

Oracle Database 10g: A Revolution in Database Technology

*An Oracle White Paper
December 2003*

Oracle Database 10g: A Revolution in Database Technology

Executive Overview	3
IT Challenges	3
The Solution	3
Storage Grid	4
Database Grid	4
Highest Quality of Service	5
Revolutionary Availability	5
Industry Leading Performance	6
Integrated Security	6
Managing the Grid	7
The Self Managing Database	7
Grid Control	8
Data Warehousing	8
Managing All Your Data	9
Data Movement	9
Application Development	10
Conclusion	10

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.

STORAGE GRID

For databases, storage is one of the key resources. In today's typical 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 provides complete storage management for the Oracle database, eliminating the need for a file system and volume manager. ASM automatically spreads database data across all disks, providing the highest I/O throughput with minimal management costs. As disks are added or dropped, ASM dynamically redistributes the data automatically. Data availability is increased with optional mirroring, and disks can be added and dropped online. ASM enables the use of low cost storage with the same robustness and high quality as expensive controllers. ASM saves money in software license fees, in DBA and system administrators' time, and in performance analysis time.

DATABASE GRID

As mentioned above, the general approach in Enterprise Grid Computing is to create large pools of servers and to then dynamically provision the servers for systems as needed. From a business perspective, provisioning gives users better performance since resources can be provisioned to the right application based on business priorities or needs.

Oracle Real Application Clusters's (RAC) unique ability to run and scale all application workloads on a cluster of servers is an essential foundation of Enterprise Grids. Oracle Database 10g makes managing a clustered database easier with two major new capabilities: integrated clusterware and automatic workload management. The integrated clusterware is available on all platforms supported by Oracle Database 10g, and includes functionality for cluster connectivity, messaging and locking, cluster control and recovery and a workload management framework. Integrated clusterware makes Oracle RAC easy by eliminating the need to purchase, install, configure, and support third-party clusterware. Automatic workload management eases dynamic provisioning of servers to workloads (services). Rules can be defined to automatically allocate processing resources to each service during both normal operations and in response to failures, and these rules can be dynamically modified to meet changing business needs. This ability to dynamically adapt resource allocation within a Database Grid is unique to Oracle RAC and essential in implementing Enterprise Grid Computing.

In addition to the improvements in cluster management, Oracle RAC in Oracle Database 10g provides automatic event notification to the mid-tier in event of cluster configuration changes. The mid-tier can then immediately adapt to instance failover or availability of a new instance. This enables end users to

continue working in event of instance failure without delays typically caused by network timeouts. In event of new instance availability, the mid-tier can immediately start load-balancing connections to that instance. Oracle Database 10g JDBC drivers have “Fast Connection Failover” functionality that can be automatically enabled to handle these events.

HIGHEST QUALITY OF SERVICE

Revolutionary Availability

Oracle Database 10g takes Oracle Database’s already market-leading high availability features another major leap forward with revolutionary availability features 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 service failover and redistribution of load over time. In addition, the integrated clusterware that comes with Real Application Clusters 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, Automatic Storage Management mirrors with low cost disks and prevents any single point of failure with the storage system. 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, allowing easy undo of any erroneous changes to the data. Flashback can be performed at the row, transaction, table or even database level. Flashback revolutionizes data recovery. In the past it might take minutes to damage a database but hours to recover it. With Flashback the time to correct errors is comparable to the time it took to make the error. It is extremely easy to use and a single short command can be used to recover the entire database instead of following complex procedures. Oracle Database 10g enhances the Hardware Assisted Resilient Data (HARD) feature introduced in Oracle Database 9i Release 2 which prevents corruption introduced in the I/O path between the database and storage. The validation is more comprehensive and all file types and blocks are protected by HARD, including database files, online logs, archive logs and backups. In addition ASM enables HARD without having to use raw disk devices. 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 real time apply for close synchronization of the standby with production.

To protect against planned downtime, the Oracle Database has existing functionality to support online schema and data reorganization. New enhancements in Oracle Database 10g include allowing tables to be redefined without invalidating stored procedures, which reduces the time required for application upgrades.

In addition to these industry leading availability features, Oracle Database 10g addresses one of the major causes of planned downtime - upgrades - with support for rolling upgrades of hardware, operating system or database versions.

Industry Leading Performance

The ability to implement Enterprise Grid Computing gives 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. In addition, performance will be automatically enhanced by many of the self-managing features of Oracle Database 10g such as the self tuning shared memory (SGA) which allows memory to be dynamically and automatically distributed based on workload.

Integrated Security

Oracle Database 10g demonstrates Oracle's continued leadership in security, providing strong server enforced security and privacy features. Oracle Database 10g goes beyond row level security with column relevant Virtual Private Database and integrates Oracle Label Security with Oracle Identity Management. Accountability increases with enhancements to fine grained auditing. Fine grained auditing can now be used with insert, update and delete operations. Oracle Database 10g introduces strong integration with Oracle's new Identity Management infrastructure to provide identity provisioning, access management, and enforcement features for large numbers of users across an integrated end-to-end environment. With Oracle Database 10g's integration with Oracle Identity Management, administrators have the ability to centrally provision database users for the entire enterprise, and integrate with existing Identity Management infrastructures. This benefits the organization in lowered user administration costs and benefits the users since they have only one password to remember. Oracle Database 10g features such as enterprise user security enable applications to manage users in the Oracle Internet Directory, reducing the number of database accounts. In addition, Oracle Database 10g reduces the cost of administering security with simplified security procedures

and enhanced database integration with Oracle Internet Directory and interoperability with 3rd party Identity Management infrastructures. Oracle recently received its 17th independent security evaluation and continues to lead the industry in information assurance.

MANAGING THE GRID

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 Database 10g itself has taken a major step forward in becoming self-managing while Oracle Enterprise Manager Grid Control provides a central infrastructure for grid-wide operations.

The Self Managing Database

Oracle Database 10g has taken a major step in simplifying and automating all the tasks in the life cycle of database management. The installation and configuration of the database has been significantly simplified. Installation time for the Oracle Database 10g server software has been cut in half to 20 minutes on a typical configuration. The most common installs can be performed from a single CD ROM eliminating the need for media changes. The installation process for Oracle Database 10g has been extended to perform pre-requisites checks to make sure the target system has the requisite OS patches, memory, disk and CPU. For developers and Independent Software Vendors (ISVs) developing applications with embedded Oracle Database, silent mode install has been made more robust and user friendly.

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 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. This provides an historical recording of performance data, which has many benefits including: reducing the need to replay workload for detailed analysis, the ability to compare performance data between two points in time and perform trend analysis. In addition, a new expert diagnosis engine called the Automatic Database Diagnostics Monitor (ADDM) uses the AWR for automatic performance diagnosis. One area examined is the top, or most resource intensive, SQL

statements, which can then be passed on to the new SQL Tuning Advisor, which provides advice on how to improve performance of the SQL. 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 inside the core database engine. ADDM and SQL Tuning Advisor provides immediate benefit to any application, including third party applications, since the recommendations are implemented transparently to the application. This ability to perform automatic performance diagnosis and produce tuning recommendations significantly enhances DBA productivity and improves performance of the database.

Grid Control

Oracle Enterprise Manager Grid Control manages grid-wide operations including management of the entire stack of software, provisioning users, cloning databases and patch management. Grid Control can even monitor the performance of all applications from the point of view of your end users and provide service level performance and availability reporting. Grid Control views the performance and availability of the grid infrastructure as a unified whole rather than as isolated storage units, databases, and application servers. IT staff can group hardware nodes, databases, and application servers into single logical entities and manage a group of targets as one unit.

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. Business Intelligence applications can be consolidated alongside transactional applications using Real Application Clusters automatic workload management 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 clause 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 can be achieved 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.

MANAGING ALL YOUR DATA

While Automatic Storage Management allows Oracle to manage storage, Oracle Database 10g builds on its existing superiority in managing all your data – relational, email, documents, multimedia, XML and spatial – in an integrated fashion.

The XML capabilities of Oracle Database 10g 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, including Microsoft Advanced Server Format (ASF) for streaming database content to Windows Media Server.

Oracle Spatial supports new spatial data models, spatial business analytics, performance improvements and other new features. These enhancements extend the range and productivity of application developers, enabling a broader range of applications and improving performance.

The ability to manage textual information in Oracle Database 10g has been extended to include the ability to “mine” text for purposes of document clustering, classification, and information extraction. This functionality can benefit many textual applications including automated email routing and response engines, and mass analysis of free form notes such as call center logs for marketing purposes.

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), increased size limits of LOBs stored in the database (128 terabytes) and by the enhanced manageability functions of Automatic Storage Management, automated space management and the new recovery area.

DATA MOVEMENT

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”, and can now operate across platforms. Among other things, this allows databases to be migrated to a new platform at the speed of file transfer. In addition, Oracle Data Pump is a new feature of Oracle Database 10g that provides high speed, parallel, bulk data and metadata movement of Oracle database content across platforms. Oracle 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 Oracle Data Pump has been significantly enhanced (1.5 to 10x) over previous import and export versions, and can be automatically scaled up with the number of available CPUs.

APPLICATION DEVELOPMENT

Oracle Database 10g provides the most comprehensive functionality for developing versatile, scalable, and high performance database applications for Grid computing. 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 enterprise infrastructure of the Oracle database. Any IT organization can easily host HTML DB environments for departments in the company that have no Oracle developer skills available.

For Java applications, Oracle Database 10g has enhanced JDBC drivers which provide comprehensive JDBC 3.0 support, faster execution, and support for Grid computing with a new Implicit Connection Cache and Fast Connection Failover.

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.

Furthermore, 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 9i Database Release 2.

CONCLUSION

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 Oracle Database 10g 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.



White Paper Title
December 2003
Author: Penny Avril
Contributing Authors:

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

Oracle Corporation provides the software
that powers the internet.

Oracle is a registered trademark of Oracle Corporation. Various
product and service names referenced herein may be trademarks
of Oracle Corporation. All other product and service names
mentioned may be trademarks of their respective owners.

Copyright © 2002 Oracle Corporation
All rights reserved.