

Digital Engineering • Universität Potsdan



Improving the Scalability and Security of Execution Environments for Auto-Graders in the Context of MOOCs

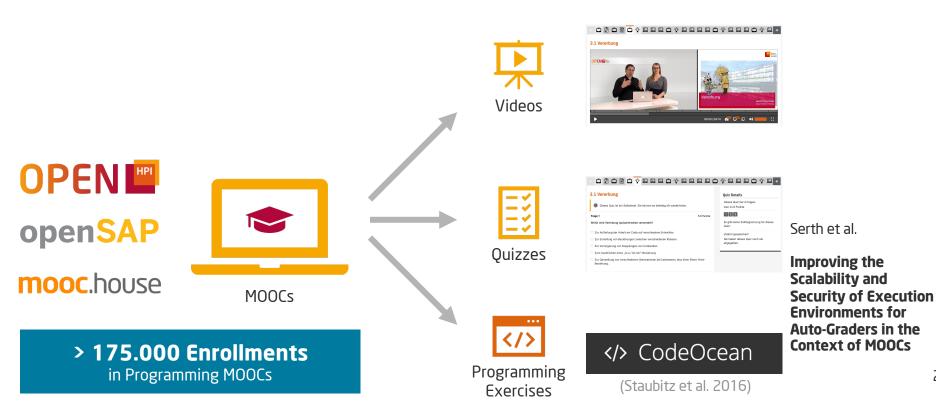
> Sebastian Serth, Daniel Köhler, Leonard Marschke, Felix Auringer, Konrad Hanff, Jan-Eric Hellenberg, Tobias Kantusch, Maximilan Paß, Christoph Meinel

> > Hasso Plattner Institute, University of Potsdam, Germany

## Programming Education with MOOCs



2



## CodeOcean Code Execution and Grading in the Browser



• • • • · ·	G 🔍 0	a codeocean.openhpi.de/exercises/717/imp	lement @	ى ھ	ů + C
	CodeOcean Administre	ation 👻	English 👻 Help	🚨 Sebastian Serth 👻	
<ul> <li>Test SPLICE - Guess The Nun omplete the game "Guess The Number".</li> <li>the computer chooses a number between 1 and</li> </ul>	mber Exercise	ild try to guess the number skillfully.			100%
he guessed number is too low, "too small" sh	ould be printed. If the guessed number is too hig	h, "too large" should be output. If the number is correct, the p	layer should be congratulated and t	then the game should be ended.	
plement the missing code snippets marked b	y a '# ToDo' in the source code. Also, pay attention	n to the programming style and satisfy the linter.			(Hide)
Collapse Action Sidebar	► Run			Collapse Output Sidebar	
Files  default.pylintrc  guess_the_number.py	1 from random import randint 2 3 4 def main(): 5 correct_number = randint(1, 100) 6 print("Welcome. Can you guess my no	number?')	Results 3 test files have been es	xecuted.	
test_theck_number.py     f     guessed_number = read_number()     while not check_number(quessed_number):			Test File 1 (test_c	:heck_number.py)	
test_read_number.py     test_style.py	<pre>10 guessed_number = road_number() 11 12 print('Would you like to play again?') 13 14 15 def check_number(guessed_number, correct_number): 16 (f guessed_number &lt; correct_number): 17 print('The guessed_number is too small.')</pre>	) in?') ect_number): r: to small.')	Passed Tests Score Feedback Error Messages	6 out of 6 5 out of 5 Well done. All tests have been passed.	
Tip 1: Input         18 - elif guessed, number > correct, number:           19   print("The guessed number is too large.")         28 - elif guessed, number - correct, number:           21   print("Congratuations, you have guessed the correct number (correct, number)         29 - elif guessed, number - correct, number		too large.") mber:	Test File 2 (test_r		
초 Download இ Reset this file	22 return True 23 return False 24 25 26 r def read number():		Passed Tests Score	3 out of 3 5 out of 5	
Rι	abij	y not a valid number.	ore	done. All tests have been passed.	

Serth et al.

## CodeOcean Interactive Elements and Detailed Score Output

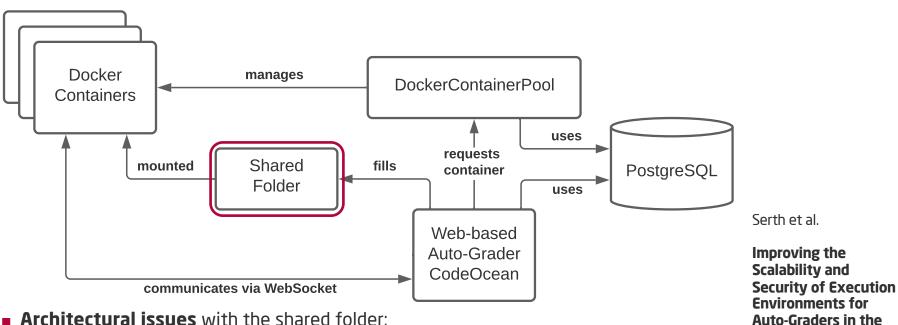


Chart 4

Your guess, plea Welcome. Can yo	see Send	
Linter Feedba Code Rating Score	ck (test_style.py) 4.64 out of 10 0.46 out of 1	
Score Feedback	0.46 out of 1 Es gibt noch einige Hinweise vom Linter.	
Messages	<ul> <li>convention: <ul> <li>superfluous-parens: Unnecessary parens after 'if' keyword (aufgabe314.py: 27)</li> <li>wrong-import-position: Import "from turtle import *" should be placed at the top of the module (aufgabe314.py: 18)</li> <li>multiple-statements: More than one statement on a single line (aufgabe314.py: 25, quadrat())</li> </ul> </li> <li>error: <ul> <li>function-redefined: function already defined line 2 (aufgabe314.py: 5, dreieck())</li> <li>warning: <ul> <li>redefined-outer-name: Redefining name 'winkel' from outer scope (line 1) (aufgabe314.py: 6, dreieck())</li> <li>redefined-builtin: Redefining built-in 'len' (aufgabe314.py: 14, dreieck())</li> <li>unreachable: Unreachable code (aufgabe314.py: 17, dreieck())</li> <li>unused-variable: Unused variable 'len' (aufgabe314.py: 14, dreieck())</li> <li>pointless-string-statement: String statement has no effect</li> </ul> </li> </ul></li></ul>	Serth et al. Improving the Scalability and Security of Execution Environments for Auto-Graders in the Context of MOOCs

## CodeOcean Current System Architecture





- Architectural issues with the shared folder:
  - Difficult horizontal scaling
  - Increased attack surface

Context of MOOCs

## Requirements Analysis Approach

- User-centered Design Thinking approach: (Meinel et al. 2016)
  - 1. Interviews with different stakeholder representatives
    - a) Learners
    - b) Teachers
    - c) Administrators
  - 2. Derive personas to visualize user needs and their pain points
  - 3. Decide on a subset of features to address
  - 4. Evaluate technical solutions for these personas
- Analysis of past executions on CodeOcean since 2015

Serth et al.



#### 7

#### Related Work Security Considerations

Assessing code submissions (Garmann 2013)

#### Dynamic code analysis

- Code under investigation is executed
  - Security implications
    - <u>Language-specific</u> using the Java Security Manager (Strickroth 2019)

#### <u>Container-based</u>

using Docker (Breitner et al. 2016)

#### Static code analysis

- Code inspection without executing the code under investigation
- →Other scalability and security considerations





# 150,000 100,000 Submissions per day 50,000 0 2016 2017 2018 2019 2020 2021

## Requirements Analysis Past Performance



Serth et al.

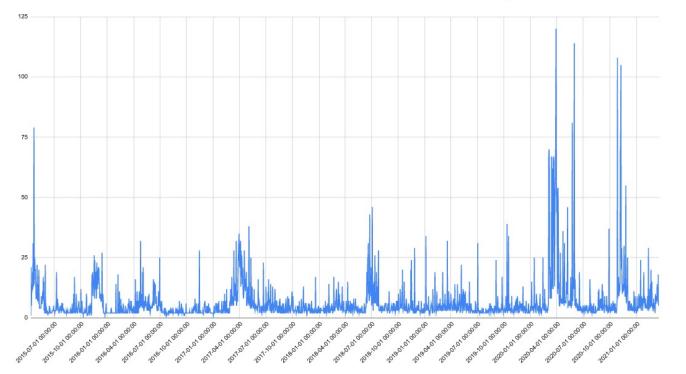
Improving the Scalability and Security of Execution Environments for Auto-Graders in the Context of MOOCs

8

## Requirements Analysis Past Performance



#### Max of total submissions per second per day



Serth et al.

#### 10

#### **Identified Requirements**

#### 1. Interactivity

- □ Synchronous channel between the execution and the users' web browser
- Real-time code executions

#### **2.** Scalability

- Current system handles up to 120 execution requests per second
- Mean execution time less than 10 seconds per execution

#### **3.** Flexibility

Compatibility to containerd ecosystem (e.g., through a Dockerfile)

Serth et al.



## **Identified Requirements**

- 1. Interactivity
- 2. Scalability
- **3.** Flexibility
- Limit resources
  - □ CPU: Wall clock time, CPU shares

 $\square$  RAM

Network access

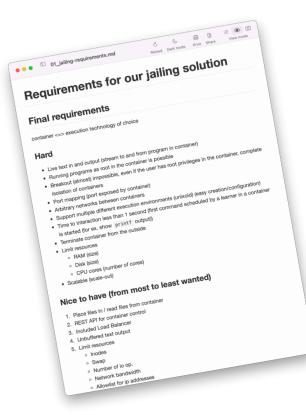
- Isolation between different executions
- Desired: Network orchestration with multiple nodes



Serth et al.

## Evaluation *Criteria*





#### Features

- Attach to stdout, stderr, stdin
- Allow superuser access within the execution

#### Metrics

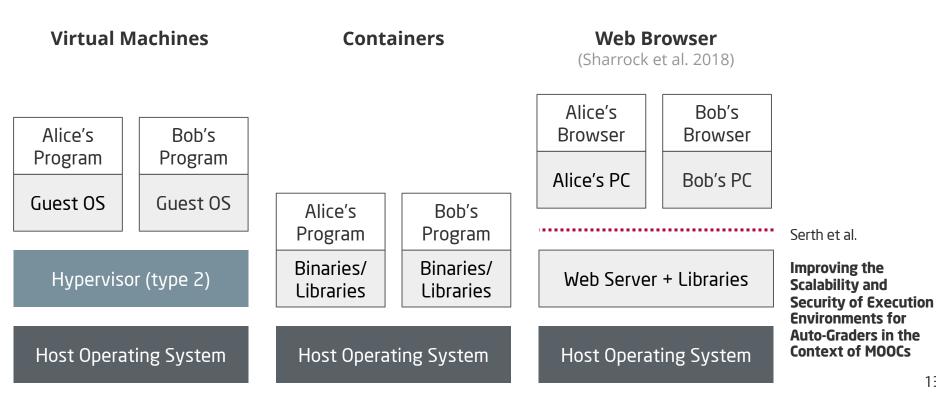
- Startup time
- RAM / CPU usage

Serth et al.

- Technical details
  - Maintenance and management
  - Isolation technologies and security implications

### Related Work Execution Environments





**Execution Environments** 



**Excluded by the requirements** 

Pouch
Pouch





VirtualBox

Approved for performance evaluation



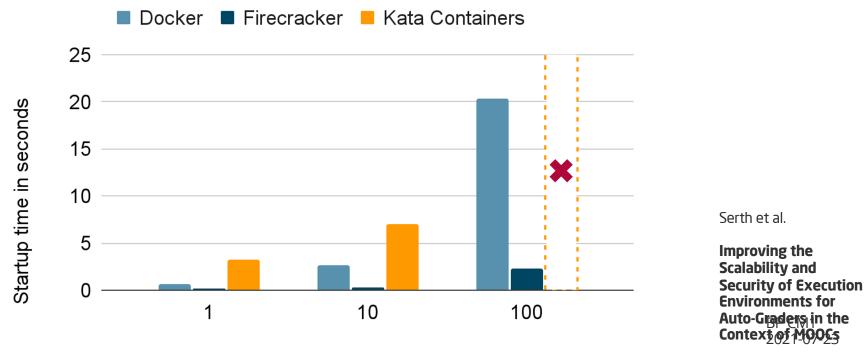




Serth et al.

## Execution Environments Startup time



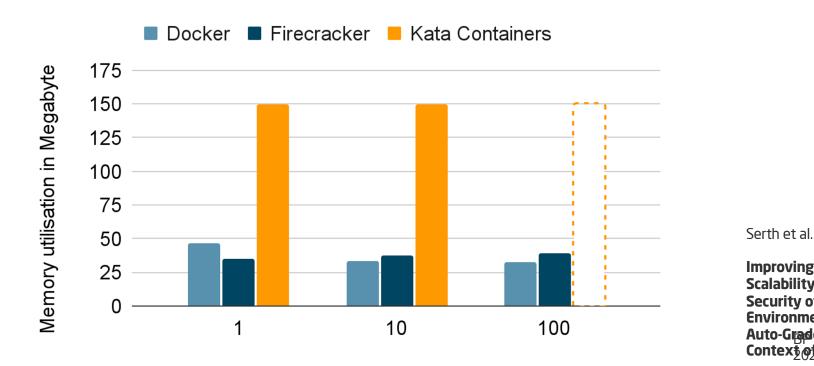


No. of executions

Chart 15

## **Execution Environments** Memory usage





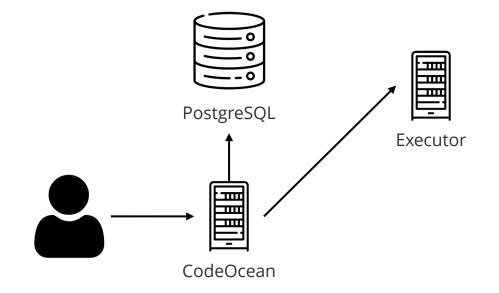
No. of executions

Improving the **Scalability and Security of Execution Environments for** Auto-Graders in the Context of MOOCs

Chart 16

# Requirement Evaluation *Scalability*

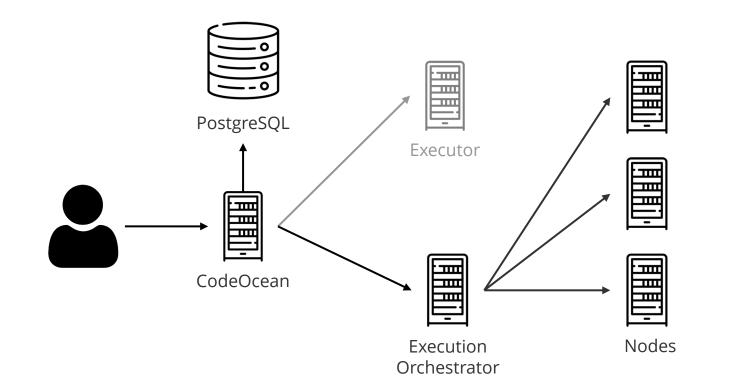




Serth et al.

# Requirement Evaluation *Scalability*





Serth et al.

**Execution Orchestrators** 

**Knative** 

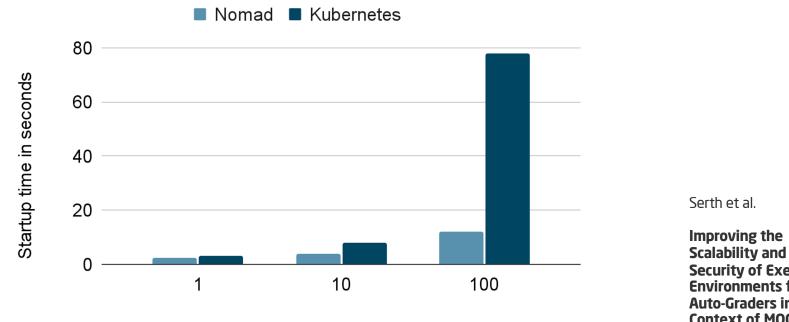


**Excluded by the requirements Approved for performance** evaluation fission Nomad **W** Kubeless **kubernetes OPENFAAS** 

Serth et al.

## **Execution Orchestrators** Nomad vs Kubernetes



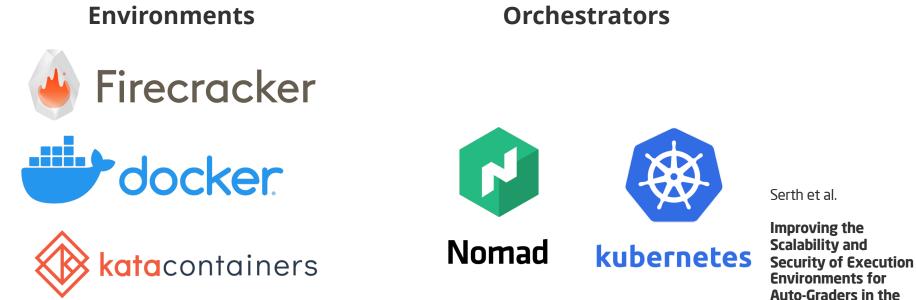


No. of containers

**Security of Execution Environments for** Auto-Graders in the **Context of MOOCs** 

Evaluation Comparison

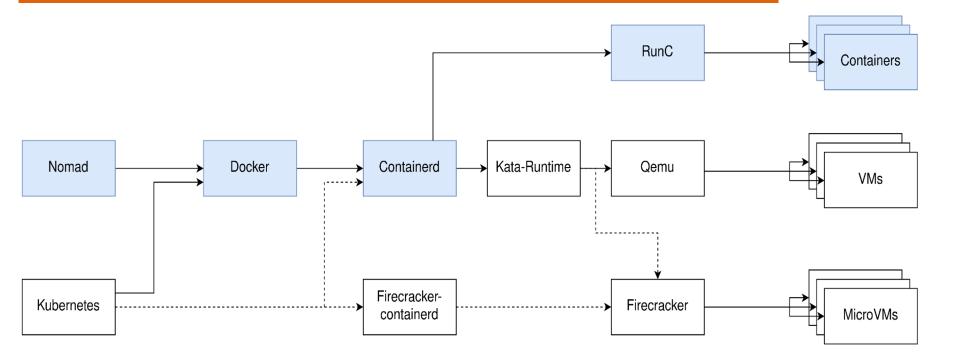




**Environments for** Auto-Graders in the **Context of MOOCs** 

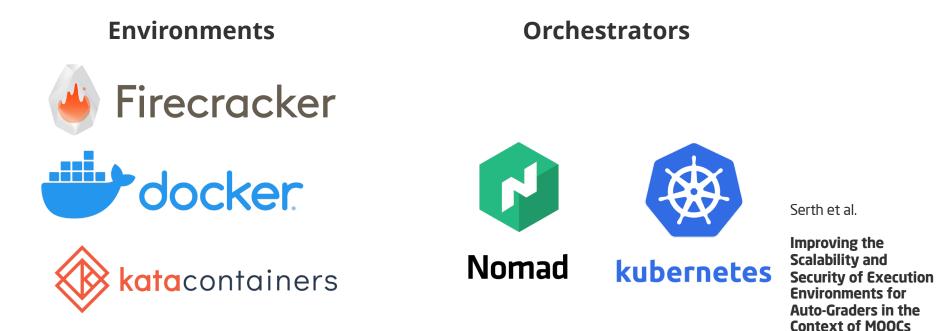
# Requirement Evaluation *Flexibility*





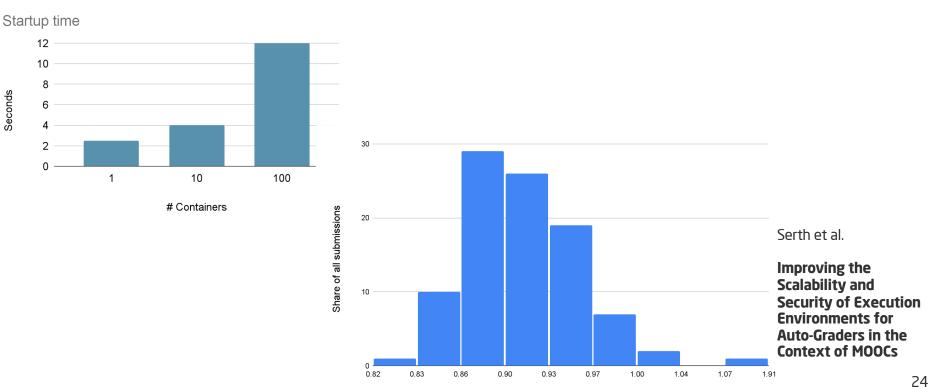
Evaluation *Comparison* 





23

## **Evaluation** Prewarming

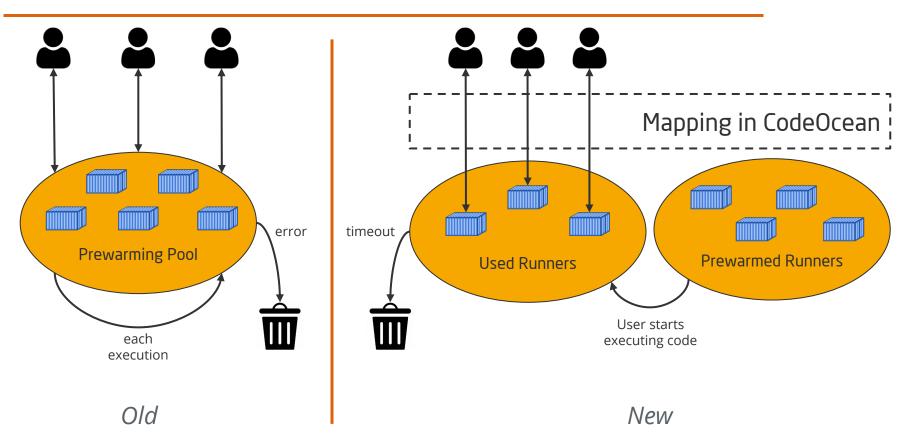


Hasso HPI Plattner Institut

Response time in seconds

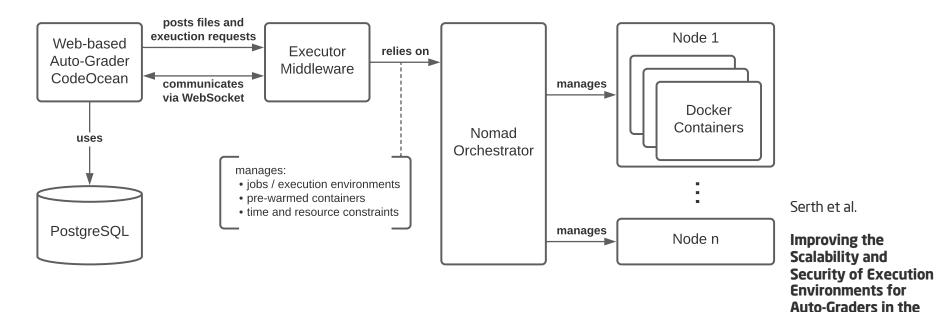
## Evaluation Container Reusing





# CodeOcean *New System Architecture*





Context of MOOCs

## Future Work





- Integration of hardware resources
- □ GPU for machine learning exercises (for teaching purposes at HPI)
- Raspberry PI for embedded smart home courses on openHPI
- Networking scenarios
- Multi-node setup and VPN access
- Long lasting executions



- Pre-warming strategies and Function-as-a-Service (FaaS)



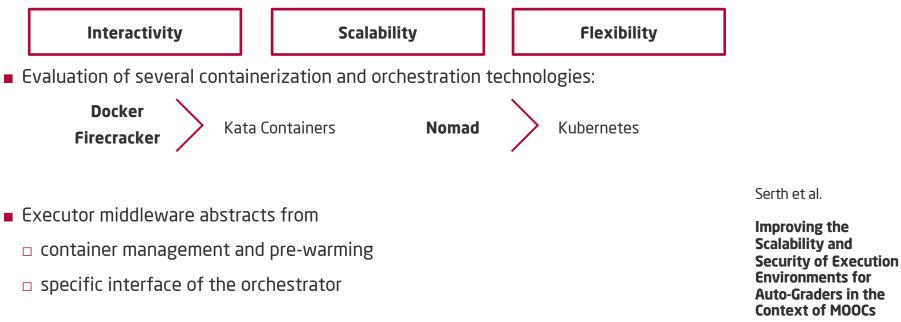
Comparison of programming languages and their security concepts for system components (especially the executor middleware)

Serth et al.



#### Conclusion

#### Identification of three main requirements:



#### → Result: Scalable and more secure architecture for execution environments



Digital Engineering • Universität Potsdan



Improving the Scalability and Security of Execution Environments for Auto-Graders in the Context of MOOCs

> Sebastian Serth, Daniel Köhler, Leonard Marschke, Felix Auringer, Konrad Hanff, Jan-Eric Hellenberg, Tobias Kantusch, Maximilan Paß, Christoph Meinel

> > Hasso Plattner Institute, University of Potsdam, Germany