White Paper
Intel® Xeon® Processor E7 v2 Family
In-Memory Computing

Changing the Way Businesses Compute...and Compete

In-Memory Computing and Real-Time Business Intelligence

"It's not going to happen overnight. But within the next 15 years, you will run your whole operation in memory."[1]

– Donald Feinberg, Gartner VP and Distinguished Analyst, April 2013

Robert Shiveley
Datacenter Software Marketing Manager,
Intel Corporation

Executive Summary

Every now and then, a technology comes along that changes the way we live and work in a fundamental way—the personal computer, mobile computing, and the Internet are classic examples. In-memory computing has the potential to drive comparable levels of game-changing innovation. By allowing businesses to filter and analyze vast amounts of data in fractions of a second, it enables real-time, data-driven decision making based on all relevant data.

Early adopters are already using in-memory computing to redefine the competitive landscape in their industries. A global enterprise closes its books in just seconds, a retailer adapts prices and inventory levels based on real-time sales trends, and an equipment manufacturer provides worldwide customers with predictive field maintenance to improve productivity. These are just a few examples out of many. Thousands of companies have implemented in-memory computing, and many are doing things that would not otherwise be possible.

Ultimately, in-memory computing will not be an alternative to traditional solutions. It will simply be the way computing is done. Real-time analysis of massive data sets will be a common-place capability. Until then, when and how companies begin their transition to in-memory computing and real-time business intelligence (BI) will be an increasingly critical business decision.

This paper explores the technologies that have come together to make in-memory computing both technically possible and economically feasible. It discusses several current and upcoming in-memory database solutions, and then focuses on Intel® processor and platform innovations—such as the Intel® Xeon® processor E7 v2 family—that are helping businesses achieve leadership performance and mission-critical availability for in-memory computing using cost-effective, industry-standard computing platforms.

Table of Contents

The "Big Data" Challenge

The term big data refers to the growing flood of data in today's hyperconnected world, and to the enormous challenges and opportunities this data presents. In a very real sense, however, data has been too big for decades. Since the first Intel processor was launched more than forty years ago, the data access speeds of bulk storage systems have not kept pace with processor capability.

Processor speed and functionality are delivered in silicon, and have benefitted from the rapid transition toward smaller and faster silicon process technologies predicted by Moore's Law.[2] Bulk storage has long been based on mechanical technologies, from punch cards to magnetic tape to the spinning magnetic disks of today's hard disk drives (HDDs). Although the cost per gigabyte (GB) of storage has declined rapidly over the years, storage performance has been limited by the mechanical nature of the data access process.

Faster, silicon-based data storage technologies have existed for many years, but have been too costly for bulk storage.[3] Instead, such technologies have been used for the main memory of computing systems and for the even faster cache subsystems that reside directly on the processor die. Although these high-speed memory subsystems ameliorate the data access problem to some degree, their limited capacity has been a performance bottleneck. Getting the right data out of bulk storage and into the right processor registers at the right time has been a tough challenge for decades.

Database vendors have done much over the years to work around and mask this performance gap, but the resulting cost and complexity have been significant. As illustrated in Figure 1, traditional HDD-dependent information infrastructure requires:

- **Separate databases for transactional and analytical applications,** along with separate infrastructure. In each case, hardware and software must be optimized to achieve acceptable performance.
- **Multiple data marts** to address specialized business intelligence needs without overloading data warehouses.
- **Constant tuning and optimization of databases and storage systems** to deliver acceptable performance, especially for analytical workloads.

Despite all the cost and effort, customers still experience long delays between the time that data is generated and the time it is available for analysis. Data must be extracted from transactional systems, transformed into required formats, and loaded into the analytics environment. In many cases, the data models in the warehouse or data mart must then be re-optimized for performance. Even with all this preparatory work, complex queries can still take many hours to complete.

Harnessing Big Data with In-Memory Computing

In-memory computing changes the computing paradigm fundamentally (see Figure 2 on the next page). All relevant data is kept in the main memory of the computing system, rather than in a separate bulk storage system.

- **Data can be accessed orders of magnitude faster,** so fast that transactional and analytical applications can run simultaneously on the same database running on the same infrastructure. The need for separate data warehouses and data marts is eliminated, along with the associated costs.[4]
- **Data is available for analysis as soon as it is generated.** Even if it is generated on a separate system, it can be replicated almost instantly into an in-memory database.

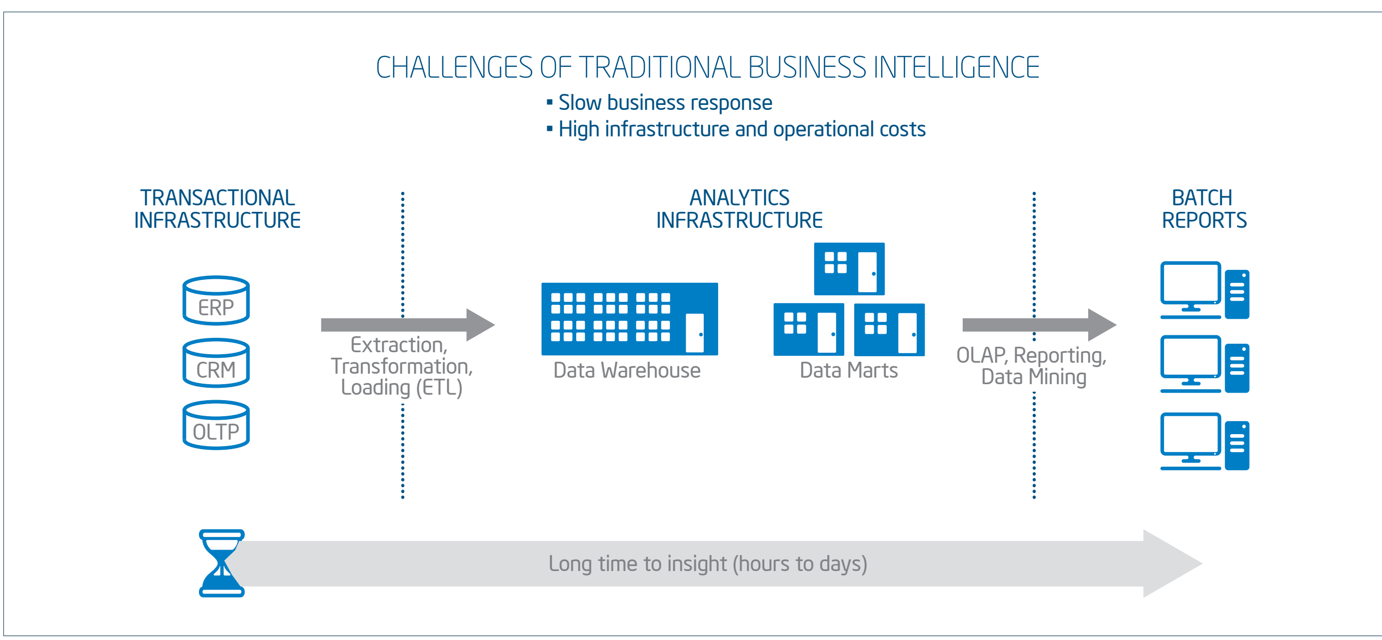

Figure 1. Traditional, HDD-dependent information infrastructure requires separate transactional and analytical infrastructure, resulting in high costs and long delays between data generation and insight.

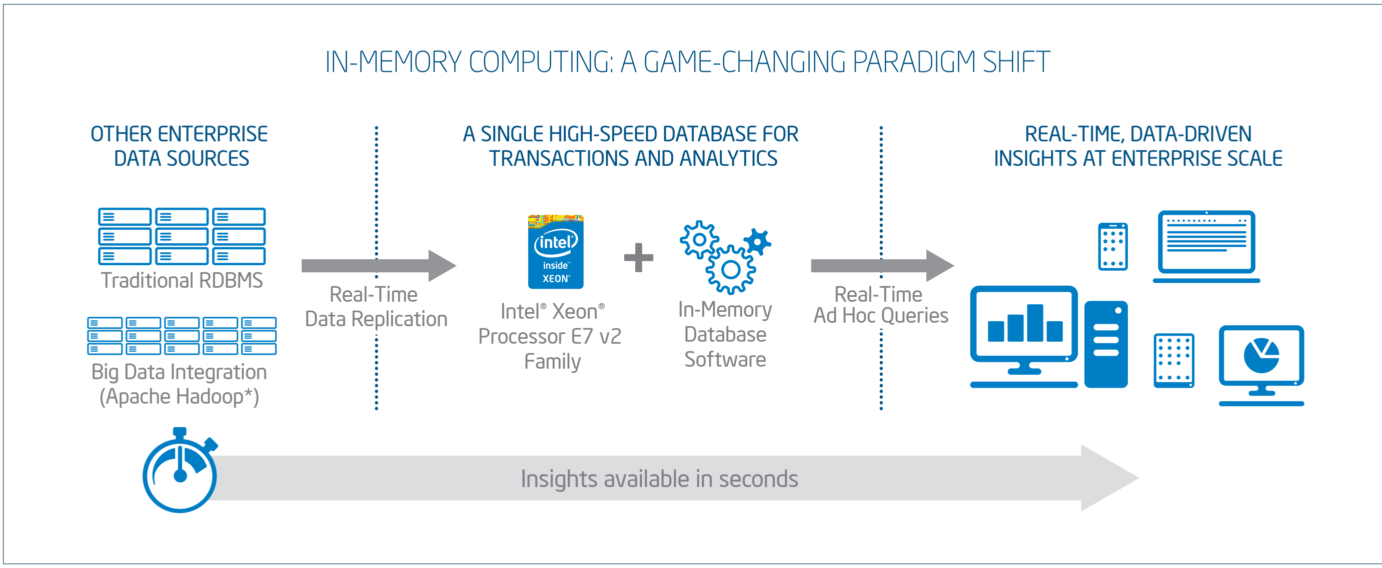

Figure 2. With in-memory computing, traditional latencies are greatly reduced to enable real-time, data-driven decision making based on all relevant data sets.

- **Data-driven business insights are available almost instantly.** Complex queries can be completed in a matter of seconds or fractions of a second, and the results can be funneled immediately back into transactional applications to improve outcomes.

A number of technology advances have come together to make in-memory computing a viable option. The first is the availability of affordable memory modules. The cost of memory has declined for years and has reached a tipping point at which the value of in-memory computing can exceed the cost in many business scenarios. According to Gartner, memory costs continue to drop by roughly 32 percent per year,[1] so the economics of in-memory computing will continue to improve.

Businesses need more than affordable memory to implement new in-memory computing models. They also need production-ready hardware and software solutions, and both are available today. Servers based on the Intel® Xeon® processor E7 v2 family provide the memory capacity, execution resources, and reliability needed to support in-memory computing solutions at enterprise-scale. In-memory database software is also available today from dozens of vendors. The transition to this new paradigm has begun and will continue to accelerate. The potential benefits are already too great to ignore for many businesses, and the associated costs will continue to fall.

The Power of Real-Time Business Intelligence

In-memory computing is already driving real-world business innovation. There are several production-ready in-memory databases available today. A prime example is SAP HANA,* which was developed jointly by Intel and SAP. More than a thousand customers across nearly 30 industries have implemented SAP HANA and have realized an average performance gain of 12,380 times.[6] For example:[7]

- **Yodobashi,** a leading Japanese retailer calculates loyalty card points for 5 million members. This process used to take three days and was performed once per month. Now it can be completed in just two seconds. The resulting insights are used to improve outcomes in real-time customer engagements by providing personalized offers based on loyalty status and product inventories.
- **Centrica,** a large utility company, rolled out smart meter technology across the United Kingdom. It used to take three days to analyze the data. The calculations are now completed in seconds to improve pricing decisions and forecast accuracy, and to provide consumers and businesses with customized information, services, and products.

> "It is now possible to store a primary database in silicon-based main memory, resulting in an orders-of-magnitude improvement in performance and enabling the development of completely new applications."[5]
>
> *– The Global Information Technology Report 2012*

- **T-Mobile,** one of the world's leading wireless providers, runs reports on 2 billion customer records in just 5 seconds to provide targeting marketing to 21 million customers. The analyses are not only much faster than the company's previous solution, but also much deeper. Instead of three months of customer records, they now analyze 24 months for better insight into customer buying patterns.

The speed and scale of SAP HANA can be applied to almost any data-intensive business challenge. It is in production use for functions as diverse as retail stocking and supply chain optimization, global financial closes, personalized medical treatments, and genomic analysis. Across all of these usage models and many more, insights are available in seconds rather than hours or days. New business models are possible and customers can analyze their data far more extensively without overloading their computing infrastructure.

In-Memory Computing is Going Mainstream

In-Memory Databases from Leading Vendors

Dozens of software companies are delivering or developing in-memory database solutions. Approaches differ somewhat, but they tend to include many of the same core capabilities.

For example, traditional databases store data in row-based tables, because that's the most efficient format for the frequent updates required by transactional applications. However, column-based tables are more efficient for analytics, which depend more on fast access to high-volume data. Current in-memory databases address this need in different ways. Some keep primary data in row-based tables and automatically generate specialized structures for analytics. Others keep primary data in columns and generate specialized structures for transactions. A few keep parallel sets of column-based and row-based tables.

Although memory prices have decreased, it is still important to use available memory as efficiently as possible, so most in-memory solutions provide some level of data compression. Compression ratios for column-based data can be as high as 10-to-1, and in some cases even higher, which greatly increases the amount of data that can be held in memory.

Many additional software optimizations are required to optimize the value of in-memory computing. Large volumes of data are available from a simple fetch operation, which eliminates the need to read and then write data from one location to another. This changes many key processes in the database engine. Instead of trying to optimize the use of limited memory resources, developers try to optimize the use of limited cache resources. They also work to parallelize queries to take advantage of the parallel execution resources in today's multicore processors.

All of the leading database vendors, including SAP, IBM, Oracle and Microsoft, are offering or will soon offer in-memory databases.

SAP HANA*

SAP HANA was the first in-memory computing platform to be available from a major software vendor. SAP has built a real-time business platform on top of its high-speed in-memory database. SAP HANA can be used to support both SAP Business Suite* applications and the SAP BW Data Warehouse* application. SAP HANA also provides a computation engine within the database, so data-intensive application logic can be executed closer to high-volume data stores. This further reduces the need for large-scale data movement and can help to reduce both application latencies and infrastructure requirements. SAP HANA can be integrated non-disruptively into existing SAP landscapes and SAP recommends it for all new SAP software deployments.

IBM DB2 with BLU Acceleration*

IBM's in-memory database option is available today to accelerate transactions and especially analytics. IBM DB2 with BLU Acceleration includes columnar data compression, parallel vector processing for improved parallel throughput, and a data-skipping capability that provides faster access to columnar data. In recent tests, Intel engineers found that running IBM DB2 with BLU Acceleration on the latest Intel Xeon processor E7 v2 family increased query performance by up to 148 times versus running the same workload on IBM DB2 10.1 and a prior-generation server processor.[8] The new in-memory solution also improved data compression by 4.5 times, holding a 10 TB data set in just 2,127 GB of memory.

Oracle Database* In-Memory Option

The Oracle Database In-memory option for Oracle 12c is currently scheduled for release later this year. This in-memory solution uses a "dual-format" approach to deliver high performance for both transactional and analytical applications, while maintaining consistency across workloads. Persistent data is stored on disk in conventional row-based format to preserve traditional, disk-optimized operational characteristics. According to Oracle, existing applications are accelerated without requiring any changes and all of Oracle's availability, security, and management features continue to work without change.[9] Support for columnar data compression is expected in a second, minor release, which is also slated for this year.

Microsoft SQL Server*

Microsoft's in-memory solution has been developed under the code name Hekaton, and Microsoft says it will be included in the next major release of SQL Server. This in-memory solution is optimized for online transaction processing (OLTP) workloads, and complements the in-memory column-store index, which is supported in Microsoft SQL Server 2012. Microsoft is implementing a new "latch-free" approach that maintains database consistency without locking data during transactions. This should help to avoid data contention and further improve performance in complex transactional and analytical environments.[10]

Cost-Effective Platforms for In-Memory Computing

Like a number of other vendors, Intel foresaw the emergence of in-memory computing many years ago. The Intel® Xeon® processor 7500 series, and the follow-on Intel Xeon processor E7 family, were designed with in-memory computing in mind and were based, in part, on the expertise Intel engineers developed while working with SAP to develop SAP HANA. The result is a processor family and server platform with the memory capacity, parallel execution resources, system bandwidth, and advanced reliability features needed to address the heavy demands of large, mission-critical, in-memory databases.

Optimized Server Platforms

The Intel Xeon processor E7 v2 family represents another leap forward in this platform evolution (Figure 3). Intel internal performance tests indicate these new servers cut query processing times roughly in half for both SAP HANA[11] and for IBM DB2 with BLU Acceleration[12] in comparison with previous-generation systems based on the Intel Xeon processor E7 family. These gains not only speed time to insight, but also enable businesses to process twice as many queries in a given time frame.

Even more importantly for some businesses, the Intel Xeon processor E7 v2 family supports major increases in memory capacity, up to three times as much as the previous generation.[13] A four-socket server can be configured with up to six terabytes (TB) of memory and an eight-socket server with up to 12 TB, enough to hold many of today's largest databases within the memory of a single server.

The Intel Xeon processor E7 v2 family also provides up to 37.5 megabytes (MB) of third-level cache per processor and up to 50 percent more cores and threads than previous-generation processors. A four-socket server provides up to 60 cores and 120 threads for rapid execution of simultaneous transactions and large, complex queries.

Although most data is held in main memory, I/O bandwidth remains important to speed communications with persistent storage and to quickly ingest streaming data from high-volume sources, such as networked sensors. The Intel Xeon processor E7 v2 family increases system bandwidth by up to 400 percent versus previous-generation processors to meet these demands.[14]

High-Performance Storage

Despite the greater reliance on system memory, the demands on persistent storage also increase. In order to maintain SQL compliance, database transactions must be atomic, consistent, isolated, and durable (ACID). The need for durability, in particular, means that transaction logs must be written to persistent storage before transactions can be completed in memory.

Intel® Solid State Drives (Intel® SSDs) and Intel® 10 Gigabit Ethernet technology complement the in-memory capabilities of the server platform by supporting low-latency, high-bandwidth persistent storage that is standards-based and cost-effective. Intel SSDs provide up to four times faster reads and writes than traditional HDDs, to help in-memory databases deliver ACID-compliance while processing enterprise-scale workloads with real-time performance. Intel 10 Gigabit Ethernet provides the speed and bandwidth to avoid network bottlenecks in this exceptionally data-intensive environment, while also help to reduce network cost and complexity.

Mission-Critical Reliability

Performance is not the only critical issue. In-memory computing systems handle more and heavier workloads per server than traditional solutions and are typically used to support mission-critical, revenue-generating processes. Data integrity and system uptime become more important than ever.

Servers based on the Intel Xeon processor E7 family are already providing levels of high availability that rival those of high-end UNIX*/RISC architectures, but with greater flexibility and superior cost models. The Intel Xeon processor E7 v2 family includes Intel® Run Sure Technology,[16] which delivers additional and even more robust reliability, availability, and serviceability (RAS) features. Advanced error detection, correction, containment, and recovery capabilities are integrated throughout the platform, with robust memory protection to help keep mission-critical systems operating at peak efficiency and without interruption or data errors.

"Taken together, these two trends—declining memory costs and the ascendency of massively multi-core processor architectures—are as transformative to enterprise software as the move to a client-server model was a generation ago."[15]

– The Global Information Technology Report 2012

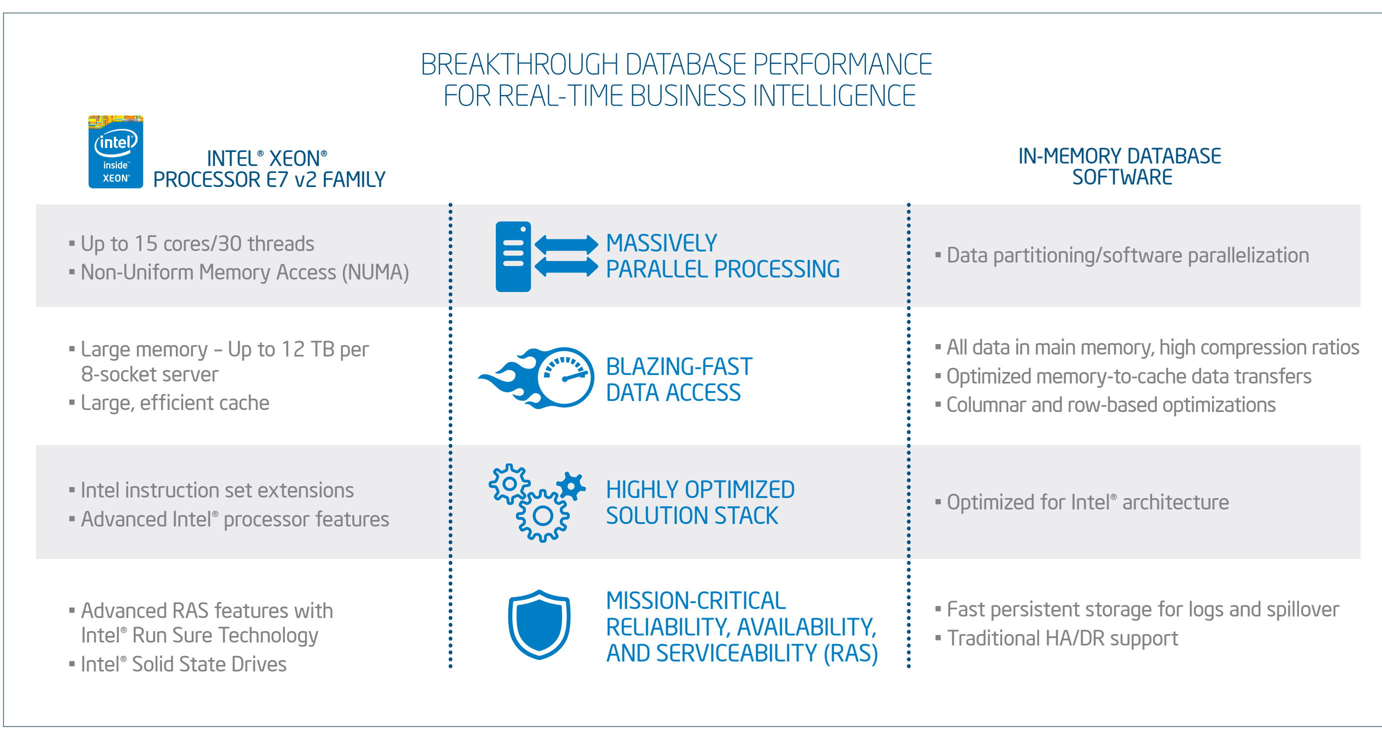

Figure 3. Breakthrough database performance for real-time business intelligence.

Big Data Integration for Massive Scalability

Many businesses already need to store and analyze terabytes and in some cases petabytes of data. It is possible to build a real-time, in-memory business platform with petabyte scalability (see the sidebar, In-Memory Computing at Petabyte Scale). However, this approach is neither financially practical nor necessary for most businesses. Apache Hadoop* offers a more cost-effective solution for integrating very large data volumes with in-memory computing environments.

Hadoop provides massively scalable storage and analytics on a distributed architecture based on affordable servers and commodity disk drives. It offers a cost-effective solution for ingesting, preparing, and storing warm data for inclusion in the real-time analytics environment. Petabytes of data, including all data types, can be stored at a cost-per-terabyte that is not only much lower than an in-memory database, but also much lower than a traditional, disk-based storage system.

Intel and SAP offer an integrated solution today based on SAP HANA and the Intel® Distribution for Apache Hadoop (IDH) software. Business users and data analysts see the data in Hadoop as an extension of the SAP HANA data set, and queries are automatically federated across both platforms. IDH also provides enterprise-class tools and capabilities for management and data security. The combined solution supports real-time analytics acting on petabytes of data.

Other database vendors are following suit. Hadoop is becoming the defacto standard for storing and analyzing massive, unstructured data sets. There may come a time when it is economical to store all data in main memory. Until then, integrating Hadoop and other massively scalable solutions with in-memory computing platforms will be key to optimizing capability versus cost across all enterprise data and all business requirements.

Where to Start

In time, all business computing will be done in memory. Today, businesses have to move forward intelligently, by balancing cost, risk and value. Companies with high-value use cases that cannot be solved using traditional tools should consider implementing in-memory computing sooner rather than later. For others, it may make more sense to wait. As vendors continue to integrate in-memory capability into their core products, costs will go down and integration will become simpler.

Regardless of your current needs and goals, now is the time to:

- **Evaluate the potential of in-memory computing in your specific business and industry.** Work with business units to identify potential, high-value use cases. Consider what you could do better, faster, or different if you could analyze large data sets, including fresh operational data, almost instantly.
- **Explore current in-memory solutions and track progress as new solutions emerge.** Given the proven business value and the maturity of the enabling technologies, in-memory computing can be expected to advance rapidly.
- **Quantify the potential business value and consider the cost and risk of implementation.** There is no doubt that a time will come when the benefits of in-memory computing exceed the cost and risk of implementation for your business. The potential benefits are huge. Be prepared so you can make the right move at the right time.

Conclusion

In-memory computing represents a paradigm shift in the way businesses manage and use data. The unprecedented speed and scale of in-memory databases allows companies to host transactional and analytical applications on the same database. Operational data is available for analysis as soon as it is generated, and complex queries can be completed in seconds rather than hours or days. Infrastructure and operational requirements are also greatly reduced, which can lead to dramatic savings in total cost of ownership.

In-Memory Computing at Petabyte Scale

In May 2012, Intel and SAP launched the SAP HANA* petascale cloud, a 100 TB in-memory system consisting of 100 servers based on the Intel® Xeon® processor E7 family. They have since expanded this cloud infrastructure to include more than 250 servers, 8,000 threads, 4,000 cores and 250 TB of RAM, all capable of running a single instance of SAP HANA.

Used for customer proof-of-concept projects and as a laboratory for Intel and SAP research teams, this petascale cloud environment has clearly demonstrated that in-memory computing can scale to deliver real-time performance while acting on massive data volumes.

For more information, see the Intel and SAP solution brief, ***"Scaling Real-Time Analytics across the Enterprise—and into the Cloud"***

In-memory solutions are available today from dozens of vendors, and all major database vendors, including SAP, IBM, Oracle, and Microsoft, offer or will soon offer in-memory options. The Intel Xeon processor E7 v2 family powers a new generation of servers that are specifically optimized for in-memory computing, delivering up to 2x higher performance than previous-generation servers[11] and providing up to 3x higher memory scalability. They are ideal for the data-intensive, mission-critical demands of in-memory computing.

Learn More

Visit the following sites for more information about in-memory computing:

VENDOR SOLUTIONS

- SAP HANA
  www.saphana.com/welcome
- IBM DB2 with BLU Acceleration
  www-01.ibm.com/software/data/db2/linux-unix-windows/db2-blu-acceleration/
- Oracle Database In-Memory Option
  www.oracle.com/us/corporate/features/database-in-memory-option/index.html
- Microsoft SQL Server In-Memory Project (code-named Hekaton)
  research.microsoft.com/en-us/news/features/hekaton-122012.aspx

INTEL SERVER PLATFORMS FOR IN-MEMORY COMPUTING

- Intel Xeon Processor E7 v2 Family
  www.intel.com/content/www/us/en/processors/xeon/xeon-processor-e7-family.html

[1] Source: Quoted in "Why In-Memory Computing is Cheaper and Changes Everything," a Business Analytics blog post by Timo Elliot, SAP Business Objects, 17 April, 2013. http://timoelliott.com/blog/2013/04/why-in-memory-computing-is-cheaper-and-changes-everything.html

[2] Moore's Law, first proposed by Intel co-founder Gordon Moore in 1965, states that the number of transistors on integrated circuits would continue to double every two years.

[3] The economics of silicon-based storage has improved with the growing maturity of flash memory, as embodied by Intel® Solid State Drives. Other non-volatile memory technologies currently in development also hold considerable promise for the future. Cost versus capacity for these silicon-based storage technologies currently precludes their use as a large-scale replacement for disk-based storage, although they have considerable value as storage accelerators for data-intensive workloads.

[4] This is true in theory, although it may be some time before the scale and economics of in-memory computing are conducive to hosting all workloads on a single in-memory database.

[5] The Global Information and Technology Report 2012, edited by Soumitra Dutta, Roland Berger Professor of Business and Technology, INSEAD, and Beñat Bilbao-Osorio, Associate Director, Economist, Centre for Global Competitiveness and Performance, World Economic Forum.

[6] Source: "Real-World Business Cases for SAP HANA: Co-Innovation from Adobe and SAP," SAP SAPHIRENOW presentation, Orlando, Florida, May 14-16. http://events.sap.com/sapphirenow/en/session/2364

[7] Source for Yodobashi, Centrica, and T-Mobile examples: The Global Information and Technology Report 2012, edited by Soumitra Dutta, Roland Berger Professor of Business and Technology, INSEAD, and Beñat Bilbao-Osorio, Associate Director, Economist, Centre for Global Competitiveness and Performance, World Economic Forum. Chapter 1.7, "Harnessing the Power of Big Data in Real Time through In-Memory Technology and Analytics," by SAP AG.

[8] Source: Intel internal tests measuring performance for a 10 TB Proof of Performance and Scalability (POPS) workload. Base configuration: Server platform with 4 x Intel® Xeon® processor E7-4870 (10 core, 30 MB Cache, 2.4 GHz, 6.4 GT/s Intel® QPI), BIOS R27 (build date Marche 29, 2011), 1024 GB memory (1067 GHz DDR3 ECC), 73 GB SAS HDD local disk to hold the operating system, IBM Gen3 XIV storage attached via SAN (111 TB total raw space on 132 disks (2TB x 12 disks/pModule x 11 XIV modules), 6 x 16 TB + 1 x 15 TB ext3 file system (only 2 x 16 TB FS were used); SuSe Linux Enterprise Server v11 SP3 64-bit, IBM DB@ 10.1 FP1 GA (non DPF setup). System under Test: Server platform with 4 x Intel® Xeon® processor E7-4890 v2 (15 core, 37.5 MB Cache, 2.8 GHz, 8.00 GT/s Intel® QPI), BIOS 44.R01, 1024 GB memory (1333 GHz DDR3 ECC), 300 GB SAS HDD local disk to hold the operating system, IBM Gen3 XIV storage attached via SAN (111 TB total raw space on 132 disks (2TB x 12 disks/pModule x 11 XIV modules), 6 x 16 TB + 1 x 15 TB ext3 file system (only 2 x 16 TB FS were used) ; SuSe Linux Enterprise Server v11 SP3 64-bit, IBM DB2 10.5 FP1 with BLU (DPF not supported at the time of the test).

[9] Source: Oracle website as of January 17, 2014. http://www.oracle.com/us/corporate/features/database-in-memory-option/index.html

[10] Source: "SQL Server In-Memory OLTP Internals Overview for CTP1," a SQL Server Technical Article, by Kalen Delaney, June 2013. http://download.microsoft.com/download/F/5/0/F5096A71-3C31-4E9F-864E-A6D097A64805/SQL_Server_Hekaton_CTP1_White_Paper.pdf

[11] Source: Intel internal measurements November 2013. Configurations: Baseline 1.0x: Intel® E7505 Chipset using four Intel® Xeon® processors E7-4870 (4P/10C/20T, 2.4GHz) with 256GB DDR3-1066 memory scoring 110,061 queries per hour. Source: Intel Technical Report #1347. New Generation 2x: Intel® C606J Chipset using four Intel® Xeon® processors E7-4890 v2 (4P/15C/30T, 2.8GHz) with 512GB DDR3-1333 (running 2:1 VMSE) memory scoring 218,406 queries per hour. Source: Intel Technical Report #1347.

[12] Source: Intel internal tests measuring performance for a 10 TB Proof of Performance and Scalability (POPS) workload running on IBM DB2 10.1 and IBM DB2 10.5 with BLU Acceleration. Platform configuration: Server platform with 4 x Intel Xeon processor E7-4870 (10 core, 30MB Cache, 2.4 GHz, 6.4 GT/s Intel® QPI), BIOS R27 (build date Marche 29, 2011), 1024 GB memory (1067 GHz DDR3 ECC), 73 GB SAS HDD local disk to hold the operating system, IBM Gen3 XIV storage attached via SAN (111 TB total raw space on 132 disks (2 TB x 12 disks/pModule x 11 XIV modules), 6 x 16 TB + 1 x 15 TB ext3 file system (only 2 x 16 TB FS were used); SuSe Linux Enterprise Server v11 SP3 64-bit, IBM DB@ 10.1 FP1 GA (non DPF setup and IBM DB2 10.5 FP1 with BLU (DPF not supported at the time of the test).

[13] On a 4-socket natively-connected platform: Intel® Xeon® processor E7 family supports 64DIMMS, max memory per DIMM of 32 GB RDIMM; Intel® Xeon® processor E7 v2 family supports 96DIMMs, max memory per DIMM of 64 GB RDIMM. This enables a 3x increase in memory.

[14] Up to 4x I/O bandwidth claim based on Intel internal estimates of the Intel® Xeon® processor E7-4890 v2 performance normalized against the improvements over dual-IOH Intel® Xeon® processor E7-4870 using an Intel internal bandwidth tool running the 1R1W test.

[15] Source for Yodobashi, Centrica, and T-Mobile examples: The Global Information and Technology Report 2012, edited by Soumitra Dutta, Roland Berger Professor of Business and Technology, INSEAD, and Beñat Bilbao-Osorio, Associate Director, Economist, Centre for Global Competitiveness and Performance, World Economic Forum. Chapter 1.7, "Harnessing the Power of Big Data in Real Time through In-Memory Technology and Analytics," by SAP AG.

[16] No computer system can provide absolute reliability, availability or serviceability. Requires an Intel® Run Sure Technology-enabled system, including an enabled Intel processor and enabled technology(ies). Built-in reliability features available on select Intel® processors may require additional software, hardware, services and/or an Internet connection. Results may vary depending upon configuration. Consult your system manufacturer for more details.

Printed in USA 0214/RS/OCG/XX/PDF Please Recycle 330164-001US