ISSN:0975-9646

An Extensive Survey of Monitoring of Load Balance in Oracle Real Application Cluster

Neha Chandrima[1], Sunil Phulre[2] Dr. Vineet Richhariya[3]

[1]*M.Tech. Scholar, Dept. of Computer Sc. & Engineering, LNCT, Bhopal (M.P.) India -462022*
[2] *Asst. Professor, Dept. of Computer Science and Engineering, LNCT, Bhopal (M.P.) India -462022*
[3] *Professor and Head, Dept. of Computer Science and Engineering, LNCT, Bhopal (M.P.) India -462022*

Abstract: -Oracle real application cluster also called Oracle RAC offers and provides very attractive and promising features for today's challenging market scenario, where safety and availability of database is the basic need and demand of clients who are working on mission critical databases. The two most alluring features which RAC provides are high availability and load balancing. Load balancing is performed by the load balancer which most of the time works as desired; But it is noticed that failure of load balancer could be a single point of failure for the entire RAC system. In the present work our aim is to have comprehensive study of literature available in the field of monitoring the load balance in Oracle RAC in order to keep a close watch on whether switchover between nodes is happening correctly and smoothly or not.

Many Researchers and developers working in this field suggested various monitoring systems for load balancing in Oracle RAC, yet there is scope to get most efficient and accurate one. The concept of relative entropy is taken as inspiration from this extensive survey of the literatures available, which will help us in comparing the randomness of the each node and also the entire RAC system. Thus entropy based monitoring system has projected as efficient system in present review work.

Keywords: Oracle RAC, Databases, Relative Entropy, Load Balancing, Randomness.

1. INTRODUCTION

Oracle RAC is one of the important clustered oracle databases and uses oracle clusterware software for the infrastructure to bind multiple servers so that they operate as a single system. A cluster comprises of multiple interconnected computers or servers that appear to be one server to end users and applications.[1] It is Oracle database option that provides a single system image for multiple servers to access one Oracle database. Oracle Cluster ware is a cluster management solution that is integrated with Oracle database. It is also a compulsorily required component for using Oracle RAC. It also enables both single-instance Oracle databases and Oracle RAC databases to use the Oracle high-availability infrastructure. Oracle clusterware assists us to create a clustered pool of storage to be used by any combination of single-instance and Oracle RAC databases. Oracle RAC databases differ architecturally from single-instance Oracle databases in that each Oracle RAC database instance also has:

- At least one additional thread of redo for each instance
- An instance-specific undo table space

Figure 1.1 depicts flow diagram of Oracle database with Oracle RAC architecture.

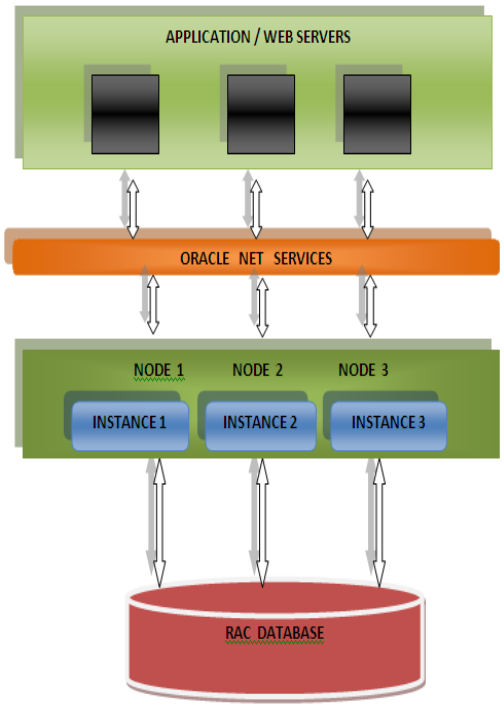

Fig 1.1 Oracle Database with Oracle RAC Architecture

Oracle RAC provides high availability and scalability for all application types. The Oracle RAC infrastructure is also a key component for implementing the Oracle enterprise grid computing architecture. Multiple instances accessing a single database, prevents the server from being a single point of failure. Oracle RAC enables us to combine smaller commodity servers into a cluster to create scalable environments that support mission critical business applications.

2. ORACLE CLUSTERWARE FOR ORACLE REAL APPLICATION CLUSTERS

Oracle Clusterware provides a complete, integrated cluster ware management solution on all Oracle Database platforms. This cluster ware functionality provides all of the features required to manage our cluster database including node membership, group services, global resource management, and high availability functions. We can install Oracle Cluster ware independently or as a prerequisite to the Oracle RAC installation process. Oracle Database features such as services use the underlying Oracle Clusterware mechanisms to provide their capabilities. Oracle Database also continues to support select third-party clusterware products on specified platforms. Oracle Clusterware is designed

for, and tightly integrated with, Oracle RAC. When we create an Oracle RAC database using any of the management tools, the database is registered with and managed by Oracle Clusterware, along with the other Oracle Database processes such as Virtual Internet Protocol (VIP) address, Global Services Daemon (GSD), the Oracle Notification Service (ONS), and the Oracle Net listeners.[2]. These resources are automatically started when Oracle Clusterware starts the node and automatically restarted if they fail. We can use Oracle Clusterware to manage high-availability operations in a cluster. Anything that Oracle Clusterware manages is known as a CRS resource, which could be a database, an instance, a service, a listener, a VIP address, an application process, and so on. Oracle Clusterware manages CRS resources based on the resource's configuration information that is stored in the Oracle Cluster Registry (OCR). Oracle Clusterware stores the information that describes the configuration of these components in the OCR.

3. ORACLE REAL APPLICATION CLUSTERS ARCHITECTURE AND PROCESSING

Oracle RAC requires Oracle Clusterware to provide the cluster infrastructure that allows multiple servers to work together. Oracle Clusterware provides group membership, communications infrastructure, event monitoring, and a high availability framework.

The following sections describe these concepts in more detail:

- Understanding Cluster-Aware Storage Solutions
- Overview of Connecting to Oracle Database Using Services and VIP Addresses
- About Oracle Real Application Clusters Software Components
- About Oracle Real Application Clusters Background Processes

3.1 Clusterware Storage Solutions

An Oracle RAC database is a shared everything database. all data files, control files, SPFILEs, and redo log files in Oracle RAC environments must reside on cluster-aware shared disks so that all of the cluster database instances can access these storage components. All database instances must use the same interconnect, which can also be used by Oracle Cluster ware. Because Oracle RAC databases use a shared everything architecture, Oracle RAC requires cluster-aware storage for all database files

3.2 Connecting to Oracle Database Using Services and VIP Addresses

All nodes in an Oracle RAC environment must connect to a Local Area Network (LAN) to enable users and applications to access the database. Applications should use the Oracle Database services feature to connect to an Oracle database. Services enable us to define rules and characteristics to control how users and applications connect to database instances. These characteristics include a unique name, workload balancing and failover options, and high availability characteristics. Oracle Net Services enable the load balancing of application connections across all of the instances in an Oracle RAC database.

3.3 Oracle Real Application Clusters Software Components

Oracle RAC databases have two or more database instances that each contain memory structures and background processes. An Oracle RAC database has the same processes and memory structures as a single-instance Oracle database as well as additional process and memory structures that are specific to Oracle RAC. Any one instance's database view is nearly identical to any other instance's view in the same Oracle RAC database; the view is a single system image of the environment.

3.4 Oracle Real Application Clusters Background Processes

The Oracle RAC processes and their identifiers are as follows:

- ACMS—Atomic Controlfile to Memory Service (ACMS)
  In an Oracle RAC environment, the atomic controlfile to memory service (ACMS) per-instance process is an agent that contributes to ensuring a distributed SGA memory update is either globally committed on success or globally aborted in the event of a failure.
- GTX0-j—Global Transaction Process
  The GTX0-j process provides transparent support for XA global transactions in a RAC environment. The database auto-tunes the number of these processes based on the workload of XA global transactions.
- LMON—Global Enqueue Service Monitor
  The LMON process monitors global enqueues and resources across the cluster and performs global enqueue recovery operations.
- LMD—Global Enqueue Service Daemon
  The LMD process manages incoming remote resource requests within each instance.
- LMS—Global Cache Service Process
  The LMS process maintains records of the datafile statuses and each cached block by recording information in a Global Resource Directory (GRD).

4. AUTOMATIC WORKLOAD MANAGEMENT

Automatic workload management enables us to manage the distribution of workloads to provide optimal performance for users and applications. This includes providing the highest availability for database connections, rapid failure recovery, and balancing workloads optimally across the active configuration. Oracle Database with Oracle RAC includes many features that can enhance workload management, such as connection load balancing, fast connection failover, the load balancing advisory, and runtime connection load balancing. Automatic workload management provides the greatest benefits to Oracle RAC environments.Automatic workload management comprises the following components:

- High Availability Framework: The Oracle RAC high availability framework enables Oracle Clusterware to maintain components in a running state at all times. Oracle high availability implies that Oracle Clusterware monitors and restarts critical components if they stop, unless we override the restart processing. The high availability framework also provides alerts to clients when configurations change. This enables clients to immediately react to the changes, enabling application developers to hide outages and reconfigurations from end users.
- Load Balancing Advisory: This is the ability of the database to provide information to applications about the current service levels being provided by the database and its instances. Applications can take advantage of this information to direct connection requests to the instance that will provide the application request with the best service quality to complete the application's processing.
- Services: Cluster-managed services provide a powerful automatic workload management facility that enables the enterprise grid vision. Services are entities that we can define in Oracle RAC databases.
- Connection Load Balancing: Oracle Net Services provides connection load balancing for database connections. Connection load balancing occurs when the connection is created. Connections for a given service are balanced across all of the running instances that offer the service using information from the load balancing advisory. We should define how we want connections to be balanced in the service definition.

5.REVIEW OF LITERATURE AVAILABLE

The RAC solution, some aspects of its implementation, its brand name, and how it is sold have evolved in the decade since the technology was originally conceived as Oracle Parallel Server. Oracle 10g RAC is aimed squarely at solving two of the most important problems facing those deploying data-centric applications: 1) application availability and 2) affordable performance. These requirements are important to data management customers, and consequently, Oracle and the industry as a whole offer features to address these requirements based on application design, requirements, and the deployment scenario. In this concern, several researchers have been working in the field of Oracle RAC and gave important findings through their research carried out. Some significant contributions of the researchers are studied and compared here and their shortcoming is projected as the problem statement of the proposed research work. Some of these are discussed here: Gong Weihua and Wang Yuanzhen [3] reported that in heterogeneous database cluster system, performance is very much close to the computing capabilities of nodes and their different categories of workloads. They introduced an innovative method to determine a load status of nodes by the weighted load values with consideration of utility of different resources and workload types in load balancer. For this, an efficient and dynamic scheme was proposed for OTLP (Online transaction processing) workloads to maximize the utility of distributed resources and better performance was achieved.

Stephane Gancarski [4] proposed a new innovative solution to balance load of autonomous applications and oracle databases. This proposed solution is similar to Distributed shared memory in that it also provides shared address space to application with distributed and replicated databases. In this research work main focus was to allow the system administrator for controlling the tradeoff of performance and consistency during of applications and databases onto nodes cluster. Jin chen et al [5] developed a method to detect outlier for fine grained load balancing in database clusters. Their work is focused towards reduction of cost of ownership in large data centers and emphasized the need for database system method for automatic performance tuning and efficient resource usage. Automatic provisioning of database servers to applications and virtualization techniques, like, live virtual machine migration was proposed as useful tools to address this problem. They applied a first step towards accurate diagnosis and selective retuning for shared database clusters. A fine-grained adjustment of resource allocation to applications, which may be more appropriate for the most common load-balancing needs, was suggested. They observed that the above mentioned shared hosting environment has a natural hierarchy in terms of both system and application structure.

Deepali Kadam et al. [6] worked on automatic failover and load balancing for Oracle real Application Clusters. RAC enables us to use clustered hardware by running multiple instances against the same database. The database files are stored on disks that are either physically or logically connected to each node, so that every active instance can read from or write to them. Oracle Real Application Clusters manages data access, so that changes are coordinated between the instances and each instance sees a consistent image of the database.

Howard Karlof and Kenneth Shirley [7] reported the concept of maximum entropy summary trees used in databases. They suggested that the best choice of which summary tree to use among those with a fixed number of nodes is the one that maximizes the information-theoretic entropy of a natural probability distribution associated with the summary tree, and also provided a pseudopolynomial-time dynamic-programming algorithm to compute this maximum entropy summary tree, when the weights are integral. The result is an automated way to summarize large trees and retain as much information about them as possible, while using and displaying only a fraction of the original node set. John McHugh and Markus Michalewicz [8] discussed various Oracle RAC deployment scenarios and provides best practices for an optimized Oracle RAC deployment in Oracle VM environments for Oracle VM versions 2.1.2 through Oracle VM 3.0.3. Version specific information is noted accordingly. Gia-Khanh Nguyen and Tim Read [9] discussed about Real Application Clusters On Oracle Solaris Zone Clusters. They explained how to deploy, Oracle RAC on Oracle Solaris Zone Clusters.

James Hamilton [10] shared disk cluster database management systems such as Oracle RAC are being discussed as a potential solution to the application scaling and robustness problem. This paper argues that the best solutions for availability have no single points of failure and support geo-clustering. RAC, with millions of lines of shared software between the DBMS and the disk that offer many single points of failure, is less suitable as an availability solution and is better used as a multinode scale-out solution.

René Kundersma [11] reported upgrading to Oracle RAC and discussed about 11gR2 Oracle Real Application Clusters. He stated that, Oracle Database 10g and Oracle Database 11g are supported with Oracle Clusterware 11g Release 2. We must upgrade Oracle Clusterware and Automatic Storage Management to 11.2.0.1 at the same time. Apply DBCA patch Bug 8288940 and then Pin the node. Markus Michalewic [12] worked on scability and agility of Oracle RAC. He reported that business continuity, high availability, scalability, flexibility and agility combined with automatic management are the pillars of successful IT infrastructure and cloud deployments. The Oracle Database with the Real Application Clusters (RAC) option has been the solution of choice for thousands of Oracle customers for their business critical database systems for more than a decade.

6. EXISTING GAPS IN LITERATURE REVIEW

An exhaustive survey reveals that there are still.scope to identify most efficient monitoring system in oracle RAC for load balancing. Entropy based monitoring system may be emerged out and developed for the same purpose.

7. CONCLUSIONS

The concept of relative entropy is taken as inspiration, from this extensive survey work, which will help us in comparing the randomness of the each node and also the entire RAC system. In this concept, a script should be designed to monitor the number of sessions in each node which should run at the expected peak hours. A threshold limit for number of sessions in each respective node is to be defined in prior, which should be near about but less than maximum session limit defined for the load balancer. As soon as the number of sessions in a particular node reaches its threshold limit, an email is sent to the DBA team which in turn alerts the team to keep a close watch on the switchover.

REFERENCES

1. docs.oracle.com/cd/E11882_01/rac.112/e41960/admcon.htm
2. docs.oracle.com/cd/B28359_01/rac.111/b28254/admcon.htm
3. Gong Weihua and Wang Yuanzhen "A new load balancing scheme on heterogeneous database cluster" Geo-spatial information science Jan, 2006; 9(3):216-222. DOI:10.1007/BF02826771
4. Stephane Ganc¸ Arski and Hubert Naacke "Load Balancing of Autonomous Applications and Databases in a Cluster System" Distributed data and structure 4, 162-164 -169.
5. Zin Chen, Gokul Soundrajan,Madlin Mihaailescu **"**Outlier Detection for Fine-grained Load Balancing in Database Clusters"ProceedingICDEW 07, Proceedings of 2007 IEEE International Conference on Data Engineering workshop, IEEE Computer Society Washington, DC, USA, 2007, 404-413

6. Deepali Kadam, Nandan Bhalwarkar, Rahul Neware, Rajesh Sapkale, Raunika Lamge "Oracle Real Application Clusters" Int. J of Sc. & Engg Research , 2, (6), June-2011, 01-05, ISSN 2229-5518
7. Howard Karlof and Kenneth Shirley "Maximum Entropy Summary Trees" Eurographics Conference on Visualization (EuroVis) 2013, Vol 32, No. 3
8. John McHugh and Markus Michalewicz "Oracle Real Application Clusters in Oracle VM Environments", An Oracle Technical White Paper March 2012
9. Gia- Khanh Nguyen and Tim Read "Running oracle® real application clusters on oracle solaris zone clusters" An Oracle White Paper July 2011
10. James Hamilton "Oracle Real Application Clusters and Industry Trends in Cluster Parallelism and Availability" Dec. 2004,18-19
11. Rene Kundersma "11gR2 Oracle Real Application Clusters / Grid Infrastructure N.F." otn.oracle.com/rac
12. Markus Michalewic "Oracle Database 12c Real Application Clusters (RAC)" An Oracle White Paper, June 2013, 15-17