

What you need to know

Q&A with database expert Richard Winter regarding business intelligence

WHEN IT COMES TO BUSINESS INTELLIGENCE (BI) ON A grand scale, Richard Winter has been there and back. He founded (and now runs) a company, WinterCorp, to help enterprises solve the thorniest BI problems in scalability, performance, availability, architecture, and manageability. He conducts extensive primary research on the practices of those who are building and managing the world's largest databases, revealing their configurations, dimensions, issues, goals, successes, and obstacles.

Earlier in his career, Winter helped create the Model 204 database engine for online complex interactive query, which was the leading product for large-scale complex query for several years, and the platform for many early systems that would now be called data warehouses. He also participated in seminal DARPA-funded research on architectures for data management and data sharing across the Arpanet, predecessor to today's Internet.

So when 24x7 went looking for an expert to discuss this rapidly evolving domain, Richard Winter topped the list. In the following interview, Winter shares his insight and perspective on the critical role of business intelligence in today's data-rich and increasingly competitive enterprise environment.

24x7: *Let's start at the beginning. How do you define business intelligence?*

Richard Winter: BI means making information accessible to people who need to make decisions on

behalf of an enterprise in a timely, accurate, consistent, integrated, and easy manner. When you can deliver the required information in response to interactive queries that may not have been anticipated, as well as through solutions that solve specific recurring problems in the business, then you have a BI capability.

Data can now be integrated across the enterprise as well as across the full range of subjects that the enterprise deals with. It's extraordinarily powerful to have information on every aspect of the business—customers, products, suppliers, finances, human resources, competition—in one integrated repository or available through one seamless user interface. The ability to do this quickly and cost-effectively and to create and deliver new business solutions that exploit the data for analytic and decision-making purposes are ideas that have become achievable for enterprise users only recently.





“BI connotes something **beyond facts and data**. The focus is on information that can be applied to **produce business insight or business action.**”

24x7: How is BI different from traditional data warehousing and decision support?

RW: BI connotes something beyond facts and data. The focus is on information that can be applied to produce business insight or business action. In the early days of data warehousing, the miracle was that retail managers could receive a set of sales reports on Monday morning that were accurate, that were consistent across hundreds or thousands of stores, and that could be aggregated to satisfy every level of management and executive interest.

But the manager's job is not just to know what happened over the weekend. It's to make the right decisions on Monday morning so that the business progresses over the coming days, weeks, and months. To make effective decisions, the manager must be able to understand what is likely to happen, and then determine the actions that will best take advantage of future opportunities. Today BI means that the system and its related applications will help the enterprise react to—and capitalize on—dynamic business conditions.



“You want a **platform that can scale**, because nearly everything in BI starts small but can get **very large.**”

24x7: *What key trends do you see in BI at this time?*

RW: One consequence of the ongoing IT revolution is that devices that capture data have become less and less expensive to produce. Radio frequency identification (RFID) is a good example. An RFID tag attached to a product can report wirelessly the location of the product as it moves through the supply chain. These tags are being attached to cartons and pallets at the point of origin, and they're allowing retailers to track the movement of goods toward the point of delivery. In the past,

this information would have been very costly to capture for every item that moved through a huge retail operation.

Of course, each transmission from an RFID tag generates a record that ends up in a data warehouse and can be used for analysis concerning the efficiency of the distribution system, how effectively stores are being stocked, how well suppliers are meeting their deadlines, and many other questions that might have been too expensive to answer in the past.

This creates a challenge in terms of very rapidly growing volumes of data in data warehouses. In fact, our research shows that in each of the last three 2-year periods, the size of the largest data warehouse in the world has tripled. In that 6-year period, we moved from an era in which the largest data warehouses were a few terabytes. In 2005 the largest database was 100 terabytes of data; soon we'll be at 1,000 terabytes, or a petabyte. Businesses need effective ways to manage all that data.

24x7: *How can enterprises deal with this explosion in data?*

RW: You need a scalable database platform and design in order to ensure that increasing data volumes—as well as growing user populations and heavier workloads—won't max out your data warehouse. Such scalability must be systematically engineered into your BI infrastructure and then proactively managed as usage grows and requirements change throughout the system life cycle. But the effort yields rich rewards. In our consulting engagements, we've seen how a scalable BI infrastructure can enable enterprises to successfully roll out a new product across thousands of stores, launch a multichannel strategy to enter new markets, and accomplish other crucial growth initiatives without disrupting existing operations.

“Today BI means that the system and its **related applications** will help the enterprise **react to**—and **capitalize on**—dynamic business conditions.”

24x7: Where do the new data warehouse appliances fit into the BI market?

RW: The idea of an appliance is to make it very simple for the customer by packaging everything you need for a data warehouse into a single unit. When the client buys an appliance, all required hardware and software arrives preintegrated, pre-installed, and preconfigured. In principle, it's wheeled into the office or computer room, plugged into the network, and the client can start loading data.

The design goal for appliances is that they be easy to acquire, install, and use. They should deliver high performance as a result of their integration and their use of commodity hardware and

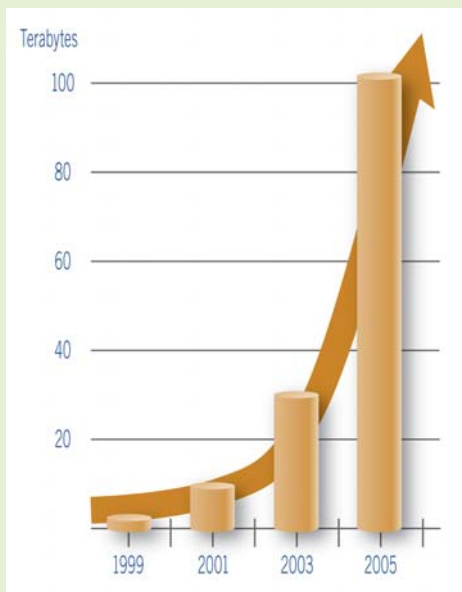
software components, and they should be easy to maintain. The movement toward the appliance concept is an effort to make BI simpler, faster, and less expensive. To the extent that these goals can be met and the promise delivered on, I think customers will be very interested.

24x7: What key factors should an enterprise consider in designing and implementing a BI system?

RW: At the root of the design challenge are the business requirements for the system. The business goals might be to improve customer satisfaction, raise revenues, acquire more profitable customers, or enhance efficiency. Whatever they are, these goals need to be clear and shared by the enterprise, and then the system must be designed to satisfy those objectives.

Concrete business goals must feed into clearly defined engineering requirements for the data warehouse. These requirements are typically quite demanding in terms of response time, performance, data freshness, scalability, and availability—and defining them requires some hard thinking and expert analysis. But by making them unambiguous and quantitative, the enterprise can ensure a design and implementation process that recognizes key risks and resolves them early and successfully.


It also pays to run tests early that realistically and sufficiently explore the engineering requirements that the system will face in production. Time and again in helping our clients conduct such tests, we've seen leading enterprises control major business risks and save millions of dollars—and months of implementation time—by doing the right type of testing early in the process.



The world's largest databases have tripled in size every two years since 1999 according to the recent WinterCorp TopTen Program, which surveys the biggest, most heavily used databases.

24x7: What are the most important characteristics of the platform that underpins an effective BI system?

RW: You want a platform that can scale, because nearly everything in BI starts small but can get very large. If you want to provide inventory management in a chain of hundreds or thousands of stores, usually you start out providing it for a few stores. The worst nightmare would be to launch a system, make a big splash with a successful 10-store pilot, and then find that your system has run out of capacity as you roll out to your remaining stores. This kind



be available; as BI has progressed to support more and more business-critical activities, we've moved into an era in which the data warehouse has to be available 24x7 in many organizations. The ability to provide continuous operation is critical.

24x7: What challenges do enterprises face in implementing a BI system?

RW: The fundamental challenge lies in managing a dynamic environment in which requirements escalate rapidly with respect to scale, complexity, availability, data latency, and other factors. For example, a financial institution might want its call center agents to start cross-selling to customers who have called in for service. These agents have the customer on the phone, have just helped the customer, know something specific about the customer's needs (two children in college, for instance), and have products available that might help.

This is a magic moment. But having a BI system that can analyze callers' specific needs and provide answers in near real time to thousands of agents is a very demanding requirement.

Data warehouse managers are sometimes asked to develop such capabilities quickly and make the new system available across an enormous enterprise. This can entail a major expansion of the data warehouse platform. So that's the kind of challenge customers face: the need to respond quickly to new business requirements that rapidly increase the scale and complexity of the data warehouse operation.

24x7: You've mentioned retail and financial services. Is BI valuable in other industries as well?

RW: Absolutely. Any enterprise, regardless of industry, will enjoy greater success to the extent that it is able to leverage its customer and operational data to improve business decision-making. In manufacturing, BI has been successfully applied in the area of supply-chain management. Having visibility all along the supply chain—with atomic-level data on every vendor, order, part or commodity, raw material, price, and delivery status—makes it possible to have the right product at the right price at the right time for manufacturing.

BI is involved in healthcare. For example, it can help providers deliver care for chronic conditions in



of thing can happen when projects are undertaken without a scalable platform underneath or without a scalable database design.

In addition to being scalable, the platform has to perform. It must be able to respond rapidly to a wide range of queries that can be simple or complex, and it must be able to incorporate new information rapidly. Of course the platform also has to

“As BI has progressed to support more and more **business-critical activities**, we’ve moved into an era in which the data warehouse has to be **available 24x7** in many organizations. The ability to provide **continuous operation is critical.**”

a consistent, medically effective, and cost-effective way. People are treated for high blood pressure, diabetes, and heart disease over decades, and there are enormous differences in the practices used from one doctor or hospital to another. BI helps employers, insurers, payers, and providers identify best practices for delivering high-quality care without wasting money.

In telecommunications, churn is a big issue. Creating strategies that enable telecommunications providers to acquire and retain profitable customers is a major focus area for BI in telecommunications. In the public sector, BI makes it possible to service citizens efficiently and cost-effectively. In law enforcement, BI enables integration of information from multiple sources to optimize resource allocation and criminal investigation. Fraud is a huge problem in public sector programs that involve payments, and BI has been very effective in these areas. There are many more examples across the full range of industry sectors.

24x7: How can an enterprise measure the success of its BI system?

RW: There are two basic domains of measurement: technical and business. On the technical side, you need key performance indicators that are aligned with the requirements for the data warehouse. These might measure such things as data latency, query response time, and availability.

Measures of the business effectiveness would relate to the business value delivered, and this can be quantified. In the example I provided earlier, where call center agents at a financial services organization present real-time offers to customers on the phone, the enterprise can track offer acceptance and revenue generation. In other words, you can measure the business effectiveness by measuring the business process that’s enabled by the BI capability.



RICHARD WINTER is president of WinterCorp, an independent consulting firm that helps enterprises increase the performance, scalability, and value of terabyte-scale databases. Founded in 1992, WinterCorp has architected many of the largest databases in production today.

At WinterCorp, Mr. Winter guides clients through database platform selection, benchmarking, design, and engineering management. He also directs the TopTen research program, the only validated survey of the world’s largest databases.

Mr. Winter is an international expert in his field with over 25 years of experience. He holds a Bachelor of Science degree in engineering from the University of Michigan and serves on the faculty of TDWI (The Data Warehousing Institute).

24x7: How would you summarize the business value of an effective BI system?

RW: Business goals range from very concrete, measurable things like increasing revenue or customer satisfaction to enhancing the quality of medical care. In general, BI is about making more accurate, timely, and effective business decisions to achieve these goals. In a vibrant BI program, the people running the business learn through contact with the data and also through examples of what the data warehouse can do for them. As they learn, they are constantly thinking up new ideas to take advantage of their BI capabilities.

The most successful BI programs enable companies to lead their industries with products and services and ways of doing business that are more effective than their competition, because they’re able to make better decisions. That, I believe, is the ultimate goal of business intelligence. ♦