

Better Late Than Never?

Selecting the Right Technology for Your
SAP Data Warehouse Environment



Overview

The SAP HANA in-memory computing platform is not just a game-changing database; it represents a whole new paradigm for enterprise computing. HANA promises to completely transform business processes and ultimately businesses themselves. With every major player in the enterprise database market now attempting to follow suit, it's no longer a question of whether your business will adopt in-memory technology: it's just a question of when.

As a consumer, you personally can afford to be a late technology adopter. Maybe you can even afford to have your business be a late technology adopter. But relying on a *technology provider* that's behind on the technology adoption curve is just too big a risk.

If you are currently using (or considering using) a Teradata solution for your data warehouse, that is exactly the risk you're taking. Teradata's partial adoption of the new paradigm only brings unwanted and distracting complexity to your environment. Moreover, Teradata's piecemeal, point-solution mentality and emphasis on technology rather than business solutions drain valuable attention and resources away from your key business priorities.

By contrast, SAP provides not only industry-transforming technology but unmatched knowledge and mastery of business process — the kind of expertise that comes from solving business problems for the world's largest, most complex, and fastest-paced organizations. Businesses that use HANA can enjoy those benefits today rather than waiting an indefinite period of time for a scaled-down and inherently compromised version of them.

I. SAP HANA Background

SAP's introduction of the HANA in-memory platform in 2011 provoked a good deal of activity among the major conventional database vendors. Each of them responded by updating, or at least announcing major modifications to, its long-established data management solution set. Microsoft SQL Server 2014 includes an in-memory option; Oracle has announced a similar option for 12C. IBM DB2 now offers BLU Acceleration to provide an in-memory boost for analytics workloads. Even Teradata came to the table with Intelligent Memory, which allows users to keep the most active parts of their data warehouses in memory.

In-memory databases are a hot topic. Gartner predicts that 35% of midsize and large businesses will include an in-memory solution in their overall architecture within the next 12 months.¹ It seems that everyone wants in on the act, but not all in-memory solutions are created equal. The major database players have introduced in-memory capability to accelerate analytics performance in a conventional environment, one where multiple copies of the data are scattered throughout the organization. These vendors have long encouraged such a model; they are, after all, in the business of selling as many database licenses as they can.

By contrast, with HANA, SAP has introduced a new paradigm for real-time computing: a single copy of the data, without the need for aggregates or special indexing, supports all transactional and analytical workloads and use cases. A HANA environment is never slowed down by traditional database performance management techniques, most of which are workarounds that the conventional database vendors have evolved over time in an attempt to keep performance in line with data growth and ad hoc access.

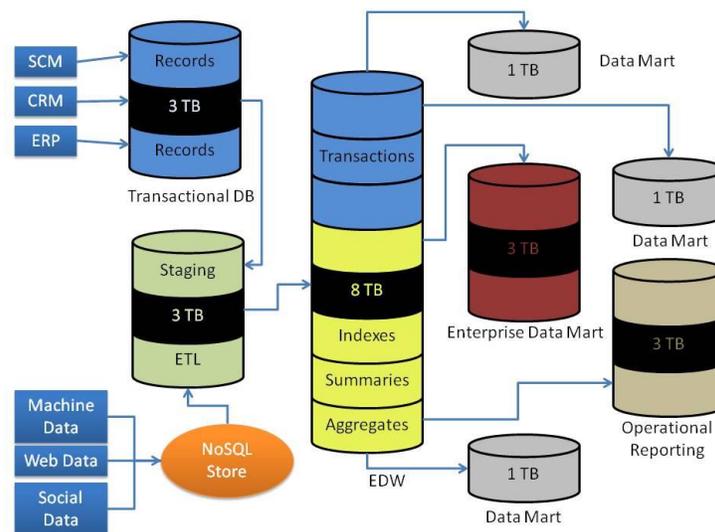


Figure 1: Three Terabytes of Enterprise Data in a Conventional Environment

¹ "Remember the titans: who will win the in-memory battle?" *Information Age*, April 16, 2014.

In developing their new solutions and bringing them to market, the mainstream database vendors are attempting to address a crucial problem, that of data latency. This is a problem that SAP took on some time ago. As detailed in Jeffrey Word's book, *SAP HANA Essentials*, SAP has through the years achieved dramatic progress in reducing the impact of the I/O (input/output) bottleneck between application (or analysis) and disk with such innovations as the Advanced Planning Optimizer and SAP LiveCache. This work culminated in 2005 with the introduction of SAP Business Warehouse Accelerator, which leveraged dual-core chips and in-memory processing to provide exponentially faster performance for SAP Business Warehouse environments.²

Having addressed nearly a decade ago the question that the conventional database vendors are only now contending with, SAP has turned its attention away from yesterday's troubles. With the advent of multicore chips and more broadly and inexpensively available RAM, SAP began to explore the possibilities for the next generation of enterprise data. As a result, they came up with a new, and much more challenging, problem to address: that of enabling the transition to real-time business.

In bringing their new solutions to market, the conventional database vendors assume that the core paradigm will remain the same. Some data will be in memory; some data will be on disk. To improve performance, you change the balance of power between the two. And any time you need to do anything new with the data, you just make another copy of the database.

SAP, on the other hand, has opened up some radically new thinking about how to design and implement an enterprise data environment. They have vastly simplified such environments by moving the entire database into RAM, which eliminates all delays associated with reading from and writing to the disk. The entire enterprise — all applications, all transactions, all analysis — can now run on that single, amazingly fast, and always up-to-date database.

Such an environment enables real-time analytical and transactional processing — something that has been long needed, but that has not been possible at the scale required to support the full enterprise. The new model does away with all of the work cycles and resources dedicated to maintaining multiple database copies and keeping them in sync with each other. It also eliminates the need to distinguish between transactional and analytical data, as well as the need to populate the organization with multiple copies of the same data. Moreover, this model does away with the need for a separate application server. A single live instance of the database replaces all that former complexity.

Although a series of fundamental shifts in database technology are core to the SAP HANA offering, HANA is more than just a new kind of database. On-the-fly materialization means that there is no need to spend time and resources creating, updating, and replacing a smaller, “digestible” subset of the data; the columnar data architecture means that the data is pre-optimized for any query without indexes or aggregates; multi-temperature data management

² Word, Jeffrey. *SAP HANA Essentials*, 2013.

means that non-active data can be seamlessly identified and moved to the appropriate storage. Nor is the new model just about the performance, although HANA delivers exponentially greater performance than any technology that has come before. SAP HANA is ultimately about the introduction of a whole new way to do enterprise computing.

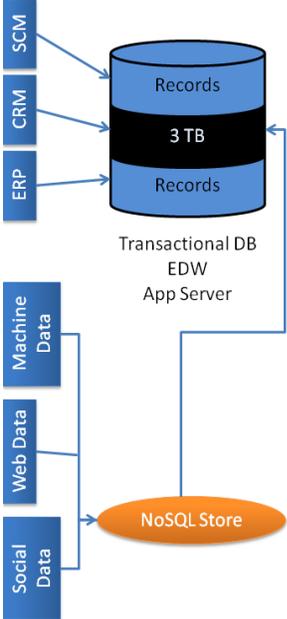


Figure 2: Three Terabytes of Enterprise Data in an SAP HANA Environment

In the same vein, it is not accurate to say that the other players have followed SAP’s example with their own “in-memory” solutions. In comparing their offerings with HANA, we have to remember that each of their solutions is firmly grounded in the old paradigm. They are not asking how to move business to real time; they are still trying to solve yesterday’s problems.

This is nowhere more evident than with Teradata, whose Intelligent Memory offering has less in common architecturally with SAP HANA than it does with SAP Business Warehouse Accelerator — which we have noted was a cutting-edge offering nearly a decade ago. SAP’s leadership in this space provides clear advantages to businesses that are ready to move to the new paradigm. But as we will see, even those businesses that are not yet ready to make such a move should be cautious when considering a solution from a technology vendor that is falling behind on the technology adoption curve.

II. SAP HANA and Technology Adoption

In 1962, Everett Rogers published the first edition of his book, *Diffusion of Innovations* (currently in its fifth edition), which provides a theoretical understanding as to how and why new technologies, and new ideas, are transmitted and adopted within a given community. Although the rate of adoption of an innovation will vary greatly depending on the characteristics of the innovation and of the community adopting it, in all cases there is an inflection point, an adoption threshold that must be met if the innovation is to succeed.

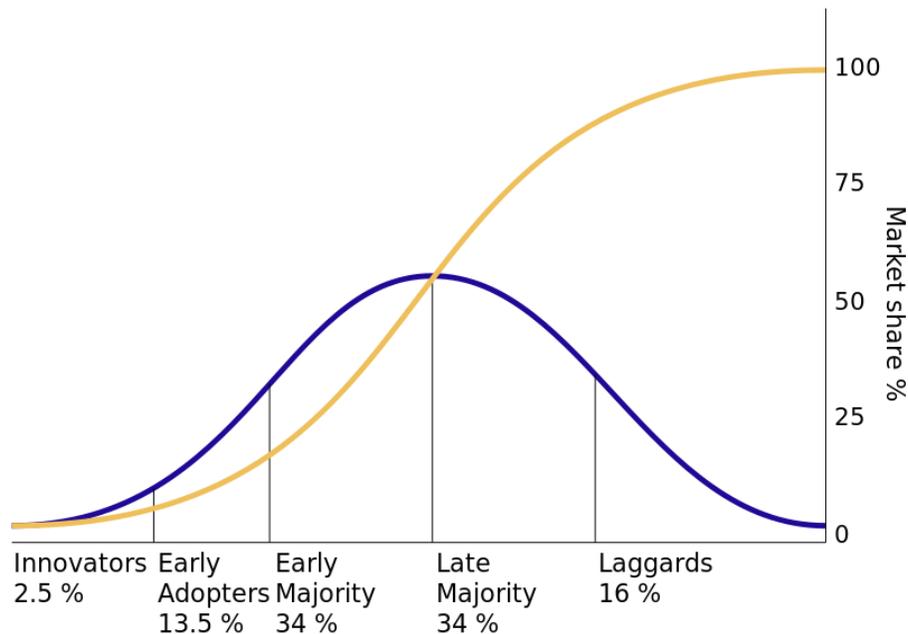


Figure 3: Everett Rogers' Diffusion of Innovations Model

Rogers introduced five categories of innovation adopters based on when in the adoption cycle individuals choose to act. As shown in Figure 3, above, the five groups form a normal distribution over the period of time the innovation is introduced. The first group to try the new idea or technology is the Innovators, often viewed as enthusiasts or trendsetters. They are followed by the Early Adopters, who see benefits in the new innovation that offset the apparent risks in adopting it. The next two categories make up the fattest part of the bell-curve distribution: the Early Majority and the Late Majority, two groups with slightly different views as to what constitutes a "safe" decision to adopt. Both view the decision in pragmatic terms, with the Late Majority adopters tending towards a somewhat more conservative approach. Finally come the Laggards, those not eager to make a change, who do so only when the decision is all but inevitable. The overall adoption rate of the new idea or technology follows an S-shaped curve that is steepest just as the Early Majority hands off to the Late Majority.

Since its introduction, Rogers' basic model has proved particularly useful in explaining how new technologies are received both in business and consumer settings. The model describes

virtually any community that is in the process of adopting an innovation. For purposes of understanding SAP HANA and how it relates to the market overall, it is instructive to apply the model both to technology providers and technology users.

Technology Providers

Consider how the diffusion model applies to the introduction of in-memory processing capability among technology providers within the data warehousing industry. As noted in the previous section, SAP is clearly an Innovator where this technology is concerned. SAP launched Business Warehouse Accelerator with in-memory capability in 2005. On the other end of the spectrum, Teradata and its Intelligent Memory solution would show up in the Late Adopter or even Laggard category. Along with Oracle, Microsoft, and IBM, Teradata has been slow to respond to the promise of in-memory processing and slow to make its benefits available to their customers.

While the diffusion model applies to in-memory processing *in general* among enterprise database vendors, we cannot currently apply it to the more disruptive innovation embedded in SAP HANA. With Business Warehouse Accelerator, SAP introduced a hybrid solution in which in-memory processing augmented traditional disk-based data access; as noted, other players in this space have eventually followed suit. With HANA, SAP introduced a platform for across-the-board real-time enterprise computing. Once again, SAP is an innovator. But to date, they are the lone innovator in this space. So far, no one else is even a Laggard.

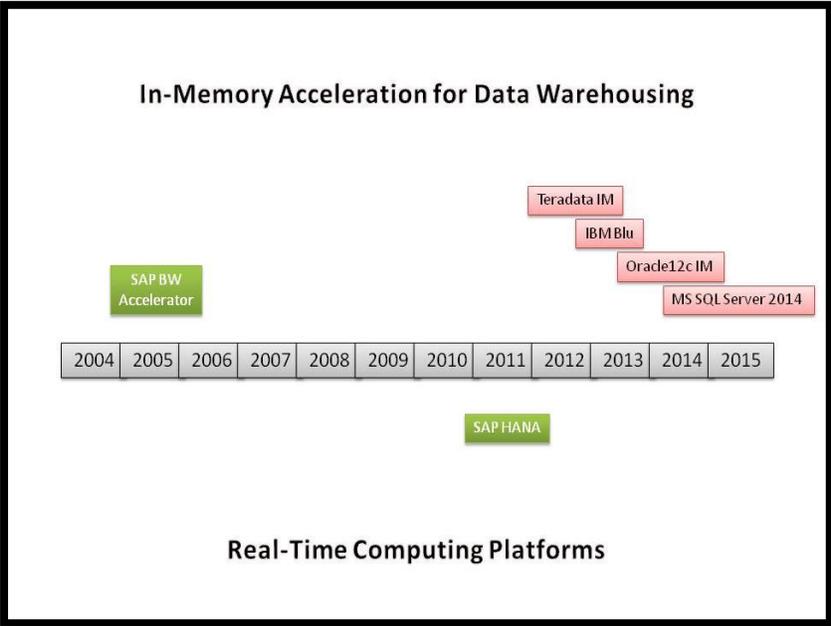


Figure 4: Technology Providers as Technology Adopters

Technology Users

The customer response to SAP HANA has been impressive from the beginning. Within a year and half of its introduction, more than 1,000 customers had adopted the HANA platform. As of this writing, more than 3,300 users have adopted HANA. These organizations use the platform to address a wide variety of business challenges. Coca Cola Enterprises implemented HANA to improve operational reporting, and has enjoyed a 725X reporting performance improvement as a result.³ Looking to enhance their supply chain management processes, Unilever implemented a HANA solution to provide real-time access to data from more than 40,000 cost centers; they also transformed their customer records reconciliation process, and now provide instant reconciliation across some 4.6 billion customer records.⁴ The NBA (National Basketball Association) has put in place a self-service statistical analysis site for fans, enabling tens of thousands of users across 200 countries to simultaneously perform analysis on the full historical record of NBA game and player stats, updated in real time.⁵

Of the 3,300-plus customers who have adopted HANA, more than 1,000 rely on the SAP Business Suite powered by HANA, which includes SAP ERP (Enterprise Resource Planning), SAP CRM (Customer Relationship Management), SAP SRM (Supplier Relationship Management), SAP SCM (Supply Chain Management) and SAP PLM (Product Lifecycle Management). Perhaps even more impressive is the large number of customers who have adopted HANA not to support an existing SAP infrastructure, but because they were facing a specific challenge and they saw SAP HANA as the best (or only) technology that could address that challenge.

In Crossing the Chasm (1991), Geoffrey Moore applies Rogers' innovation diffusion model to the introduction of disruptive technologies, exploring why some technologies catch on and ultimately become standards, while many others fall by the wayside. Moore identifies the critical juncture as occurring between the Early Adopters and Early Majority, with the "chasm" that must be crossed lying between the two. The chasm exists because of crucial differences between these two groups: differences in motivation, in how they view the technology, and ultimately in how they use the technology. While the Innovators and Early Adopters (Moore refers to them as "Technology Enthusiasts" and "Visionaries," respectively) put a premium on novelty and being ahead of the game, the Early Majority (called "Pragmatists" by Moore) are much less interested in being first or trying what's new than they are in making use of the benefits of the technology. They want those benefits, but they are also interested in being reassured that they are making a sound move. A disruptive technology can become mainstream only if it crosses that chasm. If the Pragmatists aren't convinced that the benefits of adoption outweigh the risks, the new technology will never have the momentum it needs to displace the older model. (Something else will have to come along later and do that.) But if they *are* convinced, the new technology has an excellent chance of continuing on its disruptive path and setting a new standard for the market.

³ "Coca-Cola Has Major Performance Improvement with SAP HANA." Retrieved from <http://goo.gl/3Mxnc1>.

⁴ "Unilever Uses SAP HANA to Enhance its Supply Chain Decision Making." Retrieved from <http://goo.gl/7PHEDJ>.

⁵ "NBA Enhances Fan Experience with SAP HANA." Retrieved from <http://goo.gl/RwCSKU>.

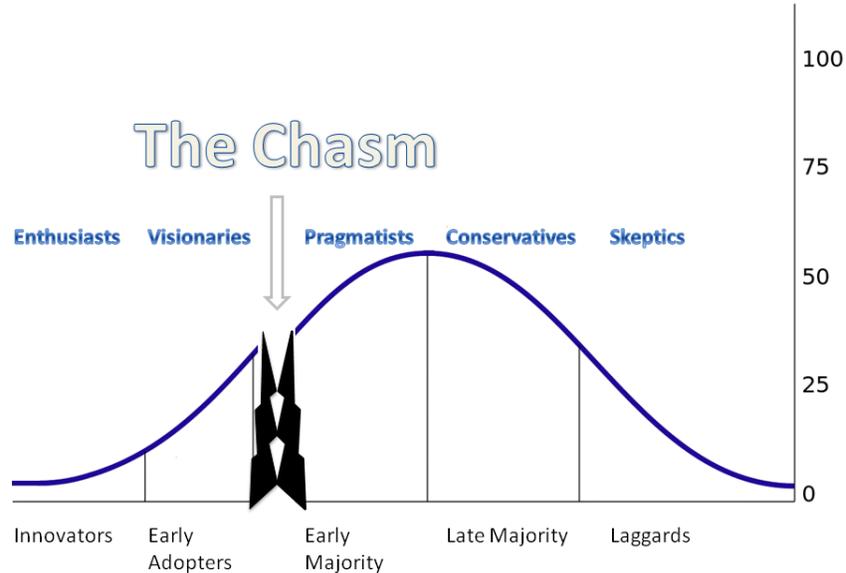


Figure 5: Geoffrey Moore's Technology Adoption Model from Crossing the Chasm

SAP introduced HANA with disruption in mind, knowing full well that the new platform would succeed only if the benefits of the technology translate into the kind of business value that will motivate users to try such an out-of-the-box approach. Geoffrey Moore recently commented on the reception that HANA has received from the SAP user community and how this transition maps to his model:

If you think about the history of enterprise IT, we've always had two domains of computing: one for online transaction processing (OLTP) and then one for online analytical processing (OLAP.) The transformative thing of HANA is you don't have to do that anymore, which means all of the ETL — all of the extraction, transformation, and loading that you would take to get things out of your OLTP environment and into your OLAP environment — the time, the errors, just the sheer magnitude of that task — that's all eliminated. What's interesting about HANA is as a technology it's clearly crossed the chasm: SAP has put the entire Business Suite on top of it. There's over a thousand customers. It's well into the mainstream.⁶

In Moore's analysis, HANA has crossed the chasm from Visionaries to Pragmatists because SAP customers see the tremendous business value that the new model represents.⁷ The differences in motivation between the various groups of technology adopters are core to the

⁶ "SAP HANA Crossing the Chasm." Retrieved from <http://goo.gl/bna6sp>.

⁷ Moore's statement that SAP HANA has "crossed the chasm" refers both to the characteristics and motivations of the customers purchasing HANA. They are implementing HANA for reasons that align closely with the Early Majority or Pragmatist categories. While thousands of customers have adopted HANA and thousands more are in the process of doing so, Moore's assessment does not map to the numeric breakdowns of the Everett model and is not a projection of adoption rates per se.

message of *Crossing the Chasm*. Moore explains that because each group has its own interests and concerns, each must be subject to its own marketing strategy. A company introducing a disruptive technology that offers the same messages and value proposition to all the various adopter groups is making a significant misstep. Likewise, technology users misstep when they entertain the idea that all technology vendors providing a similar (or similar-sounding) solution are the same. Where those vendors fall on the technology adoption curve says a good deal about how credible their innovative offering is; whether it *is*, in fact, innovative; and what the prospects are for that company going forward. After all, few businesses would want to be tied to a technology solution provided by a company that may be on its way out.

When Change Comes Too Late

In *The Innovator's Dilemma* (1997), Clayton Christensen explores the factors that prevent established companies from responding effectively to disruptive technologies that threaten their existing business. Ironically, the primary stumbling block for companies in that position is a set of management practices that are normally lauded as the ideal. In business, everyone knows that it is important to talk to the customer, find out what the customer needs, and provide products and services that meet those needs.

Disruptive innovation, however, significantly changes that dynamic. A truly disruptive technology does not just meet the customers' needs; it anticipates whole new categories of customer need and expectation. As a new technology gains prominence, the shift in customer expectation can happen in an extremely short period of time. Customers discover needs they didn't have before and will quickly turn to a provider that can meet them. Companies that have built very successful businesses addressing a particular market sector in an established way have a difficult time conceptualizing and implementing the level of change that would be required to keep them relevant in the face of this kind of challenge.

Christensen explains:

Despite their endowments in technology, brand names, manufacturing prowess, management experience, distribution muscle, and just plain cash, successful companies populated by good managers have a genuinely hard time doing what does not fit their model for how to make money. Because disruptive technologies rarely make sense during the years when investing in them is most important, conventional managerial wisdom at established firms constitutes an entry and mobility barrier that entrepreneurs and investors can bank on. It is powerful and pervasive.⁸

These companies all too often end up as victims of their own success. It is worth noting that the full title of Christensen's book is *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Teradata and the traditional database vendors have missed the opportunity

⁸ Christensen, Clayton. *The Innovator's Dilemma*, 1997.

to establish themselves as leaders, first in in-memory technology and more recently in real-time enterprise computing. Their slow adoption of the former and (to date) non-adoption of the latter show just how deeply their traditional go-to-market strategies are embedded in the old paradigm. Such a position indicates (at the very least) some difficult times ahead as the new paradigm gains momentum. Teradata will be forced to make a number of difficult and painful adjustments in order to remain viable in a market that is rapidly moving away from their model.

Contrast their position with that of SAP, which has invested both the time and resources required to establish a leadership position, and to ensure ongoing viability. SAP is clearly in a stronger position than Teradata to meet the evolving needs of customers who require a data warehouse solution, particularly one built on an SAP Business Suite environment. In the following section, we will take a closer look at the relative positions of the two companies and their offerings to gain a better understanding of the superior position that SAP occupies.

III. Business or Technology Solutions?

A core difference between how Teradata and SAP position themselves can be identified in the relative attention each of the providers places on solving technology problems versus business problems. As the late-adopting player where technology is concerned, it might be natural to expect Teradata to play up business issues rather than risk exposing the technology lag. As the technology leader, we might expect SAP to showcase their technological advantage, making it the center of their message. But the somewhat counterintuitive reality is nearly the opposite of that. The reasons for that bear exploring.

A Tale of Two Architectures

In outlining its overall value proposition in the era of Big Data, Teradata presents what it calls the Unified Data Architecture (UDA). The UDA rests on three core capabilities: data warehousing, data discovery, and a “data platform” for loading and storing data as well as preparing it for analysis. The data warehouse layer in this architecture is provided by the Teradata Integrated Data Warehouse. The data discovery layer is provided by the Teradata Aster Discovery Platform. For the data platform, users can choose between Hortonworks Hadoop or Teradata’s proprietary Integrated Big Data Platform.

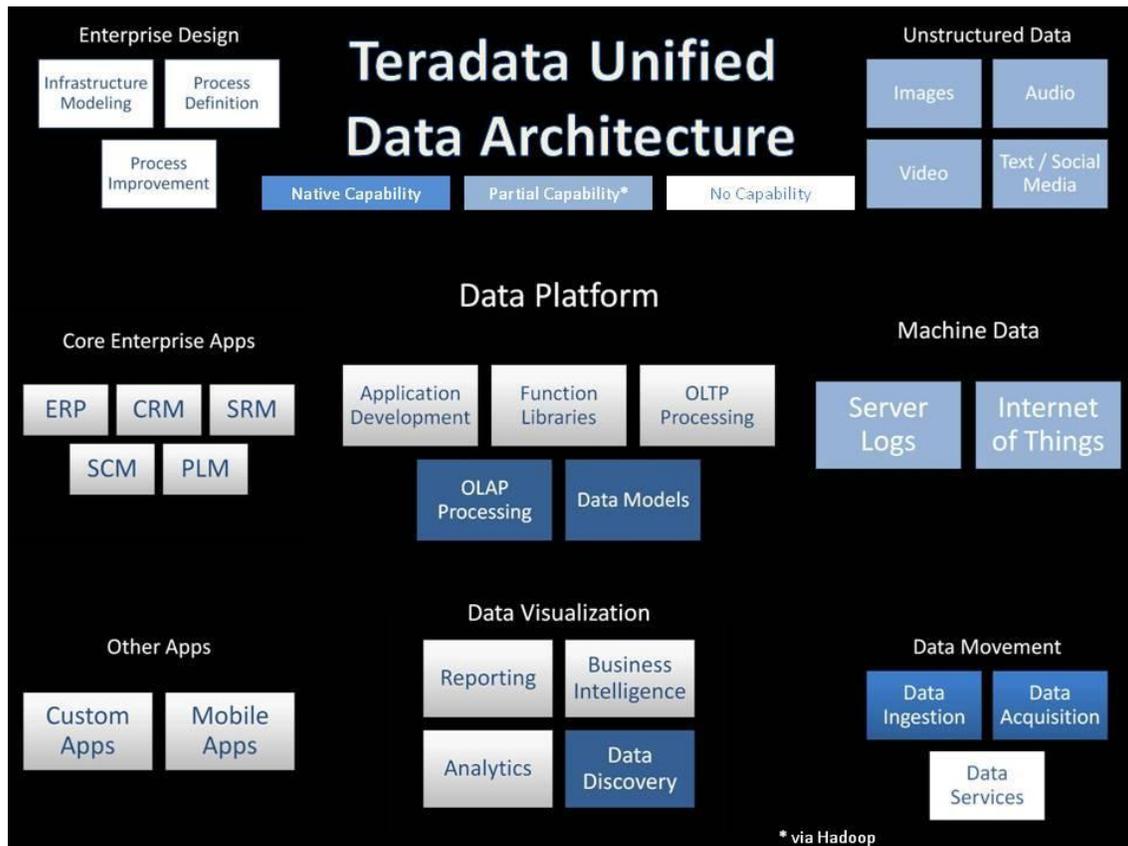


Figure 6: The Teradata UDA in Enterprise Context

Figure 6 maps the Teradata UDA against the core capabilities that an enterprise must leverage in order to manage big data effectively. The diagram distinguishes between capabilities that Teradata provides natively, those that it can make partial claim to via its use of Hadoop, and those which are not part of Teradata’s offering. Obviously, inasmuch as Teradata is a data warehousing solutions provider that makes no claim to providing capabilities such as Enterprise Applications or OLTP functionality, it is no criticism to point out that these capabilities are not native. However, it becomes immediately clear that the “unified” data architecture is in fact a fairly piecemeal construct when viewed in the enterprise context. This becomes even clearer when the Teradata UDA is contrasted with the SAP HANA In-Memory Platform.

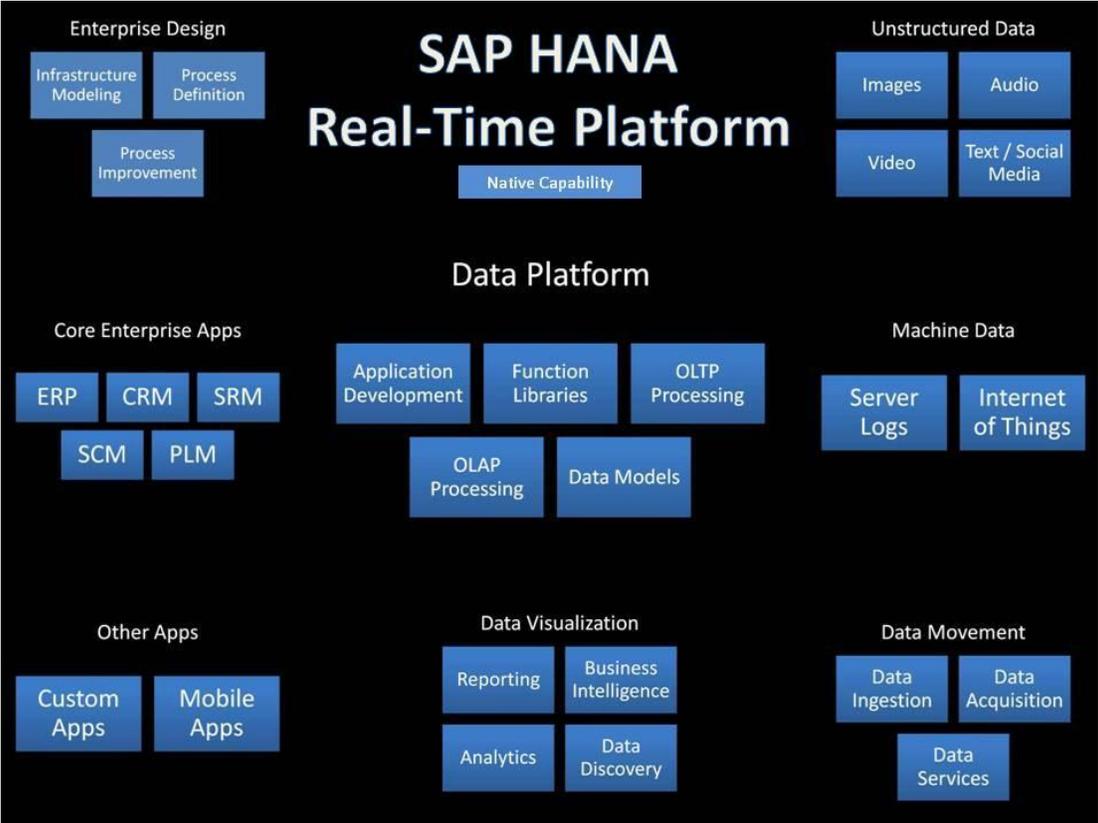


Figure 7: The SAP HANA Real-Time Platform in Enterprise Context

The jarring difference between these two diagrams (Figure 6 and Figure 7) clearly demonstrates why SAP positions itself primarily as a provider of business solutions, while Teradata positions itself as a provider of technical solutions. SAP provides technology that supports every aspect of the business. The understanding of business processes that is established in the design phase is reflected throughout the execution of those processes via transactional systems and the analysis and updating of those processes via business intelligence (BI) and analytics tools. The data structures embedded in the applications making up the SAP Business Suite are the same as those embedded in the SAP Business Warehouse. When running a true real-time enterprise on SAP HANA, one database supports both operational and analytics environments. Whenever SAP approaches any problem within the enterprise, whether it is related to organizational,

operational, or analytical issues, it does so with the benefit of this broader perspective. SAP is always able to keep business concerns where they belong: front and center.

Teradata, on the other hand, can't help but see the enterprise in a somewhat disjointed way, identifying areas where it has a direct role to play and areas where it has little or nothing to contribute. Trying to slot its offerings into the broader enterprise will always raise a host of technical issues. It is for this reason that Teradata tends to focus on point solutions. From the enterprise perspective, a Teradata data warehouse ultimately *is* just a point solution. Within the new paradigm that SAP HANA has introduced, a Teradata solution is simply another data silo.

IV. Different Perspectives; Different Solutions

Because of their fundamentally different positions on the technology adoption curve, SAP and Teradata have widely divergent perspectives both on what is happening in the market and how best to meet their customers' needs. As a result, they bring very different solutions to the table. Teradata approaches the challenges of enterprise data warehousing for SAP environments with a point-solution mentality and an evolving storyline about which technologies customers find beneficial and which they should avoid. This approach stems in part from Teradata's positioning as a mature technology in the data warehousing space. In the typical product lifecycle breakdown, a mature product is one that addresses a longstanding problem. In Teradata's view, the principal challenges related to data warehousing have been long solved. Improvements to the Teradata technology offering are therefore incremental changes around the margins. In their view, there is no big work left to be done.

By contrast, SAP has a long history of providing fully integrated business solutions for the enterprise, backed by an evolving set of value propositions that emerge from a long-term perspective and a commitment to technology leadership. SAP consistently seeks long-term and fundamental change, as demonstrated by its ongoing disruption of the data warehousing market.

Teradata and the Point-Solution Mentality

The difference in approach is vividly illustrated by the recent announcement of Teradata Analytics for SAP. Teradata claims that this new offering "does away with the need for SAP BW (Business Warehouse) altogether."⁹ The product consists of a collection of pre-packaged reports and dashboards, a data model, and a set of complex ETL scripts required to transport data from the SAP ECC (ERP Central Component). Essentially it is a scaled-down workaround for SAP BW, an add-on for SAP environments to make them more Teradata friendly. In addition to providing an opportunity to sell into SAP environments, Teradata claims that the rationale for the new product is to address performance issues associated with SAP BW and to "free" users from having to deal with the "overly complex" SAP schema.

This rationale reveals two very important characteristics about Teradata's approach to addressing challenges:

- It is short-sighted, if not firmly rooted in the past
- Its focus is technical workarounds and point solutions rather than business value

SAP takes a very different approach. As noted earlier, improving SAP Business Warehouse performance has been a priority for some years to enable users to vastly deepen their analysis of strategic data. To increase the reach and power of customer analysis, SAP developed the SAP Business Warehouse Accelerator, a fully integrated solution based on groundbreaking

⁹ "Teradata launches analytics product as SAP BW replacement," *Computer Weekly News*, 01/27/2014

technology. With their implementation of the BW Accelerator, Oracle reported results of “10-100x increased ad hoc query performance.”¹⁰ IBM, meanwhile, published scalability benchmarks showing no degradation in reporting performance of 100,000 reports per hour between systems with 5 TB of data (6 million records) and 25 TB of data (37 million records), with one IBM customer commenting, “Some reports that formerly took more than an hour now take less than a minute...”¹¹

Teradata’s claims about problems with BW performance apparently originate from a time before the introduction of BW Accelerator. Certainly any remaining performance problems have more recently been addressed by HANA, which continues to demonstrate a commitment both to providing a fully integrated solution and to opening up new technological horizons. With HANA, Business Warehouse continues to evolve, providing new options and capabilities unlike anything that has been available in the past.

For example, running on SAP HANA Live, BW data is immediately available in a highly usable format, similar in some respects to what Teradata offers via its Analytics for SAP environment. Of course, in the SAP environment, accessing this data does not require the time-consuming and complex ETL processes that Teradata requires. The HANA Live BW data can immediately feed into more than 2,000 pre-defined and customizable reports, contrasted with the 150 reports included in the Teradata package.

SAP BW on HANA addresses the broadest possible set of requirements for analysis and operational reporting for SAP ECC products. Users with needs that go beyond what is included in the 2,000 pre-defined reports will find that all the schemas for CFOs and other functions are already defined in BW. Both SAP Lumira and SAP Business Objects provide tremendous freedom for users to define analysis and reporting of results in a way that works best for their organization’s needs — and all seamlessly, natively integrated with the entire SAP environment. The SAP approach addresses the challenges of reporting and analysis as business problems to be solved at the enterprise level. The Teradata approach is a patched-on point solution requiring extensive effort even just to bring in the data, as well as additional workarounds and customizations in the very likely event that the 150 reports provided will be less than a perfect match for the organization using the product.

Teradata’s Evolving Story Line on In-Memory

After taking an outwardly skeptical and highly dismissive stance towards SAP HANA, Teradata finally announced its own hybrid in-memory solution, Intelligent Memory, with Teradata database release 14 in May, 2013. In much the same way that SAP HANA evolved (in part) from SAP Business Warehouse Accelerator, Intelligent Memory is a next-generation implementation of Teradata Virtual Storage. Using a temperature metaphor to gauge the activity

¹⁰ “SAP Business Warehouse Accelerator on Oracle’s Sun Systems.” Retrieved from <http://www.oracle.com/us/solutions/sap/infrastructure/business-intel-data-warehouse-322422.html>

¹¹ “IBM Systems solution for SAP Business Warehouse Accelerator.” Retrieved from <http://www.ibm.com/solutions/sap/us/en/landing/N367059H83793W50.html>

level of data — frequently accessed data categorized as “hot;” infrequently used data as “cold” — Teradata Virtual Storage was designed to optimize Teradata environments by moving data to the appropriate storage tier. Teradata Intelligent Memory uses the same temperature metaphor to determine which data to write to disk, and which data to keep in-memory.

Positioned as a “reasonable alternative” for in-memory computing and contrasted with the “extreme approach” that HANA represents, Teradata Intelligent Memory rests on the Pareto-inspired observation that 20% of the data in a data warehouse accounts for 80% of the activity. By keeping only that “hot” 20% of the data in memory, the Intelligent Memory solution reduces disk I/O. Using Intelligent Memory, Teradata claims, an organization gets most of the benefits of an in-memory solution without incurring the expense or complexity of committing the entire data warehouse to RAM.

Teradata has gone from rejecting in-memory processing to saying that it is appropriate for 20% of an organization’s data. It seems unlikely that this will be their last word on the subject. As the technological and market changes that enabled SAP to develop HANA become even more evident, that 20% figure is likely to go up. Sooner or later Teradata will feel pressure to flip those percentages to make the case that 80% of data should be in memory, while 20% should be on disk. Finally, the advantages of the new paradigm will be so overwhelmingly obvious that even Teradata will have to admit that having 100% of the data warehouse in memory is the only way to go. By then, SAP’s head start in this space — which is already a matter of years — will be completely insurmountable, if it is not already.

Comparing Costs and Value

The difficulty in comparing costs between a Teradata solution and SAP HANA goes back to the stark contrast between the two architecture diagrams shown in Figures 6 and 7, as well as the widely different placement of the two companies on the technology adoption timeline as shown in Figure 4. A typical Teradata success story is that of an international frozen foods company that has recently reported productivity and performance gains primarily through optimizing (and outsourcing) its ETL processes.¹² Contrast that success with the experience of a major university that has recently reported tremendous performance gains with HANA; the institution has completely *decommissioned* its ETL infrastructure along with numerous other now-superfluous BI assets.¹³ Another Teradata success story is that of a financial services institution that reduced its window for leveraging EDW data from two weeks to two hours.¹⁴ Consider how well those results stack up against those of a tools and infrastructure provider who report that with HANA they can now analyze 33 million records in three seconds.¹⁵

¹² “Teradata, IBM, and McCain Foods.” Retrieved from <http://goo.gl/sGkkQj>

¹³ “Early SAP HANA customers separate reality from the hype,” *Computerworld*, October 18, 2012

¹⁴ “Bank Internasional Indonesia Maybank: Investing in Change,” *Teradata Magazine Online*, Q2, 2014

¹⁵ “Hilti analyzes 33 million records in less than 3 seconds with SAP HANA,” Retrieved from <http://goo.gl/XflbZM>

For Teradata, success is defined by making measurable improvements around the margins of the established paradigm. For SAP HANA, it is established by helping organizations achieve unprecedented results by transforming how they do business. To put it simply, the two companies provide very different solutions with very different value propositions because they are selling very different technologies. Teradata is a data warehousing solution. As we have shown, SAP HANA is an enterprise platform for real-time computing, one that encompasses transactional processing, data warehousing, and applications.

In trying to compare costs between them, we have to take several factors into consideration:

- First, no comparison of costs between an enterprise real-time computing platform from SAP and one from Teradata (or any other vendor) is possible. *There is nothing to compare.* Teradata cannot provide that capability at any price.
- Likewise, if we focus only on the data warehouse, no comparison between fully in-memory data warehousing solutions is possible. *SAP provides one; Teradata does not.*
- It is possible to compare hybrid disk- and in-memory solutions between the two providers. Not every HANA implementation to date has been a full instantiation of enterprise-wide real-time computing. As the adoption models indicate, technology users take various routes to full adoption of a new technology. SAP supports a broad range of possible configurations, depending on specific customer circumstances and the value that a given customer is looking for.

In order to provide an apples-to-apples comparison, we have to step back from the new paradigm and allow SAP HANA and Teradata to meet on the same playing field. In this instance, that playing field is a hybrid disk- and in-memory data warehouse environment. Via its Intelligent Memory option, such a configuration represents the very cutting edge of what Teradata is able to deliver. For SAP HANA, on the other hand, such a configuration represents a substantial downgrade from its full potential performance capability.

A typical Teradata Intelligent Memory solution keeps approximately 20% of data in memory while assigning the other 80% to various storage devices, depending on how active or inactive that data may be. SAP HANA Smart Data Access will support a similar environment, with 20% of the data live in HANA and 80% on disk. Via virtualization technology, the data on disk can be in any of a number of formats: SAP ASE, SAP IQ, Hadoop, or Teradata itself.

As demonstrated in internal SAP tests and live deployments, any configuration of an SAP hybrid environment will provide considerable value relative to a Teradata Intelligent Memory environment containing the same data. A good example would come from comparing Teradata Intelligent Memory with a hybrid HANA / Hadoop environment. In such comparisons, depending on the hardware used and other deployment options, the SAP solution costs roughly the same to implement, while enabling performance that is potentially 20 times as fast as the Teradata Intelligent memory solution.

Cost at Scale and Over Time

It's important to note that the cost of implementation is only the beginning of the story. Over time, a typical Teradata environment will incur costs that do not apply to the comparable HANA environment, including:

- High-cost assets for configuration
- Software costs to upgrade to faster CPUs
- Additional storage costs
- Maintenance and overhead of duplicate data
- Maintenance and overhead of aggregate data
- Maintenance and overhead of indexes

The growth of the core data in a Teradata environment is compounded by the growth of indexes and aggregates, leading to ever-upward spiraling costs of storage and maintenance. HANA users benefit from compression of data and an architecture that eliminates the need to duplicate data around the enterprise. This core difference accounts for a widening gap in costs associated with the total environment for comparably sized solutions as the initial amount of raw data increases.

Over time and as the environments scale, HANA users benefit from a sharp drop-off in cost per terabyte calculated against the raw data on which the environment is built. From a total cost of ownership perspective, a Teradata environment is significantly more expensive than the comparable SAP HANA environment — while providing significantly slower performance.

Conclusion: SAP's Growing Value Proposition

SAP's approach to solving data warehousing challenges for its customers is diametrically opposed to the approach taken by Teradata. SAP begins with two core assumptions about what users need from a business warehouse solution:

- Users don't want another silo in their architecture; they want to expand on fully integrated business processes
- Users want lasting solutions built with today's and tomorrow's challenges in mind, not point solutions and patches that are already out of date when they are implemented

This perspective has enabled SAP, as both a technology leader and the world's leading authority on enterprise business processes, to keep up the flow of game-changing improvements for SAP customers (old and new alike.) This process has led to the introduction of HANA, which is easily the most disruptive innovation — not just for the database market but for enterprise computing generally — in the past two to three decades.

Perhaps most impressively, SAP has rolled out this multidimensionally disruptive platform in a way that is both non-intrusive to the existing user base and highly appealing to new users, primarily because of the unprecedented capabilities that it brings. And SAP has accomplished this technological feat without sacrificing its core emphasis on solving business problems. In fact, SAP has significantly expanded and deepened its support of business processes at all levels. As detailed in this paper, there are fundamental differences between SAP and Teradata in perspective, adoption of technology, and execution on a long-term strategy. A careful review reveals that the SAP approach offers tremendous advantages over the Teradata approach. .

References:

1. "Remember the titans: who will win the in-memory battle?" *Information Age*, April 16, 2014.
2. Word, Jeffrey. *SAP HANA Essentials*, 2013.
3. "Coca-Cola Has Major Performance Improvement with SAP HANA." Retrieved from <http://goo.gl/3Mxnc1>.
4. "Unilever Uses SAP HANA to Enhance its Supply Chain Decision Making." Retrieved from <http://goo.gl/7PHEDJ>.
5. "NBA Enhances Fan Experience with SAP HANA." Retrieved from <http://goo.gl/RwCSKU>.
6. "SAP HANA Crossing the Chasm." Retrieved from <http://goo.gl/bna6sp>.
7. Moore's statement that SAP HANA has "crossed the chasm" refers both to the characteristics and motivations of the customers purchasing HANA. They are implementing HANA for reasons that align closely with the Early Majority or Pragmatist categories. While thousands of customers have adopted HANA and thousands more are in the process of doing so, Moore's assessment does not map to the numeric breakdowns of the Everett model and is not a projection of adoption rates per se.
8. Christensen, Clayton. *The Innovator's Dilemma*, 1997.
9. "Teradata launches analytics product as SAP BW replacement," *Computer Weekly News*, 01/27/2014
10. "SAP Business Warehouse Accelerator on Oracle's Sun Systems." Retrieved from <http://www.oracle.com/us/solutions/sap/infrastructure/business-intel-data-warehouse-322422.html>
11. "IBM Systems solution for SAP Business Warehouse Accelerator." Retrieved from <http://www.ibm.com/solutions/sap/us/en/landing/N367059H83793W50.html>
12. "Teradata, IBM, and McCain Foods." Retrieved from <http://goo.gl/sGkkQj>
13. "Early SAP HANA customers separate reality from the hype," *Computerworld*, October 18, 2012
14. "Bank Internasional Indonesia Maybank: Investing in Change," *Teradata Magazine Online*, Q2, 2014
15. "Hilti analyzes 33 million records in less than 3 seconds with SAP HANA," Retrieved from <http://goo.gl/XflbZM>

Copyright

© Copyright 2014 SAP AG. All rights reserved

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors.

Microsoft, Windows, Excel, Outlook, and PowerPoint are registered trademarks of Microsoft Corporation.

IBM, DB2, DB2 Universal Database, System i, System i5, System p, System p5, System x, System z, System z10, System z9, z10, z9, iSeries, pSeries, xSeries, zSeries, eServer, z/VM, z/OS, i5/OS, S/390, OS/390, OS/400, AS/400, S/390 Parallel Enterprise Server, PowerVM, Power Architecture, POWER6+, POWER6, POWER5+, POWER5, POWER, OpenPower, PowerPC, BatchPipes, BladeCenter, System Storage, GPFS, HACMP, RETAIN, DB2 Connect, RACF, Redbooks, OS/2, Parallel Sysplex, MVS/ESA, AIX, Intelligent Miner, WebSphere, Netfinity, Tivoli and Informix are trademarks or registered trademarks of IBM Corporation.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.

Adobe, the Adobe logo, Acrobat, PostScript, and Reader are either trademarks or registered trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Oracle is a registered trademark of Oracle Corporation.

UNIX, X/Open, OSF/1, and Motif are registered trademarks of the Open Group.

Citrix, ICA, Program Neighborhood, MetaFrame, WinFrame, VideoFrame, and MultiWin are trademarks or registered trademarks of Citrix Systems, Inc.

HTML, XML, XHTML and W3C are trademarks or registered trademarks of W3C®, World Wide Web Consortium, Massachusetts Institute of Technology.

Java is a registered trademark of Sun Microsystems, Inc.

JavaScript is a registered trademark of Sun Microsystems, Inc., used under license for technology invented and implemented by Netscape.

SAP, R/3, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP Business ByDesign, and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries.

Business Objects and the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence, Xcelsius, and other Business Objects products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of Business Objects S.A. in the United States and in other countries. Business Objects is an SAP company.

All other product and service names mentioned are the trademarks of their respective companies. Data contained in this document serves informational purposes only. National product specifications may vary.

These materials are subject to change without notice. These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.