

## **SAP HANA**

### **By Jaqui Lynch**

#### ***Introduction***

SAP has become an integral part of many client's environments. As those SAP clients try to make the best informed business decisions they look to the analysis of data that comes from both transactional and analytical systems. Typically, this is data that is kept separate as it is a combination of massive amounts of structured and unstructured data that does not mix well. Combining this data allows businesses to mine the data in ways that could not occur when the data was separated. That mining allows businesses to gain more in-depth knowledge of their business and to better position themselves in today's competitive markets. SAP HANA is designed to deal with this kind of data – huge amounts of data that needs rapid analysis to provide the information needed by the business.

#### ***HANA***

SAP HANA was previously called "SAP High-Performance Analytic Appliance" and is an application server that includes an in-memory, column-oriented, relational DBMS (database management system). Using an in-memory platform allows the use of smart analytics that can run incredibly fast. Key advantages of SAP HANA include the ability to significantly accelerate database processing by processing transactions and analytics against the single copy of data that is loaded in-memory. This provides the ability to perform analytical processing to mine data from the past and present, and to provide deep insight into that data allowing you to make better business predictions for the future.

SAP HANA is a standard-based, open architecture consisting of three main components. The first is database services, which uses an ACID, SQL 92 compliant database that is accessible through JDBC, ODBC, ODBO or OData. There are a number of third party administration tools that you can choose from to manage the database. The second is application services, which provides for a choice of application and web servers and includes HTML5 UI libraries. Additionally, there is an Eclipse based web development tool. The third component is integration services which is responsible for data movement and federation with existing databases. Data can come from standard data sources such as DB2, Oracle, Netezza, and MS SQL (to name a few) as well as from less standard data sources such as streaming data. Streaming data could be from social media, mobile devices, sensors or any number of other entities. This is also where integration with systems like Hadoop and R is provided.

Database services within SAP HANA are accelerated by using an in-memory database that compresses the data in memory. This significantly reduces I/O time once data is loaded, as it is stored in a columnar format and can be processed in parallel across multiple processor cores. This provides for very fast data scans and allows for high-speed advanced analytics. Additionally, SAP HANA allows you to include complex business logic inside the database which

avoids caching data in application servers. The performance of SAP HANA workloads is directly affected by memory capacity available to the database. This is due to the fact that the data is stored in memory in order to reduce disk latency – the more data you can store in memory the less latency and the better the performance. Thus, SAP HANA benefits greatly from running on systems that can provide significant amounts of memory.

SAP HANA also provides for multitenant databases. By providing multitenant database containers, SAP HANA allows you to manage multiple databases as if they were one database, but database isolation is still maintained (needed for security reasons). This vastly simplifies patching, backups and ongoing maintenance. SAP HANA also provides for storage tiers. Using SAP HANA dynamic tiering you are assured that frequently accessed data will remain in memory while less active data will be moved to disk. This ensures you are taking advantage of the high speed resources for the right data and that you are not wasting resources on data that is rarely accessed. The application neither knows nor cares where the data is stored and access to the data is transparent with respect to location.

Finally, there are multiple deployment options for SAP HANA. Apart from a choice of platforms there is also a choice for how to deploy it. Options include running all the applications in the cloud, running all the applications onsite at your location and, finally, a hybrid where business priorities dictate where modules get deployed. This can be done whether SAP HANA is run as an appliance or on a combination of components that meet the SAP requirements.

### ***IBM and HANA***

Traditionally, SAP HANA was run as an appliance that was preconfigured and preinstalled and had limited flexibility. Today, IBM offers SAP HANA Business Suite on all POWER8 servers. SAP HANA BW (business warehouse) runs on SLES 11 (Suse Linux Enterprise Server) for Power with a special SP3 patch, and can be set up in a dedicated LPAR or virtualized through PowerVM. This provides you with the ability to customize the system and to choose from multiple storage options. This also allows you to integrate your SAP HANA solution into your existing Power environment, rather than bringing in a proprietary appliance. POWER8 provides for highly multithreaded computing with SMT8 (simultaneous multithreading 8), high bandwidth memory and high speed data caching along with flexibility features like COD (capacity on demand).

When you look at the RAS (reliability, availability and serviceability), performance, I/O and management capabilities of POWER8, coupled with the new SAP HANA IBM Power Solution Editions it is clear that POWER8 is a very attractive option for SAP HANA. This is especially true for those who already have POWER8 systems in their environment and who don't want to bring in yet another platform to manage. Additionally, the largest POWER8 servers provide for up to 16TB of memory which can significantly improve performance as this

means less data is on slow spinning disk. Combining this with the L2, L3 and L4 memory caches on POWER8, there is also significantly less memory latency which again improves performance. The new Power Systems Solutions Editions are designed to support SAP HANA BW on POWER8 and are architected to provide performance, reliability and flexibility combined with competitive pricing. Today you can run both SAP HANA Business Warehouse and Business Suite workloads on the same POWER8 server. This provides more flexibility and less complexity for SAP users and allows for better allocation of resources as they are needed.

### **Summary**

Today's business world is full of new challenges related to data management and analysis. Big Data and the combination of both transactional and analytical data have become a significant factor in realizing competitive advantage. SAP HANA allows customers to access and analyze important business data in real time — without having to worry about running SAP workloads on separate platforms.

Making smart business decisions requires a combination of experience and accurate information. SAP HANA is a critical method of providing the data to support good decisions based on current data, including external factors such as domestic markets, foreign markets, weather patterns, customer buying behavior, and floods of unstructured data coming from social media or the web.

The next generation of HANA is S/4HANA and it promises much faster analytics and far superior data aging, along with the ability to even more rapidly process all social, text, geo, graph and processing data. As HANA continues to innovate and take advantage of the large memory systems that are now available, customers can expect to have larger in-memory databases which will provide for better and faster access to big data analytics, providing the tools needed to compete effectively in this era of data mining. This can only improve the ability for businesses to better understand their customers and to plan for the future.

### **References**

Getting fast business insight with SAP HANA on SUSE Linux Enterprise Server and IBM Power Systems

<http://www.redbooks.ibm.com/redpapers/pdfs/redp5248.pdf>

SAP HANA Database Capabilities

<https://hana.sap.com/capabilities/database.html>

SAP HANA Analytics Capabilities

<https://hana.sap.com/capabilities/analytics.html>

SAP HANA Hadoop Integrations

<http://www.slideshare.net/SAPTechnology/sap-hana-sps10-hadoop-integration>

SAP HANA SPS11 (Support Package Stack 11) Release Notes  
[http://help.sap.com/hana\\_platform/](http://help.sap.com/hana_platform/)