SAP Innovation Center Potsdam

Develop the Digital Future!

Mission

Develop the Digital Future!

- Established February 2011
- First of its kind at SAP

Mission:

Catalyze and accelerate innovation throughout SAP by bringing together cutting edge research and practical experience to create solutions for real life problems.

Berlin / Brandenburg

A hub for technological innovation

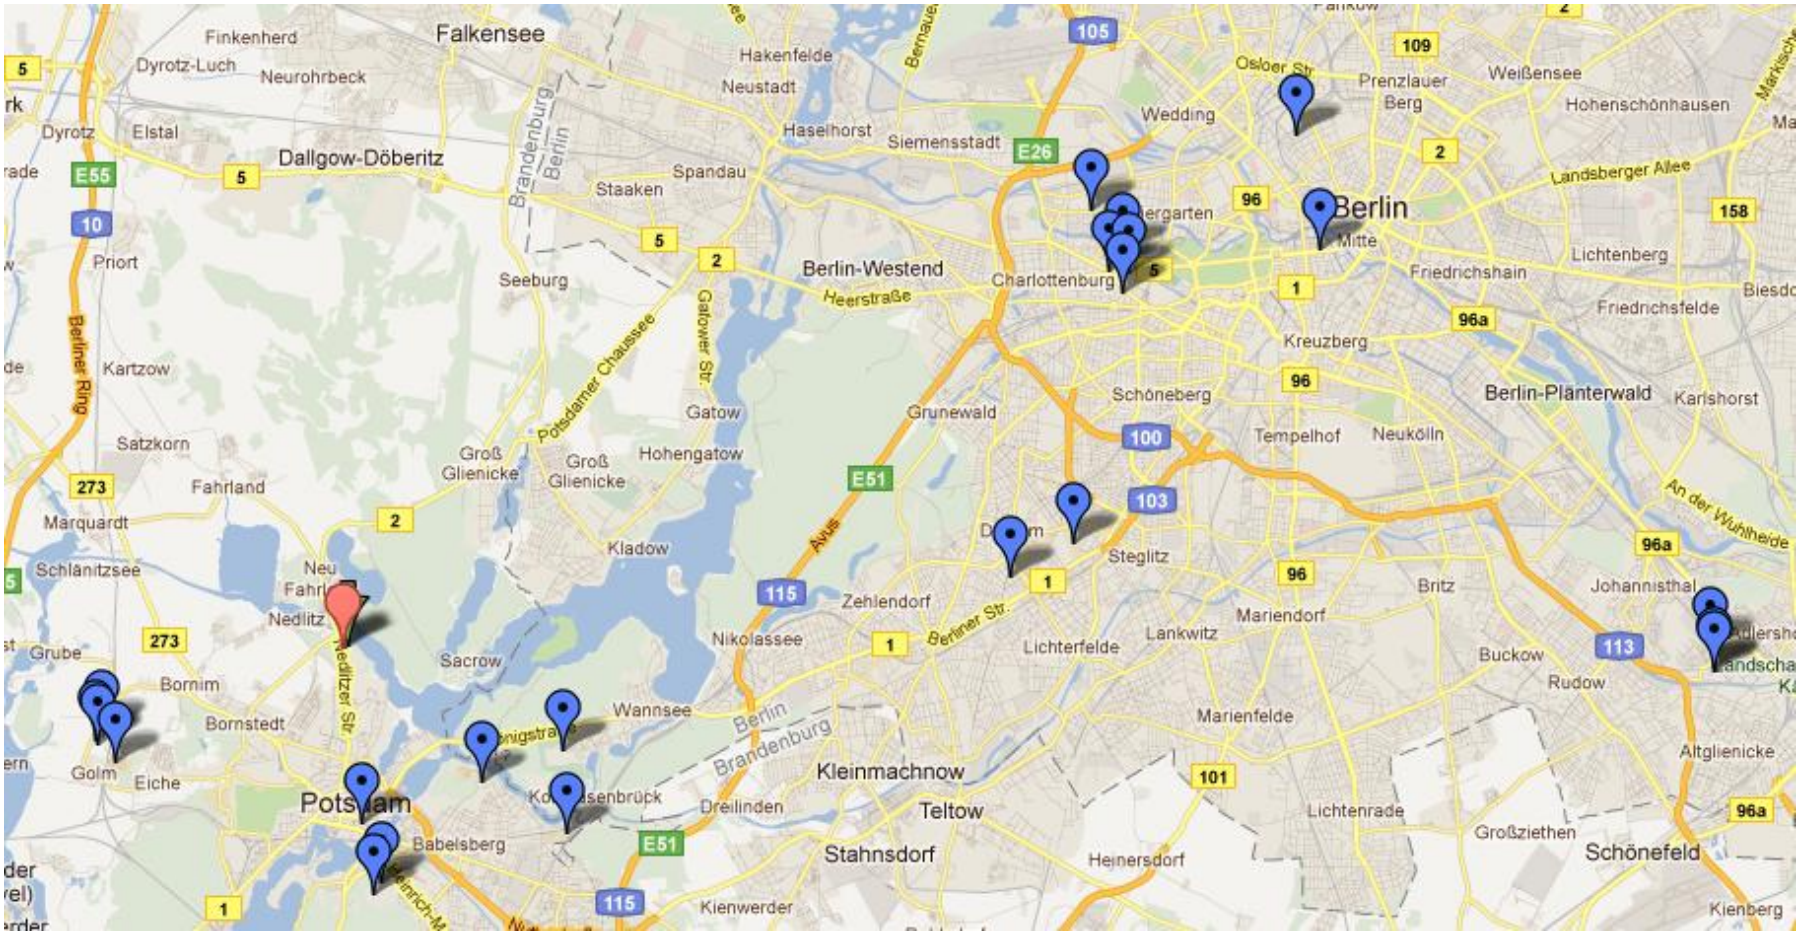

Environment Matters

ICP @ Jungfernsee

Involving students to conceptualize innovative office space

Target Size

- 100 Full-time Employees
- Up to 200 Student Workplaces

Innovation Network

Working together in Berlin/Brandenburg and beyond

Partner

SAP Innovation Center Potsdam

SAP Development

Research & Academia

Customer

Key Technology Topics

Mobility

- Mobile Apps
- User Experiences & UI

Cloud Computing

- Dependability
- Multi-Tenancy

In-Memory Data Management

- Parallelization & Scalability
- Programming Models
- Combined Structured & Unstructured Data Analysis

Innovation Cycle

Continuous Focus on Customer Value Creation

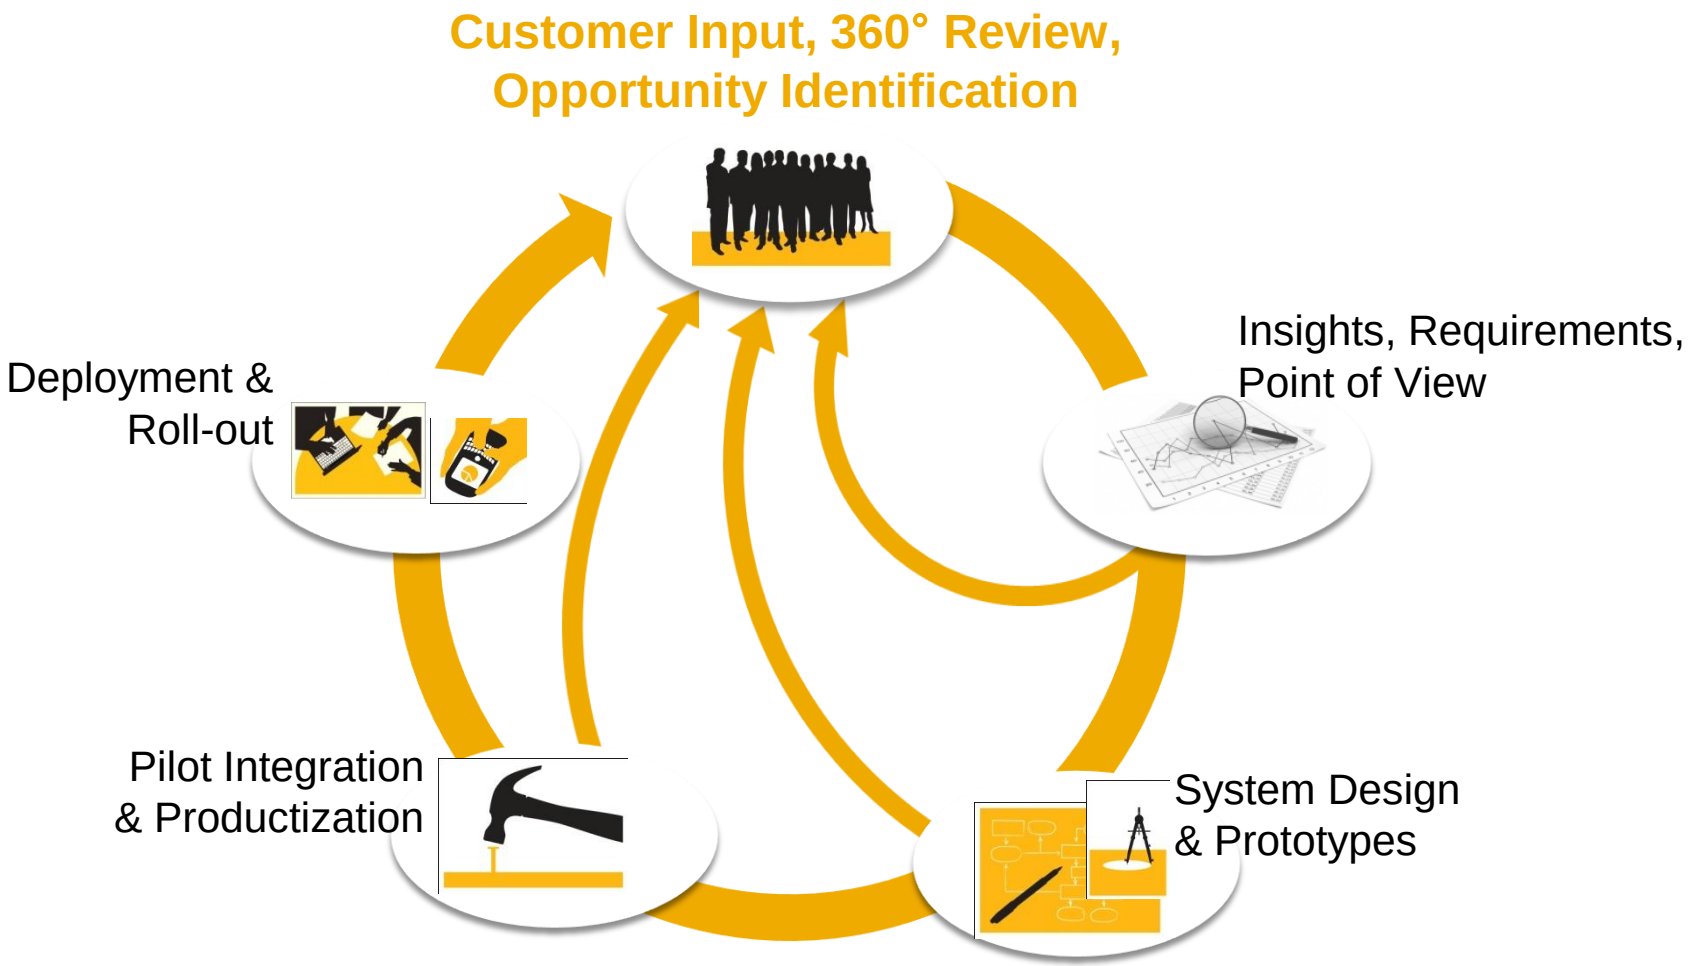

HANA Performance in a Virtual Machine Environment

HANA: High-Performance Analytic Appliance
SAP's in-memory database product

Motivation

- Businesses rely increasingly on virtual IT infrastructures (private cloud)
- HANA will become core SAP component in the future
- Tuning HANA for virtualized environments is essential

Project goals

- What is the performance loss when running HANA in a virtual machine?
- Where exactly does the performance get lost?
- Optimize HANA performance for a virtualized environment

Main Challenge

- Database scans access very large datasets (exceeding Cache and TLB capacities)

Experimental Setup

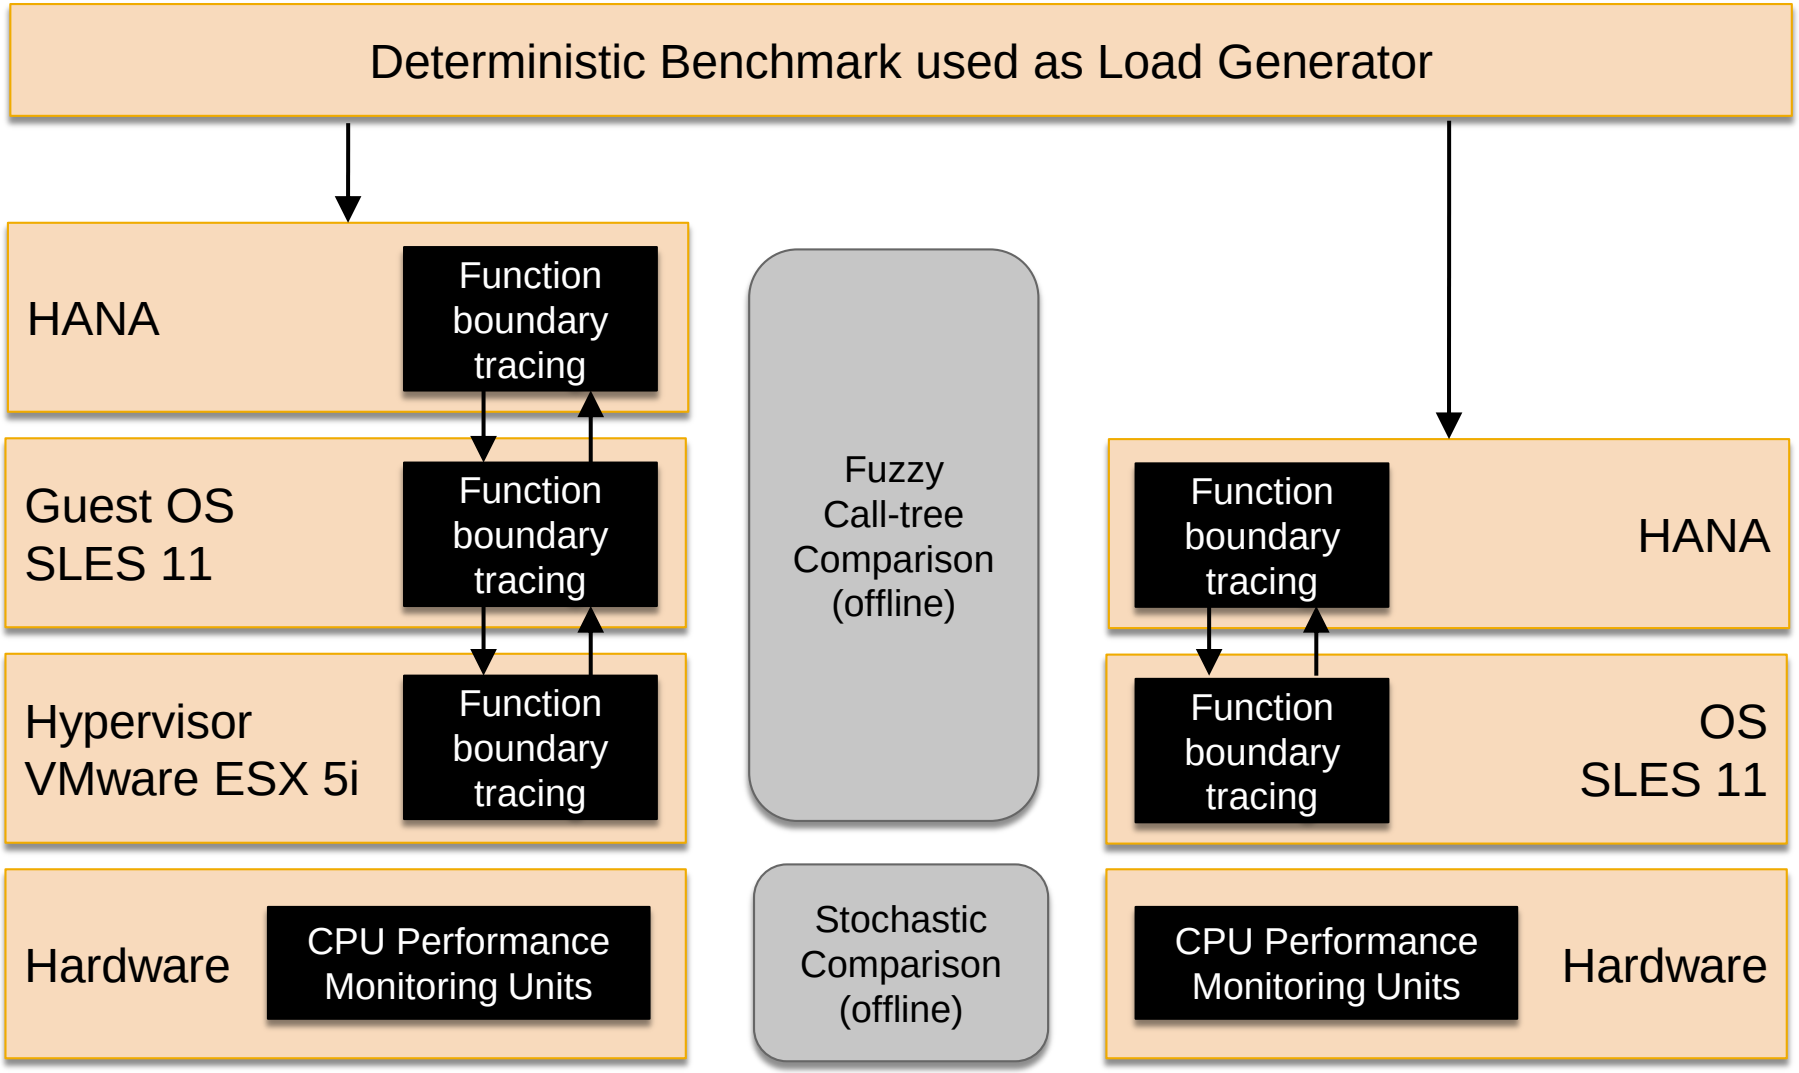

Hardware Architecture Aware Scale-Up

Introduction

- Especially analytic queries require processing of a very large amount of data
- Memory access is non-uniform in multi-processor multi-core shared memory systems
- Data layout and OS scheduling have significant impact on performance

Project goal

- Investigate data layout and scheduling policies to optimize query performance

Some numbers

- 64 billion facts (128GB data)
- 165 billion records per second processed
- Almost linear scale (tested up to 240 cores)

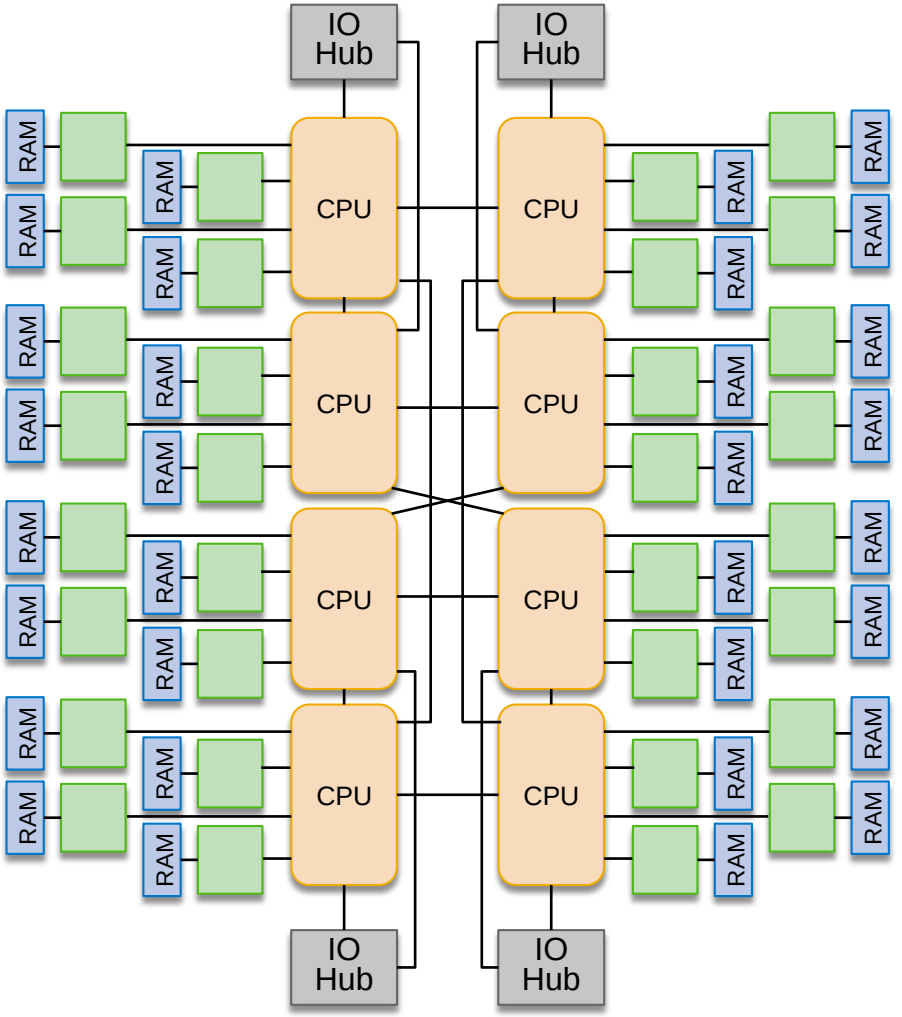

Other Operating System Related Topics

Hardware-specific database joins

Programming models for parallel computing

High-availability

Online diagnosis

Multi-tenancy and tenant migration

Load balancing

Cloud operations

Thank You!

Contact information:

Cafer Tosun
cafer.tosun@sap.com

Felix Salfner
felix.salfner@sap.com