

TDWI

CASE STUDY

FEBRUARY 2007

Driving Insights and Decisions: Data Warehousing at Hudson's Bay Company

tdwi

About TDWI Research

TDWI Research provides research and advisory services to business intelligence (BI) and data warehousing professionals worldwide. Unlike other research or analyst firms, TDWI Research staff focus exclusively on BI/data warehouse issues, teaming up with industry practitioners and TDWI Faculty members to deliver both broad and deep understanding of the business and technical issues surrounding the deployment of BI/data warehouse solutions. TDWI Research delivers commentary, reports, and inquiry services via TDWI's worldwide Membership program, and provides custom research, benchmarking, and strategic planning services to both user and vendor organizations.

About TDWI

The Data Warehousing Institute (TDWI), a division of 1105 Media, Inc., provides in-depth, high-quality education, training, research, and certification for BI and data warehousing professionals worldwide. TDWI can help your BI team stay abreast of new and emerging trends and techniques and gain the skills they need to deliver effective BI and data warehouse solutions. Through our Membership program and regional chapters, TDWI can also help you and your team establish a network of peers in the industry to whom they can turn for assistance and advice on career development or to help solve thorny technical and organizational problems.

About the Sponsor

Teradata (www.teradata.com) is the global technology leader in enterprise data warehousing, analytic applications and data warehousing services. Organizations around the world rely on the power of Teradata's award-winning solutions to get a single, integrated view of their business to enhance decision-making, customer relationships and profitability.

Data Warehousing Drives the Business

As the oldest corporate concern in North America, and one of the oldest commercial entities in the world, the Hudson's Bay Company (HBC) is a survivor.

In colonial times, HBC helped end the French stranglehold on the North American fur trade and later monopolized it in the 19th Century. Fur isn't quite the moneymaker it once was, but HBC hasn't missed a beat. In fact the HBC of today is one of the largest retailers in Canada, employing more than 70,000 people and operating more than 500 retail stores—including familiar banner brands such as the Bay, Zellers, Fields, Designer Depot, and Home Outfitters.

Key to Success. A large part of HBC's success is a data warehouse environment that provides a unified view of key business data across its divisions and product lines. The company's Teradata Warehouse provides thousands of employees detailed views of sales and inventory transactions across thousands of HBC products. According to Bill Ward, general manager of merchandise and support with HBC, this allows the data warehouse to give store managers, buyers, merchandisers, and external suppliers unprecedented insight into HBC's business performance and serves as an engine to drive decision making.

“The data warehouse is so integrated into day-to-day operations that [the impact of an outage] would be severe.” Bill Ward.

“The data warehouse is so integrated into day-to-day operations that [the impact of an outage] would be fairly severe,” says Ward. “It would take [decision makers] several days to get information that the data warehouse provides.”

Cross-departmental Resource. HBC's data warehouse environment drives a range of business processes in merchandising, inventory, and promotions. The marketing department uses the data warehouse to determine which products to advertise, the merchandising department uses it to decide how much inventory to maintain and how to allocate products across stores to meet customer demand, and store managers use the data warehouse to track sales and performance against plans.

Store managers have single-click access to right-time sales and inventory information.

More than 3,000 HBC users and over 700 suppliers currently have access to the Teradata Warehouse. That's within about 85–95 percent of “where we wanted to be at this point,” says Mary Jane Jarvis-Haig, senior manager of BI and customer relationship management (CRM). In addition, the data warehouse team has developed 100 different report templates for merchants and stores using MicroStrategy's reporting and analysis tools. Store managers, for example, now have single-click access to right-time sales and inventory information that's refreshed daily. This summer HBC started offering suppliers secure access to the same information.

“We provide users with an enormous number of prompts and filters they can apply to reports because we want to give them flexibility to see information in the way they prefer,” says Jarvis-Haig.

Driving Insights and Decisions

Jarvis-Haig says Teradata Warehouse is flexible enough to handle whatever HBC throws at it.

Better still, Jarvis-Haig says, Teradata Warehouse is flexible enough to handle whatever HBC throws at it. Monday morning is peak time: the system delivers as many as 9,000 reports to information-hungry users. During the rest of the week the system generates between 2,500 and 6,000 reports each day.

HBC's BI/DW Environment

Data Warehouse. HBC's data warehouse has grown from about 250 GB in the 1990s to nearly 6 TB of data today with 10 TB of overall disk space. It's currently powered by a 10-node NCR server running Teradata Database Version 2 R6.1. Historically HBC has added an average of two nodes every year, mostly to support retail expansion and new BI applications. Each upgrade, however, is cost-justified with quantifiable business benefits.

At present, the data warehouse stores all of HBC's detail, summary, sales, and inventory data, most of which is refreshed nightly via more than 100 ETL feeds from operational systems. Point-of-sale (POS) data is refreshed even more frequently: "We're loading constantly from the POS and querying a few hundred thousand times a day," says Jarvis-Haig.

The data in HBC's data warehouse comes primarily from the retailer's POS, inventory management, vendor management, and purchase order systems, which include several new packaged applications. "Most ERP systems don't do a good job of reporting, so we try to make the data warehouse the reporting mechanism," says Shelley Perrior, team lead and database administrator.

The data warehouse has improved the quality of HBC's data significantly.

The data warehouse has improved the quality of HBC's data significantly—although there's a definite human contribution here, too. HBC has business analysts on the QA team with strong knowledge of the business who apply regression tests and other techniques to ensure data warehouse data is valid and clean, according to Jarvis-Haig.

BI Tools. On the business intelligence tools front, HBC standardized on MicroStrategy in November of 2004. MicroStrategy in turn forms the basis of HBC's Business Information System (BIS): a collection of report templates designed to facilitate business decision making that will increase sales, manage appropriate inventory levels, and drive effective supply chain management and processes. HBC spent 18 months deploying analytic applications and finished in June of 2006.

Most of HBC's reporting takes the form of predefined, prompted reports, according to Kevin Gilbert, a senior technical analyst. This gives users the desired interaction with the data. And the number and variety of built-in prompts and filters creates a fluid and variable data warehouse workload, making it a challenge to deliver consistent performance for all reports.

“We have a lot of experienced ETL and Teradata developers who’ve created efficient code that delivers excellent performance and data quality,” says Gilbert.

ETL. Without the luxury of an ETL tool, HBC relies on SQL and mainframe jobs and Teradata’s utilities, such as MultiLoad and FastLoad. “We have a lot of experienced ETL and Teradata developers who’ve created efficient code that delivers excellent performance and data quality,” says Gilbert. HBC recently purchased an ETL tool from Informatica, which it will use initially on new projects followed by gradually replacing legacy SQL code, says Gilbert.

BI Team. To run its data warehouse environment, HBC has 21 staffers, including six full-time ETL developers, two full-time MicroStrategy programmers, and three full-time MicroStrategy administrators. These aren’t hard-and-fast positions, however. There is a lot of cross-pollination between the ETL and MicroStrategy developers, for example, with both groups chipping in as needed. In addition, the team also includes three DBAs, three QA testers, and four business analysts who gather business requirements, monitor development of these requirements, and provide training.

HBC BI Solution Center. To ensure strong user adoption and support, the data warehouse team leverages HBC’s BI Solution Center, which serves as a liaison between IT and business. Organizationally, the Solution Center is separate from the data warehouse team, but the two groups work hand in hand. The Solution Center functions as a single point of contact for users to get help with an application or submit a request for an ad hoc report. It also functions as a vehicle to evangelize the potential of data warehousing and BI to drive decision making. “The Solution Center gives us the visibility and partnership with the business to make the data warehouse successful,” Jarvis-Haig explains.

Evolution of a Solution

From Operational to Integrated Reporting

HBC began its data warehousing journey in the 1990s, building its first data warehouse with help from Texas-based consultancy Knowledge Discovery One (KD1). “We started small with a data analyst, myself, and some consultants and a Teradata box located offsite in Texas,” explains Kevin Gilbert, a senior technical analyst with HBC. KD1 consultants initially built the data warehouse, Gilbert continues, while HBC “pumped” sales and inventory data down to KD1 via FTP on a daily basis. For BI reporting, they used KD1’s proprietary tool, Basket Dynamics.

At the time, HBC’s data warehouse functioned as a peripheral system rather than a central resource for decision-makers. In fact, says Jarvis-Haig, buyers, merchandisers, and store managers still depended on reports generated from operational systems and a first-generation mainframe reporting application. At that time, its operational environment consisted of a variety of systems running on different platforms capturing data at different intervals.

Driving Insights and Decisions

“There were 17 different versions of the truth,” says Ward.

The upshot was that merchandising information was spread across several different operational systems. Custom reports took up to a week to generate, and, even then, were incomplete or misleading. As a result, buyers, merchandisers, and store managers made crucial merchandising decisions on the basis of instinct, not analysis. “There were 17 different versions of the truth. And depending on whose source system you ran, you had sales that were from Monday to Sunday or sales from Saturday to Tuesday,” Ward explains.

Consolidating and Centralizing

Enterprise Data Warehouse. In 1999, HBC moved the data warehouse—which then consisted of one data mart—to its corporate headquarters in Toronto. Over the next couple of years, the team added data marts for each business unit. In 2003 the data warehouse team convinced HBC’s line-of-business sponsors to consolidate the data marts into a single enterprise data warehouse with a unified view of all HBC’s data as well as upgrade to a newer version of the Teradata Database.

“It was a business-driven [decision.] HBC senior executives determined that our customer needed a more integrated offering and approach to detail data collection and analysis,” says Gilbert.

“We didn’t know how much [users] had integrated [a legacy BI tool] into their processes until we took it away from them.”

Getting BI Right. Also in 2003, HBC decided to convert from the original BI application, which KD1 no longer supported, to a more current BI technology platform and toolset. But changing proved more difficult than the data warehouse team expected: “There were 500 users using the tools, but we didn’t know how much they had integrated it into their processes until we took it away from them,” says Jarvis-Haig. HBC evaluated a number of BI tools. Since HBC users wanted detailed data, HBC decided to look for BI tools that would run directly against the data warehouse instead of outboard cubes. MicroStrategy eventually won out, mostly on the strength of its all-in-one architecture, support for self-service reporting, and thin (Web) client features.

Addressing Performance Issues. Consolidating data marts into an enterprise model and simultaneously implementing a new application on a new BI technology platform was overly ambitious. “To change the underlying data warehouse *and* the BI tool at the same time was a big undertaking,” says Jarvis-Haig. Beyond managing project scope, the primary challenge was to ensure adequate performance now that it would be running all queries and reports off a single data warehouse instead of multiple marts and cubes.

HBC spent many hours tuning both MicroStrategy and Teradata Warehouse to alleviate hiccups and ensure adequate performance. HBC built aggregate views, for example, because it wanted to make more effective use of Teradata’s indexing capabilities. In addition, HBC extensively tuned Teradata Warehouse by running explains, combing through statistics, and indexing to ferret out performance issues. “Those are standard tuning practices,” says DBA Perrior. “We also tried a few betas to get [to] a more advanced version of the database to help us,” she explains.

HBC uses flexible report prompts to give users virtually any view of the data they desire, which exercises the DW significantly.

One reason HBC places such emphasis on tuning, says Gilbert, is that it asks the data warehouse to do more with less. Tuning provides HBC users a lot of flexibility with report prompts, allowing them to get virtually any view of the data they desire, which exercises the Teradata environment significantly.

Driving Business Value

User Adoption. Once they made technology selections, the team’s key challenge was to ensure strong end-user adoption of the tools. The key here was to educate business executives about what they could possibly do with the tools. Many business executives and managers were not accustomed to working with integrated data and BI tools and did not understand how to use the tools to increase work efficiency and effectiveness.

Fortunately, Jarvis-Haig says, the team received a boost from HBC’s new owner, who is a strong top-down advocate for both the data warehouse and the BI initiative: “The new owner is tech savvy and wants to ensure HBC uses the latest technology to enhance our operations.” For example, HBC is accelerating the step to provide outside vendors access to their sales and inventory data. “Our new owner has said to HBC vendors [that] we are looking for a stronger partnership and we want to be more accountable and proactive in making each other more successful. So the message is, ‘we’ll give you what you need to help you be successful,’” she comments.

The DW has improved inventory, buying, and promotions, and has helped reduce lost sales and excess inventory.

ROI. Although HBC has not yet quantified the ROI of the new environment in dollars and cents, it did get a two-for-one return on invested dollars when it tracked benefits from 1999 to 2003. In addition, the combined Teradata and MicroStrategy stack has contributed to improved inventory, buying, and promotions, and has helped reduce lost sales and excess inventory. HBC continues to discover new ways to exploit the wealth of information in the data warehouse, such as identifying ethnic product opportunities by geographic region and asking “what if” questions about pricing and markdowns. Furthermore, HBC used its data warehouse and data mining software to break up a fraud ring responsible for \$26,000 in losses, which HBC completely recouped.

Support. There are plenty of intangibles, too, such as Teradata support. “The support from Teradata has been great. I think they are really good, probably better than the other database vendors I’ve worked with,” Perrior comments.

Teradata’s support helps justify—if not offset—its premium pricing.

Perrior says Teradata’s support helps justify—if not offset—its premium pricing, at least relative to competitive solutions. “The cost of their product *is* high. But the support is great. I’m on their service focus team and they’re constantly working to improve the service they provide.”

Teradata is also cheaper to manage, says Perrior. For example, she points out that in most other RDBMS platforms she spent a lot of time managing space. Not so with Teradata because it “manages its space very well by itself.” Nor does it require reorgs. Because it is massively parallel, it can perform tasks such as loading or accessing data or doing tablespace scans much faster than other

platforms. “Those features make it quicker and more efficient than other major databases,” Perrior comments.

Future Reality

HBC isn’t resting on its laurels. Ward and other officials are determined to expand the business benefits of the data warehouse. HBC plans to implement standardized business rules, exception-based reporting, and metadata management capabilities, as well as integrating customer data into the data warehouse.

Metadata management. Metadata management poses a thorny challenge. In point of fact, HBC doesn’t actually have a central metadata management solution. “Today, our documentation is somewhat fragmented,” Gilbert allows. HBC plans to use its newly acquired Informatica ETL to assist with metadata management.

The absence of a metadata tool makes it difficult for administrators to perform impact analyses when changes occur in source systems or reports. It also can be frustrating for end users who often ask, “Where does this data come from? How is it processed?” according to Gilbert.

“Right now, HBC is examining the metadata capabilities of various vendor tools, but hasn’t made any decisions yet,” Perrior stresses.

“Exception-based analysis will let users start dealing with things at a level that is actionable,” says Ward.

Exception-Based Analysis. The team is also planning to embrace exception-based reporting and dashboards. “Exception-based analysis will enable users to start dealing with things at a level that is actionable,” says Ward. Ideally, MicroStrategy will generate exception-based reports, recommend actions, access source systems at the click of a button, and perform the actual changes recommended by the system. “Whether it is automated change or user acceptance of [a proposed] change, we haven’t started that yet,” Ward says.

Many business users simply aren’t comfortable with automated—or even quasi-automated—decision making, says Haig. “The real driver to go to [MicroStrategy] was exception-based reporting, but it became apparent that our [users] weren’t ready for that. They were not prepared to trust the machine to say ‘I know what the problem is and here is where you look to fix it.’”

For now, HBC has implemented exception *highlighting*, which gives users details in a color-coded format. Whenever possible the BI development team delivers a more exception-like report to keep the evolution moving forward.

HBC plans to establish KPIs and reports that show workers how well they measure up in their domains.

KPIs. Business sponsors and BI team leaders are also hard at work developing company-wide KPIs. “One of the things that we really need to be better at is holding people accountable for the part of the jobs that they do and then measuring their performance,” Ward says. HBC plans to establish representative KPIs and build reports that show users how well they measure up in the domains for which they are accountable. “This is what the dashboard reinvention will be.

It is a common dashboard, common KPIs, common interpretative data drilled down to the level of action,” Ward explains.

CRM. HBC also hopes to achieve full integration between its separate CRM and data warehouse environments. For a long time now, the retailer has collected information for its credit card business. Some of this information—such as detailed transaction data—is stored in the data warehouse, while other data (customer credit scores, profiles, and segments for example) resides in a separate Oracle database. Each month, the data warehouse team dumps summary data from the Teradata Warehouse into a customer information file on Oracle. “It’s on Oracle because of historical decisions” Jarvis-Haig comments.” Bringing this data together will enable the company to better understand its customers and target them more effectively.

Conclusion

As confidence in the new BI platform grows, HBC stakeholders are eager to roll up their sleeves and leverage the BI applications. Ward, for his part, has even bigger plans. Like any information-hungry consumer, he’s far from satisfied.

Ward wants to transform decision making to more proactive data capture and decision making.

“We came from a paper-based environment.” He wants to transform decision making to more proactive data capture and decision making using dashboards that show performance against plan. “This way, I can build them reports so they aren’t filtering and manipulating thousands of rows of data when they only need to look at 1 or 2 percent.”

HBC’s management team has embraced many new technology options in support of its strategy to become Canada’s preeminent retailer. As Ward says, “It is just an educational thing of showing people a new way of operating. First we have to get them stable with their requirements today and efficiently provide them with the information. And we are now getting to that state where we have pretty good performance and good response time,” he concludes. The next step is to move the business toward exception-based reporting where they only need to comb through data that requires immediate action.