



An Oracle White Paper
June 2015

Oracle Clusterware 12c

Executive Overview	1
Oracle Clusterware – Overview	2
New In Oracle Clusterware 12c	3
Business Continuity and High Availability	3
Scalability and Agility	4
Cost-effective Workload Management	5
Standardized Deployment and System Management	7
Summary and Conclusion.....	8

Executive Overview

Industry adoption of cloud computing and consolidation imposes greater demands on system and application performance and availability. Disparate workloads compete for shared resources while attempting to maintain service level agreements (SLAs). Oracle Clusterware began to address consolidation and management of resource constraints in early releases of 11g introducing Quality of Service, server pool partitioning and policy management. Oracle Clusterware 12c continues to provide innovative solutions for consolidated environments. Through efficient, agile, and secure online resource management in innovative cluster topologies, Oracle Clusterware 12c insures business critical work is satisfied within defined performance and availability SLAs.

Oracle Clusterware is the preeminent clustering solution for Oracle's engineered systems. It is the established solution for database services and, increasingly, the solution for Oracle and non-Oracle application services. This combination of clustering for database and application services creates the foundation for end-to-end automated high availability (HA) and workload management for consolidated and cloud deployments. Oracle Clusterware offers the means to manage resources in all tiers, simultaneously, thereby providing the most comprehensive, HA solution available today.

This paper describes the benefits of Oracle Clusterware consistent with Oracle's overall strategy and discusses enhancements in the latest version of the Oracle Clusterware. These key benefits of Oracle Clusterware which directly contribute to business success are:

- Business Continuity and High Availability
- Scalability and Agility
- Cost-effective Workload Management
- Standardized Deployment and System Management

Oracle Clusterware – Overview

Oracle Clusterware

Oracle Clusterware is the technology which unifies servers in a server farm to form a cluster. A cluster is defined as a group of independent, yet inter-connected servers which cooperate as a single system. Oracle Clusterware is the software in this environment which coordinates the management of these servers and the applications running on them. Oracle Clusterware is a generic clustering solution tested and certified to run on all major platforms and operating systems.

Oracle Clusterware was originally introduced with Oracle Database 10g Release 1 as the underlying, generic clustering software required for running the multi-instance RDBMS, Oracle Real Application Clusters (RAC). Oracle Clusterware is shipped with Oracle Grid Infrastructure, which is a suite of software packages which includes Oracle Automatic Storage Management (ASM) for databases, Oracle Automatic Storage Cloud File System (ACFS), a POSIX compliant general purpose cluster file system, Oracle Database Quality of Service Management. Oracle Clusterware is a generic, general purpose clustering solution for all applications.

Oracle Clusterware is a complete, enterprise class clustering solution. It is a free. To get support for upgrades or patches, a valid support license for any other Oracle product such as Oracle Linux, Oracle Solaris, or any other supported Oracle application is required.

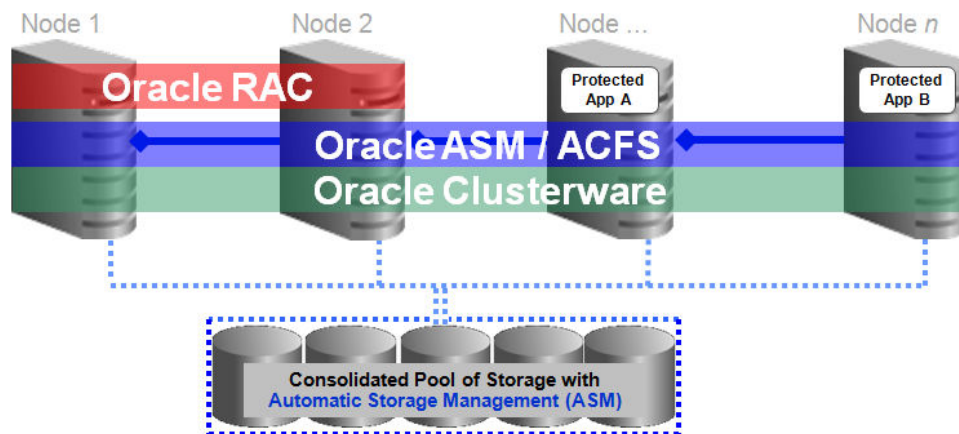


Figure 1: Oracle Clusterware 11g deployment model requiring shared storage and shared communication

Oracle Stack Management

Oracle Clusterware provides a rich management framework to integrate and manage all applications. The core of this framework is a robust policy engine for resource placement as well as start and stop dependencies between resources. This enables Oracle Clusterware to manage workload and availability through the entire stack. The ability to automate and securely manage the start, stop and placement of database and application resources for planned maintenance and workload performance is an essential ingredient for successful consolidation.

New In Oracle Clusterware 12c

Oracle Flex Cluster 12c

Oracle Clusterware, in the original design, is a tightly coupled clustering solution which requires shared storage and shared network between all servers in the cluster. This is a requirement when Oracle Clusterware is installed for Oracle databases - RAC and single instance. This tightly coupled cluster is required to preserve data integrity of multi-instance, shared cache clustered databases. Oracle 12c introduces an innovative clustering solution - **Oracle Flex Cluster** which relaxes the tightly coupled nature of Oracle Clusterware and permits the collocation of applications and the RDBMS in a common Oracle Flex Cluster. In the Oracle Flex Cluster, the database tier, composed of 'HUB' servers, preserves the tightly coupled configuration, while the application tier, composed of 'LEAF' servers, is loosely coupled and does not require shared storage and peer-to-peer communication with other servers in the cluster. This innovative tightly coupled/loosely coupled clustering solution provides far greater scalability without compromising the availability and resource management expected in high performance, and very large scale consolidations.

Business Continuity and High Availability

Oracle Flex Clusters extends the availability and management features well known in Oracle Real Application Clusters to the application tier. In the initial 12c release, Oracle Flex Cluster is designed for collocation of applications and the RDBMS in a common cluster using the HUB-LEAF deployment model described above.

The collocation of applications and the database in a common cluster offers far greater availability and ease of management between application components and the database, where start, stop ordered dependencies, and placement may be defined and automated. Specifically, applications may depend on components outside of the respective application management domain. Customers either execute the start up of these various components manually, or, they maintain cumbersome scripts in an attempt to automate. For example, a PeopleSoft application may have a WebLogic component and a Tuxedo component, each of which are managed within their respective domains. Oracle Flex Cluster allows you to manage these components collectively while maintaining secure, isolation management between the components. And if the database is not up, the dependency model can be extended between the application and database tiers allowing the application to pull up the database or, if the application tier is down, the database startup would start the application components on the correct servers in the correct order. The entire process, application tier start up with correct startup order and placement and database/application tier pull-up - is fully automated.

Oracle Flex Cluster is compatible with the Oracle Clusterware Agents. These Agents, developed by Oracle Clusterware engineering and the respective application engineering groups, are fully integrated to provide high availability (HA) and online relocation services for applications running in the Oracle Flex Cluster. The dependencies described above are implicitly configured by Oracle and can extend between the application and the database. The Agents allow for applications to define start and stop dependencies on database services or the databases directly.

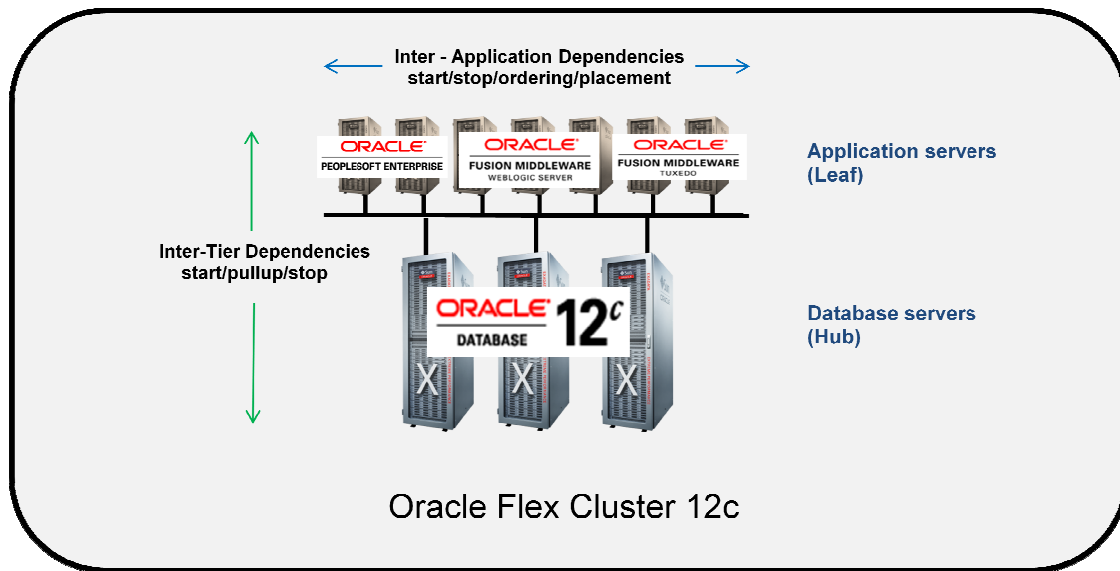


Figure 2: Oracle Flex Cluster 12c, with Database and Application Collocation and Dependencies

Figure 2 above illustrates the Oracle Flex Cluster topology where the Application servers are loosely coupled with no shared storage or peer-to-peer communication requirements with all nodes in the cluster. A LEAF node has a loose communication association with a single HUB node in the cluster.

Oracle Flex Cluster is the innovative clustering solution to preserve and protect business critical functions and insures continuous, automated availability for application and database services.

Scalability and Agility

Oracle VM and Bare Metal

As noted earlier, Oracle Flex Cluster is designed for optimal scalability where shared storage and peer-to-peer communication is restricted to the HUB tier. This design supports very large scale deployments of 100s of nodes. Another innovation introduced with Oracle Flex Cluster 12c is the ability to deploy **Oracle VM guests and Bare Metal servers in the same cluster**. Earlier releases of Oracle Clusterware support Oracle Clusterware deployments in Oracle VM environments for Oracle RAC and applications, however, all servers in the cluster are expected to be Oracle VM guest. Oracle Flex Cluster, however, allows the application/LEAF servers to be deployed as Oracle VMs while the HUB servers are Bare Metal. This deployment model allows for rapid application provisioning while preserving HUB tier performance requirements. In addition, the Oracle Flex Cluster HA and policy management capabilities are complimentary to Oracle VM. Oracle Flex Cluster Agents rapidly detect VM hangs or infrastructure failures and provide hot standby failover of the application to another VM in the cluster. The start/stop dependencies, ordering and placement driven through the Oracle Flex Cluster policy engine guarantees ordered start and stop of application components and a rich placement configuration for attraction, dispersion and hard, weak, pull-up and exclusion ordering for Bare Metal and Oracle VM servers in the Oracle Flex Cluster.

The online resource relocation service of Oracle Flex Cluster can online relocate applications or application components running inside Oracle VM guests to a hot standby Oracle VM guest in the cluster. This functionality offers fine-grain control of application and resource management within a VM and compliments the course-grained control of Oracle VM Live Migration.

The complimentary power of Oracle VM and Oracle Flex Cluster allow for very large scale deployment and agile resource management of applications running in Oracle VM guests and Bare Metal servers.

Cost-effective Workload Management

Oracle Clusterware has long provided functionality to efficiently manage concurrent workloads. Online workload management has become more important as customers deploy dense, multitenant configurations competing for shared resources. Oracle Clusterware introduced server pool partitioning in 11g release 2, where servers in the cluster may be logically partitioned for resource placement and to manage online server allocation. Server pools may be re-partitioned online to accommodate fluctuating demand for critical workloads. Originally, server pools were restricted to a set of basic attributes characterizing servers as belonging to a pool, with no way to distinguish between types of servers or specific compute attributes of servers within a pool. All servers were considered to be equal in relation to their processors, physical memory, and other characteristics.

Server Categorization

Oracle Clusterware 12c introduces **Server Categories**, where Oracle Clusterware will dynamically discover the computing attributes of a server, such as processor types, memory and other distinguishing system features and allow users to organize servers into server pools based on these categories. When `SERVER_CATEGORY` is defined in place of `SERVER_NAMES`, server membership in a pool is restricted to servers with attributes that satisfy the category. Users can define their own characteristic of servers for categorization in addition to the discoverable attributes. Oracle Flex Cluster 12c provides two default named categories, HUB and LEAF for the specific server roles configured. Use of server categories provides an efficient, automated mechanism to allocate compute resources to server pools based on pre-defined server attributes.

Cluster Configuration Policy Sets

The use of server pools and server categorization requires the cluster to be configured as a policy-managed cluster. Prior to Oracle Clusterware 12c, server pool changes were manually executed by the end-user based on fluctuating demand or as a recommendation from the Oracle Database Quality of Service Management framework (QoS Management). Oracle Clusterware 12c allows users to define a cluster configuration policy for each server pool defined in the system. These policies collectively represent the cluster configuration policy set. An administrator can set different policies to be in effect at different dates or times of the day in accordance with business needs and the system demands. The policy in effect drives the cluster topology as a whole based on server pool requirements for the target workload. Cluster configuration policy sets provide dynamic management of cluster polices across the cluster.

Cluster configuration policy sets and policy based management:

- Enable online server reallocation to satisfy workload capacity requirements
- Guarantees the allocation of required resource for critical work as defined by the policy
- Ensures isolation where necessary, so that dedicated servers are provisioned in the cluster for applications and databases
- Enables policies to be configured to change pools in accordance with changing business needs or application demand

Oracle Database Quality of Service Management

Oracle Database QoS Management¹ is a full Oracle stack development effort introduced in 11gR2 to provide effective runtime database management of datacenter SLAs by ensuring when there are sufficient compute and memory resources to meet all objectives these resources are properly allocated. Should demand or server availability change where capacity is over-subscribed, Oracle Database QoS Management insure that the most business critical SLAs are preserved at the cost of less critical work. Oracle Clusterware 12c cluster configuration policy sets are Oracle Database QoS Management aware. If you create policies with Oracle Database QoS Management the Oracle Clusterware 12c cluster configuration policy sets are disabled in deference to the Oracle Database QoS Management server pool directive overrides.

Predictability

An important element in consolidated environments, where resources and resource use fluctuate on demand, is the ability to predict the placement of resources after online relocations or resource faults. Understanding the impact of changes to the environment is important. Oracle Clusterware 12c provides a set of evaluation commands to determine the impact of certain operations before the respective operation is executed. Administrators can use the Oracle Clusterware 12c What-If command to predict Oracle Clusterware's response to a hypothetical planned or unplanned event. Oracle Clusterware 12c has the ability to react to events in the system and produce a response action plan. This action plan consists of a series of resource state transitions and/or server pool reconfigurations. What-If response plans are available for the following event categories:

- Resources: start, stop, relocate, add, and modify
- Server Pools: add, remove, and modify
- Servers: add, remove and relocate
- Policies: change active policy
- Server Category: modify

¹ For more information about Quality of Service (QoS) Management, see:
<http://www.oracle.com/technetwork/products/clustering/overview/qosmanageent-508184.html>

Oracle Clusterware 12c What-If action plans help administrators determine failover capacity for planned or unplanned events, confirm correct dependency definitions for configured resources and validate policy set deployments. It is an essential tool in automated workload management.

Standardized Deployment and System Management

The key management strength of Oracle Clusterware is the efficient management of workload in the cluster. Oracle Clusterware supports two types of workload management configurations:

1. Policy-managed which allows for flexible management by abstracting resource allocations (e.g. dynamically assigning resources to run in pools of servers in the cluster)
2. Administrator-managed which is the traditional means of managing resources where resources are statically assigned to servers in the cluster and which is an inflexible, manual management configuration

Oracle Clusterware supports both deployment configurations simultaneously; however, servers allocated to Policy Management are not available to be Administrator-managed. Support for both deployment configurations facilitates the ease of migration from Administrator-managed to Policy-managed providing optimal flexibility required for consolidation.

Policy Management and server pools offer a standardized deployment model for database and application clustering. As stated, Oracle Flex Cluster 12c provides HUB server pools for the database tiers and LEAF server pools for applications. Oracle Clusterware has a common command line interface to manage resource in both tiers either by a single GRID Administrator or, by secure, role-separated administrators for each tier. The interface is the standard Oracle Clusterware interface for all cluster resources.

Oracle Clusterware is fully integrated with Oracle Cloud Control. Oracle Clusterware configured resources - cluster, local, network, application, database, VM and storage are all exposed as Enterprise Managed targets which provides a single management 'pane of glass' for all tiers managed by Oracle Clusterware. All Oracle Clusterware management operations are available through the Oracle Cloud Control GUI.

Summary and Conclusion

Business continuity, high availability, scalability, flexibility and agility combined with stress-free management are the pillars of successful IT infrastructure, consolidation and cloud deployments. Oracle Clusterware is the solution of choice for thousands of Oracle customers for their business critical applications for more than a decade. Oracle Clusterware is a cost-effective, enterprise class high availability and resource management framework for all applications.

Oracle Clusterware 12c builds on the success of earlier releases by providing significant, innovative enhancements critical to sustaining your business success. New solutions such as Oracle Flex Cluster 12c, Server Categorization and Cluster Configuration Policy Sets all focus on flexibility, scalability and ease of management without compromising on continuous application availability.



Oracle Clusterware
June 2015

Authors: Ian Cookson
Contributing Authors:
Markus Michalewicz

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 © 2015, 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. 0113

Hardware and Software, Engineered to Work Together