

SAP HANA Platform

The platform for all applications



Disclaimer

This presentation outlines our general product direction and should not be relied on in making a purchase decision. This presentation is not subject to your license agreement or any other agreement with SAP. SAP has no obligation to pursue any course of business outlined in this presentation or to develop or release any functionality mentioned in this presentation. This presentation and SAP's strategy and possible future developments are subject to change and may be changed by SAP at any time for any reason without notice. This document is provided without a warranty of any kind, either express or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. SAP assumes no responsibility for errors or omissions in this document, except if such damages were caused by SAP intentionally or grossly negligent.

SAP HANA Platform is easy to adopt

Standard-based and open

SAP HANA PLATFORM

APPLICATION SERVICES

PROCESSING SERVICES

INTEGRATION & QUALITY SERVICES

DATABASE SERVICES

Database Services

- Standard RDBMS
- ACID, SQL 92 Compliant
- Accessible thru JDBC, ODBC, JSON, OData
- Standard security model
- Choice of third-party administration tools

Application Services

- Choice of application servers and webservers
- Eclipse-based and web development tool
- Include web application server with Java Script, Java, Node.JS, C++ runtime support
- Support git, github, maven tools
- Include HTML5 UI libraries

Processing Services

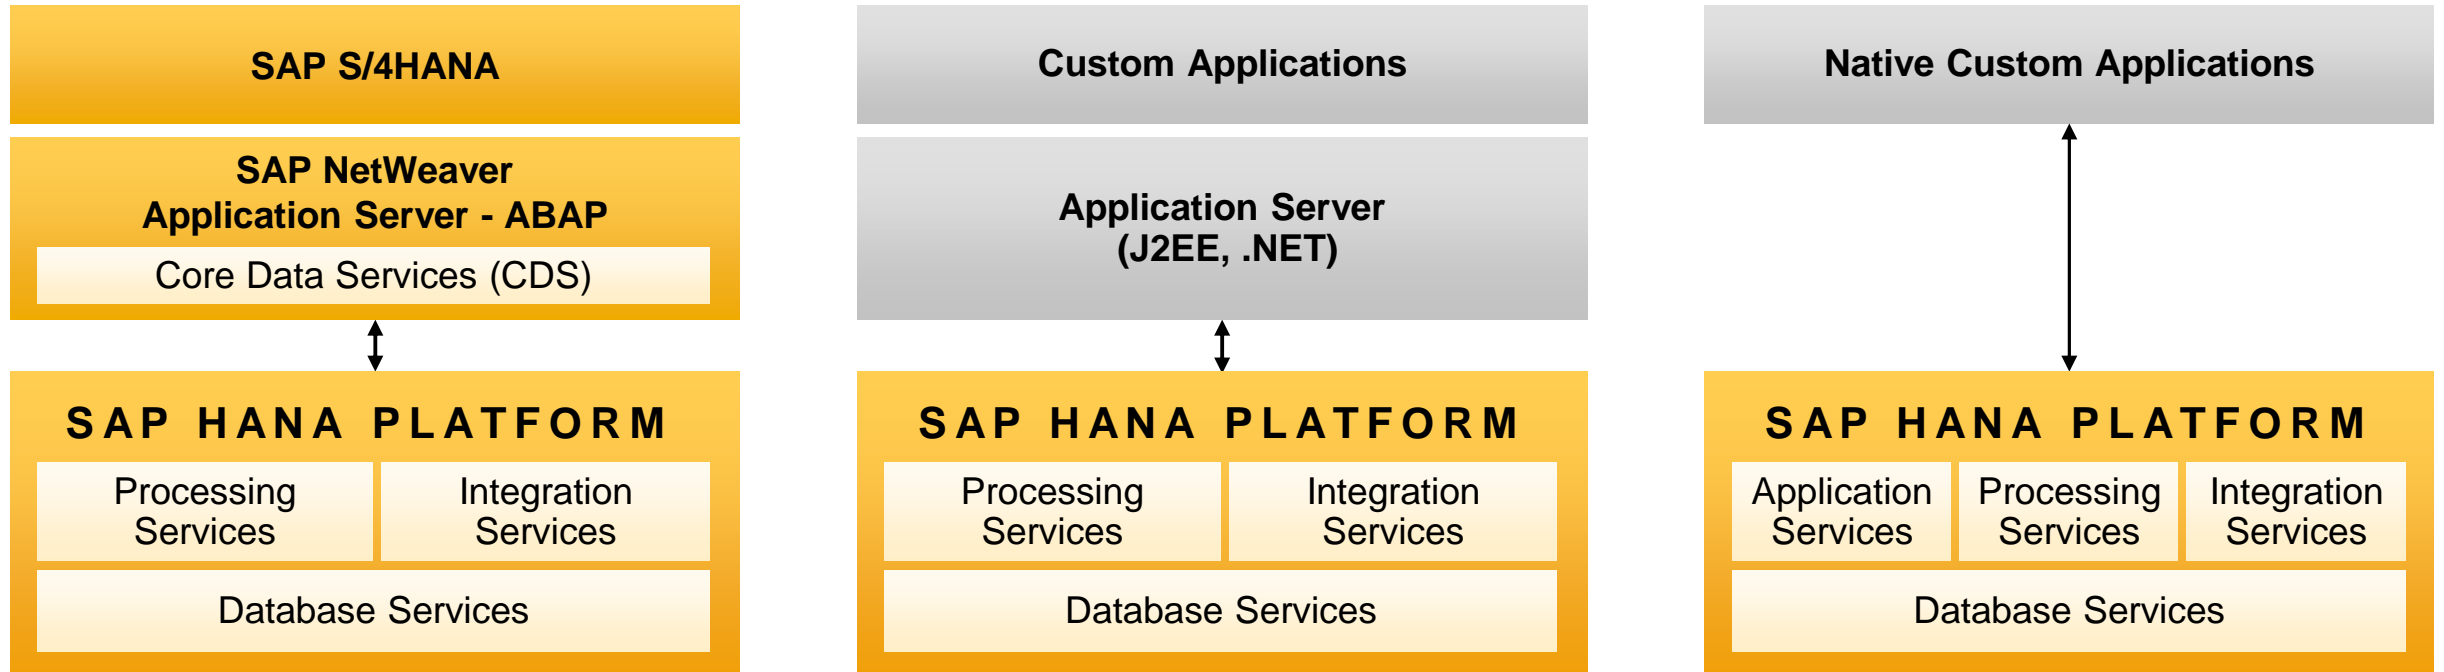
- Execute advanced data processing using SQL
- Spatial processing follows OGC standards, ISO SQL/MM, GeoJSON
- Built-in predictive libraries and supports R

Integration & Quality Services

- Data movement and federation with existing DBs
- Framework to build custom adaptors
- Integration with Spark and Hadoop

Choice of application architecture

Leverage existing assets and skills



- ABAP developers use CDS and Open SQL to leverage SAP HANA without coding SAP HANA objects
- Custom application developers choose any application server and any database interface
- SAP HANA native application developers use SAP HANA application services inside the platform

SAP HANA Platform: The platform for all applications


Simplify, accelerate, innovate

All Devices


SAP, ISV and Custom Applications

SAP HANA PLATFORM


APPLICATION SERVICES



PROCESSING SERVICES



INTEGRATION & QUALITY SERVICES



DATABASE SERVICES



ONE Open Platform

OLTP + OLAP

ONE Copy of the Data

* Graph is in controlled availability

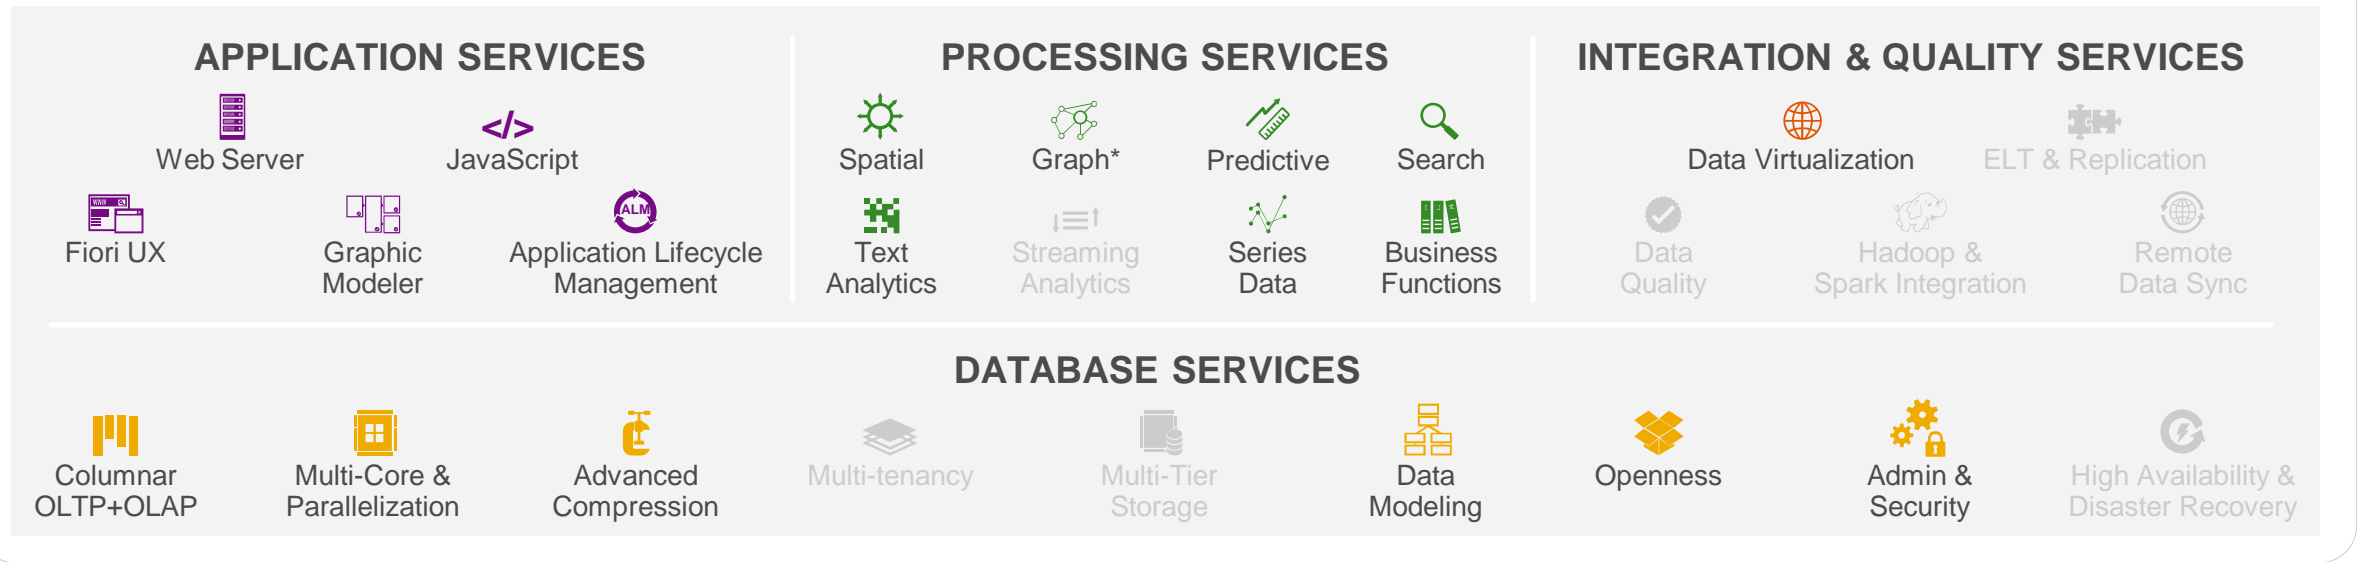
New! SAP HANA Desktop Edition For Developers*

Smaller Footprint – run SAP HANA on a laptop.

Downloadable Virtual Machine Image for SAP HANA.

- Pre-configured SAP HANA – free to download and use for development purposes.
- No need for a certified appliance – can run on a laptop.
- Limitations – 32GB RAM only.
- Community Support via SCN.
- Early Adopter version to be launched at Sapphire.

Product Capabilities



ONE Open Platform

OLTP + OLAP

ONE Copy of the Data

Choice of SAP HANA High Availability (HA) and Disaster Recovery (DR) options


Ensuring the most demanding service-levels

Supports campus, metro, and geo clusters with multiple standbys



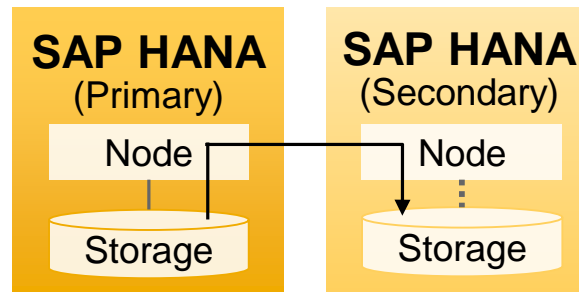
Host Auto-Failover (HA)

- Within one scale-out system
- N active nodes, M standby node(s)
- Automatically switch to standby node



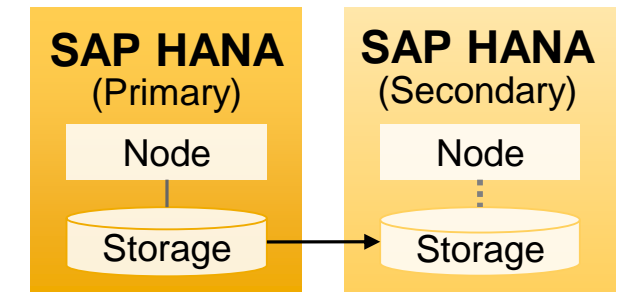
System Replication (HA & DR)

- Across multiple systems/locations
- Continuous data transfer from memory
- Fast switch-over on system failure



Storage Replication (DR)

- Across multiple systems/locations
- Transfer data using storage mirroring
- Low cost option




Secondary system can be used for Dev/QA

HANA in detail

In-memory columnar store

Faster OLTP + OLAP processing on single copy of data

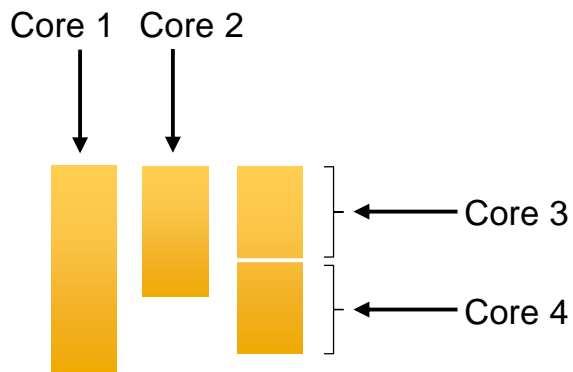
- ACID compliant
- High speed transactions support
- Aggregations on fly
- No indexes for fast access
- Process compressed data
- Optimized for multi-core parallel processing
- Single Instruction, Multiple Data (SIMD) processing support
- NUMA optimization to enable future support for very large (12TB+) nodes (CPU/Memory)



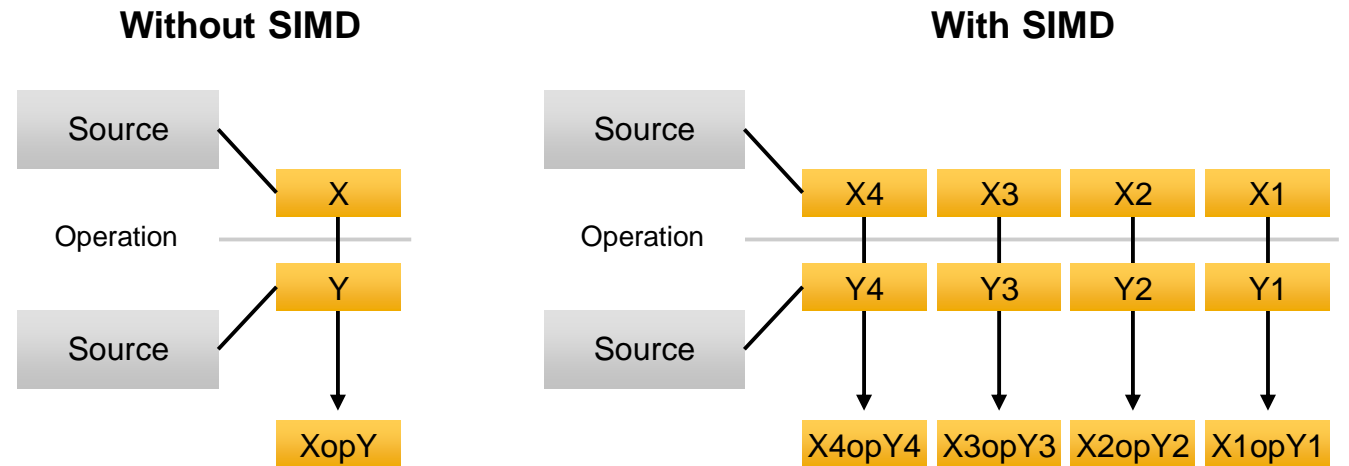
Parallel processing and SIMD

Faster data processing

Each processor can scan a column or portion of a column in parallel



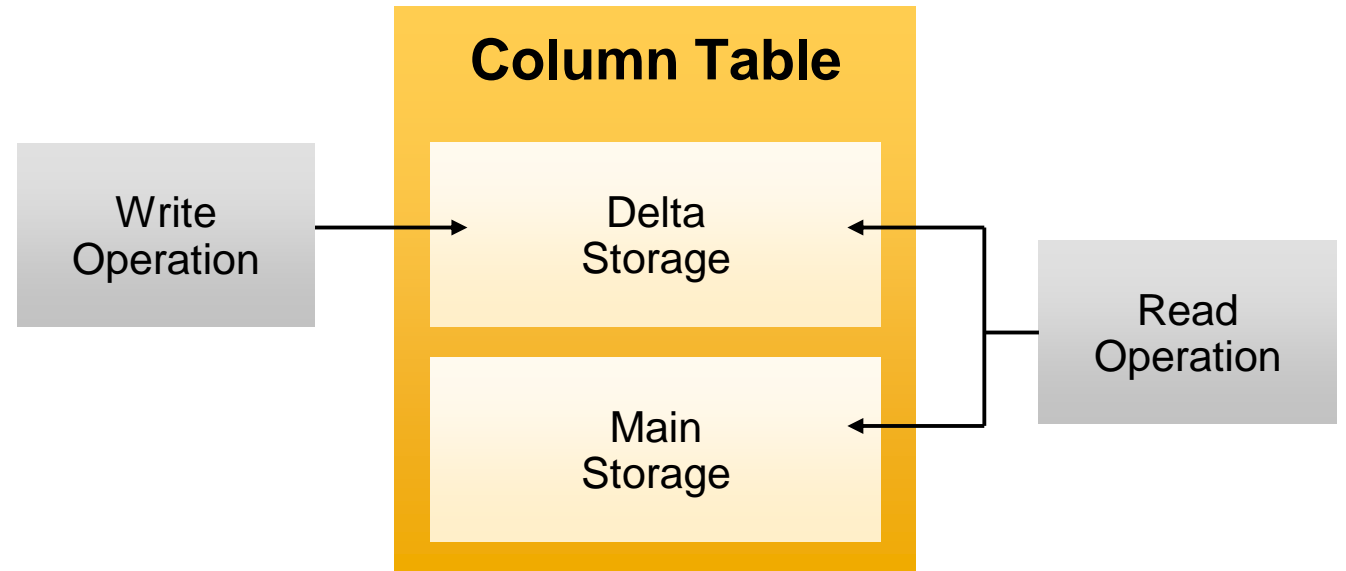
With Single Instruction, Multiple Data (SIMD) processing, the same processor operation works on multiple data blocks simultaneously



Columnar store to process transactions and queries

OLTP+OLAP on single copy of data

- Delta storage is optimized for transactions
- Delta storage is merged periodically with main storage
- No data duplication – data kept either in delta storage or in main storage

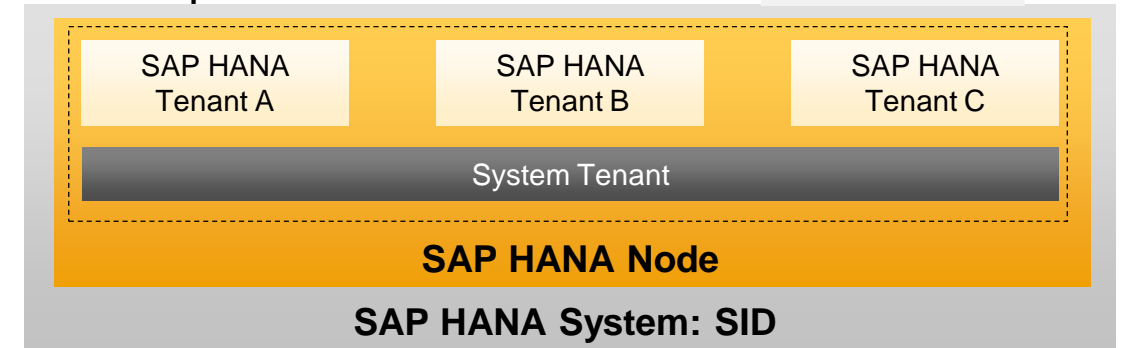


Multitenant database containers

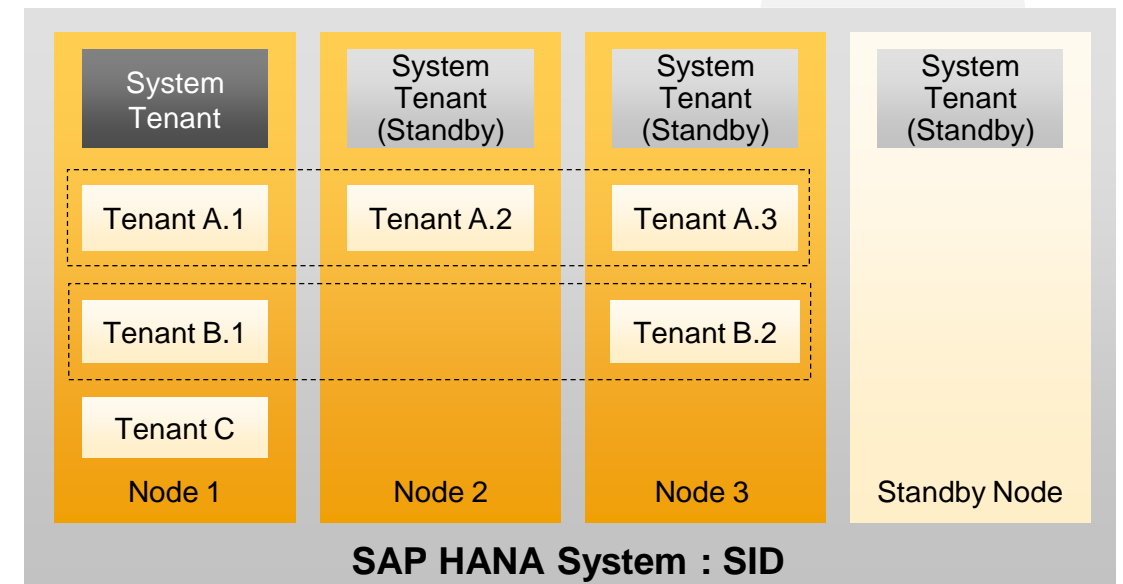
Lower capital and operating expenditure – cloud-ready

- Manage multiple databases as a unit
- Strong separation of data, resources and users among tenant databases
- Lower capital expenditure with better utilization of system resources
- Lower operating expenditure with simplified management

Scale Up



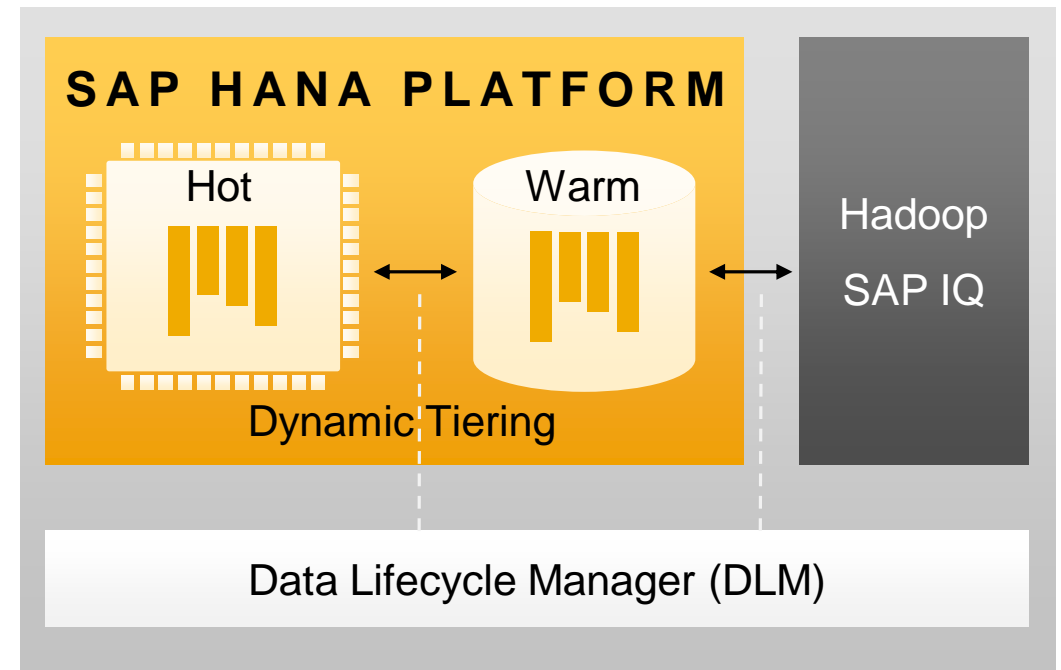
Scale Out



Dynamic tiering

Right price/performance balance between memory and disk

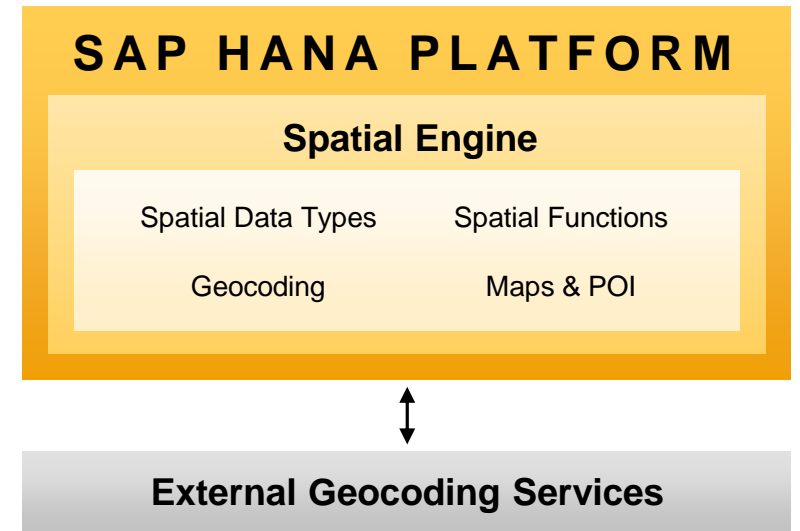
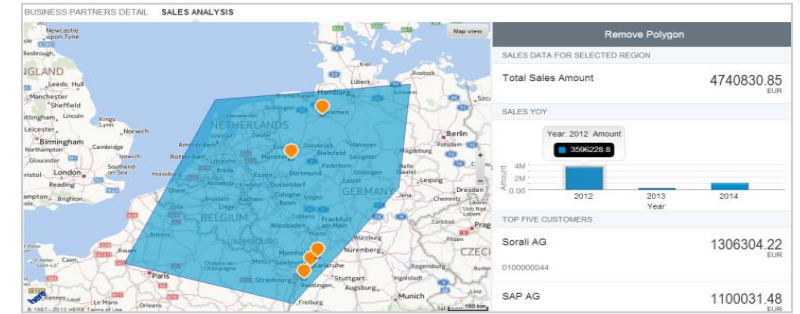
- Utilize disk-based, column-store technology to store less frequently used data
- Support petabyte scale deployment – not confined by the size of memory
- Is integral part of the single SAP HANA instance – no data duplication
- Transparently manage large data volumes by automatically moving data among memory, disk and Hadoop/SAP IQ using Data Lifecycle Manager (DLM)



Spatial processing

New insights from enriching business data with spatial data

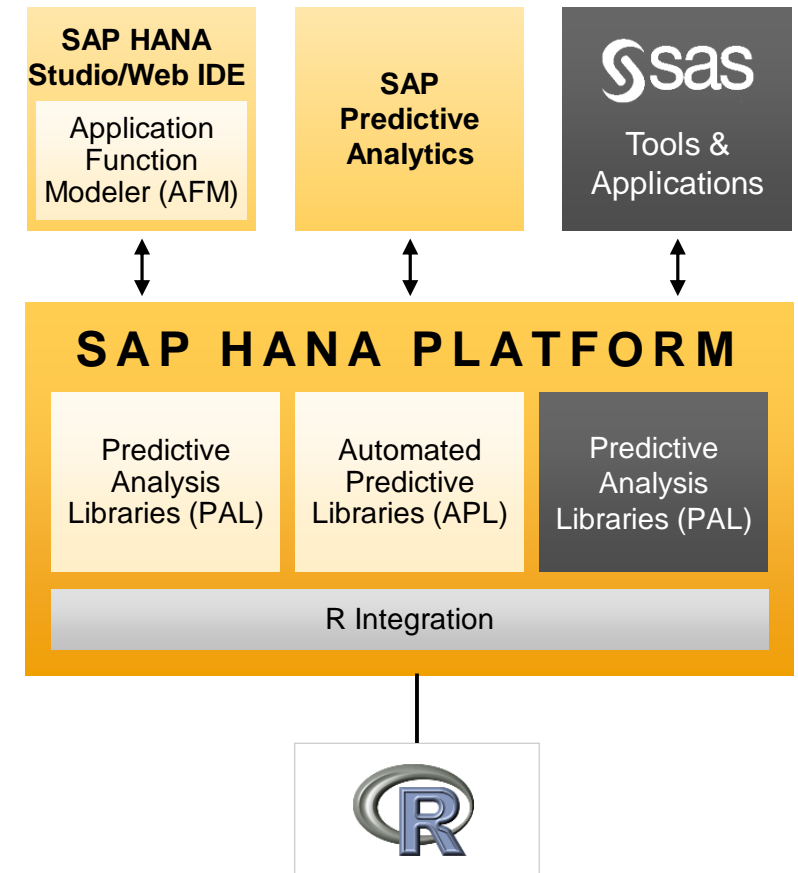
- Store and process spatial data with other data types
- No need to create spatial indexes, tessellation, etc.
- Support open standards compliance (Open Geospatial Consortium - OGC)
- Spatial data types
 - Points, lines, polygons
 - Multi-dimensional support including 3D and measurement dimension
- Spatial functions
 - Area, distance, within, touches, intersects, adjacent
- Built in geo-content
 - Maps, political boundaries, roads, Point of Interests (POI)
- Spatial join operators for SQL and Calculation Views
 - Contains, crosses, intersects, overlaps, touches, within
 - Automatically add latitude and longitude to address



Predictive analytics

Transforming the future with today's insights

- 70+ prepackaged predictive algorithms
 - Supports association, clustering, classification, regression, time series etc.
 - Supports variety of data – structured, spatial, text, streaming and series data
- AFM graphical modeling tool support to develop predictive applications using PAL and R-Script
- SAP Predictive Analytics leverage Automated Predictive Libraries (APL) libraries and PAL
- SAS processing run natively in SAP HANA eliminating data redundancy
- Leverage R advanced functions transparently
- Predictive analytics across multiple data types



SAP HANA In-Memory Predictive Analytics

Predictive Analysis Library (PAL) - algorithms supported

Association Analysis

- Apriori
- Apriori Lite
- FP-Growth
- KORD – Top K Rule Discovery

Classification Analysis

- CART
- C4.5 Decision Tree Analysis
- CHAID Decision Tree Analysis
- K Nearest Neighbour
- Logistic Regression
- Back-Propagation (Neural Network)
- Naïve Bayes
- Support Vector Machine
- Confusion Matrix
- Parameter Selection & Model Evaluation

Regression

- Multiple Linear Regression
- Polynomial Regression
- Exponential Regression
- Bi-Variate Geometric Regression
- Bi-Variate Logarithmic Regression

Cluster Analysis

- ABC Classification
- DBSCAN
- K-Means
- K-Medoid Clustering
- K-Medians
- Kohonen Self Organized Maps
- Agglomerate Hierarchical
- Affinity Propagation
- Gaussian Mixture Model
- Latent Dirichlet Allocation (LDA)

Time Series Analysis

- Single Exponential Smoothing
- Double Exponential Smoothing
- Triple Exponential Smoothing
- Forecast Smoothing
- ARIMA/Seasonal Arima
- Brown Exponential Smoothing
- Croston Method
- Forecast Accuracy Measure
- Linear Regression with Damped Trend and Seasonal Adjust
- Test for White Noise, Trend, Seasonality

Probability Distribution

- Distribution Fit
- Cumulative Distribution Function
- Quantile Function

Outlier Detection

- Inter-Quartile Range Test (Tukey's Test)
- Variance Test
- Anomaly Detection
- Grubbs Outlier Test

Link Prediction

- Common Neighbors
- Jaccard's Coefficient
- Adamic/Adar
- Katz β

Data Preparation

- Sampling
 - Random Distribution Sampling
- Binning
- Scaling
- Partitioning
- Principal Component Analysis (PCA)

Statistic Functions (Univariate)

- Mean, Median, Variance, Standard Deviation
- Kurtosis
- Skewness

Statistic Functions (Multivariate)

- Covariance Matrix
- Pearson Correlations Matrix
- Chi-squared Tests:
 - Test of Quality of Fit
 - Test of Independence
- F-test (variance equal test)

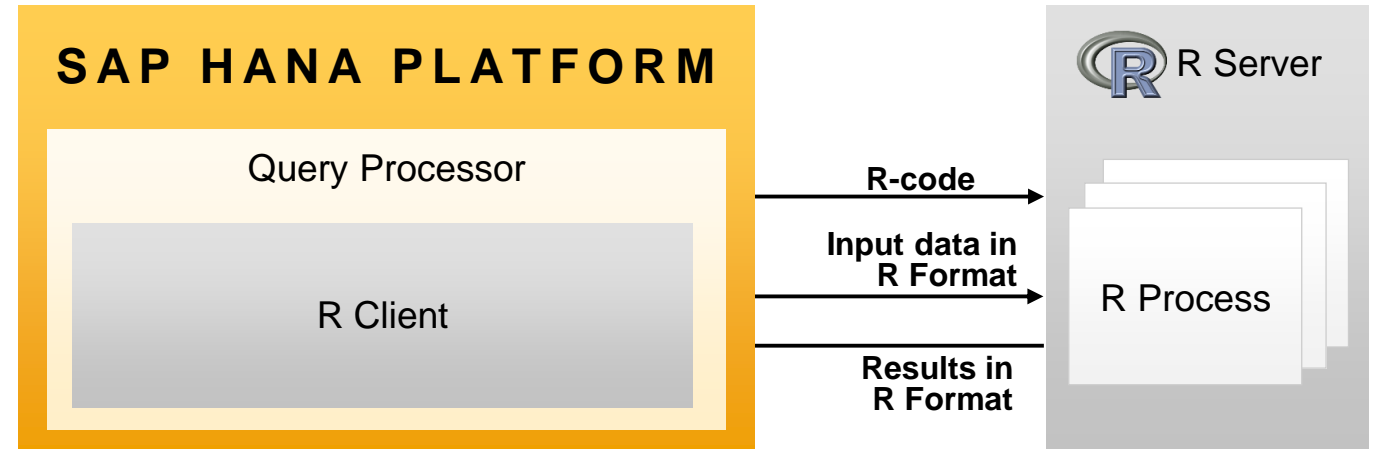
Other

- Weighted Scores Table
- Substitute Missing Values
- DenStream clustering
- Adaptive Hoeffding Tree
- Random forest
- Area Under Curve calculation
- Survival analysis (Kaplan-Meier estimate)

R integration

Use leading open source data mining software transparently

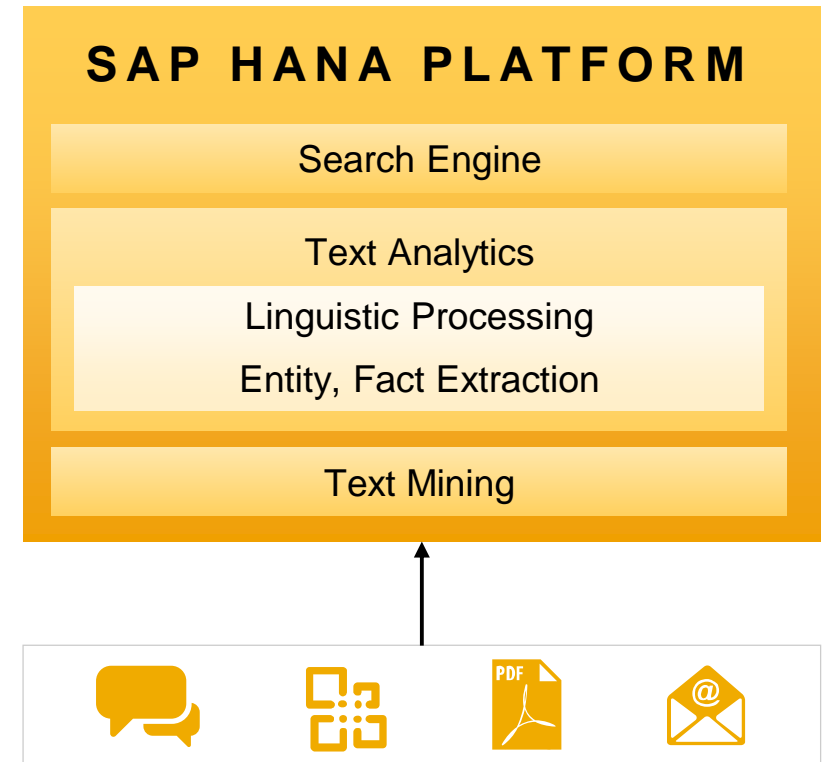
- Embed R script within SQL script
- Execute R script inside R server
- Use R vector-oriented format rather than JDBC/ODBC
- Execute multiple R processes in parallel
- Leverage 3,500+ R statistical and graphical packages



Search, text analytics and mining

Insights from unstructured data

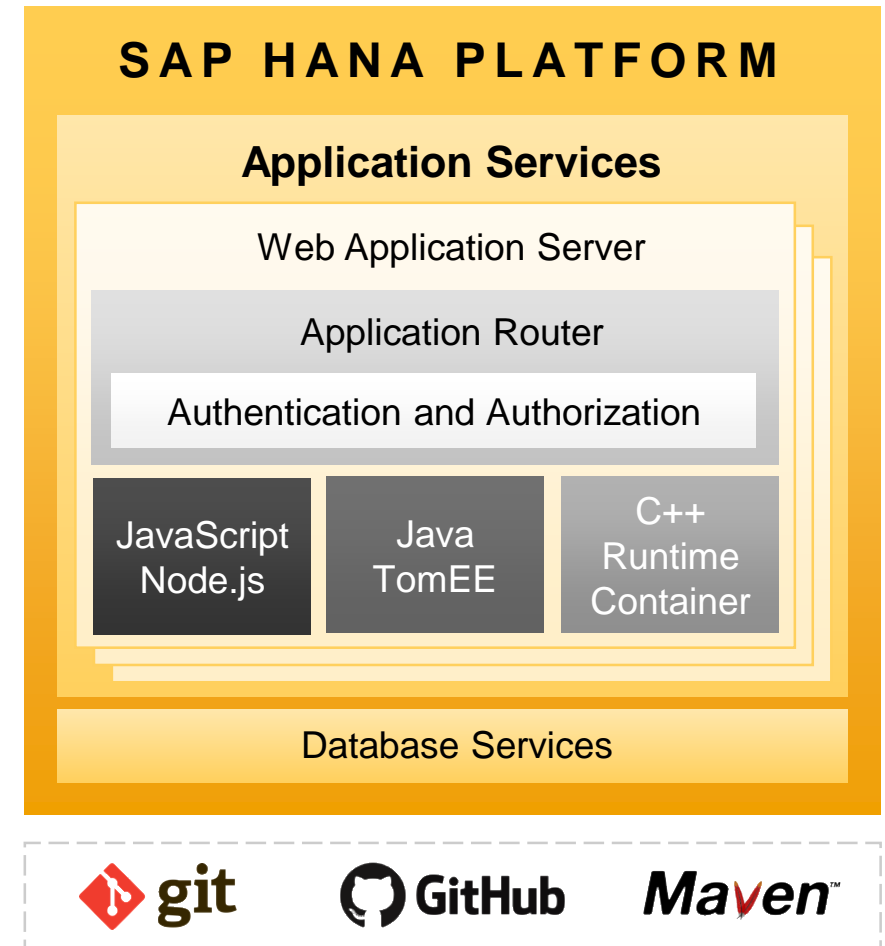
- Store text and binary files in SAP HANA for native text analysis and search
- Support various file formats (txt, html, xml, pdf, doc, ppt, xls, rtf, msg)
- Automatically detects 31 languages
- Search
 - Fuzzy, linguistic, synonymous search, using SQL
- Text Analytics
 - Extract relevant information from text (Linguistic Markup, Entity, Sentiment Extraction)
- Text Mining
 - Rank and Categorize documents by comparing with a set of pre-classified documents



Web server

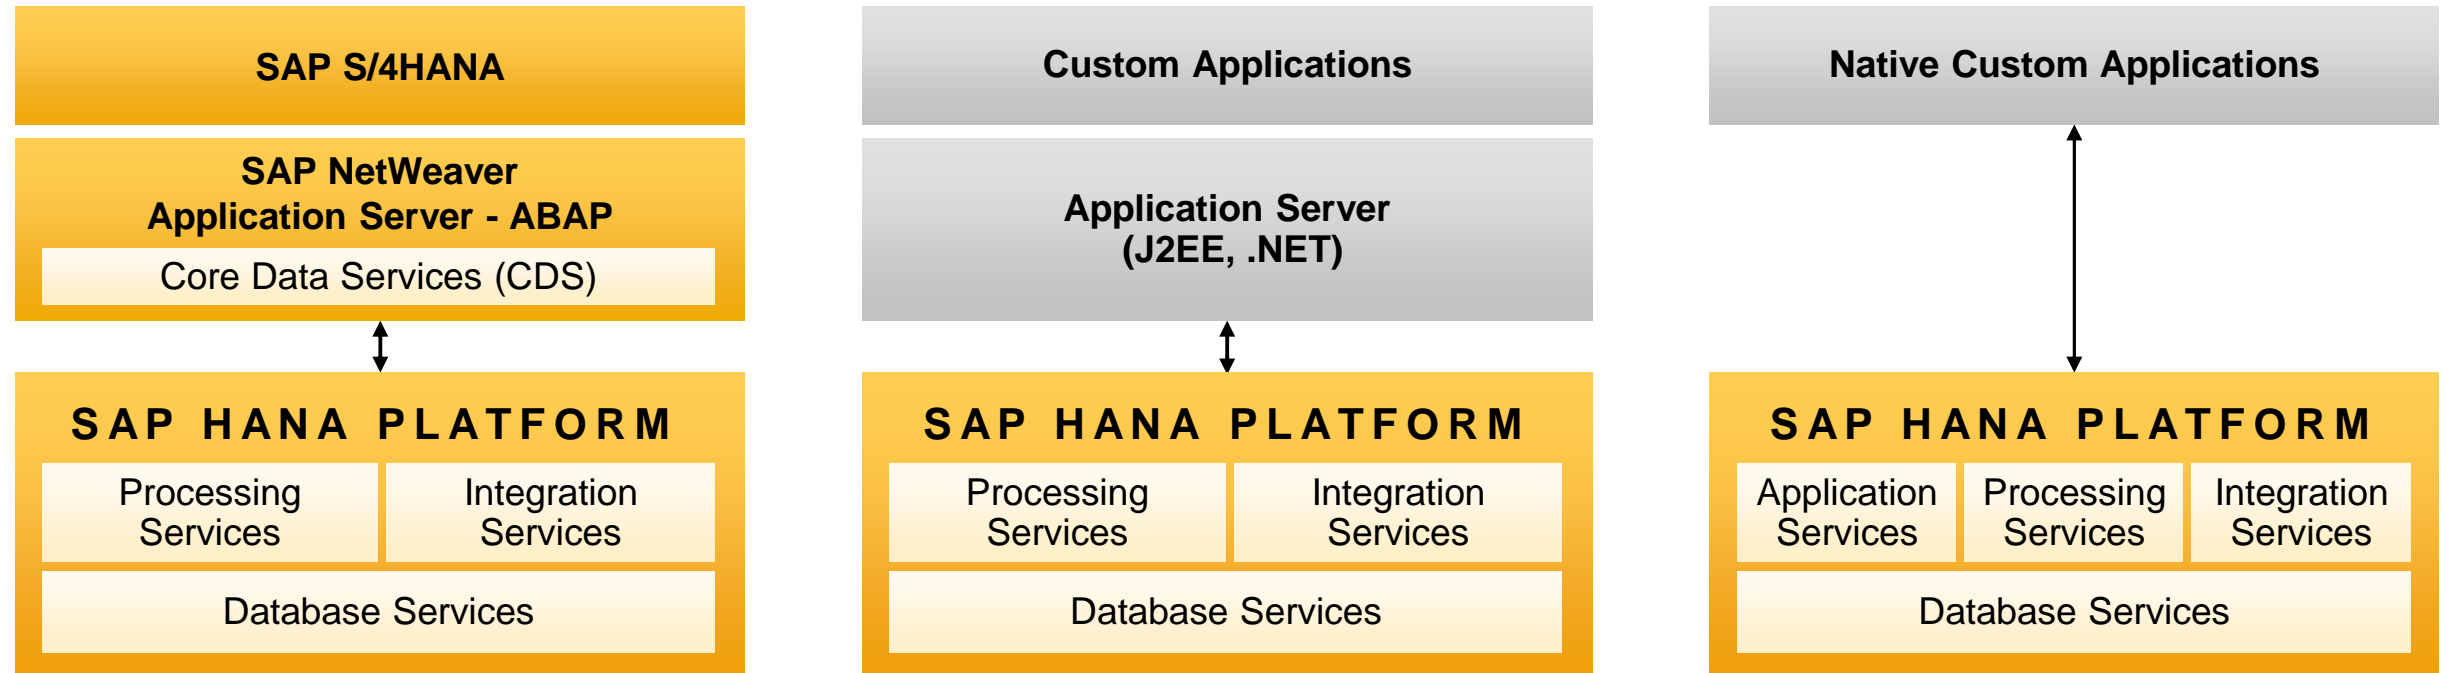
Reduced data movement – app and database services in one platform

- Scale applications independently from the database services with new web application server
- Supports choice of programming languages – Server side JavaScript on Node.js, Java on TomEE and C++ Runtime container
- Core Data Services allow developers to create database objects and relationships without SQL
- Accelerate application development with open source code management tools – Git, GitHub and Maven
- Simplify authentication and authorization with single sign-on support between application and database services
- Scheduled execution of JavaScript and SQLScript programs



Choice of application architecture

Leverage existing assets and skills

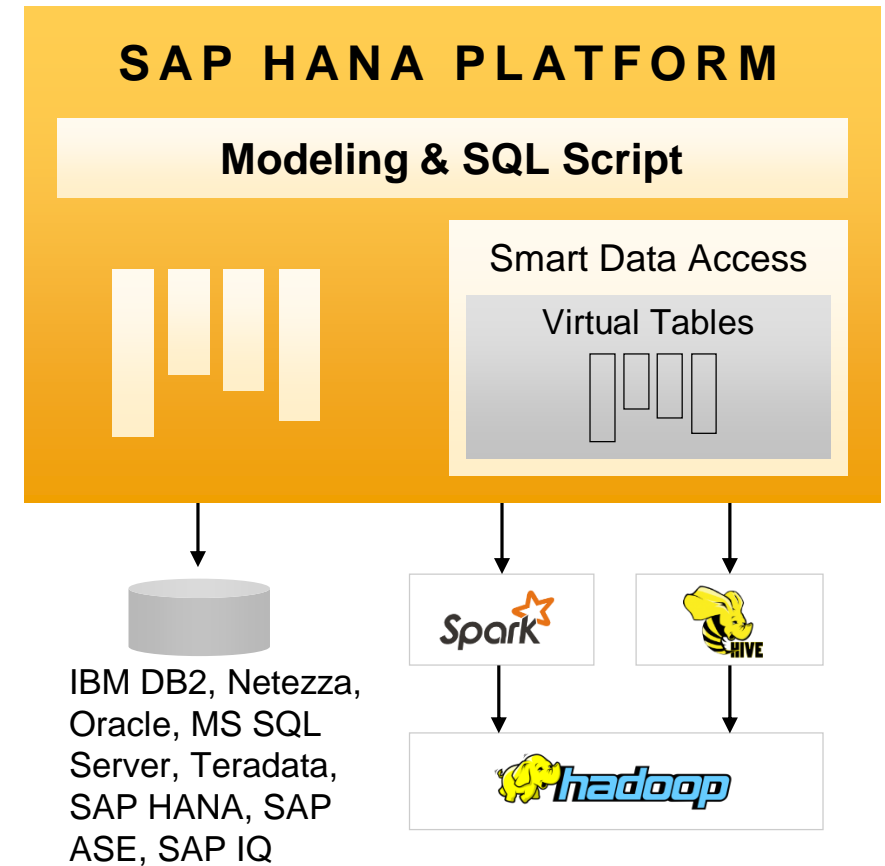


- ABAP developers use CDS and Open SQL to leverage SAP HANA without coding SAP HANA objects
- Custom application developers choose any application server and any database interface
- SAP HANA native application developers use SAP HANA application services inside the platform

Smart data access

Access any data from any source

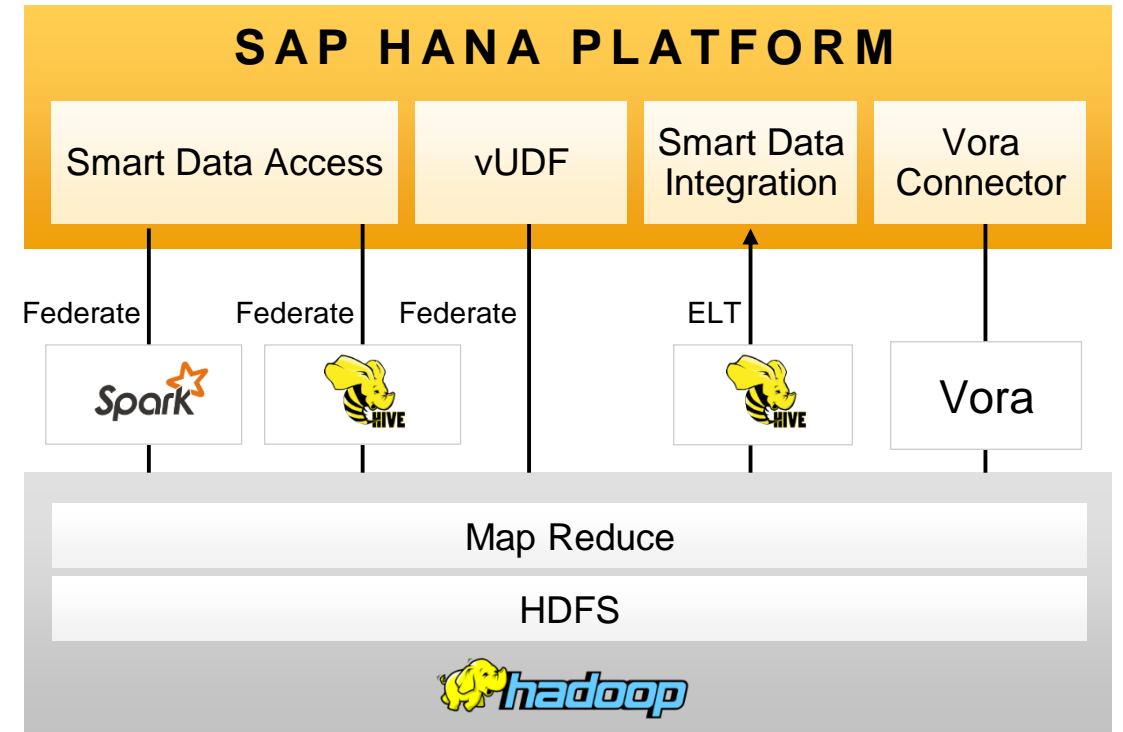
- Manage and query remote tables as local virtual tables
 - Support virtual tables in calculation view and SQL
 - Virtual tables can be combined with PAL, BFL, and Spatial
- Push query processing to remote databases
- Complement functionalities in remote database with SAP HANA capabilities
- Support remote query results caching with HIVE
- Provide SDK for adapters based on ODBC



Hadoop integration

Ad-hoc query capabilities and processing of unstructured data

- Indirect access using Spark and Hive with Smart Data Access
- Direct access using Virtual User Defined Function (vUDF)
 - Access HDFS without need for the package, mapper, and reducer specification
 - Invoke custom Map Reduce jobs
 - Embed vUDF in SQL
- Load data from Hadoop with Smart Data Integration
- Unified admin and monitoring tool for SAP HANA and Hadoop cluster
- Speed-up Hadoop data analysis with new SAP HANA Vora connector



SAP HANA Cockpit




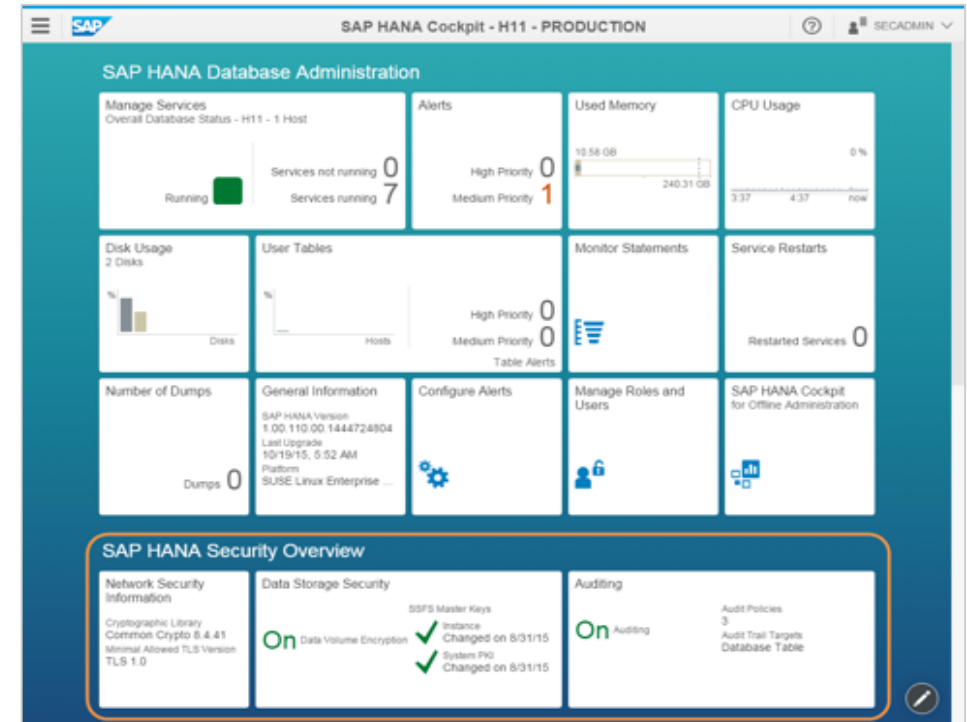
Simplify administration and monitoring

Fiori UX-based web administration tool manages SAP HANA from any device

SAP HANA Cockpit

- Catalog of Fiori tiles to manage hardware resource utilization and SAP HANA processes
- Analyze diagnostic files while the database is down for faster fault detection and correction
- Security dashboard in SAP HANA Cockpit to achieve visibility into security KPIs
- Integrated delta backup capabilities in SAP HANA Cockpit

SAP HANA Cockpit



Hana Express - DEMO

Demo

