A 101communications Publication

TDWI

REPORT SERIES

November 2003

Building the Real-Time Enterprise

Colin White

tdwi

Research Sponsors

Ascential Software Corporation

Business Objects

Celequest Corporation

DataMirror

Informatica Corporation

Teradata, a division of NCR

About TDWI

The Data Warehousing Institute (TDWI), a division of 101communications LLC, is the premier provider of in-depth, high-quality education and training in the business intelligence and data warehousing industry. TDWI offers quarterly educational conferences, regional seminars, onsite training, professional membership, leadership awards, print and online publications, and a public and private (Members-only) Web site.

About the TDWI Report Series

The TDWI Report Series is designed to educate technical and business professionals about critical issues in BI. TDWI's in-depth reports offer objective, vendor-neutral research consisting of interviews with industry experts and a survey of BI professionals worldwide. TDWI in-depth reports are sponsored by vendors who collectively wish to evangelize a discipline within BI or an emerging approach or technology.

For more information about this report or its sponsors, and to view the archived report Webinar, please visit: www.dw-institute.com/rtereport/.

TDWI
REPORT SERIES

Building the Real-Time Enterprise

by Colin White

Table of Contents

About the Author

COLIN WHITE is the President of BI Research. White is well known for his in-depth knowledge of business intelligence and enterprise business integration technologies and how they can be used to build the *Intelligent Business*. With over 33 years of IT experience, he has consulted for dozens of companies throughout the world and is a frequent speaker at leading IT events. White has co-authored several books, and writes a regular column for *DM Review* entitled "Building the Intelligent Business."

Acknowledgements

TDWI would like to thank many people who contributed to this report. First, we appreciate the many users who responded to our survey, as well as those who responded to our requests for phone interviews; second, the industry experts who reviewed report drafts, including Richard Hackathorn. We'd also like to thank our report sponsors who reviewed outlines, survey questions, and report drafts. TDWI's production team includes Dale Chinn, Wayne Eckerson, Denelle Hanlon, and Marie McFarland.

TDWI Definitions

Business Intelligence—TDWI uses the term "business intelligence" or BI as an umbrella term that encompasses ETL, data warehouses, reporting and analysis tools, and analytic applications. BI projects turn data into information, knowledge, and plans that drive profitable business decisions.

Real-Time BI—An organization's ability to react to business needs and changing business circumstances within a single day.

Scope and Methodology

Report Scope. This report is designed for technical executives who are exploring ways to capture and deliver information in near real time. The report provides an overview of "real-time" concepts and technologies and how they can be used to extend existing data warehousing and business intelligence architectures to support intra-day decision making.

Methodology. The research for this report is based on a survey that TDWI conducted in July 2003, as well as interviews with experts in the field including end-user organizations, BI consultants, industry analysts, and report sponsors.

Survey Methodology. TDWI contacted BI professionals in its database and 101communications' database. (TDWI is a business unit of 101communications.) In total, 846 people responded to the survey. Multi-choice questions and rounding techniques account for totals that don't equal 100 percent. Based on these 846 responses, 13 percent had deployed real-time projects; 37 percent were planning to deploy real-time projects; and 50 percent were not sure or had no plans to deploy a real-time project. Of the 313 people planning to deploy a real-time project, 56 percent planned to deploy within 12 months. Of the 108 people who had deployed a real-time project, 38 percent said the project was very successful, and 42 percent said it was fairly successful.

Survey Demographics. The largest portion of the qualified survey respondents (49 percent) are corporate IT professionals. The remainder are independent consultants (26 percent), business sponsors/users (19 percent), vendors (3 percent) or professors or students (3 percent.) One-third of respondents (36 percent) work at companies with revenues less than $100 million, while almost half (48 percent) work for companies with between $100 million and $10 billion in revenue. Two-thirds of respondents (64 percent) are based in North America. And respondents represent a wide variety of industries, with the majority in consulting, financial services, software, and manufacturing. (See the following illustrations for breakouts.)

Demographics

Position

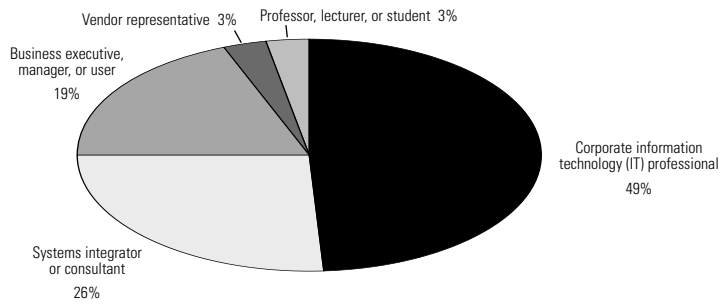

Company Revenues

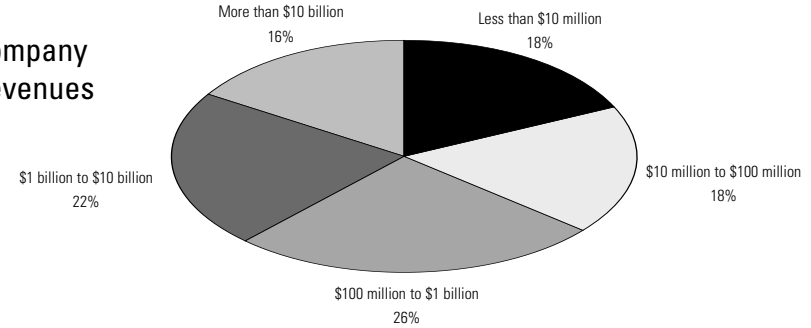

Region

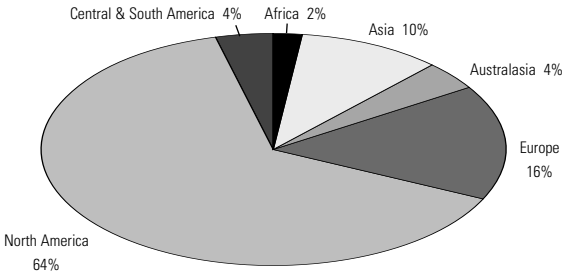

Industry

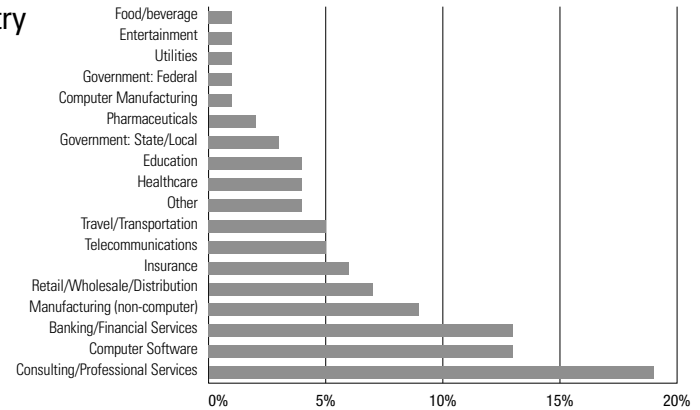

Executive Summary: Is This the Right Time for Real Time?

What Do We Mean by Real Time BI?

The term *real time* has become synonymous with an organization's ability to become more agile, and thus more responsive to changing business circumstances. Although this agility undoubtedly can provide business benefits, there is still universal confusion about what real time BI means from both a business and technology perspective.

In an ideal world, organizations would react instantly to business needs. In reality, however, this is not possible. Accurate business intelligence is required for effective decision making. And it will always take a certain amount of time to collect and deliver this information to business users, and for users to act on this information. The delay, or latency, between a business event occurring, and action being taken, depends on business user reaction time, and on the technologies used to deploy real-time processing BI applications.

Right Time Not Real Time. When developing real-time applications, organizations need to consider how responsive each business process must be to satisfy a specific business goal. From a business perspective, it's more useful to think in terms of "right time" rather than "real time."

Depending on business needs, the action time requirements in real-time environments may vary from a few seconds to several hours. A practical definition, therefore, for the meaning of real time is "the ability for an organization to react to business needs and changing business circumstances within a single day." Many organizations we spoke to agreed with the intra-day definition of real-time because this threshold is important from both a business and a technology perspective.

Business Benefits. For the business, the ability to take intra-day action can have a significant positive impact on both business finances and marketplace perception because this capability enables an organization to use business information to drive its daily business operations efficiently. Although the ROI realized by reducing action time to a specific level is the main factor in justifying real-time projects, this ROI must be balanced against the incremental IT costs of achieving the reduction.

IT Requirements. For the IT group, the technology infrastructure required to deliver business information and support business decision making and actions within a one-day period is more demanding, and often quite different, from that used to deploy applications with less rigorous action time requirements. A wide range of technologies and products support real-time processing, but this is a new and rapidly evolving field, and it is important that real-time infrastructure be designed to weather this evolution. It is vital that this infrastructure provide scalability, performance, and availability.

More Than Just Technology

Real-time processing involves more than just technology. It must be recognized that the business ROI achieved by reducing business reaction times will also depend on the ability of the organization to modify its business practices to take advantage of the improved responsiveness in the IT system.

Our research shows that success in the real-time environment requires the IT organization, from both an organizational and service level viewpoint, to eliminate the distinction between operational and decision support processing. Removing the barrier between operational and decision support IT groups eliminates "turf wars" and gives the IT organization access to both types of skills on real-time projects.

Real-time processing offers significant business benefits, and there are many examples of successful real-time projects. It must be realized, however, that this is still a leading-edge technology. And organizations must track the evolution of the real-time processing market and update real-time processing strategies and frameworks accordingly. It will also be vital to educate both business and technical staff about the benefits, technologies, and pitfalls involved in deploying a *real-time enterprise.*

1. The Real-Time Enterprise

The different types of real-time applications and technologies fall into two basic groups. Some are associated with transaction processing, whereas others are related to decision support processing and business intelligence (BI). Coupled together these applications and technologies support the concept of a real-time enterprise, defined as an enterprise that uses timely information to manage its critical business processes, improve productivity, be more responsive, and thus become more competitive.

What Is Real-Time BI?

Defining Real Time BI. There is considerable misunderstanding in industry about what constitutes real-time processing. This confusion is exacerbated by the constant introduction of new buzzwords like straight-through processing and business activity monitoring, and this confusion is reflected in our survey results. We asked respondents to pick three definitions from a list of possible business and technical meanings for real-time processing. The results showed quite a wide range of differing opinions. The top three answers were (see Illustration 1.1): real-time generation of reports and analytics (53 percent of respondents), faster decision making (52 percent), and business activity monitoring (48 percent). It is important to note that only 13 percent of respondents associate real time with system performance because this result demonstrates that organizations realize real-time BI processing has nothing to do with system performance. It is related instead to business performance and business needs.

Real-Time Actions Are Not Possible: Effective decision making requires accurate business intelligence, and it will always take a certain amount of time to collect and deliver this

Illustration 1.1

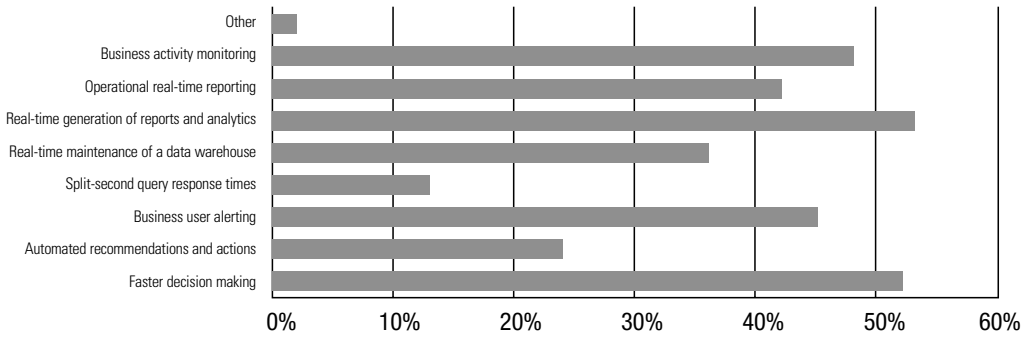

What does real-time BI mean to your team? Multi-choice question, based on 845 respondents.

information to business users, and for users to act on this information. The delay between a business event and subsequent action being taken, or *latency,* will always depend on business user reaction time, and on the technologies used to deploy real-time BI processing applications.

A Definition for Real Time. Our research shows that, depending on business needs, the action time requirements in real-time environments vary from a few seconds to several hours. Therefore, a practical definition for real time BI is, "an organization's ability to react to business needs and changing business circumstances within a single day." Many organizations we spoke to agreed with the *intra-day* definition of real-time because this threshold is significant from both a business and a technology perspective.

For the business, the ability to take intra-day action can have a significant positive impact on both business finances and marketplace perception because it enables an organization to use business information to more efficiently drive its daily business operations.

For the IT group, the technology infrastructure required to deliver business information and support business decision making and actions within this one-day period is more demanding than, and often quite different from, what is required to deploy applications with less urgent action time requirements.

Types of Real-Time Business Processing

There are four main types of real-time processing in a real-time enterprise: straight-through processing, real-time data-on-demand, real-time performance management, and real-time predictive analysis. Of the four types, real-time data-on-demand, real-time performance management, and real-time predictive analysis can be considered to be associated with decision support processing, and thus BI. Given that our survey concentrated primarily on the use of real-time processing in BI, this report therefore concentrates on these three latter types of real-time processing.

Straight-Through Processing. Organizations have employed online transaction processing (OLTP) and batch transaction applications for decades. Many of these applications are poorly integrated, however, and require manual processes to pass data between systems. This lack of integration slows business operations and the organization's overall responsiveness.

One objective of a real-time enterprise is to remove bottlenecks and errors in the processing of business transactions as they flow through the organization. The ultimate goal here is to support *straight-through processing* (STP) where a business transaction is entered once and processed to completion. STP eliminates the need to re-enter data, or to batch transactions for unnecessary processing. Thus, STP benefits a real-time enterprise by reducing delays and improving the flow of business processes. STP usually makes extensive use of application integration software.

Internet retail is an example of a market that can benefit from STP. Many online stores gather orders into a batch file and enter them into the normal order processing cycle. With this batch approach, customers can still order products that are out of stock, and the situation will not be discovered until the order is processed. STP removes this time delay, eliminating unnecessary processing and improving customer satisfaction.

Three Types of BI Real-Time Processing Will Be Discussed

Because STP in not related to BI directly, this aspect of real-time processing will not be discussed further in this report. Its benefits can be seen only in operational systems. However, several of the companies we interviewed said it is important to integrate

operational business processes in order to succeed in building a real-time enterprise. Many of the companies that had embarked on operational integration projects had taken the opportunity to build real-time event feeds into their operational applications for routing to the BI system.

Access to Consistent and Current Operational Data Is Critical

Real-Time Data-On-Demand. Business operations are becoming increasingly data driven, and both business users and business applications need access to consistent and current (sometimes called *zero-latency)* operational data from any place at any time.

A lack of integration in operational systems makes accessing consistent and accurate operational data difficult. An emerging technology known as *enterprise information integration* (EII) can provide access to data in disparate operational data stores, but this approach cannot always resolve data inconsistencies between stores. It does, however, provide access to zero-latency information, such as customer reference data.

To solve data consistency and accuracy problems, many companies capture data from multiple operational systems and integrate it into a *low-latency store.*[1] Based on our definition of real time, the latency of the information in the store (compared with operational systems) is less than one day.

Query applications can use a low-latency store to display, report on, and summarize information about business operations. The data in the store can also be used by rules-driven predictive analysis applications to fine tune predictions and recommendations. An example here would be not to make special offers to customers who have already declined similar offers. In this example, the low latency store can be used to provide a single view of all customer interactions across all channels to a predictive application that is making product recommendations and offers to Internet customers.

In addition to operational queries and reports, Our interviews revealed that several companies were also using a low-latency store to consolidate operational reference data, and in some cases, to propagate operational data between systems (from the front office to the back office, for example).

A Low-Latency Data Store Offers Many Advantages

Easy and fast access to a low-latency data store offers business users throughout the company a current and consistent view of business data and operations. How timely and how current information needs to be in such an environment will vary by company and application. Some applications will require the latency to be a matter of a few seconds (i.e., as close to zero-latency as possible), whereas a few minutes or hours may be sufficient for others.

Real-Time Performance Management. In addition to query and reporting on business operations, organizations also need to measure and manage the performance of those operations.

Start Here: Real-time performance management[2] enables a business intelligence system to tap into business events flowing through the real-time enterprise, and to measure and monitor business performance. Real-time performance management extends traditional operational

[1] Some people call such a store an operational data store (ODS), but as we shall see later in this report, an ODS is just one type of low-latency store.

[2] Various terms describe the use of business intelligence to manage business performance, including corporate performance management, enterprise performance management, business performance management, and so forth. To avoid confusion in this report, we use the term *performance management* to describe the general concept, and *real-time performance management* when discussing real-time processing.

query and report processing by relating measures of business performance to specific business goals and objectives. Business users and applications can then be informed by user alerts or application messages of any situations that require attention.

Gartner Research has coined the term *business activity monitoring* (BAM) to describe this type of processing, and several vendors now offer BAM solutions. Gartner predicts, "by 2004, in enterprises where faster reaction is key to operational effectiveness, BAM will be one of the top four initiatives driving IT initiatives and strategy." [3]

Real-Time Predictive Analysis. The three types of processing discussed so far are reactive in nature; i.e., they involve processing business events either as they occur or after they have happened. Predictive analysis, on the other hand, involves trying to predict the outcome of a business event, such as the risk of granting someone a loan, or the potential for a credit card transaction to be fraudulent. This style of processing in a real-time environment involves sending event information to a predictive analysis application and requesting a recommended action. The application must respond in a timely manner, because a customer may be on a Web site or telephone, waiting for a response. A predictive analysis application employs rules or models to predict a probable outcome or to make a recommendation. These rules or models may be generated from historical business intelligence using, for example, a data mining tool.

"The Benefits Are Unbelievable If You Do It Right"

Together, these four types of real-time processing form the main components of a real-time enterprise strategy. It is vital for organizations to understand the role of each of these components, and the business benefits they can provide. "The benefits are unbelievable if you do it right," says Danny Siegel, Senior Manager of Business Technology at Pfizer Inc. "With real-time BI you can drill down from strategic data to what happened 15 minutes ago."

Understanding Your Business Processes

A real-time enterprise cannot succeed unless the organization understands its business processes (both manual and automated), and how information flows between those processes. This is why *business process management* is receiving considerable attention at present. This software helps an organization analyze, document, and manage its business processes, whether they be simple processes like handling customer orders or more complex processes such as matching supply to demand. This also explains why many application integration software vendors are either developing process management solutions or acquiring them. Application integration software is ideally suited for moving event-based messages both within and across organizations, and for supporting STP. When coupled with process management, application integration software is an important component in the real-time enterprise technology toolbox.

Business process management also offers benefits in real-time business intelligence applications, since it can be used not only to document business process flows, but also to identify key points in those flows that should be monitored and managed.

Business activity monitoring solutions offered by application integration software and independent BAM vendors recognize the importance of process management, but so far BI vendors have been slow to understand the significance of process management for real-time performance management, and have taken little action here.

[3] D. McCoy. "Business Activity Monitoring: The Promise and the Reality." Gartner, July 2001.

2. The Business Case for Real-Time BI

Balancing Business Benefits and IT Costs

Real-time BI enables organizations to react more quickly to changing business circumstances and requirements, but the business benefits and ROI must be balanced against the IT costs of making the BI system more responsive. To understand this trade-off we need to examine the components of a real-time BI system, and how they affect the business decision-making process.

Real-Time BI Systems Reduce Action Time

"There will always be delay between a business event occurring and appropriate business action being taken," says Richard Hackathorn, President of Bolder Technology. "In a business intelligence environment, this delay has three components: data latency, analysis latency, and decision latency. Data latency is the time between the business event occurring and when the data about the event is ready for analysis in the data warehouse. Analysis latency is the time required to analyze the data and send the results to the user. Decision latency is the time required for the user to understand the information and take the appropriate action." This concept is depicted in Illustration 2.1.

Illustration 2.1

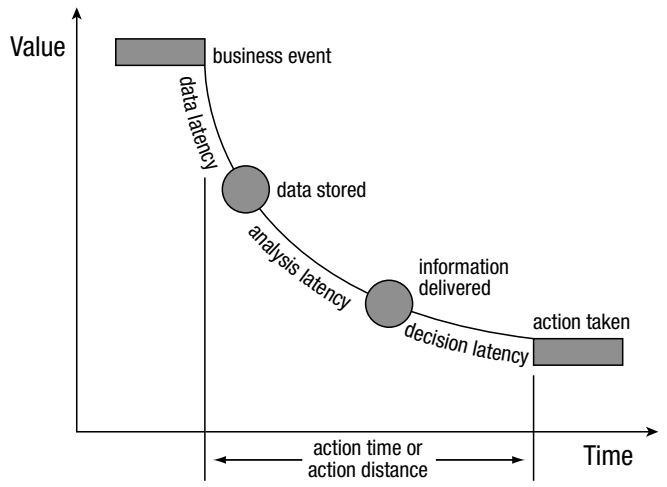

Latency in business decision making. Based on a concept [4] developed by Richard Hackathorn, President of Bolder technology. Reproduced by permission.

The sum of the three latency periods is the *action time* or *action distance,* and the objective of a real-time BI system is to reduce the action time required to respond to a business event. Data latency and analysis latency can be reduced by technology innovation, but reducing decision latency depends on the business user. Providing the user with better information will of course help reduce decision latency. Rather than simply telling users a problem has occurred, for example, it is of more value to also provide detailed analyses and suggest actions. Another approach to reducing decision latency is to replace manual user decision making with rules-driven BI processes that automatically take action on behalf of the user.

[4] Richard Hackathorn. "Minimizing Action Distance." The Data Administration Newsletter, Issue 25.0, July 2003.

Illustration 2.2

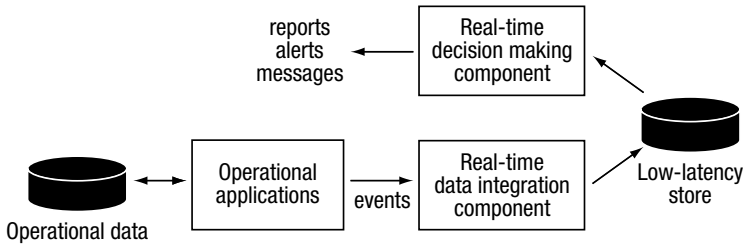

Real-time BI processing components.

Two Main Components: Real-Time Data Integration and Real-Time Decision Making

A real-time BI system has two main components: real-time data integration[5] and real-time decision making (see Illustration 2.2). The objective of the real-time data integration component is to capture business events from operational systems and integrate them into a low-latency store. This component supports the real-time data-on-demand type of real-time processing defined earlier. The real-time decision making component, on the other hand, supports real-time performance management and real-time predictive analysis.

In our survey, 50 percent of the 846 respondents either had deployed, or were planning to deploy, real-time BI projects. Of these, 21 percent were doing real-time data integration, 15 percent were doing real-time decision making, and 57 percent were doing both. Follow-up interviews indicated that many of the companies doing just real-time data integration were simply loading their data warehouses faster to alleviate performance and operational batch window problems. Organizations doing just real-time decision making, on the other hand, were usually using data in operational systems, rather than in a low-latency store, for rapid decision making. Overall, then, about a quarter of companies either had implemented, or were planning to implement, a full real-time BI system.

The survey results also demonstrate that some people still think real-time BI is about loading data into a data warehouse faster. Although this helps reduce data latency, it does nothing to reduce analysis or decision latency.

Real-Time BI Is *Not* about Loading the Data Warehouse Faster

"A real-time BI system is not about keeping data current in the data warehouse, and then telling the user to access it. It's about delivering the right information at the right time to the right business users to enable them to react rapidly to solve business problems," says Phillip Gollhofer, Manager of Business Intelligence at BNSF.

Although the ROI realized by reducing the action time to a specific level is the main factor in justifying real-time projects, this ROI must be balanced against the increased IT costs of achieving the reduction. The IT costs will depend on things like the amount and type of business information that has to be accessed and analyzed, and the BI system performance and availability requirements of business users. As shown in Illustration 2.3, the smaller the action time required, the more expensive the IT solution.

[5] The survey used the term real-time data warehousing to describe this component, but in this report we use the term real-time data integration instead, since this more accurately describes the function of this component.

Illustration 2.3

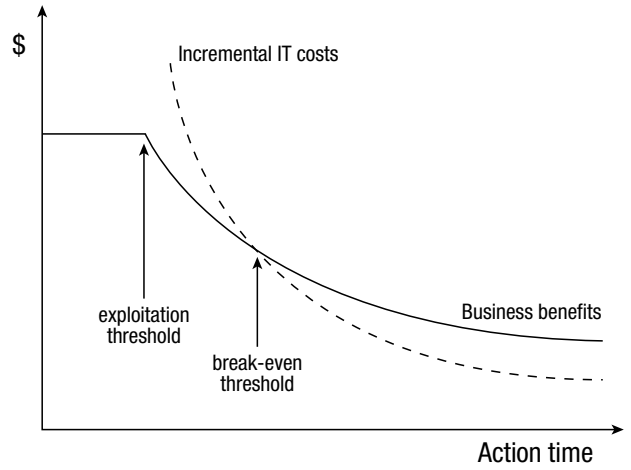

Real-time BI: Action time versus IT costs.

The Smaller the Action Time Required, the More Expensive the IT Solution

It should also be recognized that the business ROI achieved by reducing action time depends on the organization's ability to modify its business practices to take advantage of the improved responsiveness in the IT system. Several organizations we interviewed said this was an important success factor in realizing the benefits of a real-time BI system. There is a point beyond which reducing the action time any further has no business value because the increased responsiveness cannot be exploited by the business. This is shown as the *exploitation threshold* in Illustration 2.3.

"Based on discussions with business users, we chose to update the real-time component of the data warehouse every two hours," says BNSF's Gollhofer. "This is a trade-off between making rapid decisions, keeping the data stable long enough to enable it to be analyzed by users, and IT costs."

Illustration 2.4

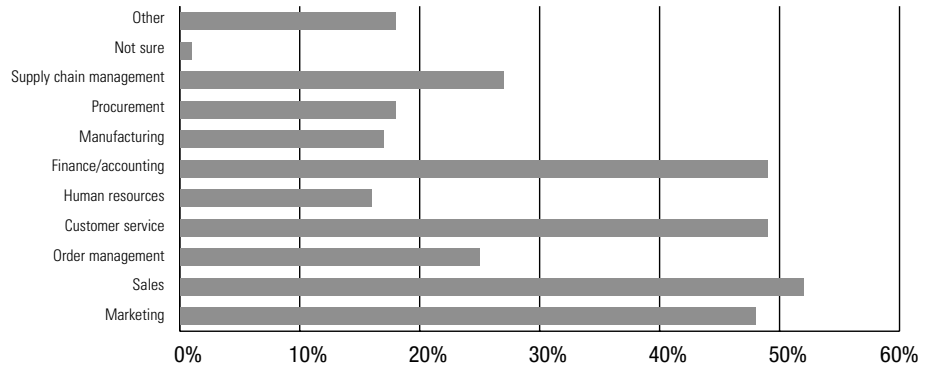

What business areas do or will your real-time project support? Multi-choice question, based on 417 respondents.

Illustration 2.5

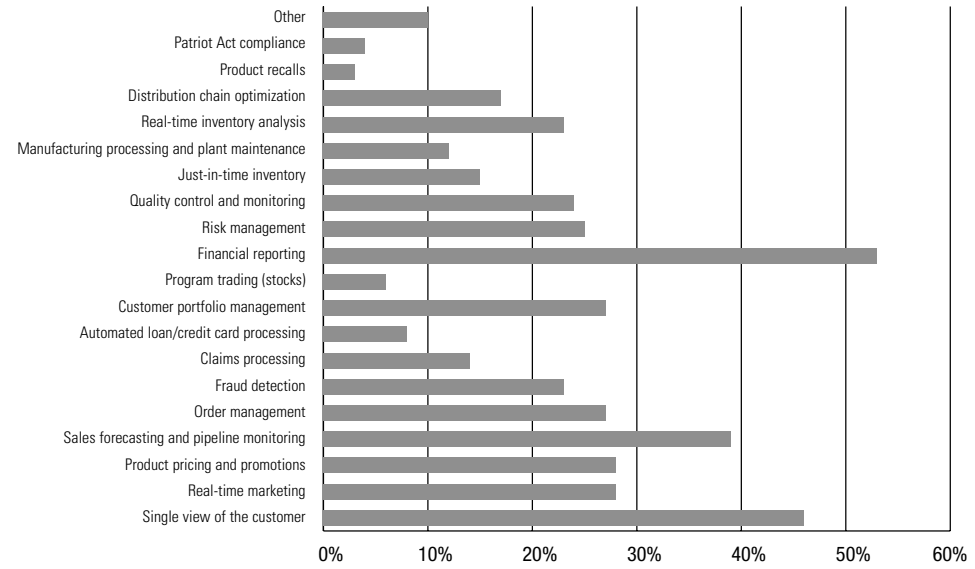

Which real-time BI applications have you implemented or do you plan to implement? Multi-choice question, based on 417 respondents.

"Organizations Should Think in Terms of Right Time"

"People should not think in terms of real time, but should consider how responsive each business process needs to be to satisfy a specific business goal," says Stephen Brobst, CTO of Teradata. "Organizations should think in terms of right time, rather than real time." This viewpoint is shared by David Randall, Pre-sales Director at DataMirror, who states that, "Right time is the appropriate term to use because each business process has its own unique requirement about how fast the organization needs to react to a particular business situation."

Applying Real-Time BI to Business Problems

In our survey we asked companies to identify the business area where they were using, or planning to use, real-time BI. As the results in Illustration 2.4 show, sales, marketing, customer service, and finance/accounting are the four main areas of focus.

Based on discussions with many people in the industry, we put together a list of potential business uses for real-time BI, and asked our survey respondents to use the list to identify real-time applications that were being implemented in their organizations. The results (see Illustration 2.5) showed that a single view of the customer, financial reporting, and sales forecasting and pipeline monitoring were by far the most common uses of real-time BI. Some 10 percent of respondents added several other real-time BI applications to the mix. Illustration 2.6 documents a final list of potential applications for real-time BI processing.

Illustration 2.6

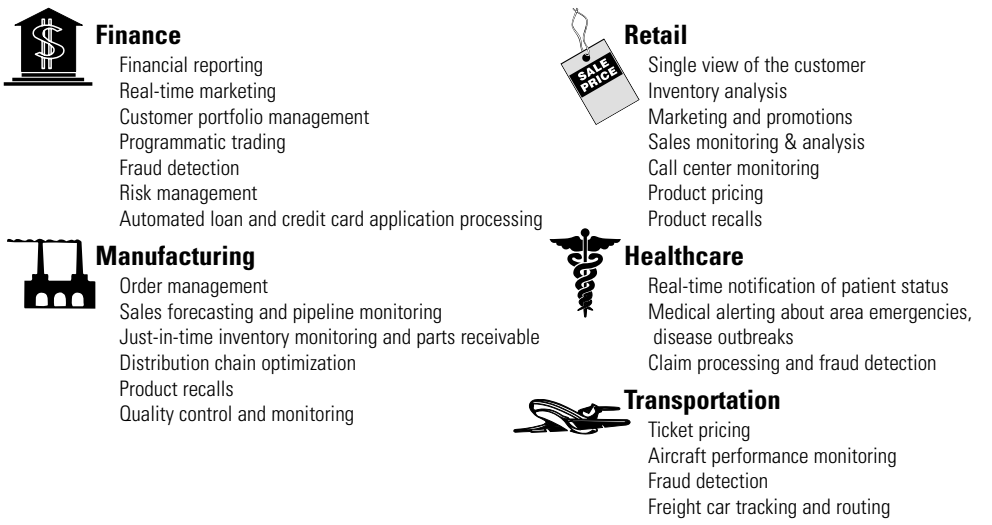

Examples of real-time BI applications.

3. Building a Real-Time Business Intelligence Infrastructure

A wide range of technologies and products support real-time processing in a business intelligence environment. This is a new and rapidly evolving field, and it is important that organizations design a real-time infrastructure that can weather the changes as this evolution continues. We will first look at each of the components of real-time BI processing, and then look at how each of these components can be integrated into a uniform BI infrastructure.

Collecting Operational Data for Real-Time Decision Making

Consistent and Accurate Data Is Vital

Consistent and accurate operational data is what drives real-time BI processing. Query applications use this data to document and display information about current business events. And real-time performance management applications extend this query capability by using information from historical data warehouses, planning applications, and so forth, to put current business events into context, and by alerting users about urgent business issues that require action (see Illustration 3.1). Real-time predictive analysis applications, on the other hand, employ business rules and models generated from historical information to assess and optimize current business events, to detect fraud, for example.

Poor Data Integration Is a Main Inhibitor of Success

After budget and business case issues, poor data integration in operational systems was identified by survey respondents as the main inhibitor (47 percent of responses) to building real-time BI applications (see Illustration 5.1 on page 29). To overcome data disintegration in their operational systems, our survey showed that many organizations (57 percent of responses) are copying operational data into a low-latency operational data store (ODS) to provide integrated data for tracking business activity (see Illustration 3.2). Some companies (18 percent) are building low-latency message-based event stores to monitor business operations. We also

Illustration 3.1

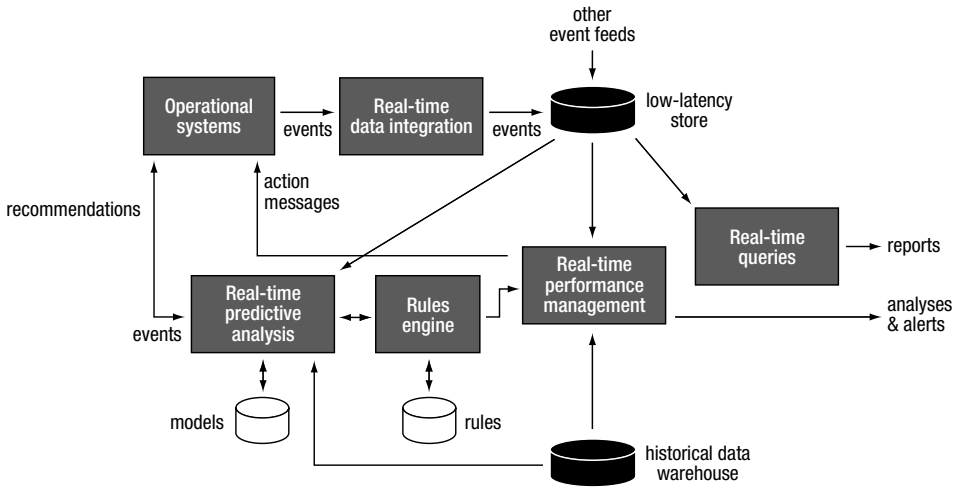

Using low-latency operational data and historical information for real-time BI processing.

asked respondents about the technology they were using to capture data and application events for loading into a low-latency store. Forty-five percent said they were using an extraction, transformation, and loading (ETL) tool; 17 percent enterprise application integration (EAI) software; and 35 percent a combination of both.

It is quite apparent from our survey, and follow-up interviews, that EAI software is likely to become a key component of a real-time BI system—30 percent of survey respondents rate both the need for ETL interfaces to EAI software, and the need for the event-driven streaming of information, as very important requirements for building a real-time BI system.

Interviews indicated that several organizations were dissatisfied with the functionality being offered by ETL vendors for streaming data into a low-latency data store. They also did not want to employ two different products to build such a store.

Illustration 3.2

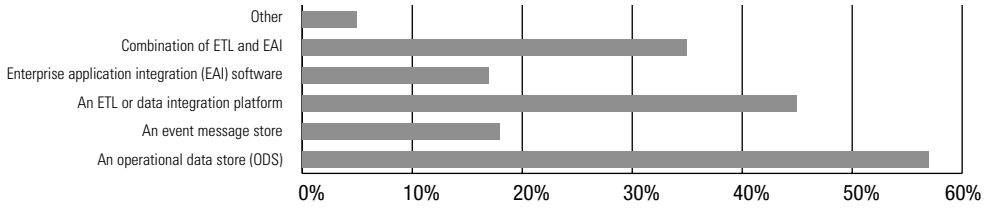

What is the basis of your current or planned real-time data integration environment? Multi-choice question based on 420 respondents.

Instead of using ETL products, these companies had built their own interfaces to operational systems and integration software. We also found that companies that were using a combination of both EAI and ETL software sometimes complained about the performance and reliability of the interfaces between the two.

Pay Careful Attention to the Functionality of New Products

Clearly, software for building event-driven low-latency stores is very new, and organizations need to pay careful attention to the capabilities, performance, and reliability of the products they use for building such stores. We anticipate as the real-time market evolves that companies will deploy products to support a standards-driven approach to building low latency stores, rather than using today's proprietary techniques. Companies are increasingly likely to deploy the application integration services offered by enterprise application server vendors, and then connect this software to data integration products using industry standards such as Web

Illustration 3.3

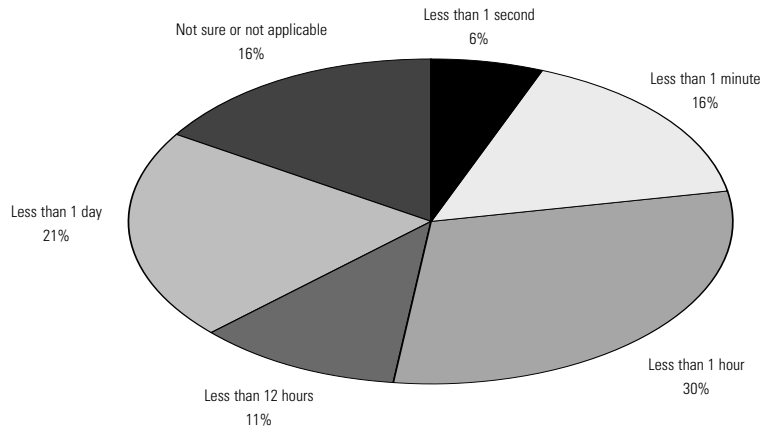

What is the latency of your low-latency store? Based on 419 respondents.

Illustration 3.4

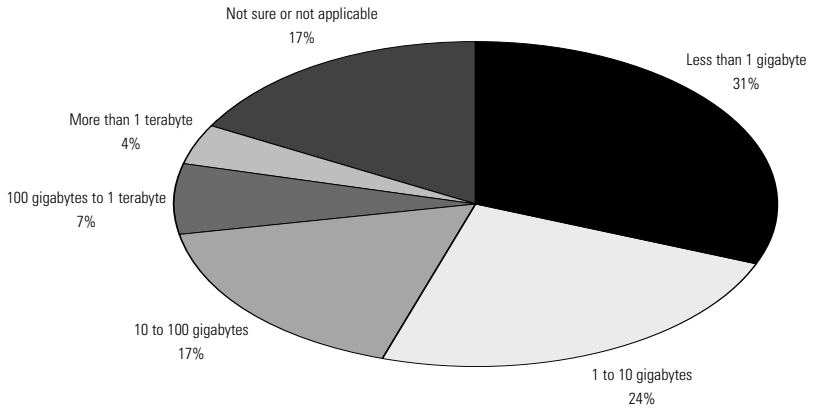

How much raw data is loaded into your low-latency store per day? Based on 418 respondents.

Services, JMS, SOAP, and XML. The application integration software will be used for event capture and routing, and the data integration software will be used for data transformation and cleanup.

"Industry standards will become very important for supporting real-time architectures and for connecting together application and data integration software," says Bob Zurek, CTO of Ascential Software. "Customers need to understand these evolving standards and should carefully evaluate both application and data integration products to determine the level of support provided for these standards."

We also asked respondents how often they were updating their low-latency stores, and how much data was flowing into them each day. Illustrations 3.3 and 3.4 show the results.

Most low-latency data is less than one hour old

Over half the respondents were loading information at a frequency of one hour or less, and 6 percent were loading data at a sub-second interval. The amount of data being loaded per day varied considerably. Some 11 percent of companies were loading more than 100 gigabytes per day, whereas 31 percent were loading less than 1 gigabyte per day. There was no correlation between how much data had to be processed and how often it was loaded into the low-latency store.

The information sources (see Illustration 3.5) used for populating low-latency stores were consistent with those documented in the TDWI report *Evaluating ETL and Data Integration Platforms,* published in March 2003. There were a few important exceptions for real-time processing, however. The sources that showed a marked percentage increase over our previous study were application event data, XML files, Web Services, EAI sources, and real-time devices.

Things to Consider When Building a Low-Latency Data Store

Debates about designing data stores for BI systems are a popular topic among analysts, consultants, and vendors. And opinions differ widely on this subject, especially when it comes to handling low-latency data. Interviews with survey respondents uncovered two key issues concerning the building of low-latency stores: 1) how and where to store low-latency data, and 2) how to maintain data integrity and quality.

Illustration 3.5

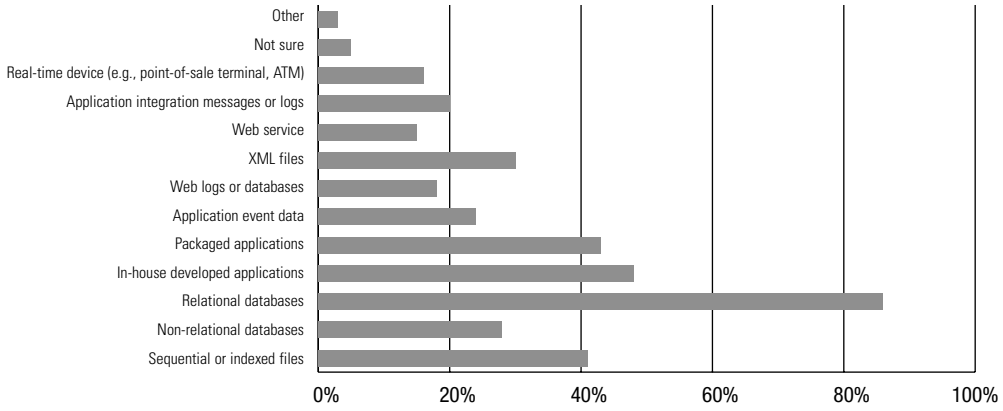

What types of operational sources are used to populate your low-latency store? Based on 421 respondents.

How and Where to Store Low-Latency Data

During follow-up interviews, we asked corporations why they kept their low latency data separate from the historical data warehouse. The reasons in most cases were that the performance and data structure requirements of the low-latency data store did not match those of the historical data warehouse. Some people commented, for example, that performance required them to use a normalized data structure for the low-latency data, rather than a star schema as in the data warehouse. We did come across some organizations that combined low-latency and historical detailed data in the same data store. The data stores in these latter cases, however, were usually normalized.

It you intend to load low-latency data onto the same DBMS server, or into the same database, as that used for your data warehouse, then you need to ensure that the DBMS can scale to support the required workloads. The DBMS should also be able to handle mixed workloads by providing resource scheduling and data management facilities that can be used to control and manage the resources consumed by different workloads, and that prevent long running tasks from interfering with real-time tasks.

It is also important to note that the loading of large-scale low-latency information has the characteristics of an operational processing workload, and a DBMS that is tuned for data warehousing may not necessarily provide good performance for such workloads. For high-volume workloads, a benchmark with live data may be required to determine the right database architecture and products.

How to Maintain Data Integrity and Quality

The design of a low-latency store is also affected by data integrity and consistency issues, and we found in our study that these factors play a large role in determining how low-latency data is handled.

Data Integrity Issues. Concerns about data integrity or consistency arise primarily because streaming data into a low-latency store means the information in the store is continually changing. This is not a problem for applications that require the continuous querying and reporting of business activity—a stock ticker, for example. But it will be an issue for applications that need to see consistent data at specific moments, like the status of a stock portfolio at midday. The frequency at which updates are applied to a low-latency store must take into account not only the latency requirements of a business application, but also its data consistency requirements.

Update Frequency Depends on Several Factors

We said earlier that BNSF updated its low-latency store every two hours. The store is used to monitor business activity, and to alert users when out-of-line business situations occur. The reason for the two-hour cycle is that when a user is alerted about a business issue, the data must remain consistent while the user investigates the problem. The two-hour update cycle is a trade-off between being able to monitor the business in real time, and the time required to handle the issues. Cost was also a factor for BNSF. Updating more frequently than the two hour interval would have required costly enhancements to the availability infrastructure of the IT environment.

Phillip Gollhofer at BNSF commented, "Some users require low-latency operational data, while others require historical information. This, and the fact that the two types of data have different information content, was the reason why the low-latency data is kept in a separate database. The data from the low-latency store is used to update the historical data warehouse at the end of each day. Additional quality checks are performed when the data is added to the historical data warehouse."

Illustration 3.6

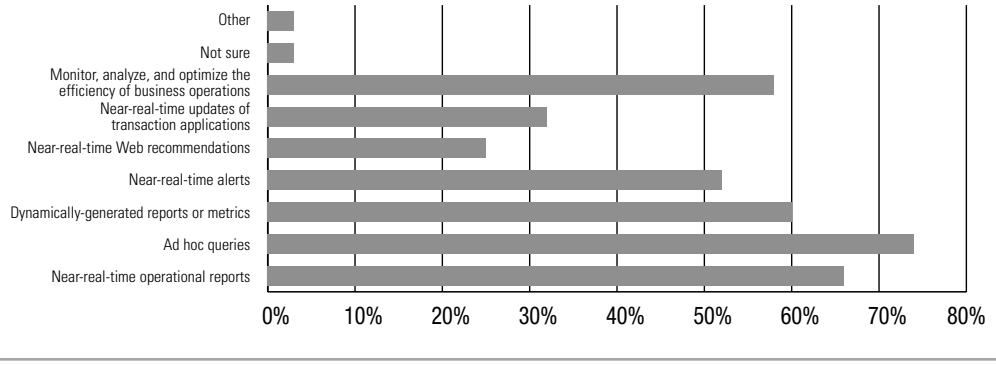

What types of real-time BI decision making has your company implemented or is it planning to implement? Multi-choice question, based on 421 respondents.

Data Quality Issues: As in any business intelligence environment, data quality was identified as a key issue in real-time BI processing. This issue was ranked fifth (43 percent of respondents) in terms of key inhibitors to deploying real-time applications. But it should be noted that there was only a 4 percent difference between this issue and the issue (non-integrated operational data) that ranked second (see Illustration 5.1 on page 29). In a related question, the ability to do real-time data validation ranked third (after performance and high availability) as a very important feature for building a low-latency data store (49 percent of respondents).

A Single Version of the Truth. One interesting item that came to light when discussing low-latency data stores was the number of companies that viewed the low-latency data store as the only place in the organization with a single version of the truth. Several organizations said the low-latency store was the only location that contained integrated data from front- and back-office applications. Surprisingly, some of these companies were updating, correcting, and enhancing the data in the low-latency store, and using it to update operational systems, or to move data from the front office to the back office. This data was also being used to reconcile errors in operational systems, and to identify data quality problems and business rule problems in those systems.

Making Real-Time Decisions

How Is Low-Latency Operational Data Being Used to Make Decisions?

One of the main objectives of our study was to understand how low-latency operational data was being used to make business decisions. Illustration 3.6 shows how various BI techniques are being employed by our respondents. We correlated these responses with those for deploying a low-latency data store and found that, depending on the BI technique being used, between 75 and 85 percent of this real-time decision making was based on information in a low-latency data store.

Operational Reporting Is Alive and Well

Query & Reporting Is the Most Widely Used BI Technique

Not surprisingly, the most widely used BI technique for real-time decision making is query and reporting, which includes *ad hoc* queries (74 percent of respondents) and canned reports (66 percent of respondents).

Real-time queries enable business users to see both granular business information (the status

Illustration 3.7

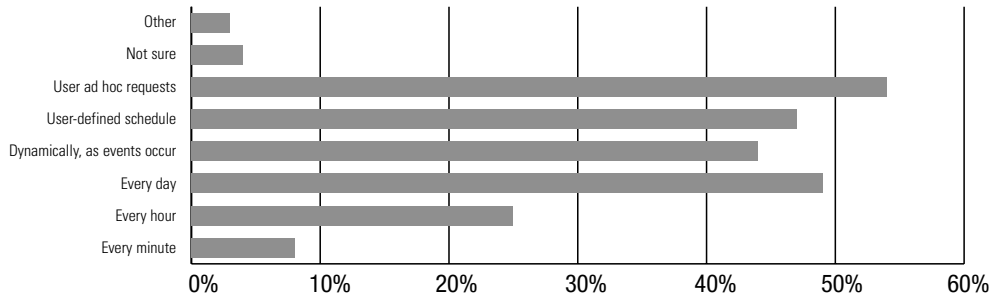

How frequently are reports and analytics generated? Multi-choice question, based on 419 respondents.

of a customer's account across all touchpoints, for example), and high-level key business measures (hourly sales volumes, for example). Some 60 percent of respondents were using dynamically generated business measures to track business operations in real time.

How frequently business reports and measures are produced depends primarily on business requirements. Those requirements also determine the update frequency of the low-latency store. Close to half of the companies surveyed (44 percent of respondents) report information dynamically as business events occur (see Illustration 3.7).

Using Real-Time BI to Manage the Business

Providing business users with more information does not necessarily lead to better business decisions. Few users have time to absorb and analyze the information produced by query and reporting applications. These users need BI that can help them do their jobs faster and more efficiently. They need information put into a business context, and often only need to see this information when some action is required; i.e., they need BI applications to help them not only to measure business performance, but also to manage it.

Manage and Measure. Real-time performance management extends operational query and reporting by relating measures of business performance to specific business goals and objectives. Coupled with alerting technology, these solutions can be used to both measure business performance and manage it.

Business goals and thresholds are defined for performance management applications in terms of business rules. These rules may be generated from historical information in a data warehouse, may come from budgeting and planning applications, or could be a component of a methodology-driven approach to improving business performance such as Balanced Scorecard or Six Sigma.

Real-time performance management offers major advantages over traditional operational query and reporting by reducing the amount of information business users have to process to do their jobs. Performance management coupled with alerts enable the BI system to drive the user, rather than the user having to drive the BI system.

"When we implemented our real-time BI system we had to convince our field staff that we weren't going to just increase their workloads," said Alicia Acebo, Director of Data Warehouse

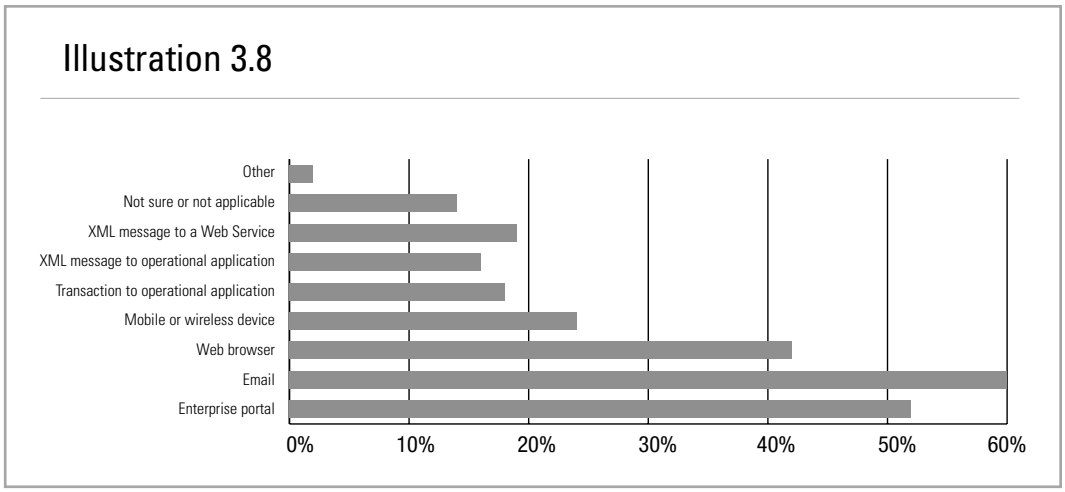

How does your team deliver alerts, actions, and messages? Multi-choice question based on 416 respondents.

"We Reduced Workloads and Increased Customer Satisfaction"

at Continental Airlines. "These people were already very busy, and we had to demonstrate that the alerts and information we were delivering to them were going to make it easier for them to handle customers who were delayed. The result was we reduced employee workloads and also increased customer satisfaction."

Alerting the User. Most performance management applications have some form of exception management capability. When a business goal is missed, or a business threshold is crossed, the exception management mechanism notifies the appropriate users. This notification may be simply highlighting the information and its business rule in a display or report, or it may involve generating an alert to send to the user. Alerting is an important feature of a performance management solution, and is being used by 52 percent of the companies in our survey (see Illustration 3.6).

Exception Management and Alerting Are Two Separate Mechanisms

It is important to consider exception management and alerting as two separate mechanisms. The objective of the exception process is to evaluate business rules against real-time business measurements. If the evaluation results in an exception, then the appropriate information has to be collected and sent to the user. To be of value, the information should not only document that an exception has been raised, but should also include details about where to find more detailed information, or even suggest appropriate actions.

The job of the alert mechanism is to route the exception message to the user. This mechanism should be able to send the message to the right user at the right time and in the right format. It should also be flexible enough that destinations and formats can be modified dynamically as users travel and use different devices.

Illustration 3.8 documents the alert destinations used by companies responding to our survey. Email (60 percent) and an enterprise portal (52 percent) represented the most common destinations. A number of users were also routing alerts to operational systems using application transactions (16 percent) and Web Services (19 percent). These figures indicate that companies are beginning to replace manual processes with automated ones to become more responsive, and to reduce the workload of business workers.

The reason why exception management and alerting need to be separate is because exceptions may be raised in a variety of tools and applications, and a common alert mechanism is required for managing and routing these exceptions to users. Two companies raised this issue

in interviews. In both cases, the common alert facility could handle exceptions raised by any operational or business intelligence application.

If an organization fails to build a common alert mechanism, managing and coordinating the exception process becomes difficult as exception management in BI and other systems increases. Unfortunately, most BI vendors combine exception and alert management into a single facility, which is often difficult to integrate into a common alert framework.

BI Vendors Are Beginning to Support Web Services

Many BI vendors are, however, beginning to support Web Services, which is one way to integrate BI alerts into a common alert architecture. From Illustration 3.8 you can see that 19 percent of companies are using Web Services to handle exceptions and alerts. When developing a real-time BI infrastructure, it is essential that the design include exception and alert management. The use of Web Services should also be evaluated for building the alert architecture.

Building versus Buying

Like most BI projects, you can build or buy your real-time decision-making applications. Some 45 percent of the organizations in our survey were building their own solutions, 61 percent were using commercial reporting and analysis tools, and 33 percent employed packaged analytic applications (see Illustration 3.9).

It was also interesting to note that 31 percent of companies were using EAI tools for deploying real-time decision-making applications. The high numbers for EAI and in-house built applications suggest that business intelligence vendors are not providing solutions that satisfy the requirements of the real-time marketplace.

Where Does BAM Fit In?

Confusion Abounds about BAM

At the beginning of this report we equated *business activity monitoring* (BAM) with real-time performance management. Our survey highlighted the industry confusion about this term. Returning to Illustration 1.1 on page 5 we can see that 48 percent of respondents associated business activity monitoring with real-time BI. Furthermore, 79 percent indicated that BAM was a key requirement (defined as *very important* or *important*) for building a low-latency data store. For real-time decision making, 84 percent of respondents indicated that BAM was a key requirement.

On the other hand, when asked if they were familiar with the Gartner concept of BAM, 71 percent said no; and only 20 percent said they had implemented, or planned to implement,

Illustration 3.9

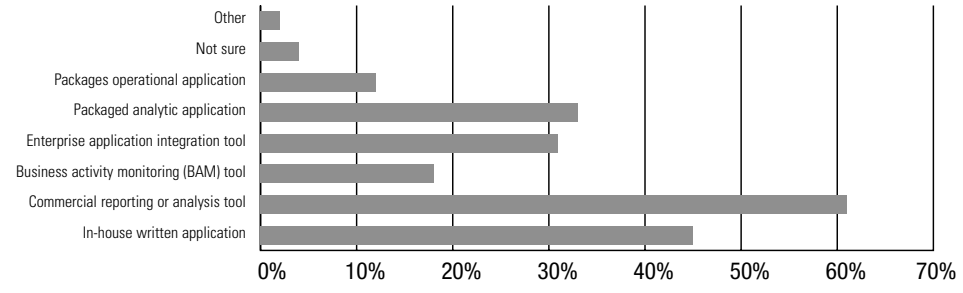

How does your team generate real-time reports and analytics? Multi-choice question, based on 421 respondents.

Illustration 3.10

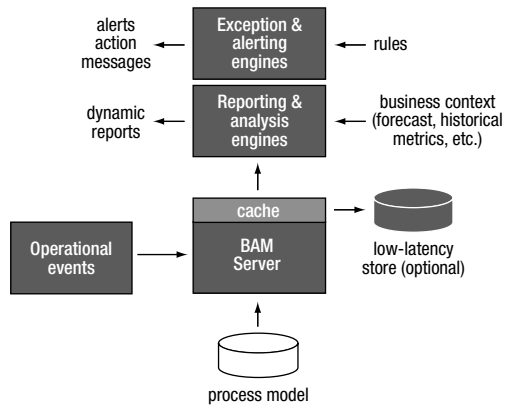

BAM architecture.

a BAM architecture. Lastly, not surprisingly, 61 percent of respondents indicated a need for more understanding and education on BAM.

Our conclusion about these conflicting results is that at a general conceptual level, business activity monitoring is considered synonymous with real-time performance management. At a more detailed business and architectural level, however, the majority of companies don't understand Gartner's concept of BAM, or the fact that BAM products often have a different architecture from BI real-time performance management products.

When to Use BAM

One of the main differences in a BAM environment, compared to a traditional BI environment, is that the server capturing the events from operational systems reports on and analyzes the events as they flow through the server (see Illustration 3.10). Existing historical business intelligence can be used to put the real-time analysis into a historical and business context. The event data processed by the BAM server can be added to a low-latency store, but this is not required. Another difference is that the view of the information being reported on and analyzed can usually be changed dynamically while the server is running. "BAM is an evolutionary extension to existing real-time BI approaches," says Diaz Nesamoney, CEO of Celequest. "BAM helps solve some of the flexibility, latency, event processing, and business process issues that occur in real-time BI applications."

BAM Is an Evolutionary Extension to BI

Events captured by the BAM server may be data events or messages like those captured by an ETL or data integration server, but they may also be events occurring in hardware devices such as an ATM or POS terminal, or application events such as a business transaction, call to another application component, and so forth. Support for industry standards such as Web Services, JMS, SOAP, and XML is crucial for BAM products.

Performance reports and scorecards produced by the BAM server can be sent to a display console, dashboard, or portal. They can also be processed by an associated rules engine that can send appropriate alerts to business users, or action messages and transactions to operational systems.

BAM applications monitor day-to-day business processes such as customer orders, insurance claims, supply chain operations, and so forth. BAM products are typically driven by process models. This is very different from data-driven ETL applications, which have little or no knowledge of business processes. "BAM bridges the gap between business intelligence and business process management," says Jim Cates, CIO and VP of IT of Brocade Communications Systems. "We have a good BI system for analytical processing, but more and more operational data is being copied into the data warehouse for operational reporting, and this is costly. I need the ability to track business processes using just-in-time information. Business process tools offer some solutions, but don't provide the ability to look at the business from multiple dimensions. BAM connects business processes to business intelligence, and gives me the ability to get the best of both worlds. It also gives me the information I need without being forced to store operational data in a data warehouse."

Cates' comments show there are two distinct real-time decision-making business requirements. The first is to monitor specific business processes in real time, and alert users when action is required. Inventory replenishment would be an example here. Often in this situation there is no need to maintain the event data in a low-latency store. The second requirement is to provide an integrated, consistent, and real-time view of disintegrated operational data for querying, reporting, and analysis. A single view of the customer is a good example of this requirement. In this situation, the creation of a low-latency store is essential.

BAM is focused on satisfying the first requirement, whereas low-latency data store technologies target the second requirement. Organizations may need to satisfy either, or both. Regardless, the technology used must fit into an overall BI infrastructure.

The Role of Business Process Management

Business process management provides a development and run-time environment for modeling and deploying enterprise applications. It provides a more productive way of building enterprise systems by replacing traditional data-centric development techniques with a process-driven development approach. Process management software allows developers to model a process, and then deploy and manage it using the services of an underlying run-time engine. Developer productivity increases due to reduced manual coding, and easier application maintenance. The use of business processes as the main building blocks in enterprise systems also makes it easier to capture business expertise and requirements, and thus makes it more likely that an application will match business needs. Process management can be used for business process analysis, integration, automation, or monitoring.

Process management solutions can be used for monitoring and reporting on business events as they happen. Several products also support exception and alert capabilities. These solutions, therefore, overlap in capability with those provided by BI vendors for real-time performance management. As mentioned above in the BAM discussion, process management software at present provides very little support for the business intelligence environment. This gap is being filled by BAM solutions.

Process Management Usage Is Becoming More Prevalent

Our survey showed that organizations are becoming increasingly interested in process management for deploying real-time applications that track business operations from a process-centric viewpoint. Business process management ranked second in our survey as a very important feature for supporting real-time decision making (34 percent of respondents).

As BI becomes more process driven than data driven, we are likely to see integration between business process management, BAM, and BI real-time performance management solutions. It is therefore essential that BI developers and architects become familiar with business process management concepts and technologies.

Process Management Software Solutions. Process management software can be purchased from independent software vendors, but there is a growing trend in the industry for application server suite and application integration vendors to add process management capabilities to their products. In some cases, this is achieved by acquiring an independent process management vendor.

To support business process management, business intelligence and BAM vendors will need to either add these capabilities to their products or create interfaces to existing process management products. Given the trend of BI and BAM vendors to create relationships with application integration and application server suite vendors (who are beginning to integrate process management into their solutions), the latter course is the more likely.

Preparing for Process Management and BAM. At present, the process management, BAM, and real-time performance management marketplaces are changing rapidly. Organizations need to track this evolution and assess its impact on BI infrastructure. It is also important to realize that process management and BAM are not only about implementing technology, but also about the ability to exploit the benefits of these technologies.

Process-driven approaches to BI development can be undertaken in a bottom-up and iterative manner. This enables organizations to gain experience in process-driven technologies, and to gradually evolve to a process-centric design approach. In some companies this is being done by simply building process models and linking them to existing data models. "Understanding your business processes is fundamental to success in real-time BI processing," says Gaz Williams, Senior Consultant at British Telecommunications. "We build models of our business processes, and then link these models to our common information model." Other organizations, however, are using full process management software implementations to drive focused projects. Frequently, this latter group uses application integration software, rather than traditional BI products, to develop real-time performance management applications.

Real-Time Predictive Analysis

Real-time performance management applications report on and analyze business events. These applications *react* to business events just after they have occurred. Another type of real-time BI decision making is predictive analysis. This technology helps organizations *predict* the outcome of a business event before it occurs. An example here is determining the propensity for someone to buy a specific product—this is sometimes called *real-time scoring*. Our survey indicates 25 percent of companies are using techniques like predictive analysis to make recommendations to Web users (see Illustration 3.6).

Predictive analysis solutions are driven by business rules and models that are developed using techniques like data mining. These models are invoked at run time by operational applications. The operational application passes event data to the model, and the model then predicts an outcome. Run-time models can be incorporated into a variety of different software, including rules engines and DBMS products. Some data mining vendors enable models to be expressed in program code that can be integrated into an operational application.

Illustration 3.11

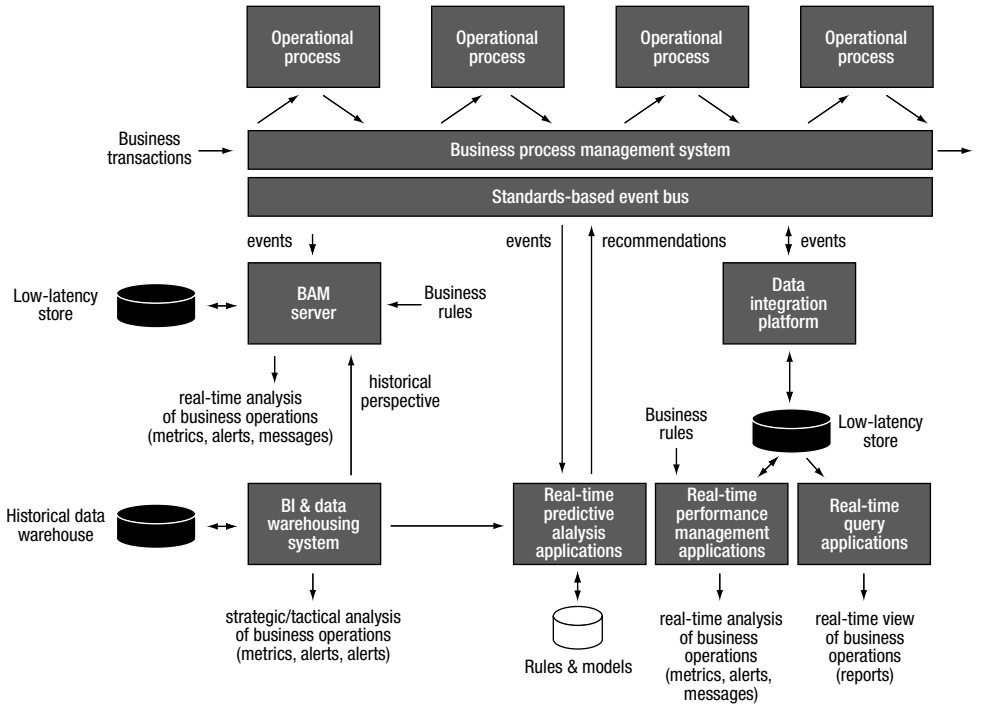

A real-time BI infrastructure. Courtesy of BI Research.

Several leading relational DBMS vendors provide the ability to store predictive models in a database, and invoke them using SQL statements. This is somewhat analogous to a stored query, and provides similar benefits. Stored query users don't have to understand SQL, and stored predictive model users don't have to understand data mining. This approach enables data mining to enter mainstream BI processing.

Predictive Techniques Yield Real Returns

An example of the value of predictive analysis can be found at telecommunications giant Sprint. The company employs a real-time architecture to integrate and analyze data scattered across many enterprise systems. A key feature of this architecture is the consolidation and real-time analysis of call detail records (CDRs). Sprint uses predictive techniques to analyze the CDRs for suspicious patterns and instances of potential fraud. Sprint has narrowed the action time, from the moment of detection to appropriate action, to within five minutes. As a result, losses due to fraud have fallen several million dollars per year.

Putting the Pieces Together

Sound Infrastructure Brings It All Together

Given the diverse components and technologies in a real-time BI system, a good BI infrastructure design is required to integrate and deploy these technologies (see Illustration 3.11 for an example). A sound infrastructure brings together key components that handle the flow of business events from operational applications into a BI system that integrates and analyzes these events, and then transforms them into business intelligence for delivery to business users. Key issues to consider here include scalability, availability, performance, integrated

security, integrated meta data, integrated administration, support for standards (XML, for example) and Web services, and data validation and transformation capabilities. "Once the real-time infrastructure is in place you can do almost any type of real-time BI processing you want," says Alan Carrick, Senior Designer at Standard Life Assurance. "For example, our XML-based messaging hub and common event-based data publishing mechanism gives us a great deal of flexibility for deploying real-time applications."

Survey respondents stated that after budget and poor data integration issues, the lack of an infrastructure for handling and managing real-time processing was the key inhibitor to deploying a real-time BI system (46 percent of respondents)—see Illustration 5.1 on page 29.

Batch Processing Is a Problem in Real-Time Processing

Several companies commented that an integrated and sound BI infrastructure is a key factor in the success of real-time projects. Some of these companies had not built a business case for real-time BI processing, but were instead modernizing their operational systems (replacing batch processing with online processing, for example), and the ability to do real-time BI was an automatic (but essential) by-product of this modernization; i.e., the infrastructure being deployed could handle both operational and BI real-time processing. Some companies also mentioned that dealing with batch processing is one of the biggest headaches in moving to a real-time BI environment.

4. Requirements for Real-Time BI Processing

Real-Time Products Fit into Two Groups

The confusing range of real-time products available gives us many things to consider when selecting real-time BI solutions. But these solutions all fit into the two real-time BI components mentioned earlier: real-time data integration and real-time decision making. Vendors may support one or both of these components.

Illustration 4.1

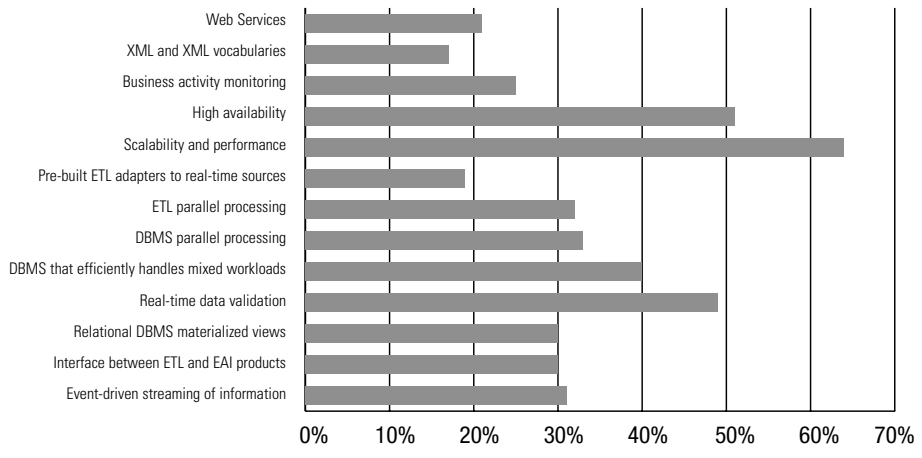

Real-time data integration requirements rated "very important." Multi-choice question, based on 421 respondents.

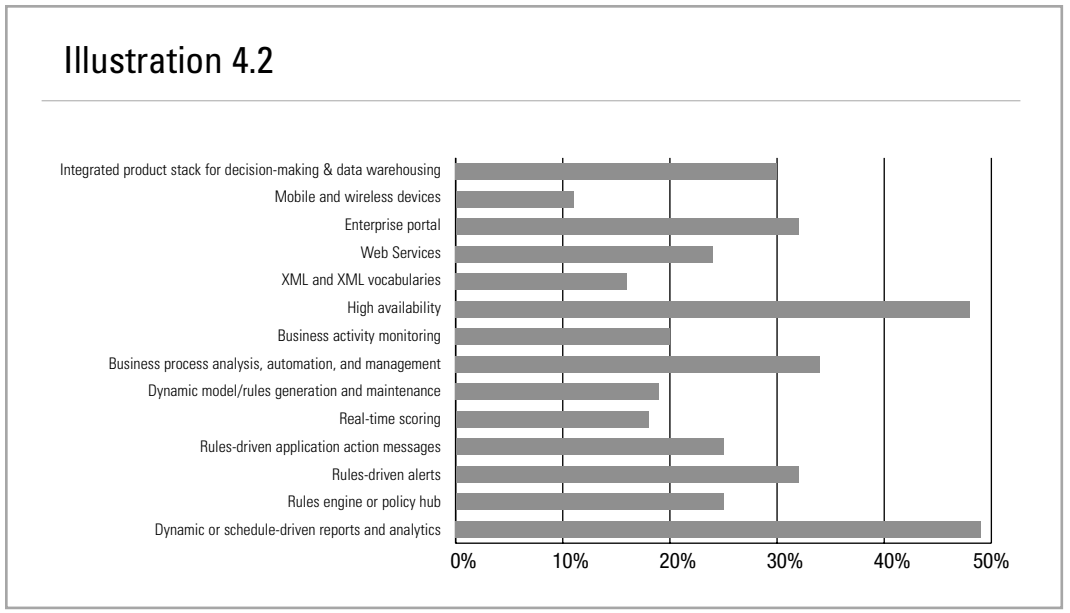

Real-time decision-making requirements rated "very important." Multi-choice question, based on 416 respondents.

The real-time data integration component supports real-time on-demand data, and supplies consistent and accurate business information for the real-time decision-making component. The real-time decision making component provides real-time performance management and real-time predictive analysis. Within real-time performance management, BAM and process management solutions offer several unique capabilities for certain types of decision making.

Business Requirements and User Needs Determine Technology

An organization should develop an infrastructure that is capable ultimately of supporting all of these options. In choosing a specific solution, however, it is important to identify the business requirements and user needs first, and then determine which technology best suits those needs. In some cases, multiple technologies may be required.

We asked companies in our survey to identify their most important requirements for both the real-time data integration and real-time decision making components. The results are shown in Illustrations 4.1 and 4.2, respectively.

Real-Time Data Integration

The report *Evaluating ETL and Data Integration Platforms* addressed requirements for building a low-latency data store. Many of the results in this study are consistent with those obtained in that report.

The top two requirements for real-time data integration in the current study were scalability and performance (64 percent of respondents), and availability (51 percent of respondents).

A Back-End DBMS Must Be Scalable and Flexible

Scalability and Performance. Business intelligence systems are managing rapidly increasing amounts of data, and the use of low-latency data stores will only accelerate this growth. To support real-time processing, the system must have a scalable and flexible back-end DBMS environment for loading and administering large amounts of data. The DBMS must also be able to handle mixed workloads, since the tasks used to update low-latency stores will need to run in parallel with real-time decision-making applications.

Availability. The 24x7 nature of business means users expect business intelligence applications to supply information at any time from any place. They also expect these applications to process data and react to a changing business climate rapidly. The move towards rapid access to business information requires that a business intelligence and data integration system provide high availability. This requires the BI system to exploit facilities in the underlying hardware infrastructure such as clustering, RAID, device redundancy, and so forth. It should also be able to automatically recover from system hardware and software failures by providing, for example, automatic fail-over to a back-up system. Facilities are also required to simplify the deployment of disaster recovery solutions.

Real-Time Decision Making

Most of the product requirements for real-time decision making have already been discussed in this report. Our survey results indicate that the main requirement for real-time decision making is dynamic or schedule-driven reporting and analytics. This result indicates that many companies are still focused primarily on real-time operational reporting, and this is consistent with the results in Illustration 3.6, which outlines the types of real-time decision making organizations are doing today.

The results (see illustration 4.2) also show a significant number of companies are planning beyond doing just real-time operational reporting. Companies now require support for business process management (34 percent of respondents), alerts (32 percent of respondents), and enterprise portals (32 percent of responses). If the results for the rating of *important* are added to those for *very important,* then all three of these requirements are needed by more than 80 percent of the companies responding to the question.

It is interesting to note that 30 percent of respondents rated the need for an integrated real-time decision-making and data integration product set as very important.

5. Planning for Real-Time BI Processing

Critical Success Factors

We have already mentioned several issues that companies view as inhibitors to success in deploying real-time BI. A summary of these is shown in Illustration 5.1.

Not surprisingly, business case and budget loom at the top of the list. As with most IT projects, a real-time BI system can be built bottom-up or top-down. Many of the companies we spoke to were building bottom-up, and in an iterative manner. These companies were deploying tactical projects that addressed specific business issues that would clearly provide a good ROI. It was quite apparent in our discussions, however, that although these companies were building their systems bottom-up, a sound infrastructure was also necessary to be successful.

Education

Education Is Also an Important Success Factor

Education of both IT staff (36 percent of respondents) and business users (46 percent of respondents) were also deemed to be success factors in deploying a real-time BI solution. A follow-up question asked respondents to identify areas where further education was needed. The top four responses (see Illustration 5.2) were real-time BI technologies and architectures (65 percent of responses), real-time data warehousing (62 percent of respondents), business activity monitoring (61 percent of respondents), and data quality (58 percent of respondents).

Illustration 5.1

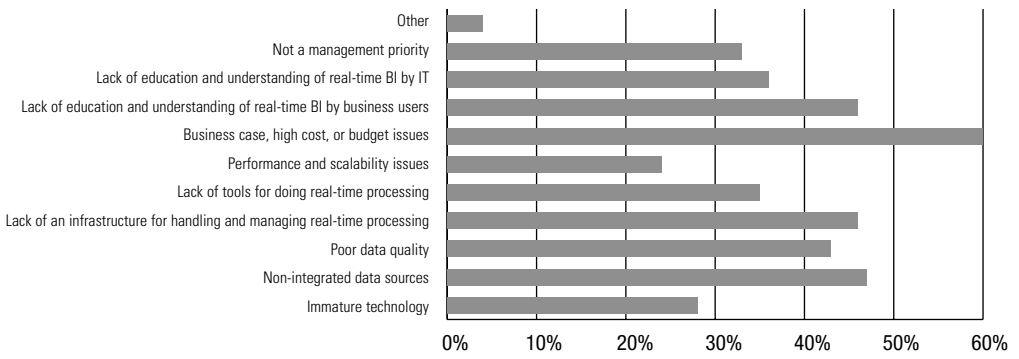

What obstacles inhibit your company's ability to deliver real-time BI applications? Multi-choice question, based on 835 respondents.

Organizational Considerations

From a business perspective, the distinction between operational and decision support processing is largely irrelevant. From an IT viewpoint, however, different groups are often responsible for operational and decision support implementations, and the two types of processing normally have different service levels. The distinction is also important from a skills and education viewpoint because many real-time business applications involve a mix of both operational and decision support processing (often called a *mixed workload*), and it is key for developers to understand the impact of this mixed workload on performance, scalability, and availability, and also on product selection.

If a decision support group is responsible for implementing real-time processing applications, it will be essential for the group to recruit staff with operational processing skills and experience,

Illustration 5.2

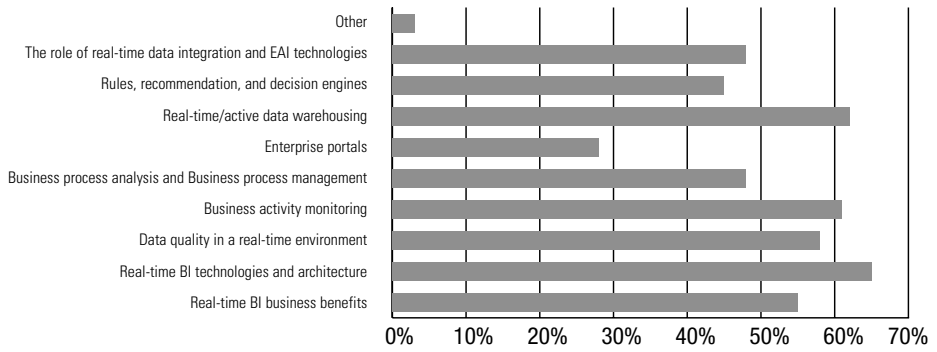

In what areas does your organization need more understanding or education? Multi-choice question, based on 820 respondents.

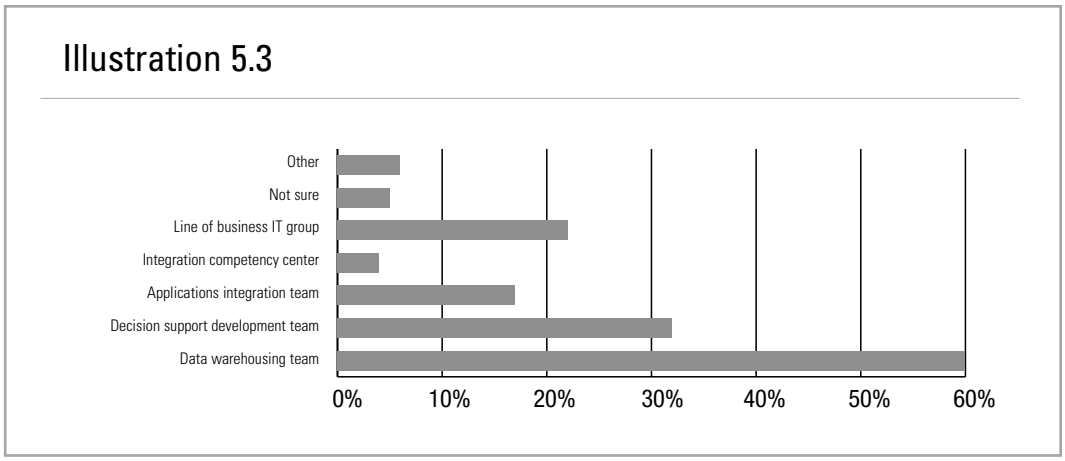

Which group has primary responsibility for real-time projects? Based on 420 respondents.

especially given how important the real-time enterprise environment is to business success. There is a strong argument in favor of setting up a real-time enterprise project office that brings together both operational and decision support staff to guide architecture development and application implementation.

The Distinction between Operational and Decision Support Processing Is Disappearing

Several companies we spoke to said the distinction between operational and decision support processing was beginning to disappear in their organizations, especially when line-of-business managers begin to realize the benefits and importance of using business intelligence to drive day-to-day business operations. In some cases, operational IT groups have merged with data warehouse and decision support groups. One consultant in a large telecommunications company said this was a key success factor to implementing real-time applications.

In our survey, 60 percent of organizations said the the data warehousing team manages the responsibility for real-time BI processing projects, while another 32 percent indicated the decision support team holds that responsibility (see Illustration 5.3). Many respondents, however, marked several options here. Subsequent interviews revealed that although the data warehousing group was driving the project, the group was very reliant on line-of-business and applications groups to provide the business expertise required for the project to be successful.

Conclusion

Given that real-time BI processing is likely to become a key component in your organization's future business intelligence framework, how do you prepare for it?

First, it is important to gain a good understanding of real-time BI technology and its business benefits. Whereas real-time BI can provide significant business ROI, it is important to realize that it can be expensive to implement, particularly in a high-volume and low-latency environment.

Of course not all real-time BI applications require latency of a few seconds. Latency of several minutes, or even hours, may be quite adequate for many implementations. Regardless, a clear business case and ROI must be established for a project. It will, therefore, be important for IT managers to discuss the benefits of real-time BI with line-of-business managers to determine where real-time BI would bring the most business benefit, for example, customer relationship management, procurement, manufacturing, distribution, etc.

Second, real-time BI applications must be integrated with existing business intelligence and enterprise application integration solutions. Without this integration, yet another information silo will be created. It will be essential to understand the facilities being offered by vendor real-time BI products, and what interfaces are provided for enterprise and BI integration. Also, given the importance of BI to a real-time enterprise, it will be crucial to ensure that the key components of the real-time BI environment can provide the required performance, scalability, accessibility, availability, and security.

Weigh Business ROI against Implementation Expenses

Third, it is important to realize that, although a number of real-time BI implementations have already become successful, this is still a leading-edge technology that will undergo rapid evolution over the next few years. It will be important to track this evolution and update real-time strategies and frameworks accordingly.

Last, it is crucial to recognize that a real-time enterprise is not just about technology. For a real-time enterprise and real-time business intelligence to be successful, organizations must modify their business practices and educate business users about real-time solutions in order to exploit and to gain maximum business benefit from real-time enterprise initiatives.

A Final Comment ...

"To be successful in real-time projects you need to understand your business processes, the roles associated with those processes, and the information that is required for each role," said Brocade's Cates. "More data doesn't lead to better decisions. I often think the reverse is true. I intend to invest in solutions that enable me to provide answers to the top 10 business questions for each business role of a business process. If I can do this, then I know the investment will pay for itself. From a technology perspective, you need to understand the impact of data on the business process before collecting it, and need to ensure that the data sources used to create information are accurate enough to make decisions in real time." ■

SPONSORS

Ascential Software Corporation
50 Washington Street
Westboro, MA 01581
800.966.9875 or 508.366.3888
Fax: 508.366.3669
Email: info@ascentialsoftware.com
Web: www.ascential.com

Ascential Software Corporation is the leading provider of enterprise data integration solutions to organizations worldwide. Customers use the Ascential Enterprise Integration Suite to integrate and leverage data across all strategic analytic and enterprise applications with confidence in the accuracy, completeness and timeliness of critical information. Ascential Software's powerful data profiling, data quality, data transformation, parallel processing, meta data, and connectivity solutions enable customers to reduce total cost of ownership and increase return on IT investment. Headquartered in Westboro, MA, Ascential Software has offices worldwide and supports more than 2,200 customers in such industries as financial services, telecommunications, healthcare, life sciences, manufacturing, consumer goods, retail, and government. More information on Ascential Software can be found on the Web at www.ascential.com.

Business Objects
3030 Orchard Parkway
San Jose, CA 95134
408.953.6000
Web: www.businessobjects.com

Business Objects is the world's leading provider of enterprise business intelligence (BI) solutions. Business Objects enables organizations to track, understand, and manage enterprise performance. Business Objects provides the industry's most integrated business intelligence suite, called BusinessObjects Enterprise 6. This suite includes the industry's best Web query, reporting, and analysis; the most advanced and complete suite of analytic applications; and the best connectivity to packaged applications. Business Objects has more than 17,500 customers in over 80 countries. Business Objects can be reached at 408.953.6000 and www.businessobjects.com.

Celequest Corporation
555 Twin Dolphin Drive
Redwood City, CA 94065
650.598.0300
Fax: 650.595.9700
Email: info@celequest.com
Web: www.celequest.com

Celequest has developed a powerful, new architectural approach for Business Activity Monitoring (BAM) that dramatically improves enterprise agility, productivity, and responsiveness for the real-time enterprise. Critical application areas for the company's technology include inventory tracking, margin optimization, risk management, logistics monitoring and compliance, enabling faster response to a wide spectrum of critical business events. Founded in 2002, Celequest was recently awarded the prestigious Investor's Choice Award at Enterprise Outlook naming it one of the Top Ten Companies "most likely to succeed." Celequest is a privately held company funded by Bay Partners and Lightspeed Venture Partners and based in Redwood Shores, California.

DataMirror
3100 Steeles Avenue East,
Suite 1100
Markham, ON
L3R 8T3, Canada
905.415.0310 or 800.362.5955
Fax: 905.415.0340
Email: info@datamirror.com
Web: www.datamirror.com

DataMirror is a leading provider of enterprise application integration and resiliency solutions that give companies the power to manage, monitor, and protect their corporate data in real time. DataMirror's comprehensive family of LiveBusiness™ solutions enables customers to easily and cost-effectively capture, transform and flow data throughout the enterprise. DataMirror unlocks the experience of now(TM) by providing the instant data access, integration, and availability companies require today across all computers in their business.

Over 1,700 companies use DataMirror software, including Debenhams, Energis, GMAC Commercial Mortgage, the London stock exchange, Oshkosh B'Gosh, Priority Health, Tiffany & Co., and Union Pacific Railroad. DataMirror is headquartered in Markham, Canada, and has offices worldwide.

Informatica Corporation
2100 Seaport Boulevard
Redwood City, CA 94063
650.385.5000
Fax: 650.385.5500
Web: www.informatica.com

Informatica Corporation is the leading provider of business analytics software that helps Global 2000 organizations monitor, manage, and optimize the performance of key business operations across the enterprise. Today, more than 1,600 companies worldwide have chosen Informatica to power their analytic solutions—resulting in millions of dollars in cost savings and productivity gains. For more information, call 650.385.5000 (800.970.1179 in the U.S.), or visit the Informatica Web site at www.informatica.com.

Teradata, a division of NCR
1700 S. Patterson Boulevard
Dayton, OH 45479
937.445.5000
Web: Teradata.com

Teradata, a division of NCR, offers powerful enterprise analytic technologies. Companies use Teradata solutions to get a single, integrated view of their businesses so they can make better, faster decisions that drive growth and profitability.

Companies can leverage this integrated view to identify opportunities to increase revenues, decrease costs, and improve relationships. At the same time, Teradata solutions are more cost efficient, simpler to manage, and capable of growing with the needs of the business.

With our proven expertise in enterprise data warehousing and best-in-class technology, Teradata can help you see your business like never before. To learn more, go to Teradata.com.

The Data Warehousing Institute
5200 Southcenter Blvd., Suite 250
Seattle, WA 98188
Local: 206.246.5059
Fax: 206.246.5952
Email: info@dw-institute.com
Web: www.dw-institute.com