.BUSINESS PROCESS TRENDS
WHITE PAPER
February 2004

Joseph M. DeFee
Senior Vice President
Advanced Systems Division
CACI

Paul Harmon
Executive Editor
Business Process Trends

Contents:

www.bptrends.com

Business Activity Monitoring and Simulation

1. Managing a Business in Real Time

Companies have always depended on processes. Historical processes many not have been as well-analyzed as they are today, but there have always been business procedures designed to turn inputs into outputs in an efficient manner. Just as there have been processes that defined how materials flowed from their arrival to assembly and then to shipping, there have always been communication and control systems that attempted to monitor the process flows and deal with events that threatened to upset the expected flow.

Consider one example of how processes have historically been managed. Imagine a small hospital of 30 years ago. As today, this hospital had a Customer Lifecycle Process that managed customers from admission through treatment to discharge. (See Figure 1.) The first subprocess or activity was probably Admissions. As patients came through the front door they were documented. The patient's medical history was determined, credit was established, and the patient was assigned to a specific ward for treatment. Over the course of time, the hospital had established expectations. In a normal week, roughly the same number of patients entered as were discharged, maintaining a predictable need for doctors, nurses, medicines, and beds.

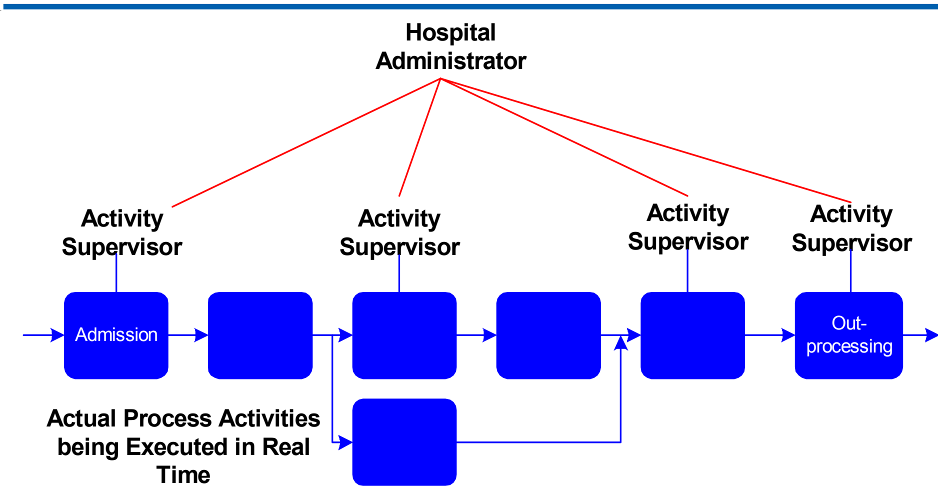

Figure 1. A manual patient lifecycle process with reports delivered by phone.

(Blue lines indicate flow of patients. Red lines indicate flow of communication and control information.)

Consider what happened when a serious local infection began to manifest itself, or when a major fire or traffic accident resulted in a large number of patients all arriving at once. When a sudden, unexpectedly large number of patients arrived in the lobby, the supervisor in charge of admissions picked up the phone and called the administrator to alert her that there was a problem. The administrator would normally ask the nature of the emergency, and then consider possible actions. If the emergency was an accident, it was easier to deal with, in the sense that a few phone calls could probably determine the extent of the accident and number of patients that would be arriving at the hospital in the course of the next hour or two. Calls to other supervisors would result in still other calls to change shifts and make more doctors and nurses available. Still other phone calls would result in shifts of bed assignments to assure that the emergency trauma ward had enough beds available. In the course of an hour or so, the hospital would adjust its activities and assign new resources to assure that the patient lifecycle process would continue to function effectively.

A harder problem would be an increase in flu patients. In this case, instead of getting a large, but known, increase in admissions over the course of a few hours, the hospital would need to deal with a number of unknowns. Admissions would begin to increase slowly, and then, as the flu spread, the increase would grow from day to day. Many variables would determine the overall course of the flu. Severity, the susceptibility of particular groups--like older people or youngsters, whether school was in session, the availability of flu shots, and many other things could limit the spread, or control the duration of hospital flu patients. Complicating matters, the flu might infect doctors and nurses, making it harder to smoothly adjust staffing schedules. Although most flu epidemics pass without serious consequences, there have been especially virulent epidemics, like the one following World War I, which killed millions of people. The alert hospital administrator has to try to plan for a variety of different scenarios, and then adjust her actions as she acquired more data on the development of the flu in the hospital's community, and in the nation as a whole.

Different industries have different kinds of problems. Most, however, have processes that are designed to run within set parameters, and those processes have communication and control systems in place to handle exceptional periods. Most exceptions are easy to understand and deal with, while some are much more challenging, involving, as they do, more complex interactions among variables over a longer period of time.

Historically, companies have relied on smart, experienced managers to gather appropriate data, interpret it correctly, and take decisions to minimize the effect of the changed circumstances on the daily functioning of company processes. As companies have become larger and processes have been disbursed over wider geographical areas, managing large business processes has become more difficult.

In the past thirty years, most companies have installed computers and used them to collect data, and, in some cases, to automate processes. Thus, for example, our hospital admission office now enters new patient data via a computer terminal and can often access data from customer databases to determine a new patient's medical history and credit. Since most supervisors have access to computers, it is often possible for ward supervisors to check the admissions database to determine how many new patients will be arriving in their ward in the next hour. Similarly, it is possible for an administrator to check historical data and generate a report that describes how many patients were admitted during the flu season last year or during the last ten flu epidemics.

In essence, computers, that were originally installed to facilitate or automate the flow of patients, parts, or assemblies through a process, can also be used to facilitate monitoring and communication, and some can even support managers who have to make decisions to maintain the efficiency of a process in unusual circumstances.

During the last few decades, most companies have also become more sophisticated in their management of processes. To counteract a tendency toward departmental functions that don't communicate as efficiently as they might, most companies have designated managers who are responsible for large-scale business processes. In product-oriented companies, these managers are often termed line-managers. In other cases, managers are assigned to coordinate processes that cross functional lines, like our patient lifecycle process. To support these managers, who are often responsible for managing processes that occur at several different locations and over long periods of time, software vendors are working to create tools that pull together all of the relevant information, highlight problems, anticipate problems, and assist making decisions to assure rapid successful adjustments of the process flow.

2. Business Activity Monitoring

In 2002, the Gartner Group coined the term Business Activity Monitoring (BAM) to refer to software products that aimed at "providing real-time access to critical business performance indicators to improve the speed and effectiveness of business operations."[1] In the past year the term *BAM* has become quite popular.

Before using the term, however, it's important to emphasize that BPM is a misleading term. The emphasis should have been on Business PROCESS Monitoring. Unfortunately, Gartner already used the acronym BPM to refer to Business Process MANAGEMENT, so Gartner apparently used "Activity" to get a unique, new acronym. BAM has caught on, and we'll use it throughout this paper, but readers should remember that the emphasis in BAM is on pulling together information about large-scale processes, rather than on monitoring small-scale activities.

BAM and Other Decision Support Technologies

As we have already suggested, using computers to help smooth the flow of items through a process is nothing new. For many years, software designers have built triggers and alerts into software applications. Thus, for example, if admissions exceed some set number, an administrative terminal may sound an alert. One only has to think of an operator at a power plant to understand how a computer can provide an operator with a wide variety of alerts and even provide diagnostic information to assist the operator in his or her job performance.

Similarly, marketing groups have analyzed data from sales for decades to determine shifts in customer preferences. In the past decade, many companies have invested in large Data Warehouses, that consolidate data from many smaller databases, and Business Intelligence (BI) applications that use special algorithms to search massive amounts of data, looking for patterns that humans might overlook. The results of these efforts usually find their way to senior managers who set strategy or design new products.

IT groups have also used ERP and EAI tools to analyze the flow of data between application components. The emphasis, in the case of IT, has been on fast, efficienct data processing and smoothly functioning middleware and not on drawing any broader meaning from the data. Still, it's easy to imagine how transaction data, relabeled, and provided within the broader context of a business process model, could help business process managers understand how a process is working.

BAM proposes something that falls in-between the immediate feedback that altert signals and triggers can provide operators and supervisors and the long term trend reports that database reports and BI can provide senior managers. (See Figure 2.) BAM aims at providing a process manager with a broad overview of a major business process. As in the case of our example, it seems to provide a hospital administrator with an overview of the current status of the patient lifecycle process. Or, it seeks to provide a factory administrator with an overview of how an entire production line is functioning.

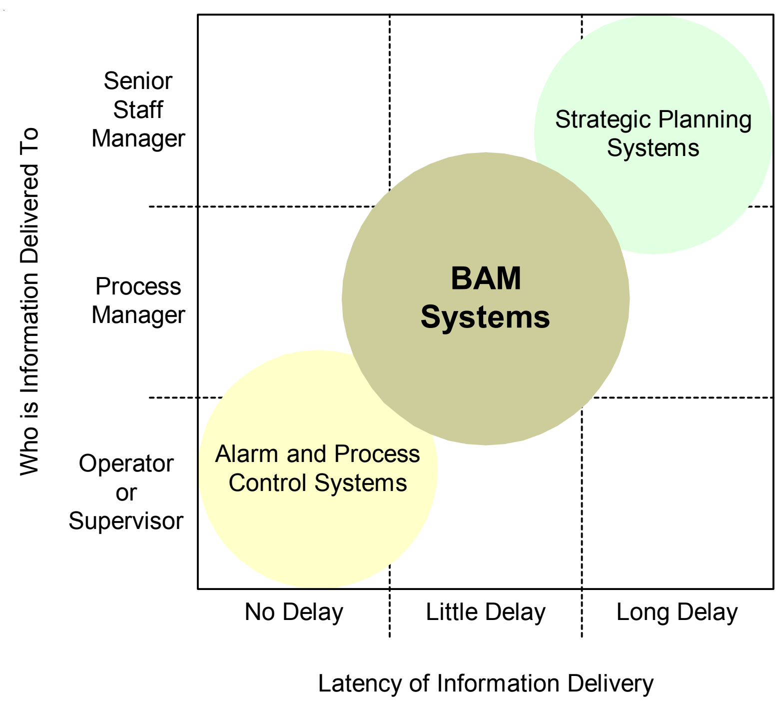

Figure 2. One way of organizing the monitoring and decision support systems in use today.

Figure 2 provides one way of summarizing the range of monitoring and decision support systems in use today. Process control systems that provide information to operators and provide alerts to supervisors are mostly real-time systems that report to employees and supervisors who are very close to a specific activity. Similarly, systems that gather data, analyze it over hours, days, or weeks and report to senior staff managers are mostly designed to aid in future planning. BAM systems are newer and less widely deployed. They aim to fill the middle-ground between activity specific and strategic planning systems by providing business process managers with near-real-time information about an entire process. Properly done, they allow the process manager to initiate changes in specific activities that keep the entire process running smoothly.

The Functions Required for an Effective BAM System

A BAM system can not simply provide the administrator with the kinds of raw data or the signals that it provides plant operators or IT managers, or the administrator would be overwhelmed with inputs. Instead, someone must design a filtering system that draws data from a wide range of sources and then massages it so that only truly significant data reaches the administrator. On the other hand, the BAM system can't spend too long in massaging the data, or it will be out of date, like the BI systems that provide trend data to strategists, and only useful for future planning. A good BAM system should provide the administrator with enough information to enable good decisions, and it should provide the information in something close to real time so that decisions, when needed, can be taken in time to actually affect the ongoing performance of the process flow.

Any effective BAM system requires a collection of modules. (See Figure 3.) There are different ways the various modules can perform their functions, but all of the functions must be present if the BAM system is to perform as its strongest advocates suggest it will. The first set of modules must convert data about actual events into digital information. In most cases this can be done by simply monitoring databases and transaction events that occur as software is used to automate a process. Thus, the same data the administration clerk enters into the computer, as he signs in a new patient, can feed a monitoring system that keeps track of the rate of patients entering the hospital.

The second level or set of modules required for BAM must provide some context for the digital data being accumulated. The BAM system may depend on an explicit model of the actual business process, as illustrated in Figure 3, or it may simply depend on a series of equations that establish relationships between data sources. One way or another, however, the system must be able to organize the data to reflect the process it is monitoring. The analysis of the relevance of the data, the generation of information about trends, and intelligent action suggestions all depend on an analysis of the process and the relationships between process elements.

iness
lution

Using its understanding of the process, the BAM system must apply some kind of logic to the data to identify problems, diagnose them, and to recommend managerial actions. For example, the BAM system might apply a set of business rules. One rule might state that whenever patient admissions increase by more 10% of the expect rate for a given period, a signal should be set. Another rule might state that whenever a signal is set as a result of a 10% increase, the patients' ward assignments should be analyzed to determine if the increase in patients was random, or if there was a significant increase of patients for a particular ward. Still another rule might say that whenever patients were assigned to a given ward in excess of some historical number, the rule should post a suggestion on the administrator's BAM monitor that specific changes are to be made in the staffing of the ward. Rule-based systems can be used to accomplish a number of different tasks. For our purposes here, they

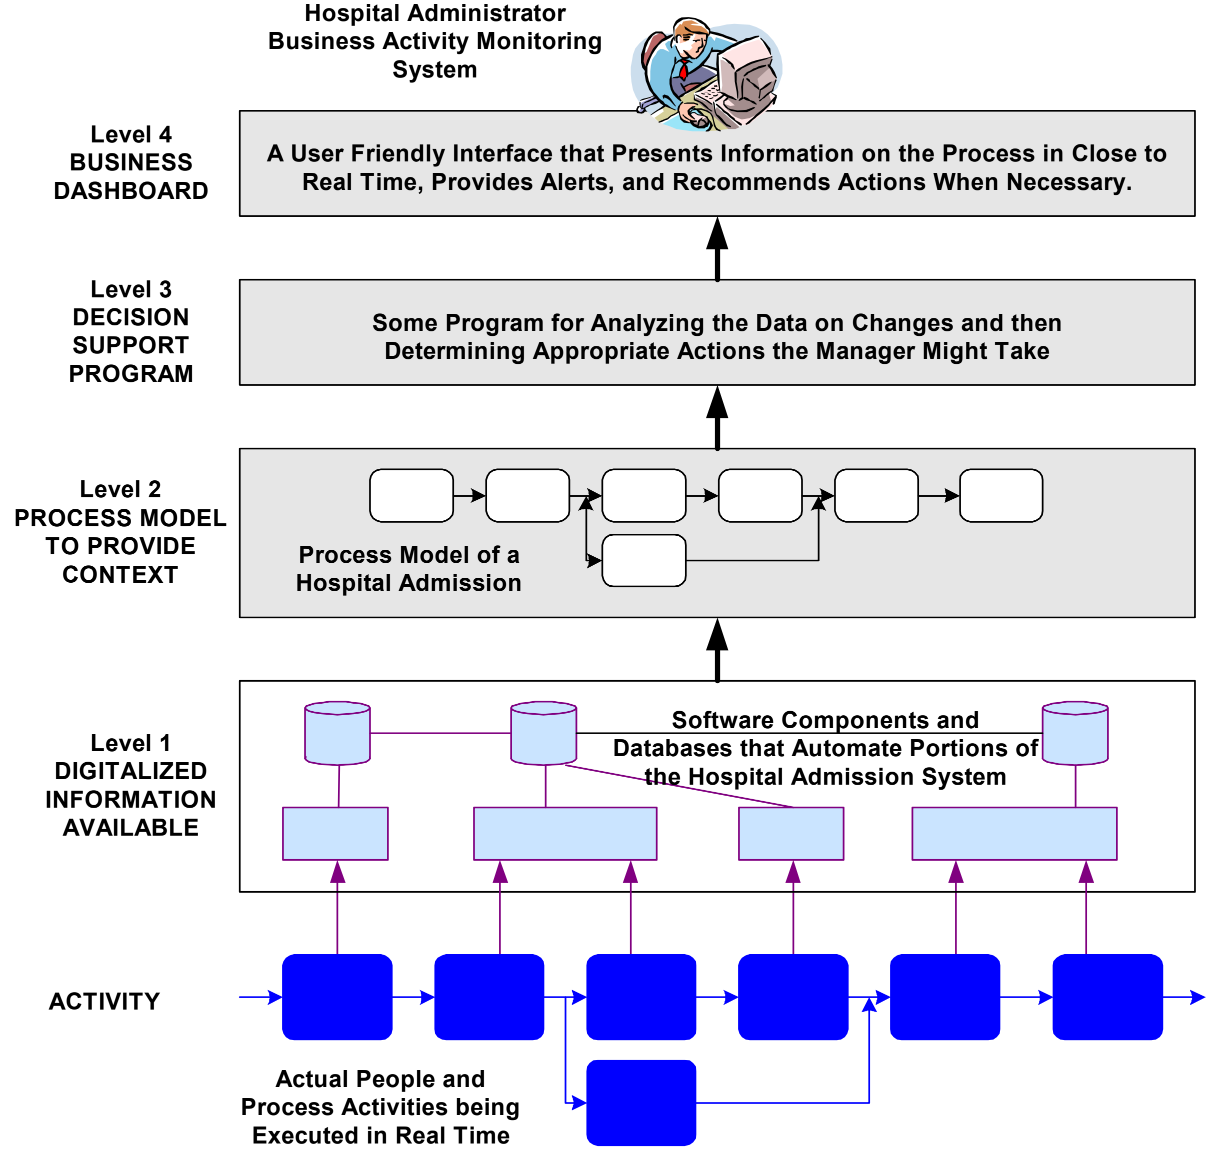

Figure 3. Modules required for a serious BAM system.

simply provide an example of one way the process data can be analyzed so as to generate action recommendations.

Finally, any BAM system needs some way of presenting the information to an administrator. Most vendors speak of these monitor displays as "dashboards." The term is meant to reflect the fact that the displays often have dials, gauges or other graphic devices that alert senior managers to changing conditions. Equally important, however, is a context for the information presented. If the manager is simply monitoring the admissions process, then several gauges might be adequate to let the manager know what is happening. If the manager is managing the entire customer lifecycle process, however, the manager is probably going to need some general picture of the process as a whole to pinpoint the problem area. For example, admissions might be normal, but discharges might drop, creating a shortage of beds for the new patients. Or, abnormal delays in obtaining lab test results may delay procedures that result in longer patient stays. So, some kind of graphic should probably help the manager pinpoint the area of concern. Then, within each area, data ought to be summarized. Finally, if alerts are displayed, information about the nature of the problem and possible corrective actions should also be presented.

3. Supporting Managerial Decision Making

At this point, let's focus on what we termed Level 3 in Figure 3, the specific approaches that are available to analyze data and generate actionable recommendations to process managers. There are, in essence, three more or less independent techniques that a BAM system might use to analyze data in order to make recommendations to a manager: Rule-Based Systems, Business Intelligence-Based Systems or Simulation Systems. (See Figure 4.)

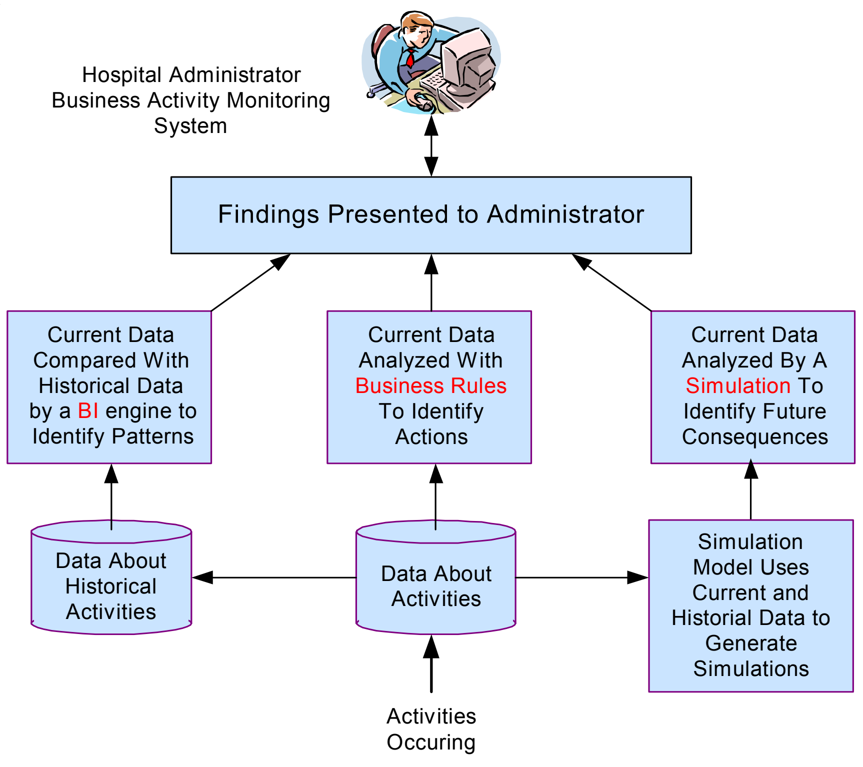

Figure 4. Sources and nature of decision analysis.

Rule-Based Systems

The most straightforward approach, as we suggested earlier, is to use a set of rules to analyze existing data. Whenever the current data triggers a rule, it fires, either generating a recommendation, or triggering still other rules that ultimately lead to a recommendation. There are different types of rule-based systems. The simplest are rule systems that are embedded in data base management programs. The more complex rely on an inference engine that processes rules which are held in a rule repository.

BI-Based Systems

An alternative approach is to rely on historical data and Business Intelligence (BI) techniques. In this case, the BAM system might compare current data to historical data in an effort to identify a pattern. Such an approach might, for example, identify a slight increase in the use of certain medicines, correlate that with the season and a slight rise in admissions, and detect the onset of a flu epidemic before the Doctors recognize that they are facing an epidemic. Using the same approach, the system might suggest to the administrator shifts in staffing and drug orders to bring today's activities in line with the staffing and drug order patterns that were relied upon during the last three years' epidemics.

BI is an umbrella term used to refer to a broad collection of software tools and analytic techniques that can be used to analyze large quantities of data. The data used by BI systems is usually stored in a data warehouse. A data warehouse consists of the data storage and accompanying data integration architecture designed specifically to support data analysis for BI. The data warehouse integrates operational data from various parts of the organization. Unlike operational databases, which typically only include current data, a data warehouse incorporates historical information, enabling analysis of business performance over time. Data warehousing is considered an essential, enabling component of most BI and analytic applications.

BI usually relies on pattern matching algorithms derived from Artificial Intelligence (AI) research, or, in some cases on specially designed rule-based systems.

Simulation-Based Systems

Simulation systems rely on a process model and a set of assumptions about how work flows through the process. The assumptions are often based on knowledge of historical flows, but can be combined with current data. In essence, a simulation system projects future states of affairs. Thus, a trend that may be too small to attract attention today, may, if unchecked, result in major problems in a few days or months. A simulation system can run repeatedly simulations with the latest data and alert managers of potential problems before they occur. It can also be used to determine how proposed changes will affect the process in the future.

Mixed Systems

In many cases, BAM products will combine various analytic techniques. Thus simulation systems could also employ rules to facilitate certain kinds of analysis. Similarly, rule systems might also employ BI and other techniques.

Table 1 provides a summary of some of the advantages and disadvantages associated with each analytic approach.

	Advantages	Disadvantages
Rule-Based Systems That Analyze Current Data	- Can be provided for very specific tasks and can operate independent of other systems. - Can be easily defined and tested. - Is a well understood approach.	- Can become complex if the range of variables are extensive - Can become complex to test and maintain over time as the business changes - Difficult to graphically validate relationships to processes
BI Systems That Use Historical Data	- Have powerful algorithms for analyzing trends and patterns - Can pull together data from EAI and ERP systems and from best-of-breed applications.	- Work best when used in conjunction with large amounts of data - Must create Data Warehouse as precondition to using BAM capabilities - Aren't designed to use a process context for reporting - Have not been designed to operate in real-time.
Simulation Systems That Project Consequences	- Provide capability to model highly complex and dynamic processes - Provides better insight into the future predicted state of the business, based on validated process flows - Can provide unique insight into longer range situations. - Notices gradual changes not so readily identified by other techniques. - Can take advantage of both BI and Rule-Based approaches to provide even better simulation results at run time.	- Simulation technology is not so well understood at the user level - Development of complex simulations requires specialized knowledge. - Requires an initial analysis and specification of a business process model.

Table 1. Comparison of Analysis Methods

Obviously there is no one best analytic technique for all situations. As a generalization, rule-based techniques are best for narrowly focused, specific analysis. BI techniques are best when there is a lot of historical data and you are reasonably sure that future situations will be like those that occurred in the past. Simulation techniques are best for more complex and changing situations.

Early Examples of BAM Offerings

Consider some early examples of vendors who are creating BAM solutions.

The ERP vendors offer application modules. Their modules are designed to store information in a common database. Most integrate the flow of information between modules with a workflow system and rules. Depending on their existing technology, most have begun to create interfaces for business managers that are driven by information from the database. They have difficulty including data about applications other than their own. A good example of this approach is SAP's BAM offering.

Similarly, workflow vendors are well-positioned to enter the BAM market. Workflow tools depend on an initial analysis of the flow of material and information between activities. They usually supplement their models of a process with rules to make it easy to control and alter the flow. Although most workflow systems are small in scope, some cover entire business processes, and some are capable of interacting with non-automated processes by providing workers with tasks and recording when they indicate they have completed the tasks. For a good example of a BAM solution being offered by a workflow vendors, see SeeBeyond's BAM offering.

EAI vendors offer systems that integrate a variety of different applications. They also depend on workflow-like systems, supported by rules, to describe how various applications are related and how to manage the flow of data between applications. EAI vendors have also begun to create BAM interfaces that allow managers to see how applications are functioning. These systems usually have trouble including information about human activities that may play a role in a large-scale process. A good example of this approach is TIBCO's BAM offering.

IBM acquired Holosofx, a business process modeling vendor, and has begun to integrate this business modeling tool into their WebSphere middleware environment. Holosofx, at the moment, relies on its strength in monitoring IBM MQSeries workflow data, but it can also monitor other middleware data flows to provide a manager with a management dashboard. Expect to see a variety of BAM offerings from Business Process Modeling tool vendors and from middleware vendors. The problem with most of these approaches is that they have to be hand-tailored for each application.

In addition, a number of Data Warehouse/Business Intelligence vendors have begun to offer BAM modules. In this case the vendors are already storing data and already have powerful BI tools to examine the data. What they usually don't have is a process model to provide context for the data, nor are they adept at providing analysis in near-real time. However, most are working on BAM extensions to their suites. A good example of this approach is the Business Objects BAM offering.

Finally, a BAM solution is being developed by at least one Simulation vendor and that is the focus of this white paper. Like workflow, EAI, and business process modeling tools, simulation tools already rely on the creation of models of processes. Simulation tools are especially flexible in their modeling capabilities, since they are often used to model large-scale processes. Simulation tools normally rely on rules that incorporate statistical assumptions and on historical data to execute their models and generate data on possible future scenarios.

Most of the current products rely on rules to offer limited decision support. The two exceptions are the DataWarehouse/BI tools that rely on their BI algorithms to identify historic patterns and the simulation tools that can use current and historical data to run scenarios and project future states based on current trends.

Obviously, there are many combinations that are also possible. Many Business Modeling tools also support limited simulation, and, increasingly, most of these tools can be integrated with offerings from more powerful rule-based tools like those from Pegasystems and Fair-Issac.

There are no mature BAM tools. All of the offerings, to date, are early products that have been assembled from the features of the vendor's current product. As the market grows and matures, more comprehensive and specialized BAM offerings will appear.

4. Simulation

At this point, since our primary focus is on the use of simulation to support BAM, let's consider simulation in more detail.

Simulation is the use of computing to mimic the behavior of a real-world system or process. Simulations are represented and executed from models that are abstractions of those real-world systems or processes. Simulation modeling is a broad topic that in its entirety is beyond the scope of this paper. We will focus on discrete-event simulation models since they are more likely to add value to business process analysis, business process management, business monitoring and decision support. Discrete event simulation models are based on events that occur within and are acted upon in a business process. By using random occurrences of those events, the simulation can mimic the dynamic behavior of the business. There are commonly two types of implementations of discrete event simulation models:

- **Probabilistic** – the use of probability distribution functions to represent a stochastic process (this type of implementation is also commonly known as Monte Carlo)
- **Deterministic** – the events that are input into the simulation will produce the same set of results over time

This section will focus primarily on probabilistic discrete event simulation models. These models are used to define the business work steps, and specifically, the entities that flow through the business, as well as the resources required to perform each work step. Once the process modeler has created the basic process diagram, a simulation expert must enter information about the flow of events. The timing and occurrence of the events are based on probability distribution functions, which reproduce the behavioral dynamics of a real world business process. The developer of the simulation must choose probability functions that reflect the behavior of a given process. The process model and the information about the events are entered into a software program, which can then "execute" the simulation. By entering initial data, and executing the system, the modeler can determine future states of the process. As the simulation executes, the events are generated, the entities flow through the process, the delays are sampled, and the resources are used – all using probability distributions to produce the real-world randomness of the business process.

Consider a simulation of our hospital patient lifecycle process. We have a probabilistic function that determines how many patients who enter Admissions are routed to the maternity ward. This function is based on historical data. If we indicate that 100 patients enter on Monday, our system will automatically assign a portion of them to the Maternity ward. If we indicate that 1000 patients enter on Tuesday, the same formula will assign a proportionally large number to the Maternity ward. As the number of patients entering the Maternity ward increases, resources, ranging from beds and rooms to doctors and nurses, must be increased. By running different simulations we can determine just what resource would run out first and develop a plan to deal with the constraint if we expect that we might one day have that number of Maternity cases.

A Detailed Simulation Study

Now consider a more detailed simulation. In this case, consider what happens to patients entering the emergency room who need operations. This will illustrate how a discrete event simulation model can successfully support business process analysis. The purpose of this model was to consider the potential business impacts of opening a new emergency room facility that increased the number of treatment rooms by fifty percent. As with any good simulation model, the primary goals of the simulation were clearly stated. They were:

- Examine the resource impacts (resource utilization, cost, etc.) from opening the new facility. Resources include physicians, nurses, support staff, supplies, and facilities.
- Examine the impacts to patient treatment cycle time. This particular hospital maintains an average treatment cycle time of 2.1 hours per patient on average. The new facility should seek to improve the cycle time at best and maintain the cycle time at worst case.
- Examine the cost of specific activities in the operation (activity based cost).

You will notice that the goals of the simulation are kept to a few key business metrics. This is one of the most important things to remember when applying simulation technology – Don't try to model the whole world all at once. Models should be built incrementally or evolved to solve increasing complex problems over time. The measurements need to have a business focus and not focus on technical problems that the management is not immediately concerned with.

The model is made up of three basic ingredients: entities, activities or work steps, and resources. The primary entities in the model are the types of patients that may arrive randomly at the door of the emergency room. The model makes assumptions about the number of individuals who will show up within any given period, and what types of problems they will have. These probabilistic assumptions are based on historical data. Some types of patients arrive more often, on average, than others. Some will have higher priority than others, based on severity of the illness or injury. The work steps include the primary activities performed, ranging from triage to initial evaluation, to treatment, to out-processing. The amount of time to perform each work step depends on the type of problem, and this, again is based on historical data and is represented as a probabilistic function that introduces a realistic variation into each patient's treatment. For example, the amount of time to perform the triage may be a random sample from an Exponential or Poisson distribution. The Normal distribution would have a mean value and a standard deviation to represent, statistically, what occurs in the real world. The resources required to perform each work step are assigned and include such attributes as the number available, cost, and planned down time. The number of each resource type may also be selected from a probability distribution in cases where a business has variations in the number of resources available. In other words, the model will consider things like employees taking leave or breakdowns in machinery.

Figure 5 illustrates a screen shot of the emergency room process model.

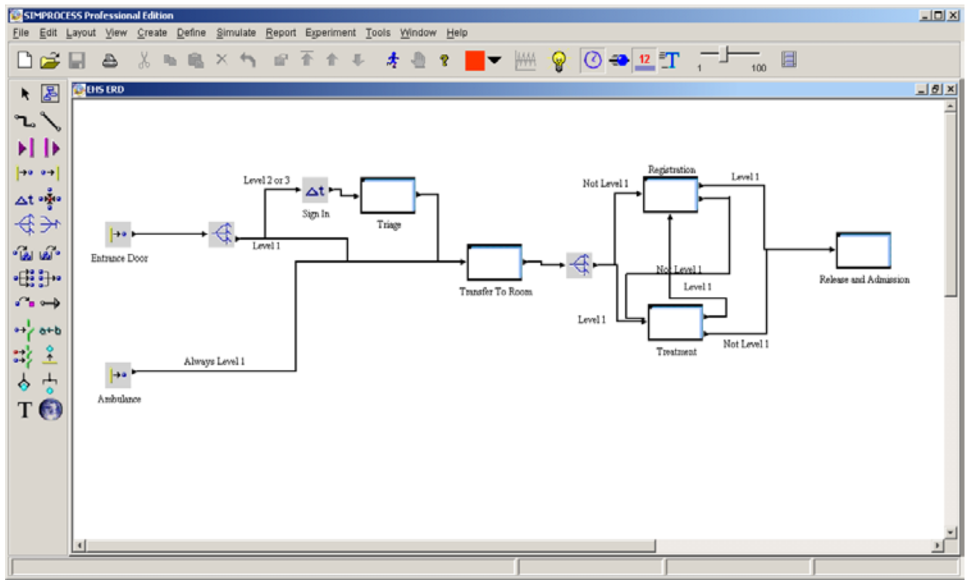

Figure 5 – SIMPROCESS™ Screen Showing Top Level Process Flow Model

Notice that the model provides a concise, high-level view of the process. Each of the major sub-processes shown in Figure 5 has been defined in more detailed process models. This is another important aspect of simulation models. Where possible, the overall model should be kept as general as possible to assure it can be easily understood and that it will generate answers to the important business questions. Models with too much detail actually make it harder to identify and answer important questions. At any point in the analysis, if more granularity is needed, you can always drill down to the lower processes and model their sub-processes and activities to get the results needed. This technique allows you to have some processes modeled in detail while others are simply pass-through boxes. It lets you see the whole process at the higher levels and drill down in later spirals of analysis, if needed, without having to change the high level model.

The patients (the key entities) in the model are categorized into three types:

- Level 1 – the most critical or severe, needing immediate attention
- Level 2 – critical patients needing attention but are probably in a life-threatening condition
- Level 3 – patients that do not need immediate attention

The three types of patients arrive at the door at varying rates. Figures 6 through 8 below show pop-up Windows that SIMPROCESS™ presents to analysts who are entering the information needed to set up a simulation. The three Windows require the entry of information about the activity that occurs at the emergency room entry door. In this case, the analyst has decided to enter three different rates to reflect the patient flow historically experienced by the hospital. Figures 6 and 7 focus on the arrival rates for Level 1 patients.

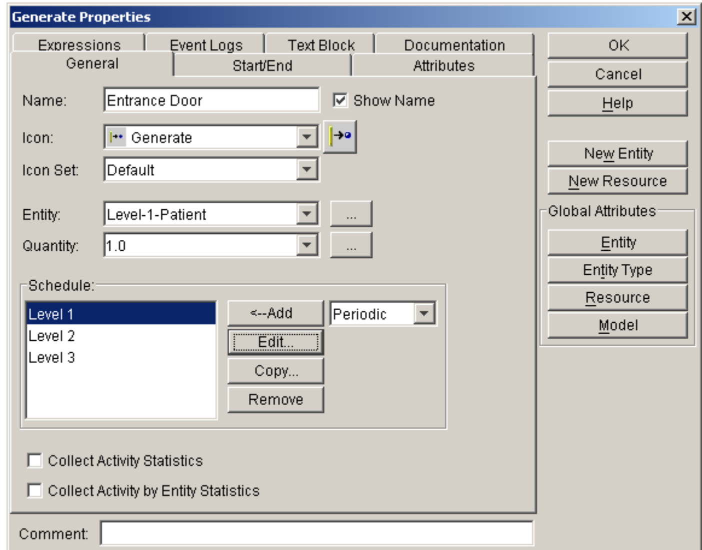

Figure 6 – SIMPROCESS™ Screen Showing Entity Definition

In Figure 6 we define the three types of patient entities that will flow through our business model . Since the patient types, Level 1, 2 and 3, arrive at the hospital at varying rates and at varying quantities over time, we must define separate inter-arrival schedules for them in SIMPROCESS™. In Figure 6 we have set up three schedules and named them Level 1, 2, and 3 to correspond to the patient entity types. By choosing the Level 1 schedule in the Figure 6 dialog and choosing the Edit function, we can now describe the inter-arrival schedule for the Level 1 patients, as depicted in Figure 7.

The dialog presented in Figure 7 allows us to create schedules within schedules. That is, the inter-arrival of Level 1 patients vary, depending on the time of day it is. In this model, we have defined three inter-arrival schedules for the patients: Day Shift, Evening Shift, and Morning Shift. By selecting the shifts in Figure 7 and selecting the Edit function, we can define how the Level 1 patients will arrive for each daily period, or shift. This concept can be extended, as necessary, to include varying the schedules for weekends, time of month, or season. This is an important and powerful capability when using simulation models since this mimics how the business really encounters patients statistically during different periods of time.

Once a shift is selected from the choices in Figure 7, the dialog box shown in Figure 8 is presented allowing the user. This dialog box allows the user to define the types of patients (entities) that could appear during that shift. Patient quantity and rate of arrival are defined as well. Since the hospital manages by shifts and keeps their metrics by shift, it only makes sense to have the simulation model be consistent with the shift schedules.
The inter-arrival rates for Level 1 patients are 0.109375 patients on average every hour. This is based on the past history at the hospital having an average of 111 total patients per day with Level 1 patients making up 2.775 of those on average. For the day shift, it is 0.875 patients over an 8-hour period resulting in the 0.109375 per hour average. You will notice the use of a probability distribution function (Poisson) to provide a representative statistical curve of how the patients arrive. If we just use the

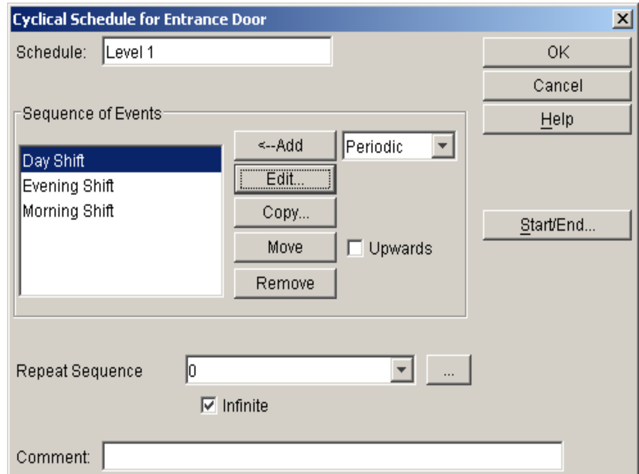

Figure 7 – SIMPROCESS™ Screen Showing Entity Schedules

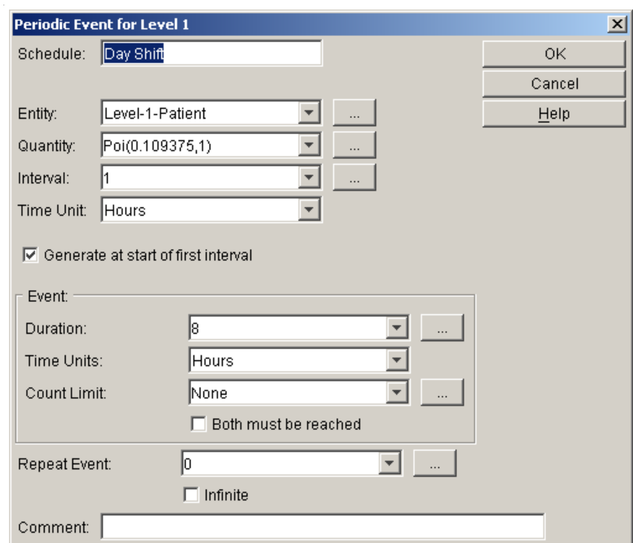

Figure 8 –SIMPROCESS™ Screen Showing Patient Arrival Rates

average, then our model would not be probabilistic and would not provide a true representation of how things really occur. If events occurred in the real world at steady states, then simulation analysis would not be needed. However, that is not the case. Not only do the patients arrive at random, resources fail (x-ray machine breaks), go down (person takes leave or sick), are occupied or busy (physician takes varying times to diagnose and treat the patient), or deplete (for consumable resources, e.g., oxygen canisters) at random rates. Likewise, the time it takes to perform work steps varies and is not accurately represented by a steady state average.

Once the entity (patient) arrival rates have been defined and the process flow diagram (see Figure 5) has been developed, the timing and resources on the work steps must be set to account for the time delay that is required and exactly what resources are assigned or consumed to perform the work steps. Figure 9 shows a pop-up Window that an analyst could use to define the delay time and the resources required to perform one of the treatment steps (Perform Treatment) in this model.

In this model, we have defined the time to perform the treatment work step as a probability distribution function. We have used the Normal distribution function with an average of one hour and a standard deviation of 0.25 hours. When the simulation is run, the entities (patients) will flow through the system and when they reach this work step, the time allotted to perform the task will be randomly sampled from the Normal probability distribution function. This will create the randomness in the time it takes to do the task just as the probability distribution function in Figure 8 defined the inter-arrival rates of patients.

Figure 10 shows the resources required to perform the actual treatment. Resources are globally defined in SIMRPOCESS™ and can be used or assigned to any task in the process model where they may be needed. The number of resources available as well as the number required to perform a certain work step can also be ramdomly sampled as described earlier using probability distribution functions.

You will notice that in this model, the number of nurses required is sampled using a probability distribution function to simulate how the real world process works since the same number of resources is not always the same, depending on the type of treatment. You will also notice that the physician resource is a shared resource since physicians bounce between five or six patients on average. This model is sampling a Poisson distribution with a mean of twenty percent of a physician's time on average being consumed by a single patient.

Additionally, resources can be required, based on different scenarios. In Figure 10, we have specified that all the resource requirements (physicians and nurses in this case) are required at the same time to perform the work step. This means the entity (patient) will wait (this is where queuing and bottlenecks are discovered) at this work

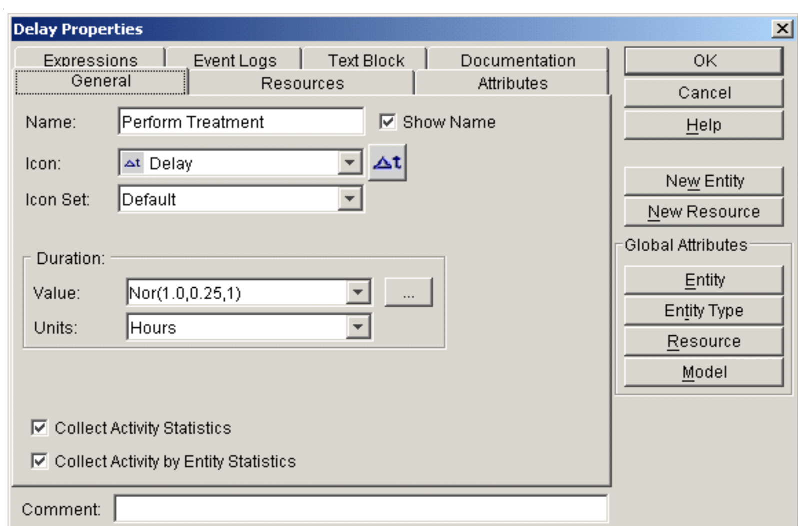

Figure 9 – SIMPROCESS™ Screen Showing Activity Delay Time

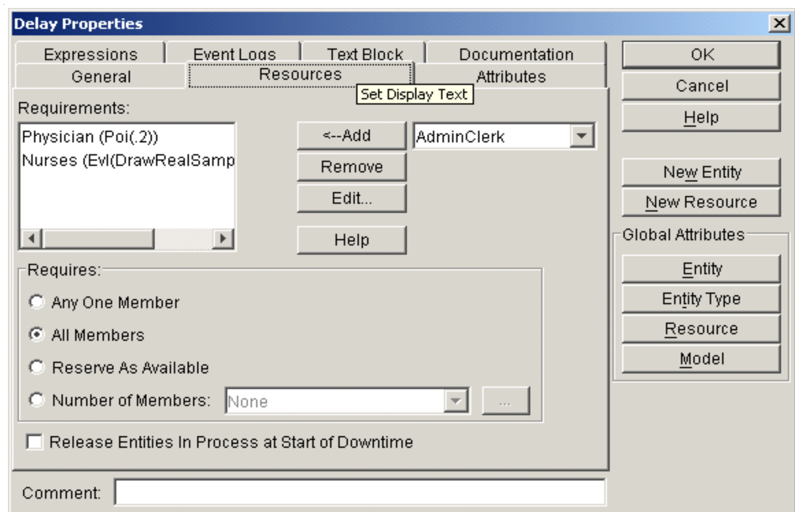

Figure 10 – SIMPROCESS™ Screen Showing Required Resources

step until both resource types are available simultaneously before the task can start. In some cases, we may list substitutable resources and set the requirements to "Any One Member" to specify that the work step can start as soon as any one of the resource types are available to do the job. The use of the "Reserve as Available" option in Figure 10 allows us to lock one of the resources needed as soon as it becomes available and wait until the others are available and lock them until all resources requirements are met.

As can be seen in the use of probability distribution functions for all the key components of a SIMPROCESS™ model (entities, activities, and resources), we can simulate even the most complex dynamic behavior of any business process.

Recall the goals of the simulation model. We used the model to do "what if" scenarios that will be influenced by the new emergency room facility and 50% increase in treatment rooms. This model looks rather simple on the surface, but is a very powerful tool to do the "what if" analysis and determine if the business metrics and goals can be met. Figure 11 is a plot of the treatment cycle time (one of their most important business performance metrics) with the old facility plotted on the left and the new (50% additional rooms) plotted on the right. The plots in Figure 11 are clipped for the first 100 hours of simulation time.

Multiple simulations of "what if" scenarios were used to ensure the hospital could achieve its treatment cycle time business goals. For example, if the number of physicians, nurses, and lab technicians remain the same, the additional rooms will cause the treatment cycle time to go above 3 hours. That is due to patients getting in the rooms without enough resources to treat them, hence waiting in the room for additional time (treatment time is counted from the time the patient gets to a room until the time they leave the treatment room). The simulations quickly uncovered these types of problems and allowed the hospital to play additional "what if" games to find the right balance of resources and costs to make optimal use of the new treatment facilities and resources. The objective of all this is being able to avoid embarrassing impacts to treatment of patients (for this particular hospital, cycle

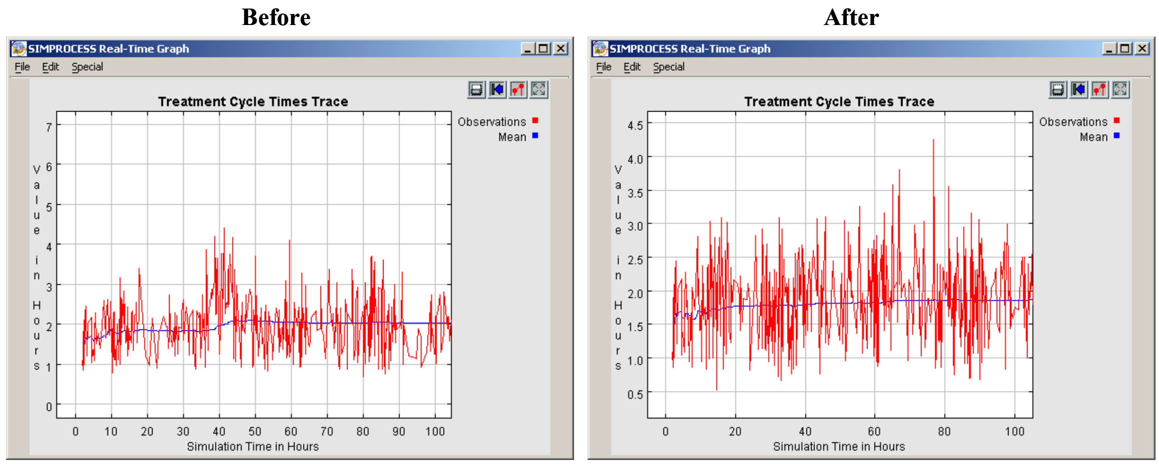

Figure 11 – Two SIMPROCESSS™ Screens Showing Patient Treatment Cycle Time: Before and After

time is one of their marketing nuggets) or worse, poor utilization of a new facility from an operating cost standpoint. As can be seen from the "After" plot in Figure 11, the treatment cycle time is actually improved when adjustments are made to the physician and nurse resources. In the old facility, adding resources would not have improved treatment time since the bottleneck was the number of room resources.

5. Using Simulation for BAM

Section 4 focused on how simulation models are used in traditional business process analysis (BPA). Models are developed and validated for an existing business process (an As Is model). Then, various changes (What If models) were imagined and tested, via simulation, to see if they would improve the efficiency of the process. In the simulation process, bottlenecks and specific inefficiencies were identified and eliminated.

Although this use of simulation is very common and valuable, most organizations use the simulation models developed in this manner during a limited improvement project and then set them aside as the new process is implemented.

It is possible, however, to use simulation to support BAM systems. In essence, the new process model is maintained in the simulation environment, and new simulations, using the latest data, are run periodically. Triggers or rules are used to identify problems. The current level of admissions at our hospital, for example, may be slightly higher than it was a week ago, but not high enough to trigger an alert. On the other hand, if a simulation is run using the past month's data, it may be determined that an underlying trend is present that will result in unacceptable rates of admissions in three weeks.

Figure 12 suggests how a Simulation Environment might be linked to software applications and databases to allow the simulation system to be run with real-time data. In effect, data generated by the actual process would be used to run simulations. The simulation system would have alerts and rules to identify problems and suggest alternatives. In some cases, the simulation system might include alternative activities.

If the system determined that the normal process would generate problems, it might try pre-packaged alternative approaches to a specific activity to see if a particular set of changes would result in an acceptable projection. This is one way in which a simulation might be able to combine alerting a manager about problems with suggested remedies.

The approach described in Figure 12 goes beyond what is available in today's simulation products. In effect, it changes simulation from something done to test alternatives, and makes it, instead, a way of dynamically determining what will happen in the future if the current state of the process is allowed to continue without change.

Obviously, this diagram greatly simplifies what is involved in using simulation for BAM. The developers need, for example, to identify the events or data items that will be monitored. Similarly, they need to insert triggers or create rules to determine when managers should be alerted. And they need to determine how frequently the simulation should be run, and how far ahead it should project. These are all decisions that will need to be made in the context of a specific company process. These decisions are not unique to the use of simulation. The same events, triggers, and

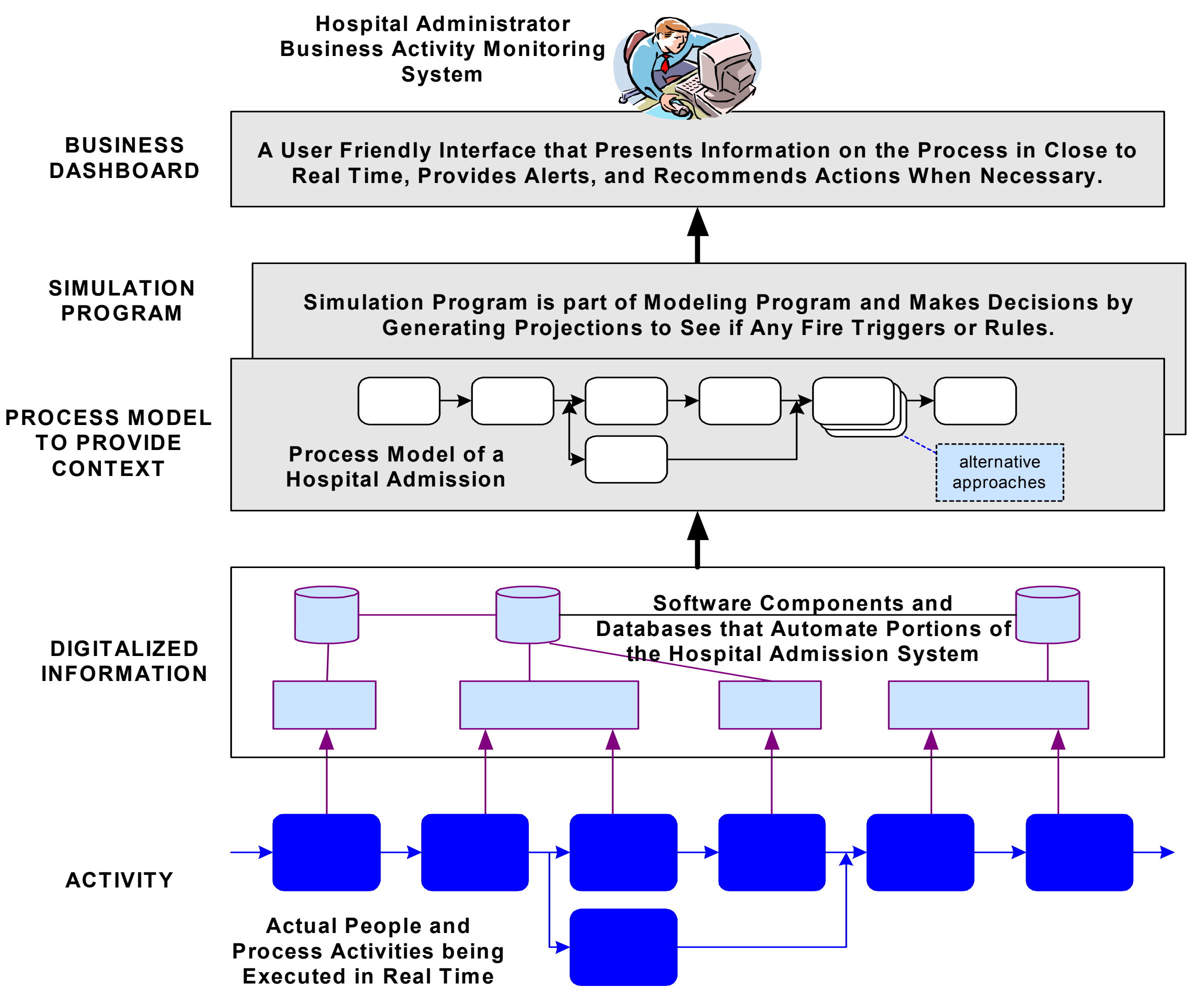

Figure 12. A BAM System Based on a Simulation System.

information monitoring will need to be defined for any BAM solution. Simulation is merely providing an additional dimension to already useful BAM solutions.

Rule-based systems, by themselves, only look at the present for problems. BI systems use historical data to look for current patterns that might suggest problems, but can only identify problems that have already occurred before, in the past. Simulation systems can combine the best of both and add the ability to look for future states that suggest problems and then dynamically try alternative assumptions to identify changes that the manager could make today to avoid the undesired future state.

Companies that have already used simulation and have been happy with the results find the possibility of reusing their simulation investment to create powerful BAM systems exciting.

CACI, for example, has used simulation in the loop for real-time decision making for customers such as the Department of Defense. For example, DOD training systems are a good example of using simulations along with feeds from of real-time operational systems to create scenarios for training purposes. The result is a hybrid tool that operates partially off validated probabilistic models along with real world events at real time. This concept, when extended to more mainstream solutions such as BAM, opens up some interesting ways of significantly improving the benefits of BAM (especially if process modeling investments have already been made with BPA models).

6. A Case Study

If we consider the hospital example described early in this paper, we can get an idea of how this capability could be put to practical use. The process simulation models were already built as part of a business process improvement project and were extended to serve as a key ingredient in a simulation-based BAM solution.

Let’s imagine a scenario of where this particular hospital -- which has operational information systems that capture the key events in the business activities such as patient sign-in, capture initial triage data, resource clock-in/out, and post treatment data capture -- can actually provide data to the simulation models in real-time. The data feeds can be done at periodic intervals to get a “look ahead” on the impacts to patient treatment cycle time based on the current resources, patients received, expected future patient arrivals (see section 5 for simulation of patient arrivals based on empirical data), and the validated standard processes documented in the simulation modeling tool. Remember, the patient inter-arrival rates were based on probability distributions, as well as the activity delays and resource assignments. These probability distributions provided us a powerful tool to mimic real world dynamics; however, if an unexpected peak occurs in any of the variables, the probability distribution function chosen may not have predicted a worse-case situation. These spikes may be statistically insignificant over a long period of time but it could skew the performance of the business for several days or weeks, impacting the business goals as well as the P&L. When using the actual real world data feed along with the probability distribution for future expected events, we get a hybrid model that is based on both real and simulated data. This is technically a pre-load of the model with current data while the simulation will used expected events as it simulates to the future, for example: 2 weeks, 1 month, 6 months, etc. The result is the capability to alarm management with data to make decisions (such as call in temporary resources or change to a crisis process) based on the simulated future.

One might ask what the benefits of this would be over traditional BAM type dashboards. The difference is that the current data and rules associated with alerting management from the traditional BAM solution may see only gradual changes in the variables of the business performance and may not have enough insight into the long range impact to know to alert management. With the simulation-based BAM, we add another dimension to the BAM solution. We can look several weeks or months into the future and predict what impacts may be experienced in overall key business performance measurements while there is still time to affect those impacts. The hospital example we used in this paper was based on the customer's very important business metric of average patient treatment cycle time. If certain run-time events drive the average up, due to unforeseen peaks, the hospital loses its ability to use its performance metric as a marketing tool. There are many other metrics that could be looked at, such as efficient resource utilization. In many cases, staffs are overworked and exhausted before management is aware of it, and exhausted staffs can create even more business performance problems such as quality of service.

Another example situation that could arise in the hospital model is if a flu outbreak starts an upward spike of Level 1 and Level 2 patient types that are normally not expected. Since Level 1 treatments can disrupt the primary process flow due to priority of patient care, and Level 2 treatment takes longer on average to administer, the BPA model would not have uncovered those scenarios unless an explicit "What If" experiment were run. With the simulation-based BAM solution, the actual data affects the simulation and, in theory, "reprograms" the simulation and gives a better future picture of the organizational performance.

The proposed simulation-based BAM does not have to have user intervention as in BPA simulation analysis. Once the model is built, it runs completely in the background and presents management dashboard data that is in a business metrics form. However, another consideration is to let the simulation run through multiple "fall back" or "crisis" process alternatives based on the real-time data, and present the user with choices that would then, in turn, be related back into BPM solutions for temporary process adjustments.

Figure 13 is an example of the type of performance metrics that could be placed on a management dashboard. The dashboard in this example is showing both the real time data that a traditional BAM would provide in the top half of Figure 13. The bottom half is a simulation of five days into the future using the validated business process model.

The dashboard information is described as follows:

- The actual real-time reporting (traditional BAM scenario) of business information is depicted in the top half of Figure 13 and includes the following graphical gadgets:
    - A meter to the left of the dashboard that depicts the running average of patients per day for the past month. The colors on the rim of the

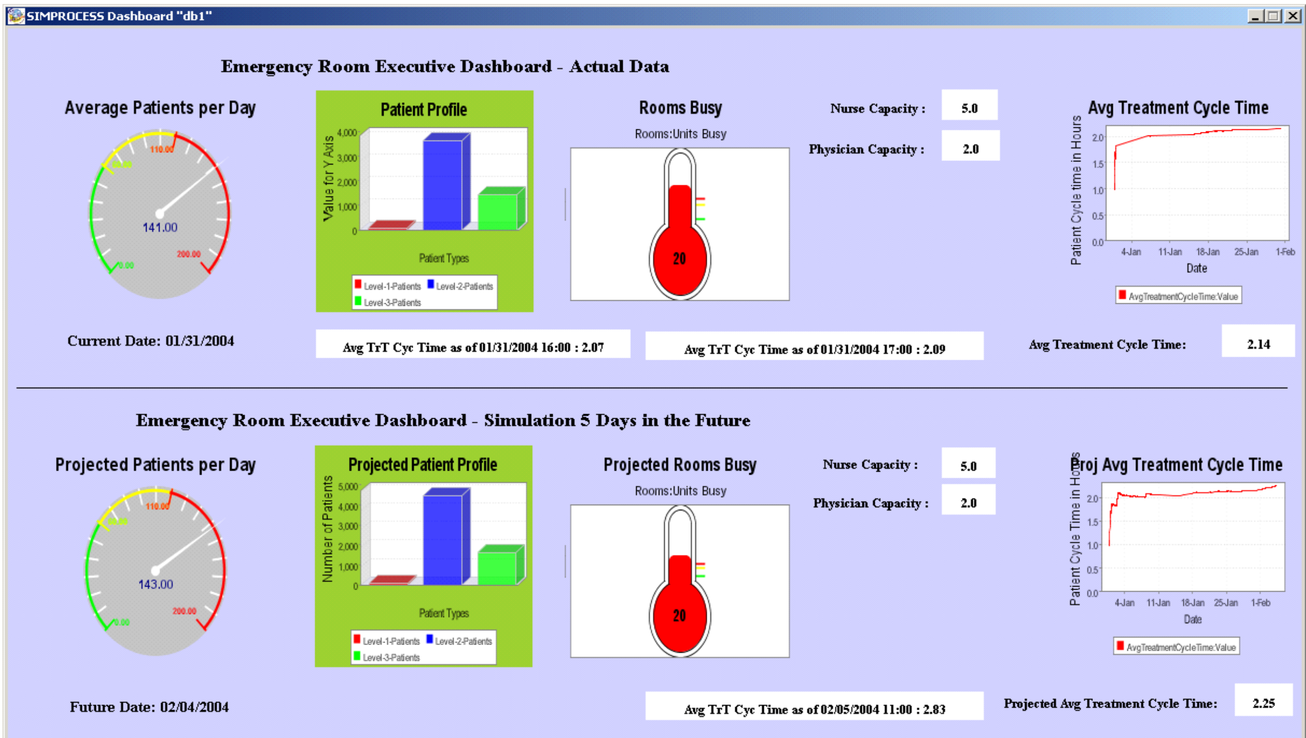

Figure 13 – Management Dashboard

meter are used to help the manager quickly see the numbers that represent critical values (i.e. red indicates a critical situation).

- A bar graph that depicts the total patients for the last month for each type of patient.
- A thermometer gadget that depicts the utilization of treatment rooms. The fill in the thermometer changes colors to indicate critical values.
- A trace plot on the far right that depicts the average treatment cycle time of all patients over the last month.
- Two text values that depict the current levels of the most critical resources – nurses and physicians.
- The text field to indicate the date of the dashboard information.
- A text field that indicates the average treatment cycle time up to the last hour of reporting. This is important to see the difference from the 30-day average and the last hour or so of activity. This field is used in conjunction with the next field.
- A text field that is used with the previous field and indicates the average treatment cycle time for the previous 24 hours in this example. Obviously, this time span can be set to a wider range based on the business being monitored.

- The projected (simulated) business information is depicted in the bottom half of Figure 13 for five days into the future (Feb 5th in this example) and includes the following graphical gadgets:

- A meter to the left of the dashboard that depicts the simulated average of patients per day up through February 5th.
- A bar graph that depicts the simulated total patients expected based on the past month and the probability distribution function in the model.
- A thermometer gadget that depicts the predicted utilization of treatment rooms.
- A trace plot on the far right that depicts the predicted average treatment cycle time of all patients up through February 5th.
- Two text values that depict the predicted levels of the most critical resources – nurses and physicians.
- The text field to indicate the date of the simulation for the dashboard information (February 5th in this example).
- A text field that indicates the expected average treatment cycle time for the last 24 hours of actual data and up through February 5th of simulated data.

The data reported in Figure 13 was based on certain events occurring in real time and as can be seen from comparing the data, the impact is minimal and probably not alarming to management based on the traditional BAM data. However, when the simulated data up through February 5th is provided, you can see the urgency of the problem as the key business metrics are affected. For example, the average treatment cycle time goes from 2.14 hours average over the last month to 2.83 hours for the last 24 hours of real-time data up through February 5th. An interesting piece of information is that if you focus only on the last 24 hours in the actual data (2.09 hours average which is less than the monthly average), you can see how not including simulated data can delay alarming management to a building problem

The events that caused the situation to become critical are as follows:

- Starting on January 15th the average arrival of patients increased by 6 per hour. This was due to a flu and cold outbreak.
- On January 30th, one of the nurses became ill and had to take emergency leave for up to 10 days.
- On January 31st at 4:00pm, one of the physicians and one additional nurse had to take emergency leave due to exhaustion and illness. Notice the number of nurse and physician resources on the dashboard have decreased from 7 and 3 to 5 and 2 respectively.
- An additional 10 patients arrived on January 30th due to two separate accidents.
- An additional 20 patients arrived on January 31st at 11:00am due to multiple accidents.

Since the traditional BAM reporting does not simulate ahead to predict the queuing theory problems of the growing arrival rates and the drop in resources, as of 5:00pm on January 31st the actual data is not showing any major impact to the business metrics. The simulated BAM data, however, predicts significant problems growing rapidly over the next several days and risks driving the average patient treatment time significantly above the goals set by management. If the problem is allowed to spike the treatment time, it could take a couple months of improved performance to pull the average back into line with management goals.

Simulation-based BAM solutions are achievable today. CACI's model-view-controller (MVC) architecture provides for separation of the simulation from the front-end analysis tool, therefore providing a server-based, GUI-less, and scalable capability. It provides connectors to outside applications through Java-based remote calls and/or XML, which are needed to feed the real-time data to the simulation tool. The example above is easily implemented with any vendor operational application with simple messaging capability.

7. Conclusion

Gartner suggests that BAM will become a major corporate concern in the next few years. Most large organizations will at least explore the possibility of improving business process management by creating systems that provide a broad overview of a process that can provide near-real-time information and advice to process managers. A variety of techniques will be used. Some "BAM" systems will, in fact, monitor subprocesses. Some will use rules to alert managers about specific real-time problems. Some will be based on simulation engines and use models that allow the system to project future events from the current state of the process and then dynamically generate alternative options to identify what changes, taken today, would maintain the process in the most efficient manner over the long term. We believe that simulation-based BAM will prove to be the most powerful and flexible approach to BAM and will increasingly be relied on by those with the more complex processes.

Note

[1] Gartner Group. *Business Activity Monitoring: The Data Perspective*. February 20, 2002.

For more information about SIMPROCESS, check www.simprocess.com

Authors

Joe DeFee is Senior Vice President and manager of the Advanced Systems Division Group at CACI. He has 24 years of information systems design, software develop-ment, enterprise architecture development, and business process reengineering experience. For the last 12 years, he has focused on business process reengineering, business process simulation technology, software reengineering, and aligning information technology to business objectives for customers. He is the co-author of CACI's RENovate methodology; a formal methodology for modernizing customers business processes and information technology.

Paul Harmon is the Executive Editor of Business Process Trends and the author of *Buisness Process Change* (Morgan Kaugmann, 2003).