



#9: Limited Brands

Unconventional Wisdom

Limited Brands finds **treasures** hidden in a data warehouse

FEW ARGUE THAT A WELL-done data warehouse is not an effective business tool, but Limited Brands discovered that it can also help defeat conventional wisdom. The \$9.5 billion apparel and lingerie giant learned that data warehouses could help their managers think in non-Limited ways.

The Limited's brands include Victoria's Secret, Bath & Body Works, Express, Lerner New York, The Limited, Henri Bendel, The White Barn Candle Co., and Aura Science.

Conventional wisdom for catalogue sales has long held that a retailer must set a strict threshold for how much a customer must buy from that catalogue in six months before the customer is automatically sent another one. That's how The Limited did things, too — until its data warehouses were merged and it could see a single picture of its customers, says Jon J. Ricker, president and CIO of Limited Brands Technology Services.

That combined image showed that many of those catalogue buyers who didn't appear to be buying anything — because the catalogue recipient never called the toll-free number and ordered — were in reality using the catalogues to pre-shop and would then drive to the local Victoria's Secret and buy plenty.

Power Data

When setting up a strong data warehouse, "it really does take a while before the power users on the data warehouse understand the smart questions to ask," Ricker says.

Making the link between those customers today is also more difficult than merely merging and analyzing databases. Retailers used to be able to do reverse-append on credit-card purchases, where the credit card company would spit back the name, address and other data after a card is swiped, but the Federal Trade Commission now bans the practice. Today, Limited relies on its private-

label cards as well as store personnel to collect the data.

Ricker's advice for creating a strong retail data warehouse? First, keep it small in the beginning. "We didn't go out and buy terabytes of storage. Our first had merely 200GBytes," he says of the structure Limited Brands built with Teradata (www.teradata.com) software sitting on an HP (www.hp.com) Unix server with Hyperion (www.hyperion.com) query tools reporting through a Windows client.

Secondly, insist that your business managers have substantial ownership. "The user community has to own it," he says. "For each one of our data warehouses, there is an officer in the business hierarchy who fully owns that asset. The responsibility in the business at Victoria's Secret, for example, is the CFO of Victoria's Secret. This guarantees that we'll clearly make a business-based decision and not a technology-based one."

Plan to Grow


Ricker isn't alone. To today's retail data warehouse manager, bigger is not necessarily better. But may be necessary. (RIS has reported on the more than 200-terabyte data warehouse developed by Wal-Mart — considered one of the best in the business.)

There are many realities that are pushing bigger. The torrent of data from e-commerce sites, tons of non-Web CRM data, marketing's — and sometimes legal's need — for historical data. Analysts project that today's multi-terabyte systems will quickly become multi-petabyte systems and then move into the exabyte neighborhood.

For many retailers well positioned in a strong space, the fear of having a data warehouse that can keep up with the company's growth is a very real fear. "You have to find some way to break the data down or you will not be able to continue to grow your data," said Gerry Claggett, the database systems administrator for the \$2.6 billion Burlington Coat Factory.

Executive Summary & ROI

Data warehouses built in silos can mask important opportunities, as Limited Brands discovered when it merged two data warehouses. They also have made the asset more strategic by placing a c-level officer over each data warehouse project. The best way to build a data warehouse is to start small, gather up all customer data and give a lot of thought to two-phase commit.



"For each one of our data warehouses, there is an officer in the business hierarchy who fully owns that asset."

- Jon J. Ricker, president and CIO, Limited Brands Technology Services

Pitfalls

One common mistake made with data warehouses is opting for the easiest way to set up and maintain the data without paying proper attention to how the data will be needed. Claggett concedes he initially made just such a mistake with Burlington's system, which manages about 20 terabytes of data.

"We did it the wrong way. We sliced it the wrong way," he said. "We initially had a table for sales and we had the sales at the day-level, the week-level, the month-level and the year-level all in one table." But when the clothing retailer started experiencing major growth, Claggett got nervous that his setup wouldn't work for long.

Claggett's new approach was better positioned for future growth, but it also brought more information to store managers right away. "We changed our layout of the data and we did this all in parallel," he notes. "The new approach was to have all types of transactions in one row for each store for each day. This meant that we were keeping our data at the style/color/size level instead of the style/color level. Our buyers were saying we could do our job better if we had this data."

Typical of Burlington Coat Factory, which constantly pushes the envelope, Claggett employed database partitioning before Oracle (www.oracle.com) even started offering it to customers.

Two-phase Commit

Another popular data warehouse approach is two-phase commit, whereby transactions are entered in multiple databases simultaneously. Although Claggett concedes that two-phase commit is needed for some retail purposes, he's not a fan. "If you can possibly avoid two-phase commit, avoid it," he says. His concern is that if the transaction fails at some stage, it may not be applied to any of the databases.

"Let's say I have a register sale and the register sale has purchased items from five different departments," Claggett says. If we had tried to apply one sales transaction to seven different databases at one moment and one failed, you might have incomplete data, but you wouldn't necessarily have incomplete data. You'd have an inconsistent state." **RIS**

Do you use two-phase commit in your databases? Drop a note to jhall@edgellmail.com.

By Evan Schuman, Contributing Editor