DATABASE
TRENDS AND APPLICATIONS
Solutions for the Information Project Team • www.dbta.com
Volume 21, Number 3 • March 2007

DW Improves Healthcare

New system provides critical capabilities.

By Walt Jordan

Harvard Pilgrim Health Care is a non-profit health plan provider that serves more than 975,000 members in Massachusetts, Maine and New Hampshire. Like many health maintenance organizations and health insurance companies, Harvard Pilgrim grew rapidly through acquisitions during the 1980s and 1990s, and then faced the daunting task of integrating the acquired companies' IT systems. Indeed, in the early 2000s, Harvard Pilgrim had 55 core application systems, including four different claims processing systems.

As a part of the restructuring of its IT systems in the fall of 2001, Harvard Pilgrim began to discuss the cost of supporting its IBM mainframe, which ran certain core applications and housed the data warehouse. During the IT consolidation, all of the company's business applications were moved to an Oracle platform. Over time, the data warehouse was the only application left on the mainframe. "We felt that the mainframe was not the most suitable for a data warehouse, and it was fairly expensive," said Ralph Miller, director of corporate information management at Harvard Pilgrim.

Since Harvard Pilgrim moved to Oracle for all its business applications, the original intent was to move the data warehouse to Oracle as well. But in the course of examining the requirements of moving to Oracle, the Teradata data warehouse offering was included in the evaluation process. "We went through an extensive evaluation," said Miller. Harvard Pilgrim talked to references for both Oracle and Teradata, and visited their benchmark centers, where they presented challenging and complex queries for execution.

"These were things that we had in production and we had huge production issues with the IBM mainframe. Teradata was able to run the queries without any changes very quickly and had much better performance," Miller recalled. When the same queries were run in the Oracle environment, the queries had to be refined and optimized to improve the performance. And even then, the performance was not as good.

Harvard Pilgrim also examined the internal support and management requirements of each platform, as well as the total cost of ownership. "The Teradata platform took fewer DBAs and was simpler to support for data warehousing," Miller said. Moreover, his team studied Teradata's scalability and ability to support companies like Wal-Mart.

After convincing management to go with Teradata - even though the company had standarized on Oracle - the team launched a multi-step migration process. The first step was to lift and drop the existing warehouse from the IBM mainframe to the Teradata platform. "We lifted and dropped the structures, but all the code to populate the structures, the ETL, had to be re-engineered and rewritten," Miller said. The ETL code had been written in COBOL. Miller wanted a more modern ETL tool to set the stage for future development and opted for Ascential Data Stage.

At this time, the IT group learned how to build an ETL process and integrate it with a metadata management process. "It was a good learning experience and allowed us to move from one platform to another," Miller said.

The second stage of the process was to build an enter-

prise data warehouse (EDW). "We wanted to build an EDW that was not built specifically for the primary transaction system," Miller said. "We wanted to be able to migrate away from the transaction system and integrate other source systems."

A small team of modelers, with assistance from Teradata, developed an enterprise logical data model, which was validated by the business experts. The next step was to build the physical structures of the EDW and then migrate the data to the new system. But the EDW was not built in isolation. "When we built the EDW, we also built a major application for finance," Miller said. The application, called Trends on Demand, generated on-demand answers to questions that had taken weeks to answer before. The company has a trend team with a staff of analysts supporting them. In the past, analysts could take two to three weeks to gather analyses to present to the trend team and additional time to answer follow-up questions. The new application generated the initial reports and offered drill-down capabilities that allowed analysts to answer 90 to 95 percent of the follow-up questions instantaneously. "We can provide information more quickly. We solved the problem," Miller said.

With that application, Harvard Pilgrim developed a constituency for the warehouse. Currently, Harvard Pilgrim is completing the migration, including the migration of SAS code, to the new system. It has also launched a formal adoption program so users feel comfortable when the old sytem is phased out. "We are looking at key groups that have influence over the user base and at opportunities to pull in people who were not using the old warehouse," Miller said.

The data can be refreshed much more frequently now. The architecture is designed for interday refreshes, but currently refreshes take place weekly (compared to monthly in the old system) and the front-end is more friendly. Miller hopes users in operating areas will use the warehouse. Executive management is another new target user community.

The EDW is key to addressing changes in the insurance industry. Large companies are increasingly self-insured and disparate data must be integrated into the data warehouse. "Every time they contract with a vendor, we have to send them a data extract," Miller said. "We source 95 percent of the extracts from the EDW."

In conclusion, Miller said, "the performance people get has been absolutely dramatic. When we did the initial lift and drop, we wiped out $1 million in infrastructure cost and improved performance. We had reports that could take days to run that now take minutes."