

Oracle Database 11g Performance and Scalability

An Oracle White Paper
November 2007

Oracle Database 11g Performance and Scalability

Introduction	3
Scalability Challenges	4
Scalability Factors	4
Scalable Execution	5
Server Result Cache	5
Oracle Call Interface (OCI) Consistent Client Cache.....	7
Database Resident Connection Pool.....	7
Cache Fusion Protocols.....	8
Highly Available Reader Farms.....	8
Optimizer Statistics	10
Real Native PL/SQL and Java Compilation	10
Scalable Storage	10
OLTP Table Compression.....	10
SecureFiles.....	11
Partitioning Enhancements.....	11
Scalable Availability.....	12
Backup and Recovery	12
Data Guard.....	13
Scalable Management.....	13
Enhanced Self-Management.....	14
Database Replay.....	14
SQL Plan Management.....	14
Conclusion.....	15

Oracle Database 11g Performance and Scalability

INTRODUCTION

Enterprises continue to leverage their investments in Information Technology (IT) to grow and improve their business. IT enables companies to better understand and interact with customers and to innovate and develop next generation products and services. The importance of IT in enterprises is illustrated by the increase in the number of applications and the demands these applications have on technology infrastructure capacity. With data retention regulations, the online distribution of rich multimedia content, and the advent of Web 2.0 (web services whose content is derived from end-user collaboration) – technology capacity demands will continue to grow at exponential rates.

The increased demand and importance of IT in combination with the increasing volumes of data provide a unique challenge to IT organizations within the enterprise. The technology infrastructure not only needs to meet the current performance requirements but also must continue to scale seamlessly as the business demand grows.

For over three decades, Oracle Database has proven to be the clear performance and scalability leader. Oracle Databases power a vast majority of the world's most demanding environments and continue to build an all-around lead across various industry benchmarks. Oracle currently owns most of the world records for TPC, SAP and other such benchmarks. Similarly, a survey of the world's largest and busiest databases as measured by the [Winter Top Ten program](#) clearly demonstrates Oracle's real world leadership. Oracle Database 11g continues to build upon the proven foundation of earlier releases to offer customers next generation performance and scalability and help them cope with the enormous demands being placed on their IT infrastructure in the most cost effective manner.

In this paper, we will further discuss the challenges a modern enterprise faces in establishing an IT infrastructure that performs and scales in accordance with the demands of the business. We will then examine the technologies available within the Oracle database, focusing on the new capabilities available in Oracle Database 11g, which enable the IT infrastructure to prepare for and meet these challenges head on.

SCALABILITY CHALLENGES

Today's enterprises rely on the Oracle database to scale and perform in their increasingly complex and growing IT environment. As stated earlier, the Oracle Database is already the natural choice for customers who are looking for the very best performance and scalability. However, the need for more computing capacity and speed continues to grow at an exponential rate due to a variety of factors. Businesses are introducing more applications and those applications are serving more and more users – both within the company and outside – resulting in an appetite for massive amounts of processing power. At the same time, developments such as long term data retention regulations, corporate mergers and acquisitions, growing volume of unstructured data (document, images, multimedia, XML), etc. are fueling unimaginable growth in the size of databases. A Terabyte sized database that used to be a novelty some years back is almost a normal occurrence today as businesses try to capture as much data they possibly can about their customers and store it for as long as they can in order to understand their business environments better. Similarly, the volume of data being generated in the scientific research and healthcare sector is literally exploding with some organizations expecting to generate more than a Petabyte (1000 Terabytes) of data every single year.

Fortunately, hardware vendors have responded to these challenges by making enormous amounts of computing power available at very affordable prices. Revolutionary developments in the multi-core processor technology are leading to a situation where some of the enterprise applications may be running on 1000 processor systems within a span of the next 3-5 years. Also, thanks to the plummeting prices of storage and memory, it won't be uncommon for customers to have databases sized in the hundreds of terabytes running on machines with multi-terabytes of memory.

Oracle Database 11g has been specifically architected to take advantage of this enormous computing capacity to provide customers with the next generation of performance and scalability. Thanks to a number of innovative enhancements to the database engine and new features, Oracle Database 11g is well positioned to scale up to thousands of processors and handle massive amounts of relational (structured) and file (unstructured) data. In addition, Oracle Database 11g incorporates technology that allows enterprises to manage massively large databases without compromising on the quality of service or inflating their IT budget.

SCALABILITY FACTORS

- **Scalability Factors**
 - **Scalable Execution**
 - **Scalable Storage**
 - **Scalable Availability**
 - **Scalable Management**

IT organizations need to focus on three primary factors when architecting their infrastructure to meet performance and scalability demands. First, application execution needs to meet the performance expectations and service level agreements with the end user population. Oracle Database 11g enables this through two main areas discussed in the following sections: *Scalable Execution and Scalable Storage*. Second, as the volume of data grows it becomes more challenging to maintain a Highly Available architecture. With this release of Oracle, advances in backup and recovery, failover (Data Guard), and clustering technologies enable *Scalable Availability* for the enterprise. Finally, as the number of applications and hardware resources increase, there must be an efficient and cost-effective method to manage the infrastructure. Oracle Database 11g's *Scalable Management* automates many

administrative activities and provides an intuitive and robust toolset for managing simple and complex environments with Enterprise Manager Grid Control.

SCALABLE EXECUTION

As discussed earlier, Oracle Database 11g fully leverages the massive amounts of hardware capacity to improve application performance and scalability. As the cost per megabyte of memory continues to decline, many customers are considering using large amounts of memory to work around storage bottlenecks and improve their application performance. Oracle Database 11g introduces new caching functionalities to better utilize the expanded memory footprint to speed up query processing.

Server Result Cache

Server Result Cache is a new component of the SGA that caches 'results' of queries and query fragments.

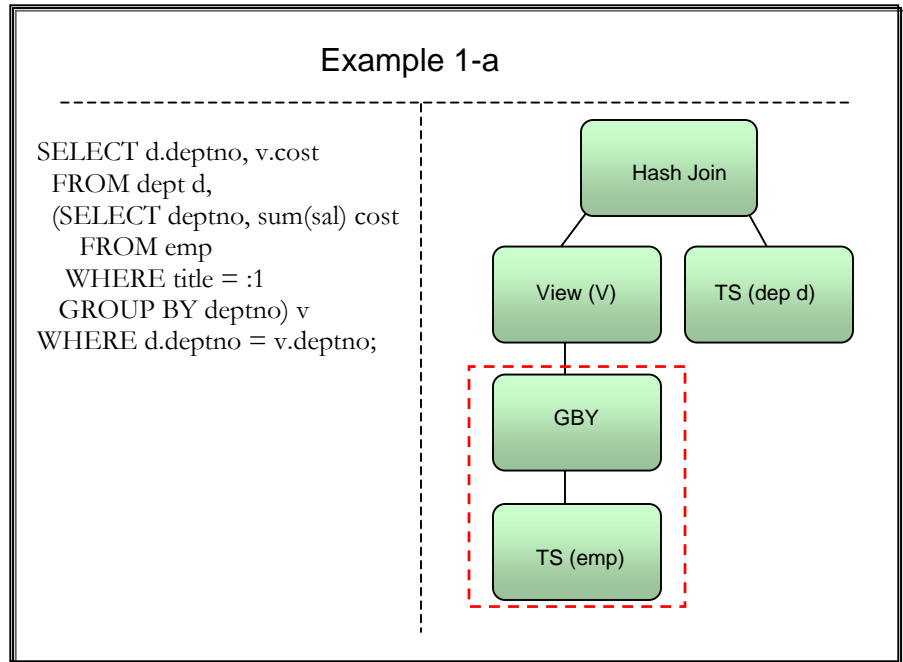
Server Result Cache, a new feature in Oracle Database 11g, enables query results to be cached in memory. These cached results are then used during future executions of similar query or query fragments to bypass regular query processing and return the results faster. Enabling Server Result Cache optimizes SQL execution by decreasing wait times for both physical and logical IO by directly fetching the results from memory. The cached results are completely shareable between sessions and SQL statements - as long as they share common execution plans, either partially or fully - and persist beyond the life of the initiating cursor. Oracle manages the cache consistency by flushing entries from the Result Cache whenever the underlying data changes. Internal tests show that the use of Server Result Cache can enhance the performance for read-intensive workloads by as much as 200%.

Server Result Cache can double the performance of your read-intensive workload.

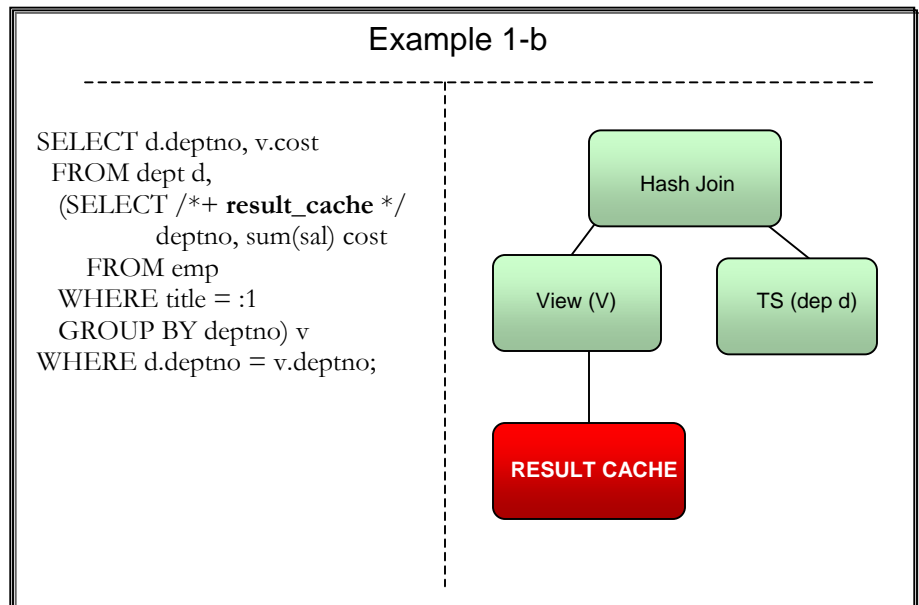
The memory dedicated to storing the results cache is located in the System Global Area and is automatically managed by the Automatic Shared Memory Management infrastructure. The result cache is enabled by setting the maximum size of the cache using the new initialization parameter `RESULT_CACHE_MAX_SIZE`. With global and fine-grained controls, there are several options available to integrate the Results Cache into your architecture. Administrators can define its usage at the system, session, table, and/or statement levels.

Administrators can monitor and manage the contents of the Server Result Cache via a new PL/SQL interface. Through this interface, administrators have the ability to flush the entire cache or individual cached result sets. Administrators can generate Result Cache memory usage reports via this interface and can utilize data dictionary views to further analyze the effectiveness of the Server Result Cache. A new optimizer hint allows users to specify the use of the Server Result Cache at the query level.

Example 1-a represents a query execution **without** the result_cache hint:



Example 1-b represents the execution of the same sql statement however, this time with the result_cache hint:



Note the new Operator 'RESULT CACHE'. Issuing this hint causes the operator to examine the Result Cache to determine if a result for this execution plan is already present in the cache. If the result is found in the cache, then the operator

bypasses the execution of the underlying execution plan and returns rows directly from the Result Cache. If the result is not present in the cache, it will fetch the rows from the underlying execution plan and store these rows in the Result Cache.

Client Cache technology reduces round-trips between the client and database server – reducing server CPU Utilization.

Oracle Call Interface (OCI) Consistent Client Cache

OCI Consistent Client Cache (Client Cache) is a complimentary caching feature to the Server Result Cache discussed in the previous section. It enables the caching of query results on the client machines. The Client Cache utilizes per-process memory on the OCI client, typically on an Application Server, to store data that affects query result sets (i.e. tables or entire result sets). The Client Cache resides in the OCI client process memory so that its contents can be shared across multiple sessions and/or threads. With the Client Cache in place, clients fetch result sets directly from the cache rather than having the server execute the query repeatedly.

Utilizing the Client Cache has an enormous performance gain by requiring fewer round-trips between the client and database server, thus reducing CPU utilization on the server as a result of executing fewer SQL calls. The Client Cache is optimal for queries of small lookup tables that are generally read-only or read-mostly. Benchmarks show that simple queries utilizing the Client Cache can have a 500% reduction in elapsed time and a 200% reduction in CPU time.

Queries utilizing the Client Cache can have a 500% reduction in elapsed time and a 200% reduction in CPU time.

Built into the Client Cache technology is an inherent mechanism that manages the consistency between the database server and the cache. Given that the technology resides within the OCI layer, the feature is automatically available to all Oracle Database 11g OCI clients, including: JDBC-OCI (Thick) Driver, ODBC, ODB.NET, PHP, and precompilers. Therefore, the Client Cache feature is transparently available to all applications that utilize these Oracle Database 11g OCI-based clients without having to change any application code.

To enable Client Caching, the database initialization parameter *CLIENT_RESULT_CACHE_SIZE* must be set to a value greater than zero. This parameter dictates the maximum size of the client per-process result set cache. In addition, the table properties, optimizer hints, and system/session parameters discussed in the previous section apply for the Client Cache. The use of a server parameter simplifies the management of the Client cache by providing a centralized way to manage the caching behavior of a large number of clients. Optionally, a client configuration file could be maintained which would override the Client Cache settings on the database server. The optional settings available on the client-side include maximum per-process cache size, maximum number of rows per cache, and the maximum size of any result set in both bytes and rows.

Database Resident Connection Pool

Enterprise applications have long utilized database connection pooling at the Application Server level. Connection pooling allows applications to use a reduced number of database sessions to service many application end-users. The ultimate objective of connection pooling is to improve application performance and scalability by limiting the overhead of database session creation and reducing database session memory utilization. Certain technologies, such as PHP, are required to use at least one database connection per web server process and thus are not candidates for Application Server based connection pooling. Furthermore, large enterprises already using Application Server level connection pooling may still

PHP applications and large Application Server Farms leverage Database Resident Connection Pools for vast increases in scalability

experience high database session counts when using 100's or 1000's of Application Servers.

The Database Resident Connection Pool is a highly scalable solution providing session connection pooling managed by the Oracle database. When Database Resident Connection Pools are enabled, clients connect to a new background process, the connection monitor (CMON), instead of a dedicated server process. CMON is responsible for managing the server side connection pool functionality. Clients accessing the database using a common username leverage previously allocated sessions. The client transparently caches persistent connections to CMON and uses them when the application requests database connections. When the application closes database connections, the dedicated server process is returned to CMON and is placed back into the pool ready for the next client request. Database Resident Connection Pools provide tremendous scalability to both applications unable to implement Application Server connection pooling and to applications that are hosted on large Application Server farms.

Cache Fusion Protocols

Oracle revolutionized clustering technology by introducing the Cache Fusion protocol that allows nodes in a cluster to communicate with each other's memory using a high speed interconnect. Cache Fusion is one of the key elements of Oracle's Real Application Cluster technology.

Oracle Database 11g introduces the next generation of Cache Fusion Protocols. The new workload aware protocols dynamically alter the inter-node messaging behavior depending on the type of workload being processed in order to maximize performance. For example, a newly introduced read optimized protocol significantly reduces the number of inter-node messages for read operations. Similarly, other new protocols optimize the messaging behavior for update and table scan operations. Oracle automatically selects the best protocol to use based on the workload profile. These optimizations to the Cache Fusion Protocols significantly reduce inter-node messaging and enhance the scalability of Real Application Clusters.

Highly Available Reader Farms

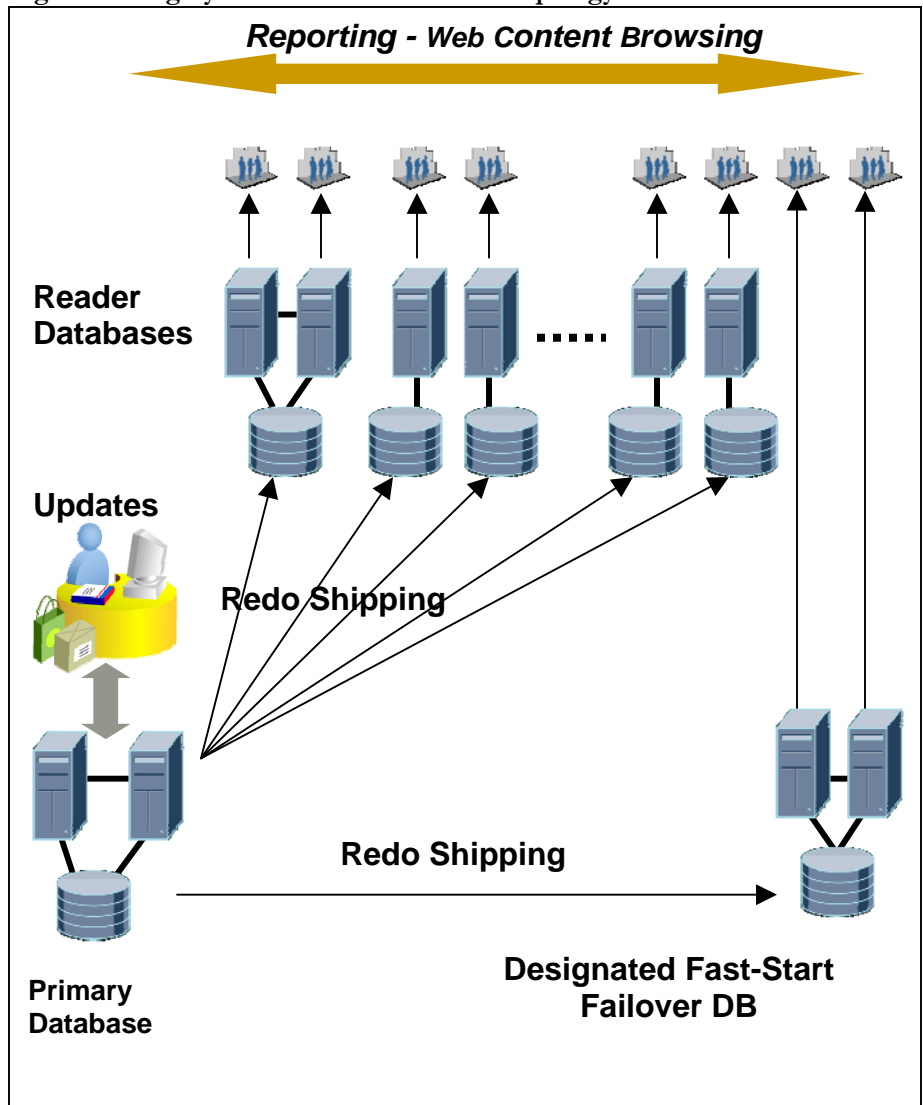
Oracle Data Guard is Oracle's premier disaster recovery technology. Using this technology, customers can protect their database against any site failures by maintaining a synchronized copy of the database remotely, called a standby database. The standby database can be activated in case of a failure or planned maintenance operation in order to ensure business continuity. Customers have an option of creating two kinds of standby database. Using Data Guard Redo Apply, customers can create a physical standby that mirrors the primary database at the database block level. Alternatively, a Logical Standby can be created using the Data Guard SQL Apply technology where the changes from the primary database are first converted into SQL and then applied on the standby database. Unlike a physical standby, a logical standby allows user to perform write operation on the standby database – most commonly to create addition indexes or materialized views to optimize reporting activities. However, due to its simplicity, a physical standby can apply the changes faster and thus has performance advantages over the logical standby.

Besides providing effective disaster protection, standby databases can also be used for a number of day-to-day operation purposes. For instance, standby databases

can be used to offload reporting or backup workloads from the primary database or for testing purposes. Prior to Oracle Database 11g, a physical standby database could not apply the changes received from the primary database while queries were being run. Oracle Database 11g lifts that restriction and enables the physical standby to be used for Real-Time query purposes. This interesting development opens several additional avenues, including Reader Farm Databases. A number of web applications that support a very large number of users employ an architecture where several read only replica of the main databases are created to support the large concurrent workload and increase application availability.

The Physical Standby with Real-Time query feature makes Oracle Database 11g the ideal foundation for a Reader Farm database architecture. All updates are made only to the primary database with the standby databases serving the read only workload. Additionally, using one of the standby databases as a designated Fast Start Failover standby further protect the primary database against disaster leading to enhanced availability. Figure 1 shows a typical Highly Available Reader Farm topology:

Figure 1 – Highly Available Reader Farm Topology




Optimizer Statistics

Oracle's Cost Based Optimizer (CBO) depends on object statistics to help determine the optimal execution plan for SQL statements. Oracle Database 11g enhancements improve Optimizer Statistics by making them faster, better, and safer. First, Oracle Database 11g performance enhancements make the collection of optimizer statistics *faster* - decreasing the total amount of time required to gather and compute optimizer statistics. Secondly, the computed optimizer statistics are more thorough, providing *better* information to the CBO by correlating statistics, such as Number of Distinct Values (NDV) and histograms, on multiple columns. Lastly, Oracle Database 11g makes gathering statistics *safer*. Newly generated statistics can remain in a private 'statistics workspace' until they can be validated and published for general use.

Real Native PL/SQL and Java Compilation

Oracle Database 11g improves the performance and scalability of both PL/SQL and Java by implementing Real Native Compilation (RNCOMP) for PL/SQL and Just In Time (JIT) compilation for Java programs stored in the Oracle database. RNCOMP improves upon the Native Compilation (NCOMP) in previous versions of Oracle by eliminating the need for customers to acquire, install, and license a C compiler. Both RNCOMP and JIT improve performance by compiling source code into the machine's native instruction set. Improvements in performance could range from 30 to 100% depending on how compute intensive the application is.

SCALABLE STORAGE



Enterprises continue to demand more physical storage for both structured relational data and unstructured data (documents, spreadsheets, etc). Hardware vendors continue to expand the per-disk capacity to meet these demands; however, the performance of disks has not improved at this same pace. With multi-terabyte databases becoming commonplace, Oracle Database 11g provides rich enhancements that empower IT organizations to cost-effectively scale their storage infrastructure, while continuing to meet performance requirements.

OLTP Table Compression

Oracle introduced data compression in Oracle Database 9i primarily for use with Data Warehouses. Limitations on how data could be manipulated in compressed tables made it unsuitable for OLTP workloads. Oracle Database 11g removes these constraints, allowing traditional DML statements (INSERT, UPDATE, and DELETE) on compressed tables. These additional compression capabilities make table compression a unique tool for DBA's managing OLTP databases. Table compression provides both a reduction in storage capacity demand and an increase in application performance.

Table compression is a strategic asset IT organizations can leverage to help manage the increasing demand in storage capacity. Typically, compressed tables consume 2 to 3 times less disk space than uncompressed versions of the same tables. Oracle operates directly on the compressed data – eliminating any application overhead associated with uncompressing data. The use of compressed tables increases performance primarily by reducing the amount of physical IO and by enhancing

Oracle's table compression has been enhanced to allow all DML operations - making compression available to OLTP applications.

cache efficiency. Table scans off disk perform up to 2 times faster for compressed tables.

SecureFiles

Unstructured data (documents, spreadsheets, images, multimedia, etc) is becoming more pervasive in the enterprise. This is due in large part to improvements in networking bandwidth enabling the sharing of rich unstructured content and the changing requirements of the regulatory landscape, such as Sarbanes-Oxley and HIPPA. IT organizations typically store unstructured content on file systems, rather than databases, due to perceived simplicity and performance benefits of file systems. Applications utilize Oracle databases to store, organize, and manage structured data, including metadata describing unstructured data. However, this paradigm, of having two independent data stores, adds additional complexity in managing, securing, and scaling the IT infrastructure. Specifically, this model introduces:

- Disjointed security and auditing models
- Fragmented backup and recovery
- Non-atomic data changes (delete files without updating the metadata)
- Complex search across disparate information stores
- Heterogeneous space management strategies


SecureFiles, a feature introduced in Oracle Database 11g, was designed to break the performance barrier preventing enterprises from storing unstructured data in Oracle databases. Similar to LOBs, SecureFiles is a data structure built to store large objects in the database. SecureFiles however, offer a much richer feature set and a vast improvement in performance when compared with LOBs or file systems.

The SecureFiles feature set includes leading edge-file system capabilities such as transparent deduplication, compression, and encryption – yet provides advanced database capabilities that far exceed that of file systems. SecureFiles are managed within Oracle’s transaction and read consistency models and are inherently supported by High Availability features such as Real Application Cluster, Readable Standby, Data Guard, Consistent Backup, Point in Time Recovery, and Flashback. Enterprises storing their unstructured data as SecureFiles will benefit from unified approaches to data security, backup and recovery, and space management. Internal tests at Oracle indicate performance parity of SecureFiles with file systems – for both read and write operations. With the benefits of the Oracle database and the increase in performance, SecureFiles provide IT organizations with a strategic advantage in storing and managing unstructured data.


Partitioning Enhancements

Several new partitioning features introduced in Oracle Database 11g provide IT organizations with simplicity in managing their partitioned objects while increasing their ability to scale and perform. **Interval Partitions**, an extension of Range Partitions, is a feature developed to ease the management of partitioned tables. Prior to Oracle Database 11g, IT organizations were responsible for creating Range Partitions using DDL statements based on their partitioning scheme. With Interval Partitions, DBA’s simply specify the interval in which Oracle will automatically create partitions when the first row meeting the interval requirements is inserted into the table.

Query/Read Performance



Insert/Write Performance



In previous versions of Oracle, **Composite Partitioning**, or the use of subpartitions, was permitted but the partition/subpartition combinations were limited. In Oracle Database 11g, these restrictions have been removed to allow for all combinations of Composite Partitions, including the new Interval Partitions. Table 1 depicts the Composite Partitioning Matrix:

Table 1: Composite Partitioning Matrix

	Range	List	Hash
Range	YES	YES	YES
List	YES	YES	YES
Interval	YES	YES	YES

Data models that incorporate data normalization in their design take advantage of the relational model to reduce the duplication of data. In Oracle Database 11g, data models utilizing Referential Integrity Constraints, or Foreign Keys, have the option to build partitions on ‘child’ tables based on the partitioning scheme of ‘parent’ tables. **Partition by Reference** creates partitions on ‘child’ tables based on the specified Referential Constraint. The ‘child’ table is not required to contain the columns that the ‘parent’ table uses as partition keys. The ‘child’ partitions are stored in the same relative partition as the ‘parent’ partitions allowing for extremely efficient partition-wise joins.

Virtual Columns, a feature introduced in Oracle Database 11g, are columns that are defined using expressions but appear as traditional columns in SQL. Virtual Columns are supported as partitioning keys, can be indexed, and can contain optimizer statistics and histograms. For example, a virtual column could be defined in a table definition such as:

```
CREATE TABLE t1 (
  c1 NUMBER,
  c2 NUMBER,
  vc AS (c1 + c2) VIRTUAL)
```

SCALABLE AVAILABILITY

Availability of applications and company information is vital to the success of the enterprise. As the importance of IT continues to grow – so too does the technology infrastructure. Therefore, IT organizations need to ensure that the technology driving their infrastructure is built with scalable High Availability (HA) features and tools. Oracle offers world-class products with built-in HA functionality. RAC, Data Guard, RMAN, and Flashback are a few examples of HA features available in Oracle. Oracle Database 11g introduces exciting new HA functionality and enhances many of the existing features for scalability and performance.

Backup and Recovery

Oracle Secure Backup is a new Oracle product that provides centralized tape backup management for both databases and file systems. Oracle Secure Backup, through its seamless integration with RMAN, backs up databases up to 25% faster than the leading competition. This is accomplished by leveraging direct calls into

Oracle Secure Backup, a centralized tape backup management system, is up to 25% faster than the leading competition.

the database engine and through efficient algorithms that skip unused data blocks. The upcoming release of Oracle Secure Backup will widen this performance lead further with intelligent functionality like avoiding the backup of unnecessary UNDO tablespace data. .

Oracle Database 10g introduced **Block Tracking** functionality to increase the performance of incremental RMAN backups. While this feature was available on a Data Guard Physical Standby – it was not actively updating the block tracking file during managed recovery. Oracle Database 11g enables block tracking on a Physical Standby during managed recovery, allowing DBAs to take advantage of the increase in performance of incremental backups while offloading the backup processing activity to the standby server.

With the size of enterprise databases continuing to grow – it has become more advantageous to take advantage of Bigfile Tablespaces. Yet another innovation from Oracle, a Bigfile Tablespace is made up of a single large file rather than numerous smaller files, allowing Oracle Databases to scale up to 8 exabytes in size. Oracle Database 11g features an enhanced Recovery Manager (RMAN) tool that can perform intra-file parallel backup and recovery operations for Bigfile Tablespaces, thereby making the backups and restore operations significantly faster.

Data Guard

As discussed earlier in the paper, Data Guard is an integral component of Oracle's High Availability solution and a key component of the Maximum Availability Architecture (MAA). Data Guard automates the process of maintaining standby databases, which are activated during planned and unplanned production outages. Oracle Database 11g enhances the Data Guard capabilities by improving the redo transport performance. Asynchronous (ASYNCH) redo transport has been improved to stream redo data without interruption by decoupling the network messaging between the standby and production database from the shipment of redo data. Having the network messaging and redo shipping be asynchronous provides greater network throughput, thus providing data protection at a lower cost. Further improvements in network throughput include the new capability to compress the redo data as it is being transferred between production and standby databases. This is particularly beneficial when the standby and production servers have been unable to communicate due to a network outage or standby server failure. The standby database automatically performs redo log gap resolution that can leverage the compression capability to increase the rate of log transport.

SCALABLE MANAGEMENT

One of the biggest challenges facing IT administrators today is managing the performance and scalability of the IT infrastructure and the applications that depend on that infrastructure. Administrators are tasked with managing performance from both a proactive and reactive perspective. Proactive management includes scaling the infrastructure for measured increases in both usage and data volume. It also requires the administrator to test and prepare for changes in the application and the infrastructure, for example Oracle software updates. Reactive management includes the day-to-day responsibilities of the administrator tasked with reactive tuning of the application or database due to unexpected changes in workload profile.

Enhanced Self-Management

Oracle Database 10g took a giant step forward toward making the database self-managing. Thanks to widely acclaimed and innovative features such as Automatic Workload Repository (AWR), Automatic Database Diagnostic Monitor (ADDM) and SQL Tuning Advisor, Oracle Database 10g provides administrators with a valuable set of tools to identify and correct performance bottlenecks quickly and easily. Oracle Database 11g continues to build on this momentum by taking database self-management to the next level. Enhancements in ADDM allow it to better analyze RAC performance. The SQL Tuning Advisor in Oracle Database 11g is now capable of being fully automated - implementing SQL tuning recommendations automatically - relieving the administrator from this previously manual task. SQL Access Advisor has been enhanced to provide recommendations on the partitioning strategy to improve application performance. Finally, a new advisor has been introduced in Oracle Database 11g to examine Streams configuration and help administrators resolve any performance bottlenecks.

Database Replay

Successful performance and capacity management requires the successful implementation of system change impact analysis. IT administrators constantly struggle in creating testing scenarios that accurately represent the production workload. Typically, administrators simulate production workloads by running automated test scripts or by employing users to test application functionality. Both of these methods have proven to be insufficient in performing an exhaustive impact analysis due to their inherent limitations in accurately representing production workloads. These deficiencies often lead to production performance degradation and production instability. Oracle Database 11g introduces Database Replay, a technology that overcomes the deficiencies of the traditional change analysis testing techniques. Database Replay, when initiated by the DBA, captures the database workload activity on the production database. The captured workload data is then replayed on a test server running a copy of the production database in order to simulate the real world production workload. Unlike other workload capture technology that artificially “simulate” production workload, Database Replay captures the entire production workload at the lowest level of detail and, for the first time, enables customers to run full-scale workloads in a test environment.

Engineered for flexibility, the Database Replay technology is extremely beneficial during Oracle database upgrade testing and platform migrations. The workload data can be captured and replayed on different versions of the database and even on different platforms.

SQL Plan Management

Oracle’s Cost-Based Optimizer (CBO) generates SQL Execution Plans based on multiple factors including: CBO version, CBO parameters, object statistics, system settings, etc. Changes in any one of these components can cause the optimizer to generate different execution plans for a particular SQL statement. Various features, such as optimizer hints and stored outlines, have been available to assist in plan stability. Oracle Database 11g takes plan stability to the next level with the advent of the SQL Plan Management capability. With this new feature, Oracle automatically maintains a history of past execution plans and uses this information to ensure that dynamic plan changes don’t affect SQL performance adversely. This functionality will help ensure that applications continue to perform consistently

The Database Replay technology records live production database activity. The captured recordings can be replayed on test databases for upgrade and performance testing.

amidst changes to data volume, data distribution, system configuration or any other such changes.

CONCLUSION

Enterprises continue to depend more and more on Information Technology. As the demand for technology increases, IT organizations are tasked with ensuring the scalability and performance of their applications, databases, and IT infrastructure. Building on over 30 years of technology innovation, Oracle Database 11g offers the next generation of performance and scalability for enterprise applications. Oracle Database 11g has been designed to fully leverage evolving technology trends to satisfy the intense demands being placed on the IT infrastructure. Oracle Database 11g will help businesses keep their users and customers satisfied, lay the foundation for business success and stay ahead of the technology curve.



Oracle Database 11g Performance and Scalability

November 2007

Author: William Hodak

Contributing Author: Sushil Kumar

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

oracle.com

Copyright © 2007, 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, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.