RC23580 (W0504-033) April 8, 2005
Computer Science

IBM Research Report

Business Performance Management System for CRM and Sales Execution

Markus Ettl, Bianca Zadrozny, Pawan Chowdhary, Naoki Abe
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Business Performance Management System for CRM and Sales Execution

Markus Ettl, Bianca Zadrozny, Pawan Chowdhary and Naoki Abe
IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA
{msettl,zadrozny,chowdhary,nabe}@us.ibm.com

Abstract

In 2004 the IBM Telesales organization launched a new customer segmentation process to improve profits, revenue growth and customer satisfaction. The challenges were to automatically monitor customer segment status to ensure results are in line with segment targets, and to automatically generate high-quality predictive analytical models to improve customer segmentation rules and management over time. This paper describes a software solution that combines business performance management with data mining techniques to provide a powerful combination of performance monitoring and proactive customer management in support of the new telesales business processes.

1. Introduction

As companies adopt more customer-centric business processes, customer segmentation strategies are the key to balancing the delivery of optimal customer experiences with the need to maximize long-term customer revenue and loyalty. However, developing an effective customer segmentation strategy is a serious challenge for many companies since they often lack the ability to track and interpret the huge volumes of customer data that exist throughout the enterprise. Many companies also lack real insight into their customer base to detect changes in overall buying behavior and customer value which makes it impossible to evaluate the long-term success of their customer segmentation efforts.

Business Performance Management (BPM) enables a company to keep up with a rapidly changing business environment. In its quest to preserve or expand competitive advantages that are threatened by intense competition and changing customer preferences, a company needs to continuously monitor critical business processes, detect existing or anticipated business issues, and respond swiftly to opportunities and threats [5,7,8]. To do so, advanced business process monitoring capabilities and decision support are essential.

In this paper we describe a BPM system developed for the IBM Telesales organization to support Customer Relationship Management (CRM) and sales execution. Our BPM solution provides an analytical and information infrastructure to support the segment strategy. It automates the demanding task of monitoring customers and telesales center performance with predefined business metrics and customer segments. Customer segments are defined by customer behavior, demographic attributes and customer value. Business metrics include profitability and sales targets, recency and frequency of purchases, and monetary value scores. Our analytical modeling techniques help telesales executives create more accurate customer segments and predict which customers are likely to churn which in turn improves profitability, customer satisfaction and revenue contributions.

The system leverages IBM's BPM architecture [2,3] to develop an end-to-end performance management framework. It comprises several components. First, a data warehouse that stores several years of historical customer purchases (laptops, desktops and options) on several thousand customer accounts and millions of sales transactions; second, an ETL (data extraction, transformation and loading) module that provides daily access to a production data mart to load new sales transactions into the data warehouse; third, a web-based management dashboard that captures key operational metrics indicating the effectiveness of the segmentation model against planned business objectives; fourth, a novel application of predictive modeling technology that is optimized for discovering recommendations for customer advancements through the segments based on previous customer spend and analytical correlations. These models help realize an early-warning alert system, flagging customers who should be moved to another segment in the near future based on their most recent buying behavior.

The paper is organized as follows. Section 2 gives an overview of the customer segmentation process. Section 3 presents the predictive modeling and mining techniques that enable detection and alerting of customer movements. Section 4 shows the solution architecture of our BPM system, focusing on functional components and the data warehouse design. Section 5 concludes this paper.

2. Customer Segmentation

Customer segmentation plays a pivotal role in providing personalized customer interactions. Segments are groups of customers that admit uniform or similar treatment, often because they have similar interests in products or services, judging from their demographics and buying behavior [4].

The segmentation strategy implemented at the IBM Telesales organization classifies customers into the three segments *acquisition*, *development* and *retention* (ADR). The segment definitions are based on a combination of soft and hard criteria that encompass RFP bid histories, revenue scores and regular buying patterns. The goal was to develop a common management and measurement system that empowers strategic planning with a single enterprise-wide customer view across the different brands, regional sales teams, and telesales operations.

In addition to providing customer-centric metrics, the measurement system defines a set of financial and operational metrics for each segment, and establishes effective and measurable goals. These metrics include:

- revenue, revenue growth and profit
- customer satisfaction
- web penetration (percent of revenue generated from web purchases)
- options clothing rate (percent of revenue generated from options)
- customer buying frequency

Our BPM system combines infrastructure components with analytics to create accurate customer segments and predict customer movements to enable faster, more relevant feedback-to-action cycles. The system provides alert-based notifications of customers that advance or retreat through the segments, based on customer behaviors and analytical correlations. It also provides alert-based recommendations of potential actions that can accelerate a customer's movement to the next segment, and detects and alerts users of dormant customers.

The main goal of predictive data mining in this system is the automatic creation of rules for suggesting when a customer should be moved to a different segment. In the past, assignments of customers to segments had been done by telesales managers based on their knowledge of the past behavior of the customer and their expectations for the future behavior of the customer. This process was slow, error-prone and, very often, customers were not moved to the appropriate segment in a reasonable time frame.

The ad-hoc nature of the segment assignments made it impossible to directly obtain exact rules for alerting when customers should be moved between segments. Instead, by using predictive data mining, we automatically extracted rules from the available data about the customers and their segment assignments, without resorting to expert's knowledge. These rules relate characteristics about the past purchase behavior of a customer to the segment that was assigned to the customer. In a sense, they mimic the segment assignments that were done manually in the past and generalize them to new situations. Thus, when there are changes in the behavior pattern of a customer, the rules can be used to identify a more appropriate segment and the system will issue an alert suggesting a segment move.

3. Predictive Data Mining

In our data mining framework, the purchase behavior of a customer is defined by the metrics listed above. For each of these metrics, we calculate variables containing moving averages (over the last 3 months, 6 months and 12 months) and trends corresponding to the difference between the average in the last 3/6/12-month period and the respective period before that.

We also include two demographic characteristics of the customers, i.e., sales region (nordics, central, etc.) and industry sector (finance, distribution, industrial, public, etc.). We included these characteristics because we expect the purchase behavior of customers in the same segment to be different depending on the region and the sector. Together, the purchase behavior variables and the demographic variables form a customer profile.

3.1. Data Mining Process

The data mining process is divided into two phases:

- **Training Phase:** Automatically obtaining segmentation rules from historical customer profiles with manually assigned segments.
- **Scoring Phase:** Applying rules to updated customer profiles and issuing alerts when the customer needs to be moved to a different segment.

In the training phase, we use a general purpose predictive data mining engine developed by IBM Research called ProbE [1]. More specifically, we use ProbE's *Stochastic Decision Tree* learning algorithm. This algorithm identifies the profile variables that are most significantly correlated with the response variable (in this case, the customer segment) and uses these profile variables to split the customers into coherent groups, creating a decision tree for assigning segments to customers. This decision tree can be translated into a set of rules for segment assignment based on the customer profile variables.

In the scoring phase, we update the customer profiles and apply the rules obtained using ProbE in the training phase. Instead of outputting hard segment assignments, the rules generated by ProbE assign a probability to each of the possible segments (A, D and R). We experimented with different thresholds and decided on the following guidelines for alerting users to move a customer from its current segment to a new segment:

- **High severity:** If the absolute difference in probability between the segment with the highest probability and the current segment is equal or greater than 0.4, we issue an alert with high severity.
- **Low severity:** If the absolute difference in probability between the segment with the highest probability and the current segment is between 0.2 and 0.4, we issue an alert with low severity.
- **Exception:** If the customer was not assigned a segment in the past, we issue an exception alert with the suggestion for the segment with the highest probability.

When issuing an alert, the algorithm provides the two most important customer metrics that appear in the rule that suggested the new segment. For these two metrics, graphs are displayed on the dashboard depicting the behavior of the metric for that customer in the past three months. The graphs help telesales managers better understand why the customer should advance or retreat to a different segment.

3.2. Evaluation

The rules obtained through the data mining process attempt to mimic the manual segment assignments. However, because the manual assignments are noisy and they possibly use information about the customers that is not available in our customer profiles, it is not possible for the rules to perfectly mimic the manual assignments.

A confusion matrix allows measuring how well the rules are mimicking the manual assignments. Each entry *(i,j)* of the matrix shows how many customers were assigned to segment *j* by the rules, when their manually assigned segment is *i*.

To quantify the performance summarized in a confusion matrix, we can use either of the following two measures:

- **Macro-accuracy:** sum the diagonal elements in the matrix (correct predictions) and divide by the total number of customers. This gives equal weight to each customer.
- **Micro-accuracy:** first divide each number in the diagonal (correct predictions) by the total number of customers in the corresponding segments, and then we average the resulting numbers. This gives equal weight to each segment.

Because the A segment is much smaller than the D and R segments, the macro-accuracy metric does not give enough consideration to the prediction performance in the A segment. For this reason, we decided to use micro-accuracy as the main performance metric for evaluating segmentation rules.

We obtained a first set of rules by applying ProbE to the original customer data. The confusion matrix for these rules is shown in Table 1. As can be seen in the matrix, the rules never predict segment A and almost always predict segment R. This is likely to be a consequence of the imbalance in the number of customers belonging to each segment in the original customer data and the noise in the manually assigned labels. From this confusion matrix, we can calculate a macro-accuracy of 59% and a micro-accuracy of only 41% (a random assignment would result in a micro-accuracy of 33%, since there are three segments).

		predicted		
		A	**D**	**R**
actual	**A**	0	67	285
	D	0	530	866
	R	0	322	1,677

Table 1: Confusion matrix for initial rules.

After removing some of the noise in the original data (for example, customers in segment D and R with no transactions in the last year) and re-balancing the data set so that the algorithm sees the same number of customers per segment, we obtained a second set of rules using ProbE. The confusion matrix for these rules is shown in Table 2. In this case, the macro-accuracy is about the same as for the initial rules, but the micro-accuracy increased from 41% to 63%, because the new rules are making better predictions for the customers in segment A.

		predicted		
		A	**D**	**R**
actual	**A**	267	48	37
	D	161	531	230
	R	343	358	885

Table 2: Confusion matrix for improved rules.

We expect that with the system in place, the users will be more consistent in assigning customers to segments because they will have more access to data about the customers. Thus, the noise in the manually labeled data is likely to decrease and we will be able to extract rules that are even more accurate.

3.3. Extensions

In the initial implementation of our system, predictive data mining was used to obtain segmentation assignment rules that generalize existing manual segment assignment. These rules are then used to detect customers with changing behavior who should be moved to a different segment, or who are otherwise assigned to obsolete or inappropriate segments.

A next step is to apply predictive data mining on historical data containing sequences of segment assignments and movements, and obtain a model that predicts likely customer segment movement in the near future. This approach will be effective in providing the type of early warning alerts that are useful for genuinely pro-active marketing and sales actions. The pro-activeness of the current approach, to an extent, relies on the same in the manual segment assignment.

The latter approach requires customer data over a longer period of time, which will be available only after the system has been in operation for some time. As a general strategy, the BPM system can employ the former approach in its initial stage, and then upgrade it to the later approach, when the system has been in operation for some time and enough historical data has been collected.

4. BPM System and Architecture

The BPM approach involves a database infrastructure that is updated frequently so that the analytical models can learn, on a regular basis, whether the segment strategy is achieving its goals and where it falls short. The ETL infrastructure of our system makes efficient use of customer transactions to support segment definitions and marketing efforts.

4.1. Data Warehouse

All the data artifacts generated by the data extraction and monitoring processes are stored in the data warehouse. The design for the data warehouse is driven by the monitoring requirements such as performance metrics, dimensions, contexts, and fact tables that bring the metrics and dimensions together. The fact tables capture the information in its most granular form. The fact table design is based on a snow flake star schema model. Figure 1 shows a simplified Entity-Relationship diagram for financial metrics and customer segments.

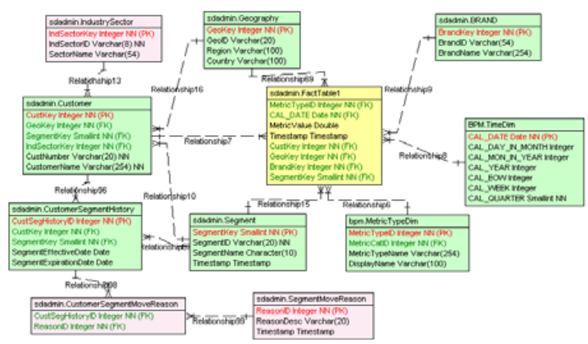

Figure 1: Simplified star schema for financial metrics and customer segments.

The fact tables are populated and updated overnight when access to the data warehouse is minimal. The enterprise dashboard displays rolled-up data that is stored in temporary tables in the data warehouse. Once the Monitoring Services complete the calculation of operational metrics, another set of queries refreshes the data in the temporary tables. The data warehouse also contains schema for storing alerts that are detected by the monitoring process.

4.2. Solution Architecture

The BPM enterprise architecture employed in this project consists of a modeler and a set of framework components that implement the modeled solution as illustrated in Figure 2.

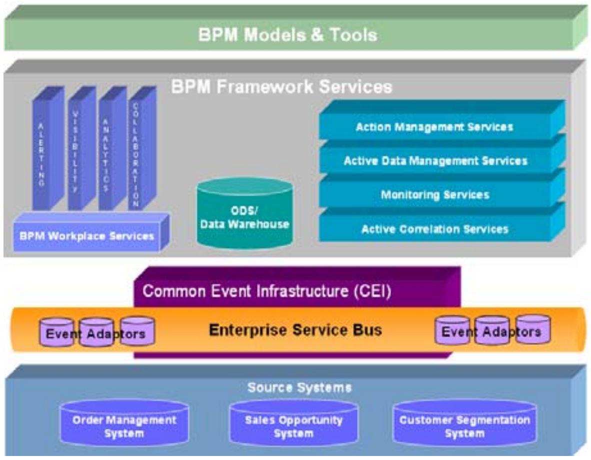

Figure 2: BPM Solution Architecture

The framework components provide capabilities to subscribe to incoming events (business events, IT events, etc.), correlate events, compute and store key performance indicators (KPI's), perform situation detections, and take actions to resolve business situations. The data artifacts generated during the monitoring process are stored in an Operational Data Store (ODS) and a data warehouse for historical analysis. The framework components and their usage in the BPM Telesales solutions are explained next.

The *Enterprise Service Bus* is a middleware layer that provides a set of infrastructure services to form the intersection between service requestors and the various service implementations. It embeds the event-driven adaptors that form data bridges. These adaptors extract all relevant transactional data elements (customer orders, sales opportunities, customer segmentation data, etc.) from the source systems and store them in the data warehouse.

The *BPM Workplace Services* provide business visualization and collaboration services for human users to participate in all BPM activities (model, deploy, monitor, analyze and act) to achieve continuous improvement. We used IBM's WebSphere Portal Server as the workplace solution [9]. We designed and implemented customized portlets to display metrics, alerts, and data mining results on the dashboard. The portlets also provide drill-down views of all metrics by sales region, brand, customer account, and segment. The screenshot in Figure 3 is a top-level view of the BPM Enterprise Workplace that illustrates the monitoring and alerting capabilities provided by the system (metrics values, customer and segment-related data were altered to protect confidential information).

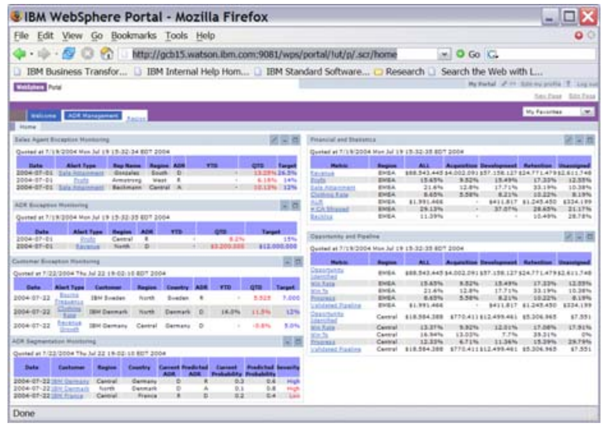

Figure 3: BPM Enterprise Workplace

The *Active Correlation Services* provide policy-driven filtering and correlation of events to detect situations of interest. Correlations are performed to filter raw transactional data from the staging tables for metrics calculation and situation detection. We used customized correlation services to filter the inputs for the data mining analysis.

The *Monitoring Services* provide capabilities to define and monitor metrics and generate alerts if a metric falls short of established goals. Alert conditions are identified during the analysis and design phase. The alert conditions are formulated as SQL queries that are executed against staging database tables. The threshold values for situation detection are gathered periodically, either via formatted csv files or direct user input via the dashboard.

The *Active Data Management Services* provide the framework components to access metrics and other business data relevant for BPM analysis and reporting. They also provide build-time components that read model data from the performance warehouse meta-model and generate the warehouse components.

The *Action Management Services* provide decision analysis and tracking of actions taken in response to BPM situations. In our system, the data mining module plays the role of decision analysis. It receives a trigger event from the Monitoring Services to reevaluate the classification of customers into segments. The current classification is displayed on the BPM Workplace, and stakeholders can review segmentation-based alerts and reclassify customers.

Finally, a set of *BPM Models and Tools* provide visual tools and metadata to capture the definition and expression of monitoring artifacts. We used the data warehouse meta-model to define the definitions and expressions of metrics and alert situations.

5. Summary

We developed a BPM software system that supports customer management and sales execution in a telesales environment. The application provides advanced analytic and infrastructure components that enable managers to optimize telesales center performance. It delivers visibility into telesales center operations and enables managers to easily measure key operational metrics and assess early-warning alerts on a personalized dashboard. We utilized BPM methods to design and implement the solution and established a BPM data model to host metrics, situations and situation targets. We applied data mining techniques to improve the segmentation rules and customer management over time. The architecture of the system can also accommodate other decision-support services like optimization, rule-based analysis, and business intelligence [6].

To further improve customer segmentation, the current modelling approaches can be extended around targeted marketing, customer lifetime value and churn management, to help telesales executives better manage all phases of the customer life cycle. New features might include determining characteristics of good customers based on expected long-term profit so that the organization can better focus their targeting; profiling customers who have bought a particular product so that the organization can focus attention on similar customers that have not bought that product (cross-selling); and profiling customers who have left so that the organization can act to retain customers who are at risk of leaving to reduce churn and attrition.

6. References

[1] C. Apte et al. “Segmentation-Based Modeling for Advanced Targeted Marketing”. In: *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 408-413, San Francisco, August 2001.

[2] K. Bhaskaran and M.-T. Schmidt. “WebSphere Business Integration: An Architectural Overview“. *IBM Systems Journal*, 43, 2, 238-254. 2004.

[3] S. Buckley et al. “Sense and Respond Business Performance Management”. In: H. Fromm and C. An (eds.) *Advances in Supply Chain Management.* Springer, 287-312. 2005.

[4] J. Dyche. *The CRM Handbook: A Business Guide to Customer Relationship Management.* Addison Wesley. 2001.

[5] S. Haeckel.. *Adaptive Enterprise: Creating and Leading Sense-and-Respond Organizations.* Harvard Business School Press, Cambridge. 1999.

[6] S. Kapoor et al. “A Technical Framework for Sense and Respond Business Management”. *IBM Systems Journal*, 44, 1, 5-24. 2005.

[7] G. Lawrie. “Preparing for Adaptive Supply Networks”. *Forrester Report*. 2003.

[8] G. Lin et al. “The Sense and Respond Enterprise”. *OR/MS Today*, 29, 2, 34-39. 2002.

[9] R. Will, S. Ramaswamy, and T. Schaeck. “WebSphere Portal: Unified User Access to Content, Applications and Services”. *IBM Systems Journal*, 43, 2, 420-429. 2004.