

## VERITAS Indepth for Oracle

### Using VERITAS Indepth for Oracle in a Data Warehouse Environment

## TABLE OF CONTENTS

TABLE OF CONTENTS .....	2
INTRODUCTION .....	3
Understanding the Data Warehouse Environment .....	3
Delivering Short- and Long-Term Performance Management.....	3
Effecting Long-Term Monitoring .....	4
Analyzing Trends .....	4
Analyzing Patterns.....	4
Analyzing Object Access .....	5
Analyzing Object Size.....	6

## INTRODUCTION

Over the past few years, interest and investment in data warehouse applications have skyrocketed. Information technology (IT) vendors have flooded the data warehouse arena with products and services that help deploy data warehouse applications in corporate MIS environments. These technologies, which help data warehouse administrators manage the daily operations of their data warehouses, run the gamut from enterprise resource modeling solutions for star schema designs to extraction and transformation utilities, data mining tools, meta data management products, and “pumping” and replication technologies that refresh the contents of data marts.

Successful data warehouses respond quickly and accurately to user queries. Therefore, in addition to demonstrating query responsiveness, which is the most important feature of data warehousing, data warehouses must demonstrate flexibility and ease of use.

## UNDERSTANDING THE DATA WAREHOUSE ENVIRONMENT

A data warehouse is a dynamic database environment, one in which dynamic SQL queries generate and execute. In the dynamic data warehouse environment, users have little control over SQL statements, and SQL statements do not typically repeat themselves in the short term. Through design and redesign, the databases that support data warehouses must reflect these unique characteristics. In fact, because data warehouses are so dynamic, they are expected to be redesigned after they are launched.

After a data warehouse has been designed, built, and launched, staff must administer it continuously. Data warehouse administration entails several tasks, including backup and recovery, data loading, and performance management. VERITAS InDepth for Oracle delivers comprehensive, proactive performance management to the data warehouse.

## DELIVERING SHORT- AND LONG-TERM PERFORMANCE MANAGEMENT

A data warehouse is an Oracle application, so it is monitored like an Oracle application. Data warehouse environments are dynamic, however, unlike standard OLTP environments. In OLTP environments, transactions repeat constantly and SQL statements execute again and again. In data warehouse environments, users are few, and the statements those users execute change often. In fact, rarely does a statement execute only once in a data warehouse environment.

Because it changes constantly, the data warehousing environment requires a special type of monitoring. Standard ad hoc monitoring, which shows the current state of the database and the SQL statement currently in cache, is inadequate for data warehousing. In the data warehouse environment, ad hoc monitoring directs administrators to problem areas based on when the monitoring was performed. Because data warehouse users constantly perform different tasks and execute different SQL statements, such monitoring may not identify the correct problem areas. For example, data warehouse monitoring during Period A may identify Statement X as the most problematic, while monitoring during Period B may identify Statement Y as the most problematic.

To be effective, monitoring in the data warehouse environment must be done throughout the data warehouse’s active period. Through such long-term monitoring, the data warehouse administrator can identify current resource users, as well as something more important: major resource consumers in

preceding months. Short-term information might direct the DBA to tune the wrong area, while long-term information helps the DBA locate performance problems and identify trends, like users who are consuming increasing amounts of resources or tables.

### **Effecting Long-Term Monitoring**

To monitor long term, a data warehouse monitor must be able to work continuously—24 hours a day, 7 days a week. The monitor must also be able to collect information as the level of database activity changes, as the number of users changes, and as the amount of consumed resources changes.

To sample, VERITAS Indepth for Oracle accesses Oracle’s system global area (SGA) directly and therefore imposes no overhead on Oracle and minimal overhead on the system. As a result, VERITAS Indepth for Oracle can sample at a very high frequency and provide the highly granular data that is crucial to performance management. With its low overhead, VERITAS Indepth for Oracle is ideal for production environments. Monitoring tools that impose high overhead cannot be used in production or at least cannot run continuously. Because it runs continuously, VERITAS Indepth for Oracle can capture all information and provide a clear, comprehensive picture of performance.

VERITAS Indepth for Oracle continuously writes its monitoring results to the Performance Warehouse, a set of Oracle tables that serves as a data repository. The Performance Warehouse provides information for any duration, regardless of when or how often Oracle was brought up. The long-term performance information in the Performance Warehouse can be viewed and analyzed using VERITAS Indepth for Oracle or any tool that accesses Oracle tables. The Performance Warehouse facilitates the analysis of:

- Trends
- Patterns
- Objects

### **Analyzing Trends**

Trending is key to understanding the constantly changing data warehouse environment. The Performance Warehouse stores information for long periods, which allows administrators to identify application resource consumption trends on many different levels, such as the user level, the program level, and the statement level. On the user level, the data warehouse administrator can view such information as the amount of resources each user consumes on different days of the week, during a certain month, and from week to week. With this information, an administrator can better plan resource growth and, when using the Oracle 8.1 resource manager, generate optimal resource distribution plans.

### **Analyzing Patterns**

Data warehouses are often accessed by statements that contain constants rather than bind variables. From Oracle’s perspective, statements like “select \* from T where salary > 100” and “as “select \* from T where salary > 200” differ. In standard monitoring, these two statements will be marked as different and resource consumption will be monitored for each statement separately.

From the data warehouse administrator's perspective, this approach is wrong. The data warehouse administrator must treat these statements in the same way and monitor the statements' total resource consumption rather than the resource consumption of each statement alone. This approach is most accurate for long-term monitoring. In that case, it is also important not to generate too much information but to help the data warehouse administrator identify patterns, not focus on specific statements.

VERITAS Indepth for Oracle identifies statements that look the same and vary only in the constants they contain. VERITAS Indepth for Oracle users can choose to treat each statement differently or to combine statements and monitor them as a combined entity. The latter approach helps the user better tune the database objects not by single statement occurrence but by statement pattern.

### **Analyzing Object Access**

Object access analysis, which focuses on database objects like tables, indexes, and columns, can be used as the sole analysis method or with SQL statement analysis. Object analysis should be used as the only method when there is no access or limited access to SQL statement text. Most tools that are used in a data warehouse environment are "point and click," allowing users to see logical schemas, point on the tables of interest, specify further conditions, and execute requests. Such tools fail to show SQL statements or to allow users to change those statements. As a result, data warehouse tuning must focus on object tuning, such as by creating optimal indexes for a table, specifying optimal columns in a table, or creating summary tables.

To perform these tasks, data warehouse administrators need information on object access and activity as well as on statement activity. The Performance Warehouse provides:

- Focus on a table and a view of all statements that access that table or all statements that access that table in a specific access path, such as in a full table scan. This option allows the administrator to redesign tables in such a way that ensures users access tables in the best ways and do not access redundant tables.
- Focus on an index and a view of all statements that access that index or all statements that access that index in a specific access path, such as in a range scan. This option allows the administrator to optimize index structure and ensure that, for load purposes, the number of indexes is minimized while the access path is optimized.
- Information on which indexes are not in use. This allows the administrator to drop unused indexes and to minimize load time.
- Focus on a column and a view of all statements that access that column or all statements that access that column based on where that column is mentioned in the statement (e.g., the "where" clause, the "order by" clause, the "group by" clause). This option also allows the administrator to focus on columns that are accessed with a specific predicate (e.g., an equal predicate or a range predicate). This allows the administrator to place columns in tables optimally and, when using Oracle 8.1, to design summary tables optimally.
- Focus on an Oracle file and view all statements that access that file. This option allows the administrator to optimize I/O per file.

## Analyzing Object Size

A data warehouse grows constantly. Soon after it is implemented, a data warehouse doubles and triples in size. In a data warehouse environment, it is crucial to monitor object size and data distribution as well as the performances of statements that access objects, because statement performance is directly related to the size of the object the statement accesses.

VERITAS Indepth for Oracle constantly monitors the sizes of all database objects as well as activities like creating or dropping and modifying objects. VERITAS Indepth for Oracle houses the resulting information in the Performance Warehouse to allow the data warehouse administrator to see such long-term information as growth rate per table, clustering factor change per index, and number of values and value distribution change per column. With this information, the administrator can do such things as better plan required disk space, identify the fastest growing data warehouse components, and compare actual growth rate with planned growth rate. When correlating object size information with the performance information of the statements that access those objects, the administrator can see such things as how statements perform while data grows and which index is chosen for a certain table size. When viewing object changes and statement performance, the administrator can see how performance of all statements that access a table is affected when an index is added or dropped.

All this information allows the administrator to better plan the forecasted performance and the changes that must be made to support performance goals.

**VERITAS Software Corporation**  
Corporate Headquarters  
350 Ellis Street  
Mountain View, CA 94043  
650-527-8000 or 866-837-4827

For additional information about  
VERITAS Software, its products, or the  
location of an office near you, please call  
our corporate headquarters or visit our  
Web site at [www.veritas.com](http://www.veritas.com).