

Oracle Database 11g: Manageability Overview

*An Oracle White Paper
August 2007*

Oracle Database 11g: Manageability Overview

Introduction	3
Manageability	3
ADDM for RAC	3
Automatic SQL Tuning	4
SQL Plan Management	5
SQL Access Advisor Enhancements: Partition Advisor	6
Automatic Memory Management	7
AWR Baselines and Adaptive Thresholds	8
Fault Diagnostic Infrastructure	10
Health checks	10
Data Recovery Advisor	10
SQL Repair Advisor	11
SQL Test Case Builder	11
Automatic Diagnostic Repository (ADR)	11
Incident packaging service (IPS)	11
Support Workbench	12
Conclusion	13

Oracle Database 11g: Manageability Overview

INTRODUCTION

The Oracle database is the market-leader and the preferred database for hundreds of thousands of enterprises as well as application developers and database administrators worldwide. Over the years, enterprises have come to rely on the Oracle database to provide unparalleled performance and reliability. In release 10g of the database, Oracle delivered a self-managing database with breakthrough manageability, dramatically reducing management costs. In this release, Oracle Database 11g has made significant advances in the areas of fault diagnostics to reduce time to diagnosis and resolution and improving the availability and reliability of your mission critical databases. Oracle has also continued to make major improvements in all manageability aspects of the database making Oracle 11g Database more self-managing than ever before.

MANAGEABILITY

ADDM for RAC

Oracle Database 10g introduced Automatic Database Diagnostic Monitor (ADDM), which was a revolutionary feature that helped create the first self-managing database. ADDM uses an integrated approach to provide database-wide performance analysis, which covered storage, system resource, space, application & SQL and backup & recovery management. It provides proactive analysis to DBAs and is available on demand to troubleshoot performance problems.

Oracle Database 11g extends ADDM by offering cluster-wide performance analysis for Real Application Clusters (RAC) databases. For RAC environments ADDM analyses the RAC cluster and reports on issues that are affecting the entire database as well as its individual instances. DBAs can now use ADDM to perform database-wide analysis of global resources, such as high-load SQL, global cache interconnect traffic, network latency issues, skew in instance response times, I/O capacity, etc. DBAs also have the ability to restrict ADDM analysis on a few specified instances of a RAC cluster as well. With ADDM for RAC, performance analysis of a RAC database becomes as simple as that of a single instance database.

In Oracle Database 11g, ADDM findings can be suppressed by DBAs using directives to filter and display only findings of interest. To better understand the

impact of the findings over time, each finding has a descriptive name that facilitates search, a link to the number of previous occurrences of the finding in last 24 hours, and affected instances.

Automatic SQL Tuning

Poor SQL performance is a common cause of badly performing databases. Many DBAs have traditionally tried to tackle this using manual SQL tuning processes. Manual SQL tuning is a complex and recurring process that presents many challenges. It is very time consuming, and requires an intimate knowledge of the schema structures and the data usage model of the application and query plans. All these factors make manual SQL tuning a challenging and resource intensive task that is ultimately very expensive for businesses.

SQL Tuning Advisor was introduced in Oracle 10g to automate the SQL tuning process by comprehensively analyzing SQL statements. The output of this analysis is in the form of recommendations, along with a rationale for each recommendation and its expected performance benefit. The recommendation relates to collection of statistics on objects, creation of new indexes, restructuring of the SQL statements, or creation of SQL Profiles. A user can review the recommendations and implement them manually as appropriate.

In Oracle Database 11g, the SQL tuning process has been further enhanced and automated to keep databases running at their peak performance. The SQL Tuning Advisor now runs automatically during the system maintenance windows as a maintenance task. In each run, it auto-selects high-load SQL queries in the system, and generates recommendations on how to tune them.

To validate the recommendation, SQL Tuning Advisor in Oracle Database 11g performs a test-execute of the SQL statements with the new execution plan for which a SQL Profile is recommended. This dramatically increases the accuracy and reliability of SQL Profile recommendations.

Automatic SQL Tuning Advisor can be configured to auto-implement SQL Profile recommendations. If you enable automatic implementation, the advisor will create SQL Profiles for only those SQL statements, where the performance improvement would be at least **three-fold**. Other types of recommendations, such as the ones to create new indexes or refresh optimizer statistics or the ones that restructure SQL, can only be implemented manually. DML statements are not considered for tuning by Automatic SQL Tuning Advisor. By default, the Automatic SQL Tuning Advisor is configured to run nightly and only report recommendations but not auto implement them.

You can view a summary of the results of automatic SQL tuning over a specified period (such as the previous seven days), as well as view a detailed report on recommendations made for all SQL statements processed. The recommendations can then be selectively implemented by a manual process. You can also view the recommendations that were automatically implemented. The Automatic SQL

Tuning Advisor can be configured to run in any maintenance window or can be disabled altogether if desired.

SQL Plan Management

SQL plan management prevents performance regressions resulting from sudden changes to the execution plan of a SQL statement by providing components for capturing, selecting, and evolving SQL execution plans. SQL performance can be affected by various changes, such as new optimizer version, changes to optimizer statistics and/or parameters, or creation of SQL profiles. SQL plan management is a preventative mechanism that records and evaluates the execution plans of SQL statements over time, and builds SQL plan baselines composed of a set of existing plans known to be efficient. The SQL plan baselines are then used to preserve performance of the corresponding SQL statements, regardless of changes occurring in the system.

Common usage scenarios where SQL plan management can improve or preserve SQL performance include:

- A database upgrade that installs a new optimizer version usually results in plan changes for a small percentage of SQL statements, with most of the plan changes resulting in either improvement or no performance changes. However, certain plan changes may cause performance regressions. The use of SQL plan baselines significantly minimizes potential performance regressions resulting from a database upgrade.
- Ongoing system and data changes can impact plans for some SQL statements, potentially causing performance regressions. The use of SQL plan baselines can also help to minimize performance regressions and stabilize SQL performance.
- Deployment of new application modules means introducing new SQL statements into the system. The application software may use appropriate SQL execution plans developed under a standard test configuration for the new SQL statements.

SQL plan baselines evolve over time to produce better performance. During the SQL plan baseline evolution phase, Oracle Database 11g routinely evaluates the performance of new plans and integrates plans with better performance into SQL plan baselines. A successful verification of a new plan consists of comparing its performance to that of a plan selected from the SQL plan baseline and ensuring that it delivers better performance.

There are three ways of evolving SQL plan baselines:

1. Manually by loading new plans verified by the user into existing SQL plan baselines.
2. Manually by using the `EVOLVE_SQL_PLAN_BASELINE` function of the `DBMS_SPM` PL/SQL package to verify new plans.

3. Automatically using the Automatic SQL Tuning capabilities of Oracle Database 11g.

SQL Access Advisor Enhancements: Partition Advisor

SQL Access Advisor has been enhanced in Oracle Database 11g to provide partition advice as a part of SQL access structure recommendations. The new and improved SQL Access Advisor now gives comprehensive advice on how to optimize schema design for best performance based on the system workload. The SQL Access Advisor takes in actual or synthetic SQL workloads as input and recommends access structures for improved performance. The recommended access structures include partitioning recommendations for tables and indexes, and materialized views as well as recommendations to create new or drop existing indexes (b-tree, bitmap, and functional indexes), materialized views and materialized view logs. SQL Access Advisor considers both queries and DML when offering recommendations.

The partition recommendations are only provided for workloads have some predicates and joins on columns of type NUMBER or DATE. Partitioning advice is only generated on the above column types and is restricted to a single column INTERVAL, HASH or RANGE partitioning. The SQL Access Advisor is sophisticated enough to identify candidates for partition and suggest partition keys and ranges for the above kinds of partitions.

Similar to SQL Tuning Advisor, SQL Access Advisor leverages the existing Cost-Based Optimizer (CBO) rules and is an easy-to-use wizard based solution. Because of the tight integration between SQL Access Advisor and the database kernel, the advisor makes the optimal recommendation for access structures based on the updated CBO rules that the kernel ships with.

SQL Access Advisor can also make recommendations for a combination of index, materialized view and partitioning solution. The factors that are considered when making SQL Access Advisor recommendations include storage (for creation and maintenance costs), whether workloads are full or partial and the overall benefit to the queries in the workload.

When processing large workloads, SQL Access Advisor can be interrupted and will offer intermediate recommendations for the set of SQL that have been processed thus far. The order in which SQL are processed by the SQL Access Advisor can be configured by the user.

Oracle Enterprise Manager displays the results of the SQL Access Advisor task by listing SQL statements in the order of highest cost improvement. DBAs have the option of executing the recommendation right away at the touch of a button. Or, in more stringent environments, DBAs can create a script with the set of executable SQL statements to implement the recommendations.

Automatic Memory Management

Oracle Database memory structures basically consist of shared memory or System Global Area (SGA) and private memory or Program Global Area (PGA). In Oracle Database 9i, Automatic SQL Execution Memory Management feature was introduced to automate management of PGA. In Oracle Database 10g, the same was done for SGA by the introduction of Automatic Shared Memory Management. This meant all the different SQL areas in PGA were auto-sized for the system workload to give best performance and all the memory pools in shared memory were similarly adjusted for size for optimal performance. The user was only required to specify the PGA and SGA target sizes and Oracle would appropriately allocate memory within these targets to give the best possible performance. PGA and SGA Advisors were also provided to help the user properly set the targets for SGA and PGA in Oracle Database 10g.

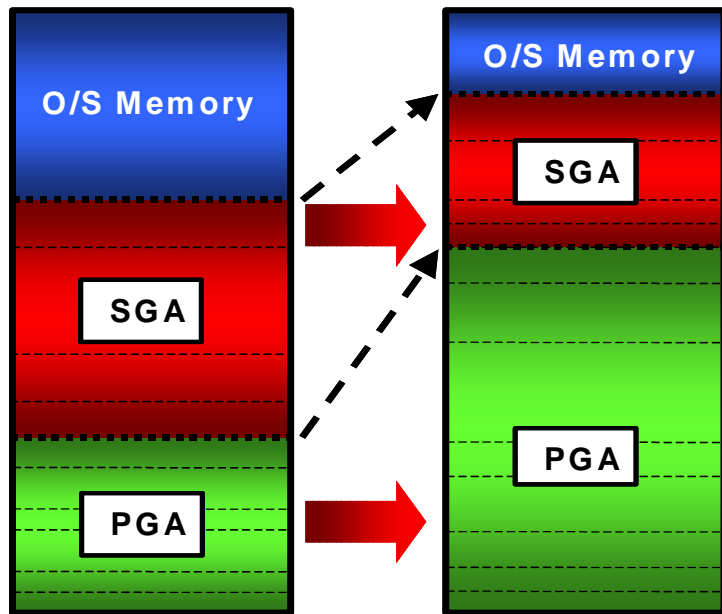


Figure 1: Automatic Memory Management

In Oracle Database 11g, memory management has been automated even further. All memory, PGA and SGA, is now managed centrally with the help of the Automatic Memory Management feature. DBAs need to specify a single parameter, `MEMORY_TARGET`, and Oracle will automatically size the Program Global Area (PGA) and System Global Area (SGA) based on the workload. Using indirect memory transfer, the database transfers memory from SGA to PGA and vice versa to respond to the load. The indirect transfer uses the operating system mechanism to free up shared memory and allocating memory to other components requesting memory, e.g., from PGA to SGA. Dynamic allocation of memory is adjusted at frequent intervals to optimize memory usage in line with workload requirements to maximize memory utilization and avoid out-of-memory errors. Users can optionally set SGA and PGA targets when using the Automatic Memory

Management feature. This ensures that SGA and PGA sizes will not be shrunk below the values specified by their respective parameter targets in auto-tuning mode. This feature is currently available in Linux, Solaris, HP-UX, AIX and Windows platforms.

First introduced in Oracle Database 10g, Memory Advisors provide graphical analyses of total memory target settings, SGA and PGA target settings, or SGA component size settings. DBAs can use these analyses to tune database performance and to perform what-if planning scenarios. Different memory advisors become available depending on the memory management mode used with the database.

For instance, if Automatic Memory Management is enabled, you can get advice for setting the target amount of memory to allocate to the entire database. This advisor provides advice for the total memory target for the instance. If Automatic Shared Memory Management is enabled, you can gain advice on configuring the target sizes of the SGA and instance PGA. If Manual Shared Memory Management is enabled, you can get advice on sizing the shared pool, buffer cache, and instance PGA.

AWR Baselines and Adaptive Thresholds

Automatic Workload Repository (AWR) was one of the most prominent self-managing features of Oracle Database 10g. The Oracle Database captures real-time and historical performance statistics in memory and in the database respectively to give DBAs the right tools and information to troubleshoot performance problems.

AWR Baselines allow DBAs to capture system performance over time periods with interesting or representative workloads. For example, if a company's current month payroll processing was slow, then the DBA can compare the system performance with last month's payroll processing to identify causes for the problems.

The AWR Compare Periods report is provided to easily make comparisons of problem time periods against saved Baselines to identify potential sources of performance deviations. In addition to performance metrics, the report also captures the configuration information, such as the total memory, and number of CPUs, which can identify external sources of problems that may have caused the reduction in performance. In case there were out-of-band changes made to critical parameters, such as COMPATIBLE, which may affect SQL performance, the report also captures database system information, such as the initialization parameters.

In addition Baselines can also be used in setting alert thresholds on system performance metrics. Most metrics can be viewed in Oracle Enterprise Manager against statistical aggregates of those same metrics observed over the Baseline period. This helps users set Baseline-informed thresholds rather than selecting thresholds without the context of actual data. In addition, Adaptive Thresholds are

available for certain key performance metrics. Adaptive Thresholds are performance alert thresholds that are automatically set and periodically adjusted by the system using the System Moving Window Baseline data as the basis for threshold determination. For customers who want to get started with Adaptive Thresholds immediately, the new “Quick Configure” option can setup a starter kit of thresholds based on common workload profiles using a few mouse clicks.

There are three types of available baselines in Oracle Database:

1. Fixed Baselines

A fixed baseline corresponds to a fixed, contiguous time period in the past specified by the user. Typically the time period chosen as baseline should represent the system operating at an optimal level so that during periods of poor performance, comparison can be made to the baselines to analyze causes of performance degradation.

2. System Moving Window

The System Moving Window is available out-of-box and defined to be all AWR data available extending from the present back a specified window size into the past (expressed as number of days). By default this window size is the current AWR retention period, which is 8 days. If you are planning to use Adaptive Thresholds consider a larger moving window (such as 35 days) to compute better threshold values over larger data samples. The System Moving Window size can be configured to be less than AWR retention in cases where customers set the latter to be very large. As a general rule the System Moving Window should optimally be between 3 and 13 weeks in size.

3. Baseline Templates

You can also create baselines for a contiguous time period in the future using baseline templates. There are two types of baseline templates, single and repeating. A single baseline template can be used to create a baseline for a single contiguous time period in the future. This is useful if you know beforehand of a time period that you want to capture in the future. For example, you may want to capture the AWR data during a system test that is scheduled for the upcoming weekend. In this case, you can create a single baseline template to automatically capture the time period when the test will take place.

A repeating baseline template can be used to create and drop baselines based on a repeating time schedule. This is useful if you want the Oracle database to automatically capture a contiguous time period on an ongoing basis. For example, you may want to capture the AWR data during every Monday morning for a month. In this case, you can create a repeating baseline template to automatically create baselines on a repeating schedule for every Monday, and automatically remove older baselines after a specified expiration interval such as one month.

AWR baselines provide powerful capabilities for defining dynamic and future baselines and considerably simplify the process of creating and managing performance data for comparison purposes.

Fault Diagnostic Infrastructure

Beginning with Release 11g, Oracle Database includes an advanced fault diagnostic infrastructure for preventing, detecting, diagnosing, and resolving problems. The problems that are targeted in particular are critical errors that can affect the health of the database. When a critical error occurs, it is assigned an incident number, and diagnostic data for the error (traces, dumps, and more) are immediately captured and tagged with this number. The data is then stored in the Automatic Diagnostic Repository (ADR)—a file based repository outside the database—where it can later be retrieved by incident number and analyzed. The extensive improvement of the fault diagnostics infrastructure in Oracle Database 11g aims to provide the following benefits:

- Respond proactively to small problems and prevent catastrophic system failure by alerting DBAs using Health Checks.
- Limiting damage and repair and interruptions after a problem is detected using Data Recovery and SQL Repair Advisor.
- Reducing problem diagnostic time through ADR and Test Case Builder.
- Simplifying customer interaction with Oracle Support using IPS and Oracle Configuration Support Manager.

The following are the key components of the fault diagnostic infrastructure:

Health checks

Health checker framework has been added in Oracle Database 11g for the purposes of performing proactive checks on system health. Upon detecting a critical error, the fault diagnostic infrastructure can run one or more health checks to perform deeper analysis of a critical error. The results of a health check are stored in a report that can be viewed as a text file or as formatted HTML in a browser. The report can be added to the other diagnostic data collected for the error. Separate individual health checks look for data corruptions, undo and redo corruptions, data dictionary corruption, and more. As a DBA, you also have the option to manually invoke these health checks, either on a regular basis or as required.

Data Recovery Advisor

You use the Data Recovery Advisor to repair data block corruptions, undo corruptions, data dictionary corruptions, and more. The Data Recovery Advisor integrates with the Support Workbench facility in Oracle Enterprise Manager and with the RMAN utility to display data corruption problems, assess their extent and impact and recommend repair options for them.

SQL Repair Advisor

The SQL Repair Advisor is a new feature that helps DBAs diagnose SQL problems. If a SQL statement fails with a critical error (e.g., an ORA-600 error), you can use the SQL Repair Advisor to analyze the problem and in many cases it can recommend a SQL patch to repair the statement. By applying the SQL patch, the SQL failure is circumvented by causing the query optimizer to choose an alternate execution plan for future executions.

SQL Test Case Builder

For many application problems, obtaining a reproducible test case is an important factor in problem resolution speed. The SQL Test Case Builder allows a user to automatically gather all the necessary information needed to reproduce the problem such as SQL text, PL/SQL, DDL, execution environment information, etc. The information gathered can then be transmitted to Oracle Support to help reproduce the problem.

Automatic Diagnostic Repository (ADR)

The ADR is a file-based repository for database diagnostic data such as traces, dumps, the alert log, health monitor reports, and more. It has a unified directory structure across multiple instances and components of the Oracle Database and it replaces the `USER_DUMP_DEST`, `BACKGROUND_DUMP_DEST`, and `CORE_DUMP_DEST` of previous releases. The diagnostic data in ADR is self-managing and is purged automatically based on predefined data retention settings. ADR also maintains meta-data for all critical errors on the database such that a user can run queries against ADR to determine what and how many critical problems occurred on the system over the last few days, months or even years. The data in ADR can be viewed using Oracle Enterprise Manager or through a command-line utility called ADR Command Interpreter or `ADRCI`.

Incident packaging service (IPS)

Incident Packaging Service automates the process of collecting all necessary diagnostic data related to one or more problems. Users no longer have to search in different directory locations trying to gather all the relevant trace files and dump files needed for problem diagnosis by Oracle Support. By invoking IPS, all diagnostic data (traces, dumps, health check reports, SQL test cases, and more) pertaining to a critical error is automatically packaged into a zip file which can then be shipped to Oracle Support.

Database Instance: database > Support Workbench >
Package Details: Pkg_database_ORA_1578_060507182254

Summary

Status	Upload File Generated	Send to Oracle
Total Size (uncompressed)	86.03 MB	
Incremental Size (uncompressed)	12.53 MB	
Created	June 5, 2007 6:23:26 PM PDT	
Description	N/A	
Problems in Package	ORA 1578	
	ORA 603	
	ORA 600 [4136]	
	ORA 600 [4137]	
Incidents Previously Excluded by User	0	
Files Excluded by User	0	

[Incidents](#) **Files** [Activity Log](#)

View

Source Name	Size (MB)	Has User Data	Timestamp
Incident b_pmon_7330_i3633.trc	2.74	No	June 5, 2007 6:16:06 PM PDT
Incident b_pmon_7330_i3633.trm	0.66	No	June 5, 2007 6:16:06 PM PDT
Incident b_m000_19538_i3633_a.trc	0	No	June 5, 2007 6:16:08 PM PDT

Figure 2: Incident package details

Support Workbench

The Support Workbench is a facility in Oracle Enterprise Manager that enables you to interact with the new fault diagnostic infrastructure of Oracle Database 11g. With it you can investigate, report, and where appropriate, repair problems, all with an easy-to-use graphical interface. The Support Workbench provides a self-service means for you to package diagnostic data using IPS, obtain a support request number, and upload the IPS package to Oracle Support with a minimum of effort and in a very short time, thereby reducing time-to-resolution for problems. Note that all automated interactions with Oracle Support, such as support number creation or IPS package upload, requires Oracle Configuration Manager to be running at the database location.

Oracle Configuration Support Manager, a proactive automated support capability included in Oracle Premier Support, offers customers a simpler way to track, manage, and support your Oracle configurations while reducing the risk of unplanned system downtime.

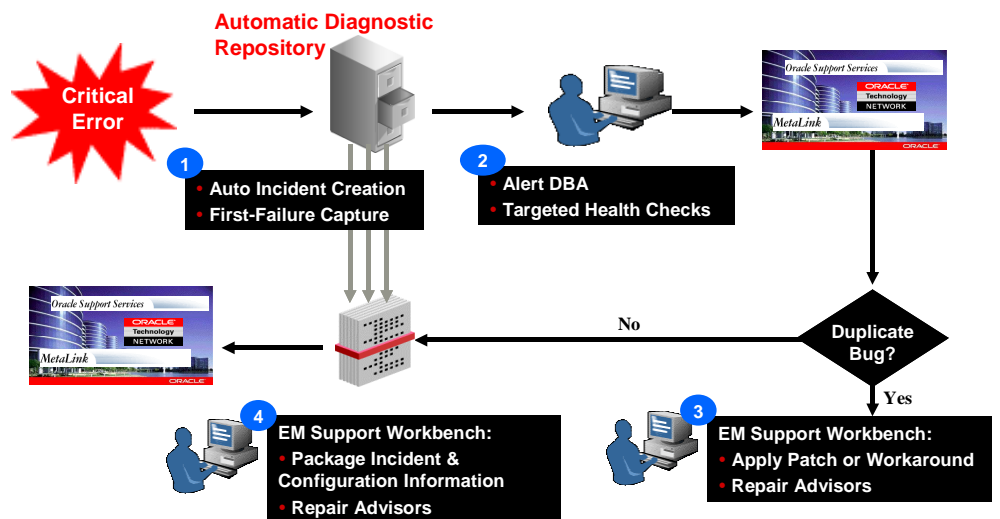


Figure 3: Support workbench workflow

The Support Workbench workflow consists of the following steps:

1. Create an incident in the database automatically based on the first occurrence of a failure.
2. Alert the DBA of the failure and run health checks in the areas where the failure was reported.
3. If it is a known issue, then recommend and apply patch to solve the problem.
4. Otherwise, package up incidents and relevant configuration information and upload to Oracle Support and run repair advisors to recover from failure.

There are many different kinds of problems that can occur in an Oracle Database and the right remedy for each problem may be different. The Support Workbench has extensive workflows that guide the user to take action that is appropriate for the problem encountered.

CONCLUSION

The manageability and diagnosability enhancements in Oracle Database 11g allows database administrators to keep their systems performant and available, while providing higher quality of service to their users.



Oracle Database 11g: Manageability Overview
August 2007
Author: Jagan R. Athreya
Contributing Authors: Mughees Minhas

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

Copyright © 2007, Oracle. 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 is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.