



Trends and Concepts in the Software Industry I

From On-Premise to the Cloud

Prof. Dr. Hasso Plattner, Dr. Michael Perscheid
Enterprise Platform and Integration Concepts
Hasso-Plattner-Institut

Goals

Trends and Concepts in the Software Industry I

Deep technical understanding of trends and concepts in enterprise computing, esp. main-memory-centric data management on modern hardware, cloud-native development, intelligent enterprises and their impact on applications.

■ Focus areas

- Principles of in-memory databases
- Characteristics and architecture of enterprise applications and systems
- Influence of cloud-native development
- Trends and challenges in enterprise computing
- Experience reports from industry
- Hands-on exercises and experiments

Course Overview

Block Week: September 2021

- General information
 - When: September (TBA, presumably hybrid format)
 - Lectures given by Prof. Plattner
 - Additional lectures by guests from industry
 - Discussions about open questions in enterprise computing with knowledgeable experts are a vital part of the lecture!
- Final grade consists of
 - Preparation quiz (mandatory)
 - Group work, presentation, and participation during the block week (40%)
 - Oral exam (60%)



Course Overview

General Information

- 6 ECTS points

- Latest enrollment: 28th of April 2021

- Modules

- IT-Systems Engineering MA

- ITSE-Konstruktion
- ITSE-Maintenance
- BPET-Konzepte und Methoden
- BPET-Spezialisierung
- BPET-Techniken und Werkzeuge
- OSIS-Konzepte und Methoden
- OSIS-Spezialisierung
- OSIS-Techniken und Werkzeuge

- Data Engineering MA

- DATA-Konzepte und Methoden
- DATA-Techniken und Werkzeuge
- DATA-Spezialisierung
- SCAL-Konzepte und Methode
- SCAL-Techniken und Werkzeuge
- SCAL-Spezialisierung

- Digital Health MA

- SCAD-Concepts and Methods
- SCAD-Technologies and Tools
- SCAD-Specialization
- APAD-Concepts and Methods
- APAD-Technologies and Tools
- APAD-Specialization

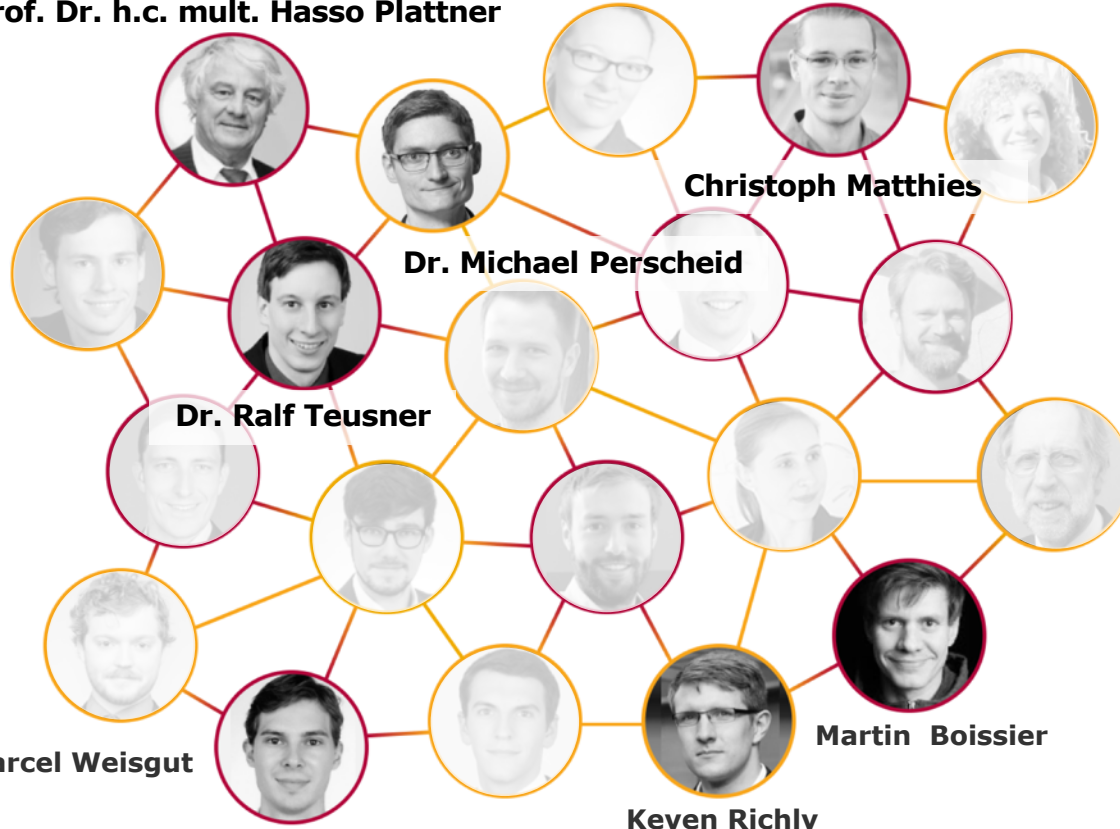
- Cybersecurity MA

- IDMG-Konzepte und Methoden
- IDMG-Techniken und Werkzeuge
- IDMG-Spezialisierung

Teaching Team

General Information

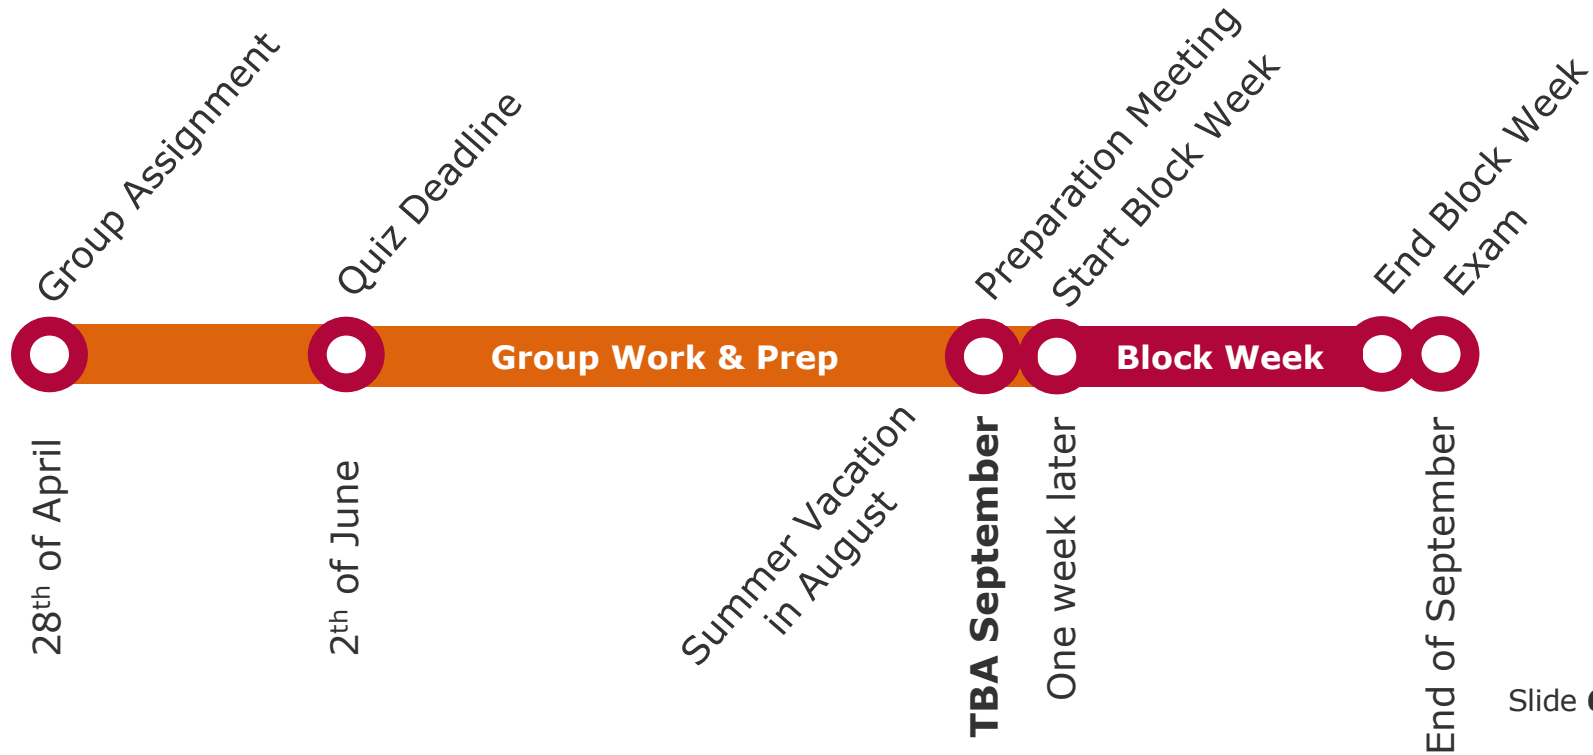
Prof. Dr. h.c. mult. Hasso Plattner



Thomas Bodner / **Marcel Weisgut**

Course Overview


Schedule



Preparation

A Course in In-Memory Data Management

- Get a solid understanding of the fundamentals of In-Memory Data Management
- Materials
 - Course book (distributed digitally)
 - openHPI course
<https://open.hpi.de/courses/tuk2021>
- Mandatory quiz
 - Start: 28th of April
 - Deadline: 2nd of June



Preparation

Further Reading

Plattner & Leukert:
The In-Memory Revolution
How SAP HANA Enables Business of the Future



Reznik, Dobson & Gienow:
Cloud Native Transformation:
Practical Patterns for Innovation

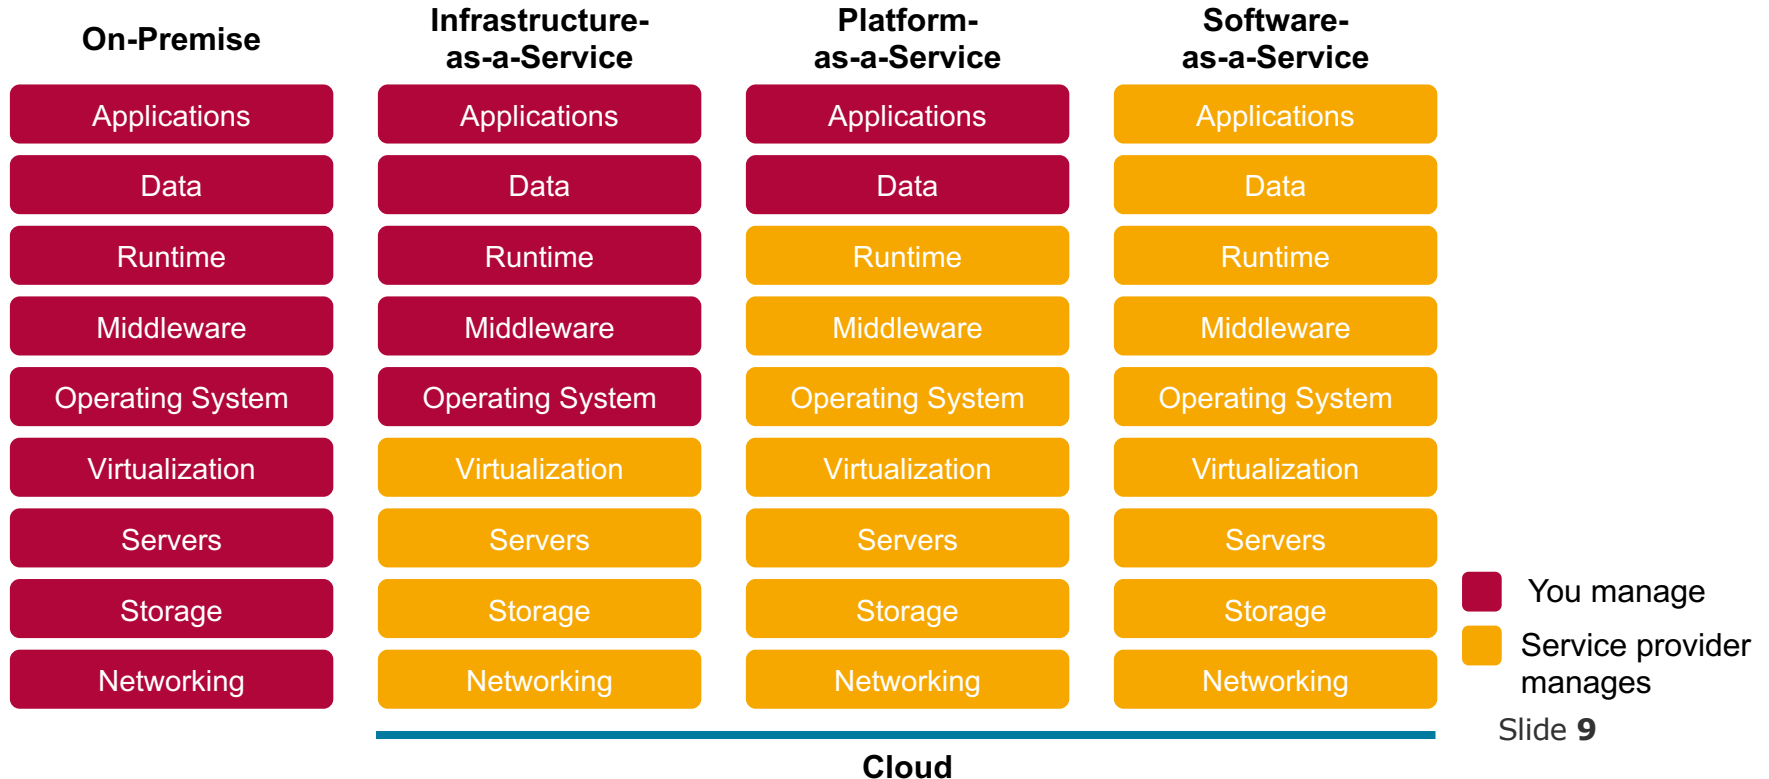



Scholl, Swanson & Jausovec:
Cloud Native:
Using Containers, Functions, and Data
to Build Next-Generation Applications



Trends and Concepts in the Software Industry I

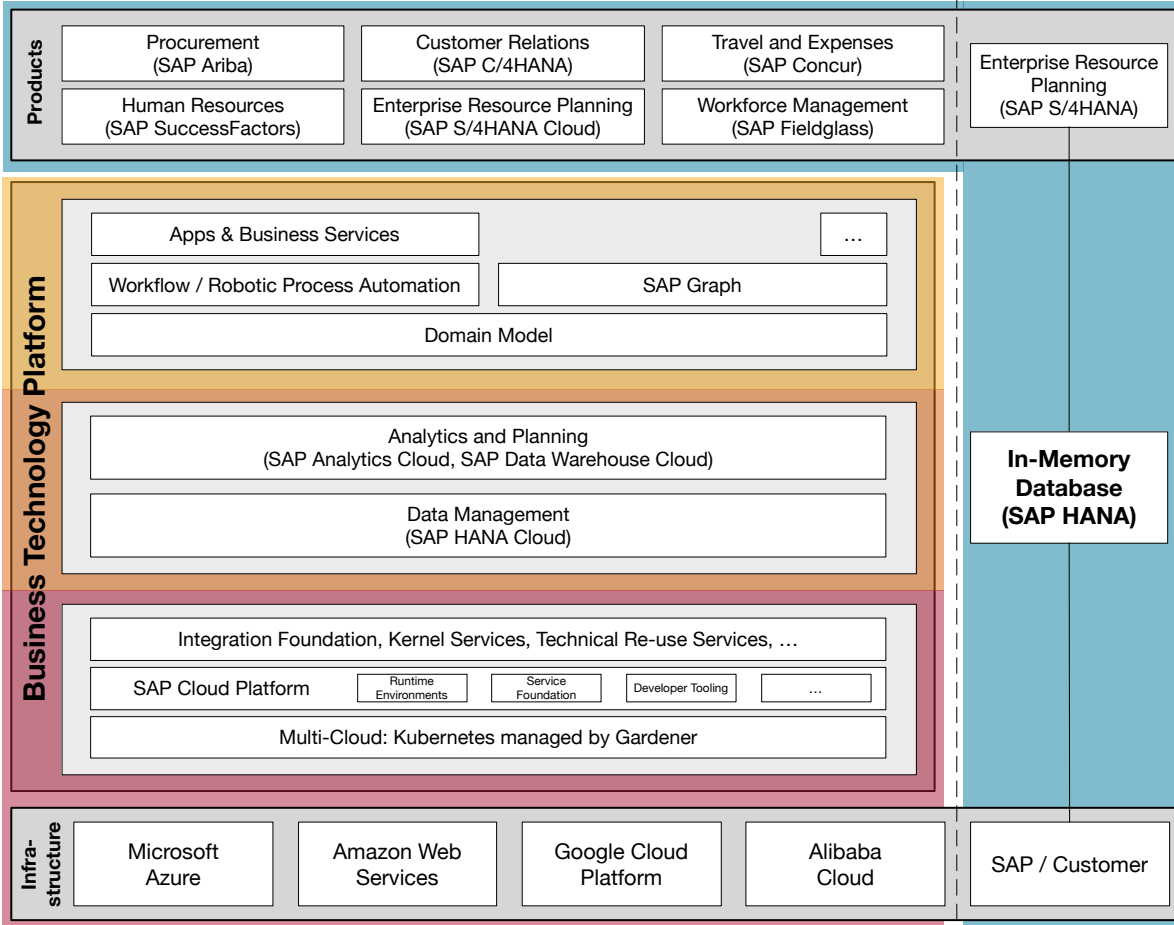
From On-Premise to the Cloud





Cloud

On Premise



Day 1: The In-Memory Revolution

Day 2: Cloud Impact on Enterprise Architecture

Day 3: In-Memory Goes Cloud

Day 4: Composing Cloud Applications

Group Work

Trends and Concepts in the Software Industry I

- Preparation of interactive group part
 - Teams of 6 to 8 students
 - Regular meetings
 - Team assignment: 28th of April
- Hands-on experiments
 - Familiarization with existing research
 - Implementation parts
 - Evaluation of the results
 - Presentation in the block week (~30 minutes)
- Tell us your topic preference:
<https://forms.gle/fRmZTTsVLGZ3fUgm9>



Group Work – Topic 1

Database Clustering

Motivation

Clustering data in *clusters with similar data characteristics* allows for various performance improvements. When searching on sorted data, *binary searches* can be used over linear searches. Moreover, clustering often enables the database to skip large parts of the data without ever looking at it. However, finding a good clustering configuration is not trivial.

Challenges


- Understand the storage layout and architecture of modern in-memory systems
- Understand where clustering might be beneficial and in which cases not
- Analyze a given workload and automatically determine a good clustering scheme

Learnings

- First looks into the in-memory database Hyrise and its storage engine
- Proper Benchmarking
- How to make automated tuning decisions

Requirements

- Basic database knowledge (partitioning, filtering, SQL)
- Mostly scripting using Python and SQL



https://en.wikipedia.org/wiki/Cluster_analysis

Group Work – Topic 2

Where Are Our Failures Located? - An Analysis of Error Reports

Motivation

The development of complex software systems is challenging and consequently not failure-free. The correction of software errors is costly and often requires time-consuming debugging sessions. Detailed knowledge about the evolution of defect reports and where failures are located (e.g., components, services, algorithms) allow us to understand our systems better and prevent future failures.

Challenges – Questions to be answered


- Understand failures and defects in real-world datasets of error tracking systems
- Develop visualization and analysis concepts to address questions such as:
 - Where are errors located?
 - How long does a failure exist?
 - How often has a ticket been reopened and how many authors are involved?
 - Are there interrelationships between failures?

Learnings

- SQL and scripting language
- Data analysis and visualization concepts

Requirements

- Basic programming skills
- Basic statistical knowledge



<https://www.mediawiki.org/wiki/Bugzilla>

Group Work – Topic 3

Cloud Analytics: Choosing your Database System in the Cloud



Motivation

For any given application workload, there today is a variety of cloud-based database offerings. The underlying database systems are based on different architectures with respective tradeoffs. These need to be understood for an educated choice between the offerings. This year's focus is on analytics.



Challenges

- Run a representative analytics workload on various cloud database offerings and interpret the results with respect to time and cost
- Tune a database system for a given workload, rinse and repeat
- Understand the architectures and tradeoffs of current database systems, i.e, be able to decide when to use what and to explain why



Learnings

- Hands-on experience with modern cloud databases
- Proper benchmarking

Requirements

- Basic database knowledge is expected

Group Work – Topic 4

Workflow Engines: Building an ERP with BPMN

Motivation

When starting (and growing) a new company, there is a need for software support. Instead of buying a complete ERP product or relying on Excel sheets, business processes can be modeled and executed using *process engines* that automate and orchestrate workflows.

Challenges

- Get to know the workings of the open-source Camunda Platform process engine
- Implement a common SAP business process using the process engine
- Build the services and endpoints that are required and interacted with (in the languages of your choice)




Learnings

- Hands-on experience with a production process engine
- Understanding of real business processes

Requirements

- Basic knowledge of BPMN is expected



Commercial Break

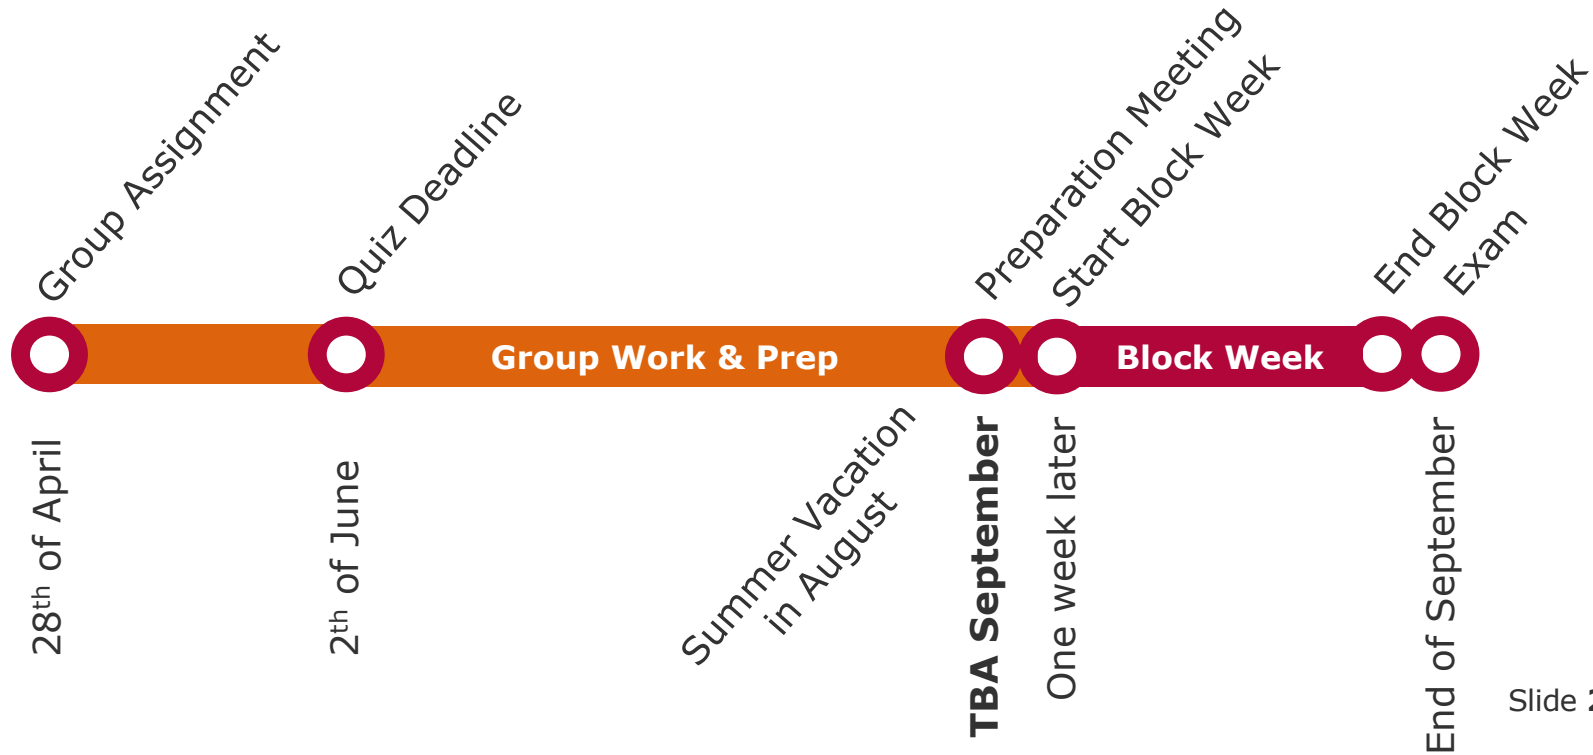
Trends and Concepts in the Software Industry II



Your task is to analyze how customers currently develop with SAP cloud technology. Based on that, you rethink how they can improve their efficiency when developing with cloud technology under the assumption that any data is available globally with (almost) zero response time.

Course Overview

Schedule



Tell Us Your Topic Preference

Trends and Concepts in the Software Industry I - Exercises



<https://forms.gle/fRmZTTsVLGZ3fUgm9>

Contact

- Dr. Michael Perscheid
 - Email: michael.perscheid@hpi.de
- Dr. Ralf Teusner
 - Email: ralf.teusner@hpi.de

HPI Campus II
Villa, Room: V-2.18

