March 2013

MANAGEMENT BRIEF

Database Economics

Comparing Value Propositions for DB2 and Oracle Database Systems

International Technology Group

609 Pacific Avenue, Suite 102
Santa Cruz, California 95060-4406
Telephone: + 831-427-9260
Email: Contact@ITGforInfo.com
Website: ITGforInfo.com

TABLE OF CONTENTS

List of Figures

EXECUTIVE SUMMARY

Challenges

The IT world remains subject to contradictory pressures. Organizations worldwide report, at best, incremental growth in IT expenditure, while cost controls continue to be rigorously imposed.

At the same time, as markets begin to recover from the impact of recession, businesses see the need to invest in new IT solutions that can create competitive advantage. New technology challenges in analytics, mobility, social media, cloud computing, big data and other areas must be addressed.

Databases are at the heart of these trends. In most organizations, they house the vast majority of business data supporting transactional, customer-facing and – increasingly – analytical applications. They must handle high double-digit annual growth in data volumes, growing structural complexity, expanding user populations and workloads, and the integration of new types of unstructured data.

There are compelling reasons to pay attention to database investments. Databases determine, in no small measure, how effectively organizations use information for business advantage. They also, to a much greater extent than is generally realized, influence overall IT cost structures.

This report deals with these issues. Specifically, it compares the costs of employing Oracle Database 11g and IBM DB2 10 for Linux, UNIX and Windows in large organizations. Comparisons are for database installations in telecommunications, financial services, health care and retail companies.

The assumptions employed in constructing these profiles were based on information supplied by 64 organizations worldwide that operate DB2 as well as Oracle databases, or have migrated from one to the other. The report also includes detailed functional comparisons of DB2 10 and Oracle Database 11g that highlight relationships between technology and costs for both databases.

Database Costs

It will come as no surprise that DB2 costs less than Oracle. IBM has consistently undercut the latter on database pricing for more than a decade. Equally, there is no serious dispute that DB2 is more efficient in its use of system resources and database administrator (DBA) time. The extent of these differences is, however, not always appreciated.

Disparities between the latest DB2 10 Advanced Enterprise, introduced in April 2012, and Oracle Database 11g Enterprise Edition are particularly striking. In the installation comparisons presented in this report, three-year license and support costs for use of DB2 10 Advanced Enterprise average 39 percent less than for equivalent Oracle software stacks. Figure 1 illustrates these results.

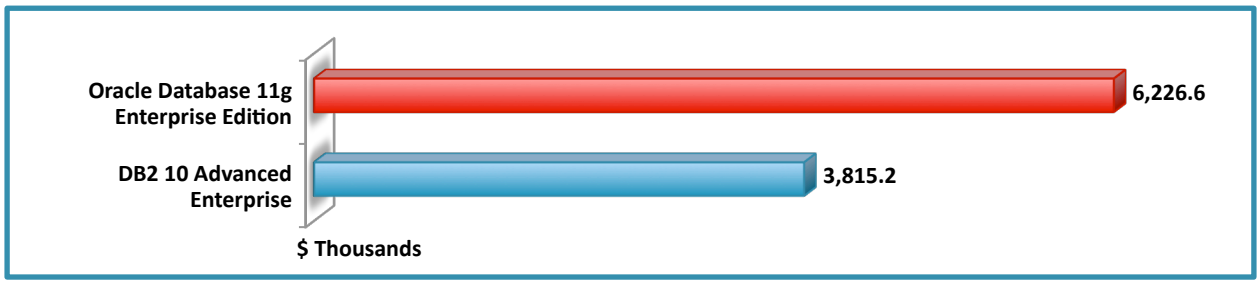

Figure 1: ***Three-year License and Support Costs for Use of DB2 10 Advanced Enterprise and Oracle Database 11g Enterprise Edition – Averages for All Installations***

Database software costs are calculated using "street" prices; i.e., discounted prices actually paid by users.

Disparities in full time equivalent (FTE) DBA staffing levels and corresponding personnel costs are even larger. For the same installations, three-year personnel costs for use of DB2 Advanced Enterprise average 59 percent less. Figure 2 illustrates these results.

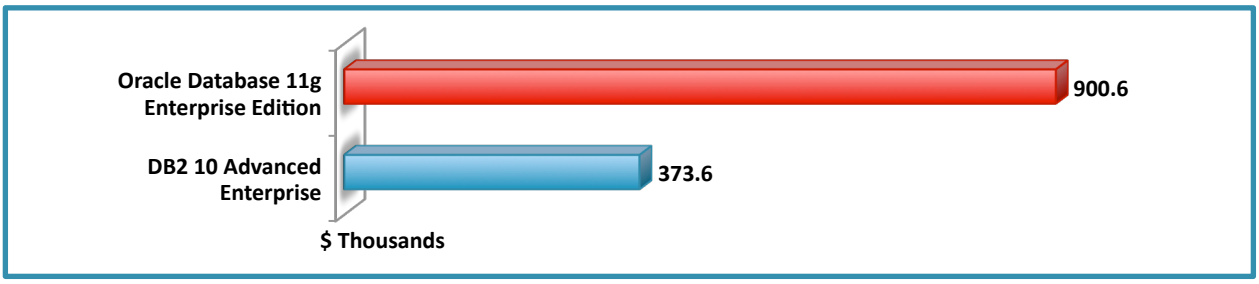

Figure 2: ***Three-year DBA Personnel Costs for Use of DB2 10 Advanced Enterprise and Oracle Database 11g Enterprise Edition – Averages for All Installations***

Although applications, task structures and job descriptions vary between organizations, these results are generally consistent with industry experience. Organizations that have conducted migrations from Oracle to DB2 routinely report two to three times reductions in FTE staffing levels.

Combined database three-year costs for database software and DBA personnel average 41 less for use of DB2 10 Advanced Enterprise.

Infrastructure Costs

Overall cost structures for database deployments extend well beyond software and DBA personnel. They include disk systems, database servers, tape systems and media, and personnel and data center costs for facilities, power and cooling. These costs are affected by differences in database technologies. This is particularly the case in two areas:

1. ***Performance***. More effective performance optimization and lower system overhead enable use of smaller database servers to execute the same workloads. Savings may be realized in hardware and maintenance costs; in license and support costs for databases and systems software priced on a per processor basis; and in server administration and data center costs.

   These effects are reinforced if effective workload management mechanisms are in place. The ability to use processor resources more efficiently may result in further savings. This will particularly be the case in environments characterized by diverse, fluctuating workloads.

2. ***Compression***. Higher levels of data compression enable savings in disk array hardware and maintenance, and in license and support fees for point-in-time copying, remote real-time replication, storage resource management and other software priced on a per terabyte basis. Further economies may be generated in backup systems and media, and in storage administrator and data center costs for these resources.

   Compression may also reduce server memory and I/O, and wide area network costs – bandwidth savings may be particularly significant for organizations that replicate data to failover sites on a real-time basis.

Marked disparities between DB2 and Oracle capabilities in these areas have been apparent since the mid-2000s. With the introduction of DB2 10, however, they have widened.

Early users report performance improvements of between 30 percent and more than two times over DB2 9.7, depending on applications and workloads. Adaptive Compression has enabled space savings of more than seven times overall, and more than 10 times for some individual tables.

In the same set of installations employed for database software and DBA personnel cost comparisons, three-year costs for disk, tape and database server infrastructures are also significantly lower for use of DB2 10 Advanced Enterprise than Oracle Database 11g Enterprise Edition. Figure 3 illustrates disparities.

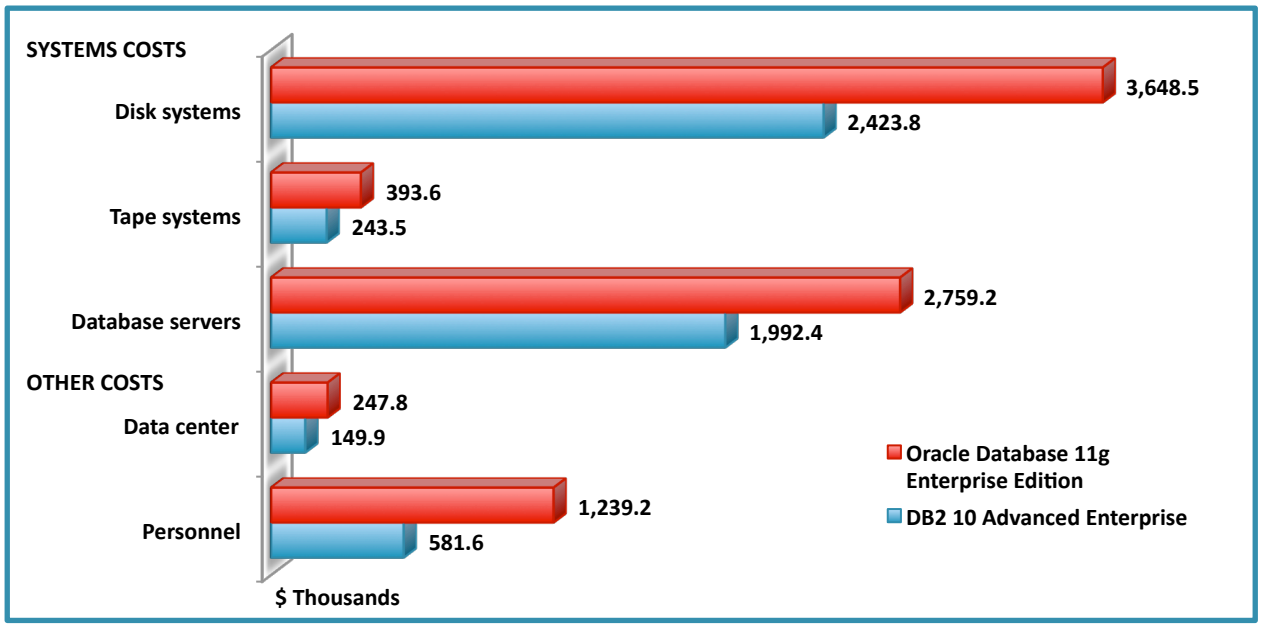

Figure 3: ***Average Three-year Infrastructure Costs for Use of DB2 10 Advanced Enterprise and Oracle Database 11g Enterprise Edition – Averages for All Installations***

In the chart above, systems costs include hardware and maintenance, systems software licenses and support, and (for tape systems) media and handling. Personnel costs are for storage and server administration, while data center costs include facilities occupancy, power and cooling.

Calculations are based on use of IBM DS8800, XIV and Storwize V7000 systems, IBM TS3500 tape systems and IBM POWER7-based database servers. Hardware, maintenance and systems software costs are again calculated using street prices.

Overall three-year costs, including database as well as infrastructure costs, averaged 37 percent less for use of DB2 10 Advanced Enterprise. Figure 4 illustrates these results.

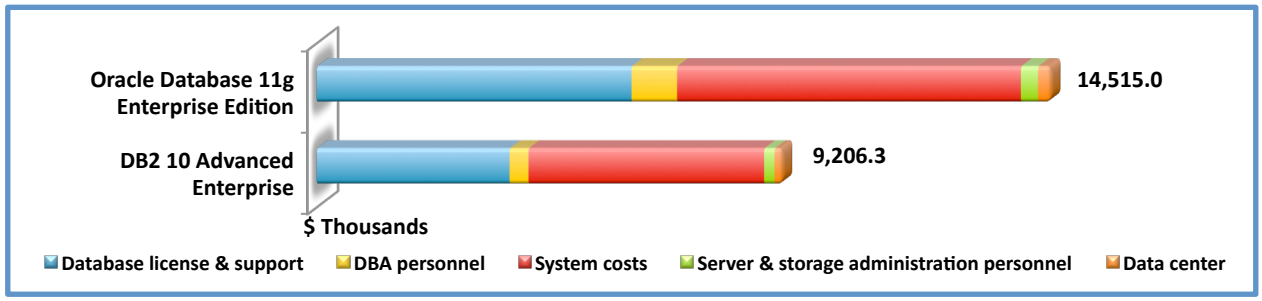

Figure 4: ***Overall Three-year Costs for Use of DB2 10 Advanced Enterprise and Oracle Database 11g Enterprise Edition – Averages for All Installations***

The Detailed Data section of this report contains further information on installations, hardware and software configurations, FTE staffing levels and methodologies and assumptions employed in calculations, along with more granular cost breakdowns.

Technology Differences

Comparing Oracle Database 11g and DB2 10

Lower costs are not all that distinguishes DB2 10 from Oracle Database 11g. There are also major differences in architecture and technology.

DB2 for Linux, UNIX and Windows is the newest of the industry's major relational databases. First introduced in 1996, it maintained compatibility and shared some features with the older IBM mainframe version of DB2, but the core design was significantly different. Successive versions equaled and then exceeded Oracle functionality. The latest DB2 10, introduced in April 2012, widened the gap further.

The Oracle database design originated in the 1970s – the first commercially available version was introduced in 1978 – and its popularity created a legacy installed base and embedded skill sets with which the company has been obliged to maintain compatibility ever since.

DB2 10 incorporates a broader range of new capabilities than any major database release, by any vendor within the last decade. Figure 5 shows examples.

FEATURES	DESCRIPTION
Performance Optimization	New features include multi-core parallelism (automatically optimizes thread to core ratios on SMP servers); index prefetching; & new optimization features for hash joins, queries across star schemas, aggregation & other processes. New zigzag join method reduces throughput time & I/O loading for complex multidimensional queries.
Adaptive Compression	New algorithms integrate table & index compression in DB2 9.7 with page-level compression. Users report an overall 30 to 40 percent increase in compression levels.
Time Travel Query	New temporal database capability distinguishes between system time (when an event is logged) & business time (an alternative date &/or time associated with the event) in maintaining & querying records. Obviates need for custom-developed applications to analyze multiple timelines. DB2 10 is the first major database to comply with temporal features of ANSI/ISO SQL:2011 standard.
Multi Temperature Data Management	New feature enables automated storage tiering for higher performance & lower overall disk costs. Obviates need for controller-based tiering for most workloads. Enables tight integration with DB2 workload management facilities.
Continuous Data Ingest	Employs new IBM parallel loading technology for extremely fast, low-overhead data transfers. Enables "real-time" data warehousing applications. Offers higher-performance alternative to conventional batch & "trickle feed" techniques.
Big Data Technologies	New ***s***upport for NoSQL standards including SPARQL & Resource Definition Framework (RDF). New Graph Store feature enables use of tripled graph stores. Overall feature set enables integration of Hadoop & MapReduce data into DB2 10 databases.
Security	New IBM technology, Row & Column-level Access Control (RCAC), enables integration of row & column-level mechanisms for Fine-Grained Access Control (FGAC).

Figure 5: *New DB2 10 Capabilities: Examples*

New DB2 10 features build upon performance optimization features introduced in earlier versions, and on established DB2 strengths in such areas as simplification and automation, workload management, clustering and SQL compatibility with Oracle Database.

Simplification and automation features extend across the entire DB2 environment, and exploit IBM autonomic technologies. Autonomic computing, meaning the application of advanced artificial intelligence technologies to IT administration and optimization tasks, has been a major IBM development focus for more than a decade. The company is the recognized industry leader in this area.

The *DB2 Workload Manager (WLM)* aligns fine-grained management of system resources to meet service level agreement (SLA) targets. Oracle Database 11g does not address service level management directly. Customized solutions are required.

DB2 pureScale feature draws upon industry-leading technologies such as Parallel Sysplex Data Sharing, which has been employed for close to 20 years to support the world's largest mainframe-based business-critical systems; and the IBM General Parallel File System (GPFS), generally recognized as the industry's highest-performing, most scalable cluster file system.

IBM has invested heavily in *SQL compatibility*. Organizations that have migrated Oracle Database applications to DB2 10 have typically found that around 98 percent of code remains unchanged, while few or no changes to development skills are required. Higher levels of DB2 10 simplification and automation also mean that DBA retraining is a comparatively simple process.

These compatibility functions are built into the core engine rather than implemented in the form of a software overlay. Users generally enjoy the same levels of performance as their native DB2 10 counterparts.

Enhancement Strategies

The widening technology gap between IBM and Oracle is not only due to differences in core designs. The two companies have also adopted significantly different enhancement strategies.

Oracle tends to add new functionality in add-ons and overlays. This approach has increased both the complexity and inefficiency of overall code structures, resulting in higher processor overhead and lower DBA productivity. In comparison, IBM has tended to redesign code.

A further difference is that performance optimization as well as simplification and automation of DBA tasks have been significantly higher priorities for IBM than for Oracle.

In a detailed comparative analysis presented later in this report, these differences are apparent "across the board" – in memory and storage management, database administration, tuning, compression, temporal capability, data loading, workload management, HA clustering, security and other areas. In most cases, DB2 advantages are greatest when databases undergo frequent changes in size, schemas, underlying data structures and workloads.

The impact of these differences on overall system efficiency and DBA productivity is cumulative. Many small differences add up to striking disparities in overall capability.

Key Database Management Solutions

Database capabilities are not only determined by software structures. Key tools employed for database development and administration, performance tuning, configuration management and other tasks, along with the skills of the staff that use them have a significant impact.

In this area too, DB2 10 enjoys a technological advantage. IBM offers a number of tools, acquired through the company's acquisition of Princeton Softech and developed internally, which are generally regarded as industry leading. Some are included in the DB2 10 Advanced Enterprise package, which is summarized in figure 6. Oracle equivalents are separately charged.

In the DB2 Advanced Enterprise package, Q Replication includes support for three servers, and InfoSphere Data Architect supports up to 10 user licenses without additional charge.

IBM DB2 10 ADVANCED ENTERPRISE	ORACLE DATABASE 11g EQUIVALENT	ORACLE U.S. LIST
DATABASE FEATURES		
DB2 10 Advanced Enterprise	Oracle Database 11g Enterprise Edition	$47,500 per processor
Table Partitioning	Partitioning	$11,500 per processor
High Availability Disaster Recovery	Active Data Guard	$10,000 per processor
DB2 pureScale Feature (separately priced)	Real Application Clusters (RAC)	$23,000 per processor
Adaptive Compression	Advanced Compression	$11,500 per processor
Time Travel Query	No direct equivalent	–
Multi Temperature Data Management	No direct equivalent	–
Continuous Data Ingest	Data Pump (included in database)	–
Multi Temperature Data Management	No direct equivalent	–
Q Replication	GoldenGate	$17,500 per processor
Label-based Access Control	Label Security	$11,500 per processor
KEY TOOLS		
InfoSphere Optim Performance Manager Extended Edition (OPM EE)	Diagnostics Pack	$5,000 per processor
InfoSphere Optim Configuration Manager	Database Lifecycle Management Pack	$12,000 per processor
InfoSphere Optim Query Workload Tuner	Tuning Pack	$5,000 per processor
InfoSphere Optim pureQuery Runtime	TopLink	$5,800 per processor
InfoSphere Data Architect	SQL Developer Data Modeler	–
Data Studio	Internet Developer Suite	$5,800 per named user

Figure 6: ***DB2 Advanced Enterprise Components and Oracle Database 11g Equivalents***

In comparing DB2 10 and Oracle Database 11g tooling, certain themes recur. IBM tools are comparatively new, while Oracle offerings such as Diagnostics and Tuning Packs date from the 1990s. Although they have been enhanced since that time, Oracle has tended, again, to implement new features in a manner that increases complexity. DBA productivity has not been a major design focus.

In comparison, IBM tools have typically been re-engineered to add new capabilities, and simplification and automation have been higher priorities. Tools have also been more closely integrated with each other and with underlying DB2 features and functions than is the case for their Oracle counterparts. The effects are again cumulative.

Conclusions

Against Oracle Database, DB2 wins in most, if not all, areas of feature-by-feature comparison. But there are also broader differentiators. DB2 is, by wide margins, better geared to handle the way in which the database world will evolve in the future.

DB2 is more effective in handling complex mixed workloads supporting large, diverse user populations. Its core design is less rooted in the transactional world than that of Oracle. Its structure more easily allows introduction of new technologies.

DB2 is, moreover, better designed to deliver higher levels of efficiency and administrator productivity across the entire infrastructure of storage systems, servers and software that support databases. As the findings of this report indicate, there is the potential for substantial cost savings. Equally, if not more important, gains may also be realized in availability, response time and other service level parameters.

Such capabilities will play an increasingly important role in determining how well organizations use information for business advantage. The choice between DB2 and Oracle, from a business as well as a technology perspective, is between leadership and legacy commitment.

TECHNOLOGY VIEW

Comparing Databases

In comparing DB2 10 and Oracle Database 11g, it is important to distinguish between what a particular feature does, and how efficiently it functions in practice. Inefficient code structures may materially increase processor, memory and disk overhead, retard throughput, increase risks of bottlenecks and outages, and cause other negative effects.

Equally, inefficient task structures – e.g., it requires 10 comparatively complex steps to perform a given task with Oracle Database 11g, and 5 simpler steps to perform the same task with DB2 10 – may materially reduce DBA productivity and undermine performance. If tuning takes too much time and effort, organizations will tend to run poorly tuned systems.

These distinctions are critical in comparing DB2 10 and Oracle Database 11g. In this section, functionality as well as efficiency are contrasted for nine key areas of capability – performance optimization, data compression, temporal data management, storage tiering, real-time data input, support for big data and XML technologies, workload management and clustering solutions.

The extent and significance of SQL compatibility for Oracle features in DB2 10 are also addressed, and IBM and Optim tooling are then compared.

Performance Optimization

Servers configured with earlier versions of DB2 have consistently outpaced Oracle databases in performance since the mid-2000s, according to a wide range of industry benchmarks. User experiences confirm this picture.

Both databases implement such capabilities as partitioning and clustering. However, DB2 10 implementations are generally recognized as more effective.

DB2 Multi-Dimensional Clustering (MDC), for example, enables flexible clustering across multiple dimensions. Query performance has typically been accelerated by around three times, and improvements of ten times or more have been reported. Oracle Database 11g implements a table-clustering feature with generally comparable capabilities, but does not match MDC performance and functionality.

DB2 10 has also taken a lead in new areas. These include multi-core parallelism (thread to core ratios on symmetric multiprocessing are adjusted automatically to optimize performance); and enhanced algorithms for index prefetching, hash joins, queries across star schemas, aggregation and other processes. A new zigzag join method can deliver sub-second response time for complex multi-dimensional queries.

Data Compression

Compression Effects

Data compression delivers a variety of potential benefits. These include savings in disk and processor capacity, and in memory and I/O costs. Server loading may be reduced, enabling use of smaller machines and lowering license and support costs for systems software priced on a per processor basis.

Not only storage hardware costs, but also outlays for software employed for point-in-time copying, remote real-time replication, storage resource management and other functions may also be reduced. On high-end disk arrays, these are typically priced on a per terabyte basis. In large installations, substantial savings may be realized in license and support fees.

Further economies may be generated in systems, media and administrator time for backups, replication, and other data management and movement processes. Costs of wide area network bandwidth may also be reduced, particularly for remote real-time replication to failover sites.

There may, however, also be drawbacks. Processor overhead tends to increase, and effective use of compression may require extensive manual intervention to collect statistics, identify compression opportunities, measure performance and set and adjust parameters.

These drawbacks may be mitigated if code structures are efficient, and if simplified and automated processes are available to administrators. This is more the case for DB2 10 than for Oracle Database 11g.

Compression Technologies

Users of earlier DB2 versions routinely experienced higher levels of compression than their Oracle counterparts. Although high levels of compression may in principle be realized with Oracle Database 11g Advanced Compression, greater performance degradation tends to occur.

Oracle Database 11g users typically realize 20 to 40 percent compression for production systems. In comparison, the DB2 9.7 norm was in the 60 to 80 percent range, even for exacting production workloads.

DB2 10 incorporates a new set of compression algorithms, Adaptive Compression, which integrates the table and index compression capabilities of DB2 9.7 with new page-level compression. Early user experiences with DB2 10 indicate that Adaptive Compression results in a further space savings compared to DB2 9.7. Results vary according to database and workload characteristics.

Multiple factors contribute to user experiences. These include:

- ***Design***. The effects of the combination of technologies in DB2 10 are cumulative. Page compression is not simply an overlay on table and index compression. The overall set of algorithms has been revised to maximize efficiency. Oracle Advanced Compression is built upon block-level techniques employed in legacy databases, and tends to be most effective in compressing indexes.

- ***Scope***. DB2 10 compression extends to a broader range of data structures – including permanent as well as temporary tables, all types of index, log files, large objects (LOBs), values, XML Data Areas (XDAs) and others – than Oracle Advanced Compression. DB2 10 automatically identifies opportunity areas.

- ***Automation***. DB2 10 requires less work by DBAs and/or storage administrators (job definitions and task structures vary between organizations) to achieve efficient, ongoing compression. Compression-specific tasks, as well as collection and interpretation of statistics, benefit from key IBM automation technologies.

Oracle Database 11g also supports Exadata Hybrid Columnar Compression (EHCC) on Oracle Exadata Database Machine. This technology is designed to compress large tables, and is most effective when these tables are processed sequentially. It would not be appropriate to compare EHCC with DB2 10. Exadata Database Machine competes primarily with high-end data warehouse appliances offered by Teradata, IBM, EMC (Greenplum) and others.

Temporal Data Management

Time Travel Query in DB2 10 distinguishes between system time (when an event is logged) and business time (an alternative date and/or time associated with the event) in maintaining and querying historical records. This feature, implements what is generally referred to as temporal database capability, responds to key informational, legal and compliance needs in a variety of industries.

Although temporal database concepts date back to the 1990s, structures in which the ability to maintain multiple timelines is a core design parameter have only recently entered the commercial mainstream. In addition to DB2 10, full temporal capabilities are otherwise offered only by the latest Teradata Database Version 14 and specialized databases.

For most organizations, temporal analysis has typically meant customizing applications to combine multiple timelines. This approach tends to be clumsy and time-consuming, and may impair database performance. Moreover, significant modifications to applications may be required if temporal parameters change. Designed-in temporal structures are a great deal more efficient and productive.

Although there is no direct Database 11g equivalent, two Oracle offerings provide similar capabilities:

1. ***Oracle Flashback Data Archive (FDA)***, also referred to as "Total Recall," was introduced as an option in Database 11g in 2007. It builds upon earlier Oracle Flashback technology. Historical transactions may be viewed in their original form, even if they have been subsequently written over or revised.

   FDA is employed primarily as a solution for long-term storage and recovery of compressed archival data for auditing purposes. Data may also be used for business analysis, although it is typically necessary to write customized code in order to do. Originally offered as a separately charged product, FDA now forms part of Oracle Advanced Compression for Database 11g.

2. ***Oracle Workspace Manager*** is a Database 11g add-on that allows developers to work with multiple timelines in dedicated "workspaces." These are, however, complex to administer, and require a great deal of DBA and developer effort over time. In addition, performance impact may be significant.

In this as in other areas, there are marked differences between IBM and Oracle approaches. IBM Time Travel Query incorporates state of the art temporal database technology, implemented in an elegant manner as part of core DB2 10 structures. Oracle solutions are overlays that require more time and effort to implement and administer, interoperate less efficiently and generate higher system overhead.

DB2 10 is the first major database to comply with the temporal features of the new ANSI/ISO SQL:2011 standard. Oracle has indicated that the company will support this standard in the future.

Storage Tiering

While both DB2 10 and Oracle Database 11g support disk controller-based automated storage tiering, DB2 10 offers a native database implementation, Multi Temperature Data Management.

The DB2 10 implementation provides full tiering capabilities, including support for solid state drives (SSDs), along with high-performance Fibre Channel (FC) and Serial Attached SCSI (SAS), and high-capacity Serial ATA (SATA) and Near-Line SAS (NL-SAS) drives in conventional and RAID configurations.

Key structures such as storage pools are built into the core database, rather than in disk controller software and microcode, while tiering functions are tightly integrated with DB2 10 WLM and highly automated. Movement of data between tiers occurs as a background process with minimal performance impact, and may be set to occur automatically or as a result of administrator intervention.

These capabilities address key weaknesses of array-centric automated storage tiering approaches. Statistics collection and data movement processes often generate high levels of system overhead, and frequent intervention by administrators may be required as workloads change.

It can be expected that, in high-volume, exceptionally performance-sensitive environments, array-based tiering approaches offered by IBM (Easy Tier) and others remain appropriate. DB2 10 Multi Temperature Data Management offers a comparatively efficient, low-cost alternative for organizations with less exacting workloads.

There is no Oracle Database 11g equivalent.

Real-time Data Input

"Real-time data warehousing," meaning BI applications that enable immediate access to continuously updated data stores, has been a major industry growth area since the mid-2000s.

Examples include "operational BI" (i.e., use of data warehouses by contact center, sales and other personnel interacting with customers); continuous analysis of sales, pricing and promotional data in marketing organizations; and continuous tracking of key performance indicators (KPIs) in a wide range of businesses.

Even higher levels of data currency have been adopted in time-sensitive e-commerce businesses, and in such industries as telecommunications and financial services for fraud detection, network monitoring and trading analysis. The general norm is that records should be accessible within five seconds.

Growth in such applications, as well as broader pressures to provide 24/7 availability of warehouse data, have caused a growing number of organizations to move away from overnight or weekly batch loading cycles. This trend is expected to accelerate.

The new Continuous Data Ingest feature in DB2 10 addresses these requirements. The feature, which employs a new IBM parallel loading technology, enables extremely fast, low-overhead transfers of data between external sources and DB2 10-based systems. It offers a higher-performance alternative to conventional batch and "trickle feed" loading techniques.

Continuous Data Ingest enables parallel loading of data from internal and/or external pipes (communications between database processes) and files into multiple DB2 10 partitions. This approach, illustrated in figure 7, avoids the delays and management overhead created by staging tables.

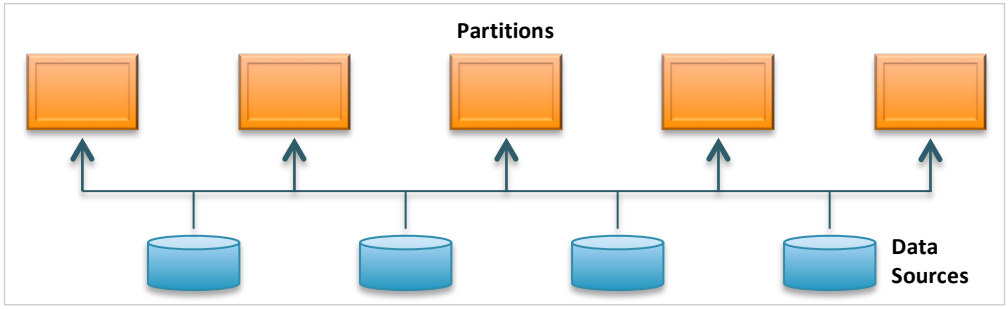

Figure 7: ***Parallel Data Loading using DB2 10 Continuous Data Ingest***

A variety of data formats are supported, and extract, transformation and load (ETL) processes may be performed concurrently.

Oracle's approach to real-time data warehousing is focused on the company's GoldenGate solution set, Continuous, high-volume input may be realized using Golden Gate, although, as in other areas of Oracle capability, throughput times are slower, and more administrator intervention is required.

GoldenGate is a separately charged and comparatively expensive feature, while Continuous Data Ingest is a standard function of DB2 10.

The no-charge Streams replication tool is also supported as part of Oracle Database 11g Enterprise Edition. Streams, however, performs poorly with large, complex data volumes and has been sidelined by Oracle in favor of GoldenGate.

Big Data Technologies

The term "big data" has been widely, and somewhat confusingly used in recent industry debates. It is used here to refer to technologies originally developed by search engine and social media companies to handle massive volumes of unstructured data that exceed the architectural limits of relational databases.

(Yahoo!, for example, manages more than 200 petabytes of unstructured data, while volumes at Facebook, Google and others are reported to be even larger.)

Key big data technologies include Hadoop and MapReduce, which were originally developed by Google and later transferred to the control of the Apache Software Foundation. The SPARQL query language and Resource Definition Framework (RDF) form part of this complex, often described as representing "NoSQL" standards.

In DB2 10, new support has been added for SPARQL and RDF data stores. The implementation conforms to standards defined by the World Wide Web Consortium (W3C), and allows integration of Hadoop and MapReduce material into DB2 10 databases.

A new DB2 10 feature, ***NoSQL Graph Store***, supports triple graph stores, which are commonly employed in big data environments to establish and illustrate relationships between different data sets. Tripled graph stores have, for example, been adopted by Facebook, LinkedIn and other social networks to establish connections between individuals on their sites.

Oracle also supports use of SPARQL and RDF on Oracle Database 11g Enterprise Edition, although the implementation is clumsier, and requires more DBA time and effort. In this as in other areas of functionality, DB2 10 capabilities are integrated into the core system structure rather than overlaid on it.

XML Support

Although IBM has moved aggressively to implement new big data technologies, the company continues to support the earlier Extensible Markup Language (XML) standard. DB2 10 users may employ either RDF and NoSQL, or native XML data models.

DB2 has supported XML since the introduction of DB2 9.1 in 2006, and capabilities have been repeatedly enhanced since that time. In DB2 10, IBM pureXML provides full support for XML storage, indexing, queries, updates and data management. The full range of DB2 10 capabilities – including database and range partitioning, multidimensional clustering, data compression, performance optimization and automation – extend to XML data content.

Oracle Database 11g supports XML and has been adopted by a number of organizations for XML database applications. Users have reported, however, that earlier DB2 versions delivered significantly higher levels of performance and functionality than Oracle equivalents, and disparities appear to have widened with the introduction of DB2 10.

Simplification and Automation

User Experiences

For most of DB2's history, higher levels of DBA productivity than for Oracle have been the norm. Core design features allow DBAs to perform tasks with fewer, simpler actions, in less time than their Oracle counterparts. This the case in most of the areas of DB2 capability described here, including performance optimization, compression and cluster management.

These strengths are reinforced by high levels of automation, which contribute to DBA productivity as well as performance (system parameters may be adjusted more rapidly and efficiently than with manual techniques) and availability (risks of performance bottlenecks and human error are reduced).

One result is that DB2 10 DBAs perform more work, more efficiently than their Oracle counterparts. One organization that contributed to this report estimated, for example, that "before" and "after" numbers of FTE DBAs for the applications migrated from Oracle to DB2 9.7 were as shown in figure 8.

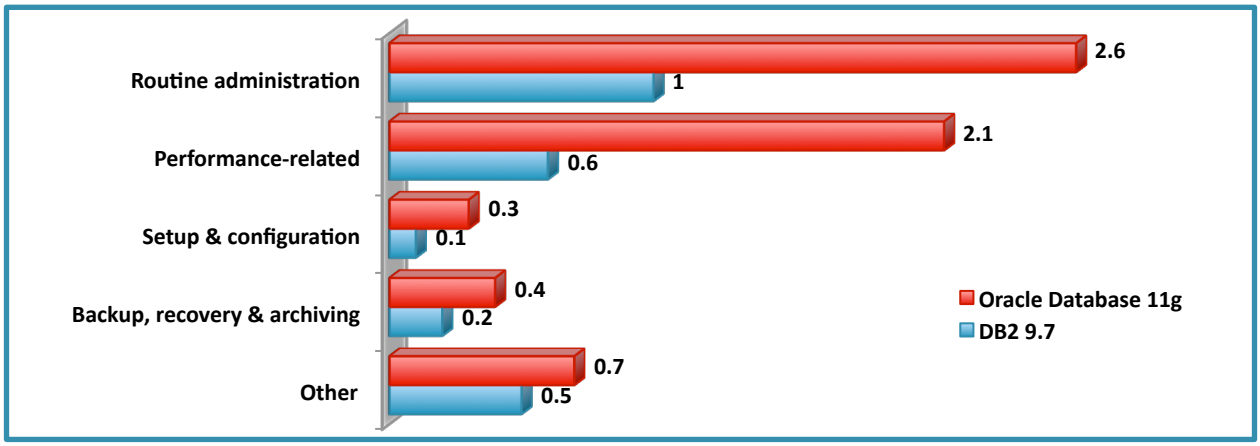

Figure 8: ***FTE DBA Staffing for DB2 9.7 and Oracle Database 11g – User Example***

A broader DB2 advantage should also be noted. In DB2 environments, tasks can often be handled by less experienced DBAs, allowing organizations to focus higher-level skills on more serious challenges. DBAs are among the most highly qualified – and most highly paid – members of most IT organizations. It makes little sense to employ them on repetitive, comparatively low value-added tasks.

DB2 10 Features

Simplification and automation are apparent in most areas of DB2 10 functionality. This is, for example, the case in the following areas:

- ***Server architecture***. DB2 instances can access multiple databases, while equivalent Oracle structures can access only one at a time (concurrent access is available only for RAC systems); new database instances may be created and managed more easily; and DB2 multithreading enables administrators to deal with only five primary processes, compared to 20+ in equivalent Oracle environments.

- ***Memory architecture***. DB2 buffer pools are more configurable, and automatic memory management is more flexible. In DB2 10, support has also been added for storage groups enabling Multi Temperature Data Management. In Oracle Database 11g, storage groups must be externally created and administered on disk arrays.

- ***Setting of parameters and environmental variables***. DB2 automatically generates parameters, and environmental variables may be set up and modified more easily; a single DB2 configuration at the instance or server level applies to all users – individual profile settings are not required, as is the case for Oracle; DB2 enables auto-configuration – there is no Oracle equivalent.

- ***Storage management***. DB2 Automatic Storage is tightly integrated into the core database design, while Oracle Automatic Storage Management (ASM) is implemented as a separate subsystem; DB2 Automatic Storage requires fewer, simpler commands; DB2 Adaptive Compression is more effective and simpler to administer than Oracle Database 11g Advanced Compression.

These strengths are reinforced by implementation of unique IBM autonomic technologies in a variety of features, including those shown in figure 9.

FEATURE	FUNCTION
Automatic storage management	Monitors & automatically creates, extends & adds storage device containers to support database growth across disk & file systems. Redefines storage paths as needed.
Self tuning memory manager	Tunes database memory settings & adjusts these in real time during run time to optimize performance for one or multiple concurrent databases. Increases DBA productivity, improves performance up to 10x, & reduces risks of bottlenecks.
Automatic maintenance	
• **Automatic/real-time statistics collection**	Determines whether statics need to be updated, & initiates collection for tables where this is the case. Real-time statistics may be generated in under five seconds if required.
• **Automatic reorganization**	Evaluates updated tables or indexes, determines whether reorganization is required, & schedules such operations during predefined periods.
• **Automatic backup**	Performs online &/or offline backups according to predefined schedules & criteria.
Workload manager	Enables fine-grain resource allocation, monitoring & management of workloads based on service classes, workload characteristics, elapsed time, time of day & other criteria. Supports up 64 primary service classes & 3,904 subclasses. Integrates with AIX, Linux & other external workload management facilities.
Silent installation	Allows DB2 installation based on application-specific setup & configuration information; i.e., no user input is required. Enables rapid startup & minimizes installation footprint.

Figure 9: *DB2 10 Autonomic Features*

The DB2 self-tuning memory manager (STMM), in particular, is one of the industry's most advanced self-tuning technologies.

Workload Management

Workload management, a longstanding DB2 strength, has been further enhanced in DB2 10.

DB2 capabilities in this area are derived from mainframe architecture, and draw upon mainframe strengths in managing diverse concurrent workloads. Organizations may realize higher levels of capacity utilization over time than with less well-optimized databases such as Oracle.

The DB2 10 Workload Manager enables highly granular, automated prioritization, along with resource allocation, queuing and real-time monitoring and management of workloads generated by hundreds or thousands of separate query jobs.

A key benefit is that service level agreement (SLA) targets may be met in a highly cost-effective and reliable manner. Priorities may be set based on user group, query type, time of day and other variables. Up 64 primary classes and 3,904 subclasses of service may be defined.

These strengths provide particular value in environments characterized by diverse transactional and/or query workload mixes. Such environments include organizational data warehouses that must handle a wide range of query types and sizes with varying degrees of time-sensitivity.

Oracle does not address workload management in Database 11g from a service level perspective. Allocation of resources to user groups and workloads is handled through the Database Resource Manager and related components of Oracle Database 11g. As described by Oracle, "resources are allocated to users according to a resource plan specified by the DBA."

In the Oracle approach, service level management is handled through separately-charged solutions offered by third parties or customized through the Oracle professional services organization and its partners.

Clustering Solutions

In this area, the principal IBM high-end offering is DB2 pureScale. DB2 High Availability Disaster Recovery (HADR) is also widely employed for less business-critical systems.

Both solutions have been enhanced in DB2 10. HADR, for example, now supports use of up to three "hot" standby servers, and allows delays to be set in the failover process to prevent replication of problems to these. PureScale enhancements include tighter integration with the core DB2 10 engine, closer alignment with WLM and support for range (table) partitions.

The equivalent Oracle offerings are Real Application Clusters (RAC) and Active Data Guard respectively. DB2 pureScale, and Oracle RAC and Active Data Guard are separately charged. HADR is included in DB2 10.

Oracle RAC and DB2 pureScale are positioned in three main roles, as enabling: (1) Tier 1 database failover and recovery in the event of a disastrous unplanned outage; (2) failover between systems in order to perform tasks that would otherwise require planned outages; and (3) realization of high levels of scalability by spreading database images across multiple physical systems.

There are, however, a number of differences in architecture and technology between the two solutions. DB2 pureScale, which was introduced in 2009, is a newer design than RAC and integrates a number of IBM technology components that are generally regarded as industry leading.

Differentiators include the following:

- ***Parallel Sysplex Data Sharing*** has supported the world's largest mainframe business-critical systems for close to 20 years. Production clusters routinely run from 2 to 10 nodes, and many are in the 10- to 20-node range. DB2 pureScale incorporates the overall design, along with distributed locking, cache management and other mechanisms derived from this architecture.

  The wide area extension of DB2 pureScale, Geographically Dispersed pureScale Clusters, also draws upon design features of its mainframe counterpart, Geographically Dispersed Parallel Sysplex (GDPS). Both solutions enable cluster failover between multiple remote sites in the event of a disastrous outage.

- ***General Parallel File System (GPFS)*** has been employed for more than a decade in commercial applications in financial services, pharmaceuticals, retail, government and other industries, and in scientific and technical computing. GPFS has demonstrated near-linear scalability in extremely large configurations. Installations of the RAC equivalent, Oracle Cluster File System, are typically a great deal smaller. The latest GPFS version, 3.5, was introduced in April 2012.

- ***Tivoli System Automation*** is another mainframe-derived solution set that automates failover and recovery processes within clusters. It has been extensively redesigned and reinforced with autonomic technologies for DB2 pureScale environments.

Differences in architecture and technology have a number of implications. For example, Oracle RAC failover is a more complex and slower process than is the case for DB2 pureScale, and tends to require more manual intervention by administrators.

Cluster overhead is also significantly higher for RAC than for DB2 pureScale. RAC clusters can typically be expanded to six to eight systems, depending on workloads, before performance degradation becomes unacceptable. In comparison, IBM tests demonstrate less than five percent performance degradation in up to 64-way clusters.

A further differentiator is that DB2 applications can be migrated "as is" to pureScale environments. In contrast, transitioning Oracle Database applications to RAC involves extensive modifications, and a great deal of testing is normally required.

SQL Compatibility with Oracle Database

IBM has placed a major emphasis on SQL compatibility with Oracle Database in order to minimize migration costs and difficulties. DB2 9.7 incorporated technologies from EnterpriseDB, an industry leader in this area, and DB2 10 expands compatibility features.

Experiences with database migrations have shown that these typically involve a range of costs, including conversion of data and applications, as well as retraining DBAs and developers to work with new tools, or the hiring of new specialists. A loss of productivity may also be expected as organizations adapt to new technologies and practices.

Migrations from Oracle databases to DB2 9.7, however, showed that these effects were minimal. This reflects a number of features, including native support for Oracle Procedural Language/Structured Query Language (PL/SQL) and the Oracle SQL dialect, along with a wide range of code, tools and functions commonly employed by Oracle developers. Examples are shown in figure 10.

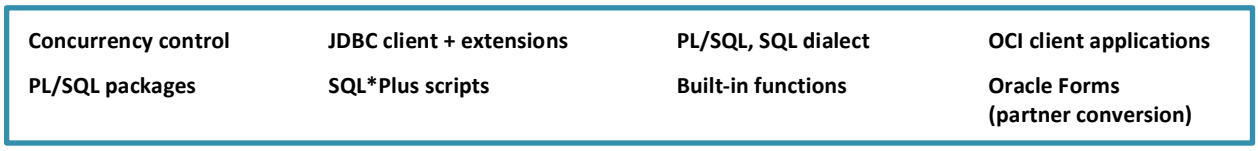

Figure 10: ***Commonly Used Oracle Features Supported by DB2 10***

Organizations that migrated Oracle applications to DB2 9.7 routinely found that more than 95 percent of code had remained unchanged, and with DB2 10, the proportion appears to be closer to 98 percent. Few or no changes to existing Oracle development tools and skills are required, and transition periods are relatively short – in some cases, less than two weeks – and non-disruptive.

A further benefit is that, in DB2 10, Oracle compatibility functions are built into the core engine rather than implemented in the form of a software overlay. Thus, organizations experience the same levels of performance as native DB2 10 users.

Although FTE DBA staffing is typically reduced by Oracle to DB2 migrations, in practice few organizations lay off staff. Higher levels of productivity make it easier to meet new challenges without increasing headcounts, and to apply the skills of key personnel to other projects and operational tasks.

Migration benefits translate into cost avoidance (it is not necessary to add personnel to deal with increasing database complexity, new application deployment and storage growth), improvements in quality of service experienced by users, or – if transitions are properly managed – both may be realized.

DB2 compatibility is not a focus for Oracle. There are few migrations from DB2 for Linux, UNIX and Windows to Oracle.

Database Management Solutions

Relational databases are complex structures with large numbers of interdependent components. Relational environments commonly include millions of lines of code. The manner in which such environments are managed by DBAs may have major impacts on performance, quality of service and costs.

IBM and Oracle both offer tools supporting a wide range of operational DBA tasks – e.g., performance tuning, database administration, backup and recovery management – as well as data modeling, and database design and development.

Operational and development-related DBA tasks are separated in many organizations, although closer integration may yield a number of advantages. For example, if causes of bottlenecks and slowdowns identified by operational DBAs are fed back into development processes, comparable problems may be avoided in the future.

IBM tools included in the DB2 Advanced Enterprise package are shown in figure 11.

COMPONENT	DESCRIPTION
DEVELOPMENT & ADMINISTRATION	
IBM Data Studio	Eclipse-based tool for collaborative database development & administration. Includes features for instance, object, data, job & connection management; defining & implementing database schema changes; centralized health monitoring; & other functions. Combines features of three previously separate IBM products: Optim Development Studio, Optim Database Administrator & Data Studio.
InfoSphere Data Architect	Eclipse-based database design, data modeling & integration tool supporting use of logical, physical & dimensional models. Integrates with other IBM Optim, Rational & InfoSphere Warehouse solutions.
PERFORMANCE OPTIMIZATION	
InfoSphere Optim Performance Manager Extended Edition	Enables rapid identification, diagnosis & resolution of database & application performance bottlenecks. Data may also be used to predict future resource requirements, to proactively identify high-cost queries & workloads for tuning & to prevent recurrence of failures in new applications.
InfoSphere Optim Configuration Manager	Enables extremely granular tracking of configuration changes initiated by DBAs, developers, end users & others, & enables immediate diagnosis & resolution of problems caused by these. Tracks & enables centralized management of client-side parameters. Provides comprehensive view of the entire database client & server topology.
InfoSphere Optim Query Workload Tuner	Provides expert, actionable advice to resolve database & application performance bottlenecks. Recommends indexes, statistics, statistical views, MDC & other parameters based on complete workload context. Supports what-if analysis, & estimates performance improvements for potential courses of action. Integrates with InfoSphere Optim Performance Manager for identification of high-cost queries & workloads. Provides performance improvement advise for these.
InfoSphere Optim pureQuery Runtime	Data access platform enabling improved performance, security & manageability of database client applications. Provides runtime environment & application programming interface (API) enabling higher performance of existing in-house applications without modifications. Also assists in rapid development of new applications.

Figure 11: ***Database Management Solutions within DB2 10 Advanced Enterprise***

The Oracle portfolio is built around an established set of offerings that include the company's Diagnostics and Tuning, along with Change Management, Configuration Management, and Provisioning and Patch Automation Packs. The last three of these were recently combined into a new Database Lifecycle Management Pack, although the individual tools were not substantially changed.

There are a number of differences between IBM and Oracle approaches. First, within the IBM portfolio, tools are better integrated across multiple operational DBA tasks, and across development and operational DBA tasks – the term "lifecycle management" is more appropriate for IBM than for Oracle offerings.

Second, IBM tends to emphasize proactive rather than reactive functionality. Capabilities for monitoring and predictive problem identification are generally better. Oracle functionality is more geared to enabling DBAs to resolve problems when they occur, rather than preventing their occurrence.

Third, in tooling as in core database design, IBM offerings are designed for higher levels of DBA productivity. Task structures are more streamlined, and more automated than for Oracle equivalents. The widespread perception that Oracle tooling is more "labor-intensive" is substantially correct.

Certain IBM tools are recognized industry leaders. InfoSphere Optim Performance Manager Extended Edition, for example, integrates monitoring, analysis and optimization tasks across application as well as database layers of DB2 environments.

InfoSphere Optim Configuration Manager enables extremely granular tracking of configuration changes by DBAs, developers and end users, and allows DBAs to immediately pinpoint problems caused by these. InfoSphere Optim Query Workload Tuner may reduce problem resolution times from days to hours. All have enjoyed strong customer reception.

DB2 10 Advanced Enterprise tools support the full range of new DB2 10 capabilities.

DETAILED DATA

Basis of Calculations

Overview

The cost comparisons presented in this report are for four installations of DB2 10 Advanced Enterprise and equivalent Oracle Database 11g Enterprise Edition software stacks in telecommunications, financial services, heath care (non-profit) and retail companies. These companies are summarized in figure 12.

Telecommunications Company	Financial Services Company	Health Care Organization	Retail Company
BUSINESS PROFILE			
Wireless & Internet access services provider $1.5+ billion sales 10+ million subscribers 3,000 employees	Diversified retail bank $30 billion+ assets 500+ branches 7 million accounts 20,000 employees	Multidisciplinary hospital & medical complex $3 billion sales 1,000 beds 20,000+ staff	Grocery retailer $15+ billion sales 300+ stores 75,000 employees
APPLICATIONS			
Billing, CRM, operational systems & data store	Core banking systems	EMR, HIS, PACS, business applications, departmental	Analytical data warehouse

Figure 12: ***Installations Summary***

Configurations reflect input from these companies, and from others with similar business profiles, application portfolios and database characteristics. Performance, compression, FTE staffing and other variables employed were based on information from users that had deployed DB2 and Oracle databases for similar applications, or had migrated from Oracle to recent versions of DB2. Altogether, 64 user organizations contributed data.

In some cases, organizations employed older servers and disk storage. Where this was the case, configurations were updated to POWER7-based servers and current-generation storage systems based on IBM rPerf ratings and on installed capacity in terabytes as reported by users.

Allowance was made for growth in disk, tape and database server capacity, and in FTE staffing over a three-year period. Resulting configurations and FTE staffing levels were as shown in figures 13 and 14.

Costs were then calculated as follows:

- ***Database software*** costs were based on use DB2 Advanced Enterprise and Oracle 11g Enterprise Edition with separately charged features, shown in figure 6, providing equivalent functionality.

  DB2 10 pureScale and Oracle RAC were employed in the telecommunications, financial services and health care companies. Calculations include initial license as well as applicable Software Subscription and Support (IBM) or Software Update License and Support (Oracle) fees. Oracle costs were calculated based on per processor prices.

Telecommunications Company	Financial Services Company	Health Care Organization	Retail Company
DATABASES & TOOLS			
Oracle Database 11g Enterprise Edition Stack, including RAC	*Oracle Database 11g Enterprise Edition Stack, including RAC*	*Oracle Database 11g Enterprise Edition Stack, including RAC*	*Oracle Database 11g Enterprise Edition Stack*
DISK SYSTEMS			
Initial Configuration			
2 x DS8800 x 144 TB *Global Mirror, FCM, TPC*	2 x DS8800 x 72 TB *Metro Mirror, FCM, TPC*	2 x XIV 132 TB, 1 x XIV 102 TB *System Software*	2 Storwize V7000 x 96 TB *System Software*
End of Period Configuration			
2 x DS8800 x 288 TB *Global Mirror, FCM, TPC*	2 x DS8800 x 115 TB *Metro Mirror, FCM, TPC*	2 x XIV 203 TB, 1 x XIV 225 TB *System Software*	4 x Storwize V7000 x 118 TB *System Software*
TAPE SYSTEMS			
Initial Configuration			
TS3500 6 x LTO-5	TS3500 4 x LTO-5	TS3500 8 x LTO-5	TS3500 4 x LTO-5
End of Period Configuration			
TS3500 11 x LTO-5	TS3500 6 x LTO-5	TS3500 15 x LTO-5	TS3500 9 x LTO-5
DATABASE SERVERS			
Initial Configuration			
2 x 780 48 x 3.92 GHz *AIX, PowerHA*	2 x 780 32 x 3.92 GHz *AIX, PowerHA*	2 x 780 32 x 3.86 GHz *AIX, PowerHA*	1 x 770 24 x 3.3 GHz *AIX*
End of Period Configuration			
2 x 795 96 x 3.72 GHz *AIX, PowerHA*	2 x 780 64 x 3.92 GHz *AIX, PowerHA*	2 x 780 64 x 3.86 GHz *AIX, PowerHA*	1 x 770 48 x 3.3 GHz *AIX*
PERSONNEL			
Initial FTEs			
3.85	2.5	3.1	1.7
End of Period FTEs			
4.45	3.05	3.7	2.35

Figure 13: ***Configurations and FTE Staffing Summary – Use of Oracle Database 11g Stack***

- ***Disk systems*** calculations were based on use of three IBM platforms:

  1. ***System Storage DS8800*** arrays were employed in the telecommunications and financial services company installations. Systems were configured for RAID 5 with 300 gigabyte (GB) high-performance SAS drives. Software included DS8800 Operating Environment License (OEL), FlashCopy Manager (FCM) for point-in-time-copy, Metro Mirror or Global Mirror for real-time replication and Tivoli Storage Productivity Center (storage management).

  2. ***XIV Storage System Gen3*** arrays configured with 3 terabyte (TB) NL-SAS drives, the XIV system software suite and IBM FlashCopy Manager were employed for the health care company. XIV systems do not employ RAID technology. The XIV systems software suite provides remote replication and storage management functions.

  3. ***Storwize V7000*** arrays configured with 600 GB high-performance SAS drives, RAID 5 and the Storwize V7000 Software suite supported the retail company data warehouse installation. Storwize V7000 Software provides storage management functions. Real-time replication for disaster recovery was not employed.

     Values shown for DS8800 and Storwize V7000 systems are for raw physical capacity, while those for XIV systems are for usable capacity – XIV systems are configured and sold by IBM based on this metric.

Telecommunications Company	Financial Services Company	Health Care Organization	Retail Company
DATABASES & TOOLS			
DB2 10 Advanced Enterprise including pureScale	*DB2 10 Advanced Enterprise including pureScale*	*DB2 10 Advanced Enterprise including pureScale*	*DB2 10 Advanced Enterprise*
DISK SYSTEMS			
Initial Configuration			
2 x DS8800 x 79 TB *Global Mirror, FCM, TPC*	2 x DS8800 x 43 TB *Metro Mirror, FCM, TPC*	2 x XIV 84 TB, 1 x XIV 55 TB *System Software*	2 x Storwize V7000 53 TB *System Software*
End of Period Configuration			
2 x DS8800 x 158 TB *Global Mirror, FCM, TPC*	2 x DS8800 x 65 TB *Metro Mirror, FCM, TPC*	2 x XIV 203 TB, 1 x XIV 225 TB *System Software*	2 x Storwize V7000 125 TB *System Software*
TAPE SYSTEMS			
Initial Configuration			
TS3500 3 x LTO-5	TS3500 3 x LTO-5	TS3500 5 x LTO-5	TS3500 3 x LTO-5
End of Period Configuration			
TS3500 6 x LTO-5	TS3500 4 x LTO-5	TS3500 9 x LTO-5	TS3500 5 x LTO-5
DATABASE SERVERS			
Initial Configuration			
2 x 780 32 x 3.92 GHz *AIX, PowerHA*	2 x 780 32 x 3.92 GHz *AIX, PowerHA*	2 x 780 32 x 3.86 GHz *AIX, PowerHA*	1 x 770 16 x 3.3 GHz *AIX*
End of Period Configuration			
2 x 780 64 x 3.92 GHz *AIX, PowerHA*	2 x 780 48 x 3.92 GHz *AIX, PowerHA*	2 x 780 48 x 3.86 GHz *AIX, PowerHA*	1 x 770 32 x 3.3 GHz *AIX*
PERSONNEL			
Initial FTEs			
2.05	1.25	1.4	0.75
End of Period FTEs			
2.35	1.5	1.65	1.0

Figure 14: *Configurations and FTE Staffing Summary – Use of DB2 10 Advanced Enterprise*

- ***Tape systems*** calculations were based on use of IBM TS3500 Tape Library systems employing Linear Tape Open Ultrium Generation 5 (LTO-5) drives and 1.5 TB Data Cartridges. Systems were configured based on size of backup data volumes, and on the frequency and duration of backup operations. Costs included handling and secure storage of cartridges.

- ***Database server*** calculations were based on use of IBM Power Systems 780, 795 and, for the retail company data warehouse installation, 770 servers. The AIX 7.1 operating system and PowerVM virtualization were employed for all installations. Calculations for the telecommunications, financial services and heath care companies include IBM PowerHA SystemMirror failover clustering software.

- ***Personnel*** cost calculations were based on annual average FTE salaries of $94,610 and $92,140 Oracle 11g and DB2 10 database administrators respectively. Annual average FTE salaries of $72,649 and $82,722 for storage and server administrators respectively were employed for DB2 10 and Oracle Database 11g calculations. All salaries were increased by 52.3 percent to allow for benefits, bonuses and related items.

- ***Data center*** cost calculations include data center occupancy and energy consumption, as well as allowance for acquisition, maintenance and operational costs for data center infrastructure equipment such as uninterruptable power supplies (UPS), power distribution systems (PDS) and cooling systems. Costs for this equipment were calculated based on prorated values per kilowatt-hour (kWh).

  Energy consumption costs were calculated for storage systems and database servers along with data center equipment supporting these based on average utilization levels (e.g., 30 percent, 70 percent) and annual hours of operation. A conservative assumption for average price per kWh was employed.

  Occupancy costs include allowance for service clearances and inactive areas. Costs were calculated using a conservative assumption for annual average cost per square foot for existing facilities (i.e., costs do not include new facilities construction).

All cost values were for the United States.

Cost Breakdowns

Detailed breakdowns are presented in figures 15 and 16.

ORACLE DATABASE 11g	Telecommunications Company	Financial Services Company	Health Care Organization	Retail Company
DATABASES ($000)				
	9,600.5	**6,573.7**	**6,573.7**	**2,158.6**
DISK SYSTEMS				
Hardware & maintenance	4,155.9	2,675.9	1,875.3	678.2
Software licenses & support	2,442.4	1,712.3	916.0	138.2
Total ($000)	**6,598.3**	**4,388.2**	**2,791.3**	**816.4**
TAPE SYSTEMS				
Hardware & maintenance	183.6	108.0	216.6	122.4
Media	224.3	203.5	343.7	172.4
Total ($000)	**407.9**	**311.5**	**560.3**	**294.8**
DATABASE SERVERS				
Hardware + maintenance	2,954.5	1,986.1	1,865.1	730.5
SW licenses & support	1,528.8	1,048.3	817.7	106.1
Total ($000)	**4,483.3**	**3,034.4**	**2,682.8**	**836.6**
DATA CENTER ($000)				
	372.5	**291.2**	**201.8**	**125.6**
PERSONNEL ($000)				
	1,653.2	**1,113.0**	**1,393.5**	**797.0**
TOTAL THREE -YEAR COSTS ($000)				
	23,115.7	**15,711.9**	**14,203.4**	**5,029.0**

Figure 15: ***Three-year Cost Breakdowns – Use of Oracle Database 11g Stack***

DB2 10 ADVANCED ENTERPRISE	Telecommunications Company	Financial Services Company	Health Care Organization	Retail Company
DATABASES ($000)				
	5,785.3	**4,071.2**	**4,339.0**	**1,065.3**
DISK SYSTEMS				
Hardware & maintenance	3,194.3	1,739.6	1,112.1	339.8
Software licenses & support	1,792.2	982.9	465.1	69.1
Total ($000)	**4,986.5**	**2,722.5**	**1,577.2**	**408.9**
TAPE SYSTEMS				
Hardware & maintenance	102.1	77.1	157.1	68.1
Media	129.2	123.4	214.8	102.2
Total ($000)	**231.3**	**200.5**	**371.9**	**170.3**
DATABASE SERVERS				
Hardware + maintenance	2,092.9	1,640.0	1,565.6	547.9
SW licenses & support	1,048.3	393.1	613.3	68.6
Total ($000)	**3,141.2**	**2,033.1**	**2,178.9**	**616.5**
DATA CENTER ($000)				
	222.8	**186.9**	**120.5**	**69.2**
PERSONNEL ($000)				
	847.6	**534.9**	**614.7**	**329.1**
TOTAL THREE -YEAR COSTS ($000)				
	15,214.7	**9,749.1**	**9,202.2**	**2,659.3**

Figure 16: ***Three-year Cost Breakdowns – Use of IBM DB2 10 Advanced Enterprise***

ABOUT THE INTERNATIONAL TECHNOLOGY GROUP

ITG sharpens your awareness of what's happening and your competitive edge
. . . this could affect your future growth and profit prospects

International Technology Group (ITG), established in 1983, is an independent research and management consulting firm specializing in information technology (IT) investment strategy, cost/benefit metrics, infrastructure studies, deployment tactics, business alignment and financial analysis.

ITG was an early innovator and pioneer in developing total cost of ownership (TCO) and return on investment (ROI) processes and methodologies. In 2004, the firm received a Decade of Education Award from the Information Technology Financial Management Association (ITFMA), the leading professional association dedicated to education and advancement of financial management practices in end-user IT organizations.

The firm has undertaken more than 120 major consulting projects, released more than 250 management reports and white papers and more than 1,800 briefings and presentations to individual clients, user groups, industry conferences and seminars throughout the world.

Client services are designed to provide factual data and reliable documentation to assist in the decision-making process. Information provided establishes the basis for developing tactical and strategic plans. Important developments are analyzed and practical guidance is offered on the most effective ways to respond to changes that may impact complex IT deployment agendas.

A broad range of services is offered, furnishing clients with the information necessary to complement their internal capabilities and resources. Customized client programs involve various combinations of the following deliverables:

Status Reports	In-depth studies of important issues
Management Briefs	Detailed analysis of significant developments
Management Briefings	Periodic interactive meetings with management
Executive Presentations	Scheduled strategic presentations for decision-makers
Email Communications	Timely replies to informational requests
Telephone Consultation	Immediate response to informational needs

Clients include a cross section of IT end users in the private and public sectors representing multinational corporations, industrial companies, financial institutions, service organizations, educational institutions, federal and state government agencies as well as IT system suppliers, software vendors and service firms. Federal government clients have included agencies within the Department of Defense (e.g., DISA), Department of Transportation (e.g., FAA) and Department of Treasury (e.g., US Mint).

International Technology Group

609 Pacific Avenue, Suite 102
Santa Cruz, California 95060-4406
Telephone: + 831-427-9260
Email: Contact@ITGforInfo.com
Website: ITGforInfo.com