

Event ID: 266016
What, Exactly, Do You Mean by Business Rules?

John Sote: Good morning, good afternoon or good evening, depending on where you are in the world, and welcome to today's webcast, What, Exactly, Do You Mean by Business Rules, brought to you by Information Week, IBM and broadcast by United Business Media, LLC. I'm John Sote and I'll be your moderator today. We want to make sure this event is as interactive and participatory as possible, so I'd like to make a few announcements before we begin.

Speaking of participation, you can join our interactive Q&A session at the end of the presentation by submitting your questions at any time during the webcast. Just type your question into the Ask A Question text area below the media player window and then click the Submit button. The slides will advance automatically throughout the event. If you'd like, you may download a copy of the slides by clicking on the Download Slide button, located below the presentation window.

If you are experiencing any technical problems, please click on the Help link below the media player and that will take you to our webcast help guide. You can also contact our live technical support helpline. Their number is located within the guide as well.

Now, on to the presentation, What, Exactly, Do You Mean by Business Rules? Joining me today are James Taylor, CEO and Principal Consultant at Decision Management Solutions and Brett Stineman, Director of Product Marketing, IBM WebSphere.

Welcome, gentlemen, and take it away, James.

James Taylor: Great. Thanks very much. So what I wanted to do was try and give you a sense of what exactly are business rules, and try and show you how a focus on understanding not just rules but decisions as well is really important.

So let's go ahead and get started. The biggest challenge with business rules is that they do, in fact, get everywhere. If you look at a typical architectural diagram with some of the key elements on -- user interface, business process, your data, your data warehouse, events and so on, what you find is that you have rules in all these different places. You have rules about user interface, what fields do I have to fill in, what values are allowed in different places?

You have rules about business process, how do I escalate to this when there's a problem? How do I route this particular transaction? How do I manage the flow of my business process? You have rules in operational data storage, you have rules about what decks can be combined with which, how many relationships (inaudible -- technical difficulty) to product components or orders to order lines.

You have rules also about your ETL, your extract transform load process for creating your data warehouse. How do I combine data from these different systems? How do I ensure data quality and data integration?

Once you get into event processing I have rules as well. I have rules that help me correlate different events together, help me identify which events are significant and which ones are not. And I have rules in what we call decision services, in components that make business decisions.

Now, at one level you could say I've got all these different rules, clearly, I need to manage them, and you do. But you don't necessarily manage them all the same way, and what we're going to try and do today is we walk through a scenario and help you understand how to tell when you should use a business rules management system, when you should really manage those business rules, and when those rules are really part and parcel of something else.

Now, I don't have time today to go through every single piece of this, so I'm going to use business processes and rules in a business process context as my example. But much of what I say is just as true if, for instance, you were talking about a business event processing system. It relies on a focus on decision services, on this idea of a separate component for decision-making. So let's talk a little bit about a business process.

If we imagine that we have a business process and it's got a number of tasks that flow through it, and then we have some kind of decision point in it, we decide should we pay this claim or not, what price should we give this customer for this product today, what discount is this supplier eligible for, given what they're doing and so on. I have these decision points within my process, so what kind of rules do I have?

Well, at one level, of course, I have a whole bunch of rules about my business process. I have rules about flow control or routing, and these rules are managed by the people who define the business process. They're often fairly technical in nature, because they relate to the underlying systems and work lists and components that I have, and crucially, they're governed as part of the process.

They have a lifecycle that's dependent on the process lifecycle. If I add new tasks, if I change the layout of my process, if I replace underlying components with new components, I have to change these rules because they are part and parcel of my business process.

But I have other kinds of rules, too. I have rules about decision-making. I have rules that tell me that this claim is complete and accurate and viable and can be paid, or that might warn me that a particular claim has a high risk of fraud and so I want to treat it differently.

I have rules about my discounts or about my product pricing, and these rules are a little different, because these are business decisions that I'm making. I'm managing them on the

business side. I'm talking about a line of business managers or business analysts working with those folks who are managing these rules.

These rules, therefore, are governed by the business and they have an independent lifecycle, and we'll come back to that concept later on. But they change when the business changes, not when the process changes, and you could replace this entire process and you wouldn't need to necessarily change the decision-making.

Imagine a claims processing process in insurance. I might go from a very manual process, one where I collect paper forms or faxes, I type that data into the system, I go ahead and decide whether it should be referred for fraud investigation, fast-tracked or referred to a claims adjuster, and then I have a whole series of paper-based processes that go ahead and assign that out to different people.

Now I come back and I completely automate that process. Now I've got web-based forms and I've got instant data capture and data feeds from my agents and handheld devices being used by my claims agents, so my process looks completely different, but my decision-making has not changed at all.

I still have to decide which claims are fraudulent or potentially fraudulent, I still have to decide which claims are complete, I still have to decide which claims I'm going to pay, which ones I'm going to fast-track, which ones I'm going to refer for fraud investigation, because my business has not necessarily changed, and that decision only changes when the business changes.

Now, one of the things that I've done in the last few years is write a book. And one of the things when you write a book you discover is that you write a lot about things you didn't expect to write about. One of the things I found I had to write a lot about was about what is a decision, exactly? What do I mean by decision?

People at one level, they know what decisions are, they have a pretty good handle on the decisions they make. Well, in the context of this kind of technology setting, what do I mean by a decision?

Well, a decision is a point in time where you gather a set of data, you consider that data in some fashion and then you make a choice or a selection from the available options. In my insurance example you've got three options: Do I refer it for fraud investigation, do I fast-track it, do I refer it to a claims adjuster?

But I might have many different alternatives here. I might have a whole range of choices to make, if like in a discount case where I'm deciding what discount to offer you, I might have a whole range of discounts.

Having made that selection I commit to some action based on that selection. I don't necessarily take that action yet, that's not part of decision-making, that's part of acting on your decision subsequently. In a technical context, your decision service, this component

that makes decisions, is going to commit you to a particular action that will probably be implemented by some kind of business process or some kind of system. And it's going to take the outcome of that decision and go ahead and make sure it actually happens.

In our claims example, I'm going to kick off one to three different processes depending on what kind of decision I have made about my claim. But what that decision does is it commits me to an action, and it turns out that understanding these decisions is really important if you're going to effectively manage business rules and put business rules to work for you.

Now, if you start looking at a business process or indeed an event processing system or any of the other elements we looked at in that first slide, you start looking for decision-making words. You'd find words like choose, select, calculate, determine, assess, validate -- all of these are decisioning words.

Choose which option you're going to present to a customer. Select the appropriate product for this customer. Calculate the discount this company is eligible for. Determine who to route this to in a business context. Assess the risk of this particular loan application. Validate that this claim is complete.

When you find these decisioning words, you can turn around and you can say, "What's the question I'm trying to answer here?" I'm trying to answer the questions like, "Is this claim complete and accurate? Is this claim at high risk of being fraudulent? Is this loan one that's eligible for my lowest rate? What rate is this loan eligible for? What cross-sell offer should I make this customer at this time in this channel?"

Questions that tell you how you're going to frame this decision and enable you then to build a decision service -- a decision-making component that answers that question for all your other components.

And when you look at a business process, what you find is you find these words and you find these kinds of decisions in a number of different places. You find them in pretty much any time you have a work list. If you think why do you put things on a work list in a business process, why do you put them into a work queue? Well, you do it so that someone could make a decision, so that someone can make a choice so that you know what action to take next in the business process.

So many work lists are decision-making components. If you want to build decision-making components, building them to replace or assist in the activities supported by a work list is often a pretty good place to start.

You find it in what I call extreme branching, and again, we'll come back to this a little bit later. When you see branches that lead to branches that lead to branches, where you have a nest of conditional branches within your business process, almost always what you are doing there is you are obscuring a decision by pretending that it's just part of your process. You might find places in your business process where you have to give the

people who are performing the process detailed procedures, loan cheat sheets -- one of the cases I'll talk about later was a great example of this.

Almost always what you're doing at that point is you're describing the decision-making process to someone so that they can make the right choice, make the right selection. You also find it when you get approvals or escalations, because many companies, they don't want people who are executing on these business processes -- their front-line staff, the first person who answered the phone, the person in the store -- they don't want those people making decisions. And so as soon as you get to a point where they have to make a decision, up the line you go. You get referred up, it gets escalated to level two, it gets promoted to my manager. Only my manager can do that, because that decision-making, instead of being explicit, instead of being managed, they just said, "Oh, that's too hard for this person to do, so we'll just refer it to somebody else, and we'll keep referring you until we find someone who's able to make this decision."

But that decision-making could have been embedded, and it could have been embedded if you had a better handle on the business rules that make up your decision-making process, because at the end of the day, your decisions are driven by business rules. That is the business rules that you have to enforce to make that decision.

If you think about where those rules might come from, they come from all sorts of places. They come from regulations that determine how you can make certain kinds of decisions. They come from policies that likewise determine how you make different decisions, though they're driven by your internal politics and procedures, not by external regulations.

But they also come from experience. What do your experienced staff know works and doesn't work, how do the best people who interact with customers interact with those customers? How do they choose what offer to make? Can you encapsulate that knowledge somehow in business rules?

They come out of your legacy applications. Most organizations are stitching together a whole range of applications with their business processes, and some of those are legacy applications. And within those legacy applications are some of the rules you need to make these decisions.

More interestingly for most groups, and one they forget most often, perhaps, is that the history, the data you have, is also a source of rules. If you use something like SPSS Modeler, you can go mine that data to see what rules you should have been using or what rules worked in the past, because there is a rich set of business rules that's hidden in the past behavior of your organization, the past successes and failures.

Taken together, all these different rules can be brought together to make decisions. And it is that ability to bring all those rules together and manage those decisions, manage the rules behind the decisions, that makes the business rules management system a powerful tool for you.

Now, obviously, there's lots of decisions in your business. And even if you start focusing in, as you should, on the operational, day-to-day business processes that drive your business, not every decision is equally suitable. Because you're obviously making a technology investment here, you're making a decision about how you're going to implement this.

So when you find decision points that have logic with a certain set of criteria, it's really important to be able to manage the business rules so those decisions are repeatable. Now, clearly, if decisions aren't repeatable, if there isn't a common thread of logic that you follow each time, it's very hard to encapsulate that decision logic in something clear, because it's a very ad-hoc, very changeable process.

So it's got to be repeatable. It's got to be -- well, perhaps it's based on lots of regulation, lots of policy. If you have to make a decision where the procedure guide or the policies that people have to follow run to many pages, a business rules management system is going to be a great tool there, because it's going to let you manage all that volume of regulation.

Similarly, if they change a lot -- it's one thing to train a bunch of people in your call center to make decisions according to a set of policies, if those policies are stable, but what happens when those policies change all the time, or when court cases or regulations or other impact drives constant change in your decision-making? Well, you're going to have to be able to manage those business rules at that point if you're going to keep that decision-making up to date.

Obviously, if the decision-making itself requires a degree of domain knowledge, and if you want to automate that decision -- and if it's repeatable, you'd like to -- if you're going to automate that decision but it involves a lot of domain knowledge, your IT department is not going to be able to do that on their own. They don't have that kind of business domain knowledge.

You need a vehicle for capturing the rules in a way that's driven by the business, and uses that business domain knowledge, but still capable of being automated. Again, that's where a business rules management system comes in.

Obviously, if the rules themselves are complex or they're highly interdependent, (inaudible -- technical difficulty) them in traditional environments, trying to build them into your business process becomes very difficult.

The last couple here, if you really have a decision where the business simply won't give up control, perhaps it's already built some little Visual Basic- or Excel-based thing that it's using and it says, you know what? We control this, and we don't want to lose control of this decision, we have to own this decision, then again, you're going to need an environment in which they can do that. You're going to need an environment where those

business rules can be managed by the people who understand them, and that can be a challenge for an IT department.

Then last but by no means least, if you need to embed analytic insight, if you need to calculate risk or opportunity as part of this decision, then almost always you're going to find that the rules for that become complex and need to be managed in a more effective way.

If you have any of these criteria or some of these criteria in combination, you're going to find that you have a real challenge, which is that you just can't use traditional tools for this. On the one hand you say, okay, these are repeatable decisions that are often high-volume. I have to make them quickly, so I want a high-volume system with low latency.

Well, normally if you need that kind of system you would go to the IT department and they would code this for you. But you also need a system, these decision-making components have to be easy for the business to inject their know-how. They have to be able to be changed rapidly, because decision-making is one of those things that is affected by lots of things outside the boundary of your organization.

So you don't know when you may need to change them. When you do need to change it, you often need to change it in a hurry. Then you have to be able to demonstrate compliance with a lot of these decisions. It's not enough to make the right decision. You can't tell regulators, well, trust us, we always make good decisions. You have to be able to show them that you've made decisions that were legal and appropriate.

All these different tensions tend to mean that, okay, while the volume and latency might lead me to code this in a traditional IT environment, these other kinds of issues might lead me to say I need something that's more business-friendly, something like an Excel or a Visual Basic, where I can make it easier for my business people to be engaged.

But if I do that, well, now that isn't going to support my operational environment. So when you get into these kinds of decisions that have this complexity of business rules and support your day-to-day business operations, you need to use a different kind of technology. You need a business rules management system that's going to let you bring these business rules together, document them, understand them, engage the business and IT in a collaborative way and then allow you to package up those rules and to deploy them as a decision service so that you can make the rest of your environment -- your business process, your event processing systems -- able to get the answers they need to these business questions.

Now, I wanted to illustrate some of my points here with some case studies, and these are IBM case studies that I've talked to or that I've worked with. This one was a mobile Telco in Europe, and they had a problem, which is that the decision which was what promotional offer to make people, was a very dynamic decision, and the businesspeople wanted to watch their competitors, listen to the radio in the morning, see the newspapers,

watch the TV ads, see what was happening in their very competitive Telco marketplace, and immediately they had to go change the way promotions were being managed.

They couldn't do it because the legacy IT system took too long to change, it was too hard-wired. So what they found is that they had to adopt a business rules management system that enabled them then to define the rules behind this decision. The decision's a pretty simple one -- what promotional offer do we make to this customer at this time?

But the rules behind that change all the time. There are restrictions on it, there are regulations and policies that have to be followed, there's a dynamic marketplace, there's a history about what's worked in the past and what hasn't worked in the past. All of this has to be brought together so they can make an effective promotional offer to retain their customers as they came up to renewal.

Classic example of a decision service supporting multiple channels, multiple business processes and ensuring that whenever they needed to know what promotional offer can and should I make to this person at this time to retain them, they got the right answer back.

Now, if you want to apply this kind of a -- if you have these kinds of decisions in your business process, what you need is a way to figure out how to manage them. When I tend to work with customers, it's a fairly simple three-step process that you have to follow. You have to manage decisions, you have to focus on this decision management context as well as on your business rules, so what is decision management?

Well, decision management is an approach, much like business process management is a way of thinking about how you manage business processes, not just a piece of technology -- so decision management is the same.

It's designed to focus you on these day-to-day decisions and let you get more value out of your existing environment. Like I say, it's got a fairly simple three-step process. You have to find these decisions -- I talked earlier about some of these decision-making words, some of these places in your business process where there are decisions, and that's just as true in an event processing sense.

Then you have to build decision services, you have to identify the components that you're going to build to implement these decisions that you can have all your other components have access to that decision-making. Then you need to be able to create a closed loop, you need to be able to perform ongoing decision analysis to make sure that you continue to assess that decision, see how well it works, see if there are changes you should and could be making.

When you apply this process, what you get is an effective tool both for legacy modernization, for improving your business processes and for establishing the difference between the kinds of business rules that drive decision-making and the kind of business rules that are required elsewhere in your systems portfolio.

Now, I'll walk through how that might look. Let's imagine we have a whole series of channels that treat our customers, like our telco example. We've got CRM system, we've got a call center, we've got our website, our e-mail marketing. Within all these different contexts I need to make decisions. I need to understand what offer to make to retain my customers.

So the first thing I do is I go ahead and I separate out some kind of decision service, and I say, okay, I have a question about this customer. What offer should I make them to retain them? All that decision service does is commit me to a particular action, it makes a selection from the available options, applies all the business rules that will help me choose between those different options, and pushes that back out as a commitment.

Now, okay, I could do that and I could do that with code, I guess, but what I really want to do is write business rules to manage that, so I can take all the business rules I've collected in these different places -- my legacy system, my policies, my regulations, my best practices -- and bring them together and focus them in this decision service.

Now over time, of course, I collect a lot of data about how my customers interact, and I feed that into some kind of historical data store, and with that, and perhaps some external data, I can do analytics. I can use this data, I can mine this data to come up with new business rules. What kinds of segments are there within my population, and what kinds of offers work better for different kinds of segments?

I can feed those rules into my process. I can get smarter about it, I can even perhaps start to make predictions -- which customers are actually at risk of failing to renew? Who do I predict as being likely to be a problem? I can go ahead and inject those predictive analytic models back into my decision service, so now I've got not just my rules, my explicit rules, I'm starting to build the rules that are implied by my historical data.

Then lastly but my no means least, of course, I have to create my closed loop, so I create an ongoing decision analysis process that enables me to assess how well I'm doing, am I getting better or worse? Are there alternative approaches I could be trying? I feed those also into my decision service.

But even though my decision service is getting more and more sophisticated, more and more fine-grained, I haven't had to change any of my other systems or processes, because they are all still just asking this component the same question -- what's the right offer to make this customer at this time? Because that's all they do, they're able to continue to operate without constant change, without constant investment, because they're all using the same decision service.

Now, I've got a great example of this again, another IBM customer in Europe, and this was a retailer. They had a problem, which is the sheer number of rules to do with discounts. So in their particular case, they had ads they ran in the newspaper, they had coupons, they had loyalty offers, they had radio ads that had a magic word, and if you

said the secret word to the person in the store you got another discount. They had in-store pricing discounts.

They had all these different ways to affect the price of goods you bought, and what they found is that the people in the stores simply were getting this wrong. There were too many rules to apply to get the price right for customers in the store. So what they did is they pushed down into their point-of-sale terminals a decision service, a business rules-based decision service, that said for this customer with this loyalty card, given the offers they've presented and what we're running in the store and what they've bought and the size of their basket and all these other factors, here's what it's going to cost them and here's what their discounts are.

Oh, and by the way, here's a second question we're going to ask at the same time, which is what's the best loyalty offer to make this person right now to get them to come back into the store and buy more stuff tomorrow, next week, next month?

Two critical decisions -- how much does this order cost this customer right now, and what should we offer this customer to get him to come back another day? Those two decisions were embedded into this point-of-sale terminal, into these work processes that are executed in the store and guarantee that no matter if it was your first day on the job, you could still get both those questions right from the point of view of your customer.

A tremendous example of empowering people at the front of your -- the very fingertips of your organization to be as smart about these decisions as anybody else in the organization.

Now, I've talked a lot about decisions, so one of the things that focusing on decisions does is really simplify business processes. That's a really important factor, because if you think about a business process, and this is again true of event processing, too, complexity is your enemy. The more complex your business process, the more difficult this is going to be.

If your business process is very complex, has lots of steps and branches and links, it's hard to make sure it's right. So it's very hard to assure your process quality. It's difficult to provide good customer service, because frankly, customers don't like long-running business processes. They prefer quick and short and simple ones.

It's also hard to find places to apply analytics and apply learning. If you've got a bunch of data about your customers and about your processes and you have this long, complicated process, how do you learn what works better in that process? It's very hard to analyze the behavior of that process if it's too complicated.

It's also much more expensive to run more complicated processes, just from a point of view of both executing a complex process, but particularly in terms of being able to maintain any complicated process and make changes to it. It's hard to make those changes safely. So that tends to mean that changes to a complex process have more of a process

wrapped around them themselves. They go through more testing and more quality assurance cycles, and that reduces your agility. It just takes you longer to change a more complicated process.

But as you think about the decision-making in a process, you can often dramatically change this dynamic. If you think about, within a process, you might have simple decisions like should we approve or decline this application? We talked about that earlier, so a choice. Or we might be combining many factors and coming up with prices or terms or conditions.

It doesn't matter which kind of these decisions we're talking about. Most processes have both. Once I start focusing on those decisions, I can often dramatically simplify the complexity of my process.

Let's take a quick example. Here's an example without decision-making. This is a business process for handling applicants for I think life insurance, let's say. Well, okay, I have a process and I know I'm going to end up with three different ways of assessing people because I've got people who are very low-risk and I've got people who are very high-risk and people in the middle, and I've got these different processes that I want to follow.

How do I decide which process to follow? Well, okay, I'll put an age branch. So I put my little age branch in, if you're young, old, middle-aged. Well, and if you're middle-aged, I'm going to start saying, okay, well, how good is your medical record, because if you've got a good one I'll treat you one way, and a bad one I'll treat you a different way.

I start to build in these different factors. Even with just a few, I've already started to get a little bit of a branching thing going here with my process, and you can easily see if I add a couple more criteria that I'm going to start to end up with a pretty messy process.

Certainly those of us who work with companies doing business process find these kind of things all the time -- a company that has a standard (inaudible -- technical difficulty) cash process with 400 country exceptions built into it, or an approval process with all these different branches so you can tell which validation routine or which approach to take to validate and approve something.

But what if I took my decision out of this? What if I looked at this branch and said, well, actually, I'm making a decision here. I'm trying to decide how risky you are. What if I had a component that would answer that question for me? What if I had a component that would let me say what type of applicant is this? Are they a high risk, a medium risk or a low risk?

Well, now I have a much simpler process, because obviously I just do that and then I run whichever process is appropriate. But more than that, I'm protecting myself against change. I've put myself into a position to go ahead and change the logic behind that approach again and again and again, without impacting my business process.

I got a great illustration recently of this in a -- I was working with an IBM client, a telco, in their call center. And they had a number of processes where they had very complex processes and complex procedure guides for their call center reps, and we started to look at these and we said, well, actually, you've got a decision in here. You've got a decision -- let's take an example.

Someone's calling up, they're going to travel, and they want to know if their phone's going to work in this foreign country, if they can -- what do they have to pay for a different service, what the dialing instructions are going to be, will their phone even work, all those kinds of questions.

To gather the data from the where are you going, what kind of services are you looking to have in those countries, and then make a decision as to what to tell this person -- do they need to upgrade their phone, upgrade their plan, is there some temporary plan we can put them on, what are the dialing instructions they're going to need to call to and from these different countries, and package all of that information up to answer that question, what do we need to tell this person so they can travel successfully to this country, and then tell that to the person on the phone.

Then there may be other processes you have to follow. You may have to sign them up for these new services and so on, but this process becomes very, very simple because I have this complex decision and then I simply pass that information on to my customer.

Every time one of these countries changes its dialing things, or every time I add a new partnership with a new service provider in one of these countries, I don't have to change my business process. I just go and I add new rules about that particular country or that particular product. So I'm able to isolate my change inside this decision-making component instead of having it clutter up and complicate my process.

This is important, because if you actually look at processes and decisions separately, when I look at all the different processes my companies I work with have used, they fall into three big categories. The first and in many ways the least common category is one in which the process and decision change in lockstep.

Every time I change the decision I have a change in the process to make, and vice-versa. Those are not terribly common. When you see them, clearly, there's not much advantage from an agility point of view of being able to change the decision separately from the process, because they're both frankly changed at the same time. But what's much more common is to have a process that doesn't change all that often, but where the decision-making logic changes all the time.

If I think about my claims processing example, I change the logic for detecting fraud all the time. I come up with new alerts, I get new advice from the FBI on fraud schemes, I find new analytic reasons to worry about claims. I continually assess how well I do at detecting fraud. But I don't change my process for handling that all that often. Perhaps if I

start working with different outsourcers or I start working differently with my claims agents, I might have some process change over time, but I change my decision constantly.

That's a much more common scenario. But when you start getting into the core operations of a business, the most common scenario is actually that you have a really stable business process. Your order-to-cash process doesn't change very often, but the decision-making to give out customer discounts, shipping route, shipper selection, those kinds of decisions that go on inside that order-to-cash process, those change often. They change all the time.

They're a very dynamic part of your process that, if you separate it out, then you're able to make those changes and keep your process stable. This is a really important aspect of decision management and of business rules that people underestimate. By separating out these decisions, not only do you simplify your process management or your event management, you also create this opportunity to have very different paces of change and different drivers of change between these different elements. You have separated these concerns and given yourself more options.

Now, I only have a couple of minutes left and I wanted to touch briefly on analytics, because there's a lot of interest in analytics out there right now. I've been to a number of events with IBM customers, and there's a tremendous amount of energy around analytics. How do we apply analytics? We have to think about what analytics means, and analytics doesn't just mean reporting or a data warehouse. Analytics is about simplifying data to amplify its meaning. It means taking all this data you have and using mathematical techniques, typically, to turn that data into something more meaningful.

So instead of having a bunch of data about which customers have or have not accepted a retention offer, I turn it into some kind of prediction that says this is the most likely offer to be successful for this customer. Instead of having a whole bunch of information about whether our fraud investigations were successful or not, whether we found out it did find fraud, it didn't find fraud, I come up with rules to help me classify high-risk claims.

If I apply these analytic models, if I think about data and analytics in this way, in order to really take advantage of that I need a place to apply that meaning. I think you've got to use that meaning somehow. What I actually need to do is I need to apply that meaning to make better decisions, because if rules constrain what might happen in your environment, if you're going to apply analytics in that context, you can understand why they're happening and when they might happen, and make, therefore, better decisions, but only if you have a place to apply that within your business processes, within your event processes.

So the other thing that comes out of separating out decision-making components, decision services and automating them using something like a business rules management system is that suddenly you have a place to go put all this insight that you're generating from your analytics.

Again, a great example of this, a customer I was with at an IBM event recently focusing on this claims idea. So what they had found is that they had got a process for fast-tracking or referring claims to an adjuster or referring them to a fraud investigation unit, and they had a whole bunch of rules for these things.

What they found is that when they started to apply analytic tools and actually mining their historical data, looking for patterns in that data, looking for fraud patterns, looking for patterns of unusual behavior, and then also looking for the patterns of people who are legitimate, they were able to dramatically improve this process.

They were able to find far more claims they could safely fast-track, they were able to refer a much higher percentage of the fraudulent claims for fraud investigation. They were able to apply this analytic insight, this meaning they got from their data, because they had a point in their process where they were making this decision. By applying that insight they dramatically improved the quality of that decision.

One of the things that often comes up in this context is we're not talking here about decisions that are macro decisions, we're talking about micro decisions. We're talking about a decision that applies to a single customer or a single product or a single order. Think about a cross-sell offer. Think about our first example with the retention offers that the telco is making.

When I make that offer you could say I've made a few decisions. I've decided what offer to make, how to phrase the marketing spiel for it, what price, what discount to give and so on. I've made a handful of decisions. But I've also made a decision each time I offer that to someone. I've made a decision to offer James this product at this price. That's what a micro decision is. It's that idea that each time you make this decision, you're going to count that.

So when you start thinking about how repeatable, how many times do I make this decision, how many decisions do I make, if you have a process that handles, as in the claims example, 700 claims a day, you're making 700 decisions a day, thousands of decisions a year, about claims.

If you can focus using business rules and analytics on making each of those decisions uniquely to that customer, that claim, that transaction, you can dramatically improve the quality of that decision-making and the precision of that decision-making. You can make them thoughtfully and explicitly.

I wanted to finish up with one last case study to illustrate exactly that. This was a European retailer, and what they wanted to do is they wanted to build loyalty to their brand no matter which store format you went to or whether you shopped online. What they had found prior to this project is that customers were loyal to their supermarkets or their hypermarkets or their corner stores or their website, but they weren't loyal to the brand as a whole.

They were using these different channels and they were pretty loyal to those channels, but they were not building a cross-channel loyalty that they really wanted. So they pulled all this data together, they used analytics to amplify the meaning of that data, and in particular, to apply to the context of a single customer as they were checking out.

They were able to deliver compelling cross-sell and loyalty offers as a check-out process of each of these different formats that turned people from being loyal to a single format and gave them reasons to start being loyal to the rest of the formats, the rest of the channels.

Again, a very successful project focusing not only on a decision point but on a micro decision, on this customer-by-customer decision. Instead of printing generic offers or standard ways to try and cross-sell between these different channels, focusing on what they know about each customer, using analytics to take all that data about the customer and focusing it, and then using rules to make sure that the right offer gets deployed.

So a quick wrap-up for me before I hand you over to Brett. Why manage decisions independently? You can make changes to the decision logic, the business rules behind these decisions, quicker, easier, faster and more reliably. You can coordinate those decisions more cleanly across products and channels. All your systems, all your processes can use that same decision service, get the same kind of answers and deliver that consistently back through all your different channels.

You get simpler processes that are easier to manage. It's easier to make changes for the process clearer, whether it's a process change or a decision change -- a much simpler process for you to go through.

You can focus your resources, your people, on decisions that are more complex, on the exceptions, on customer interaction, because your systems handle the day-to-day decision-making that drives your business.

You have a place now to take your analytic insight, and instead of just saying, "well, now we know vague things about our customer base, isn't that lovely," you can say "I know something about this kind of customer. This person, right now, is that kind of customer. I can use that knowledge now at this decision point to improve the quality of my business."

I create an opportunity now to continuously improve my decision-making separately from any process improvements, separately from any system improvements that I might make. I can focus just on my decision-making and on how good it is, and that creates an opportunity for ongoing improvement.

So final recommendations, then -- three things. Think about your decisions as well as your processes. Think about your decisions as well as your events. Make decisions one of the objects that you think about when you're designing systems. Adopt business rules management technology. Think about a business rules management system, because once you start focusing on decisions, you will need to manage the rules behind those decisions.

Then focus your analytic investments on these kinds of decisions. Find the decisions where you have history, where you have experience, where there is data and say, "how can I use this data to make a better decision at this point in this process?" And focus analytic investment, therefore, on better decision-making.

With that, I'm going to hand over to Brett.

Brett Stineman: Thanks a lot, James. That was very insightful.

What I'm going to do at this point is to go through a demonstration that will reinforce a couple of the points that James talked about, and really focus in on this idea of how to maximize decision agility through the separation of business rules and those specific, very essential decision points within a process from the overall management of the process itself.

To show you this, what I'm going to do first here is I'm going to show you an interface which is the WebSphere Lombardi Edition product, and this is their modeling environment.

Within this environment, what you can see is the full end-to-end process for an application that's focused around underwriting for automobile insurance. So obviously, we have people who are going to be applying for insurance, they may be coming in through various channels. Let's assume in this one that they're coming in through a call center channel, where various information is going to be collected about the applicant.

Then at that point we are going to have a decision point, as James talked about, so we're going to have an eligibility determination that's going to happen. That eligibility determination is either going to push the applicant through directly to pricing and from there to a notification of the applicant, or it's going to bring in a person, where we need to have a manual underwriter involved to take a look at that particular applicant's situation and their information.

Now, what we have here actually reinforces one of the points that James talked about, which was the idea that separation of decisions from other parts of the process really helps to simplify the overall process. This particular implementation we see here gives us a very nice, simple view of the overall end-to-end process, which is great.

Now, at the same time, it does one other thing, it makes that process much more flexible. So I'm going to be able to show you this in the way that we can implement changes to a decision very easily and push it out and make a change to the process without actually having to change the overall process itself.

This process model also shows a key distinction between those essential decision points, such as eligibility and pricing, from other types of rules that James talked about that really make sense to keep within the process. So if you look at the middle, where you see

that red box, which is the manual underwriter, underneath that there's a yellow box, which is an escalation rule.

So if a particular applicant is pushed into the manual queue and nobody looks at it within a certain amount of time, I want to escalate that and make sure that that isn't just sitting there and it gets followed up on. So this is a great way to have a rule within a process that's pretty static.

I know that if nobody's looked at that within six hours or a day, whatever you determine, then I need to escalate it and who it needs to get escalated to. That's great place to keep a rule within a process.

At this point, I'm going to launch the demo. I would like to let people know that if they're having any problems viewing this to hit their F5 key -- again the F5 key. That will refresh their browser window and in case there's any problems, that should clear that up.

The other thing I'd like to let you know when this demo runs, the resolution may not be optimal. This is some of the limitations of the tool that we're using, so I apologize for that. But it should give you a good idea of what we're talking about, and this idea of maximizing agility through the separation of rules and process.

Let's get this kicked off. What you're going to see first is, using the interface, I can actually simulate the process. I can actually start to drive through the various steps. And we're going to start by entering some sample information about a prospective customer, and we're going to collect some information.

We can actually see the various interfaces that a call center agent would have, where they're collecting information about the customer, about their vehicle, about different types of coverage that they have, and then they're going to go ahead and submit that.

We can go back into the process and we can actually step through it and see exactly what's going to happen. So we can step through and we can see for this particular set of data that we've collected that it's saying it's a manual underwriting that needs to happen.

From there, we can actually also see the various actors that are involved in this process. We can go ahead and see who's involved here from a manual underwriting standpoint and exactly what they would see. So we can see their interface and what you can tell from here is we get a nice snapshot of all the information about that particular prospective customer, and we can see why it was passed for manual review.

There's a specific rule, and that rule passed back some information, so this underwriter says okay, based on this we'll go ahead and approve it and push it through. That will move us now to the next step, and we can step through again and see what happens next. So based on that approval we're going to move to a pricing step, and from there we can actually see what the pricing is.

This brings us to another point about the power of separating out decisions. When you see what the pricing is for this particular customer, you'll see the number of different elements involved in pricing out this particular insurance policy. Based on all those various factors of where the client -- where the customer lives, the type of vehicle, the types of coverage, the various types of adjustments that this company has defined for those various types of data that are coming in.

So you can see here this is a great place to use business rules separated out from the process, because we can get very granular in terms of the types of pricing decisions that we can make.

Now that we've looked at that, let's look at this idea of modifying the eligibility decision that takes place after we get a submission from a prospective customer. We're going to go into an environment which is the business rules management system, and specifically what we call the rule team server. And this is where we can actually have business users who are policy managers directly implement and manage changes to the various business rules for which they're responsible.

So we have a policy manager who's focused on eligibility. He can browse through the various types of business rules associated with eligibility. He knows there's a specific rule set that's a decision table that he needs to look at and he needs to modify, and this is what's called the driver profile.

Now, looking through this, he can see in this particular table there's a set of rules that are going to determine whether we're going to just make an applicant ineligible or directly go to pricing, or whether we're going to need some type of manual review. What has been done through some business analysis is we've determined that there's some information we're collecting on the front end that's not being actually used in this particular rule set, and that information is around whether the applicant has completed a driver's education course.

What you can see here is we can easily add that element into our decision that we're making. What we're going to do, and this is going to go through very fast, is we're going to add that element about the driver's education course. And in doing that we're going to have more straight-through processing that can occur based on that change to the rule.

Once we've actually made that change to the rule our eligibility manager can actually simulate the effect of that change by running it back through some historical data. This is a simulation that we can run, and based on that simulation, what he's going to see is that the amount of straight-through processing versus manual assessment that need to occur, and what you can see here is that manual underwriting is going to comprise about 5% of the total number of cases that come through, or of applications that come through.

Now, even more powerful is he can now run a comparison of that new rule set against the previous version of the rule set against the same set of data. And what you can see here is whereas before we had manual underwriting at around 30-some percent, we've brought

that down all the way to 5%, and that's pretty powerful, the amount of straight-through processing that we can do. It's not going to be every single case, but whenever possible, based on our organizational knowledge, we want to be able to do that.

At this point, we're going to go ahead and deploy that new rule set out into the production environment. And once that's deployed, we can actually go back to the process, run it again without making any changes to that process, and actually see the effect of this new rule set.

So we see the deployment has succeeded, and now if we go back into our process view we can once again step through that process and see the result of that rule change. So going back into our process model. We're going to do what's called playback, so we're going to start from the beginning, use that same set of applicant data. And as we run through that you'll see this driver's education course that has been specified for this particular person, Cindy Davis, and you can see where the mouse is. It says, "Have you taken a driver's education course?"

Yes. Okay, we'll step through that, submit the rest of the data, and then we'll step through the process again and see how this has changed from our first running of that particular set of data.

At this point what we can see is instead of going to the manual review we've gone straight through to the pricing, and we can step through that step again, and we can see that same pricing that we saw before, which has all those various elements that need to be considered in coming up with a total price.

So I realize that was pretty fast, but hopefully that showed you some of the powerful nature of this separation that James has talked about, and why it really makes sense. To just give you a couple more points on this, again, if you're having any problems seeing the Web interface at this point, just hit F5. That will refresh your browser and you should be able to see the slide.

So really four elements that you need to think about as to why this separation is really critical for long-term process improvement initiatives. First off, you have different roles who are involved in process and rules management. You have teams that are really focused on that overall efficiency of the end-to-end process, and you've got people who are policy managers. They may be pricing managers or eligibility managers, or they may be focused on risk and compliance issues. Each role can focus on the area that their expertise is in and that they're really tasked to help improve the overall organization's way that they run their business.

Secondly, the ability to deploy processing decisions independently of each other, and we see this quite often with our customers. Depending on where their pain point is, they may want to start looking at a BPM initiative versus a business rules initiative.

Thirdly, different lifecycles. That was what this demonstration really showed you. I had a decision issue that I needed to change. I didn't need to change my process. Typically, decisions change much more frequently, as you saw from one of James's slides that talked about this, and it showed the two lines on top of each other. I'm going to have my decision changes happening much more frequently than my overall process change, and I can manage those, the lifecycles, as well as the governance requirements, independently of each other through this.

Then lastly, the idea of reuse. I may have an eligibility decision within this insurance organization that affects many different types of processes. I might have recreational vehicles, motorcycles, specialty vehicles, and maybe that eligibility determination may apply to some or all of those, or my pricing decisions, right? So I can reuse those business rules across multiple processes very easily.

To conclude, just to hit a couple points, when we talk about decision agility and effectiveness there are a couple of key technologies you need to think about, and each one has its particular strengths.

First off, business event processing. We didn't spend a whole lot of time talking about it here today, but this is a way of really improving situational awareness and response of data that's flowing in through the organization and through the various systems.

Secondly, business rules management, which is what we looked at here in this demonstration, to really improve the quality of automated decisions. Next, business process management, which is going to use the outputs of business events and business rules within the orchestrated processes that it's responsible for running.

Then lastly, analytics that James talked about at the end of his session. That's really focused on how do we continuously improve those decisions that we want to be able to automate within our various systems.

We have a couple of minutes left, and I'm going to turn it over to John for a few final items that we need to talk about before we get to the Q&A. John?

John Sote: Thanks, Brett. Before we begin with today's Q&A, please fill out the feedback form that has opened on your computer. To complete the form, please press the Submit Answer button at the bottom of the page. Thank you in advance for filling out the feedback form. Your participation in this survey helps us improve future webcasts.

Now, on to the question-and-answer portion of our event. As a reminder, to participate in the Q&A just type your question into the text box located below the media player; then click the Submit Question button.

All right, let's see. First question we have is for you, Brett, and it is "please elaborate on the difference between agile and waterfall project management methodologies, and their impact on being able to effectively define decisions and rules."

Brett Stineman: Yes, I'll touch on this very quickly. This goes back again to something James had talked about. The separation of decisions as their own element is -- agile methodology is very effective for handling those. And maybe for your overall business system adding new functionality that you need to over time, overall improvement of that business system, it may be good for a waterfall type of methodology. But by treating decisions as their own enterprise asset, it makes it much easier to implement changes. And we encourage an agile approach -- agile methodology approach to doing that.

John Sote: Okay, great. Thanks Brett. Next one is for James. How do you advise organizations on determining when rules should be managed separately?

James Taylor: Well, that's a great question, and I think the key thing is this focus on decisions. What I find is that that set of criteria I gave, rules that change a lot, rules that are complex or interact in complex ways or whether it's just a lot of rules, finding those kinds of decisions, that's where you really get the bang for your buck in applying business rules management systems, applying that separate management of business rules.

There's lots of different ways to approach it that all seem to work. But at the end of the day, if you don't really focus on the decisions that you're going to be automating with these business rules, that is decisions that the rules are for, it's very hard to effectively manage those rules, and it's hard to tell when and where to manage them. I would always start with the decisions that are important to your day-to-day business, in particular.

John Sote: Okay, great. Next one, James, for you also. Should monitoring rules be separated from processes or applications?

James Taylor: Yeah, that's back to that very first slide we had. It's this context that you have where you say "I've got rules about how I'm monitoring existing systems, I've got rules for monitoring events, I've got rules for monitoring processes, I've got routing rules within my process."

My sense is, for most of these rules, that managing them as part of the environment they're monitoring is perfectly acceptable. You're very tightly coupled with the current configuration of your systems and your processes, and they're not really there for business decision-making.

It's one of those things where you talk about it in general, it often sounds like a very finely nuanced distinction. But when you're actually working through it, typically, it's pretty clear what the business decisions are and what's much more about monitoring the current technical environment.

Because we're talking here about an IT context, and so the rules that are about monitoring are about monitoring systems and business processes and business events that are captured in your information systems.

So those things change in ways that aren't always driven by your business, so I would tend to keep those separate and just manage the real business decision-making as decision-making components, as decision services.

John Sote: Okay. How about on more -- 30 more seconds here. James, can analytics run directly within decision rules?

James Taylor: Now, that's a really good question. There are two ways you can use analytics to improve decision-making in the context we're talking about here today. The first is to mine your data to see what your rules ought to be. And you can use soft techniques like just general population analysis and distribution of your customers or your transactions, and you can use data mining, looking for things like association rules or decision trees, where you're very explicitly using mathematical techniques to find new rules.

But the other thing you can do is take more predictive models, things like an additive score card, the kind of thing used for a risk score, and embed that using rule metaphors that support describing a scorecard, because under the covers, most of those analytic models can be represented as a set of business rules. So you can certainly do that, and you can take advantage of new technology -- things like SPSS decision management, where you've got a deployment infrastructure for your models and easily integrate those right into the rule flow, the decision flow, the steps within your decision-making.

So yes, there's lots of ways to bring analytic models directly to bear on these kinds of decision services once you've isolated them, once you've got them set up as a piece of -- as Brett was saying earlier, a piece of infrastructure that you're clearly and explicitly defining and managing.

John Sote: Okay, thanks, James, and thanks, Brett -- very good presentations.

Thank everyone for attending today's webcast, What, Exactly, Do You Mean by Business Rules, brought to you by Information Week and IBM. For more information about today's webcast, please visit any of the resource links open before you.

For those who submitted questions, we'll be getting back to you with a response within the next few days, so please don't feel like we're leaving you out.

All right, next one for decision management solutions, and then finally within the next 24 hours you will receive a personalized follow-up e-mail with details and a link to today's presentation on demand. Additionally, you can view today's event on demand by visiting www.NetSeminar.com.

This webcast is copyright 2010 by United Business Media LLC. The presentation materials are owned by or copyrighted, if that is the case, by Information Week and IBM,

who are solely responsible for their content, and the individual speakers are solely responsible for their content and their opinions.

On behalf of our guests, James Taylor and Brett Stineman, I'm John Sote. Thanks for your time, and have a great day.