



---

September 2009

# Operational BI: Expanding BI Through New, Innovative Analytics

Going Beyond the Traditional Data Warehouse

Claudia Imhoff, Ph.D.

---

## Table of Contents

|                                 |    |
|---------------------------------|----|
| Executive Summary .....         | 3  |
| Operational BI Benefits .....   | 4  |
| Demands of Operational BI ..... | 4  |
| Leveraging New Analytics .....  | 7  |
| Decision Framework .....        | 9  |
| Getting Started .....           | 12 |
| Summary .....                   | 13 |

# Executive Summary

It seems that every day we read or hear something about operational Business Intelligence (BI). It has made quite a splash in our space for several years now. But what is operational BI? A widely accepted definition is that operational BI is:

**A SET OF APPLICATIONS, SERVICES, AND TECHNOLOGIES FOR MONITORING, REPORTING ON, ANALYZING, AND MANAGING THE BUSINESS PERFORMANCE OF AN ORGANIZATION'S DAILY BUSINESS OPERATIONS.**

There are several key phrases in this definition:

- A set of applications, services, and technologies – this brings business intelligence and business optimization technologies to the very heart of any corporation's operational environment.
- For monitoring, reporting on, analyzing, and managing – this means a coordinated delivery mechanism to ensure that the right information gets to the right people at precisely the right time. It is this feature that differentiates these environments from operational systems.
- The business performance of an organization's daily business operations – this means that operational BI information must be relevant for the current decision point in an operational process. Historical, low latency and real time data are all needed for an operational BI environment.

And we have learned a great deal about what can and can't be accomplished by speeding up the traditional hub and spoke architecture. Operational BI requires a significant paradigm shift in our BI architectures; it requires the implementation of new technologies, the complete integration of these with the existing traditional BI components, and a new overarching decision framework to bring it all together for a fully functional operational decision making environment. The Decision Framework presented has many advantages over traditional BI architectures: it reduces data latency inherent in traditional architectures to nearly zero, it is designed to support a significantly larger number of users, it reduces the requirement that all analytics must come

from the data warehouse, and finally it improves the consistency, reliability and accessibility of the analytics provided.

## Operational BI Benefits

Operational BI has a specific purpose in enterprises. First it enables *more informed business decisions* by directly supporting specific business processes and activities. The goal of operational BI is to provide an infrastructure from which people can make better and faster decisions – in other words, it is designed to provide analytics to everyone in the enterprise not just the analysts: e.g., front-line and customer-facing personnel, suppliers, customers, even external regulators in some cases.

Second it supports *faster business decisions* by seamlessly integrating BI with business processes to create a *closed-loop* decision-making environment. By seamlessly embedding BI into the everyday operational processes, the decision-makers can use the BI without needing to understand the complicated infrastructure, interfaces, or tools that supply the BI. In an ideal situation, the users of operational BI may not even know they are using it.

Finally operational BI provides a *more dynamic business environment* where the business can learn, adapt, and evolve based on analysis of its operational business performance. A key to operational BI success is that it must be a flexible environment. Since the business environment changes rapidly – changing financial situations, markets, customer purchasing behaviors, etc. – it is mandatory that operational BI allow the business user to change or evolve their workflows to match the changing situations.

## Demands of Operational BI

It should be obvious that operational BI has some very different demands placed upon it. Without doubt, some demands vary significantly from traditional strategic and tactical BI requirements; others remain the same. Here are some that are noteworthy.

## Increasing data velocity, volume, and variety

Operational BI requires a mixture of historical, low-latency, and real time data for its analytics. Data warehouses and operational data stores can be trickle fed integrated data in “mini” batches as frequently as every 15 minutes. This takes care of the historical and some low latency data. However, analytics requiring data that has an even lower latency or real time data necessitate a change in paradigm. It may not be practical or even necessary to consolidate real time or very low-latency data into a data warehouse or operational data store (ODS) for operational decision making. The more we drive toward real-time analytics, the more it becomes event-driven and process-centric. This change and the new form of analytics required will be discussed in the next section.

The frequency of uploads to data warehouse and ODS databases means the amount of data being stored can dramatically change. Tens of terabytes to hundreds of terabytes are not unusual storage requirements for data warehouses supporting operational BI. This growth rate is likely to accelerate given new and evolving information generating technologies including smart phones, RFID tags, sensor networks, and web information.

The increase in data volumes means that scalability is now a mandatory requirement for data warehouse technology. Seamless scalability has been achieved through numerous innovations like those found in IBM InfoSphere Warehouse offerings.

Finally the variety of data for these analytics has also changed. We now see that event data and unstructured content are viable sources of data but require new forms of analytics to be integrated with the traditional ones to support operational BI.

## Improving data warehouse management and response time

For tactical and strategic BI, the response time requirements for the data warehouse can be a bit more relaxed. If the results of these types of queries take a minute or two to return, it is hardly a matter of concern. However for operational BI, query performance must mimic or emulate the response times found in normal operational systems – that is, sub-second to just a few seconds to return results. This means that the underlying data warehouse technology must be able to handle a mixed workload gracefully and simultaneously regarding the queries submitted for low latency data. The technology must be able to prioritize queries not only according to their importance but also their response requirements. The infrastructure must ensure appropriate performance for both short operational BI queries and the more expansive, intense strategic and tactical ones.

In addition to handling a mixed workload, data warehouses must optimize data storage through mechanisms like data compression, more accurate and easier capacity planning, and data retention through multi-temperature mechanisms (e.g., actively used versus rarely used data).

Because operational BI directly supports transactional business processes and operational personnel, the environment must adhere to stringent service levels in terms of its availability. These too must mimic the operational systems' requirements in being available 24 hours a day, seven days a week, and 365 days a year. An outage of the operational BI environment will have a direct impact on the enterprise's ability to do business!

## Reaching new members of the business community

For many traditional BI environments, decision making is limited to those users with a good knowledge of the data and BI technologies involved. Less experienced users often find BI applications and tools difficult to use and many BI deployments still require significant IT involvement.

Operational BI, by its very nature, includes many more

employees involved in the day-to-day operational activities – sales execs, customer service reps, technical personnel, and so on. This means ramping up the number of users from a few hundred to potentially tens of thousands of users. And these users can be less technically savvy and less analytical in their thinking. They may also have very different interface requirements. These factors mean that BI implementers must rethink how BI is delivered to these business users. Simplified interfaces, guided decision making environments, self-service capabilities, and even mashups are now finding their way into these BI environments.

## Leveraging New Analytics

From the preceding discussion, it is clear that there are new forms of analytics needed to accomplish all the benefits stated for operational BI. Today analytics can be categorized as data, event and content analytics. Let's look at these in more detail:

**Data analytics** – This form is the most commonly used and is supported by the data warehouse, ODS, and data mart components. Traditional data analytics support intra-day, daily, weekly, monthly, and annual analytics. Typical usages are strategic BI (determining trends and patterns in data, performing predictive analytics, creating models of behavior for fraud, purchasing, etc.) and tactical BI (performing day, week, month, quarter comparisons, using multidimensional analysis to drill up, down, and around OLAP structures, determining customer segments, LTVs, next best products, etc.) Data analytics are performed on relatively static data using on demand, structured queries. We have increased customer retentions, improved supply and demand chains, determined future directions for products and the best placement for new stores through these highly effective data analytics.

**Event analytics** - Event analytics enable organizations to analyze events (business actions like ATM or POS transactions, business messages, and system events) closer to real time if not in real time. With service-driven event analytics, a service constantly monitors and analyzes real-time event data. Event

analytical results may optionally be sent to a persistent store or, if not needed, are not stored at all. Perhaps the biggest difference from traditional data analytics is in the nature of the data and the query. For event processing, the data is constantly changing as it streams real time into the enterprise while the query tends to remain static. Event analytics help detect fraudulent transactions faster, improve trading algorithms, increase web campaign effectiveness, and support location-based sales activities.

**Content analytics** – The final form of analytics for operational BI is content analytics. These are analytics performed on unstructured content obtained from web and collaborative systems like office productivity suites, email, IM, other forms of social media, as well as operational processes like purchase order systems, call centers, etc. Today, content analytics resemble data analytics in that the data is stored first (static) and then analyzed (dynamically). Content analytics add the “color commentary” to data analytics and can provide valuable insight for decision making. Examples include text analytics technology that is used to enhance data mining models, aid in more targeted customer activities, and understand “trigger” events that create at-risk customers.

Figure 1 describes the three forms of analytics needed for operational BI today.

Figure 1: Operational BI Analytics

| Data analytics                 | Event analytics                       | Content analytics                |
|--------------------------------|---------------------------------------|----------------------------------|
| Intraday, daily, monthly       | Real time, near real time, intraday   | Intraday, daily, monthly         |
| Static data                    | In-motion data                        | Static data                      |
| On demand                      | Event driven                          | On demand                        |
| Structured queries             | Services driven                       | Search queries                   |
| User centric                   | User and application centric          | User centric                     |
| Manual decision making         | Manual and automated decision making  | Manual decision making           |
| Alignment to plans and budgets | Alignment to rules and expertise      | Alignment to rules and expertise |
| Point-in-time data metrics     | Continuous process and stream metrics | Point-in-time content metrics    |



## Case Study: New York City Police Department Fights Crime Better

The New York City Police Department (NYPD) is one of the more innovative policing entities in the United States in a very challenging city. To improve what it already had, the department leaders recognized that the department had to further exploit its data resources, particularly its crime information. Unfortunately the NYPD faced the not-unusual problem of having this information fragmented across the entire department in multiple silos. Police work is heavily reliant upon information, *good* information, and the speed with which that information and resulting analytics can be delivered to the appropriate personnel. The fragmented nature of their data hindered the department from being as effective as possible and hampered informed crime fighting decisions.

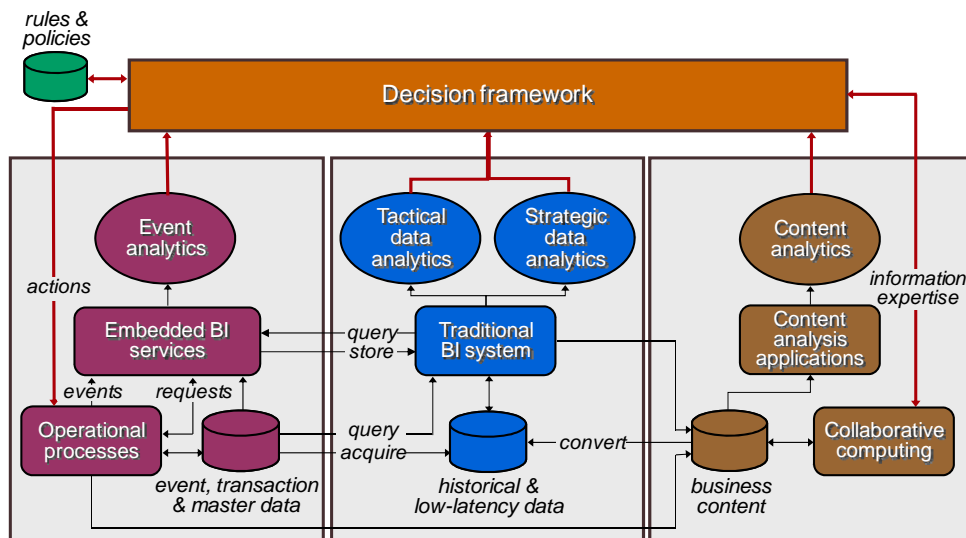
Improved analytics from an enterprise data warehouse (EDW) and operational data store (ODS) combined with operational (real-time) information were critically needed by the NYPD. They wanted an environment that would give them a more thorough and all-encompassing view of crime information. The solution had to integrate the silos of data and provide easy, yet secured, access to all crime information. With the help of IBM, the NYPD built the Crime Data Warehouse in phases on new technology that provided a state-of-the-art analytic functionality plus just-in-time information. The new platform consisted of an EDW, ODS, a web-based data entry system, and crime analytic and reporting capabilities. With the new environment, the department has removed all the hindrances it had in crime fighting and has the ability to identify crime trends, ensure officer and public safety, deploy resources more effectively, and reduce redundant work efforts.

## Decision Framework

To make operational BI a reality, implementers must embrace a new paradigm or architecture that encompasses all three

of these forms of analytics to supply the full “decision Intelligence” environment. This architecture is called the Decision Framework.<sup>1</sup> The Decision Framework brings together traditional and operational BI, operational processing, and the collaborative and business content environments (see Figure 2). All three forms of analytics are seamlessly delivered into a collaborative environment.

Figure 2: Decision Intelligence Supported by the Decision Framework



\*See Article titled "Decision Intelligence" by Claudia Imhoff and Colin White, [www.B-EYE-Network.com](http://www.B-EYE-Network.com)

The Decision Framework provides a personalized and self-service *discovery* of and *access* to business information using an *integrated* view of an organization’s business information. This information includes both structured business data and unstructured business content. The architecture ensures easy *analysis* of the business information and timely *delivery* of information via richer and more intuitive Web-based user interfaces. Finally, it supports the *sharing* of business information and expertise and a *collaborative* decision making environment.

The framework is a layered approach consisting of platforms for the DBMS, data integration and data warehouse/ODS/OLTP systems. These platforms serve as the foundation for the BI development environment and the

<sup>1</sup> See Article titled “Decision Intelligence” by Claudia Imhoff and Colin White, [www.B-EYE-Network.com](http://www.B-EYE-Network.com)

resulting BI applications which are accessed through the collaborative and web-based user interfaces.

The objective of the decision framework is to improve decision making throughout the enterprise, regardless of employee position or role. Aiding in this decision process are the underlying business rules and policies that define corporate and best practices. One benefit of decision intelligence is that it enables organizations to better understand their business processes and the factors that affect business performance. This understanding improves the quality of business rules and policies and leads to better decisions and actions. Most usage at present is user-centric decision intelligence that integrates the three analytics at the user interface level using technologies such as dashboards and portals.

## Decision Framework Technological Requirements

Operational BI environments must bring together information from many moving parts – traditional BI data analytics, event analytics applications, content management systems, and operational systems. The ability to embed BI services into the operational workflows at the appropriate points is critical to operational BI's overall success and therefore requires the use of a service-oriented architecture (SOA) to package BI capabilities as services. BI analytic results or analysis services can then be called directly by operational processes and presented to all personnel. These services produce analytics, alerts, or even automated business actions.

Second operational BI must present information in intuitive, understandable formats that seamlessly embed within operational process (applications, dashboards, portals, etc.). It should create the buffet of potential operational BI analytics and information, then give the business users access to the buffet so they can "choose" what they want at the appropriate timeframe. The environment must be easy to understand, easy to manipulate, and easy to embed in other processes. It

must also be well documented and changeable to meet the needs of the moment.

Thirdly, the infrastructure must be scalable and flexible. The infrastructure must be robust enough to withstand substantial alterations, the inclusion of new technologies and users, and the addition of enhanced capabilities. BI implementers need to constantly monitor the business community's usage patterns and build a robust architecture and technical infrastructure that can be quickly modified.

Finally operational BI must be self-contained and easy to use. It should be configurable to the individual user's needs by the user himself. Users should be able to change the metrics or thresholds being tracked, change the form and type of presentation, control the placement of metrics, and change business rules on the fly as business scenarios change during the workday.

## Getting Started

Now that you have a good feel for operational BI and its requirements, here are some tips to get you started.

The first step in any operational BI project should be to perform an honest assessment of existing data integration and delivery technologies. It is important to understand that any weaknesses discovered in assessment will only be exaggerated as you speed up the overall environment and processes. These must be addressed and corrected or the operational BI initiative could be jeopardized.

Next, determine what data management technologies are needed. In a traditional BI environment, data is integrated using ETL technology and then stored in a data warehouse before it is analyzed. Data latency can be reduced by updating the data warehouse more frequently through data propagation techniques. Data events and changes can be replicated from one data source to one or more data targets using changed data capture technologies. Data latency can be removed completely by directly accessing operational and event data using data federation technologies. Data federation can create a virtual view (a mashup) of not only operational data, but also data from a

data warehouse, content management system, or even an external data service.

It is important to pick a feasible project. Project managers may not realize that a new operational BI application may have ramifications beyond the project's immediate boundaries. While it is tempting to create just a simple application, operational BI requires that you think outside of this "box" to understand how and where your application may affect other operational systems. Standard operating procedures (SOPs) may need to be rewritten, retraining of operational personnel is just about guaranteed, and even new APIs may be needed. Alerts or alarms may be incorporated if processes are not followed or ignored.

Finally the team should get training on process modeling. A process is broken into smaller activities. Each activity is defined at a minimum in terms of what it does, when it is performed, and what data comes in and out of the activity. The implementation team can then develop an architecture that is component-based, loosely coupled, that is, SOA-oriented. And, implementers should determine if potential operational BI vendors are SOA-compliant or moving rapidly in that direction.

## Summary

The traditional BI and data warehouse architecture has served us well for nearly twenty years, but it is beginning to strain to cope with these new operational BI requirements. Pervasive integration technologies like SOA and Web 2.0 are creating a new IT environment forcing BI to fundamentally transform itself to meet higher user expectations for complete business integration.

Fortunately tools and techniques are already emerging for event analytics, embedded BI services, and content analytics. However, understanding the organizational implications for both business and IT of these new capabilities will be a key feature in operational BI's success.