

WHITE PAPER

Ensuring Competitiveness and Agility in a World of Constant Motion: Extreme Transaction Processing from SAP

Sponsored by: SAP

Carl W. Olofson

January 2014

IDC OPINION

Big Data technologies have captured a lot of attention lately with their ability to handle large amounts of data efficiently. They also have shone a light on the need of enterprises to establish a more nimble, efficient means of managing data while holding the line on cost. At the same time, the pace of business is increasing, and change requests are pouring into IT departments for more, better, faster data processing. To meet these demands, enterprises will need to do the following:

- Explore the evolving technologies that enable extreme transaction processing in a context that leverages all relevant enterprise data.
- Look at combinations of products from vendors, including database, data movement, and in-memory computing, that bring data together in real time to support extreme transaction processing at the point of action.
- Seek to bring the technology for extreme transaction processing together with that for more prosaic data management — in the same product family or architecture, if possible.
- Examine SAP Data Management offerings as examples of such a product combination.

IN THIS WHITE PAPER

This white paper considers the challenges involved in keeping pace with business as conditions change and competitors use data analytics and the agility gained from evolving data management technologies to gain business advantage. It looks at the pressures on the datacenter to evolve and the approaches available to do so. It then examines the data management platform approach offered by SAP to achieve better operational efficiency and agility with extreme transaction processing.

SITUATION OVERVIEW

The Real-Time Enterprise Challenge

Enterprises are struggling to keep up with the demands of business in a globally, electronically connected world. Not only do events that require response come at an ever-quickening pace, but business intelligence (BI) about customers and competitors may be gleaned from vast amounts of ever-changing data on the Internet, if only it can be rapidly acquired and processed. Business managers know that competitors are chasing this intelligence as fast as they are.

How can a modern enterprise keep up with the pace of business demands and also gain an edge from timely business intelligence? The key is in the ability to acquire and ingest all the necessary data, to act on that data with as little delay as possible, and to process the resulting transactions immediately.

This means dealing with streaming data from external sources and sensor and other machine-generated data from internal sources. It means adopting Big Data technologies to collect and ingest large amounts of business intelligence data quickly. And it means having a core, real-time data management platform that can enable applications to put all the pieces together and act in a timely manner. At the heart of such a platform needs to be memory-optimized database technology that operates cooperatively with extreme transaction processing.

What Is a Real-Time Enterprise?

A real-time enterprise is an enterprise that can act on events as they happen rather than wait for relevant information to be entered into and worked through the system, stored, compiled, and made available for query and reporting. A real-time enterprise needs to handle shifting external factors, such as customer demand, supplier pricing and product availability, and operational costs, as well as internal factors, such as logistics, inventory, and production rates.

This is not simply about producing an "executive dashboard." This is about marshalling IT resources based on the current situation on the ground as it changes. It is about automated decision generation and action and about supporting just-in-time tactical decisions. It affects transactional operational systems and analytic systems, including operational BI, strategic analytics, and predictive analytics.

Big Data at Rest and in Motion

Part of what drives the real-time enterprise is Big Data. Most people are familiar with Big Data at rest, which involves the collection of large amounts of data that may be either streaming, machine-generated data or unorganized collections of content; filtering, ordering, and formatting that data; and making that data available to drive decisions and actions. Hadoop is commonly used for this purpose.

Big Data in motion is different. This also involves large amounts of streaming data, but instead of ingesting it and then examining it, this approach involves recognizing events in the data and taking immediate action. Such technology generally involves the use of a complex event processing (CEP) engine that drives actions in response to defined complex and correlated events. Of course, recognizing such events often requires context, so such systems need to be able to reference facts or patterns of facts that have previously occurred, such as those collected by Big Data at rest. Thus the two types of Big Data are complementary.

Extreme Transaction Processing

While Big Data at rest is mainly used for strategic analysis and planning, Big Data in motion can drive business decisions. Data collected from such processes in a memory-optimized environment can condition application behavior and user actions to maximize business outcomes, but such functionality must be accompanied by a transactional system that can keep up with it. Formerly, it was acceptable, and even expected, that transactions such as bank deposits or store purchases would take a day or more to process. Today, we expect instant processing. This means that the number of transactions that OLTP databases are expected to handle within a very short period of time is exploding.

Extreme transaction processing addresses these needs. Extreme transaction processing is the capturing of very large amounts of concurrent transactions in a large and rapidly growing database at a rate of speed sufficient to keep up with the speed of business, yet with full data consistency and zero data loss. Instant failover and high-availability support are also generally part of the picture.

Businesses that are rapidly growing, or that must support fast-moving processes, often have extreme transaction processing requirements for various segments of the business or of business processes. These requirements include the ability to support rapid bulk processing and, simultaneously, interactive, intelligent processes, with rapid scaling in both the size of the database and the number of concurrent users.

The Right-Time Data Solution

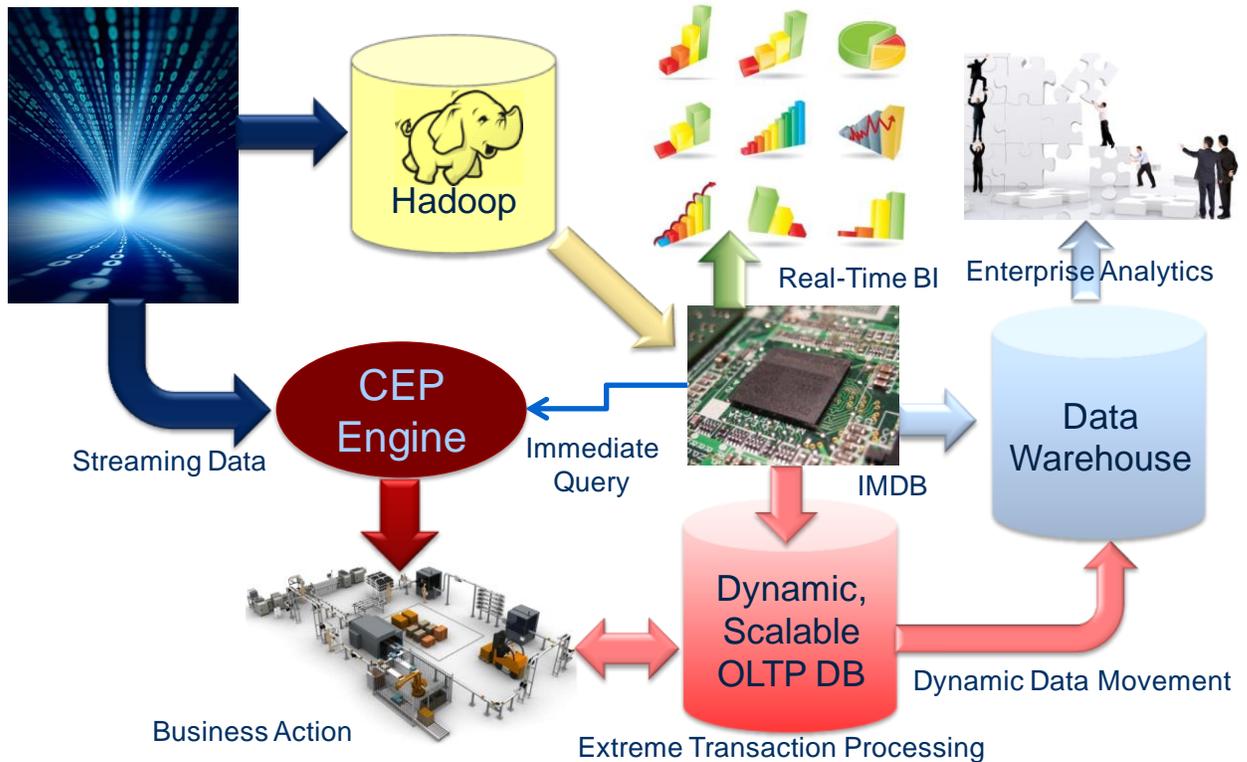
There was a time when enterprises conducted business at a pace that their computer systems, including applications and databases, could support. In this universally connected 24 x 7 world, however, business has its own pace, and enterprises that cannot keep up are left behind. It is no longer acceptable to run one's business at the pace of the application; the application must run at the pace of business.

Memory-based database technology is essential to enabling the application environment to keep up with the pace of business. A memory-optimized database, keeping the most vital data in main memory, is the only kind of database that can be used with a streaming data-driven system to process queries fast enough to ensure real-time performance. Complementing this must be transactional database technology that can process updates at that same pace, delivering extreme transaction processing. Taken together, the speed, simplicity, and nimble nature of

memory-optimized and extreme processing disk-optimized databases make them critical success factors for achieving the real-time enterprise, as illustrated in Figure 1.

FIGURE 1

Extreme Transaction Processing in the Real-Time Enterprise



Source: IDC, 2014

Memory-Optimized and Disk-Optimized Databases

Memory-Optimized Databases

A memory-optimized database uses data in memory to optimize both performance and agility. Rather than organizing the data for efficient layout on disk, the memory-optimized database organizes the data for efficient access in main memory. Often, the data is held in a columnar format, optimized for rapid resolution of complex queries. Sometimes, it is held in a row-oriented format, optimized for rapid insert and update. Some memory-optimized databases use combinations of these. The choice of structure is determined by the intended model for usage of the data.

At the core of a memory-optimized database is a data structure kept entirely in memory, sometimes called an in-memory database (IMDB). Since most of the data in a database is infrequently accessed, a memory-optimized database may offer the option of keeping only the frequently accessed data in the IMDB while swapping out the less frequently accessed data to disk. Even when it uses disk storage, however, the memory-optimized database holds the data on disk in a way that best supports its memory-optimized data organization, enabling it to satisfy the requirements of extreme transaction processing.

The memory-optimized DBMS can deliver very fast responses for complex queries and support certain classes of transaction processing as well. For the real-time enterprise that has a requirement to support continuously growing databases with extreme transaction processing, the memory-optimized DBMS complements the disk-optimized transactional database.

Disk-Optimized Databases

As should be clear by now, memory-optimized DBMS technology does not replace disk-optimized DBMS technology where extreme transaction processing applications involving very large and dynamic databases are concerned. This is because a disk-optimized DBMS delivers random read/write access to much larger data sizes than can be handled in memory in a practical way. The disk storage system provides a low-cost means of storing and, if managed properly, rapidly updating large databases. This is because, in such a DBMS, disk utilization can be optimized along multiple dimensions, including transaction and query performance as well as operational efficiency. The disk-optimized DBMS does this by mapping the data to disk locations and reducing I/O latency by various techniques such as sparse data dispersal, partitioning, collocation of related data on the same database page, and so on. Indexes are used to enable fast random access. Although this requires some time and expertise on the part of the database administrator (DBA), disk-optimized DBMSs offer the most flexibility in terms of managing rapidly changing large transactional databases. This is especially the case where complex write operations with embedded queries that span multiple database areas are involved.

Most disk-optimized DBMSs today have some in-memory assist technology, which enables rapid transaction processing. Because they are based on disk storage, they can write to a continuously growing transactional database with consistent high performance. Thus a disk-optimized DBMS with an in-memory data management capability can deliver the best of both worlds for the most demanding database applications.

Examples of the Real-Time Enterprise in Action

The following five examples of real-time enterprise applications illustrate the principles outlined previously. They are not specific use cases, but they incorporate elements of a number of use cases with which IDC is familiar.

- ☒ **Real-time retailing.** A retail firm tracks inventory in its stores by capturing RFID information from pallets as they arrive at the loading dock and sales at point-of-sale terminals (POSTs). The sales data also reveals volumes and patterns of sales moment by moment. Data about sales and inventory by store is loaded into

a memory-optimized database, and analytic software determines whether changes in sales volumes suggest a price change and whether trends compared with inventory suggest the need to restock, and if so, in which stores and from which warehouses. Simultaneously, changes in inventory regarding product availability as well as customer accounts are recorded in a disk-optimized database, immediately available for reference by sales reps and customers using the self-service Web site. Such operations minimize inventory-related costs and maximize competitiveness without ill-considered price changes up or down and also give sales and customers an exact, not approximate, view of product availability and customer account status.

- ☒ **Real-time trading.** A portfolio management company maintains portfolio holdings and rules in main memory, governed by a memory-optimized database. Such rules govern how frequently trades may occur, how much risk is tolerated, what kinds of issues are to be considered for inclusion in the portfolio, and so forth. The firm also receives streaming data about stock trades. It records data about issues of interest in the main memory database also and looks for interesting trends and correlations in share prices and patterns in share price changes. When they are found, it compares changes and the algorithmic buy-or-sell suggestion against each portfolio to determine whether, for that portfolio, a trade is warranted based on its rules. If so, the trade is executed. All this is done in milliseconds. When a disk-optimized DBMS is used for this purpose, the transactions can be verified and can be guaranteed as non-repudiable, even when distributed across multiple physical locations (assuming integrated high-speed replication), regardless of database size or complexity. This is a key feature from a compliance standpoint. This is especially significant when dealing with securities that are listed on multiple exchanges, since any trades must be instantly and permanently coordinated across all such exchanges.

- ☒ **Real-time logistics.** A trucking firm receives real-time data about the location and condition of each truck on the road as well as traffic conditions, which it maintains in a memory-optimized database, along with the contents of each truck on the road as well as its delivery route and schedule. Changes in traffic, new or canceled delivery orders, and exegetical events such as accidents or breakdowns as well as the fuel level of each truck affect orders that may be pushed out to the drivers, who change their route and schedule, all subject to moment-by-moment change. The result is more nimble pickup and delivery, faster response to problems, optimal truck routing, and fuel cost savings. As the day progresses, changes in the states of accounts, inventory, and service delivery are recorded from hundreds or even thousands of trucks in a disk-optimized database, available for immediate query by customer service reps and at self-service Web sites.

- ☒ **Real-time pay-as-you-go mobile telephony.** A mobile calling service, wishing to provide mobile calling capabilities in countries with very low personal income levels, offers rechargeable calling cards. The customer can go to any affiliated retailer (usually a small grocery store or similar shop) and pay cash; the clerk sends a message, and the card is charged with time corresponding to the amount paid (the store gets a commission for this interaction). The customer can then use his/her mobile phone immediately to make a call on that card,

regardless of which carrier supports the phone. Only by a very fast transaction processing database, and an efficient data transfer network, can the calling service log the payment and support the call immediately. The volumes in these cases can get very large, with tens of millions of customers and hundreds of thousands of card-charging requests every hour. The updates must be fast and accurate to ensure that customers can use their cards for just the amount of time they have purchased — no more, no less.

- ☒ **Real-time dynamic railroad ticketing.** A national railway in a vast, populous, and geographically diverse country, where the principal means of long distance travel is by rail, seeks to relieve station congestion and ensure better service by replacing its system of preprinted unreserved tickets with a system that issues tickets dynamically. Travelers can buy the tickets at a variety of electronic kiosks and vending machines. Ticket purchases and cancelations are processed in a coordinated manner across the country to ensure ticket validity and prevent fraud. To do this, the railroad must have rock-solid OLTP database technology, with the database servers linked by dynamic and highly reliable database replication technology. They now can process nearly 20 million ticket transactions per day across the country, with each ticket purchase requiring less than 20 seconds to complete. Formerly, passengers could only buy an unreserved ticket at the station within an hour of departure. Now, they can buy tickets for trains up to 3 days in advance. Only with an integrated extreme transaction processing data management system can this be accomplished.

Building the Right Mix of DBMS Technologies

Not all applications require support for extreme transactions. Not all applications change frequently enough to require extreme database agility. Most analytic data is accessed only occasionally and need not reside permanently in a memory-optimized database. For these reasons, disk-optimized databases remain sensible elements of the datacenter owing to the lower capital cost of keeping data on disk versus in memory.

For the agile enterprise, however, even those disk-optimized databases should be aligned with, and connected to, memory-optimized databases for areas where the data is relevant to an extreme transaction or a complex analytic problem. For this reason, it makes sense that the databases involved should have a basic level of compatibility and should be connected with high-performance replication technology.

SAP Data Management: SAP's Unified Solution for Real-Time Business

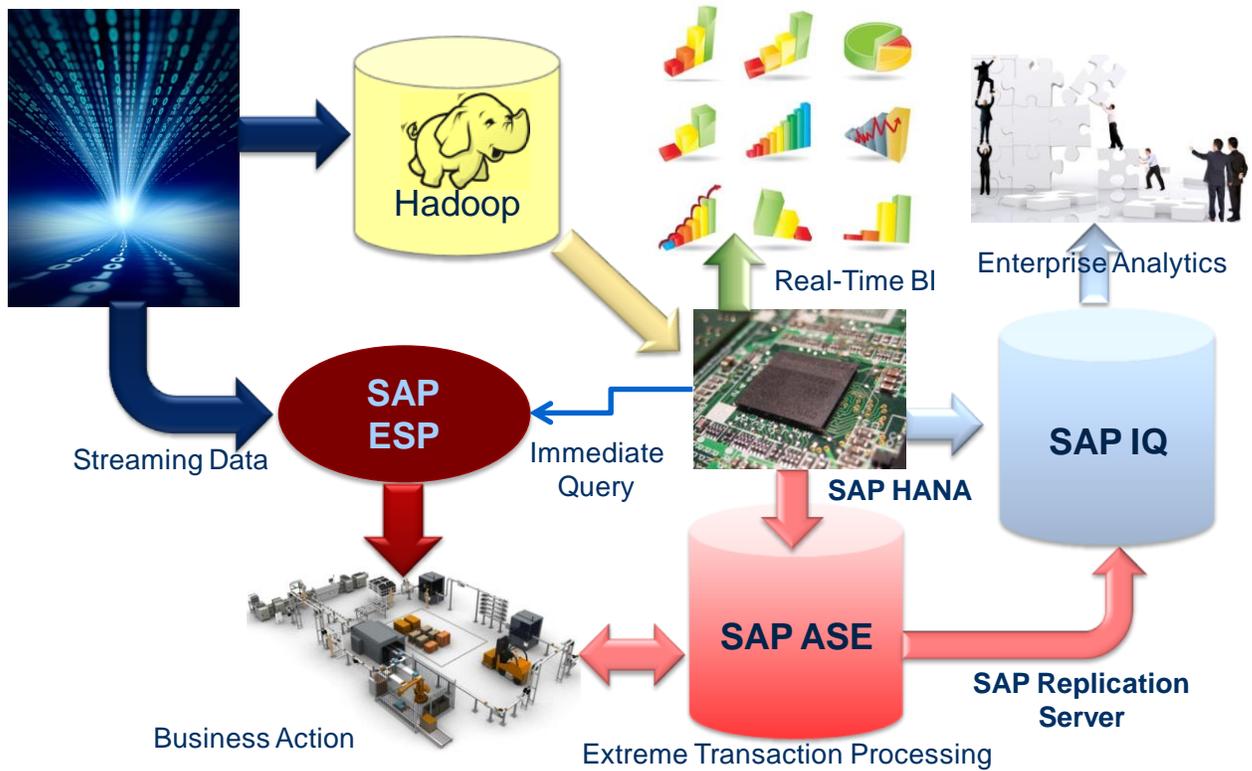
SAP offers a combination of products that meet all the requirements of this "right mix of DBMS technologies." This framework is anchored by SAP HANA, the in-memory computing platform, which provides application execution and business intelligence functionality that enables intelligent operation at the point of action, including an in-memory application platform dynamically connected to a range of fit-for-purpose, optimized database products. These are coordinated and managed by a framework called SAP Data Management, which enables data to be moved, synchronized, and made available dynamically.

Elements of SAP Data Management

Figure 2 shows the elements of the real-time enterprise extreme processing environment with the corresponding SAP products that make up SAP Data Management.

FIGURE 2

SAP Data Management Products



Source: IDC, 2014

SAP HANA

At the heart of SAP Data Management is SAP HANA, an in-memory processing platform that includes an IMDB optimized for efficient operation with the application or applications running on the platform. The HANA database is designed to act as the intelligent brain of the system, using active operational data, even data that has just been inserted, in relevant queries to support smart transactions.

SAP ASE

Working with SAP HANA to process transactions at the required speed is SAP ASE, a highly scalable disk-optimized RDBMS with support for extreme transaction processing. Its method of optimizing the storage of rapidly growing data on disk is key to delivering transactional scalability to SAP Data Management. SAP ASE provides a

number of key features delivering support for upwardly scaling database transaction processing. These include the highly nuanced use of solid state drives (SSDs) for persisting committed data very quickly and an internal architecture that supports several modes of multithreading to optimize for concurrent access.

SAP IQ

For large collections of analytic data, SAP offers SAP IQ. This RDBMS uses a columnar organization with in-memory features and various compression techniques to optimize for complex analytic queries across very large databases. SAP IQ also supports clustering for high scalability.

SAP Replication Server

Tying all the elements together is SAP Replication Server. More than a database replication tool, SAP Replication Server provides high-speed synchronization of multiple databases and can support heterogeneous sources and targets. With SAP Replication Server, relevant data can be delivered to the application database from any SAP database or many non-SAP databases at the right moment to enable real-time analytics at the point of action.

SAP Event Stream Processor

Many applications that involve real-time processing are driven by streaming data that must be handled using complex event processing technology for real-time correlations and actions. SAP Event Stream Processor (ESP) is a well-established technology leader in stream-based processing and also integrates well into the entire SAP Data Management configuration. SAP ESP has widespread adoption in very low-latency environments where instantaneous decision making is critical, like capital markets, telecommunications, and government intelligence.

Extreme Transaction Processing with SAP HANA and SAP ASE

SAP HANA delivers very high-speed responses for queries, and SAP ASE enables scalable transaction processing. Together, they provide an interesting recipe for high performance in transaction processing. How?

About 80–90% of most database transaction processing is query execution. Reference data is looked up to perform calculations that result in data that is inserted into transactional tables. The reference tables (e.g., customer, product) are relatively short. The transactional tables (e.g., sales order, inventory item, shipment) are typically very long.

Imagine that all the reference tables that are relevant to a given transaction are loaded into HANA's IMDB and RFID-enabled inventory information over a distributed warehouse environment in ASE's disk-based system database. The program queries the HANA reference tables for the necessary reference data and inventory information in ASE to perform the transaction and then, when orders are completed, inserts and updates rows in the transaction tables in SAP ASE. Consider the example of processing a customer order. The customer has picked out the items for purchase. The application queries HANA for the products and gets their product codes, looks up their prices and discount rates, checks the customer's credit line (also in HANA), and

then checks where the distributed and fast-moving inventory is available from ASE (via a proxy table registered in HANA). Finally, with all the reference and RFID inventory data assembled, the customer order is written and the inventory levels are updated to the SAP ASE database.

Of course, the processes described previously are happening with high levels of concurrency and with very low tolerance for data latency. SAP ASE is called for here because its ability to support very large numbers of concurrent sessions combined with high throughput makes it an appropriate companion to HANA. It can keep up with whatever HANA can dish out, and the result is extreme transaction processing.

Users can virtualize the data between SAP HANA and SAP ASE to increase the synergy. Also, relevant updated data in SAP ASE can be moved almost immediately into SAP HANA through SAP Replication Server. By intelligent use of the strengths of SAP HANA and SAP ASE, the application developer can achieve a level of query processing that a disk-optimized DBMS cannot deliver but also a level of scalable update transaction processing that is not practical for a memory-optimized DBMS, thereby realizing the best of both worlds.

FUTURE OUTLOOK

As memory continues to get cheaper, and chip makers explore new designs that will further enhance the value of the memory-optimized database approach, users will become accustomed to the extreme processing and agility delivered by these systems and will press for more. Similarly, optimization of database size and usage scaling will proceed apace. New architectures based on scalable hardware and virtualized resources will direct systems toward cloud architectures and create still more pressure for simplicity, flexibility, agility, and raw speed. The key to success in the DBMS world of the future will be in finding ways to combine and harmonize various DBMS technologies, taking advantage of the strengths of each, to fulfill the primary mission of the enterprise most effectively. SAP has taken a leadership position in this regard.

CHALLENGES/OPPORTUNITIES

The database world is flush with vendors scrambling to meet users' needs for better performance and faster throughput with lower operational costs. As a result, memory-optimized database technologies are being rolled out by a number of leading RDBMS vendors as well as new vendors that are just now emerging. SAP has a lead in extreme transaction processing with SAP HANA, but its real value is to be found in the synergies realized with SAP ASE and other SAP products for real-time data management. SAP will need to maintain that advantage going forward.

CONCLUSION

Businesses face growing pressure to increase their agility and respond more quickly to changing business conditions. Enterprises simply cannot rely on rigid, hard-to-change, slow databases to drive processes that interact with this rapidly accelerating world.

Combinational approaches to technology that offer the right tool for each job while interoperating to support extreme transaction processing with agility, at lower administrative cost, are the best way to address these changing conditions. Users would be well advised to consider the following:

- ☒ Determine how an agile extreme processing environment could change your ability to address your business needs.
- ☒ Look at ways of incorporating memory-optimized database technologies with scalable database transaction processing and building them into a framework that delivers intelligence at the point of action, with real-time data.
- ☒ Research the products and technologies that can deliver such benefits.
- ☒ Look at the whole of SAP Data Management products, giving special attention to SAP HANA and SAP ASE as a possible approach for architecting such extreme application solutions going forward.

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 2014 IDC. Reproduction without written permission is completely forbidden.