Management and Controlling of Time-Sensitive Business Processes with Sense & Respond

Josef Schiefer
Institute for Software Technology and Interactive Systems
js@ifs.tuwien.ac.at

Andreas Seufert
University of Applied Science Ludwigshafen
Andreas.Seufert@fh-ludwigshafen.de

Abstract

The dynamic business environment of many organizations require to monitor their business, IT and organizational processes in real-time in order to proactively respond to exceptions and to take advantage of time-sensitive business opportunities. The ability to sense and interpret events about a changing business environment or customer needs require an event-driven IT infrastructure for making fast and well-informed decisions and putting them into action. In this paper we introduce Sense & Respond loops that support a complete Business Intelligence process to sense, interpret, predict, automate and respond to business processes and aim to decrease the time it takes to make the business decisions. Our approach enables real-time analytics across corporate business processes, notifies the business of actionable recommendations or automatically triggers business operations, effectively closing the gap between Business Intelligence systems and business processes. We propose a system for executing and managing Sense & Respond loops and illustrate our approach with a supply chain business scenario.

1. Introduction

Organizational survival more and more depends on the construction and integration of knowledge fostering the adaptation to the environment. Since Business Intelligence may facilitate the connections within and between organizations, bringing real-time information to centralized repositories and support analytics that can be exploited at every horizontal and vertical level within and outside the firm, it might be very supportive in dealing with today's challenges. However, even though current Data Warehousing and Business Intelligence approaches are widely accepted as state-of–the-art decision support, so far, they do not provide sufficient support in real-time and closed loop decision-making (sense and respond). Important challenges are for example:

- Enabling quick intelligent decentralized decision-making based on current process information and process metrics and key performance indicators.
- Alerting and dynamic process adoption / improvement based on the results of process monitoring.
- Reducing decision time via automated rule-based decision-making.
- Supporting business users in rule-definition in order to easily integrate rule-based decision making models (business analytics).
- Service-oriented architecture in order to easily integrate rules with one another.

This research aims at contributing to the future development of Real-Time Business Intelligence directly addressing some of the issues above. Therefore our paper is structure as follows: Section 2 reviews related work and describes major future challenges for real-time decision support and Business Intelligence. In Section 3 we develop our approach for building a Sense and Respond infrastructure, which is finally (section 4) put into action via a brief description of a business scenario. The paper ends with a conclusion and an outlook for future research.

2. Background and Related Work

Since *Gartner* coined the term zero-latency enterprise in 1998 [16] and a lot of research related to "real-time", such as active warehousing [2], real-time warehousing [3], real-time analytics [14], real-time decision support [18], and business activity monitoring has been published.

The concept of *active data warehousing* [15] has gained great attention. It combines active mechanisms based on event-condition-action rules known from active databases with the integrated analysis capabilities of data warehouse environments to extend (passive) systems with reactive capabilities: An active data warehouse is event-driven, reacts in a timeframe appropriate to business needs, and causes operational actions rather than wait to produce periodic reports [2].

Therefore, the design of has to consider technical aspects (scalability, high availability, frequent (i.e. just-in-time or continuously) data loading, etc.) as well as the integration of active mechanisms which have to deal with two sorts of challenges: a) delays in capturing real-world events by operational systems, and b) delays in loading and integrating data into the data warehouse.

Real-time data warehousing (RTDW) refers to the technical aspects of a data warehouse that updates as data is presented to it [3]. A strict definition of real-time implies that any data change occurring in a source system is automatically and instantaneously reflected in the data warehouse. This implies that all changes in the data warehousing environment take place simultaneously with the change in the source system. RTDW concepts include physical modifications to the database schema and the database environment, movement of data across the enterprise, ETL processes, modification of downstream processes, especially alerts, creation of extracts, cubes and data marts, and the whole new methodology for designing and implementing RTDWs.

Real-time analytics [18] [14] encompasses the ability to use all of the resources in an organization to improve the operations and quality of service, at the moment they are called for. If a piece of information is created or modified in an operational system, it is sensed and acted upon by an analytical process. Real-time analytics complement real-time operational systems. Agile organizations will need to measure, evaluate and react to events with a closed-loop of telemetry-like information, rules, decisions and triggers, all in real-time.

Businesses are striving to achieve "zero-latency enterprise" [16] [19] status, where all data about a business are captured, analyzed and acted upon in real-time. Business Activity Monitoring (BAM) [4] takes business process management beyond the smooth flow of processes across multiple applications to allow analysis and monitoring of business processes and conditions as they happen. BAM applies operational Business Intelligence and application integration technologies to automated processes to continually refine them based on feedback that comes directly from knowledge of operational events. It provides more accurate information about the status and results of various operations, processes, and transactions to improve decisions, address problem areas more quickly, and to support the organization by taking full advantage of emerging opportunities [6].

As organizations interact on a real-time basis and business processes cut across multiple departments and business lines, the need for information integration leads to the adoption of Enterprise Application Integration (EAI) approaches in which message-oriented middleware forms the backbone of the enterprise integration [13].

So far, most of existing data analytical solutions work with traditional (stored) datasets. However, in some recent applications, data may take the form of continuous data streams [1] rather than finite stored datasets. Fields of application include manufacturing processes, click-streams in web-personalization, and call detail records in telecommunication [5]. There has also been some initial work addressing data streams in the data mining community. In terms of building classical data mining models over a single data streams, [10] considers frequent items sets and association rules, [7] considers clustering, and [11] considers decision trees.

3. Building a Sense and Respond Infrastructure

Organizations store data about their customers, suppliers, competitors and their business processes. Examples are customer databases, output figures, production and resource plans, or bill of materials. This data form a foundation for generating new knowledge to improve the intelligence of business processes. Generally this data fulfils certain functions within the organization over a period of time, such as

- **past-oriented** (e.g. documentation, clearly defined business processes, analysis of historical trends)
- **present-oriented** (e.g. fast response to new requirements, timely warning of risky situations) and
- **future-oriented** (e.g. detection and consideration of trends and cycles)

Process participants must have instant access to information which is relevant for the current business context. To be effective, decision support must take a broader view of the whole process of decision-making that is embedded in business processes. One of the key weaknesses of the current generation of workflow management systems and decision support systems is their lack of integration. In this section we propose a real-time Business Intelligence infrastructure called SARI (**S**ense **A**nd **R**espond **I**nfrastructure) which manages the processing of past-oriented, present-oriented and future-oriented data in order to support business processes with Business Intelligence in near real-time. SARI is geared to allow an easy integration with the operational source systems of business processes such as workflow management systems or ERP systems.

SARI's main objective is to help organizations to monitor their business processes and IT systems in

order to proactively respond to business situations or exceptions with minimal latency. SARI is able to continuously receives, processes and augments events from various source systems, and transforms in near real-time these events into performance indicators and intelligent business actions. It automatically discovers and analyses business situations or exceptions and can create reactive and proactive responses such as generating early warnings, preventing damage, loss or extensive cost, exploiting time-critical business opportunities, or adapting business systems with minimal latency. SARI processes data for Business Intelligence purposes based on the current business transactions with minimal latency. SARI takes over the integration and control of the near real-time data processing, such as:

- which data has to be captured at what location, at what time and in which format,
- which data are analysed with statistical methods,
- which analysis methods are applied,
- how data is transformed into business information,
- how business indicators are calculated,
- who are the receivers of the business information,
- which information is required for business decisions,
- which events trigger the decision making processes,
- which responses are appropriate for a decision being made.

During the continuous data processing, SARI is able to gain new knowledge and business information about the business environment with is used to initiate actions in order to optimize the business process performance. SARI is a distributed, scalable platform, which allows to model and execute various forms of Sense and Respond processes (see Figure 1).

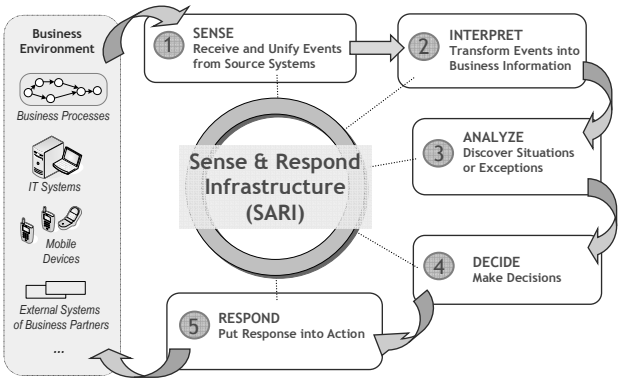

Figure 1. SARI Sense & Respond Loops

The data processing is controlled by "Sense & Response loops" which can be divided into 5 stages. Table 1 describes the function of each phase and how SARI supports it. During the processing in the 5 stages, business information is continuously generated and decisions are made to which a response follows. The response has an effect on the source systems (from which SARI originally might have received events) and consequently also on the performance and the success of the organisation. Since the 5 stages are part of the Sense & Respond loop, the events of actions as part of a response are immediately captured and processed again. With each cycle of the Sense & Respond loop the system improves the information quality and analytical results.

Stage	Description	Function in SARI
Sense	Which is the current state of the business environment?	Events are continuously captured and transmitted to SARI where they are initially unified before the actual data processing begins.
Interpret	What do the captured data indicate? What do the data mean for the current situation of the organisation?	Transformation of the captured events (raw data) into business information, such as key performance indicators, business situations and exceptions.
Analyse	Which business opportunities and risks can arise? What are the possibilities to improve the current situation of the organisation?	Analysis of key performance indicators and determination of root causes for business situations and exceptions. Prediction of the performance and assessment of the risks for changing the business environment.
Decide	Which strategy is the best to improve the current situation of the organisation? What are the actions required to successfully put a decision into action?	SARI proposes the best option for improving the current business situations and determines the most appropriate action for a response to the business environment. This step can be automated with rules or by involving humans.
Respond	Who has to implement the decision? How can the decision be put into action?	Response to business environment by communicating the decision as a command or suggestion (e.g. by e-mail) or by directly adapting and reconfiguring business processes and IT systems.

Table 1. Stages of Sense & Respond Loops

3.1. SARI Architecture

In the recent years, Service Oriented Architecture (SOA) [20] has gained popularity as new software engineering paradigm. It arose from the necessity of creating components providing clearly defined small pieces of functionality that later on can be assembled into complex (usually distributed) applications. Using the SOA approach, we model SARI as a pool of

services (system services and Sense & Respond services) and establish the infrastructure that enables a robust communication and interaction between them.

To achieve the aforementioned requirements, we propose a multi-layer, service-oriented architecture as shown in Figure 2. The underlying infrastructure offers system services, which can be universally used by the Sense & Respond services. The system services fulfill basic tasks such as event correlation, event synchronization, logging, thread management, exception handling and centralized configuration management. The Event Service Bus provides the core infrastructure that enables a robust and flexible communication between Sense & Respond services.

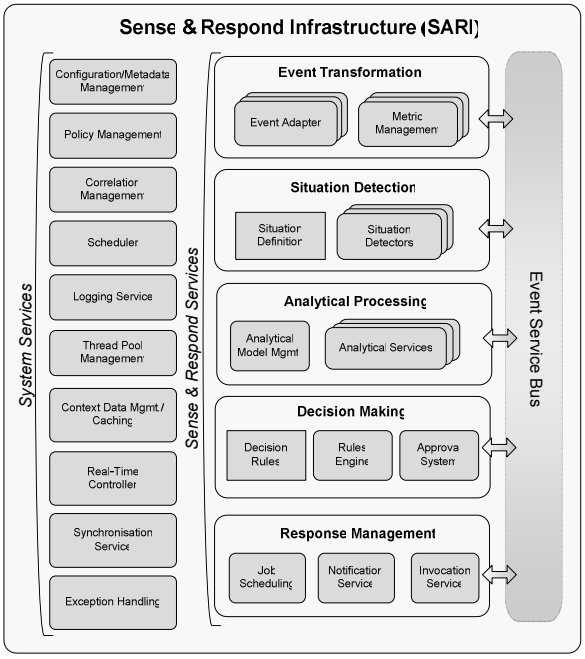

Figure 2. SARI Architecture

Each phase in the Sense & Respond loop introduced in previous section is supported by special Sense & Respond services which can flexibility interact with each other via the Event Service Bus. For instance, the Event Transformation services include Event Adapters to receive events from a business environment in order to transform them into a standardized format (→ Sense phase). Additionally, they include components to manage metrics such as the calculation of performance indicators (→ Interpret phase). The remaining Sense & Respond services correspond to the stages of Sense & Respond loops. With SARI it is possible to include user defined services for various tasks such as discovering situations, a third-party analysis tool as an analytical service or an external rule engine for making automated decisions in Sense & Respond loops.

3.2. Modeling Sense & Respond Loops

The processing steps, their relationships to each other, as well as the parameters of the analysis and data transformation processes can be defined for every organisation individually. SARI uses an event processing model (EPM) for modelling Sense & Respond loops. Similar to a construction kit, the EPM offers various building blocks for Sense & Respond services which can be used to construct a Sense & Respond loop. Dependent on the requirements and the business problem, these building blocks can be flexibly conjoined or disconnected. Links between the building blocks represent a flow of events from one service to the next. The EPM allows for example:

- Definition of the structure for the processed events and data
- Interfaces to external systems for receiving data (Sense) and also for responding by executing business transactions (Respond)
- Data transformations, data analysis and persistence
- Definition of situations and exceptions in data to which a response should be triggered
- Modelling the data and control flows for the Sense & Respond loop (e.g. calculation of a metric always occurs before data analysis)
- Declaration of Sense & Respond services for processing steps including their input and output parameters
- Definition of the relationships and dependencies of Sense & Respond services and event data (e.g. data that has to be correlated before the processing starts)

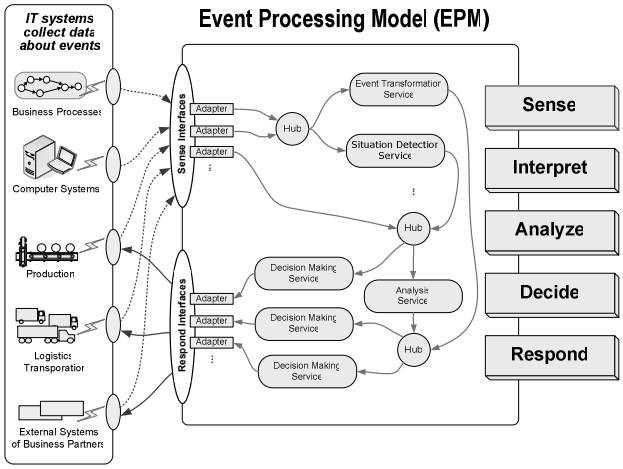

Figure 3. Event Processing Model

Figure 3 shows the EPM of SARI. Data is collected and received from source systems and continuously processed. The EPM shows the event flows of Sense & Respond processes and visualizes complex processing steps. For Sense & Respond loops it is vital that the processing steps work seamlessly together in such a way as to enable a continuous and efficient execution of all processing stages; the EPM offers this capability.

4. Supply Chain Monitoring with SARI

SARI allows a company to analyze and transform business process events in real-time. Current data and business indicators are generated through cyclical processing and the captured knowledge is then redirected to the appropriate business areas. All relevant business information is consistently available to workers and management and can be immediately used for decision making. Business processes are proactively operated, providing companies with the ability to react to business opportunities and exceptions in real-time.

In the following, we show a business scenario where SARI is used to monitor and manage supply chain processes. In our business scenario we assume that the production is very resource intensive or time-sensitive (e.g. the production of special computer chips, chemical substances or perishable goods). Therefore, the company tries to produce as much as possible based on the current demand in order to avoid running the risk of producing the wrong type or quantity of a product. Existing production planning systems determine the bill of material based on demand forecasts. The demand forecast is used to periodically generate the production plan. In our business scenario, we want to improve the accuracy of the forecasts and the production plans by continuously updating them as soon as the demand changes. Figure 4 shows how events from e-Commerce and production systems are processed. In the source systems, events are continuously captured (Sense – Step 1) and transferred to the SARI system. The e-Commerce system continuously conveys completed and cancelled orders. For example, if a customer cancels a large order the demand will change accordingly and the production system is automatically adapted by SARI (Respond – Step 2). Consequently, the required bill of material, which corresponds to the demand, will be automatically adjusted to the current level of demand.

SARI is able to optimize an enterprise's value chain by also monitoring the production processes (Sense – Step 3). Based on the current production capacity, order dates and currently available raw materials, the optimal time and quantity for the replenishment of raw materials can be determined. If SARI discovers raw material shortages, the plant management is immediately informed (Respond – Step 4) and the necessary orders are activated in the ERP system (Respond – Step 5). During the entire Sense & Respond processing the SARI system is continuously generating business metrics that can be used by supply chain managers for monitoring the supply chain.

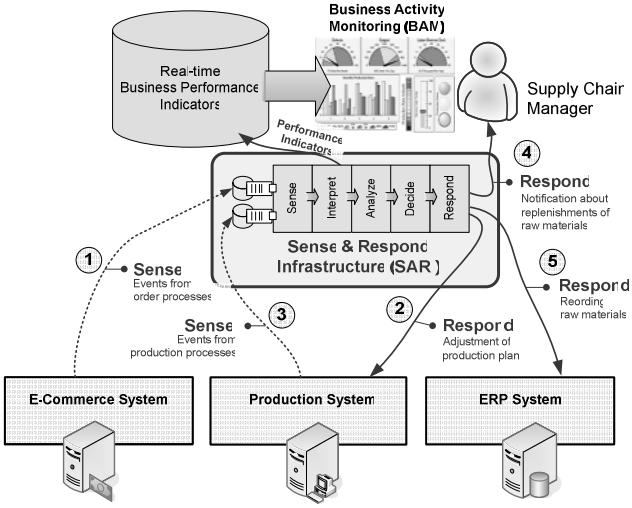

Figure 4. Supply Chain Monitoring with SARI

The least IT support in current supply chain solutions can be identified in the area of decision latency. In most cases, the interpretation of analysis results and the derivation of appropriate actions are seen as manual processes that have to be carried out by supply chain experts and are therefore time consuming. Newest advances especially in the area of Business Activity Monitoring (BAM) try to improve this situation by automating certain decision processes with the help of rule-based decision engines. Based on the real-time analysis of data from an EAI platform, the decision engine checks for predefined business rules and notifies responsible people, or triggers other tools for conducting further actions. In our architecture we follow the approach from BAM and use rule-based decision making for automating many operational and tactical decisions. Please note that the rules for these decisions are in many cases derived from strategic decisions that are made by humans. The decision rules help to intelligently respond very quickly to the current business situation.

One of the key goals of SARI is to minimize the latency for responding to changes of the supply chain. There are two kinds of response latencies: 1) the time it takes to initiate an action and 2) the time it takes to

execute and monitor the action. *Hackathorn* considers in [7] the response latency as part of the decision latency which is reasonable for strategic decision making. We consider separately in our model the time it takes to initiate and execute an action in the business environment. For this reason, SARI includes a response management module which is responsible for triggering supply chain operations and monitoring their outcome (e.g. replenishment of raw materials).

5. Conclusion and Future Work

Conventional business process management and decision support systems do not provide the low latencies needed for the management and decision making of time-sensitive business processes. This paper described SARI, a real-time Business Intelligence infrastructure with the aim of providing continuous, real-time analytics in order to enable proactive responses to a business environment for effectively managing and controlling time-sensitive business processes. We introduced a model for Sense & Respond loops and a service-oriented architecture that is able to detect situations and exceptions, perform complex analytical tasks and reflect on the gap between current situations and desired management goals. Traditional Business Intelligence architectures lack in the support of real-time BI and closed-loop decision making. A major goal of SARI is to tightly integrate Business Intelligence with business processes by monitoring various IT system or other observables and generating reactive and proactive responses such as generating early warnings, preventing damage, loss or extensive cost, exploiting time-critical business opportunities, or adapting business systems with minimal latency. The work presented in this paper is part of a larger, long-term research effort aiming to develop a service-oriented Business Intelligence platform for supporting time-sensitive business processes.

6. References

[1] Babcock, B., Babu S., Datar, M., Motwani, R. Widom, J. Models and Issues in Data Stream Systems. In Proc. of the 21st ACM Symposium on Principles of Database Systems (PODS), pp. 1–16, Madison, Wisconsin, May 2002.

[2] Brobst, S. and Ballinger, C., Active Data Warehousing. Whitepaper EB–1327, NCR Corporation, 2000.

[3] Bruckner, R. M., List, B. and Schiefer, J., Striving Towards Near Real-Time Data Integration for Data Warehouses, In Proc. of the 4th Intl. Conf. on Data Warehousing and Knowledge Discovery (DaWaK 2002), Springer LNCS 2454, pp. 317–326, Aix-en-Provence, France, Sept. 2002.

[4] Correia, J., Schroder, N. BAM: A Composite Market View, Gartner Research Brief, August 2002.

[5] Chen, Q., Hsu, M., Dayal, U. A Data-Warehouse / OLAP Framework for Scalable Telecommunication Tandem Traffic Analysis, In Proc. of the 16th Intl. Conf. on Data Engineering (ICDE), IEEE CS Press, pp. 201–210, San Diego, CA, Mar. 2000.

[6] Grigoria, D., Casatib, F., Castellanosb, M., Dayalb, U., Sayalb, M. and Shan, M.-C., Business Process Intelligence, in: Computers in Industry 53, 2004, pp. 321-343.

[7] Guha, S., Mishra, N., Motwani, R., O'Callaghan, L.: Clustering Data Streams, In Proc. of the 41st IEEE Annual Symposium On Foundations of Computer Science, pp. 359–366, Redondo Beach, CA, Nov. 2000.

[8] Hackathorn, R., Current Practices in Active Data Warehousing, http://www.dmreview.com/whitepaper/WID489.pdf accessed 10 March 2005.

[9] Hammer, M. and Champy, J., Reengineering the Corporation, Nicholas Brealey Publishing, London, 1993.

[10] Hidber, C. Online Association Rule Mining, In Proc. of the ACM SIGMOD Intl. Conf. on Management of Data, pp. 145–156, Philadelphia, PA, June 1999.

[11] Hulten, G., Spencer, L., Domingos, P. Mining Time-Changing Data Streams, In Proc. of the ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining, pp. 97–106, San Francisco, CA, August 2001

[12] Inmon, W.H., Building the Operational Data Store, 2nd edition, Wiley: New York et al. 1999.

[13] Martin, W., Business Performance Management – Efficiently Managing Business Processes. Research Bulletin, 2003, http://www.it-research.net, accessed 12 May 2003.

[14] Raden, N., Exploring the Business Imperative of Real-Time Analytics, Teradata white paper, October 2003.

[15] Schrefl, M., Thalhammer, T. On Making Data Warehouses Active, In Proc. of the 2nd Intl. Conf. on Data Warehousing and Knowledge Discovery (DaWaK), Springer LNCS 1874, pp. 34–46, London, UK, Sept. 2000.

[16] Schulte, R. W., Introducing the Zero-Latency Enterprise. Gartner Research, June 1998

[17] Schulte, R. W., Application Integration Scenario: How the War is Being Won, in: Gartner Group (Ed.): Application Integration – Making E-Business Work, London, 6-7 September 2000.

[18] Schwenk, H., Realtime analytics: the future for Business Intelligence, Ovum 2001.

[19] Young, J., Schulte, R. W., The Zero-Latency Framework, Gartner Research, April 2000.

[20] Natis, Y., Service-Oriented Architecture Scenario, Gartner Research, Number: AV-19-6751, 2003.