Oracle Database 10g: The Self-Managing Database

An Oracle White Paper
Nov. 2003

ORACLE

Oracle Database 10g: The Self-Managing Database

Introduction 3
Key Manageability Challenges 3
Oracle Rises Up To the Challenge 4
- Install and Configuration 5
  - Faster, Lightweight Install 5
  - Simplified Database Creation 6
  - Automated, Flexible and Fault Tolerant Upgrade 6
  - Fewer, Simplified Initialization Parameters 7
- Data Loading & Archiving 7
  - Data Pump 7
  - Cross Platform Transportable Tablespaces 8
- Ongoing Database Administration 8
  - Intelligent Infrastructure 9
  - Performance Diagnostic and Troubleshooting 12
  - Application/SQL Tuning 14
  - Memory Management 18
  - Space Management 20
  - Storage Management 22
  - Backup & Recovery 23
  - Configuration Management 25

What does it mean to you? 26
Conclusion 27

INTRODUCTION

The Oracle database has always been known as the industry leader in terms of performance and reliability, while at the same time providing extremely rich functionality. Most successful businesses today use the Oracle database to power their mission critical applications. However, as the business environment becomes increasingly competitive, the challenge before today's enterprise is to manage their information technology (IT) infrastructure at the lowest possible cost without compromising service level objectives. This presents an interesting dilemma since as systems provide more and more capabilities system complexity can increase, and hence these systems can become expensive to manage and sustain.

Recognizing these demands, Oracle Database 10g introduces a sophisticated self-managing database that automatically monitors, adapts, and fixes itself. This paper provides a technical overview of these self-managing solutions that allow DBAs to become more productive, help their organizations reduce management costs, and scale to manage the Enterprise Computing Grid.

KEY MANAGEABILITY CHALLENGES

In today's around-the-clock economy, the importance of an efficient and reliable IT infrastructure for the success of an enterprise hardly needs any explanation. As businesses increasingly rely on this infrastructure to service customers and partners, and disseminate information among employees, computer systems are no longer peripheral to the management of core business. Even a temporary outage of a critical application can jeopardize the viability of the entire business. Revenue and customers may be lost, penalties could be owed and, the resulting unfavorable publicity can have a disastrous affect on customer confidence and corporate stock valuation. The effective management of the enterprise IT infrastructure is therefore key to the success of modern enterprises.

Recognizing this need, companies today are investing significant amounts of financial and administrative resources to ensure the best possible management of their IT infrastructure. However, organizations today are forced to deal with more data than ever, with customers and employees needing around the clock access to this data, from more places, via more type of devices. As IT vendors deliver increasingly sophisticated solutions to meet such exacting demands, the

task of systems management has never been more complex. Hiring highly skilled administrative staff to manage such complicated environments is an expensive proposition. This coupled with frequent shortage of experienced administrative personnel often results in spiraling management costs.

While adjusting to new technologies and business practices presents numerous technical challenges, businesses today are also faced with increased competition as the Internet has collapsed the world into a small global village. An online bookseller in Seattle faces as much competition from other booksellers in the US as it does from one in China. In order to maintain business profitability amidst ever-growing competitive pressure, corporations must minimize their operating expenses with system administration costs being no exception. This creates an interesting corporate challenge. Enterprises must manage their systems much more effectively than ever -- to ensure the highest performance, scalability and availability – but at a cost significantly lower than before.

ORACLE RISES UP TO THE CHALLENGE

Continuing the tradition of leading the industry in developing solutions that address customers' business challenges in the most effective manner, Oracle has made the improved manageability of its products a strategic development focus area. With more than 200 engineers and over 50% of products architects involved in designing and implementing a comprehensive solution, manageability is the single biggest development effort that has gone into producing a revolutionary, self-managing Oracle Database 10g. A dedicated team of Oracle's senior most developers has worked hard over last couple of years talking to a diverse group of customers to learn about their most pressing manageability challenges and developing ground breaking technologies to solve these challenges in the most effective manner. The scope of these development efforts spans the entire technology stack in order to provide customers with complete, end-to-end manageability solutions.

There are two aspects to Oracle's manageability strategy. Firstly, it seeks to make each of Oracle's products, particularly the database, as self-managing as possible so that they require minimal manual administration. The Oracle Database 10g is a giant step towards Oracle's vision of creating a self-aware, self-leaning and completely self-managing database. A monumental development effort has gone into simplifying every aspect of the Oracle Database 10g administration in order to serve the dual objective of enhancing administrator productivity and helping customers reduce their operational cost by 50%.

The Second part of Oracle's manageability strategy focuses on making the management of completely data center easier, scalable and more effective. Oracle Enterprise Manager provides this solution. Enterprise Manager is Oracle's single, integrated solution for administering and monitoring applications and systems that are based on the Oracle technology stack. The

next generation HTML based Enterprise Manager provides the ability to seamlessly manage hundreds and thousands of system located across organizational and geographical boundaries from a single point of control. Built with robust functionality for managing both small and large sets of systems, Enterprise Manager 10g automates critical operations to reduce task time and the risk of errors which increases as the number of systems goes up. Its robust grouping and task automation functionality provide core features that enable reliable, rapid, and secure automation of traditionally time-consuming, error prone tasks, such as application performance management, policy-based standardization and system provisioning.

These two pieces put together provide a complete, sophisticated and effective solution to address every gamut of datacenter management. In the section to follows, we will take an in-depth look at some of the major Oracle Database 10g revolutionary manageability advances.

Install and Configuration

Faster, Lightweight Install

Problems encountered during installation can be extremely frustrating. It is no exaggeration to say that the usability of a product is often gauged by how simple or complex its installation is. Oracle has, therefore, invested significant amount of development resources in simplifying the installation, configuration and upgrade tasks.

Installation time for the Oracle Database 10g server software has been cut in half to 20 minutes on a typical configuration. The most common installs can be performed from a single CD ROM eliminating the need for media changes. Furthermore, a standard installation of the database requires only 256 Mbytes of memory and about 1 GB of disk space.

Even greater improvements can be seen on the client side: An Oracle Database 10g client install can be performed in about 1 minute and consumes only about 70 Mbytes of disk space.

The Oracle Universal Installer (OUI) has been enhanced to automate all pre and post installation tasks. In Oracle Database 10g, OUI automatically checks the system prior to beginning the installation to ensure that the OS has been configured properly, the correct patches have been applied, and enough resources are available to guarantee the success of the installation process. If any problems are detected during the pre-install check, the installer will help the administrator fix the problem by recommending corrective steps.

Also, the installation process has been made completely self-contained to automatically set up the required infrastructure for routine monitoring and administration. A fully functional, feature rich Enterprise Manager Database Management Console is now automatically configured to allow administrators

to get started with database administrative tasks without undertaking any manual post-install configuration at all. The Enterprise Manager Database Console provides all essential functionality for managing a single database, including alert notification, job scheduling and software management. In addition, all Oracle server components such as the database, listener, management framework etc. will be configured for automated startup and shutdown.

For developers and Independent Software Vendors (ISVs) developing applications with embedded Oracle Database, silent mode install has been made more robust and user friendly.

Thanks to these implications, anyone with a basic familiarity with computers should be able to install and get started with the Oracle Database 10g.

Simplified Database Creation

The database creation process has also been simplified in Oracle Database 10g by enhancing the Database Creation Assistant (DBCA). DBCA now allows the administrator to create all possible configurations of the database be it a stand-alone database, a Real Application Cluster (RAC) database, or a standby database. During the database creation process, the DBCA guides administrators in setting up an automated disk based backup and registering the database with a LDAP server, if available. Databases created using the DBCA will, therefore, be fully setup and ready to use in all respect.

DBCA can also be run in the silent mode to allow creation of database using pre-defined inputs. This capability can be used by ISVs to create database transparently as a part of their application installation.

Automated, Flexible and Fault Tolerant Upgrade

Similarly, Database Upgrade Assistant (DBUA) has been enhanced to make the database upgrade/migration process extremely simple. Using Oracle Database 10g DBUA, administrators will be able to migrate/upgrade any database configuration, including RAC and standby, just by answering a few simple questions. The tool will automatically check the system if adequate resources are available, ensure adherence to the best practices – such as backing up the database before beginning the upgrade/migration process, replacing the obsolete and deprecate initialization parameters etc – and, verify the successful completion of the operation. Starting with Oracle Database 10g, the upgrade process has been made automatically restartable allowing it to automatically resume from the point of interruption. Another significant upgrade enhancement in Oracle Database 10g is the ability to get time estimation of how long the upgrade process is likely to take. Since the time required to upgrade a database could be quite long (depending on the amount data dictionary meta data to be upgraded) and varies from one system to another, this capability will be of immense use to DBAs in planning the upgrade process.

Since the real benefit of upgrading to a newer release can only be derived by using new features and functionality, the Oracle Database 10g also helps administrators in activating the new features available in the release being upgraded to.

Fewer, Simplified Initialization Parameters

The Oracle database server provides a number of initialization parameters to optimize its operation in diverse environments. Only a few of these parameters need to be explicitly set as the default values for the rest of them are adequate in vast majority of cases. In Oracle Database 10g the parameters have been categorized into basic and advanced categories. Administrators will be able to restrict their day-to-day interaction with just with 28 basic parameters. The advanced parameters have been preserved to allow expert DBAs to adapt the behavior of the Oracle Database to meet unique requirements without overwhelming those who have no such requirements. The Oracle Database 10g, therefore, provide the right blend of simplicity and flexibility – simple enough to be used as an embedded database yet flexible enough to meet even the most challenging requirements.

Data Loading & Archiving

Data Pump

To enable very high-speed data and metadata loading and unloading to/from the Oracle Database, the Oracle Database 10g introduces a new facility called Data Pump. It automatically manages and schedules multiple, parallel streams of load or unload for maximum throughput. Data Pump infrastructure is callable via the PL/SQL package DBMS_DATAPUMP. Thus, custom data movement utilities can be built using Data Pump. Oracle Database 10g comes with four of these: New command line export and import clients (expdp & impdp), a web-based Enterprise Manager export / import interface and a custom interface for the movement of complex Data Mining models. Data Pump is also the foundation for several other key features in the Oracle server: Streams-based Replication, Logical Standby, and Transportable Tablespaces.

Data Pump dramatically decreases the elapsed time for most large export / import operations. A single thread of Data Pump's direct path data unload is about twice as fast as regular direct path export. A single thread of Data Pump data load is 15X – 45X faster than original imp. And of course, Data Pump operations can be specified with parallel threads of execution.

In addition to increased performance, Data Pump-based export and import clients support all the functionality of the original exp/imp as well as many new features such as checkpoint restart, job size estimation, very flexible, fine-grained object selection, direct loading of one instance from another and detailed job monitoring.

Cross Platform Transportable Tablespaces

The Transportable Tablespace feature, introduced in Oracle8i, allows users to quickly move a tablespace across Oracle databases. It is the most efficient way to move bulk data between databases.

Moving data using transportable tablespaces can be much faster than performing either an export/import or unload/load of the same data. This is because transporting a tablespace only requires the copying of datafiles and integrating the tablespace structural information. You can also use transportable tablespaces to move index data, thereby avoiding the index rebuilds you would have to perform when importing or loading table data.

- Oracle Database 10g provides the ability to transport tablespaces across platforms. This functionality can be used to:
- Provide an easier and more efficient means for content providers to publish structured data and distribute it to customers running Oracle on a different platform
- Simplify the distribution of data from a data warehouse environment to data marts which are often running on smaller platforms
- Enable the sharing of read only tablespaces across a heterogeneous cluster
- Allow a database to be migrated from one platform to another

The new Data Pump functionality together with cross-platform transportable tablespace feature provides Oracle DBAs with powerful, easy to use and high performance tools for moving data in and out of the database.

Ongoing Database Administration

Ongoing database administration tasks, such as performance tuning, space management, system resource tuning and backup & recovery, accounts for the biggest chunk of a database administrator's time. According to a survey conducted by Oracle, DBAs typically spend about 55% of their time performing these activities.

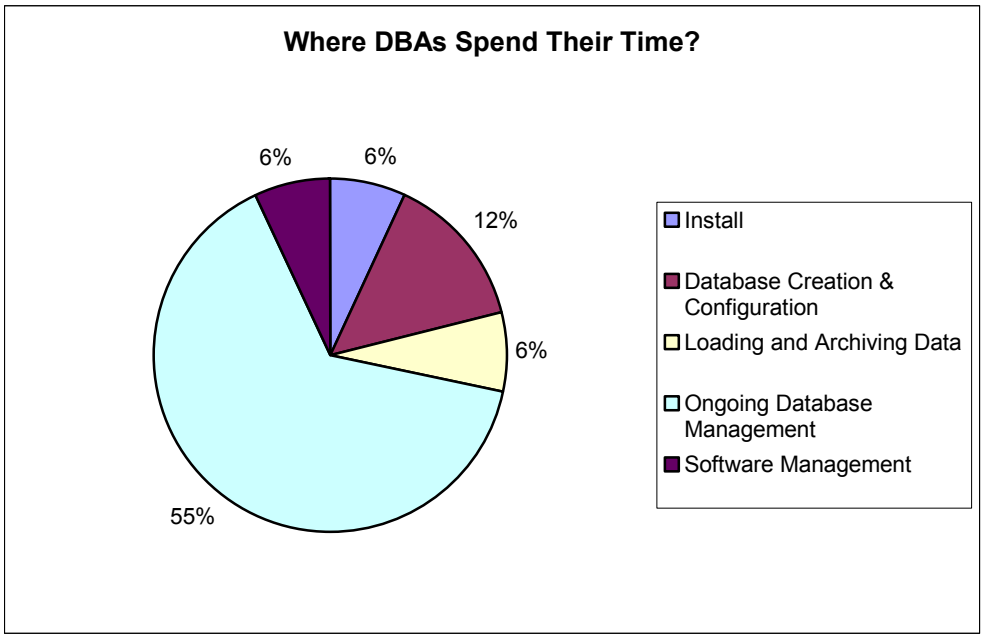

Fig 1: DBA Task-Time Distribution (Source: Survey conducted at IOUG Live! 2001)

One of the predominant development goals of the Oracle Database 10g has been to build solutions that automate and simplify the ongoing management tasks in order to dramatically reduce the time that DBAs have to spend on these activities. In order to accomplish this goal, Oracle Database 10g introduces a sophisticated self-management infrastructure that allows the database to learn about itself and use this information to adapt to workload variations or to automatically remedy any potential problem. This is one of the most crucial differentiating aspects of the manageability advances made by Oracle Database 10g. While almost every major software vendor these days claims to have simplified the use or management of their product, most of them use simple rules of thumb and oversimplify some of the difficult choices database administrators have to make on a day-to-day basis. Only the Oracle Database 10g has an intelligent self-management infrastructure integrated with the core database engine which allows it to make the self-management decisions at the same time while providing the regular services. The end result: the Oracle Database 10g can truly claim to be the only self-managing database product available today!

Intelligent Infrastructure

Oracle Database 10g's self-management infrastructure comprise of four components, Automatic Workload Repository, Automated Maintenance task infrastructure, Server Generated Alerts and Advisory Framework. Let us take a detailed look at each of these components.

Automatic Workload Repository

As may be evident from its name, the Automatic Workload Repository (AWR) is a built-in repository within every Oracle Database 10g database that contains operational statistics about that particular database and other such information. At regular intervals, the Oracle Database 10g makes a snapshot of all its vital statistics and workload information and stores them in AWR. By default, the snapshots are made every 30 minutes but the administrator has the option to change this frequency. The snapshots are stored in the AWR for a certain period of time (7 days by default) after which they are automatically purged. AWR is designed to be lightweight and manage itself completely in order to ensure that it does not impose any additional management burden on administrators.

AWR captures all of the data previously captured by Statspack and more. The data captured allows both system level and user level analysis to be performed, again reducing the requirement to repeat the workload in order to diagnose problems. Optimizations have been performed to ensure that the capture of data is performed efficiently to minimize overhead. One example of these optimizations is in the SQL statement capture. Working within the database we maintain deltas of the data for SQL statements between snapshots. These allow the Oracle Database 10g to capture only statements that have significantly impacted the load of the system since the previous snapshot in an efficient manner, rather than having to capture all statements that had performed above a threshold level of work since they first appeared in the system, as was previously the case. This both improves the performance of the SQL capture and greatly reduces the number of SQL statements that are captured over time.

AWR forms the foundation for all self-management functionality of Oracle Database 10g. It is the source of information that gives the Oracle Database 10g an historical perspective on how it is being used and enables it to make decisions which are accurate and specifically tailored for the environment that system is operating in.

Automatic Maintenance Tasks

AWR provides the Oracle Database 10g a very good "knowledge" of how it is being used. By analyzing the information stored in AWR, the database can identify the need of performing routine maintenance tasks, such as optimizer statistics refresh, rebuilding indexes, etc. The Automated Maintenance Tasks infrastructure enables the Oracle Database to automatically perform those operations. It uses the rich scheduling functionality introduced in the Oracle Database 10g (through a new feature called Unified Scheduler) to run such tasks in a pre-defined “Maintenance Window”. By default, the maintenance window starts at 10 PM every night and lasts till 6 AM next morning and throughout the weekend. All attributes of the “Maintenance Window” are customizable, including start/end time, frequency, days of the week, etc., allow it to be customized to the environment specific needs. Also, the impact of automated

maintenance tasks on normal database operations can be limited if desired by associating a Database Resource Manager resource plan to the maintenance window.

In Oracle Database 10g, the optimizer statistics are automatically refreshed using the Automatic Maintenance Task infrastructure. It will be used to automate more such tasks in the future.

Server Generated Alerts

For the problems that cannot be resolved automatically and require administrators to be notified, such as running out of space, the Oracle Database 10g contains an new self-management infrastructure component called Server Generated Alerts. As may be obvious from its name, the Oracle Database 10g has the ability to monitor itself and send out alerts to notify DBAs of any impending problem in an efficient and timely manner. Since the monitoring activities take place at the same time as the database performs its regular operation, it is much more efficient reducing the monitoring resource overhead to negligible. It also ensures that the database is aware of the problem the moment they arise enabling it to notify administrators just in time. This is in sharp contrast to the mechanism employed by monitoring tools available today, all of which regularly poll the database to evaluate the alert condition consuming a significant amount of system resources. To add to that, since the polling is done at a pre-set interval, it may cause the delay in problem detection and eventually in alert generation.

The alerts produced by the Oracle Database 10g not only notify the problem, they also provide recommendation on how the problem being reported can be resolved. This ensures quicker problem resolution and helps prevent potential failures.

Advisory Framework

The Oracle Database 10g includes a number of advisors for different sub-systems in the database to automatically determine how the operation of the corresponding sub-components could be further optimized. SQL Tuning and SQL Access Advisor, for example, provide recommendations for running SQL statements faster. Then there are memory advisers which help size the various memory components without resorting to trial-and-error techniques. There's also a Segment Advisor, which handles all space-related issues, such as recommending wasted-space reclamation, predicting the sizes of new tables and indexes, and analyzing growth trends, and an Undo Advisor, which lets you size the undo tablespace.

In order to ensure the consistency and uniformity in the way advisors function and allow them to interact with each other seamlessly, the Oracle Database 10g includes an Advisory Framework. The advisor framework provides a consistent manner in which advisors are invoked and how the results are reported. Even

though these advisors are primary used by the database itself to optimize its own performance, they can also be invoked by administrators to get more insight into the functioning of a particular sub-component. Having a uniform and consistent interface makes easier for administrators to use these advisors and use the generated information to get a better understanding of their system.

Performance Diagnostic and Troubleshooting

Building upon the data captured in AWR, Oracle Database 10g includes a self-diagnostic engine called the Automatic Database Diagnostic Monitor (ADDM). ADDM makes it possible for the Oracle Database 10g to diagnose its own performance and determine how any identified problems could be resolved. ADDM runs automatically after each AWR statistics capture and making the performance diagnostic data readily available.

ADDM examines data captured in AWR and performs analysis to determine the major issues on the system on a proactive basis and in many cases recommends solutions and quantifies expected benefits. ADDM takes a holistic approach to the performance of the system, using time as a common currency between components. The goal of ADDM is to identify those areas of the system that are consuming the most 'DB time'. ADDM drills down to identify the root cause of problems rather than just the symptoms and reports the impact that the problem is having on the system overall. If a recommendation is made it reports the benefits that can be expected, again in terms of time. The use of time throughout allows the impact of several problems or recommendations to be compared. Previously many problems have been identified based on value judgments and experience rather than quantifiable impacts. A good example of this is a system that is experiencing a high logon rate. A rule of thumb might have said that a logon rate of greater than 10 per seconds was a problem and should be fixed. However many systems can run significantly higher logon rates without it noticeably affecting performance. Using the new time distribution data in AWR, ADDM can report quantitatively that logons are taking 20% of time spent in the database. This quantified value can make it much easier to convince whoever needs to do the work to fix the problem or arrange for it to be fixed, rather than just making a statement such as 'I think that you are doing too many logons'.

ADDM starts its analysis by focusing on the activities that the database is spending most time on and then drills down through a sophisticated problem classification tree. The problem classification tree used by ADDM encapsulates decades of performance tuning experience in the Oracle Server Technologies Performance Group at Oracle HQ and other performance experts. Many of these classification rules have also been exercised in an Oracle internal tool that has been used by the Oracle Support organization for processing Statspack files for more than a year. In developing the classification tree, the prime intent was to handle the most frequently seen problems and to drill down to the root causes of

problems rather than just reporting symptoms. Some of the common problems detected by ADDM include:

- CPU bottlenecks
- Poor connection management
- Excessive parsing
- Lock contention
- IO capacity
- Under sizing of Oracle memory structures e.g. PGA, buffer cache, log buffer
- High load SQL statements
- High PL/SQL and Java time
- High checkpoint load and cause e.g. small log files, aggressive MTTR setting
- RAC specific issues

For a more complete list of problems detected by ADDM as well as to learn more details about how it works, please refer to Oracle white paper "The Self-Managing Database: Automatic Performance Diagnosis".

Besides reporting the potential performance issues, ADDM also documents the non-problem areas of the system. The sub-components, such as IO, memory, etc, that are not significantly impacting the performance of the system are pruned from the classification tree at an early stage and are listed so that the DBA can quickly see that there is little to be gained by performing actions in those areas. Again this saves time and wasted effort (both human and hardware) fixing things that will not impact the system performance overall.

The Oracle Database 10g is the first database product to introduce such a revolutionary self-diagnostic capability and will completely redefine the database administration landscape. Administrators no longer need to first collect huge volume of diagnostic data and spend endless hours analyzing them in order to find out answers to performance issues. With the Oracle Database 10g, they can simply ask the database what the performance issues and the ADDM does the rest. They can sit back, relax, follow the recommendation made by ADDM using just a few mouse click and forget about the rest. The following table technically illustrates this changed paradigm by comparing the steps required to diagnose a common performance problem, excessive hard parsing.

Task: Diagnosing and resolving hard parse problems

Before Oracle Database 10g	Oracle Database 10g
1.Examine system utilization 2.Look at wait events 3.Observe latch contention 4.See wait on shared pool and library cache latch 5.Review v$sysstat (difficult)6.See "parse time elapsed" > "parse time cpu" and #hard parses greater than normal 7.Identify SQL by • Identifying sessions with many hard parses and trace them, or • Reviewing v$sql for many statements with same hash plan (difficult) 8.Examine objects accessed and review SQL 9.Identify "hard parse" issue by observing the SQL contains literals 10.Enable cursor sharing	1.Review ADDM recommendations 2.ADDM recommends use of cursor_sharing

Table 1: Performance Diagnostic and Resolution: Before and After Oracle Database 10g

Application/SQL Tuning

Application design issues are the most predominant cause of performance problems. All the tuning wisdom of developers, DBAs and system administrators cannot make up for the slowdown caused by architectural and design deficiencies of the application. One important part of database system performance tuning is, therefore, the tuning of SQL statements.

Since the query optimizer takes crucial decisions which have a tremendous impact on the performance of a query such as, whether to use an index or not, which join techniques to use if the query involves joining multiple tables, etc., Oracle has invested considerable development effort in making the Cost Based Optimizer the industry's most sophisticated, mature, and thoroughly tested query optimizer. The cost-based optimizer is also used extensively by the major packaged applications such as Oracle Application, SAP, PeopleSoft, etc. Since the Oracle database is the platform for the vast majority of the customers using these applications, this demonstrates the success of the Oracle optimizer in a

huge number of real application settings. Consequently, in Oracle Database 10g, the Rule Based Optimizer (RBO) is no longer available and the cost-based optimizer is the only supported optimizer mode in this release. In the rest of the paper, the word "optimizer" will be used synonymously with the cost-based optimizer.

While the Oracle Database provides best possible query optimization technology which maximizes the application/query performance without any administrator intervention in majority of the cases, there may still be a few cases where the nature of the application or uniqueness of data distribution may cause certain SQL statements to consume unusually high percentage of total system resources. In such situation, the SQL tuning process involves three basic steps:

- Identify high load or top SQL statements that are responsible for a large share of the application workload and system resources, by looking at the past SQL execution history available in the system (e.g., the cursor cache statistics stored in the V$SQL dynamic view),
- Verify that the execution plans produced by the query optimizer for these statements perform reasonably well,
- Take possible corrective actions to generate better execution plans for poorly performing SQL statements.

The three steps are repeated until the system performance reaches a satisfactory level or no more statements can be tuned. Besides being extremely time consuming, the SQL tuning process outlined above also requires a high degree of expertise. Only a person with a deep knowledge of the application and database system can undertake this task.

The Oracle Database 10g completely automates the SQL tuning process. ADDM identifies SQL statements that are consuming unusually high system resources and are therefore causing performance problem. In addition, the top SQL statements in terms of CPU and shared memory consumption are automatically captured in AWR. Thus, the identification of high load SQL statements happens automatically in the Oracle Database 10g and requires ABSOLUTELY NO intervention from the administrator.

After having automatically identified the top resource consuming SQL statements, the Oracle Database10g can automatically analyze them and recommend solutions using newly added automatic tuning capability of the query optimizer, called the Automatic Tuning Optimizer. The Automatic SQL Tuning is exposed via an advisor called the SQL Tuning Advisor. The SQL Tuning Advisor takes one or more SQL statements, and produces well-tuned plans along with tuning advice. The administrator does not need to do anything other than just invoking the SQL tuning advisor. Once that is done, optimizer analyzes the problem SQL statement(s) and recommends the solution. It is important to bear in mind here that the solution is coming right from the

optimizer and not from any external tools using some pre-defined heuristics. This provides several advantages: a) the tuning is done by the system component that is ultimately responsible for the execution plans and hence the SQL performance, b) the tuning process is fully cost-based and it naturally accounts for any changes and enhancements done to the query optimizer, c) the tuning process takes into account the past execution statistics of a SQL statement and customizes the optimizer settings for that statement, and d) it collects auxiliary information in conjunction with the regular statistics based on what is considered useful by the query optimizer.

The recommendation of the Automatic Tuning Optimizer can fall into one of the following categories

Statistics Analysis: The Automatic Tuning Optimizer checks each query object for missing or stale statistics, and makes recommendation to gather relevant statistics. It also collects auxiliary information to supply missing statistics or correct stale statistics in case recommendations are not implemented. Since the Oracle Database 10g automatically gathers optimizer statistics, this should not usually be the problem unless the automatic statistics gathering functionality has been disabled for some reason.

SQL Profiling: The Automatic Tuning Optimizer verifies its own estimates and collects auxiliary information to remove estimation errors. It also collects auxiliary information in the form of customized optimizer settings (e.g., first rows vs. all rows) based on past execution history of the SQL statement. It builds a SQL Profile using the auxiliary information and makes a recommendation to create it. When a SQL Profile is created it enables the query optimizer (under normal mode) to generate a well-tuned plan. **The most powerful aspect of SQL profiles is that they enable tuning of queries without requiring any syntactical changes and thereby proving Oracle administrators and customers with an unique database –resident solution to tune the SQL statements embedded in packaged applications.**

Access Path Analysis: The Automatic Tuning Optimizer explores whether a new index can be used to significantly improve access to each table in the query, and when appropriate makes recommendations to create such indexes.

SQL Structure Analysis: Here the Automatic Tuning Optimizer tries to identify SQL statements that lend themselves to bad plans, and makes relevant suggestions to restructure them. The suggested restructurings can be syntactic as well as semantic changes to the SQL code.

Both Access Path and SQL structure analysis can be immensely useful in tuning the performance of an application under development or a homegrown production application where the administrators/developers have access to application code.

SQL Access Advisor is yet another major Oracle Database 10g manageability enhancement. It can automatically analyze the schema design for a given workload and recommend indexes and materialized views to create, retain or drop as appropriate for the workload. While generating recommendations, the SQL Access Advisor considers the impact of adding new indexes and materialized views on data manipulation activities, such as insert, update and delete, in addition to the performance improvement they are likely to provide for queries. The SQL Access Advisor provides a very easy to use interface and requires very little system knowledge. It can also be run without affecting production systems since the data can be gathered from the production system and taken to another machine where the SQL Access Advisor can be run.

The Oracle Database 10g indeed makes optimizing application performance child's play. No matter how complex the problem is or how restrictive is the environment, the Oracle Database 10g can provide the answers to SQL tuning problems. The following table depicts how the Automatic Tuning Optimizer dramatically alters the way SQL problems are detected and resolved.

Task: Tuning high load SQL statements coming from a packaged application

Before Oracle Database 10g	Oracle Database 10g
1.Examine system utilization 2.Look at wait events 3.See wait on DB scattered read 4.Determine scope – system wide, module-dependent, user-dependent? 5.Identify SQL by (difficult) •Identifying sessions with high DB scattered read waits and trace them, or •Reviewing Top Sessions in OEM 6.Get explain plan 7.Examine objects accessed (size/cardinality) 8.Review SQL statistics and/or compare to object statistics (v$sql) (difficult) 9.Identify the problem 10.Contact packaged app vendor 11.Produce test case for vendor 12.Vendor produces patch/upgrade	1.Review ADDM recommendations 2.Follow link to run Automatic SQL tuning 3.Accept SQL Profile recommendations from SQL Tuning

Before Oracle Database 10g	Oracle Database 10g
13.Patch/upgrade installed in customer's next maintenance cycle	

Table 2: SQL Tuning: Before and After Oracle Database 10g

Those interested in finding out more details about the SQL Tuning Advisor and SQL Access Advisor can take a look at the following papers:

- "The Self-Managing Database: Guided Application & SQL Tuning"
- "Turbocharge your Database: Use the Oracle database 10g SQL Access Advisor"

Memory Management

Memory is a precious system resource and administrators currently spend a significant amount of their time optimizing its use. One of the key self-management enhancements in the Oracle Database 10g is Automatic Shared (SGA) Memory management. This functionality automates the management of shared memory used by an Oracle Database instance and liberates administrators from having to configure the shared memory components manually. In Oracle Database 10g, DBAs can just specify the total amount of SGA memory available to an instance using a newly used parameter SGA_TARGET. The database server then automatically distributes the available memory among various components as required. The Automatic Shared Memory Management feature is based on sophisticated heuristics internal to the database that monitors the memory distribution and changes it according to the demands of the workload.

When the Automatic Shared Memory Management is enabled, the most commonly configured components are sized automatically. These include:

1. Shared pool (for SQL and PL/SQL execution)

2. Java pool for (java execution state)

3. Large pool (for large allocations such as RMAN backup buffers)

4. Buffer cache

There is no need to set the of size any of the above components explicitly and by default the parameters for these components will appear to have values of zero. Whenever a component needs memory, it can request that it be transferred from another component via the internal auto-tuning mechanism. This will happen transparently without user-intervention.

The performance of each of these components is also monitored by the Oracle instance. Now the instance uses internal views and statistics to determine how to optimally distribute memory among the automatically sized components. Thus,

as the workload changes, memory would be redistributed to ensure optimal performance with the new workload. This algorithm is never complacent and always tries to find the optimal distribution by taking into consideration long term as well as short terms trends.

The administrator can still exercise some control over the size of the auto-tuned components by specifying minimum values for each of these components. This can be useful in cases where the administrator knows that an application needs a minimum amount of memory in certain components to function properly.

The sizes of the automatically tuned components are remembered across shutdowns if a server parameter file (SPFILE) is used. This means that the system will not need to learn the characteristics workload from scratch each time and will pick up where it left off from the last shutdown.

The most significant benefit of using automatic SGA memory management is that the sizes of the different SGA components are flexible and will adapt to the needs of a workload without requiring user intervention. Besides maximizing the use of available memory, the Automatic Shared Memory Management feature can enhance workload performance as well. With manual configuration, it is possible that the compiled SQL statements will frequently age out of the shared pool because of its inadequate size. This will manifest into frequent hard parses and, hence, reduced performance. However when automatic management is enabled, the internal tuning algorithm will monitor the performance of the workload and grow the shared pool if it determines that doing so will reduce the number of parses required. This is one of the most wonderful aspects of Automatic Shared Memory Management feature since it provides enhanced out-of-box performance, without requiring any additional resources or manual tuning effort. Having just a single parameter to deal with simplifies the job of administrators greatly. DBAs can now just specify the amount of SGA memory an instance has its disposal and forget about the rest. They do not need to figure out the sizes of individual components any more. In addition, they can be assured of the fact that no out of memory errors will be generated unless the system has truly run out of memory.

Task: Resolving out of memory errors

Before Oracle Database 10g	**Oracle Database 10g**
1.Launch Buffer Pool Advisor 2.Examine output; check if Buffer Pool is over allocated 3.If so, reduce size of Buffer Pool 4.Launch Shared Pool Advisor 5.Examine output; check if Shared Pool is under allocated	(This space is intentionally left blank – No manual steps needed with Automatic Shared Memory Tuning)

6.If so, increase size of Shared Pool	

Table 3: SGA Memory Management: Before and After Oracle Database 10g

To find more technical details about the Automatic Shared Memory Management feature, please take a look at Oracle white paper "The Self-Managing Database: Automatic SGA Memory Management".

Space Management

Space management is one of the most time consuming tasks for database administrators. Fortunately, the Oracle Database 10g automatically manages its space consumption, alerts administrators on potential space problems, and recommends possible solutions.

Proactive Space Management

Oracle Database 10g introduces a non-intrusive and timely check for space utilization monitoring in the database server. The Oracle Database 10g automatically monitors its space utilization during its normal space allocation and de-allocation operations and alerts administrators if the free space availability falls below the pre-defined thresholds. Oracle Database 10g's space monitoring functionality is set up out of box, causes no measurable performance impact, and is uniformly available across all tablespace types. Also, the same functionality is available both through Enterprise Manager (EM) as well as SQL. Since the monitoring is performed at the same time as space is allocated and freed up in the database server, this guarantees immediate availability of space usage information whenever the user needs it.

Notification is performed using server generated alerts mechanism. The alerts are triggered when certain space related events occur in the database. For example when the space usage threshold of a tablespace is crossed, an alert is raised. Another example of an alert is when a Resumable session encounters out of space situation. An alert is sent instantaneously to the DBA to take corrective measures. The DBA may choose to get paged with the alert information and add space to the tablespace to allow the suspended operation to continue from where it left off.

The database comes with a default set of alert thresholds. The DBA may override the default for a given tablespace or set a new default for the entire database through EM.

Intelligent Capacity Planning

DBAs face a lot of challenges in managing space allocated to segments. Space may get over allocated because of the difficulty to predict the space requirement of an object or because of the inability to predict the growth trend of an object. On tables that are heavily updated, the resulting segment may have a lot of internal fragmentation and maybe even row chaining. These issues may result in

a wide variety of problem symptoms in the database from poor performance to space wastage. The Oracle Database 10g offers several new features/functionality to address these challenges.

The Oracle Database 10g can predict the size of a given table based on its structure and estimated number of rows. This is a powerful "what if" tool that allows estimation of the size of an object before it is created or rebuilt. If tablespaces have different extent management policies then the tool will help decide the tablespace that will cause least internal fragmentation. The Object Space Advisor works for tables with indexes on them as well as Index Organized Tables.

The growth trend report takes the DBA to the next step of capacity planning – planning for growth. Most database systems grow over time. Planning for growth is an important aspect of provisioning resources. In order to aid this process, the Oracle Database 10g tracks historical space utilization in the AWR and uses this information to predict the future resource requirements.

Transparent Space Reclamation

The Oracle Database 10g provides the ability of performing *an in place reorganization* of data for optimal space utilization by shrinking it. Shrinking of a segment will make unused space available to other segments in the tablespace and may improve the performance of queries and DML operations.

The Segment Shrink functionality provides the ability to both compact the space used in a segment and then deallocate it from the segment. The deallocated space is returned to the tablespace and is available to other objects in the tablespace. Sparsely populated tables may cause a performance problem for full table scans. By performing shrink, data in the table is compacted and the high water mark of the segment is pushed down. This makes full table scans read less blocks and hence, run faster.

Segment Shrink is an *online* operation – the table being shrunk is open to queries and DML while the segment is being shrunk. Additionally, segment shrink is performed *in place*. This is a key advantage over performing Online Table Redefinition for compaction and reclaiming space. The DBA may schedule segment shrink for one or all the objects in the database as nightly jobs without requiring any additional space to be provided to the database.

Segment Shrink works on heaps, IOTs, LOBs, Materialized Views and Indexes with row movement enabled in tablespaces with Automatic Segment Space Management. When Segment Shrink is performed on tables with indexes on them, the indexes are automatically maintained when rows are moved around for compaction. User defined triggers are not fired, however, because compaction is a purely physical operation and does not impact the application. Note that Segment Shrink can be performed only on tables with row movement enabled. Only applications that explicitly track rowids of objects will not be able to be

shrunk because the application tracks the physical location of rows in the objects.

In order to easily identify candidate segments for shrinking, the Oracle Database 10g also includes Segment Advisor. The Segment Advisor performs growth trend analysis on individual objects to determine if there will be any additional space left in the object in 7 days. It then uses the reclaim space target to select candidate objects to shrink. The Segment Advisor can be invoked in the *comprehensive* mode. In this mode, in addition to using the pre-computed statistics in the workload repository, the Segment Advisor performs sampling of the objects under consideration to refine the statistics for the objects. Although this operation is more resource intensive, it may be used to perform a more accurate analysis, when desired.

Oracle White Paper "The Self-Managing Database: Proactive Space and Schema Object Management" contains detailed information about the space management enhancements introduced in the Oracle Database 10g.

Storage Management

Storage layout and configuration can have a significant impact on overall database performance. Database Administrators, in conjunction of storage and network administrators, therefore have to deal with difficult tasks such as looking for hotspots that negatively affect performance and moving data files around to reduce contention. Thankfully, that scenario is about to change dramatically, with Oracle Database 10*g*'s new Automatic Storage Management (ASM) capability.

Automatic Storage Management provides a vertical integration of the file system and volume manager specifically built for the Oracle database files. ASM distributes I/O load across all available resource to optimize performance while removing the need for manual I/O tuning (spreading out the database files avoids hotspots). ASM helps DBAs manage a dynamic database environment by allowing them to grow the database size without having to shutdown the database to adjust the storage allocation.

Automatic Storage Management allows the DBA to define a pool of storage (called a disk groups) and then the Oracle kernel manages the file naming and placement of the database files on that pool of storage. The DBA can change the storage allocation (adding or removing disks) with new SQL commands (create diskgroup, alter diskgroup and drop diskgroup). One can also manage the disk groups by using Enterprise Manager (EM) and the Database Configuration Assistant (DBCA).

The Oracle Database 10*g* provides DBAs a simplified management interface for storage resources. Automatic Storage Management eliminates the need for manual I/O performance tuning. It virtualizes storage to a set of disk groups and provides redundancy options to enable a high level of protection. ASM

facilitates non-intrusive storage configuration changes with automatic rebalancing. It spreads database files across all available storage to optimize performance and resource utilization. It is a capability that saves DBAs time by automating manual storage and thereby increasing their ability to manage larger databases and more of them with increased efficiency.

Backup & Recovery

Oracle's current backup & recovery infrastructure is already considered the most robust and powerful in the industry. Oracle Recovery Manager (RMAN) is a powerful tool that simplifies, automates, and improves the performance of backup and recovery operations. It continues to mature with each new release and the revolutionary technological advances included in Oracle Database 10g make it even more simple, reliable, and automated.

Easier, Faster On-Disk Backups

As stated earlier in the paper, DBCA now provides the ability to automatically schedule an on-disk backup procedure. All that DBAs need to do is to specify the time window when they want the automatic backups to run. This automation leverages a number of new functionality introduced in the Oracle Database 10g to simplify disk-based backups. In Oracle Database 10g, a unified storage location for all recovery related files and activities in an Oracle database, called the Flash Recovery Area, may be defined using a new initialization parameter DB_RECOVERY_FILE_DEST. All files that are needed to completely recover a database from a media failure, such as control files, archived log files, flashback logs (introduced in the Oracle Database 10g), RMAN backups, etc., are part of the Flash Recovery Area. Allocating sufficient space to the Flash Recovery Area will ensure faster, simpler, and automatic recovery of the Oracle database. The Flash Recovery does a lot more than just providing a unified location for recovery related file. It actually manages the files stored in this location in an intelligent manner to maximize the space utilization and avoid out of space situations to the extent possible. Based on the specified RMAN RETENTION POLICY, the Flash Recovery Area will automatically delete obsolete backups and archive logs that are no longer required based on that configuration.

Incremental backups have been part of RMAN since it was first released in Oracle8.0. Incremental backups provide the capability to backup only the changed blocks since the previous backup. Oracle Database 10g delivers the ability for faster incrementals with the implementation of the change tracking file feature. When the block change tracking feature is enabled, Oracle tracks the physical location of all database changes. RMAN automatically use the change tracking file to determine which blocks need to be read during an incremental backup and directly accesses that block to back it up. It reduces the amount of time needed for daily backups, saves network bandwidth when backing up over a network and reduces the backup file storage.

Besides making the incremental backups significantly faster, the Oracle Database 10g also allows them to be used for updating a previously made backup. Oracle's Database 10g Incrementally Updated Backups feature merges the image copy of a datafile with a RMAN incremental backup resulting in an updated backup which contains the changes captured by the incremental backup. By providing this yet another unique functionality, the Oracle Database 10g eliminates the requirement to make a whole database backup repeatedly. It should now be possible to make a full database backup once for a given database and use incremental backups subsequently to keep the full back up updated. A backup strategy based on incrementally updated backups can help keep the time required for media recovery of your database to a minimum.

Self-Service Error Correction

According to many studies, 40% of application outages are caused by operator or user errors. Part of being human is making mistakes. But these errors are extremely difficult to avoid and in can be particularly difficult to recover from without advance planning and the right technology.

The Oracle Database 10g architecture leverages the unique technological advances in the area of database recovery due to human errors. Flashback Technology provides a set of new features to view and rewind data back and forth in time. The Flashback features offer the capability to query past versions of schema objects, query historical data, perform change analysis or perform self-service repair to recover from logical corruptions while the database is online. With the Oracle Database 10g Flashback Technology, you can indeed undo the past!

Flashback Technology revolutionizes recovery by just operating on the changed data. The time it takes to recover the error is now equal to the same amount of time it took to make the mistake. The Oracle Database 10g Flashback technology includes Flashback Database, Flashback Table, Flashback Drop, Flashback Versions Query, and Flashback Transaction Query.

Flashback Database quickly rewinds an Oracle database to a previous time, to correct any problems caused by logical data corruptions or user errors. It eliminates the down time associated with restoring a backup and makes it extremely easy to recover from unanticipated errors.

Flashback Table provides the DBA the ability to recover a table or a set of tables to a specified point in time quickly, easily, and online. Flashback Table restores the tables while automatically maintaining its associated attributes such as - the current indexes, triggers and constraints, not requiring the DBA to find and restore application specific properties. Flashback Table alleviates the need for performing more complicated point in time recovery operations.

Flashback Drop provides a safety net when dropping objects in Oracle Database 10g. When a user drops a table, Oracle automatically places it into the "Recycle

Bin". The Recycle Bin is a virtual container where all dropped objects reside. Objects remain in the Recycle Bin until either Oracle needs to reclaim the space in order to accommodate new data or the owner of the dropped objects decides to permanently remove them using the new PURGE command. As long as a dropped object remains in the recycle bin, it can be recovered using a simple SQL statement.

Task: Recovering accidentally dropped table

Before Oracle Database 10g	Oracle Database 10g
(Tablespace Point-in-time Recovery) 1. Prepare an auxiliary instance by first creating an Oracle password file 2.Create parameter file for auxiliary instance 3.Start auxiliary instance in NOMOUNT mode using SQL*Plus 4.Using RMAN interface to perform TSPITR 5.Using RMAN, connect to target database and bring tablespace in question online 6.Shutdown the auxiliary instance 7.Delete auxiliary instance data files, control files, and redo log files	1.Single Command Recovery: FLASHBACK TABLE <table_name> TO BEFORE DROP ;

Table 4: Database Recovery: Before and After Oracle Database 10g

Configuration Management

Enterprise Manager 10g has several powerful new configuration management facilities that help detect configuration changes and differences and enforce best practice configuration parameter settings. These capabilities are not limited to Oracle software only, they also encompass the underlying hosts and operating systems.

Enterprise Manager 10g continuously monitors the configuration of all Oracle systems for such things as best practice parameter settings, security set-up, storage and file space conditions and recommended feature usage. Non-conforming systems are automatically flagged to the user with a detailed explanation of the specific system configuration issue. For example, Enterprise Manager advises database administrators to use new functionality such Automatic Undo Management or Locally Managed Tablespaces if they aren't being used currently. This automatic monitoring of system configurations

promotes best practices configuration management, reduces administrator workload and the risk of availability, performance or security compromises.

Oracle Enterprise Manager 10g also significantly facilitates software patching thanks to its built-in *MetaLink integration*. First of all, Enterprise Manager automatically alerts users to new critical patches – such as important security patches – and flags all systems that require that patch. In addition, users can invoke the Enterprise Manager patch wizard to find out what interim patches are available for that installation. Using the Enterprise Manager patch wizard interim patches can be selected for download from MetaLink into the Enterprise Manager patch cache eliminating the need for repeated downloads. The user has the choice to stage appropriate patches on the destination system or systems for manual application at a later time. To further automate the patching process, the user can also provide a customizable patch application script that is executed on the destination systems at a user-defined time by the resident Enterprise Manager Agents. As patches are applied to a system, the corresponding OUI inventory is automatically updated to keep track of the systems' correct patch level.

WHAT DOES IT MEAN TO YOU?

A completely self-managing database will radically change the way enterprise data centers are managed in the future. The automation of routine administrative tasks will enable DBAs to concentrate their efforts in managing the end-user experience rather than supporting technology. As the database becomes totally self-contained, the role of database administrators will evolve from managing the database to strategic management of business information stored in it. With an increased focus on company's business goals, database administrators in the future will play an increasingly crucial role in success of the enterprise.

Businesses will reap the benefits of increased administrative productivity in form of significantly reduced operational costs. At the same time, systems will perform better, offer more reliability, and automatically scale to millions of users – all at a fraction of today's costs. The "invisible" Oracle database will thus enable enterprises to become more profitable, provide better customer service, and emerge as market leaders.

The self-management capabilities of the Oracle Database 10g will also facilitate its use as an embedded database. All the self-management functionality introduced in the Oracle Database 10g is fully accessible via open Application Programming Interface (APIs) making it extremely easy for third party application developers to create applications that automatically manage the underlying Oracle database. The unrivaled performance, availability, and scalability of the embedded Oracle database will enable such applications to offer highest class mission-critical services without exposing database management complexities.

CONCLUSION

Modern enterprises are aggressively adopting new technology solutions to enhance their competitiveness and profitability. The system management costs, however, have been steadily rising leading to eroded profit margins. The Oracle Database 10g is the answer to this challenge. It is industry first truly self-managing database featuring an intelligent self-management framework, revolutionary self-diagnostic engine, Automatic Tuning Optimizer and automatic memory management capability. The manageability advances made by Oracle Database 10g put it far ahead of its competition, will help customers in significantly lowering their management costs and enable ISVs to easily embed the industry leading Oracle Database in their application. With Oracle Database 10g, Oracle has once again lived up to its reputation of being a real pioneer!

White Paper Title: Oracle Database 10g: The Self-Managing Database
Nov 2003
Author: Sushil Kumar
Contributing Authors:

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
www.oracle.com

Oracle Corporation provides the software
that powers the internet.