ISSN 2229-5518

Oracle Real Application Clusters

Deepali Kadam, Nandan Bhalwarkar, Rahul Neware, Rajesh Sapkale, Raunika Lamge

Abstract -This paper shows automatic failover and load balancing for Oracle real Application Clusters. rac enables you 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. The Cluster Interconnect enables instances to pass coordination information and data images between each other.The architecture enables users and application to benefit from the processing power of multiple machines. Oracle RAC architecture also achieves redundancy in the case of, for example, a system crashing or becoming unavailable; the application can still access the database on any surviving instances.

Inexd Terms – Oracle, real application, clusters, rac architecture, clusterware, loadbalancing, server.

1 INTRODUCTION

In the last five years, researchers have shown that commodity clusters have the potential to provide supercomputing capabilities at a fraction of the cost of traditional multiprocessor systems . At the same time, much research has been conducted in software distributed shared memory to make it as easy to program clusters as it is to program shared memory multiprocessors.

These advances in programmability and performance are making it possible to use very large clusters as a cost-effective platform for data-intensive, long-running applications. As cluster size and application running times increase, adding failover becomes critical.At the same time, to preserve performance scalable.
In this paper, we address the problem of designing Automatic failover system, specifically targeting scalable system. The novelty of our work lies in the combination of failover techniques with memory-mapped communication, such that the resulting system can be used in very large local-area clusters. Furthermore,A common approach is, Real Application Clusters manages data access, so that changes are coordinated between the instances and each instance sees a consistent image of the database. The *Cluster Interconnect* enables instances to pass coordination information and data images between each other.

Furthermore,A common approach is, Real Application Clusters manages data access, so that changes are coordinated between the instances and each instance sees a consistent image of the database. The *Cluster Interconnect* enables instances to pass coordination information and data images between each other.
The architecture enables users and application to benefit from the processing power of multiple machines. RAC architecture also achieves redundancy in the case of, for example, a system crashing or becoming unavailable; the application can still access the database on any surviving instances.

- *DEEPALI KADAM Btech computer engineer from VJTI,MUMBAI,INDIA email :deepalisushpriya@gmail.com*
- *NANDAN BHALWANKAR Btech computer engineer from VJTI,MUMBAI,INDIA email :nandanbhalwankar123@gmail.com*
- *RAHUL NEWARE Btech computer engineer from VJTI,MUMBAI,INDIA email :rahulneware2007@gmail.com*
- *RAJESH SAPKALE Btech computer engineer from VJTI,MUMBAI,INDIA email :rajeshsapkale@hotmail.com*
- *RAUNIKA LAMGE Btech computer engineer from VJTI,MUMBAI,INDIA email :raunika24@gmail.com*
- *PRAMILA CHAWAN professor computer department VJTI,MUMBAI,INDIA email :pmchawan@vjti.org.in*
- *BANDU MESHRAM head of computer department VJTI,MUMBAI,INDIA email :bbmeshram@vjti.org.in*

ISSN 2229-5518

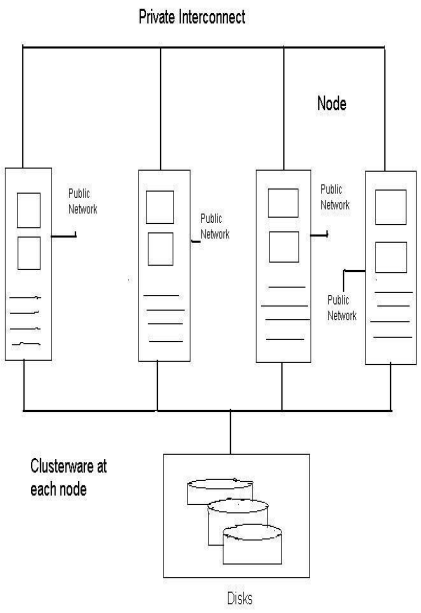

RAC ARCHITECTURE

The diagram consists of two things:
Collection of interconnected Servers, called Clusters connected together using private interconnect.
The database files storing collection of database tables which is shared among multiple instances among servers.
A cluster consists of two or more independent, but interconnected, servers. Several hardware vendors have provided cluster capability over the years to meet a variety of needs. Some clusters were intended only to provide high availability by allowing work to be transferred to a secondary node if the active node fails. Others were designed to provide scalability by allowing user connections or work to be distributed across the nodes.
Another common feature of a cluster is that it should appear to an application as if it were a single server. Similarly, management of several servers should be as similar to the management of a single server as possible. The cluster management software provides this transparency. For the nodes to act as if they were a single server, files must be stored in such a way that they can be found by the specific node that needs them. There are several different cluster topologies that address the data access issue, each dependent on the prim goal of the cluster designer. The interconnect is a physical network used as a means of communication between each node of the cluster.
A cluster comprises multiple interconnected computers or servers that appear as if they are one server to end users and applications. Oracle Real Application Clusters (Oracle RAC) enables you to cluster Oracle databases. Oracle RAC uses *Oracle Clusterware* for the infrastructure to bind multiple servers so they operate as a single system.Oracle Clusterware is a portable cluster management solution that is integrated with the Oracle database. Oracle Clusterware is also a required component for using Oracle RAC. In addition, Oracle Clusterware enables both single-instance Oracle databases and Oracle RAC databases to use the Oracle high-availability infrastructure. Oracle Clusterware enables you to create a clustered pool of storage to be used by any combination of single-instance and Oracle RAC databases.
Oracle Real Application Clusters (RAC) is a software component you can add to a high-availability solution that enables users on multiple machines to access a single database with increased performance. RAC comprises two or more Oracle database instances running on two or more clustered machines and accessing a shared storage device via cluster technology. To support this architecture, the machines that host the database instances are linked by a high-speed interconnect to form the cluster. The interconnect is a physical network used as a means of communication between the nodes of the cluster. Cluster functionality is pro vided by the operating system or compatible third party clustering software.
Oracle RAC offers features in the following areas:
Scalability
Availability
Failover
Load balancing

Oracle RAC is a cluster database with a shared cache architecture that overcomes the limitations of traditional shared-nothing and shared-disk approaches to provide highly scalable and available database solutions for all your business applications. Oracle RAC is a key component of Oracle's enterprise grid architecture. Oracle RAC support is included in the Oracle Database Standard Edition for higher levels of system uptime.

SHARED DISK ARCHITECTURE FOR ORACLE RAC:

A shared disk file system uses a storage area network (SAN) or RAID array to provide direct disk access from multiple computers at the block level. Translation from file-level operations that applications use to block-level operations used by the SAN must take place on the client node. The most common type of clustered file system, a shared disk file system adds a mechanism for concurrency control which conventional file systems intended for local storage do not have. Clients are afforded a consistent and serializable view of the file system, avoiding corruption and unintended data loss even when multiple clients try to access the same files at the same time. Shared disk file systems also usually employ some sort of a fencing mechanism to prevent data corruption in

case of node failures.

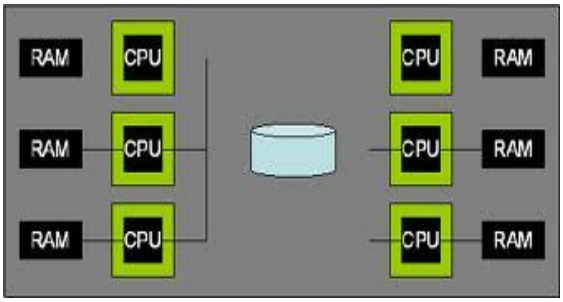

The underlying storage area network might use any of a number of block-level protocols, including SCSI, iSCSI, HyperSCSI, ATA over Ethernet (AoE), Fibre Channel, and InfiniBand.
There are different architectural approaches to a shared disk file system. Some distribute file information across all the servers in a cluster (fully distributed). Others utilize a centralized metadata server. Both achieve the same result of enabling all servers to access all the data on a shared storage device

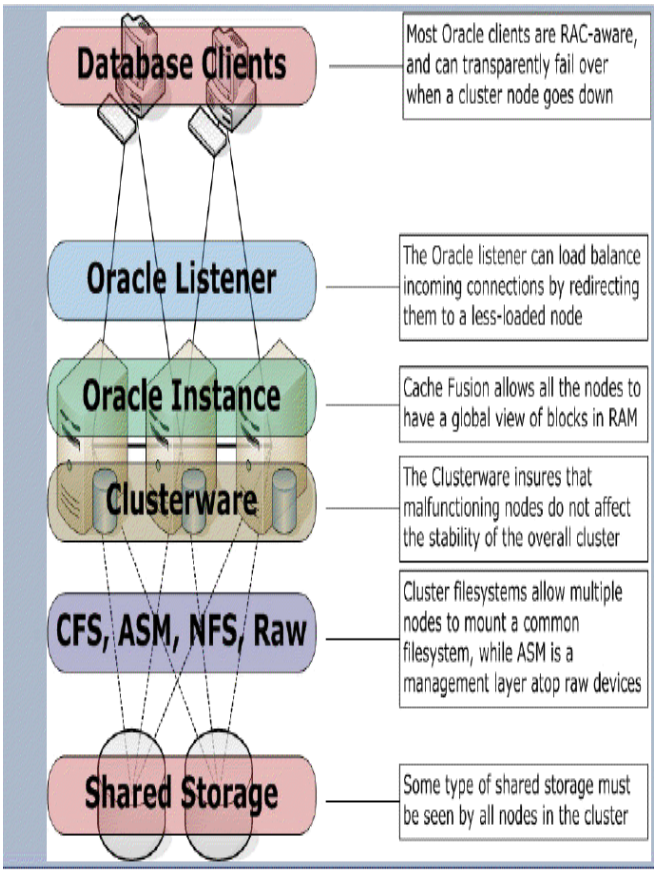

CLUSTERWARE :

Oracle Clusterware is the software, which enables the nodes to communicate with each other, and forms the cluster and makes the nodes as single logical server. Oracle Clusterware is run by Cluster Ready Services (CRS) using two key components. They are Oracle Cluster Registry (OCR), which records and maintains the cluster and node membership information. The other component is voting disk, which acts a tiebreaker during communication failures. Consistent heartbeat information from all the nodes is sent to voting disk when the cluster is running. CRS service has four components namely OPROCd, CRS Daemon (crsd), Oracle Cluster Synchronization Service Daemon (OCSSD) and Event Volume Manager Daemon (evmd) and each handles a variety of functions. Failure or death of the CRS daemon can cause the node failure and it automatically reboots the nodes to avoid the data corruption because of the possible communication failure between the nodes. The CRS daemon runs as the super user 'root' in the UNIX platforms and runs as a service in the windows platforms.

LOAD BALANCING :

Clusters have become increasingly popular as powerful and cost-effective platforms for executing parallel applications In such systems, load-balancing schemes can improve system performance by attempting to assign work, at run time, to machines with idle or underutilized resources.
The Oracle RAC system can distribute the load over many nodes this feature called as load balancing.

There are two methods of load balancing

1.Client load balancing
2.Server load balancing

Client Load Balancing distributes new connections among Oracle RAC nodes so that no one server is overloaded with connection requests and it is configured at net service name level by providing multiple descriptions in a description list or multiple addresses in an address list. For example, if connection fails over to another node in case of failure, the client load balancing ensures that the redirected connections are distributed among other nodes in the RAC.

Server Load Balancing distributes processing workload among Oracle RAC nodes. It divides the connection load evenly between all available listeners and distributes new user session connection requests to the least loaded listener(s) based on the total number of sessions which are already connected. Each listener communicates with the other listener(s) via each database instance's PMON process.
A farm of servers with the same function is the base of a load balancing cluster. To distribute the user requests to several nodes, a load balancer is useful. The load balancer checks the utilization of all nodes. The node with the estimated best performance will get the next user request. This algorithm ensures the best performance available at the time is given to the users.
A very important point is the quality of the load balancer. In this case, quality means the opportunity of the system to made a highly qualified forecast about which node will offer the best performance concerning an individual application. Another aspect to make the decision is to reconnect the client to an existing session. Without any kind of protection, the load balancer could be a single point of failure.
If one node is out of service, then the whole system will still work. The load balancer will recognize the failed node, and mark the crashed system. The total performance of the load balancing cluster will be reduced, but services will still be provided.

AUTOMATIC FAILOVER:

Automatic failover is supported in database for high-safety mode. In high-safety mode with automatic failover, once the database is synchronized, if the principal server becomes unavailable, an automatic failover occurs. An automatic failover causes the secondary server to take over the role of principal server and bring its copy of the database to the user. Requiring that the server be synchronized prevents loss to the user during failover, because every transaction committed on the principal server can also be committed on the secondary server.

Automatic failover requires the following conditions:

1. The secondary server must be running in high-safety mode
2. The secondary server must have access to the main database

How Automatic Failover Works

Under the preceding conditions, automatic failover initiates the following sequence of actions:

1. If the principal server fails, it changes the state of the principal server to DISCONNECTED and disconnects all clients from the principal server.
2. The secondary server register that the principal server is unavailable.
3. All clients from the principal server are shifted to the secondary server by clusterware.

CONCLUSIONS:

In this paper we have shown automatic failover and load balncing for oracle real Application Clusters. an oracle rac system can protect against computer failures caused by unexpected hardware failures and operating system or server crashes, as well as processing loss caused by planned maintenance.. When node failover occurs and a service connection is redirected to another node, users can continue to access the service, unaware that it is now provided from a different node. We have also explained the architecture of oracle real application cluster,oracle clusterware,shared disk architecture.

.

REFERENCES

[1] T.E. Anderson, D.E. Culler, D.A. Patterson. A Case for NOW (Networks of Workstations). *IEEE Micro*, v. 15, no. 1, pp. 54-64, February 1995.

[2] R. Bianchini, L. I. Kontothanassis, R. Pinto, M. De Maria, M. Abud, C.L. Amorim. Hiding Communication Latency and Coherence Overhead in Software DSMs. *Proc. 7th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 1996.

[3] N. J. Boden, D. Cohen, R. E. Felderman, A. E. Kulawik, C. L. Seitz, J. N. Seizovic, W. Su. Myrinet: A