EDISON Group

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

Comparative Management Costs Study

Oracle Database 11g vs. Microsoft SQL Server 2008

First Publication: July 2009
Produced by: Steve Mintz, Senior Analyst; Craig Norris, Editor; Barry Cohen, Author and Editor-in-Chief

Table of Contents

Executive Summary

This report is the latest in a series of Comparative Management Cost Studies (CMCS) comparing Oracle Database with database management systems offered by other leading enterprise software vendors. This edition of the study compares Microsoft SQL Server 2008 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 average database administrators (DBA) perform everyday administrative tasks, Oracle Database 11g holds a substantial advantage over Microsoft SQL Server 2008. Our detailed analysis reveals that:

- DBAs can perform typical administrative functions in 41 percent less time when using Oracle Database 11g compared to Microsoft SQL Server 2008.
- Oracle Database 11g requires 43 percent fewer steps for the same set of standard RDBMS tasks than Microsoft SQL Server 2008 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 $33,520.47 per year per DBA by using Oracle Database 11g rather than Microsoft SQL Server 2008.

The main areas of difference we discovered between the two products were:

- Backup & Recovery: In the case of backup and recovery tasks, Oracle Database 11g offers architectural and functional capabilities beyond those offered by SQL Server 2008. Oracle Database 11g took 53 percent less time and 60 percent fewer steps than Microsoft SQL Server in backup and recovery tasks.
- Performance Diagnostics and Tuning: In this category, Oracle Database 11g demonstrated a significant 87 percent savings in time.

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 Microsoft and Oracle have provided their customers with efficient tools for management of their respective database systems. But when the cost of operations is analyzed, Oracle Database 11g can provide organizations with 41 percent annual DBA-related cost savings over Microsoft SQL Server 2008.

About This Report

This report documents the results of a head-to-head product comparison examining the database administration functions of Oracle Database 11g and Microsoft SQL Server 2008. 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 Microsoft SQL Server 2008 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 Diagnostics and Tuning Packs

For Microsoft SQL Server 2008

- SQL Server Management Studio
- SQL Server Installation Center
- Database Engine Tuning Advisor
- SQL Server Profiler

The study measures (in quantitative and qualitative terms), the relative manageability of Oracle Database 11g and Microsoft SQL Server 2008, 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 (DBAs/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 Microsoft

SQL Server 2008 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 Microsoft SQL Server 2008 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 Installation and Setup
- 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.

- 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 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. Oracle has sponsored the studies in response to a challenge made by Edison to Oracle claims that Oracle Database was significantly easier and thus less costly to manage than the other major database management systems. Edison's research has borne out those claims for each iteration of this research. 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 following elements.

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.

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 system-affecting steps it takes to complete a given task, where a step is defined as a task component that effects a change of state to the database.

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 with the installation wizard is an increment and not a step.

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 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 Microsoft SQL Server 2008.

- 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 Administrator's Guide.
- The baseline workload weighting and task list was further checked for consistency against Microsoft documentation: SQL Server Books Online.
- Independent professional Oracle and SQL Server 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. As of this time, Edison has been unable to discover new workload studies that have been published or are otherwise available, at no charge, to the general public.

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 in 2002. 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
- Adaptive thresholds and workload comparisons

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 Microsoft SQL Server 2008 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 it has an advance copy.)

If we insert the averaged DBA compensation found in the surveys [2] into the formula below, we arrive at the following quantitative management cost (MC) saving calculation.

Median DBA Salary * (DBA time savings) = $81,757.25 *41% = $33,520.47

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 DBAs 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 Microsoft SQL Server 2008 really compare across the board. In this section, for brevity we often refer to Oracle Database 11g as "Oracle" and Microsoft SQL Server 2008 as "SQL Server"

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)	
	SQL Server	Oracle	SQL Server	Oracle
Category Sub-Total	21.58	22.92	12	5
% difference — (SQL Server - Oracle)/SQL Server	-6%		58%	
DBA Workday Savings	0%		3%	

The results for were mixed for both SQL Server and Oracle. While Microsoft SQL Server 2008 took 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, these differences are insignificant.

Qualitative Analysis

The testing results for this category may seem contradictory. Microsoft SQL Server 2008 took less time to install than what was required for installing Oracle Database 11g. On the other hand, SQL Server required more steps to accomplish the set of tasks. The main reason for this discrepancy is that SQL Server 2008 goes through extensive system and supporting software checks before beginning installation. Edison believes these tests are valuable aids in identifying potential problems and can result in a more stable installation.

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 storage models, while

subsequent installs are more predetermined. Both Microsoft SQL Server 2008 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 typical applications.

When Oracle Database 11g is installed, configuration of basic system maintenance tasks (backup, memory tuning, etc.) are incremental steps during setup. By contrast, SQL Server 2008 requires a separate post-installation task — creating a new Maintenance Plan — for these tasks. Once configured, both products perform automatic backup within scheduled maintenance windows. In Microsoft SQL Server 2008, defragmentation (called Shrink Database) and other tasks can be configured and scheduled in the same Maintenance Plan operation. The net result is that these tasks, which are evaluated here under the Day-to-Day, Backup and Recovery, or Performance Tuning task categories, become largely automatic.

In Edison's opinion, the time spent waiting for an Oracle Database 11g instance to install is less critical to the correct implementation of a database server than is the complexity of the tasks. While the time difference is real, the weighted workload factor is zero. Even the relative complexity in favor of Oracle Database 11g is not significant. Both products are quick and easy to install. The differences in complexity reflect Microsoft SQL Server's continuing feature growth with the concomitant need to control which of those features are put to use in each installation.

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)	
	SQL Server	Oracle	SQL Server	Oracle
Category Sub-Total	15.32	10.08	20	12
% difference — (SQL Server - Oracle)/SQL Server	34%		40%	
DBA Workday Savings	12%		14%	

Qualitative Analysis

Edison analysts were not completely surprised by the results of the tests in this category: While Microsoft's marketing materials for SQL Server 2008 had touted the product's improved ease of management for everyday tasks Oracle retained the edge in the tasks involved in this category. Thought it did take less time to perform the some of the tasks with SQL Server 2008 than for Oracle Database 11g, the manual actions needed for the remaining tasks quickly erased any advantage SQL Server may have enjoyed over Oracle.

There is one caveat in reviewing these results. The new task, described below, of configuring historic monitoring and performing comparative operational and performance analysis is not adequately represented by a mere step count and timing. This is because beyond merely configuring the features and reviewing the data, most Microsoft SQL Server 2008 deployments will require the creation of custom reports so that the specific data needed can be more easily accessed. One key example is how queries are tracked in the two products. Oracle's tools automatically identify and filter data so that problem SQL queries, for example, can be quickly identified. The default reporting in SQL Server 2008 includes system databases and internal queries in the list of most active queries, with the potential to obscure important data. A custom report would filter out system queries so that only application-based queries could be analyzed thereby providing a more relevant view of data to the DBA. While creation of custom reports was not included in this study but their creation would add substantially to the time and complexity of configuration. Once these reports are configured, there is a conservative estimate assigned to the time and steps needed to access these reports in order to complete the task in SQL Server.

New Task for this Study

One of a DBA's most important daily tasks is to monitor the system performance. 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 Diagnostics 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 a web-based interface in Enterprise Manager to configure Adaptive Thresholds with a few mouse clicks and interpret the 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.

Microsoft utilizes features of SQL Server Profiler and SQL Server Trace to configure and monitor SQL Server Events. SQL Data Collector is used to collect the data produced by Server Profiler and Server Trace and store it in a data warehouse. SQL Server Analysis Services and SQL Server Reporting Services features are used to provide the data warehouses for baseline and historic performance data and for running reports on that data.

SQL Server Data Collector is the central interface for configuration and management of tracked events. Once configured, collection and storage of Disk Usage, Server Activity, and Query Statistics are stored for the instance upon which it is installed. The default storage period is 14 days, but that duration can be adjusted or the data can be stored indefinitely[3]. The data warehouse can be installed on the same server being monitored or on a different server.

[3] Indefinite storage of performance data will extract a penalty in storage capacity utilization which must be understood and planned for.

When Data Collector is running, the DBA has access to a range of standard reports. Custom reports can also be created. Doing so requires installation of the free [4] Microsoft SQL Server 2008 Solution Pack. Creation of custom reports is outside the scope of this study.

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 is 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.

With Oracle Database 11g, the DBA can set up separately scheduled snapshots for multiple databases and time periods. A SQL Server 2008 DBA can set up collections for different databases and schedule snapshot periods, but reviewing history in a moving window is controlled in the reporting interface, not at capture time.

The default reports for SQL Server 2008 do not provide for the direct comparison of historic and current data. While Edison did not evaluate the tools for creating custom reports, it is possible that a report allowing the viewing of current and historic data concurrently can be created.

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.

[4] At time of writing.

Backup and Recovery	Time (min)		Complexity (steps)	
	SQL Server	Oracle	SQL Server	Oracle
Category Sub-Total	17.34	8.08	25	10
% difference — (SQL Server - Oracle)/SQL Server	53%		60%	
DBA Workday Savings	7%		8%	

Qualitative Analysis

Here Oracle won both in time and complexity, requiring only 53 percent of the time with 52 percent the number of steps Microsoft SQL Server 2008 required.

Both Oracle Database 11g and Microsoft SQL Server 2008 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 the maintenance plan operations are configured (with Microsoft SQL Server 2008). Once configured, backups run automatically.

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 two steps and one minute and ten seconds. Recovering a dropped table with SQL Server 2008 requires three times the number of steps and about four times the amount of time. These deltas under-represent the actual differences in complexity and time between the two platforms. Recovering a dropped table in SQL Server requires restoration of a whole database and copying the dropped table back into the damaged database. In our test bed, these operations were straightforward and quick. In a production system, the time required for identifying the last known good database and table, the potential need to restore dropped transactions, and so forth can significantly increase the time, complexity, and even risk involved in restoring a dropped table.
For Task 15 in this study, which assumes transaction loss due to human error and where manually restoring transactions is almost mandatory a complexity penalty was assessed. Since the time required for restoring these transactions is unknowable and may not be a task performed by the DBA, no time penalty was assessed.

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)	
	SQL Server	Oracle	SQL Server	Oracle
Category Sub-Total	18.10	2.42	10	3
% difference — (SQL Server - Oracle)/SQL Server	87%		70%	
DBA Workday Savings	23%		18%	

Of all the areas of comparison, this was the category where Oracle Database 11g revealed its greatest strength. In terms of time and steps metrics, Oracle Database 11g as compared to Microsoft SQL Server 2008 was:

- 87 percent more efficient in time, and
- 70 percent more efficient in steps.

Given that DBAs spend more than a quarter of their time conducting performance diagnostics and tuning related functions, Oracle's advantage in this category translates into DBA workday savings of:

- 23 percent in time, and
- 18 percent in steps.

This is the category where the highest-paid DBA specialists and consultants spend most of their time and, coincidentally, is the segment of the study where Oracle Database 11g demonstrated a considerable advantage over Microsoft SQL Server.

Oracle's advantage in this category was mainly due to the self-diagnostic engine, Automatic Database Diagnostic Monitor (ADDM), that proactively identifies and recommends remedies for performance problems encountered by the system, and the SQL Tuning Advisor that fully automates the complex task of application tuning. Whereas Microsoft SQL Server does have features that facilitate SQL tuning, its solution is not nearly as comprehensive as Oracle's, and multiple tools external to SQL Server Management Studio must be utilized for different aspects of application tuning.

Moreover, in the area of performance diagnostics, SQL Server does not have anything that compares directly to the self-diagnostic capabilities of ADDM. ADDM and SQL Tuning Advisor together give Oracle a significant edge over SQL Server in the performance diagnostics and tuning category, as aptly reflected in the time and complexity (steps) numbers above.

When it comes to complexity, Oracle has succeeded in automating the involved art of performance diagnostics and tuning in such a way that the adoption of Oracle Database 11g will significantly reduce the management costs for any company. Upon completion of this section of our analysis, it became clear to us that Oracle Database 11g maintains the high standard it set in the high-tech realm of automated database performance management.

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. Microsoft SQL Server 2008 provides an array of monitoring tools that can be used separately and together to monitor the server, the database engine, query performance, and workloads. Some of these tools are configured through GUI management interfaces; others require the creation of queries, stored procedures, and even scripts for implementation.

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

The time and complexity measurements reported here focus solely reviewing the standard performance oriented reports produced by the monitoring tools previously configured for other tasks in this study. The diagnostic context is for a relatively simple performance issue. In the case of Microsoft SQL Server 2008 there are several reports that need to be reviewed, accessible from several locations. For Oracle Database 11g, a DBA only needs to access the ADDM report and drill down to the appropriate finding.

Tuning Resource-Intensive SQL

Because the two products under consideration — Microsoft SQL Server 2008 and Oracle Database 11g — are not likely to be found in the simple single-server 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 can automatically tune most SQL performance issues, but for comparison sake Edison makes use of the 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. Automatic SQL Tuning run against a SQL Profile 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; code revisions can be provided to developers for future incorporation without undo disruption to projects in progress.

Oracle Database 11g automatic tuning affects only SQL code. Adding or reconfiguring indexes or partitioning of tables is not automated and must be performed manually. For this study, we used automatic tuning for our metrics. The effort required for creating indexes is treated elsewhere in this study.

Oracle Database 11g Diagnostics 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 Microsoft SQL Server 2008, tuning is separated into two main categories: Query tuning and Database Engine tuning for modifying or adding indexes and making other structural modifications.

SQL Server's Database Engine Tuning Advisor analyzes and provides recommendations for partitioning or for indexes to improve performance. Query tuning requires displaying query execution plans by using Microsoft SQL Server Management Studio and Transact-SQL SET options, and by using SQL Server Profiler event classes in traces. These tools can provide recommendations for manually tuning queries in applications.

Query execution plan guides can be used to optimize deployed applications in a manner that gives similar results to that provided with Oracle Database 11g. Plan guides influence optimization of queries by attaching query hints or a fixed query plan to them. In the plan guide, you specify the Transact-SQL statement that you want optimized as well as either an OPTION clause that contains the query hints you wish to use, or a specific query plan you want to use in optimizing the query. When the query executes,

SQL Server matches the Transact-SQL statement to the plan guide and either attaches the OPTION clause to the query at run time or else uses the specified query plan. Edison did not test this process during this evaluation but, as the above description implies, the process is much more complex than that required for Oracle Database 11g. The time required and the number of steps involved in this process can be significant and undoubtedly, in many cases, would far exceed the results reported here.

Memory Tuning

Memory tuning is another important aspect of performance tuning. Microsoft SQL Server has had automatic memory tuning for several years and versions. Oracle Database 11g now supports automatic memory tuning as well. For SQL Server there are options to manually reconfigure memory tuning, but automatic tuning is the default setting. For Oracle Database 11g 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 Microsoft SQL Server 2008 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 key factors behind Oracle's superior manageability are self-monitoring and diagnostic capabilities, and the increased automation of many otherwise manual but vital DBA tasks such as SQL tuning, space management, and performance diagnostics. These manageability advantages translate into substantial management cost savings for businesses along with higher reliability and availability of their systems.

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: Two 73.4 GB 10 K drives

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

Microsoft

- SQL Server 2008 Enterprise Edition (64 Bit)

Appendix II - What's New?

Microsoft SQL Server and Oracle Database systems support the needs of a wide variety of customer applications, ranging from 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.

Keeping up with this growth without increasing manageability costs is seen by analysts as a critical factor in the continuing success of these systems.

Microsoft SQL Server 2008

Microsoft's product collateral for SQL Server 2008 touts a range of new or enhanced management features. The new or enhanced features relevant to this study include:

- SQL Server Performance Studio
- SQL Server Profiler
- Database Engine Tuning Advisor
- Policy-Based Management
- SQL Server Agent

Of these, SQL Server Performance Studio, SQL Server Profiler and Database Tuning Advisor were the most used in this study. The descriptions below are from Microsoft product literature.

SQL Server Performance Studio

SQL Server Performance Studio is the name given to the suite of tools available for capturing, collecting and analyzing performance data. SQL Server performance data collection is used to help integrate the collection, analysis, troubleshooting and persistence of SQL Server diagnostics information. Using installed data warehousing functionality, a warehouse can be established for baseline and historical comparisons. In addition built-in reports, generated by the installed Report Services module, provide an overview of server activity, disk usage and query activity.

SQL Server Profiler

SQL Server Profiler is used to capture server events for real-time diagnosis, correlation of traces with performance counters for diagnostics.

Database Tuning Advisor

Database Tuning Advisor (DTA) can be used to tune multiple databases from the same workload. The administrator selects the databases to tune and DTA generates indexing and portioning recommendations.

The other two features listed were not critical to this study, but taken together do demonstrate Microsoft's increasing use of policy and automation for managing SQL Server and its other servers. Though outside the scope of this research, Policy-Based Management can be used to define configuration policies and apply them to servers, databases, tables and other targets. Tied with the scheduling capabilities of SQL Server Agent and the sophisticated scripting capabilities of PowerShell enable database administrators to automate regularly performed complex tasks.

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 documentation[5]. 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 Diagnostics and Tuning Packs that were relevant to this study.

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 is an advantage over Microsoft SQL Server 2008, which utilizes a fat client (SQL Server Management Studio) and utilization of additional consoles for performance tuning and other administrative tasks. In addition, there are several instances where recourse to the writing of SQL queries, stored procedures, and other manual tasks are required.

(The descriptions below are extracted from Oracle product literature.[6])

[5] Oracle® Database New Features Guide11g Release 1 (11.1) B28279-03
[6] Available at http://www.oracle.com/database/db_manageability.html

Diagnostics Pack

The key features of the Diagnostics 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.
- 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 Microsoft SQL Server 2008 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 SQL Server 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.)

In terms of disk resources, the physical definitions of what defines a database are essentially the same for both Oracle and SQL Server. From this perspective, a database is a repository of information contained in one or more data files logically organized into one or more tablespaces (in Oracle) or file groups (in SQL Server). Yet a deeper look into the logical structures that define the term "database" reveals a few interesting differences.

First of all, a distinct difference exists in how database users are defined in Oracle and in SQL Server. In Oracle, a user is an atomic entity. In SQL Server, a user has two elements: a global login and a local database user. This difference arises from the fact that Oracle is an open platform product designed to run on many operating systems, and as such cannot rely on the underlying operating system to behave in a certain way. In order to operate in a consistent manner across all supported platforms, it must implement and encapsulate all critical objects (such as users) and functions (such as authentication) that are required by the DBMS. By contrast, SQL Server has both the luxury and the liability of being totally dependent on Microsoft Windows. Given this tight coupling between SQL Server and Windows, SQL Server leverages many Windows features running in the background as operating system services (e.g., Active Directory) to accomplish various operations. Identity management of database users is the main area of significant differences between these products, which brings up the key difference in how these products logically define a database.

In Oracle, databases are equivalent to server instances (except in the case of RAC clusters). Oracle has a schema feature to match ANSI SQL standards, but practically database users can own objects, and the objects owned by a user are considered a schema. While SQL Server has a concept of database login and connection login, Oracle has a single authorization process. Oracle objects are stored, as segments, in tablespaces which are composed of files. These tablespaces are operational management facilities,

and are the nearest Oracle equivalent to the SQL Server database and/or the underlying SQL Server file groups

The preceding section has highlighted the main architectural and terminology discrepancies between the Oracle 11g Database and Microsoft SQL Server. The overarching differences in the articulation of core product terminology, and the way the vendors advise companies to manage large-scale database installations using their products, is important in understanding the Comparative Management Costs that these products will exact for an IT organization.

Appendix IV - Detailed Test Results

Task	Description	Time (mins)		Steps	
		SS08	Oracle	SS08	Oracle
Installation and Simple "Out-of-Box" Setup					
1	Install db/software/out-of-box setup	15.75	16.08	6	2
2	Create 2nd database instance	4.67	6.42	5	2
3	Setup proactive monitoring	1.17	0.42	1	1
	Category Sub-Total	21.59	22.92	12	5
	% difference -- (SS08 - Oracle)/SS08	-6%		58%	
	DBA Workday Savings (Difference * Weighting Factor)	0%		3%	
Day-to-Day Database Administration					
4	Create user with roles, privileges	0.92	1.25	2	1
5	Create tablespace/filegroup	0.67	0.92	2	1
6	Add space to database	0.45	0.92	1	1
7	Create table	1.08	0.50	1	1
8	Create index	0.45	0.83	1	1
9	Reclaim space due to fragmented data	0.92	0.58	4	1
10	Load data from text file	3.08	3.58	2	3
11	Adaptive thresholds and workload comparisons	7.75	1.50	7	3
	Category Sub-Total	15.32	10.08	20	12
	% difference -- (SS08 - Oracle)/SS08	34%		40%	
	DBA Workday Savings (Difference * Weighting Factor)	12%		14%	
Backup & Recovery Tasks					
12	Configure and perform full backup	1.42	0.00	3	0
13	Recover dropped table	4.75	1.17	6	2
14	Datafile recovery	2.58	1.92	3	2
15	Recover from erroneous transaction	1.92	1.92	5	2
16	Recover from multiple failures	6.67	3.08	8	2
	Category Sub-Total	17.34	8.09	25	10
	% difference -- (SS08 - Oracle)/SS08	53%		60%	
	DBA Workday Savings (Difference * Weighting Factor)	7%		8%	

Task	Description	Time (mins)		Steps	
		SS08	Oracle	SS08	Oracle
Performance Diagnostics & Tuning Tasks					
17	Diagnose performance problem	8.10	2.00	3	1
18	Tune resource-intensive SQL	10.00	0.42	7	2
19	Tune memory	0.00	0.00	0	0
	Category Sub-Total	18.10	2.42	10	3
	% difference -- (SS08 - Oracle)/SS08	87%		70%	
	DBA Workday Savings (Difference * Weighting Factor)	23%		18%	
	Overall Total	72.35	43.51	67	30
	% difference -- (SS08 - Oracle)/SS08	40%		55%	
	Total DBA Workday Savings	41%		43%	

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 (task 14 and task 15), 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.

Appendix VI - Detailed Task Descriptions[7]

Installation and Simple "Out-of-Box" Setup

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

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Install DB/ Software/Out-of-Box Experience							
Step 1	1	Windows Explorer: Run setup.exe from DVD	16 Min 5 Sec	Step 1	1	Windows Explorer-DVD-Setup.exe(Click>	15 Min 45 Sec
	2	11g Installation: Select Product: Accept Basic Installation Defaults / Set Database Password / Confirm Password / <NEXT>			2	System Configuration Check <select>	
	3	Installer: Product-Specific Prerequisite Checks: Verify Successful Status <Next>			3	<Installation Upgrade Advisor>	

[7] Table description: Each row in the table represents a task increment. The Key columns reference the number of steps and penalties for increments and context switches. (For space reasons, some penalties are described in the description field.) In the Task Description cells, a "/"represents a choice that needs to be made. For more information on this metric refer to the section of this study entitled The Methodology Defined.

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	4	Installer: Configuration Manager Registration – Accept Defaults <Next>			4	Upgrade Advisor Installation Wizard <Next>	
	5	Installer: Summary – Review Settings <Install>			5	Upgrade Advisor Setup \<Accept License>	
Step 2	4	Inventory: Review Installed Products <Close>		Incre-ment Penalty = 1	1	Upgrade Advisor Setup-Registration: Enter Name/ <Next>	
		Exit <Yes>			2	Upgrade Advisor Feature Selection: Accept Defaults/ <Next>	
					3	Upgrade Advisor <Install>	
					4	Upgrade Advisor <Finish>	
					5	<New SQL Server Installation>	
				Incre-ment Penalty = 1	1	Validate Configuration for Setup <OK>	
					2	Enter Product Key / <Next>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
					3	Read and Accept License / <Next>	
					4	Install Setup Support Files / <Install>	
					5	Setup Support Rules/Operation Completed / <Next>	
				Incre-ment Penalty = 1	1	Select Features <Next>	
					2	Instance Configuration/ Accept Defaults / <Next>	
					3	Service Accounts /Use Same Account	
					4	Select User or Group / Specify User / Password	
					5	Database Engine Configuration / Select <Windows Authentication Mode> / Set Password / Specify Administrator / <Add> / <Next>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
				Increment Penalty = 1	1	Analysis Services Configuration /Account Processing / <ADD CURRENT USER> <Next>	
					2	Analysis Services Configuration / Data Directories / Accept Defaults <Next>	
					3	Reporting Services Configuration / Install (But do not configure report server)/ <Next>	
					4	Installation Rules / Operation Completed /All rules passed / <Next>	
					5	Ready to install / Verify Features / <Install>	
				Increment Penalty = 1	1	Installation Process / Setup Process Complete /Review Status / Next	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
					2	Complete: Installation completed successfully / <OK>	

Metrics	Oracle	SQL Server
Time	16 Min 5 Sec	15 Min 45 Sec
Steps	2	1
Increment Penalty	0	5
Context Switch Penalty	0	0
Adjusted Steps (Complexity)	2	6

Task 2: Create additional database server/instance

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Create a new or second database instance							
Step 1	1	Start Menu / Configuration and Migration Tools / Database Configuration	6 Min 25 Sec	Step 1	1	Start Menu / SQL Installation Center	4 Min 40 Sec
	2	DB Config Assistant Welcome Screen / Next			2	SQL Installation Center / Install/ New Install/ Path to Install Disk	
	3	Create a Database / <Next>			3	Setup / Setup Support Rules / Rule Check	
	4	Select Database Template : General Purpose / Transaction Processing / <NEXT>			4	Setup Support Files / Install / <Next>	
	5	Specify Database Configuration Options: DB Naming/ DB Char Set/ DB Examples/ Keep Defaults.			5	Setup Support Rules Completed / Review: Failures = 0 (Failures would require identifying and fixing problems.)	
Inc. Pen. = 1	1	Database Identification: Enter Global DB Name / SID / <NEXT>		Inc. Pen. = 1	1	Installation Type / New SS08 <Next>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	2	Database Credentials: Use Same Admin Passwords for All Accounts / Enter PW / Confirm PW / <Next>			2	[Trial software requires Edition Choice]	
					3	Accept License	
					4	Select Features <Next>	
					5	Configure Instance Name	
				Inc. Pen. = 1	1	Review Features to Install <Next>	
					2	Disk Space Required / Review	
					3	Configure Accounts	
					4	Service Accounts / Use Same Account	
					5	Select User or Group / Specify User / Password	
				Inc. Pen. = 1	1	Authentication Mode / Accept Defaults / Add Administrator / Add Current User <Next>	
					2	Options / Messages to Microsoft / Accept Defaults / <Next>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
					3	Install Rules Complete / <Next>	
					4	Ready to Install / Review <Install>	
					5	Installation Complete <Next>	
				Inc. Pen. = 1	1	Successful Completion <Close>	

Metrics	Oracle	SQL Server
Time	6 Min 25 Sec	4 Min 40 Sec
Steps	1	1
Increment Penalty	0	4
Context Switch	0	0
Adjusted Steps (Complexity)	2	5

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

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Setup Proactive Monitoring							
Step 1	1	Enterprise Manager / Database Home Page / Metric and Policy Settings	25 Sec	Step 1	1	SQL Server Agent / New / Alert	1 Min 10 Sec
	2	Metric and Policy Settings / All Metrics/ Select Metric to Adjust / Enter Value / <OK>			2	New Alert / General/ Set: Name, Condition Alert, Object, Counter, Instance, Alert When Value, Counter	
					3	New Alert / Responses / <Notify Operator >, <email>	
					4	Option / Include alert in <email>	

Metrics	Oracle	SQL Server
Time	25 Sec	1 Min 10 Sec
Steps	1	1
Increment Penalty	0	0
Adjusted Steps (Complexity)	1	1

Metrics	Oracle	SQL Server
Time	22 Min 55 Sec	21 Min 35 Sec
Adjusted Steps (Complexity)	5	12

Day-to-Day Database Administration

Task 4: Create user, assign roles/privileges

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Create user with roles, privileges							
Step 1	1	1. Activate OEM / Server Tab	1 Min 15 Sec	Step 1	1	Management Studio / Security / New / Login	0 Min 55 Sec
	2	* Select Security/ Users link.			2	Login – New / General / Name / Search / Advanced / Find Now <OK> / <Select Name>	
	3	* Select <Create >			3	Server Roles / Select Roles / <OK>	
	4	* Enter User Name and Password / Accept other Defaults			4	User Roles / <Map>	
	5	* Enter User Name "Object has bee created Successfully" Page: <GO> / Select Username / Review Details			5	Securables <Accept Defaults>	
				Inc. Pen. = 1	1	Status <Accept Defaults>	
					2	Status <Accept Defaults> <OK>	

Metrics	Oracle	SQL Server
Time	1 Min 15 Sec	0 Min 55 Sec
Steps	1	1
Increment Penalty	0	1
Adjusted Steps (Complexity)	1	2

Task 5: Create Tablespace/Filegroup

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Create table-space							
Step 1	1	1. Activate EM Server Tab	55 Sec	Step 1	1	Management Studio /Database / Properties / Files / <Add>	40 Sec
	2	Select the Tablespaces <Create>button and Enter Tablespace Name			2	<Enter Name>, Filegroup / New Filegroup <Enter Name> <OK>	
	3	General Tab / Enter Tablespace Name / Accept Defaults <ADD>					
	4	Click the Add Button in the Datafile Section					
	5	Add Datafile Page / Enter FileName Accepting Defaults / <CONTINUE>					
Inc. Pen. = 1	1	Review Create Tablespace / Information: Accepting Defaults / <OK>					
	2	Confirmation / Tablespaces: Verify Creation					

Metrics	Oracle	SQL Server
Time	55 Sec	40 Sec
Steps	1	1
Increment Penalty	1	0
Adjusted Steps (Complexity)	2	1

Task 6: Add more space to the database

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Add more space to the DB							
Step 1	1	Activate EM Server Tab	55 Sec	Step 1	1	Management Studio / Select Database/ Database / Properties / Files / Select Data File / Initial Size / <Increase Size> / <OK>	27 Sec
	2	Select Datafiles Link.					
	3	Select an Existing Datafile / Select <CreateLike>					
	4	Datafiles / Create Datafile Page: Enter Filename / Accept Defaults <OK>					

Metrics	Oracle	SQL Server
Time	55 Sec	27 Sec
Steps	1	1
Increment Penalty	0	0
Adjusted Steps (Complexity)	1	1

Task 7: Create table

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Create table							
Step 1	1	Go to the EM Schema TAB->Tables link.	30 Sec	Step 1	1	Management Studio / Databases / Select and Expand Database / Right Click Tables / New / <Enter Name>	1 Min 5 Sec
	2	Instance / Tables: Accept Default Schema / Select <CREATE>			2	3 Columns With Data Types (number, varchar2 and date) /	
	3	Tables / Create Table Table Organization/ Standard/ <Continue>			3	Col1 / <Right Click> / <Set Primary Key> <Save>	
	4	Create Table: General/Name (Enter Table Name)/Accept Schema, Tablespace defaults/Three columns, Enter Name/Select Datatype/Size/< Ok>					
	5	Instance / Tables / Confirmation /					

Metrics	Oracle	SQL Server
Time	30 seconds	1 Min 5 Sec
Steps	1	1
Increment Penalty	0	0
Adjusted Steps (Complexity)	1	1

Task 8: Create index

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Create index							
Step 1	1	Go to the EM Schema TAB->Schema link. / Select <INDEXES>	50 Sec	Step 1	1	Management Studio / Databases / Select & Expand Databases / Select and Expand Table / Indexes / (Context Menu) <New Index>	27 Sec
	2	Select <Create>			2	New Index / <enter Name> / Columns <Add>	
	3	Enter Index Name / Accept Defaults / Enter Table Name / <Populate Columns>			3	Columns / Select <Column> / <OK>	
	4	Accept Default Sort Order / Specify Indexed Column by Entering Number in ORDERColumn / <OK>					

Metrics	Oracle	SQL Server
Time	50 Sec	27 Sec
Steps	1	1
Increment Penalty	0	0
Adjusted Steps (Complexity)	1	1

Task 9: Reclaim Lost Space due to Fragmented Data

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Determine if any tables in application tablespace need to be reorganized.							
Step 1	1	Advisor Central <Segment Advisor Recommendation >	35 Sec	Step 1	1	Management Studio / Server / Management / Maintenance Plans / (Context Menu) Maintenance Plan Wizard	55 Sec
	2	Segment Advisor Details / <Recommendation Details>			2	Maintenance Plan Wizard / Click <Next>	
	3	Segment Advisor: Recommendation Details for Tablespace Select <Segment Check Box> <Implement>			3	Select Plan Properties / Enter Plan Name / / Schedule: / < Change>	
	4	Shrink Segment: Options - Select Default: Compact Segments and Release Space <Implement>			4	Job Schedule Properties - Maintenance Properties / Frequency: Occurs: Daily / Daily Frequency: Occurs once at 10 pm / <OK>	
	5	Shrink Segment: Schedule -- Accept Defaults: Run Immediately			5	Select Maintenance Task / Shrink Table, Reorganize Index & Rebuild Index / <Next>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
				Inc. Pen. = 1	1	Select Maintenance Task Order / Click <Next>	
					2	Define Shrink Databases Task / Databases / Select <All databases> / <Next>	
					3	Define Reorganize Index Task / Select <All databases> / <Next>	
					4	Define Rebuild Index Task / Select <All databases> / <Next>	
					5	Select Report Options / e-Mail Report / <Next>	
				Inc. Pen. = 1	1	Select Report Options / e-Mail Report / <Next>	
					2	Complete the Wizard / Review Actions / <Finish>	
					3	Maintenance Plan Wizard Progress / <Close>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
				Step 2		Management Studio / Server / Management / Maintenance Plans / (Context Menu) Maintenance Plan Wizard / Select <View History>	
						Review results / <Close>	

Metrics	Oracle	SQL Server
Time	35 Sec	55 Sec
Steps	1	2
Increment Penalty	0	2
Adjusted Steps (Complexity)	1	4

Task 10: Load data from a text file

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
Task: Data Loading							
Step 1	1	Enterprise Manager <Data Movement>	3 Min 10 Sec	Step 1	1	Management Studio /Selected Database / (Context Menu) Tasks / Import Data	3 Min 5 Sec
	2	Data Movement <Load Data from User Files>			2	Import – Export Wizard / Select Data Source / Flat File	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	3	Load Data: Generate or Use Existing Control File: Accept Defaults <Continue>			3	Choose Data Source / filename / Browse Select / Set text delimiter	
	4	Load Data: Data Files: Select Data File <Next>			4	Import – Export Wizard / Choose Destination / Accept Defaults	
	5	Search and Select: File <Select>			5	Select Source Tables and Views / Destination Table / Select from List	
Step 2	1	Load Data: Data Files: Confirm Selected Import File <Next>	20 Sec	Inc. Pen. = 1	1	Review Data Type Mapping / Accept Defaults / <Next>	
	2	Load Data: Table and File Format: <SelectExistingTable> ; <Set Data File Format>;<Next>			2	Save and Run Package / Accept Defaults (Run Immediately) / <Next>	
	3	Load Data: File Format Attributes: <Set Delimiters> ; <Next> (Accepting Other Defaults)			3	Verify Choices / <Finish>	
	4	Load Data: Select <Direct Path> method <Next>			4	Monitor Process / <Close>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	5	Load Data: Options Accept Defaults <Next>					
Inc. Pen. = 1	1	Load Data: Schedule: Enter Job Name / Accept other Defaults (run immediately) / <Next>					
	2	Load Data: Review: <Control File Contents>					
	3	Load Data: Review: <Submit Job>					
	4	Job Activity: Confirmation - Page refreshes until completion					

Metrics	Oracle	SQL Server
Time	3 Min 35 Sec	3 Min 5 Sec
Steps	2	1
Increment Penalty	1	1
Adjusted Steps (Complexity)	3	2

Task 11: Configure Adaptive Thresholds (baselines) and Perform Workload Comparisons[8]

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Adaptive Thresholds and Workload Comparisons							
Step 1	1	1. OEM Home Page / <Baseline Metrics Threshold>	35 Sec	Step 1	1	Management Studio / Server / Data Collection /Configure Management Data Warehouse / Management Data Warehouse Wizard	1 Min 5 Sec
	2	* Threshold Configuration <Quick Configuration>			2	Create or Update Management Data Warehouse / <Next>	
	3	* Quick Configuration: Baseline Metric Thresholds / Select Profile (OLTP) <Continue>			3	Select Server / Database Name / <New> /	
	4	* Quick Configuration: Review Selected Profile Settings/Review/ <Finish>			4	New Database / Database Name /Enter Name /	
Step 2	1	Comparison to Baseline	30 Sec		5	New Database / Accept other defaults / <OK>	

[8] There are three different operations under this task. The operations have been given descriptive titles that reflect the differences between the products.

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
				Inc. Pen = 1	1	Configure Management Data Warehouse Wizard / <Next>	
					2	Map Logins and users / Users Mapped to this login/Select Users	
					3	Database Role Membership for "user" / Select roles / <Next>	
					4	Complete the Wizard / Review settings / <Finish>	
					5	Management Studio / Server / Data Collection /Configure Management Data Warehouse / Management Data Warehouse Wizard	
				Inc. Pen = 1	1	Select a configuration task / Setup Data Collection / <Next>	1 Min 5 Sec

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
					2	Configure Management Data Warehouse Storage / Select Server / Select Database / Enter Cache Directory (Blank uses TEMP Directory) / <Next>	
					3	Configure Management Data Warehouse Wizard Complete Wizard / Review / <Finish>	
				Step 2	1	View/Analyze Reports	5 min 30 Sec
					2	Select Data Collection / (Context Menu) Reports/ Management Data Warehouse / Query Statistics History	
					3	Analyze reports	

Metrics	Oracle	SQL Server
Time	1 Min 30 Sec	7 Min 45 Sec
Steps	2	2
Increment Penalty	0	2
Context Switch Penalty	0	0
Adjusted Steps (Complexity)	2	4

Metrics	Oracle	SQL Server
Time	10 Min 5 Sec	15Min 19 Sec
Adjusted Steps (Complexity)	12	20

Backup & Recovery

Task 12: Configure and Perform Full Database Online Backup

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Configuring for online backup							
Step 1		N/A (Oracle 11g Recommended Backups are Configured and Scheduled Automatically Upon Database creation.)	0 Min 0 Sec	Step 1	1	Management Studio / Server / Management / Maintenance Plans / (Context Menu) Maintenance Plan Wizard	1 Min 25 Sec
					2	Select Plan Properties / Enter Plan Name / Schedule / <Change>	
					3	Job Schedule Properties - Maintenance Properties / Frequency: Occurs: Daily / Daily Frequency: Occurs once at 10 pm / <OK>	
					4	Select Maintenance Task / Backup Database (Full) / <Next>	
					5	Task Order / Accept Defaults / <Next>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
				Inc. Pen. =1	1	Define Backup Task / <Select Database> <Select Database or Databases to backup>/Accept other Defaults/<Next>	
					2	Report Options / e-Mail Report <Next>	
					3	Wizard / Review Actions / <Finish>	
					4	Create Maintenance Plan <Close>	
				Step 2	1	Management Studio / Mgmt./Maintenance Plan <select plan to execute>	
					2	Popup Menu/Execute	
					3	Verify Success / <CLOSE>	

Metrics	Oracle	SQL Server
Time	**0**	**1 Min 25 Sec**
Steps	**0**	**2**
Increment Penalty		**1**
Adjusted Steps (Complexity)	**0**	**3**

Task 13: Recover Dropped Table

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task 13: Recover dropped table							
Step 1	1	Enterprise Manger <Availability>	1 Min 10 Sec	Step 1		Restore to a temporary copy of the Database	4 Min 45 Sec
	2	Availability: <Perform Recovery>			1	Management Studio / Databases/ (Context Menu) Restore Databases	
	3	Perform Recovery: RecoveryScope: <Tables> / <Flashback Dropped Tables> / <Recover>			2	Restore Database — General: Select Restore Destination/Type new database name/Select Source (dropped table has to exist in selected backup)	
	4	Perform Object Level Recovery: Dropped Objects Selection.Results <EnterSchemaName RecycleBin> <GO> /Accept Default Dropped Table In RecycleBin <NEXT>			3	Restore Database —Options: Confirm or select options/<OK>	
	5	Perform Object Level Recovery: Rename: Accept Default / <NEXT>		Step 2		Restore Table	

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Inc. Pen. = 1	1	Perform Object Level Recovery: Review: Accept Default <SUBMIT>		Cont. Switch Pen. = 2	1	Start Menu/SQL Server/Import & Export Data	
	2	Confirmation (Tables Have Been Flashed Back)			2	Wizard Welcome Screen/Next	
					3	Choose Source/ Select/<Next>	
					4	Chose a Destination/ Select/<Next>\	
					5	Select Table Copy or Query/Copy Data from One or More Tables or Views/<Next>	
				Inc. Pen. = 1	1	Select Source Table & Views / Select / <Next>	
					2	Save & Run Package / Accept Defaults (Run Immediately)/ <Next>	
					3	Verify Choices /<Finish>	
					4	Monitor Process <Close>	
				Step 3		Delete temporary copy of the Database	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
					1	Management Studio / Databases/ Temporary Database (Context Menu) Delete	
					2	Delete Object / Accept Defaults / <Ok>	

Metrics	Oracle	SQL Server
Time	1 Min 10 Sec	4 Min 45 Sec
Steps	1	3
Increment Penalty	1	1
Context Switch Penalty		2
Adjusted Steps (Complexity)	2	6

Task 14: Recover Data File

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
Task 14: Recovery: Application Data File of Size 1.2 Gigs is Lost.							
Step 1	1	Database Instance: Cannot Connect - <Perform Recovery>	1 Min 55 Sec	Step 1	1	Management Studio / System Databases/ Master (Context Menu) Backup	2 Min 35 Sec
	2	Perform Recovery: Credentials – Enter UserName / Enter Password / <Continue>			2	Backup Databases / Select Database / Backup Type: Transaction Log / Set Backup Destination: <Enter Filename><OK>	

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
	3	Database Login – Enter User Name / Enter Password / Connect As: <SYSDBA> / <LOGIN>			3	Options / Transaction Log / Backup Tail of Log and Leave DB in Restore State / <OK>	
	4	Perform Recovery: Oracle Advised Recovery <Advise and Recover>			4	Successful Backup / <OK>	
	5	View and Manage Failures: <Review Failure> (Missing DataFile) / <ADVISE>		Step 2	1	Select "Corrupted" Database / Restore / Database / From Device / Select File/	
Inc. Pen. = 1	1	Recovery Advice: <Recover>			2	Specify Backup / Add / Select File / <OK>	
	2	Failures that Will Be Resolved: Review <Submit>			3	"Corrupted" Database /(Context Menu) Tasks / Restore /Database/With Replace/With No Recovery	
	3	Processing Job Recovery Notification			4	Successful Restore / <OK>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	4	Recovery Results		Step 3	1	"Corrupted" Database / (Context Menu) / Tasks / Restore / Transaction Log / From File / Select File / <OK>	
					2	Options / Accept Default (restore with recovery) / <OK>	

Metrics	Oracle	SQL Server
Time	1 Min 55 Sec	2 Min 35 Sec
Steps	1	3
Increment Penalty	1	0
Context Penalty	0	0
Adjusted Steps (Complexity)	2	3

Note: The 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.

Task 15: Recover from erroneous transaction

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
Task: Scenario: Due to Human Error, a Transaction was Submitted by Mistake							
Step 1	1	Enterprise Manager / Schema Tab – Select <Tables>	1 Min 55 Sec	Step 1	1	Management Studio / Databases / Select Database / (Context Menu) Tasks / Restore / Select Backup Set	1 Min 55 Sec

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	2	Tables/Search – Enter Schema Name: Enter Name / <Go> / Select <TableName>			2	Options /Overwrite Existing / Allow Transaction Log Restore (No Recovery) <OK>	
	3	View Table: Actions <Flashback Table>			3	Successful Restore <OK>	
	4	Perform Object Level Recovery: Point-in-Time - <Select Flashback to Timestamp (accepting default date)> Set <Hour><Minute> <Next>			4	Management Studio / Databases / Select Database / (Context Menu) Tasks /Restore / To specific date and time / Set Restore Point Time / <OK>	
	5	Perform Object Level Recovery: Flashback Tables – Accept Defaults <Next>			5	Select Device /File / Select File / Select	
Inc. Pen. = 1	1	Perform Object Level Recovery: Flashback Tables – ERROR – Table Row Movement Not Enabled <ENABLE TABLE ROW MOVEMENT>l		Inc. Pen. = 1	1	Restore Successful / <OK>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	2	Confirmation – Table Modified Successfully <NEXT>		Step 2	1	Manually Redo Lost Transactions Ctx. Pen. = 2 Steps[9]	
	3	Perform Object Level Recovery: Review <SUBMIT>					
	4	Confirmation: Tables Have Been Flashed Back					

Metrics	Oracle	SQL Server
Time	1 Min 55 Sec	1 Min 55 Sec
Steps	1	2
Increment Penalty	1	1
Context Switch Penalty		2
Adjusted Steps (Complexity)	2	5

[9] We've accessed context switch penalties on this task. The penalty of 2 steps is being assesses for the use of manual recovery process to recover from the erroneous transactions.

Task 16: Recover from Multiple Failures

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task 16: Recover from multiple failures							
Step 1	1	Enterprise Manger <Availability>	3 Min 5 Sec	Step 1		Restore to a temporary copy of the Database	6 Min 40 Sec
	2	Availability: <Perform Recovery>			1	Management Studio / Databases/ (Context Menu) Restore Databases	
	3	Perform Recovery: RecoveryScope: <Tables> / <Flashback Dropped Tables> / <Recover>			2	Restore Database — General: Select Restore Destination/Type new database name/Select Source (dropped table has to exist in selected backup)	
	4	Perform Object Level Recovery: Dropped Objects Selection.Results – <EnterSchemaName RecycleBin> <GO> /Accept Default Dropped Table In RecycleBin <NEXT>			3	Restore Database — Options: Confirm or select options/<OK>	
	5	Perform Object Level Recovery: Rename: Accept Default / <NEXT>		Step 2		Restore Table	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	1	Perform Object Level Recovery: Review: Accept Default <SUBMIT>			1	Start Menu/SQL Server/Import & Export Data	
	2	Confirmation (Tables Have Been Flashed Back)			2	Wizard Welcome Screen/Next	
Step 2	1	Enterprise Manager / Schema Tab – Select <Tables>	1 Min 55 Sec		3	Choose Source/Select/<Next>	
	2	Tables/Search – Enter Schema Name: Enter Name / <Go> / Select <TableName>			4	Chose a Destination/Select/<Next>\	
	3	View Table: Actions <Flashback Table>			5	Select Table Copy or Query/Copy Data from One or More Tables or Views/<Next>	
	4	Perform Object Level Recovery: Point-in-Time - <Select Flashback to Timestamp (accepting default date)> Set <Hour><Minute> <Next>		Inc. Pen. = 1	1	Select Source Table & Views / Select / <Next>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	5	Perform Object Level Recovery: Flashback Tables – Accept Defaults <Next>			2	Save & Run Package / Accept Defaults (Run Immediately)/< Next>	
	1	Perform Object Level Recovery: Flashback Tables – ERROR – Table Row Movement Not Enabled <ENABLE TABLE ROW MOVEMENT>l			3	Verify Choices7 /<Finish>	
	2	Confirmation – Table Modified Successfully <NEXT>			4	Monitor Process <Close>	
	3	Perform Object Level Recovery: Review <SUBMIT>		Step 3	1	Management Studio / Databases / Select Database / (Context Menu) Tasks / Restore / Select Backup Set	
	4	Confirmation: Tables Have Been Flashed Back			2	Options /Overwrite Existing / Allow Transaction Log Restore (No Recovery) <OK>	
					3	Successful Restore <OK>	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
					4	Management Studio / Databases / Select Database / (Context Menu) Tasks /Restore / To specific date and time / Set Restore Point Time / <OK>	
					5	Select Device /File / Select File / Select	
				Inc. Pen. = 1	1	Restore Successful / <OK>	
				Step 4	1	Manually Redo Lost Transactions Context Switch Penalty = 2 steps[10]	

Metrics	Oracle	SQL Server
Time	3Min 5 Sec	6 Min 40 Sec
Steps	2	4
Increment Penalty	2	2
Context Switch Penalty		2
Adjusted Steps (Complexity)	4	8

Metrics	Oracle	SQL Server
Time	8 Min 5 Sec	17 Min 20 Sec
Adjusted Steps (Complexity)	10	25

[10] We've accessed context switch penalties on this task. The penalty of 2 steps is being assessed for the use of manual recovery process to recover from the erroneous transactions, one of the multiple failures encountered.

Performance Diagnostics & Tuning Tasks

Task 17: Diagnose Performance Problem

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Configure System to Identify Top Resource-Consuming SQL							
Step 1	1	Go to the EM home page and review the last ADDM report	2 Min	Step 1	1	Management Studio /Server/Data Collection /(Context Menu) Reports/ Query Statistics	8 Min 6 secs
	2	Drill down to the appropriate finding			2	Query Statistics History /Select Timeframe	
					3	Review Report (Top Queries by CPU) / Select problem Query	
				Step 2	1	Review and Analyze other MDW Reports (up to 15) One Increment Penalty for viewing multiple reports	

Metrics	Oracle	SQL Server
Time	2 Min	8 Min 6 Sec
Steps	1	2
Increment Penalty	0	1
Context Switch Penalty		
Adjusted Steps (Complexity)	1	3

Task 18: Fix Performance Problem - SQL Tuning

Oracle				SQL Server			
Key	**Inc.**	**Description**	**Time**	**Key**	**Inc.**	**Description**	**Time**
Task: Identify Resource-Intensive SQL							
Step 1	1	1. Enterprise Manager / <Server>	25 Sec	Step 1	1	Management Studio / Server /(Context Menu) Activity Monitor / Recent Expensive Queries	10 Min 0 Sec[1]
	2	Server / <Automated Maintenance Tasks>			2	Select Query / Edit Query Text / Save Query As (Save to a convenient location)	
	3	Automated Maintenance Tasks / <Automatic SQL Tuning>			3	Management Studio / Tools Menu / Database Engine Tuning Advisor	
	4	Automated SQL Tuning <View Report>			4	Provide Session Name (Use Default) / Select Workload File (previously saved query)/ Select Database for Workload Analysis / Select Database(s) or tables to tune/	

Oracle				SQL Server			
Key	Inc.	Description	Time	Key	Inc.	Description	Time
	5	SQL Tuning Result Details / <Implement All>			5	Tuning Options / Advanced Options / Generate online recommendations where possible /<OK>	
Inc. Pen = 1	1	Confirmation / <Implement New Profiles with Forced Matching>/<YES>		Inc. Pen. = 1	1	Tuning Options /Accept Defaults /	
	2	Automated SQL Tuning / View Results			2	Select Session (from list)/ (Context Menu) <Start Analysis>	
					3	Review Recommendations and Reports	
				Step 2 Ctx. Pen = 5 Steps	1	Implement Recommendations	

Metrics	Oracle	SQL Server
Time	35 Sec	10Min
Steps	1	1
Increment Penalty	1	1
Context Switch Penalty		5[11]
Adjusted Steps (Complexity)	2	7

[11] We've accessed context switch penalties on this task. One penalty is for the requirement to use two tools: SQL Server Management Studio and SQL Server Engine Tuning Advisor. The second through fifth penalty is for the activities within the requirement to open and manually edit the query, or otherwise take the manual tuning activities recommended. Edison believes this number of steps would be the minimum required. For complex problems, the number of steps could be much higher.

Task 19: Tune Memory

Oracle		SQL Server	
Step	**Time**	**Step**	**Time**
Task: Tuning memory			
Automatic Memory Tuning is Enabled as a Maintenance Setting Increment During System Setup	0 Min 0 Sec	Self-Tuning Memory is Automatically Enabled at installation	0 Min 0 Sec

Metrics	Oracle	SQL Server
Time	0	0
Steps	0	0
Increment Penalty		
Context Switch Penalty		
Adjusted Steps (Complexity)	0	0

Metrics	Oracle	SQL Server
Time	2 Min 25 Sec	18 Min 6 Sec
Adjusted Steps (Complexity)	3	10

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 DBAs 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 DBAs 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 DBAs 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 DBAs to manage user security in the database.

Create schema: The workload task that allows DBAs to create a new schema.

Create table: The workload task that allows DBAs 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 DBAs 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 DBAs 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 DBAs 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 DBAs 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 DBAs to optimize a poorly running query/procedure so that it executes in an efficient manner.