

BUYER CASE STUDY

Reorienting the Data Warehouse: Improving the Business Analytics Capabilities with Sybase IQ Columnar Database

Dan Vesset

IDC OPINION

A growing number of organizations of all sizes, across all industries and geographic regions, are recognizing the positive impact of business analytics and business intelligence on competitiveness. One of the firms looking to upgrade its business analytics solution was a global financial services firm whose name has been withheld at its request. To support its business analytics needs, the firm first replaced its underperforming data warehouse platform with Sybase IQ, a purpose-built, column-based analytics database. As a result, the firm was able to achieve productivity, cost reduction, and business performance improvements, such as:

- Improvement in query performance by over 90%
- Decrease of hardware and maintenance costs due to attractive data compression rates of Sybase IQ and the elimination of the need to build and rebuild data aggregates
- Ability to improve information availability and decision modeling for both internal and external stakeholders of the firm

IN THIS BUYER CASE STUDY

This IDC Buyer Case Study highlights the business and technology reasons for deploying a new purpose-built analytics database for data warehousing workloads. IDC had an opportunity to interview a global financial services firm and evaluate its business analytics solution; the results achieved as well as lessons learned by the firm are discussed in this document.

SITUATION OVERVIEW

Organization Overview

The organization is a large, global financial services firm. At the request of the company, its name cannot be disclosed in this document. Therefore IDC will refer to the company as "the firm" throughout this Buyer Case Study.

Challenges and Solution

The firm's goals for its business analytics solution were to support compliance and operational decision-making requirements of several internal groups.

For example, when the federal authorities request historical information on a trade, the firm has to be able to deliver trade information within 24 hours. In some cases, such a request may entail querying seven to eight years of historical data. However, with its existing database, when more than two to three months of trading data stored in the data warehouse was accessed, the query performance degraded to unacceptable levels.

Another issue was with data update frequency to support operational decision making. The existing database lacked the sufficient speed for intraday loading of large amounts of data from the operational database (an in-memory database deployed to support trading) to the analytics data warehouse.

In the past, the firm had used a general-purpose RDBMS for analytics workloads. The firm found that its existing database technology couldn't keep up with query processing requirements of its end users in business analytics and management roles. The massive number of records, wide record sets, and very rapid data update frequency presented a technology challenge to the existing data warehouse generation and management technology and the team assigned to support it.

As a result of these challenges and growing demands from business decision makers for a system that would support their compliance and operational decision-making processes, the firm decided to replace its data warehousing technology.

Technology Evaluation and Acquisition

There was a need for rapid acquisition and deployment of a new analytics database. The firm evaluated products from several competing data warehousing technology vendors. In addition to ensuring that the new data warehousing technology was able to address the data-loading and query performance requirements, the firm's IT group also prepared a cost justification case for the first project.

When the decision was made to replace the existing database system used for business analytics, the firm turned to Sybase and the software company's Sybase IQ product — a purpose-built column-based analytics database.

Although the original deployment was viewed as a "one-off project," the success of the initial project led to eventual global site license contract with Sybase in 2009. Importantly, during the original project, the cost savings on storage hardware alone — due to Sybase IQ compression rates — offset the spend on software license.

The Deployment

The firm invested in the Sybase IQ analytics database and deployed it first in the United States, later expanding the deployment to Asia and Europe. Today, the firm has nine Sybase IQ servers supporting about 400 front- and back-office applications with analytics workloads.

The firm uses Sybase IQ to support live reporting from operational systems as well as query and analysis workloads for applications running on the data warehouse. The analytics functionality provided by Sybase IQ to support various trading processes was first applied to equities trading and has since been applied to such systems as the swaps trading, which in the middle of 2009 was also being migrated from a different database system to Sybase IQ using Sybase Replication Server software for real-time loading.

In addition to several application-specific data marts, the firm's central enterprise data warehouse is now on Sybase IQ and includes all the reference data for the firm. This data warehouse, which includes about 2TB of data, has many different source systems, including the mainframe, which in turn receives certain data feeds after they have been analyzed from the data warehouse.

In the current architecture, data from the in-memory operational database is loaded into Sybase IQ using Sybase's data-loading tools. Today, it's a batch load. However, the firm is changing the loading process to include intraday microbatches. The largest Sybase IQ server has about 8TB of compressed data, which the firm's DBA estimated would've exploded to about 25TB in a general-purpose, row-based relational database. Data volumes in the firm's data warehouse and marts have, on average, doubled every year since 2005. There are multiple packaged and internally developed tools, including SAP BusinessObjects, used by end users to access data in Sybase IQ.

The deployment process was partially assisted by DSSI, a professional services partner of Sybase. One of the drivers for the firm's initial acquisition of Sybase IQ was the opportunity to deploy the database without any additional work on the application development side. The firm was, in fact, able to move analytics workloads off some of its older versions of Sybase ASE database to Sybase IQ without changes to the corresponding applications. However, some reworking was required for migrating other, non-Sybase, databases to Sybase IQ. In these cases, even though to an application a columnar database still looks like a relational database, some minor changes to the data model were implemented to accommodate the desired flattened data structure and elimination of aggregates.

At the firm, Sybase IQ does not have a specialized support team. Instead, shared DBAs have the responsibility for managing Sybase IQ and other databases. The firm cross-trained two DBAs on Sybase IQ and has found the system simple to maintain.

Results

According to the firm's staff interviewed by IDC, the queries that previously took 6–7 hours on only a subset of the data now run in 40 minutes on the whole data set. Although 40 minutes may still seem a long time, it is acceptable from a business process point of view, and the firm emphasized the complexity of the queries. In other words, the price-performance trade-off is acceptable to both business end users and IT staff — arguably the most important metric of success of any technology use case. Most importantly, the firm is now able to meet the required delivery window (less than 24 hours in some cases) for reporting on trade information over a multiple year time span — as required for federal reporting and compliance.

With other, less complex queries, the firm has been able to reduce queries that previously ran 4–5 hours to 4 minutes. In addition to query performance improvements, the firm has achieved other benefits from the deployment of Sybase IQ. The DBAs spend considerably less time on tuning and performance management of the database. This task required significant time investment in the past due to a large number of ad hoc queries launched by business users. Previously, there were times when a DBA might have missed an index, which resulted in one-time performance issues. Human mistakes like that are no longer a factor with Sybase IQ, which does not require the type of indexing tasks that were previously used with the row-based general-purpose RDBMS.

In addition, the firm has seen considerable improvements to data-loading speeds and reduced platform costs — partly due to the high data compression rates for Sybase IQ and partly due to the reuse of machines that already existed at the firm. For example, in the old system, trade history data was spread across three servers. With Sybase IQ, the data is on a single server with an additional server used for disaster recovery.

ESSENTIAL GUIDANCE

Technology users and purchasers looking for analytics database technology for data warehousing should consider the following lessons learned from this Buyer Case Study:

- ☒ The workload requirements for query, analysis, and reporting are different from those of operational systems. IDC research shows that over the past seven years purpose-built data warehouse technology has increased its share of the overall data warehousing market.
- ☒ Columnar database structure is becoming an increasingly popular option for organizations looking for purpose-built data warehouse technology. IDC research shows strong gains not only in awareness of columnar databases over the past two to three years but also in technology purposes and successful deployments.
- ☒ When evaluating technology for purchase and assessing the success of deployments, it is important to assess not only the technical performance characteristic but also the ongoing total cost of ownership of the technology, support provided by the vendor, and vendor's financial stability.

LEARN MORE

Related Research

- ☒ *State of the Business Analytics Market: Survey Shows Positive Buyer Sentiment Going into 2010* (IDC #221277, December 2009)
- ☒ *Worldwide Business Analytics Software as a Service 2009–2013 Forecast* (IDC #221320, December 2009)

- ☒ *Innovative Information Access Companies Under \$100 Million to Watch* (IDC #220450, October 2009)
- ☒ *Worldwide and U.S. Business Analytics Software 2008–2013 Forecast by Vertical Market and Company Size* (IDC #219834, August 2009)
- ☒ *Worldwide Business Analytics Software 2009–2013 Forecast and 2008 Vendor Shares* (IDC #219383, August 2009)
- ☒ *Decision Management Survey Shows Need for Better Collaborative Decision-Making and Knowledge Capture Support* (IDC #218277, May 2009)
- ☒ *Decision Management: A Strategy for Organizationwide Decision Support and Automation* (IDC #218353, May 2009)

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights.

Copyright 2010 IDC. Reproduction is forbidden unless authorized. All rights reserved.