An Oracle White Paper
April, 2011

Oracle Database 11g for Data Warehousing and Business Intelligence

Introduction

Oracle Database 11g is a comprehensive database platform for data warehousing and business intelligence that combines industry-leading scalability and performance, deeply-integrated analytics, and embedded integration and data-quality -- all in a single platform running on a reliable, low-cost grid infrastructure. Oracle Database 11g provides best-of-breed functionality for data warehouses and data marts, with proven scalability to 100s of terabytes and record-breaking performance. It also provides a uniquely integrated platform for analytics; by embedding OLAP, Data Mining, and statistical capabilities directly into the database, Oracle delivers all of the functionality of standalone analytic engines with the enterprise scalability, security, and reliability of an Oracle Database.

Oracle Database 11g is also a central component in the Oracle Exadata Database Machine, a complete warehouse configuration that includes servers, storage, networking, and software. Oracle Exadata delivers extreme high performance and is massively scalable, secure, and redundant.

This white-paper provides an overview of Oracle's capabilities for data warehousing, and discusses the key features and technologies by which Oracle-based business intelligence and data warehouse systems easily integrate information, perform fast queries, scale to very large data volumes and analyze any data.

Hardware

The bedrock of a solid data warehouse solution is a scalable, high-performance hardware infrastructure. One of the challenges for data warehouses has been to deliver the IO bandwidth necessary for large-scale queries, especially as data volumes and user workloads have continued to increase. The Oracle Exadata architecture provides a uniquely optimized storage solution for data warehousing that delivers order-of-magnitude performance gains for large-scale data warehouse queries. The technology driving these performance gains is a combination of a massively parallel high-bandwidth storage grid coupled with the Exadata Storage Server Software's unique "Smart Scan" capability, which performs SQL processing in the storage server, and vastly reduces the amount of IO traffic between the storage servers and the database servers.

Oracle Exadata Storage Server

The Oracle Exadata Storage Server is a database storage device powered by Exadata Storage Server Software running on Sun hardware. The hardware of the Exadata Storage Server was carefully chosen to match the needs of high performance query processing. The storage server comes preconfigured with: two Intel six-core processors, twelve disks, four flash cards, dual port InfiniBand connectivity, a management card for remote access, all the software preinstalled, and can be installed in a typical 19-inch rack.

Oracle Exadata Storage Server

Oracle Exadata Database Machine

The Oracle Exadata Database Machine is an extreme-performance data warehouse built using Exadata Storage Servers and state-of-the-art industry-standard hardware from Sun. Designed for large, multi-terabyte data warehouses with I/O-intensive workloads, the Oracle Exadata Database Machine is a complete, pre-optimized and preconfigured package of software, servers, and storage that is simple and fast to implement. The introduction of Oracle Exadata Database Machine X2-8, which comes complete with two 8-socket database servers, 14 Oracle Exadata Storage Servers and InfiniBand switches, extends the family of Exadata Database Machines, offering customers an ideal consolidation platform for very large OLTP and Data Warehousing applications.

Oracle Exadata Database Machine X2-2 (formerly known as V2) includes eight 6-socket database servers, 14 Oracle Exadata Storage Servers and InfiniBand switches, and is available in a choice of configurations that scale from a quarter rack to a full 42-unit rack

The Oracle Exadata Database Machine is built to accommodate up to 150TB of non-compressed user data and includes the following hardware:

- X2-2: Eight SunFire X4170 M2 Oracle Database Servers

- X2-8: Two SunFire X4800 Oracle Database Servers

- Both X2-2 and X2-8:

  - Fourteen Exadata Storage Servers (either SAS or SATA)

  - All the required InfiniBand infrastructure (HCAs, switches, and cables) for database server to Exadata Storage Server communication

  - Ethernet switch for communication from the Database Machine to clients or other computing systems

  - Keyboard, Video or Visual Display Unit Mouse (KVM) hardware

  - And this is all packaged in a single standard 19-inch 42U rack

Utilizing a building-block methodology, the Oracle Exadata Database Machine provides a quick and easy way to scale. As new racks of Oracle Database Machines are incrementally added to a system, the storage capacity and performance of the system grows; a two-rack system is simply twice as powerful as a single rack. Scaling out is easy; the additional Oracle Database Machine is connected to the InfiniBand interconnect in existing racks, and Oracle automatically rebalances the database to fully utilize all of the storage and processing power of all racks.

For smaller configurations, Oracle provides Oracle Exadata Database Machine X2-2 Half Rack and Quarter Rack, the Half Rack being a system that is exactly half the size of full-rack Database Machine with four database server and seven Exadata Storage Servers, while the Quarter rack is a system with two database servers and three Exadata Storage Servers.

The Exadata family delivers the scalable hardware capabilities to provide the required bandwidth for high-end data warehousing applications. This solution complements the sophisticated software of Oracle Database 11g, with its broad set of query-processing algorithms, advanced analytics, and robust data integration capabilities that will now be discussed.

Perform

Oracle provides performance optimizations for every type of data warehouse environment. Data warehouse workloads are often complex, with different users running vastly different operations, with similarly different expectations and requirements for query performance.

Oracle meets the demands of data warehouse performance by providing:

- A broad set of performance optimization techniques for every type of query and workload:

  - Advanced indexing and aggregation techniques for sub-second response times for reporting and dashboard queries

  - Star query optimizations and OLAP-based materialized views for dimensional queries

  - Scalable parallelized query access methods for processing large volumes of detail data

  - High IO bandwidth, with storage-level database optimizations, in Exadata

- A sophisticated resource manager for ensuring performance even in databases with complex, heterogeneous workloads. The Database Resource Manager allows end-users to be grouped into 'resource consumer groups', and for each group, the database administrator can set policies to govern the amount of CPU and IO resources that can be utilized, as well as specify policies for proactive query governing, and for query queuing. With the Database Resource Manager, Oracle provides the capabilities to ensure that data warehouse can meet their performance SLA's for all users.

- A unique, patented read-consistency model to ensure that data loads never impact query performance. Oracle solves the challenges of concurrent access through a technology called

multi-version read consistency; this unique technology has been the foundation of Oracle's concurrency model for over 15 years. Multi-version read consistency guarantees that a user always sees a consistent view of the data requested. If another user changes the underlying data during the query execution (such a trickle-feed update of a large data warehouse table), Oracle maintains a version of the data as it existed at the time the query began. The data returned to the query always reflects the state of the database (including all committed transactions) at the point in time at which the query was submitted regardless of what other updates may be occurring while the query is running. With this technology, Oracle is uniquely positioned to handle near real-time data loads within data warehouse environments.

- Management Packs to automate the ongoing performance tuning of a data warehouse.

With Oracle Exadata there are three new innovative performance capabilities:

Exadata "Smart Scans"

With traditional storage, all database intelligence resides on the database servers. However, Exadata has database intelligence built into the storage servers. This allows database operations, and specifically SQL processing, to leverage both the storage servers and database servers to vastly improve performance. The key feature is "Smart Scans", the technology of offloading some of the data-intensive processing into the Exadata Storage Server. Most large-scale data warehouse SQL queries operate only upon a subset of the rows and columns in the tables (for example, a query that looks at sales transactions for customers who purchased a specific set of products). With Smart Scans, only the relevant amount of the rows and columns will be sent back to the database servers to compute the final results, a normally relatively small amount of data. , Using Storage Indexes, "Smart Scans" will identify the relevant data sections without accessing the whole table at all; "Smart scans" then execute the basic row-filtering (the evaluation of where-clause predicates) and column-filtering (the evaluation of the select-list) on Exadata storage server, reading the relevant data sections either from disk or from flash. If a query requires the scan of a 1TB table, less than 1TB has to be physically scanned by the Exadata storage servers, and only the relevant data (which will often be a few GB's or less) is sent to the database servers for further SQL processing. "Smart scans" can improve the query performance of large queries by an order of magnitude, and in conjunction with the vastly superior IO bandwidth of Exadata's architecture delivers industry-leading performance for large-scale queries.

Exadata Storage Index

Completely automatic and transparent, Exadata Storage Indexes maintain column's minimum and maximum values of tables residing in the storage server. When a query is submitted to select data with a certain condition, the Exadata software examines the Storage Index to determine if

the column value from the query is outside of the minimum and maximum range to avoid a scan I/O operation. Many query operations will run dramatically faster because large numbers of I/O operations are automatically replaced by a few lookups.

Exadata Hybrid Columnar Compression

The Oracle Exadata Storage servers also enable new hybrid columnar compression technology that provides up to a 10 times compression ratio, without any loss of query performance. And, for pure historical data, a new archival level of hybrid columnar compression can be used that provides up to 40 times compression ratios.

Hybrid columnar compression is a new method for organizing how data is stored. Instead of storing the data in traditional rows, the data is grouped, ordered and stored one column at a time. This type of format offers a higher degree of compression for data that is loaded into a table for data warehousing queries. Different compression techniques can also be used on different partitions within a table.

In-memory parallel execution

Parallel execution is one of the fundamental technologies that enable organizations to manage and access tens or even hundreds of terabytes of data. It is the ability to apply multiple CPU and IO resources to the execution of a single database operation. While the Oracle Database has always leveraged memory for improved query performance via the buffer cache and other techniques, the increasing number of nodes in grid-enabled real application clusters environments and lower cost of memory is dramatically increasing the amount of memory available for large-scale data warehouses, with the total amount of memory reaching hundreds of GBs, up to TBs. Oracle's in-memory parallel execution leverages the total amount of memory available in such large systems, providing cluster-wide in-memory parallelization. At runtime, every active node in a cluster will cache a subset of the data being processed in parallel; subsequent parallel operations accessing the same large objects will leverage the cache of all active nodes. This not only minimizes the necessary IO to satisfy a given business query, it also improves the runtime for individual queries significantly, improving the overall workload a system can handle. In-memory parallel execution is complementary to the Exadata "smart scan" technology and combines the best of all worlds: the database will automatically only consider objects for in-memory parallel execution that either fit completely or to a large extent into the available memory; whenever the data volume exceeds the available memory significantly, Exadata's "smart scan" processing will take place, because smart it is more beneficial than processing the operation partially in memory.

OLAP-based materialized views

Conventional performance-tuning methods for star schemas within data warehouses have long involved the use of summary tables (or materialized views). Using this approach, the

performance of queries is improved by pre-computing one or more summaries and storing those summaries in the data warehouse. These summaries are completely transparent to the application, since the database provides 'query rewrite' capabilities in which a query against a base set of star-schema tables is rewritten to access summary data. One challenge with materialized views is that a single star schema can have thousands or even millions of possible aggregations that could be stored as summary tables. It is not practical to create a materialized view for every possible aggregation; the space utilization and maintenance costs would be too high. However, this scenario can be resolved with OLAP-based materialized views. An OLAP-based materialized view leverages the OLAP data structures, in which all possible aggregations can be quickly retrieved. Thus, a star schema with an OLAP-based materialized view provides the performance of a star schema with thousands of relational materialized views … but with a single easy-to-manage OLAP cube that is highly compressed and provides efficient update capabilities. This application of OLAP technology to general-purpose data warehouse workloads is unique to Oracle. While previous OLAP solutions were accessed via an OLAP-based API, Oracle OLAP delivers the performance benefits of OLAP to all data warehouse users using any SQL-based tool.

Scale

The size of the largest data warehouses is growing exponentially. Today, more and more business processes are becoming more completely automated and more data is collected at more granular levels, so these data volume increases show no signs of abating. Oracle Database 11g offers four key capabilities to enable scalability: partitioning, compression, clustering, and parallelism

Partitioning

Oracle Partitioning is essential for managing large databases. It enables a "divide and conquer" technique for managing the large tables in the database, especially as those tables grow.

Although your database may have twice as much data next year as it does today, your end-users are not going to tolerate their application running twice as slow, your database is not going to be given twice as much time to complete maintenance and batch processing, and your IT managers are not going to double the hardware budget for the data warehouse. Partitioning is the feature that allows a database to scale for very large datasets while maintaining consistent performance, without unduly increasing administrative or hardware resources. Partitioning breaks up large tables up into smaller pieces, and thus allows data management windows and many end-user queries to be maintained at constant performance level even as the data grows.

Oracle leads the industry with comprehensive partitioning technologies, with nine methods for partitioning tables, along with the capability for DBA's to define custom partitioning schemes; a

rich set of administrative commands for partitioned tables; and a partition adviser to guide administrators on how best to implement partitioning.

Partitioning also enables ILM ("Information Lifecycle Management") strategies within the Oracle database. A single table, when partitioned, can be distributed across multiple storage tiers. Old, less-frequently accessed data, corresponding to older partitions, can be stored on less expensive storage devices. For large databases, this approach can provide significant cost savings.

Compression

Compression capabilities are used within every large data warehouse. As customers look to store larger volumes of data, compression is a natural solution. Oracle 11g Release 2 takes database compression to the next level with new compression techniques. Oracle provides the ability to compress all tables in a data warehouse:

- "OLTP" Compression: Oracle's original table compression technology this technique delivers a typical compression ratio of 3:1 for data warehouses, with virtually no negative impact on query performance. This compression technique enables efficient updates to support compression for even data warehouse tables which are 'trickle-fed' or otherwise updated frequently.

- "Warehouse" Compression: Based upon Exadata Hybrid Columnar Compression, this technique delivers a typical compression ratio of up to 10:1 for data warehouses, with virtually no negative impact on query performance.

- Archive Compression: Based upon Exadata Hybrid Columnar Compression, this technique delivers a typical compression ratio of up to 40:1 for data warehouses, but does entail trade-offs in query performance.

Real Application Clusters

Real Application Clusters enable a single database to scale across multiple servers. For data warehousing, RAC provides a solution to 'scale-out' to grow to hundreds of CPU's. The architecture of RAC provides unique benefits for Oracle data warehousing. In addition to the capability of adding additional raw computing power to a data warehouse system, RAC also enables high-availability against node failure.

RAC also delivers a tremendous amount of flexibility for managing multiple workloads within a single database. Using 'RAC services', different applications can be assigned to separate services, which in turn can run on separate nodes of a cluster. For example, data mart users could be assigned to one set of nodes while ETL operations could be assigned to another set of nodes. A key advantage of RAC is the ability to dynamically change and adjust these workloads. One could easily and dynamically allocate more nodes to ETL processing at nighttime, for example, while re-allocating those nodes to query processing during the day.

Parallelism

Parallelism is the ability to leverage multiple processors towards the execution of a single database operation. Oracle's scalable parallel execution architecture dynamically adjusts to meet all user demands. For large operations, Oracle parallelizes database operations across all available nodes and CPU's. For smaller operations, Oracle dynamically adjust the parallelism to maximize overall system throughput while the Database Resource Manager ensure that high-priority operations continue to get the appropriate resources.

Analyze

Technologies such as OLAP, statistics, and data mining are hardly new to data warehousing and business intelligence. However, OLAP products typically have their own calculation engine, statistics products have their own data engine, and data mining products have their own mining engines. In short, an enterprise-wide business intelligence environment would maintain at least four different types of 'data engines', each requiring their own servers, their own copies of the data, their own management infrastructure, their own security administration, and their own high-availability infrastructure. Each engine has its own API's and its own set of developer tools and end-user tools. The complexity and cost of replicating entire stacks of BI technologies is significant.

Oracle Database provides a completely different approach by integrating OLAP, Data Mining, and statistics inside the database engine. Instead of moving data from a data warehouse to other analytic engines for further analysis, Oracle has instead brought the advanced analytic algorithms into its database, where the data resides. Moreover, Oracle provides SQL access to all of its analytics, so that they can be implemented with any SQL-based tool or application environment.

Beyond the considerable advantages of consolidating the back-end data architecture of an enterprise business intelligence environment, the integration of analytics within the Oracle Database provides a host of advantages unavailable to stand-alone environments. For example, does your standalone OLAP server scale to hundreds of CPU's or clusters of servers? How easily does your data-mining engine integrate into your user authentication server? And can it transparently implement all of your data security policies? How easily can you integrate the results of your statistical analysis with your data warehouse data? Within Oracle Database, all of these issues are solved simply due to the deep integration of OLAP, Data Mining and statistics.

Data Mining

Oracle Data Mining is powerful software embedded in the Oracle Database that enables you to discover new insights hidden in your data. Oracle Data Mining helps businesses to target their best customers, find and prevent fraud, discover the most influential attributes that affect Key Performance Indicators (KPIs), and find valuable new information hidden in the data. Oracle

Data Mining helps technical professionals find patterns in their data, identify key attributes, discover new clusters and associations, and uncover valuable insights.

To address these business problems, Oracle Data Mining allows companies to find new information from their data using a wide range of state-of-the-art algorithms. Data mining algorithms are machine-learning techniques for analyzing data to discover patterns and relationships. Oracle provides multiple algorithms since different algorithms are effective for different types of analysis and different business problems.

Most data mining algorithms can be broadly separated into "supervised learning" and "unsupervised learning" data mining techniques. Supervised learning requires the data analyst to identify a target attribute or dependent variable (for example, customers who bought a specific product). The supervised-learning technique then sifts through data trying to find patterns and relationships between other attributes and the target attribute (for example, the characteristics that indicate whether a prospective customer is likely to buy a specific product). Supervised learning algorithms with Oracle Data Mining include Naïve Bayes, Decision Tree, General Linear Models, and Support Vector Machines.

The other broad category of data-mining algorithms is for "unsupervised learning." In these scenarios, there is no 'target attribute'; instead the data mining algorithms seek to find associations and clusters in the data independent of any a priori defined business objective. These algorithms include Enhanced k-Means Clustering, Orthogonal Partitioning Clustering, Association Rules (market basket analysis), and Nonnegative Matrix Factorization.

Oracle Data Mining includes Oracle Data Miner, a graphical user interface for data analysts to build, evaluate, and apply data mining models. Oracle Data Miner guides the data analyst through the data mining process with complete flexibility and presents results in graphical and tabular formats. Oracle Data Miner can generate the PL/SQL code associated with a Mining Activity.

The Smart Scans feature of Exadata offloads all data mining scoring functions; as a consequence, the CPU utilization of the Database server is reduced and you can observe a performance gain up to 10x.

OLAP

Oracle OLAP is a full-feature online analytical processing (OLAP) engine embedded in the Oracle Database. Oracle OLAP enhances data warehouses by improving query performance (as discussed in the performance section) and by adding enriched analytical content.

The core feature of Oracle OLAP is cubes. Managed within the Oracle database, this data structure stores data within a highly optimized multidimensional format. Cubes provide scalable and compressed storage of dimensional data, fast incremental update, fast query performance, and the ability to compute or store advanced analytical calculations.

Oracle's strategy with Oracle OLAP is to bring these core OLAP advantages into the data warehouse. This is achieved be exposing the key capabilities of Oracle OLAP via standard SQL, so that any business intelligence tools or other SQL-based application can leverage OLAP.

The key SQL integration features include the ability to access OLAP cubes via SQL views, which expose a multidimensional cube as a set of dimension tables and fact tables; cube-based materialized views which allow OLAP cubes to be used as materialized views; data-dictionary integration so that SQL-based query tools as well as SQL developers can quickly identify all of the cubes and their attributes; and the integration of cube access paths within the Oracle optimizer.

Integrate

Today's information architecture is much more dynamic than it was just a few years ago. Businesses now demand more information, they want it sooner and they are delivering more analytics to an every-widening set of users and applications.

Oracle Warehouse Builder and Oracle Data Integrator

To address these business requirements, Oracle Database 11g includes a leading data-integration tool, Oracle Warehouse Builder (OWB). OWB is included as a no-cost database feature, and is designed to allow any Oracle customer to efficiently build a data mart or data warehouse, of any size. It includes an enterprise-ready multi-user metadata repository, data-modeling capabilities, and a wide variety of transformation and extraction techniques, and the performance and scalability of an 'ELT' architecture.

OWB's advanced features, previously part of the Enterprise ETL Option, have been merged into Oracle Data Integrator Enterprise Edition. Oracle Data Integrator provides not only 'ELT' capabilities most often required by data warehouse, but also delivers an enterprise-wide integration platform to integrate for example SOA-enabled data-services or event-driven data feeds. Today, Oracle includes Oracle Data Integrator and Oracle Warehouse Builder's advanced features as two components of Oracle Data Integrator Enterprise Edition. Going forward, these products will merge into a single unified data integration technology platform. This strategy fully preserves any existing development investments of all Oracle data integration customers and will provide a seamless, easy upgrade path from the current components to the unified platform.

Key database integration features

Oracle Warehouse Builder uses an extract-load-transform ('ELT') architecture. Rather than providing an external data-transformation engine, OWB executes all of its transformations within an Oracle database leveraging the scalability and performance of the database platform. The Oracle database has a broad set of extraction, loading and transformation capabilities. These

features are leveraged by OWB, but can be used in any environment in which data is being integrated into an Oracle data warehouse. These features include:

- Database Gateways for accessing non-Oracle systems

- Loader utility for fast data loads of flat files

- SQL extensions for data transformations: MERGE statement

- Table functions: efficient parallel user-defined transformations

- Change data capture for low-latency log-based capture from Oracle databases

Conclusion

The Oracle Database is the market leader for data warehousing, built upon a solid foundation of scalability and performance, and augmented by innovative features such as Oracle's unique read-consistency model for near-real-time data warehouses, a flexible and powerful set of table partitioning capabilities, the utilization of OLAP technology to enhance relational environments, and the unmatched support for grid architectures. The combination of the Oracle Database and an Oracle Exadata storage grid delivers the highest levels of performance for IO intensive workloads, and, with the Oracle Exadata Database Machine, Oracle delivers a complete hardware and software solution for data warehousing.

ORACLE®

Oracle Database 11g for Data Warehousing and
Business Intelligence
April   2011
Author: George Lumpkin

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

Oracle is committed to developing practices and products that help protect the environment