Reprinted from DM REVIEW

Growing Pains

Claudia would like to thank David Imhoff for his contribution to this column.

For many organizations, the first iteration of a data warehouse is like giving a child a puppy as a holiday gift. The anticipation is high, the rewards are many and costs are worth it "out of the box." What's more, given the proper care, feeding and training, the initial reward/benefit ratio is sustainable over time.

Most pet owners will tell you, though, that the care/feeding/training process requires constant monitoring and effort, especially during the growth phase. Most data warehouse owners will tell you the same.

As your implementation of the corporate information factory (CIF) grows, you will find a number of common "behaviors" hold true. First, like puppies, your first data marts will attract the attention of the others around you. If your data mart/puppy is effective/cute, others will complement you on it and want one, too. Before you know it, your implementation will include additional subject areas and data marts, and your data warehouse will grow into its enterprise data model "feet."

As the business community comes to depend on their analytical reports, the data warehouse will truly become the support mechanism for these growing marts. Its role becomes critical to the well-being of the entire business intelligence (BI) environment.

The growth is inevitable and welcomed. The downside is that as your data warehouse and data mart pet increases in size and usage, its reach increases to the kitchen counter (elsewhere in the organization). At this point, performance starts to decrease, and the cost of the entire environment becomes an eye-opener. If you don't swat it on the nose with a rolled-up newspaper, user satisfaction and overall support for the environment will begin to suffer.

Before it's too late for new tricks, you need behavior modification. First, you have to *monitor* your CIF's behavior.

Performance management technology has been around for several years. It's used to monitor the performance of servers, networks, databases and storage components. Today, there are vendors with tools available that are specifically designed to monitor BI environments. These monitoring tools are used to answer questions such as: Where are my current bottlenecks or hotspots, and why? Who is currently using the system, when and how often? What data are they (and aren't they) using? What are the response times for queries? With answers to these questions, you can *respond* to the current situation effectively.

The answers to these questions give you plenty of information about very tactical situations, but not enough about potential or hypothetical situations – somewhat like using your nose to find Fido's "surprise."

Most companies are reactive rather than proactive in terms of performance issues. Reactive performance management waits for a performance problem to appear and then deals with it in a knee-jerk fashion. These reactive modes are very costly. They cause schedule delays, lost productivity and revenue; and, perhaps most importantly, they damage the fragile relationship between IT and the business community.

It would be an improvement to *predict* that Fido needs a walk. The ability to answer questions about future performance and usage needs is critical.

- How well will the new application perform given our current setup?
- Will our current hardware configuration support additional workloads and users? If so, how many?
- What is the optimal hardware configuration to support our growing user community for the next five years?
- Are my most critical users getting the best resources in a timely manner?

The ability to provide these more holistic monitoring capabilities amounts to a competitive edge in today's economy. Proactive performance management anticipates problems and uses techniques for identifying and responding to those problems early in the process.

The problems described for reactive management can be prevented or avoided entirely by a systematic use of performance analysis and prediction techniques. The proactive approach to these problems is called strategic performance management (SPM). SPM takes into account the goals of the enterprise, looks at how each line of business supports these goals and then measures whether the IT infrastructure is, in fact, providing the right environment to support them.

The good news is that the BI monitoring tools are now featuring much more intelligent capabilities in terms of their strategic nature and ability to predict problems before they occur. For example, BEZ Systems helps companies plan, manage and scale their large data warehouse environments by analyzing all elements of the systems (data, workloads, resources, etc.), thus providing a correlated view of the entire system and the interrelationships between these elements and overall performance.

It is important to understand that a change in one area of your technological environment could adversely affect performance in another area. If you only have a very narrow view of your system's overall performance, the probability of de-optimizing one section while you optimize another is high.

To ensure that this does not happen, begin by focusing on trends and not on exceptions or point problems. SPM has the ability to determine whether an event is part of a trend or is a discrete, one-time event. We must react to a trend but may simply observe an event. Historical performance metrics are mandatory for SPM and provide the basis to determine trends from events. They also give us the basis from which to predict outcomes of possible alternative solutions.

The next step is to evaluate and understand the alternatives. Companies often face the challenge of choosing between two alternatives to solve a performance problem. Do you add additional CPU horsepower, or should you add additional disk storage across more CPUs? What about adding an additional server to split the workload? Each alternative has its own price/performance/return on investment (ROI) tradeoff, so the ability to evaluate and assess each one is important.

Finally, you can now predict the impact of proposed changes to the systems. SPM should have the ability to compare the various alternatives and suggest the best one based on factual data, not "gut feelings."

Once you understand the implications of your alternatives, expectations can be set and you can *measure* actual results once the change is implemented. This enables the IT department to have fact-based conversations with the business community about the overall performance of their BI systems.

Measuring an SPM tool's ROI can be difficult. The new capabilities of second-generation SPM tools are designed to increase an intangible area – your business community's satisfaction. However, this does not mean that these tools cannot demonstrate a positive effect on the corporation's bottom line. Preventing or diverting a potential performance problem by predicting it has definite measurable ROI: reduced downtime; higher retention rates and utilization by the business community; higher systems reliability; and reduced IT personnel for trouble-shooting.

SPM is the next logical step in the evolution of CIF performance management. It provides the foundation for creating sustainable and repeatable performance improvements to maximize IT's ROI for the entire organization.

Claudia Imhoff is the president of Intelligent Solutions, Inc. Imhoff is a popular speaker and internationally recognized expert on the Corporate Information Factory, business intelligence and CRM. She may be reached at cimhoff@intelsols.com. David Imhoff is an executive vice president of Intelligent Solutions, Inc., a leading provider of business intelligence and its architecture, the Corporate Information Factory. He may be reached at dimhoff@intelsols.com.