

Oracle Database10g vs IBM DB2 UDB 8.1 Technical Overview

An Oracle Competitive White Paper
March 2004

Technical Overview of Oracle Database 10g with IBM DB2 UDB V8.1

Introduction	3
Grid Computing	3
Manageability	7
Self-Managing Database	8
Application/SQL Tuning.....	10
Managing the Enterprise.....	11
High Availability	12
Bounding Database Crash Recovery	13
Human Error Recovery	14
Online Maintenance	15
Data Centre Disasters	15
Data Warehousing and Business Intelligence	16
Partitioning	18
Data Loading and Archiving.....	19
Operational Impact	21
Application Development.....	21
Web Based Development Tool	23
Globalization	24
Managing all Your data.....	24
Unstructured Data	24
XML.....	25
Conclusion.....	26

Technical Overview of Oracle Database 10g with IBM DB2 UDB V8.1

INTRODUCTION

Oracle Database 10g is the first database designed for Enterprise Grid computing. Oracle Database 10g reduces the costs of data management while providing the highest quality of service. It allows IT to rapidly respond to the needs of the business while greatly lowering risk. Above all, it's easy to deploy and manage.

With superior, innovative technology, Oracle's leadership continues to outpace the competition. Oracle Database 10g widens the technology gap further from IBM DB2 UDB v8.1 for Linux, Unix, and Windows. Press, analysts, and industry experts echo what Oracle customers already know--Oracle Database is the leader in performance, scalability, availability, security, and pricing. Oracle's continued leadership and innovation allows businesses of all sizes to do more while spending less.

This paper provides a technical comparison of Oracle Database 10g with IBM DB2 UDB V8.1 for Linux, Unix and Windows. It illustrates the superiority of Oracle's 'best-of-both-worlds' architecture over pure shared nothing database systems in terms of performance, scalability, resource utilization, manageability, availability, and consistency with emerging trends in technology.

GRID COMPUTING

Oracle10g is the first and only infrastructure software engineered for Grid computing. Oracle Database 10g and Application Server 10g deliver full clustering, workload management and data center automation for flexible, dynamic Grid computing. Customers using Oracle's Grids will realize higher resource utilization and lower costs. They will also benefit from Oracle's superior operational characteristics – portability, availability, security and scalability. Oracle portability ensures you get the same operational benefits on all platforms including Linux and commodity clusters. Only Oracle can truly scale, provide high availability, and dynamically provision resources on low cost commodity clusters. Oracle makes the Grid unbreakable – you cannot break an Oracle grid and you cannot break into an Oracle grid.

Oracle also has key Grid technology differentiators – such as Oracle Real Application Clusters, Oracle Streams, and Oracle Transportable Tablespaces,

Financial outfits such as Bear Stearns are citing 10g as providing a "compelling competitive advantage" over database software from rivals IBM and Microsoft. It's a competitive advantage that should persist for years, they say, and one that means Oracle is well on its way to solidifying its position as the No. 1 vendor in the RDBMS market.

<http://www.eweek.com/article2/0,1759,1549121,00.asp>

among many others. Most importantly, Oracle has a proven record of providing software for leading platforms and environments. Oracle Grid customers can feel confident their investment in Oracle technology will be leveraged as the Grid evolves.

"Oracle Grid technology provided us with the necessary features to efficiently manage our workload; thus delivering more with less." -- Rich Adelson, Enterprise Architect, The Hartford

The Hartford shares the same Oracle database resources for multiple applications - billing, reconciliation, cash, and receivables; thus achieving more with less. This was only possible with the Oracle Grid technology. Oracle technology, including Oracle Resource Manager and Oracle Streams, provided The Hartford with the capability to manage the workload and allocate resources fairly to all these applications.

Oracle Database and Oracle Application Server met the LHC Computing Grid's technology requirements.

CERN (European Organization for Nuclear Research) is building a Large Hadron Collider (LHC) Computing Grid to manage data generated by the LHC experiments. The largest experiments will generate one petabyte of data per year. Around 2000 users and 150 institutes will be involved in conducting and analyzing the data from these experiments. CERN has decided to use Oracle for the LHC Computing Grid. Databases and application servers are an important component of their solution. Oracle Database and Oracle Application Server met the LHC Computing Grid's technology requirements. Oracle offers the scalability, the information sharing, and VLDB features necessary for thousands of users to share and access data within the LHC Computing Grid. Oracle Transportable Tablespaces offers CERN an extremely fast mechanism to share large amounts of data across many sites. Unlike other methods for sharing data, Transportable Tablespaces can move large quantities of data very quickly, regardless of the amount of data involved. In addition, Oracle Real Application Clusters offers CERN a highly available and scalable database on commodity Linux clusters.

IBM has aggressively marketed its "On-Demand Computing" brand for the last year. IBM's offering consists of turnkey hosting services, modular based on open standards and custom software development from its Business Consulting Services. The key theme is integration (DB2 Information Integrator), automation (SMART) and virtualization (through hardware offerings). IBM sells large SMP servers with hardware partitioning of CPUs and term licensing of processors based on demand. This is just a recycled mainframe business model with all its costs and limitations. On Demand pieces this environment together with WebSphere as the glue. The IBM software involved in their Grid offerings are DB2 Information Integrator, DB2 Content Management Server, and WebSphere. IBM may have a library of academic white papers on Grid computing, but IBM software products are not designed for the Grid.

"Oracle 10g is the most useful infrastructure software for grid computing." Titus Mteleka, Director of Science and Technology in the Ministry of Science, Tanzania

Oracle is better for Grid computing for many reasons including:

- Oracle runs everything on commodity hardware. IBM will recommend that their database run on the more expensive SMP. Their cluster database offerings do not run real applications (for example SAP, PeopleSoft, Siebel).

"We chose Oracle because we knew it would meet our stringent uptime requirements, and it did. It has simplified deployment of mission critical applications by providing a standard for highly available database services." -- Kerry Schwab, Senior Manager of Operations for Interactive Marketing, Southwest Airlines

"One of the things I particularly like about Oracle9i Real Application Clusters is having the ability to expand quickly without all the trouble of massively upgrading hardware." -- Keith Misson, Chief Technology Officer, bluecycle.com Ltd.

"Oracle 10g with its automated storage management capability, has a vast improvement in database manageability features, and will help us save time and resources in the long run." G.M. Shenoy, senior vice president, NSE.IT Limited

"We're happy about what Advanced Queuing offers us," says Domenico Betunio, Banca IMI's manager of electronic trading. "Besides speed and scalability, we're impressed by messaging reliability, and especially auditing."

- Oracle RAC can dynamically add instances or remove instances from a database with no downtime. This means that you can change resource allocations to different databases in response to load or management priorities. IBM cluster database is shared nothing. If you add a node, you have to repartition your data, which means downtime. So, you have to size a shared nothing database for peak, and that means poor utilization.
- If you have to partition your software between blades and SMP, you have already made provisioning decisions. The software that runs on blades can't take advantage of idle SMP, and the software on SMP can't utilize idle blades. So, a heterogeneous hardware mix already limits the utilization you can achieve. To get the best resource utilization, and best response to your priorities and loads, you need a homogeneous hardware environment. A completely homogeneous blade environment is much lower cost. But, only Oracle's software stack can run entirely on blades.
- Oracle Database 10g includes Automatic Storage Management (ASM) that provides a vertical integration of a file system and volume manager specifically built for the Oracle database files. ASM distributes I/O load across all available resource to optimize performance while removing the need for manual I/O tuning. ASM helps DBAs manage a dynamic database environment by allowing them to grow the database size without an outage to the database when adjusting the storage allocation.
- Oracle Streams allows complete asynchronous information sharing with a single unified infrastructure. It has a really unique architecture, combining message queuing, replication, events, data warehouse loading, notifications and publish/subscribe into a single technology. You can share information incrementally, as it changes, any way you need to. You don't have to install, administer, and integrate a bunch of different products. IBM can not share information in all these ways without having to combine many different products, write a lot of code, and manage a lot of different components.
- Oracle Transportable Tablespaces lets you move entire portions of the Oracle database. While Oracle Streams lets you share information incrementally, as it changes, Transportable Tablespaces let you move data in bulk, quickly, even across different platforms. Of course, the two can work together.
- Sometimes, the best way to share and integrate information is just to access it remotely when you need it. Oracle has distributed SQL, remote procedure calls, and distributed transactions built in. You can query other databases for information you need, update that information when you need to, and access program logic in other servers. You can access data in other database management systems through Oracle's generic gateway, or using our Transparent Gateways. You can federate Oracle and non-Oracle data. And, again, you can use Oracle Streams with our gateways to share

"We're using Oracle Real Application Clusters to provide a consistent customer experience . . . It's given us a more reliable site, with reduced planned and unplanned downtime."
-- Jim Knight, Director of Technical Services, REI

information with non-Oracle systems. IBM requires you to buy DataJoiner if you want to update non-DB2 databases, and they don't have anything like generic gateways or Streams.

- Oracle's stack offers the high availability that Grid computing demands. Oracle RAC and Oracle Data Guard combine to insulate you from outage within a cluster, and outage across the whole cluster. This insulation is against both planned (patches, upgrades) and unplanned (server failure, disaster) outages. The whole notion of computing as a utility won't work if the computing isn't available. IBM has standby database technology, but you can't use the standby unless the primary goes down – guaranteeing you low resource utilization!
- Oracle performs and scales better than any other database in the world. That's important for Grid computing. Our read consistency model ensures that you get the highest throughput, and avoid the locking problems you can have with IBM and Microsoft. We have numerous customers running databases larger than a Terabyte. And, our RAC technology has no architectural limit to the number of nodes we can support. Oracle is currently the leader in TPC-C, TPC-H @1000GB, TPC-H @3000GB, and TPC-R @100GB as well as holding the world record TPC-C on Linux, 4-way, 16-way, and 64-way systems¹. In addition to TPC benchmarks, Oracle also runs real-world enterprise class applications benchmarks such as those on Oracle Applications and SAP². As of this publication:
 - Oracle holds the top three 2-tier Assemble to Order ATO benchmarks with a record of 34,260 orders per hour.
 - Oracle holds the top two 3-tier Assemble to Order ATO benchmarks
 - Oracle holds the top two SAP SD 2-tier results.
- Oracle is secure. Oracle endeavors to build security features and solutions into each of its products, particularly the database server, where data is stored. Security is at the core of the coding practices employed by the development staff that builds the Oracle database, resulting in the delivery of a secure product. Oracle security mechanisms have been validated by 17 independent security evaluations. IBM addresses security by delivering it outside of the database and relying on the operating system or Tivoli's product line to secure DB2 and other IBM products. IBM doesn't have features like Virtual Private Database and Label Security. And IBM has no independent security evaluation.
- Oracle portability across operating systems and hardware means Oracle can support heterogeneous grids. IBM's DB2 database doesn't have a consistent

"We use Oracle Label Security to assist us in protecting different types of health information commensurate with its level of sensitivity." --
Dixie B. Baker, Ph.D. Vice President, Chief Technology Officer Enterprise and Health Solutions Sectors, SAIC

¹ <http://www.tpc.org>

² <http://www.sap.com/benchmark/>

"Oracle Database was first choice of T-Mobile Austria. It has successfully proved its scalability and hardware portability. Today we can rely on our vital Enterprise Data Warehouse covering all aspects of the value chain at T-Mobile Austria."
Marcus Berger, CIO, T-Mobile Austria

code base across all platforms, so you can't do that with DB2 applications. Oracle runs on all major operating systems, and has the same consistent features and capabilities on all of them, since we use a single code base. Oracle's portability means that whatever direction the Grid takes, whatever operating system dominates the Grid, Oracle will work there.

- DB2 imposes a limit on the row size depending on the page size of the table space. A row cannot span multiple pages and so its size must be a little smaller – taking page overhead into account – than the page size (max 32677 bytes). Oracle allows a row to span across multiple pages or, in Oracle's terminology, blocks thereby completely avoiding this complication. The maximum database size for Oracle Database 10g is 8 exabytes(10^{18} bytes!). Alternatively, the maximum size for DB2 UDB is 8 petabytes(10^{12} bytes). Oracle's unlimited size allows flexibility for your Grid environment
- Oracle gives you all these capabilities in a single integrated stack. Our model is to sell integrated software, whether that's the E-Business suite or Grid infrastructure. Oracle gives complete integration in addition to being open and supporting all relevant standards. IBM and Microsoft give you components that can be assembled to give partial support for Grid computing. You have to assemble the components, or buy consulting to have them assembled. Oracle fully supports Grid computing, out of the box, with an integrated software stack.
- Oracle is releasing a version of the Globus Toolkit tailored to the Oracle stack. The Globus toolkit is a set of modules that support resource management, security, and discovery. IBM has shipped the Globus toolkit, but has not created a version tailored to their technology, as we have done. So, if you want to use Globus software with IBM technology, you must buy services from IBM, or spend the time and money to integrate the two yourselves.

MANAGEABILITY

Today's business environment has become increasingly competitive and the challenge before today's enterprise is to manage their information technology (IT) infrastructure at the lowest possible cost without compromising service level objectives. This presents an interesting dilemma since as systems provide more and more capabilities, complexity can increase, and hence these systems can become expensive to manage and sustain. A profitable business today not only needs reliable, high-performing applications to service its customers, but also must minimize its cost of operation to maximize profitability. This coupled with ever increasing shortage of skills in this area can quickly lead to spiraling management costs and can even hinder an organization's ability to remain competitive. Recognizing these demands, Oracle Database 10g introduces a sophisticated self-

As far as manageability goes, Oracle has it all over IBM.

<http://www.eweek.com/article2/0,1759,1549134,00.asp>

managing database that automatically monitors, adapts, and tunes itself. Oracle Database 10g simplifies every aspect of database administration in order to serve the dual objective of enhancing administrator productivity and helping customers reduce their operational cost by 50%. This extends the lead of Oracle Database over IBM DB2 UDB in ease of database management. In April 2003 Rauch Associates completed their Comparison of Management Cost of Ownership: Oracle9i Database Release 2 and IBM DB2 v8.1 which concluded that Oracle9i Database Release 2 is significantly easier to manage than IBM DB2 V8.1:

Significant savings in IT staffing and management costs are achievable due to the impressive management capabilities in Oracle9i Database Release 2.

- Oracle9i Database Release 2 required 41% fewer administrative steps than DB2 V8.1.
- Oracle9i Database Release 2 required 54% less time to perform equivalent tasks.
- Oracle9i Database Release 2 administrators require only 55% of the time used by DB2 peers to complete equivalent management workloads.
- Cost benefits from increased efficiency of Oracle9i administrators will result in savings of \$37,054 per administrator in the first year of ownership and will continue to accrue in subsequent years

"Oracle has clearly gone farther than any other vendor to instrument the database and provide expert advice to administrators."

Charlie Gary, Meta Group Source: META Group "Oracle 10g: Managing Reality", December 17, 2003

The primary factors responsible for Oracle's continued manageability superiority are self-tuning, simplification of complex management functions, increased automation of routine tasks, and more sophisticated management tools. Oracle's manageability framework assumed more of the work, and demanded less of the testers. DB2 V8.1 did make some improvements, particularly in the area of memory configuration, tuning, and health monitoring. DB2's new Health Center and its Configuration Advisor were primarily responsible for allowing DB2 to reduce Oracle's lead in time advantage (or efficiency). At the same time, Oracle has made further manageability enhancements in Oracle Database 10g, thus extending Oracle's lead even further.

Self-Managing Database

"With Oracle Database 10g, we can let the database manage itself and focus on issues that directly impact our business such as improving customer service. Rob Leaman, Manager, Business Information Management Systems Division at Deutsche Post IT Solutions GmbH (DPITS).

Oracle Database 10g's self-management infrastructure comprises of four components: Automatic Workload Repository, Automated Maintenance Task Infrastructure, Server Generated Alerts and Advisory Framework. The Automatic Workload Repository (AWR) is a built-in repository that contains information including operational statistics about that particular database. At regular intervals, the database takes a snapshot of all its vital statistics and workload information and stores them in AWR. The Automated Maintenance Tasks Infrastructure enables the database to automatically perform routine maintenance operations such as optimizer statistics refresh, rebuilding indexes, etc. It uses the rich scheduling functionality introduced in the Oracle Database 10g (through a new feature called Scheduler) to run such tasks in a pre-defined "Maintenance Window".

In DB2, it is the DBA's responsibility to decide which, how and when the statistics are to be refreshed, without any help from the database whatsoever. In other

According to DB2 Administration Guide: Performance "Ideally, you should rebind application programs after running statistics. The query optimizer might choose a different access plan if it has new statistics."³

words, A DB2 DBA has to manually track the tables that have changed since last statistics generation, be intimately familiar with data distribution to determine the columns needing histograms and, go through a trial-and-error process to determine the appropriate sample size for the data to be analyzed. That is not all. They must also rebind all the application packages in order to take advantage of new statistics. According to DB2 Administration Guide: Performance "*Ideally, you should rebind application programs after running statistics. The query optimizer might choose a different access plan if it has new statistics.*"³. When gathering statistics is complex as it is with DB2, it may deter DBAs not to do it as often as needed, resulting in poor execution plans. This perhaps is another reason why DB2 administrators rate query optimization as their most time-consuming task.

For the problems that cannot be resolved automatically and require administrators to be notified, such as running out of space, Oracle Database 10g has the ability to monitor itself and send out alerts to notify DBAs of any impending problem in an efficient and timely manner. Building upon the data captured in AWR, Oracle Database 10g includes a self-diagnostic engine called the Automatic Database Diagnostic Monitor (ADDM). ADDM makes it possible for the database to diagnose its own performance, determine how any identified problems could be resolved, and quantify expected benefits.

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

"One of the key benefits we've experienced while trialing Oracle Database 10g is better use of resources with the database's self managing features," said Suzanne Webb, project manager, NEMMCO. "This functionality will simplify and automate many of the daily and routine tasks associated with managing the database, allowing our IT professionals to put their talents to better and more strategic use. Other benefits include the improved performance and optimization of the database and resource management facilities."

Reorganizing tables or indexes in Oracle is also far less complex than in IBM DB2. Oracle administrators can analyze the database to identify objects that need to be reorganized and complete the reorganization from a single tool. DB2 requires that the use at least 2 utilities, REORCHK and REORG. Since reorganizations can be resource-intensive, they are often scheduled during non-peak hours. Using Enterprise Manager tools, reorganizations can easily be scheduled to run unattended during the maintenance window. Furthermore, EM automatically alerts the administrator when the reorganization starts, fails or succeeds. In DB2, any reorganization job is completely manual process requiring development of custom scripts and constant administrator supervision.

Oracle's reorganization process also has sophisticated error-handling capabilities. If an error occurs during reorganization, the administrator can either undo the operation or fix the problem and resume. When an error occurs during reorganization in DB2, changes will usually have to be rolled back or the database recovered.

DB2 still continues to lag behind Oracle significantly in all areas, including manageability. DB2 administrators still need to tune up to 12 parameters to

³ DB2 Administration Guide: Performance, Chapter 5, Page 120.

<http://ftp.software.ibm.com/ps/products/db2/info/vr8/pdf/letter/db2d3e80.pdf>

optimally configure SQL execution memory. Their Oracle counterparts, on the other hand, just specify the amount of query execution memory available to an instance using a single parameter and let Oracle manage it automatically to ensure the most optimal performance. Similarly, Oracle's Resumable Space Allocation feature provides a much-needed respite to administrators by allowing them to correct space related errors without risking the failure of any database operation. In contrast, DB2 administrators will continue to get woken up in the middle of the nights should they fail to accurately estimate the space requirements for routine administrative tasks like creating a new index, loading large volumes of data, etc.! Since DB2 V8, like the V7, continues to use a lock based read consistency model, they will also continue to spend a significant amount their time covering up for DB2's architectural deficiencies by monitoring deadlocks and resolving locking conflicts. Thanks to Oracle's widely applauded version based read consistency model, Oracle DBAs need not even think about these tasks! Most of the DB2 configuration parameters still can not be changed without shutting the database/instance down leading to application outages. Following are some of the other basic administrative requirements that have still not been answered by DB2:

- Dynamic configuration parameters
- Bounded Crash Recovery Time
- Recovery Time Advisor
- SQL Access to Data stored in external files
- Resumable Backup & Restore
- Automatic identification of resource intensive queries
- Automatic Identification of Tables Needing Optimizer Statistics Refresh
- Automatic Sample Size Determination for Optimizer Statistics Refresh
- Automatic Creation of Histograms
- Transportable Tablespaces
- Quality of Service Reporting

Application/SQL Tuning

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

The query optimizer makes crucial decisions which have a tremendous impact on the performance of a query such as, whether to use an index or not, which join techniques to use if the query involves joining multiple tables, etc. Oracle has invested considerable development effort in making their cost-based optimizer the

The most powerful aspect of SQL profiles is that they enable tuning of queries without requiring any syntactical changes and thereby providing Oracle administrators and customers with a unique database-resident solution to tune the SQL statements embedded in packaged applications

industry's most sophisticated, mature, and thoroughly tested query optimizer. The cost-based optimizer is used extensively by the major packaged applications such as Oracle Applications, SAP, PeopleSoft, etc. Since the Oracle database is the platform for the vast majority of the customers using these applications, this demonstrates the success of the Oracle optimizer in a huge number of real application settings.

While the Oracle Database provides best possible query optimization technology, which maximizes the application/query performance without any administrator intervention in majority of the cases, there may still be a few cases where the nature of the application or uniqueness of data distribution may cause certain SQL statements to consume unusually high percentage of total system resources. ADDM automatically identifies the top resource consuming SQL statements, and then Oracle Database10g can automatically analyze them and recommend solutions using newly added automatic tuning capability of the query optimizer, called the Automatic Tuning Optimizer. It is important to bear in mind here that the solution is coming right from the optimizer and not from any external tools using some pre-defined heuristics.

The Automatic Tuning Optimizer verifies its own estimates and collects auxiliary information to remove estimation errors. It also collects auxiliary information in the form of customized optimizer settings (e.g., first rows vs. all rows) based on past execution history of the SQL statement. It builds a SQL Profile using the auxiliary information and makes a recommendation to create it. When a SQL Profile is created it enables the query optimizer (under normal mode) to generate a well-tuned plan without any change to the application code. This unique feature is very useful for DBA's supporting 3rd Party Applications where they may not have access to the code, or may invalidate the support agreement if any code is changed.

Managing the Enterprise

With Oracle Enterprise Manager 10g Grid Control, Oracle provides a truly web-based management console allowing the DBA access from anywhere with zero client install and full firewall compatibility. The Performance Management capabilities include real-time graphical database performance monitoring through a web browser. To ease the administrative tasks of managing your Oracle software, Grid Control provides full software inventory capabilities including: hardware, operating system, operating system patches, Oracle software installed, Oracle patches installed. Grid Control tracks changes to hardware and software installations and configurations so it is easy to find "What changed?" if problems occur.

Through a direct connection to Oracle Support through Metalink, Grid Control provides the ability to quickly query for patches available for Oracle products installed across their enterprise. Patches can be found either in the context of a specific target or, if desired, the administrator can query for a specific patch. Once the necessary patch is located, Grid Control can download it from Oracle

"We now can manage hundreds of servers throughout the world from within one powerful management console."
Warwick Calkin, National Manager,
Telstra IT Infrastructure and Operations

Digex Inc. "Oracle Enterprise Manger enables us to manage literally hundreds of servers throughout the world with a small core set of resources in one location. It's an incredibly useful tool." -- Patrick Verhoeven, Manager, Product Manager Group

"Oracle Enterprise Manager 10g Grid Control will open the way for us to increase market share and sustain our competitive edge by meeting customer demands for lower cost and higher service levels--without increasing our overhead." -- John Stevenson, Vice-President, Managed Operations, Business Applications, and Communication Services, Atos Origin

MetaLink and stage it to the appropriate target hosts. Optionally, Grid Control can execute an end-user provided script to install the patch.

Duplicating software installations for development or QA purposes is a routine task in many data centers. For Oracle software, such installation cloning has become a lot easier with Grid Control. Grid Control's cloning wizard automates the duplication of database and application server installations (specifically, the directories where the corresponding Oracle homes reside). Thanks to its "multicasting" capability even multiple clones on multiple target hosts can easily be created in a single operation. Grid Control's home cloning is intelligent: host names, IP addresses and other environment dependent settings are automatically adjusted for the newly cloned home.

From one terminal, DB2 Control Center provides a common, integrated toolset for managing local and remote databases across different software and client hardware platforms. This capability although new to DB2 was introduced in Oracle way back in 1994 as a part of Oracle Enterprise Manager framework. In DB2 Version 8.1, you also have the option of using the Web Health Center to access Health Monitor information and execute new DB2 commands from a Web browser.

HIGH AVAILABILITY

Building a highly available IT infrastructure is critical to the success and well being of all enterprises in today's fast moving economy. Should a critical application, server or data become unavailable, the entire business can be placed in jeopardy. Revenue and customers can be lost, penalties can be owed, and bad press can have a lasting effect on customers and a company's reputation. It is not always easy to place a direct cost on downtime. According to The Standish Group's DARTS⁴ studies, one minute of system downtime can cost an organization anywhere from \$2,500 to \$10,000 per minute. Using that metric, even 99.9 percent data availability can cost a company nearly \$5 million dollars a year. It is important to consider causes of both unplanned and planned downtime when designing a fault tolerant and resilient IT infrastructure. Unplanned downtime is primarily the result of computer failures or data failures (corruption, human error). Planned downtime is primarily due to data changes or system changes that must be applied to the production system. Planned downtime can be just as disruptive to operations, especially in global enterprises that support users in multiple time zones. In this case it is important to design a system to minimize planned interruptions. Planned downtime includes routine operations, periodic maintenance, and new deployments.

⁴ DARTS = Demand Assessment Requirements Tracking Study, an ongoing study by the Standish Group

Bounding Database Crash Recovery

Both Oracle and DB2 can perform basic online and offline backup and recovery. Though it is feasible to determine and implement a backup plan in advance, it is difficult to anticipate all recovery scenarios. Oracle's comprehensive backup and recovery capabilities extend beyond the basic functionality provided by DB2. As a result, Oracle can handle almost any backup and recovery requirement.

Oracle Database 10g Recovery Manager (RMAN) feature set revolutionizes the recovery of critical Oracle data out of the box. At no extra cost or additional install, RMAN manages the backup and recovery of the Oracle database files. Since RMAN is tightly integrated with the Oracle kernel, it provides the insight and wisdom to efficiently recover your Oracle databases. Oracle Database 10g RMAN provides new advanced features such as

- Flash Recovery Area – organizes and manages recovery related files.
- Optimized Incremental Backups – directly backs up only the changed blocks for a backup.
- Incremental Updated Backups – merges database image copies and incremental backups to provide fast and efficient database recovery.
- Enterprise Manager and RMAN – provides point and click backup and recovery wizards

Oracle Database 10g backup and recovery also includes many other innovative capabilities including:

- Compression of backups.
- Automated failover to a previous backup when restore discovers a missing or corrupt backup.
- Automated recovery through a previous point in time recovery – recovery through resetlogs.
- Automated creation of new files during recovery.
- Automated channel failover on backup or restore.
- Automated tablespace point-in-time recovery.
- Full DB “begin backup” command for faster mirror split.
- Improved Recovery Parallelism (2 to 4 times).
- Tablespace Rename.
- Proxy (third-party) Backup for archive logs.
- Time window based throttling of backups.
- Cross Platform Transportable Tablespaces.

Split mirror backups are useful because they produce instant backups. Both Oracle and DB2 provide facilities for split mirror backups. However, Oracle can split a mirror while the database is running and writing to the disks. DB2 has to suspend database I/O to perform a mirror split, thus making the database unavailable for writers during this operation.

Should archived log files become damaged. Oracle allows damaged log files to be scavenged using the LogMiner utility, thus recovering some of the transactions recorded in the log files. With DB2, a corrupt archived log file means loss of all transactions in that particular log file plus any archived log files created after the damaged log file.

With Oracle's block-level media recovery feature, if only a single block is damaged then only that block needs to be recovered, the rest of the file, and thus the table containing the block, remains online and accessible, increasing data availability. DB2 cannot recover data in single block units, thus requiring the entire file to be taken offline, restored, and recovered.

Human Error Recovery

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

The Oracle Database 10g provides a family of human error correction technology called Flashback. Flashback provides a SQL interface to quickly analyze and repair human errors. Flashback is unique to the Oracle Database 10g and supports recovery at any level of granularity including the row, transaction, table, and database wide.

- **Flashback Query** allows the user to query data at some point-in-time in the past to reconstruct lost data that may have been deleted or changed by accident.
- **Flashback Version Query** provides a mechanism to view changes made to the database over time at the row level.
- **Flashback Transaction Query** provides a mechanism to view changes made to the database at the transaction level.
- **Flashback Database** is a new strategy for doing point-in-time recovery. It quickly rewinds an Oracle database to a previous time to correct any problems caused by logical data corruption or user error.
- **Flashback Table** provides the ability to recover a table, or a set of tables, to a specified point in time in the past very quickly and easily.
- **Flashback Drop** provides a safety net when dropping objects as you can very quickly and easily undrop a table and its dependent objects.

Incremental Disk Backups and Flashback technology eliminates the need for Split Mirror backups. DBA's no longer have to go through a lengthy restore from backups to correct errors and backups are significantly faster without expensive disk mirroring technologies.

Online Maintenance

In an ideal world, a DBA should be able to perform all administrative actions online, i.e., all business transactions can continue while the operation takes place. Oracle not only performs more operations online than DB2 but also makes their execution significantly more manageable. Virtually every reorganization operation can be performed online in Oracle, be it creating a new index, coalescing or moving an existing index or, making any kind of changes to a table including adding and dropping columns. Of these, DB2 can only coalesce an index online. All other operations in DB2 require transactions stopped on the database and all locks released⁵

The Oracle Database supports many maintenance operations without disrupting database operations or users updating or accessing data. Indexes can be added, rebuilt, or defragmented while the database is open and end users are reading or updating data. Similarly, tables can be relocated or defragmented while online. Tables can be redefined, changing table types, adding, dropping or renaming columns, and changing storage parameters without interruption to end-users who are viewing or updating the underlying data. In the Oracle Database 10g this capability has been enhanced to:

- support the easy cloning of indexes, grants, constraints, and other characteristics of the table
- convert from the long to LOB datatype online; and
- allow unique indexes instead of requiring a primary key.

Java™ and PL/SQL™ stored procedures can be updated dynamically and Oracle manages all dependencies, to properly integrate the new procedures into the database, with no impact to end-user operations. In the Oracle Database 10g this capability has been enhanced to allow many kinds of tables changes to be made without recompiling stored procedures associated with that table.

Data Centre Disasters

Oracle Data Guard, which offers the most complete and robust disaster recovery solution in the Industry, provides:

- Protection from Human Error Corruptions and Disasters

5 DB2 Command Reference, REORG command, page 470.

Data Guard customer Kevin Widdrington, IT Manager at Costcutter, explained: "The new system offers us a fail-safe solution to systems crashes when downtime can cause serious disruption. We estimate that we will now be able to recover from a system failure within a matter of minutes."

"Before we made the decision to use Oracle's standby database, we also looked at many other options. We needed to consider the safekeeping of our data, but we also needed to look at cost. Oracle Data Guard provides everything for a high availability solution at a lower cost than other alternatives." - Ann Collins, Technical Director First American Real Estate Solutions

"Oracle's integrated data intensive functionality for business intelligence is clearly the right approach. Data intensive operations such as OLAP, ETL and data mining, are best performed closest to the data, within the database, instead of moving large volumes of data to multiple servers. This approach will reduce integration costs, improve manageability by reducing complexity, and allow for advanced functionality without specialized servers and skills." Richard Winter, president, Winter Corporation

Oracle simply provides the best Business Intelligence suite on the market, both from a development perspective as well as from a result perspective." Mikael Angberg, Business Intelligence Manager, Internet Development Services CERN

- Zero Data Loss Protection
- Near Real Time Data Synchronization
- Integrated GUI-Based Management Framework

Oracle Data Guard empowers customers to survive disasters of many forms. Data Guard automates the complex tasks and provides the monitoring alerting and control mechanisms to maintain a standby operation. Plus, Data Guard reduces planned downtime by utilizing the standby server for maintenance, upgrades (hardware, OS, and database software), and routine operations in addition to reporting.

IBM does not offer a solution comparable to Oracle Data Guard. With DB2, every standby database is a custom job; tasks as basic as shipping redo logs to the standby site depends on user written log transfer callouts. DB2's user created standby database "solution" is less robust and costs more to implement.

DATA WAREHOUSING AND BUSINESS INTELLIGENCE

Good business intelligence operations are built on data repositories that handle geometrically increasing information requests because of a vastly expanded audience. It also provides customers instant information regardless of where the data resides. Whether a business is small, medium or large, brick and mortar or dotcom, business intelligence systems will consistently encounter these two requirements. In fact, AMR Research has determined that the challenges of providing a growing population of users with instant information, supported invisibly by large amounts and sources of data, are driving the consolidation of data stores into larger data warehouses at a growth rate of 40% per year. This means that data warehouses will store larger volumes of data, support more users, and require faster performance—without losing quality of service. The core requirements for data warehousing—performance, scalability and manageability—remain key factors in the successful implementation of data warehouses.

Oracle Database 10g is a single, standards-based platform that addresses all of the server-side business intelligence and data warehousing requirements, including Extraction, Transformation and Loading (ETL), Online Analytical Processing (OLAP) and data mining. Oracle Database's capabilities eliminate the need for multiple engines in the business intelligence environment. This business intelligence architecture provides numerous benefits:

- Rapid deployment: There is no longer any need to integrate multiple server components when implementing a business intelligence system.
- Reduced management costs: A single server means that all data is stored in one place, and is managed via a single tool. Moreover, there is no longer a need to manage the process of synchronizing the data stored in various servers; once new data is loaded into Oracle Database 10g, it is accessible for all business intelligence needs.

"Oracle Database was first choice of T-Mobile Austria. It has successfully proved its scalability and hardware portability. Today we can rely on our vital Enterprise Data Warehouse covering all aspects of the value chain at T-Mobile Austria."
Marcus Berger, CIO, T-Mobile Austria

- Better decision-making: With a single data store, all data in the business intelligence system is available to all business intelligence users more rapidly, since there is no need to percolate data from one server to another.
- Security: With all data in a single database, IT managers are assured of having consistent security for all business intelligence data regardless of how it is accessed.
- Availability: Oracle Database 10g provides robust availability capabilities, which far exceed the availability characteristics of any other business intelligence products.

Scalability

Oracle has long supported a variety of hardware platforms and operating systems, giving customers choice and negotiating leverage. Oracle supports the most scalable 64-bit Symmetric Multiprocessing (SMP) systems and Non-Uniform Memory Access (NUMA) systems today and has supported systems of similar architecture since the early to mid 1990s. Massively Parallel Processing (MPP) systems, supported by Oracle and IBM, require more care and feeding due to multiple operating systems and database instances to manage. The IBM SP is the MPP platform targeted by IBM DB2 UDB ESE with Partitioning (formerly EEE). The IBM solution relies on hash partitioning to provide a theoretically even distribution of data across the MPP complex. This introduces MPP management and tuning considerations to minimize data skew during real business queries largely mitigated in Oracle Real Application Clusters' combination shared nothing / shared disk approach.

IBM has delivered a type of clustering technology for some time, but it just doesn't compare to Oracle's solution.

<http://www.eweek.com/article2/0,1759,1549134,00.asp>

IBM's VLDB examples are also overwhelmingly on IBM platforms, even though IBM DB2 is offered on some other hardware platforms (HP, Sun, Linux, Windows), there is little proof of its performance and scalability on those platforms.

A few examples of Oracle multi-Terabyte (of data) customers include:

- Amazon.com with 13 TB of data on HP Superdome.
- Acxiom with 6 TB of data on HP AlphaServer cluster.
- AtosEuronext with 1.5 TB of data on RedHat Linux cluster.
- Best Buy with 1.5 TB of data on Sun E-10000.
- Colgate Palmolive with 2.6 TB of data in SAP BW on IBM p690.
- Financial Institution with 5 TB of data on multi-node 206 CPU IBM SP.
- France Telecom with 21+ TB of data on HP V-2500.
- Telecom Italia Mobile with 9.6 TB of data on HP AlphaServer cluster.

Parallelism

At the heart of Oracle's scalable, parallel decision support strategy lays a dynamic parallelism approach. This approach allows for a straightforward transparent path to parallelism without needing static table partitions. As a result, data management tasks are significantly reduced and hardware is utilized to its full potential. Where Oracle's Partitioning Option is desired for data maintenance, Oracle additionally supports parallelism across partitions including parallel DML. Oracle also supports parallelism within each partition for queries and inserts.

By comparison, IBM (DB2 UDB EEE) chose initially to use partitioning as a means to enable parallelism. IBM proposes a "UNION ALL" workaround. Interestingly, IBM chose not to take this approach for DB2 on OS/390 (Z/OS) where the approach is shared disk for the Sysplex and where partitioning (range) is used for management.

The optimizer in Oracle8i added recognition of "materialized views", a hierarchy of summary tables in the RDBMS to which queries can be transparently redirected to a summary level resulting in much better query performance. Oracle implemented this capability for fact and dimension tables. IBM took a similar approach with DB2 UDB, but only supports fact tables (since Version 7).

The costs computed by Oracle's cost based optimizer for queries can now be leveraged in Oracle's Database Resource Manager. Cost limits can be assigned to groups of users (or "query consumers") preventing queries that are not well thought out from interfering with the queries from other users who also desire performance. IBM has taken a tools based approach external to the database to solve this problem.

It is worth noting that Oracle also fully embeds an OLAP Option and Data Mining Option in the database. This enables very sophisticated analysis where the data resides instead of within a business intelligence tool and can greatly improve performance. The OLAP Option is accessible via SQL or a Java API. The Data Mining Option is accessible via a Java API.

Partitioning

Partitioning allows large database structures (tables, indexes, etc.) to be decomposed into smaller and more manageable pieces and, at the same time, improves query performance and resource utilization.

Feature	Oracle	DB2
Range partitioning	Yes	–
List partitioning	Yes	–
Hash partitioning	Yes	Yes
Composite partitioning	Yes	–

"Mark Dunlap, who manages the data warehouse program at Amazon.com, credits Oracle composite partitioning with contributing in a major way to performance. "

from Field experience with large scale Data Warehousing on Oracle by Richard Winter, Winter Corporation

Local index	Yes	Yes
Global partitioned index	Yes	–
Global non-partitioned index	Yes	–

Table 2: Partitioning options.

DB2 only supports the hash partitioning method, which has considerable limitations and weaknesses when compared to Oracle’s partitioning capabilities. By only supporting hash partitioning, DB2 does not allow for ‘rolling window’ support – a great drawback in manageability. A ‘Rolling window’ allows the data for a defined period, lets say a week or a month, to be kept online by continuously replacing the oldest data with the most recent one. DB2’s hash partitioning scheme requires data in all partitions to be redistributed whenever new data is loaded, therefore increasing the load time and decreasing data availability as the table is locked during the data redistribution process. Similarly, when the old data is archived or deleted, all partitions need to be touched. This may interfere with regular insert operations and cause space fragmentation.

Another weakness of DB2 is that it requires equi-partitioning between tables and indexes. This means that global indexes, partitioned or non-partitioned, cannot be created. This is a major problem in OLTP environments where global indexes are commonly used to offer efficient access to any individual record. With DB2, application designers have no flexibility when defining their indexing strategy in partitioned configurations.

All these limitations in DB2 partitioning options make the task of managing large volume of data quite complicated. Oracle offers far more options in this area that not only enhance manageability but also improve performance.

Data Loading and Archiving

Oracle Database 10g introduces Oracle Data Pump, a unified server-based framework for high speed, bulk data and metadata movement between Oracle database systems. Leveraging the direct path API, Oracle Data Pump provides the fastest mechanism to load and unload data from and into Oracle systems. Oracle Data Pump is an enterprise level server infrastructure that provides full parallel capabilities for load and unload, restartability, and monitoring capabilities. All interfaces to invoke Oracle Data Pump are externalized, so that you can write your own data movement utilities. With Data Pump Export and Import, any subset of database objects can be unloaded and reloaded on the target platform. If moving across platforms with different file system syntax, it is now possible to change the definition of some objects as they are created at import time, without losing any data

DB2 export allows for exporting data from only one table at a time. So, if a schema has, lets say, 20 tables, one needs to run 20 export jobs, one for each table. Besides being more time and labor intensive, this can also lead to inconsistencies in the

Ian Crocker, Performance & Storage Consultant , British Telecom comments on Oracle Database 10g Beta “We have tested Oracle Data Pump, the new Oracle10g Export and Import Utilities. Data Pump Export performed twice as fast as Original Export, and Data Pump Import performed ten times faster than Original Import. The new manageability features should give us much greater flexibility in monitoring job status”

exported data since tables with dependencies may have been exported at different times. So, if the schema contains two table ORDERS and ORDER_ITEMS, the exported data for the ORDER_ITEMS may contain entries for orders not present in the export file for ORDERS table since the ORDER_ITEMS table may have been exported after ORDERS table. Such inconsistencies may render the exported data completely useless since it cannot be used to recreate a consistent snapshot of data. Oracle, on the other hand, not only allows administrators to export data at a schema, tablespace and database level but can also ensure that the exported data for various tables is consistent to the point in time when the export job was started, irrespective of the time when a given table was exported. Therefore, in the context of the above example, the administrator needs to execute only one job to export the complete schema and does not have to worry about any inconsistencies in the exported data.

Secondly, since the DB2 export file does not extract certain table and schema attributes, it can never be used to logically recreate a schema or a database. According to “DB2 Data Movement Guide and Reference”⁶, DB2 EXPORT does not support tables with typed columns and does not retain attributes such as referential integrity constraints, check constraints, physical space allocation settings, column default values, foreign key definitions and triggers. It also does not support any non-table database objects such as views, stored procedures, etc. Therefore, unlike Oracle Export/Import utility that is frequently used to perform a logical backup and restore of a schema or a database, the DB2 tools are designed only for moving data in and out of a table. But even this may not be possible at times since DB2 import cannot be used to re-create tables with LOB columns defined to be 1 GB or larger⁷

Oracle SQL Loader is another powerful tool to load data from external files into Oracle database tables. Unlike DB2 Load, SQL Loader allows for loading data into multiple tables during the same load session, makes it possible to manipulate data using SQL functions and, leaves table data accessible after the load process completes. This is not true for DB2 unless the administrator has backed up the tablespaces and enabled the integrity constraints⁸. But going beyond simple data loading, Oracle Database 10g provides a comprehensive set of server functionality to address the typical requirements of ETL (Extraction, Transformation, Loading) processes. Using Oracle’s external table feature, data in external sources such as flat files, can be exposed within the database just like a regular database table. These “External Tables” can be accessed via SQL, so that external data can be queried directly and loaded into the database in parallel using the full power of SQL, PL/SQL, and Java. External tables significantly simplify the ETL process by

6 Chapter 2, Page 32,33,67

⁷ DB2 Data Movement Guide and Reference, Chapter 1, Page 4

⁸ Fast Path to DB2 for Experienced Relational DBAs CBT, Load Considerations, Slide 9

combining data-transformations and loading into a single step; something that will require several manual steps with DB2. With DB2, the external data first needs to be loaded into a “staging” table before any SQL based transformation and processing can be applied. The net result, more work for the DBA and a longer ETL cycle!

Operational Impact

In many organizations, one Oracle DBA typically manages several databases. For large-scale decision support, the following implementations provide an Oracle proof point regarding how few DBAs are often actually needed:

- Acxiom, 16 TB database, 2 DBAs
- Acxiom, 6 TB database (RAC), 2 DBAs
- Amazon.com, 16 TB database, 2 DBAs
- France Telecom, 29+ TB database, 2 DBAs
- Telecom Italia Mobile, 12 TB database (OPS), 3 DBAs
- WestPac, 2.3 TB database, 2 DBAs

Another consideration regarding impact on operations is the availability of skills for the technologies being considered. Since Oracle appears in far more data warehousing / decision support implementations (most analysts say about 30-50 percent), far more Oracle skilled consultants and designers exist. Oracle customers often leverage their own internal skills for such implementations. Those looking for consulting assistance can find it at Oracle, at Big Four consulting companies, and at many specialized second-tier consulting companies. As a result, Oracle project managers can leverage a wide choice of skilled resources at a variety of price levels. For example, a recent search of Oracle skills on “Monster.com” showed over 5000 entrants for individuals with Oracle skills but a much smaller number (<200) listing “DB2 UDB”.

APPLICATION DEVELOPMENT

Enterprises everywhere are developing Internet-based applications to take advantage of the ubiquity, affordability and the open standards of the Internet. However developing these applications is very challenging, placing numerous additional requirements on the development platform. Choosing the right development platform is critical to ensuring the success of these applications. Oracle Database 10g provides the most comprehensive functionality for developing versatile, scalable, and high performance database applications for Grid computing.

Surveys indicate that most major packaged application vendors support Oracle databases first. This is true across all categories, including ERP, CRM, Procurement and Supply Chain. Due to the larger Oracle installed base, there are

significantly more implementers with experience deploying packaged applications on Oracle.

Most applications involve some form of text manipulation, and these applications can benefit from the introduction of native regular expression support in SQL and PL/SQL, which revolutionizes the ability to search for, and manipulate, text within the database. With DB2, Regular expression support is provided by the PCRE library package, which is open source software, external to the database.

ANSI SQL multiset operations allow data mining applications to process market basket information natively in the database. Oracle now supports SQL Multiset operations on Nested Tables. DB2 doesn't support collection types at all!

With Oracle Database 10g, there are 2 new sorting options a case insensitive sort and an accent insensitive sort. This enables applications to easily handle data such as name and address, which may be store with accents and mixed case. IBM DB2 does not support case insensitive operations.

Staying up to date with the latest Unicode standard, Oracle character sets ATL32UTF8 and AL16UTF16 (UTF-8 and UTF-16) now support Unicode 3.2 with 4-byte supplementary characters. IBM DB2 supports Unicode 3.0 but does not have supplementary character support.

Large Objects (LOBs) are used to store unstructured data such as video, audio, and documents. Oracle Database 10g can store and retrieve a maximum of 128 Terabyte LOBs. IBM DB2 can store a maximum of 2 GB LOBs

The SQL model clause in Oracle Database 10g allows query results to be treated as sets of multidimensional arrays upon which sophisticated interdependent formulas are built. These formulas can be used in complex number-crunching applications such as budgets and forecasts without the need to extract the data to a spreadsheet or perform complex joins and unions.

For Microsoft .Net Developers, Oracle provides a customized, high performance Oracle Data Provider for .Net (ODP.Net), which is compliant with the ADO.Net Standard. ODP.Net supports native Oracle datatypes such as REF cursors and LOBS. DB2 has only a developmental release of their data provider for DB2, which supplies only a generic interface similar to ODBC.Net.

Web services are pervading every layer of enterprise computing, from packaged e-business applications (e.g., ERP, CRM) to middle-tier (e.g., J2EE, .NET) and database infrastructure. The proliferation of structured and unstructured data and data logic in databases, the increasing momentum of XML as a data exchange format, the de facto acceptance of HTTP as a ubiquitous transport mechanism -in the context of heterogeneous environments - has aroused interest in database Web services. Database Web services work in two directions: database as service provider, i.e. calling from the outside in, which lets client applications access the database via Web services mechanisms; and database as service consumer, i.e. calling from the inside out, which lets a SQL query or an application module in a

database session consume an external Web Service. Oracle allows you to turn your database into a Web service provider, thus enabling you to share data and metadata across corporate intranets and access the database operations, e.g. triggers, through SOAP requests. Also with Oracle, you can turn your database into a Web services consumer to access dynamic data. This has been available since Oracle 8i and with Oracle Database 10g has been formalized by including the new capabilities in JPublisher to support the database as Web services provider (SQL query, SQL DML, Java-in-the-database) and JPublisher support for the database as Web services consumer. DB2 has delivered web services support in Version 8 for both provider and consumer however since DB2 does not have an integrated JavaVM, they rely on either hand-crafted static SOAP messages or a Non-SOAP aware procedural library.

Web Based Development Tool

Oracle HTML DB is a browser based web development and deployment framework unique to Oracle Database 10g . It combines the ease of use and productivity of a personal database with the availability, scalability and security of an enterprise database. Oracle HTML DB is a declarative development tool and a framework for the development and deployment of database-centric web applications. Oracle HTML DB accelerates application development through built in features such as design themes, navigational controls, form handlers and flexible reports. Using only a web browser, you can quickly assemble a sophisticated database driven web application.

Many organizations are wasting valuable time by using spreadsheets and personal databases to manage information. Although easy to use, these products don't deploy well on the web nor do they scale up to handle multi-user updates to data gracefully. By consolidating lots of small workgroup databases, including spreadsheets, into a single Oracle database without sacrificing developer control, Oracle HTML DB provides the following benefits:

- Improved access by deploying data currently locked in spreadsheets and personal databases to the web, allowing concurrent updates by multiple users
- Time saved by avoiding the emailing around of spreadsheets as attachments and the merging and scrubbing of multiple copies of data
- Improved security and availability through consolidation of fragmented information into a reliable, secure, scalable Oracle database
- Build applications you never could before – even “quick and dirty” applications are immediately deployed on the web ready for an enterprise audience

Globalization

Oracle Globalization Development Kit (GDK) consists of a set of Java APIs that provide Oracle application developers with the framework to develop globalized Internet applications using the best globalization practices and features designed by Oracle.

DB2 does not ship with any Java globalization APIs, however IBM do have a Globalization library for Java called ICU(International Components for Unicode), this is open source and it provides a subset of the functionality delivered with GDK. Similar to GDK, ICU provides low level APIs for linguistic sorting, number formatting, supplementary characters support etc. Unlike GDK, it does not offer high level APIs for building global internet applications, such as User locale detection, pre-translated locale names, global application configuration file, language and character set detection, automatic redirection and fall back of localized contents etc.

MANAGING ALL YOUR DATA

Unstructured Data

Structured data represents only part of the critical data in an organization. Spreadsheets, word processing documents, video clips, newspaper articles, press releases and geographic descriptions are only a few of the unstructured objects that may need to be integrated and managed in a central repository. Oracle Database 10g builds on it's existing superiority in managing all your data – relational, email, documents, multimedia, XML and spatial – in an integrated fashion. These offerings include:

- *interMedia*
 - Image extends ability to store and retrieve images
 - Audio extends ability to store and retrieve audio clips
 - Video extends ability to store and retrieve video clips
 - Locator extends ability to retrieve data linked to coordinates
- Text
 - Text extends ability to retrieve documents and the gist of documents
- Native XML storage
- Spatial Data Option
 - Provides means to link data to coordinates
 - Typically leveraged by partner vendors of Geographic Information Systems (GIS)

IBM's approach has been to offer “extenders” to handle many of these data types. IBM stores all media and its meta data in separate tables from the associated

relational data thus requiring greater processing and I/O overhead for media access and retrieval and greater administrative overhead. Oracle Database 10g is automatically media enabled out of the box where as IBM DB2 requires the DBA to manually initiate separate processes to enable a DB2 tablespace to accept media data. Greater processing overhead is due to IBM using a trigger on each and every media object creation to update the separate administration tables that contain media objects and metadata. If using Java Object Methods with media datatypes, IBM does not provide Java class libraries for media access and they do not provide a JSP Tag library for JSP application developers. Customers must code their own. More code to develop and maintain which makes applications cost more to develop and maintain.

XML

Oracle XML DB is a high-performance XML storage and retrieval technology available starting with the Oracle9i Release 2 Database. It fully absorbs the W3C XML data model into the Oracle server, and provides new standard access methods for navigating and querying XML. With XML DB, you get all the advantages of relational database technology and XML technology at the same time. XML DB also adds an XML Repository to the database. Oracle Database 10g Intermedia is now integrated with XML DB making it possible for XML documents that contain media such as images to have metadata for those images parsed. This makes it possible to index and easily locate both the image and parent XML document.

IBM's database XML technology is supplied as the DB2 v8.1 XML extender, a loosely coupled plug-in that does not afford the benefits of Oracle's native XML support. IBM introduced XML-wrapped data integration technology with their Information Integrator 8.1. Unlike Oracle, DB2 has to be 'extended' to support XML by separate pieces of software called the XML Extender and the Text Extender. These three pieces are developed separately by different organizations worldwide and are poorly integrated, and often inconsistent with each other. Users of such systems not only have to contend with the manageability-headaches arising from poor integration, but also the loose coupling, which proves to be less functional and less performant.

Oracle's XML DB technology is superior to IBM's in terms of

- Native support for the XML data model and other XML standards
- Integrated SQL, XML, Java and Text processing that eliminates extra moving parts and provides high scalability
- Superior functionality and richness of programmatic access
- A comprehensive XML vision, including an XML Repository to support content-oriented XML.

"The Oracle Text XML search facilities, combined with our topic hierarchy, also defined in XML, enable us to quickly locate relevant and reusable questions when creating new surveys. Storing participant answers in XML format enables us to easily transform the survey results and import them into other research systems for further analysis, in particular our Clinical Genome Miner system," says Gudbjartsson.
http://otn.oracle.com/oramag/oracle/04-mar/o24xml_feature.html

A report from The Butler Group comparing the major XML data management technologies concluded: *“IBM DB2 comes out disappointingly ... overall, it does not support quite so many features as Oracle’s 9i ... Before commencing this Report, given the expertise of the people involved, an expectation was formed as to how the scoring would come out ... The one possibility that was dreaded was that there would be a tie for first place between Oracle and IBM. Not only was this not the case, but the margin of difference was so great, which came as a considerable surprise ...”*⁹

“With Oracle, we have an overall cost structure that is far more advantageous than what we had before. We are able to manage the corporate system with minimal staff, and all of our e-business capabilities are now faster, less expensive, and every bit as reliable for less money.” --Roger K. Mowen, Senior Vice President of Global Customer Services Group and CIO, Eastman Chemical Company

CONCLUSION

Oracle Database 10g is the first database designed for Enterprise Grid computing. Oracle Database 10g continues to provide industry leading performance, scalability, resource utilization, manageability, availability, and consistency with emerging trends in technology. Oracle makes GRID the solution for your enterprise.

The Oracle Database, backed by 25 years of technical innovation and experience, remains the number one database choice in the market. With the Oracle Database, not only can immediate cost savings be realized, but the investment is also guaranteed for the future.

⁹ “Database Management Systems -- Managing Relational and XML Data Structures” ISBN 0-9542845-2-6 Published by Butler Direct Ltd., The Butler Group Sep 2002. <http://www.butlergroup.com/>



Technical Overview of Oracle Database 10g vs IBM DB2 UDB V8.1
March 2004

Author: Barb Lundhild

Contributors: Sandra Cheevers, Robert Stackowiak, Penny Avril

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
www.oracle.com

Copyright © 2004, Oracle. All rights reserved.

This document is provided for information purposes only
and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to
any other warranties or conditions, whether expressed orally
or implied in law, including implied warranties and conditions of
merchantability or fitness for a particular purpose. We specifically
disclaim any liability with respect to this document and no
contractual obligations are formed either directly or indirectly
by this document. This document may not be reproduced or
transmitted in any form or by any means, electronic or mechanical,
for any purpose, without our prior written permission.

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