

Taking a Fresh Look at Business Intelligence and Data Warehousing

Part I

By Alan R. Simon

Permission is granted by the author to ***Technology Times*** (a publication of *Eastern Technology Council*) to print this article in any upcoming issue.

Introduction

The “modern data warehousing era” that began in the early 1990s can best be summed up in a single phrase: “A legacy of costly mediocrity.”

Upon reading those words you might be thinking something along the lines of “What blasphemy!” After all, *everyone* has data warehouses and data marts all over their organization, providing reporting, querying, and other forms of analysis related to customers, financial operations, personnel and human resources, manufacturing operations, supply chain operations, and many other areas important to companies and governmental organizations.

Problems or Symptoms?

So what’s the problem? Simply this: the vast majority of data warehouses and data marts actually deliver *little or no business insight and, consequently, little or no business value*, despite being up and running, filled with lots of data, and otherwise “technologically correct.”

The problems with most data warehouses' business value (or lack thereof) are typically espoused along the lines of:

- Incorrect or missing data
- Insufficient level of detail among the data warehouse's content
- Dozens or even hundreds of reports being produced, but those reports aren't of much value

...and probably a hundred more common complaints about an organization's data warehouse implementations.

So what happens the next time around? Those data warehousing practitioners bearing the "battle scars" of previous less-than-stellar systems do their best to learn from the shortcomings of those earlier experiences, and they try to ensure that the new data warehousing project plan addresses those issues. For example, they will design additional levels of data quality assurance (QA) routines into the extraction, transformation, and loading (ETL) processes to try and address data integrity issues. Or they spend additional time in the early stages of the project to delve more deeply into the organization's "real" requirements, trying to avoid reproducing existing seldom-used reports as part of the new data warehousing project.

But it's been my experience that time and time again, the next time around – that is, the next data warehousing project an organization tackles – winds up with many, of not most, of the same problems as the less-than-satisfactory data warehouse implementations that are already doing little but draining increasingly scarce computing and personnel resources and *not* doing what they were intended to do: produce actionable, high-value insight into some specific portion of an organization's processes

and functions. In fact, my contention is that the shortcomings I listed above – incorrect data, insufficient detail, and so on – actually *are not* the real problems with data warehousing and business intelligence as we know it, but rather simply *symptoms* of two “mega-problems.” Further, by focusing on these two “mega-problems” that I’ll discuss next, the chances for success of *any* data warehousing implementation are dramatically increased.

The two “mega-problems” I encounter time and time again are:

1. Misalignment between an organization’s approach to business intelligence and its organizational structure and culture; and
2. Focusing on data warehousing architecture instead of business intelligence architecture

The organizational misalignment challenge is discussed below, and in the next issue I’ll discuss the “which architecture?” quandary.

The Challenges of Organizational Structure and Organizational Culture

The best way to describe the twin issues of organizational structure and organizational culture is through a real-life example. Back in 1998, the corporate IT organization of a sizable mid-Atlantic regional bank decided to develop a bank-wide data warehousing strategy and architecture. The going-in position on the part of the bank’s IT staff members in charge of data warehousing was the typical enterprise-wide data warehouse blueprint: build a “big bang” all-encompassing system that would contain data drawn from applications in its seven different lines of business (retail banking, private banking, corporate banking, mortgage banking, and so on). Once this behemoth was

built, executives and staff members from all across the bank could do customer cross-referencing (example: find the names of senior executives from all corporate banking customers so we can market private banking services to them), various types of consolidated profitability analysis that included (and ranked) all the bank's lines of business, and so on. All in all, everything the bank's IT data warehousing staff was trying to do made sense: at least on paper.

There was one teensy-tiny problem, though: the bank's organizational structure and, as a result of that structure, the underlying culture. For starters, each of the bank's lines of business was headed by its own Chief Executive Officer (CEO), and each line of business had its own Chief Information Officer (CIO) with his or her own line of business IT staff. Then, to make matters worse, each of these line of business CEOs was basically competing with his or her counterparts to become the next overall CEO of the bank. Likewise, each of those seven CIOs had his or her eyes on the corporate CIO position. In short, despite the best intentions of the corporate IT organization to "do what was right for the bank overall" and build a data warehousing system that provided deep analysis across all of the bank's lines of business, most of the key stakeholders in the bank had little or no incentive to extend cooperation on the enterprise-wide data warehousing initiative. Then when the global financial crisis hit in the late summer and autumn of that year, the "powers that be" in the bank simply used the financial turmoil as justification to mothball the idea of bank-wide data warehousing and go back to where they should have concentrated their efforts in the first place: superior business insight and intelligence capabilities *within* the bank's lines of business, rather than a futile effort to go against the grain of the bank's organizational structure and organizational culture.

Another quick example: about a year later I worked briefly with one of Canada's largest banks on an eerily similar initiative: a bank-wide data warehouse across multiple lines of business. The difference with this Canadian bank versus the U.S. mid-Atlantic bank, though, was that the bank in Canada had already built an enterprise-wide data warehouse...and from an architectural sense, it was a pretty good job! Alas, the culture within that Canadian bank was pretty much the same as that of their American counterpart, so what we found was that – get ready! – *nobody* anywhere in the Canadian bank was using the bank-wide data warehouse! Organizations were still doing their own small-scale data marts or extract files for reports; developing their own reports drawing data directly from operational systems that they “owned;” and, basically, ignoring the bank-wide data warehouse as if it didn't exist at all.

The lesson: from a technological sense, there is nothing wrong with pursuing enterprise-scale data warehousing, but it's been my experience that *very few* organizations have *both* the organizational structure and (I think even more importantly) the organizational culture to make enterprise-wide data warehousing work. Organizations that have highly decentralized chain-of-command structures and a culture that encourages and rewards success on a “micro” level – that is, measurable success within a geographical region or line of business or some other “sub-organizational” level rather than the company as a whole – should architect their business intelligence capabilities and the underlying data warehousing implementations to reflect that fact rather than fight the inevitable losing battle.

Alan R. Simon is an independent consultant specializing in strategy and architecture for business intelligence and data warehousing. He formerly was a member of the data warehousing leadership team at Deloitte Consulting, a “Big 5” consulting firm, and previously was Vice President of Cambridge Technology Partners’ data warehousing solutions practice. He is the author of 26 books, including *Data Warehousing for Dummies*, *90 Days to the Data Mart*, and *Data Warehousing and Business Intelligence for E-Commerce*, and was a regular contributor to and the monthly data warehousing columnist for *Database Programming & Design* magazine.