

Oracle Real Application Clusters *One Node*

Technical Overview

ORACLE WHITE PAPER | MARCH 2018







Table of Contents

Executive Overview	1
Oracle RAC <i>One Node</i> – Overview	2
The Oracle RAC Stack	2
Oracle Clusterware	3
Oracle Automatic Storage Management (ASM)	3
Oracle ASM Cluster File System (Oracle ACFS)	3
Standardize on Oracle RAC (One Node)	4
Best In-Class Oracle Database Availability	5
Example: Using Online Database Relocation to perform Zero Downtime Patching	6
Optimized Online Patching in Consolidated Environments	8
Better Oracle Database Consolidation	8
Oracle RAC <i>One Node</i> Support for Oracle Multitenant	10
Better Oracle Database Virtualization	10
Additional Benefits of Oracle RAC <i>One Node</i>	12
Summary and Conclusion	13



Executive Overview

Oracle Real Application Clusters (RAC) *One Node* is an option to the Oracle Database Enterprise Edition that was introduced with Oracle Database 11g Release 2. It provides enhanced high availability for single instance Oracle Databases, protecting them from both, planned and unplanned downtime.

Oracle RAC *One Node* provides the following benefits:

- » Best in-class Oracle Database Availability
- » Better Oracle Database Consolidation
 - » Including Oracle Multitenant support
- » Better Oracle Database Virtualization
- » Oracle RAC One Node also allows customers to standardize their database deployment and management, consolidate database storage and, should the need arise, upgrade to a full multi-node Oracle RAC database without downtime or disruption..

Oracle RAC One Node – Overview

The Oracle Database with the Oracle Real Application Clusters (RAC) One Node option benefits from the same infrastructure used for Oracle RAC. Unlike Oracle RAC, Oracle RAC One Node normally runs only one instance against a shared set of data files, also known as the database. This database is fully Oracle RAC-enabled, but does not span multiple hardware systems at the same time.

Instead, the Oracle RAC One Node database instance will failover to another server in the cluster should a server, instance or a related and monitored component on this server fail. For cases of planned downtime such as OS or database patching, Oracle RAC One Node provides a unique feature, Online Database Relocation, which allows relocating a database from one server to another without interrupting the database service.

Addressing planned as well as unplanned downtime makes Oracle RAC One Node the best in-class Oracle Database failover solution. The ability to Online Upgrade to a multi-node Oracle RAC database complements its functionality and makes it an ideal infrastructure for database cloud deployments.




Figure 1: Oracle RAC One Node and Oracle RAC Stack Overview

The Oracle RAC Stack

An Oracle RAC One Node database requires Oracle Clusterware to be installed on the system prior to installing an Oracle RAC enabled database home. Oracle Clusterware (OCW) is an integral component of the Oracle Grid Infrastructure (GI) suite of products, which includes Automatic Storage Management (ASM) and the Oracle ASM Cluster File System (ACFS).

Oracle Grid Infrastructure, including Oracle ASM / ACFS and Oracle Clusterware, and the Oracle Database with the Oracle Real Application Clusters option constitute the Oracle RAC Stack.

Oracle Grid Infrastructure is the foundation for an Oracle RAC database system using Oracle Clusterware for automatic resource placement, adjustment and management, Oracle Automatic Storage Management and Oracle ACFS for efficient and reliable storage and file system management.



Oracle Clusterware

Oracle Clusterware is a technology transforming a server farm into a cluster. A cluster is defined as a group of independent, but connected servers, cooperating as a single system. Oracle Clusterware is the intelligence in this system that provides the cooperation.

Oracle Clusterware was introduced with Oracle Database 10g Release 1 as the underlying clustering software required for running Oracle Real Application Clusters databases. As part of the Oracle RAC Stack, Oracle Clusterware is also used by the clustered version of Oracle ASM and is tightly integrated into the Oracle RAC Stack, which is also used for Oracle RAC One Node.

Oracle Clusterware is a complete, free of charge clustering solution that can be used outside of Oracle RAC ¹. In such environments Oracle Clusterware is commonly used for automatic resource placement, adjustment and management of any kind of application. In both environments, Oracle Clusterware is responsible for maintaining node membership and ensuring fencing.

Oracle Automatic Storage Management (ASM)

Since Oracle RAC One Node utilizes a shared-disk architecture, the volume management and file system used for storing database data must be cluster-aware. Oracle Automatic Storage Management (ASM) is the recommended (clustered) volume manager for the Oracle Database.

Oracle ASM manages all data: Oracle database files, Oracle Clusterware files, and non-structured data such as binaries, external files and text files (via Oracle ACFS). Through its low cost, easy administration and high performance characteristics, Oracle ASM is the storage technology of choice for managing Oracle databases.

For performance and high availability, Oracle ASM follows the principle of “stripe and mirror everything”. Intelligent mirroring capabilities allow administrators to define 2- or 3-way mirrors to protect vital data. When a read operation identifies a corrupt block on a disk, Oracle ASM automatically relocates the valid block from the mirrored copy to an uncorrupted portion of the disk. ²

Oracle ASM Cluster File System (Oracle ACFS)

Oracle ASM Cluster File System (ACFS) is an industry standard, POSIX, X/OPEN and Windows compliant cluster file system. It includes advanced features such as file system snapshots, replication, tagging, security, encryption, auditing, snapshot-based replication, compression, automatic resize, support for sparse files, metadata acceleration and highly available NFS and SMB services.

Oracle ACFS is integrated with Oracle Automatic Storage Management, extending its functionalities to support general-purpose files. Starting with Oracle Database 12c Release 1, ACFS also supports Oracle Database files.

Oracle ACFS leverages Oracle Clusterware for cluster membership state transitions and resource-based high availability. Oracle ACFS, Oracle ASM and Oracle Clusterware are all bundled into the Oracle Grid Infrastructure (GI) allowing for integrated optimized management of databases, resources, volumes and file systems. ³

¹ For information about Oracle Clusterware, see www.oracle.com/goto/clusterware

² For information about Oracle ASM, see www.oracle.com/goto/asm

³ For information about Oracle ACFS, see www.oracle.com/goto/acfs

Standardize on Oracle RAC (One Node)

Oracle RAC One Node utilizes multiple features that make it the ideal standard database deployment infrastructure. From an application point of view, Oracle RAC One Node appears – for all practical purposes – as an Oracle Single Instance Database towards the application. Therefore, there is hardly ever any need to certify an application with Oracle RAC One Node explicitly, as long as the application is certified to run against an Oracle Database.

From an administration point of view, Oracle RAC One Node shares and benefits from the same architecture that has been used for Oracle RAC for more than a decade. This design allows Oracle RAC One Node to use the Online Database Relocation-feature for uninterrupted database service in case of planned downtime, typically used during patching the OS or the database home for example.

A Single Instance database, even when deployed on a cluster, needs to be taken down in this case.

These features as well as the fact that an Oracle RAC One Node Database can be online upgraded to a multi-node Oracle RAC database distinguish Oracle RAC One Node from simply running an Oracle Single Instance Database, which is a silo-approach.




Figure 2: Standardize on Oracle RAC One Node for maximum flexibility

The main difference between Oracle RAC One Node and Oracle RAC is that with an Oracle RAC One Node Database, there is only one instance running at a time under normal operation. Should this one instance be impacted by unplanned downtime, there is a short period of time during which no database instance is running.

The stack used for Oracle RAC One Node ensures that the database service will be resumed as fast as possible, but technically and as with any other failover cluster solution, there is a moment of time, during which no instance of the Oracle RAC One Node Database is available to perform any work. This would not be the case in a multi-node Oracle RAC system, because multiple instances are already running concurrently, accessing the same database.

For planned downtime Oracle RAC One Node provides a solution superior to any other failover solution on the market; "Online Database Relocation". Online Database Relocation is an Oracle RAC One Node-specific feature that enables an Oracle RAC One Node database instance to be relocated to another server in the cluster without interrupting the database service. In an Oracle RAC database this feature is not required, as databases instances are typically running on multiple servers in the cluster already. Relocating a running Oracle RAC database instance is unnecessary.

Oracle RAC One Node is a cluster-based database failover solution. Its integration with the Oracle RAC Stack, especially Oracle Grid Infrastructure, makes it the best solution in this class of solutions available on the market today.

Best In-Class Oracle Database Availability

For failure cases, Oracle RAC One Node ensures the lowest possible failover times. In addition, Oracle RAC One Node can make use of Oracle's latest Business Continuity feature "Application Continuity", which is included in the Oracle RAC, Oracle Active Data Guard and Oracle RAC One Node option.

Utilizing Application Continuity as part of the failover strategy minimizes the downtime experienced by an application running on an Oracle RAC One Node Database, but also helps to optimize planned downtime as it reduces the need to free up work on the database instance to be patched for example.⁴

Online Database Relocation completes the Oracle RAC One Node functionality making Oracle RAC One Node the best in-class Oracle Database failover solution.




Figure 3: Oracle RAC *One* Node – Utilizes Application Continuity

⁴ For more information about Application Continuity and its integration into the Oracle RAC Stack, see: <http://www.oracle.com/technetwork/database/options/clustering/applicationcontinuity/overview/index.html>

Example: Using Online Database Relocation to perform Zero Downtime Patching

Online Database Relocation allows maintenance operations such as server firmware, OS or database patching to be performed without downtime of the database service (online).

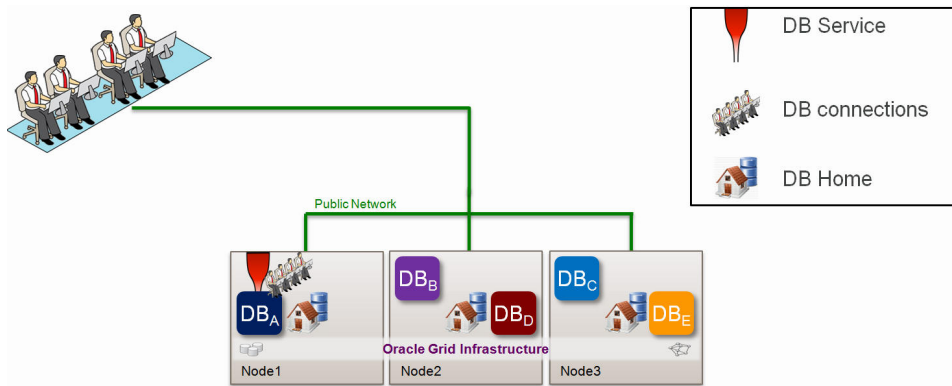


Figure 1: Zero Downtime Patching using Online Database Relocation – Step 1: status quo

While Online Database Relocation is a standard Oracle RAC One Node feature, using it for zero downtime maintenance requires some preparation. In figure 4, a three node cluster hosts various Oracle RAC One Node Databases:

- » Node 1 hosts Oracle RAC One Node Database A
- » Node 2 hosts Oracle RAC One Node Database B and D
- » Node 3 hosts Oracle RAC One Node Database C and E

The Oracle Database home for each Oracle RAC One Node database is installed on a per-server-basis (as opposed to using a shared Oracle Database home). Using an individual database home per database allows for online patching of individual databases on a server, as described below.

For simplicity reasons, the following example will only focus on Oracle RAC One Node database A. In a real world example, other databases that might be running on the same server would have to be accounted for.

Given these assumptions, online maintenance using Online Database Relocation can be performed in four simple steps (which can be further optimized to three steps depending on the scenario):

1. Initiate Online Database Relocation from source to target
2. Patch the Oracle Database home on the source server
3. Rewind (relocate back to source server) to activate the patch usage
4. Patch remaining Oracle Database home(s)

Step one in the above list makes Oracle RAC One Node superior to any other database cluster failover solution on the market today. Based on the fact that Oracle RAC One Node is an Oracle RAC enabled database, Online Database Relocation can start a second instance for the Oracle RAC One Node database for the purpose of the relocation. This is the only time during which an Oracle RAC One Node database should have two database instances running at the same time.

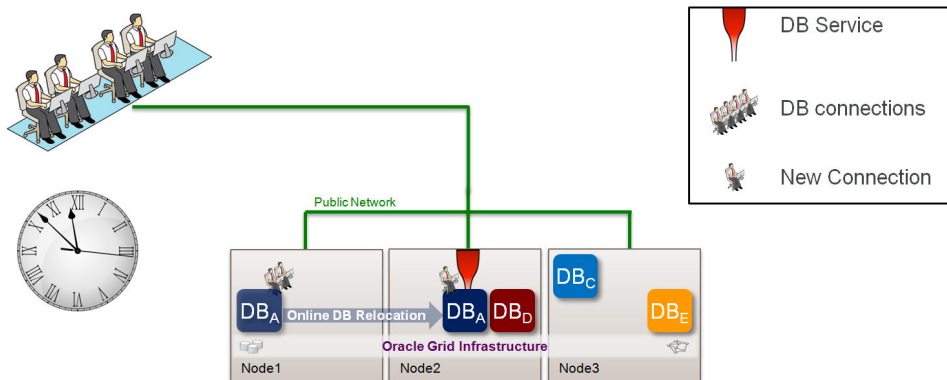


Figure 2: Zero Downtime Patching using Online Database Relocation – Step 2: Online Database Relocation

As part of the Online Database Relocation the following steps are performed:

- » The database service(s) are stopped on the source database instance (the instance that you want to stop) and then started on the target database instance (the new instance).
 - Note that Oracle RAC One Node for this reason requires at least one dynamic database service to be created as part of the DBCA-based database creation.
 - Further database services can be created as a post-installation step.
- » The database service started on the target database instance will cause new connections requesting access to the Oracle RAC One Node database go to the new target server.
- » Connections that were established at the moment of initiating an Online Database Relocation on the source server will remain on the source instance, while a shutdown transactional will be issued against the source database instance.
- » As a shutdown transactional will wait for sessions to finish the transaction and to disconnect, the Online Database Relocation feature allows for setting a timeout describing the time that it should wait before forcefully stopping the source database instance using shutdown abort.
 - The default wait time (timeout) is 30 minutes.
 - The wait time can be increased up to 24 hours.
- » After either the timeout has expired or the last session on the source database instance is closed (whichever comes first), the source database instance will be shutdown and the new (target) database instance remains as the only available instance.

As the source database instance is now stopped and assuming no other database is operated from the respective database home, the database home on the source server can now be patched (see step 3), implicitly making use of the Rolling Upgrade functionality that is inherent to Oracle RAC databases.

Step 4 foresees to rewind the operation, which means to relocate Oracle RAC One Node Database A back to the original home in order to activate the patch, as so far this database instance has only been running out of an un-patched database home. Starting the Oracle RAC One Node Database A instance on Node 1 will activate the patch for this instance and allow for patching the remaining home(s).

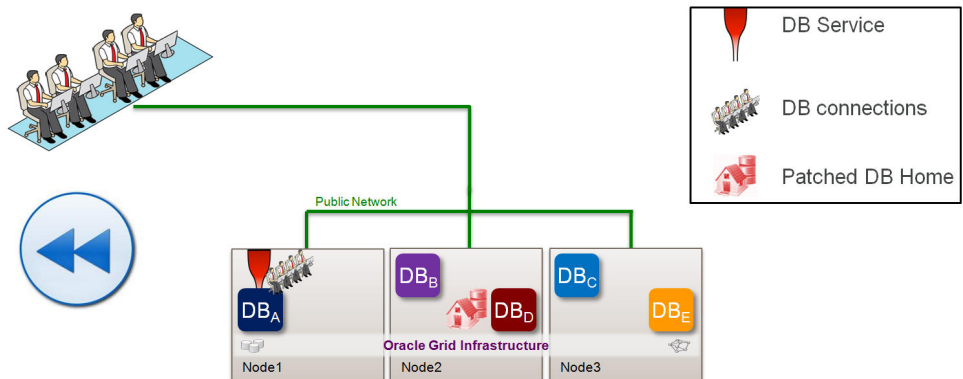


Figure 3: Zero Downtime Patching using Online Database Relocation – Step 4&5: Rewind & patch remaining

Optimized Online Patching in Consolidated Environments

The procedure outlined above called “Zero Downtime Patching using Online Database Relocation”, assumes that a particular database instance is relocated twice as part of the workflow, which leads to four steps to be executed. In consolidated environments and assuming that there is no specific server, on which a database instance needs to be hosted, the above procedure can be shortened to 3 steps:

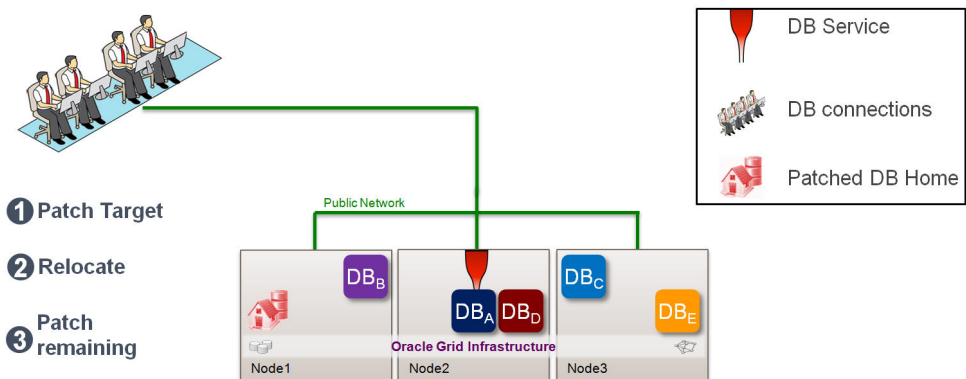



Figure 4: Zero Downtime Patching using Online Database Relocation – 3-step Optimization

Better Oracle Database Consolidation

Consolidation has become a hot topic in the IT industry. Oracle RAC One Node enables better server consolidation, enhanced protection from failures, greater flexibility, easier workload management as well as better online maintenance than virtualized environments.


Oracle RAC One Node provides superior consolidation by leveraging the benefits of a shared operating system (OS) image. This means, there is only one OS to install, configure, secure, patch, upgrade, backup, manage vs. multiple operating systems in a Virtual Machine (VM) environment.

Virtual Machines




share servers

Dedicated Databases



share servers and OS

Pluggable Databases



share servers, OS and database



Figure 5: Consolidation schemes used today

Oracle 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 of operating systems installed on a single physical server, presenting the system administrator with a dozen operating systems to install, configure, patch, secure, upgrade, back up and manage.

Using a single OS image to host multiple Oracle Databases presents the problem of isolation of resources, of which CPU utilization next to memory utilization is the most critical. Instance Caging⁵ is a feature introduced with Oracle Database 11g Release 2 that provides the required resource isolation across databases.

This feature helps delivering 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. The administrator can dynamically change the CPU allocation without taking the database offline, should requirements or demand change on the system.

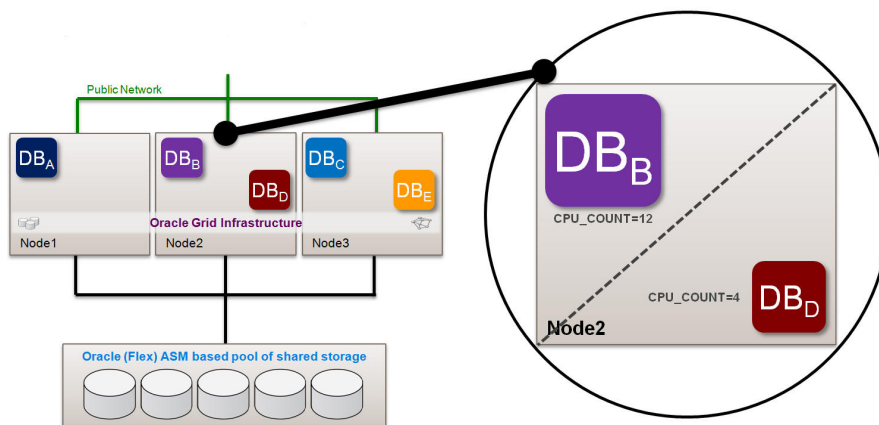


Figure 6: Oracle RAC One Node - Isolation through Instance Caging

⁵ For more information about Instance Caging, see: <http://www.oracle.com/technetwork/database/focus-areas/performance/instance-caging-wp-166854.pdf>

Oracle RAC One Node Support for Oracle Multitenant

Oracle Multitenant is a new option to the Oracle Database 12c Enterprise Edition that helps customers reduce IT costs by simplifying consolidation, provisioning, upgrades and more. It is based on a new architecture that allows a multitenant container database to hold many pluggable databases (PDBs). The idea is that an existing database can be simply adopted, with no changes in the application tier, as a pluggable database. In this architecture, Oracle RAC One Node provides local failover based high availability that is required when consolidating various business critical applications on one system.




Figure 7: Oracle RAC One Node Multitenant Container Database Support - Overview

When using PDBs with Oracle RAC One Node, the multitenant container database (CDB) is based on Oracle RAC One Node. Each pluggable database is made available on the Oracle RAC One Node CDB instance. Access to and management of the PDBs are regulated using Dynamic Database Services, which will also be used by applications to connect to the respective PDB – as they would in a Single Instance Oracle Database using Oracle Net Services for connectivity.

Better Oracle Database Virtualization

Virtualization has become a hot topic in the IT industry. This begs the question why virtualization has become so popular. One answer could be that 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 is easy to prove that the investment will have a positive ROI.

Oracle RAC One Node runs on physical servers and therefore does not suffer from potential overhead of virtual servers. 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.

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 VM-based 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 virtual machine 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.




Figure 8: VM-based HA features do not have any insight as to what the guest VM is running

A granularity that most virtualization solutions do not provide is one that goes beyond the black-box of a Guest VM. In other words, for most VM-based solutions and this includes Oracle's, the Guest VM is a black box, on which features like Live Migration or VM-based High Availability will operate. These features do not have any insight as to what the guest VM is running.⁶

It is Oracle RAC One Node and the Oracle RAC Stack that provides the required insight and functionality to fully utilize VM-based environments. Live Migration might be able to move a Guest VM from one server to another; however, it cannot move only the database between Guest VMs.

Operating on the application (database) level within the Guest VM makes a huge difference when it comes to maintenance operations. Each Guest VM still uses an OS and probably provides the Oracle Database home binaries; both entities will be subject to patching on a regular basis.

Similar applies to VM-based HA features as depicted in figure 11 (on the right). While VM-based HA works perfectly fine on Guest VM-level, which means that it will fail over a Guest VM should a failure be detected, it will neither detect nor correct a failure of an application (the database) within a guest.

It is the Oracle RAC One Node Stack that provides the missing pieces and functionality such as application-level availability monitoring and Online Database Relocation for maintenance operations.

⁶ For more information on Live Migration and Oracle VM-based HA, see page 15 of this paper: <http://www.oracle.com/technetwork/products/clustering/oracle-rac-in-oracle-vm-environment-131948.pdf>

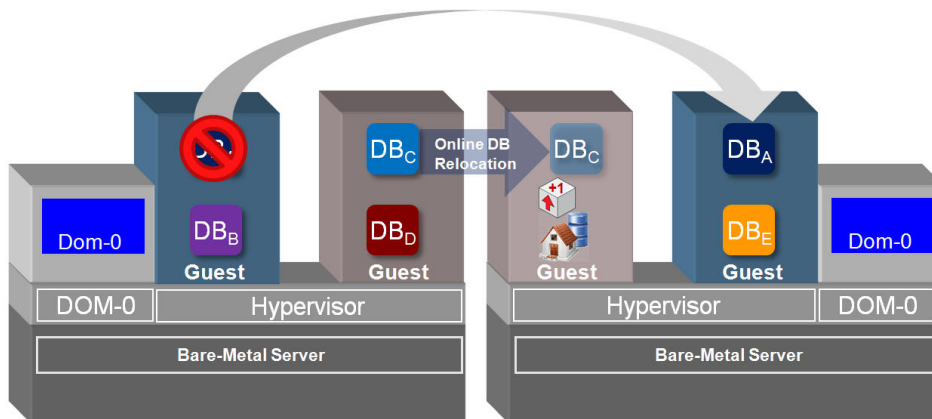


Figure 9: Oracle RAC One Node provides HA even in virtual environments

Additional Benefits of Oracle RAC One Node

Additional benefits that are unique to Oracle RAC One Node include:

- » **Storage Virtualization:** Oracle RAC One Node provides storage virtualization via Oracle Automatic Storage Management (ASM). ASM virtualizes all the storage presented to the database and automates management and tuning and seamlessly handles storage re-configurations either due to disk failure or disk add/drop events.
- » **Storage Consolidation:** In addition, 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 versus being fragmented across multiple local storage pools.
- » **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.
- » **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.

Summary and Conclusion

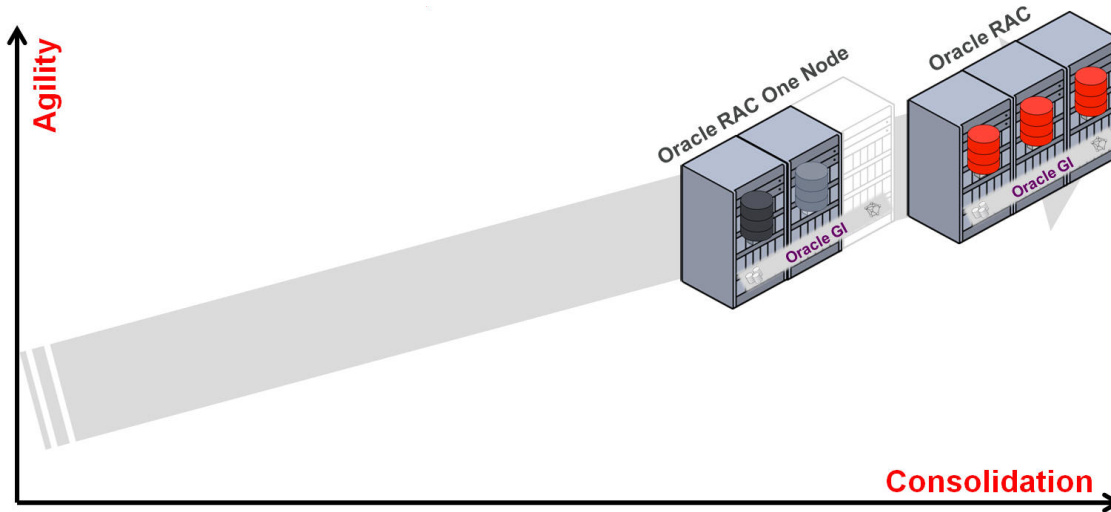


Figure 10: Oracle RAC and Pluggable Databases - The New Standard for Database Consolidation

Oracle RAC One Node provides the best in-class Oracle Database Availability. It provides features to avoid both unplanned and planned downtime. Oracle RAC One Node also improves 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 up to multi-node Oracle RAC will find that Oracle RAC One Node is an ideal solution for their databases and database cloud deployments.







Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2018, 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0318

Oracle Database 18c Real Application Clusters *One Node*
March 2018
Author: Markus Michalewicz
Contributing Authors: