

Clear the HURDLES

Changes in the role of the data warehouse create challenges on the fast track to active enterprise intelligence. *by Bill Tobey*

Everywhere you look today, established data warehouses are extending their functionality to enable active enterprise intelligence by taking on a wider range of mixed workloads,

including support for front-line systems and business processes. The transition of these quintessential back-office systems into mission-critical roles creates cultural and technical challenges.

Retraining the business users

One of the first issues an IT organization is likely to encounter as it edges a traditional data warehouse into a more operationally intensive analytics environment is the need to reset the expectations and demand reflexes of its business users. Many on the business side retain attitudes developed when analytical data was scarce, expensive and infrequently refreshed. After years of being told that the data they wanted every two hours could be delivered only once a day, they are understandably impatient and not inclined to be selective in their requests. In essence, they have been conditioned to take everything IT can give them, as fast and as often as they can get it, even if they have not yet decided how to use it.

Now that IT has the physical resources to provide more data more frequently from an active data warehouse, it needs to be sure it is delivering only what is really needed—and at an interval that correlates to the value that will be derived. All users need to understand that costs and complexity still escalate with volume and frequency, just as they did when monthly or weekly loads were the norm. Discretion is essential in determining which data attributes provide true business value. The objective should always be just the right data at just the right time. That represents a significant culture change for the business.



Photography by Getty Images

Re-synchronizing dependent workloads

On the technical side, several challenges arise that are related to shorter refresh cycles and the time variance of data. The first occurs with the transition from any load interval longer than daily to anything shorter. A number of issues must be taken into account, including the fact that the IT department can no longer count on the stability of weekend extract, transform and load (ETL) windows, so it suddenly has to compress everything into a window that is usually only four hours or less. The fact is that when people ask for something daily, they really want it ready first thing in the morning, and the IT department cannot take all night to process it.

dependencies on the data warehouse is that the system itself becomes mission-critical. However, all too often it's not adequately prepared. Many data warehouse systems were originally deployed as off-line tools for strategic analysis. The systems' primary role has been after-the-fact analysis of events and predictive modeling and forecasting based on historical information. To support mission-critical operations, the data warehouse should be engineered for high availability, current, capable of supporting the operational workloads and designed for automatic failover.

A useful definition of a mission-critical system is a system that, should it fail to operate as expected, causes an immediate

Document the return on investment of active enterprise intelligence

Finally, IT departments and their business-side customers both need to remember the importance of documenting the return-on-investment (ROI) impact of active enterprise intelligence investments. IT in particular is accustomed to running operational systems that do not require ongoing business justification. Although the advantage and necessity of moving data to the front lines for smarter, faster decision making are now being realized, active enterprise intelligence requires ongoing justification.

To maintain the developmental and investment momentum that is moving

“The point is that many organizations are creating operational system dependencies on [data] warehouses that aren't currently engineered to support the operational systems.”

—Mark Beyer, research director, Gartner

The second big hurdle occurs when load intervals are shortened to less than an hour, because all of the rules change again. Technical challenges arise with maintaining data integrity and keeping optimization layers up to date because the data warehouse starts to exhibit some of the characteristics of more volatile transactional systems. “At this point you need the capability of ensuring point-in-time comparisons of data,” explains Gartner Research Director Mark Beyer. “By that I mean, if you run a transaction now, and then run the same transaction again an hour from now, you need to know what changed.”

Hardening the data warehouse for high availability

Another challenge that develops as operational systems acquire links and

loss of revenue, an increase in costs or a degradation of the customer service experience. This system must be replaced by people performing the process or the business will suffer the consequences. If that definition is applied to the operational data warehouse, it suddenly becomes clear which high-availability attributes are required of the data warehouse platform. It also helps clarify which attributes, tables and processes need to be backed up in this failover-protected, disaster-recoverable state.

“These may not be big systems, and the revenue and customer impacts may not be large,” says Beyer, “but the point is that many organizations are creating operational system dependencies on [data] warehouses that aren't currently engineered to support the operational systems.”

the data warehouse to the front lines, senior management needs continuous documentation of the value it creates. Businesses should watch for situations where analytics have changed operating practice, and document the revenue, cost or customer service impacts. At the same time, executives should be alert for situations where the business suffers because routine intra-day decisions are delayed for lack of data or by manual analytics that could be easily automated. Every such opportunity is a chance to deliver new value to the business and increase the return on its existing active enterprise intelligence investment. **T**

Bill Tobey, a senior technology writer based in Salt Lake City, covers the business applications of information technologies.