A Comparison of Oracle Database 10gR2 Real Application Cluster Performance on Microsoft Windows 2003 Server Enterprise Edition x64 and Red Hat Enterprise Linux x86_64

By Larry Pedigo
Performance Tuning Corporation
October 2006

Contents

Introduction

Oracle® Database 10g Enterprise Edition is currently one of the most widely deployed database software packages in the world. Since its introduction in 1977, Oracle database software has been ported to a variety of hardware platforms and Operating Systems. While some database vendors only support one platform or offer different database versions on each platform, Oracle offers a consistent product across all major Operating Systems. There are several reasons that Oracle Corporation supports so many different OS versions, including:

- Oracle recognizes that their customers have different OS install bases.
- Mission-critical applications may require a particular OS.
- Experience and training of Administrators at different companies may favor different Operating Systems.

Corporations have responded by implementing the majority of their mission-critical, Enterprise class databases with Oracle database software, across a wide variety of Operating Systems. Regardless of the Operating System, there are several factors that make Oracle an attractive choice for Enterprise databases:

- Superior performance characteristics, especially for large databases.
- A variety of High Availability features.
- Oracle Real Application Cluster technology, which enhances both performance and availability, and is unique to Oracle.

Two Operating Systems in particular are being installed on increasing numbers of corporate servers every year: Linux and Microsoft® Windows Server. Both Operating Systems take advantage of readily available and relatively inexpensive server hardware based on Intel/AMD processors.

Of course, Oracle Database 10g software is available for both Operating Systems. Many people are aware that Oracle database software is available for Linux. The ties between Oracle and the Linux OS have been well publicized in the press. Less people are aware of Oracle's history and platform support for Microsoft Windows Server. Actually, the port of Oracle database software to Microsoft Windows Server (then called Windows NT) pre-dates the port of Oracle to Linux. Both Linux and Microsoft Windows Server 2003 are primary development platforms for Oracle, and both are fully supported by Oracle. A Gartner report issued in May, 2006 indicates that the shares of 2005 Oracle RDBMS revenue were approximately equal for the two Operating systems: 18% for Microsoft Windows and 18% for Linux.

There is a lot of published literature promoting the advantages of migrating from Oracle (or other databases) on legacy Operating Systems to Oracle on Linux. The Total Cost of Ownership is often cited as a reason to make the change. Other studies have examined the potential performance gains to be realized by migrating to Oracle on Linux.

Although it is less well documented, there are at least as many compelling reasons for companies to consider deploying their Oracle databases on Microsoft Windows Server 2003. Microsoft Windows Server 2003 is widely used in corporate data centers and is the primary server OS for many Oracle customers. A commonly cited reason for installing Oracle on Microsoft Windows servers is that Windows Server OS administrative expertise is readily available. Studies have shown that the Total Cost of Ownership for Microsoft Windows Server 2003 is quite attractive – primarily due to integrated support for terminal services, directory services and storage management. In addition, the Enterprise Edition of Microsoft Windows Server 2003 offers the scalability that Oracle customers need in their data centers.[1] The 32-bit version of the Enterprise edition of Microsoft Windows Server 2003 supports up to eight processor sockets and 64 GB of RAM. The 64-bit versions support up to 1 TB of RAM. Many Oracle customers are beginning to migrate their databases from 32-bit Windows Server to 64-bit Windows Server 2003, and are reporting impressive performance gains.

It is to be expected that anyone contemplating a migration from a legacy Operating System to either the Linux or Windows platform would require some hard evidence that the new platform will meet their requirements. As stated above, there is abundant literature and many published whitepapers about the performance characteristics of Oracle on Linux. Unfortunately, the availability of literature for Oracle on Microsoft Windows Server 2003 is limited, particularly for 64-bit versions of Oracle on Microsoft Windows Server 2003. While there is marketing information available for both the Linux and Windows versions of Oracle, there is little information available that would allow Managers and DBAs to evaluate the technical merits of both platforms.

This whitepaper addresses the need for technical documentation by providing an objective comparison of Oracle performance on 64-bit Red Hat® Linux and 64-bit Microsoft® Windows Server 2003. This study was designed as an "Apples-to-apples" comparison: the same type of server was used for Linux and Windows, the same processors, the same memory, the same storage configuration, and optimal tuning for each database. The tool that was used to perform these tests was SwingBench, an Oracle Real Application Cluster testing and demo tool from the Oracle UK Database Solutions Group (not officially supported by Oracle). The purpose of this whitepaper is not to promote either the Linux OS or the Microsoft Windows Server 2003 OS, but to provide Managers and Administrators with a factual basis for making decisions about which OS platform to use for implementing their Oracle databases.

32-bit Processing vs. 64-bit Processing

One of the keys to understanding the performance characteristics of Oracle on Linux and Microsoft Windows Server 2003 is to understand the impact of 64-bit processing. Both Linux and Microsoft Windows Server 2003 are currently available in both 32-bit and 64-bit versions. On a particular server, the opportunity to use a 64-bit OS version versus a 32-bit OS version is closely tied to the type of processor that is used on that server. The original generations of Intel Xeon® and AMD processors supported only 32-bit OS versions. However, the Intel Itanium® and Itanium2® processors have been available for several years now, and support 64-bit Operating Systems and processes. More recently, the AMD Opteron® and Intel EM64T® processors have been widely utilized. These processors support either 32-bit or 64-bit OS versions.

---

[1] http://www.microsoft.com/windowsserver2003/evaluation/overview/enterprise.mspx

Of course, Oracle Database 10g software is available for all "flavors" of the Linux and Microsoft Windows Server Operating Systems, including both 32-bit and 64-bit versions. However, the 64-bit versions for the Opteron/EM64T architectures have only been available for the last year or so. This includes new 64-bit versions of Oracle Database 10g R2, Oracle Database 10g R1 (Linux only), and Oracle Database 9.2 (Linux only). As a result, many Oracle implementations at customer sites are still 32-bit versions.

There are several well known performance limitations for 32-bit Oracle Database implementations. Many of these performance limitations are associated with Oracle memory usage. With a 32-bit OS, usable Oracle memory is split into "low" memory (< 3 GB) and "high memory" (up to 64 GB). High memory access requires the use of special techniques, such as Paging Address Extensions and Address Windowing Extensions. This adds overhead, making high memory access considerably less efficient and slower than low memory access. Note that approximately the same effect is observed for 32-bit Linux and 32-bit Microsoft Windows Server.

In addition to the high memory performance problems, the 32-bit Oracle Database has inherent limitations in the number of simultaneous user connections that can be achieved. Some Oracle memory types are restricted to the low memory area, such as the Shared Pool, the Large Pool, the Program Global Area, and User Process memory. Unfortunately, these are the Oracle memory types that are required to support large numbers of simultaneous user connections.

For a 32-bit Oracle Database on Linux, the impact of the user connectivity problem is that a limited, but workable number of users can be connected at any given time. Up to 3 GB of low memory is usable for connectivity. With Shared Server connections, up to a few thousand simultaneous user connections can be supported.

For a 32-bit Oracle Database on Windows, the number of simultaneous user connections is severely limited. Once the overhead for high memory access is subtracted, little memory is left for supporting user connections. By default, only 1 GB of memory for user connections is available. This will allow about a thousand simultaneous Shared Server user connections. Memory available for this purpose can be stretched to 2 GB, resulting in up to 2,000 simultaneous user connections; but only by making compromises that result in performance limitations.

Fortunately, the 64-bit Oracle Database overcomes all of these 32-bit performance limitations, and "levels the playing field" for both Linux and Microsoft Windows Server. 64-bit Operating Systems feature a flat memory model up to 1 TB, with no performance penalty for large memory usage. The number of simultaneous user connections is limited only by available memory.

As mentioned above, there are two processor architectures that currently support 64-bit processing: the Intel Itanium II processor architecture and the Intel EM64T/AMD Opteron processor architecture. The Intel Itanium II processor has been designed for high efficiency 64-bit processing (with limited 32-bit capabilities). Currently, the Itanium II processor fills a niche for high performance servers. It is particularly useful for high-end Oracle Database servers. The Intel EM64T and AMD Opteron have been designed for mixed 32-bit/64-bit environments. These processors offer good performance for both 64-bit and 32-bit applications, and are currently achieving mass-market acceptance. Figure 1 summarizes the role that the different processor architectures are filling in the migration to full 64-bit processing.

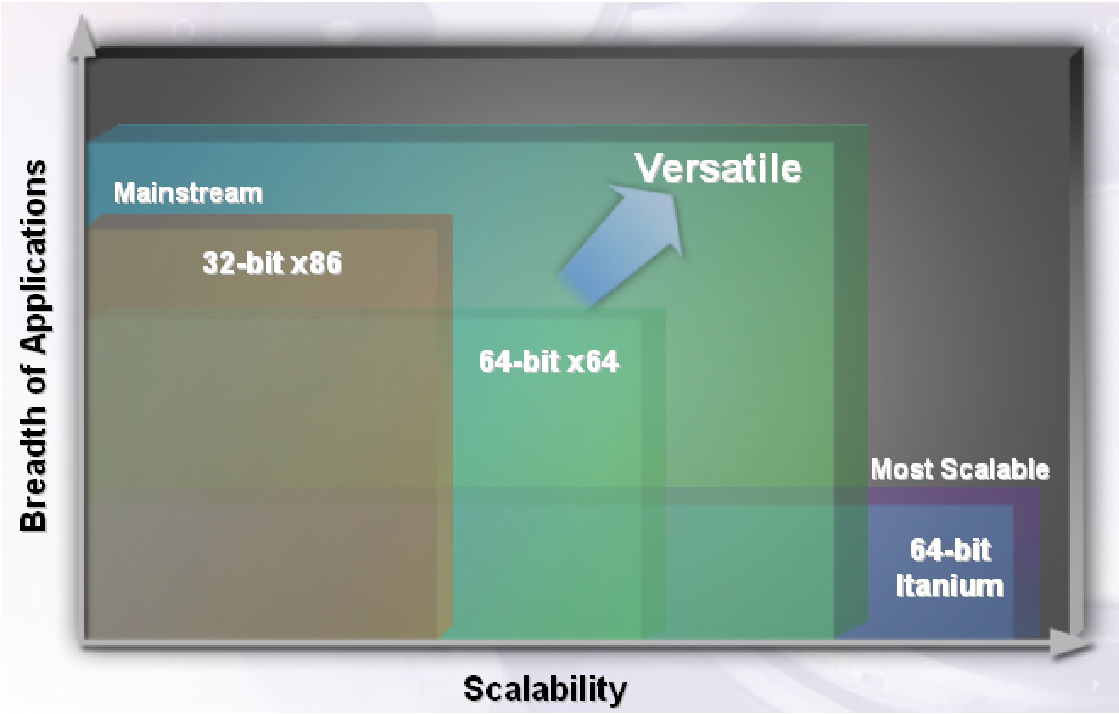

Figure 1: Migration path to 64-bit processing

For the purposes of this study, it was decided to use Intel EM64T processors to compare Oracle Database performance on Windows and Linux. These processors are representative of the architecture that the majority of Oracle migrations to the Intel/AMD platform are moving towards. In addition, this architecture is not expected to offer any clear technical advantage to either Linux or Microsoft Windows Server 2003.

It should also be noted that performance for a 64-bit Oracle Database on Linux is not necessarily exactly the same as a 64-bit Oracle Database on MS Windows. Windows uses a multi-threaded processing model. This means that only one process is implemented for Oracle, although the Oracle process contains multiple threads. In theory, a fully multi-threaded application should be highly efficient. In contrast, Linux uses a multi-process processing model. There are multiple background processes visible on the database server (PMON, SMON, etc.). Multi-threading is only utilized for certain components, such as the Multi-Threaded Server. The two processing models may behave differently in some scenarios.

Testing Environment

The goal of the setup phase of this project was to install two 2-node Oracle Database 10g Real Application Clusters, one cluster with the Red Hat Enterprise Linux Advanced Server OS, and one cluster with the Microsoft Windows Server 2003 OS. External Fibre Channel SAN storage was used for the shared disk component of the clusters. In addition to the database servers, two client machines were used for running SwingBench, the benchmark software. Every effort was made to ensure that identical hardware and software was used for both clusters. The following sections detail the hardware, software and installation procedures used for testing the two Oracle Real Application Clusters on the Red Hat Linux and Microsoft Windows Server 2003 Operating Systems

Hardware Components

Hardware components of the two Oracle clusters include:

- Database servers
- Client Computers
- Network Switches
- Storage array
- SAN Switch

Figure 2 is a schematic diagram of the hardware architecture.

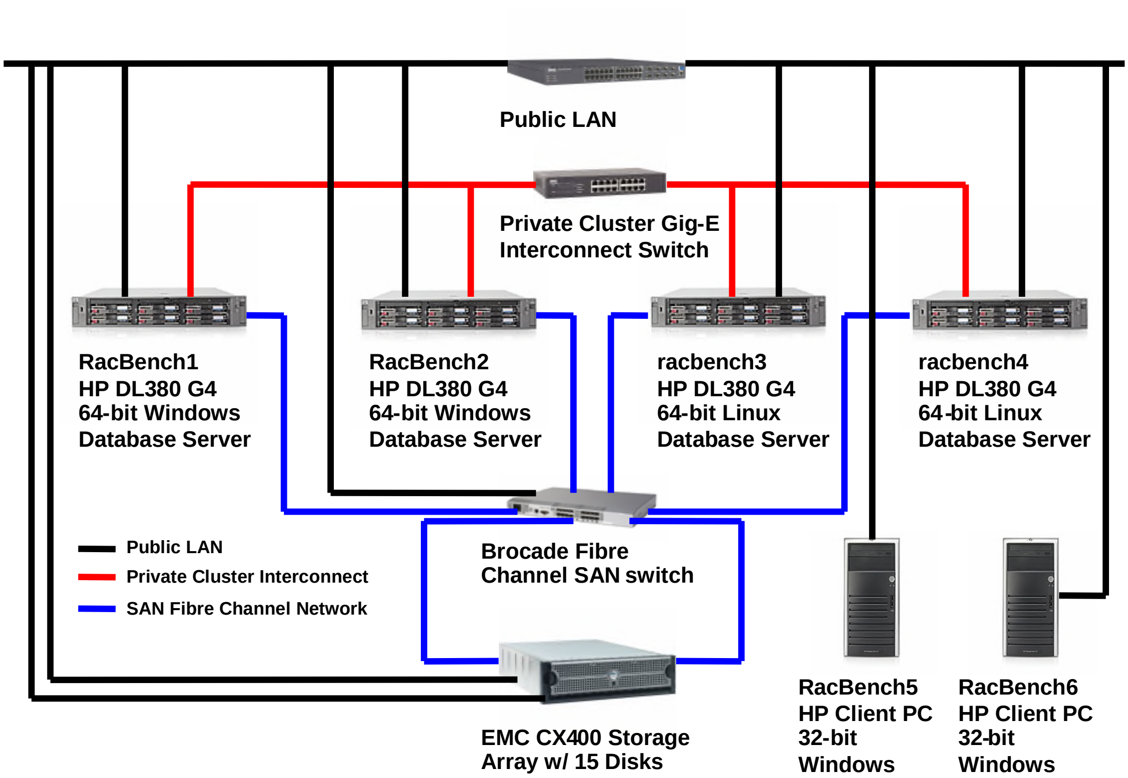

Figure 2: Schematic diagram of test hardware

Database Servers

For an Oracle Real Application Cluster system, the choice of server hardware and the choice of specific server hardware options are critical to the performance of the cluster. Before acquiring servers for these tests, a set of minimum requirements was defined. In general, server class systems with fast processors are preferred, with at least two processors per server. For the purposes of these tests, it was required that the processors support 64-bit Operating Systems. In addition, to take advantage of 64-bit memory handling, at least 8 GB of RAM was required. It is also a requirement for Oracle RAC that each server has at least one connection to external shared storage. In this case, it was decided that a single connection from each server to the SAN was feasible, since the goal was to keep the I/O subsystem configuration as simple as possible. At least two network connections are required per server: one for the public LAN and one for a private interconnect between servers. The private interconnect interface must be capable of supporting full Gig-E speed at full duplex.

For the purposes of these tests, it was important that the servers for both the Linux and Windows clusters be configured as identically as possible. The following configuration options were chosen for the RAC clusters:

- Four HP® ProLiant DL380 G4
    - Two Windows servers
        - RacBench1
        - RacBench2
    - Two Linux servers
        - racbench3
        - racbench4
- Two Intel EM64T Xeon processors per server (4 logical processors with Hyperthreading), 3.4 GHz
- Two 36 GB SCSI disks per server, configured as RAID 1
- 8 GB RAM per server
- 8 GB swap space/paging file per server
- Two Gigabit NICs per server
- One Qlogic 2340-E HBA per server

Client Computers

The SwingBench benchmark tools used for this set of tests require one or more client computers to host the benchmark software. A minimum of one processor per client computer is required, as well as a minimum of 1 GB of RAM. However, initial tests indicated that a computer configured to the minimum specifications could not support the full workload anticipated for these tests. Fortunately, SwingBench supports a distributed testing mode, where multiple clients can be coordinated to perform a single test. The decision was made to use two minimally configured clients, which proved to be sufficient for even the most demanding tests.

The client computers were configured as listed below:

- Two HP Workstations with Windows 2003 Server Enterprise Edition x86 (32-bit) installed
    - RacBench5
    - RacBench6
- A single Intel Xeon processor (2 logical processors with hyperthreading), 3.2 GHz
- A single internal 112GB ATA disk
- 1 GB RAM per server
- 2 GB paging file per server

- One Gigabit NIC per server

Network Switches

Two network switches are required for the Oracle RAC systems. One Network Interface Card on each server is connected to a switch on the Public LAN. The sub-net used for the servers can be situated behind a firewall, as long as all of the Oracle RAC servers can access each other. This is also the network interface that will be used to connect to client applications. The speed of this interface should be 100Mb or better at full duplex.

The second NIC on each sever should connect to either a dedicated private switch or to a private VLAN. This interface is used for the cluster interconnect, which is used to implement Oracle RAC Cache Fusion. No servers other than the Oracle RAC servers should be configured in this network segment. The speed of this interface must be Gig-E, at full duplex. A non-routing network address is recommended.

It is also recommended to add an additional NIC per server for the Cluster Connect. The two private interfaces per server can then be logically combined into a single private interface via network teaming (also called network bonding). This provides port redundancy and increases potential throughput.

For this test configuration, the NIC for the Public LAN interface on each server was connected to a Gig-E network switch. The NIC for the Private interface was connected to a secondary network switch, also at Gig-E speed. Although other non-RAC servers were physically connected to the secondary switch, the ports for the Private Cluster Interconnect were isolated on a Private VLAN. To prevent possible conflicts between the Interconnect interfaces for the Linux cluster and the Windows cluster, only one of the two RAC databases was allowed to be online at any given time. A third NIC was not available on these servers, so no network teaming or bonding was configured.

Storage Array

Oracle Real Application Cluster technology requires access to external shared disks from each RAC node. Typically, this is accomplished with a Fibre Channel Storage Area Network (SAN), although Network Attached Storage (NAS) or iSCSI storage may also be used for this purpose. A general rule-of-thumb is that an external storage array should contain at least ten physical disks allocated for Oracle RAC in order to achieve any reasonable performance.

The following external storage array configuration was used for this set of tests:

- EMC® Clariion CX400 Storage Array
  - Two Storage Processors, each connected with its own dedicated Fibre Channel Port to the SAN
  - Two Ethernet network interfaces, one per Storage Processor
  - Fifteen Fibre Channel Disks, 72 GB each (unformatted)

SAN Switches

When more than two servers are connected into the Storage Array, one or more SAN switches are required between the servers and the Storage Array. If multiple HBAs are installed in each server, two switches can be used to provide Fibre Channel path redundancy and failover capabilities.

In this case, only one HBA per server was configured, and only one switch was utilized. The switch was a Brocade DS8B2, with eight ports supporting 2 Gb speed.

Software Components

There is some form of software associated with each hardware component of an Oracle RAC system, whether it is database software, application software, device drivers, or even low-level components such as firmware and BIOS. One of the challenges in configuring an Oracle RAC system is that all of the software must be completely interoperable. If one application, device driver, firmware version, etc. does not work well with the other components, the Oracle RAC system may fail to function. For this reason, it is essential that all software components should be vendor certified for interoperability. One of the implications of this approach is that the newest available version of each software component should NOT be automatically installed. Rather, the latest version that has been tested and certified for Oracle RAC on the specific hardware you will be using should be installed.

To ensure a successful installation, a list of certified software versions was compiled from HP Parallel Database Cluster Kit documentation, EMC Support Matrices, and Oracle's RAC Certification lists. The following sections list the software components and versions that were used for the tests.

Database Servers

The first step in configuring software on the Oracle RAC database servers was to install and configure the Operating System. In the case of the Linux servers, the HP® Parallel Database Cluster Kit provided scripts to assist with OS installation and configuration. These scripts were augmented with additional steps from Oracle's RAC Installation manuals. In the case of the Windows servers, the OS was installed and configured according to previously published Installation Guides by HP, Oracle, and Performance Tuning Corporation.

Another key step was to update all device firmware, BIOS, and drivers to the versions recommended by the vendor Support Matrices. In most cases, this was done via a simple driver installation package or Red Hat rpm file. In the case of the SAN software for the database, the Qlogic® SAN Surfer software package was installed on both Windows and Linux servers. SAN Surfer was used to update the Qlogic HBA BIOS, modify the BIOS settings, and install the Qlogic driver on each server.

One of the concerns in planning this project was in how best to achieve identical I/O response for the Linux Oracle RAC servers and the Windows Oracle RAC servers. The decision was made to keep the I/O software stack as simple as possible. For this reason, only the Qlogic driver was used to connect the server HBAs to the EMC Clariion Storage Array. Neither the EMC Navisphere® Agent nor EMC PowerPath® software was configured on these servers. The concern was that the additional software layers could affect I/O latency differently for the Linux implementation of the software when compared to the Windows version of the software. By using only the minimum required I/O software, I/O equivalency between the two Oracle RAC systems could be assured.

The following software was configured on the Database servers:

- Server BIOS HP P51 04/26/2006
- Operating Systems
  - Red Hat Enterprise Linux Advanced Server 4.0 x86_64 Update 2, kernel 2.6.9-22.0.1.ELsmp
  - Microsoft Windows Server 2003 Enterprise Edition x64 v. 5.2.3790 R1 SP 1
  - The newest available OS versions were NOT used; OS versions chosen to comply with documented HP and industry best practices.
- HP Parallel Database Cluster Kit for Linux software
  - HP approved technique for configuring Linux RAC servers
  - Modified to work with EMC storage
  - Additional configuration guided by Oracle RAC Best Practices
- Qlogic SAN Surfer software v.2.0.30
- Qlogic HBA BIOS 1.47
- Qlogic HBA driver – qla2xxx-v8.01.00-3dkms
- SwingBench 2.3 CPU monitor agent (Linux only)
- Oracle 10g R2 Clusterware and Database software

Client Computers

The main function of the client computers is to run the SwingBench software. SwingBench software is available for either the Linux OS or the Microsoft Windows OS. To provide consistency between tests, the 32-bit version of Microsoft Windows 2003 Server was installed on both client systems.

As prerequisites to installation, SwingBench requires that the 1.5 version of the Java J2RE be installed, and that some form of Oracle client is installed. On these computers, the Oracle Instant Client for Windows was installed. The following software versions were installed on the client computers:

- Server Bios BIOS Intel Corp. BZ87510A.86A.0125.P34.0503312215, 3/31/2005
- J2RE 1.5.0.06
- Oracle Instant Client 10g R2
- SwingBench 2.3 software

Network Switches

The only software utilized on the network switches was the native VLAN software on the switch used for the private interconnect network. The ports used for Oracle RAC were isolated from the other ports with the VLAN.

Storage Array

The EMC Clariion Storage Array can be configured through web-based Navisphere Manager software, which is pre-installed on the Clariion system. The first step in software configuration was to update the EMC Flare Code version (the OS for the Storage Array). This step was performed by an EMC technician.

The next step was to use EMC Navisphere Manager software to configure the Storage Array. First, global properties such as Read and Write Cache settings were configured. For the second step, hosts were connected to the Storage Array and registered with Navisphere (this requires that Fibre Channel switch zoning be performed; see below). Since Navisphere Agent software wasn't installed on the Database servers, registration had to be performed manually, using the HBA World Wide Name to identify the hosts.

The next step was to provide RAID protection for the disks. Using Navisphere Manager, fourteen disks were used to create seven RAID 1 mirrored pairs of disks. In this case, hardware-based RAID 10 striping was not used, since striping was to be provided by Oracle ASM software. The extra disk was configured as a Hot Spare.

The RAID groups were then sub-divided into LUNs (Logical Unit Numbers). The LUNs are what the hosts see as "disks". Half of the LUNs were for Linux, and half were allocated for Windows. On each RAID group, the first LUN configured is placed on the outside of the disks. This position is preferred for best performance. To keep the I/O performance consistent between the two RAC systems, half of the disks had Windows LUNs on the outside disk sectors, and half had Linux LUNs on the outside sectors.

The final storage configuration step was to create Storage Groups. Storage groups associate LUNs with specific hosts, allowing those hosts to "see" the LUNs and denying access to other hosts.

The following list summarizes the EMC Clariion Storage Array configuration:

- Flare Code 02.19.400.07
- Navisphere Manager software
  - Seven RAID 1 groups (+ one hot spare disk)
  - Split into sixteen LUNs
    - Five LUNs for Oracle ASM data files on Linux ~ 30 GB each
    - Two LUNs for Oracle ASM REDO logs on Linux, one per RAC thread ~ 30 GB each
    - One LUN for Oracle RAC clusterware on Linux ~ 2 GB
    - Five LUNs for Oracle ASM data files on Windows ~ 30 GB each
    - Two LUNs for Oracle ASM REDO logs on Windows, one per RAC thread ~ 30 GB each
    - One LUN for Oracle RAC clusterware on Windows ~ 2 GB
    - LUN positions alternate on each disk so that half the disks have Linux LUNs on the outside sectors (preferred position for performance), and half the disks have Windows LUNs on the outside sectors
  - Two Storage groups

- Linux_Storage_Group
  - All eight Linux LUNs
  - Both Linux servers
- Windows_Storage_Group
  - All eight Windows LUNs
  - Both Windows servers

Figure 3 shows the distribution of RAID groups and LUNs on the array. Table 1 shows details for the RAID groups and LUNs.

Disk #	0000	0001	0002	0003	0004	0005	0006	0007
RAID Group	RG 0 RAID 1 60 GB		RG 1 RAID 1 60 GB		RG 2 RAID 1 60 GB		RG 3 RAID 1 66 GB	
Enclosure ID Bus_0, Shelf_0	LUN 0 30GB Linux		LUN 2 30GB Win		LUN 4 30GB Linux		LUN 6 2GB Linux	
	LUN 1 30GB Win		LUN 3 30GB Linux		LUN 5 30GB Win		LUN 7 32GB Win	
							LUN 8 32GB Linux	
Disk #	0008	0009	0010	0011	0012	0013	0014	
RAID Group	RG 4 RAID 1 66 GB		RG 5 RAID 1 66 GB		RG 6 RAID 1 66 GB		200	
Enclosure ID Bus_0, Shelf_0	LUN 9 32GB Linux		LUN 11 33GB Win		LUN 13 33GB Linux		Global Hot Spare	
	LUN 10 32GB Win		LUN 12 33GB Linux		LUN 14 33GB Win			
	LUN 15 2GB Win							

Figure 3: LUN Map for the EMC Clariion Storage Array

Raid Group	Disk Size	RAID Level	LUN #	LUN Size	SP	Storage Group	Hosts	Data Type
0	73	1	0	30	A	Linux_Storage_Group	racbench3, racbench4	ASM1 Linux +DATADG
0	73	1	1	30	B	Windows_Storage_Group	RacBench1, RacBench2	ASM1 Win. +DATADG
1	73	1	2	30	B	Windows_Storage_Group	RacBench1, RacBench2	ASM2 Win. +DATADG
1	73	1	3	30	A	Linux_Storage_Group	racbench3, racbench4	ASM2 Linux +DATADG
2	73	1	4	30	A	Linux_Storage_Group	racbench3, racbench4	ASM3 Linux +DATADG
2	73	1	5	30	B	Windows_Storage_Group	RacBench1, RacBench2	ASM3 Win. +DATADG
3	73	1	6	2	A	Linux_Storage_Group	racbench3, racbench4	Linux Quorum
3	73	1	7	32	B	Windows_Storage_Group	RacBench1, RacBench2	ASM4 Win. +DATADG
3	73	1	8	32	A	Linux_Storage_Group	racbench3, racbench4	ASM4 Linux +DATADG
4	73	1	9	32	A	Linux_Storage_Group	racbench3, racbench4	ASM5 Linux +DATADG
4	73	1	10	32	B	Windows_Storage_Group	RacBench1, RacBench2	ASM5 Win. +DATADG
4	73	1	15	2	B	Windows_Storage_Group	RacBench1, RacBench2	Windows Quorum
5	73	1	7	33	B	Windows_Storage_Group	RacBench1, RacBench2	ASM6 Win. +REDOADG
5	73	1	11	33	A	Linux_Storage_Group	racbench3, racbench4	ASM6 Linux +REDOA
6	73	1	9	33	A	Linux_Storage_Group	racbench3, racbench4	ASM7 Linux +REDOB
6	73	1	10	33	B	Windows_Storage_Group	RacBench1, RacBench2	ASM7 Win. +REDOBDG
200	73	HS	200	66	N/A			Global Hot Spare

Table 1: EMC Clariion Storage Array RAID groups and LUNs

SAN Switches

The Brocade FabricView Switch Explorer web interface was used to configure zoning on the Brocade DS8B switch. Zoning functions on a Fibre Channel switch similarly to the way that VLAN software functions on a network switch. The zoning procedure is:

- Create Aliases for the World Wide Names presented by the host HBAs and the storage array storage processor ports.
- Create a new Zone name.
- Drag and drop an alias for an HBA and a storage array port into each zone.
- Create a new Zone Configuration name.
- Drag and drop all of the defined zones into the Zoning Configuration.
- Enable the Zone Configuration.

Testing Procedures

The main tool used for these tests was SwingBench. SwingBench is a set of benchmark/test/demo Java based programs originally developed by the Oracle U.K Database Solutions group for Oracle Real Application Clusters. Although it is not officially supported by Oracle, SwingBench is freely available via download from http://www.dominicgiles.com/swingbench.php . SwingBench supports three different schemas: Calling Circle, Order Entry, and Sales History. The Calling Circle schema is designed as a simple OLTP test, usually used for single node tests. The Sales History schema is a DSS (Data Warehouse) benchmark tool. The Order Entry schema emulates a classic TPC-C OLTP benchmark, and is designed to work with Oracle Real Application Clusters. The Order Entry schema was used for the tests discussed in this whitepaper.

SwingBench offers three interfaces for testing from a single client machine. The default SwingBench interface is a Java-based utility that presents rich graphics and an interactive interface. The MiniBench interface is a simple interface with minimal graphics that conserves resources on the client machine. The charbench interface is completely command line oriented and presents output in an xml file.

Initial testing with each of the three interfaces indicated that due to memory and Java limitations, the entire test suite could not be run from one of the available client computers. To workaround this limitation, the SwingBench Coordinator (Figure 4) was used to submit tests with charbench across two client computer nodes. This produces two xml file outputs (one per RAC node, in this case) which have to be manually combined. However, the tests can be initiated and monitored with a GUI tool, ClusterOverview (Figure 5).

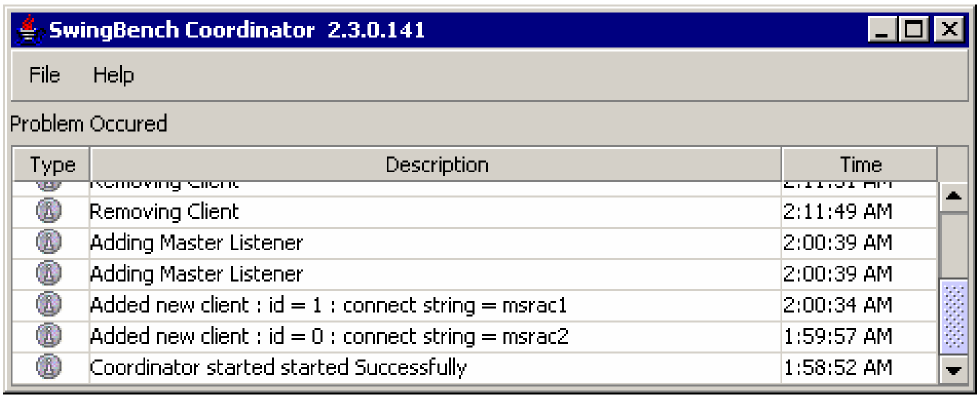

Figure 4: The SwingBench Coordinator

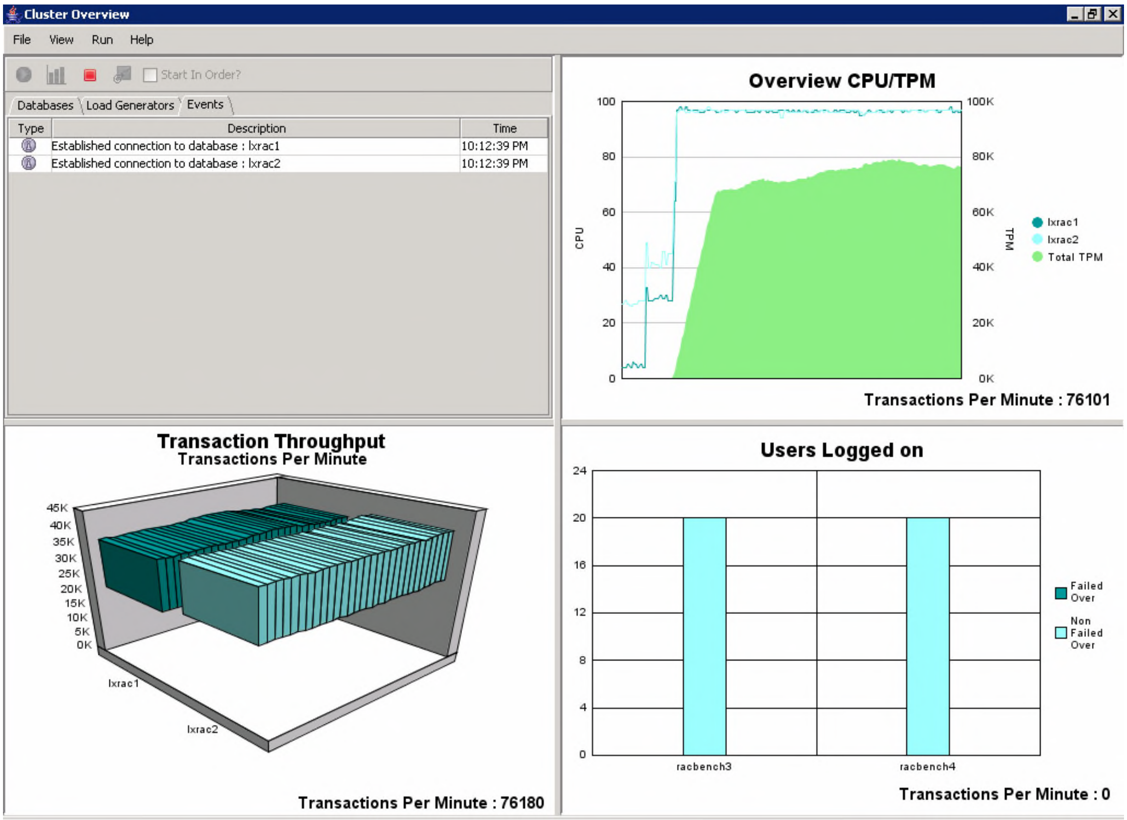

Figure 5: The SwingBench ClusterOverview tool

Two test suites were run against both the Linux Cluster and the Windows Cluster: a “Stress Test” and a “User Load Test” (both designed by the author). Both tests utilize the Order Entry schema. To make sure that I/O capabilities were tested against realistically sized tables and indexes, approximately 30 GB of data was pre-loaded with batch routines. Both tests were scheduled to run for ten minutes per test. To prevent caching of table data in memory between tests, the databases was restarted between each test.

The Stress Test was designed to use relatively small numbers of user sessions, each session “pushing” as many transactions as possible per unit of time. To accomplish this, the Maximum and Minimum Think Times were reduced to zero. Think time is meant to simulate users waiting between each transaction in order to “think” about the next step. The actual time between any two transactions is randomly chosen between the Minimum and Maximum Think Times. With Think Times set to zero, transactions are submitted back-to-back. This is guaranteed to fully stress the CPU and I/O capabilities of the system.

The User Load test was designed to simulate how a system will respond to supporting a relatively large number of user connections. For this test, a combination of a Minimum Think Time of 1 second and a Maximum Think Time of 5 seconds was used. This is a reasonable approximation of real-world work load. In addition, Shared Server connections were used to maximize the number of connections while conserving memory.

To achieve the goal of making the comparison between Oracle RAC on Linux and Microsoft Windows Server 2003 as objective as possible, Oracle tuning parameters had to be carefully managed. A target was set for the Stress Test to optimize performance for 60 users (30 users per Oracle RAC instance). Initial test indicated that above 60 users, contention levels increased, and efficiency decreased. The target set for the User Load test was to optimize performance for 3000 users (1500 users per RAC instance), which initial tests indicated was an achievable level of use activity for both Linux and Windows.

To derive an initial set of tuning parameters for each version of Oracle, the SGA_TARGET and PGA_AGGREGATE_TARGET automatic tuning parameters were set and the tests were run several times. This allowed Oracle to dynamically adjust the balance of SGA parameters to provide a first-cut at tuning. After the initial tests, the tuning was performed manually. Before and after each test, Automatic Workload Repository snapshots were created as time markers. After each test, an Automatic Workload Repository report was generated to monitor the performance efficiency of each Oracle instance. In particular, the Top 5 Timed Wait Events, the Buffer Pool Advisory, the PGA Advisory, the Shared Pool Advisory, and the SGA Advisory sections were carefully monitored. The SPFILE parameters were adjusted after each run to iteratively improve performance. After several iterations, it became obvious that the optimal parameters for both Windows and Linux were very close (this proved true for both the Stress Tests and the User Load tests). This is not surprising, since the same test is being run against databases that are physically configured exactly the same. To simplify the Stress Test procedures, a common set of optimal parameters was determined and applied to both the Linux and Windows instances. The same approach was used for the User Load tests. Table 2 lists the key SPFILE parameters for the two test types.

Oracle Parameter	Stress Test	User Load Test
db_block_size	8192	8192
db_cache_size	2466250752	734003200
db_file_multiblock_read_count	16	16
dispatchers		'….(disp=8)'
java_pool_size	16777216	16777216
large_pool_size	16777216	1073741824
max_dispatchers		10
max_shared_servers		1000
open_cursors	5000	5000
pga_aggregate_target	352321536	1394606080
processes	5000	5000
sessions	5505	5505
shared_pool_size	1023410176	2803682095
shared_servers		
streams_pool_size	0	0

Table 1: Key SPFILE Tuning Parameters for the Test Instances

Results and Analysis

In this section, the results of the SwingBench tests will be examined. Where appropriate, additional OS and Oracle performance monitoring logs will be examined to aid in understanding the root causes of the test results.

Oracle RAC Stress Test Results

The Oracle RAC Stress Test was performed for cases with 2 users – 250 users (1 – 125 users per node). Figure 6 is a graph of the results in terms of Transactions per Minute vs. the number of user sessions. Microsoft Windows Server 2003 measurements are shown in red, and Linux performance is shown in black. The Transactions per Minute measurement is an average for the entire session. (The average TPM per node is half the amount for the cluster.) The greatest gains occur between 2 users and 40 users. Past 40 users the incremental gains are less. As many as 72,000 Transactions per Minute are obtained for 100 users for both Linux and Microsoft Windows Server 2003. For 150 users and below, Linux and Microsoft Windows Server are fairly evenly matched. Microsoft Windows Server 2003 actually has a slight edge for 40 – 80 users, but the difference is minimal. For the 150 and 250 user cases, both Linux and MS Windows performance decline, but Linux performance declines more severely.

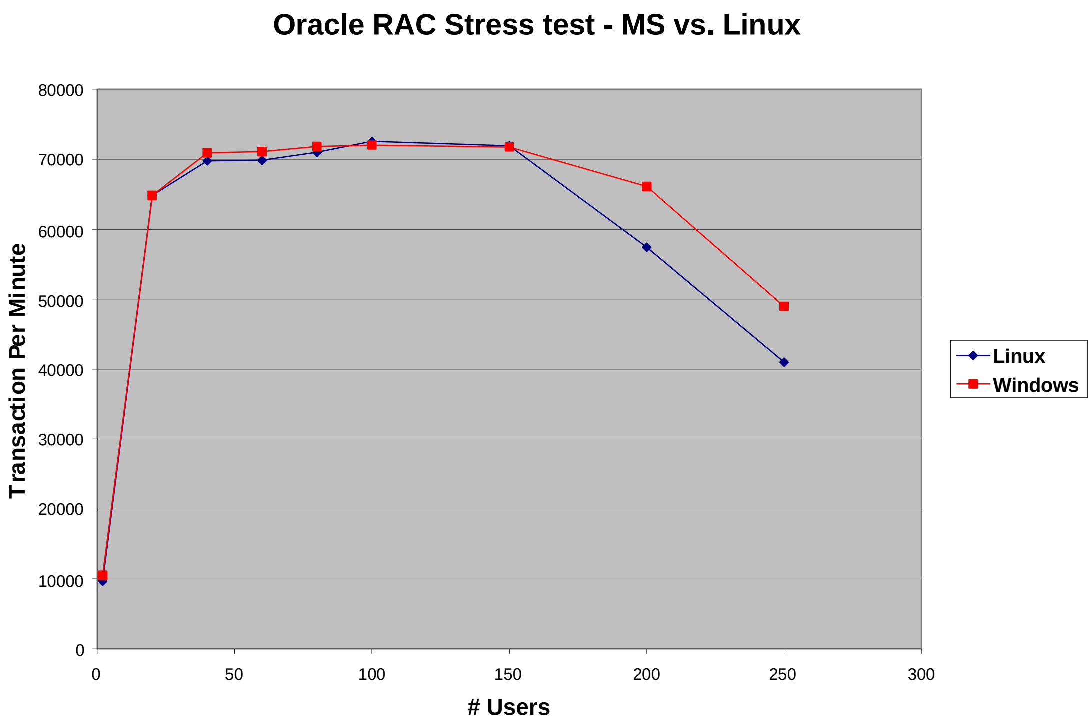

Figure 6: Oracle RAC Stress Test Transactions per Minute

One thing that is clear is that CPU utilization is consistently high for both Linux and Windows. Figure 7 shows a graph from the ClusterOverview tool. This graph shows performance for a Linux test. Note the CPU Monitor trace, showing utilization consistently above 95%. The CPU Monitor tool is not available for Windows, but similar results are obtained by viewing Windows Performance Log traces, as shown in Figure 8.

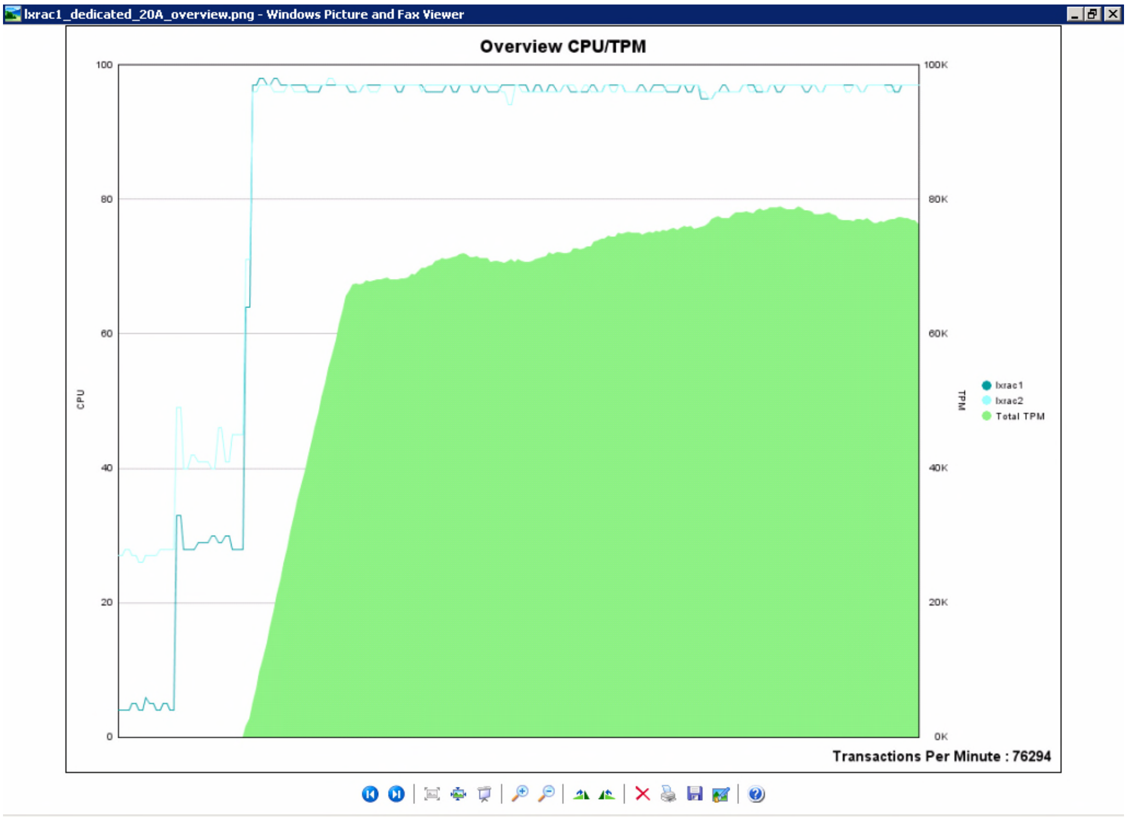

Figure 7: CPU Utilization Trace for Linux on an Overview Display

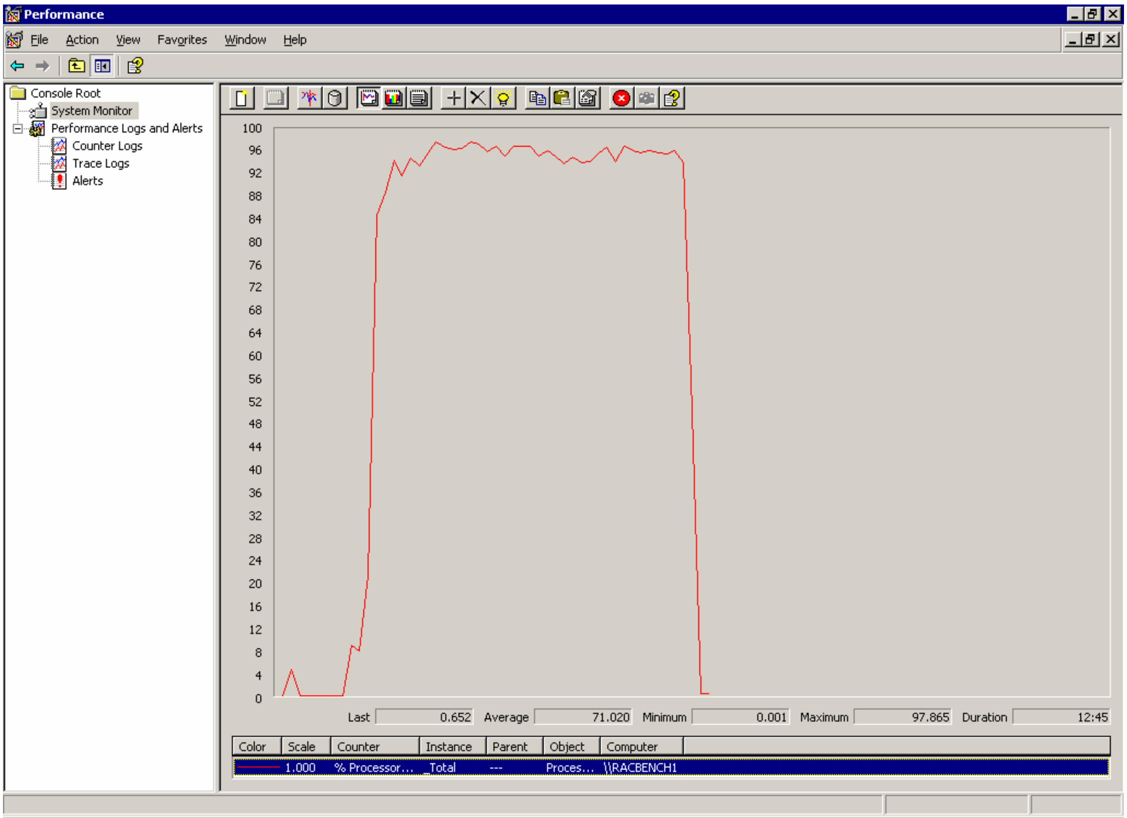

Figure 8: CPU Utilization Performance Log for Microsoft Windows Server 2003

In addition, the AWR reports show that CPU is in the top two Wait Events for the 2 – 60 user cases. This is expected and appropriate for CPU-intensive transaction activity. Above that, I/O and Oracle RAC contention wait events tend to dominate. This is inevitable as the system becomes over-saturated with user activity and user sessions begin to compete for resources. The only ways to alleviate these wait events would be to:

- Add more disk spindles for faster I/O
- Add additional private interconnect interfaces in a network bonding configuration. To increase interconnect throughput.

A breakdown of response times for individual Order Entry schema components is shown in Figure 9. For all user cases, a similar pattern is seen. Linux has noticeably faster response time for the "New Customer Registration" component. However, MS Windows has slightly faster response in all other categories. These response times roughly balance out, since overall Transactions per Minute are approximately the same for Linux and Windows for 150 users and below. Figure 10 illustrates the response times for the 250 user case. Note that the discrepancy between Windows and Linux performance seems to be associated with increased response times for the "Order Products" component of the tests.

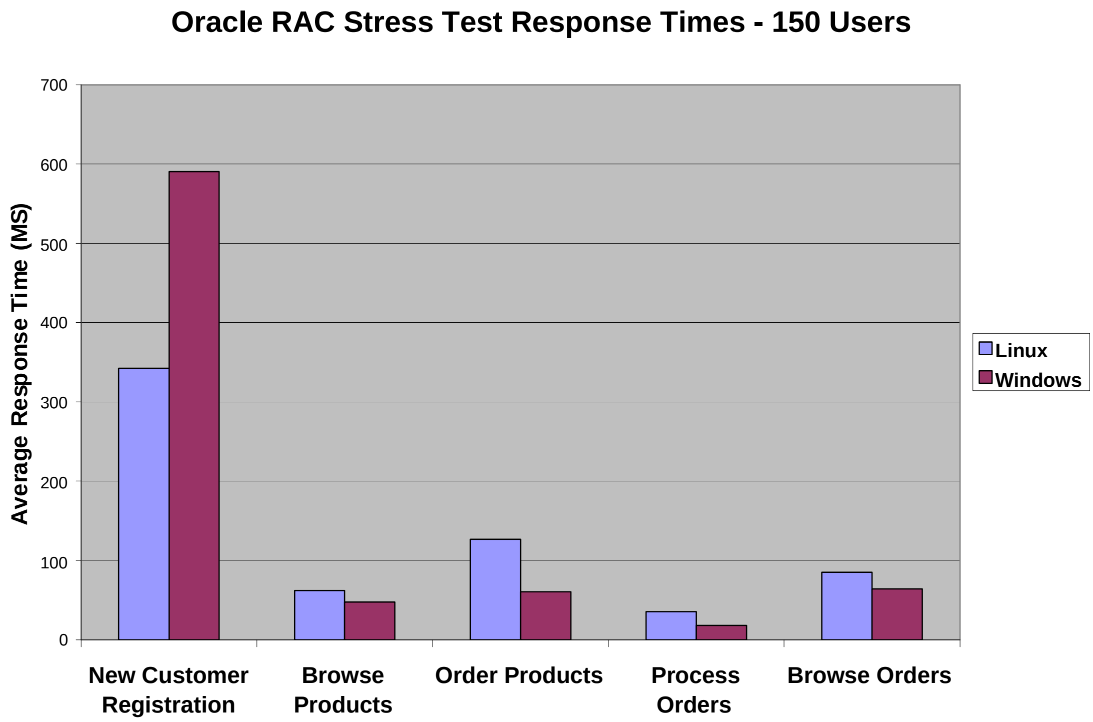

Figure 9: Oracle RAC Stress Test Response Times for 150 Users

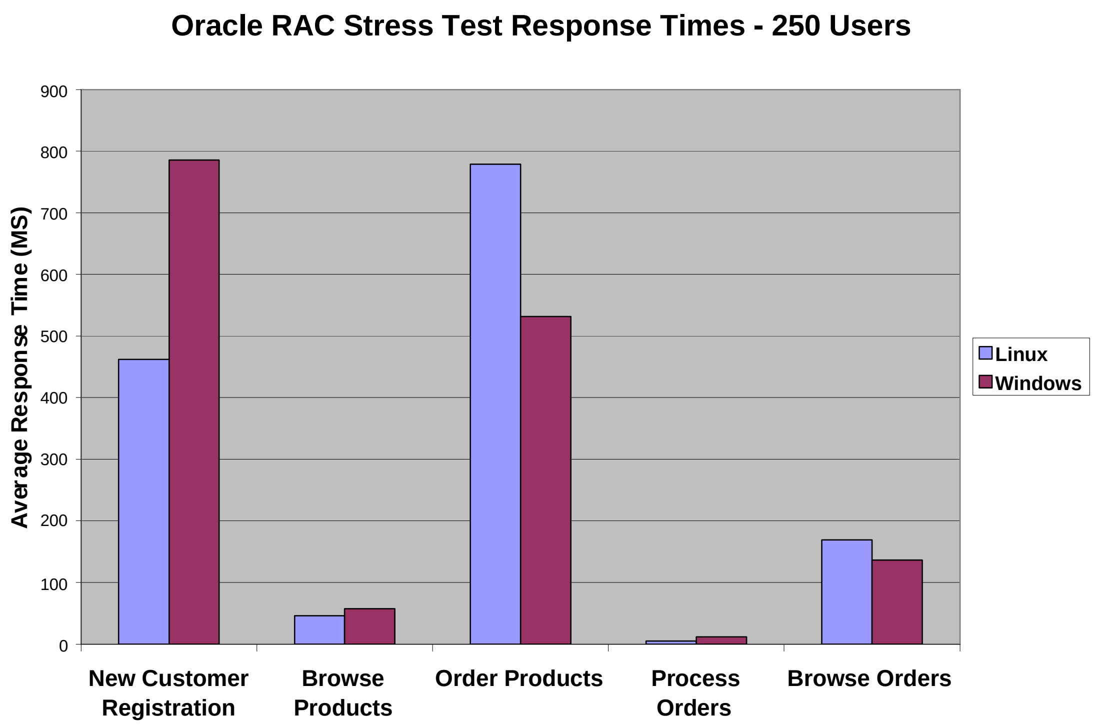

Figure 10: Oracle RAC Stress Test Response Times for 250 Users

Examining I/O trends provides some clues to what may be causing the decline in performance for the 250 user Linux test. In Figure 11, I/O data from Automatic Workload Repository Reports[2] is compiled to show total I/O (Reads + Writes) on a per tablespace basis. I/O is roughly equivalent for Linux and Windows (Linux has slightly better I/.O in this particular case). The SOEINDEX and SOE tablespaces dominate in terms of I/O, which is not too surprising. Figure 12 shows I/O for the 250 user case. The big difference here is a sharp increase in TEMP tablespace I/O, with much higher I/O for the Linux TEMP tablespace than the Windows TEMP tablespace.

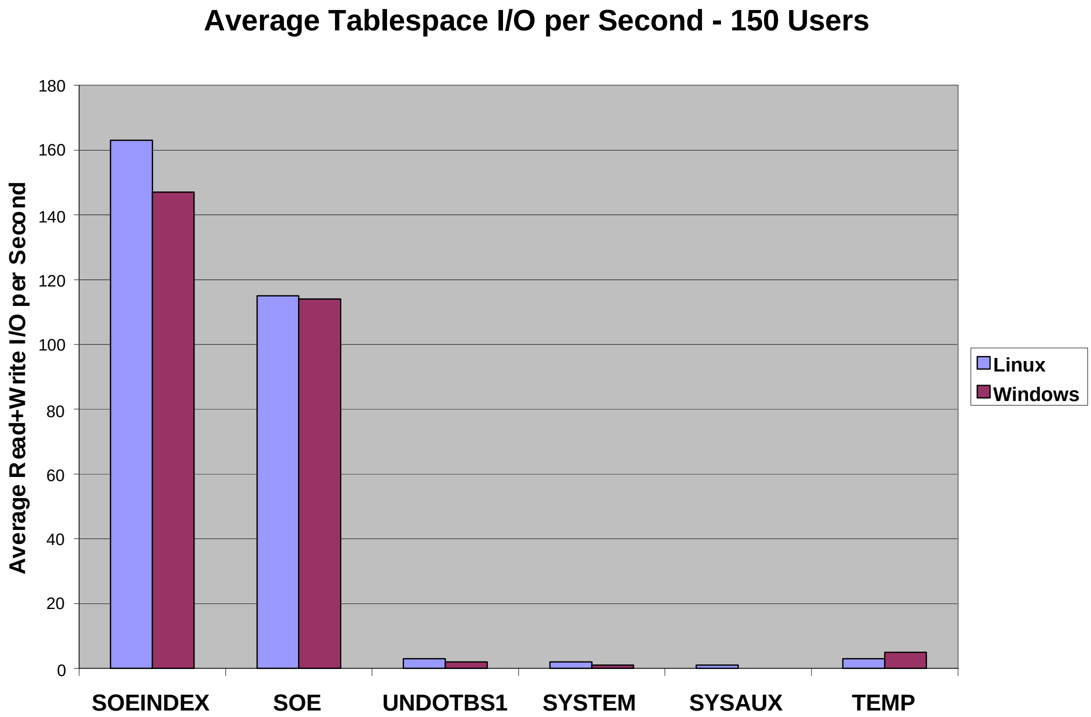

Figure 11: Tablespace I/O for 150 Users

[2] All AWR statistics are from the first node of each Oracle Real Application Cluster.

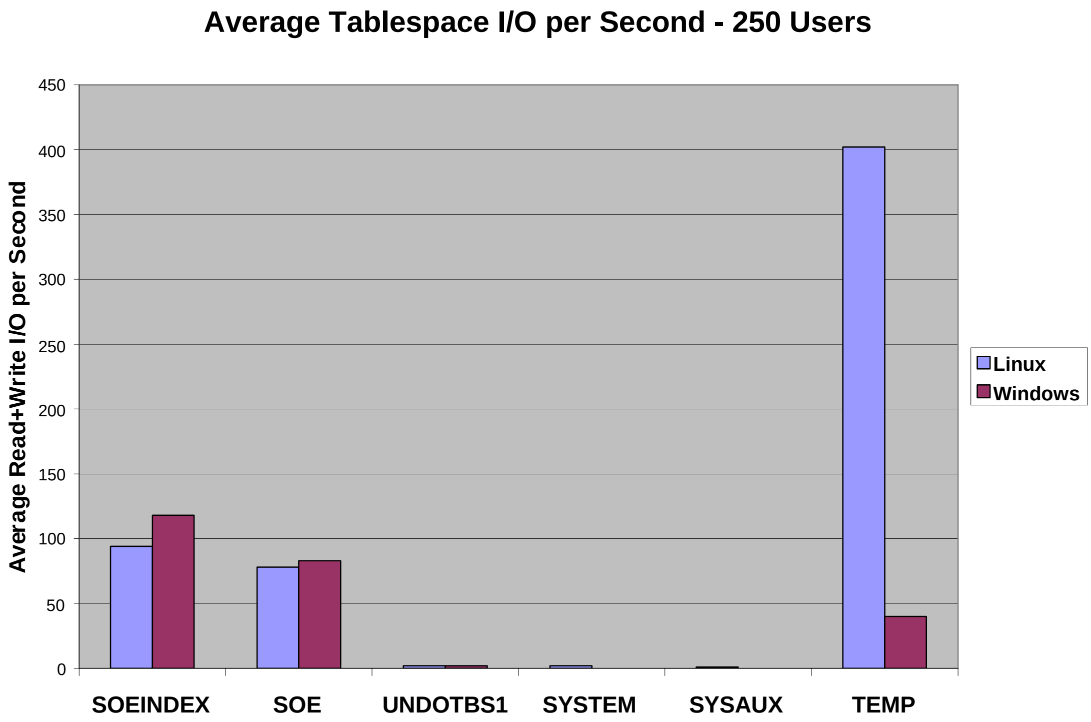

Figure 12: Tablespace I/O for 250 Users

Further examination of the AWR reports shows a possible cause for the I/O trends. As stated above, the target for optimizing tuning parameters was the 60 user case. At 250 users, the PGA_AGGREGATE_TARGET parameter is no longer adequate. Although PGA memory should be dynamically adjusted upwards through time, PGA memory is not adequate over the short term. This causes sorts to be pushed to disk (in the TEMP tablespace) rather than being performed in memory. Figure 13 illustrates this effect.

Of course, this problem can be fixed for both Linux and Windows by simply increasing the PGA_AGGREGATE_TARGET parameter. Of greater interest is attempting to understand why exhaustion of PGA memory is a much greater problem for Linux than Windows in this case. One obvious reason might be that Linux is simply using more memory than Windows. Figure 14 shows that this assumption might actually be true. Figure 14 compares peak OS memory used by Linux and Windows during these tests. The source of data was Windows Performance Logs and the Linux “top” command. In each case, Linux is using more memory than Windows to perform the same task, which may explain the PGA exhaustion.

To put this in perspective, problems with memory will likely only occur under high stress conditions. As shown during the RAC stress test, even under a high CPU load, Linux and Microsoft Windows Server 2003 will perform virtually identically.

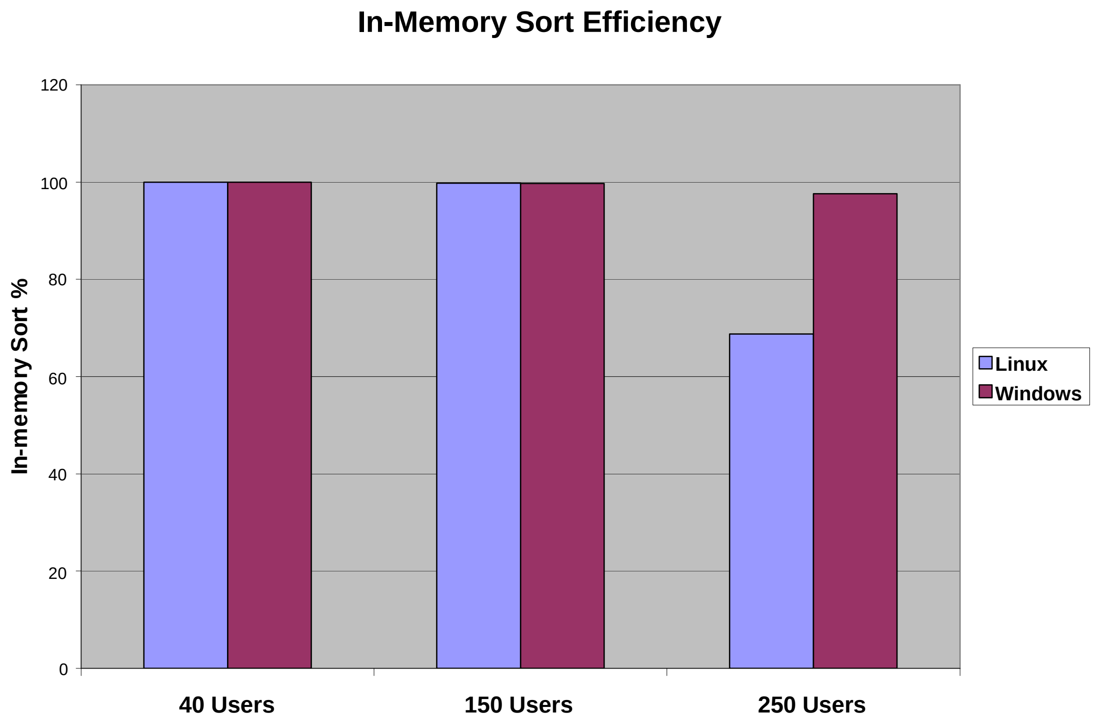

Figure 13: Sorts in Memory vs. Sorts on Disk

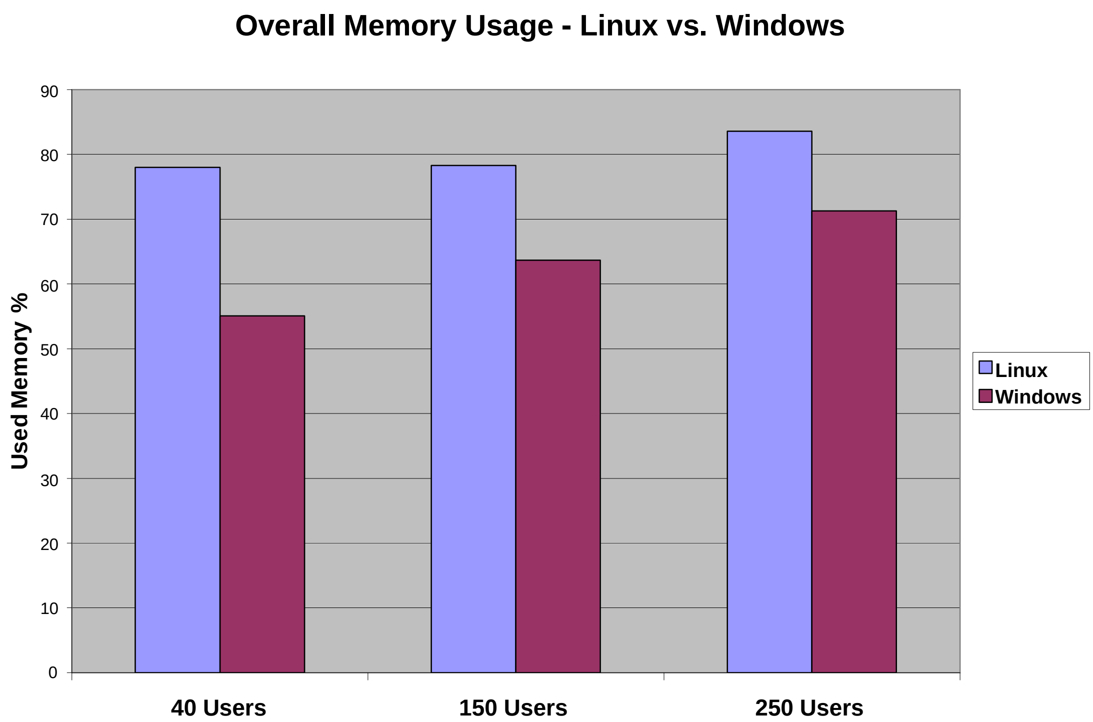

Figure 14: Used Memory by OS for the Oracle RAC Stress Test

Oracle RAC User Load Test Results

The goal of the User Load Test was to simulate the effect of supporting a large number of concurrent user sessions. Pursuing large numbers of user connections can be problematic in test like these. Oracle Net tends to reject user connection attempts if too many attempts are made in too short of a period of time. In addition, too many Java threads for the client testing software can cause sessions to crash. Another completely legitimate reason for connections to be rejected is exhaustion of system memory. Additional user connections take extra memory, whether dedicated or shared connections are used (shared server connections were used here). The number of user connections that you can support is obviously much higher if you have 16 GB of RAM vs. 8 GB of RAM.

Due to these factors, the upper limit of user connections observed in these tests is probably not the true limit. Nevertheless, over 4,000 simultaneous connections were sustained (2,000 per server). This is a fair test of user generated load, especially with moderate to high transactional activity per server.

Figure 15 shows User Load test results in terms of Transactions per Minute. The response for 500 – 2500 users is virtually identical. However, the TPM performance for 2500 – 3500 users falls off noticeably for Linux vs. Microsoft Windows Server. Windows scales almost linearly up through 4000 users. The 4500 user test failed to complete for Microsoft Windows, but it appeared to be the result of a Java problem.

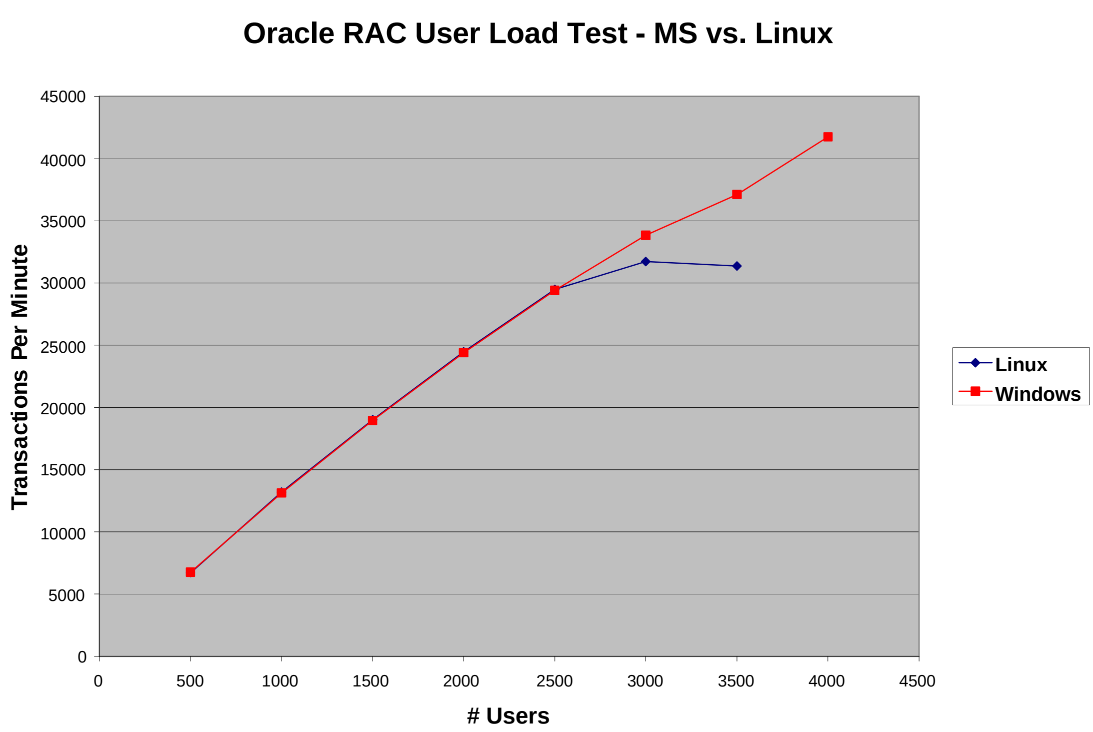

Figure 15: Oracle RAC User Load Test Transactions per Minute

Response Times for Order Entry schema components show a slight advantage for Windows for the 2500 User test case, as shown in Figure 16. For the 3500 User test case (Figure 17), Windows response time is definitely better and there appears to be some type of problem during the Linux test, based on the large gap in response times.

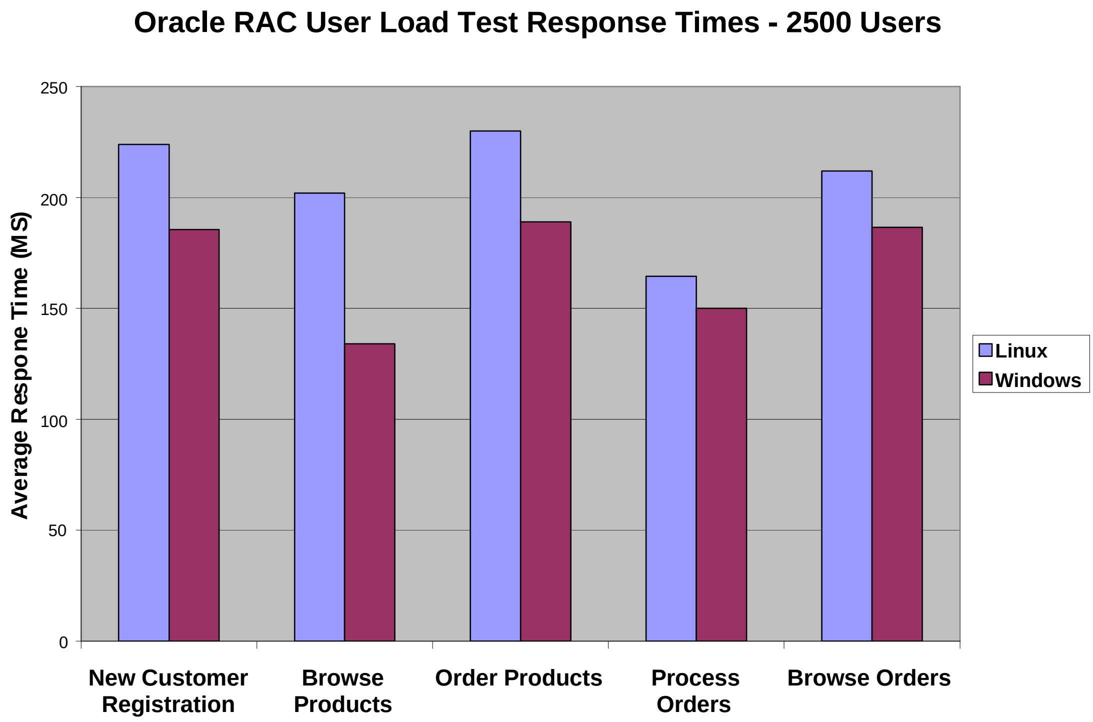

Figure 16: Oracle RAC User Load Test Response Times for 2500 Users

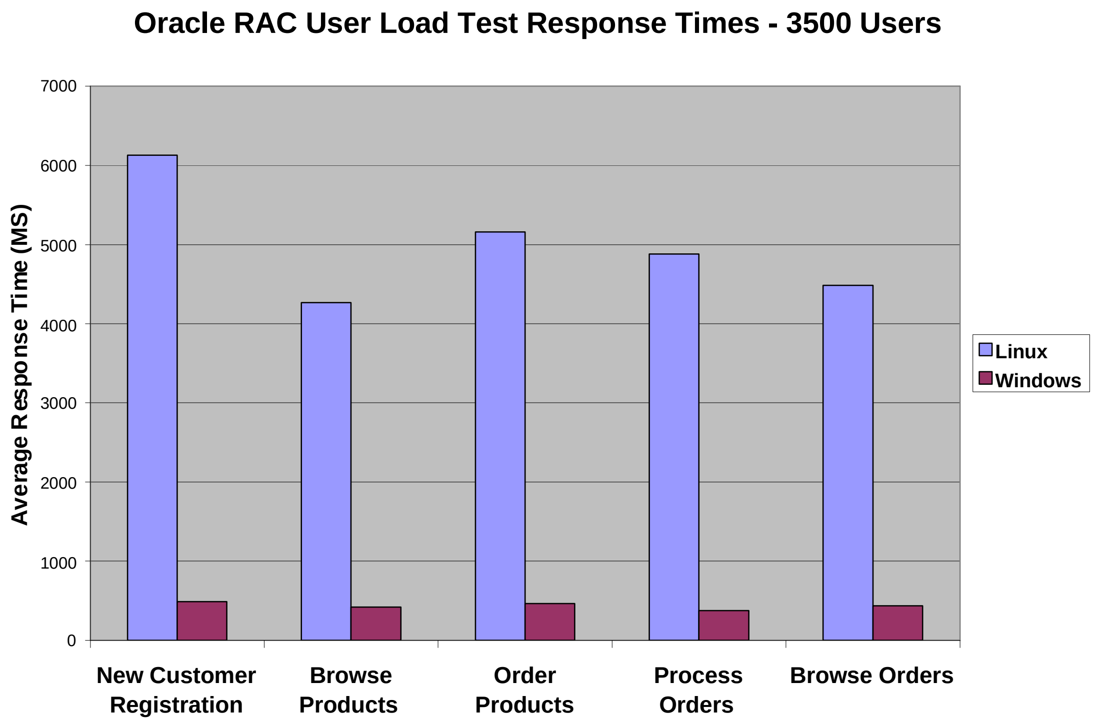

Figure 17: Oracle RAC User Load Test Response Times for 3500 Users

In the Oracle RAC Stress Test, CPU usage was near the maximum for both Windows and Linux tests. However, that is not the case here. Figure 18 shows CPU usage for the Linux test for 2500 users, averaging around 90% busy. However, Figure 19 shows that Windows Oracle RAC showed CPU approximately 40% busy for the same test. For the 3500 User test, Figure 20 shows the Linux CPUs at over 95% busy. Note that performance rapidly declines when the CPU is 100% saturated. In Figure 21, Windows CPU performance for the 3500 User case shows only 70% usage.

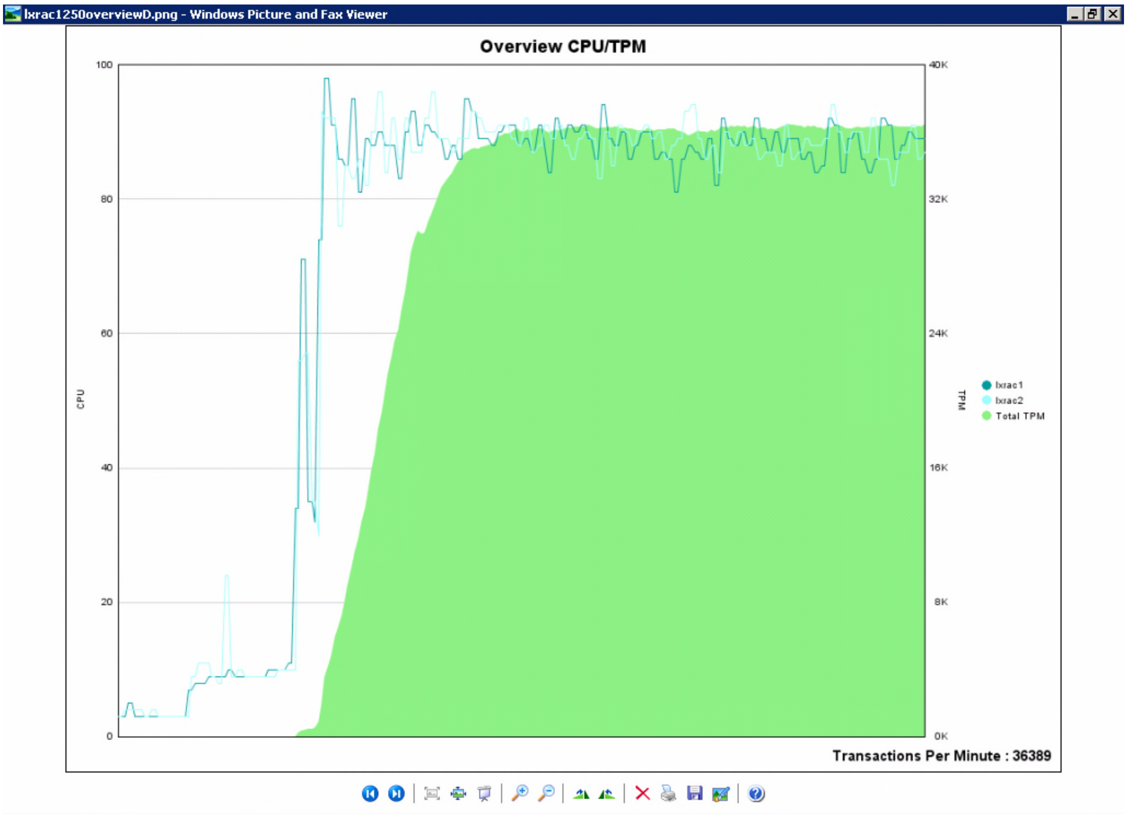

Figure 18: ClusterOverview Chart Showing Linux CPU Usage for 2500 Users

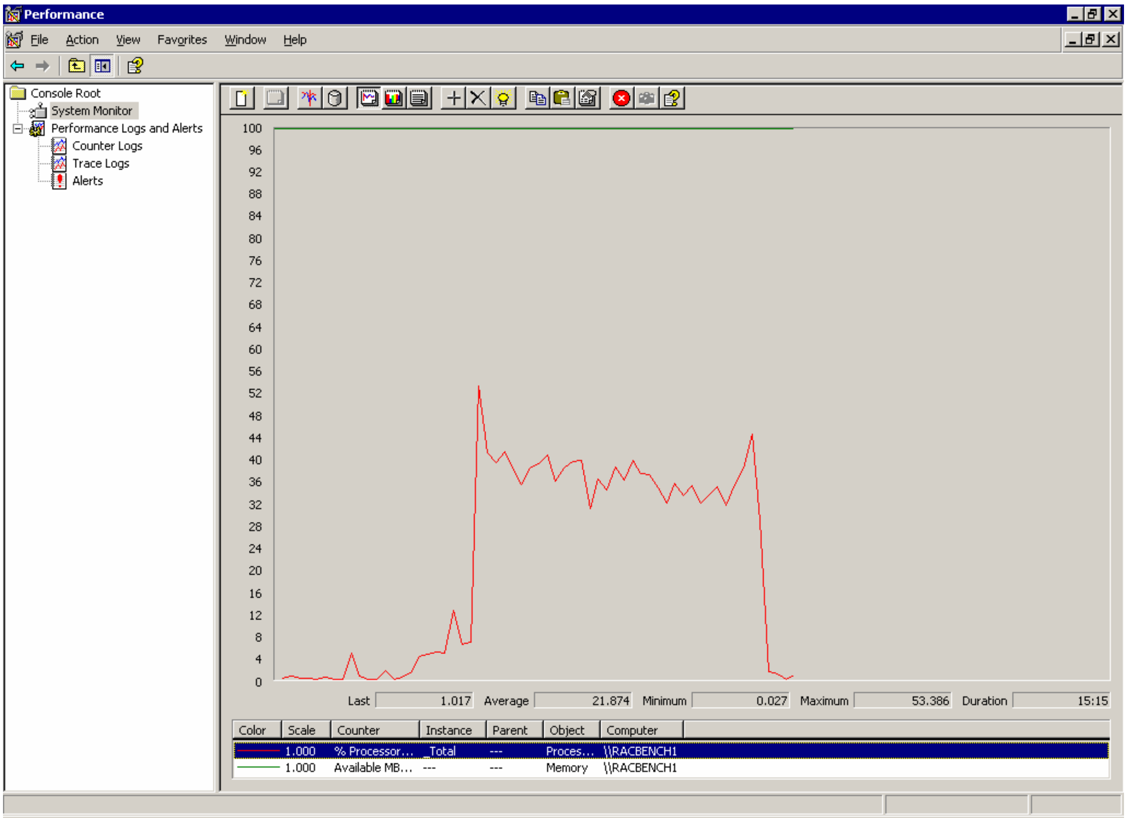

Figure 19: Windows Performance Log Showing Windows CPU Usage for 2500 Users

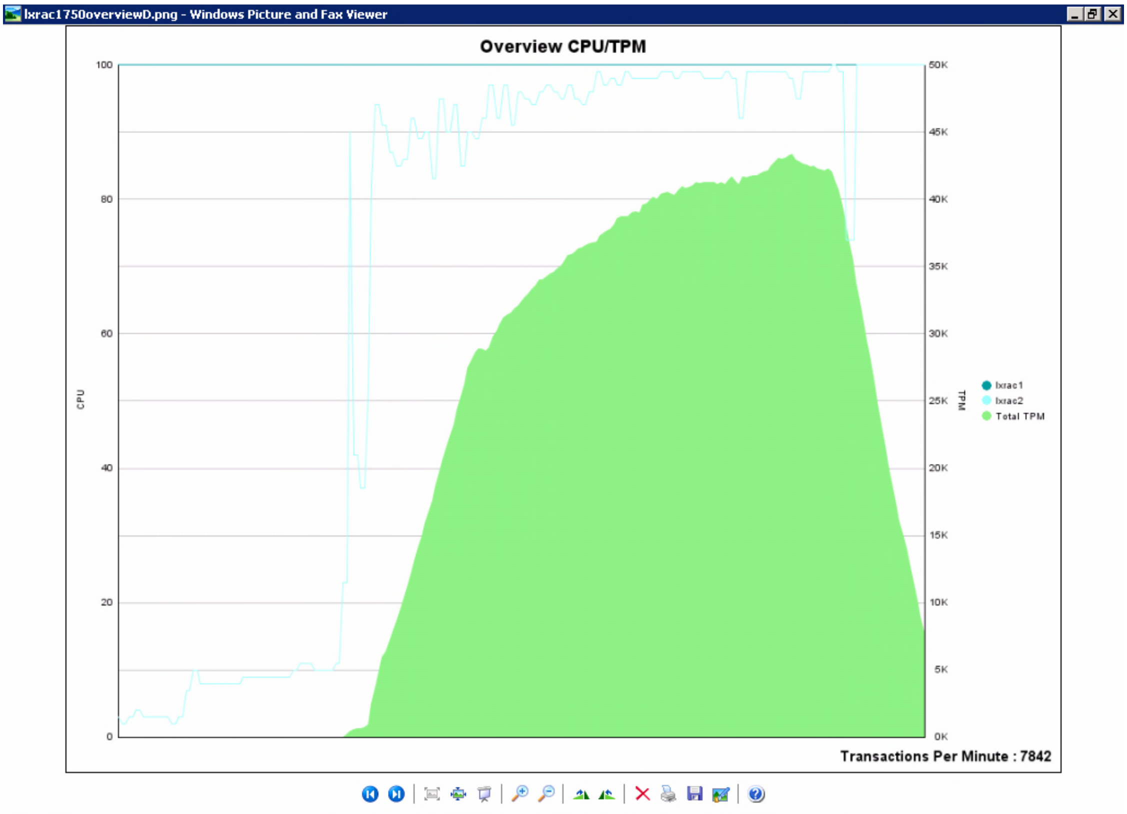

Figure 20: ClusterOverview Chart Showing Linux CPU Usage for 3500 Users

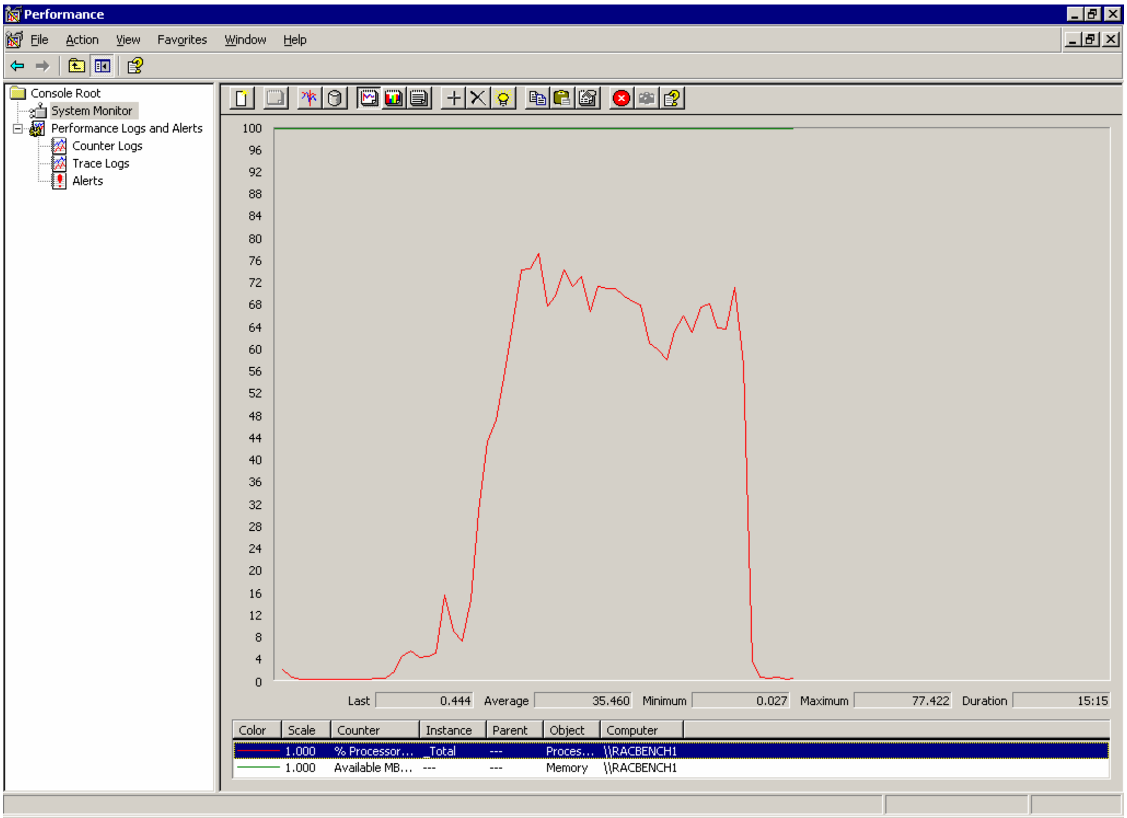

Figure 21: Windows Performance Log Showing Windows CPU Usage for 3500 Users

I/O performance shows a slightly different story. Figure 22 illustrates that tablespace I/O is actually slightly better for Linux for the 2500 user case. However, Figure 23 shows Windows tablespace I/O slightly better for 3500 users.

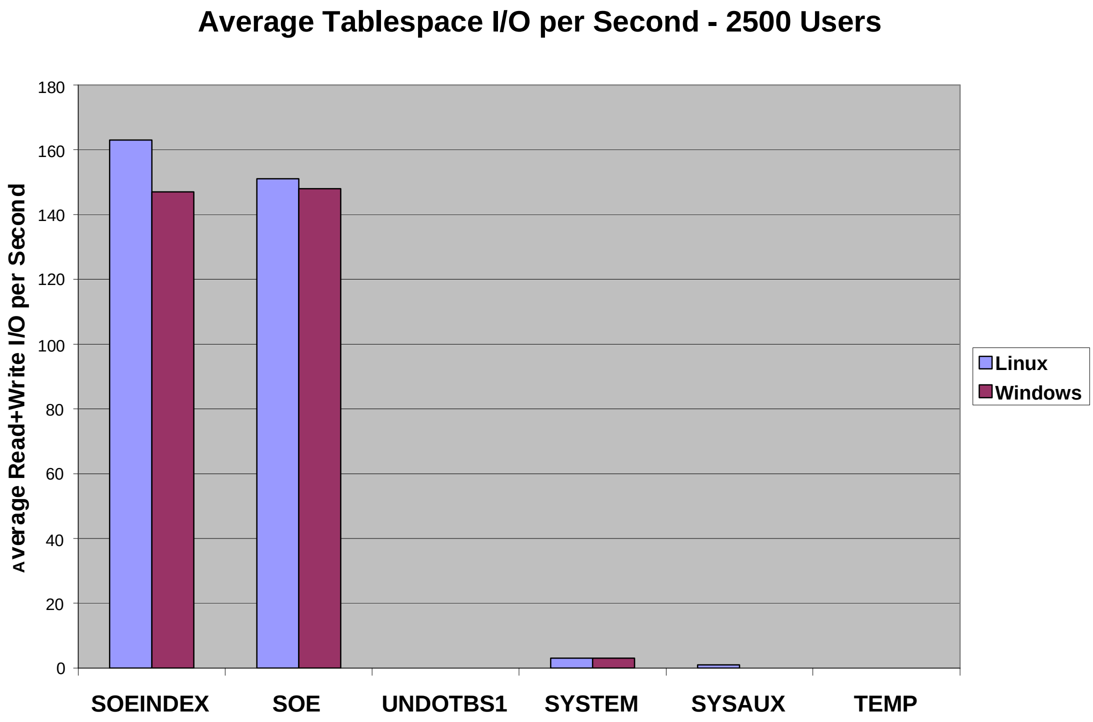

Figure 22: Tablespace I/O for 2500 Users

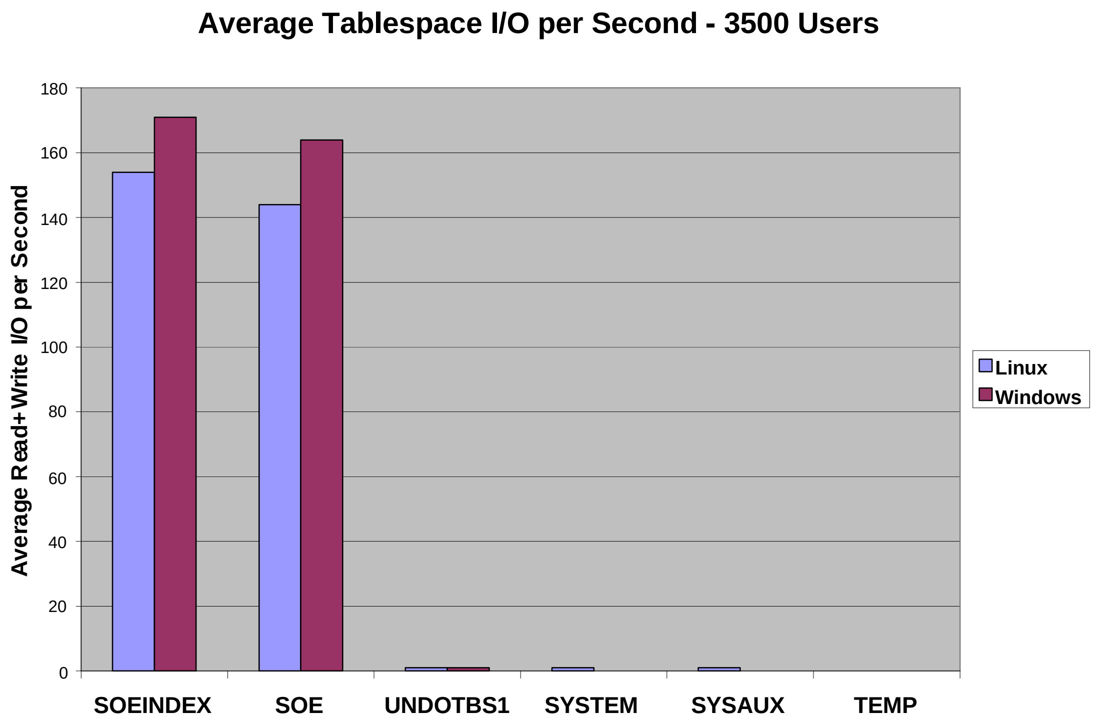

Figure 23: Tablespace I/O for 3500 Users

Based on the memory comparison seen in the Oracle RAC Stress Test, a check was made for memory usage for the Oracle RAC User Load test. Again, when compared to Linux, MS Windows shows less memory used for the same tasks, as displayed in Figure 24. It is likely that memory usage plays a role in the decline seen in performance for Linux when over 2500 users are connected. However, in this case, it is not necessarily just PGA memory that is exhausted. Figure 25 shows that for the 3500 user test, Linux begins to show significant paging to the Swap file.[3] This has an immediate negative impact on performance. This is likely due to the fact that memory paged to disk is radically slower than main memory. MS Windows seems to show a different pattern than Linux for how it uses its paging file. It appears that Windows begins using small percentages of the paging file long before RAM is completely used. Based on other these and other tests, the performance drop-off when paging occurs in Microsoft Windows Server 2003 does not appear to be as abrupt as when the Swap file is used by Linux.

One theory that might explain the CPU and memory usage anomalies relates back to how processing is performed on MS Windows as compared to Linux. Windows uses a multi-threaded model, while Linux uses a multi-process model. Theoretically, the multi-threaded model should require less overhead in the form of context switches. This could explain why Microsoft Windows Server 2003 used less CPU and memory than Linux for the same tasks.

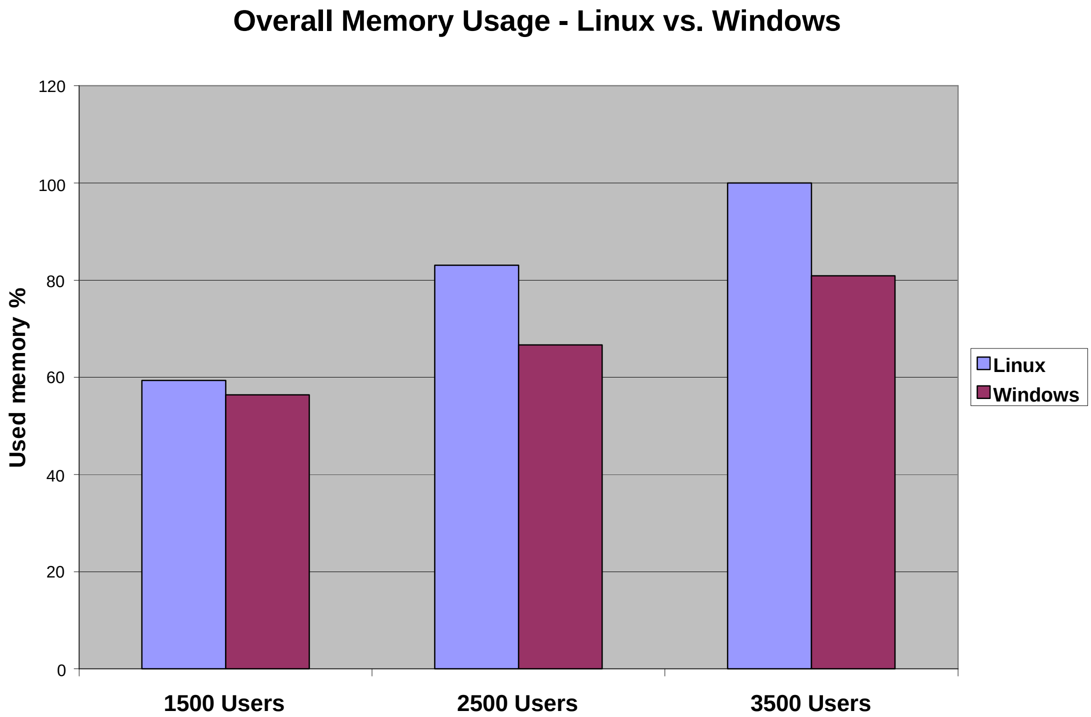

Figure 24: Used Memory by OS for the RAC User Load Test

[3] Note that the baseline % used of the page or swap file before the start of the test was subtracted from the % used during the test to derive the "Used Memory %"

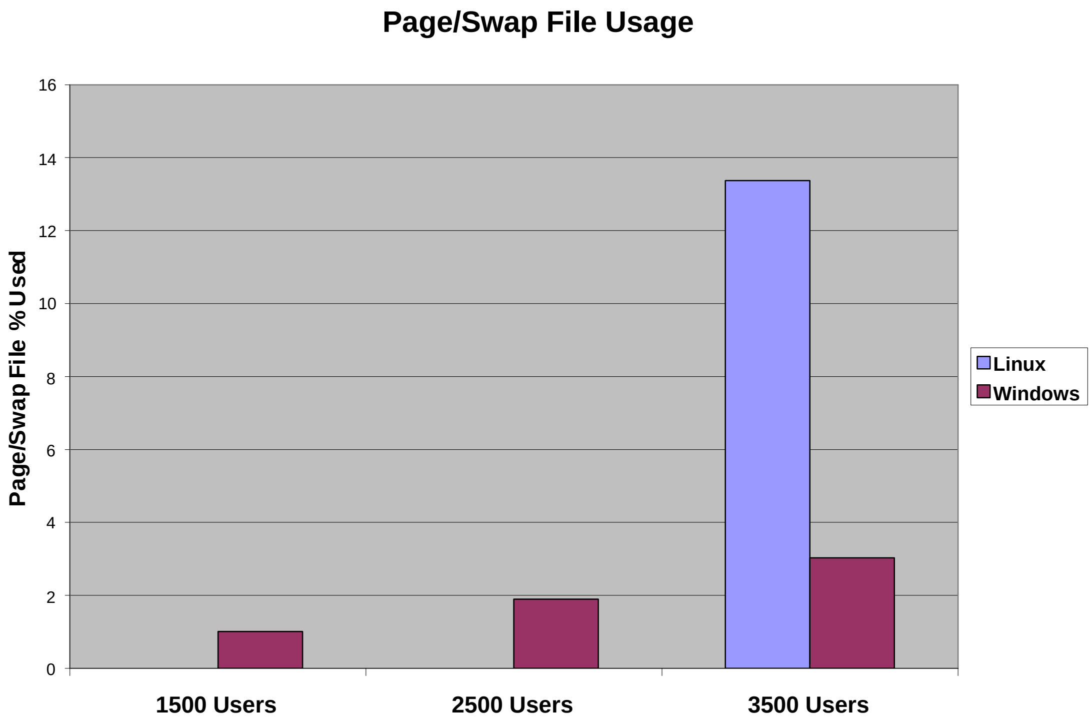

Figure 25: Page/Swap File Usage by OS for the RAC User Load Test

Conclusions

The following behavior was observed during testing of the Oracle RAC databases on Red Hat Enterprise Linux x86_64 and Microsoft Windows Server 2003 Enter:

- Oracle RAC Stress Test
  - Transactions Per Minute were roughly equivalent for 2 - 150 user sessions.
  - Transactions Per Minute were up to 16% higher for MS Windows for 150 - 250 users.
  - CPU usage was above 90% for all of these tests for both Linux and MS Windows Server.
  - The response times for the "New Company Registration" test component were up to 30% faster for Linux than MS Windows.
  - The response times for the "Order Products" test component were up to 50% faster for MS Windows than Linux, and the response times for the "Order Products" component increased substantially above 150 users.
  - Tablespace I/O throughput for the data and index tablespaces was approximately 10% higher for Linux at 150 users, but MS Windows I/O throughput for the data and index tablespaces was approximately 25% higher at 250 users.
  - Tablespace I/O for the TEMP tablespace was very high for Linux at 250 users.
  - At 250 users, Linux performed over 30% of sorts on disk instead of sorts in memory.

    - For all of these tests, Linux used 17% - 40% more memory than MS Windows Server.
- Oracle RAC User Load Tests
    - Transaction Per Minute results were virtually identical for 50 – 2500 users.
    - At 3500 users, the Transactions Per Minute for MS Windows were 18% higher than the Linux Transactions Per Minute.
    - At 4000 users, MS Windows Server TPM continued to increase while the Linux test failed to complete.
    - At 2500 users, component response times were 9% - 33% faster for MS Windows Server; while at 3500 users Windows was faster by a wide margin.
    - At 2500 users, Tablespace I/O throughput is up to 10% higher for Linux, while at 3500 users, Tablespace I/O throughput is up to 10% higher for MS Windows Server.
    - For all of these tests, Linux used 5% - 25% more memory than MS Windows Server.
    - MS Windows Server showed a gradual increase in page file usage from 1% - 3% as the number of users increased.
    - Linux showed an abrupt increase from 0% swap file usage to 13% swap file usage at 3500 users.

To summarize, for both the Oracle RAC Stress test and the User Load test, Microsoft Windows Server 2003 RAC and Red Hat Linux RAC both performed well, showing good Transactions Per Minute performance and Response Time performance as the number of user sessions and overall load was increased. Under low-moderate stress conditions (CPU, memory, and I/Os at low to moderate levels), Linux Oracle RAC and MS Windows Oracle RAC seemed to achieve almost exactly the same Transactions per Minute results. However, there appeared to be subtle differences in CPU usage, I/O performance, and memory usage, even for low-moderate stress conditions. Nevertheless, the differences seemed to roughly balance out, leading to very similar results.

Under high stress conditions (high CPU usage, high memory usage, high number of I/Os), both Windows and Linux performed well, but Windows appeared to use both CPU and memory more efficiently. For the Oracle RAC Stress tests, Linux exhausted the available PGA memory, causing sorts to be forced to disks with lowered performance. At the same number of users, Windows worked well with available PGA memory. For the User Load tests, Linux started significant paging to disk for lower numbers of users than Windows. This resulted in significantly lower performance. In addition, the CPUs were significantly less busy for the Windows Oracle RAC servers throughout these tests. I/O performance was roughly equal, with some tests showing slightly better I/O for Linux, and other tests showing slightly better I/O performance for Windows. Cases where Linux I/O was significantly slower than Windows were probably driven by the observed memory problems. A theory that may explain the observed differences in CPU and Memory usage is that the multi-threaded processing model utilized by Microsoft Windows requires less context switching than the Linux multi-process model, and is therefore more efficient.

Despite the differences in performance that were observed, the results are close enough that an Administrator could be happy with either Windows or Linux performance in a 64-

bit Production environment. It should be noted that these tests are not necessarily the last word in comparing Linux and Windows performance. With enough time and analysis, it may be possible to tune either Linux or Windows performance slightly better than was achieved with these tests. However, it is clear that the 64-bit Enterprise Edition of Microsoft Windows Server 2003 is just as valid a platform for implementing Oracle RAC as is the 64-bit version of Red Hat Enterprise Linux. Administrators who are comfortable with a Windows Server environment should not hesitate to implement Oracle RAC on the 64-bit Enterprise Edition of Microsoft Windows 2003 Server x64.

References

"ClusterOverview Reference and User Guide"; Author: Dominic Giles; Oracle Corporation UK Ltd., January 2004;
http://www.dominicgiles.com/swingbench/clusteroverview21.pdf

"Comparison of 32-bit and 64-bit Oracle Database Performance on the Dell PowerEdge 6850 Server with Microsoft Windows Server 2003", Larry Pedigo, Performance Tuning Corporation, 2005;
http://www.perftuning.com/pdf/Compare3464.pdf

"Deployment Guide: Oracle on Microsoft Windows and Dell PowerEdge Servers", Bryan Thomas & Larry Pedigo, Performance Tuning Corporation, 2004;
http://www.dell.com/downloads/global/solutions/Oracle%20on%20Windows%20Deployment.pdf

"Deployment Guide: Oracle on Microsoft Windows and the Dell PowerEdge 6850 Server", Larry Pedigo & Bryan Thomas, Performance Tuning Corporation, 2005;
http://www.perftuning.com/pdf/OracleMSWinDell..pdf

"EMC Support Matrix", October 2006 ; P/N 300−000−166, ECO 51090 Rev B64;
http://www.emc.com/interoperability/matrices/EMCSupportMatrix.pdf

"HP Parallel Database Cluster installation guide for Red Hat Linux on x86 based servers", Hewlett Packard Development Company, L.P., May, 2006;
http://h20293.www2.hp.com/portal/swdepot/displayProductInfo.do?productNumber=PDC-Kit

"Market Share: Relational Database Management Systems by Operating System, Worldwide, 2005"; Author: Colleen Graham, Gartner DataQuest Corporation, May 2006;
http://www.gartner.com/it/products/research/dataquest.jsp

"Oracle Clusterware and RAC Administration and Deployment Guide, 10g Release 2 (10.2)", Part No. B14197-03; Primary Authors: David Austin, Mark Bauer, Douglas Williams; Contributing Authors: Troy Anthony, Anand Beldalker, Carol Colrain, Jonathan Creighton, Rajesh Dasari, Yong Hu, Rajiv Jayaraman, Sameer Joshi, Raj Kumar, Ken Lee, Barb Lundhild, Venkat Maddali, Gaurav Manglik, John McHugh, Bharat Paliwal, Dipak Saggi, Sudheendra Sampath, Daniel Semler, Cathy Shea, Khethavath P. Singh, Bipul Sinha, Mike Zampiceni; Oracle Corporation, January 2006

"Oracle Database 10g Linux Administration"; Author: Edward Whalen, Performance Tuning Corporation; McGraw-Hill Osborne Media; 1st edition (October 6, 2005); ISBN 0072230533

"Oracle Database Installation Guide, 10g Release 2 (10.2) for Linux x86-64", Part No. B15667-01; Primary Authors: Apolina Das, Sanjay Sharma, Lyju Vadassery; Contributing Authors: Kevin Flood, Pat Huey, Clara Jaeckel, Emily Murphy, Terri Winters; Contributors: David Austin, Subhranshu Banerjee, Mark Bauer, Robert Chang, Jonathan Creighton, Sudip Datta, Padmanabhan Ganapathy, Thirumaleshwara Hasandka, Joel Kallman, George Kotsovolos, Simon Law, Richard Long, Rolly Lv, Padmanabhan Manavazhi, Sreejith Minnanghat, Krishna Mohan, Rajendra Pingte, Hanlin Qian, Janelle Simmons, Preeti Shukla, Roy Swonger, Douglas Williams; Oracle Corporation, September 2005

“Oracle Database Installation Guide, 10g Release 2 (10.2) for Microsoft Windows (x64)”, Part No. B15681-03; Primary Author: Patricia Huey; Contributors: Punsri Abeywickrema, Eric Belden, Phil Choi, Toby Close, Sudip Datta, Jim Emmond, David Friedman, Alex Keh, Mark Kennedy, Peter LaQuerre, Rich Long, Anu Natarajan, Mark MacDonald, Matt McKerley, Mohamed Nosseir, Bharat Paliwal, Sham Rao Pavan, Hanlin Qian, Christian Shay, Helen Slattery, Debbie Steiner, Linus Tanaka, Ravi Thammaiah, Sujatha Tolstoy, Alice Watson, and Vinisha Dharamshi; Oracle Corporation, June 2006

“Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide, 10g Release 2 (10.2) for Linux”, Part No. B14203-08; Primary Authors: David Austin, Mark Bauer, Kevin Flood, Emily Murphy, Lyju Vadassery, Douglas Williams; Contributing Authors: Jonathan Creighton, Pat Huey, Raj Kumar; Contributors: Chris Allison, Karin Brandauer, Sudip Datta, Rajiv Jayaraman, Roland Knapp, Diana Lorentz, Barb Lundhild, Vijay Lunawat, John Patrick McHugh, Randy Neville, Philip Newlan, Michael Polaski, Dipak Saggi, Sudheendra Sampath, Janelle Simmons, Clive Simpkins, Khethavath P. Singh, Nitin Vengurlekar, Gary Young; Oracle Corporation, September 2006

“Oracle Database Oracle Clusterware and Oracle Real Application Clusters Installation Guide, 10g Release 2 (10.2) for Microsoft Windows”, Part No. B14207-05; Primary Authors: David Austin, Mark Bauer, Kevin Flood, Emily Murphy, Lyju Vadassery, Douglas Williams; Contributing Authors: Jonathan Creighton, Pat Huey, Raj Kumar; Contributors: Chris Allison, Karin Brandauer, Robert Chang, Sudip Datta, Luann Ho, Rajiv Jayaraman, Roland Knapp, Diana Lorentz, Barb Lundhild, Vijay Lunawat, John Patrick McHugh, Randy Neville, Philip Newlan, Michael Polaski, Dipak Saggi, Sudheendra Sampath, Janelle Simmons, Clive Simpkins, Khethavath P. Singh, Nitin Vengurlekar, Gary Young; Oracle Corporation, September 2006

“Oracle Database Performance Tuning Guide, 10g Release 2 (10.2)”, Part No. B14211-01; Primary Author: Immanuel Chan; Contributors: Aditya Agrawal, James Barlow, Vladimir Barriere, Ruth Baylis, Eric Belden, Pete Belknap, Qiang Cao, Sunil Chakkappen, Sumanta Chatterjee, Alvaro Corena, Benoit Dageville, Dinesh Das, Karl Dias, Vinayagam Djegaradjane, Harvey Eneman, Bjorn Engsig, Mike Feng, Cecilia Gervasio, Bhaskar Ghosh, Ray Glasstone, Leslie Gloyd, Prabhaker Gongloor, Connie Dialeris Green, Russell Green, Joan Gregoire, Lester Gutierrez, Lex de Haan, Karl Haas, Brian Hirano, Lilian Hobbs, Andrew Holdsworth, Mamdouh Ibrahim, Hakan Jacobsson, Christopher Jones, Srinivas Kareenhalli, Feroz Khan, Stella Kister, Paul Lane, Sue K. Lee, Herve Lejeune, Yunrui Li, Juan Loaiza, Diana Lorentz, George Lumpkin, Joe McDonald, Bill McKenna, Mughees Minhas, Valarie Moore, Sujatha Muthulingam, Gary Ngai, Michael Orlowski, Kant C. Patel, Richard Powell, Mark Ramacher, Shankar Raman, Yair Sarig, Uri Shaft, Vinay Srihari, Sankar Subramanian, Margaret Susairaj, Hal Takahara, Misha Tyulenev, Mark Van de Wiel, Venkateshwaran Venkataramani, Nitin Vengurlekar, Stephen Vivian, Simon Watt, Andrew Witkowski, Graham Wood, Khaled Yagoub, Mohamed Zait; Oracle Corporation, June 2005

"Oracle Database Platform Guide, 10g Release 2 (10.2) for Microsoft Windows (x64)", Part No. B15688-03; Primary Author: Reema Khosla; Contributing Author: Janelle Simmons; Contributors: Ricky Chen, David Collelo, David Friedman, Sue K. Lee, Rich Long, Satish Panchumarthy, Ravi Thammaiah; Oracle Corporation, June 2006

"Oracle Performance Tuning"; Authors: Edward Whalen, Mitchell Schroeter, Performance Tuning Corporation; Addison-Wesley Pub Co; 1st edition (April 16, 2002); ISBN 0672321467

"Microsoft® Windows Server™ 2003 R2 Enterprise Edition Whitepaper - An Overview of Enterprise Edition"; Microsoft Corporation, 2006; http://www.microsoft.com/windowsserver2003/evaluation/overview/enterprise.mspx

"Recommended Configuration Matrix HP Parallel Database Cluster for Windows", Hewlett Packard Development Company, L.P., July 20, 2006; ftp://ftp.compaq.com/pub/products/servers/ha/oracle/pdc-windows-rac-config-matrix.pdf

"Support Matrix For HP Parallel Database Cluster for Oracle 10g RAC and / or Application Server (J2EE) on Red Hat Linux X86_64 and x86 32 bit with: MSA 1500, MSA1000, EVA3000, EVA5000, EVA4000/6000/8000"; Author: Kevin Lyons, Hewlett Packard Development Company , L.P.; June 12, 2006; ftp://ftp.compaq.com/pub/products/servers/ha/oracle/pdc-lx86-3264.pdf

"SwingBench 2.2 Reference and User Guide"; Author: Dominic Giles; Oracle Corporation UK Ltd., August 2005; http://www.dominicgiles.com/swingbench/swingbench22.pdf

Larry Pedigo *(lpedigo@perftuning.com) is a Senior Consultant/Instructor for Performance Tuning Corporation. He has been involved with Oracle RAC deployments, SAN deployments, and Oracle and SAN performance optimization for the past three years. Larry has a Bachelors degree in Geology from Texas A&M University, and over 16 years experience in the IT field.*