

Oracle Real Application Clusters 10g – the Foundation for Enterprise Grid Computing

*An Oracle White Paper
September 2005*

Oracle Real Application Clusters 10g – the Foundation for Enterprise Grid Computing

ORACLE REAL APPLICATION CLUSTERS – FUNDAMENTAL TO ACHIEVING ENTERPRISE GRID COMPUTING BENEFITS

Oracle Real Application Clusters (RAC) is a key enabling technology for Oracle Enterprise Grid computing. Oracle RAC 10g Release 2 offers a complete grid infrastructure for each layer of the enterprise grid - application, database and storage grids.

In the past few years, there have been a number of significant improvements in computing hardware, networking, chip architectures, data storage systems and in standards-based open-source operating systems. Oracle is the first vendor that offers a complete enterprise computing software infrastructure that leverages all of these improvements to deliver significant economic return on investment through better availability, scalability, flexibility, and manageability. Oracle Real Application Clusters (RAC) is a key element in this unique, popular and successful solution.

Whether you run an enterprise data center serving 200,000 users, a small to medium-sized operation serving 200 users or a web-based business handling 25,000 transactions per hour, you are concerned with economically maximizing the availability, scalability, flexibility and the manageability of your computing resource.

Oracle's Real Application Clusters, introduced in 2001, offers a proven alternative to traditional expensive hardware and proprietary software solutions for achieving these goals. This paper presents several scenarios representing the common questions or decisions IT executives face and how enterprise customers have used Oracle Real Application Clusters to address these business needs.

Oracle Real Application Clusters – the foundation for Enterprise Grid Computing

In June 2001, Oracle released Oracle9i with Real Application Clusters. Oracle9i RAC made it possible for a collection of database servers to cooperate transparently in the management of a single Oracle database. This allowed Oracle to deliver greater scalability and availability than had previously been possible while simultaneously reducing costs and improving flexibility. In January 2004 Oracle released the second generation of its clustered database technology, Oracle Real Application Clusters 10g release 1 and in June of 2005, 10g Release 2.

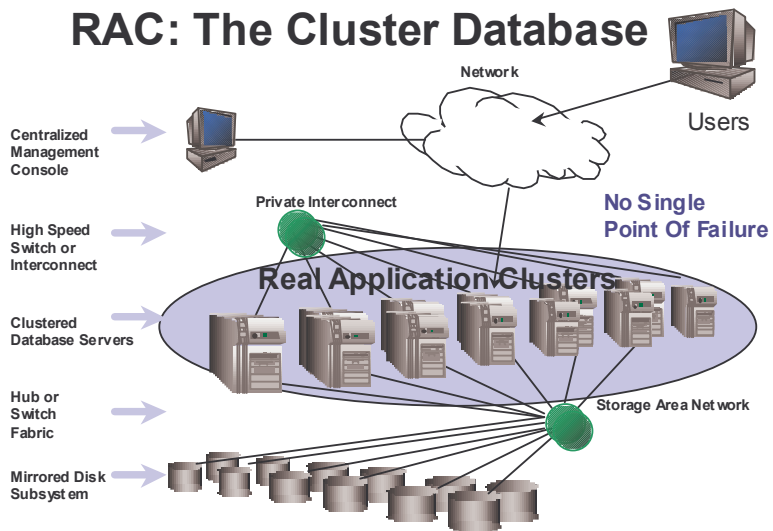


Figure 1: Oracle RAC –clustering database servers – foundation for Enterprise Grid Computing delivering high availability, scalability and flexibility.

QUESTIONS AND DECISIONS

We speak with many companies throughout the course of the year, many CIOs, CTOs and executives responsible for the computing infrastructure for a multitude of enterprises, large to small. Increasingly, they look to vendors such as Oracle for advice and guidance on the decisions and questions they face like:

- **“My SMP infrastructure reaches its financial and technology refresh point in fourteen months. What do we do? Is this an opportunity to consider an enterprise grid architecture based upon Oracle 10g Real Application Clusters?” A major online retailer made the shift and we’ve documented it here.**

- **“Our business is successful and growing – that’s the good news and the bad news. We seem to reach capacity on our critical systems and it requires a forklift upgrade to get more scalability. Expensive and disruptive. Will Oracle Real Application Clusters 10g and the Oracle Grid vision provide a less disruptive, more flexible and**

scalable solution?” **The Chicago Stock Exchange tackled this concern with Oracle RAC as told below.**

- **“Our average server utilization is well below 50% -- sure we have spikes to over 90% but that happens at different times of the month for each of the various applications we’ve dedicated to these systems.” Will Oracle’s 10g grid solution improve our utilization AND enhance the availability, scalability and flexibility of my environment? **Talk America is a big proponent of Oracle 10g Real Application Clusters because it enables a Service Oriented Architecture which virtualizes the underlying resources providing better hardware utilization, flexibility, improved SLAs and lower cost.****

Sticker Shock of nearly \$10M for a UNIX-SMP platform refresh made the decision to move to a of RAC Linux-based cluster easy for this well-known e-retailer. RAC running on low cost servers enabled them to avoid \$6 million in capital outlays.

E-RETAILER - SMP REFRESH? NO - OPTING INSTEAD FOR LINUX-BASED REAL APPLICATION CLUSTERS FOR THEIR MISSION CRITICAL APPLICATIONS.

Well-run companies manage their capital budgets and infrastructure spending very carefully. Two years ago, the core one of the largest e-retailers computing infrastructure was a large SMP system handling OLTP and Data Warehousing activities. That asset was due for refresh/replacement. This system ran a number of applications and Oracle databases from their Catalog, Customer database and their order/shipment/payment database.

What they thought was going to be a simple contract renewal wasn’t. This customer suffered severe sticker shock when presented with a nearly \$10M million quote to refresh the SMP UNIX environment that had been their mainstay. At this point, the customer had to look for other ways to achieve their high availability and scalability goals while maintaining their performance levels.

Instead, they architected a cluster of Intel-based commodity servers running Linux. Oracle Real Application Clusters with its embedded clusterware and cluster file system. Deploying on twenty-eight nodes in three clusters (one containing 16 nodes), this e-retailer was able to migrate to an available and scalable environment for a total capital cost of 25% of what was initially proposed by the hardware vendor saving them nearly \$6M in capital outlays. The annual support and maintenance fees were correspondingly reduced by 70%. Already familiar with the Oracle Database operating in a single instance mode on their SMP system, this e-retailer felt comfortable considering a Real Applications Clusters solution.

High Availability: One of the key drivers to move to Oracle RAC for this e-Retailer was its compelling “active/active” architecture and the high availability that offered

while providing better resource utilization.. With Oracle RAC, if a node in the cluster fails, the database keeps on running. All other nodes in the cluster remain up and users continue processing. With a single SMP server, if anything fails the whole machine goes down - user processing is interrupted and can take 20 to 30 minutes to resume on a cold standby (active/passive) system. In a cold/standby environment, when the primary system fails, administrators have to manually dismount and remount storage volumes, user connections have to be established to the new server and the buffer cache in the standby's memory needs to be populated with the most frequently used data. A significant part of the cost savings this e-Retailer experienced was due to not having to purchase and warehouse a cold failover SMP system.

Oracle Real Application Cluster's active/active system offers virtually uninterrupted processing for users and significantly better, more complete utilization of computing resources as compared to the traditional cold standby solutions. Oracle's RAC eliminates the need for these recovery steps – all cluster nodes are always connected to all the storage so no volume dismounting and remounting is required. Users on the surviving nodes remain connected only the users on the failed nodes need to reconnect and RAC does this:

- rapidly and transparently such that users often don't even know a failure has occurred
- automatically – no need for administrator intervention and pre-set failover preferences to specified nodes can be established through Oracle Enterprise Manager.

Oracle RAC systems can be configured to have no single point of failure, even when running on low-cost, commodity hardware and storage. In an Oracle RAC environment if database servers fail their applications simply keep running. Fail-over is frequently transparent to applications and occurs in seconds. Oracle also protects from major site failures by combining Oracle RAC with Oracle Data Guard – part of Oracle's Maximum Availability Architecture (MAA).

Our only question was: Real Application Clusters or not?" said David Milne, director of database technologies for the Chicago Stock Exchange. "Oracle's clustering technology was irresistible. Most important was its reliability and availability. Its selling points were ease and flexibility in adding large or small hardware as needed, where needed, and no changes required to applications."

THE CHICAGO STOCK EXCHANGE - MORE CAPACITY WITHOUT THE HEAVY LIFTING

The Chicago Stock Exchange (CHX) is the third most active stock exchange in the United States by volume. It is a technologically advanced exchange, providing fast, quality executions. Due to the nature of the business, continuous uptime, scalability on demand, and rock-solid performance are absolutely critical.

Facing shrinking revenues in the late 1990s, CHX's IT department needed to find a way to do more with less money. At the same time, it needed to eliminate bottlenecks created by massive data loads on heavy trading days. Seeking peak reliability and performance, CHX chose Oracle Database with Real Application Clusters on HP hardware for its optimal solution. CHX's needs were growing, most

likely requiring it to procure the next class of SMP machine. But, CHX was concerned about significantly over-buying or under-buying server capacity – trading volumes spiked by 2 times in 2000 but dropped significantly in 2002. They also were concerned that as newer technology emerged, they would not be able to take advantage of it without replacing the entire server infrastructure.

In a two node cluster, a hardware failure on one node could negatively impact all users, and result in lost trading revenues. Migrating to a four node RAC/Grid environment significantly reduced the exposure to a single hardware failure. Oracle RAC allowed the Chicago Stock Exchange to easily and quickly add more low-cost servers to the cluster. This has significantly reduced their exposure and extended the useful life of existing hardware. With Oracle RAC a hardware failure on one node is not catastrophic, the workload is automatically shifted to other hardware in the Grid. Most of the time the users are not aware of the failure. In addition, the results of a server failure are less urgent in the Grid environment, while still a critical event, it is not necessarily the “all hands on deck” scenario.

Improved Scalability, Flexibility and Manageability can be achieved

Enterprises have also had to over-provision to accommodate changes in business transaction volume and in the spikiness of web-based transaction activity. Scaling in a traditional SMP environment required long lead time advance planning due to the lead times in obtaining the hardware, preparing the environment and shifting environments.

Scale Out: Oracle RAC allows multiple servers in a cluster to transparently manage a single database. Oracle RAC allows database systems to scale-out rather than having to scale-up to meet additional workload increases. In the scale-up model of computing, once a server has been fully configured with CPUs and memory, the next step is an expensive "fork-lift upgrade", but with Oracle RAC users just plug in another server. if additional processing power is needed.

RAC also provides the opportunity to rapidly retire excess capacity and re-deploy that for a different purpose either as a node in an application server mid-tier layer or as an additional database cluster node in a different cluster.

In addition to eliminating scalability constraints, RAC has proven to scale better than SMP when evaluated in a head-to-head comparison. Given the same number and type of CPUs, the same amount of total memory, RAC has proven to offer greater scalability and processing throughput than a big SMP environment.

Project MegaGrid – Proving “scale out” as an economically attractive alternative

As a founding partner of Project MegaGrid, in phase one of this multi-phase collaboration with Dell Computers, EMC and Intel, compelling evidence of the power of a low-cost RAC environment was gathered. Starting in 2004, Oracle, Dell, EMC and Intel came together to validate Oracle’s grid computing vision and establish grid computing best practices as a way to encourage adoption and reduce

market confusion and hesitation to embrace this new model. In phase 1, a standard telecom service provisioning application from Cramer Systems, used by telecom companies such as British Telecom and Bell South, was benchmarked on both a large 72CPU SMP system and an Oracle RAC cluster of low-cost commodity Dell servers. The goal was to process as many transactions of a mixed workload per hour while maintaining a maximum four-second response time. The \$2.8 million SMP server environment was able to handle 550,000 transactions per hour. The same volume of transactions while maintaining the SLAs was achieved on a 10-node, 20 CPU Dell PowerEdge 1750 server cluster costing about \$60,000. Another test but maintaining the same SLAs was achieved on a four-node four CPU 64-bit Dell 7250 server cluster costing \$160,000.




Figure 2: Clustered, low-cost servers running Real Application Clusters can offer an economic, availability and scalability advantages over traditional architectures

“The beauty of Oracle 10g and a RAC-enabled grid is that the choice is mine. I can manage SLAs based upon my business needs.”
 Laurence Grant – Vice President of Enterprise Computing Systems – Talk America

TALK AMERICA

In business for 16 years, Talk America provides local and long distance telephone services to residential and small business customers in the United States. Over the last five years, their emphasis has shifted more to local telephone service where the marketplace for those services is very active and where customer service and where Talk America’s customer service and responsiveness is a competitive advantage – something they take so seriously they were awarded the J.D. Powers Service

Excellence award for top customer satisfaction rating in the telecommunications segment. Core to Talk America's competitiveness is their IT operation and the proprietary integrated order processing, provisioning, billing, payment, collection, and customer service information systems which enables them to offer competitively priced service plans, high-quality service and simplicity through consolidated billing and responsive customer care.

Fundamental to high-quality customer service is a highly available systems environment. Talk America's legacy environment, based upon 20 different Informix databases, was preventing them from meeting their business goals – database reliability was a problem. Additionally, Talk America was facing a number of other challenges that was stressing their IT environment:

- Rapid growth of their business and the exponential growth of call records and transaction volumes
- Need to provide a 24x7 operation
- Increasing systems dependencies – less time for maintenance
- Shrinking nightly batch windows
- Growing multitude of systems to satisfy growing end-user, internal users (sales/service) and increased reporting requirements (Sarbanes/Oxley etc.)
- Much larger data volumes growing to 60 terabytes from only 1 terabyte eight years earlier

Talk America had to make some changes – changes that mandated an upgrade of their infrastructure to a high performance, highly available, flexible environment.

Laurence Grant, Vice President of Enterprise Computing Systems at Talk America, and Vishal Anand, Director of Database Administration, were concerned that simply upgrading their silo'ed environment would be expensive and would not really solve the problems they faced. Invariably, some decisions about server and storage capacity would result in an over allocation for some applications and under allocation for others.

Talk America decided to pursue Oracle Database 10g and Real Application Clusters as their core database and high availability infrastructure. Even though the solution was very new at the time, Larry and Vishal believed that the advantages 10g brought to their problems far outweighed the risks of going with a relatively new product release.

Oracle Database 10g and Real Application Clusters – the underpinning of a Service Oriented Architecture:

Talk America was interested in the economies that could be gained if they consolidated databases and migrated to a service-oriented architecture. It meant greater hardware utilization, more flexibility to move resources to accommodate spikes in workload demands, and cost savings by not having to invest in an expensive, dedicated solution for each major system. Talk America had experience with big iron solutions and they didn't want to buy hardware every six, nine or twelve months to keep pace with the business – they wanted to avoid the business disruption and the additional costs of shuffling equipment in and out. Oracle's 10g RAC was the key that enabled the high availability and workload sharing and balancing allowing Talk America to gain better resource utilization and avoid changing the hardware footprint every few months.

Talk America wanted an environment that would enable them to manage resources to satisfy the various service level agreements (SLAs) established for their users. They took advantage of the 10g Grid capabilities and consolidated all their 20 separate databases into one 60TB Oracle database. The various services such as Order Processing, Provisioning, Leads Processing, Data Marts, CRM etc. run on the RAC cluster and share access to the single database – administration is easier, availability is higher by virtue of Oracle RAC, and flexibility is greater in terms of adding or removing processing power to match the needs of the application at any point.

Leveraging All the Hardware All the Time

The Oracle Database 10g service-oriented architecture enabled by Oracle Real Application Clusters, binds services to the underlying hardware – but in a flexible, virtual way. It is very easy within a 10g architecture to run OLTP applications and data warehousing within the same RAC cluster. And it's easy to change the allocation of resources as business processing needs change – hence leveraging all the hardware all the time.

Larry Grant explains the one of the key advantages an Oracle Real Application Clusters 10g - enabled grid provides.

“A simple example, suppose I have OLTP service allocated to four nodes. If one node goes down I'm still highly available because three nodes are up and running although you might be running with degraded performance. That may be acceptable, however, as long as you're running. But what Oracle RAC 10g says, I've got four nodes running OLTP and two nodes doing reporting in my six node cluster. Reporting is not mission-critical, however. One of my four OLTP nodes goes down now instead of running just on three nodes, I'm going to redirect some of my traffic to one of those other reporting nodes. I'm going to be running OLTP with four nodes (three surviving nodes and one re-purposed reporting node). I am consciously backing off my reporting so instead of my OLTP running degraded, its going to run at full steam and I'm going to let my reporting system degrade. The beauty of Oracle 10g and a RAC-enabled grid is that the choice is mine. I can manage SLAs based upon my business needs.”

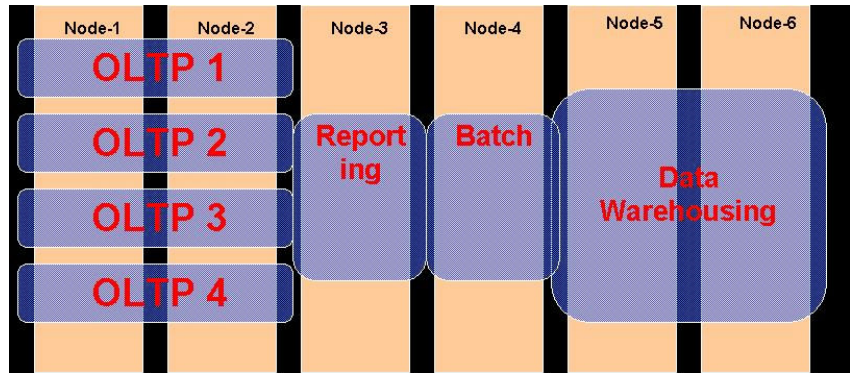


Figure 3. Talk America’s typical service allocation among the six Oracle RAC nodes implemented on a 96 CPU two HP Superdome environment. Services can be relocated and resource allocation expanded or contracted as daily, monthly, year-end processing requires.

Talk America has achieved greater hardware utilization, better SLA achievement, and significant cost savings by virtue of the full technology stack supplied by Oracle Database 10g and Real Application Clusters (no need to acquire \$500,000 of third party clustering software). Again – leveraging all the hardware all the time.

CONCLUSION

Oracle 10g Real Application Clusters offers enterprises:

- Radical economic benefits when migrating from an aging SMP environment to a modern enterprise grid of low cost servers enabled with Oracle RAC.
- An agile response to changes in business requirements through the addition or subtraction or re-purposing of server resources as transaction volumes, processing schedules, number of user sessions or application mix changes.
- Increased hardware utilization which enables your IT function to fund more strategic initiatives by re-directing budget away from under-utilized legacy hardware/software infrastructure to strategic development initiatives. This results in enhanced competitiveness and customer responsiveness.



Oracle Real Application Clusters 10g – the Foundation for Enterprise Grid Computing
September 2005
Author: R. Hietter

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
oracle.com

Copyright © 2005, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle, JD Edwards, and PeopleSoft are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.