



Spotlight on RAC (Real Application Clusters)

Product Overview



Table of contents

1	INTRODUCTION	3
1.1	PURPOSE OF THIS DOCUMENT	3
1.2	WHAT IS SPOTLIGHT ON RAC?	3
1.3	FURTHER READING.....	3
2	OVERVIEW OF ORACLE RAC.....	4
3	OVERVIEW OF SPOTLIGHT ON RAC.....	6
3.1	WHAT DOES SORAC DO?	6
3.2	SPOTLIGHT ON RAC AND SPOTLIGHT ON ORACLE	6
3.3	PLATFORM SUPPORT	6
4	SPOTLIGHT ON RAC HOME PAGE	7

1 Introduction

1.1 Purpose of this document

This document is designed to give an introductory overview of Quest's new Spotlight on RAC (SoRAC) product. It is intended to be read by Database Administrators who currently have (or will have) responsibility for administering an Oracle RAC environment.

1.2 What is Spotlight on RAC?

Spotlight on RAC is a new Spotlight which is designed to present a full diagnostic capability for an Oracle RAC (Real Application Clusters) environment.

Spotlight on Oracle (SoO) provides a basic RAC diagnostic capability via a set of cluster-related drilldowns. This functionality is limited because it essentially provided only a single instance's view of the cluster.

SoRAC takes RAC diagnostics to a whole new level by issuing queries on each instance in the cluster, and then doing appropriate calculations to provide RAC-wide diagnostic information.

As an Oracle RAC DBA, you will be interested in the following:

- 1) How is each of the instances in my cluster performing?
- 2) How is the database as a whole performing?
- 3) Is my database experiencing cluster related contention?

SoO provides some information about 1) ... how each instance in the cluster is working. SoRAC deals with both 1), 2) and 3).

1.3 Further Reading

For more technical detailed information about Oracle 9i RAC, you should acquire the book **Oracle Real Application Clusters** by Murali Vallath (ISBN 1-55558-288-5), and about Oracle 10g RAC, you should acquire the book when available **Oracle 10g RAC, Grid, Services and Clustering** by Murali Vallath (ISBN 1-55558-321-01).

For an introduction to Quest's Spotlight technology, you should read the Getting Started Guide for either Spotlight on Oracle or Spotlight on RAC.

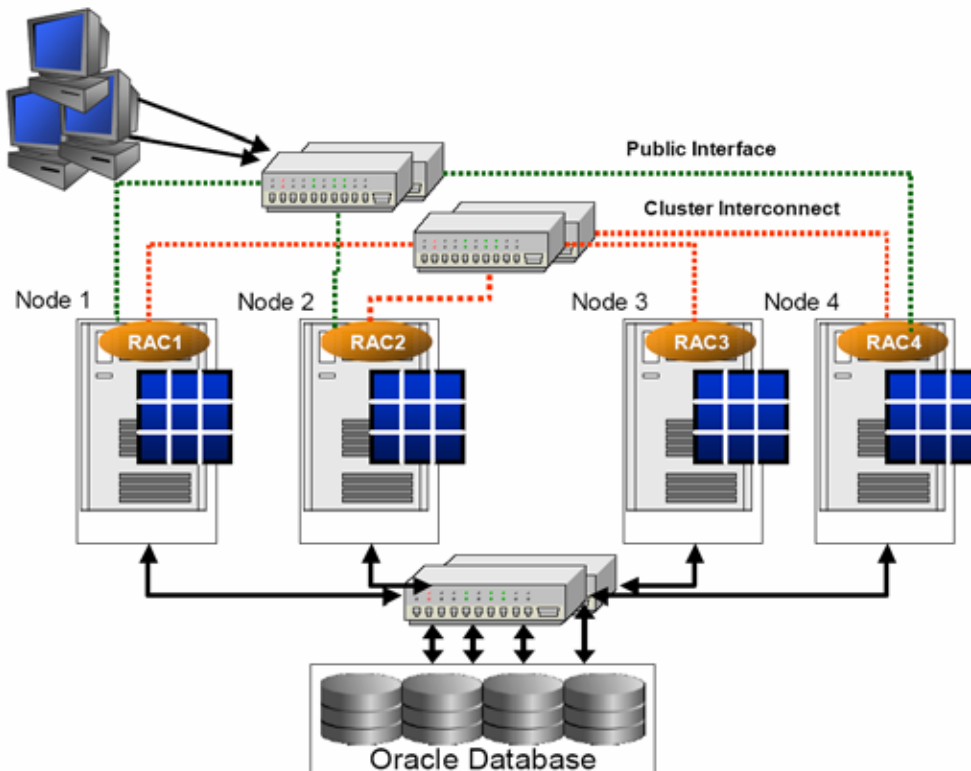
2 Overview of Oracle RAC

Real Application Clusters (RAC) is a clustering technology introduced by Oracle in release 9i. It supersedes Oracle's previous clustering technology Oracle Parallel Server (OPS). OPS technology was introduced in the early 90s but failed to achieve significant uptake because of (a) extremely challenging implementation and configuration issues and (b) inability to scale effectively for most applications.

RAC is a "shared disk" clustering solution and is significantly different from the "shared nothing" architectures offered by Microsoft SQL Server and IBM DB2. Each node in a RAC cluster has equal access to all of the database data through a shared disk subsystem. Data is neither partitioned to specific nodes, nor replicated across nodes. A high-speed network interconnect allows each node to keep its in-memory view of data consistent.

An Oracle RAC configuration will have the following architectural characteristics:

1. Many instances of Oracle running across several nodes
2. Many instances sharing a single physical copy of a shared Oracle database
3. All instances having common data and control files
4. Each instance having individual redo logs and undo segments
5. Each of the instances can simultaneously execute transactions against the same database.



Source: 10g New Features for RAC - Murali Vallath – (upcoming book from digital press)

Although Oracle claims that RAC offers scalability without changes to application architecture, this is only true if the application is correctly designed for high performance on SMP systems. Any existing contention points in an application may be amplified when such an application is migrated to RAC - leading to reduced or even negative scalability.

The key limit to RAC scalability is the maintenance of consistent in-memory views of cached data held across multiple instances of the cluster. Because Oracle "lazy writes" modified data to disk some time after the modification, it is possible that the correct version of any particular piece of data may reside in memory in one instance only and may not yet have been written to disk. Keeping these caches coherent is the major barrier to scalability in a RAC cluster.

The primary RAC-specific areas of interest in an Oracle database are:

- 1) The network interconnect between the various nodes
- 2) The shared storage subsystem
- 3) The movement of cache (cache synchronization) between the various instances.

Performance Management and Scalability

Customer expectations of RAC performance and scalability cannot be obtained or validated without visibility into load and performance bottlenecks and advice on tuning and optimization activities. Out of the box, it is quite possible that an initial RAC installation will perform more poorly than a single instance equivalent due to any of the following:

- High latency across the hardware interconnect. This can be due to poor network configuration or inappropriate configuration of the RAC instances themselves.
- Excessive inter-node communication overhead. This can occur when "hot blocks" are continuously transferred between nodes due to application/schema design issues.
- Poor load balancing across nodes. This can occur if users are allowed to connect to specific nodes, batch jobs are inappropriately distributed across nodes or if Oracle's load balancing mechanisms fail to adequately distribute load.

Providing insight into, and solutions for, these performance and scalability problems is the primary goal of Spotlight on RAC.

Visibility and Utilization

To the application, a RAC cluster appears as one logical database. However, most diagnostic tools (including SoO) allow visibility only into individual instances of the RAC cluster. This makes it impossible to determine overall load on the cluster, determine which activities are consuming resources, isolate the causes of inter-instance contention, or solve any problem that requires a consolidated view of all instances in the cluster.

Therefore the secondary goal for Spotlight on RAC is to provide full visibility into activities within the cluster.

3 Overview of Spotlight on RAC

3.1 What does SoRAC do?

SoRAC provides three classes of functionality to the RAC Database Administrator:

1. Via the SoRAC home page, the physical architecture of the RAC system is exposed, as depicted in the diagram in Section 2. The three essential components of a cluster: the individual instances, the interconnect and the I/O subsystem are displayed, including the essential data flows between them. At a glance, you can obtain an understanding of what the cluster is doing, and whether there are problems at either the cluster level, or individual instance level.
2. Oracle diagnostics are calculated on the instance as a whole. SoRAC makes a connection to each instance in the cluster, and then executes a subset of the Spotlight on Oracle queries on those instances. The outcome of those queries are then aggregated appropriately to provide cluster-wide Oracle metrics with alarm rules against them.

For example, suppose there is latch contention on *some* but not all of the instances in the cluster. This may not necessarily be a problem if there are sufficient instances that are free of latch contention. SoRAC aggregates all the latch wait percentages from all the instances to calculate a cluster-wide metric, which it then applies alarm thresholds to.

3. SoRAC evaluates several RAC-specific alarms to provide diagnostic information that is explicitly related to RAC. Most significantly, SoRAC will raise a diagnostic alarm if:
 - a. Database load is not evenly distributed across instances within the cluster
 - b. The overhead of maintaining consistency across the cluster is excessive (e.g. the database is experiencing degradation because it is part of a cluster)
 - c. There is high latency across the interconnect which links instances in the cluster
 - d. There is significant inter-instance communication (which might indicate the need for data partitioning)
 - e. Blocks are becoming corrupt or lost while being transferred between the instances in the cluster
 - f. The shared disk subsystem is experiencing contention

3.2 Spotlight on RAC and Spotlight on Oracle

SoRAC 1.0 is bundled and integrated with SoO (Version 4.0). By “bundled”, we mean that the Spotlight on RAC installer will install both the SoRAC plug-in and the SoO plug-in. By “integrated” we mean that you can launch SoO on any particular instance in the cluster, from the SoRAC home page.

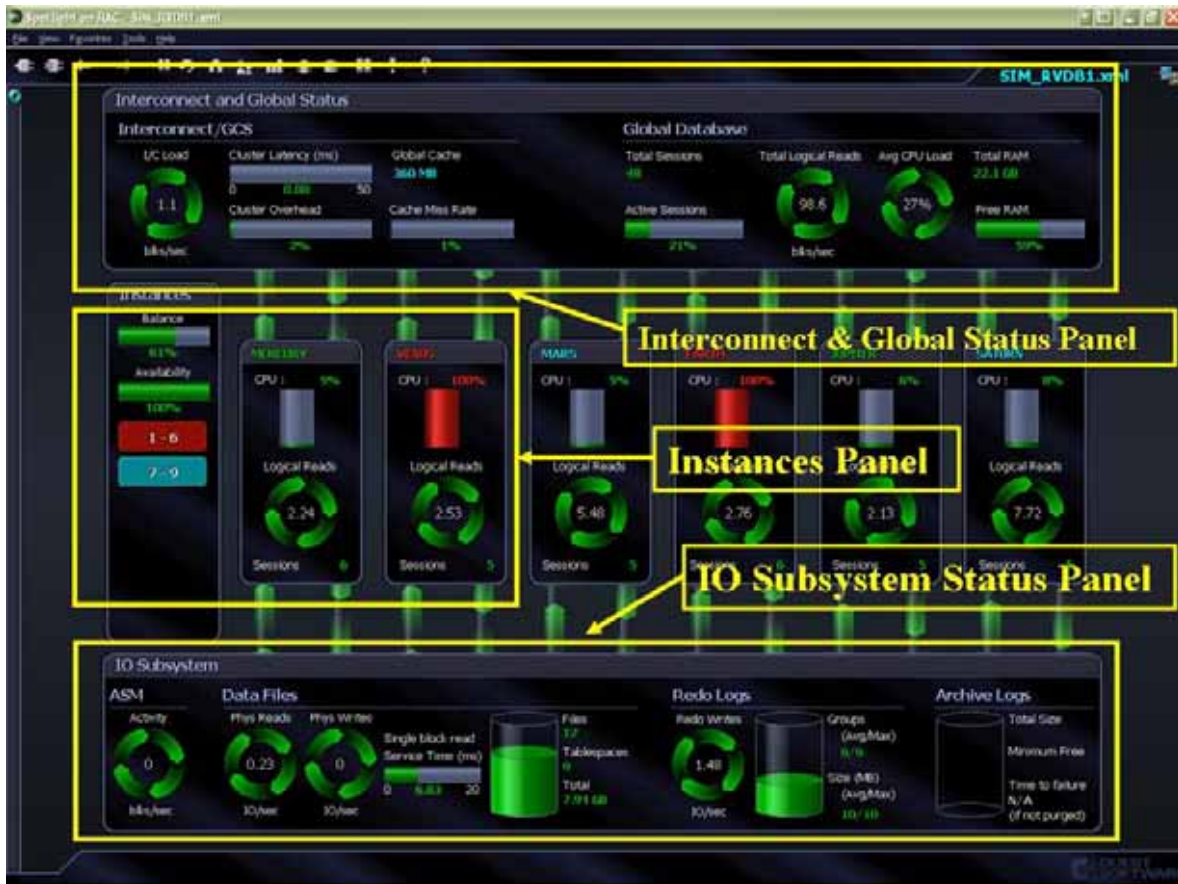
3.3 Platform Support

For the initial release of SoRAC, support will only be provided for RAC running on UNIX operating systems.

Formal support is provided for up to 16-instance clusters. The product currently supports a theoretical maximum of 36 instances in the cluster. Future releases of SoRAC will support larger clusters.

4 Spotlight on RAC Home Page

The SoRAC home page provides a detailed view of the RAC cluster as a whole:



The home page has three main panels, depicting the architecture shown in Section 2 of this document.

Instances Panel

The middle panel shows the participating instances of a cluster. SoRAC supports a theoretical maximum of 36 instances, with up to 6 instances shown on the home page. The Navigation panel (in this screenshot, showing “1-6” and “7-9”) can be used to navigate between blocks of 6 instances. In this example, we are looking at a 9-instance cluster, with instances 1-6 being shown, and 7-9 can be shown by clicking on the “7-9” component.

The color of the label on each instance (in this case “MERCURY”, “VENUS”, “MARS”, etc) represents the Spotlight on Oracle severity of the instance. So for example, instance “VENUS” is red severity because of the 100% CPU level. You can mouse over the instance name to display a list of all the alarms currently raised against that instance. The navigation buttons on the left indicate the highest severity of instances in that block; so in this screen shot, the highest severity of instances 7-9 (not currently being shown on the home page) is blue.

If you click on any of the individual instances in the Instances Panel, Spotlight on Oracle will be launched for that particular instance. This allows you to drill down further in diagnosing problems on the selected instance.

You will notice that whilst there are red alarms on two Oracle instances, the overall severity of the cluster, as shown by the color of the connection name at the top right, is blue. That is because problems with individual instances don't necessarily translate to a problem with the cluster as a whole. For example, two instances are showing 100% CPU, but the overall average CPU load (as shown by the spinner in the Global Database panel) is only 27% - not an alarm condition.

Two further metrics are shown in the instances panel:

Balance indicates the comparative load on the instances. This is described in a later section of this document.

Availability shows the percentage of instances for which a SQL*Net connection is currently open and working. Should individual instances become unavailable, SoRAC will continue to remain open and connected. Note that this is different behavior to SoO, where if the SQL*Net connection becomes unavailable, then the Spotlight connection is closed. Note also that connectivity to the OS (via the OS user) is *not* used in the Availability calculation (although obviously if the OS is down, then the instance on that OS will also be unavailable).



Interconnect and Global Panel

The top panel shows a combination of Interconnect and Global Cache Service information, and overall Database information. It consists of information aggregated from the queries sent out to the individual instances.

Some information displayed here is easily calculated: for example, the Total Logical Reads spinner is the sum of the individual Logical Reads values. Other data, such as Cluster Latency and Cluster Overhead, has a more complex calculation, which is described further in the following sections.

The flows between the Interconnect Panel and the individual instances represent the flow of data from one instance to another across the interconnect, based on the data being requested by one instance being available in the buffer cache of another instance.

I/O Subsystem Panel

The bottom panel shows information relating to the shared storage subsystem.

Some of this information is aggregated from the queries sent out to the individual instances; other information is obtained from a single query to one of the instances, for example information about data files and tablespaces.

The ASM spinner on the left of this panel is only populated if ASM is being used and has been defined in the ASM section of the connection dialog.

The flows between the I/O Subsystem Panel and the individual instances represent the flow of physical reads and writes between the instances and the physical database.