

Oracle Database 12c Release 2 with Oracle Real Application Clusters – New Features Overview

ORACLE WHITE PAPER | MARCH 2017





Table of Contents

Introduction	2
It's all about Scalability, Availability and Efficient Management	2
Oracle Database 12c Release 2 Real Application Clusters	3
Large Number of RAC Database Deployments Creates a Special Challenge	3
Better Scalability	3
Smart Mastering	4
Multitenant in an Oracle RAC environment – further optimizations.	4
Pluggable Database and Service Isolation	4
Service Isolation in a PDB Environment Improves Failover Times	4
Flex Cluster – The Extreme Scalability Architecture	6
Use Case 1: RAC Reader Nodes	6
Use Case 2: Massive Parallel Query RAC	6
Even Better Database Availability	7
Application Continuity	7
Intelligent Fencing with Node-Weighting	7
Pluggable Database and Service Isolation	8
Recovery Buddy	8
More Efficient and Intelligent Management	9
Cluster Domains and Domain Services Clusters (DSC)	9
Cluster Domain	9
Member Clusters	10
Centralized ASM Services	11
Benefits of Cluster Domain Architecture	12
Oracle's Autonomous Health Framework (AHF)	12
Introducing Cluster Health Advisor Powered by Oracle Machine Learning	13
Quality of Service Management – Support for Administrator-Managed Databases	14
Cluster Activity Log	14
Rapid Home Provisioning (RHP)	14
Conclusion	15



Introduction

Oracle Real Application Clusters (RAC) was introduced in 2001 and quickly became the de facto standard technology for **local database high availability** for many organizations' mission critical systems. Oracle has now released the eighth version of Real Application Clusters and it has grown to become the most popular database option for the Oracle Database.

It's all about Scalability, Availability and Efficient Management

As Oracle Real Application Clusters matures, the development focus has been to respond to the market requests for greater scalability, even higher database availability, and more efficient management of Oracle RAC estates.

Oracle Real Application Clusters 12c Release 2 addresses those requirements:

- Scalability and better Performance
- Even Higher Database Availability
- More efficient management
 - For large scale Oracle RAC deployments.
 - Improved diagnosability with predictive analytics powered by Oracle Machine Learning

Oracle RAC is a critical technology for on-premises deployments. RAC is also critical technology for public cloud deployments. Oracle Real Application Clusters 12c Release 2 was first released in the Oracle Public Cloud as a Database Cloud Service even before the on-premise version was made available. Oracle's public cloud is the only public cloud that offers a standard, supported, easy-to-build and easy-to-deploy Oracle RAC solution for database high availability.

With Oracle Database 12c Release 2, Real Application Clusters focuses its improvements to serve both large enterprises that have expertise in Oracle RAC and clustering technologies as well as smaller organizations that may not have any Oracle RAC experience yet.

Oracle Database 12c Release 2 Real Application Clusters

This eighth major release of Oracle Real Application Clusters offers improvements for all customers from the largest IT shops to small businesses.

Large Number of RAC Database Deployments Creates a Special Challenge

Oracle RAC has become the de facto standard for local database high availability for tens of thousands of Oracle customers worldwide. Many of these customers have deployed hundreds of standalone clusters within their organization supporting thousands of databases. With each cluster and each database more monitoring and management tools get deployed that have to be maintained, patched, and upgraded. Due to the increase management overhead, especially in large database estates, customers up until tried to avoid deploying those additional tools, often resulting in a lack of management capabilities or diagnosability data in critical situations.

With Oracle Real Application Clusters 12c Release 2, that problem is drastically reduced through these improvements:

- Better scalability and better performance
- Even better database availability
- More efficient and intelligent management for large cluster estates

For the thousands of current Oracle RAC customers, this release will significantly improve the manageability of their RAC estate – the hundreds of clusters and potentially the thousands of databases they must manage. This paper introduces the new features that enable these improvements.

For new, potential customers of RAC, this release will provide an easier way to implement RAC and detect problems when or if they occur. All customers want things that are simple (easy), and fast and that offer efficiency or enhanced productivity. This release of Oracle Real Application Clusters, Oracle Clusterware, and Automatic Storage Management delivers improvements that offer these things.

Better Scalability

Oracle Real Application Clusters 12c Release 2 offers improved scaling for the popular “all-Hub node” standalone clusters (the type of clusters most customers have implemented) and the newer “Flex Cluster” configuration introduced in Oracle Real Application Clusters 12c Release 1.

Oracle RAC scales independently of the number of nodes and does *not* require application changes unlike sharding. Oracle RAC 12c Release 2 includes many improvements to scalability, especially for singleton workloads, where a database service runs on just one instance on one node of the cluster only – a configuration that is popular in consolidated database environments.

The scalability improvements include:

What is Real Application Clusters?

Oracle Real Application Clusters is software that allows and Oracle Database to run simultaneously on two or more different physical or virtual servers. The servers which are clustered via a private network, are coordinated through another software element called Oracle Clusterware (part of the Oracle RAC license) which is provided by Oracle or third parties like Veritas, HP or IBM. Oracle RAC is used primarily as a way of ensuring database availability – that is, it protects against both physical hardware and also software outages keeping the database accessible to the applications and the end users using them.

In the early 21st century, RAC, with its unique active-active architecture, had no direct competitor. It was the only database solution that could offer a sub-30 second recovery from database instance or hardware failures. When combined with complementary technology like Application Continuity, Transparent Application Failover (TAF) and Fast Application Notification (FAN), failures can be managed so that there is no perceptible impact to the end user of database tier failures. Oracle RAC became the standard for low latency data protection with many thousands of customers implementing it since 2001.

Over time IT shops discovered that the unrivaled clustering capability of Oracle RAC gave them the ability to host multiple databases on a cluster of Linux-x86 machines thereby enabling higher levels of database consolidation. Isolated silo's of machines, software licenses and staff were now giving way to a shared environment – at least for the database tier. Sharing brought cost savings, and leverage of hardware, facilities and staff costs across multiple databases and applications. Oracle Real Application Clusters was the key technology behind this – RAC allowed multiple databases and database instances to share resources safely – allowing the most important work to have the resources it needed to maintain the service level objectives they were required to meet.

In an entire database estate, many customers deployed Oracle RAC only for their Platinum Tier 1 or mission critical systems but discovered it also enables database consolidation as well as high availability for those databases. Consequently, many customers deployed Oracle RAC in their Gold, Silver and even Bronze tiers – across their entire database estate.

- Smart Mastering
- Multitenant optimizations and Service Isolation for Real Application Clusters
- Flex Clusters for Extreme Scalability

Smart Mastering

The Service-oriented Buffer Cache Access, or “smart mastering” is a facility that enables the database to determine what objects are referenced by which service(s). Cache Fusion now maintains a “Service-to-Buffer-Cache” relationship where it tracks which service causes row(s) to be read into the buffer cache. This statistic is used to master the resource (data blocks) only on those nodes where the service is active. The result is that when a request is made to access or modify data, the data block containing that data will be available on the node requesting that service. The Global Cache Service (GCS) which is the heart of Cache Fusion, doesn’t have to find, lock, and deliver the data block across cluster nodes – the block is already there. Access happens much faster. This means that larger clusters are practical to run and can host multiple databases for greater consolidation because the inter-node cluster communication of cache fusion is reduced.

Storing this information in the database ensures that on the next startup, the formerly established affinity can be restored without the need to re-master using Dynamic Re-mastering (DRM). Unless the access patterns have changed significantly, this will lead to a consistent and smooth operation of the system right after startup.

Optimized “Resource Master” Dispersion Pre-Warms the cache during planned downtime.

One additional benefit of this tracking is the ability to determine which objects will be accessed by a given service before starting this service, and thus which data should be in the buffer cache. When migrating a service for maintenance, for example, the instance receiving the service can intelligently pre-warm its cache prior to migrating the service. This eliminates the “first-access” performance hit associated with a cold cache, eliminating any disruption to end-users affected by the relocation.

Multitenant in an Oracle RAC environment – further optimizations.

Pluggable Database and Service Isolation

With Oracle Database 12c Release 1, Oracle introduced a new database option, Oracle Multitenant. That option allows customers to host multiple Pluggable Databases (PDBs) in a Container Database (CDB). Using the Oracle Multitenant database option and Pluggable databases (PDBs) is a great way of consolidating workloads into a RAC database.

The CDB’s role was primarily to support database consolidation – running more databases but with better performance on the same or lesser hardware footprint. One of the ways this efficiency is achieved is by having one Container Database handle the writing of all redo log files for all its PDBs “plugged into” that CDB. This improved database log writing efficiency from both a hardware and software standpoint. It also improved database portability.

With Oracle Database 12c Release 2, the Multitenant option running in an Oracle RAC database is further optimized and offers much higher performance and scalability than the initial version of Multitenant running with Real Application Clusters.

Service Isolation in a PDB Environment Improves Failover Times

Given the increasing capability (CPUs and memory) of today’s servers, more customers make use of singleton PDBs or services, as they can easily run their workload in only one database instance. A singleton PDB is a single

Pluggable database or single tenant in a multitenant environment deployed in a Real Application Cluster environment. These deployments have only a single database and database instance and a single service defined and running in the database. They still rely on Oracle RAC for the availability of this singleton PDB and/or service and for the potential ability to scale out. If a failure occurs on that node or in that singleton database, that database and service is brought up on another node in the RAC cluster. It is much like a “cold failover” except with Oracle Real Application Clusters 12c Release 2, there are several improvements that make that failover occur much faster and in a more intelligent way.

Pluggable Database and Service Isolation **improves performance** by reducing Distributed Lock Manager (DLM) operations for Services not offered in all PDB instances. **Pluggable Database and Service Isolation improves availability** by ensuring that failures of instances only hosting singleton PDBs will not impact other instances of the same Oracle RAC-based CDB.

With the Oracle Database 12c Release 2 Multitenant and Oracle RAC options, the cluster now automatically locates services on the nodes to which the corresponding PDB is running. This now makes the Multitenant/RAC deployment scenario very appealing to DBAs and architects that desire the high consolidation/scalability capabilities of Multitenant but with the high availability of RAC.

Using singleton PDBs is even more compelling with Oracle RAC 12c Release 2, as they now benefit from **Pluggable Database Isolation**. Since singleton PDBs or services are only offered on one instance at a time, there is no need to share lock information for a given singleton PDB across instances. Database blocks associated with a PDB are only mastered on the nodes that the PDB is on – not on other nodes. PDB block management is more intelligent – mastering is intelligent because blocks are only mastered in the instances that reference or access them.




Figure 1: Pluggable Database and Service Isolation improves availability

In addition, Pluggable Database Isolation ensures that distributed lock management operations are only performed between those instances, in which the respective PDB(s) are open (in the example in Figure 1 above, the orange PDB is only open in two out of four instances). This reduces chatter across the nodes and improves performance and the scalability of the overall cluster and all databases running in it.

Availability is improved as well. After an instance failure, locks that were mastered by the failed instance will need to be re-mastered on another instance in the cluster. Instances running on surviving nodes who wish to access data whose locks were mastered on the failed node must pause while the affected locks are re-mastered. Since locks are mastered locally for singleton PDBs, they will be unaffected by the failure of other instances as there will be no affected locks that must be re-mastered — all are local, and thus none will have been mastered on the failed node.

Flex Cluster – The Extreme Scalability Architecture

Use Case 1: RAC Reader Nodes

Flex Clusters and the concept of Hub nodes and Leaf Nodes was introduced with Oracle Real Application Clusters 12c Release 1. Hub Nodes are the traditional cluster nodes that are interconnected via a private network and have direct access to shared storage. Leaf Nodes are loosely coupled nodes that are attached to Hub Nodes. Starting with Oracle RAC 12c Release 2, it is now possible to run a read-only workload on instances running on Leaf Nodes, which in essence are “Reader Nodes.” A Reader Node failure does not impact the overall database activity, making it easy to scale to hundreds of nodes. Customers will use this deployment option to handle Read-Only workloads running on read-mostly Leaf Node instances for ad hoc data analysis. **Read-only work can scale across hundreds of nodes** and does not result in any delay for accessing updated data and does not adversely impact OLTP performance work running on the Hub Nodes. The Leaf nodes supporting a large volume of read activity also benefits from high availability in case of a node or instance failure that Oracle RAC normally provides.

The bottom line – the Flex Cluster architecture provides for extreme scalability and significantly improved performance when both OLTP and Business Analytics workloads need to access the same data and run on the same physical infrastructure.

Use Case 2: Massive Parallel Query RAC

Many customers are now working with and deploying Hadoop Clusters for large scale data analytics. As shown in Figure 2 below, Oracle Real Application Clusters 12c Release 2, supports overlaying a Hadoop Cluster using HDFS with an Oracle Flex Cluster to access data in Hadoop via SQL and perform ad hoc analysis using standard interfaces. It is essentially marrying the traditional and well-known SQL and relational database models with the Hadoop cluster model to support analysis across a very large data set. Basically, a symbiotic relationship between RAC and Hadoop is created which supports the RAC environment retrieving data from Hadoop data sets and allowing analytics work to execute on the Leaf Node instances of the RAC cluster.

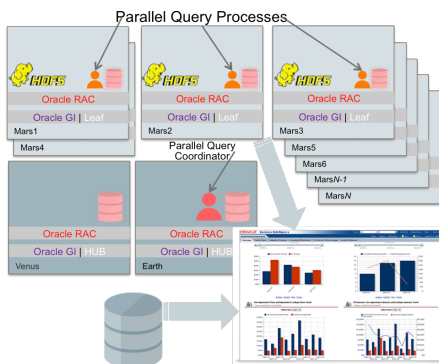


Figure 2. Hadoop HDFS combined with an Oracle Flex Cluster supporting the analysis of very large data sets but with the high availability offered by RAC.

Even Better Database Availability

Oracle RAC has always been about offering the best database availability. Its success since its debut in 2001 has established it as the de facto industry standard used to provide local database high availability. Further strides have been made to improve availability in Real Application Clusters 12c Release 2 with features such as:

- Application Continuity
- Node Weighting
- Pluggable Database and Service Isolation
- Recovery Buddy

As processors have become more and more powerful, customers are increasingly turning to solutions like Oracle RAC not just to scale workloads across multiple servers, but also to consolidate large numbers of single instance database workloads into an easy-to-manage high availability platform. This new release improves availability by reducing reconfiguration times after a failure.

Application Continuity

Application Continuity (AC) was a feature introduced with Oracle Real Application Clusters 12c Release 1. It allows for the successful completion of a database transaction even if a database instance fails before a transaction has been committed. This capability is a coordination between the database client, the client driver and the Oracle database

to mask outages from end users and applications by recovering and then completing the in-flight work for database sessions impacted by an outage. This recovery occurs beneath the application so that the outage appears to the application as a slightly delayed execution. The application or end user is completely shielded from an underlying database instance or node failure – the application does not have to replay the transaction, the Oracle database or the Oracle Real Application Clusters environment handles the completion of the transaction transparently.


Application Continuity is invoked for outages that result in recoverable errors typically related to underlying software, hardware, network, or storage layers. Without Application Continuity, it can be almost impossible for an application to mask outages in a safe and performing way. AC improves developer productivity by masking database-related outages without requiring the application developer to code elaborate error-checking and transaction replay logic.

The user experience is significantly improved with AC when handling both unplanned outages and planned maintenance. The important benefit of Application Continuity is that it strengthens the fault tolerance of systems and applications that use an Oracle database.

With Oracle Database 12c Release 2, Application Continuity is available for applications using Java JDBC, OCI and the ODP.NET unmanaged provider.

Intelligent Fencing with Node-Weighting

Prior to Oracle Real Application Clusters 12c Release 2, the outcome of a node eviction was rather predictable, but would not always lead to the most desired result. For example, in the case of a 2-node cluster where a “split in the middle” occurs, meaning the nodes cannot communicate with each other anymore using the interconnect and the voting disks, the node with the lowest node number survived. While this is a predictable outcome, it could mean that a perfectly healthy node is evicted, while the surviving node could suffer from a secondary failure or may actually not



be running any work as in the case of an Oracle RAC One Node database. In fact, the evicted node could have been running substantially more work or more important work but was evicted because of this simplistic cluster node survival logic (lowest node number survives).

Node Weighting is a new feature introduced with Oracle Real Application Clusters 12c Release 2 that considers the workload hosted in the cluster during fencing. In the example of the two node cluster where communication between nodes is interrupted, the node hosting the majority of services (at fencing time) is meant to survive. The idea is to let the majority of work survive, if *everything else is equal*.

A new tag has been created, `CSS_CRITICAL`, that can be set on various levels/components to mark them as “critical” so that the cluster will try to preserve them in case of a failure. When Oracle Clusterware makes a decision about which node to evict in the case of a split brain, `CSS_CRITICAL` will be honored if no other technical reason prohibits survival of the node which has at least one critical component at the time of failure.

The bottom line is that work running in the cluster can be protected so that availability is preserved for critical work. This is one of the key features supporting “better availability” offered by Oracle Real Application Clusters 12c Release 2.

Pluggable Database and Service Isolation

Pluggable Database and Service Isolation was discussed earlier as an improvement for RAC Scalability in a Multitenant environment. But PDB isolation and Service Isolation also improve availability by ensuring that instance failures of instances only hosting singleton PDBs will not impact other instances of the same RAC-based CDB.

After an instance failure, locks that were mastered by the failed instance will need to be re-mastered on another instance in the cluster. In prior RAC releases, instances running on surviving nodes which needed to access data whose locks were mastered on the failed node must pause while the affected locks are re-mastered. With Oracle Real Application Clusters 12c Release 2, locks are mastered locally for singleton PDBs, they will be unaffected by the failure of other instances as there will be no affected locks that must be re-mastered — all are local, and thus none would have been mastered on the failed node. This makes a RAC Database Member Cluster much more resilient and available – instance or node crashes are better isolated and contained and availability for the unaffected resources remains high.

Recovery Buddy

Instance failures are inevitable. When an instance fails, the time to recover from an instance failure in a cluster is determined by the time it takes to detect the failure and the time it takes to recover from the failure (instance recovery). **Failure detection** in a RAC Cluster is nearly instantaneous; but the **failure recovery** depends on the instance recovery time. This time can be significantly reduced by the Recovery Buddy capability available in Oracle Real Application Clusters 12c Release 2.

When an instance fails, blocks that were recently modified by that instance may be in an inconsistent state on disk and must be recovered. A surviving instance must scan the redo log to recover these blocks before new transactions can modify the data. Oracle Real Application Clusters 12c Release 2 Recovery Buddy capability creates a Buddy Instance to track modified data blocks on “buddy” nodes. If a node fails, the Buddy Instance can quickly identify which blocks require recovery, allowing new transactions on unaffected blocks to proceed rapidly and without having to wait for recovery determination.

This feature is especially significant reducing recovery time in clusters where a lot of update activity is taking place.

More Efficient and Intelligent Management

Despite a trend to consolidate workload onto fewer systems, customers are still challenged by having to monitor, manage and maintain an increasing number of systems and databases. A first-level DBA may be assigned as many as 40 databases on multiple systems to manage for availability and performance. A new Cluster Architecture and two new management capabilities have been introduced to significantly improve the manageability of a large estate of Oracle RAC-based clusters. They are:

- Cluster Domain and the Domain Services Cluster
- Oracle's Autonomous Health Framework (AHF) powered by Oracle Machine Learning
- Rapid Home Provisioning

Cluster Domains and Domain Services Clusters (DSC)

Prior to Oracle Real Application Clusters 12c Release 2, a RAC cluster was known as an Oracle Standalone Cluster. Each standalone cluster hosted all its own Oracle Grid Infrastructure services and Oracle ASM locally and required direct access to shared storage. Standalone Clusters are still supported.

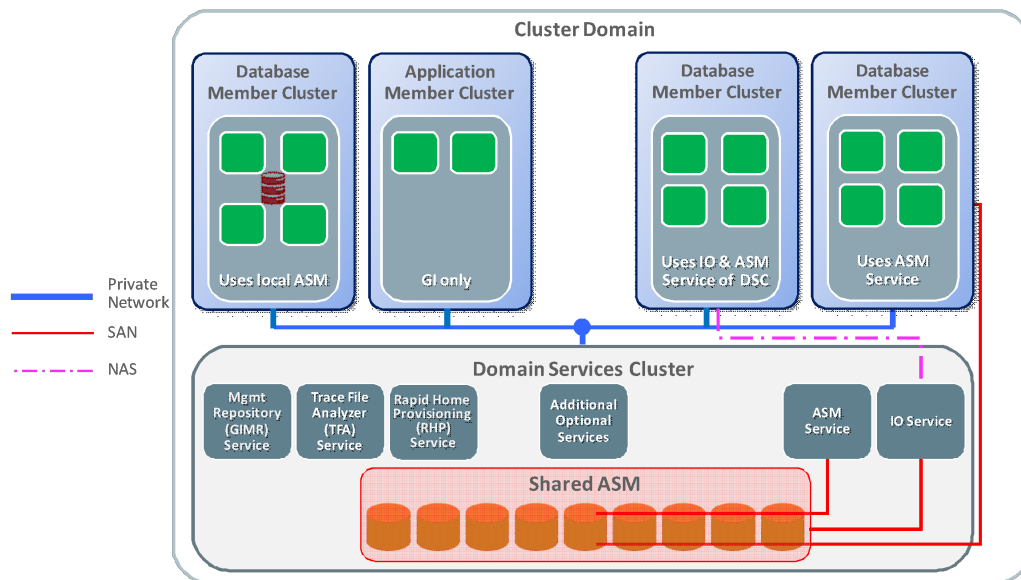



Figure 3: Cluster Domain Architecture

Cluster Domain

In Oracle Real Application Clusters 12c Release 2, a new entity has been introduced called the "Cluster Domain" shown in Figure 3. The Cluster Domain is a new architecture for clusters that centralizes common services across multiple clusters, reducing overhead and streamlining manageability. A Cluster Domain is a collection of Member Clusters all sharing a single Domain Services Cluster (DSC). All clusters use the same standard Oracle Clusterware and are deployed using the Oracle Universal Installer.



The benefit of a Cluster Domain is the more efficient management of the Member Clusters and greater flexibility in how each cluster's storage is provisioned. Another significant efficiency is that the management repository, a database that had to run in each cluster in the past is now centralized in the Domain Services Cluster. Member Clusters no longer have the additional overhead of running and managing multiple Management Repository databases (one for each member cluster) which may mean that fewer processor licenses are needed to maintain the same performance and SLAs. Overall Oracle Database and RAC licensing costs may be less as a result.

This feature will be especially attractive to large IT shops that have hundreds of clusters. For example, a large European telecommunications customer is employing Cluster Domains for the hundreds of clusters deployed in 5 data centers supporting thousands of databases. The Cluster Domain architecture and the centralized services it provides makes managing these large environments more efficient, and improves the reliability and availability of the entire cluster estate due to the reduction in sprawl and lessening of human error that might be introduced when managing many discrete clusters and management databases. With this new Cluster Domain architecture, Oracle is introducing tools to enable standardization while simplifying provisioning, patching, problem diagnosis.

Member Clusters

Member Clusters can be of two types — Oracle Member Clusters for Oracle Databases or Oracle Member Clusters for Applications. Each Member Cluster uses centralized services from a single Oracle Domain Services Cluster (DSC) belonging to that Cluster Domain.

The Oracle Domain Services Cluster enables standardization, centralization, and optimization of an Oracle Real Application Clusters deployment. The DSC is the central point of management for the cluster domain and will host a centralized Grid Infrastructure Management Repository (GIMR) database which collects key statistics and diagnosability data about the individual Member Clusters such as real-time performance data that the Cluster Health Monitor (CHM) collects and metadata that Rapid Home Provisioning requires. In prior versions of Oracle Grid Infrastructure, the Grid Infrastructure Management Repository was a standalone database running in each cluster. Now, the DSC hosts the GIMR which is a container database that runs a separate pluggable database (PDB) for each Member Cluster for which it collects data. An Oracle Member Cluster for Oracle Database always uses the remote Grid Infrastructure Management Repository (GIMR) from its Domain Services Cluster. This is one of the key benefits of the Cluster Domain architecture and Domain Services Cluster. The benefits can be significant!

First, this means that the physical resources necessary to support the GIMR in a cluster are now longer needed since the GIMR is now re-located to a Domain Services Cluster which does not require licensing as long as it only functions as a DSC. That possibly reduces the overall CPU count (and the associated licensing fees) necessary to support the RAC Databases running in the Member Cluster. Second, when problems occur and analysis is required, that work can be done in a more efficient, faster way since much or all of the necessary log and monitoring tool data is now held in a central Management Repository.

Member Clusters use the configured services running on the DSC such as:

- The Autonomous Health Framework (AHF) service for Centralized Monitoring and Management - a capability that is installed and required in all DSCs.
- Two other services that are optionally installed and configured:
 - Centralized ASM services for easier and faster storage consolidation and management. This eliminates the need for ASM and/or shared storage to be available on all the clusters in the domain.

- Rapid Home Provisioning (RHP) Service for better software management and standardization through Oracle Database home creation and patching.

The Database Member Clusters can be configured in three different ways, depending upon how storage is to be accessed. In all cases, the management repositories are stored in the Domain Management Repository. The three configurations for Database Member Cluster storage access are:

- Locally configured shared storage, with local Oracle ASM instances on the database member cluster
- Direct access to storage that is managed through remote Oracle ASM on the DSC (Figure 4)
- Indirect access to Oracle ASM storage through an Oracle ASM IO service on the DSC (Figure 5)

Centralized ASM Services

ASM Services available in Oracle Database 12c Release 2 enable Member Clusters to access data in a centralized ASM Diskgroup environment of the Domain Services Cluster (DSC). When ASM Services in a DSC are used, there are two connectivity options by which Members Clusters can access the shared Diskgroups residing in a DSC. The first is for the Member Cluster to have shared SAN attachment to the DSC storage as shown in Figure 4. The database instances in the Member Cluster rely on ASM instance(s) in the DSC for coordinating access through the SAN.

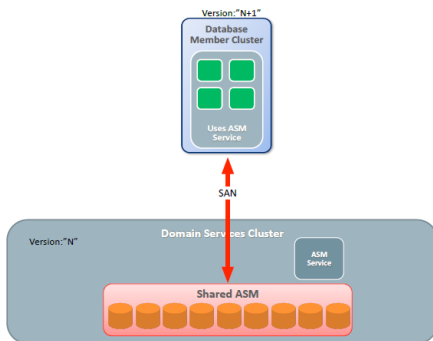


Figure 4: Member Clusters accessing ASM via ASM services in DSC

The second way in which a Member Cluster can access a DSC's Diskgroup is through ASM IO Services using the ASM Storage Network. This is shown in Figure 5.

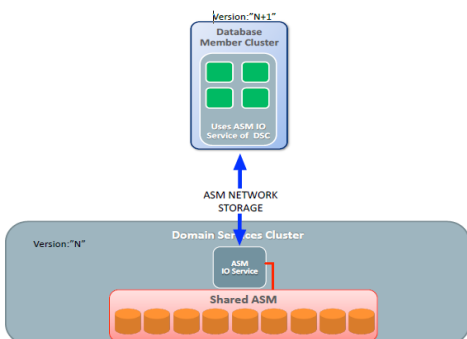



Figure 5: Member Clusters accessing storage via ASM IO service.



With this mode, there is no physical SAN connection of storage between the Member Clusters and the DSC. Member clusters do not need direct connectivity to shared disks. Using the shared ASM service running in the DSC, they can use network connectivity to the IO Service to access a centrally managed pool of storage. Member Clusters access ASM-managed data through a private network. The private network can be the same private network used by Grid Infrastructure or a separate dedicated ASM private network. This model of data access is useful for test and development configurations operating in Member Clusters for accessing production or test data that is shared through the DSC. The ASM Storage Network reduces the cost of storage with the use of network attached rather than more expensive storage attachment, such as Fibre Channel switched networks.

Using a centralized ASM environment operating in a DSC allows Member Clusters to share ASM data and access shared Diskgroups. Customers can quickly and easily create many test clusters without having a specialized or dedicated infrastructure. This eliminates an expensive and time consuming traditional obstacle many shops face of trying to provision a separate shared storage environment for every cluster they create.

Benefits of Cluster Domain Architecture

The Cluster Domain architecture and the Domain Services Cluster enable greater DBA and SysAdmin productivity because they enable rapid creation of test/dev environments and rapid and efficient management of a large cluster estate. For example, a production database managed by ASM in the DSC could be cloned using the ASM shadow copy capabilities. Then, using Rapid Home Provisioning, a new test/dev member cluster and database home could be created. No additional shared storage would be required - the new cloned database could be mounted over the network using the DSC's storage I/O service.

The Cluster Domain architecture with the Domain Services Cluster offers many benefits to IT shops managing a large cluster estate, including:

- Ease of deployment and provisioning of new member clusters
- Configuring shared services once, reusing them many times over
- Eliminating the need to provision and configure shared storage for each member cluster
- Offering flexibility for deploying Database Member Clusters with either shared storage management services (running in the DSC) or locally configured storage
- Enabling more cost effective deployment of Application Member Clusters without requiring support for a database which consumes fewer resources and does not require local shared storage
- Offers efficiencies of scale, such as centralized shared services, for managing and maintaining multiple clusters saving storage space and administrative effort.

Oracle's Autonomous Health Framework (AHF)

Oracle Autonomous Health Framework (AHF) presents the next generation of tools as components of a holistic monitoring and management framework (shown in Figure 6) which autonomously work 24x7 to keep database systems healthy and running while minimizing human reaction time. These components include both existing tools and new tools such as:

- ORAchk
- Cluster Verification Utility
- Trace File Analyzer
- Cluster Health Monitor
- Quality of Service Management, and Memory Guard
- Hang Manager
- Cluster Health Advisor (*new with Oracle Real Application Clusters 12c Release 2*)

Oracle Autonomous Health Framework components work together in daemon mode to address issues faced by Database administrators and System administrators in areas of availability and performance.

Oracle's Autonomous Health Framework (AHF) introduced with Oracle Real Application Clusters 12c Release 2, allows DBAs to focus on more important work by enabling the environment to constantly monitor and analyze all databases in the data center and send out alerts for those databases and systems requiring attention. The AHF works more efficiently in a Domain Services Cluster environment, because continuous analysis workload and data management is removed from the production member clusters and moved to the Domain Services Cluster which hosts a centralized Grid Infrastructure Management Repository as mentioned earlier.

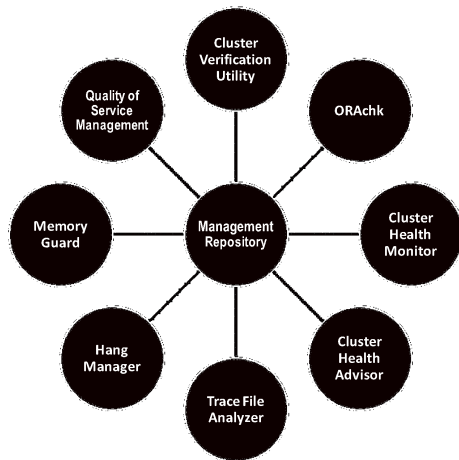



Figure 6: The Autonomous Health Framework consists of a multitude of Tools each created to address a specific aspect of the monitoring and management of a consolidated, clustered database environment

All of these facilities shown in Figure 6 are the tools DBAs and Sys Administrators use to manage a RAC environment. All of these rely on the Grid Infrastructure Management Repository as their data repository for management information. Now, with a centralized Domain Services Cluster that hosts the Grid Infrastructure Management Repository, a more efficient, easier-to-manage environment exists for large scale RAC deployments. Please consult the whitepaper **Oracle Database 12c Release 2 Autonomous Health Framework** for further details on this facility and its tools.

Introducing Cluster Health Advisor Powered by Oracle Machine Learning

New with the Oracle Real Application Cluster 12c Release 2 Autonomous Health Framework is the Cluster Health Advisor (CHA). Via the Machine Learning capability introduced with this tool, CHA is able to “learn” what the “normal” operation of a RAC cluster database environment is and then monitor its operation. Any deviations from “normal” are analyzed using the Machine Learning engine and then, if warranted, get flagged as possible indicators of an impending instance or node failure. DBAs and Sys Administrators can then take preemptive action to prevent the failure or at least manage a graceful soft landing from the failure. There has been tremendous customer interest in this because it helps anticipate and then handle the inevitable database instance or hardware failures that impact database availability.

The Oracle Autonomous Health Framework component, Cluster Health Advisor (CHA), provides system and database administrators with early warning of pending performance issues, root causes, and corrective actions for Oracle RAC databases and cluster nodes. It estimates an expected value of an observed input based on the default



model, which is a trained, calibrated model based on a normal operational period of the target system. Oracle Cluster Health Advisor then performs anomaly detection for each input based on the difference between observed and expected values. If sufficient inputs associated with a specific problem are abnormal, then the Oracle Cluster Health Advisor raises a warning and generates an immediate targeted diagnosis and corrective action.

The Oracle Cluster Health Advisor stores the analysis results, along with diagnosis information, corrective action, and metric evidence for later triage, in the Grid Infrastructure Management Repository (GIMR). Oracle Cluster Health Advisor also sends warning messages to Enterprise Manager Cloud Control using the Oracle Clusterware event notification protocol.

When Grid Infrastructure (GI) is installed for a RAC or RAC One Node database, the Cluster Health Advisor is automatically enabled by default.

By default, Cluster Health Advisor models are designed to be conservative to prevent false warning notifications. However, the default configuration may not be sensitive enough for critical production systems. Therefore, Cluster Health Advisor provides an onsite model calibration capability to use actual production workload data to form the basis of its default setting and increase accuracy and sensitivity of node and database models. Since workloads may vary on specific cluster nodes and Oracle RAC databases, Cluster Health Advisor also provides the capability to create, store, and activate multiple models with their own specific calibration data.

Quality of Service Management – Support for Administrator-Managed Databases

In Oracle Database 11g Release 2, a sophisticated monitoring and management capability was introduced for Real Application Clusters environment – Quality of Service Management. That capability allowed users to establish performance objectives (response time goals) for various workloads running in the same cluster. If goals were not met, the QoS Management facility would recommend changes in the cluster resource assignments so that the performance objective could be regained – moving resources away from less important work to the more important work. But this facility required a “Policy-Managed” database configuration instead of the traditional and default “Administrator-Managed” configuration. The vast majority of Oracle Databases are “Administrator Managed.” With Oracle Database 12c Release 2, Quality of Service Management can now operate on “Administrator Managed” databases and provide recommendations on how cluster resources could be reallocated to bring more important work (as defined by the customer) back to its performance objective.

Cluster Activity Log

The Cluster Activity Log contains information about resource activity across the cluster. The Cluster Activity Log data can be used to track resource behavior in the cluster and to understand chain reactions triggered by the planned or unplanned relocation of individual resources. This centralized log makes searching and analyzing events easier and reduces overhead on member clusters. It answers the question – “What happened – what events caused my cluster to get re-configured and services relocated within the cluster?” The Cluster Activity Log also provides support to the Autonomous Health Framework and for diagnostics.

Rapid Home Provisioning (RHP)

Rapid Home Provisioning is a feature introduced with Oracle Real Application Clusters 12c Release 1 but significantly improved for Release 2. RHP makes it easier than ever to deploy and patch large number of both RAC and single-instance databases and their underlying infrastructure. Software such as the database binaries or their patch sets, need be installed only once, then stored on the RHP server, and from there can be provisioned to any

node or cluster in the data center, whenever needed. This offers significant productivity, efficiency, standardization and security improvements for how Oracle database software is deployed and patched in a large enterprise.

ORACLE REAL APPLICATION CLUSTERS 12C RELEASE 2 KEY FEATURES RECAP

Benefit	Underlying Technology	Description
Improved Scalability	<ul style="list-style-type: none"> Smart Mastering Pluggable Database and Service Isolation 	<ul style="list-style-type: none"> Blocks are mastered on the nodes that most use them Multitenant in a RAC environment is now optimized – recovery times optimized
Even Better Availability	<ul style="list-style-type: none"> Application Continuity Node Weighting Recovery Buddy 	<ul style="list-style-type: none"> Transparent transaction completion in case of a node or instance failure Surviving nodes, in case of a split brain cluster reconfiguration are those running more work and work designated as “critica” Instance recovery is non-disruptive to instances that aren’t mastering affected blocks
Efficient Management	<ul style="list-style-type: none"> Cluster Domains, Domain Services Cluster Autonomous Health Framework <ul style="list-style-type: none"> Cluster Health Advisor Rapid Home Provisioning Cluster Activity Log 	<ul style="list-style-type: none"> Efficiently manages large number of clusters, centralizes services (storage, analytics) Proactively learns about and manages away potential server/database outages Provides centralized software deployment and maintenance Centralizes cluster activity data – makes searching and analyzing events easier - reduces overhead on member clusters

Conclusion

Oracle Real Application Clusters 12c Release 2 provides new capabilities including:

- Better Scalability through Smart Mastering and Pluggable Database and Service Isolation
- Better Availability through new and/or improved facilities such as Application Continuity, Node Weighting and the Recovery Buddy
- More Efficient Management for large scale deployments through facilities such as Cluster Domains and the Domain Services Cluster, Autonomous Health Framework, Rapid Home Provisioning, and a centralized Cluster Activity Log.

For customers that already have invested in a large number of Real Application Clusters deployments, these features will provide significant administrative and operational cost and time savings. For customers contemplating Real Application Clusters for the first time or are looking to further expand their deployments, these features make that first experience easier to implement and manage while delivering even faster and more transparent database high availability.



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 © 2017, 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. 0116

Oracle Real Application Clusters 12c Release 2 – New Features Overview
March 2017



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