

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

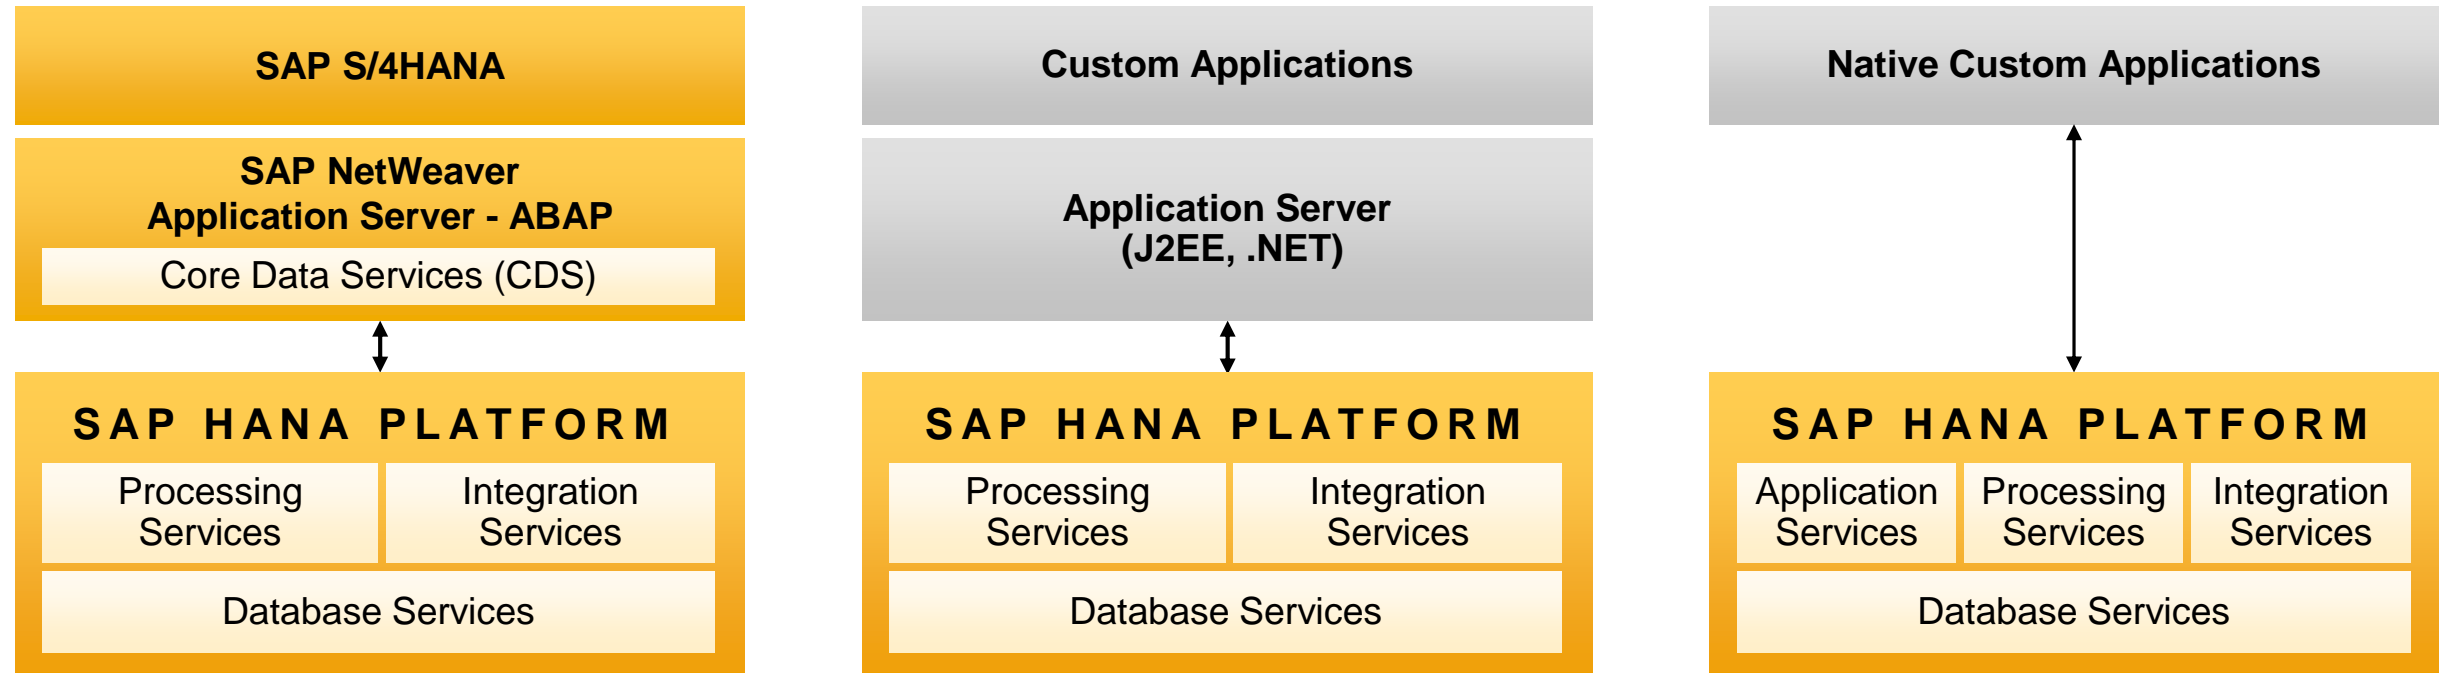
- Execute advanced data processing using SQL
- Spatial processing follows OGC standards, ISO SQL/MM, GeoJSON
- Built-in predictive libraries and supports R

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

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)



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



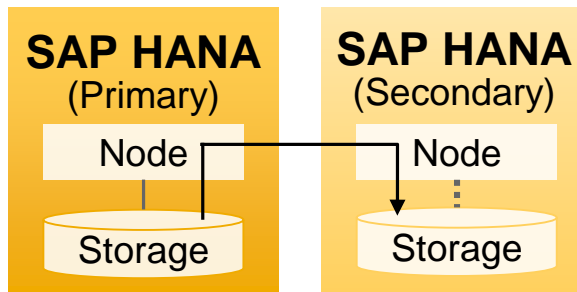
Host Auto-Failover (HA)

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



System Replication (HA & DR)

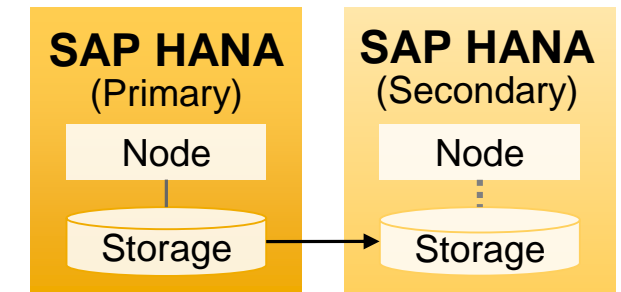
- Across multiple systems/locations
- Continuous data transfer from memory
- Fast switch-over on system failure



Secondary system can be used for Dev/QA

Storage Replication (DR)

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

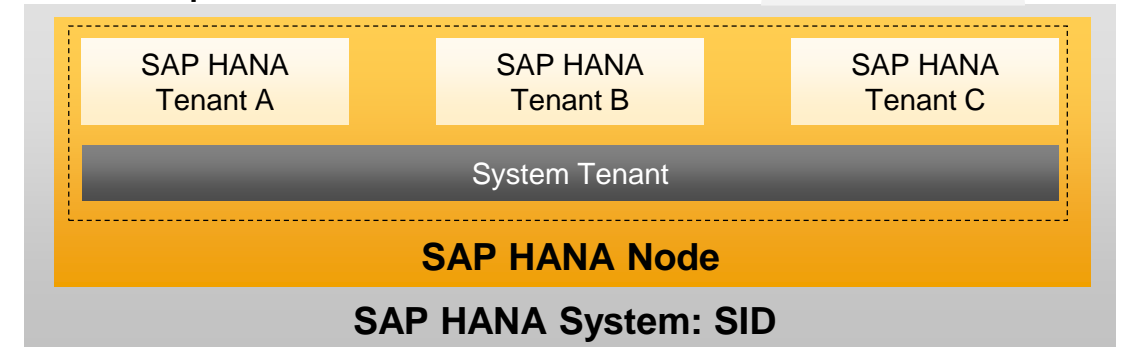


Multitenant database containers

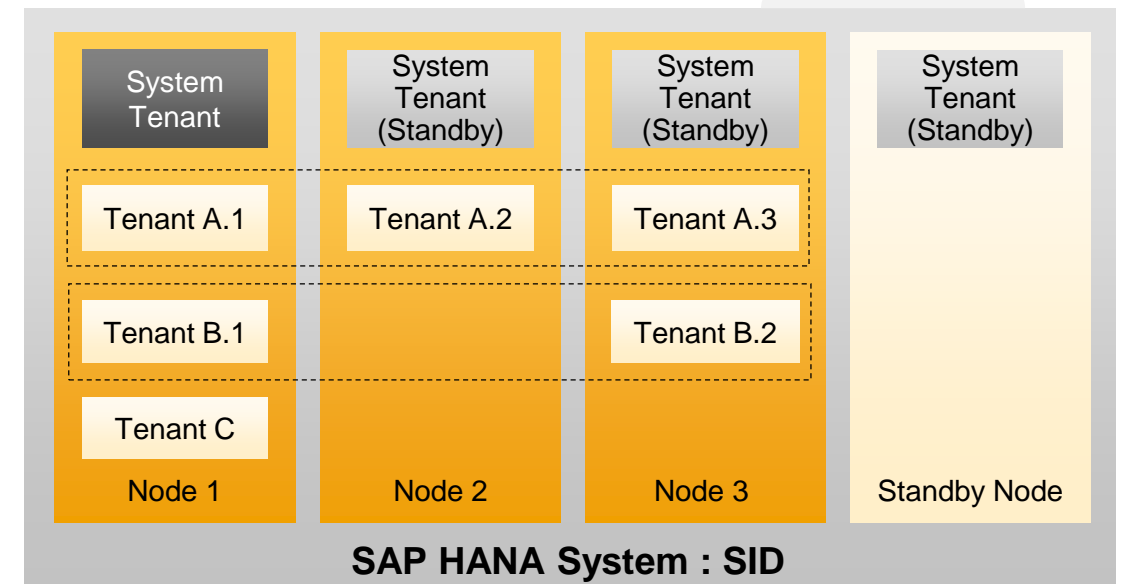
Lower capital and operating expenditure – cloud-ready

- 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



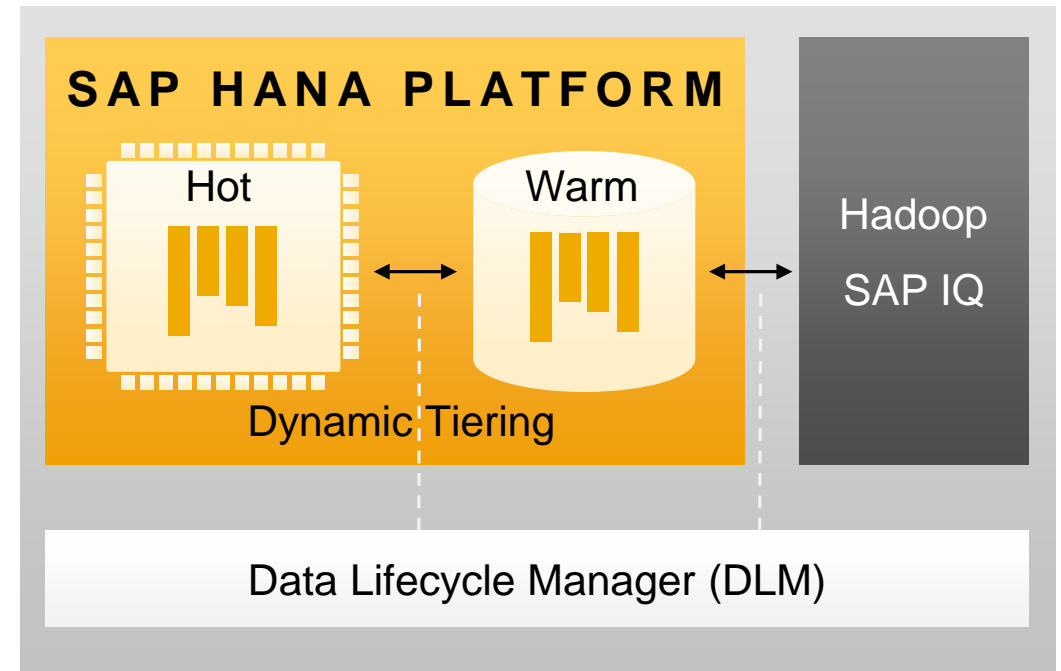
Scale Out



Dynamic tiering

Right price/performance balance between memory and disk

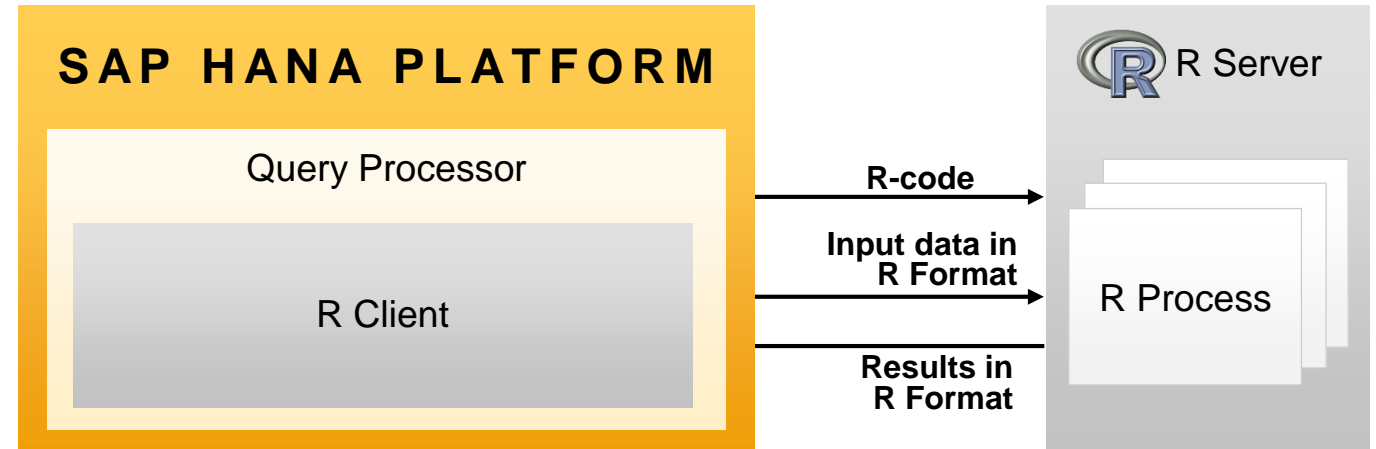
- Utilize disk-based, column-store technology to store less frequently used data
- Support petabyte scale deployment – not confined by the size of memory
- Is integral part of the single SAP HANA instance – no data duplication
- Transparently manage large data volumes by automatically moving data among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



R integration

Use leading open source data mining software transparently

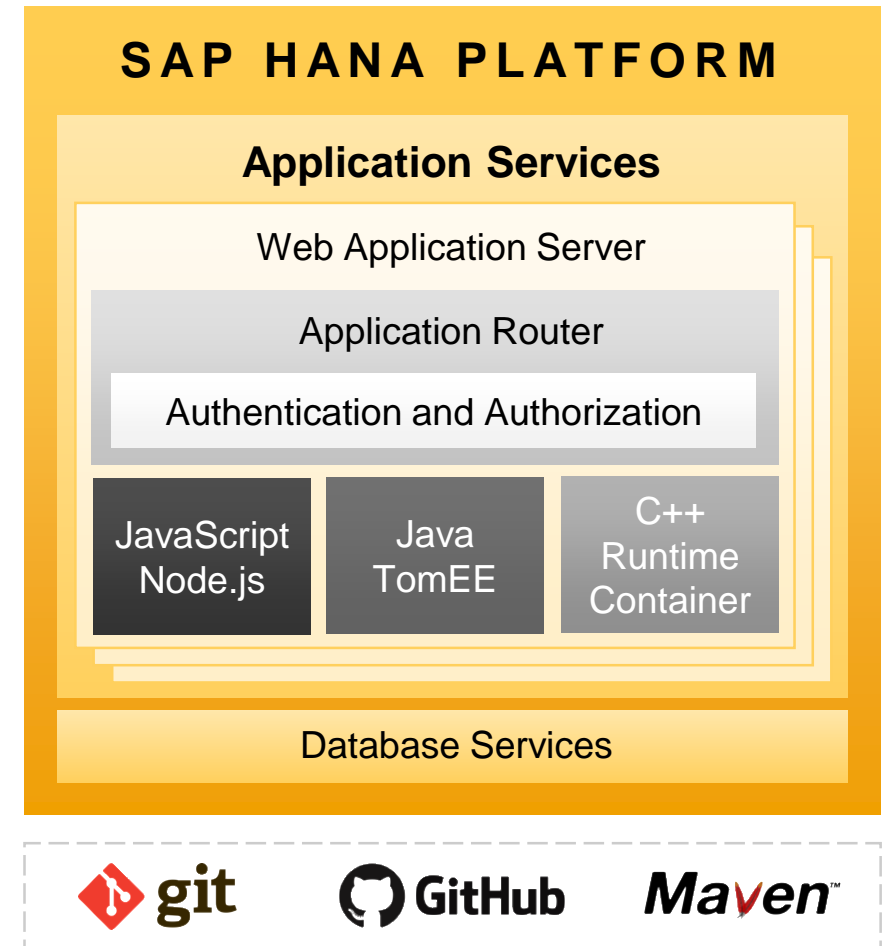
- Embed R script within SQL script
- Execute R script inside R server
- Use R vector-oriented format rather than JDBC/ODBC
- Execute multiple R processes in parallel
- Leverage 3,500+ R statistical and graphical packages



Web server

Reduced data movement – app and database services in one platform

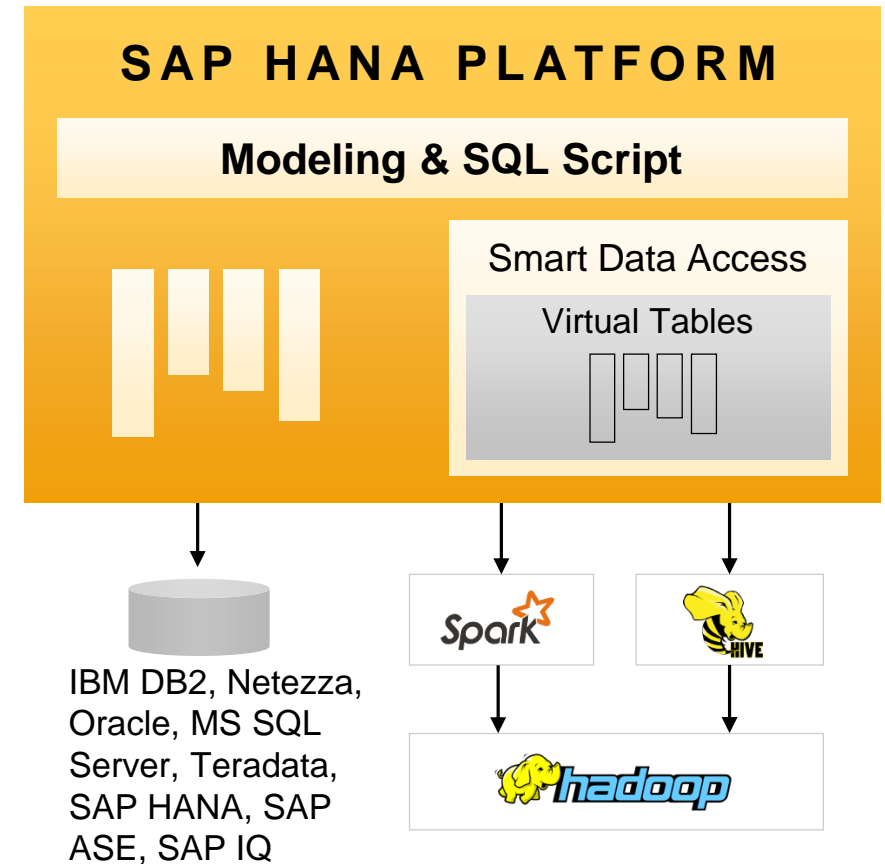
- Scale applications independently from the database services with new web application server
- Supports choice of programming languages – Server side JavaScript on Node.js, Java on TomEE and C++ Runtime container
- Core Data Services allow developers to create database objects and relationships without 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



Smart data access

Access any data from any source

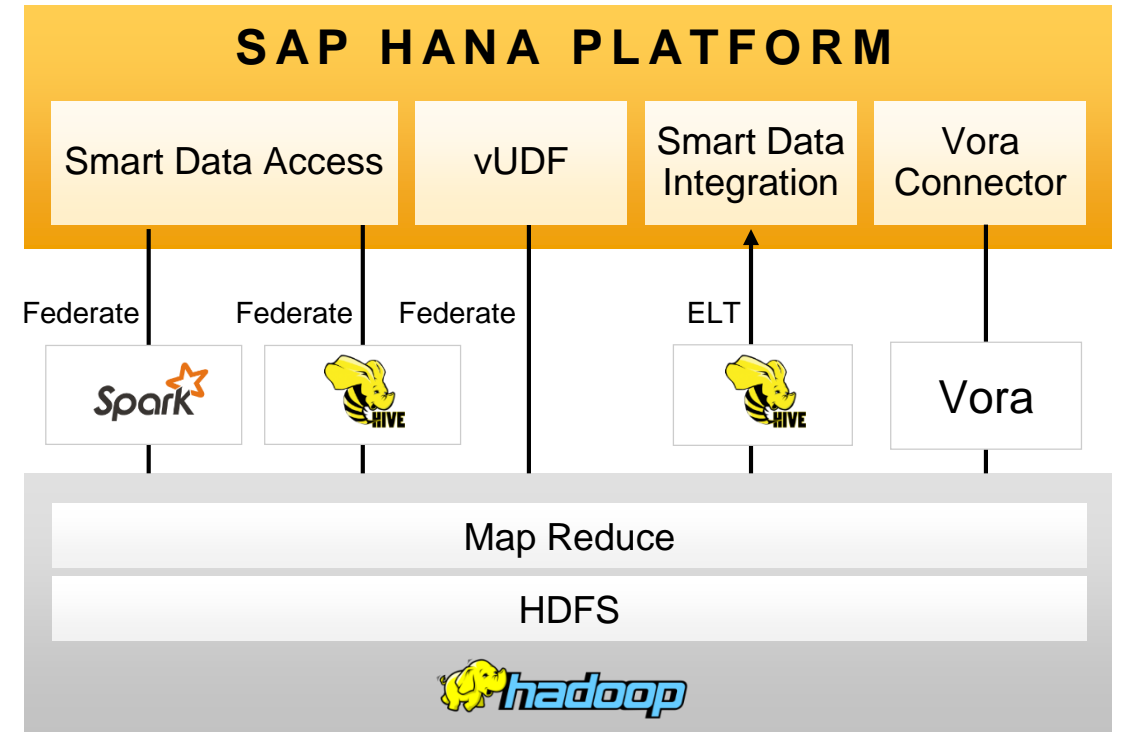
- Manage and query remote tables as local virtual tables
 - Support virtual tables in calculation view and SQL
 - Virtual tables can be combined with PAL, BFL, and Spatial
- Push query processing to remote databases
- Complement functionalities in remote database with SAP HANA capabilities
- Support 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




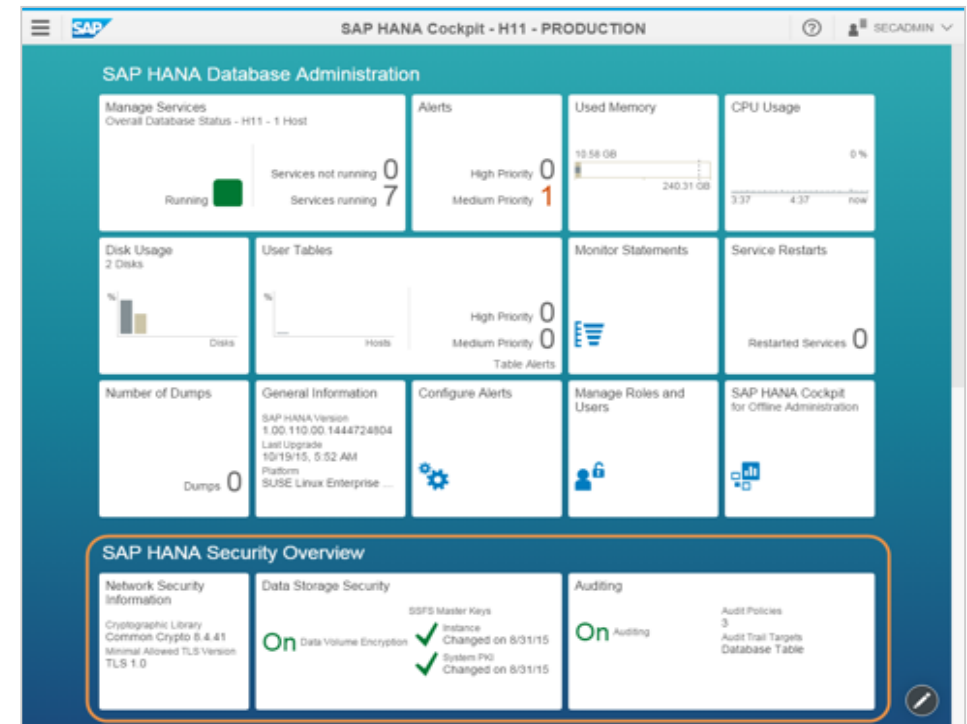
Simplify administration and monitoring

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



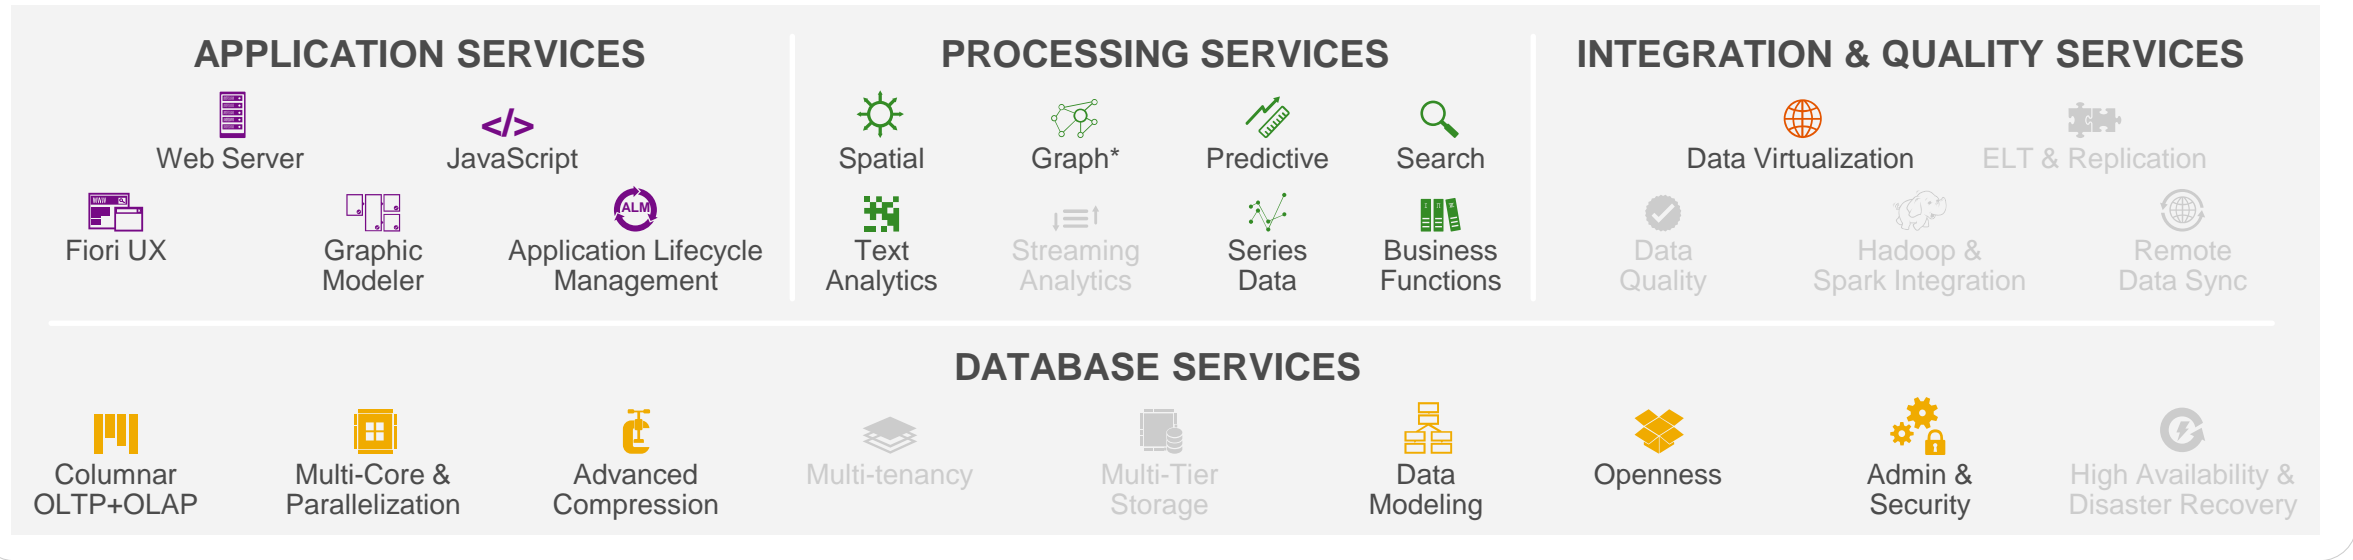
New! SAP HANA Desktop Edition For Developers*

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

Hana Express - DEMO

Demo start

1. Download software from : <http://go.sap.com/developer/topics/sap-hana-express.html>

HW requirements:

JRE8

16 GB memory

30 GB HD free space

Bios option for virtualization on

Demo Hanaexpress license

Systemuser password has to be changed:

```
hdbsql -u SYSTEM -d SystemDB -p HXEHana1  
alter user SYSTEM password LinuxLab1
```

apply the license

Read hardware key:
Hdbstudio -> right mouse button and topic license

<http://sap.com/minisap>

Check email

Demo create tenant / new database

Create database uwe SYSTEM user password LinuxLab1

Create column table mytable (counter integer);

Create table and content:

```
insert into mytable select round(rand()*1000) from dummy;  
commt;
```

```
Select * from mytable;
```

```
Hdbsql -u SYSTEM -d UWE -p LinuxLab1 -l sel.sql -o sel.out
```

Demo python

Install hdbclient

Copy python libs nach lib directory

```
import dbapi
```

```
conn = dbapi.connect('hxehost', 30015, 'SYSTEM', 'LinuxLab1')
```

```
#Check if database connection was successful or not
```

```
print conn.isconnected()
```

```
# fetch table data
```

```
stmnt = 'select * from mytable1'
```

```
cursor.execute(stmnt)
```

```
result = cursor.fetchall()
```

```
print result
```