TDWI

MONOGRAPH SERIES

JUNE 2006

# A Business Approach to Right-Time Decision Making

By Wayne W. Eckerson

Director of Research and Services, TDWI



tdwi
THE DATA WAREHOUSING INSTITUTE

## About the Author

**WAYNE W. ECKERSON** is the director of research and services for TDWI, a worldwide association of business intelligence and data warehousing professionals that provides education, training, membership, certification, and research. Eckerson has 17 years of industry experience and has covered data warehousing and business intelligence since 1995.

Eckerson is the author of many in-depth reports, a columnist for several business and technology magazines, and a noted speaker and consultant. His book, titled *Performance Dashboards: Measuring, Monitoring, and Managing Your Business*, was published by John Wiley & Sons in October 2005. He can be reached at weckerson@tdwi.org.

### *About This Monograph and Our Sponsor*

A TDWI monograph is a sponsored thought piece written by a TDWI research analyst to explore business and technical issues on a specific topic within the business intelligence and data warehousing field. This monograph was sponsored by Cognos Incorporated, which reviewed the manuscript and provided input prior to publication. The concepts and opinions expressed in the monograph as well as the text and charts were developed solely by the author and are based on many years of experience and research into business intelligence and data warehousing topics.

### *About TDWI*

The Data Warehousing Institute (TDWI), a division of 1105 Media, Inc., provides in-depth, high-quality education, training, research, and certification for business intelligence (BI) and data warehousing (DW) professionals worldwide. TDWI can help your BI team stay abreast of new and emerging trends and techniques and gain the skills it needs to deliver effective BI and DW solutions. Through our membership program and regional chapters, TDWI can also help you and your team establish a network of peers in the industry to whom you can turn for assistance and advice on solving thorny technical and organizational problems, and on career development.

 

## Executive Summary

Companies are increasingly competing on the velocity of their information delivery to business users. Being able to respond quickly to business events can spell the difference between success and failure.

But today information is locked up in multiple systems, making it hard for users to access the data they need when they need it. Although data warehouses have done a good job of pulling together large volumes of historical data, they are usually not architected to deliver just-in-time information. This information is usually maintained in operational systems that are accessed via legacy or ERP reporting tools.

The consequence is that business users often must learn several different tools for accessing critical information, which is a tall order for most time-constrained users. Alternatively, they must rely on professional developers to create reports for them. Unfortunately, custom reports are often outdated before they arrive and don't anticipate future questions that users might have, so they offer minimal reuse. With two unacceptable options, users often make decisions with incomplete information or none at all—choosing to rely solely on gut instinct.

To unlock the information in operational systems and provide users with a powerful blend of historical and current information, organizations should consider using BI products that incorporate EII technology. This technology can query data across multiple systems—including data warehouses, operational systems, Web services, and external data sources—in real time and deliver it to a report or performance dashboard for display.

While not ideal for every situation, EII-enabled BI products can simplify data access for large numbers of users, while blending historical and operational data in a transparent way. These products can also help organizations identify commonly queried data that is best managed in a data warehouse to improve performance and minimize impact on operational systems and corporate networks

## Business Context: The Need for Speed

In today's economy, the speed of information access or delivery often spells the difference between success and failure. Advertising campaigns by major software vendors and industry slogans have crept into the consciousness of corporate executives looking for a competitive advantage:

- The On Demand Business™ (IBM)
- The zero latency enterprise (Hewlett Packard)
- Information at your fingertips (Microsoft)
- 360-degree view of the customer (CRM vendors)

Today's business executives are turning these mantras into reality by deploying analytical systems that deliver the right information to the right users at the right time.

The upshot of this information arbitrage is that organizations optimize performance, because workers can correct problems and exploit new opportunities before it is too late to affect bottom-line results.

**"We leverage velocity as a competitive weapon."**

Consider Quicken Loans. They now provide mortgage consultants, managers, and systems administrators with performance dashboards that deliver up-to-the-minute information on the flow of leads, calls, channel productivity, and systems performance, among other things. Performance dashboards have improved operational efficiency, optimized sales, and better aligned operations with the company's strategy and culture.

"At Quicken Loans, we leverage velocity as a competitive weapon. Our operational dashboards help us meet the needs of our information-hungry corporate culture," says Eric Loftstrom, manager of business intelligence.[1]

**Requirement for Doing Business.** Organizations like Quicken Loans that display critical information to workers quickly in an easy-to-digest format are reaping considerable benefits. However, in the near future, the delivery of just-in-time information to all workers will change from being a lever of competitive advantage to a requirement for doing business.

## How Fast Do You Want to Go?

**The value of information changes with the timeliness of its delivery to decision makers.**

**Historical versus Operational Reporting.** The value of information changes with the timeliness of its delivery to decision makers. Most organizations deliver reports that summarize past business activity. This historical information helps managers and executives understand the effectiveness of prior strategies, plans, and initiatives and also identify emerging patterns and trends that could affect future plans.

While historical information is critical, it doesn't help organizations identify problems or opportunities as they occur, or take immediate action to optimize results. With the pace of business quickening, companies can't always afford to wait until the next day, next week, or next month to review progress. They need more timely information to monitor, manage, and optimize operational processes that affect bottom-line results.

**Proactive Intervention.** Hence, organizations like Quicken Loans are building new decision-making systems that deliver just-in-time information and insights to business users who manage key operational processes. These systems provide users with timely information gleaned from one or more applications or systems so they can intervene when necessary to optimize operational processes. The systems empower users to make better decisions and improve business performance.

[1] From Wayne Eckerson, *Performance Dashboards: Measuring, Monitoring, and Managing Your Business*, (John Wiley & Sons, 2005), p. 128.

**Business users want to view just-in-time information in context.**

**Balancing Speed and Context.** But most business users don't want either historical or just-in-time data: they want a combination of both. For instance, a sales manager may want to supplement a traditional sales report with just-in-time information to improve a forecast or understand the immediate impact of a historical trend. Conversely, managers of operational processes must compare current events to historical trends to avoid making knee-jerk reactions. For example, a store manager should avoid marking down slow-selling items before fully analyzing seasonal sales trends.

**Right-time versus Real-Time.** The question that executives and managers need to ask about just-in-time analysis is "How fast do we want to go?" In other words, how fast should information be delivered to decision makers to optimize the value to the business? There is no easy answer here. This value depends largely on individual business processes and the amount executives are willing to pay for the timely delivery of that information. In general, costs increase with the speed of information delivery.

To understand the relationship between the value of information and its latency (or delivery time), we can examine a chart created by Dr. Richard Hackathorn of Bolder Technology. (See illustration 1.) His chart shows that the time elapsed between when an event occurs and an ensuing action can be divided into three segments, each of which adds latency to the decision-making process. In other words, it takes a certain amount of time to capture a business event and turn it into data, another segment of time to deliver the information to business users, and a final segment for the user to take action in response to the data. The greater the latency between event and action, the less valuable the data becomes.

**Framework for Right-Time Reporting and Analysis**

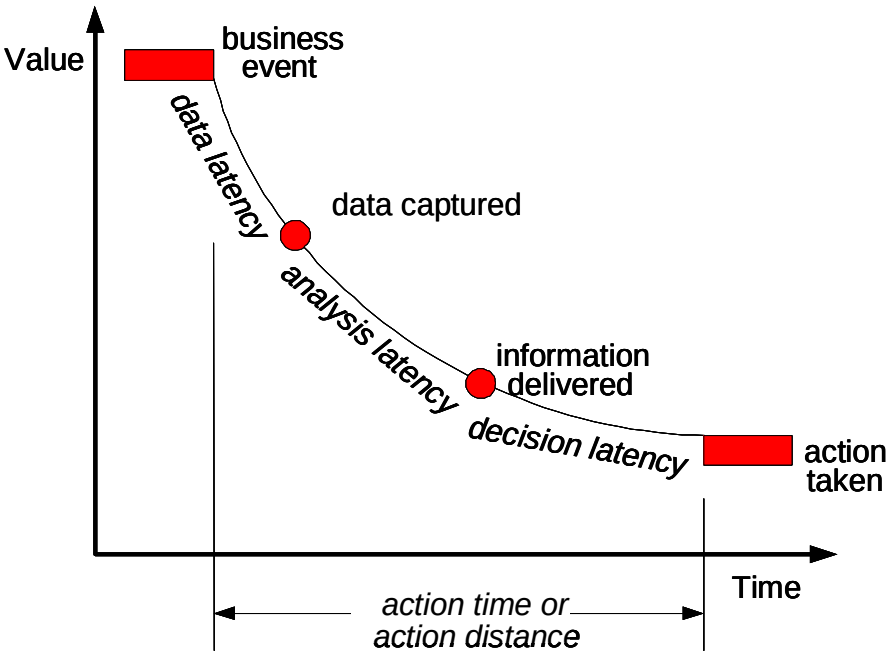


*Illustration 1. Organizations must manage three distinct processes that create latency in a right-time reporting and analysis environment. Courtesy of Richard Hackathorn, Bolder Technology.*

**Most users refer to the optimal rate of information delivery as "right-time."**

Most organizations now refer to the optimal rate of information delivery as "right-time." However, "right-time" is defined in the eye of the beholder—or more accurately, the business executives and managers who manage the processes and pay for the decision making systems.

For example, airline executives looking at operational data interpret "right time" as anything that happens within 14 minutes because the airline industry defines "on time" as 14 minutes or less from an aircraft's scheduled departure or arrival time. But this perception can change depending on what type of data they examine. When considering financial data at the end of a quarter, executives may view "right time" as data that is one hour old.

It's also important not to confuse right time with real time. From a technical perspective, real time systems deliver data to decision makers as soon as an event occurs. While securities brokers may need real-time data about trades and prices, few applications require the instantaneous delivery of information, even for a fraction of the information users require. In the end, organizations must evaluate how quickly they need to respond to different types of events and data and how much they are willing to pay for the appropriate level of delivery.

## Right-Time Reporting Challenges

### *1. Data Is Spread across Multiple Systems*

A major challenge to creating right-time reports is that information is usually scattered across a company in different systems and applications. Pulling both operational and historical data together into a single report that enables workers to monitor and optimize business operations is difficult.

**Too Complex for Business Users.** For example, a retail manager who wants to maximize sales and minimize stock outs and inventory carrying costs needs to gather up-to-the-minute data from point-of-sale, merchandising, and inventory systems throughout the company. The manager also needs to monitor marketing campaigns and track competitor announcements by accessing other sources of information to maintain the proper level of merchandise on the shelves.

**Today's challenge: users need to know where the data is, how to access it, and merge it into a single report.**

To gather all this information today, the retail manager would need to know where the data is, how to access it, and how to merge it into a single report. More than likely, the retail manager would need to learn how to use a half-dozen different tools—including potentially IT-oriented ones—just to gather the data. Then, they would need to understand the nuances of each data set and basic report writing skills before they can blend the data into a coherent whole.

Clearly, these skills and tasks are well beyond the range of any normal retail manager. By the time the manager figured out how to create such a report, the data would no longer be valid and the manager would have wasted an inordinate amount of time on tasks not germane to core responsibilities. More than likely, the manager would ask an IT developer to create the report. But the developer, too, would have to go through the same tedious steps, which may not be a good use of time either. Moreover, since many reports focus on a narrow set of data or limited ways of visualizing, grouping, or calculating information, developers may have to create new reports each time users want to investigate something new or visualize it in a different way.

**Business Users Can't Wait for Custom Reports.** Business users usually don't have time to wait for an IT developer to create a custom report. For instance, a financial analyst who needs to deliver a presentation to the executive committee wants to supplement historical data about sales expenditures found in the data warehouse with up-to-date data found in multiple general ledger systems throughout the company. To obtain the latest financial data, the analyst needs to know how to access each general ledger system, understand the nuances of each system's data, and learn how to accurately merge the data into the existing report. Again, this is a lot to ask from a business user who is not getting paid to manage information systems or create analytic applications.

**Merging historical and operational data must be brain-dead simple.**

**Dangers of Not Collecting the Data.** Unless the process of merging historical and operational data is brain-dead simple, most business users will forego the effort. They will simply not do the analysis, or make a half-baked decision based on partial information, or rely totally on intuition to make a decision or projection to the detriment of the company. The worst scenario is when a few ambitious business users try to cobble the data together but make errors along the way because they do not understand the nuances of the SQL query language or semantic discrepancies within the data itself. When this happens, the business users and the company risk making costly decisions with inaccurate data.

## *2. Re-architecting a Data Warehouse Is Costly*

**Data Warehouses to the Rescue?** The traditional way to integrate information stored in multiple systems for reporting and analysis purposes is to build a data warehouse. Unfortunately, most data warehouses extract historical information in large batch jobs, usually over night or on the weekends. Retrofitting a data warehouse to support right-time data feeds can be very complex and expensive—it's almost comparable to changing the engine of an airplane in flight!

**Retrofitting a DW to support right-time data is like changing the engine of an airplane in flight!**

However, when the value of the right-time information is high enough, executives will step in to fund the development of a new application that requires the creation of a right-time data infrastructure. For instance, Basel II and other regulations are forcing many large financial services firms to do a better job managing financial exposure and risk across multiple product lines and divisions. But rather than retrofit an existing data warehouse or data integration infrastructure, most are building new risk management applications supported by new right-time data architectures. In other words, organizations are leaving data warehouses intact to provide historic context for decision making, but developing new right-time architectures to support just-in-time analysis.

**Flexibility Required**. Another challenge is that the sources of right-time information are not constant. The right-time information that workers need today will differ from the right-time information they will want to examine tomorrow. Continually changing business requirements will require IT departments to source right-time information from different internal applications as well as external ones, such as Web Services, Web sites, and syndicated data sources (e.g. Nielsen data.) Organizations that have developed rigid architectures for delivering right-time information will be forced to continually rewrite the interfaces between systems at great expense.

## *3. Limitations of Current BI Products*

**Too Many BI Products.** When the data infrastructure is too costly to build or retrofit to support right-time information, many companies look to their analytic applications and BI products to pick up the slack. One problem with BI products is that there are too many of them! Research from TDWI shows that companies have an average of

12.9 BI tools from 3.2 different BI vendors.[2] And these numbers don't include the reporting tools that come with operational applications.

**Today, workers use different BI tools to query historical and operational data.**

Typically, workers use one type of BI tool to access historical data in a data mart or data warehouse, an ERP tool to view operational reports, a custom-built IT tool to access legacy systems data, and Excel or Access to grab everything else. This situation forces users to learn multiple sets of tools to query and analyze both historical and analytical data, which few users are capable of doing. This approach also does nothing to reconcile the semantic differences among data in different systems, placing users in the perilous position of having to translate between potentially a dozen ways of defining a customer or calculating sales and so on. Obviously, this increases the risks of making poor decisions based on inaccurate data.

**Operational Applications Lack Robust Reporting Functionality.** Most operational applications bundle in a reporting tool and canned reports. Unfortunately, these embedded reporting tools only report on data contained within the operational application, which usually only holds about 30 days of transactions. In other words, these embedded BI tools can't place operational data in historical context, leaving users with a very limited view of data. In addition, the tools aren't very flexible or easy to extend or customize.

**Limited BI Semantic Layers.** Ironically, many BI products started life as query/reporting tools that ran directly against operational systems. To shield business users from the complexity of operational systems, BI vendors provided a rich semantic layer that transformed complex operational schema into business-friendly query objects, such as "customer," "region," and "net sales," which users could select to compose complex queries.

Today, however, most reports are optimized to run against a single data source, such as a data warehouse or data mart, not multiple systems. Although the rich semantic layers in most BI products are good at producing business-friendly views of data, their servers provide little or no optimization for distributed queries that run against both historical and right-time data across multiple systems. To rub salt into the wound, BI developers also need to create and link multiple semantic layers, each of which are typically bound to a single data source. This semantic merger can be cumbersome and difficult to maintain.

## Reporting Transparency: The Key to Success

**A Better Way.** To overcome these and other challenges, we need a better way to provide business users with a quick and seamless way to combine historical and right-time data. The solution needs to be simple, repeatable, and cost-effective. Users should not be forced to ask the IT department to create a custom report, nor should

[2] See Wayne Eckerson and Cindi Howson, *Enterprise BI: Strategies and Technologies for Deploying BI on an Enterprise Scale*, TDWI, 2005. www.tdwi.org/research/reportseries

the IT department be required to re-architect a data warehousing environment to handle right-time or operational data. And users should not have to learn multiple tools to create reports that combine data from multiple systems or reconcile disparate methods of calculating critical data elements, such as customer, profit, and sale.

In short, we need to provide business users with a single BI solution that provides a single view across multiple systems—both historical and operational—and makes it easy for them to create reports that blend heterogeneous data in a meaningful way. But delivering this degree of transparency is challenging. It requires merging new federated query technology with enhanced metadata, native application interfaces, and dashboard technology.

## *Enterprise Information Integration Technology*

**Goal: "One tool, any data."**

Fortunately, these major components are now available and converging within leading BI products to deliver on the promise of "one tool, any data." One key element is enterprise information integration (EII) technology. EII consists of technologies that have been around for more than a decade, such as distributed query optimization, database federation, and virtual database view technology. The current generation of EII vendors offer greater performance by relying on XML-based technologies like XQuery and XSLT; better design tools, and greater flexibility through a services oriented architecture. Furthermore, EII tools today support a vast array of data sources, such as legacy, relational, and desktop databases, as well as native APIs to complex operational applications, such as SAP R/3 and Siebel CRM, and even data warehouses, data marts, and XML data stores.

**EII technology shields users from having to know and navigate complex back-end systems.**

Today's EII technology also delivers a set of query objects that abstracts the data held in multiple, distributed systems. BI vendors can integrate these objects into their own semantic layers, giving users a single, integrated, and business-oriented view of distributed data. Business users perceive that the data is stored in a single, local database even though it is physically distributed across multiple systems. (See Illustration 2.)

**An EII-enabled Dashboard**

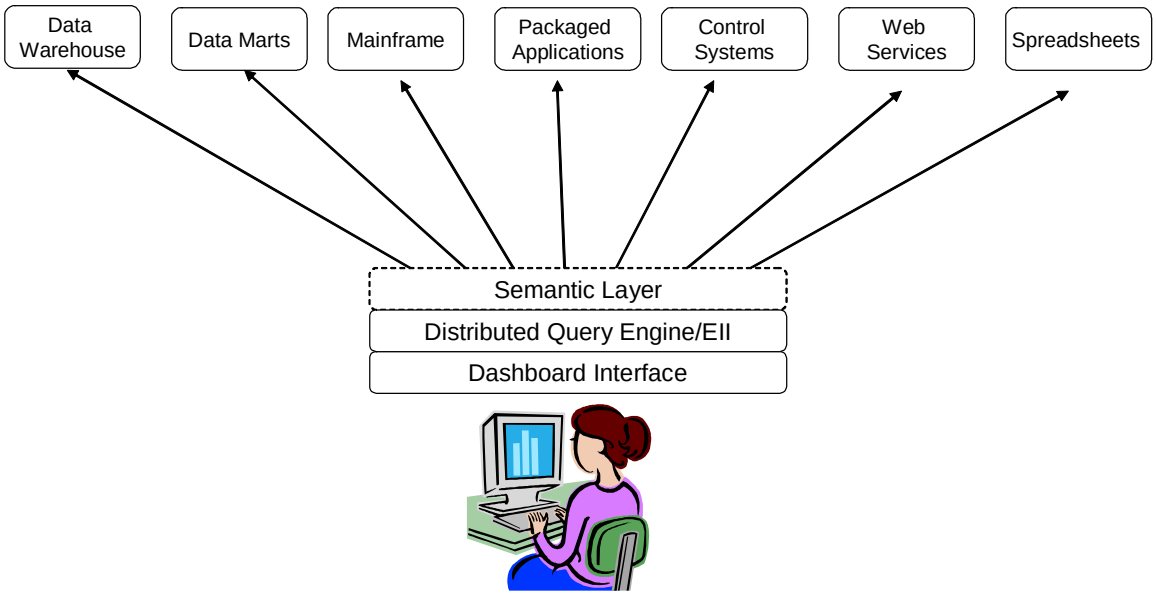


*Illustration 2. Many dashboards use EII technology to pull data from multiple distributed systems and applications in right time.*

EII tools embed intelligence that enables them to process distributed queries in the most efficient manner possible, either by moving and joining data in remote databases and files or bringing it back to the EII server for processing. They also can lightly transform data, reconciling different database and file formats on the fly where necessary. In this way, EII technology can integrate data from multiple databases and feed it to a report or dashboard.

**The merger of EII and BI semantic layers provides a single view of enterprise data.**

**EII-Enabled BI Products.** The truth is that most enterprise BI products have supported many of the pieces of EII for years but in limited form. For instance, some BI products only access data via ODBC and often can't join disparate data without dumping it first onto a disk where an application like Excel or a report is used to merge the data.

To overcome these limitations, some leading BI vendors now embed EII technology. This gives users all the benefits of BI products—the ability to craft ad hoc queries, manipulate, format, and chart the results, and navigate, filter, and reformat predefined reports—with a federated query capability that gives users seamless access to multiple, distributed data sources. Moreover, the semantic layer of EII tools can be merged into the semantic layer of BI products with minimal mapping. This merger of metadata gives users a single view of enterprise data, a truly revolutionary idea!

**Some Limitations.** EII-enabled BI products won't solve every distributed data challenge. In some cases, the data may be too full of errors or inconsistent to be delivered in a usable format on the fly. At other times, there may not be a common set of keys on which to join the data with other data sets, or the volume of data may be too great to ensure adequate performance across a corporate network. In such cases,

organizations may be required to model and process the data using classic data warehousing techniques.

**EII-enabled BI products can provide right-time data that hasn't yet been moved into a data warehouse.**

**Follow-on Benefits.** Despite these limitations, EII-enabled BI products are ideal for supplementing existing reports with right-time data from operational systems that has not yet been extracted and moved into a data warehouse. For example, an EII-enabled BI product may enable a sales manager to view a dynamic sales report at the end of a quarter that shows monthly, weekly, and daily sales by representative using data in the data warehouse as well as up-to-the minute sales by representative for the current day pulled directly from the Siebel CRM system.

In addition, data warehousing designers can use EII tools to prototype what information users may want in a data warehouse before going through the time and expense of modeling and moving the data. By tracking queries and usage patterns, administrators can identify which data to move into the data warehouse and what reports to build.

Also, developers with an EII-enabled BI product can develop new complex analytic applications much more quickly to meet the needs of fast-moving business divisions than if they were restricted to building a data mart or data warehouse as the only means of integrating back-end data. What's more, EII (whether as an autonomous tool or built into BI) is a fraction of the price of a data warehouse technology stack that includes an ETL tool, thus making it a cost-effective adjunct to—but not a replacement for—traditional data warehouses.

**BI-enabled EII products also help to "servitize" back-end data.**

**Data Services Layer.** BI-enabled EII products also help to "servitize" the data. The EII technology creates a services layer that sits between application and data layers in an analytic application, insulating each from changes in the other. This practice eliminates the need for companies to rewrite downstream reports if they swap out databases or add, delete, or revise columns in the database. This layer of abstraction creates a plug-and-play infrastructure that lets administrators upgrade applications and data sources without breaking connections between them.

## Right-Time Delivery Methods

### *Performance Dashboards*

Once an organization decides how much latency of information is optimal for the business processes they want to measure and act on, they need to define the appropriate vehicle for displaying that information to decision makers. In most cases, these displays are reports of some kind. There are many different types of reports: ad hoc, parameterized, operational, master-detail, pixel perfect (for invoices or statements), and so on. However, the most popular reporting method today to deliver right-time information is a performance dashboard.

**A dashboard is essentially an exception report that graphically displays the status of key metrics.**

A performance dashboard is essentially an exception report that graphically displays the status of key metrics using color-coded graphs and symbols as well as traditional charts and tables. Business users love dashboards because they can see at a glance what things need their immediate attention. (See Illustration 3.) They don't have to search through long reports to find the right data, fiddle with features and functions to obtain the views of information they want, or know how to query remote systems to collect the data. And if the performance dashboard is built on a BI platform, they can drill down on an alert in the graphical monitoring layer to detailed data to analyze and report on the root cause of the problem.

### Performance Dashboard Example


*Illustration 3. Performance dashboards let business users check critical activity at a glance and drill down to analyze and report on detailed information.*

**A business dashboard gives users a view of current operations in historical context.**

**Historical and Right-Time Information.** More importantly, performance dashboards are ideal for displaying both historical and right-time information. A dashboard can either blend these different types of information within a single chart or table (which is best done with EII-enabled BI) or display them in separate panes on the screen. Like an automobile dashboard, a business dashboard gives business users a comprehensive view of current operations in context of historical performance. Therefore, performance dashboards—which encompass scorecards—are powerful tools to monitor business activity and keep it aligned with strategic goals and plans.

And it should be no surprise that some performance dashboard products now embed EII technology to gather both historical and right-time information from across the enterprise.

**Alerts.** But performance dashboards aren't the only method for displaying a mix of historical and right-time information. Companies that want to compress decision latency and action time can bypass dashboards altogether and simply deliver alerts about exception conditions to the business users through delivery channels specified in advance by business users, including the Web, e-mail, pagers, or wireless devices. In other words, users don't have to stop what they're doing to get the data; the data finds them wherever they are when something critical needs their attention.

## Conclusion

To make effective decisions, business users need dashboards and reports that blend both historical and the latest transaction data. Unfortunately, few BI products today can generate views that merge data from multiple systems, including data warehouses and operational systems that contain the most up-to-date information. This situation forces users to learn how to use multiple tools or make decisions based on incomplete information.

To remedy this dilemma, organizations should consider using BI products that incorporate federated query (EII) technology that transparently merges data on the fly from multiple systems—including data warehouses, operational systems, Web services, and external data sources. While not ideal for every situation, EII-enabled BI products can simplify data access for large numbers of users and supplement a well-designed data warehouse with right-time and external information.