

## **A Comparative Investigation and Evaluation of Oracle and Sql Server with respect to Performance and Scalability.**

By: Phathisile Sibanda g05s5782@campus.ru.ac.za

Supervisor: John Ebden

Date: 30 May 2005

### **Abstract**

Performance and Scalability are important factors determining database availability and reliability. This is true especially now that computer systems, due to the advent inclusion of the internet in Online transaction processing (OLTP) and e-commerce applications, are becoming larger than ever before to accommodate large user populations concurrently accessing the database. This evolution gave birth to endearing need for high throughput, response time, good data consistency and concurrency as will be seen from this research evaluation of Oracle 9i and SQL server 2000 DBMS. Anticipated deliverables of this research will include a set of comparative graphs from performance tests and the accompanying technical explanation to support the results.

### **Introduction**

This research specifically explores and evaluates performance and scalability capabilities of SQL server 2000 and Oracle 9i. This would be accomplished through the use of a combination of standard benchmarks (TPC-C) embedded under third party software like Quest Central 2005 for Oracle and SQL server and performance tuning tools like System Monitor and Analysis Manager included with these products. Furthermore, technical features influencing performance and scalability will be examined to support the benchmark results.

The focus of this literature review is to discuss what database experts say about performance and scalability and investigate benchmarking techniques already in use by database vendors and independent organizations such as the Transaction Processing Performance Council<sup>1</sup> and ISV in identifying database bottlenecks and test subsequent performance tuning abilities for Oracle9i and SQL server 2000 database servers.

---

<sup>1</sup> <http://www.tpc.org/> The TPC defines transaction processing and database benchmarks and delivers trusted results to the industry.

First I explain the meaning of performance and scalability as viewed by database Gurus concurrently with a brief overview of platform availability for these two products. Secondly benchmarking techniques already in use are discussed simultaneously with an evaluation of performance tuning and monitoring tools. Lastly the discussion will focus on database features most likely to influence performance and scalability in modern relational database systems and thus indirectly performance tests. These will include among others: the concurrency model, indexing, partitioning, parallel execution and clustering design of the database server.

### **Terminology and platform availability evaluation for Oracle9i and SQL server**

According to the [5] Quest organisation (producers of performance tuning software for both Oracle and SQL servers) they believe performance and scalability benchmarking is best done by automation software for a number of reasons. Firstly automation allows testing teams to repeat test scenarios especially those for stress and load testing which are difficult to simulate manually. Furthermore, automation improves benchmarking accuracy and saves a lot of time allowing database administrators to concentrate on other things like scripting scenarios. [10]RPM Solutions Pty Ltd support this notion of using automation tools for performance testing and benchmarking by further breaking down performance testing into three section ,load testing , stress testing and scalability testing. They define performance testing as “tests that determine end to end timing (benchmarking) of various time critical business processes and transactions, while the system is under low load, but with a production sized database”. To enhance their performance definition they observe the following pre-requisites as necessary prior to any performance testing.

<b>Performance Test Pre-Requisites</b>	<b>Comment</b>	<b>Caveats on testing where Pre-requisites are not satisfied.</b>
Production Like Environment	Performance tests need to be executed on the same specification equipment as production if the results are to have integrity.	Low bandwidth performance testing of high bandwidth transactions where communications processing contributes to most of the response time can be tested.
Production Like Configuration	Configuration of each component needs to be production like.  For example: Database configuration and Operating System Configuration.	While system configuration will have less impact on performance testing than load testing, only substantial deviations from expected transaction response times should be reported.

Production Like Access	If clients will access the system over WAN, dial-up modems, DSL, ISDN, etc. then testing should be conducted using each communication access method.	Only tests using production like access are valid.
Production Like Data	All relevant tables in the database need to be populated with a production like quantity with a realistic mix of data. e.g. Having one million customers, 999,997 of which have the name "John Smith" would produce some very unrealistic responses to customer search transactions	Low bandwidth performance testing of high bandwidth transactions where communications processing contributes to most of the response time can be tested.

Fig 1<sup>2</sup> Some considerations for performance tests

Scalability according to the Oracle 9i documentation is a system's ability to process more workload, with a proportional increase in system resource usage. That is, in a scalable system, if you double the workload, then the system would use twice as many system resources. This sounds obvious, but due to conflicts within the system, the resource usage might exceed twice the original workload ([5] Oracle9i).

Load testing determines the system's behaviour under various workloads. Its main objective is to determine how system's components react as the workload is gradually increased. The usual outcome is the determination of the system performance in terms of throughput, response time, and CPU load and memory usage [2]Marinescu. Stress testing on the other hand determines the breaking point or unacceptable performance point of a system at maximum service. Scalability testing is used to evaluate the effects of adding additional hardware and /or software to distribute "work" among system components. Tests of this kind according to [10] RPM Solutions Pty Ltd can be performed in a variety of configurations, with such variables as network speed, number and type of server/CPU's and memory varying independently.

Another important performance variant as stated in the [3] MSDN SQL server documentation is database platform availability. Although database server portability is not directly linked to performance, software availability across a wide variety of hardware and operating systems enables users to effortlessly upgrade or replace their hardware systems without having to worry about

<sup>2</sup> Table from Loadtest.com.au – shows performance testing prerequisites  
[http://www.loadtest.com.au/types\\_of\\_tests/performance\\_tests.htm](http://www.loadtest.com.au/types_of_tests/performance_tests.htm)

changing, redesigning or rebuilding their applications. In other words, cross-platform availability helps protect the initial designs in application software and thus contributes in delivering performance consistency across multiple platforms. Database Guru Chigrik ,supported this point in his letter entitled “A *SQL Server 2000 versus Oracle 9i : focus platform availability*“, where he states that Oracle9i Database is available on a large selection of hardware and operating systems, scaling from low-end single processor servers to large symmetrical multiprocessor machines to multi-node clusters. Thus Oracle9i, users are able to upgrade hardware and operating systems without changing or rewriting their applications. SQL Server 2000 on the other hand according to [3] MSDN, 2005 only runs on Microsoft’s operating systems. Therefore customers wishing to upgrade hardware are limited to platforms running these systems and must face the cost of converting their systems completely if they ever outgrow the capacity of their platform. Although not directly but indirectly from the above view points platform availability affects performance and scalability in a big way.

**Factors influencing performance tests**

[1]Burleson writer of the article “*Database Benchmark Wars*” argues that besides platform availability it is equally important for database administrators also appreciate the influence that some DBMS design factors induce on performance tests. Five of such factors include the concurrency model, indexing, partitioning, parallel execution, and clustering. This section discusses these in more detail starting with the concurrency model. Basically in multi-user environments, concurrency control ensures that data updates made by one user do not adversely affect those made by other users. Fig 2 below summarizes features that SQL server and Oracle 9i.use to implement the concurrency model.

Oracle 9i	SQL Server 2000
Multi-version read consistency	Not available
Does not have read locks	Requires shared read locks to avoid dirty reads
No dirty reads	Dirty reads if not using shared locks
Non-escalating row-level locking	Locks escalate
Readers don’t block writers	Readers block writers
Writers don’t block readers	Writers block readers
Minimal deadlocks under load	Deadlocks can be a serious problem under load

Fig 2 Concurrency Models

As seen in Fig 2 above Oracle 9i implements the concurrency model through multi-version read consistency which always provides consistent and accurate results

through the use of undo records. Oracle uses the current information in the undo records to construct a read-consistent view of a table's data, and to ensure that a version of the information, consistent at the beginning of the uncommitted transaction, can always be returned to any user. SQL Server 2000 on the other hand, uses locking data to prevent it from changing while being read or preventing or queries from reading changed but uncommitted information. SQL Server, in contrast, uses shared locks to ensure that data readers only see committed data. These readers take and release shared locks as they read data. These shared locks do not affect other readers. A reader waits for a writer to commit the changes before reading a record. A reader holding shared locks also blocks a writer trying to update the same data.<sup>3</sup>

A consequence is that releasing locks quickly for applications that support high numbers of users is more important in SQL Server than in Oracle. Because users in SQL block each other and only write when they get the lock which means in high user populations scalability is greatly affected.

Another factor affecting performance and scalability is indexing capabilities ,indexes are basically database structures that are created to provide a faster path to data. Using indexes can dramatically reduce disk I/O operations, thus increasing the performance of data retrieval. Basically both Oracle and SQL Server 2000 support traditional B-Tree indexing schemes, which are ordered lists of key values, associated with the storage location of the table row that contains these values. Added to this they also both support clustered indexes which are basically index-organized tables that facilitate access to table data for queries involving exact match and/or range search on the primary key because table rows are stored in the leaf nodes of the primary key index. To further enhance performance gains Oracle 9i supports static bitmap indexes and bitmap join indexes, whose usage can provide huge performance benefits for load and query operations in data warehousing environments.

Partitioning allows large database structures (tables and indexes for example) to be decomposed into smaller and more manageable pieces. Although it is primarily considered a feature for manageability and availability, partitioning also provides a number of performance benefits.

Parallel execution of SQL operations can vastly affect the performance for operations involving large volumes of data. It helps reduce response time for data-intensive operations on large databases typically associated with decision support systems and data warehouses. Oracle9i will execute INSERT, UPDATE, DELETE, and MERGE<sup>4</sup> statements in parallel when accessing both partitioned

---

<sup>3</sup> Migrating Oracle Databases to SQL Server 2000”, SQL Server 2000 Resource kit, p. 57

<sup>4</sup> The MERGE statement is explained later in this document, in the section on Additional Data Warehousing capabilities

and non-partitioned database objects. With SQL Server 2000, INSERT, UPDATE, and DELETE statements are executed serially (MERGE is not supported).

Clusters are groups of independent servers, or nodes, connected via a private network (called a cluster interconnect), that work collaboratively as a single system. Clusters allow applications to scale beyond the limits imposed by single node systems when processing loads exceed the capacity of large individual servers. However only Oracle 9i provides real support for clustered configurations, where full and transparent scalability can be obtained by simply adding new nodes as the demand increases. As indicated in [3] SQL Server 2000 documentation, “SQL Server 2000 does not support this type of clustering”. Instead, users are forced to use a group of databases to achieve some kind of scalability.

A recent study<sup>5</sup> based on the above discussed factors shows that Oracle is the primary database choice for 51% of the Fortune 100 companies and that 3 out of 4 of them run their enterprise applications on Oracle. However it is important to consider more factors than only design features as suggested by the Transaction Processing Performance Council. They define a list of vendor independent performance testing steps and benchmarks which I discuss below.

### **Performance tuning and monitoring tools (for benchmarking)**

Performance tuning and diagnosis of databases has been the main reason for pursuing performance and scalability tests. This paper discusses benchmarking strategies and automation tools used by the TPC organisation, [6] PC Magazine, 2003 in collaboration with the [8] *eWEEK* publishers and the Quest Company (Producers of the automation tool Quest Central for Oracle and Microsoft databases that I will be using for workload simulation in my tests).

[6] Gambler for the PC Magazine stats that database software testing is not for the fainthearted, and it is not a project you can do alone. Some scholars believe that SQL server and Oracle databases behaviours as workload increase will not be largely dependent on the technical design of the server nor will it be dependent on the server availability to lots of platforms but on the experience. To support this view point SQL server expert Geniemani argues that it is a fallacy to say that SQL Server does not have this and that and our Oracle does all of this while forgetting that many things can be accomplished differently in different databases and environments and thus it does not make sense to compare apples and oranges. One of his examples to support this idea is that used by

---

<sup>5</sup> Database Penetration in the Fortune 100, a Usage-based Market Share Analysis-The FactPointGroup, April 2002  
[http://www.oracle.com/features/9i/index.html?tl db\\_facts.html](http://www.oracle.com/features/9i/index.html?tl db_facts.html)

Oracle expert that SQL server does not read/write files but forgetting to say that SQL Server does support extended stored procedures, which you can do file/write and virtually anything else that you can do in Windows and beyond. However [1] Burleson an Oracle expert crashes this point saying SQL server is not an enterprise solution due to the lack of scalability features and the fact that it is proprietary making it vendor specific which limits interoperability in the part of the user. These are some of the views the project will prove or disprove.

The approach taken by the three organisations is more or less similar with variations on the tools used to accomplish the same tasks. The first step in performance testing [6] PC Magazine used was the definition of input data which includes identification of scenarios, workload model and choosing suitable favourable counters for measuring performance and scalability.

Having defined your inputs [7] *Kevin Kline (Quest)* argues that the next step is to design and perform your functional test. He defines these as baseline measurements that tell you how a server behaves under normal conditions. These are done to identify problems that arise when the server is at rest, some of which may need to be corrected before any consequential load testing. For this process they used automation tools System Monitor and Quest Spotlight. Kline argues that System Monitor is the best tool to use since it allows you to view both real server statistics during server runtime and saved log statistics. Saved statistical log tables allow you to save system statistics for analysis at a later date. However when analysing a large number of counters Systems monitor according to Gambhir ([6] PC Magazine, 2003) lacks numerous graphical components required to produce visual graphs and charts to enhance analysis. As a result they used Quest Spotlight to compliment System monitor. TPC tests were largely based on the TPC-C benchmark suit which offers guidelines for defining inputs to performance test as well as tools to use. Concurrent to functional test the PC Magazine invited the respective database vendors to fine tune their applications to ensure they are in good condition before testing begins.

At the completion of functional tests the organisations embarked on performance tests (in the form of load and scalability tests) which they performed in different ways. To simulate a large number of users interacting with their Nile test dummy database and measure the performance results, they used e-TEST suite, a load/stress tool from Empirix ([www.empirix.com](http://www.empirix.com)). They used this tool to create six scenario scripts that simulated six distinct user paths. At time during testing, they gradually ramped the number of virtual users running these scripts up to 2,000, exercising the database servers to their maximum capacity. In each case, they charted the results up to the point where the throughput leveled off or declined, or the number of errors exceeded 10 percent. For this

test the compared Microsoft SQL Server 2000 Enterprise edition Service Pack 2, MySQL 4.0.1 Max, DB27.2 Fix Pack 5, Oracle9i Enterprise Edition and Sybase ASE 12.5.0.1. Measurable quantities for this load testing process for all organisations were response times, throughput rates, resource utilization (that included CPU usage, Virtual Memory, Disk I/O subsystem and Network latency utilisation).

As shown by Figs 3 and 4 below [6] PC Magazine in this case used Microsoft Excel for analysis, presentation and drawing charts for displaying to results. Fig 3 shows that Oracle 9i and MySQL tops the throughput chart while MS SQL server show almost constant throughput after 300 users.

### Oracle9i and MySQL topped throughput tests

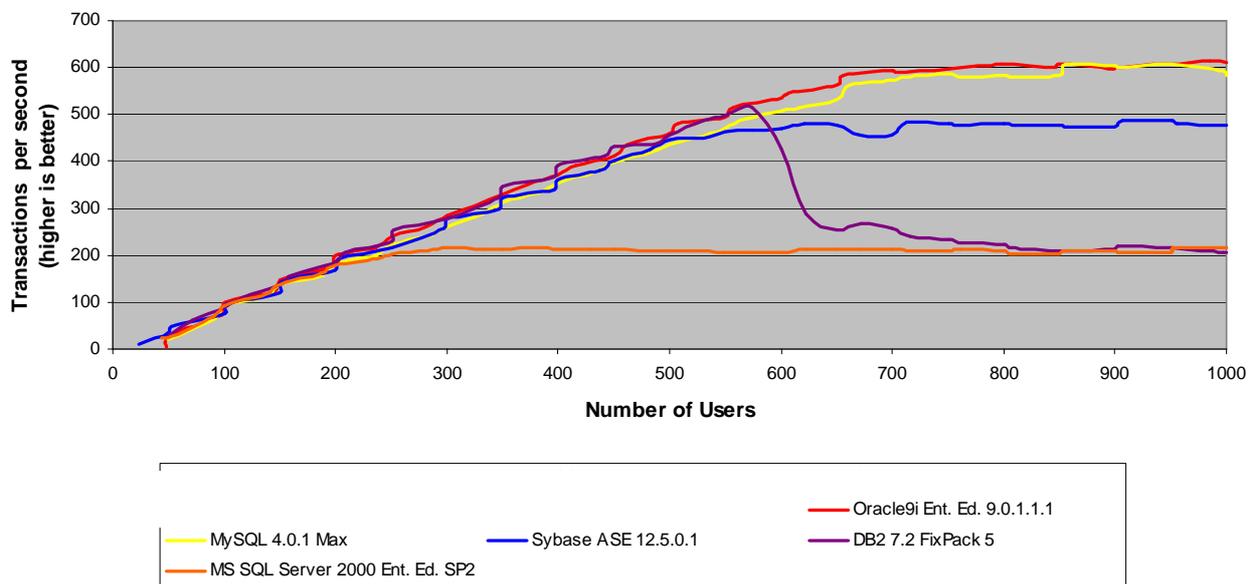


Fig 3 Throughput chart (June 2003 PC Magazine and [8] *eWEEK* publishers)

Throughput Definition: Throughput is the one of transactions per second performed by the application server. Number of users is number of concurrent Web clients driving the load.

Configuration details: Tests conducted on an Hewlett-Packard Co. HP Net Server LT 6000r with four 700 MHz Xeon CPUs, 2 GB of RAM, a gigabit Ethernet Intel PRO/1000 F Server Adapter and 24 9.1 GB Ultra3 SCSI hard disks used for database storage.

Fig 4 below shows that in Oracle9i response time increases with increasing users making it a favourite choose for Enterprise solution as Burleson would say.

## Oracle9i and MySQL offered the fastest response times

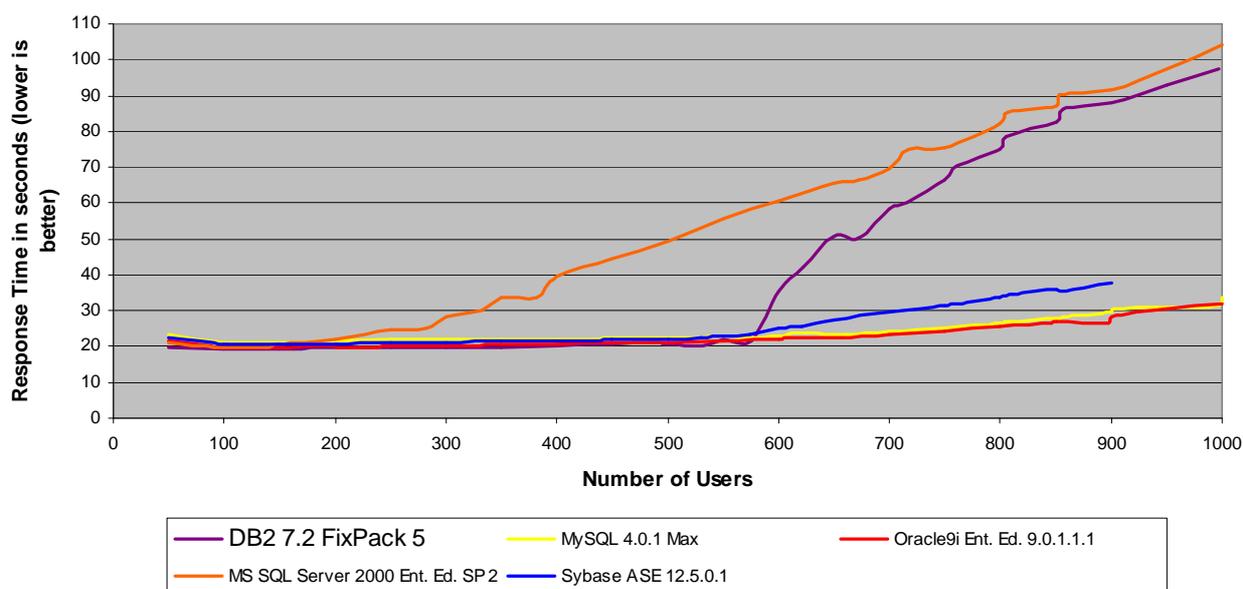


Fig 4 Response time Chart (June 2003 PC Magazine and [8] *eWEEK* publishers)

### Response time Definition used:

Response time is the time to complete the six bookstore user action sequences, weighted by frequency of each sequence in the mix. Number of users is number of concurrent Web clients driving the load.

Configuration details: Same as for throughput above.

Another approach taken by Quest to achieve load and stress testing uses a whole lot of benchmarks embedded under automation tools like Quest Central, SQL Tuning tools, Performance Monitoring toolkit provided with these databases. They used Load generator to design workload models and simulate virtual users. Measurable quantities for stress testing were synchronization issues, race conditions, memory leaks and loss of data during execution. For load testing counters they used were similar to those used by [6] PC Magazine explained above. For results analysis and presentation they used Quest produced Spotlight and Statspack provided with Oracle9i suit. Derivatives were in the form of tables, diagrams and Graphs showing throughput versus user load, response time versus user load just as done by [6] PC Magazine, 2003/ [8] *eWEEK* test above and resource utilisation versus load.

### Conclusion

Performance and scalability testing as seen in the above discussion is a very tough and expensive exercise requiring participation of many experts. Database testing as motivated by performance tuning and database diagnosis is one of the most controversial topics in the database fraternity. This is maybe due to the great need for efficiency as the user population accessing the database

increases. This project research seek to combine all these views and performance testing strategies as well as coming up with an explanation as to which of the two DBMS Oracle 9i and SQL server 2000 is the best for enterprise solution and give technical support the results.

## **LIST OF FIGURES**

Fig 1<sup>6</sup> Some considerations for performance tests

Fig 2 Concurrency Model

Fig 3 Throughput chart (June 2003 PC Magazine and *eWEEK* publishers)

Fig 4 Response time Chart (June 2003 PC Magazine and *eWEEK* publishers)

## **REFERENCE LIST**

[1] Donald Burleson, 2002, 'Oracle Guru' Database Benchmark Wars: What you need to know, Last updated 05/08/2002, Access date 15/05/05. [http://www.dba-oracle.com/art\\_db\\_benchmark.htm](http://www.dba-oracle.com/art_db_benchmark.htm)

[2] Floyd Marinescu, 2000, Tips on Performance Testing and Optimization. Access date 08/05/05  
<http://www.theserverside.com/tss>

[3] MSDN, 2005, SQL Server architecture SQL Server 2000 Documentation, Access date 12/05/05.  
[http://msdn.microsoft.com/library/default.asp?url=/library/en-us/startsql/getstart\\_4fht.asp](http://msdn.microsoft.com/library/default.asp?url=/library/en-us/startsql/getstart_4fht.asp)

[4] Mysql.com, 1995-2005. Database Server Feature Comparisons, Access date 08/05/05.  
[http://dev.mysql.com/tech-resources/crash-me.php?res\\_id=100](http://dev.mysql.com/tech-resources/crash-me.php?res_id=100)

[5] Oracle9i Database List of Books (Release 2(9.2)).Documentation  
<http://otn.oracle.com/pls/db92/db92.docindex?preference=All>

[6] PC Magazine, 2003. Performance Tests: The Big SQL Shootout. Last updated 26/03/02, Access date 12/05/05. <http://www.pcmag.com/article2/0,1759,1171144,00.asp>

[7] Quest.com, 2005. Quest Database Management Products, Access date 12/05/05.  
[http://www.quest.com/database\\_management/](http://www.quest.com/database_management/)

---

<sup>6</sup> Table from Loadtest.com.au – shows performance testing prerequisites  
[http://www.loadtest.com.au/types\\_of\\_tests/performance\\_tests.htm](http://www.loadtest.com.au/types_of_tests/performance_tests.htm)

[7] Technovations, Inc, 2005. Load Testing Products, Access date 02/05/05.

<http://www.technovations.com/loadtest.htm>

[8] Timothy Dyck (eWEEK Magazine) July 2003, Server Databases Clash. Last Updated July 2003  
Access date 06/05/05. <http://www.eweek.com/article2/0,4149,293,00.asp>

[9] Transaction processing Performance Council, 2005, Top Ten TPC-C by Performance, Access  
date 14/05/05. [http://www.tpc.org/tpcc/results/tpcc\\_perf\\_results.asp](http://www.tpc.org/tpcc/results/tpcc_perf_results.asp)

[10] RPM Solutions Pty Ltd, Last Updated July 2003 Access date 06/05/05  
[http://www.loadtest.com.au/types\\_of\\_tests/performance\\_tests.htm](http://www.loadtest.com.au/types_of_tests/performance_tests.htm)