

In-Memory Computing and the Future of Business Software

Dr. Stefan Sigg, SAP AG
February, 2011





Agenda

Technology

Business

Applications



Technology

Basic Facts About In-Memory Computing



New Computing Power for Better Business Results

“Imagine you feel hungry. But instead of just opening the fridge (imagine you don’t have one) to get hold of, say, some butter and cheese, you would have to leave the house for the nearest dairy farm. Each time you feel hungry. This is what we do today with most company data: We keep them far away from where we process them...” *



Prof. Norbert Walter


(former Chief Economist of Deutsche Bank AG)

* Quote about the new book “In-Memory Data Management” by H. Plattner and A. Zeier

A Shift of Frontiers in Computer Science


Freely Adapted from Jim Gray, Turing Award Winner 1998

- ❖ **Tape is Dead**
- ❖ **Disk is Tape**
- ❖ **Main Memory is Disk**
- ❖ **CPU Cache is Main Memory**



Orders of Magnitudes

Programming Against a New Scarce Resource...



Type of Memory	Size	Latency (~)
L1 CPU Cache	64K	1 ns
L2 CPU Cache	256K	5 ns
L3 CPU Cache	8M	20 ns
Main Memory	GBs up to TBs	100ns
Disk	TBs	>1.000.000 ns

... requires cache-conscious data-structures and algorithms.


In-Memory Computing vs. Caching

Traditional Caching

- **Use Disk structures and put them into Main Memory (Cache) for speed**
- **Runtime structures come from disk, not optimized for RAM**

True In-Memory Computing

- **Always use Main Memory structures and store them on disk for backup purposes only**
- **Runtime structures optimized / compressed for RAM**



Two Other Major Qualities


... Besides In-Memory Processing

1. Massively Parallel Processing

2. Columnar Storage

Massively Parallel Processing (MPP)

Scalability Counts



Data Distribution



- **RAM locality** – data gets spread out to all available cores
- **MPP execution** – blades share nothing when crunching large data sets
- **Failover** - Individual blades may fail without causing problems

Columnar Storage


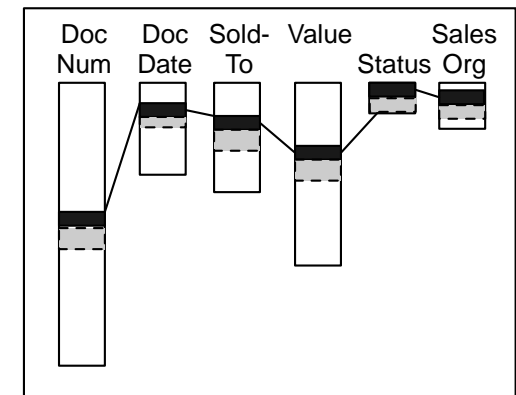
Do More With Less (Space)

- ❖ Compression (by factors)
- ❖ Read relevant data only
- ❖ Main Memory Structures
- ❖ Super-Efficient Algorithms
- ❖ Insert Only Works Fine
- ❖ Adding fields is easy

Row Store



Column Store



Benefits for Business Software

(Re-) Unification of Transactional Processing and Analytics

A Common Database Approach for OLTP and OLAP Using an In-Memory Column Database



Hasso Plattner
Hasso Plattner Institute for IT Systems Engineering
University of Potsdam
Prof.-Dr.-Helmert-Str. 2-3
14482 Potsdam, Germany
hasso.plattner@hpi.uni-potsdam.de

Categories and Subject Descriptors

H.2.0 [Information Systems]: DATABASE MANAGEMENT—*General*

General Terms

Design, Performance

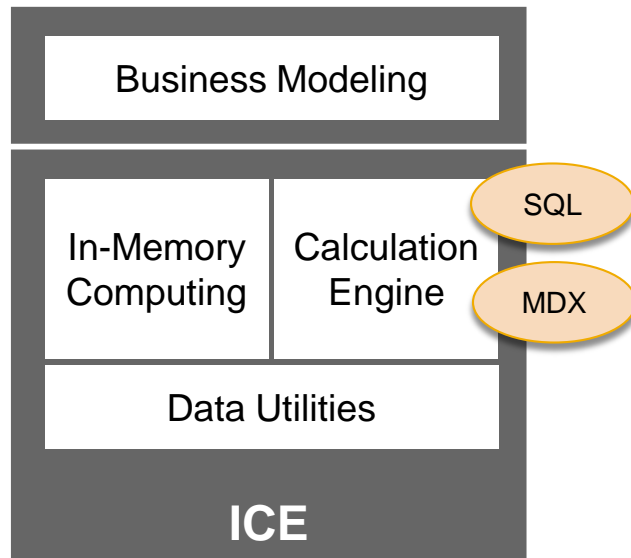
1. INTRODUCTION

Relational database systems have been the backbone of business applications for more than 20 years. We promised to provide companies with a management information system that covers the core applications, including financials, sales, order fulfillment, manufacturing, as well as human resources, which run from planning through business processes

After the conversion of attributes into integers, processing becomes faster. More recently, the use of column store databases for analytics has become quite popular. Dictionary compression on the database level and reading only those columns necessary to process a query speed up query processing significantly in the column store case.

I always believed the introduction of so-called data warehouses was a compromise. The flexibility and speed we gained had to be paid for with the additional management of extracting, and loading data, as well as controlling the redundancy. For many years, the discussion seemed to be closed and enterprise data was split into OLTP and OLAP [9]. OLTP is the necessary prerequisite for OLAP, however only with OLAP companies are able to understand their business and come to conclusions about how to steer and change course. When planned data and actual data are

Introducing the In-Memory Computing Engine

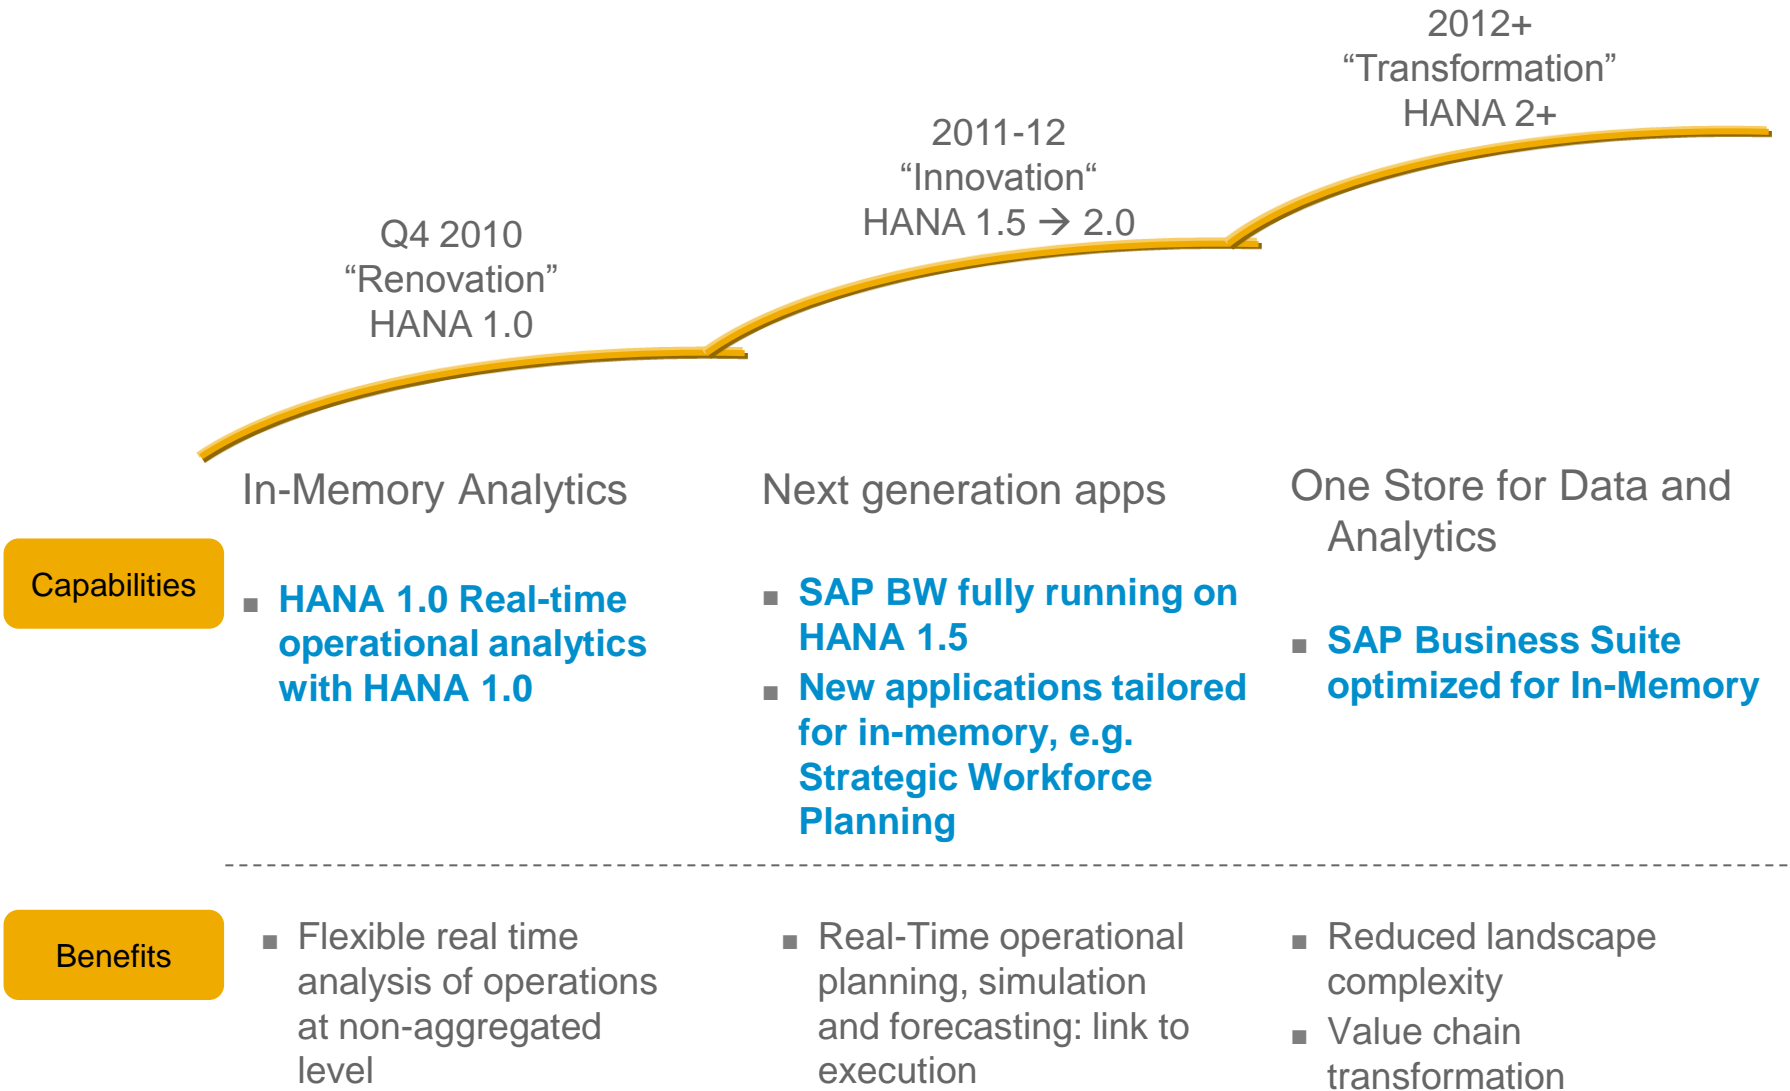


Transactional and analytic data acceleration and management in a single environment:

- a high performance in-memory computing engine
- a powerful and flexible data calculation engine
- supports SQL & MDX
- a unified information modeling design environment
- a data repository to persist views of business information

SAP In-Memory Strategy

Product Strategy



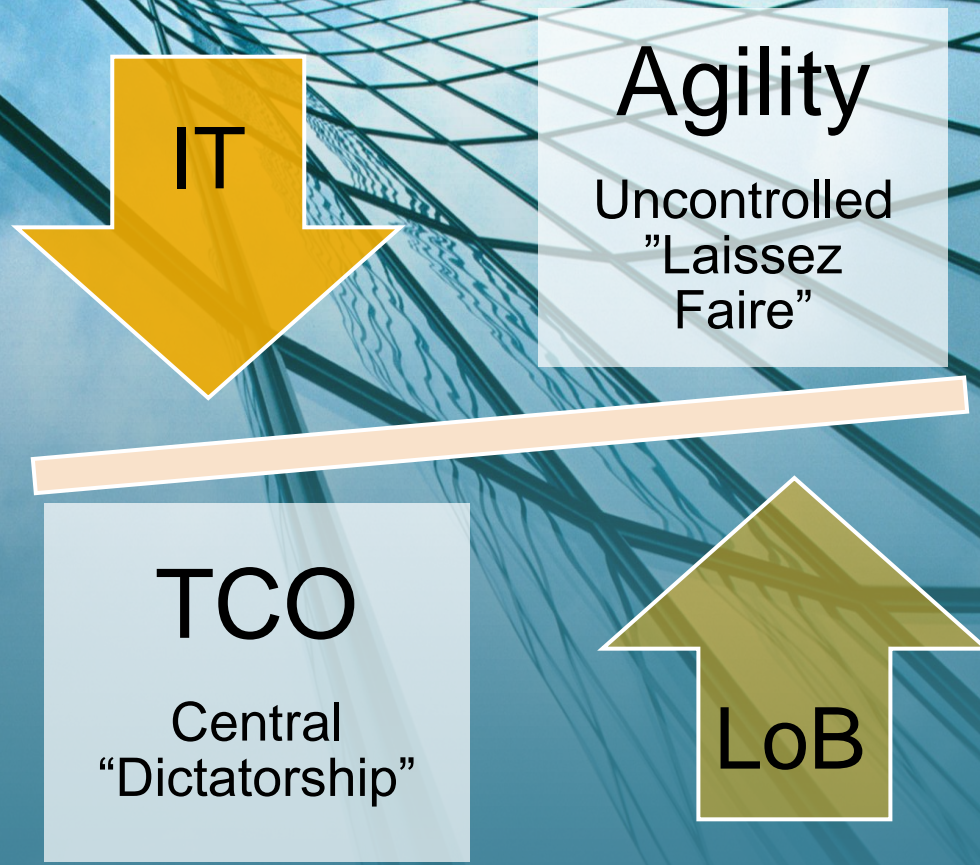


Business

Better Results for IT and Business



IT and Business – A Sensitive Equilibrium



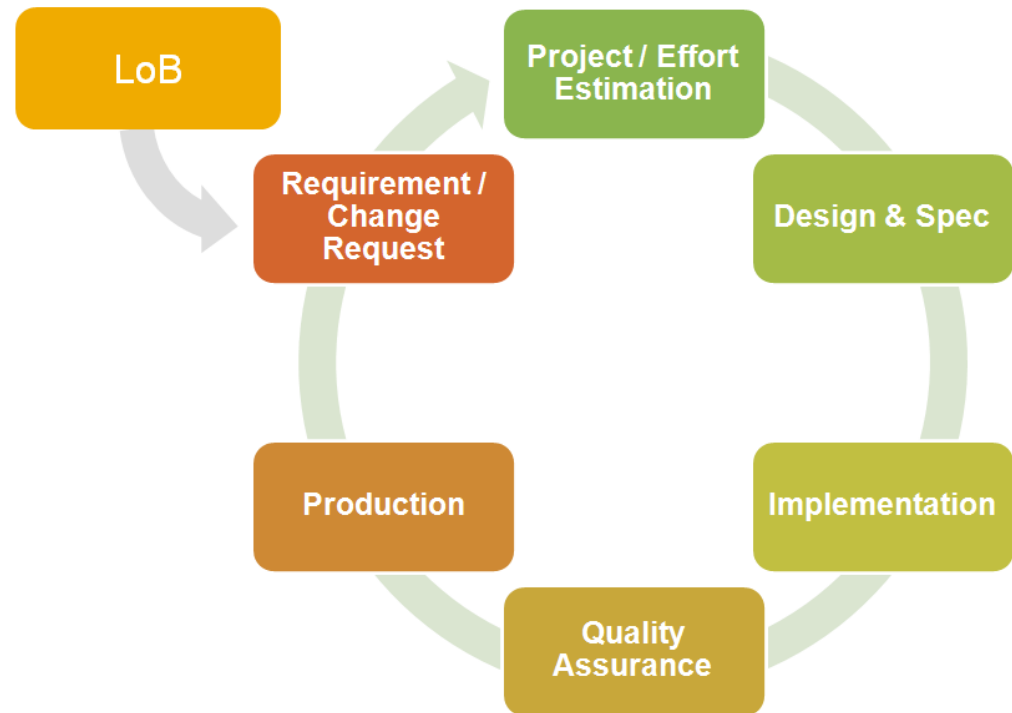
Finding the right trade-off is difficult

Self-Service BI

How Fast Can IT React on Change? – Empower Business People

How can In-Memory help?

- No aggregates
- No MOLAP cubes
- External hierarchies
- No re-alignment runs
- Virtual structures only
- Stable performance



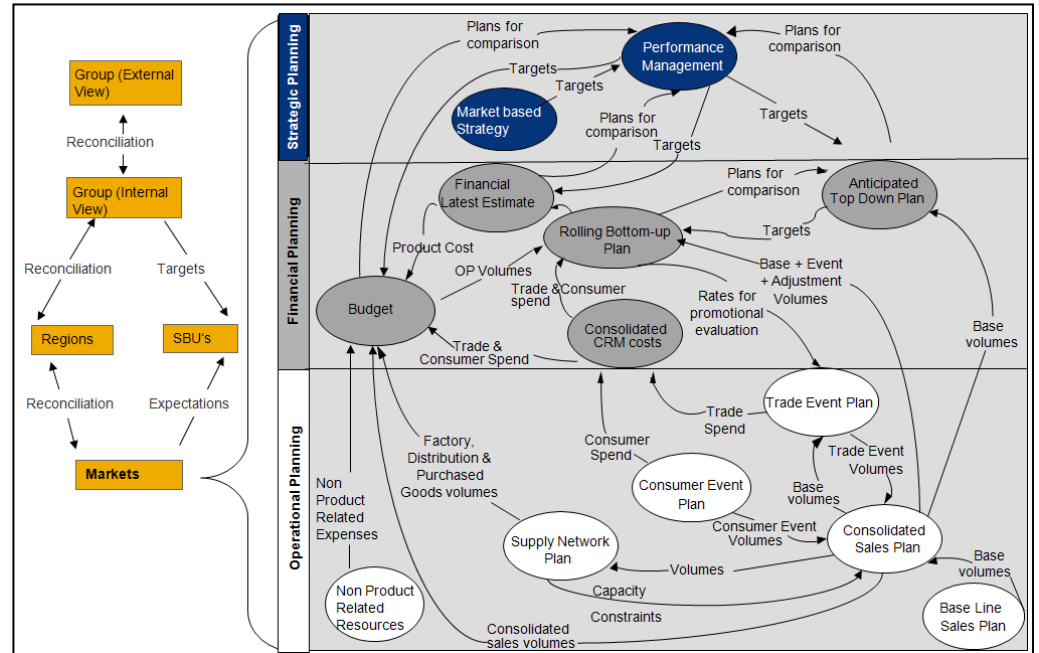
traditional change management process

Planning, Budgeting, and Forecasting

Better Planning Results, Faster

How can In-Memory help?

- Closer to real-time
- Work on entire data set
- Create local views
- Do not duplicate
- Use online forecasting algorithms
- Give design-time to LoB
- More, faster iterations
- Share master data, hierarchies, authorizations,...



Corporate planning processes

Demand / Stock Optimization

Analyze and Predict Demand – Optimize Inventory

Based on a real-life PoC project with a major CPG company



... – Replenishment – Assortment Planning – Demand Signal Analysis – ...

Entering New Dimensions How IT Can Enrich Business

“...The benefit for Hilti applying SAP’s in-memory technology is not only seen in a dramatic improvement of reporting execution speed - for example, we were able to speed up a reporting batch job from 3 hours to seconds – but even more in the opportunity to bring the way we work with information and ultimately how we service our customers on a new level. ”

Dr. Martin Petry



(CIO, Hilti AG)



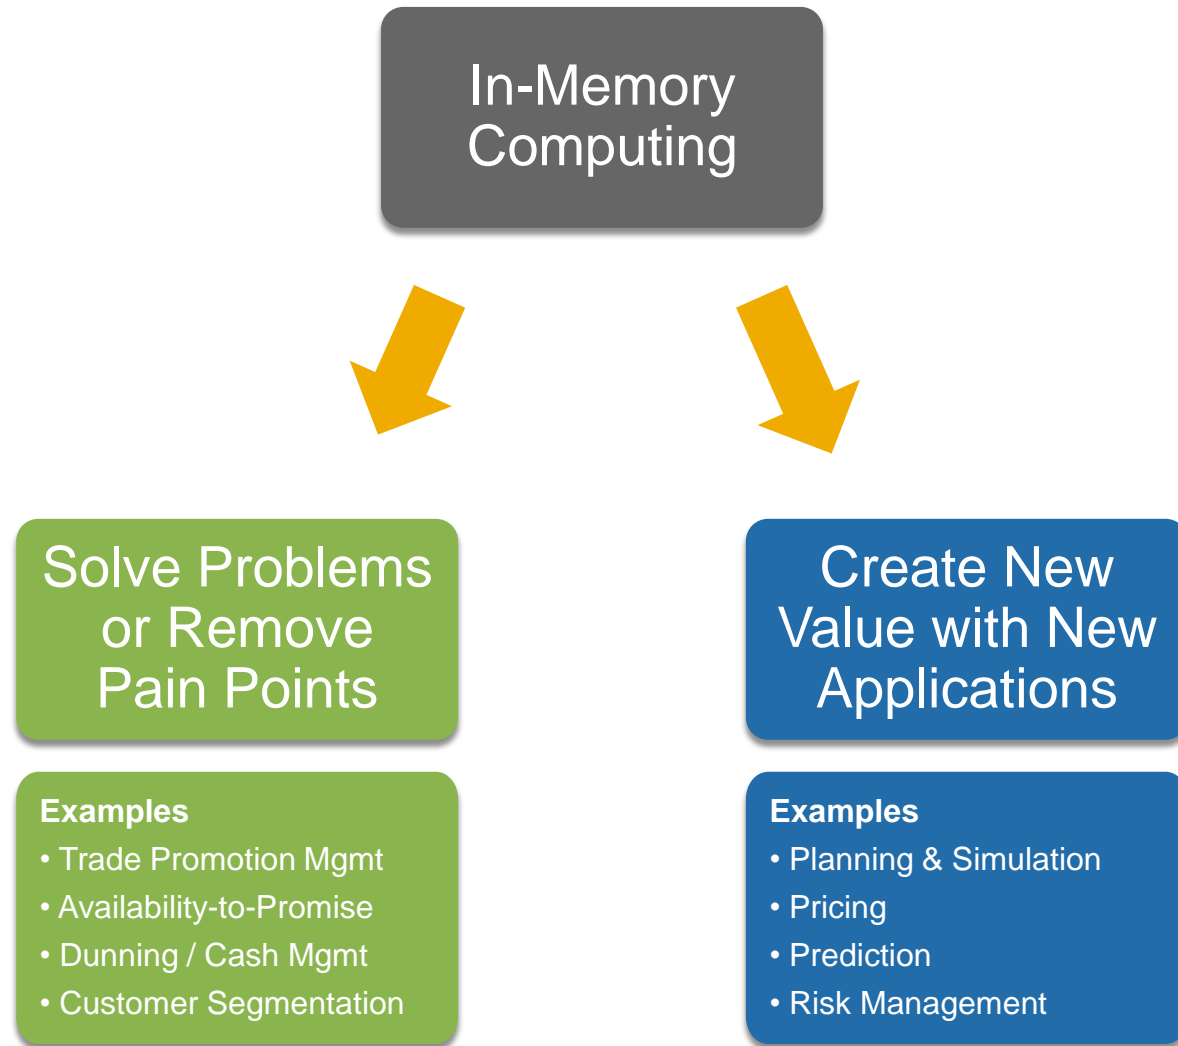
* Quote about the new book “In-Memory Data Management” by H. Plattner and A. Zeier



Applications

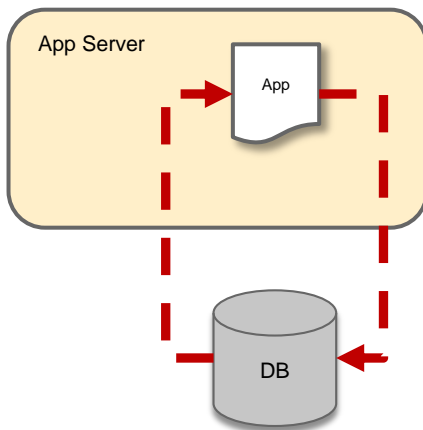
Opportunities for a New Breed of Business Software

New and Optimized Applications



Next Generation Business Applications

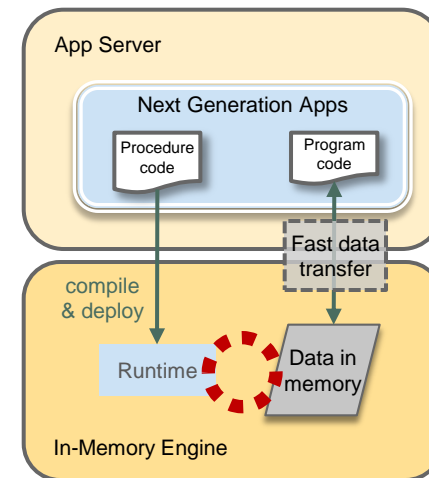
Today



Tight Coupling

- With large data volumes, reading information becomes a bottleneck
- Next generation applications will delegate data intense operations
- The runtime environment executes complex processes in memory
- In memory computing returns results by pointing apps to a location in shared memory

Tomorrow



Further Reading


“...This change in the way data is stored is having, and will continue to have a significant impact on enterprise applications and ultimately on the way businesses are run. Having real-time information available at the speed of thought provides decision makers in an organization with insights that have, until now, not existed. ”

Prof. Dr. h.c. Hasso Plattner

Co-Founder of SAP AG

Chairman of the Supervisory Board of SAP AG

Founder of the Hasso-Plattner-Institute, Potsdam



to appear soon in Springer Verlag



Thank You!

Contact information:

Dr. Stefan Sigg
SAP AG
SVP In-Memory Platform
stefan.sigg@sap.com