

COMPARING ORACLE REAL APPLICATION CLUSTERS TO FAILOVER CLUSTERING FOR ORACLE DATABASE

Customers looking to remove the server as a single point of failure for their database applications often implement a clustering solution. The cluster architecture is used by both Oracle Real Application Clusters (RAC) and Failover Clusters and this document assumes the reader understands what a cluster is. Oracle Real Application Clusters is a unique solution that provides flexible scalability and high availability for database applications. This solution is superior to the failover clustering that is available on most platforms today. Failover clustering is a solution that provides higher availability than a single server but the availability service levels are less than what you can achieve with RAC and you do not have flexible scalability or utilization of hardware resources that you can have with RAC. A failover cluster is an active-passive solution that provides high availability for database applications by monitoring for failures and restarting the database.

FAILOVER CLUSTERING FOR HIGH AVAILABILITY

Failover Clustering is a solution provided by cluster software vendors today including the Oracle Clusterware¹. The typical implementation of a failover cluster is a 2-node cluster where each node runs one or more Oracle Databases. The database runs on only one node at any time. In order not to have the second node completely idle waiting for a failure, customers often run other applications or databases on this node. The Oracle Database is under the control of the cluster software such that it is automatically started, stopped, and monitored by the cluster software. If the database instance fails, then the cluster software will recognize the failure and restart the instance often without operator knowledge. If the node fails, then the storage, database instance and any other dependent processes are failed over to the other node in the cluster and restarted. The application experiences down time for the period of time it takes the cluster to recognize the failure, reconfigure the cluster, complete the failover. Depending on the size and complexity of the environment, this can take anywhere from 1-2 minutes to 20 or 30 minutes or more. If the second node (the node to which you failed over to) is already running work, this node may become overloaded and the service level degraded for the period of the outage. To provide complete coverage for a server outage, you must buy twice the resource requirements of your applications.

ORACLE REAL APPLICATION CLUSTERS

A RAC database is a clustered database. A cluster is a group of independent servers that cooperate as a single system. Clusters provide improved fault resilience and modular incremental system growth over single symmetric multi-processor (SMP) systems. In the event of a system failure, clustering ensures high availability to users. Access to mission critical data is not lost. Redundant hardware components, such as additional nodes, interconnects, and disks, allow the cluster to provide high availability. Such redundant hardware architectures avoid single points-of-failure and provide exceptional fault resilience. A clustered database is a single database that can be accessed by multiple instances concurrently. Each instance runs on a separate server in the cluster. The advantages and the real intelligence of a cluster used within a RAC environment comes in with the database software itself.

If the database instance fails, you have not lost access to the database. The application does not experience an outage. Only a subset of users (those connected to the failed instance) is affected by the failure. This outage can be easily masked to the end user using the advanced features provided with RAC (Fast Application Notification) and the Oracle clients (Fast Connection Failover). The instance failure is recognized by another instance in the cluster and recovery automatically takes place. The Oracle Clusterware monitors the instance and automatically restarts it when a failure occurs.

Oracle Clusterware is included with Oracle Database 10g. Oracle Clusterware is a portable clusterware that is tightly integrated with Oracle RAC to provide a complete solution for your database application. Oracle RAC 10g supports up to 100 nodes in the cluster and up to 100 instances in the RAC database. The management and monitoring of the

¹ See Using Oracle Clusterware to Protect a Single Instance Oracle Database (PDF)
<http://www.oracle.com/technology/products/database/clustering/pdf/ocsingleinstance.pdf>



Oracle resources (VIP, listener, database, services) in the cluster are included out of the box. An API is included for customers to add additional processes to be managed by Oracle Clusterware to keep them highly available².

Along with the higher availability, RAC provides you with flexible scalability. Instead of having to size a single server to support your application load, RAC allows you to spread the load across multiple servers that have been clustered together. When additional resources are required, additional nodes and instances can be easily added to the cluster with no downtime. A RAC database can scale to 100 instances. Applications do not have to be changed to use RAC.

RAC scalability allows you to take advantage of the cost savings of using smaller servers clustered together to provide the resource requirements of your application. A larger cluster of smaller servers reduces the impact of a server failure. If there are 2 nodes in a cluster and one fails, you have lost 50% of your resources and 50% of your users are impacted. If you lose 1 node in a cluster of 10 nodes, you have only lost 10% of your processing power and only 10% of your users are affected. Also to provide the same level of resource to the application during a failure, you only need 10% additional resources.

CONCLUSION

Oracle Real Application Clusters has been designed for high availability and scalability. By providing protection from hardware and software failures, Oracle Real Application Clusters provides systems availability ensuring continuous data access. Its scale out and scale up features offer a platform, which can grow in any direction allowing enterprises to grow their businesses. Existing applications as well as newly developed applications benefit from the transparency Oracle Real Application Clusters provides. Application development as well as administration and change management thus become much easier allowing reduction in total cost of ownership. Oracle Real Application Clusters is unique to the market with its offering and capabilities. RAC is used by thousands of customers worldwide in all industries in mission critical and many other application environments.

For more information on using RAC: otn.oracle.com/rac

Author: Barb Lundbild, Oracle Corporation

² See Using Oracle Clusterware to Protect Third Party Applications –

http://www.oracle.com/technology/products/database/clustering/pdf/twp_oracleclusterware3rdparty%5B1%5D.pdf