ORACLE DATABASE 12c

Benchmark Results Reveal the Benefits of Oracle Database In-Memory for SAP Applications

ORACLE WHITE PAPER | SEPTEMBER 2015

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Table of Contents

Preface

For many years, SAP has published a series of standard benchmarks that allow database vendors to demonstrate the performance of their products for running SAP Business Applications. These benchmarks measure the performance of databases running SAP Business Applications like Sales and Distribution (SD), Assemble to Order (ATO), etc. By certifying results from different hardware and database vendors, SAP provides their customers with the ability to compare the performance and scalability of different platforms or databases. The benchmarks also provide guidance on how to configure and size SAP Business Suite.

Oracle and other database vendors have been running these benchmarks for many years. We are very proud of the performance of the Oracle Database running SAP business Applications, and you can see our benchmark results published on SAP.com.

Interestingly, SAP has chosen not to publish results for any of their standard OLTP benchmarks for their HANA database; instead, SAP has created a new benchmark called SAP Business Warehouse Enhanced Mixed Workload (BW-EML). This benchmark measures the performance of a database running the SAP BW analytical application. SAP only publishes BW-EML benchmark results for HANA.

Apparently, SAP does not want to disclose the performance of HANA on any of the other SAP Business Applications except for the BW application, making it impossible to tell how HANA compares to other databases running the SAP Business Applications benchmarks unless those databases run the SAP BW-EML benchmark.

This paper compares Oracle's results on the SAP BW-EML benchmark to those achieved by HANA.

Introduction

The BW Enhanced Mixed Load Benchmark (BW-EML Benchmark) was designed to measure the current demands of typical business warehouse customers, using a combination of real-time report data loads and ad-hoc queries.

This whitepaper provides an overview of the SAP BW-EML benchmark, details on the benchmark results achieved by Oracle Database 12*c*, and a description of the unique database technology used to achieve these results.

SAP BW-EML Benchmark

The BW Enhanced Mixed Load Benchmark[1] (BW-EML Benchmark) was designed to measure the current demands on a data warehouse and has two main requirements:

Demonstrate near real-time reporting – Getting instant results from analytical applications on up to the minute data is crucial for timely decision making.

Ad-hoc reporting capabilities – Provide the ability to ask any question on any data without having predefined aggregated results.

The benchmark simulates a large number of online users running reports for a 60-minute interval. During this 60-minute window new data is "trickle loaded" into the database every 5 minutes, resulting in an additional 0.1% of the total data volume over 60 minutes. The benchmark result is the total number of query navigations or decision steps that can be completed over 60 minutes.

Benchmark Schema

The database schema used in the benchmark is based on three InfoCubes and seven DataStore objects (DSO). Each DataStore contains one year's worth of data (total of 7 years of DSO) and the three InfoCubes hold the same data as the corresponding DataStore objects for the last three years. It is effectively an "Extended Star Schema", that is to say, each query joins one fact table to one or more dimension tables, and those dimension tables will join to Surrogate ID tables (SID tables).

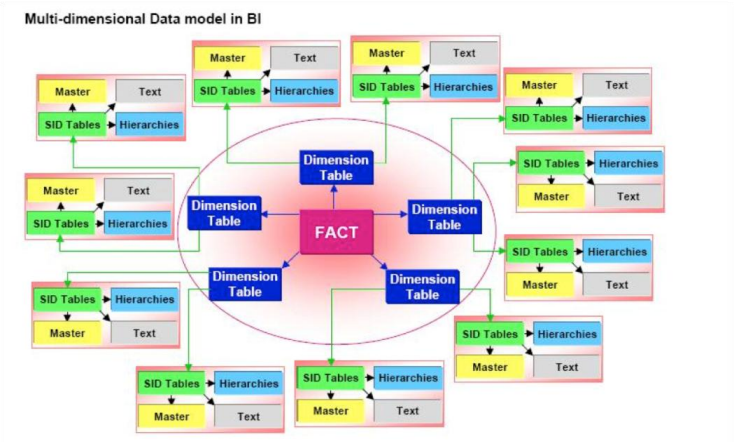

Figure 1. SAP BW-EML benchmark schema[2]

Benchmark Scale Factors

As with all standard benchmarks, the SAP BW-EML benchmark can be run at different scale factors. A scale factor represents the total number of records or rows stored in the schema, and in the case of the BW-EML benchmark the scale factors range from 500 million records to 10 billion records. To give you some idea of how much raw data that actually is, the 1 billion records scale-factor generates approximately 600GB of flat files. In this white paper we will

[1] More details on the BW-EML Benchmark can be found on sap.com
[2] More information on SAP Data Modeling can be found in their whitepaper Multi-Dimensional Modeling with BI

focus on the 1 billion record scale factor and examine that workload both on a single database server and in a scaled out solution.

Benchmark Workload

The queries are based on 8 web reports (4 on the DSO MultiProvider and 4 on the InfoCube MultiProvider). Each query is randomized both by characteristic and filter value and drill-down.

Throughout the course of the benchmark new data is loaded into the database every 5 minutes, resulting in an additional 0.1% of the total data volume over 60 minutes. That is 1 million additional records at the1 billion records scale-factor, or 100,000 new rows added to each of the ten objects.

Oracle Database 12*c* SAP BW-EML Benchmark

Motivation

In July 2014, Oracle Database 12*c* Release 1 (12.1.0.2) was released and it included a new option called Oracle Database In-Memory (Database In-Memory), which accelerates analytics by orders of magnitude while simultaneously speeding up mixed-workload OLTP. This Oracle database release including Oracle Database In-Memory has been certified by SAP for use with SAP® solutions based on the SAP NetWeaver® 7.x technology platform.

Database In-Memory leverages a unique "dual-format" architecture that enables tables to be in memory simultaneously in a traditional row format and a new in-memory column format. The Oracle SQL Optimizer automatically routes analytic queries to the column format and OLTP queries to the row format, transparently delivering best-of-both-worlds performance. Oracle Database 12*c* automatically maintains full transactional consistency between the row and the column formats, just as it maintains consistency between tables and indexes today.

There are four main reasons why accessing data in the In-Memory column format is more efficient for analytical queries:

1. Access only the columns that are needed

Analytic queries typically reference only a small number of the columns in a table. Database In-Memory minimizes work and maximizes performance by accessing only those columns needed by a query.

2. Scan and filter data in a compressed format

When data is populated into memory in the new column format it is automatically compressed using a new set of compression algorithms that allow `WHERE` clause predicates to be applied against the compressed formats. This means the volume of data scanned in the In-Memory column format for a query will be far less.

3. Prune out any unnecessary data within each column

When tables are populated into memory in the column format they are logically split into sections, and the minimum and maximum value of every column is maintained for every section of the table. This allows queries to quickly skip table sections that only contain data outside of the range of data needed by the query.

For table sections where the value we are looking for does fall within the minimum and maximum range, an additional level of data pruning is possible via the metadata dictionary created when dictionary-based compression is used. The dictionary contains a list of the unique column values within the table section. By comparing the value

in the where clause predicate to the dictionary we can easily determine if the value we are looking for is one of the distinct column values or not.

4. Use SIMD to apply filter predicates

Modern microprocessors support SIMD (Single Instruction for Multiple Data values) vector processing instructions to accelerate graphics and scientific computing. Oracle Database In-Memory can use these SIMD vector instructions to process sets of column values in a single CPU instruction. SIMD vector processing enables the Oracle Database In-Memory to scan billion of rows per second per core versus the millions of rows per second per core scan rate that can be achieved if the data is in row format.

In order to demonstrate the performance benefits of Database In-Memory and to enable SAP customers to size their systems when using Database In-Memory, Oracle decided to run the SAP BW-EML benchmark at the 1 billion records scale factor both on a single database server and in a scaled out configuration (multiple database servers).

Hardware Configuration

The hardware configuration used for the SAP BW-EML benchmark consisted of three-tiers; a driver (which simulates the users), an application tier and a database tier.

An Oracle Exalogic X4-2 server was used as the application tier. Each application server had 2 twelve-core Intel E5-2697 v2 processors (2.70GHz) and 256GB of memory.

The Database tier was an Oracle Exadata X5-2, which contains 8 database servers. Each database server consists of 2 eighteen-core Intel Xeon® E5-2699 v3 processors (2.3 GHz) and 512GB of memory.

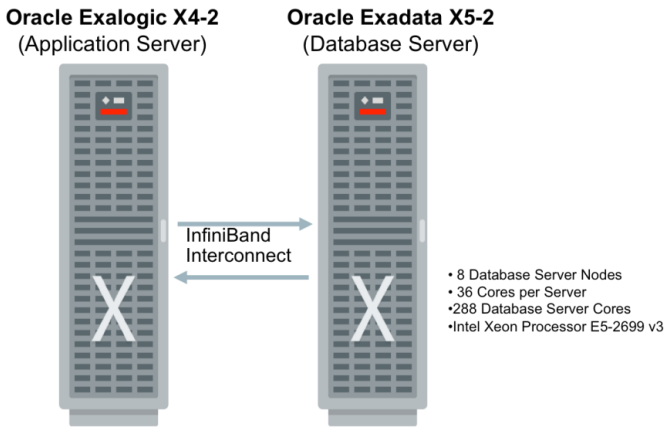

Figure 2. Oracle Database 12*c* SAP BW-EML benchmark configuration

Database Configuration

An In-Memory column store was allocated on each of the database servers, and the necessary data for the benchmark was populated into memory in the new In-Memory column format. Remember the data populated into memory is compressed, so the In-Memory column store doesn't need to be as large as the on the disk footprint of the data.

Results

1 Billion Records on a Single Database Server

The initial benchmark test was run using just one of the database servers in the Exadata X5-2 with 36 cores. Oracle Database 12*c* was able to complete over 300 thousand navigation steps in one hour. That's 84 analytic queries per second over the course of the benchmark run.

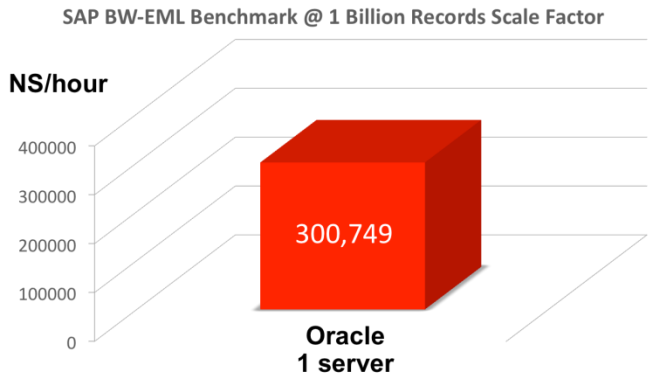

Figure 3. Oracle Database 12*c* achieves over 300 thousand navigation steps per hour on a single database server

1 Billion Records on Multiple Database Servers

The same benchmark was then repeated using 2 database servers, then 4 and finally all 8-database servers on the Exadata X5-2. Oracle Database 12*c* was able to show close to perfect scalability as the number of database servers was increased, achieving over 2.3 million navigation steps per hour with 8 database servers. That's an incredible 650 analytic queries/second.

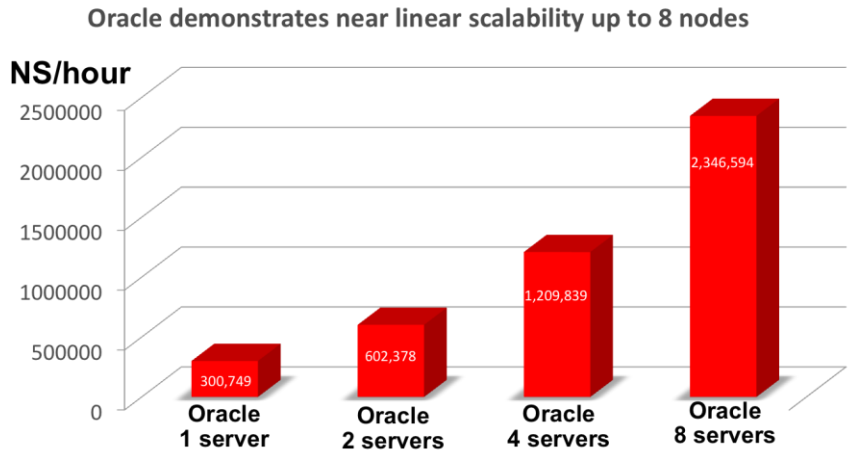

Figure 4. Oracle Database 12*c* shows linear scalability & achieves over 2.3 million navigation steps per hour on a 8 database servers

Comparison With Existing SAP BW-EML Benchmark Results

One of the goals of SAP Application Benchmarks is to provide SAP customers with the ability to compare the performance and scalability of different platforms or databases. Currently SAP HANA is the only database with published results for the SAP BW-EML benchmark at the 1 billion records scale factor[3]. Below is a comparison of the Oracle Database 12*c* results versus the best SAP HANA results, both for a single database server and for scale-out solutions.

1 Billion Records on a Single Database Server

The Oracle single server benchmark was conducted on one of the Exadata X5-2 database servers, which has 2 Intel processors (Intel Xeon Processor E5-2699), 36 cores, and 512GB of main memory. In October 2014, SAP HANA results were released on a 2-socket Dell server using the same Intel processor, with the same number of cores and the same amount of memory.

As you can see from the graph below Oracle Database 12*c* out performed SAP HANA by a factor of 2X. Oracle was able to complete over twice the number of navigation steps in 1 hour than SAP HANA on the same chipset.

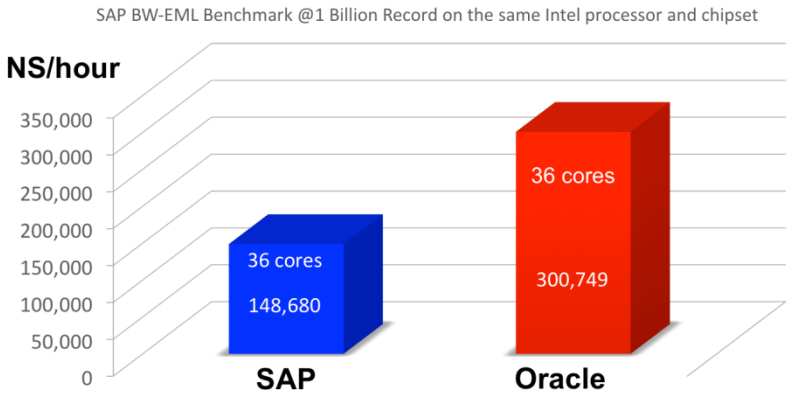

Figure 5. Oracle Beats SAP HANA by 2X at the 1 billion records scale factor on the same Intel processor & chipset.

In May 2015, a new SAP HANA benchmark result was released at the 1 billion records scale factor for a single database server. This result was again achieved using a Dell server but this time it had 4 Intel processors (Intel Xeon Processor E7-8890) with 72 cores and 1536 GB main memory.

As you can see from the graph below, on this much larger hardware platform SAP HANA narrowly beat the Oracle result achieving 320,940 navigation steps, but they needed twice the number of cores and 3 times the amount of memory to do it.

[3] The results of all SAP BW-EML benchmarks are available on SAP.com

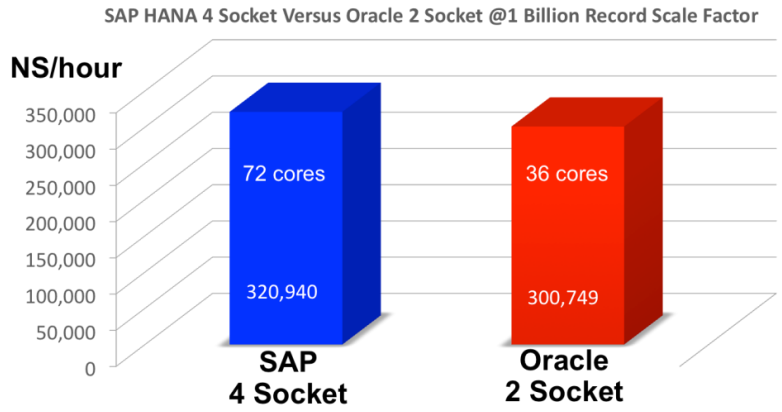

Figure 6. SAP HANA requires 2X the CPU cores and 3X the memory to narrowly beat the Oracle on a single Database server.

1 Billion Records Scale Out

Again SAP HANA is the only database vendor with published results for the SAP BW-EML benchmark at the 1 billion records scale factor with a scale out configuration but unlike Oracle they have not demonstrated linear scalability by releasing results for 2, 4, and 8 server configurations. SAP HANA has only released scale out results for a 7 server configuration.

SAP HANA's best 7-server scale out result was published in May 2015 using 7 Lenovo 3850 X6 servers, each with 4 processors, 72 cores (Intel Xeon Processor E7-8890 v3) and 1024GB of main memory.

With almost twice the number of cores (504 cores versus 288 cores) and double the amount of memory, SAP HANA still couldn't achieve the same number of navigation steps per hour that Oracle delivered on the 8-server configuration.

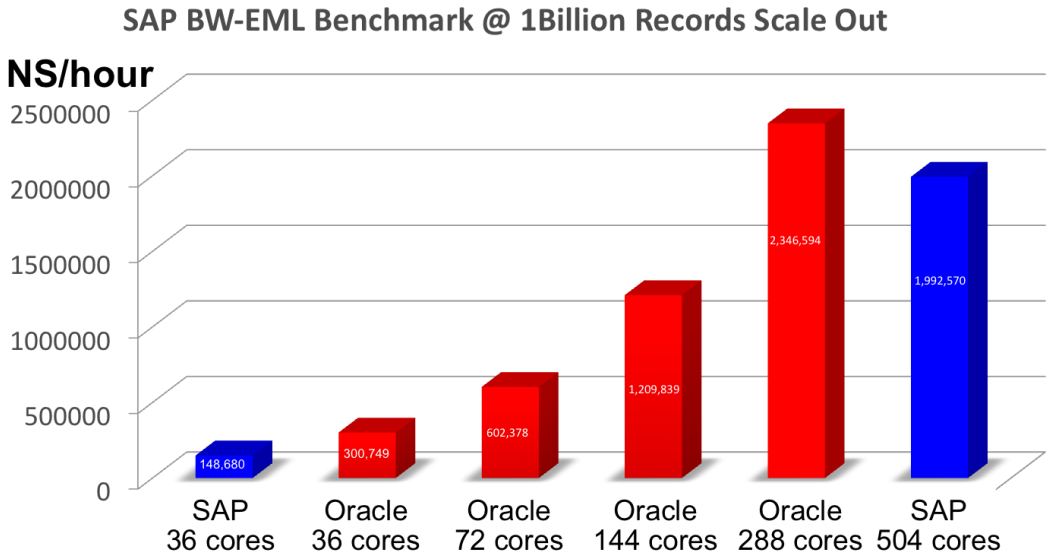

Figure 7. SAP HANA requires 4-socket servers with 2X the cores compared to Oracle's 8 2-socket servers and it still doesn't achieve the same result.

Conclusion

The two key benefits of a standard application benchmark is the ability to compare different platforms and databases in terms of performance and scale and to provide guidance on how to configure and size an SAP system.

Using the SAP BW-EML benchmark we have proved Oracle Database 12*c* with the Oracle Database In-Memory option outperforms SAP HANA by a factor of 2X on the same Intel processor for SAP workloads. Oracle Database 12*c* also shows near linear scalability as the number of database servers is increased and achieved more navigation steps per hour than SAP HANA on half the CPUs.

By sharing these benchmark results, SAP customers now have a guide to help them configure and size the best database platform for their SAP applications.

Appendix A - Detailed Benchmark Results

Oracle Database 12*c*

Date of Submission to SAP for Review mm/dd/yyyy	Technology Partner	Number of records	Ad-Hoc Navigation Steps/Hour	Operating System - Release Database Server	Database Release	SAP NetWeaver Release	Number & Type of Database/Application Servers
04/20/2015	Oracle	1,000,000,000	300749	Oracle Enterprise Linux 6	Oracle 12*c*	SAP NetWeaver 7.40	1 database server: Oracle Exadata Database Machine X5-2 Full Rack EF (1 compute node up) 2 processor / 36 cores / 72 threads, Intel Xeon Processor E5-2699 v3 2.30 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 45 MB L3 cache per processor, 512 GB main memory 4 application server: Oracle Exalogic Elastic Cloud X4-2 Exalogic Elastic Cloud Quarter Rack (4 compute nodes up) 4 x 2 processor / 4 x 24 cores / 4 x 48 threads, Intel Xeon Processor E5-2697, 2.70 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 4 x 256 GB main memory
04/20/2015	Oracle	1,000,000,000	602378	Oracle Enterprise Linux 6	Oracle 12*c*	SAP NetWeaver 7.40	2 database server: Oracle Exadata Database Machine X5-2 Full Rack EF (2 compute node up) 2 x 2 processor / 2 x 36 cores / 2 x 72 threads, Intel Xeon Processor E5-2699 v3 2.30 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 45 MB L3 cache per processor, 512 GB main memory 8 application server: Oracle Exalogic Elastic Cloud X4-2 Exalogic Elastic Cloud Quarter Rack (8 compute nodes up) 8 x 2 processor / 8 x 24 cores / 8 x 48 threads, Intel Xeon Processor E5-2697, 2.70 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 8 x 256 GB main memory
04/20/2015	Oracle	1,000,000,000	1209839	Oracle Enterprise Linux 6	Oracle 12*c*	SAP NetWeaver 7.40	4 database server: Oracle Exadata Database Machine X5-2 Full Rack EF (4 compute node up) 4 x 2 processor / 4 x 36 cores / 4 x 72 threads, Intel Xeon

							Processor E5-2699 v3 2.30 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 45 MB L3 cache per processor, 512 GB main memory 24 application server: Oracle Exalogic Elastic Cloud X4-2 Exalogic Elastic Cloud Full Rack (24 compute nodes up) 4 x 2 processor / 4 x 24 cores / 4 x 48 threads, Intel Xeon Processor E5-2697, 2.70 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 24 x 256 GB main memory
09/02/2015	Oracle	1,000,000,000	2346594	Oracle Enterprise Linux 6	Oracle 12*c*	SAP NetWeaver 7.40	8 database server: Oracle Exadata Database Machine X5-2 Full Rack EF (8 compute node up) 8 x 2 processor / 8 x 36 cores / 8 x 72 threads, Intel Xeon Processor E5-2699 v3 2.30 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 45 MB L3 cache per processor, 512 GB main memory 32 application server: Oracle Exalogic Elastic Cloud X4-2 Exalogic Elastic Cloud Full and Eighth Rack (32 compute nodes up) 32 x 2 processor / 32 x 24 cores / 32 x 48 threads, Intel Xeon Processor E5-2697, 2.70 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 32 x 256 GB main memory

SAP HANA

Date of Certification mm/dd/yyyy	Technology Partner	Number of records	Ad-Hoc Navigation Steps/Hour	Operating System - Release Database Server	Database Release	SAP NetWeaver Release	Number & Type of Database/Application Servers	Certification Number
10/20/2014	Dell	1,000,000,000	148,680	SuSE Linux Enterprise Server 11	SAP HANA 1.0	SAP NetWeaver 7.31	1 database server: Dell PowerEdge R730, 2 processors / 36 cores / 72 threads, Intel Xeon Processor E5-2699 v3, 2.30 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 45 MB L3 cache per processor, 512 GB main memory 2 application servers: Dell PowerEdge R910, 4 processors / 40 cores / 80 threads, Intel Xeon Processor E7-4870, 2.40 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 512 GB main memory	2014038
05/05/2015	Dell	1,000,000,000	320,940	SuSE Linux Enterprise Server 11	SAP HANA 1.0	SAP NetWeaver 7.31	1 database server: Dell PowerEdge R930, 4 processors / 72 cores / 144 threads, Intel Xeon Processor E7-8890 v3, 2.50 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 45 MB L3 cache per processor, 1536 GB main memory 2 application servers: Dell PowerEdge R910, 4 processor / 40 cores / 80 threads, Intel Xeon Processor E7-4870, 2.40 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 30 MB L3 cache per processor, 512 GB main memory	2015015
05/05/2015	Lenovo	1,000,000,000	1,992,570	SuSE Linux Enterprise Server 11	SAP HANA 1.0	SAP NetWeaver 7.40	7 x database servers: Lenovo x3850 X6 , 4 processor / 72 cores / 144 threads, Intel Xeon Processor E7-8890 v3, 2.50 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 45 MB L3 cache per processor, 1024 GB main memory 15 x IBM Flex System x880 X6 Compute Node, 2 processor / 30 cores / 60 threads, Intel Xeon Processor E7-8890 v2, 2.80 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 37.5 MB L3 cache per processor, 256 GB main memory	2015011

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

Integrated Cloud Applications & Platform Services

Whitepaper title: Benchmark Results Reveal the Benefits of Oracle Database In-Memory for SAP Applications

CONNECT WITH US

blogs.oracle.com/in-memory
facebook.com/oracle
twitter.com/OracleDatabase