

The value of the data warehouse is its ability to support business intelligence. To date, standard reporting and ad-hoc query and reporting applications have run directly from relational tables while more sophisticated business intelligence applications have use specialized analytical databases. These specialized analytical databases typically provide support for complex multidimensional calculations and predictive functions, however they rely on replicating large volumes of data into proprietary databases.

Replication of data into proprietary analytical databases is extremely expensive. Additional hardware is required to run analytical databases and store replicated data. Additional database administrators are required to manage the system. The replication process often causes a significant lag between the time data becomes available in the data warehouse and when it is staged for analysis in the analytical database. Latency caused by data replication can significantly affect the value of the data.

Oracle OLAP provides support for business intelligence applications without the need for replicating large volumes of data in specialized analytical databases. Oracle OLAP allows applications to support complex multidimensional calculations directly against the data warehouse. The result is a single database that is more manageable, more scalable and accessible to the largest number of applications.

Business intelligence applications are only useful when they are easily accessed. To support access by large, distributed user communities Oracle OLAP is designed for the Internet. The Oracle9i Java OLAP API provides a modern Internet-ready API that allows application developers to build Java applications, applets, servlets, and JSPs that can be deployed using a variety of devices such as PCs and workstations, web browsers, PDAs, and web enabled mobile phones.

Platform for Business Intelligence Applications

Oracle9i Database provides a comprehensive platform for business intelligence applications. The components of the platform include:

- The Oracle9i Database
- Oracle OLAP as a facility within Oracle9i Database

This platform provides:

- A complete range of analytic functionality including multidimensional and predictive functions
- Support for rapid query response times such as those that are normally associated with specialized analytical databases
- A scalable platform for storing and analyzing multi-terabyte data sets
- A platform that is open to both multidimensional and SQL based applications
- Support for Internet based applications

Oracle9i Database

The Oracle9i Database provides the foundation for Oracle OLAP by providing a scalable and secure data store, summary management facilities, metadata, SQL analytic functions, and high availability features.

Scalability features that provide support for multi-terabyte data warehouses include:

- Partitioning allows objects in the data warehouse to be broken down into smaller physical components that can then be managed independently and in parallel
- Parallel query execution allows the database to use multiple processes to satisfy a single Java OLAPI API query
- Support for NUMA and clustered systems allows organizations to effectively use and manage large hardware systems
- Oracle's Database Resource Manager helps manage large and diverse user communities by controlling the amounts of resources each user type is allowed to use

Summary Management

Materialized views provide facilities for effectively managing data within the data warehouse. As compared to summary tables, materialized views offer several advantages:

- They are transparent to applications and users
- They manage staleness of data
- They can automatically update themselves when source data changes

Like Oracle tables, materialized views can be partitioned and maintained in parallel. Unlike proprietary multidimensional cubes, data in materialized views are equally accessible by all applications using the data warehouse.

Security

Security is critical to the data warehouse. To provide the strongest possible security and to minimize administrative overhead, all security policies are enforced within the data warehouse. Users are authenticated in the Oracle database using database authentication or Oracle Internet Directory (LDAP). Access to elements of the multidimensional data model is controlled through grants and privileges in the Oracle database. Cell level access to data is controlled in the Oracle database using Oracle's Virtual Private Database feature.

Metadata

All metadata is stored in the Oracle database. Low level objects such as dimensions, tables and materialized views are defined directly from the Oracle data dictionary while higher level OLAP objects are defined in the OLAP catalog. The OLAP catalog contains objects such as Cubes and Measure folders as well as extensions to the definitions of other objects such as dimensions. The

OLAP catalog fully defined the dimensions and facts and thus completes the definition of the star schema.

Disaster Recovery

Oracle's disaster recovery features protects data in the data warehouse. Key features include:

- Oracle Data Guard, a comprehensive standby database disaster recovery solution
- Redo logs and the recovery catalog
- Backup and restore operations that are fully integrated with Oracle's partition features
- Support for incremental backup and recovery

Oracle OLAP

Oracle OLAP, an integrated part of Oracle9i Database, provides support for multidimensional calculations and predictive functions. Oracle OLAP supports both the Oracle relational tables and *analytic workspaces* (a multidimensional data type).

Key features of Oracle OLAP include:

- The ability to support complex, multidimensional calculations
- Support for predictive functions such as forecasts, models, non-additive aggregations and allocations, and scenario management (what if)
- A Java OLAP API
- Integrated OLAP administration

Multidimensional calculations allow the user to analyze data across dimensions. For example, a user could ask for "the top ten products for each of the top ten customers during a rolling six month time period based on growth in dollar sales". In this query a product ranking is nested within a customer ranking, data is analyzed across a number of time periods and a virtual measure. These types of queries are resolved directly in the relational database.

Predictive functions allow applications to answer questions such as 'how profitable will the company be next quarter' and 'how many items should be manufactured this month'. Predictive functions are resolved within a multidimensional data type known as an analytic workspace using the Oracle OLAP DML (the OLAP DML could be thought of as being the multidimensional equivalent to PL/SQL).

Oracle OLAP uses a multidimensional data model that allows users to express queries in business terms (what products, what customers, what time periods, and what facts). The multidimensional model includes measures, cubes, dimensions, levels, hierarchies, and attributes.

Java OLAP API

The Oracle9i OLAP API is based on Java. As a result it is an object-oriented, platform independent, and secure API that allows application developers to build Java applications, Java Applets, Java Servlets, and Java Server Pages (JSPs) that can be deployed to large, distributed user communities over the Internet. Key features to the Java OLAP API include:

- Encapsulation
- Support for multidimensional calculations
- Incremental query construction
- Multidimensional cursors

Performance

Oracle9i Database eliminates the tradeoff between analytical complexity and support for large databases. On smaller data sets (where specialized analytically databases typically excel) Oracle9i provides query performance that is competitive with specialized multidimensional databases. As databases grow larger and as more data must be access in order to resolve queries Oracle9i will continue to provide excellent query performance while the performance of specialized analytical databases will typically degrade.

Oracle9i Database achieves both performance and scalability through SQL that is highly optimized for multidimensional queries and the Oracle database. Accessing cells of data within the multidimensional model is key to providing query performance that is competitive with specialized analytical databases. New features in the Oracle database that provide support high performance random cell access and multidimensional queries includes:

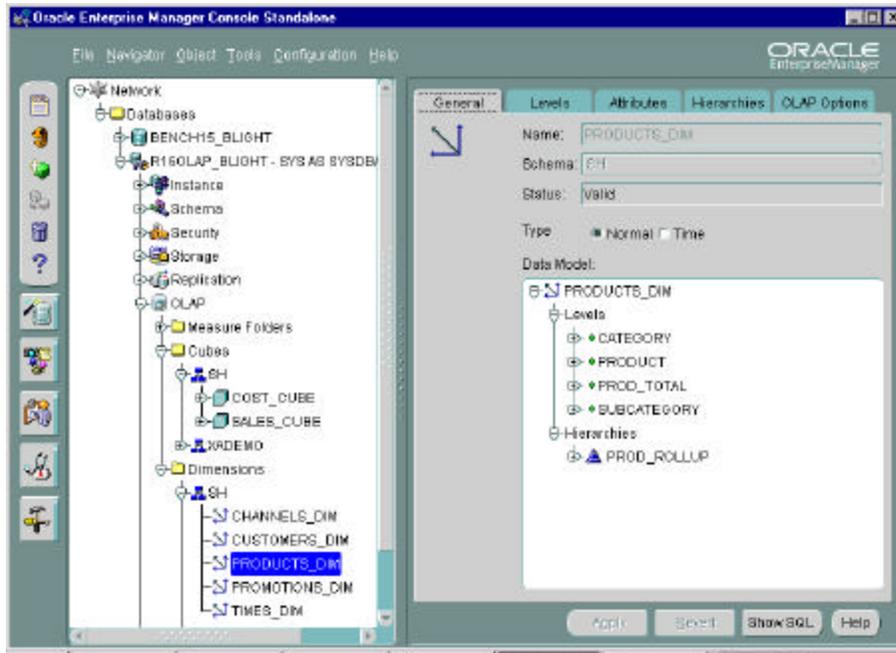
- Bitmap join indexes which are used in the warehouse to pre-join dimension table tables and fact tables and stores the result in a single bitmap index
- Grouping sets allow Oracle to select data from multiple levels of summarization in a single select statement
- The WITH clause allows Oracle to create temporary results and use these results within the query, thus eliminating the need for creating temporary tables
- SQL OLAP functions provide highly concise means to express many OLAP functions
- Automatic memory management features provide the correct amounts of memory during memory intensive tasks
- Enhanced cursor sharing eliminates the need to recompile queries when another similar query has been run.

System Management

Oracle Enterprise Manager (OEM) provides a centralized, comprehensive management tool. OEM enables administrators to monitor all aspects of the database, including Oracle OLAP.

Oracle Enterprise Manager provides management services to Oracle OLAP including:

- Instance, session, and configuration management
- Data modeling
- Performance monitoring
- Job scheduling



OLAP tool outline (left), dimension detail (right)

System Requirements

Oracle OLAP is installed as part of the Oracle9i Database and imposes no additional system requirements. Oracle OLAP can also be installed on a middle tier system. When installed on a middle tier system, 128 MB of memory is required. When analytic workspaces are used extensively, additional memory is recommended. The actual amount of memory for use with analytic workspaces will vary with your application.

