RC23002 (W0110-044) October 15, 2001
Computer Science

IBM Research Report

Lessons Learned in Designing an Expert Sales System

Shubir Kapoor, Steve Buckley, Kaan Katircioglu, Heng Cao,
Guillermo Gallego, Oyedele Oladeji
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Lessons Learned in Designing an Expert Sales System

Shubir Kapoor
Steve Buckley
Kaan Katircioglu
Heng Cao
Guillermo Gallego
Oyedele Oladeji

IBM T. J. Watson Research Center
Yorktown Heights, NY 10598

Abstract

Interest in expert sales systems is on the rise. Sellers could save significant time and travel expenses if they had an expert sales system with the knowledge of a cross-functional sales team. On the buy side, business-to-business purchasers want quicker, easier ways to find out how their business needs can be addressed by software, hardware, and service products. At IBM Research we have developed a web-based expert sales system called SEAS (Sales Enablement Across Solutions) that IBM uses in its daily sales operations. SEAS questions users about their business problems and translates those business problems into solution areas such as B2B and CRM. It also recommends specific solution vendors based on the user's industry, subindustry, revenue, and hardware platform. The purpose of this paper is to share our experiences in developing SEAS. We will relate the lessons that we have learned, illustrated by SEAS design information.

Keywords

Sales automation, electronic commerce, expert systems, return on investment

Introduction

Expert systems provide automated advice in many domains such as medicine, science, and government [2,4]. However expert sales systems are not yet in widespread use. Current use appears to be limited primarily to consumer commodities such as computers and shoes [5]. On the other hand, today's business problems are rarely solved by purchasing a consumer commodity; complex solutions are often necessary. A solution normally consists of a number of off-the-shelf and customized deliverables that in combination solve a specific customer problem. For example, a solution may combine an off-the-shelf supply chain or CRM software package with custom middleware (e.g. to bridge to existing IT infrastructure), off-the-shelf hardware, and custom services (e.g. process reengineering). Typically the sales process for such a solution is accomplished manually by a combination of general salespeople and expert consultants. Successful automation of this process by an expert sales system is considered by most people to be several years away.

Nevertheless interest in an expert sales system of this nature is growing. For example, many sellers want to reduce their SG&A costs by automating sales tasks that add little value and require the time of valuable sales people. Moreover, sellers could save significant time and travel expenses if their expert sales system had the knowledge of a cross-functional sales team. On the buy side, the Internet has created a new class of consumers that want quick access to product information. As a result, business-to-business purchasers want quicker, easier ways to find out how their business needs can be addressed by software, hardware, and service products.

At IBM Research we have developed a web-based expert sales system called SEAS (Sales Enablement Across Solutions) that IBM uses in its daily sales operations. SEAS questions users about their business problems and translates those business problems into solution areas such as B2B and CRM. It also recommends specific solution vendors based on the customer's industry, annual revenue, and other considerations. In reference [1] we described the motivation for such a system. In this paper we will share our experiences in developing SEAS. We will relate the lessons that we have learned, illustrated by SEAS design information.

Lessons Learned

1. **An expert sales system for business problems should speak in business language not in product language.**

Corporate executives control business spending but focus on financial measurements, not IT and process issues. Many business executives are confused by acronyms such as B2B and CRM. They do not want to hear that you have the best B2B or CRM system; they want to hear how your system will improve their financial measurements and provide significant return on investment. The concept of selling by focusing on financial issues is

touted by the Finlistics company [6]. Finlistics provides a tool called the *Finlistics Value Manager* [www.finlistics.com] that helps to identify weak financial areas based on competitive benchmark information. We play a related role - drilling down into IT and process issues, as illustrated in Figure 1. The SEAS sales process starts by identifying the financial areas that need improvement. The following financial areas are investigated: inventory turns, revenue growth, COGS (cost of goods sold), SG&A (sales, general, and administrative), asset turnover, accounts receivable, and accounts payable. After that the expert system drills down into IT and process issues within the user's industry that have the potential to improve the financial measurements.

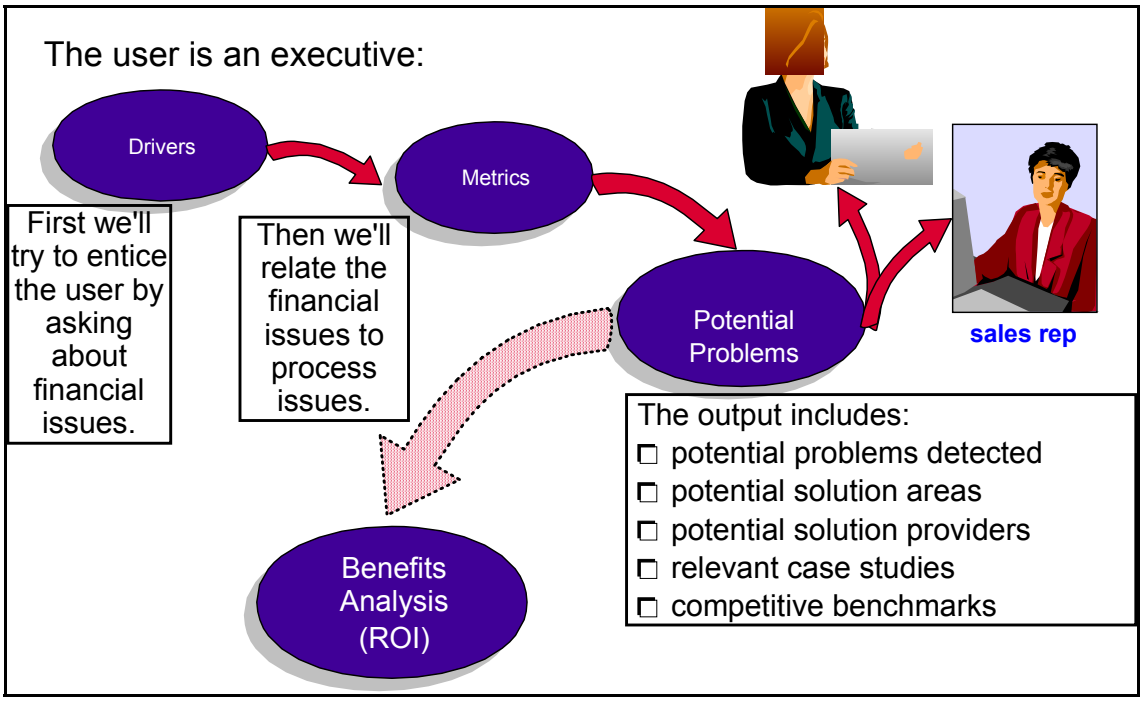

Figure 1. The SEAS sales process.

2. Industry standard metrics and best practices should be the basis for drilldowns into business needs.

In the process of developing SEAS we discovered that the best way to drill down into business needs is to ask qualitative questions about a company's performance using standard industrial performance metrics [3]. Most companies can answer these questions in a qualitative manner as long as the metrics used are widely accepted. In addition, questions such as "Are you utilizing e-marketplaces to sell off excess inventory" can be used to investigate the use of industry best practices.

In SEAS, with the assistance of IBM experts from several different industries, a list of metrics was compiled for each of the seven financial areas. A decision tree of metrics was constructed based on the observation that many metrics directly and indirectly depend on higher-level metrics. For example, the metrics Supplier Flexibility and Supplier Communication depend on the higher-level metric Supplier Performance. Similarly, the metrics Forecasting Tools and Inventory Count depend on Forecasting

Accuracy. Higher-level metrics enable the expert system to determine whether to drill down further or skip lower-level metrics. For example, if a company does not consider its Supplier Performance to be a problem, the expert system need not drill down and ask about Supplier Flexibility and Supplier Communication. As a result, only relevant questions are posed to the user. Another area where survey time is saved concerns metrics that appear in multiple sections. If the expert system detects that a particular metric has already been assigned a value, it does not repeat the question.

A sample set of metrics organized hierarchically in a decision tree is shown in Figures 2-5. Please note that the following metrics were not expanded in this decision tree due to lack of space: Revenue Growth, SG&A, Asset Turnover, COGS, Accounts Receivable, Accounts Payable, Work in Process, and Finished Goods.

Appendix I describe the design of the SEAS expert system that makes use of this structure.

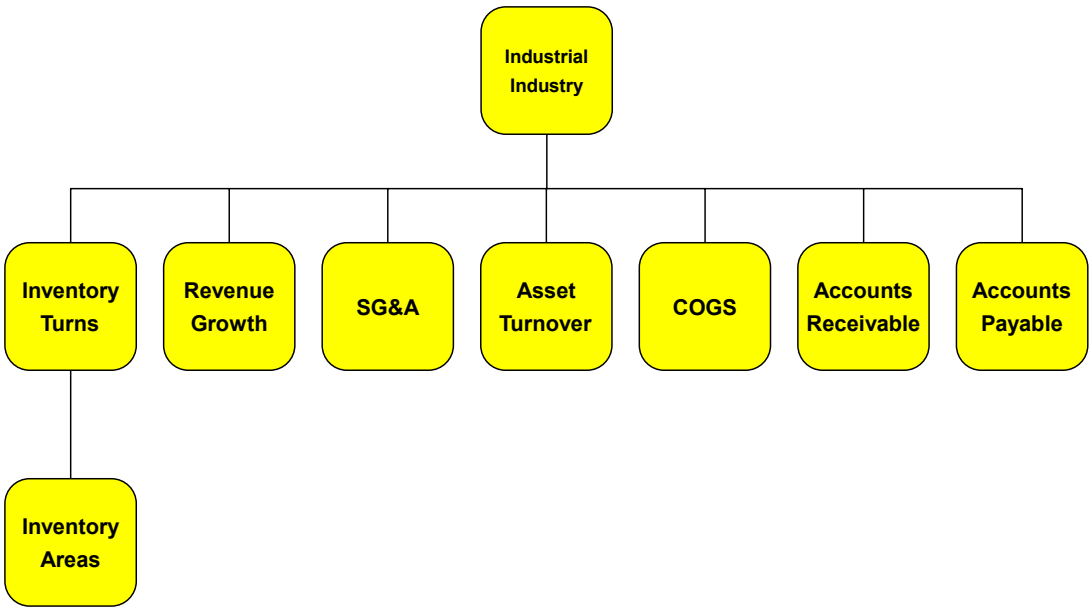

Figure 2. Top level metrics for the Industrial industry.

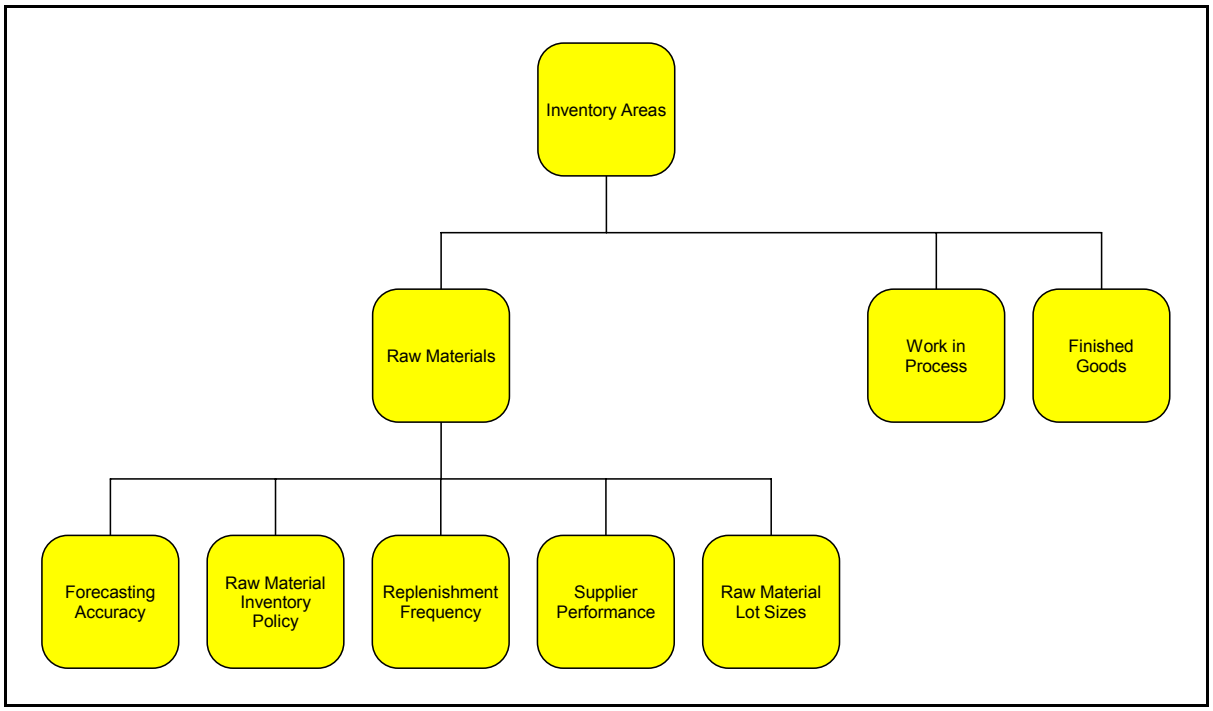

Figure 3. Lower level metrics for Inventory Areas.

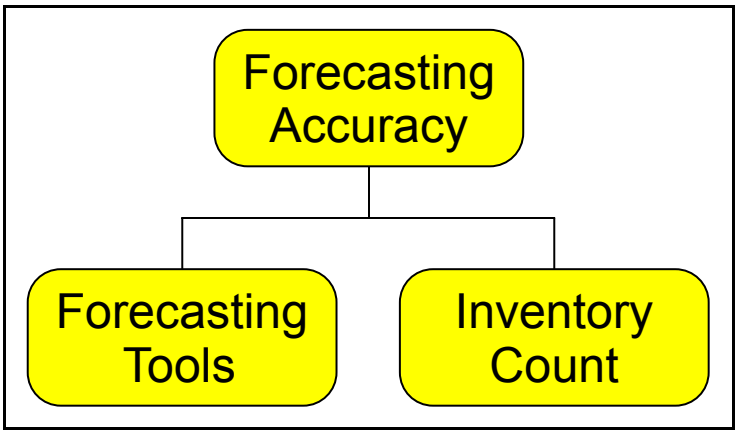

Figure 4. Lower level metrics for Forecasting Accuracy.

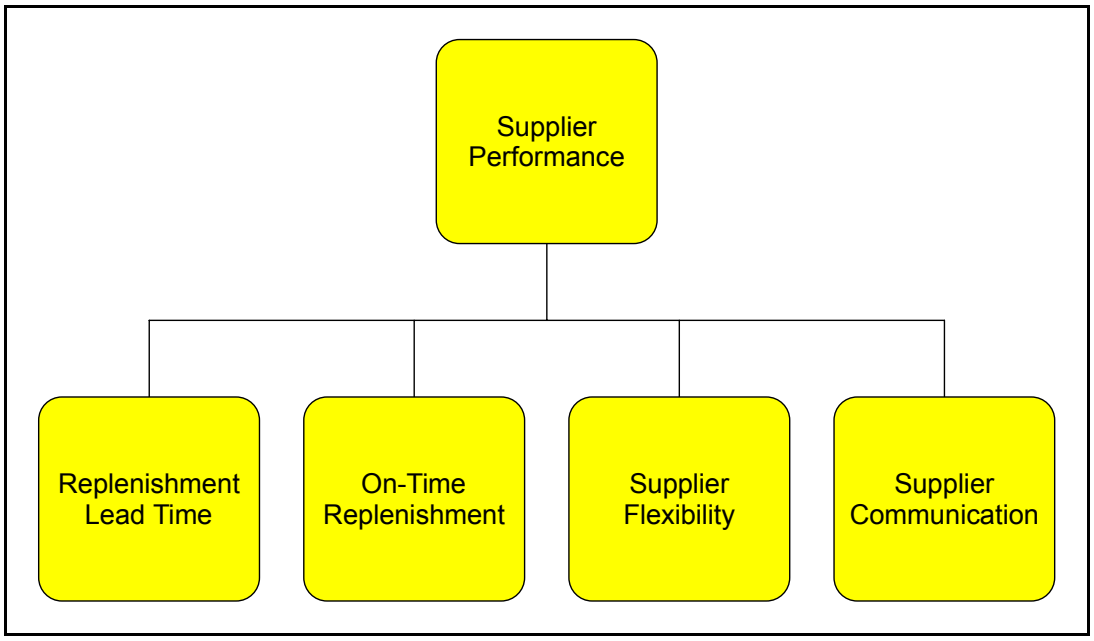

Figure 5. Lower level metrics for Supplier Performance.

3. User interest peaks when an expert sales system recommends actual products and services.

When we first began showing SEAS prototypes to audiences in 2000, SEAS did not recommend solution vendors, it merely identified potential business problems and referenced case studies where similar business problems had been solved. While this was considered by many to be a useful service, overall response to the prototype was not overwhelming. Later, in 2001, we added a capability to recommend solution areas and specific solution vendors. Response to the prototype dramatically improved. Apparently there was a distinct market need for a tool that could take users all the way from business issues to vendors who are capable of providing relevant solutions.

In SEAS, metrics such as Replenishment Lead Time that have no children are referred to as *leaf metrics*. Every leaf metric identified by SEAS as a potential problem results in execution of business rules that recommend solution areas. The recommended solution areas are then mapped onto specific solution vendors, with vendors ranked by the number of relevant solution areas they can address. At the end of the survey the user can view several reports:

- Business Assessment – List of financial drivers, business problems and potential solution areas. What if analysis and Gap analysis is also supported using key financial measurements of the clients company resulting in potential improvements in the cash flow.
- Summary Report – Lists all of the questions and answers in the current survey and references a relevant case study (see Figure 6).

- Solution Area Report – List of recommended solution areas sorted by the number of potential problems addressed (see Figure 7).
- Solution Finder - List of potential solution vendors and relevant offerings sorted by the number of relevant solution areas addressed. Parameters such as geography, platform and revenue help in narrowing down the solutions and offerings.

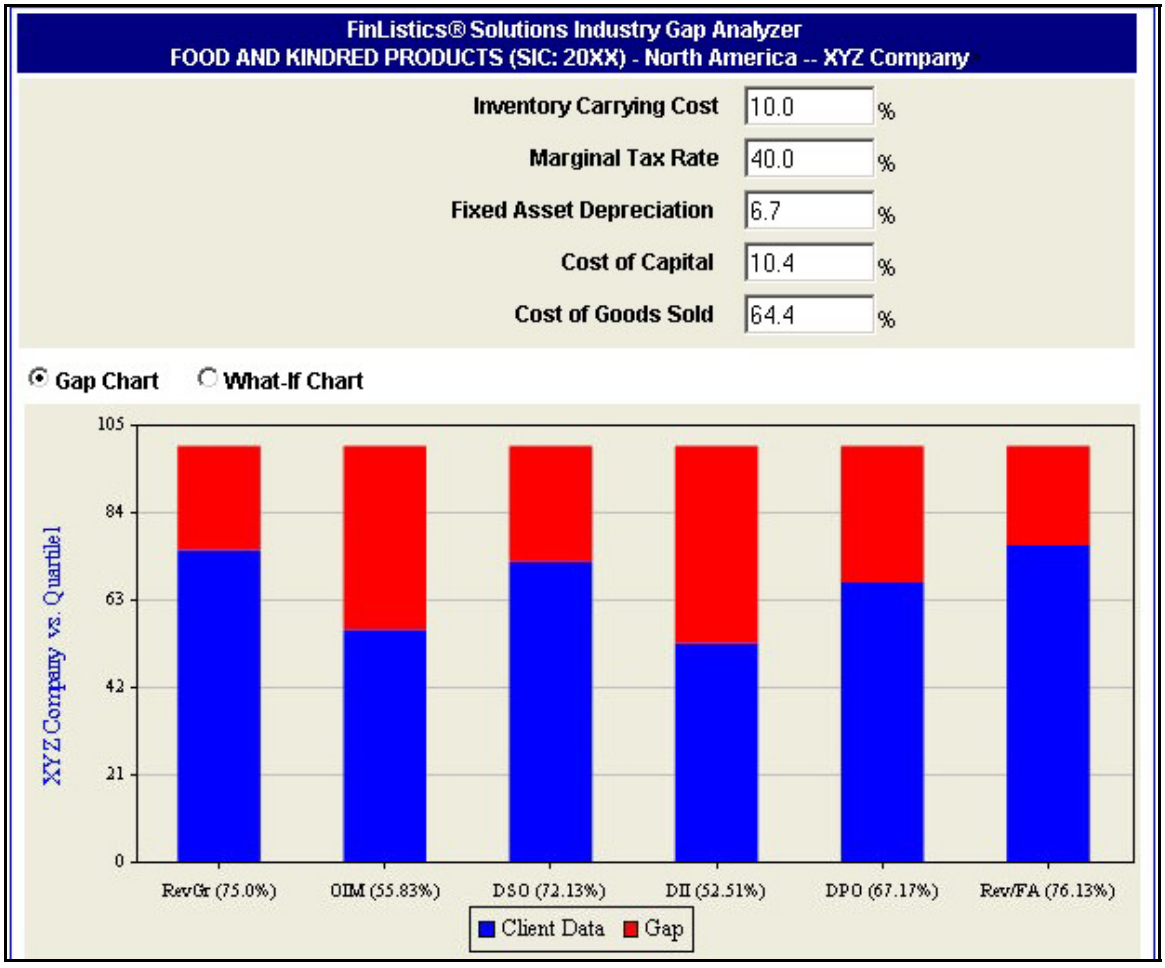

Figure 6. Gap Analysis.

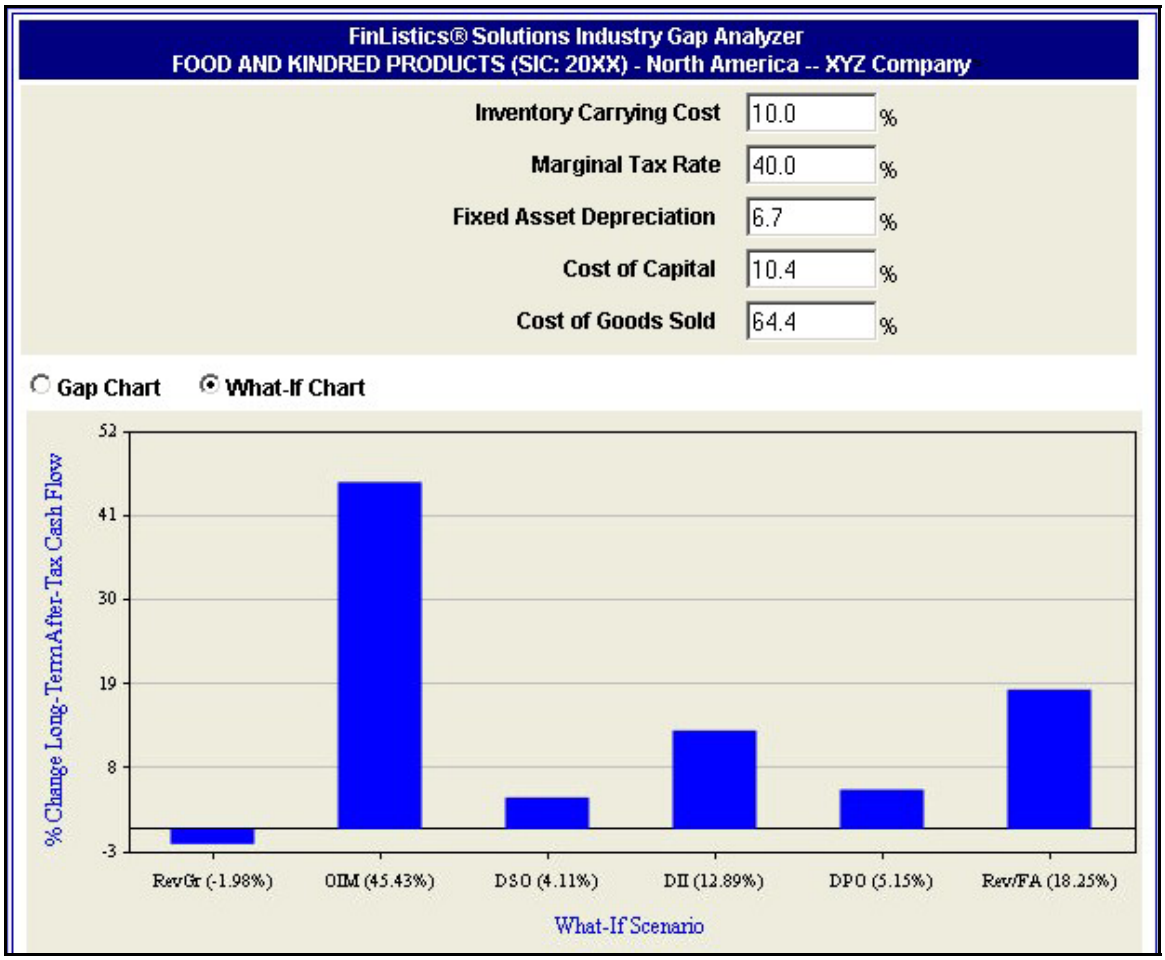

Figure 7. What – If Analysis.

10	What policies do you use to drive your raw material inventory levels?	;Vendor managed inventory;Just in time;Build to order;Continuous replenishment;Constrained MRP;Unconstrained MRP;	-
11	How often do you place orders to your suppliers?	Daily	-
12	How would you rate your supplier performance?	Room for improvement	-
13	Do you think that large lot sizes contribute to your overage of raw materials?	No	-
14	What tools do you use for forecasting?	;;Internal tools;Manual process;	Forecast error is a major cause of both stockouts and inventory buildup. By deploying advanced forecasting tools, you can reduce stockouts and improve your inventory turns significantly.There are many off-the-shelf forecasting tools available. One advanced company uses a business intelligence solution on their company intranet to help analyze customer purchase trends and forecast stock levels, allowing them to maintain a very high order fill rate
15	Do you maintain an accurate inventory count for each product and raw material?	No	In most cases, with more accurate knowledge of inventory levels, there are noticeable improvements in forecasting accuracy.
16	What is the typical lead time for raw materials?	>30	Since the lead times to replenish your raw materials are relatively high, you probably carry excess stock. In addition, you may be experiencing late deliveries or inflexible supply quantities, which might increase stockouts and have a negative effect on revenue growth
17	Do supplier shipments arrive on time?	<50	Late deliveries probably cause you to carry excess stock. In addition, they might increase stockouts and have a negative effect on revenue growth.
18	How flexible are your suppliers?	They are willing to handle emergency shipments;They accept late cancellations;They will store products for you in their warehouses;;;;	-
19	Is the majority of your communication	Yes	-

Figure 8. Summary Report.

Solution Areas-Problems

#	Solution Areas	Problem
1	Collaboration tools - (3 votes)	1. Forecasting Tools 2. Inventory Count 3. Replenishment Lead Time
2	B2B - (3 votes)	1. Forecasting Tools 2. Inventory Count 3. Replenishment Lead Time
3	Business Intelligence - (3 votes)	1. Forecasting Tools 2. Replenishment Lead Time 3. On-Time Replenishment
4	Network Management tools - (3 votes)	1. Forecasting Tools 2. Inventory Count 3. Replenishment Lead Time
5	Development tools - (3 votes)	1. Forecasting Tools 2. Inventory Count 3. Replenishment Lead Time
6	Forecasting, Demand Planning, Inventory Management - (2 votes)	1. Forecasting Tools 2. Inventory Count
7	m-Commerce - (2 votes)	1. Inventory Count 2. On-Time Replenishment
8	Logistics - (2 votes)	1. Replenishment Lead Time 2. On-Time Replenishment
9	Procurement & Replenishment - (2 votes)	1. Replenishment Lead Time 2. On-Time Replenishment
10	B2C - (1 votes)	1. Forecasting Tools
11	e-Procurement - (1 votes)	1. Replenishment Lead Time
12	e-Marketplaces - (1 votes)	1. Replenishment Lead Time

Figure 9. Solution Area Report.

4. Creating a self-service expert sales system for business problems may seem appealing, but who is going to pay for it?

Web pages that offer consumer commodities like computers and shoes are often self-service. The user logs in, selects a particular product configuration, enters a credit card number, and receives an order confirmation. However complex business problems usually require a custom solution that may involve negotiation with multiple suppliers. This is not a simple credit card transaction. Who is going to host such a web site? Ultimately, the answer may be an e-marketplace *market maker*. At present, SEAS is an IBM proprietary sales tool, not an e-marketplace tool, so we were not able to justify a self-service web site. As such SEAS is a sales aid for an IBM or IBM Business Partner salesperson. Our salespeople use SEAS to interview clients and make recommendations much like a cross-functional sales team would.

5. For business problems, an expert sales system does not replace a salesperson; it is simply a tool used by a salesperson.

Our salespeople use SEAS to generate solid leads, but there are still a lot of sales tasks to be accomplished manually – final vendor selection, configuration, customization, pricing, and contract negotiation. In the future we expect some of these tasks to be automated as well. However an expert sales system will never be able to supplant the interpersonal relationships that play such a crucial role in business-to-business sales.

6. Making an expert sales system for business problems available on the Web has some pitfalls, but there are many advantages in doing so.

One potential pitfall for a web-based sales tool is availability. If a salesperson schedules an interview at a client location, they would like to be assured that the web site will be available at that time and location, but in reality availability may be compromised due to heavy Internet traffic or faulty network connections. Another potential pitfall is that solutions and services are much more expensive and complicated than computers and shoes, hence business people may expect personalized sales attention and may not be comfortable with business advice dispensed by a machine.

Nevertheless there are many reasons why a web-based system makes sense, including:

- Global reach in multiple languages with little effort.
- Updates to the system only have to be applied once, on the web server.
- A web-based expert sales system can easily connect to other web sites to access up-to-the-minute product information, vendor information, case studies, market trends, financial benchmarks, etc.
- Leads can be easily and quickly sent to a lead management system.

All of these reasons relate to speed – the Internet clearly makes business and technological change move faster.

7. Detailed ROI analysis should be separated from the expert sales system.

Making the case for a solution sale ultimately involves some type of return on investment (ROI) estimate. However the main purpose of SEAS is to generate solid leads without spending a lot of time or money. Qualitative questions are used as much as possible to identify pain points and problems. On the other hand, detailed ROI analysis requires accurate quantitative metrics that take time for the client to collect. In addition clients may not be comfortable entering sensitive quantitative metrics over the web. As such, the ROI process is better left to a second phase after SEAS has identified the basic requirements.

We have developed an ROI tool called DIAT (Driver Impact Analysis Tool) to estimate the financial benefits of alternative solutions. DIAT is a spreadsheet-based tool that maintains a high level model of manufacturing and order fulfillment processes. It performs a quantitative analysis to show how various metrics impact customer order fulfillment. It attempts to give quick but meaningful results with minimal input data. We have been using this tool quite successfully for over a year to justify initiatives within

IBM. Using DIAT we identify the operational metrics that each proposed solution is expected to improve such as forecast error, manufacturing cycle time, and component availability. We estimate how much the solution is expected to improve these metrics based on historical benchmarks, vendor claims, and expert opinion. DIAT then calculates corresponding improvement potential in key measures such as inventory and customer serviceability through a high level mathematical model. It also provides sensitivity analysis using variability data provided for key metrics. With extensions, it could estimate savings and revenue growth due to improvement of the key measures. With such an extension an ROI calculation could be performed based on the predicted benefits and the estimated implementation cost.

Conclusion

Design of an expert sales system is a challenging task that requires careful analysis of user requirements and potential value. With these considerations in mind, clever designs are needed to implement these systems. This paper has provided important insights in each of these areas.

Acknowledgments

This work would not have been possible without advice and assistance from many people in IBM. We would like to offer special acknowledgments to Desiree Briceno, Tom Burnett, Jean Ellis, Doug Fiesel, Gino Gaenssler, Vembu Harish, Shekhar Mishra, Mary Orr, Ed Park, and Kathy Steiger.

References

[1] Buckley, S., Kapoor, S., Katircioglu, K., and Oladeji, O. (2001). Using an expert system on the web to recommend solutions and services. *2001 International Artificial Intelligence Conference*, Las Vegas, June 2001.

[2] Grove, R. (2000). Internet-based expert systems. *Expert Systems*, Vol. 17, No. 3, July 2000, pp. 129-135.

[3] Gunasekaran, A., Patel, C., Tirtiroglu, E. (2000). Performance measures and metrics in a supply chain environment. *International Journal of Operations and Production Management*, Vol. 21, No. 1-2, 2001, pp. 71-87.

[4] Huntington, D. (2000). Expert systems for online advice knowledge at your fingertips. *PC AI (USA)*, Vol. 14, No. 4, July-Aug. 2000, pp. 26-7, 49.

[5] Steinberg, D. (2001). The man with the e-commerce answer – technology supergenius Mark Lucente wants to give your web site a brain so it can have a

nice, profitable conversation with customers. *Ziff Davis Smart Business for the New Economy*, February 1, 2001, p. 100.

[6] Timme, S. The financial side of logistics. *IMRA Logistics 2001 Conference*, Orlando, Florida, 2001.

Appendix I. Design of the SEAS Expert System

SEAS is a web-based expert system built using the MVC (Model View Controller) design pattern. It makes use of Servlets, JSPs, XML and DB2. It is divided into three main areas:

- Presentation Layer – Responsible for displaying the metrics as questions that are easily understood by a business executive.

- Business Logic – Determines the next set of questions to ask based on previously answered questions. It eliminates those questions that have already been asked. When a leaf metric is reached, business rules are fired that link the metric to solution areas and rank each solution area based on the number of potential problems addressed.

- Persistence Layer – Stores metrics, questions, answers, etc. in a relational format so that simple queries result in fast data retrieval.

The business logic is based on the decision trees described earlier in this paper. Most of the business logic is stored as XML-based rules. There are three kinds of rules in the system: Flow Rules, Business Rules and Filtering Rules. Each rule has a prerequisite that must be satisfied before the rule can be executed.

- Flow Rules – These rules are organized with respect to the pages displayed in the presentation layer. Every presentation page has a set of questions associated with it. Each question corresponds to a metric. The prerequisite for each rule is to determine whether the associated metric has been answered by the user and is considered a potential problem. If the prerequisite is satisfied, the rule identifies the next presentation page.

- Business Rules – These rules fire whenever a leaf metric is identified as a potential problem. Business Rules are organized within libraries called *Rule Sets*. At the end of every presentation page a corresponding Rule Set is invoked; the rules being executed sequentially. Each rule in a Rule Set connects a single leaf metric to multiple solution areas. Since the same leaf metric could occur in multiple presentation pages, this organization allows rules to be shared between Rule Sets.

- Filter Rules – These rules fire at the end of the survey to reduce the recommended vendor list based on filters including the client's industry, subindustry, revenue, and hardware platform.

Each SEAS user must fill out a profile which includes the filters mentioned above. Based on the user profile an industry decision tree is selected and the user is presented with a series of questions. After each page of questions is answered, the Flow Rules determine the next set of questions. Since the rules are interpreted sequentially, we need to maintain the state of the rules execution. This is achieved by executing the rules in a separate thread from the Presentation Layer (see Figure 8). Whenever the Presentation Layer completes a page it passes control back to the Business Logic Layer and waits for the next set of questions. The Business Logic Layer then executes the rules of the next Rule Set, which determine the next set of questions. The Business Logic Layer then passes control to the Presentation Layer and waits until the Presentation Layer completes the page.

Interactions between the Presentation Layer and the Business Logic Layer are managed by a Rule Agent. The Rule Agent is responsible for parsing the rules sequentially and determining the next set of questions to ask the user. Every SEAS user is assigned a dedicated Rule Agent for the entire client session. Every Rule Agent runs in its own thread.

The system also maintains a pool of idle Rule Agents, managed by a Rule Agent Dispenser (see Figure 8). The Dispenser is responsible for initializing and compiling new agents and maintains a working set of agents ready to accept new clients. The Dispenser's agent queue has a defined minimum and maximum size. If the number of idle agents exceeds the maximum, the Dispenser deletes some of the idle agents. If the number of idle agents falls below the minimum, the Dispenser creates, compiles, and initializes new idle agents. This design has two main benefits:

- All of the Dispenser interactions occur in a separate thread running in the background. The user is not aware of this activity nor is the performance of the system affected in any way.

- Updates to the rule system by an administrator can occur in real time without restarting the web server. This is because the Dispenser reinitializes each Rule Agent after the associated client has completed their survey and the Rule Agent is released back into the pool. During initialization each Rule Agent reloads all rules.

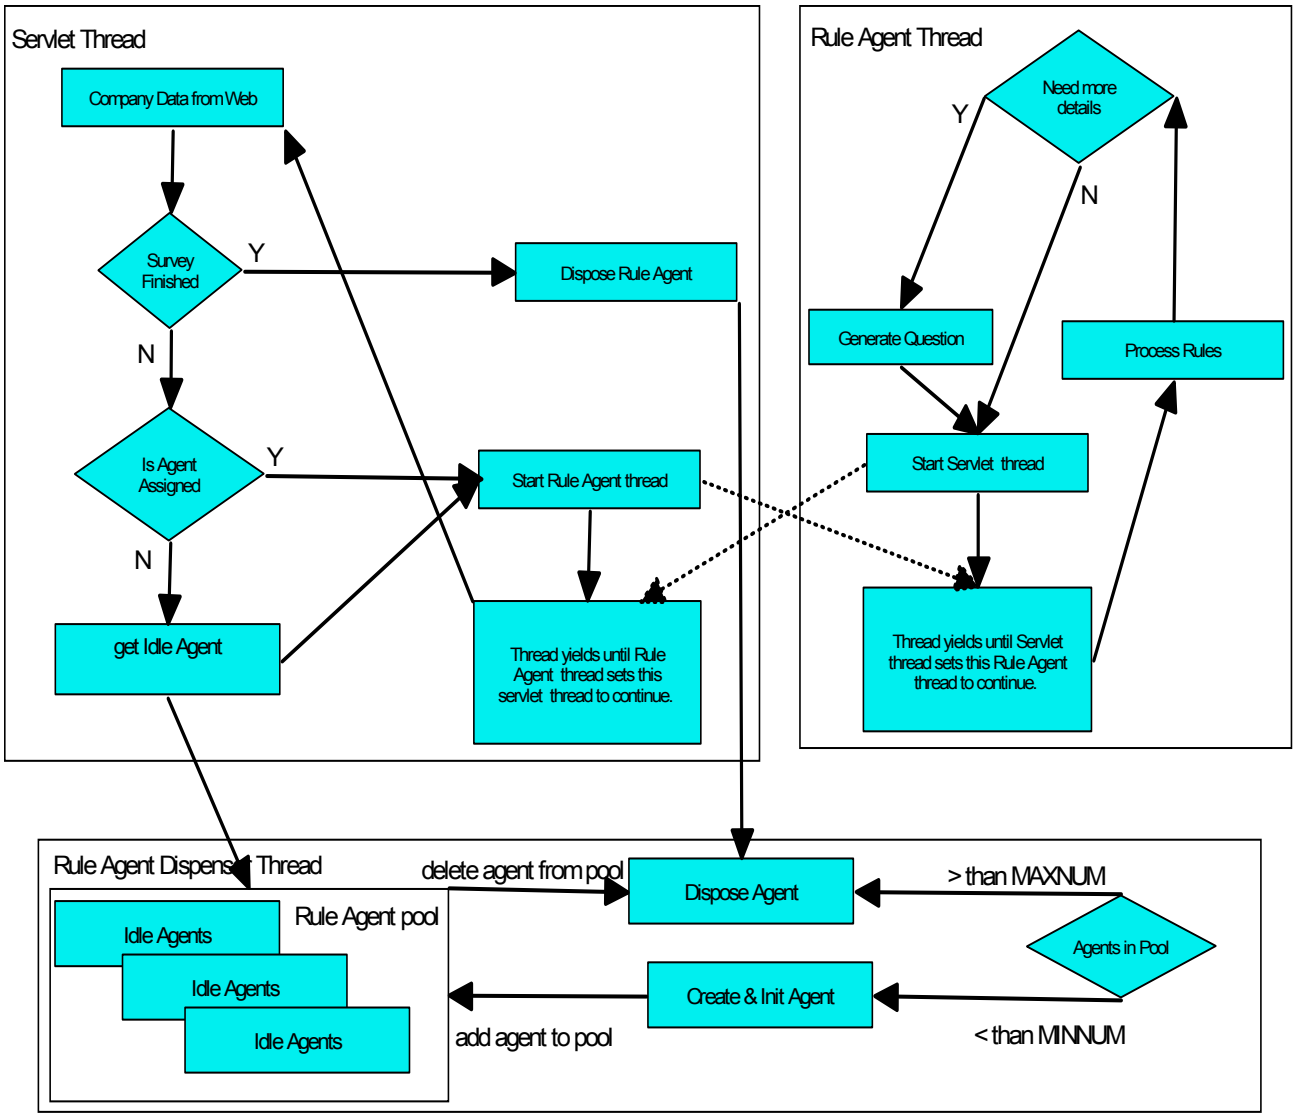

Figure 10. SEAS expert system thread interaction.