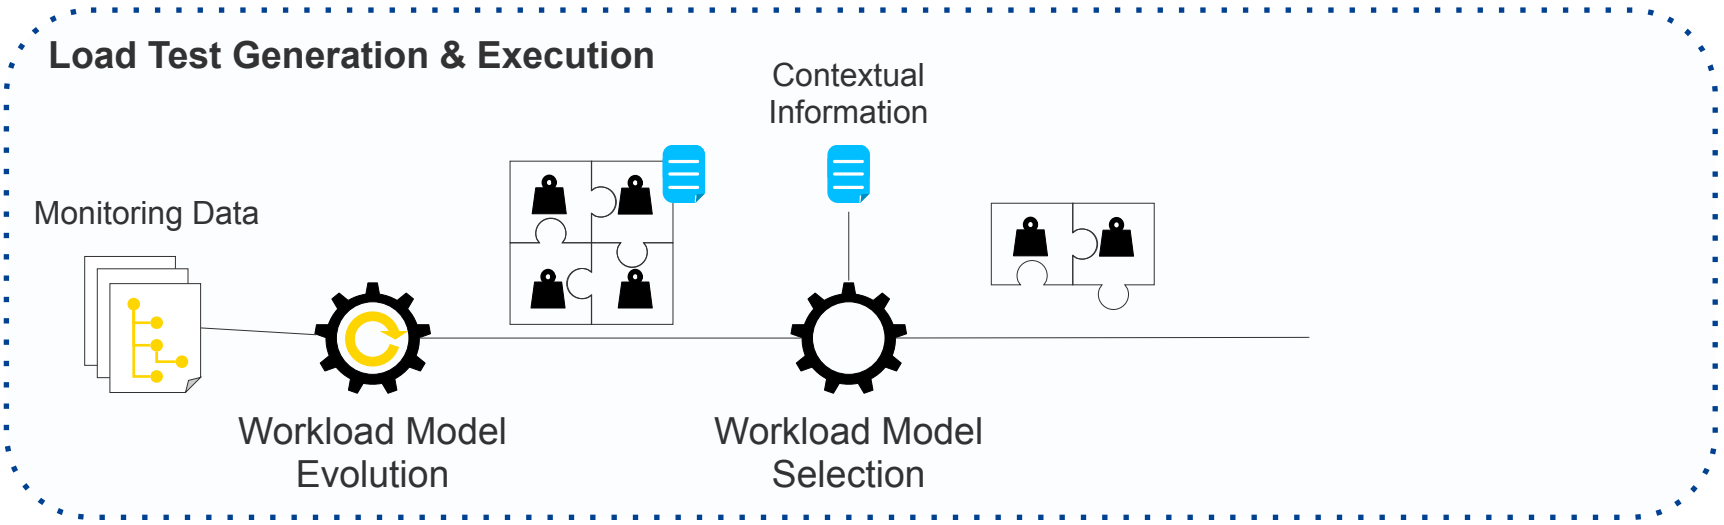

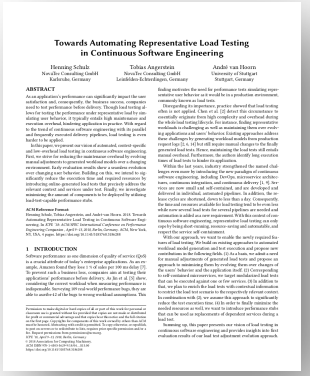

Towards Automating Representative Load Testing in Continuous Software Engineering

Henning Schulz, Tobias Angerstein, and André van Hoorn: *Towards automating representative load testing in continuous software engineering*.

In Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE 2018) Companion (7th International Workshop on Load Testing and Benchmarking of Software Systems, LTB 2018), pages 123–126. ACM, 2018

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

ContinuITy

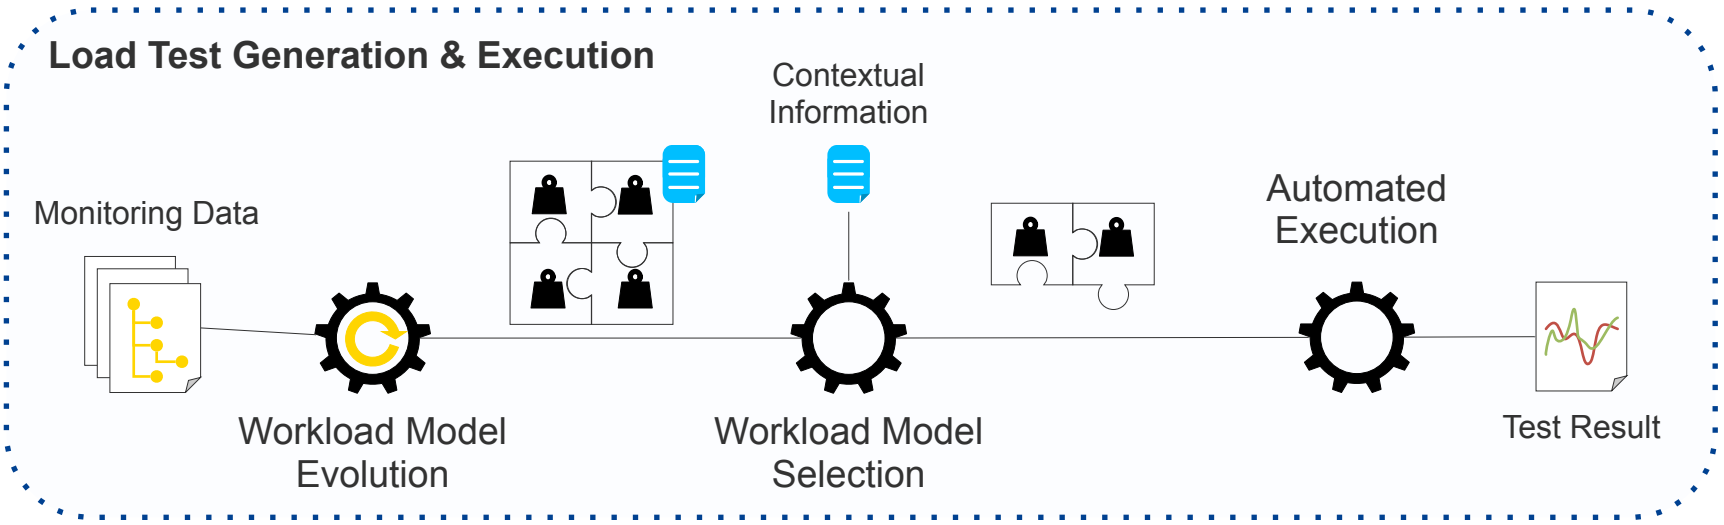

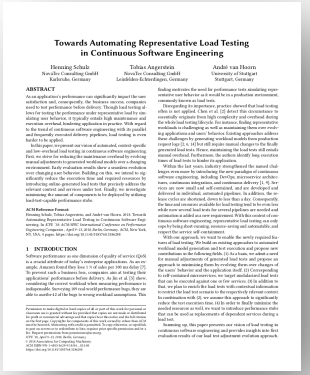

Towards Automating Representative Load Testing in Continuous Software Engineering

Henning Schulz, Tobias Angerstein, and André van Hoorn: *Towards automating representative load testing in continuous software engineering*.

In Proceedings of the ACM/SPEC International Conference on Performance Engineering (ICPE 2018) Companion (7th International Workshop on Load Testing and Benchmarking of Software Systems, LTB 2018), pages 123–126. ACM, 2018

unibz
Freie Universität Bozen
Libera Università di Bolzano
Università Liedia de Bulsan

Load test automation

Tools and Techniques for Software Testing - Barbara Russo

SwSE - Software and Systems Engineering group

What is it for?

- Identify deployment configuration(s) for which the system performs best for all workloads
- Characterize systems‘ resilience over workload
- Analyse individual service failure or degradation
- Reveal attacks
- Monitor the performance in a transition to microservices

PPTAM

Production and Performance Testing-Based Application Monitoring

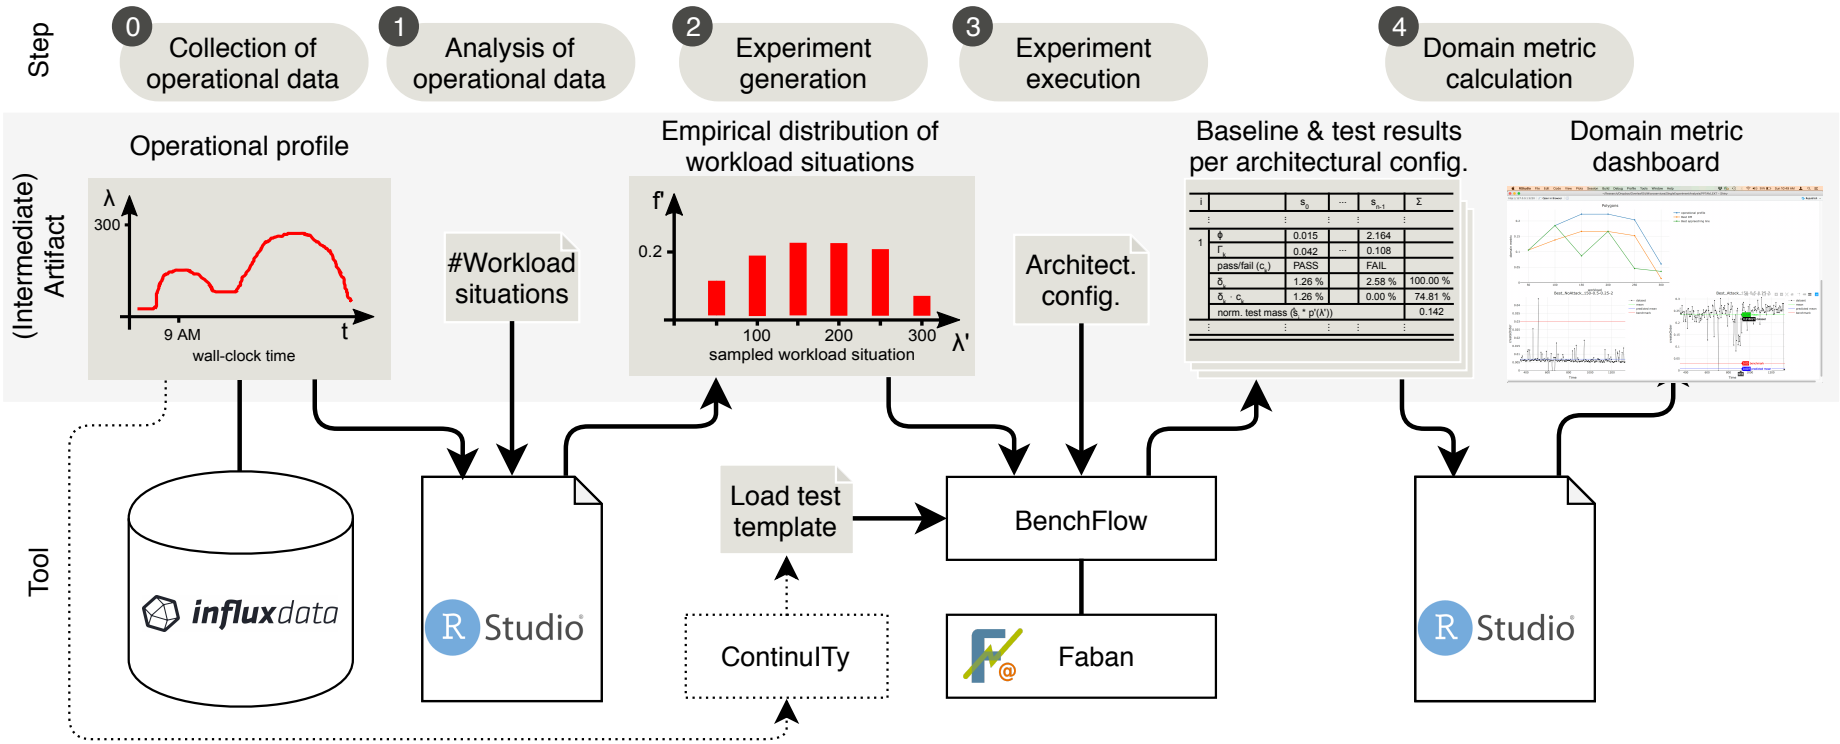

unibz

What it does?

- It **collects the operational profile** of a microservice system
- It **automatically runs a series of experiments** with given usage profiles and probability of use (R scripts)
- It **runs experiments** according to templates that control time, number of agents, and single operation max response (Faban)
- It sets the **goals of testing (e.g., resource config)** and **collects data for each experiment** or multi-experiment (Benchflow)
- It **identifies failing services** (baseline threshold R scripts)
- It **computes a total metric of performance** on non-failing services for each system configuration (CPU, memory, replicas) over workloads (via R scripts)
- It visualizes into interactive graphs (R shop, R shiny, and R plottly)

What it does

- Collects the operational data of systems
- Builds the operational profile

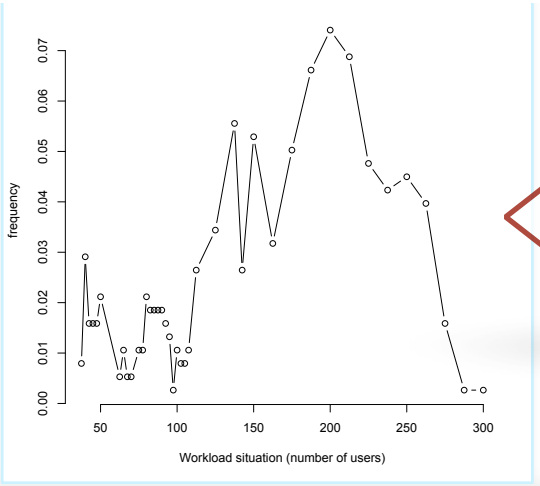

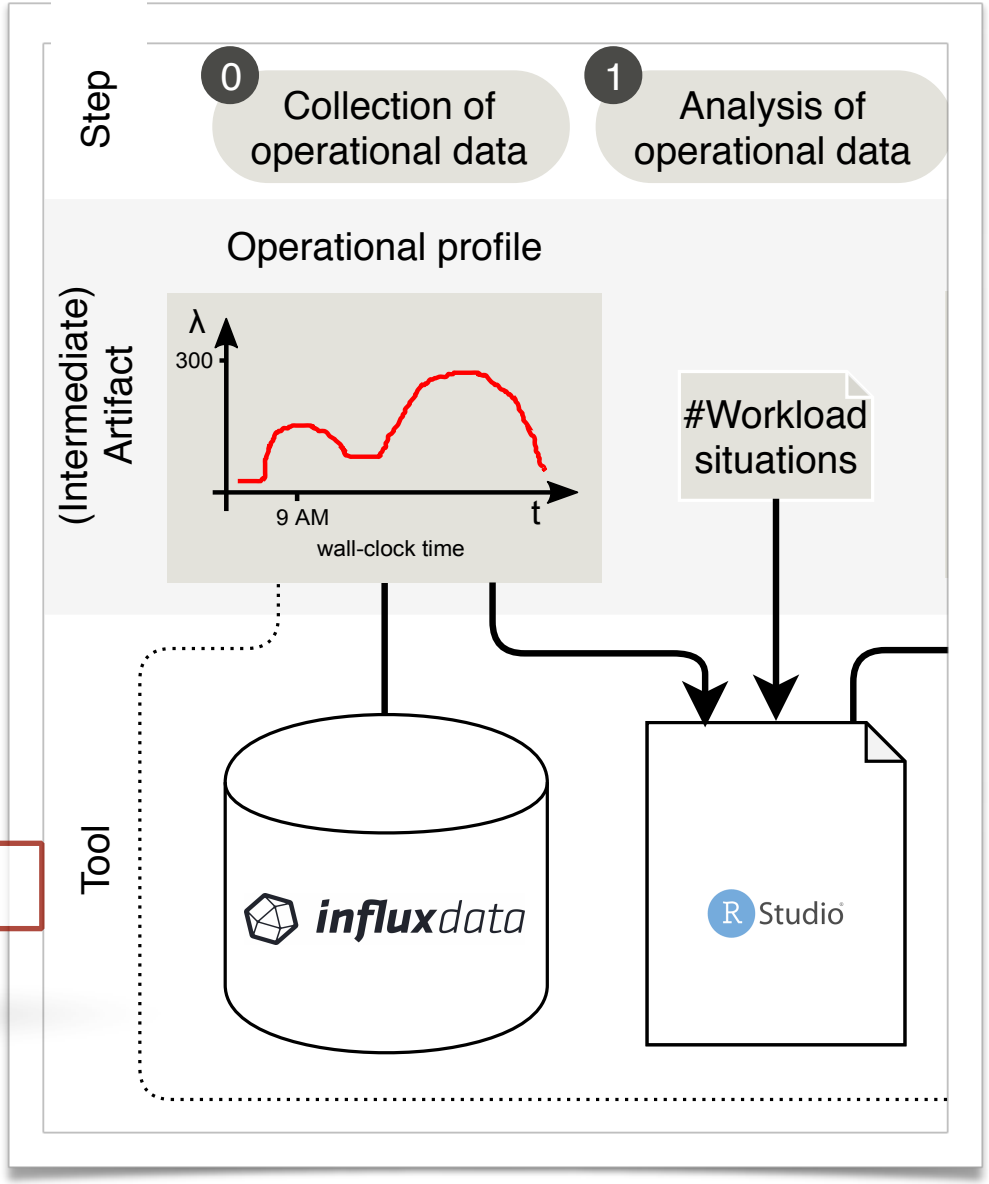

What it does

- Automatically runs a series of experiments with given usage profiles and probability of use

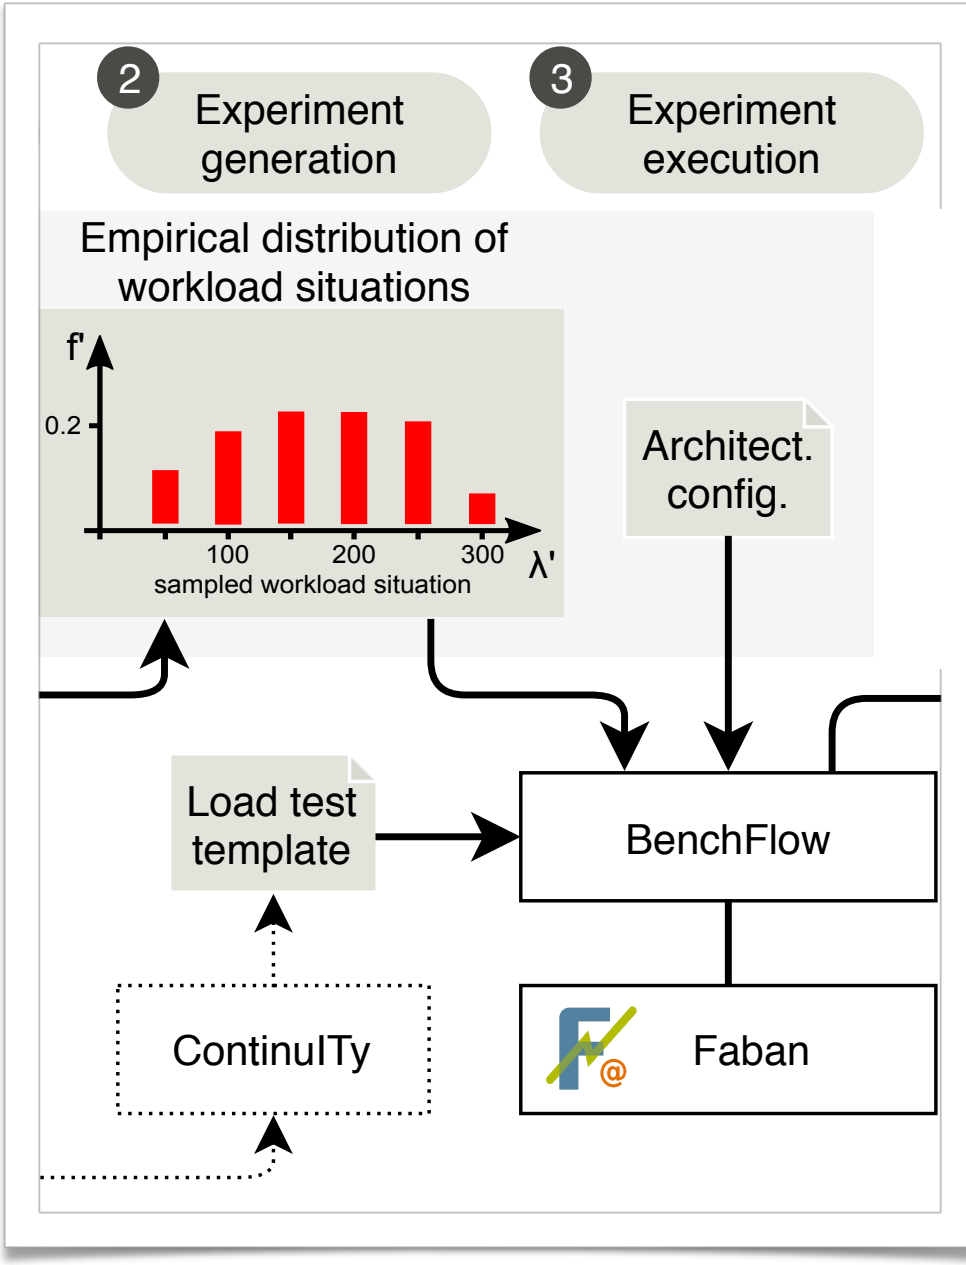

BenchFlow Automation Framework

BenchFlow Automation Framework

BenchFlow Automation Framework

BenchFlow Automation Framework

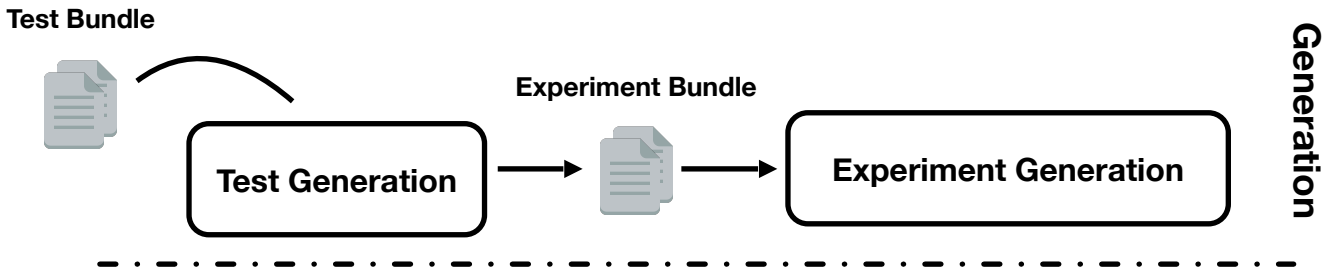

BenchFlow Automation Framework

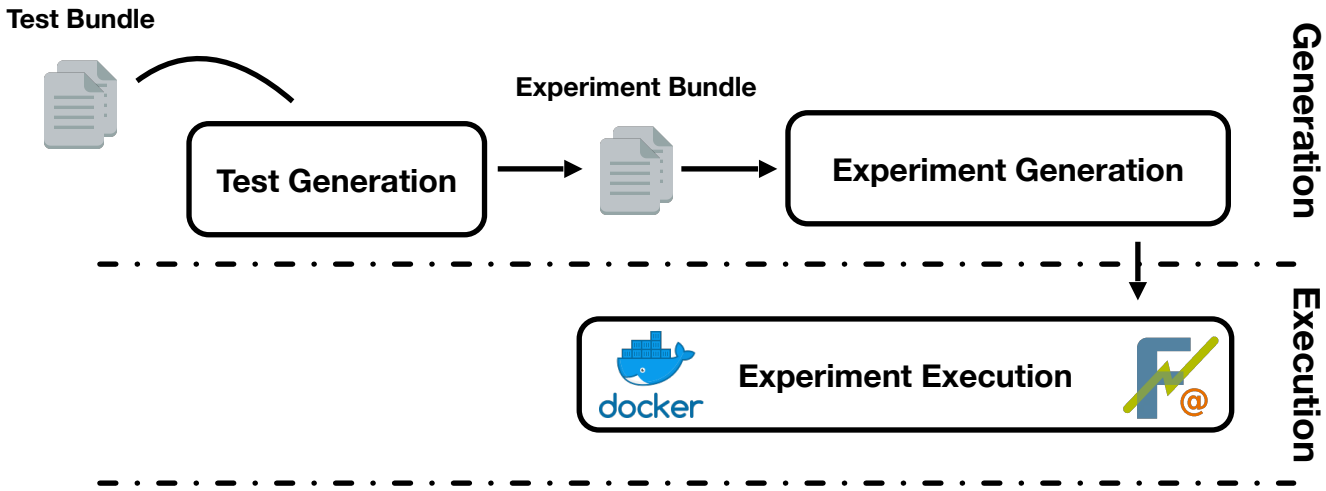

BenchFlow Automation Framework

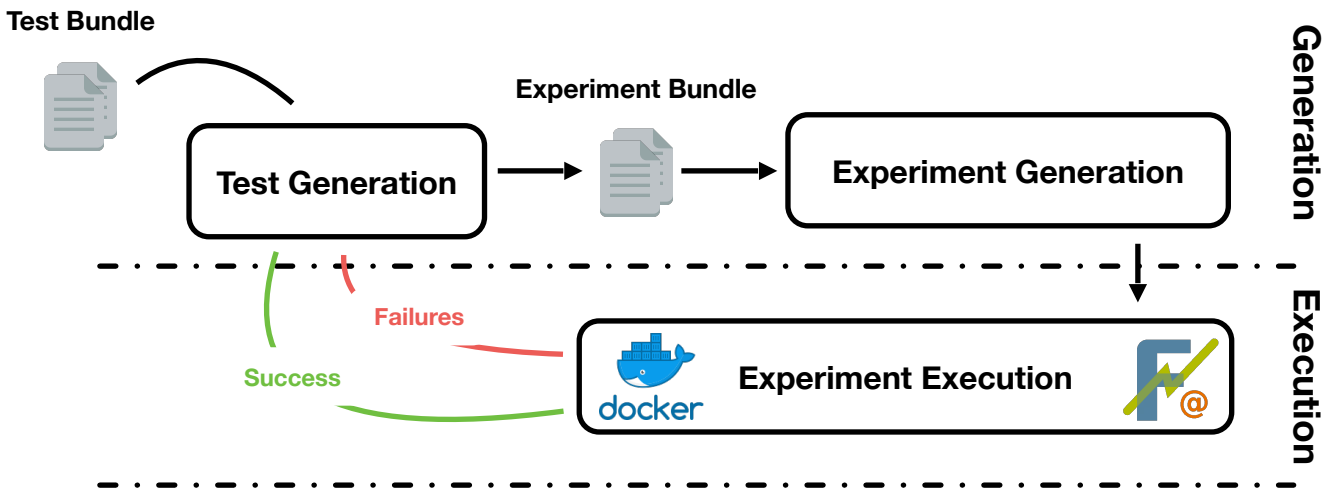

BenchFlow Automation Framework

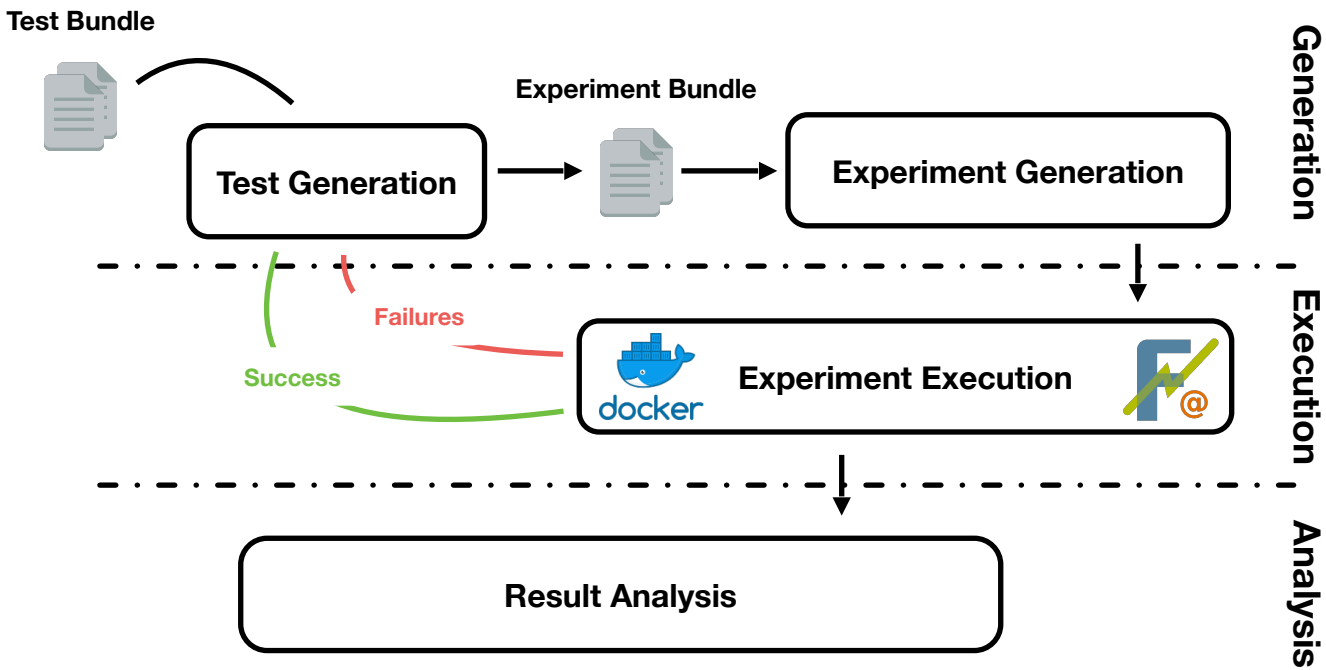

BenchFlow Automation Framework

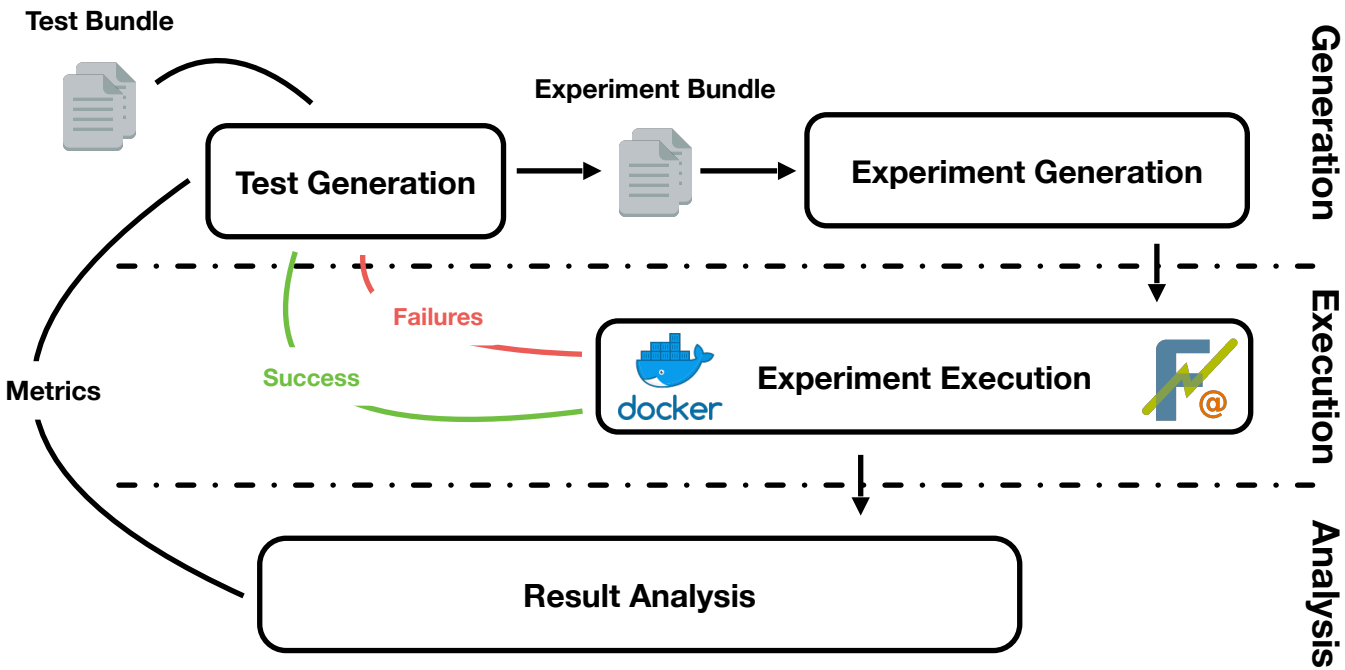

Experiment execution design (faban)

- Usage profiles (WebDriver.java)

```
30 @FlatSequenceMix (
31     mix = { 40, 30, 30 },
32     sequences = {
33         // # Visitor: View Home -> View Catalogue -> View Details
34         @OperationSequence({"home", "getCatalogue", "getCart", "home", "getCatalogue", "getCart", "catalogue", "catalogueSize",
35             "tags", "cataloguePage", "getCart", "getCustomer", "showDetails", "getItem", "getCustomer", "getCart", "getRelated"}),
36         // # Buyer: View Home -> Login -> View Catalogue -> View Details -> Add to Cart -> View Cart -> Create order
37         @OperationSequence({"home", "getCatalogue", "getCart", "login", "home", "getCatalogue", "getCart", "home", "getCatalogue",
38             "getCart", "catalogue", "catalogueSize", "tags", "cataloguePage", "getCart", "getCustomer", "showDetails", "getItem",
39             "getCustomer", "getCart", "getRelated", "addToCart", "showDetails", "getItem", "getCustomer", "getCart", "getRelated",
40             "basket", "getCart", "getCard", "getAddress", "getCatalogue", "getItem", "getCart", "getCustomer", "getItem", "createOrder",
41         // # Orders visitor: View Home -> Login -> View orders
42         @OperationSequence({"home", "getCatalogue", "getCart", "login", "home", "getCatalogue", "getCart", "viewOrdersPage",
43             "getOrders", "getCart", "getCustomer", "getItem"})
44     },
```

- Configuration file.xml

```
<!-- The rampup, steadystate, and rampdown of the driver -->
  <fa:runControl unit="time">
    <fa:rampUp>60</fa:rampUp>
    <fa:steadyState>1800</fa:steadyState>
    <fa:rampDown>0</fa:rampDown>
  </fa:runControl>
```

```
<!-- The number of agents, or host:agents pairs
        separated by space -->
<agents>10</agents>
```

What it does

- Service failures over time
- Total performance of system in use
- Per (micro)service performance time series
- Performance degradation under an attack

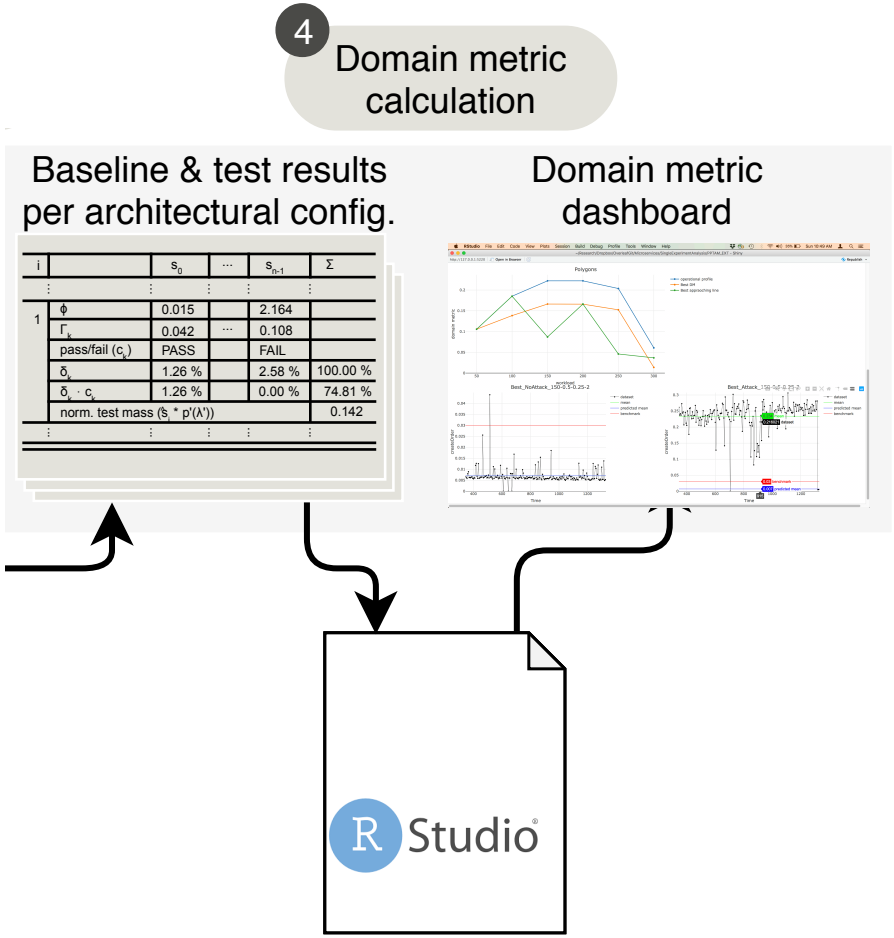

unibz

https://pptam.shinyapps.io/PPTAM_EXT/

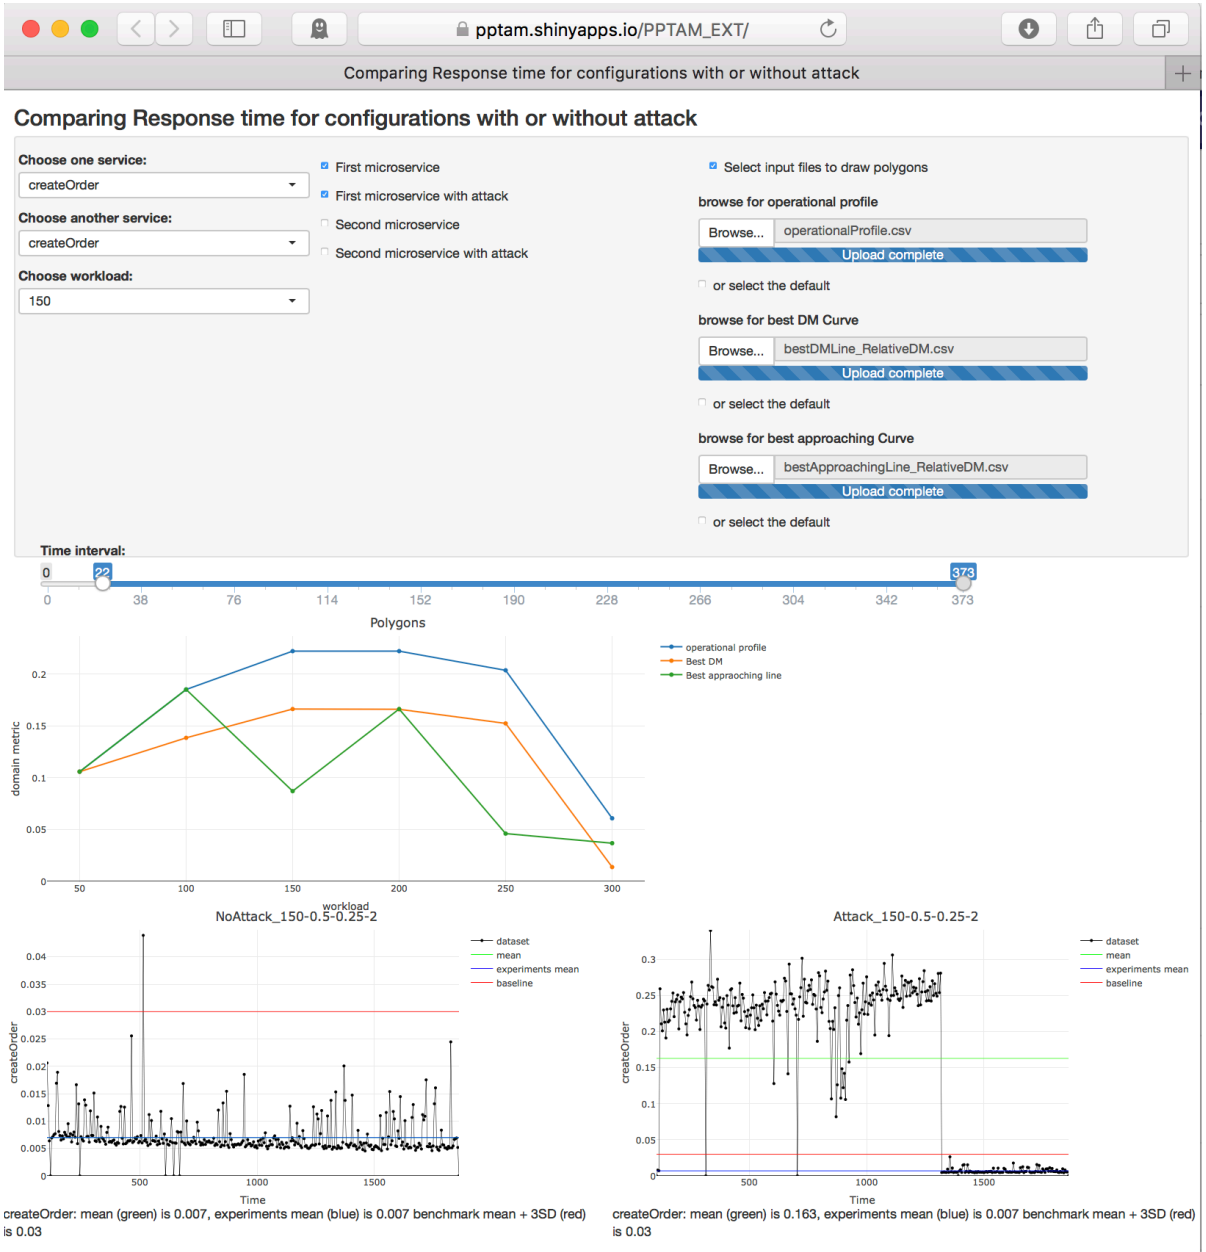

unibz

https://pptam.shinyapps.io/PPTAM_EXT/

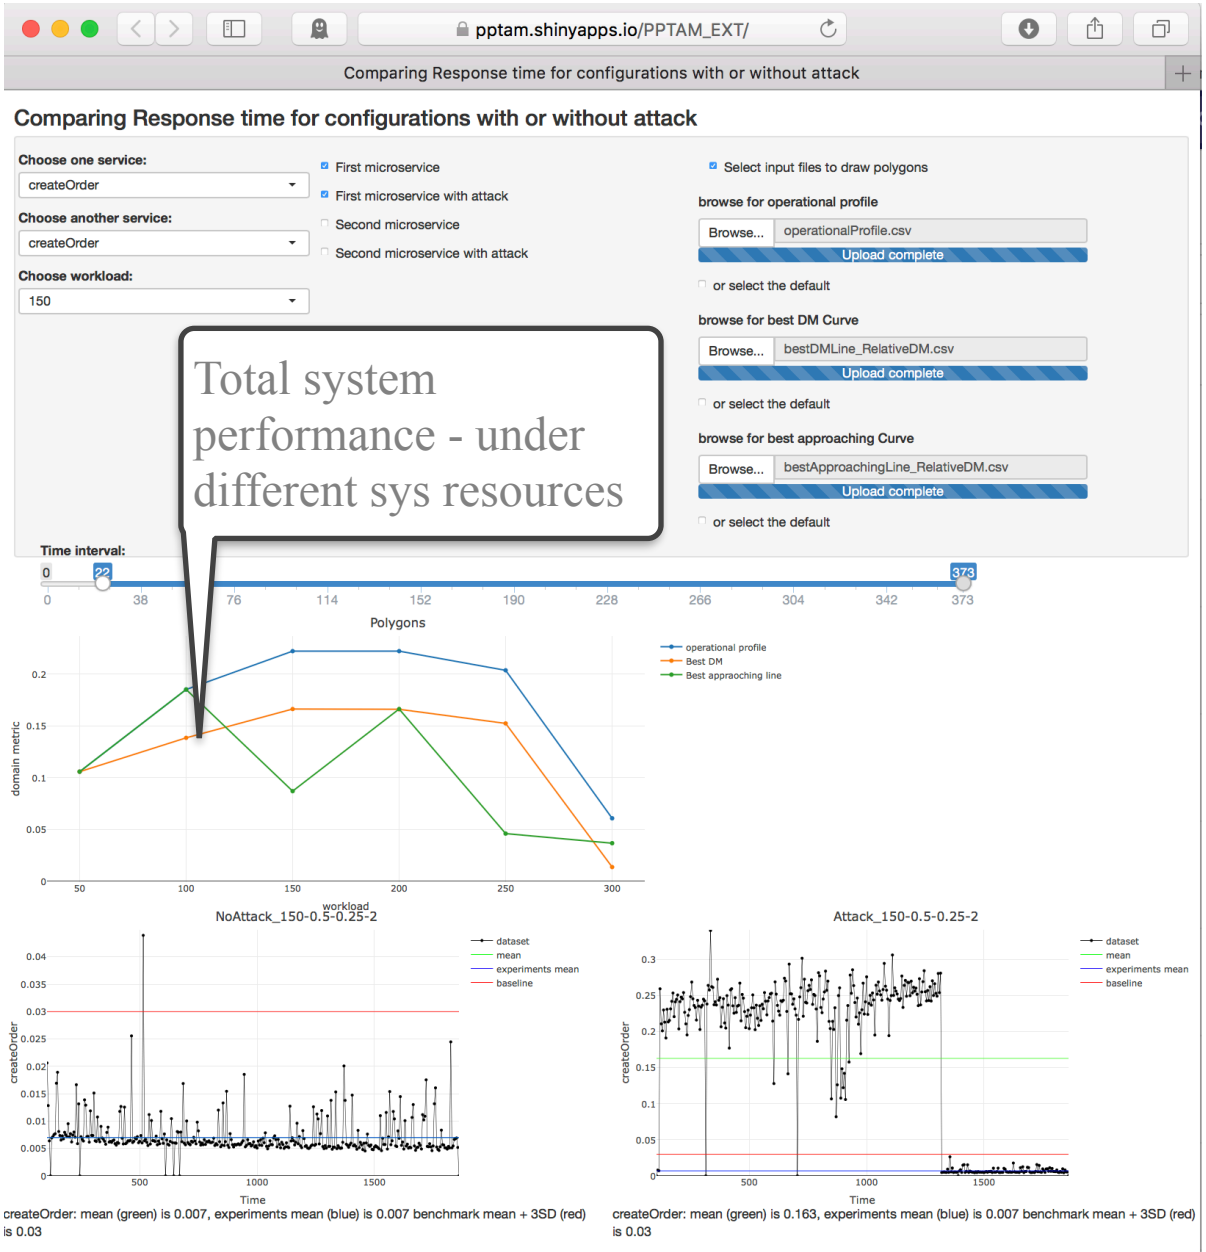

createOrder: mean (green) is 0.007, experiments mean (blue) is 0.007 benchmark mean + 3SD (red) is 0.03

createOrder: mean (green) is 0.163, experiments mean (blue) is 0.007 benchmark mean + 3SD (red) is 0.03

unibz

https://pptam.shinyapps.io/PPTAM_EXT/

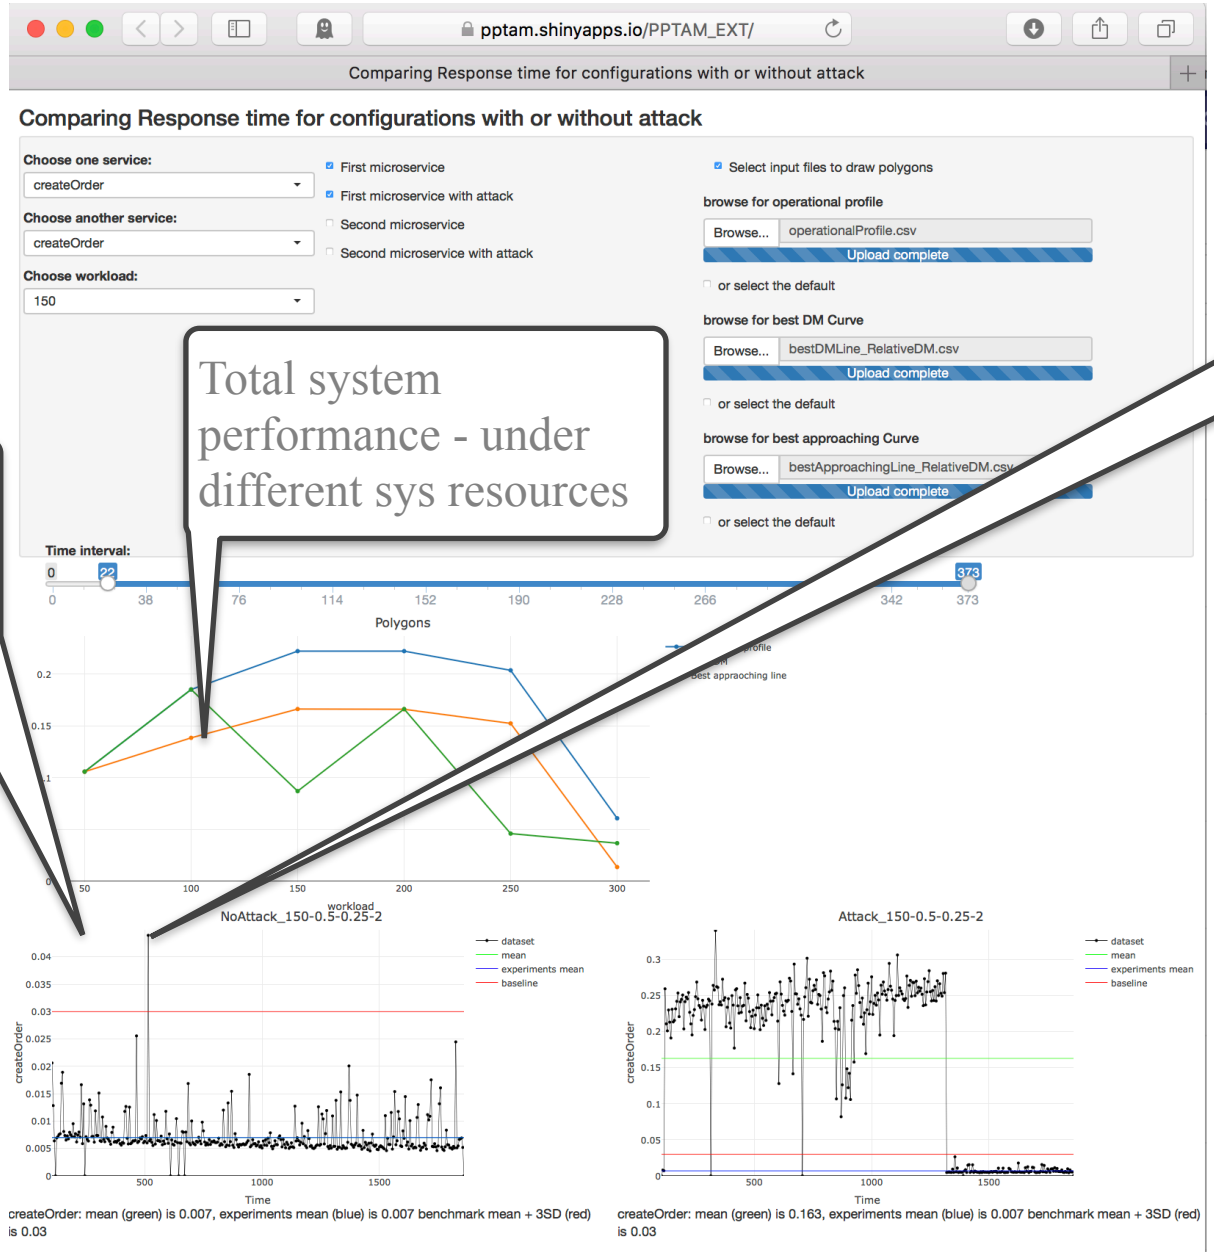

createOrder: mean (green) is 0.007, experiments mean (blue) is 0.007 benchmark mean + 3SD (red) is 0.03

createOrder: mean (green) is 0.163, experiments mean (blue) is 0.007 benchmark mean + 3SD (red) is 0.03

No attack

Total system performance - under different sys resources

Service failure

unibz

https://pptam.shinyapps.io/PPTAM_EXT/

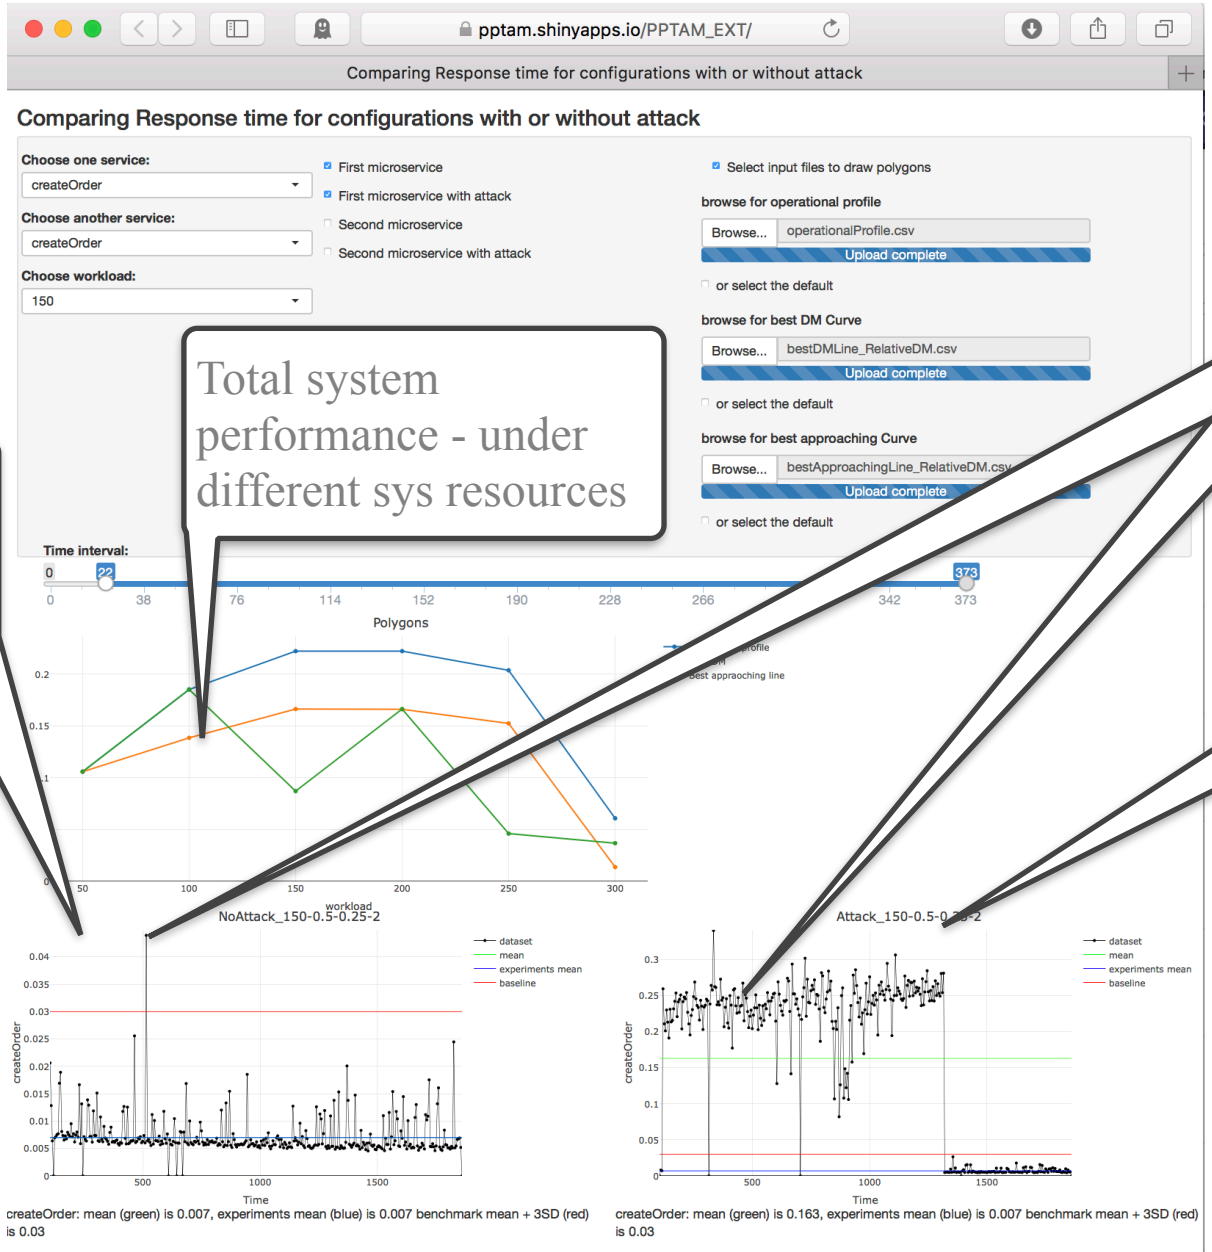

No attack

Total system performance - under different sys resources

Service failures

Under attack

unibz

Monitoring cockpit

Cockpit

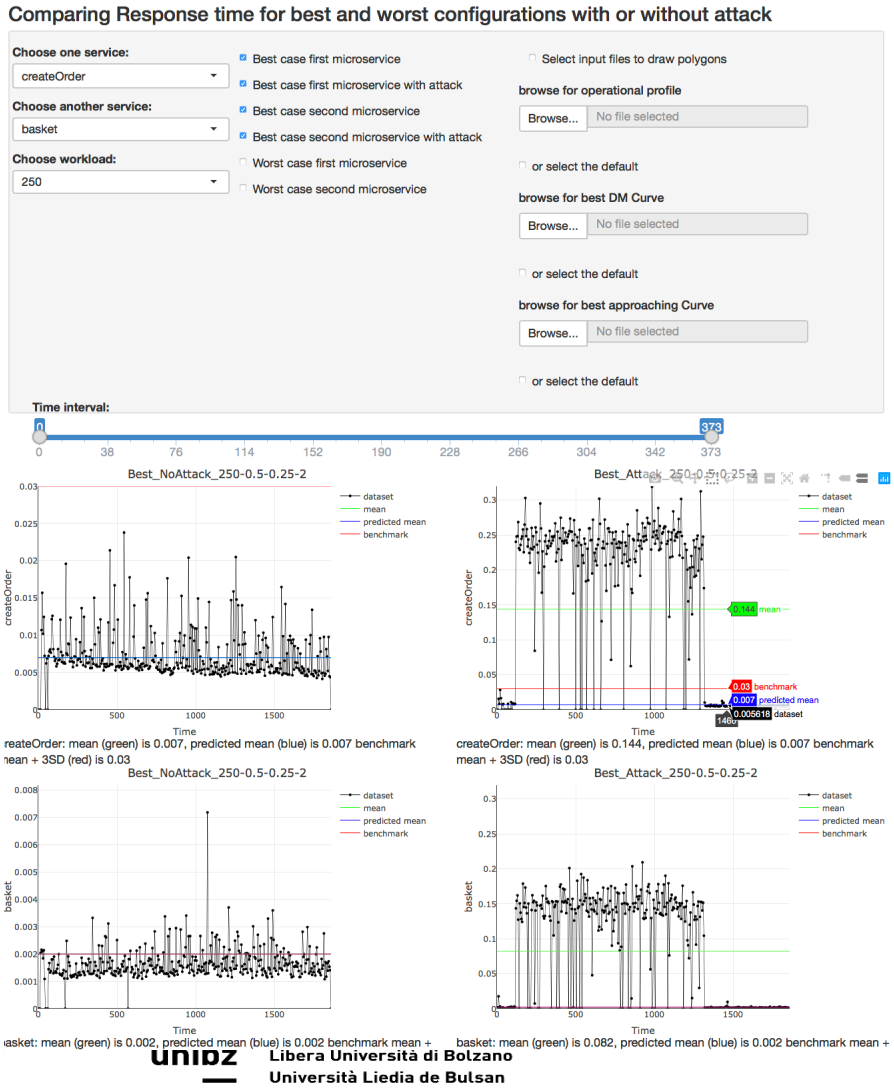

https://pptam.shinyapps.io/PPTAM_EXT/

Case studies

Microservice architecture

- In a microservice architecture, services are **fine-grained** and the protocols are lightweight rendering each micro service **loosely coupled** with the others
- Microservice architectures often use **containers** to enforce service independence

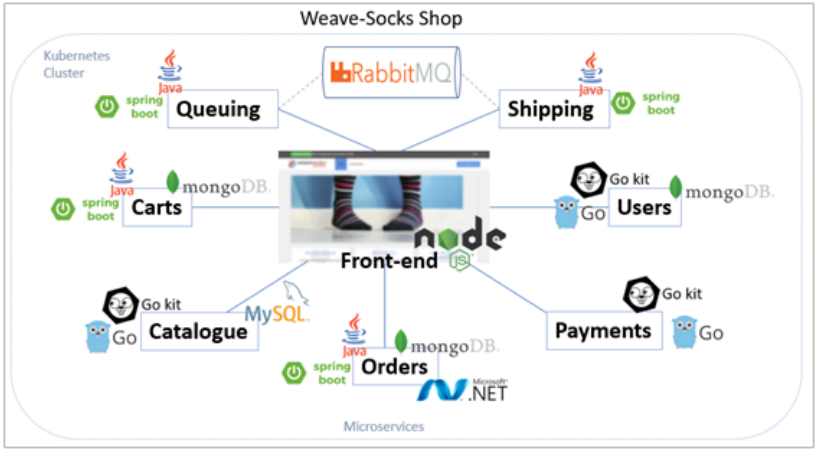

Before starting a transition …

- There are characteristics that are shared by microservice architectures:
  - Data is organized in a decentralized way: each service manages its own data makes it independently deployable
  - Teams that build systems with microservices extensively use infrastructure automation techniques (like continuous integration or continuous delivery)

Data independently deployable

Automation

Quantitative Assessment of Deployment Alternatives

- Challenge: **assess performance of architectural deployment alternatives** (e.g., number of replicas, CPU/memory allocation, technology stack) under fuzzy requirements

A Quantitative Approach for the Assessment of Microservice Architecture Deployment Alternatives by Automated Performance Testing

Alberto Avritzer[1], Vincenzo Ferme[2], Andrea Janes[3], Barbara Russo[3], Henning Schulz[4], and André van Hoorn[2]

[1] EsdaleSolutions, Inc., Princeton, NJ, USA
[2] University of Stuttgart, Germany
[3] Free University of Bozen-Bolzano, Italy
[4] NovaTec Consulting GmbH, Leinfelden-Echterdingen, Germany

Abstract. Microservices have emerged as an architectural style for developing distributed applications. Assessing the performance of architectural deployment alternatives is challenging and must be aligned with the system usage in the production environment. In this paper, we introduce an approach for using operational profiles to generate load tests to automatically assess scalability pass/fail criteria of several microservices deployment alternatives, using a baseline computation of performance requirements. We have evaluated our approach with different architecture deployment alternatives using extensive lab studies in a large bare metal host environment and a virtualized environment. We have found that in auto-scaling cloud environments careful performance engineering activities shall be executed before additional resources are added to the architecture deployment configuration.

1 Introduction

The microservices architectural style [9] is an approach for creating software applications as a collection of loosely coupled software components. These components are called microservices, and are supposed to be autonomous, automatically and independently deployable, and cohesive [2,3]. This architecture leads itself to decentralized deployment, and for continuous integration and deployment by developers. Several large companies (e.g., Amazon and Netflix) are reporting significant success with microservice architectures [5,6].

Currently, several architecture deployment alternatives are being used for microservices deployment, as for example, serverless microservices using lambdas, container-based deployment (e.g., Docker[8]), virtual machines per host, and several hosts. Of course, depending on the microservice granularity, a combination of these deployment mechanism could be used.

Microservices are supposed to be independent of each other. However, the underlying deployment environment might introduce coupling and impact the

[8] https://www.docker.com/

Alberto Avritzer, Vincenzo Ferme, Andrea Janes, Barbara Russo, Henning Schulz, and André van Hoorn:
A Quantitative Approach for the Assessment of Microservice Architecture Deployment Alternatives by Automated Performance Testing.

In Proceedings of the 12th European Conference on Software Architecture (ECSA). LNCS, Springer, 2018 (Accepted)

Quantitative Assessment of Deployment Alternatives

- Approach
  - Use operational data to generate and weigh load tests
  - Measure baseline requirements
  - Design a metric that allows quantitative comparison of deployment alternatives

Overview of Approach

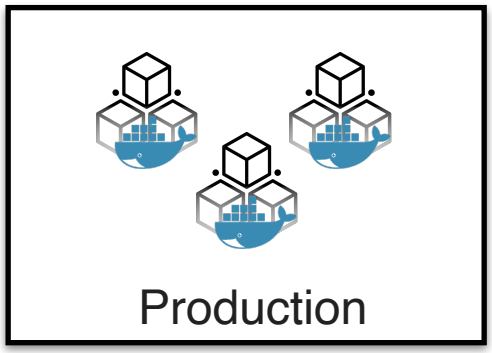

Overview of Approach

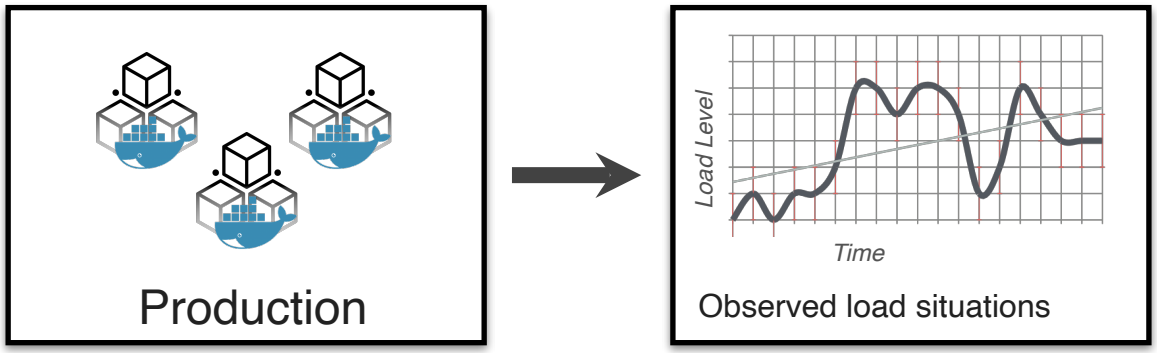

Overview of Approach

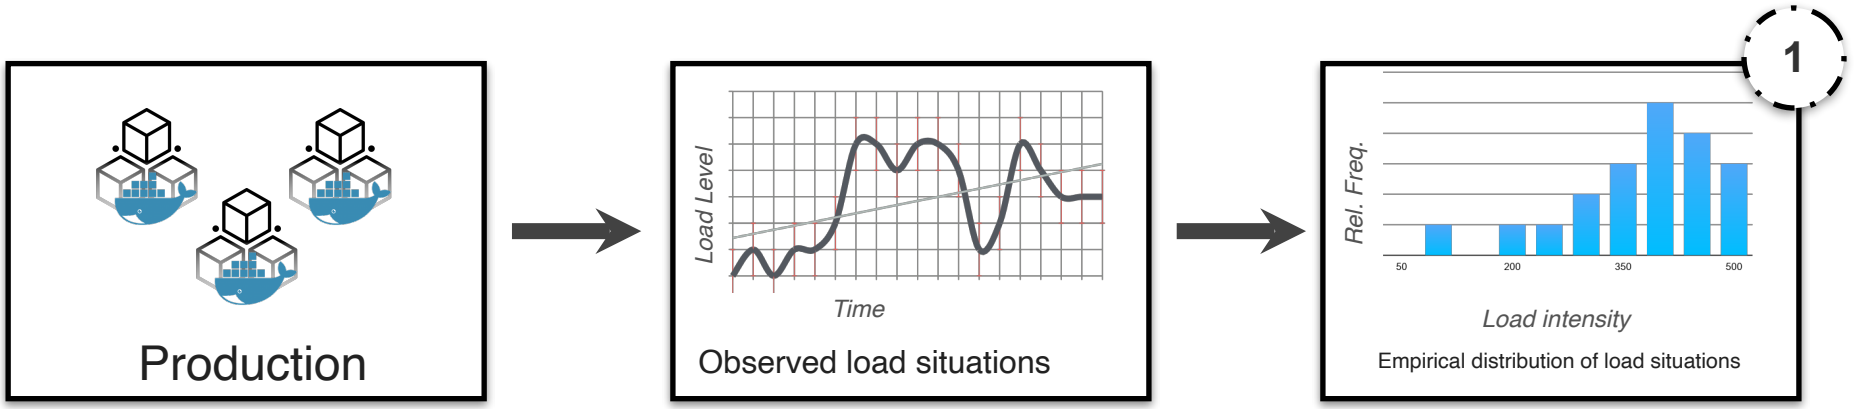

Overview of Approach

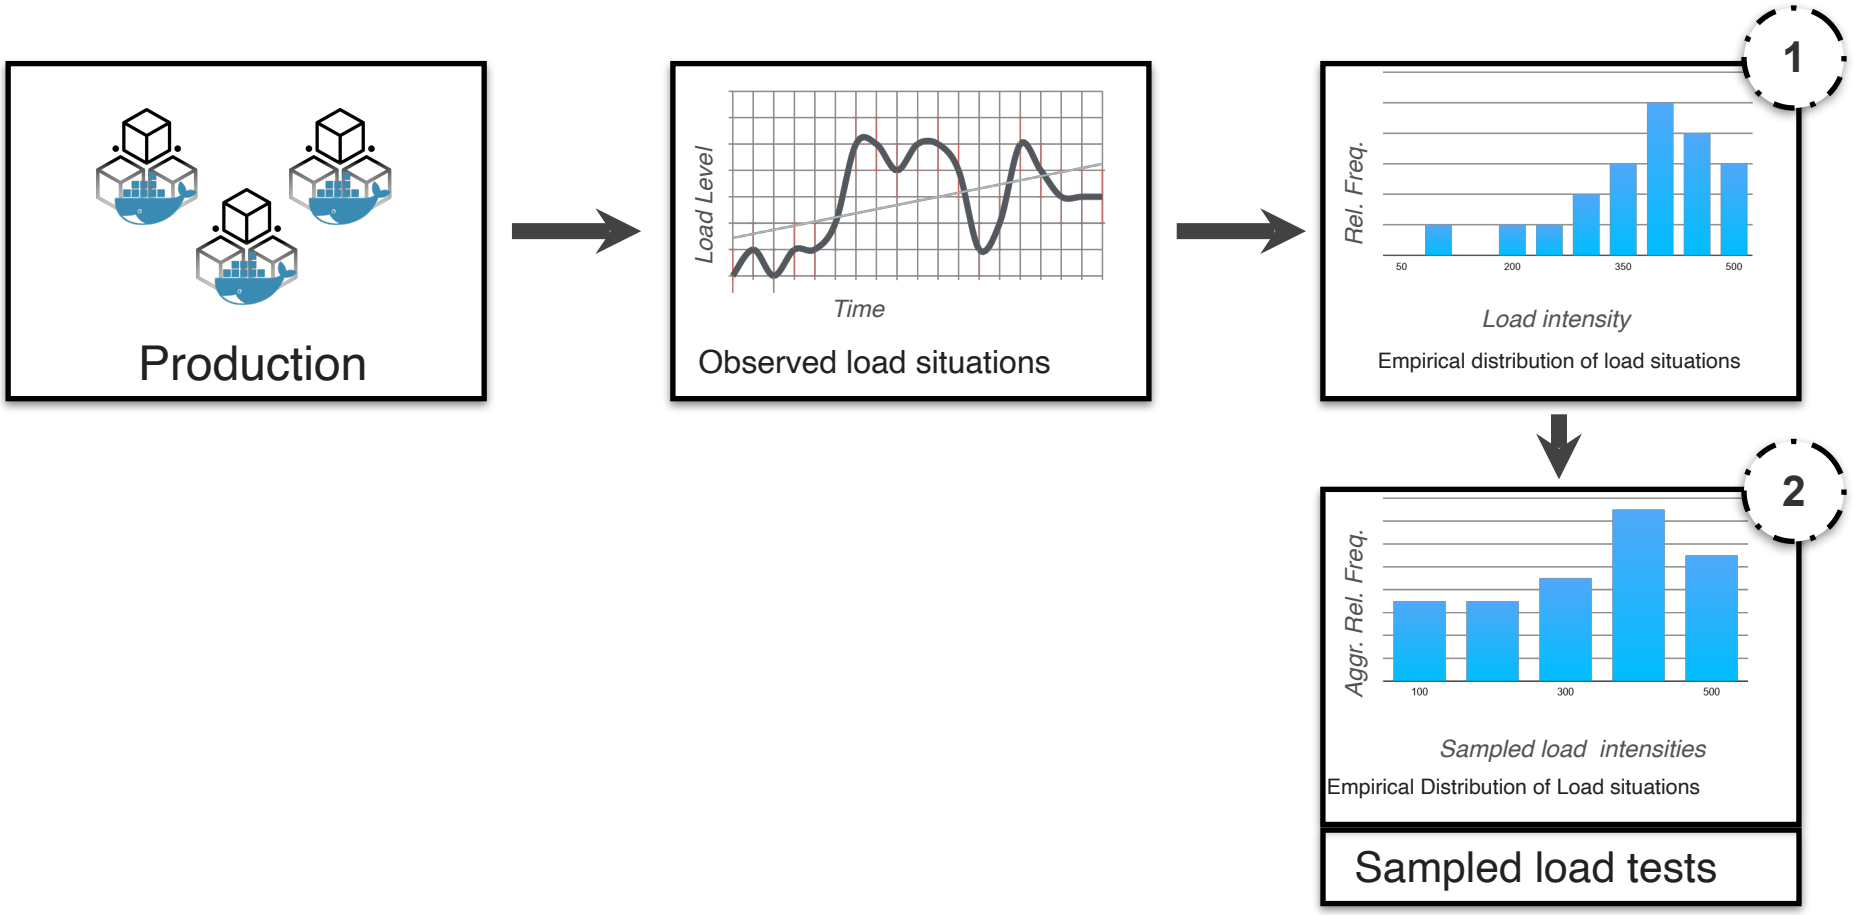

Overview of Approach

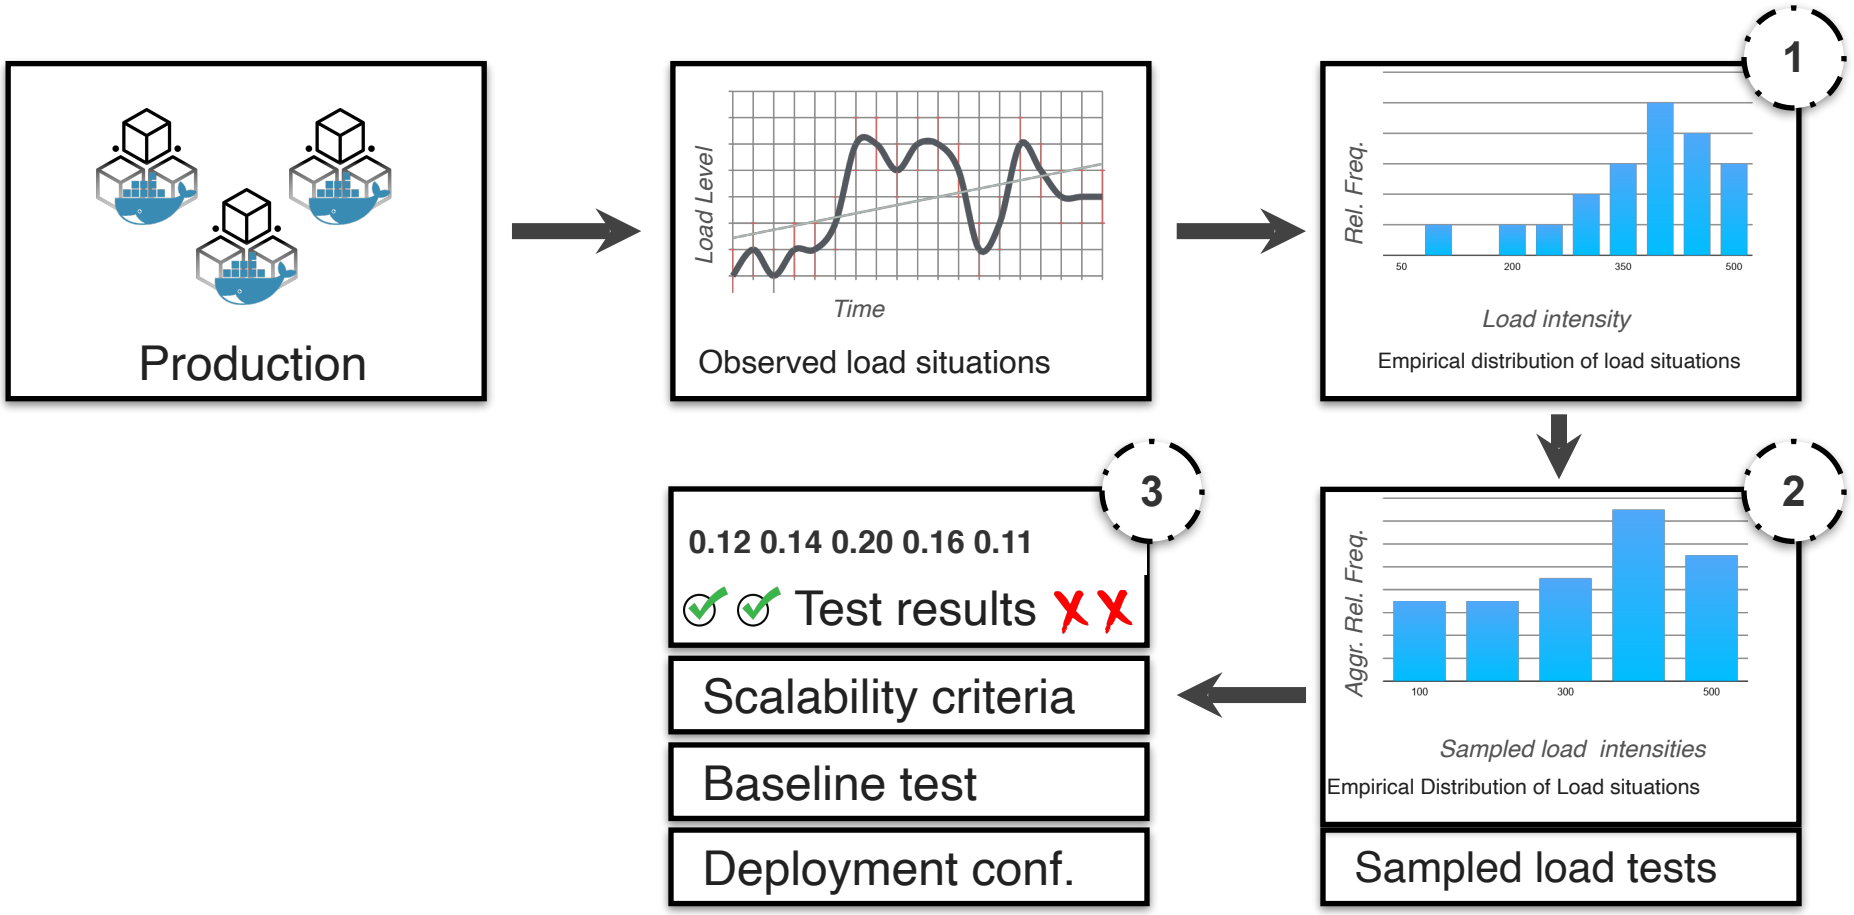

Overview of Approach

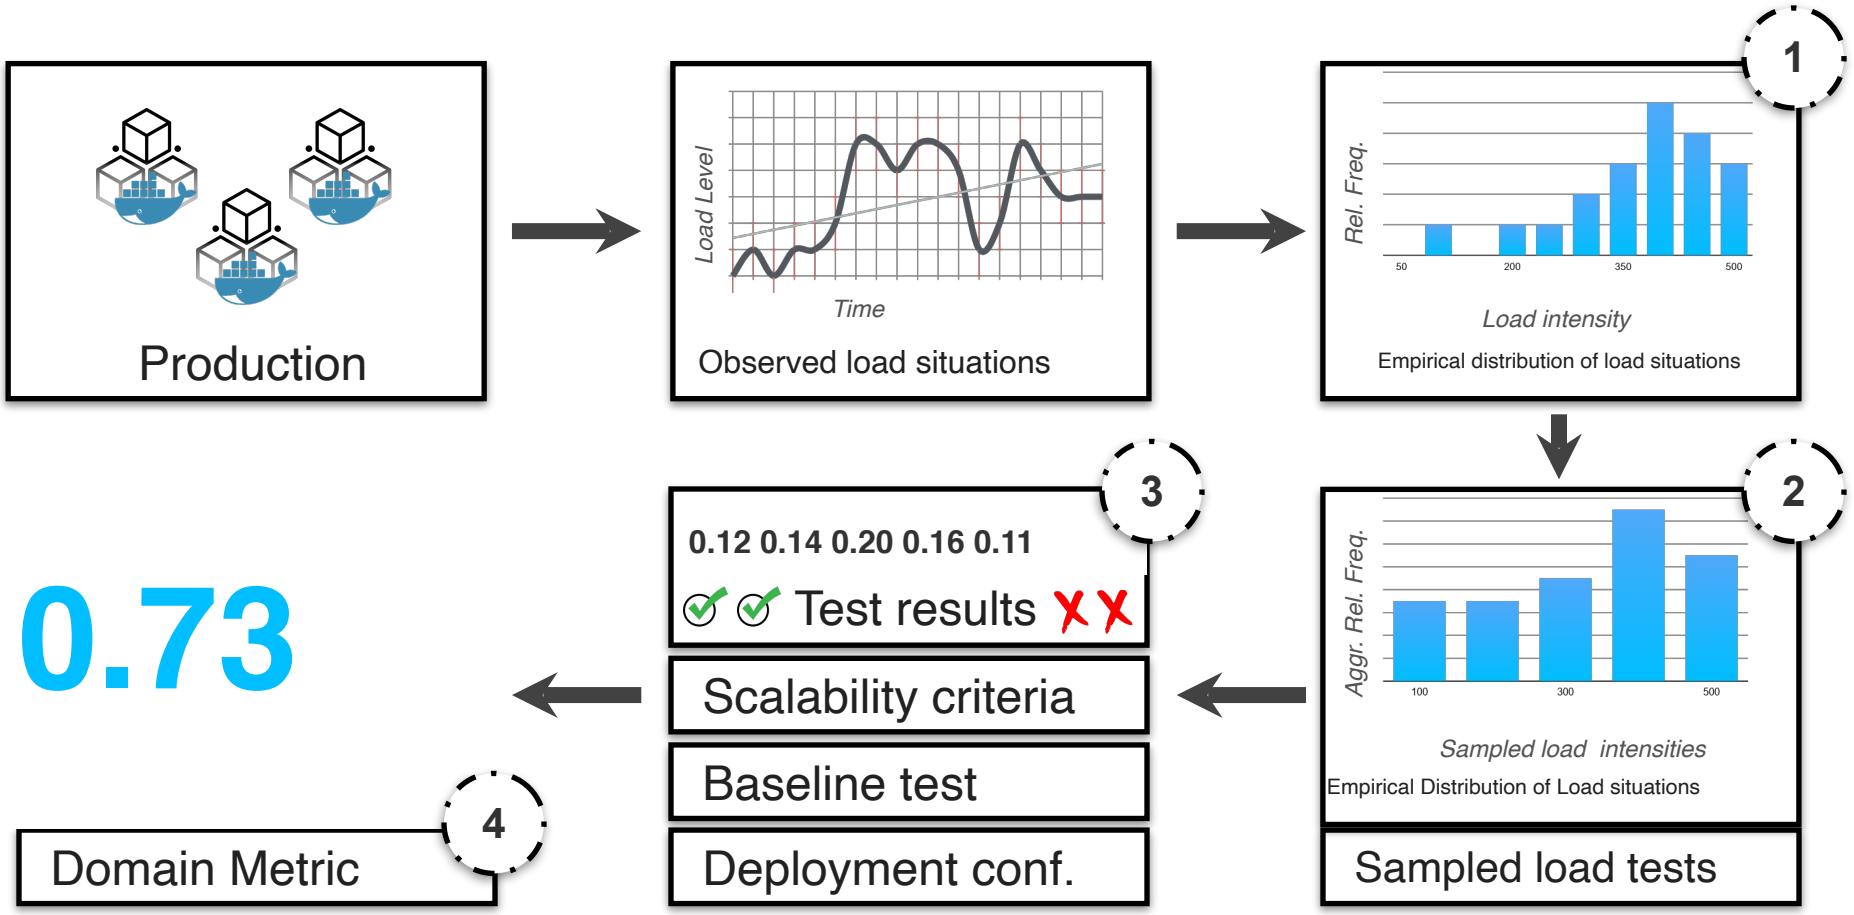

System Under Test

Sock Shop

Quickstart Docs API GitHub

Sock Shop

A Microservices Demo Application

Sock Shop simulates the user-facing part of an e-commerce website that sells socks. It is intended to aid the demonstration and testing of microservice and cloud native technologies.

Sock Shop is maintained by Weaveworks and Container Solutions.

unibz

Experiments

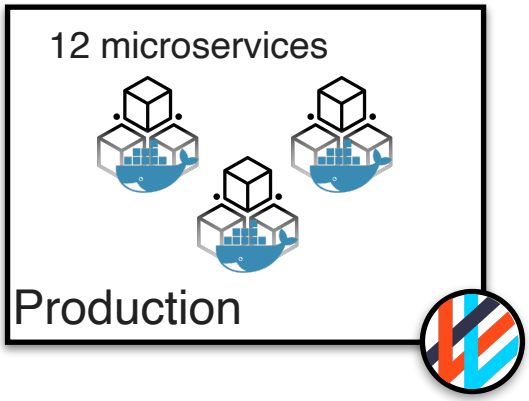

Experiments

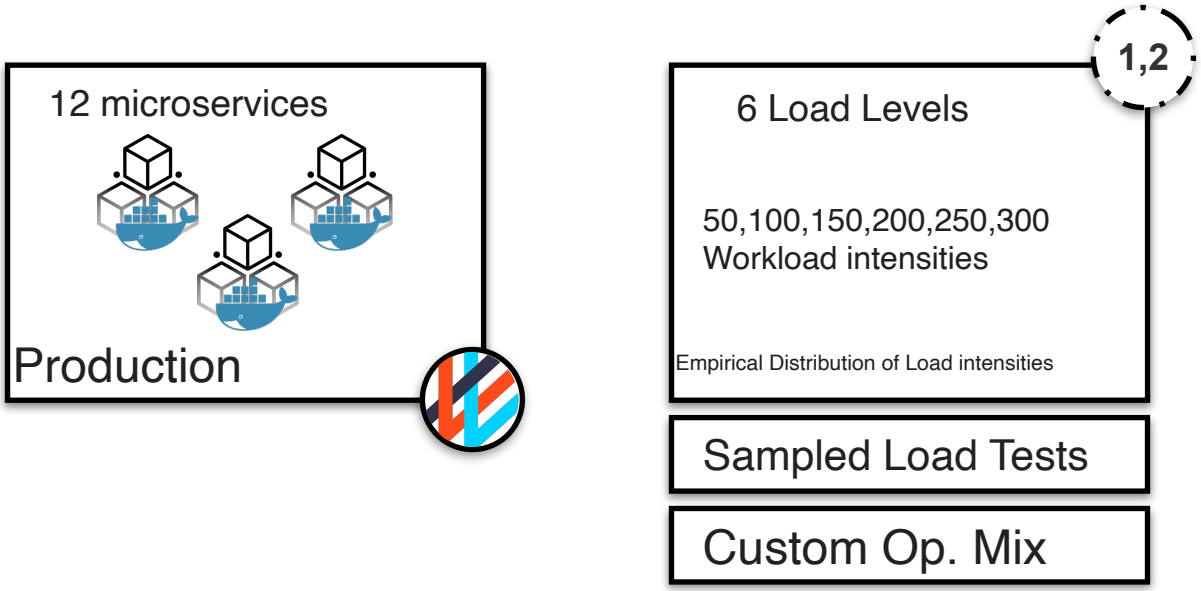

Experiments

12 microservices

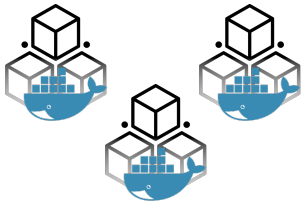

Production

6 Load Levels

50,100,150,200,250,300
Workload intensities

Empirical Distribution of Load intensities

1,2

Sampled Load Tests

Custom Op. Mix

12 configurations

RAM

CPU

Replicas

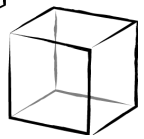

Deployment Config.

3

Scal = avg + 3σ

Experiment Results: Computation of Domain Metric (1/2)

API	Scalability Criteria
GET /	**PASS**
GET /cart	**PASS**
POST /item	**FAIL**

Users	Aggr. Rel. Freq.
50	0.10582
100	0.18519
150	0.22222
200	0.22222
250	0.20370
300	0.06085

Custom Op. Mix

Aggr. Rel. Freq.

Contrib. to Domain Metric

Deployment Configuration: 1 GB RAM, 0.25 CPU, 1 Replica

Experiment Results: Computation of Domain Metric (1/2)

API	Scalability Criteria
GET /	**PASS**
GET /cart	**PASS**
POST /item	**FAIL**

Custom Op. Mix

Users	Aggr. Rel. Freq.
50	0.10582
100	0.18519
150	0.22222
200	0.22222
250	0.20370
300	0.06085

Aggr. Rel. Freq.

Max: 0.20370

Actual: 0.13580

Contrib. to Domain Metric

Deployment Configuration: 1 GB RAM, 0.25 CPU, 1 Replica

Experiment Results: Computation of Domain Metric (2/2)

Users	Contribution
50	0.10582
100	0.18519
150	0.22222
200	0.07999
250	0.13580
300	0.04729

Contrib. to Domain Metric

Deployment Configuration: 1 GB RAM, 0.25 CPU, 1 Replica

Experiment Results: Computation of Domain Metric (2/2)

Users	Contribution
50	0.10582
100	0.18519
150	0.22222
200	0.07999
250	0.13580
300	0.04729

Max: 1

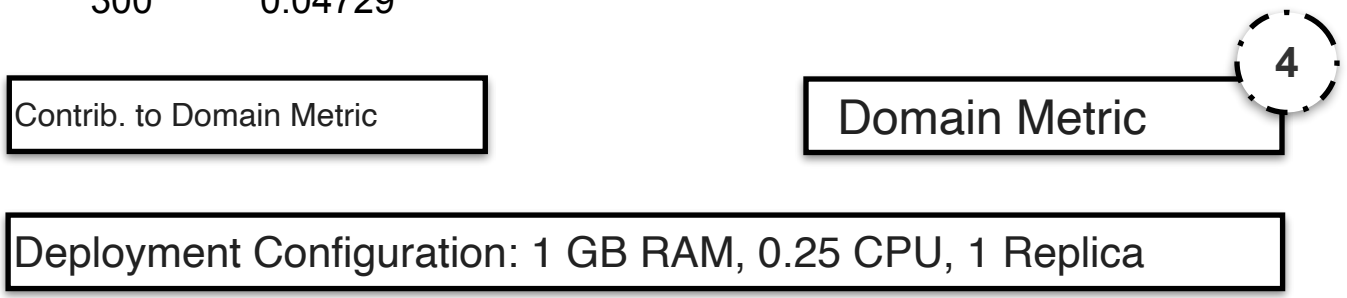

Experiment Results: Computation of Domain Metric (2/2)

Users	Contribution
50	0.10582
100	0.18519
150	0.22222
200	0.07999
250	0.13580
300	0.04729

Max: 1

Actual:

0.77631

4

Contrib. to Domain Metric

Domain Metric

Deployment Configuration: 1 GB RAM, 0.25 CPU, 1 Replica

Experiment Results: Single-Metric Comparison of Alternatives

RAM	CPU #	Cart Replicas	Domain Metric (HPI)	Domain Metric (FUB)
0.5 GB	0.25	1	0.61499	0.54134
1 GB	**0.25**	**1**	**0.77631**	0.53884
1 GB	0.5	1	0.53559	0.54106
0.5 GB	0.5	1	0.51536	0.54773
0.5 GB	0.5	2	0.50995	0.54111
1 GB	0.25	2	0.74080	0.54785
1 GB	0.5	2	0.53401	0.54106
0.5 GB	***0.5***	***4***	0.50531	***0.54939***
1 GB	0.25	4	0.37162	0.54272
1 GB	0.5	4	0.56718	0.54271

Experiment Results: Single-Metric Comparison of Alternatives

RAM	CPU #	Cart Replicas	Domain Metric (HPI)	Domain Metric (FUB)
0.5 GB	0.25	1	0.61499	0.54134
1 GB	**0.25**	**1**	**0.77631**	0.53884
1 GB	0.5	1	0.53559	0.54106
0.5 GB	0.5	1	0.51536	0.54773
0.5 GB	0.5	2	0.50995	0.54111
1 GB	0.25	2	0.74080	0.54785
1 GB	0.5	2	0.53401	0.54106
0.5 GB	***0.5***	***4***	0.50531	***0.54939***
1 GB	0.25	4	0.37162	0.54272
1 GB	0.5	4	0.56718	0.54271

Experiment Results: Single-Metric Comparison of Alternatives

RAM	CPU #	Cart Replicas	Domain Metric (HPI)	Domain Metric (FUB)
0.5 GB	0.25	1	0.61499	0.54134
1 GB	**0.25**	**1**	**0.77631**	0.53884
1 GB	0.5	1	0.53559	0.54106
0.5 GB	0.5	1	0.51536	0.54773
0.5 GB	0.5	2	0.50995	0.54111
1 GB	0.25	2	0.74080	0.54785
1 GB	0.5	2	0.53401	0.54106
0.5 GB	***0.5***	***4***	0.50531	***0.54939***
1 GB	0.25	4	0.37162	0.54272
1 GB	0.5	4	0.56718	0.54271

Experiment Results: Visual Comparison of Alternatives

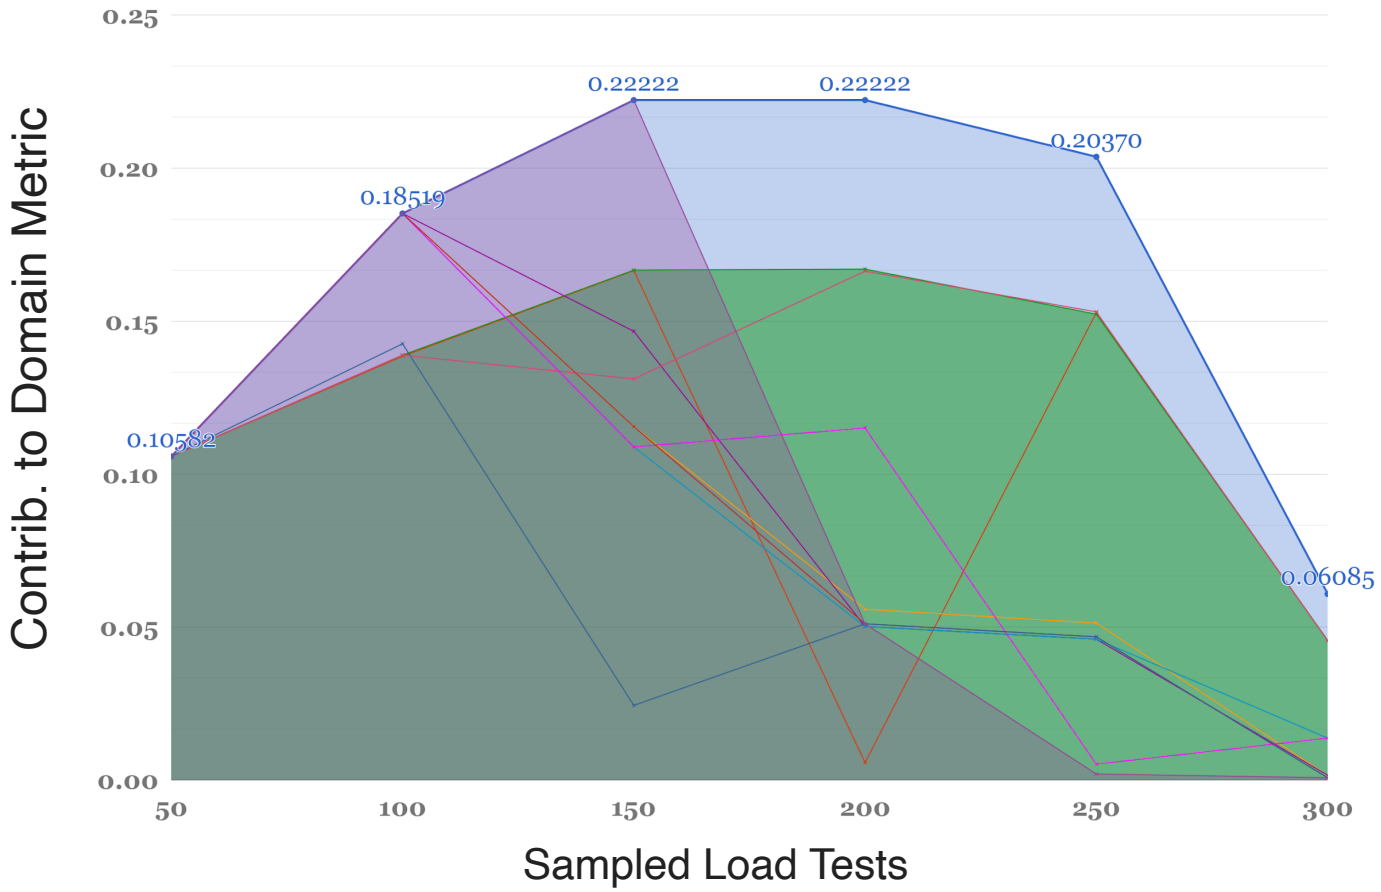

Experiment Results: Visual Comparison of Alternatives

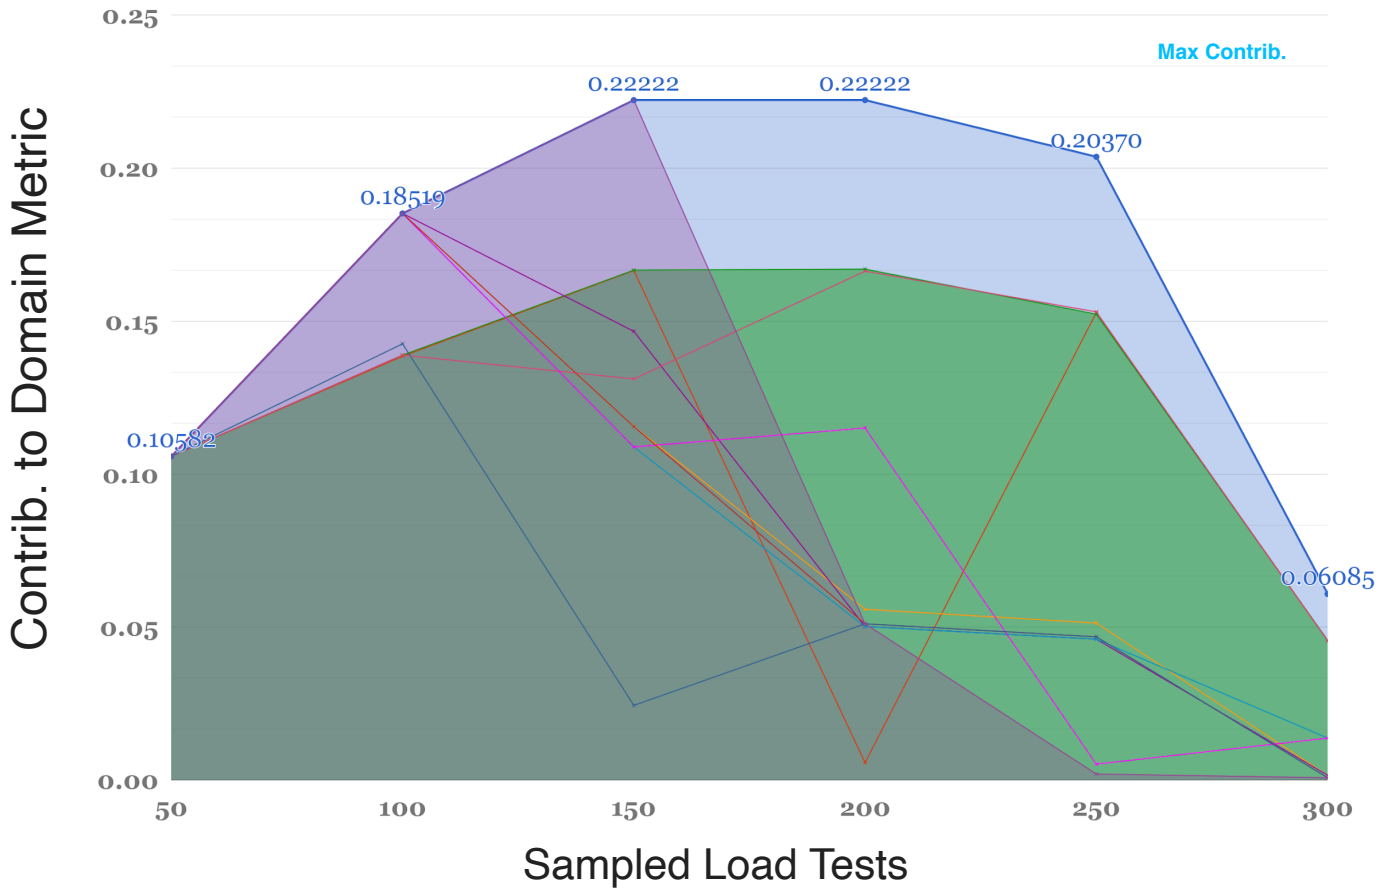

Experiment Results: Visual Comparison of Alternatives

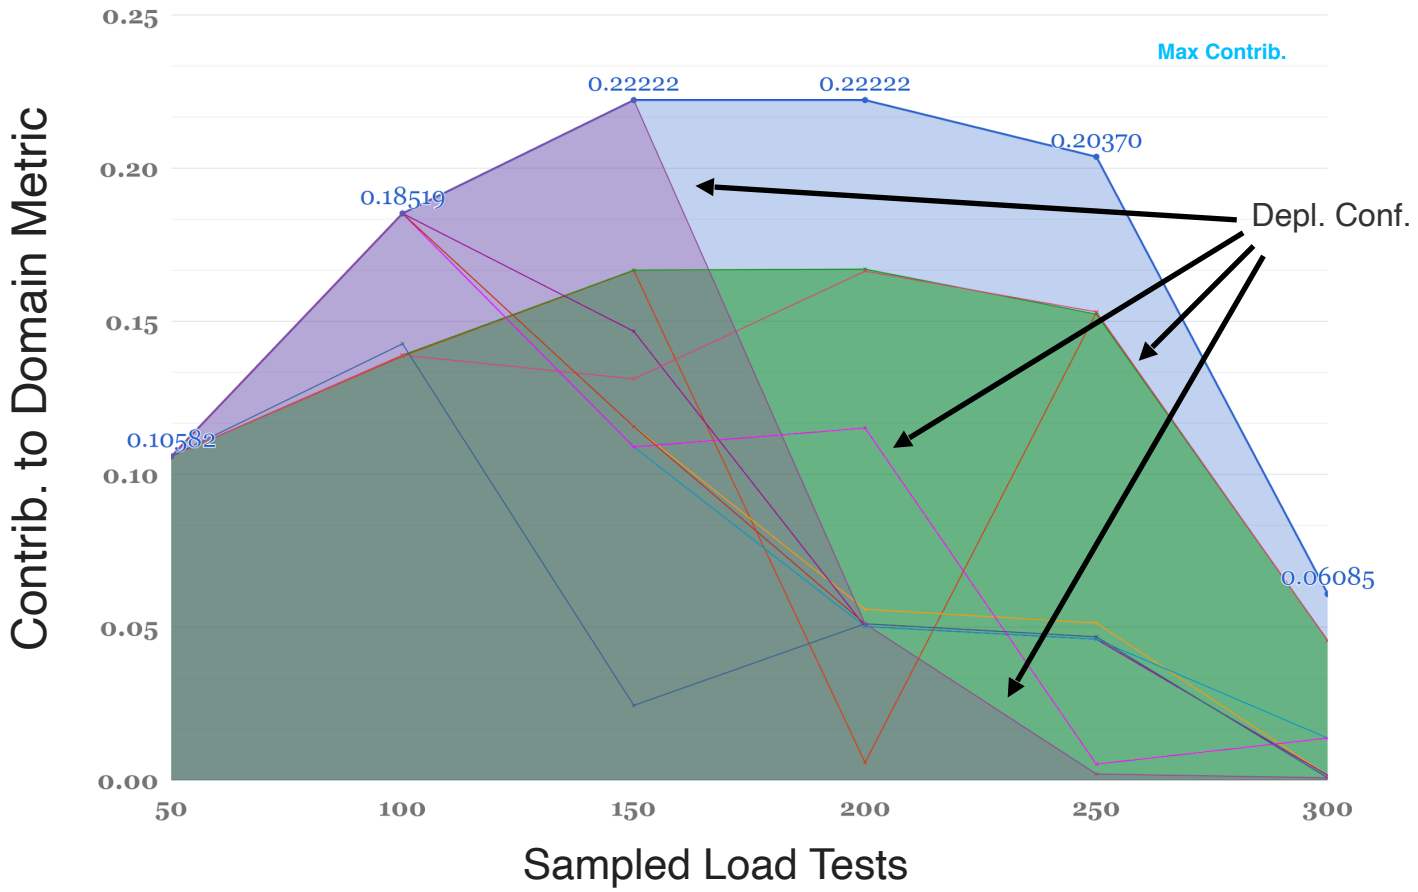

Extensions/Application

Extensions/Application

- We have tested it on an online demo-platform
- We have extended it to monitor performce degradation under attacks by incorporating Mirai

Extensions/Application

- We have tested it on an online demo-platform
- We have extended it to monitor performce degradation under attacks by incorporating Mirai
- We have designed it for monitoring performance degradation during a transition to microservices

Bare-metal versus virtualization environment

System Under Test

Sock Shop

Quickstart Docs API GitHub

Sock Shop

A Microservices Demo Application

Sock Shop simulates the user-facing part of an e-commerce website that sells socks. It is intended to aid the demonstration and testing of microservice and cloud native technologies.

Sock Shop is maintained by Weaveworks and Container Solutions.

unibz

Experiment settings

- 2 VM one for SUT and one for Test
- SUT: docker containers each for on micro service, one for DB

Bare-metal

- The containerized bare metal machines:
- Load driver server - **32 GB RAM, 24 cores** (2 threads each) at 2300 MHz and SUT server - **896 GB RAM, 80 cores** (2 threads each) at 2300 MHz
- Both machines use magnetic disks with 15 000 rpm and are connected using a shared 10 Gbit/ s network infrastructure

Virtual

- The containerized deployment in virtual machines:
- *Load driver server* - **4 GB RAM, 1 core** at 2600MHz and *SUT server* - **8 GB RAM, 4 cores** at 2600 MHz with SSDs
- Both machines use an EMCVNC 5400 series network attached storage solution12 and are connected using a shared 10 Gbit/ s network infrastructure

- We replicated with *SUT server* - **16 GB RAM, 8 cores**

Virtual

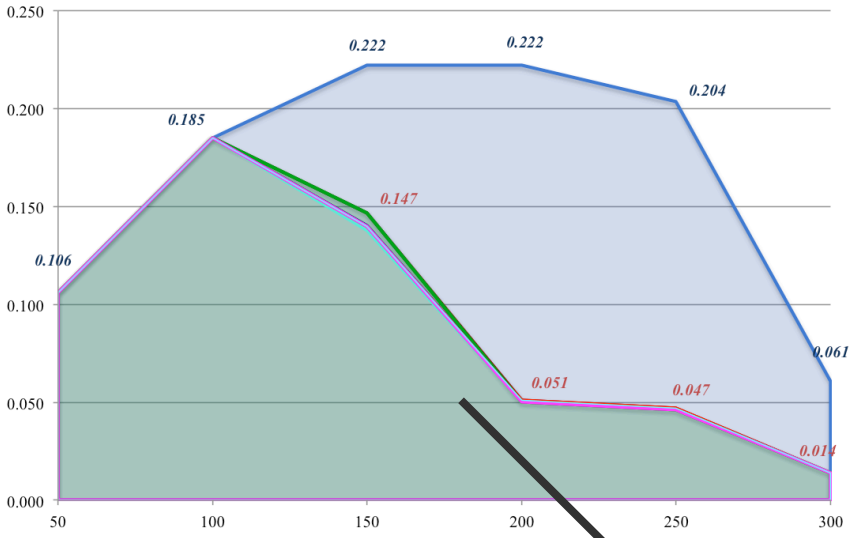

Bare metal

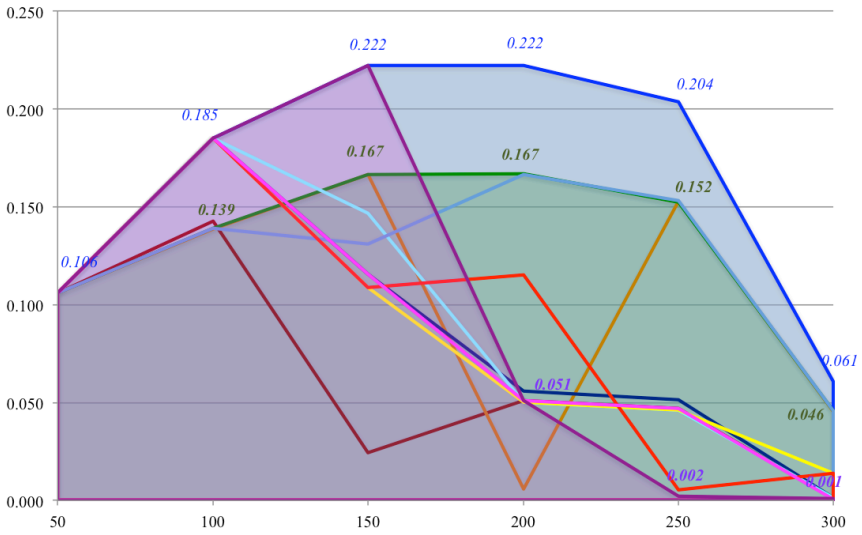

Virtual more resources for SUT

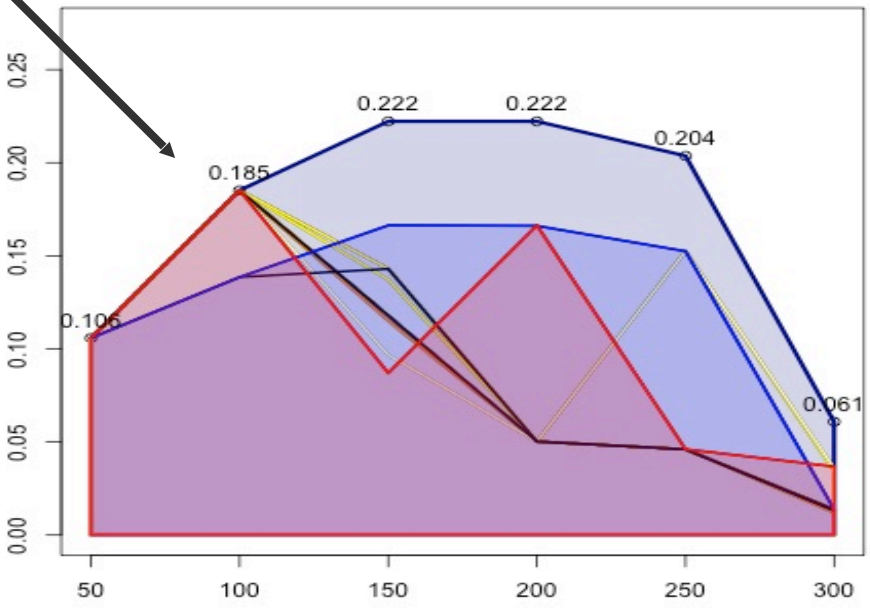

Virtual

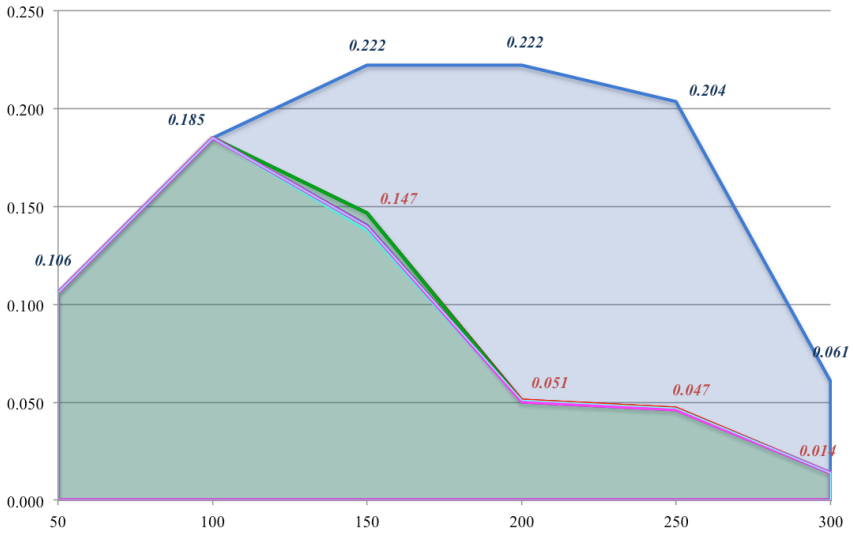

Bare metal

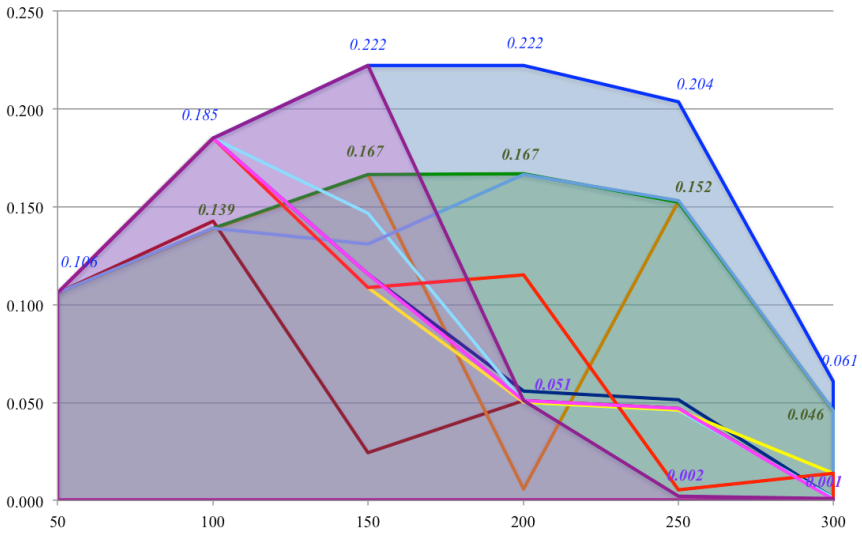

Virtual more resources for SUT

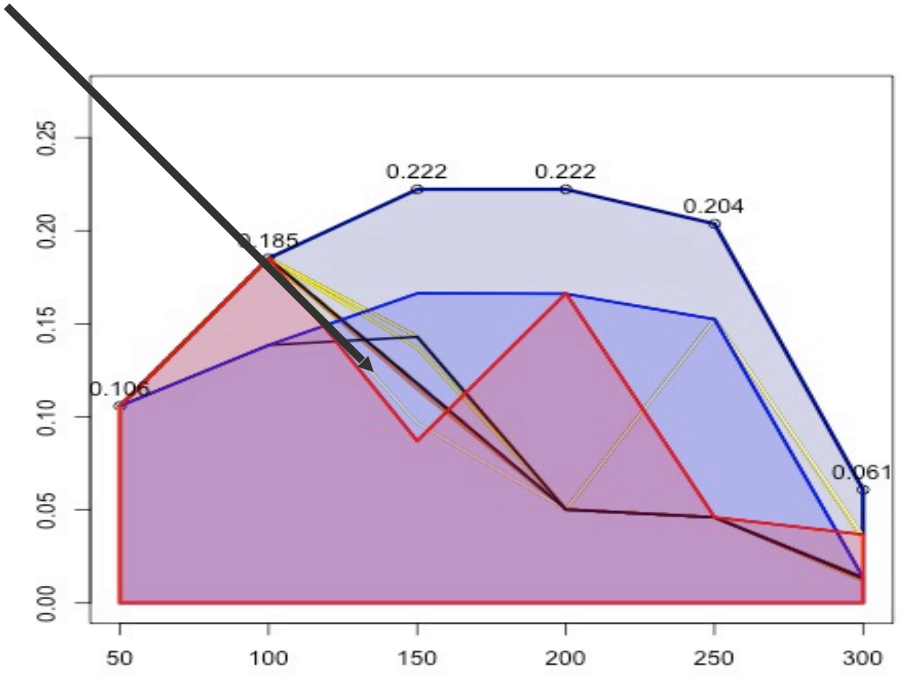

Bare-metal versus virtualization environment

Bare-metal versus virtualization environment

Monitor performance degradation under attacks

System Under Test

Sock Shop

Quickstart Docs API GitHub

Sock Shop

A Microservices Demo Application

Sock Shop simulates the user-facing part of an e-commerce website that sells socks. It is intended to aid the demonstration and testing of microservice and cloud native technologies.

Sock Shop is maintained by Weaveworks and Container Solutions.

unibz

Mirai BotNet

- Mirai is a **malware** that has been used to turn networked devices (cameras) running Linux into remotely controlled bots
- We use an academic version of it to attack the system in controlled experiments
- It can perform different types of attack. By now, we have explored http, syn, ack

Experiments with Mirai

- Attack with simple http requests (GET and POST to home - increase the load)

- Compute the metric with and without attack to understand:
  - the resilience of a system
  - the early prediction of an attack

PPTAM

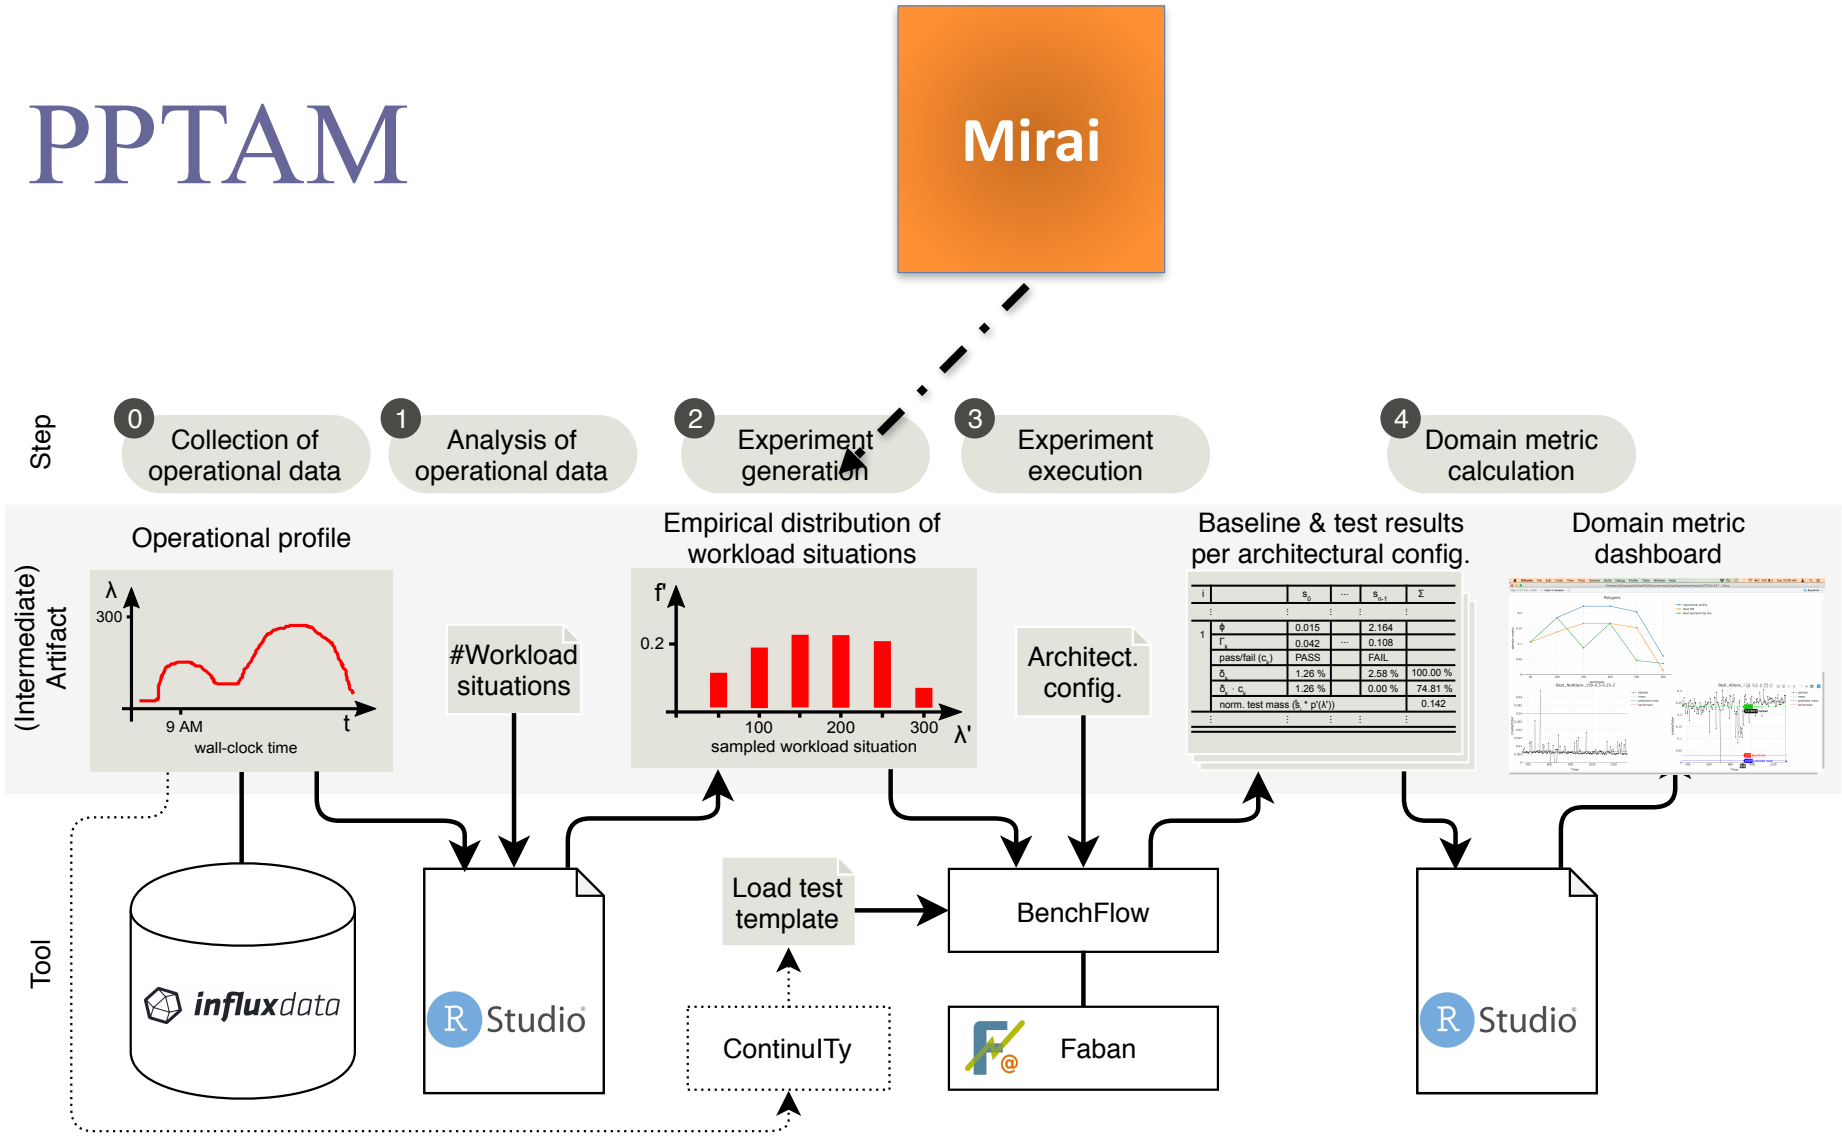
Mirai
Step
0 Collection of operational data
1 Analysis of operational data
2 Experiment generation
3 Experiment execution
4 Domain metric calculation
(Intermediate) Artifact
Operational profile
λ
300
9 AM
t
wall-clock time
#Workload situations
Empirical distribution of workload situations
f'
0.2
100
200
300
λ'
sampled workload situation
Architect. config.
Baseline & test results per architectural config.
Domain metric dashboard
Tool
influxdata
R Studio
Load test template
ContinuITy
BenchFlow
Faban
R Studio

Experiments - results

Virtual – no attack

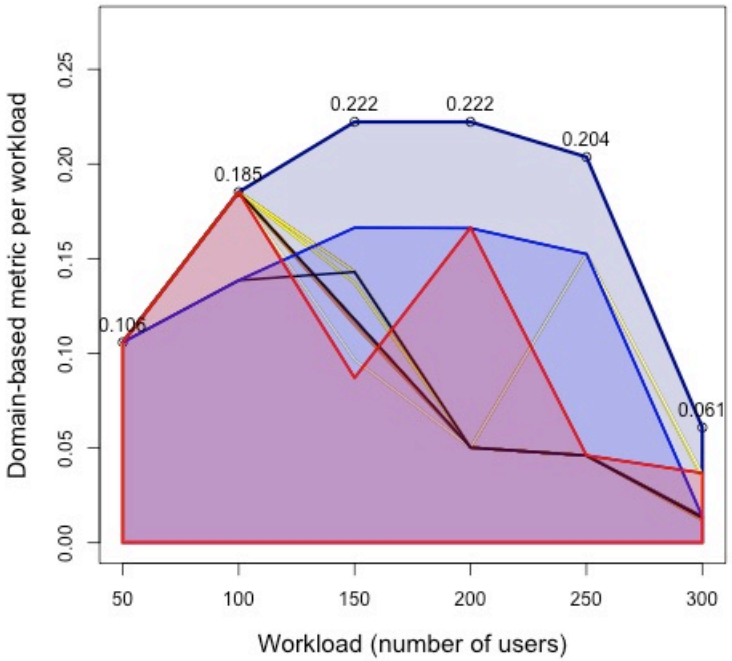

Virtual - attack

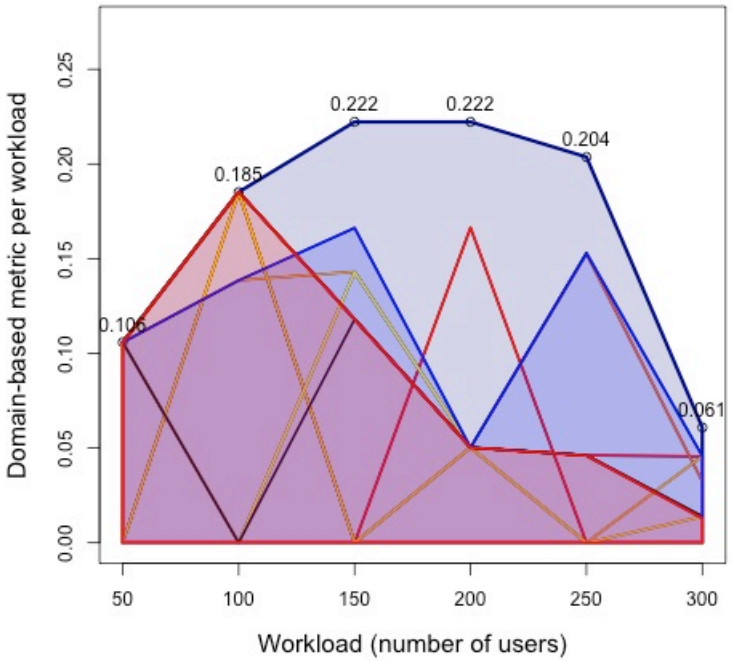

Attack design

- After few piloting attacks (5-10-20 mins)

- Duration of attack: 20 minutes (1200 seconds);
- Protocol used: HTTP;
- IP address to attack: the IP address of the SUT, i.e., the
- Machine with Sock Shop installed;
- Number of threads: 256.

Bare-metal versus virtualization environment

Monitor performance degradation under attacks

Bare-metal versus virtualization environment

Monitor performance degradation under attacks

Monitoring performance degradation during a transition to microservices

Dehghani's approach to transition

- Identify one capability in the monolith to transform it into microservice(s)

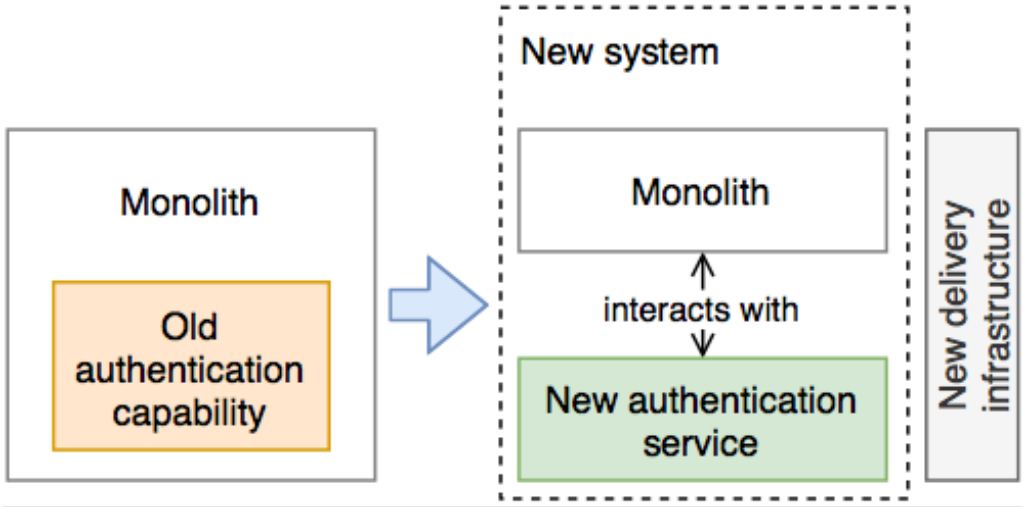

Dehghani's approach to transition

- Decouple it from the monolith into an external service
- Maintain the old monolith with all its existing functionalities
- Work incrementally: build, test, and deploy

Z. Dehghani, "How to break a Monolith into Microservices," April 2018, Fowler's page

Transition uncertainty

- Some aspects are new:
  - one has to decide on a communication infrastructure
- Other aspects that are valid when developing a monolith have to be reconsidered
  - For instance, how to keep communication between services minimal (as communication is costly and might impede scalability)

Transition uncertainty

- It requires the team to acquire new knowledge and to learn how to apply it
- New software design patterns for microservice architectures:
  - API Gateway pattern to organize how clients can access individual services

F. Pacheco, Microservice Patterns and Best Practices: Explore Patterns Like CQRS and Event Sourcing to Create Scalable, Maintainable, and Testable Microservices, Packt Publishing, 2018

Transition uncertainty

- A transition to microservices may or may not end up with the same or better performing system

- It depends on the ability of the developers to design microservices and the capability of the microservices architecture to represent the system

Main steps

Main steps

- Compute the operational profile of a monolith
- Apply PPTAM to collect individual service - individual experiment - individual workload time series
- Monitoring performance degradation over time against baseline and experiments' average performance

Main steps

- Compute the operational profile of a monolith
- Apply PPTAM to collect individual service - individual experiment - individual workload time series
- Monitoring performance degradation over time against baseline and experiments' average performance
- Analytic extension: visualize such analysis (R shiny)

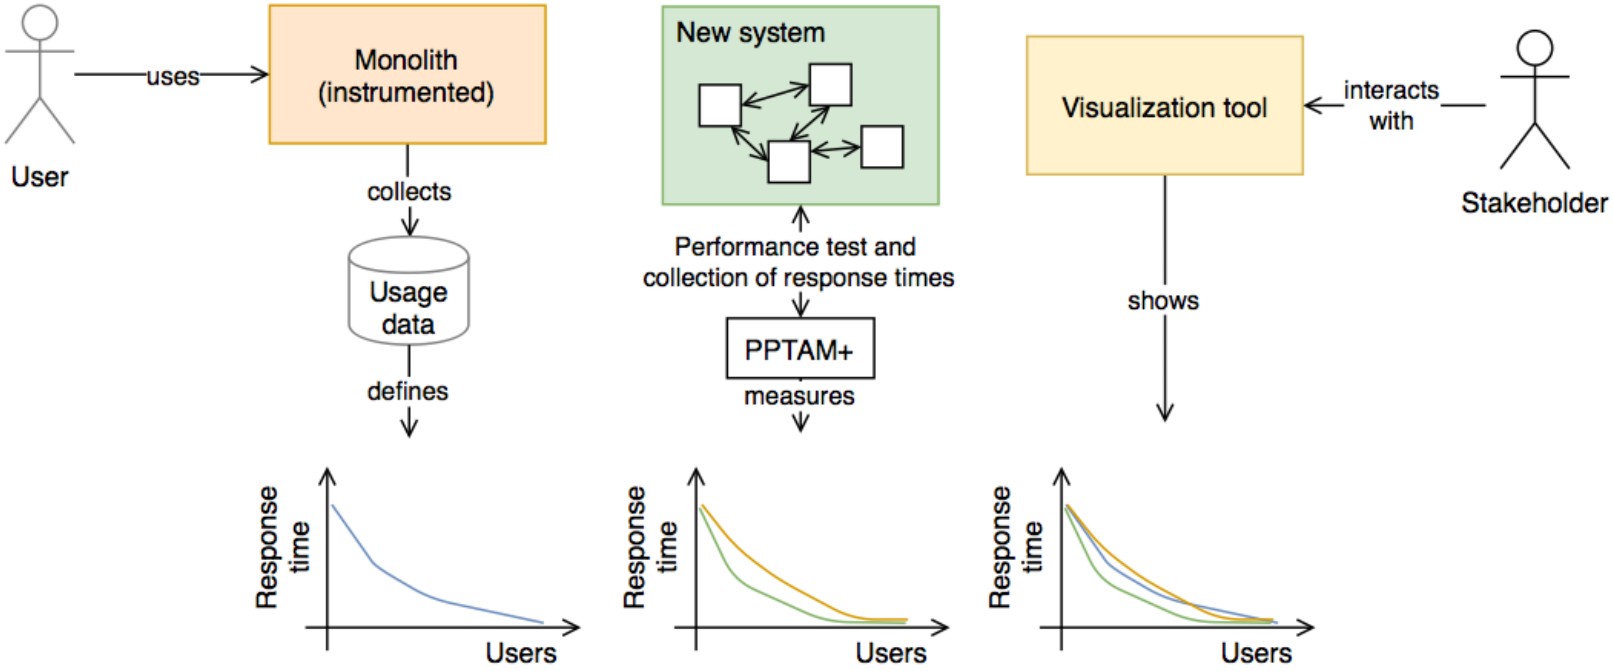
User
uses
Monolith
(instrumented)
collects
Usage
data
defines
New system
Performance test and
collection of response times
PPTAM+
measures
Visualization tool
interacts
with
Stakeholder
shows
Response
time
Users
Response
time
Users
Response
time
Users

Application to a transition

- If the new architecture **performs under a given threshold**, developers **stop and rethink** of the **architecture** or rethink the **used patterns** to guarantee that the new system - while having all advantages of a microservice architecture - does not fall short in terms of performance