# COMBINING BUSINESS ACTIVITY MONITORING WITH THE DATA WAREHOUSE FOR EVENT-CONTEXT CORRELATION

## *Examining the Practical Applicability of this BAM Approach*

Gabriel Cavalheiro, Ajantha Dahanayake

*Faculty of Technology, Policy and Management, Delft University of Technology, Jaffalaan 5, Delft, The Netherlands*

Richard Welke

*Center for Process Innovation, Georgia State University, Atlanta, GA, USA*

Keywords: Business Activity Monitoring, Data Warehouse, Real-Time Event, Historical Data, Push-Based, Performance-Indicators.

Abstract: Business Activity Monitoring (BAM) is a term introduced by Gartner, Inc to define systems that serve to provide real-time access to critical business performance indicators to improve speed and effectiveness of business operations. Despite the emphasis of BAM on the provision of low latency views on enterprise performance, literature on BAM also indicates the technical feasibility of a BAM approach, which adds context from historical information stored in a data warehouse to real-time events detected by a BAM system so as to help enterprises improving understanding of current monitoring scenarios. However, at this point, there is a lack of studies that discuss the use of this approach to tackle real-world business problems. To improve practical understanding of the potential applicability of this BAM approach, this paper will present a synthesis of existing research on BAM and data warehouse to provide an objective basis for proposing feasible business scenarios for applying the combination of both technologies. This study reveals that the noted BAM approach empowers operational managers to respond in a more precise manner to the occurrence of events by enabling a better understanding of the nature of the detected event.

## 1 INTRODUCTION

Enterprises are currently exposed to rapidly changing market conditions and increasingly competitive business environments (Luckham, 2002). To ensure competitiveness, enterprises need to be able to maximize revenue generation and cost savings when it comes to the execution of their business processes. Since the occurrence of exception events of interest in the transactional systems, such as unusual large orders, delays, situations of high risk, or unavailability of resources, have a significant economic impact on the performance of business processes, enterprises need to respond adequately to such events as they occur. However, this requires the availability of an infrastructure that provides real-time visibility into the business operations of enterprises.

To enable enterprises minimizing response times to events, software vendors started developing a new set of functions. Gartner, Inc introduced the term Business Activity Monitoring (BAM) to define this new set of functions (Cavalheiro, 2005). Fundamentally, BAM systems are push-based systems that translate input events into real-time analysis that is pushed on recipients for immediate reaction upon the occurrence of events (Govekar et al, 2002). The provision of real-time access to business performance indicators requires gathering and analysing business events from multiple and heterogeneous data sources to detect exception conditions and generate low latency alerts to enable business managers responsible for business processes to make well-informed decisions quickly (DeFee and Harmon, 2004).

Although the concept behind BAM systems emphasizes the provision of real-time operational insights, some studies acknowledge that the combination of real-time events with historical data can help enterprises diagnose problems in current monitoring scenarios (McCoy, 2002). Nonetheless, at present time, there is a lack of studies addressing the practical applicability of this business monitoring approach. The purpose of this paper is thus to contribute to fill this gap in the literature by

Cavalheiro G., Dahanayake A. and Welke R. (2006).
COMBINING BUSINESS ACTIVITY MONITORING WITH THE DATA WAREHOUSE FOR EVENT-CONTEXT CORRELATION - Examining the Practical Applicability of this BAM Approach.
In *Proceedings of the Eighth International Conference on Enterprise Information Systems - DISI*, pages 263-268
DOI: 10.5220/0002449202630268

presenting a synthesis of existing research on BAM and data warehouse to provide an objective basis for proposing the application of this approach to address real-world business scenarios. This is accomplished by examining technical characteristics of both concepts and by proposing practical combinations BAM and data warehouse functions.

This paper utilizes the insights obtained by conducting a graduation project, which involved collaboration between the faculty of Technology, Policy and Management at Delft University of Technology, Delft, The Netherlands, and the Centre for Process Innovation (CEPRIN) at Georgia State University, Atlanta, USA.

# 2 TECHNICAL CONSIDERATIONS ABOUT BAM

Essentially, BAM systems analyse real-time events to identify problems, diagnose them, and generate alerts to recommend managerial action. BAM systems can be incorporated into the IT infrastructure of enterprises as an event analytics layer at the top of existing middleware infrastructure (Govekar et al, 2002). In practice, BAM systems can be deemed as rule-based systems that detect significant real-time events. These real-time events can be characterized as instantaneous state changes in a condition from one value to another occurring in transactional systems (Delgado et al, 2004). BAM systems serve as a platform upon which event-driven applications can be developed by defining patterns to represent events with significant economic impact on business processes. To support the development of event-driven applications, BAM systems are likely to include event-modelling functions to define and validate event patterns of interest (Gassman, 2004). This makes BAM systems highly adaptable, as new event-driven applications can be developed rapidly to address new or changing business problems.

The investment committed by enterprises on the adoption of middleware technologies makes it possible that, in BAM systems, events come from multiple and heterogeneous data sources, which include both internal sources as well as external sources such as the Internet (McCoy et al, 2001). Essentially, BAM performs multi-application event analysis by correlating events originating from multiple and independent sources (McCoy, 2004).

To represent BAM systems, Gartner, Inc has proposed a three-layer model, which is comprised of an Event Absorption Layer, an Event Processing and Filtering Layer, and an Event Delivery and Display Layer (Govekar et al, 2002). In this model, the border of a BAM system is the interface with the recipients of BAM alerts. Figure 1 provides an illustration of each layer and the recipient. In the following we briefly describe each layer.

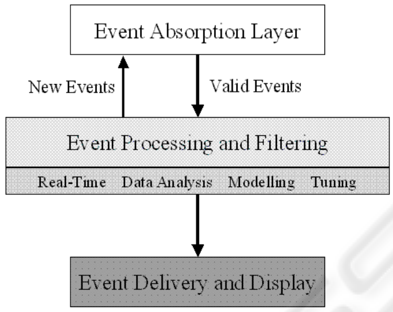


Figure 1: Layers of a BAM system.

## 2.1 Event Absorption Layer

Events are fed into the event absorption layer by an event acquisition tool. The source of event messages for BAM will most often be business or process-related, however, technical events that occur during the operations of the IT infrastructure may also be collected (Gassman, 2004). BAM systems rely on event gathering mechanisms to gather event information, in real-time, directly from transactional systems. A BAM system should have a mechanism in place to gather events occurring concurrently in different transactional systems. This is enabled by the middleware layer that has been placed over the last decade to support interoperability among disparate transactional systems.

## 2.2 Event Processing and Filtering

Basically, after gathering real-time events, BAM systems must be able to process and correlate multiple sources of independent data (McCoy, 2004). To this end, a filtering system must be in place to draw data from a wide range of sources and then compare those real-time events against rules that when met generate alerts (DeFee and Harmon, 2001). These alerts should contain information about nature of the problem associated with the occurrence of the event in order to empower the recipient of the alert, the BAM recipient, to initiate an immediate and appropriate response (Hellinger and Fingerhut, 2002).

## 2.3 Event Delivery and Display

The alerts issued by a BAM system can be sent to diverse parties that have real-time decision support needs. The alerts that are raised can populate a display or trigger an action (Gassman, 2004). In this way, alerts that are used to populate a display, are often delivered via graphical displays ("Dashboards") that are customized for use in different parts of the enterprise and for different audiences (McCoy et al, 2001). Another option is to forward the alerts to BAM recipients through other existing mechanisms such as batch office staff, emails, pagers, PDAs, and other systems to ensure that someone or something can react (McCoy, 2003).

The alert can also be used to trigger an action executed by a Business Process Management (BPM) system. In this respect, controlling the reaction cycle will probably be the role of BPM tools (McCoy, 2003). Here, BAM can be part of a BPM, where it generates an input for the BPM system that triggers a workflow corresponding to a predefined sequence of events (Gassman, 2004). The BPM tool can react to the BAM alert by running a chain of tasks that alter a running business process (McCoy, 2004).

### 2.4 BAM Recipients

The alerts issued by a BAM system can be sent to diverse BAM recipients (Govekar et al, 2002). The alerts that are delivered by a BAM system can populate a graphical display or trigger an action (Gassman, 2004). In this way, the alerts that are used to populate a display are often delivered via "Dashboards" that are customized for use in different parts of the enterprise and for different audience (McCoy et al, 2001). Another option is to forward the alerts to BAM recipients through other existing mechanisms such as batch office staff, emails, pagers, PDAs, and other systems such that someone or something can react (McCoy, 2003).

Additionally, the alert can also be used to trigger an action executed by a BPM system. In this respect, controlling the reaction cycle will probably be a role of BPM tools (McCoy, 2003). Here, BAM can be part of a BPM, where it generates an input for the BPM system that triggers a workflow corresponding to a predefined sequence of events (Gassman, 2004). As such, the BPM tool can react to the BAM alert by running a selected set of processes that alter a running business process (McCoy, 2004).

## 3 APPLICATIONS OF BAM

BAM systems are applicable to monitoring situations that require very low latency operational insights on the execution of business processes. Essentially, in BAM systems, the emphasis is on responding to events within short time-windows of opportunity (Chandy and McGoveran, 2004). Consequently, BAM systems must generate alerts in real-time, otherwise the value of the system is lost (McCoy, 2003). DeFee and Harmon (2004) define effective BAM systems as those systems that provide managers with sufficient operational insights close to real-time so managers can take decisions in time to affect the ongoing performance of the process flow. A number of examples of the need for real-time operational insights is provided by literature. Table 1 shows a set of representative examples of such needs.

Table 1: Examples of the need for real-time Operational Insights.

| Real-Time Operational Insights | Source |
|---|---|
| A retailer that needs to monitor hourly sales levels | (Hackthorn, 2004) |
| A financial system BAM tool issued an alert on a stock volume increase or decrease beyond a threshold level | (McCoy, 2003) |
| Provide a real-time view on supply chain metrics | (McCoy et al, 2001) |
| Monitor service level agreements (SLAs) | (Luckham, 2002) |

## 4 TECHNICAL CONSIDERATIONS ABOUT THE DATA WAREHOUSE

In practice, the data warehouse can be deemed as a workflow that involves periodically gathering disparate data and cleansing, transforming, and integrating that data according to business rules store in metadata repository. Then, the data is loaded into persistent data structures where it can be analyzed (Inmon et al, 2001).

The data stored in a data warehouse represents historical data in the sense that it represents events now passed (Inmon et al, 2001). The data warehouse is designed to store fact for each time period, thereby creating a historical perspective on performance (Tanler, 1997). It can create a single subject-oriented collection of information by assembling the data from heterogeneous databases (Thuraisingham, 1999). This allows business analysts to query only a single point consisting of a repository from multiple sources. In order to provide a single version of the enterprise operational data, the data warehouse need to integrate data from multiple operational systems of the enterprise and, as data warehouse environments mature, it brings in data from an ever-expanding set of sources (Inmon et al, 1999). To do this, the data warehouse typically relies on a software to carry out the extract, transform, and load (ETL) process. New or changed transactions (fact records) are moved and dimensions are captured as point-in-time snapshots for each load (Kimball and Caserta, 2004). In fact, a new snapshot is created whenever a change needs to be reflected in the data warehouse (Inmon et al, 2001). A fundamental consideration about the ETL process, though, is that it depends on the availability of windows of acceptable downtime for the transactional systems, because it places an additional load on the transactional systems during the loading process. In this way, the latency associated with the data warehouse makes it unsuitable to provide real-time operational insights.

The snapshots captured by the ETL software are, then, sent to the operational data store (ODS). The ODS was originally defined as a place where data was integrated and fed to a downstream data warehouse (Kimball and Caserta, 2004). The ODS environment has a time period identical to that of the application. The difference between the ODS and the application is that the ODS contains integrated corporate data, and the applications do not (Inmon et al, 2001).

A data mart is a collection of data that is adequate to fulfil the decision support needs of a particular department. It is a subset of a data warehouse that that is generally customized to fit the needs of a department (Inmon et al, 2001). Data Marts can be a subset of the enterprise data warehouse (Brown and Hill, 2000). The data marts that house data for various departments contain different combination and selections of the same detailed data found at the data warehouse (Sousa, 1999). They provide detailed information focused on a single area, such as marketing, sales, production, or finance (Brown and Hill, 2000). From the data warehouse, atomic data flows to various departments for their customized usage. These departmental databases are called data marts. A data mart is a body of data for a department that has an architectural foundation of a data warehouse (Sousa, 1999).

## 5 APPLICATIONS OF THE DATA WAREHOUSE

Typically, the main purpose of the data warehouse is to supply business analysts, such as managers and decision-makers, with information they need, so as not to disturb transactional systems processing transactions rapidly and reliably (Abramowicz et al, 2002). Specifically, data warehouses are generally highly adequate to fulfil decision support needs when there is a clear need for long-term trend reports. A data warehouse is designed to support exploratory analysis through adhoc data analysis, inquiry and reporting by end users. The integrated, detailed data and the robust amount of history found at the data warehouse are ideal for this comprehensive business and data analysis (Inmon et al, 2001).

A great number of examples of the exploratory analyzes based on data warehouse is presented by scientific literature. Table 2 lists some of the most interesting applications of the data warehouse.

Table 2: Examples of exploratory analyzes supported by the data warehouse.

| Example of Analyzes Based on a Data Warehouse | Source |
|---|---|
| Managers need to view sales by product and region and make correlations with advertising campaigns and marketing promotion. | (Brown and Hill, 2000) |
| Answering questions such as,<br>• What type of customer is the most profitable for our business?<br>• Over the years, how has transaction activity changed?<br>• Where has sales activity been highest in the springtime for the past three years?<br>• When we change prices, how much elasticity is there in the market place? | (Inmon et al, 2001) |
| The most frequent use of a data warehouse is to analyze historical data, discover trends and correlations and project events forward into the future. | (Meltzer, 1999) |
| Use data warehouse to analyze business data historically by focusing on planning such as trends in daily sales levels as compared with previous months. | (Hackathorn, 2004) |

## 6 LINKING BAM AND THE DATA WAREHOUSE: EVENT-CONTEXT CORRELATION

Besides analyzing real-time operational information, BAM systems can also support event-context correlation. This is an important aspect of BAM systems because the addition of contextual information to an event provides additional information about the nature of the situation that is being monitored. For example, the definition of a certain filtering criteria to detect exceptional events can be seen as a basic form of adding context to BAM alerts, which are defined by Gassman (2004) as events with context. So, by specifying a threshold for a certain performance-indicator, which is represented by an event property, the value of the threshold can be regarded as contextual information. However, more sophisticated ways of adding context to events are usually required. Since the decision support need of the BAM recipient may require the diagnose of problems to respond to events with high accuracy or event forecast data to guide anticipated response to future problems, just raising alerts indicating the occurrence of exceptional situations can be, in many monitoring situations, insufficient to ensure a proper response.

Data warehouses can be employed to add context from a historical store for business activity to a real-time alert. This can be done by building models that correlate patterns of real-time events with previous occurrences of similar events representing both operational problems and opportunities. For example, an hourly trend of expected order volume may be updated nightly by a data warehouse and used as a reference against real-time orders to detect exceptional variations in order volumes (Gassman, 2004). Additionally, BAM systems can improve long-term analysis by adding the context of real-time feeds into an Operational Data Store (ODS) (Hackthorn, 2004). This enables a more complete understanding of trends, because it supports an analysis of trends by providing understanding of current states.

# 7 APPLICATION OF THE PROPOSED BAM APPROACH

After understanding differences and similarities between BAM systems and data warehouse, we highlight a set of generic monitoring situations that can be addressed through the combination of real-time events with historical data. These monitoring situations were identified in the course of a graduation project. Table 3 illustrates the problems and analyses of historical data and real-time event information.

Table 3: Monitoring situations that can be addressed by the proposed BAM approach.

| Brief Problem Description | Data Warehouse | BAM |
|---|---|---|
| Benefit from opportunities in Customer Relationship Management (CRM) systems by selecting customers for preferred treatment based on purchase records of customers and provide a higher level of service to the type of customer that is most profitable for the business. | Identify what type of customer is the most profitable for the business. | Generate alert to notify the need to provide service privileges to the selected type of customer. |
| Health surveillance and disease control can be empowered to detect early signs of emerging epidemics. In this way, patient care decisions can be adjusted when patterns of previous epidemics are detected. | Identify patterns of patient admissions in previous epidemics | Integrate patient admission data from several hospitals and searching for pattern that matches a historical pattern of previous epidemics. |
| Refine sales forecast models by adjusting forecasts taking exception events into account. | Extrapolate sales levels from historical data. | Produce new forecast when an exceptional high order is placed that affects the forecast. |
| Support investor’s trade decisions to buy and sell shares in companies. | Calculate historical prices of shares. | Generate alerts when prices are above or below historical levels to indicate the need to buy or sell. Include news and forecast in analysis. |
| Optimize purchase of products in a supply chain. | Calculate historical market prices of products. | Generate alerts when current prices approximate historical limits. |
| Support retailer’s inventory management. | Calculate historical averages for demand for products and services | Identify peaks in demand for products or sales from historical data. |

# 8 CONCLUSIONS AND FUTURE WORK

We have presented an assessment of the practical applicability of a BAM systems and data warehouse. This paper provides substantial evidence that the combination of real-time events with historical data can help improving understanding of the nature of the current problems, which leads to a better support for rapid and adequate response. Additionally, we could identify opportunities for further research. We believe that a very important direction is the need for a thorough search for industry specific problems that are likely to require the noted BAM approach. Another area that requires further research concerns the quantitative measurement of the benefits generated by this BAM approach. Although the benefits that can be obtained through this BAM approach may seem evident, it is necessary to generate grounds for conducting cost benefit analyses.

# REFERENCES


Abramowicz, W., Kalczynski, P., and Wecel, K., 2002. *Filtering the Web to Feed Data Warehouses*, Springer, London.

Brown, J., and Hill, P., 2000. *Data Marts; Key to Reviving the Enterprise Data Warehouse*. In SCN Education, 2001. *Data Warehousing: The Ultimate Guide to Building Corporate Business Intelligence*, SCN Education, Venendaal.

Cavalheiro, G.M.C., 2005. *Defining Business Activity Monitoring: Understanding a Real-Time Event-Driven Infrastructure,* MSc Thesis, Delft University of Technology.

Chandy, M., and McGoveran, D., 2004. The Role of BAM. *Business Integration Journal, Retrieved Septembre 03, 2005 from:* http,//www.bijonline.com/PDF/chandy%20role%20of%20bam%20april.pdf

Defee, J.M. and Harmon, P., 2004. Business Activity Monitoring and Simulation*, Business Process Trends, White paper.*

Delgado, N., Gates, A.Q., and Roach, S, 2004. A Taxonomy and Catalog of Runtime Software Fault Monitoring Tools. *IEEE Transactions on Software Engineering, Volume 30, Issue 12.*

Gassman, B., 2004. How the Pieces in a BAM Architecture Work, *Gartner Document, TU-22-3754.*

Govekar, M., McCoy, D., Dresner, H., and Correia, J., 2002. Turning the Theory of BAM into a Working Reality, *Gartner Document, COM-14-9785.*

Hackathorn, R.D., 2004. The BI Watch: Who's on First. Issue of the DM Review. Retrieved July 12, 2005 from: http://www.bolder.com/pubs/DMR200403-Who%20is%20on%20First.pdf

Hellinger, M., and Fingerhut, S., 2002. Business Activity Monitoring: EAI Meets Data Warehousing. *Business Integration Journal. Retrived August 06, 2005 from:* http://www.bijonline.com/pdf/BAMFingerhut.pdf

Inmon, W.H., Imhoff, C., and Sousa, R., 2001. *Corporate Information Factory*, Wiley, New York, 2nd edition.

Inmon, W.H., Rudin, K., Buss, C.K., Sousa, R., 1999. *Data Warehouse Performance,* Wiley, New York.

Kimball, R., and Caserta, J., 2004. *The Data Warehouse ETL Toolkit: Practical Techniques for Extracting, Cleaning, Conforming and Delivering Data*, Wiley, New York.

Luckham, D., 2002. *The Power of Events: An Introduction to Complex Event Processing and Distributed Enterprise Systems,* Addison Wesley, San Francisco.

McCoy, D., 2003. Blending Business Process Management and Business Activity Monitoring, *Gartner- Strategic Planning, RU.*

McCoy, D., 2002. Business Activity Monitoring: Calm Before the Storm. *Gartner, LE-IS-9724.*

McCoy, D., 2002. Business Activity Monitoring; The Merchant's Tale, *Gartner- Case Studies, CS-16-1780.*

McCoy, D., 2004. The Convergence of BPM and BAM, *Gartner, SPA-20-6074.*

McCoy, D., 2003. Blending Business Process Management and Business Activity Monitoring, *Gartner- Strategic Planning, RU.*

McCoy, D., Schulte, R., Buytendijk, F., Rayner, N., and Tiedricht, A. 2001. Business Activity Monitoring: The Promise and Reality, *Gartner, COM-13-9992.*

Meltzer, M., 1999. *Getting Started; Building the Scalable Warehouse the Right Way*. In SCN Education, 2001. *Data Warehousing: The Ultimate Guide to Building Corporate Business Intelligence*, SCN Education, Venendaal.

Sousa, R., 1999. *Data Warehouse Performance*, Wiley, New York.

Tanler, R., 1997. *The Intranet Data Warehouse: Tools and Techniques for Building an Intranet Enabled Data Warehouse,* Wiley, New York.

Thuraisingham, B., 1999. *Data Mining; Technologies, Tools and Trends,* CRC Press, New York.