

# Baseline

HANDS ON

## Community

### The Know-It-All in the Server Room

BY TODD SPANGLER

**A high-speed data warehouse that can crunch billions of data points in less than an hour can mean the difference between getting a complex question answered today, or having to sleep on it. But is quick turn-around worth millions of dollars?**

Warner Home Video thought so. For a major release, the company has one shot to get right the number of DVDs and videotapes it ships to retailers. That's because between 40% and 45% of a video's total sales typically occur in the first week of its release, says Thomas Tileston, Warner Home Video's executive director of data management.

"If we don't ship enough and the retailer ends up with empty shelf space, that's a missed opportunity," he says.

The Burbank, Calif.-based video distributor, which sold 68 million DVDs of the three *Harry Potter* movies through the end of last year, had industry-leading sales of \$4.75 billion in 2004, according to *Video Business* magazine.

Two years ago, though, Warner Home Video couldn't accurately estimate how many copies of a video to make. On average, it was off the mark by 40% in 2003. In other words, it produced 40% more or fewer copies than the number actually bought. Unsold inventory is returned to Warner Home Video for storage or disposal; either way, the company loses money.

The problem? "We were forecasting titles one at a time," Tileston says. "And it was really, really hard to gather the data."

Previously, the company's forecasting process consisted of compiling sales figures from dozens of sources and cross-referencing data on industry trends (such as, say, how well romantic comedies have sold in Florida). Tile-

ston's group would then run the numbers through a sophisticated data-analysis program built with software from SAS—which could take as long as 22 hours to run on a desktop computer.

With millions of dollars in revenue at stake, the company needed to predict demand more precisely. So in 2003, it deployed a data warehouse system from NCR's Teradata division for about \$5.5 million.

Warner Home Video's data warehouse now holds about a trillion pieces of data, with spare capacity for eight times that amount. The system consolidates all the information that was floating around separately. That data

also is standardized, so the figures the marketing department uses are the same ones accessed by the sales staff. In addition, the system allows unit sales data from retailers to be loaded in daily, so it's never more than a day old.

And the data warehouse spits back answers quickly: Tileston says the report that once took 22 hours to generate takes 1 hour and 15 minutes on Teradata. "We're doing analytics that just weren't possible before," he says. Overall, according to Tileston, the system has helped Warner Home Video improve the accuracy of forecasts, to within 13% of actual sales last year and within 5% through the first half of 2005.

A data warehouse is a database specially designed to look at the big picture, without slowing down the main systems that keep the business run-

#### GROUP DYNAMICS: DATA WAREHOUSING

**What it is:** A system that merges data from multiple transaction-oriented databases or other sources, to provide highly detailed analysis of business information quickly.

**Key Players\*:** IBM, Microsoft, *Netezza, Oracle*, Sybase, *Teradata*

**Market Size:** \$8.7 billion worldwide, 2004 (IDC)

**What's Happening:** At the high end of the market, Teradata faces aggressively priced offerings from Netezza and other upstarts. Meanwhile, Oracle and IBM are moving up the performance curve, and SAP has put more weight behind its data warehousing products.

**Expertise Online:** DataWarehouse.com offers discussion forums, news, white papers and other resources.

COMPANIES IN *ITALICIZED ORANGE TYPE* ARE FEATURED IN DOSSIERS THIS MONTH.

#### Where businesses use analytic applications



SOURCE: THE DATA WAREHOUSING INSTITUTE, MARCH 2005  
REPORT BASED ON A SURVEY OF 473 DATA WAREHOUSING AND  
BUSINESS INTELLIGENCE PROFESSIONALS

ning. For example, a traditional database might be used to enter and retrieve customer orders, while a data warehouse would be used to analyze groups of them. Vendors include specialists such as Teradata and Netezza, and larger players like IBM and Oracle, which have versions of their transactional database systems designed for data warehousing.

Powerful data warehouses can let companies conduct more what-if scenarios with their data—and also analyze those test cases in more detail—than they feasibly could using a conventional database. A high-end data warehouse offers “faster speed to failure,” says Mike Coakley, vice president of marketing technology at Epsilon, a database marketing firm. “The key question is, how many times can I fail before I get the right answer?”

Carlson Hotels Worldwide, for example, wants to create marketing programs that increase the frequency of its customers’ visits. Jeff Heine, the company’s director of customer strategy, plans to use a Teradata system to determine, say, which pricing promotion for the top 10% of the 15 million guests who have stayed at the company’s 895 hotels will generate the most incremental revenue. “The data warehouse allows us to have more educated arguments,” Heine says.

But to others, a high-performance data warehouse is a luxury they can’t afford. Teradata, the segment’s grande dame, charges \$1.1 million for its fastest two-server system with 750 billion bytes of storage. Even Netezza, a data warehouse startup trying to make headway by undercutting Teradata’s pricing, starts at \$300,000 for a system with one trillion bytes of storage.

Quicken Loans, an online mortgage lender in Livonia, Mich., that closed \$12 billion in loans in 2004, opted to use Microsoft’s SQL Server for its data warehouse three years ago. Eric Lofstrom, manager of business intelligence, picked Microsoft for the



JOE PIGNATELLO  
SENIOR DIR., ENTERPRISE DATA SERVICES  
CINGULAR WIRELESS  
ATLANTA, GA.  
WWW.CINGULAR.COM

## VOICE OF EXPERIENCE

## Data Junkies

**MANAGER’S PROFILE:** Responsible for business intelligence systems and software, including the data warehouse infrastructure, at the largest wireless carrier in the U.S. Last year Cingular bought AT&T Wireless, and the combined company has more than 50 million subscribers and generated \$34 billion in revenue for 2004.

**HIS PROJECT:** Pignatello’s team maintains about 100 Teradata servers, which provide 8,000 employees with access to subscriber profiles, network usage patterns by geography and other information. Today the data warehouse contains 35 trillion pieces of data, and by the end of the year Pignatello expects that number to triple as Cingular incorporates even more information into the system, including detailed records about every single phone call placed on its networks.

**FEED THE BEAST:** The data warehouse is packing on the weight because Cingular’s business managers see new areas where they can merge and cross-reference data. “As you start growing the data warehouse, people start jumping on the bandwagon—they start finding out how they can use the data,” Pignatello says.

**HOUSE CALLS:** Cingular, for example, compares its own customer database with consumer information provided by Acxiom, a data reseller. The carrier uses the results to tailor promotions to an existing customer’s household, such as offering a family-based calling plan. “It’s a more cost-effective way to do marketing,” Pignatello explains.

**CHURN CHOP:** Pignatello believes the expanded data warehouse will help Cingular reduce “churn,” the rate at which customers cancel their service. For instance, managers can build forecasts for new rate plans and test the effect on churn versus profitability. On the operations side, Cingular can examine whether customers in a particular area may be canceling service because of dropped calls and use that analysis to improve coverage.

**FAST ANSWERS:** However, Pignatello says, it’s critical to maintain the data warehouse’s high performance even as Cingular crunches larger chunks of data. “The key to any kind of success with a data warehouse is going to be: How quickly can I get the data in the hands of the people making decisions in the company?” he says. “We’re information junkies.”

—TODD SPANGLER

project because he already had a full license for the database and because the software runs on inexpensive, Intel-based servers. He says the first version of the data warehouse went live for less than \$100,000.

“If you have a blank check, you might go with Teradata,” Lofstrom says. “But we didn’t have that kind of capital floating around.”

And sometimes, the main point of a data warehouse isn’t to dissect data as fast as inhumanly possible. Denver Public Schools deployed its Oracle-

based data warehouse to consolidate information from 160 student information systems and provide a single view for each of its 72,000 students.

Before, a teacher would have to run 15 different reports in the student-information system to get a complete list of grades, absences, disciplinary actions and other information, according to Kent Graziano, manager of the district’s enterprise data integration team. “Now,” he says, “we provide one-stop shopping in the data warehouse.” ◀

Reprinted from Baseline, July 2005 with permission from Ziff Davis Media Inc.  
©2005 Ziff Davis Publishing Holdings Inc. All rights reserved.