Response to "Reality behind Real Application Clusters Marketing Messages"

An Oracle White Paper
March, 2003

ORACLE

Response to "Reality behind RAC Marketing Messages"

Response to "Reality behind Real Application Clusters Marketing Messages"

EXECUTIVE OVERVIEW

This document is a response to specific allegations in Microsoft's paper **Reality behind Real Application Clusters Marketing Messages**. It is not intended as a general Oracle9*i* Real Application Clusters (Oracle9*i* RAC) versus Microsoft SQL Server 2000 comparison. Instead, we will focus on the technical aspects and assumptions made by Microsoft and provide point by point rebuttals. We will also provide both benchmark information, as well as customer references and quotes to support Oracle's claims that Oracle9*i* Real Application Clusters is for deploying all types of Real-World applications. It will become clear to the reader, after reading this paper, where the reality begins and where it is left behind.

Note. Quotes from the Microsoft paper are italicized at the beginning of each section.

SHARED CACHE VS FEDERATED DATABASES

"...SQL Server 2000 implements scale-out, by federating a group of independent databases to provide users with a view of single database... "

The advantages of a shared-cache database as implemented in Oracle9*i* RAC have been well documented. Two papers that discuss the architecture, deployment, and performance of these databases are:

- **Database Architecture: Federated vs. Clustered**
- **Technical Comparison of Oracle Database vs. SQL Server 2000: Focus on Performance**

Both of these papers can be found on Oracle's Technology Network web site: http://otn.oracle.com/deploy/performance/content.html

We briefly reiterate the important points:

- Federated databases use UNION views that combine data from separate database partitions. They do not support unique columns other than the primary key. The only way for the user to run real-world applications would be to manually modify the application by creating complex triggers. As such, SQL Server federated databases can not run real-world packaged applications such as SAP, PeopleSoft or e-Business Suites. The "single view" is provided only after modifying the application to include complex views and triggers. Even after performing these tedious and error-prone tasks, many important features such as constraints are not supported and the administrator does not have a single view of the database.
- Federated databases can not scale-out except for the most trivial applications. Any query which does not have the partitioning key in the WHERE clause will not benefit from more nodes.
- The complexity of managing a federated database increases with the number of nodes. Changes to the dictionary objects such as a new user or table being created must be manually replicated to all the nodes.

- A federated database does not improve availability by adding multiple nodes because the loss of a single node can make the entire database inaccessible. Failover clustering can remedy this situation but this not only increases the management complexity but also wastes hardware resources.

As pointed out in our above cited architecture paper, "For both the DBA as well as the Application Developer, there is a clear distinction between 'local' data, which is on the disk attached to a particular server, and 'remote' data, which is owned by another server in the federated database. Applications see a logical single view of the data through UNION ALL views and Distributed SQL – Microsoft calls this technology Distributed Partitioned Views (DPVs). The DPV is constructed differently at each node - it must explicitly consider which partitions are local and which are remote." So, for a one-time benchmark with a simple schema, such as the TPC-C,DPVs can be used, but in a real-world application environment which needs to use alternate key access, it is impossible. DPVs complicate everything from adding disks, to designing programs, to optimizing the execution plan for a query.

Microsoft's Distributed Partitioned Views are a benchmark toy. They are not applicable to real world applications.

So we are left with three simple facts about Microsoft SQL Server 2000 federated database:

- It does not provide a single view of the database to administrators, so manageability becomes an issue that will affect both scalability and availability
- It can not run real-world applications without complex modifications which are costly to maintain, as well as affecting both scalability and availability
- It does not provide for modular growth without affecting availability

A Microsoft SQL Server 2000 federated database is not comparable to a single cluster database with a single data dictionary, because it is a collection of loosely coupled databases (plural). And with that comes the cost and manageability challenges of administrating multiple databases, and handling routine system administration work. Also there is the problem of designing application programs to deal with the data partitioning that is needed. Consequently, any scale-out benchmark results published by Microsoft are meaningless because they do not apply to real-world applications.

REAL APPLICATION CLUSTERS SCALABILITY

"Oracle has not been able to back claims with proof in the form of top positions in industry standard benchmarks like TPC-C, TPC-W or LOB application benchmarks like SAP. "

In fact, Oracle has published world-record results for RAC with the TPC-C on an 8-node, 4 processor/node, Oracle9*i* RAC system on both Windows and Linux. As of December 16, 2002: Oracle9*i* Database Release 2 EE with Real Application

Clusters, HP ProLiant DL580R with 32 Intel Pentium III 900 MHz processors, 137,260.89 tpmC, $18.46/tpmC, available 09/06/02[1].

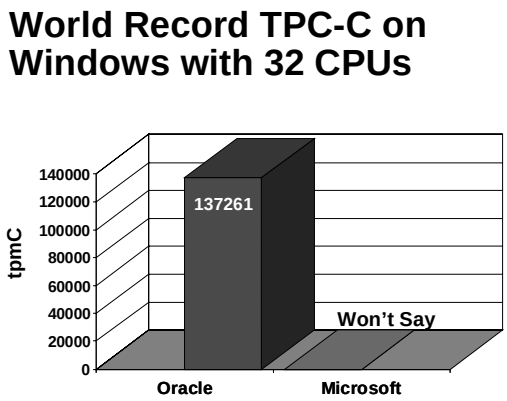

Source - Transaction Processing Council, as of December 16, 2002: Oracle9i Database Release 2 EE with Real Application Clusters, HP ProLiant DL580R with 32 Intel Pentium III 900 MHz processors, 137,260.89 tpmC, $18.46/tpmC, available 09/06/02.

Figure 1: TPC-C on Windows, 32 CPUs

Again, as of December 16, 2002: Oracle9*i* Database Release 2 with Real Application Clusters on Red Hat Linux Advanced Server, HP ProLiant DL580R,

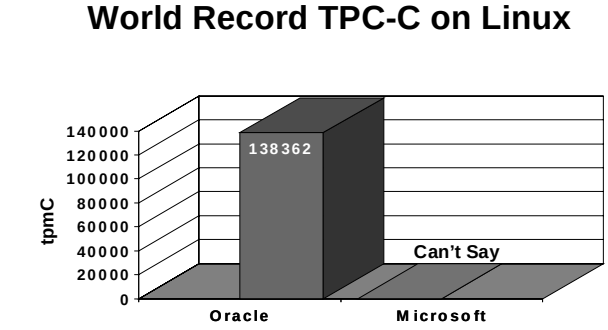

Source - Transaction Processing Council, as of December 16, 2002: Oracle9i Database Release 2 with Real Application Clusters on Red Hat Linux Advanced Server, HP ProLiant DL580R, with 32 Intel Pentium III 900 MHz processors, 138,362.025 tpmC, $17.38/tpmC, (original publication date, 9-16-02) available 03/05/03

Figure 2: TPC-C on Linux

with 32 Intel Pentium III 900 MHz processors, 138,362.025 tpmC, $17.38/tpmC, (original publication date, 9-16-02) available 03/05/03[2].

And Oracle9*i* RAC also holds the top positions for performance with the TPC-H at both the 100GB and 300 GB results. Oracle has also published numerous

[1] Source - Transaction Processing Council
[2] Source - Transaction Processing Council

benchmarks results for the TPC-C non-clustered systems holding 8 of the top 12 positions in the Top Ten Non-Clustered TPC-C by Performance (two ties).

Rank	System	tpmC	$/tpmC	Database	Operating System
2	**IBM eServer pSeries 690**	**427,760**	**17.75 US$**	**Oracle 9i Enterprise Database Server 9.2.0.1**	**IBM AIX 5L V5.2**
3	**HP 9000 Superdome Enterprise Server**	**423,414**	**15.64 US$**	**Oracle 9i Enterprise Database Server 9.2.0.1**	**HP UX 11.i 64-bit**
4	**IBM eServer pSeries 690 Turbo 7040-681**	**403,255**	**17.80 US$**	**Oracle 9i R2 Enterprise Edition**	**IBM AIX 5L V5.2**
***	**Bull Escala PL3200R**	**403,255**	**17.96US$**	**Oracle 9i R2 Enterprise Edition**	**IBM AIX 5L V5.2**
5	**HP 9000 Superdome Enterprise Server**	**389,434**	**16.41US$**	**Oracle 9i Database Enterprise Edition**	**HP UX 11.i 64-bit**
6	**NEC Express 5800/1320 Xc C/S w/Express 5800/120 Rd-**	**342,746**	**12.86US$**	**Microsoft SQL Server 2000 Enterprise Ed. 64-bit**	**Microsoft Windows .NET Server 2003 Datacenter Edit**
7	**Unisys ES7000 Orion 230 Enterprise Server**	**234,325**	**11.59US$**	**Microsoft SQL Server 2000 Enterprise Edition**	**Microsoft Windows .NET Server 2003 Datacenter Edt.**
8	**Compaq AlphaServer GS320**	**230,533**	**44.62 US$**	**Oracle 9i Database Enterprise Edition**	**Compaq Tru64 UNIX V5.1**
10	**IBM eServer pSeries 680 Model 7017-S85**	**220,807**	**29.30 US$**	**Oracle 8i Enterprise Edition v. 8.1.7**	**IBM AIX 4.3.3**
***	**Bull Escala EPC2450 c/s**	**220,807**	**34.67 US$**	**Oracle 8i Enterprise Edition v. 8.1.7**	**IBM AIX 4.3.3**

Table 1: Top Ten Non-Clustered TPC-C by Performance

*** - Duplicate results are shown with an asterisk (***) in the Rank column. Visit www.tpc.org for more details.

But with the TPC-C benchmark many of the elements are static, which is not representative of the real world. Hence, we refer to this as a synthetic benchmark and caution the interpretation of the results. Again citing the Oracle architecture paper where there is more detail, we will restate the two main reasons why we say this:

- TPC-C Schema Is Inherently Partitionable
- Most TPC-C SQL Accesses are Local

SAP Standard Application Benchmark simulates the full business workflow of how an order line item is processed.

In a real world application, Microsoft's federated database architecture doesn't scale out across multiple nodes. A good proof point of this is to take a look at a well know application benchmark. One such benchmark is SAP's Sales & Distribution Parallel Standard Application Benchmark, (SD-Parallel). The SAP SD-Parallel Standard Application Benchmark consists of a number of script files that simulate the full business workflow of how an order line item is processed:

- Creating The Order
- Creating A Delivery Note For This Order
- Displaying The Order
- Changing The Delivery
- Posting A Goods Issue
- Listing Orders
- Creating An Invoice

Although SAP has some 20,000+ tables installed (plus another 20,000+ indexes), the SD Parallel actively uses 140 hot tables in the benchmark with a high rate of concurrent inserts which are fairly long (4K and more) and causes fast growing tables. Because all changes (or requests for changes) are posted in a set of tables, these tables experience heavy update activity.

The benchmark is response time critical and it is important to get uniform response times from the database. In the benchmark, the simulated users expect their data to have been applied within 20 seconds after the request for an update has been made (or two dialog steps later). The benchmark results in a failure if this data is not present.

Likewise, a benchmark is only certifiable if the average dialog response time is less than two seconds. For this reason, a careful balance between the speed of processing both the update processes and the dialog processes is important.

Since the data distribution can significantly influence the benchmark result in a parallel environment, the SAP Benchmark Council established a rule for data

distribution in order to establish a means to reproduce and compare results within parallel benchmarks. It states that users for a certain client are equally distributed across all nodes. There is no data-to-node affinity (partitioning) with this benchmark. It causes a large amount of inter-node traffic and cache-to-cache transfers.

The SAP SD-Parallel Standard Application Benchmarks measure all performance relevant parameters such as database request times, wait times, CPU utilization, average dialog response by a given number of benchmark users (with a fixed think time of 10 seconds between each dialog step) and the achieved throughput.

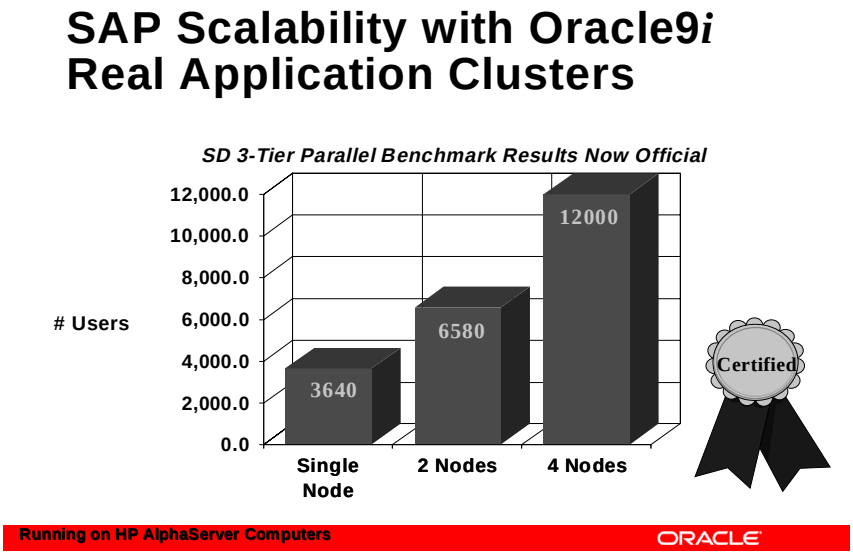

Figure 1:SAP Scalability with Oracle9i Real Application Clusters

In November 2001, Oracle and Compaq announced an initial performance test with SAP R/3 4.6C and the Oracle9*i* Database with Real Application Clusters that showed excellent scalability and throughput results. The testing environment included SAP® R/3® 4.6C and Oracle9*i* Real Application Clusters running the parallel SD workload on a four-node HP Tru64 UNIX ES45 AlphaServer cluster with 16 processors.

The robustness of Oracle9*i* RAC technology provided the foundation for near linear scaling that enables on-demand capacity planning and flexibility in deployment. The performance increased by a factor of 1.81 when going from a one-node configuration, 3640 users, to a two-node configuration, 6580 users, and by a factor of 3.30 going from a one- to four-node configuration, 12000 users. This is a world-record result in the SAP SD-Parallel benchmark. In fact, this four-node Oracle9*i* RAC result of 12,000 users is 20% better than the best 16-cpu result published with SQL Server 2000 (10,000 users). We do not intend this to

be used as a direct comparison because of hardware and software differences, but it clearly demonstrates the scalability of Oracle9*i* RAC.

The SD-Parallel benchmark puts considerable load on the highly available, clustered system. This is why this particular benchmark has become a de-facto standard in the ERP (Enterprise Resource Planning) industry. Results of benchmark tests are published by hardware partners to inform customers and interested parties in the ERP environment about the currently available performance and scalability of the mySAP.com business solution and the tested platform. As of this date Microsoft has not published a SD-Parallel benchmark.

Interconnect Traffic

"When each node attempts to process a very high number of transactions per node, as is common in any high OLTP environment, the amount of traffic traveling across the Interconnect ... will quickly flood the node's network cards..."

Oracle9*i* RAC has several built-in optimizations to reduce the number of inter-node messages.

1. Oracle9*i* RAC reduces inter-node messages by caching data in each node. After a process obtains a block either from another node or from the disk, accesses to the block do not result in inter-node communication. Thus the cost of inter-node communication to obtain a block is amortized over several transactions that subsequently access the block until a remote node needs the block in a conflicting mode. In contrast, transactions on shared-nothing and federated databases will need to send individual messages to remote-nodes even when the transactions access the same block.
2. Oracle's consistent read isolation allows transactions in one node to read a block while transactions in another node are modifying the block at the same time. This significantly reduces the frequency of communication between nodes.
3. Oracle9*i* RAC automatically optimizes the data accesses of the transactions running on different nodes to reduce inter-node communication. For example, insert transactions to the same table from different nodes access different blocks.
4. The global cache service protocol is highly optimized so that in those cases where a block needs to be obtained from a remote node, a minimum number of messages are sent.

By reducing the frequency of inter-node communication, Oracle9*i* RAC allows many demanding OLTP applications to run on clusters with commodity interconnects such as Gigabit Ethernet. Network interfaces for such interconnects are now a standard component in server hardware.

Oracle’s sweet spot for RAC is 4 CPUs per node

“Oracle's sweet spot for RAC is 4-CPUs per node. “

Oracle9*i* RAC allows a cluster database to be built out of independent nodes where the nodes are chosen not for headroom but for best price/performance. Not only does Oracle9*i* RAC allow customers to choose from the largest selection of hardware and OS vendors, but they can also choose the size of each node depending on their unique requirements. Users run Oracle9*i* RAC in production under a variety of configurations, from 2-CPUs per node to larger 36-CPUs per node SMPs. And even larger systems are in pre production testing.

Benchmark Scalability

“Specific to scale-out benchmarks, SQL Server has displayed scalability with node progression (16, 24 and 32 nodes) while Oracle RAC does (sic) has not...”

The only benchmark that SQL Server federated database can run is the TPC-C benchmark which is a trivial OLTP application.

SQL Server federated database can not run the packaged applications such as SAP, let alone demonstrate scalability. In contrast, Oracle has published benchmark results on RAC for both simple OLTP applications such as TPC-C as well as complex real-world packaged applications such as SAP, Peoplesoft and Oracle E-business Suite.

"Our customers gauge our level of commitment to the best technology from how well our site functions. We chose Oracle because we need to run both our front- and back-end operations without a hitch. Oracle9*i* Database meets our performance demands, and Oracle9*i* Real Application Clusters and Oracle Data Guard doubly ensure our availability." -- David MacDonald, President, Softchoice Corporation

One example of a real world application running Oracle9*i* RAC with Oracle E-business Suite is Softchoice, Toronto, Ontario (www.softchoice.com). Softchoice provides businesses and organizations of all sizes with a fast, flexible, and cost-effective way to research, buy, and manage software and hardware technology resources. Softchoice's dedication to simplifying software purchasing has made it the fastest growing company in its market segment in the United States and number one in Canada. North American businesses and organizations of all sizes turn to Softchoice for a fast, flexible, and cost-effective way to research, buy, and manage software and hardware resources.

David MacDonald, President, Softchoice Corporation stated, "Our customers gauge our level of commitment to the best technology from how well our site functions. We chose Oracle because we need to run both our front- and back-end operations without a hitch. Oracle9*i* Database meets our performance demands, and Oracle9*i* Real Application Clusters and Oracle Data Guard doubly ensure our availability."

The company's continuing success depends on giving customers easy access to the latest technology, therefore, its site and enterprise resource planning (ERP) operations must function flawlessly. When Softchoice found it had outgrown its legacy ERP systems, it decided to replace them with Oracle E-Business Suite, based on Oracle's integration and the economy of using its self-service functionality worldwide. At the same time, Softchoice realized a need to improve its underlying infrastructure to give the IT team the ability to focus on projects

that would win customers, instead of focusing on system management. Lastly, Softchoice wanted to improve availability beyond the 99% previously achieved with Oracle8*i* and to completely safeguard against data loss.

Using Oracle9*i* RAC to centralize application installations into one central location has reduced yearly operating costs. The company's three major call centers now use a centralized operating system, which reduces its hardware investment, database license fees, and hardware and software maintenance costs. In addition, Softchoice credits Oracle9*i* RAC for eliminating the need for additional skilled IT staff.

The Oracle9*i* Database runs on a four-node server cluster. Should one node fail, the applications will remain up and running on the other nodes. RAC's failover is transparent, so, should failure occur, another node takes over so quickly that users are not even aware of a problem. You can read more at www.oracle.com/customers/9i/softchoicerac.html

REAL APPLICATION CLUSTERS AVAILABILITY

"RAC does not provide automatic high availability out of the box."

When any of the Oracle9*i* RAC nodes fails, the client connections are automatically routed to the available nodes. The load handled by the failed node can be evenly redistributed to the surviving nodes.

In fact, it is the Microsoft SQL Server federated database architecture that does not provide automatic high availability out of the box. This is because each database partition needs to be assigned a failover node and the administrator may need to create multiple database partitions to distribute the load evenly. To avoid load skew, Microsoft SQL Server federated databases must be configured such that each surviving node takes over ownership of the same amount of data. This is achieved by creating multiple databases on each node. For example, if there are four nodes, three databases are created on each node for a total of twelve databases. If there are n nodes, for even redistribution of databases to the surviving nodes under all failure scenarios, the number of databases equals the least common multiple of n, n-1, ...,1. For each database a preferred owner or takeover list is created using the cluster software (such as MSCS) so that the databases are evenly redistributed across the nodes.

A database's availability is impacted by both planned and unplanned downtime. In a federated database, a node can neither be added to the cluster nor removed from the cluster without taking the whole set of databases off-line. This is because the administrator has to manually modify the distributed views when nodes are added or removed. This process can easily take several days for a complex application! With Oracle9*i* RAC, there is no interruption to database accesses when a node is added or removed. This reduces the need for planned downtime (to apply an OS patch, for example) and increases availability. In fact, Microsoft itself

acknowledges this: *"Oracle9i RAC does simplify the process of adding and removing nodes to a cluster."*

Raising the High Availability Bar

"RAC does not raise the bar for high availability. "

Unplanned downtime is caused by software or hardware bugs and human errors. Oracle has many features to recover from human errors which we do not describe here. Oracle9*i* RAC has raised the bar for high availability through several optimizations that reduce the time to recover from a node failure.

- Oracle9*i* RAC allows global enqueue operations to resume faster by reconfiguring the cluster faster.
- The blocks that do not need to be recovered are filtered immediately through block-written-records in the redo log.
- Global cache operations can resume faster because of recovery "claims" by the recovering instance for blocks that belong to the recovery set.
- Oracle9*i* RAC can use the contents of surviving nodes' caches to avoid disk I/Os to recover blocks.

"If you need 8-CPUs and 16GB RAM in a single server environment and you choose to deploy a 2-node cluster system for the same transaction, each node in the cluster must have the same 8-CPU and 16GB RAM configuration. "

Oracle9*i* RAC allows the user to configure the cluster in many ways and lets the user decide which services need to be failed over. Microsoft has chosen to describe only one possible configuration. In the scenario described above the user has at least two additional choices:

1. Deploy a two-node cluster system where each node has 4 CPUs and 8GB RAM. Remember that Oracle9*i* RAC runs active on both nodes so you use two lower cost SMP boxes for the same full workload. The user can configure Oracle9*i* RAC such that when a node fails, the critical services are failed over to the surviving node. With Microsoft's failover clustering the user will need to deploy at least one node with 8-CPUs and 16GB RAM for the full workload and have another 4-CPUs and 8GB RAM SMP system standing by for the failure-induced smaller workload because Microsoft does not provide a way for the load to be balanced across multiple nodes when the nodes are up.

Best of all, the user gets higher availability with Oracle9*i* RAC by buying less hardware!

2. Deploy a 3 node cluster system where each node has 4-CPUs and 8GB RAM. When a node fails, the load from all services can be handled by the two surviving nodes. If two nodes fail at the same time (which has a lower probability), the critical services can be handled by the surviving node. With this configuration, the user has the same availability as Microsoft two-node failover clustering when there is a single-node failure and better availability than Microsoft's two-node

failover clustering when two nodes fail - With Oracle9*i* RAC, the surviving node can continue to provide critical services whereas the entire database is off-line with Microsoft's two-node failover clustering. Best of all, the user gets higher availability by buying less hardware! The three-node Oracle9*i* RAC system needs 12-CPUs and 24GB RAM whereas the two- node Microsoft system needs 16-CPUs and 32GB RAM.

This is only a few of the many possible scenarios that you could deploy using Oracle9*i* RAC adaptable architecture. Oracle9*i* RAC allows a single database to be accessed seamlessly and simultaneously from all nodes in the cluster database - Microsoft does not have the architecture to do this. Hence, Oracle9*i* RAC gives much more flexibility to the user in deploying a high-availability configuration depending on the user's unique requirements, their price/performance goals, and their choice of platforms.

Maybe it is best if we let some of our customers tell you about RAC and High Availability.

BT Group PLC, London, England (www.btplc.com) is the listed holding company for the BT Group of companies, while British Telecommunications PLC (BT) is a wholly owned subsidiary. BT is one of the world's leading providers of telecommunications services and one of the largest private sector companies in Europe.

As one of Europe's leading communications companies, BT's mission is to help customers exploit the business opportunities that rapidly advancing communications technologies provide. For customers of its hosted and managed services, achieving the highest possible levels of availability for their database and applications is one of the most important criteria.

"Customers always want the highest availability possible for their core applications such as e-mail andinternet servers," says David Seddon, UNIX TP and Database Manager, BT. "Downtime affects their business and their own customers' business and results in lost revenue. Before Oracle9*i* RAC, guaranteeing high-availability could only be achieved through complex solutions that involved duplicate hardware and software, together with sophisticated levels of support."

"Before Oracle9*i* RAC, guaranteeing high-availability could only be achieved through complex solutions that involved duplicate hardware and software, together with sophisticated levels of support." -- David Seddon UNIX TP and Database Manager, BT

Without RAC, failover typically lasted up to 20 minutes, during which time users lost access to the system and had to log back on to resume operations. RAC achieved failover times of between ten seconds and one minute. During this period, processes being performed by the failed component were automatically transferred to the remaining node, with users being completely unaware of the failure. RAC environments are fully redundant because all nodes access all the disks in the cluster. There is no single point of failure in the system and the failure of one node does not affect another node's ability to process transactions. The database is available on any surviving instance following an instance failure.

'The whole set-up is simple, with many processes automated, and out-of-the-box functionality that does not require high-level management resources." --Reece Peacock, UNIX TP and Database Engineer, BT

RAC enables rapid database diagnostics by giving database administrators improved visibility of the sequence of events before and after a failure. 'Tracing every stage through to Oracle pick up and recovery was straightforward and easy to follow,' saysReece Peacock, UNIX TP and Database Engineer, BT. 'The whole set-up is simple, with many processes automated, and out-of-the-box functionality that does not require high-level management resources." You can read more at www.oracle.com/customers/9i/bt.html

RAC's scalability also reduces costs by eliminating the need to deploy standby servers that stand idle most of the time, in order to ensure that operations continue in the event of a failure. As long as the cluster has one surviving node, all database clients can continue to process all transactions, although response times may be increased due to capacity constraints on the one node.

Another proof point of RAC's high availability and scalability is Electronic Arts' The Sims Online application. Electronic Arts Inc. (EA), Redwood City, CA (www.ea.com) is the world's leading interactive entertainment software company. Its EA.com division creates and delivers about 45% of all online game playing, with about 15 million active players.

On December 17, EA launched TheSims Online, an online version of its popular PC game The Sims. EA expects about 250,000 active players with the initial launch of The Sims Online. Every character in the game is associated with a great deal of data, including physical attributes, personality characteristics, and friendship connections, and all characters can interact with one another. Given the data complexity and sheer volumes of data that supports TheSims Online, EA had to select and implement a very robust architecture and design.

The company also needed to be sure the database server wouldn't fail. Oracle9 RAC protects EA against server failure without the added expense of a standby system. "We run application clusters for high availability, because if you blow out one of the back-end databases when you have 100,000 players online, you lose a lot of customers quickly," saidMarc West, Senior Vice President and Worldwide Chief Information Officer, Electronic Arts Inc. "We have a 24/7, 99.99% mentality."

"We run application clusters for high availability, because if you blow out one of the back-end databases when you have 100,000 players online, you lose a lot of customers quickly. We have a 24/7, 99.99% mentality. Oracle9*i* RAC works as advertised ." -- Mark West, SVP & CIO, EA, Inc.

EA wanted to improve theSims game play by speeding up response time, and was able to do that using Oracle9*i*. The company also needed to be able to scale The Sims Online cost-effectively as the player community grows. While EA had an existing Unix platform that the company could have used, after extensive research and testing EA chose Oracle9*i* Real Application Clusters (RAC) running Linux on an Intel platform. The core measure was the number of queries that could be completed per second; Oracle allowed EA to meet its goal of 30,000 SQL calls per second, per small hardware cluster. EA also looked at stability and cost-effectiveness. "The database is an integral part of the game design," said West. At launch time, the system will be about 1.5terabytes, and EA expects it to expand to 6 to 8terabytes as the game grows.

West said EA could not have met its price performance objectives for The Sims Online without Oracle9*i* RAC. "When you look at the price and value of these game titles, you have to turn around and say they're worth an investment. You want to make the maximum return, and that's what this implementation is about." EA will be able to add clusters as needed without downtime or unnecessary expense. "Oracle9*i* RAC works as advertised," said West. You can read more at www.oracle.com/customers/profiles/profile8751.html

TOTAL COST OF OWNERSHIP

"There is a high likelihood that total cost of ownership will increase with RAC because the user will now need to purchase additional server(s)... "

Oracle9*i* RAC does not require any additional servers beyond what is needed to support the expected load. Oracle9*i* RAC allows the user to deploy additional server(s) in the clustered database dynamically on an as-needed basis. Unlike federated databases, you don't need to bring down the application and database because there is no need to modify any dictionary objects when a new server is added to handle the increased load.

"The user will now need to purchase networking hardware... "

Again, Oracle9*i* RAC does not require expensive networking hardware. When an interconnect does support features such as remote-DMA, Oracle9*i* RAC will fully exploit it, but it is not required. Many real-world OLTP applications can run with acceptable performance with commodity interconnects such as Gigabit Ethernet. This is because Oracle9*i* RAC has many built-in optimizations to limit the interconnect traffic. In fact Microsoft itself acknowledges that Oracle9*i* RAC can run on commodity interconnects. In the description for a typical Oracle9*i* RAC setup, Microsoft states, *"The Interconnect between all nodes uses standard Network Inteface Cards (NICs) on each node ."*

"The user will now need to purchase storage solutions (direct attached storage does not work for RAC). "

Storage is the largest cost component in the implementation and operation of mission-critical applications, as well as the key to continuous operations. And while the initial cost of direct attach storage is somewhat less expensive than current SAN solutions, a centralized, consolidated storage with universal access provides many other benefits which reduces the Total Cost of Ownership (TCO):

- SANs are easier to manage. Analysts estimate a single administrator can manage from 3 to 50 times the storage (depending on the size of the installation) when using consolidated storage. A single storage infrastructure makes it easier to control, backup and improve the overall availability of the data. SANs allow storage managers to enforce corporate storage policies that accomplish specific functions, such as preventing users from storing certain types of files or data

exceeding a certain age or size, or with a certain number of duplicates, be moved, archived, or deleted.

- Maximizing availability is also key to optimizing costs. RAID, data mirroring, and clustering reduce the potential financial impact of downtime. SANs operate over a multi-channel switched-fabric, providing continuous access to data from all the systems. Careful configuration of backup and recovery processes is crucial. Access is a critical dimension of continuous operations, but internal, bus-attached storage contributes nothing to improved access or availability. Knowing how much an outage would cost your business, how quickly you could be back online, and what policies and procedures could accelerate that process are simplified in a centrally managed storage environment.
- SANs are flexible, robust and provide unlimited scalability. Adding more storage capacity on demand as an application grows is easier. Because servers share resources, excess storage capacity can be deployed where it is needed. Scale out is only limited by the capacity of the network fabric. It is easilyreconfigureable and will tolerate changing business demands, providing high value with a low TCO.
- SANs provide better storage utilization. It is estimated that as much as 60 percent of enterprise storage capacity can remain unutilized often because spare capacity on one system can not be used on another when it is direct-attached storage. By separating storage from your servers organizations can deploy storage resources according to business needs rather than technical constraints. And it provides the ability to purchase larger, higher-capacity disk arrays which typically results in a lower cost per gigabit.

And we also practice what we preach. Oracle Corporation turned to EMC to help handle more than 240 terabytes of data storage in 43 data centers worldwide. Using multiple EMC products, including Symmetrix, Connectrix, Celerra file server, TimeFinder and PowerPath, these centers were consolidated into a dual data center environment that has cut downtime, making core applications continuously available. By centralizing server storage, Oracle has lowered power consumption and floor space by 30 percent.

"The user will need to train operators and administrators."

Oracle9*i* RAC does not require any new administration skills beyond what is needed to administer a single-node database. All administration commands for a single-node Oracle9*i* database work identically in an Oracle9*i* RAC database. Indeed, this a key distinguishing feature between Oracle9*i* RAC shared database architecture and Microsoft's federation of databases.

It is with federated databases, that an administrator needs to perform complex tasks to create the distributed views and replicate dictionary objects. With

federated databases, the complexity and number of tasks also increases with the number of nodes.

"The user will need to purchase services for migration. "

When deploying Oracle9*i* RAC, neither the application nor the database needs to be migrated or modified. Again, this is a key distinguishing feature between Oracle9*i* RAC and federated databases.

"The user will need to purchase services for certification. "

Oracle9*i* RAC is certified on all major hardware platforms and all the major operating systems. Complex packaged applications such as SAP are also certified on Oracle9*i* RAC. There is no cost for certification. In contrast, SQL Server 2000 federated database runs only on Windows and is not certified with any packaged application.

A recent customer study shows that the cost of running packaged applications (SAP, PeopleSoft, Siebel, …) on an Oracle Database is significantly lower than running those same packages on Microsoft SQL Server. The study, "**Comparative Study of Relational Databases Underlying Packaged Applications**", was conducted by INPUT, a research firm that interviewed CIOs, IT managers, development managers and DBAs at 30 organizations using 44 different application solutions. The full report is available at : www.oracle.com/features/facts/analyst_reports/DB_TCO_Packaged_Apps.pdf

The highlight of the study include:

- The annual per user cost of running packaged applications on an Oracle Database is half that of Microsoft. This includes all cost of ownership components: acquisition, implementation, and on-going costs for database, hardware, and application software. The per user costs directly attributable to the database (license, implementation, and on-going costs) are on average 43% lower with Oracle than with Microsoft.
- With Oracle, database administrators are more productive: On average, an Oracle DBA supports more than 3,000 users; a Microsoft DBA only supports 684 users.
- It is four times more expensive to secure an application environment with Microsoft SQL Server compared to doing so with an Oracle Database
- Oracle remains the preferred database for packaged applications: while customers report a similar level of satisfaction with Microsoft SQL Server and Oracle Database, Oracle is named more frequently than Microsoft as the database vendor that they will consider for future purchases.

So, while Microsoft attracts customers with a lower cost of entry, this research confirms that companies that run their enterprise applications on an Oracle Database benefit from a significantly lower cost of ownership.

CONCLUSION

This document refutes all of the claims made by Microsoft by providing you with both technical information to support our claims for Oracle9*i* RAC and examples of real world applications that demand Oracle's open cluster technology. Microsoft's claims are not supported by facts. Further, when making the claims, Microsoft has tried to confuse the reader into believing that some features are available in its cluster database, when in truth Microsoft does not have a cluster database.

Oracle9*i* Real Application Clusters' scaleable, highly available, flexible database architecture allows our customers a choice of hardware platform vendors, software system vendors and a competitive price/performance advantage. They can start small and expand at will with our revolutionary clustering technology. They can choose open source Linux, utilizing inexpensive commodity components, or they can use large SMP systems with very high speed interconnects to run their packaged application or their custom application. And Oracle9*i* RAC will even make Windows more reliable and secure than any other database.

Response to "Reality behind Real Application Clusters Marketing Messages"
March, 2003
Authors: Sashikanth Chandrasekaran, Bill Kehoe
Contributing Authors:
Editor: Mark Bauer

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
www.oracle.com

Oracle Corporation provides the software
that powers the internet.