The current issue and full text archive of this journal is available at
www.emeraldinsight.com/0263-5577.htm

Understanding management data systems for enterprise performance management

Ranjit Bose
Anderson School of Management, University of New Mexico, Albuquerque, New Mexico, USA

Abstract

Purpose – Managing enterprise performance is an important, yet a difficult process due to its complexity. The process involves monitoring the strategic focus of an enterprise, whose performance is measured from the analysis of data generated from a wide range of interrelated business activities performed at different levels within the enterprise. This study aims to investigate management data systems technologies in terms of how they are used and the issues that are related to their effective management within the broader context of enterprise performance management (EPM).

Design/methodology/approach – A range of recently published research literature on data warehousing, online analytic processing and EPM is reviewed to explore their current state, issues and challenges learned from their practice.

Findings – The findings of the study are reported in two parts. The first part discusses the current business practices of these technologies, and the second part identifies and discusses the issues and challenges the business managers dealing with these technologies face for gaining competitive advantage for their businesses.

Originality/value – The study findings are intended to assist the business managers to effectively understand the issues and technologies behind EPM implementation.

Keywords Data handling, Storage, Online databases, Business performance

Paper type Literature review

Introduction

As businesses move further into the twenty-first century, their managers increasingly need accurate and timely performance indicators to manage and lead them. In this pursuit enterprises are increasingly turning to software systems to seek support for enterprise performance measures to aid goal setting, monitor progress, identify and draw attention to financial implications of organizational decisions, facilitate internal benchmarking, identify inefficiencies in core business operations, and identify cost saving and operation improvement opportunities (Leahy, 2003a, b).

Companies today operate in an ever increasingly competitive environment. They treat their customers like royalty as they try to lure them to buy their goods and services. Finding and retaining customers is a major critical success factor for most businesses, offline and online. Customer relationship management (CRM) is a customer service approach that focuses on building long-term and sustainable customer relationships (Rowley, 2004) that add value both for the customer and the company. In today's e-business era, in order to remain competitive, enterprises are forced to respond to their customers' expectations by using the best, most powerful, and innovative systems and software (Liautaud, 2001).

Industrial Management & Data Systems
Vol. 106 No. 1, 2006
pp. 43-59
© Emerald Group Publishing Limited
0263-5577
DOI 10.1108/02635570610640988

Within the past several years, not only CRM systems, but enterprise resource planning (ERP) systems, along with many other complex software applications such as supply chain management (SCM) were deployed to record and process every detail of all of the business transactions of a corporation. These software systems capture a rich set of data that contains valuable information to the business enterprise (McAdam and Galloway, 2005). They include customer preferences and buying patterns, mission critical business process information, and a wealth of additional business data. The problem with all of these transactional data is that they are quite often not organized, integrated, and synchronized, thus making it difficult, at times even impossible, for an enterprise to determine whether it is succeeding or failing (Jr *et al.*, 2000). Hence, this mass of data often does not give the business owners, executives, or managers the appropriate indicators to chart and measure their business activities in a timely manner.

The promising idea then is to tie these operational data from enterprise-wide systems to specific strategic goals, and to provide managers with integrated visibility into performance against those goals. To ensure that each enterprise is meeting its strategic goals, the managers need to turn to business intelligence based enterprise performance management (EPM) analytics to keep tabs on enterprise actions for daily comparison to strategic goals and budget targets (Schultz, 2004). Consequently, the managers are provided with a good barometer of where their business is at any point in time, instead of learning about them at quarter's end. Capturing these current updates on company health is becoming more critical for enterprises seeking enhanced visibility into operations (Singh *et al.*, 2000). This kind of "managing in advance" or proactive stance of the EPM allows management to practice what they want to accomplish as a business rather than simply reflecting on what just happened.

EPM provides visibility into how well a company is maintaining its strategic focus. Developing a corporate strategy is a necessary step for the company in defining who it is and where it fits in the market. Once a strategy is defined, the company must need to measure how well it is executing that strategy over time. Key performance indicators (KPIs) allow a company to see in what areas it is executing well, and what areas require improvement (Reh, 2005). The identification of appropriate KPIs as well as aligning them with company strategies then becomes the key to realizing bottom line impact (Toni *et al.*, 1997).

Data warehousing and online analytic processing (OLAP) are the two fundamental technologies used by software vendors as well as enterprise IT application developers in a multitude of businesses or industries such as retail sales, telecommunications, financial services, and real estate for developing EPM systems (Mundy, 2002). A successful data warehousing system provides decision makers with consistent, timely, reliable, and accessible data, without negative impact on the operational systems from which the data is extracted. The system integrates data from various sources to describe trends in the operations of the organization and the environment in which it operates (Inmon *et al.*, 1997). OLAP refers to the technique of performing complex analysis over information stored in a data warehouse (DW) to provide status reports and decision support (Chaudhuri and Dayal, 1997). OLAP is a category of applications and technologies for collecting, managing, processing and presenting multidimensional data for analysis and management purposes (Thomsen, 1997).

The characteristics of OLAP applications are quite different from those of online transaction processing (OLTP) systems used in organizations. OLTP systems are operational systems for collecting and managing the base data in an organization, such as sales order processing, inventory, accounts payable, etc. They usually offer little or no analytical capabilities as required in EPM. The executive information systems (EIS) developed in 1980s and refined in 1990s (Basu *et al.*, 2000), are also a category of applications and technologies for presenting and analyzing corporate and external data for management purposes (Kumar and Palvia, 2001). Their characteristics include extreme ease of use and fast performance, but their analytical functionality is usually very limited thus they are not suitable for EPM.

Given the importance of EPM in today's business environment and the underlying management data systems technologies that are used to build and use them, the managers need to effectively understand these technologies in terms of how they are used and the issues that are related to their effective management within the broader context of EPM.

This study reviews the existing research literature on data warehousing, OLAP and EPM to explore their current state for addressing the above need. The paper reports the findings of the study in two parts. The first part discusses the current business practices of these technologies, and the second part identifies and discusses the issues and challenges the business managers dealing with these technologies face for gaining competitive advantage for their businesses. Interestingly, most large organizations in the United States have already implemented DWs, others are currently in the process of doing so, or are in the planning stages. Many organizations that have attempted a DW implementation have not been adequately prepared and therefore have not achieved the level of success they were expecting (Data Warehousing Institute, 2004-2005).

Data warehousing

A DW is defined as a structured extensible environment designed for the analysis of non-volatile data, logically and physically transformed from multiple source applications to align with business structure, updated and maintained for a long time period, expressed in simple business terms, and summarized for quick analysis (Jarke *et al.*, 2000).

The first key concept in the definition above is extensibility, which means that a good DW design must have built into it the ability for expansion because the demands to either include more data from the same application(s) or data from other applications arise rapidly. The second concept is that the data stored in a DW comes from one or more operational applications. Data warehousing therefore involves taking these volatile operational data and rendering them non-volatile, which is required for meaningful analyses. Although operational data stores can be the basis for limited, real-time data analysis needs, operational data stores are, however, not designed for extensive analyses. The third concept is that a sound DW design is often built around a "time dimension" and therefore the DW contains data over several periods of time. This feature allows users to perform extensive yearly, quarterly, and monthly analyses that help enable the identification of patterns and trends, which would have been much more difficult (if not impossible) to glean from operational data alone (Theodoratos and Sellis, 1999).

Figure 1 provides a general overview of an enterprise DW creation project. Data are obtained from different sources, manipulated into a common format for the warehouse, inserted into the warehouse with any necessary calculations or additional appended data, then loaded into appropriate reference tables or data marts for efficient query performance, analyses, reporting, or data mining (Forcht and Cochran, 1999) by the users through the use of commercially available tools such as business objects, web intelligence, oracle crystal reports.

An enterprise DW provides several benefits to an organization. The most important benefit is the creation of a "single source of truth", that is, a single source of organizational data. Thus it enables valid and consistent reporting and decision-support across the organization. Although an enterprise DW may start with a limited subset of enterprise data, it is designed to expand over time. Most enterprise DWs are managed and controlled by the central IT organization. An organization-wide effort to improve data quality is another important benefit that can be gained from the enterprise DW initiative (Watson *et al.*, 2001).

A data mart has a limited scope: it supports a particular region, business unit, or business function. For example, a data mart may contain sales information for a specific region or product line. In comparison, an enterprise DW contains sales data for all regions and most products, or is at least designed with this in mind. A data mart is most of the time built by central IT, but quite often managed independently by a department or workgroup. The difference between an enterprise DW and a data mart is therefore essentially a matter of scope.

Data warehousing architecture and processes

The data warehousing architecture commonly used is called a multi-tier warehouse, in which an enterprise warehouse coexists with several data marts. Several variations on this multi-tier approach have been implemented in organizations till date, namely top-down, bottom-up and hybrid. In the top-down implementation approach, data flows from the source to enterprise warehouse to data marts (Figure 1). This style of warehouse, is usually controlled by the central IT group, and improves the consistency of information in the data marts. The obvious advantage of this implementation approach is that it leads to a planned, integrated multi-tier solution. However, it usually takes more time and is relatively costly. Thus this approach often becomes unacceptable in a competitive business climate.

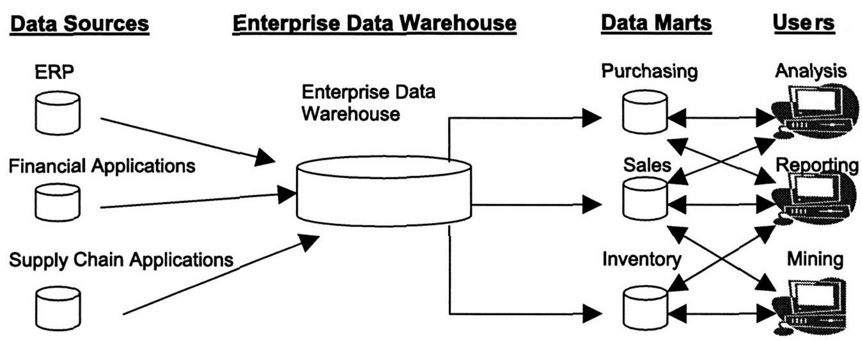

Figure 1.
Architecture of a distributed data warehouse

In the bottom-up approach, the enterprise warehouse is evolved bottom-up as a new layer on top of existing data marts. In this case, the data marts are loaded directly from source systems, and the enterprise warehouse is loaded from the data marts. Although not ideal, this approach is expedient when data marts are built before the enterprise warehouse. Although this approach gives quick results and a high return on investment, it eventually yields a disintegrated warehouse because the data marts often do not conform to a common model. The hybrid approach on the other hand may include elements of both the top-down and bottom-up approaches. In this approach the enterprise model is developed first. It is documented at a high level, so certain subject areas may be modeled in more detail as warehouse development proceeds. When the outline of the enterprise model is in place, the enterprise warehouse and the data marts can be built in parallel.

In the long run, a multi-tier warehouse is the best architecture. It provides a single source of clean, integrated data, as well as local stores tailored to the needs of specific groups. Because it is more difficult to build and manage, many organizations begin in the short term with isolated enterprise warehouses or data marts (Ma *et al.*, 2000).

The ETL and data integration processes play a vital role in the overall success of any data warehousing-based business intelligence project. Data integration problems are widespread, and they afflict just about every organization. The three major dimensions on which ETL and data integration problems are measured are huge variety, constant change, and large volumes. E in ETL stands for *extract*, which is characterized by the variety or the degree of heterogeneity of the source data present in organizations; T stands for *transform*, which is characterized by the change or the degree of data cleansing, scrubbing and standardizing that is needed; and L stands for *load*, which is characterized by the data volume or the time required to populate the warehouse.

The ETL process is carried out by a collection of relatively complex software programs. Selecting a set of reliable ETL programs is pivotal. Several studies indicate that the importance of ETL is grossly under estimated when engaging in data warehousing and business intelligence projects (Watson *et al.*, 1999). Many studies suggest that the ETL process devours 60-70 percent of the work and cost of DW development. Therefore, should an organization "make or buy its ETL"? The research literature leans towards buying whenever possible due to the complex nature of the process and that a third party vendor that specializes in this field is better able to incorporate the latest technologies in their product. In practice, however, the answer depends on several factors including the size of the organization and the composition of its IT staff (Agosta, 2000).

Generally, larger organizations with complex structures and experienced IT personnel are more likely to lean towards "making" their ETL software product while, on the other hand, the ability to "buy" a good ETL product may be the factor that offers many smaller organizations the possibility to employ DW technology in the first place. The large organizations that chose to "make" their ETL cite increased flexibility and control as the primary deciding factors; however, they concede that they might consider purchasing an ETL product in the future as the technology matures.

A large number of commercial tools support the ETL process for DW in a comprehensive way, e.g. COPYMANAGER (InformationBuilders), DATASTAGE (Informix/Ardent), EXTRACT (ETI), POWERMART (Informatica), DECISIONBASE (CA/Platinum), DATATRANSFORMATIONSERVIVE (Microsoft), and WAREHOUSEADMINISTRATOR (SAS) among others.

DWs typically use the multidimensional and relational storage structures. The multidimensional structure physically stores the data in array-like structures that are similar to a data cube. In the relational structure the data is stored in a relational database using a special schema (star or snowflake) instead of a traditional relational design.

Online analytic processing

DW and OLAP technologies are the core of modern decision support systems. They are complementary technologies because the DW makes summary data available to OLAP and ensures its timeliness, accuracy and consistency (Inmon, 1997) whereas OLAP focuses on the end-user's analytical requirements. Through an OLAP interface the decision makers access the DW to analyze corporate data on various dimensions; view corporate changes over a period of time, to obtain a macro view of business operations as well as perform a microanalysis in a specific sub-function; perform various what-if analyses; and drill-down and discover the pattern of sales of certain products in a given period of time or find how the sales performance of an individual salesperson affects the company's revenues.

OLAP supports decision making based on multi-dimensionally organized summary (aggregate) data. In multidimensional data analysis a decision maker needs summary data related to a specific subject and he/she must consider that data with respect to certain factors. Summary data are usually numerical and measurable. Therefore, the attributes representing them are often called measure attributes. The factors on the basis of which summary data is analyzed are called dimensions, represented by dimension attributes. By selecting the specific dimensions through which summary data are analyzed one can obtain a view into summary data. By changing dimensions one may construct different views. This happens through OLAP queries, which specify new multidimensional views from the basic views provided by data warehousing.

Decision makers often need to group data, e.g. they might want to consider dimensions at different levels of detail. Therefore, it is important to represent the dimensions as multilevel hierarchies. For example, the dimensions time and geography could be represented as multilevel hierarchies (days, weeks, months, quarters, years) and (cities, states, countries), respectively.

Multidimensional databases are currently developed without any widely accepted formal model. Therefore, there is no consensus on the primitives of and no established terminology in multidimensional modeling. However, a common feature of multidimensional databases is that information is represented as multidimensional arrays. Summary data are often modeled as a multidimensional data cube consisting of measure and dimension attributes. Thus multidimensional data cubes can be considered as the basic logical or conceptual model for OLAP while the operation set for data cube manipulation may vary considerably between models. At the instance level, the values of the dimension attributes are assumed unique to determine the values of all measure attributes.

A popular OLAP data model is the star schema although it is based on intuition rather than on precise formalism. In it a multidimensional data cube consisting of dimension and measure attributes is called a fact table. In addition, it contains a dimension table for each dimension attribute in the star schema. A dimension table describes the properties of the dimension at hand. The star schema is mainly a model for the logical structuring of multidimensional data.

All multidimensional models containing fact and dimension tables are variants of the star schema, the snowflake schema is probably its most famous variant. It is a star schema where the dimension tables are normalized (as in the relational model).

OLAP architectures

OLAP architectures are helpful when selecting an OLAP software product. Examples of OLAP vendors and products include Hyperion's Essbase, Oracle's Express, and Sybase's IQ. The OLAP product architectures can be classified in a matrix form, as shown by Figure 2, developed by Pendse (2004). The rows of the matrix indicate where the OLAP data are processed, and the columns indicate where the OLAP data is stored.

The rows are labeled:

- multi-pass SQL – since SQL does not have the ability to perform multidimensional calculations in a single statement, complex multi-pass is necessary;
- multidimensional server engine – popular place to perform multidimensional calculations, found in many products; and
- client multidimensional engine – since users tend to have powerful workstations.

	Multidimensional data storage options		
Multidimensional processing options	**RDBMS**	**Multidimensional database server**	**Client files**
Multi-pass SQL	*1* Cartesis Magnitude MicroStrategy		
Multidimensional server engine	*2* *Crystal Holos (ROLAP mode)* *Hyperion Essbase* Longview Khalix Speedware Media/MR Microsoft Analysis Services Oracle Express (ROLAP mode) Oracle OLAP Option (ROLAP mode) *Pilot Analysis Server* WhiteLight	*4* SAS CFO Vision *Crystal Holos* Comshare Decision *Hyperion Essbase* *Oracle Express* Oracle OLAP Option AW Gentia *Microsoft Analysis Services* PowerPlay Enterprise Server *Pilot Analysis Server* Applix TM1	
Client multidimensional engine	*3* Oracle Discoverer	*5* Comshare FDC Dimensional Insight Hyperion Enterprise Hyperion Pillar	*6* Hyperion Intelligence BusinessObjects Cognos PowerPlay Personal Express TM1 Perspectives

Figure 2.
OLAP architectures

The columns are labeled:

- RDMS – data is organized using star or snowflake schema;
- multidimensional database server – data is stored in a multidimensional database on a server; and
- client-based files – data is held on client machines on disk or in RAM.

In multidimensional OLAP (or MOLAP) databases, cubes are created and stored physically. The products in cells 4 and 5 of the matrix represent MOLAP architecture. In relational OLAP (or ROLAP), cubes are a virtual concept based on a star or snowflake schema. The products in cells 1-3 of the matrix represent ROLAP architecture.

The cell 6 of the matrix represents desktop OLAP products such as business objects and is based on multidimensional file structures. Some products such as Microsoft Analysis Services are classified as hybrid OLAP or HOLAP. HOLAP provides multidimensional analysis simultaneously of data stored in a multidimensional database and in an RDMS, and has become a popular architecture for server OLAP. The products in cells 2 and 4 of the matrix represent HOLAP architecture.

Enterprise performance management

EPM, also known as corporate performance management (CPM) and business performance management (BPM), includes the combination of planning, budgeting, financial consolidation, reporting, strategy planning, and business scorecard (Walker, 1996) tools. Most vendors do not offer the full set of these components, so they adjust their version of the definition to suit their own product set (Menninger, 2003).

The three general strategic focuses companies may employ in their EPM are described as cost-, differentiation- or growth-based. A cost-focused strategy emphasizes supplying a standard product or service that meets many customers' needs without customization at the lowest cost possible. A differentiation-focused strategy includes custom or niche products or specialized services delivered to its customers. Growth-focused companies place their emphasis on maintaining competitive economic position in the growth of the economy and industry.

Once the strategy is identified, a company must measure performance in terms of how well it is executing that strategy over time. KPIs allow the company to do that – to see in what areas it is executing well and what areas require improvement at the enterprise level or specific to departments. Therefore, within the identified strategy, KPIs help the company define and measure progress toward the company goals (Reh, 2005).

Figure 3 provides examples of KPIs for each of the three strategic focuses identified above. Note that the list of KPIs provided here is by no means exhaustive rather they are for illustration purposes. Some of these KPIs may be simple to define and measure, such as financial performance and other objective measurements. Other may be more subjective such as customer empathy or employee morale. Defining the appropriate KPIs for a corporate strategy can be as important as defining the strategy itself. KPIs are quantifiable measurements, agreed to beforehand, that reflect the critical success factors of an organization. They will differ depending on the organization. Whatever KPIs are selected, they must reflect the organization's goals, they must be key to its success, and they must be quantifiable (measurable). KPIs usually are long-term

Strategic Focus	Example Key Performance Indicators (KPIs)
Cost Leadership	• Cost measurements such as production and delivery cost, cost per unit • Cycle time such as production time, time to service customer • Conformance to product or service standards • Production or Service volume • Capacity utilization • Profitability • Quality measures such as TQM, Six Sigma
Product or Service Differentiation	• Time to market a new product or service • Product or Service customization • R&D rollouts, new patents • On-time delivery • Knowledge of customers or personalization • Customer complaint management
Growth	• Knowledge sharing such as best practices • Customer acquisition and retention • Market share • Account penetration

Figure 3.
Enterprise performance management measures

considerations. The definition of what they are and how they are measured do not change often.

An executive dashboard – also known as a manager's dashboard, an executive cockpit, digital cockpit or a business scorecard – is a software application that provides a single-screen display of relevant and critical business metrics and analytics to enable faster and more effective decision making (Menninger, 2003). In other words, a dashboard is a summary of the critical measurements required to make the daily business decisions that affect an organization's bottom line. The foundation of the executive dashboard is a set of KPIs. Revenue forecasts, gross profit, inventory levels, the list of current top customers; all qualify as KPIs as long as they are important to the business and can be measured.

In summary, an EPM system for business intelligence (BI) is an absolute must for organizations that want to keep their fingers on the pulse of their business activities. Such a system could provide an organization with the following capabilities:

- a single screen, browser-based portrait of the organization with drill-down capability on each KPI monitored;
- real-time presentation of information in chart and graph format, based on data pulled from the corporate DW, data marts, or legacy systems;
- slice-and-dice capability on KPIs that lets users perform what-if and sensitivity analysis; and
- integrated management of KPIs and issues raised by their performance levels – all based on appropriate and individual user-based security clearances.

Lessons learned from practice

Data warehouse issues and challenges

Justification and strategy. While a well-designed DW provides several advantages to an organization, getting executive management approval for its development is difficult because building and maintaining a robust DW solution is an expensive proposition and given the historically high failure rates, executive management is often apprehensive. Management would want to know what exactly the DW would do – make or save money. The CIO must therefore be able to vividly convincingly demonstrate that the data warehousing technology is a best business practice that is vital to the long-term success of the organization. Furthermore, a data warehousing project must be a collaborative effort between the IT and user communities with users providing a detailed list of requirements and visible support from executive management (Gardner, 1998).

After approval, the development team must determine whether to implement a centralized, enterprise-wide DW or a decentralized, divisional data mart solution. If possible, an organization should attempt to implement a single DW. Doing so establishes a single, reliable source for data and provides a more integrated solution for reporting and decision support across functional areas. However, the data mart solution may be well suited for highly specialized data needs. Regardless of the approach, the development team must not take on too much at once which can leave users feeling abandoned and the development team overwhelmed. Instead, an incremental approach will likely yield the best results.

Implementation. DWs face technical and organizational challenges that affect the success or failure of their implementation (Curtis and Joshi, 1998). Two approaches have been reportedly used to investigate DW implementation process. The first approach studied the factors that cause DW implementations to fail, and the second approach studied the successful DW implementations and the subsequent distilling of factors that positively impact implementation success. Based on the results obtained from these two approaches, several critical implementation factors (CIFs) were identified which were found to have significant impact on the success or failure of DW implementation. These CIFs were placed into six major categories:

(1) technical – data, technology, and expertise;
(2) management sponsorship – executive sponsorship and operating sponsorship;
(3) goals and objectives of an organization – having a business need, and having a clear link to business objectives;
(4) user-related issues – user involvement, user support, and user expectation;
(5) organizational factors – organizational resistance, and organizational politics; and
(6) system-growth-related factors –evolution and growth.

Without a proper corporate data warehousing strategy and architecture, a multitude of data mart implementations with disparate technologies and uncoordinated data models spring up creating long-term problems (Watson *et al.*, 1999). For example, due to the affordability of the technology, an increasing number of vendors have joined the marketplace with "out of a box" data mart solutions. These data mart solutions can be implemented within a matter of weeks, as opposed to the traditional methodical approaches that take months or years to implement.

The strategy and architecture should help guide organizations address issues such as best-of-the-breed versus single vendor tools, and functionality versus integration (Jarke *et al.*, 2000). Best-of-the-breed tools place functionality ahead of integration, whereas, single vendor strategy places integration among the tools ahead of functionality. Deciding on what the "best" features are can also be debatable, as the best tool from an implementation and support perspective may not be the best tool from the user interface perspective. Likewise, the best tool at low volumes may not be the best tool as volumes grow rapidly. There is also the danger that people will tend to favor tools with which they are familiar.

Data quality. Data warehousing efforts may not succeed for a variety of reasons, but nothing is more certain to yield failure than lack of concern for the quality of the data. Nevertheless, why is concern for data quality sometimes not paramount? The reason is not simple carelessness; instead, the huge amount of data that are managed by a typical DW can quickly become unruly and almost impossible to verify. The implementation teams must therefore ensure that they are able to solve the data quality problems in a DW by data cleaning and data transformation operations. The data cleaning should detect and remove all major errors and inconsistencies both in individual data sources and when integrating multiple sources. The cleaning process should be supported by tools to limit manual inspection and programming effort and should be extensible to easily cover additional sources. The data cleaning should be performed together with schema-related data transformations based on comprehensive metadata. Data transformations are needed to support any changes in the structure, representation or content of data. They are necessary in situations when one has to deal with schema evolution, migration of a legacy system to a new information system, or integration of multiple data sources.

The data from a DW is typically used for decision support – such as measuring enterprise performance and taking appropriate corrective measures when necessary – rather than for operations. A particular data set within a warehouse therefore often supports several decision processes, which complicates data management because these uses are likely to require different degrees of data quality. The data quality is typically characterized via multiple dimensions, or attributes. The dimensions of data quality that are commonly used are accuracy, completeness, consistency, and timeliness. The attributes of data quality that are commonly used from the perspective of the end-users of the data include interpretability and availability or accessibility. Hence, the types of dimensions and/or attributes that are chosen to be present in the DW data to support the varied decision processes contribute to the overall cost for maintaining the DW. Managers often need to make trade-offs in the context of limited resources available for improving the data's quality.

Data synchronization. One of the reasons why companies incur large costs for having "bad data" in various parts of their operations is that their data is out-of-sync – meaning pieces of information related to the same product or service differ between supply chain partners or between systems within the same company. It is not too uncommon for these companies to have business processes that are hampered by the lack of consistent, good data between them and their customers. The solution to these problems is data synchronization, which means achieving consistent information values for items or products within and between organizations – that is, everyone is working off the same data page. It standardizes product or service or customer information from multiple data stores (applications throughout geographies, operating units and different departments such as marketing, manufacturing, and customer service, and in many disparate places such as departmental databases and commerce-enabled web sites) into a central, continuously updated repository for use by employees and trading partners, thereby significantly minimizing business problems caused by disparate product descriptions.

Data synchronization is supported by a variety of technologies today. Database applications such as IBM's DB2 Everyplace and Oracle's Oracle Lite, for example, have their own proprietary syncing technologies. Given the large and growing diversity of applications and devices (for example, handheld, desktop, network) organizations use today, a standard synchronization technology is needed. The SyncML initiative is such a standard, which promises to enable users to buy devices that synchronize with a wide range of data and devices, and to reduce the effort and costs expended by device manufacturers, service providers, and application developers as well. The SyncML, adopted by Open Mobile Alliance, is the leading open standard that drives data mobility by establishing a common language for communications among devices, applications, and networks.

Security. To a CIO the security challenge of a DW is paramount. Unlike the majority of the other corporate assets found in an organization, the electronic information stored in its DW is not a tangible asset and thus, it is much more difficult to safeguard. The characteristics of electronic information include:

- can be given away and still kept;
- can be stolen and not missed;
- an be owned and no one can tell;
- can be distributed instantly to almost everyone; and
- cannot tell if it is "real" or not.

Ironically the very thing that makes data warehousing so valuable is also the same reason that may cause an organization to be hesitant to go forward with such a project because all the information valued by the organization is stored in a single place that can be accessed by many people. One may think that an organization can simply tighten up access to only those individuals that have a "need-to-know" reason to access the data. However, this is often easier said than done since a dichotomy exists as DW personnel are constantly on the watch for new ways to market and use their resource with executive management while trying to preserve the sensitivity of the same resource.

It is important for business managers to keep in mind that an enterprise DW could be used for numerous unplanned purposes. For example, a court could subpoena an enterprise DW, which could ultimately result in the unintended release of large amount of proprietary information. Management must therefore stay abreast of exactly what information is being stored in their DW and who has access to it. Furthermore, members of the DW development team should provide management with a summary of intended and potential unintended uses of the data contained in the DW. Answers to these questions will enable management to develop a detailed security plan for their

DW implementation so that the organization can proactively manage and protect the corporate information stored therein in the same fashion as other corporate assets are protected.

OLAP issues and challenges

Application categories. Most OLAP applications include time as a dimension, and several useful results are obtained from conducting time series analysis (Business Intelligence Ltd, 1998). Efficient time series analysis is made possible due to DW's ability to hold several years' historical data. OLAP applications are broadly categorized into: marketing and sales analysis; clickstream analysis; database marketing; budgeting; financial reporting and consolidation; controls; management reporting; balanced scorecard; profitability analysis; and quality analysis.

Particularly, the clickstream analysis is very useful in the business-to-consumer e-business environment for understanding consumer behaviors and accordingly developing effective marketing plan (Liautaud, 2001). Commercial web sites generate gigabytes of data everyday that describe every action made by every visitor to the site, this could be captured in a multidimensional framework, the different dimensions to this analysis could include where the visitors came from, the time of the day, the route they take through the site, whether or not they started/completed a transaction, and any demographic data available about customer visitors.

Database marketing applications take advantage of multidimensional analysis to determine the preferred customers, develop loyalty packages for them, determine customer profile and use them to 'clone' the best customers. These applications generally make use of statistical and data mining technologies.

Query performance. Queries used in OLAP systems perform complex multidimensional aggregation over huge amount of DW data. They require fast response time for interactive execution. A commonly used approach to process the OLAP queries efficiently is to exploit materialized views (MVs), i.e. the results of pre-selected or previously issued queries are stored in DWs (Park *et al.*, 2003). However, there is a trade-off between response times and the storage requirements of pre-aggregated data. Determining how much to pre-aggregate is an issue that is addressed by the management team on a continuing basis. This approach requires methods for selecting appropriate views to materialize in a limited space in DWs and rewriting given queries using the MVs to process them efficiently.

Product evaluation. To decide on which OLAP product to adopt, one has to evaluate the available products. For improving ease of use of OLAP systems for the organizational users, management should evaluate the following features depending on their needs for selecting an OLAP product for deployment across their organization:

- Visualization – allows users to create summary tables interactively.
- Summarization – indicates degree of aggregation of information, measured in terms of number of hierarchies allowed and the level of detail among others.
- Navigation – refers to the capability of the tool to drill-down or drill-up between levels of detail.
- Query function – is the ability of the query engines to extract data from multidimensional databases and generate outputs in 3D graphics.

- Sophisticated analysis – is the ability to perform most common types of analyses used in decision support such as statistical profiling, moving averages, cross dimension comparison, queries with self-defined formula, exception condition, and what-if analysis.
- Dimensionality – refers to the number of allowable dimensions and the ability to redefine dimension.
- Performance – refers to the response times for the basic functions such as standard report generation, customized report generation, graphic/chart generation, and data navigation.

EPM issues and challenges

KPI development. The alignment of KPIs with organization vision/mission/strategies/ objectives is the key to realizing bottom line impact. The challenge is to develop KPIs that provide a holistic and balanced view of the business. Faced with potentially hundreds (if not thousands) of candidate metrics, how does one select those that are most meaningful? One potential approach is to think of individual KPIs not just as a singular metric, but as a balanced metric that incorporates several alternative dimensions. Different businesses would have their own specific dimensions and related measures. For example, the dimensions and their related measures could be:

- productivity – sales-to-assets ratio, revenue from new customers;
- quality – customer complaints, percent returns;
- profitability – profit by segment, profit by customer;
- timeliness – percent on-time delivery, percent of late orders;
- process efficiency – yield percentage, capacity utilization;
- cycle time – processing time, time to service customer;
- resource utilization – sales per total assets, sales per channel;
- cost savings – cost per unit, cost of goods;
- innovation – R&D spend, new patents; and
- technology management – IT capital spending, web-enabled access.

Note that there are possible overlaps between one or more dimensions above.

Different organizational roles have different job responsibilities and functions to perform, thus each would be responsible for a specific subset of KPIs within the company strategy. An enterprise portal is a browser-based, enterprise wide gateway to integrated information and applications to promote information sharing, consistency, and accessibility to members of an organization (Yang *et al.*, 2005). A well-constructed dashboard should be enterprise portal based that is configured appropriately for each manager using the dashboard. The portal screen a manager sees should contain the appropriate KPIs he/she uses to manage and execute his/her role in the business. These indicators could be displayed using gauges and graphics to give an indication of the health or status of a particular KPI. Each indicator can be programmed with a drill-down capability to allow the manager to progressively drill-down to additional levels of detail to analyze performance against the particular KPI. The dashboard should also be easily customized for the manager, so he/she can manage his/her appropriate view into the KPI and metrics of the organization.

Technological and organizational. Enterprise-wide performance analytics face several technological and organizational hurdles (Leahy, 2003a, b). Enterprises should be aware of the numerous technological challenges in making these frameworks live up to their promise. They range from data acquisition, changing grid and metadata management, to aligning models, and delivering performance management business intelligence capabilities scalably, securely, and flexibly across the portals. Similarly, many organizational challenges arise when developing EPM. Organizational processes should be adequately analyzed and appropriate KPI and metrics should be developed that truly reflect the health and trends of the organization. Consensus should be reached within the management team that the dashboard KPIs and metrics are appropriate and they do effectively gauge the performance of the organization.

Security is another concern with the dashboard. The KPI and metric information on the dashboard represents the inner workings of the organization. It presents, in a visual manner, the trade secret and key competitive processes and core competency data of the organization. The dashboard identifies how well the organization is doing financially, with key customers and suppliers, and in key processes and competencies. Security of this information needs to be managed effectively.

Product evaluation. The evaluation of business intelligence vendor products is critical for selecting a product for EPM development. For improving the ease of use of the EPM system for the organizational users, management should evaluate the following functionalities of each product for selecting the best to build the dashboard.

A customizable user interface – users must be able to choose which KPIs they want to see and have those KPIs displayed with easy-to-understand graphics such as “fuel” or “speedometer” gauges to show performance against set metrics or benchmark against other units’ performances.

Exception-based reporting – when exceptions occur – that is, information is out of line with KPIs – they should be presented via “traffic lights” in red, yellow, and green, to let user know the relative impact of the event.

Proactive alerting for exceptions and milestones – the dashboard should “come to the users”. It should send alerts in the form of pages, e-mails, onscreen alert messages, and so on to make the user aware of significant events within the defined KPIs.

The ability to create detailed numerical reports – the dashboard should enable unique, user-based definition of all criteria used in developing and monitoring the KPIs.

Thin client access (via the web) – the application must give the user the ability to view powerful analytics with a thin footprint.

Exceptional security – the organization should be able to set security parameters based on user, group, and community type; sensitivity of information; geographic region; and so on. In essence, the ability to define and set user permissions must be extremely flexible because the information presented is so sensitive.

Concluding remarks

Data warehousing, OLAP and EPM concepts have been around for several years; however, the power of and the need for effective implementation of EPM are just beginning to be realized. EPM implementation in organizations is very much in the growth stage where many new strategies are being developed, tested and employed on a fairly regular basis. To reap the benefits of data warehousing, OLAP and EPM promises, it is critical to perform an economic evaluation of a project before committing to it. Like any significant IT project, implementing and deploying EPM is subject to multiple risks. Build versus buy? Technology selection? Vendor selection? Return on investment? Do we need consulting help?

Behind these questions hide the details that may spell the difference between success and failure of a project, details that can only be identified in the context of a specific EPM implementation, with specific technology, towards a specific set of goals and requirements, working with a specific team of stakeholders. The research findings attempted to assist the business managers to better understand the current practices, issues, and challenges of data systems and their management within the context of measuring enterprise performance to effectively address the above-mentioned questions.

References

Agosta, L. (2000), *The Essential Guide to Data Warehousing*, Prentice-Hall, Upper Saddle River, NJ.

Basu, C., Poindexter, S., Drosen, J. and Addo, T. (2000), "Diffusion of executive information systems in organizations and the shift to web technologies", *Industrial Management & Data Systems*, Vol. 100 No. 6, pp. 271-6.

Business Intelligence Ltd (1998), *The OLAP Report: Market Share Analysis*, Business Intelligence Ltd, London.

Chaudhuri, S. and Dayal, U. (1997), "An overview of data warehousing and OLAP technology", *ACM SIGMOD Record*, Vol. 26 No. 1, pp. 65-74.

Curtis, M.B. and Joshi, K. (1998), "Lessons learned from the implementation of a data warehouse", *Journal of Data Warehousing*, Vol. 3 No. 2, pp. 12-18.

Data Warehousing Institute (2004-2005), *TDWI Business Intelligence Fundamentals*, Data Warehousing Institute, Seattle, WA.

Forcht, K.A. and Cochran, K. (1999), "Using data mining and data warehousing techniques", *Industrial Management & Data Systems*, Vol. 99 No. 5, pp. 18-19.

Gardner, S.R. (1998), "Building the data warehouse", *Communications of the ACM*, Vol. 41 No. 9, pp. 52-60.

Inmon, W., Welch, J.D. and Glassey, K.L. (1997), *Managing the Data Warehouse*, Wiley, New York, NY.

Jarke, M., Lenzerini, M., Vassiliou, Y. and Vassiliadis, P. (2000), *Fundamentals of Data Warehouses*, Springer, Cambridge, MA.

Jr, A.H.G., Pick, R.A. and Ward, S.G. (2000), "Perennial management data systems issues in US manufacturing firms", *Industrial Management & Data Systems*, Vol. 100 No. 7, pp. 325-9.

Kumar, A. and Palvia, P. (2001), "Key data management issues in a global executive information system", *Industrial Management & Data Systems*, Vol. 101 No. 4, pp. 153-64.

Leahy, T. (2003a), "The business performance management commitment", *Business Finance*, Vol. 9 No. 1, pp. 13-16.

Leahy, T. (2003b), "5 moves towards better BPM", *Business Finance*, Vol. 9 No. 3, pp. 23-7.

Liautaud, B. (2001), *E-Business Intelligence: Turning Information into Knowledge into Profit*, McGraw-Hill, New York, NY.

McAdam, R. and Galloway, A. (2005), "Enterprise resource planning and organizational innovation: a management perspective", *Industrial Management & Data Systems*, Vol. 105 No. 3, pp. 280-90.

Ma, C., Chou, D.C. and Yen, D.C. (2000), "Data warehousing, technology assessment and management", *Industrial Management & Data Systems*, Vol. 100 No. 3, pp. 125-35.

Menninger, D. (2003), "Information on demand", *Strategic Finance*, Vol. 85 No. 3, pp. 50-3.

Mundy, J. (2002), "Relating to OLAP", *Intelligent Enterprise*, Vol. 5 No. 16, pp. 20-2.

Park, C.S., Kim, M.H. and Lee, Y.J. (2003), "Usability-based caching of query results in OLAP systems", *Journal of Systems and Software*, Vol. 68 No. 2, pp. 103-19.

Pendse, N. (2004), "OLAP architecture", available at: www.olapreport.com/Architectures.htm#ideal

Reh, F.J. (2005), "Key performance indicators (KPI)", available at: http://management.about.com/cs/generalmanagement/a/keyperfindic_p.htm

Rowley, J. (2004), "Partnering paradigms? Knowledge management and relationship marketing", *Industrial Data and Management Systems*, Vol. 104 No. 2, pp. 149-57.

Schultz, G.J. (2004), "Enterprise performance management: strategic decision analytics from deep data", *Information Strategy: The Executive's Journal*, Vol. 20 No. 3, pp. 14-25.

Singh, H., Motwani, J. and Kumar, A. (2000), "A review and analysis of the state-of-the-art research on productivity measurement", *Industrial Management & Data Systems*, Vol. 100 No. 5, pp. 234-41.

Theodoratos, D. and Sellis, T. (1999), "Designing data warehouses", *Data and Knowledge Engineering*, Vol. 31 No. 3, pp. 279-301.

Thomsen, E. (1997), *OLAP Solutions: Building Multidimensional Information Systems*, Wiley, New York, NY.

Toni, A.D., Nassimbeni, G. and Tonchia, S. (1997), "An integrated production performance measurement system", *Industrial Management & Data Systems*, Vol. 97 No. 5, pp. 180-6.

Walker, K.B. (1996), "Corporate performance reporting revisited – the balanced scorecard and dynamic management reporting", *Industrial Management & Data Systems*, Vol. 96 No. 3, pp. 24-30.

Watson, H.J., Gerard, J.G., Gonzales, L.E., Haywood, M.E. and Fenton, D. (1999), "Data warehousing failures: case studies and findings", *Journal of Data Warehousing*, Vol. 4 No. 1, pp. 44-55.

Watson, H.J., Goodhue, D.L. and Wixom, B.H. (2001), "The benefits of data warehousing: why some organizations realize exceptional payoffs", *Information and Management*, Vol. 36 No. 6, pp. 1-12.

Yang, S., Yang, M. and Wu, J.B. (2005), "The impacts of establishing enterprise information portals on e-business performance", *Industrial Management & Data Systems*, Vol. 105 No. 3, pp. 349-68.

Corresponding author
Ranjit Bose can be contacted at bose@mgt.unm.edu