Event ID 323224 - Prescriptive Analytics and Optimization for Smarter Business Decisions

Gabriela Wilson: Good morning everybody and welcome to today's Webinar, Prescriptive Analytics and Optimization for Smarter Business Decisions. My name is Gabriela Wilson and I am the moderator for the day's event.

Before we begin I would like to make a few brief announcements. This webinar is being recorded and will be available on demand for three months. You can access the archive beginning tomorrow, simply by clicking on the URL that you received to join this webinar. Our webinar is designed to be interactive so please be sure to disable any popup blockers you may have installed as they may interfere with certain aspects of our web conference.

For an enhanced view of our presentation please click the Enlarge Slide button located just underneath the slide window at the bottom of your screen. You may download a PDF version of the slide presentation by clicking the Download Slides button located underneath the slide window.

And last but not least, the Q&A session will be held during the last 15 minutes of this webinar, however, you may send a question at any time and the presenter will address it then.

Please keep questions relevant to the subject at hand. To submit a question, please type your question into the question box located in the lower left of your console and hit the Submit button.

Now, I am very excited to introduce today's speaker, Dr. Jeremy Bloom. Jeremy is IBM's Senior Product Marketing Manager for ILOG Optimization. He has more than 25 years of business experience applying operations, research and optimization. Previously he worked at the Electric Power Research Institute and General Public Utilities. His experience includes integrated resource planning, retail market tools, distributed resource and asset management.

Dr. Bloom received his undergraduate education in electrical engineering from Carnegie Mellon University and his graduate degree in operations research from MIT.

And without further ado I would like to hand over the presentation to Jeremy.

Dr. Jeremy Bloom: Thank you Gabi and welcome everyone. My purpose today is to talk about how optimization and analytics are transforming businesses and government functions across the world. And what I would like to do is cover the following topics.

I'll first of all talk about what analytics is and the analytics landscape as it exists today. I'll talk specifically about optimization, what it is and how it works. Then I'll give you a bunch of optimization examples that are -- show applications of optimization in real world problems. And finally, I'll have some concluding remarks.

Let's talk a little bit about what analytics is today. Analytics has received a great deal of attention in the trade press for the last several years and I think that's due to a couple of trends. One is the increasing power of our information technology to provide data and lots of data, and to manage large amounts of data. And people are asking, "What can I do with this data to improve my decision making processes?"

The second reason why analytics has come to prominence is because some very well-known authors studying how companies and government agencies are using information have identified analytics as a critical competency for these organizations. And on -- this chart here shows one of the results of that work. It's a taxonomy of different kinds of analytics and how they apply to providing competitive advantage and efficiency in organizations.

There are basically three categories of analytics. The one that probably most of you are familiar with are descriptive analytics. These are the various charts, reports and tables you can generate with data. There is standard reporting which tends to be repeated periodically, day-by-day, week-by-week, month-by-month, to tell people where they are and what's happening. Often that leads to further questions that need drilling down into ad hoc reporting, asking questions like, how many times has this happened, how often does it happen, where is it happening?

And then finally, we go to queries and drill downs, to find out exactly what the problem is. On the next level up are predictive analytics. And predictive analytics focus on trying to translate the history that we have in our data into the future to predict what could happen. These things -- predictive analytics include things like alerts, simulation, forecasting and predictive modeling. Forecasting is an attempt to extrapolate into the future trends that exist now. And predictive modeling is an attempt to determine how things will occur in the future, how people will behave, how systems will react based on their past history.

The final level is prescriptive analytics and prescriptive analytics includes optimization. And prescriptive analytics includes optimization. And prescriptive analytics helps to determine what to do given the information that you have, either, the descriptive analytics of the current situation or predictive analytics about what could happen in the future.

Optimization, which I'm going to focus on today, talks about how we can achieve the best outcome. And above optimization is stochastic optimization which is about risk and hedging risks. Now, as we proceed up this chart we get greater and greater efficiency and greater and greater competitive advantage. But we also increase the degree of complexity so most organizations today focus on the descriptive analytics and some of the predictive analytics.

What I'm going to try to do today is show you the value of going to the next step of prescriptive analytics. And predictive and prescriptive analytics together comprise of a category that we generally call advanced analytics.

Now, it may surprise some of you to know, it may not surprise others of you to know that there is actually a science of better decisions and that's called operations research or management science. And optimization is one of the technologies that has made the science of better

decisions a real science. Optimization can help businesses create the best possible plans, explore alternatives and understand tradeoffs and respond to changes in their business operations and their business environment.

Some examples, and we'll talk about some of these examples in much more detail later on, but some examples are airlines determining how best to allocate their aircraft and crew to fly the flight schedules that they've put in place. In manufacturing and consumer products frequently there is a tradeoff to be explored between holding inventory, which costs money, and customer satisfaction, higher inventories enable you to fill customer orders more frequently and more quickly, but at the cost of holding additional inventory. What's the tradeoff, what's the best level of inventory to hold?

In production manufacturing, supply chain environments frequently the question is what products to build and when and where, how frequently do I need to run a production lot, where do I produce it, do I produce it in a domestic market or do I produce it overseas and transport it to the domestic market?

In the finance world, the issues are around risk versus potential reward, typically you can earn greater returns by assuming greater risk. But there is a tradeoff and companies and investors have to be comfortable with that tradeoff.

Another one that's of high interest today is in energy production, the cost of energy production versus carbon emissions and as companies attempt to get their carbon footprints under control they have to wonder whether the costs are commensurate with the benefits. We'll talk about an example of that one a little bit later.

Another example, is routing package deliveries, how do I get the best utilization out of a fleet of trucks that I have in order to deliver my products on time, where they belong, at minimum cost, and I'll talk about a detailed example of that one a little bit later as well.

Let's talk now about what optimization is. Some of you may be familiar with optimization, I suspect a lot of you have heard of it but perhaps do not understand what it is so let me talk in somewhat detail about it. Optimization helps businesses and organizations make complex decisions and tradeoffs about using limited resources.

So typically in an optimization problem you have an objective that you want to optimize, you want to minimize cost, you want to maximize throughput or you want to achieve some maximum level of customer satisfaction. But you have limits, there are resources that are needed to do this and you have limits and you have multiple ways of achieving the goal. And the question is, which is the best combination of decisions to meet that goal given the limitations of the resources that exist?

Optimization can help you discover previously unknown options or approaches and the reason for that is because optimization could automatically evaluate millions of choices, choices -- many more choices than a human being could possibly assimilate. And among all of those

millions of choices, could find the ones that satisfy the constraints and achieve the optimum of the objective.

Optimization also permits you to automate and streamline decisions because it enables people to work with the data and to come up with solutions automatically involving compliance with business policies and regulations. What that does for your planners and your operation managers is it frees us their time and their attention so they can leverage their expertise across a wider set of challenges. This automating routine decisions enables people to deal with exceptional situations where human intelligence is necessary to a good solution.

Optimization also allows you to explore more scenarios and alternatives. You can understand the tradeoffs and sensitivities to various changes. And that's critical for people to accept the recommendations of a prescriptive analytics process. You can gain new insights into the data that you have and you can view results in new ways. Optimization has all of these benefits.

Now, optimization is typically used to allocate limited resources. And some of the resources that are typically allocated using optimization are capital, people, equipment, facilities, vehicles, raw materials and so on. And key words that indicate that there is an optimization problem present are words like minimize, maximize, or how many, or how much, or decide, or chose, plan, when or where, schedule, assign, route, locate or tradeoff. These words indicate that there is an optimization problem.

Optimization typically is used to do planning or scheduling in situations where there are complex operating constraints; there are limited resources, there are large volumes of data, there are complex manufacturing or design processes and in situations where there are multiple business objectives, reduce time, reduce cost, increase productivity or other key performance indicators.

And optimization enables adjustment to changes in the operating environment and what-if analysis; exploring what would happen if I did something differently or if the environment were to change.

Whether the problem is large or small, straightforward or complex optimization supports effective decision making across a wide range of issues. Firms in many industries and in government use optimization to solve business problems ranging from long-term planning to real time scheduling and rescheduling. And in fact optimization can be used in a variety of contexts where the time horizon is either near or far.

At the top level of this chart long-term planning and optimization is typically done on an annual basis or perhaps a quarterly basis or even on an occasional basis when there is a major capital decision to be made. And typically long-term planning has to do with capital decisions, investment decisions that require making investments that have long lives and have long -- far-reaching consequences for those decisions. So whether to expand a plant or open a new one, how many distribution centers to have, the value of having additional equipment over time, these are the kinds of issues that get addressed with long-term planning.

Short-term planning, usually called tactical planning, typically occurs on a monthly or weekly basis, depending on the operating environment. Questions that are addressed at this level are things like how much should we produce this week, how many shifts should we have, how much of various resources do I need, or which marketing campaigns will provide the most impact for a set budget?

Finally, the most, finest, level of detail occurs with detailed scheduling. Detailed scheduling typically occurs on a weekly basis or often on a daily basis sometimes even on an hourly basis. We have customers who do detailed scheduling on a five-minute basis. And these involve questions about specifically, which activities should be done when, sequencing activities and time, which resources should be assigned and when can maintenance or special tasks be most effectively scheduled. So optimization gets used across a wide spectrum of planning horizons.

How does optimization work? So this diagram explains in conceptual terms what optimization is. In the middle, at the top of that disk, you see a mathematical model and typically that`s where optimization starts. You have to have a mathematical model of the system or the process that you want to control and people who are skilled at building these models typically are people who are trained in operations research or management science.

Now, optimization is a data-driven activity, intensely data driven and so having formulated a mathematical model it's important to be able to populate that data, that model, with input data. And on the left-hand side of this chart you see some of the inputs that are needed. So typically there will be demands to be met, there will be resources available, there will be costs and yields and recipes for turning inputs into outputs, how many hours of labor are required to produce an automobile for example.

There are constraints, typically these things include operational constraints like the maximum number of shifts or the maximum time a worker can work. And there might be customer preferences, customers obviously prefer early delivery to late delivery but there might be a penalty for arriving too early and there might be a larger penalty for arriving too late; what's the window. And another important source of business goals -- of inputs, are business goals; what am I trying to achieve?

So what I do is I formulate a mathematical model of the system or the process that I want to control, I populate it with data, and I feed that to an optimization engine. And I'll talk in a few minutes about the optimization engines that we provide at IBM. And the result is a plan or a schedule, a detailed set of decisions that will be made, recommendations. And associated with that plan or schedule are, metrics, key performance indicators that indicate the quality of the solution. So if I'm trying to minimize costs, a plan or schedule that I develop will come with a projection of what the cost will actually be.

At IBM we provide a number of different products to help companies and government agencies use optimization. Our most basic set of tools is CPLEX Optimization Studio, this is a tool for developing optimization models. It includes both the optimization engines, and we have two, CPLEX is our -- CPLEX Optimizer is our mathematical programming engine and CP Optimizer is our constraint programming engine. We also have a full suite of development

tools that enable you to formulate models using our mathematical programming language called OPAL and it uses a graphical user interface that enables people to build and debug models.

The typical use of this tool is in an analytics department where there are specialists who know how to build optimization models. And those specialists typically are working for a line of business clients in an organization, or perhaps they`re working as consultants or independent software vendors who are trying to solve the business problems of their customers. But the key point is that CPLEX Optimization Studio is typically deployed in situations where you have optimization experts who want to use our software to build and solve optimization models.

Now, in addition to the analytics department, more and more information technology departments are adopting optimization technology. And they're doing it to serve line of business customers, business decision makers who have problems that could be solved using optimization. And frequently those line of business decision makers do not know about optimization technology and its power, but they do have business problems that have very significant value if they can solve them.

For these audiences we have a couple of other products. The first product we have is our enterprise solution platform which is called ODM Enterprise and ODM Enterprise provides for scalable enterprise deployment. So the target market is typically IT departments and line of business users. It includes the ability to create and solve optimization problems as provided by CPLEX Optimization Studio but also provides the ability to configure user, business user interfaces, that allow business users to interact with an optimization solution in ways that are natural to them.

They can create scenarios with alternative assumptions and test the solutions that they are developing for robustness against changes in the environment. They can also interface with corporate IT systems. And they can put the solution of the optimization problem which is typically a very computational intensive process, put that on a server so it doesn't bog down the other functions of the corporate IT system.

Now, on top of ODM, Enterprise, IBM has a set of vertical assets, specific optimization tools that are designed for specific types of business. And so as we encounter companies and agencies that have specific problems, we have a lot of expertise wrapped up in these vertical assets that help accelerate deployment of an optimization-based decision support application.

The final category of tools that we have, or applications we have, are supply chain optimization applications. These are specifically designed to support functions in the business of managing supply chains. We have the ability to help you optimize your distribution networks, to optimize your production, to develop transportation plans and to manage inventories. And these are specific vertical applications designed specifically for supply chain optimization.

IBM ILOG is the leader in optimization around the world. We have over 1,000 commercial customers under maintenance, over 160 of the global 500 have customized analytical decision support using ILOG Optimization tools and engines. Over 50% of the world`s largest supply chain and 50 global supply chains use ILOG Optimization, typically through supply chain --

our supply chain applications. We also have major software partners that reach thousands more, so if you're using SAP or Oracle or any of these other companies for an application that involves managing a decision process or a supply chain, chances are you've seen ILOG Optimization software embedded in there.

We also have a very large installed user base in the universities. Over 1,000 universities use our products in developing their research. And so our product CPLEX is cited in over 95% of the scientific publications that mention a mathematical optimization solver.

I'd like now to turn to talking about some of the applications of optimization, that you may encounter in your organization. As you can see from this diagram there are applications of optimization across all sectors of business and government. I'm not going to read through these but you can see that there are just a variety of problems that can be solved using optimization. And manufacturing, production planning and detailed scheduling of production, are typical applications. And transportation and logistics, fleet assignment -- assigning of vehicles, like flights to a certain demand patterns, network design, where to place warehouses and so forth, vehicle routing and delivery scheduling, these are optimization applications.

In financial services, a lot of optimization is involved in portfolio, optimization in rebalancing. Energy and natural resources, power generation scheduling. In telecom, network capacity planning, routing and adaptive network configuration. And we have applications that span multiple sectors, workforce scheduling, advertising scheduling and so forth.

Now, just in case you think these are all business applications we have multiple applications that are used in government as well. And so I've highlighted some of the types of applications that you are likely to see in a government agency, in this diagram as well. And you can see that there are many, many applications of optimization in the government sector.

Here are some examples of the return on investment that optimization applications can claim. Now, these applications actually come from an organization called The Institute for Operations Research and Management Science which is a professional society. And every year in INFORMS runs a competition called the Franz Edelman Award that identifies the best application of operations research. And one of the criteria for that award is a demonstrated benefits on the bottom line, for these applications.

So, each of these examples here are companies or organizations that have used ILOG Optimization, our product, to either become finalists or to win the Edelman Award. I'm not going to talk about these in detail right now, I will talk about some of them when I talk in more detail but you can see there is a variety of dramatic savings on the bottom line that result from using optimization. South African National Defense Force, for example, used optimization in defense force design and came up with 22% cost savings.

Moving further down, Continental Airlines, used optimization to define the repositioning of crews after a disruption such as a snow storm. They were able to demonstrate $40 million of savings in just one year.

Indeval which was the winner of this prize in 2010 does securities settlement in Mexico. And they are able to achieve a cost savings of $150 million a year in addition to reducing the settlement time from a day to a few minutes and reducing the risk of unclosed transactions.

So these are real world examples developed by a professional society, of how optimization can create real value on the bottom line.

Let's talk about some specific examples. And my first example is strategic transportation management, moving freight by trucks. The customer in this a government agency and it's one of the largest employers in the country and it operates one of the largest vehicle fleets on the continent. The business problem here is to officially operate a large and complex transportation network to accommodate many different levels of services to its customers, considering different classes of freight of various sizes and weights.

The solution in this case was developed by the client and IBM using the IBM ILOG Optimization software. They created a process called the Corridor Analytic Program that uses advanced analytics to analyze highway transportation scenarios and identify cost saving opportunities. This CAP model examines the routing and scheduling options available in order to determine how to assign shipments to vehicles, to minimize cost while meeting business goals.

So it uses specialized optimization technology built on IBM ILOG optimization technology, to identify the best allocation of freight among the various transportation resources. It enables a quick calculation of the best possible utilization of the resources, which include vehicles, personnel, time, processes, equipment, raw material, supplies, capacity and security.

And it achieves a desired business result of minimizing the cost or process time to maximize throughput for service levels. The bottom line here is that this Corridor Analytic Program identifies opportunities to consolidate truck shipments to maximize user resources without sacrificing service levels. It provides the client with the most efficient plan for utilizing existing transportation assets, accounting for routes, delivery time, truck capacity restrictions and freight classes.

A typical scenario involves approximately 15,000 business constraints and 3,200 variables which represent the decisions that have to be made. And the client realized savings of over $5 million US dollars annually from using this application.

Here's another example, this one from the private sector, cash management. In this case, the question is how to restock automatic teller machines. The customer in this case is a services company that provides financial electronic commerce services and products to financial institutions worldwide. It provides the processing systems that -- systems that process more than two-thirds of the automated clearing house transactions in the United States and it provides reconciliation, financial messaging, workflow and compliance products and services to more than 600 banks and businesses. Its customers manage 2.6 million portfolios totaling about $1.8 trillion U.S. dollars, in assets. So a lot of the -- a lot of assets are managed using the software that this company provides.

The business problem in this case is to restock automatic teller machines. So automatic teller machines are placed in locations where people want to get cash. And it's important that those machines have enough cash on hand to make the demands and of course the demands are highly variable. So if you're talking about a cash machine in Las Vegas on the Strip, during the week of a major athletic event, a lot of people are going to be in town, there is going to be a lot of demand for cash and those machines you don't want to run out of cash.

Now, on the other hand, cash in automatic teller machines has a cost. And the reason it has a cost is because it's not earning any interest so as long as their cash is in the automatic teller machine there is a lost opportunity cost from the interest it could be earning. There's also a delivery cost, and remember cash has to be delivered to these automatic cash machines using armored vehicles with crews, so there are significant costs involved with the delivery. So they want to restock these teller machines taking into account customer withdrawal habits and cash management regulations put in place by the government.

If there's too much cash you have carrying costs. If there's too little cash there are angry customers. And there are always forecasting errors, you can't predict with absolute certainty how much cash will be needed so there is volatility and there are data errors. The data is not perfect. It's static so it doesn't necessarily reflect the current situation. It's dirty, there are incorrect values, missing values and wrong values.

Now, the way cash management works is that cash is disbursed by the Federal Reserve Banks of which there are 33 zones in the United States and it's disbursed to bank vaults and typically there are two to 12 vaults per zone. And then, the cash is taken from the vaults and put into the automatic teller machines, or if there is excess cash at an automatic teller machine it might be returned to the vault.

Now, there's a series of government regulations that deal with disbursing cash from the Federal Reserve, called Cross-Shipping penalties. Basically, Cross-Shipment apply when a bank withdraws a certain denomination of cash from the Federal Reserve and then re-deposits the same denomination within a week. And so those Cross-Shipment penalties are substantial and the banks wish to avoid having to pay them. But on the other hand there is a requirement for high service levels. There's customer requirements, as I said, that are highly variable but you don't want to run out of cash.

And there is a substantial business case in this situation for an optimization based solution. So what does an optimization solution look like in this case? Here's a simple example where I have three vaults and I have a four-day period.

So one possible solution is the numbers next to the vaults indicate the net demand for cash. So plus-10 means I need, in this case, 10 bricks which is 10 bundles of cash on day one. And in this case I've got a demand for -- and excess demand of -- sorry, an excess cash of 10 bricks on day one but I'm going to need 10 bricks on day two so I handle this by carrying the cash overnight at the vault.

Another situation occurs at vault number two where there is excess cash of 10 bricks on day one, there is no net demand on day two but on day three there is a demand for 50 bricks so I can carry forward 10 bricks from day one to day three and that partially satisfies the need for cash on day three.

Another possibility is that vault one now has excess cash on day three of 40 bricks, so I can redeposit that at the Federal Reserve, provided I haven't withdrawn that same amount -- that same denomination during that week.

Another possibility occurs at vault three. Vault three has excess cash of 10 on day one and an additional excess cash of 10 on day two. So a net total of 20 on day two. I can cross transport that from vault three to vault two to help satisfy the demand at -- on day three at vault two. So now I've shipped 20 from vault three, I've carried over 10 from vault two, I'm still short 30 at that vault on day -- vault two on day three.

So I might need to make a withdrawal from the Fed, of 20 bricks, on -- in order to satisfy the demand at vault two on day three.

Another possibility is to ship cash from vault two to vault three on day four to take care of the excess cash of 20 that's going to occur at vault two on day four and to satisfy some of the demand -- I'm sorry, to satisfy the demand in vault three. In this case I only need 10 so there's 10 left that needs to be carried forward somewhere.

So this gives you some examples of a possible way to formulate this problem. The problem though is very complex because there are multiple decisions that can be made, there are multiple choices that could be made and trying to sort through all of these choices manually would be very difficult for any human being to come up with a very good solution. Optimization on the other hand can completely solve this problem automatically and determine the best way -- the best routing pattern, of cash through the vault system. And you can imagine how much more complex this is when there is more than three vaults and more than four days in the solution.

So in this case, the client implemented this solution for one of its customers. It had a customer who is disbursing $200 million of cash per day through 20,000 retail outlets -- that is branch banks and ATMs -- total cash in this system was $7 billion.

And the goals in this case were -- no change in the current schedules for replacing currency; reduce the cash level, the cash inventory levels to reduce carrying costs; reduce replenishment costs by reducing the number of deliveries that need to be made; reduce the cost of shipping -- Cross-Shipping penalties from the Fed for withdrawing and re-depositing the same denomination of cash during the same week. They wanted to improve reporting capability, they wanted to be able to piggyback shipments so that they could make multiple deliveries out of the same vehicle. And they needed to solve this problem overnight because they needed to be able to have a schedule to implement the following day.

They ran a pilot at this company, for this customer, for six months using 58 vaults and they were able to reduce cash inventories by 35%. That was attributable not only to the optimization model but to other factors including better forecasting, better operations, better people and better measurement.

They reduced replenishment costs by 55%, they reduced Cross-Shipping fees by about 63%. They were able to demonstrate that they could successfully run the optimization overnight in order to have a schedule for the following day. The project was rated highly successful by the client's internal Six Sigma Unit and they rolled it out to the entire enterprise in 2008.

The customer was very satisfied with the solution -- this is the financial services provider that implemented this solution for one of its customers. And they found that the optimization was extremely satisfactory to them. And using these advanced analytic of optimization technology, gave them a solution that met their needs and provided significant, substantial benefits on their bottom line.

Let's talk about another example. This one has to do with transportation scheduling, in particular this one has to do with transportation scheduling, in particular train timetabling. The customer in this case is Netherlands Railways, it operates the busiest national rail network in Europe. It manages more than 4,800 trains per day, it has 2,100 kilometers of track and 700 -- sorry, 279 stations.

Between 1970 and 2006 traffic doubled. And during that same period, freight transportation increased by 285%. In 2006 they carried 9 million different passengers, a total of 15.8 billion passenger kilometers, they had operating revenues of EUR15 billion and operating income of EUR200 million.

The business problem was this that they had been operating with a schedule that had been developing in 1970. And during that time, as I said, both passenger and freight traffic had doubled and they needed to reschedule their trains. They wanted to keep costs down and they wanted to keep service on time and they wanted to more accurately match trains to expected passenger traffic by developing a new automated timetabling system.

They also wanted to stimulate growth of passenger railway transport without investing lots of money in new infrastructure. They wanted to improve on-time performance and they wanted to reduce greenhouse gas emissions by reducing automobile traffic.

Now the timetabling problem is interesting because it has a number of features that make it hard to solve manually. There's a tradeoff between the number of trains and on-time performance, the more trains you have the more likely you have on time performance, but cost obviously increases with the number of trains that you operate.

They wanted to make sure that the frequency of arrivals conformed with customer needs, and the Netherlands Railway system does not operate in isolation but has connections with Germany and Belgium, that has to be observed. They wanted to use what they call a cyclic schedule, which means that they have trains departing from the same location to the same

destination, at the same minute after every hour. So if someone would know that if they wanted to go to a particular location they could always catch the train at 15 minutes past the hour.

They also have a line system with sets of origins and final destinations and frequencies which passengers make those trips, and a stopping pattern. So they wanted to develop a timetable which is the planned departure and arrival event for each train at each station served. And they have to deal with not only the timetable but also station routing, where does the train arrive, what's the inbound track, what's the platform, what's the outbound track.

They also have to schedule crews. They have to assign crews to these trains and typically they want to have a cyclic schedule for their crew, cyclic meaning that the crew returns to its home base so they can go home at night.

So they used optimization to develop a new timetable schedule, which was implemented in 2007. They were able to improve operating efficiency by 5% to 10%, which resulted in an additional profit of EUR10 million annually. And they were also able to improve their on-time performance to a record high of 87%, up from 84.8% in the previous two years. So in this case not only did they improve cost but they also improved performance and service.

Now, this is one of the examples that I cited as an Edelman Prize award winner, this -- in this case, the use of optimization to create a new timetable was recognized by the Institute for Operations Research and Management Science for its practical implications and its implementation of an advanced analytic in order to create real bottom line value.

Let's take another example, this one is operations planning, deals with dispatching power generators. In this case the customer is Red Electrica de Espana, which operates the Spanish national power grid and in 2008 this company managed the production of 280 billion kilowatt hours, it served a peak demand of 43,000 megawatts and handled 95,00 megawatts of generating capacity including 17,000 megawatts of hydroelectric and 16,000 megawatts of wind.

Spain has one of the largest wind generation bases in Europe. And it also operated 35,000 kilometers of high-voltage transmission lines. Their revenue, in 2008, was EUR1.1 billion and they had a EUR286 million after tax profit.

The business problem here is to schedule the dispatch of power generators. They need to know when to startup each generator and how much to produce each -- during each hour in order to minimize the cost of electricity consumers.

Power supply and demand must match at all times, but wind production is highly variable and so they have to be able to manage the intermittency of the wind. There is a significant cost to startup a combustion generator. There are seasonal limits on hydroelectric generation. But the reason for building a new solution to this problem was that market structure had changed. And because the market structure had changed it made the old methods of dispatching the power system obsolete.

So this is a diagram using the ILOG optimization product called ODM Enterprise that shows what a solution to this problem looks like. At the top you can see the dispatch of the power generators on an hourly basis over the course of a week and you can see that some of the generators, those at the bottom in red and blue, operate pretty much continuously whereas other ones are cycled on and off. And the lower part of this diagram shows the on-off schedule for the generators. This is the type of solution that gets produced by this optimization application.

So the bottom line, was that Red Electrica de Espana reduced their production costs by between EUR50,000 and EUR100,000 per day that`s $70,000 to $140,000 per day. And they also were able to reduce their carbon emissions by approximately 100,000 tons of CO2 annually, so in many cases people believe there is a tradeoff between cost and carbon and if you want to reduce carbon you necessary increase cost.

And this case demonstrates that that's not always the case, that in this case they were able to achieve both cost savings and carbon reductions. Now, interestingly enough, the documented cost of this project was EUR300,000 which means that it pays for itself in approximately less than a week and this is an extreme example but certainly not atypical. Optimization typically has a very high return on investment and it's quantifiable.

The customer was very satisfied with the solution which they built using the ILOG optimization software because it enabled them to build a new application that was much easier to manage and much easier and much more reliable than the previous implementation they had and also met the new market rules that were in place in Spain at the time.

Here's another example, this one is a government air service that provides air support for firefighting. In this case the customer was a Canadian provincial government that runs the air service that provides air support for firefighting, all over the world, as well as providing transport for the provincial government and medical and security transport.

The business problem they wanted to solve was to optimize the schedule of their employees. They wanted to ensure employee satisfaction, manage the work hours. In aircraft or airline operations there are very strict limits on how long a crew member can work without a break and very strict rules on the amount of rest time they need to have between work schedules. And they wanted to reduce the amount of time that it needed to produce employee work schedules. Work schedules typically are highly variable because employees get sick or they want requests for time off, they have to be accommodated. And by automating the scheduling process they were able to make the scheduling process much more flexible, and much more easier to accommodate employer requests while reducing the amount of time required to create the schedule.

In this case an IBM business partner implemented a solution using the IBM ILOG optimization software that takes into account all the needs of the air services provider. So the bottom line here is that it -- this solution evaluates numerous scenarios to enable the managers to understand the best possible solution that they can get, improve service quality by providing aircraft when and where they are needed. Reduced operating costs, reduced management time by automating the scheduling process, enables assignment of employees across a broad range of

activities while respecting the organization's regulations. Minimizing the use of overtime, creating robust schedules that can be -- accommodate employee requests for time off, and manages exceptions.

My final example has to do with a situation in which the bottom line is not measured in money value but it's measured in lives saved. And so this is one of my favorite examples, it has to do with organ transplants. So the customer in this case was a nonprofit organization that focuses on live donations of a reticular organ and in particular it arranges organ swaps among incompatible donors and recipients. I'll talk about that in a little bit more detail in a minute.

At the time this case was written this case had 70 transplant program members in 25 states in the United States and it began enrolling transplant centers -- hospital transplant centers in October of 2006. So the business problem here is that there are more than 83,000 people in the United States who are waiting for a transplant of this organ and about 12 die every day due to the lack of donors. Now, in many cases a patient has a significant other, a person who is willing to donate this organ, but because of incompatible immune systems that donor cannot give that organ to that patient.

So what this function does is it arranges paired donations, it finds pairs of donors and recipients who are incompatible with each other but who would be compatible if they were able to trade organs. And so it's also possible to create longer chains and transplants don't have to be performed simultaneously. So here's a situation, I've got a pair of donors, donor A -- sorry, a pair of donor recipients, donor A and recipient A, so donor A for example is a husband and recipient A is a wife, but they're not compatible donors, in other words a donation from donor A to recipient A would be rejected by the recipient.

I've got another pair, a similar situation, donor B and recipient B who are also not compatible. But it turns out that donor B is compatible with recipient A and donor A is compatible with recipient B. So what this program does is it finds such paired donations and it arranges for these transplants to take place. Now, it's possible that the chain of donations can extend beyond pairs. So if I find three donor recipient pairs who among them can find compatible matches, I can match them up. And in fact, in one case this organization was able to find a chain of 10 donors and recipients.

But finding these kinds of pairs -- these pairings and these longer chains is a very complicated problem that can't be done by human beings, but it can be done very easy with an optimization model. And so in this case, what the solution does is it determines organ compatibility by measuring it on a point scale that includes things such as the compatibility of the match, the blood donor, whether or not it's a pediatric case, whether there's a distance involved and so forth.

And it arranges -- it evaluates the potential transplants on this scale and then the optimization considers all the possible matches and chains of matches and it finds the chains with the largest total number of quality points. And it can also consider what we call altruistic donation which is a situation where a donor has no significant other that needs a -- that needs to be a recipient of an organ.

So the bottom line is that in the first 18 months that this program was in place they performed 18 paired donations. And in one recent month they found potential matches for 21 patients in 10 transplant centers. Between July of 2007 and March of 2008 they created a 10 donation chain that started with an altruistic initial donor. And as a result of this application they were featured on a program on NBC called Tech Now. And in 2009 IBM added this organization to its list of organizations for the employee charitable contributions campaign.

So the value of this solution was that it created a situation where transplants were performed where they could not have otherwise have been performed, they were able to do the matches much more quickly and much more accurately and they were able to save lives.

Just a final few remarks and then I'll turn it back over to Gabi. Optimization is everywhere, there are applications in multiple industries, multiple sectors. Optimization usually produces calculable returns on investment with payback in months or sometimes even weeks, and these return on investment result from capital expenses that are avoided or deferred or operating expense reductions, or improvements in revenue, revenue mix or margin.

They can improve customer satisfaction by providing better service and more customized customer service. They can improve employee satisfaction by satisfying scheduling preferences while improving productivity and ensuring better planning and scheduling processes.

They result in better decisions, faster decisions and smarter decisions for a smarter planet.

Optimization solutions can provide smarter solutions. You get better performance for lower cost because it can find non-obvious solutions, solutions that a human being would not find easily and perhaps not at all; and it produces quantifiable benefits on a bottom line. It produces faster decisions because automating the decision process increases the speed of responses in today's accelerating markets and allows your managers and planners to focus on critical complexities rather than on routine issues.

If you're building an optimization application our suite of tools called CPLEX Optimization Studio enables faster, lower cost development and maintenance. We have high-level modeling tools that enable analysts to code and validate models with less time and less effort and we can make maintaining and upgrading your system more reliable.

We provide effective enterprise decision support through ODM enterprise, it uses development of optimization based custom planning and scheduling solutions, offers powerful decision management functionality and includes out of the box capabilities to configure your business user interface, a central data server and scalable optimization.

And it enables you to turn information into action. A lot of you are making substantial investments in enterprise information technology. Advanced analytics turns that information into action.

On this presentation you will find some additional websites to go to for additional information, I encourage you to take advantage of them. And at this point I am going to turn it back over to Gabi, who will handle the question-and-answer session.

Gabriela Wilson: All right, thanks Jeremy for a great presentation and now it's time for our Q&A session. If you would like to ask a question please do so by typing your question in the question box located at the bottom left of your console and hit the Submit button.

So, Jeremy, are you ready with the next question?

Jeremy Bloom: Right now I don't see any questions so I'm still waiting for somebody to ask a question.

In the meantime, I'm often asked, what's the next step, how do I find out if I have an optimization application? Well, you can contact IBM, and in particular you can contact us through our website which is optimization on www.ibm.com and we will have a representative talk to you about what specific applications you have in mind or what business issues you're trying to deal with and we will see whether there is an optimization application that you can use.

I see a few questions coming in now.

Gabi, do you want me to address the questions?

Gabriela Wilson: Yes, that would be great, thank you.

Jeremy Bloom: Okay, so one question is, can I use your presentation to promote optimization?

We would certainly appreciate it, go right ahead and if you would like some assistance from us let us know.

Another question is, what platform does ODM run on?

Right now it runs on a Windows client and a Windows server.

Additional questions, do I use third-party value added resellers?

Absolutely, we're very pleased to work with you. Again, contact us and we'll put you in touch with one of our partner representatives.

Gabriela Wilson: Any more questions Jeremy?

Jeremy Bloom: At this point I don't see any. We'll give it another couple of minutes and if there are no further questions --

Gabriela Wilson: Sure.

Jeremy Bloom: As an IBM partner, do we have resources for training and promotion?

Yes, we do. And again, we'd be glad to talk to you about that.

Additional questions?

Well, we're at the top of the hour I don't see any more questions so Gabi, I'll turn it back to you.

Gabriela Wilson: Sure, so that's all we have time for today, ladies and gentlemen. Thank you Jeremy for a great presentation and thank you very much to the audience for joining us for today's event, Prescriptive Analytics and Optimization for Smarter Business Decisions.

Have a great day everybody.