

What Works

IN ENTERPRISE BUSINESS INTELLIGENCE

**POWERFUL CASE STUDIES, LESSONS LEARNED,
AND Q&A FOCUSING ON:**

Enterprise Business Intelligence

Search BI

Open Source BI

Data Governance

Performance Management

Predictive Analytics

FEATURE

The Myth of Self-Service BI

Wayne W. Eckerson, TDWI Research

An exploration into the myths and realities
of managing your own data warehouse.

PAGE 2

TDWI RESEARCH

Challenges to Operational BI

PAGE 36

**Issues that Influence BI Solutions
for SAP**

PAGE 48

Letter from the Editorial Director

TDWI is proud to offer you our new, topically focused *What Works in Enterprise Business Intelligence*. This collection of customer success stories and expert perspectives provides a resource for better understanding the tools, technologies, and methods that are central to enterprise business intelligence. To help you make your way through the many powerful case studies and “lessons from the experts” articles, we have arranged them into specific categories: enterprise business intelligence, search BI, open source BI, data governance, performance management, and predictive analytics.

What's inside?

CASE STUDIES

The goal of *What Works* is to provide snapshots of the most innovative business intelligence and data warehousing implementations in the industry today. The case studies included in this volume demonstrate the power of enterprise business intelligence technologies and solutions for industries ranging from banking to pharmaceuticals to online ticket sales.

LESSONS FROM THE EXPERTS

What Works also includes articles from leading experts in the services, software, and hardware vendor communities. These lessons feature perspectives about enterprise business intelligence best practices and trends.

Q&A WITH THE EXPERTS

Our Q&A section provides answers from these same experts to the enterprise business intelligence questions they hear most often, complemented by insight from an independent consultant.

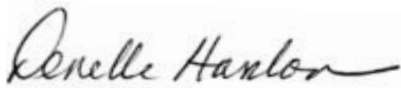
FEATURE ARTICLE

Our feature article comes from Wayne Eckerson of TDWI Research. In “The Myth of Self-Service BI,” he explores the myths and realities of managing your own data warehouse.

TDWI RESEARCH

There's more from TDWI Research. *What Works* includes excerpts from TDWI's recent best practices reports: *Best Practices in Operational BI*, from Wayne Eckerson, and *Business Intelligence Solutions for SAP*, TDWI's latest report from Philip Russom.

We hope you enjoy this collection of case studies, best practices, and expert insight that are all focused on enterprise business intelligence. We look forward to your comments; if there is anything we can do to make this publication more valuable to you, please let us know. Please join me in thanking the companies that have shared their stories and successes, their technology insights, and the lessons they have learned.



Denelle Hanlon

Editorial Director, *What Works in Enterprise Business Intelligence*
TDWI
dhanlon@tdwi.org

What Works
IN ENTERPRISE BUSINESS INTELLIGENCE

tdwi

www.tdwi.org

General Manager	Rich Zbylut
Director of Research	Wayne Eckerson
Director of Marketing	Michelle Johnson
Art Director	Deirdre Hoffman
Editorial Director	Denelle Hanlon
Production Manager	Jennifer Agee
Production Editor	Susan Stoddard
Graphic Designer	Bill Grimmer

1105 MEDIA

President & CEO	Neal Vitale
CFO	Richard Vitale
Executive Vice President Resources	Michael J. Valenti
Managing Director	Dick Blouin
Managing Director	Ellen Romanow
VP, Financial Planning & Analysis	William H. Burgin
VP, Finance & Administration	Christopher M. Coates
VP, Audience Marketing & Web Operations	Abraham Langer
VP, Information Technology	Erik Lindgren
VP, Print & Online Production	Mary Ann Paniccia
Chairman of the Board	Jeffrey S. Klein

Reaching the staff: Editors can be reached via e-mail, telephone, or mail. To e-mail a member of the staff, use the person's first initial followed by their full last name @1105media.com.

Renton office: 425.277.9126 Fax: 425.687.2842
1201 Monster Rd SW, Ste. 250, Renton, WA 98057

Corporate office: 818.834.1520; Fax 818.734.1528
9121 Oakdale Ave. Ste. 101, Chatsworth, CA 91311

Sponsorship Opportunities: Denelle Hanlon, 425.277.9130, dhanlon@tdwi.org.

Reprints: For editorial and advertising reprints of 100 copies or more, and digital (Web-based) reprints, contact PARS International at 1105reprints@parsintl.com or 212.221.9595, or visit www.magreprints.com/quickquote.asp

List rentals: This publication's subscriber list, as well as other lists from 1105 Media, Inc., is available for rental. For more information, please contact our list manager, Merit Direct. Phone: 914.368.1000; e-mail: 1105media@meritdirect.com; Web: www.meritdirect.com

© 2007 by 1105 Media, Inc. All rights reserved. Reproductions in whole or part prohibited except by written permission. Mail requests to "Permissions Editor," c/o *What Works In EBI*, 1201 Monster Road SW, Ste. 250, Renton, WA 98057. The information in this magazine has not undergone any formal testing by 1105 Media, Inc., and is distributed without any warranty expressed or implied. Implementation or use of any information contained herein is the reader's sole responsibility. While the information has been reviewed for accuracy, there is no guarantee that the same or similar results may be achieved in all environments. Technical inaccuracies may result from printing errors, new developments in the industry, and/or changes or enhancements to either hardware or software components. Printed in the USA.

TDWI is a trademark of 1105 Media, Inc. Other product and company names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Table of Contents



FEATURE

2 The Myth of Self-Service BI

Wayne W. Eckerson, Director, TDWI Research

An exploration into the myths and realities of managing your own data warehouse.

ENTERPRISE BUSINESS INTELLIGENCE CASE STUDIES AND LESSONS FROM THE EXPERTS

- 5 Enterprise BI Categories Defined

ENTERPRISE BUSINESS INTELLIGENCE

- 6 StubHub Scales to Explosive Growth
- 8 Transitioning from Successful BI Projects to Enterprise BI
- 9 U.S. Defense Commissary Agency Enriches EDW with Integration of XML Data from Point-of-Sale Systems
- 11 Redefining EDW: Data Integration Requirements for Enterprise Data Warehousing
- 12 Alpharma Pharmaceuticals LLC: Improving Pharmaceutical Sales and Marketing Effectiveness
- 13 Solving the CIO Dilemma: Breaking the Reporting Backlog with a New Formula for Self-Service
- 14 Pitney Bowes Improves Total Cost of Ownership with Enterprise Business Intelligence
- 15 Today's Enterprise Business Intelligence
- 16 Information as a Service Delivers Consistent Information for Business Intelligence
- 17 Dynamic Warehousing: The Next Evolution of Data Warehousing
- 25 Direct Marketing Company Gets It Done Faster with the Right Tool
- 26 When Enterprise ETL Slows Down ... Speed It Up
- 28 Large-Scale Retailer Gains Easy Access to Historical Data for Big-Picture Analysis
- 29 Creating a True Enterprise Data Warehouse with Grid-Enabled DW Appliances
- 30 Alvion Scales Data Analytics and Distribution
- 31 The Analytics Server: Powering On-Demand Analytics from the Ground Up
- 32 BI at Tata Teleservices Ltd.
- 33 CPM Implementation Challenges and Lessons Learned
- 34 Hospital Finds Cure for Information Overload Syndrome

SEARCH BI

- 40 University of Utah, University Health Care: BI Tips that Drive User Adoption and Increase ROI

OPEN SOURCE BI

- 42 Open Source BI: Does It Fit into Your BI Environment?
- 43 Enterprise Business Intelligence: Making the Case and Keeping the Promise
- 44 Why Open Source Was the BI Suite of Choice for ZipRealty
- 45 Pleasant Surprises: Commercial Open Source BI

DATA GOVERNANCE

- 46 Mortgage Company Drives Compliance with Data Governance
- 47 Getting Started with Data Governance

PERFORMANCE MANAGEMENT

- 52 Kelley Blue Book Improves Data Analysis and Reporting with Business Intelligence Solution
- 53 The BI Tipping Point
- 54 Operational Performance Reporting Applications Help MFS Improve Regulatory Compliance, Enhance Productivity, and Expand Global Operations
- 55 Spreadsheets in the Next Generation
- 56 Rohm and Haas: One Company, One Measure
- 57 Corporate Performance Management: Departmental or Corporate?

PREDICTIVE ANALYTICS

- 58 Innovative Analytics for Improved Performance and Consistent Decision Processes
- 59 You Can't Offer a Solution If You Don't Know the Problem
- 60 Cablecom Reduces Churn with the Help of Predictive Analytics
- 61 Enhance Enterprise BI to Drive Business Growth with Predictive Analytics
- 62 Top Retail Bank More Than Doubles Campaign Profitability Using Latest Uplift Modeling Techniques

TDWI RESEARCH: BEST PRACTICES REPORTS

- 36 Challenges to Operational BI
- 48 Issues that Influence BI Solutions for SAP

RESOURCES FOR BI & DW PROFESSIONALS

- 18 Q&A with the Experts
- 63 Best Practices Awards 2007
- 69 Solution Providers
- 76 About TDWI
- 77 TDWI Partner Members



The Myth of Self-Service BI

By Wayne W. Eckerson, Director, TDWI Research

Liberating IT?

Remember D-Day? That was the day in June 1944 when Allied troops stormed the beaches of Normandy to liberate Europe from the Nazi occupation. But there is another D-Day, one that is less well known but yet significant in its own right. That was the day in 1990 that data warehousing came to life. This is the day that information technology (IT) managers celebrate with gratitude and fondness. (Or at least they should!)

At the time, data warehouses promised to liberate IT from the drudgery of creating custom reports. Previously, the IT department was swamped. It couldn't keep pace with the demand for custom development and was slowly drowning in a backlog of requests, a good portion of which were for custom versions of corporate reports. (Alas, things haven't changed much for some organizations where IT groups are still bogged down with requests for custom reports.)

Admittedly, IT was cautious at first about creating redundant data stores—a data warehousing prerequisite and a long-time IT taboo emanating from the days of expensive disk storage. Once IT realized it could offload a huge portion of its work to end users, however, it began to evangelize the benefits of data warehousing to the business. Soon enough, data warehousing and its successor, business intelligence, became a clarion call for IT and a booming career opportunity as well.

The heart and soul of data warehousing, at least to an IT professional, is the notion of self-service reporting or self-service BI. Here,

business users create their own custom reports using end user-oriented query and reporting tools running against a data warehouse. IT steps aside as an intermediary between users and the data and gives users what they've demanded for years: complete and unfettered access to data without IT interference. All that IT needs to do is set up the data warehouse and provide query and reporting tools. It can then focus on more value-added activities, such as developing new applications.

Data warehousing created a win-win situation in which business users gained direct access to data and IT managers eliminated the need to create custom reports—or at least that was how it was supposed to work.

Downsides of Self-Service BI

Unfortunately, as many organizations have discovered the hard way, self-service BI is a myth and doesn't translate well to reality. Although the concept is valid, implementation is misguided. The result is reporting chaos.

One cause of the chaos is that organizations, keen to empower users with data warehouses and BI tools, go overboard. They give users too much responsibility for generating the information and reports they need to do their jobs. In reality, most users don't want this responsibility, and it's not part of their job descriptions. It takes too much time, and users often make mistakes and get frustrated. If they do get training, users usually forget how to use the tool by the time they need to create a report. Consequently, they either stop

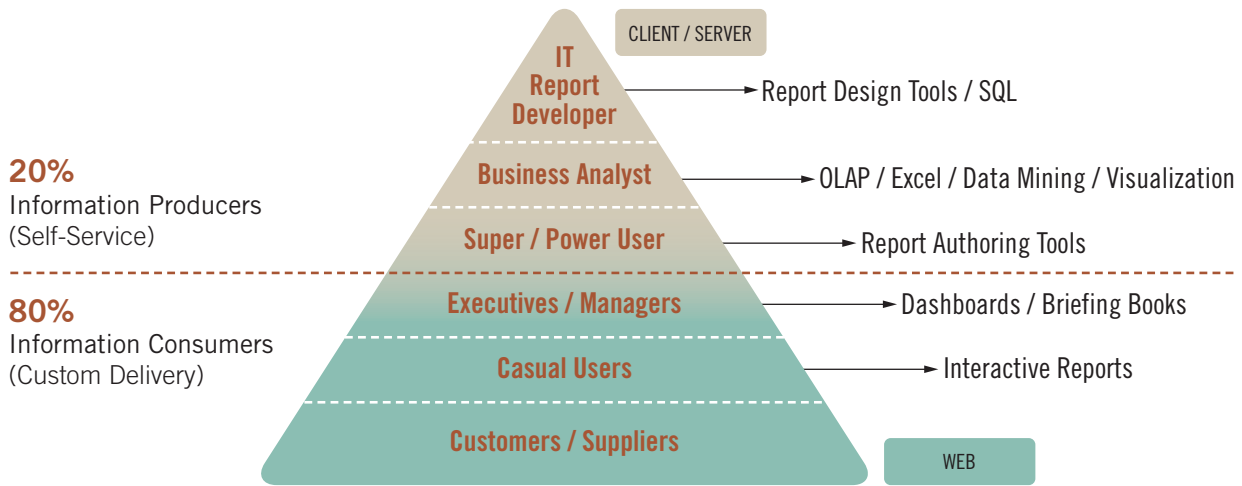


Figure 1. This diagram is one way to map user segments to BI tools.

using the tool or call IT to create the report for them. The organization then finds itself exactly where it was before it spent hundreds of thousands, if not millions, of dollars on its DW/BI program.

For example, the human resources department in one large organization I worked with discovered that it had 26,000 different reports serving 450 active users out of a potential of 3,500. Most of the reports were variations on a couple of themes, and most hadn't been used in months or years, but they were still sucking up disk space and cluttering report folders. The majority of the 3,500 employees who could benefit from the BI infrastructure found the tools difficult to use or couldn't find the right report in the directory.

There is another, more deleterious, downside to reporting chaos than lack of usage and BI tool shelfware. It occurs when the small percentage of users who do employ the new tools don't define key metrics, accounts, and terms in a consistent fashion. This makes it impossible for executives to get a harmonious view of business activity, and they may unwittingly make critical decisions based on inaccurate information.

Another problem is that overambitious business users may submit costly, long-running queries, which bog down query performance for all other users. When result sets take a painfully long time to return, business users accustomed to split-second response times for Google queries stop using the BI tools altogether. What's ironic is that most of these network-clogging, runaway queries are usually unnecessary; users typically need only a fraction of the data that they include in their queries.

The Pendulum Swings

When organizations first deploy data warehouses and BI tools, they embrace BI self-service with an exuberance bordering on delirium. Once afflicted with report chaos, however, they recognize the limi-

tations of self-service BI. Rather than ricochet back to the days of canned reporting, most organizations seek to balance self-service BI with tailored delivery of information to casual users.

Mapping Users to Tools. In order to achieve this balance, organizations realize they need to better understand the information requirements of their users and provide them with tools and reports that are optimally suited to their needs. In essence, they need to investigate what users want to achieve from a business perspective before they ask what kind of information or tool they need. Are they trying to reduce supply chain costs while increasing quality by implementing a supplier rating program? Are they trying to maximize call center productivity by tracking call times and sales volumes? Knowing their ultimate business goals focuses the discussion of BI tools on the right level.

After such investigation, most organizations discover that 80 percent of their users consume reports or information views created by the remaining 20 percent. The 80 percent of "information consumers"—or casual users, as I call them—consist of executives, managers, front-line workers, customers, and suppliers. The remaining 20 percent are "information producers," comprising business analysts, power users, and IT developers. (See Figure 1.)

There are numerous subcategories within each of these groups, varying from one organization to the next depending on the makeup of their user populations. Once organizations segment their user base, they can map each segment to one or more BI tools that meet their information requirements. This type of evaluation helps maximize BI usage and minimize BI shelfware.

Be aware, however, that individuals often play multiple roles. A power user who creates custom views for colleagues in the marketing department may be a casual user when viewing human

resources information or corporate financial data. Typically, you need to arm users with a variety of tools or logins to access different reports and sets of BI functionality.

Tools Selection Process. As simple as it sounds, many organizations skip the step of mapping users to tools. What usually happens is that information producers hijack the tool selection process and pick BI products that are suited to them but not to the majority of users. To avoid this problem, populate the BI tools selection team with a cross-section of users from multiple areas and listen carefully to their feedback. Remember, there is no tool that meets the needs of all users, so you'll need to invest in a suite of tools. Fortunately, the BI market has matured to the point where leading BI vendors now offer a comprehensive suite of integrated BI tools built largely on a single, integrated architecture (or so they claim).

IT-Driven Standard Reports

Once you catalog users and assign tools, you need to create a limited set of standard, interactive reports. These reports are geared to information consumers, the 80 percent of your employees who require periodic access to information. The remaining 20 percent of business analysts, power users, and developers can continue to use their ad hoc tools to perform deep analysis and exploration of data.

A well-designed interactive report can replace dozens, if not hundreds, of existing reports. Each report typically focuses on a specific domain and contains about 20 dimensions and 12 metrics. These standard reports are broad without being overwhelming because they present a simple, intuitive interface that makes it easy for users to apply filters to navigate and analyze the predefined data set. Once users navigate to a particularly useful view of information, they can subscribe to that view so it's available with fresh data the next time they open their reporting tool.

These standard interactive reports usually take the form of a dashboard, scorecard, or parameterized report that enables users to change a predefined view by selecting filters from a pick list or other graphical mechanism. As a result, these reports often give users the impression that they performing ad hoc queries against all the data when, in reality, their access is circumscribed by a predefined set of metrics, attributes, and filters. The BI team needs to monitor usage and examine which elements are being used and which aren't and perform periodic selective pruning to keep the reports relevant and fresh.

Report Governance. Guess who needs to create these standard interactive reports? Yes, it's the IT department or its successor in this space, the BI team. But unlike in the days of old, the BI team does not serve as a custom report development shop. Its role is to examine global requirements, create a library of standard interactive reports, and establish a governance process for adding, deleting, and modifying the reports. A governance committee,

comprising a representative group of business users and IT professionals from a single domain or multiple domains (depending on the organizational structure and DW architecture), reviews requests for new reports or modifications. The committee determines whether any of the existing reports can meet the need or if an extension or new report is required.

Effecting Change. Of course, you don't want to kill requests by committee, so before you implement the governance process, you need to work hard and fast to develop the standard interactive reports so that 80 percent of the users can get 80 percent of the information they need from them. Then, you need to shut down the older, ad hoc reporting systems. This can be politically sensitive, so you must be cautious.

Some organizations are fortunate enough to have a CEO or CFO issue a mandate that the new standard reports will be the basis for all future decisions. If you're really lucky, they'll enforce this mandate by using the reports themselves and scolding subordinates who present information created in ad hoc, nonstandard ways. Less fortunate organizations need to tread carefully, persuading groups and managers that the new reports contain more accurate and comprehensive data than their existing reporting solutions. While many groups will convert, others will hold out for a long time, even if the BI team refuses to support the system with upgrades, platform support, or help desk services.

Power Users. Remember, exceptions are the rule. Give information producers unfettered access to data and allow them to acquire the BI tools they want—because they'll acquire what they want whether you allow them to or not. In other words, self-service BI makes sense for power users, but not for everyone else. Make sure the power users create reports only for themselves, not the general audience. If they come up with a view of data that can help executives and managers make better decisions, then that view—along with data set and information tools required to produce it—needs to be adopted by the report governance committee and moved into a global report library for all to access.

Conclusion

In adolescence, we swing wildly from one passion to another, carried away by the newness and potential of our object of interest. In adulthood, we learn that there is a price to pay for our excesses, and we seek to find balance amidst the extremes.

In the world of BI, organizations often are carried away with the notion of self-service empowerment and fall prey to reporting chaos. To remedy this situation, organizations need to achieve balance between self-service BI for a select few and tailored delivery of information for the majority. This will give organizations a stronger foundation upon which to leverage information for insights and business gain.

Enterprise Business Intelligence Categories Defined

To help you make your way through the many powerful case studies and “lessons from the experts” articles in *What Works in Enterprise Business Intelligence*, we have arranged them into specific categories: Enterprise Business Intelligence, Search BI, Open Source BI, Data Governance, Performance Management, and Predictive Analytics. What do these terms mean, and how do they apply to your organization?

ENTERPRISE BUSINESS INTELLIGENCE

pages 6–17, 25–35

Enterprise business intelligence involves deploying query, reporting, and analysis capabilities to all employees who can benefit from them as well as to customers and suppliers.

SEARCH BI

pages 40–41

BI search applies keyword search capabilities to structured data and BI reports, making it possible for users to find reports, generate queries, and find related analyses with a simplified user interface.

OPEN SOURCE BI

pages 42–45

Open source BI consists of query, reporting, and analysis tools built on an open source foundation; in other words, they are free to download and use.

DATA GOVERNANCE

pages 46–47

Data governance describes the people, processes, and technology needed to ensure consistent definitions and rules for commonly used data elements.

PERFORMANCE MANAGEMENT

pages 52–57

Performance management uses business intelligence tools to monitor and manage progress toward achieving strategic objectives and goals represented in budgets and other plans.

PREDICTIVE ANALYTICS

pages 58–62

Predictive analytics uses analytical models based on statistical and machine-learning algorithms to describe or predict patterns within large volumes of data to help companies better anticipate events and behavior.

CASE STUDY

StubHub Scales to Explosive Growth

Commentary by Rob Singer
Director of Business Intelligence, StubHub

Summary

StubHub, a rapidly growing online marketplace, had limited visibility into current data, customers, and business drivers, which threatened the long-term survival of the business. The root of these problems was an inadequate technology infrastructure, but StubHub did not have the resources to make an expensive overhaul to address its IT needs.

Faced with these challenges, StubHub invested in a business intelligence (BI) and data warehouse (DW) initiative, using limited resources. The StubHub BI/DW initiative emphasized investing in specific technologies to empower end users and to free up existing staff resources, reusing existing technology investments, and employing open source technologies where possible.

With a total budget of \$750,000 (42 percent staff, 42 percent new licenses, and 16 percent training), and ongoing maintenance of less than \$100,000, StubHub realized \$22.4 million in savings in 2006—directly attributed to returns from improved processes and reduced staff. Additional savings sprang from enhanced reporting capabilities, giving StubHub executives and end users invaluable insight into current data, customer behavior, and key business drivers.

Business Impact

StubHub is an online marketplace where fans buy and sell tickets for live events. Since its founding in 2000, it experienced triple-digit year-over-year growth, becoming a market leader and putting strains on its existing information management infrastructure.

Through mid-2005, StubHub's business intelligence function relied entirely upon a Microsoft Access database—the source for about 40 Microsoft Excel pivot tables residing in a shared network file directory. The company had minimal insight into customer behavior

and no visibility into the dynamics of its Web site. (It had no Web analytics function.)

Because of the two-gigabyte storage limit for Access databases, StubHub constantly pruned the database to make room for new data, reducing the richness of its data in the process. By June 2005, StubHub devoted an average of 16 hours a week to loading the data warehouse, refreshing the pivot tables, and notifying employees via e-mail when refreshed reports were ready. With delayed and static reporting, StubHub had minimal understanding of its customers, its competitors, or its key business drivers, and decision making was based on gut feel instead of being data driven. Clearly, it had to make significant changes—fast.

StubHub assembled a steering committee and a business intelligence team to assess requirements, evaluate options, and ultimately build a scalable BI/DW infrastructure that could support StubHub's explosive growth and contribute to business value. Composed of executives, BI professionals, and end users from different disciplines at StubHub, the steering committee established key goals, which included:

1. Creating a scalable and extensible IT infrastructure to enable dynamic reporting, provide access to rich current and historical data, and reduce data latency.
2. Providing actionable customer insights, to include defining key customer segments, quantifying customer lifetime value, and creating customer behavioral profiles.
3. Providing tools to maximize site conversion rates, marketing effectiveness, and customer loyalty.
4. Identifying key business drivers and defining key company metrics.

The team also needed to accomplish these goals using creative approaches and leveraging existing resources wherever possible. It developed a detailed roadmap and project plan with clear, measurable milestones to

meet the objectives and stretch investment in the initiative. Combining open source technology with best-of-breed software tools from Business Objects, Oracle, Omniture, and SAS, StubHub used a phased approach to deploy an extensible BI/DW infrastructure, leading to the following measurable results:

1. By moving from Access and Excel to Oracle and Business Objects, StubHub migrated to a best practices model for data warehousing and report distribution that reduced data latency and switched from a manual, weekly data refresh to an automated nightly data refresh. Migrating to the automated environment, StubHub cut costs through staff reduction, resulting in \$150,000 in annual savings.
2. Creating a marketing allocation report (called NORAD) that calculates key marketing performance metrics, such as cost per order (CPO). NORAD merges transactional data with several external data feeds to associate transactions with the marketing channels that drove them. The report enables StubHub to understand the dynamics of its numerous marketing channels and enables the marketing team to optimize its marketing spend on a daily basis. In 2006, NORAD reduced CPO by 33 percent—a savings of \$15.6 million.
3. The new infrastructure permitted the development of a daily report that tracks each StubHub seller's dropped-order rate. These dropped orders—where a seller is unable to provide the purchased tickets—negatively affect customer satisfaction and increase operational costs by creating significant customer service activity. Reducing StubHub's overall dropped-order rate was a key strategic priority for 2006. Using the daily report, which includes alerts to the customer service team, enabled StubHub to work proactively with sellers to improve their performance, and resulted in a decrease in its dropped-order rate by 31 percent—saving \$600,000 in costs. Using the report to monitor dropped orders and

- alerting the customer service team to potential problem areas has enabled rapid problem resolution and, ultimately, a significant increase in customer satisfaction
4. StubHub continually monitors Web site activity using Omniture. With its real-time reporting capabilities, Omniture enables StubHub to make ongoing changes to its Web site, which increased the site-conversion rate of 8.5 percent in 2006—resulting in \$5.4 million in additional revenue.
 5. StubHub moved from an exclusively “push-based” reporting environment, where BI analysts developed all the reports, to a “push-pull” environment, where end users are more empowered. End users, including marketing, business development, finance, and customer service, now access more than 300 customizable standard reports. Further, StubHub established “sandboxes”—highly interactive subject-area reports—that business owners use to tailor reports to fit their own business needs. Because StubHub’s business users can answer their own questions, StubHub’s BI analysts could focus on answering strategic business questions. In 2006, they were able to quantify customer lifetime value, identify key customer segments, and develop high-impact reports like NORAD and the dropped-order report.
 6. Executives and business users now have visibility into key company drivers. They are engaged with the data and constantly respond to customer and competitive activity. With greater understanding of customer buying patterns, StubHub is able to anticipate demand and deploy changes to its marketplace in near real time to accommodate market-changing events.

Maturity

StubHub achieved all objectives of its BI initiative: It successfully replaced its failing Access database with a stable, scalable, and extensible data warehouse environment

that enables dynamic reporting and provides access to current and historical performance data. StubHub now has a 360-degree view of the customer, company-wide dissemination of critical company data, and the quantification and monitoring of key company metrics.

The new database and reporting environment was developed and deployed in eight months (six months of coding, two months of parallel testing). StubHub completed phases one and two of its enterprise reporting goals, migrating key Excel pivot tables to Business Objects reports, delivering sandboxes for scalability, and producing value-added reports, such as the marketing allocation report. Another key achievement was StubHub’s standardization of terms, including the definition of all metrics.

StubHub now continuously maintains a rolling two-year roadmap to build on the platform and continue driving business value through actionable intelligence. The next phases include the deployment of Business Objects Dashboard Manager to monitor StubHub’s key performance indicators. The dashboards will give executives focused, efficient delivery of mission-critical data. The long-term roadmap includes enhancing operational reporting capabilities; StubHub is building the foundation for providing business activity monitoring, balanced scorecards, and predictive analytics.

Relevance

StubHub leveraged industry best practices that can be adopted easily by similar organizations. These practices include:

- Obtaining executive buy-in upfront by securing an executive sponsor and implementing a proof of concept
- Assembling a cross-functional team to ensure full organizational buy-in and to provide an appropriate enterprisewide perspective that aligns deliverables with business goals
- Developing a detailed project plan that includes a phased approach and clear

milestones with minimal external dependencies (to maximize success rates)

- Using an ROI-driven approach for project prioritization
- Running in parallel environments until core functionality and reports are fully migrated and fully tested
- Adhering to industry-standard best practices, including using a traditional star-schema data model, a software development lifecycle methodology, and data quality techniques
- Delivering “quick wins” with key reports to demonstrate and reinforce value to the organization
- Adhering to a 12-month rollout and three-year strategic plan for implementation

Innovation

StubHub, like many midmarket companies, is focused on keeping costs as low as possible, and the team needed to find innovative ways to achieve an enterprise-class implementation on a limited budget. The team used novel approaches to that goal, including:

- Using open-source, on-demand, and free/bundled software for its commodity technology needs so that it could focus its limited software expenditures on value-added and best-of-breed tools
- Extensive use of sandboxes, minimizing BI headcount and enabling the BI team to focus on delivering value-added projects
- Repurposing hardware to reduce capital expenditures further
- Collaborating with partner groups to create new reports and functionality to ensure that the BI team does not build anything that goes to waste. All participants understand the value of what is built and are taught how to use it, resulting in organizational buy-in and a high adoption rate. ●

StubHub’s innovative solution received a 2007 Best Practices Award from TDWI (see page 63).

LESSON FROM THE EXPERTS

Transitioning from Successful BI Projects to Enterprise BI

By Timo Elliott
Senior Director, Strategic Marketing,
Business Objects

It's hard to think of an IT initiative that has as much potential to add value to today's information-rich organizations as enterprise business intelligence (BI). Although individual BI projects typically are very successful, many organizations stumble on their journey from initial project successes to enterprisewide deployments. According to a recent *Economist* study, fewer than one in 10 executives consider that they receive information when they need it, and fully 56 percent are concerned about making poor choices because of faulty, inaccurate, or incomplete data.

Enterprise BI Success Requires Relearning Some Key Skills

Just as Tiger Woods decided that he needed to relearn his swing in order to take his game to the next level, organizations must reconsider some of the basics of their BI deployments in order to move to enterprise BI. In BI, just as in golf, some of the skills and approaches typically associated with initial individual successes may actually hinder success at the enterprise BI level.

Individual BI projects typically focus on expert users who work with high-value information. This is a great place to start, but it's easy to overdeliver to a small number of vocal, technical users (generally estimated to be 15 percent of the potential user population), and miss the greater value of a broad deployment to less technically savvy users. If experts are the only people using your BI system, the temptation is to:

- Fix problems such as data quality in a patchwork, ad hoc way
- Tolerate poor ease of use and a confusing variety of different systems

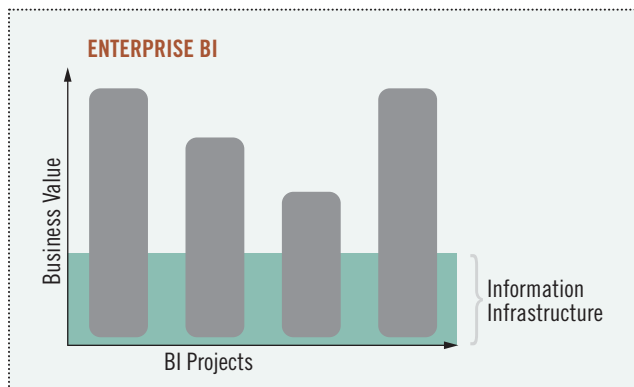


Figure 1. Moving from a series of successful BI projects to enterprisewide BI deployments involves changing some of the habits that made the initial projects successful. In particular, organizations must consciously invest in wider, simpler access for the “other 85%”—the non-expert users.

- Limit the number of self-service users

Moving from a series of successful BI projects to enterprisewide BI deployments involves changing some of the habits that made the initial projects successful. In particular, organizations must consciously invest in wider, simpler access for the “other 85 percent”—the nonexpert users.

Today's organizations should look to the success of Web 2.0-inspired initiatives—collaborative, Web-based communities such as Flickr, MySpace, and Facebook. Rather than focusing on advanced functionality, these organizations start out with simple, useful functionality—making it easy to sign up, participate, and collaborate within the community. Only after these organizations establish a community comfortable with the basics do they implement more sophisticated features and applications.

Implementing an Information Infrastructure

Enterprise BI success requires an approach similar to the easy-to-use Web 2.0 models. After several key successful projects are in place, progressive companies start filling in a broad foundation of information use throughout their organizations with the following elements:

- A systems approach to data quality and monitoring
- Simpler data access for a broader range of users (and greater investment in training)
- An emphasis on collaboration and discussion around data

- Making information actionable by embedding BI directly into operational systems at the point of decision

This information infrastructure approach puts the right framework in place for a culture of information use and fact-based decision making. It fosters a spiral of increasing numbers of employees accessing, analyzing, and sharing data—and improving operational performance.

Supporting this information infrastructure typically requires organizational and technology changes. Some form of BI competency center is essential, tasked with the optimal use of information across the organization as whole. One key task of the BI competency center is to standardize the BI products used in the organization—favoring easy-to-use, Web-based systems that have the openness, breadth, and integration necessary to support the variety of environments, applications, and profiles that exist throughout the organization.

Conclusion

Enterprise BI projects often fail because the majority of business users have little or no connection with the BI system. Organizations should invest in an information infrastructure that provides simple access to basic information for all users—not just the BI experts—before attempting to implement more complex systems. ●

For free white papers on this topic, download “The Benefits of Business Intelligence Standardization” or “Making a Strategic Business Intelligence Choice—A Checklist,” or [click here](#) for more information about Business Objects.

CASE STUDY

U.S. Defense Commissary Agency Enriches EDW with Integration of XML Data from Point-of- Sale Systems

Commentary by Stan Ferguson
Chief, Data Management Division, DeCA

The Business

The Defense Commissary Agency (DeCA) supplies more than 12 million U.S. military customers with groceries, meats, and produce at 263 commissaries around the world. Based in Fort Lee, Virginia, with more than 18,000 employees, DeCA sold \$5.41 billion in goods in 2006 at a 5 percent surcharge over cost, saving customers more than 30 percent over commercial prices. DeCA takes customer satisfaction seriously, as demonstrated by a 4.61 satisfaction rating by customers on a five-point scale.

The Challenge

An enterprise data warehouse (EDW) is central to DeCA's strategic objective of cost-effective customer service. Throughout the 2000s, DeCA has expanded the warehouse that it uses to collect and track data generated by point-of-sale (POS) systems from commissaries in the U.S. and around the world. DeCA uses the Informatica PowerCenter platform to source POS data and channel it to a Teradata-based EDW.

With the EDW, hundreds of DeCA regional, zone, and store managers have improved customer service and product offerings by analyzing sales by time and geography, promotion effectiveness, product affinity, checkout time, and other dimensions. Importantly, the EDW also supplies data to DeCA's computer-assisted ordering system to help ensure that products customers want are in stock.

In late 2006, DeCA began rolling out the first phase of a two-year, \$200 million upgrade to POS systems at its 263 commissaries. Aging NCR-based POS terminals and supporting software would be replaced by an IBM hardware and software package that DeCA called CARTS (Commissary Advanced Resale Transaction System). The CARTS system offers touch-screen and self-service checkout, as well as greater efficiency and reliability.

CARTS also introduced a challenge for the data integration processes that drove POS information to the EDW. DeCA would need to adapt its warehouse integration architecture to accommodate XML/XSD data from the new POS systems, while simultaneously integrating flat file data from legacy POS systems.

In the past, PowerCenter had sourced and delivered uncomplicated flat files generated by the POS terminals to the EDW. CARTS would use complex XML files based on an industry standard XSD (XML Schema Definition) as well as complex flat files derived from the transaction data files. Naturally, this introduced new complications and challenges on how the data would be integrated into the warehouse.

In addition, DeCA envisioned moving to a more real-time system by which POS data loads to the warehouse would occur incrementally throughout the day, rather than once a night. DeCA needed a data integration solution that could:



- Be rapidly adapted to seamlessly integrate XML/XSD and complex flat file data
- Simultaneously integrate disparate data from both legacy and new POS systems
- Support greater load frequency and reduce commissary restock time

The Solution

DeCA first selected PowerCenter in 2000 for its near-universal access to data, codeless development environment, and flexibility for future growth. By implementing PowerCenter for data warehousing, DeCA IT personnel and data integration experts from Claraview, a Virginia-based consultancy specializing in business intelligence, were knowledgeable in its capabilities. After due diligence, DeCA decided that PowerCenter was ideally suited to help it transition to the XML/XSD and complex flat file processing required by CARTS.

"Initially, we had concerns over the data integration challenges that arose as we looked at phasing in our new POS system," said Stan Ferguson, chief, data management division, DeCA. "But Informatica PowerCenter has let us quickly adapt our infrastructure to handle complex file data integration with our EDW with no disruption to our core business processes."

PowerCenter's support for complex files such as XML/XSD and complex flat files meant that DeCA could readily access and integrate

continued on page 10

CASE STUDY

continued from page 9

the new CARTS data, and would not need to identify, test, and purchase a second integration platform, nor would it need to look at custom-coding alternatives that would have been exponentially more costly, time-consuming, and prone to error. PowerCenter’s platform-neutral architecture has enabled development at a higher level of abstraction by removing physical ties to data and allowed DeCA to adapt to different source types without affecting business logic.

“With Informatica, we had a tool already in place that could read the XML files, and then validate those XML files so we can import data out of sources and targets,” said Carol Griffith, EDW system manager. “Compared to other alternatives, it gives us a quick turnaround. We don’t have to invest in any other software or integration platform, or get resources trained in a new technology.”

A team of four Claraview developers, already skilled in PowerCenter technology, was able to rapidly adapt DeCA’s existing PowerCenter solution to handle XML-based source data, as well as complex flat files. The XML/XSD source systems went live in November 2006.

Once completed in 2008, CARTS will be deployed at each of DeCA’s 263 commissaries and PowerCenter adapted to integrate the XML/XSD and complex flat file data and deliver it to the EDW, which stores two years of POS information, or about 4 TB of data (with 7.5 TB of capacity). A full six years of historical POS data is archived in storage.

PowerCenter’s performance in loading up to 15 million rows of XML/XSD data in 30 minutes is well within DeCA’s acceptable timeframes and achieved without the benefit of performance tuning and optimization.

The Results

SAVE DEVELOPMENT TIME AND REDUCE OVERALL COSTS

DeCA’s rapid adaptation of an existing data integration solution to integrate CARTS POS

data is due in large part to PowerCenter’s capacity to reuse source and target data definition objects across multiple systems. The reusability, along with PowerCenter’s drag-and-drop ease of use, has dramatically reduced the time otherwise required to develop, test, and implement mappings.

“The reusability is very beneficial,” said Griffith. “Once someone builds a mapping, you just save it and can reuse it elsewhere with some minimal tweaks.” DeCA also leverages reusable mappings to create and validate XML files from data in the Teradata warehouse and send that information to CARTS.

“Informatica gives us the ability to run the two heterogeneous POS systems in parallel at the same time by dramatically reducing the effort needed to adapt to the new data definitions. It also allowed us to move to a more real-time model thus improving service to our servicemen and women.”

Stan Ferguson, DeCA

IMPROVE RESTOCKING AND CUSTOMER SATISFACTION WITH RIGHT-TIME DATA INTEGRATION

By updating POS data every 10 to 60 minutes rather than nightly, PowerCenter enhances the business value that DeCA realizes from its EDW by accelerating ordering and restock processes, particularly important for perishable items. The accelerated load frequency can reduce by a day the time required by DeCA’s computer-assisted ordering system to restock commissary shelves with products that customers want and is an important step in DeCA’s mission of maximizing customer satisfaction. The warehouse

and a homegrown business intelligence tool called COMS (Commissary Operating Management System) enable DeCA managers to analyze customer satisfaction, product performance and promotions, seasonal sales fluctuation, and other characteristics.

ENHANCE DATA ACCURACY AND TRUST IN DECISION MAKING

The core mission of DeCA’s EDW team is to be the trusted historical central repository for data analysis. PowerCenter has helped the team build that trust with data quality technologies designed to ensure that warehouse data is accurate and reliable. DeCA takes advantage of PowerCenter error handling capabilities to screen data being loaded into the warehouse and deliver alerts on attempted duplicate data entries and other anomalies. By avoiding the “garbage in, garbage out” phenomenon, DeCA’s warehouse team has given business users confidence in the integrity of POS data.

EXTEND EDW ATOP A PLATFORM-NEUTRAL INTEGRATION ARCHITECTURE

In DeCA’s pipeline is an initiative to extend the EDW’s reach by using PowerCenter to integrate data from eight disparate instances of a 1982-vintage legacy application covering ordering, receiving, and inventory. This would give DeCA an on-demand enterprise view of core business activities, rather than manually building a large workfile of data from the eight instances to answer basic business questions. From a broader perspective, PowerCenter’s platform-neutral architecture plays a key role in enabling DeCA to comply with the Department of Defense’s Business Management Modernization Program, in part with an open platform to replace or complement aging, proprietary legacy systems. Also in the works are Power Center-based metadata services to enable users to easily search warehouse data and the establishment of a data management office. ●

Redefining EDW: Data Integration Requirements for Enterprise Data Warehousing

By Judy Ko

Senior Director, Product Management,
Informatica Corporation

Data warehousing was once a siloed activity dedicated to business intelligence and reporting at a departmental or business unit level. Over the last five years, however, the brisk pace of change coupled with growing regulatory requirements has forced data warehousing into a mission-critical, operational role. Meanwhile, technologies have matured to the point where they can support these enterprise requirements.

Data warehousing has now evolved into a strategic, enterprisewide initiative that supports multiple business applications. An enterprise data warehouse is a common data foundation that provides any and all data for business needs across applications and divisions. Enterprise data warehousing (EDW) is the process of designing, building, and managing an enterprise data warehouse to meet the requirements of consuming applications.

Key Capabilities of Enterprise Data Warehousing

To live up to its potential, enterprise data warehousing must function atop a data integration platform that powers an ongoing lifecycle of data access, discovery, quality, integration, and delivery from virtually any system. Key capabilities that characterize an enterprise data warehouse include:

- **BUSINESS RELEVANT DATA ACROSS THE ENTERPRISE**

EDW must allow organizations to access, discover, and integrate data from virtually any business system, in any format. It should deliver data throughout an enterprise via data services to support multiple

applications in a flexible manner as part of a service-oriented architecture (SOA). EDW must also provide relevant business context for the data, so that the consumers of the data are able to interpret it in a meaningful way.

- **TRUSTED, CERTIFIABLY ACCURATE DATA**
EDW must equip organizations to manage data quality in a metric-driven, programmatic fashion. Maintaining data integrity and security across extended teams throughout the data integration lifecycle is required to meet regulatory compliance and governance objectives. To ensure confidence in the validity of enterprise data, information as well as data flows and relationships must be auditable and traceable.

To live up to its potential, enterprise data warehousing must function atop a data integration platform that powers an ongoing lifecycle of data access, discovery, quality, integration, and delivery from virtually any system.

- **ENTERPRISE-CLASS DEPLOYMENT READINESS**
EDW must enable organizations to deliver data with scalability and high performance, matched to the varying needs of each business application. Adaptability and flexibility in selecting appropriate data integration methods within the EDW is also critical to optimizing performance and costs. Real-time availability of data to front-line staff as well as executives increases operational agility and on-the-ground intelligence. Finally, gathering skilled resources and best practices within an Integration competency center

accelerates the time to market for new projects, at lower implementation costs and risk.

Architectural Considerations

One of the critical considerations for deploying EDW is creating an architecture that can support the data requirements of all types of applications. To succeed in achieving enterprisewide visibility, businesses need a seamless mechanism for users to access operational data in various transactional systems and analytic data in data warehousing environments through a single abstraction layer. This mechanism is data services.

By taking a service-oriented approach, IT is able to reuse data access, transformation, and quality logic across multiple environments, reducing the time to implement new analytics functionality. IT is also able to leverage this same infrastructure to make the data in warehouses relevant and available for operational purposes, such as single-view applications for improved customer service.

Finally, the data services infrastructure is able to support users who require data at varying levels of latency—batch, real time, and near real time. As business processes evolve and businesses move from traditional decision-support systems to more operational decision making that requires operational business intelligence, the ability to support these varying latencies through an enterprise data warehouse environment built on a data services infrastructure has become a critical factor for success. ●

For free white papers on this topic, download *"Redefining Enterprise Data Warehousing (EDW): The Imperative for Taking Your Data Warehousing Practice Enterprise-Wide"* or *"Making Data Work: Addressing Data Quality at the Enterprise Level,"* or [click here](#) for more information about Informatica.

CASE STUDY

Alpharma Pharmaceuticals LLC: Improving Pharmaceutical Sales and Marketing Effectiveness

Commentary by Freddie Hannibal
Director of Sales Operations, Alpharma
Pharmaceuticals LLC

Alpharma Inc. is a global specialty pharmaceutical company with leadership positions in products for humans and animals in more than 60 countries. Alpharma Inc., through its subsidiary Alpharma Pharmaceuticals LLC, has a growing, branded pharmaceutical franchise in the chronic pain market with its morphine-based extended release Kadian product. In addition, Alpharma Inc. is among the world's leading producers of several specialty pharmaceutical-grade bulk antibiotics and is internationally recognized as a leading provider of pharmaceutical products for poultry and livestock.

Biltmore Technologies offers My Vital Signs Rx, a hosted sales and marketing data warehouse solution designed specifically for small to midsize pharmaceutical and biotech companies that lack the business and technical expertise to support their own data warehouses. With My Vital Signs Rx, which is anchored by MicroStrategy's business intelligence software, pharmaceutical and biotech companies are able to monitor new product introductions and track the sales performance of existing products, analyze the success of marketing campaigns, and determine sales force and targeting effectiveness.

Challenge

Alpharma Pharmaceuticals LLC was challenged with finding a better BI reporting solution and empowering the sales representatives in the field with the right information at the right time. Our previous BI solution lacked the robust and sophisticated analyti-

cal and ad hoc reporting capabilities needed to meet reporting needs. We also were unable to track the sales performance of the company's flagship products and other sales initiatives efficiently. The Alpharma Pharmaceuticals sales group was unable to target the right prescribers or measure the impact of our marketing initiatives in a timely fashion. Alpharma Pharmaceuticals also wanted pharmaceutical data to be made available more quickly to our user community.

“We chose Biltmore's BI solution My Vital Signs Rx, anchored by the MicroStrategy platform, to provide ad hoc query and reporting capabilities. With the Biltmore-MicroStrategy solution, more than 200 users are able to access dashboards, perform analytics, and run intuitive reports, thereby improving Alpharma's sales and marketing effectiveness.”

Freddie Hannibal, Alpharma Pharmaceuticals LLC

Alpharma Pharmaceuticals conducted an extensive review of pharmaceutical reporting applications and found that Biltmore Technologies solution exceeded all of our user requirements. We chose Biltmore's BI solution, My Vital Signs Rx, anchored by the MicroStrategy platform, to provide ad hoc query and reporting capabilities to the field sales force, sales operations, sales and marketing managers, market research analysts, and executives. With the Biltmore-MicroStrategy solution, more than 200 users are able to access dashboards, perform analytics, and run intuitive reports, thereby improving Alpharma's sales and marketing effectiveness.

Solution

Biltmore integrated Alpharma Pharmaceuticals' sales force automation data and its syndicated data and partnered with MicroStrategy to leverage its analytical and



reporting capabilities against this data. Biltmore provides this hosted Web-based solution to Alpharma, enabling the pharmaceutical company to analyze pre- and post-marketing campaign data, prescribing habits of target doctors, sales representative effectiveness, and product performance.

With the Biltmore-MicroStrategy solution, Alpharma Pharmaceuticals receives a monthly briefing book that includes a wide variety of valuable reports customized for each user across the organization. For example, an analyst accesses a report that lists prescribers and alerts decision makers of upward or downward prescribing patterns. The vice president of sales accesses executive dashboards that provide a pulse of the business with one click to follow new product introductions and sales trends of existing products. A district manager analyzes the success of a marketing campaign and learns who is performing above or below goal.

Field users receive personalized briefing books containing information about their specific geography. Internal users receive personalized briefing books and can also access and execute ad hoc queries via a robust reporting library.

Results

Alpharma Pharmaceuticals LLC has a single data source for all sales and marketing business-critical reporting to support key business strategies, including the management of its continued business growth. Users are empowered with easier access to consistent and accurate data in a more complete and timely fashion, with the MicroStrategy platform enabling greater ad hoc querying and analytical capabilities. Users can analyze data directly against the data warehouse for analysis and decision making that was not possible before. The result: more precise analysis of prescriber targets, more effective call plans, and more focused sales and marketing campaigns. ●

LESSON FROM THE EXPERTS

Solving the CIO Dilemma: Breaking the Reporting Backlog with a New Formula for Self-Service

By **Mark LaRow**

Vice President, Products, MicroStrategy

Every CIO who has invested in a data mart or data warehouse faces a daunting problem of report development backlog. Even though BI projects start off with modest reporting requirements of 20 reports for the HR department and 15 reports for Finance, the report requests quickly balloon well beyond this initial set. It seems as if every report inspires the user to request two more related reports to help answer questions raised in the original report.

To provide an idea of how serious this problem is, consider the vast number of reports that can be generated from even a modest size data mart. Imagine a data mart with the following:

- 5 dimensions
- 10 attributes per dimension
- 10 metric areas
- 20 specific metrics in each area

If we assume that all users will only ask for one simple uniform report structure containing two attributes and two metrics, and consider all the potential combinations and variations of data in the data mart, it would require a staggering 1.9 million different reports to show all the useful combinations of the data. (This even assumes that not all attributes need to be combined with all metrics.)

The traditional response to this problem proffered by BI vendors is to encourage users to design their own reports using self-service. However, this strategy has only had

limited success. The subtleties of report design and the mechanics of design interfaces limit this solution to only power users. Even then, 1,000 power users would need to create 1,900 reports each to satisfy all realistic demands.

The more practical answer lies in a lesson learned from the Internet where almost one billion people browse and find information without becoming Internet programmers. While HTTP, search, and browsing works for Internet documents, a different technology is needed for surfing through the data warehouse. That technology is relational online analytical processing (ROLAP). ROLAP allows a handful of reports to deliver the same analytic range as thousands of traditional reports using two powerful features unique to ROLAP: object prompting and drill anywhere.

Object prompting allows users to freely select any attributes and metrics to appear on a report each time they run the report. Once a user selects initial attributes and metrics and runs the report, he or she can drill down anywhere within the multi-dimensionally modeled warehouse and, thus, successively surf through combinations of data using simple right mouse click actions.

In this unique model of self-service, users see new combinations of data that might have never been explicitly designed before, because the ROLAP system can dynamically assemble any combination of attributes on the fly. Once a user finds a useful new combination of data, they can save that combination as a newly created report, just



like bookmarking a useful Web site, and can then access it directly in the future. In this way, users can create new report designs as a natural part of their investigative process without ever having to learn how to design a report. What's more, users can share those newly created report designs with other people without fear of exposing inappropriate data, because the ROLAP system automatically applies personalized security filters to every report regardless of how that report was created or by whom.

With ROLAP enabled self-service reporting capabilities, such as object prompting and drill anywhere, users are empowered to design their own reports exactly to their specifications. Streamlining the reporting process enables any business user, not just power users, to efficiently access the data they need to manage their area of the business. ●

For a free white paper on this topic, [download](#) "Enterprise Business Intelligence: Improving Corporate Performance Through Integrated Reporting, Analysis, and Monitoring," or [click here](#) for more information about MicroStrategy.

CASE STUDY

Pitney Bowes Improves Total Cost of Ownership with Enterprise Business Intelligence

Commentary by William Duffy

Data Warehouse Project Manager, Pitney Bowes Inc.

Pitney Bowes is a mailstream technology company that helps organizations worldwide manage the flow of information, mail, documents, and packages through its integrated mail, messaging, and document management solutions. Its 35,000 employees deliver technology, service, and innovation to more than two million customers worldwide. It operates four call centers and employs 1,250 sales associates and 1,500 field service representatives in North America alone.

Pitney Bowes turned to Oracle for a comprehensive business intelligence (BI) solution that not only brought cohesiveness and greater analytical capabilities to its sales and marketing operations, but would also work in the context of Pitney Bowes' multifaceted IT infrastructure. "Our technology requirement spans more than 10 legacy systems deployed on different technologies," said William Duffy, data warehouse project manager at Pitney Bowes Inc. "Oracle Business Intelligence Enterprise Edition was one of the few BI tools capable of meeting the challenge."

Over the past few years, Pitney Bowes has been on a nonstop growth trajectory, fueled by a string of acquisitions and a worldwide push to expand its service operations. To ensure that all customers—new or old—receive fast, appropriate services and offerings, the company launched a program called "The Power of One" in 2002, which sought to provide customers with a consistent experience across all business units.

Pitney Bowes receives approximately 30,000 customer calls daily. "On any given day, each of those customers has the potential to speak

to a salesperson, need a repair on a piece of equipment, or need to talk to someone about a bill," Duffy said. "Managing all that information is a huge challenge." So Pitney Bowes sought continuity and compatibility with its CRM solution and broader IT infrastructure, which includes several Oracle databases in its data warehouse environment. "Oracle is our primary database platform, with well over 600 different production databases installed. We are running Oracle Real Application Clusters in some of our high availability environments," Duffy said.

"With Oracle's business intelligence solution, we were able to deliver more than 400 reports to a large organization with just one person—now that's cost effective."

William Duffy, Pitney Bowes Inc.

Pitney Bowes needed to enable its sales representatives to associate each customer inquiry with a type of account—large or small—to ensure they were providing the appropriate level of service. Its Oracle solution now allows each representative to immediately access all information on each customer.

"We needed a business intelligence solution that would provide all customer-facing associates with the same customer information in real time so they understand who the customer is—regardless of the customer's location, history, and time with Pitney Bowes. We also wanted to be able to analyze productivity and performance, which is key to

helping us deliver the best possible service" said Duffy.

"Oracle business intelligence provides laser-like visibility into the performance of every sales rep," Duffy said. "We can better understand how agents are spending their time and what's not working. Oracle's business intelligence is key to our success in managing key performance indicators in call center productivity, field service organization, and effective campaigns."

On the marketing side, Duffy noted, "Our biggest challenge was creating campaigns and using the data effectively to understand the buying and service behaviors of the customer. Now analytic dashboards drive every important customer retention campaign, including identifying and targeting customers with expiring leases. By enabling marketing to segment customers, we can ensure that we call the right customers at the right time so we can retain those customers and generate new business."

Taking advantage of the prebuilt Oracle BI applications, Pitney Bowes was able to reduce the total cost of ownership of its BI infrastructure and also integrate BI into its business processes. Because the solution is based on a "hot-pluggable" architecture, the company can get insight from its legacy systems and third-party solutions and analyze information from historical and real-time data sources.

Other benefits achieved include enhanced sales productivity through better customer insight, increased responsiveness of sales associates, leading to greater customer satisfaction, and improved marketing effectiveness with better segmentation and targeting. ●

LESSON FROM THE EXPERTS

Today's Enterprise Business Intelligence

By José Villacís

Director, Business Intelligence Product Marketing,
Oracle Corporation

Business intelligence (BI) continues to evolve and become even more strategic, both within organizations' systems infrastructure and as part of the daily working tools used by decision makers. One of the main characteristics of enterprise BI today is greater, all-encompassing insight into business operations and the individual contribution of business processes in the overall performance of the business.

This concept of pervasive BI is marked by the delivery of intuitive, role-based intelligence to senior management and front-line employees alike, not only providing insight about their functional area, but also about the organization's performance and how they contribute to it. This is based on a set of uniform definitions recognized and agreed upon throughout the organization and also used by a variety of financial and operational systems leveraging service-oriented architectures (SOA).

Consider TIAA-CREF, a national financial services organization with more than \$414 billion in combined assets under management (3/31/07) and the leading provider of retirement services in the academic, research, medical, and cultural fields. It has more than 1,500 call center agents serving 3.2 million individual investors, and 600 field agents serving institutional investors.

TIAA-CREF embeds Oracle BI capabilities within its customer relationship management (CRM) applications to improve front-line sales, support, and case management. This way it enables call center agents to quickly identify particular situations, notify key people to take action, and trigger workflows that start or change a related business process.

"Most people don't have time to search for problems and opportunities by looking through daily reports," explains Kurt Zimmer, vice president of CRM Delivery.

"That's why we're embedding BI technology in the processes that our employees use all the time. Our strategy involves delivering BI information at the point of contact—ideally at the exact moment that it's needed." (See "Radiating Intelligence," David Baum, *Oracle Magazine*, March/April 2007.)

The concept of providing employees with a complete view of their investors is not new; but it's easier said than done if you're not using state-of-the-art BI technology.

Enterprise BI technology is marked by its capability to integrate with diverse and dispersed sources of information throughout the organization.

"It's not just about sales metrics or service metrics," Zimmer continues. "It's also about cause and effect. When you can connect marketing, service, and sales, you start to see relationships that weren't obvious before."

Another characteristic of today's enterprise BI is organizations' recognition that the technology goes beyond a particular type of query, reporting, or analytics tool typically delivering siloed information. Enterprise BI is a comprehensive portfolio of technologies, involving financial performance management, operational intelligence, transactional applications, and data repositories.

For example, at the Arnold Air Force Base, Arnold Engineering Development Center—the nation's largest complex of flight simulation test facilities—uses Oracle BI technology to help monitor center-wide activities. This includes operational maintenance of a wide variety of unique aerospace test facilities,

information technology, desktop operation and maintenance, center communications, test utility operations, environmental safety, industrial health, and quality assurance, among other performance indicators.

British Telecom is one of the world's largest providers of telecommunications solutions, with operations in 170 countries. More than 10,000 users in various divisions rely on Oracle technology to analyze customer behavior, develop customer demand, and streamline financial planning processes by adopting an enterprise approach to measuring performance.

Enterprise BI technology is also marked by its capability to integrate with diverse and dispersed sources of information throughout the organization. This is referred to as "hot-pluggable," meaning that BI implementations must be open and modular. The goal is to allow companies to use and easily leverage their existing data sources, middleware technologies, and other systems regardless of environments and state of implementation maturity.

In summary, the three key characteristics of enterprise BI today are: 1) it should provide a comprehensive view of enterprise performance; 2) it should be pervasive; and 3) it should help extend the value of existing technological investments. ●

For a free white paper on this topic, [download "Business Intelligence and Service Oriented Architectures,"](#) or [click here](#) for more information about Oracle.

CASE STUDY

Information as a Service Delivers Consistent Information for Business Intelligence

Commentary by Sean T. Crowley
Market Management, IBM Information Integration Solutions, IBM Software Group

For any retailer, sustaining growth requires ongoing innovation. At one of North America's largest retailers, enriching category and item planning and streamlining product introductions help the company get new products to market more quickly. Real-time visibility to sales and order status, in-store kiosks, and service desk enhancements help improve the customer experience and further differentiate their brand. New efficiencies in managing stock and collaborative forecasting and replenishment help ensure that the products customers want are readily available on store shelves. In order to succeed in these efforts, they needed a solid foundation that provides accurate information wherever and whenever it's needed.

Complexity Hampers Progress

However, with multiple disparate applications and multiple platforms, databases, data entry points, and transaction processing protocols, business and IT staff were concerned with the accessibility and accuracy of information.

For each new application, developers had to write code to share information between applications and supporting data sources. Even off-the-shelf applications often needed to be customized. And each time application changes were made or new software versions released, the code would need to be rewritten. These processes were time-consuming and expensive and slowed the company's ability to sense and respond to market opportunities.

Additionally, different applications often pulled and updated data to and from different sources. It meant the existence of multiple

versions of key reference, customer, and vendor information. Executives worried that the lack of consistent data could affect certification and legal reporting processes as well as the overall accuracy of decision making based on their business intelligence applications.

Information as a Service

To enable change and gain greater value from its raw business data, company executives focused their IT teams on creating information services that could be readily available to any process, person, and application as needed.

The retailer has been able to reduce its application development time dedicated to integration by up to 85 percent, with integration-related development requiring between 3 percent and 6 percent of developers' time.

In order to achieve its desired result, the company leveraged a service-oriented architecture upon which applications would be autonomous from information sources and would not share any database directly. Knowledge of interconnections would be removed from both the source application and the target data source and would be stored centrally.

By looking at information as a service and decoupling it from the actual applications, the company could simplify operations, more quickly adapt to change, and better leverage their existing investments.

Increased Responsiveness with Information on Demand

IBM Information Integration software is at the heart of this company's information

initiative and enables them to respond more quickly to market needs and more rapidly identify new opportunities.

For example, with IBM Information Server software, the IT staff has created more than 100 reusable integration objects and interfaces to enable the collection, integration, and transformation of data. These components can be plugged into new applications to enable them to share information with any of the company's database subsystems and existing applications.

As a result, the company has been able to reduce the time and cost of developing new applications and take advantage of existing legacy applications when creating new Web-based services.

Cost-Effective Delivery of Accurate, Timely Information

Through the use of IBM Information on Demand technology, the company can now deliver information as a service, making it available to those who need it, when they need it.

On the IT side, application development and integration at this company is dramatically faster than the industry average. The retailer has been able to reduce its application development time dedicated to integration by up to 85 percent, with integration-related development requiring between 3 percent and 6 percent of developers' time.

On the business side, faster development means faster time to market with new services and more time available to focus on higher-value tasks. Additionally, with access to real-time data, business staff can gain greater insight from existing information and reduce the risks that come with using out-of-date, inconsistent, or bad information for decision making and reporting. ●

LESSON FROM THE EXPERTS

Dynamic Warehousing: The Next Evolution of Data Warehousing

By Gary O'Connell
Senior Product Marketing Manager, IBM Information Platform & Solutions

Traditional data warehouses are increasingly being challenged by demands for real-time data access, analysis of structured and unstructured data, and the need to synchronize core customer and product information across operational systems to create a single view of the enterprise. These changes are the result of new business requirements to leverage enterprise information more effectively in order to:

- Identify new opportunities and deliver new products to market faster
- Optimize business processes through real-time information and analytics
- Provide increased visibility to business performance
- Meet industry compliance standards for reporting

Yet most long-standing data warehouses are designed to support a relatively small number of users who access information to support strategic decisions, financial planning, and the production of standard reports that track performance. Today, many more users need to access information in context and in-line so that critical functions are optimized to run efficiently. Information about customers—both structured and unstructured—must be analyzed and delivered wherever it is needed. Key performance indicators (KPIs) should also be available at all times to monitor performance. In short, business intelligence is becoming embedded in key business processes.

To create a true enterprise view of information that supports strategic and operational functions, enterprise data warehouses must be reinvented as a dynamic source of current and historical information. Dynamic warehousing is an approach that enables organizations to deliver more dynamic busi-

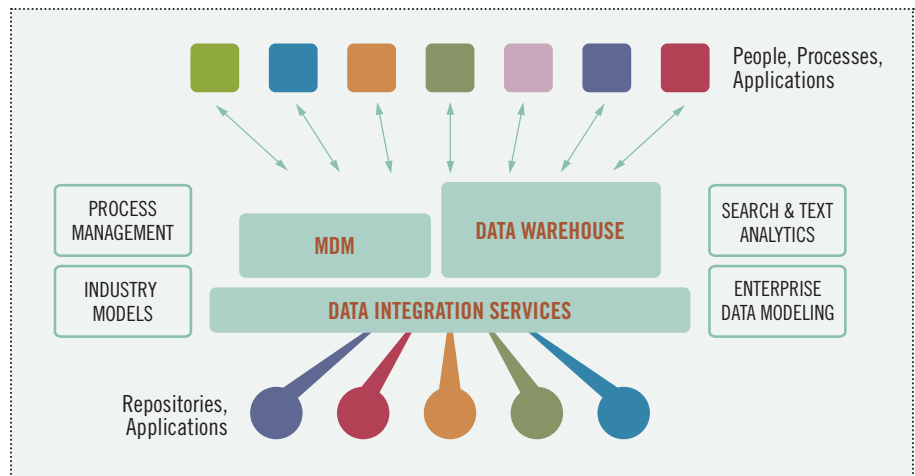


Figure 1. A dynamic warehouse has an extended infrastructure that leverages metadata to extend the data warehouse to operational systems as well as traditional business intelligence applications.

ness insights by integrating, transforming, harvesting, and analyzing insights from structured and unstructured information. Capable of processing large amounts of information, a dynamic warehousing infrastructure can enable organizations to respond on demand to unscheduled analysis requests, and as events trigger the need for information throughout the day.

A key component of a dynamic warehouse environment is the data warehouse platform. To implement a dynamic warehouse, the platform should be able to:

- Process transactions and analytical requests
- Handle varying service level agreements (SLA)
- Scale easily as the number of applications grows
- Analyze structured as well as unstructured data
- Provide real-time analytics that can be embedded in business processes
- Support advanced analytics such as data mining within the data warehouse

The requirements for dynamic warehousing go well beyond having the right data warehouse platform, however. A dynamic warehouse requires an extended infrastructure (see Figure 1) to:

- Implement changes to the business model without affecting usage

- Monitor and analyze data sources for structure and content to ensure the best data is being accessed for each application
- Provide tools that enable business users and IT staff to collaborate on data requirements and definitions
- Deliver impact analysis and data lineage reports to coordinate changes and provide visibility to critical data flows
- Deliver data that has been cleansed and harmonized to the warehouse regardless of volume or latency requirements
- Synchronize master data for key business entities across operational systems

Combining the right warehouse platform, a comprehensive industry data model with a unified platform for data integration that all share business concepts, transformation rules, and metadata will enable the deep collaboration between business analysts and IT that is required to deliver information on demand.

By creating a roadmap for a truly dynamic warehouse, organizations can meet their most pressing needs for business intelligence today, while ensuring a data warehouse environment that can help support rapid growth, significant change and increasing demand for real-time information. ●

For a free white paper on this topic, [download](#) "The Dynamic Warehouse Infrastructure: Establishing a New Foundation to Meet New Information Requirements," or [click here](#) for more information about IBM Corporation.

Q&A with the Experts

A business intelligence or data warehouse implementation can be a formidable undertaking. In these pages, leading business intelligence and data warehousing solution providers share their answers to the questions they hear often from industry professionals. Mark Hammond, an independent consultant, provides his analyst viewpoint for each Q&A.

Actuate

Q How can I accomplish true spreadsheet automation?

A Eliminate common problems at their source so that the data is correct, complete, and secure, the formulas and design structure of the document are sturdy and accurate, and the process is repeatable and reliable. Focus on:

- **The data.** Eliminate manual entry, copy and paste, and .CSV importing as much as possible; instead, manage queries through an enterprise metadata layer to ensure accuracy and security.
- **The design.** Use a “blueprint-based” design approach and tools specifically meant to improve design practices.
- **The distribution.** Automate with a server-centric architecture that manufactures pure Excel documents and does not require an Office plug-in to pull data in order to refresh the document.

ANALYST VIEWPOINT

Despite glitzier BI applications, spreadsheets remain a highly popular data analysis tool. Loyal users have been willing to tolerate the risk of errors introduced through manual data entry, as well as the time required to load and manipulate data. The business as a whole, of course, is more risk-averse and frowns on unrepeatable processes, redundancy, and duplication of “spreadmarts.” Spreadsheet automation can offer a significant improvement by centralizing spreadsheet creation and distribution in a consistent framework. Adopters should take pains to implement appropriate access and security mechanisms, and some organizations will want to examine spreadsheet automation from a compliance perspective to ensure verifiable data integrity.

Business Objects

Q Our company has seen sizable growth in our Internet sales channel over the past few years, but we know very little about the customers who use—or don’t use—this channel. How can we optimize even more sales and profitability through this channel?

A Integrate customer segmentation data from a third-party source with your transactional data from your Web site to target customers who are not yet using this channel or who abandon a purchase before completion. An online ticket broker recently did this and gained greater understanding of customer buying patterns. Now it can anticipate demand and deploy changes to its marketplace in near real time to accommodate market-changing events.

ANALYST VIEWPOINT

Third-party demographic data that matches the identity of customers who have purchased from your Web site can provide a huge advantage in building out profiles by age, gender, location, and buying patterns. Look to capture as much information as possible about buyers during your e-commerce transaction processes (short, unobtrusive surveys can be effective). Close analysis of Web site traffic patterns can pay off with revealing trends showing how surfers got to your site, and where they go next. Over time, you can build a foundation of segmented customer information that is ripe for BI analysis to determine sales channel effectiveness and customer preferences and to support targeted marketing campaigns with special offers, e-newsletters, and the like.

Collaborative Consulting

Q We have an existing data warehouse that was built to support the business, but usage is lower than expected. How can we foster adoption?

A The answer is simple, but the challenge is not. Give the business what it needs, not what you think it needs, or what they ask for the first time out. It is difficult for users to articulate exactly what they want before they see it for the first time. Low adoption or waning adoption is not uncommon once the initial enthusiasm wears off and the limitations become apparent. Revisiting the business purpose and applicability frequently with the users is key to success. Continually ask the question, “What purpose does the data warehouse serve for the business, and how well is it serving that purpose?”

ANALYST VIEWPOINT

To increase the adoption and value from a data warehouse, it's essential to first understand why usage is lower than expected. Survey the user population on what they like and what they don't. Identify the pain points that users experience. Assemble a focus group that drives collaboration between business and IT. Training sessions provide an opportunity to hear user concerns and communicate the DW value proposition. Recognize that force-feeding new technologies often generates resistance and subpar adoption. If your users are reluctant to abandon Excel, look into Excel add-ons that interface with your data warehouse to help incrementally boost usage.

DataFlux

Q What's the best way to start a data governance program?

A Any data governance program will cover three elements: people, policies, and technology. To begin any effort, companies should assess their level of maturity across each of those elements. For the “people” component, organizations can measure what level of support the governance effort has (executive, director) as well as the number of resources devoted to data stewardship. “Policies” refers to the number and complexity of business rules and controls that the company can implement within the IT environment. Finally, “technology” refers to the data profiling, data quality, and data monitoring capabilities available to enhance and maintain the value of corporate information.

ANALYST VIEWPOINT

The short answer—keep it simple! Start your data governance program by using a framework of the six basic questions of who, what, where, when, why, and how. This clearly defined structure can bring focus to your efforts and plainly delineate challenges, objectives, and roadmaps for a diverse group of stakeholders. Organizations need to be sensitive to the risk of a data governance program spiraling into unnecessary complexity and contradictory concepts from the start. By starting simply, you can build a readily understood foundation that progressively matures as participants master the nuances and details typical to data governance. Naturally, organizations will benefit by researching best practices and lessons learned from companies that have adopted data governance programs.

DATAlegro

Q We would like to encrypt our data warehouse in order to be compliant with recent legislation, but we are concerned about performance. Is it possible to encrypt data but not affect data load and data query performance?

A Absolutely. Many data warehouse administrators assume that data encryption is expensive, difficult to administer, and will create a significant degradation to performance. This is because most encryption schemes rely on software-based encryption, or they try to employ encryption of data “in flight.” With hardware-based encryption for data “at rest,” a dedicated encryption engine sits between the database and the disk storage used by each server. Since all data is decrypted before the database software starts processing it, all queries and loads work the same as before, with no significant impact on performance.

ANALYST VIEWPOINT

Identifying and implementing an ideal encryption solution can be nearly as complex as encrypted data itself. The choices typically involve trade-offs in price, performance, and degrees of security. Hardware-based solutions can sidestep the load and query performance penalty inevitable with database-level encryption by offloading encryption to a dedicated machine; however, some of these systems can be expensive and introduce an additional layer of interfaces to manage. Despite the management overhead that both software- and hardware-based encryption can impose, encryption offers the ultimate in data security and is attracting increased attention from organizations with a need to protect sensitive data.

ESRI

Q GIS seems to be deployed departmentally; can it be a strategic enterprise resource?

A ESRI has developed its products based on open standards, ensuring a high level of interoperability across platforms, databases, development languages, and applications. ESRI's GIS applications are engineered to be an integral part of an organization's strategic IT infrastructure. Our use of Web standards such as XML, SOAP, UDDI and WSDL has led to successful enterprise SOA implementations where specific location services such as geocoding, are being used throughout the organization. Our support of standard developer environments including VB, C++, .NET, and Java has led to integrations with all major BI platforms.

ANALYST VIEWPOINT

To date, most GIS deployments in mainstream business (outside of specialized utilities and energy systems, for instance) have focused on specific issues such as store location, customer clustering, and logistics. Lately, however, spatial data is increasingly featured in BI dashboards and provides business users with a visual means of understanding and analyzing information in a location-based context. Giving users an opportunity to explore the capabilities of GIS in dashboards is a good way to further adoption across the enterprise. Partnerships between GIS and BI vendors have made spatial analysis an option for many organizations, and, interestingly, Web 2.0 sites such as frappr.com are helping to put spatial data more squarely on the map.

FAST

Q What exactly is "ad hoc query and reporting" and how can I use it to improve adoption of my BI solution?

A Ad hoc tools allow users to spontaneously browse, slice and dice, and discover aggregate intelligence and detail evidence. This contrasts with more static BI models where subscribers receive reports predetermined by power users.

To provide such an application, relevant datastores (e.g., a data warehouse or internal/external sources) are connected to a data cleansing, text analytics, and high-performance index and query engine. On this platform, one or more ad hoc entry points can be provided (keyword querying, dynamically built drill-down trees, etc.) to access real-time aggregates and distribution charts.

With an ad hoc reporting tool, individuals can view the facts and figures relevant to their tasks, thereby increasing their performance and effectiveness. The combined intuitive simplicity and direct feedback are proven to dramatically increase user adoption of BI tools.

ANALYST VIEWPOINT

Theoretically, ad hoc query and reporting empowers enterprising business users to explore data by drilling down, slicing and dicing, and conducting what-if analysis. The dividends can be huge if a user discovers an unrecognized relationship that bears on efficiency or the bottom line. On the other hand, organizations need to recognize that extending ad hoc query across a broad user base can mean that some users will waste hours running queries of little business value, sometimes merely to satisfy intellectual curiosity. Organizations will profit by a judicious allocation of ad hoc query capabilities and by monitoring the usage and overall effectiveness of these analytic tools.

IBM

Q How important is metadata for successful data warehousing projects?

A Extremely important. A centralized metadata repository for information integration—one that supports analysis and active sharing of metadata artifacts across a full range of integration activities and user roles—increases the trust and understanding of the information. These artifacts include: technical metadata about the various sources of information; business metadata that describes the business meaning and usage of information; and operational metadata that describes what happens within the integration process. As companies seek to gain a more comprehensive viewpoint of their information, they need a way to keep track of how these different sources are related and how information is being pulled together and presented in reporting and analysis tools.

ANALYST VIEWPOINT

The need for regulatory compliance is a high-profile business driver that is prompting organizations to pay more attention to the management of metadata in data warehousing and operational systems. Leading metadata management tools feature robust lineage tracking, change analysis, and auditability features that can help organizations better account for the iterations and lifecycle of enterprise information at a granular metadata level. However, organizations need to be wary of the risk of architecting metadata silos and should strive towards standards-based, extensible metadata models that reduce, rather than add to, complexity in the data management infrastructure.

Informatica

Q What is the role of data integration for enterprise data warehousing?

A A critical consideration for deploying enterprise data warehousing is creating an architecture that can support the data requirements of all types of applications. This is a departure from traditional thinking, in which the data warehouse was built to support a relatively limited set of reports and analytics, and was tightly coupled to the business intelligence application.

In contrast, the enterprise data warehousing (EDW) process must function atop a data integration platform that powers the ongoing lifecycle of data access, discovery, quality, integration, and delivery via data services that can support multiple applications as part of a service-oriented architecture (SOA).

ANALYST VIEWPOINT

Because it spans multiple divisions and geographic locations and is aimed at improving enterprise performance, enterprise data warehousing puts an even greater premium on the quality, consistency, and accuracy of data. As a result, the data integration processes that support EDW are arguably the most important element in the EDW equation. Perfecting data integration technologies and processes for EDW can be a significant challenge given the breadth and complexity of enterprise data and the range of disparate source systems. However, organizations that sidestep this critical element introduce the significant risk of the “garbage in, garbage out” phenomenon and can undermine the value that EDW offers.

Information Builders

Q Use of business intelligence and analytics is not new. How have customers' requirements for these decision-support solutions changed in the last few years?

A The change occurs in two areas: people and data. Organizations access and analyze less than 20 percent of the data they collect from their business processes and systems. The nontechnical managers, operational employees, business partners, and customers demand it, and we have an opportunity to address the demands of masses of users. The interface for BI is no longer a complex tool or a manually generated spreadsheet; instead, it takes the form of custom applications, dashboards, operational scorecards, maps, and data visualization elements. Data architectures for BI must go beyond the traditional data warehouse and incorporate operational and third-party data, as well as new types of data.

ANALYST VIEWPOINT

Over the past few years, broadened BI adoption has given organizations a good read on the insights and value that may be generated by an effective BI deployment. Now customers are looking to build on success in several key areas, including: 1) enhancing the business relevance of data; 2) enabling more proactive decision-making; 3) tightening real-time BI integration with business processes; and 4) improving end-user trust in data accuracy. To achieve those objectives, customers are looking for standards-based solutions that integrate in an SOA and reduce the complexity of siloed BI deployments, transitioning towards holistic enterprise BI systems.

Microsoft

Q How does Microsoft plan to democratize performance management?

A Traditional inhibitors to the broad adoption of performance management have been high costs associated with implementations, complex tools and user interfaces that require costly and time-consuming training for employees, and confusion over disparate systems and tools for the various capabilities, including planning, budgeting, forecasting, analytics and scorecarding. Performance management tools and processes traditionally have also been siloed or stovepiped, meaning that they sit outside the day-to-day business processes of most employees. Office PerformancePoint Server 2007 was designed to address and eliminate these inhibitors, enabling performance management across the enterprise, not just for the CFO and financial analyst.

ANALYST VIEWPOINT

Giving lower-level employees tools that encourage contributions to enterprise performance is a sound idea. The concept of “democratizing performance management,” a key theme in Microsoft’s rollout of its PerformancePoint CPM solution, will prompt customer organizations to ask what performance management means for a CFO versus a call center rep versus other user roles. Organizations that extend greater performance management to lower-level employees will need to safeguard information access to ensure security of sensitive sales, productivity, and other performance-related information.

MicroStrategy

Q What's the best way for executives to consume all of the data needed to manage their businesses?

A Every day, executives and managers are faced with the challenge of reviewing volumes of data to make important business decisions. Traditionally, managers would have to find, run, and review dozens of reports and piece together similar data among those reports. With the new breed of dynamic dashboards, information from a dozen or more reports can be compressed into one expressive dashboard. Intuitive on-dashboard controls allow users to flip rapidly through interrelated information views without changing focus or context. The dashboard buttons and selectors are easy to use, so users don't need training. Managers can now find all of the interrelated data in a single dashboard.

ANALYST VIEWPOINT

Dashboards and scorecards have emerged as preferred ways for many executives to view performance information and quickly find, analyze, and explore the data they need to perform their jobs. Which solution is selected will depend on business needs, and it's worthwhile to understand the subtle distinction between the two. Dashboards are geared to inform users what the organization is doing. Scorecards emphasize how well they are doing it by tracking progress against metrics. Clarifying business objectives for a dashboard (management) or scorecard (monitoring) at the start will help your organization select and deploy the optimal solution for executives to consume and react to business information.

Oracle

Q What can my organization gain by integrating business intelligence technologies with service-oriented architectures (SOA)?

A First, it can gain an understanding of the business performance impact of business processes. Improving business results requires detailed business insight. Without enough insight into the impact and contribution of a business process on overall performance and goals, it is difficult to determine what changes should be made to improve business performance.

Second, facilitate business users taking actions as a result of insight. When business users gain insights that flag a business performance issue, the resulting corrective actions often require invoking a business process. These processes could now be started directly from the BI dashboards used to analyze performance, thereby improving efficiency and agility.

ANALYST VIEWPOINT

Business intelligence is a natural fit for the service-oriented architecture (SOA). Properly implemented, an SOA can provide a loosely coupled and extensible environment in which previously siloed BI applications can interoperate—an important step towards realizing the ideal of enterprise data warehousing and enterprise performance management. Query and reporting can be rendered as services (executing in concert with requisite data integration services) that readily extend the functional reach of analytics across a broader number of source systems. Start with small, quick-hit projects that deliver high value and build a foundation of modular components and best practices that can be incrementally extended across your infrastructure.

Pentaho

Q What can BI vendors do to become more customer-centric?

A Given the increasing adoption and strategic importance of BI, customers are now expecting greater transparency from their BI vendors. In many instances, customers and prospects want access to product documentation, product roadmaps, and unfiltered feedback from other users, or even trial software for products that they have not purchased. Customers shouldn't be afraid to ask for any of these, and vendors shouldn't hesitate to provide them. BI licensing is another area that needs more transparency. Customers hate repurchasing products or paying enablement fees for products they thought they already had paid for.

ANALYST VIEWPOINT

Steady growth in the availability and functionality of open source BI software is beginning to drive greater customer-centricity in the BI market. Traditional BI vendors are having to rethink licensing terms, free trials, upgrade cycles, and source documentation as the number of open source BI providers grows and more enterprises explore open source BI as a lower-cost, quick-hit alternative to traditional solutions. On the technology side, commercial vendors are challenged to answer increased customer demand for highly flexible, SOA-ready applications based on open standards that can be readily embedded in existing systems. As open source BI continues to evolve, transparency and flexibility in meeting customer expectations will be key.

SAP

Q Why is it important to empower business users to own and manage CPM solutions?

A Businesses change continuously. IT doesn't understand nuances of the changes, and the business users can't implement them. As a result, systems are months behind business requirements. It is important that business users are empowered to react to change and implement the changes on their own. This is most critical for corporate performance management (CPM) solutions. With budgeting or variance analysis reporting, or a strategy review, needs of users dynamic and often ad hoc; it is essential that business users model and implement changes. CPM offerings enable the business user to own, model, and maintain the solutions. This, along with the advantage of integration with ERP and built-in governance, risk, and compliance (GRC), can help business users implement changes without depending on IT.

ANALYST VIEWPOINT

The desirability of business owners using and managing corporate performance management solutions (CPM) reflects a broader trend towards greater business-side involvement in and management of BI systems. Organizations are gravitating toward a more user-centric BI model and away from the IT-led, systems-oriented data management practices of the past. Achieving this ideal requires greater collaboration between business and IT, with tech-savvy business sponsors who can bridge the gap between corporate objectives and supporting technology. CPM systems are an obvious target for business ownership because they are near the top of the BI food chain and are expected to enable game-changing insights by putting useful performance management information in the hands of executives and managers.

SPSS

Q How can I extract intelligence from textual data?

A Text is very different from the conventional, structured data held in customer data warehouses, but it is rich in context. You can mine text (for example, call center notes or responses to survey questions) to extract concepts—the most important and relevant words and phrases. These concepts, along with information on how they are linked to each other and the sentiment (positive/negative) associated with them, are effectively structured data. They can be merged with behavioral, descriptive, and attitudinal data to build a holistic customer view for predictive analytics.

ANALYST VIEWPOINT

Vendors have made significant strides in the sophistication and functionality of text analytics software in recent years. The “how” of text mining will vary by product, but the “why” is common across many organizations: unstructured and semistructured data is necessary for organizations that want to achieve a fully mature view of enterprise performance. TDWI research shows that organizations' use of unstructured data in data warehouses is on the rise, while vendors and data warehousing professionals are working to refine techniques and technologies needed to transform it into structures meaningful to a warehouse or to a reporting tool.

Sybase

Q How can I increase user scalability in a distributed analytics environment?

A A Sybase IQ multiplex configuration provides virtually unlimited user scalability. To increase the number of users, add a Sybase IQ node running on an additional server. Each node can support different users or manage different reporting or query workloads. Nodes can also partition the user base by service quality, so that higher performing nodes can support high value users or applications.

Multiplexing enables Sybase IQ to optimize workloads across multiple servers within a single environment. With one writer node and multiple reader nodes, adding additional applications and users to a database can be solved simply by adding reader nodes. The impact is minimal, as nodes do not depend upon or interfere with each other.

ANALYST VIEWPOINT

Grid computing environments are gaining attention as viable platforms to enhance scalability for large-scale data warehousing as RDBMS, BI, and data integration vendors increasingly add native functionality to take advantage of multi-node clusters. These platforms will feature load balancing algorithms to distribute queries across nodes based on resource utilization (e.g., CPU usage, memory, disk speed), and enable administrators to prioritize task execution based on user role, time of day, short- and long-running queries, and other criteria. Naturally, diligent trial-and-error application of performance tuning techniques, partitioning, and parallel processing can also pay significant scalability dividends. Apart from scalability, grid systems offer an opportunity to build in high availability and failover to better ensure system continuity.

Syncsort

Q As data volumes continue to grow, uploading data to the data warehouse has become much too time-consuming. Is there a way to speed this process?

A Uploading data can be simplified by using a data management solution with changed data capture (CDC). Also known as delta processing, the technique compares old and new files to identify changes. Only the changed data is updated in the master file, reducing the amount of data uploaded to a data warehouse, thus improving performance.

In a simple example, create a join using yesterday's transaction file as one source and today's transaction file as the other. Next, define the join key from yesterday's data to be the entire record, and define the join key from today's transaction file to be today's entire record. Keep matched records, unmatched records from yesterday's data, and unmatched records from today's data. Create three conditions (Added, Deleted, and Updated) using a tool such as Syncsort's DMExpress. Finally, add a target file and partition the data by these three conditions, creating one output file for each condition.

ANALYST VIEWPOINT

Though not new, changed data capture (CDC) has found its sweet spot with the advent of the real-time data warehouse. Available as a built-in or optional component in leading ETL tools, CDC can drastically reduce the volume of data to be moved into a warehouse. Reducing data loads by 85 percent can reduce load times at a similar rate, making multiple intraday loads practical. The additional benefits of minimizing impact on operational applications, bandwidth, and I/O make CDC worth considering by any organization that wants to freshen its warehouse data to enhance its business value.

TCS

Q How do we enhance the level of end-user confidence in the data warehouse?

A To increase end users' confidence in the data warehouse, it is imperative to build a transparent audit process using business controls for the data integration. Business controls are essentially sets of data elements at the higher levels of grain that determine the business performance, e.g., sales revenue or number of new business acquired. It is essential to establish an automated balancing control process through alternate data flows and to establish a baseline for correctness. Approaches such as general ledger balancing for financial metrics also help.

ANALYST VIEWPOINT

It's essential first to gauge the degree of end-user confidence in a data warehouse, and to pinpoint specific issues that may be undermining user confidence (and causing underutilization of the system). Depending on the gravity of the situation, this effort could be either an informal survey and follow-up, or a structured program that spans multiple divisions and extends for several months. By assessing end-user confidence in such areas as data quality, data timeliness, ease of use, and query/reporting effectiveness, organizations can build a roadmap to attack problem areas and continually enhance the business value of DW systems.

CASE STUDY

Direct Marketing Company Gets It Done Faster with the Right Tool

Commentary by Lia Szep
Senior Technical Analyst, Syncsort Incorporated

Background

In today’s complex marketing environment, most marketers have an abundance of data and good ideas but lack an efficient way to merge the two. For more than 30 years, one direct marketing company, which we will call DMC, Inc., has been building and managing customer databases for *Fortune* 1000 corporations, providing the necessary framework for organizations to aggressively apply database marketing strategies to their marketing programs.

DMC’s client base includes major corporations, many of which have databases that house data on nearly every individual in the United States. They store between 250 and 300 million names, addresses, and other demographic items, so extracting demographic data, analytics, profiles, and model scores for processing can be a cumbersome and time-consuming task.


The Challenge

With model scores for every individual in their system, DMC wanted to run decile and percentile processing at both a state and national level. With their previous process, after extracting the 240 gigabytes of data from their system, it would take another three days to process it.

The Solution

Syncsort provided DMC with a free proof of concept (POC) to demonstrate the benefits that DMExpress could provide. Using DMExpress, DMC was able to significantly decrease the elapsed time. “We extracted 240 gigabytes of data, stripped off the model scores, and ran the processing within 10 hours,” explains the senior database developer. “Now, while names and addresses are

Direct Marketing Company Shaves Days Off of Processing Time with DMExpress



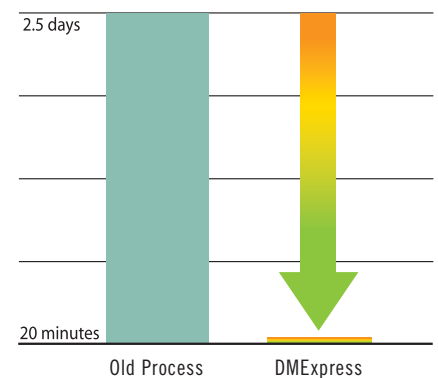
Project 1: Processing 240 GB data

being processed, I can strip off the model scores and finish the decile and percentile processing in less time than it would have taken to get the model scores.” With a 98 percent decrease in elapsed time, DMC determined that DMExpress would provide major performance benefits not only here, but across many other projects.

In another project, DMC was processing data on the mainframe. The project involved importing data to the mainframe, scheduling the job, running the processing, and sending the data to a flat file. All of this would occur while other processes were running on the mainframe. Because of this overload, the entire job would take two to three days to complete. Using DMExpress, DMC was able to completely remove the mainframe from the process and complete the job in 20 minutes.

Conclusion

Prior to DMExpress, DMC was unable to run weekly reports because the processing itself would take five business days. Now,



Project 2: Processing data on a mainframe

the development team receives the files on a Tuesday, and they’re completely done by Thursday. In addition, DMC benefited from the product’s ease of use and Syncsort’s technical support. “The ease of use has been a tremendous benefit,” the senior database developer comments, adding that “people can proficiently use the application without needing three weeks of training. Also, Syncsort’s tech support has been outstanding. From the very start—from the POC—up to my ‘weekly chats,’ I don’t think I’ve worked with a technical support group that has been as attentive and provided the level of answers that this group has. They are just fantastic.” ●

When Enterprise ETL Slows Down ... Speed It Up

By **Lia Szep**

Senior Technical Analyst, Syncsort Incorporated

Many of your BI applications rely on data gathered from your Web site. With 5 gigabytes of new Web logs a day, you've been tasked with developing an application for Web log processing. You've done all your homework, chosen the latest and greatest tools, and made your recommendation to management. The tools are bought, the development is done, and the project is in production; it seems you've made the right decision. Unfortunately, as hundreds of thousands of records turn into millions, and then billions, the system grows sluggish and reporting deadlines begin to pass unmet. Typically, the next phase of the project begins with the head-scratching question, "What was I thinking?" and ends with management's directive to "make it work." The last thing you want to do is start from scratch—so don't.

Growing Pains

The most likely reason for a processing meltdown is exponential data growth. The more successful a company becomes, the more data it generates. As companies continue to be bogged down by data growth, performance bottlenecks create pain points throughout the organization. A popular Web site can create billions of clickstream records in a single day. High-volume sales drill down to detailed financial records, increased customer data, and additional marketing demographics.

Then there's legacy, inventory, billing, etc. This wealth of data often translates to longer processing times, increased costs, and missed reporting deadlines. When an enterprisewide extract, transform, load (ETL) solution begins to slow down, the impact is felt in everything from daily business transactions to high-level business analytics.

Don't Go All Out

Code tweaks are not likely to speed your processing, and new hardware can be a costly Band-Aid that will eventually fall off. Even costlier would be a complete rip-and-replace. Instead, find a cost-effective solution that improves the performance of your current tool rather than replaces it—just make sure it fits the bill. In addition to finding a solution that works with your existing ETL, consider the following objectives.

Performance

Perhaps the most important objective will be to minimize the run time of an application. No matter what the other aspects of your system might be, unless you can minimize the time it takes to complete a job, you have no chance of meeting reporting deadlines. The key to minimizing run time is a solution that offers powerful aggregation and join tools.

High-Performance Aggregation

Aggregates are the best way to speed warehouse queries. A query answered from base-level data can take hours and involve millions of data records and millions of calculations. With precalculated aggregates, the same query can be answered in seconds with just a few records and calculations.

High-performance aggregation simplifies the creation, administration, and execution of aggregation jobs. It summarizes data much faster than other aggregation methods such as C programs, SQL statements, or third-party, multipurpose data warehouse packages. It provides the flexibility necessary to select the best aggregates for optimizing query performance.

High-Performance Join

High-performance join significantly improves the efficiency of preprocessing, retrieval, and updating in a dimensional data warehouse. With high-performance join, you can:

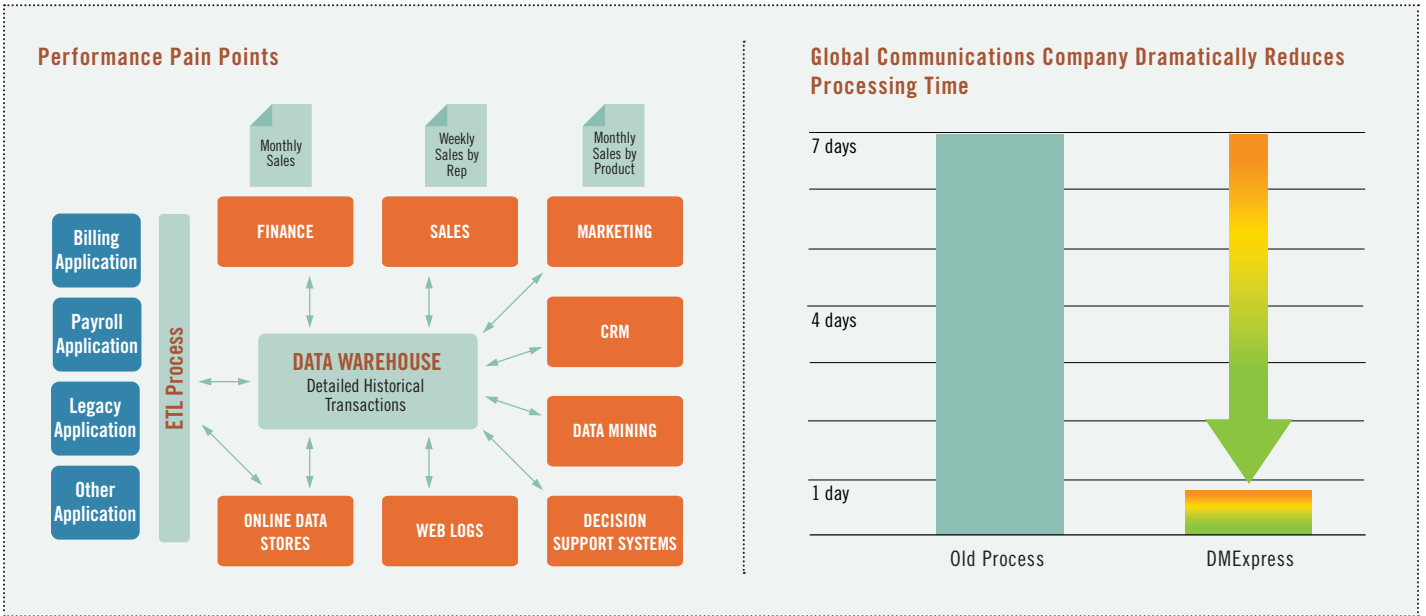
- Optimize data preparation
- Improve query performance
- Reduce the quantity of data that must be processed in response to a query
- Speed lookups and application matching
- Retrieve and summarize information more efficiently
- Minimize storage and throughput requirements
- Reduce elapsed time of changed data capture/delta processing

Scalability

The solution must also be scalable. As companies go from gigabytes to terabytes and ultimately petabytes, scalability becomes a necessary tool in managing exponential data growth. Other things to consider are ease of use and vendor reputation. If a product is easy to use, it can help control the cost and time of training staff members. If the solution is from a reputable vendor, it is more likely to have been proven reliable.

Testing

Finally, whatever other options you may be considering, testing the product in your own environment, with your own data, is absolutely the only way to determine a product's performance. Not only must the solution be



the best—it must be the best for you. This is also the only way to truly tell if a solution will work with your current tools.

Real-Life Example

Our customers are often surprised to learn that, rather than replacing their existing ETL, they can reap more value from it simply by supplementing it with a high-performance solution. In one instance, a top telecommunications company needed to convert and migrate more than one terabyte of telephone numbers, phone inventory data, and activation records from its existing system to a new one. This data needed to be converted and normalized in the old system first, implementing complex business rules, so that it could be efficiently used in the new environment, consisting of an HP9000 800 Series with 8 CPUs.

After several dry runs of the more than 100 processes using Informatica and C++ to convert the data to the new application, the company estimated the project would take a week to complete. The time factor was of critical importance: during the conversion/migration process, all retail stores would have

to be shut down—for seven days. This meant that every hour could cost the company millions of dollars in lost revenue.

Using the high-performance solution, the company was able to speed the sorting, summarization, joining for lookups, and conditional processing and transformation of the data. Five major input files were created, two for a phone number inventory scheme (CTN), and three for a phone inventory scheme (ESM). High-performance join was used for the lookups, after which the data was processed and transformed to meet the company’s requirements.

While the company still used Informatica, the high-speed solution was able to replace many of the slow-running processes, which cut the overall elapsed time by 88 percent. As a result, the project, which the company had anticipated would take one week to complete with just Informatica and C++, was finished in less than 20 hours. Without having to replace their existing system, the company was saved from a potentially heavy loss of revenue, something that certainly would have happened if the stores had been

closed for an extensive period of time to accommodate the project.

Conclusion

Ripping out a sophisticated, undoubtedly expensive, system may not be the solution to your performance problems. In your search for a fix, remember that performance is the bottom line. Minimized elapsed time, combined with scalability and ease of use, add up to increased performance. Also, in evaluating a solution’s performance, be sure you can do so in your own environment. This is the only way to determine the best solution for you and for your current system.

If you find the right tool to complement your current system, the light at the end of the tunnel may end up even brighter than you first expected. ●

For free white papers on this topic, [download "Speeding ETL Processing in Data Warehouses"](#) or ["Mitigate Business Intelligence Project Risks with Rule-Based Audits and Proof-of-Concepts,"](#) or [click here](#) for more information about Syncsort Incorporated.

CASE STUDY

Large-Scale Retailer Gains Easy Access to Historical Data for Big-Picture Analysis

Commentary by Jesse Fountain

Vice President, Product Marketing, DATAlegro

Background


A large retailer with U.S. and international subsidiaries approached DATAlegro as part of its search for a solution to data warehousing challenges. As of early 2006, the retailer operated several thousand full-line retail stores and more than a thousand specialty retail stores in the U.S., as well as several hundred full-line and specialty retail stores internationally. The company also offered its products through its Web sites.

Business Challenge

A recent merger resulted in the company owning two Teradata systems, each holding 50 terabytes of data. The retailer was struggling to analyze its data across the entire enterprise. Both systems used Teradata to support operational workload and traditional data warehousing ad hoc usage. Upgrading Teradata to accommodate the data from both chains into one system was prohibitively expensive. In addition, the high cost of storing historical data on each system prevented analysts from getting a complete picture of the business.

While the customer didn't feel it was necessary to merge all the data into one data warehouse, it did need to obtain an enterprisewide view of its customer and point-of-sale transaction data. Business needs would also benefit from the ability to store and query five or more years of data. The company had to decide if it could afford to upgrade Teradata in order to accommodate all the data it wanted to store in one data warehouse (up to 80 terabytes over time), or if there was a less expensive option that would still provide the necessary performance.

System Architecture



The retailer decided to run benchmarks with a number of leading vendors. For the proof of concept, the company selected DATAlegro, another data warehouse (DW) appliance vendor, and two other well-known, traditional data warehouse technology vendors to participate.

The Benchmark

The retailer gave each vendor one year's worth of sample data and tasked them with expanding the data to 10 terabytes (enough for 24 years of data). Vendors had two weeks to prepare and load the data into their product offerings before the retailer began running a broad range of query types, including ad hoc, complex, and long-running.

Tangible Benefits

After analyzing the query speeds, capacities, and costs of the offerings from each vendor, the customer selected DATAlegro, believing it would best support its current and future needs. Particular areas of strength were

DATAlegro's partitioning, its ability to modify DSQL to tune highly complex queries, the B25 backup unit, and the flexibility of the overall system configuration.

In addition, DATAlegro's price was less than half that of the other vendors, discounts included, without any sacrifice in performance.

The company can now reduce the amount of history on Teradata to one or two years and keep full history (five years or more) on the DATAlegro DW appliance.

By purchasing DATAlegro, the customer will save hundreds of thousands of dollars in annual maintenance costs and have a more comprehensive view of its current and future business. ●

Creating a True Enterprise Data Warehouse with Grid-Enabled DW Appliances

By **Stuart Frost**
CEO, DATAlegro

Data warehouse installations at large companies generally fall into three architectural categories:

- Centralized enterprise data warehouses (EDW)
- Decentralized collections of data marts (DM)
- Attempts at hub-and-spoke architectures that combine the two

While many organizations have obtained significant value from their data warehouse installations, few have been entirely successful at implementing full-scale enterprise data warehouses. Centralized EDW installations tend to be extremely expensive and inflexible. Consequently, business units become frustrated because the EDW won't meet their needs at sensible cost and within a reasonable timeframe. Decentralized data marts often result in many versions of the same data that are difficult to keep consistent across the enterprise. While a true hub-and-spoke architecture would address many of these issues, technical limitations with current data warehouse infrastructures have made them difficult to implement.

A Grid of Appliances

DATAlegro's DW appliances are generally used as a data warehousing "black box" with data access from a single point (the control node). However, the nodes within a DATAlegro appliance are actually self-contained Ingres database servers. Therefore, a DATAlegro appliance could be viewed as a

highly specialized grid of servers being pulled together to collectively form a DW appliance.

Taking this view, it is a small step to think of a connected set of DATAlegro systems as both a grid of appliances and a grid of nodes. Moving or loading data could be done directly between nodes in different appliances to maximize parallelism and overall transfer speeds.

Leveraging the grid concept, DATAlegro could provide a multi-temperature system that balances performance and cost across the various periods for which data must be stored.

Solving the Hub-and-Spoke Challenge

Imagine a fairly large appliance acting as the hub of a set of data mart appliances. The hub would hold detailed data (near real-time or batch-loaded) for a number of business units or perhaps the entire enterprise. ETL tools such as Informatica or SQL scripts could create star schemas. The star schemas could then be transferred to the appropriate data mart(s) via the grid at more than a terabyte per minute, depending on the number of nodes in each target appliance.

Users would connect to the independent DM appliances as usual for running queries. Each DM would be tuned according to business needs and sized to handle the required level of performance and concurrency.

Multi-Temperature Data Warehousing

Data warehouse managers are under increasing pressure to store large amounts of historical data at the same time as improving general query response times—without exceeding tight budgets.

Leveraging the grid concept, DATAlegro could provide a multi-temperature system that balances performance and cost across the periods for which data must be stored. For example, assume that the data warehouse must store seven years of data for compliance purposes. The most recent quarter (and most frequently requested data) would be placed on a very high-performance appliance. Data from three to 12 months could be stored on a standard DATAlegro appliance with very good performance, and data older than one year would be stored on one of DATAlegro's archive appliances offering up to 200 TB of user data storage per rack.

As fresh data is loaded, older data would be automatically aged (moved) across the grid. Incoming queries would be automatically broken down into the relevant date ranges and the responses from the appliances collated into a single result set before being sent back to the user.

Disaster Recovery

The grid concept could also be extended across multiple data centers to provide a highly effective disaster recovery strategy. Individual appliances could be replicated on a second site and automatically synchronized with node-to-node replication.

Summary

A high-performance hub-and-spoke architecture that is easily managed and cost effective on an enterprise scale is now a practical reality with DATAlegro's grid technology. Added benefits include support for multi-temperature and disaster recovery integrated into the grid. ●

For a free white paper on this topic, [download "Using Grid Technology to Build a Hub-and-Spoke Data Warehouse Architecture,"](#) or [click here](#) for more information about DATAlegro.

CASE STUDY

Alvion Scales Data Analytics and Distribution

Commentary by Bojan Belovic

Database Administrator, Alvion Technologies

Alvion Technologies, Inc. provides out-sourced data management services to some of the largest marketing data list owners in the world—names such as Axiom, Equifax, Experian, and Dun & Bradstreet. Alvion sits squarely in the crossroads of e-commerce and business intelligence, leveraging their exclusive technologies, extensive very large data warehousing (VLDW) expertise, and marketing know-how to provide customers with fast, reliable marketing data.

Business intelligence is the lifeblood of Alvion and companies like it. The data they collect, often from multiple sources, must be aggregated and analyzed if it is to provide value. That requires data warehousing and analytics solutions that ensure:

- The capacity to handle increased query demand, especially for ad hoc queries
- Access to hundreds or thousands of simultaneous users
- Reduced data latency, sometimes providing nearly real-time updating
- The ability to grow with increasing data volumes
- Lower hardware, administration and maintenance costs
- No downtime
- High levels of security

Alvion's customers' datasets range up to 190 million records and 200 attributes. Combined, Alvion manages about a terabyte of data for customers as an outsourcing service. Individual customers submit anywhere from 5,000 specialized records of information, to 120 gigabytes of data. Alvion then runs customer-specific data transformations and uploads it to their production servers for access by end users—widely distributed customers of the original data owners. Data update cycles vary from daily to annually.

End users of this data typically run queries and counts to determine optimal selection criteria for specific direct marketing programs. This means that a typical user session includes numerous queries, run with slight variations, to find the strongest candidates for final data selection. As the number of queries a user makes increases, so does the need for speed.

“We use IQ in production because its biggest strength as far as we are concerned are those really, really fast counts,” said Bojan Belovic, database administrator for the Alvion system. “The response time is absolutely the number one reason we are using Sybase IQ. We are talking about an order of magnitude difference.”

“Our fast query times are what defines us in this industry. We are building in more complexity with Sybase IQ as the backbone, while maintaining lightning-fast response times. That’s going to take us a long way.”

Bojan Belovic, Alvion Technologies

Alvion's success in recent years has translated into significant growth:

- More data hosted from new customers
- Growth in existing customers' record counts
- Increasing user sophistication and query complexity, and overall data mining goals
- Four-fold end-user growth

Throughout this period, the number of daily ad hoc queries has increased from 2,000

to 7,500—one-third of which are highly complex—while the number of registered users has jumped from 12,000 to 50,000. The system has scaled easily to these new levels, thanks to Sybase IQ, which has also supported the overall tremendous growth in data volume via its unique data compression capability. With Sybase IQ, disk space savings add value, especially considering the cost of the large, high-speed drives required for high-use VLDWs. Plus, the extra drives typically used to tune-up database performance were not needed.

Until now, Alvion divided its data warehouse tasks between queries and order fulfillment. Alvion is looking to consolidate all operations in a single environment, which will mean increasing the size of the production database by roughly 100 percent. Amazingly, Sybase IQ provides such effective compression that the new data will not affect the system's performance or query speed. However, this change will dramatically improve Alvion's competitive advantage by shortening data fulfillment times while maintaining the trademark query speeds so appreciated by its customers.

Belovic believes that Sybase IQ is up to that task, too.

“Our fast query times are what defines us in this industry. We are building in more complexity with Sybase IQ as the backbone, while maintaining lightning-fast response times. That’s going to take us a long way.” ●

LESSON FROM THE EXPERTS

The Analytics Server: Powering On-Demand Analytics from the Ground Up

By **Phil Bowermaster**

Worldwide Product Marketing Manager, Sybase IQ

When looking to design and implement a system that will meet their customers' requirements, companies that provide data analytics services face a number of significant challenges. They need a solution with the capacity to handle heavy and growing query demand, especially for ad hoc queries. They need a solution that can be accessed by hundreds or thousands of simultaneous users without a significant impact on performance. They need to reduce data latency, sometimes providing updates at or near real time. Of course, they also need to minimize downtime and maximize security—all while lowering hardware, administration, and maintenance costs.

These demands aren't easy to achieve, and those who provide data analytics services often run into roadblocks trying to achieve them using traditional approaches to data warehousing.

For example, a company providing this type of service will deploy a solution built on a standard relational database management system. We've seen a number of instances where companies find that, when implementing an on-demand data analytics service on a standard RDBMS, complex queries on large volumes of data run very slowly. This happens because standard relational database management system (RDBMS) solutions are designed for transactional, rather than analytic, environments. They store data in tables by row, where each row is effectively equivalent to a transaction. Querying such a database means grinding through many entire rows of data, when all that the user is looking for is a few select items found in some of the columns.

Another common approach is to use a data warehouse appliance. The typical appliance solution is designed for analytics, so it will offer a significant speed advantage over a transactional database where individual queries are concerned.

Companies that depend on fast, ad hoc query responses by many users accessing large volumes of data need a different kind of solution, one that's built from the ground up to support on-demand data analytics. This is where an analytics server comes in.

Companies using this approach can run into difficulties, too. The fact that appliances run on proprietary hardware limits a company's options in defining their own environment. They may also find that the appliance approach limits their ability to scale up to support multiple users, and that multiple queries from hundreds or thousands of users can choke the system. To make matters more difficult, database administrators are often limited in the amount of tuning they can do for specialized applications.

Companies that depend on fast, ad hoc query responses by many users accessing large volumes of data need a different kind of solution, one that's built from the ground up to support on-demand data analytics. This is where an analytics server comes in.

An analytics server, such as Sybase IQ, takes a fundamentally different approach to data analytics. It stores and accesses data in

tables by column. This serves to index automatically the entire database, because query selection criteria are defined by column. This greatly reduces the amount of data that needs to be read to respond to a query and dramatically speeds up response times.

Because of these advantages, an analytics server supports extremely fast query speeds without limiting:

- The amount of data available for analytics
- The number of concurrent users and queries
- The launching of new applications
- The complexity of queries
- The number of access windows to data
- Return on investment

In addition to its edge in query response times, an analytics server such as Sybase IQ offers a far better return on investment than other data warehousing and analytics solutions. Because of its speed, it requires fewer hardware resources than other options. Its column-based approach also reduces the need to pre-aggregate data and tune the database, saving significant administration time. All of which goes to explain why, increasingly, we're seeing the leading players in data analytics services turning to the analytics server as the platform of choice for delivering on-demand analytics solutions. ●

For a free white paper on this topic, [download "Sybase Helps Customers Deliver Value-Added Information on Demand,"](#) or [click here](#) for more information about Sybase.

CASE STUDY

BI at Tata Teleservices Ltd.

Commentary by Ramanuj Rao

Senior Manager, Business Intelligence, Tata Teleservices Ltd.

Customer Profile

Tata Teleservices Ltd. (TTSL), operates India's largest branded telecom retail chain. TTSL, incorporated in 1996, was the first company in India to launch code division multiple access (CDMA) mobile services. Together with Tata Teleservices (Maharashtra) Ltd., it pioneered the 3G1x technology platform in India and offers services to more than 19 million customers covering 3,400 towns. TTSL has partnered with Motorola, Ericsson, Lucent, and ECI Telecom to deploy a reliable, technologically advanced network.

TTSL offers an array of telecom services, including mobile services, wireless desktop phones, public booth telephony, and wire line services. Value-added services include voice portal, roaming, postpaid Internet, three-way conferencing, group calling, Wi-Fi Internet, USB modem, data cards, calling card services, and enterprise services.

The company was the first service provider in the country to launch an online outlet offering postpaid mobile connections. It also launched prepaid wireless desktop phones, public phone booths, new mobile handsets, as well as new voice and data services such as BREW games, picture messaging, and polyphonic ring tones. It also offers interactive applications such as news, cricket, astrology, etc.

On the Radar

The formerly government-run telephone system in India opened to private players in the 1990s. With liberalization and globalization, the telecommunications business is undergoing a major sea change. Telecom activities are flourishing, enhancing the infrastructure and helping technology penetrate India's diverse society. In the past two decades, the number of telephone connections in India has risen by more than 75

percent. Yet telecom market penetration is still around 20 percent.

The Challenge

Stiff competition from several other private players in India's telecom industry presents TTSL with challenges. Measuring operational efficiency and business performance is crucial, along with developing highly focused marketing strategies and financial planning. Business intelligence is the need of the hour.

"We can transform millions of daily transactions into intelligent information on the fly. This information helps our business functions for predictive and proactive approach. It acts as an enabler for TTSL to acquire and retain subscriber values."

Ramanuj Rao, Tata Teleservices Ltd.

Solution

BI in alignment with business processes provides a single version of truth with agreed definitions of key performance indicators and standardization of reports using business metadata. TTSL has partnered with Tata Consultancy Services (TCS) to develop and maintain a BI environment. TCS has implemented its telecom model, tBIDs, as the core solution.

A communication and awareness campaign plays a major role in synchronizing BI with business, which is achieved through road shows, newsletters, etc. To meet this need, business intelligence solutions are integrated with BSS/OSS applications, which help get the cluster and correlation across the captured key performance indicators.

Tangible Benefits

BI is helping to improve TTSL's operational and business efficiency. TCS achieved this using various solutions such as subscriber base management, loyalty management, OGB reduction program, revenue usage, predictive churn, customer segmentation, TAT matrix, decay analysis, and zero balance depletion program, to name a few. All these solutions give cluster and correlation across various business dimensions. Closed-loop BI was deployed in critical areas to facilitate implementation of operational action plans.

We also put into place a loyalty management program using BI—a classic example of operational and business efficiency. This program revolves around subscriber aging, recharges, and demographic patterns. It identifies subscribers per loyalty business rule and mediates credit values to an Intelligent Network Platform using TIBCO interface. Loyalty ROI is tracked in BI.

Outgoing barred (OGB) reduction program is a proactive approach for analyzing OGB customers to take corrective actions with an objective to reduce churn. This focuses on identifying key parameters on which OGB patterns can be monitored. Using closed looping, business takes corrective action to achieve reduction in OGB as well as churn.

Technology

- OLAP and Reporting: Cognos 8.2.4
- ETL: Oracle Warehouse Builder 10.1.0.4
- Current DW Size: 8 TB ●

For a free white paper on this topic, [download "5ml of BI for Growing Organizations,"](#) or [click here](#) for more information about TCS.

CASE STUDY

CPM Implementation Challenges and Lessons Learned

*Commentary by Abhijit Niyogi and Rohit Shrivastav
Tata Consultancy Services*

Client Business Scenario

The client is a large publishing corporation with business interests in education, business information, and the financial services market. The organization is structured into multiple business entities serving different niche areas or geographical regions across the globe.

The corporation's existing business information system is characterized by heavy use of Excel, a stand-alone budgeting and reporting solution, no structured workflow for routing budgets for review and approval, and a proliferation of reports and offline worksheets without any standardization across business entities.

CPM Program

Primary business objectives of the corporate performance management (CPM) program are to:

- Provide the right information to the right users at the right time to support better decision making
- Establish a platform to support management of metrics/scorecarding and analytical reporting
- Provide access to detail-level transactional data as well as integration with other internal and external data sources to increase the value of the information
- Increase transparency
- Standardize reporting and planning structures where possible
- Utilize driver-based planning and budgeting models

The CPM program was based on Cognos Solution on top of a core CPM data warehouse. All key financial data, including actual, budget, and estimate, resides in the DW. The DW is designed to integrate other internal and

external source system data—such as human resources, revenue, etc.—in the future.

Cognos Planning Solution extracts data from the DW and other source systems to facilitate driver-based detail budgeting and the monthly estimation process, along with medium-range planning. The Cognos planning system acts as a source for budget and estimate data to the DW and existing ERP system.

Cognos 8 Business Intelligence works on top of DW for facilitating Web-based standardized reports, business entities, and provides ad hoc and self-reporting capabilities to business users.

Lessons Learned

BUSINESS SPONSORSHIP

- Define clear vision and establish senior management sponsorship
- Create and strongly communicate project road map in the organization

TEAM

- Include both business and technology people in working project team, with clearly defined roles and responsibilities
- Empower team to make business and technical decisions

PROCESS

- Establish business ownership of requirements and sign off
- Manage user expectations and ensure user engagement throughout project cycle
- Prepare organization to adapt to business process changes and standardization
- Create data source ownership and understand business rules and cleansing logic for all data sources
- Plan for holistic design and staggered rollout
- Plan for iterative approach for design and development

TECHNOLOGY

- Thoroughly test and establish an appropriate and scalable technology infrastructure
- Design and test individual components of technology solution to ensure scalability and performance

DESIGN AND BUILD

- Ensure that design assumptions are clearly communicated and reviewed by business owners
- Establish and follow best practices for businesses processes and technology solution

SUPPORT

- Maintain a common pool of knowledge resources and establish a process for knowledge sharing among onsite and offshore team
- Ensure continuity by establishing a core dedicated team for the project

OUTCOMES

- Established a reliable foundation of core data, reporting, and planning capabilities. This includes process improvements, testing technology solution, and creating an environment within the organization to leverage, use, and contribute to CPM solution.
- Established a transparent and integrated driver-based budgeting and estimation system.
- Standardized and automated budget workflow and approval process.
- Designed and built a standard set of core financial reports that can be used across business entities.
- Designed and deployed self-reporting capability that is based on preverified centralized data source.
- Established data set that can be leveraged for developing dashboards and designing KPIs. ●

CASE STUDY

Hospital Finds Cure for Information Overload Syndrome

Commentary by Morie Mehyou

Assistant Vice President of IM and Decision Support,
Jefferson Regional Medical Center



ORGANIZATION

Jefferson Regional Medical Center, a health services provider located in Pine Bluff, Arkansas.

THE CHALLENGE

Enable clinicians and administrators to hone in quickly on specific information related to patients, reimbursements, payer mixes, admissions trends, staff productivity, and dozens of other business functions.

THE STRATEGY

Construct a self-service reporting environment that organizes reports into three primary groups: financial reports for the finance department, productivity reports for department managers, and bed management reports for hospital administrators.

THE RESULTS

WebFOCUS has become an integral part of the hospital's daily management activities—boosting staff productivity, improving physician deficiency rates, and cutting costs throughout the facility.

INFORMATION BUILDERS SOLUTION

FOCUS and WebFOCUS

Today's hospitals maintain an immense amount of information that must be carefully stored, managed, analyzed, and exchanged. In addition to administrative data on hospital operations, many types of patient information are entered and maintained electronically. Hospitals must streamline the flow of information—both internally and externally—even as they comply with HIPAA requirements and HL7 standards.

Maximizing reimbursements, analyzing payer mixes, tracking admissions trends, and monitoring staff productivity are all functions of having accurate information, delivered to the right place at the right time. That's why Jefferson Regional Medical Center depends on business intelligence (BI) technology from Information Builders to streamline data management and reporting activities throughout the enterprise.

"The amount of information in your average hospital can be overwhelming," admits Morie Mehyou, an assistant vice president of IM and decision support at Jefferson Regional Medical Center. "The challenge for clinicians and administrators is to quickly hone in on specific information."

As the fourth largest hospital in Arkansas, Jefferson Regional serves more than 280,000 residents in 11 counties. Licensed for 471 beds and with an average daily census of about 240 patients, the hospital employs more than 2,000 people and offers a complete range of medical services.

Treating Information Management Ailments

WebFOCUS has become an integral part of the hospital's management activities, with 75 departments now using the business intelligence software on a daily basis. "Virtually all operational data from every department—from materials management to pharmaceuticals—is now accessible via a self-service reporting environment," says Mehyou.

An accountant by training, Mehyou is no stranger to the demands for precise, auditable information. Yet before WebFOCUS came on the scene in the late 1990s, he says getting information in a timely fashion was often a hit-and-miss prospect. Most of the hospital information systems offered standard reports, but it was difficult to obtain precisely what each user needed. That meant lots of custom reporting, which dominated the available hours of the IT staff.

About that time, executive vice president Tom Harbuck asked Mehyou to develop an internal reporting system through which his staff could access information on demand—without asking IT pros for assistance. "Our managers wanted a better way to measure productivity by department," recalls Harbuck. "With WebFOCUS, we accepted that challenge head-on."

Resident in Training

Mehyou had experience with Information Builders' FOCUS reporting environment on an AS/400 computer, so learning WebFOCUS was easy. He began by creating reports for the finance department, including departmental financial statements, income statements, accounts payable reports, balance sheets, and budget reports. He went on to develop revenue, HR, and productivity-benchmarking reports—more than 150 in all, most of which are now accessible via the corporate intranet.

To give users maximum control over this information, Jefferson Regional constructed a parameterized, self-service environment that organizes reports into three primary groups: financial reports for the finance department, productivity reports for department managers, and bed management reports for hospital administrators and a few doctors. Today, this reporting system is an integral part of the hospital's daily management activities. "We have leveraged WebFOCUS for a huge number of BI projects, such as to

monitor physician deficiencies, for corporate education test scoring, and as part of our open chart review process,” Mehyou notes. “WebFOCUS helps us tabulate deficiencies, so physicians always know where they stand. For example, it tells us which patients are missing documentation, for how long, and in what areas. WebFOCUS has enabled us to provide a much more accurate and timely report relating to deficiencies.”

Meanwhile, department managers have learned to use WebFOCUS to generate productivity reports, extending back several years, which are useful for benchmarking and comparisons. “Managers might want to know, for example, how many hours it took to provide care for a certain number of patients, or how many tests they performed during a certain period,” explains Mehyou. “WebFOCUS allows them to summarize this information over any time period. It is a very flexible environment.”

Other reports have been created for diagnosis-related groups (DRGs)—the three-digit codes that describe patient conditions, from heart attacks to hypertension. Reimbursement rates are often set according to these codes, and hospitals get paid a predefined amount for each DRG. “Sometimes a number of procedures are bundled into one rate, and complicated formulas are used,” Mehyou says. “Our hospital administrators need to make sure each diagnosis is correct, the charges are justified, and that they are billing appropriately to maximize reimbursements.”

Extensive data analysis is performed every month using many FOCUS and WebFOCUS reports. “Everybody wants to see a different slice of the information,” continues Mehyou. “Some people want to see test results. Others are interested in your medical history. Still others need to know the length of your visit, review your bill, or perform statistical analysis. WebFOCUS makes it easy for these users to change parameters so they can obtain the information they need.”

Hospital personnel log into the WebFOCUS environment through a secure Web portal accessible via any Web browser. Once they enter the correct department number, they are presented with a graphical environment that directs them to the information they need. For example, department managers can review statements to make sure they were charged for the correct supplies. Other basic information includes pay reports and productivity reports. “Users can drill down to the details as required, or sort the information by user, pay period, or charge level,” Mehyou says. “Everything is protected by user and role, and the system is fully compliant with HIPAA requirements governing privacy of information.”

IT Vital Signs

To create this user-friendly reporting environment, Jefferson Regional integrated WebFOCUS with a McKesson HBOC clinical information system on an AS/400 computer. WebFOCUS accesses other AS/400 packages, such as Payroll, OR, and related business functions. These packaged applications store data in a DB2/400 relational database, which WebFOCUS can access directly. Other critical information is stored in Microsoft SQL Server data marts. “We like the way WebFOCUS can easily combine two or more types of data into the same report,” says Mehyou. “For example, we might access the AS/400 database to create a summary report for the radiology department, then compare that information to a national benchmark, which is derived from a different kind of database entirely. WebFOCUS combines information on the fly to create a cohesive presentation.”

Eventually, Jefferson Regional would like to build BI dashboards for managers, senior managers board members, and other constituencies within the hospital. “More and more patient data is now maintained electronically—everything from vital signs to medications to test results to billing information,” says Mehyou. “Much of it is stored in

a clinical application system that is easy to sort and summarize with WebFOCUS. Now that we have the proper reporting tools, we can deliver information to our chief of staff with confidence.”

For example, the hospital uses an Eclipsys Clinical Application System, which gathers many types of clinical information, such as periodic vital signs. “There are myriad data points on each admitted patient,” says Mehyou. “Ultimately, we hope to use WebFOCUS to mine this data.”

WebFOCUS has paid for itself in measurable productivity gains. For example, the labor and delivery department was operating above the 50th percentile in cost-per-delivery, on a national basis. Careful analysis of spending patterns helped the hospital reduce these costs. “We’re now in the 36th percentile, thanks to better knowledge of costs,” Mehyou says. “That’s a 28 percent improvement in productivity.”

In the face of numerous industry reports that accuse health care organizations of runaway spending and rising costs—with no noticeable improvement in quality of care—these measurable performance gains are a source of pride for Jefferson Regional. “Providing good care is our primary objective, but of course we want to do that within the bounds of what is reasonable and financially responsible,” Mehyou concludes. “WebFOCUS gives us the insight we need to create a more efficient operation. For the amount of money we have spent to purchase and maintain the software—and the ease with which we can create reports—it is one of the best technology investments the hospital has made. The cost is minimal, and the return has been phenomenal.” ●

For a free white paper on this topic, [download](#) “Worst Practices in Business Intelligence: Why BI Applications Succeed Where BI Tools Fail,” or [click here](#) for more information about Information Builders.

Best Practices in Operational BI

Challenges to Operational BI

BY WAYNE W. ECKERSON, DIRECTOR, TDWI RESEARCH

Operational business intelligence (BI) represents a turning point in the evolution of BI. Traditionally, BI has been the province of technically savvy business analysts who spend many hours with sophisticated tools analyzing trends and patterns in large volumes of historical data to improve the effectiveness of strategic and tactical decisions. But operational BI changes this equation: it moves BI out of the back room and embeds it into the fabric of the business, intertwining it with operational processes and applications that drive thousands of daily decisions. In essence, operational BI merges analytical and operational processes into a unified whole.

In addition, operational BI increases the value of BI by delivering information and insights on demand to all workers—from the shipping clerk to the CEO—so they can work smarter and faster to achieve critical business objectives. In essence, operational BI delivers the right information to the right people at the right time so they can take action. In its extreme form, operational BI encapsulates business insights into rules and models that organizations can use to automate decisions and responses, eliminating the need for human intervention. Automating decisions not only streamlines processes and reduces costs, but also improves service and gives organizations a competitive advantage in the marketplace.

Beyond operational reporting. There are many flavors of operational BI, ranging from operational reporting and process monitoring to composite applications and decision automation. While most organizations already support some form of operational reporting, many have yet to embrace more complex types of operational BI that generate greater business value. Thus, operational BI opens up a new field of endeavor for BI and gives organizations a chance to reap greater dividends from their BI investments.

Challenges

However, operational BI poses several challenges. It stretches the architectural boundaries of current BI solutions, forcing BI professionals to rethink the way they design and build systems. Queries must return in seconds rather than minutes or hours, and reports must update dynamically. Operational BI systems must capture large volumes of data in near real time without degrading the performance of existing processes and jobs on source or target systems. There is also less time to recover from a server outage, making it imperative for BI

IS YOUR OPERATIONAL BI ENVIRONMENT THE SAME AS YOUR REGULAR BI ENVIRONMENT?




Figure 1. Based on 225 respondents who have implemented operational BI according to TDWI's definition.

professionals to build resilient, highly available systems with sufficient backup and recovery.

To warehouse or not? The first technical question that an organization needs to address is whether to use a data warehousing architecture to deliver just-in-time data, or bypass it altogether. This is the most fundamental question, and the hardest to answer.

Many architects believe it's critical to adapt existing data warehouses to support operational BI. "The big showstopper is whether you are going to apply the same business rules to integrate, cleanse, and validate operational data streams as the rest of the data in your data warehouse," says John O'Brien, a BI consultant and former data warehousing architect. Pulling operational streams out of the data warehousing process undermines data quality and creates divergent data sets that may not reconcile, he claims.

Some disagree. They say that a data warehouse becomes a bottleneck if you try to load all data into it that users may possibly want to query. BI vendors such as Business Objects, Hyperion, SAS, and InetSoft, which support federated query, believe their tools provide an easy, low-cost way to capture real-time data and deliver it to users at the point of need with sufficient quality. Likewise, vendors of embedded BI, event-driven analytic platforms, composite applications, and in-memory analytics believe their offerings provide the most suitable way to meet high-end operational BI requirements. These vendors say that a data warehouse is critical for applying historical context to real-time data, but not necessary for managing the real-time data itself.

Nevertheless, data warehousing has become a well-established IT practice in corporate environments, and few organizations will jettison their DW investments without trying to adapt the architecture to support just-in-time data and operational processes. Our survey also shows that about half of organizations (51%) run both operational and analytical reporting from the same environment. Given the high number of survey respondents who equate operational reporting with operational BI, this indicates that many companies have moved operational reporting into their data warehousing environment. (See Figure 1.)

Organizations that want to deliver just-in-time data via a data warehousing environment face daunting challenges. The rest of this section describes the major challenges involved in transforming a data warehousing environment to support operational BI.

Select the right technology. There are many technologies that BI architects can use to build an operational BI environment. These technologies can be classified into three main categories that correspond to the way data flows through a just-in-time system:

- **Data acquisition.** Organizations must capture, transform, and move data from source systems into the analytical environment on a just-in-time basis. To do this, organizations can use ETL tools, replication tools, changed data capture technologies, messaging backbones, event-driven streaming engines, and BI

and EII query tools. In practice, most companies use a combination of these.

- **Data storage.** Organizations can store acquired data in a variety of database engines, including one or more layers in a data warehousing environment (i.e., staging area, data warehouse, data mart, OLAP cube database), an ODS, a low-latency database, or an event-driven analytic engine.
- **Data delivery.** To display captured data or data derived from analytic processes, organizations use custom-built applications, portals, BI tools, dashboards, composite applications, or a combination of these approaches.

Some of these technologies support multiple tasks. For instance, BI and EII tools can both acquire and deliver just-in-time data, and if you consider a report definition a storage mechanism, they support all three activities. Likewise, event-driven analytic platforms, working in concert with messaging backbones (EAI networks, enterprise service buses, publish/subscribe middleware), support all three activities. This is not surprising, since they are designed to provide end-to-end support for real-time operational BI requirements.

Our survey shows that organizations use a variety of techniques to support operational BI. The most common are ODSs, followed by more frequent batch loads and changed data capture. (See Figure 2.)

RATE THE IMPORTANCE OF THE FOLLOWING TECHNIQUES TO YOUR OPERATIONAL BI STRATEGY.

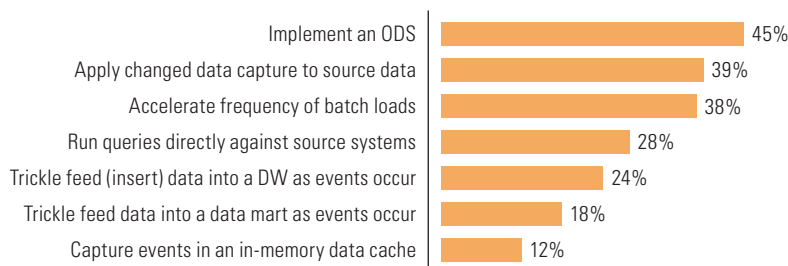


Figure 2. Based on 225 respondents who have implemented operational BI and rated the approach of “high” importance.

Increase scalability and throughput.

Another key issue is building systems that scale to support greater numbers of users, more data sources, and higher volumes of data with increased rates of throughput—while ensuring high data quality and security. Organizations need to implement high-performance data warehousing platforms from leading companies such as Sybase, Teradata, and HP, and robust data integration platforms like those from Syncsort and others. To ensure scalability and throughput, companies may

need to upgrade their networks and hardware, and parallelize key extraction, transformation, and load processes so they can adequately meet these increased throughput demands and eliminate processing bottlenecks.

Use inserts. To increase ETL throughput, some (but not all) architects recommend inserting new records (i.e., loading data) rather than adding, changing, or deleting existing records. “Avoid updates at all costs,” says O’Brien. “With large data sets, it takes too much time to find the record and update it.” Although this creates lots of duplicate data, administrators can use SQL DISTINCT and GROUP-BY statements to identify and delete duplicates, O’Brien says.

O’Brien also recommends turning off referential integrity and database logging. “Why log inserts if you have the original load file?” he asks. In addition, there is no reason to create indexes and aggregations during the ETL process, since they consume a lot of CPU and slow down the inserts. Like many architects, O’Brien now advocates using the database to perform these tasks once the data is loaded into the data warehouse to exploit the parallelism of the underlying database. This process is known as extract, load, and transform (ELT) versus the more traditional extract, transform, and load (ETL).

Increase availability and recoverability. In a just-in-time data warehousing environment there is little time to recover from errors, because batch windows are small or nonexistent. “In a real-time system, you have no time to fix problems, and messages can start piling up,” says OLAP architect Eric Lofstrom. The best way to avert these problems is to build high-availability systems to drive the acquisition and deployment process. These systems parallelize single points of failure, run in a clustered environment, support failover and backup processes, and have an offsite backup.

Microbatches. To avoid losing data, O’Brien recommends using “microbatches” instead of real-time streaming, since microbatches provide a buffer with which to recover from a network or server outage. His rule of thumb is to set the ETL batch cycle three times larger than the time required to process the data. “If it takes 90 seconds to process a file, then I’ll set the batch cycle to five minutes and have 3.5 minutes of sleep time.”

However, other IT practitioners believe it is perfectly safe to stream events into a data warehouse. Lofstrom, for example, recommends logging each message as it comes off a queue, which may slow down the speed of the streaming data, but prevents the loss of data caused by server outages. If an outage occurs, Lofstrom recommends loading accumulated messages into the warehouse in a single batch job before turning event streaming back on. In other cases, companies use two-phase

commit to update streamed messages from a queue into a staging area, so updates are never lost if there is a planned or unplanned outage in the database server.

Deliver adequate performance. A challenging problem for data warehousing designers is ensuring rapid query response times while simultaneously loading or updating a data warehouse and performing other tasks, such as monitoring events, triggering alerts, running backups, and scoring models. A mixed workload can cause RDBMS performance to degrade to the point where it blocks incoming queries or causes load processes to fail. Achieving adequate performance in a mixed workload environment can be a vexing challenge.

Complicating matters is that most users expect near-instantaneous response times in an operational BI environment. For example, customer service representatives can’t wait more than a few seconds for an integrated view of customer activity and cross-sell offers before their callers get impatient and hang up. Most data warehouses aren’t designed to support a mixed workload of tactical queries that return a few rows of data in less than a second along with complex requests that scan millions of records and take minutes or hours to process.

How do you balance these mixed workloads and maintain SLAs without causing the system to crash? There are several options, all of which have trade-offs:

- **Throttle incoming requests in response to heavy system loads.** You can schedule queries or jobs to only run at night, or queue or delete them using query governors or job managers during periods of heavy load to ensure the integrity of system updates. This approach obviously won’t sit well with users, and it’s the reason companies began building data warehouses in the first place: to provide the business with a dedicated analytical environment that would not interfere with operational processes.
- **Separate analytical and tactical queries and jobs.** Another approach is to off-load tactical queries and predictive scoring jobs to an ODS or specialized data mart and leave the data warehouse to support long-running analytical queries. Here, the ODS contains current transaction data, usually no more than 30 days’ worth, stored in third normal form models that reflect operational schemas. The downside of this approach is that it separates historical and current data into different systems, making it difficult for users to view just-in-time data in historical context.
- **Leverage RDBMS mixed-workload capabilities.** Many database management systems (DBMS) have significantly improved their ability to optimize the performance of

long-running strategic queries, short-running tactical queries, load and update processes, and other types of workloads. Teradata, for example, has made big improvements in its priority scheduler for optimizing tactical queries, says Dan Graham, marketing manager at Teradata. The priority scheduler is a real-time traffic cop that always allocates CPU to tactical queries first to guarantee sub-second response times. It continuously monitors in-flight tasks and dynamically adjusts CPU allocation to optimize performance. In addition, administrators can configure and prioritize workloads by user groups, types of activity, and other variables. HP's new Neoview data warehousing platform also boasts workload management capabilities.

- An RDBMS that supports mixed-workload processing lets organizations have their cake and eat it, too: they can load current and historical data into the same database and optimize performance across all types of queries. Without an RDBMS that supports mixed workloads, many organizations decide to avoid using a data warehouse for operational BI. However, not all mixed workload capabilities are created equal, so evaluate your database vendor's capabilities carefully. Also, running multiple workloads on a single platform may require hardware upgrades to maintain adequate performance, so calculate these additional costs when deciding whether to use a data warehouse to support operational requirements.

Avoid contention problems. When user queries and ETL processes hit the same table, there is the potential for one process to block the other and cause query or load performance to degrade. Along with mixed workload requirements, the issue of contention causes many people to maintain historical and just-in-time data in distinct data stores. There are three ways to minimize the contention problem, but they are not for the fainthearted:

- **Allow simultaneous reads and writes.** If you are inserting rather than updating data, the database won't lock the table. The only problem here is that the data can be out of sync if the user reads the table before all the inserts have completed. To avoid confusion, it's imperative to provide users with guidelines for submitting ad hoc queries against real-time data, and it helps if you dynamically time-stamp all reports. It's also important to turn off caching in BI servers, since there is no point in loading data every hour if the BI tool requests data from a cache that gets flushed once a day.
- **Insert and flip partitions.** Another approach is to create a duplicate partition of a fact table and load it with current data. Then, on a periodic basis, swap the table with the live fact table to update the data warehouse.

This approach may require you to pause the application server while the flip takes place so no new requests are initiated while the swap occurs, says Justin Langseth, CTO of Claraview, a BI consultancy in Reston, VA. Many organizations use this "insert-and-flip" approach.

- **External real-time cache.** Another approach Langseth recommends is to load data into an in-memory cache outside of the data warehouse. Requests for real-time data are fulfilled from the cache, while queries that require a combination of real-time and historical data merge the requested data in a set of temporary tables in either the data warehouse or real-time cache, depending on which database holds the largest volume of data requested. This type of just-in-time merging of information requires complex SQL and may be challenging for most BI tools to support.

Summary. There are many challenges involved in implementing operational BI. The challenges discussed here deal primarily with transforming a data warehousing architecture to support just-in-time data delivery. Although there are other ways to deliver just-in-time data to support operational BI, many architects recommend using a data warehousing environment to ensure a consistent set of data for the entire organization. ●

Wayne Eckerson is the director of TDWI Research and author of *Performance Dashboards: Measuring, Monitoring, and Managing Your Business* (John Wiley & Sons, 2005). He can be reached at weckerson@tdwi.org.

This article was excerpted from the full, 30-page report by the same name. You can download this and other TDWI Research free of charge at www.tdwi.org/research.

The report was sponsored by Business Objects, HP, Oracle/Hyperion, InetSoft, SAS, Sybase, Syncsort Incorporated, and Teradata Corporation.

University of Utah, University Health Care: BI Tips that Drive User Adoption and Increase ROI

Commentary by Jim Livingston

Director of IT Infrastructure and the Data Resource Center, University of Utah Health Care

University Health Care

HEADQUARTERS

Salt Lake City, UT, USA

CORE BUSINESS

Teaching hospital; primary healthcare and medical education.

OBJECTIVES

To serve the public by improving health and quality of life, to meet cost and efficiency targets, and to comply with all necessary regulations.

SOLUTION

FAST Radar

BENEFITS

End users such as physicians and nurses are empowered to find the data they need for decision making.

Increased adoption rates: eight times more users access data than before.

Colorful personalized charts, dashboards and scorecards make data access intuitive and actionable.

Real-time delivery of reports and personal alerts based on thresholds set by the user.

Unified architecture built on Microsoft.NET ensures rapid implementation and minimal maintenance.

Increased ROI now that end users can create their own BI reports.

Healthcare in the western states doesn't get much better than at University Health Care. The center, which manages all patient care services for the University of Utah, has been named one of America's best hospitals by *U.S. News & World Report*—for 12 years running. The center has also been listed as one of the nation's 100 most-improved hospitals for its gains in quality of care, operational efficiency, and financial performance.

With more than 850 physicians, the center cares for more than five million people in Utah and neighboring states and has an enviable reputation for academic excellence. Like almost all primary healthcare providers, though, University Health Care was under constant pressure to improve its quality of patient care, comply with a flurry of regulatory mandates, and ensure that it managed its assets and other resources as efficiently as possible.

Challenge

University Health Care had no shortage of data. Its IT teams had built robust data warehouses to house data for everything from accounts payable to procedures ordered in intensive care. IT had created a communications system that allowed data to be combined across functions—to extract financial data and combine it with clinical data, for example. However, the data was not easily usable. The hospital's business intelligence (BI) solutions were overly complicated for end users such as nurses and physicians. Two users seeking the same data could get very different answers. As such, adoption rates for the BI system were very low. The implications were clear not only for nurses anxious to avoid dangerous drug interactions but for surgeons responsible for improving

“The system has really improved our charge capture. We've found a million dollars of charges in our labor and delivery unit that we've been able to reallocate properly.”

Cheri Hunter, University Health Care

operating room utilization. “We were giving people a taste of what was possible, but we were not penetrating the organization with information,” says data warehouse manager Cheri Hunter.

Solution

University Health Care had to find a solution that would encourage non-expert users to tap into its wealth of data. “More than 70 percent of our users are medical students, nurses, physicians, and so on. Most of them will not open a system they cannot use,” says Cheri Hunter.

It quickly became clear that an easy-to-use data portal would be the answer—a single sign-on solution whose dashboards and other visual methods of data presentation would provide actionable information at a glance for data-savvy analysts and non-expert users alike. The solution had to allow decision makers to receive personally relevant information through familiar interfaces such as Microsoft Excel or via e-mail. It also had to be easy to integrate, easy to implement—and easy on the IT budget.

The hospital's IT leaders evaluated two solutions, including FAST Radar. Early proof-

of-concept projects soon confirmed what early reviews had indicated: FAST Radar was exactly what was needed. The final seal of approval? When the new BI system was demonstrated to end users, it won instant acceptance. “Everybody said, ‘This isn’t so bad—we can do this ourselves.’ Our chief nursing director even said, ‘If I can do it, anyone can do it,’” recalls Hunter.

Results

With FAST Radar on board, University Health Care managers are seeing significantly higher ROI numbers than were possible with the center’s incumbent technology. One example: today one full-time IT analyst creates reports for users, down from six previously. The new easy-to-use BI system has engaged a broader base of users—from facilities managers to nurses—which means they are making better-informed decisions and driving down costs in many departments.

FAST Radar has given the center’s professionals confidence in their data—as well as control. Its ease of use makes it simple for staff to build their own reports (no more spreadsheets needed). The Radar Write Back capability lets users update data directly to an operational database. Staff also has access to real-time data—critical for tasks such as utilization planning in Intensive Care, for instance. Physicians now use browser-based views of scorecards that track time utilization of facilities or schedule variances by site. Nurses now control what they are tasked to manage, ensuring, for example, that patient care orders are carried out in a timely manner and monitoring any alerts for potential adverse event.

“The biggest driver for us was to empower the end user. Now that we’ve done that, we’ve seen an explosion in usage,” says Director of IT Jim Livingston. Previously, the BI tool acquired in 1999 was being used by five report administrators and perhaps 80 occasional end users. Today, after just nine months, there are 400 active accounts on FAST Radar—two-thirds of them accessing the system 50 times each per month.

The new BI system also allows costs to be tracked more easily. In one case, a million dollars of charges in the labor and delivery unit have been properly reallocated. The BI tools have also helped staff to identify that anesthetic costs were not being billed correctly. “FAST has been a phenomenal success for us,” says Hunter.

“The biggest driver for us was to empower the end user. Now that we’ve done that, we’ve seen an explosion in usage.”

Jim Livingston, University Health Care

Strategic Next Steps

Given the success of BI within University Health Care, the IT group is now evaluating the strategic steps needed to apply more advanced technology for increased efficiency and effectiveness. Part of the team’s evaluation includes FAST AIW, a solution in which business intelligence is based on search. AIW provides two critical features with significant value to health organizations: advanced data cleansing capabilities, and the ability to

combine structured and unstructured data and creates associations between the two.

Like most Academic Health Centers, University of Utah Health Care has abundant information in text form but it cannot yet search the text fields and associate the results with the current BI initiatives. With the ability to do so, the center would have a 360-degree view of information in one place, saving time and providing users with more complete information. It will also enable users to uncover insights that otherwise would go undetected.

Why FAST Radar?

FAST Radar provides actionable information at a glance. Its robust browser-based dashboards and scorecards make it easy for non-expert users to get the facts they need in real-time—and to act on what they learn. It is easily scalable: with its open architecture, FAST Radar can scale up to match University Health Care’s needs, whether they involve research data for the teaching staff or patient tracking from the ER to intensive care. The solution needs minimal IT administration; built on the popular Microsoft .NET architecture, FAST Radar provides one integrated browser-based environment that requires very little ongoing maintenance. And it was very easy to implement: FAST Radar was installed in a matter of hours. ●

For a free white paper on this topic, [download “Business Intelligence Built on Search—The Adaptive Information Warehouse,”](#) or [click here](#) for more information about FAST.

Open Source BI: Does It Fit into Your BI Environment?

By John Williams

Senior Vice President, National Practice Director,
Collaborative Consulting

Open source software solutions are becoming more and more pervasive. Both large and small corporations have considered, and frequently adopted, open source software as a component of their overall software implementation strategy. Operating systems, data bases, source control, IDEs, and Web servers are some of the more common open source components.

It was inevitable that eventually open source would address data warehousing and business intelligence solutions. There are currently several open source tools in the market that address everything from ETL to BI. They claim to provide a complete capability equivalent to current market leading solutions. This brief article will address some of the pros, challenges, and considerations for open source as a part of an enterprise strategy.

First, let's start with the pros:

1. Open source is relatively inexpensive from a licensing perspective. While more traditional products cost hundreds to thousands per seat, open source software can be as inexpensive as free without support, and certainly much less than traditional products even when support is included.
2. There are open source solutions that address a majority of the software continuum for DW/BI. The current market seems to have evolved with multiple consortiums addressing each aspect of DW/BI. Some of the software vendors have aggregated these into end-to-end platforms.
3. Component-based design allows them to be retrofit into existing applications. The different subcomponents of the open

source solutions can be customized and embedded into existing or new applications.

4. Some features are more advanced than others. If your need is satisfied by these mature features, then open source may be a fit.

While the licensing costs are low, the TCO (for open source BI solutions) may not be.

Some of the challenges and considerations would include:

1. The current solutions tend to be development heavy. Many are JAVA-centric, and will require in-house development skills in order to make any real use of the tools.
2. While the licensing costs are low, the total cost of ownership (TCO) may not be. Given their development-centric nature, these tools will actually cost more for ongoing support than one might initially expect.
3. Open source tool evaluation is more time consuming than traditional tool evaluation. Even getting the vendors to provide demonstrations of their product can prove to be challenging. A free proof of concept is most likely out of the question.
4. The user interfaces are not nearly as robust as traditional software products. This is largely because these tools are still evolving, doing so at the pace of the consortium. The priorities at this point seem focused on core functionality, not advanced user interface design.

5. Complete is a relative term. In the case of open source solutions, key elements such as a semantic layer, dash boarding, data visualization, and data quality components are lacking, while other components have smaller functional gaps.

6. They are very desktop-centric. While traditional products have evolved to be heavily Web-centric, these tools are still very dependent on a rich client to provide key functionality.

The bottom line:

While open source BI tools have begun to show some promise and will force traditional software vendors to rethink their pricing and strategy, they are still a ways away from becoming enterprise solutions. In smaller organizations support will be a challenge, in larger organizations the TCO and functionality may not be there just yet. In considering open source, an organization should look at opportunities to try it out on a smaller scale, looking at specific data mart needs for a potential fit. Consider your organization's readiness carefully, and if you decide to adopt open source BI, be aware of the trade-offs. ●

LESSON FROM THE EXPERTS

Enterprise Business Intelligence: Making the Case and Keeping the Promise

By **John Williams**

Senior Vice President, National Practice Director,
Collaborative Consulting

It seems these days everyone is embarking on an enterprise business intelligence initiative. While the catalysts are varied, the basic mission statement is the same: Make information available across the enterprise in a consistent and easily accessible manner.

This seemingly innocuous and laudable mission statement is fraught with challenges and roadblocks. The first—and perhaps least obvious—challenge is answering the question, “Why?” Isn’t it obvious? Isn’t the goodness of correct, consistent, and readily available business information enough? Add the equal goodness of a more manageable technical environment with fewer packages and platforms to support, and you certainly have a compelling case, right?

Wrong. Transitioning from the typical departmental, application-centric, or fragmented BI environment of today is a nontrivial exercise involving resource costs, user tool changes, and (in some cases) business process changes. It requires executive involvement and buy-in, and in today’s age of closely scrutinized IT spending, the usual altruistic arguments just don’t wash. Even with buy-in, once the spending begins, the business will demand to see positive business impact in a reasonable timeframe or the funding will be quickly cut off.

Making the case for enterprise BI in your organization:

1. Align directly with business goals.

This sounds like an easy thing to accomplish; however, it is more of a challenge than one might think. It requires a focus on what the business is trying to achieve and alignment with those goals. BI is almost never the complete solution, but

rather an enabler, and must be positioned as such.

2. Don’t just settle for consolidation.

Many IT groups are in the process of consolidating core platforms such as enterprise resource planning and customer relationship management systems. The difference with BI systems is that users have heavily customized them to their own departmental or group needs. In order for enterprise BI to be successful, you need to consider the user needs in aggregate and show how the new solution will provide better capability than exists today, not just similar capability.

3. Investigate the reason for the current implementation.

Take a closer look at the reasons for multiple technologies and fragmented data, and you may discover things such as department-level funding, lack of corporate-level drivers, and separate lines of business as the root cause; these must be considered and addressed before an enterprise BI initiative can be successful.

4. Create alignment with key business organizations.

You need to ensure that the large, business-driving organizations’ needs are satisfied, even if it means compromising in some of the less core business groups. While not very egalitarian, this is practical and helps ensure that the initiative is adopted and funded by the business.

5. It’s not about one tool or environment; it is about the right tools and environment.

Many enterprise BI initiatives focus too heavily on an enterprise BI tool instead

In order for enterprise BI to be successful, you need to consider the user needs in aggregate and show how the new solution will provide better capability than exists today, not just similar capability.

of an enterprise BI solution. It is about an enterprise capability to deliver information in the right for efficiently and with consistency. While this may imply fewer tools, it does not imply just one or just one set.

6. Clearly identify cost savings.

Just saying that having fewer tools, environments, integrations, and platforms will reduce the IT workload is not enough to create a compelling business case. In the near term, this initiative will cost money, not save it.

In such a brief article, we can’t cover all of the bases; however, if you follow these fundamentals, you should be on your way to avoiding some of the pitfalls and solidifying your initiative’s opportunity for success. ●

For a free white paper on this topic, [download “Consolidating the Business Intelligence Infrastructure,”](#) or [click here](#) for more information about Collaborative Consulting.

CASE STUDY

Why Open Source Was the BI Suite of Choice for ZipRealty

Commentary by Salvatore Scalisi

Director of Data Administration, ZipRealty

Company Overview

ZipRealty (NASDAQ: ZIPL), a national, full-service real estate brokerage, provides home buyers and sellers with an innovative real estate solution. By using the efficiencies of the Internet, ZipRealty has streamlined the real estate process and is able to pass significant savings on to its clients. The brokerage's licensed, local ZipAgents in each of its markets allow home buyers and sellers to save thousands of dollars, without compromising on service.

The company employs 1,875 ZipAgents, serving a large client base across a growing list of diverse local real estate markets.

Business Challenge

With a large and expanding business, integrated online operations, and nationwide reach in a dynamic market, ZipRealty has long recognized the need and opportunity to integrate information from disparate sources to give business users visibility into strategic metrics. Important metrics span topics such as agent productivity, customer service, and overall return on investment of ZipRealty's marketing programs.

Technology Environment and Challenges

Critical business information for ZipRealty exists in multiple transactional systems. ZipRealty uses a combination of packaged as well as custom applications, which run primarily on an Oracle database platform. Getting a complete picture of business activity and performance across these systems requires integrating and rationalizing the data into standardized models and formats.

Prior to Pentaho, ZipRealty used an internally developed Java-based application to access these disparate data stores and to extract, transform, and load the data into a data warehouse running on the MySQL Enterprise

database. Using a custom-built solution allowed ZipRealty to deliver initial business value and results without expensive software licensing fees, while deploying an integration solution that was optimized and tailored for their specific environment.

Over time, systems changed as the business grew, and business-user analytical requirements evolved. Therefore, ZipRealty's IT team ran into challenges in maintaining the data integration environment. Making changes to business rules and integration logic became a complex and time-consuming process. It was important to ZipRealty's IT team to remain responsive to business needs, while simultaneously preserving precious IT resources for other important systems and projects.

“With professional support and world-class ETL from Pentaho, we've been able to simplify our IT environment and lower our costs. We were also surprised at how much faster Pentaho Data Integration was than our prior solution.”

Salvatore Scalisi, ZipRealty

Solution and Results

ZipRealty chose a Pentaho data integration subscription (based on Pentaho's popular Kettle open source data integration project) to address its needs. After evaluating the product, ZipRealty found that its metadata-driven architecture, graphical drag-and-drop maintenance environment, and rich library of transformations would fit well into its data warehousing environment, and would streamline and simplify maintenance.

Why Pentaho:

- Robust, graphical ETL design and maintenance environment
- Professional support and IP indemnification with no software license fees and no “vendor lock-in”
- Proven solution that was easily validated in ZipRealty's environment during the selection process
- Far lower overall costs than traditional, proprietary ETL
- Easily integrated into an environment that included a combination of open- and closed-source BI technologies

ZipRealty also was pleasantly surprised to find a three-fold increase in data loading performance, reducing its batch window by eight hours.

Furthermore, Pentaho's commercial open source model allowed it to deploy best-in-class open source ETL capabilities with professional telephone and electronic support, as well as IP indemnification with no software license fees and an overall TCO far below that of traditional, proprietary ETL tools. Demonstrating a pragmatic and successful mix of open- and closed-source BI technologies, ZipRealty continues to use MicroStrategy business intelligence tools on the front end to provide reporting and analytics to business users. ●

Pleasant Surprises: Commercial Open Source BI

By **Lance Walter**

Vice President of Marketing, Pentaho

Many organizations have already reaped the benefits of open source software in different forms—deploying low-cost, scalable Linux-based servers, or using open source databases such as MySQL or Postgres. You might be surprised how often you touch open source software in your day-to-day professional or personal life. Do you search on Google? Watch videos on YouTube? Have a TiVO at home? If so, then you're using open source software. These days, even cell phones and public transit vehicles are leveraging open source software.

Open source came to the BI market a few years back and has moved beyond the status of “market trend” or “the next big thing” to become a pragmatic choice for BI applications at a wide range of companies. Open source BI is still a new topic for many, and many “pleasant surprises” await those who invest a little bit of time to learn more.

Reality, not Hype

Organizations like yours use open source BI, whether you're a large or small company, and irrespective of your industry. Big companies such as Terra Industries, Motorola, and Unionfidi S.C., as well as small- and medium-sized companies such as iStockphoto, Boyne Resorts, and DivX use open source BI, just to name a few. Some have thousands of users or terabytes of data, while others are small departmental deployments.

“You Get What You Pay For”

No one has ever shown a positive correlation between the amount of money spent on BI software licenses and the ultimate benefits delivered by BI. One frequently cited cause of BI project failures is underbudgeting for critical services such as integration, training, and consulting, due to over-allocation of budget to software license fees.

“There's No Such Thing as a Free Lunch”

Even without proprietary license fees, business intelligence is not free. Maximizing the value of BI in your organization requires an investment of time—and almost always money (integration, training, design, hardware, maintenance, consulting). Just because you can avoid up-front license fees with open source BI doesn't mean you're getting a “free lunch.”

Reducing Open Source IP Risk by ... Using Open Source!

It's more likely than not that there is open source software on the CDs from your proprietary BI vendor. This can take the form of included application servers or Web servers such as Apache Tomcat, embedded databases such as MySQL, or embedded libraries. If you read the fine print in your license agreement or your vendor's SEC disclosures, it's likely that your vendor isn't indemnifying you. Most commercial open source BI companies include IP indemnification as part of their service agreements, which actually reduces any exposure you might have, based on use of open source software.

Transparency Rules

For years, IT organizations have asked for more transparency from their technology vendors. Open source has delivered a major leap forward in this area. In many cases, you can just go to a company's Web site and get full access to product documentation, unfiltered user feedback, product roadmaps, bug reports, and software, as well. You know

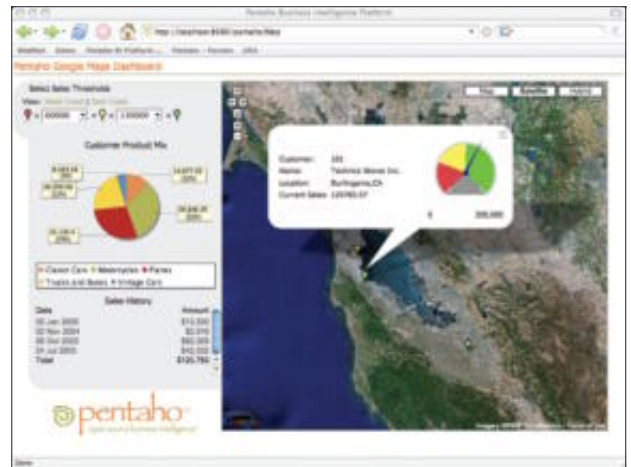


Figure 1. Open source BI is bringing new capabilities to the BI market.

what you're getting, and you know where your vendor is going.

Functional Evolution

Open source BI has come a long way over the years, and now is even providing features that aren't yet widely available from established, traditional proprietary vendors. Open source BI covers the main BI functional areas, such as reporting, analysis, dashboards, and data integration. It also offers thin-client AJAX technology, Web services access to everything, integration with the Eclipse IDE—things that many traditional vendors are still catching up on. ●

For a free white paper on this topic, [download "The Transformative Economics of Professional Open Source on Business Intelligence Technology Selection and Evaluation,"](#) or [click here](#) for more information about Pentaho.

CASE STUDY

Mortgage Company Drives Compliance with Data Governance

Commentary by Daniel Teachey

Director of Corporate Communications, DataFlux

The Business

A company that purchases, securitizes, and invests in the secondary mortgage market—having financed, since its inception, tens of millions of homes—faced a new round of regulatory pressures. Due to recent increases in scrutiny from regulatory agencies, it was relying on outmoded technology and processes to attempt to meet compliance requirements. This company saw an enterprise-wide data governance solution as the only real answer to these challenges.

The Challenge

Increased scrutiny and new regulatory requirements, such as the Sarbanes-Oxley (SOX) legislation, have only added to the compliance burden for mortgage companies. Like any financial institution, these organizations were already competing in a heavily regulated industry.

Faced with new regulatory challenges, the company found that it did not have either the technology or processes in place to implement a comprehensive data governance program. In the past, separate lines of business within the company were individually responsible for their own data quality. Each unit would work directly with the corporate IT division, creating data quality rules, which would then be implemented on each application or data source by IT.

This process had a historically slow turn-around time, with changes to the rules taking weeks or months to be implemented. Plus, the business analysts who depended on the data had very little control over it.

Seeking a system with faster implementation times that could produce more reliable data—and that would place control of data

quality in the hands of business analysts—the company turned to DataFlux.

The DataFlux Solution

The company chose DataFlux dfPower Studio and the DataFlux Integration Server to help it meet its compliance goals. By profiling their data using dfPower Studio, the company was able to seek out and correct erroneous or duplicate data to ensure that their data governance initiatives focused on the most troublesome data sources. Through statistical and numeric range analysis, business analysts verified that key metrics, such as loan-to-value ratios or total unpaid balances, fell within federally mandated acceptable ranges.

.....

With the stakes raised by Sarbanes-Oxley and other legislation, the company needed a more robust, reliable data governance system.

.....

Using dfPower Studio, the company's business analysts were able to work separately from the IT department to develop their own business rules for data governance. Once the initial data profiling had been performed, analysts were able to use dfPower Studio to transform the findings from the profiling into actionable business rules, then establish these rules as ongoing data monitoring routines and deploy them as real-time services using the DataFlux Integration Server.

The Results

The company is now much more confident about its quarterly and annual financial reports because it is confident that its

federally mandated reporting is built on high-quality data. Instead of multiple divisions having separate data quality rules, the company now has one source for data governance rules across all applications and data sources, resulting in a unified, accurate view of corporate data.

Putting data quality in the hands of the business users also allowed the company to reduce the time needed to implement data quality initiatives. Implementation times went from months to just a few hours. Because dfPower Studio uses an intuitive interface to build rules for data governance efforts, business users could act independently to analyze the data and establish real-time controls. With business users controlling data, the company gained greater flexibility to respond to changes and meet challenges as they arose.

By making data governance an ongoing part of its operations, this mortgage company is certain that it has the people, processes, and technology in place to be compliant with SOX. Furthermore, because of the adaptability of the DataFlux solution, the company also has the confidence that it will be able to meet any changing regulatory obligations in the future. ●

LESSON FROM THE EXPERTS

Getting Started with Data Governance

By **Daniel Teachey**

Director of Corporate Communications, DataFlux

As companies collect more information about their customers, products, suppliers, inventory, and finances, it becomes more difficult to accurately maintain that information in a usable, logical framework. This can severely complicate regulatory efforts, because the information within applications and databases—data pertaining to customers, products, employees, suppliers, and financial transactions—provides the foundation for audit reports and various other compliance efforts.

The data management challenges facing today's businesses stem from the way that IT systems have evolved. Enterprise data is frequently held in disparate applications across multiple departments and geographies. To address the spread of data, many companies implement enterprisewide data governance programs, which attempt to codify and enforce best practices for data management across the organization.

Data governance encompasses the people, processes, and technology that are required to create a consistent enterprise view of a company's data. By concentrating on the health of the data, organizations can address the lifeblood of their enterprise, helping create better data to support any business initiative.

The Path to Data Governance

Like many enterprise initiatives, data governance programs often start small before finding the sponsorship and support needed to transcend organizational boundaries.

Through an established data governance maturity model, organizations can identify and quantify precisely where they are—and where they can go—to create an environment that delivers and sustains high-quality information.

The data governance maturity model looks at the technology being utilized, along with the people and policies associated with the governance initiative, to ascertain the level

of data governance sophistication within that enterprise. In the first stage, Undisciplined, an organization has few defined rules and policies regarding data quality and data integration. The same data may exist in multiple applications, and redundant data is often found in different sources, formats, and records.

The danger for Undisciplined companies is the real and constant threat that the underlying data will lead to bad business decisions that may, in turn, lead to security or compliance violations. Often, a cataclysmic failure (i.e., failed compliance audit, drastic decrease in customer satisfaction) shakes the organization out of complacency and leads to further progress.

.....

Data governance is a way to address the challenges posed by an ever-changing IT environment.

.....

At the next phase, the Reactive stage, a company begins to organize a data governance program, either through grassroots efforts or, more likely, through an executive-driven effort fueled by an earlier failure. At the Reactive stage, organizations try to reconcile the effects of inconsistent, inaccurate, or unreliable data as bad records are identified. Here, the gains are often seen on a departmental or divisional level, but the company is starting to establish some best practices for data governance.

At the Proactive stage, the data governance program becomes cross-functional and has explicit executive support. To build a single view of a customer, for example, every part of the organization—sales, marketing, shipping, finance—has to agree on what attributes make up a customer record.

The final phase, the Governed stage, is where data is unified across data sources according to business rules established by an enterprise data governance team. At this final stage of the maturity model, a company has achieved a sophisticated data strategy and framework, and a major culture shift has occurred. Instead of treating issues of data quality and data integration as a series of tactical projects, these companies have a comprehensive program that elevates the process of managing business-critical data.

Although individual applications are still in use by a Governed company, the data that they access comes from a single repository that is propagated across the IT infrastructure. This provides the ultimate in control for the enterprise, as all reports and dashboards pull from the same pool of information. ●

For a free white paper on this topic, [download "The Data Governance Maturity Model: Establishing the People, Policies and Technology That Manage Enterprise Data,"](#) or [click here](#) for more information about DataFlux.

Business Intelligence Solutions for SAP

Issues that Influence BI Solutions for SAP

BY PHILIP RUSSOM, SENIOR MANAGER, TDWI RESEARCH

The issues discussed here quantify the state of BI solutions for SAP; they also describe the issues you need to consider before deciding which mix of tools, platforms, approaches, professional services, and business sponsorship are appropriate to your solution.

Applications and Data Warehouse Architecture Issues

Most of the data that populates data warehouses and reports comes from various types of operational applications, including those for ERP, CRM, and SCM. Because these are prominent sources of data that BI is based on, the number and diversity of applications strongly influences the data content and architecture of any BI solution, especially those that involve SAP.

- **SAP users don't only use SAP applications.** TDWI's survey asked, "From which vendors has your organization acquired packaged applications for ERP and other operations?" The survey population is dominated by SAP users, so it's not surprising that 71% of respondents report using SAP applications. (See Figure 1.) But many reported also acquiring applications from Oracle (25%), Peoplesoft (24%), Siebel (15%), JD Edwards (12%), and so on. Given the numerous vendors selected by survey respondents—plus home-grown solutions—it's safe to say that most SAP-centric businesses also have significant non-SAP applications.

- **The SAP-centricity of your BI solution is a matter of degree.** In other words, a BI solution for SAP may incorporate SAP applications data only, or it may also integrate data from other applications (and even non-application sources like syndicated data). The latter is necessary for a complete view of organizational performance. In fact, gaining a centralized "single version of the truth" is one of the reasons organizations deploy enterprise data warehouses. The challenge is to satisfy both application-specific and enterprise-scope BI requirements.
- **Application diversity may lead to BI solution diversity.** This may mean deploying both NetWeaver BI and an enterprise data warehouse (EDW). In most cases, NetWeaver BI supports data marts and reports specific to businesses and processes supported by SAP applications, whereas the EDW collates information from many sources for strategic decision making. To reduce the number of solutions, some SAP-centric organizations have made NetWeaver BI the equivalent of an EDW, while others rely on an EDW without any SAP BI products. Selecting one of these data warehouse architectures—or a combination of them—is a tough decision for SAP users, because it affects the cost of BI, completeness of data content, data ownership, and how BI goals are prioritized (i.e., tactical operations versus strategic planning).

FROM WHICH VENDORS HAS YOUR ORGANIZATION ACQUIRED PACKAGED APPLICATIONS FOR ERP AND OTHER OPERATIONS? (SELECT ALL THAT APPLY.)




Figure 1. Based on 476 responses from 278 respondents. Average responses per respondent: 1.7

IS THE BI ENVIRONMENT FOR SAP THE SAME AS YOUR ENTERPRISE BI ENVIRONMENT?




Figure 2. Based on 172 respondents who are SAP users.

WHO'S THE PRIMARY SPONSOR FOR YOUR BI SOLUTION FOR SAP?




Figure 3. Based on 172 respondents who are SAP users.

WHAT STRATEGY HAS YOUR ORGANIZATION FOLLOWED IN ACQUIRING ERP APPLICATIONS?




Figure 4. Based on 278 respondents who use applications from SAP and other vendors.

Organizational Issues

TDWI asked of SAP users: "Is the BI environment for SAP the same as your enterprise BI environment?" (See Figure 2.)

- **Enterprise-scope BI.** Over half (56%) said "yes," which suggests that these respondents' BI infrastructure is enterprise in scope—that is, it spans the data and BI needs of many applications and business units, not just those associated with SAP applications.
- **Application-scope BI.** Roughly one-third (37%) said "no," suggesting these users have a BI solution that is narrowly focused on the data and business processes of SAP applications. Of course, this does not preclude them from also having a separate BI solution of enterprise scope.
- **Determining scope is an important BI design decision.** A BI solution for SAP may be enterprise scope, with SAP-specific requirements satisfied by a subset of the solution. Or it may be application scope, possibly complemented by an enterprise-scope BI solution.

The scope of a BI solution for SAP is related to its ownership and sponsorship. (See Figure 3.)

- **Line-of-business managers (44%) commonly sponsor BI solutions for SAP.** Line-of-business managers often fund, own, and sponsor operational applications, so it's possible that they have extended this role to also sponsor application-specific BI. Even so, IT and BI directors sponsor BI solutions for SAP just as often (45%).
- **SAP BI sponsorship can also be shared.** Several survey respondents selected "Other" and entered comments like "we use a team approach, no single sponsor." Sponsorship is "shared by all businesses" where it involves "several stakeholders." Or BI may be sponsored by "corporate leadership as part of the SAP implementation."
- **Finance departments often sponsor BI, whether for SAP or not.** We're reminded of this by survey respondents who listed their finance organization or its CFO, VPs, and directors.

WHICH BEST DESCRIBES HOW YOU BUILT YOUR BI SOLUTION FOR SAP?




Figure 5. Based on 172 respondents who are SAP users.

Build versus Buy Issues

When it comes to ERP applications and the classic build-versus-buy decision, buying is the norm by a long shot, with most survey respondents buying and customizing (57%), followed by those who buy but don't customize much (27%). (See Figure 4.) By comparison, building an ERP system is rare (9%). This is natural, given the complexity of ERP processes and how difficult it would be to build and maintain an application to automate them. Plus, multiple software vendors offer mature ERP systems that are feature-rich, robust, and have solid track records.

To quantify the build-versus-buy approaches of BI solutions for SAP, TDWI's survey first identified respondents who are SAP users, then asked them: "Which best describes how you built your BI solution for SAP?" (See Figure 5.) A whopping 75% reported acquiring BI products from SAP, which is a high penetration of BI products for an applications vendor.

Note that the level of customization varies with SAP BI products. Relatively few use SAP BI products "as is" (10%), while most assign in-house personnel (35%) or hire consultants (30%) to customize them. Customization has ramifications for packaged solutions:

- **BI solutions are hard to package.** That's because an organization's BI data sources, its data models, and its report presentations all vary tremendously. However, if you narrow the scope of the BI solution to specific use cases, which SAP NetWeaver BI does by focusing on standard business processes, the breadth of user requirements is reduced, thereby reducing the difficulties of packaging a solution that works for multiple user organizations. But BI requirements still vary somewhat, even when focused on ERP, so organizations inevitably customize to some degree.

- **The greater the customization, the greater the difficulty of upgrading.** Some SAP customers would like to upgrade from older BW releases to recent NetWeaver BI ones, but they feel that customization makes the upgrade too difficult to attempt. Let's all recognize that customization is inevitable with BI products, so upgrades usually involve development work.
- **Avoiding customization reduces project time and cost.** A tried-and-true strategy for a packaged solution is to roll out the first phase with little or no customization, then tailor the solution to organizational requirements in subsequent phases.

"We started our implementation of SAP BW in 2002, working in parallel with the firm's implementation of the SAP ERP," said Brian Hickie, former VP of Business Intelligence at McKesson Pharmaceutical. "We only had nine months to implement the data warehouse, along with considerable ETL jobs and a number of initial analytics and reports. We made the deadline despite the fact that the SAP ERP configuration precluded the use of standard business content or standard data extractors; the ETL jobs and BW required a considerable amount of configuration to provide the data to generate analytics and reports from the ERP and legacy systems."

Developing home-grown BI solutions isn't as common as other approaches in Figure 6 (16%). However, other survey questions showed higher rates, and survey respondents who selected "Other" entered responses that involve home-grown solutions:

- **SAP BW and EDWs can coexist peacefully.** Some BI solutions for SAP involve both NetWeaver BI (specifically the BW component) and an independent enterprise data warehouse (EDW). For example, one survey respondent said: "We both use BW and

extract SAP data into a data warehouse.” Another reported using “BW and an in-house solution based on Oracle and Business Objects,” while yet another said, “[We] use SAP BI as an OLAP tool to access a Teradata EDW.”

- **Some BI solutions for SAP are a work in progress.** As one survey respondent put it, “We’re in transition from a legacy custom-developed EDW to SAP BI.”

Issues in Software Acquisition Strategies

TDWI asked of SAP users taking this report’s survey: “What was the primary reason for your organization’s implementation of a BI solution for SAP?” In the survey, the question had no pre-written answers to select; instead, each respondent typed an answer in his/her own words. Many respondents turned the question into a soapbox, because they have strong feelings either for or against SAP as a preferred software supplier and how that predetermines the products they use.

- **For many organizations, naming SAP a priority supplier leads inevitably to SAP BI products.** Some respondents stated this in matter-of-fact answers such as: “SAP is our preferred provider in all areas.” Others made a direct causal link: “Using SAP ERP was the main driver to start with an SAP BW solution” and “SAP is selected as our common platform, which includes BW.”
- **Some spoke of the positive leverage between NetWeaver BI and SAP ERP.** “BW is a good fit on top of SAP ERP and SAP CRM,” enabling users “to leverage prebuilt content and integration with the ERP.” According to other survey respondents: “Most of our [BI] information is sourced from SAP components,” and “BW provides a direct link to primary source data.”
- **A few people complained about the SAP-centric strategy.** One claimed that using SAP BI products is due to a “strategic decision to use SAP overall, even when it does not fit.” Others received SAP BI products due to management actions beyond their control: “It [NetWeaver BI] was paid for in an overall package of software” and “parent company purchased SAP solutions.” Selecting a preferred supplier is a good procurement practice, in general, but it’s frustrating for

IT professionals when the strategy doesn’t consider user requirements.

- **Most mentioned the usual benefits of BI as reasons for implementing a BI solution.** Several listed common BI goals like better decision making or improved corporate performance. Others cited specific BI needs such as financial reporting, datasets for business analysts, or operational reporting based on SAP ERP source data.

“All of my clients have multiple application modules from SAP, and for most of them, that’s a good enough reason to get their BI technology from SAP,” said a consultant who specializes in BI solutions for SAP. “Why bother with an independent data warehouse? BW infrastructure—especially with recent NetWeaver releases—is almost identical to what you’d build into an independent warehouse. So the common practice I see is to leverage what you already have, instead of reinventing the BI wheel. Also, my clients trust SAP as a large, stable, non-acquirable vendor. In terms of best of breed, other BI and data integration tools are better, but SAP’s offering is relatively complete and pretty good. And the direct integration between BI and ERP is hard to beat.”

“We have a mature, best-of-breed technology stack for data warehousing and business intelligence, but it’s currently at risk because our company’s being acquired by a company with a deep commitment to SAP applications and BI,” said an enterprise data architect. “We’re under pressure to abandon our EDW and go with SAP NetWeaver BI, but we’re not sure it will meet our requirements for strategic and financial reporting, much less scale up. The merger is barely under way, and we’ve just started investigating NetWeaver, so it’ll be a long time before we know which way to go.”

Philip Russom is the senior manager of TDWI Research and oversees many of TDWI’s research-oriented publications, services, and events. He can be reached at prussom@tdwi.org.

This article was excerpted from the full, 32-page report by the same name. You can download this and other TDWI Research free of charge at www.tdwi.org/research.

The report was sponsored by Business Objects, Cognos, HP, IBM, Microsoft, MicroStrategy, SAP, and Teradata Corporation.

CASE STUDY

Kelley Blue Book Improves Data Analysis and Reporting with Business Intelligence Solution

Commentary by Bruce Hoffman

Manager of Business Intelligence and Data
Management, Kelley Blue Book

Kelley Blue Book analyzes and publishes data on new and used vehicles. Large automotive manufacturers, as well as individual car and truck buyers, use the Kelley Blue Book Web site to find in-depth pricing information. The company gathers data from Web site visits to create trend reports for manufacturers. However, the report-generation process was time-consuming, and employees were not able to effectively analyze all the data.

In 2006, Kelley Blue Book implemented a solution based on Microsoft business intelligence technologies. With its new analytical capabilities, Kelley Blue Book has improved sales staff efficiency and created a compelling new Web analytics product that can provide its business customers with more targeted data for their marketing efforts. The company expects to see its customer base and revenues grow in response.

Situation

Kelley Blue Book provides vehicle buyers and sellers with comprehensive information on new and used cars and trucks. The company, which currently has 400 employees, publishes up-to-date vehicle pricing and values through its Web site, kbb.com. "We receive about 12 million unique visitors each month on the site," says Bruce Hoffman, manager of business intelligence and data management, Kelley Blue Book.

The company uses specific information gathered from these online visits to generate reports that it presents to automotive manufacturers. However, Kelley Blue Book had struggled to put the information into a form

its customers could easily use to predict buyer trends. Even though that data had the potential to be useful, there was too much of it for manufacturers to make sense of. "It was impossible to tell a meaningful story with the data," says Hoffman. "Our marketing team had all this great information, but they didn't have a way to package and sell it. They lacked a good software tool to make that happen."

Solution

In January 2006, a Kelley Blue Book employee alerted Hoffman to Microsoft Business Intelligence, a set of fully integrated business intelligence technologies that can help reduce the complexity of organizing and distributing information. The new solution, which the company implemented that same month, gives the Kelley Blue Book analytics team the ability to quickly and easily collect targeted data on automotive sales and trends. It includes Microsoft SQL Server 2005, Microsoft Office Excel 2003, and Office PerformancePoint Server 2007 business intelligence software.

Now, Kelley Blue Book employees can import more than 100,000 rows of Microsoft Office Excel 2003 data from the Web site into a SQL Server 2005 database each day. From there, SQL Server 2005 Integration Services moves that data into an online analytical processing (OLAP) cube, which is an arrangement of data in arrays to allow for quick analysis. The cube is viewed using PerformancePoint Server 2007 and is delivered to sales representatives via a dashboard server.



Using these business intelligence products, the analytics team can analyze data about shoppers who are looking for new vehicles.

Benefits

Kelley Blue Book now has a business intelligence solution that integrates fully with its existing IT environment, an advantage that helped the company deploy a new Web analytics sales report product, called Market Watch, in a short amount of time. With the solution, Kelley Blue Book has reduced the time it takes sales staff to create reports, has gained more in-depth analytical capabilities, and can better assist business customers with targeting their own marketing efforts. Also, because it expects its customer base to grow, the company anticipates its revenues to rise. Other advantages include: full integration with existing environment; short deployment time for new sales product; increased efficiency for sales staff; improved analytical capabilities; and more targeted marketing for business customers

Also, because Kelley Blue Book is better able to help its customers with such information, it anticipates that its base of customers using Market Watch will grow. "Our research revenues are expected to double in 2007, compared to 2006. We saw we had a unique revenue opportunity, and this solution allowed us to address that," says Hoffman. ●

LESSON FROM THE EXPERTS

The BI Tipping Point

By **Bill Baker**

Distinguished Engineer, Microsoft

I believe business intelligence (BI) is at the tipping point right now. Quantitative evidence abounds. A recent report from AMR forecasts software sales of \$6.6 billion in the general BI tools segment in North America (NA) this year. They further forecast \$5.5 billion in NA software sales for dashboards and scorecards and \$4.1 billion in NA sales for planning, budgeting, and forecasting solutions. Add in forecasts for hardware and services, and we're looking at a very substantial market for BI.

Other industry and financial analysts are providing robust forecasts for the BI market. Enthusiasm for the BI market has driven vendor actions including acquisitions, consolidation, new entrants, and new products. Competition is heating up, and that's good for customers.

That nine percent growth can come from a number of mechanisms. Net-new customers may be buying their first BI platforms and solutions. Vendors may be increasing prices. Companies with existing BI implementations may be deploying additional solutions or pushing existing solutions to new users. Freeform Dynamics and The Register recently published survey results from more than 1,100 users in 250 companies on the subject of making business performance and management information more broadly available within their organizations. Twenty percent of respondents reported that they already provide information to a broad population of users.

Despite all the numeric evidence, my strongest sense about the tipping point is the changing conversation with customers. Recently I spent a few hours reviewing all of the BI slide decks I've delivered over the 10-plus years Microsoft has been working in BI. In 1997, every presentation began with the basics. I also found plenty of slides that simply described BI as the tools and techniques for making better decisions.

One major takeaway was that our audience was very technical and at least somewhat skeptical. BI was new enough that many IT departments were having their initial BI conversations. Frankly, there were so many vendors and no end-to-end vendors that the entire segment seemed very noisy, and the technology was often very different from the day-to-day technology of the time. Ten years ago, BI was a specialty industry.

Customers are asking how to take BI to their “masses.” Microsoft’s vision has always been that pushing BI deeper and wider into companies, the so-called “BI for the masses,” results in better companies.

I meet hundreds of customers every year. I have not defined BI for a customer or prospect in several years. The whole tenor of the conversation has changed. I have more discussions that are exclusively about business. But more significantly, the tenor of the technical conversations is different too. First, customers come prepared to discuss BI, not to learn about BI. The questions I get are deeper, more thoughtful, and sometimes more challenging. The question I get most is: How do I get BI into vastly more employees’ hands?

Recently I have had deep conversations about tying scorecards to strategy maps.

Scorecards are a prime mechanism to drive BI deeper into companies as they provide both motivation and context for further analysis. By showing interest in SharePoint, BI customers further indicate interest in providing data and tools to broader populations. SharePoint is a key element, in the Microsoft stack at least, for evolving individual insights into team insights. In conversation after conversation, customers are asking how to take BI to their “masses.” Microsoft’s vision has always been that pushing BI deeper and wider into companies, the so-called “BI for the masses,” results in better companies. Quantitative market research and customers are telling me this dream is coming true. I believe that BI is well along its evolution from specialty tools to mainstream adoption. In other words, we are at the tipping point in BI adoption. ●

For a free white paper on this topic, [download “How the Intelligent Enterprise Delivers Performance Management,”](#) or [click here](#) for more information about Microsoft Corporation.

CASE STUDY

Operational Performance Reporting Applications Help MFS Improve Regulatory Compliance, Enhance Productivity, and Expand Global Operations

Commentary by Joe Piotrowski

Vice President and Development Manager, MFS

About MFS

MFS is a global asset management firm that serves investors around the world. The firm's global research team covers the world through offices in five international cities—Boston, London, Mexico City, Singapore, and Tokyo. MFS invests in securities in more than 60 countries.

THE MFS CHALLENGE

As MFS expanded its global operations, the company needed to find an easier way to deploy research without having to install and maintain applications in remote offices. MFS had integrated Crystal Reports with its desktop application, but Crystal did not offer a Web-based reporting solution at that time.

MFS did not have a centralized reporting function and had multiple reporting technologies throughout its organization. Employees manually created ad hoc reports, a process not only time-consuming but also prone to user errors and version control issues.

THE MFS SOLUTION

MFS selected Actuate for operational performance reporting based on the following criteria:

- Proven scalability
- Web-based solution
- Support for VB COM and distributed COM components

- Robust development and scripting environment

The biggest Actuate users at MFS are the equity, fixed income, fund treasury, compliance, and HR groups. According to Joe Piotrowski, vice president and development manager at MFS, "If you are not using Actuate in our key market groups, you can't do your job. Throughout MFS, all of our report development is being done with Actuate."

MFS has leveraged Actuate's open, flexible architecture to create innovative functionality for its operational performance reporting application. First, MFS has leveraged the Actuate application program interface (API) to create conditional scheduling that automatically generates and schedules reports when new data becomes available. The reports are accessible through the Web portal.

Second, MFS has developed a method to automatically generate a personalized criteria page for users. MFS leverages the Actuate API to comb through the metadata, gather the attributes of the user, and then dynamically create a personalized Web front end.

MFS utilized Actuate e.Services and three internal MFS developers to complete the implementation in less than six months. Since MFS executives support Actuate and the users endorse the product, use of Actuate is spreading throughout the company to approximately 800 users.

Benefits of Actuate

ENHANCED PRODUCTIVITY

Employee productivity has increased because of the time saved by using Actuate to automatically generate and deploy reports. "This automated process ensures important reports are delivered to management on a timely basis, saving MFS managers about one-half hour per day," states Piotrowski. "They now have more time to focus on strategic initiatives and decision making."

ENABLED GLOBAL EXPANSION

Actuate has helped MFS to expand its investment operations worldwide. Because Actuate is Web-based and secure,

investment professionals can access reports from any browser worldwide as long as they have a security ID card. Actuate has cut down on the amount of support required to expand its global operations.

IMPROVED REGULATORY COMPLIANCE

As a financial services company, MFS has to comply with increased federal and state regulations, such as Sarbanes-Oxley. According to David Domenicone, systems architect at MFS, "We have ever-increasing audit requirements, and Actuate reports have been key in meeting many of them." Actuate provides MFS a means to create accurate audit reports to monitor and comply with regulations.

CONTROLLED DEVELOPMENT COSTS

Because of Actuate's component object model (COM) support, MFS was able to leverage existing code to run reports making it an easy migration from its desktop application to a Web platform. Code written for the desktop was easily deployed to both Web and Actuate servers.

MFS can now provide reporting capabilities cost effectively to anyone in the company. There is no incremental cost to bring Actuate to an additional department apart from setting up new data in tables. MFS no longer struggles to provide Actuate to smaller departments.

INCREASED EMPLOYEE SATISFACTION

MFS employees have been happy with the Actuate technology, and usage has increased throughout the company, with 100 percent user adoption. Actuate provides a consistent set of reporting interfaces for investment areas to enhance usability.

THE MFS FUTURE

MFS will continue its migration from Crystal Reports to Actuate and is evaluating how to integrate Actuate with its PeopleSoft application.

"Actuate has enabled MFS to enhance productivity, enable global expansion, and comply with regulations," says Piotrowski. "MFS will continue to leverage Actuate and expand its usage throughout our organization." ●

LESSON FROM THE EXPERTS

Spreadsheets in the Next Generation

By **Cate Zovod**

Senior Technical Marketing Manager, Actuate

Spreadsheets are a fundamental facilitator of the way in which many financial services institutions seek to conduct their businesses in the 21st century. There is now an acceptance industrywide that spreadsheets can pose a substantial risk to the success of the business, both in terms of pure profit and loss and reputation.

However, the spreadsheet is ingrained in the fabric of the financial services industry, making it too valuable a tool to eliminate. Therefore, banks have focused their attention on how best to mitigate these risks and leverage control across spreadsheets.

In order to identify best practices that tackle the risks, Lepus Management Consultancy interviewed practitioners across a range of tier one and tier two investment banks in order to gain an insight into current industry practices. Lepus looked at how spreadsheets were used across the front, middle, and back offices throughout these organizations.

The best practices deliver live business data within a familiar spreadsheet format, with all the formatting, formulae, and graphs intact and without querying, cutting, pasting, or manual data entry errors. These practices result in instant complete information direct from source, preformatted and ready to go, ensuring accuracy, quality, and control.

Achieving these benefits can be a reality if organizations are able to follow a series of steps that utilize the existing skills of developers but introduce automation to ensure the robustness of the process. These steps are outlined below:

1. Generate automated, server-based spreadsheets that enable complete version control, eliminate human error, and ensure that the development process is auditable for compliance purposes.
2. Determine a blueprint of what an individual spreadsheet looks like and what functions it must incorporate, thus eliminating spreadsheet development errors and validating the data that exists within the spreadsheet.
3. Distribute scheduled, manufactured spreadsheets to eliminate spreadsheet changes that can occur in the dissemination process, to generate confidence in data quality, and to ensure reports and analyses are received in real time. The timely distribution of spreadsheets also ensures that business processes are not delayed by inefficient distribution procedures.
4. Centralize spreadsheet processes and environments to drive accountability and consistency in the level of control leveraged across the business. By leveraging a centralized server and page-level security, institutions are able to deliver spreadsheets to a defined group of users, however small or large, in real time.
5. Spreadsheets that access all the necessary data sources provide broader insight and better-integrated information to support decision making.
6. Management of the authorization to create spreadsheets, responsibility for signing off spreadsheets and facilitation of ongoing monitoring in a live environment are crucial to gaining control. The resulting effect is enhanced transparency in spreadsheet use.
7. Regulatory pressures, such as Sarbanes-Oxley, now make it compulsory for

organizations to take the appropriate steps to mitigate the risks inherent in the reporting and analysis performed in spreadsheets. In order to ensure compliance and accommodate the increasing regulatory scrutiny of spreadsheets in the industry, banks must standardize and rigorously manage the development process.

When spreadsheet processes remain manual and unstructured, the risks in spreadsheet use remain high. The automation fundamentally redefines the concept of spreadsheet risk. Where a repeatable development and distribution process is in place, institutions can be confident that they have gained control of, and confidence in, their spreadsheet environments. Adopting these steps will help drive compliance with the increasing regulatory scrutiny of spreadsheet use and strengthen critical business processes. ●

For a free white paper on this topic, [download "A Paradigm Shift: The Management of Spreadsheet Use in Financial Services,"](#) or [click here](#) for more information about Actuate.

CASE STUDY

Rohm and Haas: One Company, One Measure

Commentary by Mike Masciandaro

Director, Business Intelligence, Rohm and Haas Company

The Company

Rohm and Haas serves a broad segment of dynamic markets, including building and construction, electronics, food and retail, industrial processes, packaging, transportation, and water. We have operations with approximately 100 manufacturing and 35 research facilities in 27 countries with 16,500 employees.

Rohm and Haas has relied on its business intelligence program to guide and track its transformation—in eight years—from a \$3.5 billion organization to an \$8+ billion enterprise. Each year, Rohm and Haas purchases 3.5 billion pounds of raw materials, mostly petrochemicals, and 23 million BTUs of natural gas—commodities with dramatic price increases and frequently subjected to disruption and shortages. During this transformative period, the company divested more than 45 units and acquired many more. We maintained profits by making good procurement decisions, by making timely pricing decisions, and by closely watching our businesses in near real time.

None of this would have been possible without our BI program.

Our executive council—seven senior executives—visibly supported this project, was actively involved in developing the first metrics, and continues to support and use the program. We were able to provide BI functionality to the CEO and people on the shop floor with the same schema, data, tools, and metrics.

Because we operate so many businesses in so many countries, it was important that we develop a “one company, one measure” discipline in order to minimize time spent discussing the meaning of data, and to foster

agility. We can move across the hall or the world and still track business with the same tools and semantics.

At the project’s outset, we had more than 300 disparate IT systems, with IT costs accelerating at an alarming pace. Our 13 businesses, each operating in four separate geographic locations, were operating largely independently of one another.

Our challenge was to remain geographically diverse and have extraordinarily efficient and robust business processes while being flexible enough to meet local customers’ needs and expectations. This was the perfect environment for IT, finance, and business collaboration.

“We can move across town or across the world and still track our business with the same tools and terminology.”

Mike Masciandaro, Rohm and Haas

The Solution

The tools we chose, an all-SAP implementation, inherently solved some of the linkage issues between strategy and implementation because the enterprise resource planning and BI facilities are tightly bound. If we are agile in ERP, we inherit this agility in BI. We frequently change reporting hierarchies—a fluid process that requires little to no attention. Because our business is dynamic, the tight connection between our ERP and BI systems eliminates handoffs that could adversely affect our goals for accuracy and agility.

The selection of SAP tools let us quickly roll out BI capabilities to a wide audience because there was one tool for both training and support. Upgrades are synchronized and tested by one vendor, avoiding the integration complexities with tools from multiple vendors.

We had a data warehouse before, but we saw clearly that it could be improved with a more complete architecture. The advantages of moving to SAP NetWeaver BI were too obvious to ignore. About 90 percent of our data is now in SAP, so it was an easy decision.

Once we implemented NetWeaver BI, we used BEx Web for reporting and dashboards to drive usage at higher levels of the organization. Everyone now gets a direct information pipeline.

Business Benefits

The executive council describes our BI program with words like “sea change,” “pricing power,” and “confidence,” and cites dramatic increases in the speed and quality of decision making.

Business performance accountability accelerated our decision making involving critical business risks. Enterprisewide objectives successfully translated into actions at the business level. Raw material cost increases now are tracked in near real time. Gross profit is protected by better pricing decisions, materials sourcing, and timing. Managers now track results from the first day of the month, rather than 40-50 days late, and finance and IT costs have dropped from 5 percent of sales to 2.75 percent.

We have moved Rohm and Haas into an environment in which we leverage information for faster and more accurate decision making. Our high-efficiency operating platform is now the foundation for innovative worldwide growth of our businesses, organically and through acquisition. ●

For a free white paper on this topic, [download “Beyond the Basics: Accelerating BI Maturity,”](#) or [click here](#) for more information about SAP.

LESSON FROM THE EXPERTS

Corporate Performance Management: Departmental or Corporate?

By **Ryan Leask**

Product Manager, SAP CPM

Sooner or later, everyone faces the same conflict when deploying corporate performance management (CPM) applications: how to balance corporate needs against departmental requirements. In most organizations, every department wants its own individually customized solution (especially for budgeting); yet this design inherently leads to multiple versions of the truth, challenges in combining the figures at a corporate level, and total cost of ownership (TCO) problems (especially when systems from multiple vendors are used).

We see organizations generally fit into one of three situations:

- They are stuck in spreadsheet hell, with no sophisticated CPM systems in place (or only deployed to a limited user base). All the problems of security, compliance, user error, reconciliation, audit, maintenance, and all the other well-known spreadsheet issues are being dealt with.
- Each department has implemented its own CPM solution, but the IT department is driving to standardize on one vendor. Each individual department is generally happy with its solution, but to have finance/corporate reconciling all the different formats causes headaches, and the IT department is left supporting a high TCO environment.
- Corporate finance has rolled out an enterprisewide CPM system that requires departments to conform to a standard solution. Each individual department loses the flexibility needed to fulfill its unique requirements and is, therefore,

generally unhappy, and resorts to keeping spreadsheets. Corporate finance and IT are generally happy, though, because the system provides a consistent solution that can be easily and cost-effectively maintained.

So how can this situation be resolved in a way that meets the detailed needs of the individual business, satisfies the corporate requirements for visibility and consistency, and keeps the TCO low? We are generally seeing the market answer this question on its own, today: vendor standardization.

Vendor standardization is happening at a rapid pace. What change has triggered this? Perhaps the single biggest factor is the vendor consolidation taking place in the CPM market, forcing organizations to look at what systems and/or vendors they are now dealing with. Typically, in the past, each department pursued the best-of-breed solution that best met its requirements, while IT departments pursued solutions that integrated with the corporate environment. With the vendor consolidation in the market, some vendors, better known for their integration capabilities, are now in possession of what was once called a best-of-breed application. What is occurring, therefore, is the replacement of isolated departmental solutions with offerings from a single vendor that touts both best of breed and tightly integrated solutions.

It is important to clarify here what “tightly integrated” means exactly, as integration needs to take place on two fronts. The first is a technical integration (see Diagram 1), where the CPM solution needs to interact seamlessly with a customer’s ERP systems, where the core financials and transactions are held; data warehousing and BI tools, used for reporting and analysis; and GRC (governance, risk, and compliance) suite, ensuring compliance and built-in risk sensitivity.

The second aspect is business process integration. It’s not enough for just corporate finance to have a performance management solution. Every department in an organization should have its own solution, which meets its individual requirements, yet also seamlessly rolls up into a corporate view. For example, if marketing decides to run an additional campaign to improve sales on a specific product line, it will update its own performance management solution to reflect the increase in expected sales. In an integrated world, purchasing (which has its own individual solution) would now see that it has a shortfall to expected sales. Now, imagine this happening corporate-wide across all departments, with the data seamlessly rolling up into a corporate view and the systems all being supported cost-effectively by IT (see Diagram 1). At least that is how we at SAP see the Holy Grail for CPM.

Find out more at: www.sap.com/solutions/performancemanagement/. ●

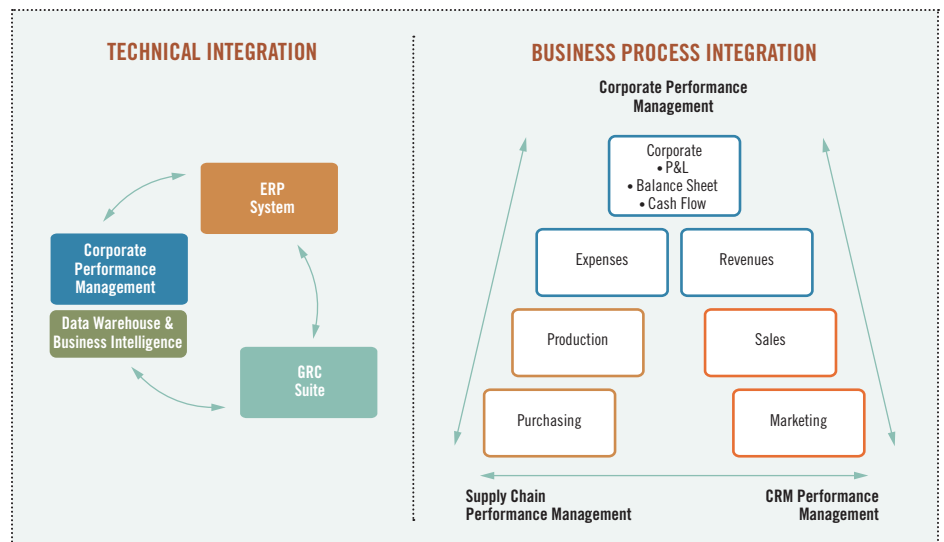


Diagram 1. Corporate performance management integration

Innovative Analytics for Improved Performance and Consistent Decision Processes

Commentary by Steve Trammell
Corporate Alliances, ESRI

The Challenge

A major food and beverage provider was experiencing inconsistent performance from its new and existing outlets throughout the U.S. Furthermore, the more competitive environment was requiring a more agile approach to managing real estate assets.

Management suspected that existing BI and geographic information systems (GIS) applications could help solve the problems, but was unsure how to leverage these assets. The challenge was to get these disparate systems working together either to arrive at a solution or to determine if additional applications were needed.

A consulting firm was engaged to evaluate the existing applications and suggest possible courses of action. It was determined that updated versions of the BI and GIS applications were needed, as well as new GIS functionality and current GIS data. The consultant brought in the existing BI and GIS providers to discuss their current capabilities and to lay out a high-level vision of how the client's challenges could be addressed.

The Solution

The BI customer analytics application was augmented with data from a new customer survey, resulting in a more accurate customer profile. This customer profile was then matched to ESRI's Community Tapestry data (demographics and purchasing potential), and a nationwide map was generated showing the distribution of people matching the profile.

The GIS could then be used to create new trade areas, and thus revenue potential, for each outlet across the nation. Trade

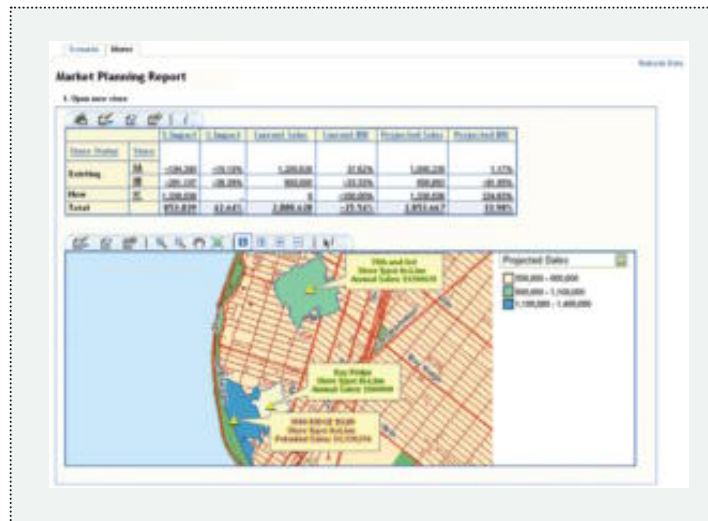


Figure 1. If the market area will support another outlet, the user can display a map showing projected revenue potential for various locations.

areas for each outlet were created based on actual travel time over the street network as opposed to the more traditional ring analysis. A further refinement of the trade-area determination model accounted for the proximity of the company's competitors and target customers. This method of modeling trade areas is a much more accurate measure of spending potential than a purely drive-time model.

Utilizing nationwide data and a more sophisticated trade-area model addressed many of the site performance inconsistencies the company experienced using the previous analysis method. These same improvements to existing systems could now be leveraged more proactively to manage the company's entire real estate portfolio using "what if" scenarios.

For example, if a market area is found to be underperforming, a new trade-area analysis for each outlet is conducted to check individual site performance. Depending on the outcome of the analysis, the user is presented with options for improving the area's performance, such as closing, moving, remodeling, or adding outlets. If the market area will support another outlet, the user can display a map showing projected revenue

potential for various locations. (See Figure 1.) If an underserved location is identified, the user can then add an outlet and conduct a trade-area analysis for the new location. The revenue potential for the new trade area is then passed back to the BI analysis package, where factors such as acquisition, construction, and operations are considered and annual net revenue is projected. If the projection is favorable, the real estate acquisition group is notified and tasked with obtaining a site in the newly identified area. Available sites are run through the trade-area analysis application to validate potential performance.

Benefits

A more rigorous and repeatable market analysis methodology gave the company greater confidence in sales forecasts. The flexible and simple-to-use, scenario-based market analysis tool greatly accelerated the site selection process, resulting in quicker returns on investment for expansion projects.

The company is now able to react to changes in the marketplace much more effectively. Developing problems can be mitigated early to minimize or eliminate losses, and newly developing markets can be served more quickly. ●

LESSON FROM THE EXPERTS

You Can't Offer a Solution If You Don't Know the Problem

By **Steve Trammell**
Corporate Alliances, ESRI

ESRI's geographic information system (GIS) applications had been used for many years by a major food and beverage provider. GIS was an indispensable part of the decision processes in its business. The end users in the company were satisfied with the functionality of their GIS applications as well as the support they were receiving. Therefore, we were surprised when we were approached by a large consulting firm to meet with the company to explore potential solutions to an increasingly critical challenge they were facing.

We also learned at this time that the company was using a number of BI applications, and the providers of those applications were being asked to join the discussions. Our account manager was completely unaware of the company's challenge, even though he was in regular contact with its GIS users.

In the conversations that followed, it became apparent that the existing GIS and BI applications could solve parts of the problem, and some additional GIS applications could solve the rest. There would also be some GIS and BI services work necessary to tailor the resulting solution to the specific needs of the company.

What was proposed was an integrated solution that provided a common point of interaction with the GIS and BI applications. Fortunately, over the past few years GIS has been integrated with virtually all major BI platforms, so a comprehensive solution could be built that leveraged many of the company's existing GIS and BI applications.

This experience showed us that today's technology environment requires account managers who are willing to engage their customers in discussions about technologies that may fall outside the account manager's area of expertise. In this case,

the GIS account manager never inquired about non-GIS related technologies being used by the company.

It was also obvious that the BI providers had never asked the same type of questions of their users at the company. Had the GIS or BI account managers proactively engaged the company's other technology users, they might have solved the company's problem sooner and avoided the cost of consulting services.

ESRI has been participating in business application conferences over the past few years, most notably BI and enterprise content management (ECM) events. Well over half of the "prospects" we engaged at these events were from organizations that were existing ESRI clients. However, many of these individuals were either unaware of the GIS implemented in their organizations or did not know it could be integrated with their departments' existing applications. This indicates that the segregation of applications experienced at the food and beverage company is a common, yet often unacknowledged, problem.

We have applied the lessons learned from these experiences in both our marketing and sales efforts. Account representatives with domain expertise in application areas other than GIS, such as BI, are no longer assigned to accounts based on location, but to accounts where their expertise will be most useful.

Our marketing efforts have spread to publications and events not targeted at traditional GIS markets. The marketing efforts are aimed at making the end users of other application platforms aware of GIS so they begin asking about GIS resources already being used in their organizations.

Had the GIS or BI account managers proactively engaged the company's other technology users, they might have solved the company's problem sooner and avoided the cost of consulting services.

Now that GIS can be deployed in service-oriented architectures (SOA), it is imperative that account representatives become familiar with their clients' entire technology landscapes. The ability of SOA to provide solutions by using services from numerous applications, including GIS, requires that account managers be cognizant of a wide variety of technologies. It also requires that technology providers make information about other technologies readily available to their account managers. ●

For a free white paper on this topic, [download "Developing and Deploying an Integrated Geoenabled SOA Business Solution: A Case Study,"](#) or [click here](#) for more information about ESRI.

CASE STUDY

Cablecom Reduces Churn with the Help of Predictive Analytics

Commentary by Federico Cesconi
 Head of Customer Insight and Retention,
 Cablecom GmbH

Situation

Cablecom GmbH is Switzerland's largest cable network operator. Broadcasting to 1.6 million homes across the country, Cablecom offers customers cable television, broadband Internet, and mobile and fixed-network telephony.

Cablecom's reputation as a market innovator is mirrored in house, where the company uses the very latest technology to better serve its customers.

Challenge

In the broadcast and telecommunications space, where churn is a major global issue, it is vital to be able to target customers efficiently with tailored marketing offers.

"Seeing the world through the customers' eyes is at the heart of our business," said Federico Cesconi, head of customer insight and retention at Cablecom. "Customer information and feedback is key to this process and is enabling us to take a proactive approach to one of our industry's most pressing problems."

"It's much easier to retain a customer than to try to win him or her back. Many win-back activities and offers are often too late for the customer looking to switch providers. In many cases, the decision has been made weeks before, and it is costly and difficult to reverse," explained Cesconi.

Solution

Cablecom recognized the key to tackling churn was to identify the point at which customers become dissatisfied—before they make the decision to switch to an alternative provider. It opted to use SPSS's best-of-breed predictive analytics technologies to assess and analyze customer feedback and to evaluate it alongside existing behavioral and demographic

data sets. Cablecom's aim is to predict future behavior and proactively meet customer demands, and therefore reduce churn.

At significant points across the customer lifecycle, Cablecom solicits detailed customer feedback. Initial results showed a peak in churn behavior between 12 and 14 months into the lifecycle, but the decision to churn was actually made around the ninth month into a contract.

As a result, the company targeted customers who had been with the company for about seven months with an online satisfaction survey. Dissatisfied customers most at risk of switching were then targeted by a dedicated customer retention team that proactively reached out to persuade them to stay.

As is typical with surveys, though, only a minority of customers responded. Using SPSS's data mining solution, Cablecom combined the survey results with other customer data and built predictive models that could score all customers on their likelihood to churn. They found more than 100 factors that indicate if a customer is likely to churn, including initial activation period, number of customer service queries, price band, and original sales channel. This insight was applied across the entire customer base to predict which customers would churn and

enable Cablecom to take proactive steps to retain them—as well as identify a segment likely to have high satisfaction as a target for cross- and up-sell campaigns.

Results

Data from the customer feedback program gives an accurate picture of the traits of satisfied and dissatisfied customers. Across the wider customer base, Cablecom can identify satisfaction levels (with a 78 percent degree of accuracy) and take appropriate action to retain or even cross- or up-sell customers.

SPSS's technology enables Cablecom to identify more accurately those customers who are likely to churn and take proactive action to improve retention. Early pilot studies show a drop in churn rate from 19 percent to an impressive two percent among treated customers.

Cablecom can also provide better leads for cross- or up-sell campaigns by targeting a customer segment most likely to accept any such offer.

As a result, Cablecom has achieved substantial results that directly affect its bottom line. ●



LESSON FROM THE EXPERTS

Enhance Enterprise BI to Drive Business Growth with Predictive Analytics

By **Colin Shearer**

Senior Vice President, Market Strategy, SPSS Inc.



Stories abound of the business improvements and impressive return on investment that organizations have achieved using predictive analytics, including dramatic reductions in customer churn, huge boosts in marketing response rates, and significant savings through better fraud detection and risk management. As C-level executives become aware of this opportunity to improve their businesses, many companies find themselves adopting predictive analytics.

Before rushing headlong into a predictive analytics project, however, you should consider a few key points to help you ensure success and maximize the value you can obtain from predictive analytics—both in your initial project and ultimately across the enterprise.

Choose your initial target. Look for areas where predictive analytics have been applied successfully in companies like yours. Across sectors, the most common initial areas for implementation are CRM—where it is used to optimize customer acquisition, value, and retention—and risk management. Seek out case studies and learn lessons that can apply to your business.

Treat it as a business project. Analysis may be a technical activity, but having Ph.D. statisticians drive your project is the wrong way to do it. Striving for technical improvements in accuracy of predictive models isn't the same as maximizing the business benefit of your analytical results. Bring together analytical skills and business skills, and ensure the focus stays on business success.

Leverage all relevant data. Predictive analytics leverages your data assets. Many

companies start with the most easily accessible data—typically, descriptive/demographic and behavioral/transactional data held in the customer warehouse. But the more you know about your customers, the closer you can come to a holistic view of them in your data—enabling you to better understand them and more accurately predict their behavior. Look for additional data that can add incremental value to your predictive analytics efforts. Consider text (e.g., transcripts of call center conversations), clickstream data from Web site visits, and attitudinal data from surveys.

Before rushing headlong into a predictive analytics project, consider a few key points to help you ensure success and maximize the value you can obtain.

Plan how to use the results. No matter how good your analysis or how accurate the predictive models you build, you only generate ROI when you do something with the results. Deployment of results is crucial and needs to be planned at the start of your project. Typically, deployment involves making intelligent model-based decisions at key decision points in the business processes you're trying to improve. This might be deciding whether to include a specific customer in an outbound mailing campaign that makes a retention offer to high-value customers predicted to be at serious risk of defection; deciding, during a contact center conversation, to make a cross-sell offer based on a propensity model's recommendation; or deciding whether

to allow a transaction based on a model's estimate of customer credit risk. As you plan how these decisions will fit into your business processes, you also need to consider how they will integrate with the operational systems that support them. For example, will model scores be used to select and rank target lists that are sent to campaign management systems, or will models be invoked in real time during contact center conversations, where decisions will be made based on data that has just become available?

Focus on early success, but keep the enterprise vision. Although individual predictive analytics solutions can generate significant ROI, organizations that apply the technology across different areas of their business gain maximum returns by taking an infrastructure approach to implementing predictive analytics. While you should focus on the success of your initial project, keep in mind the long-term benefits of sharing a common analytical platform and data views across your future applications. Plan a roadmap that will support your evolution to a truly predictive enterprise. ●

For a free white paper on this topic, [download "Predictive Analytics: Extending the Value of Your Data Warehousing Investment,"](#) or [click here](#) for more information about SPSS.

CASE STUDY

Top Retail Bank More Than Doubles Campaign Profitability Using Latest Uplift Modeling Techniques

Commentary by Mark Smith
Executive Vice President, Portrait Software

A large successful American bank that had recently launched a new high-margin product was aiming to increase new product bookings. Should the bank mail all its customers to increase sales of the new product and incur a high mailing cost while still providing a positive return on mailings? Or could the bank increase incremental sales while reducing mailing costs by focusing its mailing on those customers best influenced by the marketing?

The value of many banking products is high, so even an increase in product sales of a fraction of a percentage can provide a positive return on investment for direct mailings. Leveraging the latest techniques in uplift modeling, this bank discovered that it could increase total sales by about 10 percent while reducing mailing volumes by 60 percent. These results more than doubled campaign profitability compared with the bank's previous state-of-the-art approach to modeling.

The graph in Figure 1 shows the relative profitability of targeting different volumes with the two models, using the profit achieved by targeting the whole population as a base. As can be seen, targeting 20 percent or 30 percent of the population achieves an uplift nearly 80 percent higher than does targeting the whole population. More significantly, the profit achieved using uplift at cutoffs in the typical and desirable 20-30 percent range from 60 percent to 125 percent larger using the uplift model than with the bank's champion model.

Maximizing Returns on Direct Marketing

What is uplift modeling and how does it differ from the traditional approach of response modeling? The core benefit of uplift modeling

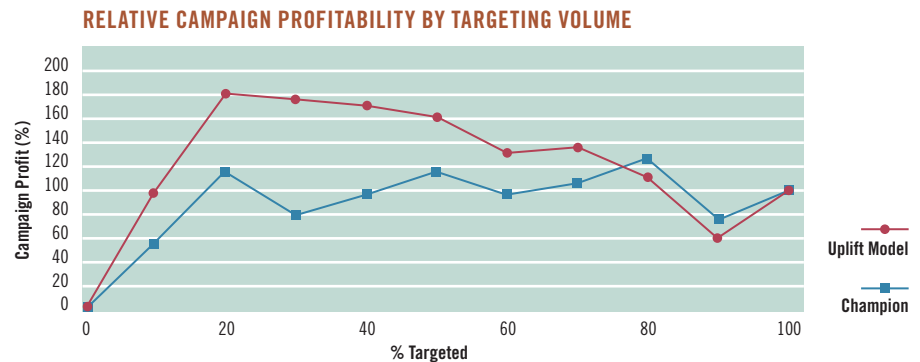


Figure 1. Targeting 20 percent or 30 percent of the population achieves an uplift nearly 80 percent higher than does targeting the whole population.

is the ability to optimize targeting to maximize the returns from direct marketing. Uplift modeling achieves this by predicting, at an individual level, the change in behavior likely to result from a particular marketing intervention. Uplift modeling enables us to target customers whose behavior we can change, rather than targeting those who would have behaved in a way even without the direct marketing intervention.

As in the case with the leading retail bank referenced above, uplift models were able to generate higher returns on marketing spend through lower targeting volumes. This constitutes a multiple win, because it means an increase in sales while simultaneously reducing targeting or mailing costs. Additionally, where incentives are used, the cost of these is usually significantly reduced because they are not “wasted” on customers who would purchase anyway. Customer satisfaction also tends to improve because customers who dislike being contacted are less likely to be targeted.

The Multiple Benefits of Uplift Modeling

Uplift modeling allows businesses to optimize their targeting of customers in such a way as to maximize the return on investment of marketing spend. The benefits from this in the areas of demand generation and customer retention include:

- Reducing the number of customers required to achieve a given level of business stimulation and thus reducing costs
- Increasing the level of business generation achieved for any given level of spend

- Lowering customer dissatisfaction by decreasing the level of negatively received material
- Enhancing understanding within the business of the effectiveness of various kinds of marketing spend
- Eliminating many or all of the negative effects associated with mistargeted campaigns
- Increasing customer retention

Uplift Optimizer: The Solution

Portrait Software's uplift modeling solution has been successfully applied at five of the world's top 10 financial services firms. Portrait's Uplift Optimizer is the culmination of more than eight years of focused R&D effort, and is the first automated solution that predicts the incremental impact (uplift) of marketing campaigns. Singled out for its innovation, Uplift Optimizer differs from traditional modeling by identifying the two most important segments for targeting:

- A. The Persuadables—people who buy (or renew) who would not have done so if the campaign had not been run
- B. The Do Not Disturbs—people for whom the campaign triggers a negative response

Portrait's Uplift modeling solution can be used as a stand-alone package or as an integrated component of an SAS package or other modeling applications. ●

[Click here](#) to request more information about Portrait Software.

Best Practices Awards 2007

TDWI's Best Practices Awards recognize organizations for developing and implementing world-class business intelligence and data warehousing solutions. Here are summaries of the winning solutions for 2007.



BI/DW ON A LIMITED BUDGET

➤ **WINNER: StubHub**

SOLUTION SPONSOR: BUSINESS OBJECTS

San Francisco–based StubHub is a unique open marketplace that enables customers to buy and sell tickets at fair market value to a vast selection of sporting, concert, theater, and other live entertainment events. Since its founding in 2000, StubHub experienced triple-digit year-over-year growth, straining its existing information management infrastructure.

StubHub's limited visibility into current data, customers, and business drivers threatened its long-term survival, and StubHub did not have the resources for an expensive overhaul. Instead, StubHub's BI/DW initiative emphasized investing in specific technologies to empower end users and free up existing staff resources—reusing existing technology investments and employing open-source technologies where possible.

StubHub successfully replaced its failing database with a stable, scalable, and extensible data warehouse environment. The new database and reporting environment was developed and deployed in eight months. StubHub now maintains a rolling two-year road map to build on the platform and continue driving business value through actionable intelligence.

StubHub leveraged industry best practices that can be adopted easily by similar organizations. For example, the initiative was designed by a cross-functional team to ensure full organizational buy-in. The company developed a detailed project plan that included a phased approach and clear milestones. The team used a traditional star-schema data model, a software development life-cycle (SDLC) methodology, and data-quality techniques. Finally, the development team delivered “quick wins” with key reports that demonstrated and reinforced value to the organization.

The StubHub team needed to find innovative ways to achieve an enterprise-class implementation on a limited budget. These included using open-source, on-demand, and free or bundled software; the extensive use of “sandboxes,” which minimized BI headcount and enabled the BI team to focus on delivering value-added projects; repurposing hardware to further reduce capital expenditures; and collaboration with partner groups to create new reports and functionality.

To read more about StubHub's innovative solution, see pages 6–7.

CUSTOMER INTELLIGENCE

▶ WINNER: Verizon Data Services

SOLUTION SPONSOR: TERADATA CORPORATION

Verizon Communications Inc. delivers communication innovations to mass market, business, government, and wholesale customers. Verizon Wireless serves 59 million customers nationwide; Verizon Business operates one of the most expansive wholly owned global IP networks; and Verizon Telecom deploys the nation's most advanced fiber-optic network.

In 2003, Verizon Telecom developed Performance Assurance Reporting Suite (PARS), a relational online analytical processing application that serves as a scalable common point of access for metrics and analytics. PARS operates atop the company's active enterprise data warehouse environment from Teradata on the same high-speed server.

Verizon Communications Inc., formed in 2000 by the merger of Bell Atlantic Corporation and GTE Corporation, has continued to grow through acquisitions. Verizon had difficulty unifying the diverse markets and product lines of the former companies and reconciling data across markets and regions. The company recognized the need for universal access to a central data repository to expedite data queries and provide valuable customer insight.

PARS has changed the way Verizon does business. With PARS, Verizon now has a single 360-degree view of the customer across all acquired brands and regions, and it operates with a single, common version of the truth for reporting and analysis.

Verizon's initiative began with an active enterprise data warehouse that served to aggregate metrics across hierarchies. Verizon sought insight through customer and product data to successfully promote its offerings.

By building its own application, Verizon has saved approximately \$2 million in resources and \$1 million in licensing fees and maintenance over a four-year period, and also reduced preparation time by 90 percent.

DATA GOVERNANCE

▶ WINNER: UMB Bank

SOLUTION SPONSOR: HARTE-HANKS TRILLIUM SOFTWARE

UMB Financial Corporation is a financial holding company that offers complete banking and related financial services to individual and business customers. Its banking subsidiaries own and operate 139 banking centers. Subsidiaries of the holding company and the lead bank, UMB Bank, n.a., include an investment services group based in Milwaukee, Wisconsin, a trust management company in South Dakota, and single-purpose companies that deal with brokerage services, consulting services, and insurance.

Years ago, UMB Bank associated its customers with certain product sets and managed its technology in the same way; however, many customers spanned both lines of business. UMB had excellent individual lines of business systems, but immature customer enterprise information management systems.

During a growth phase, UMB aligned its tools and technologies to fit its business model, which enabled UMB associates to become customer-centric. The transition required breaking down some of the organizational silos that functionally segregated customer data. In addition, UMB needed a mechanism by which all information could be aggregated.

Three key aims drove this initiative: 1) to increase front-line efficiencies by consolidating customer data, thus reducing the research time necessary to quantify the customer relationship; 2) to improve cross-sale opportunities by consolidating current product offerings and identifying complementary products; and 3) to empower associates to fulfill UMB's mission of knowing its customers. To achieve these objectives, UMB needed an enterprise CRM solution with a focus on data quality best practices. UMB evaluated available technologies that would help aggregate existing host systems to create a foundation for the overall CRM effort.

The finished design offered flexibility and logical building blocks for the future. The first phase was launched in May 2006. The project yielded CRM success, some of which UMB links directly to Trillium Software and Oracle UCM. UMB's information aggregation model is highly streamlined and has gained accolades from peer group customers and fellow commercial entities.

ENTERPRISE BUSINESS INTELLIGENCE (EBI)

 **WINNER: DaimlerChrysler AG**

DaimlerChrysler (DCX) is a global automotive manufacturer formed by the merger of Daimler-Benz and Chrysler Corporation. Procurement at DCX is part of an organization called Global Procurement and Supply (GP&S), and is extremely important to DCX because approximately \$100 billion is spent annually on parts and services. Cutting these costs by even a small percentage generates large benefits.

After the merger, operational systems for creating purchase orders and paying suppliers could not be standardized quickly. A data warehouse was built for combining the procurement summary data; the result was the Global Procurement & Supply Information System (GPSIS), which has grown far beyond the original vision.

GPSIS aids in global negotiation, facilitating achievement of lower costs and higher quality of purchased parts and services. The annual corporate goal for reducing purchased-part costs exceeds \$1 billion, and GPSIS is key to meeting that goal, as it is already credited with helping to create savings of millions of dollars.

GPSIS began as a globally available DW with structured, Web-based reporting and has evolved into a fully functional BI platform with strategic planning applications.

The original goal was merely to understand how much business was done with each supplier. This provided leverage for negotiating lower prices based on the full global value of business across commodity groups. DaimlerChrysler then needed a way to consider the other aspects of total cost of ownership, since choosing a supplier means more than finding the lowest price. The company used the GPSIS DW as a foundation for balanced scorecards to track supplier performance globally. The next step for the system was to support both Daimler and Chrysler corporations as they operated separately. Management at both companies has expressed strong desire to retain GPSIS functionality for their respective global purchasing operations.

CRM applications provide a basis for managing a company's relationship to its customers. In a similar fashion, GPSIS is a supplier relationship management system that provides a basis for managing the company's relationship to its suppliers.

ENTERPRISE DATA WAREHOUSING

 **WINNER: Ingenix**

SOLUTION SPONSOR: IBM CORPORATION

Ingenix, a wholly owned subsidiary of UnitedHealth Group (UHG), is a global healthcare information company founded in 1996 to develop, acquire, and integrate best-in-class healthcare information capabilities.

At the heart of all BI activity at Ingenix/UHG is the Galaxy enterprise data warehouse (DW). Galaxy is an atomic DW with transformations that integrate subject areas across many UGH platforms. Galaxy eliminated data silos and inconsistencies across these platforms and is now the company's single source of truth for many applications. Technologies and techniques developed for UHG are also sold to other healthcare companies, making Ingenix both an in-house analytics and IT group and a commercial solutions provider.

A centralized program-management office coordinated the transition to Galaxy from two legacy DWs and legacy systems. Requirements-gathering began in 1998, and Galaxy went into production in 2001. Its database was 2.5 TB, growing to 18 TB today. About 50 business analysts, data modelers, DBAs, and software engineers worked mostly in-house, with the exception of IBM data modeling assistance.

The consolidation saved the company several million dollars and has produced ongoing cost savings, eliminated redundancies, and dramatically increased the efficiency of integrating new data sources. Information quality is the foremost objective for Galaxy.

Ingenix has three information-quality goals: 1. Ensure the company meets business-defined quality standards and goals. Galaxy "must always be valid" to a Six Sigma level of verifiable quality. 2. Monitor information quality continuously to save operational maintenance costs. 3. Continuously improve Galaxy's overall quality to ensure UHG's overall health.

Galaxy's industry-differentiating characteristic is its value-added transformations, which derive information from multiple sources—the heart and soul of the information quality management program. These value-added transformations enable difficult processes such as finding patients to invite to drug trials. Galaxy's completeness and breadth, quantity of information, and trustworthiness give value to this process.

Galaxy's success is due to Ingenix and UHG's comprehensive and continuous focus on information quality, their organizational commitment to Galaxy across all aspects of the business, and their built-in, industry-differentiating, complex, value-added transformations.

GOVERNMENT AND NONPROFIT

▶ **WINNER: Richmond Police Department**

SOLUTION SPONSOR: INFORMATION BUILDERS

The Richmond, VA, police department needed an analytic system that would help determine the probability of particular types of crime occurring in specific areas at specific times. Combining technology expertise from SPSS and Information Builders, the Richmond Police Department (RPD) is now using predictive analysis and business intelligence technology to apply information-based policy.

The RPD system's predictive analytics capabilities help determine where a crime may occur and empower officers at all levels to take immediate actions. Overall, the system delivers accurate insight into where crimes might occur, increases public safety, and reduces the number of calls for assistance. From 2006 to 2007, major crime was down 19 percent; from 2005 to 2006, it was down 21 percent.

In 2002, RPD procured an enterprise data-mining workbench from SPSS to predict patterns of behavior that would help it use police resources more effectively. Recognizing the value in this statistical analysis system, in 2005 and 2006 RPD enhanced the system by deploying Information Builders' WebFOCUS BI tools in conjunction with GIS mapping software from ESRI. The benefits of the system have surpassed original expectations.

Three key elements contribute to the success of the project: real-time information, usability, and portability. The agility and responsiveness of the RPD system give officers fresh information at every eight-hour shift change, and some data is updated continuously.

Information Builders tied components of WebFOCUS into mapping software from ESRI to create an easy-to-understand, intuitive presentation of the law enforcement application. Dashboards serve as the key interface. The system offers point-and-click drill-down capability and visual representation of current analytic data. The new application's portability is crucial to successful deployment. With data terminals in their cars, officers receive alerts, conduct planning, and develop strategies, and supervisors are able to review performance outcomes in each area.

The predictive analysis system uses a database of arrests, past calls to police, and crime incidents. Advanced analytics from SPSS are used to examine how current crime reports relate to data on past, present, and anticipated actions. Advanced analytics include statistical, mathematical, and other algorithmic techniques.

GOVERNMENT AND NONPROFIT

▶ **WINNER: Memorial Sloan-Kettering Cancer Center**

Memorial Sloan-Kettering Cancer Center (MSKCC) is one of the premier cancer centers in the world, committed for over a century to exceptional patient care, leading-edge research, and superb educational programs. The Center has nearly 9,000 employees. In 2006, more than 21,000 patients were admitted to Memorial Hospital, and MSKCC accommodated over 430,000 outpatient visits.

The MSKCC's data warehouse, the Institutional DataBase (IDB), was launched in 1988 when the information systems senior management team formed a broad-based steering committee. This committee defined a methodology to create, maintain, and administer a data warehouse that would ultimately unite clinical, operational, and financial data. The IDB has evolved well beyond its original objective to make institutional data readily available for research, decision support, and executive information systems.

Today, IDB data provides clinicians with a patient's history of chemotherapy; it also helps clinicians study and compare the effectiveness of treatment regimens.

MSKCC's information systems' data delivery group, DataLine, has provided thousands of reports to support cancer research initiatives, leading to active protocols and published papers. The group has also provided key information in the generation of tens of millions of dollars in institutional grants.

The IDB is the primary data source for "visit processing," which supports the coding of about \$1 billion in annual ambulatory revenue. Further, IDB reports have been created to help support accreditation from the Joint Commission on Accreditation of Healthcare Organizations (JCAHO). Also, census information from the IDB supports budgetary decisions by MSKCC management.

In 1989, the IDB consisted of two subject areas. Today, 16 disparate applications with various DBMSs feed the IDB to produce more than 1,000 daily reports. In place since IDB's inception are "morning reports" for each subject area, which are monitored daily in order to identify data anomalies, growth, and exception conditions.

DataLine was created in the early 1990s, starting with one data expert; the group has grown to six, and develops more than 500 ad hoc and scheduled reports per year. Power users also provide data delivery services and run reports out of the IDB for 10 departments, accounting for thousands of reports per year.

OPERATIONAL BI

🏆 WINNER: Airlines Reporting Corporation

Airlines Reporting Corporation (ARC) is an airline-owned company offering financial settlement solutions and data and analytical services to airlines, travel agencies, airports, and travel industry analysts. Its vision was to provide a travel industry solution to customers and give them the resources of an industry data warehouse.

ARC's cross-functional Data & Analytical Services team manages ARC COMPASS, a data warehouse that captures and analyzes information related to air travel ticketing and holds 39 months of historical data. ARC has centralized and simplified industry reporting in virtually every internal decision-making and workflow function. The result is significant tangible and intangible cost reductions.

ARC's first BI application, ARC Document Retrieval Service, helped the company realize \$12.5 million in core cost reduction related to shipping fewer paper ticket coupons. ARC also mines data on behalf of external customers. The results help travel industry customers create targeted marketing plans, assess incentive programs, and optimize promotions. ARC offers insight for route and network planning, sales trends, customer behavior, demand forecasts, statistical profiles, point of sale analysis, and targeted marketing campaigns.

The Data & Analytical Products business line has experienced better than 100 percent revenue growth in each of the past two fiscal years. ARC offers more than 60 revenue-generating reporting products, and launches five new revenue-generating products each year.

ARC has implemented a comprehensive best-practice data governance framework through the assignment of data custodians and data stewards. The data provisioning team provides data integrity through modeling of the data across business, operational, and technical levels. This team also supports ongoing infrastructure initiatives. ARC develops many of its products and services through an iterative prototyping methodology. Prototypes of reports and even analytical front-end applications are developed quickly and modified easily.

ARC was the first organization in the industry to comply with the data security standard of the payment card industry, and complies with the European Union Safe Harbor program. With its entry into BI/DW solutions, ARC has expanded existing market segments and developed new ones. The company has significantly expanded its marketing list services with ARC COMPASS and can provide highly customized marketing lists for destination marketing organizations, airlines, and various service providers.

PREDICTIVE ANALYTICS

🏆 WINNER: Corporate Express US Inc.

SOLUTION SPONSOR: MICROSTRATEGY

Corporate Express, part of the Holland-based Corporate Express NV (NYSE: CXP), is one of the world's largest B2B suppliers of essential office and computer products and services, with 2006 sales of approximately \$4 billion in North America, including \$3.6 billion in the U.S.

Corporate Express's distribution infrastructure is among the industry's most advanced. It delivers an average of \$13 million in products every business day in the U.S., \$6.31 million of which is ordered through the Internet. To increase sales on all combined orders, Corporate Express built a market basket application to recommend related products to customers during their purchase.

The market basket application is based on reference products and correlated products. When a customer chooses a reference product, the application presents nine mathematically correlated products to pair with the reference product. Optimal pairing is determined by three metrics:

Support: The probability that the reference and correlated products will be purchased together.

Confidence: The probability that the correlated product will be purchased once the reference product is added to a shopping cart.

Lift: The improvement in the probability that the correlated product will be purchased once the reference product is added to a shopping cart, divided by the original probability that the correlated product will be purchased independently.

With the application in place, analysis now shows and often predicts what items might occur together in a checkout session. The BI application is having enormous impact. The company can track all online baskets and identify what SKUs are most profitable or poor performers when purchased in combinations. According to Corporate Express, the market basket application has generated the following ROI results:

- Average order size for basket purchases increased more than 2 percent
- Average order size for orders with a basket pairing is more than twice that of orders without a basket pairing

Many companies present complementary products as part of the shopping experience, often based on the professional judgment of marketing or merchandising organizations, and complements are often intuitive (e.g., pairing staplers with staple removers). Leveraging predictive analytics allowed Corporate Express to innovate past intuition and judgment to base pairing decisions on solid math.

RADICAL BI

▶ WINNER: Lawrence Livermore National Laboratory

Lawrence Livermore National Laboratory (LLNL), a premier applied science laboratory, is part of the National Nuclear Security Administration within the Department of Energy, and is responsible for ensuring that the nation's nuclear weapons remain safe, secure, and reliable through the application of advances in science and technology.

First developed in 1985, LLNL's data warehouse is now a portal-based Java and Oracle RDBMS application considered highly successful in delivering accurate and timely information. While data warehouses are now commonplace, four of LLNL's approaches and innovations combine to give its users uncommon capability and control in addressing their information needs.

A modular architecture has been employed and continually leveraged for flexibility, control, and self-sufficiency. The Enterprise Reporting Workbench (ERW) interface allows users to create formats and filters and link them to create reports. The modularity has expanded to support interaction with commercial packages, which produced the "workbench" reference in the name—a place where tools can coexist. The ERW home page is a comprehensive workbench of enterprise reporting tools.

Dimension data is separate from fact data, which allows ERW users and application a great deal of flexibility. If fact data contains the base attribute (such as account or employee number), then all relevant dimension data is dynamically available, even if absent from the reporting view. Users can easily create, maintain, and use their own virtual dimensions (URAs) or use a wealth of institutionally maintained dimensions (IRAs). This dimensioning architecture also allows for rapid deployment of completely new dimension sets without programming or creating new view structures.

Over the years, common desktop skills and capabilities have become integral in today's retrieval and analytical processes. When LLNL found some users were re-keying data from paper reports into other desktop processes, they preserved and staged their report files for FTP download. Today, all output is electronic and available in full-featured .xls, .pdf, .tsv, and HTML. Users can also integrate post-retrieval analysis with report processing within the ERW prior to distribution. The Desktop Data Integrator (DDI) allows users to create a workbook once and upload it to ERW, where its data sheets, pivots, and summary sheets will be refreshed automatically.

Solution Providers

The following solution providers have shared their enterprise business intelligence stories and successes, technology insights, and the lessons they have learned for *What Works in Enterprise Business Intelligence*.



Actuate Corporation

701 Gateway Blvd., Suite 600
 South San Francisco, CA 94080
 650.837.2000
 Fax: 650.827.1560
info@actuate.com
www.actuate.com

Actuate Corporation, the leader in business intelligence, performance management, and reporting applications, enables organizations to develop solutions that optimize corporate performance. Applications built on Actuate's open source-based platform provide all stakeholders—inside and outside the firewall, including employees, customers, partners and citizens—with information that they can easily access and understand to maximize revenue, cut costs, improve customer satisfaction, streamline operations, create competitive advantage, and make better decisions.

Actuate has more than 4,000 customers globally in a diverse range of business areas, including financial services and the public sector. Founded in 1993, Actuate has headquarters in South San Francisco, California, with offices worldwide. Actuate is listed on NASDAQ under the symbol ACTU. For more information on Actuate, visit the company's Web site at www.actuate.com.



Business Objects

North American Corporate Headquarters
 3030 Orchard Parkway
 San Jose, CA 95134
 800.527.0580
 408.953.6000
 Fax: 408.953.6001
www.businessobjects.com/company/contact_us/form.asp
www.businessobjects.com

Business Objects is the world's leading BI software company, helping organizations gain better insight into their business, improve decision making, and optimize enterprise performance. The company's business intelligence (BI) platform, BusinessObjects™ XI, offers the industry's most advanced and complete platform for reporting, query and analysis, performance management, and enterprise information management including data integration, data quality, and metadata management. BusinessObjects XI includes Crystal Reports®, the industry standard for enterprise reporting. Business Objects also has the industry's strongest and most diverse partner community, with more than 3,000 partners worldwide. In addition, the company offers consulting and education services to help customers effectively deploy their business intelligence projects.



Collaborative Consulting

70 Blanchard Road, Ste. 500
 Burlington, MA 01803
 781.565.2600
 Fax: 781.565.2700

info@collaborativeconsulting.com
 www.collaborativeconsulting.com

Collaborative Consulting is a leading professional services organization that specializes in optimizing its clients' business and technology capabilities. We combine exceptional business knowledge and market-leading technology expertise with an effective partnership approach, allowing us to understand and solve even the most complex business problems. And, by aligning business and technology initiatives, we can help clients achieve superior, cost-effective business solutions. Founded in 1999, Collaborative provides operational consulting, program management, data services, and technology services for clients across the U.S., with headquarters in Burlington, MA. Collaborative's Web site is: www.collaborativeconsulting.com.



DataFlux Corporation

940 NW Cary Parkway, Ste. 201
 Cary, NC 27513
 877.846.3589
 Fax: 919.447.3100

www.dataflux.com
 sales@dataflux.com

DataFlux enables organizations to analyze, improve, and control their data through an integrated technology platform. With DataFlux enterprise data quality and data integration products, organizations can more effectively and efficiently build a unified view of customers, products, suppliers, or any other corporate data asset. A wholly owned subsidiary of SAS (www.sas.com), DataFlux customers can rapidly assess and improve problematic data, building the foundation for enterprise data governance. Effective data governance delivers high-quality information that can fuel successful enterprise efforts such as risk management, operational efficiency, and master data management (MDM). To learn more about DataFlux, visit www.dataflux.com.



DATAlegro, Inc.

85 Enterprise, 2nd Floor
 Aliso Viejo, CA 92656
 949.330.7690
 Fax: 949.330.7691

info@datalegro.com
 www.datalegro.com

DATAlegro v3™ is the industry's most advanced data warehouse appliance utilizing an all-commodity platform. By combining DATAlegro's patent-pending software with the industry's leading hardware, storage, and database technologies, DATAlegro has taken data warehouse performance, reliability, and innovation to the next level. DATAlegro v3 goes beyond the low cost and high performance of first-generation data warehouse appliances and adds the flexibility and scalability that only a commoditized platform can offer.

Whether you have a few terabytes of user data or hundreds, DATAlegro's data warehouse appliances deliver a fast, flexible, and affordable solution that allows a company's data to grow at the pace of its business.



ESRI

380 New York Street
Redlands, CA 92373
909.793.2853
Fax: 909.307.3046

gisbi@esri.com
www.esri.com

FAST

117 Kendrick Street
Needham MA 02494
888.871.3839
Fax: 781.304.2410

info@fastsearch.com
www.fastsearch.com

IBM Corporation

1 New Orchard Road
Armonk, NY 10504-1722
1.800.IBM.4YOU

www.ibm.com/contact/us/
www.ibm.com/bi

Since 1969, ESRI has been giving customers around the world the power to think and plan geographically. As the leader in GIS, ESRI applies innovative technologies to help organizations create, analyze, and visualize information for more informed decisions. ESRI software is used in more than 300,000 organizations worldwide, including each of the 200 largest cities in the United States, most national governments, and more than two-thirds of *Fortune* 500 companies. Only ESRI provides complete technical solutions for desktop, mobile, server, and Internet platforms. ESRI's Web site is: www.esri.com

FAST creates the real-time search and business intelligence solutions that are behind the scenes at the world's best-known companies with the most demanding information challenges. FAST's flexible and scalable integrated technology platform and personalized portal connects users, regardless of medium, to the relevant information they need.

IBM is the world's leading IT provider. For nearly a century IBM has been enabling companies to make better informed decisions by helping businesses to generate, manage, and extend their enterprise data warehouse. By combining world-class database software, information integration, master data management, analytics, Web services and more, IBM is now helping organizations move beyond traditional data warehousing approaches toward dynamic warehousing—the next evolution in data warehousing. A dynamic warehousing approach makes it easier for IT organizations to support business requirements for actionable information—not just data, but all types of information with intelligence behind it to help people take action and make decisions.



Informatica Corporation

100 Cardinal Way
Redwood City, CA 94063
650.385.5000
800.653.3871
Fax: 650.385.5500
myinformatica@informatica.com
www.informatica.com

Information Builders

Two Penn Plaza
New York, NY 10121-2898
212.736.4433
Fax: 212.967.6406
askinfo@ibi.com
www.informationbuilders.com

Microsoft Corporation

One Microsoft Way
Redmond, WA 98052
425.882.8080
Fax: 425.936.7329
www.microsoft.com/BI

Informatica Corporation delivers data integration software and services to solve a problem facing most large organizations: the fragmentation of data across disparate systems. Informatica helps organizations gain greater business value from their information assets by integrating their enterprise data. Informatica's open, platform-neutral software reduces costs, speeds time to results, and scales to handle data integration projects of any size or complexity. With a proven 13-year track record of success, Informatica helps companies and government organizations of all sizes realize the full business potential of their enterprise data. That's why Informatica is known as the data integration company.

Information Builders created the industry's most widely deployed business intelligence solution and is the leader in real-time operational reporting. The WebFOCUS business intelligence platform, the company's flagship product, has the architecture, integration, and simplicity to permeate every level of the extended enterprise. It is the most scalable, secure, and flexible solution in the market and helps organizations build applications that have no barriers. Information Builders is the only vendor with its own complete integration solution.

iWay Software's ability to seamlessly integrate with WebFOCUS enables unrivaled information access.

Information Builders' award-winning technology provides quality software and superior services to more than 12,000 customers, including most of the *Fortune* 100 and U.S. federal government agencies. Headquartered in New York City with 90 offices worldwide, the company employs 1,600 people and has more than 350 business partners.

Founded in 1975, Microsoft® is the worldwide leader in software, services, and solutions that help people and businesses realize their full potential. Microsoft Business Intelligence is a complete, fully integrated offering, enabling all decision makers to drive increased business performance at strategic, tactical, and operational levels. Microsoft BI enhances the performance of all employees by providing reliable access to information they need to make informed decisions and respond appropriately to changing conditions that impact their business. With Microsoft Business Intelligence, you can take advantage of your information assets, create competitive advantages, improve customer satisfaction, and make well-informed decisions. Visit www.microsoft.com/BI for more information.



MicroStrategy

1861 International Drive
McLean, VA 22102
703.848.8600
Fax: 703.848.8610

info@microstrategy.com
www.microstrategy.com

MicroStrategy is a global leader in business intelligence (BI) technology. Founded in 1989, MicroStrategy provides integrated reporting, analysis, and monitoring software that helps leading organizations worldwide make better business decisions every day. Companies choose MicroStrategy for its advanced technical capabilities, sophisticated analytics, and superior data and user scalability.

With thousands of customer successes and a reputation for innovation and leadership, MicroStrategy is the clear choice for your business intelligence requirements. More information about MicroStrategy is available at www.microstrategy.com.

Oracle Corporation

World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
650.506.7000

www.oracle.com

Oracle's business is information—how to manage it, use it, share it, protect it. For nearly three decades, Oracle (NASDAQ: ORCL), the world's largest enterprise software company, has provided software and services that enable organizations to get the most accurate and up-to-date information from their business systems. Today, Oracle has more than 275,000 customers—including 98 of the *Fortune* 100—in more than 145 countries.

Pentaho Corporation

5950 Hazeltine National Drive, Ste. 340
Orlando, FL 32822
407.812.6736
Fax: 407.517.4575

communityconnection@pentaho.org
sales@pentaho.org
alliances@pentaho.org
www.pentaho.com

Pentaho provides a full spectrum of open source business intelligence (BI) capabilities, including reporting, analysis, dashboards, data mining, data integration, and a BI platform, that have made it the world's most popular open source BI suite. Formed by a highly experienced team of industry veterans, Pentaho has a mission to bring innovative, high quality technology and professional support to the BI market. Pentaho uses a revolutionary approach to development, distribution, and support, made possible by an open source business model. Pentaho is the primary sponsor and owner of popular open source projects including Mondrian, JFreeReport, Kettle, and Weka. Pentaho's technologies support a wide range of business initiatives from sales and profitability analysis to customer analysis, HR reporting, financial reporting, KPI dashboards, supply chain analytics, and operational reporting.



Portrait Software

125 Summer Street, 16th Floor
 Boston, MA 02110
 617.457.5200
 Fax: 617.457.5299
 marlene.jaworski@portraitsoftware.com
 www.portraitsoftware.com

SAP

3999 West Chester Pike
 Newtown Square, PA 19073
 800.872.1727
 info@sap.com
 www.sap.com

SPSS Inc.

233 South Wacker Drive, 11th Floor
 Chicago, IL 60606
 312.651.3000
 Fax: 312.651.3668
 www.spss.com

Portrait Software is a global provider of customer interaction optimization software that makes every interaction count.

Our software is used by large organizations to deliver great customer service and to dramatically improve customer revenues, profitability, and loyalty by applying insight and intelligent decisioning at the point of customer interaction.

Portrait Software offers applications and solutions to optimize these customer interactions to:

- Deliver a differentiated customer service
- Increase profitable sales
- Retain profitable customers
- Identify and prevent fraud
- Manage debt and risk
- Comply consistently with regulations

SAP is the world's leading provider of business software. More than 41,200 customers in more than 120 countries run SAP® applications—from distinct solutions addressing the needs of small and midsize enterprises to suite offerings for global organizations. Powered by the SAP NetWeaver® platform to drive innovation and enable business change, SAP software helps enterprises of all sizes around the world improve customer relationships, enhance partner collaboration and create efficiencies across their supply chains and business operations. SAP solution portfolios support the unique business processes of more than 25 industries, including high tech, retail, financial services, healthcare, and the public sector. With subsidiaries in more than 50 countries, the company is listed on several exchanges, including the Frankfurt stock exchange and NYSE under the symbol "SAP." (Additional information is at www.sap.com.)

SPSS Inc. (Nasdaq: SPSS) is a leading global provider of predictive analytics software and solutions. The company's predictive analytics technology improves business processes by giving organizations forward visibility for decisions made every day. By incorporating predictive analytics into their daily operations, organizations become Predictive Enterprises—able to direct and automate decisions to meet business goals and achieve a measurable competitive advantage. More than 250,000 public sector, academic, and commercial customers rely on SPSS technology to help increase revenue, reduce costs, and detect and prevent fraud. Founded in 1968, SPSS is headquartered in Chicago, Illinois. For more information, please visit www.spss.com.



Sybase

One Sybase Drive
Dublin, CA 94568
925.236.5000

www.sybase.com

Syncsort Incorporated

50 Tice Boulevard
Woodcliff Lake, NJ 07677
201.930.9700
Fax: 201.930.8282

info@syncsort.com

www.syncsort.com

Tata Consultancy Services

Air-India Building, 11th Floor Nariman Point
Mumbai 400 021
India
+91.22.6750.9999
Fax: +91.22.6750.9344

www.tcs.com

Sybase is the largest global enterprise software company exclusively focused on managing and mobilizing information from the data center to the point of action, providing open, cross-platform solutions that securely deliver information anytime, anywhere—enabling customers to create an information edge. Sybase addresses the need for business intelligence and data warehousing with world-class data modeling, ultrahigh-speed business analytics, and comprehensive data integration technologies. Maximize the potential of your information assets and achieve better intelligence capabilities with Sybase IQ, PowerDesigner, and our Data Integration Suite. When businesses get serious about business intelligence and data warehousing, they get Sybase. Visit www.sybase.com.

Syncsort Incorporated is a leading developer of high-performance storage management and data warehousing software. For more than 35 years, Syncsort has built a reputation for superior product performance and reliable technical support. An independent market research firm named Syncsort as one of the top data warehouse 100 vendors seven years in a row. More than 90 percent of the *Fortune* 100 companies are Syncsort customers, and Syncsort's products are used in more than 50 countries to back up and protect data in distributed environments, speed data warehouse processing, and improve performance in applications with high data volumes.

Tata Consultancy Services is an IT services, business solutions, and outsourcing organization that delivers real results to global businesses, ensuring a level of certainty no other firm can match. TCS offers a consulting-led, integrated portfolio of IT and IT-enabled services delivered through its unique Global Network Delivery Model, recognized as the benchmark of excellence in software development.

A part of the Tata Group, India's largest industrial conglomerate, TCS has more than 94,000 of the world's best-trained IT consultants in 47 countries. The company generated consolidated revenues of \$4.3 billion (U.S. dollars) for fiscal year ended March 31, 2007 and is listed on the National Stock Exchange and Bombay Stock Exchange in India. For more information, visit us at www.tcs.com.

About TDWI

TDWI, a division of 1105 Media, is the premier provider of in-depth, high-quality education and research in the business intelligence and data warehousing industry. Starting in 1995 with a single conference, TDWI is now a comprehensive resource for industry information and professional development opportunities. TDWI sponsors and promotes quarterly World Conferences, regional seminars, onsite courses, a worldwide Membership program, business intelligence certification, resourceful publications, industry news, an in-depth research program, and a comprehensive Web site (www.tdwi.org).



MEMBERSHIP

www.tdwi.org/membership

Through TDWI Membership, business intelligence and data warehousing professionals learn about the latest trends in the industry while enjoying a unique opportunity to learn, network, share ideas, and respond as a collective whole to the challenges and opportunities in the industry.

TDWI Membership includes more than 5,000 Members who are business and information technology professionals from *Fortune* 1000 corporations, consulting organizations, and governments in 45 countries. TDWI offers special Membership packages for corporate Team Members and students.

WORLD CONFERENCES

www.tdwi.org/conferences

TDWI World Conferences provide a unique opportunity to learn from world-class instructors, participate in one-on-one sessions with industry gurus, peruse hype-free exhibits, and network with peers. Each six-day conference features a wide range of content that can help business intelligence and data warehousing professionals deploy and harness business intelligence on an enterprisewide scale.

SEMINAR SERIES

www.tdwi.org/seminars

TDWI Seminars offer a broad range of courses focused on the skills and techniques at the heart of successful business intelligence and data warehousing implementations. The small class sizes and unique format of TDWI Seminars provide a high-impact learning experience with significant student-teacher interactivity. TDWI Seminars are offered at locations throughout the United States and Canada.

ONSITE COURSES

www.tdwi.org/onsite

TDWI Onsite brings TDWI courses to customer sites and offers training for all experience levels. Everyone involved gains a common knowledge base and learns in support of the same corporate objectives. Training can be tailored to meet specific business needs and can incorporate organization-specific information.

CERTIFIED BUSINESS INTELLIGENCE PROFESSIONAL (CBIP)

www.cbipro.com

Convey your experience, knowledge, and expertise with a credential respected by employers and colleagues alike. CBIP is an exam-based certification program that tests industry knowledge, skills, and experience within five areas of specialization—providing the most meaningful and credible certification available in the industry.

WEBINAR SERIES

www.tdwi.org/education/Webinars

TDWI Webinars deliver unbiased information on pertinent issues in the business intelligence and data warehousing industry. Each live Webinar is roughly one hour in length and includes an interactive question-and-answer session following the presentation.

TDWI Partner Members

These solution providers have joined TDWI as special Partner Members and share TDWI's strong commitment to quality and content in education and knowledge transfer for business intelligence and data warehousing.



VOLUME 24

What Works
IN ENTERPRISE BUSINESS INTELLIGENCE

1201 Monster Road SW
Suite 250
Renton, WA 98057

425.277.9126
info@tdwi.org
www.tdwi.org

tdwi
THE DATA WAREHOUSING INSTITUTE