SAP HANA – Real Time Computing

Chris Hallenbeck, HANA Solution Management
Christof Bornhoevd, HANA Architecture
Richard Pledereder, HANA Product Engineering

May 22nd, 2013

SAP Software Portfolio

Applications

Analytics

Mobile

Database & Tech.

Cloud

Powered by **SAP HANA**
a completely re-imagined, modern platform for real-time business

Outline

- Overview of HANA
- Leveraging New Generation Commodity Hardware
- HANA Architecture
- Summary

Modern Enterprise Business Applications

The Need for Efficient and Flexible Data Management

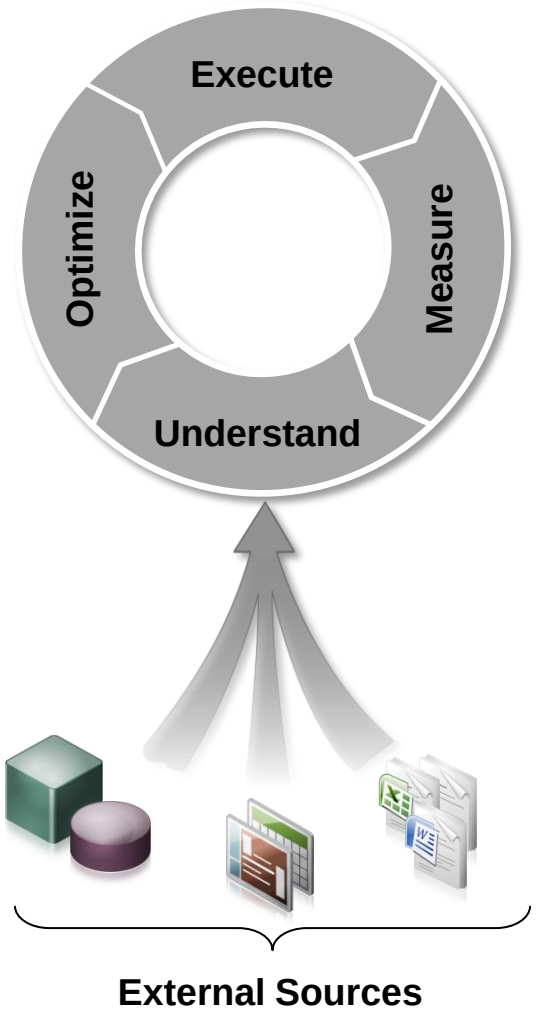

- *Combine different information access approaches: search, analysis, and exploration*
- *No clear separation between transactional and analytical parts of the application*
- *Leverage data of different degrees of structure and quality, from well-structured to irregularly structured to unstructured text data*
- *Flexibly combine internal and external data based on business decisions to be made not the set of available integrated data*
- *Are based on "real-time" current data and historical data*
- *Need to support different form factors and deployment models: on-premise, on-demand and on-device*

SAP HANA: The Challenge

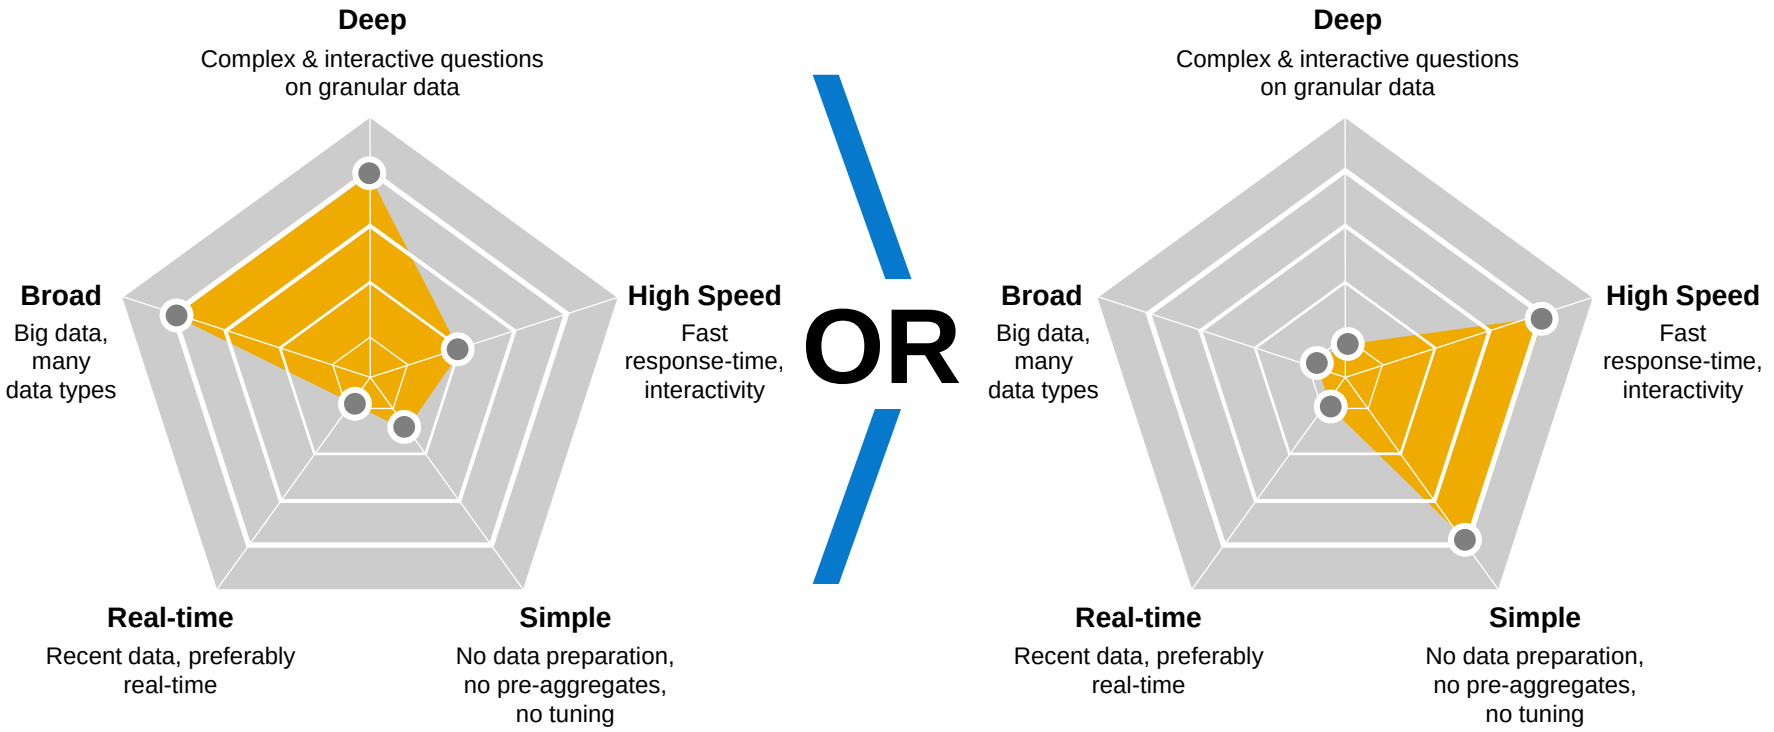

SAP HANA: The Challenge

Unify Transaction Processing and Analytics

Single System

Same Data Instance

Run Analytics in Real-Time

Run Analytics and Transactions at the “speed of thought”

SAP HANA: Single System for diverse Business Needs

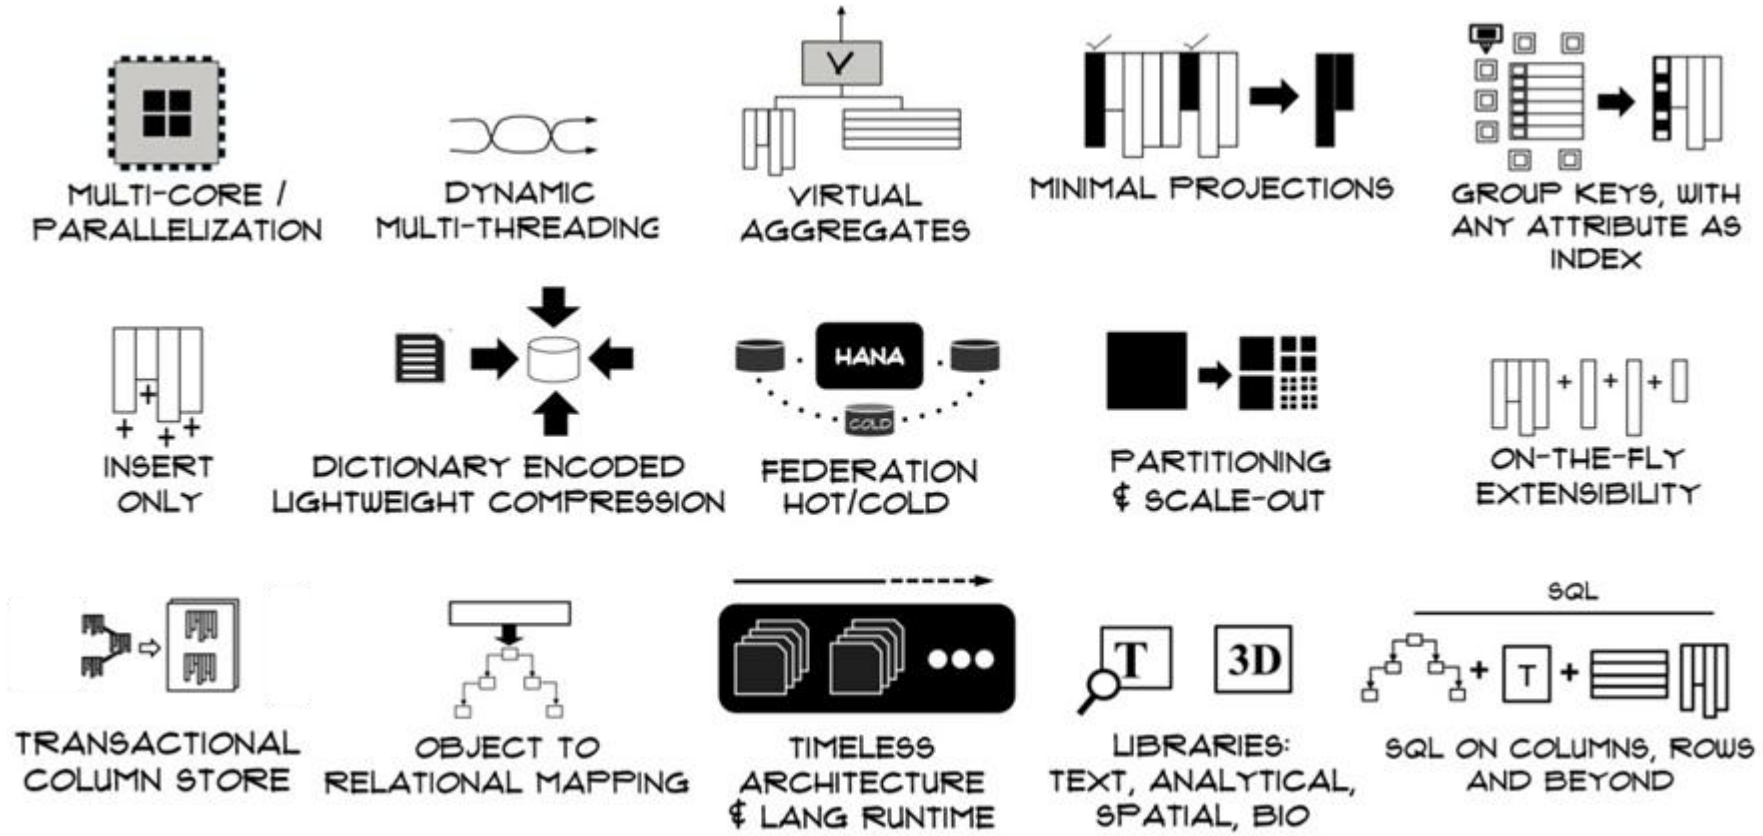

SAP HANA: A New In-Memory Data Platform

One Foundation

for

OLTP + OLAP | Structured + Unstructured Data

Legacy + New Applications

Distribution | Single Lifecycle Management

Outline

Overview of HANA

Leveraging New Generation Commodity Hardware

HANA Architecture Overview

Summary

Hardware Advances: Moore's Law - DRAM Pricing

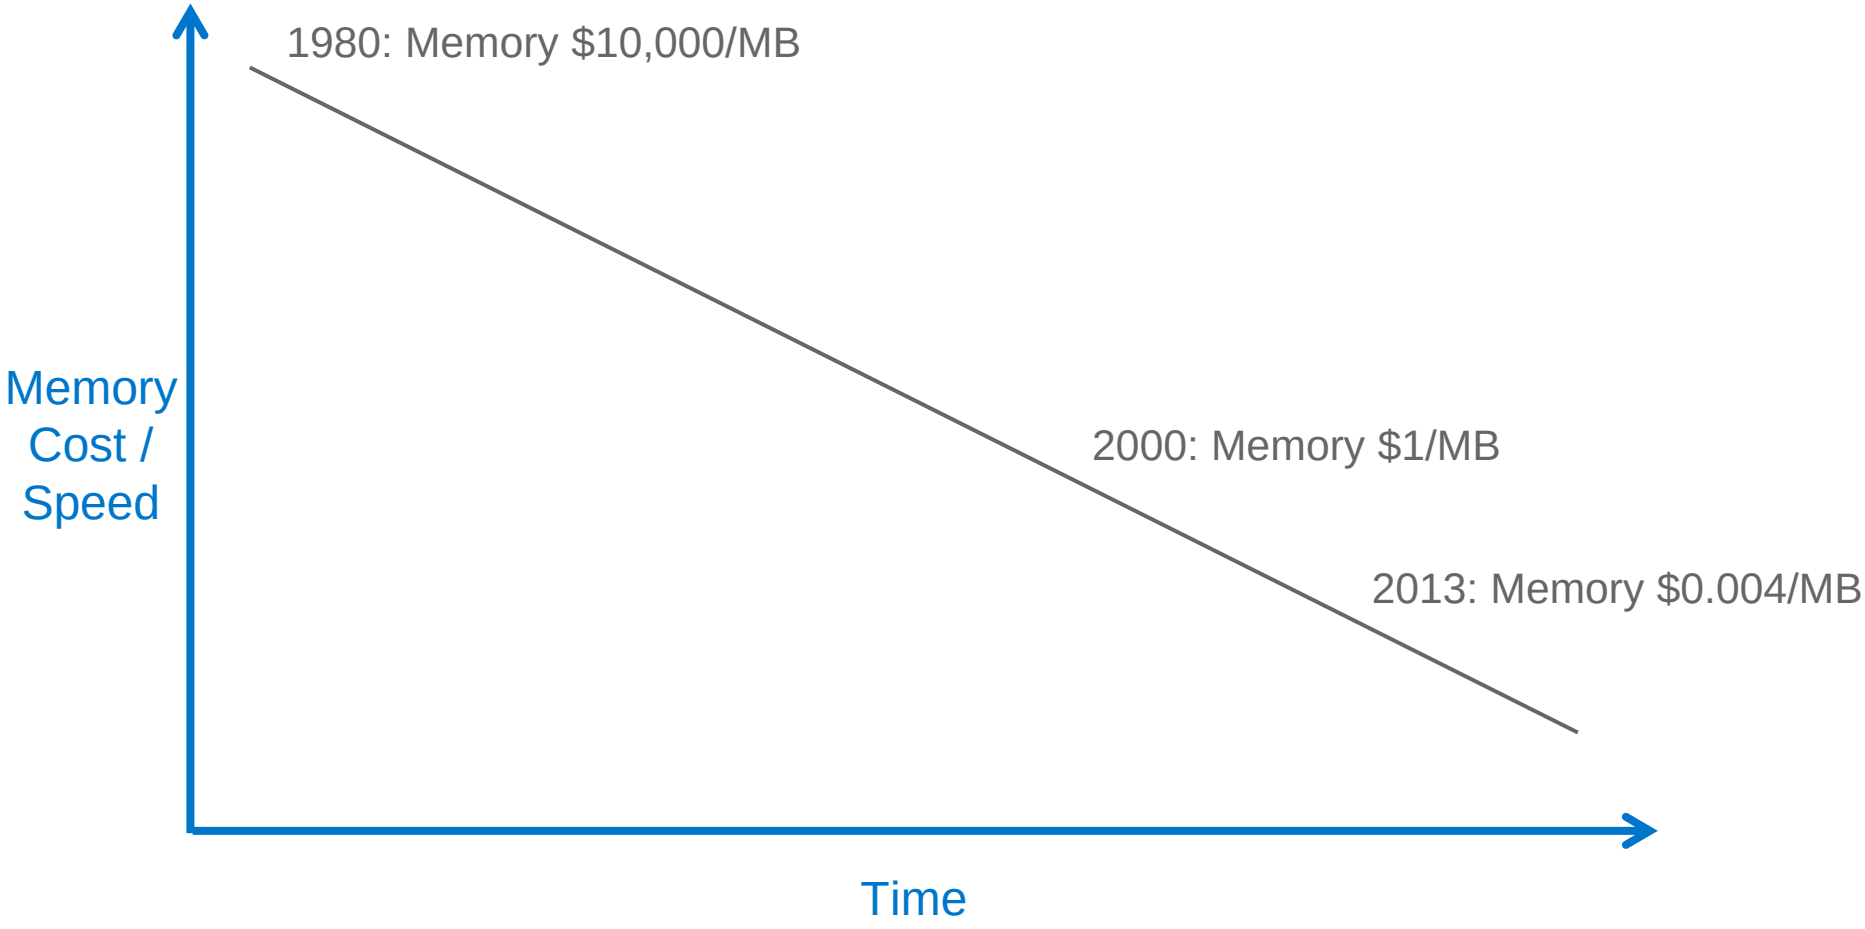

Hardware Advances: Moore‘s Law - CPUs

2002

1 core
32 bits
4MB

2007

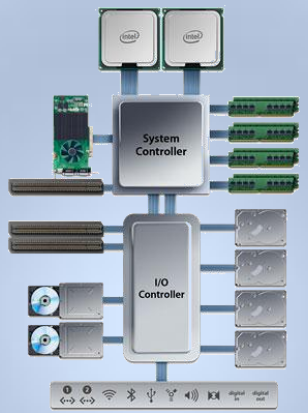

2 cores
2 CPUs per server
External Controllers

2010

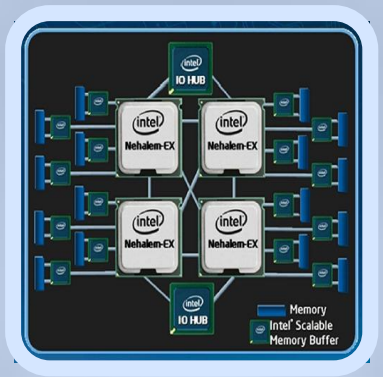

8 cores -16 threads / CPU
4 CPUs per server
On-chip memory control
Quick interconnect
VM and vector support
64 bits; 256 GB - 1 TB

2013

More cores, bigger caches
16 ... 64 CPUs per server
Greater on-chip integration
(PCIe, network, ...)
Data-direct I/O
Tens of TBs

Images: Intel, Danilo Rizzuti, FreeDigitalPhotos.net

Software Advances: Build for In-Memory Computing

Reduce Memory Access Stalls

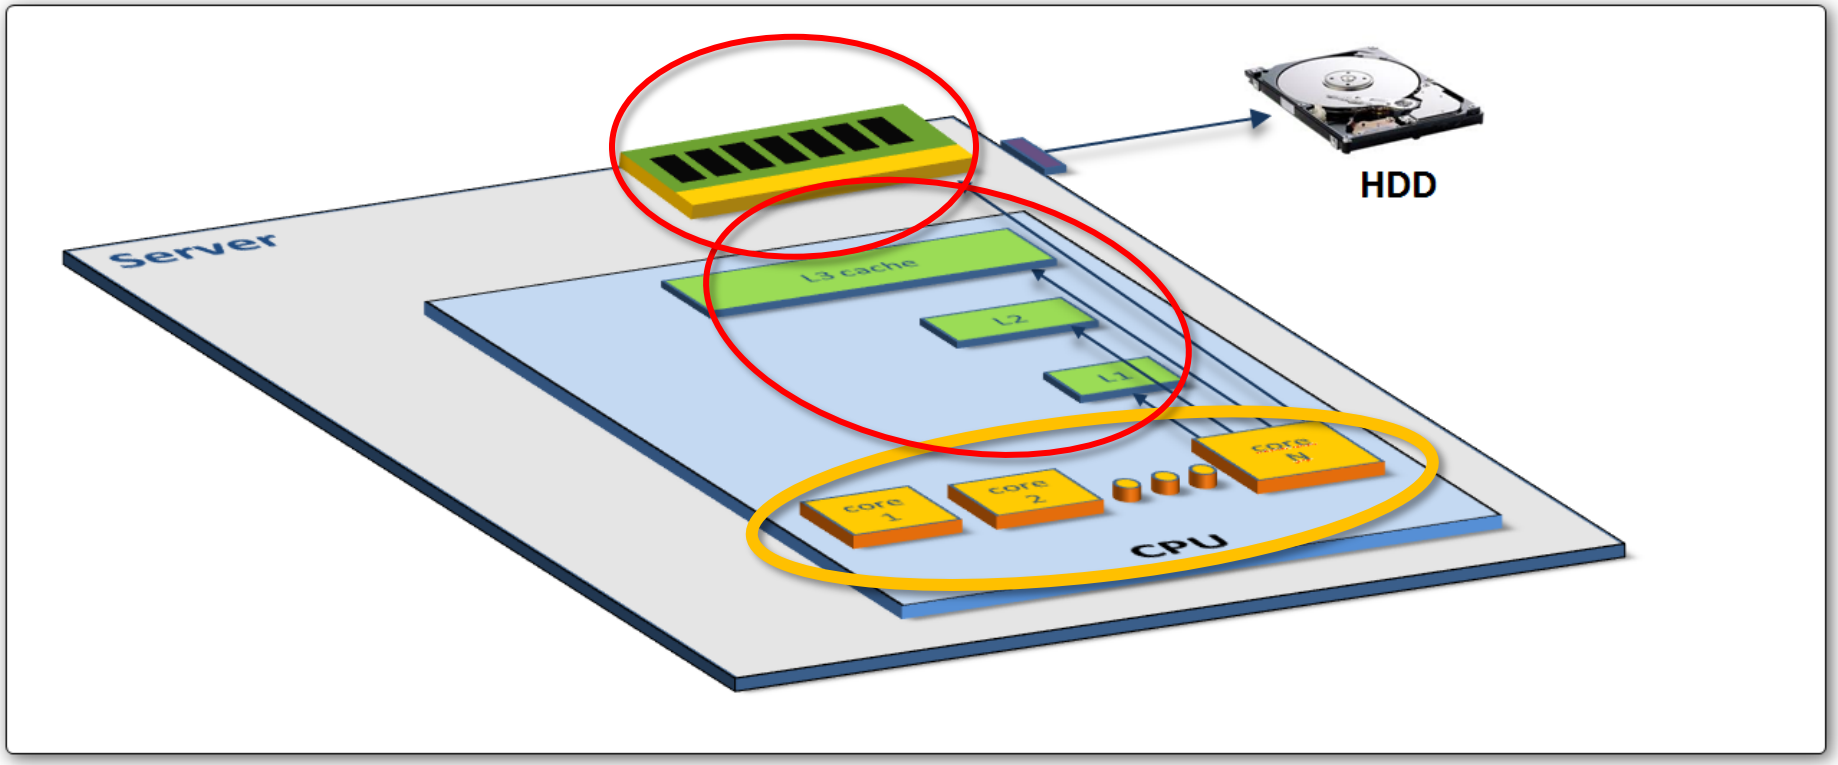

- **In-Memory Computing:** It is all data-structures (not just tables)
- **Parallelism:** Take advantage of tens, hundreds of cores
- **Data Locality:** On-chip cache awareness

In-Memory Computing

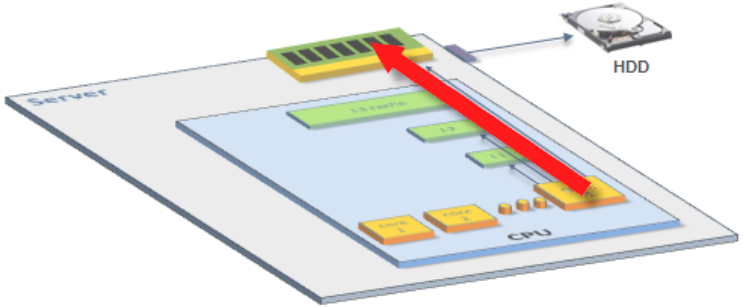

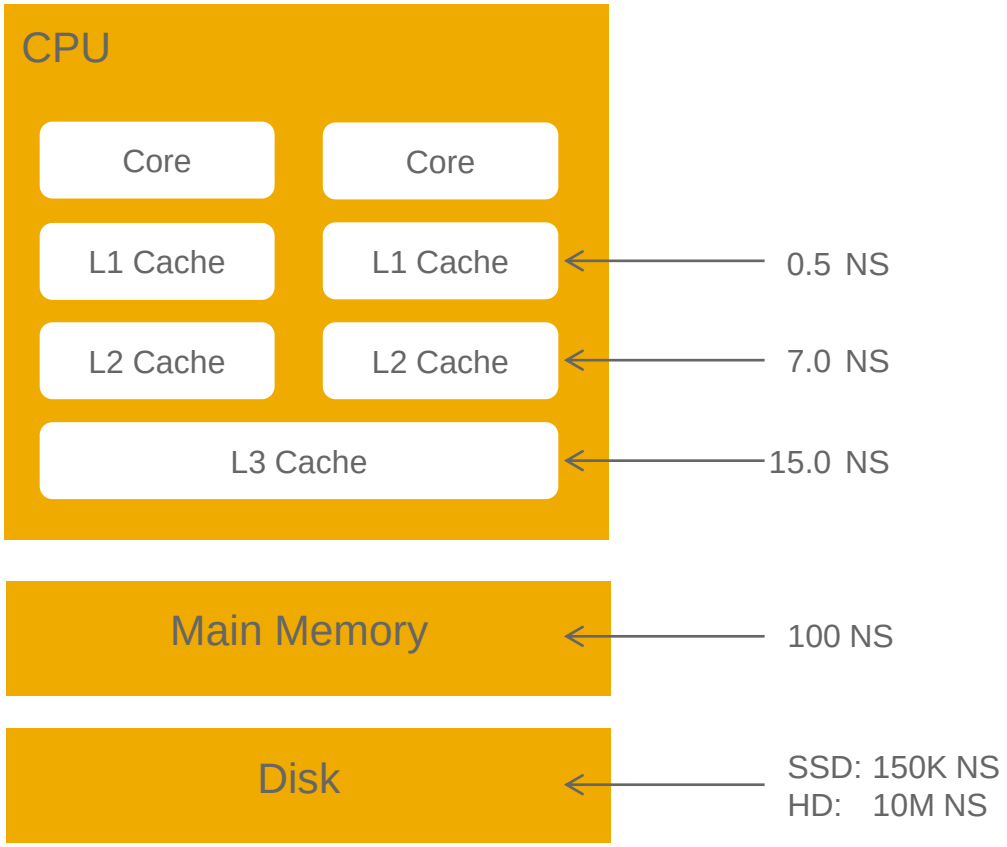

Yes, DRAM is 100,000 times faster than disk, but DRAM access is still 6-200 times slower than on-chip caches

In-Memory Computing – Data Structures

Order	Country	Product	Sales
456	France	corn	1000
457	Italy	wheat	900
458	Italy	corn	600
459	Spain	rice	800

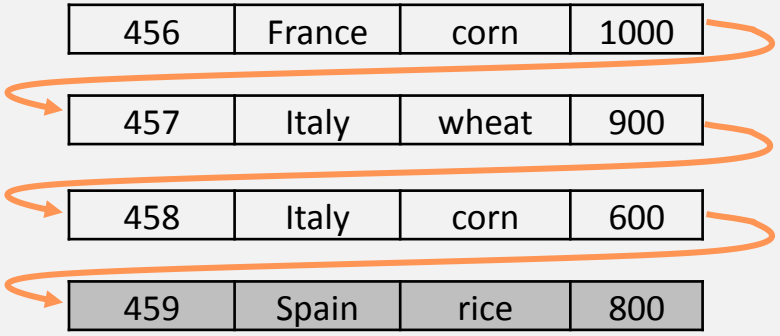

Typical Database

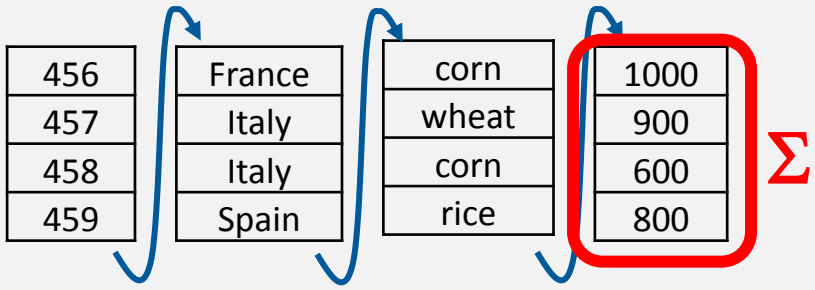

SAP HANA: column order

```
SELECT  Country, SUM(sales) FROM SalesOrders
WHERE Product = 'corn'
GROUP  BY  Country
```

SAP HANA: Dictionary Compression

Column „Name“ (uncompressed)

Miller
Jones
Millman
Zsuwalski
Baker
Miller
John
Miller
Johnson
Jones

⋰

Column „Name“ (dictionary compressed)

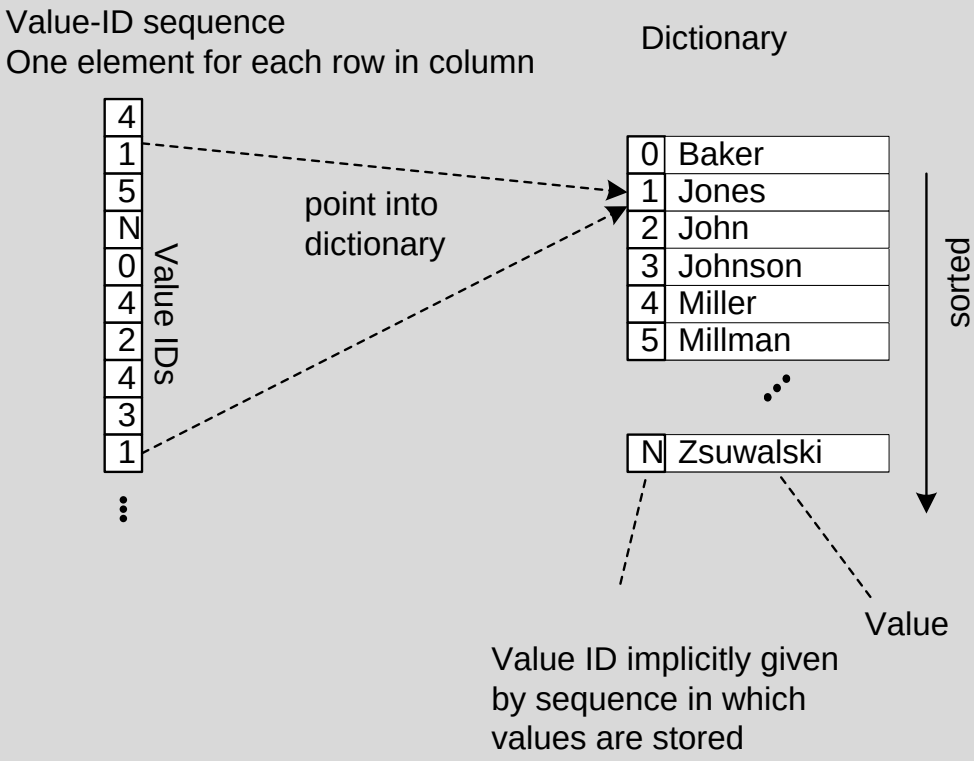

SAP HANA: Multi-Core Parallelization

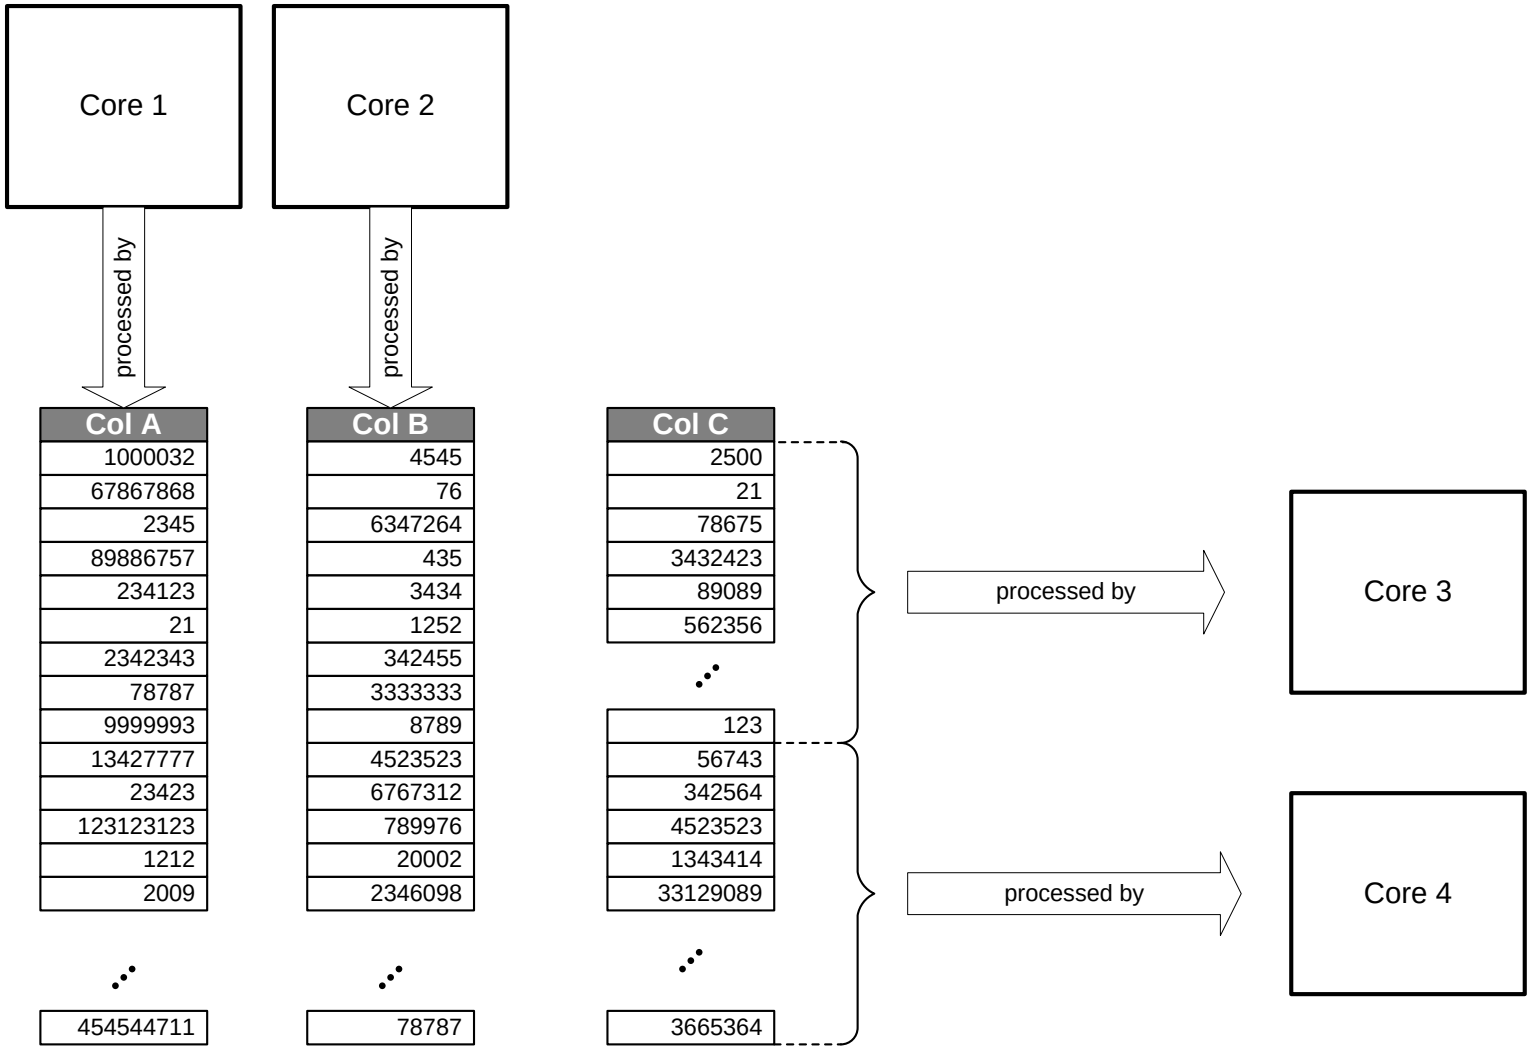

Single Instruction Multiple Data (SIMD)

- **Scalar processing**
  - traditional mode
  - one instruction produces one result

- **SIMD processing**
  - with Intel® SSE(2,3,4)
  - one instruction produces multiple results

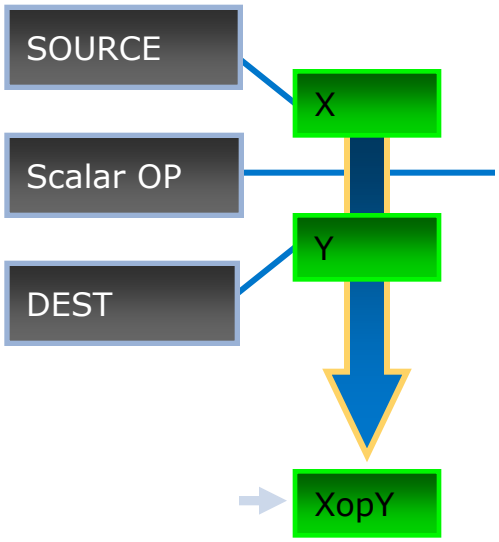

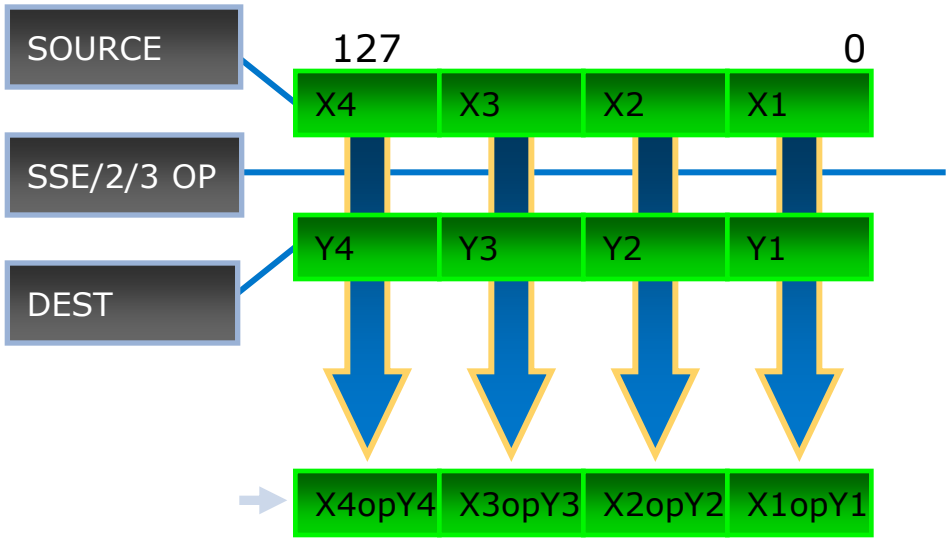

Single Instruction Multiple Data (SIMD)

128-bit wide with Intel® SSE(2,3,4)

- 2 64-bit integer ops/cycle
- 4 32-bit integer ops/cycle
- 8 16-bit integer ops/cycle
- 16 8-bit integer ops/cycle

256-bit with AVX

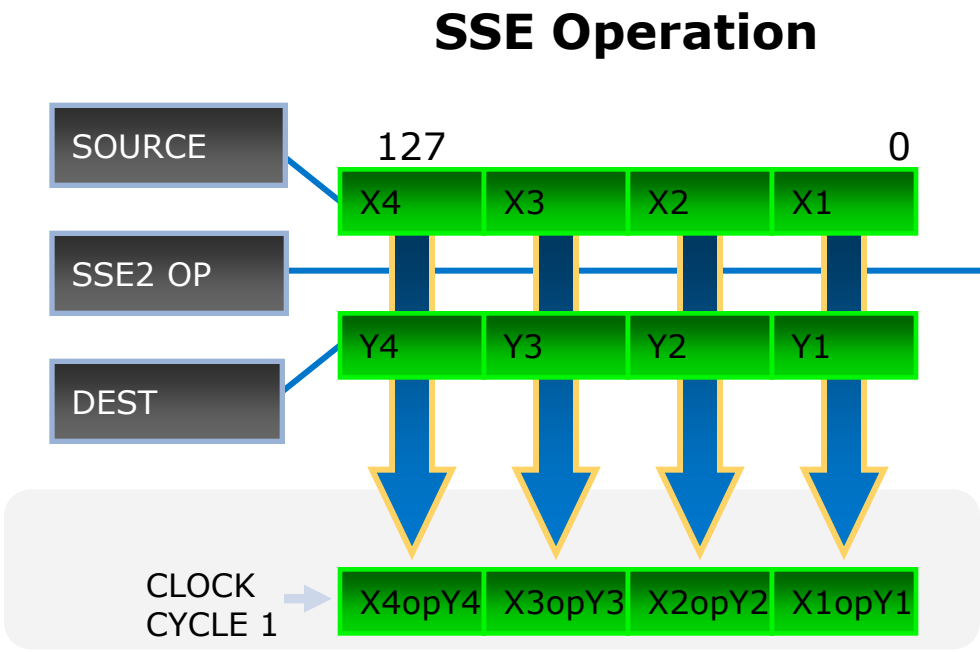

Vector-Processing Unit built-in standard processors

Parallelization at All Levels

- Multiple user sessions
- Concurrent operations within a query (… T1.A … T2.B…)
- Data partitioning on one or more hosts
- Horizontal segmentation, concurrent aggregation
- Multi-threading at Intel processor core level
- Vector Processing

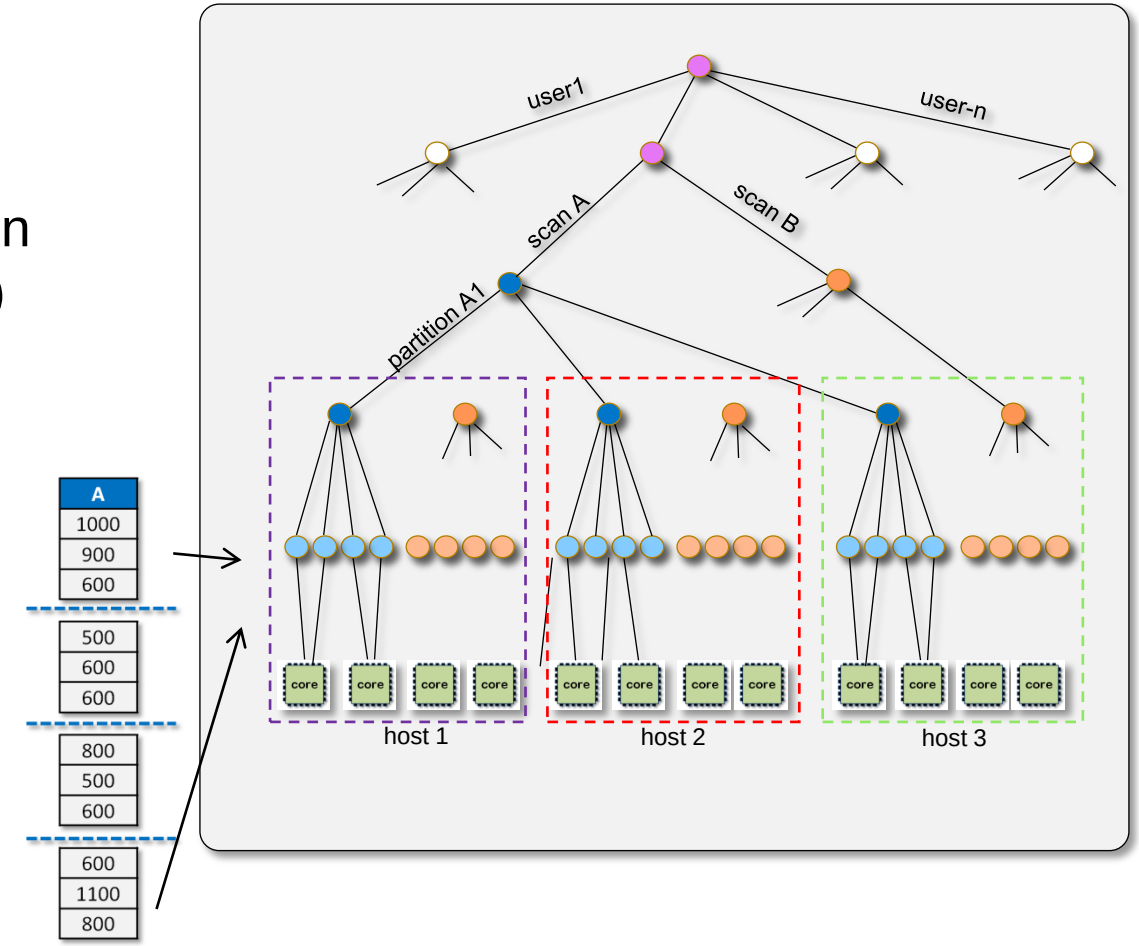

SAP HANA: True In-Memory Computing

Use Vector Based Processing (SIMD)

Leverage Data Locality

Act on Compressed Data

Cache Line (64K) Aligned Data

Hyper-Threading

3.2B Integer Scans / Second / Core

12.5M Aggregates / Second / Core

1.5M Inserts / Second

SAP HANA deployment options

Single Server

- 2 CPU 128GB to 8 CPU 1TB
- Single HANA deployments for data marts or accelerators

Scale Out Cluster

- 2 to n servers per cluster
- Each server is either 4 CPU/ 512GB or 8 CPU/ 1TB
- Largest certified configuration: 56 servers
- Largest tested configuration: 250 servers
- Support for high availability and disaster tolerance

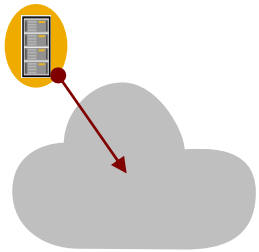

Cloud Deployment

- HANA instances can be deployed to AWS
- Free developer license
- 99 cents per hour for productive use
  (+ EUR 2.50 for the AWS machine)

Outline

Overview of HANA

Leveraging New Generation Commodity Hardware

HANA Architecture Overview

Summary

SAP HANA Technology & Features
Combined in one DBMS Platform

Common DBMS features
- SQL
- ACID: isolation (MVCC), logging and recovery
- Stored procedures

Hybrid DBMS
- Column store, row store, graph store
- In-memory and disk based
- Insert only (temporal) and updatable tables

High Performance
- Efficient compression techniques
- Massive parallelization over CPU cores and nodes
- Data structures optimized for main memory
- Data aging concept

Reduced TCO: OLTP, OLAP, search in one system

SAP HANA Database

Multi-Engine for Different Application Needs

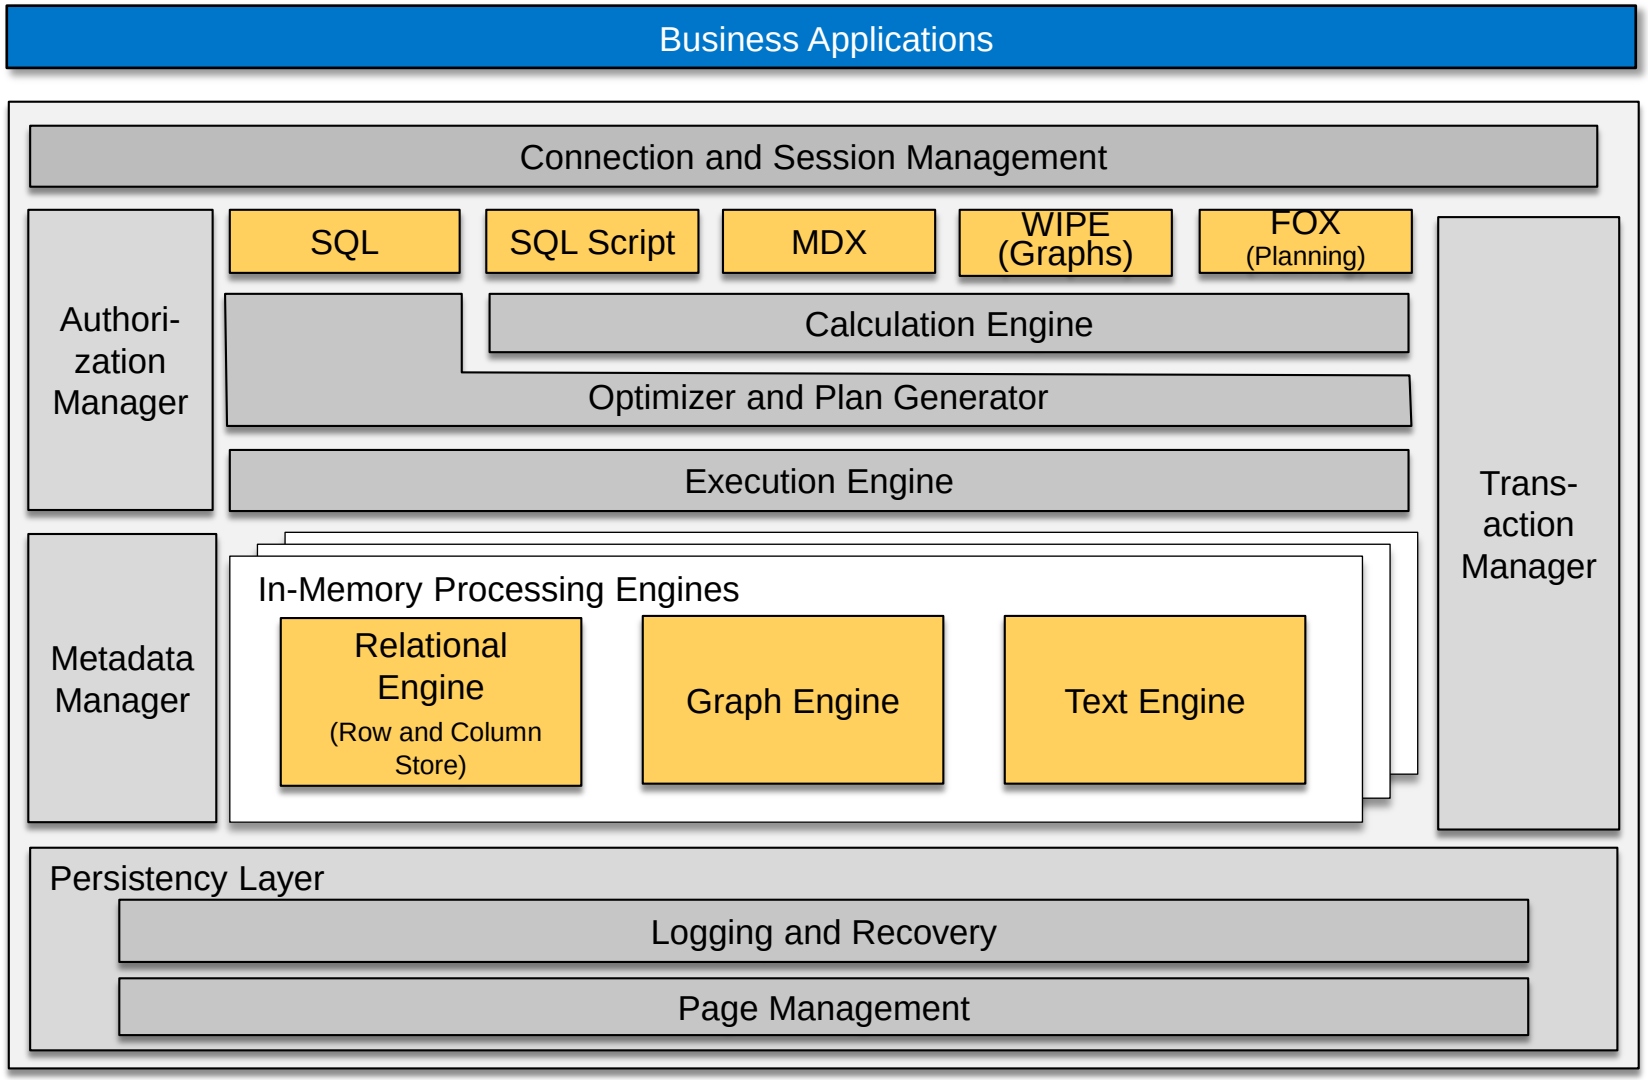

Diverse Set of Operators

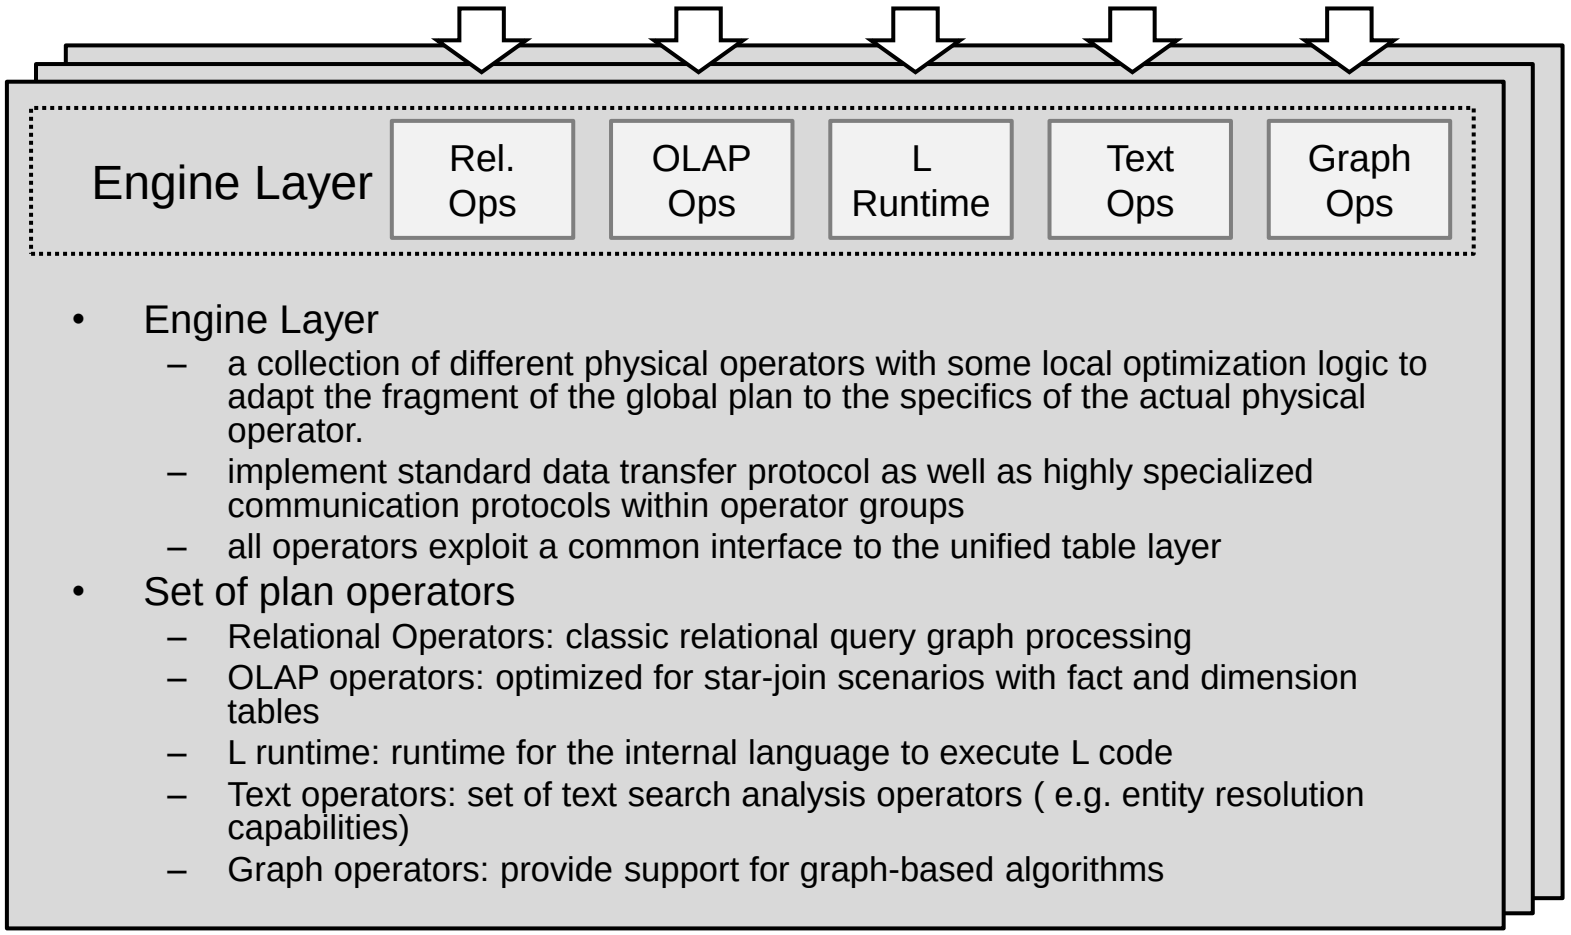

- Engine Layer
  - a collection of different physical operators with some local optimization logic to adapt the fragment of the global plan to the specifics of the actual physical operator.
  - implement standard data transfer protocol as well as highly specialized communication protocols within operator groups
  - all operators exploit a common interface to the unified table layer
- Set of plan operators
  - Relational Operators: classic relational query graph processing
  - OLAP operators: optimized for star-join scenarios with fact and dimension tables
  - L runtime: runtime for the internal language to execute L code
  - Text operators: set of text search analysis operators (e.g. entity resolution capabilities)
  - Graph operators: provide support for graph-based algorithms

SAP HANA Column Store

High data compression

- Efficient compression methods (dictionary, run length, cluster, prefix, etc.)
- Compression works well with columns and can speedup operations on columns (~ factor 10)
- Because of compression, write changes into less compressed delta storage
  - Needs to be merged into columns from time to time or when a certain size is exceeded
  - Delta merge can be done in background
  - Trade-off between compression ratio and delta merge runtime
- Updates into delta data storage and periodically merged into main data storage
  - High write performance not affected by compression
  - Data is written to delta storage with less compression which is optimized for write access. This is merged into the main area of the column store later on.

SAP HANA: Column Store

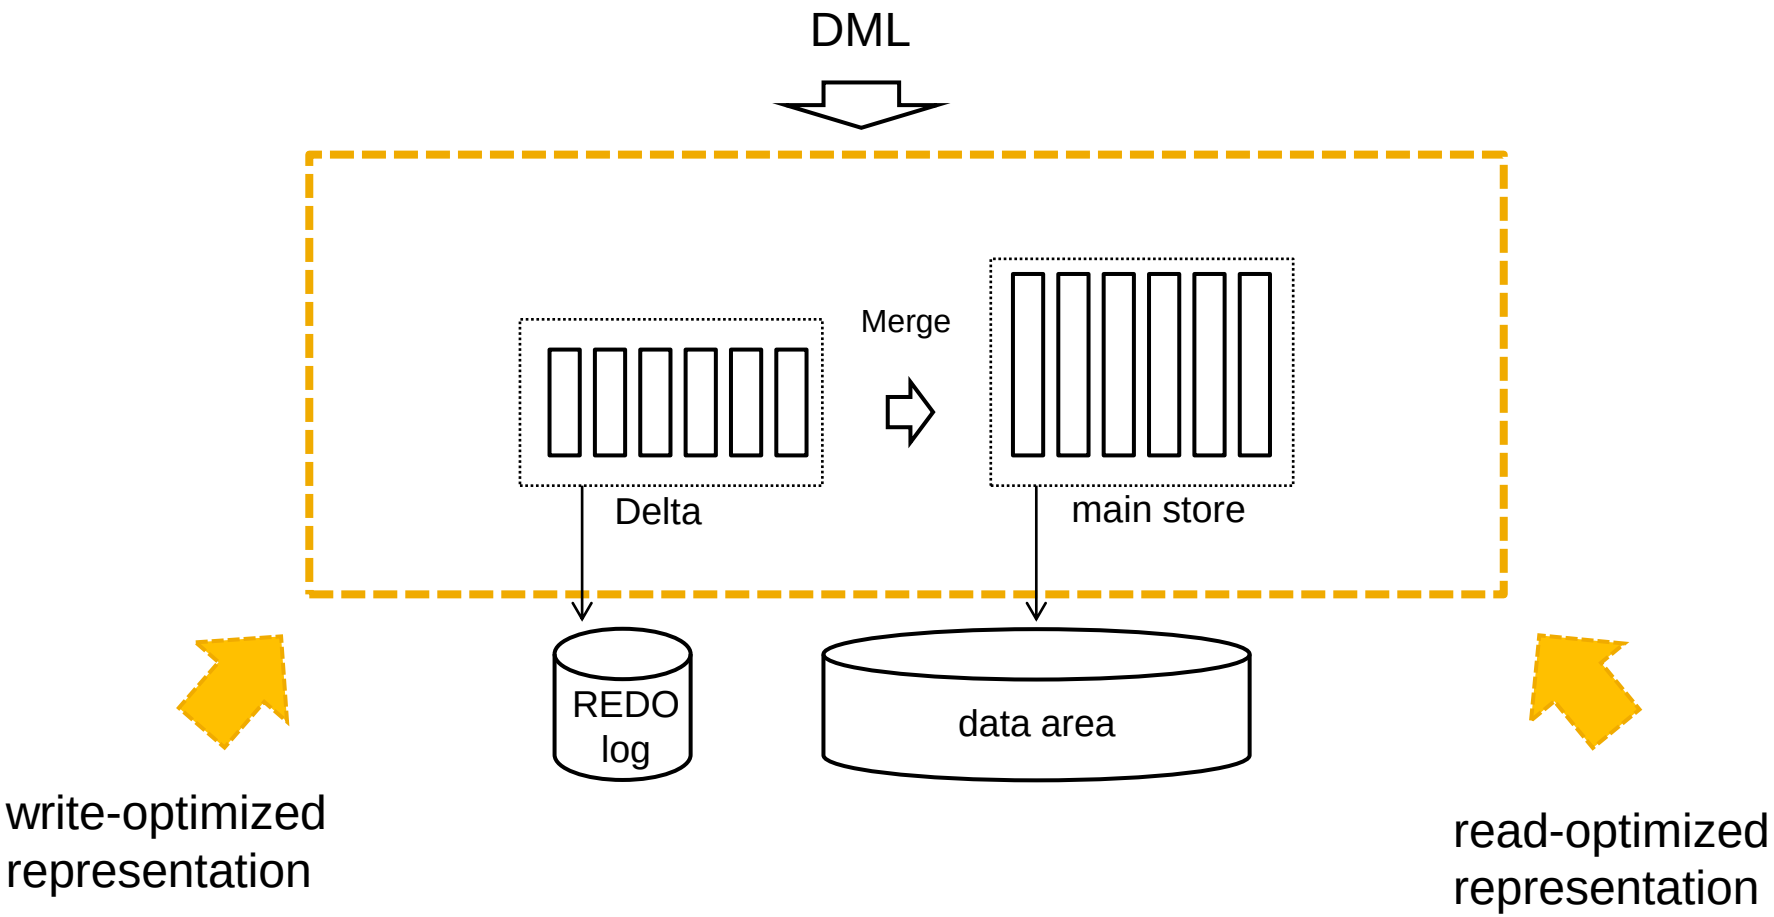

SAP HANA Database

Optimized Communication with the Application Layer

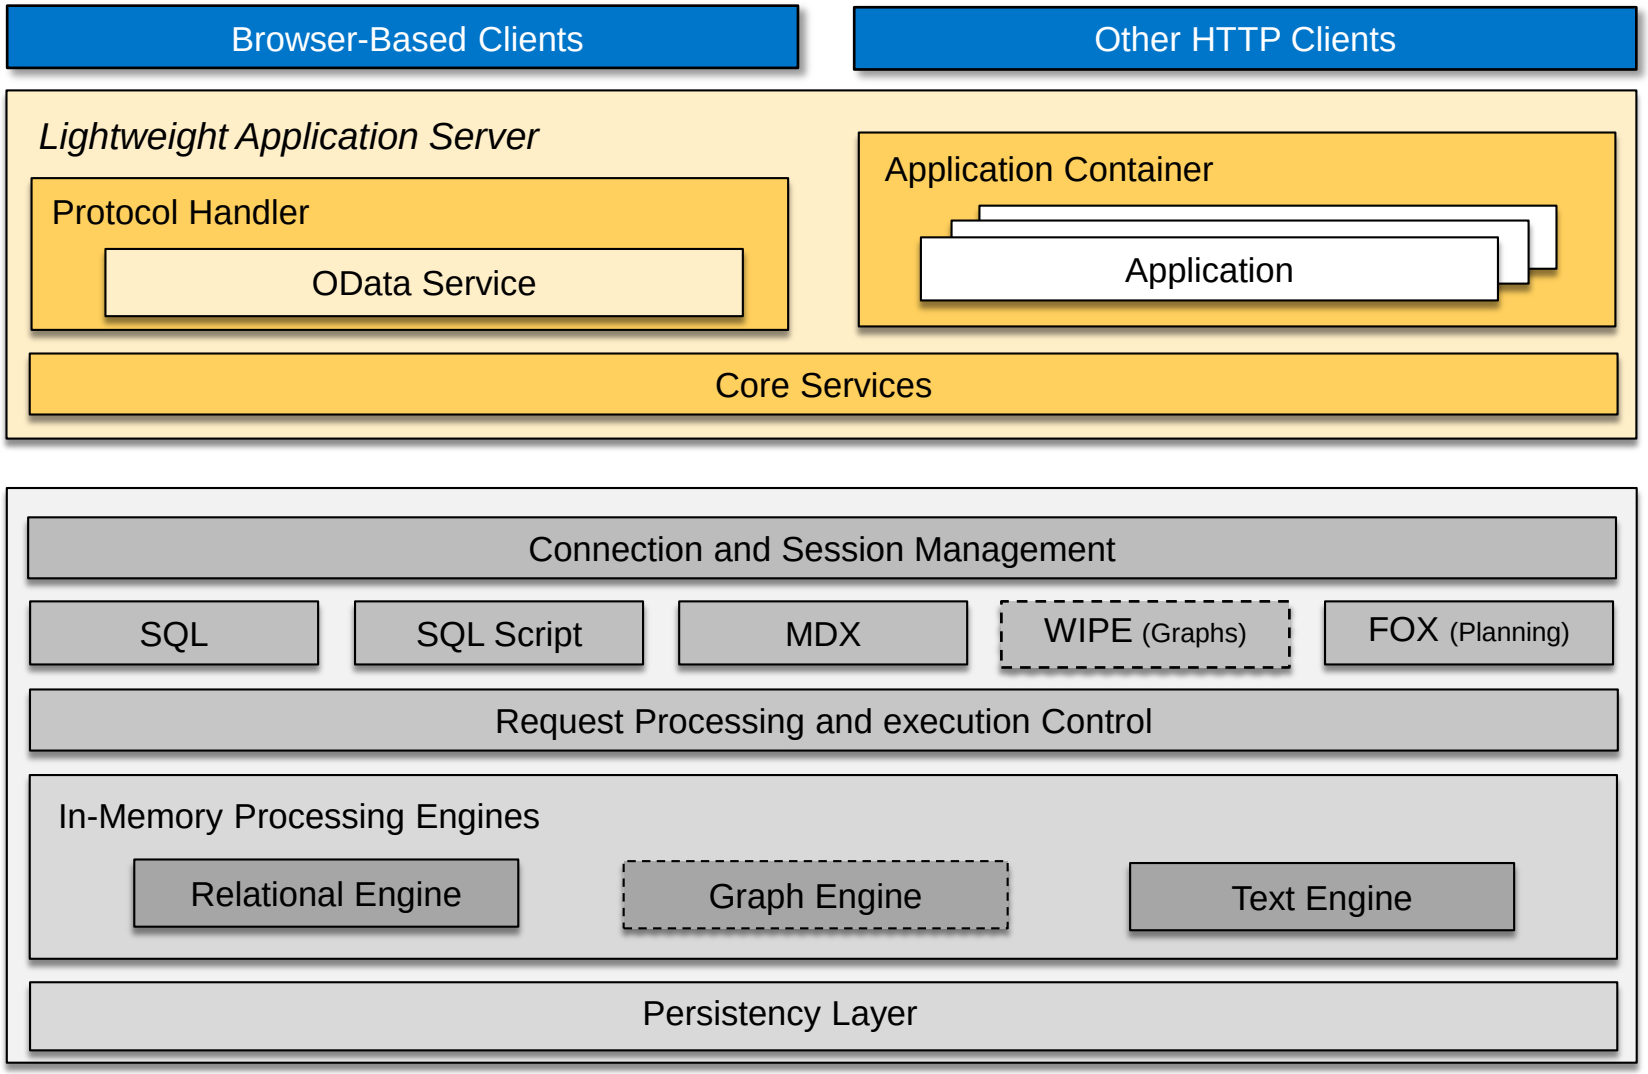

SAP HANA Database

Optimized Communication with the Application Layer

- *HANA integrates a lightweight application server component into the database cluster infrastructure to allow efficient communication and data exchange between database layer and application layer*

- *The lightweight application server component in HANA provides core application services like:*
  - *Application service runtime engine*
  - *Application lifecycle management (versioning/transport) via content repository*
  - *Programming model*
  - *standardized data exchange (OData: JSON / ATOM / XML)*
  - *Session and connection management*
  - *outbound connectivity (HTTP, SMTP) etc.)*

Outline

Overview of HANA

Leveraging New Generation Commodity Hardware

HANA Architecture Overview

Summary

A New Data Management Platform for Modern Business Applications

Much more than a Relational Database

OLTP + OLAP | Structured + Unstructured Data

Legacy + New Applications

Distribution | Single Lifecycle Management

SAP HANA – Real Time Computing

Chris Hallenbeck, HANA Solution Management
Christof Bornhoevd, HANA Architecture
Richard Pledereder, HANA Product Engineering

May 22nd, 2013