



Converging Technologies: Real-Time Business Intelligence and Big Data

Claudia Imhoff, Intelligent Solutions, Inc
Colin White, BI Research
September 2013
Sponsored by Vitria Technologies, Inc.

Table of Contents

<u>Introduction</u>	<u>1</u>
<u>Business Benefits of RT Operational Intelligence</u>	<u>1</u>
<u>RT Operational Intelligence Technology Requirements</u>	<u>2</u>
<u>Supporting RT Operational Intelligence and Big Data</u>	<u>4</u>
<u>Real-time Operational Intelligence Use Cases</u>	<u>6</u>
Customer Analytics and Revenue Generation	6
Operational Optimization and Cost Reduction	7
Asset Protection and Loss Prevention	7
<u>Vitria Operational Intelligence</u>	<u>8</u>
<u>Conclusion</u>	<u>8</u>

Introduction

Over the past few years many articles have been published describing how more and more companies are deploying business intelligence (BI) for optimizing daily business operations. So far this *operational BI* approach has been accomplished primarily by reducing the latency of data integration and data analysis tasks in traditional enterprise data warehousing systems. Whereas reducing these latencies allows closer to real-time analysis of business operations, it does not and cannot enable real-time decisions to be made on real-time data, i.e., it does not support *real-time* (RT) operational BI.

To date, few organizations have recognized their need for RT operational BI. Even where they have, the complexity and costs have often been too high. The advent of big data, however, changes the situation dramatically.

Big data is a badly defined and overused term, but, to us, it represents not only new sources of data that aid the analysis and optimization business operations, but also the advances in the sophistication and power of analytic techniques. It also reflects the significant advances made by vendors in reducing the costs and improving the performance of analytic software and hardware platforms. Big data increases the number of use cases for RT operational BI and, given the improvements in the price/performance of analytic processing, this makes real-time processing viable for a growing number of organizations.

In this paper, we explore the use cases for RT operational BI in the era of big data, examine the technologies that enable RT operational BI, and discuss how the traditional enterprise data warehouse (EDW) environment can be extended to support both real-time processing and big data. Along the way, we will also look briefly at how Vitria, the sponsor of this paper, supports the development and deployment of RT operational BI applications. A follow-on paper will provide in-depth case studies of RT operational BI projects, the technologies involved, and the lessons learned from those efforts. Note that in this paper, we use the term *RT operational intelligence* as a synonym for real-time operational business intelligence.

Business Benefits of RT Operational Intelligence

All employees are decision-makers at some point during their workday. Front line workers, managers and executives making split-second decisions about customers, products, orders and even logistics need the most current information they can get to ensure accurate and appropriate reactions and results. Timely information leads to several business benefits:

- **Faster decisions** –intervening at the right moment with a customer who is about to churn can save that relationship; shipping a replacement part *before* the old part fails means less down time.
- **More efficient business processes** – identifying performance bottlenecks in a telephone network or web store leads to better customer satisfaction; monitoring and

optimizing call center performance reduces customer wait times and enhances call center efficiency.

- **Improved revenues and reduced costs** – offering a customer the right product at the right time results in a sale that might otherwise have been lost; real time analysis can reduce inventory-carrying costs by having just the right amount in stock, and never having an outage.

RT Operational Intelligence Technology Requirements

The goal of RT operational intelligence is to enable organizations to make real-time decisions on real-time data. To understand the technologies required to enable real-time processing, it is useful to examine how business intelligence is implemented and used in organizations and then discuss how this is affected by RT operational intelligence. The BI lifecycle diagram in Figure 1 will be used to aid this discussion.

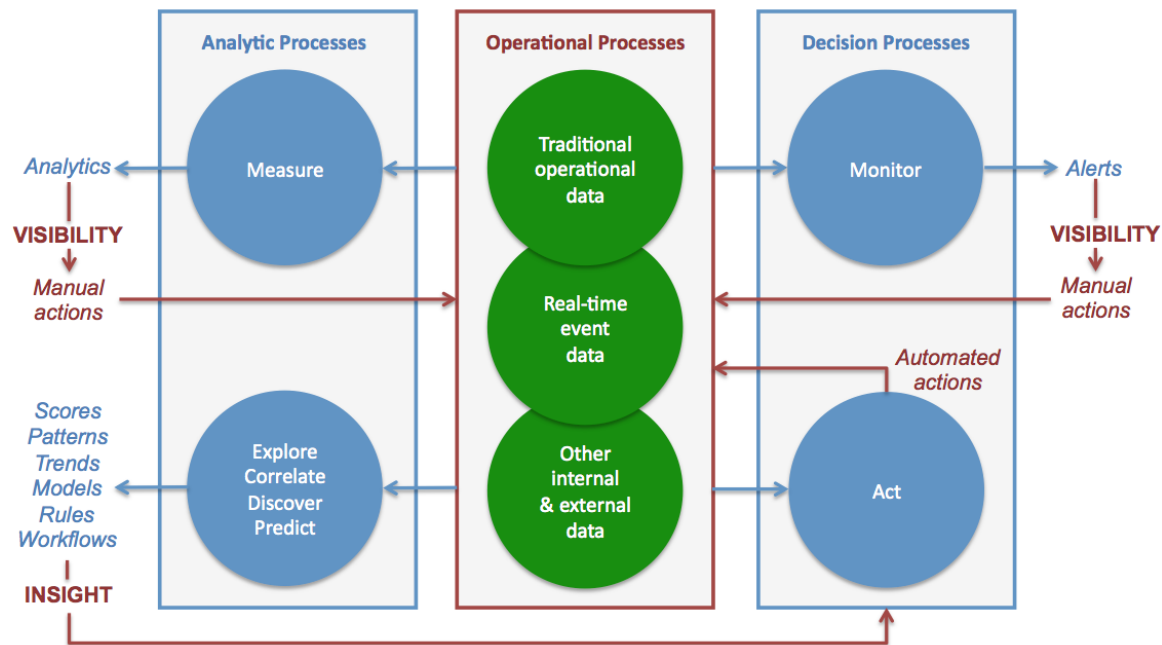


Figure 1: The Business Intelligence Lifecycle

At the center of the diagram (colored in green) are the data sources that provide information about the business operations of the organization. To date these data sources have consisted primarily of the databases and files created by business transaction processing applications. The use of big data, however, has extended these data sources to include real-time event data (from the web environment, sensors, etc.) and data from other internal and external data sources (such as call-center logs, social media sites, external information providers, etc.). These big data sources can add considerable

business value to the decision making process, but the challenge of course is extracting this value in a cost-effective and timely manner.

In a traditional EDW environment, information from operational source systems is captured, transformed and consolidated into an EDW data store ready for analysis. The data in the warehouse represents a historic snapshot of business operations at the point-in-time the data was captured from operational systems. BI applications use this historic information to produce analytics (or measures) about business operations at the time the snapshot was taken. Business users then apply their business knowledge and expertise to the analytics delivered by the EDW system to understand how the business is performing and to take any actions required to optimize business processes or correct any business problems.

The speed of business decisions and actions can be improved by using advanced analytic techniques (such as scoring, data pattern recognition, forecasting, predictive modeling, and statistical and text analysis) to create *intelligent analytic objects* (such as analytic models, business rules and workflows) that can be *embedded* in RT decision processes. This approach is illustrated in the bottom half of Figure 1. The decision processes work in conjunction with operational processes to analyze data in real time. This analysis may result in alerts to warn business users about urgent business issues that need immediate attention, or may, in business critical situations such as fraud detection, generate automated actions. The EDW environment can then be used measure the business impact of the decision processes and, if required, to update the intelligent objects.

The need to support faster business decisions and actions, and also to exploit the business benefits of advanced analytic techniques and big data, adds new technology requirements to the traditional EDW environment. These are summarized in Figure 2.

1. Capture and transform real-time event data
2. Capture and transform not only structured data, but also multi-structured data from other internal and external data sources
3. Reduce the latencies of data capture and transformation, analytic processing and user decision-making and action making
4. Support high-performance advanced analytic techniques that enable improved insight in the operation and performance of business processes through the creation of intelligent analytic objects such as scores, models, rules and workflows
5. Embed intelligent analytic objects into decision processes that can analyze data in real-time and generate automated real-time alerts and actions

Figure 2: Technology Requirements for RT Operational Intelligence

The requirements shown in Figure 2 cannot be satisfactorily met by the traditional EDW. For example, the more up-to-date the data is in the EDW with respect to operational systems, the better visibility there will be into recent business operations. There is a limit, however, to how up-to-date the data in a data warehouse can be, how fast analytics can be produced and delivered to business users, and how rapidly users can make decisions and take actions. There are inherent latencies in the traditional EDW BI lifecycle, and although these latencies can be reduced, they cannot be eliminated to enable real-time decisions to be made on real-time data. The traditional EDW environment therefore needs to be extended to support RT operational intelligence.

Supporting RT Operational Intelligence and Big Data

An extended data warehouse architecture for supporting RT operational intelligence and big data is shown in Figure 3.

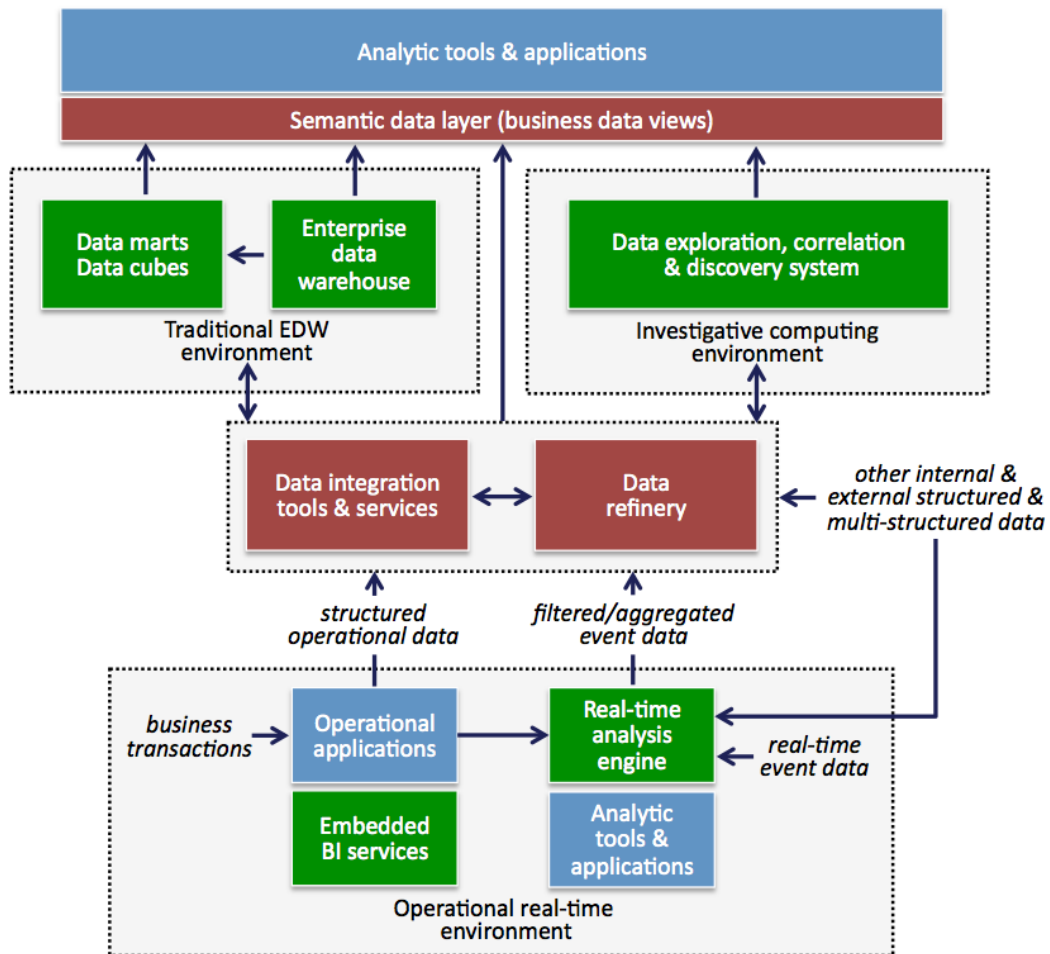


Figure 3: The Extended Data Warehouse Architecture

This architecture handles both *in-motion* data and *at-rest* data. In-motion data is real-time data flowing through the operational environment shown at the bottom of the diagram. At-rest data is a snapshot of in-motion data from operational systems, transformed and then loaded into the EDW and investigative computing environments shown at the top of the diagram. The time difference, or data latency, between the at-rest data and real-time is dependent on how often the at-rest data is refreshed or the snapshot taken.

The bridge between in-motion data systems and at-rest data systems can be built in two different ways. The traditional method is to use data integration tools and services that use an extract, transaction and load (ETL) approach to bridge the data. A growing trend, however, is to first load the captured in-motion data into a staging area or *data refinery*, transform it, and then route the transformed data into other systems as required.

Analytic processing can be done by at-rest data systems (such as an EDW or investigative computing platform) or by in-motion data systems in the operational real-time environment or both.

As already discussed, intelligent analytic objects from the at-rest data systems can be embedded in decision processes for real-time data analysis and for automating decisions and actions. These can be implemented as *embedded BI services* for use within individual operational business processes and applications as shown in the bottom left of Figure 3. One challenge with this approach, however, is that these services are often difficult to embed in existing monolithic business transaction applications. Organizations are therefore more likely to use these BI services when they develop new or reengineered operational applications using a services-oriented architecture and business process management software.

Another powerful and relatively new approach to implementing RT operational intelligence solutions is using a *real-time data analysis engine*. The technology behind many of these solutions was developed originally for the financial industry for tasks such as analyzing real-time trading. More recently, the technology has been extended to support big data sources and is being used for a variety of applications in a number of industries.

Real-time data analysis engines can process both structured data and multi-structured data. They are particularly well suited to the handling of very large volumes of event streams from sensor devices. Technology terms often associated with a real-time analysis engine are event stream processing (ESP) and complex event processing (CEP). The dividing line between ESP and CEP is blurring as vendors build solutions that support both styles of processing. As a result CEP is becoming the dominant term.

ESP is focused on data streams where the events being processed have a definite time sequence. The software provides a windowing capability that enables an application to analyze a data stream for a fixed period of time. So, for example, a financial application can analyze a trading data stream starting in five minutes for a period of ten minutes.

CEP on the other hand is targeted at correlating multiple data streams that may be unrelated to each and may not necessarily be in time sequence. An example would be correlation of retail sales with weather data to determine how the weather is affecting the purchase of certain items. This is a simplification of the ESP and CEP approaches, but the key points to note are that there are various techniques for analyzing real-time data streams and that vendor products now support a number of them.

Real-time data analysis engines can employ the intelligent analytic objects created by data-at-rest systems. They also often provide their own tools for developing such objects. Regardless of their derivation, these objects can be used to do real-time data analysis of real-time data for tasks such as fraud detection, next best customer offer, and so forth. They can also be used to filter high-volume events for use by data refineries and data-at-rest-systems. This is shown in the bottom right of Figure 3.

Companies deploying real-time data analysis engines sometimes use them in conjunction with an investigative computing platform. The real-time engine is used to filter event data for use by the investigative computing platform, which in turn creates intelligent analytic objects for use by the real-time analytic engine for analyzing real-time data. These intelligent analytic objects are built by blending and analyzing the filtered event data and data from at-rest data sources.

Real-time data analysis engines support all of the requirements outlined earlier in this paper. They are well suited for building RT operational intelligence applications using both existing data and new big data sources. They also enable a wide range of applications to be developed from churn analysis and fraud detection to equipment monitoring and traffic flow optimization.

The architecture shown in Figure 3 consists of a variety of components that enable companies to extend their traditional EDW environment to support both RT operational intelligence and big data. Most organizations are likely to evolve to such an architecture over a period of time. It is unlikely that any single vendor can supply all of these components, and therefore care must be taken when extending the EDW environment to ensure that as components are added they can be integrated easily into the existing IT infrastructure and are sufficiently open to be able to support possible future enhancements to the architecture.

Real-time Operational Intelligence Use Cases

Use cases for RT operational intelligence vary considerably, but, in general, they fit into three broad categories: customer analytics and revenue generation, operational optimization and cost reduction, and asset protection and loss prevention.

Customer Analytics and Revenue Generation

The importance of customer retention and satisfaction cannot be underestimated. Monitoring a highly profitable customers' experiences with your company, its products and personnel in real-time is critical to ensuring their continued loyalty and profitability.

Companies are now able to perform 1:1 marketing based on where customers have been, where they are now, and where they are likely to go next. They can detect anomalies in behaviors indicating customers who are likely to churn and quickly offer new or additional products and service plans that fit the customer's lifestyle and usage patterns.

The ability to take instant actions, automate escalations and begin intermediation actions to stop customers from leaving, improves the generation or retention of the revenue streams from customers. It increases customer satisfaction and gives the organization immediate insight into what works and what doesn't. It also enhances the organization's ability to correct problems before customers even know a problem has occurred.

Operational Optimization and Cost Reduction

In today's economy, corporations everywhere must squeeze inefficiency out of all their operational processes as quickly as they can to avoid higher costs and reduce wastefulness. The ability to detect these inefficiencies in real time is another good use case for RT operational intelligence. An example here is supply chain management. Often a company runs the risk of stock-outs due to a misalignment between customer demand for a product and the company's inventory resulting in significant revenue loss as well as customer dissatisfaction. Being able to detect these problems early reduces missed shipments, overages and underages and reduces costly emergency replenishment due to process inefficiency.

As another example, certain types of complex equipment and devices such as cell towers, networks, aircraft, oil wells and automobiles often provide "signals" of pending failures. Unfortunately many organizations are not able to recognize these signals, which results in avoidable outages and unhappy customers. The good news is that companies now have the technology to determine potential product outages or failures affecting their customers in real time before they actually happen.

Asset Protection and Loss Prevention

The most obvious application for RT operational intelligence in this category is being able to detect fraudulent transactions within seconds of their creation. The first step is to develop models of fraudulent behavior from a set of known fraudulent transactions. These models are used in the RT operational intelligence applications to compare to new transactions entering the systems. Those that match the fraudulent characteristics in the models are then redirected to another set of automated or manual processes for further investigation.

Financial companies are also using sophisticated algorithms and analyses for early detection of "out of process" outliers and suspicious activity. They can detect and predict transactions that will be in jeopardy of default (car and home loans, transactions missing documentation, for example) resulting in fewer late settlements, fewer fines and improved regulatory compliance. They can also use early detection analyses to determine potential cyber security breaches as well.

There are many other use cases for RT operational intelligence involving situational awareness, smart call center management and service level agreement monitoring and management. In the second of these papers, we will expand on these use cases by examining detailed customer case studies.

Vitria Operational Intelligence

The Vitria Operational Intelligence (OI) Platform supports the business intelligence lifecycle of Figure 1 and satisfies the technology requirements outlined in Figure 2. Specifically it supports both traditional analysis and advanced analysis (scoring, pattern recognition, forecasting and predictive modeling, for example) of in-motion and at-rest data from a variety of sources. These sources include web services, custom legacy applications, databases and files, cloud-computing data sources and web and social computing sites.

Included with the Vitria OI platform is its Business Process Management Server (BPMS), which provides the ability to define and manage policies (and rules), apply those policies to events, and take actions based on those policies. The policies and rules are embedded in workflow models developed using Business Process Management Notation (BPMN). These models support both automated and human-oriented workflows, and can be applied to big data sources and also existing legacy systems without disrupting existing applications and processes. These capabilities empower business analysts to define *decision processes* with automated actions, as depicted in Figure 1.

One interesting aspect of the platform is that it can be used to investigate and discover hidden business processes in existing operational systems. In addition, Vitria offers its KPI Builder App, which supports the real-time and continuous monitoring of key performance indicators and service level agreements (SLAs). The KPI Builder includes a point-and-click dashboard builder that supports both real-time alerts and notifications.

Within the extended data warehouse architecture (Figure 3), the Vitria OI Platform provides the capabilities shown in the *operational real-time environment* component. The platform includes a real-time analysis engine, real-time analytics tools and applications, and provides the ability to embed BI services within business and decision processes. The platform can be deployed both on-premises and in a cloud environment.

Conclusion

Many corporations have believed that RT operational intelligence was beyond their technological capabilities as well as their budget. The good news is that this is no longer true. Today, the complexity of the environment has been greatly simplified and its total costs reduced. We can give a fair amount of credit to the big data push in many organizations for the dramatic increase in the demand for RT operational intelligence.

RT operational intelligence has also demonstrated significant benefits to many of today's modern corporations. The ability to make faster, more accurate decisions based on

real-time data has been proven to preserve revenue streams, decrease costs and minimize significant risks. But the creation of this environment should be an evolutionary process.

Most organizations move toward RT operational intelligence through a progression of BI maturity – from simple measurements to sophisticated insight to real-time action. Many companies start by speeding up the data integration processes in their traditional enterprise data warehouse environments. The ability to reduce the latency of structured data certainly moves the organization closer to operational intelligence, but it still a far cry from RT operational intelligence.

The enterprise must embrace the extended data warehouse architecture to fully achieve the real-time analytical capabilities described in this paper. This architecture includes both the investigative computing environment and, more importantly, the analytic capabilities found within the operational environment itself – embedded BI services and a real-time analysis engine.

The use cases for RT operational intelligence fall within the broad categories of customer analytics and revenue generation, operational optimization and cost reduction, and asset protection and loss prevention. Every industry and government entity can easily generate their business cases within these to justify the implementation of their own RT operational intelligence capabilities.

About BI Research

BI Research is a research and consulting company whose goal is to help organizations understand and exploit new developments in business intelligence, data integration and data management.

About Intelligent Solutions, Inc.

Intelligent Solutions is a research and professional services firm, founded in 1992 by Dr. Claudia Imhoff, dedicated to assessing, planning, and guiding business intelligence (BI) through its services.