

Harnessing the Power of Scientific Data Warehouses

Kevin Deeb, Ph.D.

Department of Information Technology, Barry University
11415 NE 2nd Ave, Miami Shores Florida, 33161

kdeeb@mail.barry.edu

ABSTRACT

Data warehousing architecture should generally protect the confidentiality of data before it can be published, provide sufficient granularity to enable scientists to variously manipulate data, support robust metadata services, and define standardized spatial components. Data can then be transformed into information that would make them readily available in a common format that is easily accessible, fast, and bridges the islands of dispersed information. The benefits of the warehouse can be further enhanced by adding a spatial component so that the data can be brought to life, overlapping layers of information in a format that is easily grasped by management, enabling them to tease out trends in their areas of expertise.

Keywords: Warehouse, Data retrieval, Spatial data components, Oracle, ETL.

1. INTRODUCTION

Data are meaningless until they are transformed into information. Over the years many organizations have collected millions of rows of data that, when properly analyzed, translate into useful and profitable information that will assist in deciding the course of future market strategies, the creation of new products, or the retooling of existing production lines to meet consumer demands. In the scientific community, these data are often used to set new policies that can later be made into laws to protect the environment, promote advances in the medical field, or find new natural resources. Since these vast numbers of data can take many hours to process, the databases that house them must meet certain requirements that are often at odds with traditional online transaction processing databases (i.e., OLTP) which require real time access and validation. Hence the birth of the data warehouse, a database that provides mechanisms for synchronizing and updating information obtained from OLTP databases and other external sources, as well as supporting the performance requirements to process and load large volumes of data. The benefits of the warehouse can be further enhanced by adding a spatial component so that the data can be brought to life, overlapping layers of information in a format that is easily grasped by management, enabling them to tease out trends in their areas of expertise. The current necessity for database management systems (DBMS) to offer simplified yet reliable solutions for users of spatial data has driven the unification of DBMS with spatial databases. Traditional databases support regular data types, such as character, varchar, integer, datetime, etc. Nevertheless, spatial databases require more complex data types like points, lines, and polygons, that are used to represent geographical locations on Earth. Consequently the need for new technology qualified to offer these capabilities emerged.

2. DATA WAREHOUSE

Normalized data models evolved to solve the needs of online transactional processing (OLTP) systems that focused on speed of data entry, ease of editing, protection of data integrity by reducing redundancy, and immediate point of entry validation. This model served to solve the problems of the batch processing environment, in which data were updated through batches of records that included codes to add, change, or delete from the master flat file. Whenever coding changes were required, the IT group, then called "data management," had to be methodical enough to make sure all occurrences of the offending code were replaced with the new value.

Even though the old flat files were difficult to quality assure, they were very easy for most users to query. Reports generally represented little more than sums across records of data with some additional programming to convert cryptic codes into their more legible legends. Conversely, the OLTP model is perceived as a nightmare for most users who must now write multi table joins across seven or eight tables. Suddenly, the old flat file legacy systems seemed much more desirable. To compound matters, while information technology groups spent their resources designing and deploying elegant data models touting the ease of use of SQL (i.e., Structured Query Language), the user community was left unable to understand the complexities of SQL when confronted with highly normalized designs, and represented to their own management that their data was no longer accessible.

Clearly, there is a need for some compromise between the elegance and efficiency of the OLTP systems, and the needs of the end users to analyze the data that cost so many dollars to capture and ingest. The solution to this dilemma is the data warehouse. Although it raises some brows with apprehension, the data warehouse can be created using the same RDBMS used to house the OLTP database. In fact, a data warehouse is nothing more than a database that has been optimized for retrieval. The main concern of its architects is the delivery of data in a consistent, easy to access format. The emphasis of the design is speed of retrieval.

Section 2 gives an overview of data warehouses and their use in the corporate arena. Section 3 presents several factors to improve the efficiency of data warehouses. Section 4 continues to discuss current challenges in spatial databases while proposing suitable solutions and addresses techniques to populate databases with GIS data. Section 8 presents data retrieval alternative solutions to serve data to the user community. The following sections inspect popular enterprise databases that fit in the overall purpose of GIS systems and data objects.

3. DATA WAREHOUSE EFFICIENCY FACTORS

As mentioned earlier, a data warehouse must provide good performance, manage large amounts of data, and provide quick ways in which to load large volumes of OLTP data. To support these operations, the manufacturers of data warehousing software have created the following strategies:

Range, Hash, and Composite partitioning: Partitioning involves separating objects into pieces so that they can be managed more easily. Although to the database the object is logically the same, each piece is physically stored separately. Range partitioning is useful when separating data in ranges, such as by year, or by state. On the other hand, when the data cannot be separated into discrete meaningful groups, the data can be separated using a hashing algorithm that ensures that the data will be distributed evenly. The combination of the two partitioning methodologies is called composite partitioning and it uses range partitioning first, while dividing the results using hash partitioning to create sub-partitions [13]. With any of these techniques, the database designer must analyze both the kinds of data that will be stored in the warehouse, as well as the ways in which the data will be used. Partitioning data based on year may not be effective if the organization typically analyzes data aggregated by state, crossing over multiple years.

Transportable tablespaces: Tablespaces are logical units that hold database objects. Each tablespace may be composed of one or many physical data files. Since data warehouses are used to manage large numbers of data, it may be necessary to move these tablespaces to supply data to a different database or data mart, archive data, or to share data with other databases. Because the tablespaces are based on physical files, there are certain limitations to this capability. These limitations include the requirement that tablespaces be transported only between databases on the same platform, the platform's block size must be the same, and they must use the same character set [13].

The star schema: Data warehouses are frequently designed using a dimensional data model called the star schema. In this type of design, there is a central table called a fact table that is related to several look-up tables called dimensions [15]. Fact tables tend to be very large with millions of records and gigabytes to terabytes of data and contain quantitative data. On the other hand, dimension tables tend to be much smaller and contain descriptive data. Both must have a primary key, which for fact tables are usually composites of the foreign keys to the dimension tables. To access tables in this formation, a star query is used, which in turn uses what is referred to as a star join [13]. Most RDBMS (i.e., Relational Database Management System) allow developers to use a mechanism called "cost-based optimizer" to allow the computer to optimize the query performance. This is done by regularly compiling statistics on the data and then specifying the "star" hint, which if possible, will cause the database engine to position the largest table, or fact table, last in the query to reduce the number of rows read with each join [13].

Summary Management and Materialized views: Two important problems faced by data warehousing designers are the aggregation of data across multiple dimensions to hide the complexity of queries, and ensuring that the data are kept up to date when they must be available at all times. Through the use of summary management, performance can be greatly improved by keeping joined and summarized data in objects called

materialized views [13]. In general, a view is a different way in which to look at data. It can be as simple as a single table in which some columns have been renamed or dropped, or as complex as the combination of multiple tables and the creation of summaries. The problem with views is that when they are accessed, the query that they represent must be performed and, depending on its complexity, they can take a long time to execute. On the other hand, materialized views physically store data in this joined and aggregated format, so that when a query is performed, the execution time is minimized to the amount of data stored in the materialized view. With sophisticated RDBMS tools, after statistics have been gathered, the database can recommend which views should be materialized and will even re-write queries if there are materialized views to satisfy a query. The summary management component provides methods to update and propagate the data that make up the materialized view whenever the base tables are updated.

Before defining the architecture of the geospatial data warehouse it is important for analysts to locate existing stores of data. Especially in large organization, islands of information tend to form and similar data are stored in different formats making the ingestion of data a laborious and time consuming task that detracts from the business of science. Defining the scope of a design effort is critical, since it will lay the ground work for the ultimate design of the warehouse.

4. CHALLENGES OF A SPATIAL DATABASE

Spatial databases require the use of more complex data types than traditional databases, these new data types introduce an array of more complexities concerning spatial databases. These are specific details that must always be taken into consideration in order to avoid a much bigger predicament. Not providing the necessary concern to these attributes will result in more intricate problems.

The key asset in a spatial database is geometry. Geometry is what makes a spatial database unique from traditional databases. Object measurements, relationship points, lines, angles, surfaces, etc, are geometrical properties used to describe spatial data. The geometry of spatial databases then becomes more complex when objects are assembled together to make larger objects, such as neighborhoods, cities, states, etc.

This assembly of geometrical objects becomes the source of the intricacy of the distribution of objects in space. The density of the location of these objects on Earth is completely unbalanced, thus making this task more difficult. For example, the population in the Miami Beach area is much more dense compared to the population in The Everglades, which is more sparse.

The large volume of data associated with GIS is yet another challenge for spatial databases. "For example, remote sensing applications gather terabytes of data from satellites every day" (Sharma). Consequently, organizations must ensure the dependability and reliability of their spatial databases.

The scientific community is becoming increasingly dependent on Geographic Information Systems (GIS). According to J. Michael Fay, GIS experts hold the "key to the future," because by understanding the landscape they can effect positive change on the environment [1, 2, 3]. Furthermore, Jack Dangemond,

President of ESRI, states that "[t]he application of GIS is limited only by the imagination of those who use it" [4, 5]. In a nutshell, GIS involves applying layers of information [17], whether information about a location's temperature, depth, species of fish caught, number of fishing vessels, any information that can be combined to enable analysts to easily discern patterns in the data. The presentation *Geography Matters* (2002) explains that GIS is not just about maps, but it is the special relationship between data, location information, and the people that manage it. This makes the incorporation of spatial elements into the enterprise data warehouse not just a desirable but a necessary component.

Although GIS has been available for many years, acquiring spatial data layers has not always been simple. Furthermore, the existing GIS analytical tools required levels of expertise that could not be mastered easily by casual users requiring a simple geographical representation of their data. By making spatial shape files available at the enterprise level, scientists will be able to perform comparative studies and visually represent "what-if" scenarios without expending inordinate amounts of time in the process of ingesting data.

Another challenge for spatial databases is the temporary changes that objects on Earth endure. There are road changes or blocks due to construction, or flooding, which change the normal path of direction to a specific location. Then, when the road is reopened or the changes are reverted back to normal, the path of direction is changed once again.

Spatially enabling the warehouse is an ideal solution to provide GIS services to the enterprise. The many standardized layers such as bathymetry and land contours, are generally static and can be loaded one time and remain accessible to all data warehouse customers. This process will result in a great cost savings since currently, since scientists acquiring the data on their own, are loading it into their own databases, and are neither sharing the costs nor the expertise in solving the problem of data ingestion.

Unfortunately, according to Daniel R. Dolk in his paper *Introduction to Modeling Technology and Intelligent Systems Track*, GIS has been largely neglected by IT management and therefore has not benefited from MIS research (Dolk, 1999). Bringing spatial data into the data warehouse represents a unique opportunity for providing benefits to the users by enabling them to share shape files across the enterprise, but they will also benefit from the extensive expertise of their IT department that can leverage existing investments in data acquisition while optimizing the data warehouse.

Large database vendors such as Oracle with its 9i Enterprise RDMS offer a spatial option that enables the storage of spatially enabled data, that is, data that can be interpreted by native SQL and GIS tools. Line geometry can be stored in spatial data object data types that turn the spatial object into a regular field that can be queried using native SQL. For more complex shape files such as raster data, there is still no support. However, these shape files can still be stored in the data warehouse using binary data types, or, in specialized containers such as Oracle Intermedia, that allow third party products such as E-Spatial to access the data directly from the database. Data stored in binary data format can be easily read by such powerful GIS tools as those produced by the ESRI.

The reason it is important to keep data in a virtual central location (i.e., the warehouse) is that the problem of acquisition is solved a priori by the IT staff in conjunction with the data managers, and last minute research problems the week before presentations at symposia largely go away. By bringing GIS to the warehouse the entire enterprise benefits.

5. DESIGNING THE ETL

ETL or Extract, Transform, and Load procedures are the three steps necessary to populate the data warehouse. Although difficult to set up initially, once the methods are established, the business of synchronizing the transactional database with the warehouse occurs automatically, making the process transparent. Unfortunately, synchronization is slightly more complicated for the scientific community than it is in the business world. The most significant difference is that the sources of scientific data are frequently in the form of data submitted by the owners and is not immediately validated. Depending on the volume of data and the workload of the data management team, it may take from several weeks to several months to make the data ready for the warehouse. Even after the initial constraints are met, specialized analysts can identify outliers that must be corrected in the data warehouse. This presents a problem for ETL, because now the OLTP data needs to be compared with the warehoused data to ensure that the most current version is now available before it is loaded. Therefore, although older data may be archived in the warehouse, it needs to be compared periodically to ensure accuracy. Jaideep Srivastava and Ping-Yao Chen, authors of *Warehouse Creation - A Potential Roadblock*, consider the issue of data loading so important that they have named it a roadblock (1999). They have found that generally custom tools must be developed by specialized consultants to handle the ETL issue. It is important to note, however, that once the procedures are put in place, the process of synchronizing the warehouse with the OLTP database occurs automatically.

6. SERVING THE DATA: RETRIEVAL ALTERNATIVES

Using a data warehousing tool that is open and platform independent, affords architects and analysts the capability of opening the data to different tools. Once the data are warehoused, Online Analytical Processing tools can be used to access the data. It is especially important to note that a range of users can be serviced. From the less computer-savvy congressperson who needs broad summarizations to support the passing of legislation, to the expert GIS user who needs granular data to a tenth of a degree, all can benefit from the data stored in the warehouse.

One of the strengths of the data warehouse is its ability to empower end users to perform their own analysis. Historically, users were dependent on the information technology (IT) team to produce inflexible reports that sometimes took months to produce. If an analyst wanted to pivot the report (i.e., turn the report on its side so that the rows in the report become column headings), drill down to the records that contributed to summary data, or simply aggregate data in a different way, another request-for-service document needed to be written, approved, and then submitted to the IT queue until resources became

available. On-Line Analytical Processing tools enable users to access multidimensional data on their own, freeing IT resources for other more vital tasks.

Unfortunately, many data warehousing ventures fail because they lose focus on their ultimate goal, which is to serve data to the user community. Joe Kopetsky and Sally Ting of Arcplan Inc. identified three different success factors that will ultimately result in user buy-in: ability, can the stakeholders use the system; willingness, are the users willing to use the system, and knowledge, are the users familiar with how the system works (Kopetsky, 2002). Selecting an OLAP tool that meets these goals and makes the users feel confident that they can get to the data they want, the way they want it, will go a long way to ensuring customer satisfaction.

Serving GIS data poses its own challenges, and these must be met by giving access to the warehoused data at many different levels. There must be simple tools that will enable the entry of natural language queries so that geographic representation of data can be made to management without many hours of labor, while at the same time preserving the expert's ability to use the tools that they have become familiar with over the years.

7. DBMS GIS COMMERCIAL SOFTWARE

Oracle holds an 80-90% market share of the spatial database management market" (Oracle Website; IDC, Spatial Information Management: Competitive Analysis, 2002). If an organization is in need of data integrity, backup, recovery, and security, then they require the interoperability of Oracle Locator and Spatial. With the launch of Oracle Spatial, Oracle has ensured that all required to run a spatial database is encompassed in their product. "Oracle Spatial provides a completely open architecture for the management of spatial data within a database management system...Users define and manipulate spatial data through SQL and gain access to standard Oracle features, such as flexible n-tier architecture, object capabilities, robust data management utilities, [and] Java stored procedures" (Sharma).

Oracle provides the foundation required for spatial databases in GIS to produce the most reliable, efficient and accurate data needed by users of spatial data. "It takes the object-relational model as the base to store and query spatial information, and uses the principle of object oriented operation to associate several elements' information together to implement the spatial relation oriented query and other spatial operations" [18].

Oracle has created Oracle Locator and Spatial while complying with the OpenGIS Consortium Simple Features guidelines, "the new OGC Geographic Markup Language (GML), as well as Open Location Service and Web Map Server interfaces...[and] to the specifications associated with SQL92 representation of points, lines, and polygons" [10].

According to Oracle's vice president, Steve Hagan, "Oracle Spatial is used around the globe in thousands of enterprise and national GIS, digital mapping and wireless location-based services applications that are shared by and accessed through every major spatial and GIS partner technology" [7]. Oracle Locator and Spatial have the capability of being incorporated in any information systems applications, thus allowing many different organizations and its DBAs to become more effective and successful. Among numerous information systems,

Geographical Information (GIS) is one of the benefited by the advantages of using Oracle Locator and Oracle Spatial.

By making use of Oracle Locator and Spatial, companies are able to utilize data that is not currently being used to its potential. For example, disaster relief organization may be achieved by querying topographic data in areas susceptible to floods and population information in those areas. By utilizing Oracle Database with Locator or Spatial features, a company will consequently be able to respond to an emergency more quickly because they will have the tools to locate the area as well as emergency personnel closest to the area more quickly and thus respond to the emergency in a more efficient manner. GIS have the need to exploit very complex spatial databases by storing, querying and administering spatial data. "Oracle Database 10g with Oracle Locator and Oracle Spatial solves key...problems...such as how to:

- Efficiently store, retrieve, and manage location data and attribute data from a single, open database
- Improve performance for very large databases containing terabytes of location data
- Maintain versions of spatial database tables transparently, and identifying any conflicts in table values

Store location data once and make it accessible to users from multiple, heterogeneous GIS tools and e-Business applications without any modification via open, standards-based APIs" [10].

8. INTEROPERABILITY VS. DISPARATE DATA

One of the most significant features of Oracle Locator and Oracle Spatial is that Oracle has incorporated "spatial capability natively into the database" [10]. This allows storage, indexing, querying, security and accessing of spatial data all in one single database. "...Not only are all datasets maintained in a single Oracle database, but the users can also select the particular data they wish to view and query...[which] greatly improves the speed with which the individual query is completed" [1]. Consequently, redundancy is completely eliminated by allowing interoperability and the end of disparate data in separate databases.

There are several key benefits to spatial data interoperability:

1. Better data management for spatial data. Users gain access to full function spatial information systems based on industry standards with an open interface to their data (e.g. SQL)
2. Spatial data is now stored in enterprise-wide database[s], thereby facilitating spatially enabling many more applications
3. Reduced complexity of systems management by eliminating the hybrid or file based architectures of traditional GIS-based data management schemes
4. Proprietary data structures are avoided by using an open SQL platform thereby allowing for the seamless integration of e-business and location-based services. (Sharma).

Benefits of interoperability range from basically centralizing all data into one single database, to providing employees from all departments to access that data. Additionally, that data, which was spread between several departments, will consequently be managed and secured by only one department (IT). While several GIS vendors have partnered with Oracle to create interoperability kits, Oracle's objective is to achieve complete

interoperability without the need for these interoperability kits. Figures 7 and 8 illustrate the distinction between a system with interoperability and a system without interoperability. [12].

9. CONCLUSION

Data warehousing solves the problem of data accessibility. Unconcerned with transactional changes, the warehouse architects can concentrate instead on providing a streamlined pipeline to data that can be analyzed expertly. Scientists can benefit from data captured by other scientists without investing the time to locate and reformat the data to make it homogeneous. The warehouse architect does this for them. The traditional relational model, while adequate in managing online

10. REFERENCES

- [1] M. Bertolotto, J. Carswell, L. McGeown, J. McMahon. "iSmart +iSpatial Information Server: Deploying Integrated Web-Based Spatial Applications within an Oracle Database Environment", **Proceedings of the International Workshop on Web Geographical Information Systems (WGIS2001)**, IEEE CS Press, Kyoto, Japan, 2002.
- [2] D. Dolk, "Introduction to Modeling Technology and Intelligent Systems Track", **32nd Hawaii International Conference on System Sciences**, March, 1999.
- [3] ESRI, **GIS & mapping software**. Retrieved March 11, 2004 from <http://www.esri.com/>
- [4] **Geography Matters**. Geographic Information Systems. Retrieved March 11, 2004, from <http://www.gis.com>
- [5] J. Farley, J. Xie. "Managing Geographic Raster Data Using GeoRaster", **An Oracle Technical White Paper**, December, 2003.
- [6] D. Geringer. "Oracle Spatial: Best Practices", **An Oracle Technical White Paper**, December, 2003.
- [7] A. Girard. "MapInfo and Other Leading GIS Vendors Form Initiative to Make Enterprise Interoperability a Reality", **Business Wire**, September 17, 2003.
- [8] R. Kothuri. "Oracle Database 10g. Empowering Applications with Spatial Analysis and Mining", **An Oracle Technical White Paper**, February, 2004.

transaction processing, falls short when large volumes of data need to be analyzed. To improve accessibility, data warehouses have specific storage and management requirements, such as partitioning, transportable tablespaces, star schema queries and joins, and summary management. Spatial data layers can simply be added to the mix to increase the level of service to the data analysts. As a logical consequence of the need to mine these huge volumes of data, On-Line Analytical Processing tools have been developed to facilitate decision support and obtain business performance metrics. A spatially enabled data warehouse will move cooperative data sharing to a new level by allowing scientists and analysts to manipulate the data for results, rather than as a painful preliminary step to each new study.

- [9] J. Kopetsky, S. Ting. "If you build it, will they come?", **The Data Warehousing Institute**, 2002.
- [10] X. Lopez. "Oracle Spatial and Oracle Locator", **An Oracle Technical White Paper**, December, 2003.
- [11] S. Miller. "Intergraph Introduces Long Term Transaction Workflows for GeoMedia and Oracle8i", **Business Wire**, October 2, 2000.
- [12] R. Roberts. "Open Interoperability with Oracle Spatial Technology", **Intergraph Mapping and Geospatial Solutions**, September, 2003.
- [13] D. Scherer, W. Gaynor, A. Valentinsen, X. Cursetjee. **Oracle 81 Tips and Techniques**, Berkeley: Oracle Press/McGraw-Hill, Pub. 2000, pp. 143-188.
- [15] L. Silverston, W. H. Inmon, K. & Graziano. **The Data Model Resource Book**. New York: Wiley Computer Publishing, 1997, pp. 269-270.
- [16] J. Srivastava, P. Chen. "Warehouse Creation- A Potential Roadblock to Data Warehousing", **IEEE Transactions on Knowledge and Data Engineering**, 1999, pp. 118-126.
- [17] **What is GIS?** Retrieved March 11, 2004 from <http://www.gis.com/whatisgis>, 2002.
- [18] Zhang, J., Guan, J., Zhang, Ji., Chen, J. (2004). "Geographic Information Integration and Publishing Based on GML and SVG", **The Fourth International Conference on Computer and Information Technology (CIT'04)**, 2004, pp. 764-769.