

Enterprise Platform and Integration Concepts: Research Group of Prof. Dr. Hasso Plattner

# HANA - Running Enterprise Applications on an In-Memory Database

## Motivation

Earlier bachelor projects focused on how an in-memory data layer can change the way “traditional” transactional enterprise applications work. By eliminating redundantly stored aggregates, views, and indices, the applications became 10-100 times faster while extending application functionality and improving flexibility. Furthermore, the in-memory data management enables analytics on the transactional data and eliminates costly transformations to external data warehouses.

Consequently, a set of totally new applications becomes feasible with reduced total cost of ownership as a side effect. These applications enable businesses to get real-time insights, make more confident decisions, anticipate and react to changing conditions, and to become more flexible while being more effective by taking advantage of their data.

## Goal

Based on disruptive technology changes introduced by SAP's in-memory database HANA, enterprise applications are adjusted for new programming paradigms while old constraints fall away. The project will rethink the way enterprise applications are designed and how they can be built leveraging new patterns and algorithms for in-memory databases, modern hardware architectures with large amounts of main memory, and multi-core processors. For instance, the project shall evaluate the idea of transparent data life-cycle





management that allows partitioning of transactional data, based on workload usage. Furthermore, we will focus on the interface between applications and databases and investigate how object relational mappers can be used for analytical scenarios. The project team will implement a prototype application with a web-based user interface on top of our research prototype HYRISE or SAP's in-memory database HANA. During the project, the team will be supported by SAP finding a customer that can provide an use case as well as enterprise data set.

## External Partner

The project will be executed in cooperation with SAP AG and it is likely that we will stay and work for a couple of days in Walldorf.

## Setting

The project team will work on newest server hardware with in-memory and multi-core technology. Thanks to our cooperations, we are able to access new high-end hardware long before it is available on the public market. Additionally, the project will include working with data of real customers and multiple terabytes of data.

## Skills

During the project you will learn the basics of in-memory databases and how they differ from traditional disk based systems. You will get insights into SAP and how businesses use their transactional data to support decision making. We do expect you to have deep interest on databases systems as well as front-end/ web- or C++ development.

## Group Structure and Project Start

The team will consist of 6-8 students. Project start will be the 16nd of October 2013.

## Contact

You are welcome to come by at room V2.02 in the "Villa" or write us an email. Besides, on 22nd of July at 2pm there is an upfront meeting at room V2.16.

Prof. Dr. Hasso Plattner ([office-epic@hpi.uni-potsdam.de](mailto:office-epic@hpi.uni-potsdam.de))

Jens Krüger ([jens.krueger@hpi.uni-potsdam.de](mailto:jens.krueger@hpi.uni-potsdam.de))

Carsten Meyer ([carsten.meyer@hpi.uni-potsdam.de](mailto:carsten.meyer@hpi.uni-potsdam.de))

Martin Boissier ([martin.boissier@hpi.uni-potsdam.de](mailto:martin.boissier@hpi.uni-potsdam.de))