

# Data Warehouse Performance

**Beixin (Betsy) Lin**

*Montclair State University, USA*

**Yu Hong**

*Colgate-Palmolive Company, USA*

**Zu-Hsu Lee**

*Montclair State University, USA*

## INTRODUCTION

A data warehouse is a large electronic repository of information that is generated and updated in a structured manner by an enterprise over time to aid business intelligence and to support decision making. Data stored in a data warehouse is non-volatile and time variant and is organized by subjects in a manner to support decision making (Inmon et al., 2001). Data warehousing has been increasingly adopted by enterprises as the backbone technology for business intelligence reporting and query performance has become the key to the successful implementation of data warehouses. According to a survey of 358 businesses on reporting and end-user query tools, conducted by Appfluent Technology, data warehouse performance significantly affects the Return on Investment (ROI) on Business Intelligence (BI) systems and directly impacts the bottom line of the systems (Appfluent Technology, 2002). Even though in some circumstances it is very difficult to measure the benefits of BI projects in terms of ROI or dollar figures, management teams are still eager to have a “single version of the truth,” better information for strategic and tactical decision making, and more efficient business processes by using BI solutions (Eckerson, 2003).

Dramatic increases in data volumes over time and the mixed quality of data can adversely affect the performance of a data warehouse. Some data may become outdated over time and can be mixed with data that are still valid for decision making. In addition, data are often collected to meet potential requirements, but may never be used. Data warehouses also contain external data (e.g. demographic, psychographic, etc.) to support a variety of predictive data mining activities. All these factors contribute to the massive growth of data volume. As a result, even a simple query may become burden-

some to process and cause overflowing system indices (Inmon et al., 1998). Thus, exploring the techniques of performance tuning becomes an important subject in data warehouse management.

## BACKGROUND

There are inherent differences between a traditional database system and a data warehouse system, though to a certain extent, all databases are similarly designed to serve a basic administrative purpose, e.g., to deliver a quick response to transactional data processes such as entry, update, query and retrieval. For many conventional databases, this objective has been achieved by online transactional processing (OLTP) systems (e.g. Oracle Corp, 2004; Winter and Auerbach, 2004). In contrast, data warehouses deal with a huge volume of data that are more historical in nature. Moreover, data warehouse designs are strongly organized for decision making by subject matter rather than by defined access or system privileges. As a result, a dimension model is usually adopted in a data warehouse to meet these needs, whereas an Entity-Relationship model is commonly used in an OLTP system. Due to these differences, an OLTP query usually requires much shorter processing time than a data warehouse query (Raden, 2003). Performance enhancement techniques are, therefore, especially critical in the arena of data warehousing.

Despite the differences, these two types of database systems share some common characteristics. Some techniques used in a data warehouse to achieve a better performance are similar to those used in OLTP, while some are only developed in relation to data warehousing. For example, as in an OLTP system, an index is also used in a data warehouse system, though a data

warehouse might have different kinds of indexing mechanisms based on its granularity. Partitioning is a technique which can be used in data warehouse systems as well (Silberstein et al., 2003).

On the other hand, some techniques are developed specifically to improve the performance of data warehouses. For example, aggregates can be built to provide a quick response time for summary information (e.g. Eacrett, 2003; Silberstein, 2003). Query parallelism can be implemented to speed up the query when data are queried from several tables (Silberstein et al., 2003). Caching and query statistics are unique for data warehouses since the statistics will help to build a smart cache for better performance. Also, pre-calculated reports are useful to certain groups of users who are only interested in seeing static reports (Eacrett, 2003). Periodic data compression and archiving helps to cleanse the data warehouse environment. Keeping only the necessary data online will allow faster access (e.g. Kimball, 1996).

### MAIN THRUST

As discussed earlier, performance issues play a crucial role in a data warehouse environment. This chapter describes ways to design, build, and manage data warehouses for optimum performance. The techniques of tuning and refining the data warehouse discussed below have been developed in recent years to reduce operating and maintenance costs and to substantially improve the performance of new and existing data warehouses.

### Performance Optimization at the Data Model Design Stage

Adopting a good data model design is a proactive way to enhance future performance. In the data warehouse design phase, the following factors should be taken into consideration.

- **Granularity:** Granularity is the main issue that needs to be investigated carefully before the data warehouse is built. For example, does the report need the data at the level of store keeping units (SKUs), or just at the brand level? These are size questions that should be asked of business users before designing the model. Since a data ware-

house is a decision support system, rather than a transactional system, the level of detail required is usually not as deep as the latter. For instance, a data warehouse does not need data at the document level such as sales orders, purchase orders, which are usually needed in a transactional system. In such a case, data should be summarized before they are loaded into the system. Defining the data that are needed – no more and no less – will determine the performance in the future. In some cases, the Operational Data Stores (ODS) will be a good place to store the most detailed granular level data and those data can be provided on the jump query basis.

- **Cardinality:** Cardinality means the number of possible entries of the table. By collecting business requirements, the cardinality of the table can be decided. Given a table's cardinality, an appropriate indexing method can then be chosen.
- **Dimensional Models:** Most data warehouse designs use dimensional models, such as Star-Schema, Snow-Flake, and Star-Flake. A star-schema is a dimensional model with fully denormalized hierarchies, whereas a snowflake schema is a dimensional model with fully normalized hierarchies. A star-flake schema represents a combination of a star schema and a snow-flake schema (e.g. Moody and Kortink, 2003). Data warehouse architects should consider the pros and cons of each dimensional model before making a choice.

### Aggregates

Aggregates are the subsets of the fact table data (Eacrett, 2003). The data from the fact table are summarized into aggregates and stored physically in a different table than the fact table. Aggregates can significantly increase the performance of the OLAP query since the query will read fewer data from the aggregates than from the fact table. Database read time is the major factor in query execution time. Being able to reduce the database read time will help the query performance a great deal since fewer data are being read. However, the disadvantage of using aggregates is its loading performance. The data loaded into the fact table have to be rolled up to the aggregates, which means any newly updated records will have to be updated in the aggregates as well to make the data in the aggregates

4 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: [www.igi-global.com/chapter/data-warehouse-performance/10879](http://www.igi-global.com/chapter/data-warehouse-performance/10879)

## Related Content

---

### Outlier Detection

Sharanjit Kaur (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 1476-1482). [www.irma-international.org/chapter/outlier-detection/11015](http://www.irma-international.org/chapter/outlier-detection/11015)

### Examining the Validity and Reliability of the Arabic Vocabulary Achievement Instrument to Evaluate a Digital Storytelling-Based Application

Nurul Azni Mhd Alkasirah, Mariam Mohamad, Mageswaran Sanmugam, Girija Ramdas and Khairulnisak Mohamad Zaini (2024). *Embracing Cutting-Edge Technology in Modern Educational Settings* (pp. 264-284). [www.irma-international.org/chapter/examining-the-validity-and-reliability-of-the-arabic-vocabulary-achievement-instrument-to-evaluate-a-digital-storytelling-based-application/336199](http://www.irma-international.org/chapter/examining-the-validity-and-reliability-of-the-arabic-vocabulary-achievement-instrument-to-evaluate-a-digital-storytelling-based-application/336199)

### Non-Linear Dimensionality Reduction Techniques

Dilip Kumar Pratihari (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 1416-1424). [www.irma-international.org/chapter/non-linear-dimensionality-reduction-techniques/11007](http://www.irma-international.org/chapter/non-linear-dimensionality-reduction-techniques/11007)

### Modeling Score Distributions

Anca Doloc-Mihu (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 1330-1336). [www.irma-international.org/chapter/modeling-score-distributions/10994](http://www.irma-international.org/chapter/modeling-score-distributions/10994)

### Data Preparation for Data Mining

Magdi Kamel (2009). *Encyclopedia of Data Warehousing and Mining, Second Edition* (pp. 538-543). [www.irma-international.org/chapter/data-preparation-data-mining/10872](http://www.irma-international.org/chapter/data-preparation-data-mining/10872)