

An Oracle White Paper
November 2010

Oracle Real Application Clusters One Node: Better Virtualization for Databases

Executive Summary.....	1
Oracle Real Application Clusters One Node	1
Better Server Consolidation	3
Enhanced Protection from Failures.....	4
Greater Flexibility and Workload Management	6
Better Online Maintenance	7
Other Unique Benefits of Oracle RAC One Node.....	8
Complements Server Virtualization.....	9
Conclusion	9

Executive Summary

Oracle Real Application Clusters One Node is a new option to Oracle Database 11g Release 2 Enterprise Edition. It improves upon many of the benefits of server virtualization and extends them to databases running in physical server environments.

Oracle RAC One Node enables:

- Better server consolidation
- Enhanced protection from failures
- Greater flexibility and workload management
- Better online maintenance

In addition it allows customers to virtualize database storage, standardize their database environment, and, should the need arise, upgrade to a full multi-node Oracle RAC database without downtime or disruption. Further, it is fully compatible with and complementary to Oracle Virtual Machine (VM), allowing customers to build environments that leverage the strengths of both Oracle RAC One Node and server virtualization.

Oracle Real Application Clusters One Node

With Oracle Database 11g Release 2, Oracle introduces Oracle Real Application Clusters (RAC) One Node, a new feature that combines better-than-VM virtualization benefits with physical server performance.

Virtualization has become a hot topic in the IT industry (See Box 1). With Oracle RAC One Node, Oracle now provides a more sophisticated form of virtualization developed for database operations that improves upon the benefits of virtualization, and provides additional unique value. It enables better server consolidation, enhanced protection from failures, greater flexibility and workload management, and better online maintenance than VMs.

Moreover, RAC One Node runs on physical servers and does not suffer the performance handicap of virtual servers.¹ A database may incur high overhead running in many server virtualization environments due to their propensity to perform I/O and memory intensive operations—for such environments, running on physical servers will provide substantial performance improvements.

BOX 1: Benefits of Server Virtualization

Why is virtualization popular? By abstracting the physical resources from the software running on them, virtualization promises a more flexible and efficient IT environment. It provides a variety of easily attainable benefits that have a clearly measurable and understandable payback. It's easy to prove that the investment will have a positive ROI.

The most obvious benefits of virtualization are those associated with server consolidation, but more and more customers are leveraging the benefits of advanced features such as live migration and high availability. Oracle recognizes these benefits and provides its own server virtualization product, Oracle VM. Following is a summary of the key benefits driving interest in virtualization:

- **Server Consolidation:** Consolidating under-utilized servers into a single physical server can reduce the cost of servers, floor space, power, and cooling.
- **Protection from Failures:** Virtualization can be combined with monitoring capabilities that monitor the health of the virtual resource or the actual application and act to restart or re-locate the workload in the event of a failure
- **Flexibility and workload management:** Virtualization makes it possible to migrate workloads between physical servers within a cluster, balancing load across all the server resources in the pool
- **Online Maintenance:** Migrating work off a physical resource allows administrators to perform maintenance on that physical resource without affecting users of the application.

Virtualization has become synonymous with Server Virtualization, but many different types of virtualization exist. Server virtualization, the simplest form of virtualization, can provide many of the above benefits, with varying degrees of utility.

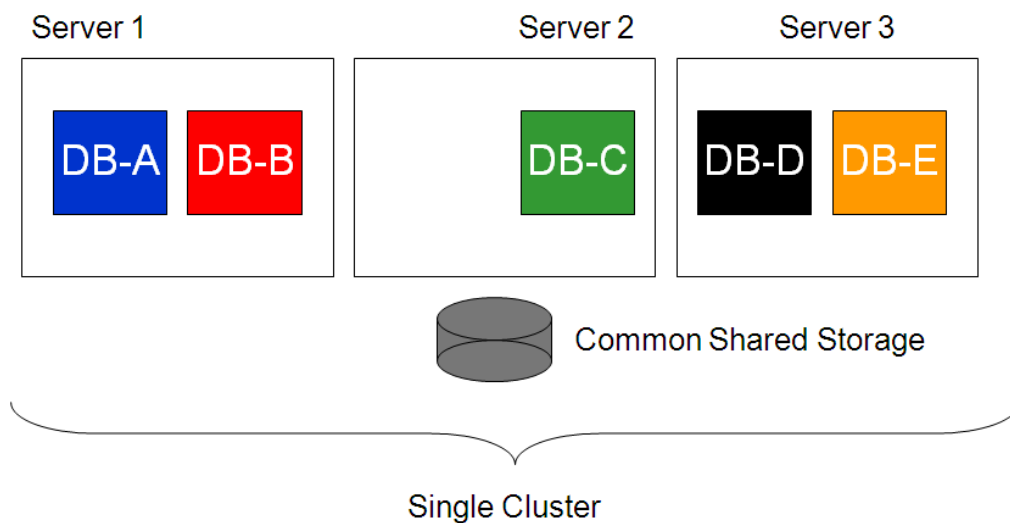
¹ Virtual servers always suffer a performance handicap vs. physical hardware. This is because CPU instructions must be translated in a virtual server, adding to CPU overhead.

Better Server Consolidation

RAC One Node provides superior consolidation by leveraging the benefits of a shared operating system (OS):

- **One OS** to install, configure, secure, patch, upgrade, backup, manage vs. multiple operating systems in a VM environment
- **Resource efficiency of running one OS vs. multiple operating systems**

A typical Oracle RAC One Node deployment will resemble the figure below:




In this example configuration, we have five single-instance Oracle RAC One Node databases running in a cluster of three servers. Server 1 is hosting Oracle RAC One Node databases DB-A and DB-B, server 2 is hosting database DB-C and server 3 is hosting databases DB-D and DB-E. Each server runs one OS. In servers A and C above, multiple databases are consolidated onto a single OS.

RAC One Node’s OS consolidation model presents the system administrator with a single OS (per server) to manage. In contrast, in a VM environment it is not unusual to have a dozen operating systems installed on a single physical server, presenting the system administrator a dozen operating systems to install, configure, patch, secure, upgrade, backup and manage in general.

Further, Oracle RAC One Node’s shared OS consolidation only requires server resources (CPU, memory) to support one OS. This conserves server resources and enables running more databases on that server. On the other hand, virtual server environments run multiple

guest operations systems on a given physical server. This model is more resource intensive and potentially reduces the number of databases that can be consolidated on a given server.



Instance Caging


Oracle Database 11g Release 2 Instance Caging feature provides resource isolation across databases. This feature helps deliver consistent service levels, without the overhead and inefficiencies of more cumbersome approaches. Instance Caging allows the administrator to limit the CPU used by an instance, preventing runaway processes in one instance from impacting others sharing that server. Should requirements change, the administrator can dynamically change the CPU allocation without taking the database offline. In contrast, VMs require a reboot when changing CPU allocation. More information on Instance Caging is available on Oracle.com.²

Enhanced Protection from Failures

² See: <http://www.oracle.com/technetwork/database/features/performance/instance-caging-wp-166854.pdf>

Oracle RAC One Node responds to both server and database failures with unattended cluster failover that relocates the impacted database service. VMs only respond to entire server failures and are unable to detect and respond to database failures.

Below is a depiction of Oracle RAC One Node failover.



In the above scenario, Server 2 failed, and database instance DB-C that was running on Server 2 is started on Server 3.

Oracle RAC One Node is integrated with Oracle Clusterware which monitors the health of the database and ensures database service availability. In the event of a failure, Oracle RAC One Node will detect the failure, and either restart the failed database, or fail it over to another server. Comparable to other 3rd-party cold failover solutions from vendors like HP, IBM, and Symantec, Oracle RAC One Node cold failover ensures unattended recovery from failures, with full services restored within 5 to 30 minutes. It is ideal for those applications that can tolerate short

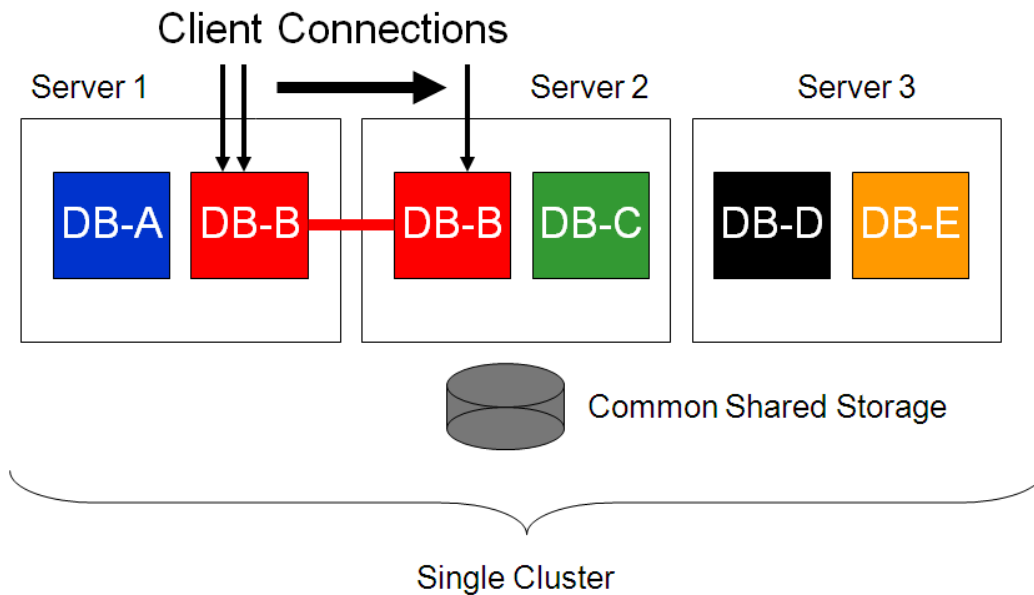
periods of downtime, but cannot tolerate waiting for an administrator to detect the problem and restore services.³

In contrast, the type of failures VMs can detect and recover from is somewhat limited. They typically only monitor the health of the VM and the operating system running within. Should the database instance crash or hang, they would not detect the failure and would not restart the VM without administrator intervention.

Greater Flexibility and Workload Management

Oracle RAC One Node online database relocation allows a running instance to be migrated from one server to another without disruption of service.

Below is a depiction of Oracle RAC One Node online database relocation:



³ Customers who need sub-minute recovery should deploy their databases on multi-node Oracle Real Application Clusters (RAC). Oracle RAC provides the best possible availability and fastest recovery from failures.

Online database relocation provides the same load balancing benefits of VMs by allowing a migration of a database from a busy server to a server with spare capacity. Online database relocation leverages the ability of Oracle Real Application Clusters to simultaneously run multiple instances servicing a single database. In the figure above, the DB-B RAC One Node database on Server 1 is migrated to Server 2. Oracle RAC One Node starts up a second DB-B instance on server 2, and for a short period of time runs in an active-active configuration. As connections complete their transactions on server 1, they are migrated to the instance on server 2. Once all the connections have migrated, the instance on server 1 is shutdown and the migration is complete.

Online database relocation does not require quiescing the environment even when the system is running at peak capacity. VMs generally require the environment to be quiesced in order for medium to heavy database workloads to be moved from one server to another. This requirement does not apply for light workloads or demos.

When migrating a VM, the VM must mirror its complete memory state across a network to the target host, recreating the state of that machine. If the database in question is highly loaded and is actively changing blocks in its database cache, it is very difficult for the memory mirroring function to keep up with the rate of changes. It becomes likely that the only way to successfully mirror the memory is to quiesce the source VM, mirror the memory, and then switch to the migrated VM. With online database relocation, highly loaded instances pose no problem, as the work is actually split between two servers during the migration. Oracle RAC One Node can therefore easily migrate even heavy workloads to another server.

Online database relocation provides the ability to move databases between servers of different processor generations. VM migration requires that the processors be identical – it's not sufficient that both processors be Intel or AMD. Both Intel (or AMD) processors must have exactly the same instruction set. Online database relocation supports migration to new processors (i.e., Nehalem), or even between different vendors (Intel and AMD).

RAC One Node provides superior scale-up and does so without service disruption. VMs allow increasing the number of CPUs available to a database instance by moving it to a larger server. However, they require a reboot of the OS for the increased CPUs to take effect. Also, many VM solutions limit VM size to a limit that is rather low for a database server. With RAC One Node, the database automatically adjusts to the larger server and takes advantage of the additional CPUs – no reboot required. Also, there is no CPU limit for RAC One Node.

Better Online Maintenance

A unique feature of Oracle RAC One Node is the ability to do online maintenance for both database and operating system. VMs do not.

RAC One Node administrators can online migrate the databases off the server to a spare server, and then perform all types of maintenance on the idle server, including hardware maintenance, OS upgrades, OS patches, and database patches. This is often a point of confusion among customers, as VMs promise a similar benefit. However, with VMs, the entire operating system and Oracle Home are migrated to the new host server, leaving no operating system or database behind to maintain (although you are free to dust your hardware). With online database relocation, a new database instance is created on a new server (running in a different operating system), and work is online migrated to the new instance. Thus, the old OS and database home remain on the former host server, and can be upgraded (OS) or patched.

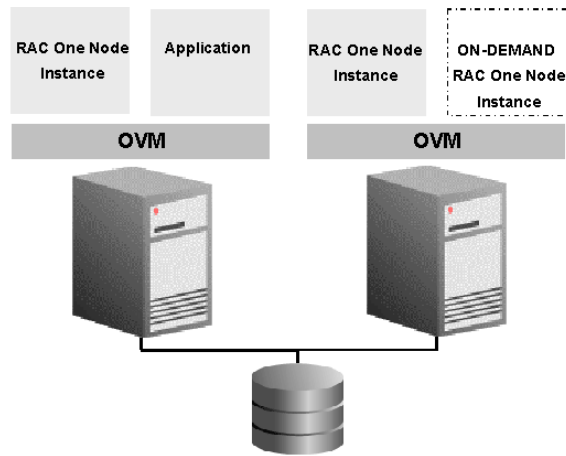
Other Unique Benefits of Oracle RAC One Node

There are additional benefits that are unique to Oracle RAC One Node. These include:

- **Storage Virtualization:** Oracle RAC One Node provides storage virtualization via a feature called Automatic Storage Management (ASM). ASM virtualizes all the storage presented to the database and automates management and tuning and seamlessly handles storage reconfigurations either due to disk failure or disk add/drop events. VMs do not provide storage virtualization.
- **Storage Consolidation:** Further, ASM increases storage utilization via a feature called storage “pooling”, where all the databases on a single server (or cluster) running Oracle RAC One Node share a single storage pool. Database disk IO is always in tune since ASM ensures that IOs are always balanced across all spindles – i.e. avoid hotspots. Free disk space is managed centrally vs. being fragmented across multiple local storage pools. VMs do not provide storage consolidation.
- **Non-disruptive scale-out:** Oracle RAC One Node can be online upgraded to Oracle RAC (with appropriate license) in order to scale beyond one server. Over time, applications may not only grow their workload, but may also become business critical. Should an application eventually require the enhanced availability provided by Oracle RAC, or scalability beyond a single physical server, Oracle RAC One Node can be upgraded to Oracle RAC with no downtime or disruption. VMs do not support scale-out period.
- **Standardized Operating Model:** By running a combination of Oracle RAC One Node and Oracle RAC (for scale-out databases), customers can embrace a single database deployment model – i.e. tools, IT processes, custom scripts, etc. – across all operating systems. This helps streamline IT, reduces education and redundant work, confusion and possible human error.

Complements Server Virtualization

Although Oracle RAC One Node improves upon many of the benefits of VMs, it can be utilized as a complementary technology to Oracle VM. Oracle RAC One Node is fully supported in Oracle VM environments—any node in the cluster could be an Oracle VM virtual machine. This extends the technology to some interesting use cases.



RAC One Node complements OVM

- **Mixed database and other workload environments:** Combining VMs and Oracle RAC One Node allows you to create a flexible environment for mixed workloads, where only some virtual servers are hosting databases on a given node.
- **Improving high availability of Oracle Database in a Virtual Machine:** Running Oracle RAC One Node inside VM extends the superior HA characteristics of Oracle RAC One Node to a database running inside a VM.
- **Online Upgrading a Single Physical Server:** If you only have a single physical server, combining RAC One Node and two virtual machines on the one physical server, enables you to online upgrade or patch your environment by using online database relocation. This is accomplished by migrating the database to a temporary virtual server for the maintenance period.

Conclusion

Simple server virtualization via VMs have inherent limitations. Oracle RAC One Node improves on VM benefits and overcomes its limitations. Those customers looking to reduce server footprint, improve availability, better manage workloads, reduce maintenance outages, streamline

database management, or ramp to multi-node Oracle RAC will find that Oracle RAC One Node provides better virtualization for their databases.



Oracle Real Application Clusters One Node
November 2010
Author: Bob Thome
Contributing Authors: Sohan DeMel

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



| Oracle is committed to developing practices and products that help protect the environment

Copyright © 2009, Oracle and/or its affiliates. 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 is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.