

# Using VMware ESX with Oracle Database and EMC CLARiiON Storage Systems

*Applied Technology*

---

## **Abstract**

This white paper provides general guidelines regarding how the Oracle Database can be deployed with VMware® Infrastructure 3 and EMC® CLARiiON® storage systems. Three features of VMware Infrastructure — Virtual SMP, VMotion, and DRS — are discussed in the context of Oracle Database performance and server manageability.

August 2008

---

---

Copyright © 2008 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on [EMC.com](http://EMC.com)

All other trademarks used herein are the property of their respective owners.

Part Number H5503

---

## Table of Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>4</b>
Audience .....	5
Terminology .....	5
<b>Oracle Database testing with VMware .....</b>	<b>6</b>
Test environment .....	6
Test methodology .....	6
Enhancing VM performance with multiprocessing .....	7
Managing server resources with VMotion.....	8
Automating resource management with DRS.....	10
<b>Conclusion .....</b>	<b>12</b>
<b>References .....</b>	<b>12</b>

---

## Executive summary

Virtualization technology is a hot topic in the IT world because of its potential for cost reduction and enhanced server manageability. The technology fits nicely with EMC® CLARiiON® midrange storage systems benefiting from its flexible shared storage options and industry-proven performance and availability. However, with applications like the Oracle Database, deployment in a virtual environment can be a challenging issue because of specific server resource and performance requirements.

In addition to existing EMC white papers detailing CLARiiON integration with VMware® ESX (refer to the “References” section), this paper describes how Oracle Database can be used with VMware ESX with respect to VMware Virtual Machine Multiprocessing, VMware VMotion and VMware Distributed Resource Scheduling (DRS). It demonstrates how these VMware features can be leveraged to manage resources and offer high availability for Oracle Database deployments.

## Introduction

The Oracle Database has traditionally been deployed in non-virtual environments because of its performance-critical nature. With VMware Infrastructure 3, the CPU overhead for the virtual environment can be minimal (typically less than 10 percent), and the Oracle Database can run at near-native performance. By leveraging CLARiiON’s shared storage, the Oracle Database can also move from one physical machine to another, which not only increases flexibility, but also increases availability of Oracle Database systems. The following outlines each component of the VMware-Oracle-CLARiiON deployment infrastructure, along with VMware’s main functional features.

### VMware ESX

VMware ESX is a “bare metal” hypervisor that partitions physical servers into multiple virtual machines. Each virtual machine represents a complete system, with processors, memory, networking, storage and BIOS. This allows multiple virtual machines to run side by side on the same server and share physical resources.

### VMware® VirtualCenter

VMware VirtualCenter is management software that delivers centralized management, monitoring, provisioning, operational automation, resource optimization, and high availability to virtualized IT environments. Multiple VMware ESX hosts and their virtual machines can be managed through a single interface that offers guided execution of VMware Infrastructure features such as VMotion, DRS, VMware High Availability, and VMware Consolidated Backup.

### VMware VMotion

VMotion provides the ability to migrate a running virtual machine from one physical VMware ESX host server to another, without any interruption of service. This functionality allows users to eliminate planned downtime for hardware maintenance, and also allows dynamic load balancing among multiple servers.

### VMware DRS

VMware DRS monitors resource allocation across a cluster of VMware ESX host servers and dynamically allocates IT resources to the highest priority applications based on policies that prioritize resources for virtual machines. The recommendations generated by policies can be applied automatically or manually.

### Other VMware features

VMware High Availability (HA) detects VMware ESX host server failures and automatically restarts virtual machines on an alternate working ESX host server. VMware Consolidated Backup creates a virtual disk snapshot to perform a full image backup of virtual machines. VMware Update Manager automates patch management and the update process for VMware ESX host servers, as well as Windows and Linux virtual machines.

---

## Oracle Database

Oracle Database is a performance and security proven enterprise software that runs on Windows, Linux, and UNIX. If deployed with Oracle Real Application Clusters, a single database can run on a cluster of machines, laying the foundation for grid computing. By deploying Oracle Database with VMware Infrastructure, server utilization can be improved and users can also dynamically move or scale server resources on multiple machines.

## CLARiiON storage system

CLARiiON CX3 series storage systems provide multiple tiers of storage accessible in Fibre Channel (FC) or iSCSI. Its support for Serial Attached SCSI (SAS) and SATA disk drives makes the storage configuration more flexible, with an opportunity to leverage lower costs with SATA drives. CLARiiON supports VMware Virtual Machine File System (VMFS) on a shared storage, which allows VMware Infrastructure and features like VMotion to run seamlessly.

## Audience

This white paper is intended for EMC customers, partners, or employees who want more information on how VMware VMotion can be used with Oracle Database and the CLARiiON storage system.

## Terminology

**Bandwidth:** The average amount of read/write data in megabytes that is passed through the storage system per second.

**Logical unit number (LUN):** A unique identifier that is used to distinguish among logical storage objects in a storage system.

**Online transaction processing (OLTP):** A type of processing in which the computer responds immediately to user requests. Each request is considered to be a transaction.

**Oracle Automated Stress Testing (OAST):** An automated test suite designed to build OLTP type workloads for systems using the Oracle Database. It creates tables, performs stress test runs, and outputs transaction-related performance data.

**Oracle Automatic Storage Management (ASM):** Storage management tool that provides cluster file system and volume management capabilities integrated into the Oracle Database. It performs automatic redistribution of storage capacity after the incremental addition or removal of storage disks.

**Redundant Array of Independent Disks (RAID):** RAID technology groups separate disks into one logical unit number (LUN). Data is distributed across these disks to improve reliability and performance. The CX3 series supports RAID level 0, 1, 1/0, 3, 5, and 6.

**Response time:** The average time in milliseconds that it takes for one request to pass through the storage, including any waiting time.

**Serial Advanced SCSI (SAS):** A high-performance I/O technology that uses a serial point-to-point interface instead of a parallel bus interface, as with parallel SCSI.

**Serial Advanced Technology Attachment (SATA):** An enhancement of ATA that uses thinner cables and provides better performance over parallel drive technology.

**Storage area network (SAN):** A Fibre Channel or iSCSI storage network used to connect servers and shared storage systems.

**Throughput:** The average number of read/write requests in I/Os that are passed through storage per second.

**Transactions per minute (TPM):** The measured number of transactions completed by an OLTP database over a 1-minute interval.

---

## Oracle Database testing with VMware

The following three VMware VI3 features have been tested with Oracle Database in this paper:

- Multiprocessing capability
- VMotion
- Distributed Resource Scheduling (DRS)

They all relate to the usage of server resources on virtual machines. From these tests, users will be able to better determine how the Oracle Database can be used under different deployment settings. VMware ESX Server 3.5.0 and Oracle Database 11g were used for all testing scenarios supporting this paper

### ***Test environment***

Server machines: 2 x Dell 2950

Hypervisor	VMware ESX 3.5.0
Number of CPUs	2 x Intel Quad-Core Xeon 1.86 GHz
Memory	16 GB
Network adapter	Intel PRO/1000PT Gigabit NIC
FC adapter	QLogic 4 GB FC HBA

Storage systems: EMC CLARiiON CX3-20f

Processors	2
Memory size	2 GB per SP
Number of disks	30 FC 73 GB @ 15k rpm
Base software	03.26.020.5.011

Virtual machines: VM1

OS	Oracle Enterprise Linux (OEL) 5 (2.6.18-8.el5)
Database	Oracle Database 11.1.0

### ***Test methodology***

In the tests, Oracle Database performance was measured using an Oracle Automated Stress Testing (OAST) tool that creates an online transaction processing (OLTP) database and continuously makes multiple-user online transactions for a set period of time. The OAST run that outputs TPM numbers indicate the performance level of the database. The actual TPM numbers will vary in real Oracle Database deployments. Oracle Database 11g was installed on a Linux virtual machine. Oracle Automatic Storage Management (ASM) was also used to manage disk groups for the database.

For the storage configuration, a Fibre Channel (FC) connection was used for the CLARiiON SAN. RAID 5 (4+1) was used as the RAID group for all LUNs. A metaLUN consisted of two 100 GB LUNs that were used as the Boot LUN that stored the OS for the virtual machines. For the OAST database, data files were spread out on four 100 GB LUNs using ASM, and log files were placed in two 50 GB LUNs. The storage group configuration captured from CLARiiON Navisphere® is shown in Figure 1. “vmware1” and “vmware2” are the two VMware ESX host servers used in this testing.




**Figure 1. Storage group configuration for VMware ESX servers**

## Enhancing VM performance with multiprocessing

Virtual machine (VM) performance depends largely on the number of CPUs and the size of memory allocated to the machine. For an enterprise-level application like the Oracle Database, multiprocessing will ensure a high performance level in a virtualized environment. It is made possible by VMware Virtual SMP, which allows a single VM to scale up to four physical processors and work in parallel.

This test analyzes how scaling of virtual CPU and memory affects VM performance and ultimately, the Oracle Database performance. Two datapoints were measured: two virtual CPUs and 2 GB virtual memory for simulation of a small-scale database server, and four virtual CPUs and 4 GB virtual memory for a medium to large scale database server. Figure 2 shows how resource change is done through the VMware VirtualCenter Edit Settings menu. A more advanced resource allocation with resource reservation and resource limits can be set in the Resources tab. The change can only be made when the VM is powered down.



**Figure 2. Virtual CPU and memory resource change in Virtual Machine Properties**

For the OAST run, 1,000 warehouses were created and 100 user workloads were stressed for 60 minutes. Given the same amount of workload, the VM configuration with four vCPUs and 4 GB memory yielded higher transactions per minute (TPM) than the VM configuration with two vCPUs and 2 GB memory. The TPM and response time results are scaled to be between 0 and 100, with 100 being the maximum value of the result. The graphs are shown in Figure 3, where the four-vCPU configuration performed about 55 percent better than the two-vCPU configuration, with about 35 percent lower response time.





Figure 3. TPM and response time comparison between two vCPU VM and four vCPU VM

### Managing server resources with VMotion


VMware VMotion works by sharing VM files between two servers and transferring only the active memory state of the VM from one server to another. Shared VMFS volume or network-attached storage (NAS) is required for VMotion to work. The CPU must also be compatible between the two servers. It is highly recommended to use a private GigE network connection for VMotion. With VMotion, network identity and connections are preserved, so there is no downtime or disruption of service.

For the test, two VMware ESX hosts were set up with a private GigE interconnect for VMotion. OAST was used for continuous database transactions, with the workload request coming in from a separate client node. A four-vCPU, 4 GB VM was used, and resource monitoring was done using the VirtualCenter Performance Monitoring chart. Figure 4 shows how VMotion is started by selecting **Migrate** from the VirtualCenter menu.




**Figure 4. Guided VMotion in VirtualCenter**

While the database in VM1 executed transactions, VMotion was initiated. The migration completed in less than two minutes, and server resource utilization levels like virtual CPU and memory went about 10 percent higher during migration. Because of this resource demand by VMotion, the OAST TPM was momentarily reduced and the response time was increased during migration. Figure 5 reflects the TPM drop of about 22 percent, and Figure 6 shows the response time rise of about 33 percent during VMotion. It is important to note that the Oracle Database did not experience any timeouts or freeze in transactions.



**Figure 5. TPM during VMotion**




**Figure 6. Response time during VMotion**

By using VMotion, the Oracle Database server does not have to be restricted to one physical machine but can migrate to different machines without interruption of service. This means that hardware maintenance can be carried out without scheduled downtime, and virtual machines can move away from underperforming servers and run on new servers without re-installation. Also during low transaction times, the database server VM can be moved to run with other applications on one server in order to save power.

**Automating resource management with DRS**


When multiple virtual machines are running on one server, it can be challenging to allocate the right amount of resources for each virtual machine. Furthermore, when multiple servers are available, moving VMs for better resource utilization can be a cumbersome task that is hard to predict. VMware DRS solves this problem by automatically and dynamically allocating VMs across a collection of servers aggregated into one logical resource pool. This ensures that an Oracle Database VM can be available at its optimal performance rate even when sharing a server with other VMs.

This test demonstrates automated resource management with DRS by creating a server cluster of two hosts that shares a resource pool at the root level. Two VMs were placed inside the cluster and within the same server. Another server had no VM. To show a more detailed DRS procedure, the Manual automation level was chosen in the cluster’s Edit Settings menu, as shown in Figure 7.




**Figure 7. VMware DRS automation level settings in a VMware ESX cluster**

The two VMs created were four vCPU, 4 GB VMs with Oracle Database installed. OAST was started for both VMs with separate clients and separate storage systems to simulate resource contention at the host level. As soon as OASTs were started for the two VMs, using almost 100 percent of their assigned CPU resources, the DRS Recommendation message was shown as in Figure 8. The message recommended that the user migrate one of the VMs to another server. The user can apply the recommendation, which triggers the VMotion migration process. When DRS is fully automated, this recommended action is executed automatically.



**Figure 8. DRS recommendations in a VMware ESX cluster**

When one VM is moved to another server, DRS allocates extra cycles from idle CPUs for VMware ESX overhead tasks and networking, resulting in a slight improvement in the remaining VM's OAST TPM numbers. The result is shown in Figure 9, where the TPM rate for the remaining VM after the migration is about 15 percent higher than when both VMs were running together in the same server.



**DRS recommendation apply**

**Figure 9. TPM performance after VM migration**

VMware DRS can be used when multiple VMs are running on one server, but the user wants to prioritize a task on a certain VM, by ensuring that the VM has enough server resources to run at its full capability. Also, this feature can be used to balance out the resource pool of multiple servers, where some VMs can be automatically moved off to any underutilized servers. When VMs are sharing the same storage system, Oracle users can also use CLARiiON's Quality of Service (QoS) Manager or the NQM, which allows users to control storage system resources for VMs. For example, a user can limit the storage resource for VMs to ensure more storage resources for a performance-critical VM, like the Oracle Database. This can be done without moving VMs to different servers.

---

## Conclusion

VMware virtualization software can greatly increase server utilization by consolidating multiple virtual machines. However, when leveraging the flexibility offered by virtualization with the Oracle Database, users need to know the impact on database performance resulting from sharing of server resources. With the help of VMware Infrastructure features like Virtual SMP, VMotion, and DRS, Oracle Database users are better able to control Oracle Database performance and manage VMs across multiple servers. This ability results in more dynamic and flexible Oracle Database deployments. CLARiiON's VMFS support and QoS Manager can also help deployments with VMware Infrastructure be more fluid and highly available.

## References

On VMware.com:

- VMware and Oracle database solutions page  
<http://www.vmware.com/partners/alliances/technology/oracle.html>

- VMware Server Virtualization Products page  
[http://www.vmware.com/products/data\\_center.html](http://www.vmware.com/products/data_center.html)

- VMware Virtual SMP product page  
<http://www.vmware.com/products/vi/esx/vsmp.html>

- VMware VMotion product page  
<http://www.vmware.com/products/vi/vc/vmotion.html>

- VMware DRS product page  
<http://www.vmware.com/products/vi/vc/drs.html>

On EMC.com:

- EMC CLARiiON CX3 UltraScale Series datasheet  
<http://www.emc.com/collateral/hardware/data-sheet/c1148-clariion-cx3-ds.pdf>
- *EMC CLARiiON Integration with VMware ESX Server* white paper  
<http://www.emc.com/collateral/hardware/white-papers/h1416-emc-clariion-intgtn-vmware-wp.pdf>
- *VMware ESX Server Using EMC CLARiiON Storage Systems Solutions Guide*  
<http://www.emc.com/collateral/software/solution-overview/h2197-vmware-esx-clariion-stor-syst-ldv.pdf>
- *Using VMware VMotion with Oracle Database and EMC CLARiiON Storage Systems* white paper  
<http://www.emc.com/collateral/hardware/white-papers/h4205-using-vmware-vmotion-oracledb-emc-clariion-stor-sys-wp.pdf>
- *Maintaining End-to-End Service Levels for VMware Virtual Machines Using VMware DRS and EMC Navisphere QoS* white paper  
<http://www.emc.com/collateral/hardware/white-papers/h4089-maint-end-end-srvc-vmware-virt-mach-vmware-drs-emc-navisphere-qos-wp.pdf>