

AUGUST 2007



Dear Readers,

What's next in business intelligence? Find out in this issue, beginning with our lead article below on three emerging uses for automated BI.

You know that a BI Competency Center is key to effective BI at your organization – but have you met the members of a typical BICC? Now you can – just take our virtual tour.

In addition, learn how Virgin Money Australia achieved a 100 percent ROI on its BI investment, how operations research helps you come up with the best business plan and more.

Thanks for reading!

A handwritten signature in black ink that reads 'Anne-Lindsay Beall'.

Anne-Lindsay Beall
Editor, *SAS Business Report*

You Can't Do That with BI – Or Can You?

Three emerging uses for automated business intelligence

By Allan Russell

Traditionally, business intelligence has been used to improve corporate performance and alignment by producing metrics that – if attained – will lead to organizational success. One common use of BI has been to improve decision making in an organization by providing data that will influence decisions within a strategic, tactical or operational framework.

Traditional types of BI activities have focused on helping decision makers improve their understanding of how the organization works, with the expectation that decision makers can effect changes that would lead to better results.

But what if you want to improve business processes themselves and use the BI system to intelligently monitor, influence, adapt and change the operational systems without manual intervention? Is that possible?

With SAS, it is – and the benefits include reduced costs, improved decisions and efficiency gains.

This column will look at three emerging uses of BI that provide automated or embedded business intelligence to business processes themselves: feedback loops, real-time BI and business-activity monitoring.

The classic BI infrastructure

A typical BI infrastructure includes operational system elements, their associated databases and possible replications of their data, such as data warehouses or operational data stores. On top of the data warehouse, you will often see managed data marts, which are tailored to support specific analytic applications.

In classic BI, an analyst uses an interface to access data from reporting and analytics engines. Warehouse data flows into these engines, which prepare forecasts, format reports, build optimization models and complete many similar tasks.

The analyst is responsible for communicating any knowledge gained from the reports or analytics to the business process owners or information consumers, who presumably know how to improve their processes based on the information provided.

Provide feedback to your BI infrastructure

Feedback loops offer tighter integration between the BI infrastructure and the applications that implement business processes. In this model, some of the results that flow to the analyst from the reporting and analytics engine also flow automatically back into the operational systems. Results from forecasts, cost models and optimizations are integrated at the data level and returned as inputs to the application databases, or possibly back into the data warehouse or managed data marts.

This technique is most suitable for applications where there is a cyclical business process such as monthly, biweekly, or even daily planning or buying. Information consumers will benefit almost immediately because the application considers the new results each time it makes a new

recommendation.

For example, a global manufacturer of bacon, sausages, pies and pâté recently implemented feedback loops into its biweekly buying cycle for pig purchases. Typical factors in the buying cycle include the different types of breeds, the price on the market, the product demand and the processing methods used. The BI system considers the results from the previous buying cycle and determines how many pigs to buy and how to process the pigs.

After each buying cycle, the optimization analyses are rerun and fed back into the manufacturer's SAP R/3 system. After 12 months of using the system, the manufacturer continues to see improved savings on every purchase cycle.

To implement feedback models, the results of analyses are yielded by a batch-type process and should be inspected by controllers or data experts before being pushed back into the operational application. The main challenges to overcome include understanding the appropriate format for the data that the operational application will expect and making any necessary modifications to the operational application so it can understand and present the incoming data correctly. The good news is that many standard applications already have the necessary hooks and APIs to allow for this way of working.

Support on-the-spot decision making

How can BI assist with decision making on the fly? Can BI respond to an external event such as the swipe of a credit card or the arrival of a telephone call? When organizations set up BI systems to support on-the-spot decisions, they can help employees make faster, more informed decisions that improve the ability to cross-sell products, halt fraudulent activity and more.

In these scenarios, the business process or operational application will need to "call out" to the BI infrastructure for help. This call will trigger the scoring or recommendation engine to execute very quickly and return a result in the form of a recommendation or suggestion. For example, the result could be a go/nogo decision for a potentially fraudulent transaction, or it could be a "next offer" suggestion based on a customer rating.

In these cases, prescoring techniques are commonly used, but prebuilt models do not respond dynamically to changes in the customer situation because the scores typically get calculated some time before the previous push cycle.

The steps of classic BI are still necessary to build the real-time models based on historical data. The difference is that now the operational application calls out to get access to the results of the model, instead of calling up a file of stored results that may have been sitting in the reporting mechanism for several weeks or months.

A number of international banks have successfully implemented this type of system. One retail banking customer uses it to generate a response of either "yes," "no" or "refer" within milliseconds of each swipe of a credit card. Refer means customers are passed on to the case management system for further processing.

In this example, the retail bank constantly updates customer changes in the customer information control system and built a data mining model to detect fraudulent transactions. The model analyzes data in the operational system that has not made it into the data warehouse yet, and the system uses

message queuing technologies to communicate the decision back to the system instantly. The result for this bank was increased fraud detection and decreased loss.

Real-time analyses involve interapplication communication, so we recommend an open infrastructure. One way to achieve this is to implement a services call from the operational application to the BI infrastructure. Another alternative is to run the model “inside” the operational database. This latter option allows for more data to be pulled together in order to make the decision. If the model runs “inside” the database, it can quickly access all the data available for analyses.

Organizations must take care to keep track of the different models that exist for analyses. They need to monitor model performance as the real world drifts or shifts against the model’s representation. As soon as such changes are identified, we recommend creating a challenger model to see whether performance improves. Depending on whether the challenger beats the champion model, a switch may be needed – all these stages of a model’s life cycle will need to be tracked and managed, and model management tools are emerging to do that work.

Send alerts – and suggestions – to improve business processes

When the BI infrastructure looks out for abnormalities within a business process and then takes action to correct those abnormalities, this is called business-activity monitoring (BAM). In a typical BAM or alert situation, the business application either sends an event message to a decision maker when something abnormal occurs or the BI infrastructure is set up to automatically listen for internal message events from the business application.

In the second scenario, the BI infrastructure monitors events in real time and applies rules to determine when the process experiences an abnormality, such as a very high load of a particular transaction type or an unacceptably long delay in processing particular transactions.

A European post and parcels organization that wanted to optimize resource and route planning for parcel deliveries found success with BAM. Resource allocations needed to be determined in real time for each shift based on the activity volumes that day. The company’s system automatically sorts and records the movement of each package. The BI infrastructure listens to the message queue and fills the data mart with volume information that is always up to date. Business rules determine allocations just before each shift, and that information is communicated directly to the operational system as alerts, and automatically schedules shifts and delivery routes for optimum speed.

To achieve business-activity monitoring, some new elements will appear in the BI infrastructure, including the alerts engine and the business rules repository, which drives the alerts engine. The alerts engine pushes events to a decision maker who may be using a dashboard that refreshes key metrics regularly or who may be notified by e-mail when a threshold is exceeded. Assuming that the operational application can consume changes, the alerts engine can also push actions to the operational application to cause change in behaviors, such as changing routing parameters based on changed loads or causing particular classes of customer transactions to receive higher priority.

Managing the emerging BI infrastructure

Today, there is a shift in BI from looking exclusively at the data to understanding the role of BI in the business process. While a number of styles exist that allow for event-driven or periodic access to the BI infrastructure, the core of that infrastructure will remain the same.

As organizations begin to use new styles of BI, however, we must recognize that simply deploying one of the styles discussed above will not suffice. The key challenge will be repeatedly building and managing many systems with a mix of styles. Areas of management and governance to focus on include:

- Classic BI activities, such as building OLAP cubes or managing multiple ETL flows.
- Model lifecycle management as the number of models increases.
- BI services, which will become more important and will be managed in the same way as other services.
- Business-activity monitoring strategies, which should be developed to manage the many alerts generated from various rules and governing multiple interlinked operational systems.

When implemented correctly, embedded BI inserts business intelligence into the processes that affect the day-to-day operations of an organization. It does not replace traditional BI but helps BI become easier to use, more effective and more pervasive.

Optimize the Colors in Your Cube

Create better strategies with operations research

By Mary Crissey

Imagine a Rubik's Cube with each turning block representing a unique business unit of operations. As decision makers strive for a bright future, they devise action plans to accomplish specific goals and to gain the best advantage in the hopes of winning the game with a solid color on each side. Real-world problem solvers know they're not just manipulating six sides of the cube with six different colors – but must in fact optimize hundreds of color combinations and thousands of different variables.

Spend too much time worrying about the blue side of the cube (inventory, perhaps), and suddenly the red side (the shipping schedule) is messed up. Turn to find a solution that puts all the red pieces in place, and the single row of yellow (purchasing) that used to be aligned is now three different colors.

Operations research (OR) makes it easier to solve the Rubik's Cube of your business operations. In particular, the latest release of SAS/OR[®] software extends predictive power with optimization features so decision makers can analytically devise the best operational game plan. The updated software from SAS simplifies and streamlines the process so that organizations can implement proactive planning quickly and efficiently.

The history of operations research

Operations research is not a new concept. During World War II, the United States and Britain called on mathematicians to study ways to improve military operations. With a scientific approach, these OR pioneers simulated battle plan scenarios, identified troop requirements, planned logistical distributions and pinpointed the best actions to weaken the German army.

With so many different types of troops, military units and potential battlefields, you can imagine the immense number of possibilities in military planning. Operations research eliminates the need to test each and every possible combination by narrowing the choice of alternatives down to just the ones that best allocate troops and equipment aligned with military objectives.

After the war, some companies began to consider preplanning strategies; however, the mathematical, statistical and algorithmic functions required to examine and cross-reference operational processes were complex and difficult to understand and master. In addition, data collection was a tedious manual process, and because computer memory was both limited and costly, most organizations could not realize the benefits of operations research.

As organizations entered the information age and began to computerize operations, the focus shifted from the accumulation of data to business intelligence (BI). BI provided insight into operations with solid ways to measure progress by showing performance trends based on comparisons to history. Most companies now have gathered tremendous amounts of data, and leveraging that wealth of information has become a priority.

Putting company information to work

Operations research helps organizations rise above the reactive, busy mode of coping with daily challenges and guides them with proactive steps to create better action plans for success. OR allows organizations to leverage BI, get a glimpse of future scenarios and devise a “best” plan for operating business affairs so they can achieve enhanced growth and productivity.

Above all, optimization recommends decisive actions that collectively fulfill identified organizational requirements to ensure the most productive use of company resources. Resources may include people, materials, silos, plants, tools, machines or time. SAS/OR provides a scientific, logical and guided approach for organizations to follow as they consider strategies for optimizing business practices.

Today, operations research is being used to streamline and enhance operations across industries in a variety of concrete ways, including:

- Resource allocation and management.
- Production and inventory planning.
- Product mix and blending.
- Staffing allocations.
- Distribution, routing, scheduling and traffic flow.
- Supply chain management and optimization.
- Capital budgeting, asset allocation and portfolio selection.

Reducing complexity

Experienced modelers must translate business scenarios into mathematical representations before any software can perform the sophisticated programming techniques involved in operations research. Once the real-world problem is expressed in terms of constraints, goals and objectives, a wide range of decisions – including resource allocation, or product mix and blending problems – can be addressed and optimized.

With SAS/OR, the complex mathematical programming is packaged and complies with industry standards so users don't have to fuss with obtuse inputs or learn fancy lingo. With the SAS platform, decision makers can launch SAS/OR directly from their favorite spreadsheets. Those who use SAS Enterprise Miner™, SAS/STAT® or SAS Forecast Server can integrate their analyses with the newly enhanced optimization product.

SAS/OR provides a single modeling language for many different types of problems. The software is smart enough to detect which type of solver, or computational algorithm, is appropriate for which business problem. As a result, users don't need to worry about declaring which mathematical technique to use (or know which niche vendor's software tool to buy). Often OR analysts are assigned to ad hoc "hot issues" teams, and with SAS, they can be confident that they are prepared to solve a wide range of optimization challenges.

Other enhancements make it easier to validate models, review them and make subsequent adjustments to better represent the real-world business cases and later rerun them with new data for similar situations.

In short, the analytics for operations research are complex. There's no way around that fact. And that's why it is important to select an OR application that is user-friendly and compatible with various BI sources.

With SAS/OR, you'll have a solid base for your analytical analyses as well as the flexibility to adapt easily to changing needs, and you'll have the optimization power to add insight, innovation and credibility to the decision-making process.

Cutting Through the Clutter

Virgin Money launched in Australia three years ago to shake up financial services. With award-winning products spanning credit cards, home loans and superannuation, Virgin Money is already achieving what the Virgin brand is renowned for worldwide – challenging traditional business models and giving consumers a better deal. For the local operations of this high-profile brand, SAS has been introduced to assist the company in maintaining the unique edge that is Virgin.

The SAS software solution

In its early years, Virgin Money Australia was a subsidiary of the more established Virgin Money U.K.. Following the rapid growth of the Australian business, the “cord” was cut in early 2005 and Virgin Money Australia became fully independent.

For the company’s Customer Insights Manager, Carlie Smee, operational independence presented the opportunity to introduce a new analytics and business intelligence solution that would better meet the local operation’s needs. “We assessed solutions from three vendors,” she says, “evaluating their ease and effectiveness in predictive modelling, data mining, reporting and data warehouse connectivity.

“The results were that SAS was the only vendor able to provide us with a total solution that met our business needs and budget. If we had opted for another solution, it would have meant introducing a multi-vendor system; and the problems associated with that in the areas of training and support simply didn’t warrant their consideration. A SAS Enterprise Intelligence Platform solution, including data mining and business intelligence, was the obvious choice.”

Proactive customer retention

“One of the core strengths of the entire range of Virgin businesses is the absolute focus on customers; and the requirement for this is essential,” Smee states. “To achieve this in Virgin Money, we need to be an analytically driven business, with data being utilised to its full potential.”

A key strategy that is being progressively introduced throughout Virgin Money and based on the SAS software solution is proactive customer retention. Smee explains: “Where it used to be a case of identifying non-profitable customers by reporting on metrics such as dormancy, the SAS solution enables us to build behavioural models that predict those customers who are likely to become dormant or churn in the near future. Essentially, we’ll be in the position of being able to identify customers who we see as being at risk of churning.”

Lowering acquisition costs

“One of the common problems facing financial services companies, particularly with credit card products, is spending resources to attract a new customer only to find out after that effort that they get knocked back based on the credit scorecard,” Smee says. “The obvious challenge – and one that we’re now squarely addressing with the SAS solution – is building predictive models to score lists based on those potential customers who are more likely to have their credit approved.

“The result is that we are able to attract a greater proportion of approved customers in relation to the number of people we contact, effectively lowering our customer acquisition costs.

It is that same approach that Virgin Money is planning to introduce with cross-sell communications to its existing customer base. Rather than utilise the typical – and costly – “scattergun” approach to customer communications, the company’s marketing team will use SAS analytics and reporting functionality to develop individual and more effective communications for reaching precise customer segments.

Easing the burden

On the issue of data warehouse integration, Smee comments that the introduction of the SAS solution has brought about a “considerable time saving” for her team of customer insight analysts.

“Using SAS software, we created a reporting portal which interfaces to our data warehouse, which provides everyone in the business access to accurate real-time reporting. Where building a single report may have previously taken one of my team members anywhere up to a full day, it’s now being done by staff in the business units in a matter of minutes; and that goes for everyone from our CEO right down the line. This allows the customer insights analysts to spend more time analysing results rather than just building reports.”

Only months for 100 percent ROI

Prior to the implementation of the SAS Enterprise Intelligence Platform, Smee and her team prepared a business case that projected a 100 percent return on investment (ROI) of barely three to four months. Smee explains: “Previously, analytic projects that required model building needed to be outsourced to a consultant, which typically cost us anywhere between AUD\$50,000 and \$100,000 dollars.

“Given we undertake around 10 such projects each year, we projected a full payback on our initial investment over the course of only two to three projects. That,” she states, “is a level of ROI that’s rarely realised, let alone even hoped for.”

The Who's Who of the BICC

What kinds of people does a Business Intelligence Competency Center (BICC) need? And what skills, knowledge and qualities should those people possess? Meet the members of a typical BICC when you take this virtual tour.

Online only at:

<http://www.bettermanagement.com/keycode.aspx?keycode=538116>

What Does It Take to Successfully Implement Performance Management?

Gary Cokins collaborated with Bob Paladino, author of the book *Five Key Principles of CPM*, to tackle this question in his most recent blog post. Check it out and leave a comment!

Online only at:

<http://blogs.sas.com/cokins/index.php/?archives/18-What-Does-It-Take-to-Successfully-Implement-Performance-Management.html>

SAS Business Knowledge Series Delivers Training for the Real World

Developed and delivered by top thought leaders from industry, SAS' Business Knowledge Series courses provide in-depth information and case studies that give professionals the tools they need to immediately benefit their organizations.

Online only at:

<http://support.sas.com/training/bks/us/>

Events:

Gartner Customer Relationship Management Summit

Sept. 17-18

Hollywood, FL

<https://www.gartner.com/it/page.jsp?id=498193&tab=overview>

IBM Business Intelligence Virtual Symposium

Sept. 18, 9:00 a.m. ET

<http://events.unisfair.com/index.jsp?eid=205&seid=28&code=V13>

Webcast: Worst Practices in Forecasting

Sept. 27, 1:00 p.m. ET

<http://www.sas.com/events/cm/147569/index.html>

M2007: Data Mining Conference

Oct. 1-2, Las Vegas

<http://www.sas.com/events/dmconf/index.html>

The Premier Business Leadership Series

Oct. 3-4, Las Vegas

<http://www.bettermanagement.com/theseries/worldwide.aspx?L=14481>