# Analyzing Business as It Happens

## SAP In-Memory Appliance Software (SAP HANA™) runs on the Intel® Xeon® processor to generate superior, real-time business intelligence.

WHITE PAPER

Intel® Xeon® Processor E7 Family
SAP In-Memory Appliance Software (SAP HANA™)

April 2011 | Version 1.0

### 1 Introduction

Increasingly sophisticated business decision models depend on extremely fast access to and manipulation of massive data stores. Insight into business operations often demands data volumes that are beyond the capabilities of traditional disk-based systems to process them in real time, which can limit access to benefits such as the following:

- **Efficiency** provided by the ability to respond in real time to the changing needs of the business
- **Flexibility** based on insight that accurately directs quick action
- **Empowerment** of business users to make and act on smart decisions

This reality often causes a separation between ongoing business needs and the analytic applications that support them. In such cases, the lag time between data gathering and interpretation delays the ability to benefit from business information, thereby limiting its value.

**SAP In-Memory Appliance Software (SAP HANA™) uses SAP In-Memory Computing technology and the performance and reliability of the Intel® Xeon® processor E7 family to dramatically improve the speed of data analysis, enabling business intelligence to be generated in real time. That capability increases the value of business data with more sophisticated decision support than has been possible with solutions based on traditional disk-based systems.**

SAP In-Memory Appliance Software (SAP HANA™) delivers next-generation in-memory computing through an ongoing engineering collaboration between SAP and Intel to provide optimized performance and reliability on Intel® architecture, including the Intel® Xeon® processor E7 family.

This paper introduces SAP In-Memory Computing technology concepts, explains how SAP HANA uses them to help deliver real-time business insight, and describes how synergies with Intel architecture-based server platforms drive solution value. The discussion begins with an examination of the hardware and software innovations that set the stage for SAP In-Memory Computing, including the ways in which the architectures of Intel Xeon processors and SAP HANA complement one another.

Next, the paper describes specific optimizations made as part of the ongoing, collaborative engineering relationship between Intel and SAP. In the context of general database concepts, the paper illustrates how SAP HANA takes advantage of Intel architecture to deliver high performance and reliability, as well as excellent support for application developers working with SAP HANA. Next, the paper describes real-world customer benefits derived from SAP HANA implementations. Finally, the paper explores how future technology may drive even further the value of joint solutions based on Intel Xeon processors and SAP HANA.

## 2 Intel and SAP Technology Advances Drive Database Evolution

The evolution of Intel architecture toward greater performance at lower cost sets the stage for the advances of SAP HANA. The SAP In-Memory Computing Database that lies at the heart of this technology is engineered to take advantage of each aspect of a high-performance, balanced compute platform:

- **Memory**. Increasingly large amounts of addressable memory, running at ever-higher speeds, combine with growing caches and other improvements to create memory subsystems robust enough to easily support the demands of SAP In-Memory Computing.
- **Processing.** Constant architectural improvements to complete more work per clock cycle extend the value of growing parallelism from multi-core designs, Intel® Hyper-Threading Technology (Intel® HT Technology), and increasingly scalable symmetric multi-processing.
- **I/O**. Very fast on-board data transfer using Intel® QuickPath Interconnect is complemented by advances in inter-node communication, to enable the low latency in data transfer within and between nodes in the SAP HANA appliance.

### 2.1 Intel Hardware Technology Innovations

Intel has continually led the industry in driving advances in commercial off-the-shelf server platforms that increase performance, scalability, and value. These innovations directly enable SAP HANA to transcend barriers to what businesses are able to do with their data, generating better business decisions in response to real-time events.

This section traces the history of Intel processor innovation, setting the stage for an understanding of how SAP HANA appliances based on Intel architecture produce synergies between the hardware and software that benefit business customers.

*"Intel and SAP engineering teams have collaborated extensively on SAP HANA™ and SAP In-Memory Computing technology. Comparing the in-memory system performance with a classical relational database, the query time was reduced from 77 minutes to 13 seconds when run on the Intel® Xeon® processor 7500 series. We are thrilled with this result, and it showcases once again the powerful impact of joint problem solving between our two companies."*

- K. Skaugen, Intel Vice President and General Manager of the Data Center Group

### 2.1.1 The Rise of Mainstream Server Parallelism

At one time, boosts in processor performance were created largely by ongoing increases in clock speed. The simple fact of providing more cycles per second retired more instructions per unit of time, allowing software to become more capable and deliver faster response times. As clock speeds grew, however, processor designs faced limitations due to growing energy consumption and heat generation.

One key solution has been to produce multi-core processors, so that each processor socket contains multiple, individual execution engines—a design concept taken from mainframe and other large-scale systems, where it had been employed for some time. Even before introducing its multi-core designs, Intel developed Intel HT Technology, which allows a single processing core to handle two strings of instructions (software threads) at a time, making better use of on-die resources that would otherwise be idle.

Together with symmetric multi-processing (multiple processors and multiple servers working together), multi-core processing and Intel HT Technology have made parallel programming the de facto industry standard. Some of the techniques used by SAP HANA to take advantage of parallel hardware resources are discussed later in this paper.

### 2.1.2 The Drive to 64-Bit Processing

Another limitation that enterprise computing systems have faced is that the 32-bit systems in general enterprise use until the mid-2000s were limited by their ability to address no more than approximately four gigabytes (GB) of system memory per server. Beginning in 2004 with the Intel Xeon processor formerly code-named Nocona, Intel began shipping hardware capable of running 64-bit OSs, which can address far larger amounts of memory. The theoretical limits of memory capacity are as follows:

**32-bit systems**: $2^{32}$ = 4,294,967,296 bytes, or **~4 GB**

**64-bit systems**: $2^{64}$ = 18,446,744,073,709,551,616 bytes, or **~18 billion GB (18 Exabytes)**

In reality, the upper limits of memory addressability for 64-bit systems are governed by the physical capabilities of the platform and to some extent by the cost of the memory itself. As platforms become more capable, system memory capacity continues to grow, and at the same time, the cost of memory on a per-gigabyte basis continues to drop. Server memory capacities are now often measured in terabytes (TB), where 1 TB = 1,000 GB.

In addition to being able to address more memory than their 32-bit ancestors, platforms based on 64-bit architectures make more and wider registers available to the compiler (theoretically enabling the compiler to generate faster code). They also have the benefit of some larger data types, a crucial factor for SAP HANA, where gigabytes of memory are handled within a single process. In fact, the port to 64-bit architecture was one of the first joint efforts by Intel and SAP related to SAP NetWeaver® Business Warehouse Accelerator (BWA, formerly known as SAP NetWeaver Business Intelligence Accelerator), a precursor in many ways to SAP HANA.

#### Intel and SAP: A History of Co-Innovation

For more than 10 years, Intel and SAP have worked together to deliver industry-leading performance of SAP solutions on Intel® architecture, and a large proportion of new SAP implementations are now deployed on Intel® platforms. The latest success from that tradition of co-innovation is available to customers of all sizes in SAP In-Memory Appliance Software (SAP HANA™), which is delivered on the Intel® Xeon® processor.

The relationship between Intel and SAP has become even stronger over the years, growing to include a broad set of collaborations and initiatives. Some of the most visible include the following:

- **Joint roadmap enablement**. Early in the design process, Intel and SAP decision makers identify complementary features and capabilities in their upcoming products, and those insights help to direct the development cycle for maximum value.
- **Collaborative product optimization**. Intel engineers located on-site at SAP work with their SAP counterparts to provide tuning expertise that enables SAP HANA and other software solutions to take advantage of the latest hardware features.
- **Combined research efforts**. Together, researchers from Intel and SAP continually explore and drive the future of business computing. Belfast Collaboratory, an excellent example, opened in October 2009 to innovate around cloud computing and the next generation of sustainable IT.

As a result of these efforts, customer solutions achieve performance, scalability, reliability, and energy efficiency that translate into favorable ROI and TCO, for increased business value.

### 2.1.3 Other Architecture Features that Benefit SAP HANA

Successive generations of Intel® server platforms introduce an enormous range of other architectural features that drive new platform capabilities. For example, the Intel Xeon processor E7 family adds the following new features and capabilities relative to its predecessor that enhance the scalable performance of SAP HANA implementations:

- **Increased core count**, from 8 cores (16 threads) to 10 cores (20 threads) per socket, provides more computation resources for SAP HANA's data manipulation demands.
- **Larger memory capacity**, from 16 GB DIMMs (1 TB for a four-way server) to 32 GB DIMMs (2 TB for a four-way server), increases the ability of SAP HANA to hold large amounts of data for in-memory computation.
- **Larger cache**, from 24 megabytes (MB) to 30 MB, increases the amount of data SAP HANA can hold for fast access without reading from DRAM.

Likewise, evolving server platforms provide new features for security, and reliability, availability, and serviceability (RAS) that support business-critical SAP HANA deployments. Again using features and capabilities introduced with the Intel Xeon processor E7 family as an example, SAP HANA benefits from Double Device Data Correction (DDDC), which allows recovery from two simultaneous memory errors, supporting greater availability of SAP HANA systems.

Platform innovation also includes the introduction of new instructions such as the series of Intel® Streaming SIMD Extensions instruction sets (Intel® SSE, Intel® SSE2, Intel® SSE3, Intel® SSSE3, Intel® SSE4.1, and Intel® SSE4.2). These instructions are designed to accelerate specific types of operations, and their impact on SAP HANA is discussed in further detail later in this paper.

Finally, Intel continues to innovate around process-technology shrink in its manufacturing operations. Moving to a smaller process technology allows the creation of smaller transistors and other on-chip features, which allows more of those features to be placed in a given amount of on-die area, for faster switching, lower per-transistor energy consumption, and lower per-transistor waste-heat production. Figure 1 shows some key advances in the development of the Intel Xeon processor E7 family.

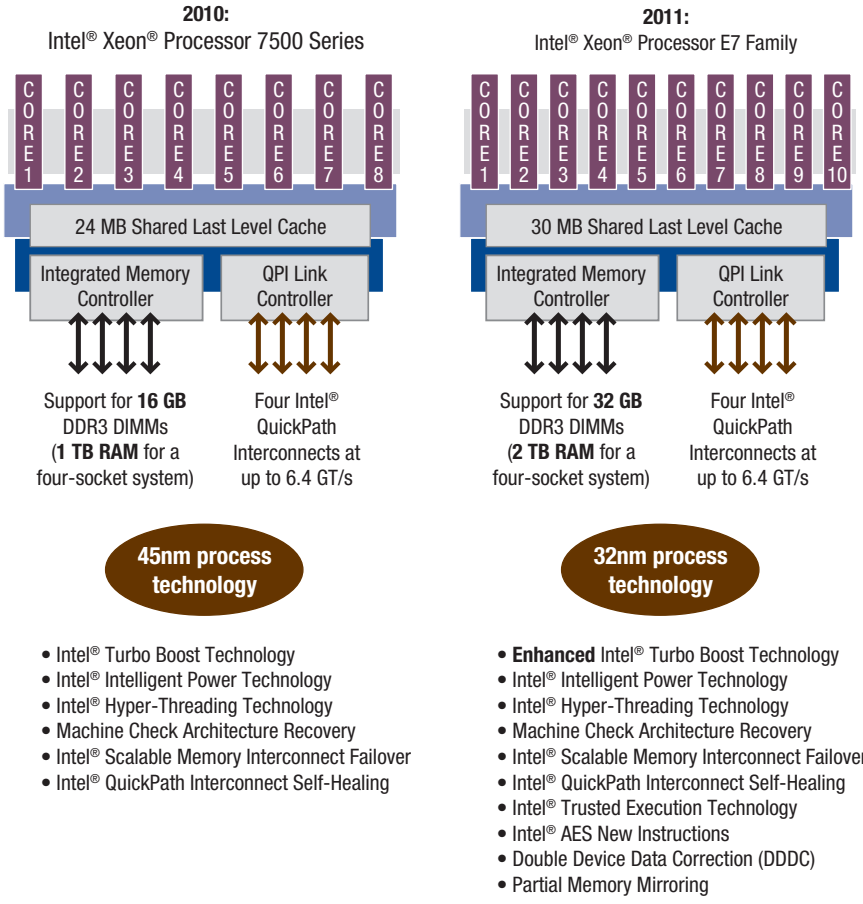


**Figure 1.** The Intel® Xeon® processor E7 family adds features and capabilities that directly benefit SAP HANA deployments.

## 2.2 SAP Software Technology Innovations

Around 2000, SAP recognized the trends toward rapidly-growing data sets owned by many companies and increasing business needs for real-time analytics. In response to those realities, SAP engineers began to re-imagine how data warehouses could handle these massive data stores to generate immediate, ad hoc analysis. Disk-access times had become a primary bottleneck in addressing those needs using conventional database queries. Research showed promise for removing that factor by conducting queries solely in memory, without having to access the physical disk.

Two key factors helped make in-memory computing viable for mainstream implementations. As discussed above, by the middle of the first decade of the 21st century, system memory had begun to fall sharply in cost, and 64-bit servers had opened the door to commercial off-the-shelf hardware that could address very large amounts of system memory. Those advances set the stage for SAP's implementation of in-memory computing for general business use.

*"For over 10 years, SAP and Intel have collaborated to create visionary solutions that enable customers to gain competitive advantage and run better. The SAP In-Memory Appliance (SAP HANA™), delivered on the Intel® Xeon® processor, provides customers with real-time results from analyses and transactions that enable better decisions and improved operations. It's just the latest example of how game-changing innovation between industry leaders can alter how enterprises do business and succeed."*

- Stefan Sigg, Senior Vice President, In-Memory Platform, SAP

### 2.2.1 Establishing a Solid Database for In-Memory Computing

In-memory computing allows the processing of massive quantities of real-time data in the main memory of the server, providing immediate results from analyses and transactions. The SAP In-Memory Computing Database (Figure 2) delivers the following capabilities:

- **In-memory computing** functionality with native support for row and columnar data stores (discussed below), providing full ACID (atomicity, consistency, isolation, durability) transactional capabilities
- **A calculation and planning engine** with a data repository, providing persistent views of business information and a unified information modeling design environment
- **A data management service** that includes SQL and MDX interfaces, data integration capabilities for accessing SAP and non-SAP data sources, and integrated life cycle management.

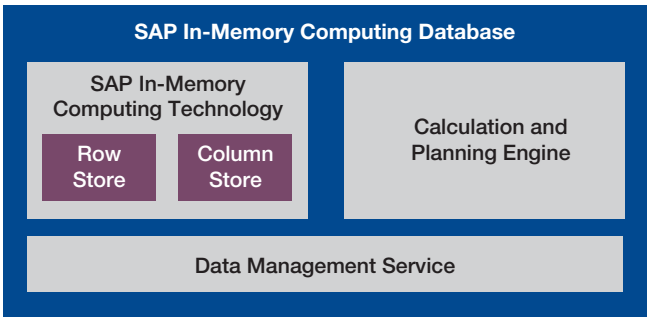


**Figure 2.** SAP In-Memory Computing supports complex calculations on massive amounts of data.

When combined, these capabilities allow the SAP In-Memory Computing Database to help decision makers explore and analyze vast amounts of information at incredible response times and with a high degree of flexibility, without the need for IT involvement.

### 2.2.2 Delivering the SAP In-Memory Computing Database into the Mainstream

SAP first made SAP In-Memory Computing available in product form with the introduction of TREX (Text Retrieval and information EXtraction), a search engine component used as an integral part of various SAP software offerings. Building on the early success of products such as SAP NetWeaver Enterprise Search, the company introduced SAP BWA in 2005.

The company's ongoing focus on business intelligence (BI) solutions led it to integrate SAP In-Memory Computing into SAP BusinessObjects to create an accelerated BI solution called the SAP BusinessObjects Explorer, Accelerated Version. SAP In-Memory Computing has also been used to enable real-time analysis in a variety of other products, including the following:

- SAP Advanced Planner and Optimizer with liveCache
- SAP CRM Customer Segmentation
- SAP Business ByDesign™ Analytics

### 2.2.3 SAP HANA: The Next Generation of Real-Time Business Analysis

SAP In-Memory Computing has evolved to create SAP HANA, a multi-purpose appliance that combines SAP software components installed on servers based on the Intel Xeon processor. The integrated software stack is highly optimized for performance and reliability on the hardware, and it includes the SAP In-Memory Computing Database, real-time replication service, data modeling, and data services, as illustrated in Figure 3.

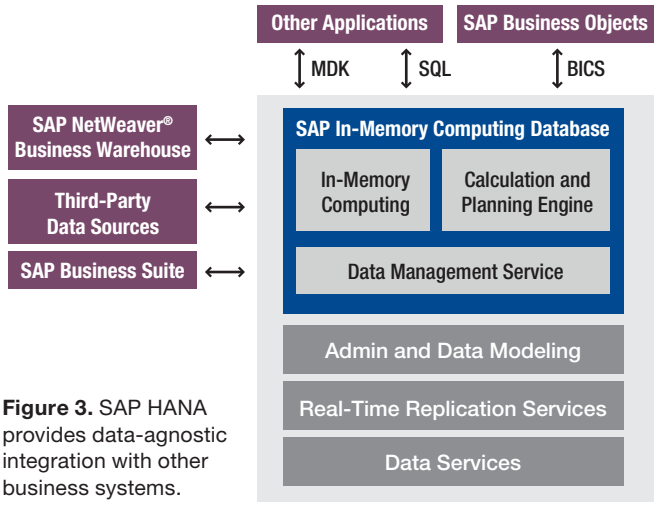


**Figure 3.** SAP HANA provides data-agnostic integration with other business systems.

SAP HANA benefits dramatically from the design features and benefits of the latest Intel Xeon processors, helping it deliver excellent value to end-customers. The performance, scalability, reliability, and energy efficiency of the hardware platform are particularly beneficial in this regard. High memory capacity and innovations in the memory and cache subsystems, as described above, support excellent results from the demands of in-memory computing. The sophisticated parallel software design, as described below, takes advantage of the platform's large-scale hardware parallelism.

SAP HANA is engineered specifically to be data-agnostic, and the appliance is designed to be implemented without affecting existing applications or systems. The solution provides a flexible, cost-effective, real-time approach for managing large data volumes, allowing organizations to dramatically reduce the hardware and maintenance costs associated with running multiple data warehouses and analytical systems.

In the future, SAP HANA will act as the technology foundation for new, innovative business applications for planning, forecasting, operational performance, and simulation. The first business application running on SAP HANA is the SAP BusinessObjects Strategic Workforce Planning, version for in-memory computing, a solution SAP developed that dramatically benefits from high-performance and real-time access to large volumes of transactional information.

## 3 Creating a Transition in Database Design

With the growth of data stores and business application complexity, solution providers have steadily developed novel approaches to optimizing performance. One key development has been the separation of online transactional processing (OLTP) and online analytical processing (OLAP) systems because both these types of relational databases have somewhat divergent design goals. In the context of BI solutions, they can be considered as follows:

- **OLTP systems** instantly record business events as they happen, such as the sale of a piece of inventory. As such, system designs focus on quickly handling large numbers of small, simultaneous transactions.
- **OLAP systems** provide analysis of the data provided by OLTP systems to support business decisions. They are designed to handle a relatively small number of often-complex transactions.

In his paper, "A Common Database Approach for OLTP and OLAP Using an In-Memory Column Database,"[1] SAP co-founder Hasso Plattner asserts the value of unifying both OLTP and OLAP in a single system based on column storage and in-memory computation. While the concept of combining the two systems is beyond the scope of this paper, Plattner's argument for the value of column storage and in-memory computing to both OLTP and OLAP has direct bearing on the capabilities of SAP HANA in those areas.

### The Intel and SAP HANA™ Value Proposition Across a Wide Variety of Business Units

SAP In-Memory Appliance Software (SAP HANA™) and the Intel® Xeon® processor enable a broad set of scenarios within various parts of the business, above and beyond transactional excellence and traditional analytics, such as the following:

- **Order management.** Available to promise, price optimization
- **Supply chain.** Transportation planning, inventory management, demand and supply planning
- **Manufacturing.** Production planning, performance management, asset utilization
- **Sales and marketing.** Analytics, customer segmentation, trade promotion management
- **Corporate finance.** Planning and budgeting, consolidation
- **Human resources.** Workforce analytics
- **Information technology.** Landscape optimization

### 3.1 Performance Implications of Column versus Row Storage

Whereas relational databases typically use row-based storage, column-based storage may be more suitable in some cases. As shown in Figure 4, a database table is conceptually a two-dimensional structure made up of cells arranged in rows and columns. Because computer memory is structured linearly, there are two options for the sequences of cell values that are stored in contiguous memory locations:

- **In row storage**, the data sequence consists of the data fields in one table row.
- **In column storage**, the data sequence consists of the entries in one table column.

Figure 4 illustrates both options.

**Table**

| Country | Product | Sales |
|---|---|---|
| US | Alpha | 3.000 |
| US | Beta | 1.250 |
| JP | Alpha | 700 |
| UK | Alpha | 450 |

**Row Store**

| | |
|---|---|
| Row 1 | US |
| | Alpha |
| | 3,000 |
| Row 2 | US |
| | Beta |
| | 1.250 |
| Row 3 | JP |
| | Alpha |
| | 700 |
| Row 4 | UK |
| | Alpha |
| | 450 |

**Column Store**

| | |
|---|---|
| Country | US |
| | US |
| | JP |
| | UK |
| Product | Alpha |
| | Beta |
| | Alpha |
| | Alpha |
| Sales | 3.000 |
| | 1.250 |
| | 700 |
| | 450 |

**Figure 4.** Unlike traditional databases, SAP HANA offers the choice between row-based and column-based data storage for each table.

Columnar storage may be more efficient in the common case where a given column includes only a relatively small number of distinct values. For example, if the table shown in Figure 4 includes a very large number of rows, one would expect that each entry in the Country column would nevertheless consist of one of perhaps a few dozen possible entries. That factor allows for very efficient data compression, which produces a number of advantages.

For example, data compression allows for faster loading from system memory into the processor cache. Since that transport tends to be the main factor limiting system performance, speeding it up may be expected to produce a performance advantage, even considering the added overhead of decompressing the data. Another example of the advantages of columnar storage is related to compression: Computing the sum of values in a column is much faster if many additions of the same value can be replaced by a single multiplication. Similarly, holding values from a column together in cache is crucial for bulk processing using Intel SSE instructions.

Columnar storage also helps reduce the number of performance-limiting cache misses. Recognizing that data is transferred from main memory to the processor in blocks of fixed size called *cache lines*, consider the case of aggregating the sum of all of the values in the Sales column of a large version of the table in Figure 4. With row storage, only a minority of the data in any given cache line would be from the Sales column and therefore be relevant to the calculation. The large proportion of irrelevant data would result in cache misses, where the processor would have to discard the irrelevant data and wait for useful data in its place. Conversely, with columnar storage, every piece of data in each cache line would be relevant to the calculation, increasing efficiency by keeping the processor supplied with useful data.

### 3.2 Dictionary Compression with Column Storage

Dictionary coding is a significant benefit when compressing column-stored data, as represented in Figure 5. Each distinct column value is assigned a sequential integer (beginning with zero) as a Value ID. An alphanumerically sorted dictionary stores all of the distinct column values, and the Value IDs are implicitly given by the sequential positions of each distinct column value. The actual column data is stored as a sequence of Value IDs, which are bit coded, meaning that they are stored using the minimum possible number of bits.

### 3.3 Efficiencies from Separation of Main and Delta Storage

As changes are committed into the database, they are stored in *delta storage*, which is distinct from *main storage*, and if new data contains column values that are not included in the main dictionary, they are stored in a *delta dictionary*. Because the delta dictionary is created without recoding the Value IDs in the main dictionary, there is no global order that encompasses the Value IDs in both the delta dictionary and the main dictionary.

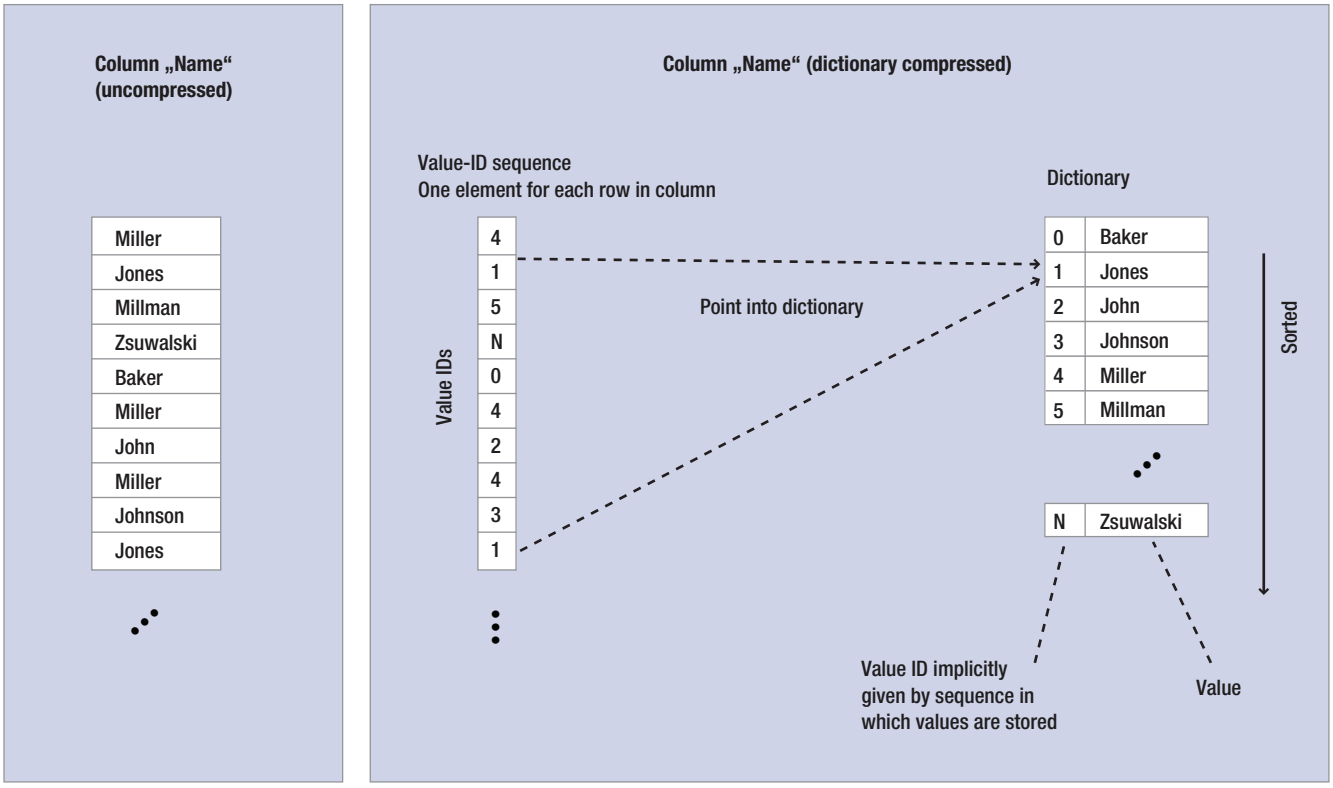


**Figure 5.** Dictionary coding can enable significant compression advantages in column-stored data sets.

Dictionary compression of columnar data is most efficient when the number of distinct column values is relatively small, compared to the size of the column (that is, when a relatively low number of Value IDs exists relative to the number of column values).

Because the dictionary is sorted, search operations and comparisons such as “greater than” and “less than” can be performed against the value IDs to return results that are meaningful with regard to the actual column values. Because those operations are conducted on integer values, they can be much faster than, for example, performing them against long strings.

The delta store supports “greater than” and “less than” queries by adding a CSB+-Tree to the delta dictionary. This tree holds the values of the delta store ordered as in the main store, only the amount of memory needed is higher than in the main store case. To establish a single, sorted dictionary that includes all Value IDs when a delta dictionary has been established, the delta dictionary and main dictionary must be merged using a *delta merge* operation. To enable even higher compression of main storage, SAP HANA selects and implements the optimal compression method for the dictionary itself during delta merge.

### 3.4 Storage of and Queries against Historical Data

As an optional feature, SAP HANA can use temporal tables to query against an historical state of the data. Write operations on temporal tables do not physically overwrite existing records. Instead, write operations always insert new versions of the data record into the database. The different versions of a data record have time stamps indicating their validity, with current data being simply the most recent version.

Historical (non-current) data stored in a temporal table is still considered to be operational data, meaning that it is expected to be used in current operations, such as transactions or reporting. By contrast, future versions of SAP HANA plan to support a category specifically for data that is not expected to be used anymore, known as *aged data*, which no longer need to be stored in memory. Instead, aged data will be eligible to be written to disk and released from memory.

Note that while aged data would not be accessed during normal operations, it would still be available if needed. To reduce access time, systems could use solid-state drives for this purpose. Alternatively, aged data could be written to flash memory, which can be accessed even more quickly. For each individual type of data, application developers will need to create rules that govern when data should be considered "aged." For example, sales data might be considered aged once the order it is associated with has been closed for at least 90 days.

### 3.5 Boosting Performance by Moving Results Instead of Source Data

The data-intensive business applications in common use often move large, detailed data sets from the data layer to the application layer, where it performs operations on that data. This arrangement is not optimal in many cases, since the result of the calculation is typically smaller than the source data needed to derive that result. Therefore, it is generally more efficient to perform calculations before moving the data and then move only the results.

To benefit from that scheme, SAP HANA incorporates a calculation engine into the data layer. The calculation engine can be manipulated programmatically by application developers. In this way, high-performance applications based on the SAP HANA architecture can delegate data-intense operations to the data layer, achieving higher efficiency and performance than would otherwise be possible.

## 4 Collaboration by SAP and Intel for Performance Optimization and Parallelization

Co-innovation and collaboration between SAP and Intel is an important prerequisite for the success of the SAP HANA software components on Intel architecture-based hardware. The optimization draws both from industry best practices and from solution-specific engineering, drawing broadly on the expertise of performance engineers from both companies. As mentioned previously in this paper, these individuals have a long history of working together, and their work on SAP HANA represents the culmination of many years' experience.

This four-part section identifies some of the specific optimizations the team has made to the SAP HANA platform and the resulting performance benefits on Intel Xeon processor-based systems. The first three subsections describe the optimizations, organized for convenience at decreasing scale within the system:

- **Among sockets.** The Benefits from Multi-Socket Topology and NUMA subsection describes the collaboration around features and capabilities at the inter-processor level.
- **Among cores.** The Optimization for Multi-Core Processing subsection describes the collaboration around features and capabilities at the intra-processor, inter-core level.
- **Within cores.** The Optimization for Core-Level Features subsection describes the collaboration around features and capabilities at the intra-core level.

The fourth subsection, Performance Benefits from Optimization and Hardware Platform Features, presents some key performance results associated with the optimizations described in the previous subsections.

### 4.1 Benefits from Multi-Socket Topology and NUMA

Because SAP HANA is a large-scale multi-socket environment, it benefits substantially from the NUMA (non-uniform memory access) shared-memory architecture.

In multi-processor systems that are not NUMA capable (or where NUMA is not enabled), each processor is connected directly to all available system memory through the system bus. The access path between any combination of processor and memory is direct and electrically equivalent to any other combination. This arrangement is therefore known as "uniform memory access" (UMA), as shown in Figure 6a. Note that all the processors must contend for bandwidth on the system bus.

In NUMA systems (Figure 6b), each processor is connected directly to its own local memory, and all of the memory banks are connected together by the system bus. Therefore, any processor has direct, relatively high-speed access to local memory, as well as access to all other memory through the system bus. Because bus-bandwidth contention only occurs on non-local memory accesses, there is a performance advantage to holding the data needed by a given processor in its own local memory; that placement is known as "processor affinity."

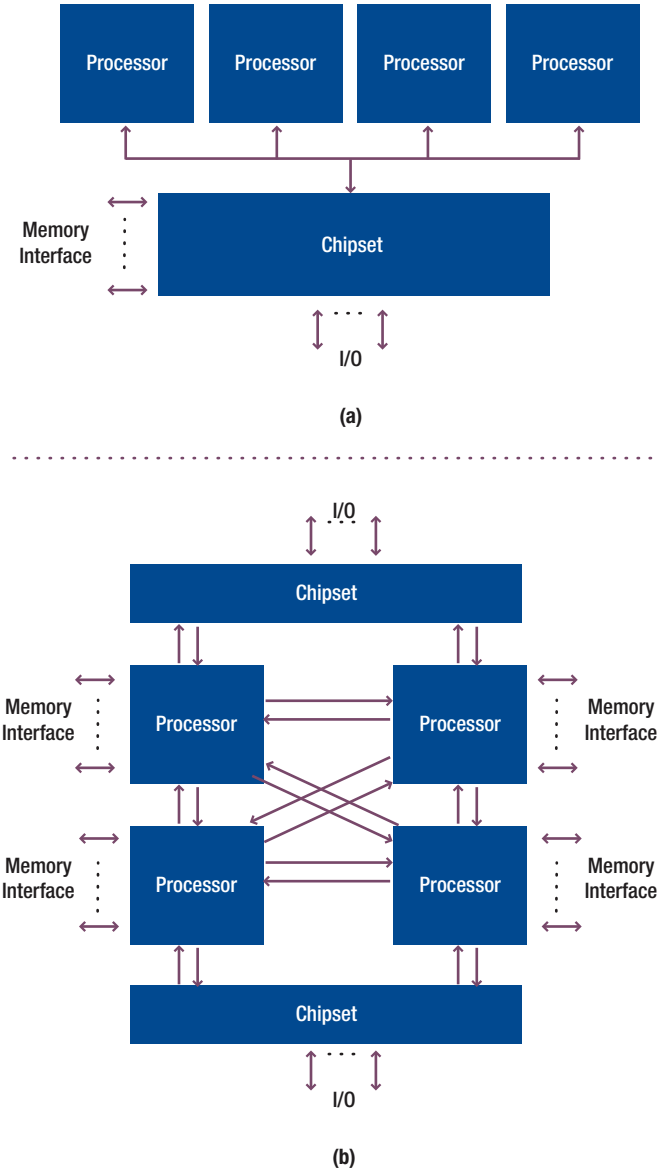


**Figure 6.** In NUMA systems, each processor has fast, direct access to its own local memory module, reducing the latency that arises due to bus-bandwidth contention.

In multi-threaded application software, the OS migrates threads between processor cores as needed to match workloads to available execution resources. Assuming that the data associated with a specific thread is already resident in memory, migrating a thread to a core in a different processor package carries with it a performance penalty.

Therefore, NUMA-capable software must be capable of working with the OS to manage the associated balance in performance implications. For example, there may be an advantage from moving a thread to a processor package where more execution resources are available that must be weighed against the disadvantage of moving data across the system bus.

### 4.2 Optimization for Multi-Core Processing

To take advantage of the increasing core count in modern processors (as many as 10 cores supporting 20 threads in the Intel Xeon processor E7 family), SAPA HANA is engineered for parallel execution that scales well with the number of available cores. Specifically, optimization for multi-core platforms accounts for the following two key considerations:

- **Data is partitioned** wherever possible in sections that allow calculations to be executed in parallel.
- **Sequential processing is avoided**, which includes finding alternatives to approaches such as thread locking.

#### 4.2.1 Parallel Aggregation

In the shared-memory architecture within a node, SAP HANA performs aggregation operations by spawning a number of threads that act in parallel, each of which has equal access to the data resident on the memory on that node. As illustrated at the top of Figure 7, each aggregation functions in a loop-wise fashion as follows:

1. Fetch a small partition of the input relation.
2. Aggregate that partition.
3. Return to step 1 until the complete input relation is processed.

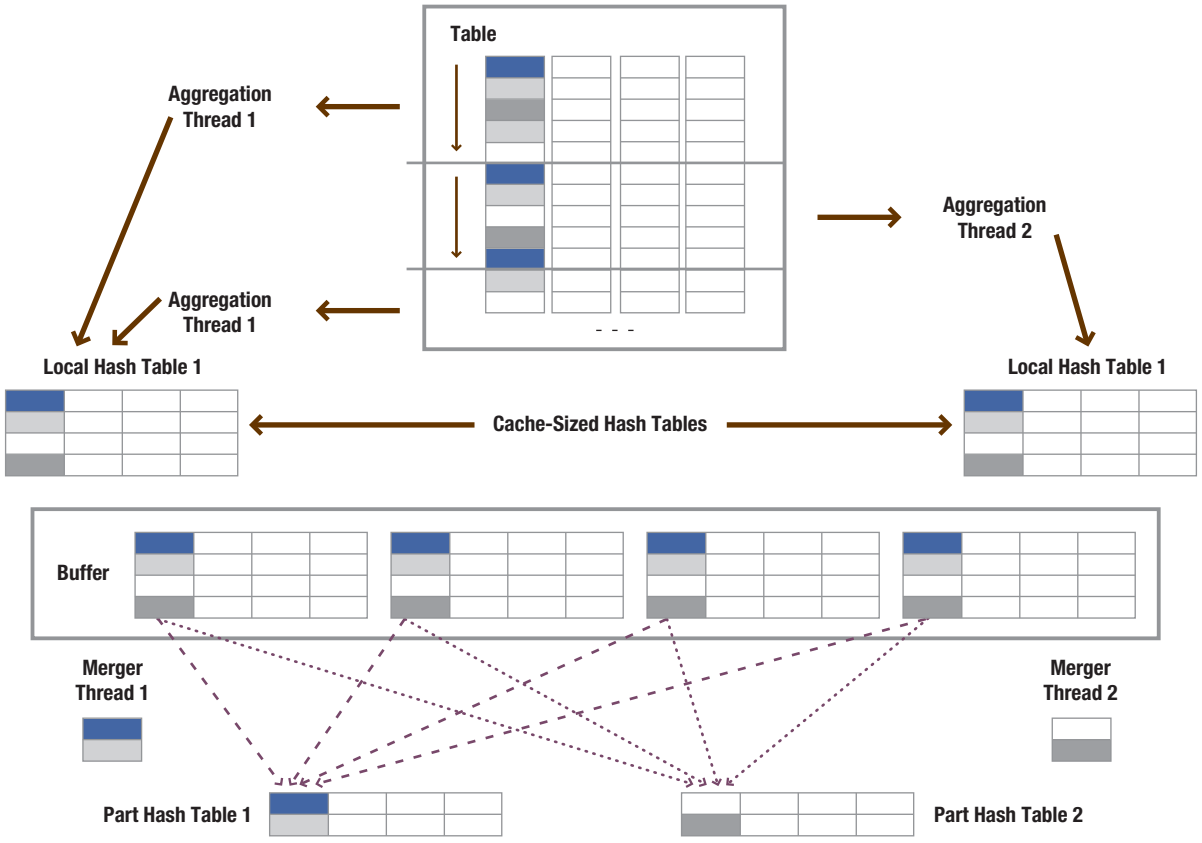


**Figure 7.** SAP HANA employs parallel aggregation to divide among threads the work of fetching, aggregating, and merging data.

Note that the approach of using small chunks and dynamically assigning them to threads distributes the computation costs over all threads relatively evenly, even if requirements differ substantially among the chunks.

Each thread has a private hash table where it writes its aggregation results. When the aggregation threads are finished, the buffered hash tables are merged by merger threads using range partitioning. Each merger thread is responsible for a certain range, which is indicated in Figure 7 by a different shade of grey. The merger threads aggregate their partition into *part hash tables*, and the final result is obtained by concatenating the part hash tables.

### 4.2.2 Processor Cache Optimizations

Taking best advantage of processor cache resources is another aspect of performance optimization that plays an important role in the development of SAP HANA. Because accessing data from cache is so much faster than accessing data in main memory, special care must be taken to improve the odds that the processor can access the data it needs at any given time from cache, rather than being forced to access memory.

As described above in the Performance Implications of Column versus Row Storage section, determining what data will be fetched together in a given cache line is an important aspect of optimizing performance within the memory/cache subsystem. Block-wise decompression of data allows decompressed data to reside in L1 cache for access by the processor.

The columnar design results in consecutive access of large amounts of data, which is an excellent fit for the hardware prefetchers built into the Intel Xeon processors. The SAP HANA development team added software prefetching hints to improve performance in the limited number of circumstances where data was not automatically prefetched (for example, when crossing page boundaries). Similarly, in Figure 7, note that the hash tables to which worker threads write aggregation results are sized to the cache, helping to reduce the incidence of cache misses in handling that data.

### 4.3 Optimization for Core-Level Features

In addition to the advantages available to SAP HANA at the levels of interactions among processors and interactions among cores, a number of features are also relevant at the level of operation within individual cores. Some of those advantages and considerations are discussed in this section.

#### 4.3.1 Optimization for Intel® HT Technology

Intel HT Technology helps improves performance through increased instruction level parallelism by enabling each processor core to accommodate two threads with independent instruction streams. While the two threads necessarily share execution resources within the core, the presence of two threads can increase utilization of the available execution units. This effect typically increases the number of instructions executed in a given amount of time within a core, as shown in Figure 8.

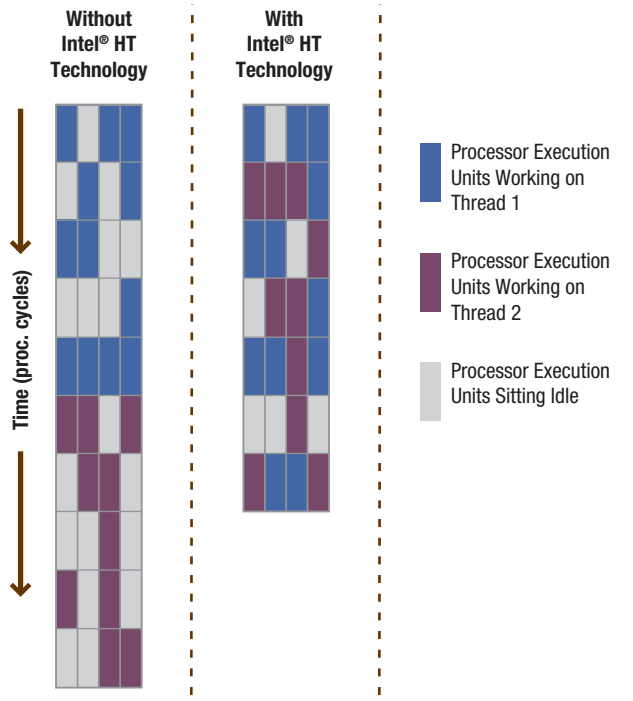


**Figure 8.** Intel® Hyper-Threading Technology (Intel® HT Technology) gives the processor access to two threads in the same time slice, reducing the level of idle hardware resources.

This greater efficiency enables higher throughput (since more work gets completed per clock cycle) and higher performance per watt (since fewer idle execution units consume power without contributing to performance). In addition, when one thread has a cache miss, branch mispredict, or any other pipeline stall, the other thread continues processing instructions at nearly the same rate as a single thread running on the core.

In large part, achieving the optimal performance benefit from Intel HT Technology is simply an extension of the optimization that must be done to take advantage of hardware parallelism generally. An in-depth discussion of issues specific to Intel HT Technology is beyond the scope of this paper. For more information, see the paper, "Performance Insights to Intel® Hyper-Threading Technology."[2]

To help ensure optimal benefit from Intel HT Technology, the development team addressed the following considerations:

- **Parallel scaling**. Intel HT Technology effectively doubles the number of software threads required to take full advantage of the execution hardware, making SAP HANA's ability to identify and spawn the correct number of threads even more important.
- **Thread balance**. Obtaining the optimal performance benefit from Intel HT Technology requires the work to be divided as evenly as possible among software threads. One approach taken with SAP HANA is to assign work in very small chunks, as described above in the discussion of parallel aggregation.
- **Avoiding threading bottlenecks**. The increased hardware parallelism associated with Intel HT Technology carries with it an increased danger from threading bottlenecks, so the team refined SAP HANA's threading model to avoid issues such as false sharing and excessive locks and synchronization.

#### 4.3.2 Standards Compliance and Accuracy for Decimal Floating-Point Math

The IEEE Standard 754-2008 for Binary Floating-Point Arithmetic defines functions that overcome important limitations in previous versions, most recently Standard 754-1985. Specifically, the standard describes multiple encoding methods for the numbers used in the calculations.

Under the previous versions of the standard, certain applications for use in regulated industries such as finance were forced to use more resource-intensive encoding to meet accuracy requirements. The accuracy issues were related to complexities in representing decimals precisely in binary format. For example, the following calculation calculated in single-precision yields the result 6.9999997504, rather than 7.00:[3]

(7.00 / 10000.0) * 10000.0

The Intel® Decimal Floating-Point Math Library provides an efficient implementation of binary integer decimal encoding that complies with Standard 754-2008, which enhances the suitability of SAP HANA for use in the full spectrum of industry settings.

### 4.3.3 Benefits from Intel® Turbo Boost Technology

Intel® Turbo Boost Technology[4] automatically allows processor cores to run faster than the base operating frequency when the OS requests the highest processor performance state (P0), if the processor is operating below power, current, and temperature specification limits. The extent of this capability is also dependent on the number of active cores in the processor package, as illustrated in Figure 9.

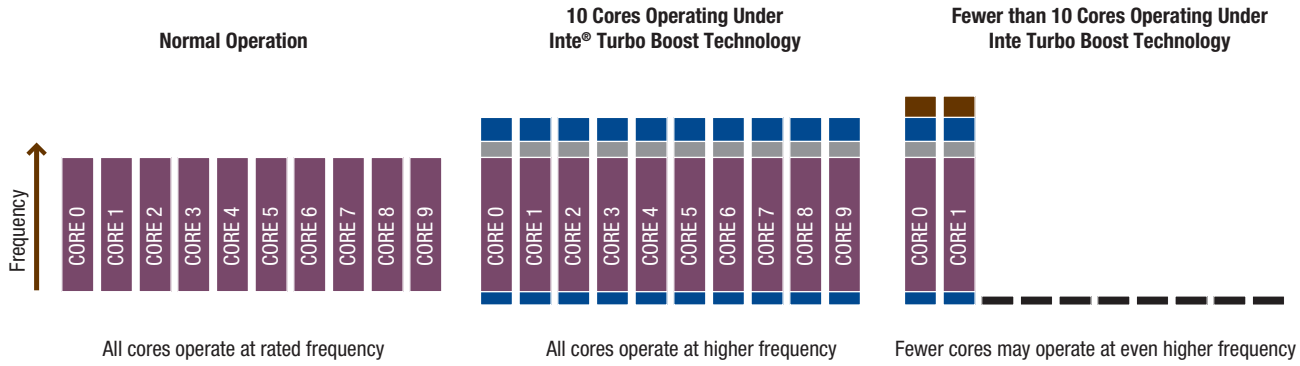


**Figure 9.** Intel® Turbo Boost Technology allows individual cores to operate briefly above the thermal design point, under certain conditions.

| Intel® Instruction Sets | | | | | | |
|---|---|---|---|---|---|---|
| **SSE (1999)**<br>70 Instructions:<br>• Single-precision vectors<br>• Streaming operations | **SSE2 (2000)**<br>144 Instructions:<br>• Double-precision vectors<br>• 128-bit vector integer | **SSE3 (2004)**<br>13 Instructions:<br>• Complex arithmetic | **SSSE3 (2006)**<br>32 Instructions:<br>• Decode operations | **SSE4.1 (2007)**<br>47 Instructions:<br>• Video accelerators<br>• Graphics<br>• Co-processor accelerators | **SSE4.2 (2008)**<br>7 Instructions:<br>• String manipulation<br>• Cyclic Redundancy Check (CRC32) | **Intel® AES-NI (2010)**<br>6 Instructions:<br>• Encryption and decryption |

**Figure 10.** Successive generations of instructions, including Intel® Streaming SIMD Extensions (Intel® SSE) and Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI) accelerate various functions on Intel® architecture.

### 4.3.4 Optimization of Bit Compression with Intel® SSE-Based Vectorization

Intel periodically introduces sets of instructions that extend platform capabilities with certain types of operations, as illustrated in Figure 10. The value of these instructions is primarily to accelerate those operations and secondarily to reduce the complexity of the implementation.

Intel SSE instructions provide the ability to support vector operations on 128-bit wide registers. Theoretically, these operations should result in a 4x speed up on 32-bit data types and a 2x speed up on 64-bit data types. SAP HANA uses this advantage for decompression and search on its base data structures. In cooperation with Intel, SAP engineers have developed a highly optimized set of algorithms. By default, SAP HANA checks on startup to determine which sets of Intel SSE instructions are supported by the processor and chooses the implementation that is best suited to maximize decompression and search performance.

## 4.4 Performance Benefits from Optimization and Hardware Platform Features

This section covers the performance increases enabled by the optimization work on SAP HANA for the Intel processor E7 family conducted jointly by SAP and Intel. It also describes the performance benefits that are available from some key hardware-platform features.

### 4.4.1 Performance Comparison between Processor Generations and Impact of Hardware Features

To demonstrate the performance benefits of implementing SAP HANA on the latest Intel Xeon processors, a test lab at Intel compared results against a reference workload between the Intel® Xeon® processor 7500 series and the Intel Xeon processor E7 family, which is summarized in Figure 11.5 The workload consisted of a single-user scenario that was provided by an SAP customer and used as the internal reference scenario for SAP. It processed 3.2 billion records of point-of-sale data in a complex fashion.

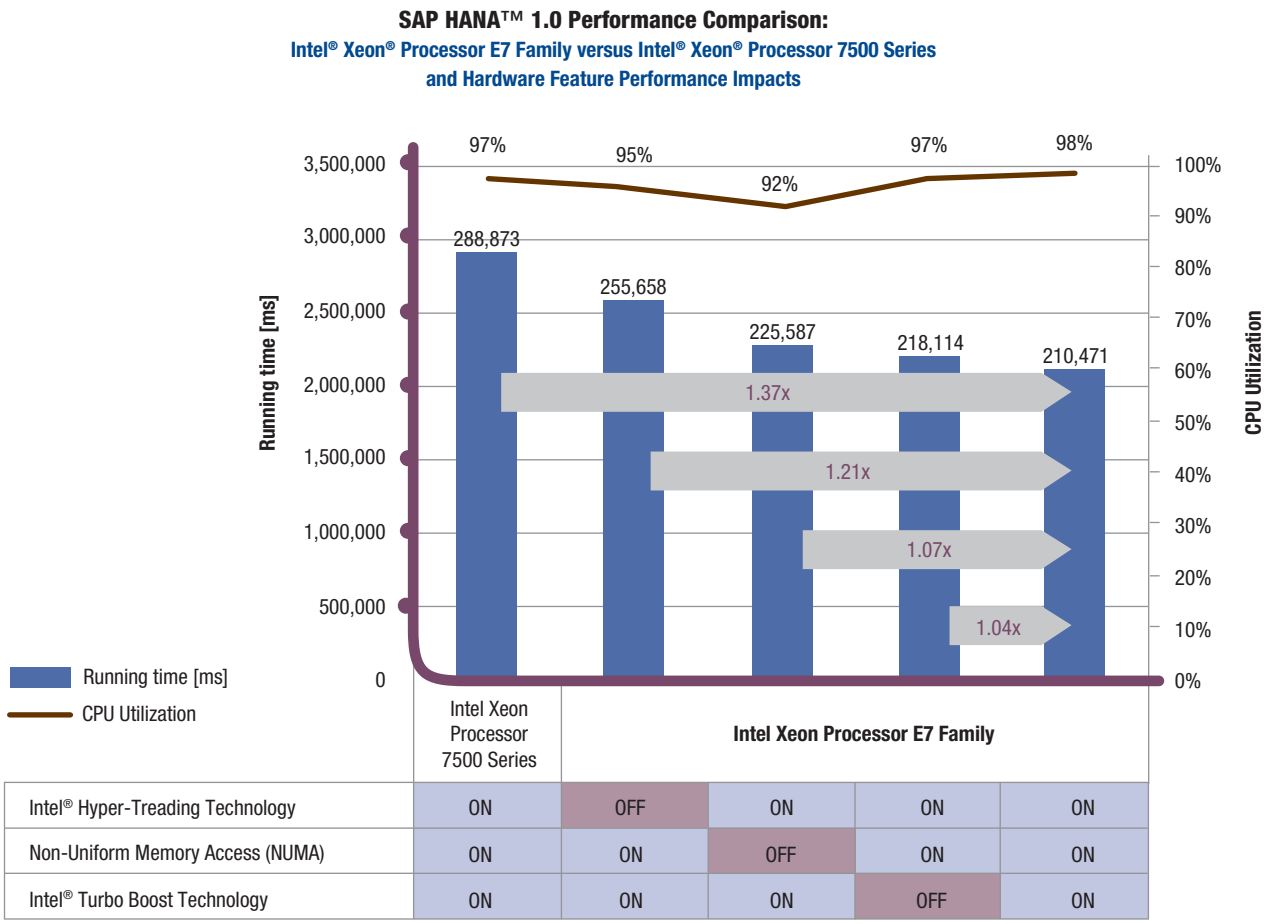


**Figure 11.** Testing demonstrated a 1.37x performance advantage by the Intel® Xeon® processor E7 family over the Intel® Xeon® processor 7500 series and benefits from key platform features.[5]

The testing summarized in Figure 11 consisted of two major aspects: testing the performance differential between the two processor generations and testing the distinct performance impacts of individual hardware-platform features (that is, Intel HT Technology, NUMA, and Intel Turbo Boost Technology. The testing approach consisted of two primary parts:

- **Performance impact of processor generation**. A direct performance comparison was made between the two processor generations under test, with all three of the platform features enabled. This result intended to show the effect of the processor generation in isolation from other aspects of the platform.
- **Performance impact of platform features**. Each of the three platform features was disabled on the Intel Xeon processor E7 family-based system, one at a time with the other two features enabled, and each of those results was compared to the case in which all three features were enabled. This result was used to show the isolated effect of the platform feature.

As reported in Figure 11, the impact of the processor generation alone was that the Intel Xeon processor E7 family delivered a 1.37x performance increase, relative to the Intel Xeon processor 7500 series.[5] On the system based on the Intel Xeon processor E7 family, Intel HT Technology produced a 1.21x performance increase,[5] NUMA technology produced a 1.07x performance increase,[5] and Intel Turbo Boost Technology produced a 1.04x performance increase.[5] Notably, these performance increases were measured under relatively high levels of processor utilization, demonstrating their value in adding to system-level performance headroom.

### 4.4.2 Performance Benefits of Intel® SSE Vectorization for Search

SAP HANA ensures flexibility and makes the most of installed system memory by compressing all base data structures using a packed bitfields approach for each of the bit cases 1 to 32 bits. (For a discussion of compression in the SAP In-Memory Computing Database, see the Creating a Transition in Database Design section.) Using the packed bitfields approach enables the system to store integer values more efficiently than as standard STL vectors.

Instead of using 32 bits per integer, SAP HANA uses only as many bits as are needed to represent the highest distinct value in a vector as a bit string. A separate implementation is required for each bit case, and over the average for all bit cases, using Intel SSE with TREX accelerates search on the compressed data by a factor of 1.6x relative to scalar code on the Intel® Xeon® processor 5500 series, as shown in Figure 12.[6]

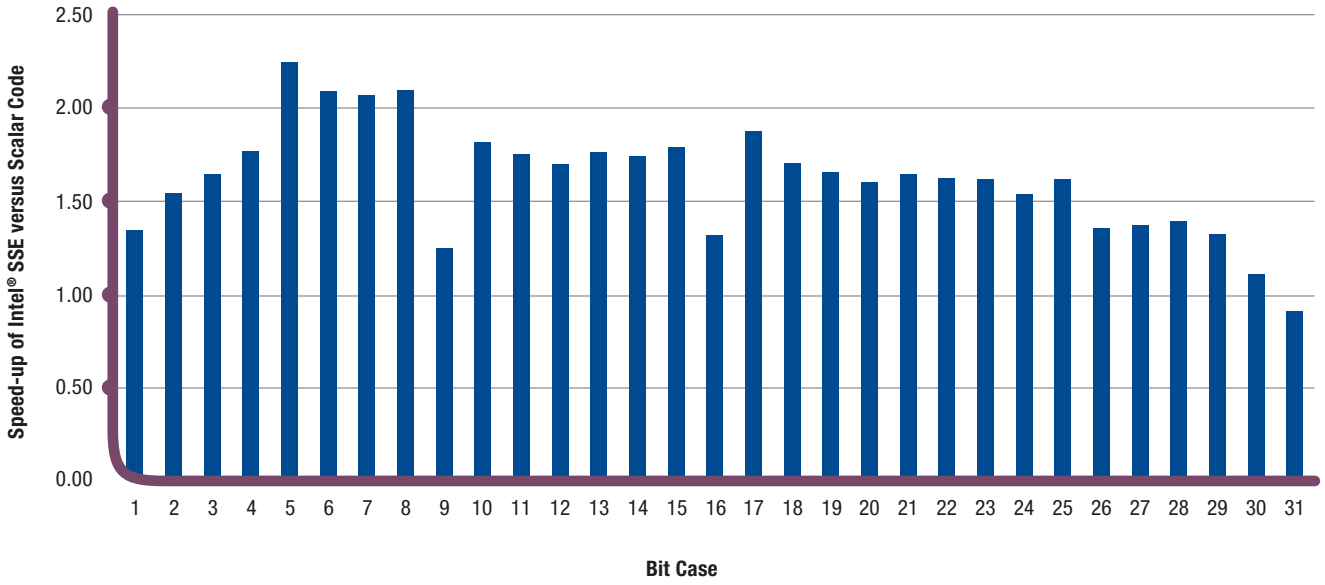


**Figure 12.** On average over all bit cases, Intel Streaming SIMD Extensions (Intel® SSE) accelerates search on compressed data using SAP In-Memory Computing Database technology by a factor of 1.6x on the Intel® Xeon® processor 5500 series.[6]

### 4.4.3 Performance Scaling with Increasing Core Count

To check performance scalability relative to increases in the number of processor cores, the team began by calculating the "ideal" case, in which the number of queries per second would increase linearly on a one-to-one basis with processor frequency (represented by the upper, dashed line in Figure 13). Next, by adjusting the core frequency using BIOS settings, they measured the actual increase (represented by the lower, solid line).

The results show that performance increases relative to increases in processor frequency at an efficiency of 91 to 94 percent.[7] While these results were obtained using SAP Business Warehouse Accelerator, a predecessor technology to SAP HANA, note that both software environments are based on similar in-memory computing technology.

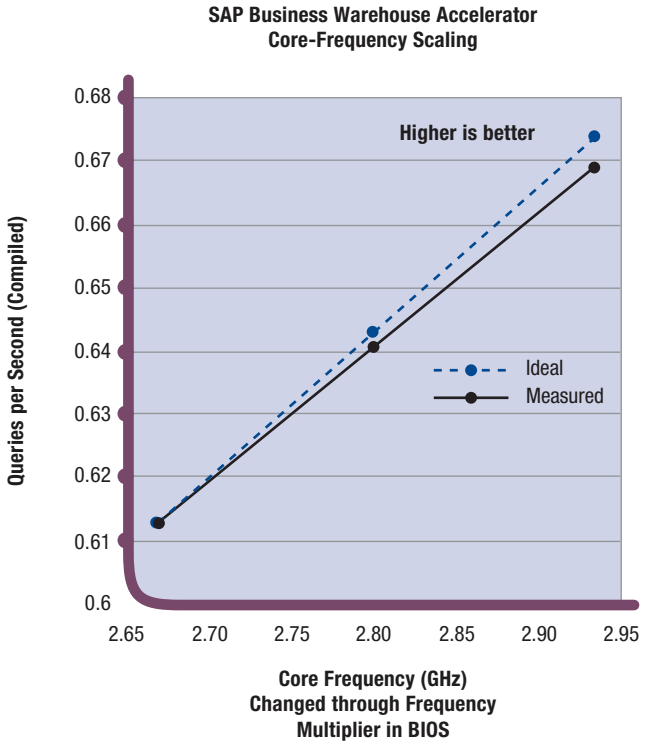


**Figure 13.** SAP Business Warehouse Accelerator workloads scale almost at the ideal rate as processor core frequency increases.[7]

## 5 Enhancing Reliability for In-Memory Computing Systems

### 5.1 In-Memory Recovery

SAP HANA includes software capabilities that take advantage of the hardware features of Intel Xeon processors to dramatically reduce the incidence of system failures. In addition, the SAP HANA appliance can fail over from one compute node to another in the event of a processor failure and to handle memory errors with as little impact to workloads as possible.

A growing scope of RAS capabilities have been introduced in the Intel Xeon processor, reaching a suitability for business-critical implementations that was previously associated with only more expensive RISC and other proprietary solutions. For example, Machine Check Architecture Recovery enables SAP HANA to work with the OS to recover from many otherwise-uncorrectable memory errors, recovering data without downtime in many cases.

### 5.2 ACID Properties

SAP HANA has been designed and refined to provide excellent compliance with the so-called ACID set of properties designed to ensure reliable processing by a database system. In general, those properties are as follows:

- **Atomicity** is that property of a database transaction that guarantees that if the entire transaction cannot be completed, no change is made to the database, meaning that "partial" transactions will not take place. Atomicity is intimately related to the reliability features of the platform described above, in the sense that those characteristics help prevent the system from failing in mid-transaction or help it recover if a failure does occur. In addition to those capabilities, SAP HANA includes programmatic safeguards to ensure Atomicity.
- **Consistency** refers to the assurance that all transactions performed by the database will be taken from one consistent state to another, which includes the requirement that only valid data will be written to the database. In the event that invalid data is written to the database in error, the Consistency property requires that the database can be rolled back to a last known good state. This capability is provided for by the presence of temporal tables described previously in this paper, and it will be enhanced by the future implementation of aged data features also discussed.
- **Isolation** is the requirement that all data associated with a transaction be inaccessible by any other operation for the duration of the transaction. This property is necessary to prevent multiple writes to a single data location during transactions, which could lead to data inconsistency and unpredictable results. Because Isolation requires data to be locked while it is in use for a transaction, SAP HANA uses sophisticated algorithms to minimize the performance impact of that locking.
- **Durability** means that once a transaction has been committed to the database it can always be recovered in the event of a subsequent hardware or software failure. This property prevents completed transactions from having to be reversed. Because in-memory databases use volatile memory as their main data store, SAP HANA relies on the use of flash memory to provide Durability by retaining a combination of snapshots and logs that together assure the recoverability of committed data in the event of a system failure.

## 6 Application Development and Tools for In-Memory Computing

SAP HANA provides a powerful application-development environment specifically tailored to the needs of organizations in building business applications. It also provides SQL Script, a language for processing application-specific code in the database layer, which has the following two advantages:

- **Transfer efficiency**. Moving calculations to the database layer eliminates the need to transfer large amounts of data from the database to the application.
- **Feature utilization**. Calculations need to be executed in the database layer to get the maximum benefit from features such as fast column operations, query optimization, and parallel execution.

### 6.1 Intel® Performance Counter Monitor for Debugging and Optimization

Traditionally, development organizations have relied on the processor utilization utilities provided with OSs to help them understand how the processor is responding to a given workload under test. Unfortunately, as platforms have become more complex, this is no longer a dependable approach, particularly in complex systems that include Intel HT Technology, multi-core processors, and multi-processor systems, all of which are the case with SAP HANA. These issues are explained in the article "Intel® Performance Counter Monitor – A Better Way to Measure CPU Utilization."[8]

To overcome this limitation and easily take advantage of the hardware-resident performance management units built into Intel® platforms, the SAP HANA development team uses Intel® Performance Counter Monitor. The source code of Intel Performance Counter Monitor is integrated into SAP HANA and available through a built-in API and also as a part of SAP HANA's internal management console. Through this functionality, HANA developers obtained a quick and easy processor monitoring facility, built directly into HANA as a standard performance tool. This practice, and the team's ability to draw excellent results from it, is a direct dividend of SAP's collaboration with Intel.

***"Intel® Performance Counter Monitor has been a very valuable tool for us to unlock the full potential of state-of-the-art multi-core processors, especially for non-uniform memory access debugging. As other tools do not allow for such quick and easy monitoring of memory accesses between different sockets, the Intel Performance Counter Monitor was a great help for SAP when developing new parallel algorithms that control the memory location on different sockets for better performance.***

***Without the tool, we were always in doubt whether the applied memory access pattern really works. Additionally, the tool has allowed us to identify unexpected problems related to cache pressure very quickly."***

– Oliver Rebholz, SAP development manager

### 6.2 Intel® Software Development Products

In the course of optimization for Intel architecture, the SAP HANA development team benefited dramatically from the use of Intel® Software Development Products and guidance from their colleagues at Intel. The following tools played particularly key roles:

- Intel® VTune™ Performance Analyzer helps take advantage of hardware counters to identify microarchitectural events, such as cache misses, and tie them back to specific code locations. This capability has played an essential role in optimizing the code for the out-of-order engine and cache hierarchy.
- Intel® Parallel Amplifier simplifies the effort to find multi-core performance bottlenecks, providing for efficient, robust optimization for scaling for the parallelism provided by multi-processor, multi-core systems with Intel HT Technology.

## 7 Real Customer Example: Building Richer Sales Support at Hilti Corporation

SAP HANA puts data at the service of business users, giving them the power to create rich, ad hoc reports and analyze operational data. Conventional data warehouse-based systems depend on canned reports that are typically run nightly, based on set parameters. New reports often require creation of materialized views and summary tables for speeding up performance, which limits the ability to do ad hoc queries and analysis.

Using main memory as the primary data store allows the SAP HANA transactional and analytic engines to be deployed together on a combined application and data layer. This integration speeds up performance and lets users perform data mining and decision-support functions directly based on their needs, drastically reducing their dependency on in-house IT departments to create reports for them.

Liechtenstein-based Hilti Corporation, a provider of equipment and technology to the global construction industry, prides itself on its dynamic sales model as a differentiator within its field. Two thirds of Hilti's 20,000 employees work directly for the customer in sales and engineering, which generates more than 200,000 customer contacts every day. IT systems based on SAP technology play a vital role in supporting those operations.

Hilti worked with SAP on the development of SAP HANA, with an eye toward the goal of maximizing operational flexibility and performance. The company's SAP HANA-based systems combine the transactional and analytic solution layers to transform the sales and support process. In early testing, these systems reduced the time required for search and analysis of 53 million customer data records from two or three hours to two or three seconds.

This advance promises to give sales employees access to robust reporting functionality in real time, including access from mobile devices in the field. Rapid analysis and business intelligence allow for very productive customer contacts based on insights that are immediately available to support business decisions where and when they are needed.

## 8 Outlook and Future Research Topics

As developments continue in both hardware and software technologies, the value of SAP HANA across industries can be expected to continue to grow. While many of these advances cannot be predicted, current trends and research topics suggest likely possibilities.

### 8.1 Evolution of Non-Volatile Memory

In the near term, existing forms of non-volatile memory, such as flash, are likely to continue to improve in terms of durability, scalability, and speed, while also becoming less expensive. Looking out a bit further, new types such as phase-change memory (PCM) are likely to become more viable for mainstream use, advancing the benefits even more. General advancement in the area of non-volatile memory is expected to benefit SAP HANA, which may make use of the new offerings for storage of aged data, data stores for recoverability, and so on.

### 8.2 Hardware Acceleration

As new Intel® hardware platforms are introduced, the number and scope of formerly software-only operations that can be assisted with hardware acceleration increases. Recent examples include the growing body of capabilities under the umbrella of Intel® Virtualization Technology and the hardware-based graphics acceleration provided by the Intel Xeon processor E7 family and other platforms based on that generation of microarchitecture. Some of these hardware-acceleration capabilities are likely to have benefits for operations performed by SAP HANA, which may help to expand the product's performance and functionality.

### 8.3 New Instructions

Another ongoing evolution in Intel® technology is the continuing development of new instruction sets that enable developers to more readily take advantage of processor features and capabilities. Intel SSE instructions, discussed earlier in this paper, are an example that presages likely future advances. Intel® Advanced Encryption Standard (Intel® AES) New Instructions are a recently introduced instruction set that accelerates encrypt and decrypt operations, making it more feasible to implement end-to-end encryption, which could potentially benefit future generations of SAP HANA. Because of the capability of new instruction sets in many cases to accelerate various functions of the hardware platform, future instruction sets are likely to benefit SAP HANA in ways that may not presently be apparent.

### 8.4 Many Cores

As the numbers of processor cores increase in future Intel platforms, the level of hardware parallelism available to SAP HANA will increase. This development will present new opportunities for dividing work within the solution, with positive implications for performance and scalability. Other development concepts include the idea of heterogeneous processor cores, where a processor could have different types of special-purpose execution units designed for specific functions, such as high floating-point performance. Such advances would require new programming techniques and in return new functionality in SAP HANA solutions with great benefits to the overlaying applications.

## 9 Conclusion

SAP HANA enables real-time decision making by bringing all the data in your enterprise within the reach of decision makers in seconds, not weeks or months, in an easy-to-use format so your company can run smarter and perform better. The optimized software components based on SAP In-Memory Computing are highly optimized to take excellent advantage of the hardware features of the latest Intel Xeon processors.

The next-generation performance available from SAP HANA running on the latest Intel Xeon processors draws on a long-standing tradition of co-engineering between the companies. Key Intel hardware advances range from increasing server parallelism to innovations in the memory subsystem such as growing caches and ever-faster access to larger amounts of system memory.

SAP software advances have built on those hardware designs to create the SAP In-Memory Computing Database, which has now evolved to support SAP HANA in delivering an unprecedented robustness in real-time business analysis. Ongoing research and development at SAP, in collaboration with their colleagues at Intel, continue to advance the state of the art in storage handling, with columnar storage and related innovations around compression and data movement.

SAP HANA is highly tuned at the inter-socket, inter-core, and intra-core levels. This work enables the solution to deliver optimal performance and scalability from systems based on Intel Xeon processors. Encompassing key performance aspects such as NUMA, multi-threading, and platform-specific instruction sets, SAP HANA has been shown to deliver as much as a 1.37x performance increase on the Intel Xeon processor E7 family, relative to the predecessor Intel Xeon processor 7500 series,[5] for excellent ROI.

In addition, solution development has carefully incorporated best practices for taking advantage of the extensive RAS capabilities of the Intel Xeon processor E7 family. Making use of hardware features such as Machine Check Architecture Recovery complements SAP HANA's excellent ACID compliance to provide a solution stack that is well suited for mission-critical implementations. Future research and development on both hardware and software capabilities are expected to add even further to this robustness.

The solution stack available from SAP HANA enables innovative new applications, combining high-volume transactions with analytics, to dramatically improve existing planning, forecasting, pricing, and other processes. It is based on proven technology that integrates directly and easily with existing data sources. A 360-degree view of business operations is available today, together with the power to make sense of it in real time, for better decision making that enables your next generation of success.

To learn more about SAP HANA, visit:

**www.sap.com/platform/in-memory-computing**

To learn more about Intel® Xeon® processors, visit:

**www.intel.com/xeon**

## 10 Further Reading


Athanassoulis, M., A. Ailamaki, S. Chen, P. Gibbons, and R. Stoica, "Flash in a DBMS: Where and How?" *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* (2010). http://www.pittsburgh.intel-research.net/~chensm/papers/flash-DEBull10.pdf

Bog, A., J. Krüger, A. Zeier, and J. Schaffner, "A Composite Benchmark for Online Transaction Processing and Operational Reporting." Hasso Plattner Institute, University of Potsdam (2009). http://epic.hpi.uni-potsdam.de/pub/Home/Sapphire09/Composite_Benchmark_for_OLTP_and_Operational_Reporting_-_Final.pdf

Bruggeman, O., T. Willhalm, and R. Dementiev, "Intel® Performance Counter Monitor - A Better Way to Measure CPU Utilization." Intel® Software Network (August 8, 2010). http://software.intel.com/en-us/articles/intel-performance-counter-monitor/

Cornea, M., and J. Crawford, "IEEE 754R Decimal Floating-Point Arithmetic: Reliable and Efficient Implementation for Intel® Architecture Platforms." *Intel Technology Journal* (February 2007). http://www.intel.com/technology/itj/2007/v11i1/s2-decimal/1-sidebar.htm

Dunning, D., R. Mooney, P. Stolt, and B. Casper,"Tera-scale Memory Challenges and Solutions." *Intel Technology Journal* (2009). http://download.intel.com/technology/itj/2009/v13i4/pdfs/ITJ9.4.7_MemoryChallenges.pdf

Intel Corporation, "An Introduction to the Intel® QuickPath Interconnect." (January 2009). http://www.intel.com/technology/quickpath/introduction.pdf

Krueger, J., M. Grund, C. Tinnefeld, J. Schaffner, S. Mueller, and A. Zeier, "Enterprise Data Management in Mixed Workload Environments." Hasso Plattner Institute for IT Systems Engineering. (2009). http://ares.epic.hpi.uni-potsdam.de/apps/static/papers/2009_JK_Enterprise_Data_Management_in_Mixed_Enviroments.pdf

Plattner, H., "A Common Database Approach for OLTP and OLAP Using an In-Memory Column Database." SIGMOD'09. (June 29–July 2, 2009). http://www.sigmod09.org/images/sigmod1ktp-plattner.pdf

Plattner, H., "Enterprise Applications – OLTP and OLAP – Share One Database Architecture." Hasso Plattner Institute for IT Systems Engineering. (2010). http://epic.hpi.uni-potsdam.de/pub/Home/InMemoryDataProcessing2010/ShareOneDB.pdf

Plattner, H., and A. Zeier, *In-Memory Data Management: An Inflection Point for Enterprise Applications.* 1st ed. April 29, 2011.

Schaffner, J., A. Bog, J. Krüger, and A. Zeier, "A Hybrid Row-Column OLTP Database Architecture for Operational Reporting." Hasso Plattner Institute for IT Systems Engineering (2009). http://epic.hpi.uni-potsdam.de/pub/Home/Sapphire09/p7-schaffner.pdf

Schroeder, B., E. Pinheiro, and W. Weber, "DRAM Errors in the Wild: A Large-Scale Field Study." SIGMETRICS/Performance'09 (June 15–19, 2009). http://www.cs.toronto.edu/~bianca/papers/sigmetrics09.pdf

Tinnefeld, C., J. Krueger, J. Schaffner, and A. Bog, "A Database Engine for Flexible Real-Time Available-to-Promise." Hasso-Plattner-Institut, University of Potsdam (2009). http://epic.hpi.uni-potsdam.de/pub/Home/Sapphire09/Real_Time_ATP.pdf

Valles, A., M. Gillespie, and G. Drysdale, "Performance Insights to Intel® Hyper-Threading Technology." Intel Software Network (November 20, 2009). http://software.intel.com/en-us/articles/performance-insights-to-intel-hyper-threading-technology/

[1]http://www.sigmod09.org/images/sigmod1ktp-plattner.pdf.

[2]http://software.intel.com/en-us/articles/performance-insights-to-intel-hyper-threading-technology.

[3]See, for example, http://www.intel.com/technology/itj/2007/v11i1/s2-decimal/1-sidebar.htm.

[4]Intel® Turbo Boost Technology requires a system with Intel Turbo Boost Technology capability. Consult your PC manufacturer. Performance varies depending on hardware, software, and system configuration. For more information, visit www.intel.com/technology/turboboost/.

[5]Test configurations:
Intel® Xeon® processor 7500 series running at 2.26 GHz, 256-GB DDR3-1066 RAM, two SAS hard disks in RAID-0 configuration, SUSE* Linux* Enterprise Server 11 SP1, HANA 1.0, Database Release 1.00.02.292791.
Intel® Xeon® processor E7 family running at 2.40 GHz, 256-GB DDR3-1066 low voltage RAM, two SAS hard disks in RAID-0 configuration, SUSE* Linux* Enterprise Server 11 SP1, HANA 1.0, Database Release 1.00.02.292791.

[6]Testing performed on two-way servers based on Intel® Xeon® processors X5560 running at 2.8 GHz.

[7]Testing performed on two-way servers based on Intel® Xeon® processors 5500 series running at 2.93 GHz and 2.8 GHz, NUMA enabled, Intel® Hyper-Threading Technology enabled, Intel® Turbo Boost Technology disabled, 24 GB RAM, SAP Business Warehouse Accelerator 720.00.169728 (DEV), two index servers pinned to memory nodes, Polestar colocto data set(< 8 GB), one user (54 different queries); CPU utilization: about 96 percent.

[8]http://software.intel.com/en-us/articles/intel-performance-counter-monitor/.



0411/SM/MESH/PDF 325214-001US