
Tune Right – Run Fast®

Comparison of 32-bit and 64-bit Oracle Database Performance on the Dell PowerEdge 6850 Server with Microsoft Windows Server 2003

By Larry Pedigo, Performance Tuning Corporation

Executive Summary

In the spring of 2005, Dell is introducing the PowerEdge 6850 server, an Enterprise-class server scalable to four processors. The Dell PowerEdge 6850 server is an update of the popular 6650 server, featuring significant new technological advances. One of the most significant new features is the adoption of the Intel EM64T architecture. With the EM64T architecture, the PowerEdge 6850 server is capable of fully native operation in either 32-bit mode or 64-bit mode, without performance penalties for “emulation” modes. Upon release, the Dell PowerEdge 6850 server will be available with several 32-bit OS options, including Microsoft Windows Server 2003. Later in 2005, the 64-bit version of Microsoft Windows Server 2003 for the EM64T processor will be released. The PowerEdge 6850 server will offer an ideal platform for migrating from 32-bit Windows to 64-bit Windows.

The Dell PowerEdge 6850 server will offer several performance advantages over previous servers for hosting Oracle databases. In 32-bit mode, faster processor speeds, faster bus performance, and faster components throughout will contribute to improved performance. In 64-bit mode, there are additional performance advantages that can be attributed to the superiority of the 64-bit processing model over the 32-bit processing model. In particular, Oracle instances that utilize System Global Area memory allocations larger than 4 GB will benefit from improvements in the handling of Very Large Memory with the 64-bit architecture. While the advantages of 64-bit architecture for large System Global Areas (SGA) are fairly well known, it is less commonly understood how the 64-bit model will benefit user scalability. With the 32-bit implementation of Very Large Memory for Microsoft Windows, it is difficult to scale Oracle to support large user populations. The 64-bit version of Microsoft Windows Server 2003 will support scalability for large user populations.

The focus of this paper will be on comparing the performance of Oracle databases on 32-bit Microsoft Windows Server 2003 and 64-bit Oracle databases on the Microsoft Windows Server 2003 version for the EM64T processor. The advantages of the Dell PowerEdge 6850 server for both performance and flexibility will be examined. Test results for scaling both 32-bit and 64-bit versions of Oracle databases for large user populations will be discussed.

New Features of the Dell PowerEdge 6850 Server

The PowerEdge 6850 server builds upon the basic architecture of the PowerEdge 6650 server. Both are scalable up to four Intel Xeon processors, and are designed for Enterprise workloads. The major change for the 6850 server is the adoption of the Intel EM64T architecture. This architecture enables both improved performance and greater flexibility. Below are some of the performance enhancing improvements for the 6850 server.

- The PowerEdge 6850 server is available with 2-4 Xeon EM64T processors, compared to the Xeon 32-bit processors previously available for the PowerEdge 6650
 - Processors are available with clock speeds in the range of 3.0 GHZ – 3.6 GHz, compared to a maximum of 3.0 GHz previously available
 - Processors are available with up to 1 MB of Level 2 cache (initial release), compared to a maximum of 250 MB of L2 cache previously available
- Contains a 677 MHz front-side bus, compared to the 400 MHz front-side bus previously available
- Utilizes 256 MB – 64 GB DDR-2 400 GHz memory, compared to the 200 MHz DDR memory previously available
- Contains four PCI Express slots, as well as PCI-X slots
- Perc4eDC-PCI Express primary and secondary controllers are available
- Windows OS Options
 - Windows Server 2003, Standard, Enterprise, or Web Edition (32-bit)
 - Windows Server 2003, Standard, Enterprise, or Web Edition (64-bit) available mid-2005

Reliability, Availability and Serviceability Features

The Dell PowerEdge 6850 server contains a number of RAS enhancements (Reliability, Availability, and Serviceability). Some of the key RAS features are listed below.

- Available with Dell Remote Access card 4 for remote management
- Redundant power supplies
- Dual on-board Network Interface Controllers (NIC)
- Embedded RAID support (ROMB)
- Up to 5 Ultra 320 SCSI internal hard drives for RAID reliability
 - 18-300 GB disks
 - 10K or 14K rotational speeds
- Support for High Availability Storage
 - PCI-Express fibre Host Bus Adapters for Dell/EMC SAN connectivity
 - Connectivity with Dell PowerVault tape libraries
 - Tape backup software from Veritas, CommVault and TapeWare

Intel EM64T Architectural Features of the PowerEdge 6850 Server

The PowerEdge 6850 server is designed to do more than simply replace 32-bit Intel Xeon processors with 64-bit capable Intel Xeon EM64T processors. Top to bottom, the 6850 server is designed to take advantage of the improved speed

and parallelism of the Intel Xeon EM64T architecture, with implications for almost every sub-system. Some of the key features of the Intel EM64T implementation are highlighted below:

- Intel EM64T processors utilize a flat address space with a single code, data, and stack space
- Large addressable memory space
 - In theory, if a full 64-bits is assigned for memory access, up to 18 Exabytes of memory are accessible
 - In the current implementation, the EM64T chip uses 40 bits to access memory, allowing access to up to 1 Terabyte of memory
- Increased number of registers and increased register width
 - 8 new registers
 - Full 64-bit width register widths
 - Leads to substantial improvements in hyper-threading performance and graphics performance
- Increased Level 2 cache
 - Up to 2 MB (1 MB on the initial 6850 implementation)
 - Brings data closer to the CPU for faster access
 - Increased memory caching leads to reduced I/O
- The Intel 7520 Chipset is another key component of the EM64T architecture
 - Provides support for Xeon processors with EM64T
 - Provides the memory controller management to deal with large amounts of memory
 - Supports up to 800 MHz Front Side Bus, while previous chipsets supported only up to 533 MHz (the 6850 has a 677 MHz front side bus)
 - Automatic retransmission of failed PCI Express packets
- DDR-2 memory
 - Higher density, enabling greater memory scalability.
 - Runs at 400 MHz, twice as fast as the DDR memory used with Xeon and Itanium-2 chips
 - Consumes less power - runs at 1.8V versus DDR which runs at 2.5V
- PCI-Express
 - PCI-Express helps move data faster between processor and memory, helping to speed up processing tasks - especially for memory-intensive applications like databases
 - PCI Express is capable of delivering three to four times the throughput of PCI-X, up to 64Gb/s of throughput
 - PCI Express technology is natively hot-pluggable
 - Doesn't require driver installation for PCI-Express device recognition
 - Enables faster device driver performance
- Demand based switching with Intel Speedstep Technology
 - Lower power consumption
 - Lower heat generation
- The Intel IOP332 I/O processor offers:
 - On-board intelligent RAID
 - Increased I/O performance and reliability
- Reliability, Availability, and Serviceability enablers (RAS)
 - Memory spares
 - Chipkill memory
 - Error correction codes (ECC)

Using the PowerEdge 6850 server in Legacy Mode

Dell PowerEdge 6850 servers may run in several different operating modes. These modes are inherent with the use of the Intel EM64T architecture. Each mode is invoked according to the OS version that the 6850 server is booted in (64-bit, 32-bit, or 16-bit), and some modes are also sensitive to individual application code.

Legacy mode is invoked when a 32-bit OS or 16-bit OS is booted. This mode offers essentially the same options as previous Xeon and Pentium processors. Legacy mode includes three sub-modes: Protected mode, Virtual-8986 mode, and Real mode.

When a 32-bit OS is booted, Protected mode is used. Under Protected mode, a 32-bit application runs in native form, just as it would on a 32-bit Xeon processor.

Virtual-8986 mode is reserved for running 16 bit applications under a 32-bit OS. It is automatically invoked on a per code segment basis. For example, this mode would be invoked when running a 16-bit application in a Command window under the 32-bit version of MS Windows 2000 Server or Windows Server 2003.

With the EM64T architecture, it is still possible to boot a 16-bit OS, such as DOS. When a 16-bit OS is booted, Real mode is invoked

Using the PowerEdge 6850 server in IA32e Mode

The IA32e mode is an extension of the standard IA32 mode for the purpose of 64-bit processing. Nonetheless, 32-bit applications are still supported, even under a 64-bit OS. IA32e mode is automatically invoked when a 64-bit OS is booted. The OS must be compiled for the EM64T processor; an OS version compiled for other 64-bit platforms will not run with an EM64T chip. Under IA32e mode, two sub-modes are available: 64-bit mode and Compatibility mode.

Under the 64-bit mode, all of the new registers are taken advantage of, and the full width of each register is used. Full 64-bit memory addressability is available, up to 1 TB of RAM. Any application compiled for the EM64T environment will run on a 64-bit OS with full access to all 64-bit features.

Compatibility mode is available for running 32-bit applications on a 64-bit OS. In this mode, 32-bit applications can run on a 64-bit OS without recompile. Compatibility mode is automatically invoked on a per code segment basis. This means that 32-bit applications and 64-bit applications can run side by side.

Limitations of the 32-bit Architecture for Oracle Databases on Microsoft Windows Server

It is well known that adequate memory is one of the keys to Oracle database performance. If a server is purchased for the purpose of hosting a mission critical database, it is not uncommon to fill all of the available RAM slots with the maximum memory possible. When an Oracle instance is started, memory is allocated as the System Global Area (SGA). In purchasing a large amount of RAM, the assumption is that performance will scale in something close to a linear fashion as the SGA size is increased. However, many Administrators are surprised to learn that this is not always the case with a system that utilizes 32-bit technology. In fact, as the SGA size is increased above 4 GB in size, performance may not increase and may actually decrease.

Addressable Memory Limitations

There is an inherent maximum of 4 GB of addressable memory for the 32 bit architecture. This is a maximum per process. That is, each process may allocate up to 4 GB of memory. This may not seem like a big limitation, because one would assume that a large number of processes may run on a server at any given time. However, Microsoft Windows uses a multi-threaded model, where Oracle runs as multiple threads under a single process. For Windows 2000 Server and Windows Server 2003, one 4 GB memory segment is shared by all user threads. By default, if you run multiple Oracle instances on the same server, or run other applications on the same server, they will share the same 4 GB of memory.

Available memory is even less, due to OS overhead. By default, on Windows 2000 and Server 2003, 2 GB of the available 4 GB of memory is reserved for the 32-bit OS and 2 GB is shared for User Threads (i.e. the Oracle SGA). One implication of this limitation is that the Oracle SGA must be smaller than 2 GB in size in order to start the database successfully. Even if the SGA is smaller than 2 GB, there may still be insufficient memory for Oracle threads. This is because all memory utilized for Oracle user sessions must also fit inside the 2 GB limit, along with the SGA.

Even without expanding available memory above 4 GB, it is possible to allow Oracle to use more memory. To expand the total memory used by Oracle above 2 GB, the /3GB flag may be set in the boot.ini file. With the /3GB flag set, only 1 GB is used for the OS, and 3 GB is available for all user threads, including the Oracle SGA. While 3 GB isn't enough for most production databases, at least it is a step in the right direction.

Workarounds for Enabling Large Memory on 32-bit Microsoft Windows Server

Workarounds are available for using memory above the 4 GB limit. Intel 32-bit processors such as the Xeon processor support Paging Address Extensions for large memory support. PAE allocates additional memory in a separate memory segment that is also assigned to the process. MS Windows 2000 and 2003 support PAE through Address Windowing Extensions (AWE). PAE/AWE may be enabled by setting the /PAE flag in the boot.ini file. The "USE_INDIRECT_BUFFERS=TRUE" parameter must also be set in the Oracle initialization file. In addition, the DB_BLOCK_BUFFERS parameter must be used instead of the DB_CACHE parameter in the Oracle initialization file. With this method, Windows 2000 Server and Windows Server 2003 versions can support up to 8 GB of total memory. Windows Advanced Server and Data Center versions support up to 64 GB of addressable memory with PAE/AWE.

One limitation of AWE is that only the Data Buffer component of the SGA may be placed in extended memory. Threads for other SGA components such as the Shared Pool and the Large Pool, as well as the PGA and all Oracle user sessions must still fit inside a relatively small memory area. There is an AWE_WINDOW_SIZE registry key parameter that is used to set the size of a kind of "swap" area in the SGA. This "swap" area is used for mapping data blocks in upper memory to a lower memory location. By default, this takes an additional 1 GB of low memory. This leaves only 2 GB of memory for everything other than the Buffer cache, assuming the /3GB flag is set. If the /3GB flag is not set, only 1 GB of memory is available for the non-Buffer Cache components.

Oracle Database Performance Implications

For sufficiently large Oracle databases, performance often increases as the SGA size is increased to the 4 GB limit (or lower, with OS overhead). This is not always true, but with appropriate performance tuning, most medium size to large databases will benefit from an increased SGA size. However, the PAE workarounds mentioned above have a fairly significant performance overhead. In fact, as the SGA size is increased to just above 4 GB (by using PAE), performance may actually decrease. As the SGA size is increased, performance may not improve above the level achieved at the 4 GB level until an 8 to 12 GB SGA size is enabled. This means that it is probably a good idea to build servers with 12 GB or more of RAM, if you wish to host an Oracle instance with a large SGA size on a 32-bit system. Even though the Advanced Server and Data Center versions of Microsoft Windows Server may access up to 64 GB of memory, there is still a performance “dip” between 4 and approximately 12 GB SGA sizes. This means that the Windows 2000 Server or Windows Server 2003 version should probably not be used for SGAs larger than 4 GB, since only 8 GB of memory is accessible with Windows Server.

When AWE is used, it causes the composition of the SGA to be highly skewed towards large Data Buffer caches, since other SGA components and user threads may not be placed above 4 GB. Increasing the Data Buffer cache size may be beneficial for systems that support long running queries that access large amounts of data, such as Data Warehouses. OLTP systems may benefit less, particularly when there are a large number of users simultaneously accessing the database. Supporting large user populations places additional burdens on the Large Pool and the Shared Pool. With only 1-2 GB to work with, this can significantly limit the scalability of Oracle in terms of number of user sessions supported. For many companies, this is the most severe limitation of Oracle on 32-bit Microsoft Windows.

Some relief may be obtained by setting the /3GB flag as well as the /PAE flag in Oracle. This at least assures that up to 2 GB of memory is available for the Large Pool, the Shared Pool, the PGA, and all user threads, after the AWE_WINDOW_SIZE parameter is taken into account. However, Microsoft recommends that the /3GB flag not be set if the /AWE flag is set. This is due to the fact that the total amount of RAM accessible for ALL purposes is limited to 16 GB if the /3GB flag is set. RAM above 16 GB simply “disappears” from the view of the OS. For PowerEdge 6850 servers that can support up to 64 GB of RAM, a limitation to only 16 GB of RAM is unacceptable.

Advantages of the 64-bit Architecture for Oracle Databases on Microsoft Windows Server

The Dell PowerEdge 6850 server, featuring the Intel EM64T architecture, is designed for hosting Oracle databases. With the PowerEdge 6850 server, Administrators can continue to host their databases on 32-bit Windows, and then make the transition to 64-bit Windows when convenient. As databases are migrated from 32-bit versions to 64-bit versions, significant performance improvements will be realized. In addition, Database Administrators that continue to run their databases in 32-bit mode may also experience improved performance, due to the advanced features of the Intel EM64T architecture.

Memory Addressability and Large SGA Sizes

In IA-32e mode, a large amount of memory is accessible: up to 1 TB in theory, up to 64 GB in practice. For servers dedicated for Oracle databases, it is not uncommon to allocate up to 75% of the available memory for the Oracle SGA (as a rough rule of thumb). The remaining 25% is used for the Operating System, and Oracle User processes contained in separate memory segments than the SGA.

The first time that Oracle accesses a particular row of data in a database table, the row must be retrieved from disk through I/O. The data is placed in the Data Buffer portion of the SGA, which is used to cache data. Once the data is cached in the SGA, it may be retrieved again directly from the SGA, for an extended period. The data is only flushed from the SGA when it has not been accessed for a relatively long period, or if more buffer space is needed for other rows. Memory is accessed about 10,000 times faster than disk drives, so replacing I/O access to data with access via memory is extremely desirable for database performance. The larger the SGA, the more data that can be buffered in memory and the higher chance that a given row is retrieved from the database buffer instead of through disk I/O.

As noted previously, the model used for extended memory access under a 32-bit Operating System entails a substantial performance penalty. However, with a 64-bit OS, a flat linear model for memory used, with no need for PAE to access memory above 4 GB. Improved performance will be experienced for database SGA sizes greater than 3 GB, due to elimination of PAE overhead.

One benefit will be improved query response times for all types of databases. In particular, databases for Decision Support Systems (DSS) are query-intensive. In addition, DSS databases tend to be large to very large, increasing the potential benefit of a large SGA size. The performance benefit for a particular query will depend on how that query is tuned to access data. Queries that access data through B-tree indexes and bitmap indexes will benefit from larger database buffers. In addition, queries that require large sorts should benefit from both faster sorting and more sorting done in memory rather than disk.

Queries that access data through full table scans may not benefit as much from large memory buffers and caching. Some tables in DSS systems may grow extremely large. It may not be possible to buffer the entire table in the SGA or to keep it there for extended periods. In these cases, access through disk I/O may actually be preferable to memory access. These queries may be tuned to direct data to the Recycle Buffer cache in the SGA, instead of the Default Buffer Cache. This will cause the table scan data to be flushed out of the SGA quickly, and avoid “flooding” the main cache, which would displace useful data. Of course, this minimizes the benefit of large memory for these queries. Nevertheless, the benefits of faster CPU performance, faster bus performance, and improved I/O efficiency may still lead to overall improvements in query speed.

Scalability Improvements for Large User Populations

It is traditional to measure database performance in terms of transactions per second or query response time. An equally useful measurement is the number of simultaneous user sessions supported. This is a particularly important measurement for OLTP (Online Transactional Processing) databases. One of the benefits of migrating to the EM64T architecture is improvement in the ability to support large user counts for Oracle databases. Certainly, the memory addressability changes mentioned above play a large role in the capability to support large user populations. This is particularly true with 32-bit Windows

2000 Server and Windows Server 2003. As discussed above, the memory used to support large user populations with AWE extended memory is only 1-2 GB. It is not uncommon for the number of simultaneous user sessions to be constrained by this limit. Under the 64-bit versions of Windows Server 2003, the memory for Oracle user sessions may grow to the bounds imposed by available memory.

OLTP systems also benefit from other attributes of the Intel EM64T architecture. These include:

- Faster context switching
- Higher frequency clock cycles
- Improved chip parallelism
- Hyper-threading improvements
- Reduced calls to I/O due to larger Level 2 cache

DSS systems, such as Data Warehouses, typically do not support as many users as an OLTP system, due to long-running resource intensive queries. However, successful Data Warehouse systems always attract more users with time, eventually straining the capacity of the system. With the Intel EM64T architecture, 64-bit DSS databases will also support more simultaneous queries, for all of the reasons noted above.

Testing Oracle Database Scalability with the EM64T Processor

To illustrate the key concepts of this paper, performance tests were conducted with two Dell PowerEdge 6850 servers, one server dedicated for 32-bit testing and one server dedicated for 64-bit testing. Each PowerEdge 6850 was equipped with the following hardware:

- 4 EM64T 3.6 GHz processors
- 20 GB of DDR-2 RAM on each server
- 2 internal 76 GB SCSI drives in a RAID 1 configuration (for the OS)
- One Gig-E network connection per server
- A CX500 array with the following configuration:
 - 10 external Fibre Channel disks in a RAID 10 configuration per server (for Oracle data and index files)
 - 4 external Fibre Channel disks in a RAID 10 configuration per server (for Oracle Redo Log files)
 - One Fibre Channel Host Bus Adapter per server direct attached via a single Fiber Optic cable to a dedicated Storage Processor

In terms of software configuration, the following components were utilized:

- 32-bit configuration:
 - Microsoft Windows Server 2003, Enterprise Edition
 - NTFS filesystem
 - Oracle Database Release 1 (10.1.0.2) for Microsoft Windows (32-bit)
- 64-bit configuration:
 - Microsoft Windows Server 2003, Enterprise x64 Edition (“Beta” version, final release candidate)
 - NTFS filesystem
 - Oracle 10g Developer’s Release for 64-bit Extended Systems

The focus of these tests was on measuring the scalability of the number of simultaneous user sessions, rather than on “benchmark” measurements. The database schema and queries for the tests were based on the standard Microsoft Nile database, version 4, downloaded from Microsoft. Previous versions of the Nile database have been used for comparison tests of database scalability by a number of vendors and IT industry magazines. The Nile database simulates the back end OLTP database for an online bookstore (a.k.a. Amazon). A 32 GB database was loaded on each 6850 server. The Nile database contains the following tables:

- Customers Table
 - CustomerID NUMBER NOT NULL
 - FirstName varchar2(50) NULL
 - LastName varchar2(50) NULL
 - Address1 varchar2(50) NULL
 - Address2 varchar2(50) NULL
 - City varchar2(50) NULL
 - State varchar2(50) NULL
 - Zip NUMBER NULL
 - Email varchar2(50) NULL
 - Phone varchar2(50) NULL
 - CreditCard varchar2(50) NULL
 - CreditCardExpiration varchar2(50) NULL
 - Username varchar2(50) NULL
 - Password varchar2(50) NULL
- Orders Table
 - OrderID NUMBER NOT NULL
 - OrderDate date NOT NULL
 - CustomerID NUMBER NOT NULL
 - NetAmount number(22,7) NOT NULL
 - Tax number(22,7) NOT NULL
 - TotalAmount number(22,7) NOT NULL
- OrdersDetails Table
 - OrderItemID NUMBER NOT NULL
 - OrderID NUMBER NOT NULL
 - BookID NUMBER NOT NULL
 - Quantity NUMBER NOT NULL
- Products Table
 - BookID NUMBER NOT NULL
 - SubjectID NUMBER NOT NULL
 - BookTitle varchar2(50) NULL
 - Author varchar2(50) NULL
 - Price number(22,7) NULL
 - Retail number(22,7) NULL
 - ISBN varchar2(50) NULL
 - QuantityOnHand NUMBER NULL
 - SpecialItem char (1) NOT NULL
- UserSession Table
 - IDSession NUMBER NOT NULL
 - CustomerID NUMBER NULL
- SessionData Table
 - Row_id NUMBER NOT NULL
 - IDSession NUMBER NOT NULL
 - BookID NUMBER NULL
 - Author varchar2(50) NULL
 - Title varchar2(50) NULL
 - Price number(22,7) NULL
 - Qty NUMBER NULL
- Subjects Table
 - SubjectID NUMBER NOT NULL
 - SubjectName varchar2(50) NOT NULL

Memory Configuration for Testing

For both 32-bit tests and 64-bit tests, an attempt was made to use the largest SGA size possible, without causing memory paging to disk or memory allocation errors. Figure 1 illustrates the three configurations used. For 32-bit tests, one set of tests was conducted using the /PAE flag in the boot.ini file. Another set of tests was run using both the /PAE and the /3GB flags in the boot.ini file. Note that the maximum addressable memory was limited to 16 GB of RAM for the latter case (as explained above). For both cases, the AWE_MEMORY_WINDOW registry parameter was set to 0.75 GB, which was the minimum size that allowed the database to start. This left approximately 1.25 GB of low memory available for the /PAE case, and approximately 2.25 GB of low memory available for the /PAE and /3GB case. In this memory area, all components of the SGA except the Buffer Cache, the PGA, and User Session memory were required to fit. To achieve the large SGA size, the bulk of the memory assignment was placed in the buffer cache, approximately 11 GB. This led to total SGA sizes of 13-14 GB for the two 32-bit cases.

For the 64-bit case, low memory constraints were not a factor. Compared to the 32-bit cases, the Buffer Cache was considerably smaller (approximately 5 GB), and other SGA components were considerably larger. For example, the Shared Pool size was over 9 GB for the 64-bit case, compared to 600 MB – 1.3 GB for the 32-bit cases. The total SGA size was approximately 15.5 GB for the 64-bit case. The SGA size does not include the actual memory used for the PGA at run time, or the memory used by User Sessions. This can easily bring the total Oracle memory usage up to the 17 GB+ level.

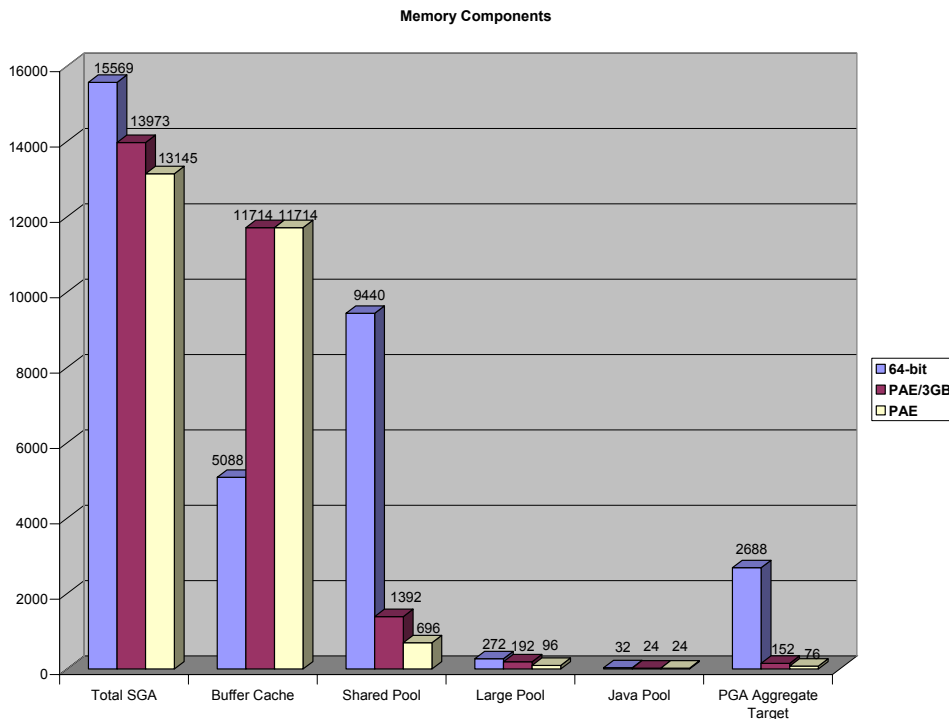


Figure 1 Memory Components of 32-bit and 64-bit Tests

User Scalability Tests

For the front end, this implementation of the Nile 4 tests uses perl scripts to launch multiple, parallel sessions on client machines, in order to simulate user connections, transactions, and queries. Two Dell PowerEdge 6650 servers with 8 GB of RAM and four Xeon processors were used as the clients; one client for the 32-bit tests, and one client for the 64-bit tests. A mix of browsing queries and “checkout” transactions was used for these tests. The number of SQL statements completed per minute and number of user connections are measured. To simulate realistic user behavior, "think" time is built in as random wait times between transactions, averaging 15 seconds. Up to 1300 simultaneous user sessions were generated on each client. Each user session was instantiated as a persistent connection, requiring the database server to support the full number of connections throughout each test.

For all of the 32-bit and 64-bit tests, the response in SQL Statements per Minute scaled linearly for increasing user sessions, as shown in Figure 1. This indicates that realistic loads on the database for web browsing queries and OLTP activity should not strain the CPU capacity of the PowerEdge 6850 server up to at least 1300 simultaneous user sessions.

The 32-bit tests were divided into four cases:

1. /PAE flag set with Dedicated Server connections.
2. /PAE flag set with Shared Server connections.
3. /PAE flag and /3GB flag set with Dedicated Server connections.
4. /PAE flag and /3GB flag set with Shared Server connections.

The 64-bit tests used Dedicated Server connections. Dedicated Server Connections offer the highest performance levels, but require more memory than Shared Server connections. Shared Server connections conserve memory by sharing connections with a multi-threaded model. However, performance per connection may be less than that experienced for Dedicated Server connections.

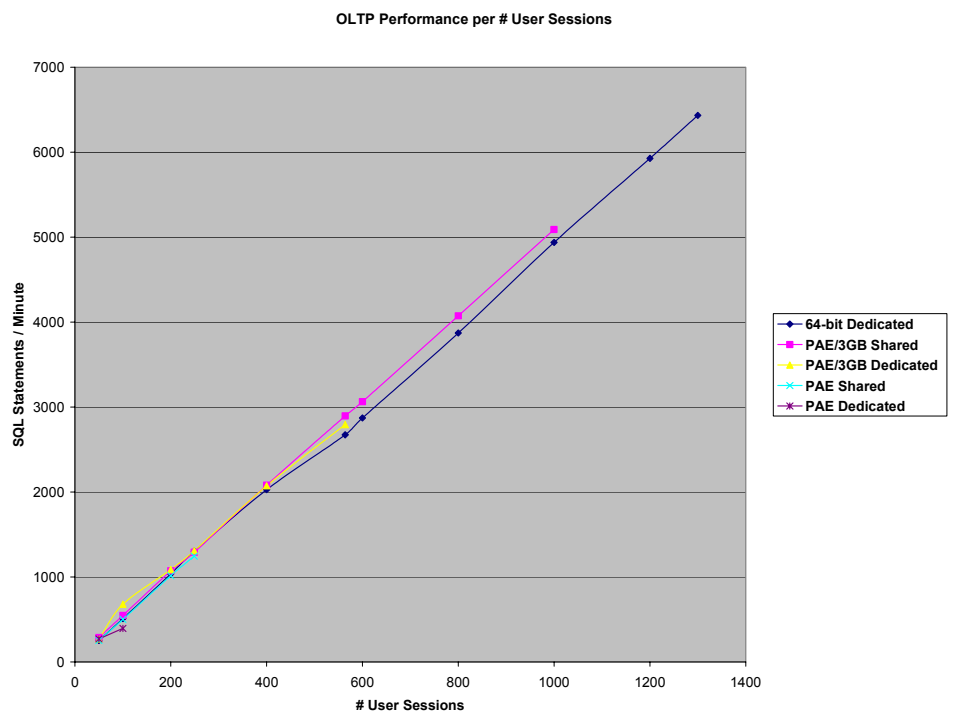


Figure 2 OLTP Performance per # User Sessions

The 64-bit Dedicated Server tests show that up to 1300 user sessions may be simultaneously before the database refuses to connect any more sessions. In Figure 1, the 32-bit tests show performance that closely tracks the 64-bit tests. However, each 32-bit test ultimately reaches a limit where no more user connections are allowed by the database. This is usually the point where Shared Pool or Large Pool resources are exhausted. This limit is displayed more clearly in Figure 2. The maximum connections limit will obviously be reached far sooner for the 32-bit database instances than the 64-bit database instance, given the 32-bit requirement to fit all non-Buffer Cache memory into the first 1-2 GB of memory. The 64-bit database has the built-in advantage that it may be tuned for OLTP activity to use relatively higher ratios of Shared Pool, Large Pool, and PGA Aggregate Target memory to Buffer Cache memory.

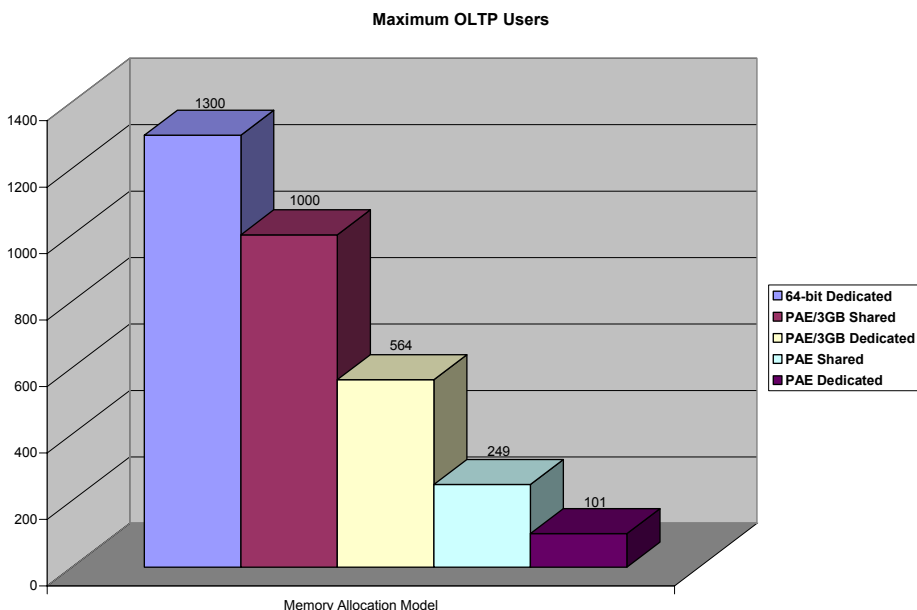


Figure 3 Maximum OLTP Users for each Memory Allocation Model

The maximum number of connections that may be made is sensitive to database tuning, in particular with regards to Shared Server connections. It is possible that more user connections may be established than shown here, with further tuning. However, a few key points should be understood:

- The trend is clear; all 32-bit cases will be limited in terms of user scalability by the amount of non-Buffer Cache memory available.
- Additional gains in the number of Shared Server connections will come at the expense of performance per connection.
- Configurations that use both the /PAE flag and the /3GB flag are ultimately counter-productive, as they limit the total addressable memory to only 16 GB.
- With 64-bit Oracle on Microsoft Windows Server 2003, user scalability may be improved by adding more RAM to the system (up to 64 GB).
- Using Shared Server connections with 64-bit Windows Server 2003 could conceivably allow several thousand simultaneous user connections to be supported.

Report Query Tests

While the main focus of these tests was on user scalability, a test of report query performance was also conducted. A typical report query was constructed that would access a large number of rows from multiple tables. The following query was used:

```
select a.state STATE, a.city CITY, b.subjectname
SUBJECTNAME, sum(c.netamount) NETAMOUNT,
sum(c.tax) TAX, sum(c.totalamount) TOTALAMOUNT from
nile4.customers a, nile4.subjects b, nile4.orders c,
nile4.ordersdetails d,
nile4.products e
where c.orderid = d.orderid and d.bookid = e.bookid and
e.subjectid = b.subjectid and c.customerid = a.customerid
group by a.state, a.city, b.subjectname
order by a.state, a.city, b.subjectname;
```

The initial test was conducted with a single user on both the 32-bit and 64-bit system. Figure 4 shows the results in terms of number of seconds required for a query response. The data shows the 32-bit system responding in approximately one half the time of the 64-bit system. This is also consistent with data loading rates, which showed the 32-bit data load completing in one half the time of the 64-bit data load.

At first glance, it may appear that the difference is due to the size of the Buffer Cache: the 32-bit system has twice the Buffer Cache of the 64-bit system. However, a closer examination of the systems revealed a different answer. This reporting query causes a substantial amount of I/O to disk. The 32-bit system performed all I/O tasks faster, including non-SQL I/O operations, such as large file copies. The performance monitor showed that the I/O pattern for the 64-bit system was “bursty”, while the 32-bit performance was continuous and efficient. The difference appears to be the Fibre Channel I/O driver. The 32-bit system utilizes a well tested driver from the HBA vendor that is known to produce good performance. The only Fibre Channel driver available for the 64-bit system is a “beta” driver supplied by Microsoft. It is clearly this driver that is the cause of the slow I/O.

Undoubtedly, the 64-bit Fibre Channel drivers will be improved substantially before the final release of Microsoft Windows Server for Extended 64-bit systems. For now, the response from the driver dominates the report query results, allowing us to draw no meaningful conclusions about the relative performance of 32-bit reporting queries and 64-bit reporting queries. Since the OLTP tests did not feature large I/O reads and writes from disks, I/O driver performance was not an issue with the OLTP tests covered in the previous section.

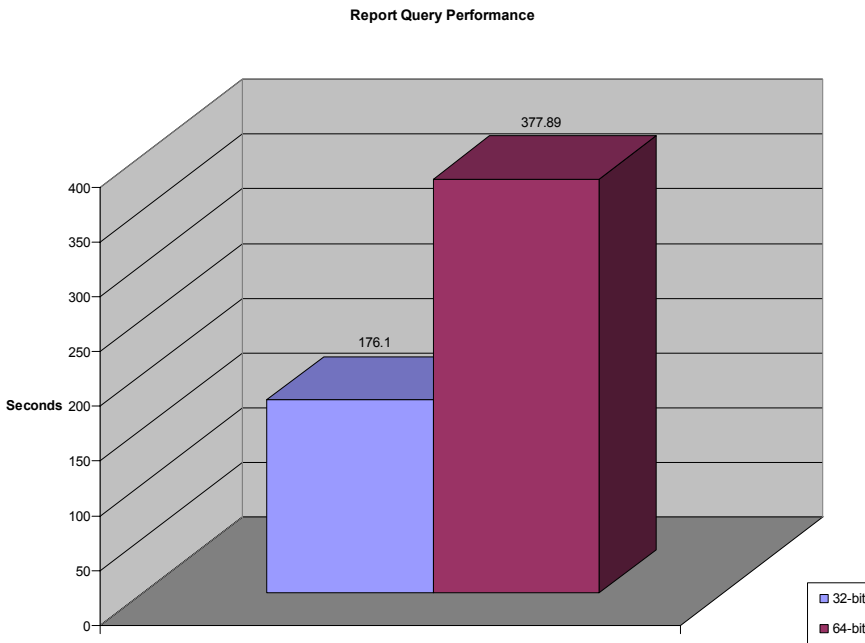


Figure 4 Time to Run a Report Query for 32-bit and 64-bit Tests

Conclusions

The Dell PowerEdge 6850 server offers a powerful platform for deploying Enterprise class applications, such as Oracle databases. Key performance improvements over previous server lines include increased processor speed, improved parallel processing performance, faster memory, the use of fast DDR-2 RAM, and the use of PCI-Express technology. In addition, Dell fully exploits the 64-bit extension technology offered by Intel's EM64T processors. From the first day of release, the PowerEdge 6850 will support Oracle databases on the Microsoft Windows 2000 Server and Windows Server 2003 32-bit operating systems. Later in 2005, when Microsoft releases Microsoft Windows Server 2003 for Extended 64-bit architectures and Oracle releases the 10g database for Extended 64-bit architectures, the PowerEdge 6850 will provide the perfect platform for transitioning to 64-bit databases.

One of the main limitations of Oracle databases running on 32-bit versions of Microsoft Windows has been the handling of Very Large Memory. Without special workarounds, it is not possible to utilize more than 4 GB of memory. With the appropriate workarounds in place, it is possible to increase the Oracle System Global Area to use more than 4 GB of memory. However, the only SGA component that may be increased in this fashion is the Buffer Cache component. While a large Buffer Cache is often beneficial for DSS or reporting queries, it is not as beneficial for OLTP activity. In particular, supporting large user populations requires extra Shared Pool, Large Pool, and Program Global Area resources, which are limited on 32-bit versions of Windows to only 1-2 GB in size. Many companies have avoided placing mission-critical Oracle databases on Microsoft Windows, due to the perception that Oracle databases on Windows do not scale to large user populations.

This paper explored the implications of using the Dell PowerEdge 6850 server to host 64-bit Oracle databases on the 64-bit version of Microsoft Windows Server 2003 for the Extended 64-bit architecture. Test results showed that there are no limitations on how Oracle memory can be configured for 64-bit databases. It is possible to support large user populations, with 1300 users

connected via Dedicated Server connections in the tests. Even larger user populations may be supported with more memory and the use of Shared Server connections. Although test results for report queries were not conclusive in this study, environments that feature long-running report queries and Data Warehouse activity should also benefit from performance improvements due to the 64-bit memory model.

Companies that are currently deploying Oracle databases on the Microsoft Windows Operating System will experience performance improvements by migrating their databases to the Dell PowerEdge 6850 Server. When the Extended 64-bit versions of Oracle and Microsoft Windows Server 2003 are released, the greatest performance gains will be realized by migrating to 64-bit databases. For those companies that are evaluating Oracle databases on Microsoft Windows, the scalability improvements that are documented here should be carefully considered. With the Dell PowerEdge 6850 server and Microsoft Windows Server 2003 for the Intel EM64T architecture, the time to implement Enterprise class Oracle databases on Microsoft Windows has arrived.

References

“The 64-bit Tipping Point - Optimizing Performance, Flexibility, and Value with Intel Itanium Architecture and Intel Extended Memory 64 Technology (Intel EM64T);” September, 2004, Intel Corporation;

http://www.intel.com/business/bss/products/server/64-bit_tipping_point.pdf

“Evaluating the Performance of 64-bit Oracle9i Database Release 2 on the Intel EM64T-based Platform;” Ramesh Radhakrishnan, PH.D., Dell Corporation; Dell Power Solutions, October 2004, Dell Corporation;

<http://www.dell.com/downloads/global/power/ps4q04-20040167-Radhakrishnan.pdf>

“Intel's Enterprise Processor Plans: Positioning the Xeon Processor and the Itanium Processor;” Author: Vernon Turner; April 2004, IDC Corporation; http://www.intel.com/technology/64bitextensions/4071_intel_xeon_rev3.pdf

“Intel Extended Memory 64 Technology FAQs;” 2005, Intel Corporation; <http://www.intel.com/technology/64bitextensions/faq.htm>

“Intel Extended Memory 64 Technology Software Developer’s Guide Volume 1 of 2;” 2004, Intel Corporation, Order Number: 300834-002; <http://www.intel.com/technology/64bitextensions/300834.htm>

“Migrating to Industry-Standard 64-bit Architectures;” John Fruehe, Dell Corporation; Dell Power Solutions, March 2004, Dell Corporation; <http://www.dell.com/downloads/global/power/1q04-fru.pdf>

“Next Generation Server Technologies;” 2005, Dell Corporation; http://www1.us.dell.com/content/topics/topic.aspx/global/products/pegde/topics/en/8g_technologies?c=us&cs=555&l=en&s=biz

“Planning Considerations for Intel Extended Memory 64 Technology;” John Coombs & John Fruehe, Dell Corporation; Dell Power Solutions, June, 2004, Dell Corporation; <http://www.dell.com/downloads/global/power/ps2q04-018.pdf>

“Product Brief - Intel Xeon Processor with 2MB L2 Cache;” 2005, Intel Corporation; http://www.intel.com/products/server/processors/server/xeon/prodbrief.pdf?iid=i_be_xeon+prodbrief_xeon&

Larry Pedigo (lpedigo@perftuning.com) is a Performance Consultant for Performance Tuning Corporation. He has been involved with Oracle RAC deployments, SAN deployments, and Oracle and SAN performance optimization for the past two years. Larry has a Bachelors degree in Geology from Texas A&M University, and over 15 years experience in the IT field.

About the Company

Performance Tuning Corporation (PTC) is a technology services company that provides database performance tuning, SAN and cluster configuration, data replication, capacity planning, benchmarking and load-testing. In business since 1997, PTC continues to distinguish itself from other database and system consulting companies through its efficient, targeted approach to solving customer problems.

PTC consultants have extensive experience with major hardware and software technologies (Dell, IBM, Sun, HP, EMC, Oracle, and Microsoft). PTC has led numerous benchmarking efforts on TPC™ and custom benchmarks. Their consultants have published more than eight books on Oracle and Microsoft database technology, as well as numerous technical papers covering database tuning, SAN deployment and clustering.

For more information, please contact PTC at:
(800) 887-4513 or
info@perftuning.com
www.perftuning.com