



VOLUME I ♦ ISSUE 1
SPRING/SUMMER 1998

HOURGLASS SYSTEMS news

Welcome

Welcome to Hourglass Systems' first newsletter and digest! We hope that you'll find this and future editions both informative and useful. Consistent with the focus of our consulting practice, we'll bring you articles on the design and implementation of data warehouses and decision support systems, as well as information on such related issues as usability and integration.

In this first issue we feature an article by Chris Silbernagel on data warehouse security. Chris is intimately familiar with the subject, having dealt expertly with it for several

clients. Security schemes are generally an anathema to warehouse performance and user needs; however, there are a number of situations where you have no choice but to implement security. Chris will point out the right reason for doing so, the techniques to use for implementation and the tradeoffs you can expect to make.

The industry is increasingly using Web technologies to deploy decision support systems. There are a number of excellent products on the market today, including MicroStrategy's DSSWeb. They provide a means to deploy powerful systems over a wide area with a minimal amount of client software and maintenance, as well as a common mechanism for customers and others outside the enterprise to access information. They also present some unique problems as the architecture and programming model for the web is different than that for conventional mainframe or client-server computing. To help meet the integration challenge, we will devote a number of future articles to education and solutions in this exciting field.

It is our practice to devote considerable effort to usability. A common theme among our clients is the need to extend the functionality of off-the-shelf DSS products with custom solutions designed to meet specific needs. Many companies are using data warehouses to feed electronic publishing applications that automate the marriage of text, numeric data and graphics

continued on page 6

Inside

Who We Are P2
Data Security Protecting the Data Warehouse from Within P2



Who We Are

Hourglass Systems, Inc. is a consulting firm specializing in the design and implementation of complex Decision Support Solutions. Simply stated we help clients turn data into actionable information.

We provide knowledge and experience. Our experience is in the analysis, design, development and implementation of large-scale applications, synthesizing complex business and technical issues into well-architected solutions. Our knowledge is in traditional, current and emerging technologies.

Hourglass is focused exclusively on developing data warehouse and decision support solutions for customers. All our energies and research are targeted on the areas of data warehouse and datamart design, data modeling, decision support project management, front-end evaluation, design and customization, multi-tier systems integration and performance consulting.

We can structure engagements ranging from multi-day "warehouse audits" to multi-year "warehouse advisors."

Our customers are leaders in a variety of enterprises, including consumer goods, manufacturing, financial services, telecommunications, health care and pharmaceuticals. Headquartered in New York, Hourglass offers services throughout the United States.



Data Security:

Though the purpose of a Data Warehouse is to exploit information, most companies seek to limit the information that a data warehouse client can access.

There are many very good and very bad reasons to restrict, or secure, the information in a Data Warehouse. All too often this decision is made well after the warehouse is designed because the requirements for security seldom surface during pilots, prototypes or even implementations of DataMarts.

The following article outlines some of the primary issues and decision points of implementing a Relational Online Analytical Processing Data Warehouse with data security. It is intended to raise awareness about the potential to underestimate the effort or price of implementing a secured data warehouse and help you think critically about the level, too high or too low, of the security you are planning.

Understanding, and Accepting, the Purpose of your Warehouse

You must clearly understand the purpose and use of your data warehouse before any security plan can be designed. Be honest. If the majority of your users are going to be viewing standard reports, even though the main "selling" point of the warehouse was data analysis, the security plan should be biased towards standardized reporting. This way, the impact of security will last far longer than the initial warehouse proposal.

In this article, we will address the three fundamental types of data warehouses.

Protecting the data warehouse from within

— Chris Silbernagel, Hourglass Systems, Inc.

Analytical

These are the data warehouses that you read about. They enjoy significant mindshare during the planning and selling of a data warehouse.

However, industry practitioners will agree that 80% of the data warehouses are built for standardized reporting. In most companies there may be a dozen individuals who understand the information, data, and have sufficient breadth of control to be effective at analyzing and therefore implementing plans based on an analytical data warehouse.

While these users are of primary concern, they are a small percentage of the overall data warehouse market.

Security or anything that restricts access to information has no place in an analytical warehouse.

Standardized Reporting

Standardized reporting has been a goal of most companies long before the term data warehouse was coined. In these cases the data warehouse has simply become the enabling technology.

Security is almost always mandated. These warehouses will have widespread usage throughout an organization and are the most sensitive to performance degradation.

Data Homogenization/Consolidation

These data warehouses combine multiple sources of information with a primary goal of integrating the information. This is one of the goals of every data warehouse. However, organizations with mature decision support systems that span many databases and systems, may choose to integrate that data into one source as the reporting and application demands outstrip their ability to “code” the functionality.

The complexity of security increases as different sources of information are combined. Financial information will have a different security plan than sales and inventory. Different areas of the organization may have their own security plan.

If the warehouse is to be successful there can only be one security plan for each type of information.

Application vs. Database Level Security

There are two schools of thought regarding the appropriate place for security.

One logical place for security to be implemented is in the application. Security is integrated into the application and can be specific to the data accessed by the application as well as the functions of the application.

The other place for security is in the data warehouse. Security is thereby provided consistently to all applications and has a single point of maintenance.

What are the Decision Points?

- Is there more than one application accessing the warehouse?
- Is the warehouse accessible via standard methods such as ODBC?
- Are there more than 100 users?
- Are there more than four dimensions?
- Are there multiple hierarchies within the same dimension?

If you answered yes to any of the above questions then implementing security at the data warehouse level is appropriate. Therefore, the following sections relate specifically to designing and implementing security in a relational database with a dimensional database design.

If the warehouse is to be successful there can only be one security plan for each type of information.

Data Security *cont.*

The Security Table

The security table contains each attribute which is to be secured along with a user identification that relates the person logged into the warehouse to the attributes.

For each user the table contains the values of each of the attributes that the user is permitted to access. If the warehouse is secured at the lowest levels of granularity this table can become quite large. It can be the largest table in the data warehouse.

This table or tables become the basis for all security views, roles or partitions that provide the physical implementation of the security plan.

Security vs. Value

The purpose of a data warehouse is to unlock and exploit a company's information. This is my information and my peers throughout the company should not be able to access it. Theory and reality.

The value of the information is protected by implementing security and lessened by restricting access. By carefully matching the purpose of the warehouse with the possibilities of the information the appropriate level of security can be determined.

Usefulness

One of the side benefits of securing the dimension in a data warehouse is the increase in usability. Your level of security only presents you with information that is permitted. List boxes and tree controls are manageable in size. Performance navigating the front-end application during query preparation may be improved over an unrestricted warehouse.

Four Types of Data Security

There are actually three definable types of data warehouse security and a fourth type that is a conglomeration of all three. If you can't describe any of the first three types without getting a "yeah, but...", you fall into the conglomeration category.

1 Individual

Example: A client can view only information about themselves.

If your warehouse is accessible by individuals outside your company this is the most common type of security.

2 Group

Example: Information for a Division can be viewed by anyone in the Division. One Division cannot view information from another Division.

For an internal data warehouse this is the least restrictive security plan. It is also the rarest.

3 Hierarchical

Example: A person at any level can view only their assignment and the assignments of their direct reports down to the sales representative level.

This is most often used in sales and marketing warehouses or in single subject area datamarts. It provides an acceptable impact to performance and meshes with the corporate cultures of most companies.

4 Conglomeration

Example: A person at any level can view details of their assignment, the assignments of their direct reports and summarized corporate level data.

Securing
a data
warehouse
can become a
larger effort
than building
the data
warehouse.

security

Performance

Performance will suffer. In our experience we have observed a 20 to 500% increase in query run times. The most important issue to consider is the threshold at which the usage of the warehouse drops.

A restrictive security plan will incur the greatest penalty with the worst performance for the least restricted users. A loose security plan will provide the best performance with the most restricted users having the worst performance.

For a secured enterprise data warehouse this is the security plan. The level of complexity varies widely and there may be more than one plan per warehouse.

These warehouses are the ones that require analysis on the feasibility and maintainability of the security plan.

Design and Construction

Schema Implications

The design and specification of your security plan must be completed before you finalize the design of your data warehouse. Since the most common methods of implementing security involve database views, roles and partitions, the underlying database must be designed in concert with these methods.

Dimensional modeling techniques in a secured warehouse will lean heavily towards a snowflake design and will have limited aggregation tables. Fact columns or rows may need to be redundant in order to provide information such as a company summary.

The list of items to consider is extensive.

Feasibility and Maintenance

Maintaining security in a database is difficult. Organization changes can require rebuilding the entire security plan. Security views can degrade the performance beyond acceptable limits. Applications and interfaces may be required to administer the rights and levels of security.

Securing a data warehouse can become a larger effort than building the data warehouse. Ensure that this reality is interjected into the design of your security plan.

Reporting Appropriateness

One of the more subtle nuances of implementing a secured warehouse is providing the recipient with a report of the information that is not there.

For example, a sales representative has five stores from a single chain. The demographic market that the sales representative is in has seven stores. When the sales representative prints a market report by chain only the five

stores will be included in the total. The report, at first glance, will lead the observer to the wrong conclusion.

We therefore, need to include information on the report about the missing components (i.e. five of seven stores). Each report needs to be correct and appropriate when viewed in isolation.

© Hourglass Systems Inc. 1997-1998

summary

This article touched on some of the main issues and concerns involving data warehouse security. Most companies will elect to secure their data warehouses. In fact most corporate cultures are biased towards protecting rather than exploiting information. Fortunately this is in alignment with the primary use of a data warehouse and current technologies enable the implementation of a feasible and maintainable security plan.

about the author

Chris Silbernagel is one of the most experienced developers of data warehouses in the industry. He is a frequent lecturer on these topics, including data modeling, database design, mapping, transformation and migration middleware, Relational/OLAP tools, front-end tool selection and web-based access. Mr. Silbernagel is Vice President, Practice Leader—Data Warehouse Implementation for Hourglass Systems, Inc., a consulting firm specializing in the design and implementation of complex Decision Support Solutions.

Lead Article *continued*

into a single document for delivery to the end-user or customer. The DSS vendors are not in the electronic publishing business, but most do provide some tools to export data to outboard applications. In future issues we'll be explaining what's there, what isn't, and how to fill in the gaps.

A related area is reporting. Many businesses have found

the data warehouse an economical alternative to the mainframe and the Web an excellent way to distribute reports. The predominance of Windows NT Server as the middle tier in many DSS implementations enables the use of desktop technologies in what has traditionally been the domain of COBOL. You can produce information for users in formats that were impractical to implement on main-

frames. You can also write some pretty sloppy systems and get into trouble quickly. We'd like to share some of our experiences in this area in future editions.

In conclusion, we're excited about the state of the industry today and its future direction. We look forward to bringing you timely information and useful suggestions.

contact us

Corporate Office
Hourglass Systems, Inc.
4 Fritz Boulevard
Albany, NY 12205

P 518.452.9062
800.466.6609

F 518.452.2824

E info@hourglass.com

Metro New York Office

P 914.987.9528

F 914.986.6594

Hourglass Systems has been added to Sun Microsystems, Inc. New York State Contact #PT00086 as an authorized reseller.

Hourglass Systems is committed to providing the best customer support, products and services. That's why we're proud of our established partnerships with the following companies:



MicroStrategy



UNIX® is a registered trademark in the United States and other countries exclusively licensed through X/Open Company. Solaris®, Sun™ systems, and SPARCserver™ systems are registered trademarks of Sun Microsystems, Inc. Windows NT®, Windows®, and Microsoft® are registered trademarks of the Microsoft Corporation. Red Brick® is a registered trademark of Red Brick Systems. Sybase® is a registered trademark of Sybase Corporation. ORACLE® is a trademark of the Oracle Corporation.

Hourglass Systems, Inc.
4 Fritz Boulevard
Albany, NY 12205

*Specialist in Data
Warehouse and
decision support
solutions.*

*Contact us at
www.hourglass.com.*