Mahindra Satyam

SAP Business Analytics and Technology

Speed of Business Change

Years It Took to Reach a Market Audience of 50 Million

Radio
38 years

TV
13 years

Internet
4 years

iPod
3 years

Facebook
2 years

"The greatest danger in times of turbulence is not the turbulence; it is to act with yesterday's logic." – Peter Drucker, 1980

Mahindra Satyam

Mahindra Satyam

In-Memory Computing

The elements of In-Memory computing are not new. However, dramatically improved hardware economics and technology innovations in software has now made it possible for SAP to deliver on its vision of the Real-Time Enterprise with In-Memory business applications

Conventional Databases
Disk Read
5 milliseconds

In-Memory Databases
Disk Read
5 nanoseconds

1 million times faster!

HW Technology Innovations

Multi-Core Architecture
(8 x 8core CPU in one server)

Massive parallel scaling with many blades

One blade ~$50.000 = 1 Enterprise Class Server

64bit address space
max at 2TB in current servers

100GB/s data throughput

Dramatic decline in price/performance

SAP SW Technology Innovations

Row and Column Store

Compression

Partitioning

No Aggregate Tables

Insert Only on Delta

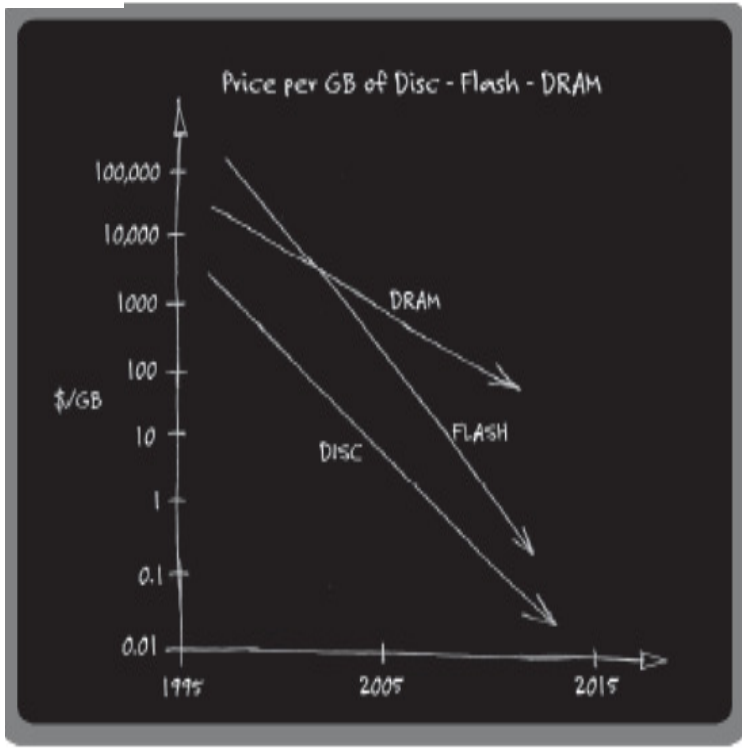

SAP -HANA – High Performance ANalytic Appliance
- In Memory Computing

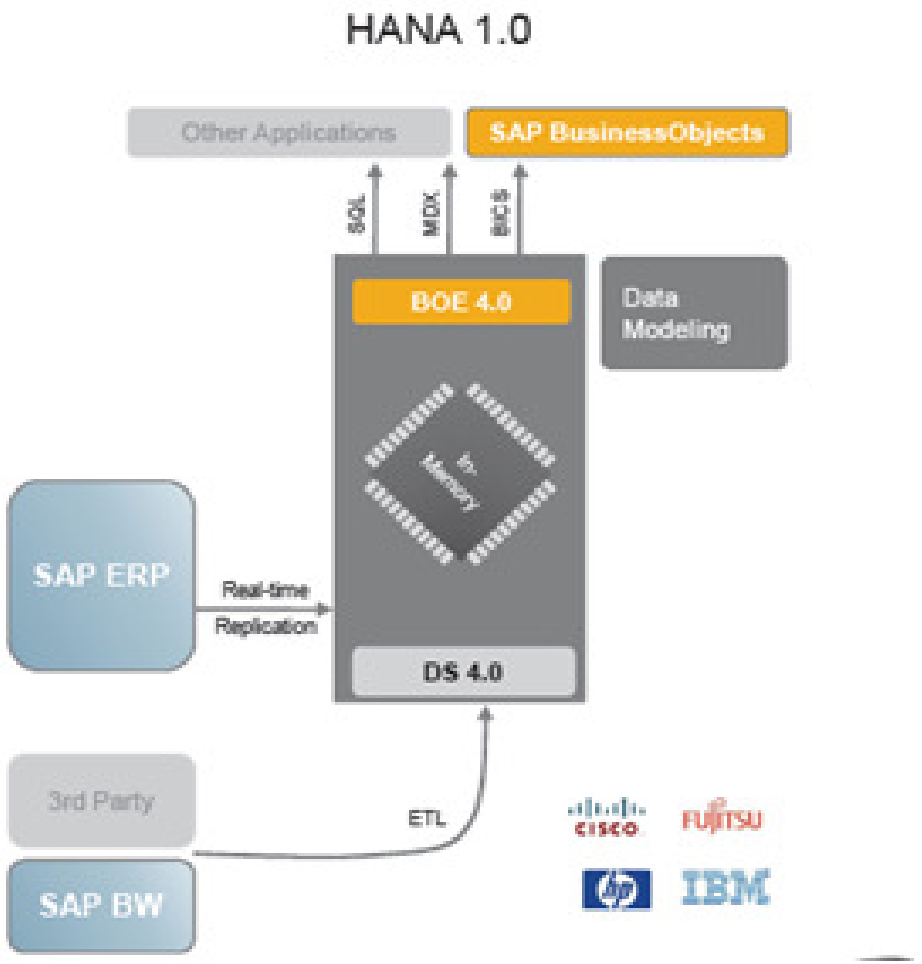

HANA Proof Points

SAP High Performance Analytic Appliance (SAP HANA)

SAP HANA is the engine of the real-time enterprise. It provides a foundation on which to build a new generation of applications, enabling customers to analyze large quantities of data from virtually any source, in real time. The example below showcases actual customer performance of a core reporting process.

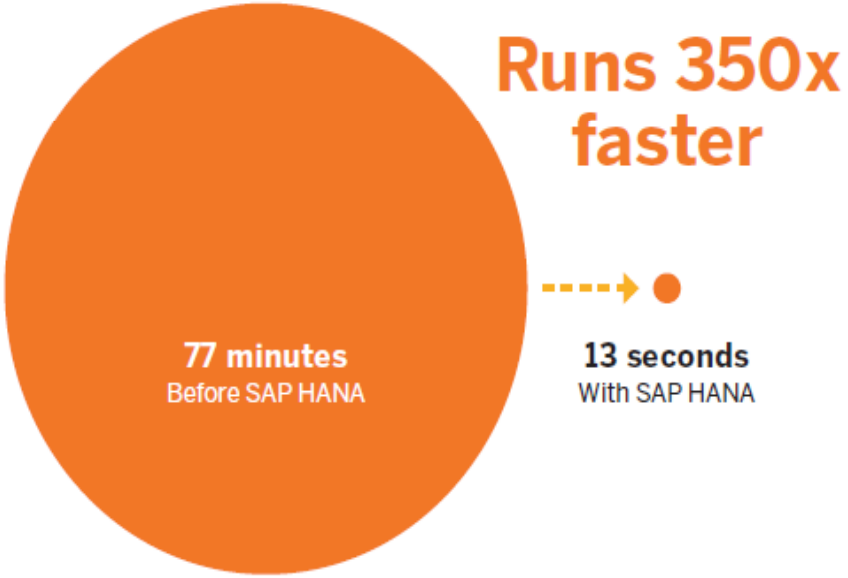

Experience the real-time Enterprise in action

A live analysis by a consumer products company reveals how SAP HANA analyzes current point-of-sale data in real time—empowering this organization to review segmentation, merchandising, inventory management, and forecasting information at the speed of thought.

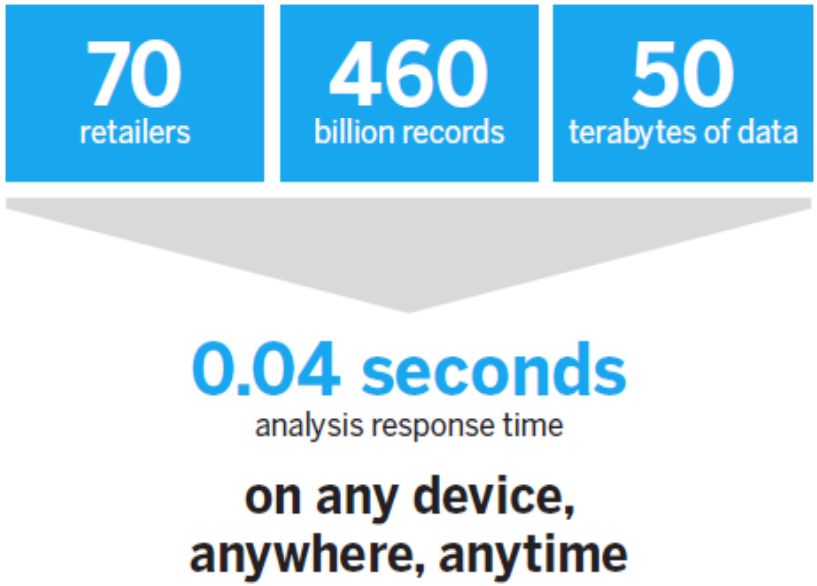

Agenda

- Components
- HANA Architecture

- Loading Data in HANA
- Modeling in HANA
- Reporting on HANA

- Request Processing and Execution
- Row Store
- Column Store
- Administration – Persistence Layer

- HANA Road Map
- Licensing Overview
- T-Shirt Sizing & Costs

Components

SAP HANA 1.0 Landscape Including External SAP BusinessObjects Servers

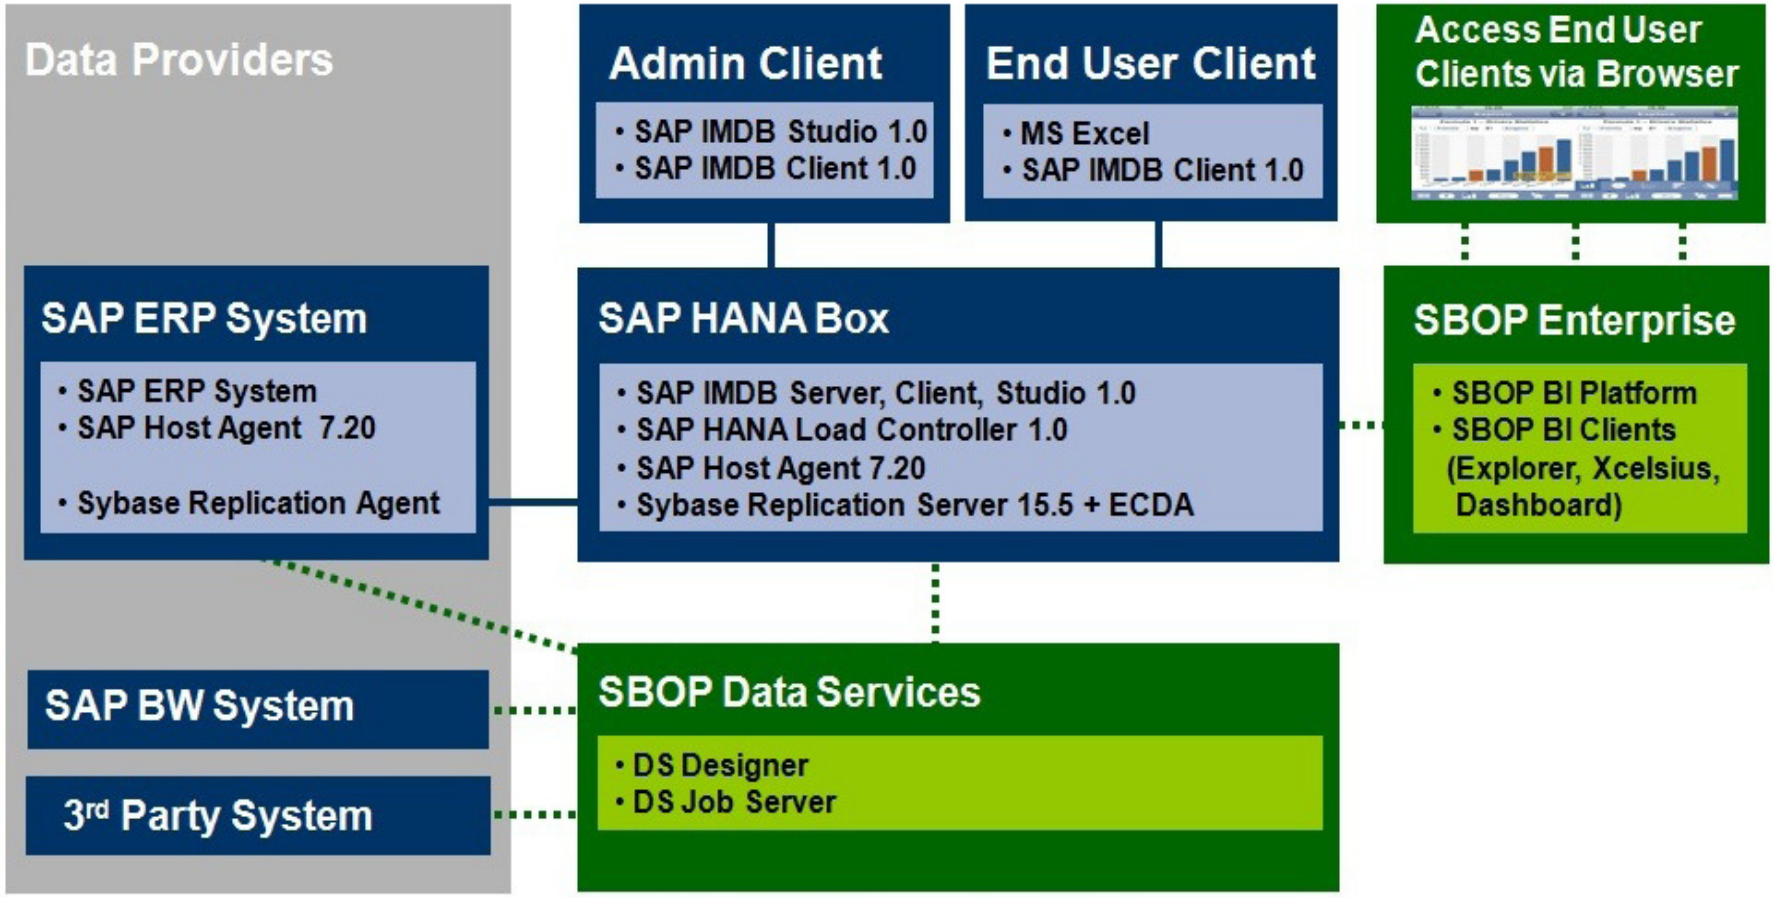

IMDB = SAP In-Memory Computing Engine
SBOP = SAP BusinessObjects Portfolio

HANA Architecture

HANA Architecture & Surrounding

Mahindra Satyam

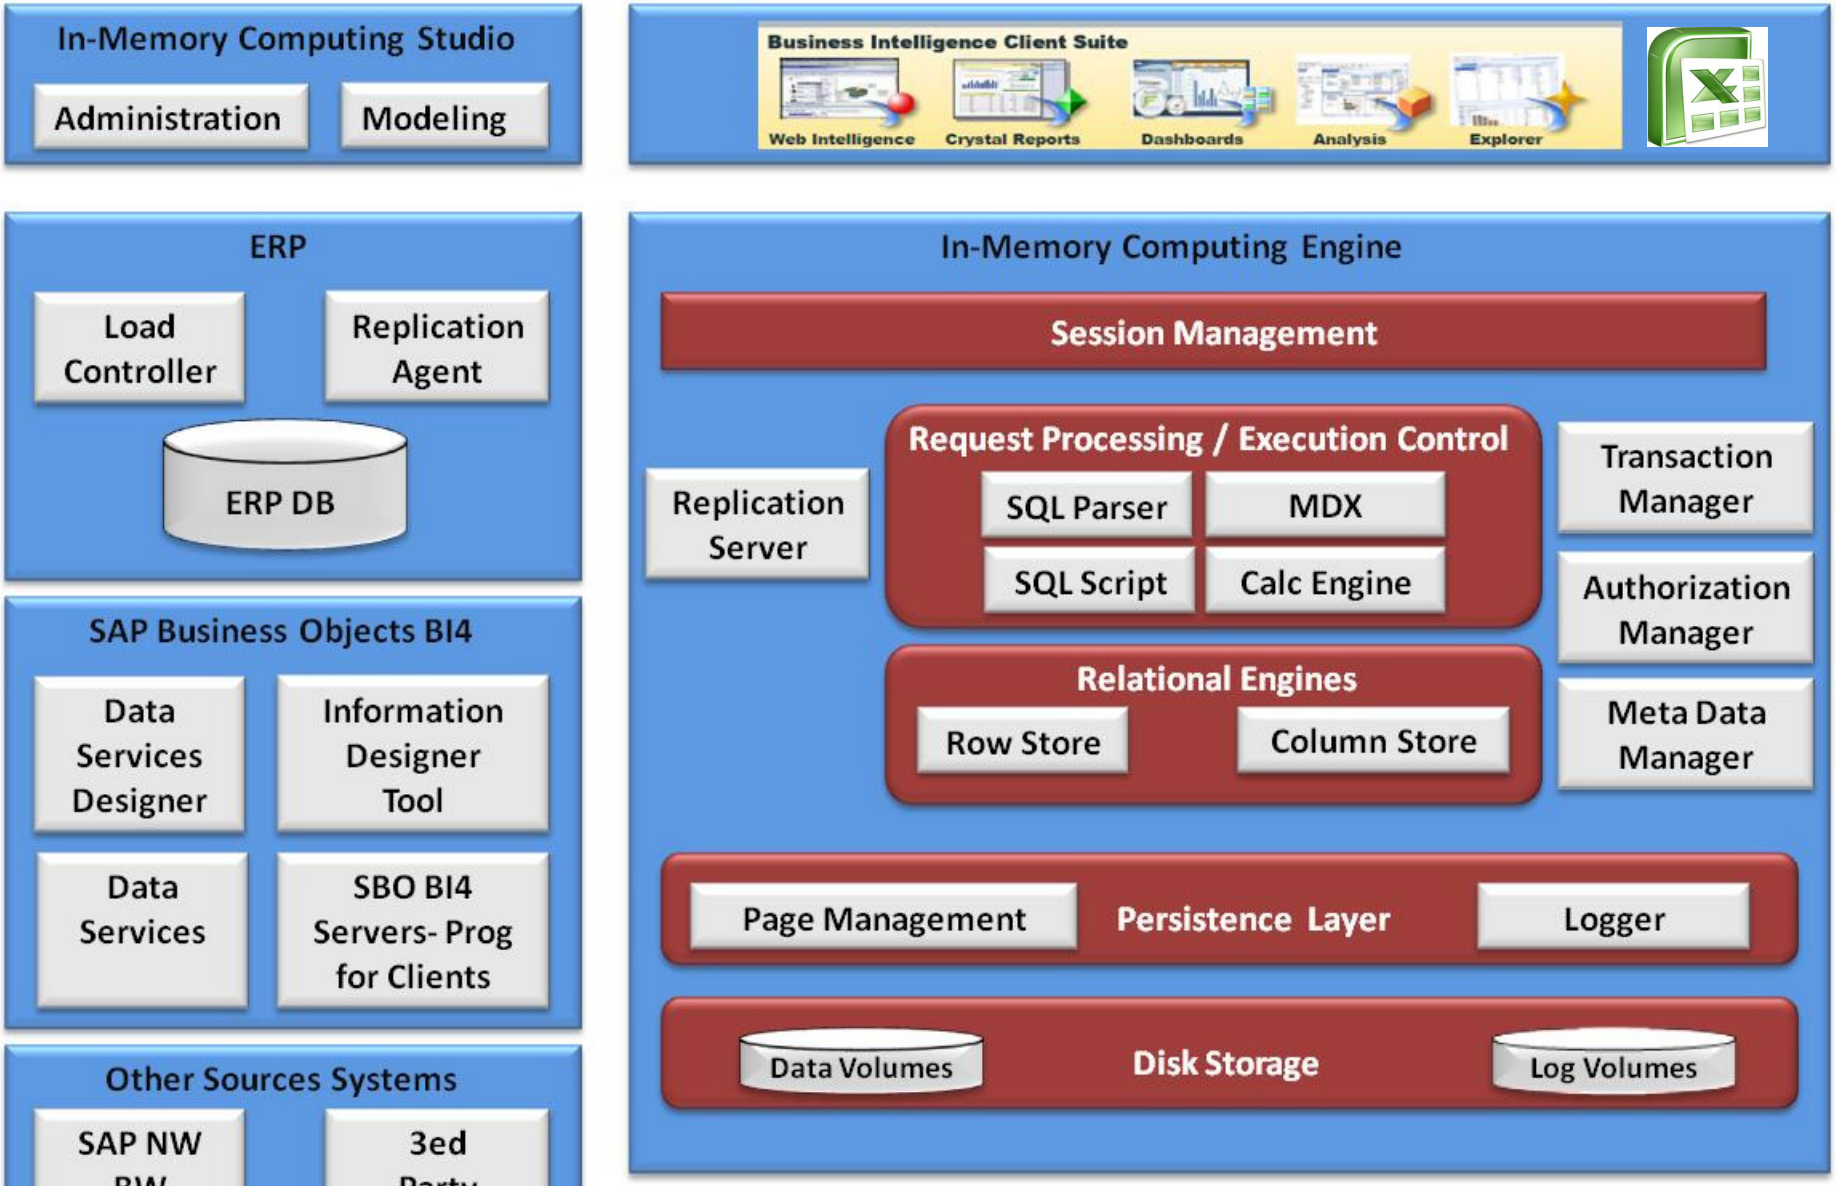

In-Memory Computing Engine

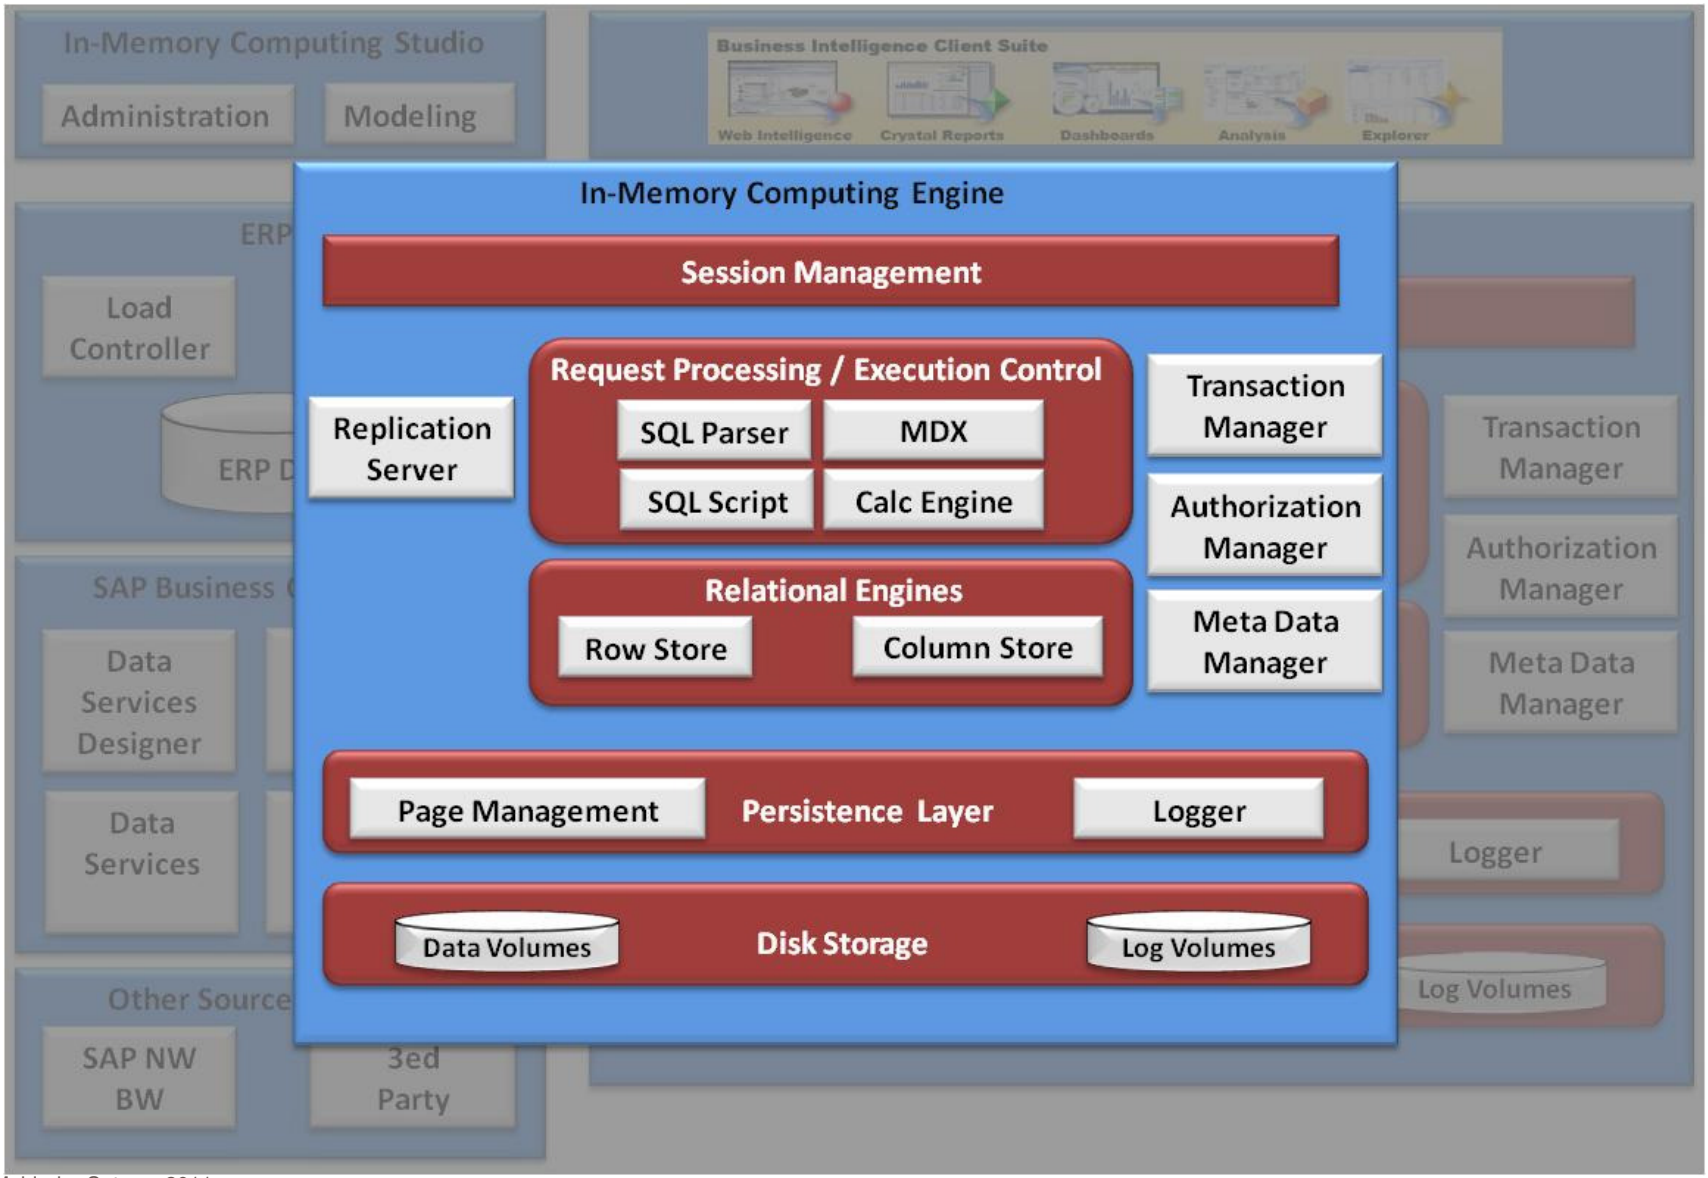

Loading Data in HANA

Loading Data into HANA
Options

Type	Tool
Trigger-Based Replication	SAP Landscape Transformation (SLT)
ETL-Based Replication	Business Objects - Data Services
Log-Based Replication	Replication Server

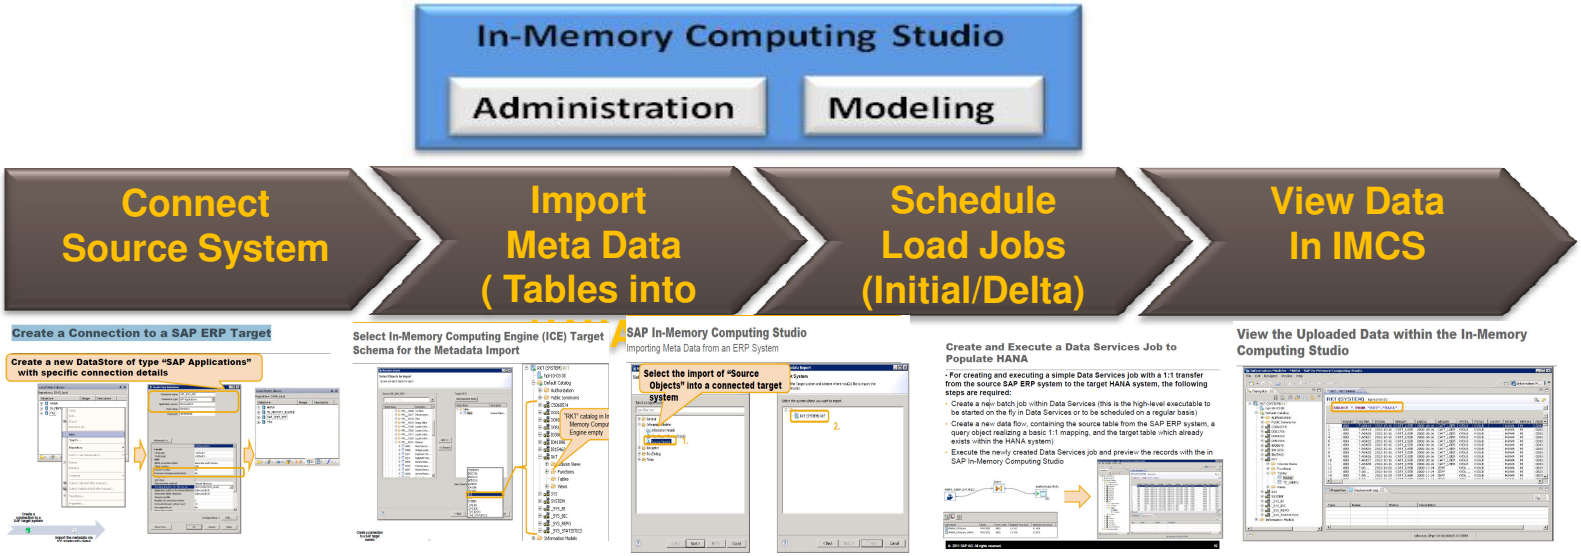

Loading Data into HANA

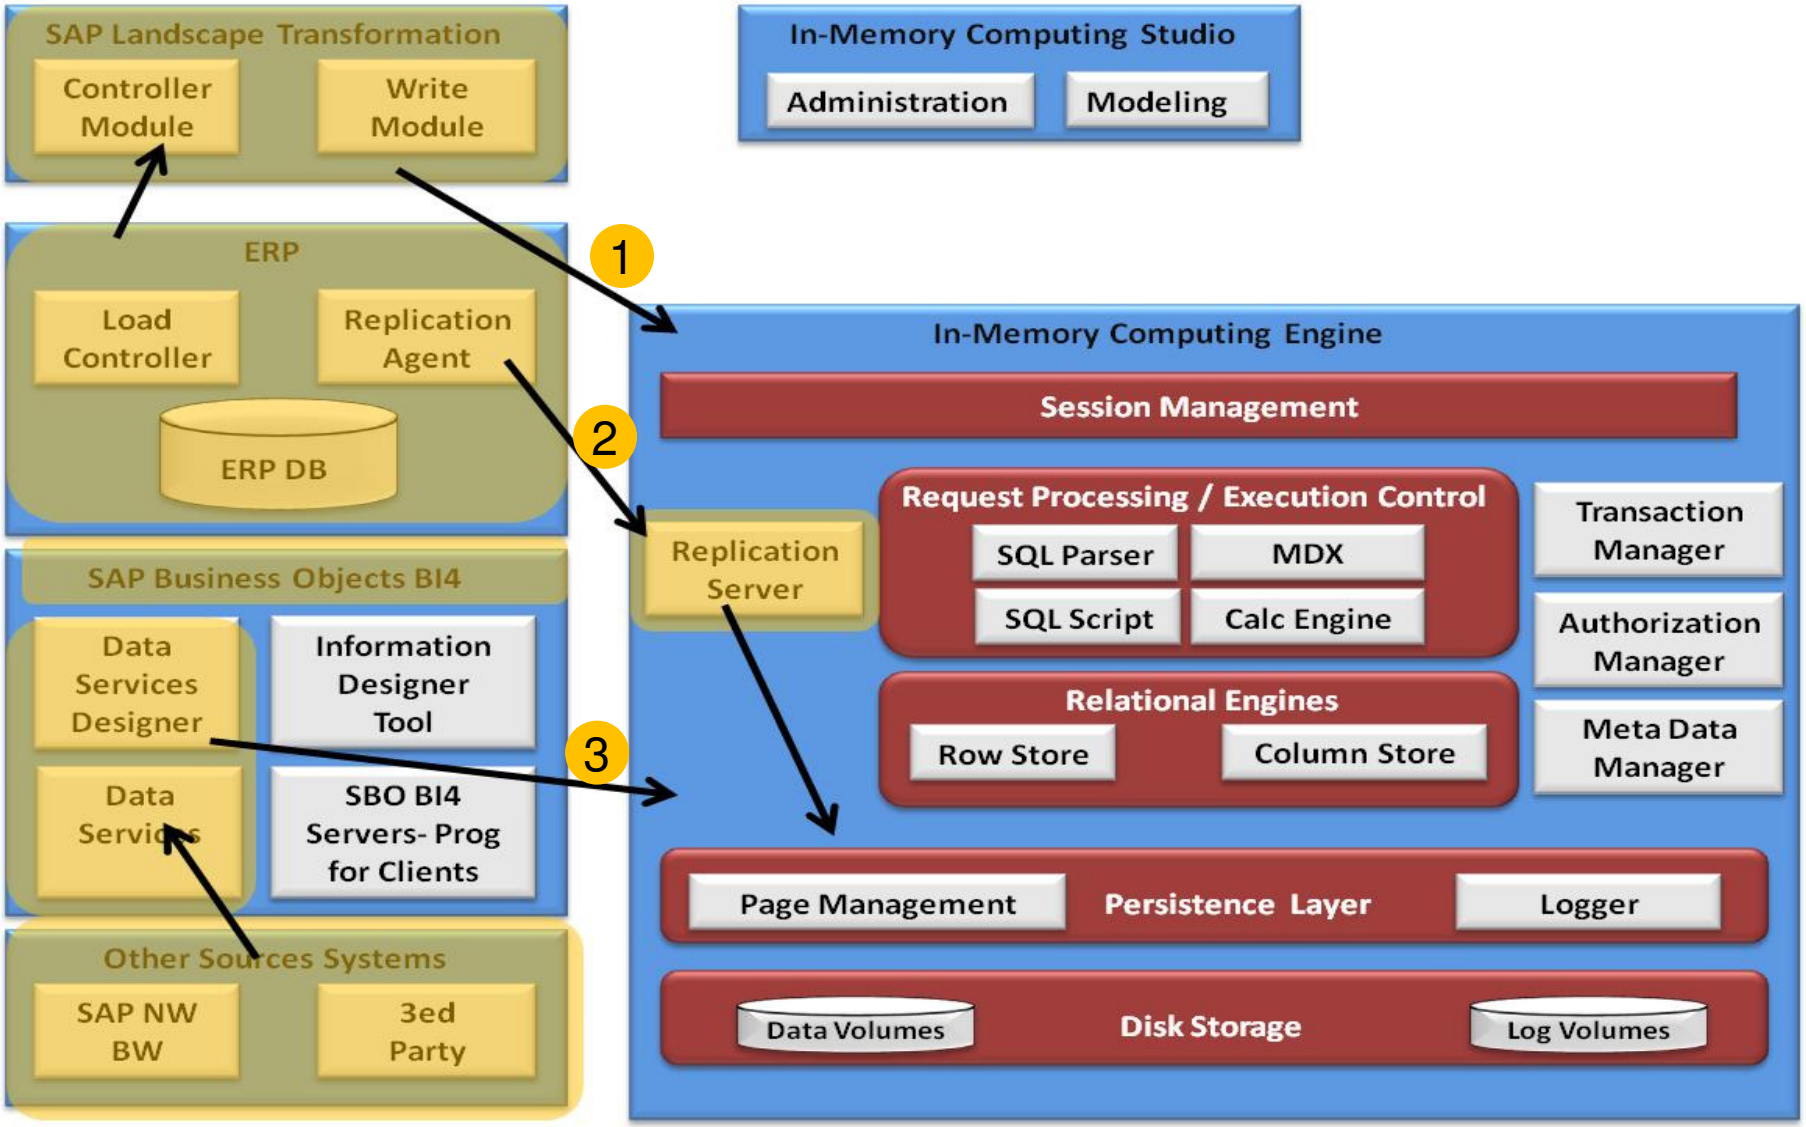

Trigger-Based Replication

SAP Landscape Transformation (SLT)

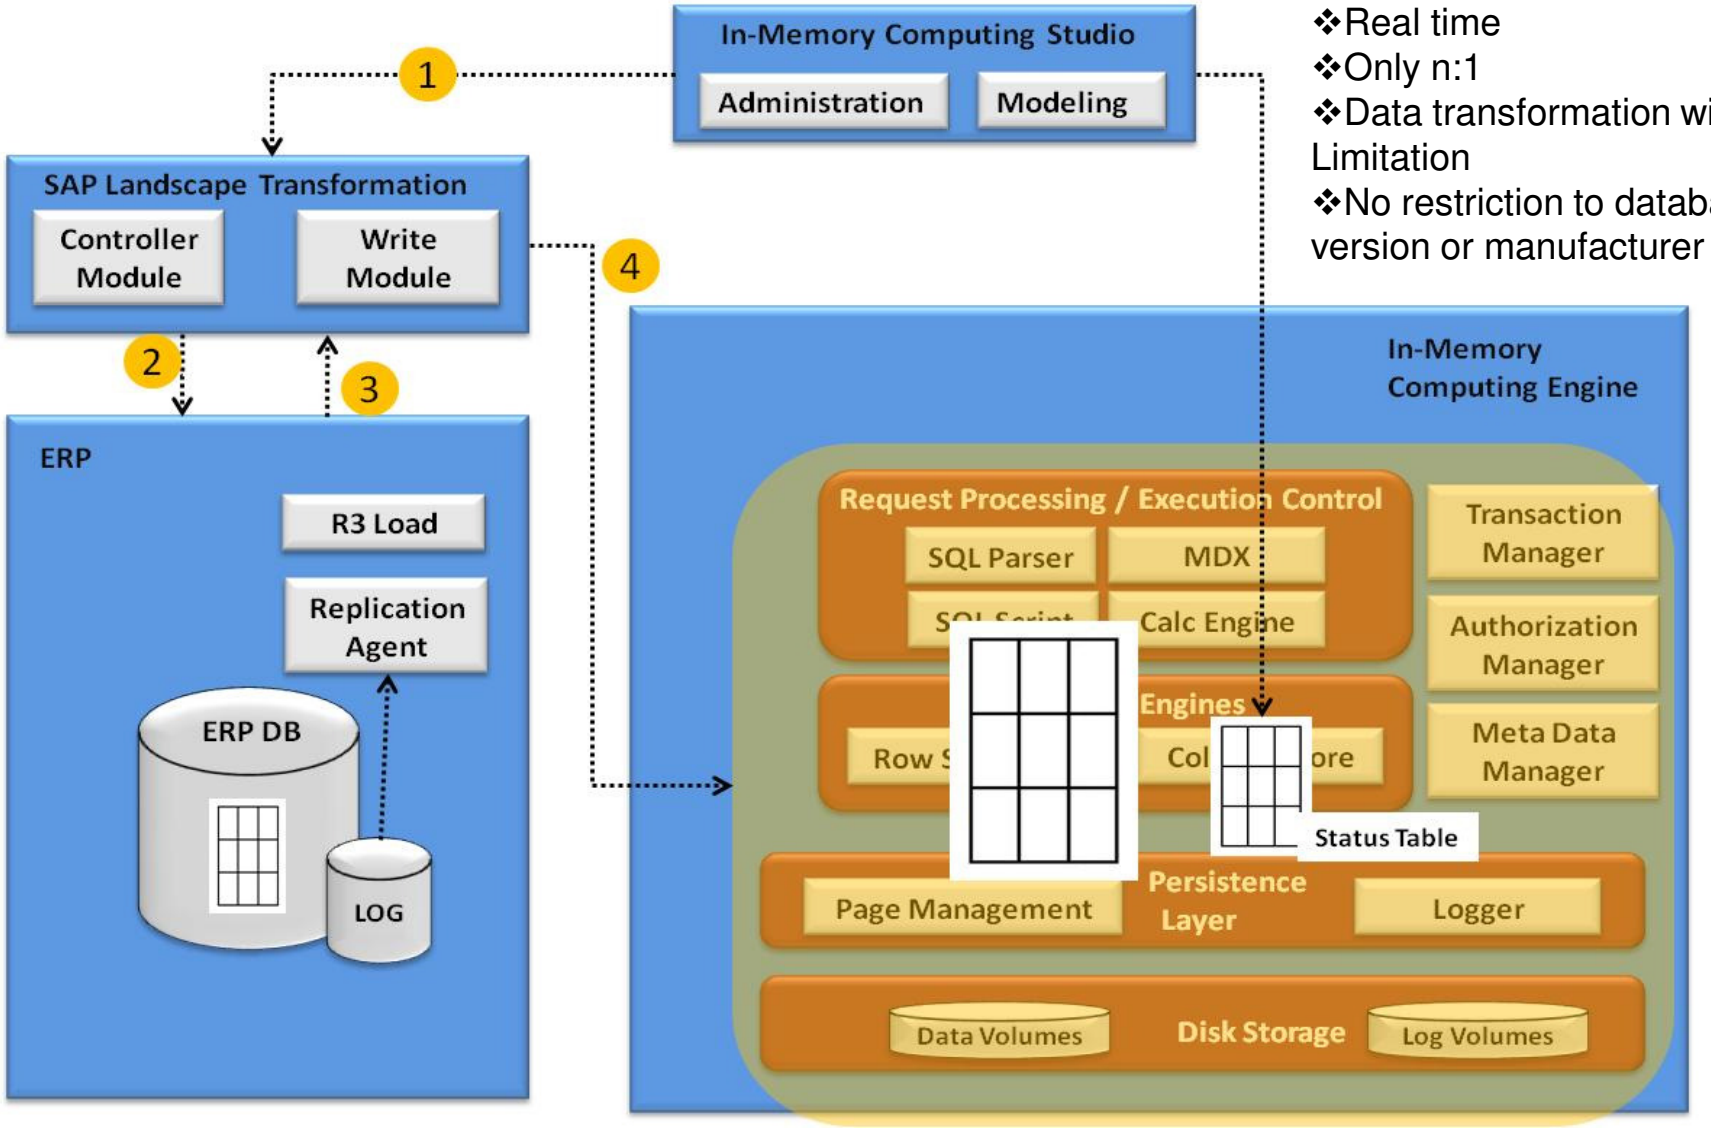

- Real time
- Only n:1
- Data transformation with Limitation
- No restriction to database version or manufacturer

ETL-Based Replication

Data Services

Mahindra Satyam

- ❖Delay more than 30 minutes
- ❖n:m
- ❖Data transformation Possible
- ❖ No restriction to database version or manufacturer

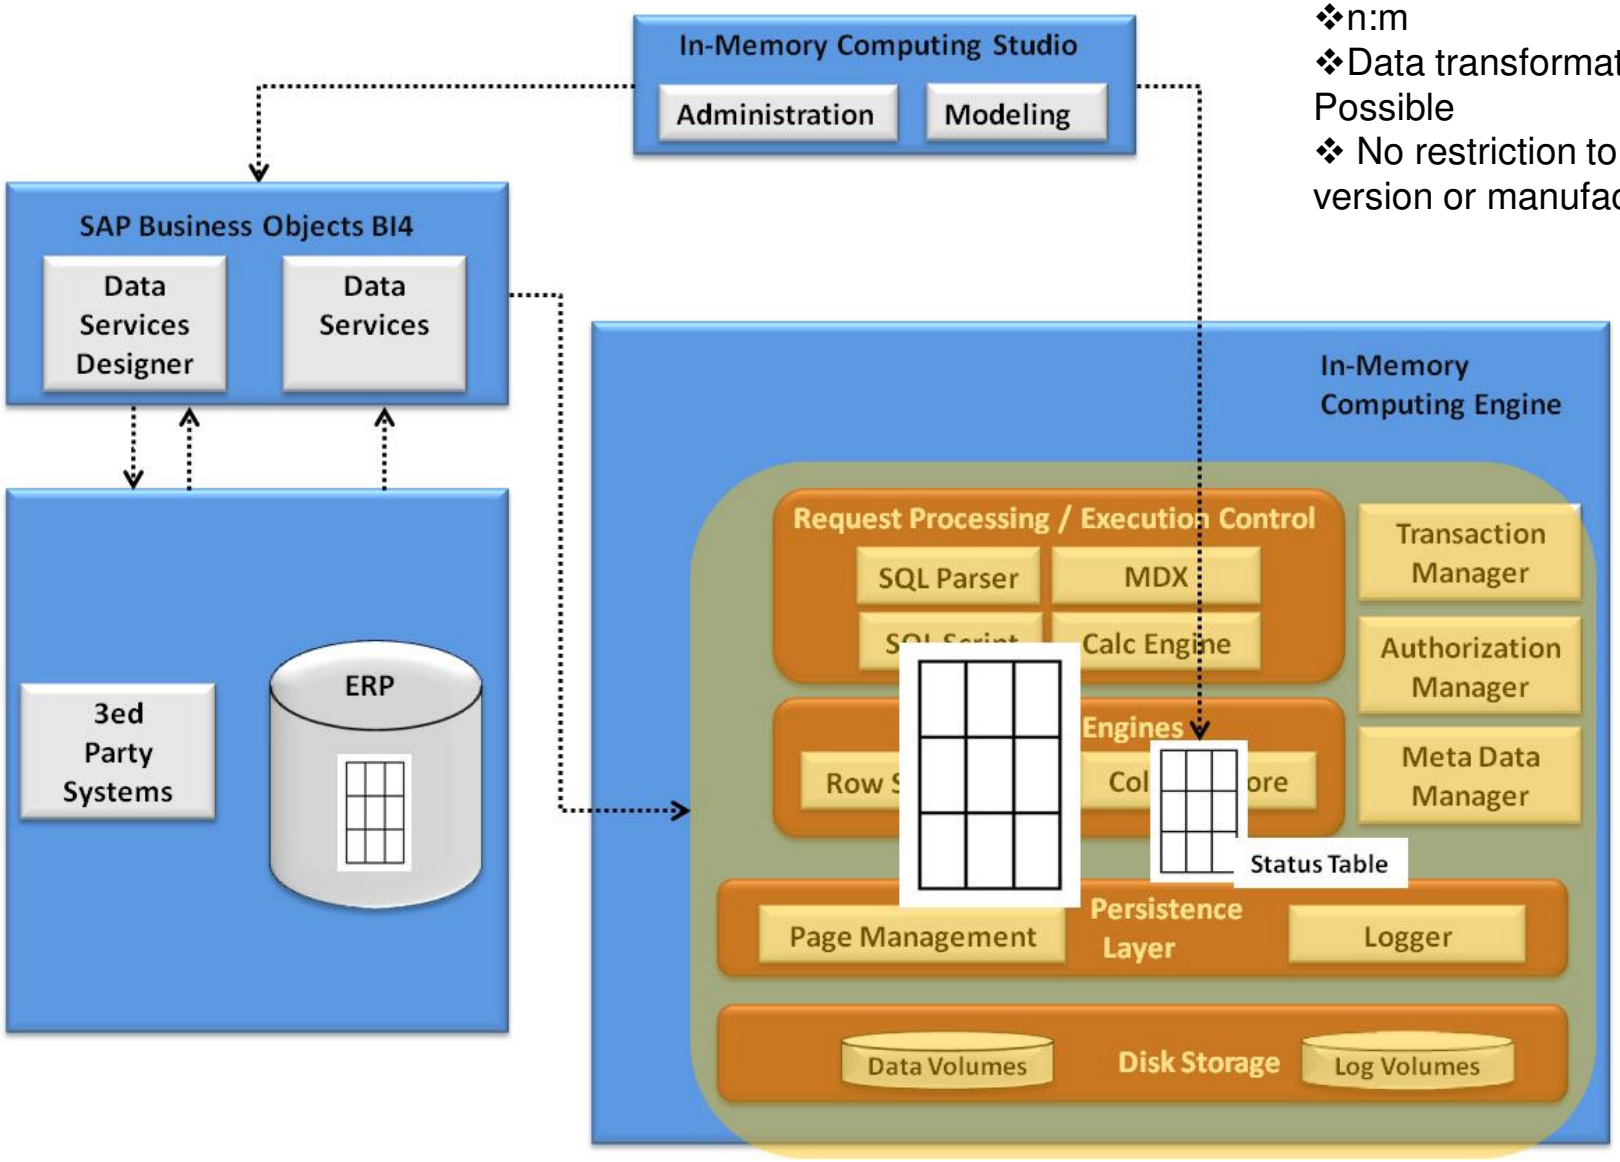

Log-Based Replication

Replication Server

- Real time
- Only 1:1
- No Data transformation
- Restriction to database version or manufacturer

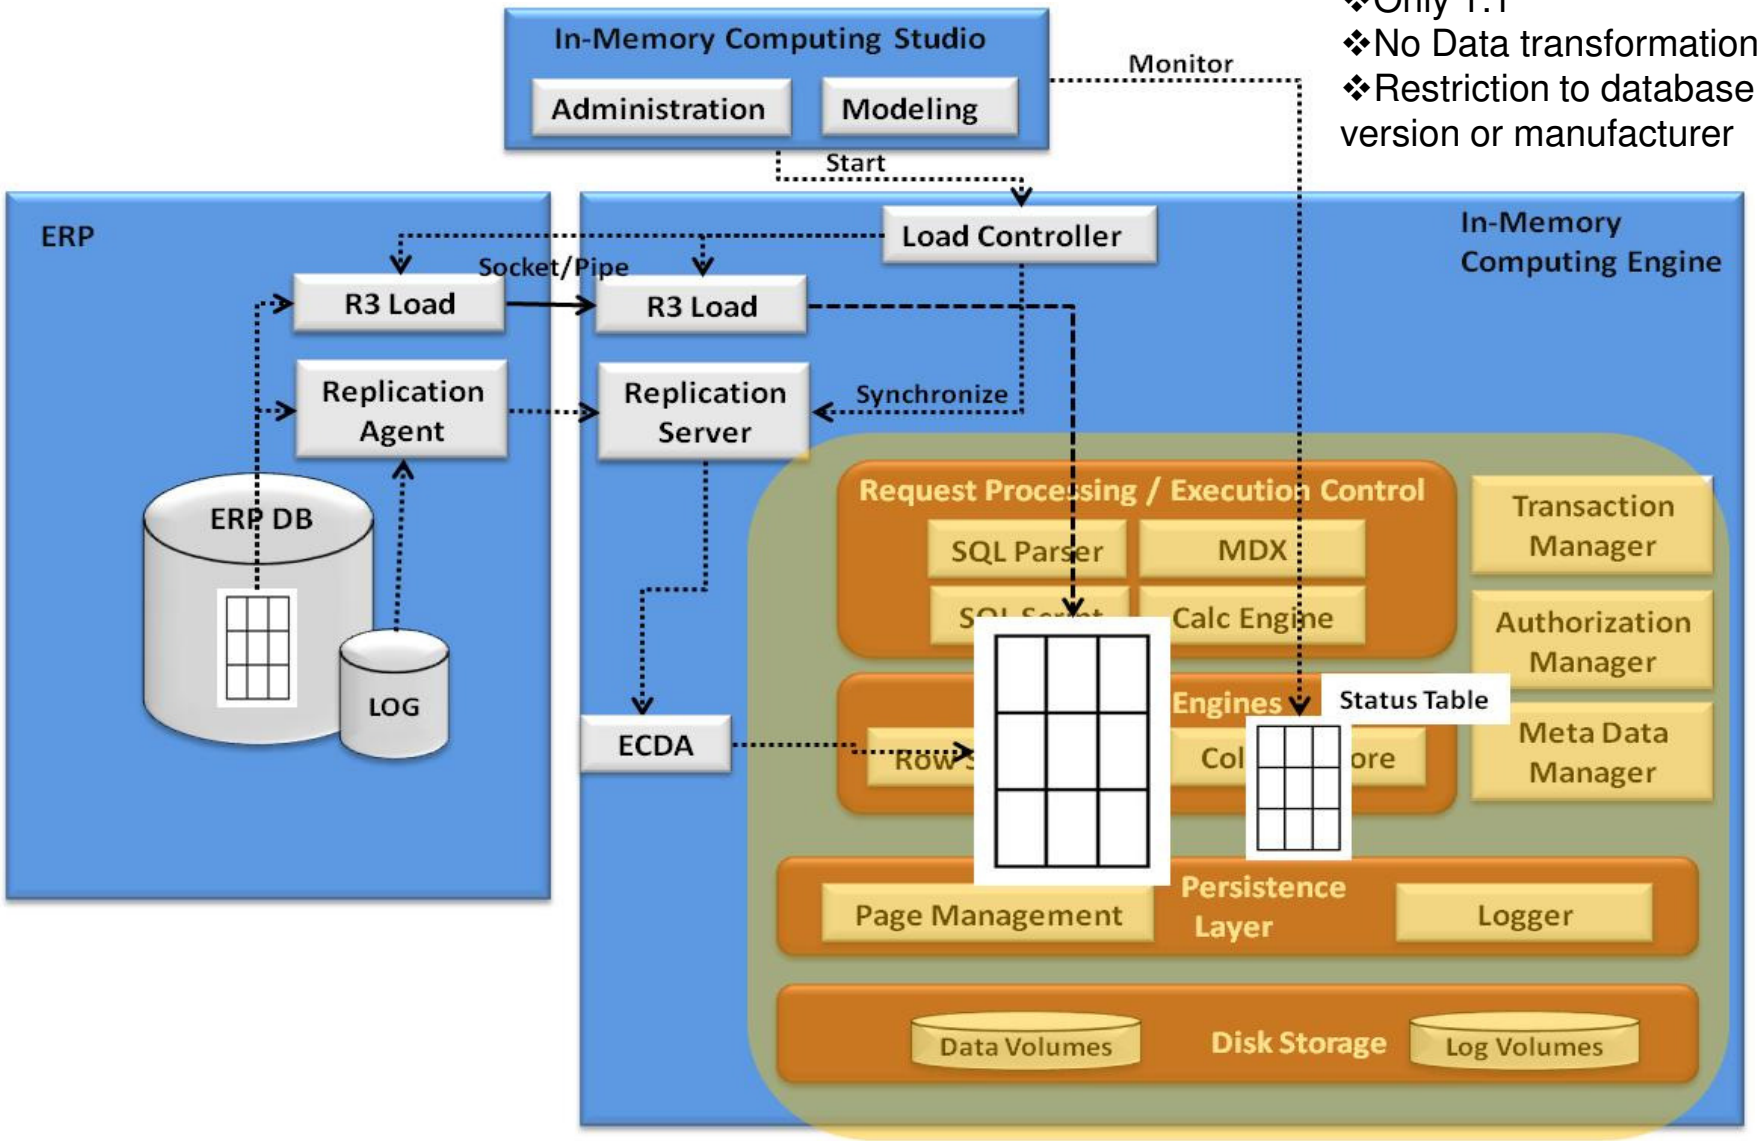

Replication Methods Comparison

Features	Trigger-Based Replication	ETL-Based Replication	Log-Based Replication
Real time or latency	Real time	Delay more than 30 minutes	Real time
n:m replication (many-to-many relationship support)	NO Only n back ends to 1 SAP HANA	YES	No Only 1:1
Table data transformation during replication (add or remove fields, conversions)	Yes, but is to be defined manually during project setup	YES	No
Free of restriction to database version or manufacturer	Yes	Yes	No
Free of restriction to operating system version	Yes (separate SLT system required if source system kernel < 720)	No NW720 required (Delta queue patch)	No
Pool /Cluster tables supported	Yes	Yes	No

Modeling in HANA

Modeling in HANA

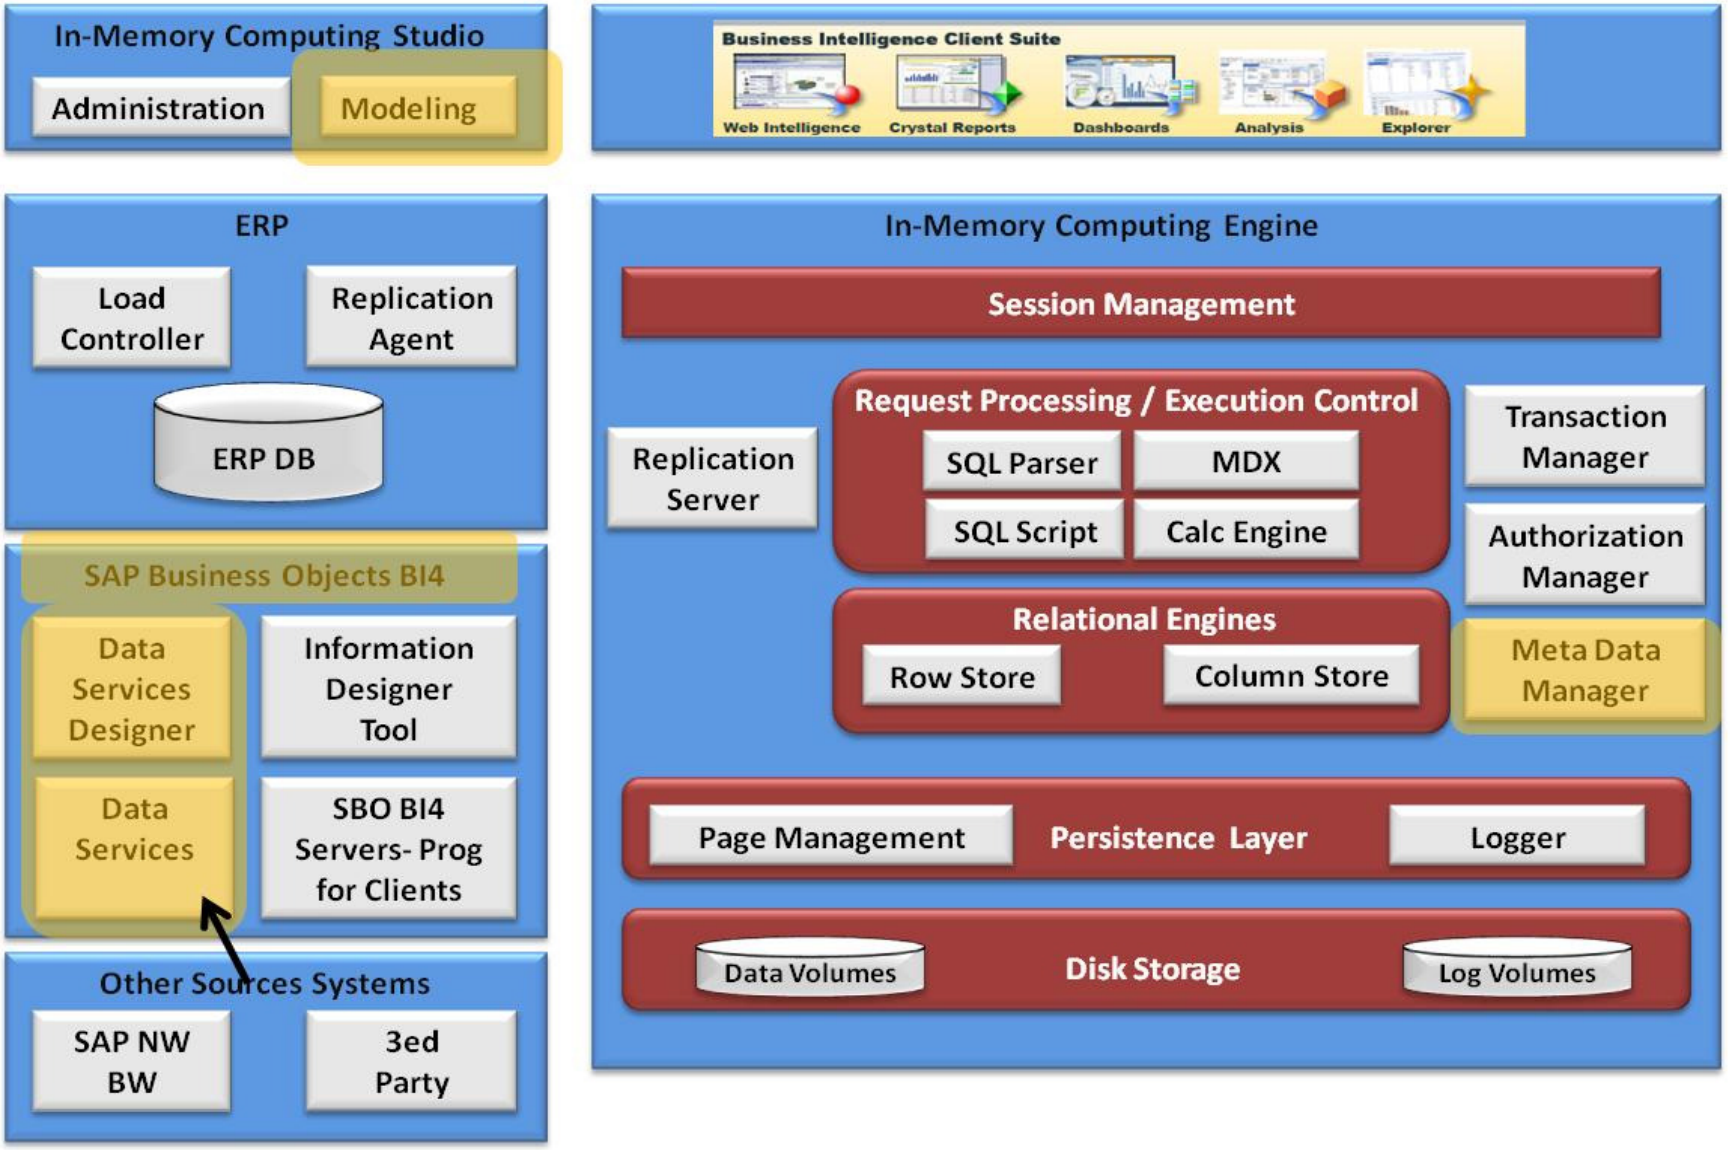

SAP In-Memory Computing Studio Look and Feel

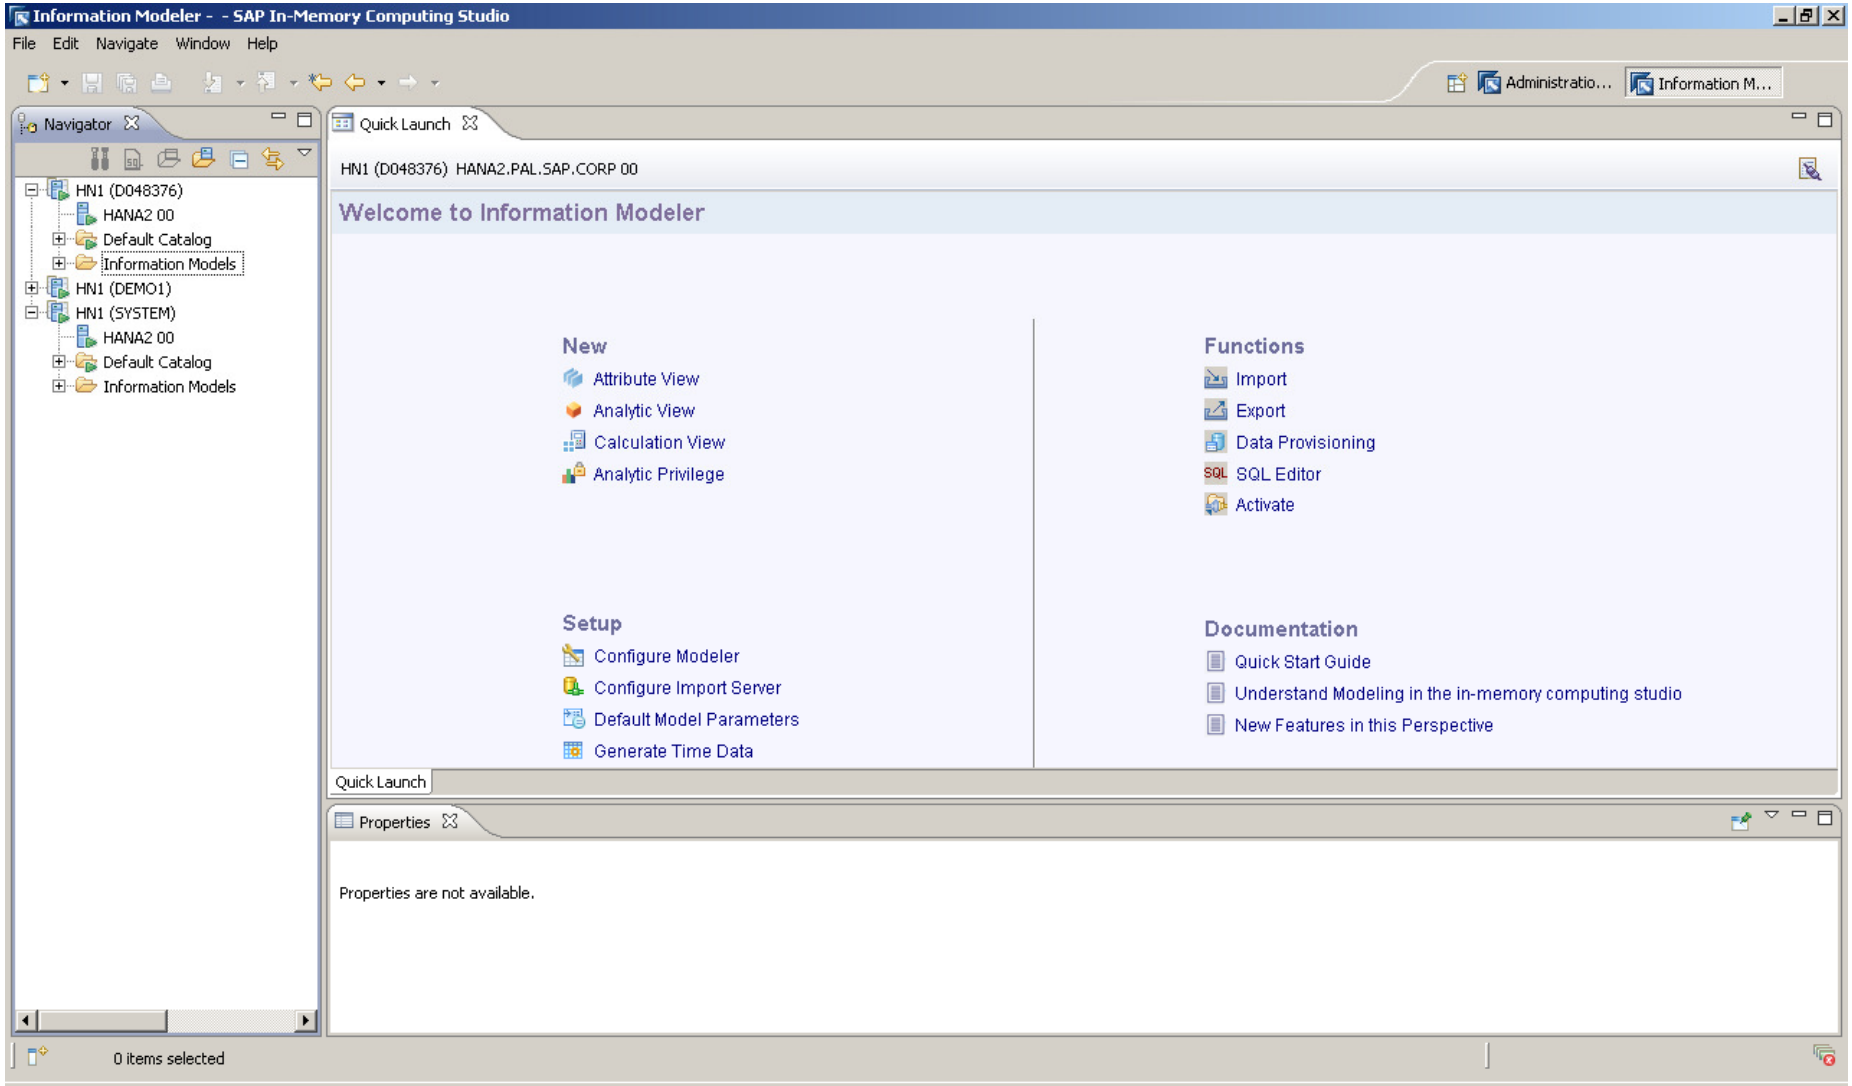

Modeling in HANA

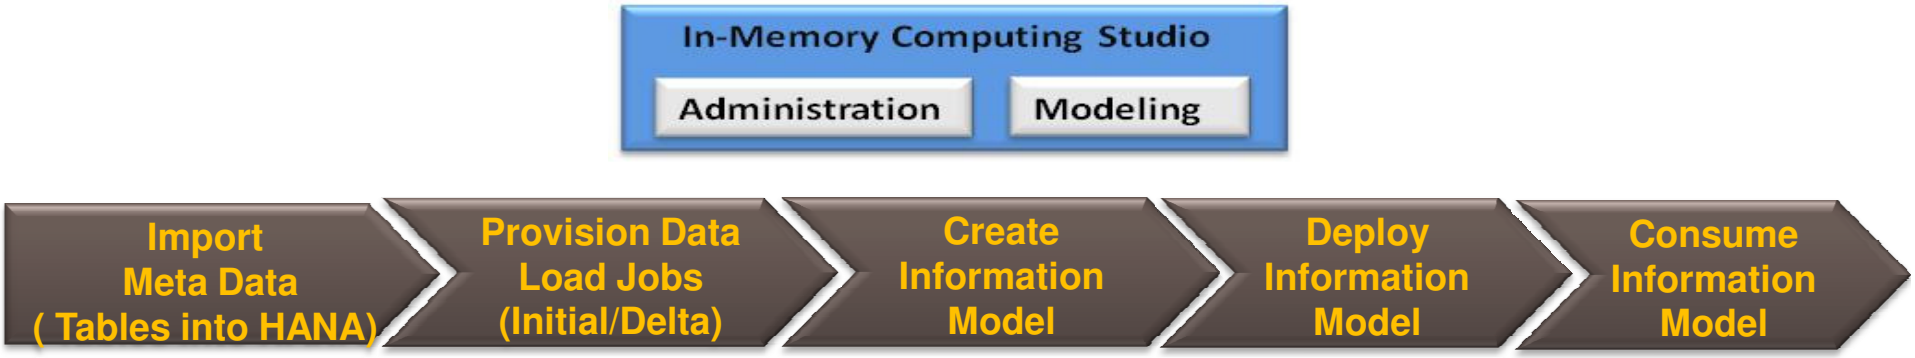

Data

Attributes : Descriptive Data (Also known as Characteristics in BW terminology)

Measures : Data that can be Quantified & Calculated . (Also known as Key Figures in BW terminology)

Views

Attribute View: i.e. Dimensions

Analytic View : i.e. Cube

Calculation View: Similar to virtual providers with service concept in BW

Hierarchies

Leveled – based on Multiple Attributes

Parent Child Hierarchy

Analytical Privileges – **Security Objects**

Modeling in HANA
Attribute View

- Attributes add context to Data
- Attributes are modeled using Attribute Views
- Can be regarded as Master data Tables
- Can be linked to fact tables in Analytical views

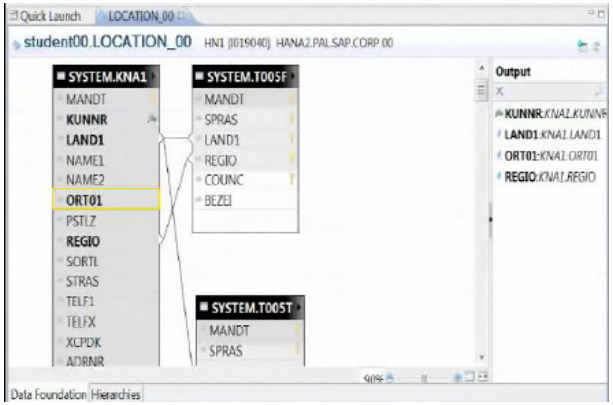

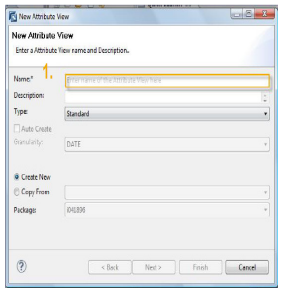

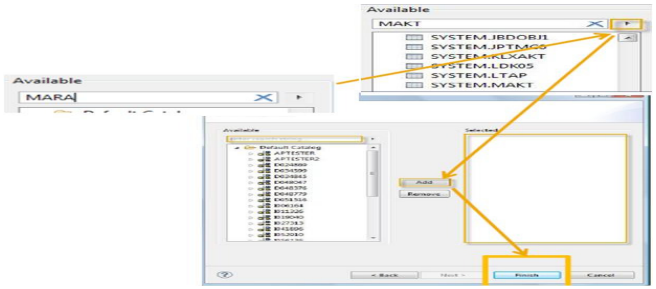

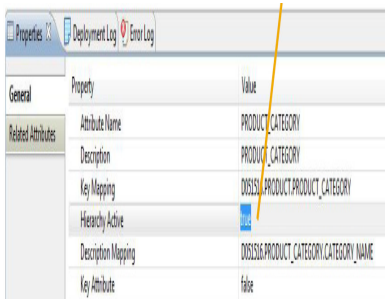

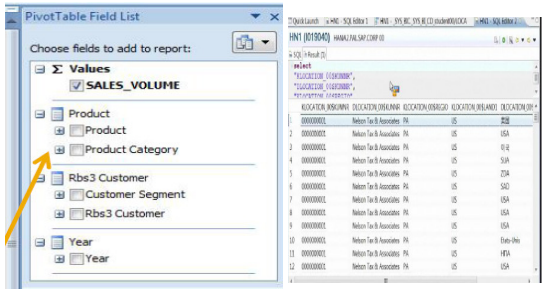

Modeling in HANA
Analytical View

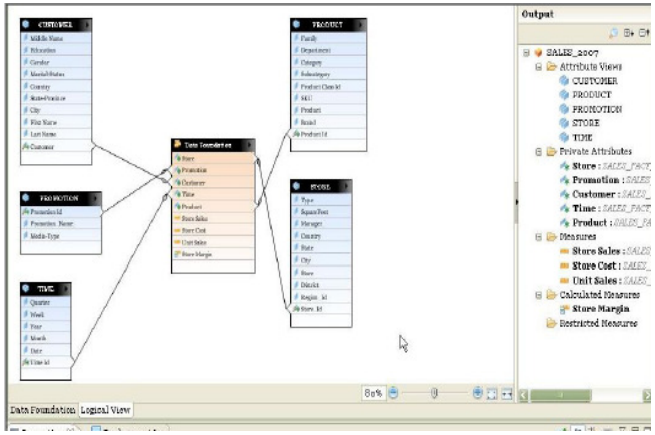

- Analytical View can be regarded as Cube
  - Multidimensional modeling
  - Fact table joined against modeled dimensions
- Analytical Views do not store data
  - Data is read from the joined tables
  - Joins and calculated measures are evaluated in runtime
  - Master data for MDX/BICS are stored in system tables

Set Parameters → **Table Selection** → **Table Join and Properties** → **Select Attributes& Measures** → **Create Restrict Measures** → **Save & Preview**

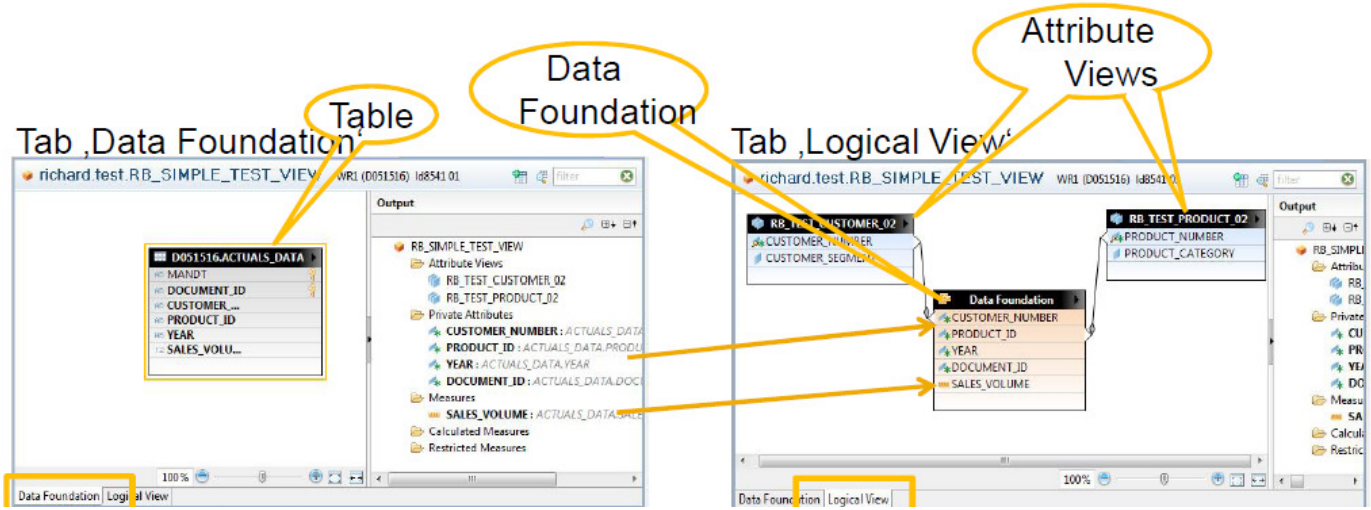

Modeling in HANA
Calculation View

Mahindra Satyam

Calculation view are used to create your own data foundation using database tables, attribute views, analytic views and calculation views to address a complex business requirement.

For example, Compare the sales of product in a particular region for the last two years.

Set Parameters → **Table Output Structure** → **Define Function (I/O Structure)** → **Write SQL Statement** → **Execute Assign Attr & Measure** → **Data Preview**

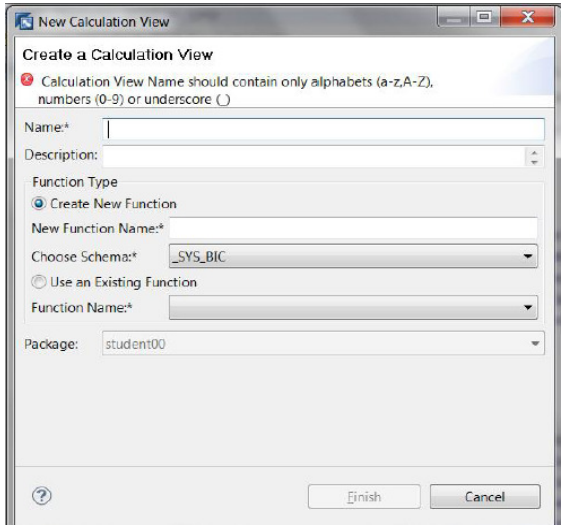

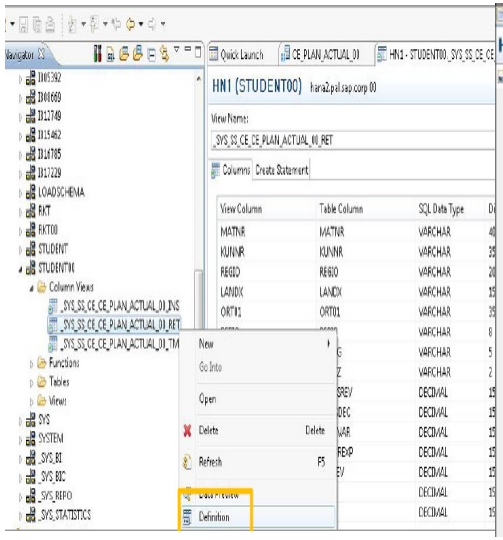

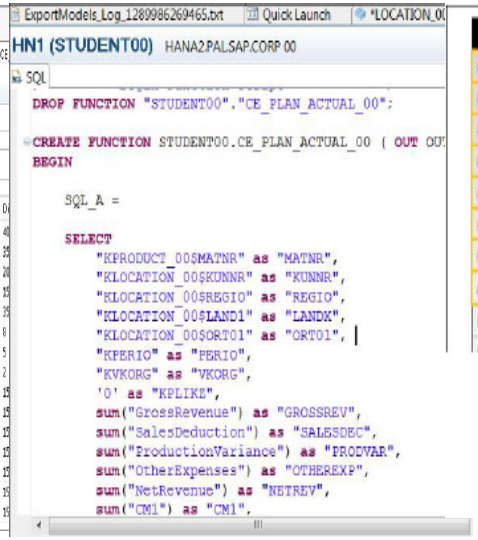

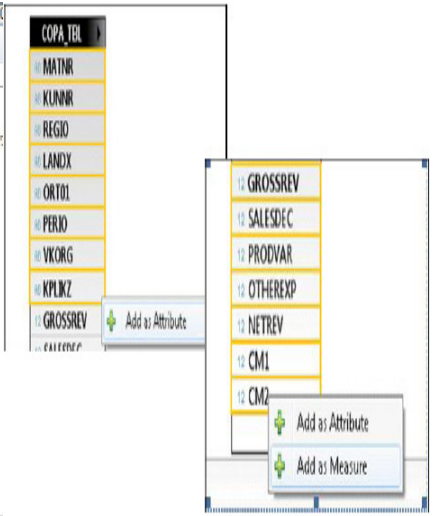

Reporting On HANA

Reporting

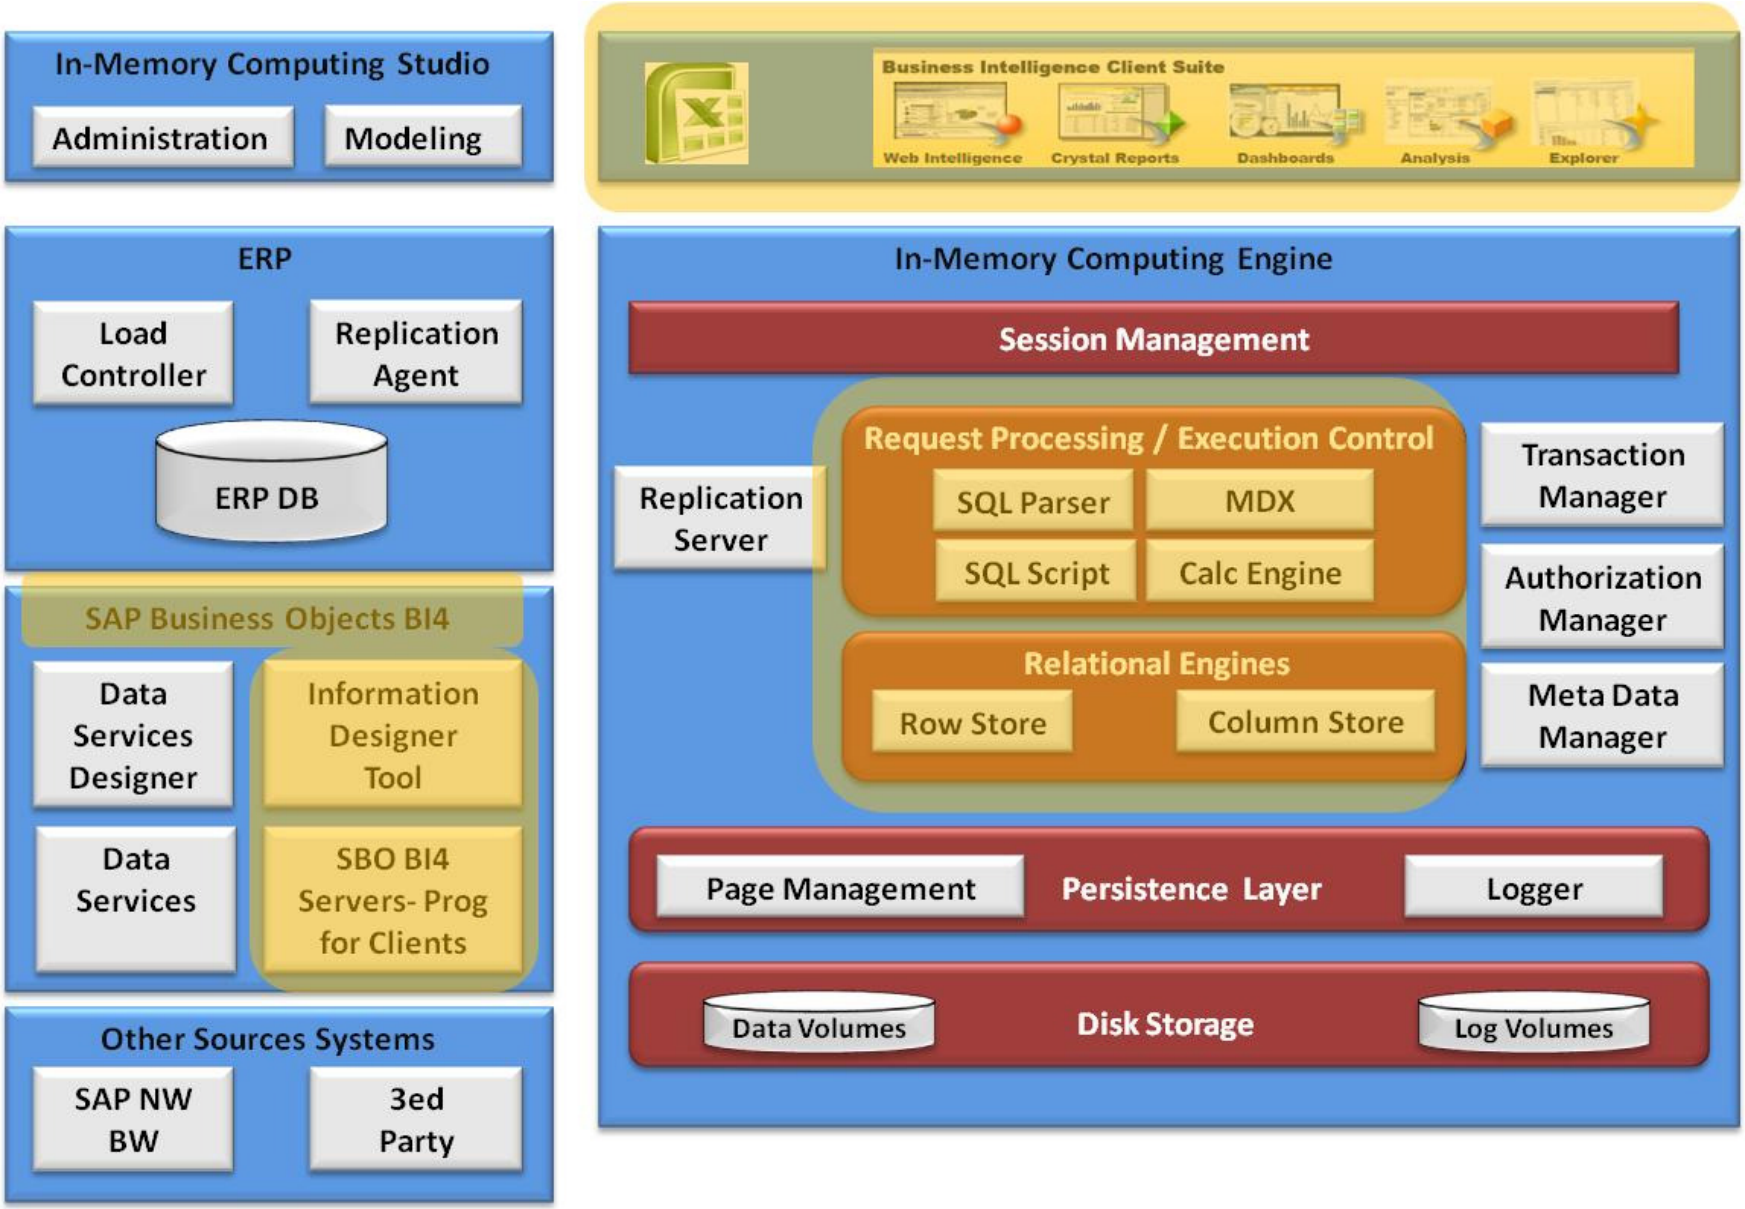

Various Interface Reporting Options

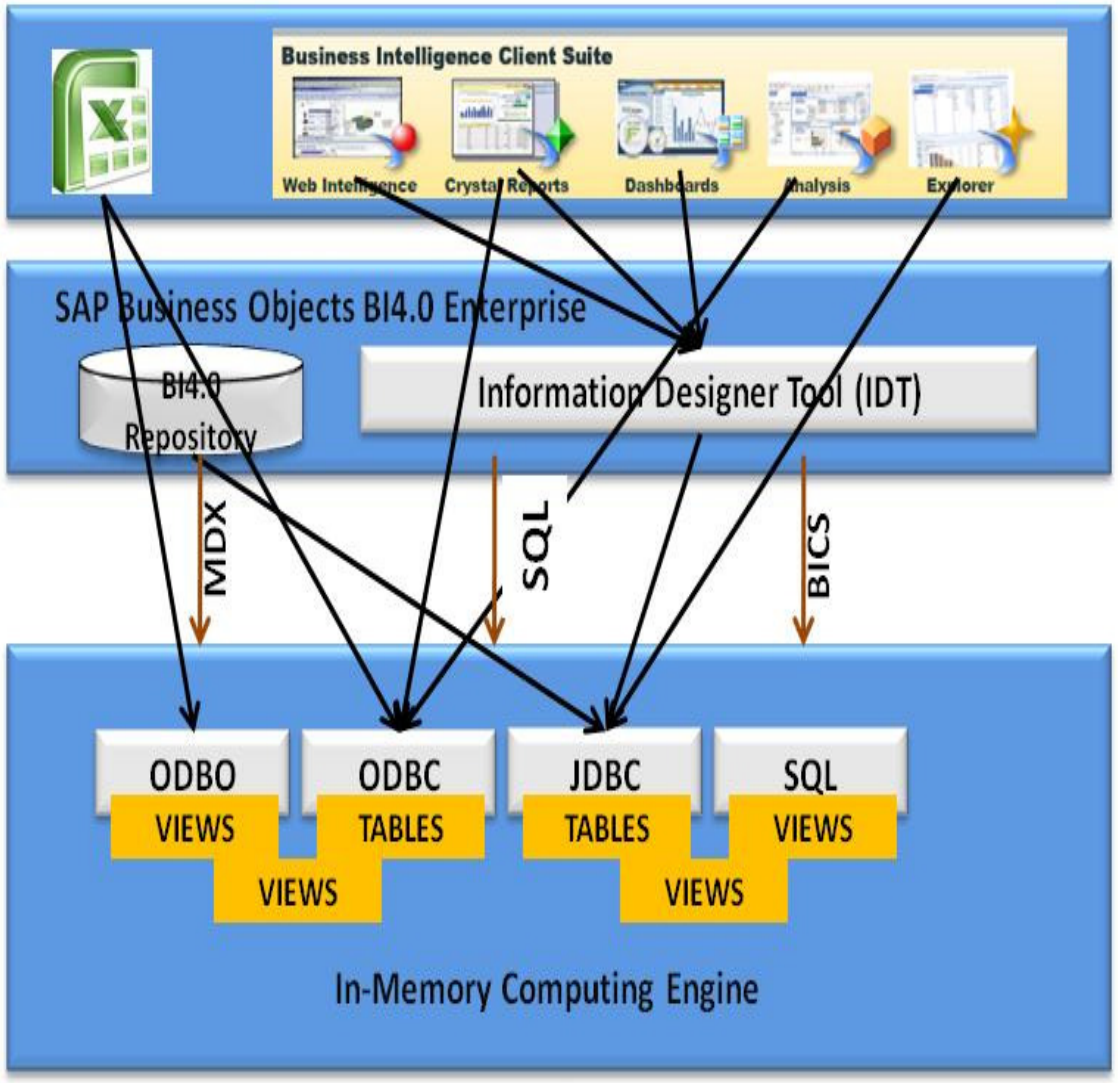

Reporting Interface Options

ODBO-(OLEDB for OLAP)

- Microsoft-driven specification for multidimensional(cross-tab style)reporting
- Requests are sent to the database via **MDX**(Multi Dimensional eXpression language)

ODBC-(Open Database Connectivity)-
Microsoft-driven specification for relational reporting .Database requests are made via SQL(Structure Query Language)

JDBC –(Java Database Connectivity)-
Relational reporting drivers specified by the Java community.

SQLDBC is SAP native database SDK

BICS-BI Consumer Services

- This is the common driver technology used by SAP Business Objects Analysis, Office Edition for connectivity to SAP NetWeaver BW

Reporting Interfaces

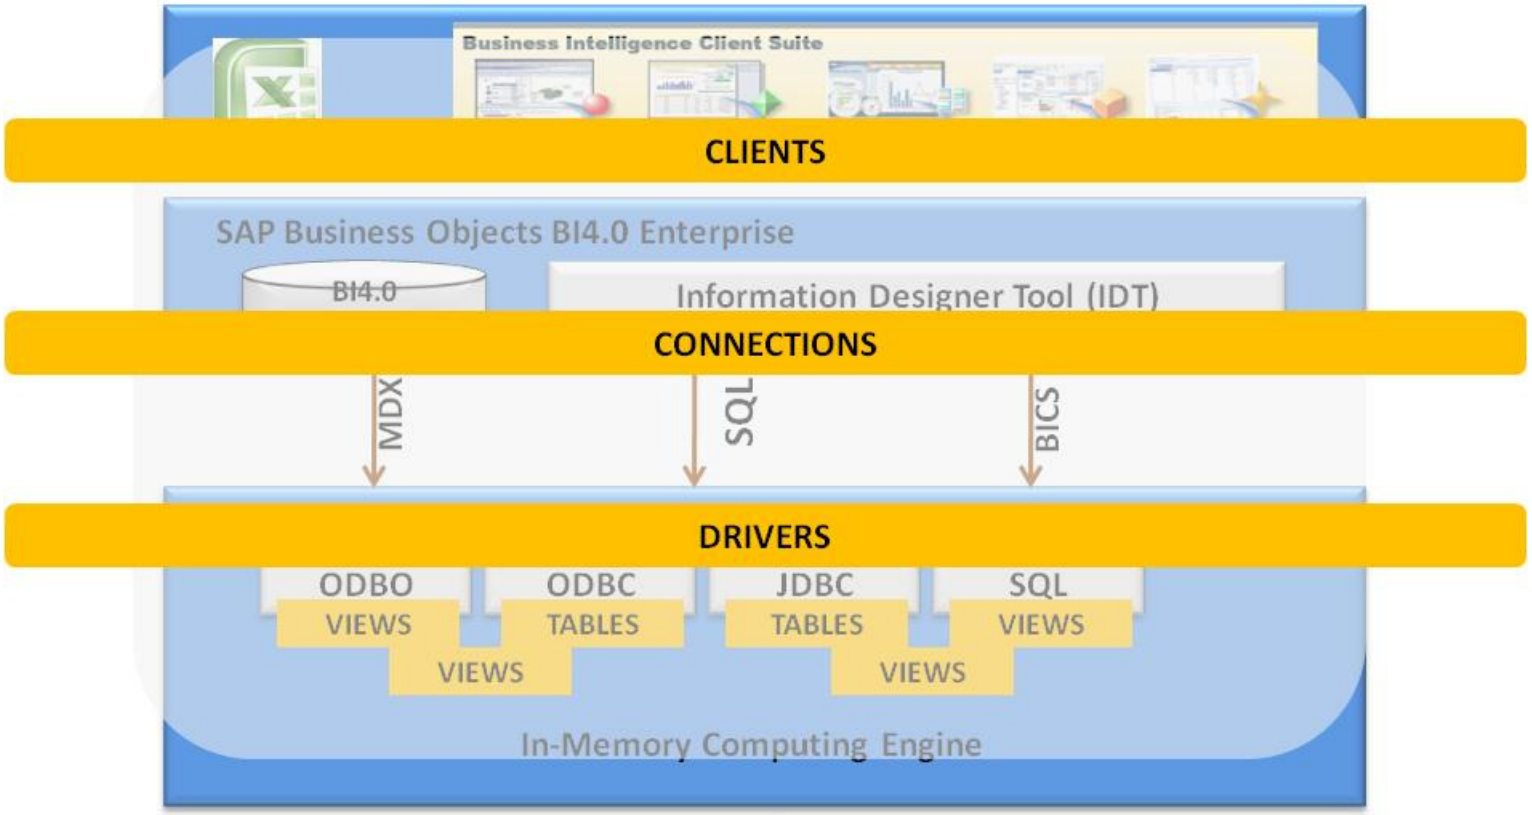

Reporting on HANA with Business Objects 4.0 Information Designer Tool

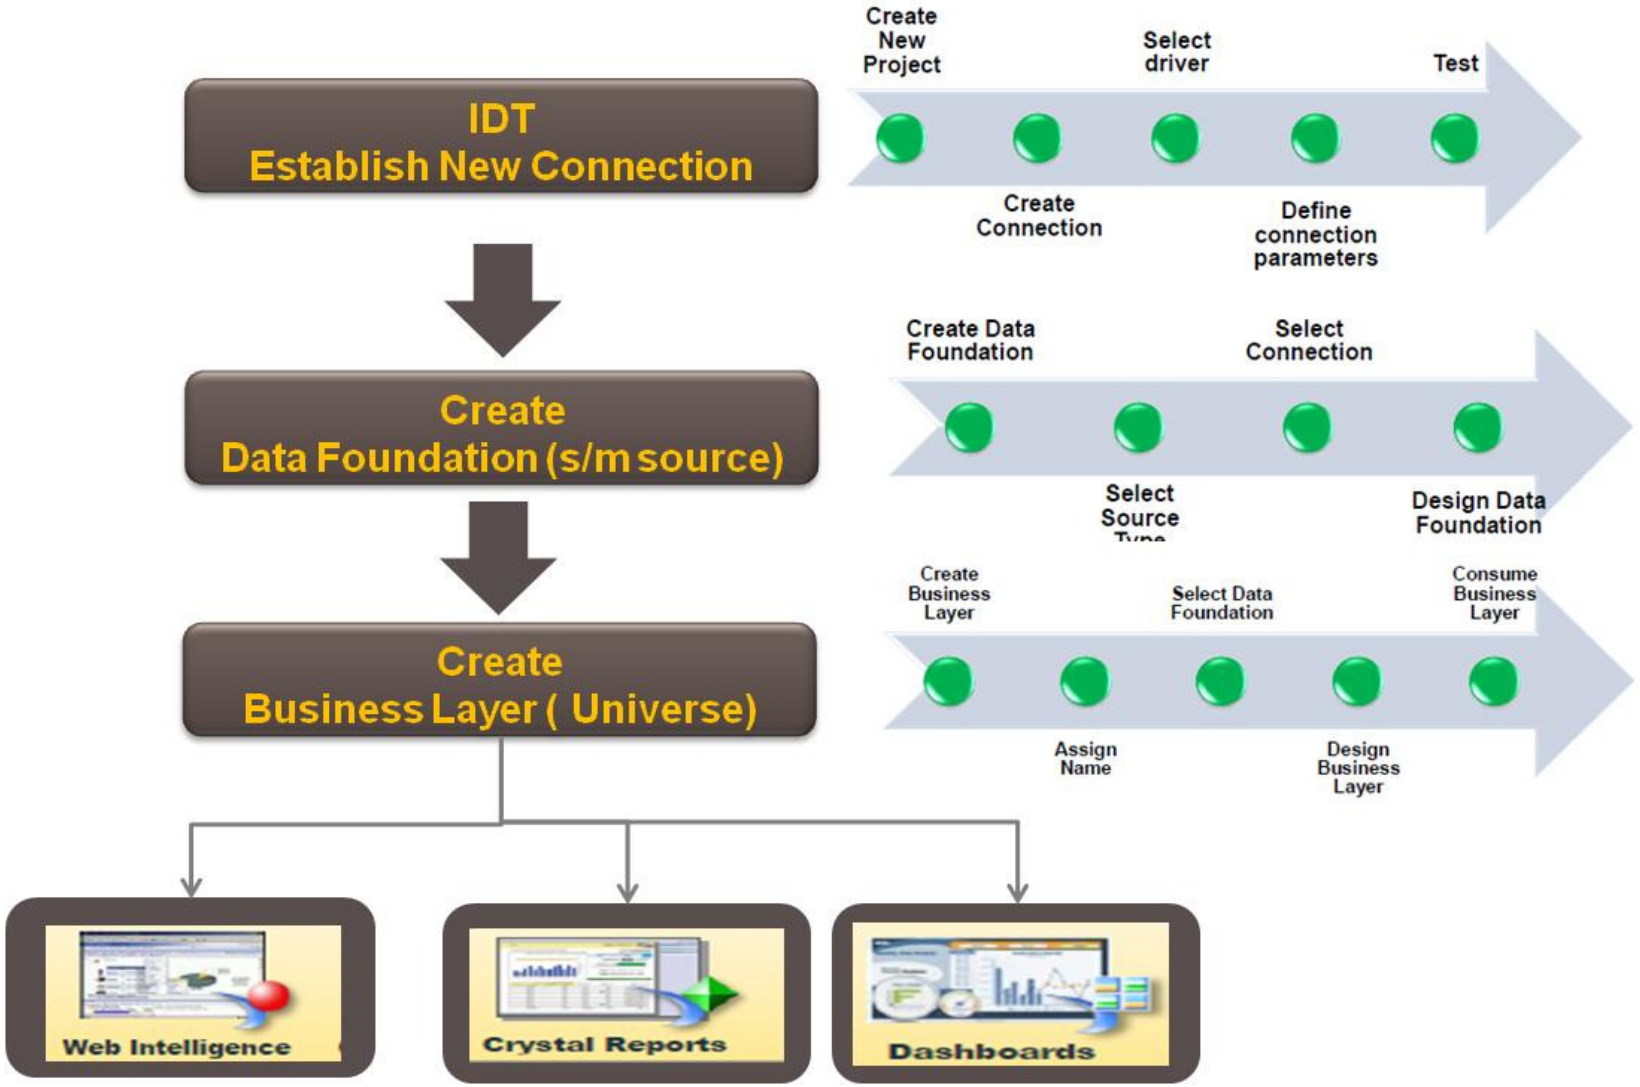

Request Processing and Execution

Request Processing and Execution

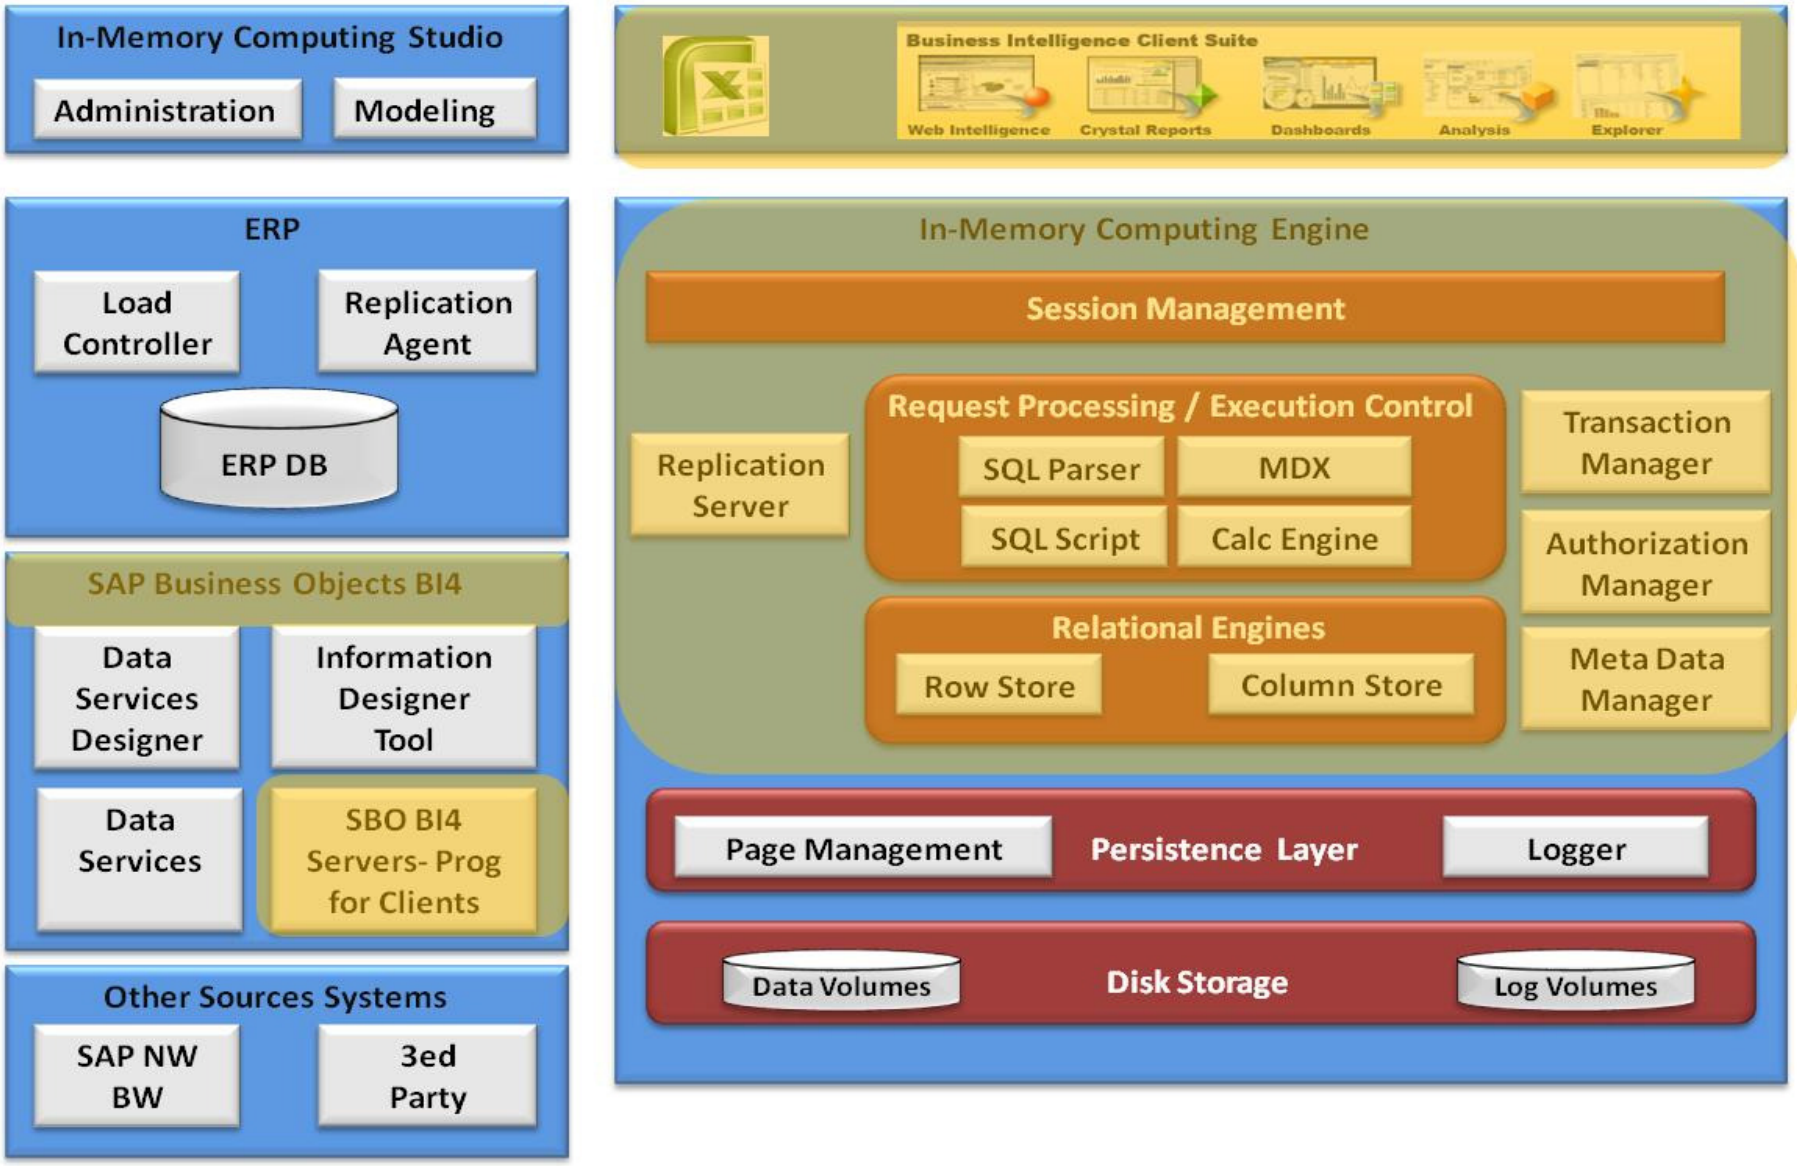

Request Execution and Processing

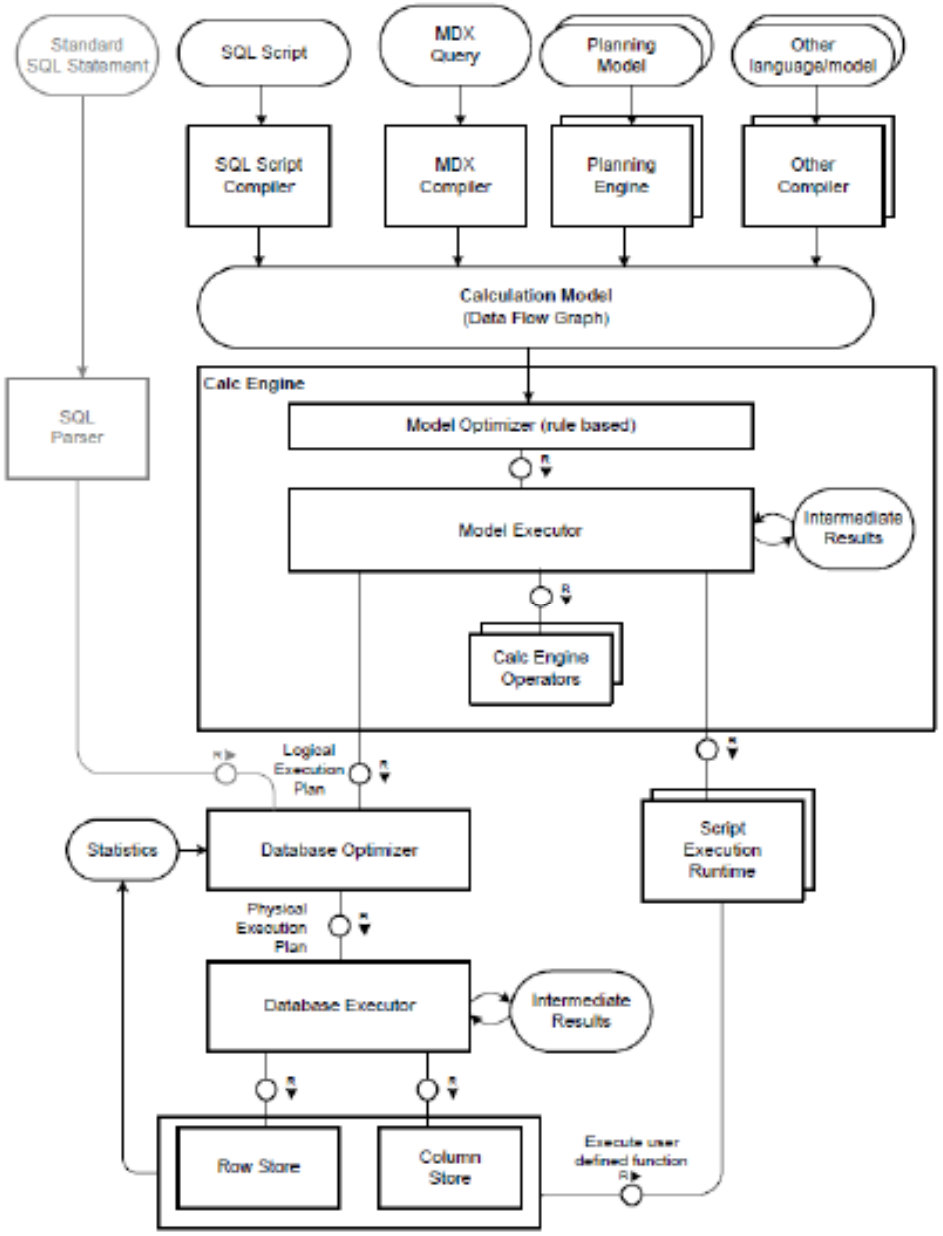

Conceptual View

•Standard SQL

Processed directly by DB engine

SQL Script , MDX and planning engine interface

Domain-specific programming languages or models

Converted into calculation models

Calc Engine

Create logical execution plan for calculation models

Execute user defined functions

Relational engine

DB optimizer produces physical execution plan

Access to row and column store

Row Store

Row Store - High Level Architecture

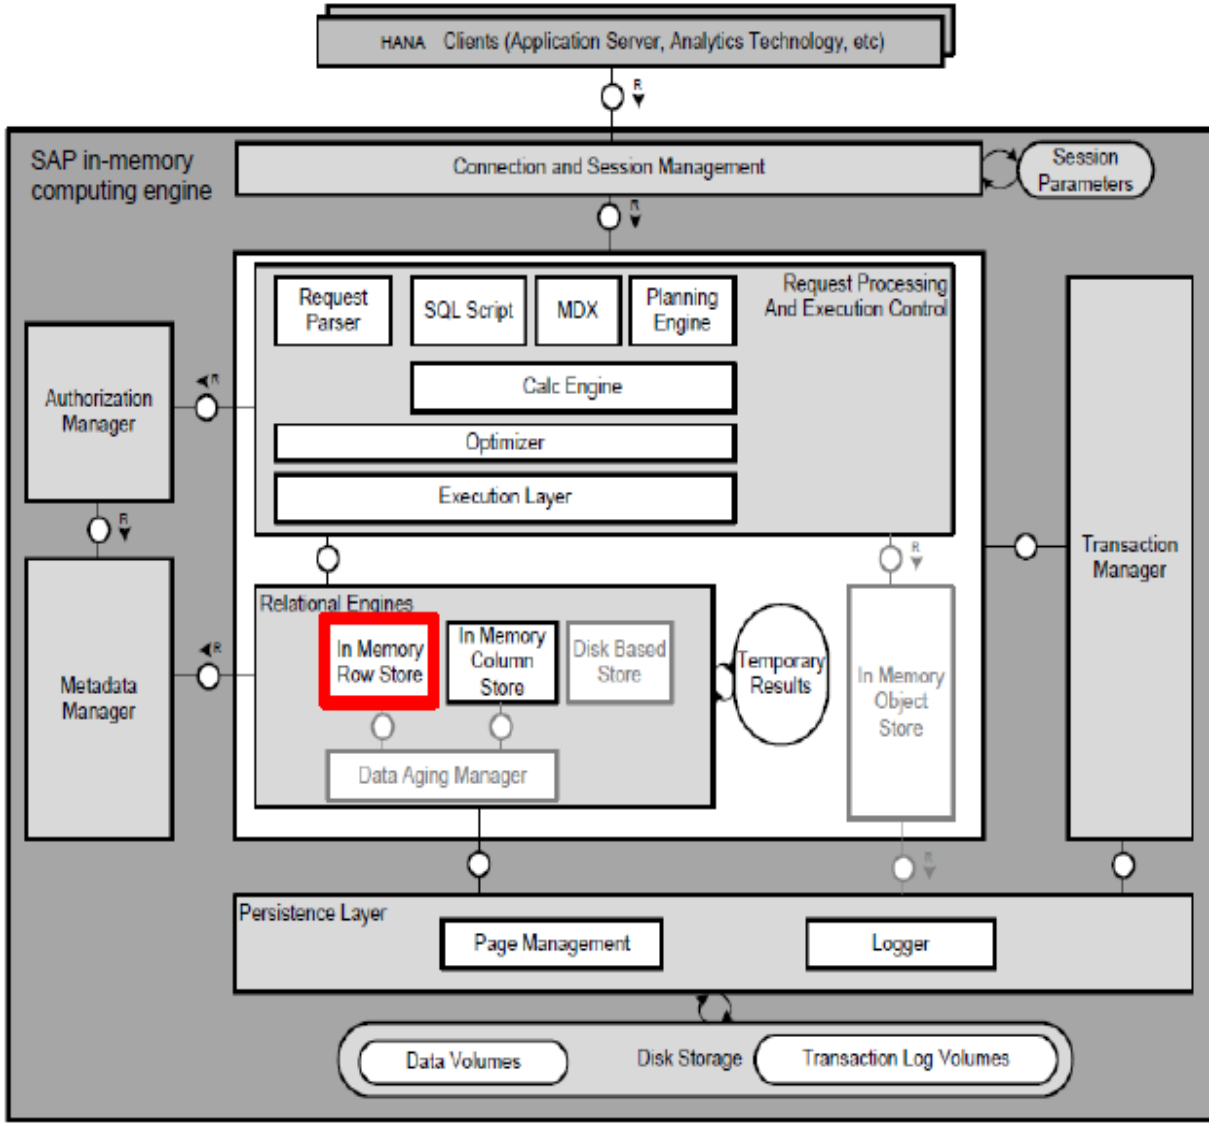

Row Store

- One of the relational engines
- Interfaced from calculation / execution layer
- Pure in-memory store
- Persistence managed in persistence layer

Row Store Architecture

Row Store - Block Diagram

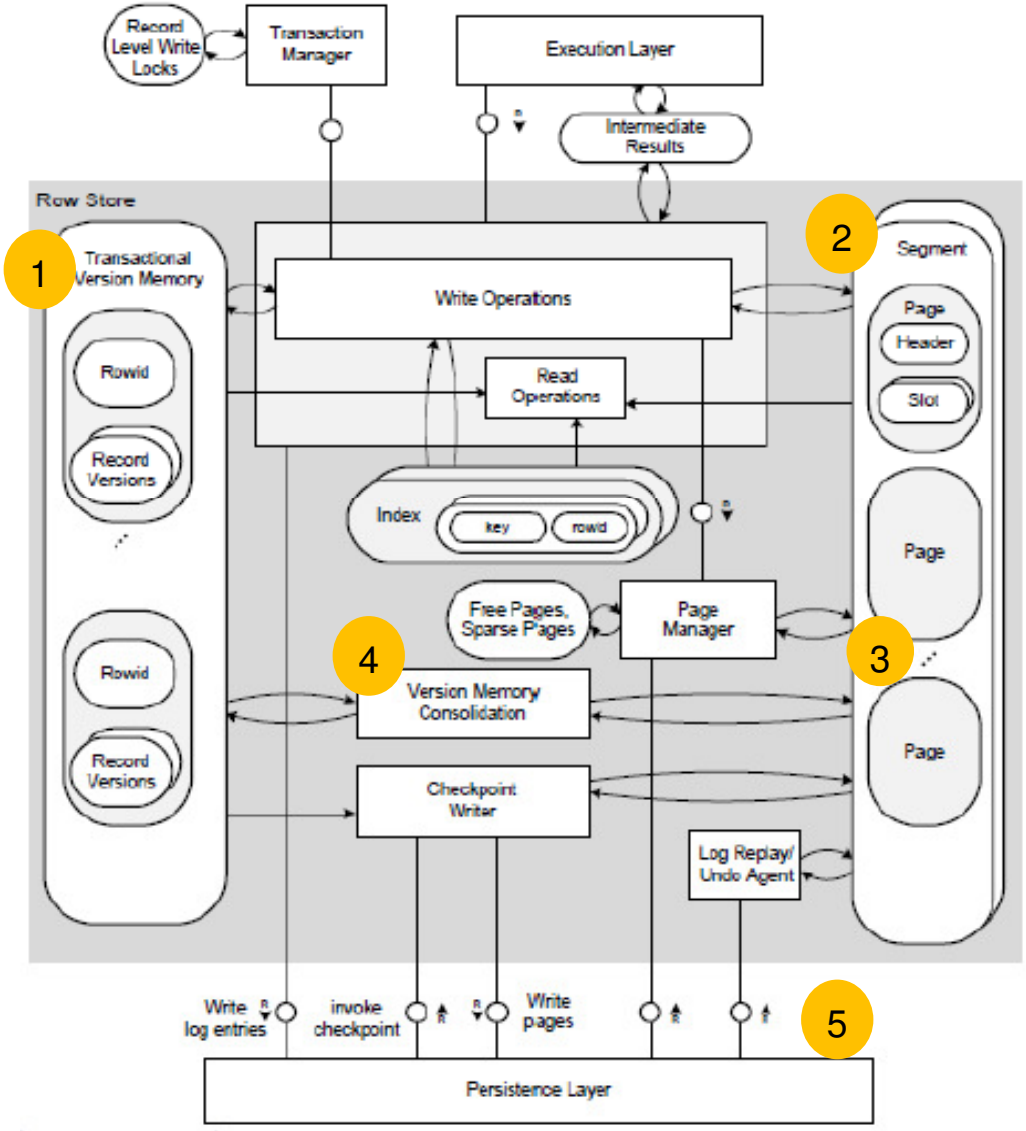

Row Store Block Diagram

1. **Transactional Version Memory**
   - Contains temporary versions
   - Needed for Multi-Version Concurrency Control (MVCC)
2. **Segments**
   - Contain the actual data (content of row-store tables) in pages
3. **Page Manager**
   - Memory allocation
   - Keeping track of free/used pages
4. **Version Memory Consolidation**
   - Think „garbage collector for MVCC“
5. **Persistence Layer**
   - Invoked in write operations (log)
   - And in performing save points checkpoint writer

Row Store Architecture – Highlights

Mahindra Satyam

Write Operations

- Mainly go into “Transactional Version Memory”
- “INSERT” also writes to Persisted Segment

Persisted Segment

- Contains data that may be seen by any ongoing transaction
- Data that has been committed before any active transaction was started)

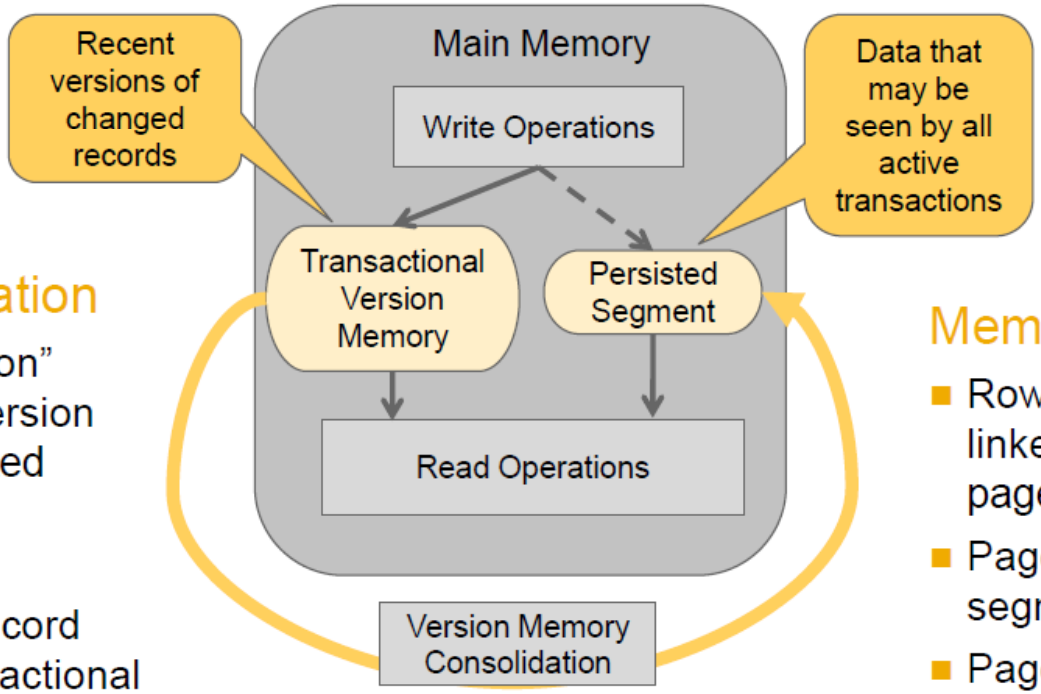

Version Consolidation

- Moves “visible version” from Transaction Version Memory into Persisted Segment (based on Commit ID)
- Clears “outdated” record versions from Transactional Version Memory

Memory Handling

- Row store tables are linked list of memory pages
- Pages are grouped in segments
- Page size: 16 KB

Column Store

Column Store

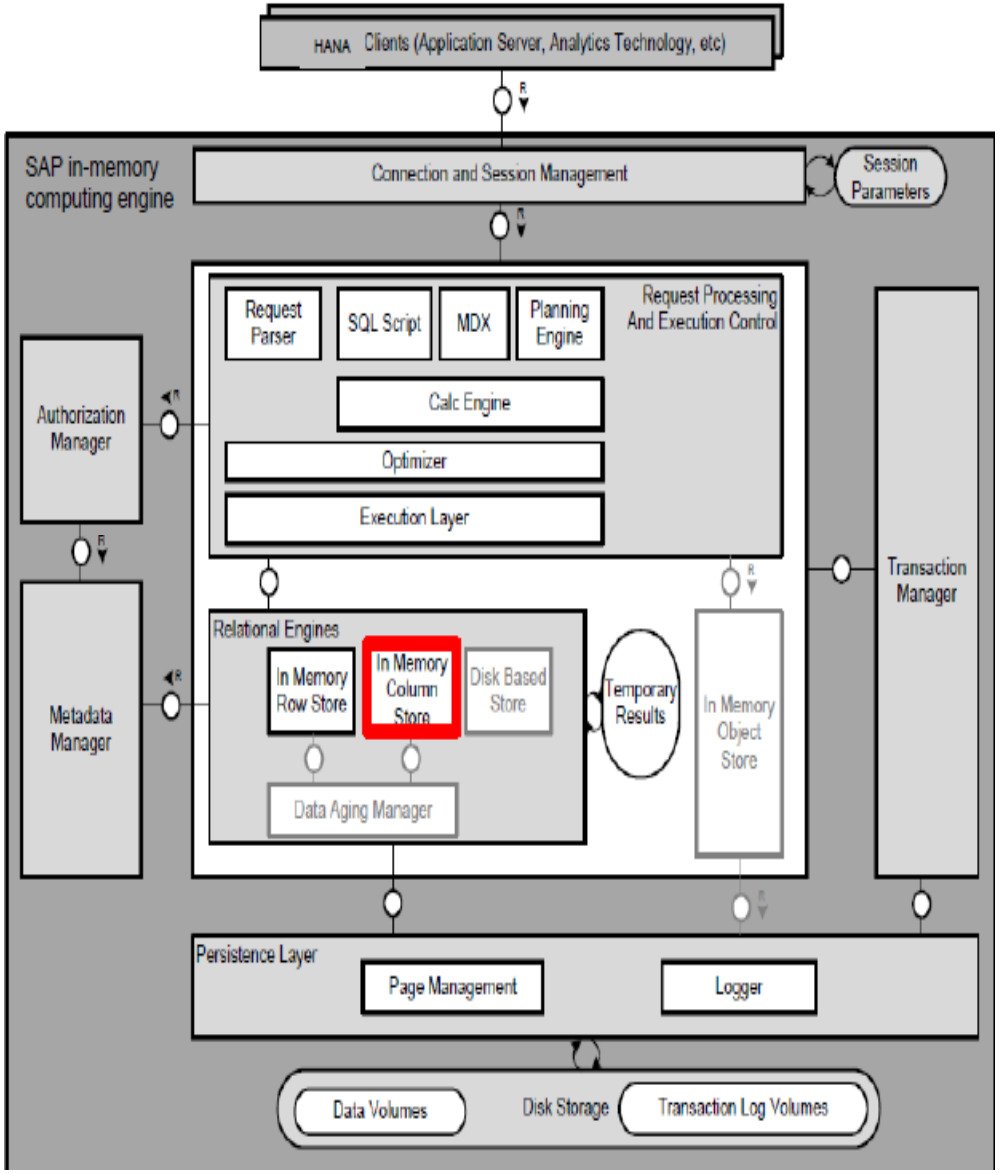

Column Store

- One of the relational engines
- Interfaced from calculation / execution layer
- Pure in-memory store
  - Persistence managed in persistence layer
- Optimized for high performance of **read operation**
- Good performance of **write operations**
- Efficient data **compression**

Column Store Architecture

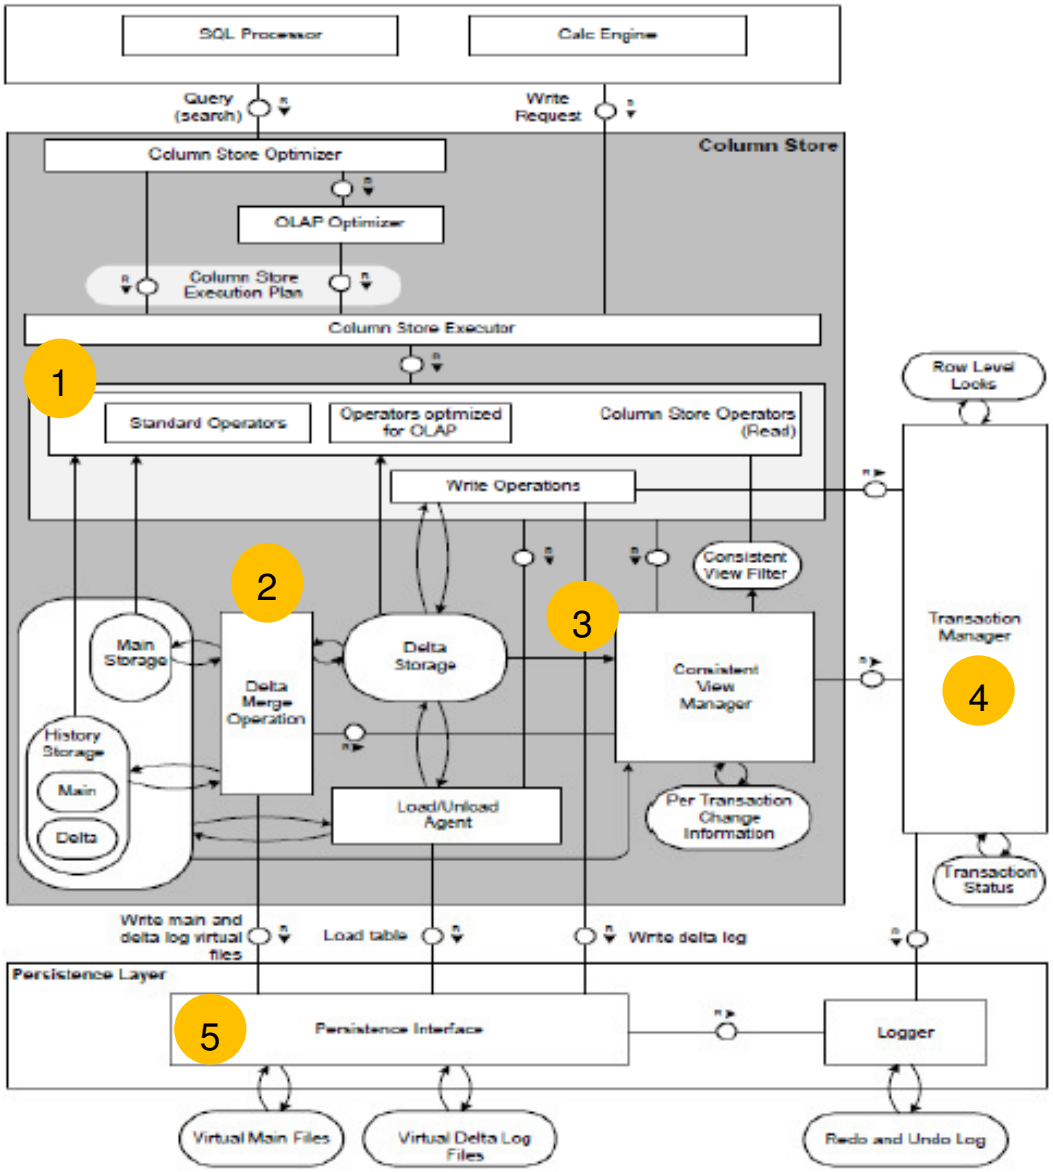

Column Store Block Diagram

Optimizer and Executor

1 Handles queries and execution plan

Main and Delta Storage

2 Compressed data for fast read
Delta data for fast write
Asynchronous delta merge

3 Consistent View Manager

4 Transaction Manager

Persistence Layer

5

Column Store Architecture – Highlights

Mahindra Satyam

Storage Separation (Main & Delta)

- Enables **high compression** and **high write performance** at the same time

Write Operations

- **Only** in **delta storage** because write optimized.
- The update is performed by **inserting** a new entry into the **delta storage.**

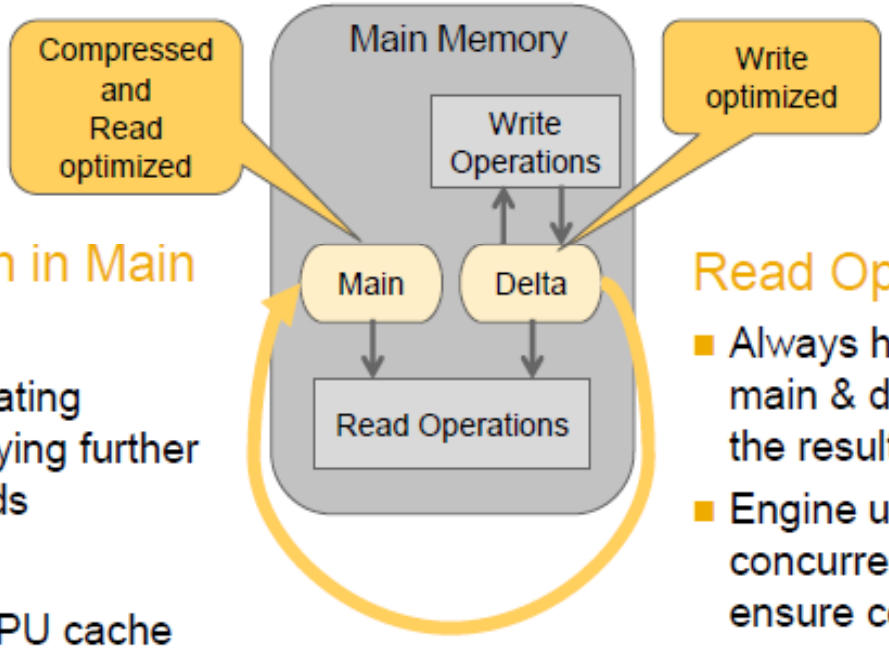

Data Compression in Main Storage

- Compression by creating **dictionary** and applying further compression methods
- Speed up
  - Data load into CPU cache
  - Equality check → Search
- The compression is computed during **delta merge operation**.

Read Operations

- Always have to read from **both** main & delta storages and merge the results.
- Engine uses multi version concurrency control (**MVCC**) to ensure consistent read operations.

Delta Merge Operation

- See next slide

Row Store Vs Column Store

Where to use which Store

Mahindra Satyam

Modeling Only Possible For Column Tables

- This answers the frequently asked question:
  "Where should I put a table – row store or column store?"
  - Information Modeler only works with column tables
  - Replication server creates tables in column store per default
  - Data Services creates tables in column store per default
  - SQL to create column table: "CREATE COLUMN TABLE ..."
  - Store can be changed with "ALTER TABLE …"

System Tables Are Created Where They Fit Best

- Administrative tables in row store:
  - Schema SYS → caches, administrative tables of engine
  - Tables from statistics server
- Administrative tables in column store:
  - Schema _SYS_BI → metadata of created views + master data for MDX
  - Schema _SYS_BIC → some generated tables for MDX
  - Schema _SYS_REPO → e.g. lists of active/modified versions of models

Administration
(Persistence Layer)

Persistence layer

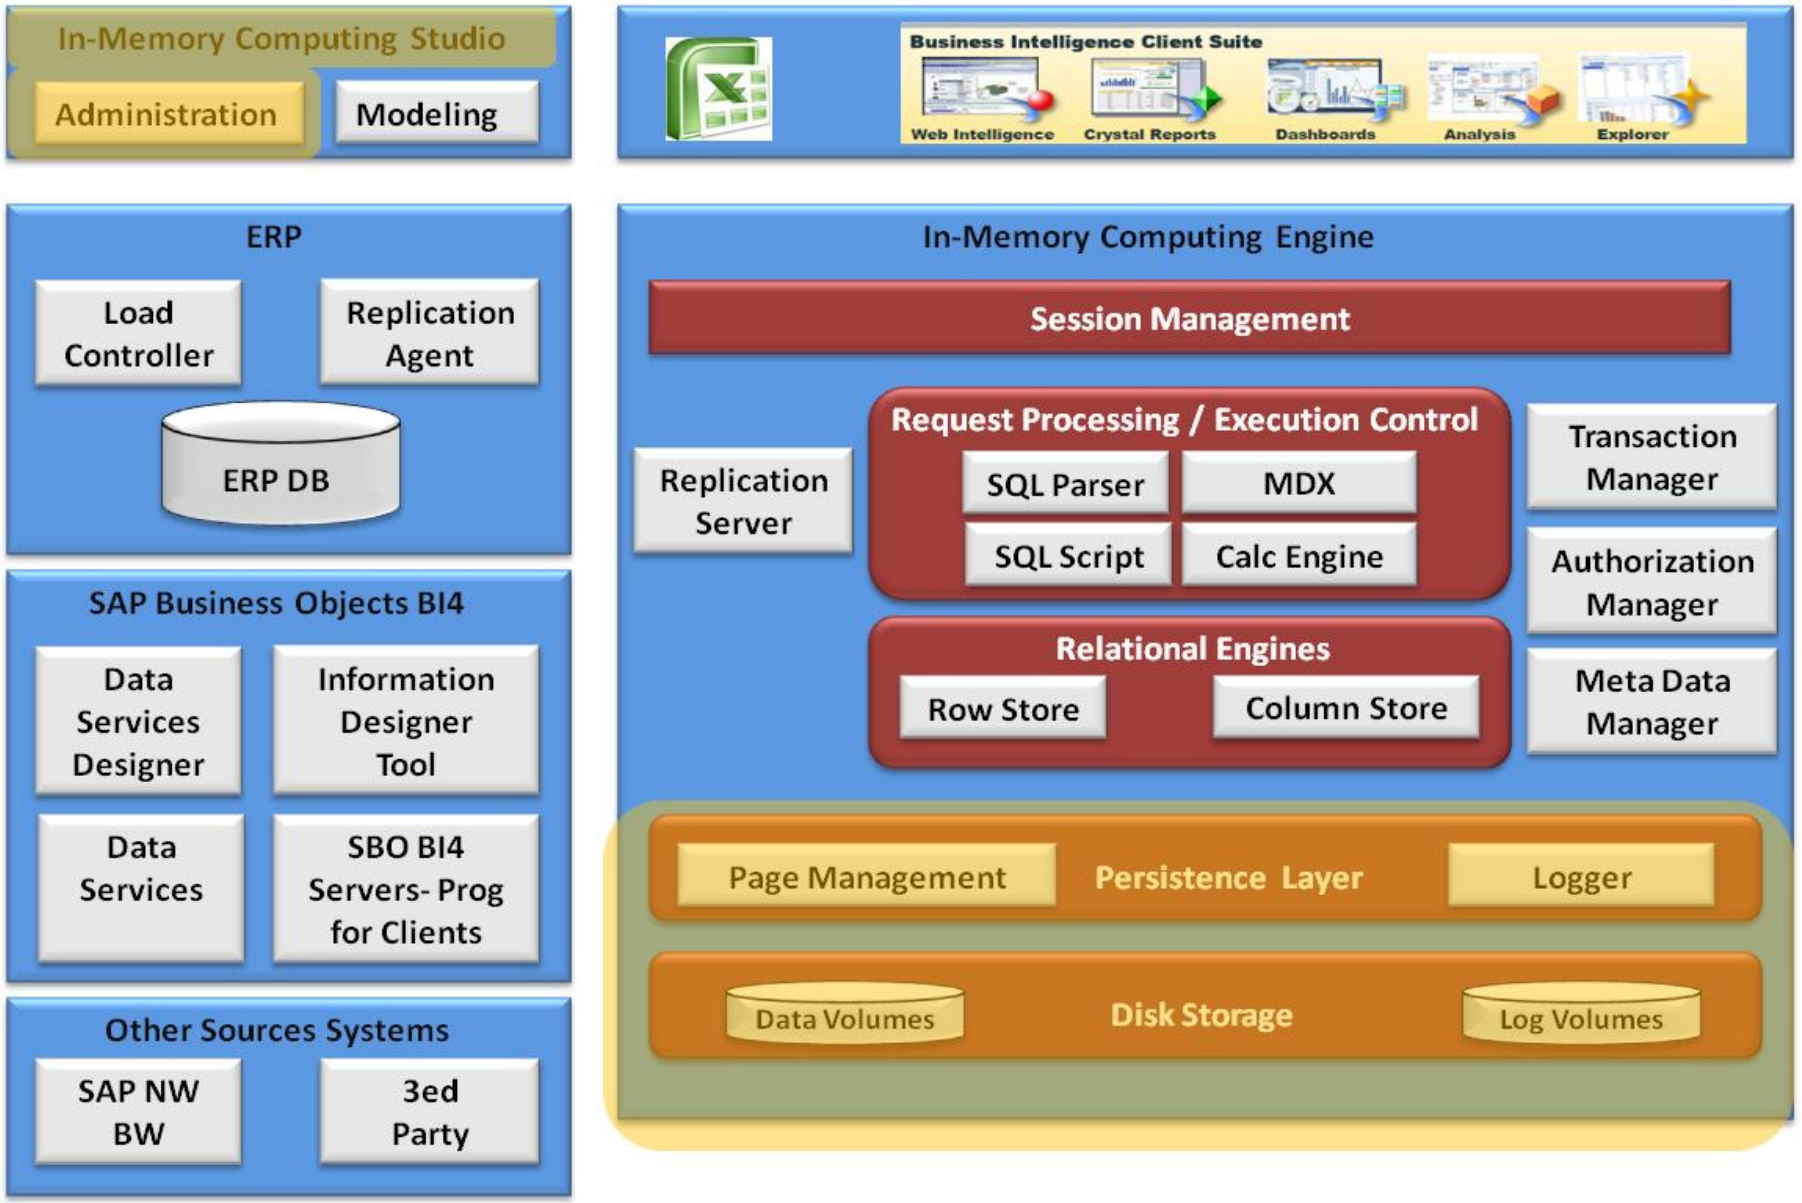

Persistence Layer

Purpose & Scope

Why Does An In-memory Database Need A Persistence Layer?

- Main Memory is volatile. What happens upon…
  - Database restart?
  - Power outage?
  - ...
  - → Data needs to be stored in a non-volatile way
- Backup and restore

- SAP in-memory computing engine offers one persistence layer which is used by row store and column store
  - Regular “savepoints”
    → full persisted image of DB at time of savepoint
  - Logs capturing all DB transactions since last savepoint (redo logs and undo logs written)
    → restore DB from latest savepoint onwards
  - Ability to create "snapshots"
    → used for backups

HANA Roadmap

HANA Road Map

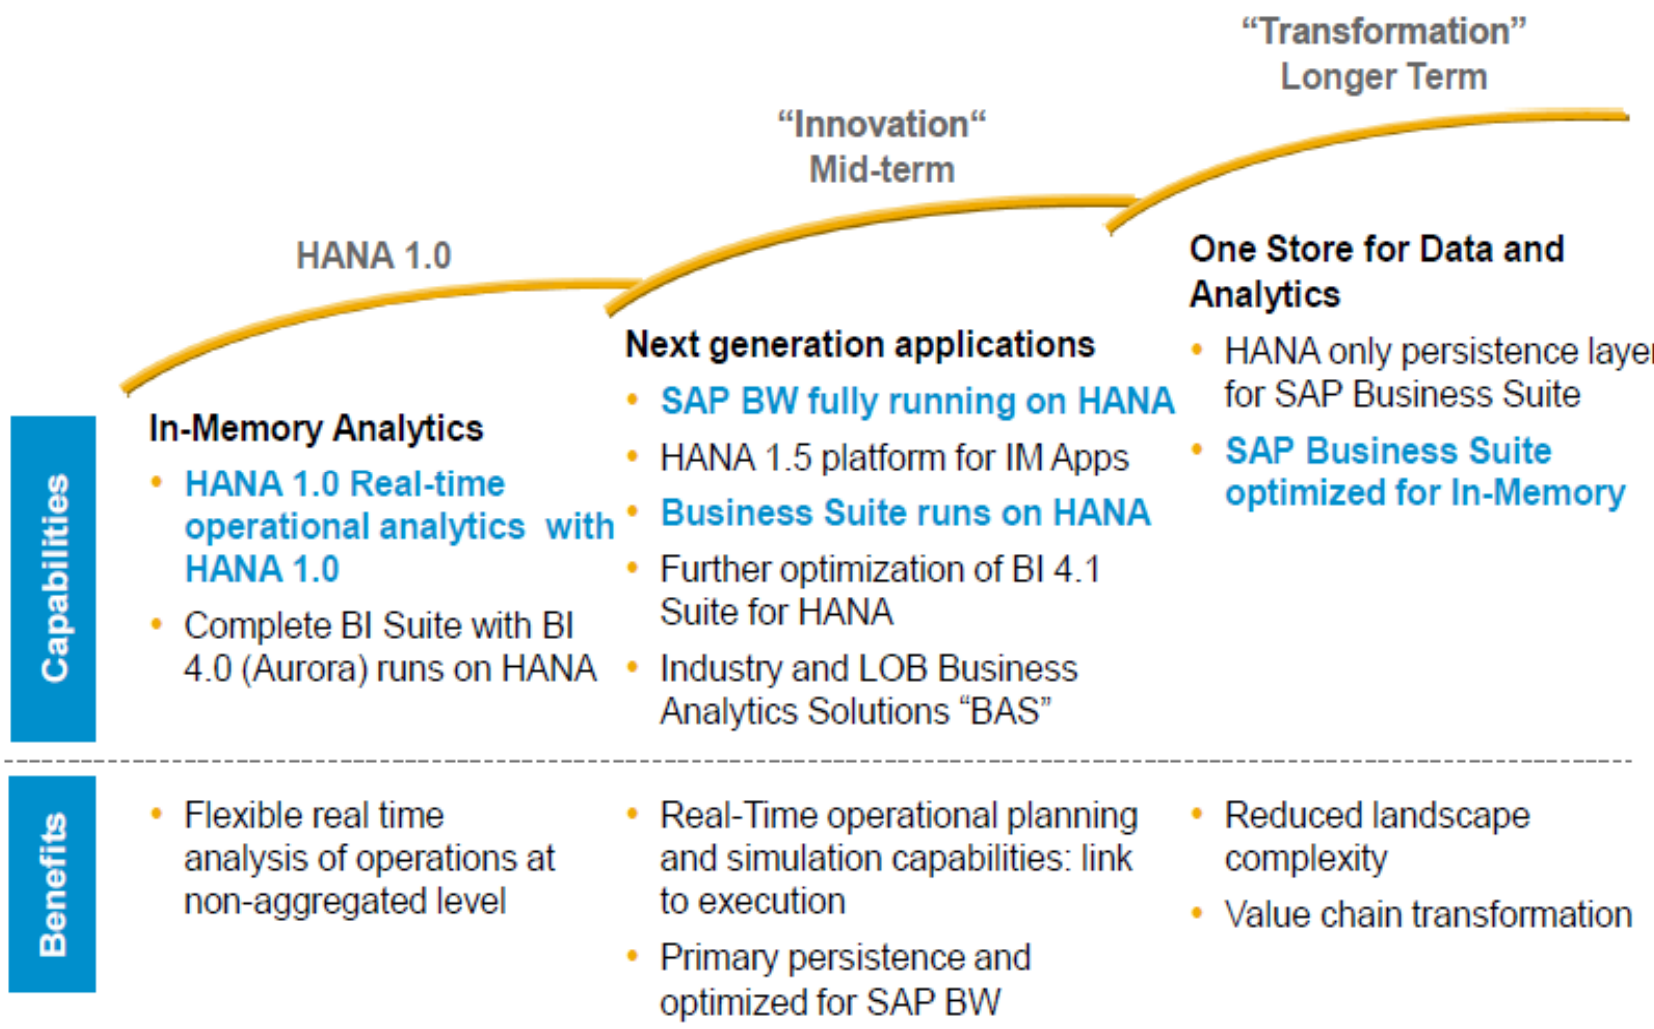

HANA Licensing
&
T-Shirt Sizing

Licensable components

A : Required Licensable components

B : Additional Licensable Components (when available)

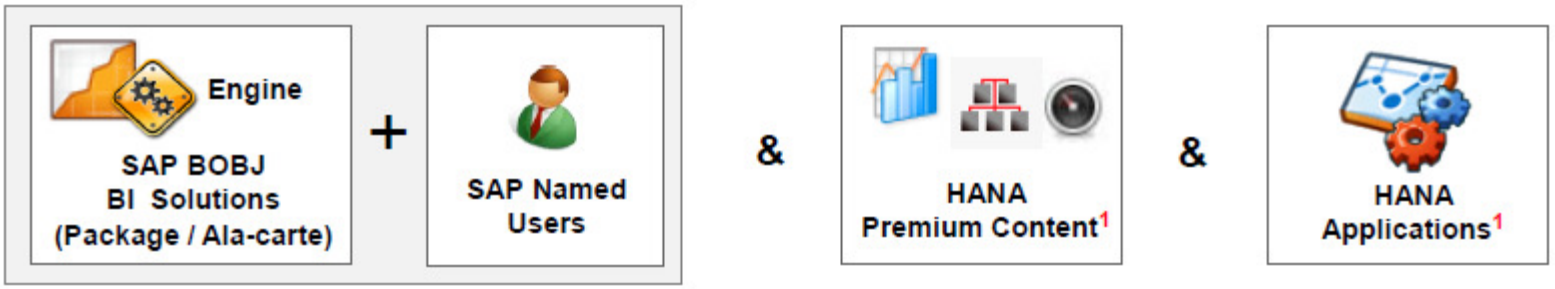

Thank you

mahindrasatyam.com

Safe Harbor

This document contains forward-looking statements within the meaning of section 27A of Securities Act of 1933, as amended, and section 21E of the Securities Exchange Act of 1934, as amended. The forward-looking statements contained herein are subject to certain risks and uncertainties that could cause actual results to differ materially from those reflected in the forward-looking statements. We undertake no duty to update any forward-looking statements. For a discussion of the risks associated with our business, please see the discussions under the heading “Risk Factors” in our report on Form 6-K concerning the quarter ended September 30, 2008, furnished to the Securities and Exchange Commission on 07 November, 2008, and the other reports filed with the Securities and Exchange Commission from time to time. These filings are available at http://www.sec.gov