Embedded Analytics in Front Office Applications

Erik Nijkamp, Martin Oberhofer

IBM Research & Development GmbH
Schönaicher Straße 220
71032 Böblingen
{nijkampe, martino}@de.ibm.com

Abstract: Today, decision making by users of front office applications happens without analytical information supporting this process. We propose as solution Embedded Analytics (EA) making analytical insight in context of the front office application available to users in an on demand fashion. "In context" means that only analytics relevant for decisions on the current User Interface (UI) screen is made available. "On demand" means that the user has the information accessible in "mouse-over" events which means the user decides when to consume which portion of the relevant analytical information. The underlying implementation uses the Unstructured Information Management Architecture (UIMA) to determine the context and then dynamically binds real-time lookup services for the delivery of the analytic insight to the application UI using Web 2.0 techniques in the UI. We explain the technologies and algorithms used for Embedded Analytics in context of business processes running in a SAP Netweaver CRM system.

1 Business Problem

Users of front office applications such as call center or customer support applications make millions and millions of decisions each day without analytical support. For example, if a support employee gets a new support ticket and needs to decide how much time should be used for problem resolution and which measures should be taken, this is done without analytical insight such as customer revenue, customer net profitability or customer potential. There are two barriers preventing the use of analytical information derived in BI systems: The first barrier is related to the constraints of the user group. Front office application users usually lack the BI analyst skills to work with complex BI tool sets to create reports. Running the daily business process operations has priority – there is no time to build and run analytic for each and every decision which needs to be made on a case by case basis s. The second barrier is the fact that the use of BI is not yet tailored to bridge the gap between the operational and analytical systems. This would require the real-time delivery of analytical information as a seamlessly consumable service within a business process. As a result, companies cannot optimize their front office departments because analytical insight derived in Business Intelligence (BI) systems or other systems is not available to users in these departments. A solution to this business problem must therefore deliver analytical information in context of the business process step the user performs without any additional skill set requirements at the time the information is needed in real-time.

2 Business Intelligence and Embedded Analytics

Business Intelligence[1] encompasses methods and technologies which support management in decision making by providing analytical insight into operational data. Looking at BI from a process perspective [GG00] then BI can be considered as a process creating analytical insight into the own company status and perspectives in the market place from heterogeneous internal and external source systems. Therefore, BI can be considered the discipline consisting of methods and processes to systematically analyze internal and external data in a electronic fashion with the goal to support management in strategic and tactical decision making supporting the overall company goals.

2.1 Business Intelligence – traditional view

Since the 1960's analytical solutions supporting decision making have been known by different names such as Decision Support Systems (DSS), Executive Information Systems (EIS), Data Warehouses (DW) and today they are known as BI solutions [HW05]. BI systems need to consume operational data from Customer Relationship Management (CRM), Enterprise Resource Planning (ERP), Supply Chain Management (SCM) and many other application systems with an overall maximum size in three digit Tera-Byte volumes at large enterprises. The BI process is usually structured as shown in Figure 1. It extracts operational data from various source systems, cleanses and transforms it and then loads it into a Data Warehouse. This first step is known as Extract-Transform-Load (ETL) process step. Once the DW is populated data analysis is done using approaches such as Online Analytical Processing (OLAP) or Data Mining.

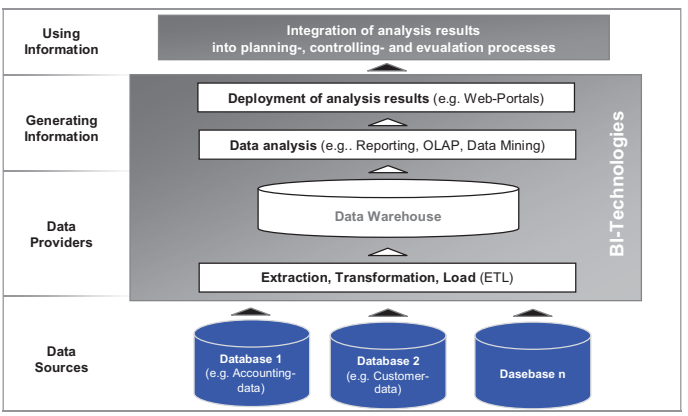

Figure 1: BI related aspects and technologies [Gen06]

The analytical results are summarized in BI reports which are sometimes deployed into intranet Portal applications used by managers to access the BI results supporting their decision making. Through the BI process operational data has been transformed into

[1] The Term „Business Intelligence" was created by the Analyst Howard Dresner in 1989.

analytical data. The key difference between operational and analytical data is that the former is used to perform a business process and the latter is used to support decision making to control the company.

Today, BI is a default component in an IT department in almost all companies and analysts project an increase in budget for BI solutions during the next year[2]. Even though BI is used in many companies, a study performed by SAP Labs concludes that BI so far reached only 30% of its potential. The study concludes in the best case only 3 out of 10 actors with a need for analytical information have access to it [Sch07]. The unresolved issue for many companies is the transfer of analytical results into their daily operations – in essence to the users of their front office applications. Simply put: The quality of the analytical information today is very high – but there is a lack of mechanisms to distribute this analytical information to the vast majority of users in need of it. This created the need for operational Business Intelligence as another sub-discipline in addition to strategic and tactical BI. The key characteristics of the three sub-disciplines are shown in Table 1.

Table 1: Subdisciplines of BI [Imh06]

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

While companies successfully use strategic and tactical BI already, operational BI is not yet there. Even worse, terms like “real-time” or “right-time” BI for operational BI solutions are imprecise from a BI process perspective as previously outlined. This means it is not clear whether or not just the ETL step is in real-time or the analytical processing within a DW or the entire process. In essence, it is not yet clear which process steps are performed in real-time which we further investigate in the next section.

2.2 Introduction of Embedded Analytics

Latency must consider regarding the overall BI process data latency, analytical latency and decision latency as shown in Figure 2. Several latencies accumulate between an original business event and the point where related decisions must be taken using information form the event.

[2] The Gartner Group projects an increase from 2.5 billion dollars in 2006 to 3 billion dollars for BI solutions in 2009. [Crm06]

The goal of operational BI is to increase agility and reduce the latency in the decision making process based on an as accurately as possible available data in order to improve agility supporting change in business requirements.

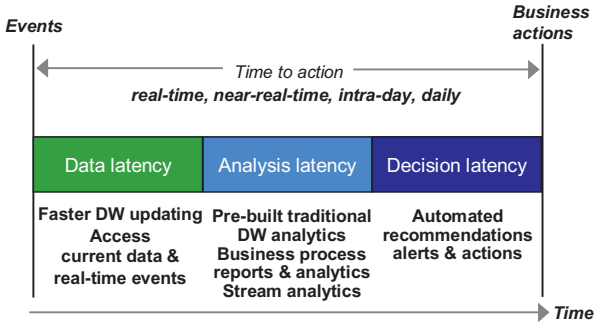

Figure 2: Latency in the BI process [WD08]

Responsible for applying the analytical insight in daily decision making are the users of the front office applications. For up to date analytical information naturally the desire arises to reduce the latency in analytical processing as much as possible with the ultimate goal to perform as close as possible to real-time. We therefore conclude that operational Business Intelligence is a trend towards a set of services, methods [WD08] and applications which monitor, analyze and optimize business processes on a daily basis. With this understanding of operational BI we can now define Embedded Analytics: Embedded Analytics (EA) is a processing environment where in a operational business process for optimization purposes real-time or near real-time analytical services support intra-day decision making by users of front office applications on a daily basis. At the core of EA is therefore the harmonization of the analytical with the operational application environment through an establishment of a direct communication channel between these two environments. This allows delivery of analytical information in context at the right time to the users of the operational applications. Separating Embedded Analytics into the term "Embedded" and "Analytics" as suggested by Eckerson [Eck06] we get a sharp focus on the essential meaning:

- Operational Process: A business process in execution
- Analytics: A process to convert data into analytical insight
- Embedding: The invocation of analytical insight from an operational process

In the following we therefore Embedded Analytics as "process-centric user analytics" with real-time latency.

2.3 The value of Embedded Analytics

The goals and therefore the potential of EA are the seamless distribution into the context of operational applications to allow "BI for the masses". Hence EA is simple to use in the context of the operational applications. A context switch from an operational application UI to a BI application UI is avoided because the analytical insight is delivered within the UI of the operational front office applications. This delivery reflects

the role and permissions of the user in the current context of the business process. Associating the BI results with operational processes enables the front office application users to make better decisions on a daily business due to improved, available business insight. This benefit materializes from two opposite directions: On the one end, potential problems can be detected before they become real problems and mitigations steps are possible preventing the real problem before it happens. On the other end, potential, temporary business opportunities can be identified on which a user can act in a daily step of a business process.

Implementing EA also requires a modularization of BI into BI analytical services. Once this step is done, these BI services can be more easily consumed in existing or new applications with the result of re-use and improved flexibility in the IT infrastructure. With EA business processes can be optimized and the operational and analytical environments become connected in a closed loop.

3 A Formal Reference Architecture for Embedded Analytics

„You will waste your investment in SOA unless you have enterprise information that SOA can exploit ...“ Andrew White

According to A. Whites assumption a service-oriented architecture should incorporate not only functionality of online transaction processing systems (abbr. OLTP) such as ERP related services, but also analytical services. Being able to exploit valuable enterprise information in orchestrated processes is crucial to fully unfold the potential of business intelligence frameworks.

The multitude of involved components and their dependencies represents a challenging problem: How to address a broad diversity of architectural and technical aspects in a holistic and integrated manner [BD07]? One approach to solve this issue is the stepwise decomposition of operational and analytical functionality into architectural layers. Within this stack a layer provides services which consume elements of the layer beneath and provide functionality to the layer above. This procedure reflects the central idea of a ***SOA***: gradually abstracted services allow the simple composition of applications and processes.

In the scope of this paper embedded analytics is characterized as a method which can be used to provide a channel delivering context relevant business information (respectively a subset of metrics) to operational environments in a customizable way. In reference to the procedure described previously the following figure 3 shows the ***embedded analytics method*** decomposed into several components.

Notice that the proposed embedded analytics solution takes two different types of users into consideration: operational members of staff being responsible for the day-to-day business and business analysts who are able to identify trends and to express economic activities as a formal model.

From an architectural point of view a set of interacting components is required to realize an embedded analytics solution.

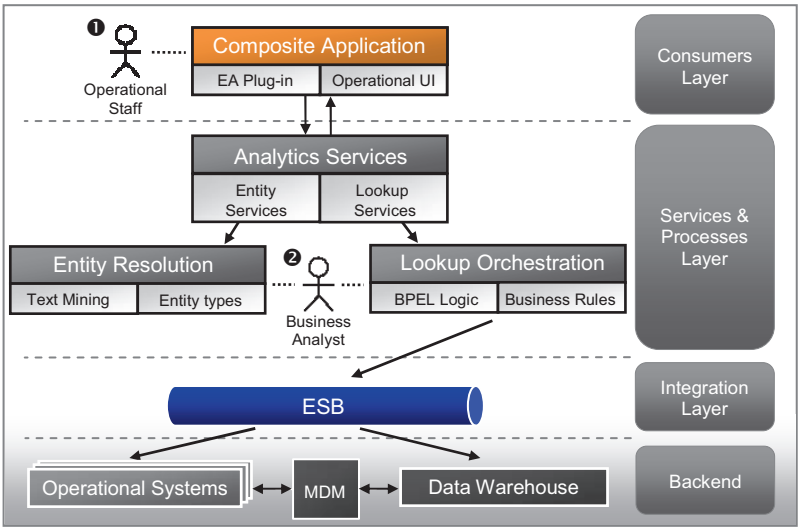

Figure 3: Proposed reference architecture for embedded analytics

Going from top to bottom, a ***composite application*** consists of the two crucial parts. The operational functionality required to manage the day-to-day business and an embedded analytics extension which provides analytical insight to support decisions instantly. In order to provide additional analytical information the embedded analytics ***plug-in*** transforms the operational user interface into a serialized representation and invokes an entity service to reveal the informative objects[3]. An entity is a type which has an identity and, if necessary, is used to transport additional attributes. Exemplary, a real-world entity such as a customer can be referred to in different ways (e.g. different name spellings). The goal of ***entity resolution*** (abbr. ER) is to resolve entities, by identifying associations between a set of varying representations and the informative object. Subsequently, ER algorithms such as text mining[4] are utilized to analyze the user interface and compare the given representation with a set of predefined entity types. Based on the retrieved list of entities the operational user interface is extended with anchors (e.g. adding mouse listeners on widgets). Whenever an event is triggered the embedded analytics plug-in calls the ***lookup service*** being associated with the given entity type. The lookup service inputs the information providing the identity of the requested entity and returns a refined representation which contains decision supporting facts. For example the lookup service could invoke a ***Business Process Execution Language (BPEL)*** process which retrieves additional customer data (e.g. historical revenue growth) delivered by a DW. In between a ***master data management*** system holds a single version of truth of master data to guarantee that the backend system can

[3] Informative in the sense of being useful for queries e.g. a product id which can be used to retrieve details.
[4] In this context we refer to the ability of text mining algorithms to determine which actual person or object a particular reference refers to, by looking at natural language text.

handle the lookup and associated entity. Viewed from a bottom up perspective operational[5] and analytical[6] systems provide services which are typically accessed in a centralized manner – this is the role, among others, of the ***enterprise service bus***.

Furthermore, the architecture incorporates a mechanism enabling business intelligence analysts to define the ***analytic logic*** in a customizable way. Business rules thus configured can be directly incorporated into reports used by the operational day-to-day business. For example figure 4 shows analytical logic defined as visual flow using BPEL as notation.

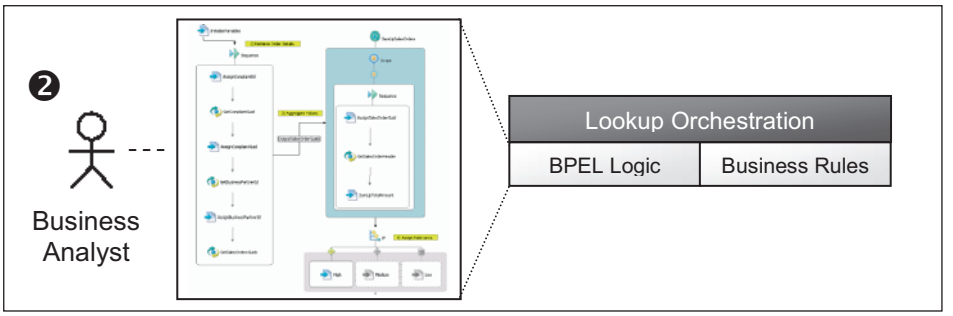

Figure 4: Analytical logic defined as a visual flow (BPEL)

Remarkably, this process of ***alignment*** between strategic/tactical guidelines and operational activities does not involve developers or administrators precisely because the logic itself can be defined by an analyst. In other words strategic and tactical goals are indirectly injected into the operational environment which affects the actual process execution and consequently optimizes the operational activities.

In addition it can be easily seen in figure 3 that the introduced layer required to define analytical logic acts as an ***abstraction layer***. Frequently, trends or substitutions are reshaping the IT landscape which affects the associated application interface. In order to react to these changes of the underlying infrastructure an analyst would ***redefine*** the analytical logic. However, the analytical output used by the embedded analytics services stays untouched. In fact this mechanism accepts mutable input parameter types while providing an immutable result – decoupling the operational or analytical systems from high level components. For instance replacing the Data Warehouse system does not affect the implementation of an arbitrary lookup service.

The advantages of this approach are as follows:

- the analytical insight of domain experts is directly injected into the daily business
- operational activities are aligned to tactical or strategic goals (e.g. business rules)
- decoupling of the underlying systems from the embedded analytics infrastructure
- a domain expert is able to modify to the analytics without involving a developer
- the visualization of analytical logic allows non-technical individuals to understand the factors which "drive the business"

[5] E.g. Enterprise-Resource-Planning or Customer-Relationship-Management Systems.
[6] Resp. Data Warehouse Systems.

4 Integrating Analytics into Complaint Management Processes

„Once a good product was all you needed. Today, the totality of experience has equal impact.“ Duane Sharp

This chapter describes how to utilize an existing operational process or user interface as vehicle for predefined embedded analytics. The key idea is to deliver target-oriented analytical information (e.g. metrics) to the responsible employee in such a way that any crucial Customer Relationship Management (abbr. CRM) process can be improved by supporting a decision at hand. In order to increase employee effectiveness the following aspects are addressed:

- analytical results are computed automatically in real-time as needed
- the analytics are available in relevant portions or as whole to the employee “on demand” in an easy to use fashion, not flooding the employee
- the analytical insight is made available in the context of the user interface where the decision processing takes places

The idea is that an operational environment implicitly defines a context which can be used to establish a link between a set of business intelligence reports and the decision to be made. Besides providing concise reports as indirect decision support, the approach can be extended to offer an appropriate set of recommendable decisions (a subset of all possible decisions). Based on the analytical logic (e.g. logic which incorporates tactical business rules) the embedded analytics plug-in could manipulate the CRM user interface to only allow decisions with are conform to the tactical goals or policies.

Figure 5 shows an embedded analytics prototype which enriches the SAP CRM web form that handles customer complaints with highlighted annotations and analytical insight. In this case a regular HTML website represents the operational user interface.

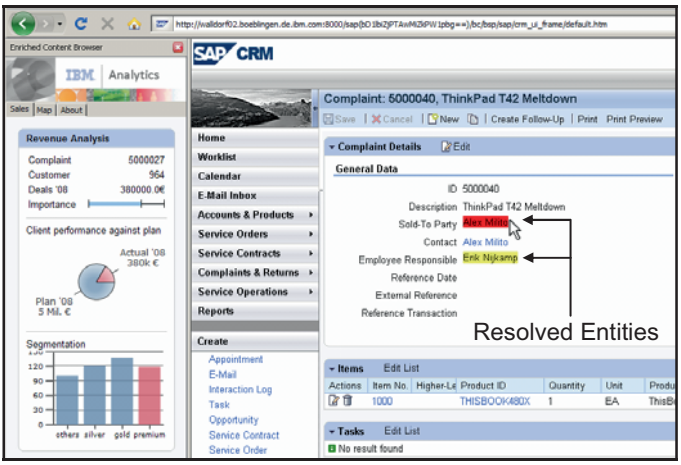

Figure 5: SAP CRM user interface extended with analytical information

The embedded analytics functionality is provided by a browser plug-in which analyzes the DOM tree and renders the analytical information. For example, in figure 3 the user pointed at the highlighted customer name and obtained additional revenue information.

Figure 6 describes the underlying method as a sequence of operations transforming a given input to output which is used by the subsequent operation.

1. The customer complaint form is analyzed for entities that are relevant to the users task (e.g. customer numbers, phone numbers)
   a. The browser plug-in draws the plain text content from the HTML document skipping all document mark-up information. However, the mark-up information can not simply be discarded as the relation between a text passage and its mark-up must be obtained. The document's plain text content is send to the entity services (UIMA annotators).
   b. The annotators process the plain text document, identify entities according to regex patterns and return a set of objects including the position within the HTML document.
2. The enriched document content is displayed in the web browser
   a. The original document and the text analysis results coming from the UIMA annotators are interwoven. The entities detected by the text analysis annotators are highlighted in different colors and extended with additional client side code that triggers one or multiple actions when the user moves the mouse over a detected entity.

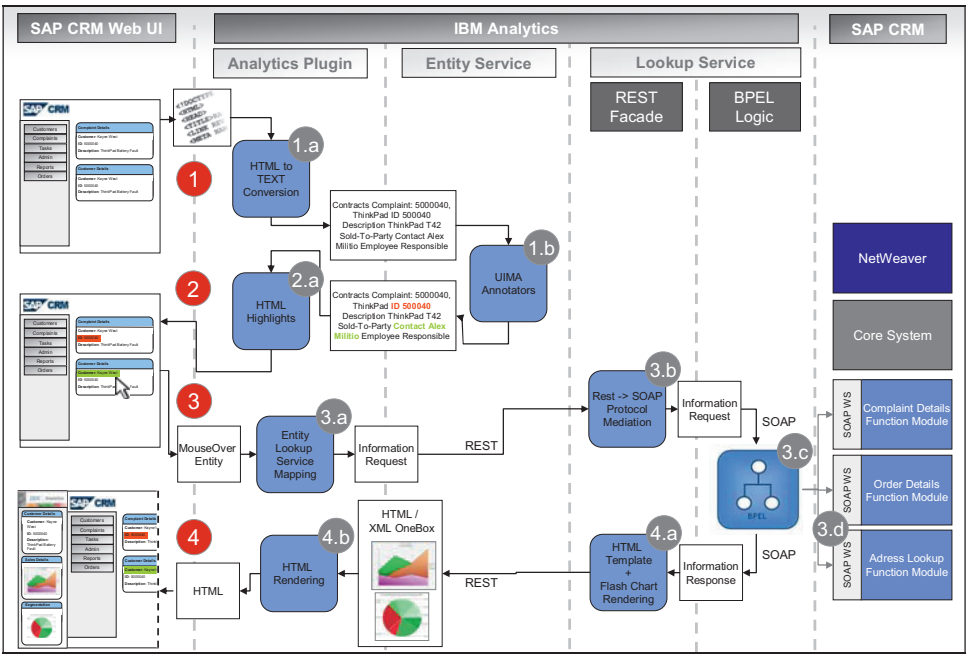

Figure 6: Analytical logic defined as visual flow

3. The browser renders the new document and the end-user can easily extract additional information from various back-end systems by simply moving the mouse over an entity.
   a. The user moves the mouse over a customer's name, which triggers a call-back function.
   b. In case the original function call was specified in the REST format, a transformation to a SOAP call is possible. This mediation can be useful as REST is very popular among web developers whereas SOAP is very common in SOA environments.
   c. The SOAP request is fed into a Workflow system where it kicks off a workflow process that could e.g. be specified in BPEL.
   d. The workflow process itself can contain some complex logic specifying which backend systems to call under what conditions and for what results.

4. The information gathered from one or multiple backend systems is displayed to the user and help's him/her to decide how to handle this customer complaint best.
    a. The output of the previously called business process is transformed into a mark-up document so that it can be displayed by a web browser (e.g. charts)
    b. The browser displays the information generated by the lookup service, the user can benefit from the additional information.

Notice that the mentioned Apache UIMA framework provides functionality which allows us to identify entities in unstructured information (e.g. HTML code). This prototype extracts the analytical information directly from the transactional SAP CRM instead of accessing a Data Warehouse for DW. Nevertheless this demonstration shows how embedded analytics can be used to optimize the complaint process by providing decision supporting information.

5 Summary and Future Work

In the context of the complaint management process in a SAP Netweaver CRM system EA delivered on its value promise which means the user get the analytical information in context at the right time on demand. As a net result, the user can now decide for which customer it is worth spending more time and efforts in resolving the complaint based on the customer revenue metrics and other analytical information. This improves customer satisfaction in the profitable and important customer segments and supports therefore key objectives of effective customer relationship management. We also conclude that Embedded Analytics will not replace the need for strategic and tactical BI solutions because it doesn't address the needs of the users with these requirements. It is in our perspective a complementary, new BI discipline enabling BI for the masses which do not have the BI analyst skill set.

If the DW provides access to accurate customer segmentation based on revenue metrics to a front office application, the DW needs to be able to provide this analytical insight in real-time. Building a DW system with seamlessly consumable services capable of returning this in a real-time fashion requires further research.

References

[GG00] Grothe, M. / Gentsch, P.: Business Intelligence - Aus Informationen Wettbewerbsvorteile gewinnen, Addison-Wesley, P. 11-13, Munich 2000.

[HW05] Humm, B. / Wietek, F.: Architektur von Data Warehouses und Business Intelligence Systemen, Springer, Munich 2005.

[Gen06] Gentsch, P.: Business Intelligence for better decisions, http://www.intelligence-group.com/downloads/BI-for-better-decisions.pdf, 2006.

[Crm06] CRM Today: Business Intelligence Software Market to Reach $3 Billion in 2009,, http://www.crm2day.com/content/t6_librarynews_1.php?news_id=117297, 2006.

[Sch07] Schubert, L.: Embedded Analytics Are Not a Myth, in: SAP Insider, P. 12-16, 9/2004.

[Imh06] Imhoff, C.: Operational business intelligence, http://www.teradata.com/tdmo/v06n03/Viewpoints/EnterpriseView/OBI.aspx, 9/2006.

[WD08] White, C. / Davis, J.: Using Embedded Business Intelligence and Analytics for Near-Real-Time Decisions and Actions, http://www.beyeresearch.com/study/7561. 6/2008.

[Eck06] Eckerson, W. (2006): Embedded Analytics, http://download.101com.com/pub/tdwi/Files/TDWI%20Monograph%20Embedded%20Analytics.pdf, 5/2006.

[BD07] Berson, A.; Dubov, L.: Master Data Management and Customer Data Integration for a Global Enterprise, The McGraq-Hill Companies, 2007.