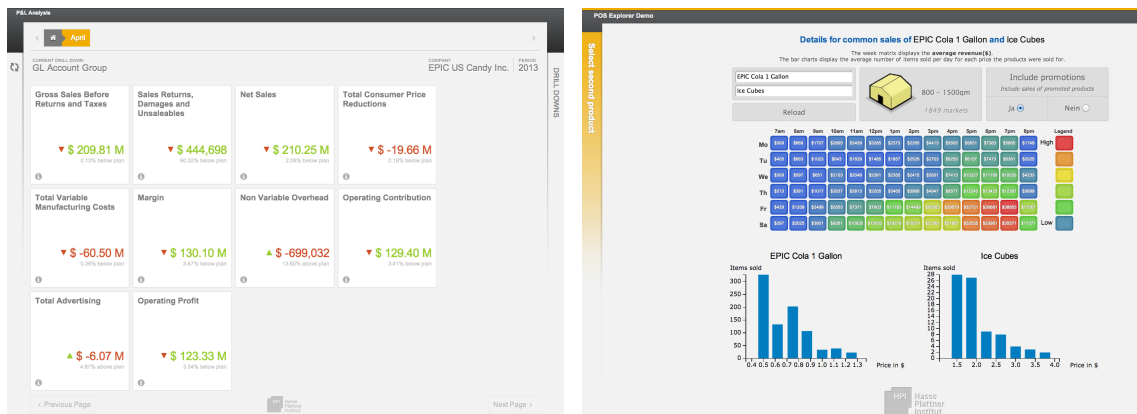


Enterprise Platform and Integration Concepts: Research Group of Prof. Dr. Hasso Plattner

Advanced Enterprise Applications using In-Memory Databases

Motivation

Previous bachelor projects focused on how an in-memory data layer can change the way “traditional” transactional enterprise applications work. By simplifying existing data models, e.g. by eliminating redundantly stored aggregates, views, and indices, the applications became 10-100 times faster. Furthermore, an in-memory database enables analytics on the transactional data and eliminates costly transformations to external data warehouses (known as ETL process).



As a result, new applications which previously were infeasible can now be developed. These applications enable businesses to get real-time insights, make more confident decisions, anticipate and react to changing conditions, and become more flexible while being more effective.

Goal

The introduction of SAP's in-memory database HANA had a disruptive impact on the industry. 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. The team will analyze existing business processes and identify areas for optimization. We strive to have a co-innovation like environment with close and frequent contact to partners and customers. With the support of SAP, the chosen business process will be rethought and the team will implement a prototype application to prove their concept. The prototype will be developed using modern web technologies and shall run on top of our research prototype HYRISE or SAP's in-memory database HANA.

External Partner

The project will be executed in cooperation with SAP AG and potentially one of their customers (will be announced in September). There might be an opportunity to travel to the headquarter in Walldorf for deeper discussions and knowledge exchange.

Setting

The project team will work on newest server hardware with in-memory and multi-core technology provided by the Enterprise Application Architecture Laboratory. The laboratory builds the foundation for activities with in-memory databases and enterprise applications at our group and the HPI. 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 fundamentals 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 not expect you to have deep knowledge about databases. Additionally, existing skills in Python/C++ or web technologies like JavaScript/CSS will be helpful.

Group Structure and Project Start

The team will consist of 6-8 students. Project start will be the 14th of October 2014.

Contact

You are welcome to come by at room V2.03 in the "Villa" or write us an email.

Prof. Dr. Hasso Plattner (office-epic@hpi.de)

Matthias Uflacker (matthias.uflacker@hpi.de)

Lars Butzmann, Martin Faust, Stefan Klauck, David Schwalb