

SAP HANA BIG DATA PLATFORM

Reinventing Businesses through Innovation, Value & Simplicity

Dr. Jan Teichmann, SAP HANA Product & Strategy

Dec 5, 2013



Agenda



**HANA Data
Platform
Journey**

**SAP HANA
Platform**

**SAP HANA
Platform Future
Evolution**

Summary

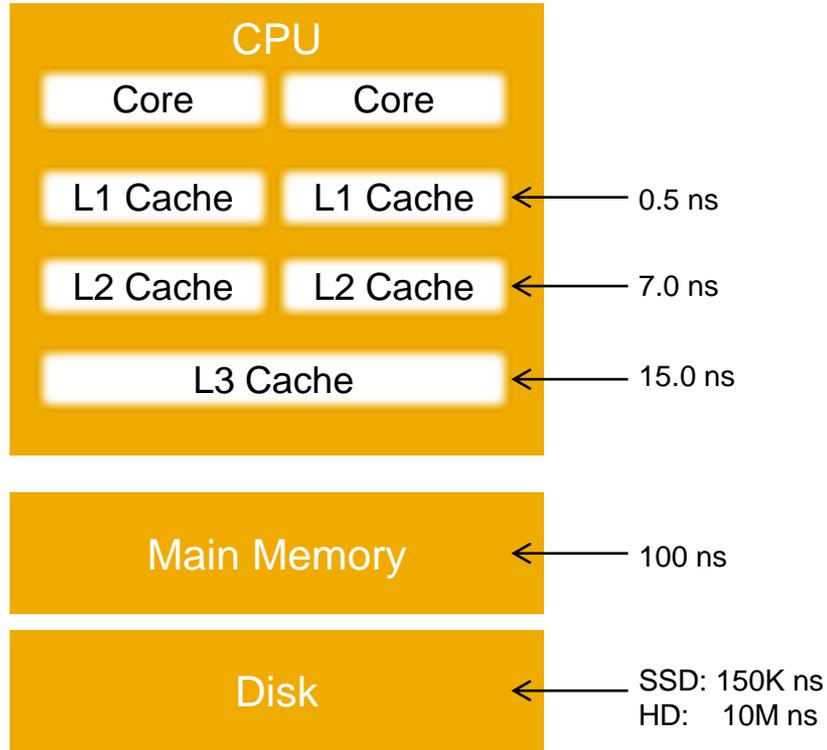
In-Memory Computing



Technology that allows the processing of **massive quantities of real time data** in the main memory of the server to provide **immediate results** from analyses and transactions

SAP HANA Database

Not *just* in-memory: Column-based Tables



Why Column Tables

- Higher Compression Rates
- Higher Performance on Column Operations
- Elimination of additional Indexes
- Elimination of materialized aggregates
- Parallelization

Optional Row Store for configuration data

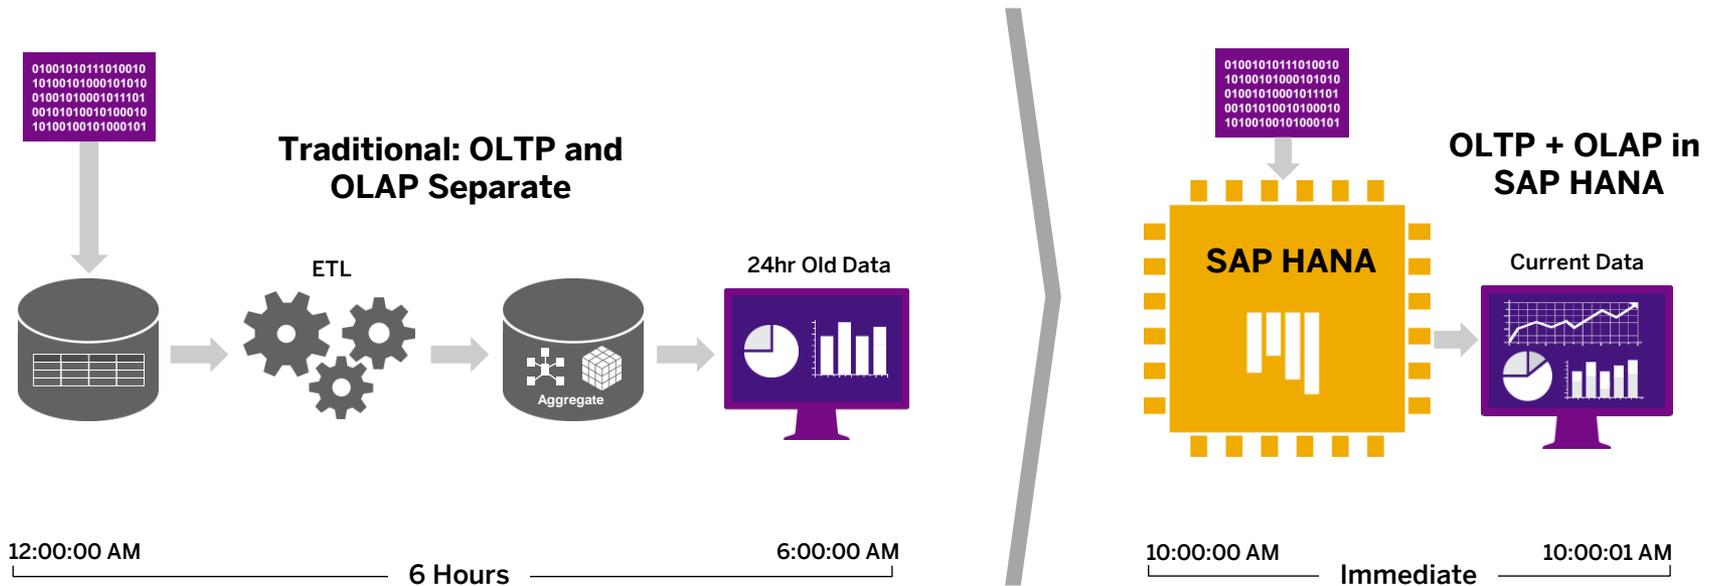
Yes, DRAM is 100,000 times faster than disk, but DRAM access is still 6 to 200 times slower than on-chip caches

SAP welcomes that other RDBMS vendors acknowledge the superior power of in-memory computing with column-store technology.

SAP will keep on striving for thought leadership and transformational innovation.

Real-time Applications, Zero Latency

Not *just* in-memory column-based: Real-time and Landscape Consolidation



Run **both transactional and analytical applications** on **one single data model** with a **single copy of data in-memory**

Simultaneous high volume/speed transactional and analytical processing without compromising data consistency (ACID compliance)

Aggregate on-the-fly with **no pre-materialization** on key figures, including current transactions

Product & Technology Convergence

TREX / BWA
P*Time
MaxDB



HANA is all about innovation



SAP Sybase ESP



SAP Sybase IQ



SAP Sybase ASE



SAP Sybase SA



MaxDB

Deep Event Streaming Integration

Hot & Cold Store tables

IQ as NLS in the HANA

IQ as an engine

Unified Table

Database federation

Development IDE integration

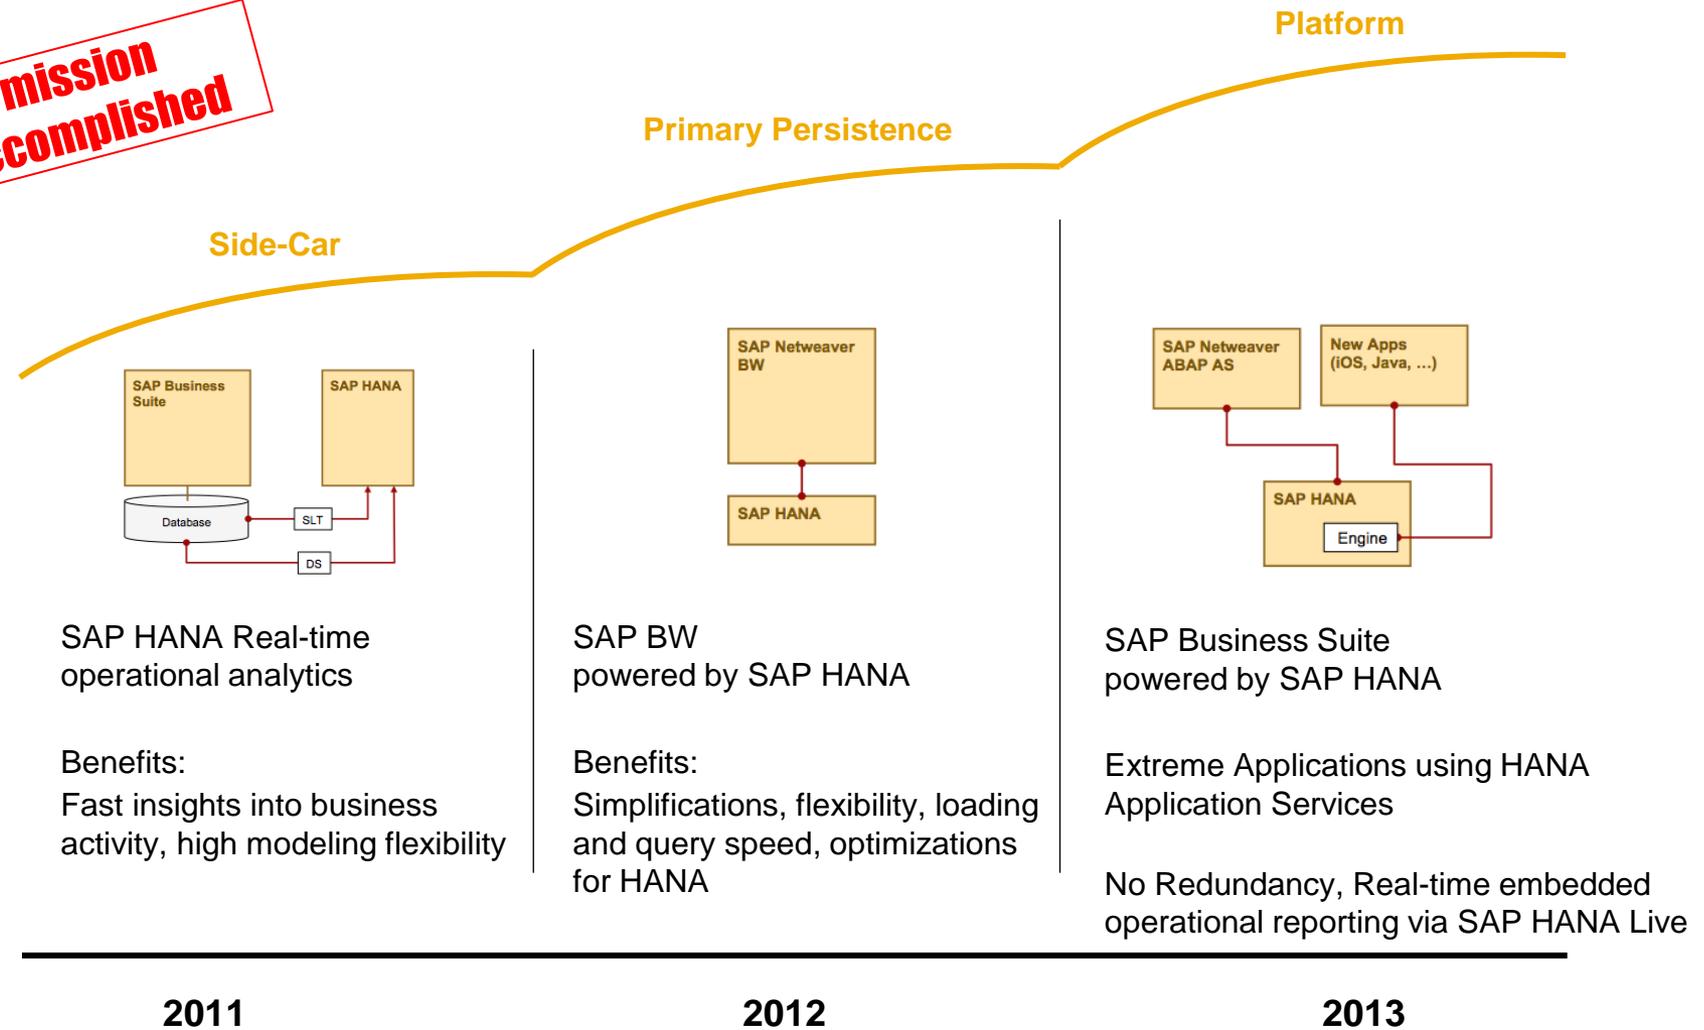
System Management

...



SAP HANA - Strategy

mission accomplished



Agenda

**SAP HANA
Data Platform
Journey**

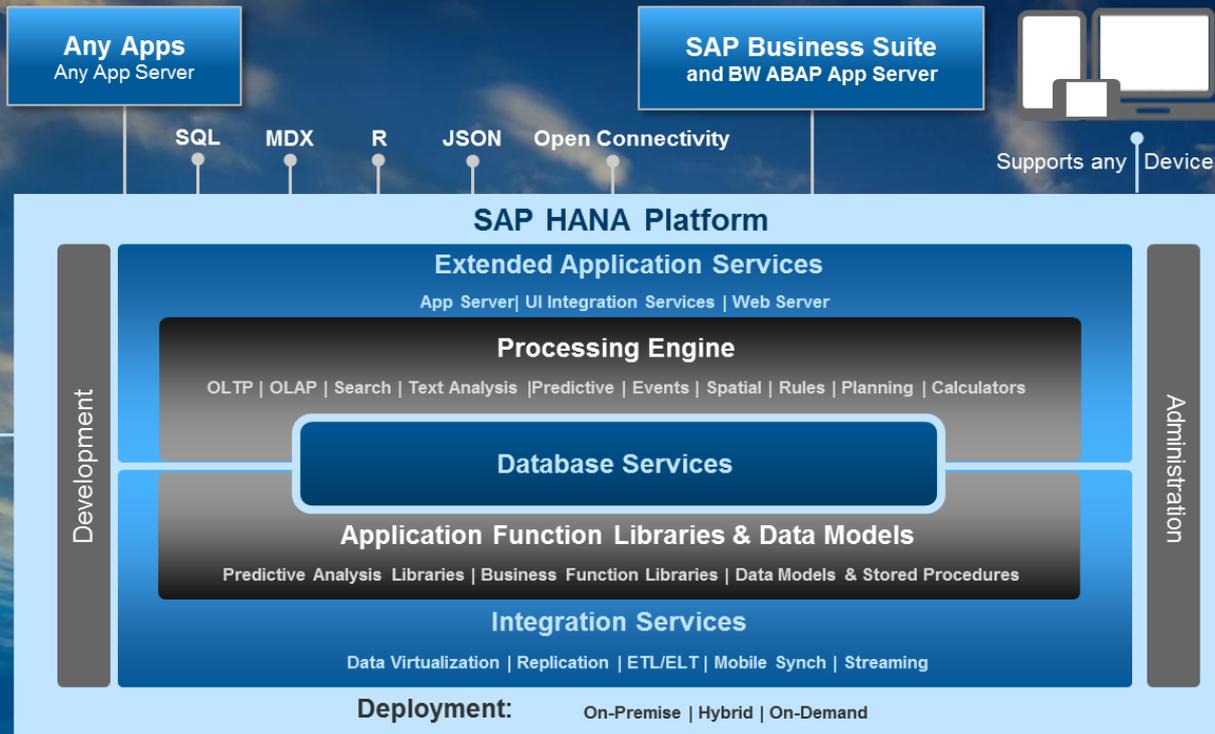
**SAP HANA
Platform**

**SAP HANA
Platform Future
Evolution**

Summary

SAP HANA Platform – More than just a database

SAP HANA PLATFORM



- Fully ACID compliant, In-memory, columnar, massively parallel processing database platform
- Includes In-memory function libraries, data models and stored procedures
- Integrated data processing for end to end analytic processing

Scan

3.2 billion billion integer/sec/core

12.5 million aggregates/sec/core

Ingest

1.5 million records/sec/node

2.5 TB data load/hour data services

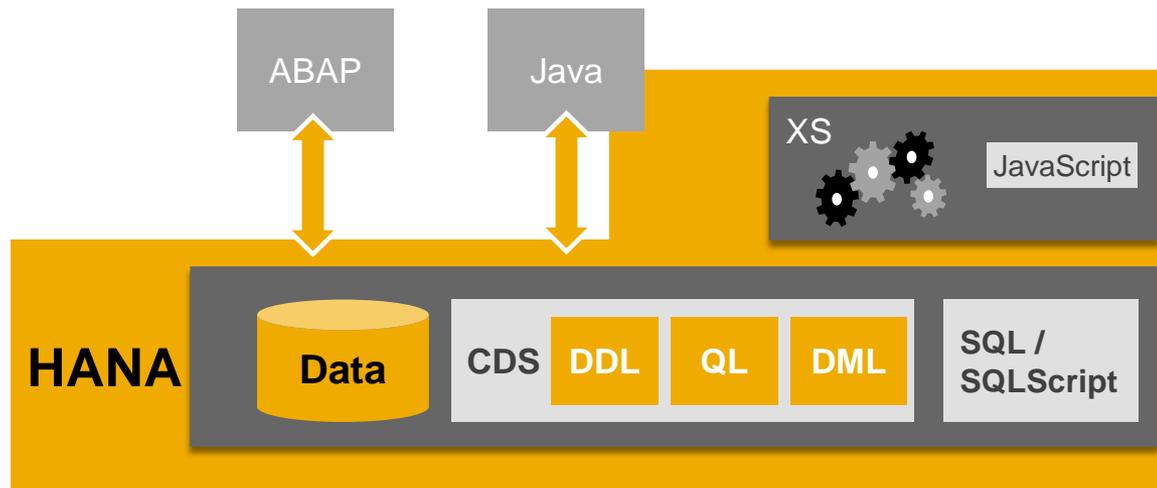
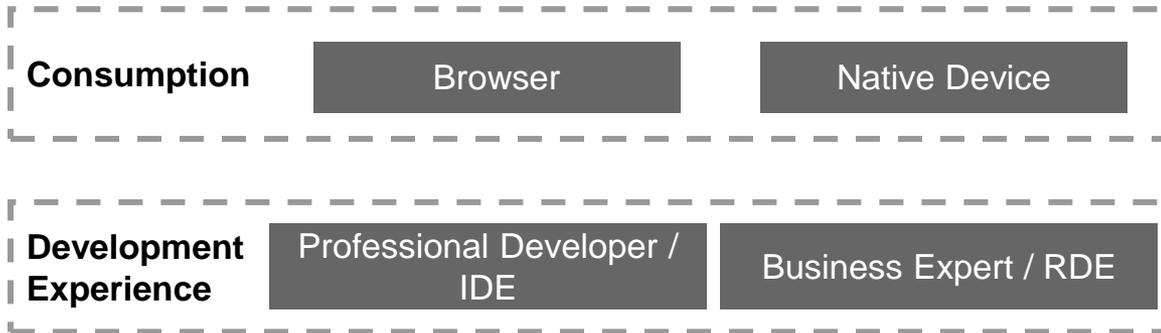
Single Atomic Data Copy
shared across multiple data processing engines to eliminate redundancy & latency

Process data and application logic in parallel (MPP)
Using all cores in a multi-core architecture, by effectively partitioning data.

Avoid unnecessary compensation during application execution
(e.g.: buffering, data duplication)

Eliminate disk I/O
by keeping all data in memory using column store, and by significantly compressing data

HANA Platform



HANA Big Data Platform (DB)



HANA Application Platform

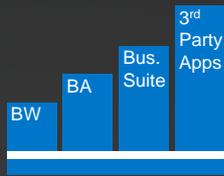


HANA Cloud Platform

Deployment Services

Provides security, privacy, and availability

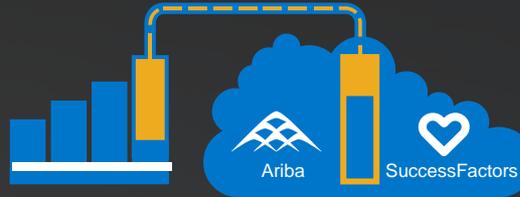
Choose and change your deployment options anytime



On Premise

Run All SAP Solutions on SAP HANA

- Build or deploy your own solutions on SAP HANA
- Maintain all within your firewall
- Upgrade or leverage existing infrastructure



Hybrid

Leverage SAP Cloud

- Migrate some solutions to the cloud
- Create or deploy new SaaS apps in the cloud
- Use cloud hosting and managed services
- Deploy via SAP HANA Enterprise Cloud or public cloud



Cloud

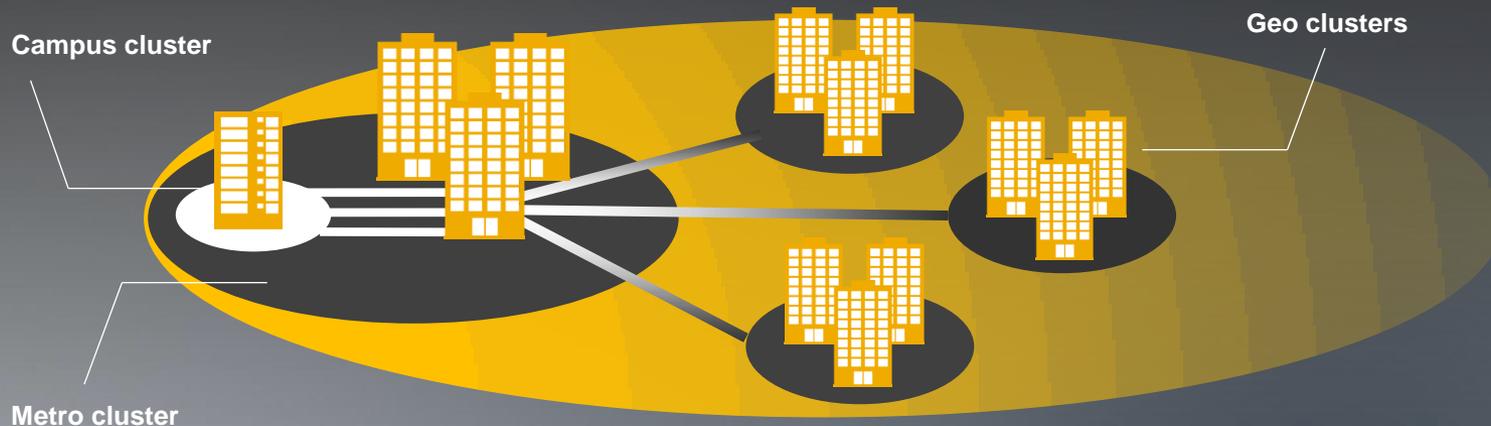
Build, Run, Deploy all Applications in the Cloud

- Consider Virtual Private Cloud option
- Enable faster innovations
- Simplify landscape
- Migrate or build new applications in SAP HANA Enterprise Cloud



Mission-critical infrastructure

Ensuring most demanding service-levels



Single Server



- 2 CPU 128GB to 8 CPU 1TB
(Special layout for Suite on HANA for up to 4 TB per host)
- Single SAP HANA deployments for data marts or accelerators
- Support for high availability and disaster recovery

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: 100+ servers
- Support for **high availability and disaster recovery**

100 TB compressed data x avg compression rate 4 – 8

400 - 800 TB raw data
+ Smart Data Access
= **Big Data Volume**

SAP Runs SAP: BW, ERP, CRM

Financial Close on SAP HANA Q3 2013

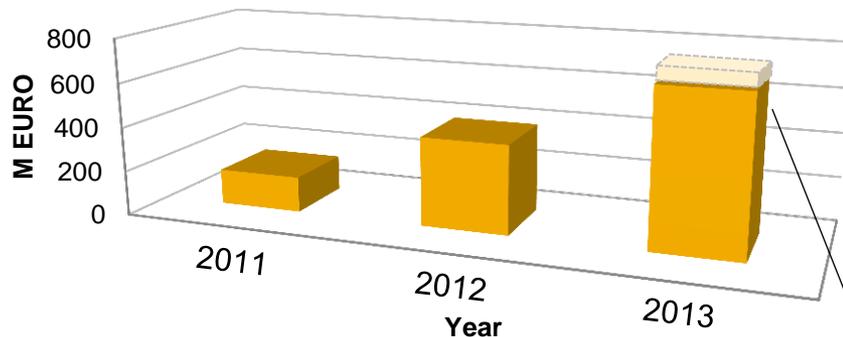
<p>Quality - real time analytics with COPA on HANA</p> <p>> 86% faster</p>	<p>Agilty - analysis of Balance Sheet and P&L</p> <p>> 29% faster</p>	<p>New Business Scenarios based on real time Sales Data</p> <p>from > 2 hours to real-time</p>
<p>Automatic Cost Accrual Postings</p> <p>> 42% faster</p>	<p>Automatic Maintenance Revenue Accruals</p> <p>> 91% faster</p>	<p>> 65.000 users</p>
<p>Cost Postings</p> <p>> 10% faster</p>	<p>Asset Depreciation</p> <p>> 74% faster</p>	<p>Batch Input Processing</p> <p>71% faster</p>
<p>Real time Error Detection through Deviation Analysis and drilldown to lowest level</p>	<p>from > 2 hours to real-time</p>	<p>Software & Maintenance Billing</p> <p>> 17 % faster</p>
		<p>Intercompany Reconciliation</p> <p>Real-time</p>

Go Live: Aug 20, 2013
First global go live of SAP ERP on HANA

More than 65,000 ERP users
Implementation in less than 5 months

SAP HANA: Financial Viability

HANA License Revenue



SAP aims to remain the fastest-growing database company

Oct 21, 2013: In-Memory Leadership – 90% Year-Over-Year Growth at Constant Currencies in HANA Software Revenue (79% at Actual Currencies)

Q4 2013 ongoing – projection based on Q3 results; published by Jim Hagemann-Snabe

500+

BW on HANA Customers

1000+

Startups on HANA

1000+

Customers of HANA in the Cloud

Unprecedented Customer Adoption

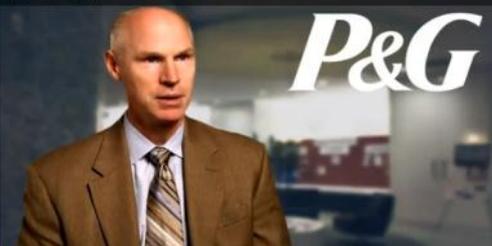
Nongfu Spring Improve Their Performance with BW Powered by SAP
by saptotechnology



Patrick Hoo | CIO
Nongfu Spring

Innovate continuously on SAP HANA platform at Nongfu Spring

SAP In-Memory Computing - Procter & Gamble Customer Testimonial
by SAPBusinessObjects



P&G rethinks architecture of their transactional systems and analytical engines with SAP HANA



Phil Shotts
Vice President & Assistant Corporate Controller
Colgate Palmolive

Colgate runs better with SAP BPC powered by SAP HANA

MKI Makes 400,000x Difference in Healthcare Industry with SAP HANA
by saptotechnology



Yukihisa Kato
CTO and Director of MKI

MKI: Real-time big data analysis with SAP HANA, R, Hadoop for genome interpretation

SAP HANA helps Chiron reduce Manufacturing and Assembly time
by saptotechnology



Frank Schlotter
Process Manager, Chiron Werke

Chiron uses SAP HANA on the assembly line to display real-time data analysis

SAP HANA Enables End-to-End Business Process at ConAgra
by saptotechnology



Gerrit Schütté
Chief Information Officer
ConAgra Foods

ConAgra enables end-to-end business with SAP BPC powered by SAP HANA

Coca-Cola has Major Performance Improvement with SAP HANA
by saptotechnology



Baris Tiryaki
Associate Director IT Finance
Coca-Cola Enterprises, Inc.

Coca Cola's performance improvement with SAP BPC powered by SAP HANA

Alliander Optimizes Customer Energy Management with Smart Meter
by saptotechnology



Jeroen Scheer
Manager Taskforce Energy Transition IT, Alliander NV

Alliander optimizes customer energy management with SAP HANA

HANA
1,500+ customers

[http://www.saphana.com/
community/learn/
customer-stories](http://www.saphana.com/community/learn/customer-stories)

eBay – Professional Service (Internet)

American multinational internet consumer-to-consumer corporation



Determine with 100% Accuracy
that a signal is positive at 97% confidence



Automated Early Signal Detection
system powered by SAP HANA

Product: Early Signal Detection System Powered by Predictive Analytics on SAP® HANA

Business Challenges/ Objectives

- Increase ability to separate signal from noise to identify key changes to the health of eBay's marketplace
- Improve predictability and forecast confidence of eBay's virtual economy
- Increase insights into deviations and their causes

Technical Challenges

- Detect critical signals from 100 PBs of data in eBay EDW
- Highly manual process because one model does not fit all the metrics hence requires analyst intervention

Benefits

- Automated signal detection system powered by predictive analytics on SAP HANA selects best model for metrics automatically; increases accuracy of forecasts
- Reliable and scalable system provides real-time insights allowing data analysts to focus on strategic tasks
- Decision tree logic and flexibility to adjust scenarios allows eBay to adapt best model for their data



“HANA is valuable in the sense that it accelerates that speed to insight. HANA, with in-memory capability, with multicore, fast, lots of data, all of that coming together is how I think analytics is going to work broadly in the future.”

David Schwarzbach, VP&CFO eBay North America at eBay Inc.

“HANA system will free up all the bandwidth right now involved in figuring out what is going. The user just has to feed in their metric, doesn't have to really worry about which algorithm is the best and be able to use the system because it is inherently intelligent and configurable.”

Gagandeep Bawa, Manager, North America FP&A at eBay Inc.



408,000x

faster than traditional disk-based systems in a technical PoC



216x faster by reducing genome analysis from several days to only 20 minutes making real-time cancer/drug screening possible

Business Challenges

- Reduce delays and minimize the costs associated with new drug discovery by optimizing the process for genome analysis
- Improve and speed decision making for hospitals which conduct cancer detection based on DNA sequence matching

Technical Implementation

- Leveraged the combination of SAP HANA, R, and Hadoop to store, pre-process, compute, and analyze huge amounts of data
- Provide access to breadth of predictive analytics libraries

Benefits

- For pharmaceutical companies, provide required new drugs on time and aid identification of “driver mutation” for new drug targets
- Able to provide a one stop service including genomic data analysis of cancer patients to support personalized patient therapeutics



Our solution is to incorporate SAP HANA along with Hadoop and R to create a single real-time big data platform. With this we have found a way to shorten the genome analysis time from several days down to only 20 minutes.

Yukihisa Kato, CTO and Director of MITSUI KNOWLEDGE INDUSTRY



5,000 events per second loaded onto SAP HANA (not possible before)



10-30% increase in revenue per year



Interactive data analysis leading to improved design thinking and game planning



At Bigpoint in the Battlestar Galactica online game, we have more than 5,000 events in the game per second which we have to load in SAP HANA environment and to work on it to create an individualized game environment to create offers for them. In this co-innovation project with SAP HANA, using Real Time Offer Management Bigpoint, we hope to increase revenue by 10-30%.

Claus Wagner, Senior Vice President SAP Technology, Bigpoint

Business Challenges

- Increase conversion rates from free → paying player
- Increase the average revenue per paying player
- Decrease churn – keep paying players playing longer

Technical Challenges

- Leverage real-time data processing in SAP HANA and classification algorithms with R integration for SAP HANA to deliver personalized context-relevant offers to players
- Analyze vast amounts of historical and transactional data to forecast player behavior patterns

Benefits

- Real-time insights
- Per player profitability analysis and increased understanding of player behavior
- Increase data volume and processing capabilities to communicate personalized messages to players

Agenda

**SAP HANA
Data Platform
Journey**

**SAP HANA
Platform**

**SAP HANA
Platform Future
Evolution**

Summary

The Pattern of Change

- The Rate of Change

“...technological change is exponential, contrary to the common-sense intuitive linear view. So we won’t experience 100 years of progress in the 21st century, it will be more like 20,000 years of progress.”



Ray Kurzweil
Born Feb 1948
Computer Science
& Literature, MIT

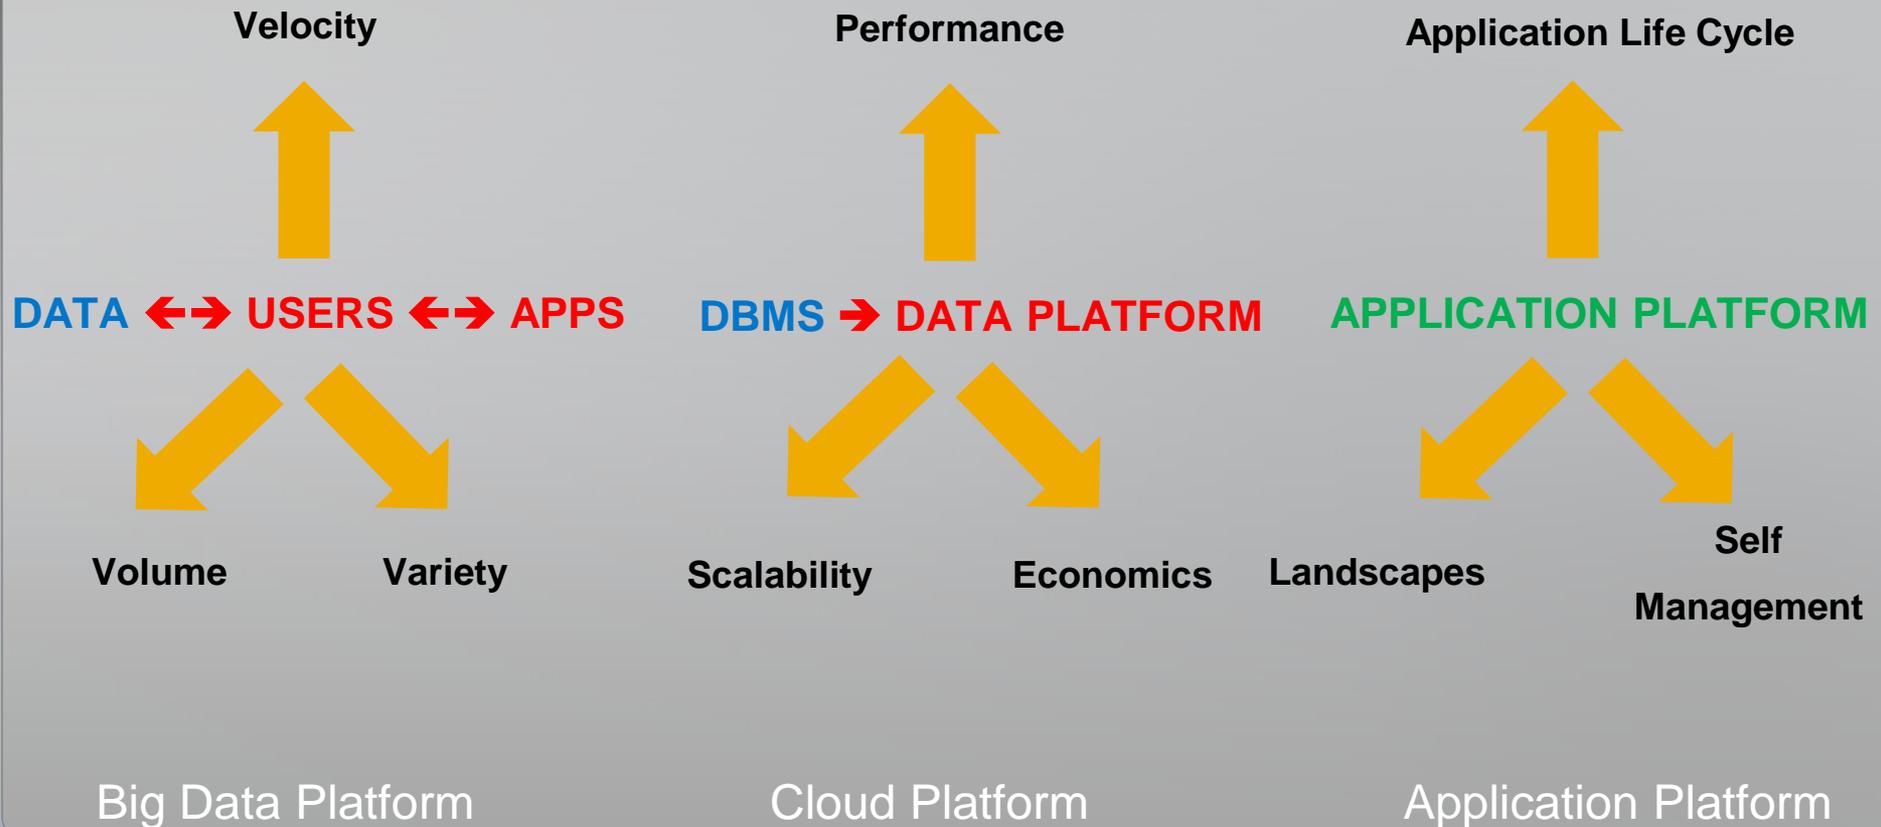
Founder, chairman
& CEO
Kurzweil
Technologies Inc

Co-founder,
chairman & CEO of
K-NFB Reading
Technology

Co-founder,
chairman & CEO of
Ray & Terry’s
Longevity
Products

The Next Frontier: Making Big Data Real

Getting past the basics



Consumer Engagement Applications

Game-changing innovation with new applications and business models



Real-time
Engagement for
consumers

Real-time
insights for
enterprise

- Generate real-time consumer insights
- Deliver personalized, engaging experiences
- Create more responsive business with precision targeting

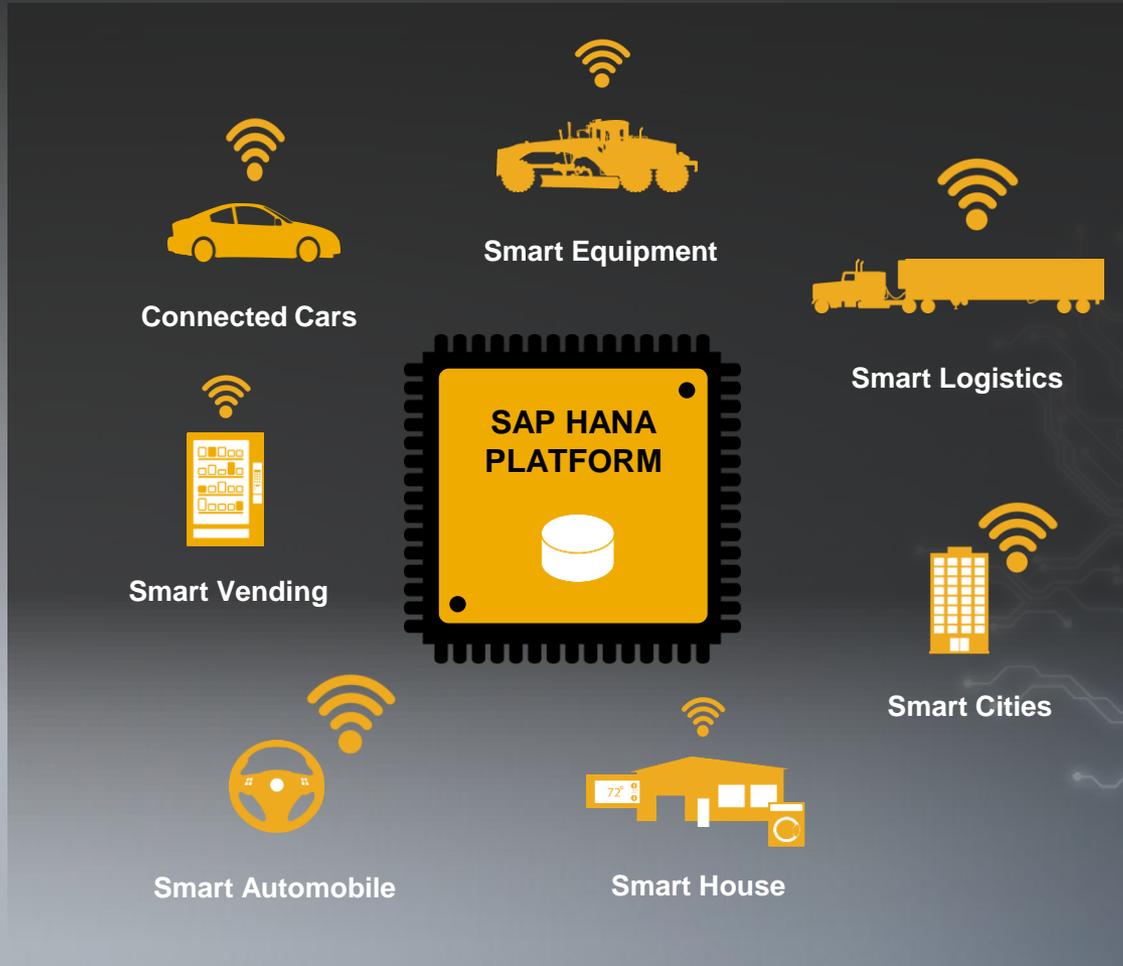


- HTML-5 support
- Mobile integration
- Consumer-grade usability

- OLTP + OLAP real-time Processing
- Sentiment Intelligence
- Predictive analytics

Sense & Respond Applications for Internet of Things

Real-time adjustments to signals in the environment



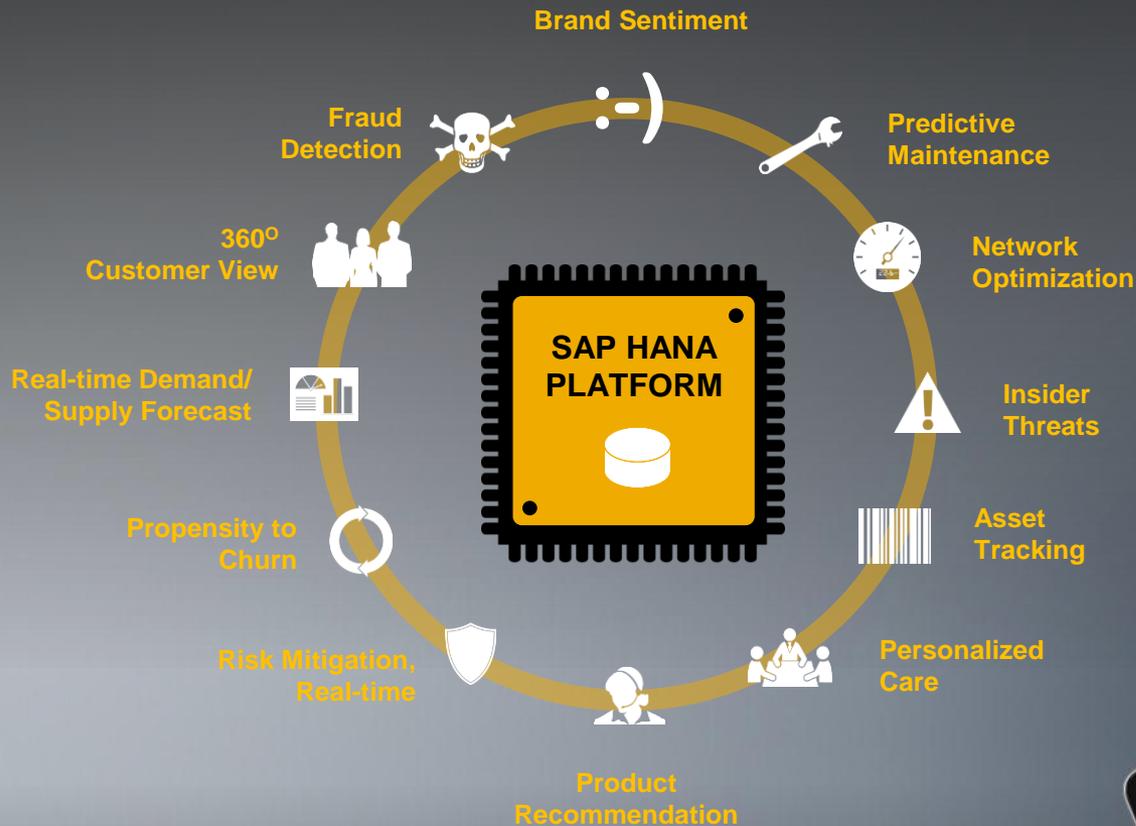
- Detect and analyze data trends by aggregating sensor data
- Enable automated M2M and machine-to-human communications
- Benefit from more energy efficient logistics, transforming retail distribution
- Transformational experiences through intelligent buildings, robots, cars and cities

- Event stream processing (ESP)
 - No data preparation
 - RFID Integration

- Embedded data processing with SQLAnywhere
 - Machine Learning
 - Spatial Processing

Big Data Warehousing drives big business

Deeper insights into customer needs and wants



- **Extract value from big data**
Bring together the power of in-memory, Hadoop, predictive, text mining and spatial analytics
- **Dramatically improved performance** for decision making from variety and volumes of structured and unstructured data
- **Create a logical data warehouse** to respond to changing business environment



- Hadoop/R Integration
- Text Search, Analytics, NLP
- Geospatial Processing

- SAP HANA smart data access
 - BW & SAP Sybase IQ
 - In-DB Predictive Analytics

SAP HANA Driving Innovation

SAP University Alliances

- SAP University Alliances (UA) provides connections between students, customers, partners, and SAP experts.
- SAP donates licenses to over 1,500 UA Program member institutions and fully outfits their professors.

Startup Program

- A 12-month global program (renewable)
- To help startups
- In the Big Data, predictive and real-time analytics space
- Develop new applications on SAP HANA
- And accelerate market traction



Agenda

**SAP HANA
Data Platform
Journey**

**SAP HANA
Platform
Evolution**

**SAP HANA Eco
System, Co-
Innovation and
Academic
Collaboration**

Summary

Summary

In-memory database technology is becoming an industry standard

SAP HANA Platform is a mature and cutting-edge data platform

SAP HANA Platform is far more than an RDBMS

SAP HANA Platform is serving a wide variety of scopes and industries – ERP, Data Warehousing and many more

Further Information

The developer guide on help.sap.com:

http://help.sap.com/hana_platform/#section6

http://help.sap.com/hana/SAP_HANA_Developer_Guide_en.pdf

The HANA Academy:

<http://www.saphana.com/docs/DOC-3017>

<http://www.saphana.com/community/hana-academy>

The openSAP course:

<https://open.sap.com/course/hana1>

HANA XS - Shortcut to the OpenSAP Tutorial Result

By Marcus Semling

<https://wiki.wdf.sap.corp/wiki/display/WebAnalytics/HANA+XS+-+Shortcut+to+the+OpenSAP+Tutorial+Result>



Further Information (II)

HANA Bluebooks

HANA Bluebook:

<https://portal.wdf.sap.corp/irj/go/km/docs/guid/506044d8-61b7-2d10-4698-ff22eb27ec23>

HANA XS Architecture Overview:

<https://portal.wdf.sap.corp/irj/go/km/docs/guid/605f4031-6d5e-3010-e19b-f6730c6a0e03>

HANA Getting Started Guide for Architects:

<https://portal.wdf.sap.corp/irj/go/km/docs/guid/41d4c6fc-b951-3010-1394-a34a38864c3d>

PUMA Sessions (2012)

<https://wiki.wdf.sap.corp/wiki/display/TIPCrossPM/Recordings+of+Previous+HANA+Sessions>

- SAP HANA 101 – Learn to use SAP HANA (Ingo Brenckmann)
- SAP HANA - Overview And Use Cases (Dennis Schwerer)
- SAP HANA Application Services – Primer for Building Applications and Application Services Directly on HANA (Thomas Jung)
- SQLScript: Leveraging the Power of HANA by Pushing Processing and Logic into the Database (Thomas Jung)
- Pushing Code Down to Hana (Stefan Dipper)



Further Information (III)

XS Community / Information

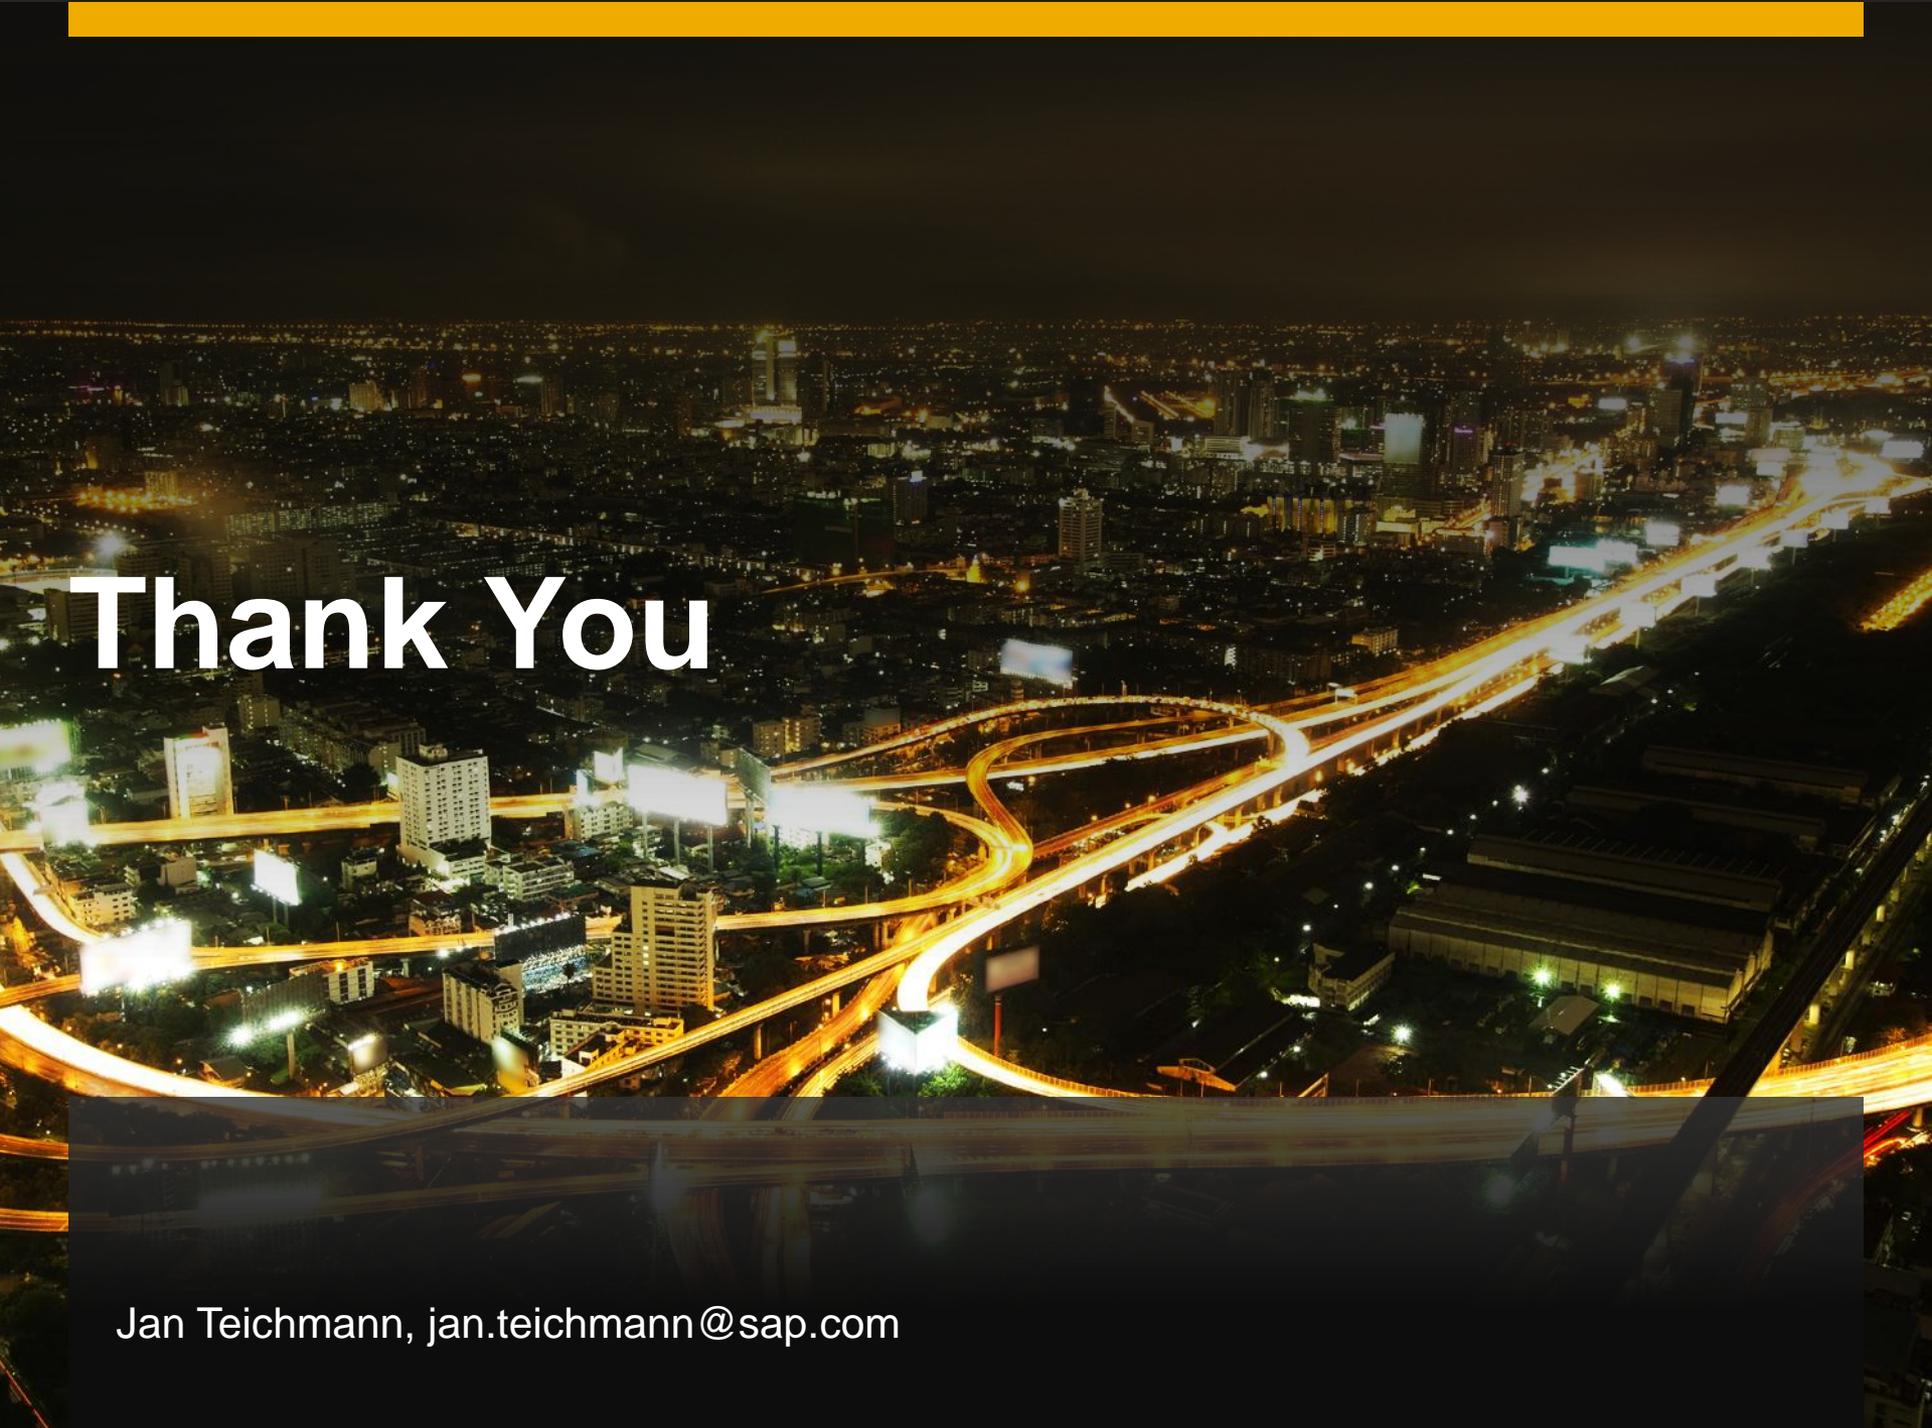
HANA XS:

<https://community.wdf.sap.corp/community/xs>

HANA XS Inner Source

<https://wiki.wdf.sap.corp/wiki/display/xs/open/XS+Inner+Source+Community>



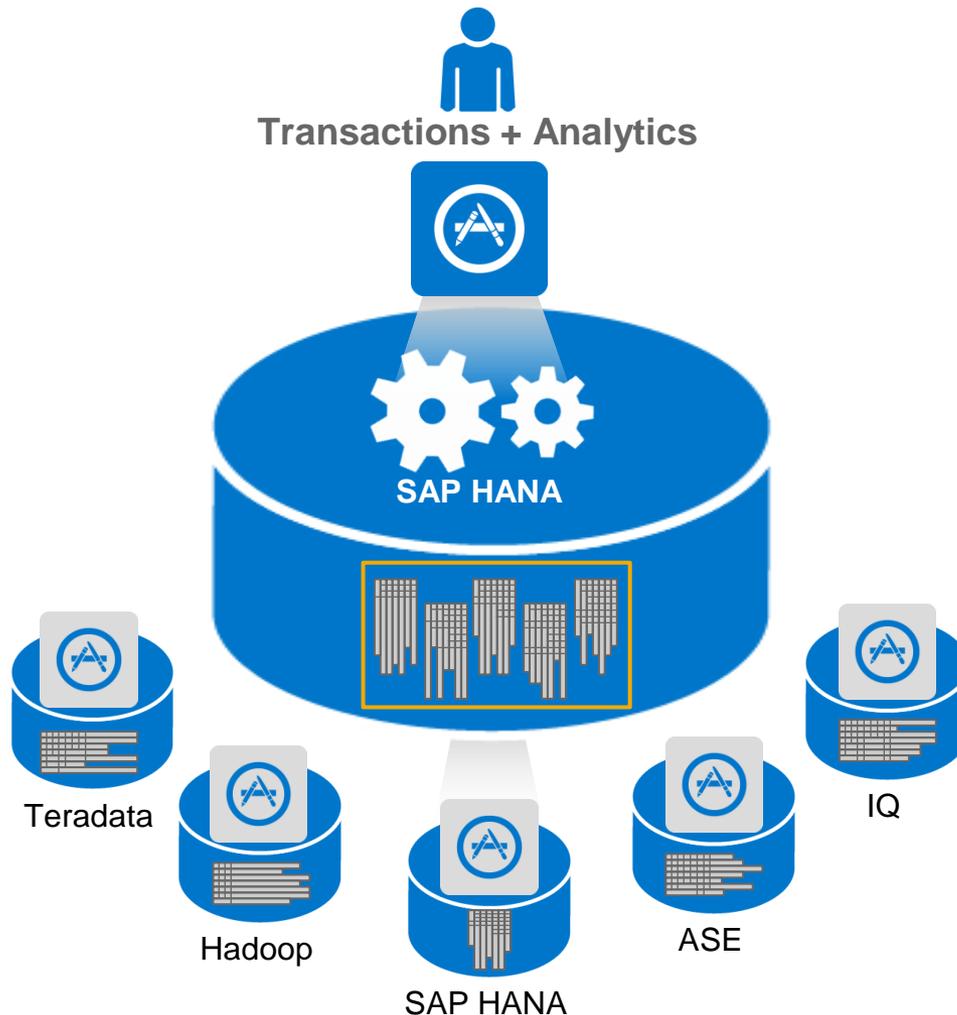
An aerial night photograph of a city, likely Singapore, showing a complex multi-level highway interchange with glowing light trails from traffic. The city lights are visible in the background, and a large, dark, semi-transparent rectangular area is overlaid at the bottom of the image.

Thank You

Jan Teichmann, jan.teichmann@sap.com

SAP HANA Smart Data Access

Data virtualization for on-premise and hybrid cloud environments



Benefits

- Remote real-time query processing
- Smart continuously self-tuning system
- Secure access to heterogeneous data sources

Heterogeneous data sources

- SAP HANA to Hadoop (Hive)
- Teradata
- SAP Sybase ASE
- SAP Sybase IQ

SAP HANA Smart data access

Differentiation

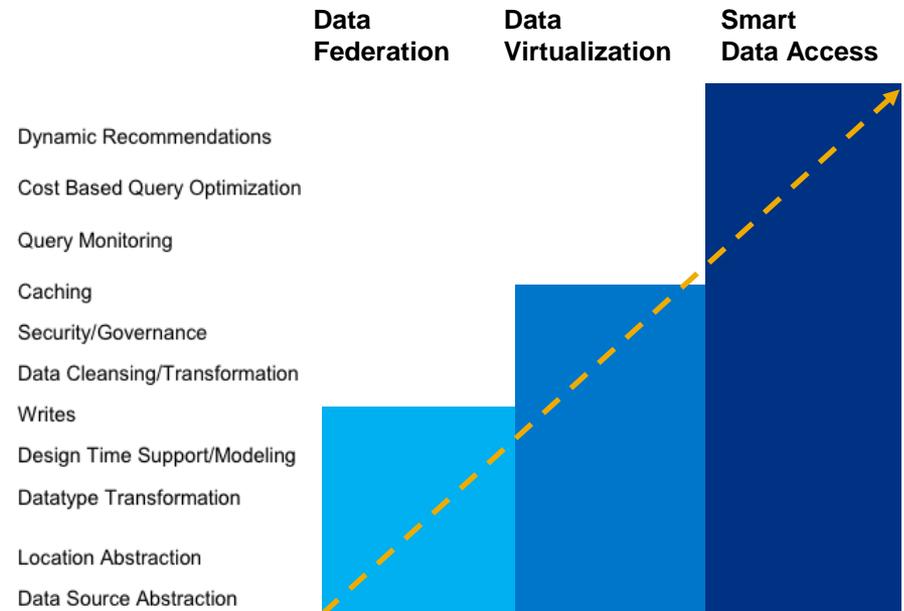
The intelligence of knowing when to delegate query processing or pull the data into SAP HANA for query processing, based on the performance windows

- Dynamic query recommendation

To return query results extremely fast. Capabilities supporting fast processing leveraging in-memory acceleration

- Cost-based query optimization
- Data pre-caching
- In-flight transformation

Converged data processing



SAP HANA - Openness

SAP is committed to a Truly Open Ecosystem for SAP HANA

- Intel partnerships for CPU optimization and Hadoop distribution
- 11 Hardware partners with > 70 available hardware landscapes, incl. Virtualization
- Open APIs for BI (MDX, SQL), WebDevelopment (HTTP/S), Dev Platforms (ODBC/JDBC)
- 3rd party Software certification for backup infrastructures, integrate SAP HANA within bigger management environments, or provide Single-Sign-On (SSO) capabilities
- Several (growing number of) Cloud Service Providers
- <http://www.saphana.com/community/blogs/blog/2013/09/24/engineering-open-appliances-for-high-performance-without-lock-in>

