

Managing the Warehouse Throughout Its Lifecycle

March 4, 1997

PLATINUM
TECHNOLOGY

The Open Enterprise Management Company

1815 S. Meyers Road
Oakbrook Terrace, IL 60181

PHONES	800-442-6861 (U.S. and Canada) 630-620-5000 (Worldwide)
FAX	630-691-0710
E-MAIL	info@platinum.com
URL	http://www.platinum.com

Table of Contents

INTRODUCTION	3
MANAGING INFORMATION ACROSS THE GLOBAL ENTERPRISE THROUGH META DATA.....	3
SYSTEMS MANAGEMENT.....	4
24 x 7 ‘Lights Out’ Operations.....	4
Chargeback.....	5
Capacity Planning	5
Enhancing Secure End-user Access.....	5
Backup and Recovery	6
PERFORMANCE MANAGEMENT.....	6
Optimizing Performance	6
Managing Data Availability	6
MANAGING CHANGE THROUGHOUT THE APPLICATION LIFECYCLE	7
Implementing and Managing Business Rules for Competitive Advantage	8
Managing Change throughout the Enterprise: Warehouse Application Vitality.....	8
Managing Database Changes	9
CHOOSING AN INTEGRATED, EASY TO MAINTAIN SOLUTION	10

Introduction

You've built a successful data warehouse. Congratulations! Now the question is how to best manage data access, maintain performance and handle expanding business needs. Some of the management areas you will need to consider include:

- How do I manage information accuracy and consistency across the enterprise?
- Can I provide 24 x 7 "Lights Out" operation of my warehouse?
- Can I easily provide financial chargeback to business units for the data marts?
- How do I plan for future capacity needs?
- How do I manage security for hundreds or thousands of end-users?
- What should my backup and recovery plans be?
- How do I maximize data availability?
- How do I optimize performance?
- How do I manage the lifecycle of the warehouse applications?
- How can I quickly reflect the changing business environment in my warehouse applications while providing stability?
- How do I manage application and database changes from the operational systems and warehouse systems?
- How can I implement these solutions without turning to multiple niche vendors?

This white paper will provide information on what questions successful data warehouse managers need to ask. It will also provide ideas on how to begin implementing management solutions within your organization.

Managing Information across the Global Enterprise through Meta Data

During the last year there has been a lot of promotion of the idea of independent, departmental data marts. However, some organizations that have built these stovepipe decision support systems are beginning to see that they are incapable of making accurate decisions across the enterprise because they have no way to consistently define data. These organizations will be slowed as they must now step back and evaluate data architecture as part of a data warehousing strategy.

PLATINUM technology believes that implementing small decision support systems designed for a specific user community, built within a short period of time, is a good thing. However, we also provide a way for our customers to build these systems as part of an enterprise solution. Here is an overview of what organizations should consider.

When defining data models for your system, select those data definitions that you know go outside your business unit. These are definitions that would make sense to roll-up into a corporate data warehouse or to integrate with another department's system. For example,

“product revenue” is something that most companies would want a common definition of so that company-wide analysis is possible. Other key terms may include “customer,” “profit,” “ship-date,” etc. Work with a corporate group to ensure that you have defined these terms appropriately for the business. Once you’ve ensured that the data definitions are consistent, there are two different ways to manage those definitions.

First, when implementing a departmental data mart, many organizations may not want to consider a full repository implementation as part of the core project. However, organizations can gain many of the benefits of a repository implementation, particularly those related to data consistency, by implementing a portion of the repository solution. This would mean:

- Scanning only applications that will be used as sources for the data mart and storing that meta data
- Using an end user access tool such as PLATINUM Data Shopper to understand meta data in business terms
- Storing target definitions in the repository for repeatability—access by other applications and departments—and reusability, enabling fast new application development
- Storing key data definitions for multi-department data consistency in the repository

This approach provides the benefits of shareable information about data, data consistency throughout the enterprise, and an excellent head start toward a full repository implementation.

As part of the initial data mart implementation, or as an add-on, organizations may choose to implement a repository. A full repository implementation can provide the data consistency benefits above in addition to the systems and applications management benefits described earlier for the data mart and for other critical applications. Repositories must support the languages and CASE tools used by your organization and should store meta data in relational database such as DB2, Oracle, Sybase, or Informix, rather than a proprietary data store.

Systems Management

24 x 7 ‘Lights Out’ Operations

A decision support system is only as good as the data that feeds it. The success of a decision support system depends in part on the reliable movement and update of data from operational systems to analytical systems. There are two separate worries: is the data being correctly transformed, and is it being reliably moved? Data movement tools should have scheduling capability themselves or should be run as jobs through a job management tool.

Let’s look first at an example of how a failed data movement job (from source DB2 to warehouse Oracle) could adversely affect an organization, and then see how that failure can be averted. Some organizations transform data for the warehouse, using either a software tool or internal applications, then FTP the data to Oracle. However, this type of movement is not easy to manage. Data is selected from a mainframe source, transformed, and moved to an Oracle database, as updates to an existing data warehouse database. If the server is unavailable, the database load fails, which makes the warehouse data out-of-date. Other data updates from different source databases may have been completed, however. This means that the warehouse data is not only incomplete, but now may be inaccurate. The problem may not even be discovered until the DBA staff arrives the next morning and nightly reports are distributed to executive management.

In 24 x 7 operations, systems management tools can not only avert disaster but help automate the solution. In this scenario, Zeke, a mainframe scheduler, schedules the data transformation

movement job. When the job fails, the global event manager notifies the network management system and the PLATINUM Enterprise Console. If a problem resolution software product such as PLATINUM Apriori were in place, a problem notification ticket can be issued so that appropriate personnel are notified and corrective action can be taken. Once the server is back online, the job is resubmitted and the data loaded. Nightly reports are run, and managers have the data for sound business decisions.

Chargeback

Once all data warehouse and data mart projects belonged solely to IT, but that time is fading. In most organizations, these projects are managed by several different departments, each of which has its own financial goals. Data warehouse managers should ensure that they can charge back appropriate costs to business units and users so that they can meet financial reporting requirements.

An integrated solution monitors IT costs by providing critical chargeback services that track information resources used organization-wide. Such a tool should provide chargeback, job accounting, activity-based costing, resource utilization reporting and system performance reporting. In organizations whose systems include not only MVS but also UNIX and other client/server systems, managers should also consider a tool that directly interfaces with mainframe usage data as well as usage data from networks, client/server, UNIX, AS/400, PBX systems and many others. An easy-to-use reporting system for point-and-click analysis of chargeback is an added feature that you may want to consider.

Capacity Planning

There is a component of warehouse planning that does not get as much attention in the industry as it gets within organizations. That is capacity planning: how will the warehouse affect my system's response time, DASD requirements, etc. A capacity planning software tool can make this effort easier and more accurate. A capacity planning tool should process enterprise-wide data and maintain a historical resource utilization and performance database. Planners should be able to review a wide variety of executive, management, and technical reports and graphs that address:

- system response performance
- task availability
- resource utilization
- operations analysis
- trends reporting
- service level reporting
- batch window analysis.

Enhancing Secure End-user Access

Because many different groups of end users will be accessing different sets of information in the analytical data store, sophisticated security for both the application and the database must be in place. Applications will access multiple servers and multiple objects, thus requiring users to have privileges on multiple servers. A security product that can define privileges through user groups should be considered. By defining privileges for whole groups of users (for example, Northeastern store managers), security administrators define privileges once, then add and delete users as they change jobs. Changes in security can be made once, to the user group definition, rather than separately to multiple users. Database security needs are similar. Look for a security tool that can help manage the DBMS security, enabling you to define user groups, analyze the

effects of changes to security, and analyze the relationship between user security and applications.

In addition to implementing a proactive mechanism that controls and prevents unauthorized access to sensitive system resources (such as files and programs) and to networks, warehouse managers should consider a tool that notifies administrators of attempted violations.

Backup and Recovery

Areas such as performance management and tuning, database administration, and security may be addressed in similar terms for OLTP systems and data warehouses, but backup and recovery needs some special consideration. Like all systems, the data warehouse should have a backup and recovery strategy that will enable the organization to recover data in case of emergency. However, it may require a different approach. For example, depending on the size of the data mart, you may choose not to do a backup, because you can refresh the data mart more efficiently—just as if it were the weekly refresh. Organizations should review the cost/benefit of each warehouse and mart, keeping in mind how often the data is updated or refreshed and how long recovery will take to implement.

Performance Management

Optimizing Performance

End-users, who include executive management, demand fast response time from a decision support system. To succeed with data warehousing, organizations must plan to address these commonly identified potential performance problems:

- Poor response time
- Slow systems due to poorly written SQL
- Resource contention
- Freespace problems

Performance may be degraded for many reasons, including poorly tuned SQL, poorly designed database objects, and poorly designed systems. Look for tools that can optimize SQL as applications are being developed as well as find poorly performing SQL and suggest ways to improve it.

An integrated database performance solution such as PLATINUM Enterprise Performance Management suite can help warehouse managers monitor, maintain, and manage warehouse performance in a distributed, heterogeneous environment. The PLATINUM Enterprise Performance Management Suite:

- Monitors and manages performance of Oracle, Sybase, Informix and DB2 databases in distributed UNIX environments.
- Monitors and manages performance of UNIX servers in distributed environments.
- Monitors and manages performance on the entire Enterprise Network
- Measures and analyzes the performance of distributed client/server transactions.

Managing Data Availability

Fragmentation and freespace problems can also impact performance and availability. For example, if the available free space falls below a value that does not allow a new segment or extent to be created, processing stops for that process. Therefore, not only does the current available free space for a tablespace need to be monitored, but also history needs to be kept for

purposes of trend analysis, to determine the effectiveness of the current allocation and storage needs for the future.

If the temporary tablespace runs out of free space during a sort operation, the operation will fail. You should put temporary segments in their own separate tablespace. This allows for easier monitoring to ensure the availability of sufficient space to support the resource needs of large, long-running jobs (for example, those that create indexes, or large sorts). Rollback segments are involved in every transaction in a database. Since rollback segments are created in a tablespace they are subject to running out of free space and reaching the maximum allowable number of extents. When either situation is encountered, the transaction forcing the rollback segment to extend will fail.

Organizations need to implement procedures for detecting and correcting fragmentation through database reorganization to ensure maximum data availability.

Managing Change throughout the Application Lifecycle

For decades, IT organizations have been plagued with problems in delivering applications on time and within budget. Application development projects typically have been difficult to deliver with consistent success. This difficulty is a fact of life for application development managers, CIOs, and end users; and statistics indicate that the problem is worsening, not improving. The Standish Group reports that only 16% of significant application development projects are completed on time and within budget. The majority of applications (53%) are completed but are delivered late and over budget—an average of 89% over budget. And a dismal 31% of all projects that are started are canceled before delivering a single line of code to end users.

This terrible track record costs businesses billions of dollars every year; more importantly, it severely limits future development projects which can lead to improvements in the competitiveness of businesses. In some organizations, these perennial problems have been accepted as the “nature of the beast”. Yet IT executives who will flourish in this new era recognize that they must improve their *business* of delivering applications so they can contribute significantly to their companies bottom line and competitive advantage.

PLATINUM *technology* helps organizations improve their application development efforts by automating and streamlining the maintenance and development process through its Application Lifecycle Solutions. This open-architected, integrated yet modular suite of tools addresses every phase of an applications lifecycle— from process and change management, to analysis and design, through construction and testing. It even supports help desk operations and software distribution. A separate white paper, “Integrating Application Lifecycle Functions to Improve the Delivery of Successful Software Systems,” describes the PLATINUM solution for smoothing and integrating phases of the application development process while remaining open to tools offered by other vendors.

The PLATINUM strategy for application development integration is four-fold:

- *Process and Methodology Integration.* Extend application development methodologies and processes to address the entire breadth of development tasks and integrate the products into the methods and processes.
- *Integrated Application Lifecycle Meta-model.* Provide repository-based sharing and management of all the components of an organization’s applications.

- *Individual Product Integration.* Build and extend bridges between individual products to share application components across PLATINUM products, as well as to integrate PLATINUM products with products from other vendors.
- *Product Operation and User Interface.* Improve the ease of learning and ease of use among products by creating a consistent user interface, a common look and feel, and common support procedures for all products. Improve tool usability by sharing components and code where possible.

Implementing and Managing Business Rules for Competitive Advantage

Traditionally, business rules have been relatively stable, and have often been coded directly into an application. However, today's global competitive environment means that business rules change frequently. Because such a large part of turning data into information involves correctly defined business rules, warehouse projects require a way to capture and manage business rules as they change so that decision support systems can be continually updated. Data warehouse applications will require an application development environment for logically complex business applications, which speed time to market and increase flexibility for competitive advantage. This translates into a demand for development environments that support the business rule focus of the data warehouse.

Let's look at the insurance industry as an example. The insurance industry is undergoing the greatest realignment of market and business forces in its history. As the lines between insurance, banking, and financial services blur, competition is becoming more fierce. Consider these factors:

- Corporate policies and government legislation are becoming more complex and are changing more quickly.
- Products change constantly and must be cross-sold in new ways.
- As claim and policy volumes rise, skilled claims reviewers and underwriters must be leveraged for only the unusual or costly cases.
- Sales agents must also act as financial service consultants.
- Managed care is redefining products and service offerings for all health insurance companies, pressuring support organizations to re-evaluate how they do business.
- Niche industries are emerging, such as high-risk auto insurance or bodily-injury specialties.
- Consumers increasingly turn to Internet services and offerings for information updates, while insurers seek new ways to reach and sustain a loyal customer base.

To keep up and gain a competitive edge, many companies organizations in every industry are applying information technology to their toughest run-the-business problemsXthose that include dynamic rules and policies.

A PLATINUM white paper, "How Insurance Companies are Automating Core Business Policies with AionDS," provides detailed information on how implementing business rules can make your organization more competitive. Please see the PLATINUM web site at <http://www.platinum.com>.

Managing Change throughout the Enterprise: Warehouse Application Vitality

Enterprise IT technology (applications, platforms, and data) must keep pace with demanding business requirements. Unfortunately, enterprise technology improvements often disrupt and impede business processes, which fosters frustration and resistance to change in IT practices. Frequently, IT's chief goal is stability, and stability chiefly means absence of change to the enterprise technology base. In contrast, PLATINUM takes the view that change is a key

ingredient in stability. The negative side-effects of change must be restricted, but change itself must be encouraged. The difference is subtle, but critical. Many IT professionals restrict change itself to control the side-effects it has on stability. Successful organizations, however, restrict only the side-effects of change and so avoid stagnation. Organizations that embrace data warehousing use their newly acquired information to make changes in their business practices that will bring them competitive advantage.

How does change become an integral part of stability? Through policies, procedures, and tools. By planning for rapid change and by using tools which can virtually eliminate negative side-effects of change, an enterprise can stay vital and fresh. Many disciplines are involved, but this paper will discuss Application change. By Application Change we mean managing the development, deployment, and support of IT software applications (usually in-house applications, but this also applies to third-party applications).

PLATINUM *technology* offers complete toolsets in the areas of development, deployment, and support. A full discussion of these individual disciplines is beyond the scope of this document. Rather, this discussion focusses on the "hand-off" between the disciplinesXhow we limit the negative effects of application change:

- between development and deployment
- between deployment and support
- between support and re-development.

We call this series of hand-offs "Application Vitality" because, implemented properly, this solution will enable business needs, not fear of change, to drive the technology that reaches the end-user's desktop.

The warehouse implementation of application change is particularly challenging because it involves not only the transaction-based operational systems, but also the decision support systems and the applications that tie the operational systems with the decision support systems. The complexity of the warehouse environment makes the challenge of application vitality even more critical to a business that relies on changing business practices for competitive advantage.

Please see the white paper, "PLATINUM Application Vitality," on our Web site for detailed information on this concept.

Managing Database Changes

One of the more complex areas to manage is the area of database change. Because an enterprise data warehouse system may incorporate multiple operational systems, multiple operational data stores, and multiple analytical data stores across multiple platforms, a single database change can affect dozens of other database objects on multiple servers. This is one area, in particular, where strong tools can significantly reduce risk.

Repository technology can help manage not only reference data, but also all data within the data warehouse (or even within the enterprise). If data definitions are stored in the repository, there are two ways to manage database change throughout the enterprise, using tools like PLATINUM Enterprise DBA, an enterprise-level control center for managing heterogeneous distributed databases and database objects.

The preferred method is to make the changes to the logical model (repository) first. Once a change to a database object is recorded in the repository, PLATINUM Repository issues a change

description and reference list that is passed to Enterprise DBA. Enterprise DBA then creates a strategy (and the DBA can review and edit it, if desired) to implement the changes to the physical databases, across the enterprise, for mainframe and client/server databases.

However, many organizations will independently make changes to physical databases throughout the enterprise without notifying the repository or data warehousing group. In this second scenario, the independent database administrators use Enterprise DBA to make changes to their DB2, Oracle, Microsoft, Sybase, or Informix database objects. PLATINUM Repository stores information about the database objects, including server information, and how different data elements relate to one another. Once a change is made to a physical object, the information in the Repository can be used by Enterprise DBA to migrate the changes to the appropriate databases.

Choosing An Integrated, Easy to Maintain Solution

When choosing a data warehousing vendor, it's important to choose a vendor who:

- understands the business and technology issues
- provides an experienced-based methodology
- provides data transformation and movement solutions
- provides data consistency through repository technology
- provides ongoing systems and database management solutions
- has a consulting force that can provide senior level project managers as well as education, training, and implementation consultants for end-to-end solution.

PLATINUM *technology* has the products, the people, and the methodology to help your organization succeed, including solutions for all the warehouse management challenges presented in this paper.

To obtain additional information about PLATINUM **data warehouse, change and systems management solutions**, please contact PLATINUM *technology, inc.*, at:



The Open Enterprise Management Company

1815 S. Meyers Road
Oakbrook Terrace, IL 60181
PHONES 800-442-6861 (U.S. and Canada)
630-620-5000 (worldwide)
FAX 630-691-0710
E-MAIL info@platinum.com
URL <http://www.platinum.com>

© 1997 by PLATINUM *technology, inc.* All rights reserved.

The PLATINUM products referenced in this document are trademarks of PLATINUM *technology, inc.* and its subsidiaries.