

Oracle RAC on VMware with DRS and HA

A DBA's Perspective

By Kathy Gibbs, Product Manager and Janis Griffin, Senior DBA

Confio Software
4772 Walnut Street, Suite 100
Boulder, CO 80301

www.confio.com

Introduction

Though virtualization has become a key strategy to rein in the costs and complexity of data environments, and increasingly even for business-critical applications such as Oracle RAC, Oracle has been slow to fully support it. However, with at least 50% of database servers virtualized as of 2011ⁱ, and that number only increasing, Oracle has been forced to reconsider, and has finally issued policies explicitly supportive of running Oracle products, including Oracle RAC, on VMware—the hypervisor in use for server virtualization at more than 67% of shops that virtualizeⁱⁱ.

For the Oracle DBA, this opens up significant new possibilities. In fact, VMware advocates that the next-generation DBA will be what it terms the vRAC DBA—the custodian of the “dataverse”—fully skilled in optimizing performance of Oracle RAC in a virtualized environment—and in exceptionally high demandⁱⁱⁱ.

To realize the full potential of virtualization, the Oracle DBA must become familiar with the critical issues and best practices for running Oracle RAC on VMware. In this paper we discuss the configuration changes needed for installing Oracle RAC on VMware, considerations for using the DRS Host Affinity Rules and High Availability (HA) configurations, and the performance parameters that must be watched.

Yes, Oracle will really support Oracle RAC on VMware

Beginning with an official policy on Oracle products running on VMware vSphere^{iv}, issued in November, 2010, Oracle officially extended its support to Oracle RAC a short time later^v. Despite this, a key clause in Oracle’s official policy creates concern for some:

“Oracle will only provide support for issues that either are known to occur on the native OS, or can be demonstrated not to be as a result of running on VMware.”

This is simply a standard Oracle practice for isolating a problem. In reality, once you install a hypervisor on a physical host and create a guest, the operating system is fully contained in the guest, and there is no modification to the operating system itself. It is exactly as if the guest is running on a physical host, and there is no implication of the virtual machine (VM) in resolving any database issues that may occur.

The worst case would be that Oracle might ask you to recreate the problem on a physical host. Again, this isn’t a new approach for technical support. We constantly monitor numerous support boards and communities—and issues involving Oracle RAC on VMware are not showing up, and we have not seen reference made to organizations actually being asked to reproduce problems in this way.

Yes, you can control licensing costs for Oracle RAC on VMware

Unlike most other major vendors, Oracle doesn’t currently support virtual CPU-based licensing and instead licenses its products by physical core or socket. This approach, deplored by many industry experts as being excessively anti-virtualization or punitive to virtualization technologies other than Oracle’s, means that careful consideration must be made for how to effectively install Oracle RAC to

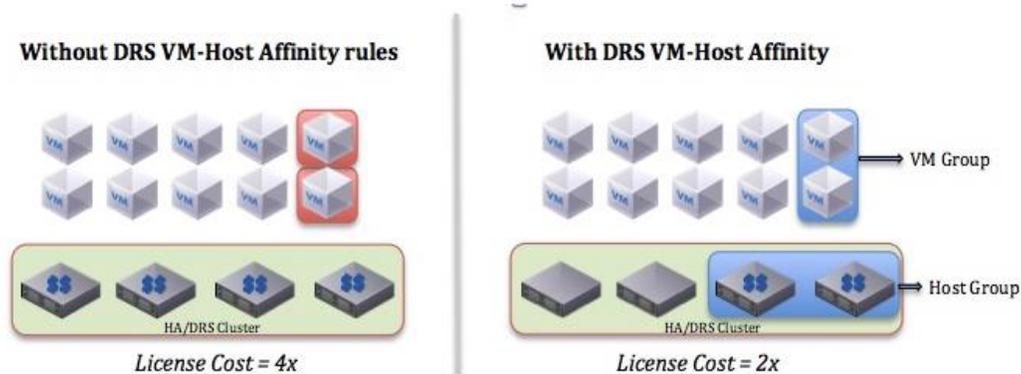
minimize licensing fees. The good news is that you can use the very capabilities of VMware virtualization to keep the costs down.

NOTE: As an Oracle DBA, you should become familiar with the licensing provisions for your specific configuration, and discuss any issues or concerns with your Oracle representative.

Consider this typical licensing scenario. Often, when an organization decides to virtualize, it acquires one very large server and installs everything on it. In the example of a box with 128 cores, Oracle licensing would require you to license all 128 cores, regardless of how many cores the database was running on.

A workaround to this big box approach would be to restrict movement of the database application across hosts in the virtualized environment. VMware's DRS Host Affinity rules (DRS) allow you to restrict movement of VMs between a subset of hosts within a cluster, so that not all hardware in the cluster is "available" to the Oracle database. DRS is a clustering technology, rather than a mechanism for soft or hard partitioning of the servers.

For example, consider a configuration of Oracle RAC on VMware set up to provide failover for disaster preparedness (see figure below). Without using DRS, the two VMs could run on all four hosts. With DRS, you could pay to license just two hosts and restrict the VMs to only run on one of the licensed servers. In this case, your licensing costs could be as much as 50 percent less.



Using DRS host affinity rules to minimize licensing costs

Deploying Oracle RAC on VMware with DRS and High Availability (HA) isn't hard at all

From a DBA point of view, deploying Oracle RAC in a VMware environment brings significant advantages. We can use DRS to load balance and mitigate licensing costs, while VMware High Availability (HA) provides availability without the costs and complexity of failover solutions dependent on an operating system or specific applications. VMware HA, a feature of vSphere, leverages multiple VMware hosts configured as a cluster to provide rapid recovery from unplanned outages as well as cost-effective high availability for applications running on VMs.

We tested this Oracle RAC VMware with DRS and HA configuration at our labs. For this purpose, we had the configuration shown in the table below.

Host	VM	Instance	Service
lindert	OracleRAC1	ORCL1	ProdWK
lado	OracleRAC2	ORCL2	ProdWK
lowden	OracleRAC3	ORCL3	Report

Test Oracle RAC on VMware configuration

While we found that installing Oracle RAC into the VMware environment was almost the same as into a physical environment, there were a few things we needed to do differently.

To perform VMware HA functions (such as vMotion, a critical tool), there are a few additional steps you must take at the beginning of the install process (for detailed, step-by-step instructions on set up, see <http://tiny.cc/bngjgw>), including:

- Disable simultaneous write protection provided by VMFS using the multi-writer flag.
- Eager-zeroed thick disk must have all space allocated and wiped clean of any previous contents on the physical media at creation time. Note that such disks may take longer to configure during creation.
- Configure SCSI adapter to Paravirtual.
- In the case of third-party cluster-aware applications in which the applications ensure that writes originating from two or more different VMs will not cause data loss, you must manually enable the multi-writer flag using these instructions: <http://tiny.cc/i5ahgw>

Overall, the installation and configuration of Oracle RAC into the VMware environment was nearly identical as into a physical environment.

Does Oracle RAC on VMware really perform as well as on a physical host?

From a DBA point of view, we wanted to know how Oracle RAC on VMware actually performed. To do this, we replicated elements of a published VMware test of performance^{vi}, and achieved nearly identical results as VMware—all at least as good or better than on a physical host.

For this test, we monitored 1,000 OLTP users over 24 hours. As in the VMware test, we saw the workload stabilize in the last 12-15 hours of the run, and all the performance indicators leveled out.

The workload was evenly distributed across four Oracle RAC nodes, and each node showed similarly high CPU utilization (about 90%). Looking at these types of numbers can be a startlingly clear reminder of the power of virtualization—unlike when running on a physical host, you want to see really high utilization!

A few other results that we considered:

- Redo LUN latency was less than 10 milliseconds, after the workload leveled out. The 10-millisecond threshold is a guideline.
- IOPS of approximately 0800-9000.
- The approximate interconnect traffic was 242MB/sec. The network counters from esxtop (VMware Utility) reflected similar numbers
- The top five wait events for the 24-hour period demonstrates that latency due to the Oracle RAC interconnect was less than 5%.
- Overall throughput of SwingBench transactions did not degrade.

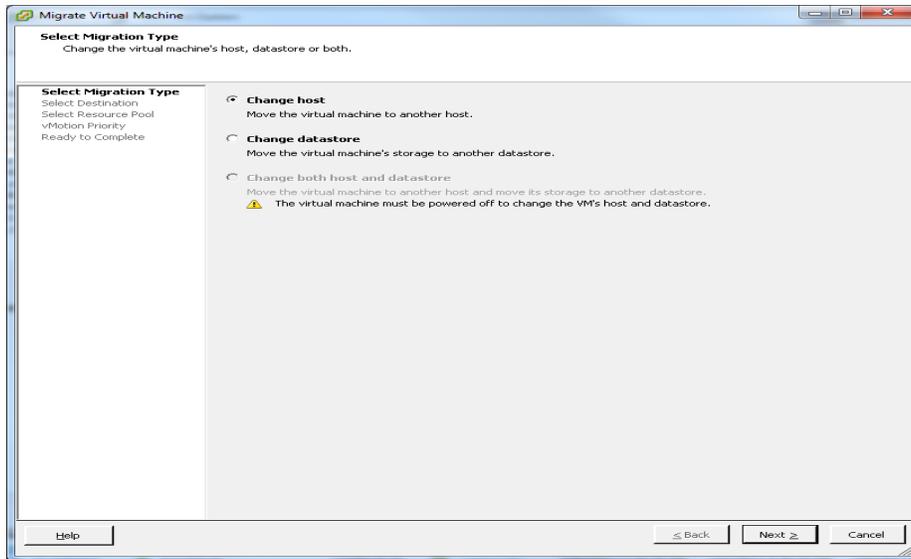
Our performance test results clearly indicated that Oracle RAC on VMware performed as well or better than on a physical host, and greatly enhanced utilization and flexibility.

Keeping Oracle RAC on VMware up, even with unplanned down time, using vMotion

For business-critical applications like Oracle RAC, downtime and disruption are unacceptable. One of the more powerful tools available to the DBA in the VMware environment is vMotion, a tool for live migration of a virtual machine (VM) from one host to another, with no loss of connectivity. With vMotion, you can quickly migrate a VM from, say, a failing physical host, without any loss of uptime or performance hit.

In our test, we assumed a memory failure on our physical host lindert. Without the ability to perform a live migration, we would have had to complete a series of time-consuming and business-disruptive steps, including: scheduling down time for the physical host, bringing the host down; completing the work to sway out the physical memory; bringing the host back up. This process could often take hours or days to fully complete, and the downtime could be very costly to the business.

Running on VMware, however, we simply used vMotion to move the VM from host lindert to host lado (see figures below).



Using vMotion to Select Migration Type

There was no business disruption. And the process took minutes, instead of hours or days.



Using vMotion to move a Virtual Machine from one host to another

Quickly expand a cluster and deploy new database instance for business needs

Another powerful ability for DBAs deploying Oracle RAC on VMware is the rapid set up or take down of new instances of the database when and where you need them, with minimal time and disruption. This type of task could take days or weeks in a physical environment; in a virtualized environment, the DBA can accomplish this in minutes.

For example, we wanted to see how easy it would be accommodate a new and different group of database users. Our first and primary group of users needs access to the database for OLTP. A second, new group needs access for just quarterly reporting. Without virtualization capabilities, accommodating a new and different group of users would typically be a time-consuming process that might take days or weeks. You would have to acquire and build a physical host, including the install, the patches, the configuration, provisioning, testing and deploying. Once up and running, it would likely not have the exact same configuration as the instance the first group is using, as new builds typically have different patches.

With the capabilities of running in a VMware environment, however, we were able to add an instance onto the cluster in just hours, using a copy of the guest we made after first setting it up. At a high level, the process was quite straightforward:

1. Copy the host.
2. Copy the guest.
3. Copy the CRS home (tar up existing home; untar it on new instance).
4. Run the clone CRS.
5. Run the addNode script [see <http://tiny.cc/wekjgw> for more details].
6. Clone the ASM and database homes (by script or DBCA). This can also be done easily if you have OEM grid set up and it is running on all the nodes.

Note that while there is documentation for cloning or restoring copy from VMware, this isn't tested officially and is not a preferred method.

The ability to respond quickly to business needs in this way is tremendously valuable and well worth exploring.

What should a DBA do differently to monitor performance of Oracle RAC on VMware?

The ability to run Oracle RAC on VMware clearly brings significant benefits. To do so most effectively, it's important that the Oracle DBA work closely with the VMware, system and storage administrators to ensure optimal database performance. When issues do arise, it's especially important that the Oracle DBA has appropriate tools available to clearly see:

- What changed recently that affected end users
- What layer (VM or DB) is causing the problem
- Who and how the problem should be addressed

Armed with the right tools, such as SolarWinds Database Performance Analyzer (DPA) and DPA VM Option, or given access to VMware’s vSphere tool, the Oracle DBA can closely monitor and quickly identify issues, and be able to quickly identify who and how issues can be addressed.

Getting a single view across the database, virtual and physical layers

As a DBA working with Oracle RAC on VMware, you will need to understand what is happening in the virtual and physical environments and how that impacts performance of the database. DPA VM Option is one tool that gives you a powerful, visual way to correlate the database and VMware activity so that you can see through the entire stack and understand where issues are happening and how to solve them.

Below, you can at the top the SQL hash IDs currently running; and you can see what is happening in the guest, the physical host and the storage levels. In this example, you can see latency happening at the physical host level, but not at the guest.

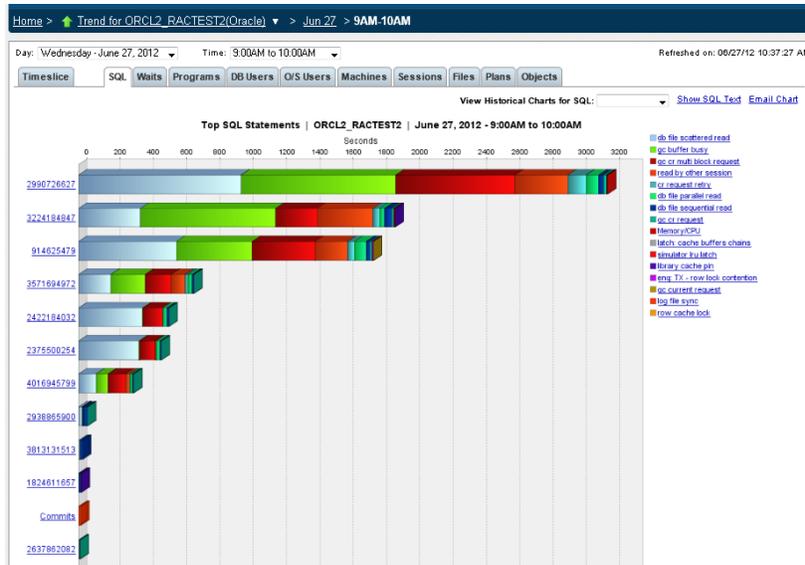


Using DPA VM Option to see the whole stack

Monitoring Oracle RAC wait events

Just as with Oracle RAC on a physical server, the ability to see wait events can help you quickly pinpoint a problem when users start complaining about slowness.

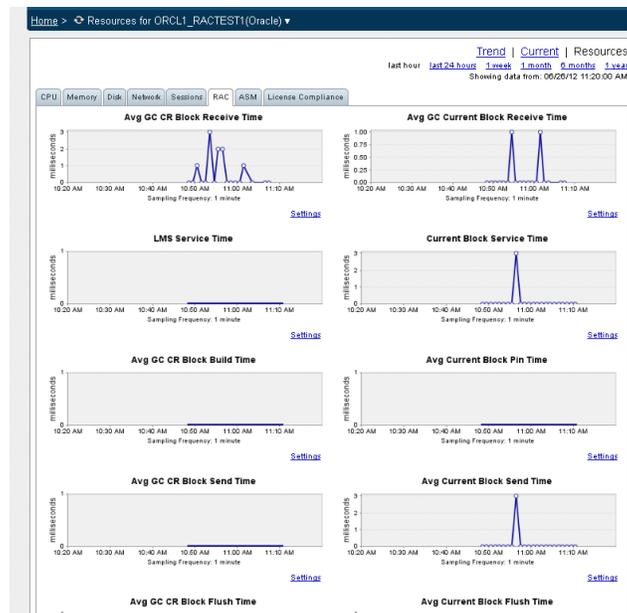
In the figure below, DPA lets you review all SQL statements and see that there are several with a long “gc buffer busy” time, letting you know quickly which ones are slow due to wait events. With this information, you can isolate problems that are due to database performance instead of VM, server or storage issues.



Using DPA to monitor wait events

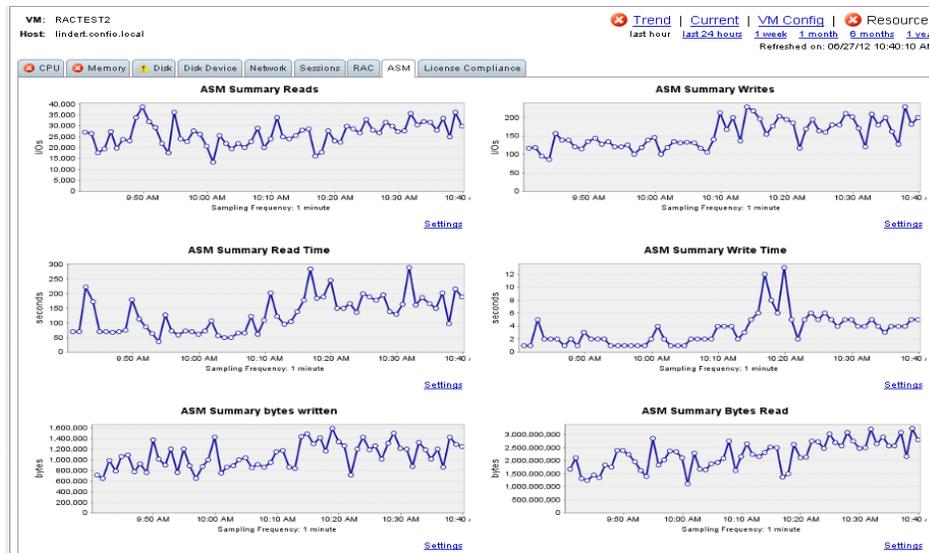
Monitoring Oracle RAC resource information

You will also want to monitor Oracle RAC information to monitor how the Interconnect and cluster are performing. SolarWinds Database Performance Analyzer (DPA) VM Option gives you this specialized view so that you can see any critical events that might be happening.



Using DPA VM Option to monitor Oracle RAC interconnect and cluster performance

DPA VM Option also lets you monitor latency with ASM (see figure below), so that you can monitor latency in reads and writes and ensure that issues are addressed quickly.



Using DPA VM Option to monitor ASM latency

Conclusion: realize the benefits of Oracle RAC on VMware with no risk

Oracle’s evolving support and licensing policies for Oracle RAC on VMware means you can realize the benefits of virtualization-- reducing IT costs while increasing the efficiency, utilization, and flexibility of your existing assets—without risk. And most available data, including our own performance testing, indicates that Oracle RAC runs as well or better on a virtual machine than on a physical host.

As an Oracle DBA, your knowledge of the issues and best practices for running Oracle RAC on VMware, including use of DRS and HA, is critical. From understanding how and when to use live migration tools, to how to provide adequate failover protection while minimizing licensing costs, you can make a significant difference in how effective your organization is in taking advantage of all the benefits of virtualization, and developing the skills of the next-generation DBA.

About Configo Software

Configo Software, now a part of the SolarWinds family, builds award-winning database performance analysis tools for DBAs and developers. SolarWinds Database Performance Analyzer (formerly Configo Ignite) improves the productivity and efficiency of IT organizations. By resolving problems faster, speeding development cycles, and squeezing more performance out of expensive database systems, Database Performance Analyzer makes DBA and development teams more productive and valuable to the organization. Customers worldwide use our products to improve database performance on Oracle, SQL Server, Sybase and DB2 on physical and virtual machines.

For more information, please visit: <http://www.confio.com/performance/oracle/rac/>

Glossary of Common Oracle RAC and VMware terms

- ASM—A volume manager and a file system for Oracle database files that supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations.
- Clusterware—Oracle’s software that allows clustering of independent servers so that they cooperate as a single system.
- DBCA—An Oracle utility used for creating, configuring and removing Oracle databases.
- Distributed Resource Scheduler (DRS) – an Oracle tool used to set up, balance and isolate workload automatically across hosts
- ESX and ESXi – VMware’s hypervisor and foundation for VMware products
- Guest—The operating system that runs in a virtual machine. There can be multiple guests per physical machine, but only one per virtual machine.
- Interconnect – VMware’s high-speed, low-latency communication link between nodes in a cluster
- Oracle RAC—Oracle Real Application Clusters (RAC) supports the transparent deployment of a single Oracle database across a cluster of servers, providing fault tolerance from hardware failures or planned outages.
- Physical host – underlying hardware where VMware ESX is installed
- vCenter Server – VMware’s centralized management tool for the VM
- vMotion –VMware’s tool for live migration of a VM from one host to another with no loss of connectivity
- VMware High Availability (HA)—VMware’s solution for providing high availability to virtual machines
- vSphere Client – VMware’s admin and monitoring tool for the VM
- Virtual Machine (VM) –a container inside the host that looks like a physical machine

Additional Resources

- [IOUG Virtualization Group](#)
- From VMware: [VMware’s Virtualizing Oracle Kit](#)
- From Oracle: [Oracle 11g Release 1 RAC On Linux Using VMware Server](#)

ⁱ 2010 IOUG Enterprise Platform Decision Survey: Toward a Smarter Information Foundation.

ⁱⁱ Vanson Bourne, Survey of IT Decision-makers at 578 businesses with at least 1,000 employees in the United States and Europe, 2012

ⁱⁱⁱ <http://www.dbta.com/Articles/Columns/My-View/Last-New-Horizon-of-Computing---the-Dataverse-82638.aspx>

^{iv} MyOracleSupport Document ID #249212.1; in November 2010, Oracle expanded this support policy to include Oracle RAC on vSphere.

^v MyOracleSupport Document ID #249212.1; in November 2010, Oracle expanded this support policy to include (RAC) on vSphere.

^{vi} VMware Workload Characterization Study <http://tiny.cc/pffkgw>