



89 Fifth Avenue, 7th Floor
New York, NY 10003
www.TheEdison.com
212.367.7400

Comparative Management Costs Study

**Oracle Database 11g vs.
IBM DB2[®] Enterprise 9.5**

Printed in the United States of America.

Copyright © 2008 Edison Group, Inc. New York. Edison Group offers no warranty either expressed or implied on the information contained herein and shall be held harmless for errors resulting from its use.

All products are trademarks of their respective owners.

First Publication: January 2009

Produced by: Steve Mintz, Senior Analyst, Craig Norris, Editor; Barry Cohen,
Senior Analyst, Editor-in-Chief

Table of Contents

Executive Summary	1
About This Report	3
Who Should Read This Report	3
Methodology Overview.....	3
Contents of this Report	4
Rationale Behind this Comparison	6
The Methodology Defined.....	7
Workload Weighting.....	7
The Study.....	7
Tasks.....	7
Workload Weighting.....	10
Test Administration Workload.....	11
Test Results	13
Comparative Management Cost Savings.....	13
Workload Category Results Synopsis	14
Install DB/Software/Out-of-the-Box Setup (Results Synopsis)	14
Day-to-Day Database Administration (Results Synopsis).....	16
Backup and Recovery (Results Synopsis)	19
Performance Tuning (Results Synopsis)	23
Conclusion.....	27
Appendix I - Test Platform Details.....	28
Appendix II - What's New?.....	29
Appendix III - Architectural & Terminology Discrepancies Requiring Clarification .	35
Appendix IV - Detailed Test Results	36
Appendix V – Complexity Calculation Formula.....	38

Table of Contents - continued

Appendix VI - Detailed Task Descriptions	40
Installation and Simple “Out-of-Box” Setup	40
• Task 1: Install database and management software, create starter database	40
• Task 2: Create additional database server/instance	43
• Task 3: Set up proactive monitoring for performance and space utilization	45
Day-to-Day Database Administration.....	47
• Task 4: Create user, assign roles/privileges.....	47
• Task 5: Create Tablespace/Filegroup.....	48
• Task 6: Add more space to the database.....	49
• Task 7: Create table.....	50
• Task 8: Create index.....	51
• Task 9: Reclaim Lost Space due to Fragmented Data	52
• Task 10: Load data from a text file.....	53
• Task 11: Configure Adaptive Thresholds, Perform Workload Comparisons	55
Backup & Recovery	59
• Task 12: Configure and Perform Full Database Online Backup.....	59
• Task 13: Recover Dropped Table	59
• Task 14: Recover Data File	61
• Task 15: Recover from erroneous transaction.....	63
• Task 16: Recover from Multiple Failures.....	65
Performance Diagnostics & Tuning Tasks.....	68
• Task 17: Diagnose Performance Problem	68
• Task 18: Fix Performance Problem - SQL Tuning	69
• Task 19: Tune Memory.....	71
Appendix VII – Glossary of Task Areas/Tasks	72

Executive Summary

This report is the latest in a series of Comparative Management Cost Studies (CMCS) comparing database management systems offered by other leading enterprise software vendors. This edition of the study compares IBM DB2 Enterprise 9.5 with Oracle Database 11g.

Following the same approach taken in the earlier studies, Edison Group set up a laboratory environment to analyze a suite of standard RDBMS administrative tasks, and measured their respective management efficiency (time taken to complete tasks) and their complexity based on a proprietary manageability metric. Using the management efficiency results, Edison Group calculated the annual savings for businesses due to the enhanced DBA productivity that would result from using the product with superior manageability.

The study results show that, while both the vendor products have features to help the average database administrator (DBA) perform everyday administrative tasks, Oracle Database 11g holds substantial advantage over IBM DB2 Enterprise 9.5. Our detailed analysis reveals that:

- DBA's can perform typical administrative functions in 40 percent less time when using Oracle Database 11g compared to IBM DB2 Enterprise 9.5.
- Oracle Database 11g requires 49 percent fewer steps for the same set of standard RDBMS tasks than IBM DB2 Enterprise 9.5 using Edison's metric for complexity assessment.
- Benefiting from increased DBA productivity due to lower complexity and higher efficiency cited above, businesses could save up to \$35,155 per year per DBA by using Oracle Database 11g rather than IBM DB2 Enterprise 9.5.

The main areas responsible for the manageability difference between the two products were:

- **Performance Diagnostics and Tuning** – In this category, where DBAs spend considerable amount of time, Oracle Database 11g demonstrated a significant 88 percent savings in time and 93 percent fewer steps. Oracle's Automatic Database Diagnostic Monitor (ADDM) and Automatic SQL Tuning features were mainly responsible for this advantage.
- **Day-to-Day Administration** – In this category Oracle Database 11g provided 47 percent time savings and 37 percent lower complexity over IBM DB2 Enterprise 9.5.

The main contributor to this difference was DB2's lack of an easy-to-use tool for auto-configuring of thresholds for proactive system monitoring.

Oracle Database's management console – Oracle Enterprise Manager – is the sole management tool for managing Oracle Database. It can be accessed via a web browser from any authorized location. This was a key advantage over IBM DB2 Enterprise 9.5, which has several management consoles, each requiring a separate client installation on every machine to be used for database administration.

The study also highlights the incremental manageability changes in the two products when compared to the previous study. In our previous study comparing Oracle Database 10g R2 and IBM DB2 UDB 9.1, Oracle came ahead as well by a similar wide margin. In that study, Oracle Database 10g Release 2 had a 38 percent time advantage and 35 percent complexity differential over IBM DB2 UDB 9.1, compared to the 40 percent time and 52 percent complexity differential that Oracle Database 11g holds over IBM DB2 Enterprise 9.5.

The two key factors responsible for Oracle improving its lead over DB2 in this latest study were:

- Oracle Database 11g's improvement in managing day-to-day administration tasks,
- Oracle's further automation in the important tasks of SQL tuning and recovering from human induced failures.

There are numerous ways to interpret the significance of these savings, depending upon a variety of factors. These include the size of the organizations involved as well as the relative importance attached to higher productivity in these organizations and the need for superior performance for specific tasks and classes of tasks.

What is clear is that both IBM and Oracle have provided their customers with efficient tools for day-to-day management of their respective database systems. But when the cost of day-to-day operations is analyzed, Oracle Database 11g can provide organizations with 43 percent annual DBA-related cost savings.

About This Report

This report documents the results of a head-to-head product comparison examining the database administrative functions of Oracle Database 11g and IBM DB2 Enterprise 9.5. The study focuses on the use of human resources. Its objective is to reveal the comparative database administration costs of operating the two products.

Common database management tasks were performed in Oracle Database 11g and IBM DB2 Enterprise 9.5 and compared for their ease of use. For both products, their native management tools were used in the study as follows:

For Oracle Database 11g

- Oracle Enterprise Manager Database Control, supplemented with the Diagnostic and Tuning Management Packs

For IBM DB2 Enterprise 9.5

- Control Center
- Command Center
- Health Monitor / Health Center, supplemented with the DB2 Optimization and Recovery Expert Features

The study measures (in quantitative and qualitative terms), the relative manageability of Oracle Database 11g and IBM DB2 Enterprise 9.5, and projects the expected savings in management cost over the course of a year due to the administrative efficiency of one product over the other.

Who Should Read This Report

This report will be useful for corporate decision makers, technical end users (DBA's /System Administrators), and independent software vendors (ISVs). It will also be of particular interest to small and medium businesses with critical database requirements but with limited IT resources to manage them.

Methodology Overview

This Comparative Management Cost Study (CMCS), conducted by Edison Group, compares the ease of use or manageability of Oracle Database 11g with that of IBM DB2 Enterprise 9.5 and assesses their relative cost of management to a business. It represents

a product-specific application of a proprietary, general-purpose methodology developed by Edison Group for making product management cost comparisons. The result is a summary definition of the annual costs that will be incurred by any corporate IT department or ISV running either of these two products.

In the course of this study, Oracle Database 11g and IBM DB2 Enterprise 9.5 were compared against a set of methodology metrics in order to determine which of the two products is easier to operate for businesses with real-world database management requirements. The Test Administration Workload Task Areas that we used to perform this study fall into the following four categories:

- Database Setup and Configuration
- Day-to-Day Database Administration
- Backup and Recovery
- Performance Diagnostics and Tuning

Task categories were divided into individual tasks that logically map into their respective areas. To determine the overall manageability of a given task for a given product, each task was broken down into steps to assess the complexity and usability involved. Each task was then measured for time and number of steps required for successful completion.

Next, tasks were weighted against workload weighting constants. These weightings were used to determine the relative importance of a given task as measured against all of the tasks required to manage the entire product administration lifecycle. In other words, simple tasks that occur relatively infrequently were given a proportionately lower weighting than complex tasks that occur on a regular basis.

Finally, the results were tallied and the CMCS metrics for each product were substituted into manageability cost formulas to determine the projected human resources cost of operating both products, based on median DBA salary.

Contents of this Report

The following is a brief overview of the sections contained in this document, to provide for quick reference.

- **Rationale Behind this Report** – a discussion of the reasons Edison Group engaged in this research.
- **What's New?** - a discussion of the new features in Oracle Database 11g and IBM DB2 Enterprise 9.5 that are relevant to this white paper.

- **The Methodology Defined** – explains the criteria used in the study, including how we weighted and calculated the results and a description of the workloads evaluated.
- **Test Results** – presents the results of each set of tests, providing summary findings and a discussion of their relevance to business operations.
- **Conclusion** – summarizes our findings.
- **Appendices** – provide details on the test platform, clarification of architectural and terminology issues, the detailed test results, a discussion of our Complexity Calculation Formula, a detailed list of the actual steps performed and their individual timings, and a glossary of task areas.

Rationale Behind this Comparison

Over the last several years, Edison Group has developed an approach to Total Cost of Ownership analysis that recognizes the need for all companies to more easily manage the increasing rate at which change takes place in their data center environments. The approach focuses on day-to-day management, based on our belief that user-friendly management interfaces can lower administrative training costs and dramatically improve administrator efficiency. Edison's approach also takes into account the trend in which lower-skilled, lower cost IT generalists are used for day-to-day management tasks, thereby reserving more experienced, specialized, and highly paid staff for mission-critical tasks and more complex challenges.

The methodology utilized for the Comparative Management Costs Studies for Relational Database Management Systems is based upon reviews of published documents on database administrator tasks and interviews with administrators for the various products under study. For the series of studies, a simple laboratory environment was created, consisting of a single server onto which the respective database management systems were installed and configured. Once this was done, Edison Group's analysts performed and documented the evaluations contained in this paper. The conclusions in this report are our own, based upon the research we performed.

This series of research studies, begun in 2004, compares the manageability of RDBMS solutions from the three top vendors: IBM, Microsoft, and Oracle. The main thrust of this study is to independently test on a point-by-point basis Oracle's claims that its 11g database product is significantly easier and thus less costly to manage than IBM DB2 Enterprise 9.5. Edison's research has borne out those claims for the latest database releases of IBM and Oracle at the time of writing this report. The current paper demonstrates, once again, that Oracle offers its customer significant advantages in management costs over its competitors.

The Methodology Defined

For purposes of this study, the methodology is defined as a product manageability cost evaluation process, whereby the two products in question are compared against a set of task-oriented objective and subjective metrics in order to derive an accurate set of analytical results. The outcome of this study determines the Comparative Management Cost (CMC) incurred by managing and operating either of these products in a production environment. The methodology employed to conduct this comparison consists of the elements discussed below.

Workload Weighting

The workload weighting is a set of constants that define the relative importance of a single task area in the workload, based on frequency of execution and measured against the entire set of task areas that compose this study.

The Study

The study is the baseline checklist of standard database administration tasks routinely performed, which are quantitatively and qualitatively compared in order to objectively determine, on a task-by-task basis, which product is superior. This is measured primarily in terms of ease of administration and secondarily (for certain tasks only) in terms of system speed of execution — the wall clock time it takes for the system in question to complete a job once it has been submitted by a DBA. The function of this study is to apply a set of quantitative metrics, developed by Edison Group, to a list of tasks typically regarded as qualitative in nature, in order to derive a meaningful set of CMCS statistics that can reveal the real difference in management costs for the two products in question.

Tasks

A task is defined as a complete logical activity, composed of one or more steps, all of which effect a significant alteration on the state of the database that accomplishes a specific work goal. Each task is measured for time and complexity. Time and complexity, as measured in the study, are defined as follows:

Time

Defined as the amount of time it takes to perform a given task. For certain (asynchronous) tasks, when a job can be run in the background so that the DBA can use

the time for accomplishing other tasks, time is measured strictly in terms of the time it takes the DBA to perform the steps to configure, initiate, and submit a given task.

For other (synchronous) tasks in the study that demand the DBA's full attention and prevent the accomplishment of other tasks (as in performing a hot recovery operation on a live database), time is measured to include both the time it takes for a DBA to configure/execute the task in question as well as the time it takes the system to complete the task. All time metrics are measured in wall clock time and are rounded up to the nearest ten seconds.

Complexity

For the purposes of this study, complexity is measured using a proprietary metric devised by Edison Group. It is defined as the number of steps it takes to complete a given task, where a step is defined, within a wizard interface, as a major decision point, such as clicking the "Next" or "Finish" button.

Because not all steps have the same inherent complexity, each step is further broken down into increments to account for the difference. An increment is a decision point that the user must make to complete a step. Increments are technically defined as a part of a step that will have a measurable effect on the state or execution path of that step in the task process, but which in and of itself does not affect a change upon the underlying database state until the step being executed is complete. For example, selecting Basic vs. Advanced Install in the Oracle installation wizard screen is an increment and not a step, which effects an incremental change on the flow of the database installation process but does not change the state of the database.

Complexity is then measured in terms of number of steps, but taking into account the following factors:

- The number of increments it takes to complete each step.
- Whether or not instrumentation for a given step is GUI-based or requires the use of a command line/scripting interface.
- Whether or not the task requires a context switch between multiple interfaces in order to be completed. If a context switch exists, then additional steps will be added to the total step count for a given task.
- The above factors affect the complexity calculation as follows:
 - The primary measure is steps. If a step has many increments, it is considered several steps. The metric allows each step five increments, and thereafter we add steps for each additional five increments rounded up. So if a step has between 0–5 increments, it remains unchanged; if it has between 6–10 increments, it is

increased by one; between 11–15 increments, it is increased by two; and so on. We decided to do this because, while increments are secondary to steps in determining complexity, they do modify the relative complexity of a given step in the course of completing a task. In other words, steps with a low number of increments are simple, and steps with a high number of increments are complex.

- The other modifiers (instrumentation and context switching) occur very infrequently in the products under review, but they were significant enough a factor that we needed to account for them in some meaningful way in order to generate a measure of complexity that accurately reflects our experience of using the two products.
- Regarding instrumentation, if an operation could be executed entirely within a GUI interface, then the complexity/step value for that task would remain unmodified. If, on the other hand, a step required the use of a command line interface, this would increase the step count. For a simple single-line command operation, the step count was increased by one, whereas if the operation required the user to write a script, the step value was increased by two or more, depending on how much work was required to write the script in question.

Lastly, we come to the matter of context switching. If a context switch was encountered during the course of completing a given task, then two or more steps were added to the step count for that task. The possible addition of more than two steps was allowed for as a judgment call on the part of the analyst performing the task under consideration. The reason tasks containing context switches were penalized is that we regard as inherently more complex the need to understand the dependencies of relating and performing a single operation in two different environments in order to complete a single task, as opposed to performing a similarly complex task in a well-integrated environment where all the operations can be accomplished in one place.

The workload for this CMCS was reduced to the basic set of atomic maintenance operations that effectively fulfill all fundamental database administration procedures. The reasoning behind this approach is that enterprise-class database configuration and administration is a non-trivial matter; we therefore set out to develop a (relatively) simple yet comprehensive evaluation process, establishing a CMCS methodology benchmark that we feel is realistic in its technical assessment, yet accessible to the large audience of non-technical decision makers who will read this document.

The evaluation of each task in the study workload was executed by measuring the time that a typical DBA spends in a given workload task area using the methodology metrics. We counted compute time for utility tasks only in situations where the DBA is expected to “baby sit,” such as instance recovery.

This was the process we used in determining whether Oracle Database 11g is less expensive to operate than IBM DB2 Enterprise 9.5.

- **Licensing and Packaging:** License costs were not considered in this study. There are several reasons for this choice, the most important of which is that licensing costs are usually the least expensive factor in cost of ownership. For example, the management costs analyzed here are usually of greater significance.
- CMCS methodology has been derived from the following sources:
 - The initial baseline workload task list for this paper was based upon research performed by Oracle. This initial baseline was then modified by Edison Group and certain tasks were adjusted in order to compare both products on equal terms. The task lists have been further updated to reflect changes in the products since the first study was published in 2004.
 - The workload task weighting was based on a survey published by Database Trends.¹
 - The baseline workload task list was checked for process consistency by Edison Group analysts against the Oracle Database 11g Administrators Guide.
 - The baseline workload weighting and task list was further checked for consistency against the IBM DB2 Enterprise Administration Guide: Implementation, the IBM DB2 Enterprise Guide: Performance, and the IBM DB2 Enterprise Utilities Guide and Reference.
 - Independent professional Oracle and DB2 database administrators and engineers were consulted as anonymous third-party verifiers of the methodology and workload tasks employed in the course of conducting this CMCS.

Workload Weighting

To view these results in terms of management costs, we recognized that the tasks in the workload have different levels of importance and complexity, and are performed at differing levels of frequency. For example, tuning a database or creating a new table is performed more frequently than creating a new database. In order to accurately account for this, we have used a weighted average of the workload test areas to measure each set of tasks according to their typical degree of use.

¹<http://databasetrends.com/> - This study is no longer available for review. Edison has made enhancements to the original model based on interviews with experts and research performed by Edison staff.

Here are the weightings used:

Database Administration Workload Weighting	
Setup and Configuration	5%
Day-to-Day Administration	34%
Backup & Recovery	14%
Performance Tuning	26%
Other	21%
Total	100%

The Database Workload Weighting metrics in the table above were derived from an article published in Database Trends and Applications Online. The “Other” category represents tasks that were not included in the study, such as software license maintenance and database upgrades. We left this category in the weighting in order to acknowledge in the Comparative Management Costs (CMC) calculations that such activities are a necessary part of day-to-day business. Furthermore, in performing the CMC calculations for this analysis, it was presumed that both products required the same degree of “Other” tasks so as not to favor one product over the other. This approach renders the 21 percent of the workload weighting that falls into the “Other” category irrelevant to the outcome of this study.

Test Administration Workload

Installation and Simple “Out-of-Box” Setup Tasks

- Install db/software/out-of-box setup
- Create second database server/instance
- Setup proactive monitoring

Day-to-Day Database Administration Tasks

- Create user with roles, privileges
- Create tablespace/filegroup
- Add space to database
- Create table
- Create index
- Reclaim space due to fragmented data

- Load data from text file
- Set adaptive thresholds and compare workload performance to baselines

Backup & Recovery Tasks

- Configure and perform full backup
- Recover dropped table
- Recover data file
- Recover from erroneous transaction
- Recover from multiple failures

Performance Diagnostics & Tuning Tasks

- Diagnose performance problem
- Tune resource-intensive SQL
- Tune memory

Test Results

Comparative Management Cost Savings

The core premise of any Comparative Management Cost Study is that the true cost of owning and operating complex systems like Oracle 11g and IBM DB2 Enterprise 9.5 only start to accrue after the product has been purchased. In most real-world business environments, the management costs will far outweigh the licensing and support costs throughout the life of the product. With this in mind, we estimated the annual costs that businesses can expect to save due to the DBA-related time savings that result from one product being easier to administer and operate than the other.

In order to compute cost savings, we used DBA salary figures published by Enterprise Systems and Global Knowledge. Further information on these salary figures can be obtained from: <http://esj.com/News/article.aspx?EditorialsID=2769&pg=8> and <http://www.globalknowledge.com/training/generic.asp?pageid=1732&country=United+States> (The Global Knowledge link is for the 2007 report. Edison used the 2008 edition, of which we have an advance copy.)

If we insert the averaged DBA compensation found in the surveys² into the formula below, we arrive at the following quantitative management cost (MC) saving calculation.

$$\text{Median DBA Salary} * (\text{DBA time savings}) = \$81,757.25 * 43\% = \$35,155.62$$

This result can be interpolated to match to your company's DBA salary expenses by applying the above formula.

When multiplied across all of the DBA's in an organization, these management cost savings quickly grow into a figure that dwarfs the one-time licensing fee required to acquire a product of this nature.

² Edison averaged the salaries from the two reports as follows: Relevant job titles and salaries from the Global Knowledge report were: Database Manager, \$87,261; Database Administrator, \$78,468. Enterprise systems differentiated DB2 administrators at \$81,700 and Oracle administrators at \$79,600. The average salary used for this study, therefore, is: \$81,757.25.

Workload Category Results Synopsis

This section of the report delineates the Results Synopsis for each of the workload categories tested in this study. It will help the reader acquire a deeper understanding of how Oracle Database 11g and IBM DB2 Enterprise 9.5 really compare across the board. In this section, for brevity we refer to Oracle Database 11g as “Oracle” and IBM DB2 Enterprise 9.5 as “DB2”

Install DB/Software/Out-of-the-Box Setup (Results Synopsis)

This workload category addresses tasks relating to software installation and default out-of-the-box setup. Three tasks were performed in this area. The quantitative results synopsis for this area is documented in the table below.

Installation and Simple “Out-of-Box” Setup	Time (min)		Complexity (steps)	
	DB2	Oracle	DB2	Oracle
Category Sub-Total	12.67	22.92	54	37
Performance % difference — (DB2 - Oracle)/DB2	-81%		31%	
DBA Workday Savings	-4%		2%	

The results for this task were mixed for both, DB2 and Oracle. While IBM DB2 Enterprise 9.5 took much less time to install, the installation process was more complex than for Oracle Database 11g. When looked at from the perspective of weighted DBA Workday savings, the differences — though real — may not be significant for most organizations.

Qualitative Analysis

IBM DB2 Enterprise 9.5 took about 80 percent less amount of time to install than Oracle Database 11g. On the other hand, Oracle Database 11g required over 32 percent fewer steps than IBM DB2 Enterprise 9.5 to accomplish the same set of tasks. Reasons for these odd differences include the different manner in which the two products are sold or packaged to customers. Oracle Database 11g code comes on a single DVD, while IBM DB2 Enterprise 9.5 requires separate installation tasks to deliver the same database and manageability utilities. These separate installation tasks are the main contributor to DB2’s added complexity. At the same time, the additional software for DB2 was installed from files on the local hard drive, which is a generally accepted practice and contributed significantly to the lower installation time.

Because installation is not a frequently performed DBA activity, the major concern is not how long a product takes to install, but rather the complexity of choices involved. A first install often raises concerns about choices of such issues as block sizes or storage models, while subsequent installs are more predetermined. Both IBM DB2 Enterprise 9.5 and Oracle Database 11g have reduced the number of required installation parameters, and this category dramatizes the success of these reduced administration efforts. In both cases, the default installation parameters generate a reasonably sound database instance that can be used for common day-to-day administrative tasks.

It is important to note that the Oracle Database 11g installation includes the installation of Oracle Management Packs from the same installation media – in this case a DVD. Installation of the Performance Optimization Feature for IBM DB2 Enterprise 9.5 required additional installation and configuration steps; the software was installed from a downloaded file on the local disk, which is much quicker than installing from DVD media.

In large-scale implementations, installation of the DB2 Performance Optimization Feature is usually performed on only one server, and the DBA client software is often installed on DBA workstations rather than on the server. Edison chose this installation process and accounting to demonstrate the extra tasks involved for setting up the server and administrator tools for IBM DB2 Enterprise 9.5. In contrast, the administrator tools for Oracle Database 11g are installed and integrated into the database server at system installation; access to these tools takes place via a web browser, so no additional installations of client software are required.

When installing Oracle Database 11g, configuration of basic system maintenance tasks (backup, memory tuning, etc.) are accounted as increments of the appropriate install steps. Additionally, after installing an Oracle 11g database, automatic SQL Tuning and defragmentation (Segment Advisor) tasks can also be configured to run in maintenance windows using the Server ->Automated Maintenance Tasks link in Enterprise Manager. By contrast, IBM DB2 Enterprise 9.5 only requires a separate post-installation task – Configure Maintenance, for configuring optimizer statistics, automatic backup and data defragmentation tasks. If we set aside the one-time mechanism of how tasks are configured, the main difference between the two products is that Oracle Database 11g supports automated SQL tuning task, while IBM DB2 Enterprise 9.5 does not. The net result is that these tasks, which are evaluated under the Day-to-Day or Backup and Recovery or Performance Tuning task categories in previous studies, are now accounted for as part of Installation and Setup.

In Edison's opinion, the time spent waiting for an Oracle Database 11g instance to install is less critical in the bigger picture when compared to the correct configuration of the

database server. Post installation, the benefits of such packaged out-of-the-box configuration with Oracle result in significantly less complexity for routine database administrative tasks than required for IBM DB2 Enterprise 9.5.

Day-to-Day Database Administration (Results Synopsis)

The Day-to-Day Database Administration Workload task category relates to routine database object maintenance operations such as creating users, tables, and indexes, as well as reorganizing data and loading information into the database from external sources. Eight tasks were performed in this category. The quantitative results synopsis for this category is documented in the table below.

Day-to-Day Database Administration	Time (min)		Complexity (steps)	
	DB2	Oracle	DB2	Oracle
Category Sub-Total	21.50	11.33	106	67
% difference — (DB2 - Oracle)/DB2	47%		37%	
DBA Workday Savings	14%		13%	

Qualitative Analysis

The tasks under this category were performed in 47 percent less time with Oracle Database 11g than with IBM DB2 Enterprise 9.5. Recovering space due to fragmented data in IBM DB2 Enterprise 9.5 is now one of the several automatic maintenance tasks accounted for, in this study, during initial installation and setup of the system. As of this writing, configuring defragmentation in Oracle Database 11g is still a task requiring separate intervention by a DBA. The greatest difference for this category is the time required and the complexity of performing a task new for this version of this study: Set Adaptive Thresholds and Compare Workload Performance to Baselines.

New Task for This Study

One of a DBA's most important daily tasks is to monitor performance and identify anomalous events proactively. Doing so requires capturing of performance baselines for comparison with current activity. When a performance problem occurs, a DBA can use these baselines to compare the current status with past performance in order to determine the severity of the problem and the frequency with which it occurs, and to perform other analyses that can be helpful in diagnosis and prioritization of tasks for resolving the problem.

In order to proactively monitor anomalous events or performance, DBAs (across the database vendor spectrum) traditionally set static thresholds based on appropriate metrics to be alerted. This is so that they can react to the situation before it snowballs into a catastrophic event. However, this technique is highly ineffective because DBAs end up with the situation of managing too many false alerts, too many alerts, or worst yet, not being alerted when it is most required. There are many reasons for the poor success rate of static threshold-based alerting:

- Alert thresholds do not adjust to the time of the day
- Alert thresholds do not adjust to the type of workload, some metrics may not be relevant to a certain workload (e.g., TXN/sec for DSS workload)
- Too many metrics to choose from, typically a few hundred or more
- No easy mechanism to help DBAs choose the correct metric based on the time and type of the workload and the actual value to set it to based on a known normal operating reference period.

Oracle approaches this task through the use of Automatic Workload Repository (AWR), a key feature of the Oracle Database Diagnostic Management Pack. AWR is a warehouse of key workload performance data stored within the database that is used self-tuning and self-healing of the database. AWR tracks the type and time of the workload running on the system and provides ability to perform sophisticated statistical analysis required for accurate alerting. Oracle's Adaptive Thresholds feature overcomes the deficiencies of static threshold-based alerts. Oracle also provides DBAs with EM support for configuring Adaptive Thresholds with a few mouse clicks and interpreting data associated with the anomalous events. The number of metrics a DBA has to manage with Oracle Database 11g has been reduced to about a dozen from a few hundred. These metrics are sufficient to define the shape and volume of the workload running on the system.

Additionally, DBAs can create a baseline for a known normal time period, and can compare performance data for a problem that is currently happening to the baseline. Thus, DBAs can easily identify the cause of the problems looking at performance data from anomalous events. DBAs can also use an existing baseline for setting the adaptive thresholds on the system. A seven-day system moving window baseline is available to DBAs out-of-the-box and additional user specified baselines could be created as necessary.

IBM DB2 Enterprise 9.5 does not have easy to use Adaptive Thresholds and Workload comparison features. DB2 DBAs rely on the functionality of Performance Expert/Performance Optimization Feature. Performance Expert is a tool for enhanced

activity and health monitoring that stores workload information in Performance Warehouses (PWH). There are 25 predefined queries that may be run against the information stored in the Performance Warehouse. DBAs can use these queries to manually analyze workload type and patterns and write scripts to adjust the thresholds to reflect the observed workload.

Edison's evaluation of these sets of tools is focused solely on configuration and use and is not intended to be an evaluation of the design philosophies behind their creation and use. Within this task there are three sub-tasks:

- Set adaptive thresholds
- Configure seven-day moving window and a fixed period single day baseline
- Compare workload performance to moving window and fixed baselines

A real-world example for the above tasks would be a workload that encompasses daytime Online Transaction Processing (OLTP) and nighttime batch processing. In this scenario, you would want to monitor the regular peak utilization period and be alerted if an anomalous event, such as average transaction response, increased significantly. You would also like to compare the current average transaction response time with the seven day moving average and also to a fixed time period (say Monday 8AM-12PM peak or month-end close activity) on a regular basis.

Setting and scheduling these baselines was dissimilar for the platforms. Setting the equivalent to a moving window in IBM DB2 Enterprise 9.5, required repeating the weekly report configuration for each day a report was to be run. A DBA will also need to go through the configuration wizard each week, as the weekly setting is by calendar day rather than a default time period. Once the data is gathered for the week, the DBA would then have to manually analyze all the data and come up with the appropriate metric and its values based on workload type and pattern. The DBA then needs to write a script that adjusts the metrics and their values appropriately to reflect the data analyzed. This whole process of accomplishing an adaptive threshold alert equivalent to Oracle in IBM DB2 Enterprise 9.5 is very complex in nature and generally involves many iterations.

Another important difference arises when the DBA needs to compare baseline data with current activity. With Oracle Database 11g, baseline and current activity can be accessed with a few mouse clicks and viewed in line charts for immediate review and analysis. With IBM DB2 Enterprise 9.5, the data in a PWH can be analyzed through the use of any of several default queries, or customized queries can be run. The data from these queries can then be viewed as text or exported for analysis and charting in a spreadsheet program such as Microsoft Excel.

Edison performed several evaluations of workload comparisons in IBM DB2 Enterprise 9.5, including exporting of query results for analysis in Excel, printing out of results for paper-based review, and onscreen comparisons between multiple windows. The time required to perform these tasks was considerable, ranging from a two or three minutes to more than five minutes. Extracting near-active data and comparing it to snapshot data in Excel was even more time-consuming.

Actually analyzing problems in this manner was outside the scope of this study, but the reader can readily perceive the difference — manually combining, formatting, and analyzing data from several sources, as opposed to viewing a graphic representation comparing past and current activity as the result of a single procedure. Edison added five minutes to this task as a reasonable average time for generating reports and importing the data into an analysis tool. This time was added to the comparison step.

Performance workload comparison to moving or fixed window baselines using Oracle Database 11g only required a few clicks and is relatively simple. In contrast, IBM DB2 Enterprise 9.5 allows queries to be run over the entire warehouse of data, but the results are not directly comparable with live data, nor are there moving windows for performance monitoring available in Performance Expert.³

Edison has found that Oracle Database 11g provides DBA's with a powerful and immediately accessible interface for monitoring current activity and comparing it with baselines. We believe that, within the context of immediacy and ease of use, these tools are far superior to those offered by IBM DB2 Enterprise 9.5.

Backup and Recovery (Results Synopsis)

This task category addresses tasks relating to database backup and recovery management. Four tasks were performed in this category. The quantitative results synopsis for this category is documented in the table below.

Backup and Recovery	Time (min)		Complexity (steps)	
	DB2	Oracle	DB2	Oracle
Category Sub-Total	35.50	8.83	67	30
% difference — (DB2 - Oracle)/DB2	75%		55%	
DBA Workday Savings	11%		8%	

³ IBM DB2 9.5 Activity Monitor provides limited movable window analysis of current activity, but does not provide comparisons to baselines captured in PWH.

Qualitative Analysis

For this category, Oracle Database 11g performed significantly better than IBM DB2 Enterprise 9.5, requiring only 25 percent of the time and about half the number of steps required for IBM DB2.

Both Oracle Database 11g and IBM DB2 9.5 offer excellent tools for automating or simplifying backup and recovery tasks for their respective databases. Configuring automatic backups can be performed either during installation itself – with Oracle Database 11g – or when configuring the maintenance window operations – with IBM DB2 Enterprise 9.5. Once configured, backups run automatically. Our appendices have a place holder for the backup configuration task for consistency with earlier versions of this report, but we did not recount the steps required for configuration. The steps can be considered as part of Installation and one-time setup of the database.

Of course, backups can also be run manually when the need arises. Though Edison does not compare the steps involved in running a manual backup, for both products the task is wizard-driven and requires very few steps. Backup time is, of course, dependent upon the amount of data being backed up and the performance of the underlying hardware.

Performing a dropped table recovery with Oracle Database 11g takes only seven steps and one minute and ten seconds. According to the documentation, recovering a dropped table with IBM DB2 Enterprise 9.5 and DB2 Recovery Expert (RE) only requires nine steps.

Further, dropped table recovery is only possible if DB2 Recovery Expert versioning repository is up-to-date and a full offline back up is available before the table is dropped. This puts too much onus on DBA doing the tasks correctly and hence error-prone. Unfortunately, DB2 Recovery Expert was not yet compatible with the 64-Bit Windows version of DB2 Enterprise Edition we were running for the study. We therefore could not run Recovery Expert for recovering a table.

To show the steps involved when using DB2 Recovery Expert, we have provided the step descriptions from IBM's excellent Recovery Expert documentation. The following description is very technical in nature. It is included to illustrate the complexity of performing the recovery task manually.

The process to perform dropped table recovery is command line driven, and is superficially simple. It requires the following command to be run:

DB2 alter tablespace <table space> ⁴ dropped table recovery on

⁴IBM uses the term “tablespace” for the command and “table space” for the object when describing the storage location where the actual data underlying database objects can be

For this command to run, though, the DBA must enable:

DROPPED TABLE RECOVERY

This should be set at system setup.

While entering this simple command does not take long, it takes some time to perform the necessary steps, which include rolling forward, identifying and reapplying transactions just prior to the point-in-time when the table was dropped, and verifying the number of transactions. A safe time estimate for the small databases in our test bed would be a minimum of 10 minutes.

Even with DROPPED TABLE RECOVERY enabled, the DBA must restore the entire table space, not a single table. This requirement is fraught with danger — according to IBM's documentation:

“...recovering a dropped table space using native DB2 methods can be very disruptive and costly.

The documentation continues:

DB2 ... does not allow restore of a table space that has already been dropped. The table space has to exist in the database before we can do a table space level restore into it. Therefore, a table space restore from a table space backup is not possible even if you have table space data in a backup image. This means that the only option you have to recover table space is to restore the entire database and roll forward to a point-in-time before the table space was dropped. This can be very disruptive, as the entire database would become unavailable during the duration of recovery. Restoring a very large database can take several hours. In DPF environments you would have to restore the database on all nodes

Moreover, with database restore and rollforward, you must stop rollforward before the table space was dropped. There could have been many transactions done against other tables and table spaces since the time of drop table space. Once you stop at this point in time, you lose all the transactions that were executed against the database after this point-in-time.

This assertion is disturbing, to say the least. For our scenario of a single dropped table, the concern of there being no existing tablespace is not completely relevant, so that issue may be ignored. But the potential need to restore to a point in time, rollforward, and perhaps re-perform transactions still exists.

kept. Oracle uses the term “tablespace” for both usages. This paper uses “tablespace” except when quoting IBM documentation.

In the backup/recovery scenarios, it is significant to point out the Oracle flashback features are an enormous benefit in specific scenarios. Oracle provides a recovery window in flashback, allowing logical recovery — a conceptual leap over the traditional physical recovery — that can reduce problem resolution efforts by orders of magnitude. The Oracle recovery through flashback also seems qualitatively safer than the DB2 options of restoring a database and merging data or losing data through point-in-time recovery. This flashback technology was the main reason behind Oracle’s substantial advantage over DB2.⁵

Since the tasks of recovery of dropped table and recovery of deleted rows required extensive manual intervention for DB2, not usually not required for Oracle Database 11g, Edison assigned each of these two tasks, in DB2, a ten-minute time penalty to make the two efforts more comparable from a workload perspective. Without this penalty, an IBM DB2 Enterprise 9.5 DBA would need less time to perform the tasks. In reality, the need for manual intervention in a real dropped table or rows scenario would probably require far more than the ten-minute penalty.

For performing various recovery scenarios with IBM DB2 Enterprise 9.5 (even with Recovery Expert operational), the onus is on the DBA to have a full offline backup, archive log meta-data, and in-sync versioning repository available. This is in contrast to Oracle Database 11g that places no such requirements and onus on the DBA. Oracle’s flashback database technology makes it the only vendor that provides database built-in Continuous Data Protection (CDP) capability. Additionally with Oracle Database 11g’s Data Recover Advisor (DRA) capabilities, DBAs can auto-diagnose multiple failures of physical and logical nature, understand user impact of such failures, and implement suggested recovery actions. IBM DB2 Enterprise 9.5’s RE does not provide auto-diagnosis of failure — one first has to know which recovery action to take to begin with. In summary, Edison believes that while both Oracle Database 11g and IBM DB2 Enterprise 9.5 have made major strides in simplifying data restoration of their products, Oracle Database 11g is clearly ahead in this category providing easy-to-use and compelling solutions to eliminate human induced errors.

⁵ As mentioned previously, IBM does offer an add-on product, IBM DB2 Recovery Expert, which might have affected the results of this study in the area of database and table recovery. Recovery Expert was excluded from consideration for one very interesting reason: the version of IBM DB2 Recovery Expert available at the time of testing was not compatible with the Windows 64-Bit version of DB2 9.5 we tested. We did review the documentation for DB2 Recovery Expert and determined that, had it worked, it would have reduced the complexity and — most significantly — the time required for recovery, but with the caveats mentioned in this section

Performance Tuning (Results Synopsis)

This task category addresses tasks relating to manual and automated systems performance diagnostics and tuning. Two tasks were performed in this area. The quantitative results synopsis for this category is documented in the table below.

Performance Diagnostics and Tuning	Time (min)		Complexity (steps)	
	DB2	Oracle	DB2	Oracle
Category Sub-Total	2.75	.33	44	3
% difference — (DB2 - Oracle)/DB2	88%		93%	
DBA Workday Savings	23%		24%	

Qualitative Analysis

Oracle Database 11g is 88 percent more efficient in time and 93 percent less complex than IBM DB2 Enterprise 9.5 for diagnosing and tuning the database. This significant difference is mostly due to Oracle’s automated nature of performance diagnosis and SQL tuning. Additionally, Oracle’s interface for performance diagnosis and for SQL tuning are so well designed that this takes very little time from an end-user perspective, while the IBM DB2 Enterprise 9.5 interface required more time to be spent reviewing and making selections from the choices available.

This category is the one presenting the greatest leeway in testing. There are a bewildering number of ways in a modern RDBMS to inadvertently create performance issues, as well as to identify and diagnose them. Outside of design, the most rarified stratum of skills within the DBA world that can most affect utility is the tuning of applications and configurations. The complexity of these tasks has driven vendors towards increased automation of performance tuning.

Due to the multiplicity of means for performance tuning, we believe that the best way to examine the results for this category is task-by-task. This study was originally designed with three tasks: diagnosis, SQL tuning, and memory tuning. As we discuss below, this list has been reduced to only two tasks; memory tuning is now automated for both platforms.

Diagnose Performance Problems

For the task of diagnosing performance problems, the Oracle Database 11g product installation process — combined with the selection of appropriate defaults when creating tables/indexes, gathering workload performance data, and other tasks —

enables automatic diagnostics to function with no further interaction on the part of the DBA. IBM DB2 Enterprise 9.5 Health Monitor can identify problems when certain thresholds are exceeded and bring them to the DBA's attention. Diagnosis of a problem can now be performed using DB2 Activity Monitor, from which previously identified or other problem SQL can be selected and diagnosed. An explanation of the problem is provided (in terms of the observed symptoms), including presentation of a resolution plan.

As for Oracle Database 11g, unique features such as AWR (Automatic Workload Repository), ASH (Active Session History), and ADDM (Automatic Database Diagnostic Monitor) completely automate the process of performance diagnosis, alleviating the need to replay workload or to enable monitor events/tracing.⁶

Tuning Resource-Intensive SQL

Because the two products under consideration — IBM DB2 Enterprise 9.5 and Oracle Database 11g — are not likely to be found in the single-server simple database scenario that we used in this study, it sometimes becomes necessary to recognize and consider the real-world manageability challenges frequently faced in the complex enterprise-class data centers more likely to deploy them. Tuning resource-intensive SQL is one such challenge.

Oracle Database 11g provides a SQL Tuning Advisor which can be run against Oracle's diagnosis report. The Advisor returns a tuning recommendation including suggestions for new or revised indexes, auxiliary information about the query (called SQL Profile) that transparently improves performance by tuning the SQL execution plan, and modifications to SQL code. For this study the recommendation of SQL Profile was accepted and the tuning went forward automatically. This choice demonstrated both an automatic and interactive approach to tuning. For this study, we used the automatic mode of SQL Tuning and auto implementation of SQL Profiles by configuring it to run in system default maintenance windows. Tuning SQL statements through SQL Profiles recommended by Automatic SQL Tuning feature can also be performed without application code being changed, helping improve performance on packaged applications where a vendor upgrade may be unlikely, untimely, or not relevant to a particular customer's needs. This feature also applies to internal development lifecycles

⁶ Oracle ADDM reports many kinds of problems (Resource Intensive SQL, literal usage, log sizing, locking contention, I/O Issues, Striping issues, etc.), prioritized in descending order of impact and quantified by DB time consumed. This information helps a DBA determine the benefits of making specific tuning decisions.

Obtaining similar information from IBM DB2 requires performing the diagnosis process several times — once for each datum.

— code revisions can be provided to developers for future incorporation without undo disruption to projects in progress.

Oracle Database 11g Diagnostic Pack and the Tuning Pack can be utilized in conjunction with each other to directly identify problems and compare them to baseline performance, greatly simplifying root-cause analysis.

With IBM DB2 Enterprise 9.5, deeper digging for root-cause analysis may require the initiation of appropriate monitoring events, gathering snapshot data, and (at times) enabling CLI tracing mechanisms. Based on the problem at hand, users have to decide what appropriate events or level of tracing to enable, both of which incur performance overhead. Additionally, most often these problems are not known beforehand and often might require replaying the problem workload. These extra steps add to the time and complexity of the diagnostic process in IBM DB2 Enterprise 9.5.

Tuning with IBM DB2 Enterprise 9.5 has been improved over past versions Edison has evaluated. As mention above, problems can be identified and diagnosed through the DB2 Activity Monitor. From the monitor, DB2 Design Advisor — a tool for tuning SQL statements — is called, and the problem is further analyzed. In a manner similar to that provided by Oracle Database 11g, DB2 Design Advisor provides the DBA with a resolution plan for manual or automatic resolution. It goes still further with automated repair, providing the DBA with the choice of options for automatically tuning the SQL code plus automatically adding or restructuring indexes. The DBA can also choose from table partitioning options, when applicable.

There are two types of steps measured in the performance of this task: setting up the monitoring or data capture for the problem SQL and performing the tuning itself. With Oracle Database 11g, the database automatically identifies the high load SQL statements and tunes them in the maintenance window, and implements SQL Profiles where possible to transparently improve application performance. A DBA needs only to look at the system generated tuning report for information. With IBM DB2 Enterprise 9.5, both setup and tuning needs to be run each time its required. To illustrate this, our metric assumes that tuning is performed twice. For Oracle Database 11g, this means that a DBA needs to only look at the system-generated tuning report twice. For IBM DB2 Enterprise 9.5 all of the steps are counted twice.

Edison Group therefore feels that when SQL tuning is needed for problematic queries Oracle has an advantage over DB2. In the real-world data center, the need for the tuning of complex queries is a common and time-consuming DBA task: this is why it carries the strongest weighting in our workload metric.

Memory Tuning

Memory tuning is another important aspect of performance tuning. Both IBM DB2 Enterprise 9.5 and Oracle Database 11g now support automatic memory tuning. All that is required is enabling automatic memory tuning during setup. While there is an option to enable it later, within the context of this study there are neither steps nor time required to tune memory.

Conclusion

This CMCS study represents our detailed analysis of Oracle Database 11g and IBM DB2 Enterprise 9.5 in the area of manageability. The study demonstrates not only that DBAs can perform routine administrative functions more quickly with Oracle, but that these functions are quantifiably easier to perform.

The differences between the two platforms have slightly increased over the years that Edison has been performing these studies. While IBM has done a good job in addressing many of the criticisms Edison has leveled in the past, most significantly in the areas of general system maintenance, Oracle has also significantly improved its offerings in these areas.

The key factors for Oracle's superior manageability are self-monitoring and diagnostic capabilities, and the increased automation of many otherwise manual but vital DBA tasks such as recovering from human errors, SQL tuning, and performance diagnostics. These manageability advantages translate into substantial management cost savings for businesses along with higher reliability and availability of their systems.

Oracle Database 11g demonstrated significant strengths in the more heavily weighted areas of Performance Diagnostics and Tuning and Backup and Recovery. These strengths directly translated into the 43 percent cost savings calculated herein.

A final point, perhaps not immediately apparent from our report: Oracle Database 11g's management console — Oracle Enterprise Manager — is the sole interface required for access to all the management features of the system. This console can be accessed via a web browser from any authorized location.

IBM DB2 Enterprise 9.5, on the other hand, presents a single central management console, Command Center, through which the separate management consoles for each discreet tool or management application can be accessed. In addition, the interfaces of these several applications and consoles are not completely consistent. We have taken account of these design differences in assessing complexity, but it is important to recognize the impact that switching interfaces and multiple consoles can have on productivity. A deeper discussion of these factors lies outside the scope of this document.

Appendix I - Test Platform Details

Below are the technical specifications of the lab computers used to conduct this CMCS.

Hardware Platform: OS/Hardware Platform Data Points

- **OS:** Microsoft Windows Server 2003 R2 X64
- **Processors:** Four Intel Xeon MP Dual Core 3.4 GHz CPUs
- **Memory:** 4 GB DDR, 4x1 GB
- **Storage:** 73.4 GB 10 K drive

Software Platforms

Relational database management software platform installed for this study.

Oracle

- Oracle Database 11g Enterprise Database (64 Bit)
- Oracle Database Diagnostics and Tuning Packs
- Oracle Enterprise Manager - Database Control
- Database Configuration Assistant
- SQL*Plus

IBM

- DB2 — DB2 Enterprise Server Edition for Windows (64-bit), v9.5
- DB2 Control Center
- DB2 Command Center
- DB2 Health Monitor / Health Center
- DB2 Command Line Processor
- DB2 Optimization Feature – Performance Expert

Appendix II - What's New?

IBM DB2 and Oracle Database systems support the needs of a wide variety of customer applications, ranging from those of small through medium to large businesses, which need the capabilities of a state-of-the-art database management system. This past decade has seen the growth of these database systems in capability, functionality, and scalability.

Analysts see keeping up with this growth without increasing manageability costs as a critical factor in the continued success of these systems.

IBM DB2 Enterprise 9.5

IBM has addressed the management ease of use challenge by focusing on autonomic enhancements. In the most recent edition, version 9.5, the enhancements include several changes to product architecture and features that have some degree of relevance to the Edison CMCS process. Among these changes are (content excerpted from IBM literature⁷):

Multithreaded Architecture

DB2 data servers now have a multithreaded architecture on all platforms. Prior to Version 9.5, DB2 data servers running on UNIX® and Linux® operating systems used the process-based model, in which each agent ran in its own process.

The move to the multithreaded architecture provides the following benefits:

- Dynamic self-tuning memory is now fully enabled on all platforms.
- Memory configuration is simplified through more automatic and dynamic memory configuration parameters.
- Agent and process model configuration is simplified. These enhancements eliminate the need for regular DBA interventions to adjust the process model-related parameters and reduce the time and effort required to configure them. You do not need to shut down and restart DB2 instances to have the new values take effect.
- New dynamic configuration parameters enable the database manager to automatically tune your system without manual intervention.

⁷ IBM Information Center at: <https://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.db2.luw.doc/welcome.html>

- Performance is improved because context switching between threads is usually faster than between processes (depending on the application and platform).
- Sharing of resources, such as file handles, is more efficient and provides system resources savings because all agents working on the same database file share the same file handle.
- Memory footprint is reduced.
- A consistent threading architecture provided across all operating systems reduces the complexity and maintenance of your data servers.

Automatic Storage Management Enhancements

Automatic storage automatically expands the size of your database across disk and file systems. It eliminates the need to manage storage containers while taking advantage of the performance and flexibility of database-managed space. Version 9.5 introduces tablespace handling enhancements that ensure that you are able to reduce the size of a tablespace to an accurate High Water Mark (HWM). This allows you to reclaim unused space automatically.

Additional Automatic Configuration Parameters

Version 9.5 includes more tuning parameters that the data server automatically handles, without requiring you to stop and restart your instance or database.

Automated Maintenance Improvements

If you want to take advantage of automatic maintenance but need to maintain tight control over the process and policies, you can use the new system stored procedures for collecting automated maintenance configuration information for the following areas:

- Maintenance windows
- Automatic backups
- Automatic table and index reorganizations
- Automatic table RUNSTATS operations

Since many of the above changes are at the command line level, their effect may not be directly evident with Edison's test methodology. Other items, shown as new or revised commands, will have an effect because they reflect default settings accessible from DB2's GUI. There are many other changes to the product that are listed on IBM's web site. Refer to that site if you wish to learn more.

In addition to the above built-in functionality, IBM DB2 Enterprise 9.5 also takes advantage of add-in features relevant to this study. Key among these is the DB2 Performance Optimization Feature. Requiring a separate installation task as compared to Oracle's inclusion of its similar management packs at initial installation, this feature provides advanced capabilities for managing workloads and optimizing performance across a database environment. It includes DB2 Workload Management, DB2 Performance Expert, and DB2 Query Patroller.

- DB2 Workload Management customizes your data server to support a spectrum of users and workloads on the same system, using either a predictive or reactive model.
- DB2 Performance Expert is a performance analysis and tuning tool for managing a heterogeneous mix of DB2 systems with a single end-user interface.
- DB2 Query Patroller is a powerful query workload management tool that controls submission and execution of queries to better manage DB2 data server workloads.

Oracle Database 11g

Oracle has made significant enhancements in the areas of manageability, availability, and performance. It has introduced innovative capabilities for change assurance, fault diagnostics, and performance tuning. The new capabilities introduced in Oracle Database 11g are packaged to customers in different ways: as part of the base product, as new options or as Database Management Packs. Further discussion on product packaging in Oracle Database 11g is beyond the scope of this document, but from a manageability perspective, it is important to note that no additional effort is needed by a DBA. A comprehensive list of new capabilities is available in the Oracle Database 11g documentation⁸. Detailed coverage of these features is beyond the scope of this paper; therefore we highlight the features we considered interesting or those belonging to Database Diagnostic and Tuning Packs that were relevant to this study.

(The descriptions below are extracted from Oracle product literature.⁹)

Diagnostic Pack

The key features of the Diagnostic Pack utilized in this study include:

- **Automatic Database Diagnostic Monitor (ADDM)** - this self-diagnostic solution enables the Oracle Database 11g to automatically diagnose its performance problems, thereby completely liberating administrators from this arduous and complex task. ADDM also diagnoses RAC specific performance problems and provide recommendations on how best to resolve them.

⁸ Oracle® Database New Features Guide 11g Release 1 (11.1) B28279-03

⁹ Available at http://www.oracle.com/database/db_manageability.html

- **Automatic Workload Repository (AWR)** - in order for ADDM to accurately diagnose performance problems, it is important that it have detailed knowledge of database activities and the workload the database is running. AWR is the built-in repository that contains the operational statistics on the database it is running as well as other relevant information. At regular intervals (hourly by default) the database takes a snapshot of all its relevant vital statistics and workload information and stores them in the AWR. The data is stored in AWR for a given period of time (eight days is the default) before being purged. AWR is designed to be lightweight and to automatically manage its use of storage space, ensuring that it does not put an additional burden on administrators.
- **Performance Monitoring** - AWR also supports the creation of performance baselines for normal operating time periods (such as peak or month-end workloads) and configuring adaptive thresholds. The performance baselines can be compared to a problematic time period to help quickly determine the cause of performance regressions. By default, a seven-day moving window baseline is available to the users for performance comparison. EM interface is easy to use and intuitive and supports comprehensive (hosts, database, instances, listeners) performance monitoring.

Tuning Pack

- **SQL Tuning Advisor** – automates the SQL tuning process by comprehensively exploring all the possible ways of tuning a SQL statement. The database engine’s significantly enhanced query optimizer performs the analysis and tuning. Four types of analysis are performed by SQL Tuning Advisor:
 - **Statistics Analysis** – in this analysis, objects with stale or missing statistics are identified and recommendations are made to remedy the problem.
 - **SQL Profiling** – eliminates the need for manual processes and tunes SQL statements without requiring any change to the application code. This ability to tune SQL without changing the application code also helps to solve the problem of tuning packaged applications. With SQL profiling the tuning process is automatic and immediate.
 - **Access Path Analysis** – new indexes that can significantly enhance query performance are identified and recommended.
 - **SQL Structure Analysis** – relevant suggestions are made to restructure selected SQL statements for improved performance.
- **Automatic SQL Tuning Advisor** – When SQL Tuning Advisor is run in automatic mode, high-load SQL queries are automatically selected and recommendations are generated on how to tune them. The Automatic SQL Tuning Advisor can also 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 performance improvement is significant. Other types of recommendations, such as to create new indexes, refresh optimizer statistics, or restructure SQL are also reported to DBAs.

- SQL Access Advisor — for a given workload provides comprehensive tuning recommendations, such as adding /dropping indexes, partitioning advise (type and columns), materialized view/logs to improve workload performance.
- Object Reorganization Wizard — Managing the space usage of your tablespaces efficiently by removing wasted space is not only a good space management practice, it also enhances performance by reducing unnecessary disk I/O. Reorganization is used for:
 - Recreating objects with optimal storage attributes
 - Rebuilding indexes and tables that are fragmented
 - Relocating objects to another tablespace

Automatic Diagnostic Repository and EM Support Workbench

Automatic Diagnostic Repository is new system managed repository for storing and organizing trace files and other diagnostic data. ADR provides comprehensive view of critical errors encountered by the database and maintains all relevant data needed for problem diagnostics and their eventual resolution. Automatic Diagnostic Repository (ADR) provides a uniform and consistent mechanism to store, format, and locate all database diagnostic information. As a result, customers can now correlate errors across various components such as Oracle RAC, Oracle Clusterware, OCI, Net, and processes. ADR also automatically generates incidents for serious errors and provides incident management functionality. EM Support Workbench provides an easy-to-use interface for packaging first-failure diagnostic information to support and thereby reducing problem resolution time for customers.

Real Application Testing

- Real Application Testing - is a new option in the database that provides businesses with extremely cost-effective and easy-to-use solution to fully assess outcome of a change in a test environment, and take corrective actions if necessary. The change can be then introduced safely to production systems thereby minimizing any undesirable impacts of the change. Real Application Testing consists of two features, Database Replay and SQL Performance Analyzer. Database Replay makes it possible to capture production workload including information about timing, synchronization and concurrency and replay the workload on a test system with production characteristics to assess the impact of change. SQL Performance Analyzer allows fine-grain impact analysis of database environment change on SQL execution

plans and performance. It integrates with other features such as SQL Plan Management and SQL Tuning Advisor to remediate regressed SQL statements. Together, SQL Performance Analyzer and Database Replay provide realistic workload testing for system changes. Real Application Testing is available for database releases 9.2.0.8 and above to help customers smoothly transition to higher releases.

Appendix III - Architectural and Terminology Discrepancies Requiring Clarification

Both Oracle Database 11g and IBM DB2 Enterprise 9.5 are mainstream relational database management systems. Yet, relative to how these products have been designed to execute their respective database management functions, a few key architectural and terminological discrepancies should be clarified so that both Oracle and DB2 users reading this report will have a clear understanding of the synonymous terms and tasks comprised within this comparison. (See Appendix VI for a Glossary of Terms.)

Oracle database users are either created only within the database or both in the database and in the operating system; DB2 users must have an Operating System ID (or corresponding IBM Tivoli identification).

Oracle schemas are user names, while in DB2 they are distinct entities. DB2 has coarse grain point-in-time recovery; recovering a single tablespace is the lowest logical recovery point in DB2, while in Oracle the flashback mechanism provides any detail of transactional history with a clear and powerful interface. DB2 SQL is generally precompiled into plans and packages of compiled access paths, while Oracle mostly uses dynamic SQL, relying on cache for performance.

Appendix IV - Detailed Test Results

Task	Description	Time (min)		Steps	
		DB2	Oracle	DB2	Oracle
Installation and Simple "Out-of-Box" Setup					
1	Install db/software/out-of-box setup - Incl. P.E. install and configure	10.33	16.08	35	12
2	Create 2nd database instance	1.08	6.42	9	22
3	Setup proactive monitoring	1.25	0.42	10	3
	Category Sub-Total	12.67	22.92	54	37
	% difference -- (DB2 - Oracle)/DB2	-81%		31%	
	DBA Workday Savings (Difference * Weighting Factor)	-4%		2%	
Day-to-Day Database Administration					
4	Create user with roles, privileges	0.83	1.25	6	7
5	Create tablespace	0.92	0.92	9	7
6	Add space to tablespace	0.67	0.92	5	6
7	Create table	2.00	1.75	16	5
8	Create index	1.17	0.83	7	5
9	Defragment tablespace	0.00	0.58	0	6
10	Load data from text file	2.42	3.58	5	16
11	Set adaptive thresholds and compare workload comparison to baselines ¹⁰	13.5	1.50	57	16
	Category Sub-Total	21.50	11.33	106	67
	% difference -- (DB2 - Oracle)/DB2	47%		37%	
	DBA Workday Savings (Difference * Weighting Factor)	14%		13%	

¹⁰ Includes setting up a recurring daily report and a weekly moving window (or equivalent) report in either platform. See body of white paper for details on this task.

Task	Description	Time (min)		Steps	
		DB2	Oracle	DB2	Oracle
Backup & Recovery Tasks					
12	Configure and perform full backup	0.00	0.00	0	0
13	Recover dropped table	11.33	1.17	15	7
14	Datafile recovery	0.75	1.92	12	7
15	Recover deleted rows: FB transaction	11.33	1.92	16	8
16	Recover from multiple failures	12.08	3.83	24	8
	Category Sub-Total	35.50	8.83	67	30
	% difference -- (DB2 - Oracle)/DB2	75%		55%	
	DBA Workday Savings (Difference * Weighting Factor)	11%		8%	
Performance Diagnostics & Tuning Tasks					
17	Diagnose performance problem	1.00	0.33	8	3
18	Fix performance problem: SQL Tune	1.75	0	36	0
19	Memory tuning (SGA/PGA)	0.00	0.00	0	0
	Category Sub-Total	2.75	.33	44	3
	% difference -- (DB2 - Oracle)/DB2	88%		93%	
	DBA Workday Savings (Difference * Weighting Factor)	23%		24%	
	Overall Total	74.42	43.42	271	137
	% difference -- (DB2 - Oracle)/DB2	40%		49%	
	Total DBA Workday Savings	43%		46%	

Appendix V – Complexity Calculation Formula

The following is the complexity formula utilized throughout these studies. If our testing identifies a need for modification of the approach, the section will be modified appropriately. Any modifications will be annotated for reader reference.

Complexity is defined as the number of computed steps it takes to complete a given task. The formula used to compute complexity for each of the tasks in this study is as follows:

- For every five increments contained in a step we increase the step value by one. For example:
 - If a step has 0–5 increments, step value = step + 0,
 - If a step has 6–10 increments, step value = step + 1
 - If a steps has 11–15 increments, step value = step + 2
 - ... and so on.
- The type of instrumentation offered to perform a given task modifies the task's complexity.
- If a task can be performed completely with a GUI, then step count = step count + 0.
- If a task requires the use of a command line interface, then step count is modified as follows:
 - If the command line operation consists of a single-line command, then step count = step count + 1.
 - If the command line operation requires writing a script, then step count = step count + 2 or more steps, depending on a subjective assessment of the complexity of the script.
- If a task requires a context switch between different environments, then step count = step count + 2 or more steps, depending on a subjective assessment of the complexity of the context switch.
- In the case of point-in-time recovery operations, we exercised the discretionary option of assessing time penalties that we consider exceptions to these rules, due to the open-ended subjective and analytical nature of the operations required to successfully complete those tasks.
- The tasks labeled Adaptive Thresholds and Workloads Comparisons compares features that are very dissimilar between the platforms. Oracle Database 11g has a specific set of tools that automate and streamline this very important functionality.

These tools are not available in IBM DB2 Enterprise 9.5. Therefore the Edison analysts identified the closest similar tools in IBM DB2 Enterprise 9.5 that provide equivalent functionality. Since DB2 lacks the same automation features, several iterations of the steps needed to be performed to deliver equivalent reports.

Appendix VI - Detailed Task Descriptions¹¹

Installation and Simple “Out-of-Box” Setup

Task 1: Install database and management software, and create starter database

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Install DB/ Software/Out-of-Box Experience			
Windows Explorer: Run setup.exe from DVD	16 Min 5 Sec	Windows Explorer Run setup.exe from DVD	10 Min 20 Sec
11g Installation: Select Product: Accept Basic Installation Defaults / Set Database Password / Confirm Password / <NEXT>		Welcome Screen / <Install A Product>	
Installer: Product-Specific Prerequisite Checks: Verify Successful Status <Next>		Install A Product / Enterprise Version / <Install New>	
Installer: Configuration Manager Registration – Accept Defaults <Next>		Continue <NEXT>	
Installer: Summary – Review Settings <Install>		Software License Agreement / Accept / <NEXT>	
Database Configuration Assistant: Backup Encryption Key and Password Management Notification <OK>	(1 Context Penalty)	Select Installation Type / Typical <Next>	

Task Description Table Key:

- Each row counts as a step (except where noted)
- “/” in the Step column represents an increment for the step.
- The button or link that ends a step is usually indicated by brackets “<>” surrounding the label.
- Tasks with increment, context or other penalties are indicated within parentheses in the adjacent time column.
- There are explanatory comments and footnotes as required.

For more information on the metrics refer to the section of this study entitled The Methodology Defined and Appendix V – Complexity Calculation Formula.

Oracle		IBM DB2	
Step	Time	Step	Time
Installer: End of Installation (Success) <Exit>		Confirm Install <YES>	
Exit: <Yes>		Installation / Response File Options / Configure /<Next>	
Inventory: Review Installed Products <Close>		Specify Installation Folder / (Accept Default) / <NEXT>	
Exit <Yes>		User Info For Admin Server / Specify / <NEXT>	
		Instance Creation / Accept Default / <NEXT>	
		Prepare Tools Catalog / [Select] / <NEXT>	
		Setup Notifications / Specify Server / Accept Defaults / <NEXT>	
		Specify Health Monitor Contact / <NEXT>	
		Enable OS Security on DB Objects / Accept Defaults / <NEXT>	
		Begin Install / Review Settings / <Finish>	
		Setup Complete / <Finish>	
		PES Install: Start <OK>	(1 Context Penalty)
		PES Install: Ready to install <Start>	
		PES Install: Confirm <Next>	
		PES Install: License <Next>	
		PES Install: Location - Accept Default or Browse to New Location<Next>	
		PES Install: Install <Next>	
		PES Install: Startup Preferences - Auto (Default), Manual <Next>	

Oracle		IBM DB2	
Step	Time	Step	Time
		PES Install: Configure Now y/n	
		PES Config: Start PE Server (From Start Menu) Open Start Tool	(1 Context Penalty)
		PES Config: Open Configure PE Server GUI	
		PES Config: Connect (Local Instance or Remote Instance) - Local Selected and Discovered Automatically	
		PES Config: Add Database - Right-Click Monitored Database - Select Add Database or Add Database from Catalog	
		PES Config: Enable Monitoring – Right-Click Instance, Select Enable Monitoring	
		PES Config: Close Configuration	

Metrics	Oracle	IBM DB2
Time	16 Min 5 Sec	10 Min 20 Sec
Steps	10	31
Increment Penalty		
Context Switch Penalty	2	4
Adjusted Steps (Complexity)	12	35

Task 2: Create additional database server/instance

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Create a new or second database instance			
Start Menu / Configuration and Migration Tools / Database Configuration	6 Min 25 Sec	Start Menu/DB2/First Steps	1 Min 5 Sec
DB Config Assistant Welcome Screen / Next		Database Creation/Create Your Own Database	
Create a Database / <Next>		Specify Name - Accept Defaults / <NEXT>	
Select Database Template : General Purpose / Transaction Processing / <NEXT>		Specify Storage - Accept Defaults / <NEXT>	
Specify Database Configuration Options: DB Naming/ DB Char Set/ DB Examples/ Keep Defaults.		Select Maintenance Strategy / Offline Maintenance Window / <NEXT>	
Database Identification: Enter Global DB Name / SID / <NEXT>		Offline Auto Maintenance / Schedule / <NEXT>	
Database Credentials: Use Same Admin Passwords for All Accounts / Enter PW / Confirm PW / <Next>		Notification Contacts / Specify / <NEXT>	
DB File Locations : Accept Defaults / <NEXT>		Review Actions / <Finish>	
Set Backup Job Credentials With the Username and Password of a User with at Least Backup Operator Privileges.		Monitor Progress	
Recovery Configuration: Accept Defaults <Next>			
Database Content – Accept Defaults (No Sample Schemas or Custom Scripts) <Next>			

Oracle		IBM DB2	
Step	Time	Step	Time
Initialization Parameters : Memory/ Sizing/ Character Sets/ Connection Mode/ Accept Defaults <Next>			
Security Settings: Accept Defaults – Keep enhanced 11g Defaults: <Next>			
Automatic Maintenance : Accept Defaults (Enable) <Next>			
Database Storage: Accept Defaults <Next>			
Creation Option: Accept Defaults (Create Database) <Finish>			
Review / Confirm Selections <OK>			
Perform <Create Database>			
Review Final Messages <EXIT>			
Specify Database Schema Passwords / Use Same Password for All Accounts / <Install>	(1 Context Penalty)		

Metrics	Oracle	IBM DB2
Time	6 Min 25 Sec	1 Min 5 Sec
Steps	20	9
Increment Penalty	0	
Context Switch	2	
Adjusted Steps (Complexity)	22	9

Task 3: Set up proactive monitoring for performance and space utilization

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Setup Proactive Monitoring			
Enterprise Manager / Database Home Page / Metric and Policy Settings	25 Sec	Health Center /Configure /Settings	1 Min 15 Sec
Metric and Policy Settings / All Metrics/ Select Metric to Adjust / Enter Value / <OK>		Indicator Configuration Launch Pad <Object Settings>	
		Select a Database Object <Selected Database Name><OK>	
		Object Health Indicator Configuration "Database Name" / <Database Heap Utilization>	
		Configure Health Indicator / Review Threshold Settings / <OK>	
		Settings Configuration Message: "Settings Adopted" <CLOSE>	
		Select Object Dropdown / Select <Tablespace Name> / <OK>	
		Object Health Indicator Configuration (For Tablespace) / <Tablespace Utilization>	
		Configure Health Indicator / Review Threshold Settings / <OK>	
		Settings Configuration Message: "Settings Adopted" <CLOSE>	

Besides the above three tasks in Install and Simple “Out-of-the-box” Setup category, additional configuration for maintenance activity was needed for both Oracle Database 11g and IBM DB2 Enterprise 9.5. Typical maintenance activity involves tasks such as

gathering optimizer statistics, defragmentation of tables, automatic SQL Tuning, and backup/recovery.

For IBM DB2 Enterprise 9.5, the following steps were needed on a one-time basis for configuration of maintenance tasks:

Control Center -> View (Advanced) -> Configure Automatic Maintenance, Select Change Automation Settings, -> Specify Time Window -> Change Offline Window (Create new offline window) -> Specify Activity (Enable Backup and Recovery, Defragment data to offline window)-> Specify Notification -> Review, Ok

For Oracle Database 11g, the following steps were needed on a one-time basis for configuration of Automatic SQL Tuning task:

EM Home Page -> Server Tab, Automated Maintenance Tasks, Click on Configure Button, Click on Configure button for Automatic SQL Tuning ->"Yes" on Automatic Implementation of SQL Profiles, Apply

Since the maintenance configuration only needs to be done once after installation and comparable for both Oracle Database 11g and IBM DB2 Enterprise 9.5, we have not included "Maintenance Configuration" as a separate task in this study. However, an important difference to point here is that in IBM DB2 Enterprise 9.5, the maintenance tasks such as backup/recovery, defragmentation of tables require offline mode. In contrast, Oracle Database 11g can perform backup, defragmentation of segments, and automatic SQL Tuning in online mode, so businesses don't have schedule system down time on a regular basis. This means business using Oracle Database 11g will gain from improved business productivity due to online nature of the maintenance tasks.

Metrics	Oracle	IBM DB2
Time	25 Sec	1 Min 15 Sec
Steps	2	10
Increment Penalty	1	
Adjusted Steps (Complexity)	3	10

Installation and Simple "out-of-box" Setup Tasks Results Summary

Metrics	Oracle	IBM DB2
Time	22 Min 55 Sec	12 Min 40 Sec
Adjusted Steps (Complexity)	37	54

Day-to-Day Database Administration

Task 4: Create user, assign roles/privileges

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Create user with roles, privileges			
1. Activate OEM / Server Tab	1 Min 15 Sec	Control Center /Databases / Select Database / <User and Group Objects > / Add	50 Sec
Select Security/ Users link		Add User / User / <Select From List > / Authorize Connect / Tablespace Tab / <Add Tablespace > ¹²	
Select <Create >		Add User / Add Tablespace / Select Tablespace / <OK>	
Enter User Name and Password / Accept other Defaults		Tablespace Tab / <Grant All>	
Activate Popup for Default Tablespace / Select Tablespace		Select Grant Option <OK>	
Activate Popup for Default Temporary Tablespace / Select Tablespace Object from List / <Select>		Add User / <Apply> / <OK>	
Enter User Name "Object has Been Created Successfully" Page: <GO> / Select Username / Review Details			

Metrics	Oracle	IBM DB2
Time	1 Min 15 Sec	50 Sec
Steps	7	6
Increment Penalty	0	
Adjusted Steps (Complexity)	7	6

¹²On configuration tested, user must be a Windows user.

Task 5: Create Tablespace/Filegroup

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Create table-space			
1. Activate EM Server Tab	55 Sec	Control Center / All Databases / Selected Database / Tablespaces / <Create>	55 Sec
Select the Tablespaces <Create>button and Enter Tablespace Name		Create Tablespace / Enter Name / Select Storage Management <Next>	
General Tab / Enter Tablespace Name / Accept Defaults <ADD>		Create Tablespace / Specify Type <Next>	
Click the Add Button in the Datafile Section		Create Tablespace / Specify Buffer Pool <Next>	
Add Datafile Page / Enter FileName Accepting Defaults / <CONTINUE>		Create Tablespace / Specify Extent and Prefetch Sizes <Next>	
Review Create Tablespace / Information: Accepting Defaults / <OK>		Create Tablespace / Describe Hard Drive <Next>	
Confirmation / Tablespaces: Verify Creation		Create Tablespace / Specify Dropped Table Recovery Option <Next>	
		Create Tablespace / Review <Finish>	
		Create Tablespace / Command Completed Successfully / <CLOSE>	

Metrics	Oracle	IBM DB2
Time	55 Sec	55 Sec
Steps	7	9
Increment Penalty		
Adjusted Steps (Complexity)	7	9

Task 6: Add more space to the database

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Add more space to the DB			
1. Activate EM ServerTab	55 Sec	Control Center / Database / Manage Storage / Add Automatic Storage	40 Sec
Select Datafiles Link		Activate Enterprise Manager <Server> Tab ¹³	
Select an Existing Datafile / Select <CreateLike>		Add Storage (Database) / Add New Storage Path /<ADD>	
Datafiles / Create Datafile Page: Enter Filename / Accept Defaults <OK>		File Dialog / Select Path <OK>	
Datafiles / Edit Datafiles / Review Data File Properties / Accept Defaults / <APPLY>		Add Storage (Database) / Add New Storage Path /<OK>	
		Add Storage / Completed Successfully <Close>	

Metrics	Oracle	IBM DB2
Time	55 Sec	40 Sec
Steps	5	6
Increment Penalty		
Adjusted Steps (Complexity)	5	6

¹³DB2 DB with Automatic Storage: Only need to add a new file path —DBMS will add Tablespace Container

Task 7: Create table

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Create table			
Go to the EM Schema TAB->Tables link	1 Min 45 Sec	Control Center / Database / Tables / <CREATE>	2 Min 0 Sec
Instance / Tables: Accept Default Schema / Select <CREATE>		Create Table: Identify Schema and Table Name / Select Schema from Drop Down / Enter Table Name <NEXT>	
Tables / Create Table Enter Table Name / 3 Columns With Data Types (number, varchar2 and date) / <OK>		Column Definitions for New Table: <ADD>	
Instance / Tables / Confirmation / Enter Search String for New Table / <GO>		Column Definitions for New Table: Enter Name / Select Data Type / INT Data Type Properties <APPLY>	
Select Table Name from Search Results / Verify Details		Column Definitions for New Table: <ADD>	
		Column Definitions for New Table: Enter Name / Select Data Type / Varchar Data Type Properties <APPLY>	
		Column Definitions for New Table: <ADD>	
		Column Definitions for New Table: Enter Name / Select Data Type / Timestamp Data Type Properties <APPLY>	
		Change Column Definitions <Next>	
		Define Data Partitions <Next>	
	Specify Space for Sorting Table Data <Next>		

Oracle		IBM DB2	
Step	Time	Step	Time
		Define Keys <Next>	
		Set Up Clustering <Next>	
		Define Check Constraints <Next>	
		Review Create Table Actions <Finish>	
		Message: Completed Successfully <Close>	

Metrics	Oracle	IBM DB2
Time	1 Min 45 Sec	2 Min 0 Sec
Steps	5	16
Increment Penalty		
Adjusted Steps (Complexity)	5	16

Task 8: Create index

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Create index			
Go to the EM Schema TAB->Schema link. / Select <INDEXES>	50 Sec	Control Center / Database / Index / Create	1 Min 10 Sec
Select <Create>		Specify Table Schema and Name: Select Schema / Select Table / <Next>	
Enter Index Name / Accept Defaults / Enter Table Name / <Populate Columns>		Specify Schema and Index Name <Next>	
Accept Default Sort Order / Specify Indexed Column by Entering Number in ORDER Column / <OK>		Specify Index Keys: <Set Focus on Column> / <Move Column to Selection List> / Set Sort Order / <Next>	
Confirmation Page: Accept Default Search By Schema.TableName: Enter Table Name to Review Details		Specify Index options / Accept Default (Do not Customize for Performance) / <NEXT>	

Oracle		IBM DB2	
Step	Time	Step	Time
		Review Actions / <Finish>	
		Success Notification Message / <Finish>	

Metrics	Oracle	IBM DB2
Time	50 Sec	1 Min 10 Sec
Steps	5	7
Increment Penalty		
Adjusted Steps (Complexity)	5	7

Task 9: Reclaim Lost Space due to Fragmented Data

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Determine if any tables in application tablespace need to be reorganized.			
EM Home Page: Space Summary, Problem Tablespaces link	35 Sec	Defragmenting of Tablespaces is Configured as Part of Setting up System Maintenance. This task is Being Treated as Part of the Installation and Setup Procedures	0 Min 0 Sec
Segment Advisor Details / <Recommendation Details>		Defragmentation Runs During Scheduled Maintenance Windows in offline mode	
Segment Advisor: Recommendation Details for Tablespace Select <Segment Check Box> <Implement>			
Shrink Segment: Options - Select Default: Compact Segments and Release Space <Implement>			
Shrink Segment: Schedule -- Accept Defaults: Run Immediately			
Scheduler Job: Job Running			

Metrics	Oracle	IBM DB2
Time	35 Sec	0 Min 0 Sec
Steps	6	0
Increment Penalty		
Adjusted Steps (Complexity)	6	0

Task 10: Load data from a text file

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Data Loading			
Enterprise Manager <Data Movement>	3 Min 10 Sec	Control Center / Databases / Selected Databases / Table / Import	40 sec
Data Movement <Load Data from User Files>		Import Table / Select File / File Dialog / <OK>	
Load Data: Generate or Use Existing Control File: Accept Defaults <Continue>		Import Table / File Type / Delimiter Options /Set Column Delimiter / <OK>	
Load Data: Data Files: Select Data File <Next>		Import Table / Log File / File Dialog / <OK>	
Search and Select: File <Select>		Import Table / Accept Defaults / <OK>	
Load Data: Data Files: Confirm Selected Import File <Next>	20 Sec	Import Table / Progress / Run Until Completion <OK>	1 Min, 45 Sec
Load Data: Table and File Format: <SelectExistingTable> ; <Set Data File Format>;<Next>			
Load Data: File Format Attributes: <Set Delimiters> ; <Next> (Accepting Other Defaults)			
Load Data: Select <Direct Path> method <Next>			
Load Data: Options Accept Defaults <Next>			

Oracle		IBM DB2	
Step	Time	Step	Time
Load Data: Schedule: Enter Job Name / Accept other Defaults (run immediately) / <Next>			
Load Data: Review: <Control File Contents>			
Review: Control File Contents: Modify SQL Function for Statement	(1 Command Line Penalty ¹⁴)		
Load Data: Review: <Submit Job>			
Job Activity: Confirmation - Page Refreshes Until Completion			

Metrics	Oracle	IBM DB2
Time	3 Min 35 Sec	2 Min 25 Sec
Steps	15	6
Command Line Penalty	1	
Adjusted Steps (Complexity)	16	6

¹⁴ While not requiring exiting to an editor, modifying SQL code is being considered using the command line.

Task 11: Configure Adaptive Thresholds (baselines) and Perform Workload Comparisons¹⁵

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Adaptive Thresholds and Workload Comparisons			
Set adaptive thresholds (based on seven day moving window)	35 Sec	Weekly Metrics¹⁶	1 Min 5 Sec (The Weekly Metrics Steps are repeated and counted five times to represent multiple iterations)
OEM Home Page / <Baseline Metrics Threshold>		Configuration	
Threshold Configuration <Quick Configuration>		Control Center/ Performance Expert/ Performance Warehouse / Expert / Processes / Select Performance Report / <Copy>	
Quick Configuration: Baseline Metric Thresholds /		Copy Process / Select Folder / Update Report Name /	

¹⁵ There are several different operations under this task. The operations have been given descriptive titles that reflect the differences between the products and are in bold. The titles are not counted as steps. For purposes of clarity, there are also subtotal rows listed.

¹⁶ Weekly analysis is date specific. The time period has to be configured every time data is needed. To create the equivalent of a Seven Day Moving Window (an Oracle Database 11g feature), this set of tasks must be repeated for each day the data is to be reported. For purposes of this white paper, we are assuming configuring five weekday reports, changing the dates for each instance when run. The total time and total step counts reflect these iterations; the steps and time shown in the table itself is for a single iteration, with a subtotal shown as well.

Oracle		IBM DB2	
Step	Time	Step	Time
Select Profile (OLTP) <Continue>		<OK>	
Quick Configuration: Review Selected Profile Settings/Review/ <Finish>		Select Copied Report In Selected Folder /Steps/ <Properties> / <Options> / Set Date From = Process Execution date -"8 " / Set Date To = Process Execution Date -"1" /) Accept Default Start/End Times / Select <Database> to monitor / <OK>	(1 Inc. Penalty = 1 step)
Create Baseline for known period (single day)		Select Report / <Execute >	
OEM Home Page /<Server>		Warehouse Process Scheduler/<Start Tab> / <Execute Repeatedly According To Schedule >	
(Server Statistics)		Warehouse Process Scheduler / <Schedule Tab> / Execute Every Monday at "00:05" / <Finish>	
		Configuration Subtotal	6 Steps, 1 Increment Penalty Step times 5 iterations = 35 Steps
		Review and Analysis	
		Review analysis reports and manually adjust thresholds periodically to reflect workload type and patterns	1 Step plus 2 Context Penalties or 4 Steps = 5 steps
Server / Statistics Management / <AWR Baselines >	30 Sec	Daily Metrics	1 Min 5 Sec
AWR Baselines / <Create>		Control Center	(1 Inc.

Oracle		IBM DB2	
Step	Time	Step	Time
		/Performance Expert / Performance Warehouse / Expert / Processes / Select Performance Report / <Copy>	Penalty = 1 step)
Create Baseline / Select Type / <Single>/<Continue>		Copy Process / Select Folder / Update Name / <OK>	
Create Single Baseline / <Enter Name>/ <Set Interval> / <Finish>		Select Copied Report In Selected Folder /Steps/ <Properties> / <Options> / Set Date From = Process Execution Date / Set Date To = Process Execution Date / Set From Time = 08:00 / Set To Time = 12:00 / Select <Database> to Monitor / <OK>	(1 Inc. Penalty = 1 Step)
AWR Baselines / Review and Confirm		Select Report / <Execute >	
Workload Comparison to Baselines (single day and seven-day moving window) <small>¹⁷</small>	25 Sec	Warehouse Process Scheduler/<Start Tab> / <Execute Repeatedly According To Schedule >	
		Warehouse Process Scheduler / <Schedule Tab> / Execute Every Monday at "00:05" / <Finish>	
		Subtotal Daily Metrics	6 Steps + 2 Increment Penalty Steps + 4 Context Penalty Steps = 12 Steps

¹⁷ Comparison to Baseline and Review and Analyze reports steps are not measured in this analysis. See the body of the report for more information.

Oracle		IBM DB2	
Step	Time	Step	Time
		Review and Analysis	5 Min
		Review analysis reports and manually adjust thresholds periodically to reflect workload type and patterns	1 Step plus 2 Context Penalties or 4 Steps = 5 Steps

Metrics	Oracle	IBM DB2
Time	1 Min 30 Sec	13 Min 30 Sec
Steps	12	38
Increment Penalty	3	7
Context Switch Penalty		12
Adjusted Steps (Complexity)	15	57 ¹⁸

Day-to-Day Database Administration Tasks Results Summary

Metrics	Oracle	IBM DB2
Time	11 Min 20 Sec	21 Min 30 Sec
Adjusted Steps (Complexity)	67	106

¹⁸ The “Weekly Metrics” steps are counted five times, therefore the context penalties and increment penalties are also counted five times.

Backup & Recovery

Task 12: Configure and Perform Full Database Online Backup

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Configuring for online backup			
N/A (Oracle 11g Recommended backups are configured and scheduled automatically upon database creation. backup is online)		Configuration is Part of Setup/Maintenance/Backups Automatically Scheduled Backup is Offline	

Metrics	Oracle	IBM DB2
Time	0	0
Steps	0	0
Increment Penalty		
Adjusted Steps (Complexity)	0	0

Task 13: Recover Dropped Table

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Recover dropped table			
Enterprise Manger <Availability>	1 Min 10 Sec	Not Counted as steps: Prerequisite: Tablespace Created with "Drop Table Recovery Option" - Default with Create Table Wizard Prerequisite: Database Logging Set to Archival	11 Min 20 Sec
Availability: <Perform Recovery>		Control Center / Command Editor (A) / Enter "Select Statement to Get ID of Dropped Table"	(Command Line Penalty - 2 steps)
Perform Recovery: RecoveryScope: <Tables> / <Flashback Dropped Tables> / <Recover>		Control Center / <Select Database to Restore> / <Right Mouse Down> / <Restore Database>	

Oracle		IBM DB2	
Step	Time	Step	Time
Perform Object Level Recovery: Dropped Objects Selection.Results – <EnterSchemaName RecycleBin> <GO> /Accept Default Dropped Table In RecycleBin <NEXT>		Restore Data Wizard / Confirm Details (selected database) / <Next>	(Context Penalty – 2 steps)
Perform Object Level Recovery: Rename: Accept Default / <NEXT>		Restore Data Wizard / Replace Selected Tablespace / Select Tablespace w/Deleted Container / <Next>	
Perform Object Level Recovery: Review: Accept Default <SUBMIT>		Restore Data Wizard / Select Backup Image / <Next>	
Confirmation (Tables Have Been Flashed Back)		Restore Data Wizard / Change Tablespace Containers / Accept Defaults / <Next>	
		Restore DB and Roll Forward + Roll Forward to End of the Logs / <Next>	
		Restore Data Wizard / Select Performance Options for Restore / Accept Defaults / <Next>	
		Restore Data Wizard / Schedule Task Execution / Run Now / <next>	
		Restore Data Wizard / Review Actions / <Finish>	
		DB2 Message: Restore Successful / <Close>	

Metrics	Oracle	IBM DB2
Time	1 Min 10 Sec	11 Min 20 Sec
Steps	7	11
Context Switch Penalty		2
Command Line Usage Penalty		2
Adjusted Steps (Complexity)	7	15

Task 14: Recover Data File¹⁹

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Recovery: Application Data File of Size 1.2 GB is Lost.			
Database Instance: Cannot Connect - <Perform Recovery>	1 Min 55 Sec	Control Center / <Select Database to Restore> / <Right Mouse Down>/<Restore Database>	45 sec
Perform Recovery: Credentials – Enter UserName / Enter Password / <Continue>		Restore Data Wizard / Confirm Details (Selected Database) / <Next>	
Database Login – Enter User Name / Enter Password / Connect As: <SYSDBA> / <LOGIN>		Restore Data Wizard / Restore Tablespace /Select Tablespace <Next>	
Perform Recovery: Oracle Advised Recovery <Advise and Recover>		Restore Data Wizard / Select Backup Image / <Next>	
View and Manage Failures: <Review Failure> (Missing DataFile) / <ADVISE>		Restore Data Wizard / Change Tablespace Containers	
Manual Actions (Recommendation) : Restore File <Continue with Advise>		Restore Data Wizard / Roll Forward Options / Restore and Roll Forward to End of Logs	
Recovery Advice: <Recover>	25 Sec	Restore Data Wizard / (no option) – Complete Restore and Return to Active State <Next>	
		Restore Data Wizard / Restore Options /Accept Defaults (Availability: Online / No History File Replace , No Log File Restore) / <Next>	

¹⁹The time to actually perform a datafile recovery in a real-world scenario varies greatly based on file size and the speed of the system performing the operation.

Oracle		IBM DB2	
Step	Time	Step	Time
		Wizard / Select Performance Options for Restore / Accept Defaults / <Next>	
		Wizard Schedule Task execution / Run Now / <next>	
		Wizard / Review Actions / <Finish>	
		DB2 Message: Restore Successful / <Close>	

Metrics	Oracle	IBM DB2
Time	1 Min 55 Sec	45 Sec
Steps	7	12
Increment Penalty		
Adjusted Steps (Complexity)	7	12

Task 15: Recover from erroneous transaction

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Scenario: Due to Human Error, a Transaction was Submitted by Mistake			
Enterprise Manager / Schema Tab – Select <Tables>	1 Min 55 Sec	Control Center / <Select Database.Tablespace to Restore> / <Right Mouse Down> / <Restore Tablespace>	11 Min 20 Sec
Tables/Search – Enter Schema Name: Enter Name / <Go> / Select <TableName>		Restore Data Wizard / Confirm Details (selected database) / <Next>	
View Table: Actions <Flashback Tableby Timestamp> Go -> Identify /Input timestamp to recover from		Restore Data Wizard / Restore Tablespace /Select Tablespace <Next>	
Perform Object-Level Recovery: Point-in-Time - <Select Flashback to Timestamp (accepting default date)> Set <Hour><Minute><Next>		Restore Data Wizard / Select Backup Image / <Next>	
Perform Object Level Recovery: Flashback Tables – Accept Defaults <Next>		Restore Data Wizard / Change Tablespace Containers	
Perform Object Level Recovery: Flashback Tables – ERROR – Table Row Movement Not Enabled <ENABLE TABLE ROW MOVEMENT>		Restore Data Wizard / Roll Forward Options / Restore and Roll Forward to Point in Time / Set to Time Prior to Error Transcript	
Confirmation – Table Modified Successfully <NEXT>		Final State of Database / Complete Restore and Return to Active State <Next>	

Oracle		IBM DB2	
Step	Time	Step	Time
Perform Object Level Recovery: Review <SUBMIT>		Restore Data Wizard / Restore Options /Accept Defaults (Availability: Online / No History File Replace, No Log File Restore) / Next>	
Confirmation: Tables Have Been Flashed Back		Wizard / Select Performance Options for Restore / Accept Defaults / <Next>	
		Wizard Schedule Task Execution / Run Now / <next>	
		Wizard / Review Actions / <Finish>	
		DB2 Message: Restore Successful / <Close>	
		Manually Rerun the Transactions Lost Due to Point-in-Time Recovery in Affected Tablespace (not accounted for in time measurement)	

Metrics	Oracle	IBM DB2
Time	1 Min 55 Sec	11 Min 20 Sec
Steps	8	12
Increment Penalty	0	0
Context Switch Penalty	0	4
Adjusted Steps (Complexity)	8	16

Task 16: Recover from Multiple Failures

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Recovery scenario:			
Instance Fails to Start on Re-activation	3 Min 50 Sec	Control Center / <Select Database.Tablespace to Restore> / <Right Mouse Down> / <Restore Tablespace>	12 Min 5 Sec
Perform Recovery: Credentials – Enter UserName / Enter Password / <Continue>		Restore Data Wizard / Confirm Details (selected database) / <Next>	
Cannot Connect to Database Instance: <Perform Recovery>		Restore Data Wizard / Restore Tablespace /Select Tablespace <Next>	
Database Login – Enter User Name / Enter Password / Connect As: <SYSDBA> / <LOGIN>		Restore Data Wizard / Select Backup Image / <Next>	
Perform Recovery/Oracle Advised Recovery: <Advise and Recover>		Restore Data Wizard / Change Tablespace Containers	
Manual Actions: Review Manual Action Recommendations / <Continue with Advise>		Restore Data Wizard / Roll Forward Options / Restore Only	
Review Auto Repair Resolution: <Submit Recovery Job>		Final State of Database / Leave Roll forward Pending <Next>	
Successful Recovery Results / Recovery Succeeded <Open Database>		Restore Data Wizard / Restore Options /Accept Defaults (Availability: Online / No History File Replace , No Log File Restore) / <Next>	
		Wizard / Select Performance Options for Restore / Accept Defaults / <Next>	
		Wizard Schedule task execution / Run Now / <next>	

Oracle		IBM DB2	
Step	Time	Step	Time
		Wizard / Review Actions / <Finish>	
		DB2 Message: Restore Successful / <Close>	
		Control Center / <Select Database.Tablespace to Restore> / <Right Mouse Down> / <Restore Tablespace>	
		Restore Data Wizard / Confirm Details (selected database) / <Next>	
		Restore Data Wizard / Restore Tablespace /Select Tablespace <Next>	
		Restore Data Wizard / Select Backup Image / <Next>	
		Restore Data Wizard / Change Tablespace Containers	
		Restore Data Wizard / Roll Forward Options / Restore and Roll Forward to End of Logs	
		Restore Data Wizard / (no option) – Complete Restore and Return to Active State <Next>	
		Restore Data Wizard / Restore Options /Accept Defaults (Availability: Online / No History File Replace, No Log File Restore) / <Next>	
		Wizard / Select Performance Options for Restore /Accept Defaults / <Next>	
		Wizard Schedule Task Execution / Run Now / <next>	

Oracle		IBM DB2	
Step	Time	Step	Time
		Wizard / Review Actions / <Finish>	
		DB2 Message: Restore Successful / <Close>	

Metrics	Oracle	IBM DB2
Time	3 Min 50 Sec	12 Min 5 Sec
Steps	8	24
Increment Penalty		
Context Switch Penalty		
Adjusted Steps (Complexity)	8	24

Backup & Recovery Task Summary

Metrics	Oracle	IBM DB2
Time	8 Min 50 Sec	35 Min 30 Sec
Adjusted Steps (Complexity)	30	67

Performance Diagnostics & Tuning Tasks

Task 17: Diagnose Performance Problem

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Configure System to Identify Top Resource-Consuming SQL			
Enterprise Manager / <Server>	20 Sec	Control Center / <Selected Database> / <Activity Monitor>	1 Min 0 Sec
Go to the EM home page and review the last ADDM report		Activity Monitor / Introduction / Confirm DB info / <NEXT>	
Drill down to the appropriate finding		Activity Monitor / Monitoring Tasks / <Resolve Performance Degradation> / <Next>	
		Activity Monitor / Filter / Monitor All Applications / <Next>	
		Activity Monitor / Summary / Confirm Monitoring Task Settings / <Next>	
		Activity Monitor / Application Degradation / Select Report: Top Time-Consuming SQL	
		Top Time-Consuming SQL / Review SQL Statement Metrics / Select Query To Analyze / Explain Query /	
		Access Plan Graph / Review Query Explain Analysis	

Metrics	Oracle	IBM DB2
Time	20 Sec	1 Min 0 Sec
Steps	3	8
Increment Penalty		
Context Switch Penalty		
Adjusted Steps (Complexity)	3	8

Task 18: Fix Performance Problem - SQL Tuning

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Identify Resource-Intensive SQL			
Automatic SQL Tuning is enabled as part of Installation and Setup Category	0 Sec	Control Center / <Selected Database> / <Activity Monitor>	1 Min 45 Sec
		Activity Monitor / Introduction / Confirm DB Info / <NEXT>	
		Activity Monitor / Monitoring Tasks / <Resolve Performance Degradation> / <Next>	
		Activity Monitor / Filter / Monitor All Applications / <Next>	
		Activity Monitor / Summary / Confirm Monitoring Task Settings / <Next>	
		Activity Monitor / Application Degradation / Select Report: Top Time-Consuming SQL	
		Top Time-Consuming SQL / Review SQL Statement Metrics / Select Query To Analyze / Optimize Workload /	
		Design Advisor / Introduction / Review Features / <Next>	
		Design Advisor / Features / Configure / Accept Defaults / <NEXT>	
		Design Advisor / Workload / Define New Workload / Enter Name / <Next>	

Oracle		IBM DB2	
Step	Time	Step	Time
		Design Advisor / Update Catalog Statistics / Select All Tables / <Next>	
		Design Advisor / Options / Specify Recommendation Options / Accept Defaults / <Next>	
		Design Advisor / Calculations / Accept Defaults / <Next>	
		Design Advisor / Recommendations / Select Recommendations / Accept All / <NEXT>	
		Design Advisor / Unused Objects / Review -- Keep All / <Next>	
		Design Advisor / Schedule Recommended Task Execution / <Run Now>	
		Design Advisor / Recommendations Summary / Review / <Finish>	
		DB2 Message / Review Completion Status (All of the DB2 tasks are performed twice for illustration purposes.)	

In Oracle Database 11g, Automatic SQL Tuning is enabled as part of Installation and Setup Category. A DBA can review at system tuning if necessary by the following navigation: EM Server -> Automated Maintenance Tasks -> Automatic SQL Tuning -> View Automatic SQL Tuning Report. In general, this review step is not required since system tuning is done after careful validation of the tuning recommendations.

Metrics	Oracle	IBM DB2
Time	0Sec	1 Min 45 Sec
Steps	0	36
Increment Penalty		
Context Switch Penalty		
Adjusted Steps (Complexity)	0	36 ²⁰

Task 19: Tune Memory

Oracle		IBM DB2	
Step	Time	Step	Time
Task: Tuning memory			
Automatic Memory Tuning is enabled as part of the installation	0 Min 0 Sec	Self-Tuning Memory is Automatically Enabled (for non-partitioned databases) when a New Database is Created	0 Min 0 Sec

Metrics	Oracle	IBM DB2
Time	0	0
Steps	0	0
Increment Penalty		
Context Switch Penalty		
Adjusted Steps (Complexity)	0	0

Performance Tuning Task Summary

Metrics	Oracle	IBM DB2
Time	20 Sec	2 Min 45 Sec
Adjusted Steps (Complexity)	3	44

²⁰ As noted, for illustration purposes the total time and step count is for two iterations of the task.

Appendix VII – Glossary of Task Areas/Tasks

The following glossary will be updated as required to reflect changes in the products and tasks performed.

Database setup and configuration (workload task area): A workload task area that encompasses all of the operations that a DBA would need to perform in order to accurately install and configure either product prior to using it in a real-world application.

Install db/software/out-of-box setup: The workload task that identifies the process of installing and configuring either of these products for the first time.

Create new database server/instance: The workload task that allows a DBA to create a second database server/instance on a computer already running one or more instances of the RDBMS.

Set up proactive monitoring: Proactive monitoring is a tool that allows a DBA to identify problems with the RDBMS before they become performance or operational issues. Setting up monitoring is crucial in providing a highly reliable system.

Day-to-day database administration (workload task area): The workload task area where all of the routine DBA operations occur, such as creating database users and objects (for example, tables, indexes, triggers, procedures), as well as granting and revoking roles/privileges, plus database table/tablespace/data file sizing.

Create index: The workload task that allows DBA's to create an index on a set of columns in a table that are heavily queried in order to speed the execution of queries run against that table.

Create index on partitioned table: The workload task that allows DBA's to create local indexes on the partitioned table to help speed the execution of queries run against the table.

Create Partitioned Table: The workload task that allows DBA's to create a partitioned table object to store information inside a schema in the database.

Create user with roles, privileges: The workload task that allows DBA's to manage user security in the database.

Create schema: The workload task that allows DBA's to create a new schema.

Create table: The workload task that allows DBA's to create a table object to store information inside a schema in the database.

Reclaim wasted space from tables with fragmented data: The workload task that allows DBA's to pack/shrink the database after prolonged use in order to consolidate space and optimize performance.

Load data from text file: The workload task that allows a DBA to load information from an external source such as a flat file or spreadsheet into one or more tables in the database.

Create tablespace: The workload task that allows DBA's to build a new tablespace for use by a particular database instance or schema.

Add space to tablespace: The workload task that allows a DBA to add data files to a tablespace, thus increasing the amount of space available to all of the objects contained in that database.

Backup and recovery tasks (workload task area): The workload task area where all tasks pertaining to database backup and recovery are performed.

Configure and perform full backup: The workload task that allows DBA's to schedule and execute regular system backups as part and parcel of standard system fault tolerance operations.

Recover dropped table: The workload task that allows a DBA to recover a table that has been inadvertently dropped from the database by a DBA, developer, or power user in the course of working with the database.

Recover data file: The workload task that allows the DBA to recover a data file from a backup copy in the event of a media failure.

Recover from erroneous transaction: The workload task that allows a DBA to undo a mistakenly executed transaction in order to recover objects ill-affected by this transaction to the state they were in before the transaction was executed.

Performance diagnostics and tuning tasks (workload task area): The workload task area where all performance-related diagnostic, tuning, and optimization tasks are performed.

Diagnose performance problem: The workload task that allows a DBA to analyze a poorly performing system in order to assess the performance problem as a prerequisite to performing the systems optimization functions required to bring the level of system performance to an optimal state (such as creating additional indexes, tuning a query/procedure, defragmenting tablespaces, or adjusting the server's memory configuration).

Tune instance memory: The workload task that allows DBA's to adjust server/instance system memory configuration parameters in order to efficiently support the load placed on the system by all applications that access the database.

Fix performance problem (tune SQL statement): The workload task that allows DBA's to optimize a poorly running query/procedure so that it executes in an efficient manner.