

## ORACLE DATABASE 10g: A REVOLUTION IN DATABASE TECHNOLOGY

*Andy Mendelsohn, Oracle Corporation*

### EXECUTIVE OVERVIEW

Oracle Database 10g is the first database designed for Enterprise Grid Computing. Oracle Database 10g cuts costs 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.

### IT CHALLENGES

For a variety of reasons, most enterprises today have dedicated storage and servers for each major application system. By segregating systems on their own hardware and software, it has traditionally been easier to deploy an individual system. However, the result of this practice in an enterprise is an enormous number of systems to maintain, which is more costly as well as more prone to failure, due to the increased failure points.

In addition, individual systems get sized for peak capacity, and it is difficult to move resources from one system to another. The result is that enterprises today often have underutilized storage and CPU. Analysts have described typical storage utilization at 50% and CPU utilization at 15-20%.

Underutilization is exacerbated by continual growth in demand. The end result is not only higher costs, but an increasingly complex infrastructure, which diminishes the ability of IT to respond to the rapidly changing needs of the business.

### THE SOLUTION

The solution to these problems is Enterprise Grid Computing. The focus of Enterprise Grid Computing is to create large pools of industry standard, modular storage and servers. With this architecture, each new system may be rapidly provisioned from the pool of components.

There is also no need to size systems for peak workloads, as capacity can be easily added or reallocated from the resource pools as needed.

Enterprise Grid Computing drastically cuts costs in a number of ways. It lowers hardware costs by leveraging low cost components and by greatly increasing resource utilization levels. There is great potential to lower labor costs as each system can be built following the same blueprint, and centralized management and management tools reduce the burden of maintenance and monitoring. With this greater efficiency, IT can rapidly respond to the needs of the business, and overall risk is greatly reduced.

### EFFICIENT STORAGE FOR ALL YOUR DATA

For databases, storage is one of the key resources. In today's standard environment, the process of laying out data across the disk storage for maximum throughput can be a time consuming process, which is repeated each time disks are added or removed from the storage area for the database.

Oracle Database 10g's new Automatic Storage Management (ASM) feature addresses these challenges. ASM allows the database to directly manage raw disks. It eliminates the need for a file system and volume manager to manage both data files and log files. ASM automatically stripes all database data across all disks, providing the highest I/O throughput without any management costs. As disks are added or dropped, ASM maintains striping automatically. Data availability is increased with optional mirroring and disks can be added and dropped online. In addition, the possibility of data loss due to human error is removed, since ASM handles the removal of files that are no longer part of the database. ASM even automatically detects hot spots and rebalances data to alleviate them. ASM saves money, since it works well with low cost, modular storage. ASM makes storage underneath the Oracle Database 10g fast,

inexpensive, and easy to manage.

Whilst ASM manages all the disks, Oracle Database 10g manages all your data – relational, email, documents, multimedia, XML and spatial. The XML capabilities of Oracle10g have been expanded to focus on the effective management of large amounts of XML data, which can be stored natively in the database with enterprise class performance and scalability. There is support for additional native interMedia types, and the 4GB restriction on LOBs stored in the database is removed, raising the limit to 128 Terabytes and enabling the storage of larger multi-media documents.

With the consolidation of databases that comes with an Enterprise Grid Computing environment, there is a growing requirement to support larger and larger databases. Oracle Database 10g meets and exceeds any capacity requirements with support for Ultra Large Databases (ULDB), which can be as large as millions of terabytes (exabytes). The database size limit has been raised by the introduction of ultra large files (terabytes in a single file), unlimited size LOB columns (terabytes) and by the enhanced manageability functions of Automated Storage Management, automated space management and the new recovery area.

To assist in migrating to Enterprise Grid Computing, Oracle Database 10g enables very high-speed movement of data from one database to another across platforms and database versions. For highest performance, transportable tablespaces move data at the datafile level for fast “plug-and-play”. Transportable tablespaces can now operate across platforms. Among other things, this allows databases to be migrated to a new platform at the speed of file transfer. Oracle Data Pump is a new feature of Oracle Database 10g that provides high speed, parallel, bulk data and metadata movement of Oracle database contents across platforms and database versions. Data Pump jobs can run directly across the network or via disk and if stopped for any reason, can be restarted at a later point in time, without losing any data. Even single thread performance of import and export performance has been significantly enhanced (1.5 to 10x) over previous versions and can be automatically scaled up with the number of available CPUs.

### **OPTIMAL RESOURCE UTILIZATION**

As mentioned above, the general approach in Enterprise Grid Computing is to create large pools of servers and then to dynamically provision the servers to systems as needed. From a strict technical viewpoint, provisioning might not result in performance gains. However, from a business perspective, provisioning will give users better performance. With the same resources, users can get more performance as resources can be provisioned to the right application based on business priorities or needs.

Oracle Real Application Clusters (RAC) is unique in its ability to provision servers since it is the only database technology that supports the ability to run and scale all application workloads on a cluster of servers. Oracle Database 10g makes managing a cluster database easier with two major new capabilities: Automatic Service Provisioning and integrated clusterware. Automatic Service Provisioning offers hands-free allocation and re-allocation of servers to workloads (services). Clients login to services and are automatically routed to the appropriate server with the lowest load. On server failure, surviving servers are re-allocated to services. Integrated clusterware makes RAC easy by eliminating the need to purchase, install, configure, and support third-party clusterware. Servers can be easily added and dropped to a RAC cluster with no downtime. RAC is the only database technology to truly support this flexible capability, which is essential in implementing Enterprise Grid Computing. Real Application Clusters are deployed in over 500 production systems, as of August, 2003.

### **HIGHEST QUALITY OF SERVICE**

Oracle Database 10g takes Oracle Database's already market-leading high availability features another major leap forward with comprehensive support to minimize or eliminate all causes of both unplanned and planned downtime.

Oracle Database 10g addresses unplanned downtime from either computer or data failures. Real Application Clusters protect against node failure in a Grid with automatic node failover and redistribution of load over time. The integrated clusterware that comes with Oracle Database 10g reduces risk by deploying a single vendor solution stack.

With Oracle Database 10g, efficient utilization of inexpensive disk provides fast recovery from any type of data failure. At the disk level, Automated Storage Management mirrors with low cost disks and prevents any single points of failure with the concept of failure groups. Human errors are generally considered the biggest single cause of downtime, and Oracle Database 10g is a breakthrough release for human error correction by means of its greatly enhanced flashback

capabilities. Flashback is a capability that lets an administrator “rewind” activity on the database. Flashback can be performed at the row, transaction, table or even database level. Point-in-time recovery of an error performed a few minutes ago can be performed in a few minutes; no restore of a backup and roll forward is required. Oracle Database 10g enhances the Hardware Assisted Resilient Data (HARD) feature introduced in Oracle Database 9i Release 2 which prevents corruption introduced in the IO path between the database and storage. All database file types are checked and HARD is automatically enabled with ASM. Automatic backup and recovery can maintain a recovery area on disk that is incrementally “rolled forward” by updating on disk only those blocks changed since the last backup. The recovery area provides much faster and more reliable backup and restore compared to traditional tape based mechanisms. This ability to very quickly recover a system after a media failure results in much higher availability of the system.

Oracle Data Guard, introduced in Oracle Database 8i, protects against site failure by maintaining standby databases and has been further enhanced in Oracle Database 10g with optional compression and encryption of the log traffic and additional data type support for SQL Apply mode.

To protect against planned downtime, Oracle Database 10g has established functionality for allowing online redefinition of table and indexes. Enhancements have been made to allow tables to be redefined without invalidating stored procedures.

One of the major cause of planned downtime - upgrades - are addressed in Oracle Database 10g with support for rolling upgrades of hardware, operating system or database versions.

The ability to implement Enterprise Grid Computing will give users better performance. With the same resources, users can get more performance as resources can be provisioned to the right application based on business priorities or needs. Furthermore, Oracle Database 10g continues its record of database performance leadership through new performance features as well as database optimizations while expanding Oracle Database’s platform coverage to include 64-bit versions of Windows and Linux. Notably, a new PL/SQL compiler has been introduced that provides a framework for efficient and ongoing optimization of compute-intensive PL/SQL programs, with performance gains of about 2 times over Oracle Database 9i Release 2.

Oracle Database 10g continues to provide the highest level of security assurance, and integrates with Oracle’s new Identity Management infrastructure to provide identity provisioning, access management, and enforcement features. With Oracle Database 10g’s integration with Oracle Identity Management, administrators have the ability to delegate the database user password management to the users. This benefits the organization in lowered password administration costs and benefits the users since they have one password to remember. Applications using Oracle Identity Management have the ability to leverage Oracle Database 10g security features including Fine Grained Audit and generic Auditing, Virtual Private Database, Oracle Label Security features.

## **EASE OF MANAGEMENT**

Management is one of the largest contributors to the overall cost of ownership for software systems. One the major value propositions of Oracle Database 10g is the significant reduction in the management cost of deploying and maintaining an Oracle-based solution. Oracle Enterprise Manager Grid Control manages grid-wide operations including management of the entire stack of software, provisioning users, cloning databases and patch management. Oracle Enterprise Manager can even monitor the performance of all applications from the point of view of your end users.

Oracle Database 10g itself has taken a major step in simplifying and automating all the tasks in the life cycle of database management.

Oracle Database 10g has significantly simplified the installation and configuration of the database. The install has been made much faster and lightweight. Oracle Database 10g can be installed from a single CD, which holds binaries for the Personal, Standard and Enterprise Edition along with a pre-created seed database. The installation process for Oracle Database 10g has been extended to perform pre-requisites check to make sure the target system has the requisite OS patches, memory, and CPU. Oracle Database 10g installation can be run in a silent mode for installing and de-installing Oracle.

The number of relevant initialization parameters is drastically reduced. For example, there are now only two memory

size parameters for a DBA to set across the entire database: SQL execution (PGA) memory and shared (SGA) memory. Most Oracle Database 10g customers will only have to use 28 basic initialization parameters. Oracle Database 10g enhances the evolution of automated space management functions with the addition of online segment shrink to eliminate segment fragmentation.

To further ease ongoing management, Oracle Database 10g performs automatic performance diagnosis and produces tuning recommendations. An Automatic Workload Repository (AWR) has been introduced which periodically gathers and stores system activity and workload data. A new expert diagnosis engine called the Automatic Database Diagnostics Monitor (ADDM) analyzes this data. One area examined by ADDM is the top, or most resource intensive, SQL statements. Once identified, ADDM passes the top SQL onto the new SQL Tuning Advisor, which provides advice on how to improve their performance. The advice is in the form of one or more recommendations with a rationale and an estimated benefit obtained when implemented. ADDM and SQL Tuning Advisor encapsulate the expert knowledge of Oracle performance engineers into the core database engine. ADDM and SQL Tuning Advisor can provide immediate benefit to any application, including third party applications, since the recommendations are implemented transparently to the application. An example where the recommendations provide significant business benefit is when third party SQL runs sub-optimally due to real world customer data volumes and skew. Traditional approaches to tuning such third party SQL would involve significant time and labor cost since a test case would need to be produced, a patch developed, and the patch installed in the next maintenance schedule. With ADDM and the SQL Tuning Advisor, the recommendation on how to improve the SQL performance is automatically presented to the DBA who would then simply accept the recommendation to be transparently implemented.

## **APPLICATION DEVELOPMENT**

Applications running on a Grid computing infrastructure typically interface through common Web services. Oracle Database 10g has the capability to act as a Web services provider and consumer. In addition, Oracle Database 10g includes a powerful native development tool with HTML DB. HTML DB is a hosted, development environment for non-programmers that allows the creation of database-centric web applications. It has the developer productivity of a desktop database but deploys those applications on the Oracle database. Any IT organization can easily host HTML DB environments for departments in the company that have no Oracle developer skills available.

## **DATA WAREHOUSING**

Oracle Database 10g has also enhanced its data warehouse and business intelligence capabilities, which results in further reduction of the total cost of ownership while enabling customers to derive more value from their data and supporting real time data feeds.

Consolidation and integration of traditionally disparate business intelligence systems into a single integrated engine is further enhanced in Oracle Database 10g. Database size limits have been raised to millions of terabytes. Business Intelligence applications can be consolidated alongside transactional applications using Real Application Clusters automatic service provisioning to manage resource allocation. This consolidation means analysis can be performed directly against operational data and resource utilization can be maximized by reallocating servers to workloads as the business needs change. The value of data is increased with the ability to perform even more diverse analytic operations against core data with enhanced OLAP analytics, a data mining GUI and a new SQL model feature. The SQL model cause 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.

Real Time Warehousing is enabled either by consolidating business intelligence with operational applications, or by new change data capture capabilities based on Oracle Streams which produce low or zero latency trickle feeds with integrated ETL processing.

## **SUMMARY**

Oracle Database 10g is the first database designed for Enterprise Grid Computing. Oracle Database 10g gives you the flexibility you need to implement Enterprise Grid Computing, even using low cost servers and disks. The dramatically

lowered management burden of Oracle10g Database also helps to lower costs while providing highest quality of service. Oracle Database 10g provides the dual benefits of making IT more responsive to the changing business environment and lowering risks with lower costs, greater scalability and predictability and the highest levels of availability.