

WHITE PAPER

Blending Transactions and Analytics in a Single In-Memory Platform: Key to the Real-Time Enterprise

Sponsored by: SAP

Carl W. Olofson

Henry D. Morris

February 2013

IDC OPINION

Business applications started out decades ago as systems of record, capturing sales orders, trouble tickets, journal entries, and other records of business operations. Line-of-business personnel needed access to this information to decide how to price, plan, and set priorities as examples. IT then built separate analytic decision support systems on completely different platforms. Separate systems for separate workloads on separate platforms sounds logical in theory. And for years, separate platforms has been the accepted approach to ensure good performance for transactions and analytics, given the different requirements for capturing data versus analyzing data. In practice, however, the requirement to transfer data between the two platforms underlying the two types of systems causes information delays and reduces the ability to connect insight to action because of a disconnect from analytics to source data. This leaves the enterprise at risk when decision makers are forced to rely on stale or insufficient data, or when they cannot connect analytic insights to source data.

But it doesn't have to be this way. Today's businesses can couple transactions (as in ERP or CRM applications) with analysis in real time in a single blended environment to determine the best way to get live insight about a fast-breaking situation. Rather than using separate transactional and analytical applications built on separate platforms, a single data management environment for both systems of record and systems of decision (assuming good performance can be achieved for both) would yield the following benefits:

- ☑ Users could access and analyze the latest data as soon as it was captured rather than waiting for data transfer — eliminating a major source of information delay. And by reducing the overhead of multiple platforms, IT's data management task and the business' data governance function would be simplified.
- ☑ With instant access to data, business personnel could make business decisions faster, based on the latest information. Business processes could also be accelerated. As an example, during an interaction, contact center employees would have the latest customer data available and understandable. A cross-sell recommendation could be made on the most recent customer orders, or even on an order that is in process.
- ☑ More innovative business processes and practices could be developed, because when analytics can be applied directly to systems of record, the enterprise can be nimble enough to adjust to changing business conditions and take advantage of opportunities as they occur, even transforming and innovating its business models.

METHODOLOGY

This white paper on a single platform for transactions and analytics was prepared based on the following:

- ☒ Extensive background information gathered by IDC over a number of years, including the most recent innovations in database technology
- ☒ Detailed information provided by SAP regarding the technology and use cases for SAP HANA, an in-memory platform that provides a single data management environment for systems of record and systems of decision
- ☒ Data gathered from a survey executed on behalf of SAP that polled 759 respondents from every major geographic region, including 405 IT managers and 352 business managers, regarding the business problems and technology challenges they see with separate platforms for transactional applications (systems of record) versus analytic applications (systems of decision) with data managed separately in transactional versus analytical databases

IN THIS WHITE PAPER

This white paper discusses the issues involved in the traditional practice of deploying transactional and analytic applications on separate platforms using separate databases. It analyzes the results from a user survey, conducted on SAP's behalf by IDC, that explores these issues. The paper then considers how SAP HANA, with its combination of in-memory data management and its ability to handle both transactions and analytics in real time, can resolve these issues. It explores how businesses may find opportunities for innovation (such as the ability to engage in a richer dialog with a customer based on analysis of the latest transactional information), for speed (with the ability to provide faster access to information to make timely decisions), and for simplification of the IT landscape with a single in-memory platform.

SITUATION OVERVIEW

The Challenge: The Divide Between Systems of Record and Systems of Decision

IDC's survey of IT managers found that over 40% required more than 2 days to prepare financial data for reporting. On the line-of-business side, it's therefore not surprising that 25% of end users indicated that faster access to information would have a significant impact on their organization. This is consistent with IDC findings over the past 10 years that over 70% of the effort in an analytics project is dedicated to preparing the data for analysis.

Part of the reason for this information delay is that there can be multiple systems of record that have captured data relevant to making a decision. Some are run within the enterprise, but in departments separate from enterprise IT. Think of a supply chain that must coordinate activities across multiple departments and multiple firms. Supply chain planning requires information about demand and inventory levels that cut across enterprises. So visibility into multiple systems is needed, and optimal ways to define the relationships among the data sources and then bring the information together are needed.

But it is also the case that even information captured within the enterprise, and under the control of central IT, can face delays getting to decision makers. This is often caused by the issue of separate platforms underlying the transactional application (system of record) and the analytic application (system of decision). Bridging this divide requires data transfer, delaying the decision maker's access to new information. The ability to deploy both transactional applications (systems of record) and analytic applications (systems of decisions) on the same data management platform would address a key cause of information delay.

Technical Issues Causing Information Delay

Separate data management platforms for systems of record and systems of decision are key causes of information delay. While transactional applications need data to be organized in a way that helps serve business transactions, analytic applications need the data to be organized in ways that help serve business intelligence purposes such as historical trend analysis. Over time, it became too difficult to optimize a disk-based database to serve both purposes at once, so separate database architectures emerged — one for transactions and one for analytics. Since enterprise data had become badly fragmented by this point, such data needed to be collected from the various transactional databases, converted to a common form, and loaded into this analytic database for analysis and reporting.

There can be a role for a separate enterprise data warehouse that brings together relevant information from within and outside the enterprise. But real-time operational reporting and analytics can be streamlined, addressing information delay, by running an enterprise application suite and reporting applications on a single data management platform. This provides the business user visibility into data from the moment the transaction is captured.

Business Problems and Costs Resulting from Information Delay

- ☒ Because data in a separate analytics platform is usually not current, it is not possible to run analyses on such data that can usefully inform decisions made in the midst of a business process. The information is only as current as the last time the data was moved to the analytics platform and transformed for analytic usage. The inherent latency of the data inhibits the business in the following ways: The speed of business processes is confined to the performance characteristics of the transactional database, which, in turn, is dependent on the arcane tuning of its disk-based operations. If some analytics are introduced into the process, then the performance of the analytical database and of the data movement software can further slow down business processes.
- ☒ If analytics data is used in interactions with customers and suppliers, such interactions can be complex and cumbersome as sales, support, and other staff must switch awkwardly between their transactional applications and the analytic applications that provide relevant reports and visualizations. The ideal solution would be to integrate real-time analytics into the transactional application, but this is problematic when the transactional and analytical data is kept in different places. For example, a contact center agent would have to look away from the transactional application during an interaction and try to find relevant analytics. Even if this could be done, it's likely that the analysis will have been based on data gathered prior to the current interaction episode.
- ☒ Without the integration of analytics into transactional applications, there is little opportunity to build into such applications the flexibility and dynamism to adjust to shifting business circumstances. Building applications that can optimize business operations, or recognize cross-sell and upsell opportunities, is extremely difficult, and business innovation is inhibited.

What Users Told Us

Issues Associated with Separate Applications Running on Separate Database Platforms


The survey indicated that users encountered substantial waiting periods in moving data from transactional to analytical databases. For instance, the survey asked the 407 IT manager respondents, "On average, what is the estimated time to process the data and make it available for analysis by the business user?" As indicated in Figure 1, there were statistically significant responses in each of these areas: finance, sales, marketing, customer service, manufacturing, supply chain/logistics, and R&D.

Most of these areas showed that such movement required more than a day in over 50% of cases. Some areas showed over 50% indicating that such movement required more than 2 days. A large percentage in each case indicated that intervals of a week to 10 days are not uncommon. It is hard to imagine any nimbleness in business execution with such long delays between the setting of transactional data and the time it is available for analysis. By the time users analyze and determine a course of action, that course of action may be irrelevant.

FIGURE 1

Time Required to Move Data from Transactional to Analytical Database by Business Area

Q. On average, what is the estimated time to process the data and make it available for analysis by the business user?



n = 407

Source: IDC's SAP Survey, 2012


Business Users Settle for the Unacceptable

The survey found that the majority of business users found the slow and fragmented nature of their data systems to be perfectly OK. For instance, as shown in Figure 2, most were satisfied with the time it took to get special reports.

FIGURE 2

Satisfaction with Ad Hoc Reports

Q. When you need a change made to an existing report or need to get an ad hoc report, how satisfied are you with your IT department's response time?



n = 352

Source: IDC's SAP Survey, 2012


Consider the results illustrated in Figure 3 — from IT in the same survey of the same organizations — indicating how long such special requests took to fulfill.

Since the majority of IT responded that a report request required at least a day to fulfill, and in many cases multiple days, it seems clear that either the business users polled had very low expectations in terms of report delivery or they don't make report requests often. It also seems clear that few business users expect to be able to use such reports to help with in-the-moment business decisions.

FIGURE 3

Time Required to Fulfill Operational Report Request

Q. On average, how long does it take to complete an operational reporting request for a business user?



n = 407


Source: IDC's SAP Survey, 2012

Business users don't expect to have the data they need to make decisions in a timely manner. Over 60% of those polled reinforced this view. How bad is the problem? We asked IT how long it took to move data from transactional to analytical databases where business users could employ the data to help make key decisions; their responses are indicated in Figure 4.

FIGURE 4

Data Movement from Transactional to Analytical Databases

Q. On average, what is the estimated time to process the sales, finance, or R&D data and make it available for analysis by the business user?



n = 181

Source: IDC's SAP Survey, 2012


Over 20% said that it took more than 2 days to move sales data from transactional to analytical databases. For finance data, the total was over 40%, and nearly 10% said that it took *over 11 days*. No wonder there's not much flexibility in developing analyses and reports, because nearly 50% had to wait at least 2 days for R&D data, and 20% had to wait almost a week.

With the complexity, inefficiency, and inflexibility of current data systems, any analytic activities that pertain to in-flight business processes are largely unavailable, especially those that involve complex analysis. This is illustrated in Figure 5.

FIGURE 5

Business Tasks Made Unmanageable by Technology Limitations

Q. Please select all of the business tasks that your users cannot manage today because of technology limitations.



n = 407

Source: IDC's SAP Survey, 2012


Well over 40% of IT polled said that users could not use predictive analysis or simulations or work with real-time data because of these limitations. Big Data workloads and diving down into more granular data than is routinely available were also off the table for almost as many respondents.

Benefits of a Unified Data Management Platform

The survey asked the 352 business manager respondents which of the benefits associated with a unified data management platform are the most significant for their organization or business unit. Their responses are illustrated in Figure 6.

FIGURE 6

Changes Likely to Have Greatest Positive Impact on the Business



Source: IDC's SAP Survey, 2012

It is significant to note that the three top benefits have to do with speed and accuracy, which are associated not only with a unified data platform but also with in-memory data management.

The Solution: A Unified System of Record and Decision

What It Means: Making the Enterprise More Agile to Manage

Bringing together systems of record and systems of decision means that as business processes execute, business managers can make decisions based on the current state of the business and can bring in a range of supporting facts that can make those decisions stronger. This requires that transactional and analytical data be managed together. It also means that queries, reports, and visualizations for decisions makers must be generated not in days or hours but in seconds. It means that the system must support large amounts of supporting facts and be able to bring those facts to bear immediately upon demand.

In response to a number of our survey questions, over 60% of business users seemed OK with an environment that could not deliver timely information. We wondered how this could be until we realized that this is the normal state of play in most enterprises most of the time. Business users don't complain that information is old or incomplete because this is what they are used to; it's hard to imagine a fully integrated system that can apply a wealth of up-to-the-second facts in support of decisions that they make every hour of every day.


What's harder to imagine is how such capability could change, even dramatically transform, current business processes, policies, and practices. Most of us confine or schedule business events based on the readiness of our systems to support them. So, when inventory is searched, only immediate inventory is available. When orders are taken, we wait until that data is moved to the inventory management system before starting the actions to fulfill the order. When interacting with a customer, we have no expectation of associating that customer's patterns of type or behavior with those of thousands of others, or associating those patterns with current patterns of consumption, so that we may make more informed recommendations to the customer regarding products or services that the customer might find appealing.

The business users surveyed did see the possibility that access to the right information at the right time could have a transformative effect on business processes, making them more dynamic. This fact is demonstrated clearly in Figure 7.

FIGURE 7

Most Critical Objective from Timely Information

Q. *What is the most critical objective that your organization could achieve by allowing business users to always get the right information at the right time?*



n = 352


Source: IDC's SAP Survey, 2012

As Figure 7 shows, given the opportunity to act on very timely data, over 50% of the respondents indicated that they would be able to transform business processes more dynamically. This shows an awareness that real-time data access can actually change the way business is done. Dynamic business processes are not achieved simply by making databases run faster. They are achieved by having a data platform that is both unified and highly flexible. IT is well aware that business users need more of this flexibility, but IT is constrained by the existing technology. These factors weigh on the minds of IT much more than pure speed, as is illustrated in Figure 8.

FIGURE 8

Main IT Issues Faced by Organizations

Q. *What is the main IT issue your organization faces with regard to your current transactions (OLTP) and analytics (OLAP) systems to support the business?*



n = 407

Source: IDC's SAP Survey, 2012

The top concerns were fairly divided among the issues of inflexibility, complexity, data volume, data movement from transactional to analytical databases, and security risks, with the top issue being that of complexity (21%). This suggests that IT is frustrated by its inability to fully meet the needs of the enterprise but understands that those needs revolve around flexibility which, in turn, is dependent on a much simpler data environment.

It should be noted that changing a major core system of any kind results in such fundamental change. The results are not simply incremental improvements in the quality of current processes and decision making, but in dramatically altered processes and new classes of decision that bring the enterprise greater business opportunity and operational efficiency than can be imagined currently.

How It Is Done: Building a Unified High-Speed Data Management Platform

The obvious solution to the problems associated with running separate transactional and analytical databases is to merge them into a single database environment. For reasons already discussed, however, this has not been technically feasible until now.

The reason this problem can be solved now is that, because of the dramatic increase in processor power, fall in memory prices, and commonplace support for 64-bit memory systems, the relational database management system (RDBMS) world is moving rapidly toward memory-based databases. Such databases treat main memory, rather than disk, as the home for the data they manage. Some memory-based databases swap the data to compressed disk spaces, but others manage the data entirely in memory.


In-memory databases (IMDBs) offer key advantages over disk-based RDBMSs. For one thing, they are much faster. They are also more flexible, handling schema changes in minutes, whereas disk-based systems require days and sometimes weeks to make such changes because the data needs to be restructured on disk. This is important because businesses frequently need to make changes to the kinds of data they process and analyze but are usually constrained by their databases' inability to support such changes in a timely manner.

The IT managers surveyed indicated the benefits they saw from IMDBs, which centered around simplicity of management and ease in handling large and complex data tasks, as indicated in Figure 9. Interestingly, the top 3 improvements cited were better utilization of IT staff, improved management of Big Data, and simplification of the IT infrastructure. Such improvements are best realized by unifying data and delivering it on a more flexibly managed and high-performance platform.

FIGURE 9

Anticipated Benefits of IMDBs

Q. Which one of the following improvements would you expect to have the most positive impact on your organization?



n = 407

Source: IDC's SAP Survey, 2012

The better utilization of IT staff, as acknowledged by 24% of respondents as a top benefit, is the result of the fact that an IMDB does not require constant tuning of data allocations to disk, index refinement, storage tuning and reallocation, unloading and reloading of the database, and other time-consuming routine tasks that chew up so much staff time with disk DBMSs. One reason that the IMDB affords better handling of Big Data (and other requirements that can cause schema changes), as acknowledged by 23% of respondents, is that with an IMDB, schema changes can be done in minutes rather than days, as is commonly the case with disk-based databases. Another key advantage that an IMDB offers is that it can flexibly handle both transactional data and analytical data by optimizing the way it holds each kind of data while it is being altered or queried. This means that the IMDB offers a technology

base for accomplishing the kind of unified data platform that was not possible in the past. Such a platform promises to yield the following business benefits:

- ☒ Faster access to real-time information
- ☒ Faster simulation, planning, and analysis
- ☒ Detailed access to granular information
- ☒ More contextual information
- ☒ Less dependency on IT for business users

SAP HANA: A Unified In-Memory Data Platform

How SAP HANA Unifies Transactional and Analytical Workloads

SAP has introduced its in-memory database platform in several steps, from analytics to business transactions:

- ☒ SAP HANA was introduced 18 months ago and primarily used for real-time analytics.
- ☒ SAP HANA powers the SAP NetWeaver Business Warehouse as enterprise data warehouse.
- ☒ SAP HANA powers the SAP Business Suite as another database choice for SAP customers, optimized for real-time business.

SAP HANA is not just a database. It is an application execution platform that offers integrated in-memory database management and application management, designed to run together in a way that optimizes performance. It handles both analytical and transactional workloads so users may deploy online transaction processing (OLTP) and online analytical processing (OLAP) applications in the same environment. It runs on a cluster, which gives it scalability via the ability to add nodes when needed. It can also act as a SQL database for applications that do not run within the SAP HANA platform.

What this means is that the moment a business event takes place, its data is available for analysis, driving immediate in-process decisions, thus unifying systems of record and decision. If, for example, a sales order is placed, the details about that order will be included in the next analytical query regarding today's sales, even if the order happened only a millisecond ago. It also means that analytical results can inform the application as it runs, enabling it to run "smarter." For instance, that same sales application might advise upselling the customer with an additional offer while the sale is still underway, because the sale in question fits a particular pattern that suggests that this customer might be interested in an additional product that was also ordered by customers whose orders fit the same pattern.

SAP HANA can now also be used as the platform for SAP Business Suite applications (although configurations based on other database products will continue to be fully supported), and it includes a host of analytic capabilities as well. As such, it is designed to optimize the various business process tasks of the application suite while also presenting relevant real-time analytics data for enhanced decision making. This results in an application system that is not only much faster than before but also

much smarter, enabling business users to better understand and take charge of situations as they develop in real time.

The resulting system also integrates well with the SAP NetWeaver Business Warehouse as the enterprise data warehouse and also powered by SAP HANA.

Sample Use Cases Involving the Unified Workload Approach

The following are examples of use cases and their benefits arising from the use of SAP Business Suite powered by SAP HANA:

- ☒ **Insight-driven service.** Most enterprises find it difficult to analyze large volumes of service data to identify root causes of product failure and to detect potentially fraudulent warranty claims. They need to collect and organize information from the field in real time and connect the associated assets with statistical analyses of historical data in order to make smarter decisions. They also need real-time monitoring of product usage to which they can apply predictive analysis to build better service business models. SAP HANA provides fast access and analysis of real-time service data from multiple sources, and supports the use of predictive analysis on sensor data from connected assets, enabling real-time monitoring of service operations, service levels, and field resource utilization. The result: new business models leveraging machine data from sold devices, real-time scheduling of field resources, and early information about potential product failures.
- ☒ **Real-time material planning.** Often, material planning decisions are driven by outdated information from scheduled materials requirement planning (MRP) reports. Also, it is generally not possible to change plans in response to rapidly changing business conditions, and holistic production planning across productions sites is usually not possible. With SAP HANA, MRP runtimes are cut significantly. Stored procedures are used to accelerate calculations such as the "stock requirements list" and the "where used list." The system can instantly update the supply network collaboration with the latest demand information. These improvements result in faster replanning and allow more frequent MRP runs in shorter cycles. They also enable "what if" scenarios in real time, and they better synchronize demand and supply data for global cross-plant planning, which reduces the so-called "bullwhip effect."
- ☒ **Insight-driven marketing.** At present, it is simply not possible to achieve a 360-degree view of the customer in real time; there's just too much data in too many sources to bring together at any given time. Making matters worse, customer connectivity through mobile devices and social networks has led to an explosion of data at speeds and volumes impossible for current systems to capture and leverage. What's needed is the ability to capture such data and discover market opportunities as they occur, responding at just the right moment as demand shifts. SAP HANA provides a central location for collecting all customer data from various sources and delivers the flexibility to slice and dice that detailed data. It also enables new analytical data mining capabilities on Big Data. The result is near-real-time monitoring and reporting on marketing performance, rule-based and event-triggered marketing automation that can orchestrate inbound and outbound activities, and precision marketing that can personalize every interaction with every customer.

FUTURE OUTLOOK

IMDBs are the future of all database technology. The question will be this: To what extent will other vendors take the comprehensive mixed workload approach of SAP HANA both driving and optimizing combined transactional and analytical applications? Already there are many IMDB products just starting to become available, but most of these are designed for specialized workloads. If customers choose to implement multiple IMDBs for specialized workloads, they will still have the problem of moving data between databases for those cases where live transactional data requires analysis and subsequent action. SAP HANA is in a strong position to offer a unified alternative to address a number of key business scenarios.

The integration of information into unified systems of record and decision is an inevitable development and necessary to enable the real-time enterprise. SAP HANA as the platform for SAP's transactional and analytical applications is a first step toward this goal. As enterprises adopt the unified in-memory goal and race toward it, those that are successful in transforming their businesses from the fragmented, ill-advised, and delayed processes of today to the integrated, complete, real-time decision-driven processes enabled by unified systems of record and decision will have a clear competitive advantage going forward.

CHALLENGES/OPPORTUNITIES

IDC research has shown that many organizations have difficulty building their business case. Given this state of affairs, SAP needs to be proactive in helping customers build their business case for an in-memory platform that underlies both analytics and transactions. IT operations factors should also be considered. In addition, SAP should assist organizations in measuring business benefits derived from business innovations, faster access to real-time data, and IT simplification.

Since it seems likely that other vendors will eventually offer technology to compete with SAP HANA, providing unified functionality for key business workloads, pressure will remain on SAP into the indefinite future to continue innovating in order to maintain SAP's leadership position in this important technology area.

CONCLUSION

The changes to business processes and practices that will result from the adoption of unified systems of record and decision based on IMDB technology are likely to be just as sweeping, profound, and difficult to predict as the changes wrought by the U.S. Interstate Highway System.

The Interstate Highway System was originally intended both to give U.S. defensive forces the ability to rapidly mobilize and, on a more day-to-day basis, to provide a flexible, quick, and cheap means for shipping products between cities and for delivering perishable farm goods from the country to urban markets. But the changes that resulted included the change of passenger rail service from a profit to a cost center as people chose to drive distances rather than take the train, a decline in rail freight business as shippers chose the slightly more expensive but much more flexible option of shipping by truck rather than train, a massive migration of the middle

classes from cities to suburbs as highways made it easy to commute without taking the train, and the development of suburban shopping malls and office parks located closer to consumers and employees.

It is not unreasonable to expect a similar set of profound changes to the way we do business as the result of an IMDB-driven unified system of record and decision. Beyond simply making applications run faster and enabling more detailed, analytic-driven in-process decisions, this change opens the door to profound changes in the ways in which business is done, and many, if not most, of these changes are as difficult to predict as were the effects of the Interstate Highway System.

SAP HANA is an example of such a system. With it, the closely related business transaction areas covered by SAP Business Suite may be maintained together in a single system, and additional data used for various kinds of analysis, including trends and context provided by social media data, may be integrated with it and leveraged so that both transactions and analytics can be optimized and enriched.

With this in mind, users should consider the following:

- ☒ Examine your business systems, especially the data-intensive systems, and consider how closely they are operationally interrelated and interdependent yet managed completely separately and coordinated only through extensive inefficient and often labor-intensive processes.
- ☒ Consider the extent to which your enterprise is a slave to the fixed processes and definitions of IT, without the ability to dynamically change and innovate to meet shifting business demands.
- ☒ Imagine what your business would look like if you could process any related data together, could transact business while simultaneously using analytic data about transactions in process, and could adjust business processes and even automate business decisions based on the actual facts of the moment rather than make decisions later, based on data that takes hours or even days to collect and distribute.
- ☒ Estimate the benefits that could arise from adopting a unified in-memory data platform to drive business processes, combining both transactions and analytics in real time.

For too long, users have formed their expectations for business process integration and informed decision making based on the limitations of their existing IT systems. It is time to demand that those systems apply the full potential of IT to transform the business. Business should never shape its processes to conform to the limitations of IT, but rather, IT should stretch its capability to meet the needs of the business.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2013 IDC. Reproduction without written permission is completely forbidden.