

Event ID: 346200
Best Practices for Designing Decision Management Architectures

Brett Stineman: Good morning, good afternoon, or good evening, depending on where you are in the world. Welcome to today's webcast -- Expert Panel, Best Practices for Designing Decision Management Architectures brought to you by Information Week, IBM, and broadcast by United Business Media Limited.

This is Part 3 of 4 of the IBM Good Decision webcast series, Architects and Developers Crash Course in Decision Management. I'm Brett Stineman. I'll be your Moderator today, and we have just a few announcements before we begin.

You can participate in the Q&A session by asking questions at any time during this webcast. Just type your questions into the Ask a Question text area below the video window and then click the Submit button. At this time we recommend you disable any pop-up blockers. The slides will advance automatically throughout the event. You may also download a copy of the slides by clicking on the Download Slide button located below the presentation window.

If you're experiencing any technical problems, please visit the Tech Webcast Help Guide by clicking on the Help link below the media player. In addition, you can contact our Technical Support Helpline, which is also located in the webcast Help Guide.

Now on to the presentation, Expert Panel Best Practices for Designing Decision Management Architectures.

Joining me today for our discussion are three extremely knowledgeable and experienced Technologists from IBM. We have Steve Demuth. He's a Senior Technology Staff Member for Decision Management. Steve is a member of the WebSphere Business Performance Optimization Team. He also leads the WebSphere Business Events in Rules, Products, Architecture Team, and is responsible for WebSphere's technical strategy for Decision Management, which combines rules, events, and predictive analytics.

As past Director of Product Management and Architecture for ILog, which IBM acquired in 2009, Steve led the creation of the ILOG BRMS product suite from 2005 to 2010, as well as the integration of the ILOG Business Rule Management System with WebSphere Business Events to create a new product that came out in 2010 called WebSphere Decision Server. He's a recognized expert in business rules technologies in the emerging decision management discipline.

Prior to his work at IBM and ILOG, Steve was Chief Technology Officer for FIDIC Consulting Firm that specialized in business rule applications where he consulted on architecture and design for a number of Enterprise Business Rule projects.

Next, we have Steve Lyons. He's a Senior Technology Staff Member for Business Events at IBM. Steve is currently the lead architect for the events processing capabilities of IBM's WebSphere's Decision Server offering. Previously, he was Vice President of Engineering at

AptSoft Corporation in 2008. And prior to AptSoft, he was CTO at Parallel DB Inc., a real-time database replication company, whose products were designed to enhance disaster recovery and information availability.

Before that Steve was responsible for architecture and deployment at OrderTrust LoyalNet application, a massive transaction processing engine, as well as a Senior Technologist in the Advanced Products Group at BroadVision. He's also been -- worked on development for Lotus Domino and led the team at ported Domino to the Linux platform.

Our third Panelist today is Jean Pommier, distinguished Engineer and CTO of WebSphere Software Services at IBM. Jean has been in charge of service engineering, which covers implementation methods, best practices and consulting offerings for IBM's WebSphere Software Services Team.

Jean joined IBM from the ILOG acquisition, where he was ILOG's VP of Methodology. And he was at ILOG from its creation in 1987 through the acquisition by IBM in 2009. He's been involved in R&D for rule-based systems, and he moved into consulting and then management. And he's led ILOG's Worldwide Professional Services Group. Overall, Jean has contributed to more than 400 successful customer implementations, decision support systems, and decision management systems, helping customers across a wide range of industries to improve the flexibility of their operational system, and the overall agility in responding to change.

Now before we begin with the discussion with these great panelists, I'd like to recap quickly what was covered in the first two sessions of this series. And here, in this next slide. In session one, which was titled, What do we mean by Decision Management, and What Does it Mean for Business?, we heard from IBC Technology Analyst, Steve Hendrick, who explained how organizations are shifting from the art of the decision to the science of the decision using technology to provide greater support at decision making as a management plan.

Steve discussed three types of decisions within organizations, strategic, operational and tactical, and he looked at different technologies used to help organizations improve how they handle each type. Steve described the decision lifecycle, which you see here, and the essential associated technologies with that lifecycle. And as you can see he called out three key technologies that bring business value by increasing automation and efficiency in day-to-day operations. Those are the items that you see listed in red.

Also in session one, we heard from Bill Looney at Horizon Healthcare Services. Bill shared how Horizon is using decision and project management technologies not only to drive efficiencies through automation, but also to bring improved transparency and adaptability to their business systems. In addition, Horizon Healthcare is using these technologies to improve decision support, enabling intelligent interaction with their brokers, customers, and medical providers through a customized guided user experience.

In session 2 which was called Business Process Optimization, Achieving Better Outcomes with BPM and Decision Management, Pierre-Henri Clouin and Asit Dan from IBM looked at how decision and process management technologies are used to improve both front office and back

office interaction. They also looked at a phased approach implementing these technologies and the patterns of implementation at each phase. Now all registered attendees can display archived versions of these presentations by using the link that is provided in the confirmation email that they received.

So now let's move on to why we're all here. Now, I've already told you about the panelists, and so let's get right into our questions. We have a set of topics that we're planning on covering today, but we'd also like to hear from audience members on topics and questions that relate to business -- to decision management. As I described before, you can participate in this discussion by typing your questions into the Questions text box located below the media player, then click Submit Question -- or the Submit Question button.

So now let's get started with the panel discussion. And to start let's -- I'd like you to direct these questions to Steve Demuth. Now Steve, based on what we've heard in the first two sessions, can you summarize the role of managing operational decisions and the various decisioning touch-point across an organization?

Steve Demuth: Sure. When we talk about managing operational decisions, what we're really looking at is the kind of decisions that your organization needs to make repeatedly, as in many times a day, and repeatably, as in according to some specified methodology or some specified set of rules or goals that govern the decision. And the repeatability is sometimes determined because of audit requirements, and sometimes just for the purpose of consistency.

But again, when we're thinking operational decisions, we're thinking repeated, and repeatable. And these kinds of decisions live everywhere in the modern organization, and I'll give you just a couple of examples to sort of get your head thinking about this.

If you're in the insurance industry, you need to make underwriting decisions. Or if you're a mortgage provider, you need to make loan underwriting decisions. And you need to be able to make these decisions quickly, reliably, predictably. And you need to make them repeatedly many times a day and repeatably. And most importantly, they are in the heart of your business. They are the thing that you absolutely have to do well in order to be successful.

Let's take another completely different example. If you are medical care provider, a hospital for example, and you need to make a decision on all of the patients in your ICU. Every hour you need to make the decision as to whether a particular patient should be moved to a Med Surg ward, you don't want to make those decisions in an ad hoc sort of fashion. You want to them based on the best possible prediction of what is the best outcome for the patient. And that's a decision that can at least be made at the recommendation level using the kinds of technologies we're talking about.

So the interesting characteristic of these kinds of decisions is they live inside your core business processes and they enable those core businesses to work more intelligently. So again, you go back to the loan underwriting example. You have a loan origination process, and there're some specified point in that process at which you need to make an underwriting decision. Perhaps you have a retail website and someone is about to checkout at your cart on your retail website. That's

a point which you need to make a decision right now as to what would be the best possible offer to make to this customer to induce them to increase the size of their shopping cart, for example. Could you offer them a discount? Could you offer them a coupon of some sort for a larger purchase order? Again, that's a decision you need to make there in the context of that checkout process.

So when we look at these kinds of decisions, what we're really talking about is, how do I characterize those operational decisions? Understand my goals for the decision. Understand the policy of the business rule that governs those decisions. And also understand the goals for the decision. What makes good decisions versus bad decisions? And bring to bear a combination of technologies, business rule technology, predictive analytics, and business event process and technologies to make sure that each time I make one of these decisions, I make it the best possible decision for that particular case, that particular loan applicant, that particular patient, that particular customer. I want to make the decision that leads to the best next course of action within the context of that business process. Brett? Does that help?

Brett Stineman: Okay. Yes, that's great. Thanks a lot, Steve. Now Jean, one thing we've heard throughout these sessions, and including what Steve just talked about, is that decisions and processes tend to interact frequently. So something I hear quite often is that many people tend to think that everything in an organization -- or that an organization does can be defined as a process. So can you explain when and why decisions should be managed separately from processes?

Jean Pommier: Yes. I can give you -- the decisions are always made in a context of a business process, as part of a business process, meaning whenever the decision is made by a machine or a person, we're always interacting with others. So it's eventual to -- and for many organizations, to start with a process modeling, a process management. But what we have seen over the years is that some organizations and many of our clients actually, have seen actually faster improvement of their business performance through the automation of decisions, and not necessarily the automation or instrumentation of all the processes. And it takes some organizations years before gathering all the processes and modeling events.

So what we are seeing is, for instance, and if you cannot make some manual processing with a rule-based decision service, it is not efficient to spend time modeling this part of the process in details, especially the logic or the policy or the regulation that keeps changing.

So on this chart here you see an example, of course, a process. And it's a claim management process. And what's important is to identify a few tasks which are three better implemented as a decision point. So this has been our experience in automating some tasks like that, meaning claim processing, underwriting and -- in our clients.

Now when we are modeling processes, there are ways to actually capture rules. And we call them -- we want to capture only the process and how to -- the writing of rules and at least how to process rules. In our latest release of DPM platform we even provide users with the same English-like syntax and same editor as in our Decision Server product to make this capture of rules more consistent and easier to occur. But if you have too many of these rules in your

process definition, if you -- if rules keep changing, if they are sort of complex condition, that's definitely something you want to externalize into a decision point or decision task. So that's really is a key thing.

And now, conversely, if you have many decisions, many independent decisions, it's important to also be able to orchestrate them and to call them in the right decision and the right time. So that's kind of the risk and interplay and intercalation between the DPM and decision management platform -- that people have to go to.

One more thing, by the way, in many organizations, actually the people managing or handling or responsible for processes are different from the ones who are the subject matter -- subject matter experts who are actually implementing the policies and making sure that the policies are properly followed.

So that's another kind of explaining, triggered to make sure that each each of these people are equipped with the right tool. In also a few organizations, we see many of these cases, so we came up with an offering to actually help organizations making the best decisions about which technology they should start with. Is it between pure decision automation, decision management or a purely process management, process integration? So we have a very light, very short discover workshop that organization execute and help organizations out. But definitely those decisions process there is strong interplay and very powerful interplay, I would say.

Steve Demuth: So Brett, this is Steve again. If I might talk just a little bit more about the process interaction here. I put up a picture here that I call the operational triangle that tries to put some very concrete context around what we're talking about when we're talking about the interaction of processes and decisions. And the idea behind this is to think about what goes on in your organization that needs to be automated from an operational point of view? And where do decisions and managed decisions fit into that? It's a very abstract diagram but you could think about this as illustrating a canonical use case for business. And it works a little bit like this.

In the upper left-hand corner, you have what we call Situational Awareness, and that's business event processor. It's knowing what is a cause for action in your business. When you have to do something. And that can be as simple as the loan application came across the desk, or as complex as something was detected in a business event processing environment that consists of a series of unsuccessful customer interactions, followed by a customer action that indicates potential defection of that customer, for example.

Or it could be as complicated as a series of readings of a train axle indicate -- a train axle monitor indicating that the axle is overheating. And some sort of -- something has to be done in order to prevent an equipment failure.

Situational Awareness triggers processes. Do you want a process to respond to this cause for action that occurred inside your business? The processes meanwhile push information back into Situational Awareness to say, there are these processes going on, and there are the specific states.

The really interesting thing happens when you get in the process. One of the first things the process often has to do is to decide what the best course of action is. So, if, to use a customer, example, customer relationship management example, you've got an alert that's triggered by the Situational Awareness that says, this is a customer who is at risk of defection.

Then the first thing you want to do in your process is to decide how important it is that you intervene with that customer. How valuable is that customer to you. Are they a premium customer that you should have somebody call in the next 24 hours to try to set things right? Or are they a customer that technically you don't really care whether they're a customer because they are a low value customer and maybe you'll send them an email or a coupon but you're not going to do anything extravagant about it. So you call a decision to say what should be my best next action?

And then things get really interesting because you want the best next action to be very cognizant of everything you know about that customer. And so you can bring in predictive models. You can have a predictive model that you created from your data warehouse that says the potential future value of this customer is \$10,000 over the course of the next 3 years. Well if that's the case, then it's worth spending the time of having someone from your call center try to get ahold of that customer and try to make sure that they do something to keep that customer happy.

On the other hand, if your predictive model says, based on the demographics and the past buying history and the past interacting history of that customer, that their potential value is \$100 over the course of the next 3 years, well you're not going to put a person onto that problem because it's too expensive. You've got to send them an email or offer them a very low-value coupon in an email and do something like that.

So a decision becomes the smarts of the process. And by dragging in this predictive model that says, I can actually predict things about this interaction that will allow me to make a smarter decision to determine which direction my process should go, I can manage that decision as a key enabler of agility and possibility within my organization.

Then the last little piece is those same predictive analytics can feed back up into the Situational Awareness to tell you what situations you should be looking for. So if you're looking at customers who defect, there's probably a pattern of activity with that customer or a pattern of non-activity with that customer, that precedes their defection. And a predictive analytics model can -- a statistical model can actually tell you what pattern of activity and non-activity to look for to be that cause of intervention, and sort of close the loop on the whole thing.

Brett Stineman: Great. You know there're some really interesting ideas around analytics, and I'm going to come back to that in a moment. But there's also some interesting ideas here about kind of how all these systems fit together in an enterprise architecture.

So I'm going to actually open up a question here to Steve Lyons. And that's, in terms of enterprise architectures, service-oriented architecture has become an established concept and framework for building and integrating these types of modular applications and systems that need to work together, and need to be able to be connected together.

Now we're hearing many people talking about an idea of event-driven architecture as being critical for establishing aware responsive systems such as what Steve was talking about, this idea of Situational Awareness. So Steve Lyons, can you explain what enterprise-driven architecture is and how it relates to SOA.

Steve Lyons: Sure, Brett, absolutely. So as Steve mentioned in the last slide, as we get into this really agile world, events really end up taking a very prominent place. And I think what we find is a lot of organizations have spent a lot of time establishing a service-oriented architecture. They had this service as established and now they're faced with this new challenge, this new architecture called an event-driven architecture, and there's lots of questions about why they -- what's new about this, and how does this relate to what I've already done and what I've already invested my time into?

And there's good news there because an event-driven architecture actually can utilize a service-oriented architecture in a lot of ways. Before we go there, let me just take a couple of minutes here to talk a little bit about what an event-driven architecture is so you can recognize one when you see one. An event-driven architecture, just as a definition, is really about an architecture that can do -- that can look at patterns, patterns in production, detection, and consumption, and then reaction to those events. So they have certain characteristics, as called out on the slide, that are very, very different than some of the traditional SOA architectures that you might be used to.

So one of the characteristics is that it has one-way communication. So, these events get fired off into the ether by a particular application and then there is no request response, it's not sitting around like in an SOA architecture where you may be waiting for a response. There is no response. The application expects nothing except that it notified the ether that this event has happened, which kind of falls into the second characteristic which is functional loose coupling.

Again, the events are happening. Everybody is notified of what events are actually going on. But there isn't any expectation that your event is going to be processed or that some system is actually even listening. Many times in an event-driven architecture, hundreds and thousands of events get fired off just in the normal course of business flow by the business systems. And most of those events are ignored by anybody. They are just fired off and they're not interesting, or they're not interesting at that particular time to any of the systems that happen to be listening, if any systems are listening at all. But the applicants aren't aware. They just fire those events off to let whoever happens to be listening knows that there's loose coupling there.

The other thing is really about real-time processing. So traditionally what we've seen is when we have these types of events, these types of transactions or whatever that occur, there's some kind of a batch processing that might happen that we process all the deposits. For example, in a financial institution at the end of a night, we kind of see what happens. And in an event-driven architecture it's all about real-time processing. So there's some kind of an event processing runtime that's sitting there listening and it's evaluated all the events that are coming in based on the time they come in, based on things that are happening outside and otherwise in the enterprise. And it's tying all these things together and then reacting to them.

The fourth thing is that they -- an event-driven architecture is very often business event-centric. So these events are not just IT events that an IT person would recognize and understand what they are. They're really things that a business person can understand. They could take a look at an event and they could see it's a new deposit event or a new customer event. And it can be tied to a context.

So there's all of the events that pertain to a certain customer, for example, or a certain customer ID. So again, back to the financial service example. Maybe all the deposits, withdrawals, change of PINs, all those things could be events, business events, that can be grouped by customer ID, and that we want to react to those as we see them or we want to do things based on the fact that a certain customer has a pattern of events happening. Or a certain branch has certain things happening in it that are out of the ordinary or that we want to take advantage of whether positive or negative.

The fourth -- the other term I can explain to you about quite a bit is this term of CEP or complex event processing. And complex event processing is no more than instead of doing layer-type processing with events and watching for events that are happening, complex event processing is really about patterns of events, and patterns of events over time. So monitoring -- if this is the fourth time that I've seen -- I think Steve was talking about in the first example -- if I see 3 or 4 times that a customer has a bad experience, and you can define what events make up a bad experience. And then you see some kind of an event within a certain time within 24 hours that indicates that customer might be withdrawing his or her money from your institution, that's an example of a pattern you might want to be looking for with CEP. You would then go out and maybe want to kick off some kind of a process to take advantage of that.

Event-driven architectures also very often have physical device connectivity. So it's not only systems that you want to be able to react to these and be a part of these eventing systems, you also want to be able to have devices, RFID, other types of devices in your organization be able to submit events. And you want to be able to look at those in the context of all the other business events that might be happening and take advantage of that.

And finally, push communication. So again, these things -- the event processing system is very frequently if you're looking at an event or an architecture, sitting and waiting for events to happen. It expects events to come to it. So you're not very often in an event-driven architecture seeing a polling model or where the event runtime is actually going out and soliciting events from different systems that it expects. Usually it's sitting and waiting and if it receives an event or if it doesn't receive event, which is sometimes just as interesting, it will then take some kind of action.

And finally, there are often ties in an event-driven architecture to some kind of a BAM tool to do business activity monitoring. Because in the end, all these things are happening, and it's very critical to be able to watch and monitor what's going on. Again, because they're business event-centric we want to do it with some kind of a BAM tool where they could go off and react to what's happening. Change the patterns on the fly of the events as you're looking for, and change the reactions that's going to happen.

So that's what an event-driven architecture is. How that ties into SOA is really very straightforward and what we find and when we go out to talk to the customers, EDA and SOA, far from being similar architectures that are repeats of one another, they're really very complementary as you go out and look at forces of events and targets of event pattern matching.

So what you're going to do once you find one of those event patterns, very frequently the best thing to do is to be able to use services that have already regenerated or processes that have already been generated. So those services can actually act as an event source or act as an event target. They could be used to process events. Once you see a certain pattern you can actually call about to a service and have it do the correct reaction.

And likewise events can also just straight out trigger services. So sometimes there isn't a pattern at all. You just have an event happen and when you see an event happen, like a new customer, you want to go out and call the new customer processing service.

So they really come together in a pattern we like to call actually event-driven SOA, which to us really indicate the fact that combining those two, an EDA and an SOA, are really what's required often to address the full business needs as you start getting into these really agile different processes. Brett?

Brett Stineman: Okay. That's interesting. So is that a term we're going to be hearing more often now about, event-driven SOA?

Steve Lyons: I think it is. It's definitely one of the terms that's being bandied about in the industry, especially when you're talking about, kind of this marriage of where event-driven architectures are and where they map to service-oriented architectures.

Brett Stineman. Great. Now this actually, this topic was in our first audience question. And I'd like to encourage people in the audience to submit some questions based on what you're hearing. If something is causing you to think differently about how this all works together and the value of these technologies, please let us know so we can bring the audience into this conversation. And I'm going to -- I think this also might relate to where my next question goes. And it says, one of the things that tends to affect firms are things they don't notice until it's too late. So how does EDA or any other methodologies deal with events and facts that we simply aren't looking for? And so I'll open that up. Maybe Steve Demuth or Steve Lyons can talk on that.

Steve Demuth: Yes. So this is Steve Demuth. I'll take a shot at that. Just an interesting observation. Firms tend to be affected or damaged or killed by things they didn't anticipate. And of course at one level, that's unavoidable because it's impossible to predict everything that's going to happen in the market. But for most everyday purposes, we can, in fact, have business metrics that tell us the health of our businesses. And if we watch them, we will know when the health of our businesses is going off the rail. And it certainly requires some thought and analysis to determine what those metrics are.

So let's assume for the moment that you have some pretty good idea, and it could be customer retention. It could be average value per customer. It could be average time to resolve a contract.

There's lots of different things. But let's suppose you've figured out the metrics that determine whether or not you are successful as the business, and you're able to watch them because you've got the business intelligence tools to watch them. The interesting question then is, once those metrics start to go off the rails, once they go in the direction you don't like, how do you know what to do about them? And here's where I think the technology I think we're talking about fit in.

Because one of the things you can do with predictive analytics and statistical analytics is, having noticed that one of your metrics is going awry -- and let's say it's customer retention, that your customer rollover is increasing -- is to go back to your data warehouse and start looking at the collection of customers you have lost as the customer rollover numbers have gone up. And start to build predictive models from that subset of the customers that say, these are the things that indicate a customer at risk. This is the sequence of the interactions or as Steve Lyons said, non-interactions, expected non-interactions with a customer, that typically precede that customer stopping doing business with you entirely. And you can use that predictive model then either in a batch fashion or in a process-driven fashion to start to try to retain those customers, to take active interventions to retain those customers.

Now I want to give you an example what you could peel out of the news recently. If you've been listening you know that Netflix is struggling because they are in the middle of a transition from a business model that says, I make money by delivering DVDs through the mail, 3 at a time to customers who pay me a monthly subscription to a business where I'm both in that business but I'm also in the business of delivering content over the Internet, also for that same fixed monthly fee.

And Netflix has bobbled this. Because when Netflix introduced streaming video content, they clearly weren't watching carefully enough the segmentation of their customers into a group that was happy with the DVD model, and needed to be nurtured in the DVD model, versus those that maybe were experimenting with a combination. Versus those that had just been waiting for streaming video had pretty much abandoned the DVD model and frankly felt like, well okay guys, so now I'm paying for this service that I'm not using.

And I would submit that had Netflix been using proper business intelligence software to monitor those metrics and the proper analytics to segment that population, then they could have used an event aware architecture such as Steve's been talking about or the picture I had shown you earlier to look for the patterns of activities in customers that said, yes, this customer fits in this segment of happy with DVDs. This customer fits in this segment of experimenting with streaming and then falling off the planet because maybe they decided they liked somebody else's streaming better. Or this customer clearly likes our streaming and we need to do something to retain their loyalty.

You could detect that from the pattern of interactions. Netflix interacts with its customers multiple times per week. So they've got good data to mine on these things. And having done that, and then maybe if you combine it with the types of content they watch on streaming versus the type of content they get on DVD, you can drive point-in-time decisions about customer

interactions. How do I keep these customers loyal? And you can drive longer-term strategic decisions about how do I segment my offering to fit in to these emerging models?

Brett Stineman: Okay. I want to, I was going to say, I want to elaborate a little bit more on this idea of analytics that you're talking about. And I want to push out a slide, I think, that you can talk to related to this. And the point I want to ask you about is, there's various forms of analytics and this idea's been brought up several times, but typically analytics are thought of as being an offline activity, mining data for insights and intelligence. Now in terms of decision automation, how can analytics fit into this idea of -- these ideas we're talking about, of decision management and automation?

Steve Demuth. So we're talking here about -- let me take a broad overview of the analytics. In the broadest sense, there's a kind of analytics which are descriptive analytics in which you go back and you do some data mining, and you discover patterns in the data and you surface those patterns through reports. And that's the sort of thing that business intelligence does. In the IBM world we have a very robust Cognos line of products that manage all of that.

There's another kind of analytics, and the ones that we're more focused on in the decision management space is what we call predictive analytics. And it starts at the same place. It starts at going back to your data warehouse, to your historical data, and mining that for intelligence. But instead of generating a report that says, the population of our customers are shifting this way, Mr. Executive, what do you want to do about that? We generate things we call predictive models.

And the predictive model is an algorithm that derived from the statistical data about -- in your data warehouse, can be applied to a given transaction and make a prediction about how that transaction will play out. In the sort of easy to describe example of this is your credit score. What your credit score is a predictive model is is a predictive model that FICO Corporation or Equifax or TransUnion or whoever you get your credit score from, they've applied a statistical model to everything they know about you that tries to predict how likely you are to repay a loan if someone loans you money, or someone puts you on a cell phone plan that requires recurring payments over the course of 2 years.

And in the case of the FICO score, that a number between 200 and 800 that's normalized similar to the SAT scores are. But it could just as easily be a number that says, you're 90% likely to repay a loan of the following type. But it could also be a predictive model that says, statistically speaking, based on everything we know about our entire customer base, your potential value to me over the course of the next 3 years is about \$10,000 or it's about \$5,000. And those are the sorts of models you can build.

Now, the actual statistics behind that can be anything from multi-linear regression to general regression models to cluster models. There's all sorts of statistical niceties going on beyond. The thing to understand in the context of decision management is these models allow you to look at a particular instance, a particular loan application, a particular patient in the ICU, a particular -- any instance that you could be processing transactionally in an OLTP fashion, and make a very concrete prediction about the likelihood of this, the value of that. And then use rules in the

context of the decision management to say well, you know what? If this is a customer at risk, based on the analysis that Steve Lyons described in the event processing system, and this is a customer at risk and their potential value to me is \$10,000 over 3 years. And in the past they have typically talked to me through this particular broker that lives in their hometown, I'm going to send an alert to the broker to say, you've got to call Joe up because Joe looks to be at risk and he's particularly valuable to us. So let's make sure that he stays in place.

And in terms of the chart that Brett is showing here now, what we're trying to show, and particularly on the left-hand side of the chart, is that when you're creating a decision model that you're going to use for these transactional operational decisions, in some cases those decision models are going to be purely rule based. So if you're determining Medicare, Medicaid eligibility in the State of Minnesota, the rules were already written by the legislature, and you have no choice but to implement those rules as they're written.

But if you're making a decision about whether -- what the interest rate you should offer someone on a mortgage application is, or what price you should offer a potential cell phone customer, those things are not strictly rule based. Those are things that you can get at better by using these predictive models to say, well the potential value of this customer given everything I know about all my customers is \$75 a month. So I should be pitching them a cell phone plan that fits into that budget. I shouldn't be pitching them a \$200 a month plan because they're not going to buy it. And I shouldn't be pitching them a \$20 a month plan because that's going to undercut their potential value to me.

And so what we show in the diagram is rules and analytic models coming together to make decisions. The other thing we show there is very important and that is the improvement of the decisions. And one of the nice things about once you put your decision points into these managed decision artifacts, you can start to ask questions about how I can improve that decision. And you can change the rules or tweak the models and run a simulation using data you've got from the last quarter or the last month or the last year and discover wow, you know what? If I'd have made that decision differently, my outcomes would have been better or my outcomes would have been worse. So I'm going to use this tweaked model or I'm going to use these tweaked rules. That's what the simulation is meant to indicate in that diagram.

Brett Stineman: Okay. So I'm going to do a switch now from technologies to organizational dynamics, and specifically how line of business make key functions interact in terms of the systems that organizations depend on.

This question is for Jean. And many technology vendors, including IBM, have been providing capabilities to allow line of business users to have more direct control of their application. For example, with business rules management, being able to access and maintain those rules using tools design for people who aren't technical. And for IT this can cause concern around things such as governance and control. So how do you explain the role of business user participation to the IT teams you work with in your client engagements?

Jean Pommier: That's a good point, Brett, because and actually it pertains to what Steve was talking about earlier about changing, and the way to anticipate change which is always difficult.

So we are for more than 10 years, but actually 20 years, we have provided features, the technical features for business people to edit business logic. But it took way longer for many organizations to actually open the door to business users. And there are many reasons, good reasons for that. The good news is that in the tool mix of technology and also very innovative governance processes, governance, you can control the level of involvement you want to give the business users. And actually I would highlight that there is already some good involvement if you just allow business users to read the business logic. So just in a concentrative mode, not for modification.

Every year are at our user conference and we see, we hear more and more organizations giving some control of the business logic to their business users and this is very encouraging. It was used to be a few exceptions here and there and now it's, I would say over 50% of our engagement get some control to the business. So I would say it's becoming the norm and finally (inaudible).

So now the thing which is really key to make that happen and successfully happen is the importance of modeling because without modeling then it's difficult for users to understand their logic, which is behind the rules, the events and the analytical models.

So a good model, a good modeling exercise would make the system more maintainable because the rules are well organized and people know where to add the rules or remove rules. And then on top of that there is also testing, which is very important. So people are kind of scared on both sides. On the IT side, on the business side you have potential changes they can bring to the system. And again back to the point of anticipating change, decision management is all about that, about providing the right platform for people to test new things, new business models, new channels, new products and they are going to push to the market.

There was a good example. I like the example of Netflix that Steve brought up but AT&T and T-Mobile -- AT&T was kind of fighting against T-Mobile because T-Mobile was so quick to release new offerings and definitely decision management assembling all these products together give you this agility that the usual IT will offer doing everything in Java or C-plus-plus or even COBOL doesn't allow you.

So again, the modeling is very key to exist -- to bring these business users in the loop, and the technology really helps. But it's all about also processes, governance processes that you have to equip people with and that the platform supports. But there was a lot of change management as well, and making sure that the IT organization kind of let go in the control that they were used to. So I think it's a very positive move that we see on actual engagements and that we are going to see more and more now that the businesses are more educated about the potential that this technology is bringing to them.

Brett Stineman: Great. So this brought in a question from the audience. The question was, as you build out such a system as you're describing, how do you manage the impact of a rule change? And it says, change management and control would seem to be something that could create battles between IT and line of business. And I think that Jean or whoever else would like to respond to that.

Jean Pommier: So that's definitely very important. What we use actually as a theme in our consulting engagement is that we have to prepare the platform for change in the system for -- and in the change that organizations see. I briefly mentioned that this team and validation capability that you have to have on this platform that would provide on the IT side, that's -- so we have achieved tooling which other people to build their own use cases and test cases and so they can do what if simulation.

The other thing which is very important is to identify KPIs and I know Steve Demuth mentioned KPI before. So if you have KPIs you can have -- you can relate the changes to the impact on the business side. And so, very important. And then the governance, of course, when you don't want everybody to change anything without some control. So you have to be aware of that and we have again, the right technology to do that. And then we fine-tune the governance process to every organization. Our clients do that on their own with the tooling we provide them.

Brett Stineman: Okay. I think probably kind of related to what you're talking about here, it is a collaborative approach. It's not about handing it over to business. But it's about business and IT being able to work together, hand in hand.

Jean Pommier: Yes, very much so. And which is one of the key best practices would seem -- I know I'm going to bring in something which some people don't like in the agile programming approach. For us it's very key, especially in the case of decision management because -- for two reasons.

First, decisions keep changing. People change the way they are going to decide about things, and we talked about change before. A new thing is that we, when you talk about decisions this evening, it's not easy to grasp and to really describe in writing like we were used to do requirements for a year, in different systems or the, now, space programs. So in our case we deal with something which is pretty volatile. So the sooner you get the IT and the business to talk, and then to the agile principle, the fourth principle the drive is actually to make sure that the IT and business users interact every day.

So even if you do waterfall, I encourage you to consider this very strong interaction and make sure that you validate things. Why we are pushing iterative approach is that it gives a way to deliver value incrementally and very soon, as opposed to wait for 2 years before the project is completed.

And the second thing is that users like -- when they are going to see how what they say is translated into the tool, they may come back to you and say well, it's what I said but it is not what I meant. So we do a lot of refactoring with business users at this level. And again not to mention the policies keep changing, the competition keeps bringing more constraints. So you also want to re-write the content itself.

Steve Lyons: Great. And one of the keys -- so I'm sorry, Brett, if I can just jump in right off the top there. One of the keys to that, because it's very important, that we often go -- we have tools that are designed that we advertise that we market as designed to a line of business. And

certainly, they are useable by line of business, we have a line of business that use them on a daily basis.

But even in organizations where IT says our line of business aren't going to touch these tools because they're going to go into production, the value of being able to have IT use those tools and have it in a business-oriented language where they can share the rules or the patterns or whatever it is they're developing, kind of takes them into an email, show them or send it to a business person or whoever happens to be at the same screen and look at the rules that are being developed and to have a discussion and say, is this right, in English and not in code or some other language is really valuable and it really takes everything to the next step. Even if, as I said, even in those organizations where IT is against having a line of business actually being involved in the tool. And we really see the gamut where it really runs all over the map whether or not depending on the organization.

Brett Stineman: That's a great point, Steve. And actually, I have a -- coming back to some of these technologies been talking about, I want to ask you, Steve Lyons, another question. Are there best practices in terms of topologies for building out decision management applications and solutions? And as part of that, there was a question that came in, how do you bring in things such as social media and other types of information and data that's coming and going across an architecture that maybe in the past hasn't been able to be incorporated into these systems?

Steve Lyons: Sure. So the social media question is actually extremely interesting. I think Steve Demuth talked a little bit about analytics and doing predictive analytics. There's tooling, and there are many ways to bring it in. We have lots of sources of information, lots of sources of events in our organization. One of the things, one of the ways to solve the social media problem that I've seen is, we have tooling that does things called streaming analytics. And streaming analytics can be used for text. It can be used for binary information. It can be used for things like monitoring blood pressure, minute changes in blood pressure and things like that.

Now what it really is is you're actually taking a stream of data which might be text in a social media from Twitter or something like that, and you're doing the same thing that we're doing on the business side, on the business event side. You're looking for patterns, right. So maybe you're looking for mentions of your product, your new product that you've just started marketing in Twitter feeds. Maybe you're looking for slight changes in temperature for an infant in the NICU, or blood pressure or whatever. But streaming analytics allow you to do, that's kind of another phase, and that allows to take those patterns of events that you see, the streaming events if you will, and maybe what you do is you raise a business event to your business event system and you combine that with other information that you might have.

So we've be doing a lot of work very recently in the last year or two with energy companies, for example. And a few of the energy companies were using used tools like streaming analytics to watch energy consumption. Lots of meter readings, smart meter readings from all over the place.

They react to those, they send a business event off to the business event system that might say, we're seeing an outage or we're seeing drops in this area. And then the business event system can then take that information, take that event that says there may be an issue in this area. Take a

look at other business factors, like do I have any trucks in the area to repair trucks that I could rollout to that area? If they're cheaper, they may make a business decision at that point to say, yes I have one that's right in that area. I'll send one over to investigate. Or no, I don't have anything that would actually cost some serious money to send a truck over there. I'll wait a period of time and see if it's an anomaly and it goes away, or whatever.

So there's streaming analytics. There are other IT eventing tools. There's, I had mentioned RFID before. There are lots of different ways that you can take other information that's happening in your organization, turn those into events and then send those -- and combine all those together and then take intelligent action based on that.

Steve Demuth: Steve, I would add to that specifically for the question on social media where we're looking at text, there are tools, and we have them in the XPFS brand, specifically for deriving meaning via text analytics. You saw some of this if you saw Watson playing Jeopardy. That's sort of the high-end version of text analytics. You have versions of that technology that can monitor Twitter feed, monitor Facebook, email streams, and actually derive specific meaning as related to the specific interactions and throw events based on that. So you can actually get solid meaning from these social media sorts of things.

Brett Stineman: Absolutely. And some of those actually, we got some new questions coming in but we are running out of time for today's session. Yes, I'd like to just put that one last question in to Steve Demuth, maybe you can answer this.

In terms of organizations moving from approaches dealing with individual decision in challenges to a broader adoption of decision management, and how they can be successful in that. Any final thoughts? And while you answer that I'm going to push out a feedback form that people can respond to while you're answering that last question.

Steve Demuth: Sure. So I think the critical thing here, and it's probably familiar from the adoption of other new technologies, is to focus on getting your business and technology groups simultaneously and jointly aware of the value of identifying and managing specific decisions from which you can derive operational value by organizing those decisions. And I've talked about lots of examples of what they are. And then set up a center of competency or a center of excellence or work group, whatever works in your business situation, to pull together the various threads because we're talking about a lot of different threads here. By the time you bring in events, rules and analytics, process and business intelligence and monitoring, in order to make this smarter approach, you need someone who's looking at the big picture. And often what works in that case is to have in your enterprise architecture team, and your enterprise business team, someone, a center of excellence that focuses on identifying opportunities related to decision management, and determining which pieces are required for an individual piece.

And I'll give you a very simple example. We have a very large client who's in the business of managing 401Ks for other companies. And there are many, many opportunities for improved customer or participant alerting, improved interactions with the program sponsors, meeting their requirements under the new consumer finance laws as to how they inform customers about purchase and all those things.

What they're doing is setting up a group that is responsible for looking at specific processes and specific decision points and then recommending how to use the enterprise platform to bring together analytics and rules, events and rules, for specific applications to drive very particular business value. And I think you'd have to approach it that way. These are not technologies for technology's sake. The idea here is to make your business decisions smarter so your business is more successful. And so you start from the business problem, and you bring a technology center of excellence to bear on those business problems.

Brett Stineman: Great. Now that takes us to the end of this hour. I just am showing the audience some information on the next steps, and you can get this by downloading the slides. And to close, I'd just like to say, visit these links to get more, and within the next 24 hours you will receive a personalized follow-up email 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. And we have one final webcast next week. I encourage people to attend that one, and it will take place live on September 27, 11 a.m., Pacific time, 2 p.m. Eastern time. And this concluding session will be a live demonstration called Architectural Differences and Similarities between Rules and Events.

I'd like to thank everyone for attending today's webcast, Expert Panel Best Practices for Designing Decision Management Architectures brought to you by Information Week and IBM. This Webcast is copyrighted 2011 by United Business Media, LLC. Presentation materials are owned by, or copyrighted as the case by Information Week and IBM are solely responsible for its content and the individual speakers are solely responsible for their content and their opinions.

On behalf of our guests, Steve Demuth, Steve Lyons, and Jean Pommier, I'm Brett Stineman. Thanks again to our panelists again. And thank you, our audience for your time. And have a great day.