SAP MaxAttention Summit 2013

An Introduction to the SAP Platform powered by SAP HANA

TABLE OF CONTENTS

1 INTRODUCTION

Today, most large enterprises use the SAP Business Suite as the data foundation for running their business. The business data is stored in the databases of the SAP Business Suite systems – partly spread over different systems and databases. The data is often duplicated to provide different views e.g. for online operations or analytics purposes. Business data availability and consistency is very important to our customers, and they would like to use it in additional use cases and broader analysis scenarios. They would also like to combine the business data with information which is currently not stored inside SAP, and use the combination as the basis of more advanced decision support.

Why was this not possible in the past?

In the past, business data had to be replicated into a separate system for analysis, because of the technical limitations of the transactional systems. The CPU and memory capacity of transactional systems did not allow flexible analytics without impacting the transactional performance. This required many copies of data, and long running batch jobs for data load and aggregation. This has led to restrictions in analytical functions: accruals, profitability analysis, planning and analytics were based on duplicated, summarized and not current data.

The SAP platform powered by SAP HANA removes these technical limitations. You can now run transactional applications and analytics on the same set of data, and in the same database. Data loading for separate analytical systems and pre-calculated and separately stored aggregates are no longer required. Business process operation becomes much simpler as well: Operating on a single set of real-time business data eliminates the manual work which is necessary today. Further, working on original business data allows drill-downs to line item level. The SAP platform powered by SAP HANA is therefore the foundation for a new type of decision support – real-time and with full precision. Finally the overall administration cost is also reduced significantly, bringing you into the position to provide much smarter IT solutions with less cost.

The SAP platform powered by SAP HANA establishes a “single source of truth for data”. In addition to consolidating the SAP data, this platform provides rich capabilities to include data from different sources – structured and unstructured – and perform analytics across different data formats. You can for instance now efficiently integrate data from your suppliers and customers, or from social media. Fully leveraging this platform will optimize your value chain end-to-end, and will give you completely new insights into your customers.

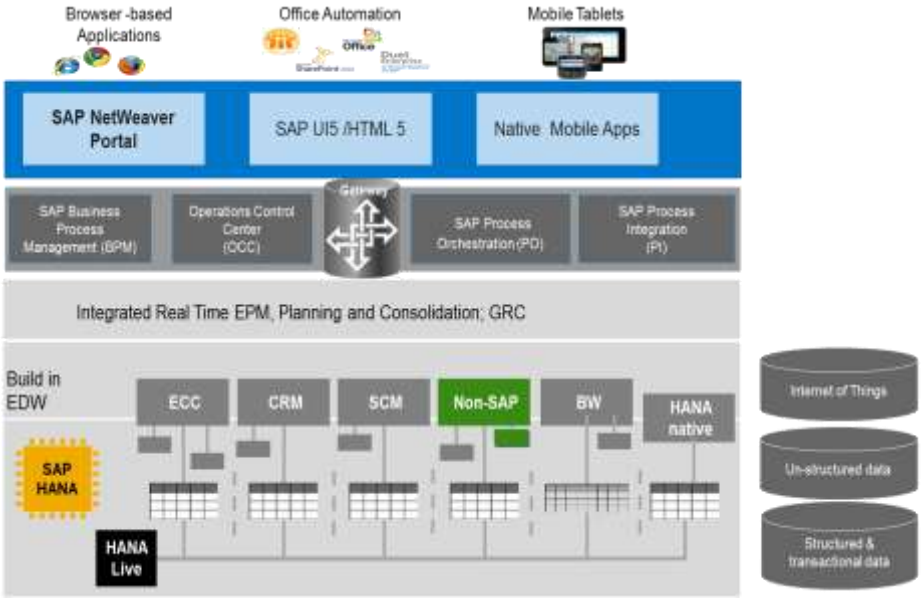

As a result, the IT landscape becomes much simpler. Different components are re-integrated into one consistent platform. The solution landscape is homogenized and standardized step by step, reducing IT effort, operational costs and risks. The SAP platform is based on SAP HANA, but also integrates other

component technologies and standards, e.g. for integration and orchestration tasks and for the provisioning of a state of the art user experience.

With the development of application innovations in all areas of the application space, SAP leverages the capabilities of this new platform. The majority of future SAP innovations will be a combination of both new applications and data-intensive execution of business functions pushed down into the core of the SAP HANA platform.

In this white paper, we will provide an overview of the different layers and components –from the database foundation at the bottom, to user experience at the top. Chapter two introduces the SAP HANA platform capabilities, giving a broad perspective from data foundation to user experience. Chapter three presents deployment options.

2 SAP PLATFORM – MAIN CAPABILITIES

SAP HANA replaces the traditional database in the technology stack, but it is much more than just a database – it is a platform providing database services, advanced computing capabilities, built-in high-performance functionality, and real-time data access between different applications and systems. In this chapter we will explain the main elements of this platform.

2.1 Software and hardware innovations

Hardware innovations and SAP HANA Software Innovations

In recent years a fundamental change in hardware innovation has taken place. In previous years, CPUs became faster every year, by increasing the clock rate. But we have reached the limits of physics – such rates of increase of clock speed are no longer possible. The chip manufacturers had to look for other ways of increasing computing power, and their answer is to increase the number of CPUs, and cores per CPU. In addition the price for main memory has lowered significantly. To efficiently leverage this new computing power, applications need to be adjusted:

- The applications need to be enabled to make use of many CPU cores – they need to be enabled to efficiently handle parallel processing.
- Data management needs to be moved from disk to main memory.

SAP HANA has been designed from the very beginning to use the new hardware architecture effectively:

- SAP HANA is an “in memory database”. Data processing runs entirely in main memory. Referencing data in main memory takes about 100ns, which is 100000 times faster than the average seek time for hard drives.
- SAP HANA is fully enabled for parallel processing, efficiently leveraging many CPU cores. The internal architecture is based on multiple levels of parallelization. For example, with SAP HANA, each column of a database table, or even partitions of a single column, can be analyzed by a separate CPU core. Many CPU cores can run in parallel, delivering the result much faster. SAP HANA uses hardware resources much more effectively compared to traditional database management systems, or applications relying on parallelization technology offered by the underlying operating system.

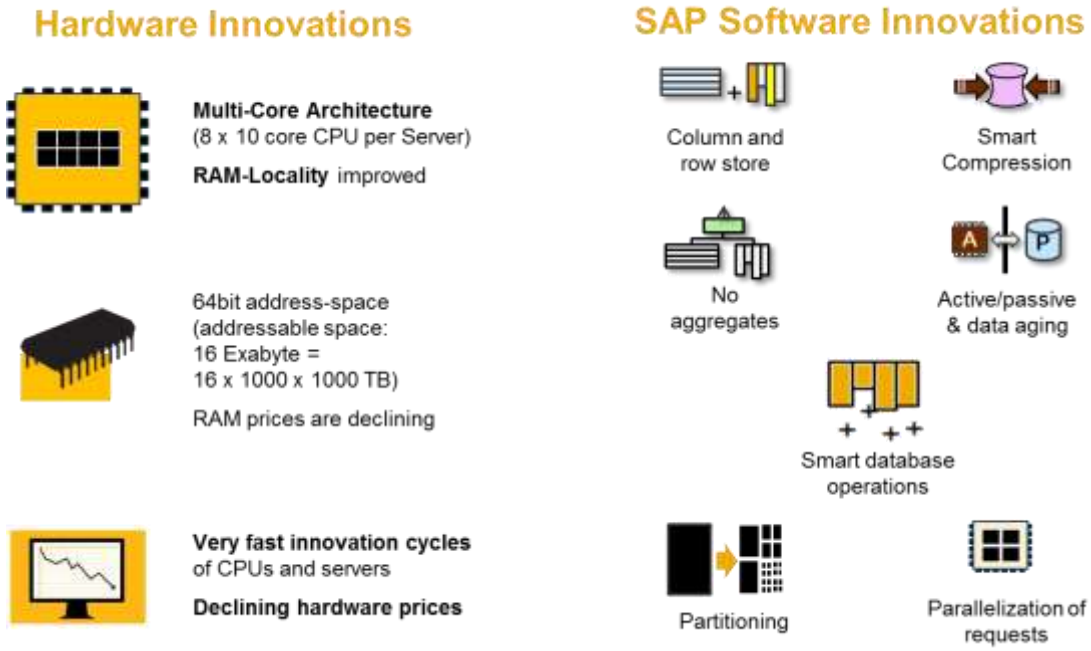

With the advanced parallelization and clustering capabilities of SAP HANA, it is possible to build high-end systems using commodity hardware. SAP HANA can achieve a performance on Intel commodity servers, which is much higher than the performance of today's high-end Unix servers. SAP HANA uses the commodity components to build a platform providing ultimate performance, scalability and state-of-the-art high availability and disaster recovery functions.

Placing the database into memory, just like other database providers do, is not enough. To allow software solutions to truly leverage the possibilities of the new hardware, the whole architecture of the platform needs to be designed accordingly. At the moment, only SAP HANA supports SAP applications as well as non-SAP applications to leverage the new hardware capabilities effectively.

SAP HANA server certification

The hardware servers for SAP HANA are certified by SAP. Why is this required? Because SAP HANA is aware of the new hardware technologies and makes explicit use of advanced architecture elements – e.g. location of data in the different types of CPU caches – it is important to ensure that only CPUs of the specified technology and types are used. For the overall processing, a defined ratio of memory size to CPUs also is important. SAP certification avoids issues of not-suitable hardware.

At the same time, the SAP certification is an open process, we have certified 9 server vendors so far.

2.2 SAP HANA architecture

This picture below shows the SAP HANA architecture. We will look at it from the bottom up.

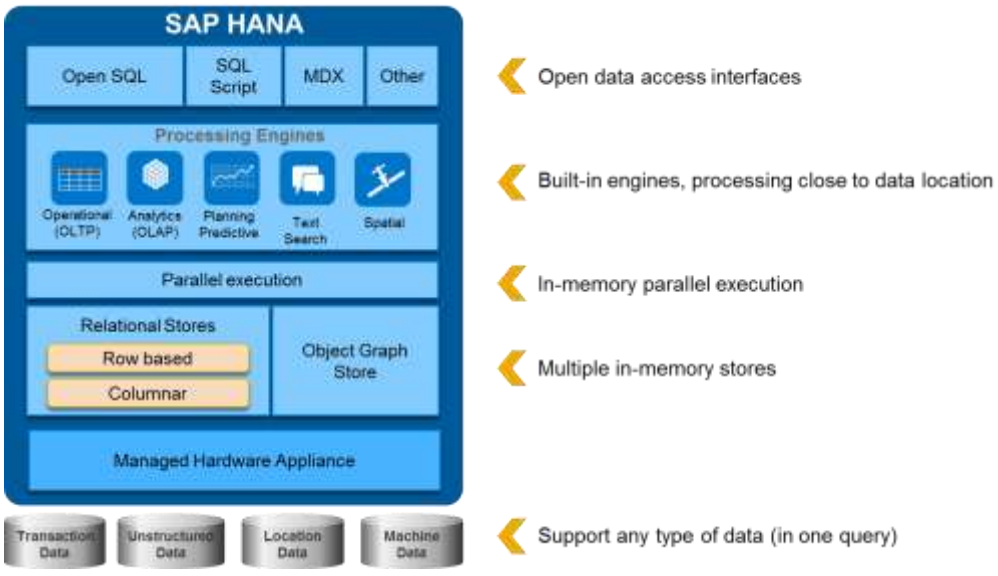

Support any type of data.

SAP HANA is an in-memory platform, but it can handle and import any data format. For many use cases – especially in the area of "Big Data" and "Internet of Things" – this is very important. Specialized import libraries import a very broad variety of structured and unstructured data formats, efficiently and quickly. The processing engines included in the platform allow an efficient analysis and manipulation of these data formats. For example, it is possible to execute a query across multiple data formats in one single SQL statement.

Database layer

From the experience of tuning the high-end systems of the top companies in the world, SAP has determined that column storage has many advantages compared to traditional row storage, in most use cases. The main SAP HANA storage model is based on column store,, which accesses the data much faster, especially in analytical scenarios. SAP HANA still has row storage as well, for a small number of use cases.

SAP has implemented smart compression based on column store, which is much more effective than the compression algorithms of traditional databases. SAP compression is based on a combination of a dictionary and of sorting the data in the columns, not on an algorithm, which would consume computing resources. This reduces data volumes by at least a factor of four, compared to a traditional database.

This storage concept also eliminates the need for nearly all indexes. This reduces memory size and CPU load (the permanent update of the indexes, as in today's databases, is no longer needed).

Another element is the elimination of aggregates. Because of the speed of SAP HANA it is no longer necessary to create aggregates or otherwise modified copies of the data, and to store them explicitly in the database in advance. In an ERP system based on a traditional database, there are at least 3 copies of each data element. SAP HANA makes these calculations at runtime, which reduces the data volume considerably. It also provides a real-time view of the original data for all applications – a "single source of truth" for data. We no longer need batch jobs to create aggregates and modified data copies.

This also simplifies development – all functional modules can access the original data directly (enriched views of the data are just database views).

Parallel execution

Another important element of SAP HANA is efficient parallel execution and computing. Column storage plays an important role in these capabilities as well. A separate CPU core performs database operations on each column, so a query can run fully in parallel on multiple columns/database fields. For larger tables, SAP HANA also sub-partitions the columns to enable multiple CPU cores to work simultaneously on a single column. SAP HANA has a smart built-in workload management – it makes efficient use of the large number of CPU cores available.

The SAP HANA architecture is based on parallel usage of completely independent services/processes. These processes have no interdependencies – SAP HANA uses a so called "shared nothing" approach, with which very efficient scaling is possible. This is different from the approaches of traditional databases. Currently, SAP HANA configurations with 100 server nodes – 100 hardware servers, each of them running multiple SAP HANA Index Server processes – have been tested. A limit of scalability has not yet been detected.

There is a long list of SAP HANA innovations driving performance and scalability. It also contains features to handle inserts, updates, etc. intelligently.

Integrated processing engines

In addition to providing database services, SAP HANA also provides the possibility to run code directly in the platform. The principle is to run data-intensive operations close to the data. Before SAP HANA, calculations have typically been done on application server level (and there was no other chance to do so, because otherwise the database server would have been the bottleneck of the entire system). Sometimes this means that large amounts of data need to be transferred to the application server, where the analysis and calculation take place. And often, in the end, a very small result set is generated. With SAP HANA, calculations, which are "data-intensive", are executed directly in the SAP HANA database where the data resides. This dramatically reduces data movement, hardware capacity on application server level and boosts the overall performance of an application.

SAP ships so called "engines" representing re-use functionality, following this approach. There are already a number of such SAP HANA engines, including engines for text analysis, predictive analytics, planning, OLAP, etc. SAP will develop further engines over time, but you can also create custom code and push it down into the SAP HANA database. Future SAP applications and custom code will be a combination of

business logic on the application server level and re-use engines and code fragments executed directly in the SAP HANA database.

One example of such a SAP HANA engine is the **SAP Planning Engine**, which works directly on current data and plan data, and combines it with master data. This enables faster planning cycles and a new quality of decision support. Through real-time access to plan, budget, forecast and actual data, plus the integrated simulation capabilities, business users and managers gain a new level of transparency – across all layers and in real-time. While in the past decisions have been based on the results of planning runs from the previous day (or at least from several hours ago), decisions can now be based on real-time data.

Open data access interfaces

SAP HANA contains numerous innovations and optimizations which make it a new generation of platform. At the same time, it provides a broad range of open standard interfaces for applications. Starting from Open SQL or ODATA/JDBC/ODBC/JSON for data access, there are interfaces for .Net connections, for the Java world and for the World Wide Web.

So SAP HANA is not only a platform for SAP applications. You can run any 3rd-party or custom applications on SAP HANA. These applications will benefit immediately from the SAP HANA capabilities, and over time they can be optimized to leverage all HANA capabilities even more (e.g. by push down code, engines and any type of data source).

2.3 Real-Time Platform

We have discussed the scaling and the data duplication avoidance capabilities of an SAP system. We Now move up one level and look at another very important capability of SAP HANA – the ability to function as a central platform running multiple SAP and non-SAP systems in parallel.

Using the so-called MCOD feature (Multiple Components One Database), it has also been possible with traditional databases to install multiple SAP systems on one database instance, but only for small and non-production systems. The performance and scalability capabilities of SAP HANA change the situation, and SAP HANA also offers one key feature which makes this deployment option very attractive for business applications: with SAP HANA it is possible to access data of different systems – in a strict security controlled way.

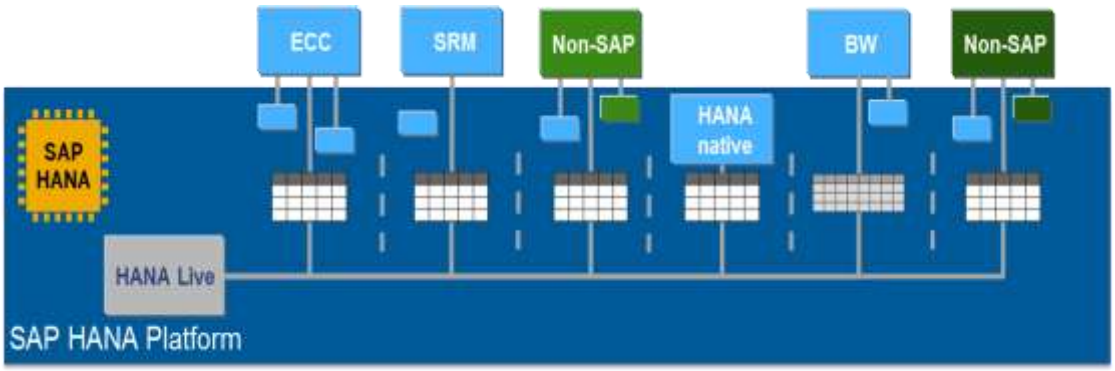

This feature eliminates data duplication between systems. For instance it is now possible to access the tables of an ERP system directly from a “built-in” BW system running on the same SAP HANA instance. This eliminates the ETL processes, but most importantly, it enables a new quality of reporting, which will be explained later in this chapter

Looking at business processes crossing multiple systems, another area of benefits becomes obvious. With SAP HANA, it is now possible to write applications providing an end-to-end business process view, leveraging the original data in real-time.

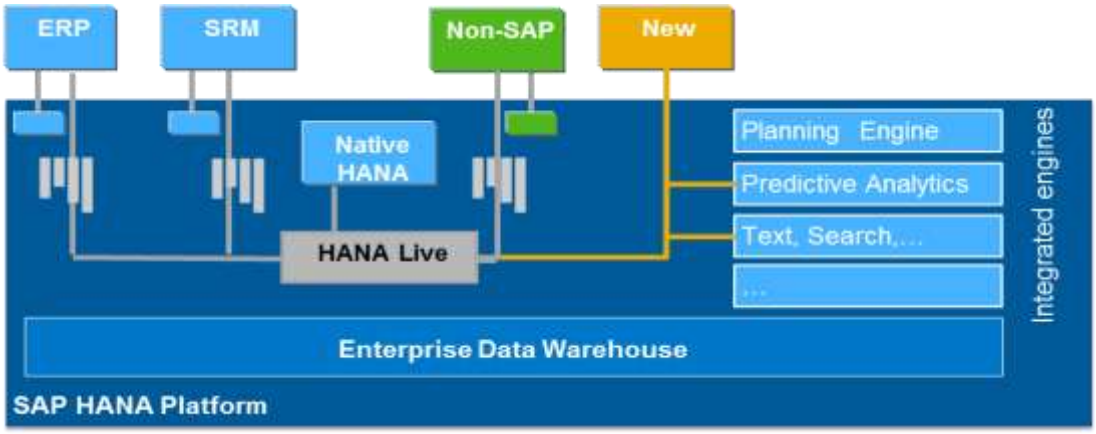

Multiple systems running a cross-system business process can now reside in one SAP HANA instance. This enables a very efficient integration and orchestration of processes. These applications can access the data of all the systems involved in a given business process in real-time and use the integrated SAP HANA engines to execute analysis, planning or simulation activities, in real-time.

One example is a new SAP standard application for cross-company reconciliation, which is currently being developed. It automates and controls the end-to-end intra-group invoicing, payment and settlement activities. It can now access the original data in real-time, so it systematically avoids inter-company differences. It also increases transparency across all units and speeds up financial closing. This is just one example among many.

Another new capability is Real-time Integrated Decision Support – providing an immediate answer to any question.

The left-hand side of the picture below shows the traditional data flow when reporting is required. There are many layers and steps involved. To fully build and implement such a process it not only costs time and effort but requires as well multiple systems, multiple technologies, additional applications and increased TCO. If it is required to react to new business requirements later, multiple layers and steps may need to be adjusted. This limits the agility of business and IT.

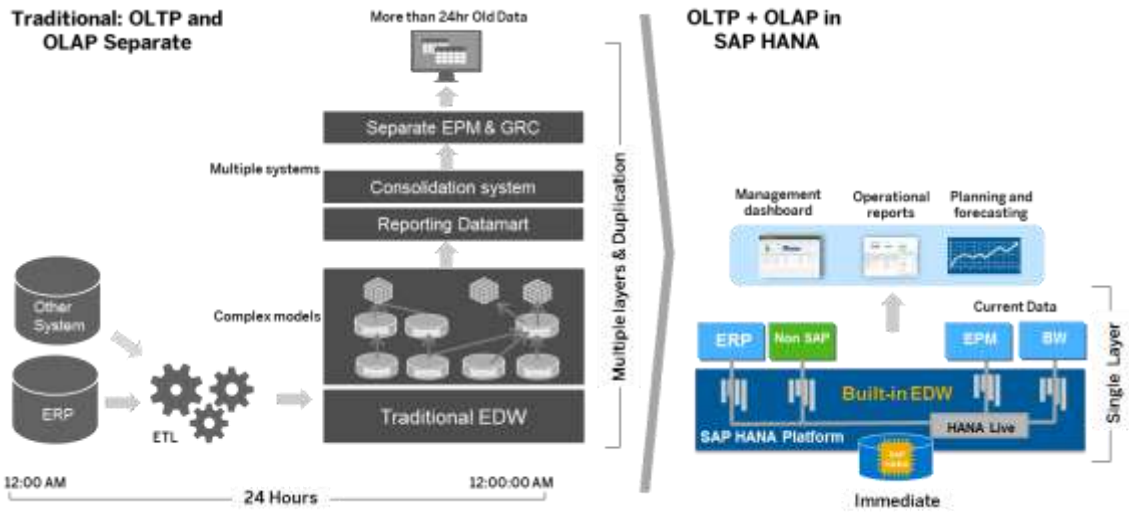

SAP HANA greatly simplifies the overall architecture. All the data resides in one SAP HANA, and the BW system can access it without loading or duplication. The data model is also simplified – there is no longer any need to pre-calculate and store intermediate results and aggregates. The fast execution capabilities of SAP HANA allow calculation in real-time.

Operational reporting can also be implemented directly in the SAP HANA platform, leveraging “HANA Live”.

The most important effect is on decision support. It is now possible to report on the original data in real-time, so business decisions are based on current status, not on the status of the previous day[s]. It is now also possible to drill-down to line item level. With the traditional BW ETL process, only a subset of the data (usually summarized) was duplicated to BW, which limited the drill-down capabilities. You can now drill down to the full details.

So decision support is not only faster, there is a new quality of decision support – real-time and with full precision.

2.4 Big Data and the Internet of Things

SAP HANA can manage huge amounts of data at high speed, based on the SAP HANA Big Data processing framework, which covers the complete collection, storage, processing and presentation chain.

- **Collection**. SAP HANA collects data from various interfaces, from various sources and in different formats, including stream data via the SAP Sybase Event Streaming Processor, and data from *Hadoop* in the background.
- **Storage.** The heart of data storage is SAP HANA. The SAP HANA data storage capabilities – column store plus effective parallelization of all requests –handle huge amounts of data efficiently. An intelligent data-aging concept moves old data to Near Line Storage, based on Sybase IQ, completely transparently for the applications.
- **Processing**. Data processing uses the capabilities of SAP HANA. Running data-intensive operations directly in the platform, close to the data, and the massive parallelization capabilities, processes data of different formats and in huge quantities, efficiently.
- **Presentation.** The SAP Business Objects BI platform and real-time dashboards primarily handle presentation.

Internet of Things scenarios are based on efficient Big Data handling, and integrated functions and engines like the Sybase Event Streaming Processor ESP and Machine-to-Machine Services and SQL Anywhere for synchronization of mobile devices and machine data”. The following drawing depicts the high-level architecture of such a scenario. The core of the solution is the SAP HANA platform in which all the data is collected, stored and analyzed.

Various design tools support the development of custom-specific scenarios. Open interfaces and agents connect a broad variety of devices and present the results on devices with any type of UI– mobile, webbing or tablet.

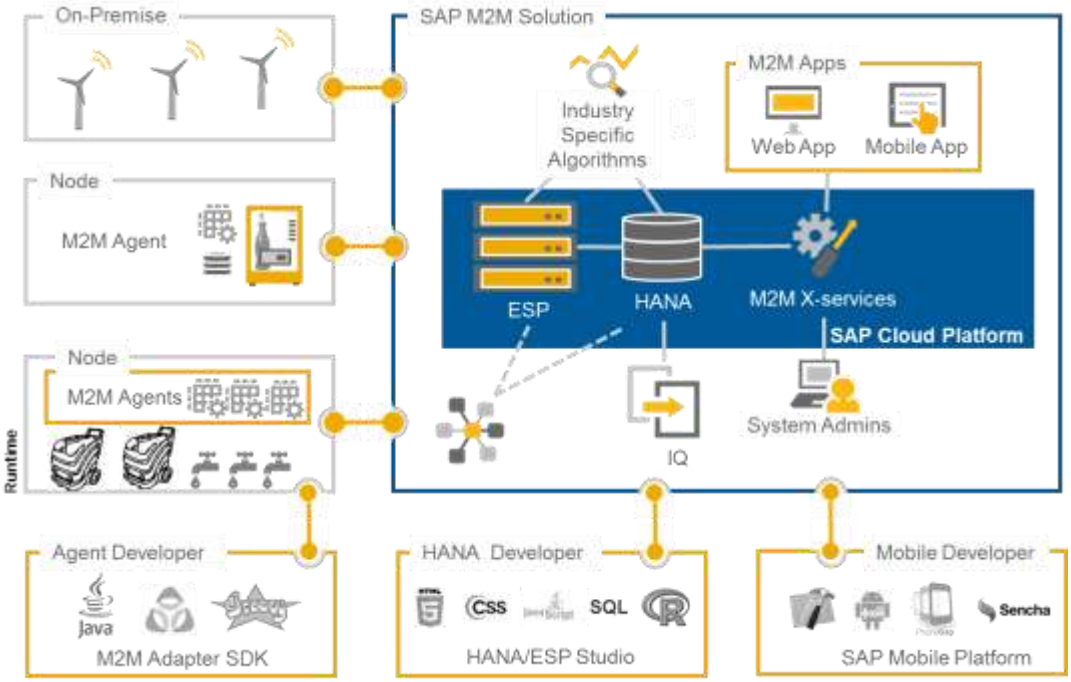

The number of machine-to-machine integration scenarios is exploding. Many of them combine business data with physical data to create a basis for decision-making that was not previously available. Others rely on the ability to manage large amounts of data in real-time and to correlate data of different sources and formats in single real-time operations. These new possibilities are not only about making existing business processes faster, but they allow new business functions and opportunities to evolve, which were previously not possible or unimaginable.

Similar to developing standard scenarios and best practices for many industries and LOB's, SAP is engaged with many companies in developing new integration scenarios under NDA – scenarios seen by these companies as differentiating them highly in the market.

2.5 Orchestration Layer

Today's solution landscape consists of a multitude of individual components. SAP application components interplay with legacy systems, services and individually developed applications. Organizational boundaries and deployment strategies and legal entity considerations increase the need for a distributed landscape, but all business processes still need to be designed and managed end-to-end.

SAP provides the respective tools and best practices to manage such a solution landscape and all its business processes. The following integration paradigms highlight the different types of integration.

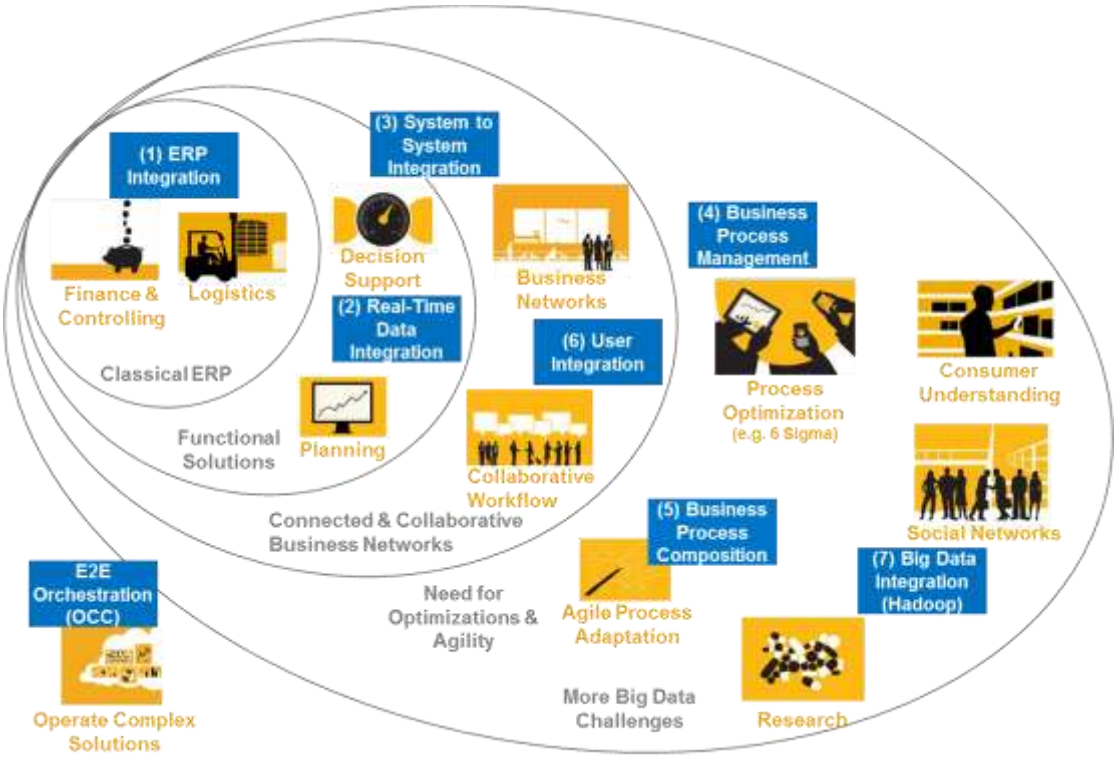

Efficient, scalable, flexible and stable handling of the different integration tasks is key to success of the overall solution! With the SAP platform, the different integration types are **efficiently covered in the layer where they belong**:

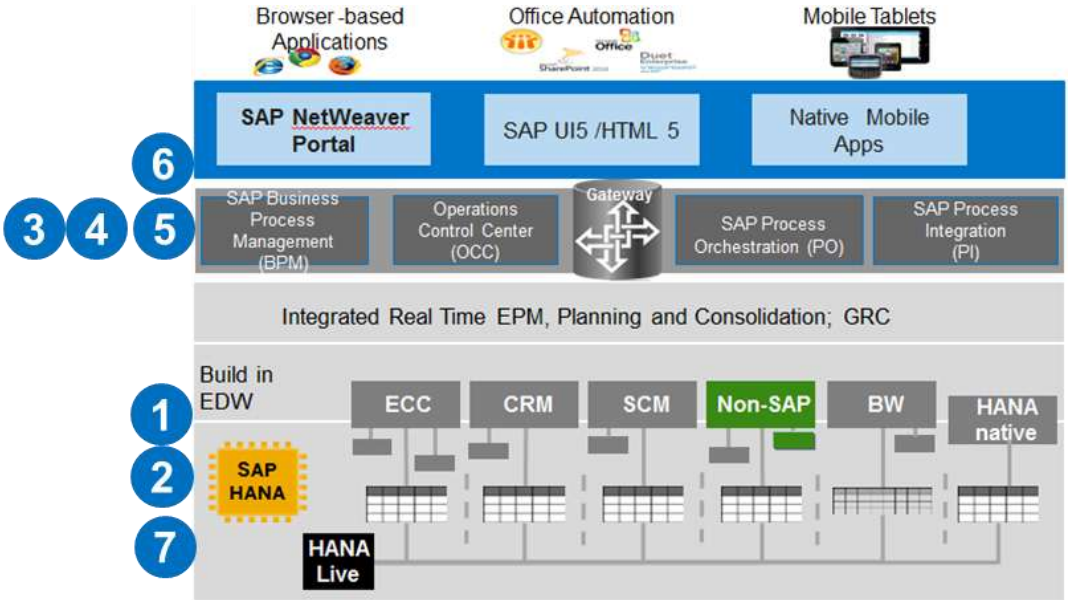

(1) ERP Type Integration. Built-in integration is one of the key advantages of SAP application solutions like ECC, CRM, SCM and SRM. Every posting, e.g. a posting a production order, consists of the posting in the dependent function modules of the same system (automatically) and in related systems (via automatically generated messages). Related postings are triggered automatically, without further manual or paper-based interaction. The principle that a transaction either completes

in all or none is fundamental to all transactional systems, but in a distributed environment, transactions can succeed partially, so aborted or incomplete postings need to be recorded and either re-processed or completely removed.

(2) Real-time data integration within the SAP HANA platform. While today SAP and non-SAP systems are integrated via messages, the SAP HANA platform integrates them seamlessly, directly within the platform. This type of integration replaces data duplication step-by-step, and reduces the volume of messages exchanged between the systems dramatically. The first use cases of this integration within the platform are in analytical scenarios – using HANA Live to access the data of other systems in the same SAP HANA instance – and in new applications and engines, like the SAP Planning Engine, leveraging the access to the original data to process cross-system functions.

(3) System to system integration. Integration scenarios which can't be handled by the built-in integration in the ERP systems, or by the platform itself, need to be addressed by an explicit orchestration layer. Consider a warehouse management solution which interacts with the ECC-based order to cash solution. In today's solution landscapes, there are many examples in which orders are sent back and forth numerous times. In contrast to the transaction above, these interfaces are message-based, e.g. sending an IDOC to another system, SAP as well as non-SAP. SAP PI manages the flow and delivery of messages efficiently. Interface adapters integrate into standard application packages and custom developed applications based on standards..

(4) Business process Management. A typical business process executes on a variety of solution components. Even within one system, several departments (LOB) may be part of the business process. Workflow and sign off procedures need to complement the business process execution. Consider, for example, a contact or payment release. It is mandatory to record who released what, when and where. Any delay or other exception that occurs during execution must be reported as visibly and quickly as possible. SAP Business Process Management (BPM) and SAP Solution Manager are the primary management systems, where you can integrate workflow processes, SOA-Services, comprehensive business process monitoring and guided procedures for efficient exception management, including task, incident and problem analysis and resolution.

(5) Composition across several systems or components. To better manage risk and performance aspects of a business transaction, some industries split the business transaction into a number of sub-entities executing on different components. The banking and insurance companies are typical examples. Consider a banking solution with several systems implementing loan, deposit management and customer master data. Transactions need to be managed across applications, and perhaps even cross LOBs. A business transaction is mapped to a sequence of services offered by the different solution components. The process orchestration layer (POL) owns the definition and execution environment for the mapping and execution sequence. It also manages all postings in an all or nothing fashion, and all exceptions that arise with partially completed sequences.

(6) User integration – Portal and Workspaces. A process orchestration layer would be incomplete without integration of the SAP systems interfaces into one comprehensive workplace. Users connected to a solution want to observe their key indicators and to-do lists at a glance. For each task item, they need to know what to do, and how to do it. A landing page, with key indicators prominently visible, an intuitive list layout of tasks to do, and clear and easy guidance to how to perform them, are today's working paradigm. The SAP portal is a good example of how landing pages, task and guidance based on the individual's role and capabilities, deliver a unique and pleasant user experience. Standard UI technologies, such as HTML5 web pages with JavaScript enablement, allow the development of user interfaces that run on multiple system platforms and form factors. SAP FIORI is a good example of how such applications can be developed to not only provide a pleasant user experience, but also to deal intelligently with the various form factors like smart phone, tablet and desktop.

(7) Big Data integration. SAP HANA is an in-memory platform. At the same time, SAP HANA provides the possibility to handle and import any data format. For many use cases – especially in the area of "Big Data" and "Internet of Things" – this is very important. Specialized import libraries allow for a fast import of a broad variety of structured and unstructured data formats. The processing engines

included in the platform support an efficient analysis and manipulation of these data formats – in one single SQL statement you can retrieve and combine data of different formats from different sources. These capabilities are the foundation for better understanding your customers – e.g. by leveraging social media data.
In addition, with the SAP Big Data processing framework, SAP HANA offers a seamless integration with Hadoop.

To achieve the highest level of business process orchestration and realize the greatest business value for end users, SAP delivers the technology foundation, the design, deployment and management tools, and a set of best practices for deploying and operating such an environment. Key indicators for successful setup are the cost of implementation and the cost of operations, and the effort to manager data consistency and maintenance.

2.6 User Experience Layer

User experience is concerned with all aspects of the end user's interaction with products, systems, processes and services. The way users work with software solutions have changed dramatically over the last few years. Addressing different device form factors and providing simplified and consumer-grade experiences are among the top priorities for application development. A key requirement for implementing improvements of the user experience is that it can be done non-disruptively.

On top of the orchestration layer, the SAP platform will also help you to achieve a superior user experience based on the latest industry standards. This section will briefly outline the core elements of this important layer.

- **HTML5 Industry Standard and device-agnostic development**. At the heart of the user interface is the industry standard HTML5. SAP UI5 is SAP's HTMPL5 UI technology best suited to building user interfaces across different devices. It is the strategy chosen by SAP for supporting multiple devices and form factors. As a device-agnostic development approach, all applications built on it benefit from common elements such as universal HTML client, landing pages, event-driven working model, and common search and decision support facilities.
- **OData and Gateway**. The SAP gateway is the interface to the application, providing a user interface to support cloud, social, mobile and desktop applications at the same time. Based on the open industry standard communication protocol ODATA, it is especially easy to integrate 3rd-party applications. ODATA provides a lean text Interface for structured transactional and analytical data.
- **Fiori Experience Builder**. Fiori applications are HTML5-based applications built on top of the SAP Business Suite, enabled for OData consumption through the SAP Gateway. A set of common design methodologies, such as the model view controller paradigm and predefined UI controls, tracing and logging support, enable an efficient, device-agnostic UI development strategy.
- **SAP Portal, SAP Business Client, Personas**, **Worklist and guided procedures**. SAP offers a rich set of tools and extensions to current products to help to build a personalized workspace. The side panel in the SAOP business client, and floor plan managers for Web Dynpro and WEB UI frameworks are the technological basis. SAP Screen Personas help to redesign SAP GUI screens, beyond the simple hiding of fields, enabling a focused and slicker user experience.

SAP recommends that customers define priorities for their own user experience strategy, based on what they require from SAP software. Business and IT departments are responsible for jointly developing this roadmap. Detailed guidelines and documentation are available from SAP.

3 SAP PLATFORM – TECHNICAL DEPLOYMENT OPTIONS

3.1 Cloud Deployment Options

Before we discuss the new possibilities for cloud computing with SAP HANA, we would like to clarify what is meant by the term “cloud” in this context. Cloud computing provides the means through which the main IT Infrastructure capabilities — from computing power to infrastructure, development platforms, applications, business processes to personal collaboration — can be delivered as a service, wherever and whenever it is needed. Cloud computing has some essential characteristics like elasticity – the ability to scale up and down, depending on the requirements – resource pooling and automation.

The main objective of cloud computing is to provide a scalable, flexible and yet cost-effective infrastructure, ensuring the best possible service for the application components, while using the underlying resources efficiently. Another view of cloud usage is the location and ownership of the cloud. The following picture provides an overview on the existing models.

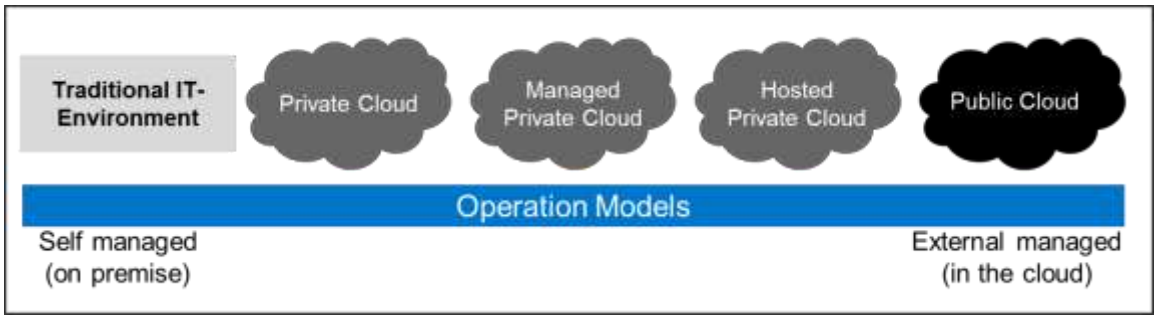

At a high level, there are five patterns, characterized by location and management responsibility:

- **Traditional IT Environment:** This is the environment from which most of our customers come – a static landscape installed on physical servers and storage landscapes, running fix sized DBs, managed by the local IT organization. Advanced implementations already deploy a virtualized server and/or storage landscape, but there is generally limited flexibility to react to quickly changing business needs.
- **Private Cloud:** Private cloud has applied the full set of elasticity and flexibility options in the customer’s data center, and is fully under the control of the customer. The cloud is managed by the customer’s IT operations team, and all implemented resources are reserved for the customer’s operations.
- **Managed Private Cloud:** The technology is a private cloud. The landscape is still in the customer’s data center, but an external partner runs daily operations and landscape evolution and planning.
- **Outsourced Private Cloud:** The outsourced private cloud is the cloud extension of hosting. The private cloud is moved to the outsourcing partner’s data center and is managed by the partner.
- **Public Cloud:** The public cloud is a fully shared environment. In certain cases, several users or customers may share the same physical resources. It is highly standardized and usually used in a pay-per-use model.

SAP offers several cloud deployment scenarios for SAP HANA:

- **SAP HANA One:** SAP HANA One is offered via public cloud providers. The target customers for this offering are start-ups, ISVs and teams implementing native cloud scenarios. Currently, appliance memory sizes of up to 64 GB are available with SAP HANA One. The SAP HANA usage is offered in a pay-per-use subscription model – for application development and to run production systems.
- **SAP HANA Enterprise Cloud:** The SAP HANA Enterprise Cloud is a hosted private cloud offering. Customers can run their mission-critical systems based on SAP HANA, in a private cloud, hosted in

one of the SAP data centers. SAP onboarding services ensure smooth and efficient migration of your systems into the SAP HANA Enterprise Cloud.

- **Private Cloud for SAP HANA:** For various reasons, some customers are not willing to run their business systems in a hosted environment. To enable these customers also to leverage the benefits of cloud, like elasticity, flexibility and efficient resource usage, we have defined a concept, together with our hardware and technology partners, for a private cloud based on SAP HANA. Best practices methodologies and tools support the transition of SAP and non-SAP systems into this private cloud. If required, this process can also be managed entirely by SAP. SAP also provides detailed operational practices – empowering your IT organization to run the overall solution safely.

This white paper focuses on the private cloud for SAP HANA. This approach provides the flexibility and cost savings of cloud computing based on SAP HANA, while keeping the data and control of the solution inside your data center. SAP can also manage this cloud – providing a managed private cloud as a service.

3.2 Private Cloud based on SAP HANA

SAP HANA is a new technology. We recommend that you take this opportunity to define and build a state-of-the-art technical infrastructure solution – for SAP HANA and for all other applications. A key element in this concept is **simplification** – simplifying the overall solution landscape, replacing the large number of individual technology components by an integrated platform. The second aspect is **standardization** – standardizing the underlying hardware infrastructure. This simplifies scaling, management and reduces cost. The third element is **scalability**. The architecture needs to be flexible, to be able to react to new and changing requirements. It remains important to use hardware resources efficiently – keeping cost under control, and business continuity – high availability and disaster recovery – need to be ensured, fulfill your SLAs. Last, but not least, it is required to deploy the solution to the SAP HANA cloud quickly and safely. To ensure this, we apply proven methodologies and best practices. The methodologies are pre-defined and based on templates delivered via the SAP Solution Manager.

3.3 Architecture Overview

The main objective of our concept for the private cloud based on SAP HANA is to greatly reduce the risks for you to move to this new platform. The architecture combines the SAP HANA platform capabilities with industry standard cloud operations. Based on the experience of working with over 40.000 customers around the world, SAP Active Global Support has defined a concept for a private cloud, reflecting proven best practices. Together with our hardware and technology partners, we can deliver the solution. The SAP private cloud comprises all required components in the solution landscape – not just the SAP HANA database. The figure below outlines the main principles and building blocks.

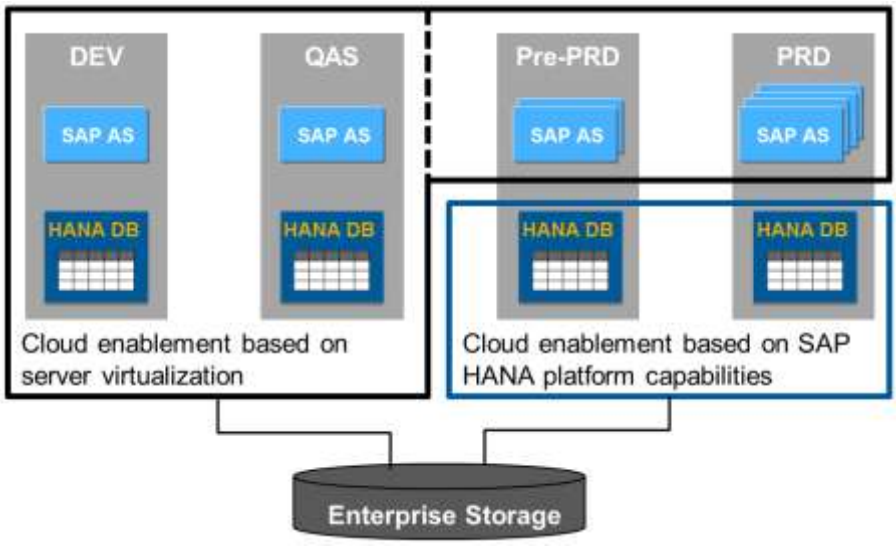

Enterprise Storage

The basis of data storage is a powerful enterprise storage system. The storage is not integrated directly into the SAP HANA servers and appliances. Many SAP HANA systems share one central enterprise storage system, which provides features to simplify the operations of the cloud, like snapshots, seamless scale or data replication functions.

How does this fit to the SAP HANA appliance concept?

By default, SAP HANA appliances come with integrated storage, but since SAPPHIRE 2013, SAP HANA appliances can be deployed via a tailored data center integration approach. Certified servers can be connected to existing enterprise storage systems. SAP has defined KPIs for the storage subsystems – like I/O throughput and latency – ensuring that the overall infrastructure forms a SAP HANA-ready environment.

Enterprise storage systems offer a number of functions which are important for efficient cloud operation – like snapshots, backup connections, or system provisioning functions, so for deployment in a private cloud we recommend a powerful enterprise storage infrastructure, not using SAP HANA appliances with integrated storage but leverage a powerful enterprise storage infrastructure. In many companies such infrastructures already exist – they can be extended and reused.

Operating System level Virtualization

For the application servers we use a state-of-the-art virtualization technology allowing efficient use of server resources, and providing application server relocation functionality. The priorities for non-productive instances are different from production instances. A production instance requires availability, performance and scalability. In non-production instances it is more important to copy, move, rebuild or change systems. These requirements are supported effectively by virtualization technologies.

We recommend virtualization technologies for the non-productive systems – application servers and the underlying SAP HANA database instances – as well. SAP has a close partnership with VMware. VMware virtualization can also be used for the non-productive SAP HANA instances, and these virtualization technologies ensure that the underlying hardware resources are used efficiently.

Management tools

One aspect of an effective private cloud is the availability of flexible and powerful management tools which support the extension or adoption of the cloud structure according to changing needs. SAP uses a tool portfolio which combines the VMWare tools with the SAP Landscape Virtualization Manager (LVM) and SAP Solution Manager.

SAP HANA platform capabilities for virtualization and resource management

For productive SAP HANA instances, we use another technology which reflects their requirements. SAP HANA has integrated capabilities that can manage multiple application systems in one single SAP HANA instance – using a workload management, which is application-aware and leads to very efficient virtualization of the underlying hardware resources.

SAP HANA platform capabilities in more detail

SAP HANA is much more than just a database engine – it is a platform for application development and also for running mission-critical high-performance applications in a in a scalable and cost-effective way.

As shown in the following graphic, the SAP HANA platform provides an integrated cloud layer, running database services for SAP and non-SAP applications, in parallel. SAP HANA leverages a pool of servers and an underlying storage infrastructure, and runs multiple applications of different types with intelligent workload management – on platform level.

This is the SAP vision. First elements of this vision are already available; others will follow in the next support packages. There are already a number of customers implementing/running BW with ECC or CRM, on one SAP HANA instance, leveraging the real-time data access on platform level.

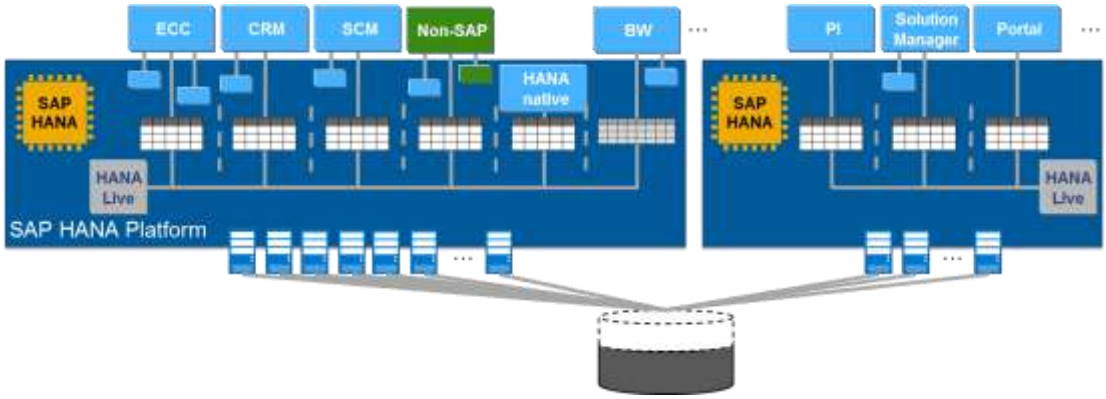

SAP HANA provides cloud characteristics like elasticity and flexibility on platform level. It enables much more efficient virtualization compared to the classical approach in which a hypervisor acts on the operating system level, underneath the applications.

3.4 High Availability (HA) and Disaster Recovery (DR)

The SAP HANA platform provides integrated features for HA and DR. Combining these features with the capabilities of the underlying IT infrastructure and the storage subsystem allows you to build powerful solutions, reflecting the requirements of a company cost-effectively.

How do we define the right HA/DR solution for your requirements?

Discussions often start with features and comparisons of individual functions straight away, but this is a mistake. A structured approach is required to define an effective HA and DR concept, otherwise a solution may be implemented which does not fulfill the requirements, or over-fulfills them, leading to much higher cost.

The first step is to understand and document the requirements per system. This process requires considerable effort, including the need for confirmation by the business departments.

It is Then important to understand the boundary conditions with respect to the data center strategy – how many data centers there are existing, and what are their roles.

In the last step, we look at the existing capabilities for HA and DR, and map them onto the requirements. This defines an infrastructure concept that fulfills the requirements cost-effectively, and documents the decision path – explaining why the solution was selected.

Definition of requirements and KPIs

As a starting point, categorize the existing systems into different HA and DR classes. No company will implement the most expensive HA/DR solution for all of their systems. Define in a second step the KPIs for each of the classes differently for high availability and disaster recovery. High availability addresses the most frequent cases of system outages – hardware failures. Typically, HA is implemented in one data center and the KPIs are high – very short system fail-over times (RTO), and no data loss allowed (RPO = 0).

Looking at the hardware capacity that needs to be available in case of a hardware failure, normally the full capacity is required. This means that after a hardware failure the business processes just continue to function without any limitations. The disaster recovery situation is different. A physical disaster happens much less often. In a DR case, companies can typically accept longer outage times, and some data loss for some of the systems is acceptable (data can in most cases be reconstructed from other sources). It is

important that the result of this decision is documented, agreed with the business and clearly communicated to all parties involved.

Boundary condition: Data center setup

The data center strategy is an important boundary condition, determining which technical solution can be implemented. The most common cases are listed below:

a) **Two data centers located close to each other (<~50 km).** This allows synchronous replication of data between the data centers, but does not protect against a metro-level major physical disaster.

b) **Two data centers far away from each other.** This setup provides a protection against physical disasters, but in such an environment it is more difficult to keep the systems in synch – because of the distance and the related latency. Synchronous data replication is not possible.

c) **Three data centers (two close to each other and one further away).** This data center setup offers the best failure protection. The two proximate data centers can be leveraged for fast fail-over, while the third one provides protection against major physical disasters. There are different ways of setting up a 3rd data center, depending on recovery time requirements:

   i. Basic data protection: backup devices/tapes, e.g. electronic shipment of backups and log files

   ii. Online data protection: The 3rd data center contains storage systems and the data is permanently replicated there asynchronously. In case of a major physical disaster, the data is protected but the systems need to be re-established in the 3rd data center, which leads to a considerable downtime.

   iii. Data and infrastructure protection: The 3rd data center contains both storage and server hardware. Data is replicated asynchronously.

There are a number of possible technical solutions for each of these cases. We will describe some common solutions later.

SAP HANA capabilities for HA and DR

All features of SAP HANA for HA/DR are in the platform and available to all customers with no additional cost.

1) Local HA functionality – SAP HANA host auto fail-over

This function provides protection against server hardware failures. SAP HANA permanently monitors the status of all processes. If one server node fails, SAP HANA automatically shuts it down and makes its data available to one of the configured standby nodes. The standby node reads the data of the failing node from the shared storage infrastructure, and takes over the role of this node. There is no loss of data or of committed transactions.

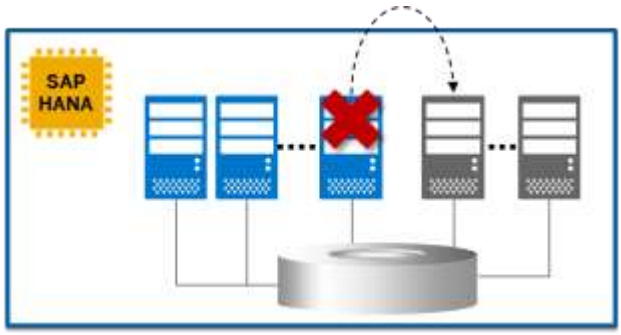

This HA function works only for scale-out configurations. This functionality does not protect from failures of the underlying storage subsystem.

2) SAP HANA system replication – for HA and DR

The SAP HANA system replication feature is a database-level replication implemented inside SAP HANA. A primary SAP HANA system replicates all changes to a standby SAP HANA system (a complete system, including storage), This replication is possible in synchronous as well as in asynchronous mode. In synchronous mode, it is guaranteed that there is no data loss in any hardware failure. If the primary system goes down, the standby system takes over and only not-committed transactions are rolled back. Synchronous replication is possible with 2 proximate data centers.

For a replication between 2 distant datacenters, the asynchronous mode needs to be used. With an asynchronous replication a small data loss is possible (depending on the network latency this can be very small).

This approach achieves very short fail-over times, because the standby SAP HANA instance is already preloaded during the replication process. Fail-over times are are typically less than 5 minutes, independently of the size of the database. This approach also provides protection against block corruption in the storage system.

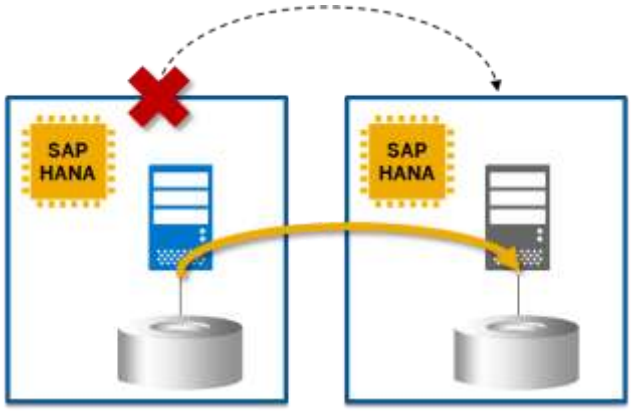

So this approach combines the shortest possible fail-over times with the highest possible protection. It is integrated into SAP HANA, it is the SAP strategic, and will be further enhanced and extended in the future. Depending on the configuration, the standby database server hardware can be reused for non-production systems.

3) Storage-level replication

Storage replication features are also often used for DR, because replication on database level (like SAP HANA system replication) is an expensive add-on feature for traditional databases (with SAP HANA there is no additional cost), and complex to set up.

Storage level replication has some limitations compared to SAP HANA system replication. It does not protect from block corruption on storage level, and fail-over times are considerably longer – because in case of a failover the data needs to be loaded from disk into main memory. A typical motivation for using this feature would be to re-use an existing storage replication infrastructure in a company. The hardware of the standby database server can be reused for non-production systems.

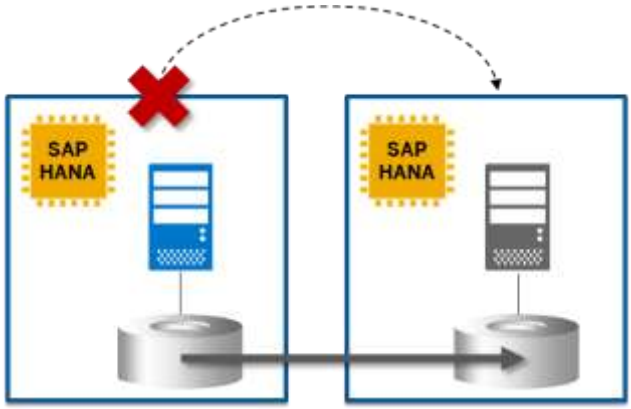

Other approaches are also possible, e.g. disaster recovery by means of data and log backup shipping. We suggest to conducting a dedicated workshop to evaluate the possible approaches and select the best solution for your company.

3.5 Typical HA/DR solutions

Typical configuration with 2 proximate data centers (< ~50 km apart)

This setup uses two proximate data centers with a stable network connection. It is based on synchronous SAP HANA system replication, covering both HA and DR. It achieves very short fail-over time (RTO), with no data loss (zero RPO), and it is a very cost effective solution.

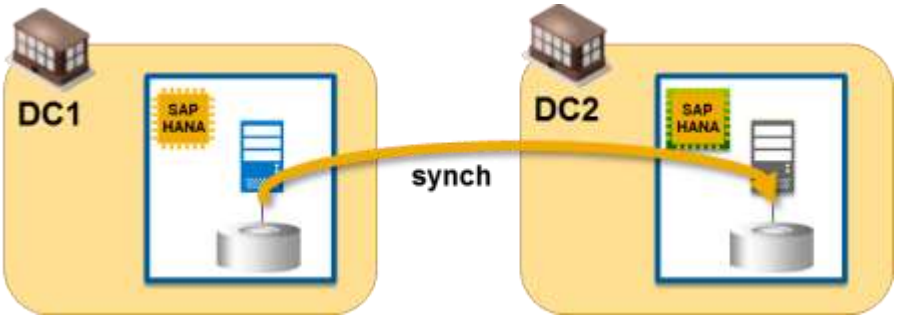

Typical configurations with 2 proximate or distant data centers

The examples below can be implemented in two proximate data centers or two data centers which are far apart. The difference between the two cases is the replication mode. While with 2 proximate datacenters it is possible to use a synchronous replication mode, for 2 remote datacenters only an asynchronous replication is possible.

a) **Separate HA and DR with storage replication**. This approach provides very good HA in the primary data center (very short fail-over times, protection against server hardware and storage failures, zero data loss). Data is replicated to the secondary data center by storage replication. The assumption for this selection is that there already exists a working storage replication between the data centers.

In the case of two distant data centers, the storage level replication functions in asynchronous mode what leads to a certain amount of data loss in case of a physical disaster. Asynchronous storage replication is not yet supported, it is currently being tested with an SAP hardware partner.

b) **Separate HA and DR with SAP HANA system replication**. In this configuration, SAP HANA system replication is used for both HA in the local data center and DR to the distant data center. As with option a), for the local HA, the best possible protection is provided. SAP HANA system replication is used for DR as well, providing a higher level of safety than storage replication.

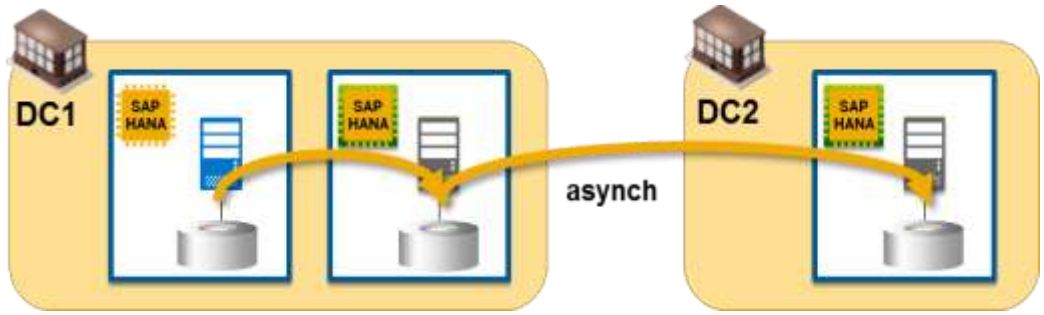

SAP HANA will support two replications (one local and one distant) from SP7, at the end of 2013. This is the best solution. It provides optimal HA in the primary data center (DC) and the built-in SAP HANA features are used for DR – very short fail-over time and protection against storage failures and block corruption.

Example configuration with 3 data centers

Technically, all of the above solutions for 2 data centers can be extended to integrate a third one. The example shown below is a pragmatic approach, leveraging the advantages of having two proximate data centers and using the third data center just for data protection.

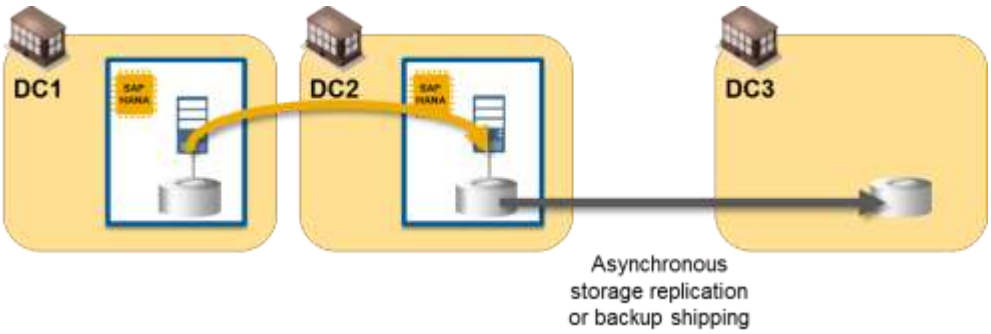

For the 2 proximate data centers, the HA functionality is again stretched across the two locations – combining HA with local DR in a pragmatic way. The short distance between the data centers ensures that fail-over times for HA are very short, and it is possible to use a synchronous replication avoiding any data

loss, so for HA there is nearly no limitation. In case of a serious problem in one of the data centers (e.g. fire), the second will take over seamlessly, providing full functionality to the end users.

In this example, the distant data center is used only to save the data to protect it against a major physical disaster destroying both DC1 and DC2. If this happened, the outage time would be considerable – systems need to be installed, but some customers do follow this approach. They typically have an agreement with a hardware vendor who will provide the hardware components in a very short time.

4 SAP PLATFORM – TRANSITION AND OPERATIONS

4.1 Transitioning to the SAP platform

Before operation, there is transition. SAP Active Global Support (SAP AGS) offers an engineering service which addresses all aspects of the safe transformation of the SAP system landscape to the SAP HANA platform in a private cloud, holistically. It combines best practices, remote and on-site activities into one end-to-end methodology.

The SAP Active Global Support offering is delivered as a flexible package consisting of the methodology implemented in SAP Solution Manager, with the benefit of a consistent methodology used by all roles (TQM, consultants, customer experts, partners), and a consistent approach and quality across all parties involved.

Best practices templates and guidelines, on-site services and remote offerings, such as Expert Guided Implementations and remote expert sessions, and guidelines for the reduction of customer code, provide a solid foundation for a successful SAP HANA transition.

The methodology, and all related templates and best practices documents, are part of the SAP Solution Manager, and you can use it free of charge. If required, we can support you with tailored services in each of the project phases. You decide for which phase and for which task you would like to involve SAP.

Methodology Overview

The methodology starts in the planning phase. It offers templates and methods to evaluate the main influencing factors in your company, and to define your transition project in detail – the main project phases, the systems used during migration, the length of the different project phases, and the prerequisites.

SAP has developed migration tools and procedures for the transition to SAP HANA, ensuring a minimum downtime of very few hours, independently of the database size. We provide different tools to define a migration method which fulfills your requirements. As well as supporting and guiding the project, this methodology will ensure end-to-end quality management of the project.

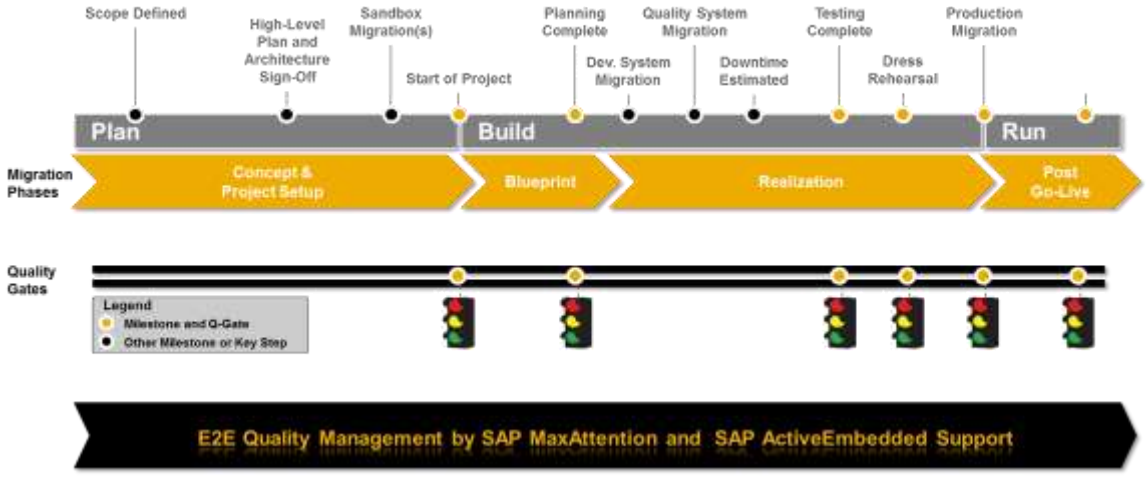

For SAP, this is the most important element of the methodology. Our objective is to ensure a safe transition to HAHA, independently of who carries out the migration and implementation tasks.

This approach provides various benefits to you:

- **Risk mitigation:** Leveraging the breadth of customer experiences, risks are identified early, so that appropriate action can be taken proactively.
- **Fast and predictable time to value:** The migration methodology in the SAP Solution Manager can establish a solid foundation for the project, and track all migration activities.
- **Quality assurance:** The methodology includes comprehensive quality assurance, avoiding critical risks.
- **Cost control:** Careful planning prior to migration, and regular evaluation and assessment of all aspects of the project, help to keep costs under control.

If required, SAP Active Global Support can take over entire transition projects. Alternatively, partners can be empowered to execute the migration and transition tasks, and the SAP team supports them with best practices documents, methodologies and back-office support.

4.2 Managing a private cloud

Operations aspects have to be considered as well as the transition and deployment. Establishing a cloud and realizing it's benefits for the solution landscape has been the focus of the previous sections. But what is necessary for operations?

There are many aspects of managing a cloud that also apply to managing any data center and application solution landscape. Core implementation processes, from solution planning to requirements collection, to project planning and execution, to design and test, and integration validation standard adherence, are required. Core support processes, from service request ticketing to monitoring and alerting, to exception handling, to incident and problem management, to change management, to continuous improvement, are also required.

The SAP Solution Manager platform has been the central management platform for these processes for many years. Nevertheless, it needs to be complemented for comprehensive cloud operations.

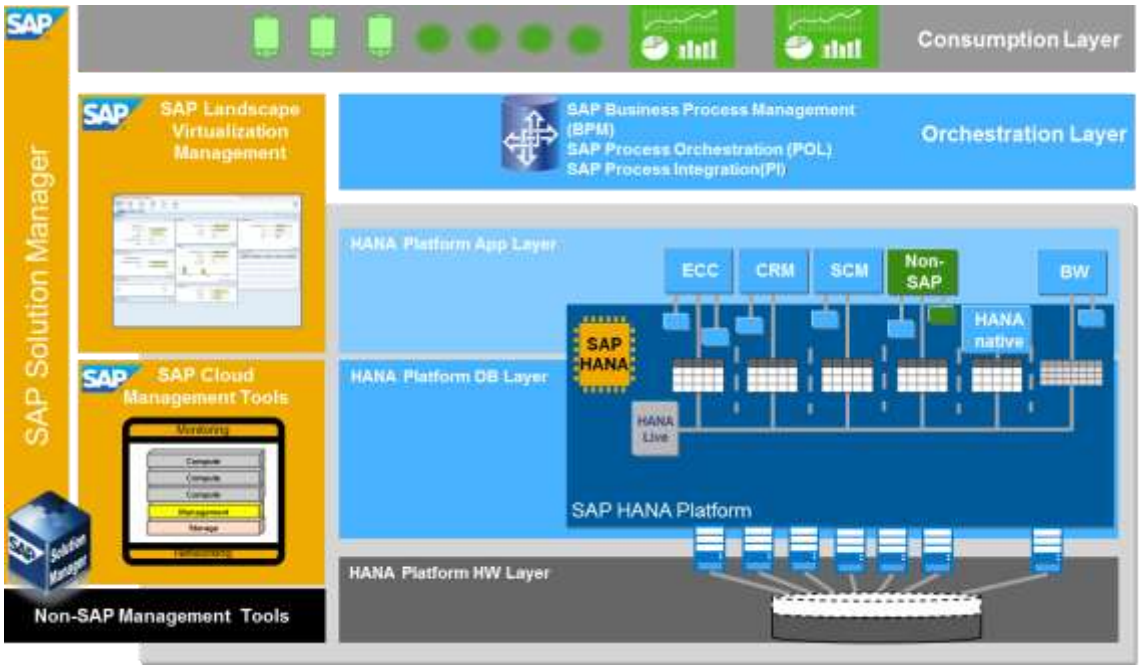

Cloud management needs to address a few more critical topics. In addition to the traditional availability, stability and performance functions, additional cloud operations processes and functions are required, to support higher demands for precise monitoring across all layers, flexibility of change, automation, and system provisioning.

- **Fast solution provisioning.** A key characteristic of a cloud is the on-demand provisioning of application or platform instances. The IT stack layers need to be well-defined and isolated and standardized, to facilitate rapid provisioning. Manual and automated provisioning procedures to manage applications, SAP HANA and the virtualization layer, are a key management requirement..
- **Flexible resource assignment**. The overall data center architecture needs to support growth and shrinkage within the IT infrastructure layers. Tool-managed and self-managed resource assignment capabilities, and their monitoring, are another key function in our management tool platform.
- **Scalability in all layers.** The overall architecture rests on technical capabilities to extend the cloud landscape smoothly, with minimal operational impact. The management tool platform needs to support infrastructure and application growth, and redistribute the load.
- **Availability of the solution landscape.** Depending on the application, the application and system architecture needs to support state-of-the-art fail-over and recovery capabilities within the cloud, or even across clouds for disaster recovery. The management platform needs to support manual and automated fail-over management and recovery functions.
- **Transition into the cloud.** To facilitate the smooth transition of a solution into the cloud environment, there needs to be a standardized infrastructure layer with the core implementation and operations processes mentioned above. The SAP Solution Manager platform offers a rich set of best practices, guidelines and functions to support this important first step into the cloud.
- **Cloud Supportability**. Supportability of the landscape and applications is ensured through the extension of our management platform with cloud-specific best practices, guided procedures and tool extensions, across all layers and management process.

The picture in this section provides a high level overview of the SAP cloud management tool-set. On the lowest technology layer, we partner the providers of the infrastructure components. For management on server and system level, SAP has developed a tool to support and automate functions like server and system provisioning and relocation. The SAP Landscape Virtualization Manager LVM is used on application level.

SAP Solution Manager provides the required transparency across all layers, end-to-end, for all business processes.

4.3 Reduction of IT cost through the SAP platform

Private cloud infrastructure capabilities, with a comprehensive management platform, ensure stable operations, and reduce overall IT cost significantly. In this section, we will look at how SAP HANA and the deployment in a private cloud reduce IT cost.

Firstly, there is a general reduction of service and infrastructure cost.

- **Cheaper Hardware.** SAP HANA runs on modern, efficient and "commodity" Intel-based server systems. SAP HANA allows you to run on these generally cheaper systems, and simply scale to higher levels via SAP HANA, instead of installing more costly equipment.
- **Efficient usage of servers through virtualization.** Rapid deployment of SAP HANA solutions depends to a great degree on the capability to provide server systems with the right performance and capacity instantly. For a number of years, SAP has had a strategic partnership with VMware to make efficient use of available hardware capacity – for non-productive systems and application servers. Advanced automation techniques to manage the virtualized environment use resources efficiently, and reduce management effort significantly.
- **Less storage required.** One of the many breakthrough innovations of SAP HANA is intelligent compression, which is much more effective than compression in traditional databases. SAP HANA technological innovations greatly reduce the volume of indices, and require no aggregates, no data duplication, and consequently less data.
- **Green IT Strategy contribution**. SAP HANA with deployment in a private cloud can be part of an energy-conscious data center strategy, for the following reasons:

  Virtualization technologies allow SAP HANA to maximize the usage of the hardware components in a data center

  By strongly compressing data and eliminating data duplication, explicit aggregates and most of the indices, SAP HANA reduces the overall disk space required. This reduces the mumber of devices needed by the storage subsystem landscape, and overall power consumption.

  The reduction of systems achieved by consolidation based on SAP HANA also reduces power consumption.

  SAP HANA's integration into SAP Landscape Virtualization Manager allows for a largely automated and efficient Green IT data center management. Changes in the workload behavior can be dealt with efficiently by moving, starting, stopping, pausing, extending, or reducing whatever resource is needed.

Landscape simplification is also important, as well as the significant saving potential.

- **Fewer SAP Systems**. SAP HANA integrates OLAP and OLTP loads. SAP HANA, can consolidate systems, with a direct impact on the number of systems, interfaces and replication.
- **Simplified technical operations**. SAP HANA built-in administration and close integration with SAP Solution Manager further reduce the operations effort. Combined with best practices and guided exception management, two operators per shift becomes a reality!

Finally, there are business process operational aspects.

- **End-to-end business process management**. SAP HANA simplifies and reduces business process management effort. Reduced cycle times, and reduced time and effort to close the books, are typical examples.
- **Reduced business operations effort**. Reduced effort and time, in different areas, like reduction of time for sign-off procedures, reduction of report reconciliation effort through real-time reporting, are other examples of the benefits realized.

5 CONCLUSION

The SAP platform powered by SAP HANA is a new generation of platforms, enabling new types of applications and interaction models. SAP HANA is much more than just a fast in-memory database, and it is complemented by powerful orchestration and user experience components. With this platform SAP can develop business applications which have not been possible before. And it enables you to re-think your own business process implementations – leveraging the SAP innovations and building your own custom applications based on the new SAP platform.

SAP Active Global Support accompanies you on this journey – with best practices and methodologies, with customer-tailored empowering initiatives, with remote and on-site services, and with a global organization, making the top experts available to you when you need them. The combination of the Innovation Control Center, the Operations Control Center and the central Mission Control Center at SAP is the foundation for ensuring the success of your projects, and the efficiency and stability of your solutions.

www.sap.com