



Creating Business Value with SAP HANA

Bram Coolen and Aart-Jan Boor



B. Coolen MSc RE
is director at KPMG Advisory.
Coolen.Bram@kpmg.nl



A.K. Boor MSc
is consultant at KPMG Advisory.
Boor.Aart-Jan@kpmg.nl

Bangalore, India, December 1, 2010. SAP announced the first version of SAP HANA, a new product directed toward “in-memory computing.” The product was positioned as a revolutionary innovation. It would be a “game changer” that would entrench data analysis at the heart of the organization. In the meantime, we are almost three years and more than 250 press releases on HANA further. Various success stories have come to the fore about organizations that carried out the first HANA implementations. At the same time, there is still an overriding need for a good business case with regard to new investments in the application landscape. An accompanying challenge is the fact that technology alone has no business value: it is all about the way in which technology is applied and the value it can generate for an organization. This article deals with the opportunities offered by SAP HANA, and we give concrete examples about the way HANA can be applied to create business value.

Introduction

ERP systems were originally developed to support daily operations within organizations. The integrating nature of the system, in particular, provides benefits in the execution of processes. Information needs to be entered only once and is then directly available in other processes that can make use of it in various ways.

However, soon after the introduction of the ERP system, there arose an increasing need to formulate reports on the collected data in the system. In addition to operational reports, there was also a demand for management reports for tactical and strategic use by management and staff departments. SAP responded to this need by launching the SAP Business Information Warehouse (now: SAP BW) in 1997. This data warehouse allowed SAP data to be retrieved, processed and published in reports.

In 2007, SAP decided to take over Business Objects (BO). Because this company was strong in the field of analysis and reporting, this acquisition formed a good supplement to the BW software, which was, in turn, strong in the field of extraction, storage and transformation of data.

Since then, market alterations, product developments and technologies have succeeded one another at an ever-increasing rate. As a consequence, the demand for data analysis and reporting tools to cope with and capitalize on these changes has grown. At the same time, data volumes

are doubling approximately every year ([GANT12]), which presents many challenges to many organizations with regard to the performance of their SAP BW systems. In this context, these organizations are faced with the expanding issue of being able to combine internal data sources with external data (such as Twitter). Finally, many organizations are also struggling with the rigid structure of their ERP and BW systems. This means that the development of new applications and reports requires a great deal of time. Rapid response to this situation therefore becomes extremely expensive.

SAP now presents the new SAP HANA technology as a response to these developments. Although it was originally launched as a fast database, the product has quickly evolved into an application platform. This article first elucidates the functional components that make up HANA. It then covers the relationship between HANA and other SAP products, and the way organizations can create business value with this new technology.

What is SAP HANA?

Figure 1 provides a functional overview of a typical HANA environment. An important property of HANA is that it is an open platform. Both the database and the platform components offer various technical possibilities to integrate this with other software.

The analytical capacities of HANA are characterized by their simplicity, but they are exceptionally powerful

In-memory database

The core of SAP HANA lies in the in-memory database. Computers contain two types of memory. The first type is the harddisk on which the data is stored. This is relatively slow due to the use of mechanical components, but it is also cheap. The second type is the fast working memory (RAM) in which data that are being used by active programs are temporarily stored. The disadvantage of this type of memory is that all data are lost if a power outage occurs.

In SAP HANA, the complete database is loaded into working memory, so that it can be consulted extremely quickly. In terms of hardware, this is very expensive because large quantities of RAM memory have to be installed on the server. As a result of the declining costs of this type of memory, however, this solution is progressively coming within reach of many organizations. Moreover, this process is accelerated by the use of increasingly advanced data compression techniques that enable an abundance of data to be stored in certain amounts of RAM memory. Furthermore, advanced backup techniques prevent data loss in cases of power outage.

Data sources

An important component in a HANA environment is data acquisition from the source systems. The distinctive feature of this capacity is that HANA makes use of the SAP Landscape Transformation (SLT) component to retrieve (almost) real-time data from SAP ERP source systems. An improvement in relation to SAP BW is that no separate development of extractors is needed for this function – the tables to be retrieved can be simply configured without complex coding.

In addition to ERP data, HANA can also retrieve data from various source systems, including SAP BW, SQL Server databases and Hadoop. In these scenarios, data in the source system is copied to the HANA in-memory database.

An alternative is to use the SAP HANA “smart data access technology,” which makes it possible to analyze data that is not physically located in the HANA database. Instead, a request is made to the source system to retrieve the necessary data *ad hoc* during the execution of the analysis. Although this may delay the execution of the analysis, it offers organizations the possibility of combining fast memory with slower and cheaper memory.

Analytics

The analytical capacities of HANA are characterized by their simplicity in application, but they are exceptionally powerful in use and provide many more possibilities than SAP BW or SAP BO. For instance, it uses a structure with reusable “views” that can combine data from various tables with calculations and aggregations. In addition to graphic tools to support the construction of views, it can also use the well-known SQL language to execute queries.

In addition, HANA also contains the Predictive Analytics Library (PAL). This is a function library with a set of frequently used, advanced analysis algorithms. Regression and cluster algorithms are examples of such facilities. These functions can be used to create views and to carry out analyses. Being integrated in the HANA database, these algorithms perform extremely well and they can also be used in simulation applications or when applied to real-time data.

In the cases in which these algorithms do not suffice, HANA offers the opportunity to set up integration with an R server. R is a statistical programming language much used in science. It is an open-source product that has already made several thousand packages available for free. These packages contain diverse algorithms for various

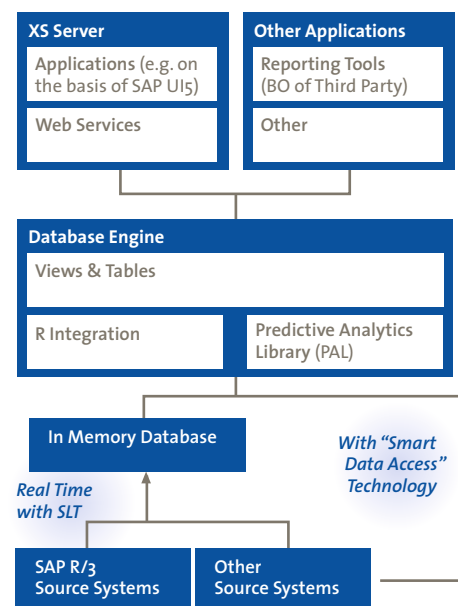


Figure 1. Overview of SAP HANA.

applications. Due to the integration with HANA, these algorithms can be directly applied to data in the HANA database, and the output of the algorithms can again be stored in the HANA database. Although the performance is generally slightly inferior to that of PAL in similar circumstances, the diversity of the algorithms does enable interesting analyses.

Big Data

SAP HANA is often positioned as an application for processing Big Data. The above-mentioned components indeed make HANA a suitable product for this purpose. The possibility of using external data and various advanced analysis algorithms bring new Big Data applications within arm's reach. In addition, it is also important that the performance is also sufficient when dealing with large data sets, and that the solution offers adequate scalability for processing such volumes of data.

Application server

Besides a database and analysis platform, SAP HANA also provides a built-in (web) application server, the XS Server.

Although this server only offers simple functionality, direct integration with the HANA database makes it possible to run very powerful applications directly from the HANA server. The benefits of this include a simplification of the landscape and good performance as a result of the direct connection with the HANA database.

The applications should be developed through HTML5, CSS and JavaScript. By using these web technologies, the applications can be run in the browser. To create the applications, use can be made of generic tooling from various suppliers. SAP itself also offers a toolkit, namely the SAP UI5 library, by means of which simple, consistent user interfaces can be built, in which data from SAP HANA can be used.

In addition, the XS Server can also access data from the HANA database in the form of web services that can be used for selected external applications. These applications can make use of the data from HANA or may lean entirely on HANA for the execution of their activities.

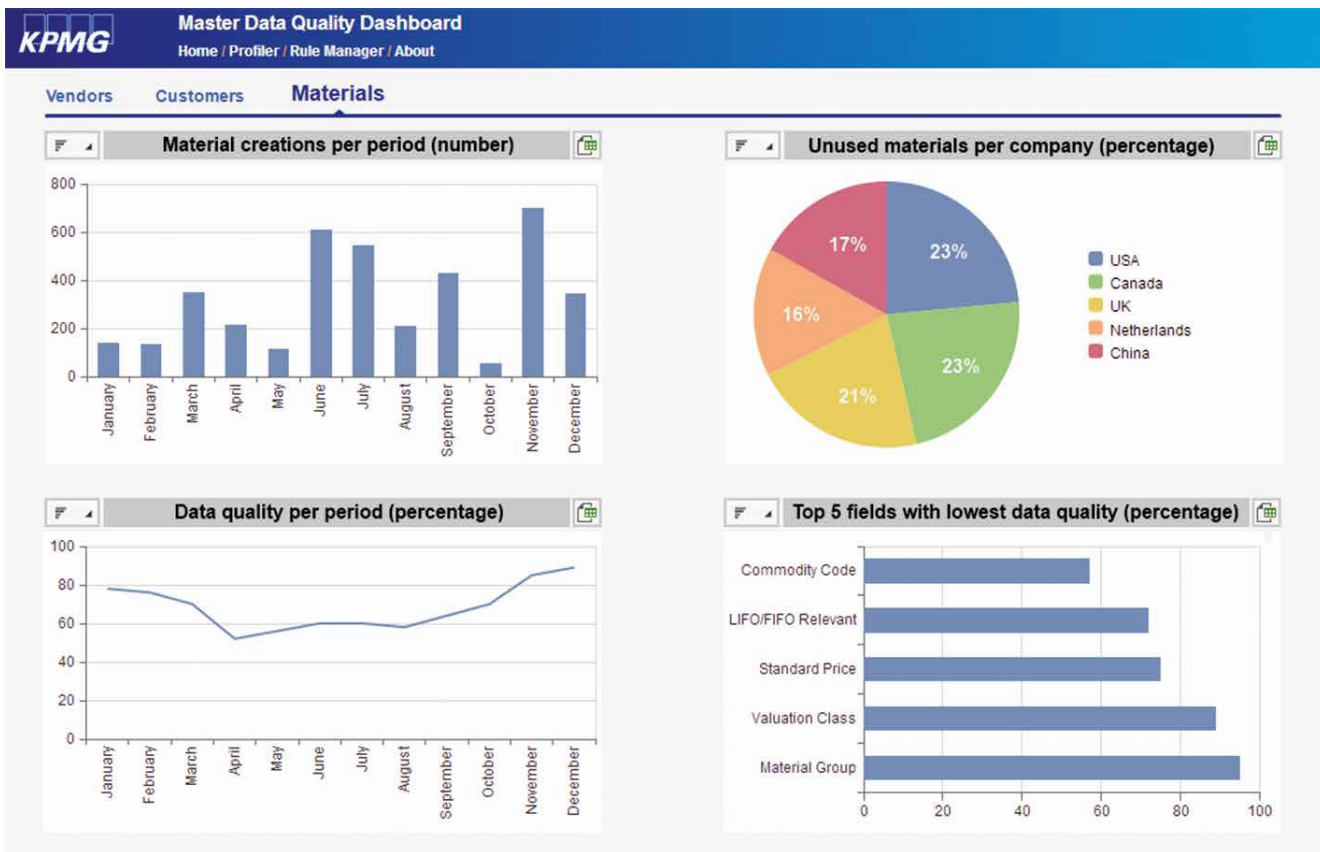


Figure 2. HANA Master Data Dashboard on XS Server.

Not the technology in itself but the applications generate business value

Reporting

SAP has ensured good integration between SAP HANA and SAP BO, so that it is more or less self-evident that the BO tools should be used for reporting. HANA views can thus be directly consulted in many BO applications. In addition, the open character of the HANA database also provides the opportunity to connect it to other reporting tools. Various interfaces and connecting options are available for this purpose, such as ODBC, ODBO and MDX. This gives organizations the possibility of integrating the reports in the applications that are already in use. Finally, it is possible, of course, to access dashboards and reports via the XS Server. Figure 2 illustrates an example of a dashboard running on the XS Server, reporting the quality of master data.

Cloud

As with other software, SAP places a great deal of faith in the Cloud with the introduction of SAP HANA. There are already various suppliers from whom SAP HANA systems can be accessed from the Cloud through flexible settlement models. A well-known supplier is Amazon, which offers SAP HANA servers as part of their regular assortment. These can be hired at an hourly rate. SAP too is engaged in setting up a Cloud from which HANA servers can be hired on the basis of flexible contracts.

HANA in the SAP landscape

Discussion often centers on the way in which SAP HANA is related to other products within the SAP landscape, particularly SAP BW and SAP BO. Since SAP announced, in spring 2013, that the SAP ERP system can also run fully on HANA, other SAP databases such as Sybase ASE and IQ have come into existence.

With regard to this last point, an organization has to decide which product is the best option. It is expected that, in many cases, a combination of databases will be chosen, with the aim of limiting the costs. As a result of the HANA “smart data access” technology, data from one application can be physically distributed among various databases without the application having to be specially adapted for

this purpose. Data that is often consulted can be placed in the HANA database, for example, while archived data and log data can be stored in a secondary (cheaper) database. Moreover, specialized databases will probably be used, such as Hadoop or databases for the storage of geographical or localized information.

The choice between HANA and SAP BW will have to be evaluated in each situation. At this moment, BW contains more functionality than HANA in the domain of data warehousing and modeling. Specific functionalities such as Time Dependent Master Data and hierarchies enjoy better support from BW. However, the specific advantages of HANA, such as real-time data access, performance, speed of development and advanced data analyses, are difficult to realize through BW. Each organization will have to deliberate as to the particular product in which they wish to invest for a specific application.

Making a specific decision for a specific situation also applies to BO. With BO’s self-service reporting tools, HANA enables the simple graphic consumption of information. However, many of these tools are not capable of altering data, for example, so that they are ultimately not suited to building dynamic applications. In addition, extra licenses are required for the use of BO with regard to the built-in XS Server (which is included in the HANA license). Nevertheless, BO does provide much more functionality for the (end) user than HANA does, with respect to the reporting layer.

Business value

One of the major barriers to working with HANA is the sizable investment required to purchase HANA (with regard to both hardware and software). Although HANA provides powerful functionality, there is often doubt as to whether or not, in practice, the business value can offset these considerable costs. The challenge actually lies in the fact that the technology in itself does not create any value; it only enables the realization of valuable applications that can generate business value.

Responding to questions from the business

Many globally operational organizations have set up a Business Intelligence (BI) architecture directed toward

the delivery of reliable reports that are consulted by end users every day, week or month, and where it is crucial that the figures are 100% correct. As Figure 3 shows, this data often consists of information from various sources and layers, depending upon the place of the end user in the organization (locally in a (business) unit, as opposed to centralized). Although this architecture meets the need for reliable data and reports, it may cost much in terms of effort to develop new analyses.

Case: Alliander

Alliander is the largest network operator in the Netherlands, with 3.5 million customers. Alliander has opted to use SAP HANA in order to be able to analyze enormous amounts of user data in real-time. On the basis of these analyses, customers can receive tips to reduce their energy usage.

Case: Suiker Unie

The Suiker Unie (literally “Sugar Union”) is a production company that makes various sugar products from sugar beet. The Suiker Unie opted for HANA in order to obtain real-time insight into the results of quality samples during the short sugar-beet harvesting period. In this way, direct interventions can be made in the supply and processing of the beet, so that the quality can be improved and production costs kept to a minimum.

the business. Roughly speaking, these fall into two categories:

1. Business(-process)-specific questions involving the primary processes of organizations, such as operational reports in the field of goods movements, stock levels and order backlog. The most important feature of these kind of questions is that an answer is almost always required immediately, and is directly used to make decisions and to initiate follow-up activities (concerning, for example, which sales order generated delivery problems or which delayed purchasing orders could have an effect on delivery reliability). In this context, these type of questions are usually not (well) covered by the operational reports in SAP ERP.
2. Theme-driven questions, often raised by current developments (such as the unexpected rise in the required working capital of an organization, and the wish to know the cause of this, and the determination of actions to rectify the situation). The most important feature of these types of questions is that the possible answers should follow and cohere thematically: with which suppliers have short payment periods been agreed? Are payment discounts for short payment periods being used? etc.

The properties of SAP HANA are extremely suited to providing quick answers in both categories on the basis of one integrated system (see also the case studies of Alliander and Suiker Unie in this framework). In addition to savings on management costs, cost savings can be realized in the development of solutions through the use of Agile

For such reasons, organizations are increasingly searching for ways to find rapid answers to specific questions from

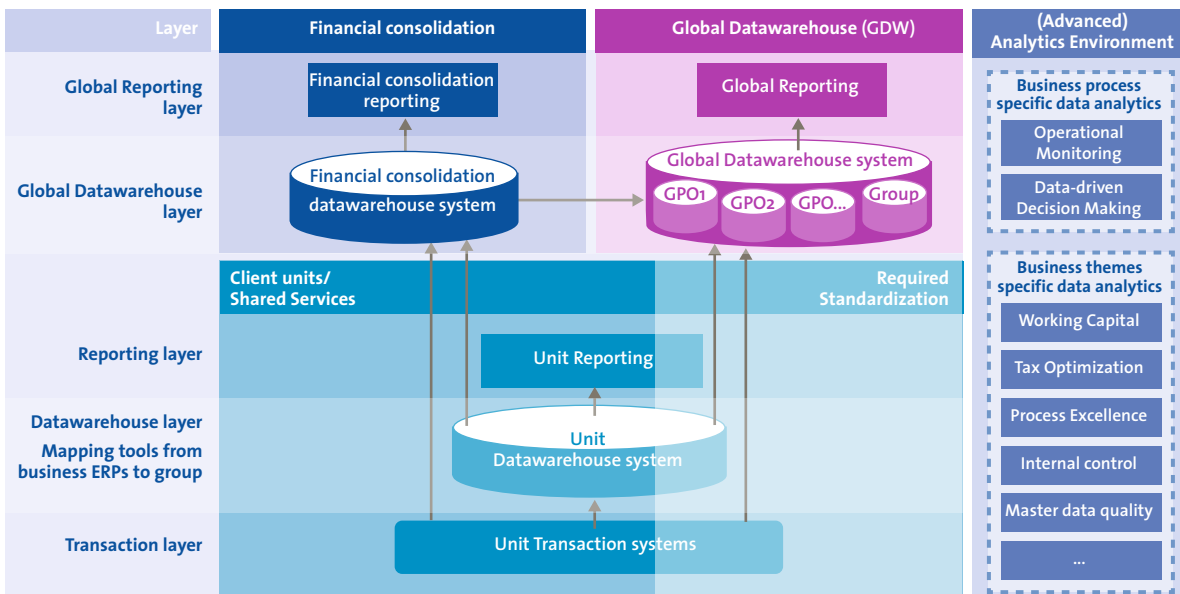


Figure 3. Example of Business Intelligence architecture.

Because development in HANA is low-threshold, it is relatively easy to involve the customer in the development

Development and the SAP HANA Marketplace.

Agile Development

Many organizations have already switched from traditional developmental models to Agile methodologies. A feature of this is that such models work with short developmental iterations, each of which delivers a working product. This gives the customer the opportunity to make rapid use of (partial) functionality and to make minor adjustments (if necessary).

Development in HANA is low-threshold by nature and supports a high degree of modularity. As a consequence, it is relatively easy to work with an Agile methodology and to involve the customer in the development process to an optimum extent. It is also easy to initiate pilot projects in order to test whether or not a certain functionality or report can deliver business value without having to elaborate this to the fullest extent. The value to the business is that results can regularly be evaluated interim, and can be adjusted during the process. This ought to produce a better result in less time.

Marketplace

The spring of 2013 witnessed the launch of the SAP HANA Marketplace. This is an online platform on which, analogous to the Apple “Appstore,” various ready-made solutions can be provided by different suppliers. Examples of such solutions include an application for the analysis of operational processes, a security monitoring application, and an application for enterprise performance management. Due to the fact that HANA is an open platform and has a built-in application server, these applications can often be easily installed on the HANA system of the organization. Although the assortment is still rather limited, it will probably grow swiftly in the near future. It is expected that the transparent price-forming and ease of use will lead to lower purchasing costs of HANA applications that make use of data from the SAP system. It may be interesting to organizations to purchase a standard app for a certain application, for use in a specific project, for example. In addition to cost-saving in relation to tailor-made solutions, the short lead-time gives an important benefit in many business scenarios. In many cases, the applications can be

installed reasonably easily on the existing HANA system, which means that the business can capitalize more rapidly on developments in the market and in the organization.

Conclusion

This article has provided various examples of how to use SAP HANA to create new business value, such as the positioning of HANA next to the traditional BI environment in order to create specific BI applications to support projects and to respond to *ad hoc* questions from the business. Another possibility is the realization of new applications that are technically difficult to develop by means of traditional systems (such as predictive analytics and the fast deployment of HANA apps).

In addition, business value can also be sought in the reduction of developmental costs within the BI domain. By making use of Agile developmental methods and off-the-shelf BI from the SAP HANA Marketplace, costs can again be cut.

In summary, we hope that, with this article, we have been able to give you insight into recent developments. We also trust that we have provided points to consider in deciding whether or not investment in SAP HANA may deliver business value for your organization.

References

[Gant12] J. Gantz and D. Reinsel, *The Digital Universe in 2020: Big Data, Bigger Digital Shadows, and Biggest Growth in the Far East*, IDC, 2012.

About the authors

B. Coolen MSc RE is director at KPMG Advisory in the field of process improvement, business intelligence & reporting, master data management and ERP/BI strategies. He is particularly engaged with organizations running with SAP systems, helping them to identify and realize improvements.

A.K. Boor MSc is consultant at KPMG Advisory and works within KPMG's SAP Advisory group. He has gained experience with a broad range of ERP-related advisory and audit assignments. He is certified to work with SAP HANA and is experienced in the development of HANA applications related to creating business value.