Evaluating the Performance of 64-bit Oracle9*i* Database Release 2 on the Intel EM64T-based Platform

Intel® Extended Memory 64 Technology (EM64T) is designed to allow the 64-bit register extensions available in Intel Xeon™ processors to enable a larger flat memory address space, compared to 32-bit processors. This article describes testing conducted by Dell engineers to evaluate and compare the performance of database workloads on a Dell™ PowerEdge™ 2850 server running in legacy 32-bit and extended 64-bit modes.

BY RAMESH RADHAKRISHNAN, PH.D.

Eighth-generation Dell PowerEdge servers incorporate Intel Xeon processors (previously code-named Nocona) that support Intel Extended Memory 64 Technology (EM64T).[1] EM64T extends the 32-bit Intel architecture by allowing the Intel Xeon processor to run 64-bit code and address larger amounts of memory than a 32-bit processor running in 32-bit mode. This feature can be beneficial for database applications that require large contiguous blocks of memory to cache data and support large numbers of concurrent users. To evaluate the impact of EM64T on database performance, Dell engineers compared the performance of 32-bit and 64-bit modes by running database workloads on a Dell PowerEdge 2850 server configured with EM64T-capable processors.

Understanding 64-bit and 32-bit memory access

A 64-bit processor can theoretically address up to 18 exabytes[2] of contiguous memory, as compared to 4 GB addressable by a 32-bit processor. The EM64T-enabled processors are designed to provide extended memory addressing that can allow 64-bit computers to run larger computational problems and process more applications concurrently, as well as to reduce I/O latency by keeping additional data in physical memory, compared to 32-bit computers. The Intel Xeon processor with EM64T, when used with the Intel E7520 chip set, supports up to 12 GB of RAM on the eighth-generation PowerEdge platforms. Because the operating system sees all of this memory as flat memory space, 64-bit applications are not limited by the 4 GB barrier imposed by a 32-bit architecture.

When running on x86 processors, 32-bit applications cannot address more than 4 GB of memory and, in general practice, cannot use more than 2 GB or 3 GB of memory without special extensions. This limitation means that 32-bit systems cannot hold large databases in memory or support large numbers of users without

[1] For more information about the eighth-generation Dell server platforms, visit http://www.dell.com/servers.

[2] One exabyte = 2^{60} (1,152,921,504,606,846,976) bytes, or 1,024 petabytes.

Reprinted from *Dell Power Solutions,* October 2004. Copyright © 2004 Dell Inc. All rights reserved.

	64-bit configuration	32-bit configuration
Operating system	Red Hat Enterprise Linux AS 3 (64 bits)	Red Hat Enterprise Linux AS 3 (32 bits)
Database server	Oracle9*i* Database Release 2 Enterprise Edition Release Candidate 2 (RC2) for Linux x86-64	Oracle9*i* Database Release 2 Enterprise Edition RC2 for Linux x86
Server platform	Dell PowerEdge 2850	Dell PowerEdge 2850
CPU	Dual Intel Xeon processors at 3.6 GHz with 1 MB L2 cache	Dual Intel Xeon processors at 3.6 GHz with 1 MB L2 cache
Memory	12 GB DDR2 at 400 MHz	12 GB DDR2 at 400 MHz
Internal disk	One 37 GB	One 37 GB
Network interface card	Two 1000 Mbps* (internal)	Two 1000 Mbps (internal)
Disk controller	PowerEdge Expandable RAID Controller 4, Dual Channel (PERC 4/DC)	PERC 4/DC
Disk enclosure	PowerVault 220S	PowerVault 220S
External disks	Fourteen 37 GB Ultra320 SCSI	Fourteen 37 GB Ultra320 SCSI

*This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

Figure 1. Comparison of 64-bit and 32-bit configurations

resorting to segmentation or memory page swapping to the slower disk subsystem. When used with a 64-bit OS and applications, Intel EM64T helps overcome this limitation, enabling larger memory addressability by removing the 4 GB barrier.

Establishing the test environment

To determine the effect of enabling EM64T, in June 2004 Dell engineers examined the performance of database workloads running on a Dell PowerEdge 2850 server configured with dual EM64T-capable Intel Xeon processors at 3.6 GHz and 12 GB of double data rate 2 (DDR2) memory at 400 MHz. The tests were performed on the PowerEdge 2850 server using the processors in 32-bit mode and in 64-bit mode with a 32-bit application and a 64-bit application, respectively. Because the workloads were identical in both test cases, the difference in performance can be attributed to EM64T.

To properly evaluate the performance advantage of Intel EM64T, the test team kept differences in hardware and software for the 32-bit and 64-bit configurations to a minimum (see Figure 1). A 32-bit configuration was obtained by installing a 32-bit version of the Red Hat® Enterprise Linux® AS 3 operating system, which forces the processors to run in 32-bit mode, also known as Legacy mode. In this configuration, Dell also installed a 32-bit version of Oracle9*i*™ Database Release 2 on the PowerEdge 2850. Similarly, 64-bit versions of the operating system and Oracle9*i* database server were installed on the PowerEdge 2850 and were used to test the 64-bit configuration.

Both the 32-bit and 64-bit systems were connected to a Dell PowerVault™ 220S storage subsystem with several hundred gigabytes of storage to host the database. Identical large databases (approximately 100 GB) were built for the 32-bit and 64-bit systems. The database was modeled after an online DVD store, and applications that simulated online order entry and report generation were run against the database.

Various workloads were run against both 32-bit and 64-bit configurations to determine the scaling behavior of both systems. The database cache size was varied from small to large (up to 75 percent of the installed system memory) to evaluate the advantages of the flat memory addressing capability for the database workloads.

Database server setup

Oracle9*i* Database Release 2 for Linux x86 and Linux x86-64 was used on the two configurations as the database server. These releases were downloaded from http://www.oracle.com, and the Enterprise Edition was installed in both configurations.

The database tablespaces for both configurations were identical (see Figure 2). To obtain the best performance, raw devices were used instead of a file system to store the database files. One ten-drive, RAID-0 logical unit number (LUN) was mounted on /u01 to hold the data, index, undo, and temporary tablespaces. A RAID-0 LUN was mounted on /u05 to hold the log files. RAID-0 was used to reduce any potential read/write latencies to the disk subsystem and to allow for better memory addressability and scaling at the higher workloads.

Workload simulation

The Oracle9*i* Database Release 2 Client for Windows 98/NT/2000/XP was installed on the client machine. Driver applications were then run on the client machine to generate workload requests to the database server.

The online transaction processing (OLTP) workload modeled the back end of a large online DVD store database. The database consisted of a set of data tables organized according to a schema, in addition to a set of stored procedures. The stored procedures managed the database based on requests generated by the driver applications.

Tablespace	Description	Location	Space used/available
CUSTTBS	Customer table	/u01	34 GB/45 GB
INDXTBS	Index table	/u01	24 GB/30 GB
ORDERTBS	Order and order detail tables	/u01	15 GB/20 GB
DBMISC	Miscellaneous	/u01	0.5 GB/1 GB
UNDOTBS	Undo tablespace	/u01	0.4 GB/1 GB
TEMP	Temporary table	/u01	18 GB/30 GB

Figure 2. Database tablespaces

Separate multithreaded driver programs were used to model order-entry (OLTP) and report-request workloads. For the OLTP driver, each thread of the driver application connected to the database and modeled a series of customers logging in, browsing, and placing orders. The driver measured the transaction rates and the average time for completing the transactions. The number of threads was varied during the tests to place different load levels on the server. Each thread in the report-request driver connected to the database and called a stored procedure that calculated the total sales for a DVD category for the previous month, quarter, and half-year. Numerous report requests were made to place different loads on the server. In one test scenario, OLTP and report requests were run in parallel to model typical enterprise database usage.

> Eighth-generation Dell PowerEdge servers incorporate Intel Xeon processors that support Intel Extended Memory 64 Technology. EM64T extends the 32-bit Intel architecture by allowing the Intel Xeon processor to run 64-bit code and address larger amounts of memory than a 32-bit processor.

Examining the effect of Intel EM64T on database workloads

To test the performance scaling of the database workloads at different levels of system load, OLTP and report requests were made on both the 64-bit and 32-bit systems. For all database workload tests, the Dell test team compared the performance of both systems configured with a range of database cache sizes, from the default size of 350 MB to 9 GB. The database cache buffer is part of the system global area (SGA), and increasing cache size typically allows the database to cache more data.

To support database cache sizes greater than 3 GB on the 32-bit system, very large memory (VLM) was enabled on the Linux operating system. VLM allows an application to address memory outside the 4 GB boundary. The extended buffer cache functionality available in Oracle9*i* Database Release 2 was then used to enable database cache sizes from 3 GB to 9 GB. In VLM mode, the SGA is split into two pieces. The buffer cache is moved outside the normal address space, so the 32-bit address space restriction does not apply. Everything else in the SGA (shared pool, log buffer, variable SGA, and so forth) remains in the 4 GB address space.

Enabling large database caches on the 64-bit system was more straightforward. Because the 64-bit architecture does not have a 4 GB limit per process, the database cache size can be increased without using VLM or splitting the SGA into two pieces. The 64-bit system incurs no overhead for memory access over the 4 GB limit that exists for 32-bit systems.

Throughput comparison for a medium OLTP workload

Figure 3 shows the OLTP throughput (in orders per minute) that the 64-bit and 32-bit configurations sustained when servicing five simultaneous OLTP connections (a medium workload). The 32-bit system completed more orders per minute for database cache sizes ranging from 350 MB to 3 GB. For larger database cache sizes (over 6 GB), the 64-bit configuration performed better, which indicated that it took advantage of the flat memory addressing capability of Intel EM64T.

Performance scaling for a medium OLTP workload

Figures 4 and 5 show the improvements in OLTP throughput for the medium workload (five OLTP connections that resulted

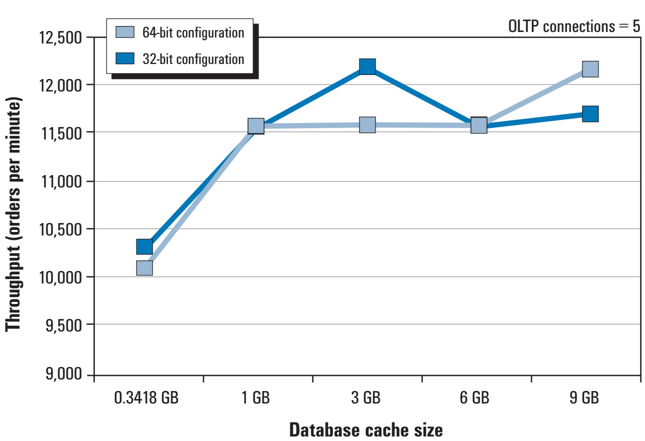

Figure 3. Performance for a medium OLTP workload on the 32-bit and 64-bit configurations

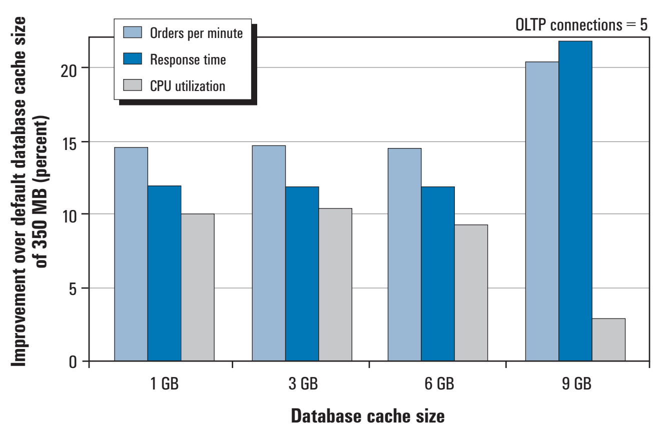

Figure 4. Performance improvements for a medium OLTP workload on the 64-bit configuration

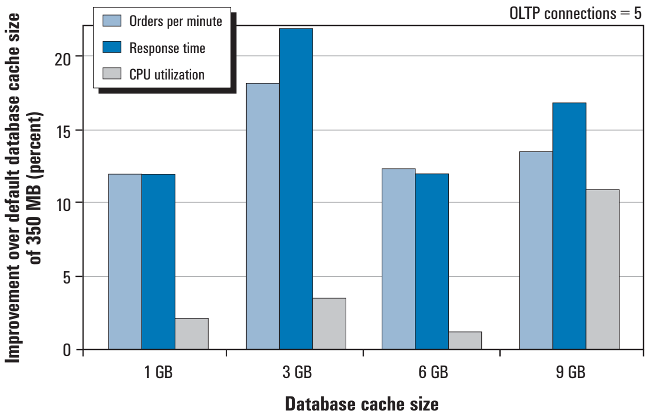

Figure 5. Performance improvements for a medium OLTP workload on the 32-bit configuration

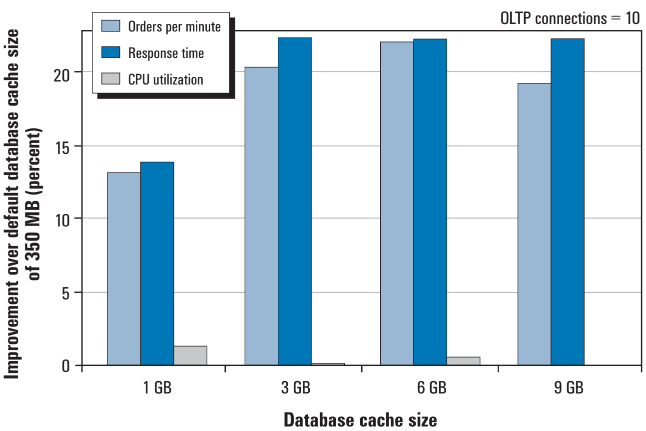

Figure 6. Performance improvements for a heavy OLTP workload on the 64-bit configuration

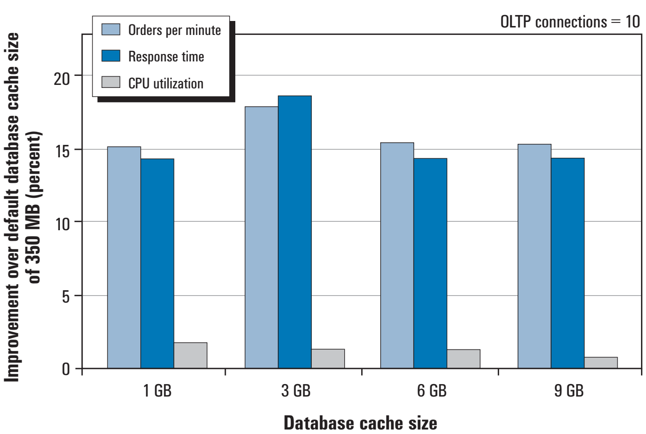

Figure 7. Performance improvements for a heavy OLTP workload on the 32-bit configuration

in 55 percent to 60 percent CPU utilization for all database cache sizes) as compared to the default 350 MB database cache size. As the database cache sizes increased, the 64-bit system scaled better than the 32-bit system because EM64T enabled the 64-bit system to address a larger flat memory space.

A 64-bit processor can theoretically address up to 18 exabytes of contiguous memory, as compared to 4 GB addressable by a 32-bit processor.

In terms of orders per minute, the 64-bit system, which had EM64T enabled, scaled 20 percent better at a database cache size of 6 GB and 50 percent better at a database cache size of 9 GB compared to the 32-bit system, which incurred additional overhead to address memory over 4 GB.

Performance scaling for a heavy OLTP workload

Figures 6 and 7 show the improvements in OLTP throughput for a heavy workload (10 OLTP connections resulting in 92 percent to 95 percent CPU utilization for all database cache sizes). Similar to the medium OLTP workload test, as database cache sizes increased, the 64-bit system under a heavy workload scaled better than the 32-bit system.

On the 64-bit system, as the database cache size increased from 6 GB to 9 GB, the performance (in terms of orders per minute) degraded by 15 percent. On the 32-bit system, the performance remained flat. This effect was observed because the heavy workload required more memory than the medium workload to support a higher number of simulated users.

For the heavy OLTP workload, the 64-bit system using EM64T scaled 44 percent better (for the 6 GB database cache size) and 25 percent better (for the 9 GB database cache size) in terms of orders per minute, compared to the 32-bit system. A higher improvement for the response times in the 64-bit system occurred when the database cache sizes were increased. On the 32-bit system, increasing the database cache size more than 3 GB did not provide any additional benefits in response time and throughput for the heavy workload.

Performance scaling for a report-generation workload

The Dell test team set the client to make five simultaneous report requests to the database server, to create reports for five movie categories. The average completion time (amount of time for the report generation to finish) was approximately 30 minutes, and the average CPU utilization was between 85 percent and 90 percent for all database cache sizes.

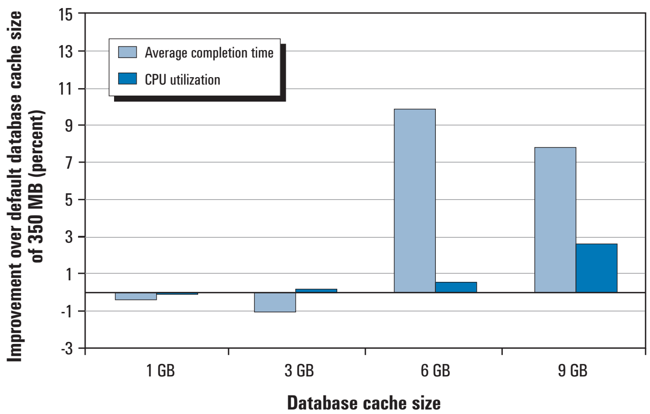

Figure 8. Performance improvements for a report-generation workload on the 64-bit configuration

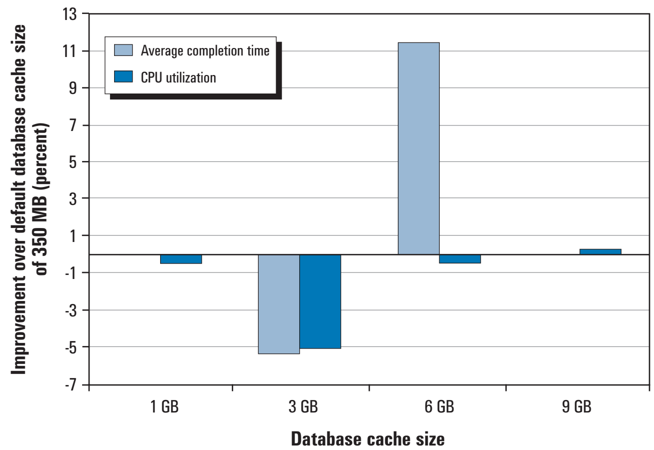

Figure 9. Performance improvements for a report-generation workload on the 32-bit configuration

Figure 8 shows the effect of performance on a 64-bit system when the database cache size for the report-creation workload was increased. Increasing the database cache size to 1 GB and 3 GB did not provide any benefit for this workload. Going to larger sizes of 6 GB and 9 GB provided an 8 percent to 10 percent improvement in report completion time for this workload.

Figure 9 shows the performance improvements over the default database cache size for the 32-bit configuration. Database cache sizes of 1 GB and 3 GB did not improve performance, but a database cache size of 6 GB reduced (improved) the report-completion time by 11.5 percent. Increasing the database cache size to 9 GB did not help performance.

OLTP performance on the 64-bit configuration was observed to scale 25 percent to 50 percent better for a database cache size of 9 GB in both medium and heavy workloads, compared to the 32-bit configuration.

Helping to improve the scalability of database workloads

As indicated by the Dell tests conducted for this study, the PowerEdge 2850 server using processors enabled with Intel EM64T can host database applications that require large contiguous blocks of memory and support large numbers of concurrent users. OLTP performance on the 64-bit configuration was observed to scale 25 percent to 50 percent better for a database cache size of 9 GB in both medium and heavy workloads, compared to the 32-bit configuration. Similarly, for the 64-bit configuration, the report-generation workload scaled better for larger database cache sizes, such as 9 GB, than the 32-bit configuration.

Acknowledgments

The author would like to thank Dave Jaffe and Todd Muirhead of the Dell Technology Showcase team for developing and sharing the DVD store database tests.

Paul Rad of the Dell Database and Applications Engineering Department helped solve Oracle® installation problems and provided assistance with enabling large database cache sizes for 32-bit systems. Van Okamura of Oracle provided assistance with SGA issues on both the 32-bit and 64-bit configurations. Nicholas Wakou and Philip Wong of the Dell Server and Storage Performance team helped with general issues regarding Oracle.

Ramesh Radhakrishnan, Ph.D., is an engineer on the Dell High-Performance Computing Cluster team. Ramesh has a Ph.D. in Computer Engineering from The University of Texas at Austin.

FOR MORE INFORMATION

Dell PowerEdge 2850:
http://www1.us.dell.com/content/products/productdetails.aspx/pedge_2850?c=us&l=en&s=gen

Intel Extended Memory 64 Technology:
http://www.intel.com/technology/64bitextensions