

Enterprise Business Intelligence

May 2006

Prepared by

Dr. Claudia Imhoff
Intelligent Solutions, Inc.



Intelligent Solutions, Inc.

Solve your business puzzles with Intelligent Solutions

What is Business Intelligence?

Most corporations today are inundated with data – from their own internal operational systems, their vendors, suppliers and customers, and from other external sources like credit bureaus or industry sales data. The problem with understanding where your company is going is not in the amount of data coming into it. The problem is that this tidal wave of data is not in a form that can easily be digested, comprehended, or even accessed. Ask simple questions like who are your best customers or what are your most profitable products and you will most likely get as many answers as there are employees. Not a comforting position to have in today's era of economic stress.

This is where Business Intelligence or BI comes in. The goal of BI is to provide the enterprise with a repository of “trusted” data – data that can be used in a multitude of applications to answer the questions about customers, products, supply and demand chains, production inefficiencies, financial trends, fraud, and even employees. It can be used to flag anomalies via alerts, provide visualization and statistical models, and understand the cause and effects of decisions upon the enterprise. Just about every aspect of an enterprise's business can benefit from the insights garnered from BI.

The formal definition for BI is that it is an environment in which business users receive data for analysis that is:

- **Reliable** – The data has been documented as the certified or approved data for the enterprise. The business users are confident that the data is the best that it can be and that it suits their decision-making purposes.
- **Consistent** – The processes that deliver the data to the business community are well documented; there are no surprises like missing or inaccurate data in the mix, analytics that won't run, response times that are unpredictable.
- **Understandable** – The data has been defined in business terms; calculations and algorithms are easily accessed for comprehension. These are documented in a data dictionary or metadata repository that is easy to access and understand.
- **Easily manipulated** – You don't need a PhD in statistics to get sophisticated analytics delivered to your fingertips. And it is just as easy to change the question or set different parameters to twist and turn the data in ways unimaginable just a few years ago.

Why do we build such an environment? The answer is so that enterprises can understand:

- **Where the business has been** – Historical perspective is always important in determining trends and patterns of behavior.
- **Where it is now** – Current situations are critical to either modify if not acceptable or encourage if they are trending in the right direction.
- **And where it will be in the near future** – Being able to predict with surety the direction of the company is critical to sound planning and to creating sound business strategies.

Evolution of BI from Strategic to Operational

In today's BI environment, we now find three forms of BI – strategic, tactical, and operational. It was not always so. BI started off as a strictly strategic or tactical process in nature. Operations were completely separated from the analytic BI environment. And with good reason! Massive strategic queries involving hundreds of thousands of records would have destroyed the performance of the critical operational systems, doing single record OLTP queries and updates. Until recently, technology was not available to support a “mixed work load” – that is, the ability to do massive analytical queries and short, transaction queries in the same environment.

BI implementers found a ready audience in business analysts. These business users had tried for years to create an environment in which they could determine trends, study patterns, and analyze period over period data points. Operational users already had their order entry systems, general ledger and billing applications, transaction processing environments and did not see a need for BI yet.

Therefore BI vendors focused on the analysts, with great success. In the past decade and a half, we have overcome many of the very difficult, onerous problems associated with integrating data for analytic purposes. We now have mature technologies to extract, transform, and load (ETL) data into data warehouses and marts. We have sophisticated data quality tools to clean up and help develop the elusive “one version of the truth” for our enterprises and we have established BI applications to support just about any strategic or tactical BI purpose.

As enterprises begin to see the success of these BI environments, they are demanding that the same capabilities – integrated data, easy to use access tools, affordable, storage and database capabilities – be made available for the operations of the business. This is a major shift from traditional BI and yet another hurdle for most technologists and BI vendors. The ability to handle a “mixed workload” such as this is not easy to accomplish.

BI is being used by more and more of the business community members who have a very operational focus. It is only natural that they need data closer to “real time”. Traditional strategic and tactical BI applications are certainly still needed but operational BI is emerging as an important new capability, particularly in ensuring that the business strategies are aligned with their execution. Table 1 illustrates the differences between strategic, tactical and operational BI. It is very important for the BI implementers to ascertain carefully what data must be a part of the operational BI environment. Otherwise, we may burden the BI environment unnecessarily with difficult ETL processing for data that may not need it.

	Strategic BI	Tactical BI	Operational BI
Business focus	Achieve long-term organizational goals	Conduct short-term analysis to achieve strategic goals	Manage daily operations, integrate BI with operational systems
Primary users	Executives, analysts	Executives, analysts, LOB managers	LOB managers, operational users and systems
Timeframe	Months to years	Day(s) to weeks to months	Intra-day
Data	Historical metrics	Historical metrics	Right-time metrics

Let’s Have a Reality Check, Please

The more we speed up the data acquisition and integration process to support right-time and alignment goals, the more complex the environment becomes. Let’s face it – putting up-to-the-second data into the hands of the entire corporation is expensive, creates a major burden for IT, and should require doing a cost-benefit analysis of how fast is fast enough. We need to

take another look at the real-time scenario with a more comprehensible and pragmatic approach. What corporations need is not universal real-time data delivery but rather “right-time” or “on-demand” data delivery.

The right-time data delivery process is a continuum that is a mix of instantaneous, rapid intermittent, and longer batch-type processes – each yielding a different delivery timeframe: sub-second to other intervals such as a few seconds to several hours to overnight or longer intervals. For example, a securities trader will need immediate access to stock market data. Credit card approval may appear to have immediate access to customer data but in reality it takes a few seconds for approval to occur. Order fulfillment information may be generated once or twice a day and mailing lists may be generated once a month or longer depending on the timing of marketing campaigns. These are right-time deliveries of data that are completely appropriate for their particular processes. Yet they may appear to the employees using them to be real-time. In reality, the applications generating these bits of data most likely use a mixture of real-time, sporadic, and longer-cycle data delivery processes.

If you accept the premise that right-time data delivery is the correct way to go for your enterprise, the challenge becomes to properly identify the time continuum for all business processes, that is, which processes need to be accelerated and why. For example, you may need to speed up data delivery because your enterprise is at a competitive disadvantage due to its decision making processes being too slow. For example, the time it takes from order to fulfillment may be off the mark and the enterprise loses business to companies that can deliver customer products faster. The first step then is for you to perform a baseline assessment of your existing data delivery capabilities (e.g., available technologies, maturity of the BI architecture, existing personnel, etc.) combined with a solid understanding of the business requirements for right-time data. It is also important to understand which weaknesses discovered in the assessment will be *exaggerated* as you speed up the enterprise.

The technology assessment consists of documenting what is in place that can be increased in velocity versus what new technology you will need to put in place to accomplish this increase. ETL tools, DBMS characteristics, data quality software, networks, even query and data access tools should be included in this assessment.

The determination of the business need may be more difficult. It is recommended that you start with a good definition and a common understanding of what right-time data delivery is to ensure business community understanding. From this understanding, you should develop scenarios showing how operational BI can be combined with operational processing to create a smarter enterprise. Let’s look at some examples of these types of work flows.

- Implementation of corporate strategy – How many of us have been told by telecom or utility company customer service representatives that an installer will be at our residence within a four to six hour window? It is certainly not my choice to sit around for 4 to 6 hours waiting for the installer and, even worse, sometimes they don’t even show up! With GPS tracking systems and coordination with the operational systems, these companies can do a much better job of scheduling these resources to meet promises made to customers. They know where the installer is, know about problems like traffic delays, and narrow the waiting timeframe to perhaps less than an hour.
- Capture the customer’s attention at the appropriate times – The window of opportunity to offer customers the next best product is while they are standing in front of your sales clerk or visiting your web site. It is not mailing them an offer long after they have departed either channel. Operational BI must be integrated into the operations processes of the enterprise to enable this form of immediate feedback like suggestive selling to the POS clerk or to the ad slots on the website.

- Align all employees with changing directions of the company – It does no good if analysts determine which product or products will sell well during the Holiday season if the dock worker doesn't get the rush order message and loads other products onto delivery trucks instead. He may continue to do things the same way he always has without realizing that he is jeopardizing significant revenues for the corporation simply because he does not have access to the appropriate shipping information.
- Compliance problems go undetected – These can be breaches of service level agreements or violations of legislative compliance. In either case, front line workers may be approaching such a violation and may not be aware of it without proper operational BI being accessible to them.

BI implementers can use these types of case studies to determine which components of their BI architecture should be altered, replaced, or rebuilt to ensure faster delivery and integration of BI. We are fortunate today to have a variety of integration technologies (EII, ETL and EAI) to help deliver data and BI information to front line workers and to have database engines (mixed work load technologies) to support their requirements.

Benefits from Right-Time BI

Richard Hackathorn of Bolder Technologies developed an interesting graph to demonstrate the value of operational BI. Figure 1 shows the three latencies that impact the speed of decision making. These are data, analysis, and decision latencies. **Data latency** is the time duration to make data ready for analysis, i.e., the time for extracting, transforming, and cleansing the data, and loading it into the database. All this can take time depending on the state of the operational data to begin with.

Analysis latency is the time from which data is made available to the time when analysis is complete. Its length depends on the time it takes a business to do analysis. Usually, we think of this as the time it takes a human to do the analysis, but this can be decreased by the use of automated analytics that have thresholds. When the thresholds are exceeded, alerts or alarms can be issued to appropriate personnel, or they can cause exception processes to be initiated with no human intervention needed.

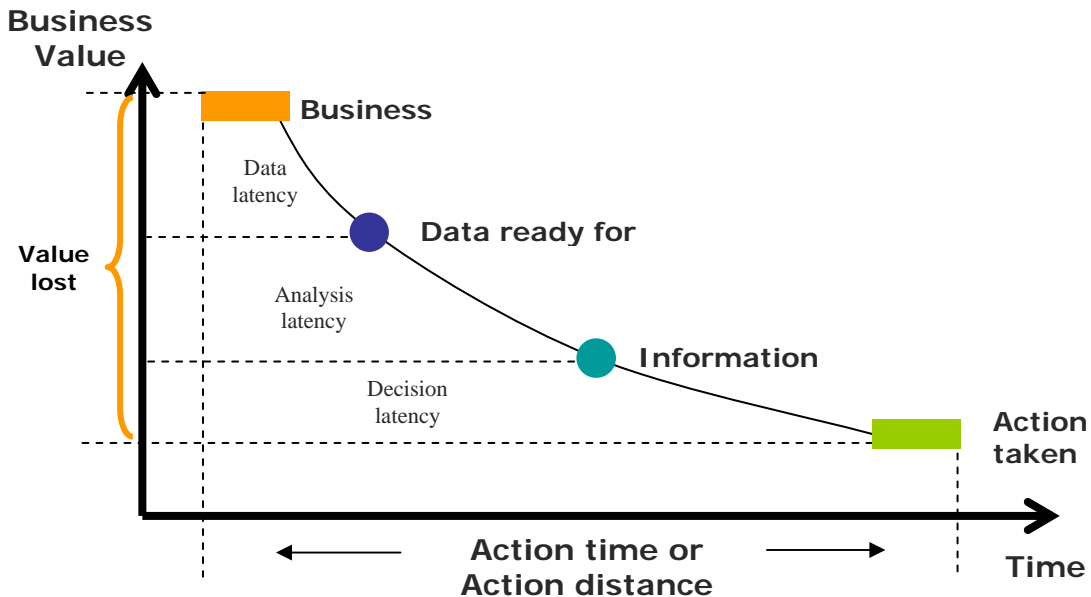
The last latency is **decision latency** which is the time it takes a human to comprehend the analytic result and determine an appropriate action. This form of latency is very difficult to reduce. Certainly education of the front-line business users about proper courses of action can help reduce this latency. Also the creation of rules-based suggested actions can help guide the business user to a quicker response time as well. The ability to remove the decision making process from the human and automate it will greatly reduce the overall decision latency. Many forward thinking companies are doing just that. For example, rather than send a high value customer a letter informing them of a bounced check (which takes days to get to the customer), an automated system can simply send them an immediate email or voice message informing them of the problem.

The key is to shorten these latencies so that the time frame for opportunistic influences on customers, suppliers, and others, is faster, more interactive, and better positioned. As mentioned above, the best time to influence a customer is not after they have left the store or the website. It is while they are still in the store or still wandering around your website. For example, a customer who is searching your web site for travel deals is far more likely to be influenced by appropriate messaging actions then and there. Actions taken immediately, while the customer is still in the site, might include:

- Offering the customer an appropriate coupon for the trip they showed interest in while searching for cheap airfares.

- Being given information about their current purchase such as the suggestion that visas are needed.
- Congratulating them on reaching a certain frequent buyer level and giving them 10% off an item.

Figure 1: The Latency Between a Business Event and an Action Taken from Richard Hackathorn, Bolder Technologies



A website represents another great opportunity to influence a customer, if the interactions are appropriate and timely. For example:

- A banner could offer the next best product to offer right after the customer puts an item in her basket.
- The customer could receive an offer for a product he just removed from his shopping basket.
- Appropriate instructions for the use of a product could come up on the customer's screen; perhaps warnings if the product should not be used by children under three, if the customer has such a child.

The move toward operational BI is inevitable; it requires rethinking the architectures and technologies used to support decision making and a study of how these are integrated into the everyday operational processes.

Impact of Operational BI on the BI Environment

This shift from traditional strategic BI to operational BI has some dramatic ramifications for most BI environments – the most obvious being the increased number of users. Traditional BI rarely had to support more than a few hundred, maybe a thousand or so users – there just weren't that many analysts in most enterprises. Opening BI up to operational personnel though means ramping up into the thousands even tens of thousands of users, which is a significant change for most BI environments. These users also have very different interface requirements. They are used to the operational systems that are menu-driven, have very clear pathways of navigation, and promote quick comprehension through visual and simple mechanisms. These requirements

mean that BI implementers may have to rethink the way BI is delivered to the business user. In many cases, the operational systems and user interfaces can remain the same, but they are augmented with strategic insights. For example, an airline agent checking people in might have a screen for seat rebookings (e.g., bumping some people into first class). The decision about who goes in what class seat doesn't require any new user interfaces per se. The decision for the BI implementer is to determine how to deliver strategic information to many more people. It means a tighter and faster connectivity of the enterprise decision support environment to the rest of the company. It may also require communicating in both directions, that is, more "events" coming into the system, and more "decisions" going out.

Secondly, the volumes of data needed to support operational BI will increase significantly. Detailed intraday snapshots of data are now being loaded or trickle-fed into data warehouses substantially increasing the amount of data stored. Tens to hundreds of terabytes are not unusual storage requirements for these expanding BI environments. Scalability is now a mandatory requirement in any BI technology, whether it is in the processing and integration of data, storage of massive volumes, or retrieval of query responses.

And finally, for operational BI to be useful in operational situations, its query performance must mimic or emulate response times in operational systems – that is, sub-second to just a few seconds to return data from a query. The ability to prioritize queries not only according to their importance but also their response requirements is a difficult but mandatory success criterion. It is this last feature that stumped so many BI implementers and BI vendors. The ability to handle this mixed workload gracefully and simultaneously is not for the faint-hearted. It takes a revolutionary rethinking of the overall infrastructure to successfully service these requirements.

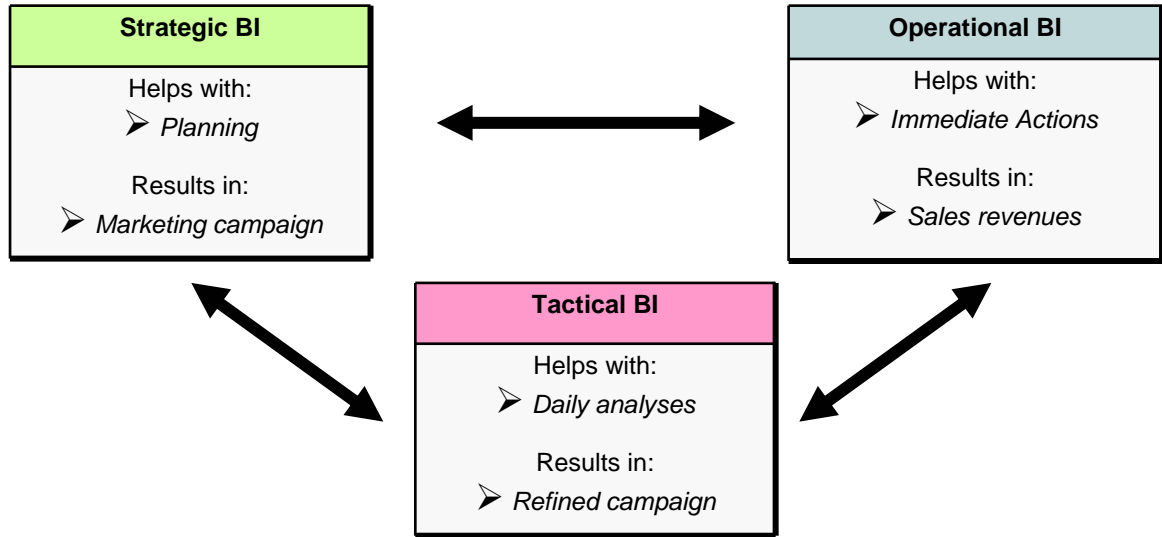
These three impacts cause many technologies to run for cover. The ability to handle such a mixed workload – operational response times, short tactical queries, massive analytical queries, thousands of concurrent users, and volumes of data unheard of before – takes careful thought and planning on the part of the technology vendor. A successful mixed BI environment depends on how well the technologies have embraced this new paradigm.

Bringing BI Altogether

The three forms of BI are not performed in isolation from each other. It is important to understand that they must work with each other, feeding results from strategic to tactical to promote better operational decision making. Figure 2 demonstrates this synergy. In this example, strategic BI is used in the planning stages of a market campaign. The results of these analytics form the basis for the beginnings of a new campaign, targeting specific customers or demographics, for example. The daily analyses of the campaign are used by the more tactical form of BI to change the course of the campaign if its results are not tracking where expected. Perhaps a different message is needed, or the inventory levels are not sufficient to maintain the current sales pace so the scope of marketing might be changed. These results are then fed into the more operational BI for immediate actions – offering a different product, optimizing the sale price of the product, or changing the daily message going to selected customer segments.

For this synergy to work, it should be obvious that the three forms of BI must be tightly integrated with each other. Minimal time should be lost transporting the results from one technological environment to another. Seamlessness in terms of data and process flow is a must.

Figure 2: The three forms of BI must work toward a common goal.



Impact of Operational BI on Operations

For operational BI to be at its most effective, it must be integrated into operational systems and operational processes to yield an ideal right time decision making environment. This requirement can pose some difficult problems if your operational environment is not ready or capable of interfacing with this new form of BI. There are several reasons for this incompatibility, not the least of which is the age of the technology used in operations. Older technologies may have proprietary software or fragile processes that make any sort of interface difficult or disruptive to the operational system. In these cases, you may have to replace the older operational system with a newer one before a truly functional workbench can be created.

In other cases, you may use Enterprise Information Integration (EII) technologies to tie the operational systems and BI insights together, displaying operational data along side the operational BI recommendations in a virtual fashion. One caveat here is to carefully monitor the impact the EII technology may have on the operational system. If operational performance is impacted by the intrusion of the EII technology upon the operational system, then this virtual environment will not be suitable.

Another possibility would be for the operational BI results to be fed directly into the operational system itself through some message brokering mechanism. For example, a list of next best product offers for specific customers could be transferred into the CSR's contact management system where it is displayed along with the other customer data when a customer is contacted. Perhaps the most straightforward answer for many environments would be to evolve what you already have. For example, the CSR screens could simply be augmented with new fields and panels to hold the BI results, not replaced with another display mechanism.

Another alternative could be to use web services from the operational system (if it is modern) to access the data warehouse, retrieving the BI information for usage in operational decision making. In this case, no other technologies would be needed.

Creating a World Class BI Environment

In building up your operational BI applications, there are a few project management things to be aware of that can greatly impact the success of your BI projects:

- The scope of an operational BI project is important. Many project managers of such projects do not realize that creating an operational BI application may have ramifications beyond the project's immediate boundaries. In choosing your first operational BI project, pick one that is "strategically" needed – that is, one that is not only recognized as having a significant impact on operations but can be accomplished in a reasonably short time frame, e.g., 3 to 6 months. Look for gaps in your current operational processing where operational BI could help. Don't try to make big changes to the operational processes; just speed up or make the processes you already have in place more efficient.
- Once you have chosen a project scope, perform a feasibility study of the proposed operational BI – operations interface. Understand what impacts the project may have on the operational system's performance, database design, and other factors. These technical aspects can impact the time frame of your project and increase its scope.
- The implementation of operational BI will most likely require retraining of operational personnel. How they make decisions, how they access and use the information flowing from the application, and how they monitor the impact of their decisions may change dramatically.
- Finally, operational BI applications generally cause changes to operational procedures or processes as well. The changes are usually needed to ensure that the operational BI information is being used optimally. Since these are new and different applications, the traditional processes may have to be rethought or rewritten to ensure proper execution.

What Has Made Operational BI Possible Now?

So what is different today that makes operational BI feasible? Teradata, a division of NCR has been working hard to make operational BI possible. Several technological advances, including their Active Enterprise Intelligence, have focused on:

1. Storage – Reduction in the cost of data storage has been a major factor facilitating operational BI. The massive increase in the volume of data needed to do operational decision making historically put the cost of these applications out of the range of most organizations until data storage technologies dropped in price tag. It is possible today to affordably and easily store tens to hundreds of terabytes of data.
2. Workload Management – Perhaps the biggest break-through has been the development of mixed workload technology. For years, we always separated online transaction processing from online analytic processing believing that the two environments were simply incompatible with each other. Through clever partitioning capabilities, indexing schemes and prioritization of query processing, Teradata® makes it feasible to support a mixed work load environment in the active data warehouse that can indeed combine into a single environment all three forms of BI – strategic, tactical, and operational – each having reasonable performance and response attributes for each business community.
3. Systems Management – A key to this success is, of course, the ability to actively manage the environment in mission-critical mode. Mixed workload environments require constant monitoring and vigilance to meet service level requirements for all users. Tools that give IT personnel quick insight into the performance of the technology are mandatory. Environments without these advanced tools are bound to run into trouble or cause sub-optimization for all user groups.
4. The ability to create alert-centric dashboards within workbenches spanning mixed BI / operational environments has been a great boon to the integration of operational BI in support of day-to-day decision making. For any BI – to be more than just an academic exercise, it must be actionable. This is particularly critical for operational BI where there must

be minimal delay between an event happening and an appropriate action being taken. The advent of operational dashboards and the creation of workflow-oriented workbenches make the ability to act quickly to changing circumstances not only feasible but mandatory for today's enterprises.

Summary

Business Intelligence has evolved from the analyst-only environment of strategic and tactical analytics. Today, operational BI is front and center supporting the operational decisions as well. To be truly effective, all three forms of BI must be implemented in a cohesive and manageable fashion. With the new technological advances from Teradata, the ability to integrate BI with operations is now possible. Mixed workload environments with both transactional and analytic attributes can be supported on a single database.

However, it must be stressed that operational BI projects must be carefully scoped and the ramifications to the operational processes and systems of the business be thoroughly thought out. New procedures, interfaces, perhaps even retraining of operational personnel may be needed to take full advantage of these new capabilities. Once created though, the advantages of making faster, more accurate, and more informed operational decisions are numerous. Better competitive advantage, stronger customer relationships, more efficient operations are just a few of the major benefits from integrating BI into your operations.